THE UNIVERSITY OF CHICAGO

SUPERSPEC: DEVELOPMENT AND CHARACTERIZATION OF KINETIC INDUCTANCE DETECTOR-BASED ON-CHIP SPECTROMETER FOR MILLIMETER AND SUB-MILLIMETER OBSERVATIONS

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

DEPARTMENT OF PHYSICS

BY
RYAN ALEXANDER MCGEEHAN

CHICAGO, ILLINOIS
AUGUST 2023

I would like to dedicate this thesis to Erik Shirokoff who passed away in early 2023. The work presented here was inspired by his vision, and I hope it lives on.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

First, I want to acknowledge my advisor Prof. Erik Shirkoff, who sadly passed away earlier this year. I will forever remember how kind and caring he was as an advisor. He took a chance on me, letting me join as his first graduate student when I had limited lab experience, and I'm glad he did.

Thank you to my committee, Prof. Stephan Meyer, Prof. Craig Hogan, and Prof. Juan Collar for their guidance and support. Prof. Clarence Chang, thank you for your support in these last few hard months, and for teaching my favorite class that I took graduate school. To Kirit Karkare, thank you for taking on the role as advisor and helping me cross the finish line. I've learned so much in the last couple years tinkering in the lab with you. I'll be forever thankful for how much I've learned from you and your support over the last several years. Thank you Peter Barry, you came in and revolutionized the lab and helped to create a great lab environment for a lot of the work presented here. It was a lot of fun in those early years setting up the lab and discovering how to do fabrication. To the SuperSpec team, particularly Matt Bradford, Steve Hailey-Dunsheath, Joe Redford, and Jordan Wheeler, you've all provided incredible discussion and feedback and helped shape me as a scientist. I am looking forward to what you all accomplish with SuperSpec and beyond. To all of Shirolab, a big thank you for all the good times in the lab over the years.

To my Mom, Dad, and Kelly, it's been a long ride, and I'm so thankful you've supported me all the way.

ABSTRACT

SuperSpec is an ultra-sensitive, moderate resolution ($R \sim 300$), on-chip spectrometer for mm and sub-mm wave observations of high redshift dusty galaxies. The device employs a filter-bank architecture in which Titanium Nitride kinetic inductance detectors (KIDs) are coupled to mm-wave resonant filters along a single microwave feedline. The detector's small size and and natural multiplexability compared to traditional grating spectrometers will enable future multi-object spectrographic measurements as well as high redshift line intensity mapping observations. This thesis presents the testing and characterization of these detectors, as well as preparations for on-sky deployment. Additionally, we present the progress on improvements to the SuperSpec filter bank technology utilizing thin film aluminum KIDs with a newly designed mm-wave filter bank structure.

CHAPTER 1

INTRODUCTION

1.1 mm and sub-mm Astrophysical Observations

The mm and sub-mm radiation bands contain a wealth of information about the early universe that we have yet to fully explore. The cosmic infrared background (CIB), consisting primarily of light from high redshift infrared galaxies, contains a substantial amount of total integrated power of radiation, which is comparable to the integrated optical and UV power in the universe, as we see in Figure 1.1[Dole, H. et al., 2006]. These early star forming galaxies radiate strongly in the UV and optical, which is absorbed by interstellar dust, and re-irradiated in the far infrared and sub-mm bands. In fact, by some estimates [Hauser and Dwek, 2001], nearly half of all star generated radiation has been down-converted into these observation bands, which is why there is such interest in these measurements.

This field is relatively immature compared to UV and optical detection counterparts due to some unique technological challenges in the mm and sub-mm observation bands [Pardo et al., 2001]. First, for ground-based observations, the atmospheric absorption due to water vapor limits the frequencies that are observable. From the ground, the atmosphere permits particular windows where water vapor is largely transparent, and SuperSpec targets one of those windows in the 185 GHz to 315 GHz range. Even in these windows, atmospheric transmission suffers due to smaller levels of water vapor absorption, and thus, observations greatly benefit from dry and high-altitude observation sites. Some of these sites include the Atacama desert [Menanteau et al., 2010], Sierra Negra [Schloerb and Carrasco, 2004], or the South Pole [Benson et al., 2014]. Additionally, mm and sub-mm wave photo-detectors are newer technologies being actively developed. Only recently, the field reached background limited detectors for ground-based observations at excellent observation sites. A significant effort, including the detector design considerations of SuperSpec, include the desire to create

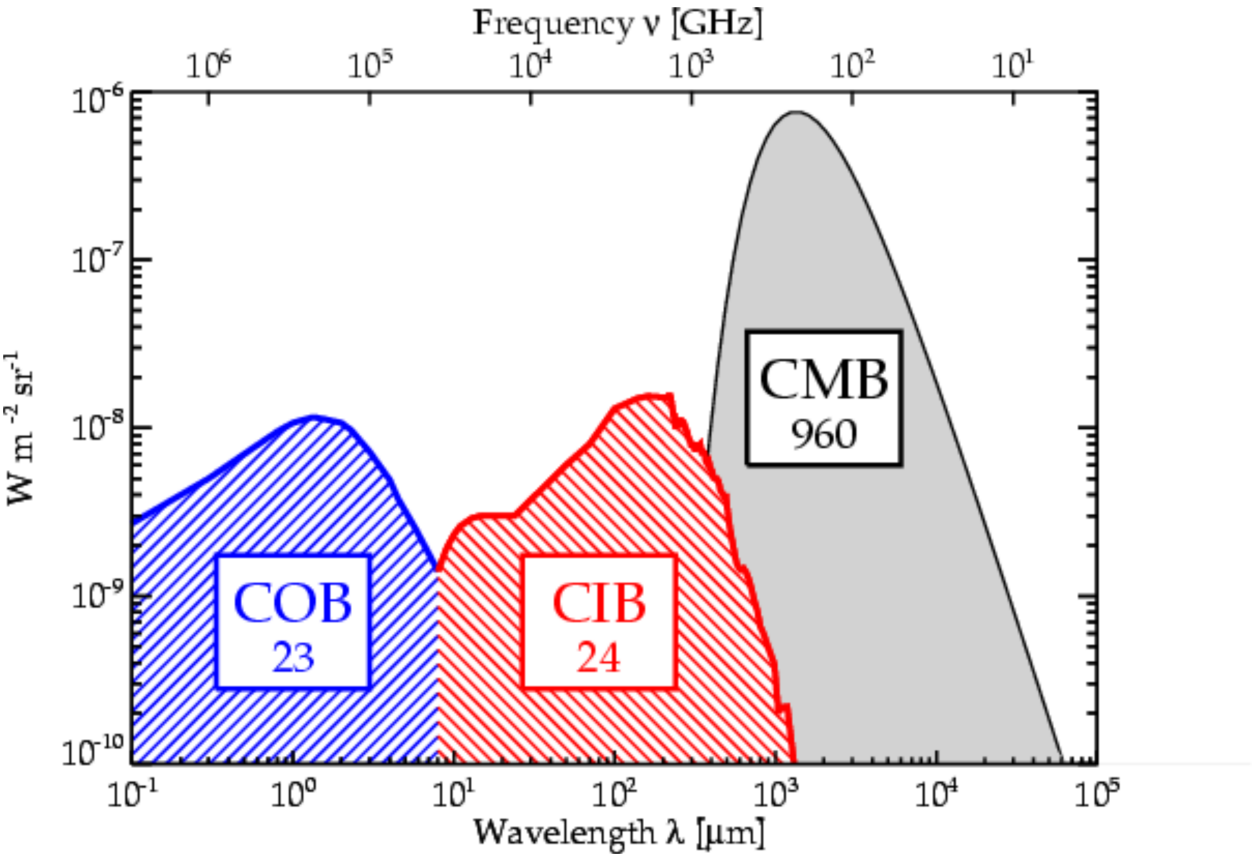

Figure 1.1: The relative intensities of the three most prominent radiation backgrounds, including the cosmic optical, infrared, and microwave backgrounds. (Figure Dole et al. 2006)

large arrays of detectors with high pixel counts to achieve efficient mapping of large sky areas.

1.1.1 Star Formation History

Approximately 400,000 years after the Big Bang, the density of the Universe decreased enough to allow the hot dense fog of ionized gas to cool and combine into neutral hydrogen and helium. At this moment, photon decoupling occurred, generating the cosmic microwave background radiation relic (CMB). It took another 400 million years for the neutral matter in the Universe to begin to gravitationally clump, cool, and collapse, forming the first stars and galaxies. Over the next several billion years, star formation rose, reaching a peak at

approximately $z = 2$, decreasing by a factor of 30 to the present day, which can be seen in Figure 1.2.

The young stars strongly emit in the UV, which would suggest observations in the UV are ideal tracers of star formation. Yet, UV observations come with their own challenges. The main factor is that UV radiation is absorbed by interstellar dust, and therefore the dust must be well-modeled and corrected for to obtain correct star formation rates. But these dust-extinction corrections suffer from high uncertainties [Madau and Dickinson, 2014]. On the other hand, the UV radiation absorbed by the dust largely is reirradiated in the far infrared, and additionally, this dust emission and Active Galactic Nuclei (AGN) account for the majority of the total IR luminosity.

Another phenomenon known as the negative K correction, is a curious alignment that allows us to observe distant galaxies [Blain and Longair, 1993]. Essentially, the thermal blackbody spectral energy distribution (SED) peak shifts into the observable mm-wave band as the galaxy distance increases. And happily, this corresponds to an increased luminosity that nearly exactly cancels out the cosmological dimming due to the increased distance [Casey et al., 2014].

At wavelengths longer than the blackbody peak, we are in the Rayleigh-Jeans (RJ) limit where the SED scales as $\nu^{2+\beta}$, where $\beta \approx 1.5$ to 2 is the dust emissivity index. The observed flux density will then scale as $S_\nu \sim \nu^{2+\beta}/D_L^2$, where D_L is the luminosity distance. It can be approximated that $D_L \sim (1+z)^2$. Finally, due to redshift, $S_\nu(z) \sim \nu_{rest}^{2+\beta}(1+z)^{2+\beta}/(1+z)^4 \sim (1+z)^{\beta-2}$, which is approximately constant. Therefore, across a wide range of redshifts where the RJ applies ($z \approx 1$ to 8), the observed flux density of galaxies is approximately constant, or slightly increasing with redshift. This brightness in the far-IR dust emission at high z makes this observation band a great candidate for studying high redshift objects.

There are other observations in the mm and sub-mm bands that can inform astrophysical processes within galaxies. The dust continuum emission described above traces the star

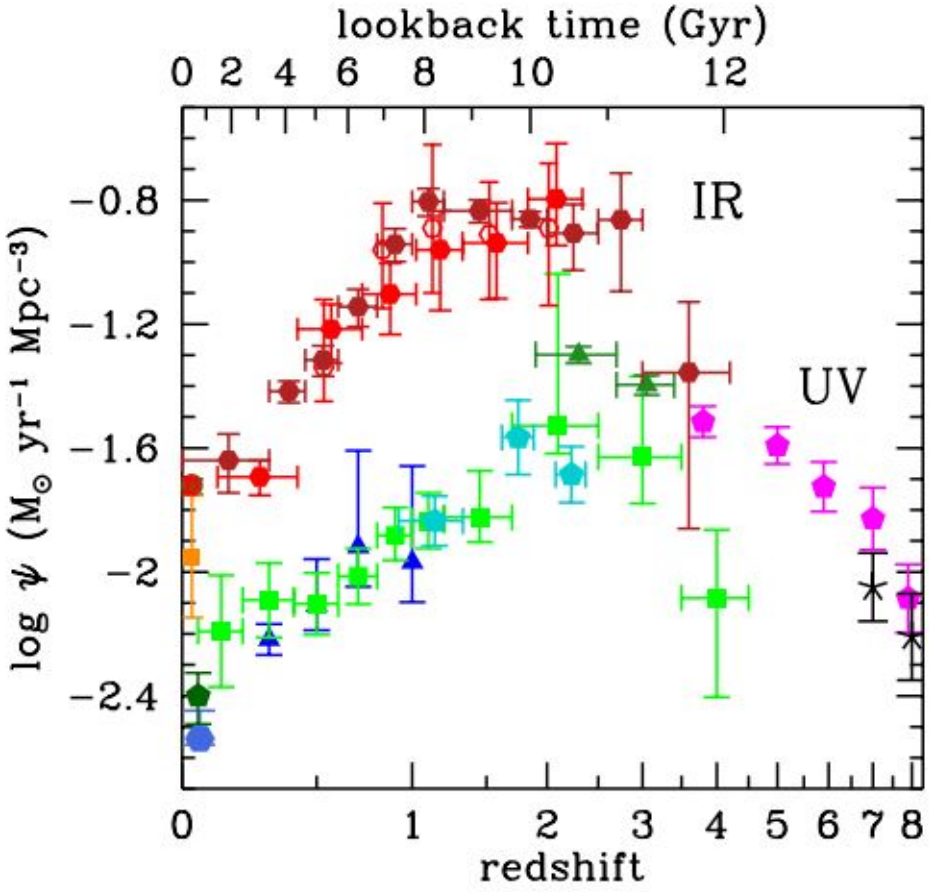

Figure 1.2: Measured star formation rate densities as a for different bands. We see peak at redshift $z = 2$. Note that the IR measurements are unmeasured above $z = 3$, which is the gap that SuperSpec aims to fill. (Figure from Madau and Dickinson [2014])

formation rate by directly mapping onto the total radiation output of the young hot stars, but does not actually describe anything about how the stars formed. The star formation process can be studied by observing the cool molecular gas around stars, revealing information about processes such the gas kinematic. The spectral lines associated with various molecules found in these photo-dissociative regions around young stars provide useful probes into formation and dynamics [Hollenbach and Tielens, 1997, 1999]. One such molecular gas commonly used is carbon monoxide (CO). CO is the most abundant molecule in these cool molecular gases after Hydrogen, and has a dipole that is well matched to typical energy scales and conditions in molecular gas clouds [Carilli and Walter, 2013]. Spectroscopic CO measurements can

reveal where in a galaxy stars are forming as well as with high resolution spectroscopy, one can measure the kinematics of the gas flow.

CO is a linear polar molecule, and contains angular momentum states with corresponding quantum number $J = 0, 1, 2, \ldots$ such that J^2 has eigenvalues of $J(J+1)\hbar^2$. The frequency of an angular momentum transition, with $\Delta J = \pm 1$ due to momentum conservation, $\nu_{J \to J-1} = J(\hbar/2\pi I)$, where I is the moment of inertia of the molecule. We can see that the transition generates an excitation proportional to the angular momentum number J so that we observe a ladder of emission lines. Measuring two adjacent emission lines at redshift z will be separated by $\Delta\nu = (\hbar/2\pi I)/(1+z)$, giving the redshift of the source.

Another bright line SuperSpec will target is the [CII] hyperfine line in ionized carbon, accounting for approximately 0.3% of the total IR luminosity [Gong et al., 2011]. [CII] acts as a tracer of star formation across a wide range of redshifts. As well, in photo-dissociate regions, a key mechanism for gas heating and cooling is the photoelectric effect, where UV photons are obsorbed by dust, which then eject electrons, heating the surrounding gas. The brightness of the [CII] line intensity shows that it plays a particularly important role in this heating and cooling mechanism in young stars. Most importantly, the rest frame line lies at 158 µm, which makes it observable in the SuperSpec band, the 1 mm atmospheric window, for approximately $z = 5-9$. Spectral lines do not benefit from the above mentioned negative K-correction, and thus the incredible brightness of [CII] has established the line as a reliable high-redshift observable.

1.2 The Future of Spectroscopy

Heterodyne receivers such as ALMA [P.C. Cortes, 2023] employ the use of large spatial arrays of receivers to observe mm-wave frequencies, and then down convert to IF while preserving phase and amplitude of the observed signal. This measurement technique is capable of high-resolution spectroscopic measurements with unrivaled sensitivities when compared to a

single dish telescope. ALMA is a fantastic instrument for observing single sources. However, by building a wide-band multi-object spectrometer, one could observe hundreds of galaxies simultaneously, primarily for galaxy surveys. This observation strategy, although not as sensitive as ALMA, can identify many more objects and their redshifts. An instrument like ALMA is not optimized for these surveys, as heterodyne instruments are lacking in instantaneous bandwidth, and therefore could miss the detection of galaxies if the redshifts are not well constrained, meaning the spectral lines could fall outside of the observation band. A multi-object, wide-band spectrometer would actually feed into an ALMA observation strategy by identifying many objects and constraining their redshifts for further single object observation follow-up.

Large-bandwidth mm-wave spectrometers have been successfully deployed in the past, such as Z-Spec or TIME-Pilot, which is a cryogenic, broadband, mm-wave grating spectrometer [Earle et al., 2006, Crites et al., 2014]. Taking Z-Spec as an example, this groundbreaking instrument has moderate resolution ($R = 200 - 400$) and observed in the desired 1 mm atmospheric band of 190 GHz to 305 GHz. Looking towards the future, though, this detector archetype lacks scalability. One spectrometer instrument is approximately 1 m in scale. With background limited detectors, the way to scale up observations in the future is to integrate with more detectors, and for that, we must miniaturize. As well, as the field has moved towards cryogenic systems to achieve low noise, a therefore cryogenic cold space in the receiver is precious.

SuperSpec attempts to solve these scaling issues by creating a moderate resolution, photon-noise limited spectrometer on a chip. A SuperSpec chip for comparison is 3.5 by 5.5 cm in size, theoretically allowing for 2 orders of magnitude more devices packed into the same space as an instrument like Z-Spec. The second large challenge of scaling to large arrays of spectrometers are the readout concerns. Each spectrometer pixel will contain hundreds of individual detectors to be read out, one for each spectrometer channel. Future

line intensity mapping (LIM) experiments will require two dimensional arrays or optimally filling focal planes, and it is predicted that there will be hundreds or thousands of detectors per spectrometer pixel [Karkare et al., 2022b]. To overcome this challenge, SuperSpec uses the innately multiplexable kinetic inductance detector (KID) as a power detector on each spectrometer channel [Day et al., 2003].

1.3 Microwave Kinetic Inductance Detectors

Simply put, a photo-detector measures some physical property of the detector system that changes as a function of number or flux of photons incident on the detector. A KID makes use of the kinetic inductance effect, where superconducting electrons, called Cooper pairs, reacting to AC fields carry non-zero inductance due to the inertia of the charge carriers. This kinetic inductance depends on the Cooper pair density, and thus the metal film impedance will be modified by changing the Cooper pair density. The KID detection mechanism then arises as a pair breaking phenomenon, wherein a photon of energy $E > 2\Delta$ is absorbed by the superconducting metal and breaks a Cooper pair, or equivalently, creates a quasiparticle above the band gap. We then design superconducting circuit around this phenomenon to be able to measure this change in impedance as a function of pair-breaking photon power incident on the detector. If we pattern our superconducting metal film into a microwave resonator with quality factor Q_r and resonant frequency f_0, a photon flux with sufficiently high energy will modify these quantities. We can then use microwave readout techniques to measure the phase and amplitude of a microwave signal interacting with the resonator, and in effect, measure this change in resonant frequency. Through careful measurement, we quantify how the resonator changes as a function of the optical power with well understood sources mm-wave radiation, which allows us to calibrate the detector response.

As we begin to make detectors that are background limited, the only way to get high signal to noise is to increase the total number of photons integrated. This can be achieved

by simply observing for longer times, or by observing with more detectors. As focal planes scale to hundreds or even thousands of detectors, readout becomes a significant factor in the design, and this is where KIDs excel. The ability to lithographically define the resonant frequencies of the KIDs, it is possible to pattern thousands of detectors at different resonant frequencies. Furthermore, these detectors can all be read out on a single microwave feedline. This greatly reduces the complexity, cost, and size of the readout scheme, as well as other microwave components, such as the use of a single cryogenic amplifier for all the detectors on a chip. That said, there are many technical challenges in developing these detectors. For example, It has been observed in many systems that KIDs exhibit excess low frequency noise, likely as a consequence of two level systems (TLS).

There are broadly two classes of KIDs, which are commonly referred to as lumped element or distributed. A lumped element KID such as SuperSpec geometrically separates out the capacitive and inductive elements of the KID resonator with an interdigitated capacitor and an inductive meander. As the pair-breaking mechanism affects the inductance, the mm-wave radiation is directed towards the inductor. Commonly, these inductive structures are used as free-space absorbers, directly absorbing the incoming radiation. Distributed resonators are lengths of transmission line with cavity frequencies set by the length as well as the grounding conditions, commonly using $\lambda/4$ or $\lambda/2$ resonances.

KIDs typically operate in the range of 0.1 GHz to 5 GHz, with SuperSpec residing in the lower edge of this range. Lower frequency KIDs offer the advantage of having access to more octaves per unit bandwidth, but it can prove difficult to achieve high enough capacitence and inductance to drive the frequencies so low.

KIDs are a relatively new technology with promising simplicity in fabrication and readout scaling. Although the technology is new, there are several experiments that have deployed background-limited arrays consisting of thousands of KIDs including BLAST-TNG Lourie et al., 2018, TolTec [Bryan, 2018], NIKA-2 [Adam et al., 2018], and MUSCAT [Brien et al.,

2018]. In fact, two of these experiments, TolTec and MUSCAT, have or are currently deployed at the 50 m Large Millimeter Telescope Alfonso Serrano, which is the proposed site for SuperSpec deployment.

1.4 Thesis Outline

Chapter 2 will describe the theoretical underpinnings of why and how a KID works. We present the basic modeling of superconducting resonators, microwave circuits, and KID noise sources. Chapter 3 contains the bulk of the information about SuperSpec device description, testing, and characterization. In Chapter 4, we highlight several other projects undertaken on the path towards deployment, including readout development and resonator surgery. Finally, in chapter 5, we describe the design, fabrication, and testing of a newly re-designed aluminum version of the SuperSpec concept.

CHAPTER 2

PHYSICS OF MICROWAVE KINETIC INDUCTANCE DETECTORS

In this chapter, I will describe the physics governing the kinetic inductance detector (KID). We must first take a look at the electrodynamics of superconductors, as the detection mechanism of a KID comes from the understanding of superconducting phenomena. Below the superconducting critical temperature T_c, DC currents flow with no loss, with this supercurrent being comprised of a pair of bound electrons with binding energy $2\Delta \approx \frac{7}{2}k_B T_c$, known as Cooper pairs. However, in response to AC fields, these Cooper pairs have nonzero impedance, and it is fundamentally this impedance that will change a function of photon flux that can be used to make a photodetector. Thus, it is imperative to understand and model the charge carrier in a superconductor as well as the impedance of thin metal films.

2.1 Superconductivity

2.1.1 The Two Fluid Model

The two fluid model gives a simple description of the flow of charge carriers in a superconducting system, and as its name suggests, there are two types, namely normal electrons and superconducting Cooper pairs [Gorter and Casimir, 1934]. Below the superconducting transition temperature T_c, the electrons separate into these two fluids, where the population of Cooper pairs carry charge with zero dissipation. However, at non-zero frequency, the inertia of the massive electron results in a non-zero complex impedance. One can image in a crude model, that the supercurrent electrons gain kinetic energy due to the electron mass as they are accelerated in an electric field. However, when the field is reversed, and the supercurrent direction reversed, that kinetic energy must be transferred to the field as the electrons slow and ultimately reverse direction. This inertia of the electron causes it to lag behind the

electric field, acting identically to energy stored in an inductor. One can derive the complex conductivity as

$$\sigma_n(\omega) = \frac{\sigma_0}{1 + j\omega\tau} \tag{2.1}$$

where $\sigma_0 = n_e q_e^2 \tau / m_e$, with τ being the characteristic scattering time of the electron in the metal. We see that at low frequencies and low scattering times, meaning $\omega\tau \ll 1$, the complex part of the conductance vanishes. However, at low temperatures, and with low levels of impurities in the metal, characteristic times can be many orders of magnitude larger than at room temperature. Thus, the imaginary part of the conductance becomes appreciable at low temperatures and high frequencies.

2.1.2 Mattis-Bardeen Theory

A more mature theory of superconducting electrodynamics was developed by Mattis and Bardeen [1958] that well models the behavior we can observe in KIDs. This thesis will not cover the details of these derivations, but for more detail, refer to Mazin [2004], Zmuidzinas [2012], and Gao [2008]. In this theory, we have a generalized expression for the complex conductivity of a superconducting metal $\sigma = \sigma_1 - i\sigma_2$

$$\frac{\sigma_1}{\sigma_n} = \frac{2}{\hbar\omega} \int_{\Delta}^{\infty} dE \frac{E^2 + \Delta^2 + \hbar\omega E}{\sqrt{E^2 - \Delta^2}\sqrt{(E + \hbar\omega)^2 - \Delta^2}} [f(E) - f(E + \hbar\omega)] \tag{2.2}$$

$$\frac{\sigma_2}{\sigma_n} = \frac{1}{\hbar\omega} \int_{\Delta}^{\Delta + \hbar\omega} dE \frac{E^2 + \Delta^2 - \hbar\omega E}{\sqrt{E^2 - \Delta^2}\sqrt{\Delta^2 - (E - \hbar\omega)^2}} [1 - 2f(E)] \tag{2.3}$$

where σ_n is the normal state conductivity, Δ the superconducting band gap energy, and f is the distribution function for the quasiparticles, which for thermal equilibrium is the Fermi-Dirac distribution $f(E) = 1/(\exp(E/k_b T) + 1)$. For this distribution, for the density of electron states at the Fermi energy N_0, it then follows that the quasiparticle density can

be calculated as

$$n_{qp} = 4N_0 \int_{\Delta}^{\infty} dE \frac{E}{E^2 - \Delta^2} f(E) \tag{2.4}$$

We can obtain simplified expression by taking the low temperature limit of $k_B T, \hbar\omega \ll \Delta_0$, where $\Delta_0 = \Delta(T = 0)$, and calculate the following expressions from the above integrals,

$$\frac{\sigma_1}{\sigma_n} \approx \frac{4\Delta}{\hbar\omega} \exp(-\Delta/k_B T) \sinh \xi K_0(\xi) \tag{2.5}$$

$$\frac{\sigma_2}{\sigma_n} \approx \frac{\pi\Delta}{\hbar\omega}[1 - 2\exp(-\Delta/k_B T) \exp(-\xi) I_0(\xi)] \tag{2.6}$$

In these simplified expressions, $\xi = \hbar\omega/2k_B T$, and I_0 and K_0 refer to the Bessel function of the first and second kind respectively.

$$n_{qp} \approx 2N_0 \sqrt{2\pi k_B T \Delta} \exp(-\Delta/k_B T) \tag{2.7}$$

2.1.3 Surface Impedance of Thin Films

So far, we have described the complex impedance of general superconductors in particular frequency and temperature limits. In detector systems like the one described in this thesis, we are directly measuring not the conductivity, but the impedance of thin films. Therefore the goal of this section is to link the previous conductivity calculations to the impedance of the thin film.

We must first consider several different limits related to film thicknesses and physical length scales according to material parameters [de Visser, 2014, Zmuidzinas, 2012]. First, for a material with electron Fermi velocity v_f, we consider the coherence length $\xi_0 = \hbar v_f/\pi\Delta_0$. This length scale concerns the coherence length of Cooper pairs, and can be roughly interpreted to be the physical extent of the electron pair. This coherence length will be degraded

due to non-idealities in the superconductor's physical structure, such as impurities in the lattice, such that the effective coherence length will be

$$\frac{1}{\xi} = \frac{1}{\xi_0} + \frac{1}{l} \tag{2.8}$$

where $l = \tau v_f$ is the mean free path of the Cooper pair and τ the collision timescale of the charge carrier with the impurities. We refer to the clean limit being where the mean free path is much larger than the coherence length, $l \gg \xi_0$. The physical interpretation here is that the length scale of the Cooper pair dynamics is much smaller than the scale of the impurities, so that the material acts like a pure material with $\xi = \xi_0$. In the dirty limit, taking the opposite $l \ll \xi_0$, then $\xi = l$.

The final length scale we consider is the penetration depth $\lambda_L = \sqrt{mc^2/4\pi n_s q_e^2}$, which can be physically interpreted as the length scale to which the magnetic field can penetrate the surface of the superconductor. For a film thickness much smaller than the penetration depth, the current distribution as a function of position within the superconductor is constant. In the thin film limit, due to the diffuse reflection condition at the surface of the film, it can be seen that for film thickness $t \ll \lambda$, then $t = \xi = l$. One can then derive for a general surface impedance $Z = R_s + j\omega L_s = R_s + jX_s$.

$$Z_s = \sqrt{\frac{j\omega\mu_0}{\sigma}} \coth \frac{t}{\lambda}\sqrt{1 + j\frac{\sigma_1}{\sigma_2}} \approx j\mu_0\omega\lambda \left(1 + \frac{j\sigma_1}{\sigma_2}\right)^{-1} \tag{2.9}$$

One can see that with $\sigma = \sigma_1 - j\sigma_2$, then we can re-arrange, and find

$$Z_s = \frac{1}{\sigma t} \tag{2.10}$$

$$\frac{\delta Z_s}{Z_s} = -\gamma\frac{\delta\sigma}{\sigma} \tag{2.11}$$

Therefore, a fractional change in the surface impedance is related to a fractional change in the conductance. The equation 2.11 is a generalized form of this relationship, parameterized by γ, where in the thin film limit, we have calculated $\gamma = 1$. This calculation can be repeated in other cases [Gao, 2008], where one can obtain $\gamma = 1/2, 1/3$ in various other limits which do not apply to the films considered in this work. This relationship between surface impedance and conductivity allows us to model the measured impedance changes in a superconducting resonant circuit, and map those measurements to the physical parameters in the model. For example, in section 3.3, we will use this relationship to measure the superconducting transition temperature T_c by measuring the change in resonant frequency as a function of temperature.

2.1.4 Quasiparticle Lifetime and Responsivity

At non-zero temperature, the quasiparticle density will undergo statistical fluctuation. This process, known as generation and recombination, occurs when due to random fluctuations, Cooper pairs will be broken, and a quasiparticle generated. This quasiparticle will exist on average a time τ_{qp}, which we call the quasiparticle lifetime, before the electrons once again condense into a Cooper pair. As we can see from the calculations in the previous section, these fluctuation will affect the surface impedance, and thus introduce an intrinsic noise source for superconducting films.

We seek now to relate the change quasiparticle density due to a change in optical power by using a previous calculation of N_{qp} and a rate equation. For thermal recombination rate Γ_{th}, recombination rate Γ_r, optical generation rate Γ_{opt}, we find

$$\frac{dN_{qp}}{dt} = 2(\Gamma_{th} + \Gamma_{opt} - \Gamma_r) \tag{2.12}$$

As well, from Kaplan et al. [1976], we get an expression of τ_{qp}, which can be related to the quasiparticle density by

$$\tau_{qp}^{-1} = \frac{\sqrt{\pi}}{\tau_0}\left(\frac{2\Delta}{k_B T_c}\right)^{5/2}(T/T_c)^{1/2}\exp(-\Delta/k_B T) = \frac{2\Delta^2}{N_0(k_B T_c)^3}\frac{n_{qp}}{\tau_0} \tag{2.13}$$

We can then define a useful recombination constant $\tau_{qp} n_{qp} = 1/R$, and we find from Wilson and Prober [2004], that the recombination rate can be expressed simply in terms of this constant, with $\Gamma_r = RVn_{qp}^2/2$, where V is the volume such that $V = N_{qp}/n_{qp}$. By further noting that the generation and recombination rates are equal in thermal equilibrium, and setting $\Gamma_{opt} = \eta_e P_{opt}/2\Delta$, we can use equation 2.4 and the rate equation to find $n_{qp} = \sqrt{n_{th}^2 + n_{opt}^2}$ and

$$n_{opt}^2 = \frac{N_0\tau_0(k_B T_c)^3}{\Delta^3 V}\eta_{pb} P_{opt} \tag{2.14}$$

Here, $\eta_{pb} \approx 0.57$ is the pair breaking efficiency, which is less than 1 due to inefficiencies in the conversion of photon energy into quasiparticles due to other possible excitations such as phonons [Day et al., 2003, de Visser et al., 2012, Chang and Scalapino, 1977] . If the quasiparticle density is dominated by optical quasiparticle generation, then we can directly compute the change in quasiparticle density as a function of optical power as

$$\frac{dn_{qp}}{dP_{opt}} = \frac{N_0\tau_0(k_B T_c)^3\eta_{pb}}{4\Delta^3 V}P_{opt}^{-1/2} = \frac{\eta_e\tau_{qp}}{\Delta V} \tag{2.15}$$

We note from equation 2.15 the scaling relationship between quasiparticle density, lifetime, and incident optical power, with $n_{qp} \sim P_{opt}^{-1/2} \sim \tau_{qp}$.

Another form for τ_{qp} was proposed by Zmuidzinas [2012], accounting for other limiting mechanisms in the quasiparticle lifetime. From the previous equations, it would seem that at low temperature, the the quasiparticle density decreases exponentially, and therefore the lifetime should increase exponentially. It has been observed that there is a limiting τ_{max}, and thus we should consider the relationship between the lifetime and quasiparticle density as

$$\tau_{qp} = \frac{\tau_{max}}{1 + n_{qp}/n^*} \tag{2.16}$$

where n^* sets the scale where the inversely proportional scaling begins to dominate.

