

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

XKID: The MKID Camera for High Contrast Direct Imaging at the Magellan Clay Telescope

Permalink

<https://escholarship.org/uc/item/6dc715dc>

Author

Swimmer, Noah Jacob

Publication Date

2023

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

XKID: The MKID Camera for High Contrast Direct Imaging at the Magellan Clay Telescope

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

Doctor of Philosophy
in
Physics

by

Noah Jacob Swimmer

Committee in charge:

Professor Benjamin A. Mazin, Chair
Professor Timothy D. Brandt
Professor S. Peng Oh

June 2023

The Dissertation of Noah Jacob Swimmer is approved.

Professor Timothy D. Brandt

Professor S. Peng Oh

Professor Benjamin A. Mazin, Committee Chair

June 2023

XKID: The MKID Camera for High Contrast Direct Imaging at the Magellan Clay
Telescope

Copyright © 2023

by

Noah Jacob Swimmer

For my friends and family,
Thank you always.

Acknowledgements

I want to use this space to thank everyone who has helped me get to where I am today. Personally and professionally I've been incredibly lucky to be surrounded by an exceptional group who have provided me with endless support and patience to help me grow. These acknowledgements are my attempt to recognize as many people as I can who have been contributors and collaborators through this journey. Science and life don't exist in a vacuum, and my experience is certainly no exception. To everyone who has been with me through my life, my Ph.D. work, and beyond: from the bottom of my heart, I thank you and I am so grateful.

To begin I want to thank my advisor, Ben Mazin, for all of the support and wisdom through my early scientific career. Without his extensive efforts working on and promoting MKID technology none of the work in this thesis would have been possible, and his efforts to create a positive working environment and group atmosphere for all of his students never goes unappreciated or unnoticed. The opportunities I've had working in this lab have been incredibly unique and rewarding.

To all the members of the Mazin group, no matter how long or short our paths here have crossed, you've contributed immensely to creating an exceptional experience here. Jeb Bailey, you've been an invaluable teacher and guide, and you've helped me become much stronger in my craft than I would otherwise be. Sarah Steiger, who has been with me through grad school from the week we started: thank you not only for your wisdom and brilliance here at work, but for your amazing friendship beyond. I'm thrilled to see all of the exciting things you're going to do going forward. To Seth Meeker, Paul Szypryt, Clint Bockstiegel, Alex Walter, Miguel Daal, Greg Coiffard, Kristina Davis, Rupert Dodkins, Isabel Lipartito, Neelay Fruitwala, Nick Zobrist, Jenny Smith, Hawkins Clay, Majid Mohammad, Crystal Kim, Josh Breckenridge, Aled Cuda, Joe Redford, and

Giovanni Strampelli: I could not have had any of the successes I've had here without all of your help. You have all been exceptional teachers, mentors, guides, and friends. Truly, I stand here on the shoulders of giants; you are all exceptional at what you do.

I have also been granted the opportunity to work with several exceptional groups of collaborators at SCEExAO, the University of Arizona, Las Cumbres Observatory (LCO) and UMass Lowell.

From SCEExAO: Thayne Currie, Olivier Guyon, Julien Lozi, Vincent Deo, Sebastian Vievard, Kyohoon Ahn, Nour Skaf, Jeff Chilcote, Tim Brandt, Nemanja Jovanovic, and many more whose work has enabled amazing science with MEC that I've been so lucky to be able to be a part of.

At the University of Arizona: Jared Males, Laird Close, Sebastiaan Haffert, Joseph Long, Jay Kueny, Eden McEwen, Jialin Li, Avalon McLeod, Warren Foster, and the rest of the MagAO-X team who are the masterminds behind an incredible instrument and welcomed me (and XKID) with open arms as a part of the team down at LCO.

At LCO: Alan Uomoto, Povilas Palunas, Roberto Bermudez, Francesco DeMille, Perla Pommer, Jorge Bravo, Victor Meriño, Mauricio Flores, and all of the staff on the mountain and in La Serena who helped me and the instrument that makes up the bulk of this thesis get to and survive up at LCO.

Finally, at UMass Lowell: Chris Mendillo, Thad Potter, Supriya Chakrabarti, Tim Cook, and the PICTURE-C team for letting us be a part of such a unique project. Although we weren't able to get everything integrated in time, I'm incredibly appreciative of the opportunity to visit CSBF and get to work on a camera for such an ambitious project. To all of my collaborators, thank you.

Outside of work, I would never have survived without family and friends – who are my family in their own right. Starting with all of the friends I've made in grad school, you have been essential pieces of my growth here, you've all helped to make me a much

stronger, more thoughtful person. Josh Straub, I'm incredibly lucky to have you as a friend and confidant, I cannot ever thank you enough for all of your support, wisdom, and kindness. Will Schultz, Sarah Steiger, Iman Salty, Josh, Quark and Aramis: thank you for being an absolute foundation, for letting me make your house my second home, going on camping and hiking trips all over the state, running the best picobrewery in Santa Barbara, and so much more; My time here wouldn't be close to the same without you. Joe Costello, Seamus O'Hara, Aleksei Khindanov, Sergio Hernandez Cuenca, Demi Kouvaris, Jon Dorsey, Paige FitzGibbon, Vaughan Andrews, Maddie Julian, thank you for keeping me sane outside of work. Thank you for the book clubs, trips to wine country, beach bonfires, hiking trips (sans mosquitoes), and countless other moments of joy, love, and happiness. You are all wonderful and I cannot wait to see all the great things you do going forward.

From my time at Michigan, I was exceptionally lucky to have found such a strong group of people. Elliott Metzler, who I met less than an hour after moving into my freshman dorm: thank you for being by my side through everything for all of these years. I'd love to recap all of the exceptional memories we have, but that would take a long novel and my space here is limited, so I hope this thank you will suffice. Jillian Goldberg, who I met not long after that: you're one of the sweetest, most genuine people I know and although I'm not always the best at being in touch, when we can hang it's like not a second has passed. Jordan Berman, thank you for your kindness and empathy. You're selfless and want only the best for those around you, and I hope the positivity and grace you put into the world comes back to you tenfold. Eric Stemper, Nick Lowe, Adam Gabalski, Jordan, and Elliott, I couldn't have asked for a better group to have surrounded me during our time in Ann Arbor. As roommates, friends, and – most importantly – brothers, I love you all. It has been a privilege to watch all of you grow, seeing your successes, and I'm even more excited to see where the future takes you.

Going back even further to two of my oldest friends, Karl Kratkiewicz and Adam Rozenberg. I met you both in middle school and it's hard to believe that I'm sitting here now and you're still both such big parts of my life. Karl, watching you grow up, start a family, and have such a fulfilled life fills me with joy. Thank you for letting me stay a part of it for this long. Adam, it's amazing how far we've come. From our closest times to our furthest, you've been one of my rocks. Thank you for being in touch almost daily, whether it's about amazing life-changing news or a new recipe you're trying out. It's fantastic to see all the strides you've taken, and I'm so proud of you.

To my partner, Zoë O'Connor, I could never thank you enough for all of your support and love, but I aim to try. Thank you for sticking with me through several long-distance years between Michigan and California, between Santa Barbara and Los Angeles, and finally making it to the same city once again. Thank you for helping me grow into a much stronger, kinder, and more thoughtful person, and thank you for every little thing that I don't have the space to list but I assure you never goes unnoticed or unappreciated. You've been such a wonderful source of strength, kindness, and love in my life and I hope one day I'm able to match that.

Finally, my family. To my grandfather, "Papa": you are and have always been an inspiration for me. You're the strongest and most generous person I've ever known, and I would be lucky to inherit even a fraction of those traits. I've already been so lucky for all of the time I've gotten to spend with you and hope that continues for as long as humanly possible. To Uncle Howard, thank you for always being such a vocal supporter of what I do, always checking in with me from across the country and keeping tabs on what I'm up to. To my mom and dad, thank you for your unwavering support and guidance. I wouldn't have gotten here without you. I don't have words to express my gratitude for all you've done to help me get to where I am, and I hope that you can tell from all that I do I am deeply grateful. To my sister, Allison, thank you for always being in my corner.

You've always had my back and been one of the most impressive people I've ever had the luck of knowing. I've said it before and I'll say it again, you are by far the best of the Swimmer twins and I always look to you for your excellent advice and perspective. I will always look up to you and I'm so proud of all the great things you've done and the even better things you'll do.

To everyone who has touched and been a part of my life, thank you again. None of this would be possible without your support. Although I can never say it enough, I am so lucky to have been surrounded by such outstanding people. I am forever grateful.

Curriculum Vitæ

Noah Jacob Swimmer

Education

- 2023 Ph.D. in Physics with Astrophysics Emphasis (Expected)
University of California, Santa Barbara
- 2020 M.A. in Physics with Astrophysics Emphasis
University of California, Santa Barbara
- 2017 B.S. in Physics
University of Michigan, Ann Arbor

Selected Publications

Swimmer, N., W.H. Clay, N. Zobrist, B.A. Mazin. “Characterizing the Dark Count Rate of a Large-Format MKID Array.” *Optics Express*, 31, 10775 (2023)

Steiger, S., T.D. Brandt, O. Guyon, **N. Swimmer**, et al. “Probing Photon Statistics in Adaptive Optics Images with SCEXAO/MEC.” *The Astronomical Journal*, 164, 186 (2022)

Swimmer, N., T. Currie, S. Steiger, et al. “SCEXAO and Keck Direct Imaging Discovery of a Low-mass Companion Around the Accelerating F5 Star HIP 5319.” *The Astronomical Journal*, 164, 152 (2022)

Steiger, S., J.I. Bailey III, N. Zobrist, **N. Swimmer**, et al. “The MKID Pipeline: A Data Reduction and Analysis Pipeline for UVOIR MKID Data.” *The Astronomical Journal* 163, 193 (2022)

Abstract

XKID: The MKID Camera for High Contrast Direct Imaging at the Magellan Clay
Telescope

by

Noah Jacob Swimmer

High contrast direct imaging is one of the most technologically challenging techniques used to search for exoplanets and circumstellar disks. Despite its difficulty, it is incredibly powerful since it enables further characterization of these companions. Ground-based high contrast observations require Adaptive Optics (AO) systems to correct wavefront aberrations caused by the atmosphere and starlight suppression techniques like coronagraphy to reject on-axis starlight. The last two decades have seen several extreme AO systems come online, such as SPHERE at ESO, GPI at VLT, SCExAO at Subaru, and MagAO-X at LCO. As these AO systems have matured they have significantly improved achievable contrast ratios, but they are still limited by diffracted and scattered light that causes an interference patterns called speckles in astronomical images. To search for fainter sources these systems will need to improve their ultimate contrast ratios via improved wavefront control and novel data processing techniques.

MKIDs are photon counting detectors capable of measuring a photon's arrival time with microsecond accuracy, its energy (or color), and the position it struck within an image. When used as astronomical cameras MKIDs can generate images with spectral information tens of thousands of times per second without readout noise associated with their conventional counterparts. This ability enables novel noise mitigation techniques and higher signal to noise ratios in astronomical images.

This thesis will present the design, development, and commissioning of XKID, an

MKID-based instrument behind the Magellan Extreme Adaptive Optics (MagAO-X) system at the Las Campanas Observatory (LCO). It is a near-infrared integral field spectrograph (IFS) whose goal is to search for and characterize protoplanets and exoplanets and to eventually operate as a real-time Focal Plane Wavefront Sensor (FPWFS). The first-light commissioning run for XKID took place in March 2023 and is slated to return to LCO during future MagAO-X runs. It is the third MKID camera commissioned for high contrast imaging and the first to be used behind a visible-light AO system.

I will also discuss the discovery and characterization of a potentially substellar companion that was found using the MKID Exoplanet Camera (MEC) at the Subaru Observatory on Maunakea. The discovery was enabled as part of a targeted direct imaging search and was made in conjunction with 3 other instruments on the Subaru and Keck Telescopes: CHARIS, VAMPIRES, and NIRC2.

Next, I will introduce an algorithm that was developed for the identification and rejection of cosmic rays in MKID datasets that has since been successfully integrated into the MKID Data Reduction Pipeline.

Finally, I will report the results of an experiment measuring the dark count rates measured by large-format MKID arrays. This is an essential task for better understanding and characterizing the performance of MKIDs when they are unilluminated which will be integral as they push forward to more photon-starved regimes. This experiment also showed that the nominal performance of MKIDs in the dark is comparable to that of cutting-edge conventional detectors at their quietest.

Contents

| | |
|--|------------|
| Curriculum Vitae | x |
| Abstract | xi |
| List of Tables | xvi |
| 1 Introduction | 1 |
| 1.1 Overview | 1 |
| 1.2 Exoplanets | 3 |
| 1.2.1 Scientific Motivation | 3 |
| 1.2.2 Detection Methods | 5 |
| 1.2.3 High Contrast Direct Imaging | 11 |
| 1.2.4 Adaptive Optics | 14 |
| 1.2.5 Coronagraphy | 18 |
| 1.2.6 Integral Field Spectrographs | 20 |
| 1.2.7 Observing Methods | 20 |
| 1.3 Microwave Kinetic Inductance Detectors | 24 |
| 1.3.1 Operation Principle | 24 |
| 1.3.2 Large-Format MKID Arrays | 32 |
| 1.3.3 MKID Photon Measurement | 32 |
| 1.3.4 Energy Resolution | 34 |
| 1.4 Previous MKID Instruments | 36 |
| 1.4.1 ARCONS | 36 |
| 1.4.2 DARKNESS | 36 |
| 1.4.3 MEC | 38 |
| 1.4.4 The PICTURE-C MKID Camera | 41 |
| 2 XKID: Design | 45 |
| 2.1 Overview | 45 |
| 2.1.1 Magellan Telescopes | 46 |
| 2.2 MagAO-X | 50 |

| | | |
|----------|--|------------|
| 2.3 | Cryostat Design | 54 |
| 2.3.1 | Pulse Tube | 57 |
| 2.3.2 | Adiabatic Demagnetization Refrigerator | 59 |
| 2.4 | Footprint at Magellan Clay Telescope | 62 |
| 2.5 | Mounting Frame | 65 |
| 2.5.1 | Earthquake Stability | 68 |
| 2.6 | Optical Design | 70 |
| 2.6.1 | Foreoptics: MagAO-X to XKID | 71 |
| 2.6.2 | Calibration Snout | 72 |
| 2.6.3 | Cryostat Optics | 75 |
| 2.6.4 | Image at MKID Array | 76 |
| 2.7 | Readout Electronics and Hardware | 77 |
| 2.7.1 | Telemetry and Cryostat Readout | 79 |
| 2.7.2 | MKID Readout | 82 |
| 2.8 | Microwave Signal Path | 85 |
| 2.8.1 | FLAX Cables | 86 |
| 3 | XKID: Operation | 91 |
| 3.1 | Introduction | 91 |
| 3.2 | Instrument Setup | 92 |
| 3.2.1 | Cooling from 300 K to 3 K | 92 |
| 3.2.2 | Cooling from 3 K to 90 mK | 93 |
| 3.2.3 | Finding the MKID Resonators | 97 |
| 3.2.4 | Beammapping | 99 |
| 3.2.5 | Optimal Filtering | 100 |
| 3.3 | Instrument Control Software | 101 |
| 3.3.1 | Devices | 101 |
| 3.3.2 | Agents | 102 |
| 3.3.3 | Service Files | 103 |
| 3.3.4 | Redis | 104 |
| 3.3.5 | Running the XKID Control Software | 107 |
| 3.4 | Instrument Control GUI | 109 |
| 3.4.1 | Initgui | 109 |
| 3.4.2 | Hightemplar | 110 |
| 3.4.3 | MKID Control | 112 |
| 4 | XKID: Commissioning | 114 |
| 4.1 | First Light Run | 114 |
| 4.2 | Calibration | 114 |
| 4.2.1 | Wavelength Calibration | 116 |
| 4.2.2 | Conex to Pixel Mapping | 116 |
| 4.2.3 | Plate Scale and Device Rotation | 118 |

| | | |
|----------|---|------------|
| 4.3 | Data Reduction | 120 |
| 4.4 | Observational Data | 122 |
| 4.5 | Future Work | 126 |
| 5 | The Discovery and Characterization of HIP 5319B | 128 |
| 5.1 | Attribution | 128 |
| 5.2 | Introduction | 128 |
| 5.3 | Stellar Properties and Observations | 130 |
| 5.3.1 | HIP 5319 A Basic Properties | 130 |
| 5.3.2 | Observations and Data Reduction | 133 |
| 5.3.3 | Detections | 138 |
| 5.4 | Analysis | 141 |
| 5.4.1 | Characterization of HIP 5319 A as a Single Star | 141 |
| 5.4.2 | Non-detection of Ca HK Emission | 145 |
| 5.4.3 | Spectrum of HIP 5319 B | 148 |
| 5.4.4 | Orbit and Dynamical Mass | 152 |
| 5.5 | Summary and Discussion | 157 |
| 6 | An Algorithm for Cosmic Ray Rejection | 160 |
| 6.1 | Cosmic Rays | 160 |
| 6.2 | Identifying Cosmic Rays | 163 |
| 6.3 | Cosmic Ray Rejection and Removal | 166 |
| 7 | Characterization of the Dark Count Rate in MKID Arrays | 169 |
| 7.1 | Attribution | 169 |
| 7.2 | Introduction | 169 |
| 7.3 | MKID Photon Measurement | 170 |
| 7.4 | Experiment Overview | 173 |
| 7.4.1 | Optics | 173 |
| 7.4.2 | MKID Array | 174 |
| 7.4.3 | Array Calibration | 175 |
| 7.4.4 | Electronic Readouts | 176 |
| 7.5 | Digital Readout: Many-detector Measurement | 177 |
| 7.5.1 | Experimental Setup | 177 |
| 7.5.2 | Data Collection and Reduction | 178 |
| 7.5.3 | Analysis | 182 |
| 7.6 | Analog Readout: Single-detector Measurements | 187 |
| 7.6.1 | Experimental Setup | 187 |
| 7.6.2 | Data Collection and Reduction | 188 |
| 7.6.3 | Analysis | 189 |
| 7.7 | Discussion and Conclusions | 195 |

List of Tables

| | | |
|-----|---|-----|
| 2.1 | XKID Instrument Design Summary | 46 |
| 2.2 | XKID Selectable Optics Parameters | 72 |
| 4.1 | XKID Commissioning Targets | 122 |
| 5.1 | HIP 5319 Observing Log: Coronagraphs and Photometric Bands | 133 |
| 5.2 | HIP 5319 Observing Log: Observations and Post-processing Strategies | 133 |
| 5.3 | HIP 5319 LCOGT Observing Log | 135 |
| 5.4 | HIP 5319 B Detection Significance and Astrometry | 141 |
| 5.5 | HIP 5319 B Detection Photometry | 141 |
| 5.6 | HIP 5319 B Spectra | 149 |
| 5.7 | HIP 5319 B Orbit Fitting Results and Priors | 153 |
| 5.8 | HIP 5319 B Orbit Fitting Results For Different Priors on Secondary Mass | 155 |
| 7.1 | MKID Resolving Powers | 176 |
| 7.2 | Dark Count Rates With and Without Reduction Steps | 183 |
| 7.3 | Dark Count Rates Between Analog and Digital MKID Readouts | 190 |

Chapter 1

Introduction

1.1 Overview

The focus of this thesis is developing, using, and characterizing astronomical instruments enabled by Microwave Kinetic Inductance Detector (MKIDs; Day et al., 2003). MKIDs are powerful photon-counting detectors that are well suited as a tool for low-count rate applications including high contrast astronomical direct imaging at ultraviolet, optical, and infrared (UVOIR) wavelengths, biophysics, dark matter direct detection experiments, and more (Day et al., 2003; Mazin et al., 2012; Meeker et al., 2018; Walter et al., 2020; Swimmer et al., 2020, 2022b). It is predominantly concerned with using MKID cameras for astronomy, specifically in the pursuit of discovering exoplanets (sometimes called extra-solar planets) and other faint celestial objects that orbit much brighter host stars.

The first chapter is an introduction that will serve as a foundation for the rest of the thesis. It will present the scientific motivation for high contrast direct imaging and the modern technology, techniques, and instruments used to enable it. Next it will describe the operational principles behind the MKID technology used in the cameras developed at

the Mazin Lab at UCSB before introducing several of the MKID cameras that have been commissioned at several telescopes around the Northern Hemisphere. It will provide a broader context for the work presented in the remaining chapters.

Chapter 2 will describe the design and construction of the XKID instrument. XKID (Swimmer et al., 2022b) is a new MKID camera that will sit behind the Magellan Extreme Adaptive Optics (MagAO-X; Close et al., 2018) instrument on the Magellan Clay Telescope at the Las Campanas Observatory (LCO) in the southern Atacama Desert in Chile. Its goal is to perform direct imaging observations of protoplanets, which are planets in the process of forming, and to eventually serve as a Focal Plane Wavefront Sensor (FPWFS; see Fruitwala et al., 2018; Walter et al., 2020) for improved atmospheric correction. The chapter will touch on the MagAO-X mission design and goals, the XKID cryostat and instrument design and construction, and its optical design.

Chapters 3 and 4 will report on the successful XKID commissioning run that took place between February and March of 2023. Chapter 3 will describe the operation procedure for cooling the cryostat and properly configuring the instrument to take astronomical data. Next, it will present an entirely new suite of MKID instrument control software that was developed for the XKID instrument. The goal for this control software was to make something that is both backward- and forward-compatible so as to be easily extensible for other MKID instruments, to be very modular, and to be fairly simple to add to or modify as needed. In Chapter 4 the first on-sky run is detailed and include a discussion of the on-sky calibrations performed that are necessary for astronomical science output, as well as several of the initial observations taken with the instrument.

In chapter 5, the discovery of the potentially substellar companion HIP 5319 B is presented. This was the second companion discovered using the MKID Exoplanet Camera (MEC; Walter et al., 2020) that lives at the Subaru Telescope on Maunakea in Hawai'i and is now commissioned for use by the wider astronomical community.

Chapter 6 focuses on the development and use of a cosmic ray identification and rejection algorithm to mitigate contamination from non-astrophysical sources in MKID datasets. Cosmic rays are high energy particles originating from space that are well known sources of noise in astronomy images, both in conventional astronomy cameras that use CCDs and CMOS detectors and in novel detectors such as MKID instruments. The ability of MKIDs to count single photons, combined with very fast readouts, enables a unique ability to identify and remove these sources from MKID datasets in a way that other types of cameras cannot. The algorithm described within was tested and successfully integrated into the MKID Data Reduction Pipeline (Steiger et al., 2022a).

Lastly, chapter 7 describes the characterization of MKID arrays when they are not illuminated. As MKID technology moves forward they will start to be used for a wider range of targets and harder science goals. The majority of direct imaging targets observed using MKID cameras are still sufficiently bright so as to be limited by noise in the astronomical images themselves, but as projects move toward more photon-starved regimes it will be essential to understand the noise and background in the detectors and ambient environment. This work focuses on characterizing the background and noise sources measured by MKIDs while they are not observing any photon sources in order to better understand their limitations and baseline performance.

1.2 Exoplanets

1.2.1 Scientific Motivation

Since the first discovery (Wolszczan & Frail, 1992) and confirmation (Wolszczan, 1994) of two exoplanets around the Pulsar PSR1257+12 in 1992 and 1994, respectively, the field of exoplanet science has been one of the fastest-growing in astronomy. This




Figure 1.1 The cumulative number of exoplanets discovered by year starting in 1992 and running through 31 March 2023. Each color refers to the method by which a given exoplanet was detected and corresponds to the legend in the top left. Accessed via <https://exoplanetarchive.ipac.caltech.edu/>.

is due in large part to ground-based observatories with 8-meter class telescopes and space-based missions such as Kepler (Borucki, 2016) and TESS (Transiting Exoplanet Survey Satellite Ricker et al., 2015) coming online, new and more sensitive detector technologies, novel noise reduction algorithms and data reduction techniques, and has led to the discovery and confirmation of over 5,300 exoplanets and counting (Figure 1.1). Detecting exoplanets is an extremely difficult challenge due to their distances from Earth and relatively small size and faintness compared to the stars that they orbit.

The increased ability to detect and confirm the existence of exoplanets has led to the realization that exoplanets are abundant and the systems that they exist in are incredibly diverse. This has enabled many statistical analyses of exoplanet populations using their masses, orbital separations, and host star properties (Brandt et al., 2014; Foreman-Mackey et al., 2016; Bowler & Nielsen, 2018) which represents a major step




Figure 1.2 Confirmed exoplanets as a function of planet mass (in Jupiter masses) and orbital period (days). Each color represents a different detection method. Transits and RVs (the two most common) are sensitive to higher-mass planets at relatively short periods, while direct imaging tends to detect planets at wider separations and similarly high mass. Accessed via <https://exoplanetarchive.ipac.caltech.edu/>

toward gaining a deeper understanding of the physics of exoplanets and the systems they form in¹. As detection methods and telescopes become more sensitive to these faint sources, the ability to utilize light that is reflected off of or emitted by an exoplanet may also be used to characterize it more deeply. This means that in addition to learning about the configuration of the star-planet(s) system it is becoming possible to characterize the composition and properties of the atmospheres of some of these planets as well.

1.2.2 Detection Methods

There are several methods that astronomers use to look for exoplanets. The majority of these methods are “indirect”, meaning that the presence of a planet is inferred based

¹The Kepler Orrery (<https://exoplanets.nasa.gov/resources/1018/kepler-orrery-ii/>) provides an illustration showing the exoplanet systems discovered by the Kepler Mission (Borucki, 2016).

on measured properties of its host star, while the “direct” method aims to measure light from the planet itself. While it would be impossible to fully survey each method, a brief description of each along with its merits and drawbacks presented here to provide context for the work in the rest of this thesis. Figure 1.2 shows all of the confirmed exoplanets as a function of their mass and orbital periods, while the color of each shows what method it was detected with. The fact that the different detection methods tend to “clump” together in Mass-Period space is suggestive of the fact that each method thrives in detecting different types of planets and planetary systems.

Radial Velocities

If a star hosts a planet the star will move around the system’s center of mass (called the barycenter) in response to the gravity of the planet. This will cause periodic variations in the radial velocity (RV) of the star as it moves toward or away from Earth. Ultimately, the changing radial velocity manifests as the spectral emission or absorption lines of the star being periodically blue- or redshifted via the Doppler Effect. By determining the the periodicity and magnitude of the blue- and redshifting astronomers can determine system parameters such as the orbital period P and the minimum planet mass, $M\sin(i)$. If the inclination i of the system is known, a more exact mass may be measured.

Although the RV method is the second most successful detection method based on number of planets found it suffers from observational biases due to geometric and detection-driven effects. Since the velocity of the star must be changing along the line of sight to Earth, this method is sensitive to planetary systems that are close to edge-on as viewed from Earth. Additionally, the Doppler shifts created by a planet are small and require sensitive spectrometers to measure accurately. Because more massive planets will create larger Doppler shifts (which also decrease as the system becomes more face-on relative to Earth), the RV method is primarily sensitive to massive planets that orbit close

in to their stars.

Transits

The transit method is by far the most prolific exoplanet detection method, discovering over half of the confirmed exoplanets known today. This is largely because of the Kepler (Borucki, 2016) and TESS (the Transiting Exoplanet Survey Satellite; Ricker et al., 2015) missions. The transit method looks for periodic dips in the brightness of a star as a planet passes in front of it, blocking some of its light. To measure this, light curves of a star are measured, showing its brightness as a function of time. The light curve will then be analyzed for the presence of periodic dips that may indicate the presence of a planet. Additionally, if a transiting planet has a “puffy” atmosphere then as starlight passes through it, its chemical signatures are imprinted on the light in the form of extra emission or absorption lines not present in the stellar spectrum. This enables transit spectroscopy, which is one of several ways a planetary atmosphere may be analyzed.

Like the RV method, the transit method is susceptible to selection bias. Since a transit requires that a planet pass in front of a star, only planetary systems that are nearly edge-on relative to Earth can be observed this way. Because a transit detection also requires the planet to block a significant portion of starlight and multiple iterations of the brightness dip in the light curve, the transit method is more sensitive to massive and large planets that are very close to their stars and have short orbital periods.

Microlensing

Gravitational lensing occurs when the gravity of an object that is closer to an observer acts as a lens, magnifying the light of an object along the same line of sight but further away. This creates a light curve with a characteristic shape, causing a large spike in brightness as the lensing event occurs before falling back to the original brightness.

However, if the foreground star hosts an exoplanet then planet's gravity can contribute a smaller – but still detectable – modification to the characteristic microlensing light curve.

The main drawback to the microlensing method is its exceeding rarity. It doesn't suffer from the same biases as the transit and RV methods, but because it requires a foreground and background star to be almost perfectly aligned with respect to Earth, it does not happen often and when it does occur, it cannot be repeated. However, microlensing tends to be sensitive to planets around less massive stars at wider separations than the previously discussed methods.

Direct Imaging

Direct imaging refers to collecting light from an exoplanet itself, essentially taking a photograph of the planet. Unfortunately, planets are extremely faint light sources relative to the brightness of their host stars. The brightest planets may be $\lesssim 10^{-4}$ times as bright as their hosts while planets like Earth are $\sim 10^{-10}$ as bright as the Sun. This creates a massive challenge for an observer to directly image a planet. Typically, telescopes will use adaptive optics (AO) systems and coronagraphs to overcome some of the challenges associated with blocking starlight and removing noise sources in the images that would otherwise prevent the planet's detection. So far, nearly all of the commissioned MKID instruments have aimed to perform direct imaging observations.

Aside from the technical difficulties associated with imaging very faint sources, the direct imaging method is limited by the inner working angle (IWA) of the telescope, AO, and coronagraph system and the contrast from the astronomical detectors and data reduction algorithms. However, it does not suffer from geometric selection biases because as long as the planet is not directly behind or in front of the star, it is possible that an image may be taken of it. It is also more sensitive to planets at far-flung, long period orbits from their host stars relative to other detection methods.

Timing Variations

The timing variation (or Transit Timing Variation, TTV) method works on the principle that a planet orbiting a star should have a relatively constant periodicity. However, the presence of a second (or beyond) planet may cause gravitational perturbations in the system which cause the periodicity of the first planet's transit to vary. This provides a method to detect successive planets in a system that were not originally measured by the transit or RV method, either because they did not also transit or the signal-to-noise of the brightness dip or Doppler shift were insufficient for a detection. The first exoplanets ever discovered (Wolszczan & Frail, 1992) were found using this method around a pulsar.

Orbital Brightness Modulation

Orbital brightness modulation occurs when some facet of the star-planet system causes the brightness of the system to change as the planet orbits. Two examples of this are (1) tidal forces from a planet deforming the star, causing a smaller or larger portion of it to be visible at different phases of its orbit or (2) starlight causing a close-in planet to drastically heat up, leading to high thermal emission that is unresolvable from the star itself but contributes to the star appearing brighter when the planet is visible and dimmer when it goes behind the star. As with many of the other detection methods, orbital brightness modulation is much more sensitive to massive, close-in planets that exert large forces and are highly influenced by their host stars.

Astrometry

Astrometry involves measuring the precise position of stars in the sky and observing how they change over time. Most stars move in a predictable way through the sky, but when a star hosts a planet that planet may exert a gravitational tug on the star

throughout its orbit that causes the star to “wobble”. Essentially, the orbiting planet imprints a small, periodic perturbation to the position of a star in the sky. This method requires precise position measurements over long periods of time and for the changes in position of the star to be detectable. This means that it is sensitive to stars that are close to Earth (the wobble is more perceptible) and to low-mass stars hosting more massive planets (the tug on the star is more significant).

Astrometry has not been responsible for the detection of many exoplanets, but it is one of the earliest exoplanet detection methods used and one of the only ones that was historically done visually, using hand-drawn records. It is also an incredibly useful tool for pre-selecting targets that *may* host planets to be explored using other, more sensitive methods. This will be discussed in Chapter 5.

Ultimately, each detection method has strengths, weaknesses, and selection biases that lead to the clustering in Figure 1.2. This is to say that none will likely ever prove to be sensitive to all types of planets in all star-planet configurations and that conclusions drawn about exoplanet populations must carefully consider the selection biases and limitations of the detection method in question. Moreover, many of the detection methods are complementary to one another, such as how observing a transit of a planet discovered using RVs may offer an independent measurement of the inclination and therefore provide a more exact planet mass or a planet discovered using RVs may be directly imaged, enabling a dynamical mass to be calculated. It is important to remember that each method has its strengths and all provide a great deal of value as more and more exoplanets are discovered and the detectable parameter space is increased.

In the following sections high contrast direct imaging will be introduced in greater detail, including some of the tools and techniques that are currently being used and developed to search for fainter planets closer in to their host stars.

1.2.3 High Contrast Direct Imaging

Direct imaging is technically demanding but incredibly useful detection method that enables the characterization of faint stellar companions such as exoplanets, low-mass stars, protoplanetary disks, and debris disks. It is also complementary to many of the other more developed exoplanet detection methods such as transits and RVs. While those methods are sensitive to planets that are close in to their host stars, the strength of direct imaging lies in the characterization of exoplanet systems from the outside in, searching for self-luminous planets at wide separations. The next several sections will explore the motivation for direct imaging and several of the most common tools used in modern observatories to search for planets in this way. For a more complete review of the state of the field the reader is referred to the excellent review article and chapter from Bowler (2016) and Currie et al. (2022a), respectively.

The ultimate goal of high contrast imaging (and the majority of exoplanet science) is to detect and characterize rocky, Earth-like planets in the habitable zone of their stars. Figure 1.3 shows the enormous difficulty faced in this task. To directly image a planet like Earth around a sun-like star, one would need to achieve a contrast ratio deeper than $\sim 10^{-10}$ at $\lambda = 1\mu\text{m}$. For even the brightest targets, contrast ratios of $\lesssim 10^{-4}$ are required from the ground at narrow angular separations ($\lesssim 1''$).

Future thirty-meter class telescopes may be capable of detecting these planets around low-mass, nearby stars by enabling deeper contrasts and tighter Inner Working Angles (IWAs) but current ground-based systems are typically limited to contrast ratios of $\sim 10^{-6}$. Although detecting an Earth analog around a sun-like star may ultimately require a space-based telescope (where deeper contrast ratios are easier to achieve) there are still enormous regions of exploration space accessible from ground-based telescopes. Indeed, ground-based high-contrast imaging has already played a critical role in expanding the




Figure 1.3 From Seager & Deming (2010). The modeled blackbody flux of several solar-system objects (Jupiter, Venus, Earth, and Mars) compared to that of a hot Jupiter and sun-like star. For each planet the peak at shorter wavelengths comes from reflected starlight and the peak at longer wavelength comes from its own thermal emission.

understanding planetary architectures by probing separations out to several hundred AU (astronomical units²) and masses greater than $1 M_{\text{Jup}}$ (Bowler, 2016).

Additionally, ground-based telescopes are ideal testing ground for new technologies that provide higher contrasts and better control of the light entering the telescope and hitting the astronomical detectors (adaptive optics systems, section 1.2.4), tighter IWAs (coronagraphs, section 1.2.5), sensitivity to detect and characterize these very faint sources (integral field spectrographs, section 1.2.6), and observing algorithms (section 1.2.7).




Figure 1.4 A generalized diagram of an AO system. Telescope light is sent into the AO system, bouncing off of the deformable mirror (DM) before being split between the WFS and the science camera. The control system takes in WFS data, calculates the wavefront aberration, then tells the DM to change its shape to correct the aberrations. While this is happening, the WFS makes a subsequent measurement of the errors in the corrected wavefront, and the loop runs through again. From the Lyot Project at <https://lyot.org>.

1.2.4 Adaptive Optics

High contrast observations are typically limited by imperfections in the wavefront that ultimately hits the detector. The ideal wavefront is perfectly flat but atmospheric turbulence and imperfections in the telescope and instrument optics lead to the light scattering and diffracting in an uncontrolled manner which manifests as a “speckle halo” in the final image. When the aberrations in the wavefront are not corrected, the diffraction will lead to a seeing limited image that has a point-spread function (PSF) of characteristic size $\sim \lambda/r_0$, where λ is the wavelength of observed light and r_0 is the Fried parameter. The Fried parameter describes the characteristic size of cells of atmospheric turbulence, meaning it is highly dependent on what the atmosphere is like at a given observatory location and in given weather conditions.. A higher r_0 – meaning less turbulence – will correspond to better seeing and a smaller speckle halo, and vice versa.

The size of the speckle halo in a seeing limited image is typically much larger than the core of a diffraction limited image (i.e. the ideal image produced with perfect optics and no aberrations). This happens because the characteristic size of the speckle halo, λ/r_0 is typically larger than the diffraction limit of the telescope, λ/D , where D is the diameter of the telescope. Since $r_0 \ll D$ at most observatories, not correcting for the wavefront aberrations will cause the telescope to have the same angular resolution as a telescope with a diameter equal to r_0 . Moreover, the observing conditions at many of the world’s foremost observatories that are used for direct imaging lead to halos of up to $\sim 1''\text{--}2''$. Since many of the exoplanets targeted by direct imaging surveys are at separations less than $1''$ this speckle halo can completely overpower the signal from an exoplanet, rendering it undetectable in the final image.

To improve from the atmospheric seeing and approach the diffraction limit, Adaptive

²1 AU \approx the distance between the Earth and Sun \approx 150 million km \approx 93 million miles.

Optics (AO) instruments are used as control systems to correct for wavefront aberrations and concentrate the light back into the diffraction limited PSF. Modern AO systems can typically get up to 90% of the diffracted light back into the core of the diffraction limited PSF (Currie et al., 2019). The success of the AO system correcting for these aberrations is measured by the Strehl Ratio (SR), which is the ratio between the intensity of the measured PSF and the maximum possible intensity of the PSF in a perfect optical system. An SR of 1 (or 100%) means that the AO system has perfectly corrected the aberrations in the image.

In general, an AO system like the one in Figure 1.4 consists of several components. There will be a wavefront sensor (WFS) camera to measure the aberrations in the wavefront, a deformable mirror (DM) to apply a correction to the wavefront, and a control system. The control system is typically a fast computer that takes in data from the WFS, crunches it, calculates an appropriate correction, and applies the correction the DM. The DM is usually a thin sheet of glass attached to hundreds to thousands of actuators that push and pull on the glass to change its shape in response to the error measured by the WFS. There may also be a beamsplitter that sends some light to the WFS and allows the rest to pass through to the science camera. The WFS, control system, and DM are connected in a feedback loop so as the correction improves and the wavefront aberrations are tamped down, the AO system will only measured the residual wavefront errors that weren't corrected in previous iterations of the loop.

In practice modern AO and extreme AO (xAO) systems are much more complicated but operate on the same principle. For example, some systems like the Magellan Extreme Adaptive Optics instrument (MagAO-X; Males et al., 2020) have two DMs in a woofer/tweeter configuration, where the “woofer” DM corrects aberrations at low spatial frequencies and the “tweeter” operates at higher spatial frequencies. They may also operate at different speeds depending on how quickly the aberrations evolve in time. One

may also use a Focal Plane Wavefront Sensing (FPWFS) architecture where the science camera doubles as the WFS so there is no difference in the optical path between the two. The MEC instrument has operated as a FPWFS behind SCEXAO at Subaru (Walter, 2019; Fruitwala, 2021). Most systems for high contrast imaging are extreme AO instruments (Jovanovic et al., 2015; Males et al., 2020; Ahn et al., 2021) which refers to the number of DM actuators; they also usually operate at faster speeds than AO systems.

Wavefront sensors come in many flavors, each with strengths and drawbacks. Here we will describe two styles that have been used with MKID instruments: Shack-Hartmann (with DARKNESS and PICTURE-C) and Pyramid WFSs (with MEC and XKID). A Shack-Hartmann WFS consists of a lenslet array placed in the pupil plane of an imaging system, leading to light being focused into an array of spots in the focal plane. The displacement of the centroid of each spot is proportional to the slope of the wavefront at the lenslet array. By integrating over all of the spots a map of the aberrated wavefront can be computed and used to apply a DM correction (Platt & Shack, 2001). A Pyramid WFS uses a pyramid-shaped optic where the telescope beam is focused onto the point of the pyramid (in the focal plane). This will split the light into 4 images of the pupil that are measured by a camera. Aberrations in the wavefront will cause the intensity pattern in each of the four images to change, then the pattern of relative intensities between the images can be used to reconstruct the aberrated wavefront (Esposito & Riccardi, 2001). Shack-Hartmann WFSs are more stable but can suffer from low sensitivity to certain aberrations while Pyramid WFSs tend to be less stable and harder to align, but have much greater sensitivity.

Speckles

After correcting for most of the wavefront errors, the residual errors that persist will manifest as a speckle background. A series of simulated speckle images is shown in




Figure 1.5 Simulated speckle fields in an astronomical image. Each row shows the realized speckle field at a different time, the first three columns show the field at 1000, 700, and 400 nm, while the final column shows the combined, broadband speckle field. Originally from Hall & Jefferies (2022).

Figure 1.5. These speckles may also masquerade as companions due to their relatively similar (or higher) brightnesses and characteristic size (the diffraction limit, $\sim \lambda/D$). There are two types of speckles, each characterized by their relative lifetimes. These are quasistatic speckles with $\tau \sim$ seconds to hours and atmospheric speckles with $\tau \sim$ tens of milliseconds, depending on local observing conditions.

Quasistatic speckles stem from mechanical changes and imperfections in the telescope and instrument that are not measured and corrected by the WFS. These are called non-common path errors (NCPE) and tend to vary over long timescales, from several seconds to hours. Often they are attributed to the telescope moving or flexure in the optical system/components. In Figure 1.4 these would occur between the beamsplitter and high resolution camera. In reality they can occur in more places than this, but it is essential to try to reduce the non-shared path length between the WFS and science camera. This is also why focal plane wavefront sensing is so useful, as one doesn't need to split the

light between WFS and science camera meaning that NCPEs are reduced.

Atmospheric speckles come from aberrations induced in the wavefront by the turbulent atmosphere above the telescope that are not corrected by the AO system. The typical lifetime of an atmospheric speckle is tens of milliseconds, though it depends on how quickly the atmosphere evolves in front of the telescope and other local weather conditions. The AO system may be unable to correct wavefront aberrations for several reasons such as photon noise in the WFS, spatial aliasing, inability for the system to keep up with the atmospheric turbulence, chromaticity between WFS and science camera, and the error in the wavefront being too large for the DM to correct (Guyon, 2005).

Both types of speckles prevent systems from achieving the deepest possible contrasts, but there are many tools and observing methods that aim to mitigate speckle noise.

1.2.5 Coronagraphy

The job of the AO systems is to improve the image quality by reducing the wavefront errors and moving light from the seeing halo back into the ideal diffraction pattern. Once this job is done there is often still too much starlight that will overwhelm any companions in the image. To suppress on-axis starlight and allow off-axis signals to pass through to the science camera, a coronagraph is typically used.

Most coronagraphs utilize a combination of masks and filters in various pupil and focal planes of the optical path. The earliest successful coronagraph and most common design in current use is called a Lyot coronagraph (Figure 1.6). It operates by placing an occulting spot – usually an optic with an opaque dot – in the first focal plane on-axis at the center of the Airy pattern. The light that is not blocked by the spot will be diffracted toward the edges of the image in the subsequent pupil plane. To block that light there will be a Lyot stop, a mask that is transparent in the center, opaque at the

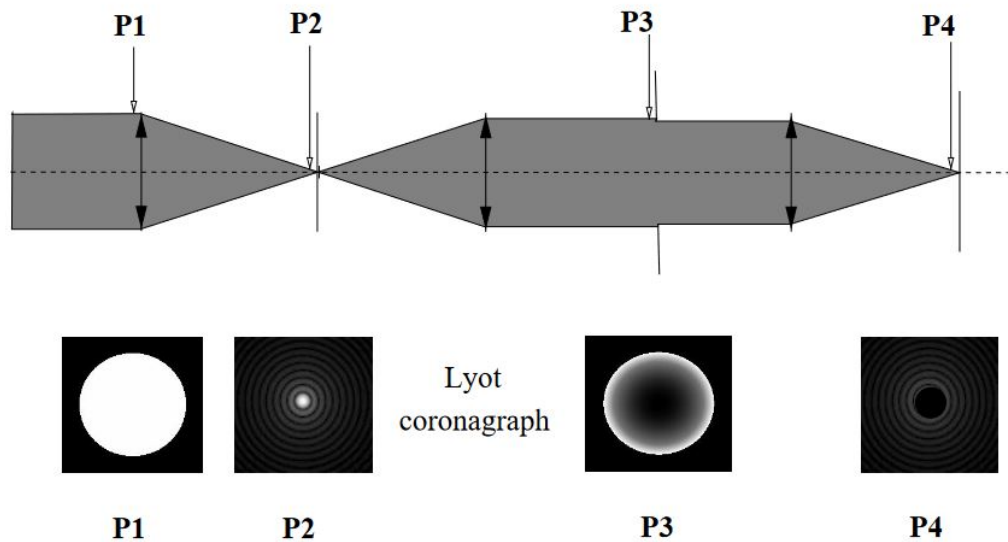


Figure 1.6 Diagram of a Lyot coronagraph. Light enters from the left, passing through the entrance pupil plane (P1) before a focal plane mask at P2 blocks the core of the PSF. At P3 the light passes through a Lyot stop where the diffracted ring at the edge of the pupil is blocked before being focused onto P4, the final focal plane, where the (ideal) final image is shown. Reproduced from Guyon et al. (1999).

edges, and sized slightly smaller than the beam of light itself to block the diffracted light. The remaining light is focused onto a detector in the focal plane. Ultimately, the central PSF can be attenuated to a small fraction of its original intensity while the off-axis light (containing light from companions of interest) remains unblocked (Guyon et al., 1999).

