

University of California
Santa Barbara

MEC: The MKID Exoplanet Camera for High Speed Focal Plane Control at the Subaru Telescope

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

Doctor of Philosophy
in
Physics - Astrophysics Emphasis

by

Alexander Bryce Walter

Committee in charge:

Professor Benjamin A. Mazin, Chair
Professor Timothy Brandt
Professor Philip Lubin

September 2019

The Dissertation of Alexander Bryce Walter is approved.

Professor Timothy Brandt

Professor Philip Lubin

Professor Benjamin A. Mazin, Committee Chair

August 2019

MEC: The MKID Exoplanet Camera for High Speed Focal Plane Control at the Subaru
Telescope

Copyright © 2019

by

Alexander Bryce Walter

Dedicated to my parents
for their unconditional support

Acknowledgements

The work presented in this thesis would have been impossible without the advice, guidance, support, and direct collaboration of many people. First I will acknowledge those involved in some major technical components required for building MEC.

The readout hardware was designed by Ben Mazin and Paschal Strader with our collaborators at Fermilab. The readout firmware was developed by Sean McHugh, Paschal Strader, and Neelay Fruitwala. The readout crate was assembled by Giulia Collura. The FPGA chips and ADCs were generously donated by Xilinx and Analog Devices, respectively.

The MKID arrays were fabricated by Bruce Bumble at JPL, and Paul Szypryt and Gregoire Coiffard at UCSB. A multitude of testing on MEC arrays or their precursors was performed by Seth Meeker, Clint Bockstiegel, and Nicholas Zobrist.

Incredibly detailed CAD drawings of MEC and the mechanical interface to the SCEXAO bench were made by James Massie. James was also responsible for building the HEMT power supply box.

The LabView temperature control GUI for MEC was appreciatively developed by Clint Bockstiegel.

MEC's integration at the Subaru Telescope was made possible with support from Subaru staff including Oliver Guyon, Julien Lozi, Nemanja Jovanovic (now at Caltech), Matthew Wung, Lucio Ramos, and others. A special thanks to Subaru Day Crew for the many days of ferrying me to the summit, summit safety training, and miscellaneous help.

Night on-sky observations were only possible with support from the SCEXAO team: Olivier Guyon, Julien Lozi, Sebastien Vievard, and Ananya Sahoo. Remote observations from UCSB were made in part with Ben Mazin, Neelay Fruitwala, Isabel Lipartito,

Nicholas Zobrist, Sarah Steiger, Noah Swimmer, Jeb Bailey, Kristina Davis, Rupert Dodkins, and Jenny Smith.

With great gratitude I thank the Japanese Society for the Promotion of Science - JSPS (Grant-in-Aid for Research 26220704) for the bulk of funding for the MEC instrument. We were also assisted by technological developments in the MKID and flexible cable technologies supported by various NSF and NASA grants.

I also want to take this opportunity to thank the many people who personally supported me throughout my graduate school adventure.

First my parents, whom I remember watching complete their own dissertations when I was a child, and for the anecdotes they shared with me that I did not remember. Without their unconditional love and support, leading example, or advice during my grueling days of thesis writing, I would not have made it here.

A special thanks is due to my sister, Ayla Walter, for her endless enthusiasm and interest in my work. I hope to make her proud. Ayla kindly made and provided the clip-art drawing of a sea turtle blocking sunlight like a coronagraph that is at the beginning of each chapter.

To my thesis adviser, Ben Mazin: thank you for the last seven years of guidance, advice, and support. This project would not be possible without your ambitious vision and raw proclivity for innovation. Your leadership during the ARCONS and DARKNESS commissioning runs was crucial, indispensable, and inspirational. I could not ask for a better mentor.

For my friends, from physicists to soccer players, from California to around the world, thank you for your love and support. For Jimmy Chen, Claire McLellan, Jack Bradmiller-Feld, Adam Dishaw, and Shankari Rajagopal you made the first summer of living in California most enjoyable and the following years of dedication worthwhile. A special

thanks to Shankari from which I drew endless inspiration, encouragement, and moral support.

To Jessica Couture, Paul Ore, Rachel Grillot, and the other friends I made on the soccer field (of whom only Jessica has been around from the beginning), thank you for many enjoyable hours on and off the field.

I want to thank the many friends and roommates that allowed me to sleep on their couch or use their rooms during the months before moving to Hawaii.

To the friends I made in Hilo including, but not limited to, Sebastien Vievard, Ananya Sahoo, Amy Durham, and Jon Kunz: thank you for your ceaseless empathy, enduring encouragement, and sharing your passionate worldviews. I want to especially thank Sebastien for his selfless kindness.

Finally, I want to thank my current roommates Andrew Dunsworth, Christina Knapp, Kevin Satzinger, Sam Boland, and Emma Kroloff, for allowing me to live with them and putting up with me during these stressful last months.

Curriculum Vitæ

Alexander Bryce Walter

Education

- 2019 Ph.D., Physics - Astrophysics Emphasis, University of California, Santa Barbara.
- 2015 M.A., Physics, University of California, Santa Barbara.
- 2012 B.S., Physics with Honors, University of Minnesota - Twin Cities
- 2012 B.S., Mathematics, University of Minnesota - Twin Cities

Honors and Awards

- 2019 NASA Postdoctoral Program Fellowship
- 2018-19 Physics Department Service Award
- 2014 Worster Summer Research Fellowship Mentor

Selected Publications

"Stochastic Speckle Discrimination with Time-Tagged Photon Lists: Digging Below the Speckle Noise Floor", **A. B. Walter**, C. Bockstiegel, T. D. Brandt, B. A. Mazin, *PASP*, (accepted).

"Laminated NbTi-on-Kapton Microstrip Cables for Flexible Sub-Kelvin RF Electronics", **A. B. Walter**, C. Bockstiegel, B. A. Mazin, M. Daal, *IEEE Trans. Appl. Supercond.*, 28(1), 2018.

"SCEXAO, an instrument with a dual purpose: perform cutting-edge science and develop new technologies", J. Lozi, O. Guyon, N. Jovanovic, et al., *Proc. SPIE*, 10703, 2018.

"DARKNESS: A Microwave Kinetic Inductance Detector Integral Field Spectrograph for High-contrast Astronomy", S. R. Meeker, B. A. Mazin, **A. B. Walter**, et al., *PASP*, 130(988), 2018.

"The ARCONS Pipeline: Data Reduction for MKID Arrays", J. C. van Eyken, M. J. Strader, **A. B. Walter**, et al., *ApJ Suppl. Ser.*, 219(1), 2015.

Abstract

MEC: The MKID Exoplanet Camera for High Speed Focal Plane Control at the Subaru
Telescope

by

Alexander Bryce Walter

Direct Imaging of exoplanets from the ground is one of the most technically difficult techniques used to study exoplanets, but holds immense promise for not just detecting but characterizing planets around the nearest stars. Ambitious instruments at the world's largest telescopes have been built to carry out this science: the Gemini Planet Imager (GPI), SPHERE at VLT, SCEXAO at Subaru, and the P1640 and Stellar Double Coronagraph (SDC) at Palomar. These instruments share a common archetype consisting of an extreme AO system feeding a coronagraph for on-axis stellar light rejection followed by a focal plane Integral Field Spectrograph (IFS). They are currently limited by uncontrolled scattered and diffracted light which produces a coherent speckle halo in the image plane. A number of differential imaging schemes exist to mitigate these issues resulting in star-planet contrast ratios as deep as $\sim 10^{-6}$ at low angular separations. Surpassing this contrast limit requires high speed active speckle nullification from a focal plane wavefront sensor (FPWS) and new processing techniques.

MEC, the Microwave Kinetic Inductance Detector (MKID) Exoplanet Camera, is a J-band IFS module behind Subaru Telescope's SCEXAO system. MEC is capable of producing an image cube several thousand times a second without the read noise that dominates conventional high speed IFUs. This enables it to integrate with SCEXAO as an extremely fast FPWS while eliminating non-common path aberrations by doubling as a science camera. Key science objectives can be further explored if longer wavelengths

(H and K band) are simultaneously sent to CHARIS for high resolution spectroscopy. MEC, commissioned in March 2018, is the second MKID IFS for high contrast imaging following DARKNESS' debut at Palomar in July 2016.

MEC is designed to follow up on young planets and debris disks discovered in the SEEDS survey or by Project 1640 as well as discover self-luminous massive planets. The increased sensitivity, combined with the advanced coronagraphs in SCExAO which have inner working angles (IWAs) as small as $0.03''$ at $1.2 \mu\text{m}$, allows young Jupiter-sized objects to be imaged as close as 4 AU from their host star. If the wavefront control enabled by MEC is fully realized, it may begin to probe the reflected light of giant planets around some nearby stars, opening a new parameter space for direct imaging targeting older stars. While direct imaging of reflected light exoplanets is the most challenging of the scientific goals, it is a promising long-term path towards characterization of habitable planets around nearby stars using Extremely Large Telescopes (ELTs). With diameters $\gtrsim 30$ m, an ELT can resolve the habitable zones of nearby M-type stars, for which an Earth-sized planet would be at $\sim 10^{-7}$ contrast at $1 \mu\text{m}$. This will complement future space-based high contrast optical imaging targeting the wider habitable zones of sun-like stars for $\sim 10^{-10}$ contrast earth analogs.

This thesis presents the motivation, design, and commissioning of MEC including initial lab and on-sky characterization and results. We include detail on technological advances in cryogenic microwave engineering and a new algorithm for discriminating speckles from exoplanets in post-processing.

Contents

Curriculum Vitae	viii
Abstract	ix
List of Figures	xiii
List of Tables	xv
1 Introduction to Exoplanets and Motivation	1
1.1 Introduction to Direct Imaging	2
1.2 Overcoming the Speckle Barrier	4
1.2.1 Speckle Nulling	5
1.2.2 Linear Dark Field Control	5
1.2.3 Ultimate Contrast Limits at Subaru Telescope	7
1.3 Science Enabled by a Higher Contrast Limit	9
1.4 MEC: The MKID Exoplanet Camera	11
2 The MKID Exoplanet Camera: Design	12
2.1 Microwave Kinetic Inductance Detectors	12
2.1.1 Operating Principle	12
2.1.2 MEC Device Performance	14
2.2 MEC Footprint at Subaru Telescope	15
2.3 Cryostat Mechanical Design	17
2.4 MEC Electronics Rack	20
2.5 Optical Design	22
2.5.1 Specifications of Incoming Beam	22
2.5.2 MEC Fore-Optics	22
2.6 Readout Electronics and Hardware Interface	22
3 MEC Microwave Signal Path	27
3.1 300 K To 4 K Copper Flex Cables	28
3.2 Superconducting Microwave Wiring	28
3.2.1 Introduction to Cryogenic Microwave Wiring	28

3.2.2	Superconducting Flex Cable Design	32
3.2.3	Transmission and Cross Talk Simulations and Measurements . . .	34
3.2.4	Heat Flow Calculations	37
3.2.5	Conclusions on the Microwave Wiring	39
3.2.6	Comments on Performance in MEC	41
3.2.7	Preliminary Work for MEC Upgrade	42
4	Instrument Control, Setup, and Readout Software	43
4.1	Temperature Control of Cryostat	43
4.2	Pixel Calibration and Pre-Setup	44
4.2.1	Resonator Frequency Identification and Power Optimization . . .	44
4.2.2	Choosing the LO Frequency	45
4.2.3	Calculating Pixel Optimal Filters	45
4.2.4	Determining Pixel Location in the Array (Beammapping)	46
4.3	Instrument Setup, Control, and Observing Strategy	46
4.3.1	Readout Board Initialization	46
4.3.2	High Templar: Readout Setup	47
4.3.3	Dashboard: Instrument Control	50
4.3.4	Wavelength, Flat Field, and Pointing Calibration Data	50
5	MEC On-Sky Observations and Data Processing	53
5.1	Calibration of Raw Data	53
5.2	Image Processing	56
6	Stochastic Speckle Discrimination	58
6.1	Introduction to the SSD Algorithm	58
6.2	Simulating Photon Arrival Times	60
6.2.1	Modeling the Stellar Speckle Intensity	60
6.2.2	Modeling an Incoherent Source	61
6.2.3	Generating Mock Photon Lists	62
6.3	SSD with Short Exposure Images	63
6.3.1	Maximum Likelihood Model for Discrete Light Curves	65
6.3.2	Performance of Millisecond Imaging SSD	66
6.4	SSD in the Photon Counting Regime	67
6.4.1	Maximum Likelihood Model for Photon Arrival Times	67
6.4.2	Maximum A Posteriori Estimation	71
6.4.3	Performance on Simulated Telescope Image	73
6.5	Discussion of SSD Simulations	75
6.6	SSD Algorithm Conclusions and Future Work	77
A	Deriving Speckle Noise in a Long Exposure	79
	Bibliography	81