To summarize so far, we have calculated the change in quasiparticle lifetime as a function of optical power, but we must then relate this to changes in resonator parameters.

2.2 Resonant Microwave Circuits

The previous sections have generally described the properties of superconductors and superconducting films, but in order to make a photodetector, we pattern these films into high quality resonators in order to read out the properties of these resonances and their changes in response to photons. Therefore, we must understand how resonant circuits respond to changes in surface impedance due to pair-breaking photons.

We begin by reviewing the properties of a resonant parallel RLC circuit. The total impedance of the RLC circuit, Z_{RLC}, calculated by

$$Z_{RLC} = \left(\frac{1}{j\omega L} + j\omega C + \frac{1}{R} \right)^{-1} \tag{2.17}$$

Then, defining $\omega_0 = 1/\sqrt{LC}$ and the fractional detuning $x = \Delta\omega/\omega_0 = (\omega - \omega_0)/\omega_0$, then we can calculate for small perturbations $\Delta\omega \ll 1$,

$$Z_{RLC} = \frac{R}{1 + 2jQ_i x} \tag{2.18}$$

where $Q_i = \omega_0 RC$ is the intrinsic quality factor.

The previous equations are for an isolated RLC circuit, but in order to read it out, we couple capacitively to a feedline with characteristic impedance Z_0 so that the RLC circuit acts as a shunt to ground, as we see in Figure 2.1. In this case the the coupling capacitor C_c is in series with the RLC circuit such that $Z = \frac{1}{j\omega C_c} + Z_{RLC}$. The condition $\Im(Z) = 0$ is

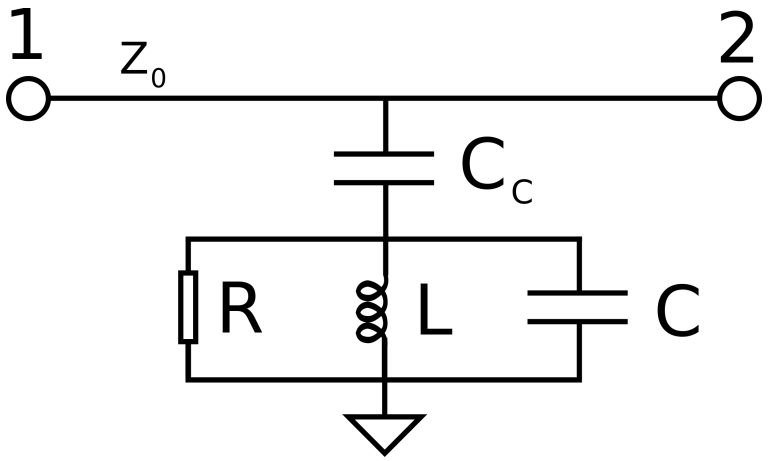

Figure 2.1: Circuit diagram of parallel RLC circuit capacitively coupled to feedline of characteristic impedance Z_0.

the definition of the resonant frequency, and thus we find that the addition of the coupling capcitor shifts the resonant frequency to ω_r such that

$$\frac{\omega_r - \omega_0}{\omega_0} \approx -\frac{\omega_r C_c}{2\omega_0 C} \tag{2.19}$$

Additionally, the capacitive coupling to the feedline introduces another loss mechanism, where power leaks out of the resonator and is lost to the terminations in the feedline. The power dissipated in this way [Barry, 2014], can be calculated by

$$P_{diss} = \frac{I^2 Z_0}{4} = \left|\frac{V_c}{Z_{res}}\right|^2 \frac{Z_0}{4} = \frac{|V_c|^2 \omega_0^2 C_c^2 Z_0}{4} \tag{2.20}$$

where $Z_{res} = 1/\omega_0 C_c$ is the resonator impedance, and V_c is the peak voltage across the coupling capacitor such that the current through the series of coupling capacitor and resonator is $I = V_c/Z_{res}$.

Additionally, the total energy stored in the circuit oscillates between the electric and magnetic field, and the maximal energy stored in the E field can be described by the peak voltage across the resonator capacitor [Pozar, 2011] such that $E = C|V_c|^2/2$. Combining this with the general coupling quality factor definition, we arrive at

$$Q_c = \frac{\omega_r E}{P_{diss}} = \frac{2C}{\omega_r Z_0 C_c^2} \tag{2.21}$$

It is also useful to quantify the resonator in terms of the scattering matrix, S. In particular, we calculate the transfer function S_{21} by

$$S_{21} = \frac{2}{2 + \frac{Z_0}{Z}} \tag{2.22}$$

Using the relationship in equations 2.19 and 2.21, and combining with the total impedance

$$Z = \frac{1}{j\omega C_c} + \frac{R}{1 + 2jQ_i x} \tag{2.23}$$

one can show that the transfer function simplifies to

$$S_{21} = 1 - \frac{Q_r}{Q_c} \frac{1}{1 + 2jQ_r x} \tag{2.24}$$

where $Q_r^{-1} = Q_r^{-1} + Q_c^{-1}$ is the total quality factor of the resonator.

2.2.1 Fitting Asymmetric Resonances

In the case where there are impedance mismatches on the input or output transmission lines, there will be a distortion in the resonator shape. From Geerlings [2013], one can derive a modified expression for the resonator transfer function as

$$S_{21} = \frac{Q_c + jQ_c Q_i \left(2x + \frac{2\delta\omega_{asym}}{\omega_0}\right)}{Q_i + Q_c + 2jxQ_i Q_c} = 1 - \frac{\tilde{Q}\tilde{Q}_c/}{1 + 2i\tilde{Q}x} \tag{2.25}$$

The new parameter $\delta\omega_{asym}$ accounts for the asymmetry caused by the impedance mismatch, leading to a slightly shifted resonance frequency $\omega_0 + \delta\omega_{asym}$. Equivalently, instead of considering this as a resonant frequency shift, this mismatch can be absorbed by letting Q_c be complex, such that $1/Q_c = \Re(1/\tilde{Q}_c)$.

2.2.2 Kinetic Inductance Nonlinearity and Bifurcation

The KID resonator is fundamentally limited in the probe tone power by what is known as KID bifurcation. This phenomenon arises due to a nonlinearity in the kinetic inductance as a function of current [Swenson et al., 2013]. There is symmetry in the direction of the current such that a power series expansion takes the form

$$L_k(I) = L_0(1 + \frac{I^2}{I_*^2} + \cdots) \tag{2.26}$$

The power dissipated in the resonator for a given signal power P_g as

$$P_{diss} = P_g(1 - |S_{11}|^2 - |S_{21}|^2) = \frac{2Q_r^2}{Q_i Q_c} \frac{P_g}{1 + 4Q_r^2 x^2} \tag{2.27}$$

Then, observing that $Q_i = \omega_r E / P_{diss}$, we can simplify such that

$$x = x_0 + \frac{2Q_r^2}{Q_c} \frac{1}{1 + 4Q_r^2 x^2} \frac{P_g}{\omega_r E_*} \tag{2.28}$$

where $E_* \sim 2I_*^2/\alpha$ is an electric field scale associated with the scale of the current based nonlinearity. With the definition $y = Q_r x$, and choosing simplifying parameters, we can show

$$y = y_0 + \frac{a}{1 + 4y^2} \tag{2.29}$$

$$a = \frac{2Q_r^3 P_g}{Q_c \omega_r E_*} \tag{2.30}$$

One can solve this cubic implicit equation in y, which depends on the value of the parameter a, which we designate the non-linearity parameter. A bifurcation, in the mathematical sense, occurs in the solutions for y at $a = \frac{4\sqrt{3}}{9}$. For $a < \frac{4\sqrt{3}}{9}$, there is only one real solution for y, but in the regime where $a > \frac{4\sqrt{3}}{9}$, that is at higher signal power, there are suddenly

three real solution, two of which are stable. The behavior of this system is an example of a Duffing oscillator [Strogatz, 2015].

The effect is that at high tone powers, the KID can be driven into the state where there are two stable solutions. When sweeping the tone across the resonance, runaway feedback once the power begins to dissipate in the KID will cause the solution to 'snap' to the low frequency solution. This in effect will cause what looks like a sharp discontinuity in the transfer function of the KID. There is advantage in driving the KID with high power though, as this nonlinearity can improve small signal response, but avoiding bifurcation is imperative to avoid the discontinuity.

2.2.3 Resonator Response

With our description of the microwave resonant circuits, we can now link the change in film impedance to changes in resonator parameters. The resonant frequency of an RLC resonator scales as $\omega \sim L^{-1/2}$, with $L = L_k + L_g$ the sum of kinetic and geometric inductance, and thus we compute the fractional change in resonant frequency

$$\delta x = \frac{\delta\omega}{\omega_0} = -\frac{\alpha\delta L_k}{2L_k} = -\frac{\alpha\delta X_s}{2X_s} \tag{2.31}$$

Where $\alpha = L_k/L$ is the kinetic inductance fraction and $X = \omega L$ is the reactance. Additionally, with $Q_i = \omega_0 L/R$

$$\delta Q_i^{-1} = \alpha\frac{\delta R}{X} \tag{2.32}$$

We can then relate these to the conductivity equations 2.5 and 2.6, and derive [Zmuidzinas, 2012] a relationship between the fraction resonant frequency shift and the quasiparticle density as

$$\frac{\delta x}{\delta n_{qp}} = \frac{\alpha\gamma}{4N_0\Delta_0}\left(1 + \sqrt{\frac{2\Delta_0}{\pi k_B T}}\exp(-\xi)I_0(\xi)\right) \tag{2.33}$$

and as well, the intrinsic quality factor, which can be calculated by $Q_\sigma = \sigma_2/\sigma_1$, changes as

$$\frac{\delta Q_i^{-1}}{\delta n_{qp}} = \frac{2\alpha\gamma}{4N_0\Delta_0}\frac{2}{\pi}\sqrt{\frac{2\Delta_0}{\pi k_B T}}\sinh(\xi)K_0(\xi) \tag{2.34}$$

where K_0 is the modified Bessel function of the second kind.

Finally, we relate the change in quasiparticle density to optical power. Similar to arguments in section 2.1.4, we recognize that for a given detector volume V, the change in quasiparticle density for a given absorbed optical power P_{opt} can be computed as

$$\frac{\delta n_{qp}}{\delta P_{opt}} = \frac{eta_{pb}\tau_{qp}}{2\Delta V} \tag{2.35}$$

The resonator response, R_{det}, can then be calculated as

$$R = \frac{\delta x}{\delta P_{opt}} = \frac{\delta x}{\delta n_{qp}}\frac{\delta n_{qp}}{\delta P_{opt}} = \frac{\eta_{pb}\tau_{qp}}{2\Delta V}\frac{\delta x}{\delta n_{qp}} \tag{2.36}$$

which can be further combined with equation 2.33 to arrive at a final expression. Similarly,

$$\frac{\delta Q_i^{-1}}{\delta P_{opt}} = \frac{\eta_{pb}\tau_{qp}}{2\Delta V}\frac{\delta Q_i^{-1}}{\delta n_{qp}} \tag{2.37}$$

2.3 Noise in Kinetic Inductance Detectors

2.3.1 Generation-Recombination Noise

At nonzero temperature, the average number of quasiparticles is well described in a superconductor by equation 2.7, but there will be stochstic variations about this average. These

fluctuations about the equilibrium result from the generation of quasiparticles, and recombination thereof into Cooper pairs. Thus, we also see a stochastic fluctuation in the film impedance of our resonator circuits, which we denote as generation-recombination noise (GR noise). If we assume that the GR processes have a flat spectral density (frequency independent), and follows Poisson statistics, the one can show [Gao, 2008, Wilson and Prober, 2004] that the spectral density of the GR process is

$$S_n(f) = \frac{4N_{qp}\tau_{qp}}{1 + (2\pi f \tau_{qp})^2} \tag{2.38}$$

Typically, the sensitivity of photodetectors are characterized by the noise equivalent power, which is the signal power giving signal to noise ration of 1 in 0.5 second integration. We can then compute this NEP due to the GR noise by calculating the equivalent optical power due to these thermal fluctuations

$$NEP_{gr} = \sqrt{S_n}\frac{dn_{qp}}{dP_{opt}}^{-1} = \frac{2\Delta}{\eta_{pb}}\sqrt{\frac{N_{qp}}{\tau_{qp}}}\frac{1}{\sqrt{1 + (2\pi f \tau_{qp})^2}} \tag{2.39}$$

This noise source presents spectrally as white noise (constant) up until the quasiparticle lifetime rolloff, where the time scales become much smaller than where quasiparticle dynamics play a role. Furthermore, the white noise level due to the GR contribution will rise as N_{qp} rises changes as a function of temperature, with an approximate scaling of $NEP_{gr} \sim \exp(-\Delta/k_B T)$.

2.3.2 Photon Noise

All photodetectors are fundamentally limited by the noise inherent in the photon signal due to its discrete nature and arrival statistics. The variance in the number of photons goes as $< \Delta n^2 >= n + n^2$. Then taking $E = nh\nu$ and integrating, with the assumption that the photon spectrum is white, we get

$$NEP_{\gamma}^{2} = \int 2h\nu P(\nu)d\nu + P(\nu)^{2} \approx 2h\nu_{0}P + P^{2}/\Delta\nu = 2h\nu_{0}P(1+\bar{n}) \tag{2.40}$$

where the NEP is calculated for absorbed power P across bandwidth $\Delta\nu$, centered at frequency ν_0, and $\bar{n}$ deontes the average photon occupation number. In this equation, we can recognize that the first term corresponds to shot noise due to Poisson arrival statistics, and the second term corresponds to the Bose statistics, commonly referred to as the "bunching" noise term.

2.3.3 Two-Level System Noise

KIDs have been known to exhibit excess frequency noise [Day et al., 2003, Gao, 2008], meaning, there is a time dependent resonant frequency jitter observable in the frequency noise power spectrum. It has been demonstrated that this noise corresponds to effects at the surface layer of the devices, and is believed to be caused by two-level systems (TLS) found in amorphous dielectrics found on this surface layer [Neill et al., 2013, Gao et al., 2008]. These electric dipole moment of these TLS systems couple to the electric field of the KID, and these coupling in turn change the effective dielectric properties of the material. As a result, due to random energy fluctuations between the two energy levels, this can equally be viewed as a noisy dielectric constant material, which then affects the capacitence and resonant frequency of the KID [Barends et al., 2008].

A precise microscopic model is not fully understood for for this source of noise, but particular scalings have been well characterized such that the power spectrum of the TLS noise contribution scales as

$$S_{TLS} \sim f^{-1/2}P_{r}^{-1/2}T^{-2}\frac{\int |E|^{3}d^{3}r}{(\int |\epsilon E^{2}|^{2}d^{3}r)^{2}} \tag{2.41}$$

where P_r is the readout power of the microwave tone, and E the electric field in the KID,

usually in the capacitor for lumped element KIDs.

From a microwave perspective, the KID coupling to an external dipole acts as a source of loss in the resonator, and therefore has a corresponding associated Q, which can be calculated [Zmuidzinas, 2012] by

$$\frac{1}{Q_{TLS}} = F_{TLS}\delta_0 \tanh(\frac{\hbar\omega}{2k_BT})\frac{1}{\sqrt{1+(\frac{P_r}{P_*})^2}} \tag{2.42}$$

where the δ_0 is the intrinsic dielectric loss tangent and F_{TLS} is a geometric factor describing the distribution of the electric field contained within the dielectric material containing the TLSs. As well, we see a power dependence as at high microwave power, the TLSs are saturated, described by the parameter P_*.

A KID with significant TLS contributions can be recognized in particular by the resonant frequency shift as a function of temperature. The Mattis-Bardeen theory of superconductors show that as the temperature of the device is increased the resonant frequency should always decease as we generate more thermal quasiparticles. One can calculate the fractional frequency shift due to TLS effects as

$$\frac{\delta f}{f_0} = \frac{F_{TLS}\delta_0}{\pi}\left[\Re\left(\Psi\left(\frac{1}{2}+\frac{\hbar\omega}{2\pi j k_BT}\right)\right) - \ln\frac{\hbar\omega}{k_BT}\right]\frac{1}{\sqrt{1+(\frac{P_r}{P_*})^2}} \tag{2.43}$$

where Ψ is the complex diagamma function. Note that this change in resonant frequency can be positive, and commonly at low temperatures, for TLS affected devices, one can see the resonant frequency rise with temperature before decreasing as the thermal quasiparticle generation becomes dominant in the frequency shift.

CHAPTER 3

SUPERSPEC

3.1 The SuperSpec Concept

SuperSpec is an on-chip spectrometer operating in the mm and sub-mm wave bands. The SuperSpec project aims to address one of the fundamental limiting factors of previous detector methods. Heterodyne detection experiments, such as ALMA, lack the ability to cover large detection bandwidths, whereas SuperSpec operates simultaneously across the 180 GHz to 310 GHz band. Classical grating spectrometers are physically too large, and therefore lack the ability to scale efficiently for putting large numbers of detectors in a single focal plane, although they are able to achieve high bandwidth and high spectral resolution. SuperSpec addresses both of these concerns by creating a lithographically patterned filter bank on a silicon chip connected to a broadband on-chip antenna. For more detail on the particular design considerations, refer to Barry et al. [2012], Kovács et al. [2012], Shirokoff et al. [2012].

3.1.1 Filter Bank

After incoming radiation is absorbed by the dual slot antenna, the radiation then propagates down a mm-wave microstrip feedline. Along this feedline, resonant filters, consisting of $\lambda/2$ half wave resonators, absorb the radiation corresponding to their resonant frequency. The mm-wave filter resonant frequencies are varied so to cover the desired detection band. Coupled to each mm-wave resonator is a KID, which essentially measures the power that has coupled into the mm-wave resonator.

With optimal coupling, half of the propagating power from the antenna will be absorbed by each filter, but this can be improved. If we space the filters along the feedline with $\lambda/4$ or $3\lambda/4$ spacing, the radiation that was not absorbed will partially be reflected by the next

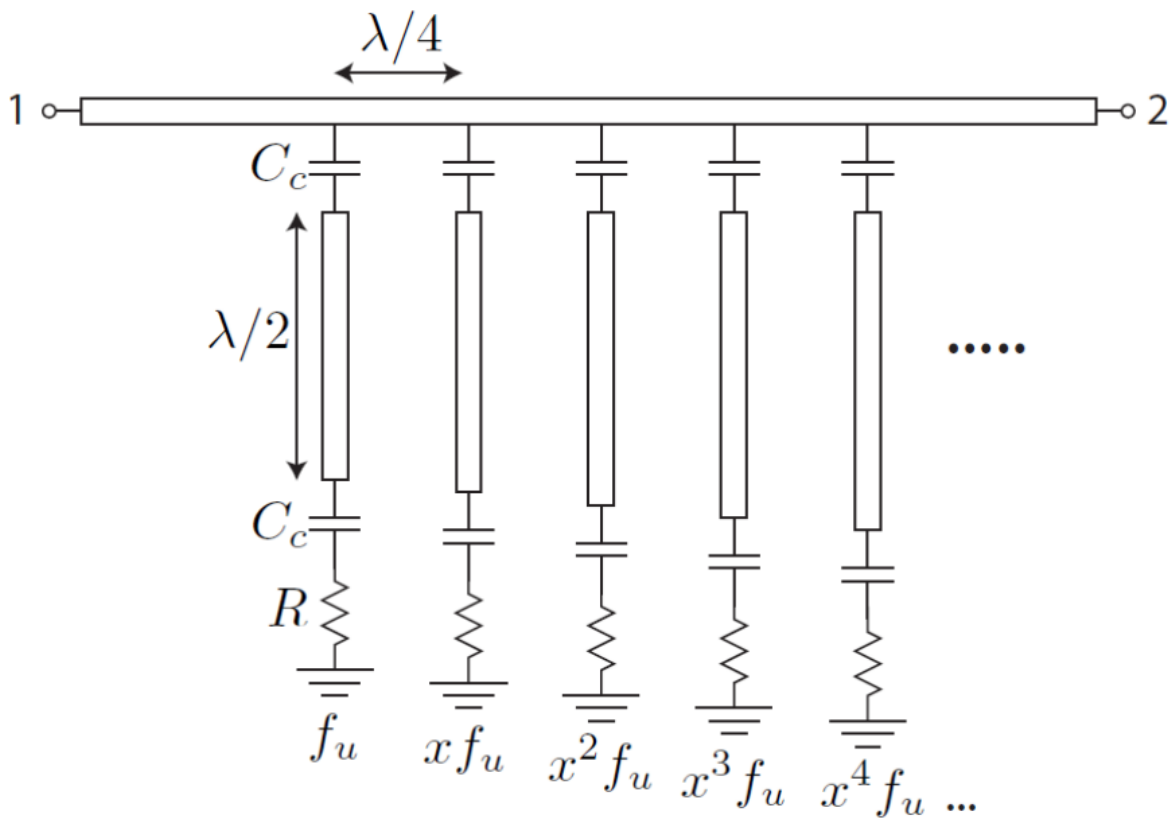

Figure 3.1: Schematic of the SuperSpec cochlear filter bank. The ports 1 and 2 label the input and output of the mm-wave feedline. A series of resonant filters tuned in decreasing frequency are coupled along the feedline, picking off the signal corresponding to its resonant frequency. These mm-wave resonators are each coupled to an individual power detector to determine the strength of the signal that coupled into the filter bank resonator. The filter quality factor is determined by the strength of these couplings along with the dielectric loss, which is quantified in Q_{loss}.Zmuidzinas, 2012

filter structure, and will then back-propagate and constructively interfere. One must also consider though further reflection from filters further down the line. The second filter will be either $\lambda/2$ or $3\lambda/2$ down the feedline, which are perfectly deconstructively interfering when reflecting back to the original filter. We have found, from Che [2018], Wheeler [2019], that the optimal spacing is less than $\lambda/4$, although $\lambda/4$ performs well. As is expected, $\lambda/2$ performs the worst, as all reflections are destructive. The size of the SuperSpec capacitors makes it difficult to achieve smaller than $\lambda/4$ spacing without starting to run into issues where the KID capacitors are too close physically, which may induce unwanted KID to KID

couplings, as well as potential ground issues. Therefore, we have employed $\lambda/4$ spacing.

While the frequency of the filters is simply set by the length of the half wave resonator, there are other important parameters to optimize. The resonator Q_{mm} defines the bandwidth of the filter, and hence, the spectrometer resolution, R_{spec}. The total quality factor of the filter is given by a combination of several quality factors such that

$$R_{spec} = \frac{1}{Q_{mm}} = \frac{1}{Q_i} + \frac{1}{Q_c} + \frac{1}{Q_{loss}} \tag{3.1}$$

Here, Q_c describes the coupling strength from the resonator to the feedline, Q_i the internal quality factor describes the coupling strength from the resonator to KID, and Q_{loss} encompasses the radiation loss mechanisms. Q_c is set by the electromagnetic proximity coupling of the mm-wave feedline to the filter half wave resonator. The resonator physically is bent into a symmetric 'U' shape, which we describe as a "staple," and the long arm of the staple is set parallel to the feedline. The length of this arm combined with the distance to the feedline determines the coupling strength. Through a combination of EM simulation and design iteration with empirical corrections, we can have this coupling factor well defined across the whole observing band. Q_i can be thought of as another loss in some sense, but it is a loss from the filter to the power detector, the KID. The sensitive part of the KID is the TiN inductor, where we want all of the pair breaking to occur. Therefore, we proximity couple the KID inductor by placing the TiN inductive meander inside of the staple filter resonator. The majority of the coupling strength occurs at the tips of the staple resonator [Shiu] due to the voltage anode at that point, and the coupling is predominantly capacitive. Q_{loss} is an amalgamation of several distinct loss mechanisms which all have a similar dissipative effect. For example, there is radiative loss, conductor loss in the Nb, and losses in the dielectric SiN. We have found that this loss is dominated by loss in the dielectric, which is typically characterized by the loss tangent of the material $\tan\delta = 1/Q_{loss}$, with a value of approximately 8.5×10^{-4} for our detectors [Hailey-Dunsheath et al., 2015, Wheeler et al.,

2018].

The ratio of Q_i and Q_c has a large effect on the efficiency of the detector, where maximal coupling of 50% occurring where $Q_i = Q_c$. When Q_c dominates, more than 50% of the incident power is reflecting, and when Q_i, a majority of the power is transmitted. We always seek to maximize the number of integrated photons, so it is imperative to hit the $Q_i = Q_c$ target as well as possible. As well, we can overcome this loss of photons by oversampling within our detector band. Suppose we define our lower and upper detector band limits as f_l and f_u. We can then construct our oversampled filter bank consisting of N detectors by logarithmically spacing them across the band in decreasing order of f_u, xf_u, ..., $x^{N-1}f_u$, where

$$x = \exp -\frac{\ln f_u - \ln f_l}{N-1} \tag{3.2}$$

From this spacing, we can define an oversampling factor Σ, and given detector resolution R_{det}, we have

$$N = \Sigma R_{det} \ln \frac{f_u}{f_l} \tag{3.3}$$

Higher oversampling will increase the overall absorption of photons, but it does come at the slight cost of less photons absorbed per detector. As well, high oversampling induces some complexity by simply increasing the number of detectors required to cover the same band, which comes with associated design, fabrication, and readout costs. The parameter Σ is the oversampling factor, that is how many resonators per linewidth are packed into the filter bank bandwidth, with $\Sigma = 2$ found to be the optimal value.

3.1.2 Kinetic Inductance Detectors

The KIDs employed in the superspec detector are lumped-element, meaning inductor and capacitor are geometrically and physically defined in separate physical structures, and in this case, with different materials. The inductor, which is the photosensitive structure of the KID is made from thin titanium nitride (TiN) film. This material has been chosen for a couple reasons. First, with its high innate kinetic inductance, the theoretical noise floor for a detector made with such a material is very low, less than $1 \times 10^{-19}\,\mathrm{WHz}^{-1/2}$ [Leduc et al., 2010]. These theoretical sensitivities approach what could even be required for low-background space platforms. Additionally, the T_c of the material is tune-able and lies in the region where mm-wave pair breaking works. Stoichiometric TiN, that is where there is maximal nitrogen in the metallic structure with high N flow rate during deposition, exhibits a superconducting transition temperature $T_c = 4.2\,\mathrm{K}$. However, by tuning the nitrogen flow rate lower during the Ti sputter, one can lower the T_c of the material, as low as $0.5\,\mathrm{K}$. As detector response scales favorably with lower T_c, but one must be careful to consider the operating temperature of the cryostat. In this case, with a He-10 adsorption fridge, the cryostat operates at $210\,\mathrm{mK}$. If the operating temperature is too close to the T_c of the TiN, the detector noise and quality factor will be dominated by the thermally generated quasiparticles. Thus, SuperSpec has generally sought to be in the range of $T_c = 1.0\,\mathrm{K}$ to $1.2\,\mathrm{K}$ so that we are operating at approximately $T_c/5$.

The TiN inductor meander is proximity coupled to the mm-wave resonant filter by placing the inductor inside of staple shaped $\lambda/2$ resonator. This is enabled by the ultra-low detector volume of $2.6\,\mu\mathrm{m}^3$. This low volume also equates to higher response, as the change in quasiparticle density per pair-breaking photon is inversely proportional to the volume of the detector. Consequently, the NEP also then scales with responsivity so that SuperSpec can incredibly low NEPs, meeting the photon-limit requirements for the best ground-based observation sites.

The inverted microstrip design utilized has the KID layer directly on the Si substrate, with the dielectric and ground plan placed on top. This has the added advantage of a clean Si surface to put the sensitive TiN material first. The lumped-element Nb capacitor attached to the TiN inductor in the same layer. Once the SiN dielectric is deposited on top, the inductor is well protected from any other fabrication concerns, such as contamination or etchants from other layers affecting the film quality. The ground plane is deposited on top, but importantly, the ground plan and dielectric are etched away above the KID capacitor in order to limit the presence of two-level system noise in the dielectric.

3.1.3 Device Layout

The first iteration of the prototype SuperSpec filter bank chips contain 50 spectral channels with spectrometer resolution of $R = 300$ [Wheeler, 2019]. We see the a cartoon version of the overall structure and layout of the 50 channel SuperSpec chip in figure 3.2. Using a hyper-hemispherical silicon lens with anti-reflection coating, the mm-wave radiation is focused onto the on-chip broadband antenna. The radiation is absorbed by the antenna, and the excitations propagate down the mm-wave feedline. Along the feedline, the tuned resonant filters pick off particular radiation frequencies, with monotonically decreasing frequency order. This ordering mitigates losses in the dielectric as the radiation propagates along the feedline because higher frequency radiation loses more power per unit length. The mm-wave filter has the TiN meander proximity coupled by putting the meander inside the 'U' shape of the resonator. The coupling to the KID allows the mm-wave power to break cooper pairs in the TiN inductor, inducing a shift in the phase and amplitude of the KID, which is then measured through a readout microwave tone.

We then scaled up to a 300 channel device, with essentially the same overall structure [Redford, 2023]. A cartoon representation can be seen in figure 3.3. This design was more aggressive though in the frequency spacing, having all 300 resonators between 100 to 500 MHz.

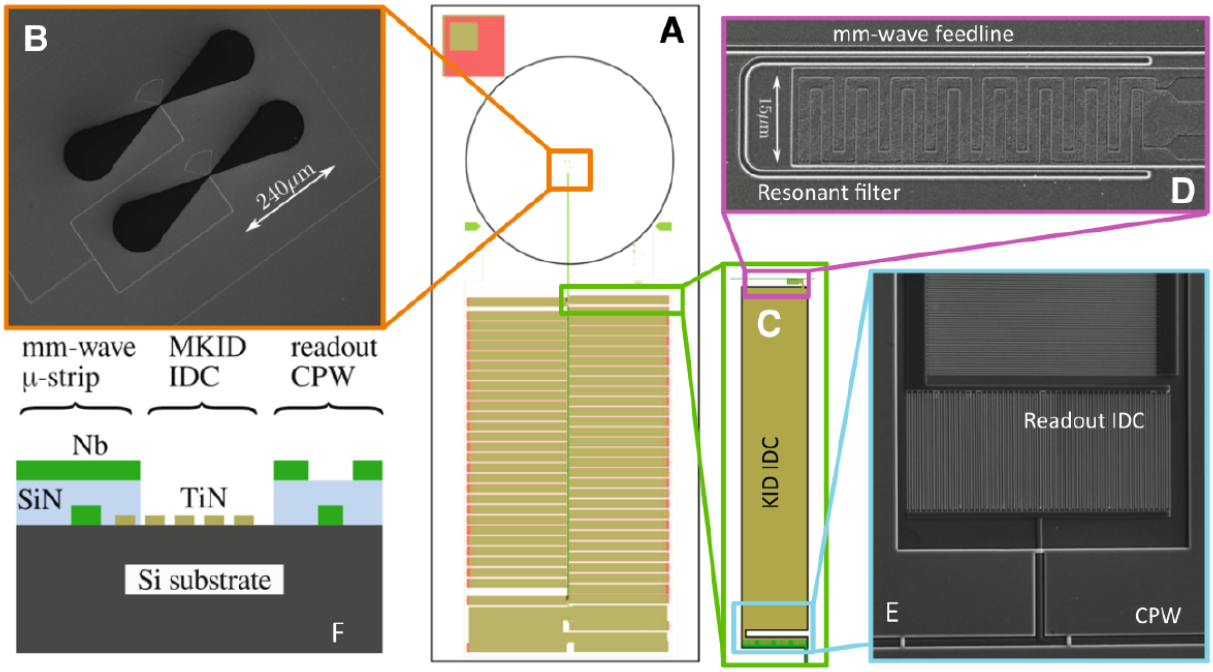

Figure 3.2: *A* The overall die layout of the 50-channel device. The large circle is the lens footprint with the antenna in the center and the mm-wave feedline going vertically down. *B* Dual slot bowtie antenna. *C* Laoyout of the KID. The large long rectangular block is the KID interdigitated capacitor (IDC), and smaller readout IDC towards the bottom. The mm-wave filter and KID inductor are tiny at toward the top right of this part of the image. *D* Zoomed micrograph of the mm-wave feedline, half wave filter bank resonator, and KID inductor meander within the filter. Note that these features seen are in relief, as they are under dielectric and ground plane. *E* Readout IDC and microwave coplanar waveguide readout line. *F* Fabrication layer stackup of the device. The KID inductors are TiN, but all other conductors are Nb. The dielectric material is silicon nitride (SiN).

The reason to pack them so tightly and at low frequency is that the readout bandwidth is limited to 512 MHz. This led to some constraints on the physical design. First, with fixed ultra low volume inductors of 2.6 μm^3 and $\lambda/4$ spacing, the KID capacitors become extremely long and narrow. This design led to stray capacitence and inductance in the long rails of the KID IDC. More importantly, these devices had abnormally low optical response, a couple order of magnitude lower than had been previously measured in the 50 channel prototype arrays. The issue was suspected to be the lensed antenna system coupling, as it was demonstrated that there was some broadband response bypassing the filterbank.