Lyot coronagraphs remain the most common in use today, but other novel styles such as Vector Vortex Coronagraphs (VVC; Mawet et al., 2009) and Phase-Induced Amplitude Apodization (PIAA) Coronagraphs (Pluzhnik et al., 2005) aim to achieve better performance by increasing the throughput of off-axis light, reducing diffraction caused by hard edges of the stops, and decreasing the inner working angles set by the size of the occulting spots.

1.2.6 Integral Field Spectrographs

An Integral Field Spectrograph (IFS³) is an instrument capable of taking astronomical images *and* providing spectral data at each point in those images. This results in output data cubes with one spectral and two spatial dimensions. These are useful tools for high contrast imaging because they capture the light from a companion and enable immediate spectral characterization without the need for extra follow-up. Additionally, because the position and behavior of speckles are wavelength dependent, an IFS may enable the user to mitigate the speckles more effectively in pre- or post-processing. At UCSB, five MKID (Section 1.3) IFSs have been developed. The first four are discussed briefly in section 1.4 and the development and deployment of the fifth is covered in Chapters 2, 3, and 4.

1.2.7 Observing Methods

The previous sections were concerned with the hardware and instrumentation tools that enable high contrast imaging. However it was seen that even using AO systems and coronagraphs does not provide a perfect image and there are still sources of noise that contaminate it and prevent the observer from identifying a companion. The following section will describe several common observing methods and data processing algorithms that are to reduce speckle noise in astronomical images to dig down to lower contrasts.

Reference Differential Imaging

Reference Differential Imaging (RDI; Ruane et al., 2019; Xie et al., 2022) is a technique where a reference star is observed in addition to the target star that may host a companion. The reference star is typically of a similar spectral type and brightness to the target star and is nearby in the sky. An image of the reference star will then be

³Sometimes also called an Integral Field Unit (IFU).

subtracted from the image of interest, where a companion will hopefully be revealed.

RDI promises to work at narrow separations where other differential imaging techniques suffer. It also works at wider separations and can be used in concert with other techniques to decrease the noise in the image further than either technique alone.

Spectral Differential Imaging

Spectral Differential Imaging (SDI; Marois et al., 2000) leverages the wavelength dependence of the speckles to detect faint sources in IFS data cubes. The characteristic size of speckles and the speckle halo are both proportional to the wavelength of light they are observed at. This means that as one looks at astronomical images with a speckle halo in narrow wavelength bands, as the wavelength increases, so too will the separation of the speckle halo from the central star. However, the position of companions and real astronomical objects is *not* wavelength dependent, which means they will not move.

SDI can be performed one of two ways. First, a series of images successively increasing in wavelength can be analyzed and if there are any sources that do not move, they are likely to be real sources. Second, one can magnify/shrink the subsequent images so that the speckle halo remains the same size while everything else shrinks or grows. Taking the median of all images will provide a reference image of the speckle halo that can then be subtracted from each resized image, removing the speckle light but leaving some companion light in each. The resized images can then be returned to their original size and combined, revealing the companion while removing the speckle halo.

SDI can suffer from self-subtraction where light from real objects is removed because it was accidentally included in the reference image that was subtracted from each frame. It is also only able to be used when the astronomical images also have spectral data, so conventional cameras without spectral resolving power cannot perform SDI.

Angular Differential Imaging

Angular Differential Imaging (ADI; Marois et al., 2006) is another technique that aims to create a reference image from the observing dataset to remove noise from the final product. In contrast to SDI which used the chromaticity of the speckles, ADI uses the rotation of the sky. As a telescope tracks objects through the sky, the sky will rotate due to the motion of the Earth. An alt-az⁴ telescope will typically have a component called a “derotator” which rotates the telescope at the same rate as the sky so objects stay in the same place on an astronomical detector.

With the derotator turned off the sky rotates through the image, while the speckle pattern caused by imperfections in the telescope optics stays static and does not rotate. To perform ADI a sequence of many images is taken with the derotator off, allowing the sky to rotate between each image. Since the speckle pattern here is dominated by the quasistatics from the telescope, a median reference image can be generated from all of the frames in the sequence before being subtracted from each of these frames, leaving only light from the sky. Then the images of the residuals from this subtraction are derotated to counteract the rotation of the sky and recombined into a final image. By derotating and combining the image, any source that was rotating with the sky would not have been caught in the reference image and will now show up in the final image stack.

ADI works poorly when there is little sky rotation or when a companion is at a narrow separation from the star. Both cases will result in self-subtraction, where the companion light is caught in the reference image and so is removed during that stage.

⁴Altitude-Azimuth telescope, where the two pointing axes track the height of the target and the telescope’s direction along the horizon.

Stochastic Speckle Discrimination

Stochastic Speckle Discrimination (SSD; Walter et al., 2019; Steiger et al., 2022b) is a novel technique that uses cameras with fast exposure times or photon counting detectors to measure the difference in photon arrival time statistics between a planet/extended source and an off-axis speckle.

The intensity distribution governing the statistics of coherent light with a speckle pattern was originally developed in Goodman (1975) for laser light. Cagigal & Canales (2001) verified experimentally that the probability of seeing an instantaneous intensity I in an image with a speckle pattern is described by a Modified Rician (MR) distribution:

$$\rho_{\text{MR}}(I) = \frac{1}{I_S} \exp\left(-\frac{I + I_C}{I_S}\right) I_0\left(\frac{2\sqrt{II_C}}{I_S}\right) \quad (1.1)$$

where I_0 is the zero-order modified Bessel function of the first kind, I_C is the coherent component of the light (i.e. the stellar PSF without the presence of the atmosphere) and I_S is the variable component that describes the intensity of the speckle halo.

While all points in an astronomical image will have a total intensity I made up of components I_C and I_S , the ratio of I_C/I_S will vary depending on whether the source is predominantly coherent light (a real source) or predominantly time-variable, stochastic light generated from diffraction in the atmosphere and telescope optics (a speckle).

When a real astronomical object is present in an astronomical image it will have a stronger I_C component than I_S components. In comparison, when there is no astronomical object at a location the I_C component will be weaker relative to the I_S component. This ratio affects the shape of the MR distribution, with a high value of I_C/I_S leading to a distribution with less skew and a lower I_C/I_S component leading to a distribution with a more positively skewed distribution (Figure 1.7). By creating a map of I_C/I_S at each point in an image one may determine if sources of similar intensities are more likely

to be companions (high I_C/I_S values) or speckles (lower I_C/I_S values).

This technique has been performed on sky and has led to the direct imaging discovery of a previously unknown stellar companion (HIP 109427 B; Steiger et al., 2021) in addition to directly imaging an extended disk around a nearby star without using typical polarization imaging techniques (AB Aurigae; Steiger et al., 2022b).

Since atmospheric speckles tend to evolve over tens of milliseconds and SSD requires that the intensity distribution is sampled faster than the speckle lifetime (i.e. $t_{\text{exposure}} \lesssim \tau$) it is necessary to use an astronomical detector that can take images very quickly and maintain a reasonable signal-to-noise value. One such technology that enables these types of fast cameras is MKIDs, which will be introduced in the following section and discussed through the rest of this thesis for their utility in astronomical instruments.

1.3 Microwave Kinetic Inductance Detectors

1.3.1 Operation Principle

Microwave Kinetic Inductance Detectors (MKIDs) are superconducting resonators capable of measuring single photon events with microsecond timing resolution and resolving the energy of the incident photon to within a few percent.

The charge carriers in superconductors, Cooper Pairs, are bosons formed by two interacting electrons. When cooled below its critical temperature ($T < T_C$) the Cooper pairs in a superconductor flow as a superfluid with zero DC resistance. However, in an alternating electric field the Cooper Pairs can have a high reactance. This occurs because the electrons in the Cooper Pairs have inertia when they move, meaning they will oppose the electromotive force from the changing electric field and take a finite time to accelerate in the opposite direction. The resulting lag in voltage can be modelled as

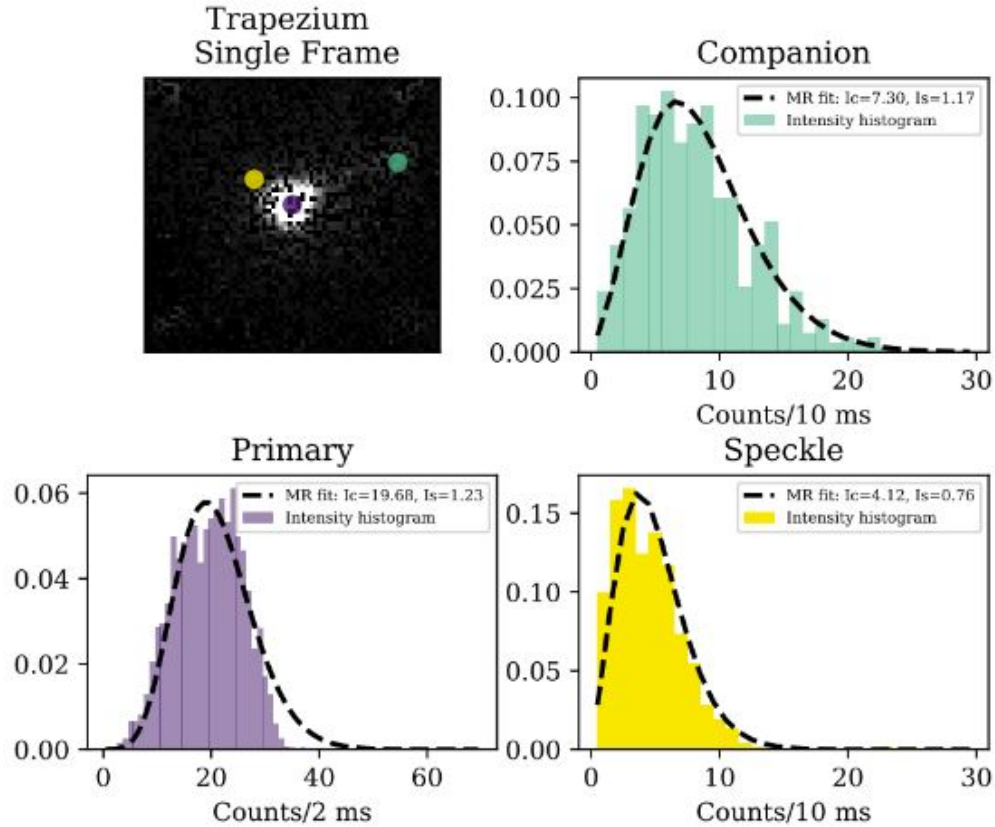


Figure 1.7 Reproduced from Steiger et al. (2022b). (Top left) Shows an astronomical image of Θ 1 Orionis (centered on the purple circle) along with its faint companion (below the green dot) and a speckle pattern (with one speckle highlighted in yellow). (Bottom Left) The MR probability density function measured at the central star. One can see that it has a high I_C/I_S ratio and is nearly symmetrical. (Top Right) The MR PDF measured at the companion, while the total intensity is lower, the I_C/I_S ratio remains relatively high and is reflected by a slightly positively skewed distribution. (Bottom Right) The MR PDF measured on top of a speckle of similar intensity to the companion itself. The I_C/I_S ratio is much lower than in the panel above which is reflected in the much greater positive skew of the distribution.

an inductance term - called kinetic inductance - in the resonator circuit in addition to the normal geometric inductance of the superconductor. This variable inductance term is inversely proportional to the number of Cooper Pairs in the superconductor (or directly proportional to the number of quasiparticles created when Cooper Pairs are broken). An MKID is a superconducting resonant circuit that can measure the changing kinetic inductance when a photon is absorbed in the superconducting material.

For a photon to be absorbed by a superconductor, its energy ($h\nu$) must exceed twice the superconducting gap energy (2Δ), which is a function of the superconductor's critical temperature T_C ,

$$2\Delta = 3.52k_B T_C, \quad (1.2)$$

where k_B is the Boltzmann constant. When a photon with sufficient energy is absorbed by a strip of superconducting material, many Cooper Pairs will be broken into quasiparticles and cause the kinetic inductance to increase. The number of quasiparticles created by an absorbed photon is

$$N_{qp} = \frac{\eta h\nu}{\Delta} < \frac{h\nu}{\Delta} \quad (1.3)$$

where η is the efficiency of quasiparticle creation. This value will always be below 1 and vary from photon event to photon event because the creation of quasiparticles is a statistical process and energy may be lost to the substrate beneath the superconductor.

To monitor and measure the change in inductance when a photon strikes an MKID resonator, the resonator is excited using a probe tone at the resonant frequency of the circuit. Figure 1.8 demonstrates how an MKID detects an incident photon. Panel **(b)** shows the resonator modeled as a lumped-element tank circuit with a variable inductor that is capacitatively coupled to a microwave feedline. The variable inductor shows that the inductance varies when the MKID is struck by a photon. It also shows a photon of energy $h\nu$ striking the inductor. Panel **(a)** shows the same photon breaking Cooper Pairs,

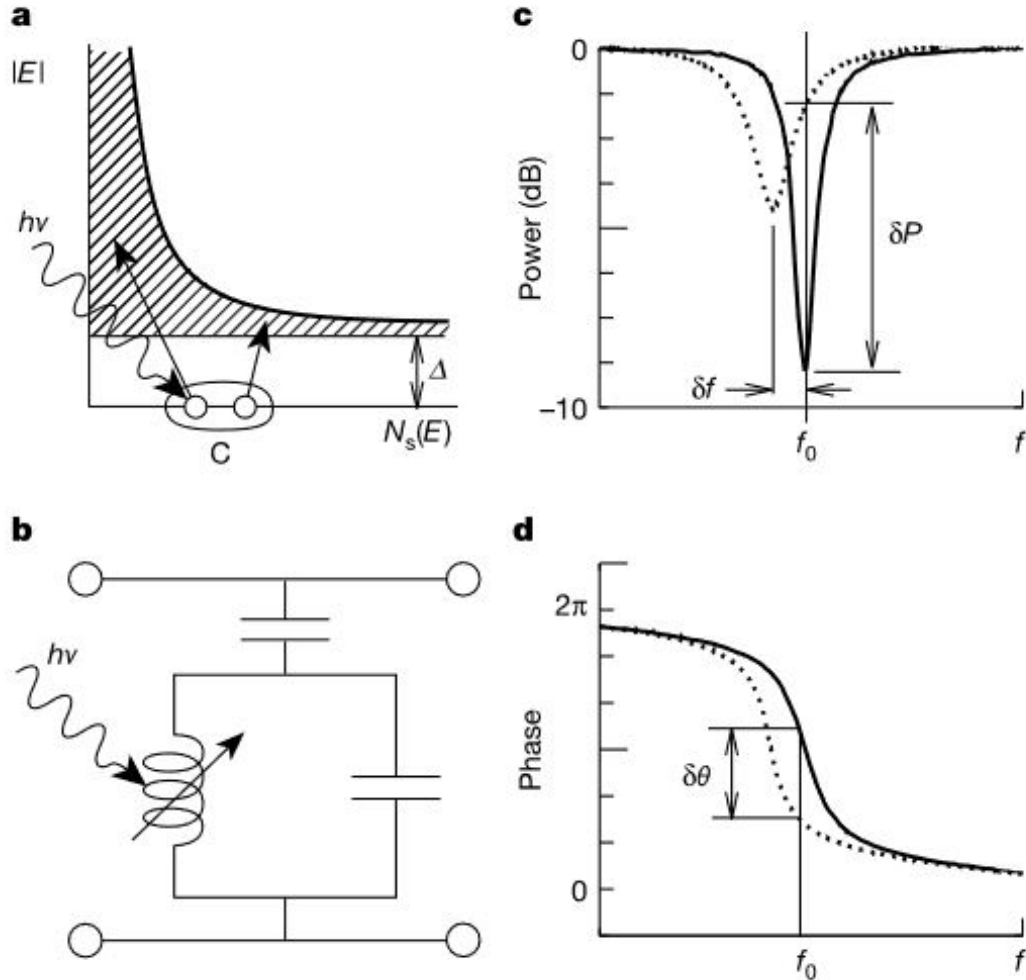


Figure 1.8 Reproduced from Day et al. (2003). The panels in this figure describe the MKID operation and photon detection principle. Panel (a) shows that a photon with energy $h\nu \geq 2\Delta$ may be absorbed by a thin superconducting film, breaking Cooper pairs (C) into a number of quasiparticle excitations, N_{qp} . Panel (b) shows a single MKID resonator modeled as a lumped-element LC circuit with a variable inductance and capacitively coupled to a through line. When the photon is absorbed, the surface inductance (L_S) and surface resistance (R_S) increase. The changed surface inductance increases the total inductance (L) while the increase in surface resistance will make the resonator slightly lossy. Ultimately this manifests as a temporary change in the resonator's resonant frequency (f_0) along with the amplitude (c) and phase (d) of a microwave probe tone that is being sent through the circuit.

creating N_{qp} quasiparticles. Panel **(c)** shows that a probe tone sent down the feedline at the MKID's resonant frequency (f_0), there is a transmission dip at f_0 . The quasiparticles created by the incident photon cause the inductance to increase, temporarily moving the resonance to a lower frequency. Panels **(c)** and **(d)** show the amplitude and phase of the probe tone, respectively, varying with the changing inductance.

After the initial deposition of energy in the MKID breaks many Cooper Pairs into quasiparticles, the phase and amplitude change rapidly over the course of a few microseconds. The inductance - and therefore the phase and amplitude - returns to its quiescent value as the quasiparticles recombine into Cooper Pairs, which typically occurs at $\sim 100\mu\text{s}$. This will manifest as a phase 'pulse' (Figure 1.9) that has a characteristic shape, a sharp increase followed by an exponential tail as the resonator returns to the unexcited state. The phase pulse also allows the determination of the incident photon's energy. This is because the number of Cooper Pairs broken by an incident photon is proportional to that photon's energy. Since the number of Cooper Pairs broken is also proportional to the change in frequency and the phase of the probe tone, the height of the resulting phase pulse will also be proportional to the energy of the incident photon.

A typical MKID resonator is read out using room-temperature readout electronics that sample the phase of the probe tone at MHz rates, which means that when a photon event occurs in an MKID it can be read out with microsecond timing resolution.

Figure 1.10 shows the transmission of several MKID resonators through a single microwave feedline. Since the resonators are designed to have high quality factors (Q) they have near-perfect transmission away from resonance. This thousands of MKIDs to be read out on a single feedline using a scheme called frequency domain multiplexing (FDM). Typical resonant frequencies for MKID resonators are in the microwave range from 4 to 8 GHz. The resonant frequency of a given MKID is controlled by its capacitance (Figure 1.8, panel **(b)**). In practice, an MKID is fabricated using a combination of




Figure 1.9 A typical timestream from an MKID resonator showing the measured phase of the probe tone through the resonator as a photon is absorbed. The quiescent phase sits at ~ 0 degrees until a photon strikes the photosensitive inductor. As Cooper Pairs are broken into quasiparticles, the impedance changes and causes a spike in the measured phase. The quasiparticles then start recombining into Cooper Pairs and the phase returns to its normal value. The peak height of this pulse is proportional to the energy of the absorbed photon. Reproduced from Swimmer et al. (2023).

photolithography and etching to create a thin-film resonant circuit that consists of a meandered inductor where photons are absorbed and an interdigitated capacitor to set the resonant frequency. A close-up view of a single MKID can be seen in the right panel of Figure 1.11. The red square at the top shows the inductor and the blue square below shows the interdigitated capacitor. Just beneath the interdigitated capacitor there is also a coupling capacitor which couples the resonator to a microwave feedline, allowing the resonator to be excited by a probe tone sent down the feedline and read out at the output. The left panel of Figure 1.11 shows many MKIDs coupled to a microwave feedline. Each MKID on a given feedline will have a different capacitance to ensure that it has a unique resonant frequency so it can be read out without colliding with another in frequency space and rendering both unusable.




Figure 1.10 Reproduced from Mazin et al. (2012). This figure shows an illustration of the frequency domain multiplexing scheme that allows many MKID resonators may be read out on a single microwave feedline. (Top) A circuit diagram showing many MKID resonators capacitively coupled to a single feedline. In practice, the probe tones that excited the resonators are sent into port 1 and read out at port 2. (Bottom) Transmission data of many MKIDs on a single feedline. Their high quality factors (Q) mean that each resonator has a deep, narrow transmission dip at its resonant frequency and near perfect transmission away from resonance, letting resonators with unique resonant frequencies to be read out on the same microwave feedline.

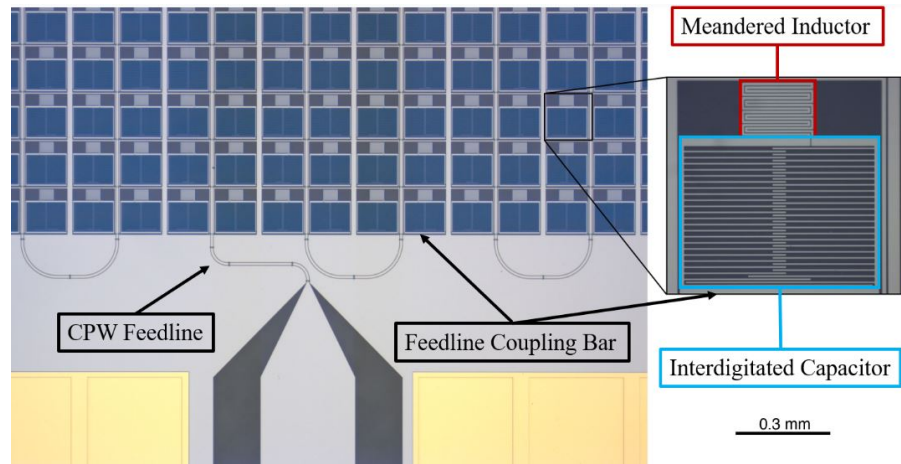


Figure 1.11 (Left) A zoomed in image of a large-format (20,000 pixel) MKID array. A coplanar waveguide (CPW) is seen in the bottom center, leading to the CPW feedline that meanders between columns of MKID resonators. Gold bond pads can be seen at the bottom left and right, which are used to thermalize the MKID array. (Right) A single MKID resonator. The red square encloses the photosensitive inductor and the blue square encloses the interdigitated capacitor that is used to control the resonant frequency of each MKID. The feedline coupling bar that is used to couple each MKID to the feedline is also called out in both panels. Reproduced from Walter et al. (2020).




Figure 1.12 Two large-format MKID arrays. (Left) An array consisting of 5 feedlines and 10,000 MKIDs. This style of array has been used in the DARKNESS and XKID instruments. (Right) A 10 feedline, 20,440 MKID array installed in its readout box. This style of array is used in the MEC instrument. The array itself is the colorful square in the middle and measures $\sim 1'' \times 1''$. The box that it is mounted in measures $\sim 2'' \times 2''$. A zoomed in image of the array in the left panel is seen in Figure 1.11.

1.3.2 Large-Format MKID Arrays

The ability to utilize a frequency domain multiplexing scheme to read out many MKIDs from a single microwave feedline enables the design and fabrication of large-format arrays. Each feedline of an MKID array can read out $\approx 2,000$ detectors with unique resonant frequencies between 4-8 GHz. To create a full MKID array the same feedline design will be repeated side-by-side 5 or 10 times, resulting in a rectangular or square array made up of 10,000 or 20,440 pixels, respectively. By patterning the same feedline over and over, the complexity of developing and fabricating a full array is greatly reduced. Two examples of large-format MKID arrays are shown in Figure 1.12.

The reason for creating these large-format arrays is that when MKID detectors are used to detect photons they will be able to measure when the photon struck a detector, the photon's energy, and *where* the photon was absorbed on the array. Ultimately this means MKID instruments can be used as IFSs, measuring spectral data at each point of the images they generate.

1.3.3 MKID Photon Measurement

Each MKID resonator in an array – when used as an IFS – serves as a pixel. Throughout the rest of this thesis “detector”, “resonator”, and “pixel” will be used interchangeably, with “MKID” also sometimes referring to a single MKID resonator as well. This section will briefly expand on the theoretical description from Section 1.3.1 of what a single-photon event looks like when measured by an MKID and how it is read out, which will be addressed further in Section 2.7.2 and Chapter 7.

MKID cameras utilize a room temperature readout system to measure and record photons. The current system to read out an MKID array is called the 2nd Generation MKID Readout (Gen2; Fruitwala et al., 2020; Fruitwala, 2021), which uses hardware de-

veloped by the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER; Hickish et al., 2016) team to control field programmable gate arrays (FPGAs) that generate the signals sent into the MKIDs and process the output.

For simplicity we will follow the signal through a single MKID, understanding that in a full array this process done for thousands of resonators over several microwave feedlines in parallel. To begin, a microwave probe tone is sent through the MKID, exciting it and measuring a dip in the measured transmission. After passing through the MKID the phase probe tone is measured on the output side. The quiescent phase of the probe tone is measured to find a baseline phase value $\hat{\phi}$ and the phase noise σ_ϕ (the “jitter” on the signal) is also calculated. The phase noise is then used to create a threshold that serves as the lowest phase value a photon can be detected at. Typically the Gen2 readout system will be set so that a photon cannot be measured unless it causes the phase to deviate greater than $4\sigma_\phi$ or $5\sigma_\phi$ from $\hat{\phi}$.

In addition to the threshold the phase pulse must meet certain criteria. These criteria are that the phase must first increase in magnitude for two consecutive measurements (a sharp spike) before decreasing over seven subsequent measurements (a longer decay time back to its normal value). For a given MKID, a template can also be created by measuring many photons and taking their average which enables the construction of an “optimal filter” that the readout can use to convolve with the measured phase data to improve the readout’s ability to recognize photons.

Figure 1.9 showed a single photon event measured by an MKID. This data has not used a filter and was part of a dataset that used a $6\sigma_\phi$ detection threshold. There is a rapid rise time where the magnitude of the measured phase increases drastically before exponentially decaying back to its normal value. The average phase value is subtracted from all data points so that $\hat{\phi} = 0$ radians. A filter using the templates of many photons would serve to smooth out this data and make it less noisy.




Figure 1.13 Reproduced from Zobrist (2021). This shows a single resonator using a low-noise readout scheme to measure the MKID’s resolving power between 254 and 1310 nm. The histograms show the distribution of phase pulse heights for each of the lasers used. The central peak of each corresponds to the wavelength λ and the width to $\Delta\lambda$. The wider the distribution, the lower the \mathcal{R} .

Finally, after the MKID readout has detected a photon the time that it struck at is recorded (using a UNIX timestamp) along with the pulse height and an identifier telling the user which MKID was struck by that photon which will enable the determination of its position on the array. Since the readout is being sampled at 1 MHz, this means that each photon’s arrival time will be measured to microsecond timing resolution. Ultimately this results in a time-tagged list of photons where each photon in the list contains a time, phase, and position value. These values can then be used to generate images, movies, spectra (from the fact that the phase measured is proportional to the energy of the photon), and more data products from large-format MKID arrays.

1.3.4 Energy Resolution

As discussed in the previous section, the pulse height from an incident photon is proportional to the energy of the incident photon. However several different effects can

cause the MKID’s phase response to vary slightly to photons of the same energy. This may be due to noise in the phase measurement, the deposition of energy into the MKID being a nondeterministic process (i.e. photons of equivalent energy do not always break the same number of Cooper pairs), energy loss into the array substrate, and two level system (TLS) noise in the detectors.

This means that to determine the resolution of an MKID a wavelength calibration or “wavecal” is performed. During a wavecal, all of the detectors are illuminated by monochromatic lasers of known wavelengths. A given detector’s response to many photons from each of the different lasers (shined one at a time) is then measured. The distribution of these pulse heights is then fit for each laser, which enables the resolving power to be calculated using the formula

$$\mathcal{R} = \frac{\lambda}{\Delta\lambda} = \frac{E}{\Delta E} \quad (1.4)$$

where λ is the central peak of the distribution and $\Delta\lambda$ is the Full Width at Half Maximum (FWHM), which describes the spread in measured values. A larger spread in values will correspond to a lower \mathcal{R} , meaning the pixel has less ability to distinguish photons of that wavelength from those of another. Figure 1.13 shows a wavecal dataset for a resonator with $\mathcal{R} \sim 20$ at 800 nm.

In MKID instruments typical values of \mathcal{R} range from 3-7, while in-lab measurements can yield $\mathcal{R} \sim 6-8$. Using traveling wave parametric amplifiers (TWPAs; Eom et al., 2012) and low-noise readout systems measurements of \mathcal{R} as high as 35 have been seen in recent experiments (Zobrist et al., 2021).

1.4 Previous MKID Instruments

Prior to the deployment of the XKID instrument (Chapters 2, 3, 4) there have been 3 MKID instruments successfully commissioned (ARCONS, DARKNESS, and MEC) and 1 cancelled MKID mission (The PICTURE-C MKID Camera). This section is intended to be a survey of the previous MKID projects and to provide context of the history of MKID instruments as UVOIR astronomical cameras.

1.4.1 ARCONS

ARCONS, the **A**rray **C**amera for **O**ptical to **N**ear-**I**R **S**pectrophotometry, was an MKID IFU that was successfully deployed in 2011 at the Coudé focus of the Hale 200-inch telescope at Palomar Observatory. It used a 44x46 pixel array (2024 pixels total) and had a bandpass of 400-1100 nm. To operate the MKIDs at 100 mK, the ARCONS instrument used a cryostat with a base temperature of ~ 3 K and operated the MKIDs at a temperature of 100 mK using an adiabatic demagnetization refrigerator (ADR). ADR operation is discussed in depth in Section 2.3.2.

The ARCONS instrument was the first optical-to-near-IR MKID instrument used as an astronomical camera. It had a 10" x 10" field of view and primarily targeted pulsars and x-ray binary systems, with an added goal of measuring the redshifts of galaxies using their spectral shape and Lyman- α emission (Mazin et al., 2012). Figure 1.14 shows the ARCONS instrument mounted at the Coudé focus at Palomar.

1.4.2 DARKNESS

The next MKID camera to be commissioned was DARKNESS, the **D**ark-speckle **N**ear-Infrared **E**nergy-resolving **S**uperconducting **S**pectrophotometer. The initial deployment came in July 2016, also on the Hale 200-inch telescope at Palomar Observatory.



Figure 1.14 The ARCONS cryostat mounted at the Palomar 200" telescope Coudé focus. Also shown are several of the foreoptics (Pickoff M4, Off Axis Paraboloid, Reflective Slit) that enable the light to be directed into the instrument. Accessed via <http://web.physics.ucsb.edu/~bmazin/projects/arcons.html>.

The DARKNESS instrument was the subject of Seth Meeker's thesis (Meeker, 2017).

DARKNESS used an 80x125 (10,000 pixel) MKID array with a $\sim 1.6'' \times 2.5''$ field of view and operated between 800-1400 nm, in the near infrared (NIR). In contrast to ARCONS, DARKNESS sat behind Palomar's PALM-3000 extreme AO system (Dekany et al., 2013) and Stellar Double Coronagraph (SDC; Bottom et al., 2016). Its goal was to perform high contrast direct imaging observations utilizing real-time speckle control (Fruitwala et al., 2018) and post-processing speckle suppression at high frame rates. DARKNESS can be seen installed on the Hale 200" Telescope in Figure 1.15. To operate the MKIDs at cryogenic temperatures, the DARKNESS cryostat was cooled using liquid nitrogen (LN₂) and liquid helium (LHe) to drop to 4.2 K, then used an ADR to further drop the MKIDs to 90 mK (Meeker et al., 2018).

DARKNESS was successfully deployed at Palomar Observatory several more times in November 2016, April 2017, and October 2017 for science verification. This marked the first use of a 10,000 pixel MKID instrument on sky. It produced diffraction limited images and removed contaminating speckle noise from the on-sky data (Meeker et al.,




Figure 1.15 The DARKNESS instrument (right) behind the Stellar Double Coronagraph (SDC, left) both attached to the P3K bench. The configuration is taken when installed in the Cassegrain cage at the Hale 200” telescope at Palomar.

2018). An image of a resolved binary from DARKNESS is shown in Figure 1.16.

1.4.3 MEC

Following DARKNESS, the next commissioned MKID instrument was MEC, the **MKID Exoplanet Camera**. MEC uses a 140x146 (20,440 pixel) MKID array, operates in the same bandpass as DARKNESS (800-1400 nm), and was designed for use behind the Subaru Coronagraphic Extreme Adaptive Optics instrument (SCEXAO; Jovanovic et al., 2015; Ahn et al., 2021) on the Subaru Telescope on Maunakea. The thesis (Walter, 2019) and instrument paper (Walter et al., 2020) by Alex Walter provide a comprehensive description of the MEC instrument, which saw first light in October 2018. Figure 1.17 shows MEC at the Naysmyth platform at Subaru. MEC, like ARCONS, uses a pulse tube to cool the cryostat to ~ 3 K before using an ADR to operate the MKIDs at 90 mK.




Figure 1.16 Reproduced from Meeker et al. (2018). A J-band image of 10 Uma, a spectroscopic binary system. The large panel shows the binary with a coronagraph focal plane mask blocking light from the primary. The inset shows the system without a coronagraph installed, which reveals the primary. The separation of the stars is $0''.42$.




Figure 1.17 Reproduced from Walter et al. (2020). The MEC instrument (center) sitting behind SCEXAO and the AO188 AO instruments (left) on the Naysmyth platform at Subaru on Maunakea. The panel of the MEC electronics rack is also seen to the right.




Figure 1.18 Reproduced from Steiger et al. (2021). The first stellar companion directly imaged using MEC, HIP 109427 B. (Left) A total intensity image of the HIP 109427 system. The central star is blocked by a coronagraph, the companion is circled in red, and a “speckle halo” of diffracted light can be seen around the coronagraph. (Right) The HIP 109427 system reduced using the SSD noise reduction algorithm. The companion is clearly picked out, while the noise from the speckle halo is almost completely removed.

Unlike ARCONS and DARKNESS which were visitor instruments at Palomar Observatory, MEC has been permanently installed at the Subaru Telescope since its first light and is currently available for community use. This makes it the first permanently deployed near-IR MKID instrument. It operates as an IFU and can act as a FPWFS in a feedback loop with SCExAO (Fruitwala et al., 2018; Fruitwala, 2021). As an IFU, the photon-counting ability with microsecond timing resolution also enables Stochastic Speckle Discrimination (SSD; Walter et al., 2019; Steiger et al., 2022b, Section 1.2.7).

MEC has been the MKID camera responsible for many MKID instrument “firsts”. It enabled the first direct imaging discovery of a stellar companion using SSD, HIP 109427 B (Figure 1.18; Steiger et al., 2021), the first discovery and characterization of a potentially substellar companion, HIP 5319 B (Figure 1.19, Chapter 5; Swimmer et al., 2022a), and performed the first on-sky demonstration of SSD on an extended object, the circumstellar disk around the star AB Aurigae (Steiger et al., 2022b), to name a few.

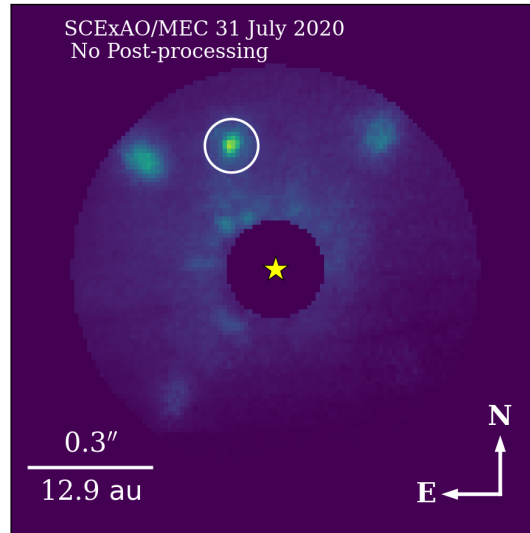


Figure 1.19 Originally published in Figure 1 from Swimmer et al. (2022a). The second stellar companion (HIP 5319 B) imaged using MEC, and the first of potentially substellar mass. As in Figure 1.18, the central star is obscured by a coronagraph while the companion is circled in white. No post-processing noise reduction was required in this image due to the favorable contrast and signal-to-noise ratios.

1.4.4 The PICTURE-C MKID Camera

The Planetary Imaging Concept Testbed Using a Recoverable Experiment - Coronagraph (PICTURE-C; Cook et al., 2015; Mendillo et al., 2022) is an experimental test bed whose astronomical goal is to image debris disks in a narrow bandpass (540-660 nm) around nearby stars ($d \lesssim 10$ pc). The mission was a collaboration between University of Massachusetts - Lowell, NASA, and UC Santa Barbara. PICTURE-C is a NASA balloon-borne telescope mission that originally consisted of a 0.6-m primary mirror, a modest AO system, a Vector Vortex Coronagraph (VVC; Mawet et al., 2009), and an MKID camera during its second flight. The first flight (September 2019) was a calibration and engineering test to evaluate system performance using a CCD (Charge Coupled Device) camera and the second flight (originally scheduled for September 2021) was to collect science data using an MKID instrument. Each flight would fly from the Columbia



Figure 1.20 The PICTURE-C gondola being tested at CSBF in Fort Sumner, NM prior to its first flight in September 2019. The gondola is the large white structure holding the telescope that is covered by Mylar to reflect stray light at night and solar radiation during the day. (Left) A lift test being performed, hoisting the gondola up with the balloon mounting harness beneath an overhead crane. (Right) A pointing test with the PICTURE-C gondola where the telescope travels through its full range of motion.

Scientific Balloon Facility (CSBF) in Fort Sumner, New Mexico.

For each flight, a short duration helium balloon would lift PICTURE-C to an altitude >35 km, above $\sim 99\%$ of the atmosphere for a single night of observations. At the end of the night the gondola, the support structure beneath the balloon that carries the telescope and other instrumentation (Figure 1.20), is designed to be detached from the balloon and fall to earth using a parachute and crash pads. From there it can be recovered on the same day and brought back to Fort Sumner.

The first flight flew in late September 2019. Upon recovery of the telescope after the flight, the telemetry was analyzed where it was found that the gondola was suffering from extensive vibrations in addition to the primary mirror requiring repolishing and replacement. This necessitated extensive work and further development on the system that was expected to push the integration of an MKID camera to a proposed 3rd flight. However, in early 2020 the COVID-19 pandemic significantly slowed operations and work on all phases of the project, where further delays related to the pandemic and funding resulted

in the MKID camera being struck from the project for the foreseeable future. In September 2022, roughly a year delayed from its original scheduled flight date, PICTURE-C had its second flight. Analysis of the flight data is ongoing and initial reports indicate that it was successful, leaving the door hopefully open for a future collaboration to integrate the existing PICTURE-C MKID camera with the now-working rest of the telescope system.

The PICTURE-C MKID Camera itself is a near-clone of the DARKNESS instrument with several key differences. Whereas DARKNESS was only able to be cooled using liquid cryogenics (LN₂ and LHe), the PICTURE-C MKID Camera was designed to be able to be cooled to a base temperature of ~ 4 K using either a pulse tube or liquid cryogenics, depending on the instrument configuration. As with the previous instruments, the MKID array was cooled from 4K down to an operational temperature of 100 mK using an ADR. The two possible configurations (Figure 1.21) are the “lab” configuration and the “balloon” configuration. The “balloon” configuration, when the cryostat is mounted on the gondola, uses liquid cryogenics that cool the fridge much more quickly to 4K. This requires more maintenance to periodically refill the cryostat and presents additional concerns such as the formation of ice plugs that can damage the cryostat. Liquid helium is also a non-renewable, scarce resource that makes it relatively expensive to acquire and use. The “lab” configuration is used when the cryostat is being tested at UCSB. It uses a pulse tube for the initial cooling of the fridge rather than liquid cryogenics. It takes the pulse tube more time to cool to 4K but is logistically much simpler, requires less maintenance, and is significantly less expensive. It also essentially negates the possibility of ice plugs forming, meaning that will not be a concern during lab tests.

The initial design and development are covered robustly in Clint Bockstiegel’s thesis (Bockstiegel, 2019). The continued development, along with the eventual integration and deployment of the PICTURE-C MKID camera was intended to be the focus of this thesis. However, due to the aforementioned delays stemming from instrument development and




Figure 1.21 (Left) PICTURE-C MKID Camera in its lab configuration. (Center) CAD design view of the PICTURE-C cryostat in the lab configuration with the 300K shell removed to see the internal structure of compressor shells. (Right) CAD design view of the cryostat in the balloon configuration, flipped upright with the pulse tube removed. Images originally from Bockstiegel (2019).

the COVID-19 pandemic, the PICTURE-C MKID Camera project was put on hold indefinitely. Several components from the PICTURE-C MKID Camera found their way into XKID (Chapter 2), where they will be discussed in greater detail.

Chapter 2

XKID: Design

2.1 Overview

This chapter will focus on the development of the MagAO-**X** **MKID** (XKID) instrument, the newest MKID IFS. XKID uses a modified version of the ARCONS (Mazin et al., 2012) cryostat to operate a 5-feedline, 10,000 pixel MKID array. It uses PtSi MKIDs designed to detect 800-1400 nm light (Szypryt et al., 2017) and has previously been measured with a resolution of $\mathcal{R} \sim 6 - 8$ (Meeker et al., 2018). The pixel pitch is 150 μm , similar to DARKNESS and MEC. Table 2.1 shows the design and measured parameters of the XKID instrument.

XKID is designed to interface with the Magellan Extreme Adaptive Optics (MagAO-X) instrument at the Magellan Clay Telescope at the Las Campanas Observatory (LCO) near Vallenar, Chile at the southern end of the Atacama Desert. This represents the first foray of an MKID camera for direct imaging in the Southern Hemisphere, a boon for exoplanet detection due to the increased number of young star-forming clusters visible at southern latitudes. A younger population of stars should in turn host younger exoplanets that tend to be at more favorable contrasts for direct imaging.

Table 2.1. XKID Instrument Design Summary

| Parameter | Values |
|-----------------------------------|---------------------------|
| Device Format | 80x125 |
| Number of Microwave Feedlines | 5 |
| Pixel Pitch | 150 μm |
| Plate Scale | 21.15 mas/pixel |
| Field of View | 1.77" x 2.77" |
| Bandpass | 800-1400 nm (Z, Y, J) |
| Resolving Power (\mathcal{R}) | 3-5 |
| Operating Temperature | 90 mK |
| Cryostat Hold Time | 12 to 24+ hours |
| Maximum Count Rate | 5000 counts/s |
| Pixel Dead Time | 10 μs |
| Readout System | MKID Gen2 Readout |

2.1.1 Magellan Telescopes

The Magellan Telescopes are a pair of 6.5-meter telescopes at the Las Campanas Observatory and were built on behalf of the Magellan Project, a collaboration between University of Arizona, the Carnegie Institution, University of Michigan, Harvard, and MIT (together called the Magellan Consortium). The twin telescopes, called the Walter Baade Telescope (“Baade”) and the Landon T. Clay Telescope (“Clay”) sit 60 meters apart from one another atop Cerro Manquí at an elevation of 2,516 meters (8,255 feet). Each of the telescopes have principal foci at two Nasmyth ports that each output an F/11 beam and one Cassegrain port that outputs an F/15 beam. Figure 2.1 shows the two telescopes at dusk prior to an observing night.

Walter Baade Telescope

Baade houses 4 facility class instruments. These are: the Inamori Magellan Areal Camera and Spectrograph (IMACS; Dressler et al., 2006), a wide-field imager and multi-object spectrograph; FourStar, A wide-field near-IR camera for use on the Nasmyth East



Figure 2.1 (Left) The Magellan Telescopes at dusk from the southeast. The Baade Telescope can be seen in the process of opening for the night, and to the right the Clay Telescope is rotating into position before opening. Between the two is the Auxiliary Building where the mirrors are polished and instrument maintenance is performed. (Right) Courtesy of Overbye (2023). A closer view of Baade Telescope at night.

