Instrumentation for Kinetic-Inductance-Detector-Based Submillimeter Radio Astronomy

Thesis by

Ran Duan

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California

2015
(Defended June 21, 2013)

Acknowledgments

The past five and a half years at Caltech have been an important part of my life. Looking back on these years, standing both at an endpoint and a starting point in my life, I can confidently and proudly say that I have enjoyed every moment at Caltech, particularly the research life within our group and the many people I have met who have become close friends. It has been 5 years full of discovery, excitement, and achievement, all of which I appreciate with a grateful heart.

I finally have the opportunity to express my appreciation to Sunil, the best advisor I can imagine, who was one of the most important aspects of my time at Caltech and one of the kindest people I have ever met. I feel extremely lucky and grateful to have worked with him.

I greatly appreciate the support of and discussions with Jonas Zmuidzinas, my co-advisor at Caltech, who is always full of knowledge and wisdom. He is the role model I want to become one day.

I have benefitted greatly from Sandy Weinreb, who was the first person to contact me from the Electrical Engineering Department and who hired me. I will never forget the moment when I received my offer letter from him on my birthday in 2008.

I would also like to thank my PhD committee member Dr. David Rutledge, who has been extremely helpful and has offered useful suggestions about my work, and Dr. Ali Hajimiri for his support and guidance, as this would not be possible without his kind support and understanding. I would also like to thank Dr. Di Li, who shared ambitions and dreams with me.

I would also like to express my sincere appreciation to my undergraduate advisor, Prof. YiLong Lu, from Singapore, for his heartfelt support. I also greatly appreciate the help and advice that Jason Glenn and Jamie Bock have offered along the way. Glenn Jones has been a good resource for discussing readout problems. With him, I have learned a great deal from one of the best people in the field.

The development of an instrument is a significant undertaking. I have learned a considerable amount from our group members: Seth Siegel, Nicole Czakon, Tasos Vayonakis, Dave Miller, Dave Moore, Jack Sayers, Tom Downes, Matt Hollister, Shwetank Kumar, Phil Maloney, Roger OBrient, Loren Swenson, Chris Mckenney, Jacob Kooi, Matthew Sumner, Rick Leduc, Peter Day, James Schlaerth, Clint Bockstiegel, Spencer Brugger, Amandeep Gill, and Rick Raffanti.

I would like to thank Dr. Ben Mazin, Dr. Sean McHugh, and Dr. Bruno Serfass for not only for working with me on electronics problems but also for the ping-pong games. I would particularly like to thank the members of the CASPER Group for their international support of my work and for sharing their experiences with me. This collaboration inspired me to develop the open-source readout platform and to propagate the collaborative spirit.

I would like to offer my gratitude to my friend Jiansong Gao, who provided suggestions as a true friend. I would also like to thank my officemate and good friend, Omid Noroozian, for his camaraderie.

I would like to thank our secretaries, Kathy Deniston, Sheri Stoll, Susan McCurdy, and Barbara Wertz, and my friends in the group, Larry Beirich and Mona Delitsky, who have been an important part of my life here.

My family has been my primary source of encouragement and inspiration. I would like to thank my family for supporting me in the pursuit of my ambitions.

Abstract

A substantial amount of important scientific information is contained within astronomical data at the submillimeter and far-infrared (FIR) wavelengths, including information regarding dusty galaxies, galaxy clusters, and star-forming regions; however, these wavelengths are among the least-explored fields in astronomy because of the technological difficulties involved in such research. Over the past 20 years, considerable efforts have been devoted to developing submillimeter- and millimeter-wavelength astronomical instruments and telescopes.

The number of detectors is an important property of such instruments and is the subject of the current study. Future telescopes will require as many as hundreds of thousands of detectors to meet the necessary requirements in terms of the field of view, scan speed, and resolution. A large pixel count is one benefit of the development of multiplexable detectors that use kinetic inductance detector (KID) technology.

This dissertation presents the development of a KID-based instrument including a portion of the millimeter-wave bandpass filters and all aspects of the readout electronics, which together enabled one of the largest detector counts achieved to date in submillimeter-/millimeter-wavelength imaging arrays: a total of 2304 detectors. The work presented in this dissertation has been implemented in the MUltiwavelength Submillimeter Inductance Camera (MUSIC), a new instrument for the Caltech Submillimeter Observatory (CSO).

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Scientific Motivation

The results of the Cosmic Background Explorer (COBE) indicate that 50% of the luminosity and 98% of the photons emitted by the Big Bang fall within the submillimeter and far-infrared (FIR) ranges, as shown in Figure 1.1. The National Research Council's Decadal Survey (New Worlds, New Horizons in Astronomy and Astrophysics 2010) emphasizes that "a high priority in the coming decade will be to undertake large and detailed surveys of galaxies as they evolve across the wide interval of cosmic time" (pp. 218) and to develop, in effect, "a 13-billion year-long movie that traces the build-up of structure since the universe first became transparent to light" (pp. 2–13).

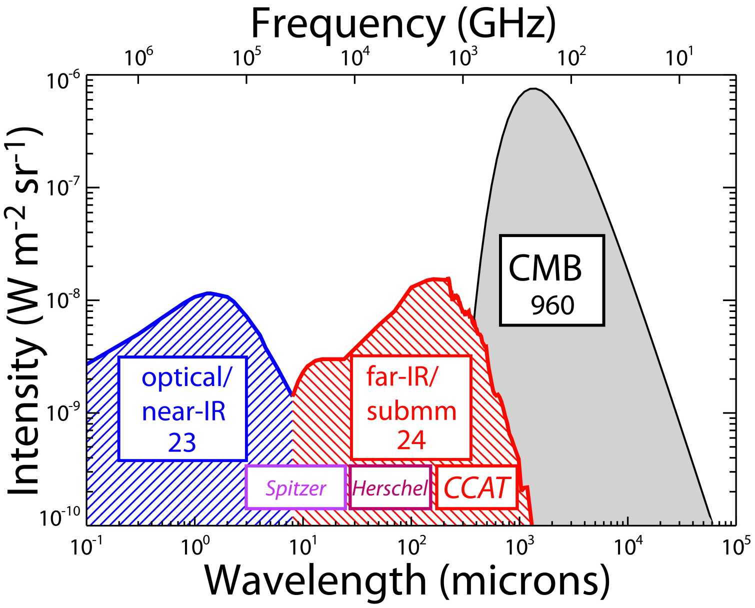

Figure 1.1: Integrated brightness in cosmic optical, infrared, and microwave background radiation. Reprinted from reference [1].

Considerable effort has been devoted to studying this topic over the years, including ground-based, balloon-borne, airborne, and space observatory measurements. Detection technology has also rapidly developed for constructing and upgrading instruments. The MUSIC instrument that we designed and constructed (Figure 1.2 shows the major components of the instrument, which will be discussed in detail in the following sections) is ideally suited to address the priorities identified in the report and may serve as a guide for future submillimeter- and millimeter-wavelength telescopes such as the Cerro Chajnantor Atacama telescope (CCAT).

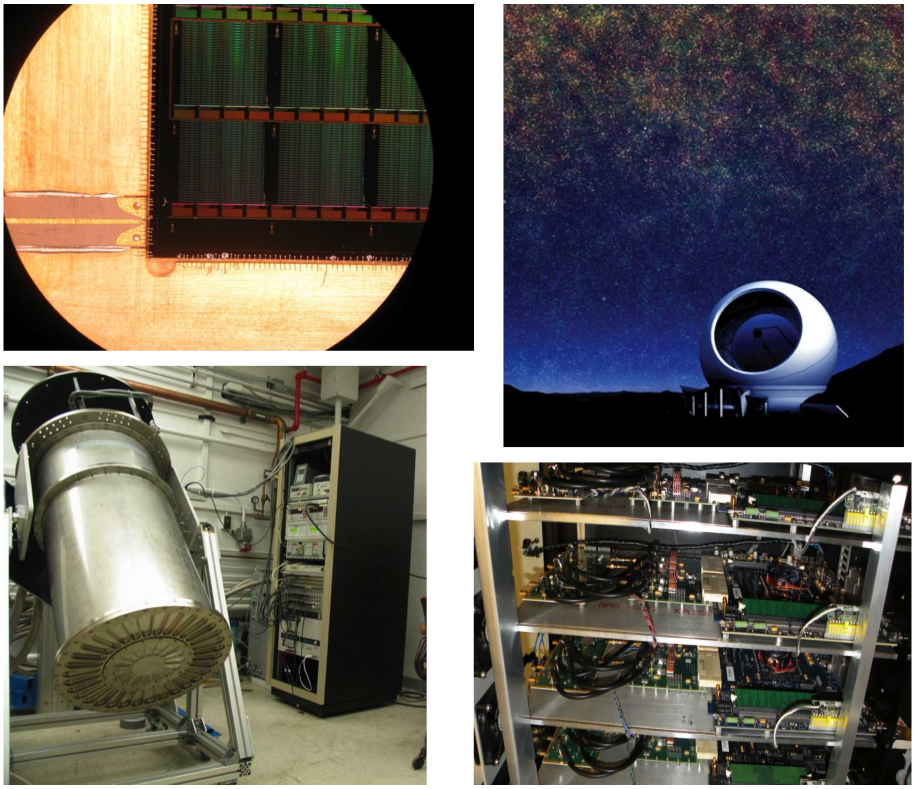

Figure 1.2: Photographs of the various components of the telescope instrument. (Upper left) The detector wafer that serves as the camera of the telescope instrument. (Lower left) The cryostat that houses the detector wafer and cold electronics. (Lower right) Room-temperature readout electronics that read the signal from the detector wafer. (Upper right) The telescope in which the instrument is used for astronomical observations.

MUSIC is designed to have a high mapping speed, a high angular resolution, and four frequency bands probed by kinetic inductance detectors (KIDs). As the newest instrument in the Caltech Submillimeter Observatory, MUSIC is among the first microwave KID cameras and is the camera with the largest number of detectors sensitive to the submillimeter-wavelength range. MUSIC can be used to observe the Sunyaev–Zeldovich (SZ) effect in galaxy clusters, dusty star-forming galaxies, and dark matter halos to address fundamental questions regarding the large-scale structure of the universe and the history of star formation over cosmic time.

1.2 Submillimeter Astronomy Instruments

Over the past few years, submillimeter/millimeter astronomical instruments have revolutionized our understanding of the formation of stars, galaxies, and clusters by measuring the fundamental parameters of our universe. Thus, these instruments are critical to helping us gain an understanding of nature and to continue to make new discoveries.

One commonly used detector technology for submillimeter wavelengths is the transition edge sensor (TES), which is a cryogenic sensor based on the temperature-dependent resistance of the superconducting phase transition. To read the signal from a TES, a superconducting quantum interference device (SQUID) is paired with the detector. TESs have been used for the detection of submillimeter/millimeter wavelengths in many types of instruments; however, their complex fabrication process and readout method make them difficult to scale to larger arrays. Another relatively new technology is the KID, which was developed at the Caltech/Jet Propulsion Laboratory (JPL) in the early 2000s. KIDs can be easily fabricated on a two- to three-layer wafer and frequency-domain multiplexed; all of the readout functions are performed by room-temperature electronics, with the exception of one cryogenic amplifier. KIDs are ideal for large-array implementation, which will be necessary for future telescope development; for instance, the CCAT will require millions of detectors, as shown in Figure 1.3, to fully exploit its large field of view and small arcsecond resolution.

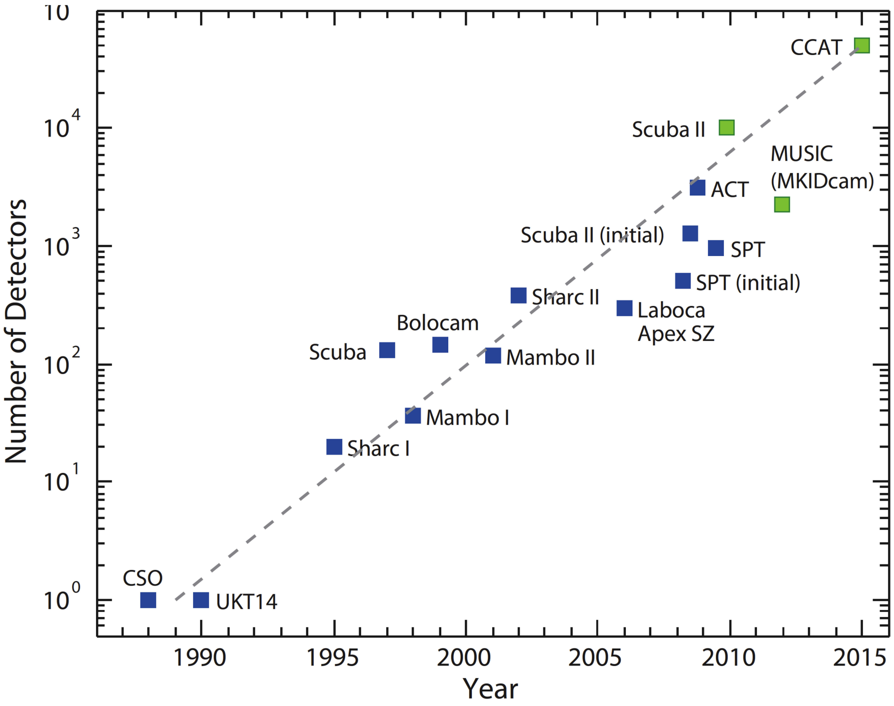

Figure 1.3: The growth in the size of submillimeter detector arrays as a function of time over the past two decades. The blue points represent existing instruments, and the green points are projections at the time the plot was created. Plot courtesy of Jonas Zmuidzinas.

The instrument that we constructed using KID technology, MUSIC, covers wavelengths of 0.87, 1.04, 1.33, and 1.98 mm. From a technological perspective, the MUSIC design successfully demonstrates a photolithographic focal plane, as illustrated in Figure 1.4. This figure provides an overview of the MUSIC focal-plane wafer, which contains a broadband phased-array antenna for beam definition, a bank of four bandpass filters (BPFs) for band selection, and a microwave KID (MKID) for power detection. The details of the KID and BPFs will be discussed in Chapters 1.3 and 3, respectively. MUSIC contains the largest number of detectors of any existing KID instrument (a total of 2304 detectors) and offers large-array readout.

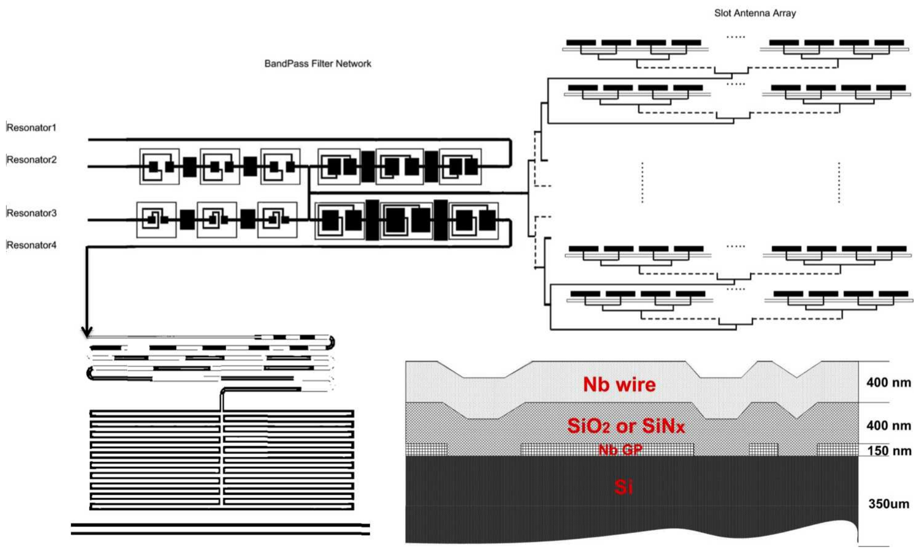

Figure 1.4: Scaled version of the antenna, BPF network, and KID layout. (Upper right) A phased-array antenna with a binary summing tree is used to capture the signal. (Upper left) Four bandpass filters split the signal from the summing tree into four different frequency bands. (Lower left) The kinetic inductance detector layout. (Lower right) Substrate layers of the detector wafer.

MUSIC provides compelling opportunities for scientific applications including studying galaxy clusters at all redshifts at which they exist, searching for the highest-redshift contributors to the FIR background by looking for objects that are bright in the MUSIC bands but not in the SPIRE bands, and surveying protostellar cores and young stellar objects. MUSIC is also complementary to existing telescopes such as ALMA, Herschel, and Planck. It provides long-wavelength spectral-energy distribution measurements for the establishment of stricter constraints on the luminosities and dust temperatures of the high-z galaxies detected by Herschel, offering a unique probe of the high-z galaxy population and providing finding charts to the astronomical community. By providing multicolor measurements and through the combination with data from ALMA and Herschel, MUSIC survey data will stringently test models of the formation and evolution of galaxies and the dark-matter halos in which they form.

1.3 Multicolor and Multiplexing Capabilities of MUSIC

Two important features of MUSIC are its multicolor detection, which is achieved through the lithographic design of the antenna and BPF bank, and its frequency-multiplexed imaging array, which is achieved through multiplexed KIDs and a sophisticated readout system. These two topics are the primary focus of this dissertation.

1.3.1 Multicolor Detection with a Bandpass Filter Bank

The MUSIC imaging wafer is sensitive to four frequency bands centered at 150, 225, 288, and 345 GHz. These frequency bands are defined by a BPF bank and are optimized using a dusty-galaxy signal source.

The multiwavelength capabilities of MUSIC provide unique advantages for radio astronomy. The large-scale surveys enabled by the multicolor nature and rapid mapping speed of MUSIC will allow for major advances in high-redshift astrophysics, complementing extensive ongoing surveys. Simultaneous consistent flux calibration across all four bands will be essential for identifying high-z galaxies based on their colors. MUSIC surveys will provide optimized catalogs for ALMA by observing equatorial fields with extensive multiwavelength coverage, particularly the Herschel Multi-tiered Extragalactic Survey (HerMES) fields.

The accuracy of the BPF frequency response is critical for high-sensitivity submillimeter applications such as MUSIC. An error of several percentage points will seriously degrade the signal source throughput and increase the level of irremovable crosstalk with the atmosphere and among the bands. This dissertation will elaborate on the filter design for MUSIC in Chapter 3.

1.3.2 Working Principle of MUSIC

To understand the camera functions, we must discuss the wafer design illustrated in Figure 1.4. A signal is first received by the phased-array antenna and then split into four frequency bands defined by the BPF network. The output of each filter is fed into one LC resonator, which is coupled to the feedline.

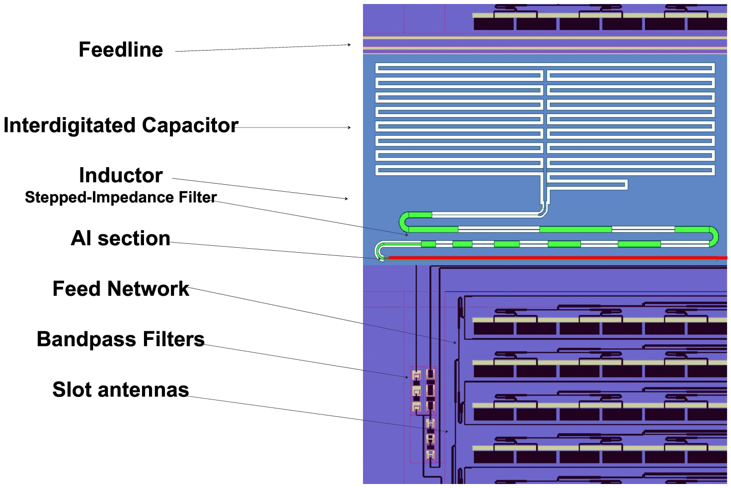

Figure 1.5: Mask design: layout of the major components on the wafer. Plot courtesy of the MUSIC group.

Figure 1.5 depicts the fabrication layout of a KID wafer (cf. Figure 1.4; Figure 1.5 shows the correct relative sizes of the components). This figure shows one interdigitated capacitor and one inductor, which together form one resonator that is coupled to the feedline in the figure. The aluminum section (the red-colored section in Figure 1.5) is also part of the inductor. Along the inductor microstrip line, we use a stepped-impedance filter (SIF) rather than a straight stripline. The SIF is a high-pass filter that prevents radiation above 70 GHz (near the Al gap energy) received by the capacitor from being routed to the Al portion of the inductor. This figure shows only part of the phased-array antenna; each feed stub of the antenna is summed together by the binary summing tree feed network. The feed-network binary tree for each slot antenna ends in one microstrip line, which carries the entire signal and feeds it into a BPF bank (this figure shows three BPFs), whose output is connected to the LC resonators.

The phased-array antenna and BPFs are designed to receive signals within the submillimeter-frequency range (100 to 400 GHz), and the LC resonator is designed to resonate within the microwave-frequency range (3 to 4 GHz). They are linked by a section of aluminum, which receives the signal from the antenna and is part of the inductor.

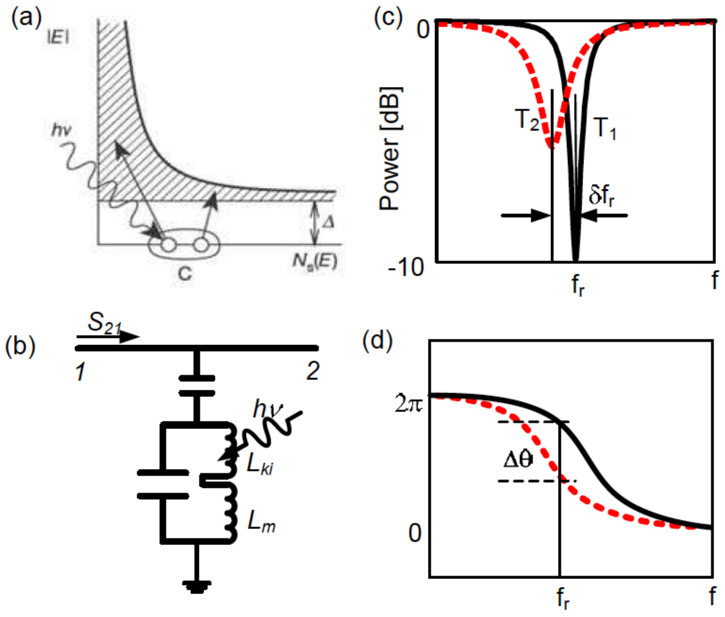

Figure 1.6: Schematic illustration of MKID operation [2, 3]. (a) A photon breaks Cooper pairs and creates quasiparticles in a superconducting strip. (b) The increase in the quasiparticle density changes the surface impedance. (c) The transmission through the resonant circuit exhibits a narrow dip at the resonance frequency f_r, which shifts when the surface impedance changes. (d) The microwave probe signal experiences a phase shift when f_r changes. Reprinted from [4].

As shown in Figure 1.6, when the aluminum superconductor receives photons, the photons break Cooper pairs and create quasiparticles, which change the surface impedance of the aluminum superconductor. Because the aluminum section is part of the meander inductor of the LC resonator, the photons cause the inductance of the LC resonator to change, thereby affecting its resonance properties including the phase and power (i.e., a measure of how the resonator couples to the microstrip line). By monitoring the resonance properties, we can detect signals at submillimeter wavelengths. We call such an LC resonator a microwave KID.

Overall, the KID-based instrument that we developed is an ideal choice for use in submillimeter instruments because of the following features:

1. The detectors are frequency-domain multiplexed with hundreds of detectors coupled through one feedline.

2. The process for fabricating a large detector array is relatively simple, consisting of two to four layers of a lithographic process on a silicon wafer.

3. The readout system for this large detector array has been developed and demonstrated to satisfy large-array requirements.

These advantages not only make the KID-based instrument scalable to larger arrays but also allow for cost control on a per-pixel basis.

In addition to the BPF design, this dissertation will discuss all aspects of the readout system to provide a detailed description of the instrument's development in the next chapter.

Chapter 2

Readout for Kinetic Inductance Detectors

This chapter will present a systematic study of the design, implementation, and performance analysis of an open-source readout (OSR) system for a superconducting microresonator array. The OSR system performs frequency-domain multiplexed real-time complex transmission measurements to monitor the instantaneous resonance frequency and dissipation of the KIDs. With a total of 16 readout units, our OSR system can read more than 3000 complex frequency channels simultaneously. All hardware, software, and firmware were successfully installed, tested, and optimized at the Caltech Submillimeter Observatory with the first MKID camera, MUSIC, in 2010 and 2012. The system demonstrated its ability to satisfy the requirements for detector readout, data acquisition, and telescope operation. As part of the MUSIC instrument, the OSR has been in use at the CSO for scientific observation since the summer of 2013.

2.1 Background

The direct multiplexing of TES detectors, the other detector technology that is widely used at submillimeter and millimeter wavelengths, requires several biasing wires and one amplifier per detector. In contrast, only one input and output transmission line and one cryogenic amplifier are required for the readout of an entire KID array, thus simplifying the design of the readout circuits at the cold stage.

Figure 2.1 shows the S_{12} measurements at the two ends of the transmission line for a KID array, where each dip corresponds to one resonance frequency.

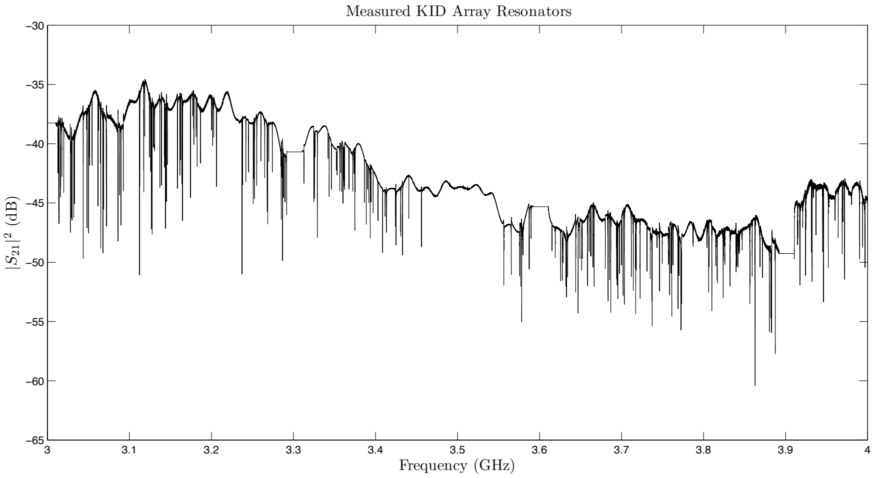

Figure 2.1: Multiplexed KID concept: S_{12} of the measured KID wafer exhibiting multiple resonance frequencies. Plot courtesy of the MUSIC group.

We multiplex many detectors using a single transmission line by lithographically setting slightly different resonance frequencies for each KID. The signal source (e.g., submillimeter, X-ray, FIR, ultraviolet [UV], and optical signals) modifies the microwave frequency response of the resonator. This response can be measured using a single cryogenic amplifier and sophisticated room-temperature electronics, thus shifting the complexity and challenge of the detector-array readout to the room-temperature electronics.

Since it was first proposed at Caltech/JPL, KID technology has developed rapidly as a result of its numerous advantages and potential applications [2, 5–11]. The readout system has been demonstrated by a number of groups [11–15]. In general, the procedure for the KID OSR is as follows:

1. A drive tone or carrier tone (which generally ranges from approximately 100 MHz to a few gigahertz) is sent through the transmission line on the device with which the detector is coupled. The drive tone is then modulated by the detector as it responds to the power of the astronomical signal.

2. If the modulated carrier tone is at a frequency that is too high to be directly digitized, it is down-converted using a mixer.

3. The received tone is digitized and processed. This step includes digitizing the analog signal and processing the digital signal using a field-programmable gate array (FPGA), a graphics processing unit (GPU), a central processing unit (CPU), or some combination thereof. The signal from the detector generally requires real-time processing to capture the signal source, reduce the data rate, and extract useful information.

4. Auxiliary information such as a timestamp or information regarding the telescope is stored on the data acquisition (DAQ) computer.

5. Frequently, the raw data must be subjected to additional computer-based processing steps prior to use. These steps may include serialization of the data stream, noise reduction, transformation of the data format or units, calibration, or mapmaking.

Using the MUSIC readout as an example, this chapter will attempt to encompass most aspects of the readout design and highlight various interesting aspects of the detector readout electronics.

2.1.1 Basics of OSR Development with MUSIC

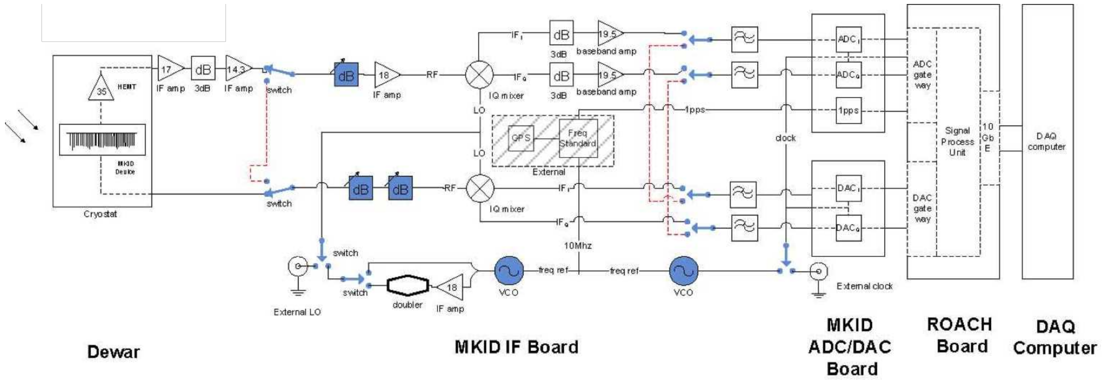

Figure 2.2: Block diagram of the full readout system.

Figure 2.2 presents an overview of the MUSIC readout system. The basic concept of resonator readout is to use IQ homodyne mixing, which is essentially a dual-phase lock-in detection technique. In general, the signal forms a closed loop with an FPGA to be the start and end points; along one path in the FPGA, the readout electronics send frequency tones to the device in the cryostat [the lower half of the diagram, which includes the digital-to-analog converter (DAC), IQ mixer, and digital attenuators], and along the other path [the upper half of the diagram, which includes the amplifier, attenuator, IQ mixer, and analog-to-digital converter (ADC)], the readout electronics receive the signal from the cryostat output and process the signal. Both signal paths utilize the DAQ computer, the signal processing board (an FPGA board known as Reconfigurable Open Architecture Computing Hardware [ROACH] [16–21], an ADC/DAC board, and an intermediate frequency board are shown in the diagram, which will be explained in detail in the following sections), and a cryostat.

To read out the signal from the detector wafer, a comb of probe frequencies is generated for all resonators coupled to a single transmission line. This frequency comb is sent through the KID array, where the array modifies the amplitude and phase of the comb signals on the basis of the change in the surface impedance of the superconductor caused by the incident photons.

After amplification by a cryogenic amplifier such as a high-electron-mobility transistor (HEMT) or silicon–germanium bipolar-junction transistor, the comb frequencies are transmitted to room-temperature electronics for digitization and analysis. There are no cryogenic components other than the cryogenic amplifier, attenuators, and KID wafer. For a resonator operating at a few gigahertz, a microwave probe signal can be generated by up-converting (mixing with a local oscillator) the baseband signal produced by sending a preprogrammed waveform stored on a memory block through a fast DAC card. Similarly, on the receiver side, the microwave signal is first down-converted by a mixer and then digitized by a fast ADC card. Recent advances in software-defined radio technology have provided additional options for fast signal processing. For example, the signal processing for a digitized signal can be performed digitally using algorithms on an FPGA.

Prior to 2009, detector readout in our group was performed using off-the-shelf equipment, and only a few detectors could be read out. Between 2009 and 2010, we developed an open-source FPGA-based readout. We constructed prototypes of the DAC and ADC boards (Section 2.3.2) and demonstrated the simultaneous readout of 126 detectors. From 2010 to 2012, we developed intermediate frequency (IF) boards (Section 2.3.3), which helped us integrate the electronics and improve stability. Other advances included the second generation of the combined DAC–ADC board, an improved version of the firmware on an FPGA, expansion of the scale from 1 to 16 boards, and the production of a full set of DAQ software. The objective was to develop an OSR that could handle all of the tasks required for KID readout with a high level of automation.

This chapter summarizes the work our group has performed with respect to readout electronics since 2009. In the summers of 2010 and 2012, we conducted two successful engineering runs at the CSO. The readout electronics have been intensively tested and are now in operation at the CSO for scientific observation.

2.2 Readout Design and Development

2.2.1 Noise Requirements and Calculation

For the majority of KID readout applications, the signal power at the detector is less than −70 dBm. Therefore, amplification is required in the signal-receiving chain before the room-temperature electronics. To ensure that the entire signal-receiving chain provides good signal-to-noise performance, a high-gain and low-noise component must be placed at the front of the signal-receiving chain. Cryogenic amplifiers (HEMT or SiGe) are the best components available for this purpose. In general, a cryogenic amplifier has a noise temperature of 2–5 K across the frequency band inhabited by our resonators. Therefore, the noise of the readout electronics is designed such that the white noise of the HEMT amplifier is the dominant contribution relative to the rest of the electronics noise.

From the readout power of the carrier tone for each detector, we can calculate the signal-to-noise ratio (SNR) requirement at the HEMT. In our case, each resonator has a readout power between 10 and 30 pW at the device before the HEMT, meaning that the 144 resonators have a total power of −58.4 dBm. Inside the cryostat, the HEMT gain is 35 dB, and a conservative value of the HEMT noise temperature, which places the most stringent requirements possible on the remainder of the electronics, is 2 K. The SNR at the HEMT output is approximately 56.68 dB, as calculated in Section 2.5.1. We carry this SNR throughout the following procedures.

To develop the MUSIC readout system, we carefully designed the components to minimize the SNR degradation between the cryostat and the ADC to 1–2 dB, as discussed in Section 2.8.5. The SNR requirement in front of the ADC is approximately 55.94 dB for a 2-K HEMT noise temperature. We obtained suitable commercially available ADC chips that satisfied the SNR and bandwidth requirements with a margin of a few decibels.

2.2.2 Bandwidth and Center-Frequency Considerations

The bandwidth of every OSR system is a tradeoff between the wafer design (how compactly the resonators can be packed with respect to the frequency) and the electronics limitations due to the SNR, the bandwidth, and the maximum number of detectors that the readout electronics can process. In our case, each resonator has a bandwidth of approximately 200–400 kHz[1]. We established a 2-MHz separation between the resonators, which is an intentionally large margin to account for shifts in the resonator position caused by fabrication errors. The relative positions of the resonators will not change during observation or cool-down cycles, but the fabrication variation can cause them to shift in position relative to one another.

In the previous section, we determined that the ADC should have an SNR greater than 55.9 dB. If we consider only quantization noise, then this SNR would theoretically require an ADC with at least 10 bits. However, the noise of fast ADCs is greater than the quantization noise expected on the basis of the number of bits; thus, in practice, we require a 12-bit ADC. The best 12-bit chip commercially available in 2009 has a 64-dB SNR and a sampling rate of up to 550 MSPS (chip model number TI ADS5486). Once we confirmed this choice of ADC chip, we expanded the separation between the resonators to 2.5 MHz.

The center-frequency requirement for the readout is similar to the center-frequency requirement for the KIDs themselves and commonly ranges from 2 to 8 GHz. For MUSIC, the initial KID design had a resonator frequency of approximately 3–4 GHz, and the OSR system we developed is therefore also centered in that range. More recent KID designs can operate at frequencies as low as approximately 100 MHz. A lower frequency is attractive because it provides greater responsiveness

---

[1] On the basis of the quality factor required for MUSIC, the coupling quality factors range from 60k to 110k. We found that Qc = 60k maximizes the mapping speed. Moreover, on the basis of a resonance frequency of 3 GHz (which is the KID resonance frequency that we knew how to design at that time), we determined that the full width at half maximum (FWHM) of the resonator was approximately 100 kHz [22].

in the frequency direction and thus lowers the two-level system (TLS) noise. Our group is also conducting extensive research on KIDs of this type.

2.2.3 Readout of Both In-Phase and Quadrature-Phase Components

This section explains why we use two ADC and two DAC chips (which we often refer to as I and Q, respectively, because these chips are used with IQ mixers) and the relationship between the in-phase and quadrature-phase components.

Regardless of the signal-processing method used, we are limited by the Nyquist sampling theorem. We use two chips to cover a wider bandwidth, thereby providing us with a useful bandwidth that is equal to the full sampling rate of 550 MHz rather than merely the Nyquist bandwidth of 275 MHz. If we only needed to read the amplitude or the phase information of the frequency tone, various DSP methods could be implemented, such as recovering the amplitude information from the data stream (ignoring the phase component) to enable the full 550-MHz bandwidth to be read using only one DAC and one ADC chip. However, in this case, phase information would no longer be available. In summary, two ADCs and two DACs are employed to obtain both the amplitude and phase information of the full sampling bandwidth.

For detectors resonating at a few gigahertz, a mixer is required to convert the frequency from the baseband to the gigahertz band. An IQ mixer is a natural alternative to two single-sideband mixers for the simultaneous conversion of two signals. The use of an IQ mixer to utilize the full sampling bandwidth suggests that the baseband signal can be generated as DAC I and DAC Q, where I and Q have a phase difference of 90 degrees. In a complex data stream, I and Q serve as the real and imaginary parts of the complex data. If we perform a Fourier transform on a complex IQ data stream in the time domain, we obtain the full 550-MHz bandwidth in the frequency domain.

The electronics emit (from the DACs) and receive (by the ADCs) the drive tones with both the amplitude and phase information[2]. The submillimeter signal from the sky is converted into a change in the surface impedance of the superconducting inductor. We obtain the signal by monitoring the drive tones that couple with the resonators.

2.3 Hardware

2.3.1 Overview of the OSR Hardware, Software, and Firmware

The OSR system can be divided into three parts: hardware (Section 2.3), firmware (Section 2.4), and software (Section 2.6). In general, the hardware includes a customized ADC/DAC board, an IF board, an FPGA-based signal processing board, and an auxiliary system [e.g., a frequency standard

---

[2] The amplitude and phase information can be captured from the in-phase and quadrature-phase components, respectively

or global positioning system (GPS)]. Firmware refers to the program that runs on the FPGA chip, whereas software includes all of the programming implemented to control and automate the readout. Each part of the OSR system will be discussed in detail in the following sections.

2.3.2 ADC and DAC Boards

2.3.2.1 Chip Selection

There are several qualified ADC chips on the market, and the SNR and sampling-rate requirements for ADC chips have been discussed in previous sections. In addition to these requirements, the spurious-free dynamic range (SFDR) and intermodulation distortion (IMD) must also be considered to prevent harmonics or spurs from impacting the SNR of the resonator, particularly as the number of tones being read increases. We selected chips with a random spur frequency power level well below any level that could affect our resonators. For MUSIC, the signal considered here is located in the range of 0.1–10 Hz near the carrier tones. The resonators occupy a very small fraction of the entire radio-frequency (RF) bandwidth; therefore, the probability of a spur occurring within the resonator signal bandwidth is very low. However, for some applications such as the KID dark-matter detector developed at Caltech, the resonators must be monitored at higher frequencies (of a few kilohertz) for pulse detection. In this case, the harmonics or intermodulated frequencies are within a power level and frequency range that could affect the detection results of the resonators; therefore, the IMD and harmonics must be considered when generating carrier tones, e.g., by designing the drive tones to avoid these harmonics and IMD frequencies.

The choice of DACs is more flexible than the choice of ADCs; 16-bit and 1-GSPS DACs are readily and commercially available. We ultimately chose DAC5681 from Analog Devices, Inc. (Norwood, MA, US) and ADS5486 from Texas Instruments (Dallas, TX, US). The DAC can operate at up to 1 GSPS with a measured SNR of 75 dBFS, and the ADC can operate at up to 550 MSPS with a measured SNR of 64 dBFS. The 16-bit DAC and 12-bit ADC were both evaluated to confirm that they satisfied the SNR, SFDR, and IMD requirements. As a result of the rapid development of new semiconductor chips, faster ADCs and DACs have already appeared on the market, such as the 12-bit 3-GSPS ADC chip announced by Analog Devices, Inc. These semiconductor advances will be taken into consideration for future ADC/DAC development.

The ROACH boards that we selected use Zdok connectors (as shown in Figures 2.4 and 2.5) to connect to the ADC/DAC boards, allowing us to develop the ADC/DAC board independently and to use the ADC/DAC boards for different cameras and applications.

2.3.2.2 Clock and 1-PPS Signals on the Hardware Boards

ADC and DAC boards require a clock frequency to operate. This clock frequency can be provided by either the FPGA or by an external clock source such as a standalone frequency synthesizer. An external clock provides the flexibility to choose any clock frequency and more stable performance (to satisfy high timing requirements) than that provided by an FPGA clock. We decided to use an external clock for both the ADC and DAC chips. We also chose to use the ADC's external clock for the FPGA to achieve better phase performance and synchronization with the ADC, DAC, and FPGA. We integrated the clock-frequency generation and local oscillator (LO) frequency generation on our IF board (which will be discussed in detail in Section 2.3.3). The frequency stability originates from the 10-MHz reference, which is fed into the IF board and used to lock the clock and LO.

To synchronize the system, we added a synchronization port to the ADC board. This port provides a link from an external one-pulse-per-second (PPS) signal generated by a GPS device to the FPGA fabric through the same Zdok connector used by the ADC and DAC boards.

2.3.2.3 Two Generations of ADC/DAC Boards

The low-frequency noise in the system is extremely important because we desire stability of the electronics system for 10 s or longer, which is not an important consideration for many room-temperature semiconductor manufacturers.

The first-generation board used independent voltage references for the ADCs and DACs, which added low-frequency noise arising from shifts in the two voltage references relative to one another. After testing the first-generation board, we also determined that it would not be possible to use the direct-current (DC) supply from the FPGA boards for the ADC/DAC board and that, instead, an external clean power supply would be necessary; moreover, the temperature of the ADC chips was found to increase quickly, which caused the low-frequency noise to worsen at high ADC chip temperatures.

We developed a second-generation ADC/DAC board with a common voltage reference (chip model ADR441), which substantially improved the low-frequency stability. The layout of this board is illustrated in Figure 2.3. In addition, the new board was designed to have substantially improved thermal stability. Other advances included an external voltage supply with a regulator as well as modification of the positions of the SMA connectors to facilitate mating the ADC/DAC board to the IF board.