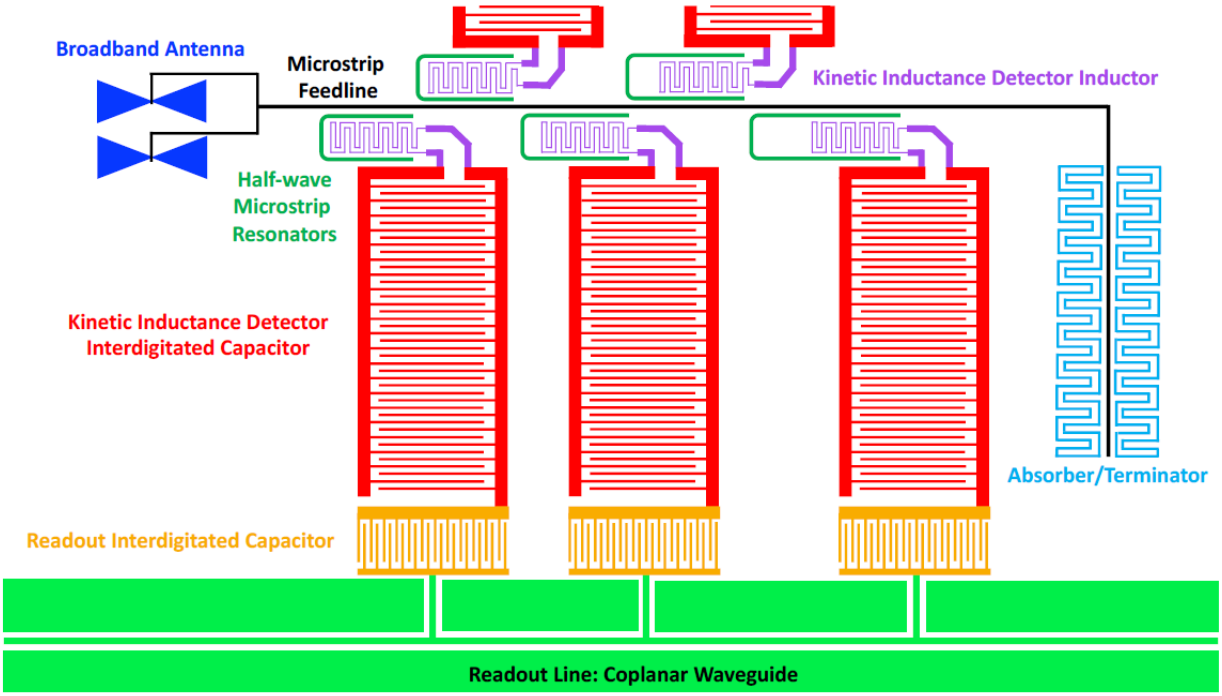

Figure 3.3: Cartoon depiction of the SuperSpec device chip layout (Figure from Wheeler [2019])

A decision was made to scale back, and generate an intermediate device with 110 channels Redford [2023]. This device would maintain similarity to the more successful 50 channel devices. By utilizing more physical space on the larger 300 channel chip package, it was possible to use wider, shorter capacitors, with $3\lambda/4$ mm-wave spacing. This spacing allowed for the reduction of parasitics that effected the 300 channel design. The lower number of KIDs meant that the KID resonant frequencies could be raised while maintaining a manageable frequency scheduling and staying within the required readout bandwidth. The 110 channel dies end up containing KID resonant frequencies in the 200 to 450 MHz range.

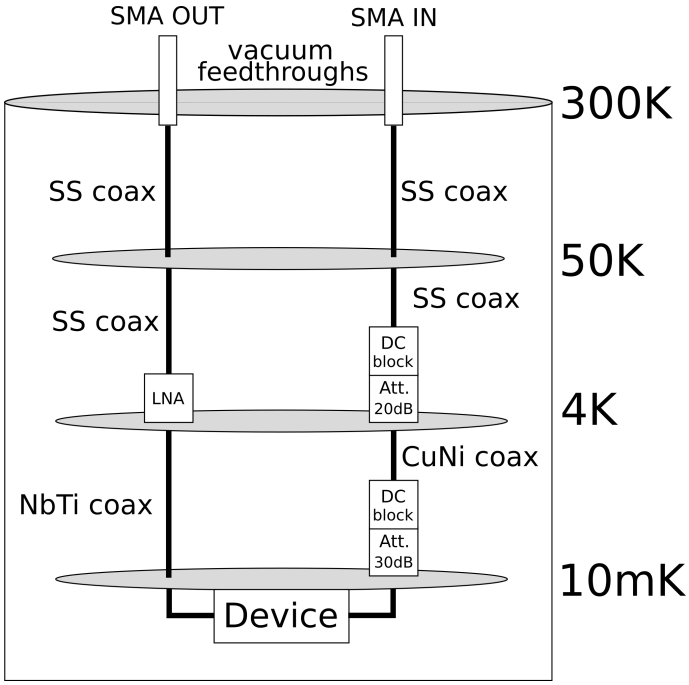

Figure 3.4: Cartoon diagram of the cryogenic hardware in the readout measurement chain within the BlueFor fridge.

3.2 Experimental Setup

3.2.1 Cryostat Testbed

All measurements performed at the University of Chicago were done in a BlueFors He-3 dilution refrigerator[1]. This cryostat can regularly achieve base temperatures of 10 mK in a dark configuration, and down to 30 mK in an optical configuration, where the cold stage has light path open to the room through a set of various filters. With quick cooldown and warm up cycles on the order of 16 to 24 hours, this crysotat has proved to be a workhorse of a test-bed, as we have cooled down over 150 times in my graduate career.

The coaxial connections for measurements are fed in through vacuum tight SMA connectors on input KF-40 flanges, which then connect to stainless steel coaxial cables. These cables feed down to the 4 K flange and connect to cupronickel coaxial cable down to the mix-

1. https://bluefors.com/products/ld-dilution-refrigerator/

ing chamber flange at base temperature. We use flexible copper coaxial cables to connect to the device mounted on the cold plate, and then run superconducting NbTi coaxial cable to up to the cryogenic amplifiers on the 4 K flange. Finally, stainless steel coax takes the signal up and out through the top of the cryostat similar to the input. There are also various interspersed cryogenic attenuators and DC blocks added for filtering, power adjustment, and thermal isolation. We can see a cartoon diagram for this coax measurement chain in Figure 3.4.

3.2.2 Chip Housing

Before mounting, if an optical test will be performed, an AR coated silicon lens is mounted to the backside of the chip using a custom lens mounting jig, seen in Figure 3.5. The chip is secured to the jig via two stainless steel dowel pins that are etched through the wafer and aligned so that the lens is centered on the antenna on-chip. The lens is secured to the chip using GE varnish thinned by adding acetone. A minimal amount of this thinned varnish is applied around the edge of the lens, making sure not to accidentally adhere the chip to the jig itself. If this is the case, a small amount of acetone sprayed at the edges of the lens can loosen the chip enough to dislodge it from the jig, but not entirely dislodge the lens.

The chip is then mounted in a gold plated OHFC box, secured by flexible clips around the edges of the chip. As well, we dab a small dab of rubber cement at the corners of the chip. At one corner, the input and output microwave microstrip radout lines terminate in a Nb pad that use to wirebond from the chip to SMA laucher boards. These copper on duroid launcher boards are a matched transition from microstrip to CPW, and the tip of the SMA connection is soldered to the CPW side of the board. There are additional wirebonds from the top layer groundplane to copper strips mounted and grounded with silver epoxy to the box, which ensures proper grounding.

The box housing includes a snout blackened with liquid eccosorb, which allows for an

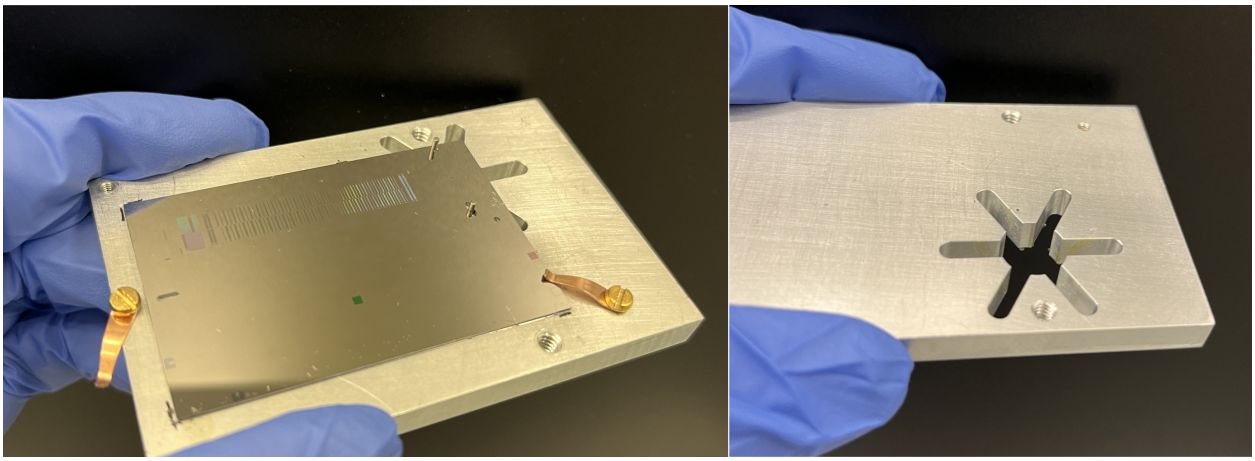

Figure 3.5: *Left* Lens mounting jig with 110 channel device attached. Note 1 mm dowel pin aligning with through etched holes in the chip. *Right* The other side of the jig, where the lens is pushed through the hole, with star pattern cutout to apply thinned GE varnish to the edge of the lens.

approximate 25 degree half beam angle as seen by the antenna beam. Mounted at the top of the snout, we mount copper wire mesh band-pass filters, with cutoffs at 180 and 320 GHz. Some photographs of the box with device mounted can be seen in figure 3.6.

3.2.3 Readout

We employ standard microwave measurements in the work presented here with a vector network analyzer, as well as a homodyne RF transmission measurement setup that we refer to as the single tone measurement setup. This is in contrast to the multi-tone measurement system, which essentially replicates the homodyne measurement for multiple probe tone frequencies simultaneously using a FPGA ROACH-2 readout.

3.2.4 Single Tone Homodyne Measurement

We employed a homodyne measurement scheme to take frequency sweeps as well as measure detector noise spectra. A sinusoidal signal is generated in a synthesizer, which is then split into two identical signals. One half of the signal power is directed into the cryostat, down

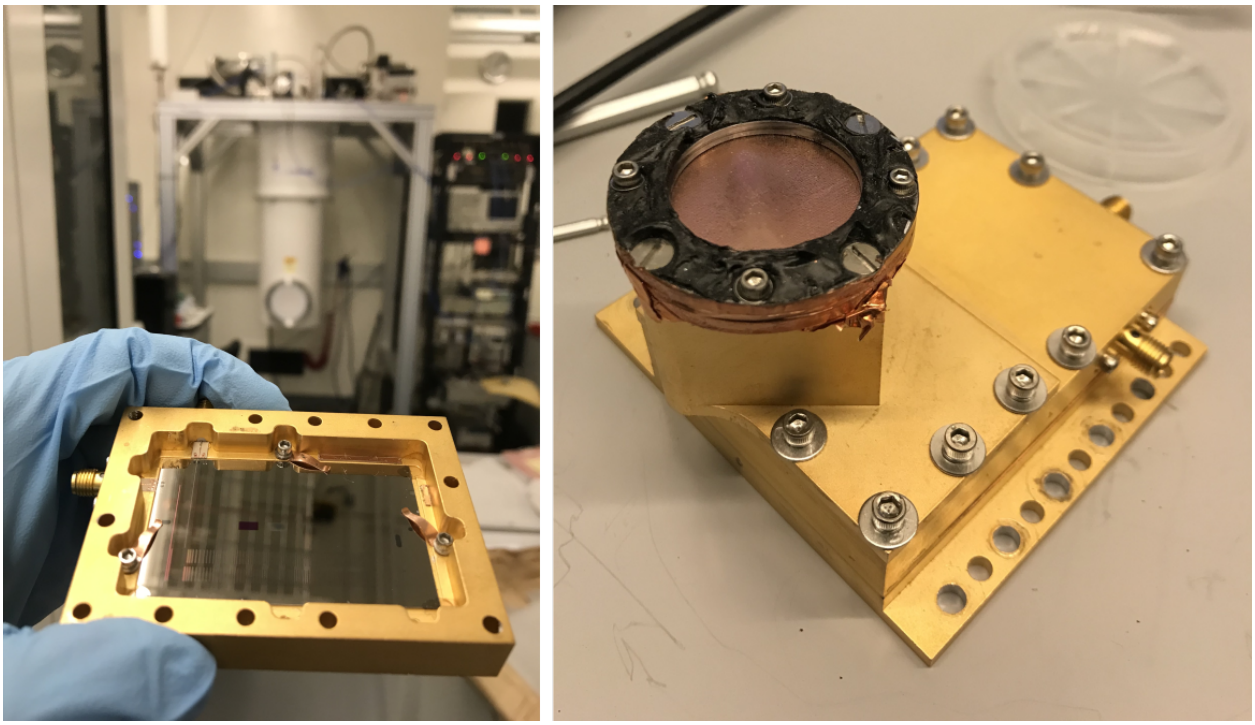

Figure 3.6: *Left* A 100 channel die secured in gold plated Al box with copper clips and rubber cement in the corners. Dilution fridge in optical configuration open to the room is in the background. *Right* The same device with the lid and snout with band-pass defining filters mounted on the tip if the snout.

through a series of coaxial cables and attenuators, and interacts with the detector. The signal is amplified with a high gain low-noise cryogenic amplifier and carried out of the frige. We further provide warm amplification, and then feed the signal into the RF port of a quadrature demodulator. The second half of the tone power is fed into the local oscillator (LO) port of an IQ demodulator. In this demodulator, the sine wave carried through the fridge is demodulated, extracting the in-phase and quadrature compononents I and Q. We do a further stage of low-noise amplification and provide a low pass anti-alias filter cutoff, typically around 200 kHz. The signal is then digitized in an ADC with 1 Gs/s sample rate and recorded digitally. See Figure 3.7 for a visual schematic of this system. In order to take a resonance sweep, the frequency of the synthesizer is varied across the resonator bandwidth to trace out the shape of the resonance. When taking data for measuring noise spectra, the

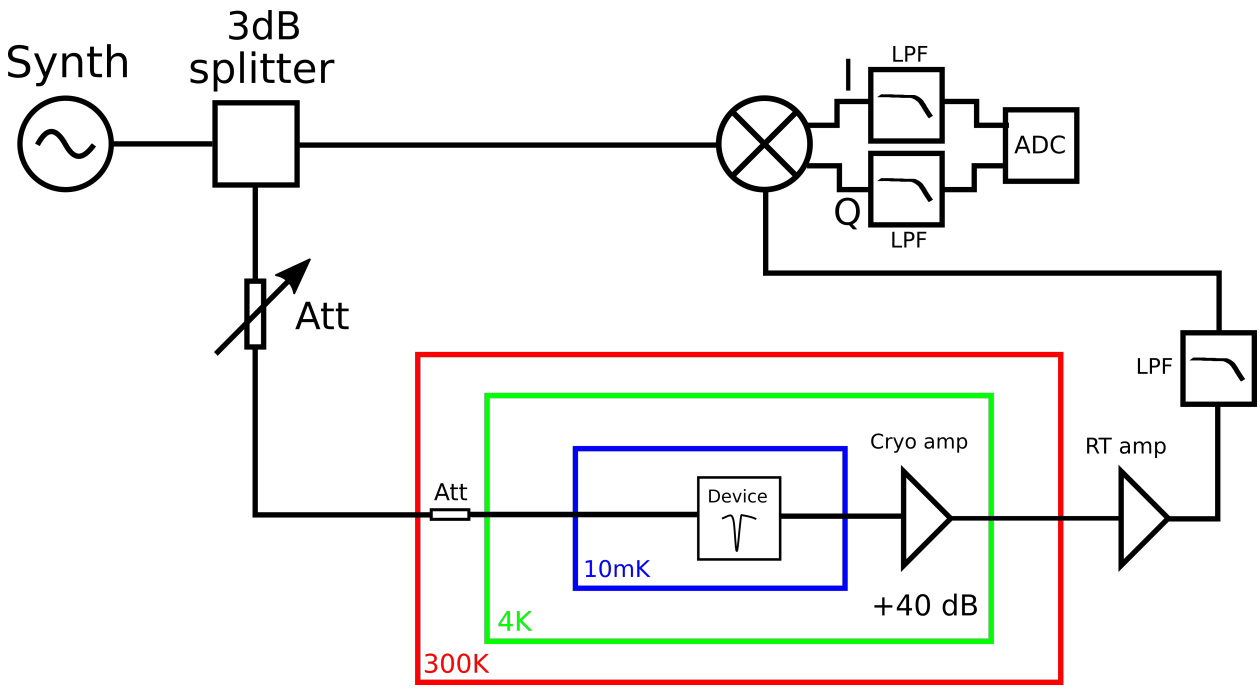

Figure 3.7: Schematic of single tone readout system with standard homodyne microwave measurement. For more detail in RF chain, see Figure 3.4.

synthesizer tone is placed on resonance, and the data timestreams are recorded digitally.

3.2.5 Multi-tone Readout System

When reading out arrays of many hundred or thousands of KIDs, we need the ability to place a tone on each resonance, and measure the phase and amplitude changes of every tone. The system developed to complete this task [Gordon et al., 2016, Gordon, 2019], consists of a ROACH FPGA board along with a DAC/ADC board developed for the MUSIC instrument [Duan et al., 2010]. More detail for this system will be discussed in Chapter 4, but essentially, the DAC generates many tones, which are mixed up to the desired frequencies, which are then mixed back down to baseband and digitized in the ADC. The FPGA board processes the digitized raw waveforms for many simultaneous tones, measuring the phase and amplitude for each tone. We can see this diagrammatically in Figure 3.8.

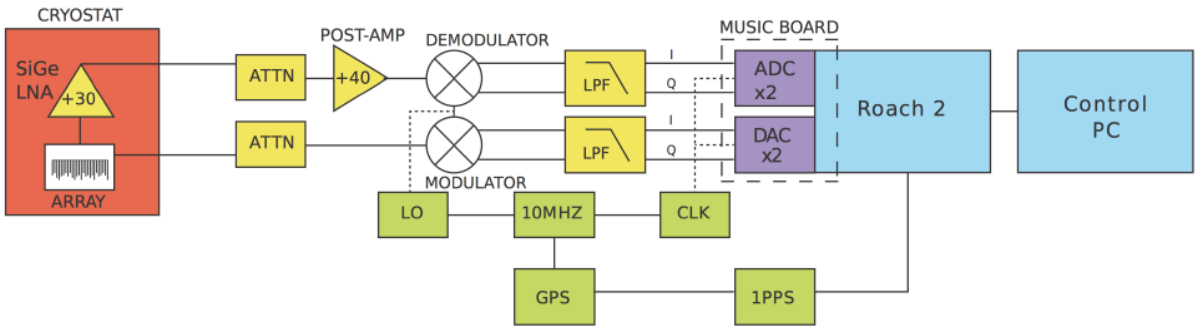

Figure 3.8: Schematic for the ROACH-based multitone readout system. Tones are generated in the DAC at baseband from -256 MHz to 256 MHz, which are then fed through low pass filters. These baseband frequencies are mixed up using a local oscillator (LO) to the desired tone frequencies. The signal is then fed into the fridge with attenuation or amplification as needed. After travelling through the fridge, along with cold amplification, warm amplifiers are also used at room temperature to maximize signal to noise in the ADC. The digitized signal is then processed on the ROACH board, providing the in phase and quadrature (I,Q) components for every tone. (figure from Gordon et al. [2016])

3.2.6 Noise Measurements and Analysis

The following is a description of the general analysis methodology for measuring KID noise spectra. See Gao [2008] for more details.

The purpose of these measurements is to capture features of the device noise, such as the low frequency noise, in-band noise level, and high frequency roll-off. In order to do this, we need to calibrate the detector noise timestream to a physical quantity tied to the resonator. Typically, we do this by converting our detector signal, consisting of $I(t), Q(t)$, to fractional frequency units, $\delta f/f$, as seen in figure 3.9. In order to do this conversion, first a set of frequency sweeps tracking out the shape of the resonance need to be performed. We perform one sweep at low resolution to identify the resonant frequency, which is the highest response point in the resonance. Once that point is identified by calculating the maximum of $\sqrt{dI/df^2 + dQ/df^2}$, which we call the speed of the IQ sweep, we then take a higher resolution sweep centered at the resonant frequency. We then have all the data we need to do our calibration, and can begin streaming with out tone set to the resonant frequency,

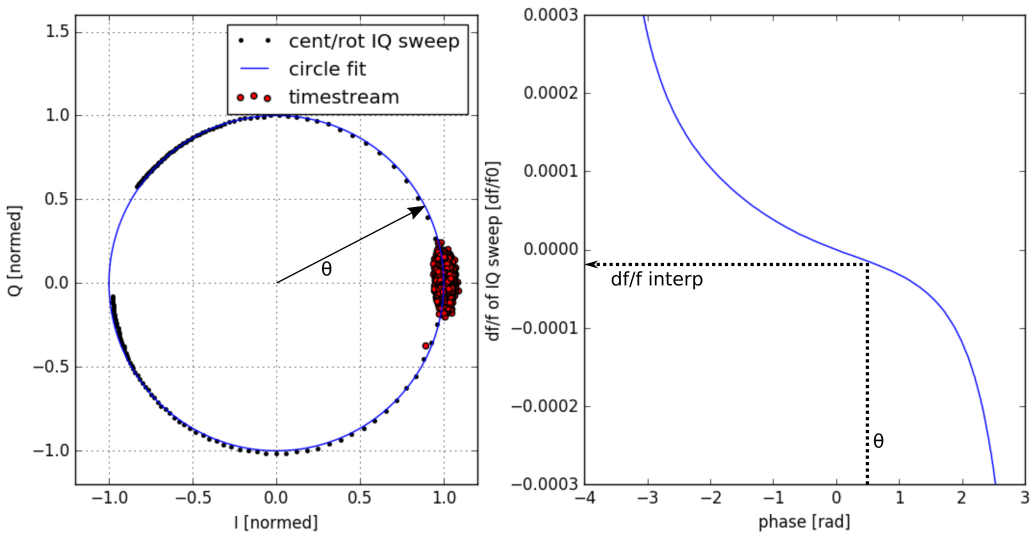

Figure 3.9: *Left* Centered and rotated IQ sweep of KID resonance. Overplotted in red, we see the noise ball from the noise timestream. The calibration to *df* units is done by calculating for a given θ, what is the equivalent *df*, the frequency difference from the resonant frequency according to the sweep. *Right* The phase of the IQ sweep and the equivalent *df*, which is used to interpolate the (I,Q) timestream to a *df* timestream.

recording the $I(t), Q(t)$ timestreams. We repeat the noise timestream process after moving our probe tone frequency off-resonance, typically a couple line widths away from the resonant frequency where no other resonances lie. The purpose of this is to compare to on-resonance and get a measure of the system noise that does not include the device, which is typically dominated by the cryogenic amplifier.

With the single-tone setup, we take two separate noise timestreams with high and low sampling rates, typically 100 kHz and 0.1 MHz. The sampling rate is proportional to the maximum accessible frequency in the power spectral density of the noise, so in order to measure higher frequency noise components, one needs higher sampling rate. But, in the other direction, the time length of the noise timestream determines the frequency resolution and therefore the lowest accessible frequencies in the noise spectrum. So we take the approach of taking low sampling rate, long noise timestreams, and short high frequency timestreams. A long, high sampling rate timestream would be inefficient in terms of data storage and size

for the required low frequency information needed. These noise spectra can later be stitched together. With the multi-tone ROACH system, we are fixed to a sampling rate of 488 Hz, so there is only noise spectrum information below 244 Hz. Note, with the multitone system, the off-resonance noise is taken by changing the LO, which shifts the entire RF comb of frequencies by the same amount. One can also consider for a system like this to add 'blind tones' to the RF comb, which would be tones at frequencies where no resonators are located, and using this as a measurement of the off-resonance noise.

Once the $I(t), Q(t)$ timestreams are converted to $dx = df(t)/f_0$, the noise power spectral density is then computed as the power spectrum of this dx timestream, which is commonly labeled S_{xx}. The noise spectra have a few general features. At low frequency, it is common to see $1/f^n$ behavior that can arise for several reasons such as temperature or voltage fluctations, or even TLS contributions. At high enough frequency, the $1/f$ hits the white noise floor. This frequency sometimes referred to as the 'knee' typically occurs in the range from 0.1 Hz to 10 Hz. This white noise level is dominated by generation-recombination noise, as described in 2.3. Eventually, the white noise rolls off. This can be caused by either the resonator ringdown time or the quasiparticle lifetime depending on which timescale is shorter. For a detector timestream that exhibits white noise above the system noise, usually dominated by the cryogenic amplifier, one will see this white amplifier noise, and finally another rolloff due to anti-alias filters. These last two features can be directly observed in the off-resonance noise spectrum. All of this behavior can be modeled and well fit by

$$S_{xx}(f) = \Big(\frac{A + Bf^{-n}}{1 + (2\pi f\tau)^2} + C\Big)\frac{1}{1 + (2\pi f\tau_{filt})^2} \tag{3.4}$$

and we can see an example of the noise PSD and fit in the *Right* plot of figure 3.12.

3.3 Temperature Response and Film Critical Temperature

One of the most important parameters of the TiN inductor to target correctly is the superconducting transition temperature T_c. The T_c of TiN at high N flow rates is approximately 5.6 K, but we can drop this value significantly by using a low N flow rate while sputter in order to make sub-stoichiometric films. These films where the ratio of Ti:N can vary down to 1:0.6, have lowered transition temperatures. Many of the devices presented in this work have T_c in the range of 0.9 K to 1.2 K.

This method of generating low T_c films also comes with its challenges. It can be difficult to control chip to chip due to variations in fabrication conditions in the Ti sputtering chamber and gas flow rates. The TiN deposition process requires constant re-tuning and characterization if the chamber environment drifts over time. Even more troubling, we observe gradients in T_c at the chip level, where there is a variation based on the physical position of the TiN inductor. This is likely caused by non-uniform nitrogen gas in the sputtering chamber volume, resulting in different levels of Ti:N stoichiometry, and therefore variable T_c.

The measurement of T_c is important, as T_c is linearly related to the superconducting band gap, $\Delta = 1.76 k_B T_c$, which sets the low frequency cut-off for our detectors. For example, a $T_c = 1$ K corresponds to 73 GHz, so that any radiation below that value will not break cooper pairs. As our lower frequency of the observation band is 180 GHz, we are well into the regime where small T_c variation will not be an issue for this particular consideration. As well, the responsivity of the detectors scales with $T_c^{-3.5}$, and is of vital importance to measure and understand for calibration. As well, T_c affects the kinetic inductance L_k, and therefore shifts the resonant frequency, scaling as $\sqrt{T_c}$. This resonant frequency shift causes many resonator collisions, which will be further addressed in 4.2.

We measure T_c by utilizing the fractional frequency shift of the resonance as a function of array temperature. We know from Gao [2008], the fractional frequency shift can be expressed in terms of temperature T, transition temperature T_c, kinetic inductance fraction L_k, and

γ, a parameter related to the surface impedance and conductivity by $\delta Z_s/Z_s = \gamma\delta\sigma/\sigma$.

$$\frac{\delta f}{f} = -\frac{\alpha\gamma}{2}\frac{\delta\sigma_2}{\sigma_2} = -\frac{\alpha\gamma}{2}(\frac{\sigma_2(T)}{\sigma_2(0)} - 1) \tag{3.5}$$

Combining this with equation 2.6 for σ_2, we calculate the fractional frequency shift as a function of temperature T

$$\frac{\delta f}{f} = -\frac{\alpha\gamma}{2}(1 - \sqrt{\frac{2\pi k_B T}{\Delta_0}}e^{-\frac{\Delta_0}{k_B T}} - 2e^{-\frac{\Delta_0}{k_B T}}e^{-\xi}I_0(\xi)) \tag{3.6}$$

with $\Delta_0 = 1.76k_BT_c$, the superconducting band gap energy at zero temperature, $\xi = \frac{hf}{2k_BT}$, and I_0 is the zeroth order Bessel function of the first kind. We also must identify γ, which as in section 2.1.3 takes different values based on the thickness of the film and the electron mean free path. In this case, with TiN film thicknesses of 20 nm, we are in the thin film where the penetration depth large compared to the superconductor thickness, and the current distribution is approximately constant, and we can calculate $\gamma = 1$. The kinetic inductance fraction $\alpha = L_k/L_{tot}$, is the fraction of the inductance from the kinetic inductance effect compared to the total inductance in the circuit including geometric inductance. The kinetic inductance of TiN is large, so that alpha is approximately unity. An example of the temperature sweeps for a single resonator along with the fit function for T_c are seen in Figure 3.10

3.3.1 Critical Temperature Variation in TiN KIDs

There is a clear systematic cause of resonator collisions when we plot the measured T_c against the KID readout frequency. As in Figure 3.11, we see a clear trend of the four readout banks with lower T_c at the end of the bank. This is consistent with a T_c variation with position on the wafer. There variation is suspected to arise from the nitrogen concentration in the sputter chamber varying with position during the deposition. Furthermore, it has been

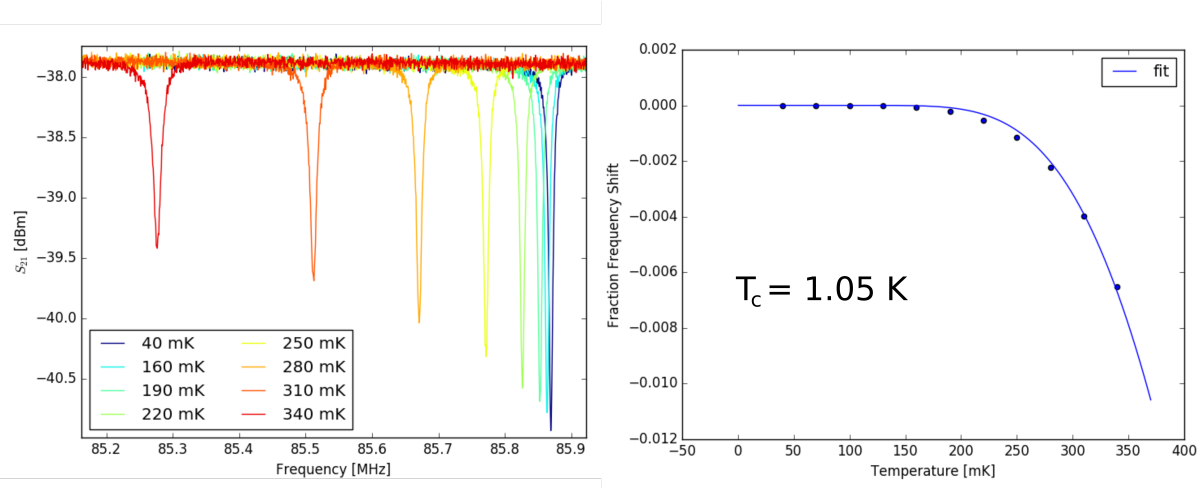

Figure 3.10: *Left* Frequency sweep of resonator as a function of stage temperature. We see the resonant frequency move lower as the quasiparticle density changes due to the rise in stage temperature. *Right* Extracted from the left plot, we fit the resonant frequency as function of temperature to estimate the film T_c.

shown [Wheeler, 2019] that this positional variation is strongest in the radial direction when accounting for the position of the chip on the wafer before dicing.

This effect has a few consequences for the performance of these detectors. Most importantly, the lower T_c equates to higher kinetic inductance, meaning the resonant frequency of the KID will be shifted lower as the T_c drops. The designed mask does not take into account this variation, and this causes the resonators at the end of the bank with low T_c to bunch together and collide. The process through capacitor trimming is discussed in section 4.2.

3.4 Low Temperature Noise Performance

The results presented in the following text can be also seen in McGeehan et al. [2018]. The device under test was one of a 55-channel prototype devices tested in a dark setting with a helium dilution refrigerator, allowing for measurements at temperatures well below the intended operating temperature of 210 mK in the deployment cryostat. The data presented were taken with the standard homodyne single-tone measurement discussed in previous sections.