(NASE) platform (Persson et al., 2013); the Folded port InfraRed Echellette (FIRE; Simcoe et al., 2013), A moderate resolution near-IR echellette spectrograph; and the Magellan Echellette (MagE) Spectrograph, a moderate resolution optical echellette spectrograph (Marshall et al., 2008). Each of these instruments operates nearly continuously as facility instruments, meaning – like MEC – they are permanently installed on the telescope and available for use so long as they are operational.

Landon T. Clay Telescope

The Clay Telescope uses both Nasmyth ports, Nasymth East (NASE) and West (NASW), and its Cassegrain port. The NASW platform is home to the Low Dispersion Survey Spectrograph (LDSS3-C¹), a high efficiency wide-field multislit spectrograph that is a facility instrument. The Cassegrain port can house MEGACAM (McLeod et al., 2015), a mosaic CCD instrument that has a 24'x24' field of view (FOV). The NASE port

¹<https://astro.uchicago.edu/research/ldss-3c.php>




Figure 2.2 Noah Swimmer standing beneath the Clay 6.5-meter telescope, preparing to ascend the ladder to NASE. From (Overbye, 2023). The primary mirror is held in the white cylinder (center) and the secondary mirror is suspended from the telescope spiders above (top). NASW is to the left and NASE is to the right, both just out of frame.

at Clay operates as a visitor port where “PI instruments” (non-permanently installed) are used. The PI instruments that can be used on NASE are:

- The Magellan Inamori Kyocera Echelle (MIKE; Bernstein et al., 2003) spectrograph, a high-throughput double echelle spectrograph.
- The Planet Finding Spectrograph (PFS; Crane et al., 2010), a high resolution optical echelle spectrograph for exoplanet detection.
- The Michigan/Magellan Fiber Spectrograph (M2FS; Mateo et al., 2012), a powerful wide-field, multiplexed spectroscopic survey instrument.
- The Parallel Imager for Southern Cosmology Observations (PISCO; Stalder et al., 2014) a simultaneous multi-band visible imager.
- The Magellan Adaptive Optics (MagAO; Morzinski et al., 2014) system, an AO instrument that uses an adaptive secondary mirror and control system before feeding light into VisAO (for visible-light observations) and Clio (for IR observations). This instrument is the predecessor of MagAO-X.
- The Magellan Extreme Adaptive Optics (MagAO-X; Close et al., 2018) instrument, an extreme AO system optimized for faster correction at shorter wavelengths. MagAO-X will be explored further in Section 2.2 and is the instrument that XKID is designed to accept light from.

Each of the PI instruments can be removed from the NASE port and allow others to be installed, enabling many disparate observing projects to take place at this port of the telescope. Figure 2.2 shows a partial view of the Clay Telescope with myself for scale standing beneath it, preparing to go up to the NASE platform.

2.2 MagAO-X

MagAO-X is the Magellan Extreme Adaptive Optics instrument, which is being developed at the University of Arizona (UA) by groups led by Laird Close and Jared Males. It uses a Woofer-tweeter architecture that is fed by the F/11 output from the static secondary mirror on the Clay telescope (Close et al., 2018). It uses 2040 actuators and can operate at up to 3630 Hz and promises to achieve high Strehl ratios ($\gtrsim 70\%$) at $H\alpha$ ($\lambda = 656$ nm). The main goal of MagAO-X is to perform high-contrast imaging of accreting protoplanets at $H\alpha$, with secondary goals of using an IFS to perform broadband high-contrast imaging. It is an extreme AO testbed designed to be shuttled between LCO during and observing runs and UA for continued local development.

The fact that MagAO-X promises high Strehl ratios at visible wavelengths is noteworthy since typical extreme AO systems work in the near IR. AO correction becomes significantly more challenging at shorter wavelengths because the number n of turbulent cells across the telescope will scale as $n \propto 1/r_0^2$, where r_0 is the Fried parameter discussed in Section 1.2.4. Since $r_0 \sim \lambda^{6/5}$ (Fried, 1965) this means the number of turbulent cells in front of the telescope will be $n \propto \lambda^{-12/5}$. This scaling shows that there is significantly more turbulence as λ decreases and instruments move from the IR to the visible regimes.

Woofer-Tweeter Architecture

MagAO-X's woofer-tweeter architecture is a common one for AO control. It is a two-stage control system wherein the woofer, an Alpao DM-97 deformable mirror, is a part of the Low Order Wavefront Sensor (LOWFS) and the tweeter, a Boston Micromachines Corp. (BCM) 2k deformable mirror, is used in the High Order Wavefront Sensor (HOWFS). The LOWFS aims to correct wavefront aberrations with low spatial frequencies, such as tip and tilt errors. It will make coarse corrections to the wavefront entering

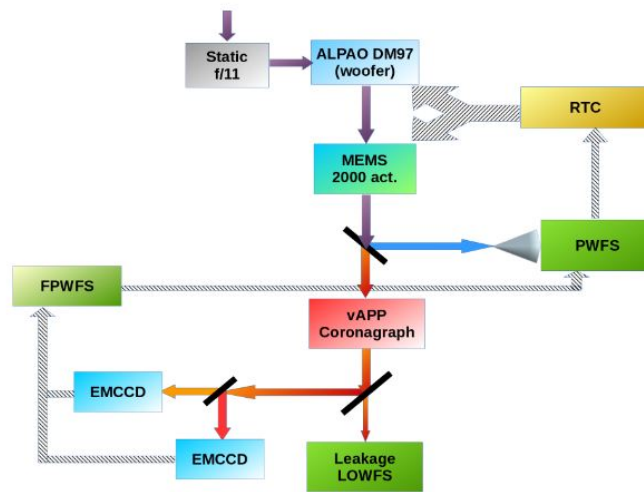


Figure 2.3 The MagAO-X woofer-tweeter architecture for the instrument’s Phase I configuration (without an MKID camera). Reproduced from Males et al. (2018).

the telescope and reduce the wavefront error sufficiently so the tweeter DM does not saturate trying to make a correction beyond the limits of its actuators. The HOWFS then cleans up the remaining wavefront errors with high spatial frequencies. The high order control is comprised of a Pyramid WFS and 2000-actuator deformable mirror.

Phase I

MagAO-X Phase I (Figure 2.3; Close et al., 2018; Males et al., 2018) involved the development of the instrument itself and initial on-sky demonstration. This phase of the project utilized the woofer-tweeter architecture before sending light through a vector Apodizing Phase Plate (vAPP; Doelman et al., 2021) coronagraph, a high-efficiency coronagraph that operates by creating two PSFs with opposite circular polarizations that destructively interfere, removing light from the central star. Light is then passed to two identical Electron Multiplying Charge Coupled Device (EMCCD) cameras to perform simultaneous differential imaging. This architecture allows the EMCCDs to act as focal plane wavefront sensors (FPWFSs), providing feedback to the real time computer (RTC).

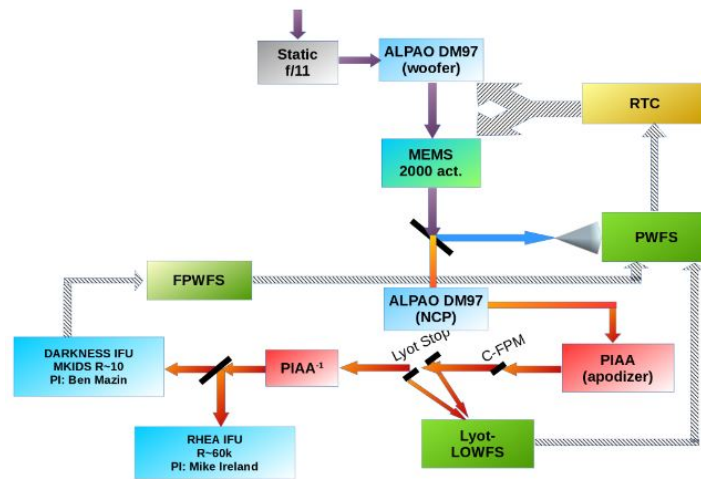


Figure 2.4 The MagAO-X woofer-tweeter architecture for the instrument’s Phase II configuration, with an MKID camera a new visible IFU, and coronagraph. Reproduced from Males et al. (2018).

First light was in December 2019 and a second commissioning run took place in April 2022. The scheduled runs in 2020 and 2021 were cancelled because of the COVID-19 pandemic. The science goal for Phase I was to search for and characterize accreting planets in $H\alpha$ around nearby T Tauri and Herbig Ae/Be stars (Males et al., 2022). T Tauri stars are variable stars younger than 10 Myr, less than $2M_{\odot}$, and of spectral type F, G, K, and M. Herbig Ae/Be stars are also young stars (typically $\lesssim 10$ Myr) of spectral type A or B and range from 2-8 M_{\odot} that tend to be found embedded in gas envelopes, possess circumstellar disks, or both. Each of these classes of star are ideal to search for forming planets around due to their young ages, leading to favorable contrast ratios and high likelihoods of hosting protoplanets still in the process of formation.

Phase II

Phase II (Figure 2.4) represents a significant upgrade to the MagAO-X instrument. First, the control and compute computers underwent an upgrade to reduce latency and

enable more effective real-time processing, further improving wavefront control. There is also a new DM added after the woofer-tweeter to measure non-common path (NCP) errors prior to sending the light to a phase induced amplitude apodization complex mask coronagraph (PIAACMC, or PIAA for short). The PIAA also requires a Lyot coronagraph in the system, which was installed during this upgrade. Finally, the EMCCDs were to be replaced by an MKID IFS and a visible-light IFS (Males et al., 2022).

The science goals for Phase II remained largely similar to those from Phase I, although the hardware upgrades are expected to lead to increased AO performance, enabling high-contrast imaging around even fainter stars than before (down to a limiting magnitude of ~ 12 in *I*-band ($\lambda \approx 800nm$). The incorporation of an MKID IFS will also hopefully lead to real-time wavefront control using MKIDs as a simultaneous science camera and FPWFS (Males et al., 2020; Males et al., 2022), which has so far been demonstrated using MEC at the Subaru Telescope (Fruitwala et al., 2018) in the near IR.

The original Phase II plan was to bring DARKNESS (Meeker et al., 2018) to LCO as the MKID IFS, but due to the increased cost of using liquid cryogenics operating DARKNESS would cause a significant increase in operational costs. This was the impetus for the decision to send a different MKID camera that could be operated without liquid cryogenics. When this decision was made, the PICTURE-C MKID Camera (Section 1.4.4) was undergoing the most active development and was a near-clone of DARKNESS, save for its ability to operate using a pulse tube for its initial cooling. However, it was designed to be slightly larger than DARKNESS to accommodate more liquid cryogenics for a longer hold time. The heat load from the excess material proved too great for the installed pulse tube, which wasn't able to cool it to a low enough temperature to operate the MKIDs. This discovery led to a third and final decision, to modify the ARCONS cryostat (Section 1.4.1) to accommodate a modern 10,000 pixel MKID array.

The required modifications to the original cryostat and use of a new-generation MKID

digital readout (Fruitwala et al., 2020) led to the transformation of ARCONS to XKID, now christened the MagAO-X MKID instrument. The remainder of this chapter will focus on XKID’s design and Chapters 3 and 4 will focus on its use and commissioning.

Phase II began following the return of MagAO-X to UA after its April 2022 run, where it started to receive several of the upgrades described above. It returned to LCO in February 2023, where XKID was incorporated with it at the telescope for MagAO-X’s third commissioning run and XKID’s first light.

2.3 Cryostat Design

The XKID cryostat is a pulse tube cooled adiabatic demagnetization refrigerator (ADR) originally purchased from Janis Research Inc. that can reach temperatures as low as ~ 50 mK. The outer vacuum shell is a cylinder that measures 16 inches in diameter and 41 inches tall, including the pulse tube. The pulse tube is a Cryomech 5W PT-405 cryocooler. The cryostat itself has 5 RF input ports and 5 RF output ports to read out the 10,000 pixel MKID array, thermometry to monitor temperatures at each stage, and ports to provide power to cryogenic amplifiers. There are 3 shells at 300 K, 50 K, and 3 K. The cryostat is shown both opened and closed in Figure 2.5.

The outer shell is at 300 K (room temperature) and holds a vacuum for the cryostat, without which it could not operate. It prevents air and other gases from entering the cryostat and freezing out on the cold stages, creating potential thermal touches between stages, excess heat loads, and hindering instrument performance.

The 50 K stage and shell are cooled by the first stage of the pulse tube. It has heat sinks for wires and cables that terminate on lower temperature stages, which provides an intermediate step down in temperature so there is not a direct path from 300 K to 3 K.

The 3 K stage holds most of the active parts in the cryostat. There are 5 cryogenic




Figure 2.5 (Left) The open XKID cryostat with the 50 K plate shown at top to the device stage at the bottom. (Right) The closed XKID cryostat with the 300 K outer shell installed.

HEMT (High-electron-mobility transistor) amplifiers, the ADR unit, a mechanical heat switch for ADR operation, and the device stage (or detector package).

The device stage (Figure 2.6) is held off from the 3 K stage itself by two copper struts so it does not interfere with the ADR but remains in thermal contact with the 3 K stage. The intermediate 1 K stage is supported by carbon fiber posts to provide sturdy, robust structure with low thermal conductivity. The MKID array hangs below the 1 K stage using Vespel SCP-5050 support posts that are similarly rigid to the carbon fiber supports and have low thermal conductivity, thermally isolating the array from the rest of the fridge. Typically there will also be a black baffle mounted to the 1 K ring to prevent any off-axis light from hitting the MKID array while allowing on-axis light through. There is also a high-permeability Amumetal magnetic shield (Figure 2.5, bottom left) that is mounted over the 1 K stage and MKID array to prevent magnetic fields from interacting with the MKIDs.

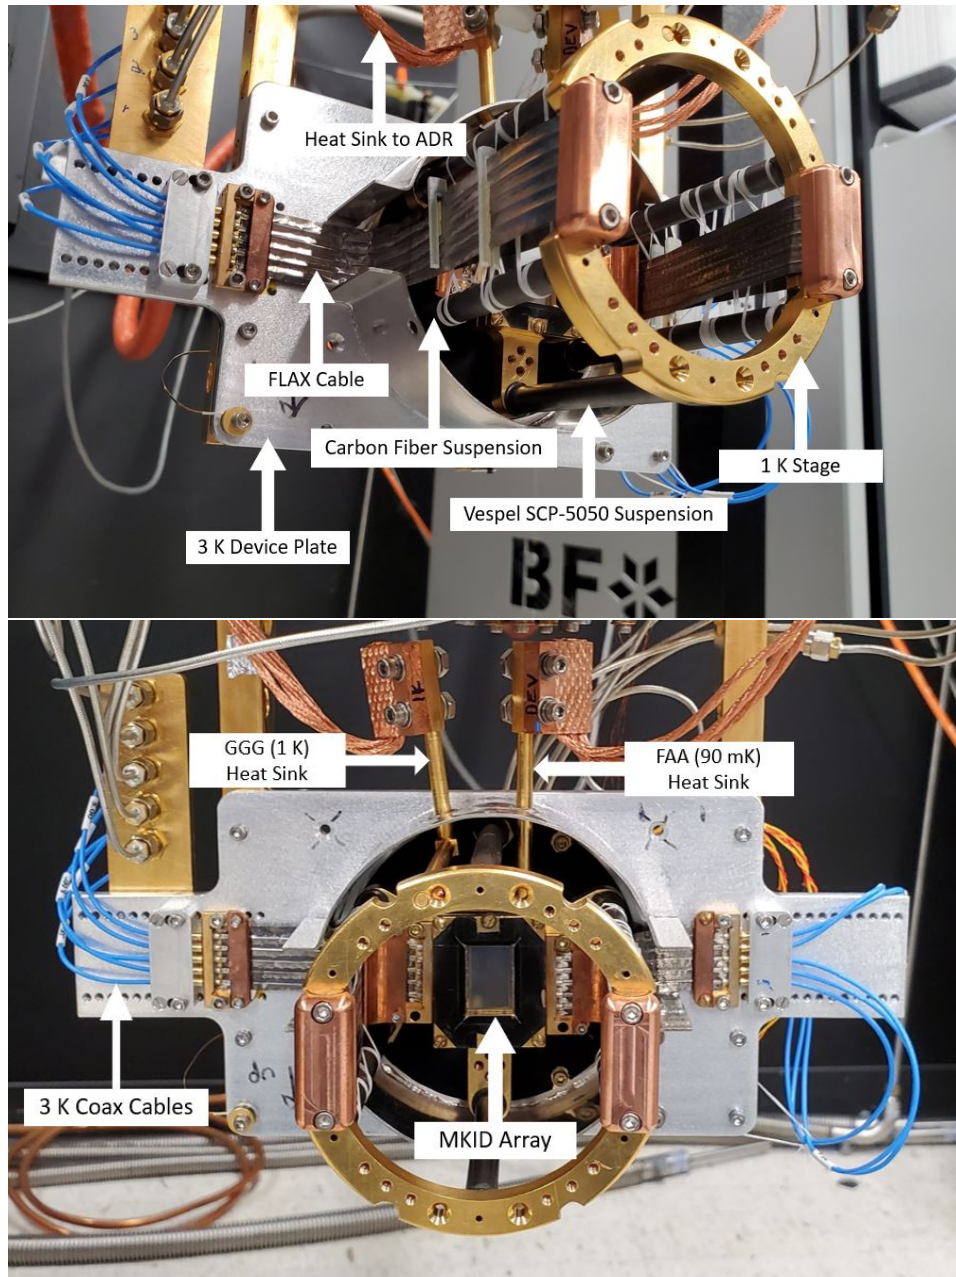


Figure 2.6 The XKID device stage. The aluminum plate at the back is at 3 K and is the structural base for the device stage. (Top) Side view showing the carbon fiber suspension holding the 1 K ring off from the 3 K plate and Vespel SCP-5050 suspension holding the MKID array below the 1 K ring. Also shows the FLAX cables (Section 2.8) traveling from the 3 K stage, being heat sunk at the 1 K ring, and attaching to the MKID array. (Bottom) Face on view more clearly showing the MKID array, the copper rope heat sinks to the ADR, and coax cables that carry signals to and from the 3 K plate. The HEMTs are mounted on the opposite side of the 3 K plate.




Figure 2.7 The 10,000 pixel XKID array in its microwave mounting box. Not shown is the lid with MLA that prevents stray light from hitting the array and focuses light onto each of the MKIDs.

The MKID array itself is mounted in a microwave package (Figure 2.7) that measures roughly $1.5'' \times 1.25'' \times 0.325''$. It has external ports to allow microwave signals in and out and – when in use – has a lid with a microlens array (MLA) to focus light onto the photosensitive inductor of each MKID pixel. The box is attached via a copper rope heat sink to the ADR which enables cooling to its operating temperature.

2.3.1 Pulse Tube

The pulse tube (Figure 2.8) used in the XKID fridge is a PT-405 from Cryomech Inc. It has a cooling capacity of 0.5 W at 4.2 K and 25 W at 65 K. It has two stages at a high and low temperature, which in XKID typically sit at 50 K and 3 K, respectively.

The pulse tube capacity curve in Figure 2.9 shows that the pulse tube is absorbing about 14 W from the 50 K stage and roughly 0.25 W from the 3 K stage, meaning that it is operating well below its maximum capacity and is not being overtaxed or strained.




Figure 2.8 CAD drawing of the Cryomech PT-405 pulse tube. The two stages (shown here in copper) allow the cryostat to step down in temperature from 300 K to 50 K (the upper copper plate) and from 50 K to 3 K (lower). The ports at the top allow compressed Helium into and out of the pulse tube as the working material for the closed cycle refrigerator. Reproduced from <https://www.cryomech.com/products/pt405/>




Figure 2.9 The PT-405 heat capacity curve. The temperature of each pulse tube stage show the relative heat loads on each. The XKID cryostat regularly gets to between 46-50 K at the high temperature stage and 3 K at the low temperature stage. This means the pulse tube is absorbing ~ 14 W and ~ 0.25 W at the high and low temperature stages, respectively. Reproduced from <https://www.cryomech.com/products/pt405/>

Things like thermal touches or excess thermal loads (too much material to cool, light leaks allowing hot blackbody radiation to hit a low temperature stage, etc.) may be extreme enough that the pulse tube will not be able to get the fridge to cold enough temperatures and make it impossible to operate the MKIDs.

2.3.2 Adiabatic Demagnetization Refrigerator

The ADR in XKID came installed from Janis Research Inc. Figure 2.10 shows the ADR when it was removed from the cryostat for repairs and modification. The purpose of the ADR is to cool the MKIDs from 3 K to their operating temperature of 90 mK. To do this it uses two paramagnetic salt pills made of gadolinium gallium garnet (GGG) and ferric ammonium alum (FAA) in the bore of a ~ 4 Tesla solenoid shaped superconducting magnet and a heat switch. The operation principle relies on the fact that in the presence of a strong magnetic field the paramagnetic salt pills will have low entropy because the majority of their spins in the crystal lattice will align, and therefore become more ordered, as the field strength gets higher. The process of using the ADR to cool the MKIDs from 3 K to 90 mK is as follows:

1. The mechanical heat switch is closed, bringing the pills into thermal contact with the 3 K stage which acts functionally as an infinite heat bath.
2. Current is sent through the superconducting magnet, ramping up at a constant rate until the magnetic field reaches its maximum value of ~ 4 T (at a current of ~ 9.25 A). One must take care not ramp at a rate that would cause back-emf across the magnet of $\gtrsim 250$ mV, which can damage the ADR.
3. The current is left at its maximum value for the pills to “soak” in the magnetic field for 1 hour, allowing the spins to align and the pills to thermalize, letting any excess heat from the current ramp to flow into the 3 K stage.




Figure 2.10 The XKID ADR removed from the cryostat. It uses a 4 Tesla superconducting magnet (not shown, internal to the copper shields and metal bore) to align the spins of two paramagnetic salt pills, release them from thermal contact with the rest of the fridge using the heat switch, then draw heat from the MKID array and support structure, providing ~ 100 mJ of cooling power. The gold-plated circular plate is used to mount the ADR to the cryostat. The ADR can cool the MKIDs to their operating temperature of 90 mK from 12 to 24+ hours, depending on the heat load in the fridge.

4. The mechanical heat switch is opened, taking the salt pills out of thermal contact from the rest of the fridge.
5. The current is decreased at a constant rate to 0 A. As the field decreases it no longer forces the magnetic spins in the pill into alignment so heat can flow into the pill to increase the entropy of the spins and cause them to become more disordered.

The operating principle for both salt pills is identical since both are paramagnetic, but their cold temperatures and heat capacities are different. The GGG pill will get to temperatures between 550-700 mK and can absorb ~ 1 J of heat. The FAA pill can fall as low as 50 mK but has a heat capacity below 100 mJ. The GGG is used as a buffer between the FAA and low temperature stage and the hotter 3 K stage to reduce the heat load on the MKID array, letting it stay colder for longer. The GGG pill is physically connected to the 1 K stage via a copper rope heat sink between the stage and its support strut. The FAA pill is connected to the device stage in much the same way. The copper rope heat sinks and posts can be seen in the bottom panel of Figure 2.6. To intercept heat at the 1 K stage the microwave wiring is routed from the 3 K stage, over the 1 K stage where it is heat sunk using copper blocks, then back down to the MKID array.

Although the FAA pill can drop to temperatures as low as 50 mK, it is important to keep the MKID array (and FAA pill itself) at a fairly constant temperature. This is because the number of thermally generated quasiparticles increases with temperature, which can lead to decreased quality factors Q_i and changed resonant frequencies f_0 , which prevent resonators from being optimally read out. In XKID a feedback loop is used to maintain a temperature of 90 ± 0.03 mK. The feedback loop consists of a thermometer at the cold stage, a proportional-integral (PI) controller, and a superconducting magnet power supply. First, the Ruthenium Oxide RTD (resistance temperature detector) measures the temperature of the MKIDs. Then the PI controller sends a small analog voltage

to the magnet power supply, which adjusts a small current to correct for any differences between the currently measured temperature and the temperature setpoint of 90 mK. This is a continuous feedback loop that can be tuned to offer very precise temperature control down to $\pm 30 \mu\text{K}$, although it has been empirically shown that it is incredibly sensitive to electromagnetic interference and things like walkie-talkies being used in the vicinity can cause the PI control loop to go haywire.

After the spins in the paramagnetic salt pills have become maximally disordered and their cooling capacity has been exhausted, another cycle of the ADR magnet must be run to realign the magnetic spins and cool the device back down. For this reason the ADR is called a “single-shot” magnetic cooler: for each magnet cycle it can drop to its cold temperature once before needing to go through the entire cooling process again.

2.4 Footprint at Magellan Clay Telescope

The XKID instrument was designed to have a small footprint, reducing space required at LCO and on the NASE platform at Clay. Figures 2.11 and Figure 2.12 show a schematic diagram and CAD model of XKID and MagAO-X on NASE.

The XKID cryostat and mounting frame sit next to MagAO-X on NASE, roughly 2” from the edge of the optical bench. This allows the calibration snout that hangs off the front of the cryostat to align with the XKID foreoptics “diving board” that hangs off of MagAO-X (see Section 2.6 for further details). This spacing will allow a person to have clearance to access components on the opposite side without bumping into it. The frame has one 120 V/60 Hz power cable that attaches to an outlet near the electronics rack and 1 USB cable which attaches to part of the foreoptics contained on MagAO-X itself.

The electronics rack is also stationed on NASE in the corner closest to XKID. Between the two there are 2 24-pin Fischer connector cables for thermometry and powering the




Figure 2.11 Schematic block diagram showing the overhead layout of XKID on NASE at Clay. Red shows power connections, blue shows operational connects (glycol for cooling and helium for pulse tube operation), green shows communication connections. The yellow arrow shows the position of the output F/11 beam from the telescope.

HEMT cryogenic amplifiers, 10 coax cables (1 input and 1 output for each array feedline), 1 USB extension cable that runs between the control computer and a USB hub on the frame, and the cable that provides current to the ADR. It also has 2 120 V/60 Hz power outputs that plug into outlets at the edge of the platform. An ethernet cable is also routed from the rack to a port on the platform so XKID can be accessed on the local network. The rack is connected to the system glycol in series with MagAO-X.

The only component of the instrument that does not live on NASE is the compressor that enables the pulse tube to cool to 3 K. The compressor is situated in the pier, which is a tunnel below Clay and Baade. In the pier, the compressor is connected to 3 phase, 220 V/60 Hz power. 2 40-meter helium hoses with aeroquip connects and 1 40-meter power cable are routed from the pier and up through a cable wrap (Figure 2.13) where they terminate at NASE. There, a second set of helium hoses and power cable extension are attached to complete the connection of the compressor to the pulse tube.




Figure 2.12 CAD model showing XKID (orange frame with blue dewar) on NASE. The F/11 Nasmyth port is the black ring at the top left which feeds into MagAO-X, the grey optical table. The Gen2 MKID readout electronics rack (green) is shown here without side walls or readout crate installed. The ladder down from NASE is at the top right next to the electronics rack. The remainder of the platform safety railing and the MagAO-X electronics rack are not shown.



Figure 2.13 Compressor hoses being installed at Clay. (Left) Hoses being routed from the pier to NASE, the blue platform at the top center. (Right) Into the cable wrap after installation.

XKID is designed to be mobile while in its frame and to have a small footprint so it can be moved up to NASE before MagAO-X so it can start cooling and getting set up (the process of which takes about 3 full days) but still move around NASE to allow any instrument that is already up there to be taken off and for MagAO-X to be moved up to the platform without stopping any of XKID's operations.

2.5 Mounting Frame

The XKID mounting frame (Figure 2.14) is shown in orange in Figure 2.12. It is made up of four components:

- An aluminum frame from structural support and to suspend the cryostat.
- An aluminum assembly to attach the cryostat to the frame and provide linear motion control in the X and Y directions.
- A steel base with wheels and lifting jacks to move the full assembly and provide control in the Z direction.
- A steel ballast structure to lower the center of gravity to prevent swaying and improve earthquake stability.

The aluminum frame was fabricated by the Rettig Machine Shop in Redlands, California and the other three components in the UCSB Physics and Chemistry Machine Shops. This section will briefly detail each component of the frame.

Frame

The purpose of the frame is to provide a structure to mount XKID in that will prevent swaying, twisting, and other unwanted motion of the cryostat, allow access to the the cryostat if maintenance is required, all without obstructing the optical path.



Figure 2.14 XKID mounted in its frame for the first time at LCO with Noah Swimmer (myself, left) and Jeb Bailey (right). The base can be seen at the bottom on its casters with the front lifting jacks visible. The ballast is the large mass held up by the cross beams attached to the base. The frame is the tall white vertical structure, and the linear motion control assembly is mounted at the top, which the cryostat is hanging from.

The frame consists of 4 legs splayed out forward and backward to increase stability while maintaining a tight side-to-side profile to ensure the ability to walk past it on NASE. There are horizontal cross braces in the back and on each side to provide further structural stability, with none on the front to prevent optical collisions with the incoming beam, and diagonal cross braces on each side to resist twisting modes in the frame.

Linear Motion Assembly

The linear motion assembly (Figure 2.14, top center) is designed to allow fine control in the X and Y directions when XKID is on NASE. Each stage consists of 3 linear rails and runners from Bosch Rexroth to allow for smooth movement, a 1/2"-20 UNF bolt to drive the motion, and a brake to lock the stage in place once alignment is reached.

Base

The base is comprised of two pieces, a bottom U-shaped section that the frame bolts onto and two vertical tubes to mount the lifting jacks on. There is also a front section that closes the U-shaped piece into a closed rectangle and also has two vertical tubes for mounting the opposing lifting jacks to balance the frame on. On the bottom there are 4" fully-locking swivel casters mounted for easy movement throughout the observatory and coarse alignment adjustment. It was made of steel to add extra weight to the bottom (in addition to the ballast) to make the sure the structure is more stable.

Ballast

The ballast weight was included to ensure the stability of the instrument to potential earthquake activity at the telescope. It is comprised of two tubes that mount to the U-section of the base and support the ballast weight that hangs below. The weight hangs beneath the support tubes to ensure the center of gravity is as low as possible.




Figure 2.15 Figure showing the angle between the center of gravity of the XKID frame and instrument and the pivot points in each direction.

The weights consist of 10 50-lb steel plates (12" x 8" x 1") and 1 100-lb weight (12" x 16" x 1"). The heaviest weight sits at the bottom where the 10 smaller plates are stacked side-by-side and sit on top of it. 4 bolts are then passed through holes in the plates and threaded into the holes in the large steel plate to hold the assembly off the ground.

2.5.1 Earthquake Stability

The theory and design for making the instrument stable to earthquakes is adapted from the explanation passed to the team at UCSB by Jeff Crane at Carnegie.

An object will tip over if its projected center of gravity (CG) moves past its ground contact points. If one projects the CG straight down to the ground, the tipping threat

will be greatest toward the nearest footprint border since it has the shortest distance to go to go over that contact point. The ground contact point is the pivot point for tipping and should be used for calculating the destabilizing and restoring moments.

At Magellan the assumption is that during an earthquake there will be a 0.15g upward vertical force (destabilizing) and that the remaining 0.85g downward force of gravity will have a stabilizing effect. The stabilizing moment itself will then be

$$\mu_{\text{stabilizing}} = 0.85 \times w \times h_{\text{CG}} \quad (2.1)$$

where w is the weight of the frame and h_{CG} is the height of the CG from the ground. The horizontal destabilizing moment to get the CG over its pivot point will then be

$$\mu_{\text{destabilizing}} = 0.4 \times w \times \Delta x \quad (2.2)$$

where w again is the weight and Δx is the horizontal distance between the CG and the pivot point in question. The stabilizing moment should ultimately be made to be greater than the destabilizing moment, ideally with a margin for error.

Using this formulation, one can find the minimum angle between the imaginary line vertically down from the CG and the line connecting the CG and ground contact point. The minimum required angle to be stable to tipping is then

$$\theta_{\min} = \arctan\left(\frac{\mu_{\text{destabilizing}}}{\mu_{\text{stabilizing}}}\right) = \arctan\left(\frac{0.4w\Delta x}{0.85wh_{\text{CG}}}\right) \approx \arctan\left(\frac{0.4}{0.85}\right) \quad (2.3)$$

which, when evaluated leads to $\theta_{\min} \approx 25.2^\circ$. Figure 2.15 shows a schematic drawing with the CG and the closest ground contact point in each direction. It shows that even in the narrowest direction the angle θ between the CG vertical line and the line between CG and ground contact is 25.7° , well above the required tipping threshold of 25.2° .

2.6 Optical Design

The optical design of the XKID instrument is relatively simple and largely consists of foreoptics to bring light out of MagAO-X and into XKID, a calibration snout, a flipper mirror to switch between the two, several windows at the 3 K and 50 K shells with IR filters, and a microlens array (MLA) to focus light onto the MKIDs.

For typical instruments that have small form factors MagAO-X has a 12" x 28" visitor port that consists of a breadboard in the corner of the optical bench itself that is enclosed by the dust cover for a user to mount their instrument in, where the beam position is well understood and can be picked off to send to the detector. The input to the visitor port from MagAO-X was reported by the team at UA to be an F/69 beam and the choice of optics in the path reflects this value. The focus of the beam into the visitor port is designed to be 18" from the short edge of the breadboard nearest the outside of the bench and 6" from either of the long edges. It is also designed to be 4.5" above the visitor port breadboard (or 5" above the optical table itself).

Using XKID in the visitor port was unfeasible for two reasons. First, the cryostat is too large to fit on the optical table and would have required one of the dust covers to be significantly modified to let the cryostat to poke out, leading to stray light entering MagAO-X. Second, the foundation of MagAO-X is an optical table that has actively controlled air cushions to keep it leveled with extreme precision and to isolate it from vibrations on the NASE platform. With this being said, the XKID cryostat is physically attached to NASE via the helium hoses that enable pulse tube operation which are not vibration isolating and could introduce uncontrolled vibration to the optical table, degrading MagAO-X's ability to generate the exquisite AO corrections that it promises. For these reasons the decision was made for XKID to sit alongside MagAO-X and use a foreoptics system to bring the light from the visitor port to the MKID array.




Figure 2.16 (Left) CAD design of the foreoptics rail from MagAO-X to XKID. The thin brown lines show the optical axis from the meniscus lens and the optical axis to the XKID instrument. (Right) Photo of the first XKID rail installation on MagAO-X in the cleanroom at LCO. The meniscus lens is not shown but was installed internally.

2.6.1 Foreoptics: MagAO-X to XKID

The foreoptics that bring light from MagAO-X’s visitor port consist of several optics mounted to a 750 mm rail that has one side mounted to the MagAO-X optical and the other end in a “diving-board” configuration off its edge. All of the optics on the rail are off-the-shelf parts from ThorLabs or Newport. Figure 2.16 shows a CAD model of the foreoptics rail and what it looks like when they are integrated into MagAO-X.

The first optic on the rail is an $f = -300$ mm N-BK7 meniscus lens at the visitor port focus position. The focus position is not well-determined because it has a depth of focus of $\sim 1/4$ ”, so the meniscus lens is placed on a linear-motion stage with a 50 mm range of motion. An attempt was made to mount the meniscus lens in such a way that the focus point was in the middle of that range so the user may adjust it in either direction. The meniscus lens sends the beam to the output port of the MagAO-X dust cover and is the only XKID optic physically contained in MagAO-X.

Table 2.2. XKID Selectable Optics Parameters

| Parameters | Values |
|--|---|
| Flipper Mirror | Up (Light from MagAO-X) Down (Light from calibration sources) |
| Calibration Laser Wavelength | 808 nm, 904 nm, 980 nm 1120 nm, 1310 nm |
| Filter Wheel Diffusers and COlor Filters | Closed, Y, z_s Manuakea J Diffuser (high grit) Diffuser (low grit) Open |

Next, there is not an optic but a modification plate made to the MagAO-X dust cover to enable the linear motion control power and communication cable to be routed out to the control computer and a 2" port for the light to leave the MagAO-X bench. The new plate was custom made by Ben Mazin and is designed to fit in a hole that was cut into the MagAO-X dust cover next to the visitor port for this purpose. When it is in use, all of the holes are blocked off and when not in use, the plate is removed and the MagAO-X team blocks out the hole using dark masking tape and other light insulation.

The dust cover modification plate has a port for a 2" optical tube to be screwed into it. The next part of the foreoptics is a 2" diameter, 12" long optical tube to baffle the light path outside of the MagAO-X dust covers. It will prevent stray light from entering and ambient light from contaminating the telescope beam outside of the optical bench.

The final optic is a Conex Piezoelectric driven mirror mount with 2-axis control. It serves as a tip-tilt mirror that can place the beam on different parts of the array.

2.6.2 Calibration Snout

The second piece of the foreoptics is the calibration snout, shown in Figure 2.17. The calibration snout is designed so that the XKID instrument can be fully set up and calibrated without manual intervention. The design stems from lessons that have




Figure 2.17 (Left) CAD model of the XKID calibration snout showing with parts labelled. Not shown is the cap piece that covers the flip mirror to reduce dust and to prevent stray light from getting in. (Right) A picture of XKID with the calibration snout mounted to the cryostat and aligned to the foreoptics rail that comes off of MagAO-X.

been learned with the MEC and DARKNESS instruments, each of which require the instrument to be in separate physical configurations for setup and operation.

The calibration snout consists of a laser calibration source and diffuser, a “beammapper” phone, a beamsplitter to bring in light from either, a variable aperture, a 1:1 imaging relay lens, a fold mirror, a flipper mirror to select between light from MagAO-X or from the calibration sources, and a 2” filter wheel. Several of these components have selectable parameters that are shown in Table 2.2.

There are two inputs to the XKID calibration snout: the laser calibration source and the beammapper phone (calibration processes will be described in Section 3.2). The laser calibration source is input into a 2” optical tube using a standard ThorLabs fiber port that flood illuminates a diffuser 140 mm away. The diffuser is 90 mm away from a 50:50 beamsplitter. The beammapper phone is also mounted so its screen is 90 mm from the

beamsplitter. The beamsplitter is an off-the-shelf part and is oriented to send light from both sources further down the snout to XKID.

Following the beamsplitter there is a variable aperture ~ 390 mm farther down the optical path, which in turn is followed by a 1:1 reimaging relay lens ~ 45 mm away. The variable aperture is typically opened to 8 mm diameter. This value can be increased to allow more light into the system in cases such as one of the calibration lasers being too dim, but it can also significantly degrade the image quality and so is not recommended other than for flood illumination purposes. After the reimaging relay lens there is a 2" fold mirror that allows the optical path to take a 90° turn so that it is travelling horizontally rather than vertically.

The fold mirror is followed by a flipper mirror that can be moved down “into” place, allowing light from the calibration snout to enter XKID, or up “out of” place, allowing in light from MagAO-X. The flipper mirror is remotely controllable so a user may easily switch between calibration and on-sky sources without requiring reconfiguration to attach or detach optics. This lets the calibration snout stay permanently attached to XKID, even during normal on-sky operation.

The final optic in the calibration snout is a Finger Lakes Instruments (FLI) 2" filter wheel with 7 usable filters (Table 2.2). It fits snugly onto the cryostat’s 300 K window holder and is the last optic outside of the cryostat.

Beammapper Phone

The XKID beammapper is a Samsung Galaxy S6 that has a “Calibrator” application designed by Seth Meeker to aid in the beammapping process for DARKNESS (Meeker et al., 2018) and has also been used with MEC (Walter et al., 2020). It is mounted in an aluminum box that mounts directly to the 2" optical tubes of the calibration snout.

Calibration Laser Source

The calibration laser source (or “laser box”) consists of an ArduinoUNO, 5 laser diodes, an integrating sphere, and a fiber port output from the integrating sphere.

The 5 lasers are at 808 nm, 904 nm, 980 nm, 1120 nm, and 1310 nm. They can be controlled by physical buttons on the outside of their aluminum housing box or via the Arduino itself. The On/Off physical power buttons only allow for binary control, with the laser either being fully off or powered at 100%. The Arduino gives more granular control where the user can step between 0-100%, allowing the user to set laser powers to intermediate values when desired. Each laser is mounted so it shines into the integrating sphere. The laser box output is a fiber port that is attached to the input port of the calibration snout by a 2-meter optical fiber ($\lambda=400-1400$ nm).

2.6.3 Cryostat Optics

In the cryostat the only optics are 3 windows, an off-the-shelf ThorLabs for the 300 K vacuum shell and 2 custom filters from Custom Scientific for blocking IR light at 50 K and 3 K and a custom microlens array.

Starting at 300 K there is an uncoated 2” diameter, 12 mm thick window to let light into the cryostat. The window holder has two O-rings, one situated around the window and one beneath it to ensure the cryostat’s outer shell holds vacuum.

Next there is another 2” window at the 50 K stage. It is 10 mm thick and has a coating designed to maximally allow 800-1400 nm light through. There is a sharp drop off at $\lambda > 1400$ nm and a more gradual tail dropping off below $\lambda < 800$ nm. This short wavelength tail is acceptable because the majority of contaminating photons are low-energy, long-wavelength thermal photons from an IR blackbody. The final window is at the 3 K stage and has an identical coating to the window at the 50 K stage for further

IR blocking and is 20 mm thick. The 50 K and 3 K IR blocking windows are mounted in custom holders that hold them at 3° angles from the optical axis (i.e. the vector normal to their front face is 3° away from the optical axis) to prevent any ghosting. Ghosting occurs when one sees fainter copies of an object due to reflections between planar optics.

After the 3 K filter window there is a baffle mounted to the 1 K stage. The baffle was painted using non-specular black paint and is corrugated on the inside to prevent stray light from entering the optical path. Finally, there is a microlens array (MLA) that was manufactured by Advanced Microoptic Systems GMBH. It is 15 mm x 21.75 mm x 1 mm and made of S-TiH53. The lenslets have a focal length of $f=0.95$ mm, and they are mounted at that height above the array itself. The MLA is mounted in the lid of the MKID array microwave package which was designed by Seth Meeker for DARKNESS (Meeker et al., 2018) to hold the MLA at exactly this height.

Changing Filter Windows

The window holders at 50 K and 3 K are designed to be easily removed from the shells so they can be removed and replaced with a separate filter stack if desired. XKID also has a set of Asahi Supercold 400-1100 nm filters for observing runs where visible-light astronomy is desired. Reinstallation requires warming up to room temperature, removing the vacuum shells, unscrewing the window holders from the shells, replacing the old filter windows with whatever new 2" filter windows are desired, then reversing the process to install the new filters and close the cryostat back up.

2.6.4 Image at MKID Array

The optical path was designed and modeled using Zemax OpticStudio to verify the expected behavior of the optical system, the results of which are shown in Figure 2.18.

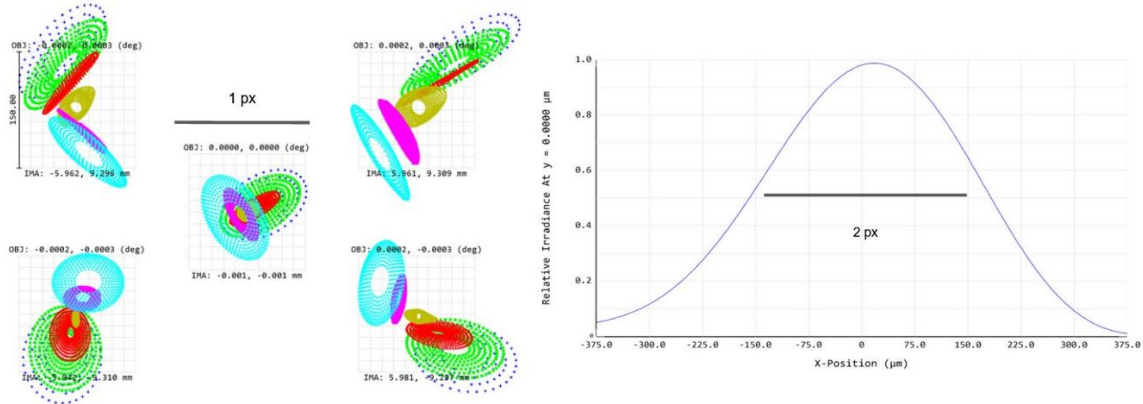


Figure 2.18 (Left) Focus spots shown at 5 wavelengths between 800-1400 nm at the center of the array and at each corner. The scale bar shows the size of one pixel. At the field center the focus is not wavelength dependent but at the field edges, there is a wavelength shift of ~ 1 pixel across the bandpass. (Right) The shape of the PSF at the corner of the field. The optical design is for the PSF to take up ~ 2.15 pixels at field center, and this shows that the PSF also spans about 2.15 pixels at field edge, so there is not field-dependent distortion.

At the center of the MKID array, the beam is focused onto a single pixel at all wavelengths in the bandpass. At the field edge there is a shift of ~ 1 pixel as one goes through the bandpass, meaning that at low wavelength end of the bandpass the spot will be focused at one edge of a pixel and at the high wavelength end the spot will be focused at the opposite end. However, the PSF core remains constant in size between the field center and field edges. This means that although there is some wavelength dependence to the focus of the optical train, the ultimate broadband image should not vary over super-pixel scales across the field of view.