Figures 2.4 and 2.5 show the layout of the MUSIC ADC/DAC board. The corresponding components are labeled in the layout. The figures show the following components of the board: two DAC chips, two ADC chips, a clock for the ADC, a clock for the DAC, the 1-PPS input, the Zdok connectors, and the common voltage reference.

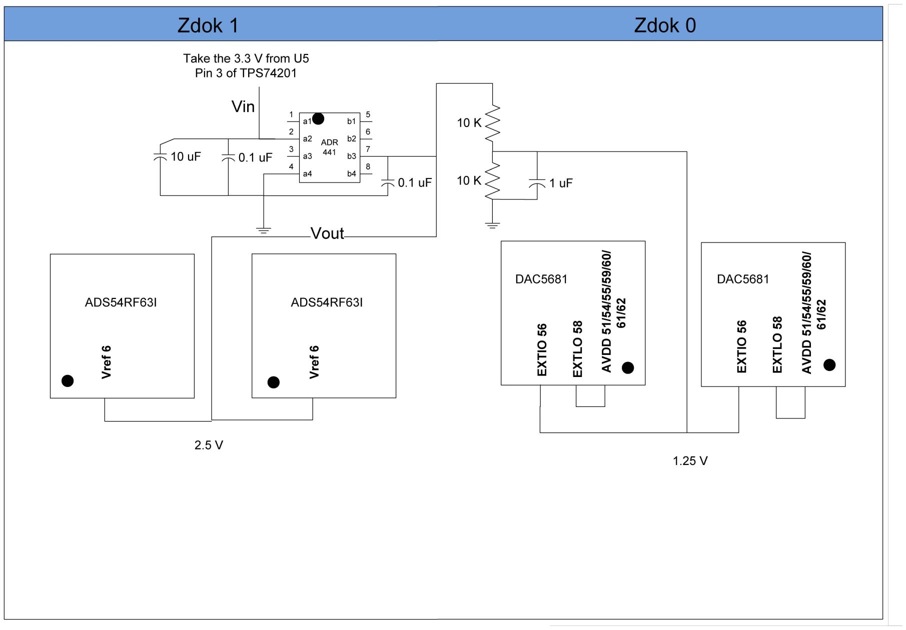

Figure 2.3: Layout of the common voltage reference for the ADC/DAC board.

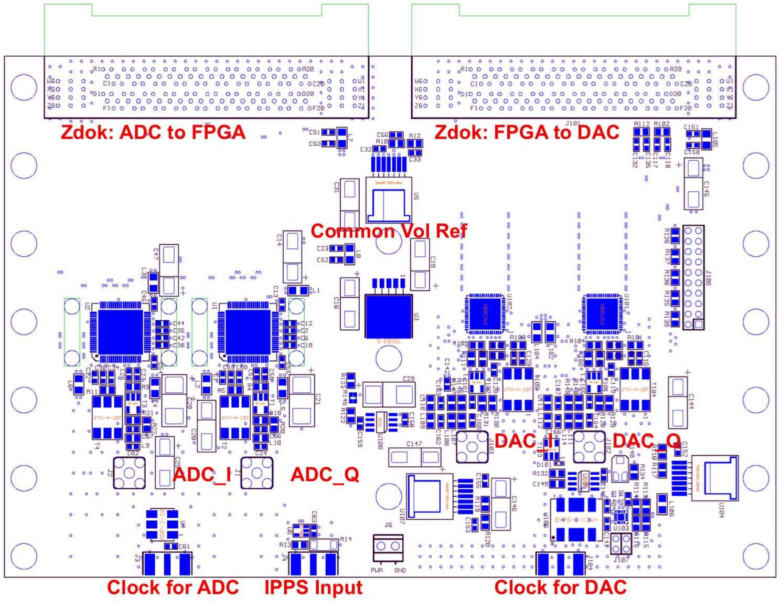

Figure 2.4: MUSIC ADC/DAC board layout design.

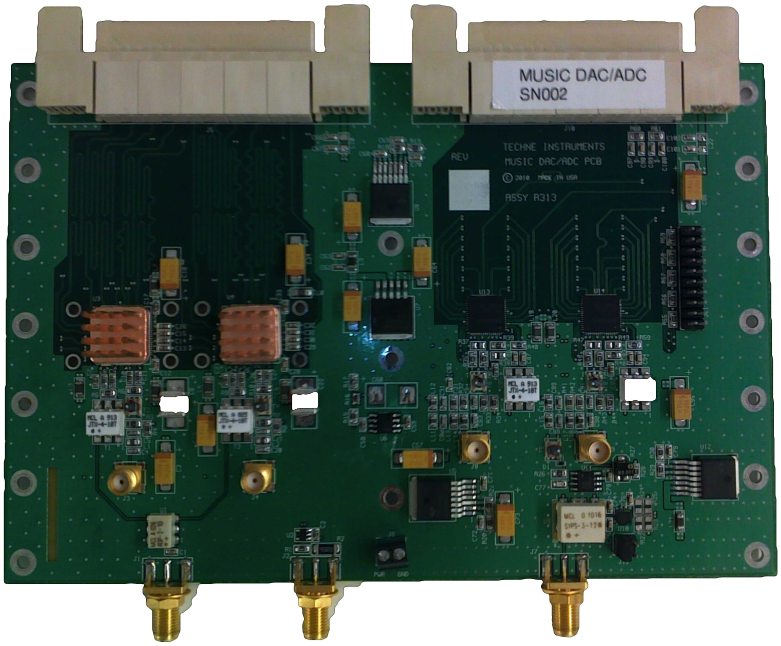

Figure 2.5: Photograph of the MUSIC ADC/DAC board.

2.3.2.4 DAC Roll-Off

The DAC output experiences frequency-dependent roll-off as a result of the following factors:

1. The anti-aliasing filter, a Mini-Circuits low-pass filter (LPF) with a cut-off at 250 MHz and a DAC sampling rate of 500 MSPS. Its S_{21} decreases from 0 dB at 8 MHz to -0.4 dB at 200 MHz.

2. The DAC chip itself. The output from even an ideal DAC suffers from the fact that the output is held constant for each clock cycle rather than smoothly changing. This causes the programmed DAC output to be modified by a sinc function in the frequency space. For a 500-MHz sampling rate, the sinc function causes -0.004 dB of attenuation at 8 MHz and -1.516 dB of attenuation at 200 MHz, with decreased attenuation for higher sampling rates.

3. The transformer (which functions as a high-pass filter) on the DAC board, which causes the outputs to roll off below 5 MHz.

The combination of the above factors means that the DAC output may experience a roll-off of approximately 2 dB across the full bandwidth. To output well-behaved frequency tones to optimally drive the resonator on the camera wafer and produce the optimal sensitivity, the roll-off factors listed above were considered and compensated for by programming the carrier-tone power in the generation buffer accordingly.

2.3.3 IF Board

Because the DAC and ADC work with baseband signals (DC to the Nyquist sampling rate) and because our resonator operates at microwave frequencies (a few gigahertz), the baseband signals must be up-converted to or down-converted from the resonator frequency band. The basic hardware components of the OSR system are described in Figure 2.2 and Section 2.1.

After the DAC board emits the frequency carrier signal, the signal passes through an LPF, an IQ up-converting mixer, and a digital attenuator before entering the dewar. On the receiver side, the carrier signal from the device is amplified by the HEMT and by room-temperature low-noise amplifiers (LNAs) before passing through a digital attenuator and being down-converted by another IQ mixer to the baseband. The down-converted carrier signal is then amplified, Nyquist filtered, and finally digitized by the DAC.

Many integrated circuit (IC) components are used in the intermediate stage to convert the frequency, control the power level, and perform a loopback test. The purpose of an IF board is to integrate all of the hardware components between the ADC/DAC board and the dewar.

In 2011, we designed and fabricated the IF board. Figures 2.6 and 2.7 present the IF board fabrication layout, which contains the following:

- two voltage-controlled oscillators (VCOs), which provide the FPGA clock signal (ranging from 137.5 to 4400 MHz) and the IQ mixer LO signal (ranging from 2.2 to 4.4 GHz or, with the doubler, from 4.4 to 5 GHz);
- two IQ mixers, which convert between the baseband and resonator frequency ranges;
- two digital attenuators (each with attenuation from 0 to 31.5 dB in 0.5-dB steps) to set power levels for transmission into the dewar; one digital attenuator and five amplifiers that serve the function of setting power levels for reception by the IQ mixer, and reception by the ADC.
- nine digital switches that allow the signal to be looped back at various points in the signal chain (e.g., baseband loopback, which bypasses the entire up-conversion/down-conversion and dewar chain, and RF loopback, which bypasses the dewar) and provide the option of choosing an external clock, an external LO, or doubling the LO frequency;
- one frequency doubler to double the LO frequency for higher resonant-frequency applications; and
- four anti-aliasing LPFs.

All of these functions of the IF board can be digitally controlled or programmed by the ROACH board. The control signal is sent from the FPGA and connected to the GPIO connector on the IF board via a 20-pin ribbon cable. The control firmware is designed to be independent and can be

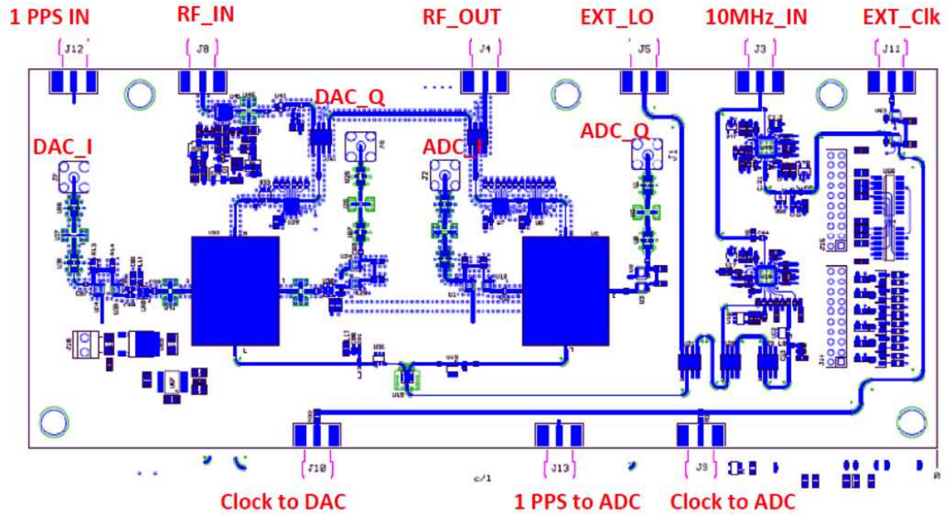

Figure 2.6: MUSIC IF board layout.

added to any existing firmware; thus, the IF board can be fully controlled or reprogrammed while the channelizing firmware is running. Each component on the IF board was carefully selected and configured to ensure that the noise level that reaches the ADC is dominated by the HEMT noise, and all other components in the system, such as the amplifiers and ADC, add a negligible amount of noise. By virtue of the IF board and a carefully designed DAC buffer, the probe signal that reaches the MKID device is optimized for each individual resonator across the entire readout bandwidth in both frequency and amplitude.

A noise study of the IF board will be presented in Section 2.8.5. The details of the IF board test are described in Appendices A and B.

2.3.4 FPGA Board

The FPGA processing board that we used was developed by the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) and is a standalone FPGA board known as ROACH (three connected circuit boards are shown in Figure 2.8; the ROACH board is on the left, the ADC/DAC board is in the middle, and the IF board is on the right. All three boards are mounted together on an Al plate to prevent relative vibration and flexing of the connections and to provide a good heat sink for the system).

The central component of the ROACH board is the Xilinx Vertex 5 FGPA. An independent PowerPC (PPC) running the Linux operating system is used to control the FPGA. In addition to the memory on the FPGA, one double data rate (DDR) dynamic random-access memory (DRAM) module and two quad data rate (QDR) static RAM modules are used to provide extra memory capacity to store the lookup table (LUT; this is the buffer played back by the DAC), provide the fast

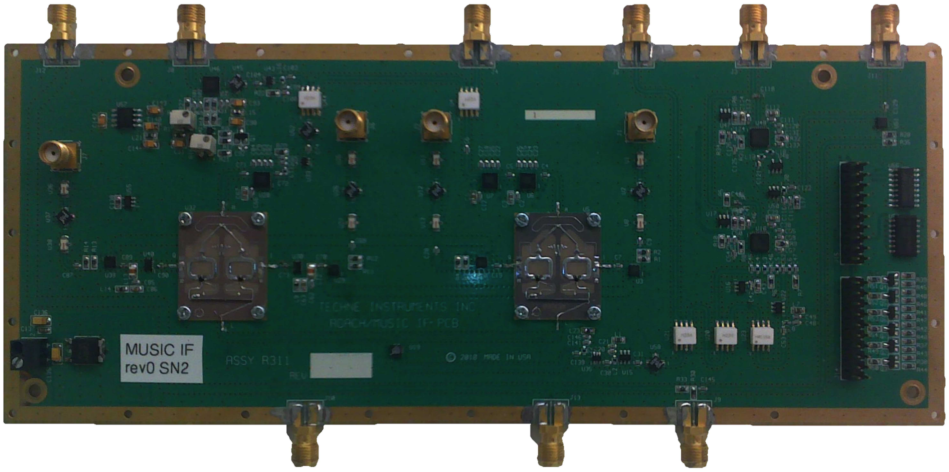

Figure 2.7: Photograph of the MUSIC IF board.

Fourier transform (FFT) coefficients, and buffer intermediate data during FPGA signal processing; two Zdok connectors allow the DAC, the ADC, or other interfaces to attach to the FPGA; four CX4 input and output connectors provide a data rate of up to 40 Gbps to transfer data or link multiple ROACHes together; and a 1-Gbps Ethernet port can be used by the DAQ computer for data communication and control of the PPC on the board [23].

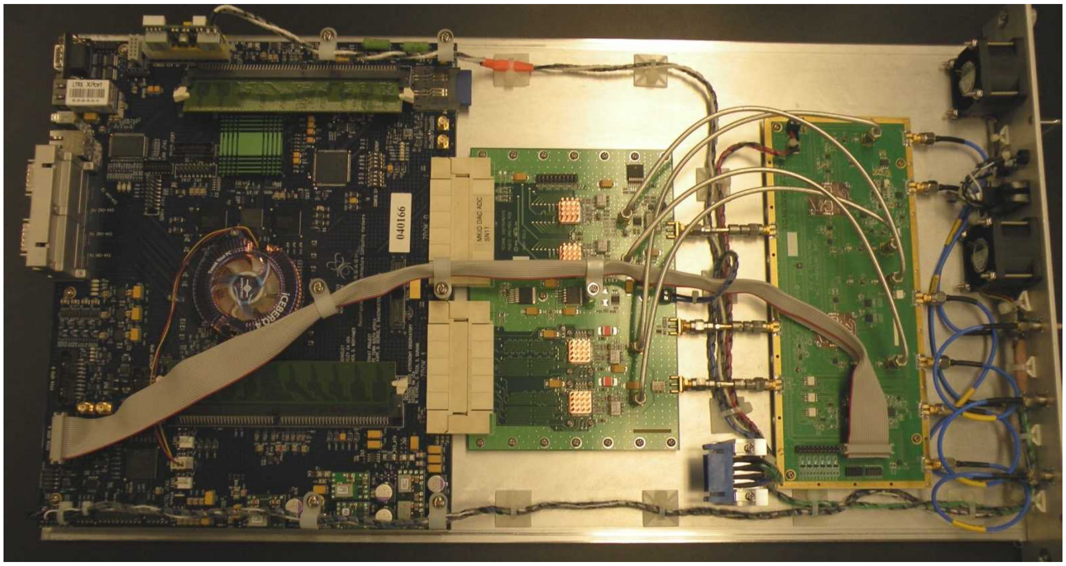

Figure 2.8: Photograph of the readout hardware mounted on the Al plate.

2.3.5 Auxiliary System

As previously mentioned, the readout electronics require two auxiliary signal inputs: a 1-PPS signal and 10-MHz frequency reference. The IF board directly uses the 10-MHz standard to generate the LO and FPGA/ADC/DAC clocks. The 1-PPS signal is routed through the IF board and ADC/DAC board to the ROACH board to be used for time synchronization. A GPS receiver (used by the entire CSO; model number Garmin LVC 18x) generates a 1-PPS signal. This 1-PPS signal is received by a 10-MHz rubidium frequency standard (Stanford Research Systems [SRS] FS725). The 10-MHz standard is locked to the 1-PPS signal. It averages over a long timescale and outputs a stabilized 1-PPS signal (as well as the 10-MHz reference). We use distribution amplifiers (SRS FS735/1/1 for the 10-MHz signal and SRS FS735/3/3 for the 1-PPS signal) to generate 16 copies of the 1-PPS and 10-MHz signals, which are routed to the individual OSR units. We use the GPS-sourced 1-PPS signal to ensure precise synchronization with the telescope pointing timestream and precise knowledge of the absolute time. Details of the synchronization of the time samples to an absolute time reference will be discussed in Section 2.4.5.

2.4 FPGA Firmware

The previous section described the hardware in the system. This section will discuss the development of the firmware for the OSR (i.e., the program running on the FPGA chip).

A flow chart that provides an overview of the firmware behavior is presented in Figure 2.9. The lower diagram shows the DAC LUT and DAC gateway, which is the direction in which the FPGA plays back the LUT buffer and sends the signal to the DAC board. This pathway includes a serialization step to provide the DAC with four DAC clock cycles of data during each FPGA clock cycle. The upper diagram shows the signal receiving and processing chain on the FPGA. On the receiving side, the ADC outputs the data in parallelized (deserialized) format, i.e., four ADC clock cycles for each FPGA clock cycle. The corresponding data streams from the I and Q ADCs are added to obtain a complex number, and the four complex streams are then processed by four parallel FFT cores (housed in two Biplex FFT blocks). Of the 2^{16} bins output by the four parallel FFTs, we are interested only in the 192 that have carrier tones; therefore, the data stream passes through a frequency-bin selector, and then, the 192 data streams pass through a 192-channel FIR filter to decimate the data rate by a factor of 75, which makes the final data rate exactly 100 Hz. Before the data are sent out of the FPGA, we package the data stream into a transmission control protocol (TCP) packet with a timestamp.

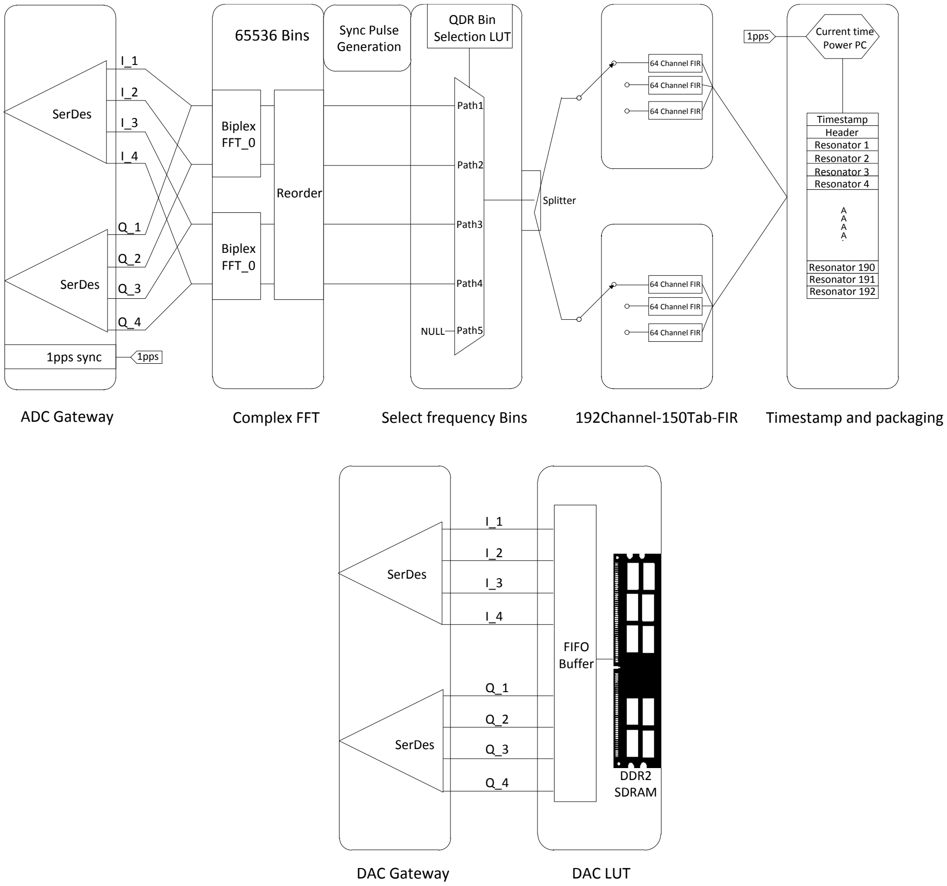

Figure 2.9: Flow chart of the FPGA firmware: (top) from the ADC input to the FPGA output and (bottom) lookup table and DAC.

2.4.1 Signal-Tone LUT Generation

The DAC plays back a stored LUT to send out a waveform that consists of approximately 200 sine waves with predefined frequencies and phases. This section describes how this LUT is generated.

2.4.1.1 Frequency-Tone Resolution and LUT Storage

The length of the LUT is defined by how precisely we need to drive the resonator tones. The criterion is that the drive tone should match the resonance frequency to a precision of 10% of the resonator FWHM bandwidth. The on-sky resonator quality factors of 60k to 110k combined with resonance frequencies of 3–4 GHz imply a resonator FWHM of approximately 100 kHz; therefore, a precision of 10 kHz is required.

To obtain consistent phase information for each readout tone, the DAC frequency-bin width must be equal to or an integer multiple of the FFT channelizer resolution. In our case, a bandwidth of approximately 500 MHz with a 2^{16}-bin FFT yields a resolution of approximately 7.5 kHz. Because we will later decimate the data to the CSO-telescope pointing update rate of 100 Hz, the output data rate and thus the frequency resolution of the FFT should be a multiple of 100 Hz; hence, we choose a 7500-Hz FFT resolution rather than $500/2^{16}$ MHz. A total of 2^{16} FFT bins yield a 491.52-MHz bandwidth.

The most convenient place to store the LUT is in the FPGA fabric. However, there is limited memory space on the FPGA, and all of that space is required for the FFT. Instead, we successfully stored the LUT in the DRAM while satisfying the LUT synchronization performance requirements by carefully tuning the latency between the DRAM input and the output addresses.

2.4.1.2 Roll-Off Pattern Compensation

The full transfer function of the readout chain, including the effects of the DAC intrinsic sinc function, LPFs, IQ mixers, transformers, impedance mismatches, standing waves, cable delay, etc., must be considered when setting the optimal power level for each resonator. We used the network analyzer mode (described in Section 2.4.6) to record the roll-off pattern of each frequency bin for the full sampling bandwidth in the current system setup. We compensated for the roll-off by adjusting the DAC LUT and digital attenuator to optimize both the power level and the frequency across the entire readout bandwidth for each resonator in the MKID array.

2.4.1.3 Real-Time Reprogramming

During KID camera observation, the resonance frequency changes when the optical power on the resonator changes (e.g., when the telescope points toward a different part of the sky or when the atmospheric opacity changes). To compensate for these changes, we designed the firmware to allow for rapid updating of the DAC LUT using a new buffer with new drive tones while the channelization firmware on the FPGA continues to run (it takes less than one second for a new set of LUT buffers to upload).

2.4.1.4 Anticlipping Algorithm

The MKID LUT waveform is the summation of approximately 200 different carrier frequencies of similar but different powers. In an actual DAC, this summation poses a clipping problem, namely, the sum of the carrier values at a certain waveform position might exceed the maximum range of the DAC, even when the average power is within the DAC limits and when the carrier phases are randomly defined. Simply scaling the waveform amplitude to fit within the DAC limits will result in a reduction in the carrier and average power by the square of the scaling factor. This problem is often described in terms of the peak-to-average-power ratio (PAR). The SNR of the DAC is not effectively utilized when the PAR is significantly greater than unity.

Under the assumption of a sinusoidal wave output from a DAC with N bits, the theoretical maximum SNR can be calculated using the following quantization noise formula: $\mathrm{SNR} = (1.761 + 6.02 \times N)$ dB. From experience, if we simply use random phases for the 200 carrier tones with an LUT length of 2^{16}, we will obtain an SNR that is 1–2 dB below the theoretical maximum SNR defined by the effective number of bits. However, if we use a random phase, the summation of 200 carrier tones in the time domain may overflow the maximum range that can be defined using the number of bits available. Using MATLAB's random function to generate the phases, we must conduct many trials to ensure that no clipping occurs (occasionally as many as 100 trials). The indeterminacy of such an algorithm is undesirable for a scientific instrument, and this algorithm does not scale well with the number of carriers.

Another possible solution is to simply "clip" the waveform at the DAC limits, regardless of the consequences. Clipping the waveform does not strongly affect the average power, but it scatters power from the carrier frequencies into other frequency bins. This effect can be understood by equating the clipping of the waveform to the addition of a scaled delta function to the time sample at which the clip occurs and from the fact that such a pulse has a white frequency spectrum. The PAR in this scheme is better than that in the scaling scheme, but the power that remains in any given carrier will differ from what was originally designed.

This discussion suggests an alternative approach: at time samples at which the waveform would clip, subtract off a pulse that is designed to reduce the waveform amplitude below the DAC range but whose Fourier spectrum has had the carrier-frequency bins removed. In this manner, a signal (or noise) is introduced into the noncarrier bins in a manner similar to the introduction of a higher white-noise level, but this noise does not affect the SNR for the carrier bins. The applicability of this algorithm ultimately requires a high ratio of noncarrier bins to carrier bins. This clipping avoidance program can generate waveforms with a PAR approaching unity, making essentially full use of the dynamic range of the DAC.

2.4.2 ADC Digitization

A clock rate of approximately 500 MSPS is too fast, even for a state-of-the-art FGPA. To process the data from the ADC and DAC in real time, we first deserialize the data on the FPGA to reduce the FPGA clock rate by a factor of four (e.g., the clock rate on the FPGA becomes 491.52/4 MHz) at the expense of requiring the FPGA to simultaneously process four parallel data streams [24].

Furthermore, FFT algorithms have been designed to handle this parallelization of the data streams, as will be described in the next section. If we consider the 2^{16} nominal output bins of a nonparallelized FFT and number them sequentially as 1, 2, 3, etc., the parallelized FFT algorithm will then output four parallel streams with the frequency bins numbered 1, $2^{14} + 1$, $2 \times 2^{14} + 1$, $3 \times 2^{14} + 1$ output on one clock cycle; bins 2, $2^{14} + 2$, $2 \times 2^{14} + 2$, $3 \times 2^{14} + 2$ output on the next clock cycle; and so on.

Table 2.1: Deserialization of the data stream into four parallel paths for the FFT in the FPGA

FPGA Clock Cycle	Cycle 1	Cycle 2	...	Cycle 16383	Cycle 16384
FFT O/P 1	Bin1	Bin2	...	Bin $2^{14}-1$	Bin 2^{14}
FFT O/P 2	Bin $2^{14}+1$	Bin $2^{14}+2$	...	Bin $2\times2^{14}-1$	Bin 2×2^{14}
FFT O/P 3	Bin $2\times2^{14}+1$	Bin $2\times2^{14}+2$	...	Bin $3\times2^{14}-1$	Bin 3×2^{14}
FFT O/P 4	Bin $3\times2^{14}+1$	Bin $3\times2^{14}+2$	...	Bin $4\times2^{14}-1$	Bin 4×2^{14}

As the requirements for the bandwidth and ADC sampling rate increase, the data streams and FFTs can be further parallelized if necessary at the expense of requiring larger or multiple FPGAs.

2.4.3 Design of the Frequency Demultiplexer

There are several methods for implementing a demultiplexer in an FPGA: a digital down-converter (DDC), a single-stage FFT, a polyphase filter bank (PFB), an FFT zoom (two cascaded FFTs), a combination of a PFB and an FFT, a combination of an FFT and a DDC, and a combination of coaddition and an FFT. These options were all considered and compared for the design of the demultiplexer. For example, a DDC would not be an appropriate choice for an increasing number of resonators because it requires a long storage waveform (compared with the length of a single-stage FFT) to play back for the readout of each carrier tone. An FFT zoom program will suffer significant power leakage in the first FFT stage, particularly when the frequency does not fall at the centers of the bins of the first FFT.

In 2010, we implemented a combination of a 2^{11}-sample polyphase filter bank and a 2^{6}-sample FFT for our first DemoCam run. We obtained better performance using the final design that was implemented for the 2012 commissioning runs, which used a single-stage 2^{16}-sample FFT.

2.4.3.1 Design of the First-Version Firmware: PFB and FFT

The PFB was developed to account for the fact that the discrete Fourier transform (DFT) is an imperfect approximation of a set of BPFs. Ideally, each sample in a DFT would represent the response of a perfectly rectangular BPF centered on an integral multiple of F_n/N (the Nyquist frequency divided by the size of the DFT) with a width of exactly F_n/N. In actuality, each sample in a DFT has a sinc frequency response. The PFB squares off this filter shape by first convolving the input with a sinc function over a long time window and then applying a DFT. The length of this time window determines the steepness of the edges on the filter; typically, it may be four to eight times (four to eight taps of the filter) the length of the DFT.

In the combined PFB and FFT design used in the 2010 DemoCam run, we used a four-tap 2^{11}-point Hamming-window complex polyphase filter bank followed by a transpose function implemented on the QDR and then a 2^{6}-point complex FFT. This design is similar to an FFT zoom, but the use of the PFB substantially reduces the bin-to-bin power leakage present in the first FFT of an FFT zoom. A 2^{11}-point PFB together with a 2^{16}-point FFT yields 2^{17} frequency bins.

Compared with a direct 2^{17}-point FFT, the combined PFB and FFT solution required considerably less space on the FPGA; a 2^{17}-point FFT would not have fit. This design would have allowed us to expand the size of the channelizer up to 16 million points on a single ROACH (or even larger by combining ROACHes together through a 10-Gbit connector). The cascaded design shifts the FPGA fabric requirement for the FFT into the external QDR memory by dividing the large FFT into two parts.

The combined PFB and FFT design suffers intrinsic power leakage between bins in the first stage unless an infinitely long filter bank (impossible in practice) is implemented in the first stage. In a real FPGA implementation, the first-stage PFB is very expensive in terms of resources because it must sum and buffer the full raw timestream X times, where X is the number of taps in the PFB. The value of X used is a tradeoff between the amount of FPGA fabric used and the quality of the response function.

In the design that we used in 2010, we implemented a four-tap polyphase filter bank in the first stage (buffering the raw data stream four times). Because the PFB and FFT are both well-defined linear calculations, the bin-to-bin leakage is known and can be corrected for. However, doing so increases the amount of FPGA fabric required because additional frequency bins must now be carried along for every carrier bin. This method has been implemented and validated for a situation in which the number of carriers is not large. However, as the number of carriers increases, even the summation of only four additional bins for each carrier bin incurs considerable computational effort. In principle, this problem can be solved through matrix inversion, but this is also computationally challenging. One method to simplify the summation is to use a large first-stage PFB and a small second-stage FFT. However, this approach causes the two-stage design to become more similar to a large single-stage FFT, which conflicts with the original intent of conserving computing resources by using the two-stage design. One effect that must be accounted for in the PFB–FFT design is that the PFB induces a phase rotation of the Fourier coefficients. This rotation is deterministic and can be corrected for.

For radio astronomy applications other than KID readout that do not involve a large number of predefined carriers tones, the aforementioned power leakage problems are less important. For KID readout, however, they are an issue of considerable concern, and the difficulties will only increase as the number of tones per unit bandwidth increase. Ultimately, we chose to focus on implementing a large single-stage FFT.

2.4.3.2 Design of a Large Complex FFT on a Xilinx FGPA

We evaluated FFT designs from both the CASPER group and Xilinx. The CASPER FFT uses our four parallel timestreams exiting the ADC [termed "Decimation In Time" (DIT)]. The Xilinx FFT is more logic-cell efficient than the CASPER FFT for a large FFT size (greater than 2^{15}) and a large output bit width (greater than 18 bits). We considered the advantages of both designs and ultimately developed an improved FFT design by recognizing the following deficiencies of the CASPER FFT block:

1. For digital processing, the number of bits required to store the number inside the FPGA for each stage must be determined to ensure that no information is lost during digital signal processing. We can address this issue from an SNR perspective. We know that the maximum possible SNR that can be represented by N bits is $(1.761 + 6.02 \times N)$ dB. For example, if we use a 2^{16}-sample FFT and a 64-dB SNR at the ADC, the SNR at the FFT output will then be 112 dB. Thus, the output must be at least 18.39 bits. Here, we use 20 bits for the FFT output and 19 bits for the calculation coefficients within the FFT process. A 20- or 19-bit output and coefficient calculation does not mean that no number can be greater than 20 bits during the FFT calculation. The intermediate numbers inside the FFT calculation are floating points and can thus be scaled on the basis of the calculation requirements.

   For simplicity, the CASPER FFT assumes that the output bit width is identical to the input bit width. However, the ADC does not have a significant number of bits (e.g., our ADC has 12 bits), and this method wastes FPGA resources if we use 19 bits for both the input and output of the FFT.

2. The CASPER FFT is constructed around biplex blocks that process two parallel input timestreams. For four parallel streams, it uses two biplex blocks. The coefficients used by all of these biplex blocks are identical, but each block possesses its own copy of the coefficients by default. By sharing a single copy of the coefficients, substantial FPGA fabric can be saved.

3. Some of the coefficients for the FFT calculation are stored in read-only memory (ROM) as an LUT for the CASPER FFT. However, those numbers can be generated on the fly using FPGA logic, thereby conserving a significant amount of resources.

To overcome the aforementioned deficiencies, we implemented an improved FFT block to solve the problems and to enable a single 2^{16}-bin single-stage FFT design on a Vertex 5 FPGA. Table 2.2 summarizes the FPGA logic-cell utilization of the 2^{16}-bin single-stage FFT alone (without other firmware functions running on the FPGA)[3].

---

[3] The default number of flip-flops on the FPGA is limited, but it can be replaced with slices when additional flip-flops are required.

Table 2.2: FPGA logic-cell utilization summary

Logic Utilization	Used	Available	Utilization Rate
Number of slice registers	7348	58880	12%
Number of slice LUTs	5928	58880	10%
Number of slices as memory	2881	24320	11%
Number of route-thrus	329	117760	1%
Number of flip-flop pairs used	7575	7575	100%
Number of bonded IO buffers	448	640	70%
Number of RAM/FIFO blocks	20	244	8%
Number of DSP48Es	62	640	9%

2.4.4 Bin Selection and Decimation

2.4.4.1 Bin-Selection Design

Because we are only interested in the information contained in the approximately 200 bins that carry the resonator tones, we wish to extract these 200 bins from the full FFT. When the DAC LUT is generated, we also generate a position table for the FFT bins that carry resonator tones, which is automatically updated as the DAC LUT is updated while the FPGA channelization program is running.

As previously discussed in the context of the ADC output streams (Section 2.4.2), the signal timestream is deserialized into four parallel streams, and the FFT on the FPGA also has four parallel outputs (as summarized in Table 2.1).

We select bins by masking all output cycles with an LUT of $2^{16}/4$ in size. For each FPGA clock cycle, we determine whether there were any resonators within this clock cycle and, if so, which of the four parallel outputs are contained the resonator bin.

Because there are four parallel simultaneous outputs for each clock cycle, we cannot select bins from two or more of the four parallel FFT outputs at once without buffering the four parallel output data streams. To bypass this difficulty, we ensure that no more than one FFT bin is selected per clock cycle. This scheme requires that two resonator tones not be separated by an exact multiple of 2^{14} bins (e.g., bin 1 and bin $2^{14}+1$ cannot be selected at the same time).

2.4.4.2 Decimation for Astronomical Signals

After selection of the resonator bins, approximately 200 data streams carrying resonator information are further processed and stored. The data rate after the FFT is 7500 Hz (the sampling rate divided by the FFT length). To reduce the data rate, rather than using simple coaddition or averaging, each selected resonator data stream is processed with a 150-tap Hamming-window FIR filter. The output timestream is then decimated by a factor of 75 to 100 Hz to match the update rate of the CSO-telescope pointing timestream. The use of an FIR filter prior to decimation provides better frequency–space characteristics than does coaddition or averaging. The signal flow for the decimation process is presented in Table 2.3. The selected carrier tones are assigned to FIR filters 1, 2 and 3 in a cyclical manner.

Table 2.3: Signal flow of the 192 resonators feeding into three FIR filters

-	Channel 1	Channel 2	...	Channel 63	Channel 64
I/P to FIR 1	*Res*_1	*Res*_4	...	*Res*_187	*Res*_190
I/P to FIR 2	*Res*_2	*Res*_5	...	*Res*_188	*Res*_191
I/P to FIR 3	*Res*_3	*Res*_6	...	*Res*_189	*Res*_192

As previously mentioned, there are 144 resonators within one 6 × 6 antenna array. Hence, there are 144 on-resonator frequency bins that contain resonator information. The extra readout capacity enables us to generate 50–60 tones in off-resonance frequency bins and fully read out those off-resonance carrier bins. We use these data to measure the common-mode electronics 1/f noise (HEMT and room-temperature electronics; this noise is discussed in detail in Section 2.8.3.).

We eventually chose 48 off-resonator carriers because each FIR filter block accepts 64 streams of data; thus, the three FIR blocks that are required to handle 144 resonators can process an additional 48 off-resonance streams without requiring additional FPGA resources.

For each channel, the 150 coefficients are stored using 22 bits because the SNR that is required at the FIR output for each channel is 130.9 dB[4]. At least 21.46 bits are needed for the coefficients to provide this output SNR.

Here, we used the Xilinx FIR block. In the future, when a larger number of resonators must be read out and when the required number of FIR blocks cannot fit on the FPGA, the external memory (e.g., QDR or DRAM) can be used to store the FIR coefficients and buffer the data streams.

2.4.5 Synchronization

2.4.5.1 Timestamps and Data Packaging

The 1-PPS TTL signals are received on the IF board and routed through the ADC/DAC board and ROACH board to the FPGA, where they provide a logic signal with a rising edge on the exact-second boundary with respect to the absolute time of day. Both the DAC and channelizer begin at the same

---

[4] As discussed in Section 2.4.3.2, the SNR requirement at the FFT output is 112 dB; after the FIR, with a decimation rate of 75, the SNR requirement becomes $112 + 10 \times \log_{10}(75)$ dB.

second boundary to ensure that we obtain a consistent phase for all carrier bins. For synchronization with the absolute time of day, a C program running on the PPC transfers the current Linux time from the PPC [which is locked to a network time protocol (NTP) server] to the FPGA. A counter in the FPGA that is locked to the 1-PPS signal begins counting in integer seconds from that time. Another counter on the FPGA serves to count the number of FPGA clock cycles since the last second boundary. Thus, the two counters combined provide timestamp information accurate to 1/(FPGA clock rate) $\approx$ 10 ns.

Each output data packet contains a timestamp (both seconds and fractional seconds), a header (a predefined number to separate the timestamp and data), and 192 complex resonator signal values (as summarized in Table 2.4). The data are transmitted via 1-Gbit Ethernet at 100 Hz to a DAQ computer using a TCP data server that we constructed for this purpose. The first-version firmware used 10-Gbit user datagram protocol (UDP) Ethernet, but this Ethernet card is expensive, and 16-port 10-Gbit Ethernet switches are not readily available. In addition, the UDP packet, although simple to set up, does not have handshake functions to protect against data loss. Thus, we switched to 1-Gbit Ethernet with TCP packets. After optimization of the connection parameters, we observed no lost packets when using TCP at the required data rate.

Table 2.4: Output packet data format for MKID readout

Position	1	2	3	...	194
top 32 bits	Unix seconds	header	res1 real part	...	res192 real part
bottom 32 bits	fractional second	header	res1 imaginary part	...	res192 imaginary part

2.4.6 Network Analyzer Mode and IQ Sweep Mode

In addition to the mode described above, which is used for observations, we designed several other modes of firmware operation. The following are two commonly used firmware operation modes:

- IQ sweep mode. This mode is a variant of the normal DAQ mode. Normally, the LO is held fixed while data are acquired during an observation. Instead, we acquire short blocks of data (on the order of 0.5–1 s) in this mode to characterize the system's complex transmission as a function of the frequency near each resonance, and between blocks, we step the LO away from its nominal value over a range that is sufficiently large to characterize each resonator's complex resonance circle. On the DAQ personal computer (PC), each block of data is averaged in time to yield a single mean I and Q data point for each carrier tone. As the LO is stepped, the IQ points trace a trajectory in the complex plane for each resonance. From the IQ-plane trajectory, we can determine the optimal power and frequency at which to drive each individual resonator.

- Network analyzer mode. The DAC plays back a white noise signal or any other signal that has equal power in all frequency bins. Because we are interested in the information in all FFT bins in this case and cannot send that information in its entirety to the DAQ PC at the 7.5-kHz FFT cycle rate, the time-domain data are coadded prior to the application of an FFT, and the FFT output is then sent to the DAQ PC. The coaddition preserves the signal but averages down the noise; this is because the phases of the transmitted and received signals do not change from one LUT playback cycle to another, but the phase of the noise does vary and is therefore integrated down. The network analyzer mode allows us to perform the following tasks:

    1. Rapid identification of approximate resonator locations.
    2. Calculation of the cable delay of the readout setup through a comparison of the phase of each frequency bin.
    3. Examination of the current ADC dynamic range to adjust the digital attenuation on the IF board.
    4. Verification of system connections.

A network-analyzer frequency scan can also be implemented using the IQ sweep mode; in this case, the full 491.52-MHz bandwidth is separated into 192 subbandwidths of equal length. We use 192 drive tones to simultaneously scan these subbandwidths and then combine the data to obtain the full 491.52-MHz bandwidth.

Using the OSR network analyzer mode, we can completely avoid the use of other network analyzers, positioning the OSR as a complete solution for detector readout.

2.5 Noise Study of the Readout System

2.5.1 Resonator Output Power and HEMT Noise-to-Carrier Ratio

In the previous section, we discussed the noise requirements of the OSR system. In combination with the resonator power level, we can calculate the required noise-to-carrier ratio for the OSR system. Each resonator has a power of 10 to 30 pW (−80 to −75 dBm) at the device.