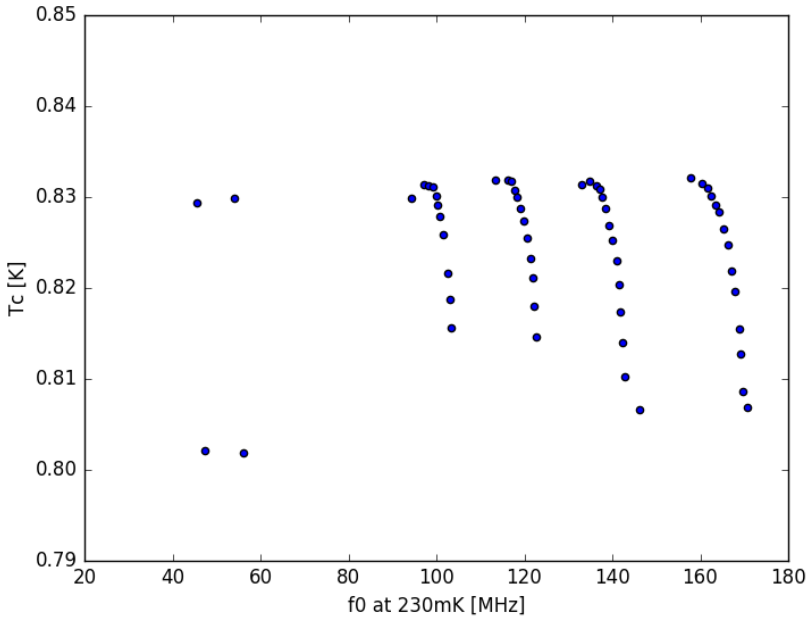

Figure 3.11: KID T_c plotted as a function of KID resonant frequency f_0. There are four clear readout banks, where within each bank, the resonators span across the chip. The KIDs close in T_c are physically adjacent on the chip. The four lowest frequency KIDs are broadband absorbers at the beginning and end of the mm-wave feedline.

3.4.1 Dark Detector Noise and NEP

We report the measurement of dark detector noise and estimated noise equivalent power (NEP) for this $T_c = 0.93\,\mathrm{K}$ detector. We performed standard detector noise measurements with a single-tone system.

We see in Figure 3.12 an example of the resonator transfer function in the complex plane, along with the noise timestream points and corresponding power spectral density (PSD) of the fractional frequency noise. The demodulated noise timestream was mapped tot he equivalent fractional frequency shift in the resonator using the phase-mapping technique as described in section 3.2.6. A Savitsky-Golay filter was used to smooth the high resolution sweep in order to ensure smooth mapping from complex phase to δf. The same procedure was performed for several different resonators across a wide range of device temperature. The power spectral density of the fractional frequency noise was fit in accordance with equation

3.4. In this fitting model, the important parameters we extract are the resonator white noise level, τ the roll-off frequency associated with the resonator ring-down time, and a low frequency noise term.

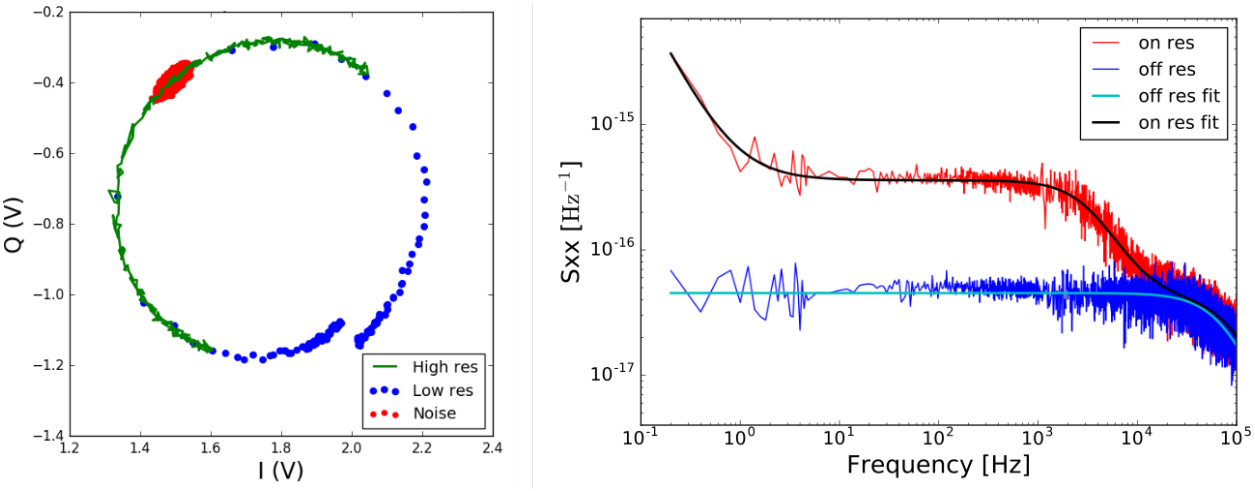

Figure 3.12: *Left* The low resolution and high resolution sweeps in blue and green. The low resolution sweep is taken to lock onto the resonant frequency, and the high resolution takes a much finer frequency spacing sweep close to the resonant frequency for I, Q to df calibration. In red, we see the noise ball. Note that the asymmetric shape in the dissipation direction (tangent to the circle) indicates we are observing device noise. *Right* The red on-resonance power spectral density corresponds to the timestream noise ball in the left plot. We see the characteristic $1/f$, white noise, and roll off. The blue PSD is the off-resonance noise, which is a noise timestream taken well off-resonance to measure the system noise, dominated by the cold amplifier.

Because the Q of the resonator drops as a function of increasing temperature, the drive power of the probe tone needs tuned for each temperature point, increasing at higher temperatures. This is especially true for SuperSpec detectors because of the very small detector volume of $2.6\,\mu\text{m}^3$, and the resultant bifurcation power on the order of $-110\,\text{dBm}$ at the detector. Then, if the probe tone power is too low, the noise of the cold amplifier dominates. There is a window in between the upper limit of bifurcation and lower limit of amplifier noise where the device noise can be seen. This power window can be small for these devices, sometimes on the order of 3 dB to 5 dB, and therefor careful tuning at each temperature is imperative to see the device noise above the off-resonance noise floor. This is achieved in the

single-tone case by, at each temperature, sweeping around each resonance for several tone powers. We then identify the power that corresponds to bifurcation and then dial back by 1 dBm, which was proven to be sufficient to maintain signal above the cold amplifier noise.

The noise measurements were taken in a dark configuration, but this device was also measured in a different cryostat in an optical configuration. That data set determined that the median responsivity for optically -coupled channels was $1.4 \times 10^{10}\,\mathrm{W}^{-1}$ [Wheeler et al., 2018], but it should be that the stage temperature for this measurement was at temperatures greater than 210 mK. If we adopt this value in order to get an estimate of the device noise equivalent power corresponding to our measured dark device noise. We note that the optical responsivity was measured at a higher base temperature, however there is evidence that it remains approximately constant with temperature below 230 mK, as see in the inset of Figure 3.13. Combined with the dark data presented, we report a median NEP at operating temperature of 210 mK of $2.7 \times 10^{-18}\,\mathrm{WHz}^{-1/2}$. The device noise is below the expected photon noise level for expected SuperSpec requirements as a low or moderate resolution ground-based instrument at an excellent mm-wave observing site.

For this optimized small-volume KID under dark conditions, we expect the white noise near the operating temperature to be dominated by fluctuations in the equilibrium quasiparticle density, that is, the generation and recombination (GR) of thermal quasiparticles. We fit $S_{xx,0}$, which accounts for other unknown sources of noise that are independent of temperature and optical loading. The white-noise level as a function of temperature is shown in Figure 3.13, along with the fit to the model given by

$$\mathrm{S}_{xx} = \left(\frac{\alpha_k \gamma S_2}{4N_0\Delta_0}\right)^2 \frac{4n_{qp}\tau_{qp}}{V}\left(1 + \frac{\tau_{qp}}{\tau_{max}}\right) + S_{xx,0}, \tag{3.7}$$

where S_2 is the Mattis-Bardeen conductivity factor and $\tau_{qp} = \tau_{max}(1 + n_{qp}/n^*)^{-1}$, with τ_{max} being a maximum observed quasiparticle lifetime and n^* being the crossover density [Zmuidzinas, 2012]. We assume the values $\alpha_k = 1$, $\gamma = 1$, which are appropriate for thin film

TiN. In addition, we fix $\Delta_0 = 1.76 k_B T_c$, $N_0 = 4.0 \times 10^{10}\,\mu\text{m}^{-3}\text{eV}^{-1}$, and $V = 2.6\,\mu\text{m}^3$.

In our fit, we find n^* is unbounded, meaning $\tau_{qp} = \tau_{max}$ across the entire temperature range with $\tau_{max} = 5.2\,\mu\text{s}$. These results are consistent with previous measurements, but extend the model to lower temperatures [Wheeler et al., 2016]. Furthermore, we find $S_{xx,0} = 9.7 \times 10^{-17}\,\text{Hz}^{-1}$, which corresponds to a low temperature NEP of $7.0 \times 10^{-19}\,\text{WHz}^{-1/2}$.

Next, we performed the same measurements, but with lossy Eccosorb CR-110 filters on the coaxial input and output of the detector at the cold stage. As seen in Figure 3.13 in red, these filters reduced the white noise floor down to $S_{xx,0} = 6.3 \times 10^{-17}\,\text{Hz}^{-1}$, corresponding to a low temperature NEP of $5.7 \times 10^{-19}\,\text{WHz}^{-1/2}$

The source of the white noise floor at low temperatures is currently unknown and further work is being done to done to investigate its origin. The effects of the Eccosorb filter suggest that possible contributions are chip heating, microwave heating, two level system noise or excess photon leaks. For example, the equivalent optical load that would give rise to the observed $S_{xx,0}$ is small. Considering light above the pair-breaking gap energy of 68 GHz, photons on the order of ν =100 GHz with 1 fW loading would account for the excess noise.

3.5 Optical Characterization

3.5.1 Filter Bank Performance

We measure the spectral profiles of the filter bank through Fourier transform spectroscopy (FTS) [Lambert and Richards, 1978]. The FTS used in these measurements has throw of 140 cm, with a PID controlled filament source that can reach 1100 K, and generates a collimated beam. This beam can be further focused by HDPE lenses if needed, but the source is bright and SuperSpec detectors are responsive, so typically signal to noise is high enough to see the interferogram above the noise by eye without the need for a focusing lens.

Before we take FTS data, it is important to set up the system to make the data taking as

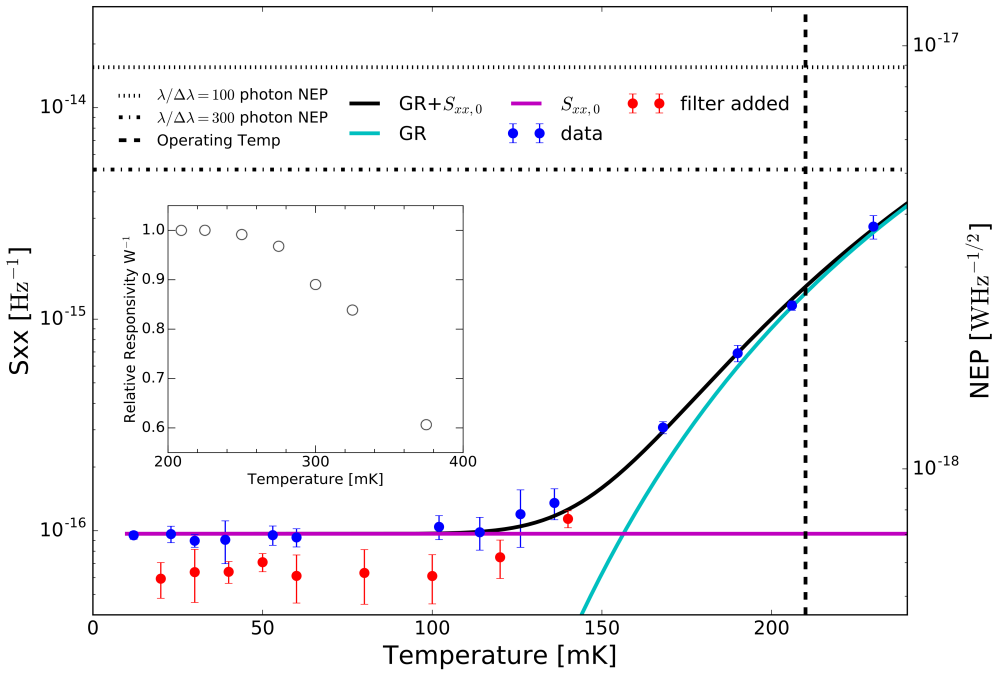

Figure 3.13: The blue data are the white noise level as a function of temperature averaged over three KIDs in the 116 MHz to 141 MHz range. A fit was performed on each data set with the noise model as in equation 3.7. The fit finds $\tau_{qp} = 5.1\,\mu$sec and $S_{xx,0} = 9.7\times10^{-17}\,\mathrm{Hz}^{-1}$. The right axis NEP is calculated by taking the median KID responsivity at 210 mK, as seen in the inset figure. We see the responsivity as a function of temperature flattens to an approximately constant value at temperatures below this value. At the operating temperature of 210 mK, we find an NEP of $2.7\times10^{-18}\,\mathrm{WHz}^{-1/2}$, which then decreases to an approximately constant value of $7.0\times10^{-19}\,\mathrm{WHz}^{-1/2}$ below 130 mK. The noise measurements were repeated with added low-pass Eccosorb CR-110 filters on the coaxial input and output of the detector. The filters reduced the noise at low temperatures, resulting in a noise floor of $S_{xx,0} = 6.3\times10^{-17}\,\mathrm{Hz}^{-1}$ and an NEP of $5.7\times10^{-19}\,\mathrm{WHz}^{-1/2}$.

easy as possible. With the multi-tone system, it is sometimes difficult to track KIDs because they move out of the sweep window under large load changes. For this reason, we attempt to get an accurate list of resonant frequencies with a VNA sweep under an equivalent load to what the detector will experience with the mirror at one extremum of the throw. We have found that for our system, a beam filling room temperature load is works well. This list of resonant frequencies, that we call a toneslist, is then fed into the multi-tone system. With the FTS warmed up (typically 1000 °C), roughly aligned, and mirror sent to one end of the

throw, a sweep with the multitone system should have all of the KIDs roughly land within a window of 500 kHz of the toneslist values.

At this point, the power tuning must occur, where we optimize the individual tone power for each resonator. This ensures that we are driving each resonator with sufficient power to see the detector signal above the amplifier noise floor. For full detail about the ROACH readout system refer to Chapter 4.

Once the powers have been set and tone frequency placement aligned to the resonant frequency of each KID, we are ready to start streaming. We begin recording the data and saving to disk the IQ timestreams from each tone. After streaming has started, we begin sweeping the FTS mirror, typically with scan speeds at around 1 mm/sec across a throw of 150 mm. Several scans are taken in order to average them in analysis for better signal to noise. Note that in this system, we are not simultaneously recording the mirror position, but log the endpoint times of the sweeps, and align using python timestamps in the mirror control script and the ROACH system.

In Figure 3.14, we see a resultant interferogram from a FTS sweep. The mirror position has been converted into optical path difference (*opd*), with $opd = 0$ corresponding to the white light fringe where the two arms of the interferometer are equal in length and therefore all frequencies are constructively interfering incident on the detector. The small scale oscillations correspond to the center frequency of the mm-wave resonance, with a broadband envelope across the whole mirror path. The IQ timestreams have been converted to df units as previously described, and the interferogram was baseline removed by fitting a high order polynomial. The spectrum associated with this interferogram is calculated by taking the magnitude of the Fourier of these data. Note that this is only one interferogram, but one interferogram is generated for each tone or resonator in the measurement simultaneously. On point to note it that for this FTS system, we are limited by the throw, that is the total length of the mirror sweep. We can relate the resolution of the FTS of the resultant spectrum

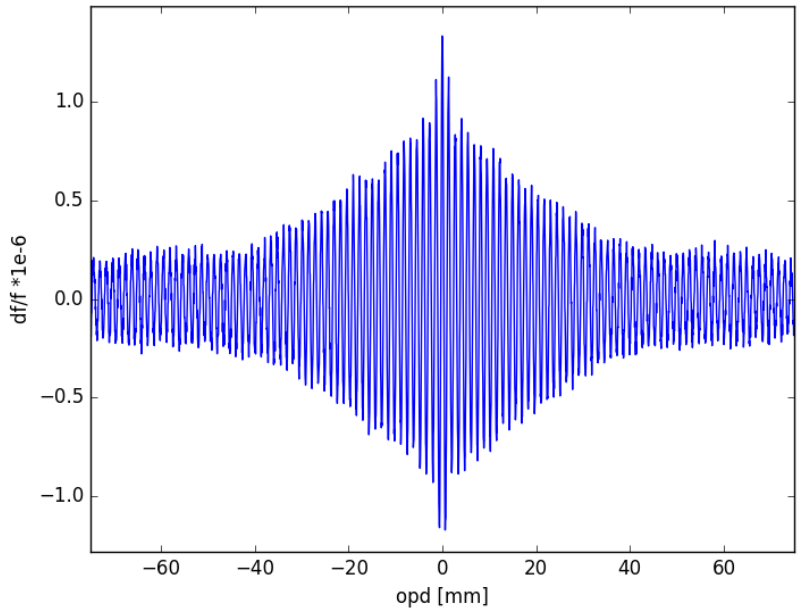

Figure 3.14: Interferogram of single spectral channel from FTS measurement. The IQ timestream has been converted to df units, and the mirror position converted to optical path difference (opd). We can clearly see the small scale sinusoidal behavior, corresponding to a narrow spectral response, as well as a prominent white light fringe where the opd is zero.

by the minimum frequency separation between data points as

$$\Delta\nu_{FTS} = \frac{c}{2(OPD_{max} - OPD_{min})} \tag{3.8}$$

which in this case is approximately $\Delta\nu_{FTS} \approx 1\,\text{GHz}$. We are limited in the ability to resolve the spectral profiles with this system, as for a spectrometer resolution of $R = 300$, we get a full-width half-max of $\Delta\omega_{FWHM} = \omega/Q = 0.75\,\text{GHz}$, meaning we only see power above the noise in several frequency bins for our expected resolution

In Figure 3.15, we can see the sweep of the resonator in S_{21} with the vertical . The resulting spectrum calculated from the interferogram is plotted on the right hand side of the figure, with an overplotted Lorentzian fit, giving $Q_{mm,tot} \approx 275$, centered at 224.8 GHz.

In Figure 3.16 [Karkare et al., 2020], we see a full filter bank spectrum of a 50 channel

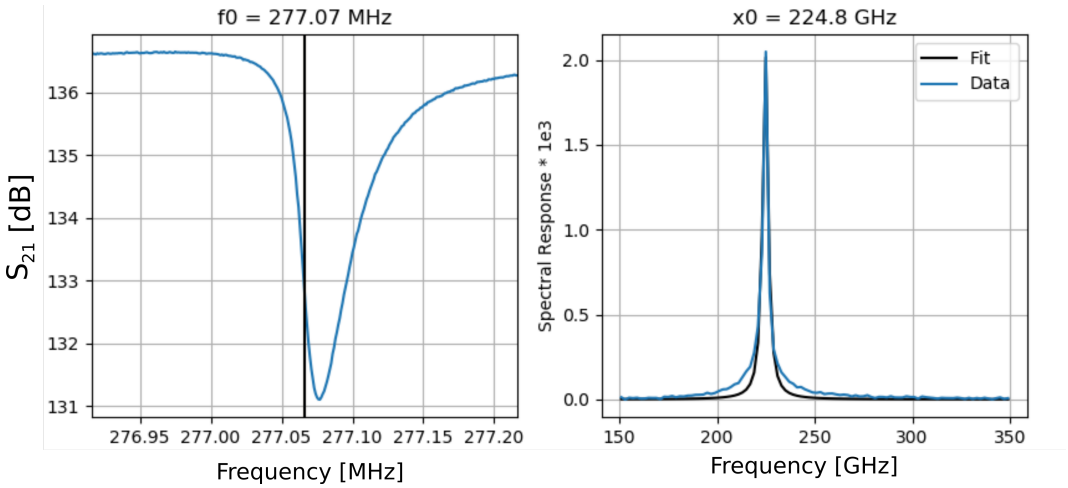

Figure 3.15: *Left* S_{21} sweep of the KID, with vertical line representing the placement of the readout tone. *Right* Resultant spectrum from FTS sweep with Lorentzian fit.

device. The FWHM for these profiles is approximately 1 GHz corresponding to $R \approx 275$. One can also note the interweaving of the KID vs mm-wave resonances by the color of the spectrum curve mapping to a KID readout frequency.

3.5.2 Optical Response and NEP

We quantify the detector responsivity by exposing the detector to a known blackbody source, which have well describes optical load. The simplest way to perform this measurement is what we refer to as a hot-cold measurement. With the cryostat in the optical configuration where the detector is open to the room through a series of filters, shaders, we place a beam-filling load in front of the window. This standardized load is a styrofoam bucket, transparent in the mm-wave, which has several layers of eccosorb foam that acts as the temperature load. This bucket can be kep at room temperature to serve as a standard 300 K load, or can be filled with liquid nitrogen (LN2) at 77 K.

A standard VNA sweep is performed across all the resonators at both temperature loads. It is ideal to perform the 300 K sweeps first, as it can take several hours for the LN2 to evaporate and the eccosorb to come back to room temperature. If the detectors are optically

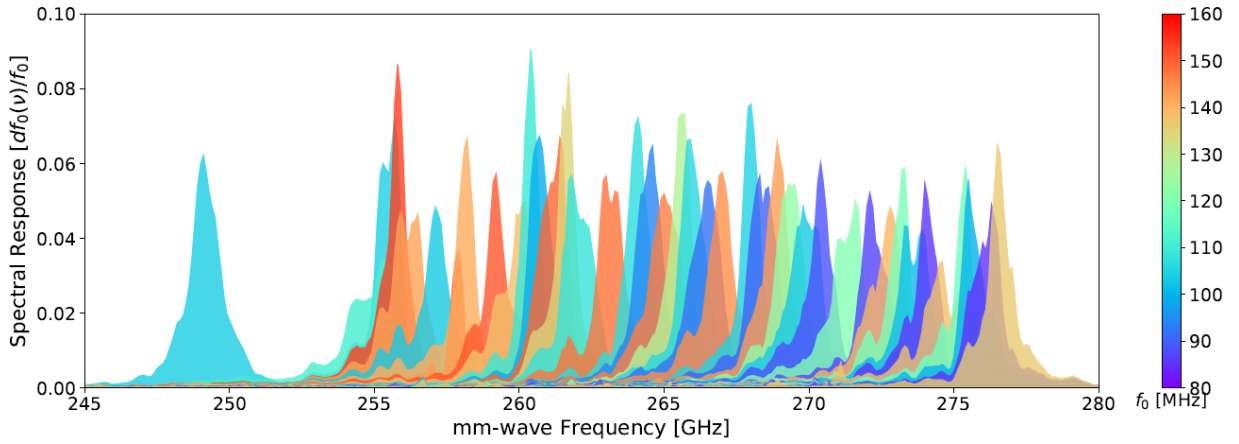

Figure 3.16: Full filter bank spectrum of a 50 channel die. The colorbar refers to the readout frequency of the KID coupled to the filter channel represented in the spectrum. These devices have an approximate $Q_c = 325$ and $Q_i = 225$ with $Q_{loss} = 465$. (Karkare et al., 2020)

responsive, one should see that the resonant frequency shifts between the two temperature loads. From this change in resonant frequency, we can calculate a responsivity, a change in fractional frequency per Watt by calculating the limit $h\nu \ll k_b T$, then the total black body power given detector bandwidth $\Delta\nu$, is given by $P = \eta\lambda^2 B(\nu, T)T\Delta\nu$, where $\eta = 0.65$ accounts for optical filter efficiencies between the source and the device. The detector bandwidth $\Delta\nu$ is determined directly from the detector profiles

$$\Delta\nu = \frac{\int (S(\nu)d\nu)^2}{\int S^2(\nu)d\nu} \tag{3.9}$$

The responsivity R_{det} is interpreted as the slope of f_0 versus P relationship such that given a change in resonant frequency Δf_0 we calculate the responsivity as

$$R_{det} = \frac{\Delta f_0}{f_0} k_b \Delta T \Delta\nu \tag{3.10}$$

where $\Delta T = 298 - 77 = 221\,\mathrm{K}$, The bandwidth is derived from FTS measurements previously described.

In Figure 3.17, we can see the hot/cold measurement for a 110 channel device, showing

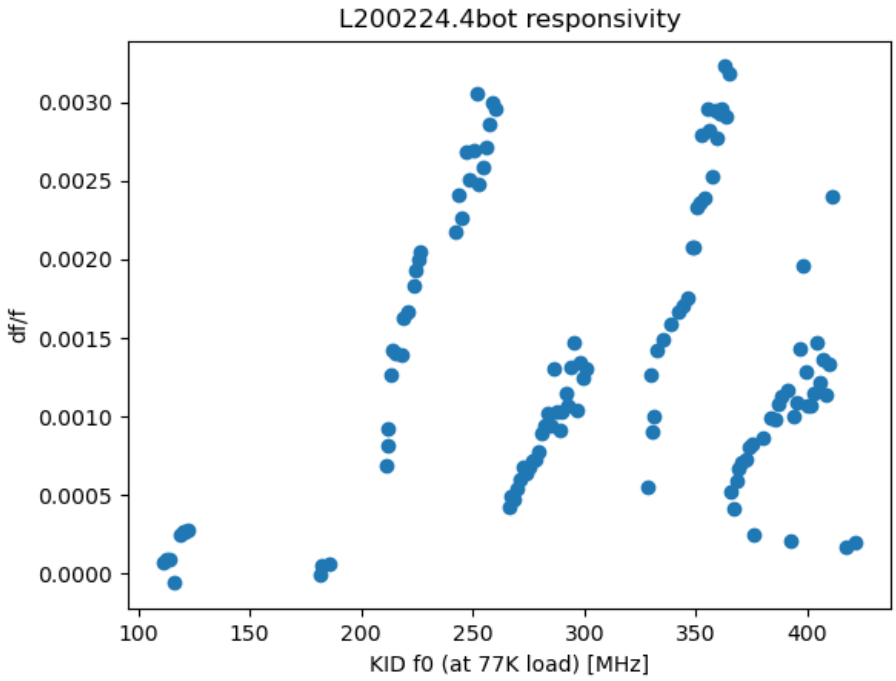

Figure 3.17: Hot/Cold responsivity measurement with beam filling Eccosorb load at room temperature and LN2 temperature. We can see the four resonator banks clearly as the four separate groupings of data points in steep vertical lines. The 'end' of the banks are closest to the antenna, and therefore experience less dielectric loss, so that their response is larger. As well, the T_c variation across the chip worsens this responsivity discrepancy. Furthermore due to a mask error, the filter bank order was incorrect in a way that caused banks B and D to receive less power than A and C.

the measured $\Delta f_0/f_0$ as used in equation 3.10. We observe a variation in frequency response across each individual bank, which can be accounted for by dielectric losses ($Q_{loss} \approx 465$) along the mm-wave feedline as well as T_c variation discussed previously. Additionally, due to an unfortunate mask error, two of the banks, second and fourth, were slightly misaligned with respect to the feedline, and thus received less power than the other two banks.

As in Karkare et al. [2020], we can also perform this responsivity measurement with a cryogenic blackbody. The idea of this measuremnt is equivalent to the hot/cold, but we are using a cryogenic load that is thermally controlled via a heater while it is weakly thermally in contact with the 4 K temperature stage within the fridge. Changing the temperature of

this blackbody load, we can measure more points along the df_0/f_0 vs P curve to linearly fit for the response, as seen in Figure 3.18.

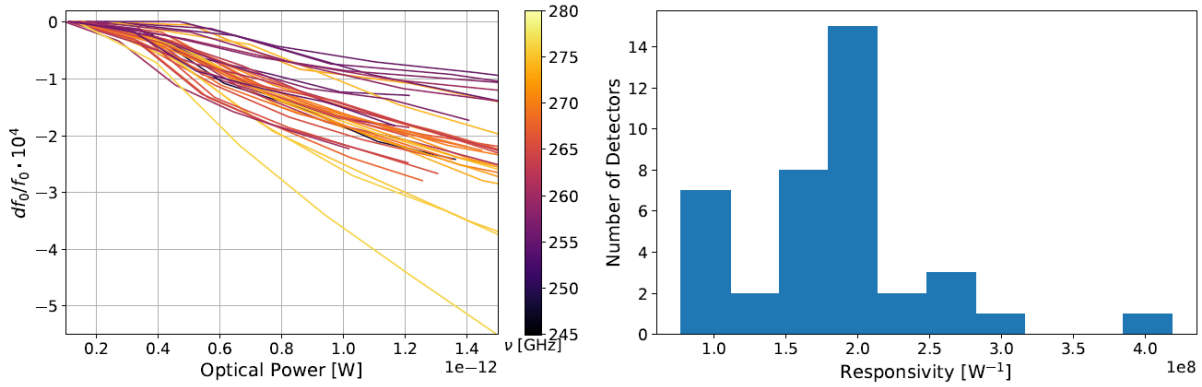

Figure 3.18: *Left* Fractional frequency shift vs optical load using a cryogenic blackbody source, with the curve color indicating the mm-wave filter resonant frequency. *Right* Histogram of calculated responsivity from left hand plot. [Karkare et al., 2020]

Finally, using the measured GR noise level S_{xx} under some defined optical load, we can combine the responsivity and noise measurements to arrive at the noise equivalent power calculation,

$$NEP = \frac{\sqrt{S_{xx}}}{R_{det}} \tag{3.11}$$

For the 50 channel device data presented here, we observe the white noise levels $S_{xx} = (3.4 \pm 1.3) \times 10^{-16}$−1 at 0.7×10^{-12} W loading and $S_{xx} = (4.2 \pm 1.6) \times 10^{-16}$−1 at 1.2×10^{-12} W loading, with the distribution of measured NEP seen in Figure 3.19.

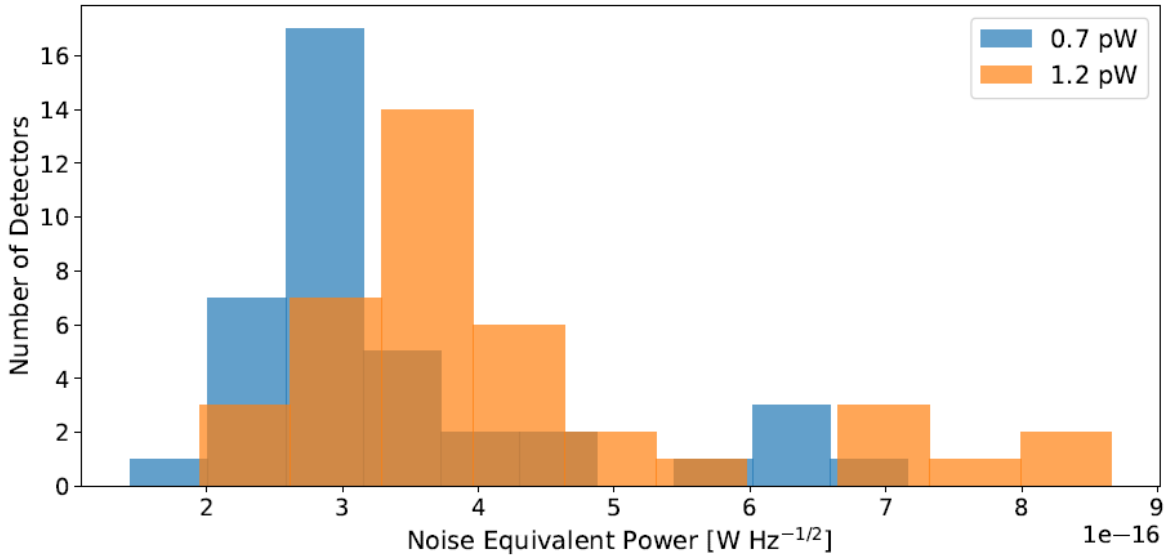

Figure 3.19: Histogram of device NEPs for a 50 channel device. These NEPs include losses from the lens, antenna, and filterbank. [Karkare et al., 2020]

CHAPTER 4

PROGRESS TOWARDS SUPERSPEC DEPLOYMENT

4.1 KID Readout Development

One of the main advantages to using KIDs is their innate frequency multiplexing capability. To read out hundreds, or even thousands, of KIDs at the same time, one must develop hardware systems in order to simultaneously perform these amplitude and phase measurements on every KID. The approach taken here is to create a digital waveform containing the summed contribution of many tones, then generate this analog signal in a DAC to be fed into the cryostat. Then after the tones have interacted with the detector, they are fed back into an ADC and then processed in an FPGA system to do the digital demodulation and extract the phase and amplitude corresponding to each tone. SuperSpec has worked with BLAST-TNG [Gordon et al., 2016] and MUSCAT [Rowe et al., 2023, Castillo-Dominguez et al., 2018], who developed a readout scheme using the CASPER Reconfigurable Open Architecture Computing Hardware (ROACH-2). The Virtex-6 FPGA based board allows for the high-speed processing, which is interfaced with a MUSIC based ADC/DAC board [Duan et al., 2010] used in the first generation of the ROACH.

