

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Scaling Energy-Resolving Microwave Kinetic Inductance Detector Readout

Permalink

<https://escholarship.org/uc/item/5m22t48n>

Author

Smith, Jennifer

Publication Date

2024

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Scaling Energy-Resolving Microwave Kinetic Inductance Detector Readout

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Physics

by

Jennifer Pearl Smith

Committee in charge:

Professor Benjamin A. Mazin, Chair
Professor Maxwell Millar-Blanchaer
Professor Omer Blaes

June 2024

The Dissertation of Jennifer Pearl Smith is approved.

Professor Maxwell Millar-Blanchaer

Professor Omer Blaes

Professor Benjamin A. Mazin, Committee Chair

May 2024

Scaling Energy-Resolving Microwave Kinetic Inductance Detector Readout

Copyright © June 2024

by

Jennifer Pearl Smith

For my teachers

Acknowledgements

What a joy it's been getting a PhD in physics in Santa Barbara! I'm so grateful for all of the amazing people who helped me get here, supported me throughout my PhD, and have and will continue to influence and support me in my next adventure.

To start, I'd like to thank my mom and dad, who from a young age nurtured the scientist in me. Whether it was bringing bugs into the house or asking why the sky was blue *like actually* (for the tenth-time), they had endless patience for my antics and solidified my love for the natural world. I'd also like to thank my sister for beating me in every math challenge growing up. She toughened me up for the constant failure in research. On a more serious note, her mentorship has been invaluable and I'm so grateful to have someone like her paving the way. It was a great joy having Caleb join the family while I was in graduate school. I'd also like to thank my extended family for their support and for bearing with me every holiday as I try to explain my research.

I've been blessed to have amazing teachers who have supported and mentored me in my research journey. John Garnevicus was my high school physics teacher and deserves credit for piquing my interest in the subject. Ben McKinley is another amazing teacher who I was fortunate to have for multiple years in high school math.

The entire Harvey Mudd College physics department deserves all the credit in the world for cementing my foundation in the field. Peter Saeta, Thomas Donnelly, Nicholas Breznay, Ann Esin, Jason Gallicchio, Sharon Gerbode, Theresa Lynn, Liz Connolly, Gregory Lyzenga, and John Townsend: You all broke my brain and with it, broke down the way I think about and approach problems. Then you built it back up the right way—the physics way! You gave me an amazing foundation for graduate school. Michael Storrie-Lombardi is an incredible professor, mentor, and friend and I was so overjoyed when you and Lisa moved to Santa Barbara. Thank you for always being willing to

bounce ideas with me and for the lovely dinner parties. It's hard for me to think of someone more supportive, enthusiastic, and encouraging than you.

I'd also like to thank my undergraduate research advisors. Karl Haushalter at Harvey Mudd fostered an early passion for instrumentation. Gray Rybka and Leslie Rosenberg at the University of Washington let an undergraduate characterize a HEMT. Zeesh Ahmed at SLAC is the reason I chose my field of research. Thank you for encouraging me to apply to UCSB and seeding my eventual placement in Ben Mazin's lab.

Thank you Ben for teaching and mentoring me all these years. I have so much admiration for your skills and breadth of knowledge. I learned so much and got many opportunities as a graduate student in your group. Beyond research, you are a compassionate, caring, and interesting person—all of which I appreciate, and I am beyond thankful to have had a cool advisor. It's no surprise with all of your talents that we have a cool research group. Isabel, Neelay, Sarah, and Noah all came before me and were so friendly and supportive and showed me the ropes. I'd like to give a special shout out to Nic for being a stellar researcher and friend. I learned so much from you and I'm grateful to have stayed in touch with you all these years. I'd also like to acknowledge the new cohort of Mazin lab graduate students: Hawkins, Crystal, Majid, Josh, Aled, Michael, and Nikki—the lab is in good hands with you all. It's been a joy to work with you and I am excited for your future. I have to acknowledge Jeb Bailey, who was like my second advisor. Thank you for sticking with me on the crazy, exciting, and challenging projects that we tackled together. You are an amazing researcher and your passion for quality really resonates with me. Your work ethic is also inspiring albeit a bit terrifying. I'd also like to thank my committee members Omer Blaes and Max Millar-Blanchaer. I had the great privilege of taking two courses with Omer. Your notes are prolific and you are an incredibly kind and compassionate person. Max, I so appreciate the Astro coffee you organized this year. My Thursday maple doughnut (and the great company) got me

through the last bit of thesis writing.

Professionally, I'd like to thank Graham Schelle at AMD/Xilinx for taking a chance on a physics student for his intern. I learned so much working with you at Xilinx and I'm forever grateful for your mentorship. Thank you also to Damon Russel at JPL / Amazon for sharing your knowledge in microwave engineering with me and encouraging me to branch out in my research. I'd also like to thank Daniel Sank at Google for encouraging me to learn more about quantum computing. Thank you to Evan Jeffrey at Google for advising my research project, answering every question I had about qubit readout, and helping me hone my FPGA skills. I'd also like to give a shout out to Richard Boone for tutoring me in Verilog. You all were so helpful to me as I explored the professional side of research and I am very grateful for my experiences.

In addition to all the academic and professional support, I had an amazing cadre of friends rooting for me. Sara, Annabel, Mary, and Christina are life-long friends. You all are amazing people and despite all our growth, I still feel like we get one another on a deep level. Kim and Elise carried me through college with their hilarity and awoke my passion for cheese, wine, and ice cream. I'd also like to thank the CMS volleyball team. Winning a national championship with you guys is one of my proudest accomplishments. Special shout out to Mariah Rigg who continually inspires me with her kindness and depth.

I realized once I got to graduate school that I had taken for granted having so many close female friends. I am grateful to the UCSB Women in Physics / Women and Gender Minorities in Physics and Women in Science and Engineering groups that introduced to me to amazing women in STEM. I am particularly grateful to Molly Wolfson, who succumbed to Starbucks addiction with me. Our daily walks to the UCEN kept me sane these past few years and despite all of our differences, it always felt like we were going through something similar. I'd also like to thank Ana who has been a bright spark of

friendship and encouragement these past few years. Thank you also to Amara for your friendship and for sharing your family Challah recipe with me—I will treasure it!

In my cohort I'd like to thank: Mirek Brandt, for the long phone calls; Sean McBride, for all the crazy stories and fun BBQs; Minghan Chen, for the early morning football; Craig Pellegrino, for always being down for anything; Alex Potts, for being incredibly thoughtful and kind; Fernando Iniguez: for always being so open and genuine, and Matt McEwen: for the cocktail/tea nights. I'm looking forward to staying in touch with you all.

I'd also like to thank Charlie and Alina, who were invaluable mentors and friends at the start of grad school. Thank you for sharing your home, friends, and life advice with me. Thank you also to Jamie Burke for your infectious laughter and for organizing so many fun trips. Caddying your journey though every elevation in the continental United States was a highlight.

Lastly I'd like to thank Alex and Cosmo. You both have seen it all: the triumphs, the tears, and everything in-between. Thank you Alex for your unending patience and love. Your rock-solid support allowed me to flourish in spite of the stress and pressure of graduate school. I will forever be grateful for the personal tour of the Compact Muon Solenoid at CERN and I am so excited for our future adventures together.

Curriculum Vitæ

Jennifer Pearl Smith

Education

2024 Ph.D. in Physics, University of California, Santa Barbara
2016 B.S. in Physics, *with high distinction*, Harvey Mudd College

Awards and Honors

Zonta International Amelia Earhart Fellowship Award
NASA Space Technology Research Fellowship

Internships

Google Quantum AI
Xilinx, Inc.

Selected Publications

J. Smith, J. I. Bailey, III, A. Cuda, N. Zobrist, B. A. Mazin, “MKIDGen3: Energy-Resolving, Single-Photon-Counting MKID Readout on an RFSoc Platform,” *ArXiv pre-print*, June 2024, doi:10.48550/arXiv.2406.09764.

J. Smith, B. A. Mazin, A. Boaventura, K. Thompson, M. Daal, “Improved Flexible Coaxial Ribbon Cable for High-Density Superconducting Arrays,” in *IEEE Transactions on Applied Superconductivity*, vol. 34, no. 2, pp. 1-6, March 2024, doi: 10.1109/TASC.2024.3350516.

J. Smith, J. I. Bailey, III, B. A. Mazin, “Highly-Multiplexed Superconducting Detector Readout: Approachable High-Speed FPGA Design,” in *IEEE 30th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pp. 1-2, May 2022, doi: 10.1109/FCCM53951.2022.9786140.

J. Smith, M. Ruiez, G. Schelle, “100 Gigabit Ethernet for RFSoc-PYNQ Overlays,” in *PYNQ: Python Productivity Learn, AMD Xilinx*, 20 Sept. 2021, <https://discuss.pynq.io/t/100-gigabit-ethernet-for-rfsoc-pynq-overlays/3053/5>. Accessed 2024.

J. Smith, J. I. Bailey, III, J. Tuthill, L. Stefanazzi, G. Cancelo, K. Treptow, B. A. Mazin, “A High-Throughput Oversampled Polyphase Filter Bank using Vivado HLS and

PYNQ on a RFSoc,” in *IEEE Open Journal of Circuits and Systems*, vol. 1, no. 1, Feb. 2021, doi: 10.1109/OJCAS.2020.3041208.

J. Smith, B. A. Mazin, A. Walter, M. Daal, J. I. Bailey, III, C. Bockstiegel, N. Zobrist, N., Swimmer, S. Steiger, N. Fruitwala, “Flexible Coaxial Ribbon Cable for High-Density Superconducting Microwave Device Arrays,” in *IEEE Transactions on Applied Superconductivity*, vol. 31, no. 1, pp. 1-5, Jan. 2021, doi: 10.1109/TASC.2020.3008591.

A. Walter, N. Fruitwala, S. Steiger, J. I. Bailey, III, N. Zobrist, N. Swimmer, I. Lipartito, **J. Smith**, S. Meeker, C. Bockstiegel, G. Coiffard, R. Dodkins, P. Szypryt, K. Davis, M. Daal, B. Bumble, G. Collura, O. Guyon, J. Lozi, S. Vievard, F. Martinache, T. Currie, B. A. Mazin, “The MKID Exoplanet Camera for Subaru SCEXAO,” in *Publications of the Astronomical Society of the Pacific*, vol. 132, no. 1018, Dec. 2020, doi: 10.1088/1538-3873/abc60f.

N. J. Swimmer, B. A. Mazin, C. Bockstiegel, J. I. Bailey III, G. Coiffard, M. Daal, K. Davis, N. Fruitwala, I. Lipartito, **J. Smith**, S. Steiger, N. Zobrist, T. Cook, S. Chakrabarti, C. Mendillo, J. Martel, and K. Hewawasam, “The PICTURE-C MKID camera”, in *Ground-based and Airborne Instrumentation for Astronomy VIII*, edited by C. J. Evans, J. J. Bryant, and K. Motohara (Dec. 2020), p. 79.

R. Kothari, V. Jones, D. Mena, V. Bermudez, Y. Shon, **J. Smith**, D. Schmolze, P. Cha, Y. Fong, M. Storrie-Lombardi, “Raman Spectroscopy and Artificial Intelligence to Predict the Bayesian Probability of Breast Cancer,” in *Biophysical Journal*, vol. 118, no. 3, Feb. 2020, doi: 10.1016/j.bpj.2019.11.355.

Abstract

Scaling Energy-Resolving Microwave Kinetic Inductance Detector Readout

by

Jennifer Pearl Smith

Superconducting detectors are well-suited to serve as wide field-of-view, energy resolving, single-photon-counting cameras for ground and space-based instruments. These cameras have a wide range of sensitive imaging applications in biology, astronomy, particle physics, cosmology, and quantum information. However, science results have been delayed by challenges in scaling superconducting detectors into large arrays with sufficient pixel resolution to create sharp images. Superconducting Microwave Kinetic Inductance Detectors (MKIDs) are promising detectors for these applications because they are inherently multiplexable, providing a feasible way to create large-format arrays. In this thesis, I present two major improvements made in scaling MKID array technology inside and outside the cryogenic system. First, I discuss a new RFSoc-based digital readout that provides a dramatic reduction in the weight, volume, and power of the room temperature electronics. This compact digital readout enables scaling to megapixel array formats and increases the feasibility of future space-based deployment. Next, I share a new superconducting coaxial ribbon cable that improves readout signal integrity inside the cryogenic system and reduces cryogenic heat load while supporting a dense wiring format. In addition to furthering MKID technologies, these systems may support scaling of other superconducting detector or qubit systems in the near future.

Contents

Curriculum Vitae	ix
Abstract	xi
List of Figures	xiv
List of Tables	xvi
1 Introduction	1
1.1 Superconducting Detectors	1
1.2 Microwave Kinetic Inductance Detectors (MKIDs)	10
1.3 MKID Readout	11
1.4 Field Programmable Gate Arrays	18
2 Gen3 MKID Readout	22
2.1 Introduction	23
2.2 MKID Readout	24
2.3 Requirements	29
2.4 System Design	35
2.5 System Performance Characterization	53
2.6 Discussion	67
2.7 Conclusion	72
3 Oversampled Polyphase Filter Bank	74
3.1 Introduction	75
3.2 Channelizer Design	85
3.3 Implementation	86
3.4 Hardware Testing and Verification	93
3.5 Conclusion	98

4	Cryogenic Cabling	108
4.1	FLAX 1.0	109
4.2	FLAX 2.0	120
5	Conclusion	137
	Bibliography	139

List of Figures

1.1	Semiconductor vs. Superconductor Band Gap	3
1.2	Superconducting Detector Imaging Area Advantage	4
1.3	MKID Operating Principle	10
1.4	Frequency-Division Multiplexing	12
1.5	Twenty-Kilopixel MKID Array	13
1.6	MKID Frequency Response	14
1.7	Homodyne Readout	15
1.8	MKID Photon Trajectory and Phase Response	16
1.9	Frequency-Multiplexed Readout Waveform	17
1.10	AMD/Xilinx Zynq™ Architecture	21
2.1	MKID Frequency Response	25
2.2	Frequency-Division Multiplexing	26
2.3	MKID Photon Readout	28
2.4	Photo of Data-Taking Setup	36
2.5	Gen3 vs. Gen2 Scale Photo Comparison	38
2.6	Digital Signal Processing Overview	43
2.7	Matched Filter Generation	44
2.8	Trigger Operation	45
2.9	System Block Design	50
2.10	Capture Subsystem Block Design	51
2.11	Device Chip Area Resource Utilization	52
2.12	Dilution Refrigerator Experimental Setup	55
2.13	Loopback Comb Transmission	57
2.14	Single Channel Phase Noise Power Spectral Density	58
2.15	Phase Noise Power Spectral Density Floors Across All Channels	59
2.16	Matched Filter Performance	61
2.17	Resolving Power with Increasing Tone Number	62
2.18	On-Resonance Phase Noise Power Spectral Density with Increasing Tones	63
2.19	DAC Dynamic Range vs. Number of Tones	64
2.20	Resolving Power Degradation Experimental Summary	68

3.1	Channelization Overview	77
3.2	Windowing Data Before a Fourier Transform	79
3.3	Photon Phase Pulse and Detector Multiplexing	82
3.4	Critically Sampled vs. Oversampled Channelizer Response	83
3.5	Oversampled Channelizer Spectral Leakage	84
3.6	Milti-Lane OPFB Block Diagram	87
3.7	HLS Delay + Interleave Block	88
3.8	HLS Reorder Block	89
3.9	Multi-Lane Polyphase Filter Bank Coefficient Decomposition	90
3.10	Oversampled Polyphase Filter Bank Example Project Vivado Block Design	91
3.11	Python Productivity for ZYNQ Overview	92
3.12	Implemented OPFB Channelizer Response	96
3.13	OPFB Device Area Utilization	97
3.14	Polyphase Decomposition	100
3.15	Abstracted Polyphase Decomposition	101
3.16	Data Movement in an Oversampled Polyphase Filter Bank Channelizer .	102
3.17	Data Reordering in an Oversampled Polyphase Filter Bank Channelizer .	106
3.18	Actual Gen3 OPFB Channelizer Response.	107
4.1	FLAX 1.0 Images	110
4.2	FLAX 1.0 Connectorization Exploded View	111
4.3	FLAX 1.0 Experimental Setup	113
4.4	FLAX 1.0 Transmission and Crosstalk	114
4.5	FLAX 1.0 TDR	116
4.6	FLAX 1.0 Thermal Conductivity	118
4.7	Maybell Quantum FLAX Prototype	123
4.8	Maybell Quantum FLAX Manufacturing Data	125
4.9	FLAX 2.0 Connectorization Exploded View	126
4.10	FLAX 2.0 End Connectorization	128
4.11	FLAX 2.0 Experimental Setup	129
4.12	FLAX 2.0 Transmission and Crosstalk	130
4.13	FLAX 2.0 TDR	132
4.14	FLAX 2.0 Heatload	134

List of Tables

2.1	Gen3 Requirements	31
2.2	Gen3 vs. Gen2 Specifications	37
2.3	FPGA Resource Utilization	50
3.1	FPGA Channelizer Method Resource Comparison [1]	79
3.2	OPFB FPGA Resource Utilization Comparison	94
4.1	FLAX 1.0 Summary of Properties	115
4.2	FLAX 2.0 Summary of Properties	127

Chapter 1

Introduction

1.1 Superconducting Detectors

Superconducting devices are lauded as an up-and-coming technology with broad applications to highly-sensitive measurements across a variety of scientific disciplines. Excitement is fueled by the promise of high-sensitivity and low noise as well as an impressive range of application-specific advantages. In quantum computing, superconducting qubits are the forefront qubit technology in part because the quantum processor is naturally isolated from the environment and the nanolithography fabrication process provides some control over the qubit, coupler, and wiring design [2, 3]. To the dismay of superconducting qubit researchers, these devices also behave as amazingly sensitive detectors [4, 5]. Ironically, this detector technology has been harnessed by competing quantum computing and quantum information platforms which use trapped ions or photons as qubits and superconducting detectors to measure the quantum state [6, 7]. Despite growing interest in quantum computing, the main development and deployment of superconducting detectors has historically been for highly sensitive photon detectors [8, 9, 10, 11, 12]. This application has the broadest reach with demonstrated utility in bio-imaging, particle physics,

astronomy, cosmology, and other high-sensitivity imaging science [13, 14, 15, 16, 17]. Going forward, we will focus on this imaging application as it pertains to the main goals of this work.

1.1.1 Superconducting Detectors for Sensitive Cameras

Superconducting detectors provide several key advantages for sensitive imaging applications. They are capable of single-photon-counting across a broad energy range with intrinsic energy resolution and precise timing information. They are also radiation tolerant. These attributes make them well-suited to serve as wide field of view, energy resolving, single-photon-counting cameras for ground and space-based instruments. A qualitative discussion of specific advantages superconducting detectors can offer over current state-of-the-art semiconductor technologies is presented below.

Single-Photon-Counting

Superconducting detectors get their sensitivity from Cooper pairs which are pairs of electrons bound together in an energy well with a defined superconducting gap energy¹. This is an important contrast from semiconductor-based detectors whose sensitivity is determined by the semiconductor band gap which is the energy required to promote an electron to the valence band. Superconducting band gaps are typically a thousand times smaller than semiconductor gaps meaning photon signals are orders of magnitude larger in superconductors (see Fig. 1.1). Semiconductor readout also requires flowing current to transport the released charge(s) to a detector and it's all but impossible to determine the total number of photons incident on the semiconductor. There can also be a noise penalty for reading out the detector too quickly. In the case of superconducting

¹For a more complete, quantitative treatment of superconducting detector physics see [18] p.15-28.




Figure 1.1: Cartoon representation of the difference in sensitivity between semiconductor and superconductor detector technologies. The small superconducting band gap makes superconductors more sensitive and responsive to a broader incident wavelength range.

detectors, individual photons are continuously resolved with effectively no noise. These differing detection sensitivities give superconducting detectors a natural advantage in photon-limited applications where the noise per photon is an important metric.

Wavelength Sensitivity

The small superconducting gap also means that superconducting detectors are sensitive to a wide range of photon energies. In practice, these detectors can detect photons from x-ray to radio wavelengths. Depending on the specific noise requirements of the application, accomplishing the same feat with semiconductors might require multiple




Figure 1.2: Cartoon representing the detector area advantage offered by superconducting detectors that are inherently energy resolving. Similar science can be done with fewer superconducting pixels.

separate detector technologies. Even in narrow-band applications, superconductors are suitable for many different detection regimes with little re-engineering.

Energy Resolution

Another boon for superconducting detectors is the photon response is proportional to the photon energy. A higher energy photon breaks more Cooper pairs. By quantifying the detector response, each superconducting detector can act as a spectral pixel. Semiconductor detectors are not so fortunate and must use a wavelength dispersing optical element, such as a grating, to diffract different wavelengths across multiple pixels. Superconducting detectors have the potential to greatly reduce the number of pixels needed to achieve the same energy resolution (see Fig. 1.2). For applications requiring very high spectral resolution, such as chemical spectroscopy, a grating can be paired with an energy-resolving detector for a multiplicative improvement in energy resolution (this concept is explored in detail in [19]).

Timing Precision and Count Rate

We've already alluded to an effectively continuous readout mode for superconducting detectors which does not carry a read noise penalty like semiconductor readout. For superconductors, the limit on photon arrival time precision is set by the detector rise time and is related to the rate at which photoelectrons from the initial photon impact are down-converted to phonons that break Cooper pairs. Similarly, the maximum count rate is based on the time it takes the broken Cooper pairs to recombine. In practice both the rise time and recombination time are order microseconds [20]. This timing resolution gives superconducting detectors an advantage in ground based astronomy where the speed is fast enough to couple with adaptive optics systems and correct for atmospheric aberrations in real time. The precise photon arrival statistics also allow for more advanced post-processing techniques to further correct for atmospheric effects [21, 22].

Radiation Hardness

For space-based missions, radiation hardness is an important consideration for mission success and return on investment. Semiconductor detectors are made of specially doped materials whose properties can be irreversibly damaged by exposure to charged particles. One high-profile example is the Chandra X-ray semiconductor detectors which suffered rapid degradation early on in the mission from exposure to low energy protons each time the instrument passed through Earth's radiation belt [23, 24]. The situation was mitigated somewhat by storing the detectors in a shielded position every radiation belt crossing but still it's estimated the charge-transfer-inefficiency² degrades 2.3% a year [26].

Superconducting detectors do not suffer from the same radiation sensitivity due to

²Poor charge-transfer-inefficiency in this type of semiconductor detector can cause streaks or smearing in an image, reducing energy resolution and in some cases losing signals entirely. Refer to [25] for more details.

their simple construction. Typically, they consist of one or two layers of metal patterned onto a high-quality crystalline dielectric and the structure is not easily disrupted. Indeed this has been studied with a 160 MeV proton beam source and researchers found no degradation in the superconducting detectors after an equivalent of 5 years of worst case scenario radiation exposure at the Lagrange L2 point [27].

1.1.2 Superconducting Detectors at Work

With all the previously mentioned technical advantages, it's no wonder there is a large body of work developing and demonstrating superconducting detectors. A large development effort has been under way for Transition Edge Sensors (TESs), which use the superconducting transition to measure incident energy [12]. TESs can be single-photon-counting and energy-resolving but they've gained the most traction in the sub-millimeter where they act as integrating energy detectors for measuring the Cosmic Microwave Background (CMB) polarization [28, 29]. Most recently, they were deployed in a suite of such experiments: Keck Array (2011), BICEP3 (2015), and the balloon SPIDER (2015) [17]. Based on the success of those experiments, 32,000 TESs were deployed in the latest iteration, BICEP Array, which began observations in 2020 [30].

Superconducting nanowire single-photon detectors (SNSPDs) are skinny meandering wires that, similar to TESs, rely on incident photons to momentarily disrupt superconductivity in the device [31]. SNSPDs have made inroads in single-photon counting applications where low timing jitter is critical, such as quantum photonics [32, 33]. Commercial SNSPDs are available for precisely this application [34, 35]. High-count rate applications also benefit from SNSPDs. A 1,024-pixel array was demonstrated in lab in 2019 [36], and most recently a 400,000-pixel SNSPD camera was demonstrated by NIST [37]. SNSPDs have found use in variety of applications from biological imaging to remote

sensing [38] and many more proposals exist for future science applications [39, 40].

We will save the full introduction of Microwave Kinetic Inductance Detectors (MKIDs)—the main subject of this work, for Sec. 1.2, but it’s worth noting now these superconducting detectors have been successfully demonstrated as well. The MKID Exoplanet Camera (MEC) is the premiere energy-resolving, single-photon-counting MKID camera with 20,000 pixels [41]. MEC has been used to detect a binary star [42] and a brown dwarf [43] and is still observing at Subaru Telescope.

Finally, we will give Superconducting Tunnel Junctions (STJs) an honorable mention. Even though they haven’t seen much interest since a 120-pixel array was developed in 2006 [44], STJs were one of the first superconducting detector technologies and laid the foundation for modern understandings of superconducting detector physics [8]. They found use in radio astronomy in the 80s but interest ultimately petered out partly due to the readout complexity imposed by large arrays [45, 46].

1.1.3 Modern Challenges

Despite advancements over the past few decades, there are still several technical hurdles that must be cleared before the full benefits of superconducting detectors can be widely utilized. Many of the problems boil down to one central issue: scaling.

Scaling is uniquely challenging for superconducting arrays because it is fundamentally at odds with the limitations of the cryogenic systems needed to keep the devices cold and superconducting. Not only does the cryogenic aspect slow the rigorous cold testing of fabrication recipes needed to improve device yield, but it also puts strict heat load requirements on cryogenic electronics—the likes of which the commercial electronics industry has never seen before. Presently, researchers are left to design, engineer, and test many cryogenic components themselves.

Cryogenic device readout in particular has been a pain point felt from quantum computing to imaging projects where systems need hundreds of thousands if not millions of superconducting devices to meaningfully operate [47, 48]. This scale is hindered by the cryogenic system’s cooling capacity which limits the number of wires carrying readout signals from room temperature to the ~ 1 K detectors. This head load requirement motivates the need to multiplex as many superconducting devices per readout line as possible to facilitate large arrays. This in turn puts enormous pressure on room temperature electronics to separate out and process all the data in real time. Readout demands are especially high for projects looking to use superconducting devices as single-photon-counting, energy-resolving detectors because the data volume requires large data bandwidths and intensive, real-time digital signal processing. While it’s true large TES arrays have been fabricated for CMB readout, the science goals permitted the TESs to be operated in a simpler mode; they did not time-tag photons or discern their energy, leading to relaxed throughput requirements. Even so, there was an enormous effort to develop a complicated, power-intensive, multiplexed readout for TESs [49, 50]. All this development in the sub-millimeter for CMB TESs has not translated to the single-photon-counting, energy resolving version and the largest such TES array is still only 36 pixels [51]. More progress has been made with MKIDs which have demonstrated an array of 20,000 energy-resolving, single-photon-counting pixels in the MEC instrument; however, the densely-packed cryogenic wires had issues with signal integrity and the digital electronics rack is excessively large, heavy, and power-intensive by modern standards. The readout system is also totally unfit for spaceflight where minimal weight, volume, and power are critical requirements.

Overall, challenges associated with scaling cryogenic readout systems have proved to be a bottle-neck for superconducting array development. Fortunately, many of these problems can be solved by applying engineering principles guided by the underlying

physics. This thesis aims to do just that and take a step towards overcoming some of the scaling challenges both inside and outside the cryogenic system. It's my hope this work will push us closer to leveraging all of the advantages of superconducting detectors for ground and space based missions. Specific contributions are summarized next.

1.1.4 This Work

This PhD thesis is focused on improving the scalability of single-photon-counting, energy-resolving Microwave Kinetic Inductance Detectors (MKIDs). Many results are also applicable for other superconducting devices previously mentioned. Two main contributions are a new digital readout (Chapter 2 and 3) and a new cryogenic wiring solution (Chapter 4). The cryogenic wiring solution is notable for fixing some early signal integrity issues in MEC and helping unleash the first science results. The technology is now under commercial development for quantum computing, proving it's broad utility. The new digital readout provides a dramatic reduction in the weight, volume, and power of the readout electronics and paves the way for future scaling to megapixel arrays. It also progresses the technology road-map towards space-based missions.

It's been exciting and challenging to contribute to building, deploying, and fixing some of the largest superconducting detector arrays in the world. Part of the opportunity to push this scale comes from Microwave Kinetic Inductance Detectors themselves which in some ways are the most inherently scalable superconducting detector technology. Next we discuss their unique properties and provide context for this work.




Figure 1.3: Overview of MKID operating principle and photon readout. When a photon strikes the inductor, the resonance frequency shifts which creates a pulse in the phase of the probe tone.

1.2 Microwave Kinetic Inductance Detectors (MKIDs)

Microwave Kinetic Inductance Detectors (MKIDs) work by detecting changes induced in the kinetic inductance of superconducting materials by incident photons³. When a photon strikes the inductor, it breaks Cooper pairs and generates quasiparticles (unbound electrons) which raise the inductance with magnitude proportional to the energy of the photon. After some time, the quasiparticles recombine and the inductor returns to its quiescent state. The photo-sensitive inductor is fabricated in a superconducting resonant circuit where the properties can be continuously monitored (see Fig. 1.3).

MKID detector design, fabrication, and materials selection are all active areas of research. The MEC array uses a platinum silicide (PtSi_x) superconducting film and a lumped element resonator design (see Fig. 1.4, lower left and Fig. 1.5). MKIDs have also been fabricated and characterized with titanium nitride, aluminum, and hafnium, and more. A more complete guide to MKID materials including their fabrication process, properties, and performance is given in Chapter 4 of [18] and reference [54].

³For more details on MKID physics, refer to the following dissertations: [18, 52, 53]

Frequency Division Multiplexing

The beauty of the MKID resonator circuit is that it provides a natural way to multiplex and scale detector arrays. This is one of the main advantages MKIDs have over TESs, SNSPDs, and other superconducting detector technologies with no straightforward approach to scaling. By varying the capacitor in each MKID, each resonator can be made with a separate resonance frequency (see Fig. 1.4). This allows each device to be addressed and read out independently despite sharing the same readout line. This technique is called frequency division multiplexing (FDM). Using FDM, we can create 2,048 detectors between 4-8 GHz on a single microwave feedline. These feedlines can be replicated to form kilopixel arrays. The 20,000-pixel MEC array consists of 10 of these microwave feedlines and is shown in Fig. 1.5. Before observing, a micro-lens array is fitted over the MKID array which focuses light on the photosensitive inductor portions of each circuit. This brings the photosensitive fill factor to over 90%.

1.3 MKID Readout

MKID readout is central to the design and operation of MKIDs. We've already covered individual MKID readout and frequency-division-multiplexing but here we will fill in a few more details and practical considerations.

Overview

MKID readout works by using an individual probe tone at the resonance frequency of each resonator. The phase of the probe tone remains roughly constant when there are no photons present, but when a photon strikes the MKID inductor, we observe a pulse in the phase of the probe tone. The pulse rise encodes the impact time and the pulse height encodes the photon energy (summarized in Fig. 1.3).




Figure 1.4: Overview of frequency-division-multiplexing. Top: Circuit representation of frequency-division-multiplexed MKIDs. Each MKID resonator has a unique resonance frequency. Bottom: Microscope image of a PtSi_x MKID array. The microwave feedline (bottom triangle) connects to many MKID pixels. Four MKIDs are shown in a zoomed view on the right. The large structure in the middle of each MKID is the interdigitated capacitor. This capacitance value is unique to each MKID on a feedline. The smaller light purple rectangle above each MKID is the meandered photo-sensitive inductor. Together each capacitor and inductor make a resonance circuit with a unique resonance frequency. Each MKID is horizontally and vertically spaced $150 \mu\text{m}$.




Figure 1.5: A 20,000-pixel MKID array. The array consists of 10 microwave feedlines each servicing 2,000 MKID resonators. Each of the barrel connectors on one side of the gold box feed in a superposition of probe tones between 4 and 8 GHz. The other side routes the readout tones back out to the readout electronics where the phase of each tone is monitored. Before observing, a micro-lens array is positioned over the detectors which focuses light onto the photo-sensitive inductor portion of each circuit.

MKID Setup

When an MKID feedline comes back from fabrication, the exact resonance frequencies are unknown due to process uncertainties. The first step in reading out an MKID is identifying the resonance frequency and optimal drive power. These can be identified from the frequency response of the circuit. A simulated resonator frequency response is shown in Fig. 1.6⁴. The transmission minimum is the resonance frequency of the circuit. A frequency near the resonance frequency is used for the readout tone because this is where the phase is both linear and most responsive. While the transmission and phase response are more intuitive representations of a resonator, the IQ loop contains the most information. In practice, the IQ loop is what we measure and use to characterize the

⁴This simulated data is generated from the open source `MKIDReadoutAnalysis` package available on GitHub. The physical model used in the package is derived from microwave engineering principals in [18] Appendix B1 and is based on work presented in [55, 56, 57].