2.7 Readout Electronics and Hardware

The XKID electronics rack (Figure 2.19) sits at the corner of NASE separated slightly from the XKID cryostat to allow a path past the instrument but close enough to run cables between the two. It requires ~ 1.5 kW at maximum power split between two 120

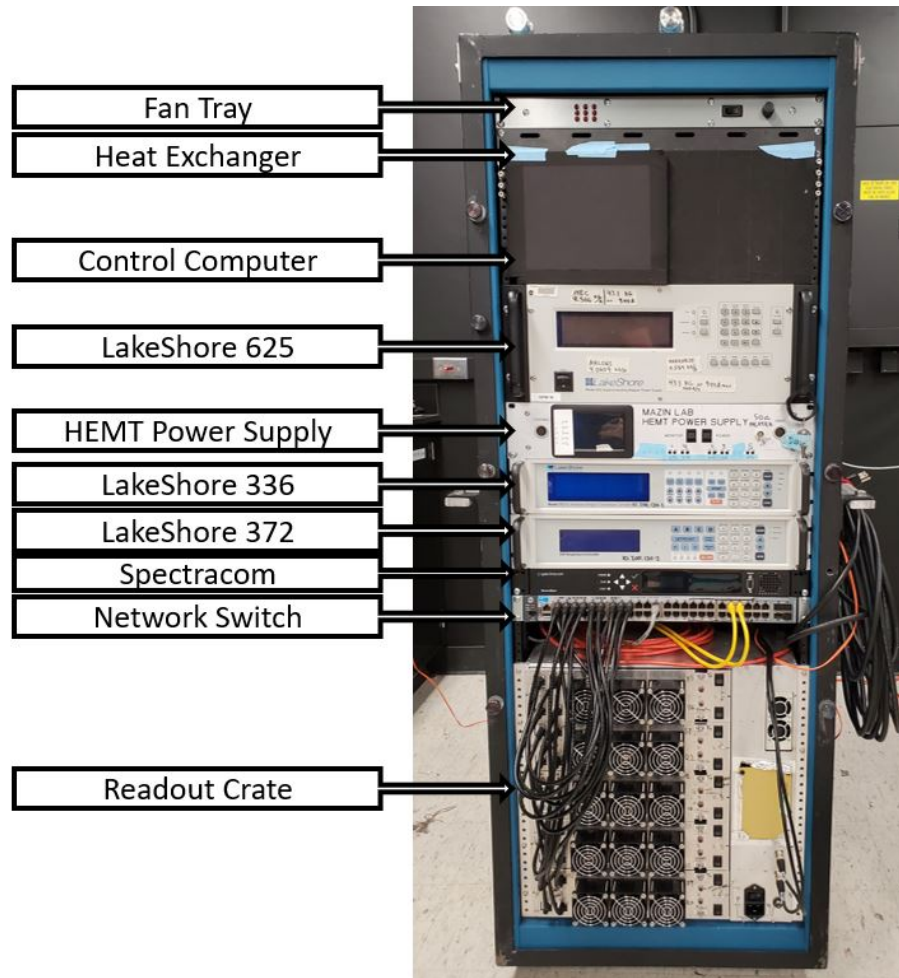


Figure 2.19 Labelled XKID electronics rack. The image at the right shows the rack in the process of being wired up. The labels at the left point to each of the different components in the rack. The glycol ports can be seen at the top.

V/60 Hz single phase power cords. It can be hooked into the Clay glycol cooling supply directly or in series with MagAO-X. Nominally it is connected after MagAO-X in series and the output from the XKID rack goes to the NASE return port. The glycol flows at a rate of several gallons per minute in the back of the rack and draws heat from the heat exchanger at the top (for the “housekeeping” electronics) and at the back of the readout crate (for the readout boards). The fan tray at the top of the rack pulls hot air from the rack up through the top heat exchanger. The readout crate has its own fan system (seen at the front bottom of Figure 2.19) that draws in ambient air and pushes it over the readout boards and through the heat exchanger at the back of the readout crate.

At Clay each instrument is required to emit ≤ 50 W of heat into the dome to prevent local turbulence and dome seeing. Previous tests with the XKID rack measuring the change in temperature between the ambient input air and exhaust have made it challenging to provide an exact measurement of the heat dumped from the rack due to there being no discernible difference in temperature between the two. With this in mind the MagAO-X and Clay teams approved the use of the XKID rack at NASE as is.

The E-rack is broken into two parts, the telemetry (“housekeeping”) and the MKID readout. The former is used in service of monitoring and operating the cryostat itself while the latter is responsible for monitoring the MKIDs themselves.

2.7.1 Telemetry and Cryostat Readout

The telemetry portion of the XKID E-rack runs from the Network Switch and above (in Figure 2.19). This section will briefly detail each component and its purpose. In totality, the telemetry portion of the rack consumes $\lesssim 300$ W at its peak output when the ADR magnet is soaking at its nominal soak current of 9.25 A and will consume markedly less while regulating at current values of $\lesssim 150$ mA.

The network switch is an HP 2920-48G configured to have a virtual internal network on $\sim 1/2$ of the ports where the Gen2 readout boards can communicate with the control computer. The rest of the ports can be used on external networks. It has a 10 GbE fiber port in the rear that is used to send the MKID data to the readout computer.

The Spectracom is a GPS-locked Rubidium clock that has 10 MHz and 1 pulse per second (PPS) outputs used to condition the MKID readout boards. The 1 PPS signal gives the readout boards a reference signal once per second to synchronize to and the 10 MHz signal provides a reference tone so the boards can create clean tones to be sent into the cryostat to excite and read out the MKIDs. The spectracom can be configured to act as a follower of an NTP server or – if its antenna can see the sky – get a GPS lock to provide exquisitely clean 1 PPS and 10 MHz signals.

The Lake Shore 372 AC Resistance bridge is used to monitor the temperature of two resistance temperature thermometers (RTDs) at the device stage and 1 K stages. It can read them out at low powers (\sim femtowatts) greatly decreasing the likelihood of self-heating. It also provides the ability to run Proportional-Integral-Derivative (PID) loops to maintain the fridge at a certain temperature. To do this the Lake Shore 372 then uses the input signal from the device stage thermometer, conditions it and sends it to an output BNC port. This output is sent to the ADR magnet power supply which sends current into the ADR, raising or lowering the temperature accordingly. The thermometer's response is fed back into the Lake Shore 372 which adjusts the current appropriately and tries to get as close to the user-specified temperature setpoint as possible. Typically, it holds the MKID array at 90 ± 0.03 mK.

The Lake Shore 336 is another thermometry unit, though not an AC resistance bridge like the model 372. This module is responsible for reading out two diode thermometers stationed at the 3 K and 50 K stages. These stages do not require control and so the PID functionality of the Lake Shore 336 is not used in the XKID system architecture. It

also has the ability to read out RTDs in addition to diode thermometers which means that it can read out any of the 4 thermometers in the fridge.

The HEMT Power Supply is a custom module built originally for the DARKNESS instrument. It uses two Low Noise Factory (LNF) PS_EU2 constant current HEMT power supply boards, each of which are capable of powering 4 cryogenic HEMT amplifiers. Since there are only five in the fridge, one of the boards is used to power three of the HEMTs while the other powers the remaining two.

The Lake Shore 625 is the superconducting magnet power supply and is the one piece of the readout rack that is markedly different from its counterparts in MEC and DARKNESS (a Kepco BOP 20-10ML). The Lake Shore 625 is capable of pushing up to 120 A through a superconducting magnet and generating a user-specified current, a current proportional to an external signal (such as a PID control signal), or a sum of the two. While cycling the ADR to soak the salt pills and prepare them for MKID operation the it is controlled via the control computer, which tells it how quickly to ramp up to its soak current, how long to soak, and when to ramp back down. Once the magnet cycle is done and the MKIDs are at operational temperatures the PID loop kicks in and the output from the Lake Shore 625 is controlled via the output from the Lake Shore 372.

The control computer is a NUC 11 Pro (NUC11PAHi7) kit that contains a 4 TB solid state drive (SSD) and two 32 GB RAM cards. It is responsible both for the cryostat monitoring/control and the MKID readout. Previous instruments have had separate computers for these tasks, but with developments in the intervening years leading to increased processing power and efficiency these tasks can now be completed on the same computer. The control computer is attached to several USB hubs and one thunderbolt-to-10 GbE adapter. There is one 7-port powered USB 2.0 hub that stays in the rack and enables the NUC to connect to all of the different fiducial devices in the E-rack. The second 7-port powered USB 2.0 hub is on the XKID frame and connects the computer

to all of the electronically controlled components there (the heat switch, flipper mirror, laser box, focus slider, etc.). Finally, the thunderbolt-to-10 GbE adapter connects via an optical fiber to the 10 GbE port on the network switch to talk to and read from the Gen2 readout boards. It runs a standard Ubuntu LTS operating system and can be controlled via an SSH connection and the command line or via Remote Desktop Protocol (RDP).

The fan tray is a Kooltronic KT3X502 and the heat exchanger is a 723SN00A02. These are not controlled or monitored by the control computer. At power-up, the fans are switched on and left on in until the final powering down for a given observing run. The heat exchanger is also passive and works to draw heat the rack and pass it into the glycol that runs through it when the rack is hooked into the NASE glycol system.

2.7.2 MKID Readout

The readout for XKID is identical to that of DARKNESS, which is detailed in depth in Strader (2016), Fruitwala (2021), and Steiger et al. (2022a). This section will provide a brief overview of the system and recognizes that the design and characterization of this complicated system has been subject of several theses on its own, namely those of Paschal Strader and Neelay Fruitwala that were previously cited.

The 2nd Generation MKID readout for XKID uses ten ROACH2 (Reconfigurable Open Architecture Computing Hardware) boards developed by the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER; Hickish et al., 2016) each of which connects to an ADC/DAC² board and an RF/IF board, both designed by collaborators at Fermilab. Each readout board (Figure 2.20) can reading out 1024 MKIDs over a 2 GHz band. Since each feedline on the XKID array has 2000 resonators, each feedline requires two readout boards. Each set of two readout boards is combined into a readout “cartridge”, five of which are installed in the readout crate (Figure 2.19).

²ADC = Analog-to-Digital Conversion, DAC = Digital-to-Analog Conversion




Figure 2.20 A readout cartridge showing 1 readout board (the bottom side when flipped over is the other identical board). The ROACH-2 (right) is connected to the ADC/DAC board (left) via 2 ZDOK connectors. The RF/IF board is mounted above the ADC/DAC board using SMP connectors for signals and GPIO pins for programming. Figure and caption reproduced from Strader (2016).

The ROACH board has a Xilinx Virtex-6 FPGA (Field Programmable Gate Array) for channelization and pulse detection. The ADC/DAC board uses two ADCs, two DACs, and a Xilinx Virtex-7 FPGA for control. The RF/IF board is used for IQ modulation, converting the MKID probe tones to the proper frequencies to be sent into the cryostat and downconverting the outputs to values that can be monitored by the board.

Figure 2.21 shows the general outline of the readout procedure for a single readout board. The procedure goes as follows:

1. On the ADC/DAC board an optimal comb of probe tones from -1 GHz to +1 GHz (IF band) is generated and then digitized using the 2 DACs. Each probe tone corresponds to a resonant frequency for an identified resonator.
2. After being passed to the RF/IF board, the probe tones are upconverted to the RF band via IQ mixing the tones with a local oscillator (LO) at ~ 5 GHz or ~ 7 GHz,

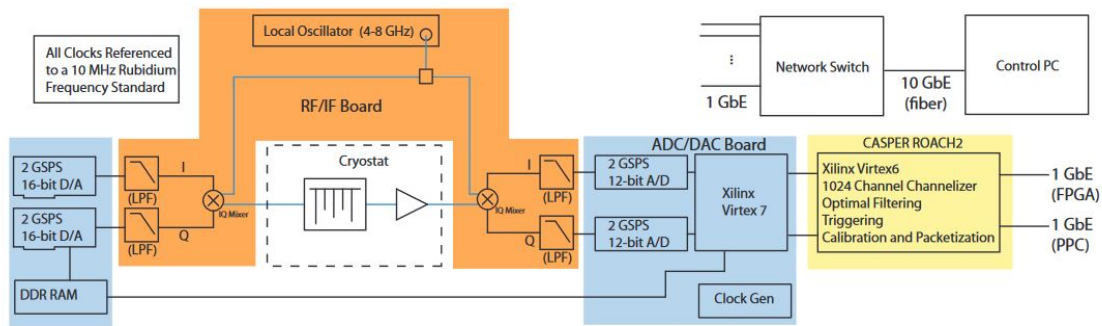


Figure 2.21 Schematic diagram of the Gen2 MKID readout boards used to read out up to 1024 pixels over a 2 GHz span. The ADC/DAC board is shown in light blue, the RF/IF board in orange, and the ROACH in yellow. Reproduced from Meeker et al. (2018).

- depending on whether the board will be reading out resonators from the MKID feedline between 4-6 GHz or 6-8 GHz.
3. The tones are sent into the cryostat, attenuated as they pass to lower temperature stages, and excite the MKID resonators at each of their resonant frequencies.
 4. After passing through the MKID, the tones are amplified at 3 K by the HEMT amplifiers and sent back out of the cryostat.
 5. The signal is sent back into the RF/IF board where they are downconverted from their RF band frequencies back to the IF band using the IQ mixer.
 6. The IF frequency tones are then digitized using the 2 ADCs.
 7. After passing through the ADCs, the data is streamed to the Virtex-6 FPGA on the ROACH board to be channelized, filtered, and to trigger on photon events.
 8. When a photon event is measured, the arrival time and pixel phase pulse height are streamed over ethernet to the network switch.

9. The network switch sends the photon events from all boards to the control PC, where they are stored in binary files.

Each photon event is sent as a 64-bit word that contains the x and y coordinate (10 bits each) of the pixel it struck, the arrival timestamp (8 bits), the pulse height of the MKID's phase when the photon struck (18 bits), and the MKID's baseline phase value (18 bits). On the control computer, each of the photon events is collected and written to disk as the raw observation files. Each raw file contains 1 second of data from all of the MKIDs from all of the readout boards that were being read out during that second. These binary files with the raw observation data will form the basis of the postprocessing efforts to perform analysis of the data collected from the MKIDs.

Additionally, as the photons are being written and stored to disk, they are being histogrammed on the time axis in parallel and those histograms are written to shared memory, which can be used to view the MKID array in real time. The real-time MKID array viewer and instrument control dashboard has historically been a GUI built in a PyQt and LabView framework, but for XKID a new web-based GUI has been developed which will be discussed further in Chapter 3.

2.8 Microwave Signal Path

The XKID cryostat requires 5 feedlines to read out the MKID array. The path starts and ends at the output and input of the readout cartridges. The major components of the microwave signal path are shown in Figures 2.22 and 2.23.

The output from the readout cartridge is sent to a 4-meter SMA cable that attaches to a hermetic SMA bulkhead in the 300 K shell to the fridge. Once in the fridge the signal travels from the 300 K stage down to 3 K through semi-rigid hand formable SMA-to-SMA coax that are heat sunk at the 50 K stage. The output side of these cables

terminate at SMA-to-SMA bulkheads in the 3 K stage to provide additional heat sinking and break up the signal path. On the opposite side of the 3 K stage SMA bulkheads there is a 20 dB and 6 dB attenuator in series, for 26 total dB of attenuation. Next the signal travels through a flexible SMA-to-SMA coax that terminates at the 3 K portion of the device stage. The signal then passes into an SMA-to-G3PO coax that connects directly to a custom flexible coaxial (FLAX) Nb47Ti cable (Section 2.8.1). The FLAX cable starts at 3 K, is heat sunk on the 1 K stage, and terminates at the 90 mK stage where it connects directly to the MKID microwave package through 5 G3PO barrels. The G3PO sockets in the MKID boxes are soldered to gold-plated copper on duroid coplanar waveguide (CPW) transition boards. The traces of the transition board are connected to the MKID array itself using Aluminum wirebonds. Finally, the signal passes into the device, exciting the resonators to be read out.

After passing through the MKID array the reverse path is followed, with the exception that at the 3K stage the signal passes through a Low Noise Factory (LNF) High Electron Mobility Transistor (HEMT) cryogenic amplifier rather than through cryogenic attenuators. The signal path is completed after it passes out through hermetic SMA bulkheads at the 300 K shell, through a 6-meter room temperature coax, and back to the input of the readout cartridge.

2.8.1 FLAX Cables

The flexible coaxial (FLAX) ribbon cables were designed and fabricated at UCSB by Jenny Smith to enable a high density of microwave wiring to pass from 3 K to 90 mK while minimizing its heat load relative to conventional NbTi cryogenic cables. Their development was published in Smith et al. (2021) and will be highlighted here.

The FLAX cables consist of 5 feedlines made of 0.076 mm (0.003”) diameter NbTi




Figure 2.22 Interior of the XKID cryostat with several pieces of the microwave signal path labeled. Not shown are the FLAX cables, the section between 300 K and 50 K, and the HEMTs.




Figure 2.23 The XKID device stage with microwave signal path pieces labelled. (Left) Shows a front view with the flexible SMA-to-G3PO connectors and heat sunk flax cables. (Right) Shows the HEMTs and the flexible cables that attach to their input and outputs.




Figure 2.24 (Top) S_{21} transmission through the FLAX cables through various traces at 4 K. The rippling effect suggests there was a slight impedance mismatch at some point in the cable which has been fixed in updated versions. (Bottom) S_{41} , the cross talk between physical nearest neighbor traces in the FLAX cable. Taken from the same cables as the top panel, also at 4 K. In both panels, the lines labeled MEC FLAX come from Smith et al. (2021) and XKID FLAX come from the follow-up publication to Smith et al. (2021), which is in preparation as of the writing of this thesis.

center conductors that are shielded with 0.28 mm (0.011”) PFA³ wire from Supercon Inc. The outer conductor that joins each of the inner conductors is 0.025 mm (0.001”) thick Nb47Ti foil. Two pieces of the foil are sandwiched around the wires, which are held in place by semicylindrical crimps in the foil. The two pieces of foil are then mechanically and electrically bonded together with spot welds that run the length of the cable. These crimps serve dual duty of holding the wires in place and giving a $\sim 50\Omega$ characteristic impedance designed to match that of the other microwave components in the signal path to prevent reflections because of impedance mismatches. At each end, the inner conductors are exposed and soldered to a transition board similar to that in the MKID microwave package (Smith et al., 2021).

The performance of the FLAX cables is shown in Figure 2.24. The transmission is

³PFA \approx Teflon

nearly flat across the measured bandpass (4-8 GHz). The cross talk is also roughly flat across the bandpass and is nearly 60 dB lower than the transmission through each of the traces, meaning there is very little signal that can transfer between traces (Smith et al., 2021). This was an issue that plagued the original flexible coax stripline cables used in MEC and DARKNESS that had cross talk at values closer to -20 to -30 dB, only 20 to 30 dB lower than the transmission in its traces (Walter et al., 2018). This led degraded signal and increased in those instruments, which is now greatly mitigated by the use of the FLAX cable design.

Finally, the FLAX cables are designed to have a lower heat load than comparable superconducting coax cables and other flexible coaxial ribbon cables like those in Walter et al. (2018). The thermal conductivity of the FLAX cable was computed by summing the thermal conductivity of each constituent material weighted by its cross-section then integrated along the length of the cable. The resulting thermal conductivity for the FLAX cables and MKID flex cables originally used in MEC are shown in Figure 2.25.

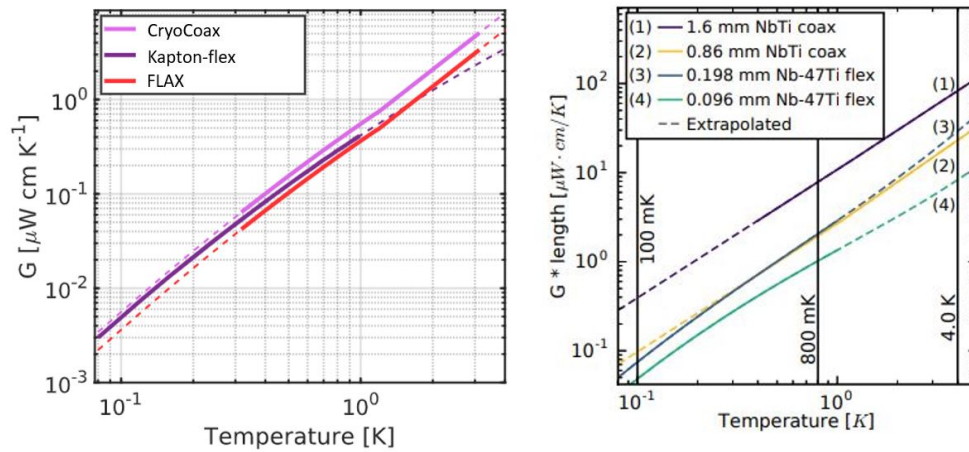


Figure 2.25 The thermal conductivity of the FLAX cables (left, from the follow-up to Smith et al. (2021), in preparation as of the writing of this thesis) and MEC/DARKNESS flex cables (right, from Walter et al. (2018)) compared to other standard superconducting coax materials. It can be seen that the flex cables had superior thermal conductivity to other materials when it was made, but the FLAX cable has since improved on that and remains the lowest thermal conductivity coax cable available for use in an MKID camera.

Chapter 3

XKID: Operation

3.1 Introduction

XKID was shipped out from UCSB on February 1, 2023 and arrived at LCO about two weeks later. Initial uncrating and reassembly of the instrument commenced on February 23, 2023. The initial cooldown and testing of XKID at LCO began on February 26, 2023. After initial tests the instrument was brought back to room temperature to reopen the cryostat and make several modifications before cooling it back down on March 1, 2023. Once cold the instrument was able to be configured and set up for observing over the next several days. The affiliated MagAO-X commissioning run began on March 3, 2023 and ran for two weeks until the 17th. XKID had its initial commissioning run as a part of the larger MagAO-X run, where it had its successful first light and was able to spend parts 5 nights gathering data on sky.

This chapter presents the XKID control software, how the software and fridge operate, and several of the calibration and on-sky results gathered during the first-light run.

3.2 Instrument Setup

3.2.1 Cooling from 300 K to 3 K

With a pulse tube backed cryostat (sometimes called a “dry” cryostat for its lack of liquid cryogens), cooling from 300 K to 3 K is a relatively simple process. First, the fridge is pumped out using a scroll pump to go from atmospheric pressure to below ~ 10 mbar. Once below 10 mbar, a turbo pump will be started to pump to $\lesssim 5 \times 10^{-2}$ mbar before starting the pulse tube helium compressor. Here it is noted that one may choose to pump the cryostat down to as low a pressure as possible prior to starting the pulse tube compressor and the fridge left under vacuum until it is ready to be cooled via the pulse tube. To start the XKID compressor, the “Power On/Off” button must be pressed manually which will begin the cooling process.

Once the pulse tube compressor has been turned on and the temperatures in the fridge begin to drop, the vacuum port can be closed and the pumps turned off. The fridge can then be left for ~ 18 - 22 hours, upon which the lowest temperature stages in the fridge should all be at its base temperature of ~ 3 K. There is one possible failure mode during this phase of the cryostat cooldown. Since there are several O-rings to seal the vacuum beneath the hermetic RF and cable ports and the external 300 K vacuum shell, a leak to atmosphere through one of these O-rings will cause the temperatures to stop falling before getting to their base temperature as air enters the cryostat. To fix this issue the pulse tube must be stopped and the fridge allowed to return to room temperature. Once at room temperature it is necessary to find the leak, which can be done using a standard leak checker or by inspecting the O-rings for obvious damage or debris, then the process must be started again, pumping out the cryostat and starting the pulse tube.

Warming to 300 K

After an observing run is complete or when it is necessary to warm the fridge back up, the “Power On/Off” button on the pulse tube compressor should be pressed to stop its operation and let temperatures rise back to room temperature. Warming to room temperature usually takes around 36 hours without external intervention.

Several options are available to speed up the warming process. One option is to soften the vacuum and allow exchange gas from outside the fridge into the cryostat. The other method that has proved very useful is to wrap the cryostat in a heated blanket or heating pad and turn it on to as high a temperature as possible. The heated blanket method has been shown to speed up the warm-up time to ~ 24 hours.

3.2.2 Cooling from 3 K to 90 mK

Cooling from 3 K to 90 mK requires the use of the ADR unit, the superconducting magnet power supply, and the mechanical heat switch. The ADR cycle is broken into 4 parts: ramp up, soak, ramp down, and regulation. The heat switch must be closed prior to cycling, keeping the 1 K and device stages in thermal contact with the 3 K stage.

Cycling the ADR

To ramp up a command is sent to the Lake Shore 625 in the electronics rack specifying the desired current ramp rate in Amps/second. The nominal value for ramping the current in the ADR is 0.005 A/s (5 mA/s), with values up to 50 mA/s being used for fast ramps when required. The ramp rate should never exceed a value that would generate a back EMF above 250 mV across the leads of the superconducting magnet, which may cause hardware damage. A second command is then sent to the Lake Shore 625 specifying the desired current that the magnet will “soak” at. This current corresponds

to a magnetic field in the bore of the solenoid-shaped superconducting magnet in the ADR. The magnet's field constant is 4.0609 kG/A and the maximum field allowed is ~ 38.33 kG, which corresponds to a maximum allowable soak current of 9.44 A. Trying to push more current through the magnet may cause irreparable damage. The default soak current for the XKID ADR is 9.25 A. After the second command is set, the Lake Shore 625 will ramp the current through the magnet at the constant rate that was originally specified. The ramp rate may be changed during the ramp up process, but should only be changed in small increments.

At any point if there is current in the magnet there is the risk of a "quench". A quench occurs when part of the superconducting magnet current carrying wires goes normal, or is no longer superconducting. When this happens the leads go from having 0 DC resistance to some non-zero resistance. This will cause the magnet to heat up since a current I traveling through material with a resistance R will dissipate power as $P = I^2R$. This will cause the rest of the superconducting magnet to heat up, making it *all* go normal and causing the current through the magnet to drop to 0 A. The Lake Shore 625 has built in quench detection, so it will quickly register a quench and stop trying to push any current through the magnet. A quench will cause the 3 K stage to quickly heat up by an amount proportional to the amount of current that was in the magnet when the quench occurred. After this happens the operator should wait for the excess heat to dissipate and for all temperatures to return to their nominal ~ 3 K value.

The ramp up process is exothermic, meaning that each low temperature stage will slightly increase in temperature as current is increased in the magnet.

The soak stage begins when the current reaches highest value, the soak current. During this process the magnetic spins in the GGG (gadolinium gallium garnet) and FAA (ferric ammonium alum) salt pills are allowed to maximally align, decreasing their overall entropy, which will allow them to draw heat from the 1 K and device stages,

respectively. During this stage there are no commands that are required to be sent to any of the control hardware or software. If the user desires a faster or slower ramp down speed than the current was originally ramped at, it may be modified here. Otherwise, once the soak current is reached a timer is set and will signal to begin the ramp down when the duration has elapsed. By the end of the soak stage the temperatures of the low temperature stages should have returned to ~ 3 K, with the excess heat from the ramp out having been pulled out via the pulse tube.

The minimum recommended soak time is 30 minutes, while the default setting for the XKID ADR cycle is 60 minutes. Between soak times 30 minutes and 4 hours, a longer soak will result in a longer hold time at 90 mK.

After the soak duration has elapsed the fridge must be prepared to ramp the current down to nearly 0 A, allowing the temperature of the salt pills in the magnet to drop. The first step before ramping down is to open the heat switch. An “Open” command is sent to the Zaber T-NM17 stepper motor. As it begins opening the heat switch, the 1 K and device stages will come out of thermal contact with the rest of the cryostat. Once the heat switch is sufficiently opened a command is sent to the Lake Shore 625 specifying the current ramp down rate in A/s. Typically this is the same as the rate the current was ramped up at, but that is not required. The default value is 5 mA/s for XKID. Finally, once the ramp down rate has been set a command is sent to the Lake Shore 625 to set the current to 0 A. It will then decrease the current through the magnet at the current rate specified by the user. During this stage there may be a slight spike in temperature when the heat switch initially opens before the decreasing current starts to drop in temperature of both stages.

Before the current drops to 0 A, the current will typically have dropped below the default 90 mK regulation setpoint. The Lake Shore 625 is required to be configured to be controlled using its “Sum” feature, meaning it can be either manually controlled by a

user or externally controlled by a PID controller. Using the sum feature allows a seamless transition between manual control of the magnet to PID control using the Lake Shore 372 AC Resistance Bridge that measures the temperature of the device stage.

Regulation

The final stage of an ADR cooldown is regulation. This occurs once the temperature of the device stage has dropped low enough that the Lake Shore 372 can apply a correcting voltage to send to the Lake Shore 625 which will generate a small current to heat up the device stage to the desired 90 mK value. The sole command that must be sent during this stage is to the Lake Shore 372 to ensure that the PID output is turned on. However, there are several parameters that the user may tune to improve the stability of the PID loop. First are the P, I, and D parameters themselves. The PID loop operates via a feedback loop where the controlling voltage the Lake Shore 372 output is

$$\text{OutputVoltage} = P \left[\epsilon + \frac{1}{I} \int \epsilon dt + D \frac{d\epsilon}{dt} \right] \quad (3.1)$$

where $\epsilon = \text{Measured Temperature} - \text{Setpoint Temperature}$. Adjusting P, I, and D is how the user can get the fridge to control the temperature to within ± 0.030 mK.

Additionally the user may control the heater output range, which is the amount of current allowed to be output by the Lake Shore 372 when it applies the control voltage to the Lake Shore 625. If the value is too low the PID loop will never bring the temperature up to the setpoint and if it is too high the PID loop will not be able to control its output voltage with fine enough resolution, leading to a noisier feedback loop and worse control. Each of these parameters may be tuned manually by the user while regulation is ongoing. Once the PID loop is tuned, it is typically stable and does not require further tuning unless there are major hardware changes such as a new thermometer being used.

A well tuned PID loop – in addition to an appropriately long soak time – will allow the device stage of the fridge to stay at the MKID operational temperature for 8-24+ hours, depending on how well heat sunk the hardware going to the device stage is.

Finally, the ADR cooldown cycle can end in one of two ways. First, if the current in the magnet runs out and it can no longer regulate the temperature then the user must recycle the magnet to begin the cycle again. In the second case, if the user has an issue with the cooldown and wishes to begin again they may send an “abort cooldown” command, which will immediately close the heat switch and bring the current in the magnet to 0 A in preparation for a new cooldown when desired.

3.2.3 Finding the MKID Resonators

The process of identifying each resonator’s resonant frequency and ideal excitation power as described in detail in Fruitwala et al. (2021) and Fruitwala (2021).

The standard procedure treats resonant frequency identification and drive power tuning independently. To make the required measurements, each step uses the complex frequency response $S_{21}(f)=I(f)+iQ(f)$ of the MKID array across the instrument’s bandpass of $\sim 4-8$ GHz. When S_{21} is plotted in I vs Q space (the complex plane) a resonator will manifest as a “loop”. When the magnitude $|S_{21}|$ (i.e. transmission) is plotted as a function of frequency each resonator will appear as a sharp dip. Figure 3.1 shows a single resonator’s response both in the complex plane and the magnitude. The response data is taken at many powers over all frequencies to properly ID the resonators.

Frequency Identification

First, each resonator’s resonant frequency must be read out. To do this a comb of 2048 evenly spaced probe tones at a single power will be generated over the 4-8 GHz




Figure 3.1 Reproduced from Fruitwala et al. (2021). A single resonator’s frequency response over a ≈ 200 kHz band. (Left) The magnitude $|S_{21}|$ vs frequency shows a transmission dip at the resonator’s resonant frequency. (Right) In I vs. Q space, a sweep in frequency will cause the resonator to appear as a “loop” when appropriately powered.

band. The comb is then be swept ± 2 MHz from its original frequency values in 7.7 kHz steps. This creates nearly complete coverage of the frequency response over the feedline.

The magnitude $|S_{21}|$ is then calculated for the full frequency range and a standard peak-finding algorithm such as SciPy’s `signal.find_peaks()` function¹ is combined with a spatial bandpass filter to identify where each resonator falls in frequency space.

Power Tuning

To maximize the SNR when detecting a photon it is necessary to drive each resonator at the highest possible power. If a resonator is *too* highly powered it will bifurcate and become unusable. To find the ideal power for each resonator, the frequency sweep described is performed at 31 different powers (a 30 dB range in steps of 1 dB, ends inclusive), with the highest power being at a point above where all of the resonators bifurcate. Figure 3.2 shows a single resonator’s IQ loop at 4 different powers going to

¹https://docs.scipy.org/doc/scipy/reference/generated/scipy.signal.find_peaks.html




Figure 3.2 Reproduced from Fruitwala et al. (2021). A resonator is shown in IQ space at 4 different powers. At the lowest power (red) the resonator is a smooth loop and at the highest power (blue) there is a discontinuity in the loop where it becomes bifurcated.

a sufficiently high power and beginning to bifurcate. The data from the 31 sweeps at different powers is then fed into a neural net classifier which estimates the appropriate power (Dodkins et al., 2018; Fruitwala et al., 2021) for each individual resonator.

After this tuning process lists of resonators are stored that contain each one’s optimal frequency and readout power which will be loaded into the Gen2 readout to read out each resonator. Once loaded in, a user can manually inspect the resonators and retune them individually to improve performance. Each list is comprised of rows that contain a 5- or 6-digit resonator ID number, a metadata flag for ingestion in data reduction pipeline steps, the resonant frequency of the pixel, and its ideal readout power.

3.2.4 Beammapping

Once the frequencies and readout powers of each resonator has been found it is necessary to figure out where each resonator is physically located on the array. This involves

a process called “beammapping” using the Samsung Galaxy S6 phone in the calibration snout (Section 2.6.2). The process and design for reimaging the screen of the phone on the MKID array is described further in Bottom et al. (2018). The flipper mirror for the calibration snout will first be moved to its “down” position and the filter wheel will be set to its “Open” position so light from the calibration optics reaches the detector.

The process for beammapping the array is fairly straightforward. With the custom calibration app, a bar of light ≈ 1 pixel wide is swept across the screen in the “x” and then the “y” direction. By measuring when each pixel sees the bar of light its relative location in the array can be determined. To get a more accurate measurement, several sweeps in each direction can be taken to provide a stronger signal from each pixel.

After this data is taken, it is processed through reduction code in the `mkidreadout` repository² which generates a list that has the ID of each resonator along with its (x, y) position on the array and a metadata flag used by the readout boards to determine if the pixel should be read out or if it is “dead”/unable to be calibrated.

3.2.5 Optimal Filtering

The last setup step required is to generate optimal filters that improve the SNR of the real-time photon triggering in the readout. To do this the MKID array is flooded with monochromatic light (usually from the 808 nm laser in the laser box) and the phase timestream for each resonator is sampled at 1 MHz for 5 seconds. This is done to measure a sufficient number of photon events to average them into a template and measure the phase noise in the resonator so that a Weiner optimal filter can be calculated uniquely for each resonator. After being calculated, the optimal filter coefficients are stored for later use when they are loaded into the readout firmware during array initialization.

²<https://github.com/MazinLab/MKIDReadout>

3.3 Instrument Control Software

The instrument control software suite is designed to be a distributed, modular system that is forward compatible with for new instruments, backwards compatible to replace LabView-based control software used in previous instruments, and relatively easy to customize or modify. It is an open source repository (called “mkidcontrol”) that can be found at <https://github.com/MazinLab/mkidcontrol>. It is written in Python, C++, C (and Cython), HTML, and JavaScript and comes in an easy-to-install package. The package includes software for controlling the instrument and reading it out. In previous instruments these tasks have been decoupled but the new architecture combines the two to reduce complexity and enable more seamless integration which in turn reduces the overhead required to set up the control system to read out the MKID array. It was developed and tested on Ubuntu LTS releases and has not been validated on other operating systems. The GUI was developed using Flask³, a web framework written in Python that allows users to create custom web apps that may be hosted locally or more broadly on the internet. The control system is detailed in the following section.

3.3.1 Devices

In the parlance of the mkidcontrol software, a “device” is a piece of hardware that requires control through the software. For XKID the devices are as follows:

- Lake Shore 336 Temperature Controller
- Lake Shore 372 AC Resistance Bridge
- Lake Shore 625 Superconducting Magnet Power Supply
- CONEX-AG-M100D Piezo Motor Mount

³<https://flask.palletsprojects.com/en/2.3.x/>

- Finger Lakes Instrumentation CFW2-7 Filter Wheel
- Thorlabs TDC001 + MTS50/M-Z8 Positioning Stage
- Zaber T-NM17 Heat Switch Motor
- ArduinoUno (for controlling calibration lasers and flip mirror for calibration snout)

Each device is controlled via a python class named after the device (e.g. the Lake Shore 372 is backed by a python class called `LakeShore372`). Each device class is responsible for the serial communication with the physical hardware it is responsible for controlling and for exposing its useful functionality. For example, the `LakeShore372` class can query temperatures, configure thermometry input channels, configure its PID settings, etc. Since many of the hardware devices share serial communication protocols and other functionality, they may be subclasses of more generalized device classes. For example, each Lake Shore device utilizes the same command syntax and serial communication structure, so there is a generalized `LakeShoreDevice` class that handles the more basic functions that are shared between devices. This structure reduces clutter, verbosity, and repetition in the code base.

3.3.2 Agents

To control devices or complex processes the software utilizes “agents”. Each agent is a python program tasked with instantiating and connecting to a device, monitoring it, gracefully handling any errors during operation, and processing any commands sent by the user or internally to modify the device’s settings or state.

In addition to controlling specific devices, there are agents responsible for controlling processes. These processes include starting/stopping data collection from the Gen2 readout, operating the ADR cycle, and controlling the GUI (Graphical User Interface).

The naming convention for an agent is `<controlled unit>Agent.py`. For example, the agent responsible for controlling the Lake Shore 625 is called `LakeShore625Agent.py`

Each agent can be run in one of two ways. The first is to run it as a regular python program via the command line. This is a useful method for testing individual devices and monitoring them in an isolated manner. However, it is not recommended for normal use. In normal use, the agents are started, stopped, and restarted using Linux systemd unit configuration files, here called “service files”.

3.3.3 Service Files

Systemd is a software suite used by Linux operating systems as a system and service manager. Its main goal is as an `init` system that can bootstrap user space and manage user processes, which is to say that it is used to generate the user’s custom environment in a secular fashion that is separate from the underlying kernel.

In the XKID software architecture the reason for using systemd service files is to create daemons that run each agent program (and any other required programs) in the background without requiring user intervention to start and stop them or to restart them in the case of a fatal error. The systemd service files can be enabled to run at system startup, after another program it is dependant on has run, or following other system events. They may also be configured to attempt to restart the program they control in the event of its failure. In the XKID software the unit files are contained in the repository, copied to the appropriate system directory upon installation, and “enabled”, meaning the OS will run them automatically per the rules specified in each file.

A user may also assume manual control of the unit files via the command line. The generic command is `sudo systemctl <command> <servicename>.service`. Several common commands are `start`, `stop`, `restart`, and `status`. The `<servicename>`

for each of the unit files will match the name of the agent it is responsible for running to prevent ambiguity where possible. For example, the service file responsible for controlling the `LakeShore625Agent.py` program is called `lakeshore625.service`.

3.3.4 Redis

The connective tissue and foundation of the XKID control architecture is redis⁴ (the Remote Dictionary Server). It is an in-memory data store that can be used as a database, cache, streaming engine, and message broker. Redis is open-source and written in C, but has robust Python wrappers that enable effective use in the control software. The control software suite primarily uses redis’s fast storage capability as a database for storing telemetry and instrument status information and its message brokering “pubsub” (publish and subscribe) functionality to communicate between programs. It is installed automatically when the `install.sh` script in the `mkidcontrol` repository is run, which also contains links and notes regarding its installation and configuration.

Data Storage

In XKID the data storage takes two forms, time series data (temperature, resistance, current measurements, etc.) or non-time series data (statuses, instrument states, configuration parameters, etc.). Redis is capable of storing complex data structures such as dictionaries but in the control software is restricted to storing “flat” data types such as `strings`, `floats`, and `ints`. If a `dict` is required, it will be stored as a `JSON string` that can be converted back to a `dict` when extracted from the database for use.

For storage in the redis database, a “key” and a “value” are required. The key is the identifier that will be used to store or retrieve the value in question. The key will always be a `string` and the value may take the form of any of the types from the previous

⁴<https://redis.io>

paragraph. When storing a non-time series value the previous value stored beneath that key will be overwritten. In contrast, when a time series key is stored the value will be stored along with UNIX and local time zone timestamps in a tuple and appended to a list of all previous values that were stored for the time series key. Since the time series key is meant for measurement data, the only values it accepts are `ints` or `floats`.

The schema for storing instrument data is found in the `configuration` directory of the `mkidcontrol` repository in the file called `mkidcontrol.schema.yml`. This schema creates a set of rules to generate keys in a systematic, readable way that allows the user to quickly identify what piece of information they are storing/reading.

When retrieving data from the redis database, one solely needs the name of the key they want information from. If the user queries a non-time series key, then the redis client will return the most recently stored value. If the query is for a time series key, the user has the options for more granular control of the returned value. For a time series key the user can ask for the most recently stored value (by default), the entire time series, or a portion of the time series (by giving start and stop UNIX timestamp stamps).

The database itself is stored to a file on disk which can be copied, moved, or wiped by the user. It is also configurable and the user may set it to automatically wipe itself or create a compressed copy once it has hit a specific size or on a certain cadence (e.g. once per month). Finally, each store action to the redis database is atomic and the database supports millions of store/read operations per second.

Pubsub

While the fast data storage and retrieval is effective for letting many programs manipulate data simultaneously, it does not offer a straightforward way to let separate programs communicate with one another. However, redis has another utility called Pubsub, short for publish and subscribe, that allows it to act as a message broker between programs.

The first element of the Pubsub system is publishing. A program does not require extra setup to publish a message. If redis is set up in a program (which it is set up in all `mkidcontrol` programs), it simply has to call the redis `publish` function with a “channel” and “value”. These are analogous to the key/value pairs from the data storage section. The naming rule for publishing a command from a program is to append the prefix `command` to an existing key from the data storage namespace. The allowable command channel names are also contained in the `mkidcontrol_schema.yml` file.

On the opposite end, programs must subscribe to specific channels to hear commands and messages from them. This is achieved by using the redis `subscribe` function, which takes a list of the desired channels to listen to. Typically a program listens for channels that trigger a state or configuration change, although they may listen for any channel the user desires. Once a program has subscribed to the required channels, it will register each time that a message is published to that channel. This is the foundation for the communication between programs and also allows the user to control the different components in the system from a high level by sending commands from the control GUI.

One drawback of the redis Pubsub architecture is that there is no built in handshaking, so if a program sends a message it does not automatically know if it was received by the desired listener. However, the pubsub architecture does let the program know *how many* programs heard the message. With this in mind, the system has been designed such that it knows how many programs should be listening to a given channel and – after sending a message on a channel – checks to see whether or not the correct number of programs heard it. If too few heard the message then it will report a “failure” to the user, who may choose what course of action to take to rectify the situation, such as restarting the listening program or simply re-sending the message.

3.3.5 Running the XKID Control Software

The previous sections have discussed the architecture and design of the MKID instrument control software system but do not provide specific instructions to download and run the control software. In the `mkidcontrol` repository there is a `README`, an `install.sh` script that is highly detailed and will allow the user to install all of the requisite packages and create a python environment capable of running the XKID instrument, and several files containing notes on the system architecture and how to run the code.

To begin, the user should be working from a clean install of a Debian Linux distribution using Linux kernel 5.x.x+. The most common Debian-based distribution is Ubuntu, which is the one that was used to develop and test the XKID software. It requires a version with a kernel of 5 or greater due to several modifications that must be made to low-level drivers that are incompatible with earlier kernels. Once the computer OS is configured and running, the user will clone the `mkidcontrol` repository from `github.com`. As of the writing of this thesis it is maintained at <https://github.com/MazinLab/mkidcontrol.git> on the `develop` branch.

After cloning the `mkidcontrol` repository, the user should navigate into the repository's top-level directory and run the `install.sh` script. While this will work from a fresh OS install, it is recommended that the user also follow along with each step in the installation script to understand the setup via the notes for each step in the process. Following the completion of the installation script, the next step will be to check that a Python environment has been properly set up to run the `mkidcontrol` software. This can be done by running the command `conda activate control`. If successful this means that the `anaconda` package management system has been properly installed to the computer and the `control` environment has been appropriately created.