List of Figures

1.1	Contrast of Solar System objects	3
1.2	Speckle nulling demonstration at SCEXAO	6
1.3	Simulation of the ultimate contrast limits at Subaru Telescope	8
1.4	Reflected light contrasts for known RV exoplanets	10
2.1	MKID operating principles	14
2.2	PtSi MKID microscope image	15
2.3	MKID noise measured in lab	16
2.4	MEC arrival to Subaru Telescope	16
2.5	MEC floorplan on Nasmyth IR	17
2.6	Picture of MEC mounted on the Subaru Telescope	18
2.7	MEC cryostat open in the lab	19
2.8	MEC MKID device	20
2.9	MEC electronics rack	21
2.10	Image of MEC's optics box	23
2.11	MEC readout boards	25
2.12	Firmware block diagram for MEC readout boards	25
3.1	MEC cryostat internal microwave wiring and signal path	29
3.2	Mounting an MKID device	30
3.3	Copper flex cable microwave properties	31
3.4	Flex cable design	33
3.5	Flex cable measurement diagram	35
3.6	Measured transmission and crosstalk for flex cable compared to simulation	36
3.7	Equivalent circuit for flex cable	37
3.8	Thermal conductivity of flex cable	38
3.9	Flex cables surrounded by Al tape	41
3.10	Preliminary NbTi flex cable design	42
4.1	MEC Temperature Control GUI	44
4.2	High Templar: setup software GUI	48
4.3	Dashboard: MEC instrument control GUI	51

5.1	MKID pipeline flowchart	54
5.2	Theta ¹ Orionis B in y-band taken by MEC	57
6.1	Simulating lightcurves and calculating the most likely I_c and I_s	64
6.2	ROC curve performance of the millisecond imaging SSD algorithm	68
6.3	ROC curve performance of the photon-counting SSD algorithm	72
6.4	The photon-counting SSD algorithm using simulated telescope images	74

List of Tables

2.1	MEC Instrument Summary	13
2.2	MEC Optics	23
2.3	MEC Neutral Density Filter Wheel	24
3.1	Summary of thermal, mechanical, and microwave properties for the Nb-47Ti flex cable	39
6.1	Companion SSD Signal-to-Noise Ratio	75

Chapter 1



Introduction to Exoplanets and Motivation

Exoplanet observation programs and population surveys over the past decade have led researches to discover that exoplanets are ubiquitous and their masses, sizes, and system architectures are diverse. The exoplanet community seeks to understand which processes and environments inform this diversity of planet/system types and how this relates to their atmospheric composition and structure. Ultimately, we are searching for the signatures of life and the prevalence of conditions conducive to life.

Two exoplanet discovery techniques are able to probe for exoplanet atmospheric composition and structure; transits and direct imaging. Through the fortuitous but rare event of an exoplanet transiting in front of its host star, the absorption of stellar light through the thin ring of the planet's atmosphere can be measured. This is called transit absorption spectroscopy. Although high precision is required, the secondary transit in which the exoplanet travels behind its host star allows a measure of the reflected light albedo. Similarly, if the stellar light can be suppressed well enough, an exoplanet spatially resolved from its host star can have its reflected light directly observed. While high contrast instruments have yet to directly observe an exoplanet in reflected light, the thermal emission of a handful of young, hot exoplanets have been measured with

the direct imaging technique. Absorption, emission, and reflected light spectroscopy can provide different and complementary information on the composition and structure of clouds or the photochemistry of hazes (Morley et al. 2015). The transit and direct imaging techniques are also biased but complementary in their targeted exoplanet populations. Transiting surveys are sensitive to small planets, with small separations, over a large volume of space. Direct imaging surveys are sensitive to large planets, at wide separations, and only around nearby stars.

This thesis is focused on high contrast imaging for the discovery and characterization of exoplanets with the direct imaging technique.

1.1 Introduction to Direct Imaging

Direct imaging is a challenging exoplanet discovery and characterization technique due to the extreme contrast ($<10^{-4}$ for ground based targets) and small angular separations ($\lesssim 1''$) between the planetary companion and its stellar host (for a review see Bowler (2016)). For example, Earth is 10^{10} times fainter than the sun in the visible (see Figure 1.1). Sophisticated ground-based high contrast instruments have been or are being built including: the Gemini Planet Imager (GPI, Macintosh et al. 2008; Macintosh et al. 2014), SPHERE at VLT (Carillet et al. 2011), SCEXAO at Subaru (Jovanovic et al. 2015), P1640 (Crepp et al. 2011) and the Stellar Double Coronagraph (SDC, Mawet et al. 2014b) at Palomar, the Keck Planet Imager and Characterizer (KPIC) at Keck (Mawet et al. 2016), and MagAO-X at the Magellan Clay Telescope (Males et al. 2018). Adaptive optics (AO) and coronagraphy have enabled the discovery of planets up to $\sim 10^6$ (15 magnitudes) times fainter than their host stars which has allowed us to survey formation conditions in primordial disks and the temperature of a handful of giant young exoplanets in the stellar neighborhood (Marois et al. 2008; Lagrange et al. 2010;

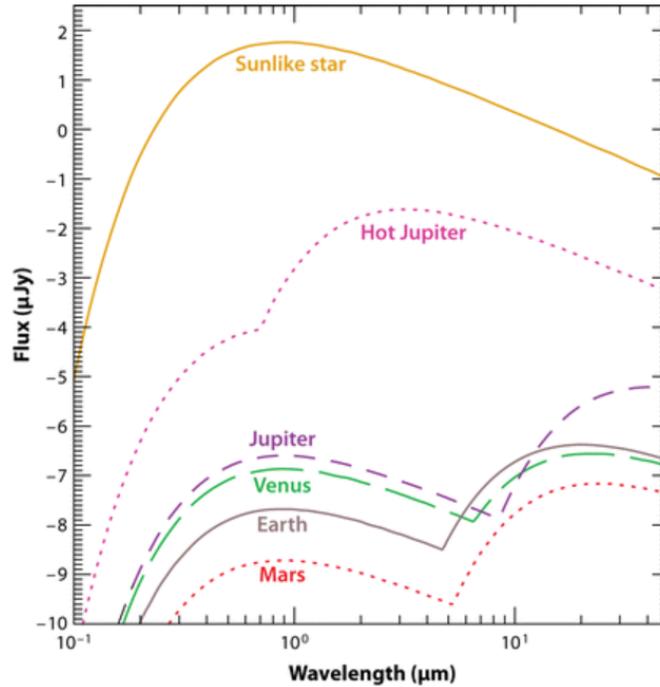


Figure 1.1: From Seager and Deming (2010). Black body flux of several Solar System bodies and a putative hot Jupiter, as seen from 10 pc away. The planets have two peaks in their spectra. The short-wavelength peak is due to sunlight scattered from the planet atmosphere and is computed using the planet’s geometric albedo. The long-wavelength peak is from the planet’s thermal emission and is estimated by a black body of the planet’s effective temperature.

Kuzuhara et al. 2013; Macintosh et al. 2015; Keppler et al. 2018).

High-contrast imaging is limited by uncontrolled scattered and diffracted light, which produces a coherent speckle halo in the image plane. Speckles can arise from atmospheric aberrations or from the non-ideal nature of the instrument which suffers from noisy near-IR detectors, time lag in the AO correction, speed of the AO control loop, non-common path errors, telescope vibrations, and chromaticity between the wavefront and science cameras (Guyon 2005; Lozi et al. 2018a). Fast atmospheric speckles average down over an observation, while slower, quasistatic speckles must be removed using post-processing techniques. Angular differential imaging (ADI, Marois et al. 2006) exploits the rotation of the Earth, and hence the field-of-view of an altitude-azimuth telescope, to distinguish

diffraction speckles from astrophysical sources. Spectral differential imaging (SDI, Racine et al. 1999; Marois et al. 2000; Sparks and Ford 2002) uses the scaling of diffraction speckles with wavelength. Since the initial development of ADI and SDI, a variety of post-processing algorithms have refined their approaches to dig deeper into the stellar PSF (e.g. Lafrenière et al. 2007; Soummer, Pueyo, and Larkin 2012; Marois et al. 2014; Gomez Gonzalez, C. A. et al. 2016).

The time variability and chromaticity of quasi-static speckles limit the performance of ADI and SDI (Gerard et al. 2019). Both techniques also suffer at small separations where exoplanets are more likely to hide. The speckle spectral dispersion used by SDI is proportional to the separation: close to the star, it becomes smaller than the planet’s PSF. For ADI, the arclength traced by the companion’s sky rotation is proportional to the separation. Furthermore, the precision of the background estimate for PSF subtraction is limited by low counting statistics at small separations (Mawet et al. 2014a). Even without these issues, the variability induced by speckle fluctuations can dominate the photon noise and be well above the shot noise expected from the total number of photons.

1.2 Overcoming the Speckle Barrier

Speckles that act as noise by masquerading as astrophysical sources must be overcome by improving the raw contrast of the starlight suppression in the image plane as well as improving post-processing techniques.

Coronagraphs that block on-axis stellar sources have matured to the point of reaching very high performance and small inner working angle (IWA, Guyon et al. 2010; Mawet et al. 2012; Ruane et al. 2018). Yet, we are not reaching their full potential because of our limited ability to sense wavefront errors. In short exposures, these result in fast-evolving atmospheric speckles (due to the inability of the AO system to sense fast enough and

chromatic effects due to the wavefront sensor typically operating at a different wavelength than the science camera), as well as more slowly evolving speckles. Both errors can be addressed by focal plane speckle control, an approach which is by construction free of non-common path errors. Residual atmospheric speckles can be corrected if using a high-speed photon-counting detector at the science focal plane. Such instruments also enable our ability to improve our effective contrast (beyond the raw contrast) through post-processing techniques.

1.2.1 Speckle Nulling

Usually it is impossible to measure the electric field phase from light observed as intensity in the image plane. We can use the deformable mirror (DM) to purposely reflect star light with known phase into the image on top of a speckle and measure the coherent interference. In this way we can determine the speckle's phase and null it with an artificial speckle with the opposite phase. This is called speckle nulling. With this technique we can create a dark hole to improve the contrast limit in half the image plane (see Figure 1.2). This technique has been demonstrated on-sky at low speeds (<10 Hz) where it is able to remove static and some quasi-static speckle (Martinache et al. 2014; Jovanovic et al. 2015; Bottom et al. 2016). To remove fast atmospheric speckles we would need a control bandwidth of at least 100 Hz which would require kHz camera frame rates.