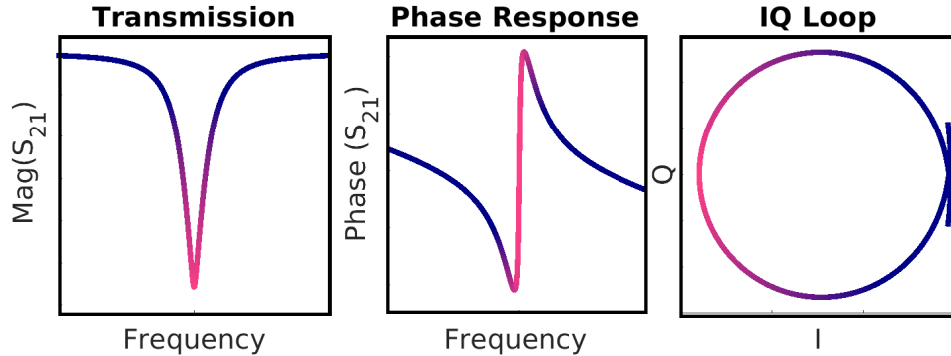


Figure 1.6: Superconducting detector frequency response magnitude (left), phase (center), and complex representation (right). The difference between the drive frequency and resonance frequency is shown in a pink to purple gradient on each plot. The optimal readout frequency is near the resonance frequency (pink)⁴.

detector. Beyond identifying the optimal readout power and frequency, the IQ loop can be used for debugging potential issues such as frequency collisions with nearby resonators. Next we'll discuss how to measure this loop.

Readout Electronics: Single MKID

MKID devices are typically read out using a homodyne scheme which means the signal of interest is a modulation of a single frequency. In this case, we are reading out phase and amplitude modulations of the readout tone. Fig. 1.7 shows a basic experimental setup for MKID readout. An oscillator is used to generate a sine wave at the readout frequency, f_{ro} . The MKID device modulates the amplitude and phase of the readout tone with a time-dependent signal $A(t)e^{i\phi(t)}$. Next, we want to eliminate the effects of the readout tone in order to isolate the modulation. This is also referred to as down conversion because we are converting the carrier frequency down to DC or 0 Hz. Down conversion is accomplished mathematically by multiplying the MKID output by the complex conjugate of the readout signal,

$$A(t)e^{i(2\pi f_{ro}t + \phi(t))}e^{-i(2\pi f_{ro}t)} = A(t)e^{i\phi(t)}. \quad (1.1)$$




Figure 1.7: Overview of homodyne MKID readout. An oscillator on the left creates a single-frequency waveform at the readout frequency. The MKID device modulates the amplitude and phase of the readout tone. The same oscillator frequency is phase-shifted by 0° and 90° and multiplied by the MKID output to effectively cancel out the the readout tone. The remaining in-phase, $I(t)$, and quadrature, $Q(t)$ signals can be represented as a time-dependent trajectory on the complex plane.

In electronics hardware, the complex multiply in Eq. 1.1 is realized by two real multiply operations: one with an in-phase signal to produce the real part, $I(t)$, and the other with a 90° phase-shifted signal to produce the quadrature part, $Q(t)$. The two resulting signals $I(t)$ and $Q(t)$ can be represented as a time-dependent imaginary number where $I(t)$ is the real-part and $Q(t)$ is the imaginary-part.

In practice, the oscillator on the left in Fig. 1.7 is re-programmable. To characterize a device, the oscillator is swept, producing a range of frequencies around the MKID resonant frequency. During the sweep, $A(t)e^{i\phi(t)}$ traces out the resonator loop (see Fig. 2.3, left). After the loop has been characterized and the optimal readout frequency identified, the oscillator is fixed at this frequency. In this mode, the down converted signal, $A(t)e^{i\phi(t)}$, is relatively stationary⁵ in the complex plane. Now we are ready for photon detection. When a photon strikes the detector, $A(t)e^{i\phi(t)}$ follows a trajectory shown in Fig. 2.3. The modulation $A(t)e^{i\phi(t)}$ can be represented by two coordinates related to the phase and amplitude of the signal. In practice, due to computational

⁵In reality there is some noise from the readout electronics and MKID device. See Chapter 5 in [18] for a detailed description of readout and detector noise.

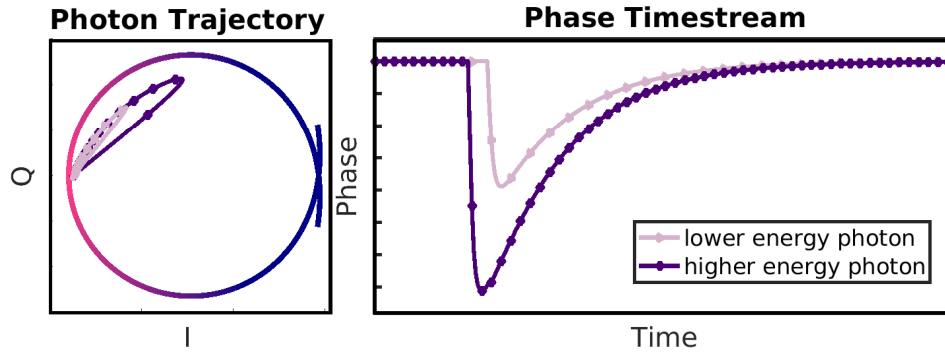


Figure 1.8: Simulation of MKID photon readout trajectory in the complex plane and phase time series. Left: the MKID resonator loop is shown in a pink to purple gradient representing the frequency difference from the resonant frequency. Each point on the loop corresponds to a different readout frequency with pink being the resonant frequency. The light purple and dark purple lines show the trajectory the resonant frequency takes when a low energy and high energy photon hits the MKID, respectively. Right: The phase of the resonant frequency vs time for the two photon impacts. The depth of the phase pulse corresponds to the energy of the photon. The low and high energy photons were constructed to hit at different times for visual clarity. Data was generated using the `MKIDReadoutAnalysis` package (see footnote 4).

resource limitations, we only use the higher signal-to-noise phase signal (Fig. 2.3, right) and throw away amplitude information. This final conversion to phase is accomplished by evaluating

$$\phi(t) = \tan^{-1} \left(\frac{Q(t)}{I(t)} \right), \quad (1.2)$$

where the loop has been centered and rotated in the complex plane such that the dark phase is zero and the maximum pulse size is less than 180° .

Readout Electronics: Many MKIDs

Previously, we discussed the basic electronics needed to set up and read out a single MKID but we need a succinct way to do this for all 2,000 MKIDs on a feedline. Thankfully, many of the signals involved can be combined and extracted using filters, Fourier methods, and digital signal processing approaches.

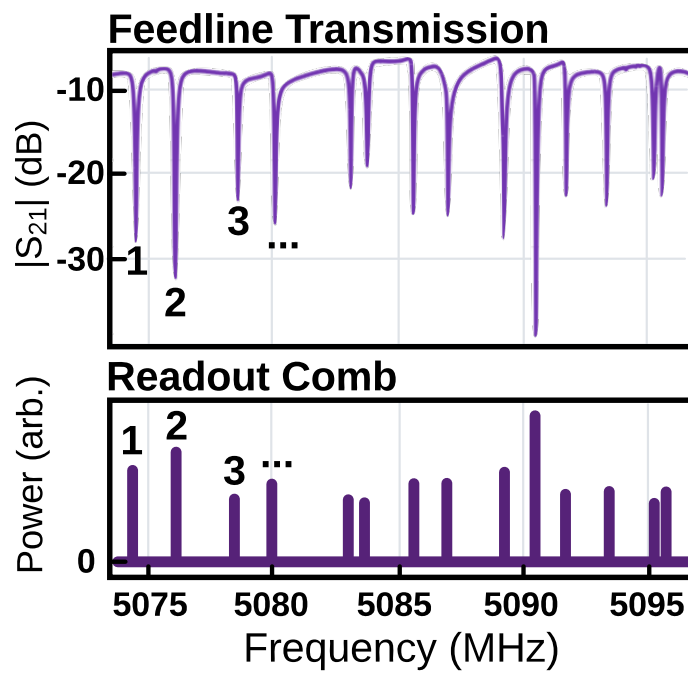


Figure 1.9: Overview of signal generation and readout strategy for a frequency-division-multiplexed feedline. Top: MKID feedline transmission. Each dip is a different resonator and the first three are numbered. Bottom: Corresponding DAC waveform in Fourier space that would be used to read out this feedline. Every MKID has a corresponding tone in the readout comb.

As one example, we can combine all MKID readout frequencies into a single waveform which is the superposition of all the readout tones. This readout comb can be stored in digital memory and played by a digital-to-analog converter (DAC). A sample feedline transmission and corresponding DAC output waveform is shown in Fig. 2.2.

On the input side, an analog-to-digital (ADC) converter is used to sample the frequency comb. Tracking the photon signals from all 2,000 resonators at this point requires highly-performant algorithms and custom digital signal processing (DSP). The technology of choice to implement DSP for this application and similar is a Field Programmable Gate Array which will be explained in the next section.

1.4 Field Programmable Gate Arrays

Field Programmable Gate Arrays (FPGAs) are custom digital computing platforms with many advantages for high-data-rate, low-latency, computation-intensive tasks. Unlike CPUs and GPUs, FPGA contain many I/O pins and are highly-parallel which allow them to ingest and process large data bandwidths. FPGAs are also fully customizable and re-programmable, allowing users to implement whatever specific algorithms they want and to update them as needed. An alternative to FPGAs is Application Specific Integrated Circuits (ASICs) which are also custom logic designs but they are specific to one task and are not re-programmable. Since they are not editable, ASIC designs require significantly more verification, simulation, and optimization. ASIC design cycles tend to be too long and costly for academic applications but many of the resources in an FPGA are actually embedded ASICs. The main goal in an FPGA design is choose the right resources and connect them in the right way such that the FPGA compiler can route all the wires and realize the circuit. In this section, I'll provide a lightning recap of FPGA design fundamentals to provide context for the digital readout design.

FPGA Resources

An FPGA has a set amount of computing resources which the user can connect in different ways to implement different algorithms. The type and number of FPGA resources are critical considerations for a FPGA design.

- DSP48 [58]

The DSP48 or simply “DSP” is the base resource used for digital signal processing. It includes an efficient structure which can be configured for adds, multiplies, and accumulation functions common to DSP algorithms. Depending on the specific design of the algorithm and the FPGA fabric clock speed, a DSP may be shared amongst multiple math operations. As a general rule of thumb, 1 DSP tile will be used per multiply per clock.

- LUT [59]

Look-Up-Tables can act as memory or logic. In logic mode, the LUT implements the equivalent logic circuit’s truth table. In memory mode, LUTs can be used as a distributed RAM and can cache small amounts of data and implement delay lines as needed throughout an FPGA design. This is one of the fastest but most scarce memory resources on an FPGA.

- BRAM [60]

Block RAM is a larger chunk of memory which can be used to store data. Each BRAM resource stores 18 Kbits. BRAMs have two read and two write ports which can be operated with different clocks/widths and can act as a true dual port memory. A typical use of BRAM is to store coefficients for a filter.

- URAM [61]

UltraRAM is similar to BRAM but is optimized and simplified for large data storage that does not need simultaneous read and write operations. UltraRAM cascades data, address, and control signals when connecting multiple URAM instances together, leading to a resource efficient implementation of a large memory. BRAM only cascades data lines and uses logic resources such as LUTs to connect multiple instances together. This optimization makes URAM a good choice for storing the output waveform used to drive the MKID resonators.

Timing

Another critical consideration for FPGA designs is timing. The digital circuits and logic operations in the FPGA rely on signals and systems driven by a clock. Many signals are synchronous and so must both arrive at a given logic gate within a narrow time interval or else the logic cannot function properly. If the timing constraint is not met, the design fails timing and will not work. Higher clock speeds and larger distances between operations can strain timing. In general, FPGA designers use the slowest clock speed possible and minimal resources so that the FPGA tools have lots of options when routing the signals and can find a solution that meets timing constraints.

1.4.1 System-On-Chip

Modern FPGA devices are increasingly including embedded ASICs to efficiently perform tasks ubiquitous in FPGA designs. The most prominent recent example is the System-On-Chip or SoC paradigm where a hardened CPU is embedded in the FPGA programmable logic. The SoC can run a full Linux operating system and is tightly coupled to the FPGA to facilitate control and data movement. Previously, developers had to implement a soft-core microprocessor, such as a MicroBlazeTM, using FPGA fabric

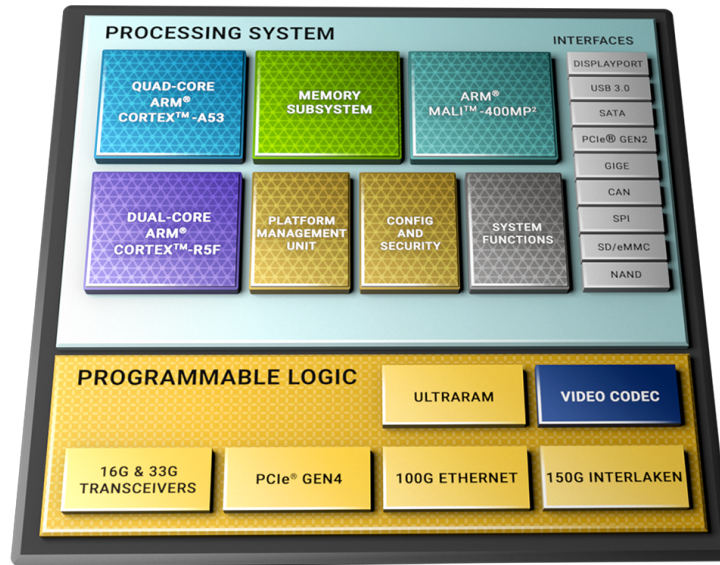


Figure 1.10: Overview of the AMD/Xilinx Zynq™ architecture. Figure is reproduced from [62].

resources; now, the SoC provides a more performant solution and reduces system cost, complexity, form factor, and power. The leading commercial SoC product is the AMD / Xilinx Zynq™ architecture which features an ARM core CPU embedded in the programmable logic (see Fig.1.10).

The communications, defense, and automotive industries continue to push advancements in high-performance, low-power, real-time computing platforms. Like multiplexed superconducting detector readout, these industries must also scale to handle larger data bandwidths in smaller form factors with less power. The theme is integration. Going forward, we will see how integrated electronics are enabling the next generation MKID digital readout platform and will hopefully continue to provide scaling advancements in years to come.

Chapter 2

Gen3 MKID Readout

Gen3 Preface and Attributions

This chapter is dedicated to the design, implementation, and performance of the new MKID digital readout system. This work is reproduced from an article submitted to AIP Review of Scientific Instruments. An arXiv pre-print is available [63]. The Gen3 MKID readout platform is the result of a collaboration with John I. Bailey, III, Aled Cuda, Nicholas Zobrist, and Benjamin A. Mazin. I would also like to acknowledge our collaborators at Fermilab: Gustavo Cancelo, Leandro Stefanazzi, Ken Treptow and Ted Zmuda who designed and fabricated the original IF board and provided DSP consultation for Gen3. I would also like to thank Jack Hickish, Mitch Burnett, and Dan Werthimer along with the rest of the CASPER collaboration for their early advice on FPGA-based signal processing approaches and for providing insight on RFSoc clocking and overall performance.

2.1 Introduction

Microwave Kinetic Inductance Detectors (MKIDs) are promising superconducting detectors with applications for highly-sensitive photon measurements such as those needed in quantum computing, biological imaging, and astrophysical observations [6, 64, 13, 14, 15, 65]. MKIDs are able to detect the energy and arrival time of single photons using the changes induced in the kinetic inductance of superconducting materials by incident photons [66]. This operating principle allows for single-photon counting with zero read noise or dark current across ultraviolet, optical, and infrared (UVOIR) wavelengths with microsecond timing precision [67, 68]. MKIDs also natively support multiplexing many detectors per readout line using frequency-division multiplexing (FDM)—a technique that will be discussed more in Sec. 2.2.2. FDM provides a feasible means of scaling MKIDs to large cryogenic arrays suitable for high-resolution imaging.

The focus of our group is to develop UVOIR MKIDs into fast, energy-resolving, single-photon-counting cameras for scientific imaging and spectroscopy. We have successfully deployed a twenty-kilopixel UVOIR MKID array in the MKID Exoplanet Camera (MEC) [41] and have demonstrated several high-contrast imaging astronomy results [42, 69]. We are developing an echelle spectrograph testbench [19] and are also pursuing science applications in quantum information and biological imaging.

Progress has been slowed in part by the current, second-generation MKID digital readout system (hereafter Gen2) which is excessively large, cumbersome, and power-hungry by modern standards [70]. Gen2 will not scale for larger ground-based arrays and is not suitable for future space-based missions. Gen2 also relies on obsolete hardware and tools, making it exceedingly difficult to update and to integrate modern advances in photon signal detection.

To alleviate these challenges, we have created the next generation MKID digital read-

out which is capable of reading out twice as many MKIDs per board with a fraction of the weight, volume, and power of the previous system. A key advancement is the migration to the Xilinx RFSoc platform with FPGA-integrated, high-speed analog-to-digital and digital-to-analog converters (ADCs/DACs). The integrated platform provides a dramatic reduction in power and device footprint that will enable future high-altitude missions and ease system scaling. We also leveraged modern FPGA programming tools including Vitis High-Level Synthesis (HLS), which synthesizes C++ code to FPGA-specific hardware description language, and Python Productivity for Zynq (PYNQ), which facilitates interacting with the FPGA through Python. These tools are more accessible to scientists and astronomers without specialized knowledge of FPGA design and create a system which is easier to maintain and upgrade.

In this work, we discuss the system design, implementation, and performance. We include a brief introduction to MKID readout in Sec. 2.2 to provide context for the requirements outlined in Sec. 2.3. Next, we describe our approach to system design in Sec. 2.4 and describe the FPGA implementation including our use of HLS and our timing closure strategy. We provide performance characterization for the system in loopback in Sec. 2.5 and include detailed studies of actual cryogenic MKID readout performance in Sec. 2.5.4. We conclude with a discussion of the system performance in Sec. 2.6 and describe future directions and next steps. The MKID digital readout described here is fully open-source and available on GitHub under a GPLv3.0 license¹.

2.2 MKID Readout

MKID readout is a rich topic with a large body of prior work dedicated to the underlying physics and improving detector performance [52, 53, 18]. Here we provide a

¹<https://github.com/MazinLab/MKIDGen3>




Figure 2.1: Simulated MKID frequency response magnitude (left), phase (center), and complex representation (right). The difference between the drive frequency and resonance frequency is shown in a pink to purple gradient on each plot. The optimal readout frequency is near the resonance frequency (pink).

brief review of core MKID readout signals and the strategy underpinning the digital design. We include a high-level description of MKID setup and operation to clarify goals, terminology, and analysis used in the remainder of this work.

2.2.1 MKIDs as Superconducting Resonators

Each MKID pixel is a superconducting LC resonator. Our designs target resonance frequencies in 4-8 GHz due to commercial availability of cryogenic low-noise amplifiers in this band. MKIDs are read out using a homodyne scheme where the signal of interest is a modulation on the readout tone and can be represented as a complex signal. This signal is acquired by down-converting the readout tone to 0 Hz. A simulated MKID frequency response is plotted in Fig. 2.1. The IQ loop (Fig. 2.1, right) contains all the information needed to determine the correct readout frequency, phase offset, and loop center coordinates for photon readout, and as such characterizing each MKID IQ loop is an important setup function for MKID readout systems.

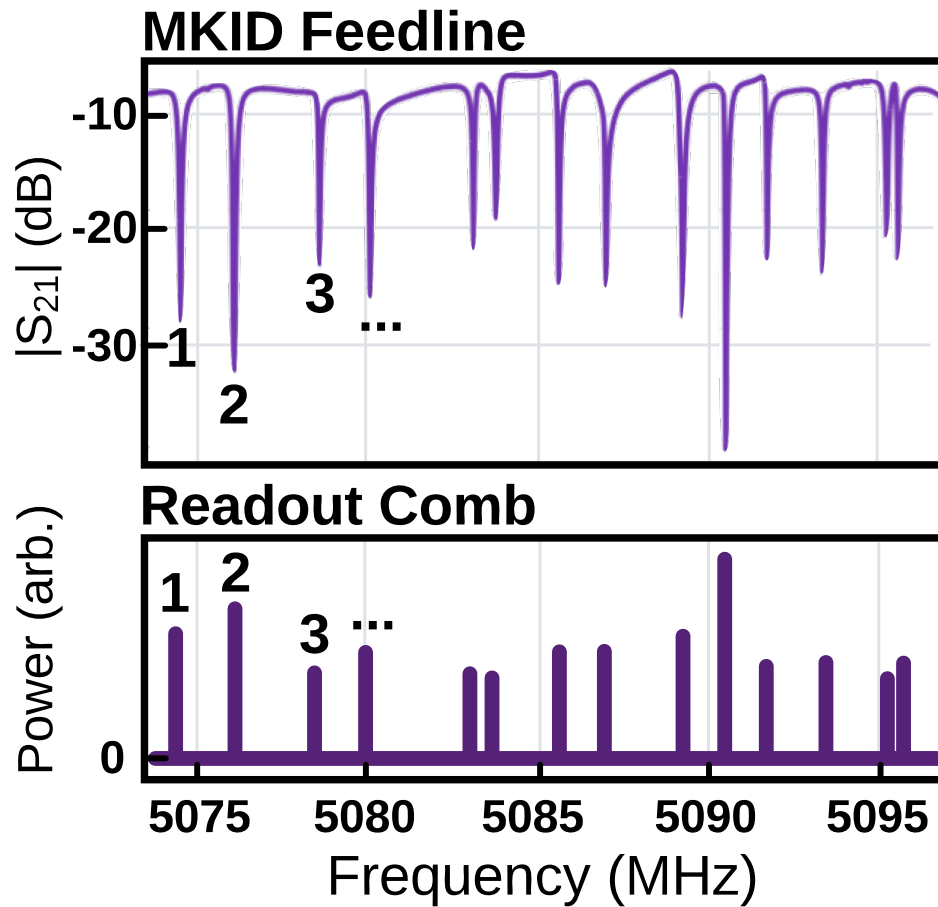


Figure 2.2: Overview of readout strategy for a frequency-division-multiplexed MKID feedline. Top: MKID feedline transmission. Each dip in transmission is a different MKID resonator. Bottom: Cartoon representation of corresponding DAC waveform in Fourier space that would be used to read out this feedline. Every MKID has a corresponding tone in the readout comb near its resonance frequency. The first three MKIDs shown in the feedline and their corresponding readout tones in the frequency comb are numbered.

2.2.2 Frequency-Division Multiplexing

MKIDs use frequency-division multiplexing to thread many superconducting resonators on a single readout line. This technique allows MKID readout signals to share a microwave feedline, conserving cryogenic heat load and promoting the creation of large arrays. A sample transmission from a real MKID feedline is shown in Fig. 2.2, (top). Each resonator is fabricated with a different resonance frequency, allowing each resonator to be addressed by its unique resonance. We target 2048 resonators with a spacing of 2 MHz but lose 10% to 40% of pixels due to resonator collisions and other fabrication process uncertainties. More details on MKID array fabrication are provided in references [71, 72, 54]. Each feedline represents a base unit for a full MKID array with each feedline being an imperfect copy of the others. In this work, references to multi-MKID readout refer to a single feedline with the understanding that a high-level control program can aggregate data from each feedline to create array-level data products.

2.2.3 MKID Feedline Setup

After fabrication, the precise locations of the pixels and their optimal readout parameters are unknown. The digital readout uses a reprogrammable local oscillator (LO) to sweep a comb of uniform tones across the 4-8 GHz feedline while collecting transmission data, similar to a vector network analyzer. The frequency sweep is repeated at different powers and the resulting IQ loops are used to determine the optimal readout frequency, power, phase offset, and loop center coordinates for each MKID. The readout powers and frequencies are used to generate a customized DAC output consisting of a superposition of the readout tones (see Fig. 2.2, bottom). The phase offsets and loop center coordinates are used to rotate and center each IQ loop in order to standardize the photon response curve. After the optimal readout tones and powers have been used to generate




Figure 2.3: Simulation of MKID photon readout. Left: the MKID resonator loop is shown in a pink to purple gradient representing the frequency difference from the resonant frequency. The light purple and dark purple lines show the trajectory the resonant frequency takes when a low energy and high energy photon hits the MKID, respectively. Right: The phase of the resonant frequency versus time for the two photon impacts. The depth of the phase pulse corresponds to the energy of the photon. The two photons were constructed to hit at different times for visual clarity.

the readout waveform and the IQ loops have been rotated and centered, the device is ready for photon readout.

2.2.4 Photon Readout

When a photon strikes the detector, the down-converted readout tone follows a trajectory through the IQ plane simulated in Fig. 2.3, left. To simplify readout, we compute the phase of the trajectory, $\tan^{-1}(Q/I)$, and use this one-dimensional signal to characterize the incident photon (see Fig. 2.3, right). The start of the phase pulse signifies when the photon hit the detector and the pulse height corresponds to the photon energy.

2.2.5 Resolving Power

Our ability to determine the energy of the incident photon is an important science metric. The resolving power, R , is defined as

$$R = \frac{E}{\Delta E} \simeq \frac{\lambda}{\Delta \lambda}, \quad (2.1)$$

where E is the photon energy and ΔE is the energy uncertainty. The relation can be expressed in wavelength-space assuming ΔE is small, but in this work we do all calculations and analysis in energy-space. R is fundamentally limited by the detector design but can also be negatively impacted by phase noise in the probe tone from readout electronics.

R is characterized using a series of lasers with known wavelength. Each laser produces counts with varying pulse heights corresponding to a single photon energy. The full-width-half-max of the distribution of pulse heights is the energy uncertainty. The measured mean pulse height is fit to known laser energy to produce a mapping from phase to energy. Full details of R characterization are available in reference [18].

2.3 Requirements

The next generation MKID digital readout system must perform several key functions and adhere to science-driven noise requirements while meeting hardware constraints. The overarching goal is to serve as the digital readout system for a 20,000+ pixel UVOIR MKID instrument. The system is expected to both set up and read out the MKID array. Key objectives, requirements, performance criteria, and constraints are explored in the following sections.

2.3.1 Key Objectives

Several objectives and philosophies played into the system design. Most importantly, we wanted to maximize the number of detectors read out while minimizing the weight, volume, and power of the readout electronics. We also wanted to create a system that is straightforward to migrate to future platforms. This will allow us to continually take advantage of advances in analog-integrated electronics driven by commercial industry.

Another goal was to create a system that is easy for scientists and others without prior FPGA design experience to use, maintain, and improve. In service of this goal and to maximize future collaboration and scientific utility, the system design favors open-source tools and uses more-accessible high-level-synthesis wherever possible.

2.3.2 Functional Requirements

There are several functions the system must perform to effectively read out an MKID feedline. To begin the setup and readout process, the system must be able to generate a 4-8 GHz readout waveform with up to 2048 run-time-programmable frequencies and powers. Each of the 2048 user-defined MKID readout frequencies constitutes a readout channel. For every independent channel, the system must be able to do the following continuously and in real-time:

- Down-convert the channel.
- Apply a translational and rotational coordinate transform to the down-converted signal.
- Calculate the channel phase.
- Apply a custom, user-defined filter to the channel phase.
- Continuously monitor the channel phase, searching for and recording photon events: the arrival time, pulse height, and channel number.

For setup and calibration, the system must be able to capture the following data in finite intervals:

- The full data rate input waveform.
- The IQ loop of every channel.

Table 2.1: Summary of signal and noise performance requirements and achievements.

Signal or Noise Metric	Requirement	Req. Section	Achieved	Perf. Section
Readout Tone Frequency Resolution	≤ 7.813 kHz	Section 2.3.3	7.813 kHz	Section 2.4.3
Readout Tone Amplitude Control	≤ 1 dB	Section 2.3.3	0.25 dB	Section 2.4.1
Channel Cross Talk	≤ -30 dB	Section 2.3.3	≤ -30 dB	Section 2.5.2
Intermodulation Spurs	≤ -30 dB	Section 2.3.3	≤ -30 dB	Section 2.5.2
Resolving Power, Single MKID	$R \geq 7$ at 808 nm	Section 2.3.3	$R = 7$ at 808 nm	Section 2.5.4
Resolving Power, 2048 MKIDs	$R \geq 4$ at 808 nm	Section 2.3.3	$R = 4$ at 808 nm	Section 2.5.4
Timing Resolution	≤ 1 μ s	Section 2.3.3	1 μ s	Section 2.4.3
Absolute Timing Precision	≤ 1 μ s	Section 2.3.3	≤ 500 ns	Section 2.4.3

- The phase time stream of every channel.

These captures must each be able to occur simultaneously with photon monitoring as they can provide valuable debug information during an observing run. To debug photon triggering, the system must also be able to simultaneously record a subset of user-defined IQ and phase time streams for a short interval around each photon trigger.

2.3.3 Signal and Noise Requirements

Signal and noise requirements for UVOIR MKID readout can be complex, often involving application-specific trade-offs. In this section, we list the the main performance criteria that are common to most use cases and drive design decisions. Requirements are summarized in Table 2.1.

Tone Generation

Readout tone frequency precision will impact the detector’s perceived responsivity and linearity. Errors in tone amplitude affect responsivity as well as the phase noise floor. The magnitude of these effects depends on device characteristics including the resonator quality factor and asymmetry. We demonstrated 7.813 kHz frequency resolution and 1 dB amplitude control is sufficient to avoid noticeable performance degradation in the Gen2 readout [70].

Cross Talk and Intermodulation Distortion

Cross talk between channels and spurious signals from intermodulation distortion (IMD) cause line noise in the phase time stream that can distort pulse heights and ultimately degrade resolving power. The worst line noise comes from cross talk in the cryogenic MKID device and images generated by gain and phase imbalance in the IQ analog signal chain. Both noise sources can produce spurious signals 20 to 30 dB down from the read out tones. Presently, the IQ mixers are necessary to interface the cryogenic and room temperature electronics and so we require the digital readout crosstalk and IMD to be no worse than -30 dB as referenced to the readout tones.

Resolving Power

The resolving power is an important figure of merit used to characterize the detector's ability to discern the energy of the incident photon (see Sec. 2.2.5). Pinning down a resolving power requirement specific to the digital readout can be challenging because many other factors including the detector design and fabrication, the cryogenic system and signal chain, the setup and calibration quality, and the off-line data analysis methods can all impact the measurement. To develop our requirement, we bound our system between two extremes: a best-case scenario and a worst-case scenario.

The best-case scenario performance should about match our in-lab, analog-based, MKID readout which can only read out a single MKID at a time and is nominally used to provide feedback for detector design and fabrication cycles. This system uses a dilution refrigerator (20 mK) with a first-stage, quantum-noise-limited amplifier (TWPA), commercial analog electronics, and Python analysis package with extensive device modeling to measure the best-possible resolving power of individual MKID devices. Zobrist *et al.* [67] demonstrated a resolving power of 8.9 at 808 nm on a PtSi MKID [73] with

this system. However, the analysis methods included a wavelength-specific, 500-point+ matched filter and resonator-specific, two-dimensional, quadratic coordinate transform [74]– both of which are not feasible for an FPGA-based instrument at our target scale. To compensate for the difference in analysis methods, we relax our best-case scenario requirement to $R \geq 7$ when reading out a single PtSi, array-style MKID with 808 nm photons.

The worst-case scenario performance should be better than or equal to what was achieved by the previous Gen2 system in the field. Gen2 was deployed to the summit of Mauna Kea with the MKID Exoplanet Camera (MEC) which features an adiabatic demagnetization refrigerator (90 mK) with a first-stage, Low Noise Factory HEMT amplifier. The Gen2 Python code uses a machine-learning-based approach to semi-automate the set up and read out of 1024 pixels per board. MEC achieved a median R of 4 at 850 nm across the PtSi MKID array. This sets our worst-case scenario requirement to $R \geq 4$ at 808 nm² for a PtSi array-style MKID when all 2048 readout channels are active in Gen3.

Timing Precision

Millisecond photon arrival resolution is sufficient for most astronomical sources with the notable exception of pulsar timing studies, which benefit from microsecond precision [75]. In quantum photonics applications, microsecond or better timing precision is preferred [76]. In UVOIR MKID devices, the pulse relaxation time is typically on the order of microseconds. As demonstrated in the Gen2 system [70], a sample rate of 1 microsecond is sufficient to resolve the pulse peak to within a microsecond and provides enough bandwidth to facilitate signal processing techniques aimed at improving the resolving

²We set our requirement at 808 nm despite the fact MEC was characterized at 850 nm due to difference in availability of laser sources between Subaru Telescope and our lab at UCSB.

power.

2.3.4 Constraints

In addition to meeting the performance goals, the new digital readout must adhere to several practical constraints.

Cryogenic Interface and MKID Bandwidth

The UVOIR MKIDs are designed with resonance frequencies in the 4-8 GHz band because this is where cryogenic low-noise amplifiers and other specialized components are commercially available. The readout system must be able to supply, sample, and process signals in this bandwidth.

Hardware

Due to budget and time restrictions, we are unable to support designing or commissioning a custom analog or digital board and must use existing technology. There are presently no suitable commercial options for the analog interface between the DAC output and cryogenic input, so we must re-purpose existing Gen2 IF boards [70].

Data Rate

The system must be able to support a count rate of 5,000 counts per second on every pixel. The system should not drop any registered photon data less than this level and should notify users if photons are being dropped.

Setup Time and Stability

In the past, MKID instruments have been deployed with adiabatic demagnetization refrigerator (ADR) cryogenic systems. ADRs are well-suited for observatory environments, but impose a time constraint because they can only keep the MKID array cold for a finite hold time. The readout system must be able to perform all setup and calibration steps within this hold time with enough time left over to conduct science observations. In practice, the readout must be able to set up an array from scratch in a few hours. Similarly, the readout must be stable and able to continuously observe and record photons for a typical observing night (at least eight hours).

2.4 System Design

The Gen3 readout system is based on a Xilinx, analog-integrated FPGA (RFSoc). The integrated platform removes the need for external ADCs and DACs, greatly reducing power consumption and device footprint. At present, the fastest available integrated ADC sampling rate is 5.9 GSPS in the Xilinx RFSoc DFE device [62]. To directly sample the 4-8 GHz readout band, the DFE ADC would have to be operated in the third Nyquist zone which is not naively supported [77]. With direct RF seemingly still out of reach, we opted to build the system around the slower Gen3 RFSoc part which is available in a small, cost-effective academic board: the RFSoc4x2.