Since the `mkidcontrol` environment is highly customized and under continuous devel-

opment it is necessary to perform a second step for configuring the environment. While in the `control` environment, the user must also navigate to the top-level `mkidcore` directory and run the command `pip install -e .`, which is an editable install of the custom package that will compile the C code required to read out the MKID array.

After running the installation script and the editable install, the control software will be set up and ready to use. The service files will have been enabled and started by the installation script, meaning that all appropriate programs should already be up and running. To ensure that this is the case, the user can run the command `sudo systemctl restart mkidcontrol`. This command will restart the `mkidcontrol` application program which is dependent on getting all of the other programs up and running. If a program had not been started, this command will attempt to start it before starting the application itself. Finally, the user can open the control GUI and access the functionality that comes with it. Using the instrument control GUI will be discussed in Section 3.4.

Enabling Photon Capture

The final step before reading out the MKID array is starting two programs, one called `packetmaster` and the other called `send_photons`. These two programs are not started along with the rest because they require manual intervention from the user.

Once the control GUI is up and running, the user should navigate to the “Redis Commander” page in the “Utilities” menu. The Redis Commander is a web interface for the redis database that allows the user to easily modify different values. There the user should modify the value for the key named “gen2:dashboard-yaml” to be the fully qualified path to the readout configuration file called “dashboard.yaml” for that day’s observations. This file contains relevant storage paths, IP addresses, and other configuration info for the readout boards to use. Typically it will live in a directory whose format is `/data/XKID/ut<current date>/config/dashboard.yaml`.

Once the configuration file has been appropriately updated, copied into the desired directory, the value for the path updated in the redis database, and the readout boards initialized (see Section 3.4) the user can start the final two programs. To do this the command `sudo systemctl restart packetmaster send_photons mkidcontrol` is run.

3.4 Instrument Control GUI

There are three GUIs used to set up and configure the instrument control. They are:

- `initgui`: Responsible for initializing the readout boards
- `hightemplar`: Responsible for configuring the readout boards
- `MKID Control GUI`: The newly developed Python Flask application to control and read out the XKID instrument from a web-based app.

The first two are part of the 2nd Generation MKID readout and have been used extensively with previous MKID instruments (Meeker et al., 2018; Fruitwala et al., 2018; Walter et al., 2020). The last was newly developed to integrate real-time array viewing and instrument control, which replaces the old Python-based Qt `dashboard` program previously used as a real-time array viewer and LabView based instrument control GUI.

3.4.1 `Initgui`

Whenever the readout boards are powered up they must be initialized using the `initgui.py` program (Figure 3.3, left). This program will first connect to each board and verify it has booted and can talk to the control computer over the local private network in the electronics rack. The next step uploads the firmware to and then programs the Vertex V6 FPGA on the ROACH2. After this, the Vertex V7 FPGA is initialized by enabling

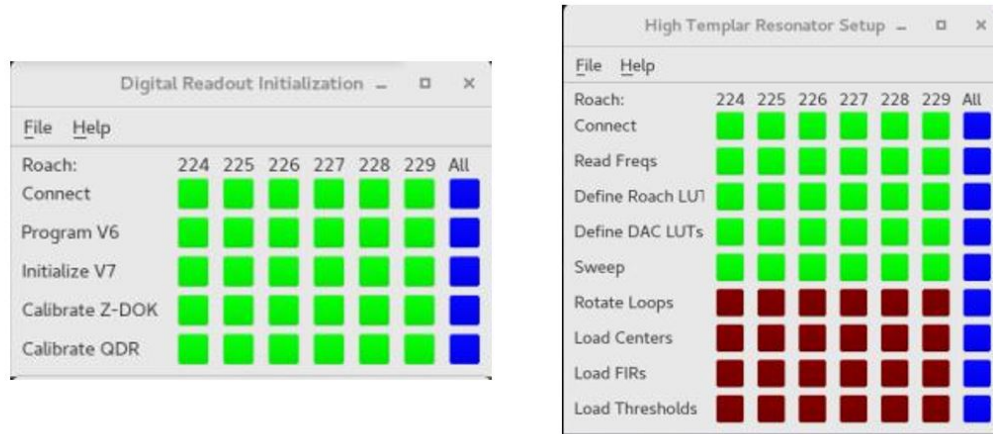


Figure 3.3 Screenshots showing the `initgui` (left) and `hightemplar` GUI (right). Each column represents a different readout board, with the 3 numbers at the top corresponding to the final 3 numbers of the board’s IP address. Green squares indicate a process that has been completed, red squares have not yet been run, and blue squares are used to run a step for all of the boards at once.

UART communication on the ADC/DAC board before configuring the ADC, DAC, LO chip, and on-board attenuators. Finally, the QDR on the ROACH2 is formatted so a DDS LUT (Direct Digital Synthesis Lookup Table) can be loaded later.

Once the steps are completed for each board, the GUI may be closed. However, if a board experiences issues or it seems there is too much noise in some of the pixels on a board it is usually helpful to power cycle the board and re-initialize it.

3.4.2 Hightemplar

The next step for setting up the readout is to run the `hightemplar` program. In contrast to `initgui`, which must only be run when the readout boards are powered, this program should be run each time the cryostat is cooled from 3 K to 90 mK. The control GUI for this program is also shown in Figure 3.3 on the right.

The first step is for the control computer to connect to the boards and verify the firmware was properly initialized. Next, the pixel frequencies and drive powers are read

in from text files generated during the setup.

The third step is to define a Roach LUT. To do this, an FFT (Fast Fourier Transform) bin is calculated for each resonator's probe tone, which is then uploaded into the firmware. Next, the DDS tone for each resonator can be calculated and loaded to the DDS LUT in the QDR. This enables digital downconversion from RF to IF frequencies.

Next, the "define DAC LUT" step is run. In this step the ideal Local Oscillator (LO) frequency is calculated for each board and loaded into the LO chip on the RF/IF board. Then a frequency comb is generated and loaded onto the ADC/DAC board so it can be used to probe all of the resonators for that board. Next, the DAC attenuators on the board are set to maximize the DAC's dynamic range while each probe tone is kept close to the ideal power. Finally, the ADC attenuators are optimized so that the output signal from the resonators to the board also uses the maximal dynamic range without overflow. If the ADC attenuation is too low and the dynamic range is exceeded it will cause significant noise in the resonators.

After both LUTs are loaded, the resonators may finally be probed. The following step is to sweep the resonators, which involves stepping the LO tone in 10 kHz increments over a 500 kHz band. This lets the board measure each pixel's IQ loop (Figure 3.1, right). If there is significant noise on the resonators here, the user can either re-initialize the board or manually change the ADC attenuation (for all resonators on the board) or DAC attenuation (for that resonator specifically) to mitigate the noise.

Once the pixels have been swept, the phase of the each pixel's probe tone is measured and its negative value is applied to the tones from the DDS LUT to set the resonator's IQ loop to have a baseline phase of 0 radians. Following the loop rotation, the center of each pixel's IQ loop is measured and loaded into the firmware so the phase may be easily calculated in the photon-triggering block.

The next step is to load the Finite Impulse Response (FIR) filters, which consists

of reading the optimal filter coefficients (Section 3.2) for each pixel into the firmware. The optimal filters significantly smooth out the phase response of each resonator which in turn will increase the SNR of any photon detections. Finally, thresholds are loaded into the photon-triggering block in the firmware. This is done by calculating the phase for each pixel and measuring the noise on the phase. A 5σ threshold is calculated for each pixel individually which is then loaded into the firmware. This threshold sets the minimum phase response that can trigger a photon event.

The `hightemplar` GUI is typically left up and running throughout an observing night so a board can be re-swept and re-thresholded if for some reason it becomes noisy but does not require a restart. After `hightemplar` has been run, the commands to start the `packetmaster` and `send_photons` programs can be started using the command from Section 3.3. These may NOT run before `hightemplar` has finished, as they assume the readout boards have been set up, configured, and are ready to send photons.

3.4.3 MKID Control

The MKID Control GUI does not need to be started after the `hightemplar` steps have been run, and in fact it often runs concurrently since it enables real-time MKID array visualization, instrument monitoring, and control. Prior to running `packetmaster` and `send_photons` the array viewer will be completely blank and will not show anything until those programs are manually started by the user. To check that the control GUI is up and running, the command `sudo systemctl status mkidcontrol` can be run. If its status is “active” then the GUI is ready to be accessed and used.

The webpage for the GUI can be accessed via the machine’s external IP address. At LCO the name for the IP is `xkid.lco.cl` although the numerical label for this address is statically assigned and will change each time XKID connects to their local network.

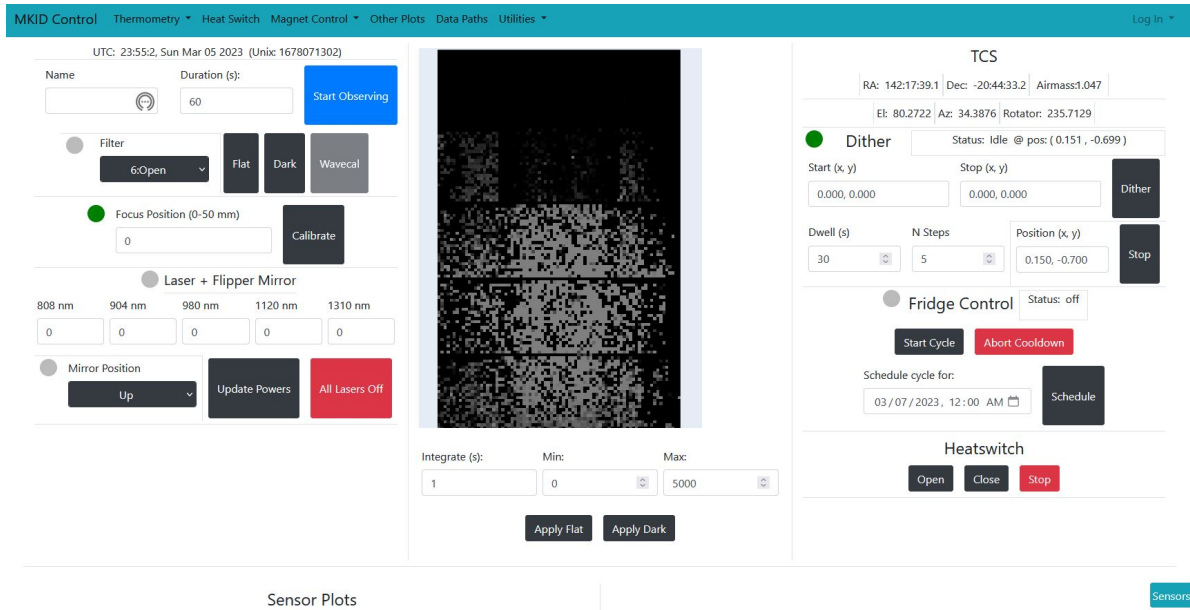


Figure 3.4 A screenshot of the control GUI homepage taken at the moment of first light (2023/03/06 03:06 UTC). The central panel is a live array view (prior to finding the focus). The left and right panels have real time information and control widgets.

At UCSB, the address is `xkid.physics.ucsb.edu` and the numerical label is dynamically assigned using DHCP. If creating an SSH tunnel or using RDP is undesirable, the GUI can also be accessed by going to the web address `http://xkid.lco.cl:8000` or `http://xkid.physics.ucsb.edu:8000`, depending on what network the instrument control computer is on. The home page that the user will land on is shown in Figure 3.4.

The MKID control software was developed as a robust, modifiable, Python-based alternative to previous MKID instrument control GUIs. It was successfully used to control XKID during its first-light run at the Magellan Clay telescope in northern Spring 2023. There are also future plans to migrate existing MKID instruments such as MEC at Subaru to use this software suite as a replacement for its current LabView-based control software.

Chapter 4

XKID: Commissioning

4.1 First Light Run

The first light night for XKID was on 5 March 2023 at the Magellan Clay Telescope at LCO behind MagAO-X. The observing run ran until 10 March 2023, after which XKID was turned off and removed from the telescope. Figures 3.4 and 4.1 each show screenshots 1-second exposures from the live array viewer during the first light night. During the run XKID was used for 5 half-nights and left inactive for 1 night in which an internal MagAO-X instrument was allocated time. While the instrument had down time on that night, it was still used to take dark and flat observations for future calibrations.

This chapter will report several of the most important instrument calibration steps and some of the early observational results.

4.2 Calibration

Several of the key calibration results will be discussed here. These are the wavelength calibrations (or “wavecal”), the conax-to-pixel mapping, which enables proper mapping




Figure 4.1 Screenshots showing 1-second exposures in the live array viewer from XKID during its first light run. (Left) A single star at all observable wavelengths (~ 800 - 1400 nm) while the focus was being found. (Right) A single star in J -band with Lyot coronagraph in and focus determined. The bright spots above and below are “sparkles”, dimmer copies of the stellar PSF intentionally placed by the DM for calibration. The horizontal gaps in each image are an artifact of the beammap that was used by the live viewer during the observations. When reducing the data, a subsequent beammap with the gaps fixed can be used.

of dither observation sequences for stacking+adding and mapping to sky coordinates, and measuring the plate scale and device orientation on sky.

4.2.1 Wavelength Calibration

During each observing night it is required that a wavelength calibration (see Section 1.3.4 for more details) is performed. In this process each of the 5 lasers (at 808, 904, 980, 1120, and 1310 nm) is subsequently shined on the array, flood illuminating all of the pixels. The wavecal datasets that were taken each night have been reduced and currently report that the median resolving power across the 808-1310 nm bandpass is $\mathcal{R} \approx 3.5$ and remains fairly flat as a function of wavelength.

The resolving power of 3.5 is lower than that of comparable instruments such as MEC and DARKNESS which each measured resolving powers closer to $\mathcal{R} \sim 5-7$ (Meeker et al., 2018; Walter et al., 2020). It is believed that the IR background is creating a significant number of counts in each pixel and causing degradation of the measured resolving power. Removing and calibrating the IR background to analyze how it degrades the detector resolving power is a problem undergoing investigation as of the writing of this thesis.

4.2.2 Conex to Pixel Mapping

The next calibration step is to measure how the change in conex tip-tilt mirror position changes the PSF's location on the MKID array. This enables reconstruction of dithered datasets by appropriately shifting each frame to the proper location so the successive frames can be added in the correct place. An incorrect mapping will lead to blurring of features in the image, hindering the SNR of any potential astronomy detections and make identification of real sources more challenging, if not impossible.

The two requirements for a dataset to generate a conex-to-pixel mapping are:

1. The dataset must be a dithered sequence.
2. There must be “sparkles” (or “satellite speckles”) generated by the DM.

The observation must be dithered because the goal is to discover what a change in conex position looks like as a function of pixel position. In essence the aim is to measure $dpixel_x/dconex_x$ (dp/dcx) and $dpixel_y/dconex_y$ (dp/dcy).

There must be sparkles because the code will attempt to fit a model PSF including sparkles to the data. This is because (1) most observations with sparkles also use a coronagraph, covering the central star and (2) the sparkles are designed to be astrometric calibration tools. If one draws a cross connecting sparkles that are opposite one another, the intersection point provides an accurate measurement of the target center.

The process for generating a measurement of dp/dcx and dp/dcy is as follows. The user will provide a set of initial guesses for several parameters that describe the shape of the coronagraphic PSF. These include the sparkle separation, the sparkle and background fluxes, the rotation of the cross made by the sparkles, and an initial guess for the central position of the PSF itself. These are used to generate an initial guess for a model of the PSF which is fit to the data from a single dither frame. Using a Markov Chain Monte Carlo (MCMC) fitter, the model is modified until the χ^2 metric describing the goodness of fit between model and data is minimized. The best-fit model will then provide an accurate estimate of the (x, y) pixel position of the center of the coronagraphic PSF while the metadata from the dither will give the conex position.

After having fit a model to the first frame, the conex to pixel code automatically will repeats this process for all of the subsequent frames in the dither sequence. Once a model PSF has been fit for each dither step, the values for dp/dcx and dp/dcy can be calculated. This is done by (in each direction) performing a least squares regression fit to find the slope of the PSF central coordinate vs conex position data. The slope in each




Figure 4.2 A 30-second XKID image of the HD 96220 binary system used to calibrate the platescale and device angle for XKID. (Left) An uncalibrated image as seen by the array with each element of the binary circled in white. (Right) The calibrated image, properly rotated on sky and with an accurate scale bar.

direction is therefore the measured dp/dcx or dp/dcy value, which can then be used in the data reduction steps to properly manipulate dithered data.

4.2.3 Plate Scale and Device Rotation

The final required calibration is to determine the plate scale and device rotation. The plate scale is a measure the angular size of a pixel on sky and is usually reported in milliarcseconds per pixel (mas/pix). The device rotation is how far away the vertical axis of the MKID array is from “true North” on the sky.

To measure both quantities, a binary system with a known separation and position angle (PA) was measured with XKID. The known separation enables the calculation of the plate scale and the PA – the orientation of the binary system measured East from North – will help generate the device rotation. The binary system that was used to measure these quantities, HD 96220, has a separation of $0''.60$ and a PA of 132° .

The central pixel of each member of the binary must first be found to help measure each quantity. This is achieved using a Python source-fitting function based on the `photutils` package's photometry tools. The results of this fit are shown without any rotation or scaling (i.e. as the array saw it) in the left panel of Figure 4.2.

The plate scale is calculated by dividing the known separation of the binaries by the pixel separation. From the source fitting routine it was found that the center-to-center distance of the PSF of each star is ~ 26.37 pixels. Combining this with the known separation of $0''.60$ it is determined that XKID has a plate scale of 21.15 mas/pixel.

The process for finding the device angle is slightly more involved. The goal is to measure how far the vertical axis of the array must be rotated so when an image is generated and rotated by that amount it will correspond to how the object is actually oriented on sky and the vertical axis of the image will correspond to on sky North. The first step to determine the device angle is to find the known PA of the system, which says that the dimmer component of the binary is 132° East of North. Second, the image the array sees is used to find the angle between a the vertical axis and a line connecting the binary stars, measured here to be 175.31° . The difference between these represents the device angle and is $\approx 40.31^\circ$, or the global offset that must be applied when generating images so that the image orientation the image matches what is expected on sky.

The right panel of figure 4.2 shows the same image as in the left panel now calibrated to reflect real on sky parameters. Finding the PA and plate scale are the final calibration steps required before the data reduction pipeline can produce images that are true to reality so an end user can use the output products from the instrument to search for exoplanets in the appropriate regions of those images.

4.3 Data Reduction

The data from the XKID observations were reduced using the MKID Science Data Pipeline (Steiger et al., 2022a). Ingesting and processing XKID data was possible with only minimal modifications to the existing MKID pipeline software package. The data reduction steps are enumerated and briefly described below.

1. **Bin2HDF:** The raw output data from MKID instruments are stored as binary (.bin) files that each contain 1 second of data from the camera. The first step of the pipeline reads the photons from these .bin files into Hierarchical Data Format (HDF5, or .h5) files. The .h5 files are structured tables that contain information from all of the photons measured over its duration and observation metadata. Each photon in the .h5 file is a line consisting of the resonator the photon struck, the timestamp when it arrived, a “weight”, and the peak of the phase pulse.
2. **Metadata Attachment:** During observations the instrument records information from the instrument, MagAO-X, and the telescope. After constructing the .h5 file, the metadata from the appropriate time span is also attached to the file header.
3. **Cosmic Ray Rejection:** Cosmic rays are short-duration, non-astronomical events that cause the full MKID array to “flash”. A cosmic ray rejection algorithm was developed (see Chapter 6) to remove spurious counts from these events.
4. **Wavelength Calibration:** The wavelength calibration measures the relationship between each pixel’s phase response and the wavelength of an incident photon. Once a calibration has been made, it is applied to each of the photons in the .h5 file, where the phase values are converted to wavelength values.
5. **Pixel Calibration:** The pixel calibration step identifies a pixel as “hot”, “cold”, or “dead”. A “dead” pixel is never responsive or marked as badly calibrated in the

beammap. A “hot” pixel is significantly more responsive than other pixels in its vicinity while a “cold” pixel is significantly less responsive to its neighboring pixels.

6. **Linearity Calibration:** The MKID Gen2 firmware sets a finite dead time for each pixel where the readout will not trigger on a photon for $10 \mu\text{s}$ after it has already triggered on an earlier photon. This leads to a non-linear response from MKID pixels at higher count rates as they start missing more photon events due to the dead time. In van Eyken et al. (2015) an empirical correction was found to be $(1 - N \times \tau/T)^{-1}$, where N is the number of photons measured in time T for a pixel whose dead time is τ . This step may be implemented but the MKID response at the count rates in XKID is typically low enough that it does not become non-linear.
7. **Flat-field Calibration:** The flat-field correction is designed to achieve a uniform response across the array. To perform this calibration the array is flood illuminated then pixel responses are evaluated relative to the average flux per pixel and assigned a weight based on its response. More responsive pixels have a low weight and less responsive have a higher weight. The weights are recorded in the .h5 file.
8. **Astrometric Calibration:** The astrometric calibration determines the World Coordinate System (WCS) transformation to convert an image from pixel (x,y) to on-sky (RA, Dec.) coordinates. This is done by correlating the on-array (x,y) position of a PSF of a star with a known (RA, Dec.) value. Performing this at multiple positions allows a WCS be generated along with a conex-to-pixel mapping.
9. **Data Product Generation:** After performing all (or the desired subset) of the calibration steps, the desired outputs will be generated for the user’s analysis.

Table 4.1. XKID Commissioning Targets

| Name | Dates Observed | Total Duration (s) | Description |
|-------------------------------|----------------|--------------------|---|
| HD 96220 | 20230306 | 70 | Binary for instrument calibration |
| Trapezium B4 | 20230307 | 1132 | Binary for instrument calibration |
| β Pictoris ^a | 20230307 | 1195 | Star with 2 known exoplanets, |
| – | 20230308 | 5274 | β Pic b and β Pic c |
| – | 20230310 | 3580 | |
| HD 32297 | 20230311 | 2640 | Star with a bright edge-on disk |
| HD 72946 | 20230311 | 4400 | Star with a close-in brown dwarf (BD) companion |

Note. — (a) The observations of β Pic on 20230307 and 20230308 were structured to try to uncover β Pic b while the observations on 20230310 are oriented at imaging β Pic c.

4.4 Observational Data

Several of the initially observed target datasets have been reduced through the MKID science data pipeline (Steiger et al., 2022a). The astronomy analysis will get underway in earnest after validating that each step from the pipeline works as intended with a new instrument and configuration. The stars targeted are listed in Table 4.1 and several of them will be discussed and shown in the remainder of this section. Over the course of the observing run, the best seeing achieved was below $0''.40$, with median seeing between $0''.50$ - $0''.70$ and occasional deviations up to $\sim 1''.25$.

β Pictoris

The β Pictoris (β Pic for short) system is one of the most studied exoplanet systems and is host to two super-Jupiters. The outermost planet, β Pic b, has a mass of $11 \pm 2 M_{\text{Jup}}$ and orbits at ~ 10 au from its host star. It was originally discovered with direct imaging at the Very Large Telescope (Lagrange et al., 2009). Since its discovery it has been

followed up to characterizing its orbital parameters and physical properties.

Roughly a decade later a second planet, β Pic c, was discovered orbiting closer to the host star. Initially discovered using the Radial Velocity (RV) method at La Silla Observatory (Lagrange et al., 2019), it is noteworthy for being one of the few exoplanets confirmed using three separate techniques. In addition to an RV detection it was found using direct imaging (Nowak et al., 2020) and astrometry (Brandt et al., 2021). β Pic c is also $\sim 10M_{\text{Jup}}$, but orbits at approximately 2.7 au from the host star.

Figure 4.3 shows a VLT image of the β Pic system after the discovery of the first planet, an artists rendering of the system after the second planet’s discovery, and an image of the star using MagAO-X/XKID.

The XKID image shown comes from a 1080-second dithered observation that covers a roughly $0''.60$ -x- $0''.60$ patch of sky without any noise-reduction algorithms applied. This observation used a coronagraph with diameter $d \approx 80$ mas that is obscured in the image by a Spot of Arago, a bright point of light that appears in the center of a circular object’s shadow because of effects from Fresnel (“near-field”) diffraction. Also seen are the first 4 Airy rings from the diffraction pattern of the optical system. Neither companion is immediately visible, although the brightness of β Pic c is expected to be comparable to the brightness of the Airy rings, making it an excellent candidate for SSD.

HD 32297

HD 32297 is not a known planetary system, although it still provides an exciting direct imaging target due to the bright disk that it hosts. The Gemini Planet Imager (GPI) team was able to directly image the bright edge-on debris disk in H -band (1475-1785 nm) down to $0''.15$ in both polarized and total intensity (Duchêne et al., 2020). At tighter angular separations from this, however, the disk has not been imaged since it was blocked by the GPI coronagraph (Figure 4.4, panels **(a)** and **(b)**).




Figure 4.3 The β Pic system. (Left) A coronagraphic image taken at VLT. The central star is obscured by a coronagraph and β Pic b is seen northeast of the star. Also shown is the circumstellar debris disk the planets are embedded in. Credit <https://eso.org> and A.-M. Lagrange. (Center) An artist rendering from <https://www.sci.news/astromy/beta-pictoris-c-direct-image-08926.html> of the β Pic system with b, c, and the debris disk. The dashed white box shows the XKID field of view. (Right) A coronagraphic image of β Pic from MagAO-X/XKID. The central peak is a “Spot of Arago” and the first 4 Airy rings with several “pinned” speckles can be seen surrounding it.

In principle, if the disk does extend further in toward the central star it will be visible to MKIDs with sufficiently good seeing and exposure time. Additionally, it has been demonstrated in Steiger et al. (2022b) that SSD works on extended sources so it may also be used to dig beneath the speckle halo in the MagAO-X/XKID data for this star (Figure 4.4, panel (c)) and potentially uncover structure that isn’t visible in raw images.

HD 72946

The final science target observed during the XKID commissioning run is HD 72946, a fairly bright G5V star the same mass and brightness as the Sun that is orbited by a brown dwarf companion. The companion was originally found using RV data and confirmed via direct imaging in 2019 (Maire et al., 2020). It has a mass of $72.5 \pm 1.3 M_{\text{Jup}}$ and orbital radius of approximately 6.5 au (Brandt et al., 2021).

In the YJ photometric bands, the brown dwarf companion has a contrast of roughly




Figure 4.4 GPI H -band total intensity images, reproduced from Duchêne et al. (2020) for a single frame (a) and a complete observational sequence (b) shown on the same logarithmic stretch. Each of these frames is $2''.5$ on a side. The white dashed square in (b) shows the approximate field of view from the XKID observation. Panel (c) shows a coronagraphic image using a linear stretch from MagAO-X XKID that extends $\sim 0''.60$ per side in the same North-East orientation as panel (b).




Figure 4.5 (Left) Reproduced from Maire et al. (2020). A frame showing a coronagraphic image of HD 72946 B in YJ band. (Center) A coronagraphic MagAO-X/XKID dithered dataset of HD 72946 oriented similarly to the left panel. The red circle shows the approximate region of the predicted position of HD 72946 B in March 2023. The 4 elongated “blobs” are the sparkles used as an astrometric reference, slightly smeared out by rotation during the observation sequence. (Right) The predicted orbit of HD 72946 B around its host star, also from Maire et al. (2020).

5×10^{-5} from its host star at $0''.27$ (Maire et al., 2020), making it the most technically demanding of the targets discussed to view due to the combination of its relatively high contrast level and narrow separation. As with the previous targets, it is a good SSD target – although potentially worse than the previous two since its brightness is lower relative to the speckle halo. In addition, HD 72946 was observed through its transit meaning that there is significant rotation through the duration of the observation so it will be a good candidate to reduce using ADI (Angular Differential Imaging, Marois et al. (2006) and Section 1.2.7).

The left panel of Figure 4.5 shows the initial SPHERE instrument (Beuzit et al., 2019) observations of the BD companion in YJ bands. In the rightmost panel, the predicted orbit from Maire et al. (2020) is also shown for reference. The central panel shows a MagAO-X/XKID image without any post-processing or noise mitigation techniques applied. It is possible to see the cross made by the sparkles in the XKID image, slightly smeared rotationally from the rotation of the sky during the observation. The predicted position of HD 72946 B is also shown by the red oval in the picture. It can be seen that within the predicted region there are two fairly bright spots which look like they *may* be PSFs. Further analysis is required to validate whether these are speckles that are pinned to the edge of the coronagraph or if they are actual astronomical sources.

4.5 Future Work

The previous discussion of the three main science targets all showed raw MagAO-X/XKID images (Figures 4.3, 4.4, and 4.5). While there are no clear and obvious detections from these sequences, there are still several techniques available to reduce noise in the images such as ADI and SSD. Analysis using these techniques is ongoing and will be presented in either the XKID commissioning paper or science papers, or both.

XKID is also slated to return to LCO and the Magellan Clay telescope during the April 2024 MagAO-X run. Although MagAO-X is nominally scheduled to return to LCO multiple times each year, the period following the March 2023 run coincides with a major hardware upgrade in the MagAO-X optical bench, which includes the installation and characterization of a new deformable mirror.

During this time XKID will return to UCSB. While it is at UCSB there are plans to modify the filter stack to reduce IR transmission into the fridge. The design of the new filter to mitigate IR photons in the cryostat is ongoing.

Additionally, the XKID camera will be used as part of a bioimaging project collaborating with the Streichan Lab, also at UCSB to demonstrate the utility of MKID instruments in other photon-starved regimes, such as the imaging of fluorescent proteins in fruit fly (*Drosophila melanogaster*) embryos to explore the insect's morphogenesis.

Chapter 5

The Discovery and Characterization of HIP 5319B

5.1 Attribution

The work in this chapter was originally published in Swimmer et al. (2022a) and has been modified for content and flow in the context of this thesis.

5.2 Introduction

Over the past two decades, both facility AO and now extreme AO systems have provided numerous images of planets and low-mass brown dwarfs around nearby stars (e.g. Marois et al., 2010; Carson et al., 2013; Kuzuhara et al., 2013; Currie et al., 2014; Macintosh et al., 2015; Konopacky et al., 2016; Chauvin et al., 2017; Cheetham et al., 2018; Keppler et al., 2018; Currie et al., 2022c). The majority of discoveries draw from “unbiased” surveys, where targets are selected based on age and distance (e.g. Desidera et al., 2021). However, these same surveys show that occurrence rates of detectable

moderate-to-wide separation planets and brown dwarf companions is low, \sim a few percent around FGK stars (Nielsen et al., 2019; Vigan et al., 2021; Currie et al., 2022a).

Recent work has demonstrated the success instead of dynamics-selected direct imaging surveys, specifically using precision astrometry from the *Gaia* and *Hipparcos* satellites in the *Hipparcos-Gaia Catalog of Accelerations* (HGCA) to identify stars showing a proper motion anomaly – i.e. an astrometric acceleration – likely due to an unseen low-mass companion (van Leeuwen, 2007; Brown et al., 2018; Gaia Collaboration et al., 2021; Brandt, 2021). Direct imaging of targets showing an acceleration from HGCA have revealed white dwarfs (Bonavita et al., 2020), low-mass stars (Steiger et al., 2021; Chilcote et al., 2021), moderate-to-low mass brown dwarfs (Currie et al., 2020; Bowler et al., 2021; Bonavita et al., 2022; Kuzuhara et al., 2022), and now planets (Currie et al., 2022b).

Jointly analyzing absolute astrometry of the star from HGCA and relative astrometry of the imaged companion with Markov-Chain Monte Carlo (MCMC) codes like *orvara* (Brandt et al., 2021) can provide strong constraints on the companion’s dynamical mass and orbit (e.g. Brandt et al., 2021). To derive these constraints, MCMC codes require input priors for the orbital parameters, primary mass, and companion mass(es). Typical orbital priors include a log-normal distribution in semimajor axis ($p(a) \propto 1/a$), uniform prior in inclination ($p(i) \propto \sin(i)$), gaussian prior in primary mass, and log-normal prior in companion mass ($p(M_2) \propto 1/M_2$) (e.g. Kuzuhara et al., 2022).

While the above orbital priors are long regarded as standard in MCMC modeling (e.g. Blunt et al., 2020), the most appropriate companion prior may differ. The initial mass function for companions near the substellar to stellar boundary exhibits a more gaussian-like distribution (e.g. Chabrier, 2003): i.e. a turnover in the mass function near the hydrogen-burning limit. Ancillary system properties – e.g. age, primary and companion spectral type, etc. – also are often used to inform adopted priors but may derive from heterogeneously-sourced data.

Here, we report the direct imaging discovery of HIP 5319 B: a low mass – potentially substellar – companion around the F-type star HIP 5319 A using the Subaru Coronagraphic Extreme Adaptive Optics system (SCExAO; Jovanovic et al., 2015; Ahn et al., 2021) coupled with the MKID Exoplanet Camera (MEC; Walter et al., 2020), the Visible Aperture Masking Polarimetric Imager for Resolved Exoplanetary Structures (VAMPIRES; Norris et al., 2015), the Coronagraphic High Angular Resolution Imaging Spectrograph (CHARIS; Groff et al., 2016), and the NIRC2 camera on the Keck II telescope.

HIP 5319 B illustrates the sensitivity of adopted priors for companion mass for parameters derived from jointly modeling direct imaging and astrometric data and the need to verify ancillary information about the system – e.g. binarity, age, rotation – in direct imaging + astrometric surveys.

5.3 Stellar Properties and Observations

5.3.1 HIP 5319 A Basic Properties

HIP 5319 (*78 Psc) is an F5IV spectral class star (Boro Saikia et al., 2018) at $d = 42.93 \pm 0.06$ pc (Prusti et al. (2016), Brown et al. (2018)). Banyan- Σ (Gagné et al., 2018) shows no evidence that the system is a member of any moving group or young association. It has previously been identified as an RS CVn binary star by Fleming et al. (1989), who measured a projected rotation rate of $v \sin(i) = 68 \pm 20.5$ km/s and x-ray luminosity of $L_x = 9.2 \pm 3.7 \times 10^{28}$ erg/s.

System Age

Evidence informing the HIP 5319 system’s age is complex. On one hand, HIP 5319 has an extreme level of chromospheric activity ($\log(R'_{HK}) = -4.016$) as measured by

Calcium II H and K lines, which tracks the strength of the emission at the cores of the 2 lines (Boro Saikia et al., 2018). The chromospheric index easily exceeds values for stars in the Pleiades and Hyades associations and is comparable or higher to the stars in the Scorpius-Centaurus (Sco-Cen) association (Mamajek & Hillenbrand, 2008; Pecaut & Mamajek, 2013). Its Hertzsprung-Russell diagram position in *Gaia* color-magnitude space (M_G vs $G_{BP} - G_{RP}=2.97, 0.54$) lies between the Pleiades and Hyades, which is consistent with either a main sequence star between 115 and ~ 800 Myr, respectively (Gossage et al., 2018), or a pre-main sequence star much younger than the Pleiades. Based on its activity, Stanford-Moore et al. (2020) estimate a young age of $75^{+492}_{-63.5}$ Myr.

HIP 5319 was also observed by the Transiting Exoplanet Survey Satellite (TESS; Ricker et al., 2015) and has 2-minute cadence photometry for one sector. This observation may be too short to show spots reliably, but it does show pulsations with a period of just less than 1 day¹. It was also observed once by the International Ultraviolet Explorer (IUE) during IUE Program ID: CB401 (Stellar Chromospheres; Blanco et al., 1982). In the spectrum from IUE², HIP 5319 A shows strong emission from the Lyman α line. These two data points show signs that the primary might be chromospherically active, though follow up observation is required to determine the nature of this activity.

On the other hand, RS CVn binaries – of which HIP 5319 is claimed to be an example – typically have orbital periods less than 14 days and show high levels of chromospheric activity via strong emission in Calcium II H and K lines, and have a hotter component of spectral type F or G (Montesinos et al., 1988). Multiple sources have reported $v \sin(i)$ values with significant scatter, which may suggest binarity: 125 km/s (Danziger & Faber, 1972), 68 ± 20.5 km/s (Fleming et al., 1989), 36.4 ± 4.8 km/s (de Medeiros & Mayor, 1999), 35 km/s (Nordström et al., 2004), and 41.5 km/s (Głębocki & Gnaniński (2005), Głębocki

¹Accessed via <https://mast.stsci.edu/portal/Mashup/Clients/Mast/Portal.html>

²Accessed via <https://archive.stsci.edu/iue/obtaining.html>

& Gnacinski (2005)). The fractional x-ray luminosity of the star is $\log(L_x/L_{bol}) \approx -4.9$ (Gioia et al., 1990; Favata et al., 1995), almost two orders of magnitude less than a typical pre-main sequence star, which would have values of $\log(L_x/L_{bol}) \approx -3.2$ for fractional x-ray luminosity (Preibisch et al., 2005), respectively. Other authors have estimate the age of the star using isochrones and have found values of $1.6_{-0.4}^{+0.3}$ Gyr (Holmberg et al., 2009) and 1.07-1.23 Gyr using Padova and BASTI models (Casagrande et al., 2011).

Ultimately, the conflicting identifications of the HIP 5319 primary as either a young, chromospherically active star or an older star whose Ca II HK emission is due to a close binary has significant implication on the understanding of the stellar system and interpretation of any of its companions' properties. If there is not significant HK emission and little evidence of binarity then the higher age estimate is likely the correct one, which anchors the interpretation of its companion. In addition to performing a direct imaging search for such a binary companion, a study of the primary using a high resolution spectrograph is necessary to disentangle the possible identities of the star and settle on the correct interpretation. This will be discussed further in sections 5.4.1 and 5.4.2.

Evidence for An Astrometric Acceleration

The *Hipparcos-Gaia* Catalog of Accelerations reports a $\chi^2 = 171.04$, evidence of a $12.9\text{-}\sigma$ significant acceleration of the primary with 2 degrees of freedom (Brandt, 2021). The statistically significant acceleration of HIP 5319 is suggestive of the presence of a previously unseen low-mass companion at a $\gtrsim 10$ au scale. HIP 5319 was not known to have a wide-separation binary companion that could plausibly be source of this acceleration.

Therefore we chose to observe this target in an attempt to uncover any previously unimaged low-mass companions around this accelerating star, following a similar method of target selection as in Currie et al. (2020) and Steiger et al. (2021).

Table 5.1. HIP 5319 Observing Log: Coronagraphs and Photometric Bands

| UT Date | Instrument | coronagraph | Seeing (") | Passband | λ (μm) ^a |
|----------|-----------------|-------------|------------|------------|--|
| 20200731 | SCEXAO/CHARIS | Lyot | 0.4–0.6 | <i>JHK</i> | 1.16–2.37 |
| – | SCEXAO/MEC | Lyot | – | <i>Y</i> | 0.95–1.12 |
| 20210911 | SCEXAO/CHARIS | Lyot | 0.5–0.6 | <i>JHK</i> | 1.16–2.37 |
| – | SCEXAO/VAMPIRES | – | – | 750nm | 0.75 |
| 20220115 | Keck/NIRC2 | none | 0.6 | L_p | 3.78 |
| 20220119 | SCEXAO/MEC | Lyot | 0.7 | <i>YJ</i> | 0.95–1.4 |

Note. — a) For CHARIS and MEC data, this column refers to the wavelength range. For broadband imaging data, it refers to the central wavelength.

Table 5.2. HIP 5319 Observing Log: Observations and Post-processing Strategies

| UT Date | Instrument | t_{exp} (s) | N_{exp} | ΔPA ($^\circ$) | Post-Processing Strategy |
|----------|-----------------|----------------------|---------------------|--------------------------------|--------------------------|
| 20200731 | SCEXAO/CHARIS | 30.98 | 14 | 5.3 | RDI-KLIP |
| – | SCEXAO/MEC | 5.0–10.0 | 61 ^a | 4.6 | none |
| 20210911 | SCEXAO/CHARIS | 30.98 | 8 (32) ^b | 9.9 | none |
| – | SCEXAO/VAMPIRES | 12.8 | 48 | 11.2 | ADI-ALOC1 |
| 20220115 | Keck/NIRC2 | 30 | 30 | 9.1 | RDI-KLIP |
| 20220119 | SCEXAO/MEC | 15 | 49 | 3.8 | none |

Note. — a) Total integration time is 430 s. b) In total, we obtained 32 exposures but only 8 were retained due to substantial PSF core splitting from low-wind effect.

5.3.2 Observations and Data Reduction

HIP 5319 was observed over three epochs in July 2020, September 2021, and January 2022 at the Subaru Telescope on Maunakea using SCEXAO coupled with the CHARIS, MEC, and VAMPIRES instruments. During these epochs, the seeing conditions at the Subaru Telescope ranged between $\theta_V=0''.4\text{--}0''.7$. Observing conditions were photometric each night³. It was observed for a fourth epoch in January 2022 at the W.M. Keck Observatory on Maunakea using the NIRC2 instrument coupled with the Keck Adaptive Optics system. The seeing during this epoch was $\theta_V=0''.6$. The observations from these runs are summarized in Tables 5.1 and 5.2.

³The observing conditions during the January 2022 epoch were photometric, but due to instrument constraints there was no appropriate energy calibration of the MEC instrument, disallowing the measurement of a meaningful photometric data point.

All of the observations taken with SCEXAO used its “vertical angle”/pupil-tracking mode which enables ADI (Marois et al., 2006). Each set of data also used the Lyot coronagraph ($0''.113$ radius occulting mask) to suppress light from the primary star. The data in both epochs also utilized satellite spots for precise astrometric and spectrophotometric calibration (Jovanovic et al., 2015; Currie et al., 2018a).

The MEC data in July 2020 was taken in Y band ($0.95\text{-}1.12\ \mu\text{m}$) with a spectral resolution $\mathcal{R} \sim 4.0$ simultaneously with CHARIS broadband data. The CHARIS data in both epochs was taken in its low-resolution broadband mode covering JHK passbands ($1.16\text{-}2.37\ \mu\text{m}$) at $\mathcal{R} \sim 18$. VAMPIRES data were taken at $750\ \text{nm}$ concurrently with CHARIS in broadband mode in September 2021. In addition to the SCEXAO observing mode allowing for ADI, the CHARIS spectral coverage enables SDI (Marois et al., 2000). The NIRC2 data were taken in the L_p filter ($\lambda_c = 3.78\ \mu\text{m}$). Later in January 2022 more MEC data were taken covering YJ bands ($0.95\text{-}1.14\ \mu\text{m}$) with resolution $\mathcal{R} \sim 2.4$.

HIP 5319 was also observed for spectroscopic characterization of the primary during January and February of 2022. Spectra were obtained using the Network of Robotic Echelle Spectrographs (NRES) 1-m instrument operated by the Las Cumbres Observatory global telescope network (LCOGT; Brown et al., 2013) at the Wise Observatory in Mitzpe Ramon, Israel over the course of 9 nights from 20 January to 12 February 2022. These were taken using fiber-fed optical ($0.38\text{-}0.86\ \mu\text{m}$) echelle spectrographs with a spectral resolution $\mathcal{R} \approx 50,000$ and an $\text{SNR} > 200$ for all but two of the spectra. The spectroscopic observations from the LCOGT NRES instrument are summarized in Table 5.3.

CHARIS

We extracted CHARIS data cubes from the raw data using the standard CHARIS pipeline (Brandt et al., 2017) to perform basic reduction steps – image registration and spectrophotometric calibration. We did not obtain sky frames for sky subtraction. For

Table 5.3. HIP 5319 LCOGT Observing Log

| BJD | t_{exp} (s) | SNR ^a | RV (km/s) | $v\sin(i)$ (km/s) |
|--------------------------|----------------------|------------------|------------|-------------------|
| 2459600.268 | 1000 | 230 | 17.30±1.80 | 95.24±1.65 |
| 2459601.266 | – | 237 | 17.94±1.59 | 95.63±1.59 |
| 2459605.194 | 1500 | 316 | 18.12±1.81 | 93.37±1.64 |
| 2459607.221 | 1000 | 227 | 16.07±1.23 | 95.89±1.64 |
| 2459608.227 | – | 246 | 16.18±1.13 | 94.50±1.64 |
| 2459608.246 | 1500 | 218 | 14.04±1.28 | 94.49±1.64 |
| 2459609.221 | – | 281 | 17.31±2.56 | 92.69±1.60 |
| 2459609.202 ^b | 1000 | 227 | – | – |
| 2459610.220 | 1500 | 277 | 17.24±1.41 | 94.51±1.61 |
| 2459610.242 ^c | 1000 | 249 | – | – |
| 2459612.185 | – | 189 | 16.31±1.88 | 93.08±1.67 |
| 2459614.192 | – | 170 | 16.53±1.39 | 92.68±1.79 |
| 2459622.193 | – | 256 | 17.63±2.29 | 94.50±1.60 |
| 2459623.194 | – | 203 | 18.41±2.63 | 92.22±1.76 |

Note. — BJD 2459600 corresponds to UT Date 20220120. All data taken from $\lambda = 0.38 - 0.86\mu\text{m}$.