1.2.2 Linear Dark Field Control

Linear dark field control (LDFC) is a technique that can stabilize the dark hole in one half of the image using the bright field in the other half (Miller, Guyon, and Males 2017). It is based on the idea that the aberrations that produce the speckle pattern in one half of the image are the same aberrations that produce the speckle pattern in the other

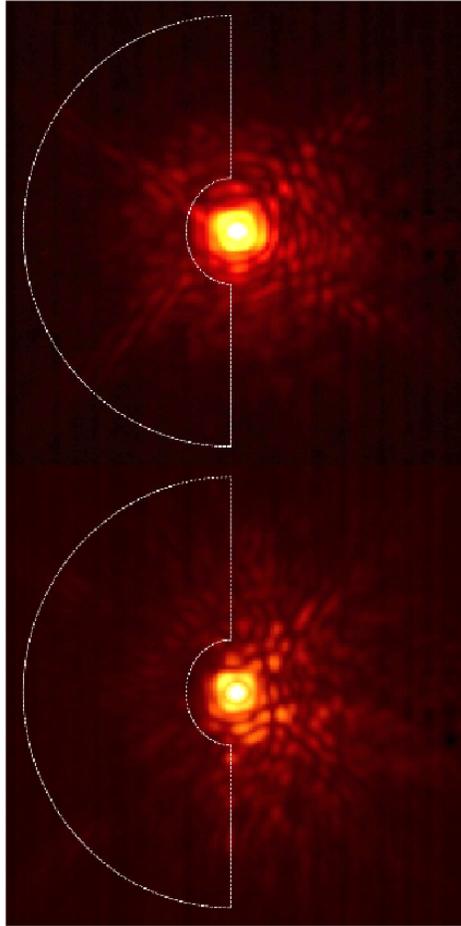


Figure 1.2: From Jovanovic et al. (2015). Top: RX Boo from SCEXAO with no speckle nulling applied. Bottom: RX Boo with speckle nulling performed on the region enclosed by the white dashed line. Each image is a composite of 5000, $50 \mu s$ frames which have been shifted and added together.

half. Since the bright field has high signal to noise and the aberrations are assumed to be linear this control algorithm is relatively simple. Deployment on a real system becomes challenging due to interactions with other control loop(s) and the deviation from linearity in the presence of real atmospheric turbulence and noise. The high order correction modes and possibly some lower order modes could be non-linear depending on the atmospheric conditions. LDFC needs to select the modes that are linear in order to implement a stable control loop.

1.2.3 Ultimate Contrast Limits at Subaru Telescope

There is often confusion between the raw contrast delivered by an AO system and coronagraph, typically around 10^{-3} to 10^{-5} , with the quoted effective contrast after post-processing, around 10^{-6} in the best direct imaging instruments. What this means is that in the focal plane the atmospheric speckles have typical photon fluxes around 10^{-4} of the host star. A typical target star at 5th magnitude in J at Subaru and a conservative 3% estimate of the total efficiency of the atmosphere, telescope, AO system, and coronagraph results in 50 million photons/second from the host star. At a raw contrast level of 10^{-4} this means an atmospheric speckle has a flux of 5,000 photons/second. The non-speckle background due to sky and scattered light will be very low, so we expect to be able to detect a speckle in roughly 10 ms and cancel it within 50 ms. Given an atmospheric speckle lifetime of around 1 second, we should be spending around 95% of the observing time in any given region of interest within the dark hole integrating with highly suppressed speckles. The speckle-suppressed dark hole data will then be further processed to increase the final contrast ratio using more standard high contrast image processing like ADI and SDI.

By applying the focal plane wavefront techniques outlined above we should be able to approach photon shot noise limited performance. Figure 1.3 shows the expected high contrast imaging performance of SCExAO operating at the photon noise limit on a typical nearby M-type star with $m_I = 8$ and $m_H = 6$. The optimal wavefront sensing wavelength is approximately $1 \mu\text{m}$ for early to mid M-type stars due to a compromise between photon count and wavelength-dependent optical gain. The raw contrast then dips below 10^{-4} at small separations (1 to $5 \lambda/D$), and contains slow speckles due to atmosphere chromaticity and AO loop time lag. By adding near-IR focal plane speckle sensing and assuming that the sensing is done at the photon noise limit, the final detection contrast is $\sim 10^{-7}$; very

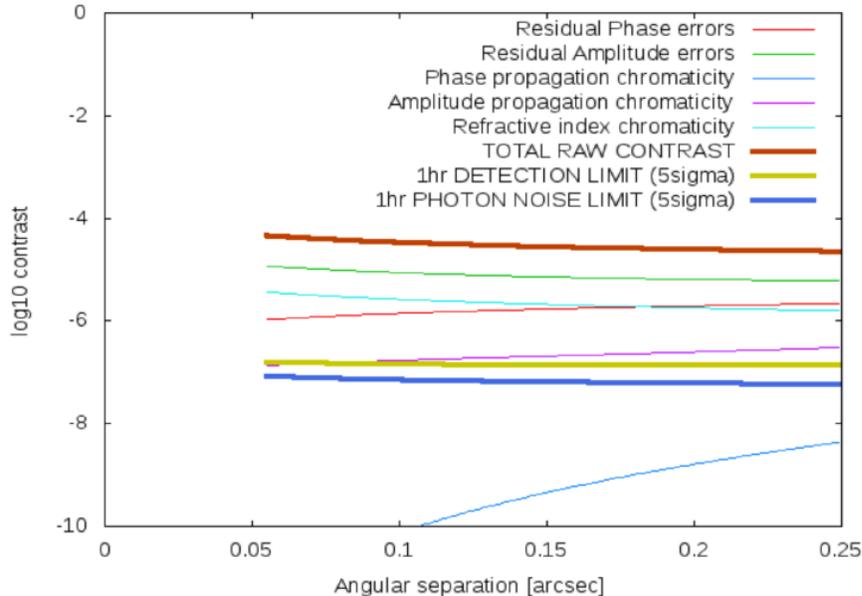


Figure 1.3: Simulation of the performance of SCEXAO with a diffraction limited (pyramid) wavefront sensor and MKID focal plane for speckle control on a $m_I = 8$ and $m_H = 6$ star. Assumptions in this simulation include $1.3 \lambda/D$ inner working angle, 30% system efficiency, 40% bandwidth in WFS and MKIDs, 1.5 WFS frames time lag, and median Mauna Kea seeing. Plot from Olivier Guyon private communication. See Lozi et al. (2015) for simulation with TMT.

close to the photon noise limit in the science images. Even deeper raw and detection contrasts are theoretically possible with AO predictive control (Guyon and Males 2017; Males and Guyon 2018) and focal plane speckle control at the science wavelength. The focal plane sensing loop should operate at ~ 300 Hz, requiring \sim kHz frame rate: slower sensing yields poorer detection limit due to temporal speckle noise.

As a caveat, these results should be taken as an optimistic limit. Reliable derivation of the ultimate performance limits are notoriously challenging given the large gap between current achieved performance and our predictions. While first-order derivations suggest that 10^{-8} detection contrast could theoretically be reached, the required techniques have not been validated on-sky, and have only been explored through semi-analytical derivations. On-sky testing is essential to validate approaches and develop the

practical system-level experience which will guide development of instruments on current and future (ELT) telescopes.

1.3 Science Enabled by a Higher Contrast Limit

The increased sensitivity, combined with small inner working angle coronagraphs (0.03" at 1.2 μm), allows young Jupiter-sized objects to be imaged as close as 4 AU from their host star. Spectroscopy on these objects can distinguish between rocky or icy bodies and thick atmospheres, and constrain their physics properties and chemical composition. Since the mass-luminosity relationship is extremely model dependent, characterizing the atmospheres of massive planets will reduce the systematic uncertainty in the estimated mass of directly imaged planets. Masses from radial velocity (RV) measurements can be difficult to obtain due to the large semi major axes involved and large variability in young stars. These young systems are critical for understanding giant planet formation and the initial conditions for exoplanetary system evolution of which there are precious few empirical constraints.

If 10^{-7} contrasts are truly achieved we will begin to probe the reflected light of giant planets around some nearby stars, a feat once thought only achievable in space. This opens a new parameter space for direct imaging of older stars. Figure 1.4 shows the estimated reflected-light contrast vs. angular separation for known RV exoplanets. Several should have observable contrasts for direct imaging, such as the following nearby (<5 pc) RV targets accessible from Mauna Kea:

- GJ876b, a 2.3 Jupiter-mass planet at 45 mas separation, contrast $\sim 9\text{e-}7$
- GJ687b, a 19 Earth-mass planet at 36 mas separation, contrast $\sim 2\text{e-}7$
- HD219134d, a 16 Earth-mass planet at 35 mas separation, contrast $\sim 1\text{e-}7$

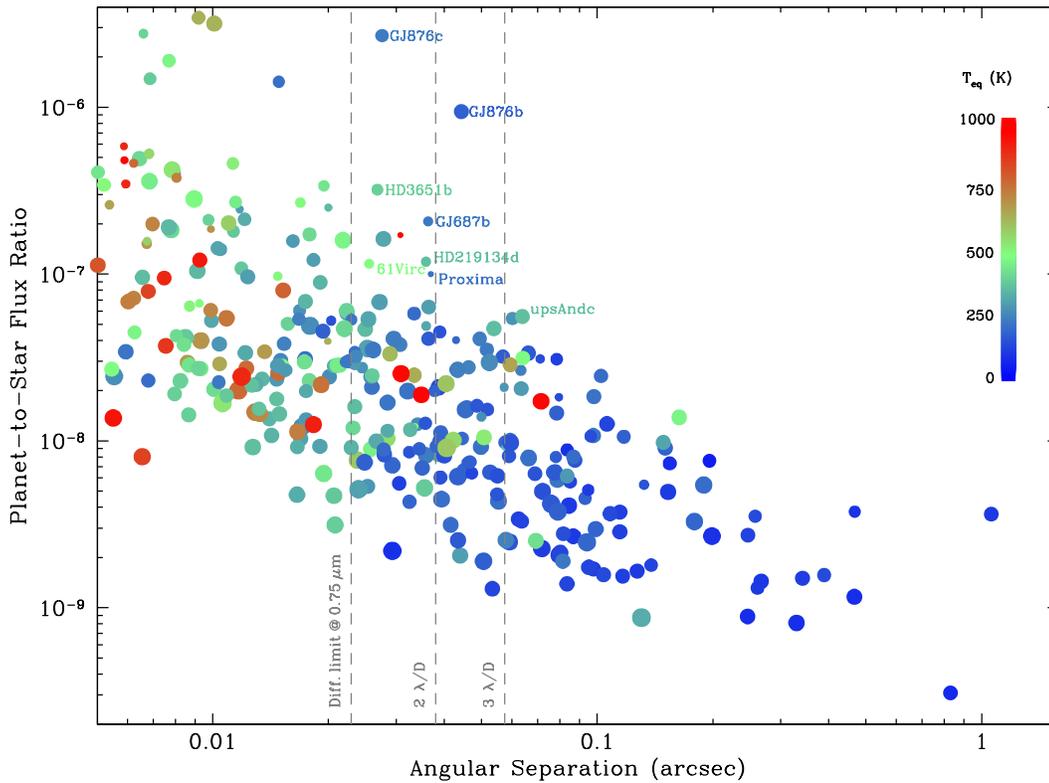


Figure 1.4: Estimated planet-to-star contrast in reflected light for known exoplanets as a function of angular separation from their host star. Dot size is proportional to the logarithm of planet mass, while the color scale represents equilibrium temperature (assuming a Bond albedo of 0.3). Vertical dashed lines indicate the diffraction limit, $2 \lambda/D$ and $3 \lambda/D$ thresholds for the 8.2-m VLT at 750 nm. Figure reproduced from Lovis et al. (2017).

- GJ876e, a 15 Earth-mass planet in the outer part of the habitable zone, 72 mas separation, contrast $\sim 7e-9$ (a challenging target!)

We can exceed RV sensitivity in nearby stars: in the habitable zone of ϵ Eridani, the expected contrast detection limit corresponds to a 2 Earth radius planet, which would have escaped both astrometric and RV detection attempts. Imaging of orbital parameters, timed photometry, and high resolution spectroscopy of reflected light planets will allow probes into their radius (albedo), atmospheric composition, and rotational period.

While direct imaging of reflected-light exoplanets is the most challenging of the scientific goals, it is a promising long-term path to characterization of habitable planets

around nearby stars like Proxima Cen, which becomes possible with the next generation of Extremely Large Telescopes (ELTs). With diameters of about 30-m, an ELT equipped with low-IWA coronagraph can resolve the habitable zones of nearby M-type stars, for which an Earth-sized planet would be at a $10^{-7} - 10^{-8}$ contrast (Skidmore et al. 2015).

1.4 MEC: The MKID Exoplanet Camera

To breach the speckle contrast barrier in high contrast imaging, we have constructed the MKID Exoplanet Camera (MEC) for Subaru Telescope’s SCEXAO system. MEC uses a high-speed, noise-free Microwave Kinetic Inductance Detector (MKID) array designed to implement a fast feedback loop for focal plane wavefront control as well as improved post-processing. MEC is currently following up on known young planets and debris disks discovered in the SEEDS survey or by Project 1640 as well as searching for new self-luminous massive planets. Even if only small improvements are made to the AO performance at small inner working angles, it will enhance current observations, boost detection efficiency, enable observations in worse seeing conditions, and may find new exoplanet candidates. With the same improvements scaled up to a 30 m telescope, Earth-sized habitable planets around nearby M-dwarfs may be feasible. This makes MEC an important technological and algorithmic test-bed for future ELTs.