The RFSoc4x2 data converters generate and sample the MKID readout tones using quadrature sampling. An intermediate frequency (IF) mixing board is used to translate the quadrature signals to and from the microwave readout band. The FPGA programmable logic implements the main digital signal processing pipeline and calibration capture functions. The RFSoc4x2 processing system is used to command and configure the programmable logic as well as peripherals on the IF board. All system communication




Figure 2.4: Photo of MKID digital readout during MKID readout. MKIDs (not shown) are housed at 90 mK inside the dilution refrigerator. Light is passed through the window on the dilution refrigerator front plate. Microwave multiplexed readout signals are driven and acquired by the RFSoc and IF boards via a cryogenic signal chain (not shown). Photon data is saved to network storage over 1 Gigabit Ethernet.

Table 2.2: Gen2 vs. Gen3 system specifications and capabilities.

Specifications	Gen2: ROACH2	Gen3: RFSoc4x2
Dimensions	24"x12"x4"	10"x6"x2"
Weight	10 kg	2 kg
RF Bandwidth	2 GHz	4 GHz
MKID Pixels	1024	2048
Data Rate	40 MiB/s	80 MiB/s + 16 GiB/s
Cost	\$10/pixel	\$3/pixel
Power	175 mW / pixel	25 mW / pixel
Network	1 GbE	1 GbE + 100 GbE
Design System	ISE/Simulink/CASPER	Vivado/PYNQ/HLS
Control	Python 2.7	Python 3.11 / ZeroMQ

and data transfer happens over 1 Gigabit Ethernet.

The readout platform can be run via a Jupyter notebook, acting as a lab-based MKID diagnostic platform, or via a client/server software that facilitates full feedline setup and multi-board photon readout [78]. Subsystems are detailed in the following sections.

2.4.1 Hardware

A picture of the Gen3 readout hardware and experimental setup is shown in Fig. 2.4. Migrating to the RFSoc4x2 platform resulted in a dramatic reduction in the weight, volume, and power of the readout electronics. In addition to using only one fifth the power as the previous Gen2 system, the data converters are twice as fast, allowing us to double the number of detectors read out per board. A to-scale picture showing the size reduction from Gen2 to Gen3 is shown in Fig. 2.5. Key specifications between the two systems are tabulated in Table 2.2.




Figure 2.5: Scale photo comparison of previous system (Gen2) ROACH2 hardware (left) and Gen3 RFSoc4x2 hardware (right). The same IF board is used in both systems with a few modified components. The Gen3 RFSoc4x2 is significantly smaller, lighter, and more power efficient and is capable of reading out 2x as many detectors as the ROACH2.

IF Board and Carrier

The IF board is responsible for translating the ± 2 GHz readout waveform to the 4-8 GHz microwave band. It is a modified version of the same board used to convert signals in the Gen2 system [70]. To accommodate the RFSoc's increased bandwidth, we swapped the 630 MHz (LFCN-630+) anti-aliasing filters on the I and Q DAC output *IF to RF* paths for the 2 GHz version (LFCN-2000+). We also swapped 800 MHz (LFCN-800+) and 1800 MHz (LFCN-1800+) anti-aliasing filters on the I and Q *RF to IF* ADC input signal paths with the 2 GHz (LFCN-2000+) and 3.8 GHz (LFCN-3800+) versions, respectively.

The input and output RF ports each have two programmable attenuators which together can be adjusted from 0 to 63 dB in steps of 0.25 dB. This allows adjusting the waveform power going into the fridge, to optimally drive the MKID resonators, and into the RFSoc, to utilize the full ADC dynamic range. The IF board receiver path has an amplifier chain totaling 88 dB of gain to boost the small cryogenic signals. The board also features a TRF3765 programmable LO and complex mixers which can be used to sweep RFSoc-driven frequencies to different parts of the microwave readout band.

The IF board is powered by a simple 12 V carrier board (see blue board in Fig. 2.4 under the IF Board). The programmable components are controlled by the carrier's Arduino Nano which is commanded by the RFSoc over USB.

RFSoc4x2

Both RFSoc4x2 DACs and two of the ADCs are run at 4.096 GSPS to create a quadrature sampled readout waveform with frequencies in ± 2 GHz. The FPGA implements a high-throughput digital signal processing pipeline to down-convert each readout tone and monitor the phase.

The RFSoc integrated ARM core runs Ubuntu-flavored Linux and serves as the command and control hub for each feedline readout. The processing system (PS) and programmable logic (PL) each have 4 GiB of available DDR4 SDRAM. The PL DDR4 serves as a data buffer for setup and calibration captures while the PS DDR4 is used for continuous photon capture and pixel monitoring. All FPGA capture, configuration, and control signals are managed by a Python package running on the PS. The PS is also able to program the IF board LO and attenuators to facilitate sweeps.

The RFSoc4x2 platform has several quirks which impacted system design. The first was data converter synchronization. We found the phase mismatch between the DAC tile clocks disrupted the quadrature phase relation of our signals enough to cause -20 dB image tones. This did not meet our -30 dB IMD requirement (Sec. 2.3) and forced us to implement Multi-Tile Synchronization (MTS) [79]. MTS suppresses the image tones to -40 dB but adds complexity and constraints to the clocking architecture.

The other quirk was the platform's default MPSoc clock configuration led the 1 Gigabit Ethernet to perform at roughly one-third capacity. We found the current 4x2 board support package (BSP) configures the full power domain main clock slower than the low power domain switch. We modified the BSP to correct this and recovered near line rate 1 GbE performance. More information on the Zynq UltraScale+ MPSoc is provided in reference [62].

2.4.2 FPGA Design

The FPGA design followed a similar strategy to the Gen2 system but with several optimizations to manage double the number of channels and speed up data capture. The MKID readout waveform is computed in software and written to a large FPGA memory buffer where it can be replayed from the DACs in a loop. After the waveform passes

through the feedline, it is sampled by the ADCs which feed into an intensive real-time processing pipeline. The digital signal processing steps are described in the following section.

Digital Signal Processing

An overview of the key digital signal processing (DSP) steps is shown in Fig. 2.6. After the 4 GHz waveform is injected into the programmable logic, it must be channelized into 2048 MKID channels. Due to fabrication uncertainties, the channelizer must be able to down-convert 2048 channels from arbitrary position in the microwave band. We use a two-stage oversampled polyphase filter bank (OPFB) channelizer followed by a direct digital-down-converter (DDC). The OPFB produces 4096, 2 MHz coarse channels which overlap 50%. This allows every MKID readout tone to pass through the channelizer un-attenuated. The OPFB was highly-optimized as it is the most resource-intensive step and it is the subject of its own publication [80].

After the OPFB, the Bin Select core selects the 2048 user-defined channels which contain MKID readout tones, copying bins with multiple tones as needed, and feeds them to the DDC. The DDC multiplies each channel by the complex conjugate of the readout tone to do the final down-conversion. The DDC core is also responsible for applying each channel's custom coordinate transform. The DDC complex multiplier includes an optional phase offset which can individually rotate each MKID loop. The core also implements a complex subtraction which serves to center each MKID loop. After each readout tone has been fully down-converted and the coordinate transform has been applied, each channel is low-pass filtered to remove other MKID readout signals that may fall in the same OPFB bin. The channels are decimated to produce 2048, 1 MHz fine channels each with one MKID readout tone in the center.

The next step is to convert each channel to phase by evaluating $\tan^{-1}(Q/I)$. This

operation results in 2048 phase time series each sampled every 1 microsecond. To further reduce phase noise, we apply a custom matched filter to every channel. The matched filter generation process is summarized in Fig. 2.7. For every pixel on the feedline, we average unfiltered laser photons from the middle of the energy band to characterize the detector signal. We also estimate the channel phase noise by averaging the noise power spectral density when there are no photons arriving. With the measured signal and noise, we construct a matched filter for every channel.

After the matched filter enhances photon signals and suppresses channel-specific noise, we are finally ready for photon triggering. The trigger behavior is shown for a filtered pulse in Fig. 2.8. Once the pulse crosses the threshold, the channel triggers. The minimum phase value is continually updated for a number of samples specified by the holdoff parameter. After the holdoff counter expires, the photon energy is recorded as the minimum value in the window and the channel is free to trigger again. The photon time, energy, and channel are recorded to disk.

2.4.3 Implementation

The FPGA design was implemented using Vivado 2021.1 in an IP integrator-based project flow. A system block design is shown in Fig. 2.9. Resource utilization is summarized in Table 3.2 and device area utilization is highlighted by block in Fig. 2.11. The implementation is notable for realizing an intensive signal processing pipeline, requiring a large percentage of the FPGA resources at high clock rates, while using high-level synthesis tools. In this section, we will describe our approach to implementation, starting with our choice of tools.




Figure 2.6: Summary of key digital signal processing steps for frequency-division-multiplexed MKID readout. Signals are first channelized with a two-stage oversampled polyphase filter bank channelizer (OPFB) and direct digital-down-converter (DDC). The MKID tones are then filtered by custom matched filters before a trigger records photon time, energy, and location across all 2048 channels.

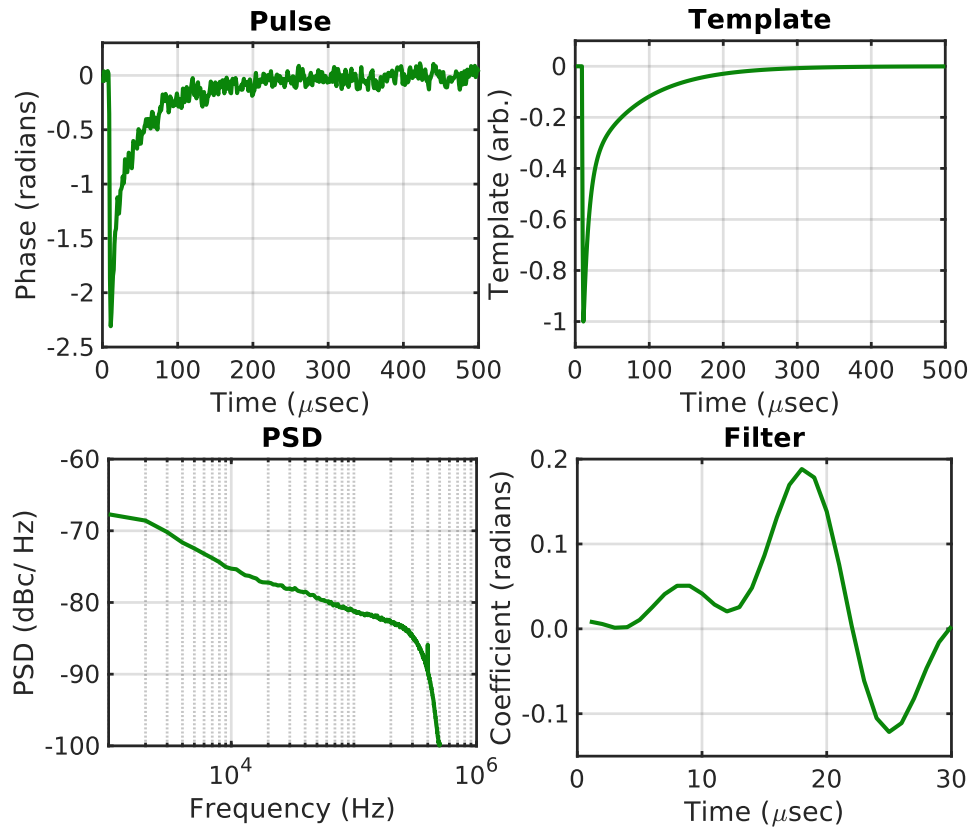


Figure 2.7: Overview of MKID-specific matched filter generation. Photon pulses are recorded for each MKID pixel (upper left). Pulses are averaged in time to remove random noise and create a pulse template (upper right). Phase time streams from segments without photon pulses are averaged in the Fourier domain to estimate the phase noise in each channel (lower left). A Wiener filter is computed from the template and the noise and a final low pass filter is applied to produce each channel's matched filter (lower right).

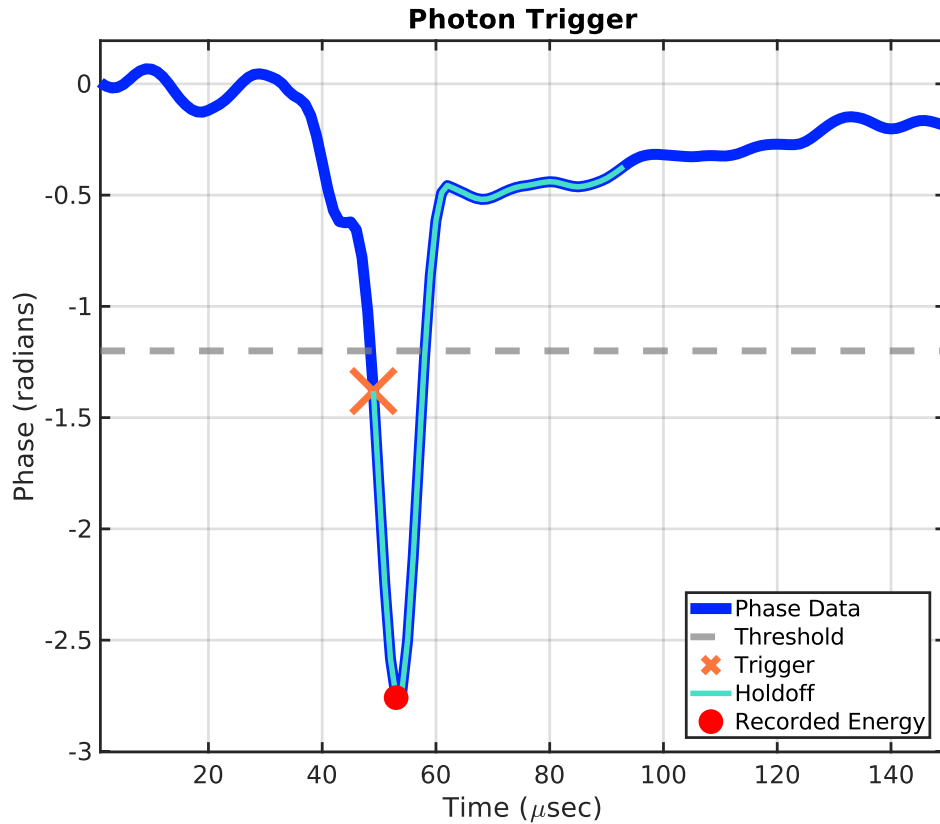


Figure 2.8: Trigger operation on a filtered pulse. Once the phase crosses below the pixel's threshold, the channel triggers. During the trigger state, the minimum phase value is updated for some number of samples specified by the holdoff value. The minimum phase value in the holdoff window is recorded as the photon energy. After the holdoff expires, the channel is eligible to trigger again. Both the threshold and holdoff settings are calibrated per pixel.

Tools

All custom IP blocks, barring a few exceptions discussed below, were implemented using Vitis High-Level Synthesis (HLS) which synthesizes high-level C/C++ code to low-level hardware description language (HDL). HLS has the advantage of being more accessible to scientists without experience in hardware description language (HDL) and can be more flexible and easier to port to new hardware. However, it is notorious for using more resources and increasing timing strain. We developed several strategies for HLS usage to improve resource utilization and timing closure:

- Small, single task, HLS blocks are preferable to complex blocks.
- Clean, small internal functions can significantly improve generated HDL.
- Partitioned temporary variables in unrolled loops is preferable to automatic inference within a loop.
- The `ap_ctrl_none` pragma directive can significantly improve control and logic optimization.
- Manually picking bits instead of using the `DATAPACK` pragma directive can be vital.
- It is better to place memory resources such as BRAM and URAM manually using the IP generator in Vivado than to infer them in HLS.
- External AXIS to AXI conversion is necessary for continuous writes.

To help close timing in the full design, we used Vivado Intelligent Design Runs (IDR). IDR uses ML-based strategy predictions and incremental compile to help close timing in Vivado in tight designs [81]. We used quality of result (QOR) suggestions and IDR-implemented runs to identify optimizations and achieve timing closure in the full design. Subsystem-specific strategies are discussed in context below.

Waveform Replay

The DAC Replay block streams a repeating waveform consisting of 2^{19} complex samples from a look-up-table (LUT) to the 4.096 GSPS DACs. The LUT size and DAC sample rate achieve 7.813 kHz frequency resolution for each superimposed tone in the readout comb. The table stores the 32-bit complex samples in a 2 MiB URAM buffer. An external URAM was used in Vivado as opposed HLS inference because HLS would not properly cascade the URAM blocks—a necessary optimization for reducing logic routing resources.

Channelizer

The OPFB, Bin Select, and DDC subsystems implement the channelizer. The OPFB is implemented using two HLS blocks that route and reorder data, sixteen parallel Xilinx FIR cores that implement the filter, and the Xilinx SSR FFT block exported from System Generator / Model Composer. Full implementation details are provided in Smith *et al.* [80].

The Bin Select block caches eight copies of the preceding 4096-point FFT frame and routes them into eight parallel output lanes in any order (i.e., the pathological case of 2048 resonators in a single channel is supported). The channel order is run-time user programmable via AXI4-Lite.

The DDC uses external BRAMs to store phase increments, offsets, and centers which are used with an internal cosine LUT to down-convert the channels and apply the complex rotation and translation coordinate transform. The DDC values are channel-specific and run-time user programmable over AXI4-Lite. A Xilinx FIR filter implements a low-pass and decimate operation to produce the fine channels.

Convert to Phase and Filter

The complex, demodulated signals are converted to phase through parallel Xilinx CORDIC [82] cores which implement $\tan^{-1}(Q/I)$. The matched filters are implemented using re-programmable Xilinx FIR compiler cores [83]. The filter coefficients are run-time user programmable over AXI4-Lite.

Photon Event Trigger

The trigger is user-programmed at run-time with 2048 thresholds and holdoffs unique to each channel. The threshold is encoded as an 8-bit signed value allowing the phase threshold to be set with 0.02 radian precision. The holdoff value indicates the number of samples (microseconds) until a subsequent trigger is permitted. The holdoff takes integer values between 8 and 254. The maximum value of 254 corresponds to 254 microseconds which exceeds the typical filtered pulse recovery time (see Fig. 2.8).

Photon Capture

Photon capture is implemented using a pair of 800 KiB buffers. The processing system reads from one buffer while the other records photon data. The buffer is rotated either when full or when two photons arrive a user-defined interval ($\sim 0.0005 - 1$ s) after the first buffered photon.

Postage Capture

The postage capture system is capable of recording a 127-microsecond IQ time series window of up to 8000 total trigger events across any of 16 user-selected channels.

Capture Subsystem

The capture subsystem manages data capture from various places in the signal processing pipeline to facilitate MKID feedline characterization and setup. An overview of the capture hierarchy is shown in Fig. 2.10. The system is backed by the PL DDR4 which provides 4 GiB of SDRAM storage. It is implemented using parallel HLS blocks that select user-requested groups of channels or, in the case of ADC capture, combine I and Q signals. A small, open-source RTL core [84] is used to translate the AXI4-Stream interface to AXI4 transactions because HLS generated AXI4S-AXI4M interfaces do not support continuous write at the gigabyte level. AXI4 glue logic is used to cross clock domains and buffer the data before it is written to DDR4. The memory interface generator (MIG) is clocked at 333 MHz with a 512-bit bus, providing around 16 GiB/s offload bandwidth.

Routing congestion caused by I/O pins to the PL DDR4 being proximate to the RFDC I/O pins caused trouble closing timing in the capture subsystem. A successful IDR run suggested the problem could be alleviated by constraining the MIG to the right side of the chip to create open lanes for the DAC output routing. The resulting square Pblock is visible in the placement view (see Fig. 2.11).

Timestamps

Photon timestamps are provided by a time keeping core that supports several different modes. In basic operation, timestamps are generated using a 19-hour, 1 microsecond counter. All RFSoc boards are synchronized using the pulse-per-second (PPS) input linked to a GPS signal, providing a pulse each time the second rolls over. The current second as determined by UTC accessed from an NTP server by the processing system is synchronized with the PPS signal. The 36-bit, UTC-based timestamp is associated with

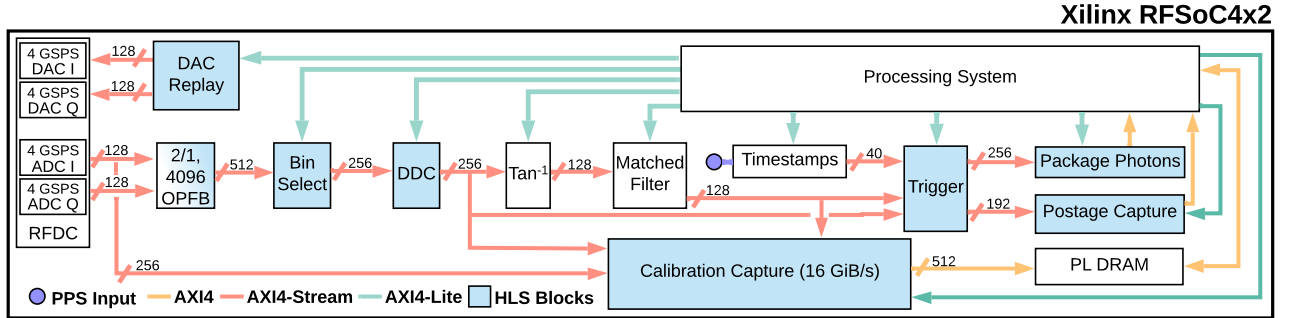


Figure 2.9: Programmable logic system block design. Two 4.096 GSPS DACs produce the output waveform via a URAM lookup table. Two 4.096 GSPS ADCs feed data to a signal processing pipeline which monitors the phase of 2048 channels for photon events. A calibration capture subsystem routes streams to the PL DRAM, providing snapshots of data that can be used for setup and calibration tasks, such as identifying the MKID resonant frequencies, optimal readout powers, and phase biases. The processing system is used to configure blocks over AXI4-Lite and acts and the system command and control. Blocks implemented in Vivado/Vitis High-Level-Synthesis are shown in blue. Protocols and data-widths are labeled for select pathways.

Table 2.3: Resource utilization.

Resource	Number	% of ZU48DR
CLB	36012	68%
LUT (Logic)	79809	19%
LUT (Memory)	98368	46%
BRAM	262	24%
URAM	64	80%
DSP	475	11%

each phase sample in the trigger block and recorded when a photon event is triggered. The resulting system provides $1 \mu\text{s}$ timing resolution. Absolute timing error between boards was measured to be less than 150 ns^3 .

2.4.4 Software

The project software is based primarily on the Python Productivity for ZYNQ (PYNQ) framework. PYNQ is an open-source, AMD/Xilinx maintained project that facilitates

³PPS signal was generated by a Stanford Research FS725 10 MHz Rubidium Frequency Standard.

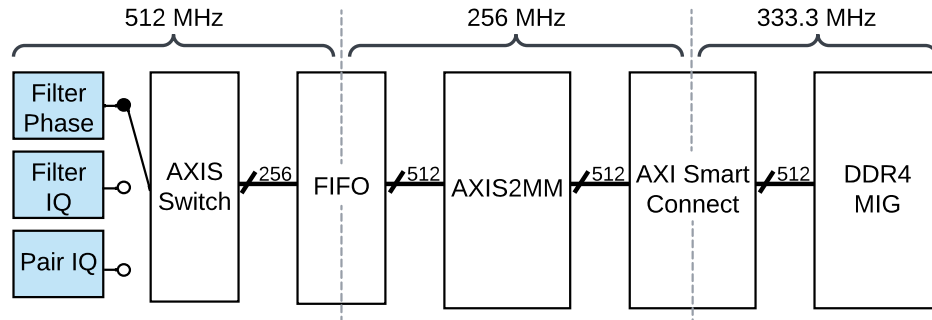


Figure 2.10: Capture subsystem block design. Three HLS blocks (blue) package different streaming data formats from different places in the signal processing pathway. An AXIS switch selects the data stream to be captured. A dual-clocked FIFO feeds the selected data stream to an open source RTL core which produces AXI transactions for off-chip communication. The data is sent to the PL DDR4 via an AXI Smart connect and the Memory Interface Generator (MIG). The system is able to capture up to 4 GiB at 16 GiB/s.

interacting with and programming a ZYNQ architecture-based FPGA through Python [85]. PYNQ provides a Linux-based image, Jupyter notebook server, and extensible Python packages all running on the MPSoC. The `mkidgen3` package includes Python drivers for the FPGA design IP, MKID-specific setup and calibration functions, and diagnostic plots. In the following subsections we detail our adoption and adaptation of the PYNQ framework to support MKID readout.

Deviations from the PYNQ Image

The `mkidgen3` system image has two main modifications from the Ubuntu 22.04-based PYNQ 3.0.1 stock image. First, we reserved a segment in the device tree for the PL DDR4. Editing the device tree allowed more robust, full integration of the PL DDR4 with the PYNQ memory model and recommended `pynq.allocate` method. The second adaptation was to patch the `xrfdc` Bare Metal C RFDC driver to include MTS-related functions. These deviations are documented in the main project repository.

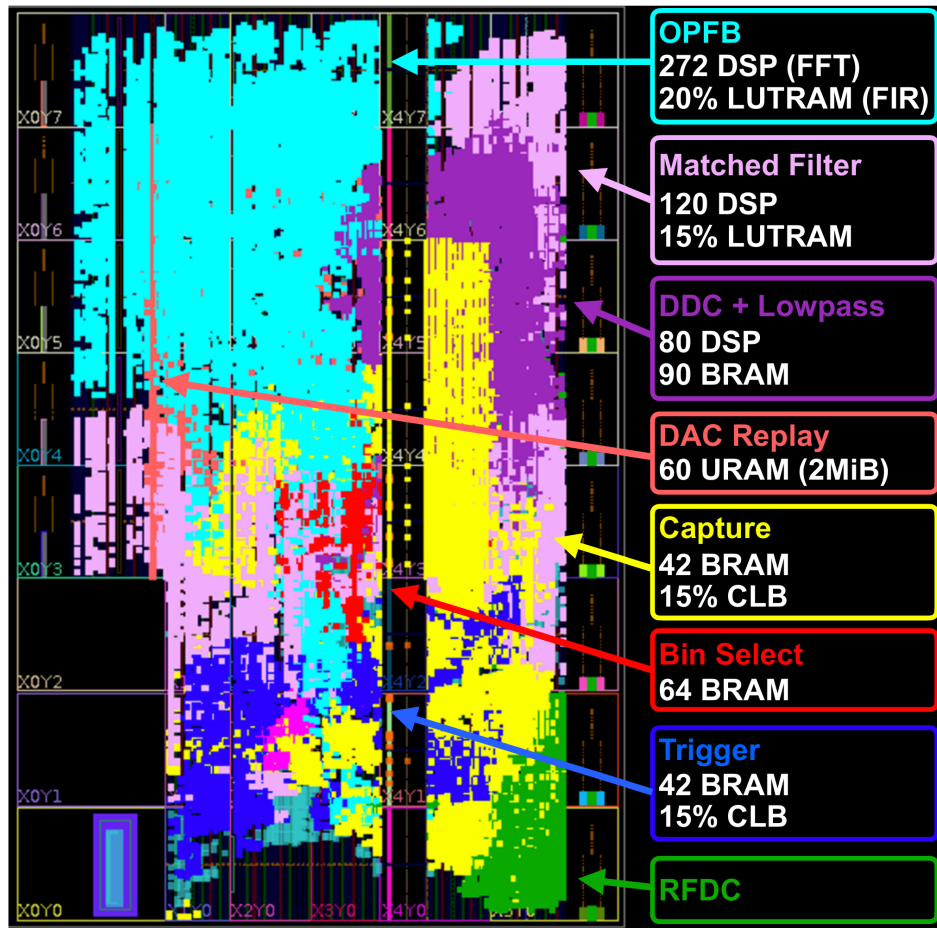


Figure 2.11: Overview of the RFSoc ZU48DR device chip area utilization. Key signal processing subsystems are highlighted with different colors. Select FPGA resources are summarized for the highlighted systems on the right.

Python Drivers

PYNQ provides a simple methodology to create Python drivers for IP. The framework parses the Vivado-generated hardware hand-off (`.hwh`) file, effectively auto-discovering IP, and provides the user with a `pynq.Overlay` object that can be used to command and control design IP through its associated register map. We use this approach to write Python drivers for all configurable IP in the design. The PYNQ drivers can be used directly to set up and read out MKIDs using the on-board Jupyter notebook server or through remote Python execution, i.e., over `ssh`.

Server Architecture

The software ecosystem also includes high-level client and server software that facilitate array-level set up, capture, and observations and is intended for use in a multi-board MKID instrument environment. Commands are based on the `zmq` request-reply pattern with each RFSoc running a feedline server that schedules capture requests from the client and guarantees all captures occur with the required FPGA state. Resulting capture data and status updates are broadcast via the `zmq` pub-sub pattern, allowing any network device to monitor photon or engineering data. At the time of writing, the server is still undergoing testing and will be fully detailed in a future publication.

2.5 System Performance Characterization

The Gen3 MKID readout system has been successfully demonstrated in lab and meets all performance requirements outlined in Sec. 2.3. The system was characterized in three configurations:

1. **RFDC Loopback:** The RFSoc4x2 I and Q ADCs are directly connected to the

I and Q DACs, respectively, with short, matched cables.

2. **IF Loopback:** The IF board *RF Out* is connected to the IF board *RF In* with a 30 dB fixed attenuator in-between.
3. **MKID Measurement:** The full system is connected to a cryogenic UVOIR MKID device (see Fig. 2.12).

In all configurations, a Stanford Research FS725 10 MHz rubidium frequency standard was used to synchronize the RFSoc4x2 and IF board clocks and provide the PPS signal.

2.5.1 Multi-MKID Readout Approximation

A pseudo-random comb was used to simulate the simultaneous readout of 2048 MKIDs. The comb consists of a superposition of readout tones where each tone is randomly placed in every other 2 MHz coarse channel in the ± 2 GHz readout band for a total of 2048 approximately-evenly-spaced tones. The tones are designed to be no closer than 300 kHz to their 2 MHz channel edge which mimics the OPFB overlap region and simplifies channel assignment in the FPGA. The resulting waveform is similar to what we find with modern UVOIR MKID devices and provides realistic conditions for evaluating multi-MKID readout performance.

2.5.2 Output Comb Performance

A 2048-tone, uniform-amplitude, pseudo-random-frequency (see Sec. 2.5.1) comb was generated and played out the DACs to simulate an MKID readout waveform. The output tone and image powers for all frequencies in IF loopback and in RFDC loopback are shown in Fig. 2.13. In RFDC loopback, the comb is fairly uniform, showing only a few dB variation across the ± 2 GHz readout band. The image tones are at most -30

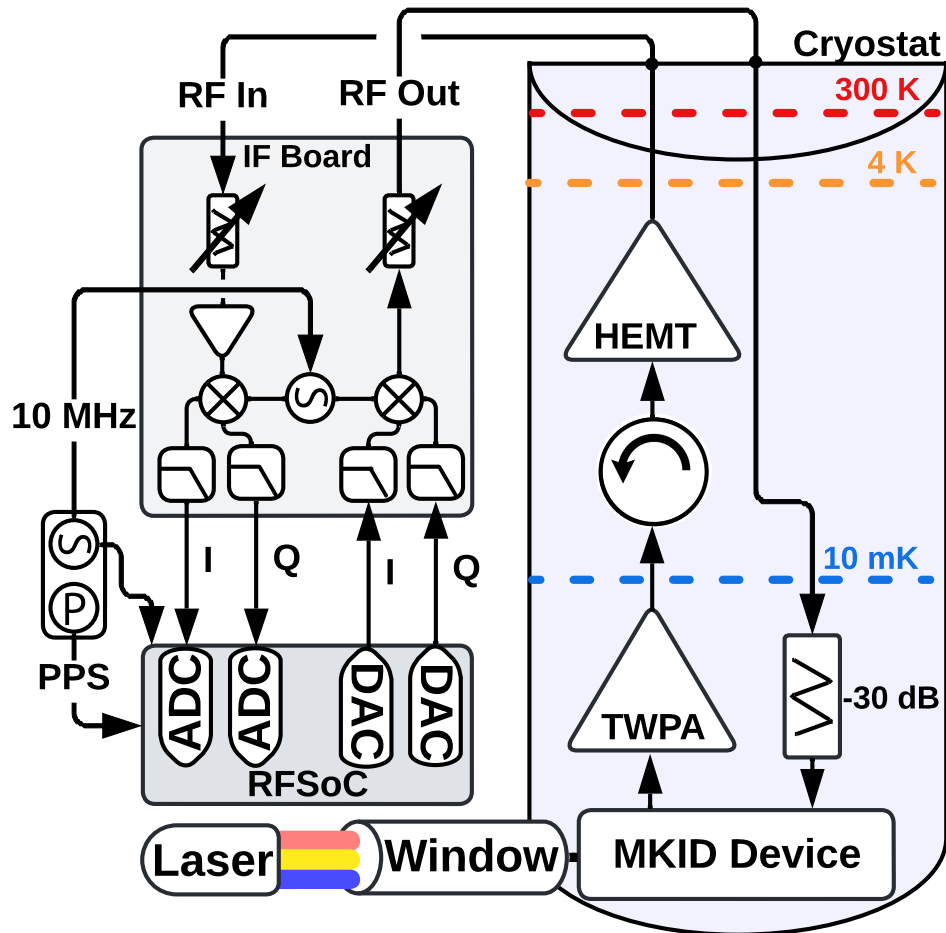


Figure 2.12: Experimental setup for system performance characterization. RFDC loopback measurements were taken with the RFSoc DAC *I* and *Q* connected to the RFSoc ADC *I* and *Q* ports via external SMA cables (not shown). IF loopback measurements were taken with the IF board *RF Out* connected to the IF board *RF In* with a 30 dB fixed attenuator in-between the ports (not shown). MKID measurements were taken with the setup shown and include the cryogenic wiring and amplifier chain. During MKID testing, lasers are shone through a window in the cryostat to illuminate the device with known-energy photons.

dB at the band edges and meet our performance requirement (Sec. 2.3.3). In the IF loopback, the output tone power is more attenuated and non-uniform across the band. The attenuation at the edges of the 4 GHz band are caused by imperfect roll-off in the IF board analog anti-aliasing low-pass filters (see Sec. 2.4.1). Ripples in the tone power are caused by impedance mismatches which launch reflections and creating standing waves in the IF board. The real-valued frequencies (0, 2 GHz) are more attenuated than the imaginary-valued frequencies (-2 GHz, 0), likely due to excess attenuation in the IF board I signal path. We hypothesize the manual rework of the anti-aliasing filters in-house resulted in poor microwave hygiene and wide variation between signal paths in the IF board, leading to reduced performance.