- a) Values reported are SNR per resolution element at $0.518\mu\text{m}$.
b, c) 2459609 and 2459610 both have 2 spectra. For each night, both spectra are combined to measure RV and $v\sin(i)$ signal.

spectrophotometric calibration, we adopted a Kurucz stellar atmosphere model appropriate for an F5IV star. HIP 5319 B is easily visible in the raw data for both CHARIS observations, but the September 2021 data suffered chronic PSF splitting due to low-wind effect, leaving only 8 exposures totaling just over 4 minutes of integration time. The July 2020 data were stable, so we consider the September 2021 data only for astrometry and employ PSF subtraction to yield a high-quality spectrum for the July 2020 data.

To subtract the PSF in the July 2020 data, we followed steps as in Steiger et al. (2021), using a full-frame implementation of reference differential imaging (RDI) using the *Karhunen-Loe'Ve Image Projection* (KLIP; Soummer et al., 2012) algorithm as in Currie et al. (2019). Since the companion around HIP 5319 was easily visible, we adopted a conservative approach, truncating the KLIP basis set at one mode ($KL = 1$) and correcting for minor throughput losses using KLIP forward-modeling as in Pueyo (2016).

VAMPIRES

For VAMPIRES data, we subtracted dark frames and then aligned each sub-exposure within the 12.8 second data cubes, removing outliers. Subsequent steps used the general purpose high-contrast ADI broadband imaging pipeline from Currie et al. (2011). To calibrate the VAMPIRES photometry an appropriate PHOENIX model stellar spectrum⁴ (Husser et al., 2013) for an F5IV star was obtained and then normalized to the reported J band flux value for the HIP 5319 primary from the Two Micron All Sky Survey (2MASS; Skrutskie et al., 2006). Once the model stellar spectrum had been calibrated, the flux density at 750 nm was found to be 13.18 Jy. For PSF subtraction, we found the best results with a full-frame implementation of ALOCI (Currie et al., 2012, 2015). Following Currie et al. (2018a), we used forward-modeling to correct for throughput losses.

⁴<http://phoenix.astro.physik.uni-goettingen.de/>

MEC

Y band images were created using the MKID Science Data Pipeline (Steiger et al., 2022a) to apply calibrations to the raw MEC data that include cold-, dead-, and hot-pixel masking, along with wavelength, astrometric, and spectrophotometric calibrations. There was no PSF subtraction performed for the data from MEC in this analysis.

The spectrophotometric calibration follows the treatment in Steiger et al. (2021) in which the flux from the elongated satellite spots in the image was measured using a “racetrack” aperture (Millar-Blanchaer et al., 2016) before being converted to the stellar flux behind the coronagraph using the relationship between satellite spot contrast and bandpass described in Currie et al. (2018a). The stellar flux in the observation is matched to the calibrated model spectrum from section 5.3.2 to find a spectrophotometric solution, which is applied to the image to convert from counts per second to flux density units.

NIRC2

Our reduction steps followed ones outlined in Steiger et al. (2021). Briefly, we used a well-tested general purpose high-contrast ADI broadband imaging pipeline (Currie et al., 2011) to perform sky subtraction, image registration, and photometric calibration. To subtract the PSF, we used a full-frame implementation of RDI using the KLIP algorithm as in Currie et al. (2019). The star BD+54 408 was used as a reference PSF. Following Pueyo (2016), we used forward-modeling to correct for throughput losses.

NRES

All spectra from the LCOGT 1-m NRES observations are automatically reduced using the BANZAI-NRES data reduction pipeline⁵. After reduction, each spectrum was fit to

⁵Accessible at <https://github.com/lcogt/banzai-nres>

the same F5IV star model stellar spectrum used by MEC and VAMPIRES. Both RV and $v\sin(i)$ values for the primary were extracted using the H α and H β spectral lines (nominally at $\lambda_\alpha = 0.656\mu\text{m}$ and $\lambda_\beta = 0.486\mu\text{m}$) which are shown in Table 5.3.

The RV and $v\sin(i)$ values were calculated iteratively. For each spectra, an RV offset was fit via cross correlation with the PHOENIX model spectrum convolved to a first guess $v\sin(i)$ of 100 km/s. At that RV offset a $v\sin(i)$ is then calculated by minimizing χ^2 between the model and NRES spectra, convolving over a grid of $v\sin(i)$ values between 50 and 150 km/s. This process is then iterated until the values for RV and $v\sin(i)$ converge, meaning that the scatter between the value of the most recent iteration and the previous is less than the formal error. The formal errors on the $v\sin(i)$ values are calculated using standard χ^2 statistics. The formal error on the RV values are from the 1σ confidence interval of the bootstrap probability density of the radial velocity.

5.3.3 Detections

Over all epochs the seeing conditions and data quality led to strong detections of the companion in each observing dataset. To calculate flux density measurements in each channel, we performed aperture photometry sized to $1 \lambda/D$. The SNR was calculated in the standard fashion, summing each pixel within an aperture, computing the robust standard deviation of these summed pixels as a function of angular separation and dividing by the stellar flux (Marois et al., 2008; Currie et al., 2011). Our spectrophotometric errors and SNR values consider finite-element corrections (Mawet et al., 2014). SNR values range from 15.7 in the Y band image from MEC to 763 in the broadband wavelength-collapsed CHARIS data taken in January 2022 and July 2020, respectively. Following previous work, we use the IDL function `cntrd.pro` to estimate companion centroids: the error budget considers the intrinsic SNR of the detection, uncertainties in the plate scale




Figure 5.1 Detection of HIP 5319 B from SCEXAO coupled with MEC, CHARIS, and VAMPIRES and Keck II Adaptive optics coupled with NIRC2. The MEC and VAMPIRES images retain some residual signal from satellite spots used for spectrophotometric and astrometric calibration. In MEC data, these spots appear with different brightnesses due to vignetting from the optics in MEC and dead pixels on the array, both of which have since been corrected. The NIRC2 image also retains some signal from the primary that was not removed by RDI-KLIP. The CHARIS data do retain some residual signal although the signal is so low that it cannot be seen without drastically lowering the maximum value of intensity in the image and saturating the PSF.

and north position angle, and astrometric biases from processing (Pueyo, 2016).

In the July 2020 data, HIP 5319 B is located at $[E,N]''=[0''.124, 0''.311]\pm[0''.004, 0''.004]$ and $[0''.119, 0''.314]\pm[0''.010, 0''.010]$ in the CHARIS and MEC data, respectively. The errors in position take into account centroiding precision, the uncertainty in true north position angle, and pixel scale of each instrument following Currie et al. (2020).

The September 2021 data from CHARIS and VAMPIRES show the companion at $[E,N]''=[0''.133, 0''.287]\pm[0''.004, 0''.004]$ and $[0''.132, 0''.287]\pm[0''.004, 0''.004]$. The measurements taken by multiple instruments in both epochs are the same within error. The detections from each instrument are shown in Figure 5.1.

In January 2022, the NIRC2 and MEC data show the companion at $[E,N]''=[0''.133, 0''.275]\pm[0''.003, 0''.003]$ and $[0''.131, 0''.273]\pm[0''.010, 0''.010]$, where the MEC data were taken 4 days after the NIRC2 observations.

Based on the proper motion of the primary between July 2020 and September 2021, a background star would have moved north-west by $\sim [-0''.23, 0''.03]$, which is inconsistent with the measured companion offset of $[0''.009, -0''.024]$.

In standard Maunakea Observatory filters, the photometry for HIP 5319 B from the CHARIS broadband data is found to be $J = 10.88 \pm 0.02$, $H = 10.31 \pm 0.02$, and $K = 10.07 \pm 0.03$ from the July 2020 data. These values are within 1σ uncertainty for H and K band and 2σ uncertainty for the measured J band photometry points measured in September 2021. The MEC Y band photometry is found to be $Y = 11.3 \pm 0.1$, and VAMPIRES measured a flux density of 18.83 ± 0.83 mJy at 750 nm⁶. Note that these measurements do not consider an absolute spectrophotometric uncertainty – i.e. a multiplicative factor in flux density, additive in magnitude – of 5% due to uncertainties in the mapping between the deformable mirror modulation amplitude used to produce satellite spots and the resulting spot contrast at our fiducial wavelength of 1.55 μm .

⁶For discussion of VAMPIRES photometry and its conversion to a pseudomagnitude see section 5.4.3.

Table 5.4. HIP 5319 B Detection Significance and Astrometry

| UT Date | Instrument | SNR | [E,N]('') |
|----------|-----------------|------|-------------------------------------|
| 20200731 | SCExAO/CHARIS | 763 | [0.124, 0.311] \pm [0.004, 0.004] |
| 20200731 | SCExAO/MEC | 22.8 | [0.119, 0.314] \pm [0.010, 0.010] |
| 20210911 | SCExAO/CHARIS | 48 | [0.133, 0.287] \pm [0.004, 0.004] |
| 20210911 | SCExAO/VAMPIRES | 23 | [0.132, 0.287] \pm [0.004, 0.004] |
| 20220115 | Keck/NIRC2 | 16.1 | [0.133, 0.275] \pm [0.003, 0.003] |
| 20220119 | SCExAO/MEC | 15.7 | [0.131, 0.273] \pm [0.010, 0.010] |

Table 5.5. HIP 5319 B Detection Photometry

| UT Date | Instrument | Passband | Photometry |
|----------|-----------------|------------|---|
| 20200731 | SCExAO/CHARIS | <i>JHK</i> | J = 10.88 \pm 0.02 , H = 10.31 \pm 0.02, K = 10.07 \pm 0.03 |
| 20200731 | SCExAO/MEC | <i>Y</i> | Y = 11.3 \pm 0.1 |
| 20210911 | SCExAO/CHARIS | <i>JHK</i> | J = 11.02 \pm 0.06 , H = 10.38 \pm 0.05, K = 10.09 \pm 0.06 |
| 20210911 | SCExAO/VAMPIRES | 750nm | 18.83 mJy \pm 0.83 mJy |
| 20220115 | Keck/NIRC2 | L_p | $L_p = 9.39 \pm 0.067$ |
| 20220119 | SCExAO/MEC | <i>YJ</i> | — |

Note. — There is no photometry point measured during the 20220119 SCExAO/MEC observation. The CHARIS photometry do not consider an additional 0.05 magnitude uncertainty drawn from the mapping between the deformable mirror modulation amplitude (used to produce satellite spots used for spectrophotometric calibration) and the resulting satellite spot contrast with respect to the star.

(Currie et al., 2018b). In the Keck II Telescope filters the photometry from the NIRC2 data is found to be $L_p = 9.39 \pm 0.07$. The full summary of the HIP 5319 B detection significance, astrometry, and photometry is found in Table 5.4 and 5.5.

5.4 Analysis

5.4.1 Characterization of HIP 5319 A as a Single Star

Before the properties of the companion can be determined it is first necessary to identify whether the primary is a binary or single star. Using the RV and $v\sin(i)$ values we look for periodic variations in time to help infer the presence of a companion or lack

thereof. The top panels in Figure 5.2 show the measured values of each quantity and the best fit to a constant velocity.

The search for $v\sin(i)$ is motivated due to large scatter in this quantity's previously reported values in the literature, which range from 35 km/s (Nordström et al., 2004) to 125 km/s (Danziger & Faber, 1972) at the low and high ends, respectively. In the collection of stars discussed in Fleming et al. (1989) HIP 5319 has the greatest uncertainty on its $v\sin(i)$ value, nearly double the next highest uncertainty and almost 1/3 of its reported rotation rate. This wide scatter in reported rotation rates along with the high uncertainties reported on these measurements led us to consider whether there may be a binary companion where both objects contribute to the spectrum whose individual signals have not been teased out. Since we can obtain $v\sin(i)$ from the NRES spectra we use this opportunity to search for any signal in the data which may indicate the presence of a second, unseen companion contaminating the signal from the primary star.

The bottom panels in figure 5.2 show periodograms of the residuals from the RV and $v\sin(i)$ data. The peak values of each periodogram are 0.559 and 0.535, respectively. Assuming there is no periodic signal in the data, this means that a peak this high or higher will be seen 79.6% of the time in the RV data and 67.6% of the time in the $v\sin(i)$ data. Also shown are the required peak heights to attain a 1% false alarm probability for each measurement. For the radial velocity data a peak would have to have a power of 0.888 to attain a false alarm probability below 1%, while the $v\sin(i)$ peak would need to have a power of 0.894 to meet the same criterion. The height of the 2 peaks from the periodograms combined with the high peak values needed to attain a 1% false alarm probability show that there is no obvious periodic signal, meaning the time series RV and $v\sin(i)$ data are not consistent with oscillatory behavior caused by a close-in companion.

Both sets of measurements are consistent with constant values to within 1 standard deviation except for a single point: the radial velocity measured from the second spectrum

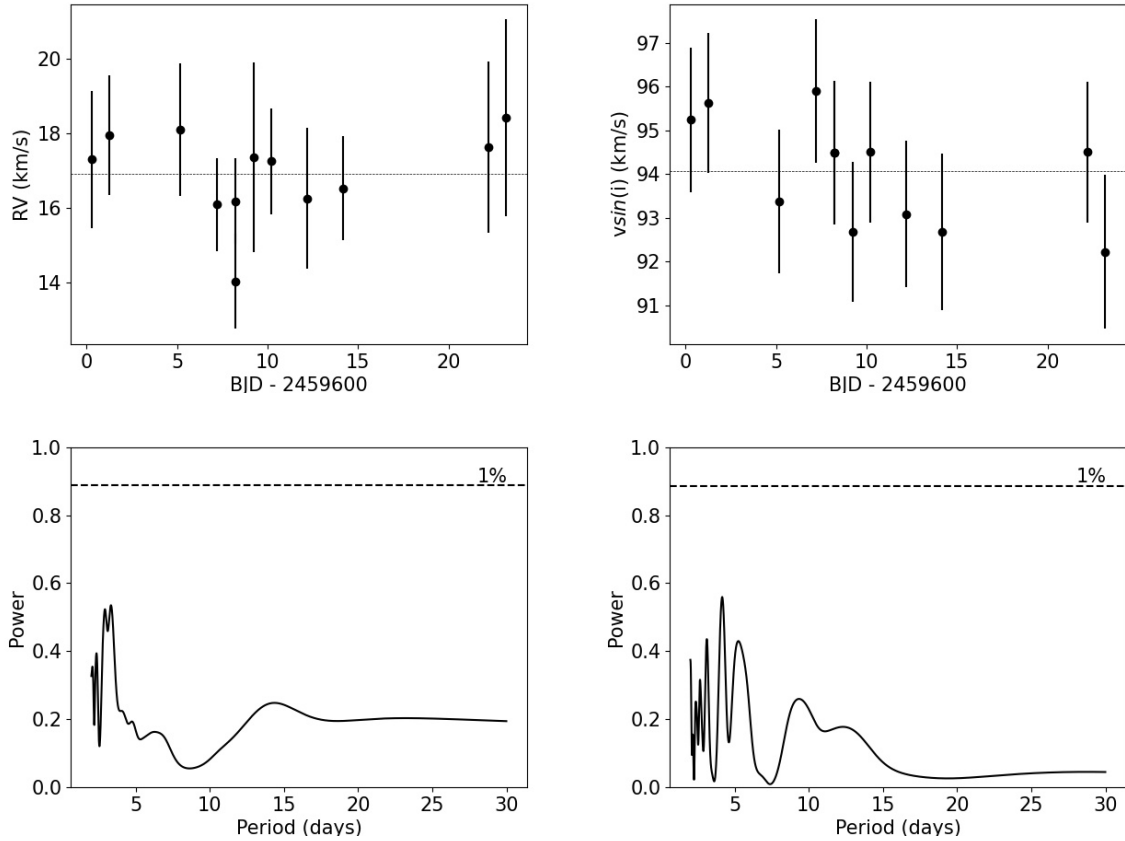


Figure 5.2 (Top) Radial velocity (left) and $v \sin(i)$ (right) values measured for HIP 5319A. The dotted lines in each panel are the best fit constant velocity to the data, where $RV=16.71$ km/s and $v \sin(i)=94.21$ km/s. Neither metric shows significant variation in time or obvious periodicity. (Bottom) Periodograms of the residuals from the radial velocity (left) and $v \sin(i)$ (right) values. The residuals are calculated by taking the measured data and subtracting the best fit constant velocity. The false alarm probability of 1%, calculated using bootstrap randomization, is shown by the dashed lines.




Figure 5.3 (Left) Minimum detectable binary companion mass for various periods (P) and inclinations (i). For a given combination of period and inclination, the reported ‘minimum detectable mass’ can be found, which corresponds to the lowest mass a companion would have that would result in a semi-amplitude $K > 3$ km/s. Any companions less than that mass would be undetectable in the spectroscopic data and any companions that were more massive would have been detected. (Right) Minimum detectable mass as a function of period for selected inclinations.

on BJD 2459608. In both cases we see that we would be sensitive to any periodic signal with a semi-amplitude $K \gtrsim 3$ km/s, while any signal that has $K \lesssim 3$ km/s may still be hidden within the measurement error.

Using equation 5.1 - which relates the semi-amplitude K to the orbital period P of a companion of mass M_2 around a host of mass M_1 with eccentricity and inclination e and i - one can estimate the detectable companion mass for a given set of P , i , and e values.

$$K = \left(\frac{2\pi G}{P} \right)^{1/3} \frac{M_2 \sin(i)}{(M_2 + M_1)^{2/3}} \frac{1}{\sqrt{1 - e^2}} \quad (5.1)$$

For this estimation $K_{max} = 3$ km/s and e is assumed to be equal to 0. We then vary P and i and calculate the smallest mass that would generate an RV semi-amplitude $K > K_{max}$ for each (P, i) combination. The results of this are shown in Figure 5.3 for $2 \leq P \leq 30$ days and $30^\circ \leq i \leq 90^\circ$.

The choice to restrict this analysis to periods between 2 and 30 days is due to the

cadence of observations and the duration of the survey. A companion with a shorter period may still have been detectable although without being able to accurately measure the period. We would not have enough data to detect a companion with a period $P \gtrsim 30$ days since there would be insufficient time to see periodicity in the signal; however, our data do cover the range of expected periods for an RS CVn system ($P \lesssim 14$ days). With regards to the inclination the analysis is not performed below 30° due to the difficulty of detecting companions in RV signals for near face on orbits. The original claim was of this star as a spectroscopic binary, meaning the system would not have been face on.

At the extreme values of the analysis we find that a binary companion with $P = 2$ days and $i = 90^\circ$ would be detectable if it had a mass greater than $24 M_J$ whereas for a companion with $P = 30$ days and $i = 30^\circ$ the minimum mass that would be detectable via an RV signal would be $122 M_J$. This tells us that in the spectroscopic data taken on this star we would have seen the signature for a binary companion above $122 M_J$ at worst and $24 M_J$ at best.

Further spectroscopic data taken at higher precision and over longer times will aid in ruling out lower mass and longer period binary companions, but current data suggest there is no companion with mass greater than $122 M_J$ with a duration less than 30 days, which is sufficient to refute previous evidence of this star being a spectroscopic binary.

5.4.2 Non-detection of Ca HK Emission

We also reassess evidence that HIP 5319 has a high chromospheric activity from Boro Saikia et al. (2018) which claimed to have measured a value of $\log(R'_{HK}) = -4.016$. The methodology behind this claim was to measure the surface flux R_{HK} by co-adding all available spectra for the target into a template spectra that was then normalized to a PHOENIX model atmosphere in order to convert to absolute flux units. The photospheric




Figure 5.4 Data from LCOGT NRES spectra of the primary star HIP 5319 A compared to a PHOENIX model spectrum for an F5IV star surrounding the Ca II H line at its vacuum wavelength $\lambda = 3969.5\text{\AA}$. The model spectrum has been broadened by 100 km/s to match the best-fit $v\sin(i)$ value for the Ca II H line from the LCOGT spectra. (Left top) Model spectra plotted over data from the 12 LCOGT spectra between $\lambda = 3955 - 3985\text{\AA}$. (Left bottom) The O-C (Observed-Computed) plot showing the residuals between measured data and model. Grey points are the residuals from each of the 12 spectra, while the red points are rebinned to the original NRES spectral resolution. (Right) The same data and residuals between $\lambda = 3968 - 3972\text{\AA}$. In both cases there is no excess flux beyond the 1% percent level in the spectrum at any point near the Ca H line.

flux contribution $R_{phot} = F_{phot}/\sigma T_{eff}^4$ was then subtracted from the integrated flux of the Ca II H and K line cores from the PHOENIX model atmosphere. The excess that was seen after this subtraction interpreted as being from emission at the H and K lines.

By comparing the high resolution LCOGT spectra (Section 5.3.2, Table 5.3) and the model PHOENIX spectrum for an F5IV star used for photometric calibration (Section 5.3.2) we find no evidence to support the claim of any excess flux around the Ca II H or K lines beyond the 1% level.

To compare the difference between the model PHOENIX spectrum and LCOGT spectra, each nightly spectrum was individually normalized using a scale factor, slope, and offset. Figure 5.4 shows the result of this comparison for the Ca II H line at $\lambda = 3969.5\text{\AA}$.

The top panels in Figure 5.4 show the data from all of the spectra in Table 5.3 compared to the model PHOENIX spectrum, while the bottom panels show the residuals between the data and model. The residuals from each spectra compared to the model are shown as grey points, while the red points show the residuals when the data are rebinned to the original NRES $R \sim 50,000$. This rebinning was performed because each spectra that makes up the combined dataset (made of 12 individual spectra) samples slightly different rest-frame wavelengths because of the evolving barycenter velocity over the 23 days where spectra were collected. This means that there is roughly 12 times as much data since the same wavelengths are not sampled multiple times. By rebinning to the original NRES resolution this has the effect of demonstrating what a single spectra would look like for ~ 12 times as much observation time as one of the individual spectra.

The data collected in the 14 observations match the model without any significant deviation around the cores of the Ca II HK lines. Along with the non-detection of an time-varying signal in the RV and $v\sin(i)$ data this refutes the evidence that the primary is an RS CVn binary which is expected to have high chromospheric activity and a period below 14 days, meaning it is likely a single star. This is in good agreement with the report




Figure 5.5 Combined SCEsAO/CHARIS spectra, SCEsAO/MEC photometry, SCEsAO/VAMPIRES and Keck/NIRC2 photometry of the low mass companion HIP 5319 B taken on July 31, 2020 (CHARIS and MEC), September 11, 2021 (CHARIS and VAMPIRES), and January 15, 2022 (NIRC2) at the Subaru and Keck II telescopes. The reddest CHARIS channel has substantially higher uncertainty in our spectrophotometric calibration, because we did not obtain sky frames.

of HIP 5319A from the Gaia Early Data Release 3 (Gaia eDR3; Brown et al., 2021) as being well fit by a 5-parameter single star solution whose Renormalized Unit Weight Error (RUWE) is 1.01, which effectively rules out stellar-mass companions greater than $\sim 0.4M_{\odot}$ and a period between 1 and 10 days.

5.4.3 Spectrum of HIP 5319 B

Figure 5.5 shows the 2020 and 2021 CHARIS spectra (whose data can be found in Table 5.6) as well as MEC⁷, VAMPIRES, and NIRC2 photometric points. The MEC photometry and CHARIS spectra are flat in F_{ν} units except for a broad peak in H band. Formally, the SNR of HIP 5319 B in each spectral channel is extremely high (SNR

⁷Although the MEC data has a median spectral resolution $\mathcal{R} \sim 4.0$, we bin our spectral data to a single Y band photometry point for comparison with the standard photometric band.

Table 5.6. HIP 5319 B Spectra

| Wavelength (μm) | 31 July 2020 | | | 11 September 2021 | | |
|------------------------------|---------------|----------------------|-------|-------------------|----------------------|-------|
| | F_ν (mJy) | σF_ν (mJy) | SNR | F_ν (mJy) | σF_ν (mJy) | SNR |
| 1.160 | 69.197 | 1.720 | 61.8 | 67.313 | 3.076 | 52.7 |
| 1.200 | 61.625 | 1.612 | 56.7 | 65.107 | 2.733 | 55.5 |
| 1.241 | 65.347 | 1.490 | 72.5 | 65.612 | 2.589 | 55.9 |
| 1.284 | 67.608 | 1.490 | 84.7 | 68.917 | 2.856 | 47.0 |
| 1.329 | 65.402 | 1.407 | 82.0 | 70.522 | 2.650 | 58.4 |
| 1.375 | 62.158 | 1.237 | 111.0 | 55.005 | 1.672 | 78.1 |
| 1.422 | 64.376 | 1.308 | 106.7 | 66.510 | 2.453 | 77.4 |
| 1.471 | 69.387 | 1.408 | 118.3 | 67.856 | 2.367 | 93.6 |
| 1.522 | 71.086 | 1.419 | 141.1 | 71.960 | 2.449 | 85.6 |
| 1.575 | 74.866 | 1.518 | 150.9 | 74.057 | 2.521 | 71.3 |
| 1.630 | 80.553 | 1.703 | 130.8 | 83.695 | 2.936 | 76.6 |
| 1.686 | 78.667 | 1.716 | 116.0 | 78.616 | 2.461 | 81.6 |
| 1.744 | 74.888 | 1.802 | 101.4 | 81.785 | 2.802 | 78.7 |
| 1.805 | 68.638 | 1.745 | 83.4 | 65.169 | 2.220 | 66.7 |
| 1.867 | 64.104 | 1.793 | 67.1 | 63.105 | 2.272 | 74.0 |
| 1.932 | 61.858 | 1.793 | 71.5 | 68.272 | 2.493 | 117.6 |
| 1.999 | 57.205 | 1.672 | 77.4 | 62.256 | 2.167 | 94.3 |
| 2.068 | 61.378 | 1.770 | 102.2 | 64.435 | 2.329 | 89.4 |
| 2.139 | 59.341 | 1.688 | 109.1 | 62.607 | 2.294 | 87.4 |
| 2.213 | 63.070 | 1.926 | 93.5 | 65.636 | 2.625 | 83.0 |
| 2.290 | 59.136 | 3.066 | 64.8 | 61.504 | 2.955 | 87.8 |
| 2.369 | 67.170 | 12.091 | 68.5 | 67.050 | 5.416 | 53.3 |

Note. — Throughput-corrected HIP 5319 B spectra from July 2020 and September 2021 CHARIS data.

> 77). Outside of the H -band peak, consecutive wavelength channels show a “wavy” pattern, which may indicate the impact of spectrally correlated noise (see below). The two CHARIS spectra show broad agreement: due to the higher SNR for the 2020 epoch spectrum, we focus on it for subsequent analysis.

HIP 5319 B’s broadband near-IR colors ($J-H \sim 0.57 \pm 0.03$; $H-K_s \sim 0.24 \pm 0.03$) resemble those of early to mid M dwarfs (Pecaut & Mamajek, 2013). HIP 5319 B is substantially fainter than the primary in VAMPIRES 750 nm data ($\Delta m \sim 7.110$). The VAMPIRES filter does not correspond to a standard photometric bandpass with a published zeropoint flux density but lies between the Johnson-Cousins R and I bands. Adopting standard colors from Pecaut & Mamajek (2013) and R band optical photometry for the primary from the *Simbad* database, we estimate a pseudomagnitude of ≈ 13 at 750 nm.




Figure 5.6 The CHARIS HIP 5319 B spectrum (black) compared to those of field brown dwarfs (magenta) with spectral types M0, M5, and L0 from the Montreal Spectral Library binned to CHARIS’s resolution.

We compare HIP 5319 B’s CHARIS spectrum with other low-mass objects in the Montreal Spectral Library⁸ (e.g. Gagné et al., 2015). Only the CHARIS spectrum was used because the wavelength range for the Montreal Spectral Library covers *JHK*, but is rather non-uniform otherwise. Following the methods described in Greco & Brandt (2016), we find that the CHARIS spectrum shows noise that is highly spatially and spectrally correlated ($A_\rho \sim 0.69$, $A_\lambda \sim 0.22$). HIP 5319 B is best matched by an M3–M7 dwarf: earlier M dwarfs and L dwarfs fail to reproduce the CHARIS spectra, especially in the *J* and *K* bands (Figure 5.6).

Following similar analysis in Steiger et al. (2021), we compared the MEC, VAMPIRES, and NIRC2 photometry and CHARIS spectrum to the BT-Settl atmosphere models (Allard et al., 2012) with the Asplund et al. (2009) abundances and solar metallicities. We focus only on the CHARIS channels unaffected by telluric absorption and

⁸<https://jgagneastro.com/the-montreal-spectral-library/>




Figure 5.7 (Left) BT-Settl model for solar metallicity with $T=3200$ K and $\log(g)=5.5$. CHARIS spectra is shown in dark blue, VAMPIRES, MEC and NIRC2 photometry in cyan compared to the model-predicted CHARIS spectrophotometry in light green, and predicted VAMPIRES/MEC/NIRC2 photometry (dark green crosses). Although the spectrum SNR is quite high, the CHARIS data’s spectral covariance is also high, leading to a large χ^2 . (Right) Corresponding contour plots for χ^2 as a function of temperature and surface gravity. The best-fit solution is shown with a red diamond while the 1σ , 2σ , 3σ , and 5σ contours are shown in white, magenta, blue, and purple, respectively.

also remove the first CHARIS channel, whose high flux density is not reproduced in any empirical spectrum in the Montreal Library. We define the fit quality for the k th model using the χ^2 statistic, considering the spectral covariance.

Figure 5.7 shows the best-fit solar metallicity model and associated χ^2 contours. An atmosphere with a temperature of $T_{\text{eff}} = 3100\text{--}3200$ K and a high gravity ($\log(g) = 5.5$) fits the data the best⁹, although the family of solutions drawn from high gravity models ($\log(g) = 5\text{--}5.5$) at 3100 K and 3300 K and those at 3200 K and a lower gravity of $\log(g) = 4\text{--}4.5$ fall within 5σ of the best-fit model. The radii that minimize χ^2 are $3.25\text{--}3.62 R_J$, yielding a luminosity of $\log(L/L_\odot) = -1.94 \pm 0.04$. The best-fitting atmospheric models ($\log(g)=5.5$, $R_{\text{sec}} = 3.4 - 3.59 R_J$) correspond to a companion whose mass is $\sim 448 - 1675 M_{\text{Jup}}$, or $0.427 - 1.60 M_\odot$. Some of these values would be significantly higher than those for a typical M3-M7 star (Pecaut & Mamajek, 2013): potentially greater than

⁹Fits at 3100 K and 3200 K are almost numerically equivalent.

the mass of the primary itself. However, 5- σ confidence interval containing lower gravity solutions implies masses down to $44 M_J$ and includes a wider range of radii (3.25–3.62 R_J). Thus, while the temperature of HIP 5319 B is well constrained to 3100–3300 K , the companion’s poorly constrained surface gravity results in poor mass limits.

Using isochrones from Baraffe et al. (2015) we find that using the age estimate of 1.07–1.23 Gyr from the Padova and BASTI models in (Casagrande et al., 2011) and adopting the luminosity of $\log(L/L_\odot) = -1.94 \pm 0.04$ from the atmospheric models we estimate the mass of the secondary would fall between roughly 0.3 – $0.35 M_\odot$. Considering the *widest* possible range of ages of 8 Myr to 2 Gyr (the lowest end predicted by Stanford-Moore et al. (2020) and the highest predicted by Holmberg et al. (2009)) we find that the range of masses extends from $40 M_{Jup}$ to $0.35 M_\odot$. Both possible ranges include typical masses of M dwarfs from Pecaute & Mamajek (2013), while the low end of the range suggest masses down to $40 M_{Jup}$, which does not disagree with either the dynamical mass (Section 5.4.4) or the mass estimated from the atmospheric models above.

5.4.4 Orbit and Dynamical Mass

We used the open-source code `orvara` (Brandt et al., 2021) to fit for the mass and orbit of HIP 5319 B. `orvara` uses a combination of radial velocity, absolute astrometry of the primary, and relative astrometry of the companion to measure orbital parameters even when the observations of the companion only cover small fractions of an orbit.

Results Using a $1/M_p$ Prior for Companion Mass

For this companion, we used HGCA absolute astrometry measurements for the star and three epochs of relative astrometry from CHARIS, MEC, and NIRC2. There is no archival RV data for this target and so it is not included in the `orvara` fits. A

Table 5.7. HIP 5319 B Orbit Fitting Results and Priors

| Parameter | Fitted Value | Prior |
|--------------------------------|---------------------------|--------------------------|
| $M_{pri} (M_{\odot})$ | $1.397^{+0.050}_{-0.052}$ | Gaussian, 1.4 ± 0.05 |
| $M_{sec} (M_{Jup})$ | 31^{+35}_{-11} | $1/M_{sec}$ (log flat) |
| Semimajor axis a (au) | $18.6^{+10}_{-4.1}$ | $1/a$ (log flat) |
| Eccentricity e | $0.42^{+0.39}_{-0.29}$ | uniform |
| Inclination i ($^{\circ}$) | $69.4^{+5.6}_{-15}$ | $\sin i$ (geometric) |

Note. — Posterior distributions for the secondary mass and semimajor axis are both positively skewed and favor low mass, low separation distributions. The eccentricity is not well constrained using only 2 relative astrometry points and no RV data, though future astrometry for this target should serve to better constrain this value.

Gaussian prior of $1.4 \pm 0.05 M_{\odot}$ was chosen based on literature values for the primary mass (Casagrande et al., 2011), while a log-flat ($1/M$) prior was chosen for the mass of HIP 5319 B. This choice is motivated by the shape of the initial mass function for low-mass objects and planets, which says that low-mass objects are expected to occur more frequently than high-mass ones (Chabrier, 2003; Nielsen et al., 2019).

Figure 5.8 shows the posterior distributions for the primary and secondary masses along with select orbital parameters. The fit parameters are also summarized in Table 5.7. The primary mass of $1.397^{+0.050}_{-0.052} M_{\odot}$ is nearly the same as the adopted prior and the secondary mass best fit value is $31^{+35}_{-11} M_{Jup}$. The companion has a best-fit semimajor axis of $18.6^{+10}_{-4.1}$ au with an eccentricity of $0.42^{+0.39}_{-0.29}$ and inclination of $69.4^{+5.6}_{-15}$ degrees.

From the corner plot and inset in Figure 5.8 it is clear that the low-mass solutions favor less eccentric orbits at shorter semimajor axes. We also note the bimodal behavior of the distribution of eccentricities with peaks at $e \sim 0.13$ and ~ 0.81 . Continued monitoring in follow-up observations will serve to further constrain the best-fit values for the orbit of this companion as greater fractions of its orbit are observed.

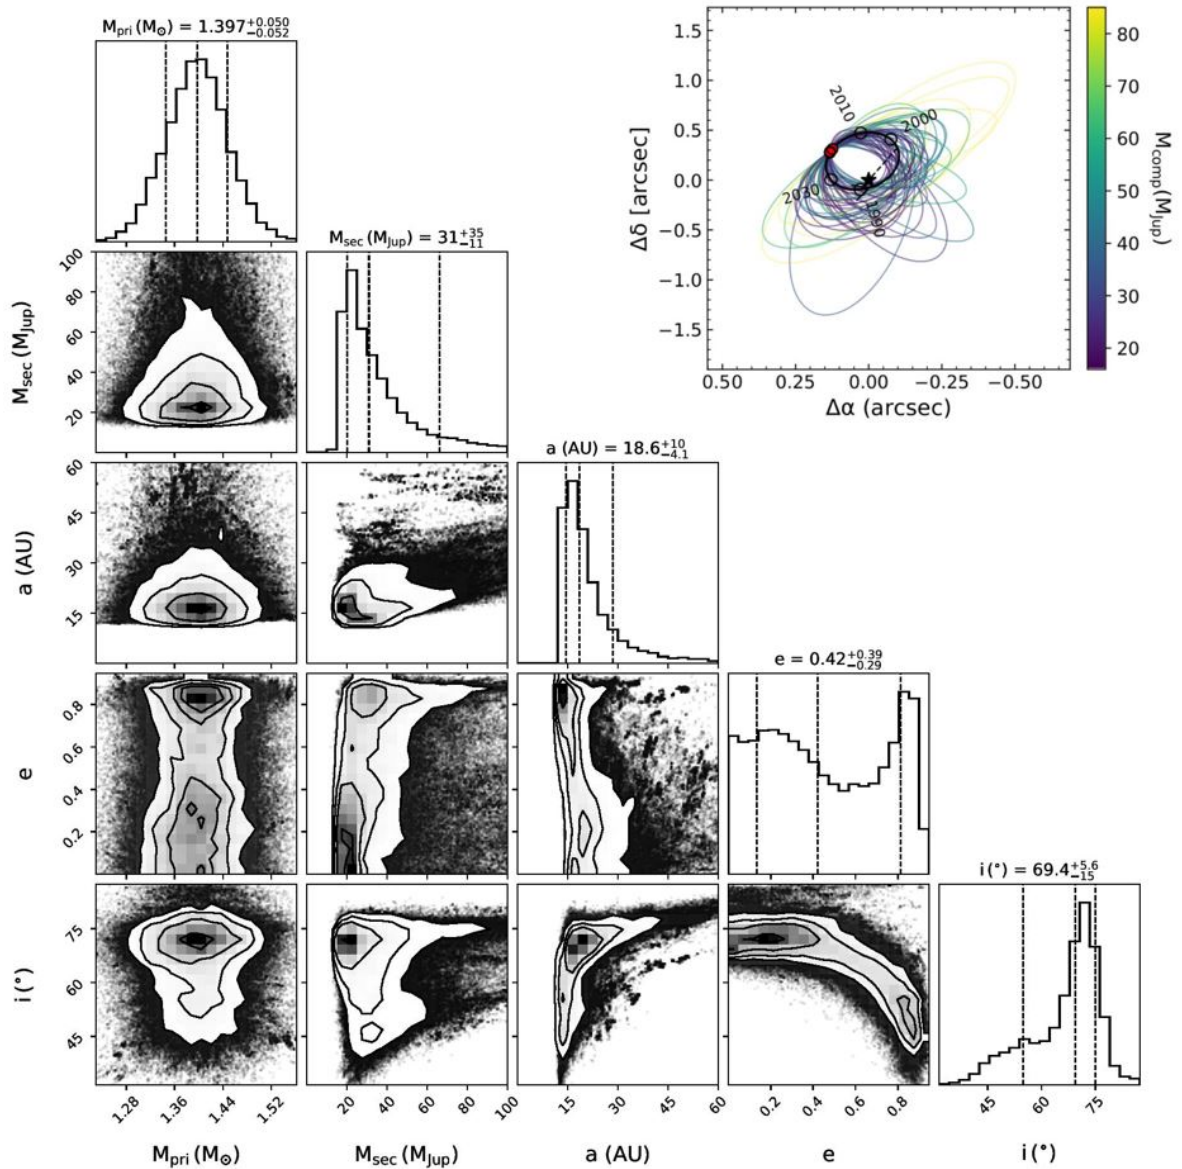


Figure 5.8 Corner plot showing the posterior distributions of selected orbital parameters using a log-normal ($1/M$) prior for the mass of the secondary companion. The orbit fits used HGCA absolute astrometry and relative astrometry from SCExAO/CHARIS and MEC data. The inset in the figure shows the best fit orbit (black) with 50 random orbits drawn from the MCMC fits color coded by the mass of HIP 5319 B. The red-colored points in the orbit represent relative astrometry points from the 3 epochs where data were taken, and the unfilled circles show the predicted location of the companion at different past and future epochs. The companion is orbiting counterclockwise.

Table 5.8. HIP 5319 B Orbit Fitting Results For Different Priors on Secondary Mass

| Parameter | log-flat (1/M) | Gaussian ($0.2 \pm 0.1 M_{\odot}$) |
|--------------------------------|---------------------------|---|
| M_{pri} (M_{\odot}) | $1.397^{+0.050}_{-0.052}$ | $1.399^{+0.051}_{-0.050}$ |
| M_{sec} (M_{Jup}) | 31^{+35}_{-11} | 128^{+127}_{-88} |
| Semimajor axis a (au) | $18.6^{+10}_{-4.1}$ | 36^{+17}_{-17} |
| Eccentricity e | $0.42^{+0.39}_{-0.29}$ | $0.33^{+0.38}_{-0.24}$ |
| Inclination i ($^{\circ}$) | $69.4^{+5.6}_{-15}$ | $75.5^{+3.9}_{-9.0}$ |

Note. — Posterior distributions for 2 different priors on the secondary mass. The priors on all other parameters being fit remain unchanged between the simulations and can be found for reference in Table 5.7.

Results Using a Gaussian Prior for Companion Mass

We have focused on the `orvara` fits using a log-flat prior for the secondary mass. However, the mass function near the hydrogen-burning limit exhibits a turnover, where lower-mass objects are less common (Chabrier, 2003). To investigate how the choice of prior may affect the posterior distribution for companion mass, we reran `orvara` using a Gaussian prior of $M_{sec} = 0.2 \pm 0.1 M_{\odot}$ ($210 \pm 105 M_{Jup}$), comparable to the implied masses for M3-M7 stars (Section 5.4.3). It is also similar to the turnover in the binary mass function from Chabrier (2003). Assuming this companion is on the main sequence, the upper limit of its mass would be $M_{sec} \sim 0.3-0.4 M_{\odot}$. This prior therefore encompasses these potential values by creating a Gaussian where the expected values of the secondary mass will fall between $+2\sigma$ and -2σ .

Table 5.8 lists the resulting best-fit posterior values; Figure 5.9 displays the corner plot showing the posterior distributions. The eccentricity and inclination distributions – $e = 0.33^{+0.38}_{-0.24}$, $i = 75^{+3.9}_{-9.0}$ degrees – agree with earlier analyses. However, compared to

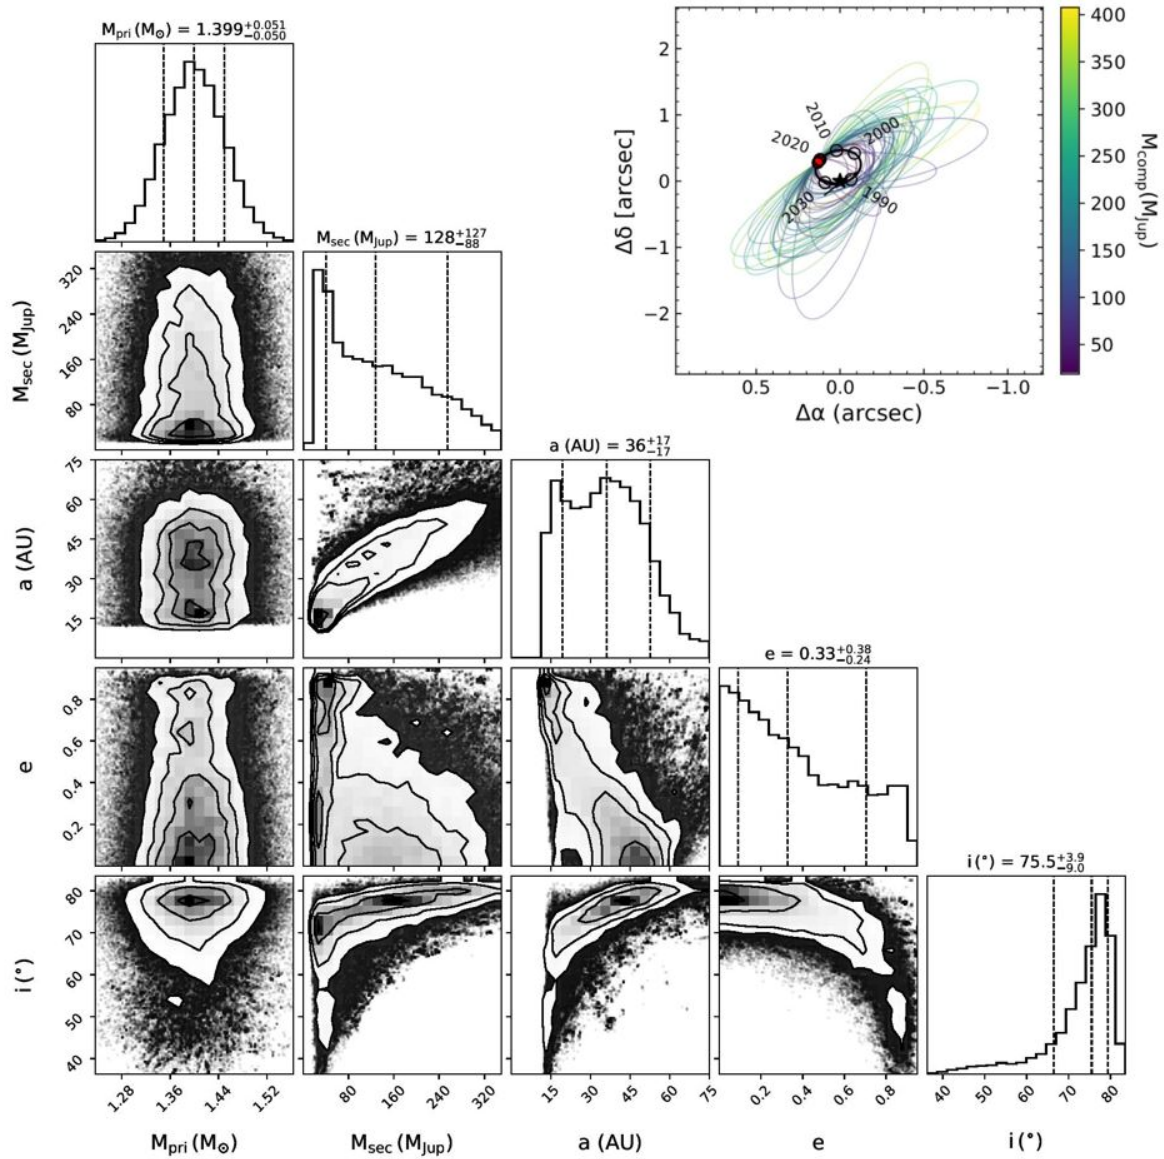


Figure 5.9 Corner plot showing the posterior distributions of selected orbital parameters using a Gaussian prior of $0.2 \pm 0.1 M_{\odot}$ for the mass of the secondary companion. The orbit fits used HGCA absolute astrometry and relative astrometry from SCExAO/CHARIS and MEC data. The inset in the figure shows the best fit orbit (black) with 50 random orbits drawn from the MCMC fits color coded by the mass of HIP 5319 B. The red-colored points in the orbit represent relative astrometry points from the 3 epochs where data were taken, and the unfilled circles show the predicted location of the companion at different past and future epochs. The companion is orbiting counterclockwise.

results for a log-normal companion mass prior, the median of the posterior distributions for HIP 5319 B’s mass and semimajor axis have shifted to larger values: $128_{-88}^{+127} M_{\text{Jup}}$ and 36_{-17}^{+17} au. For companion mass, the posterior distribution peak is $\sim 20\text{--}40 M_{\text{J}}$: comparable to values derived assuming a log-normal companion mass prior. But the posterior distribution includes a tail of far higher mass solutions, out to $\sim 350 M_{\text{J}}$, resulting in a far larger median value. The semimajor axis posterior distribution contains two peaks – one near 18 au and a second near 35-40 au.