This thesis describes MEC. Chapter 2 provides the detailed mechanical design of MEC. Chapter 3 describes the microwave readout signal path through MEC which required significant technological advances over the course of this work. Chapter 4 summarizes the software development, algorithms, and working knowledge of the instrument control software. Chapter 5 describes the basic post-processing calibration procedures for MKID-like data. Finally, Chapter 6 describes a new post-processing technique that discriminates between speckles and planets using only the photon arrival time statistics.

Chapter 2



The MKID Exoplanet Camera: Design

We introduce MEC, the MKID Exoplanet Camera, commissioned at Subaru Observatory in early 2018. MEC is designed as a y , J-band IFU for high contrast imaging and as a focal plane wavefront sensor for a multi-kHz feedback loop with the Subaru Coronagraphic Extreme Adaptive Optics (SCEXAO) system. The enabling technology are Microwave Kinetic Inductance Detectors (MKIDs) which are capable of color discriminating near-infrared photons with read noise free microsecond photon timing. The unique capabilities of MKIDs coupled with SCEXAO provide a powerful means of overcoming the limits of ground based high contrast imaging and acts as a testbed for technologies aimed at future 30 m class telescopes. Important instrument parameters are summarized in Table 2.1.

2.1 Microwave Kinetic Inductance Detectors

2.1.1 Operating Principle

MKIDs work on the principle that incident photons change the surface impedance of a superconductor through the kinetic inductance effect. The kinetic inductance effect occurs because energy can be stored in the supercurrent (the flow of Cooper Pairs) of

Table 2.1: MEC Instrument Summary

Parameters	Values
Device Materials	PtSi on Sapphire w/ Nb ground plane
Device Format	140x146 (10 feed lines)
Pixel Pitch	150 μm
Plate Scale	10.4 mas/pixel
Field of View	1.4" x 1.5"
Wavelength Band	800-1400 nm (y, J)
Spectral Resolution ($\lambda/\Delta\lambda$)	5-7
Max Count Rate	5000 cts/second
Pixel Dead Time	10 μs
Readout Frame Rate	>2 kHz
Operating Temp.	90 mK
4 K Stage Base Temp.	3.1 K
60 K Stage Base Temp.	57 K
Cryostat Hold Time	17 hours

a superconductor. Reversing the direction of the supercurrent requires extracting the kinetic energy stored in it, which yields an extra inductance term in addition to the familiar geometric inductance. The magnitude of the change in surface impedance depends on the number of Cooper Pairs broken by incident photons, and hence is proportional to the amount of energy deposited in the superconductor. This change can be accurately measured by placing a superconducting inductor in a lithographed resonator. A microwave probe signal is tuned to the resonant frequency of the resonator and any photons which are absorbed in the inductor will imprint their signature as changes in phase and amplitude of this probe signal.

Since the quality factor, Q , of the resonators is high and their transmission off resonance is nearly perfect, multiplexing can be accomplished by tuning each pixel to a different resonant frequency with lithography during device fabrication. A comb of probe signals is sent into the device and room temperature electronics recover the changes in amplitude and phase. This makes a device capable of measuring the arrival time (to a microsecond) and energy (to 5-10%) of each arriving photon without read noise or dark

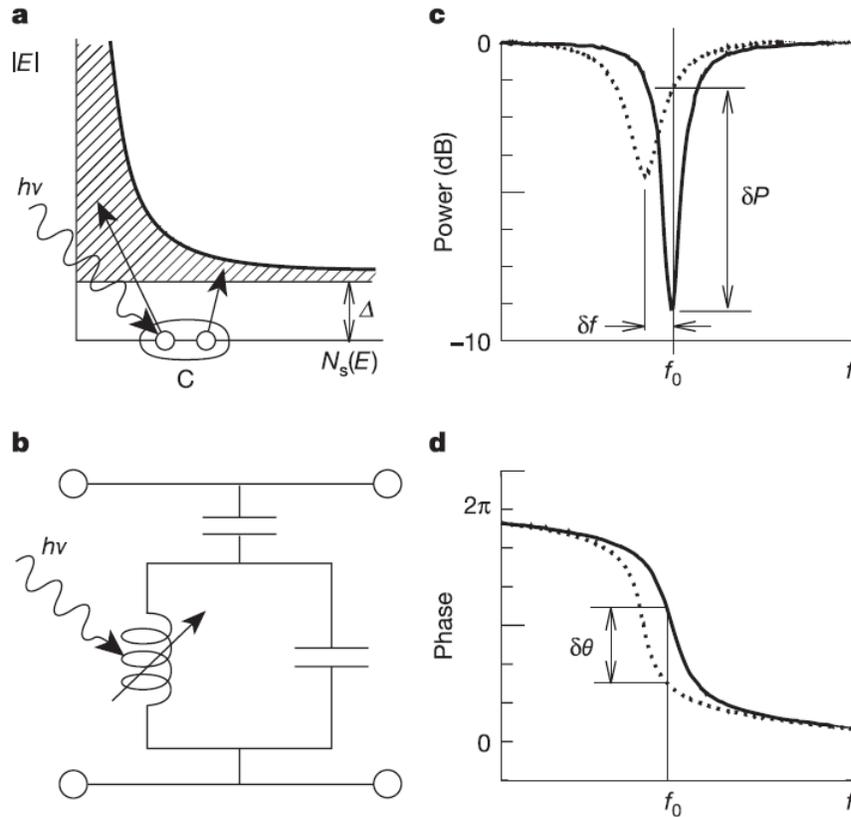


Figure 2.1: The basic operation of an MKID, from Day et al. (2003). (a) Photons with energy $h\nu$ are absorbed in a superconducting film producing a number of excitations called quasiparticles. (b) To sensitively measure these quasiparticles the film is placed in a high frequency planar resonant circuit. The amplitude (c) and phase (d) of a microwave excitation signal sent through the resonator. The change in the surface impedance of the film following a photon absorption event pushes the resonance to lower frequency and changes its amplitude. If the detector (resonator) is excited with a constant on-resonance microwave signal, the energy of the absorbed photon can be determined by measuring the degree of phase and amplitude shift.

current; an optical/near-IR analog of an X-ray microcalorimeter. More details on MKIDs can be found in Mazin et al. (2012) and Mazin et al. (2013) and in Figure 2.1.

2.1.2 MEC Device Performance

The MEC MKID array has 10 feedlines that each probe 14×146 pixels for a total of 20440 pixels in the array. This makes it the largest superconducting detector in the

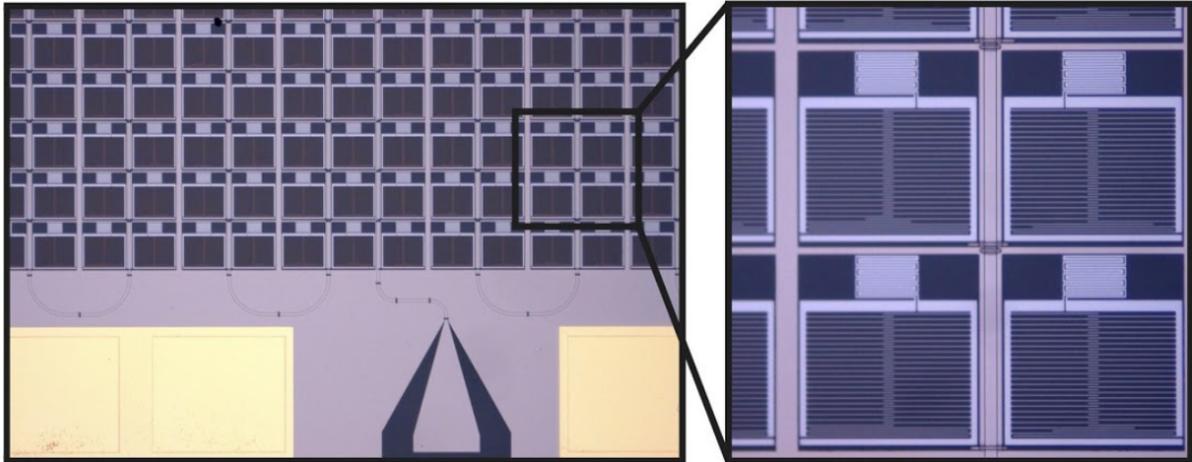


Figure 2.2: Microscope image of a PtSi MKID made for the DARKNESS instrument. Left: A feedline connects to a grid of MKID pixels. Right: zoomed in picture of an MKID showing the large interdigitated capacitor and meandered inductor.

world. It uses the PtSi MKIDs described in Szypryt et al. (2017) which are optimized for 800-1400 nm light with a spectral resolution of about $\lambda/\Delta\lambda = 6$. The pixels are on a $150 \mu\text{m}$ pitch. See Figure 2.2. The device currently in MEC has 7/10 feed lines with good transmission and on those there is about 80% pixel yield.

Measurements of the typical pixel's phase noise are shown in Figure 2.3. Noise from the HEMT amplifier is equally important to noise from two level systems (TLSs) that arise from the resonator's electric field coupling to dipole moments from impurities in the substrate surface layer (Gao 2008). Both noise sources need to be addressed for there to be fundamental improvements in phase noise and thus MKID spectral resolution.

2.2 MEC Footprint at Subaru Telescope

MEC arrived at Subaru on Mauna Kea on March 19, 2018 (Figure 2.4). During the first week it was inspected for damage in sensitive equipment including the MKID array, internal microwave wiring, ADR suspension, and electronics boards. All components arrived safely. The cryostat is currently cooled down to a base 3.1 K where it will stay

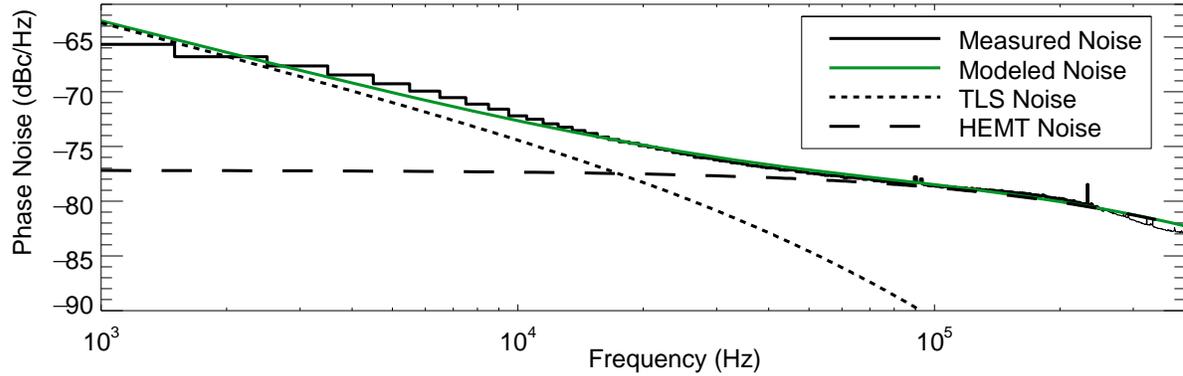


Figure 2.3: Lab measurement of a typical MEC MKID pixel noise versus frequency. The modeled TLS noise plus HEMT amplifier noise are fit to the measured noise showing that both sources are important.



Figure 2.4: Ben Mazin (left) and Julien Lozi (right) pull the MEC cryostat and components crate underneath the Subaru Telescope 8.2 m primary on the way to Nasmyth IR port.

unless upgrading the MKID array.