2.5.3 Phase Noise

The phase power spectral density from a down-converted readout tone is shown in Fig. 2.14. The phase noise floor is fit with the red dashed line and is approximately -80 dBc/Hz. There is some line noise present around 150 kHz, possibly from electromagnetic interference in the lab or a nearby image tone. Line noise can also be generated from rounding errors in the DSP pipeline which uses truncation as the default rounding mode to conserve FPGA resources. Truncation can introduce a DC bias, for example in an OPFB channel, which may be up-converted to an image tone during the DDC multiply operation.

Fig. 2.15 shows the PSD-floor fit for all 2048 channels in the pseudo-random-tone waveform. The PSD floors follow a pattern inverse to the IF loopback tone power measured in Fig. 2.13, with the real-valued frequencies showing more noise. This pattern is indicative of the same issue with the IF board microwave hygiene and poor transmission performance in the *I* signal path.




Figure 2.13: Power level of 2048 driven readout tones taken with the system in RFDC loopback (solid purple) and IF loopback (dashed blue). Power level of non-ideal images tones created by IQ imbalance are shown for IF loopback (dotted gold) and RFDC loopback (solid red). As expected, the IF loopback traces show more tone power loss and higher image tones due to IQ imbalance and impedance mismatches. The effects of the IF board anti-aliasing filters are also visible as attenuation at the edges of the band in the IF loopback tone trace.




Figure 2.14: Phase noise power spectral density in a single channel. No matched filter is applied. The data was collected with the system in IF loopback with 2048 pseudo-random tones. The white noise floor is fit with the red dashed line.




Figure 2.15: Phase noise power spectral density noise floor fit in all channels. The data was collected with the system in IF loopback with 2048 pseudo-random tones running. The white noise floor was fit and recorded for each channel as shown in Fig 2.14. The phase noise drifts ~ 10 dB across the band, likely due to IQ imbalances in the system.

2.5.4 MKID Readout

The Gen3 readout system was also characterized in lab using a real MKID and a series of lasers. These studies provide a concrete means to study the system’s energy-resolving capability and help disentangle effects of readout-specific noise on scientific utility.

Experimental Setup

The MKID device was a PtSi-on-sapphire, array-style chip⁴ with the same design and manufacturing process used for the device in MEC and in Zobrist *et al.* [67]. The MKID was cooled in a dilution refrigerator (20 mK) with a first-stage, quantum-noise-limited parametric amplifier [86]. Lasers at 405.9 nm, 663.1 nm, and 808.0 nm illuminated the MKID through windows in the the dilution refrigerator. The experimental setup is shown in Fig. 2.12.

A single 4.5 GHz MKID was characterized, biased, and prepared for photon readout using the `mkidgen3` Python package on the RFSoc-hosted Jupyter Notebook. The resonator was set up manually using a series of sweeps and was ultimately biased at the correct frequency using the IF board programmable LO, and the correct power using the IF board *RF Out* programmable attenuator. This technique is not suitable for setting up multiple MKIDs on one feedline because the LO and programmable attenuator provide global frequency and power offsets; however, it provides a decent approximation for the full-setup algorithms and server utility that are currently are under development. To characterize the extra readout noise generated during multi-MKID readout, excess readout tones were pseudo-randomly generated (see Sec. 2.5.1) and propagated through the MKID device and full readout chain.

For maximum flexibility in the data analysis, the raw phase time series was captured

⁴Full fabrication details are provided in [73].




Figure 2.16: Summary of matched filter performance in the case of 2048 pseudo-random tones. Top: Phase time series taken with a red 663.1 nm laser illuminating the MKID without (light pink) and with (dark pink) the 30-tap matched filter applied in software. Bottom: Resolving power achieved without (left) and with (right) the matched filter applied. The resolving power, R , is annotated above each energy histogram. Different wavelength photons are shown in different colors. The matched filter reduces noise in the phase time series and greatly improves resolving power.

with a unity matched filter (no effect) and the matched filter and trigger algorithms were applied in software. The matched filter and trigger performance were verified on the FPGA for a subset of experimental settings.

2.5.5 Matched Filter Performance

The matched filters are chiefly responsible for boosting the signal-to-noise ratio in the phase time series and improving resolving power. The test MKID's 30-tap filter was computed as described in Fig. 2.7 using pulse data from the middle (red, 633.1 nm)




Figure 2.17: Single MKID response to different wavelength photons as more pseudo-random tones are added to the readout comb. The resolving power, R , is annotated above each energy histogram. Different wavelength photons are shown in different colors. The resolving power degrades slightly as more tones are added.

laser wavelength. The detector resolving power was characterized with and without the matched filter. In each case, the readout system used a 2048-tone, pseudo-random comb to simulate the noise environment of a full-feedline readout. Results are summarized in Fig. 2.16 and show a dramatic improvement in resolving power with the matched filter applied.

2.5.6 Resolving Power Results

The detector resolving power, R , was measured with different numbers of pseudo-random tones in the readout waveform to characterize the best-case and worst-case scenario performance. Results are summarized in Fig. 2.17. In the case of only one readout tone in the waveform, we find excellent resolving power that is consistent with Zobrist *et al.* [67] and meets our best-case scenario performance requirement (Sec. 2.3.3). In the worst-case scenario where we are driving all 2048 tones, we see some degradation in the resolving power but the results are within the acceptable limits set by the Gen2 MEC deployment (Sec. 2.3.3).

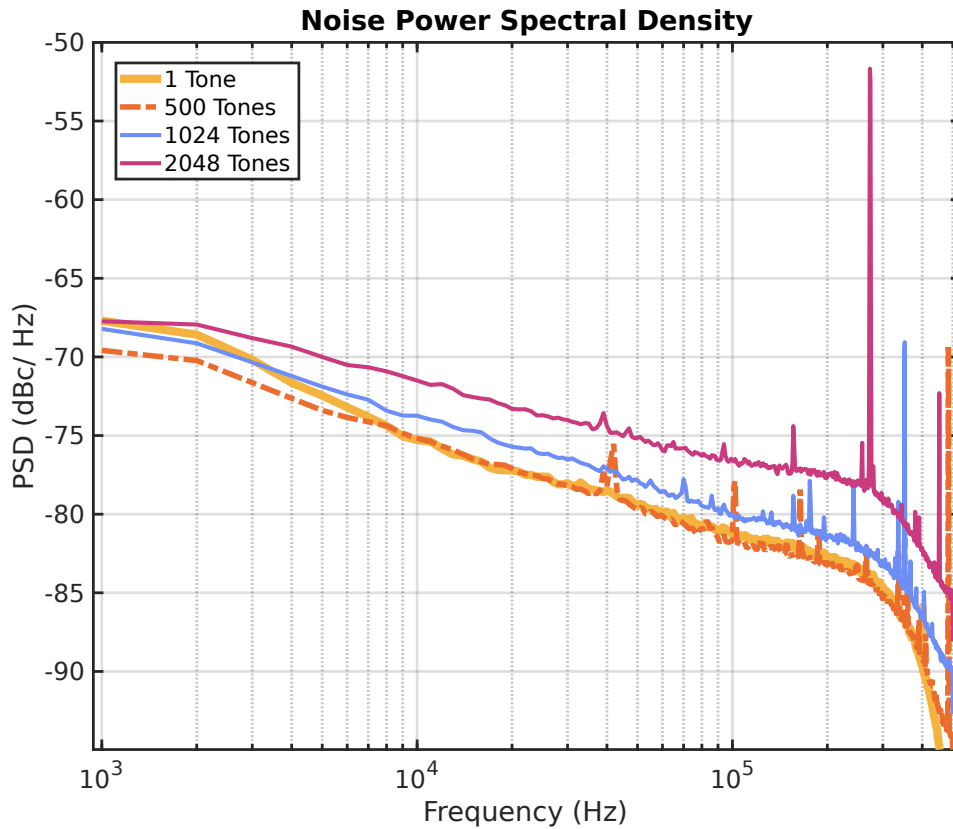


Figure 2.18: Phase noise power spectral densities taken with the system connected to the fridge as shown in Fig. 2.12. The data was collected with the readout tone biased to the optimal MKID resonance frequency and power and includes MKID device two-level system noise as well as noise from the cryogenic amplifier chain. As more tones are driven, the phase noise floor and spurious line noise both increase. When 2048 tones are driven, a large tone appears around 280 kHz at the -52 dB level; it may be an image tone or intermodulation product from one or more of the other driven tones.




Figure 2.19: Bits of DAC dynamic range available per readout tone as a function of the number of tones assuming all tones have equal amplitudes and random phases. The light blue dashed line represents \log_2 of the average individual tone quantized amplitude. The actual number of bits is the floor of this result and is shown in solid navy. There's a large fall-off in DAC dynamic range between a single tone and 500 tones. At our maximum number of channels, there are six bits available per tone.

2.5.7 Resolving Power Degradation with Tone Number

While the resolving power requirements are met, future MKID devices are expected to have better intrinsic resolving power. Understanding the nature of the resolving power degradation is useful for planning system upgrades and drafting requirements for future systems. The MKID phase noise power spectral density for each number of tones is shown in Fig. 2.18. The data was collected with the readout tone biased to the optimal MKID resonance frequency and power and includes the MKID device two-level system noise as well as noise from the cryogenic amplifier chain. Fig. 2.18 shows that as the number of tones goes up, both the noise floor and line noise increase.

We hypothesize there are three main ways in which increasing the number of tones in the waveform can raise system noise and decrease resolving power.

1. Reduced DAC Dynamic Range

For best noise performance, the DAC maximum voltage output is typically mapped to the maximum value of the readout waveform. The IF board *RF Out* programmable attenuator can be used to reduce the power level to the optimal MKID drive power (around -100 dBm) as needed. As more tones are added, the maximum value of the superimposed readout waveform grows and there is less available DAC dynamic range per readout tone. The reduction in number of bits available per tone is shown in Fig. 2.19. This effect increases the quantization noise in each channel.

2. Reduced ADC Dynamic Range

Similar to the DAC, there is also a reduction in ADC dynamic range per tone as more tones are added. Typically, the IF board *RF In* programmable attenuator is adjusted so the maximum waveform value will be close to the maximum ADC input

voltage. The more tones superimposed in the waveform, the fewer bits available per tone. A similar relation to Fig. 2.19 applies except the RFSoc4x2 ADC maximum dynamic range is 12-bits, not 14.

3. Intermodulation Effects

Finally, there are spurs, intermodulation products, image tones, and reflections all caused by the proliferation of readout tones in the analog domain. These signals may also impact dynamic range in the digital domain, for example, more noise power will cause more signal resolution to be lost during truncation rounding.

We devised two additional studies to isolate the ADC and DAC dynamic range effects and determine the leading contributor to the resolving power degradation we observe with increasing tone number. In the first study, a single tone was played from the DAC but the waveform was artificially made to use different DAC dynamic ranges in accordance with the values expected for different numbers of tones (see Fig. 2.19). For each DAC dynamic range setting, the *RF Out* IF board programmable attenuator was adjusted to ensure the MKID device was still being driven at the correct level. The ADC input attenuator was fixed conservatively at a level consistent with what would be compatible with 2048 tones. In this way, we were able to probe the effects of the loss in DAC dynamic range absent intermodulation products and changes in ADC dynamic range utilization. To study the effects of ADC dynamic range, we again played a single tone from the DAC but this time at a fixed dynamic range, again conservatively in-line with what would be available with 2048 tones. The IF board *RF In* programmable attenuator was adjusted to purposefully lower the ADC dynamic range used to what would be available per tone for different numbers of tones. Results for these studies is compared with the data collected for the multi-tone experiment in Fig. 2.20. The experiments show that the degradation in DAC and ADC dynamic range, and consequently the increase in quantization noise, does not

significantly impact device resolving power. In all cases, the measured resolving power is similar to what was measured in the best-case scenario with a single tone and maximal ADC/DAC dynamic range utilization. We conclude that spurious tones, intermodulation products, image tones, and reflections caused by the multitude of readout tones is the primary factor driving resolving power degradation with increasing number of tones.

2.6 Discussion

The Gen3 system characterized in this work is capable of reading out 2048 UVOIR MKIDs across 4 GHz of bandwidth using an RFSoc and IF board. While the system meets all performance requirements outlined in Sec. 2.3 and provides a major improvement in UVOIR MKID readout scalability, there are several areas that could be improved and questions that warrant further investigation. Performance-based trade-offs and recommendations are discussed in context below.

2.6.1 IF Board IQ Imbalance and Images

We use quadrature sampling in the RFSoc along with IQ mixers in the IF board to access the 4-8 GHz UVOIR MKID band. Quadrature sampling has the benefit of reducing the required data converter speed by half in order to access a given band. At present, this technique is necessary to reach our desired readout band using an RFSoc platform. Unfortunately, quadrature sampling comes at the cost of being very sensitive to gain/phase imbalance between the two quadratures which results in undesired image tones. These imbalances can be corrected at a single frequency but the calibration complexity grows significantly with bandwidth and in practice it is not feasible to calibrate more than about 400 MHz of instantaneous bandwidth [87].

In the RFSoc, using Multi-Tile Synchronization [79], we achieve enough consistency

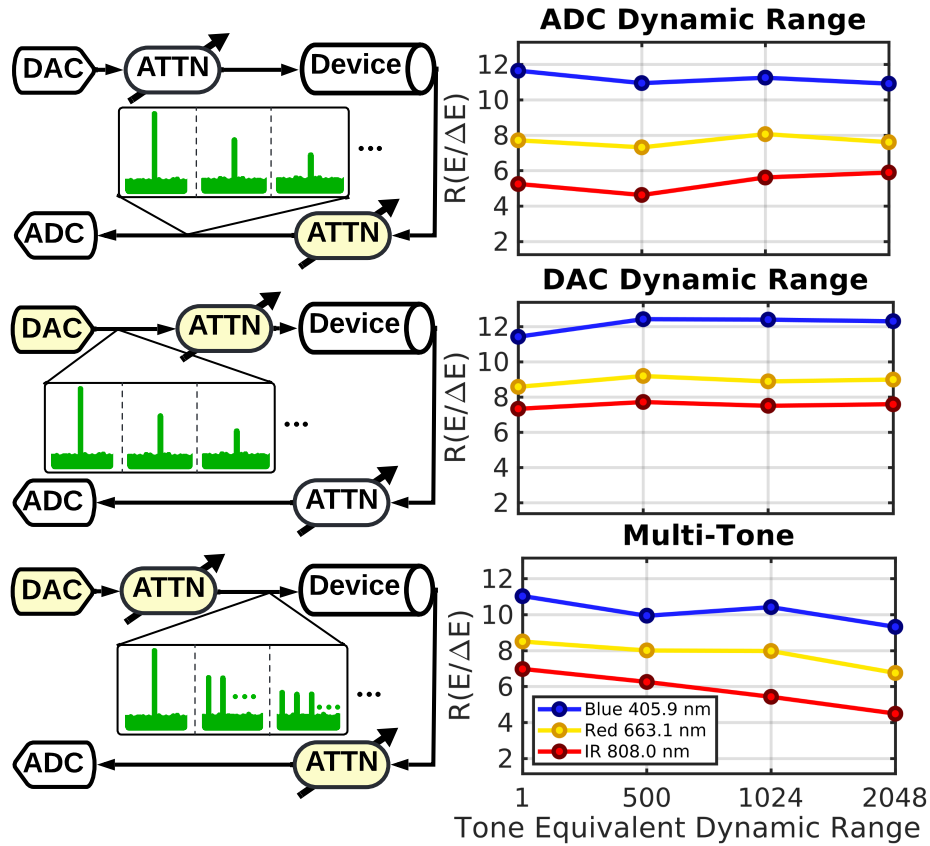


Figure 2.20: Three experimental setups (left) and results (right) intended to disentangle the ADC and DAC dynamic range from intermodulation product effects on detector resolving power. In each experiment, changing components are highlighted in yellow. (Top) A single readout tone is output by the DAC at a dynamic range equivalent to that available when driving 2048 tones (see Fig. 2.19). The ADC programmable attenuation is increased to lower the ADC dynamic range to that available per tone with 1, 500, 1024, and 2048 tones. The ADC dynamic range does not appear to have a large impact on resolving power. (Middle) The ADC input attenuator is fixed at a level compatible with 2048 tones while the DAC outputs a single tone with various dynamic range utilization as indicated by the number of tones. The DAC programmable attenuator is changed with each step to compensate for the reduced waveform amplitude and keep the readout tone power at the correct bias level for the MKID device, isolating the DAC dynamic range. The DAC dynamic range does not show a large impact on R at the levels tested. (Bottom) The impact of driving more tones on resolving power. The DAC and ADC programmable attenuators are adjusted with each number of tones to use the full dynamic range of both data converters. Histograms for this study are also shown in Fig. 2.17. We observe some degradation in R with increasing tone count, suggesting that intermodulation products and image tones arising from the multitude of driven tones are more damaging to resolving power than the increase in quantization noise from data converter dynamic range constraints.

between the I and Q paths that the image tones are within the -30 dB requirement (Sec. 2.3.3); however, this is not the case when the IF board is included (see Fig. 2.13), with images rising to within -25 dB and even -15 dB in some cases with respect to the driven tones. As discussed in Sec. 2.5.2 and Sec. 2.5.3, the IF board anti-aliasing filters were manually reworked in lab, resulting in poor signal integrity and imbalance between the I and Q paths (see Fig. 2.13 and Fig. 2.15). While the current IF board prototype provides a means to measure MKIDs in the lab, it does not meet the performance criteria across the full readout band and is likely not suitable for instrument deployment. With a modest redesign and professional PCB manufacturing, we expect the IF board loopback measurement can achieve ≤ 30 dB image tones and uniform -80 dBc / Hz phase noise floor across the band.

2.6.2 Direct RF Sampling

Even better than redesigning the IF board would be to do away with it entirely and use a direct RF approach. This eliminates the need for quadrature sampling with separate I and Q signal paths, potentially greatly improving RF signal integrity, readout phase line noise, and ultimately detector resolving power. Current RFSoc devices do not have the data converter speed required to sample 4-8 GHz directly in either the first or second Nyquist zone. It may be possible to use the 5.9 GSPS DFE device [62] at the maximum rate in the third Nyquist zone but more testing is needed to verify performance. Another approach is to use the second Nyquist zone (2.95-5.9 GHz) and modify the MKID readout band, though this would require significant device re-design and may not be compatible with commercial, low-noise cryogenic amplifiers required to boost the device signal. A faster RFSoc may be available soon in the Versal AI RF part which is advertised as having 36 GHz of analog bandwidth (4 GHz IBW per channel) and would likely be able

to directly access the full readout band using the first or second Nyquist zone [88].

With direct-RF-capable platforms on the horizon, the Gen3 system design prioritized platform-upgradability. The modular approach disentangles I/O from core DSP blocks and is backed by a `C Makefile` and Vivado `tc1` scripts that should greatly simplify re-targeting the design to a new RFSoc. The HLS DSP blocks expedite updates by exposing bit-widths, DSP parallelization, and other mutable parameters as top-level `C` directives. We've already seen some success in this approach; we were able to quickly re-target our original ZCU111-based design to the cost-effective, RFSoc4x2 platform shortly after its release. Going forward, these techniques are expected to continue helping us quickly leverage advances in the RFSoc product line.

2.6.3 UVOIR MKID Resolving Power

Resolving power is a critical UVOIR MKID capability. As discussed in Sec. 2.3.3, the measured resolving power is a conglomeration of many factors and it can be challenging to identify performance requirements that can meaningfully inform design trade-offs. Several recent superconducting detector digital readouts have derived noise requirements from the cryogenic signal chain [70, 50]. Fruitwala *et al.* [70] showed the Gen2 UVOIR MKID readout system IF loopback phase noise floor was around -90 dBc/Hz and was lower than the expected phase noise floor contributed by the first-stage, low-noise cryogenic HEMT amplifier (\simeq -85 dBc/Hz) [70]. Here, the measured phase noise floor in the Gen3 system is 5-10 dB higher than this calculated HEMT noise floor (see Fig. 2.14 and Fig. 2.15). This is not surprising given Gen3 runs twice as many tones in twice the bandwidth as compared to Gen2. We showed increasing the number of tones raises the system noise in Sec. 2.5.7 and the larger bandwidth is expected to harbor more intermodulation products, reflections, and other spurious signal distortion.

The HEMT phase noise floor is not even the relevant cryogenic amplifier limit in this work because we used a quantum-noise-limited TWPA for the first-stage. Integrating a TWPA before the HEMT lowers the effective cryogenic amplifier noise temperature from around 5 K to close to 0.5 K [67], further increasing the gulf between the expected phase noise contributions of the cryogenic signal chain and the Gen3 readout electronics. Despite having higher phase noise than the cryogenic amplifiers, Gen3 achieves respectable resolving power—competitive with dedicated analog electronics in the case of one tone [67] and comparable to Gen2 in the case of 2048 tones [89].

Our results support the idea that the resolving power is predominately limited by the detector itself. Zobrist *et al.* [67] showed the addition of the first stage TWPA in the analog readout system did not improve resolving power as much as the reduction in phase noise would have predicted, indicating a detector limitation. We see a similar pattern in this work where the phase noise in Gen3 is higher than what was achieved in Gen2, yet the resolving power is comparable.

This detector resolving power limitation has been attributed to a variable amount of the incident photon energy escaping into the non-photosensitive detector substrate, leading to a variable detector response. This problem, known in some literature references as “hot phonon escape” [90, 20, 91, 92, 68], is not captured in the phase noise power spectral density and is not improved by reduced phase noise in cryogenic or room temperature electronics. Zobrist *et al.* [68] fabricated a UVOIR MKID using a bilayer process that reflects energy back into the photosensitive region, recovering $R = 20$ at 814 nm. Unfortunately, this bilayer process has not been demonstrated with large, array-style MKID devices and may be impracticable due to difficulties with indium processes. Solving hot phonon escape is an active area of UVOIR MKID design and fabrication research, but for now it presents a fundamental limitation in resolving power.

Our results coupled with previous work suggests UVOIR MKID readout systems

should strive to achieve detector-limited resolving power as opposed to holding the phase noise under the contributions of the cryogenic amplifiers. With the addition of quantum-noise-limited TWPAs, this requirement becomes challenging and may come at the cost of increased system complexity and decreased scalability only to have the resolving power limited by the detectors in the end. In this work, we achieve detector-limited resolving power in the best-case scenario with one tone but we see degradation when running all 2048 tones through the readout chain (see Fig. 2.17). The best-case scenario performance is useful as a tool in lab to characterize devices and inform design and fabrication updates; however, going forward we will strive to achieve detector limited resolving power at scale. Such a solution may require direct RF and or integration of more complex signal processing techniques such as tone tracking [50], a resonator-based coordinate transform [74], and a longer matched filter (Fig. 2.7).

2.7 Conclusion

In conclusion, we have built a new MKID digital readout based on the RFSoc platform. The system fits twice as many MKID pixels onto a board that can be run with a fifth of the power of the previous system. Gen3 is also less than half the size and a tenth the weight of the Gen2 system. Gen3 utilizes high-level tools which simplify interacting with and programming the FPGA and make it easier for scientists to use, maintain, and upgrade. The system was demonstrated in lab and achieved high-fidelity, detector-limited resolving power in the case of few resonators and showed minimal degradation in R with more tones. Experiments showed the degradation is linked to intermodulation products, image tones, and spurs in the analog domain which may be improved with future iterations of the IF board or a switch to a direct RF approach.

Presently, we are continuing to update and improve the digital domain. We plan

to adjust the rounding mode to an unbiased algorithm and include intelligent scaling to maximize dynamic range in the programmable logic. We are also working on new algorithms and techniques to expand the taps in the matched filter for improved noise rejection. Overall, we believe this readout system provides a maintainable, scalable platform to work towards megapixel MKID arrays and future space-based deployment.

Chapter 3

Oversampled Polyphase Filter Bank

Preface

The Oversampled Polyphase Filter Bank (OPFB) is the most resource-intensive, digital-signal-processing step in the MKID FPGA design. While Gen2 had an OPFB FPGA implementation, it was too inefficient to handle the 2x data rate needed for Gen3 and did not make timing at our required FPGA fabric speed.

The work presented in this chapter was a collaboration with John I. Bailey, III, John Tuthill, Leandro Stefanazzi, Gustavo Cancelo, Ken Treptow, and Benjamin A. Mazin and was originally published in IEEE Open Journal of Circuits and Systems in a special section on circuits, systems, and algorithms for beyond 5G and toward 6G [80]. I am sincerely grateful to our Fermilab colleagues, who were amazing hosts for a few productive weeks—especially Leandro Stefanazzi, who shared his experience and expertise in polyphase filter bank design and implementation. I am also particularly grateful to John Tuthill, who answered an email from a stranger and provided invaluable insight into the inner workings of this algorithm. Lastly, I would like to thank Ross Martin at Bit by Bit Signal processing for sharing his expertise on high-performance FFTs and providing

us with a power-efficient FFT that will enhance future design flexibility. This project includes an open-source Github project¹, which includes a Jupyter Notebook-based walk-through of the algorithm.

3.1 Introduction

Large data bandwidths are an important theme underpinning the development of new technologies including 5G/6G devices for wireless communications [93, 94], real-time video processing for self-driving cars [95], and microwave device arrays for quantum computing and astronomy [96, 97, 98, 99]. Digital signal processing systems built on Field-Programmable Gate Arrays (FPGAs) are well-suited to high-performance applications that demand rapid, real-time processing of wide data bandwidths as these devices uniquely offer powerful and highly customizable parallel computing architectures, large IO bandwidths, and low power consumption.

Despite the performance advantages, widespread FPGA adoption has been slowed by high engineering costs associated with programming, verifying, and maintaining FPGA systems [100, 101]. Traditionally, FPGAs have been programmed and verified using specialized Hardware Description Languages (HDL) and vendor entrenched FPGA development programs known only by professional electrical engineers. Xilinx², is addressing these drawbacks by creating new programming tools to increase FPGA accessibility and cost-effectiveness for industries including machine learning, RF engineering, and scientific instrumentation where the technical workforce is trained in mainstream software development languages such as C++ and Python and not HDL. The work presented here leverages two of these modern tools to create an adaptable, high-performance FPGA

¹https://github.com/MazinLab/RFSoc_OPFB

²Xilinx, Inc., 2100 Logic Dr., San Jose, CA.

design: Vivado High-Level Synthesis (HLS) [102], which synthesizes C++ code to HDL, and Python Productivity for Zynq (PYNQ) [103], which facilitates FPGA programming and testing through Python.

Beyond lowering the entry barrier through improved tools, Xilinx has increased FPGA utility for expanding RF sectors including 5G/6G wireless by incorporating wide bandwidth analog devices into FPGAs to create products such as the Xilinx RF System on a Chip (RFSoc). The RFSoc targeted in this work incorporates 4 GSPS ADCs and DACs into the Zynq Ultrascale+ architecture [104]. This advancement eliminates the need for separate analog devices and infrastructure circuitry reducing power consumption and device footprint around 50% [105, 106]. Consequently, RF-integrated SoC architectures have become the focus of our group working to develop readout schemes for frequency-division-multiplexed microwave resonator detectors and our collaborators and colleagues whose work includes 5G/6G wireless [93, 94], quantum computing [97, 96], and software-defined-radio [107].

3.1.1 Polyphase Filter Banks

Many applications targeting new RFSoc devices must divide the large, instantaneous RF bandwidth into sub-bands where data can be processed in parallel by the programmable logic. The process of translating sections of the incoming bandwidth to lower frequencies, filtering, and down sampling is known as channelization and is summarized in Fig. 3.1. Polyphase filter banks (PFBs) are a popular choice for the channelization algorithm because they minimize spectral leakage across sub-bands by efficiently applying a windowing filter prior to taking the discrete Fourier transform (see Fig. 3.2) [108, 109]. A derivation of polyphase filter bank structure first published in [110] is reproduced in the Appendix.

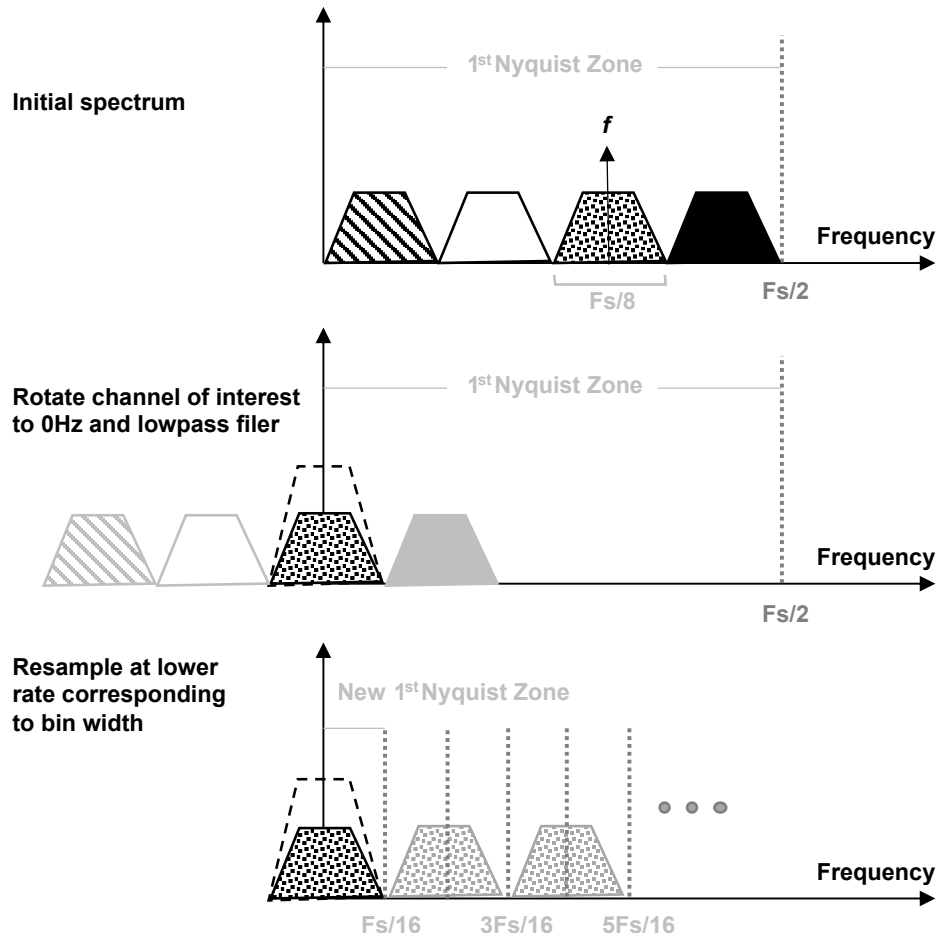


Figure 3.1: Basic channelization scheme. The sub-band of interest is migrated to 0 Hz. A Low-pass filter is applied to remove out-of-band signals. The sub-band is re-sampled at the new Nyquist rate. Images of the sub-band of interest appear in higher Nyquist zones but can be ignored.

Another attractive feature of PFBs is their inherently parallel structure which facilitates resource-efficient implementations on parallel computing architectures. Alternative channelizer algorithms include a tree-structure which uses a cascade of digital down converters and down samplers to repeatably halve the incoming bandwidth and a per-channel approach which uses a dedicated digital down converter, lowpass filter, and down sampler for each output channel. Reference [111] provides a comparison of these channelization schemes and finds the polyphase channelizer yields 4x reduced computational complexity as compared to the tree structure and 20x improvements over the per-channel approach. A comparison of FPGA resource utilization for the three channelizer structures is reproduced from [1] in Table 3.1.

The advantages of polyphase channelizers are well understood by the multirate digital signal processing community and there are many published examples of polyphase channelizers implemented on FPGAs [112, 113, 114, 115, 116]. Xilinx and MATLAB have both published detailed examples of resource efficient PFB implementations on FPGAs which include source code and may be used as a guide for groups to create their own [117, 118]. This work hopes to serve as a similar guide for how to create an efficient *oversampled* polyphase filter bank channelizer which is introduced in the next section. Specifically, this design provides the first example of an ultra-wideband, high-performance oversampled polyphase filter bank channelizer implemented on modern architecture using open-source code which readily allows customization and re-use by other groups.