The DAC/ADC has limited bandwidth, only able to generate tones in the range from $-256\,\mathrm{MHz}$ to $256\,\mathrm{MHz}$, where the sign is referring π phase shift of the tone, for a total effective bandwidth of $512\,\mathrm{MHz}$. This full bandwidth can be utilized by further modulating the output of the DAC with a standard microwave modulator[1]. For example, if our detectors lie in the 400 to $500\,\mathrm{MHz}$ range, we can generate base band tones from 100 to $200\,\mathrm{MHz}$ and mix up in the modulator using an LO set at $300\,\mathrm{MHz}$, and then demodulate by the same frequency.

The DAC can generate waveforms with a 1 GS/s sample rate, and the ADC then captures

1. Polyphase AM0350A, AD0105B

the returned signal at 550 MS/s, which is later channelized and downsampled in firmware to 488 Hz. This downsampled rate of 488 Hz is our set sample rate for all testing and deployment, barring some fundamental change in the firmware. The system in total allows for up to 1024 filter bank channels to be measured simultaneously.

4.1.1 pcp Operation

The pcp[2] readout system is a Python-based ROACH control software suite for large arrays of KIDs. The code was largely developed by Pete Barry (Uchicago/Cardiff) and Sam Rowe (Cardiff), while building on old scripts and interfaces written by Sam Gordon (ASU). The purpose of this software is to develop software routines for use in on-sky observations, as well as lab testing and characterization. The software supports the implementation of multiple ROACH subsystems running in parallel. At deployment, SuperSpec will taking deploy 6 devices, requiring the readout and control of 6 ROACH systems. In the following sections, we will describe some of the important functionalities of the software.

Sweeping and Tuning

One of the fundamental files for taking measurements is the toneslist which defines the number of tones, their identifiers, and other parameters. These toneslists are generated typically from a standard VNA sweep, and using scripts for automated resonator identification. We are able to indicate particular tone power by directly modifying the amplitude of the corresponding frequency component in the digitally generated waveform, as well as other flags on particular tones useful for lab testing. The toneslist, if taken at appropriately similar loadings and temperatures, should generate tones where the resonators actually lie within a few linewidths. But, fine-tuning of each tone must be done in order to lock onto the maximally responsive point of each resonator. In order to adjust the tone frequencies, we first perform

2. https://github.com/shirolab/roach2pcp

a resonance sweep. This is done by sweeping the LO in a band around its true frequency, and then matching up the LO frequency with the binned recorded data. In this way, each tone sweeps a fixed bandwidth about the center placement of the original tone, producing the transfer function of the device about each KID.

We have developed a visualization of these sweeps, as we can see in figure 4.1. We can see the phase, amplitude, and speed of the resonance sweep, as well as the sweep plotted in IQ. This visualization allows for quick visual verification that the tuning make sense. The software then automatically calculates the max speed point and saves that information, that can then be written to the FPGA, which adjust the tones to the correct frequencies. One last sweep should be taken once the tones are in the correct place for calibration purposes, as we write the sweep data to file to be attached to any future data timestreams taken with under this calibration.

4.1.2 Readout Power Optimization

When reading out our KID detectors with a probe tone, we set the tone frequency to the resonant frequency of the KID to maximize the response signal. But we also have another tone parameter to optimize, namely, the probe tone power. If this tone power is too low, the resultant signal will be dominated by noise sources, which in most cases will be the cryogenic amplifier located on the 4 K stage. There is also a natural upper limit of this probe tone power that is set by the bifurcation power of the KID. This nonlinear behavior arises from the non-linear kinetic inductance effect, where the kinetic inductance of the superconducting film L_k scale with the current flowing through the film.

We would like to drive the KID as hard as possible without running into bifurcation. We must first consider the output of the ROACH. The total output power of the ROACH is approximately 0 dBm. That power is in principle evenly split amongst the generated tones in the digital waveform, which means that the power per tone decreases as the number of

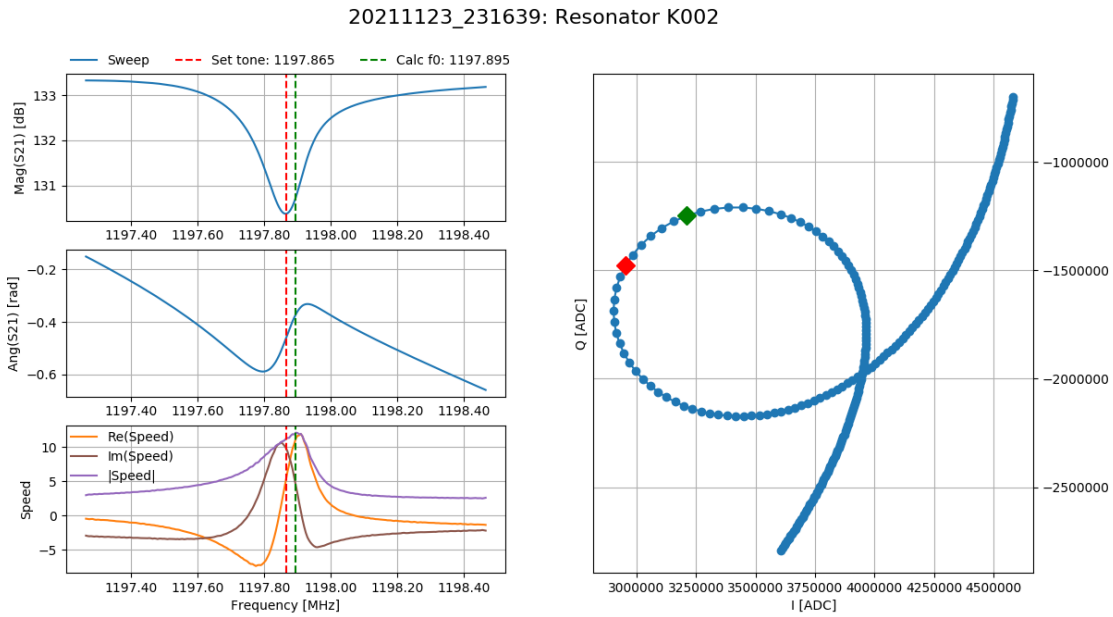

Figure 4.1: Sweep interface for pcp resonator tuning. Having completed an S_{21} sweep measurement, we plot the magnitude, phase, and speed ($\frac{dI}{df}^2 + \frac{dQ}{df}^2$) in order to tune to the maximally responsive point on the resonator. The red vertical dashed line defines the middle of the sweep, and we calculated maximum speed in green and set the tone to this value. This right hand plot shows the sweep in IQ space with the old and to be written tone frequencies in red and green points.

tones increases. It is possible that warm amplification before the cryostat input is needed if there are many resonators. For many of the 110 channel chips presented in this work, this was not necessary, but could be for future larger arrays.

As well, there are non-idealities which result in non-uniform power per tone as a function of baseband frequency, largely due to conversion loss in the modulator. As well, cryostat components have their own frequency dependent losses. All of this adds up to the fact that the power that each KID sees for its readout tone is not individually well-controlled. But, all of these can be calibrated out in a single step, in the following way.

First, we bring the system to a state of high power, where all of the resonators are bifurcated. One can usually verify this by eye from the discontinuity in the IQ loop. We then take many sweeps while lowering the amplitude of the digital waveform. This does

increase the noise in the sweep, as we are not maximizing the output of the DAC, but it is in practice not an issue if the warm amplification is optimized, as we are able to get good fits of the S_{21} sweeps. Eventually, when the power has dropped low enough, all of the KIDs should no longer be bifurcated. While this does lower the signal to noise output of the ROACH, we have not seen digitization noise as a significant noise source. We then fit the resonator shape with a non-linear fit, accounting for the bifurcation parameter, as a function of tone attenuation. Bifurcation turns on when the bifurcation parameter $a = \frac{4\sqrt{3}}{9}$, and thus we arbitrarily chose to interpolate the desired power to $a = 0.5$. In practice, this typically results in tone powers approximately 1 dB below bifurcation. We parameterize the fit in the following way [Siegel, 2016], for $y_g = Q_r x_g = Q_r(f - f_0)/f_0 = y + a/(1 + y^2)$,

$$|S_{21}|^2 = b_0 + b_1 x_{lin}|1 - \frac{Ae^{j\phi}}{1 + 2jy} + \frac{Ae^{j\phi} - 1}{2}|^2 \tag{4.1}$$

where b_0 and b_1 account for the DC and frequency dependent gain variation, $A = Q_r/Q_c$, ϕ is a phase mismatch parameter, and $x_{lin} = (f - f_{lin})/f_{lin}$ is the fractional frequency for with $a = 0$. This 8 parameter fits requires careful tuning of the initial guesses for the parameters for fast convergence, especially considering that this fit will be performed approximately 10 times per resonator for the full power tuning.

Note that at this point, we have measured the relative bifurcation between all of the KIDs. For example, in figure 4.2, we find that the attenuation of 1.7 dB is optimal. But once we re-scale the waveform to maximize the output of the DAC, that 1.7 dB does not refer to the absolute dBm. We perform another set of sweeps with the tuned relative bifurcation powers adjusted for each tone, but this time, we keep the tone power constant, and adjust the overall power with a variable attenuation on the cryostat input. In this way, every KID should undergo bifurcation at the same attenuation setting on this input attenuator. Fitting again for the optimal $a = 0.5$ below that bifurcation level, we have now measured both the relative and absolute bifurcation powers. This methodology also accounts for all of the power

dependence in the microwave components on the input, as we are using the KID bifurcation, which is relative to the power incident on the device, as our absolute power calibration.

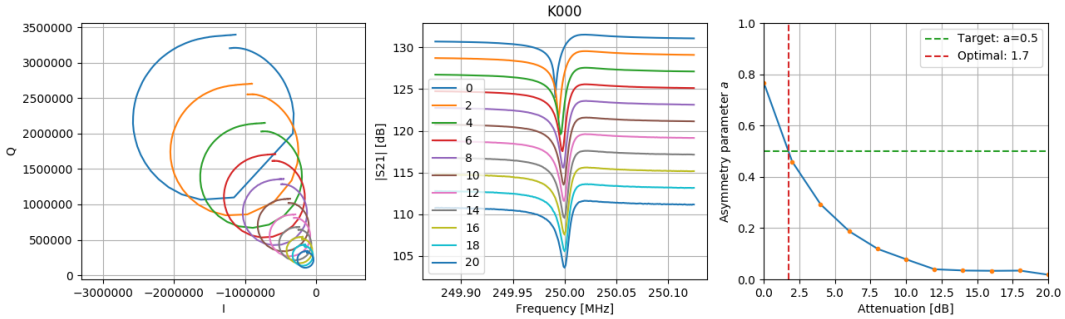

Figure 4.2: We see the result of a power sweep for a single KID. The plots from left and middle plots show the resonator sweeps in the IQ plane and amplitude. The sweeps are fit for non-linearity parameter a, which is then plotted as a function of the tone attenuation on the right hand side. We optimize to $a = 0.5$, which is just below bifurcation, typically around 1 dB for SuperSpec detectors.

4.2 Capacitor Trimming for KID Resonant Frequency Placement

One of the innate qualities of KIDs is their frequency multiplexing simplicity. It is possible to read out hundreds, or even thousands of KIDs on a single microwave coaxial line. In an ideal case, we have high resonator quality factors ($> 1 \times 10^4$), and good control of resonant frequencies placement so that we can read out many resonators within a given bandwidth. In practice though, fabrication consistency is difficult to achieve and results in arrays with collisions, where two or more KIDs have a resonant frequency difference within a few linewidths [Wheeler, 2019]. These clashes can prove to be problematic.

In the case of SuperSpec, the KID inductors are identical in their mask dimensions, and the resonant frequency is tuned by changing a combination of the tine density and widths of the KID capacitor. There are non-systematic variations in resonant frequency attributable to many factors that result in scattering from the design. For example, in fabrication, there will always be some natural variation in quantities such as deposition rates, etch rates, or

tine widths from lithography. Fabrication processes can slightly change due to multitudes of factors that can be difficult to control day to day. We strive to maintain a process that is as immune as possible to these variations, but when all factors are combined, it can be very difficult to achieve perfectly fabricated arrays of hundreds of detectors. The fact that the TiN inductors are so low in volume for our detectors also makes our design particularly susceptible to these variations, as a small defect in the inductor dimensions will have a relatively large effect on the resonant frequency compared to a defect in the massive Nb capacitor.

Additionally, as discussed in Chapter 3, we observe systematic scattering of the TiN KIDs due to T_c variations across the wafer. This shift in T_c, and consequently a shift in L_k, causes a change in resonant frequency. This effect is strongest in the radial direction of the wafer, having lower T_c farther from the center of the wafer, likely due to non-uniform nitrogen flow during the Ti sputter to create the TiN inductors. Due to the location of the chip on the wafer, and the physical ordering to the KIDs on the chip, this causes the KIDs at the end of each of the four readout banks to have resonant frequency systematically lower than expected, resulting in a 'crunch' at the end of each bank. We can see this effect in a particular device for the second bank B in figure 4.3.

The clashed resonators in their current state are not feasible for use in observation. Additionally, when scaling up to hundreds or thousands of detectors over multiple pixels, it will be necessary to automatically tune the probe tone locations and powers by sweeping and fitting the resonance to lock on to each resonator. When there are multiple collided resonators within one sweep window, this process becomes difficult due to KID-KID interactions and the resulting complicated fitting models. As well, the KID-KID interactions cause the filter-bank profiles to become jumbled, with multiple filter resonances corresponding to each KID in the collision.

These large numbers of clashed resonators due to T_c variation then represent a large

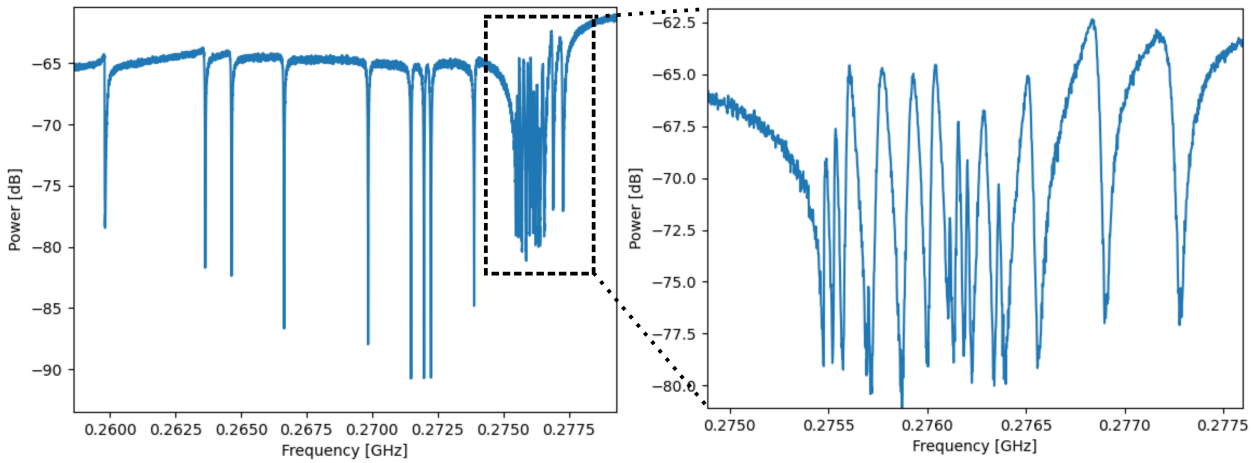

Figure 4.3: *Left* S_{21} sweep for a bank B for a SuperSpec 109 channel device. In an ideal case, this bank would contain 27 resonators spaced in the 260 MHz to 285 MHz range. We can clearly see that there in addition to natural scatter, the systematic T_c shift towards the end of the bank 'crunches' the end of the bank. The higher T_c towards the end of the bank causes the resonant frequency to come out lower than expected, resulting in a multitude of clashes. *Right* Zoomed in section of dashed area, showing in detail of the highly clashed end of bank. There are approximately 16 resonators packed into this small 2 MHz band.

roadblock in deployment, as it drastically reduces the effective detector yield, up to approximately 50% on some particular chips. But the non-systematic scattering of KID resonant frequency placement also contribute to resonator clashing. For the frequency scheduling for the 300 channel chips, with the T_c effects removed, Figure 4.4 shows the expected number of clashes due to natural variations stemming from fabrication.

4.2.1 FTS Measurements for KID identification

The task at hand is to identify each resonator with a physical KID on the chip. There are other experiments that have performed this task [Liu et al., 2017, McKenney et al., 2019], albeit with broadband detectors. By illuminating particular pixels in the focal plane, for example, with a flashing cryogenic LED, one can then see the corresponding resonators respond attached to that pixel. The alignment of the LEDs for each pixel then gives the mapping from readout resonator to physical location of the pixel on the chip.

The SuperSpec device is only a single pixel on a chip, but its function as a filter bank

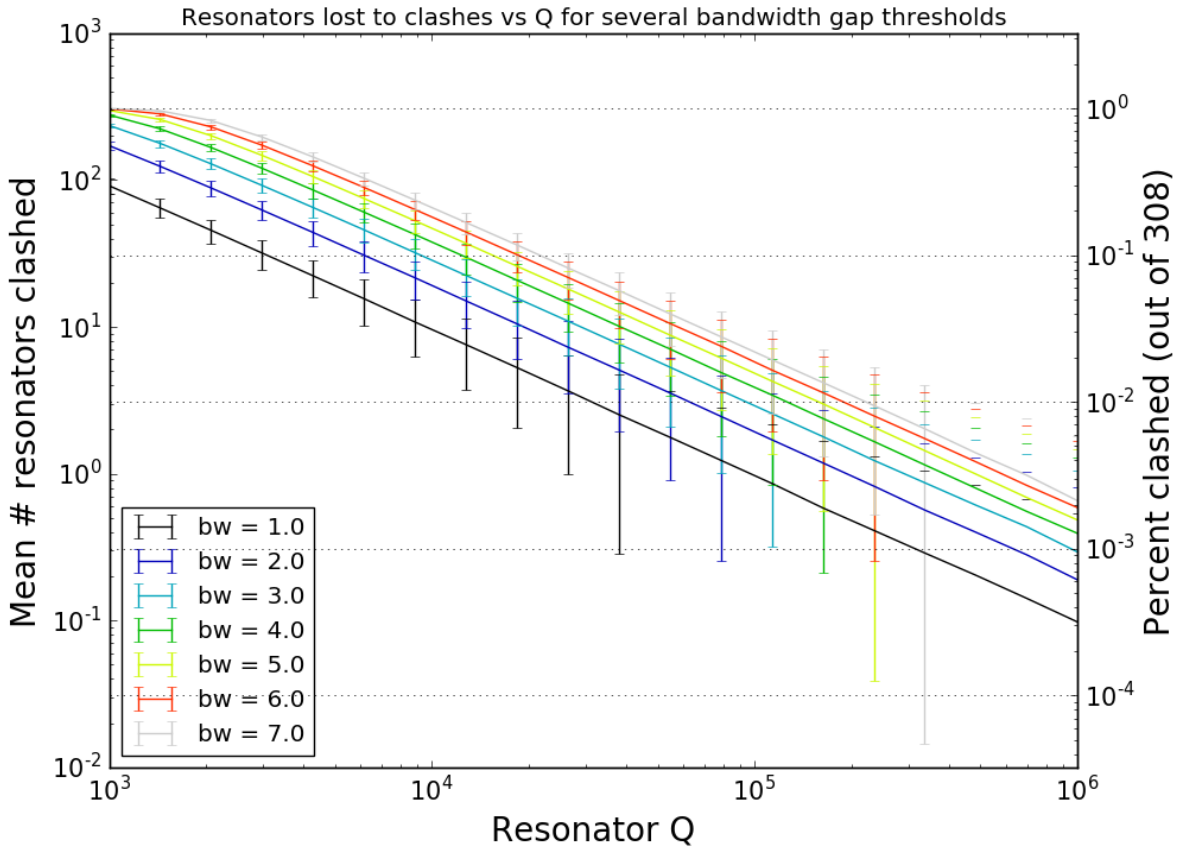

Figure 4.4: For the 300 channel die frequency scheduling, accounting for a correction in the T_c variation, we show the expected number and percentage of clashes expected for various thresholds of what is considered clashed. The value 'bw' refers to bandwidths, so that two resonators are considered clashed if their resonant frequencies differ by some number of bandwidths.

inherently carries its own identification mapping. If we can verify and trust our filter bank channel frequency spacing, combined with some clever choices of mm-wave and KID frequency arrangements and detective work, the unique filter bank channel should give us the mapping we require. Starting physically on the side of the chip closest to the antenna, the filter bank starts at the highest mmw-wave frequency and monotonically decreases as the feedline moves away from the antenna. Combining this with the separation of the KIDs into four distinct banks, we can look within each readout bank, measure the spectral profile of each detector, and sort by the spectral peak frequency to identify each resonator on the chip.

After sorting through the identification process, our fundamental pieces of data are two lists containing KID frequencies and the corresponding physical location on the chip. We must then identify which KIDs we would like to move, and where we should move them. Generally, the approach taken here is to keep the KIDs in the same increasing f_0 order, but to space them out so as to keep them at least 5 linewidths apart after moving. The way that we achieve this resonant frequency change is by performing a lithography and etch on the KID capacitors. We remove some fraction of the interdigitated capacitor tines, decreasing the capacitence, and therefore increasing the resonant frequency of the KID. The amount that we etch away is determined by a simple interdigitated capacitor model derived from electromagnetic simulations. The final data product in the end is some amount of tines to etch away from the capacitor on each KID on the chip, from which we generate a mask, and then perform the lithography and etches in the fabrication facility.

4.2.2 Capacitor Trimming Calculation and Fabrication

Calculating Expected Resonant Frequency

The capacitors are trimmed in order to decrease the KID capacitence, and hence, raise the resonant frequency. The amount that we expect the resonant frequency f_0 to change follows the simple scaling of $f_0 \sim 1/\sqrt{C}$. Because we know this scaling, we do not need to measure the capacitor exactly, but just to understand how it scales with the number of tines. For small changes in the number of tines, the capcitence of an interdigitated capacitor scales linearly with the number of tines, or in other terms, with the length of the capacitor. If the KID at resonant frequency f_0 should be moved to f_0', by changing KID capacitor with capacitence C to C'

$$f_0'/f_0 = \sqrt{C/C'} = \sqrt{l/l'} \tag{4.2}$$

where l' is the length of the KID capacitor after clipping, having a fixed width and tine density after before and after the clip.

In practice, this scaling does not perfectly match measured results. We believe there is a parasitic capacitence, likely contained in the KID capacitor rails, whose lengths are of order 3 mm. We can then model the capacitence of the KID scaling as $C \approx C(l) + C_0$, where $C(l)$ is the capacitence of the KID capacitor, scaling as $C(l) = kl$ for capacitence per unit length k, and C_0 being the parasitic capacitence. From the above, we can then determine that

$$\frac{f_0'}{f_0}^2 = \sqrt{\frac{kl + C_0}{kl' + C_0}} C_0/k = \frac{l' \frac{f_0'}{f_0}^2 - l}{1 - \frac{f_0'}{f_0}^2} \tag{4.3}$$

where C_0/k can be thought of as the 'equivalent length' capacitor added by the parasitic capacitence. As every quantity on the right hand side of the equation is measured, we can compute this value after performing the clipping, and feed back in to the fabrication cycle for future devices. This quantity has stayed constant with the same mask design on different devices, which strongly suggests its capacitor geometry origins, and luckily this also means this quantity only needs to be computed once for a particular design. For the capacitor geometries in the data presented in this work, the change in capacitor length with this added effect is significant, increasing the length to be clipped by 18 % to 23 % to hit the desired resonant frequencies. The way that we define our KID resonant frequencies causes this correction factor to vary as a function of resonator bank. We have four defined banks with 27 resonators each, and each bank has a different width, but same length, with the lowest frequency bank having the widest capacitors. Within the bank, the resonators are sorted by varying the tine density. So, for a toy example, a KID in bank A that needs to clipped from 100 MHz to 101 MHz, with capacitor length 3 mm, we will trim $1.18 \times (3\,\text{mm} - 3\,\text{mm}\frac{100\,\text{MHz}}{101\,\text{MHz}}^2) = 70\,\mu\text{m}$ of length off of the capacitor to hit the target resonant frequency.

Fabrication

After the chip has been removed from the test bed, the lens is removed (adhered by GE varnish thinned by acetone) by soaking in an acetone bath. The chip can then be fully cleaned in acetone and IPA with sonication. The fabrication is simple, including a single lithography, etch, and strip before the chip is ready to be cooled down again to verify the results of the clipping. The lithography is performed with positive photoresist AZ 1512 in the Heidelberg maskless aligner using the automatically generated mask file from the analysis of the FTS data.

After developing and baking in the vacuum oven, we perform two plasma gas etches. The capacitors for SuperSpec are comprised of a thin 20 nm layer of TiN on the bottom and thick Nb 250 µm on top. The Nb is etched first in the flourine based gas etch, and stops on the TiN. We then do a further chlorine based gas etch to remove the TiN. This final etch is very fast, usually less than 10 seconds long, and should be timed carefully as to not accrue black residue on the exposed silicon surfaces in between the capacitor fingers. We can see a micrograph before and after clipping procedure for a single capacitor in figure 4.5.

This procedure will leave behind silicon trenches with silicon pillars where the metal fingers once stood. We can verify that the metal has been fully removed by noting that the color of the finger structures matches the silicon. These silicon features left behind should not affect the capacitence of the IDC, as the electric fields should be very small in this region.

Fabrication Risk

This fabrication process does not come without risk. Chips can be dropped and machinery during fabrication can malfunction, giving some nonzero chance the the device will break somewhere in this process. There is an argument to be made that the clipping procedure may not be worth doing if the chip will only marginally improve the yield. This is not the case for the SuperSpec chips presented here, as the clashes reduce the yield by nearly 50%

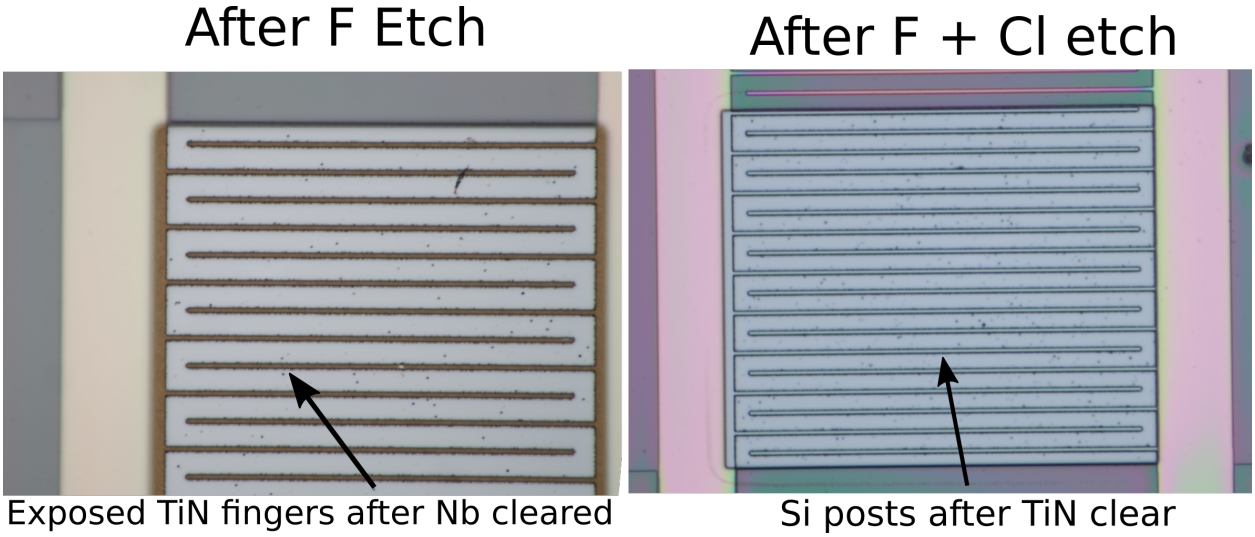

Figure 4.5: Micrographs of the etched area during the clipping procedure. Both images taken are of the bottom of the KID capacitor closes to the readout line, with patterned resist exposing just a small portion of the tines to be etched. Two etches are performed to first remove the thick Nb with a F based gas etch, and then the TiN is removed with a Cl based gas etch.

in some devices.

While the goal is to maximize the yield, we must also take into special consideration that a gap in the observation band, corresponding to several sequential mm-wave channels, could be very detrimental. Our observations are looking for narrow band signals, and if a chunk of our detector observation band is lost, we could entirely miss a detection. In this case, due to the linearly increasing resonant frequency of the filter channels going away from the antenna, we would lose higher frequency channels of our observing band if this clipping were not possible.

Another risk associated particularly to SuperSpec is that our lens mounting procedure has resulted in a broken device before. The jig we have used to align the lens to the chip antenna is not forgiving, and in the future could be designed in a safer way. This process requires the an additional lens mounting after the clipping, which is another added risk. With a bit of luck and good technique, none of the devices presented here have broken due to the clipping procedure.

4.2.3 Clipping Example

I will demonstrate here some of the details of the identification process and how to sort out common situations. For the devices presented here, on the first cooldown, we mounted the device in the optical box with the fridge configuration open to the room for FTS measurements. We perform VNA sweeps with a 300 K load in front of the window to identify all of the KIDs and record each resonant frequency. We use this temperature load to approximate the loads that will be incident on the detectors during the FTS measurements. This identification process can be automated, but in areas where there are many clashed KIDs, the identification is done manually. The resultant S_{21} sweep can be very messy, as seen in Figure 4.3, so it requires some careful observation to count all of the KIDs in the highly clashed sections. As we will see, there may be sections so highly clashed that there will be hidden KIDs in VNA sweep, but it is still possible to tease these out of the data.

From this list of resonant frequencies, we then set up and perform FTS measurements, as described previously in Chapter 3. From the interferograms, we calculate the spectral response of each detector. For a well-separated KID, this spectral response should have a narrow profile, such that it is simple to fit the data with a Lorentzian curve shape to extract a filter-bank resonant frequency. In this simple case, we have a clear mapping from KID readout frequency to mm-wave frequency, and will later use this in concert with the rest of the data set to gather the mapping to physical location on the mask.

If two or more KIDs are clashed, typically having the frequency separation on the order of the resonator line widths, the situation becomes more complicated, as the resultant spectral profile will have multiple peaks. We can see this type of behavior in figure 4.6, in which there is a "double", where two KIDs are clashed together. On the double, we place two probe tones, one for each KID at their respective maximally responsive points. The two KIDs are overlapping, and therefore, generally we will see two peaks in the FTS spectrum. The heights of these peaks allow us to distinguish which KID corresponds to which peak.

In an ideal case, we will see one strong peak, and one weak, and for the paired tone, the relative strengths of those peaks will swap. This swap occurs because the tone is optimally placed for only one of the two resonators in the pair so that the resonator furthest from the tone gives a lower spectral peak. This gives us an obvious interpretation for the KID identification, with the maximal peak belonging to that tone frequency resonance. It may be the case, however, that one KID in the double pair exhibits a much stronger spectral peak. This may be caused by many factors, including loss along the mm-wave feedline, detector response, resonance Q, and tone power optimization. In this case, though, we commonly see that one peak remains dominant, but the relative heights of the peaks changes. We can see this behavior in Figure 4.6, where the peak at 247 GHz in the spectrum remains the strongest signal for both readout tones. However, the secondary peak at 280 GHz changes in relative height to the dominant peak. As labeled in 4.6, the measured relative peak height for Tone 1 spectral response is larger than that of Tone 2. We can then assign the mm-wave frequency 247 GHz to the resonator at Tone 2, and 280 GHz to the resonator at Tone 1.

In extreme cases, there may be triple clashes, but the logic in sorting out the identification is the same as the doubles. Look for peaks in the mm-wave spectrum swapping, and if a resonator dominating the triple, look to the relative peak heights for the secondary peaks.

Once we have assigned each KID a corresponding mm-wave frequency, we can then look to the plot in Figure 4.7 for a visualization of each bank of resonators and their corresponding mm-wave peaks. We can use this plot to determine if particular KIDs have not yielded, and finalize the mapping from KID to physical location on the mask.