Therefore, the noise-to-carrier ratio per hertz at the HEMT will be

$$\left(\frac{N}{C}\right)_f = \frac{k_B \times T_N}{\text{Readout Power}} = \frac{1.38 \times 10^{-23} \times 2}{30 \times 10^{-12}} = 9.2 \times 10^{-13}/\text{Hz} = -120 \text{ dB/Hz} \tag{2.1}$$

Before the HEMT, the 144 resonators will each have a power of −58.4 to −53.4 dBm. Inside the cryostat, the HEMT (gain of 35 dB; noise temperature of 2 K) and coaxial cable (loss of 3 dB) result in a timestream noise-to-carrier ratio at the HEMT of −56.7 to −51.7 dB at a noise bandwidth of 245.76 MHz.

2.5.2 Noise-to-Carrier Ratio of the Detector

The noise study is an important part of the OSR. A complete study of the noise is presented in Appendices A and B, but in this section, an important aspect of the noise study related to the OSR system will be discussed: how the $\left(\frac{N}{C}\right)$ ratio specifications should be translated for the interpretation of measurements.

As indicated in the previous section, the $\left(\frac{N}{C}\right)_f$ ratio leaving the HEMT is −120.63 dB/Hz. We refer to this quantity as the frequency–space $\left(\frac{N}{C}\right)$ ratio: $\left(\frac{N}{C}\right)_f$ in units of decibels per hertz. We provide a short calculation here to explain the meaning of this value in terms of a power-spectrum measurement.

2.5.2.1 Single Detector

A value of $\left(\frac{N}{C}\right)_f = 1 \text{ x } 10^{-12}$/Hz implies that the timestream noise-to-carrier ratio will be

$$\left(\frac{N}{C}\right)_t = \left(\frac{N}{C}\right)_f \times \text{Nyquist Bandwidth} \tag{2.2}$$

where the Nyquist bandwidth is 245.76 MHz in our case[5].

Now, if we perform a measurement using a spectrum analyzer, we cannot directly measure $\left(\frac{N}{C}\right)_f$. This is because the spectrum analyzer treats the carrier and noise in the same manner, either providing the Fourier coefficients of the power or the Fourier spectral density of the power for both, as shown in the following equation:

$$\left(\frac{N}{C}\right)_{PSD} = \frac{PSD_{noise}(f)}{PSD_{signal}(f)} = \frac{PSD_{noise}(f)}{\text{Signal Power}(f)/\Delta f} = \left(\frac{N}{C}\right)_f \times \Delta f \tag{2.3}$$

where Δf is the FFT bin width

$$\Delta f = \frac{\text{Nyquist Bandwidth}}{\text{Number of PSD Bins}} \tag{2.4}$$

and we have

$$\text{Noise Power} = PSD_{noise}(f) \times \text{Nyquist Bandwidth} \tag{2.5}$$

---

[5] Equation 2.2 only holds for N/C in linear units and not in decibels.

combine the Equation 2.3 and 2.5, we have

$$\left(\frac{N}{C}\right)_f = \frac{PSD_{noise}(f)}{\text{Signal Power}(f)} = \frac{\text{Noise Power}}{\text{Signal Power} \times \text{Nyquist Bandwidth}} \tag{2.6}$$

and combine the Equation 2.4 and 2.6, we have

$$\left(\frac{N}{C}\right)_{PSD} = \left(\frac{N}{C}\right)_f \times \Delta f = \frac{\text{Noise Power}}{\text{Signal Power} \times \text{Number of PSD Bins}} \tag{2.7}$$

Here, PSD_{noise} and PSD_{signal} are the power spectrum densities of the noise and signal, respectively; the former is determined by the HEMT noise temperature, and the latter is determined by the KID readout power and the frequency bin width. Therefore, the measurement we acquire using a spectrum analyzer depends on the frequency-bin width Δf that is used. In our case, $\Delta f =$ 7.5 kHz, which corresponds to a 491.52-MHz sampling rate and a 2^{16}-point FFT. If we use a wider frequency-bin width, we will integrate the HEMT noise over a larger bandwidth without changing the signal power and will thus obtain a larger value of $\left(\frac{N}{C}\right)_{PSD}$ for the same $\left(\frac{N}{C}\right)_f$.

2.5.2.2 Multiple Detectors

For multiple tone readouts, let N_c be the number of detectors. The signal power increases linearly with N_c, but the number of detectors does not affect the HEMT noise power. Therefore,

$$\left(\frac{N}{C}\right)_{fN_c} = \frac{\left(\frac{N}{C}\right)_f}{N_c} \tag{2.8}$$

The noise-to-carrier ratios $\left(\frac{N}{C}\right)$ were used in noise studies of both digital and analog systems for the OSR. Three different units are used throughout this dissertation. In Table 2.5 below, the definitions of $\left(\frac{N}{C}\right)$ are summarized in these different units, namely, time-domain, frequency-domain, and spectrum-analyzer units, and the conversions among the different units are presented.

Table 2.5: Interpretation of various noise-to-carrier ratios

Interpretation of Various Noise-to-Carrier Ratios		
Time-Domain $\frac{N}{C}$	$\left(\frac{N}{C}\right)_t = \frac{\text{Noise Power}}{\text{Signal Power}}$	dB
Frequency-Domain $\frac{N}{C}$	$\left(\frac{N}{C}\right)_f = \frac{\text{Noise Power}}{\text{Signal Power} \times \text{Nyquist Bandwidth}}$	dB/Hz
Spectrum-Analyzer (PSD) $\frac{N}{C}$	$\left(\frac{N}{C}\right)_{PSD} = \frac{\text{Noise Power}}{\text{Signal Power} \times \text{Number of PSD Bins}}$	dB
Conversion Among Ratios	$\left(\frac{N}{C}\right)_f \times \text{Nyquist Bandwidth} = \left(\frac{N}{C}\right)_t$ $\left(\frac{N}{C}\right)_{PSD} = \left(\frac{N}{C}\right)_f \times \Delta f$	Δf $= \frac{\text{Nyquist Bandwidth}}{\text{Number of PSD Bins}}$
Noise Power $= PSD_{noise}(f) \times$ Nyquist Bandwidth		

2.6 Data Acquisition System and Telescope Operation

Heretofore, we have discussed both hardware and firmware. In addition to these components, we also developed a software package to fully automate the process of multiboard readout, hardware control, communication with the telescope, and computer-controlled data acquisition. We call this software package the DAQ system, and a block diagram of this system is presented in Figure 2.10.

The DAQ functions are listed in the block diagram surrounding the hardware and firmware block in the figure.

The preparatory functions for use before the start of observation are on the left-hand side, which include the following:

1. Initialization of the FPGA firmware. This commands the FPGA to stand by and be ready to process a signal. The required IF board control function and LUT programming are run on the FPGA.
2. IF board setup. The DAQ program sets the parameters for the IF boards, including digital attenuations, digital switches, the FPGA clock, and the LO frequency.
3. LO sweep and generation of resonator frequencies. On the basis of current sky and observation conditions, the IQ sweep mode (Section 2.4.6) is used to determine new drive-tone frequencies.
4. LUT and bin selection. On the basis of the newly generated drive-tone frequencies, the LUT and bin-selection buffer on the FPGA are reprogrammed. A second IQ sweep is performed using the new carrier frequencies, which allows us to analyze the resonator properties during the noise-removal and analysis processes[6].
5. Timestamp generation. An accurate timestamp is generated (Section 2.4.5.1) and programmed into the FPGA such that every data packet output by the FPGA contains the correct absolute time-of-day information.
6. Readout system-alive diagnostic. To ensure that the readout system is ready to observe, we run a final system-alive diagnostic to verify that the data output by the FPGA are reasonable by checking that we receive the correct number of carrier tones, that the carrier tones are in the correct order, that the timestamp is correct compared to the clock, that the power level is as expected, and that the on and off resonator carriers behave as expected.

After all of the functions on the left-hand side have been completed, we can begin observing. Thee follow-up functions to be executed once we have raw data from the FPGA output are on the right-hand side and are as follows:

---

[6] The second IQ sweep is necessary because the transfer function of the system is a function of both the LO and baseband frequencies and not just of the RF frequency; therefore, the transfer function must be remeasured at the same LO and baseband frequencies that will be used during data acquisition.

1. Acquisition of CSO pointing information. The CSO generates 100-Hz pointing information, which is concurrently and separately recorded.

2. Interpretation of the output raw data packets. The output raw data must be converted into an easy-to-read format including separating the on and off resonators, linking the data to the resonators, and combining the timestamps with the telescope pointing information.

3. Noise-removal and data-analysis pipelines. Noise removal is an important part of the system because the astronomical signal contains inherent noise (e.g., sky noise), and the readout electronics introduce additional noise into the raw data. This noise must be removed to obtain the desired SNR, as is discussed in detail in Appendices A and B.

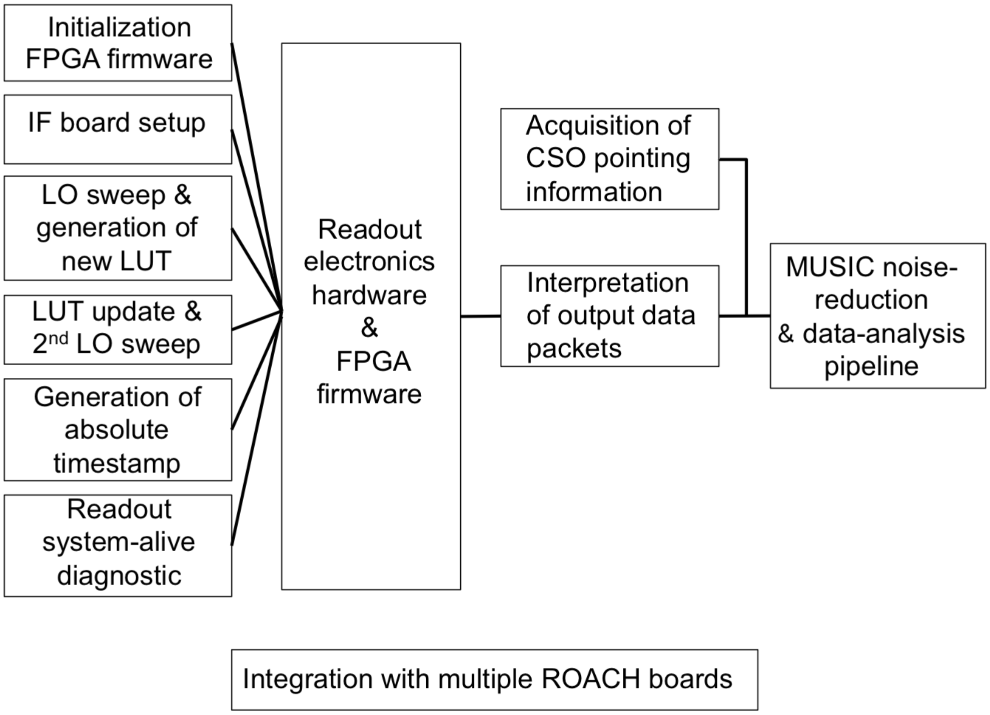

Figure 2.10: Block diagram of the DAQ functions and processes.

2.7 Summary of the Specifications of the Open-Source MKID Readout System

The final OSR specifications that we used for the CSO telescope are summarized in Table 2.6 below.

Table 2.6: Summary of the specifications of the open-source MKID readout system

Noise-to-carrier ratio	Less than -123 dB/Hz
Number of tones	192
Bandwidth	491 MHz
DAC frequency step size	7500 Hz
DAC waveform buffer	Continuous; Fast reprogramming
Channelizer resolution	7500 Hz
Channelizer size	65536 bins
Output data format	Complex output with 2×32 bits
Channelizing speed	Real time
Output data rate	100 Hz/Channel
Output protocol	1-Gbit or 10-Gbit Ethernet; 16 OSR unit switch
Timestamp	Absolute time of day (precision of up to 1×10^{-8} s)
OSR synchronization	All electronics components are locked together
Telescope synchronization	Telescope, electronics and DAQ PC are all locked together
Operating mode	Many modes available, making the OSR a complete KID readout solution
Telescope operation	Fully developed and tested
Data acquisition	Fully developed and tested

2.8 Discussion of Special Topics

In addition to the hardware, firmware, and software, we also made several interesting discoveries regarding the OSR system during the development process.

2.8.1 Multiple-Tone Readout Considerations

Single-tone readout is highly similar to a commercial network analyzer. When ports 1 and 2 are connected to a device, a network analyzer can be used to sweep all frequency bins at a predefined bandwidth and resolution (the S_{21} parameter). Asking the network analyzer to examine only one bin rather than all bins in the bandwidth is similar to reading a single resonator; the multitone readout that we developed is analogous to having 192 network analyzers simultaneously examine 192 bins.

The readout of many tones from the same transmission line is an extension of the signal-tone readout with some additional criteria:

1. SNR limitations. If all tones are read using the same DAC and ADC chip, we must ensure that both the individual tone and overall power satisfy the SNR requirements. After the cryogenic amplifier, the ADC is typically the next most important limiting factor for the SNR of the entire readout system.

2. Better adjustment for each tone is required. When we increase the number of tones and begin to utilize all available bandwidth, we must consider the roll-off and ripples across the bandwidth, particularly when we want each drive tone to be optimized at the device, with no loss in power level after propagation along the transmission chain. Thus, there is a need to leave a sufficient margin beyond the theoretical values, which are based on a PAR equal to one.

3. Increased occurrence of spurious tones. The ADC and DAC commonly have an SFDR and IMD that indicate the highest harmonic or spur that the chip will generate. In addition, certain components such as mixers will also introduce spurs. Some spur frequencies are well known, such as the IMD and harmonics. However, other spurs may be randomly distributed across the bandwidth and defined by the SFDR of the chip. To minimize the effects of random spurs, we selected ADC and DAC chips with good SFDRs. For the third-harmonic spurs, because their frequencies are well defined, we can design the carrier frequency bins and program the LUT to ensure that the tones and their harmonics will not overlap.

4. An increased power level may affect the performance of electronic components. For amplifiers and mixers, we must avoid driving the components into their compression ranges. Thus, we should maintain operating limits far away from the 1-dB compression point. For MUSIC, we chose components such that the power levels would be at least 7 dB away from the compression range. For applications such as MUSIC that require longer-timescale stability with the need to probe signals less than 1 Hz, driving the components excessively hard may introduce additional 1/f noise, as discussed in Section 2.8.3.

In summary, an OSR for KIDs is analogous to connecting multiple network analyzers on a single transmission line, except that the OSR must read at a predefined time and frequency resolution. Each readout unit (including one ADC/DAC board, one IF board, and one FPGA board) can read out a limited number of resonators (192 carrier tones in the case of MUSIC). We must use a combination of many readout units because there are several thousand resonators in the final telescope instruments. For MUSIC, we integrated eight readout units into a single crate (Figure 2.11). This electronic crate not only houses all of the individual readout units together but also provides power through a combined power supply, which significantly facilitates the setup of the readout electronics.

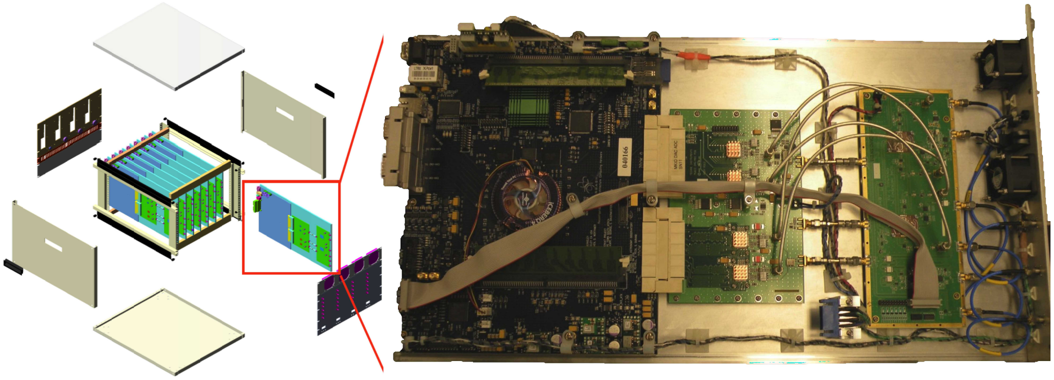

Figure 2.11: Electronics crate for eight readout units.

2.8.2 Discussion of Different Types of DSP Hardware: GPUs and FPGAs

There are several options for the types of signal processing hardware that can be used for KID readout, including ROACH I, ROACH II, mini-ROACH, or custom-developed boards. There is a tradeoff among the FPGA speed, the readout bandwidth, the FFT length, the number of detectors to read out, and the cost per pixel. We may also need to consider the size of the KID array on the wafer.

For the MUSIC readout, we implemented all of the required functions using ROACH I. The Zdok connectors are common to ROACH I and II; therefore, the ADC/DAC and IF boards we developed are compatible with both. The yellow blocks[7] for the MKID ADC/DAC would require only a minor modification of the pinout to match a new FPGA on ROACH II. All firmware and software will also operate on ROACH II with an automatic update from the Xilinx compiling tools.

For KIDs operating at a few gigahertz (as is the case for the MUSIC system), we pack 192 detectors into the 491.52-MHz bandwidth and use a channelizer resolution of 10 kHz. In this case, we may not be able to fully utilize the computing ability of ROACH II, which has a higher-grade FPGA chip. ROACH I is more cost efficient than ROACH II and satisfies our requirements.

For KIDs operating below 500 MHz, 1000–3000 resonators could fit within this bandwidth. In this case, the readout system will require a larger FFT size (ranging from a few million to a few hundred million bins) on the FPGA. The requirements of a higher FFT size, better frequency resolution, or higher I/O rate indicate that ROACH II may be a more suitable choice.

Another interesting recent development is that of graphics processing unit (GPU) readout, which combines the functions of an FPGA and a GPU. A GPU is based on a computer operating system and works with a computer motherboard. It offers many advantages including lower cost, ease of installation and programming, and the ability to input and output directly using the computer.

---

[7] The input and output simulation blocks in MATLAB Simulink are known as yellow blocks.

GPUs are ideal for applications that require signal processing that is more complex than a normal FFT or correlator or nonlinear calculations; however, the speed and input–output (I/O) capability of a GPU are not as good as those of an FPGA. Hence, a combination of an FPGA frontend and a GPU backend is a popular solution. The major calculations that must be performed in a KID application lie in the linear FFT and FIR filter, and these calculations can be largely performed using an FPGA. We have not explored the GPU option further, but the use of a GPU backend would be a suitable choice if more post-FFT analysis becomes necessary in future large KID arrays.

There has been increasing interest in and effort directed toward the development of an inexpensive customized FPGA/GPU board for KID applications to reduce the cost per pixel. This topic is of interest because the number of KIDs on each telescope will increase in the future to utilize the frequency multiplexing capability of KIDs. ROACH boards are not customized for KID readout, as they contain unnecessary components and features that should be improved upon, and they still have a high cost. This may be a promising direction for future KID arrays.

2.8.3 System 1/f Noise

Noise in the camera system originates from the readout electronics, HEMT, sky, and MKID device itself. Some of these sources produce uncorrelated noise across different resonators, such as two-level system noise, whereas others produce common-mode noise such as electronics noise.

We carefully studied the noise performance of the readout electronics, particularly in the low-frequency range. In the 100-Hz audio stream output of the readout system, we observed an increase in noise in both the amplitude and phase directions for frequencies below 10 Hz. Some components of the readout electronics, including the voltage regulator, exhibit such 1/f noise behavior. Clock jitter from the signal synthesizer or voltage-controlled oscillator (VCO) and aperture jitter from the ADC will also appear as low-frequency phase noise.

To minimize the 1/f noise, we used the frequency-standard locked low-phase-noise option of the VCO on the IF board and used the same VCO as the LO for both up- and down-mixing. We designed a second-generation ADC/DAC board with a low-noise common voltage regulator, clean external DC supply, and better heat sink to reduce low-frequency noise.

In the OSR setup, we observed a high correlation among all 192 tones in both the amplitude and phase directions (with average correlation coefficients greater than 0.93), as shown in Figure 2.12. We can recover a signal that suffers from 1/f readout electronics noise by comparing the on and off resonator tones that were sent out and processed simultaneously, as they experience the same 1/f noise.

A plot of the noise power spectrum density (PSD) is presented in Figure 2.13. In this plot, the red line indicates the measured noise floor of the readout system, which is in agreement with the theoretical calculation of -135 dBFS, and the green line indicates the signal noise floor after the removal of 1/f noise, which can reach a theoretical noise floor above 1 Hz.

When the electronics 1/f noise is combined with the resonator's two-level system noise and the sky noise, the nature of the noise becomes too complex to solve through the simple subtraction of common-mode noise. A detailed and comprehensive study of noise removal performed by our group was integrated into the data-processing pipeline in the DAQ.

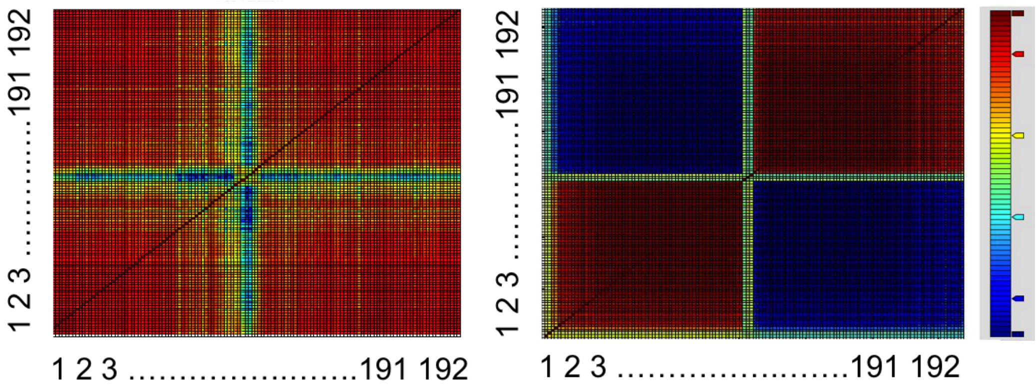

Figure 2.12: Noise correlation coefficients for the amplitude and phase components.

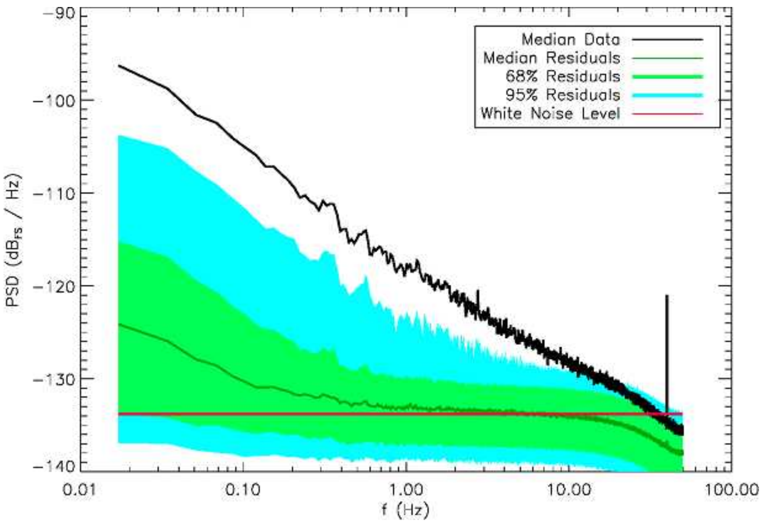

Figure 2.13: Noise power spectral density ($\sqrt{\mathrm{I}^2 + \mathrm{Q}^2}$) for the off-resonance tones in laboratory data acquired with a stationary dewar. The thick black line represents the raw data, the thin green line represents the median post-removal data, and the green and cyan bands represent the 68% and 95% ranges for the post-removal data. The horizontal red line indicates the inferred white-noise level, which should be dominated by the HEMT white noise. The expectation for the true HEMT noise level is somewhat uncertain because of the uncertainty in the system gain from the HEMT to the ADC, but it is in the range of −140 to −135 dBFS/Hz, consistent with the measurements. Plot courtesy of Seth Siegel.

2.8.4 Spurious Frequency Sources

Spurious frequencies are detrimental to most KID readouts. The possible causes of spurious frequencies include the following:

1. Nonlinearity of the IQ mixer. The DAC commonly exhibits good IMD performance; thus, the problem does not originate solely from the DAC output. However, the typical IQ mixer that we use has an input intercept point (IIP) of 10–16 dBm, which means that if we send out a carrier with an amplitude of -10 dBm, we should obtain an IMD spur that is approximately 40 dB lower than the signal at the IQ mixer output. These spurs will appear at the intermodulation products of the input frequencies. The density of intermodulation products is sufficiently high that they may appear to be white noise.

2. DAC quantization. When using MATLAB or Python for sinusoidal calculations, we round the number to the nearest integer, which is a typical source of quantization noise. The theoretical SNR can be calculated to be $6.05N + 1.76$, where N is 16 for a 16-bit DAC. This calculation yields an extremely low quantization noise, but a periodic pattern will still be introduced into the LUT timestream, and there will be spurs in the frequency domain that are greater than the theoretical noise level. We have checked this result using a MATLAB simulation and found that this type of spur will be lower than the ADC SNR and will therefore not contribute noise to the system. Although we are still limited by ADC noise, we should be able to visualize this type of spur in the DAC output using the spectrum analyzer.

3. IQ imbalance. When we check the full frequency spectrum of ADC digitized I and Q timestreams (without bin selection and FIR), in addition to the bin with the expected carrier ($BIN_{carrier}$), sometimes a second peak is observed. We define the bin in which this second peak appears at BIN_{2nd_Peak}. We find this peak exhibits the following behavior:

   $$\mathrm{BIN}_{carrier} + \mathrm{BIN}_{2nd\_Peak} = \text{total FFT bin number} \tag{2.9}$$

   For example, for carrier bin 16382, the 2nd peak appears at bin 114690 for the 2^{17} FFT bins.

   This result can be explained by the slight difference in the digitized data between the ADC I and Q inputs. This IQ imbalance can originate from either the amplitude or phase mismatch between the I and Q signals. For the discussion presented below, take the amplitude imbalance as an example, and let Amp_I and Amp_Q be the amplitudes of the digitized data for the in-phase and quadrature-phase inputs, respectively. For the cases of $Amp_I = Amp_Q$, $Amp_I = 1.05 \times Amp_Q$, and $Amp_I = 1.2 \times Amp_Q$, we studied the effect of this IQ imbalance, as shown in Figure 2.14.

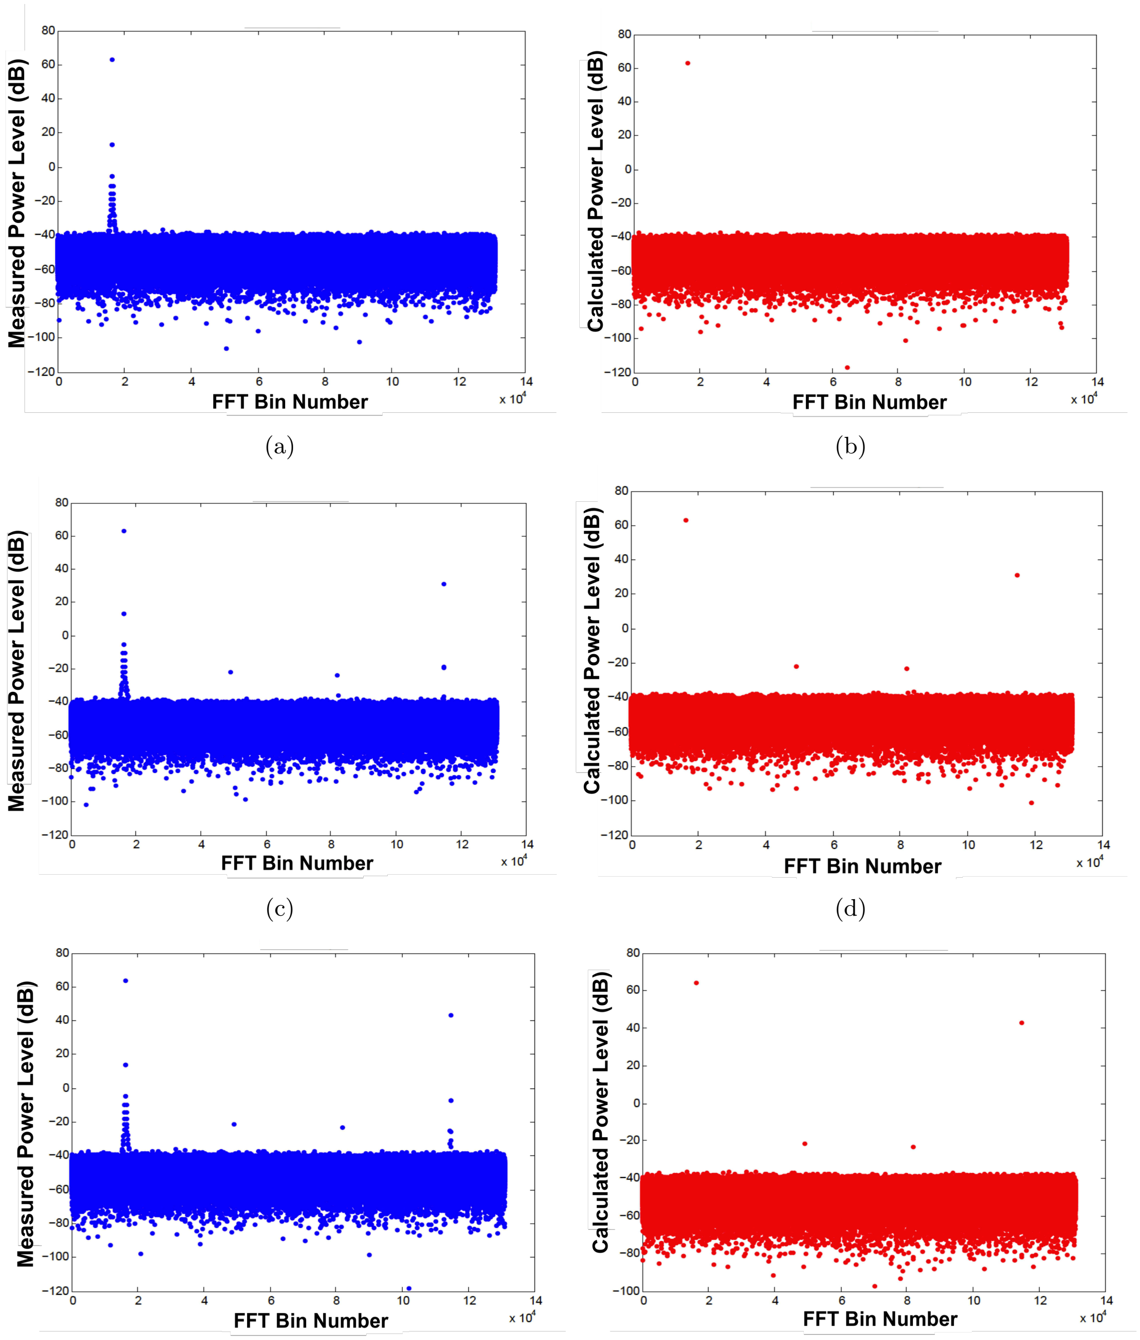

(a) (b)

(c) (d)

(e) (f)

Figure 2.14: Study of the power leakage caused by unbalanced I and Q power levels. The x axis represents the 2^{17} frequency bins; the y axis represents the power level relative to a full-scale of 100 dBm. (a) and (b) $Amp_I = Amp_Q$, (c) and (d) $Amp_I = 1.05 \times Amp_Q$, and (e) and (f) $Amp_I = 1.2 \times Amp_Q$. (a), (c), and (e) The blue plots represent FPGA complex FFT output of ADC digitized timestreams of DAC/ADC loopback: (FFT(Complex(Digitized_I, Digitized_Q))). We store the full FPGA complex FFT output without bin selection and FIR. The amplitude difference between I and Q is achieved by programming the DAC with different LUTs; (b), (c), and (f) the red plots represent the simulation obtained using MATLAB.

From Figure 2.14, we can conclude that the measured power leakage due to IQ imbalance agrees with the simulation. The spurious frequency caused by IQ imbalance is well defined in frequency, and can be avoided by carefully picking the carrier tones that do not fall on the leaked frequencies.

4. In addition to the issues mentioned above, other factors (e.g., the delay between the up and down IQ mixers, the IQ mixer DC offset) may also introduce spurs that should be considered in combination with the IIP level; however, these spurious frequencies are more random.

2.8.5 SNR Table

Figure 2.15 presents a table that summarizes the parameters that propagate through the IF board, including the detailed signal power level, noise level, IMD level, accumulated SNR, equivalent noise temperature, and component specifications (model number; gain, loss, and noise figures).

The table in Figure 2.15 was adapted from one provided by an engineer from Omnisys, a vendor that provided an earlier version of the readout system, prior to our OSR development. Under the assumption that the signal power, as the output of the cryogenic HEMT amplifier, is between −77.5 and −55.9 dBm (for the cases corresponding to 1 and 144 resonators, respectively), we used a 2-K HEMT noise temperature to define the noise level and a 245.76-MHz bandwidth to define the SNR bandwidth as the initial values for the calculation of this SNR table. The remainder of the signal chain could then be designed to inflict minimal SNR degradation.

Consider the first LNA on the IF board (model number HMC753) as an example; the table indicates that it has a gain of 17 dB, a noise figure of 1.5 dB, and a third-order intercept point output of 30 dBm. The table demonstrates that the first-stage LNA will degrade the SNR by 0.29 dB for both the 144-resonator case (SNR of 56.01 dB) and the one-resonator case (SNR of 34.41 dB) relative to the previous stage (the splitter[8]). The corresponding increase in the noise temperature with respect to the HEMT input is from 2.18 to 2.33 K.

This table provides not only the power, noise, and intermodulation values at each stage but also some indication of how the accumulated SNR or noise temperature propagates through the receiving chain. We used this table to select the components, arrange the sequence of components, and optimize the overall SNR and IMD on the IF boards. Once the components were set, we tuned the variable attenuators in the chain to fine-tune the performance. For 144 resonators, the overall SNR degradation is approximately 0.7 dB; thus, we consider the impact of the designed IF board to be sufficiently low.

---

[8] Each readout unit (which includes one ADC board, one IF board, and one FPGA board on the Al plate) covers a bandwidth of approximately 500 MHz, and we use two readout units to read out the signal from one feedline that contains resonator information in two separate 500 MHz bandwidths. The splitter was used to separate the signal into two readout units. The splitter in the system is shown the upper right plot in Figure 4.4

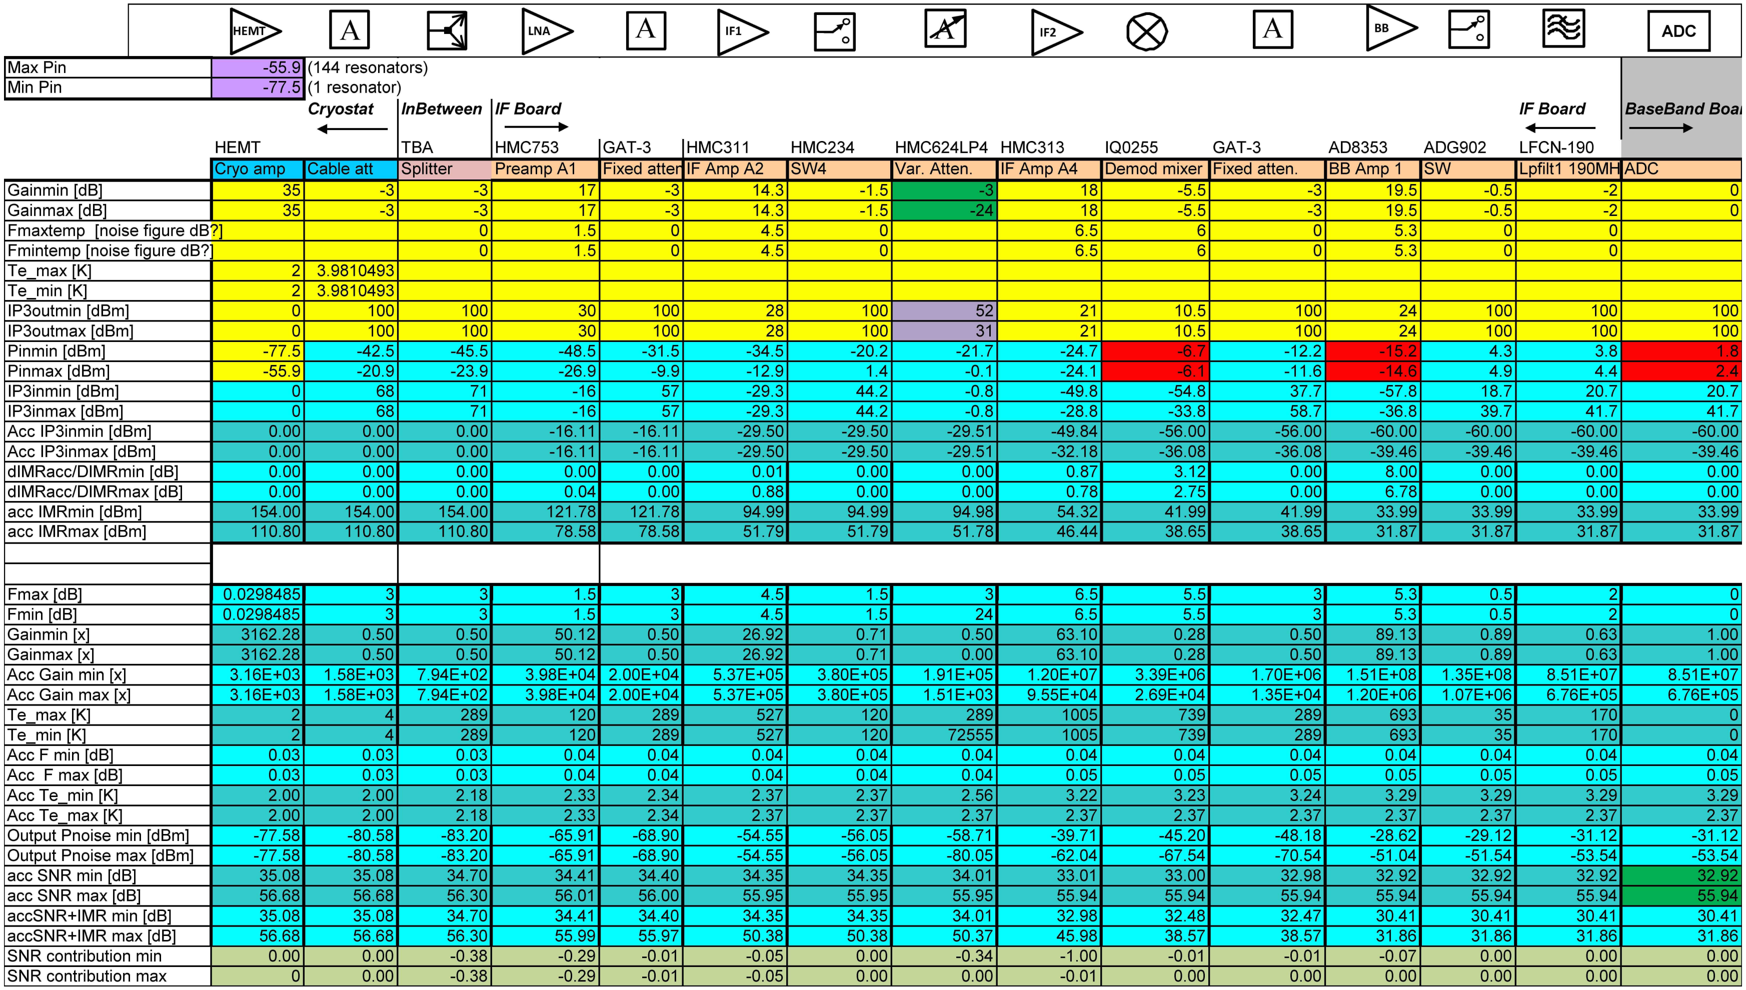

Max Pin	-55.9	(144 resonators)
Min Pin	-77.5	(1 resonator)

		Cryostat ←	InBetween	IF Board →										IF Board ←	BaseBand Boar →
	HEMT		TBA	HMC753	GAT-3	HMC311	HMC234	HMC624LP4	HMC313	IQ0255	GAT-3	AD8353	ADG902	LFCN-190	
	Cryo amp	Cable att	Splitter	Preamp A1	Fixed atten	IF Amp A2	SW4	Var. Atten.	IF Amp A4	Demod mixer	Fixed atten.	BB Amp 1	SW	Lpfilt1 190MH	ADC
Gainmin [dB]	35	-3	-3	17	-3	14.3	-1.5	-3	18	-5.5	-3	19.5	-0.5	-2	0
Gainmax [dB]	35	-3	-3	17	-3	14.3	-1.5	-24	18	-5.5	-3	19.5	-0.5	-2	0
Fmaxtemp [noise figure dB?]			0	1.5	0	4.5	0		6.5	6	0	5.3	0	0	
Fmintemp [noise figure dB?]			0	1.5	0	4.5	0		6.5	6	0	5.3	0	0	
Te_max [K]	2	3.9810493													
Te_min [K]	2	3.9810493													
IP3outmin [dBm]	0	100	100	30	100	28	100	52	21	10.5	100	24	100	100	100
IP3outmax [dBm]	0	100	100	30	100	28	100	31	21	10.5	100	24	100	100	100
Pinmin [dBm]	-77.5	-42.5	-45.5	-48.5	-31.5	-34.5	-20.2	-21.7	-24.7	-6.7	-12.2	-15.2	4.3	3.8	1.8
Pinmax [dBm]	-55.9	-20.9	-23.9	-26.9	-9.9	-12.9	1.4	-0.1	-24.1	-6.1	-11.6	-14.6	4.9	4.4	2.4
IP3inmin [dBm]	0	68	71	-16	57	-29.3	44.2	-0.8	-49.8	-54.8	37.7	-57.8	18.7	20.7	20.7
IP3inmax [dBm]	0	68	71	-16	57	-29.3	44.2	-0.8	-28.8	-33.8	58.7	-36.8	39.7	41.7	41.7
Acc IP3inmin [dBm]	0.00	0.00	0.00	-16.11	-16.11	-29.50	-29.50	-29.51	-49.84	-56.00	-56.00	-60.00	-60.00	-60.00	-60.00
Acc IP3inmax [dBm]	0.00	0.00	0.00	-16.11	-16.11	-29.50	-29.50	-29.51	-32.18	-36.08	-36.08	-39.46	-39.46	-39.46	-39.46
dIMRacc/DIMRmin [dB]	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.87	3.12	0.00	8.00	0.00	0.00	0.00
dIMRacc/DIMRmax [dB]	0.00	0.00	0.00	0.04	0.00	0.88	0.00	0.00	0.78	2.75	0.00	6.78	0.00	0.00	0.00
acc IMRmin [dBm]	154.00	154.00	154.00	121.78	121.78	94.99	94.99	94.98	54.32	41.99	41.99	33.99	33.99	33.99	33.99
acc IMRmax [dBm]	110.80	110.80	110.80	78.58	78.58	51.79	51.79	51.78	46.44	38.65	38.65	31.87	31.87	31.87	31.87
Fmax [dB]	0.0298485	3	3	1.5	3	4.5	1.5	3	6.5	5.5	3	5.3	0.5	2	0
Fmin [dB]	0.0298485	3	3	1.5	3	4.5	1.5	24	6.5	5.5	3	5.3	0.5	2	0
Gainmin [x]	3162.28	0.50	0.50	50.12	0.50	26.92	0.71	0.50	63.10	0.28	0.50	89.13	0.89	0.63	1.00
Gainmax [x]	3162.28	0.50	0.50	50.12	0.50	26.92	0.71	0.00	63.10	0.28	0.50	89.13	0.89	0.63	1.00
Acc Gain min [x]	3.16E+03	1.58E+03	7.94E+02	3.98E+04	2.00E+04	5.37E+05	3.80E+05	1.91E+05	1.20E+07	3.39E+06	1.70E+06	1.51E+08	1.35E+08	8.51E+07	8.51E+07
Acc Gain max [x]	3.16E+03	1.58E+03	7.94E+02	3.98E+04	2.00E+04	5.37E+05	3.80E+05	1.51E+03	9.55E+04	2.69E+04	1.35E+04	1.20E+06	1.07E+06	6.76E+05	6.76E+05
Te_max [K]	2	4	289	120	289	527	120	289	1005	739	289	693	35	170	0
Te_min [K]	2	4	289	120	289	527	120	72555	1005	739	289	693	35	170	0
Acc F min [dB]	0.03	0.03	0.03	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
Acc F max [dB]	0.03	0.03	0.03	0.04	0.04	0.04	0.04	0.04	0.05	0.05	0.05	0.05	0.05	0.05	0.05
Acc Te_min [K]	2.00	2.00	2.18	2.33	2.34	2.37	2.37	2.56	3.22	3.23	3.24	3.29	3.29	3.29	3.29
Acc Te_max [K]	2.00	2.00	2.18	2.33	2.34	2.37	2.37	2.37	2.37	2.37	2.37	2.37	2.37	2.37	2.37
Output Pnoise min [dBm]	-77.58	-80.58	-83.20	-65.91	-68.90	-54.55	-56.05	-58.71	-39.71	-45.20	-48.18	-28.62	-29.12	-31.12	-31.12
Output Pnoise max [dBm]	-77.58	-80.58	-83.20	-65.91	-68.90	-54.55	-56.05	-80.05	-62.04	-67.54	-70.54	-51.04	-51.54	-53.54	-53.54
acc SNR min [dB]	35.08	35.08	34.70	34.41	34.40	34.35	34.35	34.01	33.01	33.00	32.98	32.92	32.92	32.92	32.92
acc SNR max [dB]	56.68	56.68	56.30	56.01	56.00	55.95	55.95	55.95	55.94	55.94	55.94	55.94	55.94	55.94	55.94
accSNR+IMR min [dB]	35.08	35.08	34.70	34.41	34.40	34.35	34.35	34.01	32.98	32.48	32.47	30.41	30.41	30.41	30.41
accSNR+IMR max [dB]	56.68	56.68	56.30	55.99	55.97	50.38	50.38	50.37	45.98	38.57	38.57	31.86	31.86	31.86	31.86
SNR contribution min	0.00	0.00	-0.38	-0.29	-0.01	-0.05	0.00	-0.34	-1.00	-0.01	-0.01	-0.07	0.00	0.00	0.00
SNR contribution max	0	0.00	-0.38	-0.29	-0.01	-0.05	0.00	0.00	-0.01	0.00	0.00	0.00	0.00	0.00	0.00

Figure 2.15: Signal and noise propagation along the analog readout receiving chain.