3.1.2 Oversampled Polyphase Filter Bank

Oversampled polyphase filter banks (OPFB) are an extension of the PFB algorithm where data is recycled, effectively implementing overlapping channels while still achieving the required stopband rejection for each channel. This oversampled structure is required

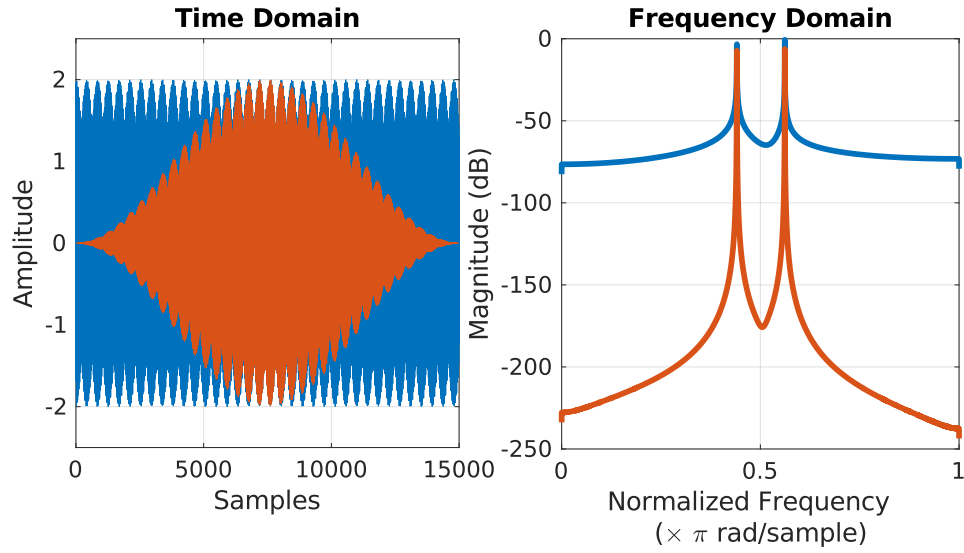


Figure 3.2: Left: Time-domain representation of a superposition of two frequencies shown with a sine envelope windowing function applied (red) and without (blue). Right: DFT output of the left time stream. The signal to noise is vastly improved in the windowed time stream (red).

Table 3.1: FPGA Channelizer Method Resource Comparison [1]

Channelization Method	Number of Channels	LUTs	Memory (bits)
Per-Channel	256	317,498	436,224
	512	650,114	876,544
	1024	1,336,754	1,761,280
Tree Structure	256	27,930	3,840
	512	32,270	6,529
	1024	36,610	10,625
Polyphase Filter Bank	256	8,070	4,608
	512	9,169	4,793
	1024	10,341	5,345

for applications demanding uniform response over wide bandwidths containing many separate channels such as a spectrometer, applications looking to utilize the full spectrum including 5G/6G wireless, and applications multiplexing large numbers of microwave resonators as is done in quantum computing and astronomy. A presentation of the data movement required to achieve oversampling is presented in the Appendix.

Despite strong theoretical foundations presented over a decade ago in [110, 119], there are no previously published examples of resource efficient OPFB FPGA implementations which provide enough information or source code to reproduce or reuse the design. Reference [120] describes a detailed OPFB FPGA implementation; however, the design uses more FPGA resources than the design presented in this paper to process half the bandwidth (2 GHz) into half the number of channels (2048) and does not run at processing clock speeds over 250 MHz. Reference [99] describes an efficient OPFB FPGA design but does not provide enough details to recreate the implementation and the authors are unable to share the source code (J. Tuthill, personal communication, January 2020). Electronics described in [121] make use of an oversampled polyphase filter bank but do not include any implementation details and the authors also declined to share the source code (J.M. D'Ewart, personal communication, August 2019). There are many more examples of OPFB FPGA designs referenced in publications with no implementation details [122, 123, 124]. As of this publication, OPFB is not supported by any of the Xilinx block set available in Vivado and MATLAB System Generator versions 2020.1 and prior.

The dearth of OPFB FPGA implementation examples can be partially explained by the previous requirement of custom HDL code resulting in designs too inflexible and application-specific to be relevant for other groups. This paper contributes the first example of an oversampled polyphase filter bank implemented without writing a single line of custom HDL. The design is inherently flexible as the custom digital signal processing blocks can be modified by making high-level changes in the C++ code. Un-

like HDL, the Vivado HLS compiler automatically generates the appropriate pipeline stages, greatly simplifying the task of porting this firmware between devices and speed grades. Several groups have compared Vivado HLS and custom HDL implementations in terms of FPGA resource utilization and maximum fabric speed with mixed results [125, 126, 127, 128, 129]. While there are no published FPGA resource utilization details for OPFB channelizers with comparable input bandwidths or number of channels, the authors will attempt to cross compare with smaller channelizers and show the resource utilization is better than or consistent with custom HDL approaches. It is also worth noting the digital design presented here has a larger input bandwidth (4 GHz) and greater FPGA fabric speed (512 MHz) than any previously published FPGA design incorporating Vivado HLS. Regardless of tool-flow, the authors know of only one other example of an oversampled polyphase channelizer that can process as much bandwidth into as many channels referenced in [121, 130]; direct resource comparisons are not possible as this full channelizer has yet to be demonstrated on an FPGA.

Overall, this paper demonstrates how to leverage modern FPGA programming tools to create a complex, high-performance digital design on modern architectures without custom HDL. We expect this work to be highly relevant for groups working in 5G/6G wireless, software-defined-radio, and scientific instrumentation who could benefit from using FPGAs to process ultrawide bandwidths and cannot expend the resources to create a custom HDL design. As the first group to make our oversampled polyphase filter bank source code freely available, we hope the results of our efforts will be adapted, reused, and built upon for future technological endeavors.

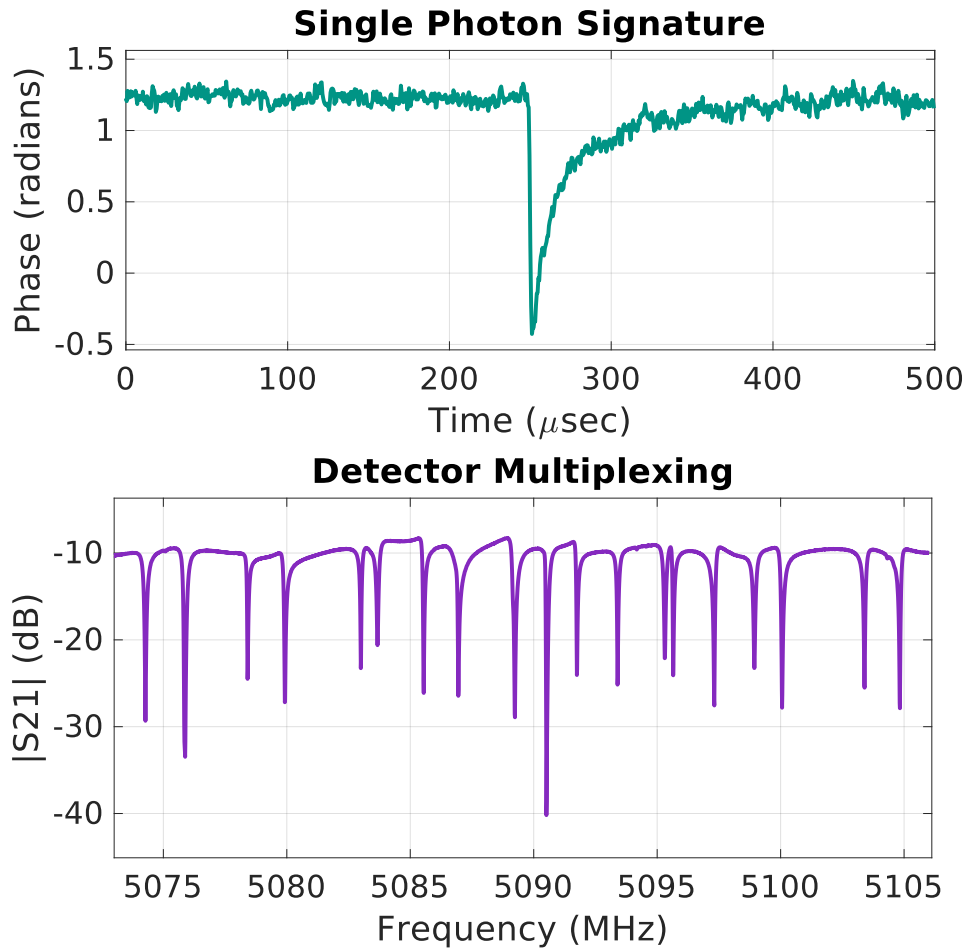


Figure 3.3: Top: Phase time stream of a single MKID resonance tone during a photon event. The pulse time gives the time of the event, the pulse depth gives the energy of the photon, and the resonator can be mapped to a location on-sky. Bottom: Transmission across a representative MKID feedline showing 19 frequency-division-multiplexed resonators. Each resonance dip corresponds to one MKID detector.




Figure 3.4: Sub-band spacing for adjacent sub-bands in the case of a critically sampled polyphase filter bank (top) and 2/1 oversampled polyphase filter bank (bottom). Only the oversampled PFB yields a uniform spectral response with the required 600 kHz of flat overlap.




Figure 3.5: Sub-band (blue) and band image (black, dotted) showing the alias band leaking into the 1 MHz sub-band edge (red, dotted) below 100 dB.

3.2 Channelizer Design

The oversampled polyphase filter bank was designed to serve as the first stage channelizer in a system which records the time, energy, and location of photons hitting a detector array. The array consists of ten microwave feedlines each servicing 2,000 superconducting microwave kinetic inductance detectors (MKIDs) and is optimized for fast time-domain optical astronomy [11, 131]. Each MKID pixel consists of a superconducting inductor and capacitor which together form a microwave resonator. When a photon strikes an MKID pixel, the resonance frequency momentarily changes and the photon signal is imprinted as a pulse in the phase time stream of the original resonant tone (see Fig. 3.3, top). The channelizer operates on a single 4 GHz microwave feedline which carries up to 2048 MKIDs, each with a unique resonance. A sample transmission of an MKID feedline showing nineteen resonators is reproduced in Fig. 3.3, bottom. Accurate characterization of the photon pulse requires each MKID resonant frequency to be migrated into the ± 1 MHz band while preserving ± 300 kHz of flat spectrum around each tone³. MKID fabrication procedures are unable to control the exact resonance location which requires the channelizer to be able to process a tone at any location in the band. The situation is analogous to a system with 2048 carrier waves which are randomly placed in a 4 GHz band and each modulated by a 300 kHz signal. Figure 3.4, top illustrates that a critically sampled polyphase filter bank channelizer is not suitable as any tone that appears in-between sub-bands will not have 300 kHz of flat spectrum on either side as required. We accomplished the task by instead using a non-maximally decimated polyphase filter bank with an oversampled rate of 2/1 and a 32,768-tap equiripple FIR prototype filter (see 3.4, bottom). This combination ensures that no matter where an MKID resonance is in the 4 GHz band, it is guaranteed to fall into a 2 MHz bin with 300

³Negative frequencies indicate the signals of interest are complex.

kHz of flat spectrum on either side. The resulting design provides up to 100 dB image rejection at the 1 MHz sub-band edges with floating-point computations (see Fig. 3.5).

3.3 Implementation

The polyphase channelizer was implemented on the Xilinx ZCU111 RFSoc evaluation board⁴ with four main signal processing blocks to achieve a high-throughput, multi-lane design similar to the technique used by the Australian Square Kilometer Array Pathfinder stage 1 polyphase channelizer (see Fig. 3.6) [99]. In this modular approach, the number of lanes L is set by the sampling frequency, f_s , the output channel spacing f_c , and the number of polyphase branches per lane, M_L , and is given by

$$L = \frac{f_s}{M_L f_c}. \quad (3.1)$$

For the OPFB presented here, $f_s = 4096$ MHz, $f_c = 1$ MHz, $M_L = 256$, and $L = 16$ with the total number of channels, $M = 4096$. The OPFB is designed to channelize complex samples formed from two ADCs; one sampling an in-phase channel, I , and the other sampling the quadrature channel, Q . The resulting Nyquist sampled bandwidth is 4 GHz. The design processes 16-bit numbers and truncates data as needed to maintain a 16-bit word length.

3.3.1 Digital Signal Processing Blocks

HLS Delay + Interleave

Two on-board ADCs each run at 4 GSPS and package their samples into eight parallel data lanes, resulting in eight sequential samples from each ADC simultaneously arriving in the programmable logic at a rate of 512 MHz [132]. These eight-sample-wide, 512 MHz

⁴P/N: xczu28dr-ffvg1517-2-e

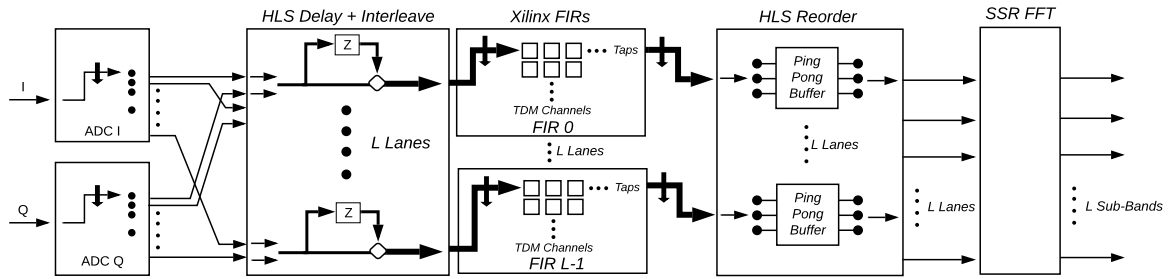


Figure 3.6: Overview of the multi-lane oversampled polyphase channelizer structure. Arrows indicate the flow of data between the four main signal processing blocks. All blocks operate on a 512 MHz clock domain.

streams obey the AXI4 standard adopted by Xilinx to manage high-speed data movement through FPGA devices [133]. Each delivery of eight samples will be referred to as a single AXI4-Stream beat. The *HLS Delay + Interleave* block first accepts two AXI4-Stream beats each from the in-phase and quadrature ADCs to create a sixteen-wide AXI4-Stream of complex samples updated at 256 MHz. The block has an initiation interval of 1 which is required to continually process the streams. To achieve an oversampling ratio of 2/1, every sample must be processed twice. This is accomplished by copying and delaying each stream by $M_L/2 = 128$ clocks, then interleaving the delayed stream with the non-delayed stream. The resultant output is sixteen data lanes containing two copies of each complex sample, each updated at 512 MHz (see Fig. 3.7).

Xilinx FIR

The Xilinx FIR core was chosen to execute the filter because it is widely available, highly configurable, and can operate at clock speeds over 500 MHz. One FIR core is instantiated per lane and configured to apply $M_L = 256$ coefficient sets to $2M_L = 512$ time-division multiplexed channels.

Following the standard polyphase decomposition, the prototype filter can be decomposed into M coefficient sets belonging to the M branches of the polyphase filter (see

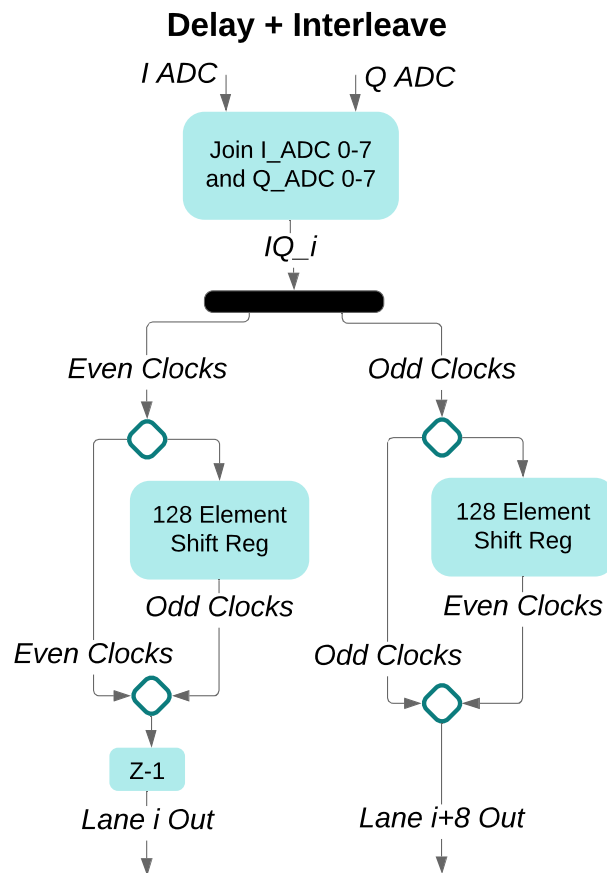


Figure 3.7: Algorithmic overview of the functions performed by the HLS delay + interleave block. The block creates complex samples by joining the I and Q streams from the two ADCs then delays and interleaves the resulting AXI4-Streams to create sixteen lanes.

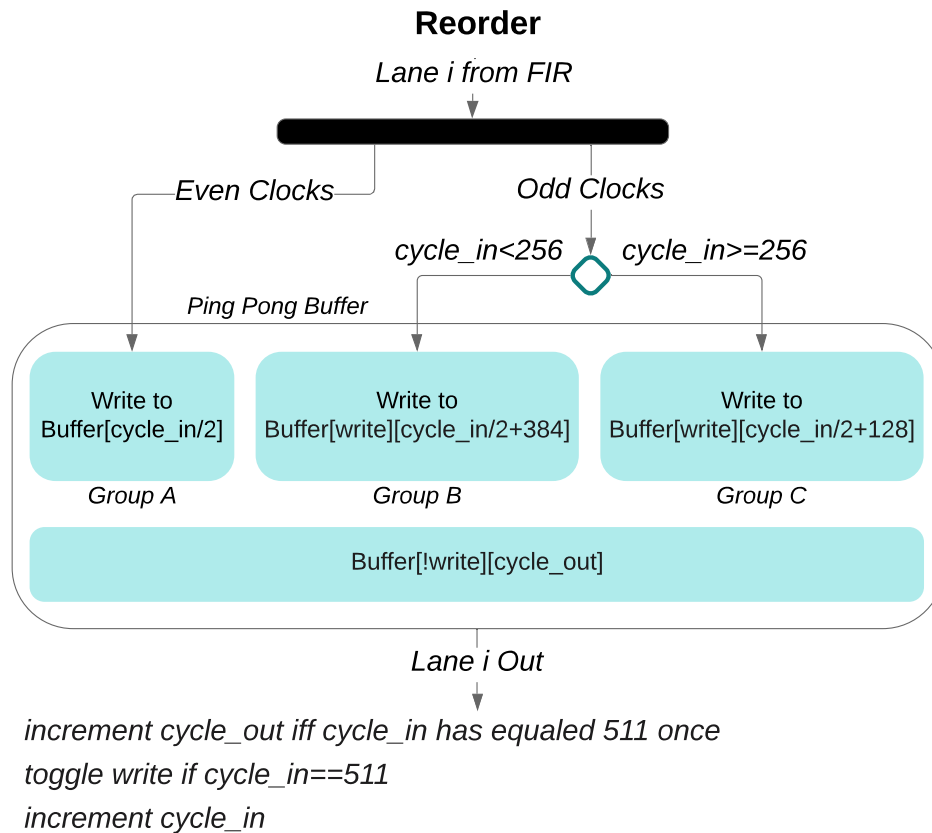


Figure 3.8: Algorithmic overview of the functions performed by the HLS reorder block. The incoming stream is broken into two groups each receiving every other sample. The even samples are denoted group A and the odd samples are further grouped into a the first 128 samples (group B) and the second 128 samples (group C). Samples are stored in an internal ping pong buffer and then sent to the vector FFT in order from A, then C, then B, thereby permuting the channels every other FFT frame to correct for the phase rotation induced by oversampling.




Figure 3.9: Illustration showing how the original polyphase decomposed coefficient sets are arranged in the multi-lane structure and the order they are used by the FIRs.

Appendix). We first consider coefficient set 0 as the set that is applied to samples $0, M, 2M, \dots$; coefficient set 1 is applied to samples $1, M+1, 2M+1, \dots$; coefficient set 2 is applied to samples $2, M+2, 2M+2, \dots$; and so on. Following the movement of the samples implemented by the *Delay + Interleave* block, lane 1 must apply sets $0, M, 2M, \dots$; lane 2 sets $1, M+1, 2M+1, \dots$; lane 3 sets $2, M+2, 2M+2, \dots$; and so on. Through careful consideration of the delayed, interleaved data sequence, one can derive the FIRs must apply their coefficient sets according to the channel configuration $1, 1, 2, 2, 3, 3, \dots$ (see Fig. 3.9). This has the effect of processing each data sample through two coefficient sets 180° out of phase with respect to the original polyphase decomposition which achieves $2/1$ oversampling.

HLS Reorder

It is well-known that non-maximally decimated PFB channelizers incur a phase rotation in each output sub-band (see Appendix)[110, 119, 99]. In the case of $2/1$ oversampling, this manifests as every other M -point FFT frame is cyclically permuted by $M/2$. This is corrected prior to the FFT by the HLS Reorder block which routes data into even (A group) and odd samples, with the odd samples grouped into a group of the first 128

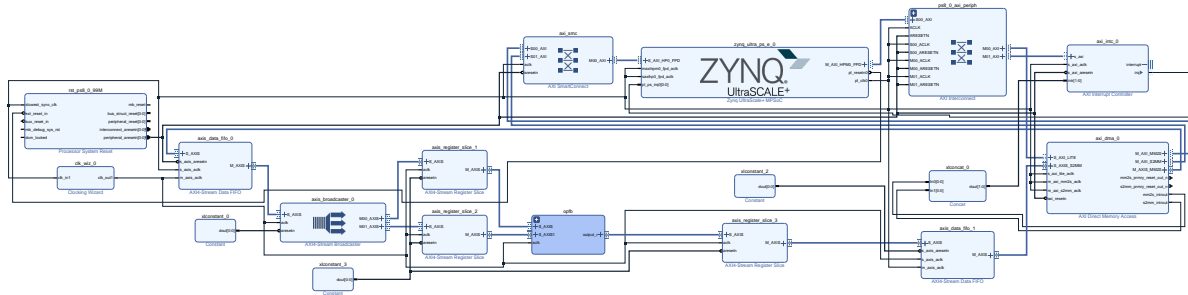


Figure 3.10: PYNQ Overlay as shown in Vivado HLx. The oversampled polyphase filter bank is contained within the OPFB hierarchy (dark blue). Data is delivered to and from the OPFB by Direct Memory Access (DMA) across two FIFOs. The Zynq processing system runs the Ubuntu-based Linux kernel and hosts the Jupyter Notebook server which creates the test data, triggers DMA transfers, and visualizes the channelized data. The clocking wizard outputs a 512 MHz clock used to drive the data FIFOs, AXI4-Stream Broadcaster, and OPFB. All other regions run at 100 MHz.

(B group) and the second 128 (C group). This block also has an initiation interval of 1. Samples are stored in an internal ping pong buffer before being sent to the vector FFT in order from A, then C, then B. This alternating pattern permutes the channels every other FFT frame, correcting the phase rotation induced by oversampling and presenting the data to the vector FFT in the expected, natural order (see Fig. 3.8).

Xilinx Vector FFT

The Xilinx Vector FFT was chosen to perform the discrete Fourier transform operation because it accepts parallel inputs and can operate at clock speeds over 500 MHz. The Vector FFT block was exported from the Xilinx Super Sample Rate (SSR) blockset available in System Generator versions 2019.1+ [134]. This block is configured such that $SSR = L$ and the transform length is M .

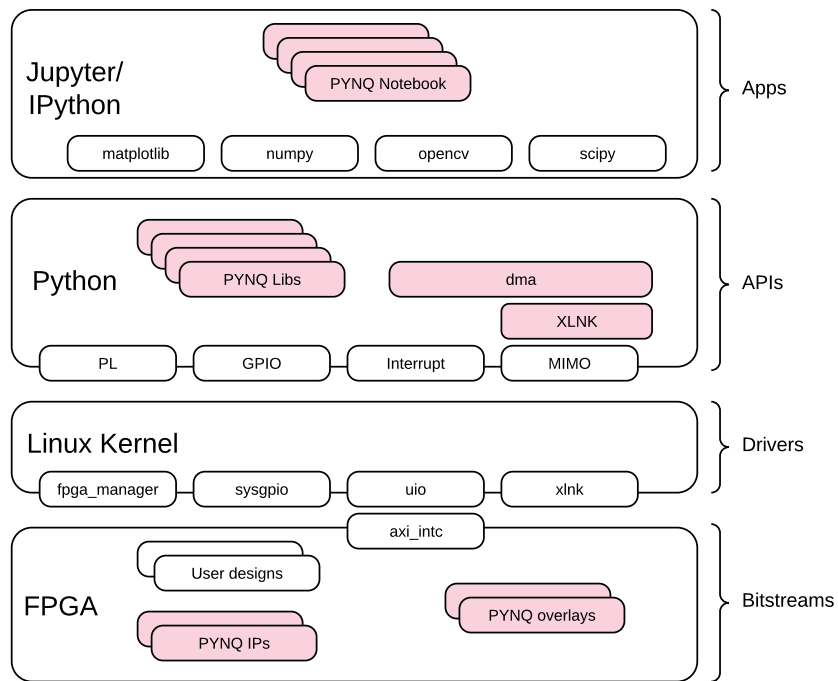


Figure 3.11: Overview of the Python Productivity for Zynq (PYNQ) system. PYNQ overlays are created using hardware design tools and are loaded onto the FPGA. The overlays are wrapped with the PYNQ Python API allowing software developers to easily program the FPGA, send data to the programmable logic, and plot intermediate results using common tools such as `matplotlib`.

3.4 Hardware Testing and Verification

The design was tested on the ZCU111 using the PYNQ system [103]. PYNQ is an open source Xilinx project which facilitates easy programming and interacting with Zynq FPGA devices through Python. PYNQ leverages the Zynq architecture which features hard-core ARM⁵ processing systems (PS) embedded in the FPGA programmable logic (PL). The framework consists of a PYNQ overlay which is a firmware design incorporating the blocks described in Section III built using Vivado and a Jupyter Notebook which is pure Python. The notebook server runs in an Ubuntu-based Linux kernel on the PS and can be used to program and configure the PL as well as move data through the PL and visualize the result. An overview of the PYNQ system modified from [103] is depicted in Fig. 3.11.

PYNQ was chosen to facilitate data movement, visualization, and testing on the FPGA because it provides a broad range of familiar python libraries including `matplotlib`, `numpy`, `scipy`, and more. Structuring the FPGA interaction in Python offers greater accessibility to a larger community of developers and promotes code adaptation and reuse. Data visualization could be accomplished by writing a bare-metal or Linux C/C++ application or moving data off the Zynq chip to a server where it can be processed; these options are not as straightforward, quick to implement, or accessible as the PYNQ framework. In this work we contribute another example of a FPGA design created and tested using PYNQ which joins a plethora of published examples [135, 136, 137]. A more in-depth discussion of the PYNQ framework is provided in [103, 135].

⁵ARM, 110 Fulbourn Rd., Cambridge, UK.

Table 3.2: OPFB FPGA Resource Utilization Comparison

Reference	Input BW ⁶ [MHz]	Filter Taps	N Output Channels	Oversample Ratio	Architecture
[139]	625	128	16	2/1	Virtex 6
[140]	2,000	40,960	4,096	2/1	Virtex 6
[141]	400	14,336	512	32/27	Kintex Ultrascale
[142]	1,000	567	32	2/1	Zynq Ultrascale+
This Work	4,000	32,768	4,096	2/1	Zynq Ultrascale+

⁶ Nyquist sampled input bandwidth.

Reference	LUT	BRAM ⁷	DSP	Max Speed
[139]	21,169	N/A	371	160 MHz
[140]	N/A	30	609	250 MHz
[141]	3,375	17	84	240 MHz
[142]	13,772	18	348	500 MHz
This Work	73,427	102	400	512 MHz

⁷ 36-Kbit BRAM.

3.4.1 Hardware Verification

The channelizer overlay was synthesized in Vivado and loaded onto the FPGA using the PYNQ `Overlay` class. The overlay block design as seen in Vivado is shown in Fig. 3.10. A 4 GHz Nyquist-sampled data stream from a 4 GHz I ADC and 4 GHz Q ADC was simulated using Python. Real data samples from each quadrature are cast to a 12-bit integer to simulate the ZCU111 ADC resolution and packed into 16-bit samples using the Python `fpbinary` library [138]. This C-backed library enables accurate characterization of quantization effects through the digital signal processing chain. The input stream is fed through the OPFB channelizer in packets using the Xilinx Direct Memory Access (DMA) tool and the PYNQ DMA driver. The resulting data can be manipulated and visualized in the notebook through standard Python tools.

To test the channelizer, two tones of equal amplitude were generated and the resulting sampled stream was fed through the OPFB using the test Jupyter Notebook. The tones were chosen such that one tone (300.5 MHz) would land exactly in-between two sub-bands and the other tone (302 MHz) would be centered in a sub-band. The resulting spectrum

of the three sub-bands containing the tones is plotted in Fig. 3.12. As expected, the tone in-between two bins appears in both overlapping bins with zero attenuation with respect to the band-centered tone and the tones appear at the correct frequencies with no observable harmonics in other bins. This confirms the OPFB is functioning as expected. Note if the channelizer was not overlapping, we would expect the tone located in-between sub-bands to be attenuated with respect to the band-centered tone (see Fig. 3.4). The implemented design provides -60 dB of spectral image rejection. Software floating-point models suggest the image rejection could be improved to -100 dB by increasing the computational precision at the cost of additional FPGA resources (see Fig. 3.5).

3.4.2 Resource Utilization

Resource utilization is reported for the dark-blue OPFB hierarchy pictured in Fig. 3.10 and does not include the Zynq processing system or data movement infrastructure surrounding the block. The design fits comfortably on the ZCU111 hardware target and uses 24% of the LUTs, 9% of the DSPs, 11% of the BRAM, and 35% of the LUTRAM. Device utilization footprint by block as pictured in Vivado Hlx is shown in Fig. 3.13. As expected, the compute intensive vector FFT and FIRs make up the bulk of the footprint. Resource utilization numbers are summarized in Table 3.2 along with results for several oversampled polyphase filter bank FPGA implementations published in the past five years. Direct comparison with the channelizer presented here is not possible as there are no previously published results with comparable parameters. Reference [142] achieves similar timing performance on the same architecture but is using comparable DSPs to process a quarter of the bandwidth into $1/128^{\text{th}}$ the number of channels with a filter that is sixty-times shorter. Without details on where the various resources are used or the ability to re-synthesize the design, we cannot attempt to extrapolate the resources to




Figure 3.12: The power spectral density of three OPFB sub-bands from the FPGA showing the frequency content of each sub-band. The 300.5 MHz tone is directly between two bins and the 302 MHz tone is band-centered. As expected, the tone in-between two bins appears in both bins and shows no attenuation with respect to the tone in the center of its own bin. This confirms the function of the oversampled polyphase filter bank.

match the parameters presented in this paper. Reference [140] has the same number of output channels and a larger filter but only processes half the bandwidth as the design presented here using 50% more DSPs and does not make timing past 250 MHz. Resource and timing differences across architectures cannot fully be attributed to differences in the FPGA device family because the maximum speed of the DSP units in the Ultrascale family is only 150 MHz higher than in the Virtex 6 family (750 MHz vs. 600 MHz) [143, 144]. These observations advocate the oversampled polyphase channelizer presented here is resource efficient as compared to previously published smaller channelizers, all of which were implemented using custom HDL approaches.




Figure 3.13: Device utilization footprint depicting the areas of the device used for specific functions. The full overlay pictured in Fig. 3.10 is shown with key DSP blocks internal to the OPFB hierarchy highlighted and labeled.

3.5 Conclusion

We have implemented a resource-efficient, high-performance oversampled polyphase channelizer on a Xilinx RFSoc capable of processing 4 GHz of instantaneous bandwidth. The modular, multi-lane design can be adapted to accommodate various channel bandwidths and is applicable to a wide range of fields including 5G/6G wireless, software-defined-radio, and superconducting microwave device arrays. Unlike most high-performance digital designs, this work was developed without custom HDL. The design leverages modern FPGA programming tools and mainstream software development languages and serves as an example for how a broader community of technologists might utilize FPGAs. Programming digital signal processing architectures with these tools allows digital designs to be modified and adapted with simple, mainstream software languages and promotes code reuse. The work presented here is available on GitHub⁶ under a GNU General Public License (GPLv3). We hope that our design will be recycled and expanded by other developers.