4.2.4 Discussion of Results

The set of deployment devices will have 6 chips in total, consisting of two 50 channel devices and four 110 channel devices. The collisions described in this work were mostly confined to the 110 channel devices, and thus the four to be deployed had this clipping procedure

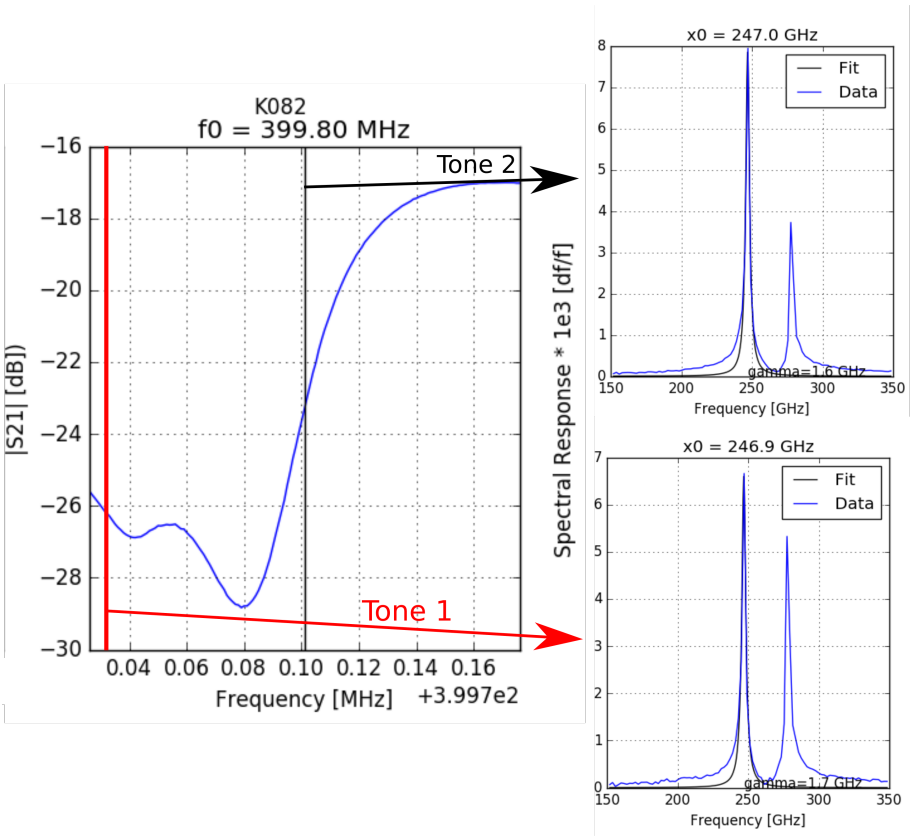

Figure 4.6: For the double indicated by the S_{21} plot on the left hand side, we place two readout tones, one for each resonance. We then take FTS data, and the resultant spectrum for each tone contains the same two peaks. But, in this case, one peak remains the 'dominant' peak with high power in both cases. So in order to distinguish which KID belongs to which mm-wave peak, one can simply look at the relative peak heights, and determine that tone 1 corresponds to the right hand peak at 280 GHz, and the resonator at tone 2 belong to the 247 GHz peak.

performed on them. We can see a summary of the results in terms of yield per device before and after the procedure in Table 4.1. The yield before clipping indicates the percentage of resonators that were observed and were not clashed, meaning this yield does not include resonators missing due to issues from not yielding due to fabrication. The resonator was considered to be clashed if another KID was located was within a 5 linewidth band around its resonant frequency.

We can see the same array of resonances before clipping in Figure 4.3, modified into Figure 4.8. It is clear that we have transformed this bank of resonator from possibly 6 useful

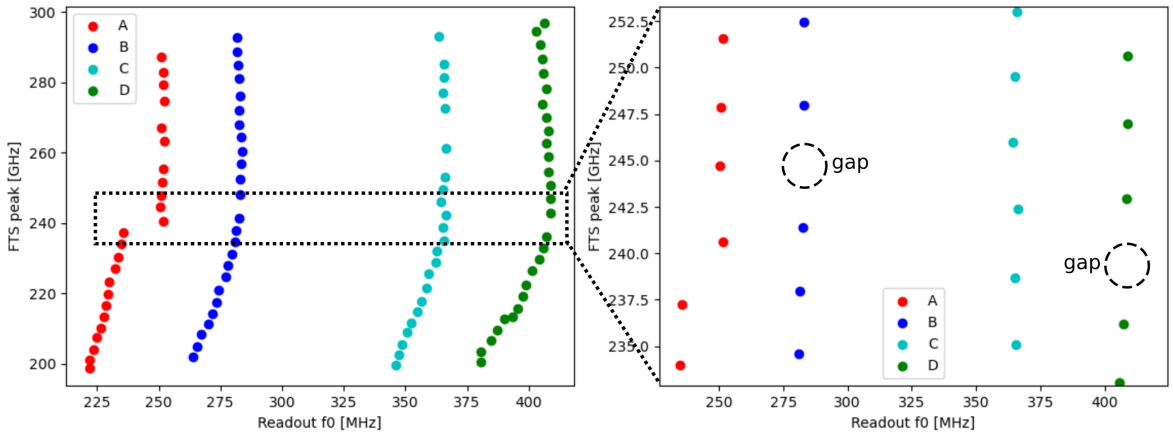

Figure 4.7: Once we believe we have correctly identified each mm-wave peak with a KID readout frequency, we can begin to do the mapping to physical location through plotting visualization. We see the four frequency banks, and within the banks, there are clear gaps as the mm-wave frequency steps upwards, which indicates that a particular KID has not yielded.

Device Name	Yield Before Clip (%)	Yield Post Clip (%)
L190919.4top	38.5	93.5
L200224.1mid	31.1	94.3
L200224.4bot	40.4	94.4
L200224.4top	36.1	95.4

Table 4.1: Table of four 110 channel devices that were clipped in preparation for their use during deployment. We reference here the yield in terms of what percentage of existing resonators observable in S_{21} are able to be used, meaning not clashed.

resonators, to yielding 26 in total.

It is possible to incorrectly count the number of KIDs prior to clipping because when performing the FTS measurements, we place one tone per identified KID. For particularly extreme clashed areas, there may be a KID hidden in the messy S_{21} sweep. Looking at Figure 4.8, we can see a unexpected resonance at 284 MHz that breaks the nice resonator spacing. It turns out that this resonator existed at this spot before the clipping, but was not identified, and thus was never clipped. It remained at that resonant frequency, causing a near clash. This error can be avoided by counting the number of spectral peaks when

over-plotting all FTS spectra in a bank, and making sure that number matches with the number of readout tones. If this is not the case, and there are more mm-wave peaks than tones, there is a hidden KID that must be teased out, for example, by observing a set of two nearby tones that contain three unique mm-wave peaks.

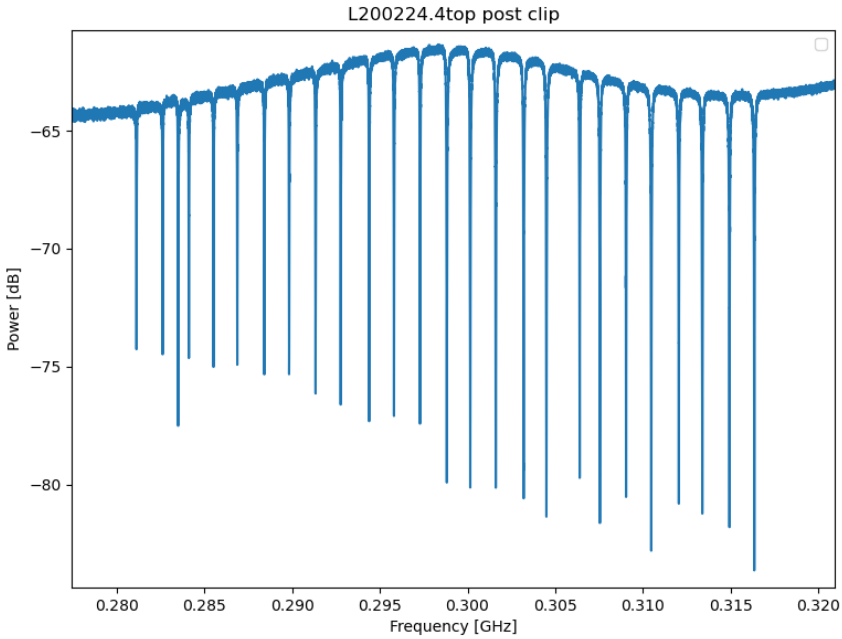

Figure 4.8: The bank presented in this figure is the result of chip surgery for the set of resonators seen in the highly clashed bank of figure 4.3.

There are several factors limiting the fundamental accuracy of this process. First, there is limited accuracy on the knowledge of the true location of the KID resonance before clipping. When two resonant frequencies are in close proximity, the interaction between the resonators will cause small shifts in the measured resonant frequency of each compared to the innate resonant frequency of each given no interactions. The double pair after clipping will be split, and thus the interaction shift will no longer affect the resonant frequency. The current modeling does not account for this interaction.

In order to quantify the accuracy of clipping while avoid systematics, we use a data set from which a number of KIDs were clipped that were not heavily clashed in order to avoid

the aforementioned uncertainty. As well, this data comes from a 50 channel device that did not exhibit the systematic parasitic capacitence mentioned in section 4.2.2. We can see an example of one of these etched resonators, including the target f_0 in the left panel of Figure 4.9. Note that another systematic shift was observed between cooldown for all resonators, likely due to the magnetic environment in the fridge during the superconducting phase transition. This shift was corrected by taking nearby non-trimmed resonators, measuring their shift from one cooldown to the next, and correcting the measured clipped resonator by this amount. Typical shifts for non-trimmed resonators were on the order of 25 kHz. With these corrections, we then quantify the clipping error with

$$\epsilon_{clip} = \frac{|f'_{0,meas} - f'_{0,exp}|}{f'_{0,meas}} \tag{4.4}$$

where $f'_{0,meas}$ is the measured resonant frequency after clipping and $f'_{0,exp} = f_0\sqrt{C/C'}$ is the expected resonant frequency after clipping. With this definition, we report a mean $\epsilon_{clip} = (4.1 \pm 2.4) \times 10^{-4}$

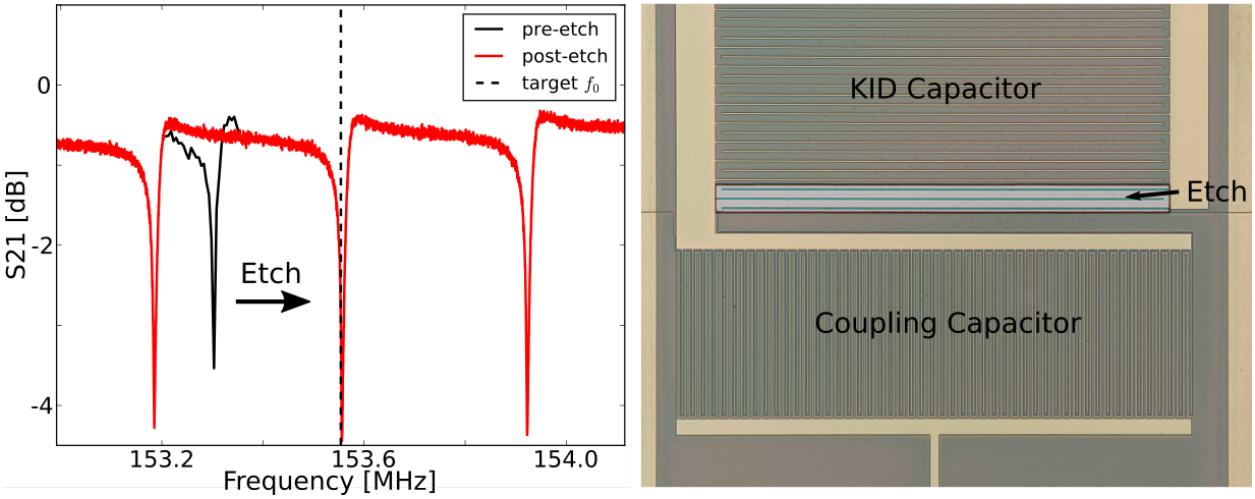

Figure 4.9: *Left* KID resonance before and after clipping, with target f_0 shown compared to measured resonance location. *Right* The clipped capacitor during the clipping fabrication processing showing the amount of tines etched.

Enabling Clipping Procedures for Future Arrays

The fundamental reason this procedure works is that we are able to identify banks of KIDs by their grouped frequency placement in readout space, and then within that bank, we trust the ordering of the mm-wave filter resonators. The monotonically increasing mm-wave filter frequency corresponds directly to physical position on the chip. If, for some reason, the scatter in the mm-wave filter frequencies became larger, on the order of 5 GHz, this method could no longer be trusted.

This basis reveals some challenges in implementation and scaling up this process for future instruments. For instance, for these devices, we have a cushion of empty readout space between each KID bank so that we can cleanly identify each resonance with a particular bank. As we move towards higher resonator packing densities in future instruments, this may no longer be possible, and then disentangling the beginning and end of resonator banks may prove to be difficult. Similarly, future technologies may employ higher spectral resolution filter banks compared to the moderate spectral resolution of current SuperSpec around $R = 300$. This will lead to much tighter filter bank packing in the mm-wave, necessitating high accuracy of mm-wave filter resonances. These problems can be mitigated by introducing a larger number of banks, which will increase the mm-wave spacing of sequential detectors within a bank. As well, moving to a narrow coherent mm-wave source instead of FTS measurements would aid in the mm-wave peak identification with much greater resolution on the spectral measurements.

Improvements may be made to the clipping precision by etching partial tines. The work presented here etched only full discrete tines, but by modeling the capacitence of shortened tines, one could greatly increase the precision of the target resonant frequency. There are roughly 600 tines, which implies that the minimum fractional frequency shift when clipping a full tine is approximately $1 - \sqrt{599/600} = 8.3 \times 10^{-4}$, when only considering non parasitic capacitence. Given that the most narrow the capacitor tines are of order 100 µm, and

lithography capabilities can easily define etching boxes with tolerances of 1 µm, then, one could effectively clip 0.01 tines to get a minimum fractional frequency shift of 8.3×10^{-6}.

For further applications, it is noted that this clipping technique was only implemented on the KID capacitor. But in principle, we should be able to employ this technique to modify the readout coupling capacitor. By removing tines in the coupling capacitor, we raise the coupling quality factor Q_c of the KID. For resonators that have low Q_c, one could not only move resonances, but also change the dip depth of the resonators to make well-spaced and optimally coupled arrays of KIDs.

4.3 Chopping Mirror Control Development

When on-sky at the telescope, it is common to use an observational technique called chopping, wherein a mirror in the optical path from oscillates between two mirror positions at some frequency. These two mirror positions correspond to two positions that the detector sees on the sky, and when pointing at an object, we can designate them as on-source and off-source. Assuming that the atmospheric contribution to the loading is constant between the two mirror positions, we can them remove background load by subtracting the off-source from the on-source, leaving the just source load remaining. We can see in Figure 4.10 how the chopper interfaces with the hardware and readout scheme at the telescope [Redford, 2023].

The SuperSpec focal plane consists of 6 chips total, separated into two groups of 3 by reflection or transmission through a polarizing grid. The group of three pixels are aligned linearly along a single axis. Therefore, when pointing at an object, to maximize integration time on-source overall, the chop will align the beam of two adjacent pixels on source in each respective position. The combination of their physical spacing in the focal plane, combined with the optics setup, which consists of several warm mirrors, determines the required chop angle.

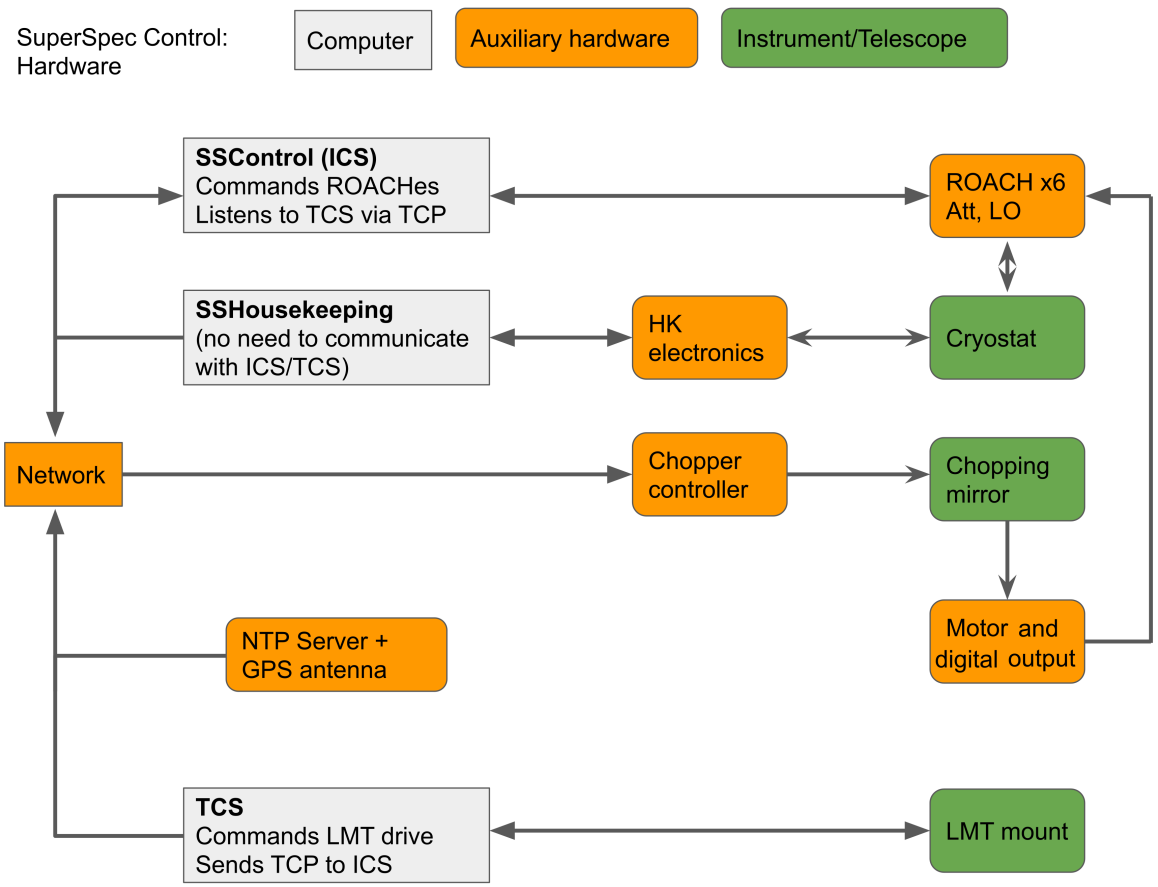

Figure 4.10: Block diagram of the hardware interfaces at the LMT

The output of the chopping motor stage consists of four digital output pins in a current sourcing configuration that supply current when switched to the on state, and zero current when off. The state of these pins can be set by sending commands to the motor controller using a python interface. We use these four pins to record information about the state of the chopper while simultaneously streaming data on the ROACH readout system. The ROACH contains 8 GPIO pins that use 3.3V TTL logic, and the state of the logic level on each pin is then recorded as a bit in a single 8 bit number which is then recorded in data packet. The goal is to then pipe the four digital output pins of the chopper controller into four GPIO pins, and then later decode the 8 bit number to assess the state of the chopper system.

During the chopping motion, we continually query the chopper motor controller to mea-

sure the instantaneous position and error on its position measurement. These values are used to set the digital output bits to their on or off state. The bits in the chopper contain information regarding these queries values. The first bit is simply the left/right position. The second and third bits give information on when the chopper is close to its desired position, within some set thresholds. The last bit determines when the error on the measured position is below some error threshold.

The interface between these two systems was not as trivial as connecting the outputs to the inputs via cables. The current sourcing outputs of the motor control system are more commonly used to drive relays, and not for voltage triggered logic input pins. In order to define the voltage across the ROACH input load, a voltage divider circuit was implemented on each digital output channel. The resistance of the two legs of the voltage divider were tuned to keep the voltage across the input pins to approximately 3V, while also keeping the current output of the digital outputs less than 10 mA as to not overcurrent the gpio input circuitry. A simple schematic and picture of the box can be seen in figure 4.11.

4.3.1 Chopper Parameter Optimization

The acceleration curve defines how the chopper operates at the beginning and end of travel when issued a move command. This acceleration profile must be optimized in order to maximize the time that the mirror is on-source within some specified positional error. The strategy is to chop between two positions while maximizing the duty cycle such that the time spent at the end of travel within some positional error is maximized. The optics hardware designs require the chopper to be able to operate in the range where the angle chopped is from 0.972 to 1.2 degrees, accounting for non-idealities when the hardware is assembled at the telescope. As well, we set the position error requirements to be 1/20 of the expected beam angle, which equates to 0.0071 degrees, although 0.01 degrees, which equates to 1/14 of beam may be an acceptable relaxed error requirement.

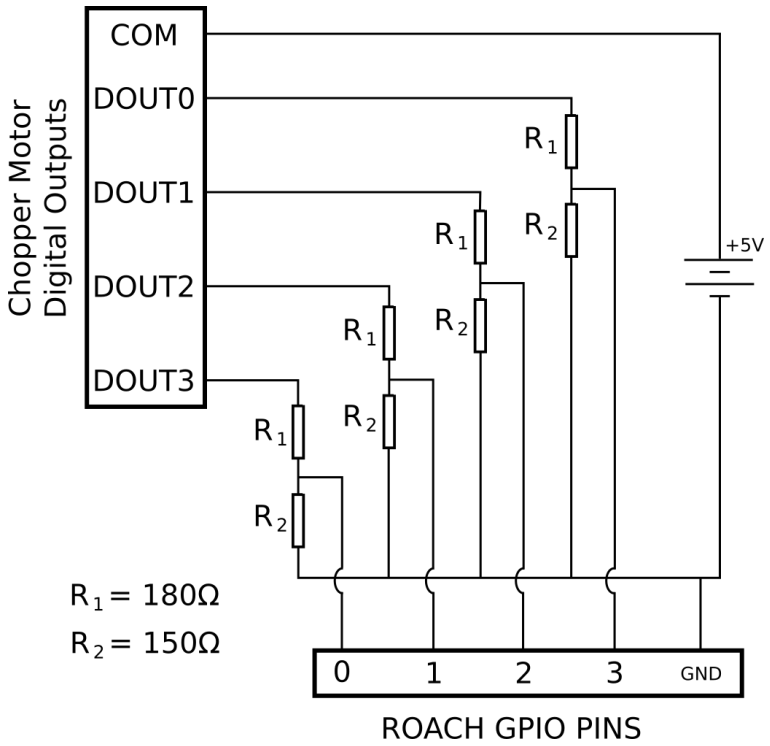

Figure 4.11: Circuit diagram of circuit connecting the chopping motor output to the ROACH GPIO pins. This connection allows us to quantify the position of the chopping mirror during operation, recorded directly in the data timestreams in the ROACH readout system.

As seen in Figure 4.12, we are able to achieve duty cycles above 0.8 below 1 Hz chop frequency with the stated requirements. At higher frequency chop, the duty cycle quickly falls off at 0.65 for a 1.6 Hz chop.

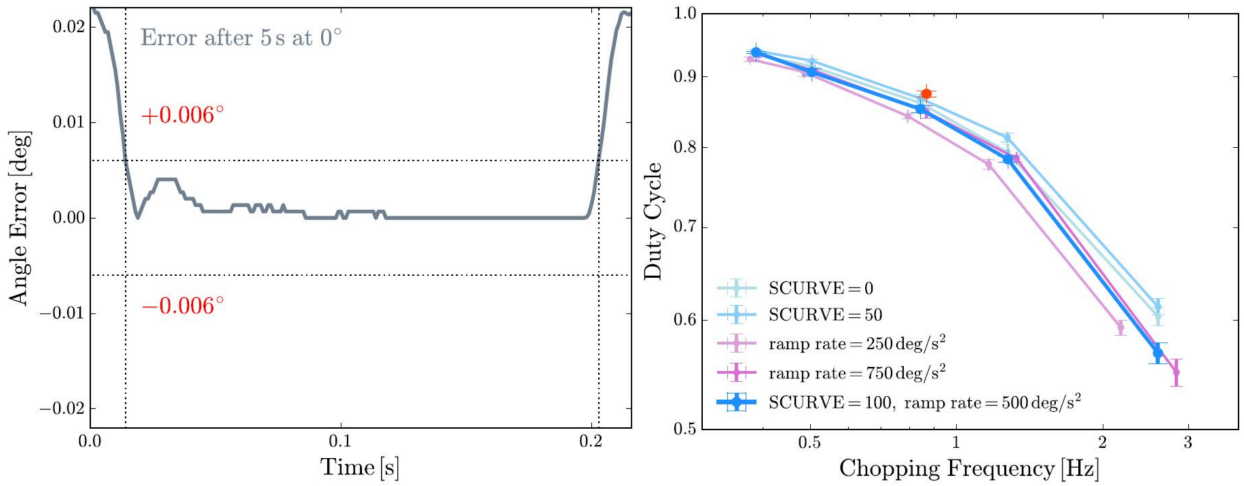

Figure 4.12: *Left*Plot of mirror position at end of travel for 0.2 second dwell time and positional error threshold of 0.006 degrees. *Right* Measurement of Duty cycle, meaning the fraction of time spend within the position error threshold at the end of travel during dwell. A duty cycle of 1 would be perfect chopper operation. Here, we show several SCURVE values, which change the acceleration ramp to be more trapezoidal, and ramp rate sets the speed of the ramp. [credit Ryley Hill]

CHAPTER 5

ADVANCED SUPERSPEC

5.1 Enabling Future Far-IR Missions

The SuperSpec project has shown that in principle, ultra-sensitive on-chip spectrometers operating in the mm and sub-mm bands are set to transform the field of wide-band, wide-field astrophysical spectroscopy. Looking towards the future, this miniaturization is a key technology for the next generation of cryogenic far-IR missions, such as the Origins Space Telescope (OST) [Cooray and Origins Space Telescope Study Team, 2018]. In space missions, where mass and volume become important concerns, SuperSpec shines compared to traditional grating spectrometer systems. As well, the antenna-coupled transmission line approach gives a comparative advantage to tradition spectrometers with direct absorption detectors. The direct absorption area scales as λ^2, whereas for the SuperSpec design, one is free to create a very compact absorption region, translating to a sensitivity advantage. For example a larger volume KID direct absorber translates to lower response and higher NEP, or for a bolometer, the larger island has higher heat capacity and thus reduced speed of response. SuperSpec solves these main issues of size and sensitivity at longer wavelengths.

There are large gains to be had where, in space, the fundamental photon noise is dominated no longer by the atmosphere, but by the solar system and galactic dust. As the fundamental noise in the signal drops, it is then required to design detectors with better sensitivity so that the detector remains photon noise limited. Above the atmosphere, the observation band is no longer limited by atmospheric absorption, opening up frequency regimes that are not possible to observe on the ground. As can be see in figure 5.1, a future SuperSpec-like detector could achieve two to three orders of magnitude higher sensitivity to optimized balloon spectrometers or SOFIA, a telescope on a modified plane in the stratosphere [Becklin and Moon, 2002].

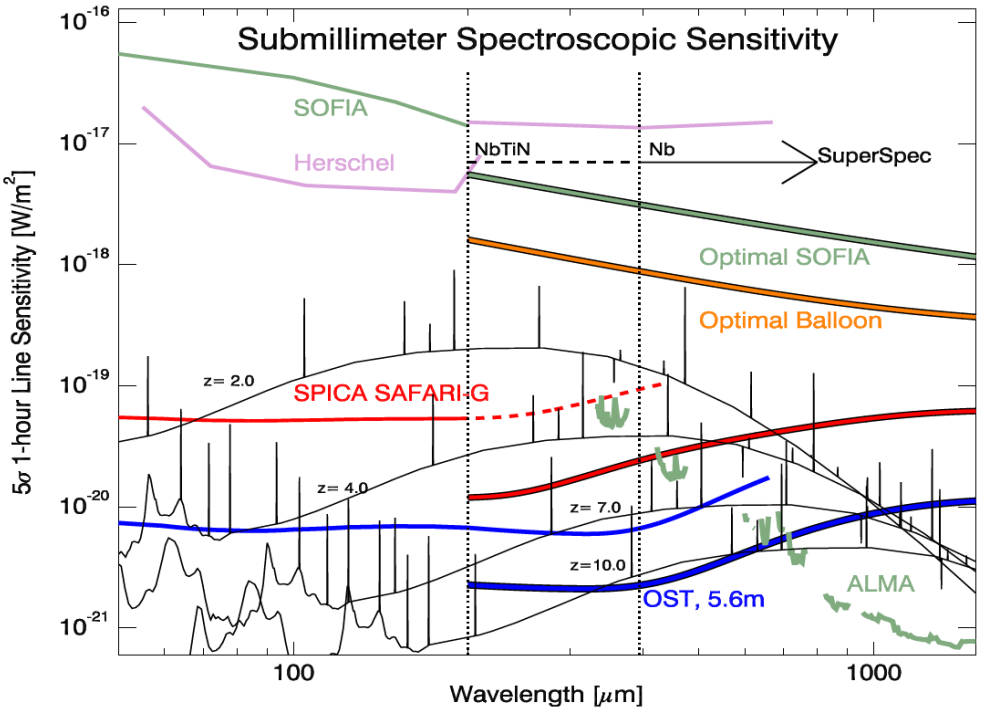

Figure 5.1: Various spectral sensitivities in the far-IR, with the thick red and blue curves showing the possibilities for SuperSpec performance. The blue curves show the smaller of two concepts for the Origins Space Telescope (OST), with assumptions for future SuperSpec missions achieveing $R = 3000$ spectral resolutions and NEP of $2 \times 10^{-20}\,\mathrm{WHz}^{-1/2}$. For reference to non-space missions, included are sensitivity projections for low-emissivity balloon (2.5 m aperture at 2% emissivity)and SOFIA (assuming 20% total emissivity) spectrometers. ALMA achieves comparable or better sensitivity when observing single objects, but lacks the wide-bandwidth of a SuperSpec-like architecture, and is thus not equipped for large surveys. The vertical dotted lines indicate the low-frequency cutoff for SuperSpec devices, largely due to material concerns where photons will begin to break cooper pairs in the mm-wave transmission line circuitry. Moving to higher T_c materials, such as NbTiN could extend observations down to 200 µm.

There are several challenges that must be met in order to deploy a SuperSpec style detector on a space platform. In order to push to higher spectrometer resolutions, as well as higher frequencies, dielectric material losses must be lowered by approximately an order of magnitude compared to the current SiN. As well, to push to higher mm-wave frequencies, the current Nb wiring will not suffice, as its T_c is too low, and photons will begin to break Cooper pairs in the Nb before ever reaching the KIDs. The focus of this work is in the third improvement that must be made, the increase in detector sensitivity in order to meet the requirements of the low-background space platforms.

5.2 Advanced SuperSpec Detector Design

In order to achieve background limited detectors compatible with space platforms, we must improve detector sensitivity beyond the current SuperSpec capabilities. Pushing for even more sensitive detectors becomes very challenging with SuperSpec without changing fundamental qualities of the TiN KID detectors. The high L_k, tuneable T_c TiN, combined with ultra-low volume inductors of $2.6\,\mu m^3$, allowed the SuperSpec concept to push to extremely low sensitivity, demonstrating some of the most sensitive detectors ever built with an NEP of $7 \times 10^{-19}\,\mathrm{WHz}^{-1/2}$ [McGeehan et al., 2018].

To maximize sensitivity, we must attempt to optimize detector optical response, $R_{det} = (df/f)/dP$, which scales with detector volume V and quasi-particle lifetime as $R_{det} \sim \tau_{qp}/V$. As we are already pushing the limits of what is achievable with these low detector volumes, it is difficult to obtain another order or magnitude in sensitivity by simply lowering detector volume. As well, measurements have shown that our sensitivity is fundamentally limited by the quasi-particle lifetimes in our TiN films, measuring lifetimes on the order of $\tau_{qp} \approx 5\,\mu s$. Meanwhile, it has been shown that is is possible to create kilo-pixel arrays of Aluminum KIDs with extremely high quality Al films boasting $\tau_{qp} \approx 1.3\,ms$ with detector volumes of $1000\,\mu m^3$ [Baselmans, J. J. A. et al., 2017]. We argue that pursuing high quality Al films with

order 1 ms lifetimes combined with a low volume inductor design will allow us to approach a detector NEP of $3 \times 10^{-20}\,\mathrm{WHz}^{-1/2}$, which is sufficient for background limited operation for low-background space observations. In a sense, we want to get the best of both worlds, utilizing high quality Al films, while leveraging the low detector volume SuperSpec design.

In the change to Al, there are many targets, considerations, and constraints put on our design choices, and in the following text, we attempt to plainly lay out the how we converged on a new design. First, we imposed fabrication constraints in order to improve output and efficiency of device fabrication scheduling. The employment of mask-less lithography processes, using machines like the Heidelberg Maskless Aligner MLA 150[1], enable significantly reduced cycle times because design modifications can be implemented instantly compared to traditional stepper systems. We were able to consistently pattern features larger than 1.5 µm with the aforementioned MLA 150 across a 4 inch wafer, but this then sets a lower bound on the minimal feature size of our design. Allowing for fabrication tolerance margins, we decided to set our minimal feature size of our inductors to 2 µm. The smallest features on the mask were the KID inductors and antenna slot probes, and setting these feature sizes consistently across the wafer was important to creating uniform detector response across our filter bank band.

This fabrication constraint leads to a roadblock when trying to exactly copy over the original SuperSpec coupling scheme. The tine width of the TiN indictors is 250 nm, and therefore the new Al inductors 2 µm wide need to occupy a much larger space, and will no longer neatly proximity couple inside the $\lambda/2$ Nb mm-wave resonators. This is compounded by the fact that Al is a much lower kinetic inductance material, leading to more Al inductor length needed to achieve the same inductances as TiN.