Chapter 3

Filtering for Submillimeter and Millimeter Imaging

In this chapter, I will discuss the design, optimization, and measured results of the lumped-element BPF network used for submillimeter and millimeter imaging. Careful calculations are required during the filter design process to achieve a filter response that meets our needs. Circuit and layout models are studied to represent the filter transfer function with a physical layout. The measured results exhibit good frequency response and strong agreement with the predicted results. The proposed BPF network is used as part of the multicolor superconducting microresonator imaging array instrument MUSIC for the CSO.

3.1 Introduction of Submillimeter-Wavelength Filtering Technology and the MUSIC Filtering Design

3.1.1 Metal Mesh Filter and Lithographic Filter

Quasi-optical metal-mesh filters are widely used for filtering in far-infrared and submillimeter-wavelength applications [25]. With the rapid development of new antenna-coupled bolometers for submillimeter and millimeter detection [2, 5–11, 26], it has become possible to implement large arrays of sensitive detectors on a silicon wafer, enabling the use of a lithographic microstrip filter to define the frequency bands and achieve multicolor imaging. Compared with metal-mesh filters, this approach provides significant advantages:

- Compact ($\approx 100\ \mu$m) lithographic filters are easy to fabricate and provide considerable flexibility in frequency selection.
- The combined filter bank can be used to feed multiple detectors from a broadband antenna to increase the instrument throughput and to enable multicolor imaging.

- The polarization-separation stage can be placed before the filters, which is important because antenna-coupled technologies are inherently polarization selective.

For these reasons, submillimeter and millimeter lithographic filters are being pursued for a wide variety of projects [27–32] including the MUSIC instrument that we constructed.

3.1.2 Wafer Design and Filtering Requirements of MUSIC

The BPF network is part of the camera wafer. The components of the wafer include the broadband phased-array antenna, BPF network, and KID resonators, as shown in Figure 3.1.

The astronomical signals are first captured by the broadband antenna, after which the microstrip line from the antenna binary summing tree passes the wide-band signal to the BPF network to define four different frequency bands (the four colors in the sky map). The output of each BPF will feed into one MKID. We read out the astronomical signal through the transmission line coupled to the detector.

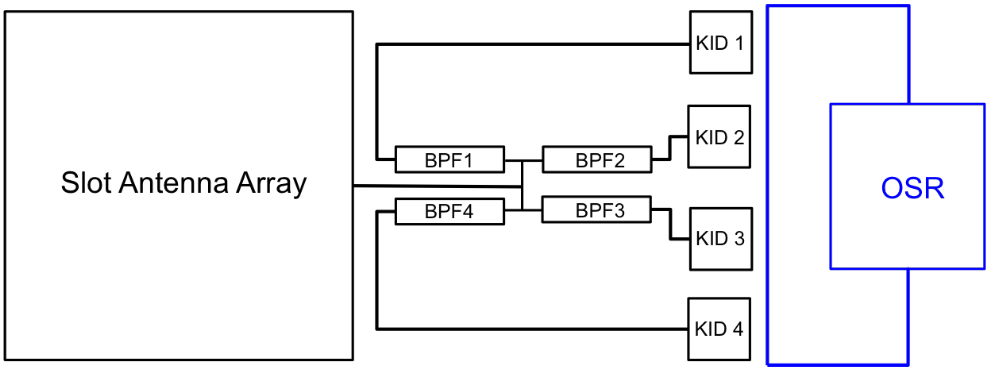

Figure 3.1: Block diagram of the phased-array antenna, BPFs, KIDs, and OSR.

Unlike most filter design processes, there are no strict requirements for the cut-off frequency, bandwidth, or other typical filter response parameters when designing the BPF network. Instead, because the flux density from a dusty galaxy is the primary intended source type for the instrument, maximization of the sensitivity of each band for this type of signal source serves as the primary design goal. We determined all filter parameters with this goal in mind.

The following basic requirements can be used to determine the frequency bandwidth range:

- We are interested in submillimeter wavelengths.
- The superconducting aluminum that we used has an energy gap of 88 GHz.
- The broadband phased-array antenna also has certain bandwidth limitations (Figure 3.2 presents a typical antenna efficiency plot).

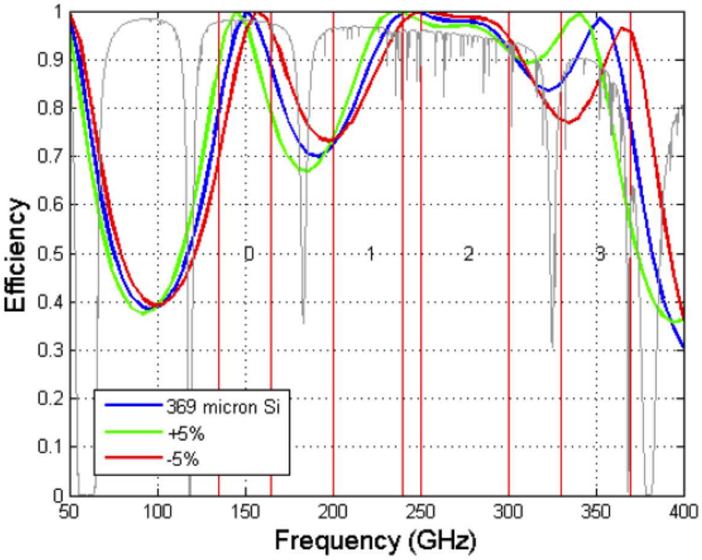

Figure 3.2: Simulated antenna transmission with ±5% variants in the substrate thickness. The blue line represents the simulated antenna transmission for 369-μm-thick silicon. The green and red lines represent the simulated antenna transmission for +5% and −5% silicon thicknesses, respectively. The gray line represents the water-vapor transmission at the CSO. The vertical red lines represent several possible transmission windows allowed by the antenna and the water-vapor transmission. Plot courtesy of Peter Day.

These requirements together provide a general frequency range from 88 to 450 GHz. In addition, the water-vapor line has a significant effect on the sensitivity of ground-based cameras in the submillimeter frequency range. In our case, the MUSIC instrument was to be installed at the CSO (located at the summit of Mauna Kea), where the water-vapor line defines several transmission windows. In Figure 3.2, the gray line in the background represents the water-vapor transmission, and the vertical red lines indicate several possible windows for the BPF network. We now have a basic idea of the frequency response required for the BPFs: each BPF should be centered at 100–400 GHz with a 20–40% bandwidth[1]. To minimize the loss of each BPF on the wafer, we employed Nb as the ground (150-nm-thick) and wiring (400-nm-thick) layers with 400-nm-thick silicon dioxide or silicon nitride as the dielectric layer fabricated on a 369-μm-thick silicon wafer (as shown in Figure 3.3). Nb is an excellent superconducting microstrip material that provides low loss for the microstrip line under superconducting conditions (the BPF wafer operates at 230 mK).

---

[1] During the MUSIC development process, we investigated two different antenna bandwidth coverages: 120–375 GHz, including a frequency band centered at 150 GHz, and 175–425 GHz, including a frequency band centered at 405 GHz. Figure 3.2 shows the coverage for the first case.

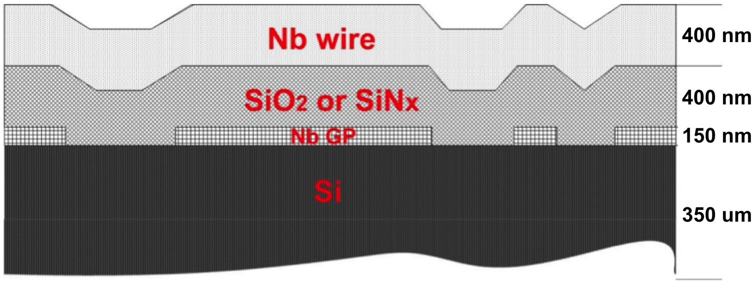

Figure 3.3: Dielectric layers for BPF fabrication.

3.1.3 Lumped- and Distributed-Element Design

For a lithographic BPF centered at ∼ 300 GHz, we studied both lumped-element and distributed-element designs.

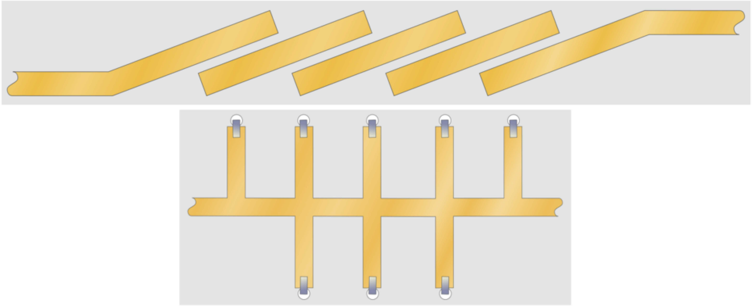

Figure 3.4: Two common designs for distributed BPFs: (top) coupled lines and (bottom) a stub via an edge.

For the distributed filter design, two commonly used designs are depicted in Figure 3.4:

- Coupled-line designs include a capacitive gap, a hairpin, or an interdigitated structure. Figure 3.5 shows a design with coupled lines. In the left-hand plot, the design layout is presented in the upper-left corner, and its response is shown as a red line. Spurious resonance frequencies are evident at harmonics of 100 GHz. If we wish to design a BPF with a single passband, one solution is to cascade another LPF after the coupled-line BPF, as shown in the right-hand plot, where the LPF design is illustrated in the upper-right corner (this figure is reprinted from [33]).

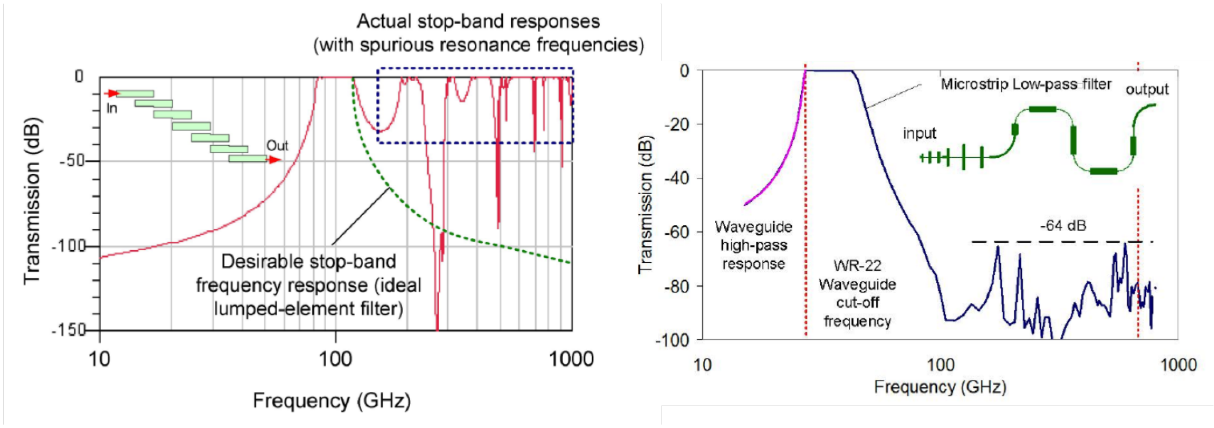

Figure 3.5: Example of spurious frequencies arising in distributed structures: a coupled-line filter. Reprinted from [33].

- Stub designs include a via an edge or a short circuit with an impedance transformer. The upper-left plot in Figure 3.6 depicts a design with a via edge. The frequency response is indicated by a black line on the right-hand side of the figure. Periodic harmonic passbands are evident and centered at 2.5, 7.5, 12.5, and 17.5 GHz. If we wish to design a BPF with a single passband, one solution is to use an electronic–magnetic band gap (EBG) structure in the ground plane; as shown in the lower-left plot, the ground-plane holes serve as LPFs. The frequency response of the design with the EBG structure is represented as a red line in the plot at the right (this figure is reprinted from [34]).

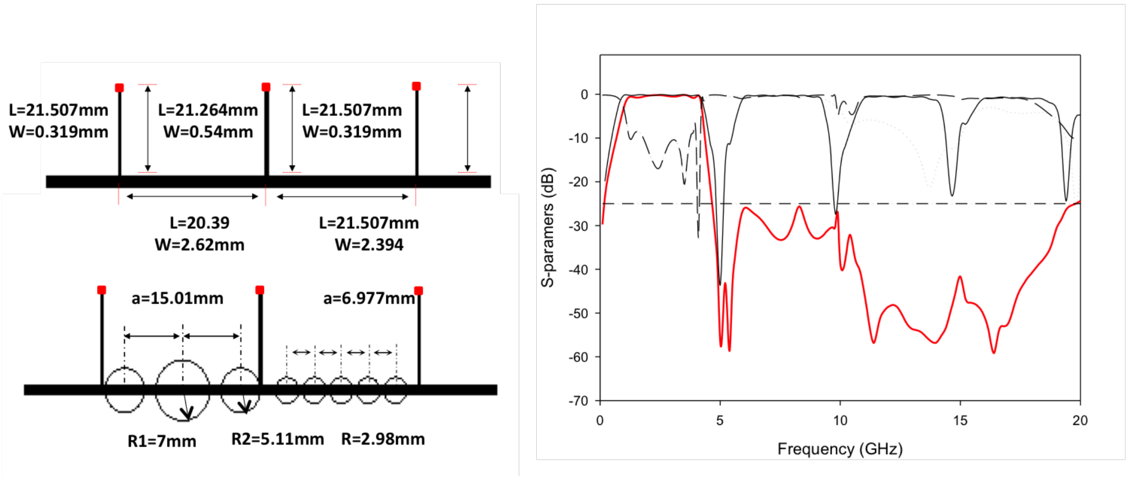

Figure 3.6: Example of spurious frequencies arising in distributed structures: a via-edge filter. Reprinted from [34].

We find that the lumped-element design offers clear advantages for our application:

- First, the lumped-element design is more compact in size than a filter based on half-wave resonant structures. For example, a distributed BPF centered at 300 GHz is approximately 0.5–1 mm in size, whereas the lumped design can be as small as 0.1 mm; in other words, the lumped-element design is approximately one-fifth of the size of the distributed design.

- Second, there are no "natural" re-entrant modes for lumped designs, which is important for our filter network (or "filter bank") application because the harmonics from one filter will affect the other filters connected to it. There are several potential adaptations of distributed filters (e.g., a shorted edge, an electromagnetic band gap, or an additional LPF) that could be adopted to suppress the undesired passbands (as shown in Figures 3.5 and 3.6), but this will increase the complexity of the fabrication process (e.g., via holes through layers) or introduce an additional structure.

- Finally, lumped elements occasionally offer more flexibility regarding the ranges of inductances and capacitances, which may result in better filter response (e.g., a wider passband bandwidth in our case). However, we must consider the parasitic impedance, particularly for inductors; the fabrication tolerance is also a problem in some cases because of the smaller size.

3.1.4 MUSIC BPF Design Steps

To implement the lumped-element lithographic BPF network, we utilized the following steps (each step will be presented in greater detail in later sections):

1. BPF circuit model: The circuit model we used was derived by our group for previous projects (Figure 3.7) [30]. We relied on this basic circuit model and improved our calculation and optimization approaches.

2. Calculation of the ideal inductances and capacitances: The initial inductances and capacitances were calculated to maximize the overall sensitivity for each frequency band[2]. The atmospheric transmission windows limit the width of the filters, and the filters should not overlap because of the risk of confusing the spectral information. We derived a set of optimal inductances and capacitances for each filter on the basis of this calculation.

3. Translation of the ideal values into the layout: We implemented a spiral inductor and parallel-plate capacitor on the wafer.

4. Optimization of the design through simulation: After determining the geometry of the inductors and capacitors, the inductors and capacitors were fine-tuned at the layout level. All

---

[2] This calculation was performed to obtain the optimal noise-equivalent flux density and noise-equivalent power, NEFD_1mm and NEP, respectively. The NEFD of a system is the level of flux density required to equal the noise present in the system. It is a measure that is commonly used by astronomers for determining the accuracy of observations. The NEFD can be related to the light detector's NEP for a certain collection area and photon bandwidth.

substrate and fabrication parameters were studied, and their effects on the filter performance will be discussed.

5. Consideration of the feed network: We found that the feed network that connects the four BPFs together has a nonnegligible effect on the frequency response of each BPF. We accounted for the feed network by combining the S-parameter matrices together: four 2×2 matrices for each filter and one 5×5 matrix for the feed network.

6. Construction of a Sonnet model for accurate prediction: As mentioned in step 3, rather than considering the inductors and capacitors as ideal components, we must consider many practical factors (detailed in Section 3.5). A Sonnet model was developed to yield reliable predictions.

3.2 BPF Calculations

3.2.1 BPF Circuit Model

A filter circuit model can be specified by its transfer function, e.g., Butterworth, Chebyshev, Bessel, Elliptic, or Gaussian. The circuit model we used was developed by our group for previous projects (as shown in Figure 3.7) [30]. The structure is a modification of the Butterworth filter and consists of eight lumped components including three inductors (L), three series capacitors (C1 and C2), and three parallel capacitors (C3).

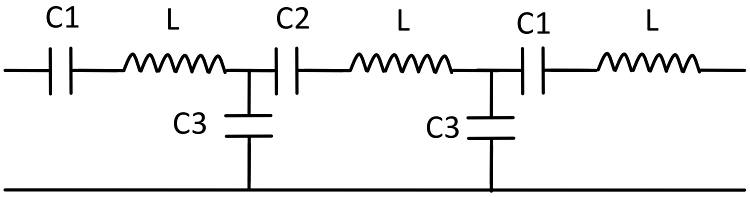

Figure 3.7: BPF circuit model. Reprinted from [30].

3.2.2 BPF Design Criteria for Astronomical Signals

In our MUSIC application, the lithographic broadband phased-array antenna covers a wide frequency range, and we wish to use BPFs to divide the bandwidth into four frequency bands that correspond to the four different “colors” desired for MUSIC.

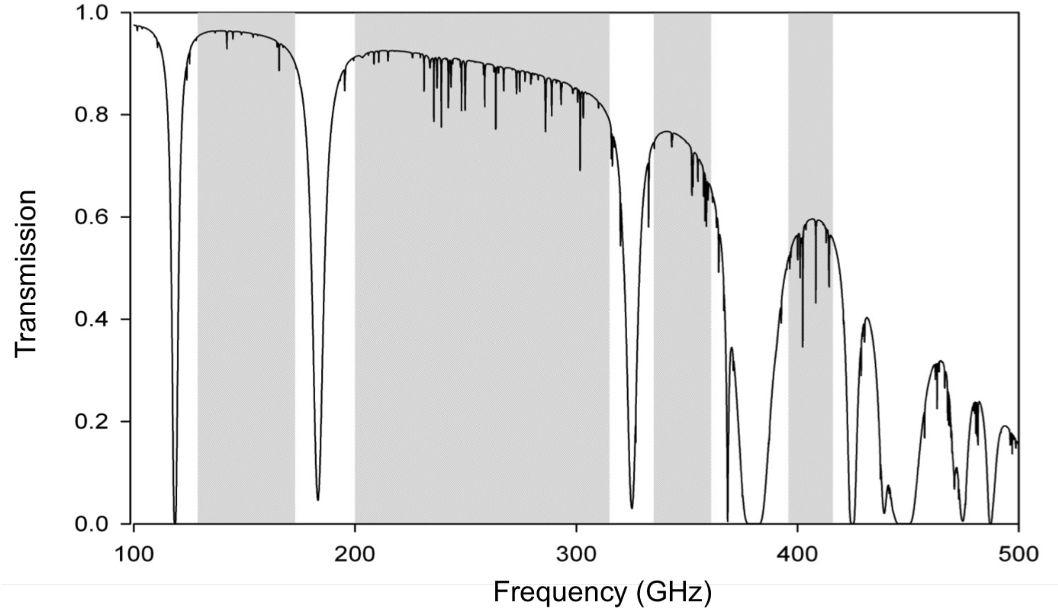

Figure 3.8: Atmospheric transmission for 1.5 mm of perceptible water vapor.

As shown in Figure 3.8, we will need four BPFs centered at 225, 280, 350, and 405 GHz (or use a 150-GHz window rather than a 405-GHz window) with two BPFs sharing the same water-vapor window from 180 to 325 GHz to address the four transmission windows (shadowed areas). In addition to the atmospheric transmission shown in Figure 3.8, we also optimized the filter parameters by considering the following factors to determine the filter response that maximizes the sensitivity: the phased-array-antenna transmission (as shown in Figure 3.2) and the optical loading from the dewar and telescope, among other components. We also defined four different bands with controllable overlap among them to avoid undesired correlations between bands (the detailed calculations are presented in Appendix C).

3.2.3 Mathematical Model and Optimization

During the calculation, each filter transmission was generated from the combined impedance of one set of L, C1, C2, and C3 (as shown in Figure 3.7) using the 20-Ω characteristic impedance of the microstrip transmission line (the value of 20 Ω corresponds to a 2.5-μm-wide microstrip line). The transmission coefficient S_{21} was calculated from the filter transfer function as follows:

$$|S_{12}| = \frac{40}{C_3^2 S^2 \left(\frac{A+B*C}{C_3 S} + C_3 S * A \left(A + B * (C + \frac{1}{C_3 S}))\right)\right)} \tag{3.1}$$

$$A = \frac{20 + LS + \frac{1}{C_1 S}}{C_3 S} \tag{3.2}$$

$$B = 20 + \frac{1}{C_1 S} + \frac{1}{C_3 S} + LS \tag{3.3}$$

$$C = \frac{1}{C_2 S} + LS \tag{3.4}$$

where S = jω and $\omega = 2\pi f$.

The inductors and capacitors were tuned in the ranges shown in Table 3.1, where the initial start and end values represent the range that could yield the desired frequency response. We also required that the inductances and capacitances could be physically implemented as lumped elements on our wafer.

Table 3.1: Summary of the inductance and capacitance tuning ranges

Parameter	Start	End	Step	Units
L	70	150	0.1	pH
C1	1.6	8	0.01	fF
C2	1	11	0.01	fF
C3	16	60	0.1	fF

From the optimization using lumped-element parameters, we investigated the effects of the lumped-element values on the final BPF response:

- Parallel C: shifts the upper cut-off frequency of a BPF without affecting the center frequency.
- Inductance L: shifts the center frequency without affecting the bandwidth or passband shape.
- Series C: both C1 and C2 (as shown in Figure 3.7) can be used to optimize the passband shape.

3.2.4 Bands Sharing the Same Frequency Window

For the case in which two frequency bands (band 1 and band 2) share the same water-vapor windows, we calculated the desired filter response with two objectives:

- Minimizing the difference between the band 1 NEFD and the band 2 NEFD.
- Achieving a small controllable power overlap between bands 1 and 2[3].

---

[3] An overlap of less than 5% with respect to the transmission power was allowed. Although this choice was somewhat arbitrary, it was driven by the detrimental effect that a tighter criterion of 1% would have exerted on the sensitivities by requiring the filters to be too narrow.

Two variables were defined to facilitate the optimization of bands 1 and 2: the cut-off frequency between the bands and the separation from the cut-off frequency to the right edge of band 1 or to the left edge of band 2 (where the 50% transmission point was used to define the edge).

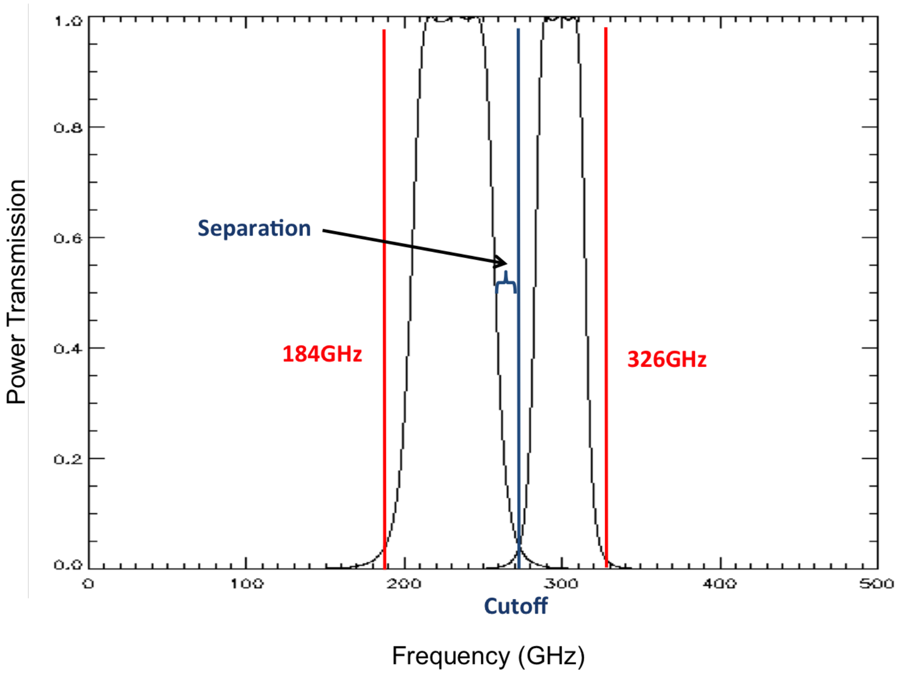

Figure 3.9: Two variables for BPF optimization in the same transmission window: cut-off and separation.

For a given separation value, we optimized band 1 and band 2 for all cut-off frequencies and sought equal NEFD_1mm values for bands 1 and 2. The vertical axis of Figure 3.10 represents the resulting NEFDs when we achieved equal NEFD_1mm values for bands 1 and 2; the horizontal axis represents the power overlap fraction of bands 1 and 2 under this condition. On the basis of this plot, we can make the following statements:

- The NEFD_1 mm values will not substantially improve when the overlap is reduced below 2%.
- The equal NEFD_1mm values of bands 1 and 2 that we can obtain range from 0.04 to 0.05 $\mathrm{Jy}/\sqrt{\mathrm{Hz}}$.

Additionally, during the optimization process, we also find the following:

- To achieve equal NEFD_1mm values for bands 1 and 2, the cut-off frequency (one of the two variables that we defined) must be approximately 265–275 GHz.
- For a wider separation between the bands, a lower cut-off frequency at which the bands achieve equal NEFD_1mm values is obtained.

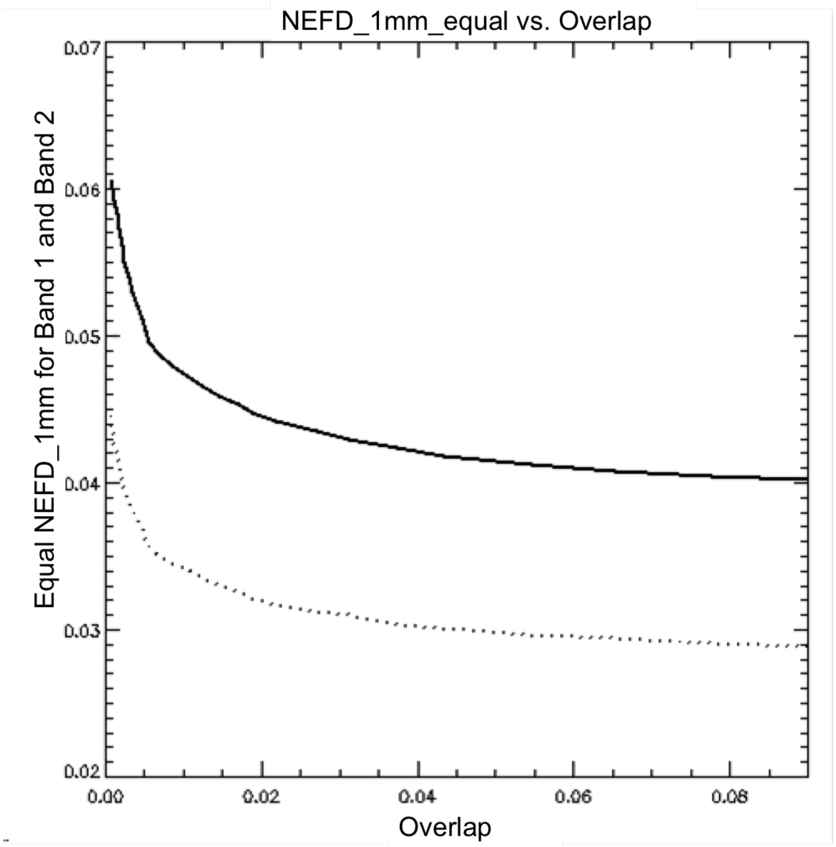

Figure 3.10: Power overlap fraction for bands 1 and 2 and the corresponding NEFD_1mm values (in units of Jy/$\sqrt{\text{Hz}}$): (dotted line) total NEFD_1mm for bands 1 and 2 together; (solid line) NEFD_1mm of band 1, which is equal to the NEFD_1mm of band 2.

3.2.5 Scaling Factor for the Flux Density from Dusty Galaxies

To explicitly optimize the filter design, we used IDL software, which included the calculation of the filter transfer function, the calculation of the NEP and NEFD values with the related transmission coefficients (e.g., antenna, optical transmission), and the sensitivity optimization. The plots presented in Figure 3.11 provide an example of the actual frequency response and the corresponding dusty-galaxy power of each band (the detailed calculations are presented in Appendix C). Subplots (a) and (b) show the frequency response of each band, and subplots (c) and (d) show the dusty-galaxy energy that survives through BPF filtering and atmospheric transmission in units of janskys (Jy). This plot illustrates that there is no significant power overlap of the BPF response with the atmospheric water-vapor transmission (Figure 3.8), which suggests that the left edge of band 1 and the right edge of band 2 are well defined for the calculation.

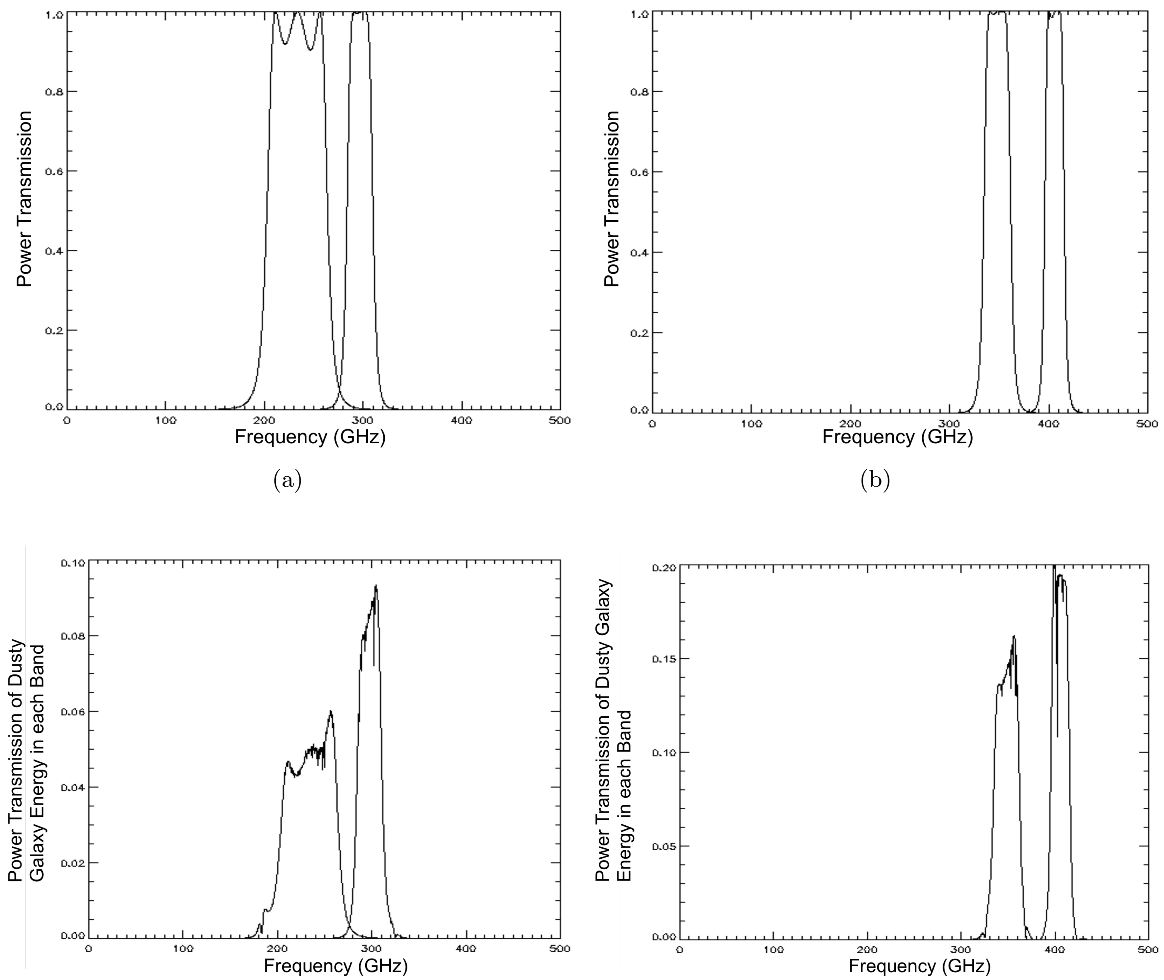

(a) (b)

(c) (d)

Figure 3.11: Frequency response and dusty-galaxy energy in each band. (a) and (b) Power transmission of the four BPFs, as calculated in IDL. (c)and (d) Dusty-galaxy energy (Jy) in each frequency band.

In subplots (c) and (d) of Figure 3.11, a $frequency^{3.0}$ scaling of the flux density from dusty galaxies is the default assumption used for the astronomical signal source. We recalculated the results for the NEFD_1mm values and the overlap using the power-law exponents of 2.5 and 3.5. The sensitivity results (NEFD_1mm, in units of Jy/$\sqrt{\text{Hz}}$) for each BPF are summarized in Table 3.2. We found that changing the frequency scaling to $frequency^{2.5}$ or $frequency^{3.5}$ did not change our optimized filter designs; thus, we used the $frequency^{3.0}$ scaling as the default assumption for all calculations.

Table 3.2: NEFD_1mm (Jy/$\sqrt{\text{Hz}}$) for each BPF with various frequency scalings for the astronomical source flux

	Band 1	Band 2	Band 3	Band 4
$frequency^{2.5}$	0.0468	0.0420	0.0517	0.0739
$frequency^{3.0}$	0.0470	0.0472	0.0480	0.0546
$frequency^{3.5}$	0.0472	0.0531	0.0445	0.0546

From the optimization of each individual BPF, we also found that the optimization tends to favor a sharp cut-off over flat transmission across the passband. In other words, the optimized sensitivity does not necessarily require smooth and flat passband performance of the BPF; instead, given the constraints placed on the available spectral bandwidth by the atmospheric transmission, wider passband coverage results in better sensitivity.

3.3 Implementation of the Calculated BPF Values in the Layout

In the previous section, we used IDL to optimize the inductances and capacitances. However, IDL uses only numerical values to perform calculations; it cannot account for the many practical factors that also affect the filter response, e.g., the kinetic inductance, the coupling of the inductor to the ground plane, and the parasitic capacitance. In this section, we begin to construct the layout of the BPF network.

As shown in Figure 3.12, we implemented each inductance using a spiral inductor and each capacitance using a parallel-plate capacitor on the wafer. The series capacitances C1 and C2 were split into two equal-valued capacitors sandwiching a spiral inductor for ease of fabrication. The dielectric constants of silicon dioxide ($\epsilon_r = 3.9$) and silicon nitride ($\epsilon_r = 7.0$) were estimated on the basis of the measurement of the wafers our group developed for the filter, antenna, and detectors. These values were the best estimations that we could obtain for low-temperature submillimeter and millimeter applications.

To implement the BPF network, we connected four BPFs (as shown in Figure 3.13).

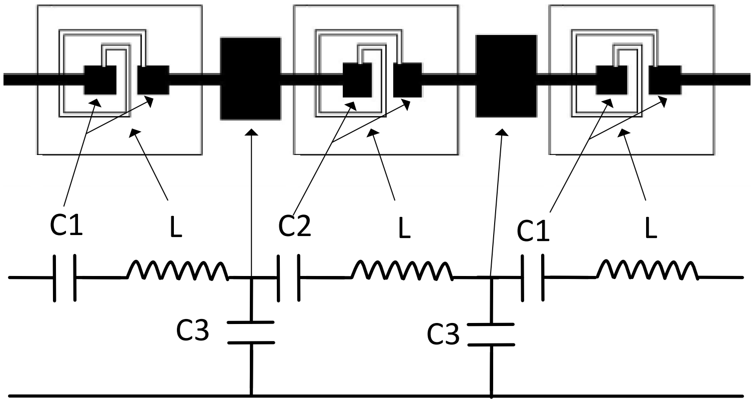

Figure 3.12: BPF circuit diagram and the corresponding layout.

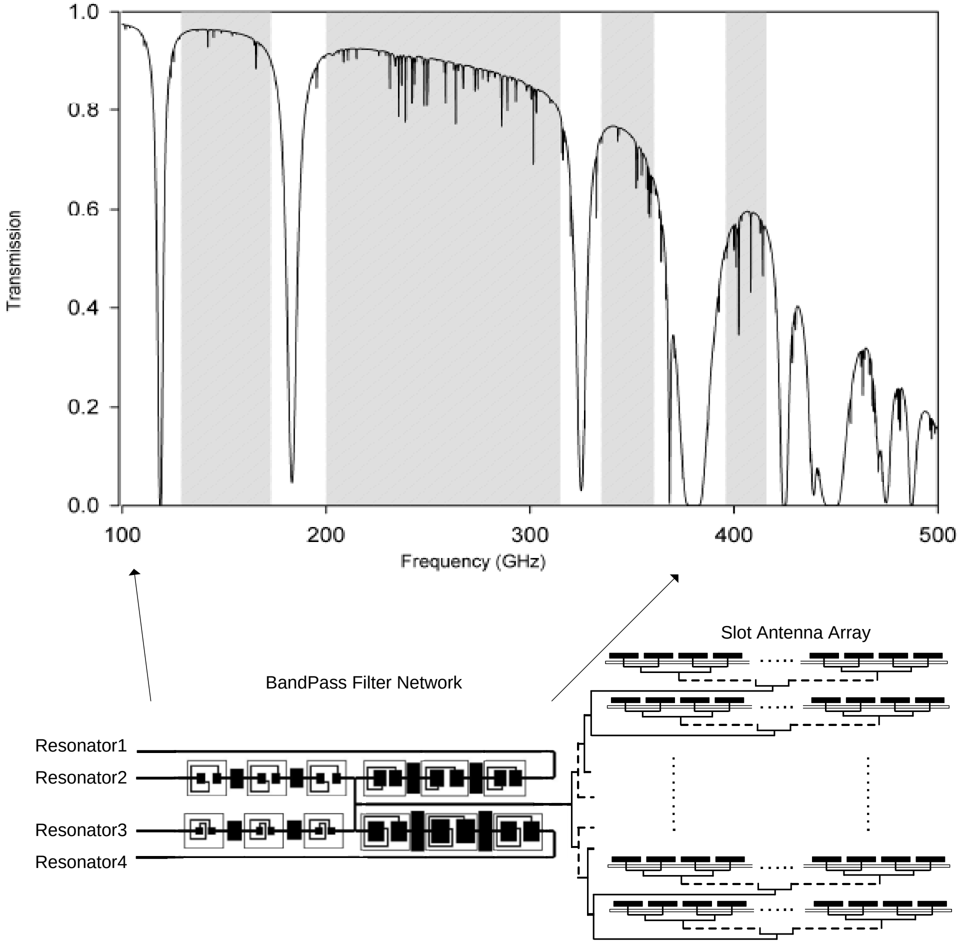

Figure 3.13: Layout of the BPF network with a slot-antenna summing tree and an atmospheric transmission window at 1 mm of water vapor.