In future work, we hope to extend the design to accommodate arbitrary fractional decimation rates. Moving to other oversample ratios will require replacing the Xilinx FIRs with a custom polyphase filter designed using Vivado HLS. The current approach is resourced-constrained by the Xilinx FIR LUTRAM. Implementing the filter operation using Vivado HLS should also yield greater control over resource trade-offs, improving overall chip utilization. We also hope to build a mirror system which reverses the algorithm to produce a flat spectrum from the base-band channels; a technique otherwise known as polyphase synthesis.

$$H(Z) = h(0) + h(1)Z^{-1} + h(2)Z^{-2} + \dots + h(N-1)Z^{-(N-1)} \quad (3.2)$$

$$H(Z) = \sum_{r=0}^{M-1} Z^{-r} \sum_{n=0}^{(N/M)-1} h(r+nM)Z^{-Mn} = \sum_{r=0}^{M-1} Z^{-r} H_r(Z^M) \quad (3.3)$$

$$\begin{array}{rcccc}
H(Z) = & h(0) & + & h(M+0)Z^{-M} & + & h(2M+0)Z^{-(2M+0)} & + & \dots \\
& h(1)Z^{-1} & + & h(M+1)Z^{-(M+1)} & + & h(2M+1)Z^{-(2M+1)} & + & \dots \\
& h(2)Z^{-2} & + & h(M+2)Z^{-(M+2)} & + & h(2M+2)Z^{-(2M+2)} & + & \dots \\
& h(3)Z^{-3} & + & h(M+3)Z^{-(M+3)} & + & h(2M+3)Z^{-(2M+3)} & + & \dots \\
& \vdots & & \vdots & & \vdots & & \vdots \\
& h(M-1)Z^{-(M-1)} & + & h(1M-1)Z^{-(2M-1)} & + & h(3M-1)Z^{-(3M-1)} & + & \dots
\end{array} \quad (3.4)$$

$$H(Z) = H_0(Z^M) + Z^{-1}H_1(Z^M) + Z^{-2}H_2(Z^M) + \dots + Z^{-(M-1)}H_{(M-1)}(Z^M) \quad (3.5)$$

Appendix

Polyphase Decomposition

The derivation presented here is adapted from work presented in [110, 99]. To understand polyphase decomposition we begin with the more familiar mathematical description of a polynomial filter H as a sum of delayed polynomials in Z^M (3.2). Here N represents the length of the filter and must be an integer multiple of the FFT size M . This sum can be rearranged into a sum of M sums, each of which is a delayed polynomial in Z^M (3.3). The delayed polynomials are denoted by $H_r(Z^M)$. This formulation maps the one dimensional array of filter coefficients and index markers to a two dimensional array as shown in (3.4). This mapping is also used in the Cooley-Tukey fast Fourier transform (FFT) and allows the array to be loaded by columns but processed by rows. The first row in the matrix is a polynomial in Z^M and corresponds to H_0Z^M . Factoring Z^{-1} from the second row of the array also yields a polynomial in Z^M meaning the second row of the array corresponds to $Z^{-1}H_1(Z^M)$. In this way, the r^{th} row of the array can be

⁶https://github.com/MazinLab/RFSoc_OPFB

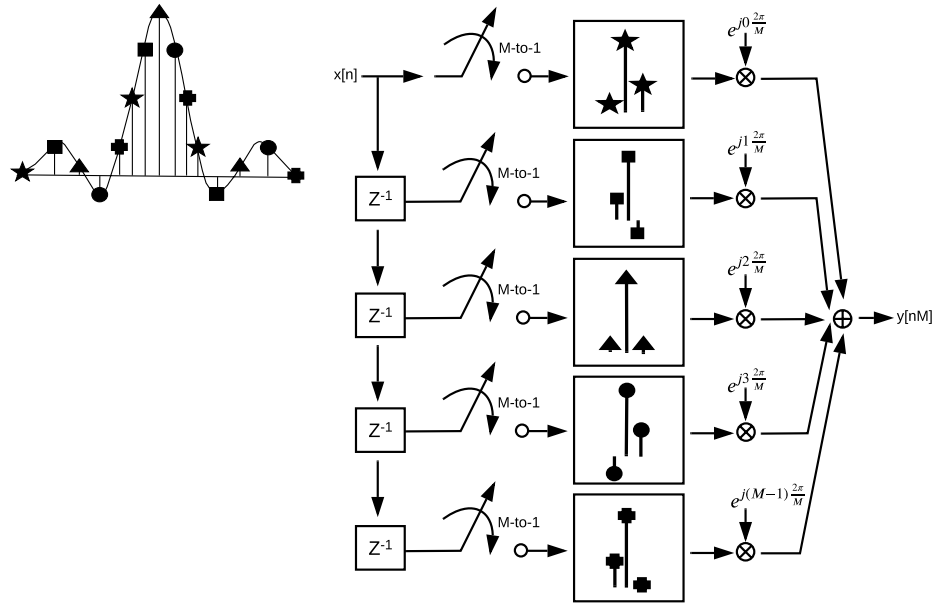


Figure 3.14: Block diagram illustrating the signal processing steps required to perform polyphase channelization into $M = 5$ sub-bands with a 15-tap prototype filter. Coefficients in the original filter are partitioned into the polyphase branches as indicated by the shape groupings. Data samples $x[n]$ are sequentially delivered to the five branches as indicated by the delays and M -to-1 down-samplers. Each branch applies the polyphase sub-filter by outputting the sum of the products of the incoming data samples and the three coefficients. Each incoming data sample will be multiplied by all three coefficients, moving forward one coefficient each time the branch receives a new sample. Complex multiplies align the phases such that the final sum results in destructive cancellation of the aliased terms introduced by the M -to-1 down sample operation. The resulting structure sequentially outputs M channels indicated by $y[nM]$.




Figure 3.15: Efficient FPGA implementation of a critically sampled polyphase channelizer. The input commutator delivers samples sequentially to the various branches, effectively implementing the delays shown on the input in Fig. 3.14. Each branch is processed by a filter structure which applies the polyphase components explicitly shown in Fig. 3.14. The phase coherent summation is performed by an M-point FFT.

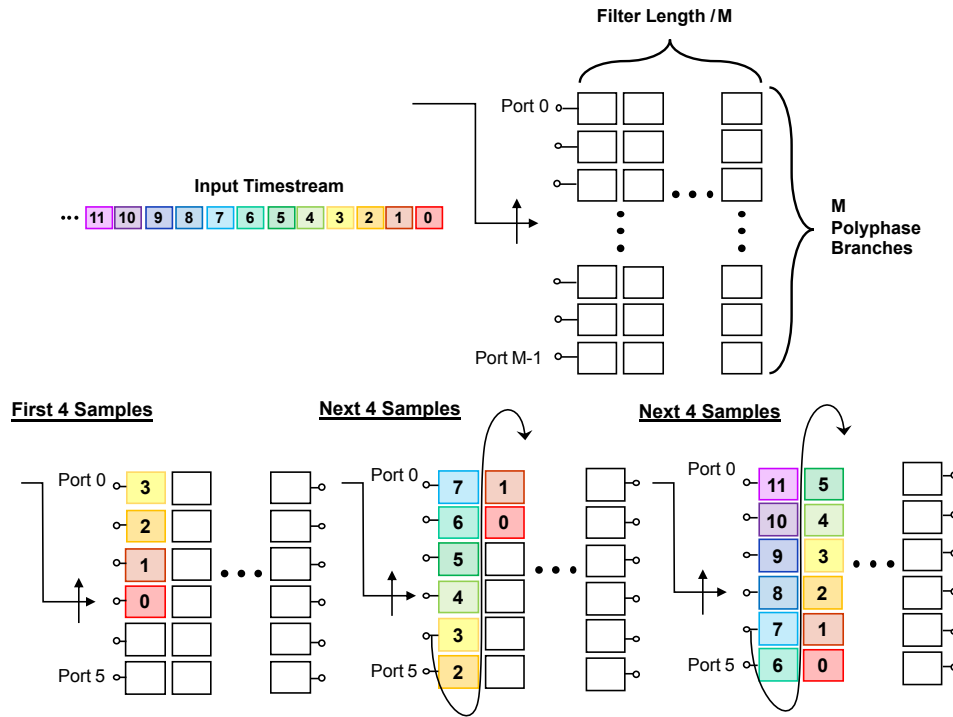


Figure 3.16: Overview of input data snake movement through the polyphase filter in the case $D = 4$ and $M = 6$. Four samples are loaded in per cycle. The previous data moves through in a snake-like fashion indicated by the arrows. Each cycle, every row in the polyphase filter outputs the sum of the products of the data samples in the boxes and the coefficients represented by the boxes to create six outputs for every four inputs.

described by $Z^{-r}H_r(Z^M)$. This notation can be used to recast (3.4) in a way that can be implemented as M parallel sub-filters or branches (3.5).

Equation (3.5) can be efficiently implemented in hardware through a structure sketched in Fig. 3.14, 3.15. A commutator is used to sequentially deliver data samples to the sub-filters and naturally implements the branch dependent delay corresponding to Z^{-r} in (3.3). The commutator implements an M -to-1 down-sample operation which breaks the Nyquist criterion and allows M -multiples of the output sample rate to alias down to 0 Hz, effectively translating each of the M bins to 0 Hz. As each successive sample comes in, the phase of the center frequency of the r th bin rotates by $2\pi r/M$. The aliased center frequencies incur an additional phase delay relative to one another which is the product of the center frequency and the branch-dependent path delay Z^{-r} . It is these branch dependent, differential phase effects which allow the bin centers to be separated and monitored after all aliasing to the same spectral zone.

When the branches are multiplied by phase rotators and summed as shown (3.6), the aliased terms which were not explicitly aligned destructively cancel,

$$y(nM) = \sum_{r=0}^{M-1} y_r(nM)e^{\pm j\frac{2\pi}{M}r}. \quad (3.6)$$

This process of aligning the bin centers and summing the result is an example of a phase coherent summation and is identical to applying an M -point FFT. Note if the direction of the commutator is reversed, the phase rotators incur a minus sign and the phase coherent summation is identical to applying an M -point IFFT. The final critically sampled polyphase filter bank structure with the input commutator, branch-dependent polyphase sub-filters, and M -point FFT is sketched in Fig. 3.15.

Oversampled Polyphase Filter Banks

The polyphase channelizer presented in the previous section delivers M input samples along the M branches and computes the output for each channel at a rate F_s/M . In this section, the design is modified to compute M output samples after only D ($D < M$) inputs have been loaded. This results in the FFT bins being over sampled by a factor of M/D , meaning they can be widened by a factor of M/D , resulting in overlapping without risk of aliasing. The trick is to recycle some of the previously loaded data much in the same way a moving average function computes an average of many data points after only a few new points have been loaded. The ratio of the output to input sample rate of the polyphase channelizer can be made to be any rational number by recirculating the data as next described.

In the non-overlapping structure, data is delivered sequentially to all M polyphase branches. Oversampling by D/M is accomplished by delivering only D samples to the first D ports. On the next data-load cycle, the samples snake through the full polyphase filter. This is equivalent to applying a linear shift to the 1-D filter expressed in (3.2) prior to the polyphase decomposition. The serpentine movement is explicitly shown in Fig. 3.16 for the case of $M/D = 6/4$.

Oversampling introduces a variable phase offset because the data windows length D are smaller than the FFT length M . This can be seen in Fig. 3.16 by observing how the same data sample cycles between a fixed number of rows which will result in the same data sample being fed into different FFT input ports. There are N different phase offsets where N is the numerator of the reduced fraction D/M . This is corrected by cyclically permuting the data before it enters the FFT by,

$$R(n) = (nD) \pmod{M}, \quad n = 0, 1, \dots, N - 1. \quad (3.7)$$

The cyclic permutation correcting the phase offset is explicitly shown for the case of $M/D = 6/4$ in Fig. 3.17. The permutation ensures the same data sample enters the same FFT port, correcting the phase rotation introduced by oversampling. A high-performance, resource-efficient FPGA implementation of an oversampled polyphase filter bank with $M/D = 4098/2048$ is the subject of this paper.

Gen3 OPFB Addendum

To make timing in the full Gen3 design, we reduced the OPFB filter size from eight taps per branch to four. This changes the actual channelizer response in the Gen3 system from what is shown in Fig. 3.4 to the response shown in Fig. 3.18.

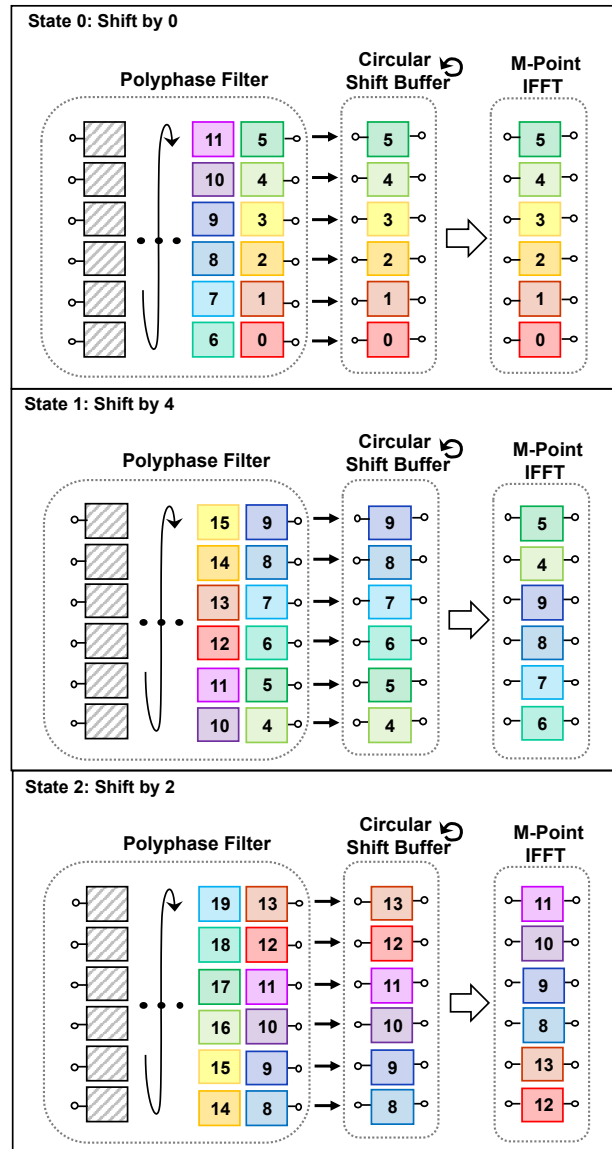


Figure 3.17: Overview of how to correct the phase shift induced by oversampling in the case $M = 6$ and $D = 4$ which yields $N = 3$ shift states. The cyclic shift up ensures the same data sample is fed into the same FFT port. The state increments forward every data-load cycle (four new samples enter the polyphase filter) and loops back to state 0 after state 2.




Figure 3.18: Gen3 4-tap OPFB channelizer response. Compare with Fig. 3.4, bottom.

Chapter 4

Cryogenic Cabling

Preface and Attributions

Previously, we motivated the need to multiplex many detectors per readout line with the cryogenic system heat budget. This heat load limits the number of wires between the room temperature electronics and the cryogenic device array. In this chapter, I will present work developing a new cryogenic cable that out-performed all commercially available options for MKID readout. The original version of the cable, presented in Sec. 4.1, was a collaboration with Benjamin A. Mazin, Alex B. Walter, Miguel Daal, J. I. Bailey, III., Clinton Bockstiegel, Nicholas Zobrist, Noah Swimmer, Sarah Steiger, and Neelay Fruitwala and was originally published in *IEEE Transactions on Applied Superconductivity*[145]. The second iteration was a collaboration with Benjamin A. Mazin, Alirio Boaventura, Kyle J. Thompson, and Miguel Daal and was also published in *IEEE Transactions on Applied Superconductivity*[146].

4.1 FLAX 1.0

4.1.1 Introduction

Superconducting devices are revolutionizing a wide range of research and technological fields including quantum computing [147, 148, 149, 150], nanowire single-photon detectors [151], X-ray microcalorimeters [152], submillimeter bolometers [153], and Microwave Kinetic Inductance Detectors (MKIDs) [154, 155, 156, 157]. These applications require increasingly large superconducting arrays, which present the common technical challenge of transporting microwave signals from the cold device stage to room temperature without losing or corrupting the signal or conducting excess heat to the cold stage. Low thermal conductivity is especially important for detector arrays in the field or in space using adiabatic demagnetization refrigerators (ADRs) which have less cooling power than dilution refrigerators but offer smaller form factors and simpler operation.

Commercially available superconducting coaxial cables are often used below 4 K; however, they are either semi-rigid and cumbersome to use in small cryogenic volumes, have large cross-sections yielding excessive heat loads, or both. Another option is flexible superconducting circuits fabricated using lithography techniques. These laminated cable technologies boast low thermal conductivity and high-density interconnects but lack the length, durability, and signal isolation needed for many applications [158, 159, 160, 161, 162].

An optimal solution should be made from superconducting material with a transition temperature well above 4 K to maximize transmission with an encompassing ground shield to minimize cross talk and pickup. It must have a small cross-section and be made from a low thermal conductivity material. Lastly, it should be flexible, durable, and ideally cheap and easy to manufacture. Such a structure is difficult to realize because few materials have the desired properties and often are difficult to work with and interface




Figure 4.1: Photographs showing a.) Close-up of cable end where NbTi center conductors connect to center trace of GCPW transition board via stainless steel capillary tubing. b.) Fully assembled cable end with protruding micro spot-welded, shared Nb47Ti ground shield. c.) Fully assembled cable spanning the 3.4 K stage and the 90 mK cold ADR stage with a thermal sink at 800 mK halfway down the length of the cable in the MKID Expoplanet Camera (MEC) experiment [161, 163].




Figure 4.2: Exploded view of cable-end assembly diagram with key dimensions shown in mm. Drawing is not to scale. From top/back to bottom/front: G3PO half-shell connectors are soldered to the transition board. Two ground tabs with via borders and an intervening signal trace create a 50Ω grounded coplanar waveguide. The FLAX cable center conductors are crimped into stainless steel capillary tubing and soldered to the center traces. The FLAX ground shield is spot welded to the ground tabs. The cable cross-section shows the PFA (blue) insulated NbTi (grey) wire set in semicylindrical crimps made in the shared Nb₄₇Ti foil ground shield. The two sides of the shield are mechanically and electrically bonded with micro spot welds less than $\lambda/16 \simeq 2$ mm (at 8 GHz) apart which run in-between the traces down the length of the cable.

with connectors.

In this paper we present a superconducting FLeXible coAXial ribbon cable (FLAX) which uniquely satisfies the aforementioned criteria. We developed this solution to carry broadband signals for 10 000+ pixel multiplexed Microwave Kinetic Inductance Detector (MKID) arrays for exoplanet detection operating at 90 mK [163, 164]. We expect this technology to be especially relevant for superconducting technologies requiring high detector isolation and low thermal load.

4.1.2 FLAX Design and Manufacture

The FLAX cables are fabricated using 0.076 mm [0.003"] O.D. NbTi center conductor insulated with 0.28 mm [0.011"] O.D. PFA wire obtained from Supercon¹. The shared outer coaxial conductor is formed with 0.025 mm [0.001"] Nb47Ti foil purchased and rolled by ATI² and HPM³. The wires are held in ten, 0.28 mm O.D. semicylindrical crimps made 3.56 mm apart in the foil to achieve a $\sim 50 \Omega$ characteristic impedance and 3.56 mm standard trace pitch density used by G3PO connectors available from Corning Gilbert⁴ (compatible with SMP-S) (see Fig. 4.7, 4.2). The two sides of the ground shield are mechanically and electrically bonded by micro spot welds which run the length of the cable between each trace. The welds are approximately every 2 mm which is less than $\lambda/16 = 2.3$ mm at 8 GHz (see Fig. 4.7, 4.2).

At the ends of the cable, the protruding center conductors are threaded into 1.6 mm O.D., 0.13 mm thick stainless steel capillary tubing. The tubing is crimped onto the center conductor before the assembly is soldered to the center traces of the transition board using a stainless steel soldering flux (see Fig. 4.7a, 4.2). The transition board is a 0.25 mm thick RT/Duroid6010LM PCB with 50 Ω grounded coplanar waveguide (GCPW) geometry for increased signal isolation. Between each trace, the Nb47Ti outer conductor foil is micro spot welded to the ground tabs of the transition board while surface mount coaxial G3PO push-on connectors are soldered to the other end of the GCPW (Fig. 4.7a). The cable end assembly is clamped in a 3×7 cm gold-plated copper box which provides strain relief and allows for easy push-on connection of all ten traces with G3PO blind-mate bullet connectors (Fig. 4.7b).

¹Supercon Inc., 830 Boston Turnpike, Shrewsbury, MA.

²ATI Specialty Alloys & Components, 1600 Old Salem Rd., Albany, OR.

³Hamilton Precision Metals, 1780 Rohrerstown Rd., Lancaster, PA.

⁴Corning Optical Communications, 4200 Corning Place, Charlotte, NC.




Figure 4.3: Schematic diagram depicting FLAX attachment to the G3PO push-on connectors via a capillary tube soldered to a coplanar waveguide transition board and the device under test (DUT) circuit at 4 K.

4.1.3 Performance Characterization

Transmission loss (S_{21}), cross talk (S_{41}), and time domain reflectometry measurements were performed in a dilution refrigerator under vacuum at 4 K with a Keysight N9917A network analyzer. The device under test circuit consisted of the assembled FLAX cable with a 3 dB cryo-attenuator obtained from XMA⁵ and a 25 cm nonmagnetic SMA-to-G3PO adapter coaxial cable obtained from Koaxis⁶ on either end (see Fig. 4.11). A Crystek⁷ braided, semi-rigid coax through line was used as a calibration reference. Repeated handling through the testing process revealed the cables have a minimum inside bend radius close to 2 mm and are robust to cryogenic cycling.

⁵XMA Corporation-Omni Spectra, 7 Perimeter Road, Manchester, NH.
P/N: 2082-6040-03-CRYO

⁶Koaxis RF Cable Assemblies, 2081 Lucon Road, Schwenksville, PA.
P/N: AO10-CC047C-YO18

⁷Obtained through Digikey. P/N: CCSMA18-MM-141-12




Figure 4.4: Top: S_{21} (transmission) measurement of sample FLAX traces from various cables at 4 K. Bottom: S_{41} (nearest neighbor forward cross talk) measurement of sample FLAX traces from the same cable at 4 K. The average cross talk level is given by the dashed red line.

Transmission

Ripples in the FLAX transmission suggest standing wave modes are present on the traces which is indicative of an impedance mismatch between the FLAX cable and the $50\ \Omega$ circuit (see Fig. 4.12). The transmission ripples are not uniformly harmonic which suggests the impedance is changing with length along each trace. This could be explained by flaws in micro spot welding placements along the cable which determine the distance between the inner and outer coaxial conductors and therefore the characteristic impedance. The characteristic impedance of the traces were probed using a time domain reflectometry measurement adjusted for loss (see [165] for details on loss correction) which confirmed the impedance varies from $55\text{--}65\pm 3\ \Omega$ along the traces (see Fig. 4.13). This mismatch at various points in the cable launches reflected waves which contribute to the

observed ripple.

We hypothesize an additional factor contributing to the impedance mismatch originates in the intermediate regions of the cable ends where the center conductor exits the foil sheath and transitions onto the GCPW transition board (see Fig. 4.7, a.). After exiting the ground shield, the exposed wire can act as an inductor. Previous work done by our group shows inductance on the input and output of a transmission line causes ripples which increase in magnitude at higher frequencies [161]. This is because the impedance of a perfect inductor grows linearly with frequency, i.e., $Z_L = j\omega L$. With each successive cable iteration, manufacturing techniques improved, the length of exposed wire was shortened, and the frequency-dependent ripple amplitude diminished. The use of a capillary tube to pin the hair-like center conductor close to the transition board dramatically reduced the cable end inductance.

Using the peak of the ripple, we report the loss of the 30 cm cable at 8 GHz to be roughly 1 dB which is slightly higher than the 0.5 dB/m loss reported by commercially available superconducting coaxial cables [166, 167]. This difference cannot be explained by a difference in cable materials or geometry [168]. Likely, the source of our additional loss is the impedance mismatch caused by manufacturing imperfections which produce reflections in the cable and off the ends as described above.

Table 4.1: Summary of thermal, mechanical, and microwave properties of superconducting coaxial ribbon cable, laminated microstrip cable, and best commercially available superconducting coaxial cables

Cable	Thermal Load ⁸ per trace [nW] 100mK to		Mechanical All Dimensions [mm]				Microwave Values at 8GHz Attenuation [dB]
	1 K	4 K	Trace Pitch	OD (\varnothing)	Min. Inside Bend Radius	Cross Talk [dB]	
FLAX	16	800	3.556	0.376	2	-60	1
CryoCoax	26	1400	>13	0.900	3.2	N/A	< 0.5
KEYCOM	34	1800	>13	0.860	8	N/A	< 0.5
Nikaflex	16	460	3.556	0.198 ⁹	6.4	-25	1

⁸ Computed using dimensions available from [166, 167, 161] and assuming a cable length of 30 cm.

⁹ For the microstrip geometry this is the total cable thickness.




Figure 4.5: A typical time domain reflectometry (TDR) measurement of the cryogenic signal path showing the characteristic impedance at lengths along the signal path. Commercially available 50 Ω standard coaxial cables border the FLAX cable highlighted by the double arrow. Note the TDR measurement is accurate to $\pm 3 \Omega$.

Cross Talk

We found the average nearest-neighbor forward cross talk to be -60 dB (see Fig. 4.12). This is roughly 30 dB lower than what we previously achieved using flexible laminated NbTi-on-Kapton microstrip cables [161]. Since the cable's installation in the MKID Exoplanet Camera (MEC) at Subaru Observatory, this enhanced isolation has increased our pixel yield $\sim 20\%$ [163]. We suspect this large improvement is because the exposed microstrip geometry allows trace-to-trace coupling whereas the coaxial nature of the FLAX shields the center conductors thereby preventing signal corruption. In early iterations of the cable, we found infrequent or failed micro spot welds in the ground shield lead to much higher levels of cross talk. This leads us to conclude incorporating micro spot welds less than $\lambda/16$ apart between the traces reduces electromagnetic coupling.

Thermal Conductivity

Following previous convention, a cable thermal conductivity, $G(T)$, was computed by summing literature values of constituent materials weighted by their cross-sections (see Fig. 4.14) [169, 161]. We compare our superconducting coaxial ribbon cable to two commercially available superconducting coaxial cables as well as our lab's previously developed laminated NbTi-on-Kapton microstrip cables [161]. We estimate the thermal conductivity of the PFA dielectric present in the flexible coaxial ribbon cables using PTFE; the same dielectric used in the two commercial solutions [168]. The smallest commercially available superconducting coaxial cables from KEYCOM¹⁰ and CryoCoax¹¹ were chosen for comparison. The electrical and thermal properties of the cables are summarized in table 4.2.

¹⁰KEYCOM Corp. 3-40-2 Minamiotsuka, Toshima-ku Tokyo.

P/N: NbTiNbTi034

¹¹CryoCoax - Intelliconnect, 448 Old Lantana Road, Crossville, TN.

P/N: 5139-P1NN-611-100P

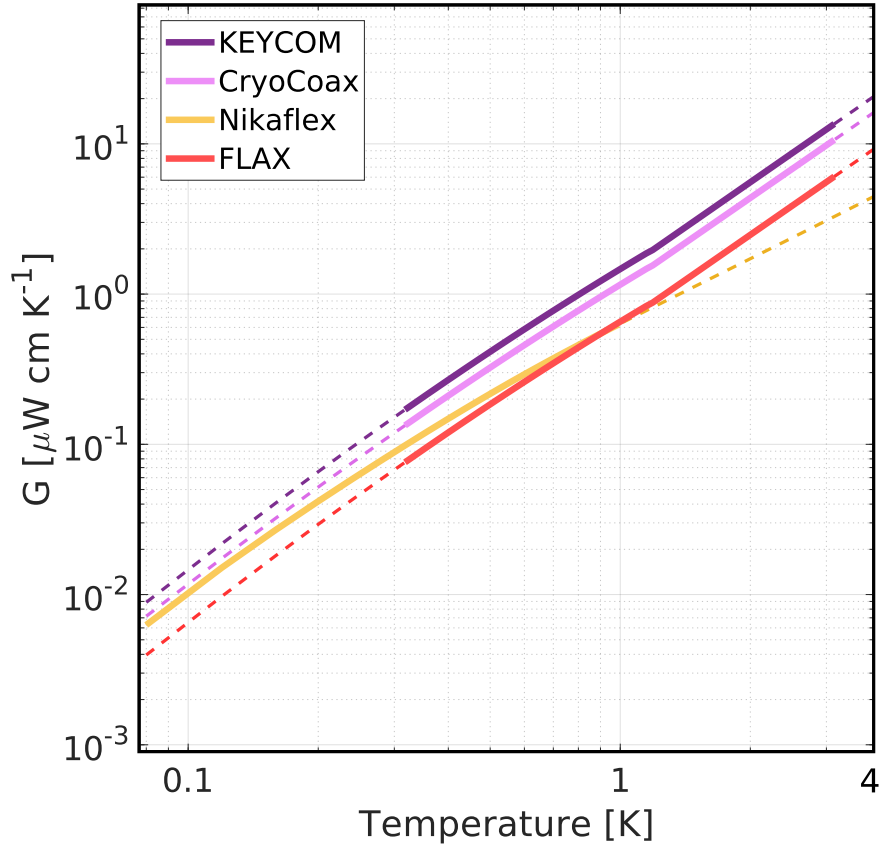


Figure 4.6: We computed a cable thermal conductivity $G(T)$ in units of μWcmK^{-1} by summing the thermal conductivity of each constituent material weighted by the cross-section [169, 161]. The cable previously developed by our lab (Nikaflex, gold) is compared with the subject of this paper (FLAX, salmon), and two commercial options by KEYCOM (P/N: NbTiNbTi034, burgundy) and Cryocoax (P/N 5139-P1NN-611-100P, pink). Solid lines are computed using literature values for Nb47Ti [170, 171], PTFE [169], Nikaflex (Kapton polyimide film) [161, 172], and Pyralux [170]. PTFE values were used to estimate the PFA dielectric in the FLAX cable [168]. Dashed lines indicate extrapolation.

The heat load from one temperature stage to another can be computed by integrating values in Fig. 4.14 from T_1 to T_2 ($T_1 < T_2$) and dividing by the cable length. The ten-trace FLAX cables are currently installed in the MEC experiment where they span 33 cm from the 3.4 K stage to the 90 mK cold ADR stage with a thermal sink at 800 mK about halfway down the length of the cable [164]. We estimate they generate a thermal load of ~ 200 nW on the 90 mK cold ADR stage. This is about equivalent to the thermal load created by the Nikaflex cables and approximately half the computed heat load of either commercial option.

4.1.4 Conclusion

We have manufactured a superconducting flexible coaxial cable capable of delivering microwave signals between temperature stages with minimal loss, cross talk, and heat conduction. Strong signal isolation is especially important for our application of moving 4-8GHz servicing 10 000+ multiplexed sensors across temperature stages. The FLAX cable represents a 30 dB improvement in cross talk as compared to our group's previously developed NbTi-on-Kapton microstrip cables. This enhanced isolation facilitated a $\sim 20\%$ increase in MKID pixel yield in the MEC experiment [163]. We expect these results will be especially useful for high-density microwave superconducting detector arrays requiring strong signal isolation.

The cable technology presented in this paper also has very low thermal conductivity. For a given thermal budget, the FLAX cables allow for twice as many detectors as the leading commercial option. The reduced heat load combined with the push-on, small form factor connectors and reduced trace pitch allow for increased detector density in a cryogenic system.

We found an attenuation of 1 dB at 8 GHz with ~ 3 dB ripples which is at worst 2x

more loss than commercial options. This magnitude of ripples and loss do not impact our array on the input side as we can drive microwave resonators (MKIDs) located at transmission dips with higher power than their frequency neighbors. However, these features degrade the overall signal to noise ratio on the output. Ripples and loss may become prohibitive for systems operating at frequencies over 8 GHz or systems constrained by amplifier dynamic range. Insertion loss and ripples can be reduced by improving manufacturing precision in the forming of the NbTi foil crimps and location of micro spot welds. Alternative methods to join the push on connectors and traces, e.g., brush plating the NbTi center conductor with an easily solderable material such as nickel may also improve the impedance match.

Lastly, we note these cables are relatively easy to fabricate. Many components, most notably the fine, NbTi center conductor wire, are commercially available. All cable iterations were manufactured in-house at the University of California, Santa Barbara. Ten trace FLAX can be assembled in two days. Overall, we find this cable technology to be superior to commercial options for our applications building high-density superconducting detector arrays.

4.2 FLAX 2.0

After the initial cable iteration, we partnered with Maybell Quantum Industries to create a commercial solution through the DARPA SBIR program [173]. With funding and support from Maybell, we improved on the previous design and published the second iteration, reproduced below [146]. At the time of writing this thesis, Maybell is continuing research and development on the wires and is planning to offer them as an option in their dilution refrigerator systems [174].

4.2.1 Introduction

Superconducting devices are driving technology advancements in quantum computing [2, 175], single-photon-counting sensors [41, 67], low-noise amplifiers [176], and more. These technologies all require specialized cryogenic wiring capable of transporting electrical signals across temperature stages with minimal loss, crosstalk, and thermal conductivity. The development of high-density cryogenic wiring has proved to be a bottle-neck in scaling up superconducting systems that need potentially thousands of cryogenic I/O channels to meaningfully operate [177].

Progress has been slowed in part by stringent performance requirements that limit wiring solutions to materials that are difficult to work with. Niobium-titanium (NbTi) is one such material that has very low thermal conductivity and is superconducting under 11 K, providing excellent cryogenic transmission and isolation with a fraction of the heat load afforded by other metals [178]. This unique combination of electrical and thermal properties has led NbTi to become the conductor of choice for cryogenic wiring despite the fact it cannot be soldered and is otherwise difficult to work with and connectorize [179]. Similarly, Teflon-type materials make ideal dielectric mediums because they are flexible with low thermal conductivity and dielectric loss. Unfortunately, these materials come with a host of manufacturing challenges including extremely high melting temperatures, non-trivial phase changes, potential to produce highly-toxic byproducts, and inability to adhere well to other materials [180, 181].

Despite the fabrication challenges, vendors have managed to fabricate individual cryogenic coax lines by extruding NbTi tubes threaded with Teflon-coated NbTi wire [182, 183]. These cables boast excellent transmission and isolation but are semi-rigid and cumbersome to use in small cryogenic volumes. Furthermore, the extrusion process constrains their size, resulting in relatively large cross sections and consequently excessive

heat loads. More recently, processes have been developed that involve sputtering NbTi onto a polyimide material such as Kapton to form stripline traces [184, 185, 186]. NbTi-on-polyimide cables are more flexible and have lower heat load than the aforementioned semi-rigid cables but have higher dielectric loss that scales with frequency. This loss can be prohibitive to applications carrying RF or Microwave signals from cryogenic devices before amplification where the signals are small and vulnerable to noise. Neither of these leading commercial options met the transmission, crosstalk, and heat load requirements for our application delivering 4-8 GHz signals to an array of 20,000+ microwave kinetic inductance detectors (MKIDs) at 10 mK [187, 41].