MEC is mounted behind SCEXAO in the place of HiCIAO. An Electronics rack (E-rack) is stationed next to MEC on the Nasmyth floor. A He compressor is located in the

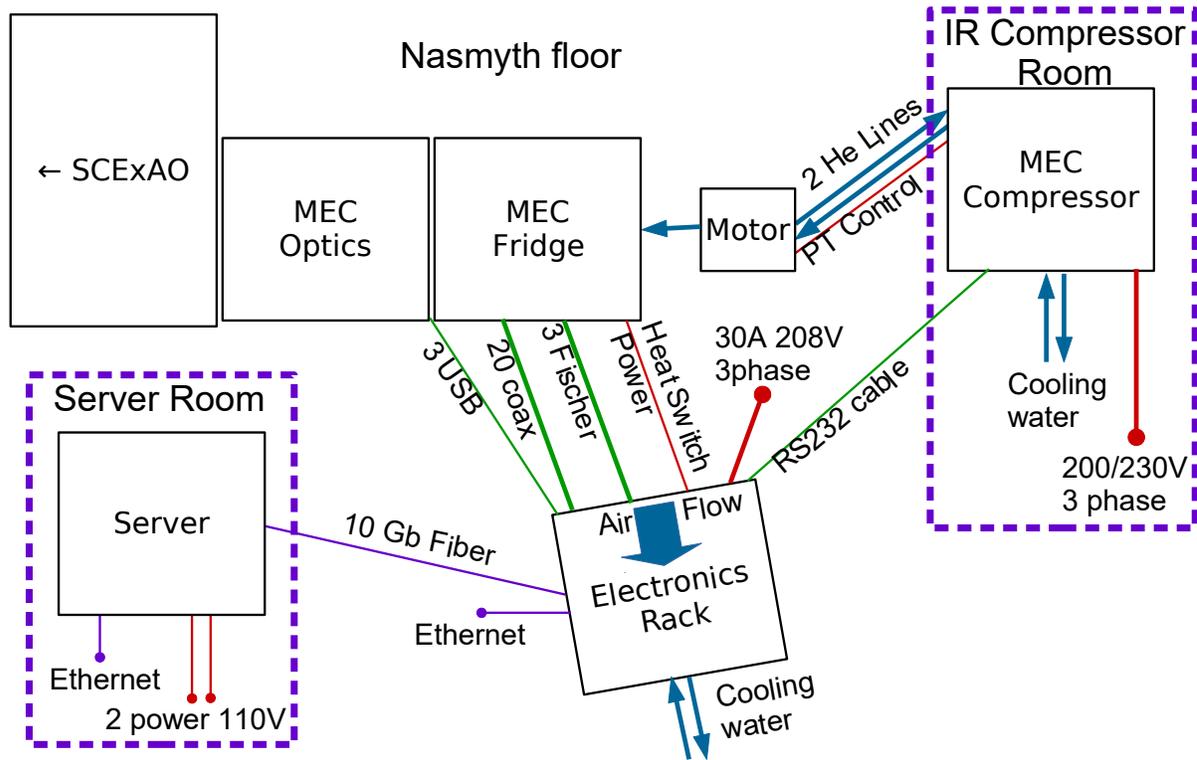


Figure 2.5: A diagram showing the MEC floor plan and interconnects at the Subaru Telescope.

IR compressor room on the observation floor and a linux server is in the server room on the first floor of the control building. Figure 2.5 diagrams how MEC is situated on the Nasmyth floor. Figure 2.6 is an image of MEC mounted at the telescope.

2.3 Cryostat Mechanical Design

MEC's cryostat is a pulse tube cooled Adiabatic Demagnetization Refrigerator (ADR) capable of reaching temperatures below 50 mK. The outer vacuum shell 300 K enclosure measures roughly 22 cm wide x 33 cm deep x 96 cm tall. It contains RF gaskets to reduce radio frequency interference, a 7.5 kW Cryomech PT-407 pulse tube, and all the vacuum and wiring connections to accommodate a 20 kpix MKID array. The open MEC cryostat can be seen in Figure 2.7.

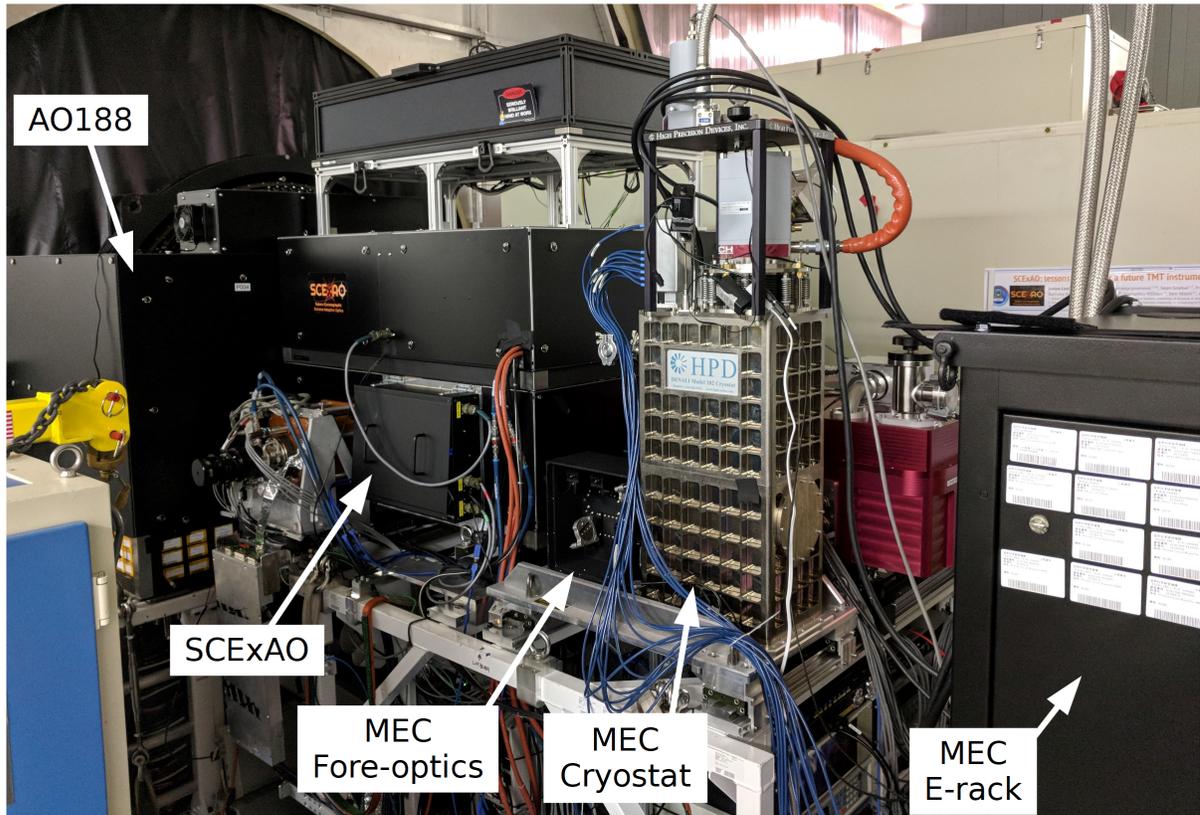


Figure 2.6: MEC is mounted behind the SCExAO bench and AO188 facility adaptive optics systems on the IR Nasmyth port at Subaru Telescope.

The 60 K shell is cooled by the first stage of the Cryomech pulse tube. It contains heat sinks to cool the wiring to reduce heat load on the 4 K stage. The 4 K stage is cooled by the second stage of the pulse tube, and has 0.75 Watts of cooling power at 4 K.

Attached to the 4 K plate are: the ADR unit with mechanical heat-switch, the ten cryogenic HEMT amplifiers from Low Noise Factory, and the detector package. The detector package is mounted to the 4 K plate and is enclosed in an Amumetal 4 K magnetic shield. Carbon fiber support structures hold the intermediate 1 K stage above the 4 K base, and the MKID array hangs from the 1 K stage by Vespel SCP-5050 suspension. This design places the MKID array far from the magnetic shield opening where field leakage will be strongest and also allows us to install a 1 K, 9 cm long black



Figure 2.7: Image of Alex Walter standing next to the open MEC cryostat

baffle to further reduce off-axis scattered light and 4 K black body radiation. The MKID array is mounted in a microwave package 5 cm x 5 cm x 2 cm in size, as shown in Figure 2.8. This box is attached to a gold-plated copper rod that sticks out of the base of the magnetic shield and connects to the ADR unit by a copper strap.

The ADR acts as a single-shot magnetic cooler which brings the MKID array down to 90 mK where the temperature is stabilized with a feedback loop to the ADR magnet power supply. We achieve a 90 mK hold time of >17 hours on the telescope, more than sufficient for the 13-14 hours required for a night of calibration and on-sky observations.

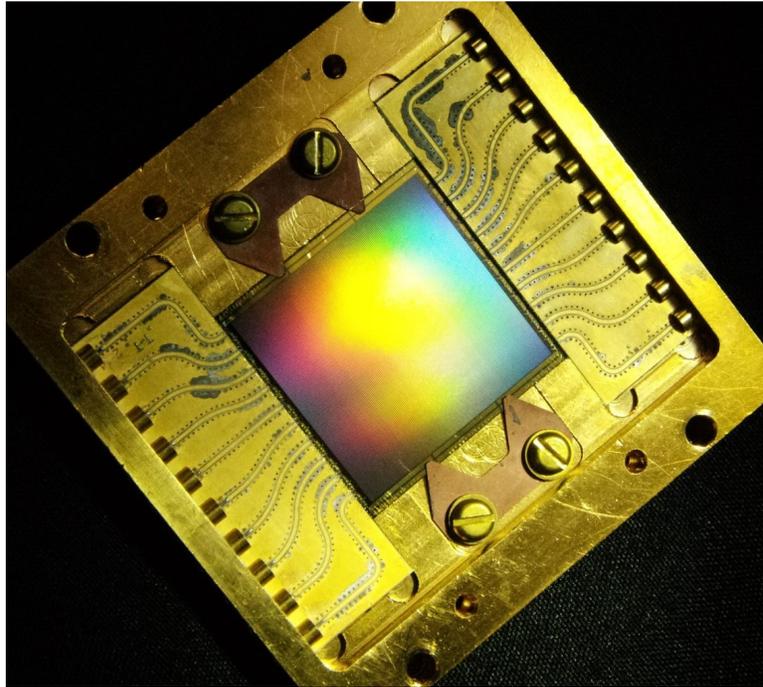


Figure 2.8: Image of a 20440 pixel MKID device designed for MEC.

2.4 MEC Electronics Rack

The electronics rack (E-rack) is situated on the Nasmyth platform next to the MEC instrument (Figure 2.9). It requires 3.5 kW max of power supplied to the Eaton PC975-LT PDU with a 30 A, 208 V 3-phase power cord. Cooling water/glycol is supplied to the E-rack by Subaru. The rack has 3/8 inch QC swagelok bulkhead fittings. The inlet connector is male in order to mate with the Subaru facility female supply connector. The outlet side has the opposite gender. Pressure provided is 40-60 PSI and a flow monitor records coolant flow through E-rack. There should be unobstructed air flow to two readout crates in the electronics rack. The front (facing SCEXAO) draws air in, and the back expels air through a heat exchanger into the Nasmyth room.