3.4 BPF Feed Network

In this section, I will discuss the analysis of the four filters jointly with the feed network, which was performed to ensure that the behavior was acceptable and that the artifacts created by the interactions among the filters were acceptable. The feed network design is illustrated in Figure 3.14, and each section of the microstrip line (labeled as Mline in the figure) was studied.

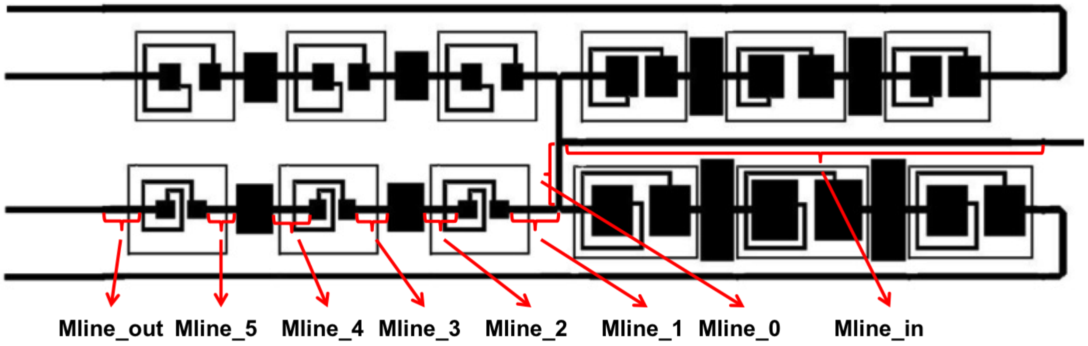

Figure 3.14: Mask layout of the BPF network. "Mline" stands for microstrip line, and the sections of the microstrip line in the BPF network are labeled Mline 0 to Mline 5, Mline input, and Mline output.

Advanced Design System (ADS) software was used to study the BPF network. ADS can construct a simulation model by treating the inductors and capacitors as lumped elements. The microstrip lines (all labeled Mline sections in Figure 3.14) between BPFs were also taken into account. In the ADS simulator, the inductors and capacitors were specified in terms of their inductances and capacitances, respectively, and the microstrip lines were specified in terms of the microstrip-line substrate, width, length, and losses.

First, we simulated each individual BPF's power transmission, the results are shown in Figure 3.16. This is the results simulated in ADS without considering any microstrip lines and treating the inductors and capacitors as ideal lumped elements. When we design BPF network, the power transmission of each BPF might change, because the BPFs are designed to have a matched impedance for the passband and a high impedance for the out-of-passband frequency range. When the BPFs are connected to the same feed junction point, each BPF's out-of-band response will affect the other bands. In our case, the microstrip line from the phased-array antenna has an impedance of 20–25 Ω. The simulation of input impedance for each BPF without considering the microstrip lines[4] is shown in Figure 3.17.

---

[4] The input impedance can only be simulated when the BPFs are not connected with feed network.

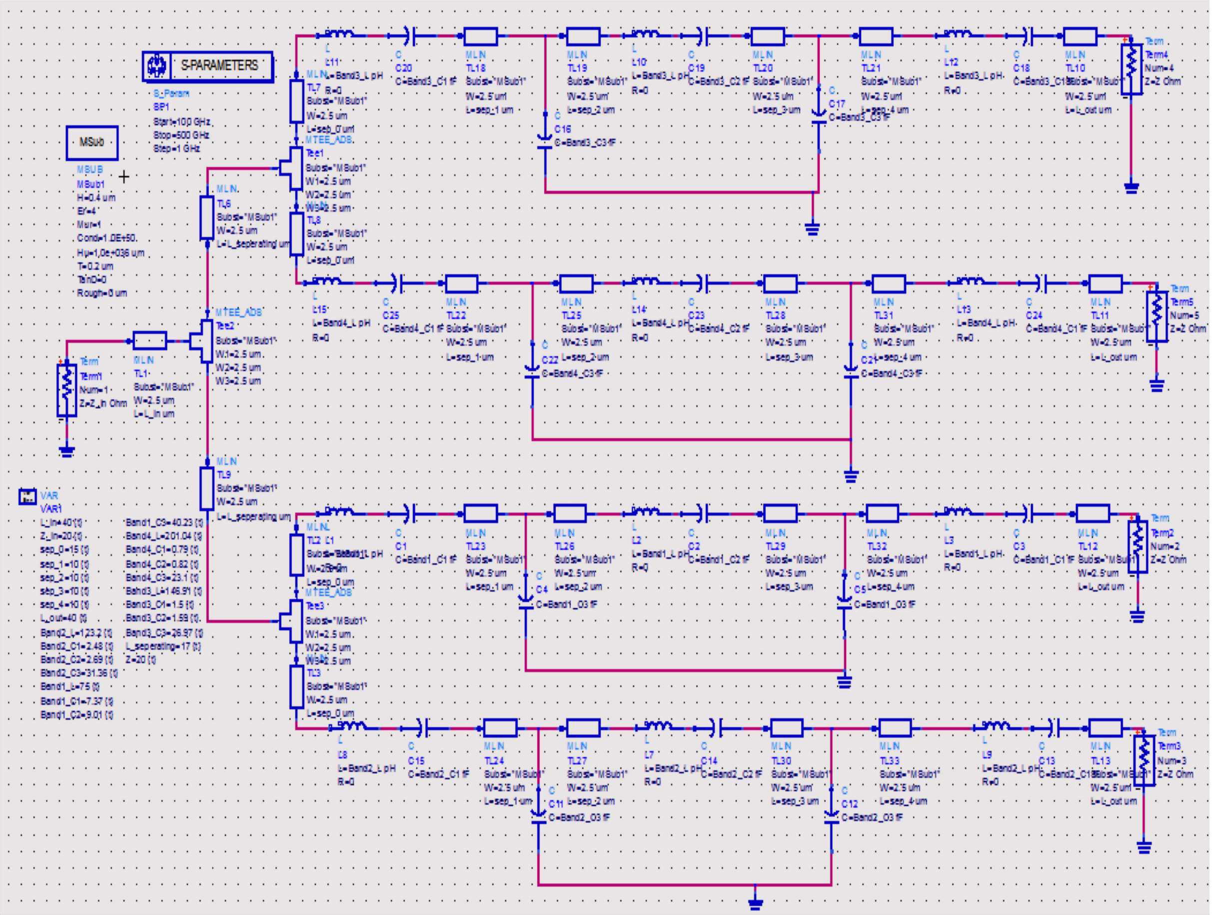

Figure 3.15: Simulation setup for the BPF network in ADS.

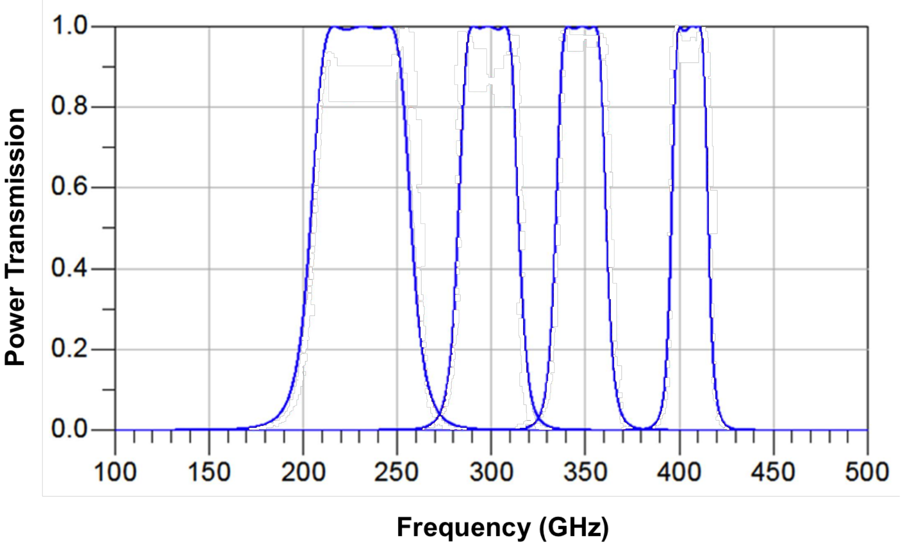

Figure 3.16: ADS simulations of four BPFs without connections to a common input port.

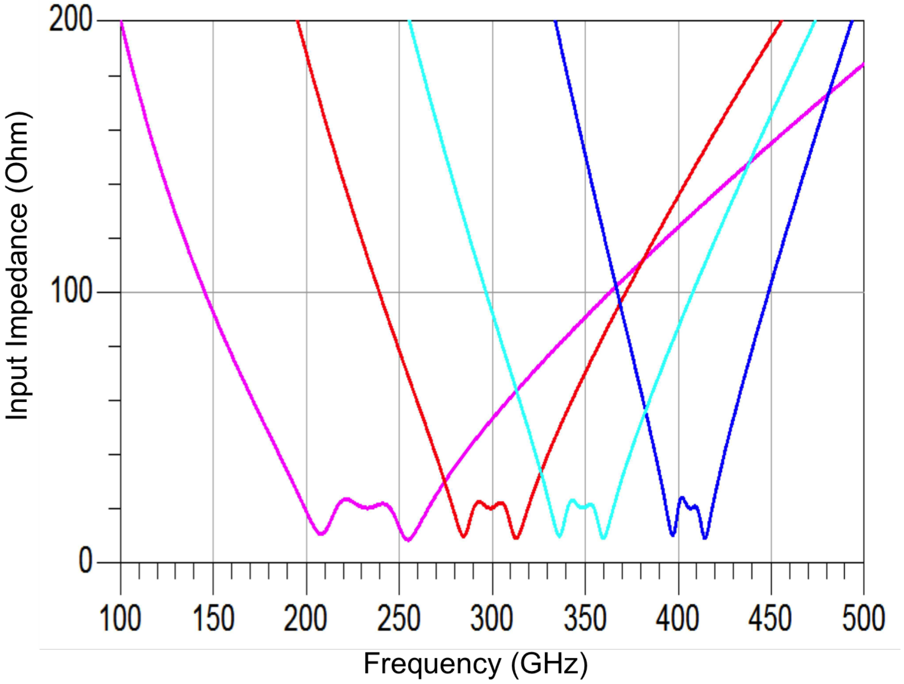

Figure 3.17: ADS simulations of the input impedance for each BPF without feed network.

Next, we study the effect of various microstrip-line lengths between the BPFs in the ADS simulation. The frequency responses (S_{12}) and impedances (Z_{in}) of the BPFs for various combinations of microstrip-line lengths are summarized in Figure 3.18. In subplots (a), (c) and (e), the blue lines represent the power transmission of each BPF for comparision, and the red dotted lines represent the power transmission when four BPFs are connected to the feed network. It is apparent that passband transmission is significantly degraded when the filters are connected together. In subplots (b), (d) and (f), the input impedance of each BPF with different microstrip-line lengths was simulated when BPFs were not connected to the feed network. In subplots (a) and (b), all the microstrip lines are ignored; in subplots (c) and (d), the nominal values (the microstrip line Mline0 = 17 μm; Mline1 = 15 μm) were used; in subplots (e) and (f), the microstrip line Mline0 = 17 μm and Mline1 length range from 7 μm to 16 μm.

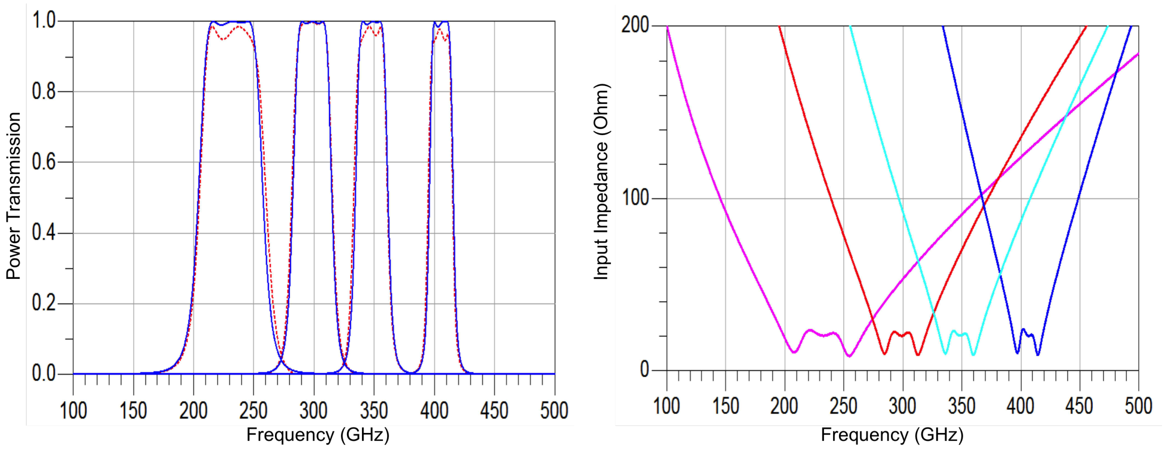

(a) S_{12}: Mlines = 0 (b) Z_{in}: Mlines = 0

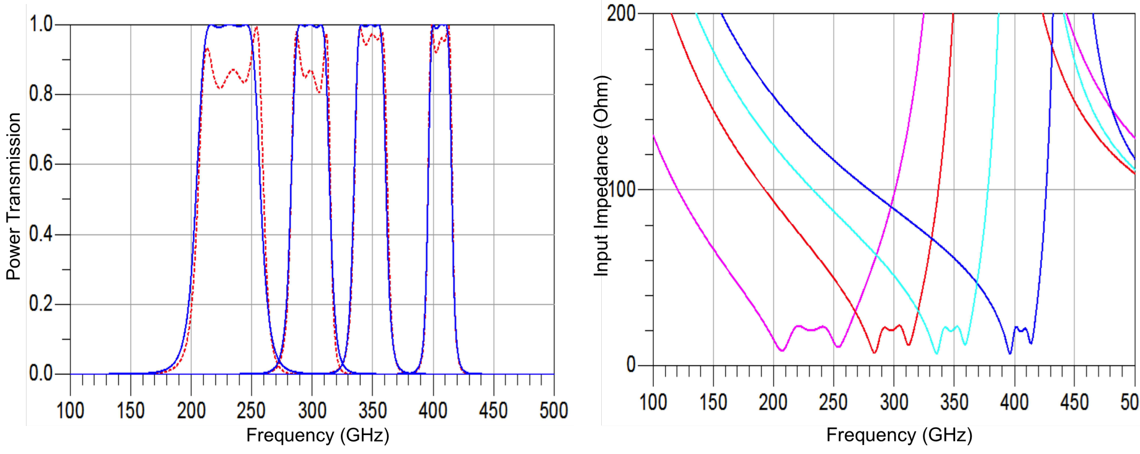

(c) S_{12}: Mline0 = 17 μm; Mline1 = 15 μm (d) Z_{in}: Mline0 = 17 μm; Mline1 = 15 μm

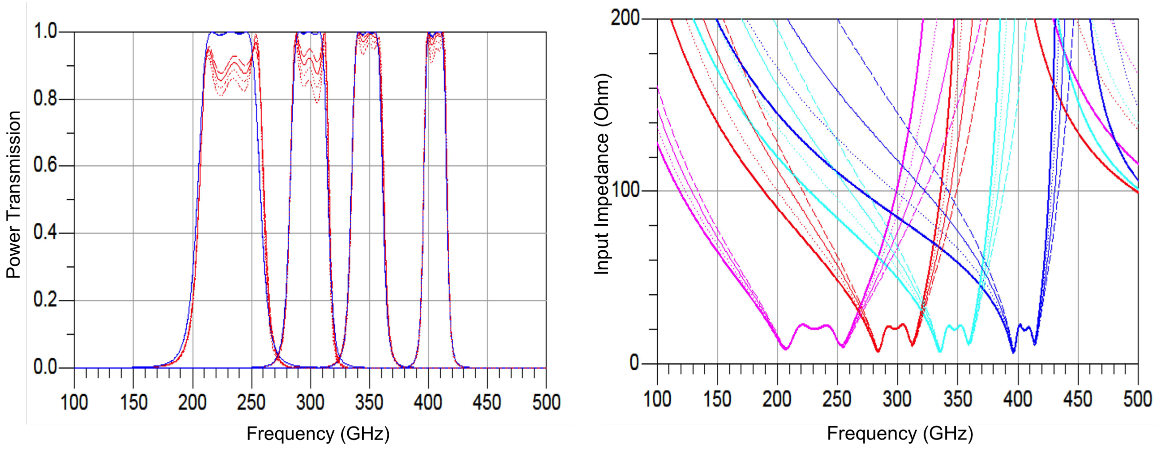

(e) S_{12}: Mline1 = 7, 10, 13 and 16 μm (f) Z_{in}: Mline1 = 7, 10, 13 and 16 μm

Figure 3.18: Study of the frequency response (S_{12}) and input impedance (Z_{in}) for various microstrip-line lengths in the BPF network: (a), (c), and (e) the vertical axis represents the power transmission, and the horizontal axis represents the frequency (GHz); the blue solid line represents the four filters considered independently, and the red dotted line is the microstrip line length given in the caption; (b), (d), and (f) the vertical axis represents the input impedance (Ω), and the horizontal axis represents the frequency (GHz).

As we can observe from subplots (c) and (d) in Figure 3.18, BPF1 (the pink line in the plot) has an input impedance of approximately 25 Ω in the band 1 passband frequency from 210 GHz to 255 GHz. However, the other three bands (the red, cyan, and blue lines) also have nonnegligible impedances at 210–255 GHz, which will degrade the performance of BPF1 and the other bands. For the passband range of BPF4 (approximately 405 GHz), the impedances of the other three bands are higher than that of BPF1. Thus, BPF4 experiences less degradation than BPF1, as is evident in the simulation results presented in Figure 3.16.

On the basis of the simulation results presented in Figure 3.18, we can make the following statements:

- The feed network will not change the cut-off frequency of each filter.
- A longer feed network transmission-line length will yield worse passband performance.

We also observe the following during the simulation and optimization of the feed network in ADS:

- The impedance between the bands does not depend on each BPF's first inductance, first capacitance, internal transmission line, port impedance, or output transmission line.
- The input impedance for a BPF network application depends on the input transmission-line length (Mline0 and Mline1 in Figure 3.14). Thus, although each filter alone is designed to have high out-of-band impedance, this impedance is decreased by the microstrip between the two filters.
- Increasing the transmission-line length inside each filter will slightly change its bandwidth.

This mutual impedance loading effect can be minimized by shortening the feed-network transmission-line length between filters as much as possible. We analyzed the effects of the feed network by simulating the feed network and modeling it as a 5×5 S-parameter matrix cascaded with four 2×2 filter S-parameter matrices. We then reinserted the S-parameter matrices into the IDL program and performed the optimization for all four filters again. It is difficult to completely eliminate the feed-network effect by tuning the feed network microstrip line lengths or BPF lumped element values, instead, we minimized the loading effect between the filters through the use of the feed line length as short as possible (the feed network dimensions are summarized in Table 3.6). The final sensitivity and overlapping power fractions of the BPF network are summarized in Tables 3.3 and 3.4.

Table 3.3: Sensitivity and power overlap in each band

	NEFD_1mm (mJy/$\sqrt{\text{Hz}}$)	NEP($\times 10^{-17}$) (W/$\sqrt{\text{Hz}}$)
Band 1	239	5.6
Band 2	69	7.8
Band 3	44	8.2
Band 4	50	8.9

Table 3.4: Power overlap between different bands

Power Overlap	Band 1	Band 2	Band 3	Band 4
Band 1	1	0.289 %	0.006 %	0.00002 %
Band 2	-	1	0.071 %	0.002 %
Band 3	-	-	1	0.032 %
Band 4	-	-	-	1

3.5 Layout Simulation

A circuit model cannot account for many of the practical factors that may introduce nonnegligible errors in the simulations and predictions, including the kinetic inductance (the superconducting effect), coupling effect (between inductors and the ground plane), parasitic capacitance, and dielectric loss, among others. All of these factors were taken into account when we converted the lumped-element values into a physical layout and simulated this layout using the Sonnet software package, which we used to simulate the filter geometry and the substrate.

3.5.1 Kinetic Inductance

Nb has a superconducting transition temperature of 9.2 K. At the operating temperature of our camera wafer, 230 mK, Nb has a penetration depth of approximately 100 nm. The thin strip line (in our BPF design, the inductor is 150 nm thick and 1 μm wide) will experience changes in the surface impedance as a result of the kinetic inductance (KI = $\mu_0 \lambda_{penetration}$), which will affect its overall impedance. Thus, when we developed the capacitor and inductor layout, we considered Nb to be a thick superconducting metal, particularly for the inductor structure.

A. R. Kerr has developed a superconducting model[5] for transmission-line simulations using tools

---

[5] The developed model considers the surface impedance of the superconductor thin film as a pair of thin sheets with a predefined impedance. The impedance that is seen by a transverse electric magnetic (TEM) wave traveling along a transmission line is essentially the same as the surface impedance that is seen by an electromagnetic wave

such as Sonnet [35]. Our group experimentally studied this superconducting transmission-line model [36]. However, the model is suitable only for the simulation of straight microstrip lines. In our case, we must consider the kinetic inductance for the entire filter structure, which includes discontinuous sections inside the filter structure. A simple two-layer model of a microstrip line with feed ports on both layers cannot accurately represent discontinuities between the two ends of a two-port network. To apply the simulation model from [35], we made the following modifications:

- At the input and output ends of the BPF network, we placed an excitation port on both layers.
- For the discontinuous sections inside the BPF network, we inserted a via between the two layers and prevented current from flowing in the vertical direction.

In this manner, we obtained a two-layer-plus-via kinetic-inductance model in Sonnet (Figure 3.19).

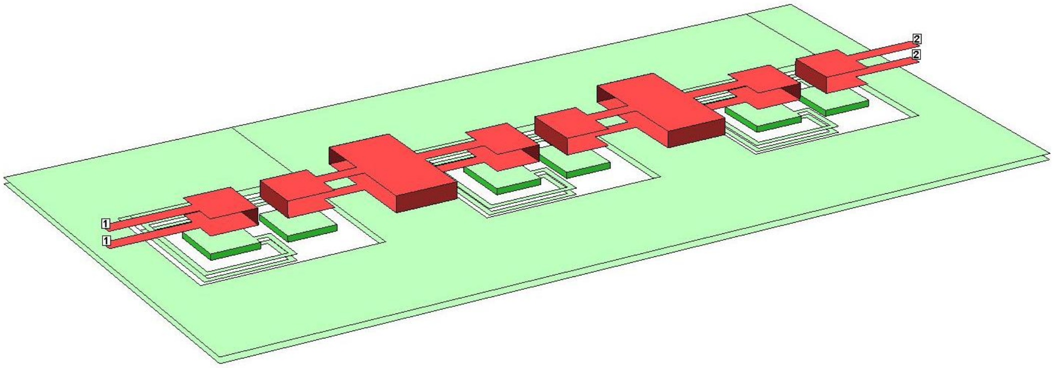

Figure 3.19: Three-dimensional (3D) structure of the simulation model used in Sonnet.

We drew the following conclusions from the simulations performed using the two-layer-plus-via kinetic inductance model:

1. There are several ways to place excitation ports and via holes/edges in Sonnet, as shown in Figure 3.20; for the excitations, we can either place a port on the lower layer (and connect the two layers with a via) or place ports on both layers. These two options proved to be identical, suggesting that ports in Sonnet with the same number are essentially connected to each other.
2. The filter response differs considerably for simulations with and without the via on the edge (the edge of the far end of the capacitors[6]). Because the potential on the surface of the conductor should be the same, we placed the via on the edge in the final model.

---

normally incident upon a conductor. For a two-layer model, the space between the two metal layers is modeled as an inductance corresponding to the magnetic-field energy stored between the sheets.

[6] We found that placing vias on all edges of the capacitor yields the exact same result as the case with only a via on one edge of the capacitor, as long as the two layers are connected. Additionally, in all of the simulations, we only use a via on the edge of the far end of the capacitors.

3. The simulation accuracy of this model can be significantly improved by considering a thick metal and kinetic inductance. Three- and five-layer models from [36] were also investigated and yielded results very similar to those of the two-layer model.

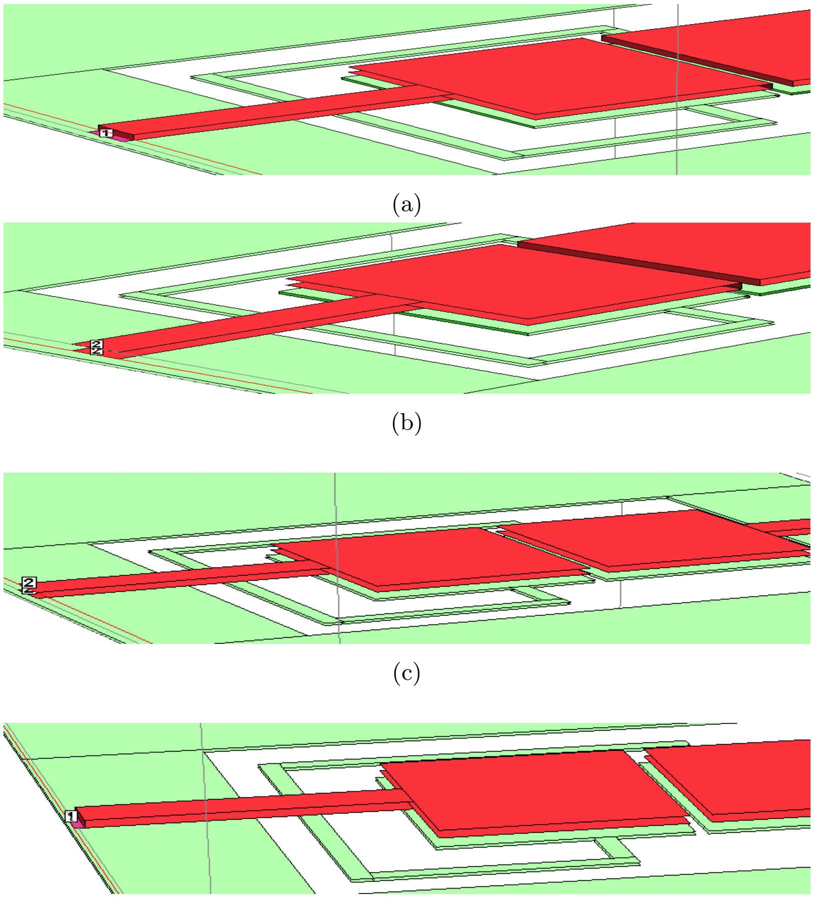

(a)

(b)

(c)

(d)

Figure 3.20: Various methods of setting up ports and edge vias in Sonnet for the superconducting BPF simulations. (a) Place a port on the lower layer and an edge via on both layers. (b) Place ports on both layers and an edge via on both layers. (c) Place ports on both layers but do not connect the layers with an edge via. (d) Place a port on the lower layer but do not connect the layers with an edge via. In Sonnet, the port is represented by a number (1 or 2 in this figure) at the end of the microstrip line. The edge via is located between two wire layers (red-colored layers) and on the edge (opposite to the microstrip line) of the rectangular capacitors.

3.5.2 Current Distribution Simulation

Layout simulation in Sonnet also provides current-distribution information on the BPF structure, which is very helpful to understand the relation between physical layout and filter simulation results. For example, on the ground plane in Figure 3.21, it is evident that the inductor exhibits a coupling effect with the edges of the ground-plane holes.

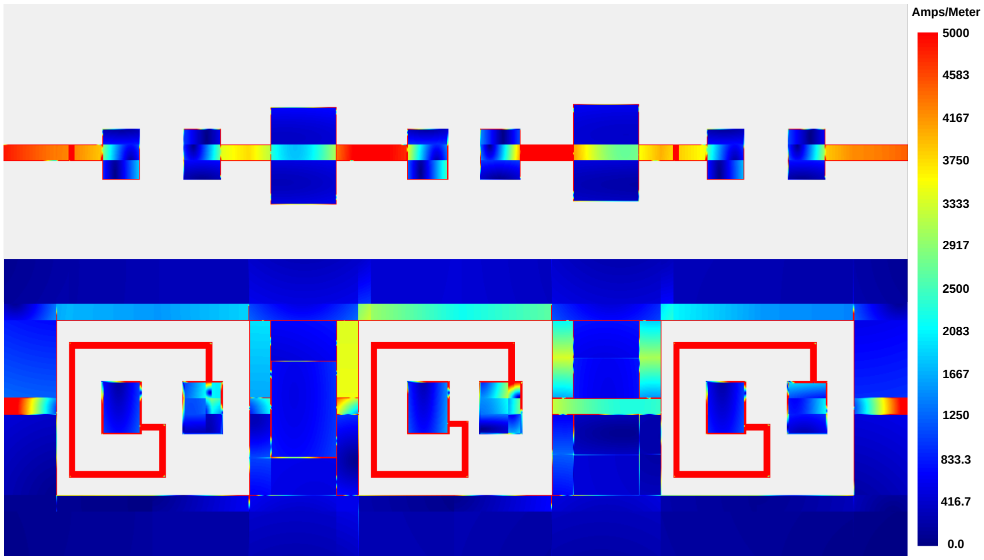

Figure 3.21: Current distributions on (top) the BPF wiring and (bottom) the ground plane in units of amps/meter.

3.5.3 Single-Capacitor and Single-Inductor Simulations

Using the kinetic-inductance simulation model of the BPFs, we compared single-capacitor and single-inductor simulations with theoretical calculations of the capacitances and inductances to test the simulation model.

In Figure 3.22, the simulation layout for a single capacitor in Sonnet is shown on the left; in the plot on the right, we compare the simulated capacitances with three different simulation models:

1. a two-layer model without a via on the edge of the capacitor and without kinetic inductance applied to the metal,
2. a two-layer model with a via on the edge of the capacitor and without kinetic inductance applied to the metal, and
3. a two-layer model with a via on the edge of the capacitor and with kinetic inductance applied to the metal.

From these simulations, we can make the following observations:

- The capacitor simulation performed in Sonnet is very similar to the parallel-plate capacitor formula, Capacitance $= \epsilon A/s$, where A is the area, and s is the separation between the capacitor plates.
- The kinetic inductance has a very minor effect on the capacitance because the structure of the capacitor is not a thin strip line.

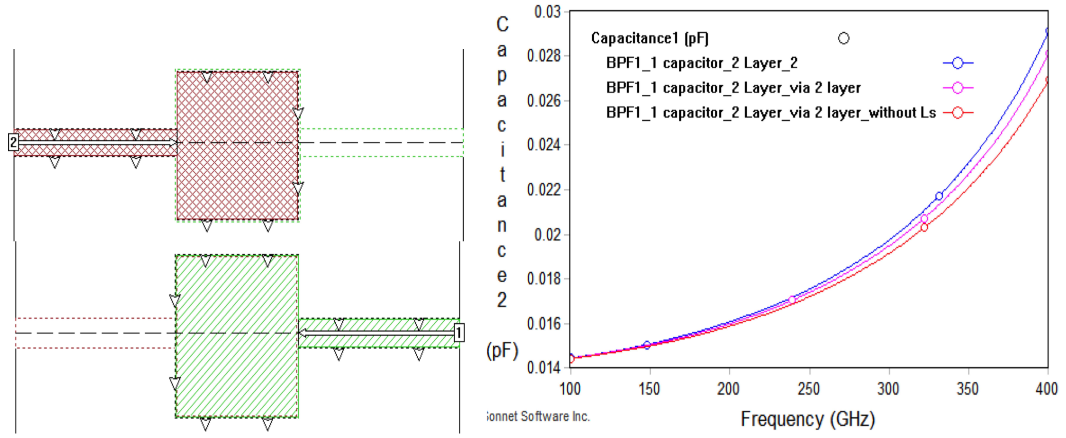

(a) (b)

Figure 3.22: Simulation of a single capacitor in Sonnet. (a) Capacitor layout in Sonnet. (b) Effects of various simulation methods on the capacitance: (blue) a two-layer model without an edge via and with kinetic inductance (L_S) added, (pink) a two-layer model with an edge via and L_S, and (red) a two-layer model with an edge via but without L_S.

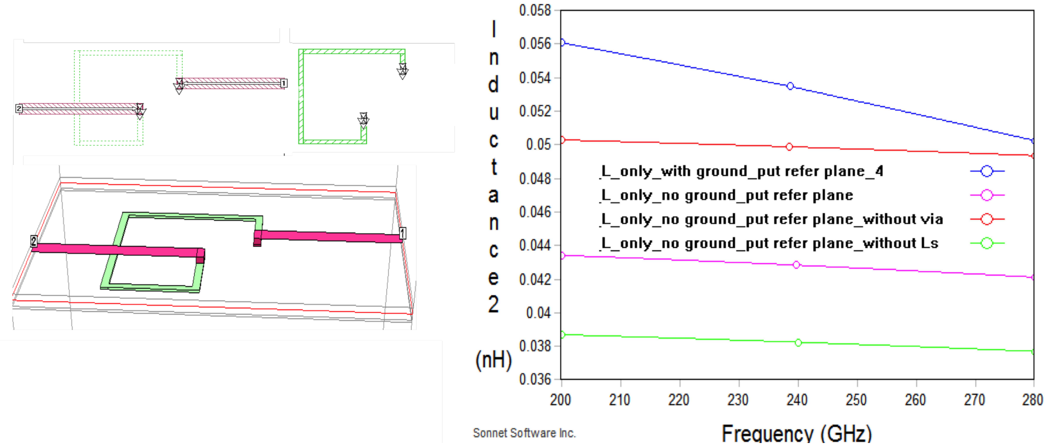

(a) (b)

Figure 3.23: Simulation of a single inductor in Sonnet. (a) Inductor layout in Sonnet: (upper left) the microstrip feed line layer; (upper right) the inductor layer; (lower) 3-D view of the simulation setup. In order to simulate the inductor, we introduced a vertical perfect conductor that connect the feed microstrip line and two layers of the inductor together, which has similar effect to edge via. (b) Effects of various simulation methods on the inductance: (blue) a simulation with a ground-plane hole surrounding the inductor, (red) a simulation without a ground plane but with an edge via and L_S, (pink) a simulation without a ground plane or edge via but with L_S, and (green) a simulation without a ground plane or L_S but with an edge via.

A similar test for an inductor is presented in Figure 3.23. We observed that the inductance calculated with kinetic inductance applied was not consistent with the planar spiral inductance formula [37]. This difference can be explained by the inclusion of the kinetic inductance and the coupling with the ground plane near the inductor (as shown in Figure 3.21). For the simulation without the inclusion of the ground plane and kinetic inductance[7], the inductance was similar to that calculated using the formula.

3.5.4 Effect of the Wafer Fabrication Process on the BPF Frequency Response

From the studies presented above, we obtained a well-defined Sonnet model with the desired lumped-element values and the desired frequency response. During the actual fabrication process, we found that it was also necessary to consider the fabrication alignment and substrate material in the simulation:

- During the fabrication process, the lower side of the parallel-plate capacitor in the BPF must be 0.2 μm larger for proper alignment. This extra plate area must be considered to obtain accurate simulation results.
- Different substrate materials (silicon or sapphire) will also have a nonnegligible effect on the filter performance. This effect is a result of the location of the lumped inductor between the dielectric and substrate layers and because the electromagnetic field generated by the square inductor is related to the dielectric parameter of the material surrounding it.

3.6 Measured Results for the BPF Network

Over the past 5 years, we have fabricated a number of camera wafers with BPFs using different dielectrics (silicon dioxide or silicon nitride), frequency bands (from band 0 to band 5, ranging from 90–450 GHz), and combinations of BPFs (two-, three-, and four-BPF networks). Figure 3.24 presents a representative measurement result (blue lines) for a four-BPF network obtained via Fourier transform spectrometry (FTS). The water-vapor transmission (red line) and predicted values (dotted black line) are also plotted for comparison. The measurement results indicate that the constructed model provides reasonable control over the filter performance.

The dimensions of each BPF component and the feed network are summarized in Table 3.5 and Table 3.6 (refer to Figure 3.12 for the inductor and capacitor definitions; refer to Figure 3.14 for the microstrip line definitions).

---

[7] When simulating the inductor alone without connecting to a capacitor, we cannot place a via on the edge of the capacitor; instead, we place vias on all edges of the inductor.

Table 3.5: Summary of each BPF inductors and capacitors dimensions in units of micrometers

	L (length of each section)	C1	C2	C3
Band 0	3.9/28.7/22.6/20/14/3.3	12.9 × 15	14.9 × 18.2	10.5 × 32.4
Band 1	2/23.3/17/12.3/2.2	9.9 × 12.8	11.4 × 13.4	10.5 × 24.8
Band 2	4.3/23/17.7/15.5/6.7/3.8	6.4 × 8.4	6.9 × 8.4	10.5 × 15.5
Band 3	6.2/16.5/15.2/13.5/12.1/5.5/3.1	6 × 5.4	5.4 × 6.4	11.5 × 15.5

Table 3.6: Summary of the BPF feed network dimensions in units of micrometers

	Mline_0	Mline_1	Mline_2	Mline_3	Mline_4	Mline_5
Band 0	17	10	3.5	6	3.5	7.5
Band 1	16.5	17	4.5	8.5	5	9.5
Band 2	16.5	13.5	11	8.5	11	8.5
Band 3	17	15.5	11.5	10.5	12	10

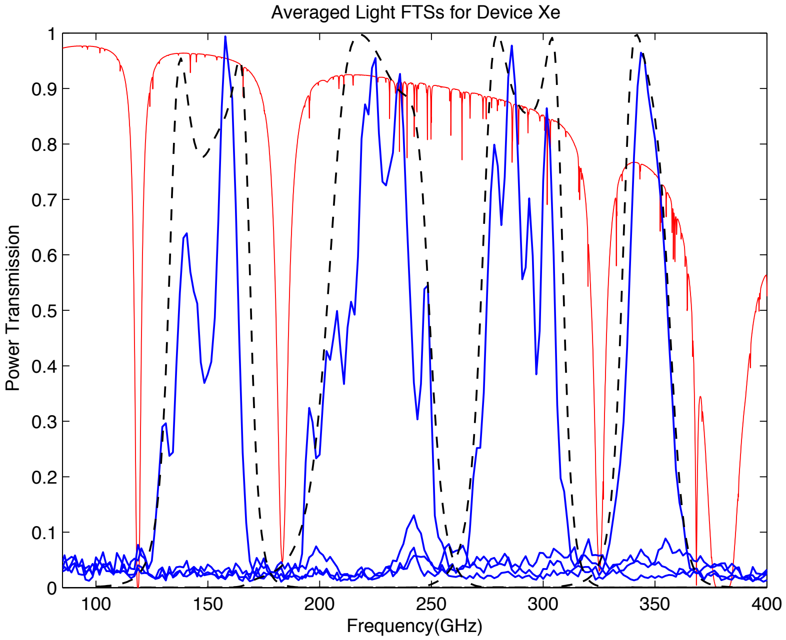

Figure 3.24: Measured and simulated BPF performance with respect to 1.5-mm water-vapor windows. Dotted line: simulation; solid blue line: measured FTS results; solid red line: water-vapor transmission.

However, we also observed variations in the measured filter response, even for the same fabrication mask design, which may be attributable to variations in the thickness of the dielectric layer across the wafer. It should be possible to resolve this issue by improving the fabrication process.

Another phenomenon observed from the measurements was fringing in the passband, which was not expected from our simulations. To identify the source of the fringing in the measured filter response, we scaled the same BPF network design to the microwave-frequency range (as shown in Figure 3.25) such that we could directly test the BPF network design alone using microwave equipment without including the optical chain in front of the focal plane, the standing wave between the antenna and the devices on the wafer, or any potential errors introduced by our FTS equipment. The measured results for the microwave version of the BPF network did not exhibit any fringing, suggesting that the BPF design itself does not introduce fringing[8].

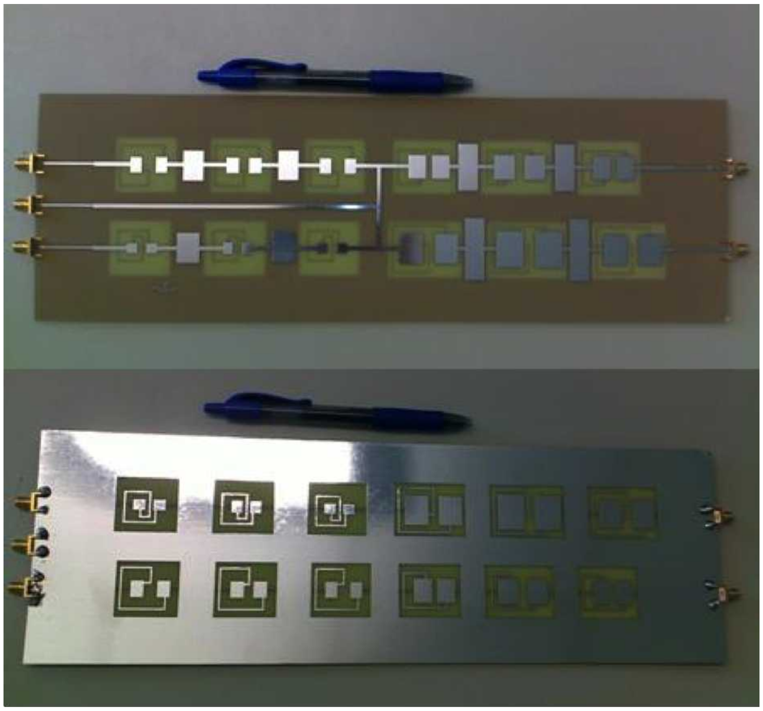

Figure 3.25: Microwave-frequency-scale model of the BPF network.