In practical terms, our analyses cannot conclusively clarify whether HIP 5319 B is a brown dwarf or a low-mass star. Dynamical modeling assuming a log-normal companion mass prior favors a brown dwarf at 18.6 au, while modeling adopting a gaussian prior admits a much wider range of companion masses, including those on both sides of the hydrogen burning limit. The implied masses from masses from atmospheric modeling admit a wide range of possible values: $44 M_{\text{J}}$ to $1675 M_{\text{J}}$. However, the orbit insets to Figures 5.8 and 5.9 suggest that future astrometric monitoring of HIP 5319 B should clarify the companion’s nature.

5.5 Summary and Discussion

Spectroscopy from SCExAO/CHARIS and photometry from SCExAO/VAMPIRES, Keck/NIRC2, and SCExAO/MEC have enabled the identification of a candidate substellar companion to the young F5IV star HIP 5319. Comparisons of the SCExAO/CHARIS spectra to the spectra of objects in the Montreal Spectral Library show this companion to be best matched with M3-M7 dwarfs, with earlier-type M and L dwarfs failing to match the CHARIS spectra measured in *J* and *K* bands. By combining measurements from *Hip-parcos* and *Gaia* with our relative astrometry from CHARIS/MEC/VAMPIRES/NIRC2 we can constrain the dynamical mass and orbit of HIP 5319 B.

Assuming a log-normal prior, we find a dynamical mass of $31_{-11}^{+35} M_{\text{Jup}}$ for the companion, suggesting that HIP 5319 B is a brown dwarf. The posterior distributions from the fits for dynamical mass show a bimodal distribution in possible eccentricity values, where high-eccentricity solutions are favored at more edge-on inclinations and low-eccentricity solutions are favored for more inclined orbits. However, adopting a Gaussian prior for the companion mass yields a higher mass of $128_{-88}^{+127} M_{\text{J}}$ which favors the interpretation of the companion as a low mass star although the distribution’s peak still falls in the substellar range. Future RV measurements, relative astrometry from direct imaging, and more precise astrometry from *Gaia* will contribute to further constraining this companion’s mass and orbital parameters, providing deeper clarity on this companion’s identity.

Atmospheric models of the companion best fit to an atmosphere with solar metallicity at $T=3200$ K with a surface gravity $\log(g)=5.5$, though solutions with comparably good fits exist with temperatures from 3100 K to 3300 K and lower surface gravities ($\log(g)=4-4.5$). Best-fit models show radii between $3.25-3.62 R_{\text{J}}$ and $\log(L/L_{\odot})=-1.94 \pm 0.04$. The mass inferred from atmospheric modeling is poorly constrained.

This work highlights the need for an updated inventory of system measurements when interpreting companions imaged around accelerating stars. While older data suggested that HIP 5319 is a RS CVn (short-period) binary, our RV data rule out stellar companions with an orbital period less than 30 days whose presence would affect our conclusions about HIP 5319 B’s mass and orbital properties. Similarly, our HIP 5319 spectra found no evidence for Ca HK emission that could reveal evidence of HIP 5319’s youth. Other system measurements whose values may impact derived companion masses and orbits include spectral type/luminosity, rotation rate, lithium abundances, x-ray activity, etc.

Finally, this work demonstrates the importance of priors in dynamical models used to estimate companion masses and orbits from direct imaging and astrometry. When a small fraction of a companion’s orbit has been observed - as is the case with HIP 5319 B -

the selection of prior for a given parameter may influence the final shape of the posterior distributions and the reported values of the dynamical mass and orbital parameters. The chosen prior should not cause the fitted values to change significantly (see also Currie et al., 2022b). Performing multiple fits for orbital parameters using disparate priors (e.g. Gaussian, log-normal, uniform, geometric, depending on the parameter of interest) can confirm that the extracted masses and orbital parameters are robust. If the results from multiple fits are in good agreement with one another - the values within the 95% or 68% confidence interval overlap with one another, for example - one may say conclusively that the derived dynamical mass is robust. Otherwise, the data are not sufficiently constraining: more of the orbit must then be observed before one can make a definitive claim regarding the fitted orbital parameters and masses of the system.

This direct imaging detection was - in part - made due to the identification of the system as having statistically significant astrometric acceleration in the HGCA. Previous works which include – but are not limited to – Brandt et al. (2019), Kervella et al. (2019), Currie et al. (2020), Bonavita et al. (2020), Bowler et al. (2021), Chilcote et al. (2021), Li et al. (2021), Steiger et al. (2021), Currie et al. (2022b), Kuzuhara et al. (2022), Miskovetz et al. (2022), and Salama et al. (2022) have also used the HGCA to select targets that have been found to host previously unidentified companions. This discovery further demonstrates the efficacy of using astrometry to select direct imaging targets instead of conducting blind searches. As more HGCA targets are observed, future *Gaia* data releases yield more precise astrometry, and direct imaging capabilities improve, this survey approach will only become more powerful in discovering substellar companions, including numerous planets (Currie et al., 2021).

Chapter 6

An Algorithm for Cosmic Ray Rejection

6.1 Cosmic Rays

Generally speaking, cosmic ray events occur when a high energy particle strikes an astronomical detector and causes non-astrophysical counts to be measured. In practice, a cosmic ray is likely to be something like a muon, proton, α particle, or other exotic heavy particle that come from the sun, galactic and extragalactic sources, or are *secondary* particles that come from initial high energy particles interacting with Earth's atmosphere and causing an "air shower" of heavy particles raining down toward Earth. Regardless of their nature or provenance, they are sources of noise in astronomical images that must be well understood and – if possible – removed. This is especially important for very low count rate experiments which require exquisite understanding of the environmental background so it may be calibrated out without removing whatever signal may exist.




Figure 6.1 Reproduced from https://www.eso.org/~ohainaut/ccd/CCD_artifacts.html. (Left) A raw CCD image showing several stars (trailing and smeared out in this image) with several cosmic rays (sharp white dots/lines). The only real object that is not trailing or elongated is circled in red. The white square shows the rough extents of the image on the right. (Right) Detail from the previous image on a smaller scale. Each of the white dots/lines are cosmic ray events.

In CCDs

These cosmic ray events have been known to plague astronomical detectors such as CCD-, EMCCD- and CMOS¹-based cameras (Niedzwiecki et al., 2019; Harun et al., 2020). In semiconductor cameras, one of these high energy particles will hit the atoms in the semiconductor chip itself and knock into their electrons, giving them energy and liberating them from the potential well they live in. As these electrons are excited they will be counted as part of the astronomical signal, causing bright spots or lines across the semiconducting camera that contaminate the images. An example of a CCD camera with cosmic rays is shown in Figure 6.1.

The problem with cosmic ray contamination in semiconducting cameras is that once they strike the detector and read out, they are part of the image. While they may be identified and masked, this process can also remove real astronomy sources of interest.

¹Complimentary Metal-Oxide Semiconductor




Figure 6.2 A timestream showing a cosmic ray event on a 20,440 pixel MEC-style array over $\sim 250\mu\text{s}$. At $T=0$, the particle strikes the array, depositing energy which spreads through the array and excites resonators as it does. Eventually, the energy is either absorbed by the resonators it passed or passes out through the array substrate and the array returns to quiescence.

In MKID Cameras

In contrast to semiconductor cameras that require traditional exposures where the camera is left recording photons for a certain amount of time before reading out all the photons that struck in that time, MKID cameras' photon counting ability and timing resolution enables a unique power to identify and remove cosmic rays from MKID datasets. When a cosmic ray is incident on an MKID array it excites phonons in the detector substrate. As the energy is carried outward will cause a majority of the pixels to register a photon in a very short duration. Figure 6.2 shows frames from an animation of a cosmic ray striking a 20,440 pixel array. Starting when the cosmic ray strikes one can see that as the phonons carry the energy outward across the array each resonator is triggered once before the energy dissipates and returns to quiescence roughly $150\mu\text{s}$ later.

The ability to visualize these cosmic ray events in such a way may provide unique measurements of the phonon sound speed in different materials *and* characterization of the natural cosmic ray background, which is of great interest to the quantum computing community. Additionally, the behavior of the MKID array during a cosmic ray event also helps to identify and fully remove these non-astronomical contaminating sources. This enabled the development of an algorithm to find and remove cosmic rays in MKID




Figure 6.3 A photon timestream showing the count rate across the MKID array in broadband without removing any measured photons based on their wavelength (blue) and over a typical wavelength band ($\lambda=800\text{-}1375$ nm) observed by an MKID instrument (orange).

datasets which has been incorporated into the MKID Science Data Pipeline described in Steiger et al. (2022a) and will be detailed in this chapter.

6.2 Identifying Cosmic Rays

Figure 6.2 shows that a frame-by-frame animation of the MKID array timestream enables easy identification of a cosmic ray event when a “fireball” spreads out as the energy from a cosmic ray travels through the array and causes the pixels to light up. However, this is impractical since it requires a slowed-down animation and is not automated.

The cosmic ray rejection algorithm that was developed for automatic identification comes in two flavors for determining the threshold of counts across the array needed to be considered a cosmic ray event. The first uses the Poisson statistics count rate over the array as a function of time while the other calculates a sensible measure of the variation in the timestream, both with the goal of using a statistical metric to determine how many




Figure 6.4 The narrow-band ($\lambda=800\text{-}1375$ nm) timestream from Figure 6.3, shown with red arrows pointing out time bins with more counts than the Poisson threshold.

counts in a given time constitute a cosmic ray event to be removed.

The “Poisson Method” is optimized for use with narrow band data where the user has made a wavelength cut, meaning photons outside of the calibrated bandpass are not included. The “Peak Finding Method” is optimized for broadband data where no wavelength cut has been made. Although both have similar ability to identify cosmic ray events with low false-positive and false-negative rates, the peak-finding method takes more time and computational power since it has to analyze many more photons.

Each method begins by splitting the observation of interest into $\sim 10\mu\text{s}$ time bins. The reason for choosing this duration is that they must be sufficiently short so the number of counts across the array follows Poisson (counting) statistics. A typical MKID array timestream that includes cosmic rays is shown in Figure 6.3.

Narrow band: Poisson Method

This approach operates under the assumption that the count rates across the MKID array follow Poisson statistics. First, a cumulative density function (CDF) is determined




Figure 6.5 The broadband timestream from Figure 6.3, shown with red arrows pointing out where time bins with more counts than the Peak-Finding threshold are found.

by defining the number of standard deviations away from the average that a given count rate must be so a time bin may be considered to have a cosmic ray. Next a Poisson percent point function (PPF) is evaluated on the original CDF using the mean count rate as the expected value, which will return the appropriate threshold value. Above this threshold a time bin may be classified as containing a cosmic ray. Figure 6.4 shows the narrow-band timestream from Figure 6.3 with arrows pointing out each of the bins identified using the Poisson method as cosmic ray events.

Broadband: Peak Finding Method

The second method does not require that the statistics of the count rate be Poisson. The so-called Peak-Finding method first finds value of the standard deviation of counts across the array, excluding any outliers that fall above or below 3 standard deviations from the mean. The threshold is then defined as being N standard deviations above the mean number of counts in the timestream. The number of standard deviations N is a user input parameter and can be tuned according to the dataset. A higher N is less

sensitive and removes fewer cosmic rays and a lower N is more aggressive and identifies more cosmic rays. However, a value of N that is too low may identify small spikes in the noise as events and lead to a more significant portion of the observation being removed. Figure 6.5 shows the broadband timestream from Figure 6.3 with arrows showing the time bins identified as having cosmic ray events from the Peak-Finding method.

Compared to the Poisson method this will tend to find more cosmic ray events, as some of them disappear from the timestream after making a wavelength cut(Figure 6.6). This may be due to them being low energy cosmic rays that do not excite many resonators or different noise sources altogether. In the second case their origin does not strictly matter what will be identified and can be removed from the broadband data and they do not appear at all in the narrow-band data. If, however, they aren't actually cosmic rays and their nature is fundamentally important to the analysis of the MKID data then it may be useful to explore their origins in greater detail. This idea will be explored further in Chapter 7.

Regardless of the method that was used to determine the cosmic ray threshold the next step of the process, identification of the cosmic rays. To identify which bins contain cosmic rays, a simple peak-finding algorithm from the `scipy.signal` package is used.

6.3 Cosmic Ray Rejection and Removal

The final step in this process is to actually flag and remove the time ranges affected by the cosmic rays. To do this, the affected intervals are recorded in addition to the total and average number of counts across the array, and the peak count rates measured during the cosmic ray event. This data is then stored as associated metadata with the rest of the observation dataset.

The time interval over which a cosmic ray interacts with MKIDs in the array is




Figure 6.6 The timestreams from Figure 6.3 now including the cosmic ray events identified by each method. Those found by the Poisson method are shown in red and by the Peak-Finding method are in green. There is good agreement between the two. Most events that appear in the narrow band timestream are also found in the broadband data.

determined by the energy of the cosmic ray and the phonon sound speed in the substrate. For this reason, there is a user defined removal time which can be changed when looking at cosmic rays measured by under different circumstances. The default range to remove is $50 \mu\text{s}$ prior to the peak of the event and $100 \mu\text{s}$ after. In an average MKID dataset, this will lead to the removal of $\lesssim 0.1\%$ of the duration of the dataset in question, leaving the majority of data still usable and now uncontaminated by these noise sources.

As opposed to CCD and CMOS cameras where a cosmic ray will bounce around the substrate and excite many electrons that jump their potential wells which causes it show up as a very bright spot in an image, a cosmic ray in an MKID camera will result in 1 photon per pixel as the wave of energy passes through it. This is also useful because – in addition to the aforementioned algorithm for removing them – if a cosmic ray is *not* removed then it will only contribute a single extra photon count to the dataset. During observations with many thousands of counts, a missed cosmic ray event will not contribute

significantly to the noise background. However, in low-count rate environments such as dark matter direct detection or X-ray astronomy experiment this will become a problem and so it will become more and more important to be able to accurately identify and remove cosmic rays when they occur. The following chapter (and Swimmer et al. (2023)) will go into further detail of the characterization of MKID arrays in low-count rate environments, which relies heavily on the removal and understanding of the measured cosmic ray background noise sources.

Chapter 7

Characterization of the Dark Count Rate in MKID Arrays

7.1 Attribution

The work in this chapter was originally published in Swimmer et al. (2023) and has been modified for content and flow in the context of this thesis.

7.2 Introduction

Current generation MKIDs are designed to be sensitive to different energy ranges such as for ultraviolet, optical, and near-infrared (UVOIR) astronomy and X-ray detection. They also offer straightforward ways to tune their sensitivity to higher- or lower-energy bandpasses. This ability opens opportunities for MKIDs to be used as detectors for new physics applications such as the search for dark matter. In regimes where sources emit very few photons it is of high importance to characterize the performance of the detectors to ensure that each photon detected is ‘real’, i.e. the photon comes from the observed

source rather than errant sources such as blackbody radiation from inside a cryostat.

In this experiment we aim to characterize “dark counts” measured by a large-format MKID array. Dark counts are events registered as photons by the detector when it is not exposed to a light source. These dark counts differ from those of conventional semiconductor detectors because of their origin. Semiconductor detectors register false counts due to dark current and read noise. Dark current is the generation of thermal electrons in the material that are captured by the detector’s potential well and counted as part of the signal while read noise is the noise that is added to the measured signal from charge photon-to-voltage conversion and signal processing such as analog to digital conversion. In MKIDs false triggers may stem from noise in the room-temperature readout (Fruitwala et al., 2020), blackbody photons from the environment, or more complex sources.

7.3 MKID Photon Measurement

MKID operation was discussed at length in Section 1.3. Here we will not rehash their operation principle but will go into further detail of how photons are measured.

For an MKID to detect a photon, the photon event must cause the phase of the detector to increase beyond a minimum threshold. For a given pixel this threshold is calculated by measuring the phase (in radians) while not illuminated. This allows a phase noise to be measured, which we find does not typically exceed $\sigma_\phi \sim 0.15$ radians in good-performing pixels. The threshold for the given resonator is then set to be $6\sigma_\phi$ away from the average phase value – typically 0 radians. Mathematically,

$$\phi_{\text{Threshold}} = \hat{\phi} + 6\sigma_\phi \approx 6\sigma_\phi \quad (7.1)$$

where the right hand approximation holds when $\hat{\phi} \approx 0$. This is the method by which

each pixel’s photon detection threshold is calculated for both electronic readout systems described in this paper (Sections 7.5 and 7.6). Since the phase noise is Gaussian and a $6\sigma_\phi$ threshold is used, that means that there is less than a 1 in 500,000,000 chance that the phase will fluctuate above the threshold during any phase measurement. The process for determining the threshold for a pixel is described in Steiger et al. (2022a). Each pixel’s phase response to photons of different energies is also measured when the array is calibrated (Sections 1.3.4 and 7.4.3). For the MKID resonators in this experiment we find that the typical phase response for the lowest-energy photons in the bandpass (0.946 eV, $\lambda=1310$ nm) is $\phi_{0.946\text{eV}} = -1.4 \pm 0.2$ radians and the typical phase response for the highest-energy photons (1.534 eV, $\lambda=808$ nm) is $\phi_{1.534\text{eV}} = -2.2 \pm 0.3$ radians. This demonstrates that the phase noise is unlikely to ever swing sufficiently high to cause a resonator to trigger on an event within the calibrated bandpass.

It is also instructive to calculate the signal-to-noise ratio (SNR) for photon pulses seen by each resonator. The resolving power $\mathcal{R} = E/\Delta E \sim E/2.355\sigma_\phi$, where E is the energy of a photon event and is proportional to the phase pulse height, meaning that the $\text{SNR} \sim E/\sigma_\phi$ (Section 7.4.3). Rearranging, one finds that the $\text{SNR} \sim 2.355\mathcal{R}$. This means that for a resonator with an $\mathcal{R} = 4.5$ – a typical value in this experiment – the SNR for a photon pulse will be about 10.6, again demonstrating that a resonator’s phase response to a photon will be significantly higher than the phase noise itself.

A single photon event is shown in Figure 7.1. When a photon is measured the readout records the time the photon struck and the height of the peak of the pulse. The height of the pulse is related to the energy of the incident photon. By performing an energy calibration with lasers of known wavelength, these measured pulse heights can be converted to energies of the incident photons within the calibrated range.

Figure 7.2 shows several types of events measured by a single detector that can contaminate a dataset. The first occurs when noise in the system causes the readout




Figure 7.1 A single photon event measured by an MKID pixel with readout sampling at 0.8 MHz. The x-axis shows time in microseconds and y-axis shows the detector phase response measured in radians.




Figure 7.2 Two possible of ‘bad’ photon events from the MKID readout. (Left) A phase timestream taken when noise in the data caused the readout to trigger as if a photon hit the detector although clearly none did. (Right) A second photon arriving on the tail of a previous one. This contaminates the analysis of the first and, depending on the dead time, the readout may not trigger on it.

to register a photon when there is none present. This may occur when there is a slight spike in the phase data or a very low energy photon hits the detector. These noise triggered events are removed in post-processing as their peak pulse heights fall outside of the calibrated energy bandpass and will be removed when cutting photons outside that range. The second is when multiple photons are caught riding on the tail of the initial photon registered by the readout. This can be mitigated by decreasing the detector dead time. In the digital readout (Section 7.5) the first photon is counted while the second (and beyond) photon is ignored completely, while in the analog readout (Section 7.6) one can manually calibrate these by splitting up timestreams that have multiple photon events, manually including both if desired or treating them similarly to the digital readout and not registering the later arriving photons.

7.4 Experiment Overview

7.4.1 Optics

The optical path to put light onto the MKID array is relatively simple. One of five lasers within the sensitivity range of the devices (808 nm, 920 nm, 980 nm, 1120 nm, 1310 nm) can be inserted into an integrating sphere which has an output to a multimode fiber at room temperature. The fiber is then inserted into a port in the dilution refrigerator where it is routed to directly to a collimator at the same temperature as the MKID device. The fiber has transmission greater than 90% at all of the laser wavelengths above. The collimator is oriented so the collimated light shines directly onto a microlens array (MLA) which focuses the light onto the photosensitive inductor of each MKID pixel. The collimator creates a spot greater in size than the MKID array, meaning that the distribution of light from the lasers is spatially uniform across the MKID pixels. A




Figure 7.4 A 10-feedline MEC-style MKID array with 20,440 pixels. The lid with microlens array has been removed so that the MKID array itself can be seen. The MKID digital readout (Section 7.5) combined with the dilution refrigerator setup (Figure 7.3) allow for up to 1 of the microwave feedlines (2044 MKID pixels) to be read out at a time. In an MKID instrument all 10 can be read out simultaneously.

resonator. Figure 7.4 does not show either the lid+MLA so that the array itself may be seen for clarity or the fiber collimator designed for this experiment.

7.4.3 Array Calibration

In this experiment, only half of a single MKID array feedline was used. This was due to the available microwave RF ports in the BlueFors LH Dilution Refrigerator and the bandpass of the Traveling Wave Parametric Amplifier (TWPA, Section 7.5.1). 819 MKID pixels were initially identified before any data collection. An energy calibration dataset was taken prior to and after data acquisition to assess the stability of each pixel's response to photons of a given energy.

We require that each resonator was successfully energy calibrated in both datasets. This means that each pixel was marked by the energy calibration software as being successfully calibrated at each of the 5 laser energies. This cut removed 40 pixels, leaving 779 (95.1%) of the original 819. This cut was made for two reasons. The first being

Table 7.1. MKID Resolving Powers

| Energy (eV) | Wavelength (nm) | Median \mathcal{R}_1 | Median \mathcal{R}_2 |
|-------------|-----------------|------------------------|------------------------|
| 0.946 | 1310 | 4.5 | 4.9 |
| 1.107 | 1120 | 4.4 | 4.5 |
| 1.265 | 980 | 4.4 | 4.4 |
| 1.348 | 920 | 4.4 | 4.5 |
| 1.534 | 808 | 4.6 | 4.5 |

Note. — The median resolving power of all resonators measured at each calibration laser energy. Columns 1 and 2 show energy and wavelength of each laser. Columns 3 and 4 show the median resolving power of all resonators taken prior to and after data collection.

that if the MKID Data Reduction Pipeline (Steiger et al., 2022a) is not able to identify a resonator in one of the two calibration datasets, we cannot confirm that it stayed stable through the duration of the data collection. The second is that if a resonator was identified in both calibration datasets, it must be successfully energy calibrated at all 5 laser energies. We have seen that resonators that only pass 3 or 4 of the laser energies are worse performing and typically see many low-energy ‘noise triggered events’ that make such pixels unreliable. The cuts led to 14 and 26 of the 40 pixels removed, respectively. Table 7.1 shows the median \mathcal{R} values for each calibration dataset. The best performing pixels that remained after these cuts are those whose data will be used for the dark count analysis in Section 7.5.2. In Section 7.6 a single one of these resonators will be used to characterize the nature of the photons that are seen.

7.4.4 Electronic Readouts

This investigation used 2 separate readout systems. The first readout system is the Gen2 MKID Digital Readout Fruitwala et al. (2020) to read a substantial number of

MKID pixels simultaneously. The second system uses an analog readout scheme similar to that described in Zobrist et al. (2019) and Zobrist et al. (2022) to read out a single MKID pixel. This system ensures lower readout noise and the ability to record photon phase timestreams. We review the former in Section 7.5 and the latter in Section 7.6.

7.5 Digital Readout: Many-detector Measurement

7.5.1 Experimental Setup

The full-array readout in this experiment used the second generation MKID digital readout (Fruitwala et al., 2020) which is divided between 3 temperature stages where most of the large electronics are at room temperature. The remainder were housed in a BlueFors Dilution Refrigerator which cools them to 4K or 100 mK, depending on the component. The internal schematic for the fridge is seen in Figure 7.3.

The room-temperature components are discussed in detail in Fruitwala et al. (2020) and Section 2.7.2. Here we note that although each readout cartridge contains two sets of readout boards for reading out MKIDs from 4-6 GHz and from 6-8 GHz, for this experiment we chose to only read out the higher-frequency half of an MKID feedline from 6-8 GHz to best match the bandpass of the parametric amplifier.

Two 4-meter RF coaxial cables are then attached between the readout cartridge and the dilution fridge to send the signal to the MKID array and carry the output back to the readout cartridge to be read out.

Internally, the signal is sent from room temperature to 100 mK using cryogenic RF coax cables that are heat sunk at intermediate temperature stages to reduce heat flow to the MKIDs. At the 4K stage there is a 20 dB attenuator to attenuate Johnson noise from the room temperature input. A second 20 dB attenuator is added at the 100 mK stage

to further reduce Johnson noise from the 4K stage and to account for the amplification of the signals on the output side of the array. The probe tones are then sent through the MKID device, exciting the resonators of a single MKID array microwave feedline. On the output side the signal is sent through a traveling wave parametric amplifier (TWPA), a wideband, high-power amplifier capable of reading out photon events with near quantum-limited amplifier noise (Eom et al., 2012; Zobrist et al., 2019). The signal is further amplified by an LNF HEMT amplifier at 4K. The microwave signal is then sent routed the room temperature readout cartridge so the MKIDs may be measured.

Using the readout to monitor half of one MKID feedline enabled up to 1024 pixels to be read out during this experiment. The number of functioning pixels compared to the total number of possible pixels is called the pixel yield. Current-generation MEC arrays have yields of about 80%. Each pixel in an MKID array is unique and can be rendered non-functioning for different reasons such as a piece of dust or residue landing on a pixel and shorting it to ground. Variation in the thickness of the resonators or their chemical composition can also cause them to move to unpredictable frequencies, colliding with another resonator in frequency space and making one or both unusable (Walter et al., 2020). The pixel yield of about 80% leads to 819 of the possible 1024 resonators being read out. These pixels will form the basis of our analysis of the dark count rate measured by the MKIDs. We note that this was the first simultaneous readout of many resonators on a large MKID array with a parametric amplifier that we are aware of.

7.5.2 Data Collection and Reduction

Using the Gen2 readout (Fruitwala et al., 2020) each resonator from half of an MKID feedline collected data with no light on the detectors for 86,750 seconds from 10-13 December 2020 while the dilution fridge regulated the device temperature at 100 mK.




Figure 7.5 A timestream showing the number of counts across the 819 resonators using the Gen2 Readout over a 300-second span. The quiescent count rate is approximately 0 counts/second, but cosmic rays hitting the array and other potential system noise cause flashes where many pixels register a photon. The inset shows a single cosmic ray event (marked by the red arrow) for 200 μs before and 600 μs after the peak.

Notably, the data taken using the digital readout is useful because it allows many resonators to be read out simultaneously. This enables useful characterization of the data that can only be inferred via bulk properties of the array such as identification and rejection of cosmic ray hits and ‘flashes’ across the device.

The 86,750 seconds of data taken using the MKID Digital Readout were processed and reduced using the MKID Science Data Pipeline (Steiger et al., 2022a). The data were split into smaller chunks of time in order to create photon lists of manageable size.

During this reduction, a cosmic ray calibration (Chapter 6) is performed by making a time-based coincidence veto based on significantly more photons being detected across the array in a short burst above the quiescent count rate. A full-array timestream with many cosmic events can be seen in Figure 7.5. To prevent any possible contamination

from cosmic events, 5000 μs and 10000 μs are removed from before and after each event, respectively (the asymmetrical nature is due to there being a ‘tailing off’ behavior while the energy spreads through the array). This veto leads to a short time range surrounding each cosmic ray event being removed, ultimately removing $\sim 0.3\%$ of the total time each resonator was taking data.

For the last quality cut we remove any “hot” pixels remaining. Qualitatively, a pixel is considered hot if it counts significantly more photons than other pixels after cosmic ray rejection. These pixels may become hot for reasons such as having a very low phase threshold, being under- or overpowered, electrical noise in the readout electronics causing the probe tone for that individual resonator to become noisy, the resonator being adversely affected by local electrical or magnetic fields and shifting slightly in frequency, causing the resonator to go out of calibration, or other more complicated phenomena leading to a specific resonator triggering more frequently than its neighbors. The hot pixels cause the distribution of the total number of photons measured per resonator to be highly positively skewed. To properly characterize the shape of this distribution and catch hot pixels, a metric that is robust to outliers must be chosen. The median of the number of counts (\hat{c}) from each resonator is taken to be the expected value of the distribution and the spread is measured using the `astropy mad_std`¹ function to calculate a robust standard deviation ($\sigma_{c,MAD}$) using the Median Absolute Deviation (MAD). Equation 7.2 shows that a pixel is “hot” if the number of counts c it sees is greater than the median number of counts measured plus 15 times the MAD standard deviation.

$$c \geq \hat{c} + 15\sigma_{c,MAD} \quad (7.2)$$

This cut leaves 590 of the 779 pixels that remained from Section 7.4.3. Ultimately

¹https://docs.astropy.org/en/stable/api/astropy.stats.mad_std.html




Figure 7.6 The average count rate values in photon counts per pixel per second in each energy bin with 1σ error bars measured from the ensemble of count rates from all remaining pixels.

this represents a cut of 24.3% of the pixels that were energy calibrated at all 5 laser energies (or 26.7% of the initial 819 pixels).

The median resolving power of the remaining 590 MKID pixels across all wavelengths is $\mathcal{R} \sim 4.54$. For each pixel the photons which fell between the calibrated energy values and were not removed by the cosmic ray calibration were divided into equal-width energy bins and the count rate of photons per second was then calculated at each energy. The measured count rates can be seen in Figure 7.6 and Table 7.2.

7.5.3 Analysis

In Section 7.5.2 the data collection and reduction was discussed in detail. After the final subset of pixels was determined the dark count rate could be calculated. To do this, the count rate in each pixel was measured and the error on each value calculated using Poisson statistics. The energy bins were determined by the calibrated bandpass (0.946-1.534 eV, or 1310-808 nm) and the median resolving power $\mathcal{R} \sim 5$. This resulted in 5 energy bins of 0.118 eV centered at 1.005, 1.123, 1.240, 1.358, and 1.476 eV.

The measured count rate for the unilluminated MKID array ranges from $(2.73 \pm 0.02) \times 10^{-4}$ photons/pixel/s at 1.476 eV to $(6.26 \pm 0.04) \times 10^{-4}$ photons/pixel/s at 1.005 eV after cosmic ray rejection and other cleaning steps, respectively. This corresponds to an MKID pixel seeing a low energy photon roughly every 1600 ± 10 seconds and a high energy photon every 3660 ± 30 seconds. Over the full bandpass the count rate is measured at $(1.847 \pm 0.006) \times 10^{-3}$ photons/pixel/s. In flux units, this is $(3.14 \pm 0.01) \times 10^{-3}$ photons/pixel/s/eV or $(3.68 \pm 0.01) \times 10^{-6}$ photons/pixel/s/nm.

Although the dark count rate measured by an MKID pixel is not directly analogous to the dark current measured by CCDs and EMCCDs (Electron-Multiplying Charge Coupled Devices) due to the differing origin of the events it is useful to compare the event rates seen by each. Above, we reported that the number of dark count events measured per MKID pixel is $(1.847 \pm 0.006) \times 10^{-3}$ photons/pixel/s. State-of-the-art CCDs have measured dark count rates of 1.66×10^{-3} electrons/pixel/s (Castelló-Mor, 2020) while EMCCDs have measured dark current rates as low as 1×10^{-3} electrons/pixel/s (Daigle et al., 2012). Assuming that the CCD and EMCCD gain (the ratio of photons needed to generate 1 electron in the CCD/EMCCD) is 1 – i.e. 1 electron in the detector corresponds to 1 photon event – then the dark current generation in these state-of-the-art CCDs and EMCCDs are comparable to the dark count rates measured by an MKID pixel.

Table 7.2. Dark Count Rates With and Without Reduction Steps

| Bin Center (eV) | With Reduction Steps | | No Reduction Steps | |
|--------------------|--|---------------------------|--|---------------------------|
| | Rate ($\times 10^{-4}$) (photons/pixel/s) | Total Counts (photons) | Rate ($\times 10^{-3}$) (photons/pixel/s) | Total Counts (photons) |
| 1.005 | 6.26 ± 0.04 | 54.1 ± 0.3 | 5.0 ± 0.1 | 436.7 ± 0.9 |
| 1.123 | 3.68 ± 0.03 | 31.8 ± 0.2 | 3.83 ± 0.09 | 332.2 ± 0.8 |
| 1.240 | 3.06 ± 0.02 | 26.5 ± 0.2 | 3.17 ± 0.08 | 275.2 ± 0.7 |
| 1.358 | 2.75 ± 0.02 | 23.7 ± 0.2 | 2.73 ± 0.07 | 237.1 ± 0.6 |
| 1.476 | 2.73 ± 0.02 | 23.6 ± 0.2 | 2.44 ± 0.07 | 211.6 ± 0.6 |

Note. — The average count rate per pixel and total number of counts per pixel for the data with the different calibration and cleaning steps (see also Figure 7.6) compared to the same quantities without. Without reduction steps, each pixel took data for 86,750 seconds. After cleaning, each pixel was left with 86,491 seconds of data.

Comparison to Raw Data

In Section 7.5.2 the data reduction algorithm was discussed. This includes the removal of photons from cosmic ray events and excluding ‘hot’ pixels that live in a state where they count significantly more photons than their physical neighbors. To show the improvement these cuts bring we also calculate the dark count rate with none of the removal steps performed. The results are shown in Table 7.2.

Prior to any data cleaning each resonator saw an average of 1493 ± 2 counts across the calibrated bandpass. Afterwards this number was reduced to 159.5 ± 0.5 , an improvement of nearly a factor of 10. While we were able to reduce the number of counts by almost a factor of 10 for each energy band, the time cut was not particularly aggressive and only about 0.3% of the duration from each pixel was removed via the cosmic ray cuts. This shows that the majority of the counts seen come from spurious events that can be calibrated out without removing significant time from the dataset in question.

Potential Photon Sources

The steps taken during the data reduction in Section 7.5.2 aimed to mitigate effects from noisy pixels and well-understood noise sources such as cosmic ray events. However there is still the possibility that there are more complex or uncalibrated sources that are not well characterized in this system. This can may include noise from cryogenic amplifiers or room temperature readout electronics, secondary photons from cosmic rays causing fluorescence in the fiber optic cable, and simple blackbody radiation.

Work has been done to characterize the noise characteristics of the MKID Digital Readout (Fruitwala et al., 2020) although it is not well understood how electrical noise in the system translates to spurious triggers on non-photon events. However, well-behaved MKID pixels are partially characterized by showing low noise, meaning they will be less susceptible to pixel-specific noise causing a false photon trigger.

If a cosmic ray is absorbed in the fiber optic cable or the rest of the optical path it may deposit its energy in that material, exciting electrons which will then release secondary photons from this particle being absorbed. In this case, individual photons may be generated in the optical path which can then be transmitted to the MKID detectors. This would be a ‘true’ photon detection from an unintended physical source.

The previous two sources of photon detections are both issues that may contaminate sensitive data in a photon-starved environment but are currently challenging if not impossible to mitigate. For the first, reducing electrical noise by preventing ground loops, using low-noise power sources, and working in an isolated environment will reduce false triggers from electronic noise but in practice this is nearly impossible to eradicate. We note however that large scale electrical noise typically affects all resonators simultaneously and can therefore be removed (using a similar time-coincidence veto as a cosmic ray) or causes single resonators to become ‘hot’ or ‘cold’ which may also be handled




Figure 7.7 MKID spectrum (blue) compared to a 300 K blackbody spectrum (black) scaled to the central value of the MKID spectrum shown on a logarithmic scale. In this bandpass the blackbody spectrum varies over 7 orders of magnitude while the MKID spectrum remains relatively flat. The central point of the 300 K spectrum is normalized to the central point of the MKID spectrum. Error bars on the MKID spectrum are sufficiently small that they are contained within the points themselves.

gracefully in the data reduction pipeline (Steiger et al., 2022a). For the second, a secondary photon from a cosmic ray may be removed in the data reduction if its energy is sufficiently far outside the calibrated bandpass of the detector, but if its energy is within the bandpass then it will be impossible to remove as it is a single photon event and therefore not subject to the same veto as when a cosmic ray strikes the detector directly.

We explored the possibility that the photons that the MKIDs are seeing while purportedly unilluminated are coming from blackbody radiation in the dilution refrigerator. Although all precautions have been taken to prevent stray photons from hitting the de-

tector photons are incredibly difficult to insulate against, so we examine the possibility that the photons measured in the dark environment come from a thermal source.

First, we generate blackbody spectra for each of the 4 potential temperature stages which may be generating blackbody photons that could possibly hit the MKIDs. These are the 100 mK stage where the array itself and fiber collimator are mounted and the array is directly exposed to, the walls of the 4 K and 50 K intermediate stages that are used to step down from room to operating temperature and that are nested around the 100 mK stage, and 300 K, which represents the blackbody radiation from the ambient environment or the inner face of the outermost temperature stage of the dilution refrigerator.

Using Planck's Law

$$B_{\lambda}(T) = \frac{2hc^2}{\lambda^5} \frac{1}{\exp\left(\frac{hc}{\lambda k_B T}\right) - 1} \quad (7.3)$$

we find that the flux density of 100 mK, 4 K, and 50 K blackbody radiation between 0.9-1.6 eV is sufficiently small that a blackbody at any of these temperatures would not produce photons there over the duration of the experiment (86,750 seconds) so the blackbody spectra at these temperatures are not included in the analysis.

The 300 K blackbody spectrum can be seen plotted against the spectrum measured by the MKID pixels in Figure 7.7. The spectra are shown in units of $\log_{10}(\text{Flux Density})$, where the flux density was measured/calculated in $\text{ergs/s/cm}^2/\text{\AA}$. The plot is shown in log scale because while the MKID spectrum remains relatively flat across the bandpass, ranging from $(5.22 \pm 0.04) \times 10^{-14} \text{ ergs/s/cm}^2/\text{\AA}$ to $(3.07 \pm 0.02) \times 10^{-14} \text{ ergs/s/cm}^2/\text{\AA}$, the 300 K blackbody spectrum varies over 7 orders of magnitude.

The massive discrepancy in shape of the two spectra show that the photon hits that are still being measured by the MKID pixels are not generated by a 300 K blackbody. With this and the fact that the 100 mK, 4 K, and 50 K stages will not generate any blackbody photons over the calibrated bandpass it is possible to say that the source of

the remaining photons measured using the MKID digital readout do not come solely from blackbody sources in the environment.

7.6 Analog Readout: Single-detector Measurements

7.6.1 Experimental Setup

The analog readout utilizes the same internal electronics of the fridge, but externally the 6 foot SMA cables attach to a homodyne readout system consisting of two Anritsu MG37022A Signal generators, a Weinschel Attenuator box 8310 Series, an National Instruments-ADC/DAC, and an IQ mixer box. The function of these devices is the same as in the digital readout case. The schematic for the analog readout system, and the parametric amplifier, is shown in Figure 7.3. Unlike the digital readout, the analog readout supplies individual frequencies from an Anritsu Synthesizer to probe single resonators on the array. The analog readout has less noise associated with it compared to the digital readout which has to make compromises so that it is able to issue many probe tones while also dealing with limited dynamic range in the ADC attenuators and precision in its firmware computations.

The primary reason for taking a set of data in the dark with the analog readout is to analyze the nature of the photon pulses. Despite the fact that MKIDs are not susceptible to read noise and dark current (Day et al., 2003; Mazin et al., 2012; Fruitwala et al., 2020), empirical evidence has shown that they do indeed still measure photon-like events when they are not illuminated that may be triggered by noise sources such as those noted in Section 7.4.2. Since the MKID Analog Readout saves photon pulse timestream data (as in Figures 7.1 and 7.2) it is possible to explore the nature of these events to determine if these ‘dark counts’ look the same as ‘true’ photon events, or if they are demonstrably

different. By measuring several thousand ‘dark’ photons this way we aim to assess them and determine if we can assign them any explainable origin.

7.6.2 Data Collection and Reduction

A second set of data was collected in addition that from the MKID Digital Readout while the MKID was unilluminated where an MKID pixel was read out using an analog readout system designed for low noise, single-pixel characterization.

We chose a single MKID resonator that was also read out using the digital readout that had above-average resolving powers at all calibration energies. In this configuration there is only one MKID pixel read out, so we are no longer able to leverage the bulk properties of the MKID array for cosmic ray rejection. However, the analog readout saves the phase timestream of the resonator surrounding each photon event which allows inspection of each pulse (see Figures 7.1 and 7.2) to determine whether it is characteristic of a ‘real’ photon or whether noise caused the trigger.

In contrast to the digital readout which continuously takes data until the user decides to stop, the analog readout accepts the number of desired photon counts to measure before stopping. In this case the quiescent count rate of photon counts in the dark was first measured and found to be ~ 0.03 Hz, which likely consists of predominantly cosmic ray events. With this in mind we chose to register 8000 photon counts and expected this should take roughly 3 days. The primary goal of this investigation is to see if the photons which are being triggered on look ‘real’ or if they look like noise, although it is also possible to ascertain a dark count rate.

The analog readout system saves its data in a slightly different structure than the digital readout of Section 7.5. The digital readout saves time-tagged lists of photons along with the pixel location and height of the photon event but due to computational

constraints do not save further information about the photons. In contrast, the analog readout operates in a way so that the recent phase data from each MKID pixel being read out is kept in memory so when a photon is measured the readout system can save a timestream of that phase data from the pixel from the time surrounding the photon event. An example of a timestream saved by the analog readout is shown in Figure 7.1. The duration and sample rate of this phase timestream can be tuned by the user. In this experiment the resonator’s phase was sampled at 0.8 MHz and each phase timestream saved the 5000 μs surrounding the photon event (2500 μs before and after).

Analogously to the MKID Science Data pipeline, we first reject any photons which are outside of the calibrated bandpass (i.e. the peak of the phase is too high or too low) and any timestreams that contain more than 1 photon hit. The second criterion is the closest proxy we have to a time-coincidence veto in lieu of using bulk statistics from many pixels. This leaves 1118 photons, $\sim 14\%$ of the total observed counts. In comparison, the cuts from the digital readout left us with $\sim 10.7\%$ of the total observed counts (94262 remained of the initial 880855 from the analyzed pixels). The percentage for the analog readout is higher due to the lower noise from the readout system resulting in fewer counts below the calibrated region (i.e. triggers on noise) and the inability to make simultaneity cuts on cosmic ray events in the analog case.