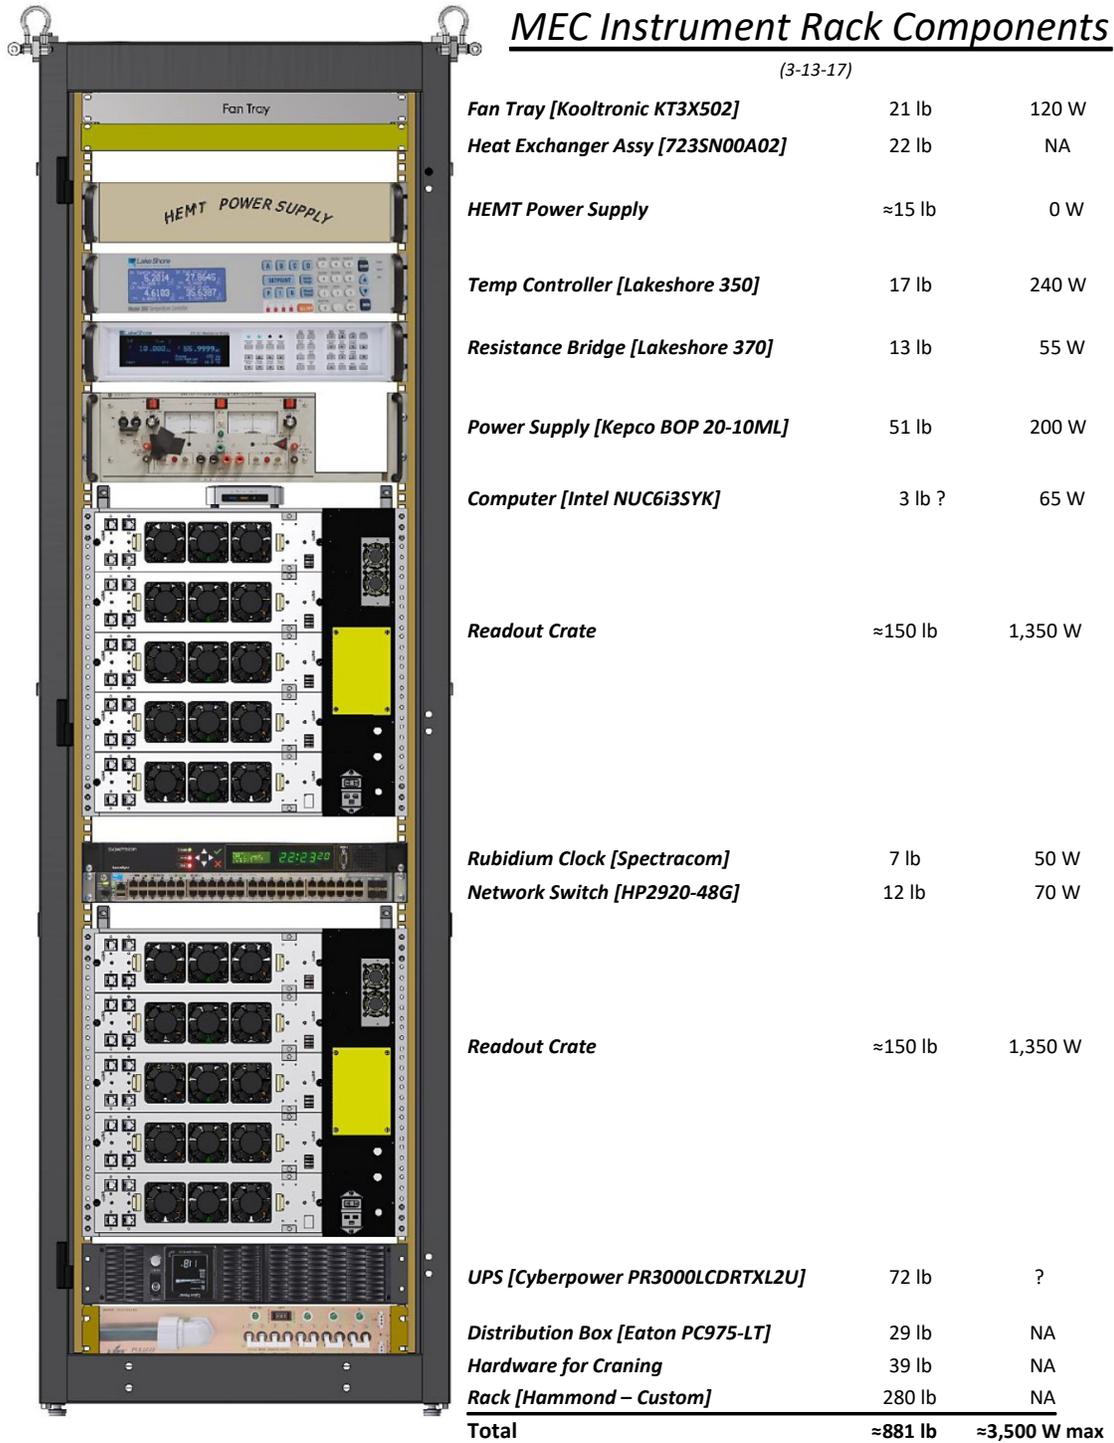


Figure 2.9: CAD drawing of MEC electronics rack (E-rack) and list of components.

2.5 Optical Design

2.5.1 Specifications of Incoming Beam

The SCExAO exit beam is located 293 ± 1 mm from the top of the top mounting rail. The beam comes off the bench at about 58 mm from the left edge of the bench. The SCExAO exit beam is collimated. The pupil position is about 54 ± 20 mm from the back of the bench. The pupil size is 8.67 ± 0.15 mm. The beam is expected to exhibit <10 nm of chromatic RMS optical wavefront error over the 0.8-1.4 μm bandwidth.

2.5.2 MEC Fore-Optics

The MEC optics box has a two inch hole for the incoming SCExAO beam (Figure 2.10). A set of three lenses creates a telecentric beam designed to reimaging the focus onto the MKID array as simulated in Zemax. The beam leaves the optics box with an f-number $f/377.7$, travels through the neutral density (ND) filter wheel (Table 2.3), and enters the cryostat through the front vacuum window. There are two IR bandpass filters at the 60 K and 4 K shield to block black body radiation and finally a microlens array focuses the beam onto the light sensitive part of the MKID pixels. The MEC optics are listed in Table 2.2.

2.6 Readout Electronics and Hardware Interface

The readout for MEC is based on the ARCONS readout, which is detailed in McHugh et al. (2012) and is identical to the DARKNESS readout detailed in Strader (2016). The MEC readout employs twenty ROACH2 (Reconfigurable Open Architecture Computing Hardware) boards produced by the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) that connect to ADC/DAC and RF/IF boards designed

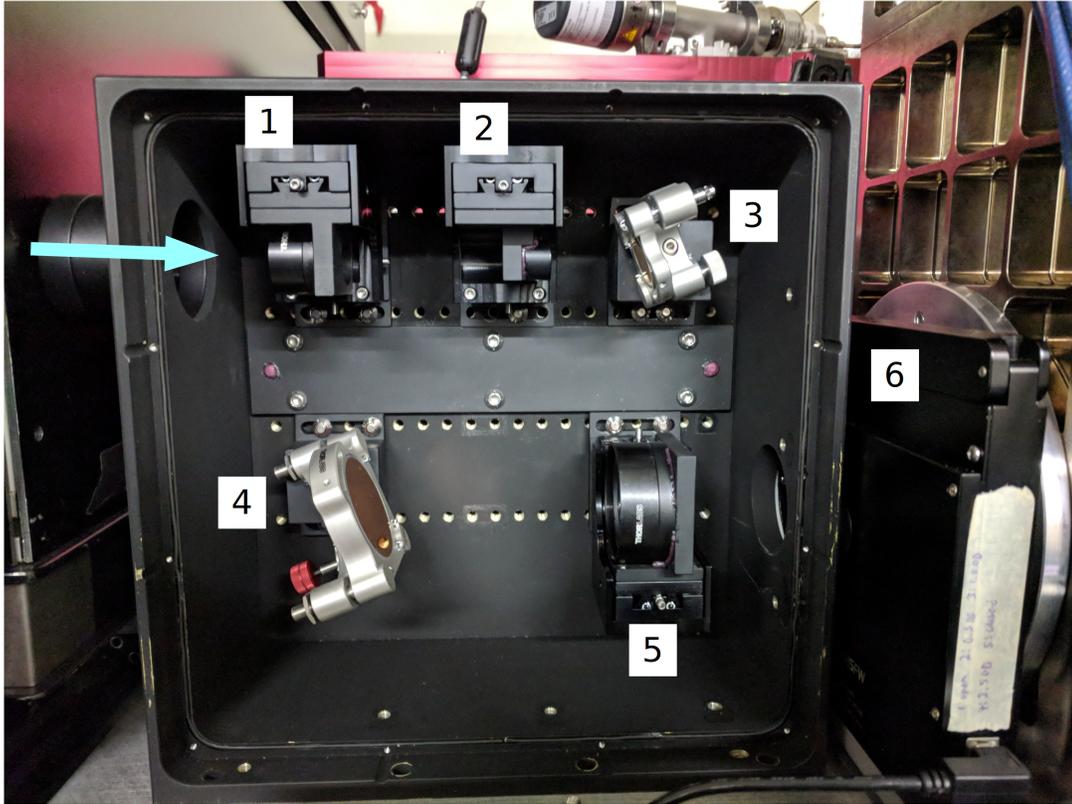


Figure 2.10: The MEC foreoptics box reimages the beam from SCEXAO (arrow from left) onto the MKID device. The optics components are numbered and correspond to Table 2.2.

Table 2.2: MEC Optics

#	Optic	Description
1	100 mm FL lens	Thorlabs AC254-100-c
2	9 mm FL lens	Edmunds 45-783
3	2.54 mm gold mirror	Mounted on CONEX-AG-M100D piezo stage
4	5.08 mm gold mirror	Fixed mount
5	300 mm FL lens	Thorlabs AC508-300-c
6	2 in ND Filter wheel	See Table 2.3
7	Front Window	Edmund 48-130
8	60 K IR filter	10 mm thick N-BK7
9	4 K IR filter	20 mm thick N-BK7
10	Microlens Array	μ s APO-GT-P150-R0.8

at Fermilab (See Figure 2.11 for one board set). Each set of boards reads out 1024 pixels in 2 GHz of bandwidth. Each Roach2 houses a Virtex6 FPGA, which processes and

Table 2.3: MEC Neutral Density Filter Wheel

Slot	Filter	Part #
1	OPEN	
2	0.5 OD	Edmund 62-673
3	1.0 OD	Edmund 62-676
4	2.5 OD	Edmund 62-680
5	CLOSED	

channelizes the signal from the ADC using a scaled up version of the ARCONS firmware. This firmware is written in Simulink and is compiled using the Xilinx System Generator. The ADC/DAC board houses a Virtex7 FPGA with firmware written using Xilinx Vivado. The Roach2 and ADC/DAC board are connected by two Z-Dok connectors. One of the Z-Dok pins allows the ADC/DAC firmware to communicate with the ROACH2 firmware with SPI (Serial Peripheral Interface). The ADC/DAC firmware employs a Xilinx MicroBlaze soft microprocessor core to run a C program which accepts and executes commands from the ROACH2. The ROACH2 boards were developed by CASPER for general astronomical signal processing and the RF/IF and ADC/DAC boards have been developed at Fermilab specifically for the MKID arrays developed by UCSB.

The hardware interface uses the updated Casper library (`casperfpga`) for the ROACH2 boards. The ROACH2 boards are on a private LAN network as moderated by the network switch. We run all readout preprocessing software on the MEC data server, however, any computer (with the required Python libraries) that is connected to the LAN can perform this task. Each ROACH2 board boots up with an operating system and file system stored in Flash memory and runs a TCP BORPH server (provided by CASPER) to decipher commands. The firmware for the ADC/DAC board is initially loaded by a USB to JTAG interface and is stored on flash memory. The next time the board is booted it will load the firmware from memory automatically. We use NTP/PTP to sync the clocks for each ROACH2 board. A Telnet command programs the ROACH2's FPGA with the firmware

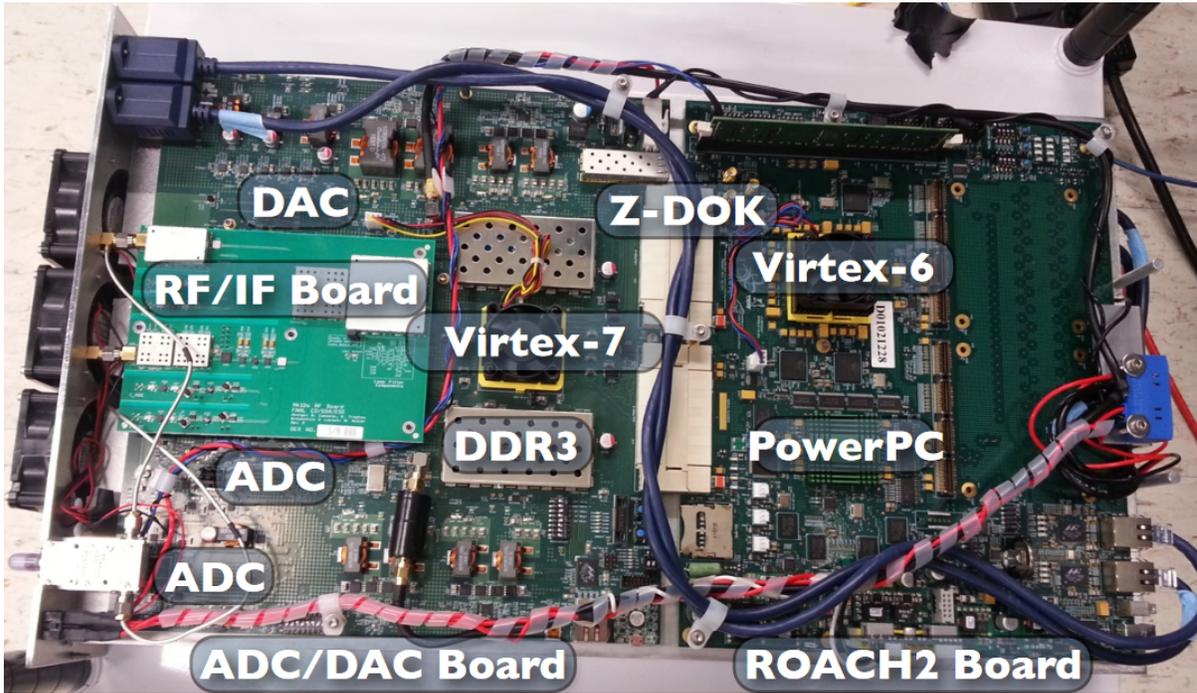


Figure 2.11: The ROACH2 board is connected to the ADC/DAC board by two Z-DOK connectors. The RF/IF board is mounted on the ADC/DAC board using SMP blind-mate connectors for signals and GPIO pins for programming. Another set of three boards are mounted to the underside of this cartridge. Figure and Caption reproduced from Strader (2016).