---

[8] The scale model of the BPF network was not scaled in terms of the relative dimensions between the capacitors and inductors; rather, the BPF performance was matched in terms of the impedance and resonance by modifying the dielectric material and thickness. Consequently, the simulated and measured center frequencies differ considerably, but the results clearly demonstrate that no fringing was present in the BPF frequency response, which is the purpose of this test.

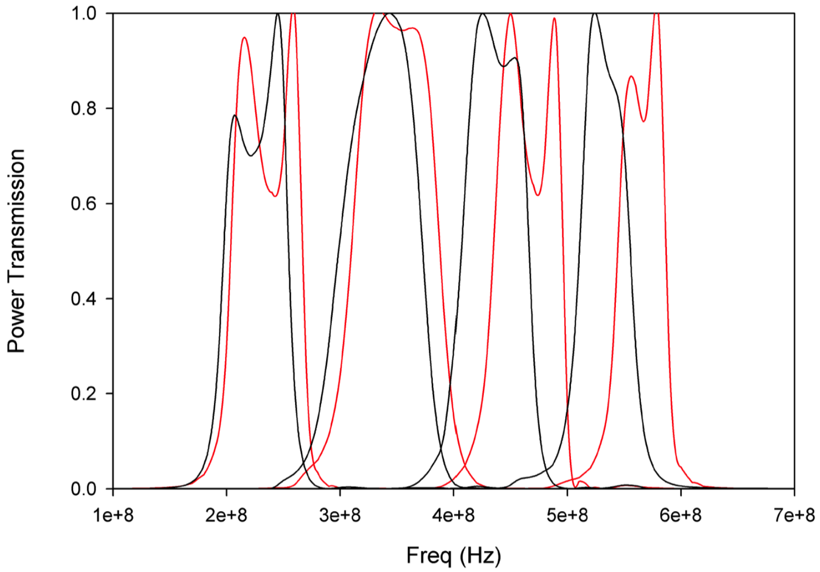

Figure 3.26: Results of the microwave-frequency-scale BPF simulations and measurements: (red line) simulation results for the microwave-frequency-scale BPF network and (black line) measurement results for the microwave-frequency-scale BPF network.

A previous measurement performed by our group using only the antenna and the detectors on the wafer revealed similar fringing in the frequency response, which also suggests that the BPF network is not the source of fringing. On the basis of the fringing frequency, we believe that the fringing originates from the optical path of the dewar rather than the standing wave on the wafer. In future work, we will study and attempt to completely eliminate the fringing effect of the camera.

Chapter 4

Two Engineering Runs of MUSIC with the Telescope

In 2010, we tested our MUSIC prototype on the sky at the CSO. At that time, we could read out only 126 resonators; additionally, the system still used many off-the-shelf components including the DC power supply, signal generator, and mixer and used only one ROACH board. As shown in Figure 4.2, we even had a computer rotating with the telescope at all times.

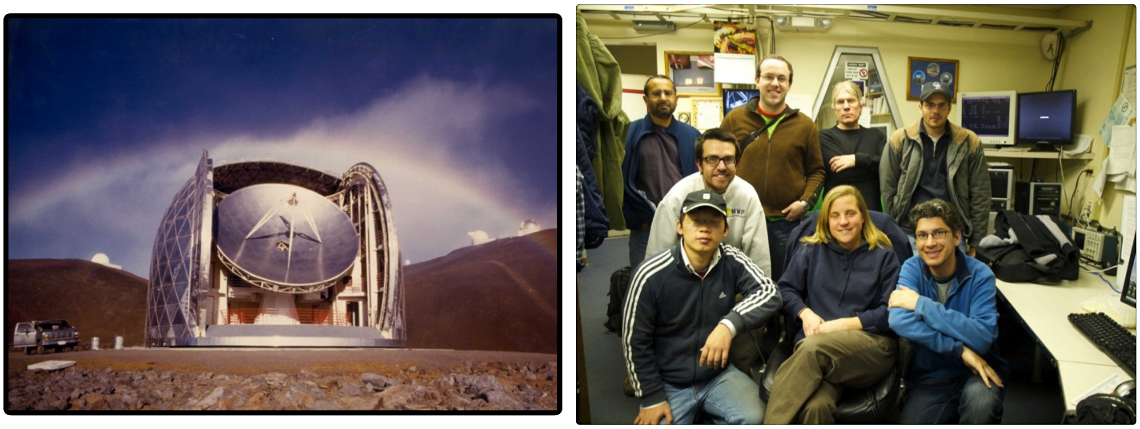

Figure 4.1: Engineering run at the CSO during Summer 2010: (left) the CSO and (right) the MUSIC team in the CSO control room.

The 2010 telescope engineering run was extremely successful: we tested our MKID wafer, the multicolor imaging, the first version of the FPGA firmware, and most of the DAQ system using the telescope.

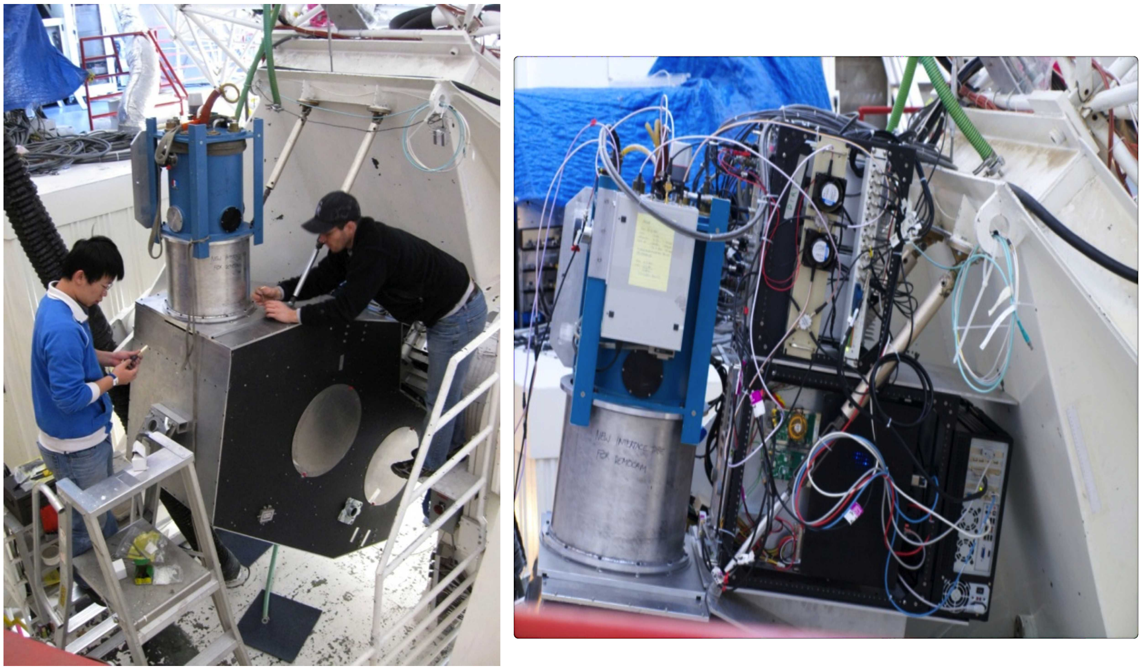

Figure 4.2: Engineering run at the CSO during Summer 2010: (left) dewar and optical box installation and (right) electronics and dewar mounted at the CSO.

After the first engineering run, we finalized the MUSIC wafer design and completed the final version of the readout system. In 2012, we performed a second engineering run at the CSO. In addition to a new wafer with four-BPF networks and the associated electronics, we implemented a new optical box and cryostat. Figure 4.3 shows the installation of the new optical box, new dewar, and new electronics crate, and Figure 4.4 shows the final electronics crate that we used. For the OSR system, we implemented two electronics crates, each containing eight OSR units, which allowed us to read out more than 3000 carrier tones. The readout electronics demonstrated their suitability for large-KID-array applications. Figure 4.5 shows two subarrays of resonator maps projected onto the sky that were collected during the 2012 run.

Following the 2012 engineering run, we attempted to improve, complete, and better understand the MUSIC instrument and prepare it to serve as a permanent instrument at the CSO. Since 2013, MUSIC has provided various astronomy groups with submillimeter observations and has started to become widely used by the public. Our group will continue to develop new instruments for submillimeter and millimeter detection. The work that we performed in developing MUSIC at the CSO serves as an excellent and valuable guide for the technological challenges of future submillimeter astronomy.

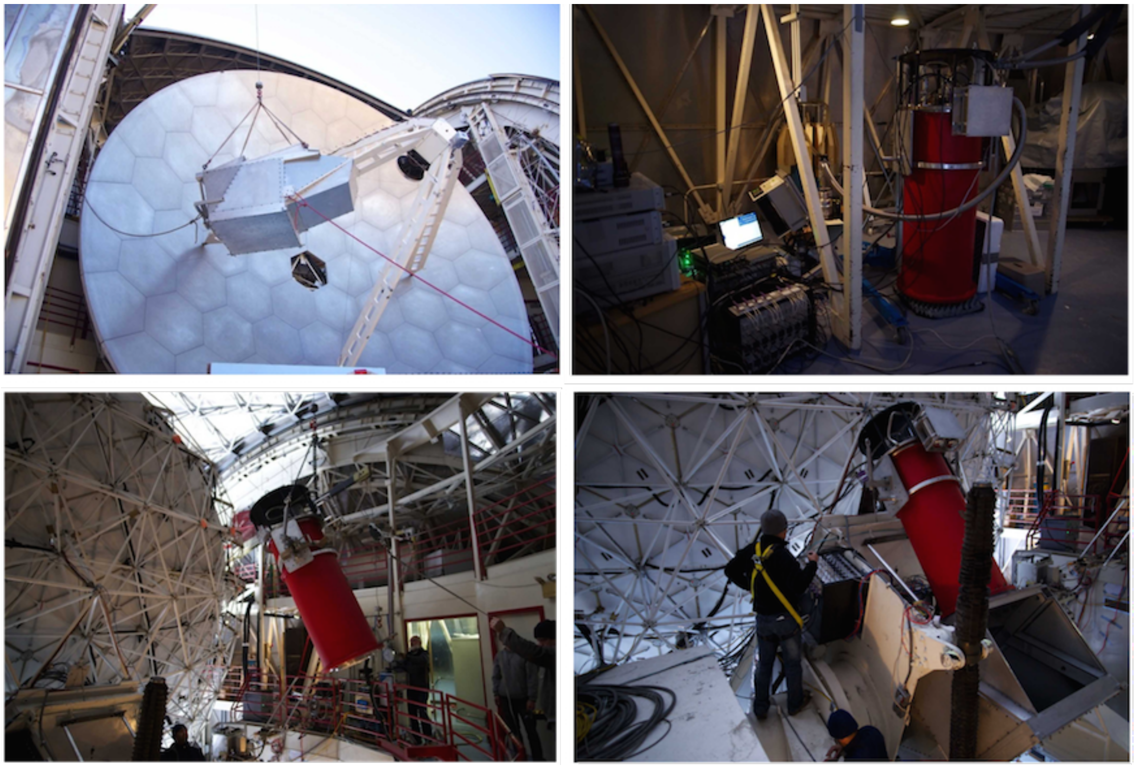

Figure 4.3: Engineering run at the CSO during Summer 2012: (upper left) lifting and installation of the new optical box, (upper right) preparation test of the electronics with the dewar on the third floor of the CSO, (lower left) lifting and installation of the new dewar, and (lower right) installation and mounting of the electronics.

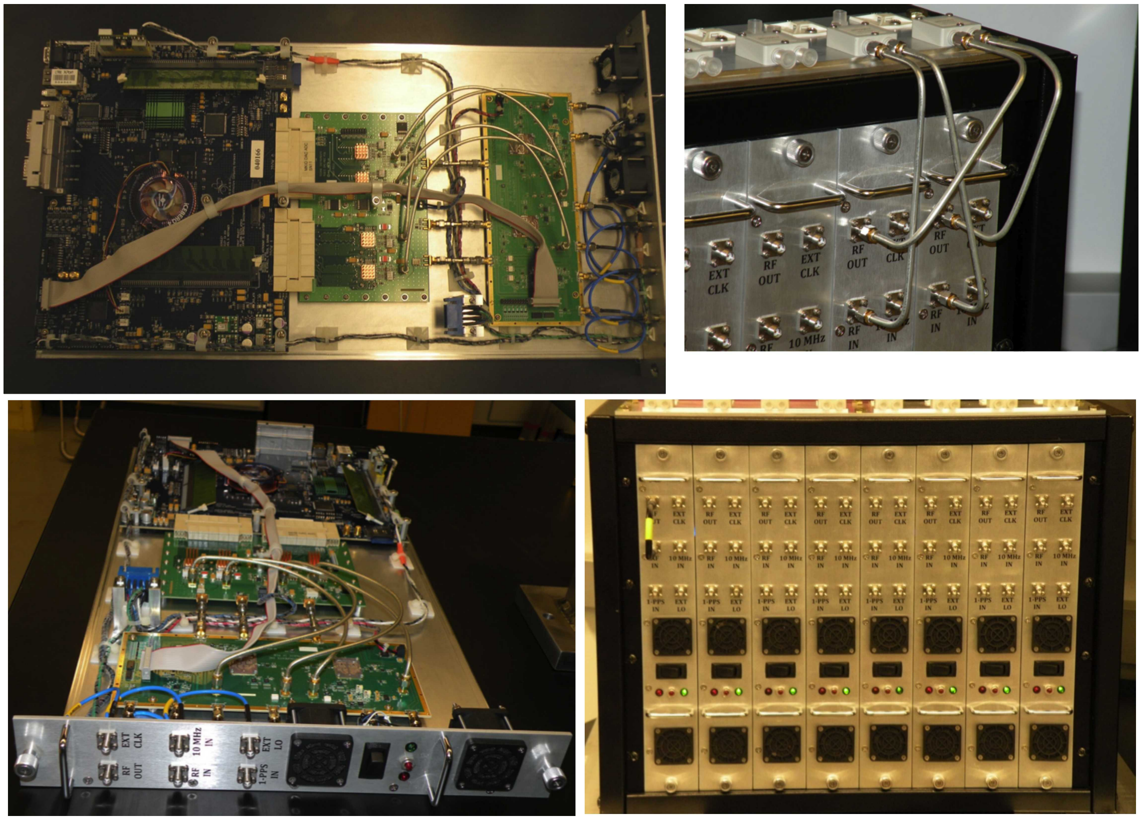

Figure 4.4: OSR system used in the engineering run at the CSO during Summer 2012: (upper left) and (lower left) one unit of OSR electronics, (upper right) the splitter connection on the electronics crate, and (lower right) one complete electronics crate with eight OSR units.

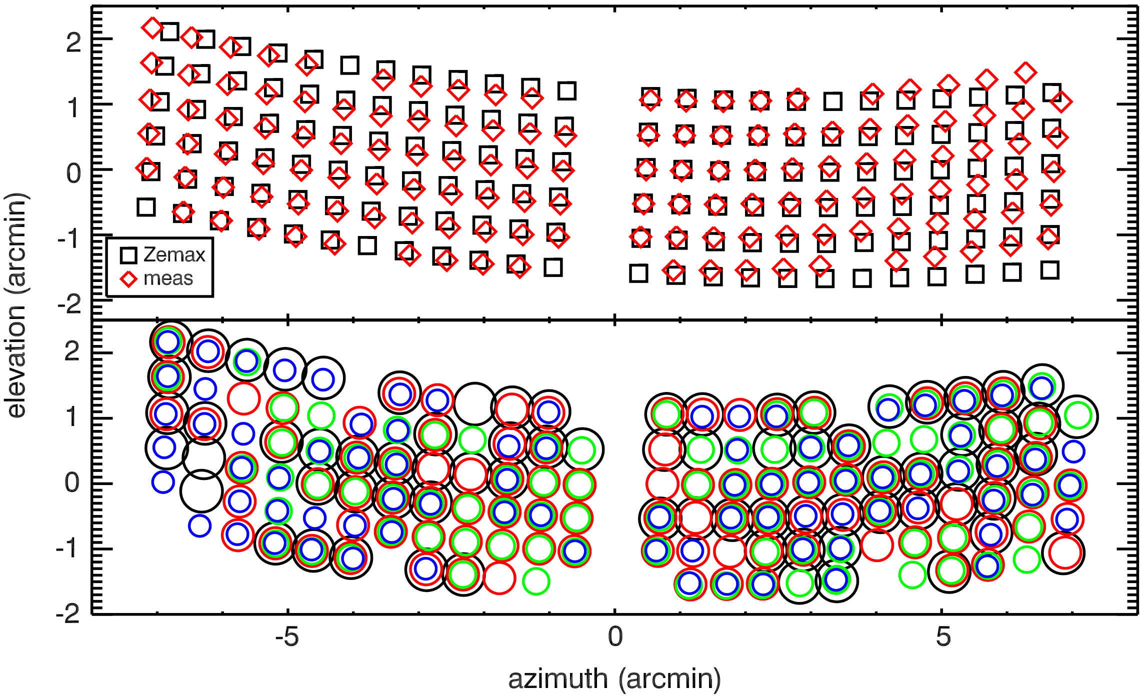

Figure 4.5: Maps projected onto the sky that were collected by the two subarrays (25% of the full focal plane) employed during the May 2012 commissioning run based on maps of bright point sources such as Uranus. The top plot compares the observed and expected positions of the beam centroids. The bottom plot shows circles centered on the reconstructed positions with diameters that indicate the reconstructed beam FWHMs and colors that identify the millimeter-wave bands. In the first analysis of these data, all four optical bands were functional in 26 of the 144 total pixels, and three of the four bands were functional in 44 pixels. Plot courtesy of the MUSIC group.

Chapter 5

Conclusions and Future Work

Overall, the MUSIC instrument including the BPF and readout electronics discussed in this dissertation represents a major step forward in the demonstration of KID technology for millimeter/submillimeter radio astronomy. The BPF bank clearly fulfilled its purpose for multicolor imaging, and the methods used for BPF bank calculation and simulation may be useful for other superconducting lithographic band-defining applications. The OSR system proved to be suitable for large-KID-array readout. The flexible FPGA-based architecture of the digital signal processing technique provides an excellent test bed for not only KID-based readout but also for applications that require fast I/O access and real-time digital signal processing. With support from the MUSIC research group, CASPER, Techne Instruments Inc., and Omnisys Inc., we have summarized the readout development work performed in our group since 2009 on an OSR platform website (www.its.caltech.edu/∼rduan/readout.html). This website includes over 400 reports and notes and addresses essentially all aspects of KID-based readout.

The receiver and backend constitute one of the most important elements of telescope construction; oftentimes, the performance of the receiver system will be the key factor or limiting factor in the overall telescope performance. As a next step, I will continue my research on readout technology for radio astronomy.

- Backend for an ultrawide-band receiver. As part of a science project in China, the Five-hundred-meter-Aperture Spherical radio Telescope (FAST) is under construction and will be completed by 2016. A low-frequency ultrawide-band receiver will play a significant role in allowing this telescope to study the evolution of the universe, discover new molecules and pulsars, and make many other significant contributions to research. We plan to design a wide-band receiver system in preparation for the FAST telescope implementation. This receiver is projected to span a range from 70 MHz to 3 GHz using a 5-GSPS 10-bit ADC or a 3-GSPS 12-bit ADC on ROACH II.
- KID-based readout for future large arrays. As the number of detectors increases for future

KID-based instruments, the design of future readout systems will focus on lowering costs and increasing the level of multiplexing. Some possible topics of study include the following:

1. Use ROACH II to explore the possibility of using one FPGA board to simultaneously cover multiple readout bandwidths. For example, we will study the possibility of using a 3-GSPS ADC to cover six 500-MHz readout bandwidths.
2. Test the low-cost mini-ROACH (a simplified version of ROACH manufactured by Techne Instruments) technology for its capability to satisfy detector readout requirements.
3. Development of customized FPGA boards for KID-based readout.

- COsmology and Molecule Explorer (COME). A small lunar satellite spanning a range of 40 to 400 MHz, preferably with a superconducting amplifier, high-sampling-resolution ADC, space-grade FPGA, and reliable channelization method. This prototype study will test the stability of FPGA-based receivers and their potential for space applications.

5.1 Acknowledgments

The projects in this dissertation were supported by NSF grant AST-0705157 provided to the University of Colorado, a NASA grant provided to the California Institute of Technology, the Gordon and Betty Moore Foundation, and the JPL Research and Technology Development Fund.

Appendix A

OSR System Digital Noise Analysis

In this chapter, starting from time-domain data streams, I will discuss the digital noise of the OSR system in detail to explain the following[1]:

1. what signals we output from the DAC,
2. how the resonator signal appears in the digitized ADC signal stream, and
3. how the carrier amplitude and noise propagate through the various calculations done by the FPGA.

A.1 IQ Conversion and Digitized Signal

Let $I_D(t)$ and $Q_D(t)$ denote the signals output by the in-phase and quadrature-phase outputs (I and Q outputs, respectively) of a DAC:

$$I_D(t) = A_0 \cos(2\pi f_{BB} t) \tag{A.1}$$

$$Q_D(t) = -A_0 \sin(2\pi f_{BB} t) \tag{A.2}$$

where $f_{BB} = f_{resonator} - f_{LO}$. The IQ mixer performs the following calculation:

$$a(t) = I_D(t) \cos(2\pi f_{LO} t) + Q_D(t) \sin(2\pi f_{LO} t) \tag{A.3}$$

Using the above equations for $I_D(t)$ and $Q_D(t)$, we obtain

$$a(t) = A_0 \cos(2\pi (f_{LO} + f_{BB}) t) \tag{A.4}$$

This choice of $I_D(t)$ and $Q_D(t)$ produces this cosine at $f_{LO} + f_{BB}$. Different choices will result in different expressions for $a(t)$; some contain signals only at $f_{LO} + f_{BB}$, whereas some contain signals

---

[1] The calculations in this chapter were prepared together with Tom Downes.

at $f_{LO} - f_{BB}$[2].

Assume that the resonator, dewar, and cables have complex transmission $S_{21}(f,t)$ at the frequency f and at time t, where the time variation of $S_{21}(f,t)$ is assumed to be slow compared to $1/f$ and to the FFT cycle time in Equation A.11. The effect of $S_{21}(f,t)$ is then

$$b(t) = \left(\frac{A_0}{2}\right)[e^{2\pi i(f_{LO}+f_{BB})t}S_{21}\left(f_{LO}+f_{BB},t\right) + e^{-2\pi i(f_{LO}+f_{BB})t}{S_{21}}^{*}(f_{LO}+f_{BB},t)] \tag{A.5}$$

Now, we apply the IQ down-conversion. Additionally, there are LPFs at the outputs to eliminate the signals that will appear near $2 \times f_{LO}$,

$$I(t) =< b(t)\cos(2\pi f_{LO}t) > \tag{A.6}$$

$$Q(t) =< b(t)\sin(2\pi f_{LO}t) > \tag{A.7}$$

where <> indicates low-pass filtering. Using our formula for $b(t)$, we obtain

$$I(t) = (\frac{A_0}{2})[\mathrm{Re}\{S_{21}(f_{LO}+f_{BB},t)\}\cos(2\pi i f_{BB}t) - \mathrm{Im}\{S_{21}(f_{LO}+f_{BB},t)\}\sin(2\pi i f_{BB}t)\}] \tag{A.8}$$

$$Q(t) = -\frac{A_0}{2}[\mathrm{Re}\{S_{21}(f_{LO}+f_{BB},t)\}\sin(2\pi i f_{BB}t) + \mathrm{Im}\{S_{21}(f_{LO}+f_{BB},t)\}\cos(2\pi i f_{BB}t)\}] \tag{A.9}$$

Finally, let us calculate the complex quantity $Z = I - iQ$[3], which we will use to perform the FFT, as follows:

$$Z(t) = \tfrac{A_0}{2}S_{21}(f_{LO}+f_{BB},t)e^{2\pi i f_{BB}t} \tag{A.10}$$

By calculating the Fourier coefficient, we obtain

$$Z_{fft}(f_{BB},t) = (A_0/2)S_{21}(f_{LO}+f_{BB},t) \tag{A.11}$$

Recall that the time variation of $S_{21}(f,t)$ is assumed to be slow compared to the timescale on which the FFT is calculated, which is why the t dependence of S_{21} is not Fourier transformed. And note that there is a factor of $\frac{1}{2}$ in the amplitude introduced during IQ mixer up-converting and down-converting. In the DAC-to-ADC loop setup for the following sections, the amplitude of $I(t)$ or $Q(t)$ is still A_0.

---

[2] As discussed in Section 2.2.3, $I_D(t)$ and $Q_D(t)$ from the DAC output will range from DC to the Nyquist frequency, and the RF output of the IQ up-mixer will range from (f_{LO} − Nyquist frequency) to (f_{LO} + Nyquist frequency).

[3] The choice of $Z = I - iQ$ (rather than $Z = I + iQ$) is made such that one properly recovers $S_{21}\left(f_{LO}+f_{BB},t\right)e^{2\pi i f_{BB}t}$. If $+i$ had been used, one would have obtained a different result.

A.2 Expected Signal for Looping the DAC Directly Back to the ADC with Cable-Delay

In the readout setup, the signal from the DAC outputs will go through different hardware (e.g., IF board, dewar, and KID device) before it can be digitized by the ADC. The physical distance of this travel introduces a time delay between the generated and digitized signals, termed the “cable-delay”, t_0[4]:

$$I(t) = I_D(t - t_0) \tag{A.12}$$

$$Q(t) = Q_D(t - t_0) \tag{A.13}$$

Using the expressions provided above for $I_D(t)$ and $Q_D(t)$, we obtain

$$I(t) = A_0 \cos(2\pi f_{BB}(t - t_0)) \tag{A.14}$$

$$Q(t) = -A_0 \sin(2\pi f_{BB}(t - t_0)) \tag{A.15}$$

By comparing these expressions to the expressions for $I(t)$ and $Q(t)$ given above for a generic S_{21}, we find that this scenario is equivalent to

$$S_{21}((f_{LO} + f_{BB}), t) = e^{-2\pi i f_{BB} t_0} \tag{A.16}$$

The resulting FFT is therefore

$$Z_{fft}(f_{BB}, t_0) = A_0 e^{-2\pi i f_{BB} t_0} \tag{A.17}$$

That is, the cable delay results in a phase factor that is dependent on f_{BB} but otherwise constant in time if t_0 is constant in time.

A.3 Expected Signal-to-Noise for DAC-ADC Loopback Through the FFT

We repeat the above calculation for looping the DAC directly to the ADC, except that we now add digitizer noise:

$$I(t) = A_0 \cos(2\pi f_{BB}(t - t_0)) + i_n(t) \tag{A.18}$$

$$Q(t) = -A_0 \sin(2\pi f_{BB}(t - t_0)) + q_n(t) \tag{A.19}$$

---

[4] As described in Section 2.4.5, the DAC, ADC, and channelization are started synchronously by the firmware. Thus, the only delay present is the physical delay described here.

$i_n(t)$ and $q_n(t)$ have the same statistical properties with a variance that is determined by the digitization noise and a white spectrum. Because the ADC has an SNR of 64 dBFS and 12 bits, we have a full-scale voltage of 2^{12} steps (for a 12-bit ADC, the full-scale sine wave runs from -2^{11} to $2^{11}-1$). We then multiply by the fraction corresponding to -64 dB in units of power. The variance of $i_n(t)$ and $q_n(t)$ should be

$$\begin{aligned} &< |i_n, q_n|^2 >= \tfrac{(2^{11})^2}{2\times SNR_{ADC\_t}} \\ &= \tfrac{(2^{11})^2}{2}/10^{(\frac{64dBFS}{10})} \\ &= 0.83 \text{ bits}^2 \end{aligned} \tag{A.20}$$

$$\text{rms}[i_n, q_n] = 0.91 \text{ bits} \tag{A.21}$$

A factor of 2 is included in the denominator in Equation A.20 because the power of a sine wave is half the square of its peak amplitude.

Let us calculate what we expect to find in the frequency space. We apply an FFT to the complex timestream, $z_n = i_n - iq_n$. The variance of this quantity will be twice the variance of i_n or q_n individually:

$$< |z_n|^2 >= 2 \times \tfrac{(2^{11})^2}{2}/SNR_{ADC\_t} = 1.66 \text{ bits}^2 \tag{A.22}$$

The white noise in the FFT of this quantity can be obtained using Parseval's theorem with 2^{16} bins:

$$\begin{aligned} &< |\text{FFT}(z_n)|^2 >= \tfrac{<|z_n|^2>}{2^{16}} = \tfrac{(2^{11})^2}{2^{16}\times SNR_{ADC\_t}} \\ &= 2.6\times 10^{-5} \text{ bits}^2 \end{aligned} \tag{A.23}$$

where we have used 2^{16} bins rather than 2^{15} bins because the negative and positive frequencies are independent in the FFT of a complex timestream.

Additionally, as demonstrated above, the following should hold for the bin that contains the carrier (neglecting the noise in that bin):

$$\begin{aligned} &Z_{fft}(f_{BB}, t_0) = A_0 e^{-2\pi i f_{BB} t_0} = 2^{11} e^{-2\pi i f_{BB} t_0} \\ &\text{or } |Z_{fft}(f_{BB}, t_0)|^2 = 2^{22} \end{aligned} \tag{A.24}$$

Note that the factor of 1/2 we had to apply in Equation A.20 is not needed here because we are now considering the complex FFT obtained from $I - iQ$, which sums the signal power in the I and Q timestreams.

The SNR in the FFT spectrum is

$$\begin{aligned} SNR_{FFT} &= \tfrac{|Z_{fft}(f_{BB})|^2}{<|\mathrm{FFT}(i_n - iq_n)|^2>} \\ &= \tfrac{(2^{11})^2}{(2^{11})^2/(2^{16}\times SNR_{ADC\_t})} \\ &= 2^{16} \times SNR_{ADC\_t} \\ &= 1.7\times 10^{11} = 112 \text{ dB} \end{aligned} \tag{A.25}$$

Next, let us consider the noise in the final audio-frequency timestream. This timestream consists of the sequential values of the FFT in the bins that contain carriers, FFT cycle by FFT cycle. Since the FFT cycle frequency is 7.5 kHz, the audio-frequency timestream is sampled at this frequency. We denote this timestream by $Z_{fft}(f_{BB}, t_a)$ where f_{BB} indicates the FFT bin being considered and t_a is the time of the FFT cycle (sampled at 7.5 kHz). This timestream consists of the expected value of the FFT in that bin plus the ADC noise in that bin:

$$Z_{fft}(f_{BB}, t_a) = A_0 e^{-2\pi i f_{BB} t_0} + i_{fft}(f_{BB}, t_a) - iq_{fft}(f_{BB}, t_a) \tag{A.26}$$

where $i_{fft}(f_{BB}, t_a)$ and $q_{fft}(f_{BB}, t_a)$ are the FFTs of the noise timestreams $i_n(t)$ and $q_n(t)$ in the bin f_{BB} evaluated during the FFT cycle that takes place at t_a. As calculated in Equation A.23

$$< |i_{fft}(f_{BB}, t_a) - iq_{fft}(f_{BB}, t_a)|^2 >= \frac{(2^{11})^2}{2^{16} \times SNR_{ADC\_t}} = 2.6 \times 10^{-5} \text{ bits}^2 \tag{A.27}$$

This is white noise that is added to the carrier amplitude. With resonators, we would determine the resonator amplitude and phase components and apply a rotation to this timestream. With only the ADC, the noise in the real and imaginary parts are equal.

Let us now take the audio-frequency FFT of $Z_{fft}(f_{BB}, t_a)$ so we can see what the signal-to-noise in audio-frequency space is. We will consider the FFTs of the real and imaginary parts of $Z_{fft}(f_{BB}, t_a)$ separately first, then consider their combination. The DC bin of the real part will be

$$\text{DC bin} =< \mathrm{Re}(Z_{fft}(f_{BB}, t_a)) >= \mathrm{Re}(A_0 e^{-2\pi i f_{BB} t_0}) = A_0 \cos(2\pi f_{BB} t_0) \tag{A.28}$$

To calculate the noise on the DC bin, we use Parseval's theorem along with the known variance of the noise in the audio timestream and the fact that it is white. If we consider an audio timestream that is N_a samples long, we divide the variance $< |i_{fft}(f_{BB}, t_a) - iq_{fft}(f_{BB}, t_a)|^2 >$ by 2 because we are considering only the real part here, and then we divide by $N_a/2$ because we are considering the FFT of a real quantity and so the positive and negative frequencies are not independent and thus there are $N_a/2$ frequency bins. Therefore, the noise power in the FFT of the real part of the audio timestream is:

$$< |\mathrm{FFT}(\mathrm{Re}(Z_{fft}(f_{BB}, t_a)))|^2 >= \frac{< |i_{fft}(f_{BB}, t_a) - iq_{fft}(f_{BB}, t_a)|^2 >}{2 \times (N_a/2)} \tag{A.29}$$

The SNR of the power spectrum of the real part of the audio-frequency timestream will be

$$\begin{aligned} & SNR_{real} = \frac{[A_0 \cos(2\pi f_{BB} t_0)]^2}{<|\mathrm{FFT}(\mathrm{Re}(Z_{fft}(f_{BB}, t_a)))|^2>} \\ & = [\cos(2\pi f_{BB} t_0)]^2 \times SNR_{FFT} \times \frac{N_a}{2} \times 2 \\ & = [\cos(2\pi f_{BB} t_0)]^2 \times SNR_{ADC\_t} \times 2^{16} \times N_a \\ & = 1.7{\times}10^{11} \times [\cos(2\pi f_{BB} t_0)]^2 \times N_a \end{aligned} \tag{A.30}$$

We can handle the imaginary part of $Z_{fft}(f_{BB}, t_a)$ as we did the real part above. First the DC bin:

$$\text{DC bin} =< \mathrm{Im}(Z_{fft}(f_{BB}, t_a)) >= \mathrm{Im}(A_0 e^{-2\pi i f_{BB} t_0}) = -A_0 \sin(2\pi f_{BB} t_0) \tag{A.31}$$

The imaginary part of the non-DC bins have the same noise as the real part, and so the SNR of the imaginary part is given by:

$$\begin{aligned} & SNR_{imag} = [\sin(-2\pi f_{BB} t_0)]^2 \times SNR_{ADC\_t} \times 2^{16} \times \frac{N_a}{2} \times 2 \\ & = 1.7{\times}10^{11} \times [\sin(2\pi f_{BB} t_0)]^2 \times N_a \end{aligned} \tag{A.32}$$

In Appendix B, we will show power spectra of $\sqrt{[\mathrm{FFT}(\mathrm{Re}(Z_{fft}(f_{BB}, t_a)))]^2 + [\mathrm{FFT}(\mathrm{Im}(Z_{fft}(f_{BB}, t_a)))]^2}$. The appropriate SNR for those power spectra is given by the ratio of the quadrature sum of the signal power spectra in the the real and imaginary timestreams and the quadrature sum of the noise power spectra in the real and imaginary timestreams. This is:

$$\begin{aligned} & SNR_{total} = \frac{S_{total}}{N_{total}} \\ & = \frac{[A_0 \sin(2\pi f_{BB} t_0)]^2 + [A_0 \cos(2\pi f_{BB} t_0)]^2}{<|\mathrm{FFT}(\mathrm{Re}(Z_{fft}(f_{BB}, t_a)))|^2> + <|\mathrm{FFT}(\mathrm{Im}(Z_{fft}(f_{BB}, t_a)))|^2>} \\ & = \frac{{A_0}^2}{<|\mathrm{FFT}(Z_{fft}(f_{BB}, t_a))|^2>} \\ & = SNR_{ADC\_t} \times 2^{16} \times N_a \\ & = SNR_{FFT} \times N_a \\ & = 1.7 \times 10^{11} \times N_a \end{aligned} \tag{A.33}$$

On the basis of the above discussion, the magnitude of the amplitude of the DC bin should be $A_0 = 2^{11}$. The magnitude of the total noise (sum of noise in the real and imaginary parts) of the final timestream will be:

$$\begin{aligned} & < |\mathrm{FFT}(Z_{fft}(f_{BB}, t_a))|^2 >= \frac{(2^{11})^2}{SNR_{FFT} \times N_a} \\ & = \frac{(2^{11})^2}{SNR_{ADC\_t} \times 2^{16} \times N_a} \end{aligned} \tag{A.34}$$

A.4 Expected Signal-to-Noise for DAC-ADC Loopback Through FIR Filtering and Decimation

If we download N_a samples of the raw output of the FFT, the SNR of the audio-frequency output is, as given above in Equation A.33:

$$SNR_{audio} = SNR_{ADC\_t} \times 2^{16} \times N_a \tag{A.35}$$

where $SNR_{ADC\_t} = 10^{(64\ \mathrm{dBFS}/10)}$ is the timestream SNR of the ADC. Recall that the timestream variance of the complex audio-frequency timestream before filtering or decimation is shown in Equation A.27. Let us suppose that we decimate by a factor of N_d. For example, for a 491.52-MHz input timestream and a 2^{16}-point FFT (these parameters correspond to the final MUSIC instrument setup), the audio-frequency timestream will be at 7.5 kHz, and so we must filter at 50 Hz and then decimate by a factor of $N_d = 75$ to obtain an output timestream at 100 Hz with a Nyquist frequency of 50 Hz. We decrease the noise bandwidth by a factor of N_d and thus decrease the timestream variance by the same factor:

$$\begin{aligned}
&< |\mathrm{LPF, decimated}\ Z_{fft}(f_{BB}, t_a)|^2 > \\
&=< |\mathrm{LPF, decimated}\ i_{fft}(f_{BB}, t_a) - iq_{fft}(f_{BB}, t_a)|^2 > \\
&= \tfrac{<|\ i_{fft}(f_{BB},t_a) - iq_{fft}(f_{BB},t_a)|^2>}{N_d} \\
&= \tfrac{(2^{11})^2}{SNR_{ADC\_t} \times 2^{16} \times N_d} = 3.5{\times}10^{-7}\ \mathrm{bits}^2
\end{aligned} \tag{A.36}$$

Let us calculate the noise and SNR of the real part of the audio timestream. Assuming there are $N_{LPF} = 2^{10} = 1024$ samples of the low-pass-filtered and decimated audio timestream[5], the noise is divided by 2 because we consider only the real part and then divided by $N_{LPF}/2$ because we are considering the FFT of a real quantity and so the positive and negative frequencies are not independent and thus there are $N_{LPF}/2$ frequency bins. Therefore, the noise power spectrum of the real part of the low-pass filtered, decimated audio timestream is:

$$\begin{aligned}
&< |\mathrm{FFT(Re(LPF, decimated}\ Z_{fft}(f_{BB}, t_a)))|^2 > \\
&= \tfrac{<|\mathrm{LPF,decimated}\ Z_{fft}(f_{BB},t_a)|^2>}{2\times(\frac{N_{LPF}}{2})} \\
&= \tfrac{<|\mathrm{LPF,decimated}\ Z_{fft}(f_{BB},t_a)|^2>}{N_{LPF}} \\
&= \tfrac{(2^{11})^2}{SNR_{ADC\_t} \times 2^{16} \times N_{LPF} \times N_d} = 3.4{\times}10^{-10}\ \mathrm{bits}^2 \\
&\text{for } N_d = 75 \text{ and } N_{LPF} = 1024
\end{aligned} \tag{A.37}$$

---

[5] Note the data rate of this low-pass-filtered and decimated audio stream is 100 Hz.

The effective number of frequency bins is then $N_d \times N_{LPF}$. The SNR of the power spectrum of the real part of the audio-frequency output is then

$$\begin{aligned} SNR_{real} &= \tfrac{[A_0 \cos(2\pi f_{BB} t_0)]^2}{<|\text{FFT}(\text{Re}(\text{LPF,decimated } Z_{fft}(f_{BB},t_a)))|^2>} \\ &= [\cos(2\pi f_{BB} t_0)]^2 \times SNR_{ADC\_t} \times 2^{16} \times N_{LPF} \times N_d \\ &= 1.7 \times 10^{11} \times N_{LPF} \times N_d \times [\cos(2\pi f_{BB} t_0)]^2 \\ &= 1.3 \times 10^{16} \times [\cos(2\pi f_{BB} t_0)]^2 \end{aligned} \tag{A.38}$$

The SNR of the power spectrum of the imaginary part of the audio timestream is similarly calculated:

$$\begin{aligned} SNR_{imag} &= 1.7 \times 10^{11} \times N_{LPF} \times N_d \times [\sin(2\pi f_{BB} t_0)]^2 \\ &= 1.3 \times 10^{16} \times [\sin(2\pi f_{BB} t_0)]^2 \end{aligned} \tag{A.39}$$

Finally, let us consider the total SNR of the low-pass-filtered, decimated audio timestream. Analogous to Equation A.33, we consider the power spectrum of the timestream given by $\sqrt{|\text{FFT}(\text{Re}(\text{LPF}, \text{decimated} Z_{fft}(f_{BB}, t_a)))|^2 + |\text{FFT}(\text{Im}(\text{LPF}, \text{decimated} Z_{fft}(f_{BB}, t_a)))|^2}$. The SNR of the power spectrum of this timestream is given by the ratio of the total signal power (real and imaginary parts of the audio timestream combined) to the total noise power spectrum (real and imaginary parts combined). That is:

$$\begin{aligned} SNR_{total} &= \tfrac{S_{total}}{N_{total}} \\ &= \tfrac{[A_0 \cos(2\pi f_{BB} t_0)]^2 + [A_0 \sin(2\pi f_{BB} t_0)]^2}{<|\text{FFT}(\text{Re}(\text{LPF,decimated } Z_{fft}(f_{BB},t_a)))|^2> + <|\text{FFT}(\text{Im}(\text{LPF,decimated } Z_{fft}(f_{BB},t_a)))|^2>} \\ &= \tfrac{{A_0}^2}{<|\text{FFT}((\text{LPF,decimated } Z_{fft}(f_{BB},t_a)))|^2>} \\ &= SNR_{ADC\_t} \times 2^{16} \times N_{LPF} \times N_d \\ &= 1.7 \times 10^{11} \times N_{LPF} \times N_d \\ &= 1.3 \times 10^{16} \\ &= 161 \text{ dB} \end{aligned} \tag{A.40}$$

If we decimate without FIR filtering, we will simply preserve all noise that was present in the undecimated timestream. The noise will not be decreased by a factor of N_d as it was above. Then, the formula will be

$$\begin{aligned} &< |\text{FFT}(\text{Re}(\text{decimated } Z_{fft}(f_{BB}, t_a)))|^2 > \\ &= \tfrac{(2^{11})^2}{SNR_{ADC\_t} \times 2^{16} \times N_{LPF}} = 2.5 \times 10^{-8} \text{ bits}^2 \end{aligned} \tag{A.41}$$

$$\begin{aligned} SNR_{real} &= 1.7 \times 10^{11} \times N_{LPF} \times [\cos(2\pi f_{BB} t_0)]^2 \\ &= 1.74 \times 10^{14} \times [\cos(2\pi f_{BB} t_0)]^2 \end{aligned} \tag{A.42}$$

$$\begin{aligned} SNR_{imag} &= 1.7 \times 10^{11} \times N_{LPF} \times [\sin(2\pi f_{BB} t_0)]^2 \\ &= 1.74 \times 10^{14} \times [\sin(2\pi f_{BB} t_0)]^2 \end{aligned} \tag{A.43}$$

And the total SNR will be:

$$\begin{aligned}
&SNR_{total} = SNR_{ADC\_t} \times 2^{16} \times N_{LPF} \\
&= 1.7 \times 10^{11} \times N_{LPF} \\
&= 1.74 \times 10^{14} \\
&= 142 \text{ dB}
\end{aligned} \tag{A.44}$$

Appendix B

OSR System Analog Noise Analysis

B.1 Summary of the OSR System Components and Their Noise Specifications

The major OSR hardware components and their noise and signal specifications are summarized in Table B.1.