Unable to buy suitable cryogenic wiring, we previously developed our own in-house cable manufacturing process at the University of California, Santa Barbara that involves spot welding NbTi foil around PFA-coated NbTi wires to form a flexible coaxial ribbon cable (FLAX) [145]. This technology aims to replicate the near-perfect isolation and transmission offered by semi-rigid, superconducting coax but in a smaller, more flexible format that reduces heat load and allows for increased I/O density. In [145], we demonstrated competitive heat load and isolation but the transmission loss could prove to be a dominant noise source for highly-sensitive experiments like those requiring quantum-noise-limited amplifiers. Another drawback was the highly variable, manual, and labor intensive manufacturing process [145].

In this paper, we present improvements to the cable design and show 3x improved cable transmission without negatively impacting the isolation or heat load. We share initial progress by Maybell Quantum Industries on a roll-to-roll manufacturing system to streamline and standardize cable fabrication in preparation for commercialization. Lastly, we compare the latest performance data with the smallest commercially available semi-rigid superconducting coax from CryoCoax¹² [182] and the NbTi-on-polyimide

¹²CryoCoax - Intelliconnect, 448 Old Lantana Road, Crossville, TN, USA. P/N: BCB012




Figure 4.7: Eight-trace Maybell cable prototype manufactured using automated roll-to-roll laser welding system. Two sheets of 1 mil thick NbTi foil are welded around PFA-coated NbTi wires with a continuous laser weld. The cable appears uniform, flexible, and durable.

Cri/oFlex[®]3 from Delft Circuits¹³ [186].

4.2.2 Manufacturing Improvements

The previous iteration of the FLAX cable was manufactured using individual, hand-placed spot welds to act as vias between the cable traces [145]. While effective, this process was slow and error-prone because any welds placed too close to the center wire result in an electrical short, rendering the cable useless. Maybell has developed a laser welding cable manufacturing system and is developing a roll-to-roll apparatus that will greatly simplify and speed up cable manufacturing as well as increase overall uniformity. A prototype sample produced using a linearly actuated mechanical stage is shown in Fig. 4.7. A close up of the continuous laser welds surrounding two wires is shown in Fig. 4.8. We expect cables manufactured with this system to show enhanced durability and better impedance control due to the precise, continuous laser welds. We also note this manufacturing system allows cables to be made with arbitrary length, giving more flexibility to cryogenic routing and design.

¹³Delft Circuits, 2627 AN, Delft, Netherlands. P/N: Cri/oFlex[®]3

4.2.3 Design Improvements

Two main design improvements were made to correct for impedance mismatches in the cable and improve the microwave transmission. First, the characteristic impedance of the cable was adjusted closer to 50Ω . Second, the connectorization was modified to reduce reactive impedance reflections off the connector ends.

Characteristic Impedance

The characteristic impedance of the cable was originally designed in [145] assuming a perfect circular cross section for the inner and outer conductor geometry. In reality, the manufacturing process produces an outer conductor profile more parabolic than semicircular (see Fig. 4.8). This results in the cable cross section appearing more elliptical (see Fig. 4.9) which causes the outer conductor to be effectively further from the inner-conductor, elevating the cable characteristic impedance.

We modeled the actual coaxial shape as a spline in Ansys HFSS and determined increasing the center conductor diameter from 0.003 inches to 0.005 inches (without changing the dielectric thickness) would effectively bring the inner and outer conductors closer together and bring the impedance back down to the target 50Ω .

End Inductance

When traveling out of the cable and into the connector, the electric field transitions from a coaxial mode in the cable to a coplanar waveguide mode in the transition board, and then back to a coaxial mode in the connector. The center conductors, which behave inductively, are relatively continuous in the transition; however, the surrounding ground planes, which provide a balancing capacitance, undergo a large geometry shift from 3D to 2D and back. The authors theorize that the disruption to the 3D ground plane

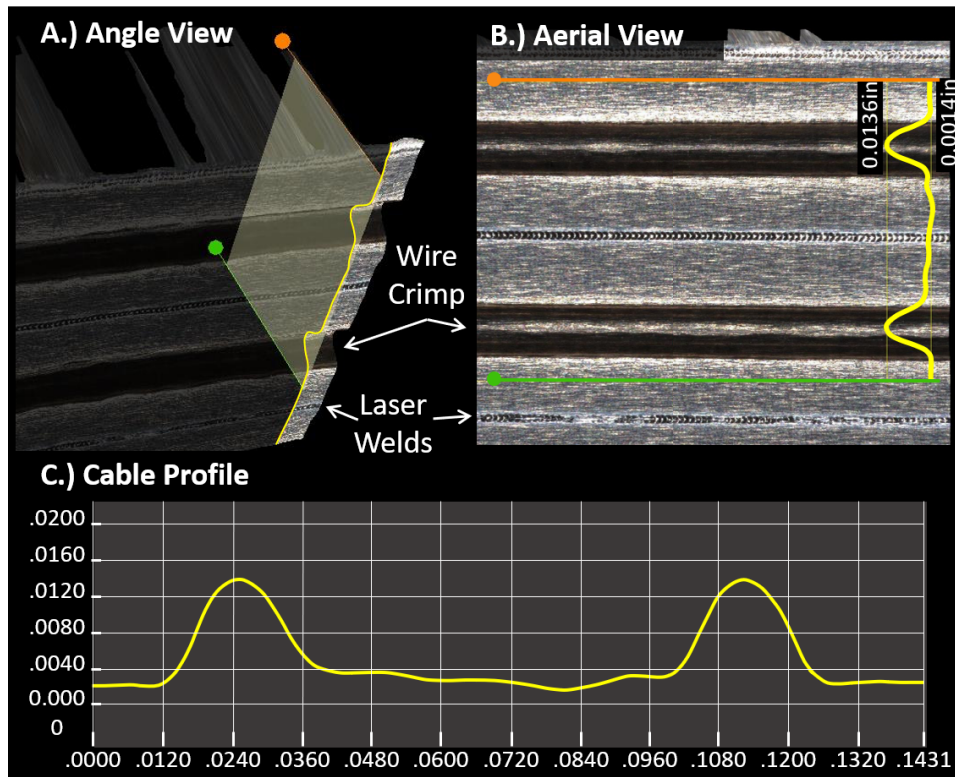


Figure 4.8: Manufacturing data obtained from JT Automation during a prototype cable manufacturing run. A.) Angle view showing a single sheet of NbTi foil with two crimps made to house wires with continuous laser welds in-between. B.) Aerial view of the NbTi foil showing the same welds and crimps pictured to the left. C.) Cable profile measured in inches. In all views the cable profile is highlighted in yellow. The profile shows good uniformity and repeatability across the length of the cable. The laser welds appear high-quality.




Figure 4.9: Exploded view of cable-end assembly with key dimensions shown in inches. Drawing is not to scale. From top/back to bottom/front: G3PO half-shell connectors are soldered to the transition board. The transition board features $50\ \Omega$ grounded coplanar waveguide structures with the upper ground planes connected in front to create capacitance between the upper ground planes and signal traces. The FLAX cable center conductors are crimped into stainless steel capillary tubes and soldered to the center traces. The FLAX ground shield is spot welded to the upper ground planes. The cable cross section shows the PFA (blue) insulated NbTi (grey) wire set in parabolic crimps made in the shared NbTi foil ground shield. The two sides of the shield are mechanically and electrically bonded with micro spot welds less than $\lambda/16 \simeq 2\ \text{mm}$ (at 8 GHz) apart that run in-between the traces down the length of the cable. Figure is modified from [145].

reduces capacitance in the transition and causes the connector ends to behave inductively. This can be seen in the transmission spectrum from [145], which shows ripples in the transmission consistent with the fundamental mode of the cable growing with increasing frequency. This is consistent with reflections off inductive features in the cable ends because the impedance of a perfect inductor grows linearly with frequency, i.e., $Z_L = j\omega L$.

The authors used Ansys HFSS to simulate the connector transition and determined the inductance in the cable ends could be mitigated by including a capacitive feature in the transition board. This was accomplished by modifying the upper coplanar waveguide ground planes on the transition board, bringing them together in front of the signal trace with a 0.006 inch gap (see Fig. 4.9 and Fig. 4.10).

4.2.4 Performance Characterization

Table 4.2: Summary of thermal, mechanical, and microwave properties of superconducting coaxial ribbon cable, commercial semi-rigid superconducting coax, and commercial superconducting polyimide flex cables

Cable	Thermal Load ¹⁴ per trace [nW] 100mK to		Mechanical All Dimensions [inches]				Microwave ¹⁵ 8 GHz Attn. [dB]
	1 K	4 K	Connector Pitch	OD (\varnothing)	Min. Inside Bend Radius	S41 [dB]	
FLAX [This Work]	5	245	0.140	0.02	0.08	-50	1.5
FLAX [145]	5	238	0.140	0.02	0.08	-50	7
Semi-Rigid ¹⁶	7	370	> 0.5	0.900	0.13	< -60	0.5
Polyimide-Flex ¹⁷	TBD	TBD	0.3	0.01 ¹⁸	0.04	-60	7

¹⁴ Computed using a cable length of 1 foot.

¹⁵ Values reported for \sim 1-ft long cables.

¹⁶ CryoCoax 0.034 inch NbTi/NbTi, P/N: BCB012 [182].

¹⁷ Delft Circuits Cri/oFlex[®]3 with NbTi conductor [186].

¹⁸ For the microstrip geometry this is the total cable thickness.

Design modifications were tested using a 5-trace, 1-ft long cable manufactured in-house at the University of California, Santa Barbara using the same process described in [145]. We compare our results with the previous iteration of the cable as well as the smallest commercially available superconducting coax from CryoCoax (P/N: BCB102)




Figure 4.10: A.) Box end housing and grounded coplanar waveguide transition board design used in [145] that yielded excess end inductance. B.) Box end and modified grounded coplanar waveguide design with capacitive coupling used in this work. C.) Connectorization scheme showing how cable is connectorized in this work. The cable ground is spot welded to the transition board ground. The cable center wire is exposed by cutting the foil ground and then crimping a stainless steel tube onto the wire. The tube is soldered to the transition board. The transition board and G3PO connectors are soldered into the box housing.

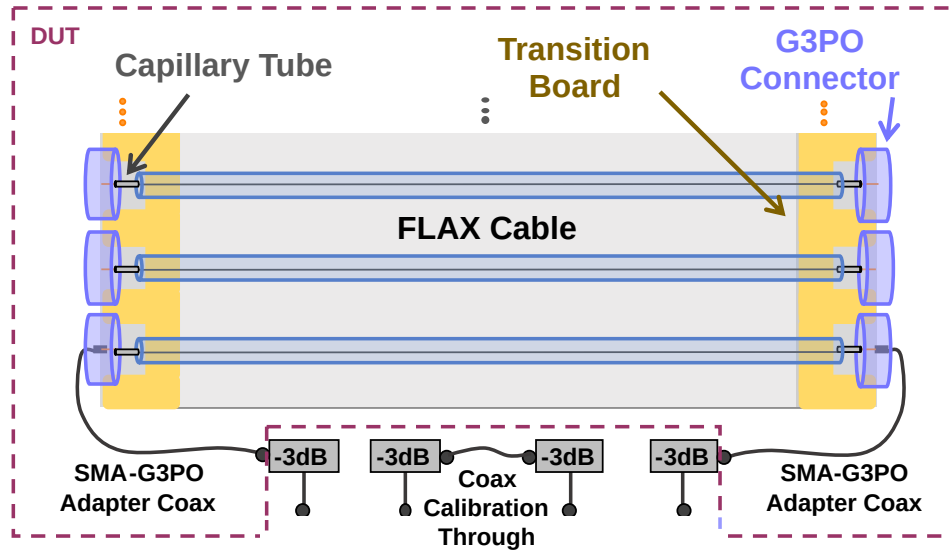


Figure 4.11: Schematic diagram depicting the device under test circuit at 4 K. Figure is modified from [145]

[182] and the only commercially available NbTi-on-polyimide flex cable from Delft Circuits (P/N: Cri/oFlex[®]3) [186]. Key thermal, mechanical, and electrical properties are summarized in Table 4.2.

Transmission loss (S_{21}), crosstalk (S_{41}), and time domain reflectometry were measured at 4 K with a network analyzer. The device under test circuit consisted of the assembled FLAX cable with 3 dB cryo-attenuators obtained from XMA¹⁹ and 25 cm nonmagnetic SMA-to-G3PO adapter coaxial cables obtained from Koaxis²⁰ on either end (see Fig. 4.11). A non-superconducting commercial coax through line was used as a calibration reference for fridge wiring as in [145].




Figure 4.12: Top: S_{21} (transmission) measurement of sample traces from [145] (dashed) and this work (solid) at 4 K. The transmission is improved 3x at high frequencies in the new design. Bottom: S_{41} (nearest neighbor far-end crosstalk) measurement of sample FLAX traces from the same cables at 4 K. The cross talk values are largely unchanged between design iterations.

Transmission

Transmission measurements (S_{21}) at 4 K for two traces from the same test cable are plotted in Fig. 4.12 along with sample traces from the previous iteration of the cable first published in [145]. The transmission is improved from at most 7 dB of loss in [145] to 1.5 dB of loss at 8 GHz in this work. The loss is now on-par with commercial semi-rigid superconducting coax cables and 3x better than commercial NbTi-on-polyimide cables (see Table 4.2).

Previously, the authors had hypothesized ripples in the transmission spectrum were caused by inductance in the connector transition that produced reflections off the cable ends [145]. The transmission curves from this work do not show the same ripples increasing with frequency as the previous iteration, suggesting the capacitive tuning in the transition board successfully mitigated the inductive connectorization.

An updated time domain reflectometry measurement is shown in Fig. 4.13 and confirms the larger center conductor wire lowered the impedance from $\sim 60 \Omega$ to $\sim 53 \Omega$. The impedance being closer to 50Ω is likely contributing to the improved transmission at lower frequencies where the insertion loss dominates over the inductive ripple. The improved characteristic impedance and transmission spectra are consistent with 3D electromagnetic simulations the authors used to model design modifications to the center wire and transition board.

Crosstalk

The authors measured both the near-end crosstalk (S_{31}) and the far-end crosstalk (S_{41}) and found the crosstalk levels to be the same. This suggests all crosstalk is hap-

¹⁹XMA Corporation-Omni Spectra, 7 Perimeter Road, Manchester, NH.

P/N: 2082-6040-03-CRYO

²⁰Koaxis RF Cable Assemblies, 2081 Lucon Road, Schwenksville, PA.

P/N: AO10-CC047C-YO18




Figure 4.13: Time Domain Reflectometry (TDR) measurement of the cryogenic signal path showing the characteristic impedance at lengths along the signal path at 4 K. Traces are shown from [145] (dashed, light blue) and this work (dark blue). This work shows improved impedance match with the cable characteristic impedance being closer to $50\ \Omega$ than the previous iteration. Commercially available $50\ \Omega$ coaxial cables border the FLAX cables highlighted by the double arrow. Note the 4 K TDR measurement is accurate to $\pm 3\ \Omega$.

pening in the transition board and connector ends. This is consistent with the idea that the flexible coaxial ribbon cable provides a 3D superconducting ground around each wire making crosstalk in the cable itself virtually impossible. Only (S_{41}) is shown in Fig. 4.12 for comparison with [145]. We find the crosstalk remained around -50 dB which is competitive with commercial semi-rigid and Kapton-flex cables. We hypothesize the crosstalk may be further reduced by adding 3D shielding on the surface of the transition board.

Thermal Conductivity

Following previous convention, we compute $G(T)$, a length-dependent cable thermal conductivity, by summing literature values of cable materials weighted by their cross sections (see Fig. 4.14) [188, 145, 184]. The heat load between temperature stages can be computed by integrating values in Fig. 4.14 from T_1 to T_2 ($T_1 < T_2$) and dividing by the cable length. Results of this calculation are presented for select temperatures and cables in Table 4.2. $G(T)$ is plotted for the superconducting coaxial ribbon cable in this work, the previous iteration [145], and the smallest commercially available superconducting coax from CryoCoax (P/N: BCB012) [182] in Fig. 4.14. Delft Circuits declined request to share cross-sectional area information for the superconducting polyimide Cri/oFlex[®]3 cable (Delft Circuits, personal communication, May 2023) and so we instead show a hypothetical NbTi-on-polyimide cable where the polyimide²¹ cable dielectric and conductor have the same cross sectional areas as the dielectric and conductor in this work, respectively. In this way, the differences in heat load between the three superconducting coaxes in Fig. 4.14 are purely geometrical whereas the difference between this work and the polyimide/Kapton is due solely to materials differences in the dielectric.

We find the heat load increased slightly between [145] and this work due to the

²¹approximated as the Nikaflex DuPont/Nikkan Kapton film [190] used in [184].




Figure 4.14: We computed a cable thermal conductivity $G(T)$ in units of μWcmK^{-1} by summing the thermal conductivity of each constituent material weighted by the cross section [188, 184, 145]. The smallest commercially available superconducting coax from CryoCoax (P/N: BCB012) labeled “semi-rigid” (pink) is compared with this work (salmon). The authors were unable to obtain cross-sectional dimensions for the Delft Cri/oFlex[®]3 cable and so instead display a fictitious NbTi/Kapton cable, labeled “Kapton-flex” (purple), where the Kapton and NbTi materials have been assigned the same cross-sectional area as the dielectric and conductor in this work, respectively. The differences in heat load between the semi-rigid (pink), [145] (gold), and this work (salmon) are due purely to geometrical differences whereas the difference between the Kapton-flex (purple) and this work (salmon) is due purely to difference in dielectric materials. Solid lines are computed using literature values for Nb47Ti [178], PTFE [188], and Kapton [184, 189]. PTFE values were used to estimate the PFA dielectric in the FLAX cable. Dashed lines indicate extrapolation. Figure is modified from [145].

increased center conductor diameter (see Table 4.2). Despite the increase, the cable presented in this work remains roughly half as thermally conductive as the smallest commercially available superconducting coax. While the heat load of the Cri/oFlex[®]3 cables with NbTi conductor has not been directly measured, Delft Circuits suggests it may be $\sim 10\%$ of the heat load of the same polyimide-flex cables with Ag conductor, i.e. 6x more thermally conductive from 1 K to 100 mK than the cable presented in this work (Delft Circuits, personal communication, May 2023) [186]. This may be because the polyimide material is the dominant cross section and is approximately 5 times more thermally conductive than NbTi below 1 K [178]. In this work, we hypothesize the cable heat load is dominated by the NbTi conductor because [178] suggests NbTi is approximately ten-times more thermally conductive than PTFE below 1 K. To minimize heat load in cryogenic cables, we suggest NbTi/PTFE/PFA solutions minimize conductor cross section whereas NbTi/Kapton/polyimide technologies minimize dielectric cross section.

4.2.5 Conclusion

We have improved the transmission 3x without negatively impacting the heat load or isolation and made progress on commercial manufacturing of the superconducting flexible coaxial ribbon cable presented in [145]. We adjusted the center conductor diameter and re-designed the transition board to improve the transmission loss while retaining strong signal isolation and low heat load.

The cable technology presented in this work has been deployed in two superconducting-array-based instruments: the MKID Exoplanet Camera (MEC) [41] and the MagAO-X MKID Instrument (XKID) [43]. In both cases, FLAX cables replaced our group's previously developed laminated NbTi-on-Kapton microstrip cables [184]. In MEC, the improvement in signal integrity increased the amount of usable superconducting detectors

in the array by $\sim 20\%$ [145]. In XKID, the average MKID energy resolution— a key superconducting detector performance metric, increased by a factor of 4.

This cable technology has already demonstrated promising performance gains in superconducting detector array instruments. With the modifications in this work along with future improvements and commercialization by Maybell Quantum Industries, we see the potential for this technology to be widely applicable to superconducting device systems. Future improvements may include adding 3D shielding in the transition board which would allow for higher connector pitch density, lowering the head load per trace and potentially improving the crosstalk.

Overall, we find this cable technology to be superior to commercial options for our applications building high-density superconducting detector arrays. The high electrical and thermal isolation coupled with low loss and connectorized with high-density, coaxial push-on connectors makes for an easy-to-use cryogenic wiring solution. We are looking forward to future improvements led by Maybell Quantum Industries and we expect these results will be especially promising for groups looking to operate high-density, superconducting microwave device arrays.

Chapter 5

Conclusion

This dissertation has reviewed superconducting detector utility with a focus on applications to fast, single-photon-counting, energy-resolving cameras for scientific imaging. In order to fully realize this scientific utility, superconducting detectors must be fabricated into large arrays. MKIDs are uniquely suited to tile into arrays because of their frequency-division multiplexing ability, however large MKID arrays still face many engineering challenges.

This work details a new digital readout system which is capable of reading out twice the number of pixels with a fraction of the weight, volume, and power as compared to the previous system. The RFSoc-based digital readout makes use of modern FPGA tools and is more flexible and maintainable than the previous system.

This work also featured a new superconducting ribbon cable for high-density, cryogenic wiring. The cable had half the heat load of the best commercially available option at the time of publication and is now under commercial development. Together, these technologies are expected to pave the way towards megapixel arrays, enabling new science and positioning MKIDs as a technology target for future space missions.

Overall, it has been a joy and a privilege to work on these exciting, up-and-coming

detectors. I am looking forward to continuing to develop experimental electronics solutions for research projects in my future career. I will be closely tracking MKID research progress and am eagerly anticipating the next wave of technology-enabled science.

Bibliography

- [1] J. Lillington, “Comparison of Wideband Channelisation Architectures”, in *Int. signal processing conf.* Dallas, 2003. [Cited on pages xvi, 78, and 79]
- [2] F. Arute et al., “Quantum supremacy using a programmable superconducting processor”, *Nature* **574** (October, 2019) 505–510, doi:10.1038/s41586-019-1666-5. Publisher: Nature Publishing Group. [Cited on pages 1 and 121]
- [3] P. Krantz et al., “A quantum engineer’s guide to superconducting qubits”, *Applied Physics Reviews* **6** (June, 2019) 021318, doi:10.1063/1.5089550. Publisher: American Institute of Physics. [Cited on page 1]
- [4] M. McEwen et al., “Resolving catastrophic error bursts from cosmic rays in large arrays of superconducting qubits”, *Nature Physics* **18** (January, 2022) 107–111, doi:10.1038/s41567-021-01432-8. Publisher: Nature Publishing Group. [Cited on page 1]
- [5] A. P. Vepsäläinen et al., “Impact of ionizing radiation on superconducting qubit coherence”, *Nature* **584** (August, 2020) 551–556, doi:10.1038/s41586-020-2619-8. Publisher: Nature Publishing Group. [Cited on page 1]
- [6] S. Todaro et al., “State Readout of a Trapped Ion Qubit Using a Trap-Integrated Superconducting Photon Detector”, *Physical Review Letters* **126** (January, 2021) 010501, doi:10.1103/PhysRevLett.126.010501. Publisher: American Physical Society. [Cited on pages 1 and 23]
- [7] S. Paesani et al., “Generation and sampling of quantum states of light in a silicon chip”, *Nature Physics* **15** (September, 2019) 925–929, doi:10.1038/s41567-019-0567-8. Publisher: Nature Publishing Group. [Cited on page 1]
- [8] E. Burstein, D. N. Langenberg, and B. N. Taylor, “Superconductors as Quantum Detectors for Microwave and Sub-Millimeter-Wave Radiation”, *Physical Review Letters* **6** (February, 1961) 92–94, doi:10.1103/PhysRevLett.6.92. Publisher: American Physical Society. [Cited on pages 1 and 7]

- [9] P. L. Richards, “Bolometers for infrared and millimeter waves”, *Journal of Applied Physics* **76** (July, 1994) 1–24, doi:10.1063/1.357128. [Cited on page 1]
- [10] A. Peacock et al., “Single optical photon detection with a superconducting tunnel junction”, *Nature* **381** (May, 1996) 135–137, doi:10.1038/381135a0. Publisher: Nature Publishing Group. [Cited on page 1]
- [11] P. K. Day et al., “A broadband superconducting detector suitable for use in large arrays”, *Nature* **425** (oct, 2003) 817–821, doi:10.1038/nature02037. [Cited on pages 1 and 85]
- [12] K. Irwin and G. Hilton, “Transition-Edge Sensors”, in *Cryogenic Particle Detection*, C. Enss, ed., pp. 63–150. Springer, Berlin, Heidelberg, 2005. doi:10.1007/10933596_3. [Cited on pages 1 and 6]
- [13] F. Xia et al., “Short-Wave Infrared Confocal Fluorescence Imaging of Deep Mouse Brain with a Superconducting Nanowire Single-Photon Detector”, *ACS Photonics* **8** (September, 2021) 2800–2810, doi:10.1021/acsp Photonics.1c01018. Publisher: American Chemical Society. [Cited on pages 2 and 23]
- [14] K. Niwa, T. Numata, K. Hattori, and D. Fukuda, “Few-photon color imaging using energy-dispersive superconducting transition-edge sensor spectrometry”, *Scientific Reports* **7** (April, 2017) 45660, doi:10.1038/srep45660. Publisher: Nature Publishing Group. [Cited on pages 2 and 23]
- [15] J. Zmuidzinas, “Superconducting Microresonators: Physics and Applications”, *Annual Review of Condensed Matter Physics* **3** (March, 2012) 169–214, doi:10.1146/annurev-conmatphys-020911-125022. Publisher: Annual Reviews. [Cited on pages 2 and 23]
- [16] J. A. B. Mates et al., “Simultaneous readout of 128 X-ray and gamma-ray transition-edge microcalorimeters using microwave SQUID multiplexing”, *Applied Physics Letters* **111** (August, 2017) 062601, doi:10.1063/1.4986222. [Cited on page 2]
- [17] P. A. R. Ade et al., “Antenna-Coupled TES Bolometers Used in BICEP2, Keck Array, and SPIDER”, *The Astrophysical Journal* **812** (October, 2015) 176, doi:10.1088/0004-637X/812/2/176. Publisher: The American Astronomical Society. [Cited on pages 2 and 6]
- [18] N. Zobrist, “Improving the Resolving Power of Ultraviolet to NearInfrared Microwave Kinetic Inductance Detectors”. PhD Thesis, University of California, Santa Barbara, Santa Barbara, CA, June, 2022. [Cited on pages 2, 10, 13, 15, 24, and 29]

- [19] J. I. B. Iii et al., “M-MOST: The MKID Multi-Object Echelle(tte) Spectrographic Testbench”, in *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation V*, volume 12188, p. 121884Z. SPIE, August, 2022. doi:10.1117/12.2629519. [Cited on pages 4 and 23]
- [20] A. G. Kozorezov et al., “Electron energy down-conversion in thin superconducting films”, *Physical Review B* **75** (March, 2007) 094513, doi:10.1103/PhysRevB.75.094513. Publisher: American Physical Society. [Cited on pages 5 and 71]
- [21] A. B. Walter, C. Bockstiegel, T. D. Brandt, and B. A. Mazin, “Stochastic Speckle Discrimination with Time-tagged Photon Lists: Digging below the Speckle Noise Floor”, *Publications of the Astronomical Society of the Pacific* **131** (October, 2019) 114506, doi:10.1088/1538-3873/ab389a. Publisher: The Astronomical Society of the Pacific. [Cited on page 5]
- [22] S. Steiger et al., “Probing Photon Statistics in Adaptive Optics Images with SCEXAO/MEC*”, *The Astronomical Journal* **164** (October, 2022) 186, doi:10.3847/1538-3881/ac922f. Publisher: The American Astronomical Society. [Cited on page 5]
- [23] G. Y. Prigozhin et al., “Radiation damage in the Chandra x-ray CCDs”, in *X-Ray Optics, Instruments, and Missions III*, volume 4012, pp. 720–730. SPIE, July, 2000. doi:10.1117/12.391612. [Cited on page 5]
- [24] G. Y. Prigozhin et al., “Characterization of the radiation damage in the Chandra x-ray CCDs”, in *X-Ray and Gamma-Ray Instrumentation for Astronomy XI*, volume 4140, pp. 123–134. SPIE, December, 2000. doi:10.1117/12.409106. [Cited on page 5]
- [25] “Dictionary: CTI: Charge Transfer Inefficiency - CIAO 4.16”. [Cited on page 5]
- [26] S. L. O’Dell et al., “Managing radiation degradation of CCDs on the Chandra X-ray Observatory III”, in *UV, X-Ray, and Gamma-Ray Space Instrumentation for Astronomy XV*, volume 6686, pp. 19–30. SPIE, September, 2007. doi:10.1117/12.734594. [Cited on page 5]
- [27] K. Karatsu et al., “Radiation Tolerance of Aluminum Microwave Kinetic Inductance Detector”, *Journal of Low Temperature Physics* **184** (August, 2016) 540–546, doi:10.1007/s10909-016-1523-y. [Cited on page 6]
- [28] J. N. Ullom and D. A. Bennett, “Review of superconducting transition-edge sensors for x-ray and gamma-ray spectroscopy*”, *Superconductor Science and Technology* **28** (July, 2015) 084003, doi:10.1088/0953-2048/28/8/084003. Publisher: IOP Publishing. [Cited on page 6]

- [29] B. Cabrera et al., “Detection of single infrared, optical, and ultraviolet photons using superconducting transition edge sensors”, *Applied Physics Letters* **73** (August, 1998) 735–737, doi:10.1063/1.121984. [Cited on page 6]
- [30] L. Moncelsi et al., “Receiver development for BICEP Array, a next-generation CMB polarimeter at the South Pole”, in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy X*, volume 11453, pp. 189–206. SPIE, December, 2020. doi:10.1117/12.2561995. [Cited on page 6]
- [31] C. M. Natarajan, M. G. Tanner, and R. H. Hadfield, “Superconducting nanowire single-photon detectors: physics and applications”, *Superconductor Science and Technology* **25** (April, 2012) 063001, doi:10.1088/0953-2048/25/6/063001. Publisher: IOP Publishing. [Cited on page 6]
- [32] S. Gyger et al., “Reconfigurable photonics with on-chip single-photon detectors”, *Nature Communications* **12** (March, 2021) 1408, doi:10.1038/s41467-021-21624-3. Publisher: Nature Publishing Group. [Cited on page 6]
- [33] F. Grünenfelder et al., “Fast single-photon detectors and real-time key distillation enable high secret-key-rate quantum key distribution systems”, *Nature Photonics* **17** (May, 2023) 422–426, doi:10.1038/s41566-023-01168-2. Publisher: Nature Publishing Group. [Cited on page 6]
- [34] “SNSPD technology”. [Cited on page 6]
- [35] “SNSPD”. [Cited on page 6]
- [36] E. E. Wollman et al., “Kilopixel array of superconducting nanowire single-photon detectors”, *Optics Express* **27** (November, 2019) 35279, doi:10.1364/OE.27.035279. [Cited on page 6]
- [37] B. G. Oripov et al., “A superconducting nanowire single-photon camera with 400,000 pixels”, *Nature* **622** (October, 2023) 730–734, doi:10.1038/s41586-023-06550-2. Publisher: Nature Publishing Group. [Cited on page 6]
- [38] Emma Wollman, “New applications for superconducting nanowire single-photon detectors”, 2023. [Cited on page 7]
- [39] B. Hampel, R. Mirin, S. W. Nam, and V. Verma, “Superconducting Nanowire Single-Photon Detector Arrays for the Near- to Mid-Infrared”, *NIST* (October, 2023). Last Modified: 2023-11-15T22:11-05:00 Publisher: Benedikt Hampel, Richard Mirin, Sae Woo Nam, Varun Verma. [Cited on page 7]

- [40] H. Hao et al., “A compact multi-pixel superconducting nanowire single-photon detector array supporting gigabit space-to-ground communications”, *Light: Science & Applications* **13** (January, 2024) 25, doi:10.1038/s41377-023-01374-1. Publisher: Nature Publishing Group. [Cited on page 7]
- [41] A. B. Walter et al., “The MKID Exoplanet Camera for Subaru SCEXAO”, *Publications of the Astronomical Society of the Pacific* **132** (November, 2020) 125005, doi:10.1088/1538-3873/abc60f. Publisher: The Astronomical Society of the Pacific. [Cited on pages 7, 23, 121, 122, and 135]
- [42] S. Steiger et al., “SCEXAO/MEC and CHARIS Discovery of a Low-mass, 6 au Separation Companion to HIP 109427 Using Stochastic Speckle Discrimination and High-contrast Spectroscopy*”, *The Astronomical Journal* **162** (July, 2021) 44, doi:10.3847/1538-3881/ac02cc. Publisher: The American Astronomical Society. [Cited on pages 7 and 23]
- [43] N. J. Swimmer et al., “An MKID camera for use behind MagAO-X”, in *Ground-based and Airborne Instrumentation for Astronomy IX*, volume 12184, p. 121843Z. SPIE, August, 2022. doi:10.1117/12.2629573. [Cited on pages 7 and 135]
- [44] P. Verhoeve et al., “S-Cam 3: Optical astronomy with a STJ-based imaging spectrophotometer”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **559** (April, 2006) 598–601, doi:10.1016/j.nima.2005.12.198. [Cited on page 7]
- [45] D. Woody, R. Miller, and M. Wengler, “85 - 115-GHz Receivers for Radio Astronomy”, *IEEE Transactions on Microwave Theory and Techniques* **33** (February, 1985) 90–95, doi:10.1109/TMTT.1985.1132955. Conference Name: IEEE Transactions on Microwave Theory and Techniques. [Cited on page 7]
- [46] T. Buttgenbach et al., “A broad-band low-noise SIS receiver for submillimeter astronomy”, *IEEE Transactions on Microwave Theory and Techniques* **36** (December, 1988) 1720–1726, doi:10.1109/22.17405. Conference Name: IEEE Transactions on Microwave Theory and Techniques. [Cited on page 7]
- [47] P. Scowen, “Large Focal Plane Arrays for Future Missions”,. [Cited on page 8]
- [48] C. Gidney and M. Ekerå, “How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits”, *Quantum* **5** (April, 2021) 433, doi:10.22331/q-2021-04-15-433. Publisher: Verein zur Förderung des Open Access Publizierens in den Quantenwissenschaften. [Cited on page 8]