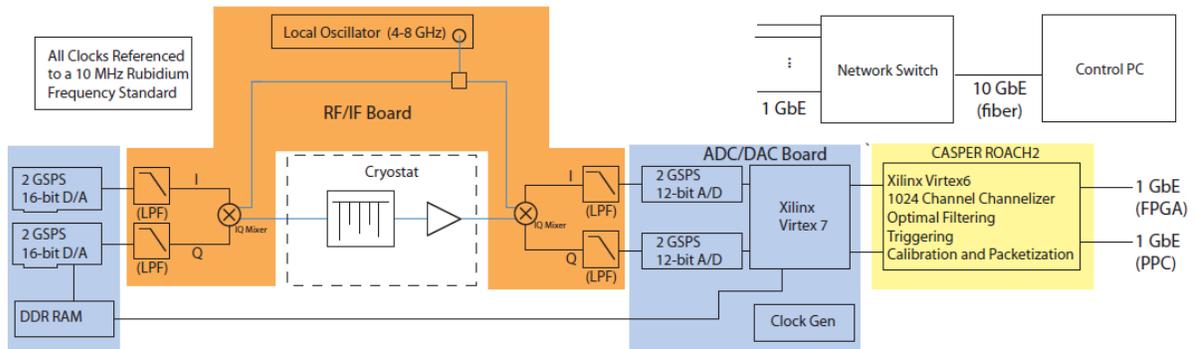


Figure 2.12: Firmware block diagram for MEC readout boards. This diagram is for 1 set of boards which can read out 1024 pixels at a time. Figure reproduced from Strader (2016).

(See Figure 2.12).

The initial setup defines the frequency comb that probes the resonators. Our python

software computes the optimum frequency comb and sends it to the ROACH2, which then gives it to the ADC/DAC board via the SPI connection, to be stored in DDR3. This is then played to the DAC in a loop while the system runs. The signal is passed through the RF/IF board where an IQ mixer up-converts the signal to the right frequency range (4 to 8 GHz) using a local oscillator. After traveling through the cryostat, MKID array, and HEMTs the RF/IF board mixes down the returning signal to baseband where it is digitized by the ADC. Then the Virtex7 on the ADC/DAC board transmits the signal to the ROACH2 over the Z-Doks. The ROACH2 firmware processes the signal, performing the following actions:

1. Separate the comb of frequencies into individual pixel frequencies using an FFT (Fast Fourier Transform) and Direct digital conversion
2. Sample the phase of each resonator's signal every microsecond
3. Run a pixel-wise unique optimal filter over each pixel's phase signal
4. Trigger on photon events and store photon packets in buffer
5. Send buffer to disk every 0.5 millisecond (or faster if full) over 1-Gbit Ethernet via UDP

The photon packets are 64 bit words with the following breakdown: 10 bit x-coordinate, 10 bit y-coordinate, 8 bit timestamp, 18 bit wavelength, 18 bit baseline phase. A C program on the data server collects the bundles of photon packets sent from the ROACH2 boards and writes them to disk as observation files. Additionally, it histograms the photons into images for real time display.

Chapter 3



MEC Microwave Signal Path

MEC requires ten feedlines to read out an entire array. The signal paths begin with hermetic SMA bulkhead connectors which bring the signals in through a box near the top of the cryostat as shown in Figure 3.1. The signals then go through hand-formable SMA to G3PO coax, which feeds a custom 10 conductor copper flex stripline cable that brings the signals down to 4 K with an intermediate heat sink at 60 K. The input side of the cables has integrated 30 dB attenuators which connect directly at 4 K through G3PO barrels to a custom 10 conductor superconducting Nb-47Ti flex cable. We have fabricated custom microstrip Nb-47Ti/Kapton/Nb-47Ti flex cables to allow for a high density of feedlines while minimizing heat load from 4 K to 100 mK as compared to ten individual NbTi coax. These flex cables then connect to the MKID box through ten small G3PO barrels allowing for a much more compact detector package than standard SMA connectors. The box-mounted G3PO connectors are solder connected to gold-plated copper on duroid CPW transition boards. The input side of the transition boards has a 20 dB Pi attenuator. The transition board traces are then connected to the MKID chip with Al wire bonds (Figure 3.2). After passing through the MKID array the ten signals are brought out through the same series of CPW board, G3PO connectors, custom Nb-47Ti microstrip flexcable, Cu flex cables, and G3PO-to-SMA coax. At 4 K each feed

line is amplified by a Low Noise Factory High Electron Mobility Transistor (HEMT) amplifier.

The flex cables were custom designed in Sonnet and laid out in L-Edit. They are described in more detail in the following sections.

3.1 300 K To 4 K Copper Flex Cables

Between 300 K and 4 K we use the Rogers Corp. LCP dielectric film Ultralam with 1/4 oz Cu cladding to make a 10 trace stripline flexible cable. The ends have G3PO push-on coax connectors. The input side has an integrated 30 dB Pi attenuator at the 4 K end made with 1 mm NiCr chip resistors. The input cable is 38 cm and the output cable is 25 cm. We estimate the 4 K heat loading per cable to be about 2.7 W-cm. The transmission and crosstalk are shown in Figure 3.3.

3.2 Superconducting Microwave Wiring

The work in this section was published in Walter et al. (2018).

3.2.1 Introduction to Cryogenic Microwave Wiring

Microwave interconnects that span from room temperature to sub-Kelvin temperatures are required for large arrays of superconducting devices. Low transmission loss before amplification and digitization is required for highly sensitive detectors. As device arrays increase in size, thermal conduction along the electronic leads dominates the heat load on the cold stage of the cryogenic refrigerator. It is common to use superconducting coaxial cables below 4 K but these quickly become unwieldy in large numbers within a cramped cryostat and their large cross sections result in unacceptable heat loads for

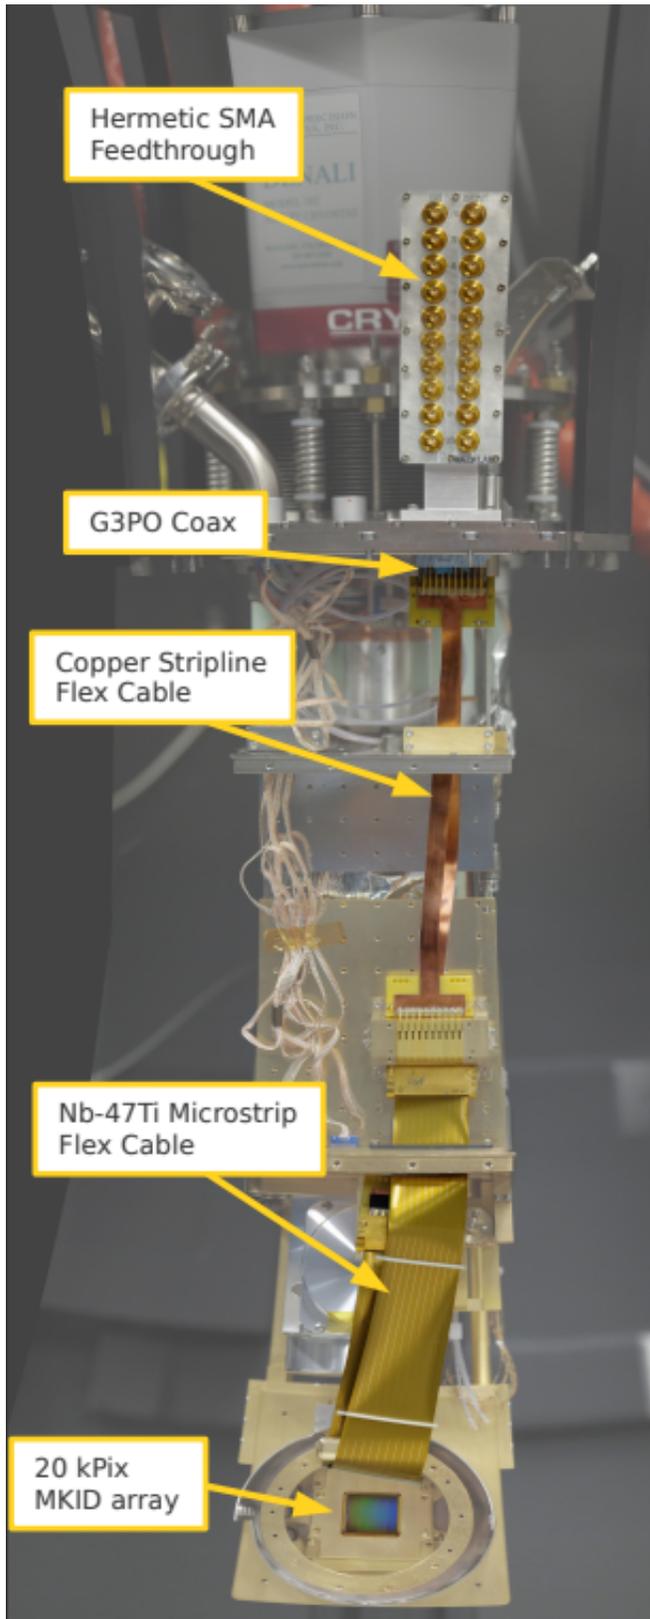


Figure 3.1: MEC cryostat internal microwave wiring and signal path. The input copper flex cable has a 30 dB attenuator and the input device box has 20 dB attenuation. After the signal passes through the MKID array there is 35-40 dB of gain from the HEMT amplifiers.