Instead of proximity coupling the KID to the mm-wave resonator, we opted to couple the two via a Nb microstrip feedline. This coupling scheme can be seen in Figure 5.2. In

1. https://heidelberg-instruments.com/product/mla150/

essence, this idea was to use the same Nb resonator style as SuperSpec, but to make the output coupling to the KID independent of the KID design. There would be a Nb microstrip proximity coupled to the KID instead, that would then bring the mm-wave excitations along the microstrip to the KID. The first conception of this design placed the Nb resonator and coupler on the same bottom layer of the inverted microstrip design. Through simulation, it became clear that the gaps necessary to achieve the desired resolutions were too small. For example, $R = 300$ required fabricating a gap of 0.4 µm, which is very difficult our maskless lithographic system.

The idea came to instead move the coupling to a different layer. This coupling strategy is an effective one when worrying about fabrication constraints and tolerances. While it is difficult to achieve maskless lithographic structures below 0.5 µm, it is simple to measure the deposition rate of SiN films, and control the thickness of these films within 10s of nms. This control, combined with layer-to-layer alignment within 200 nm, allows for finer control of the coupling between these components. This design has its drawbacks though, particularly when it comes to step coverage issues. The comparable thicknesses of the Nb structures and two SiN layers of 250 µm each resulted in thin sections of SiN near the mm-wave Nb structures, leading to potential shorts and coupling strengths not well modeled by the geometry of ideal depicted in Figure 5.2.

In further refining this design, we borrowed from OMT-coupled style leKIDs from Barry et al. [2018], seen in Figure 5.3. In this design, the mm-wave radiation from the antenna couples to the resonant $\lambda/2$ filter capacitively through an isolated Nb rectangle etched out of the ground plane on the top layer. Similarly, the mm-wave filter couples to another capacitive structure on the output, which then connects to an output microstrip to the KID inductor. There are several advantages to this choice of coupling geometry. First, it makes the KID design and mm-wave couplings relatively independent, while also allowing for much simpler simulation. Second, the couplings $Q_{i,mm}$ and $Q_{c,mm}$ of the mm-wave filter are easily tunable

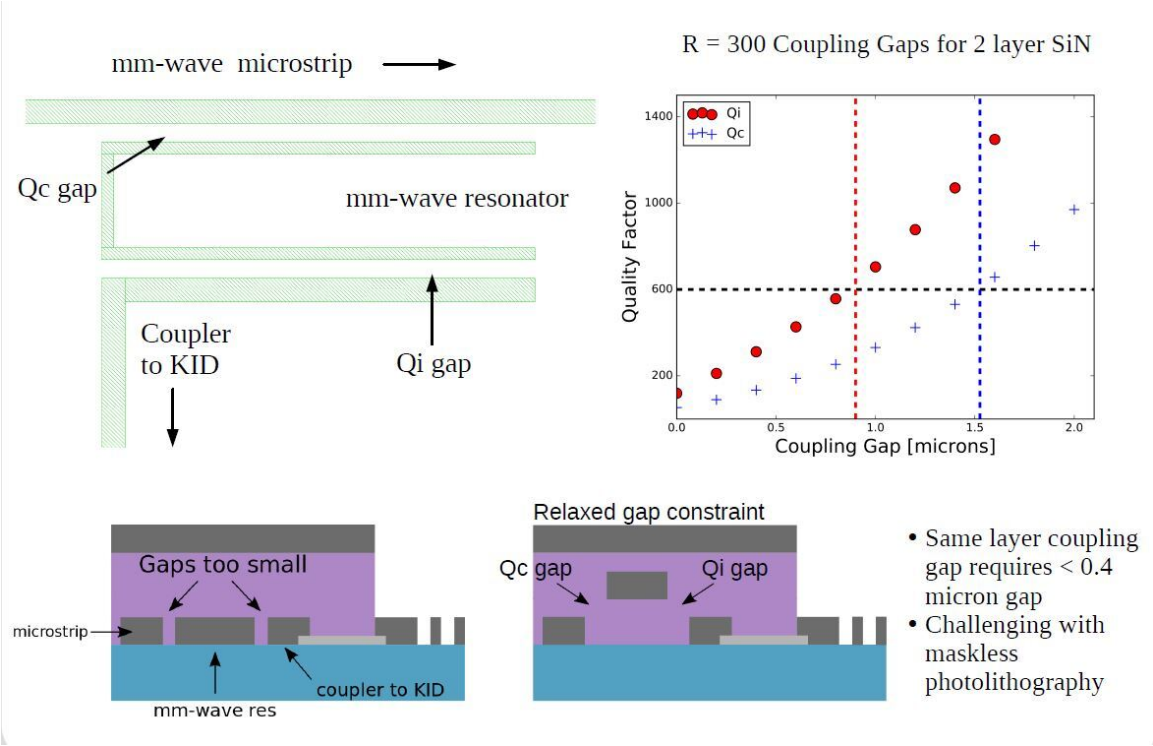

Figure 5.2: First attempt at filter bank coupling re-design. Keeping the original shape of the mm-wave resonators, we use a microstrip coupler to bring the filter excitations to the KID. Placing the coupler on the same layer was not feasible due to small gaps, so the coupler was moved to an intermediate layer to allow coupling out of the plane of the silicon. There were inconsistency issues in fabrication related to step-coverage, which led to a more mature and variable design seen in figure 5.3

through simple geometry changes by simply changing the size of the coupling probes on the resonator layer. As well, $Q_{c,mm}$ only has strong dependence on the input coupling from the feedline, and $Q_{i,mm}$ for the ouptut coupling, so that simulation and interpolation becomes much easier due to less cross-correlation between several geometric parameters and their effect on coupling strenghts. For example, we can see in Figure 5.3 that $Q_{i,mm}$ depends strongly on the output capacitor coupler width, whereas the is very small dependence of $Q_{i,mm}$ on the input coupler length. Lastly, the strength of these couplings can be made smaller or stronger based on the dielectric thickness, which enables this design to be used for a wide range of spectrometer resolution requirements.

With the filter-bank geometry set, we must then consider the consequences for the design of the KID. Unlike traditional directly absorbing lumped element KIDs, which are matched to the impedance of free space, the film resistivity of the Al inductor must be matched to the impedance of the of the microstrip output from the filter. In this case, a 4.5 µm wide Nb microstrip under 500 µm of SiN ($\epsilon_r = 7$), is well matched to a T-junction at the midpoint of an Al trace of 2 um with sheet resistivity of 1 $\Omega/\square$, corresponding to thickness of 30 nm. This junction will not be perfectly matched though, as the Al inductor acts as a lossy transmission line. As the mm-wave radiation arrives from the filter microstrip coupler, the excitation propagates through the $T_c = 1.2$ K Al, and is absorbed, breaking Cooper pairs. This fact that the Al is lossy at mm-wave means that the imaginary part of the impedance will be non-zero, while the Nb microstrip remains lossless, causing reflections at the T-junction. Minimizing the film resistivity of the Al will lower this reflection, but comes with considerable trade-offs. Surface resistance is proportional to kinetic inductance, and therefore the kinetic inductance fraction α_k will decrease. As well, lower surface resistance requires thicker films, which results in higher detector volume, and consequently lower detector responsivity. Finally, the lossy transmission line must be long enough so that all of the radiation is absorbed in the inductor before the Al terminates on the Nb KID capacitors.

$$\alpha_c = 8.686 \times \frac{R_s}{2Z_0 w} \tag{5.1}$$

Equation 5.1 describes the attenuation length for microstrip conductor loss [Barry et al., 2018], which for $R_s = 1\,\Omega/\square$, $Z_0 = 23.8\,\Omega$, and trace width $w = 2.0\,\mu\text{m}$, comes out to approximately 91 dB/mm. With an inductor length of 2 mm, essentially all of the radiation in the Al will be absorbed before propagating to the KID capacitors.

A final consideration in KID design is that resonators with lower resonant frequency are preferred. For reference, SuperSpec KIDs lie typically in the 100 MHz to 400 MHz range. Lower resonant frequencies gives access to more octaves in frequency space multiplexing, essentially meaning we are able to read out more KIDs per unit of readout bandwidth. As well, lower operating frequencies can reduce two level system noise [Zmuidzinas, 2012]. Yet, lower frequency resonances requires larger KID inductors and capacitors. We should minimize the KID inductor size in as this is the responsive detector volume V, and as in the case with SuperSpec, the interdigitated capacitors become nearly macroscopic, reach lengths of 3 mm.

The final dimensions of the Al KID inductors come to 2.0 µm width, 30 nm film thickness, and 2 mm length, yielding detector volume of $120\,\mu\text{m}^3$. The KID inductors are identical for each detector, and the resonant frequency is varied by changing the width of the KID capacitors. The capacitors have fixed lengths of 3 mm, with tine widths of 2 mm and 3 mm spacing, and the widths are varied from 104 µm to 162 µm.

5.2.1 Simulation

Electromagnetic simulation was performed to optimize mm-wave couplings as well as the KID inductor and capacitor geometries. These simulations were performed using Sonnet[2] 3D planar software. As seen in figure 5.3, the resonant filter transfer functions were extracted

2. https://www.sonnetsoftware.com/

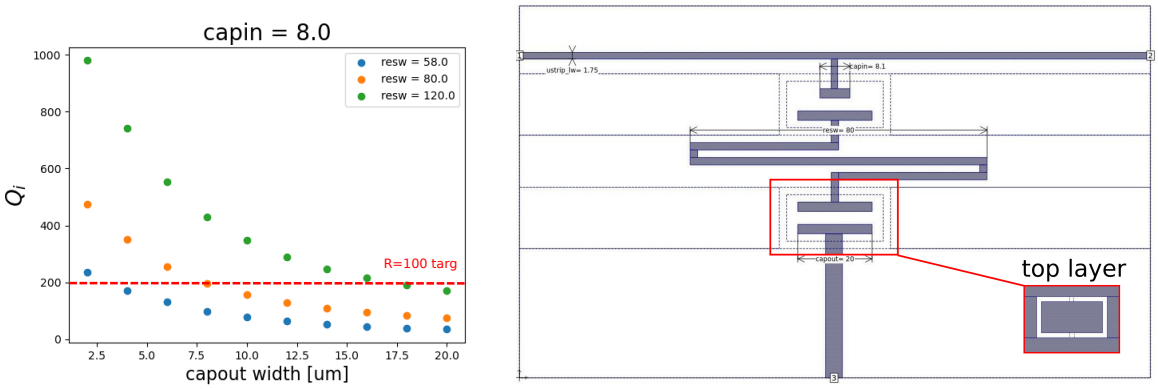

Figure 5.3: *Left* Results of simulation of mm-wave coupler on the right hand side. The plot shows the mm-wave filter Q_i as a function of the output coupler width for a fixed input coupler width of 8 µm. The three colors of the data represent the length of the mm-wave resonant filter, corresponding to filter frequency at the lower, middle, and high end of the band. The red dotted line indicates the requirements for $R = 100$ for this particular fixed value of input coupler width. *Right* A top down view of the mm-wave filter in Sonnet EM simulation software with the top ground plane obscured to reveal the filter structure on the bottom layer.

as a function of physical parameters, such as the coupler tine widths. These filter resonances were then fit for quality factor parameters Q_i and Q_c, and run through a multi-dimensional optimization of geometric parameters across the band, optimizing for $Q_{i,mm} = Q_{c,mm} = 200$, which equates to $1/Q_{mm,tot} = R = 100$ target.

On the KID side, with the inductor geometry set by factors discussed in the previous section, the size, width, and pitch of the interdigitated capacitors was modeled and interpolated to create a spacing of KIDs across the 1.3 GHz to 1.6 GHz range.

5.3 Fabrication

The fabrication was performed in the Pritzker Nanofabrication Facility[3] on the University of Chicago campus. The devices were made on 4 inch high resisitivity, $> 10\,\mathrm{k\Omega\,cm}$, silicon

3. This work made use of the Pritzker Nanofabrication Facility part of the Pritzker School of Molecular Engineering at the University of Chicago, which receives support from Soft and Hybrid Nanotechnology Experimental (SHyNE) Resource (NSF ECCS-2025633), a node of the National Science Foundation's National Nanotechnology Coordinated Infrastructure

wafers, approximately 500 µm thick.

The silicon wafer is first cleaned by sonication in acetone, followed by isopropyl alcohol (IPA), then rinsed in deionized water (DI). The surface is prepped by submerging the wafer in Nano-Strip and hydroflouric acid (HF), with a DI rinse after each dunk. Then the wafer is dehydrated in a 110 °C vacuum oven, and immediately placed in the vacuum load-lock of the Al deposition chamber. To ensure high quality Al films, it is important to achieve low chamber pressures in the sputtering chamber. To aid in this effort, we pre-sputter titanium for 10 minutes in the chamber with the shutter closed to act as a getter, which lowers the chamber pressure. As well, simply waiting several hours, or pumping overnight and depositing in the morning, can produce chamber pressures below 1×10^{-8} torr. We then deposit the 35 nm Al KID inductor layer in the AJA ATC 2200 Sputtering System. Here the inverted microstrip design gives a great fabrication advantage, as we are able to create a clean surface for the deposition of the most sensitive structure of our detector. We then pattern resist, using the maskless lithography writer, and perform a Chlorine (Cl) gas etch to pattern the Al KID layer.

After stripping, cleaning, and dehydration bake, we load the wafer into the sputtering chamber and perform an Argon ion mill to etch away the nature surface aluminum oxide layer. This ion mill was calibrated by measuring the DC resistance of an Al-Nb test structure as a function of ion mill time, and found the resistance of the superconducting Al-Nb interface goes to zero for ion mill etch of 90 s. After ion milling, we deposit 150 nm of Nb, then similarly pattern and perform a flourine-based gas etch to pattern the Nb mm-wave structures, KID capacitors, and microwave readout wiring.

Next we pattern 7 µm thick lift-off resist, deposit and lift off 500 nm of SiN. The final 150 nm Nb ground plane layer is deposited. We spin resist and pattern for a final flourine-based etch to define the ground plane structures for the microstrip readout, filter-bank coupling capacitors, and dual-slot antenna.

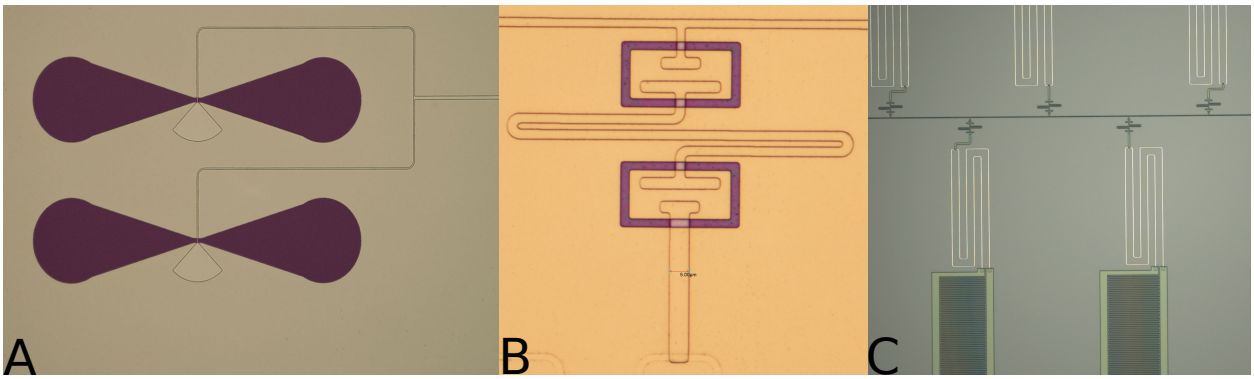

Figure 5.4: (A) Dual-slot bowtie antenna. The bowtie shapes are holes etched into the topside ground plane, exposing the SiN underneath. A small probe across each slot is then combined into the main mm-wave feedline, which then travels to the filter-bank. (B) A micrograph of an individual filter-bank channel resonator. The purple rectangular outlines are the etched ground plane areas that define the input and output coupling parallel plate capacitors for the filter-bank channel. The probe features couple out of the page, through the SiN dielectric, to the parallel plate on the ground plane layer. At the bottom of the image, the Nb-Al connection to the Al inductor can be faintly seen. (C) An image of several mm-wave channels and their associated KIDs tiled along the mm-wave feedline before the SiN and Nb ground plane have been deposited. The thin, bright white meanders are the Al KID inductors, connectected to large Nb interdigitated capacitors. The smaller mm-wave Nb resonators can be seen close to the mm-wave feedline near the midpoint of each KID inductor.

At this point, the device is finished, but we must etch through holes in the silicon for lens alignment. The silicon etching step is also used to simultaneously dice the chip. We spin a thick 20 µm resist to be able to etch all the way through the 500 µm thick Si wafer, achieving an etch selectivity in the range of 30 to 35. The wafer is mounted on a dummy wafer of the same size with silicone bonding oil and processed until the bottom dummy wafer is exposed. After cleaning the device, it is then ready for lens mounting and testing.

In figure 5.4, we see micrographs of the device at various stages of the fabrication process.

5.4 Experimental Setup

In order to use existing testing infrastructure, such as chip housing boxes and silicon AR-coated lenses, the footprint of the chip was kept to be identical to that of the SuperSpec

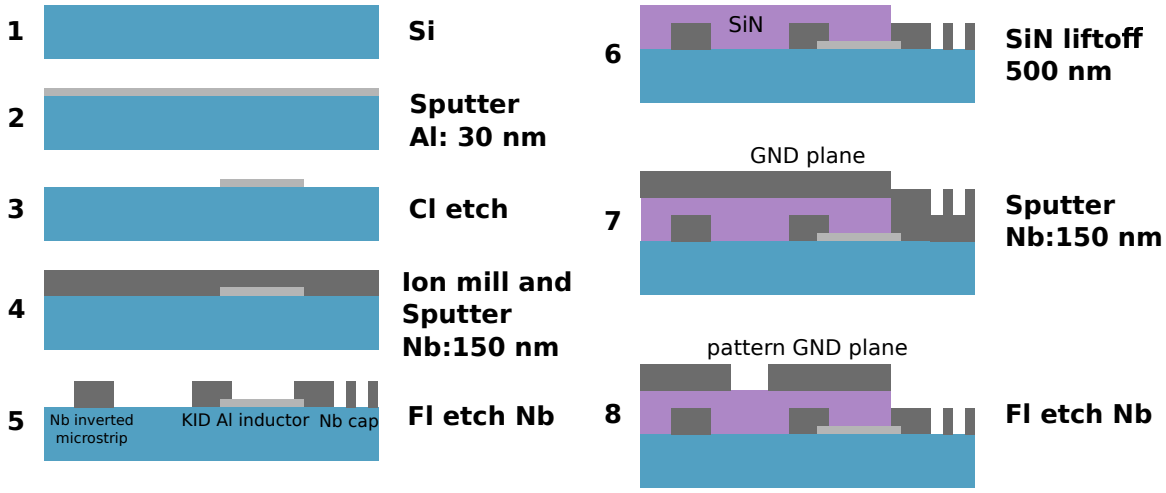

Figure 5.5: Cartoon representation of fabrication layers and main process steps. Not pictured is the final silicon etch through the wafer for dicing and alignment pins during the lens mounting process.

devices described in Chapter 3. As well, this device targets the same mm-wave band, so that the hardware including boxes and filters are identical. Refer to section 3.2 for details. The only hardware change is from the readout frequency of the KIDs increasing from 100 MHz to 400 MHz up to 1.3 GHz to 1.6 GHz, which requires a higher frequency mixer on the warm readout side.

5.5 Results

5.5.1 Film Properties

The film thickness is determined from carefully calibrated deposition rates and ion mill etch rates during the fabrication process. The Al films are estimated to be 30 nm thick. We measure the resistance of these films with via a long Al wire with two wirebond pads at either end, and probe the resistance with a known number of squares in the wire. We measure the sheet resistance of the film to be 1.1 $\Omega/\square$ at 4 K. We typically do not use this method to measure the film T_c as the temperature stability of the fridge worsens above 1 K, with the potential of losing the dilution cooling cycle entirely above approximately 1.4 K.

By taking sweeps with bath temperature, as described in section 3.3, we are able to fit to the temperature dependence of the resonant frequency of the KID. The important distinction in the fit from the previous section is that in the Al case, the kinetic inductance fraction, that is, the fraction of inductance that is determined by the material, is much lower than in TiN. In the case of SuperSpec, $\alpha = 1$, whereas here, we must introduce α as a fit parameter. We measure the film T_c in this way to 1.2 K, uniform across the device. Simultaneously, we fit for the kinetic inductance fraction $\alpha = 0.15$, which is in agreement with simulation.

5.5.2 Al KID Resonators

This prototype device yielded 35 KIDs with resonant frequencies ranging from 1.33 GHz to 1.63 GHz with $T_c = 1.2$ K and median Q_i of 2.8×10^4 in a dark configuration. The internal quality factor came out lower than expected as we designed for coupling quality factor $Q_c = 1 \times 10^5$, with optimal operation at $Q_i = Q_c$. The median coupling quality factor was measured to be $Q_c = 8.1 \times 10^4$. This mismatch with $Q_c > Q_i$ can also be seen in the small dip depth of the KID , as seen in the left panel of figure 5.8, with an optimal depth of 6 dB. The internal quality factor was also measured under 77 K and room temperature loads, as seen in figure 5.6. The Q_i has dropped to an average of $Q_i = 2.2 \times 10^4$, when compared to the dark setting. With the current resonator packing density in the readout space, these quality factors are sufficient for a current 100 to 300 channel device.

5.5.3 Noise Performance

We collected noise spectra using the single-tone system described in Chapter 3, with a resulting noise spectrum in Figure 5.7. The noise spectrum contains the classical features, including the $1/f$ noise at low frequency, a GR noise dominated white level, and a rolloff at high frequency. The average on-resonance white noise level was measured to be $1.2 \times 10^{-17}\,\mathrm{WHz}^{-1/2}$.

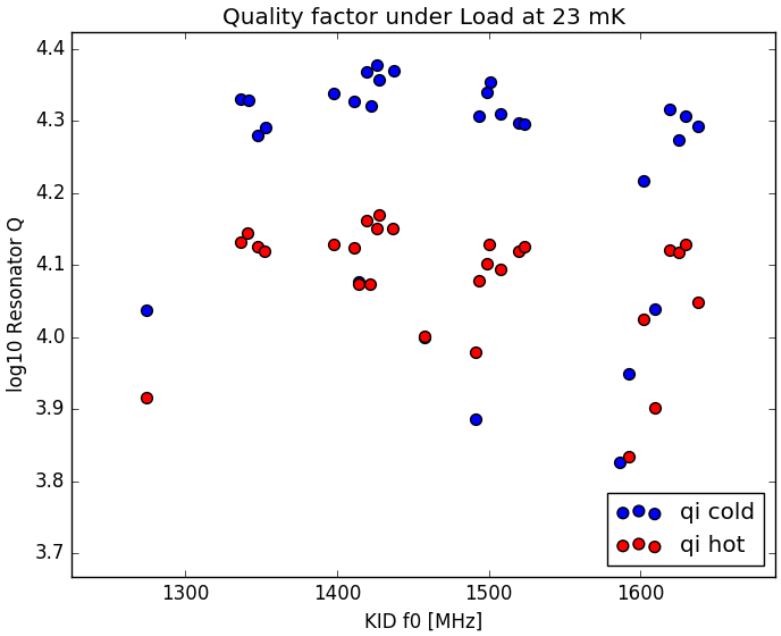

Figure 5.6: The internal quality factor of spectral channel KIDs when exposed to temperature loads of 77 K and 293 K.

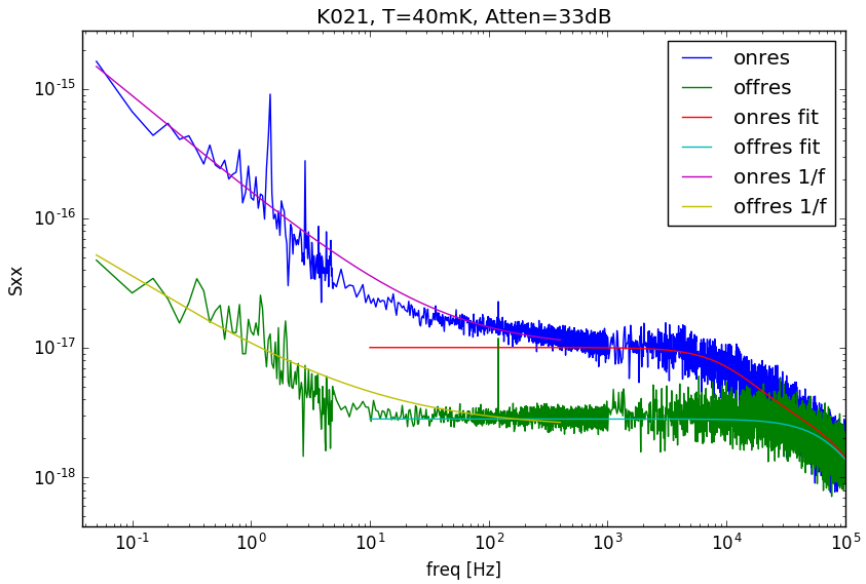

Figure 5.7: Noise power spectral density of KID in dark setting with single-tone noise measurement setup. Included are fits for the various features, including the white noise level on and off resonance and the $1/f$ noise at low frequency.

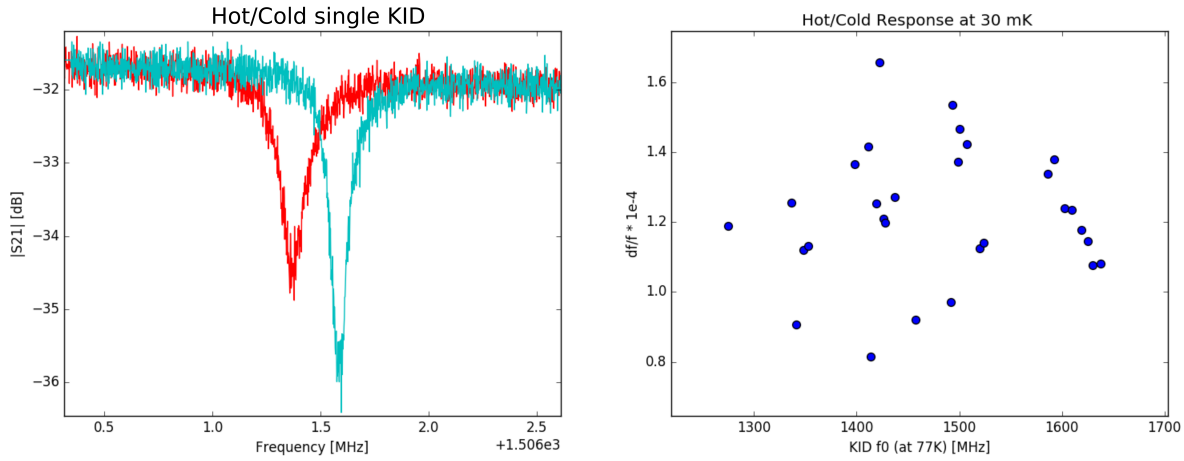

Figure 5.8: *Left* A single resonator transfer function under hot and cold loads. Under the hot load (red), the resonator Q_i and resonant frequency drop. *Right* For the 35 KIDs yielded, we plot the fractional frequency response, calculated from fitting for the resonant frequency in the hot and cold transfer functions from the left panel.

5.5.4 Responsivity

The optical performance of the device was measured in the same described in Chapter 3, with an AR coated silicon lens mounted on the device.

Figure 5.8 displays the results of a hot/cold response test. A beam filling temperature load is placed in front of the window at room temperature and liquid nitrogen, and we perform VNA sweeps to measure the change in fractional resonant frequency change. The spectral channel KIDs shift in resonant frequency to this change in optical load, and can be converted to power, $P = \eta\lambda^2 B(\nu, T)\Delta\nu$. Here, $\eta = 0.65$, accounting for the absorption and scattering due to the optical filters between the source and the device. As well, we must account for the single-polarization nature of the antenna, and we measure median fractional frequency response of $2.5 \times 10^7\,\mathrm{W}^{-1}$ with stage temperature of 30 mK.

5.5.5 Filter Bank Performance

The spectral profiles of the filter bank channels were measured through Fourier transform spectroscopy. We depict this measurement in 5.9 for a single KID. In the right panel, we see

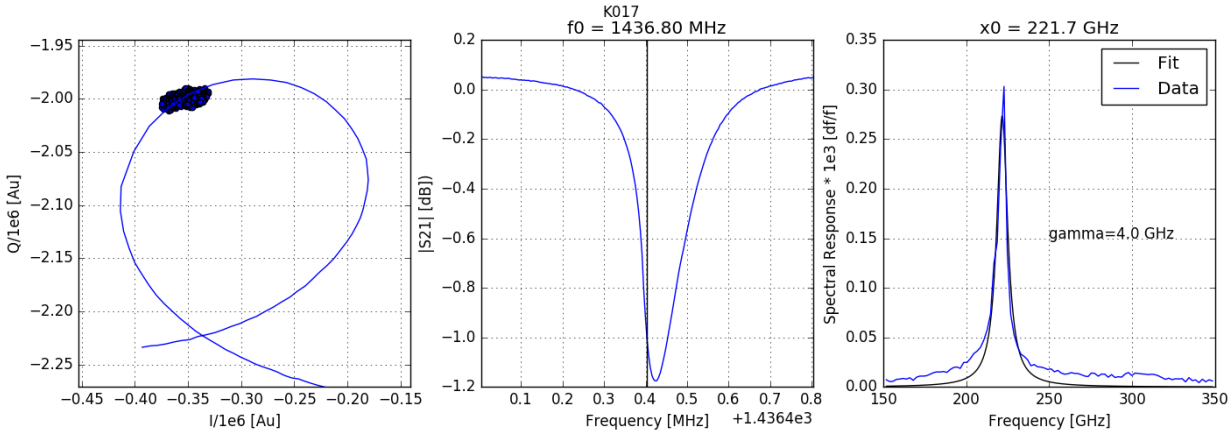

Figure 5.9: *Left* IQ sweep with over-plotted IQ points associated with interferogram taken with FTS. *Middle* S_{21} sweep, of resonator labeled K017. The vertical line indicates the maximally responsive frequency where the readout probe tone is placed during measurement. *Right* Resultant spectrum corresponding to interferogram taken with FTS. This spectral channel is centered at 221.7 GHz with a Lorentzian width of 4.0 GHz. Note that there are some excesses in the wings compared to a true Lorentzian line shape.

there is peaked spectral response in individual spectral channels that fall within the desired bandwidth of 180 GHz to 310 GHz. From these spectra, we calculate a channel resolution by fitting to the Lorentzian peak, and extract the total quality factor of the resonance such that $R = 1/Q_{mm,tot}$, with the pictured peak measuring approximately $R = 57$, centered at 221.7 GHz. This measurement is simultaneously done on all KIDs, and we report a median resolution of $R = 45$.

The gap of spectral channels in the 250 GHz to 280 GHz range offers a clue to the low yield. We have identified the liftoff step as a potential candidate to introduce shorting from the ground plane layer to the KID layer. Typically, when liftoff fails, there are physical sections of the chip that do not lift off as cleanly as others. For this design, the mm-wave ordering corresponds directly to physical location on the chip, so that a chunk of spectral channels not yielding in mm-wave space means an isolated section of the chip did not yield, which is consistent with the liftoff hypothesis.

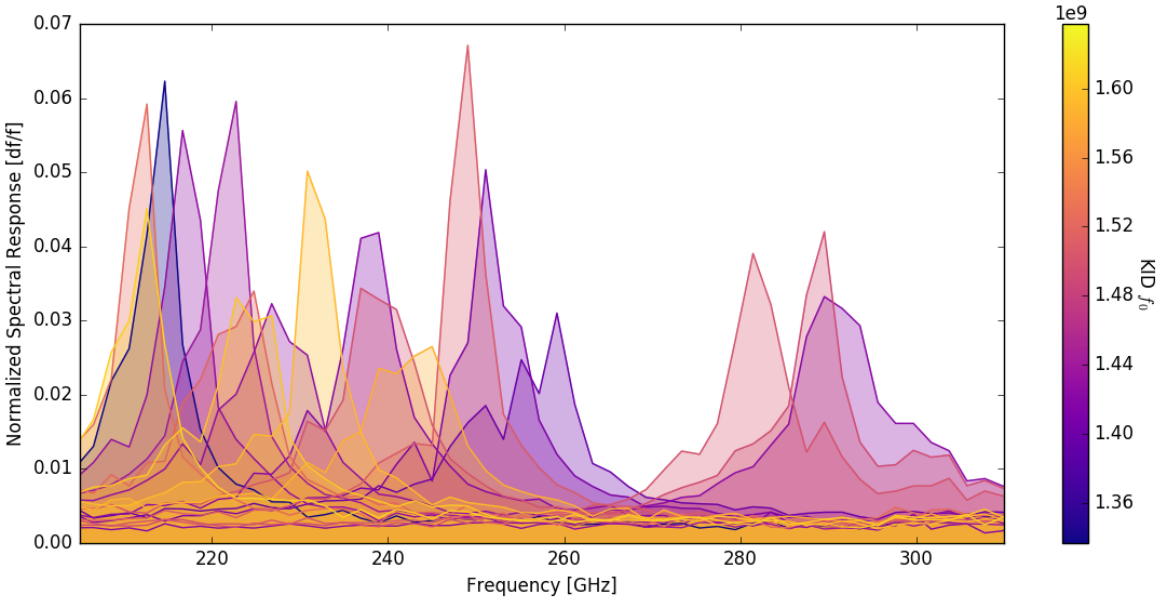

Figure 5.10: Spectral response of filter bank channels. The color of each individual curve represents the readout resonant frequency of the associated KID, as indicated by the colorbar on the right. The curves have been normalized so that the area under the curve for each is unity. The median resolution of the spectral channels is approximately $R = 45$, although there are several channels with much lower resolution, around $R = 30$. There is also a gap in the filter bank where a section of KIDs did not yield, which is further discussed in the text.