Table B.1: Summary of the OSR system components and noise specifications[1]

Summary of OSR System Components and Noise Specifications	
Main FPGA	Vertex-5XC5VSX95T-1FF1136
Memory	QDRII and DDRII
I/O Connectors	Ethernet and CX4
DAC Chip	DAC5681 / 16 bits / 550 MSPS
ADC Chip	ADS5486 / 12 bits / 550 MSPS
Channelization	491.52 MHz, 2^{16}-channel FFT, and 192 carriers for the final version firmware in telescope; 340.7872 MHz, 2^{17}-channel FFT, and 126 carriers for the first version firmware in Appendix B test
HEMT	Noise temperature 2 K
Readout power	30 pW for each resonator
HEMT $\frac{N}{C}$	$\left(\frac{N}{C}\right)_{PSD} = -104.9$ dB, $\left(\frac{N}{C}\right)_{f} = -120$ dB/Hz, $(\frac{N}{C})_t = -51.7$ dB
ADC SNR	64 dBFS
ADC $\frac{N}{C}$ (assume full-scale carrier)	$(\frac{N}{C})_{2^{16}Bin\_PSD} = -112$ dBFS, $(\frac{N}{C})_{2^{17}Bin\_PSD} = -115$ dBFS, $(\frac{N}{C})_{f\_491MHz} = -148$ dBFS/Hz, $(\frac{N}{C})_{f\_341MHz} = -146$ dBFS/Hz, $(\frac{N}{C})_t = -64$ dBFS
DAC SNR	75 dBFS
DAC $\frac{N}{C}$ (assume full-scale carrier)	$(\frac{N}{C})_{2^{16}Bin\_PSD} = -126$ dBFS, $(\frac{N}{C})_{2^{17}Bin\_PSD} = -129$ dBFS, $(\frac{N}{C})_{f\_491MHz} = -159$ dBFS/Hz, $(\frac{N}{C})_{f\_341MHz} = -157$ dBFS/Hz, $(\frac{N}{C})_t = -75$ dBFS

The FPGA firmware used in Appendix B was the first version of the firmware[2] which use 2^{17}-bin FFT inside FPGA without FIR filter and decimation. Referring to Equation A.25, the SNR_{FFT} for 2^{17}-bin FFT will be 115 dB. And similarly the $(\frac{N}{C})_{PSD}$ of ADC is -115 dBFS.

B.2 Various Test Configurations for the OSR System

We defined the following configurations for the OSR tests, as shown in Figure B.1:

- External RF loop: connects the ADC/DAC and IF board together and connects the IF board's

---

[1] Where the FFT parameters are not specified, 491.52 MHz, 2^{16}-channel FFT with 192 carriers should be assumed.

[2] The first version of the firmware use clock frequency of 340.7872 MHz and the audio-frequency timestream data rate is 2.6 kHz (340.7872 MHz / 2^{17}).

RF output and RF input together with attenuators and external cables;

- RF loop: without an external cable, uses the RF switches on the IF board for loopback;
- Baseband (BB) loop: without an external cable, uses the BB switches on the IF board for loopback; and
- DAC/ADC loop: loopback without passing through the IF board that uses a cable to connect the DAC outputs to the ADC inputs through Nyquist-cut-off LPFs.

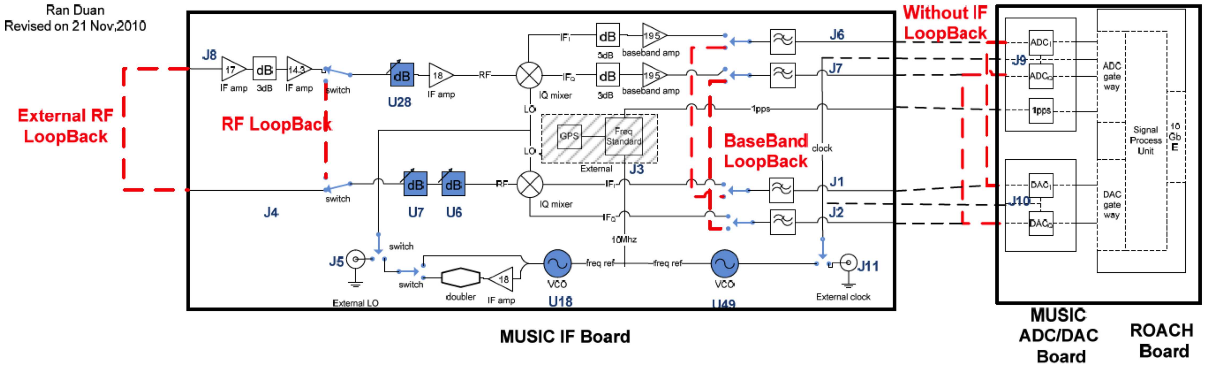

Figure B.1: Block diagram of the readout system with four different loop options (red dashed lines): external RF loopback, RF loopback, BB loopback, and DAC/ADC loopback (the LPFs for the DAC/ADC loopback are not shown in the Figure).

B.3 Noise Performance of DAC and Analog Components of OSR

We studied the noise floor using various loopback configurations and compared the theoretical calculations to the results measured using a spectrum analyzer.

The noise floor calculations were performed using the SNR Excel spreadsheet (Section 2.8.5) and under the assumption that the SNR of the DAC, assuming full-scale signals, is 75 dB (Section 2.3.2.1). From the calculation with Excel spreadsheet, the noise power ($PSD_{noise}(f) \times$ Nyquist Bandwidth) we get for BB, RF, and external RF loopback are -57.8, -57.0 and -51.3 dBm, respectively. We divide the noise power by the 6 dBm full-scale signal power of the ADC and calculate that the $(\frac{N}{C})_t$ for BB, RF, and external RF loopback are -63.8, -63.0 and -57.3 dB, respectively. Then we convert the $(\frac{N}{C})_t$ to $(\frac{N}{C})_{PSD}$ with 2^{17}-bin FFT (refer to Section 2.5.2 for the conversion). The calculated results are summarized in Table B.2.

The measured noise floor results are summarized in Table B.3. We replaced the ADC input with a spectrum analyzer in the measurement. The spectrum analyzer is set with same 2.6 kHz bin-width as in the calculation and give us noise floor in dBm per bin-width. Then we divide the noise by the full-scale signal power to get the measured power spectrum of $(\frac{N}{C})_{PSD}$ relative to full-scale carrier.

Table B.2: Calculated $(\frac{N}{C})_{PSD}$ for various loopback configurations for 2.6 kHz bin width and without ADC noise included

Loopback Setup	Calculated $(\frac{N}{C})_{PSD}$ relative to a full-scale carrier (dBFS)
BB Loopback	-114
RF Loopback	-112
Ext RF Loopback	-106

Table B.3: Measured $(\frac{N}{C})_{PSD}$ for various loopback configurations, as measured with a spectrum analyzer with 2.6 kHz bin width

Loopback Setup	Measured $(\frac{N}{C})_{PSD}$ relative to a full-scale carrier (dBFS)
BB Loopback	-114.9
RF Loopback	-114.1
Ext RF Loopback	-108.4

The results of the noise floor measurements are also presented in Figure B.2. The BB, and RF loopback each has a noise floor of approximately -114 dBFS, and the external RF loopback has a noise floor of approximately -109 dBFS. These values are in agreement with our calculations. The results of the noise tests of the various loopback configurations confirmed that the DAC and IF board of OSR system (without the ADC) performs as expected in terms of the noise floor levels.

B.4 1/f Noise Study

The 1/f noise is briefly discussed in Section 2.8.3. In this section, I will present a detailed summary of the measured results for the OSR 1/f noise and noise floor for various loops and carrier power levels.

B.4.1 1/f Noise for Different Carrier Power Levels

Below is a brief description of Figure B.3:

- Plots (a) and (b), (c) and (d), (e) and (f), and (g) and (h) correspond to the DAC/ADC loopback, BB loopback, RF loopback, and external RF loopback, respectively.
- Plots (a), (c), (e), and (g) present the power levels received for all 126 frequency bins, where the x axis represents the carrier number, and the y axis represents the power level relative to a

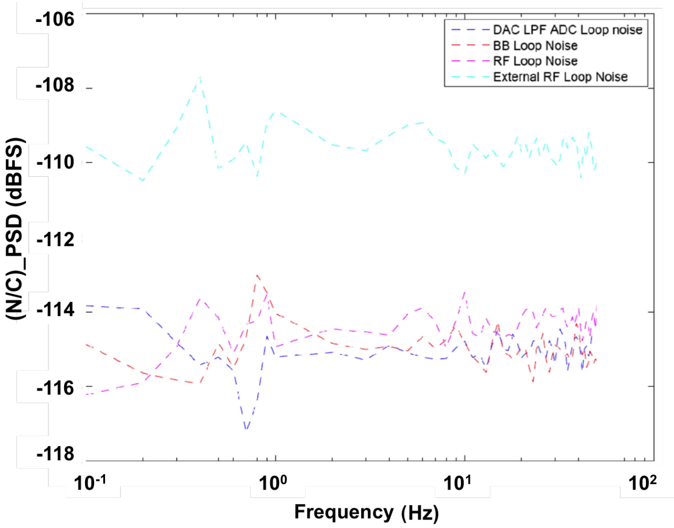

Figure B.2: Measured noise-to-carrier PSD relative to full-scale carrier for the OSR system with spectrum analyzer: (dotted blue line) DAC/ADC loopback, (dotted red line) BB loopback, (dotted pink line) RF loopback, and (dotted cyan line) external RF loopback.

full-scale carrier. The power level is provided by the $\sqrt{[mean(I)]^2 + [mean(Q)]^2}$ where I and Q are the timestreams directly output by the FPGA's FFT algorithm (effectively, Equations A.28 and A.29). Because we send out only a single frequency from the DAC, for plots of the same color, only a single tone has a power higher than that of all other bins; all 125 remaining tones correspond to off-carrier tones. The noise floors shown in plots (a), (c), (e), and (g) have the same units and scale as those shown to the right in plots (b), (e), (f), and (h) and should be around -115 dBFS.

- Plots (b), (c), (f), and (h) present the power spectrum[3] of $\sqrt{I^2 + Q^2}$, where the x axis is the frequency in hertz, and the y axis is the $(\frac{N}{C})_{PSD}$ assuming a full-scale carrier. The dashed line corresponds to the off-carrier data, representing the noise floor, whereas the solid line corresponds to the on-carrier data, representing the 1/f noise. Data presented in the same color were collected during the same measurement. These plots use the real (I) and imaginary (Q) data streams from the ROACH complex FFT outputs.

- Dividing by X corresponds to dividing the DAC LUT buffer value by X. For example, dividing by one means that the full range of the DAC is used, whereas dividing by two means that the DAC buffer is half of the full amplitude, which is 6 dB lower in power.

- The power level presented here is the received power level at the ADC relative to a full-scale carrier.

---

[3] Note that the power spectrum of $\sqrt{I^2 + Q^2}$ is the same as the sum of the power spectra of I and Q.

On the basis of the plots of the 1/f noise as a function of the carrier power level, we can conclude that the 1/f noise scales with the carrier power level. In other words, higher 1/f noise is obtained for a higher power level passing through the system. This is consistent with the expectation that 1/f noise is dominated by transfer function and gain fluctuations, which are multiplicative, rather than an additive noise component.

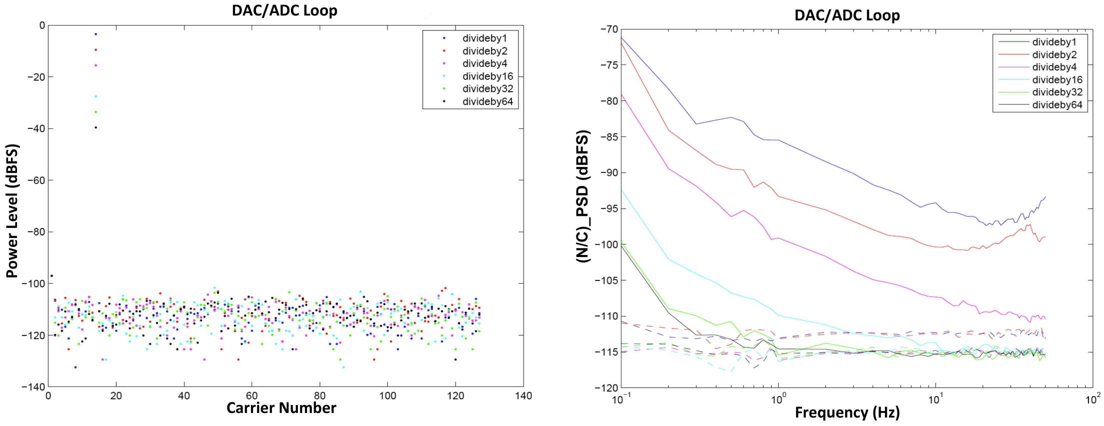

(a)

(b)

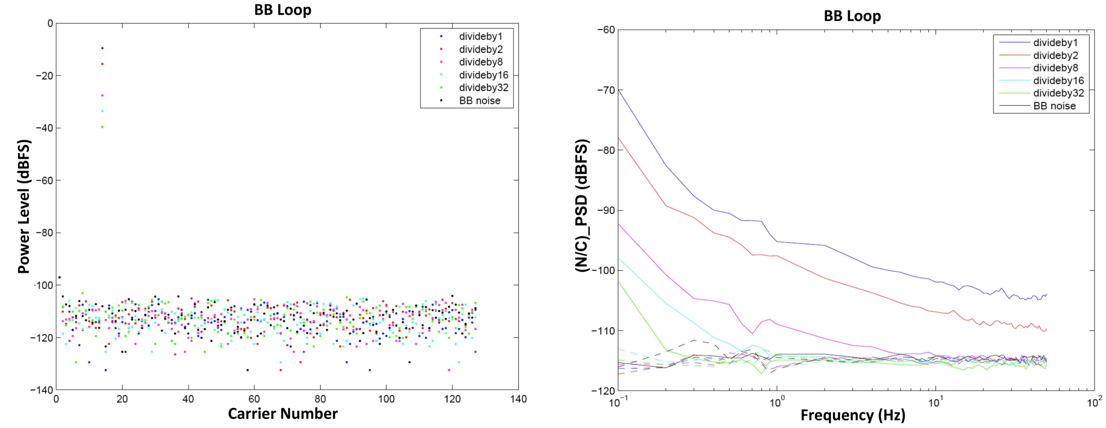

(c)

(d)

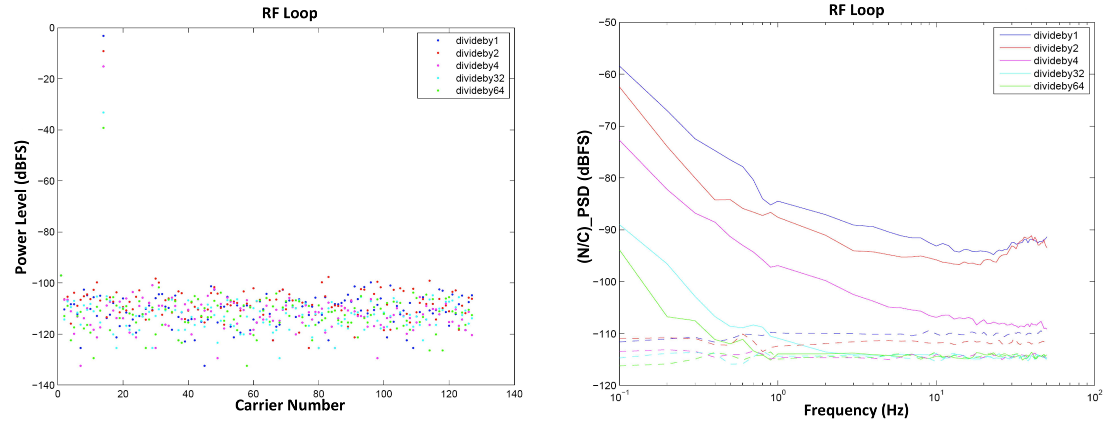

(e)

(f)

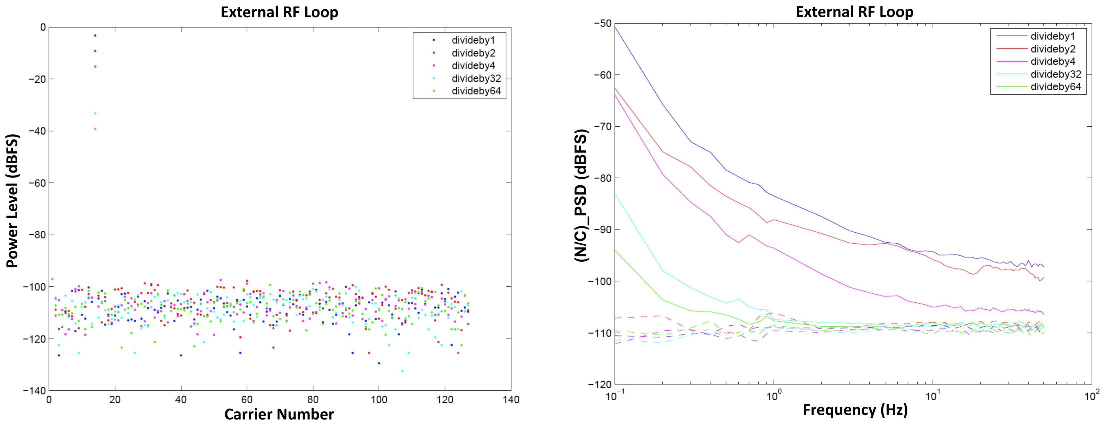

(g) (h)

Figure B.3: Study of the 1/f noise at various power levels: (a), (c), (e), and (g) power levels of the carriers relative to a full-scale carrier; (b), (d), (f), and (h) noise power spectrum $(\frac{N}{C})_{PSD}$ for carrier with power and for a representative off-carrier tone. The power level and noise power spectrum plots are made with the data from first firmware FPGA's FFT complex output: audio-frequency timestream I and Q, which is sampled at 2.6 kHz. We scale the I and Q timestreams to the full-scale carrier. For the power plot, we plot $20 \log(\mathrm{abs}(\mathrm{Complex}(\mathrm{mean}(\mathrm{I}), \mathrm{mean}(\mathrm{Q}))))$ (dBFS); for the power spectrum plot, we plot $20 \log(\mathrm{abs}(\mathrm{FFT}(\mathrm{Complex}(\mathrm{I}, \mathrm{Q}))))$, which give us the $(\frac{N}{C})$ ratio per 2.6 kHz bin width[4]: $(\frac{N}{C})_{PSD}$ (Table 2.5) (dBFS).

B.4.2 1/f Noise with Various Loopback Configurations

In Figure B.4, we rearrange the same data from Figure B.3 to compare different loopback configurations; the x and y axes have the same units as in Figure B.3. For DAC/ADC loopback, RF loopback, and external RF loopback, the same received power level also indicates the same power level sent out from the DAC. In contrast, for BB loopback (because there is a fixed 6-dB attenuation between the DAC and the ADC), it is not possible to increase the received power at the ADC to the same values achieved for the other three loops [which is why only three lines are shown in the 3-dBFS plot presented in subplot (b)]. This feature of the design also implies that, for a given received power at the ADC, the DAC output power for BB loopback is 6 dB higher than for the other configurations. Nevertheless, we observe that the DAC/ADC loopback and BB loopback configurations show the same noise, implying that the difference in DAC output power has no effect on the fractional size of the 1/f noise. Again, this is consistent with the 1/f noise being due to a multiplicative rather than additive noise.

---

[4] Refer to Equation A.33 for audio frequency timestream without FIR and decimation. In the power spectrum plot, we truncate the plot above 50 Hz because we will eventually decimate the data to the CSO-telescope pointing update rate of 100 Hz (Section 2.4.4.2).

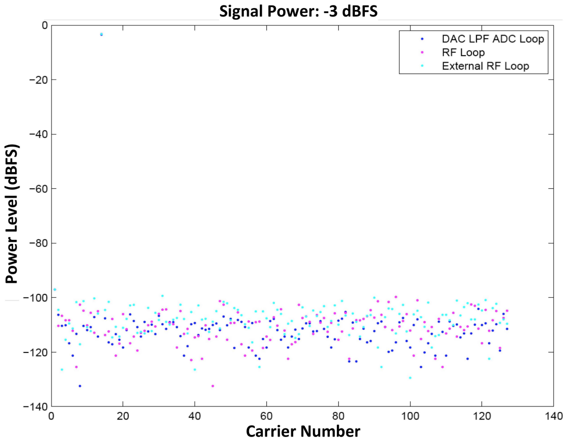
Signal Power: -3 dBFS
Power Level (dBFS)
Carrier Number
DAC LPF ADC Loop
RF Loop
External RF Loop
0
−20
−40
−60
−80
−100
−120
−140
0
20
40
60
80
100
120
140

(a)

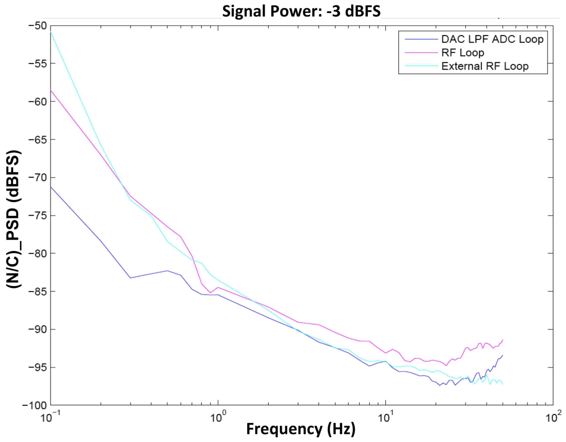
Signal Power: -3 dBFS
(N/C)_PSD (dBFS)
Frequency (Hz)
DAC LPF ADC Loop
RF Loop
External RF Loop
-50
-55
-60
-65
-70
-75
-80
-85
-90
-95
-100
10^{-1}
10^{0}
10^{1}
10^{2}

(b)

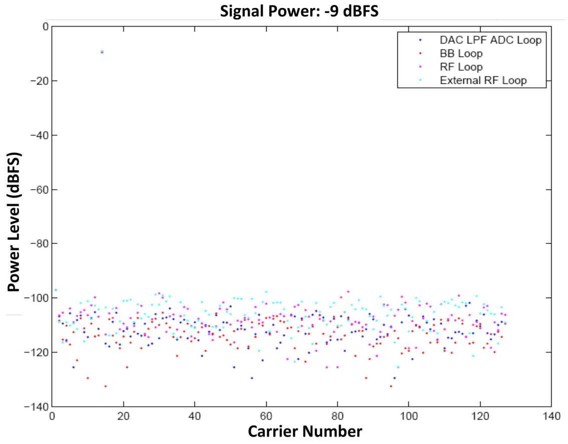
Signal Power: -9 dBFS
Power Level (dBFS)
Carrier Number
DAC LPF ADC Loop
BB Loop
RF Loop
External RF Loop
0
−20
−40
−60
−80
−100
−120
−140
0
20
40
60
80
100
120
140

(c)

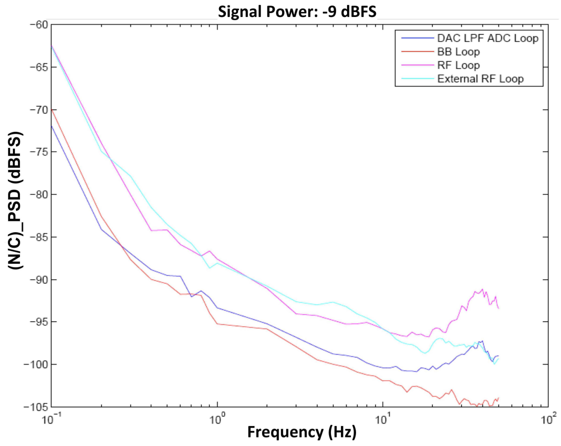
Signal Power: -9 dBFS
(N/C)_PSD (dBFS)
Frequency (Hz)
DAC LPF ADC Loop
BB Loop
RF Loop
External RF Loop
-60
-65
-70
-75
-80
-85
-90
-95
-100
-105
10^{-1}
10^{0}
10^{1}
10^{2}

(d)

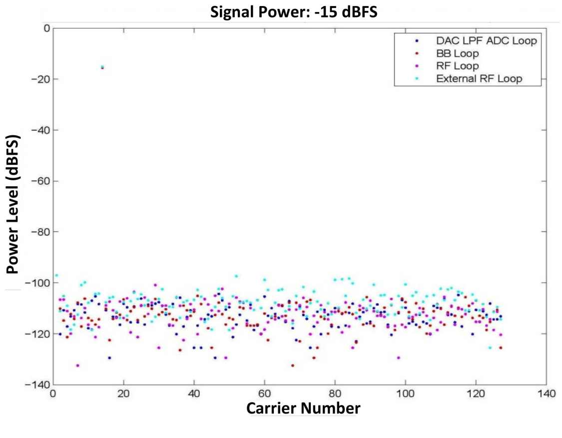
Signal Power: -15 dBFS
Power Level (dBFS)
Carrier Number
DAC LPF ADC Loop
BB Loop
RF Loop
External RF Loop
0
−20
−40
−60
−80
−100
−120
−140
0
20
40
60
80
100
120
140

(e)

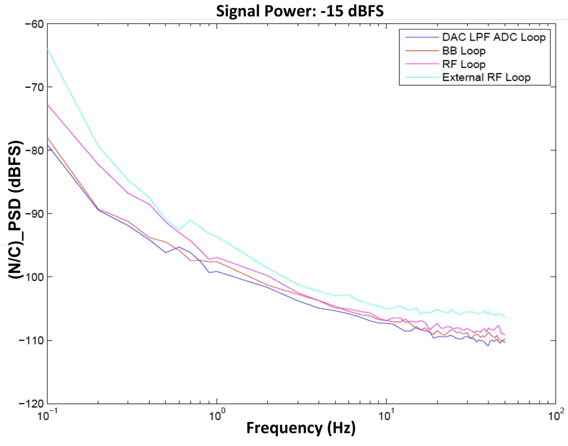
Signal Power: -15 dBFS
(N/C)_PSD (dBFS)
Frequency (Hz)
DAC LPF ADC Loop
BB Loop
RF Loop
External RF Loop
-60
-70
-80
-90
-100
-110
-120
10^{-1}
10^{0}
10^{1}
10^{2}

(f)

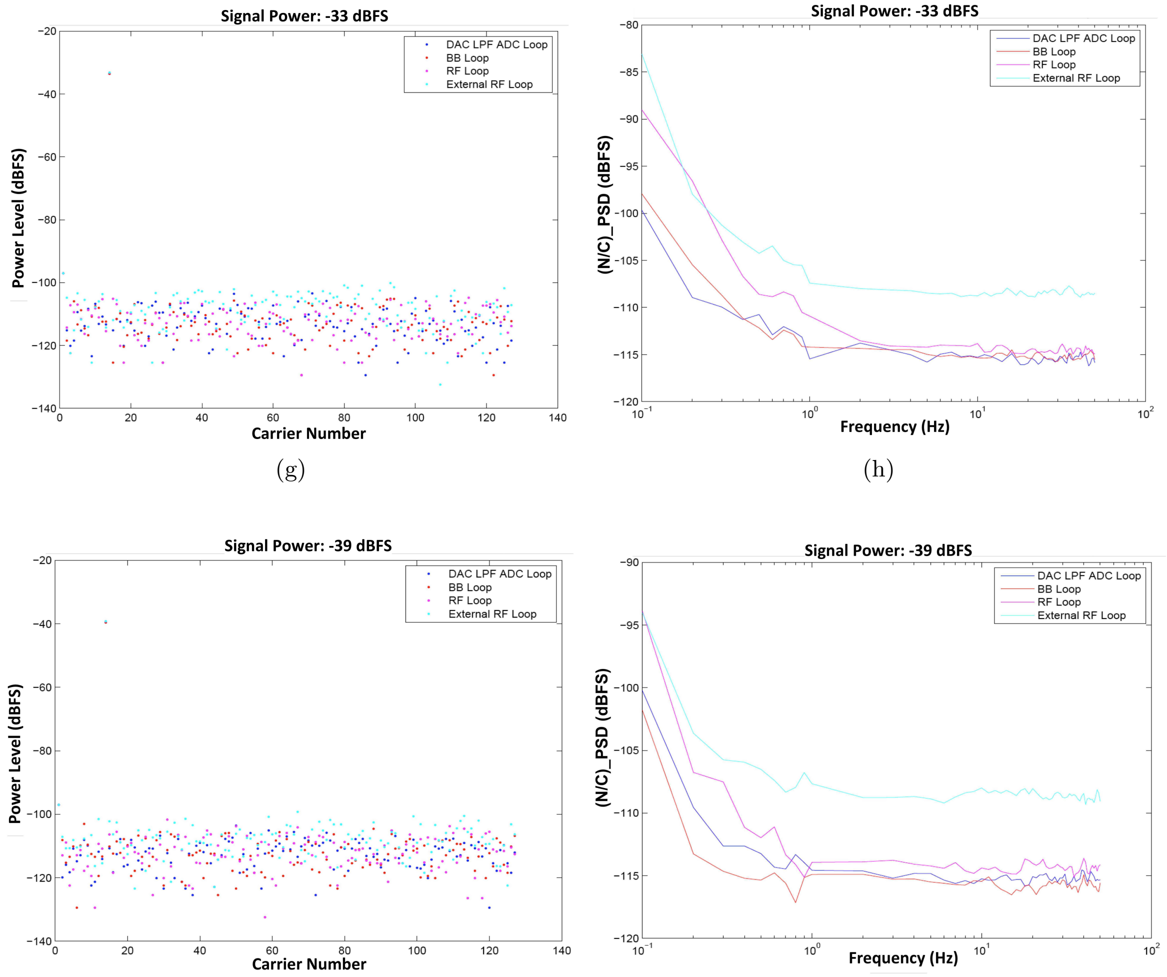

(g) (h)

(i) (j)

Figure B.4: Study of the 1/f noise for various loop configurations: (a), (c), (e), (g), and (i) power levels of the carriers relative to a full-scale carrier; (b), (d), (f), (h), and (j) noise power spectrum $(\frac{N}{C})_{PSD}$ for carrier with power (refer to Figure B.3 for definition of $(\frac{N}{C})_{PSD}$).

On the basis of the 1/f noise observed for the different loopback configurations, we concluded the following:

1. The 1/f noise performance of the BB loopback is similar to that of the DAC/ADC loopback, which indicates that the 1/f noise using the antialiasing filters on the IF board is comparable to the 1/f noise when using in-line antialiasing filters.

2. The fact that there is 1/f noise in the DAC/ADC loopback indicates that the use of a common voltage reference does not completely solve the 1/f noise problem. Relative jitter of the ADC and DAC timing could in principle be a source of residual 1/f noise, but we show in Section B.6.1 that this jitter is too small.

3. The 1/f noise performance of the RF loopback is similar to that of the external RF loopback,

which indicates that the two additional LNAs in the external RF loopback do not contribute appreciable additional 1/f noise.

4. The 1/f noise performance of the BB or DAC/ADC loopback is superior to that of the RF or external RF loopback, which indicates that the DAC/ADC 1/f noise is subdominant to other sources of 1/f noise.

5. The tests are not able to distinguish what the individual contributions of the VCO, the IQ mixers, the last RF amplifier, the baseband amplifiers, and the various attenuators are.

B.5 Gain and Phase Noise

Given the strong evidence that the 1/f noise of the OSR is dominated by multiplicative noise, we are motivated to model it in terms of the gain noise (δ_g) and the phase noise (δ_φ). We also must account for the additive noise (δ_a). Assume that the complex data stream we obtain from the OSR output is as follows:

$$a = \overline{a}(1+\delta_g)e^{i\delta_\varphi} + \delta_a \cong \overline{a}(1+\delta_g+i\delta_\varphi)+\delta_a \tag{B.1}$$

and

$$\text{Noise} = \frac{\overline{a}^*(a-\overline{a})}{|\overline{a}|^2} \cong \delta_g + i\delta_\varphi + \frac{\overline{a}^*\delta_a}{|\overline{a}|^2} \tag{B.2}$$

Then, we have

$$\text{Gain Noise} = \text{Re(Noise)} \tag{B.3}$$

and

$$\text{Phase Noise} = \text{Im(Noise)} \tag{B.4}$$

The phase and gain power spectrum for the various loopback configurations and signal power levels are summarized in Figure B.5; the phase $(\frac{N}{C})_{PSD}$ is shown on the left, and the gain $(\frac{N}{C})_{PSD}$ is shown on the right. Note from the separated gain and phase $(\frac{N}{C})_{PSD}$ plots, since we plot power spectrum of only real (corresponding to gain $(\frac{N}{C})_{PSD}$) or imaginary part (corresponding to phase $(\frac{N}{C})_{PSD}$), the noise floor is 3 dB lower than the $(\frac{N}{C})_{PSD}$ of $\sqrt{I^2+Q^2}$.

From this study of the phase and gain noise, we can conclude the following:

1. The 1/f gain noise is similar for different loops, but the phase noise of the DAC/ADC and BB loopback is much less than that of the RF and external RF loopback.

2. The 1/f gain noise (for all four loopback configurations) is of approximately the same order as the phase noise of the BB loopback.

3. For the RF and external RF loopback, the 1/f noise is dominated by the phase noise rather than the gain noise.
4. The IF board adds 1/f noise predominantly to the phase component and adds 10-15 dB less 1/f noise to the gain component.

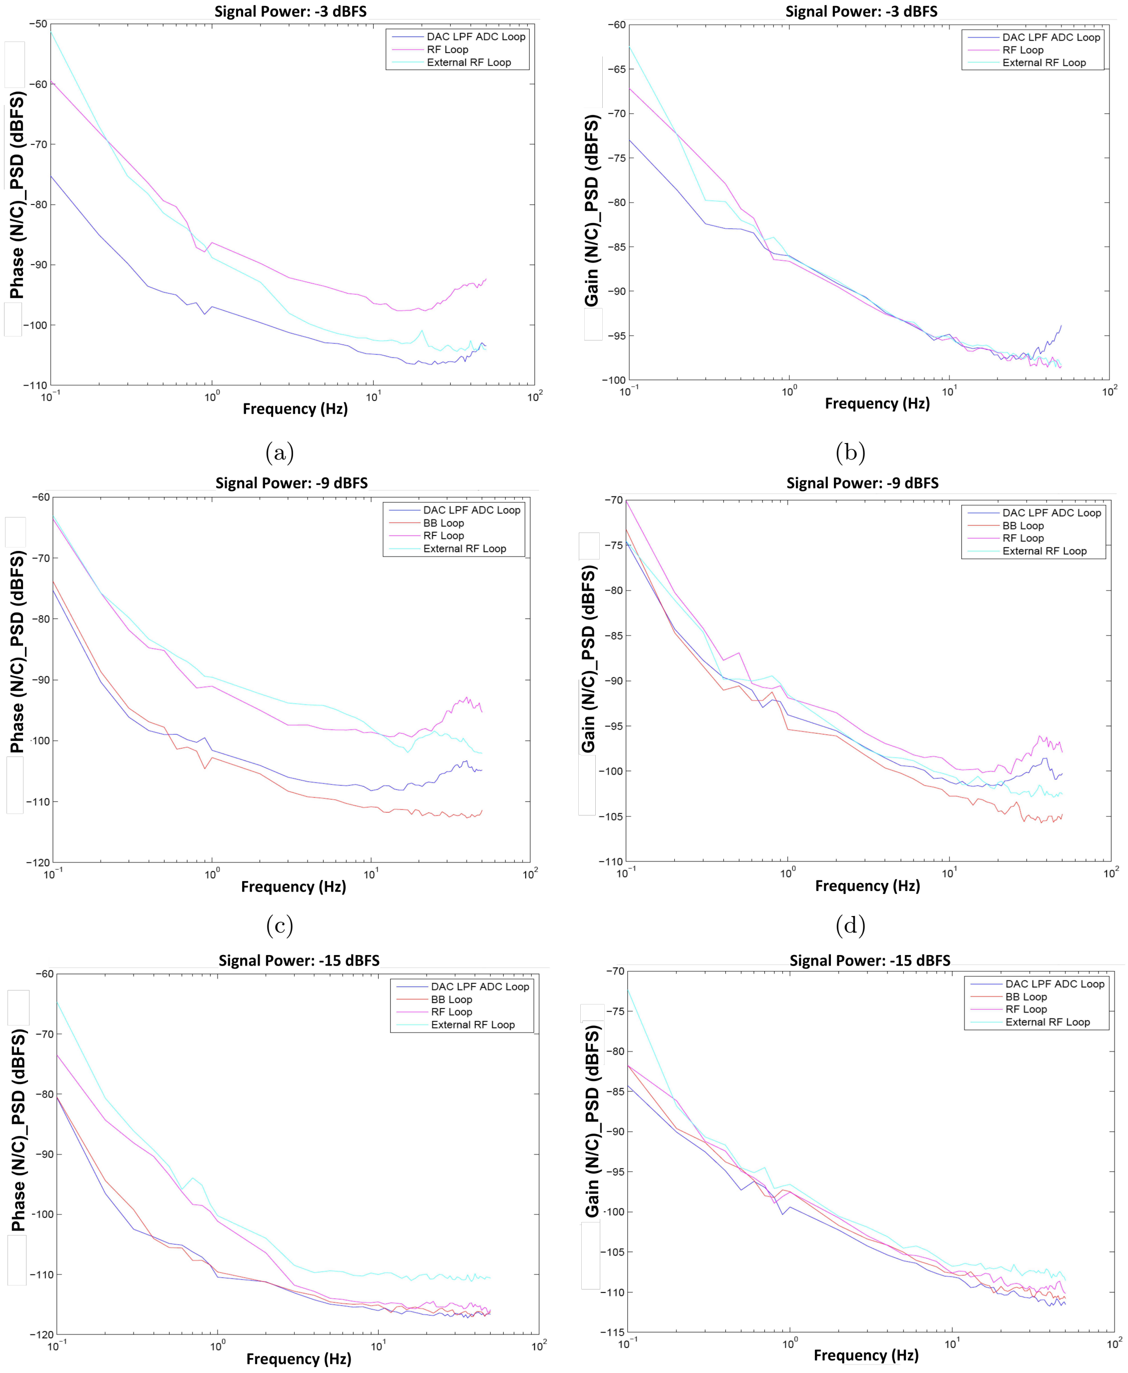

(a) (b)

(c) (d)

(e) (f)

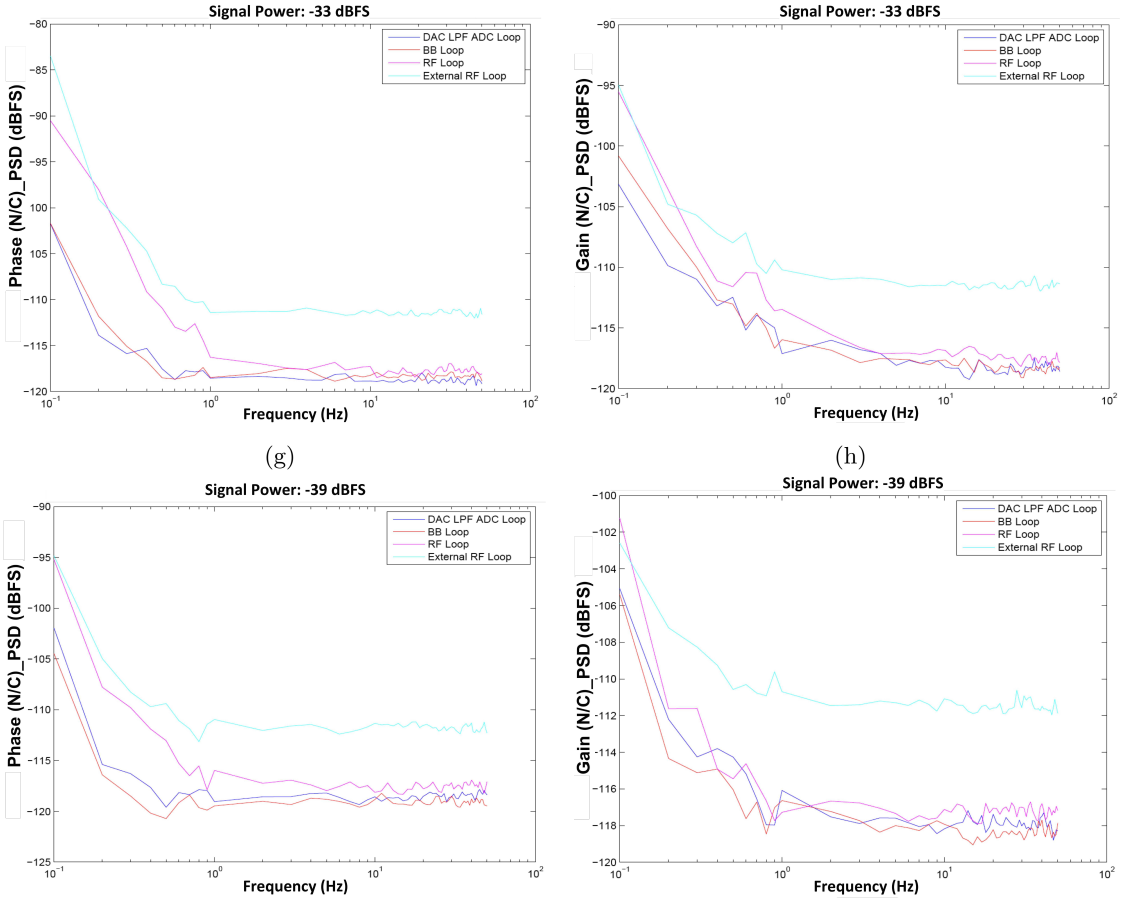

(g) (h)

(i) (j)

Figure B.5: Power spectrum of separated phase and gain directions: (a), (c), (e), (g), and (i) phase power spectrum $(\frac{N}{C})_{PSD}$ for different carrier power of various loopback configurations; (b), (d), (f), (h), and (j) gain power spectrum $(\frac{N}{C})_{PSD}$ for different carrier power of various loopback configurations (refer to Figure B.3 for definition of $(\frac{N}{C})_{PSD}$).