- [49] S. W. Henderson et al., “Advanced RFSoc readout for space-based superconducting sensor arrays”, in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*, volume 12190, pp. 339–355. SPIE, August, 2022. doi:10.1117/12.2630412. [Cited on page 8]
- [50] C. Yu et al., “SLAC microresonator RF (SMuRF) electronics: A tone-tracking readout system for superconducting microwave resonator arrays”, *Review of Scientific Instruments* **94** (January, 2023) 014712, doi:10.1063/5.0125084. [Cited on pages 8, 70, and 72]
- [51] J. Burney et al., “Transition-edge sensor arrays for UV-optical-IR astrophysics”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **559** (April, 2006) 525–527, doi:10.1016/j.nima.2005.12.060. [Cited on page 8]
- [52] J. Gao, “The Physics of Superconducting Microwave Resonators”. PhD thesis, California Institute of Technology, 2008. [Cited on pages 10 and 24]
- [53] B. A. Mazin, “Microwave Kinetic Inductance Detectors”. PhD thesis, California Institute of Technology, 2005. [Cited on pages 10 and 24]
- [54] B. A. Mazin, “Superconducting Materials for Microwave Kinetic Inductance Detectors”, April, 2020. arXiv:2004.14576 [astro-ph, physics:cond-mat]. doi:10.48550/arXiv.2004.14576. [Cited on pages 10 and 27]
- [55] F. W. Carter, T. S. Khaire, V. Novosad, and C. L. Chang, “scraps: An Open-Source Python-Based Analysis Package for Analyzing and Plotting Superconducting Resonator Data”, *IEEE Transactions on Applied Superconductivity* **27** (June, 2017) 1–5, doi:10.1109/TASC.2016.2625767. Conference Name: IEEE Transactions on Applied Superconductivity. [Cited on page 13]
- [56] K. L. Geerlings, “Improving Coherence of Superconducting Qubits and Resonators”. PhD thesis, Yale University, 2013. [Cited on page 13]
- [57] M. S. Khalil, M. J. A. Stoutimore, F. C. Wellstood, and K. D. Osborn, “An analysis method for asymmetric resonator transmission applied to superconducting devices”, *Journal of Applied Physics* **111** (March, 2012) 054510, doi:10.1063/1.3692073. [Cited on page 13]
- [58] “Overview DSP Macro LogiCORE IP Product Guide (PG323) Reader AMD Technical Information Portal”. [Cited on page 19]
- [59] “LUT6 Vivado Design Suite 7 Series FPGA and Zynq 7000 SoC Libraries Guide (UG953) Reader AMD Technical Information Portal”. [Cited on page 19]

- [60] “Block Memory Generator”. [Cited on page 19]
- [61] “UltraRAM Resources Versal ACAP Memory Resources Architecture Manual (AM007) Reader AMD Technical Information Portal”. [Cited on page 19]
- [62] “Zynq UltraScale+ RFSoc ZCU670 Evaluation Kit”. [Cited on pages 21, 35, 40, and 69]
- [63] J. P. Smith et al., “MKIDGen3: Energy-Resolving, Single-Photon-Counting MKID Readout on an RFSoc”, June, 2024. arXiv:2406.09764. [Cited on page 22]
- [64] M. Mariani et al., “Implementing the Quantum von Neumann Architecture with Superconducting Circuits”, *Science* **334** (October, 2011) 61–65, doi:10.1126/science.1208517. Publisher: American Association for the Advancement of Science. [Cited on page 23]
- [65] B. A. Mazin, “Microwave Kinetic Inductance Detectors: The First Decade”, *AIP Conference Proceedings* **1185** (December, 2009) 135–142, doi:10.1063/1.3292300. [Cited on page 23]
- [66] P. K. Day et al., “A broadband superconducting detector suitable for use in large arrays”, *Nature* **425** (October, 2003) 817–821, doi:10.1038/nature02037. Publisher: Nature Publishing Group. [Cited on page 23]
- [67] N. Zobrist et al., “Wide-band parametric amplifier readout and resolution of optical microwave kinetic inductance detectors”, *Applied Physics Letters* **115** (July, 2019) 042601, doi:10.1063/1.5098469. [Cited on pages 23, 32, 60, 62, 71, and 121]
- [68] N. Zobrist et al., “Membraneless Phonon Trapping and Resolution Enhancement in Optical Microwave Kinetic Inductance Detectors”, *Physical Review Letters* **129** (July, 2022) 017701, doi:10.1103/PhysRevLett.129.017701. Publisher: American Physical Society. [Cited on pages 23 and 71]
- [69] N. Swimmer et al., “SCEXAO and Keck Direct Imaging Discovery of a Low-mass Companion Around the Accelerating F5 Star HIP 5319*”, *The Astronomical Journal* **164** (September, 2022) 152, doi:10.3847/1538-3881/ac85a8. Publisher: The American Astronomical Society. [Cited on page 23]
- [70] N. Fruitwala et al., “Second generation readout for large format photon counting microwave kinetic inductance detectors”, *Review of Scientific Instruments* **91** (December, 2020) 124705, doi:10.1063/5.0029457. Publisher: American Institute of Physics. [Cited on pages 23, 31, 33, 34, 39, and 70]
- [71] P. Szypryt et al., “High quality factor platinum silicide microwave kinetic inductance detectors”, *Applied Physics Letters* **109** (October, 2016) 151102, doi:10.1063/1.4964665. [Cited on page 27]

- [72] G. Coiffard et al., “Characterization of sputtered hafnium thin films for high quality factor microwave kinetic inductance detectors”, *Superconductor Science and Technology* **33** (May, 2020) 07LT02, doi:10.1088/1361-6668/ab8d99. Publisher: IOP Publishing. [Cited on page 27]
- [73] P. Szypryt et al., “Large-format platinum silicide microwave kinetic inductance detectors for optical to near-IR astronomy”, *Optics Express* **25** (October, 2017) 25894–25909, doi:10.1364/OE.25.025894. Publisher: Optica Publishing Group. [Cited on pages 32 and 60]
- [74] N. Zobrist et al., “Improving the dynamic range of single photon counting kinetic inductance detectors”, *Journal of Astronomical Telescopes, Instruments, and Systems* **7** (February, 2021) 010501, doi:10.1117/1.JATIS.7.1.010501. Publisher: SPIE. [Cited on pages 33 and 72]
- [75] M. J. Strader et al., “EXCESS OPTICAL ENHANCEMENT OBSERVED WITH ARCONS FOR EARLY CRAB GIANT PULSES”, *The Astrophysical Journal Letters* **779** (November, 2013) L12, doi:10.1088/2041-8205/779/1/L12. Publisher: The American Astronomical Society. [Cited on page 33]
- [76] S. Guo et al., “High-timing-precision detection of single X-ray photons by superconducting nanowires”, *National Science Review* **11** (April, 2023) nwad102, doi:10.1093/nsr/nwad102. [Cited on page 33]
- [77] “RF-ADC Multi-Band Operation Zynq UltraScale+ RFSoc RF Data Converter v2.6 Gen 1/2/3/DFE LogiCORE IP Product Guide (PG269) Reader AMD Technical Information Portal”. [Cited on page 35]
- [78] John I. Bailey III et al., “MKIDGen3: a scalable readout for next-generation kilopixel UVOIR MKIDs”, SPIE, Yokohama, Japan. [Cited on page 37]
- [79] “Multi-Tile Synchronization Zynq UltraScale+ RFSoc RF Data Converter v2.6 Gen 1/2/3/DFE LogiCORE IP Product Guide (PG269) Reader AMD Technical Information Portal”. [Cited on pages 40 and 67]
- [80] J. P. Smith et al., “A High-Throughput Oversampled Polyphase Filter Bank Using Vivado HLS and PYNQ on a RFSoc”, *IEEE Open Journal of Circuits and Systems* **2** (2021) 241–252, doi:10.1109/OJCS.2020.3041208. Conference Name: IEEE Open Journal of Circuits and Systems. [Cited on pages 41, 47, and 74]
- [81] “Intelligent Design Runs Vivado Design Suite User Guide: Design Analysis and Closure Techniques (UG906) Reader AMD Technical Information Portal”. [Cited on page 46]
- [82] “CORDIC”. [Cited on page 48]

- [83] “FIR Compiler”. [Cited on page 48]
- [84] “MazinLab/wb2axip: Bus bridges and other odds and ends”. [Cited on page 49]
- [85] “PYNQ”. [Cited on page 51]
- [86] F. Faramarzi et al., “A 4-8 GHz Kinetic Inductance Travelling-Wave Parametric Amplifier Using Four-Wave Mixing with Near Quantum-Limit Noise Performance”, June, 2024. arXiv:2402.11751 [astro-ph, physics:quant-ph]. doi:10.48550/arXiv.2402.11751. [Cited on page 60]
- [87] “Image Due to IQ Imbalance Co-location Deployment Considerations for Direct RF Sampling Transceivers (WP541) Reader AMD Technical Information Portal”. [Cited on page 67]
- [88] “Aerospace and Defense Platforms”. [Cited on page 70]
- [89] A. Walter et al., “MEC: the MKID exoplanet camera for high contrast astronomy at Subaru (Conference Presentation)”, p. 31. July, 2018. doi:10.1117/12.2311586. [Cited on page 71]
- [90] A. G. Kozorezov et al., “Quasiparticle-phonon downconversion in nonequilibrium superconductors”, *Physical Review B* **61** (May, 2000) 11807–11819, doi:10.1103/PhysRevB.61.11807. Publisher: American Physical Society. [Cited on page 71]
- [91] A. G. Kozorezov et al., “Phonon Noise in Thin Metal Films in an Advanced Energy Down-Conversion Stage”, *Journal of Low Temperature Physics* **151** (April, 2008) 51–57, doi:10.1007/s10909-007-9612-6. [Cited on page 71]
- [92] P. J. de Visser et al., “Phonon-Trapping-Enhanced Energy Resolution in Superconducting Single-Photon Detectors”, *Physical Review Applied* **16** (September, 2021) 034051, doi:10.1103/PhysRevApplied.16.034051. Publisher: American Physical Society. [Cited on page 71]
- [93] S. Ahmadi, “Toward 5G Xilinx Solutions and Enablers for Next-Generation Wireless Systems Toward 5G Xilinx Solutions and Enablers for Next-Generation Wireless Systems”, technical report, Xilinx, 2016. [Cited on pages 75 and 76]
- [94] G. Sklivanitis, A. Gannon, S. N. Batalama, and D. A. Pados, “Addressing next-generation wireless challenges with commercial software-defined radio platforms”, *IEEE Comm. Mag.* **54** (1, 2016) 59–67, doi:10.1109/MCOM.2016.7378427. [Cited on pages 75 and 76]
- [95] Xilinx, “FPGAs in the Emerging DNN Inference Landscape”, technical report, Xilinx, San Jose, 2019. [Cited on page 75]

- [96] J. Pfau, S. P. D. Figuli, S. Bähr, and J. Becker, “Reconfigurable FPGA-based channelization using polyphase filter banks for quantum computing systems”, in *Lect. Notes in Computer Science*, volume 10824 LNCS, pp. 615–626. Springer Verlag, 2018. doi:10.1007/978-3-319-78890-6{_}49. [Cited on pages 75 and 76]
- [97] C. Guo et al., “Control and Readout Software for Superconducting Quantum Computing”, *IEEE Trans. on Nuclear Science* **66** (7, 2019) 1222–1227, doi:10.1109/TNS.2019.2920337. [Cited on pages 75 and 76]
- [98] S. A. Kernasovskiy et al., “SLAC Microresonator Radio Frequency (SMuRF) Electronics for Read Out of Frequency-Division-Multiplexed Cryogenic Sensors”, *J. of Low Temp. Phys.* **193** (11, 2018) 570–577, doi:10.1007/s10909-018-1981-5. [Cited on page 75]
- [99] J. Tuthill et al., “Compensating for oversampling effects in polyphase channelizers: A radio astronomy application”, in *2015 IEEE Sig. Proc. and Sig. Proc. Edu. Workshop*, pp. 255–260. Institute of Electrical and Electronics Engineers Inc., 12, 2015. doi:10.1109/DSP-SPE.2015.7369562. [Cited on pages 75, 80, 86, 90, and 99]
- [100] H. D. Foster, “FPGA Verification Challenges and Opportunities”, technical report, Mentor, 2018. [Cited on page 75]
- [101] A. Bhutani and S. Yadav, “Field Programmable Gate Array (FPGA) Market Share, 2019-2026 Forecasts”, 11, 2019. [Cited on page 75]
- [102] Xilinx, “Vivado High-Level Synthesis”. [Cited on page 76]
- [103] “PYNQ - Python productivity for Zynq”. [Cited on pages 76 and 93]
- [104] B. Vaz et al., “A 13b 4GS/s Digitally Assisted Dynamic 3-Stage Asynchronous Pipelined-SAR ADC”, in *IEEE Intern. Solid-State Circuits Conf.*, pp. 276–281, Xilinx. San Francisco, 2017. [Cited on page 76]
- [105] A. Finnerty and M. Lee, “Integrated SD-FEC in Zynq UltraScale+ RFSocS for Higher Throughput and Power Efficiency (WP498)”, technical report, Xilinx, 2018. [Cited on page 76]
- [106] Xilinx, “An Adaptable Direct RF Sampling Solution”, technical report, Xilinx, 2019. [Cited on page 76]
- [107] R. G. Machado and A. M. Wyglinski, “Software-defined radio: Bridging the analog-digital divide”, *Proc. of the IEEE* **103** (3, 2015) 409–423, doi:10.1109/JPR0C.2015.2399173. [Cited on page 76]

- [108] F. Harris, “On the relationship between multirate polyphase FIR filters and windowed, overlapped, FFT processing”, in *Conference Record - Asilomar Conference on Circuits, Systems & Computers*, volume 1, pp. 485–488. Publ by Maple Press, Inc, 1989. doi:10.1109/acssc.1989.1200838. [Cited on page 76]
- [109] E. Hewitt and R. E. Hewitt, “The Gibbs-Wilbraham phenomenon: An episode in fourier analysis”, *Archive for History of Exact Sciences* **21** (6, 1979) 129–160, doi:10.1007/BF00330404. [Cited on page 76]
- [110] F. J. Harris, C. Dick, and M. Rice, “Digital receivers and transmitters using polyphase filter banks for wireless communications”, *IEEE Trans. on Microwave Theory and Techniques* **51** (2003), no. 4, 1395–1412, doi:10.1109/TMTT.2003.809176. [Cited on pages 76, 80, 90, and 99]
- [111] A. Ghandour, A. Mansour, H. Alasadi, and W. Ghandour, “Design and Implementation of Polyphase Fast Fourier Transform Channelizer”, in *Int. Wireless Comm. and Mobile Computing*, pp. 613–618. Institute of Electrical and Electronics Engineers Inc., Cyprus, 6, 2020. doi:10.1109/IWCMC48107.2020.9148292. [Cited on page 78]
- [112] S. Berner and P. DeLeon, “FPGA-Based Filterbank Implementation for Parallel Digital Signal Processing”, technical report, New Mexico State University, Las Cruces, 1999. [Cited on page 78]
- [113] S. A. Fahmy and L. Doyle, “Reconfigurable polyphase filter bank architecture for spectrum sensing”, in *Reconfigurable Computing: Architectures, Tools and Applications*, volume 5992, pp. 343–350. Springer, Berlin, Heidelberg, Bangkok, 2010. doi:10.1007/978-3-642-12133-3_{_}32. [Cited on page 78]
- [114] H. Chen, C. H. Xiong, S. A. Zhong, and H. Wang, “FPGA-based efficient programmable polyphase FIR filter”, *J. of Beijing Institute of Technology (English Edition)* **14** (3, 2005) 4–8. [Cited on page 78]
- [115] P. Fiala and R. Linhart, “High performance polyphase FIR filter structures in VHDL language for Software Defined Radio based on FPGA”, in *Int. Conf. on Applied Electronics*, number January, pp. 83–86. IEEE Computer Society, 1, 2015. doi:10.1109/AE.2014.7011674. [Cited on page 78]
- [116] M. U. R. Awan, M. M. Alam, P. Koch, and N. Behjou, “Design and implementation of an FPGA-based multi-standard software radio receiver”, in *25th Norchip Conf.* 2007. doi:10.1109/NORCHP.2007.4481026. [Cited on page 78]
- [117] Xilinx, “Polyphase Filter Bank Channelizer”, technical report, Xilinx, San Jose, 2013. [Cited on page 78]

- [118] MATLAB & Simulink, “High Throughput Channelizer for FPGA”. [Cited on page 78]
- [119] F. Harris, “Multirate Signal Processing for Communication Systems”. Prentice Hall, San Diego, 2004. [Cited on pages 80 and 90]
- [120] M. Strader, “Digital Readout for Microwave Kinetic Inductance Detectors and Applications in High Time Resolution Astronomy”. PhD thesis, University of California, Santa Barbara, Santa Barbara, 2016. [Cited on page 80]
- [121] S. W. Henderson et al., “Highly-multiplexed microwave SQUID readout using the SLAC Microresonator Radio Frequency (SMuRF) electronics for future CMB and sub-millimeter surveys”, in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, J. Zmuidzinas and J.-R. Gao, eds., volume 10708, p. 43. SPIE, 7, 2018. doi:10.1117/12.2314435. [Cited on pages 80 and 81]
- [122] X. Chen, F. J. Harris, E. Venosa, and B. D. Rao, “Non-maximally decimated analysis/synthesis filter banks: Applications in wideband digital filtering”, *IEEE Trans. on Signal Processing* **62** (2, 2014) 852–867, doi:10.1109/TSP.2013.2295549. [Cited on page 80]
- [123] W. E. Wilson et al., “The Australia Telescope Compact Array Broad-band Backend: Description and first results”, *Monthly Notices of the Royal Astronomical Society* **416** (9, 2011) 832–856, doi:10.1111/j.1365-2966.2011.19054.x. [Cited on page 80]
- [124] G. Hampson et al., “Xilinx FPGAs Beam Up Next-Gen Radio Astronomy”, *Xilinx Xcell* (2011), no. 75,. [Cited on page 80]
- [125] L. Guo et al., “Analysis and Optimization of the Implicit Broadcasts in FPGA HLS to Improve Maximum Frequency”, in *Int. Symp. on Field-Programmable Gate Arrays*, pp. 311–311. Association for Computing Machinery, New York, NY, USA, 2, 2020. doi:10.1145/3373087.3375332. [Cited on page 81]
- [126] H. M. Makrani et al., “Pyramid: Machine learning framework to estimate the optimal timing and resource usage of a high-level synthesis design”, in *Proc. 29th Int. Conf. on Field-Programmable Logic and Applications*, pp. 397–403. Institute of Electrical and Electronics Engineers Inc., 9, 2019. doi:10.1109/FPL.2019.00069. [Cited on page 81]
- [127] S. Summers et al., “Fast inference of Boosted Decision Trees in FPGAs for particle physics”, *J. of Instrumentation* **15** (2, 2020) doi:10.1088/1748-0221/15/05/p05026. [Cited on page 81]

- [128] T. Marc-André, “Two FPGA Case Studies Comparing High Level Synthesis and Manual HDL for HEP applications”, in *IEEE NPSS Real Time Conf.*, pp. 1–3. Williamsburg, 6, 2018. [Cited on page 81]
- [129] M. A. Mansoori and M. R. Casu, “Efficient FPGA Implementation of PCA Algorithm for Large Data using High Level Synthesis”, in *15th Conf. on Ph.D. Research in Microelectronics and Electronics*, pp. 65–68. Institute of Electrical and Electronics Engineers Inc., 7, 2019. doi:10.1109/PRIME.2019.8787782. [Cited on page 81]
- [130] A. Cukierman et al., “Microwave Multiplexing on the Keck Array”, *J. of Low Temperature Physics* **199** (5, 2020) 858–866, doi:10.1007/s10909-019-02296-2. [Cited on page 81]
- [131] P. Szypryt et al., “Large-format platinum silicide microwave kinetic inductance detectors for optical to near-IR astronomy”, *Optics Express* **25** (2017), no. 21, 25894, doi:10.1364/oe.25.025894. [Cited on page 85]
- [132] Xilinx, “Zynq UltraScale+ RFSoc RF Data Converter v2.3 LogiCORE IP Product Guide”, technical report, Xilinx, 10, 2020. [Cited on page 86]
- [133] Xilinx, “AXI Reference Guide”, technical report, Xilinx, Inc., San Jose, 3, 2011. [Cited on page 87]
- [134] Xilinx, “Vivado Design Suite Reference Guide: Model-Based DSP Design Using System Generator”, 2019. [Cited on page 91]
- [135] J. Goldsmith et al., “Control and Visualisation of a Software Defined Radio System on the Xilinx RFSoc Platform Using the PYNQ Framework”, *IEEE Access* **8** (2020) 129012–129031, doi:10.1109/ACCESS.2020.3008954. [Cited on page 93]
- [136] E. Wang, J. J. Davis, and P. Y. Cheung, “A PYNQ-Based Framework for Rapid CNN Prototyping”, in *Proc. 26th IEEE Int. Symposium on Field-Programmable Custom Computing Machines*, p. 223. Institute of Electrical and Electronics Engineers Inc., 9, 2018. doi:10.1109/FCCM.2018.00057. [Cited on page 93]
- [137] H. Bingo, “Development of a Control Target Recognition for Autonomous Vehicle Using FPGA with Python”, in *Proc. 2018 Int. Conf. on Field-Programmable Technology*, pp. 422–423. Institute of Electrical and Electronics Engineers Inc., 12, 2018. doi:10.1109/FPT.2018.00089. [Cited on page 93]
- [138] “Binary fixed point library for Python”. [Cited on page 94]
- [139] S. Cappello et al., “Flexible Channel Extractor For Wideband Systems Based On Polyphase Filter Bank”, *J. of Theoretical and Applied Information Technology* **31** (2017), no. 16,. [Cited on page 94]

- [140] T. Koehn, “FPGA Based Polyphase Filter Bank Channelizers”, in *GNURadio Conf.*, Virginia Tech. Washington, D.C., 9, 2014. [Cited on pages 94 and 96]
- [141] G. Comoretto et al., “Signal Processing Firmware for the Low Frequency Aperture Array”, *J. of Astronomical Instrumentation* **6** (2, 2020) doi:10.1142/S2251171716410154. [Cited on page 94]
- [142] N. Karcher et al., “SDR-Based Readout Electronics for the ECHo Experiment”, *J. of Low Temperature Physics* **200** (9, 2020) 261–268, doi:10.1007/s10909-020-02463-w. [Cited on pages 94 and 95]
- [143] Xilinx, “Virtex-6 Family Overview”, technical report, Xilinx, Inc., San Jose, 20015. [Cited on page 96]
- [144] Xilinx, “Kintex UltraScale FPGAs Data Sheet: DC and AC Switching Characteristics”, technical report, San Jose, 2020. [Cited on page 96]
- [145] J. P. Smith et al., “Flexible Coaxial Ribbon Cable for High-Density Superconducting Microwave Device Arrays”, *IEEE Transactions on Applied Superconductivity* **31** (January, 2021) 1–5, doi:10.1109/TASC.2020.3008591. Conference Name: IEEE Transactions on Applied Superconductivity. [Cited on pages 108, 122, 123, 124, 126, 127, 128, 129, 130, 131, 132, 133, 134, 135, and 136]
- [146] J. P. Smith et al., “Improved Flexible Coaxial Ribbon Cable for High-Density Superconducting Arrays”, *IEEE Transactions on Applied Superconductivity* **34** (March, 2024) 1–6, doi:10.1109/TASC.2024.3350516. [Cited on pages 108 and 120]
- [147] R. Barends et al., “Superconducting quantum circuits at the surface code threshold for fault tolerance”, *Nature* **508** (2014), no. 7497, 500–503, doi:10.1038/nature13171. [Cited on page 109]
- [148] R. Barends et al., “Digitized adiabatic quantum computing with a superconducting circuit”, *Nature* **534** (2016), no. 7606, 222–226, doi:10.1038/nature17658. [Cited on page 109]
- [149] N. Ofek et al., “Extending the lifetime of a quantum bit with error correction in superconducting circuits”, *Nature* **536** (7, 2016) 441–445, doi:10.1038/nature18949. [Cited on page 109]
- [150] X. Gu et al., “Microwave photonics with superconducting quantum circuits”, 11, 2017. doi:10.1016/j.physrep.2017.10.002. [Cited on page 109]
- [151] E. E. Wollman et al., “Kilopixel array of superconducting nanowire single-photon detectors”, *Optics Express* **27** (11, 2019) 35279, doi:10.1364/oe.27.035279. [Cited on page 109]

- [152] G. Ulbricht et al., “Highly multiplexible thermal kinetic inductance detectors for x-ray imaging spectroscopy”, *Applied Physics Letters* **106** (6, 2015) 251103, doi:10.1063/1.4923096. [Cited on page 109]
- [153] W. S. Holland et al., “SCUBA-2: The 10 000 pixel bolometer camera on the James Clerk Maxwell Telescope”, *Monthly Notices of the Royal Astronomical Society* **430** (2013), no. 4, 2513–2533, doi:10.1093/mnras/sts612. [Cited on page 109]
- [154] M. Calvo et al., “The NIKA2 Instrument, A Dual-Band Kilopixel KID Array for Millimetric Astronomy”, *Journal of Low Temperature Physics* **184** (8, 2016) 816–823, doi:10.1007/s10909-016-1582-0. [Cited on page 109]
- [155] “ARCONS: A 2024 Pixel Optical through Near-IR Cryogenic Imaging Spectrophotometer”, *Publications of the Astronomical Society of the Pacific* **125** (11, 2013) 1348–1361, doi:10.1086/674013. [Cited on page 109]
- [156] S. R. Meeker et al., “Design and development status of MKID integral field spectrographs for high contrast imaging”, in *Adaptive Optics for Extremely Large Telescopes Conf. Proc.*, volume 1. Los Angeles, 2015. doi:10.20353/K3T4CP1131701. [Cited on page 109]
- [157] S. R. Meeker et al., “DARKNESS: A Microwave Kinetic Inductance Detector Integral Field Spectrograph for High-contrast Astronomy”, doi:10.1088/1538-3873/aab5e7. [Cited on page 109]
- [158] C. G. Pappas et al., “High-Density Superconducting Cables for Advanced ACTPol”, *Journal of Low Temperature Physics* **184** (7, 2016) 473–479, doi:10.1007/s10909-015-1454-z. [Cited on page 109]
- [159] D. B. Tuckerman et al., “Flexible superconducting Nb transmission lines on thin film polyimide for quantum computing applications”, *Superconductor Science and Technology* **29** (7, 2016) 084007, doi:10.1088/0953-2048/29/8/084007. [Cited on page 109]
- [160] S. Zou et al., “Low-loss cable-to-cable parallel connection method for thin-film superconducting flexible microwave transmission lines”, *Superconductor Science and Technology* **32** (5, 2019) 075006, doi:10.1088/1361-6668/AB1825. [Cited on page 109]
- [161] A. B. Walter, C. Bockstiegel, B. A. Mazin, and M. Daal, “Laminated NbTi-on-kapton microstrip cables for flexible sub-kelvin RF electronics”, *IEEE Transactions on Applied Superconductivity* **28** (2018), no. 1, doi:10.1109/TASC.2017.2773836. [Cited on pages 109, 110, 115, 117, and 118]

- [162] V. Gupta et al., “Thin-film Nb/Polyimide superconducting stripline flexible cables”, *IEEE Transactions on Applied Superconductivity* **29** (8, 2019) doi:10.1109/TASC.2019.2904203. [Cited on page 109]
- [163] A. B. Walter et al., “The MKID Exoplanet Camera for Subaru SCEXAO”, *In Prep.* (2020). [Cited on pages 110, 111, 117, and 119]
- [164] A. B. Walter, “MEC: The MKID Exoplanet Camera for High Speed Focal Plane Control at the Subaru Telescope”. PhD thesis, University of California, Santa Barbara, 2019. [Cited on pages 111 and 119]
- [165] F. Gisin, “Characterizing Lossy PCB Interconnects Using a TDR Instrument”, Technical Report March, Multek, New Territories, Hong Kong, 2017. [Cited on page 114]
- [166] CryoCoax, “Cryogenic Cable and Cable Assemblies — CryoCoax”. [Cited on page 115]
- [167] KEYCOM, “Superconducting coaxial cable assemblies(0.085” type NbTi-NbTi superconducting coaxial cables)”. [Cited on page 115]
- [168] Emerson, “PTFE and PFA similarities and differences”, technical report, Emerson Automation Solutions, Shakopee, 2017. [Cited on pages 115, 117, and 118]
- [169] A. Kushino, M. Ohkubo, and K. Fujioka, “Thermal conduction measurement of miniature coaxial cables between 0.3 and 4.5 K for the wiring of superconducting detectors”, *Cryogenics* **45** (2005), no. 9, 637–640, doi:10.1016/j.cryogenics.2005.07.002. [Cited on pages 117 and 118]
- [170] M. Daal et al., “Properties of selected structural and flat flexible cabling materials for low temperature applications”, *Cryogenics* **98** (3, 2019) 47–59, doi:10.1016/j.cryogenics.2018.10.019. [Cited on page 118]
- [171] J. R. Olson, “Thermal conductivity of some common cryostat materials between 0.05 and 2 K”, *Cryogenics* **33** (1993), no. 7, 729–731, doi:10.1016/0011-2275(93)90027-L. [Cited on page 118]
- [172] N. Kellaris et al., “Sub-kelvin thermal conductivity and radioactivity of some useful materials in low background cryogenic experiments”, in *Journal of Low Temperature Physics*, volume 176, pp. 201–208. Springer New York LLC, 1, 2014. doi:10.1007/s10909-013-1048-6. [Cited on page 118]
- [173] “Scaled Production of High-Density Cryogenic Flexible coAXial (FLAX) RF Ribbon Cables | SBIR.gov”. [Cited on page 120]
- [174] “Flexlines - The best RF wiring for scaled quantum systems”. [Cited on page 120]

- [175] R. Acharya et al., “Suppressing quantum errors by scaling a surface code logical qubit”, *Nature* **614** (February, 2023) 676–681, doi:10.1038/s41586-022-05434-1. Number: 7949 Publisher: Nature Publishing Group. [Cited on page 121]
- [176] T. White et al., “Readout of a quantum processor with high dynamic range Josephson parametric amplifiers”, *Applied Physics Letters* **122** (January, 2023) 014001, doi:10.1063/5.0127375. Publisher: American Institute of Physics. [Cited on page 121]
- [177] “IBM promises 1000-qubit quantum computer—a milestone—by 2023”. [Cited on page 121]
- [178] M. Daal et al., “Properties of selected structural and flat flexible cabling materials for low temperature applications”, *Cryogenics* **98** (March, 2019) 47–59, doi:10.1016/j.cryogenics.2018.10.019. [Cited on pages 121, 134, and 135]
- [179] “Working With Niobium - Ganoksin Jewelry Making Community”. [Cited on page 121]
- [180] “COWIE TECHNOLOGY - PTFE: High Thermal Stability”. [Cited on page 121]
- [181] “Teflon (PTFE) Thermal Decomposition Products. Fluoride Action Network Pesticide Project.”. [Cited on page 121]
- [182] “Cryogenic Particle Detection”. [Cited on pages 121, 122, 127, 129, and 133]
- [183] “SC-086/50-NbTi-NbTi | COAX CO., LTD.”. [Cited on page 121]
- [184] A. B. Walter, C. Bockstiegel, B. A. Mazin, and M. Daal, “Laminated NbTi-on-Kapton Microstrip Cables for Flexible Sub-Kelvin RF Electronics”, *IEEE Transactions on Applied Superconductivity* **28** (January, 2018) 1–5, doi:10.1109/TASC.2017.2773836. Conference Name: IEEE Transactions on Applied Superconductivity. [Cited on pages 122, 133, 134, and 135]
- [185] G. A. Hernandez et al., “Microwave Performance of Niobium/Kapton Superconducting Flexible Cables”, *IEEE Transactions on Applied Superconductivity* **27** (June, 2017) 1–4, doi:10.1109/TASC.2016.2645683. Conference Name: IEEE Transactions on Applied Superconductivity. [Cited on page 122]
- [186] “Cri/oFlex® 3 | Flexible RF Multi Channel | Cryo Multi RF”. [Cited on pages 122, 123, 127, 129, and 135]

- [187] S. Meeker et al., “Design and Development Status of MKID Integral Field Spectrographs for High Contrast Imaging”, *Adaptive Optics for Extremely Large Telescopes 4 – Conference Proceedings* **1** (2015), no. 1, doi:10.20353/K3T4CP1131701. [Cited on page 122]
- [188] A. Kushino, M. Ohkubo, and K. Fujioka, “Thermal conduction measurement of miniature coaxial cables between 0.3 and 4.5K for the wiring of superconducting detectors”, *Cryogenics* **45** (September, 2005) 637–640, doi:10.1016/j.cryogenics.2005.07.002. [Cited on pages 133 and 134]
- [189] N. Kellaris et al., “Sub-Kelvin Thermal Conductivity and Radioactivity of Some Useful Materials in Low Background Cryogenic Experiments”, *Journal of Low Temperature Physics* **176** (August, 2014) 201–208, doi:10.1007/s10909-013-1048-6. [Cited on page 134]
- [190] “DuPont Pyralux® Flexible CCL”. [Cited on page 133]