5.6 Discussion and Summary

We have demonstrated a prototype Advanced SuperSpec device with a novel mm-wave filter bank structure combined with a new Al KID design. This device exhibits moderate noise and response, with median white noise levels in a dark configuration at approximately $1.2 \times 10^{-17}\,\mathrm{WHz}^{-1/2}$ and a responsivity or $2.5 \times 10^{7}\,\mathrm{W}^{-1}$. Additionally, the KIDs have moderate quality factors, averaging approximately 3×10^{4} in a dark setting, with limited decay to 2.3×10^{4} under 77 K optical load. The device exhibits spectral response, with a spectrometer resolution of approximately $R = 45$ median. Unfortunately, there was low yield, 34%, which is likely attributed to a liftoff step in the fabrication processing. This can be avoided in future iterations with taller liftoff resist.

This basic design was developed in parallel and adopted into the South Pole Telescope summertime line intensity mapper (SPT-SLIM) experiment [Karkare et al., 2022a]. That the SPT-SLIM detectors follow essentially the same mm-wave coupling structures attached to Al KID detectors in the same style presented in this work, but operating at 120 GHz to 180 GHz, with the experiment set to deploy in the Antarctic summer of 2023.

The Advanced SuperSpec prototype has proven that a SuperSpec-like detector is possible to make with Al, but this design will need to be pushed further to hope to achieve low-background space platform limits. There is hope still here, as optimizations are at hand in the design of the Al inductor. Further reduction in inductor volume is possible by pushing the lithographic constraints to thinner tine widths. Future design iterations will allow for the tuning of the filter bank geometries to better align with the desired spectrometer resolution. Higher resolutions are readily available, which require only making the coupling stubs smaller in size.

REFERENCES

R. Adam, A. Adane, P. A. R. Ade, P. André, A. Andrianasolo, H. Aussel, A. Beelen, A. Benoît, A. Bideaud, N. Billot, O. Bourrion, A. Bracco, M. Calvo, A. Catalano, G. Coiffard, B. Comis, M. De Petris, F. X. Désert, S. Doyle, E. F. C. Driessen, R. Evans, J. Goupy, C. Kramer, G. Lagache, S. Leclercq, J. P. Leggeri, J. F. Lestrade, J. F. Macías-Pérez, P. Mauskopf, F. Mayet, A. Maury, A. Monfardini, S. Navarro, E. Pascale, L. Perotto, G. Pisano, N. Ponthieu, V. Revéret, A. Rigby, A. Ritacco, C. Romero, H. Roussel, F. Ruppin, K. Schuster, A. Sievers, S. Triqueneaux, C. Tucker, and R. Zylka. The NIKA2 large-field-of-view millimetre continuum camera for the 30 m IRAM telescope. *A & A*, 609: A115, January 2018. doi:10.1051/0004-6361/201731503.

R. Barends, H. L. Hortensius, T. Zijlstra, J. J. A. Baselmans, S. J. C. Yates, J. R. Gao, and T. M. Klapwijk. Contribution of dielectrics to frequency and noise of NbTiN superconducting resonators. *Applied Physics Letters*, 92(22):223502, 06 2008. ISSN 0003-6951. doi:10.1063/1.2937837. URL `https://doi.org/10.1063/1.2937837`.

P. S. Barry, E. Shirokoff, A. Kovács, T. J. Reck, S. Hailey-Dunsheath, C. M. McKenney, L. J. Swenson, M. I. Hollister, H. G. Leduc, S. Doyle, R. O'Brient, N. Llombart, D. Marrone, G. Chattopadhyay, P. K. Day, S. Padin, C. M. Bradford, P. D. Mauskopf, and J. Zmuidzinas. Electromagnetic design for SuperSpec: a lithographically-patterned millimetre-wave spectrograph. In Wayne S. Holland, editor, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*, volume 8452, page 84522F. International Society for Optics and Photonics, SPIE, 2012. doi:10.1117/12.927089. URL `https://doi.org/10.1117/12.927089`.

P. S. Barry, S. Doyle, A. L. Hornsby, A. Kofman, E. Mayer, A. Nadolski, Q. Y. Tang, J. Vieira, and E. Shirokoff. Design and performance of the antenna-coupled lumped-element kinetic inductance detector. *Journal of Low Temperature Physics*, 193(3-4):176–183, may 2018. doi:10.1007/s10909-018-1943-y. URL `https://doi.org/10.1007%2Fs10909-018-1943-y`.

Peter Barry. *On the Development of SuperSpec: A Fully Integrated On-Chip Spectrometer for Far-Infrared Astronomy*. PhD thesis, Cardiff University, 2014.

Baselmans, J. J. A., Bueno, J., Yates, S. J. C., Yurduseven, O., Llombart, N., Karatsu, K., Baryshev, A. M., Ferrari, L., Endo, A., Thoen, D. J., de Visser, P. J., Janssen, R. M. J., Murugesan, V., Driessen, E. F. C., Coiffard, G., Martin-Pintado, J., Hargrave, P., and Griffin, M. A kilo-pixel imaging system for future space based far-infrared observatories using microwave kinetic inductance detectors. *A&A*, 601:A89, 2017. doi:10.1051/0004-6361/201629653. URL `https://doi.org/10.1051/0004-6361/201629653`.

E.E. Becklin and L.J. Moon. Stratospheric observatory for infrared astronomy (sofia). *Advances in Space Research*, 30(9):2083–2088, 2002. ISSN 0273-1177. doi:https://doi.org/10.1016/S0273-1177(02)00591-4. URL `https://www.sciencedirect.com/science/article/pii/S0273117702005914`.

B. A. Benson, P. A. R. Ade, Z. Ahmed, S. W. Allen, K. Arnold, J. E. Austermann, A. N. Bender, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, J. F. Cliche, T. M. Crawford, A. Cukierman, T. de Haan, M. A. Dobbs, D. Dutcher, W. Everett, A. Gilbert, N. W. Halverson, D. Hanson, N. L. Harrington, K. Hattori, J. W. Henning, G. C. Hilton, G. P. Holder, W. L. Holzapfel, K. D. Irwin, R. Keisler, L. Knox, D. Kubik, C. L. Kuo, A. T. Lee, E. M. Leitch, D. Li, M. McDonald, S. S. Meyer, J. Montgomery, M. Myers, T. Natoli, H. Nguyen, V. Novosad, S. Padin, Z. Pan, J. Pearson, C. Reichardt, J. E. Ruhl, B. R. Saliwanchik, G. Simard, G. Smecher, J. T. Sayre, E. Shirokoff, A. A. Stark, K. Story, A. Suzuki, K. L. Thompson, C. Tucker, K. Vanderlinde, J. D. Vieira, A. Vikhlinin, G. Wang, V. Yefremenko, and K. W. Yoon. SPT-3G: a next-generation cosmic microwave background polarization experiment on the South Pole telescope. In Wayne S. Holland and Jonas Zmuidzinas, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, volume 9153 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 91531P, July 2014. doi:10.1117/12.2057305.

A. W. Blain and M. S. Longair. Submillimetre cosmology. *Monthly Notices of the Royal Astronomical Society*, 264(2):509–521, 09 1993. ISSN 0035-8711. doi:10.1093/mnras/264.2.509. URL https://doi.org/10.1093/mnras/264.2.509.

Thomas L. R. Brien, Peter A. R. Ade, Peter S. Barry, Edgar Castillo-Domìnguez, Daniel Ferrusca, Thomas Gascard, Victor Gómez, Peter C. Hargrave, Amber L. Hornsby, David Hughes, Enzo Pascale, Josie D. A. Parrianen, Abel Perez, Sam Rowe, Carole Tucker, Salvador Ventura González, and Simon M. Doyle. MUSCAT: the Mexico-UK Sub-Millimetre Camera for AsTronomy. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, volume 10708, page 107080M. International Society for Optics and Photonics, SPIE, 2018. doi:10.1117/12.2313697. URL https://doi.org/10.1117/12.2313697.

Sean Bryan. The TolTEC Camera for the LMT Telescope. In *Atacama Large-Aperture Submm/mm Telescope (AtLAST)*, page 36, January 2018. doi:10.5281/zenodo.1159073.

C.L. Carilli and F. Walter. Cool gas in high-redshift galaxies. *Annual Review of Astronomy and Astrophysics*, 51(1):105–161, 2013. doi:10.1146/annurev-astro-082812-140953. URL https://doi.org/10.1146/annurev-astro-082812-140953.

Caitlin M. Casey, Desika Narayanan, and Asantha Cooray. Dusty star-forming galaxies at high redshift. *Physics Reports*, 541(2):45–161, 2014. ISSN 0370-1573. doi:https://doi.org/10.1016/j.physrep.2014.02.009. URL https://www.sciencedirect.com/science/article/pii/S0370157314000477. Dusty star-forming galaxies at high-redshift.

E. Castillo-Dominguez, P. Ade, P. S. Barry, T. Brien, S. Doyle, D. Ferrusca, V. Gomez-Rivera, P. Hargrave, A. Hornsby, D. Hughes, P. D. Mauskopf, P. Moseley, E. Pascale,

A. Perez-Fajardo, G. Pisano, S. Rowe, C. Tucker, and M. Velazquez. Mexico-UK submillimeter camera for astronomy. *Journal of Low Temperature Physics*, 193(5-6):1010–1015, jul 2018. doi:10.1007/s10909-018-2018-9. URL https://doi.org/10.1007%2Fs10909-018-2018-9.

Jhy-Jiun Chang and D. J. Scalapino. Kinetic-equation approach to nonequilibrium superconductivity. *Phys. Rev. B*, 15:2651–2670, Mar 1977. doi:10.1103/PhysRevB.15.2651. URL https://link.aps.org/doi/10.1103/PhysRevB.15.2651.

George Che. *Advancements in Kinetic Inductance Detector, Spectrometer, and Amplifier Technologies for Millimeter-Wave Astronomy*. PhD thesis, Arizona State University, 2018.

Asantha Cooray and Origins Space Telescope Study Team. Origins Space Telescope. In *American Astronomical Society Meeting Abstracts #231*, volume 231 of *American Astronomical Society Meeting Abstracts*, page 103.01, January 2018.

A. T. Crites, J. J. Bock, C. M. Bradford, T. C. Chang, A. R. Cooray, L. Duband, Y. Gong, S. Hailey-Dunsheath, J. Hunacek, P. M. Koch, C. T. Li, R. C. O'Brient, T. Prouve, E. Shirokoff, M. B. Silva, Z. Staniszewski, B. Uzgil, and M. Zemcov. The TIME-Pilot intensity mapping experiment. In Wayne S. Holland and Jonas Zmuidzinas, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, volume 9153, page 91531W. International Society for Optics and Photonics, SPIE, 2014. doi:10.1117/12.2057207. URL https://doi.org/10.1117/12.2057207.

Peter K. Day, Henry G. LeDuc, Benjamin A. Mazin, Anastasios Vayonakis, and Jonas Zmuidzinas. A broadband superconducting detector suitable for use in large arrays. *Nature*, 425:817–821, 2003. doi:10.1038/nature02037.

P. J. de Visser, J. J. A. Baselmans, P. Diener, S. J. C. Yates, A. Endo, and T. M. Klapwijk. Generation-recombination noise: The fundamental sensitivity limit for kinetic inductance detectors. *Journal of Low Temperature Physics*, 167:335–340, 2012. doi:10.1007/s10909-012-0519-5.

Pieter de Visser. *Quasiparticle dynamics in aluminium superconducting microwave resonators*. PhD thesis, Delft University of Technology, 2014.

Dole, H., Lagache, G., Puget, J.-L., Caputi, K. I., Fernández-Conde, N., Le Floc'h, E., Papovich, C., Pérez-González, P. G., Rieke, G. H., and Blaylock, M. The cosmic infrared background resolved by spitzer - contributions of mid-infrared galaxies to the far-infrared background. *A&A*, 451(2):417–429, 2006. doi:10.1051/0004-6361:20054446. URL https://doi.org/10.1051/0004-6361:20054446.

Ran Duan, Sean McHugh, Bruno Serfass, Benjamin A. Mazin, A. Merrill, Sunil R. Golwala, Thomas P. Downes, Nicole G. Czakon, Peter K. Day, Jiansong Gao, Jason Glenn, Matthew I. Hollister, Henry G. Leduc, Philip R. Maloney, Omid Noroozian, Hien T. Nguyen, Jack Sayers, James A. Schlaerth, Seth Siegel, John E. Vaillancourt, Anastasios Vayonakis, Philip R. Wilson, and Jonas Zmuidzinas. An open-source readout for

MKIDs. In Wayne S. Holland and Jonas Zmuidzinas, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 77411V, July 2010. doi:10.1117/12.856832.

L. Earle, P. Ade, J. Aguirre, R. Aikin, J. Battle, J. Bock, C. M. Bradford, M. Dragovan, L. Duband, J. Glenn, G. Griffin, V. Hristov, P. Maloney, H. Matsuhara, B. Naylor, H. Nguyen, M. Yun, and J. Zmuidzinas. Z-Spec: a broadband direct-detection millimeter-wave spectrometer – instrument status and first results. In Jonas Zmuidzinas, Wayne S. Holland, Stafford Withington, and William D. Duncan, editors, *Millimeter and Submillimeter Detectors and Instrumentation for Astronomy III*, volume 6275, page 627510. International Society for Optics and Photonics, SPIE, 2006. doi:10.1117/12.672309. URL `https://doi.org/10.1117/12.672309`.

Jiansong Gao. *The Physics of Superconducting Microwave Resonators*. PhD thesis, California Institute of Technology, 2008.

Jiansong Gao, Miguel Daal, John M. Martinis, Anastasios Vayonakis, Jonas Zmuidzinas, Bernard Sadoulet, Benjamin A. Mazin, Peter K. Day, and Henry G. Leduc. A semiempirical model for two-level system noise in superconducting microresonators. *Applied Physics Letters*, 92(21):212504, 05 2008. ISSN 0003-6951. doi:10.1063/1.2937855. URL `https://doi.org/10.1063/1.2937855`.

Kurtis Lee Geerlings. *Improving Coherence of Superconducting Qubits and Resonators*. PhD thesis, Yale University, 2013.

Yan Gong, Asantha Cooray, Marta Silva, Mario G. Santos, James Bock, C. Matt Bradford, and Michael Zemcov. Intensity mapping of the [cii] fine structure line during the epoch of reionization. *The Astrophysical Journal*, 745(1):49, dec 2011. doi:10.1088/0004-637X/745/1/49. URL `https://dx.doi.org/10.1088/0004-637X/745/1/49`.

Samuel Gordon. *Highly multiplexed superconducting detectors and readout electronics for balloon-borne and ground-based far-infrared imaging and polarimetry*. PhD thesis, Arizona State University, 2019.

Samuel Gordon, Brad Dober, Adrian Sinclair, Samuel Rowe, Sean Bryan, Philip Mauskopf, Jason Austermann, Mark Devlin, Simon Dicker, Jiansong Gao, Gene C. Hilton, Johannes Hubmayr, Glenn Jones, Jeffrey Klein, Nathan P. Lourie, Christopher McKenney, Federico Nati, Juan D. Soler, Matthew Strader, and Michael Vissers. An open source, FPGA-based LeKID readout for BLAST-TNG: Pre-flight results. *Journal of Astronomical Instrumentation*, 05(04), dec 2016. doi:10.1142/s2251171716410038. URL `https://doi.org/10.1142%2Fs2251171716410038`.

C.J Gorter and H Casimir. On supraconductivity i. *Physica*, 1(1):306–320, 1934. ISSN 0031-8914. doi:https://doi.org/10.1016/S0031-8914(34)90037-9. URL `https://www.sciencedirect.com/science/article/pii/S0031891434900379`.

S. Hailey-Dunsheath, E. Shirokoff, P. S. Barry, C. M. Bradford, S. Chapman, G. Che, J. Glenn, M. Hollister, A. Kovács, H. G. LeDuc, P. Mauskopf, C. McKenney, R. O'Brient, S. Padin, T. Reck, C. Shiu, C. E. Tucker, J. Wheeler, R. Williamson, and J. Zmuidzinas. Low noise titanium nitride KIDs for SuperSpec: A millimeter-wave on-chip spectrometer. *Journal of Low Temperature Physics*, 184(1-2):180–187, dec 2015. doi:10.1007/s10909-015-1375-x. URL https://doi.org/10.1007%2Fs10909-015-1375-x.

Michael G. Hauser and Eli Dwek. The Cosmic Infrared Background: Measurements and Implications. *Annual Review of Astronomy and Astrophysics*, 39:249–307, jan 2001. doi:10.1146/annurev.astro.39.1.249.

R. Herrera-Camus, A. D. Bolatto, M. G. Wolfire, J. D. Smith, K. V. Croxall, R. C. Kennicutt, D. Calzetti, G. Helou, F. Walter, A. K. Leroy, B. Draine, B. R. Brandl, L. Armus, K. M. Sandstrom, D. A. Dale, G. Aniano, S. E. Meidt, M. Boquien, L. K. Hunt, M. Galametz, F. S. Tabatabaei, E. J. Murphy, P. Appleton, H. Roussel, C. Engelbracht, and P. Beirao. [cii] 158μm emission as a star formation tracer. *The Astrophysical Journal*, 800(1):1, feb 2015. doi:10.1088/0004-637X/800/1/1. URL https://dx.doi.org/10.1088/0004-637X/800/1/1.

D. J. Hollenbach and A. G. G. M. Tielens. Dense photodissociation regions (pdrs). *Annual Review of Astronomy and Astrophysics*, 35(1):179–215, 1997. doi:10.1146/annurev.astro.35.1.179. URL https://doi.org/10.1146/annurev.astro.35.1.179.

D. J. Hollenbach and A. G. G. M. Tielens. Photodissociation regions in the interstellar medium of galaxies. *Rev. Mod. Phys.*, 71:173–230, Jan 1999. doi:10.1103/RevModPhys.71.173. URL https://link.aps.org/doi/10.1103/RevModPhys.71.173.

S. B. Kaplan, C. C. Chi, D. N. Langenberg, J. J. Chang, S. Jafarey, and D. J. Scalapino. Quasiparticle and phonon lifetimes in superconductors. *Phys. Rev. B*, 14:4854–4873, Dec 1976. doi:10.1103/PhysRevB.14.4854. URL https://link.aps.org/doi/10.1103/PhysRevB.14.4854.

K. S. Karkare, P. S. Barry, C. M. Bradford, S. Chapman, S. Doyle, J. Glenn, S. Gordon, S. Hailey-Dunsheath, R. M. J. Janssen, A. Kovács, H. G. LeDuc, P. Mauskopf, R. McGeehan, J. Redford, E. Shirokoff, C. Tucker, J. Wheeler, and J. Zmuidzinas. Full-array noise performance of deployment-grade SuperSpec mm-wave on-chip spectrometers. *Journal of Low Temperature Physics*, 199(3-4):849–857, feb 2020. doi:10.1007/s10909-020-02407-4. URL https://doi.org/10.1007%2Fs10909-020-02407-4.

K. S. Karkare, A. J. Anderson, P. S. Barry, B. A. Benson, J. E. Carlstrom, T. Cecil, C. L. Chang, M. A. Dobbs, M. Hollister, G. K. Keating, D. P. Marrone, J. McMahon, J. Montgomery, Z. Pan, G. Robson, M. Rouble, E. Shirokoff, and G. Smecher. SPT-SLIM: A

line intensity mapping pathfinder for the south pole telescope. *Journal of Low Temperature Physics*, 209(5-6):758–765, mar 2022a. doi:10.1007/s10909-022-02702-2. URL https://doi.org/10.1007%2Fs10909-022-02702-2.

Kirit S. Karkare, Azadeh Moradinezhad Dizgah, Garrett K. Keating, Patrick Breysse, and Dongwoo T. Chung. Snowmass 2021 cosmic frontier white paper: Cosmology with millimeter-wave line intensity mapping, 2022b.

Attila Kovács, Peter S. Barry, Charles M. Bradford, Goutam Chattopadhyay, Peter Day, Simon Doyle, Steve Hailey-Dunsheath, Matthew Hollister, Christopher McKenney, Henry G LeDuc, Nuria Llombart, Daniel P. Marrone, Philip Mauskopf, Roger C. O'Brient, Stephen Padin, Loren J. Swenson, and Jonas Zmuidzinas. SuperSpec: design concept and circuit simulations. In Wayne S. Holland, editor, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*, volume 8452, page 84522G. International Society for Optics and Photonics, SPIE, 2012. doi:10.1117/12.927160. URL https://doi.org/10.1117/12.927160.

D. K. Lambert and P. L. Richards. Martin-puplett interferometer: an analysis. *Appl. Opt.*, 17(10):1595–1602, May 1978. doi:10.1364/AO.17.001595. URL https://opg.optica.org/ao/abstract.cfm?URI=ao-17-10-1595.

Henry G. Leduc, Bruce Bumble, Peter K. Day, Byeong Ho Eom, Jiansong Gao, Sunil Golwala, Benjamin A. Mazin, Sean McHugh, Andrew Merrill, David C. Moore, Omid Noroozian, Anthony D. Turner, and Jonas Zmuidzinas. Titanium nitride films for ultrasensitive microresonator detectors. *Applied Physics Letters*, 97(10), sep 2010. doi:10.1063/1.3480420. URL https://doi.org/10.1063%2F1.3480420.

X. Liu, W. Guo, Y. Wang, M. Dai, L. F. Wei, B. Dober, C. M. McKenney, G. C. Hilton, J. Hubmayr, J. E. Austermann, J. N. Ullom, J. Gao, and M. R. Vissers. Superconducting micro-resonator arrays with ideal frequency spacing. *Applied Physics Letters*, 111(25): 252601, 12 2017. ISSN 0003-6951. doi:10.1063/1.5016190. URL https://doi.org/10.1063/1.5016190.

Nathan P. Lourie, Peter A. R. Ade, Francisco E. Angile, Peter C. Ashton, Jason E. Austermann, Mark J. Devlin, Bradley Dober, Nicholas Galitzki, Jiansong Gao, Sam Gordon, Christopher E. Groppi, Jeffrey Klein, Gene C. Hilton, Johannes Hubmayr, Dale Li, Ian Lowe, Hamdi Mani, Philip Mauskopf, Christopher M. McKenney, Federico Nati, Giles Novak, Enzo Pascale, Giampaolo Pisano, Adrian Sinclair, Juan D. Soler, Carole Tucker, Joel N. Ullom, Michael Vissers, and Paul A. Williams. Preflight characterization of the BLAST-TNG receiver and detector arrays. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, volume 10708, page 107080L. International Society for Optics and Photonics, SPIE, 2018. doi:10.1117/12.2314396. URL https://doi.org/10.1117/12.2314396.

Piero Madau and Mark Dickinson. Cosmic star-formation history. *Annual Review of Astronomy and Astrophysics*, 52(1):415–486, 2014. doi:10.1146/annurev-astro-081811-125615. URL https://doi.org/10.1146/annurev-astro-081811-125615.

D. C. Mattis and J. Bardeen. Theory of the anomalous skin effect in normal and superconducting metals. *Phys. Rev.*, 111:412–417, Jul 1958. doi:10.1103/PhysRev.111.412. URL https://link.aps.org/doi/10.1103/PhysRev.111.412.

Benjamin A. Mazin. *Microwave Kinetic Inductance Detectors*. PhD thesis, California Institute of Technology, 2004.

R. McGeehan, P. S. Barry, E. Shirokoff, C. M. Bradford, G. Che, J. Glenn, S. Gordon, S. Hailey-Dunsheath, M. Hollister, A. Kovács, H. G. LeDuc, P. Mauskopf, C. McKenney, T. Reck, J. Redford, C. Tucker, J. Turner, S. Walker, J. Wheeler, and J. Zmuidzinas. Low-Temperature Noise Performance of SuperSpec and Other Developments on the Path to Deployment. *Journal of Low Temperature Physics*, 193(5-6):1024–1032, December 2018. doi:10.1007/s10909-018-2061-6.

Christopher M. McKenney, Jason E. Austermann, James A. Beall, Bradley J. Dober, Shannon M. Duff, Jiansong Gao, Gene C. Hilton, Johannes Hubmayr, Dale Li, Joel N. Ullom, Jeff L. Van Lanen, and Michael R. Vissers. Tile-and-trim micro-resonator array fabrication optimized for high multiplexing factors. *Review of Scientific Instruments*, 90(2), feb 2019. doi:10.1063/1.5037301. URL https://doi.org/10.1063%2F1.5037301.

Felipe Menanteau, Jorge González, Jean-Baptiste Juin, Tobias A. Marriage, Erik D. Reese, Viviana Acquaviva, Paula Aguirre, John William Appel, Andrew J. Baker, L. Felipe Barrientos, Elia S. Battistelli, J. Richard Bond, Sudeep Das, Amruta J. Deshpande, Mark J. Devlin, Simon Dicker, Joanna Dunkley, Rolando Dünner, Thomas Essinger-Hileman, Joseph W. Fowler, Amir Hajian, Mark Halpern, Matthew Hasselfield, Carlos Hernández-Monteagudo, Matt Hilton, Adam D. Hincks, Renée Hlozek, Kevin M. Huffenberger, John P. Hughes, Leopoldo Infante, Kent D. Irwin, Jeff Klein, Arthur Kosowsky, Yen-Ting Lin, Danica Marsden, Kavilan Moodley, Michael D. Niemack, Michael R. Nolta, Lyman A. Page, Lucas Parker, Bruce Partridge, Neelima Sehgal, Jon Sievers, David N. Spergel, Suzanne T. Staggs, Daniel Swetz, Eric Switzer, Robert Thornton, Hy Trac, Ryan Warne, and Ed Wollack. The atacama cosmology telescope: Physical properties and purity of a galaxy cluster sample selected via the sunyaev–zel'dovich effect. *The Astrophysical Journal*, 723(2):1523, oct 2010. doi:10.1088/0004-637X/723/2/1523. URL https://dx.doi.org/10.1088/0004-637X/723/2/1523.

C. Neill, A. Megrant, R. Barends, Yu Chen, B. Chiaro, J. Kelly, J. Y. Mutus, P. J. J. O'Malley, D. Sank, J. Wenner, T. C. White, Yi Yin, A. N. Cleland, and John M. Martinis. Fluctuations from edge defects in superconducting resonators. *Applied Physics Letters*, 103 (7):072601, 08 2013. ISSN 0003-6951. doi:10.1063/1.4818710. URL https://doi.org/10.1063/1.4818710.

J.R. Pardo, J. Cernicharo, and E. Serabyn. Atmospheric transmission at microwaves (atm): an improved model for millimeter/submillimeter applications. *IEEE Transactions on Antennas and Propagation*, 49(12):1683–1694, 2001. doi:10.1109/8.982447.

P.C. Cortes. *ALMA Technical Handbook*. ALMA Doc. 8.4, ver 1.0, 2023. ISBN 978-3-923524-66-2.

David M. Pozar. *Microwave Engineering*. John Wiley & Sons, Inc., 2011.

Joseph Redford. *SuperSpec: Development and Characterization of a Full Band On-Chip Spectrometer and Preparations for On-Sky Observations*. PhD thesis, California Institute of Technology, 2023.

S. Rowe, M. Tapia, P. S. Barry, K. S. Karkare, A. Papageorgiou, P. A. R. Ade, T. L. R. Brien, E. Castillo-Domínguez, D. Ferrusca, V. Gómez-Rivera, P. Hargrave, J. L. Hernández-Rebollar, A. Hornsby, J. M. Jáuregui-García, P. Mauskopf, D. Murias, E. Pascale, A. Pérez, M. W. L. Smith, C. Tucker, M. Velázquez, S. Ventura, D. H. Hughes, and S. Doyle. The MUSCAT Readout Electronics Backend: Design and Predeployment Performance. *Journal of Low Temperature Physics*, 211(5-6):289–301, June 2023. doi:10.1007/s10909-022-02868-9.

F. Peter Schloerb and Luis Carrasco. The Large Millimeter Telescope. In Jacobus M. Oschmann Jr., editor, *Ground-based Telescopes*, volume 5489, pages 754 – 762. International Society for Optics and Photonics, SPIE, 2004. doi:10.1117/12.551935. URL `https://doi.org/10.1117/12.551935`.

Erik Shirokoff, Peter Barry, Charles Bradford, Goutam Chattopadhyay, Peter Day, S. Doyle, S. Hailey-Dunsheath, Matthew Hollister, Andras Kovacs, Christopher McKenney, Henry Leduc, Nuria Llombart, Daniel Marrone, Philip Mauskopf, Roger O'Brient, Sara Padin, T. Reck, Loren Swenson, and Jonas Zmuidzinas. Mkid development for superspec: An on-chip, mm-wave, filter-bank spectrometer. *Proceedings - Society of Photo-Optical Instrumentation Engineers*, 8452:84520R, 11 2012. doi:10.1117/12.927070.

Corwin Shiu. Design considerations to improving sensitivity in superspec: an onchip kid-based, mm-wave spectrometer.

S. R. Siegel. *A Multiwavelength Study of the Intracluster Medium and the Characterization of the Multiwavelength Sub/millimeter Inductance Camera*. PhD thesis, California Institute of Technology,, 2016.

G. J. Stacey, N. Geis, R. Genzel, J. B. Lugten, A. Poglitsch, A. Sternberg, and C. H. Townes. The 158 Micron [C ii] Line: A Measure of Global Star Formation Activity in Galaxies. *The Astrophysical Journal*, 373:423, June 1991. doi:10.1086/170062.

Steven H. Strogatz. *Nonlinear Dynamics and Chaos*. Westview Press, 2015.

L. J. Swenson, P. K. Day, B. H. Eom, H. G. Leduc, N. Llombart, C. M. McKenney, O. Noroozian, and J. Zmuidzinas. Operation of a titanium nitride superconducting microresonator detector in the nonlinear regime. *Journal of Applied Physics*, 113(10):104501, 03 2013. ISSN 0021-8979. doi:10.1063/1.4794808. URL https://doi.org/10.1063/1.4794808.

Qing Yang Tang. *Development of Kinetic Inductance Detectors for mm and Sub-mm Observations*. PhD thesis, University of Chicago, 2021.

J. Wheeler, S. Hailey-Dunsheath, E. Shirokoff, P. S. Barry, C. M. Bradford, S. Chapman, G. Che, J. Glenn, M. Hollister, A. Kovács, H. G. LeDuc, P. Mauskopf, R. McGeehan, C. M. McKenney, R. O'Brient, S. Padin, T. Reck, C. Ross, C. Shiu, C. E. Tucker, R. Williamson, and J. Zmuidzinas. SuperSpec: development towards a full-scale filter bank. In Wayne S. Holland and Jonas Zmuidzinas, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*, volume 9914, page 99143K. International Society for Optics and Photonics, SPIE, 2016. doi:10.1117/12.2233798. URL https://doi.org/10.1117/12.2233798.

Jordan Wheeler. *Millimeter-Wave Galaxy Spectroscopy: SuperSpec On-Chip Spectrometer Technology Development and ALMA Observations of Molecular Gas in the Arp 220 Nucleus*. PhD thesis, University of Colordado, 2019.

Jordan Wheeler, S. Hailey-Dunsheath, E. Shirokoff, P. S. Barry, C. M. Bradford, S. Chapman, G. Che, S. Doyle, J. Glenn, S. Gordon, M. Hollister, A. Kovács, H. G. LeDuc, P. Mauskopf, R. McGeehan, C. McKenney, T. Reck, J. Redford, C. Ross, C. Shiu, C. Tucker, J. Turner, S. Walker, and J. Zmuidzinas. SuperSpec, The On-Chip Spectrometer: Improved NEP and Antenna Performance. *Journal of Low Temperature Physics*, 193 (3-4):408–414, November 2018. doi:10.1007/s10909-018-1926-z.

C. M. Wilson and D. E. Prober. Quasiparticle number fluctuations in superconductors. *Phys. Rev. B*, 69:094524, Mar 2004. doi:10.1103/PhysRevB.69.094524. URL https://link.aps.org/doi/10.1103/PhysRevB.69.094524.

Jonas Zmuidzinas. Superconducting microresonators: Physics and applications. *Annual Review of Condensed Matter Physics*, 3(1):169–214, 2012. doi:10.1146/annurev-conmatphys-020911-125022. URL https://doi.org/10.1146/annurev-conmatphys-020911-125022.