7.6.3 Analysis

By binning photons using the same energy bins as in the MKID Digital readout we are able to compare the improvement that is gained when reading out a single pixel using significantly less noisy readout electronics. The count rates within the calibrated bandpass are shown in Table 7.3. As in the digital readout analysis the errors on photon counts and ultimate count rates are calculated using Poisson statistics.

Table 7.3. Dark Count Rates Between Analog and Digital MKID Readouts

| Bin Center (eV) | Analog Readout | | Digital Readout (No Reduction) | |
|--------------------|--|---------------------------|--|---------------------------|
| | Rate ($\times 10^{-4}$) (photons/pixel/s) | Total Counts (photons) | Rate ($\times 10^{-3}$) (photons/pixel/s) | Total Counts (photons) |
| 1.005 | 1.8 \pm 0.1 | 189 \pm 14 | 5.0 \pm 0.1 | 436.7 \pm 0.9 |
| 1.123 | 1.4 \pm 0.1 | 148 \pm 12 | 3.83 \pm 0.09 | 332.2 \pm 0.8 |
| 1.240 | 1.4 \pm 0.1 | 145 \pm 12 | 3.17 \pm 0.08 | 275.2 \pm 0.7 |
| 1.358 | 1.4 \pm 0.1 | 141 \pm 12 | 2.73 \pm 0.07 | 237.1 \pm 0.6 |
| 1.476 | 1.2 \pm 0.1 | 121 \pm 11 | 2.44 \pm 0.07 | 211.6 \pm 0.6 |

Note. — Comparison of the total counts in the calibrated bandpass between 0.946-1.534 eV (1310-808 nm) using the analog readout and digital MKID readout systems.

Photon Rise and Fall Times

As previously discussed, the analog readout system saves photon timestream data which allows us to examine the characteristic rise and fall times of the MKID pixel's phase when a photon event is triggered. For a baseline measurement an 808 nm (0.946 eV) laser and a 1310 nm (1.534 eV) laser are each shined on the pixel until it has registered 20,000 photon events.

A photon event is characterized by a fast exponential rise time in the measured phase as a photon strikes the detector, depositing its energy and breaking Cooper pairs into quasiparticles followed by a slower exponential tail as the quasiparticles recombine. To fit the rise and the fall times for a given photon event, the timestream is first split into a rising portion from the start of the timestream to the peak and a falling portion from the peak to the end. The rise and fall times are then each calculated from their respective sections of the timestream by fitting an exponential of the form

$$\phi = Ae^{-(t/b)} + c \quad (7.4)$$

where ϕ is the phase, t is the corresponding time, b is the time constant (which we call the rise or the fall time depending on which part of the event we are fitting), and A and c are constants to account for any offset or scaling differences between pulses.

We fit the rise and fall times for all of the photons from both the 808 and 1310 nm lasers as well as the dark-count photons measured by the pixel while it was unilluminated. Figure 7.8 shows how the rise and fall times of the dark counts compare to those of known photons from the two lasers as a function of wavelength. It can be clearly seen that the 1σ (68%) confidence intervals of the rise times of the dark count photons overlaps with the 1σ CI of the rise times of the 808 and 1310 nm photons. This is relatively unsurprising as phase spikes so sharply when energy is deposited into the pixel that it is effectively instantaneous even with the microsecond timing resolution.

On the other hand, when the fall times of the dark counts are compared to those of the photons from the lasers one can see a stark distinction between the two distributions. The fall times from the laser photons are significantly faster than those of the dark counts for photons of similar energy (and phase response). In theory the phase decay back to its quiescent value is governed by the quasiparticle recombination time, an effect solely due to properties of the superconducting material. Since the phase response is proportional to the number of quasiparticles generated in the material, photons that cause similar phase responses should have similar fall times since a roughly equivalent number of quasiparticles need to combine back into Cooper pairs. We can see in the bottom panel of Figure 7.8 that this is not the case in this experiment. Figure 7.9 shows an example of each of these photons. In the top panel, a photon that was measured when a laser was incident on the pixel can be seen with a measured fall time of $32.07 \mu\text{s}$. Below that a second photon of the same energy can be seen, this time from when there was no light incident on the array. The fall time in this case is measured to be $79.89 \mu\text{s}$, consistent with the difference in the distribution of the fall times from each population.

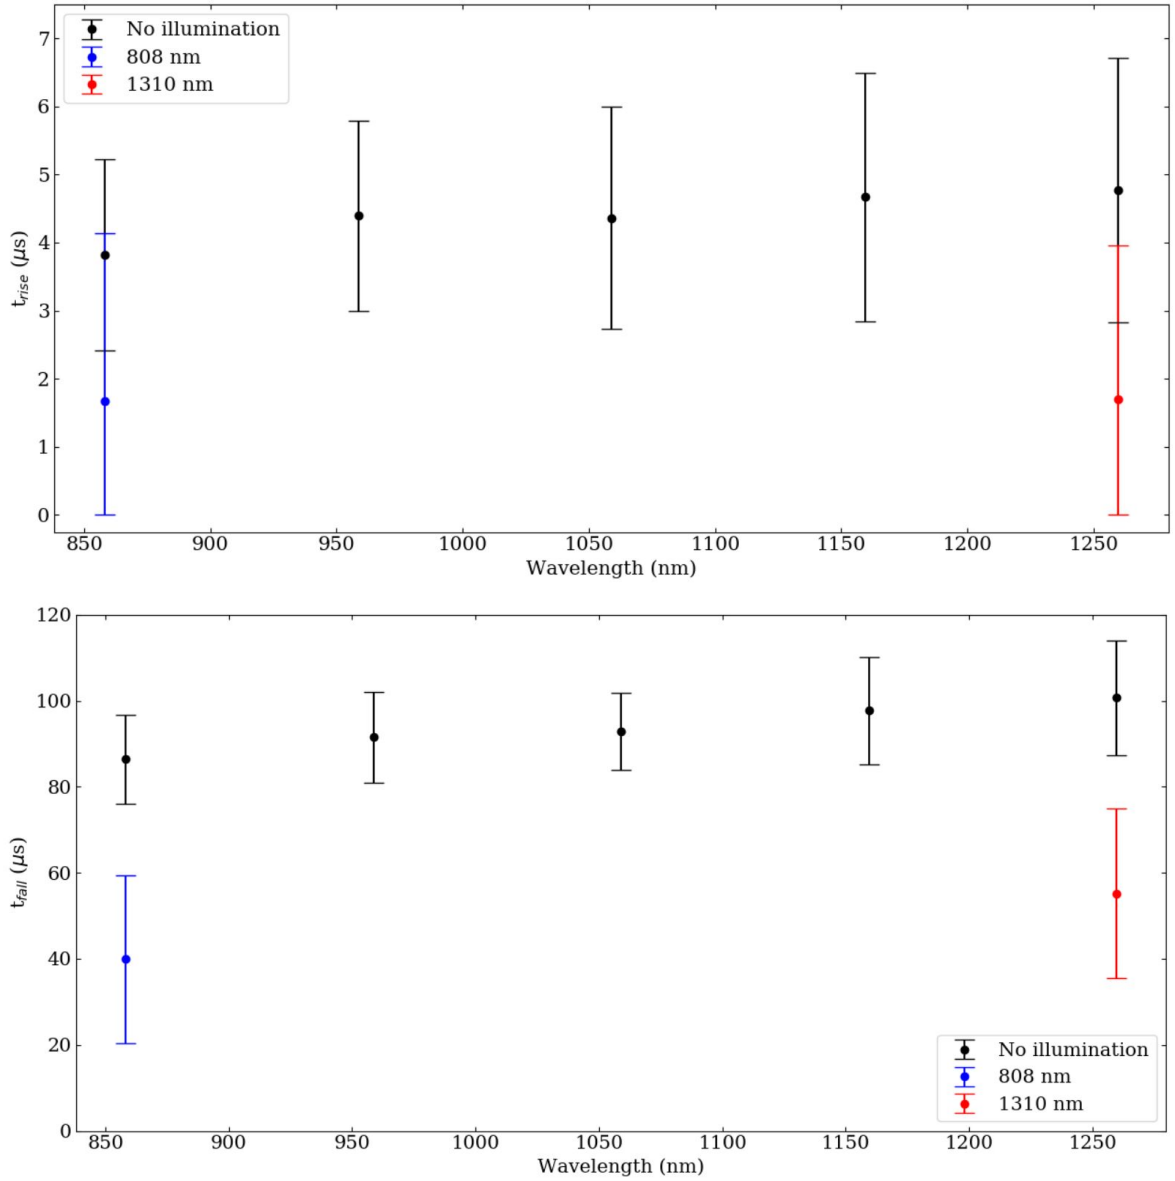


Figure 7.8 Comparison of the rise times (top) and fall times (bottom) as a function of photon wavelength between dark counts - shown in black - and photons from 808 and 1310 nm lasers - shown in blue and red, respectively. The error bars for each point are the 1σ -errors for each measurement. In the top panel it can be clearly seen that the rise times of dark count photons and laser photons overlap significantly. In the bottom panel, the fall times for the dark count events do not overlap the photon events from the lasers.




Figure 7.9 Comparing the calculated rise and fall times of photons of the same energy measured by the MKID resonator. (Top) An 808 nm photon measured with an 808 nm laser shining on the pixel. (Bottom) An 808 nm photon measured when there is no light incident on the pixel. This corroborates the significant difference in fall times shown in Figure 7.8.

Although the dark count photons that remain after all previous cuts *qualitatively* appear similar to the photons measured when lasers were being shined on the MKIDs, there is a quantitative difference in the two populations. The photons that are measured when a laser is being shined on the MKID show significantly shorter fall times than dark count photons measured by the same pixel. The explanation for this behavior stems from the sources of the different families of photons and how they deposit energy into and are subsequently measured by the pixel. The rise and fall times of a photon event correlates with how quickly energy is deposited into the resonator and how long it takes for that energy to dissipate and allow the resonator return to its unexcited state. When energy is deposited into the resonator over a short time it will rapidly break many Cooper pairs which then begin to recombine right away, allowing the resonator to return to its unexcited state in a short time. If energy is deposited into the resonator over a longer time scale it will still break many Cooper pairs quickly but as the energy remains in the system before dissipating it will prevent those Cooper pairs from combining as fast, leading to a longer fall time back to quiescence. Real photons from a light source fall into

the first category; the photon is absorbed and all of its energy is immediately deposited into the MKID, leading to a fast rise and fall time. Dark counts from sources such as cosmic rays are in the second family. Photon events triggered by a cosmic ray occur when the energy from the incident energetic particle is down-converted into a cloud of phonons that spread the energy through the array substrate. As these phonons move past MKID pixels they deposit some energy into the pixel over the time it takes for them to move through the resonator. This means that dark count photons from cosmic ray events are non-instantaneous process and will consequently have longer fall times.

A separate potential explanation for the discrepancy in fall times between the laser photons and dark count photons is that while illuminating the array with a laser the photon is significantly higher than when the laser is off. This higher flux may lead to slight heating of the film of the MKID array. In turn, this would lead to shorter quasiparticle recombination times and ultimately shorter fall times for each photon event.

Currently there is insufficient evidence to conclusively say that these are two completely different populations of photons due to the current gap in understanding of the noise sources in the MKIDs and inability to simultaneously read out an array using both the digital and analog readouts to correlate cosmic ray events (digital) to single photon traces (analog) so we cannot calibrate these photons out based solely on the difference in their fall times than is expected. If the assumption is made that these counts do come from cosmic rays and a calibration cut is made, the number of counts would decrease significantly from 121 to 10 photons in the high energy bin and from 189 to 27 photons in the lowest energy bin, with each bin seeing a reduction by about a factor of 8. Across the full bandpass the count rate would fall from $(7.1 \pm 0.3) \times 10^{-3}$ photons/pixel/s to $(9.3 \pm 0.9) \times 10^{-4}$ photons/pixel/s. In wavelength space this corresponds to $(1.42 \pm 0.05) \times 10^{-5}$ photons/pixel/s/nm and $(1.9 \pm 0.2) \times 10^{-6}$ photons/pixel/s/nm for the ‘uncalibrated’ and ‘calibrated’ rates, respectively. The final value with long-fall

time events removed is just slightly below the $(2.77 \pm 0.02) \times 10^{-6}$ photons/pixel/s/nm from the digital readout in Section 7.5.3.

7.7 Discussion and Conclusions

In this experiment we showed that across the calibrated MKID bandpass from 0.946 to 1.534 eV (1310-808 nm), the count rate seen by the detectors in a large format array is $(3.14 \pm 0.01) \times 10^{-3}$ photons/pixel/s/eV or $(3.68 \pm 0.01) \times 10^{-6}$ photons/pixel/s/nm. It is also demonstrated that by using a relatively light calibration cut for cosmic ray events we are able to reduce the number of spurious photon events by nearly a factor 10 while removing less than 1% of the duration of the data collection.

Using the MKID analog readout system and recording the shape of photon pulses in a single pixel we first show that the count rate across the calibrated bandpass is $(1.42 \pm 0.05) \times 10^{-5}$ photons/pixel/s/nm without any further data cleaning steps, demonstrating that a quieter system does lead to lower count rates in an unilluminated MKID device.

While investigating the shape of the photon pulses using the analog readout it was found that the exponential tail of the dark counts corresponds to a significantly longer fall time than from photons generated from a laser at the same wavelengths. The long fall times are indicative of energy taking a long time to dissipate from the resonators which points to events causing these triggers happening in the substrate rather than the pixels themselves. An example of a known event that takes place in the substrate and causes contaminating photon events is a cosmic ray hitting the MKID array. Making MKID pixels atop membranes is an ongoing field of research that offers a straightforward way to minimize the potential for substrate absorptions to cause contaminating photon events. Because of the added complexity of making MKIDs on membranes in addition to the existing difficulty in fabrication this is not feasible for the large-format MKID arrays

currently in use. However, in future experiments requiring much fewer (1 to ~ 100) MKIDs, making them on membranes may be a reasonable path forward that will help mitigate contamination from substrate absorptions. Additionally, the current generation of MKID readout hardware does not allow for side-by-side simultaneous readout of many pixels while still capturing the photon phase timestreams and so there is no way to determine if these long fall time photons are coexistent with cosmic ray events at present. The 3rd Generation MKID Digital Readout is currently under development (Smith et al., 2022) and promises to allow both capabilities at the same time. Further investigation of the source of these dark photons and if they are in fact generated cosmic rays will be explored in future work with the 3rd Generation MKID Readout. We note here that if future work does find that these long fall time photons are from contaminating sources it offers a straightforward way to calibrate them out of MKID datasets.

For a future dark matter detector experiment that would use a 100 pixel MKID array with 10 nm energy bins, the maximum dark count rate in the detector would be $\sim (3.68 \pm 0.01) \times 10^{-3}$ photons/s if the current style arrays and generation of MKID digital readout were used. With this said, any future MKID dark matter direct detection instrument will have several key upgrades to mitigate noise in the system. First, a new generation of MKID readout is currently under development which promises to be a significantly less noisy system than the one used at present. The continued development of Traveling Wave Parametric Amplifiers (TWPA; Eom et al., 2012; Zobrist et al., 2019) will also significantly reduce system noise compared to the more commonly used HEMT amplifiers. Finally, this instrument itself will use an array that has anti-reflection (AR) coating on the MKID devices and will not have optics that allow visible light to enter the cryostat. Both of these upgrades will prevent more stray photons from entering the fridge and causing spurious, unattributable counts on the detector.

Bibliography

- Ahn, K., Guyon, O., Lozi, J., et al. 2021, in *Techniques and Instrumentation for Detection of Exoplanets X*, ed. S. B. Shaklan & G. J. Ruane, Vol. 11823, International Society for Optics and Photonics (SPIE), 9 – 21, doi: 10.1117/12.2594963
- Allard, F., Homeier, D., & Freytag, B. 2012, *Philosophical Transactions of the Royal Society of London Series A*, 370, 2765, doi: 10.1098/rsta.2011.0269
- Asplund, M., Grevesse, N., Sauval, A. J., & Scott, P. 2009, *ARA&A*, 47, 481, doi: 10.1146/annurev.astro.46.060407.145222
- Baraffe, I., Homeier, D., Allard, F., & Chabrier, G. 2015, *A&A*, 577, A42, doi: 10.1051/0004-6361/201425481
- Bernstein, R., Shtetman, S. A., Gunnels, S. M., Mochacki, S., & Athey, A. E. 2003, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 4841, *Instrument Design and Performance for Optical/Infrared Ground-based Telescopes*, ed. M. Iye & A. F. M. Moorwood, 1694–1704, doi: 10.1117/12.461502
- Beuzit, J.-L., Vigan, A., Mouillet, D., et al. 2019, *Astronomy & Astrophysics*, 631, A155, doi: 10.1051/0004-6361/201935251
- Blanco, C., Bruca, L., Catalano, S., & Marilli, E. 1982, *A&A*, 115, 280
- Blunt, S., Wang, J. J., Angelo, I., et al. 2020, *AJ*, 159, 89, doi: 10.3847/1538-3881/ab6663
- Bockstiegel, C. 2019, in *UCSB Electronic Theses and Dissertations*. <https://escholarship.org/uc/item/87w8q68w>
- Bonavita, M., Fontanive, C., Desidera, S., et al. 2020, *MNRAS*, 494, 3481, doi: 10.1093/mnras/staa590
- Bonavita, M., Fontanive, C., Gratton, R., et al. 2022, *MNRAS*, 513, 5588, doi: 10.1093/mnras/stac1250
- Boro Saikia, S., Marvin, C. J., Jeffers, S. V., et al. 2018, *Astronomy & Astrophysics*, 616, A108, doi: 10.1051/0004-6361/201629518

- Borucki, W. J. 2016, *Reports on Progress in Physics*, 79, 036901, doi: 10.1088/0034-4885/79/3/036901
- Bottom, M., Neat, L. S., Morrissey, P., et al. 2018, in *High Energy, Optical, and Infrared Detectors for Astronomy VIII*, ed. A. D. Holland & J. Beletic (SPIE), doi: 10.1117/12.2312335
- Bottom, M., Shelton, J. C., Wallace, J. K., et al. 2016, *Publications of the Astronomical Society of the Pacific*, 128, 075003, doi: 10.1088/1538-3873/128/965/075003
- Bowler, B. P. 2016, *Publications of the Astronomical Society of the Pacific*, 128, 102001, doi: 10.1088/1538-3873/128/968/102001
- Bowler, B. P., & Nielsen, E. L. 2018, in *Handbook of Exoplanets* (Springer International Publishing), 1967–1983, doi: 10.1007/978-3-319-55333-7_155
- Bowler, B. P., Cochran, W. D., Endl, M., et al. 2021, *AJ*, 161, 106, doi: 10.3847/1538-3881/abd243
- Brandt, G. M., Brandt, T. D., Dupuy, T. J., Li, Y., & Michalik, D. 2021, *The Astronomical Journal*, 161, 179, doi: 10.3847/1538-3881/abdc2e
- Brandt, G. M., Dupuy, T. J., Li, Y., et al. 2021, *AJ*, 162, 301, doi: 10.3847/1538-3881/ac273e
- Brandt, T. D. 2021, *The Astrophysical Journal Supplement Series*, 254, 42, doi: 10.3847/1538-4365/abf93c
- Brandt, T. D., Dupuy, T. J., & Bowler, B. P. 2019, *AJ*, 158, 140, doi: 10.3847/1538-3881/ab04a8
- Brandt, T. D., Dupuy, T. J., Li, Y., et al. 2021, *AJ*, 162, 186, doi: 10.3847/1538-3881/ac042e
- Brandt, T. D., McElwain, M. W., Turner, E. L., et al. 2014, *The Astrophysical Journal*, 794, 159, doi: 10.1088/0004-637x/794/2/159
- Brandt, T. D., Rizzo, M., Groff, T., et al. 2017, *Journal of Astronomical Telescopes, Instruments, and Systems*, 3, 048002, doi: 10.1117/1.JATIS.3.4.048002
- Brown, A. G. A., Vallenari, A., Prusti, T., et al. 2018, *Astronomy & Astrophysics*, 616, A1, doi: 10.1051/0004-6361/201833051
- . 2021, *Astronomy & Astrophysics*, 650, C3, doi: 10.1051/0004-6361/202039657e
- Brown, T. M., Baliber, N., Bianco, F. B., et al. 2013, *PASP*, 125

- Cagigal, M. P., & Canales, V. F. 2001, *Optical Engineering*, 40, 2690, doi: 10.1117/1.1417495
- Carson, J., Thalmann, C., Janson, M., et al. 2013, *ApJL*, 763, L32, doi: 10.1088/2041-8205/763/2/L32
- Casagrande, L., Schönrich, R., Asplund, M., et al. 2011, *Astronomy & Astrophysics*, 530, A138, doi: 10.1051/0004-6361/201016276
- Castelló-Mor, N. 2020, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 958, 162933, doi: 10.1016/j.nima.2019.162933
- Chabrier, G. 2003, *ApJL*, 586, L133, doi: 10.1086/374879
- Chauvin, G., Desidera, S., Lagrange, A. M., et al. 2017, *A&A*, 605, L9, doi: 10.1051/0004-6361/201731152
- Cheetham, A., Bonnefoy, M., Desidera, S., et al. 2018, *A&A*, 615, A160, doi: 10.1051/0004-6361/201832650
- Chilcote, J., Tobin, T., Currie, T., et al. 2021, *The Astronomical Journal*, 162, 251, doi: 10.3847/1538-3881/ac29ba
- Close, L. M., Males, J. R., Durney, O., et al. 2018, *Optical and mechanical design of the extreme AO coronagraphic instrument MagAO-X*. <https://arxiv.org/abs/1807.04311>
- Cook, T., Cahoy, K., Chakrabarti, S., et al. 2015, *Journal of Astronomical Telescopes, Instruments, and Systems*, 1, 044001, doi: 10.1117/1.JATIS.1.4.044001
- Crane, J. D., Shtetman, S. A., Butler, R. P., et al. 2010, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7735, *Ground-based and Airborne Instrumentation for Astronomy III*, ed. I. S. McLean, S. K. Ramsay, & H. Takami, 773553, doi: 10.1117/12.857792
- Currie, T., Biller, B., Lagrange, A.-M., et al. 2022a, *arXiv e-prints*, arXiv:2205.05696. <https://arxiv.org/abs/2205.05696>
- Currie, T., Brandt, G. M., & Brandt, T. D. 2022b, submitted to *Science*, under final revision
- Currie, T., Cloutier, R., Brittain, S., et al. 2015, *ApJL*, 814, L27, doi: 10.1088/2041-8205/814/2/L27
- Currie, T., Daemgen, S., Debes, J., et al. 2014, *ApJL*, 780, L30, doi: 10.1088/2041-8205/780/2/L30

- Currie, T., Burrows, A., Itoh, Y., et al. 2011, *ApJ*, 729, 128, doi: 10.1088/0004-637X/729/2/128
- Currie, T., Debes, J., Rodigas, T. J., et al. 2012, *ApJL*, 760, L32, doi: 10.1088/2041-8205/760/2/L32
- Currie, T., Brandt, T. D., Uyama, T., et al. 2018a, *AJ*, 156, 291, doi: 10.3847/1538-3881/aae9ea
- Currie, T., Kasdin, N., Groff, T., et al. 2018b, *Publications of the Astronomical Society of the Pacific*, 130, doi: 10.1088/1538-3873/aaab41
- Currie, T., Guyon, O., Lozi, J., et al. 2019, Performance and Early Science with the Subaru Coronagraphic Extreme Adaptive Optics Project. <https://arxiv.org/abs/1909.10522>
- Currie, T., Marois, C., Cieza, L., et al. 2019, *ApJL*, 877, L3, doi: 10.3847/2041-8213/ab1b42
- Currie, T., Brandt, T. D., Kuzuhara, M., et al. 2020, *The Astrophysical Journal Letters*, 904, L25, doi: 10.3847/2041-8213/abc631
- Currie, T., Brandt, T. D., Kuzuhara, M., et al. 2021, in *Techniques and Instrumentation for Detection of Exoplanets X*, ed. S. B. Shaklan & G. J. Ruane, Vol. 11823, International Society for Optics and Photonics (SPIE), 22 – 40, doi: 10.1117/12.2595001
- Currie, T., Lawson, K., Schneider, G., et al. 2022c, *Nature Astronomy*, doi: 10.1038/s41550-022-01634-x
- Daigle, O., Djazovski, O., Laurin, D., Doyon, R., & Artigau, É. 2012, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 8453, High Energy, Optical, and Infrared Detectors for Astronomy V, ed. A. D. Holland & J. W. Beletic, 845303, doi: 10.1117/12.926385
- Danziger, I. J., & Faber, S. M. 1972, *A&A*, 18, 428
- Day, P. K., LeDuc, H. G., Mazin, B. A., Vayonakis, A., & Zmuidzinas, J. 2003, *Nature*, 425, 817, doi: 10.1038/nature02037
- de Medeiros, J. R., & Mayor, M. 1999, *A&AS*, 139, 433, doi: 10.1051/aas:1999401
- Dekany, R., Roberts, J., Burruss, R., et al. 2013, *The Astrophysical Journal*, 776, 130, doi: 10.1088/0004-637x/776/2/130
- Desidera, S., Chauvin, G., Bonavita, M., et al. 2021, *A&A*, 651, A70, doi: 10.1051/0004-6361/202038806

- Dodkins, R., Mahashabde, S., O'Brien, K., et al. 2018, *Astronomy and Computing*, 23, 60, doi: <https://doi.org/10.1016/j.ascom.2018.03.001>
- Doelman, D. S., Snik, F., Por, E. H., et al. 2021, *Applied Optics*, 60, D52, doi: [10.1364/ao.422155](https://doi.org/10.1364/ao.422155)
- Dressler, A., Hare, T., Bigelow, B. C., & Osip, D. J. 2006, in *Ground-based and Airborne Instrumentation for Astronomy*, ed. I. S. McLean & M. Iye, Vol. 6269, International Society for Optics and Photonics (SPIE), 62690F, doi: [10.1117/12.670573](https://doi.org/10.1117/12.670573)
- Duchêne, G., Rice, M., Hom, J., et al. 2020, *AJ*, 159, 251, doi: [10.3847/1538-3881/ab8881](https://doi.org/10.3847/1538-3881/ab8881)
- Eom, B. H., Day, P. K., Leduc, H. G., & Zmuidzinas, J. 2012, *A Wideband, Low-Noise Superconducting Amplifier with High Dynamic Range*, arXiv, doi: [10.48550/ARXIV.1201.2392](https://doi.org/10.48550/ARXIV.1201.2392)
- Esposito, S., & Riccardi, A. 2001, *A&A*, 369, L9, doi: [10.1051/0004-6361:20010219](https://doi.org/10.1051/0004-6361:20010219)
- Favata, F., Micela, G., & Sciortino, S. 1995, *A&A*, 298, 482
- Fleming, T. A., Gioia, I. M., & Maccacaro, T. 1989, *ApJ*, 340, 1011, doi: [10.1086/167454](https://doi.org/10.1086/167454)
- Foreman-Mackey, D., Morton, T. D., Hogg, D. W., Agol, E., & Schölkopf, B. 2016, *AJ*, 152, 206, doi: [10.3847/0004-6256/152/6/206](https://doi.org/10.3847/0004-6256/152/6/206)
- Fried, D. L. 1965, *Journal of the Optical Society of America (1917-1983)*, 55, 1427
- Fruitwala, N. 2021, in *UCSB Electronic Theses and Dissertations*. <https://escholarship.org/uc/item/71t3j3dk>
- Fruitwala, N., Walter, A. B., au2, J. I. B. I., Dodkins, R., & Mazin, B. A. 2021, *End-to-end Deep Learning Pipeline for Microwave Kinetic Inductance Detector (MKID) Resonator Identification and Tuning*. <https://arxiv.org/abs/2104.01282>
- Fruitwala, N., Meeker, S., Mazin, B., et al. 2018, in *SPIE 2018: Adaptive Optics Systems VI*, 57, doi: [10.1117/12.2311402](https://doi.org/10.1117/12.2311402)
- Fruitwala, N., Strader, P., Canelo, G., et al. 2020, *Review of Scientific Instruments*, 91, 124705, doi: [10.1063/5.0029457](https://doi.org/10.1063/5.0029457)
- Gagné, J., Faherty, J. K., Cruz, K. L., et al. 2015, *ApJS*, 219, 33, doi: [10.1088/0067-0049/219/2/33](https://doi.org/10.1088/0067-0049/219/2/33)
- Gagné, J., Mamajek, E. E., Malo, L., et al. 2018, *ApJ*, 856, 23, doi: [10.3847/1538-4357/aaae09](https://doi.org/10.3847/1538-4357/aaae09)

- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2021, *A&A*, 649, A1, doi: 10.1051/0004-6361/202039657
- Gioia, I. M., Maccacaro, T., Schild, R. E., et al. 1990, *ApJS*, 72, 567, doi: 10.1086/191426
- Głębocki, R., & Gnaciński, P. 2005, in *ESA Special Publication*, Vol. 560, 13th Cambridge Workshop on Cool Stars, Stellar Systems and the Sun, ed. F. Favata, G. A. J. Hussain, & B. Battrick, 571
- Glebocki, R., & Gnacinski, P. 2005, *VizieR Online Data Catalog*, III/244
- Goodman, J. W. 1975, *Statistical Properties of Laser Speckle Patterns* (Berlin, Heidelberg: Springer Berlin Heidelberg), 9–75, doi: 10.1007/978-3-662-43205-1_2
- Gossage, S., Conroy, C., Dotter, A., et al. 2018, *The Astrophysical Journal*, 863, 67, doi: 10.3847/1538-4357/aad0a0
- Greco, J. P., & Brandt, T. D. 2016, *ApJ*, 833, 134, doi: 10.3847/1538-4357/833/2/134
- Groff, T. D., Chilcote, J., Kasdin, N. J., et al. 2016, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9908, *Ground-based and Airborne Instrumentation for Astronomy VI*, 99080O, doi: 10.1117/12.2233447
- Guyon, O. 2005, *The Astrophysical Journal*, 629, 592, doi: 10.1086/431209
- Guyon, O., Roddier, C., Graves, J. E., et al. 1999, *PASP*, 111, 1321, doi: 10.1086/316445
- Hall, R., & Jefferies, S. M. 2022, *Journal of the Astronautical Sciences*, 69, 581, doi: 10.1007/s40295-022-00306-2
- Harun, M., Rashid, A., Megat Ahmad, M. H., et al. 2020, *Detection of ionizing radiations using CMOS sensor from consumer camera device*, doi: 10.36227/techrxiv.13705414.v1
- Hickish, J., Abdurashidova, Z., Ali, Z., et al. 2016, *A Decade of Developing Radio-Astronomy Instrumentation using CASPER Open-Source Technology*, arXiv, doi: 10.48550/ARXIV.1611.01826
- Holmberg, J., Nordström, B., & Andersen, J. 2009, *A&A*, 501, 941, doi: 10.1051/0004-6361/200811191
- Husser, T.-O., Wende-von Berg, S., Dreizler, S., et al. 2013, *Astronomy & Astrophysics*, 553, A6, doi: 10.1051/0004-6361/201219058
- Jovanovic, N., Guyon, O., Martinache, F., et al. 2015, *ApJL*, 813, L24, doi: 10.1088/2041-8205/813/2/L24

- Jovanovic, N., Martinache, F., Guyon, O., et al. 2015, *Publications of the Astronomical Society of the Pacific*, 127, 890–910, doi: 10.1086/682989
- Keppler, M., Benisty, M., Müller, A., et al. 2018, *Astronomy & Astrophysics*, 617, A44, doi: 10.1051/0004-6361/201832957
- Kervella, P., Arenou, F., Mignard, F., & Thévenin, F. 2019, *Astronomy & Astrophysics*, 623, A72, doi: 10.1051/0004-6361/201834371
- Konopacky, Q. M., Rameau, J., Duchêne, G., et al. 2016, *ApJL*, 829, L4, doi: 10.3847/2041-8205/829/1/L4
- Kuzuhara, M., Tamura, M., Kudo, T., et al. 2013, *ApJ*, 774, 11, doi: 10.1088/0004-637X/774/1/11
- Kuzuhara, M., Currie, T., Takarada, T., et al. 2022, arXiv e-prints, arXiv:2205.02729. <https://arxiv.org/abs/2205.02729>
- Lagrange, A. M., Gratadour, D., Chauvin, G., et al. 2009, *A&A*, 493, L21, doi: 10.1051/0004-6361:200811325
- Lagrange, A.-M., Meunier, N., Rubini, P., et al. 2019, *Nature*
- Li, Y., Brandt, T. D., Brandt, G. M., et al. 2021, *AJ*, 162, 266, doi: 10.3847/1538-3881/ac27ab
- Macintosh, B. A., Graham, J. R., Barman, T. S., et al. 2015, *Science*, 350, 64
- Maire, A. L., Baudino, J. L., Desidera, S., et al. 2020, *A&A*, 633, L2, doi: 10.1051/0004-6361/201937134
- Males, J. R., Close, L. M., Miller, K., et al. 2018, MagAO-X: project status and first laboratory results. <https://arxiv.org/abs/1807.04315>
- Males, J. R., Close, L. M., Guyon, O., et al. 2020, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 11448, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 114484L, doi: 10.1117/12.2561682
- Males, J. R., Close, L. M., Haffert, S., et al. 2022, MagAO-X: current status and plans for Phase II. <https://arxiv.org/abs/2208.07299>
- Mamajek, E. E., & Hillenbrand, L. A. 2008, *The Astrophysical Journal*, 687, 1264–1293, doi: 10.1086/591785
- Marois, C., Doyon, R., Racine, R., & Nadeau, D. 2000, *PASP*, 112, 91, doi: 10.1086/316492

- Marois, C., Lafreniere, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *The Astrophysical Journal*, 641, 556–564, doi: 10.1086/500401
- Marois, C., Lafrenière, D., Macintosh, B., & Doyon, R. 2008, *ApJ*, 673, 647, doi: 10.1086/523839
- Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., & Barman, T. 2010, *Nature*, 468, 1080–1083, doi: 10.1038/nature09684
- Marshall, J. L., Burles, S., Thompson, I. B., et al. 2008, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7014, *Ground-based and Airborne Instrumentation for Astronomy II*, ed. I. S. McLean & M. M. Casali, 701454, doi: 10.1117/12.789972
- Mateo, M., Bailey, J. I., Crane, J., et al. 2012, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 8446, *Ground-based and Airborne Instrumentation for Astronomy IV*, ed. I. S. McLean, S. K. Ramsay, & H. Takami, 84464Y, doi: 10.1117/12.926448
- Mawet, D., Serabyn, E., Liewer, K., et al. 2009, *The Astrophysical Journal*, 709, 53, doi: 10.1088/0004-637x/709/1/53
- Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, *ApJ*, 792, 97, doi: 10.1088/0004-637x/792/2/97
- Mazin, B. A., Bumble, B., Meeker, S. R., et al. 2012, *Optics Express*, 20, 1503, doi: 10.1364/oe.20.001503
- McLeod, B., Geary, J., Conroy, M., et al. 2015, *Publications of the Astronomical Society of the Pacific*, 127, 366, doi: 10.1086/680687
- Meeker, S. R. 2017, in *UCSB Electronic Theses and Dissertations*. <https://escholarship.org/uc/item/1hj1q7s1>
- Meeker, S. R., Mazin, B. A., Walter, A. B., et al. 2018, *Publications of the Astronomical Society of the Pacific*, 130, 065001, doi: 10.1088/1538-3873/aab5e7
- Mendillo, C. B., Hewawasam, K., Martel, J., et al. 2022, in *Space Telescopes and Instrumentation 2022: Optical, Infrared, and Millimeter Wave*, ed. L. E. Coyle, S. Matsuura, & M. D. Perrin, Vol. 12180, *International Society for Optics and Photonics (SPIE)*, 1218022, doi: 10.1117/12.2630237
- Millar-Blanchaer, M. A., Perrin, M. D., Hung, L.-W., et al. 2016, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9908, *Ground-based and Airborne Instrumentation for Astronomy VI*, ed. C. J. Evans, L. Simard, & H. Takami, 990836, doi: 10.1117/12.2233071

- Miskovetz, K., Dupuy, T. J., Schonhut-Stasik, J., & Stassun, K. G. 2022, AAS Meeting Abstracts, doi: [10.48550/ARXIV.2205.06899](https://doi.org/10.48550/ARXIV.2205.06899)
- Montesinos, B., Gimenez, A., & Fernandez-Figueroa, M. 1988, Monthly Notices of the Royal Astronomical Society, 232, 361, doi: [10.1093/mnras/232.2.361](https://doi.org/10.1093/mnras/232.2.361)
- Morzinski, K. M., Close, L. M., Males, J. R., et al. 2014, in Astronomical Telescopes and Instrumentation
- Niedzwiecki, M., Rzecki, K., Marek, M., et al. 2019, Recognition and classification of the cosmic-ray events in images captured by CMOS/CCD cameras. <https://arxiv.org/abs/1909.01929>
- Nielsen, E. L., De Rosa, R. J., Macintosh, B., et al. 2019, AJ, 158, 13, doi: [10.3847/1538-3881/ab16e9](https://doi.org/10.3847/1538-3881/ab16e9)
- Nordström, B., Mayor, M., Andersen, J., et al. 2004, Astronomy & Astrophysics, 418, 989–1019, doi: [10.1051/0004-6361:20035959](https://doi.org/10.1051/0004-6361:20035959)
- Norris, B., Schworer, G., Tuthill, P., et al. 2015, MNRAS, 447, 2894, doi: [10.1093/mnras/stu2529](https://doi.org/10.1093/mnras/stu2529)
- Nowak, M., Lacour, S., Lagrange, A. M., et al. 2020, A&A, 642, L2, doi: [10.1051/0004-6361/202039039](https://doi.org/10.1051/0004-6361/202039039)
- Overbye, D. 2023, New York Times. <https://www.nytimes.com/2023/04/18/science/astronomy-telescopes-magellan-chile.html>
- Pecaut, M. J., & Mamajek, E. E. 2013, ApJS, 208, 9, doi: [10.1088/0067-0049/208/1/9](https://doi.org/10.1088/0067-0049/208/1/9)
- Persson, S. E., Murphy, D. C., Smee, S., et al. 2013, PASP, 125, 654, doi: [10.1086/671164](https://doi.org/10.1086/671164)
- Platt, B. C., & Shack, R. 2001, History and principles of Shack-Hartmann wavefront sensing, SLACK Incorporated Thorofare, NJ
- Pluzhnik, E. A., Guyon, O., Ridgway, S., et al. 2005, in American Astronomical Society Meeting Abstracts, Vol. 207, American Astronomical Society Meeting Abstracts, 110.17
- Preibisch, T., Kim, Y.-C., Favata, F., et al. 2005, The Astrophysical Journal Supplement Series, 160, 401, doi: [10.1086/432891](https://doi.org/10.1086/432891)
- Prusti, T., de Bruijne, J. H. J., Brown, A. G. A., et al. 2016, Astronomy & Astrophysics, 595, A1, doi: [10.1051/0004-6361/201629272](https://doi.org/10.1051/0004-6361/201629272)
- Pueyo, L. 2016, ApJ, 824, 117, doi: [10.3847/0004-637X/824/2/117](https://doi.org/10.3847/0004-637X/824/2/117)

- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *Journal of Astronomical Telescopes, Instruments, and Systems*, 1, 014003, doi: 10.1117/1.JATIS.1.1.014003
- Ruane, G., Ngo, H., Mawet, D., et al. 2019, *The Astronomical Journal*, 157, 118, doi: 10.3847/1538-3881/aafee2
- Salama, M., Ziegler, C., Baranec, C., et al. 2022, *The Astronomical Journal*, 163, 200, doi: 10.3847/1538-3881/ac53fc
- Seager, S., & Deming, D. 2010, *Annual Review of Astronomy and Astrophysics*, 48, 631, doi: 10.1146/annurev-astro-081309-130837
- Simcoe, R. A., Burgasser, A. J., Schechter, P. L., et al. 2013, *PASP*, 125, 270, doi: 10.1086/670241
- Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *AJ*, 131, 1163, doi: 10.1086/498708
- Smith, J. P., Bailey, J. I., & Mazin, B. A. 2022, in 2022 IEEE 30th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM), 1–2, doi: 10.1109/FCCM53951.2022.9786140
- Smith, J. P., Mazin, B. A., Walter, A. B., et al. 2021, *IEEE Transactions on Applied Superconductivity*, 31, 1, doi: 10.1109/tasc.2020.3008591
- Soummer, R., Pueyo, L., & Larkin, J. 2012, *ApJL*, 755, L28, doi: 10.1088/2041-8205/755/2/L28
- Stalder, B., Stark, A. A., Amato, S. M., et al. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9147, *Ground-based and Airborne Instrumentation for Astronomy V*, ed. S. K. Ramsay, I. S. McLean, & H. Takami, 91473Y, doi: 10.1117/12.2054933
- Stanford-Moore, S. A., Nielsen, E. L., De Rosa, R. J., Macintosh, B., & Czekala, I. 2020, *ApJ*, 898, 27, doi: 10.3847/1538-4357/ab9a35
- Steiger, S., Bailey III, J. I., Zobrist, N., et al. 2022a, arXiv e-prints, arXiv:2203.01406. <https://arxiv.org/abs/2203.01406>
- Steiger, S., Currie, T., Brandt, T. D., et al. 2021, *AJ*, 162, 44, doi: 10.3847/1538-3881/ac02cc
- Steiger, S., Brandt, T. D., Guyon, O., et al. 2022b, *The Astronomical Journal*, 164, 186, doi: 10.3847/1538-3881/ac922f
- Strader, M. J. 2016, in *UCSB Electronic Theses and Dissertations*. <https://escholarship.org/uc/item/243538qd>

- Swimmer, N., Clay, W. H., Zobrist, N., & Mazin, B. A. 2023, *Opt. Express*, 31, 10775, doi: 10.1364/OE.485003
- Swimmer, N., Mazin, B., Bockstiegel, C., et al. 2020, in *Ground-based and Airborne Instrumentation for Astronomy VIII*, 79, doi: 10.1117/12.2561770
- Swimmer, N., Currie, T., Steiger, S., et al. 2022a, *The Astronomical Journal*, 164, 152, doi: 10.3847/1538-3881/ac85a8
- Swimmer, N. J., Mazin, B., III, J. I. B., et al. 2022b, in *Ground-based and Airborne Instrumentation for Astronomy IX*, ed. C. J. Evans, J. J. Bryant, & K. Motohara, Vol. 12184, International Society for Optics and Photonics (SPIE), 121843Z, doi: 10.1117/12.2629573
- Szypryt, P., Meeker, S. R., Coiffard, G., et al. 2017, *Optics Express*, 25, 25894, doi: 10.1364/oe.25.025894
- van Eyken, J. C., Strader, M. J., Walter, A. B., et al. 2015, *ApJS*, 219, 14, doi: 10.1088/0067-0049/219/1/14
- van Leeuwen, F. 2007, *A&A*, 474, 653, doi: 10.1051/0004-6361:20078357
- Vigan, A., Fontanive, C., Meyer, M., et al. 2021, *A&A*, 651, A72, doi: 10.1051/0004-6361/202038107
- Walter, A. B. 2019, in *UCSB Electronic Theses and Dissertations*. <https://escholarship.org/uc/item/73t506w9>
- Walter, A. B., Bockstiegel, C., Brandt, T. D., & Mazin, B. A. 2019, *Publications of the Astronomical Society of the Pacific*, 131, 114506, doi: 10.1088/1538-3873/ab389a
- Walter, A. B., Bockstiegel, C., Mazin, B. A., & Daal, M. 2018, *IEEE Transactions on Applied Superconductivity*, 28, 1, doi: 10.1109/TASC.2017.2773836
- Walter, A. B., Fruitwala, N., Steiger, S., et al. 2020, *Publications of the Astronomical Society of the Pacific*, 132, 125005, doi: 10.1088/1538-3873/abc60f
- Wolszczan, A. 1994, *Science*, 264, 538, doi: 10.1126/science.264.5158.538
- Wolszczan, A., & Frail, D. A. 1992, *Nature*, 355, 145, doi: 10.1038/355145a0
- Xie, C., Choquet, E., Vigan, A., et al. 2022, *Astronomy & Astrophysics*, 666, A32, doi: 10.1051/0004-6361/202243379
- Zobrist, N., Clay, W. H., Coiffard, G., et al. 2022, *Physical Review Letters*, 129, doi: 10.1103/physrevlett.129.017701

Zobrist, N., Eom, B. H., Day, P., et al. 2019, Applied Physics Letters, 115, 042601, doi: 10.1063/1.5098469

Zobrist, N., Klimovich, N., Eom, B. H., et al. 2021, Journal of Astronomical Telescopes, Instruments, and Systems, 7, doi: 10.1117/1.jatis.7.1.010501

Zobrist, N. R. 2021, in UCSB Electronic Theses and Dissertations. <https://escholarship.org/uc/item/0rw66345>