B.6 Understanding the Source of the Residual 1/f Noise

In the test described in the previous section, we found that the external RF loopback and RF loopback have gain noise that is similar to that of the BB and DAC/ADC loopback; in contrast, the phase noise of the former is 10–15 dB higher than that of the latter. This difference suggests that the 1/f noise of the MUSIC OSR is dominated by phase noise. In this section, we will attempt to

- identify the source of the phase noise, and
- determine the properties of the phase noise.

To fully understand the noise, we divided this study into four aspects:

1. Calculation of ADC phase noise due to clock jitter,

2. Testing of the MUSIC IF board using an external LO rather than an on-board VCO serving as the LO,
3. Use of an external synthesizer as an RF input into the IF board, and
4. Comparison of the measured phase noise for the OSR, VCO, and synthesizer.

B.6.1 ADC Phase Noise and Clock Jitter

The clock and ADC aperture jitter from the ADC can be converted into phase noise. For our ADC (ADS54RF63), the aperture jitter (rms) δt_{ADC} is 150 fs (listed on the data sheet). We know that the theoretical SNR, limited by jitter, is

$$\text{SNR} = -20 \times \log_{10}(2\pi f_{in} \delta t) \tag{B.5}$$

where f_{in} is the analog input frequency (not the clock frequency), and $\delta t = \sqrt{{\delta t_{ADC}}^2 + {\delta t_{clock}}^2}$.

For example, at an input frequency of 230 MHz, an SNR of 64.1 dB requires a maximum clock jitter of no greater than 405 fs. In our case, the same frequency source provides clock for both DAC and ADC, so δt_{clock} is common between them and should cancel out (differential clock jitter can be ignored because the cable delay between DAC and ADC is very small). The DAC chip (DAC5681) usually has much lower jitter than ADC, so the total DAC and ADC jitter $\sqrt{{\delta t_{DAC}}^2 + {\delta t_{ADC}}^2}$ will still be small and should not introduce additional noise to the system.

B.6.2 Testing of the MUSIC IF Board Using an External LO Rather than an On-Board VCO Serving as the LO

We wanted to check whether that the additional phase noise added by the IF board originated from the VCO phase noise on the IF board. We know that the phase noise of VCO chip is 10 dB better than external synthesizer at offset 10 Hz (Section B.6.4). To this end, we used an external synthesizer to generate an LO for the IF board and tested the full system using the RF and external RF loopback configurations at various signal power levels. In this way, we can see how the phase noise of the OSR system change with the amplitude of the LO's phase noise. First we plot both power spectrum of $\sqrt{I^2 + Q^2}$ and phase direction only in Figure B.6 and B.7.

From Figure B.6 and B.7, we can observe that the phase $(\frac{N}{C})_{PSD}$ is dominating the overall $(\frac{N}{C})_{PSD}$. Given this, it is sufficient to plot the total $(\frac{N}{C})_{PSD}$ rather than the phase $(\frac{N}{C})_{PSD}$. Next, we consider the power spectrum corresponding to LO frequencies of 2.2 and 4.4 GHz. Figure B.8 presents the power spectrum of $\sqrt{I^2 + Q^2}$ obtained using an external synthesizer as the LO, and Figure B.9 presents the same results obtained using an on-board VCO as the LO.

We conclude from these figures the following regarding the 1/f phase noise due to the IF board:

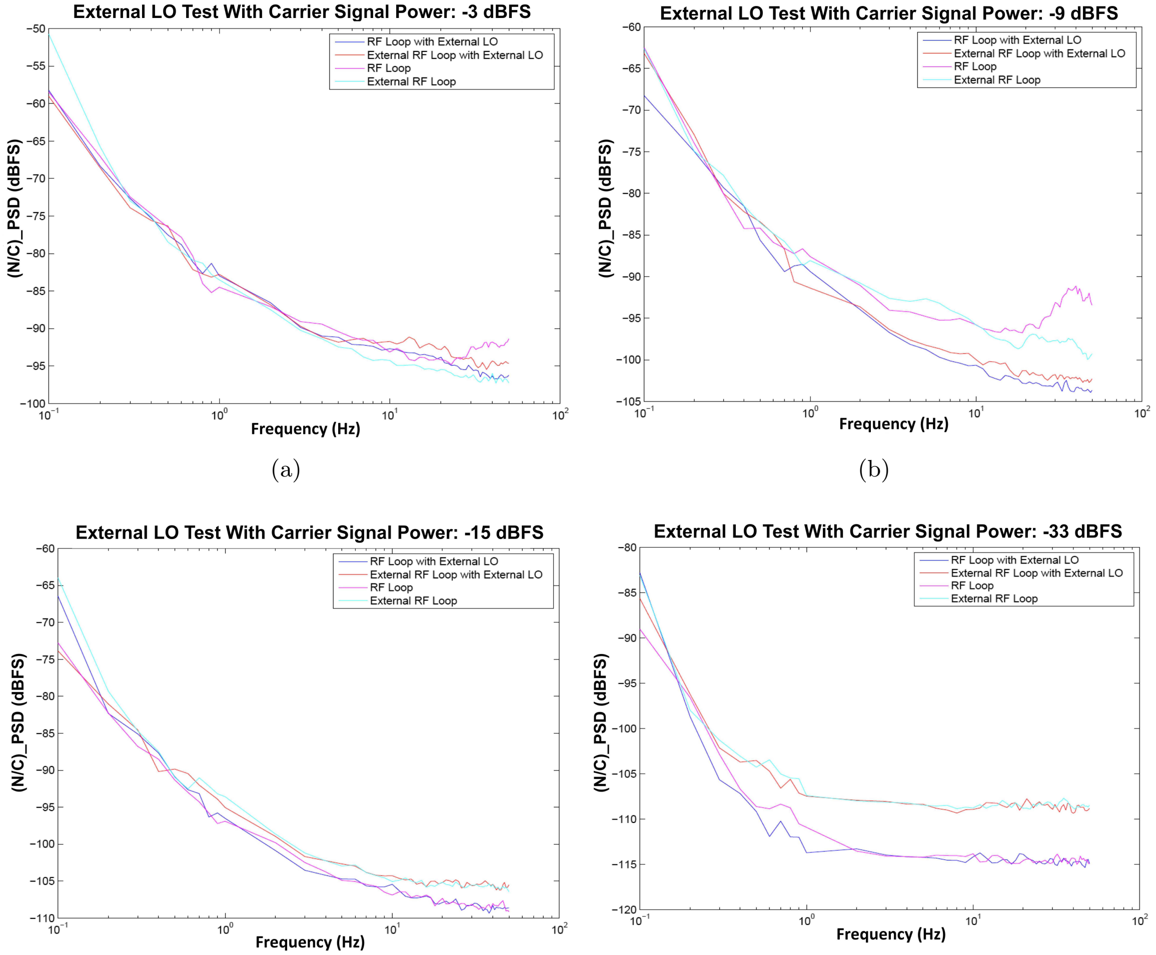

(a)

(b)

(c)

(d)

Figure B.6: Power spectrum study with an external LO: the x axis represents the frequency in hertz, and the y axis represents the $(\frac{N}{C})_{PSD}$ relative to a full-scale carrier (refer to Figure B.3 for definition of $(\frac{N}{C})_{PSD}$). Plots (a)–(d) correspond to carrier power levels of -3, -9, -15, and -33 dBFS.

- Figures B.6 to B.9 show that there is no significant difference between using an external synthesizer as the LO and using an on-board VCO as the LO. These two LO sources have a 10 dB difference in phase noise at 10 Hz offset frequency. Therefore, if imperfect cancellation of phase noise in the LO between the two IQ mixers were the dominant source of the residual 1/f noise, we would expect this phase noise to increase when we use the external LO. We do not observe such a difference, leading us to conclude that this type of imperfect cancellation is not the dominant source of residual 1/f noise.

- Changing the LO frequency (which indicates different resonator frequency in device readout) has no effect on the phase noise in the OSR.

- The LO frequency power level does not affect the phase noise as long as the power level meets

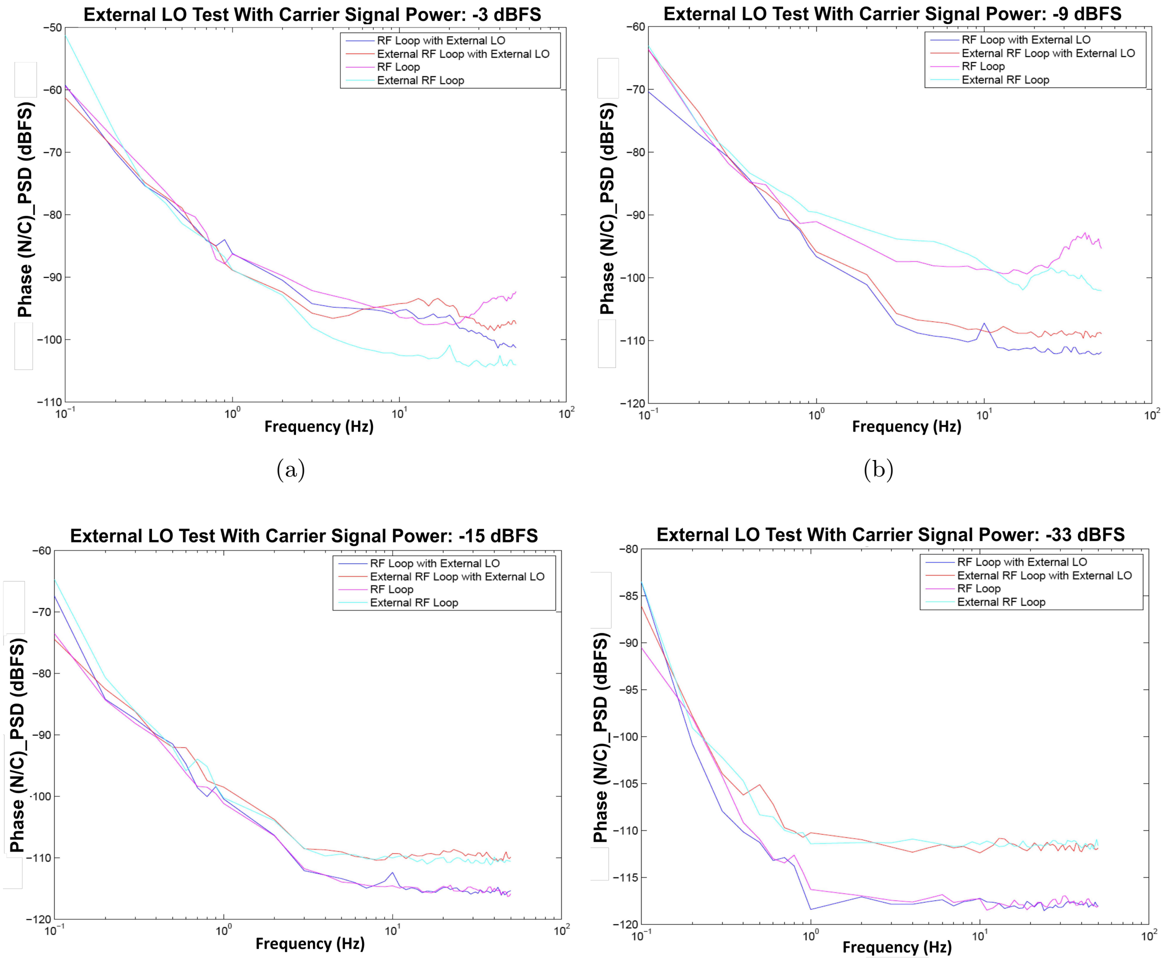

(a) (b)

(c) (d)

Figure B.7: Phase power spectrum test with an external LO: the x axis represents the frequency in hertz, and the y axis represents the phase $(\frac{N}{C})_{PSD}$ relative to a full-scale carrier (refer to Figure B.3 for definition of $(\frac{N}{C})_{PSD}$). Plots (a)–(d) correspond to carrier power levels of -3, -9, -15, and -33 dBFS.

the IQ mixer's requirement.

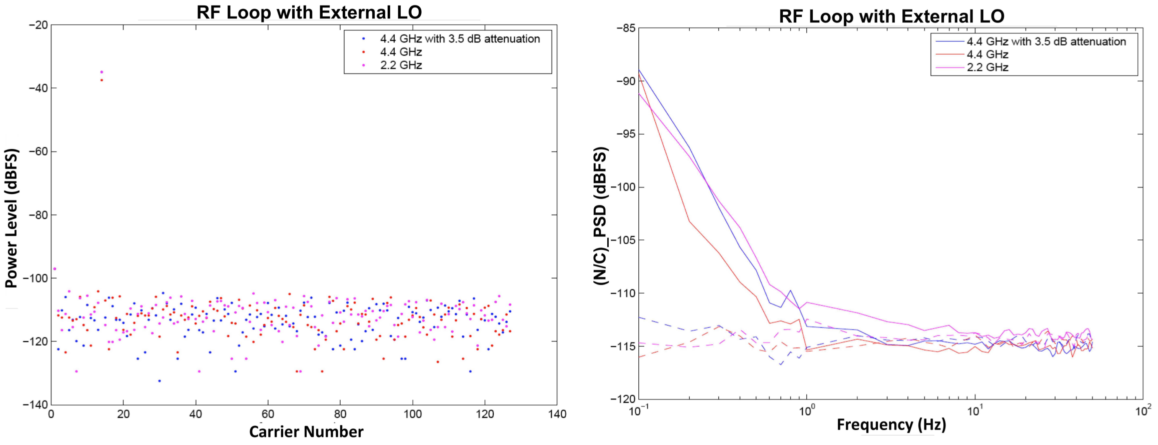

(a) (b)

Figure B.8: Comparison of the power spectrum for different frequencies of the external-synthesizer LO: the blue curve corresponds to an LO frequency of 4.4 GHz with a 3.5-dB decrease in power, the red curve corresponds to an LO frequency of 4.4 GHz, and the pink curve corresponds to an LO frequency of 2.2 GHz. (a) Power level of the received carrier tones, where the y axis represents the power level relative to a full-scale carrier, and the x axis corresponds to the carrier-tone number; and (b) power spectrum $(\frac{N}{C})_{PSD}$, where the y axis represents the $(\frac{N}{C})_{PSD}$ relative to a full-scale carrier (refer to Figure B.3 for definition of $(\frac{N}{C})_{PSD}$), and the x axis represents the frequency.

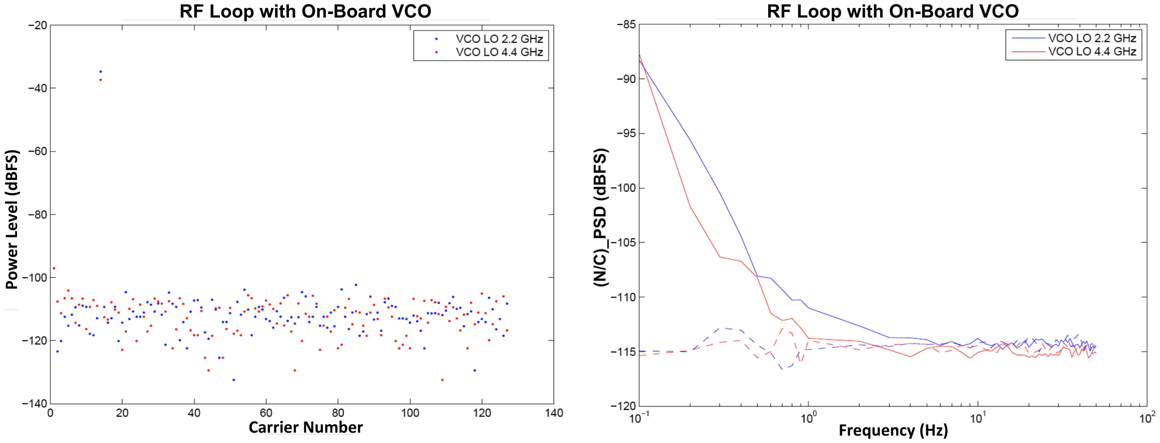

(a) (b)

Figure B.9: Comparison of the power spectrum for different frequencies of the on-IF-board VCO: the blue curve corresponds to an LO frequency of 2.2 GHz, and the red curve corresponds to an LO frequency of 4.4 GHz. (a) Power level of received carrier tones, where the y axis represents the power level relative to a full-scale carrier, and the x axis corresponds to the carrier-tone number; and (b) power spectrum of $\sqrt{I^2 + Q^2}$, where the y axis represents the $(\frac{N}{C})_{PSD}$ relative to a full-scale carrier (refer to Figure B.3 for definition of $(\frac{N}{C})_{PSD}$), and the x axis represents the frequency.

B.6.3 Use of an External Synthesizer as an RF Input to the IF Board

In this section, we present a test of the OSR system that was performed using an external signal source rather than the loopback mode of the OSR (in which the signal is generated from the DAC).

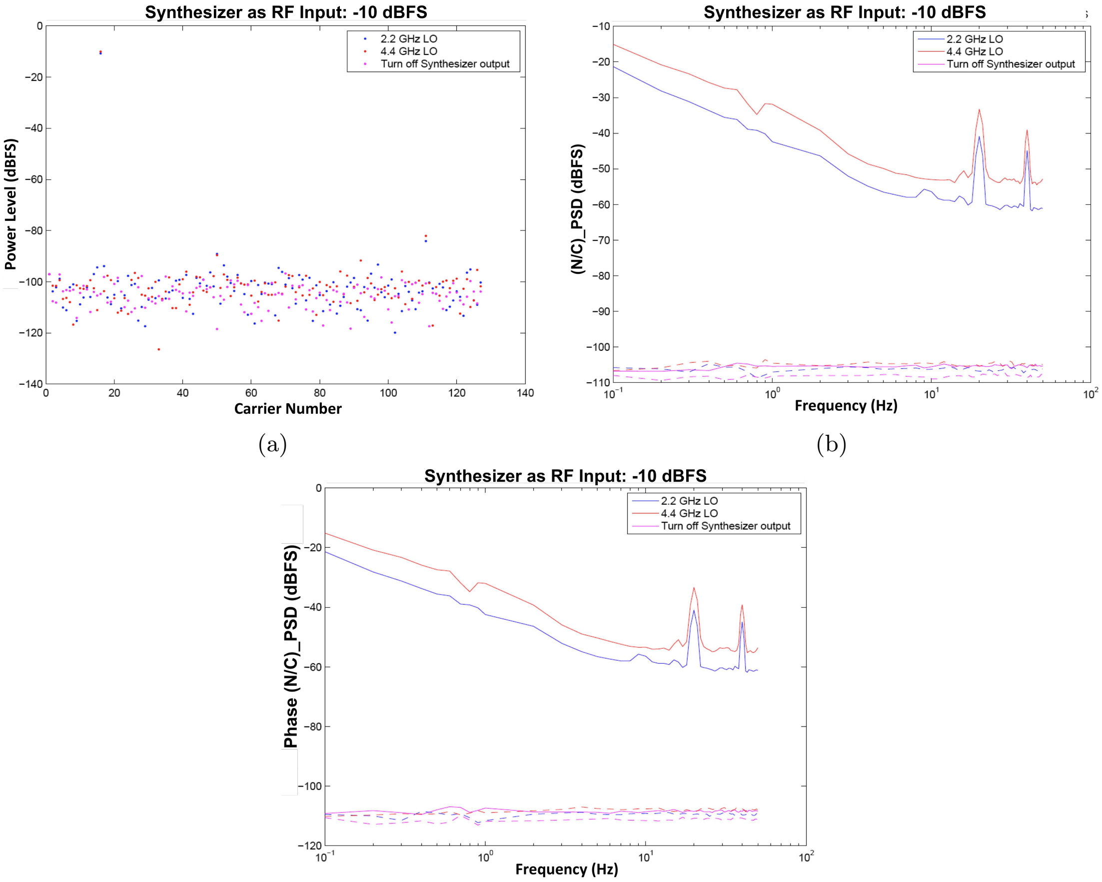

(a) (b)

(c)

Figure B.10: Power spectrum study with an external synthesizer serving as an RF input with a signal power level of -10 dBFS: (a) Power level of received carrier tones, where the y axis represents the power level relative to a full-scale carrier, and the x axis corresponds to the carrier-tone number; (b) power spectrum $(\frac{N}{C})_{PSD}$ (refer to Figure B.3 for definition of $(\frac{N}{C})_{PSD}$), where the y axis represents the $(\frac{N}{C})_{PSD}$, and the x axis represents the frequency; and (c) phase power spectrum $(\frac{N}{C})_{PSD}$, where the y axis represents the phase $(\frac{N}{C})_{PSD}$, and the x axis represents the frequency. In sub-plot (b) and (c), the dashed line corresponds to a representative off-carrier data, whereas the solid line corresponds to the on-carrier data. The blue and red curve correspond to an LO frequency of 2.2 GHz and 4.4 GHz respectively, and the pink curve represents the case in which the external signal to the IF board is turned off.

We used a synthesizer to generate a signal tone and sent it to the IF board. Figure B.10 and B.11 show the noise power spectrum with an external synthesizer serving as an RF input for an ADC received signal level of -10 dBFS and -17 dBFS, respectively.

We drew the following conclusions from this test:

1. The phase noise observed when an external synthesizer is used as the RF input and the on-board VCO is used as the LO is considerably higher than the phase noise observed when either a synthesizer or a VCO is used as both the up- and down-mixer LO, even when the VCO and the RF-input external synthesizer are both connected to the same 10-MHz reference. Because:

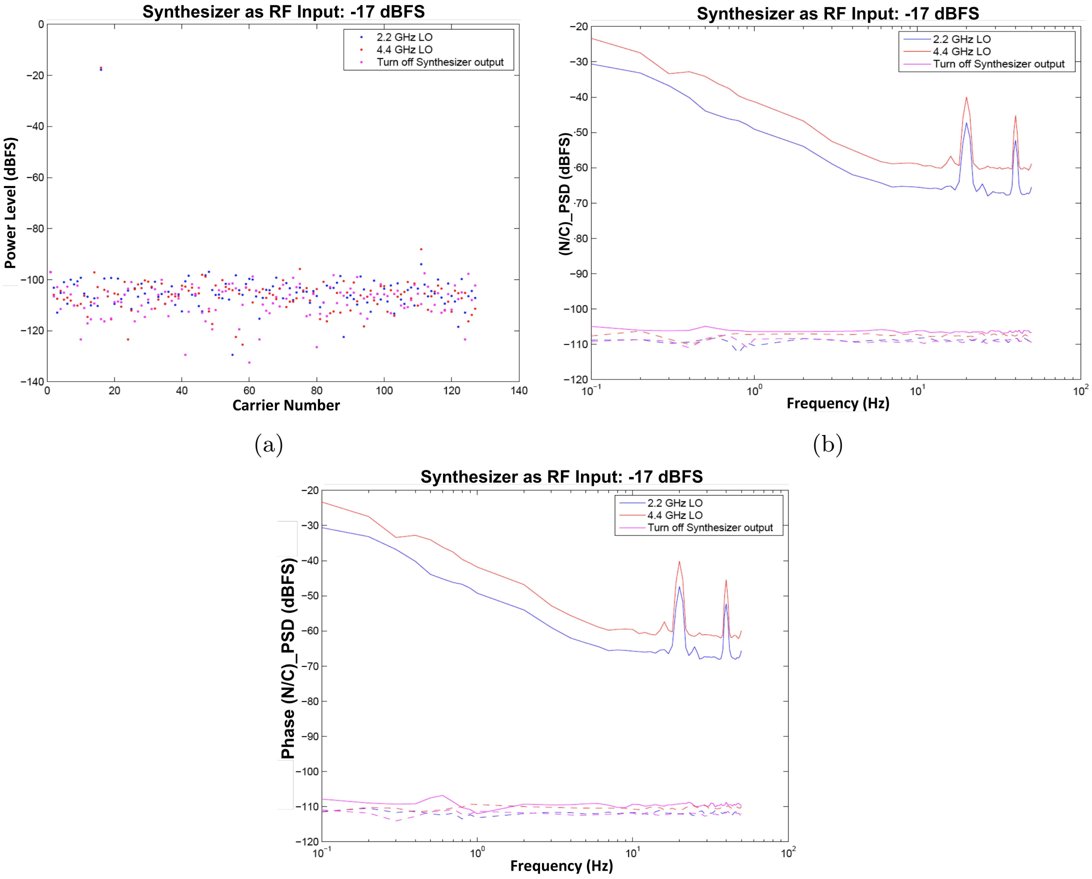

(a) (b)

(c)

Figure B.11: Power spectrum study with an external synthesizer serving as an RF input with a signal power level of -17 dBFS: same units as in Figure B.10.

1) the VCO has phase noise comparable to the external RF synthesizer; 2) the cancellation of LO phase noise between the two IQ mixers cannot occur in this setup (only one IQ mixer is involved); and 3) the phase noise is much higher than in the case where there can be cancellation, this test suggests that cancellation of LO/VCO phase noise between the two IQ mixers is significant.

2. For the phase noise, if they are sourced by a common time jitter, $\delta\phi = 2\pi f \delta t$, we will have

$$\frac{\text{Phase Noise}_{freq\_1}}{\text{Phase Noise}_{freq\_2}} = 20 \times log_{10}(\frac{freq\_1}{freq\_2}) \tag{B.6}$$

The phase noise observed when an external synthesizer is used as the RF input and an on-board VCO is used as the LO scales with the frequency and is consistent with the phase noise results predicted using the above Equation, which is a 6-dB difference in noise for a factor of two in frequency difference. This phase noise could be explained by a time jitter between the VCO and external RF source set by differential time jitter relative to the common 10-MHz reference.

3. There are spurious signals at 20 Hz and 40 Hz. These signals may originate from the VCO or external synthesizer.

B.6.4 Comparison of the Phase Noise of the VCO, Synthesizer, and OSR System

In Table B.4, we summarized the phase noise for the VCO[5] and external synthesizer (data from the data sheet) and the measured phase noise for the OSR system.

Table B.4: Phase-noise comparison

Offset Frequency	0.1 Hz	1 Hz	10 Hz
OSR System (3 GHz)	-55 dBFS/Hz	-84 dBFS/Hz	-97 dBFS/Hz
VCO on IF board (3 GHz)	-50 dBFS/Hz	-60 dBFS/Hz	-70 dBFS/Hz
External Synthesizer (3 GHz)	-	-	-60 dBFS/Hz

From the comparison table, we can observe that the phase noise of the OSR system is lower than the phase noise of the VCO chip and the phase noise of VCO is 10 dB better than the external synthesizer phase noise at 10 Hz offset freuqency[6].

B.6.5 Conclusion on the Cause of OSR Phase Noise

We summarize here the conclusions we have been able to draw about the residual 1/f noise.

1. The 1/f noise is multiplicative, indicating it is due to transfer function or gain fluctuations (Section B.4.1).
2. It is not dominated by either the first two RF amplifiers or by the DAC and ADC (Section B.4.2).
3. It is primarily a phase noise (Section B.5).
4. It cannot be explained by ADC and DAC clock jitter (Section B.6.1), consistent with (2) above.
5. It does not increase if the LO phase noise is increased, implying it is not due to imperfect cancellation of LO phase noise between the two IQ mixers.

---

[5] The VCO chip's data sheet does not provide phase noise information below 1 kHz. In order to compare the noise down to offset 0.1 Hz, we simulated the phase noise with the VCO chip and same PLL setup on the IF board.

[6] The synthesizer does not provide phase noise information below 10Hz.

Given the above, we can conclude that the 1/f noise is dominated by a phase noise caused by one, some, or all of the following sources: the remaining RF amplifiers; the baseband amplifiers; the attenuators, and some sort of other phase noise contribution from the IQ mixers that does not scale with the LO phase noise. We are not able to isolate the individual contributions of these sources given the data in hand.

B.7 Comparison of OSR Phase Noise to Expected TLS Noise

In this section, we compare the phase noise from the readout system with the two-level system noise from the resonator [38]. The relationships among the phase, frequency, and fractional frequency noise are derived by the following equations:

Consider a frequency synthesizer has an output frequency of f_{LO}. The phase of that signal as a function of time is

$$\phi = 2\pi f_{LO} t \tag{B.7}$$

The rate of change of ϕ with t is

$$\frac{d\phi}{dt} = 2\pi f_{LO} \tag{B.8}$$

Consider the fluctuation in f_{LO}, (δf_{LO}). It yields a fluctuation of $\frac{d\phi}{dt}$:

$$\delta(\frac{d\phi}{dt}) = 2\pi(\delta f_{LO}) \tag{B.9}$$

Since fluctuations and derivatives are both linear, we exchange the fluctuation δ with the derivative $\frac{d}{dt}$:

$$\frac{d}{dt}(\delta\phi) = 2\pi(\delta f_{LO}) \tag{B.10}$$

Then take the Fourier transform of the above equation, which converts $\frac{d}{dt}$ to $(2\pi i f)$

$$(2\pi i f)(\delta\phi) = 2\pi(\delta f_{LO}) \tag{B.11}$$

$$i f(\delta\phi) = (\delta f_{LO}) \tag{B.12}$$

where f is the noise frequency (or offset frequency).

For the PSD of the phase and frequency, we then have:

$$S_{freq}(f) = f^2 \times S_{phase}(f) \tag{B.13}$$

Fractional frequency noise:

$$\frac{S_{freq}(f)}{{f_{LO}}^2} = \frac{f^2 \times S_{phase}(f)}{{f_{LO}}^2} \tag{B.14}$$

The measured fractional frequency noise of the OSR system is shown in Table B.5:

Table B.5: Conversion of the measured carrier phase noise to the fractional frequency noise

OSR phase noise measured from the electronics			
Offset Frequency	0.1 Hz	1 Hz	10 Hz
Measured Phase Noise of the OSR	3.2×10^{-6}/Hz	4.0×10^{-9}/Hz	2.0×10^{-10}/Hz
Convert the phase noise into fractional frequency noise using 5.56 GHz resonator frequency			
Measured Fractional Frequency Noise of the OSR	1.0×10^{-27}/Hz	1.3×10^{-28}/Hz	6.3×10^{-28}/Hz

We wish to compare the measured readout electronics fractional frequency noise due to the resonator TLS fractional frequency noise. The noise spectrum of TLS noise for KID devices with the same capacitor design as for MUSIC is shown in Figure B.12. We see that the measured fractional frequency noise in Table B.5 is below the TLS fractional frequency noise in Figure B.12 for the audio frequency range over which it was measured, 0.1-10 Hz. This coincides with our astronomical signal band and thus assures us that TLS noise, and not readout noise, will dominate at these frequencies.

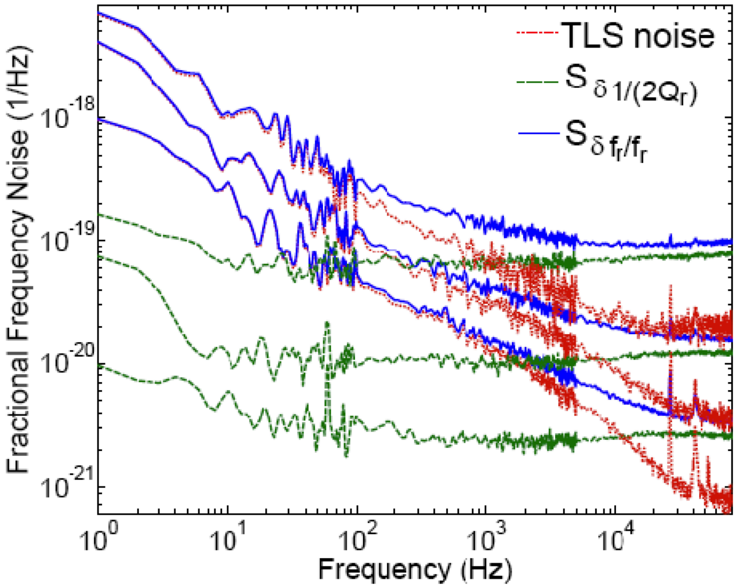

Figure B.12: Measured TLS noise level and fractional frequency noise for KID devices. The plot corresponding to 5.56 GHz resonator frequency with readout powers of −109, −103, and −97 dBm. Reprinted from [39].

Appendix C

Calculation of the BPF NEP and NEFD Values

The main steps in the calculation of the NEP and NEFD_1mm values are as follows:

$$Q_{dewar} = 2 \times k_B \times T_{dewar} \times eta \times \delta\nu \tag{C.1}$$

$$Q\_tel = 2 \times k_B \times T_{tel} \times eta \times \delta\nu \tag{C.2}$$

$$Q = P_{atm} + P_{cmb} + Q_{dewar} + Q_{tel} \tag{C.3}$$

$$P_{atm} = \frac{Q}{2} \tag{C.4}$$

$$NEP = \sqrt{2 \times h \times \nu_0 \times P_{atm} + 2 \times \frac{{P_{atm}}^2}{\delta\nu}} \tag{C.5}$$

$$NEFD\_1mm = \frac{NEP \times f_{\nu\_gal\_1mm}}{(P_{gal\_filt}/2)/\sqrt{2}} \tag{C.6}$$

$$NEFD = NEFD\_1mm \times (\frac{\nu_0}{300 \times 10^9})^{f_{\nu\_gal\_\exp}} \tag{C.7}$$

$$NEFD\_total = 1/\sqrt{\frac{1}{NEFD\_band1^2} + \frac{1}{NEFD\_band2^2}} \tag{C.8}$$

$NEFD\_1mm$ has units of $Jy/\sqrt{Hz}$

eta: optical efficiency constant 0.207 (provided by Jack Sayer)

$\delta\nu$: bandwidth of filter (Hz)

P_{atm}: total atmospheric power (W)

P_{cmb}: total CMB power (W)

T_{dewar}: dewar load (K)

T_{tel}: telescope load (K)

$f_{\nu\_gal\_1mm}$: flux density at 1 mm assumed for a galaxy source [Jy]

P_{gal}: total galaxy power for a point source of flux density $f_{\nu\_gal\_1mm}$ (W)

$f_{\nu\_gal\_exp}$: power-law exponent assumed to yield f_{ν_gal} in the band,
where $f_{\nu\_gal} = f_{\nu\_gal\_1mm} \times (\nu_{ghz}/300\text{ GHz})^{f_{\nu\_gal\_exp}}$

Bibliography

[1] H. Dole and et al., “The cosmic infrared background resolved by spitzer. contributions of mid infrared galaxies to the far-infrared background,” *Astronomy and Astrophysics*, 2006.

[2] P.K.Day, H. LeDuc, B. Mazin, A. Vayonakis, and J. Zmuidzinas, “A broadband superconducting detector suitable for use in large arrays,” *Nature*, vol. 425, p. 817, 2003.

[3] J. Zmuidzinas, “Superconducting microresonators: Physics and applications,” *Annual Review of Condensed Matter Physics*, vol. 3, pp. 169–214, 2012.

[4] J. Gao, *The Physics of Superconducting Microwave Resonators.* PhD Thesis, 2008.

[5] J. A. Schlaerth and et al., “Mkid multicolor array status and results from democam,” *SPIE Proceeding*, 2010.

[6] P. R. Maloney and et al., “Music for sub millimeter astrophysics,” *SPIE Proceeding*, 2010.

[7] N. G. Czakon and et al., “Optimization of mkid noise performance via readout technique for astronomical applications,” *SPIE Proceeding*, 2010.

[8] J. Sayers and et al., “Optics for music: a new (sub)millimeter camera for the caltech submillimeter observatory,” *SPIE Proceeding*, 2010.

[9] M. I. Hollister and et al, “The cryomechanical design of music: a novel imaging instrument for millimeter-wave astrophysics at the caltech submillimeter observatory,” *SPIE Proceeding*, 2010.

[10] P. K. Day and et al, “A slot array antenna for a millimeter/submillimeter-wave focal plane,” *SPIE Proceeding*, 2010.

[11] B. A. Mazin, P. K. Day, K. D. Irwin, C. D. Reintsema, and J. Zmuidzinas, “Digital readous for large microwave low-temperature detector arrays,” *Nuclear Instruments and Methods in Physics Research A*, vol. 599, pp. 799–801, 2006.

[12] S. J. C. Yates, A. M. Baryshev, J. J. A. Baselmans, B. Klein, and R. Gusten, “Fast fourier transform spectrometer readout for large arrays of microwave kinetic inductance detectors,” *Applied Physics Letter*, vol. 95, p. 042504, 2009.

[13] A. O. Benz, P. C. Grigis, V. Hungerbuhler, H. Meyer, C. Monstein, B. Stuber, and D. Zardet, "A broadband fft spectrometer for radio and millimeter astronomy," *Astronomy and Astrophysics*, vol. 442, pp. 767–773, 2005.

[14] K. D. Irwin and M. Halpern, "Time-division squid multiplexers," *TDM white paper*, 2008.

[15] M. Dobbs and A. Lee, "Mhz freuqency domain multiplexed readout," *CMBpol Instrument Technologies Whitepaper*, 2008.

[16] Parsons, A. Backer, D. Chang, C. Chapman, D. Chen, H. Crescini, P. de Jesus, C. Dick, C. Droz, P. MacMahon, D. Meder, K. Mock, J. Nagpal, V. Nikolic, B. Parsa, A. Richards, B. Siemion, A. Wawrzynek, J. Werthimer, D. Wright, M., "Petaop/second fpga signal processing for seti and radio astronomy," *Asilomar Conference on Signals and Systems, Pacific Grove, CA*, 2006.

[17] J. S. V. Korff, *Astropulse: A Search for Microsecond Transient Radio Signals Using Distributed Computing.* PhD Thesis, 2010.

[18] A. Parsons, *Low-Frequency Interferometry: Design, Calibration, and Analysis Towards Detecting the Epoch of Reionization.* PhD Thesis, 2009.

[19] G. Jones, *Instrumentation for wide bandwidth radio astronomy.* PhD Thesis, 2009.

[20] P. McMahon, *Adventures in Radio Astronomy Instrumentation and Signal Processing.* Master Thesis, 2008.

[21] "Million channels spectrometer http://casper.berkeley.edu/svn/trunk/projects/roachmspec/."

[22] J. A. Schlaerth, *Microwave Kinetic Inductance Detector Camera Development for Millimeter-Wave Astrophysics.* PhD Thesis, 2010.

[23] "http://casper.berkeley.edu/wiki/roach."

[24] P. A., "The symmetric group in data permutation, with applications to high-bandwidth pipelined fft architectures," *IEEE Signal Processing Letters*, 2009.

[25] P. A. Ade, G. Pisano, C. Tucker, and S. Weaver, "A review of metal mesh filters," *Proc. of SPIE Millimeter and Submillimeter Detectors and Instrumentation for Astronomy*, vol. 6275, 2006.

[26] R. Duan and et al., "An open-source readout for mkid," *SPIE Proceeding*, 2010.

[27] C. Kuo, J. Bock, J. Bonetti, J.Brevik, G. Chattopadthyay, P. Day, S. Golwala, M. Kenyon, A. Lange, H. LeDuc, H. Nguyen, R. Ogburn, A. Orlando, A. Trangsrud, A. Turner, G. Wang,

and J. Zmuidzinas, “Antenna-coupled tes bolometer arrays for cmb polarimetry,” *Proc. of SPIE, Millimeter and Submillimeter Detectors and Instrumentation for Astronomy*, vol. 7020, 2008.

[28] M. Myers, K. Arnold, P. Ade, G. Engargiola, W. Holzapfel, A. Lee, X. Meng, R. O. Brient, P. Richards, H. Spieler, and H. Tran, “Antenna-coupled bolometer arrays for measurement of the cosmic microwave background polarization,” *J. Low Temp. Phys.*, vol. 151, pp. 464–470, 2008.

[29] R. O. Brient, P. Ade, K. Arnold, G. Engargiola, W. Holzapfel, A. Lee, M. Myers, X. Meng, E. Quealy, P. Richards, H. Spieler, and H. Tran, “A multi-band dual-polarized antenna-coupled tes bolometer,” *J. Low Temp. Phys.*, vol. 151, pp. 459–463, 2008.

[30] S. Kumar, A. Vayonakis, H. G. LeDuc, P. K. Day, S. Golwala, and J. Zmuidzinas, “Millimeter-wave lumped element superconducting bandpass filters for multi-color imaging,” *IEEE Transactions on Applied Superconducitity*, vol. 19, 2009.

[31] S. Hu, L. Wang, Y.-Z. Xiong, J. Shi, D. Z. Bo Zhang, T. G. Lim, and X. Yuan, “Millimeter-wave/thz passive components design using through silicon via (tsv) technology,” *Electronic Components and Technology Conference*, 2010.

[32] L. Nan, K. Mouthaan, Y.-Z. Xiong, J. Shi, S. C. Rustagi, and B.-L. Ooi, “Design of 60- and 77-ghz narrow-bandpass filters in cmos technology,” *IEEE Transcations on Circuits and System*, vol. 55, pp. 738–742, 2008.

[33] E. Wollack, “Planar transmission line technologies,” *NASA/Goddard Space Flight Center White Paper*, 2008.

[34] R. Duan and Y. LU, “Optimization of broadband bpfusing tapered electromagnetic bandgap structures,” *Asian-Pasic Microwave Conference Proceeding*, 2008.

[35] A. R. Kerr, “Surface impedance of superconductors and normal conductors in em simulators,” *NRAO Electronics Division Internal Report*, 1996.

[36] A. Vayonakis, H. LeDuc, C. Luo, and J. Zmuidzinas, “Precision measurements of the millimeter-wave properties of superconducting thin-film microstrip lines,” *Internal Report of MUSIC group*, 2009.

[37] S. S. Mohan, M. del Mar Hershenson, S. P. Boyd, and T. H. Lee, “Simple accurate expressions for planar spiral inductances,” *IEEE Journal of Solid-State Circuits*, 1999.

[38] J. Gao, *The Physics of Superconducting Microwave Resonators.* PhD Thesis, 2008.

[39] O. Noroozian, *Superconducting Microwave Resonator Arrays for Submillimeter/Far-Infrared Imaging.* PhD Thesis, 2012.