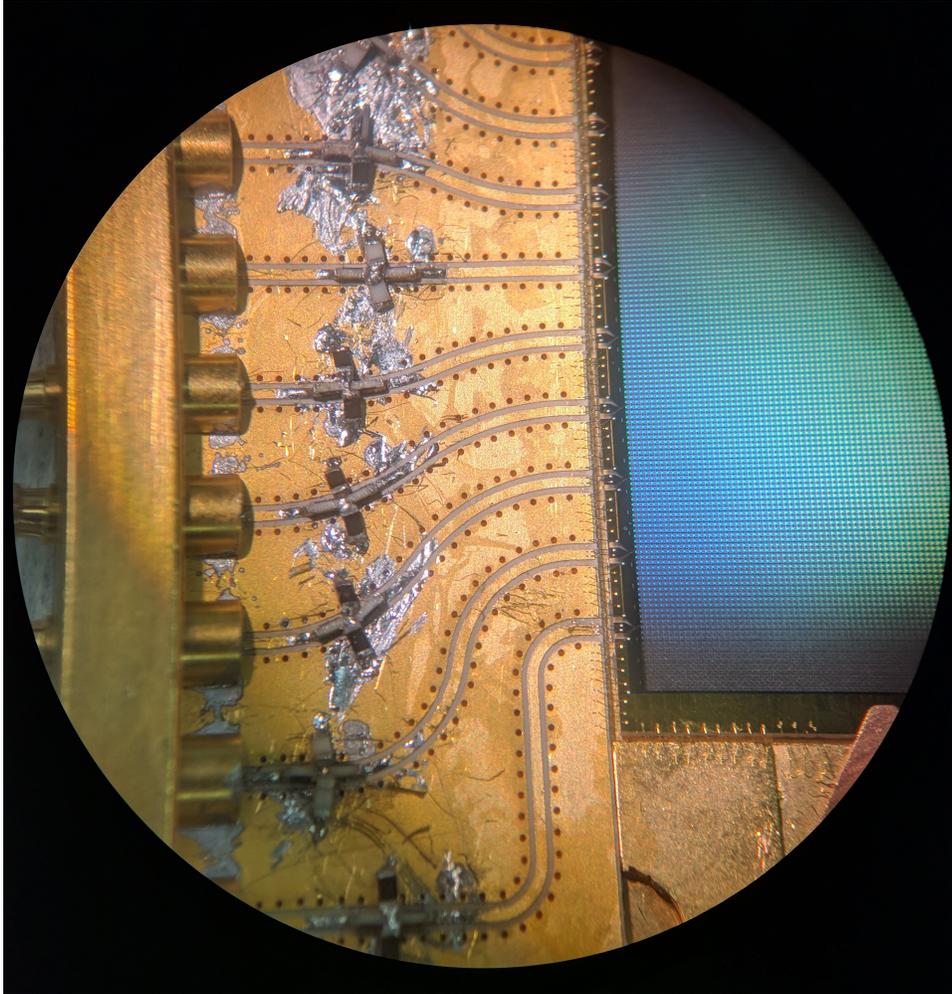


Figure 3.2: The MEC MKID array is wirebonded into its device holder. The probe signal plugs into G3PO connectors on the left. The CPW transition board has 20 dB Pi attenuators on the input side to reduce thermal noise.

adiabatic demagnetization refrigerators (ADRs). Planar transmission line circuits on a flexible substrate naturally organize multiple lines and their shared ground plane reduces the thermal cross section. There are several examples of cryogenic ribbon cables for DC voltages at millikelvin temperatures (Woodcraft et al. 2010; Yung and Moeckly 2011) and for frequencies spanning 1–20 GHz down to 4 K (Harris et al. 2012; McGarey et al. 2014). A large effort has been made in depositing polyimide and metal on a substrate wafer before performing a liftoff to create flexible circuits with Al, Nb, NbTi, and YBCO

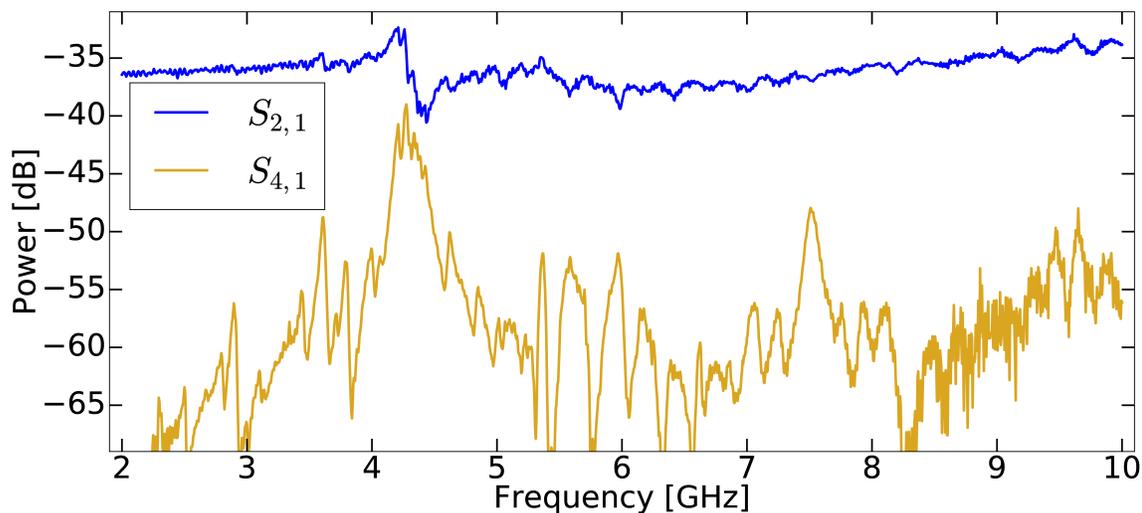


Figure 3.3: The transmission and crosstalk of the input copper flex cable including the integrated 30 db Pi attenuator.

superconductors (Paik and Mogro-Campero 1994; Allen, Franz, and Moseley 2006; Weers et al. 2013; Pappas et al. 2016; Tuckerman et al. 2016). Scalable, high density, low loss, broadband microwave interconnects for millikelvin temperatures are motivated by large arrays of sensitive cryogenic devices currently under development such as quantum computing qubits (Barends et al. 2016), X-Ray microcalorimeters (Ulbricht et al. 2015), sub-millimeter bolometers (Holland et al. 2013), nanowire single-photon detectors (Verma et al. 2014; Bellei et al. 2016), and broadband Microwave Kinetic Inductance Detectors (MKIDs) (Calvo et al. 2016; Schlaerth et al. 2012; Mazin et al. 2013; Meeker et al. 2018).

We developed superconducting NbTi microstrip flex cables specifically as broadband signal carriers for multiplexed 10,000+ sensor MKID arrays for astronomy operating at 100 mK (Meeker et al. 2018; Cook et al. 2015). These cables are unique in that they target low loss up to 10 GHz, working temperatures between 0.1–4 K, reduced complexity with push-on connectors, and lengths up to 40 cm limited by the standard commercial panel size. Our process laminates a 53 wt% Nb - 47 wt% Ti (Nb-47Ti) alloy foil to an

adhesive/polyimide panel stackup, etches a microstrip pattern into the top metal layer, and laser cuts the flex cables from the panel. This design and fabrication is detailed in Section 3.2.2 along with the G3PO coaxial interface. We present transmission loss and crosstalk measurements at 30 mK in Section 3.2.3 and compare with an analytical model and simulations made with Sonnet Suites planar EM software. While we have not directly measured the thermal conductivity, in Section 3.2.4 we use published data to estimate the thermal load on a 0.1 K stage from 0.8 K for microstrip flex cables in an ADR and compare to commercially available superconducting coaxial cables.

3.2.2 Superconducting Flex Cable Design

The flex cables were fabricated by Tech-Etch¹ using a proprietary thermal pressure lamination and wet etch process. They consist of Nb-47Ti foil sandwiching a dielectric stack up of Nikaflex CISV2525 and Hanwha HGB E250LP bisphenol epoxy. Nikaflex is a flexible DuPont/Nikkan coverlay film composed of 0.025 mm HFB-E250YG adhesive and 0.025 mm Kapton 100V. Nb-47Ti is superconducting below 10 K, is non-magnetic, and has relatively low thermal conductivity below 4 K compared to metals like Al, Nb, CuNi and SS alloys (Daal et al. 2019). The first iteration uses 0.048 mm Nb-47Ti purchased and rolled by ATI² and 0.051 mm of the Hanwha HGB epoxy for a total stackup of 0.198 mm and is pictured in Figure 3.4. In our second manufacturing trial Rikazai³ rerolled the same Nb-47Ti to 0.010 mm thickness and we use only 0.025 mm thick Hanwha HGB epoxy for a total stackup of 0.096 mm. After lamination, the top metal layer is etched into parallel traces with a width of 0.2 mm each for a 50 Ohm impedance. Several designs were fabricated having up to 40 cm in length and up to 10 traces per cable. The flex cable tested for this paper is from our first iteration of cables having a thickness of 0.198 mm,

¹Tech-Etch, Inc., 45 Aldrin Road, Plymouth, MA 02360 USA

²ATI Specialty Alloys & Components, 1600 Old Salem Rd., Albany, OR 97321 USA

³Rikazai Co., LTD., 1810 Shimonumabe Nakahara-ku, Kawasaki City, Kanagawa Pref., Japan

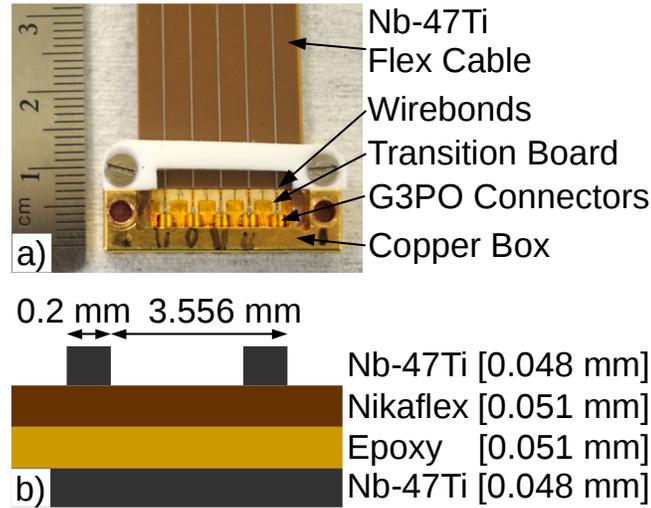


Figure 3.4: a) Image of assembled Nb-47Ti flex cable. b) Flex cable fabrication stackup with various thicknesses and dimensions indicated. Not to scale in horizontal direction.

22 cm length, and 5 traces. We use a 3.556 mm trace pitch in order to match the standard package density of Corning Gilbert's G3PO connector (compatible with SMP-S).

Hand flexional stress tests indicate the cables will flex elastically for bends as small as 6.4 mm in radius and are robust to multiple cryogenic cooling cycles. Sharper bends, when repeated, can cause fractures in the Nb-47Ti traces. Thinner materials allow a smaller bending radius.

Flex cable traces are connected to traces on a copper transition board with two 0.025 mm Al wire bonds. The transition board is a 0.254 mm thick RT/Duroid 6010LM PCB with 50 Ohm traces that start as microstrips and transition to a grounded coplanar waveguide (GCPW) geometry for increased signal isolation. G3PO surface mount connectors are soldered onto the GCPW side of the transition board and the whole end assembly is supported by a copper box. The flex cable is glued to the copper box with EPO-TEK EE129-A silver epoxy which is still conductive at cryogenic temperatures. The transition board and copper box with G3PO connectors allows for easy push-on coax interconnections between electronics and supports the end of the flex cables for rigidity. Directly

soldered connections between the flex cable and G3PO connector would be preferable, but tin/lead solder does not stick to Nb-47Ti. There has been some success with copper or aluminum platings but these techniques are still under development (Daal 2015).

3.2.3 Transmission and Cross Talk Simulations and Measurements

Transmission loss and cross talk measurements were performed in a dilution refrigerator under vacuum at 30 mK with a Keysight N9917A network analyzer. Modified Radiall R573423600⁴ switches at the mixing chamber allowed for in-situ calibration to remove effects of coaxial feedlines with a 127 mm coax as a through reference (See Figure 3.5). The components of the Device Under Test (DUT) consist of two 203 mm Koaxis CC047c non-magnetic SMA-to-G3PO adapter coaxes and the assembled flex cable. Figure 3.6 plots typical S-parameter measurements showing the S_{21} insertion loss of 2.5 dB and S_{41} cross talk of -25 dB at 8 GHz. This is an upper bound of 11.4 dB/m loss for the 22 cm cable tested.

Large transmission dips appear at multiples of 420 MHz. We hypothesize power trapped in standing waves caused by the inductive impedance mismatch of the wirebonds is dissipated by dielectric losses. If this is true, the effective relative dielectric constant $\epsilon_{eff} = (c/(\lambda \cdot 420 \text{ MHz}))^2 = 2.6$ where the standing wave's fundamental mode wavelength, $\lambda = 2 \cdot l$, is twice the flex cable length, $l = 22 \text{ cm}$, and c is the speed of light. This is consistent with empirical approximations (Wadell 1991) from the microstrip geometry $\epsilon_{eff} \approx \frac{1}{2} \cdot (\epsilon_r + 1) + \frac{1}{2} \cdot (\epsilon_r - 1) \cdot (1 + 12 \cdot \frac{t}{w})^{-0.5} = 2.6$ where t is the dielectric's thickness, w is the trace width, and ϵ_r is the relative dielectric constant. At room temperature, Nikaflex has $\epsilon_r = 3.5$ and the adhesive used in the dielectric stackup has $\epsilon_r = 3.25$.

⁴Microwave switches were custom modified to work at cryogenic temperatures

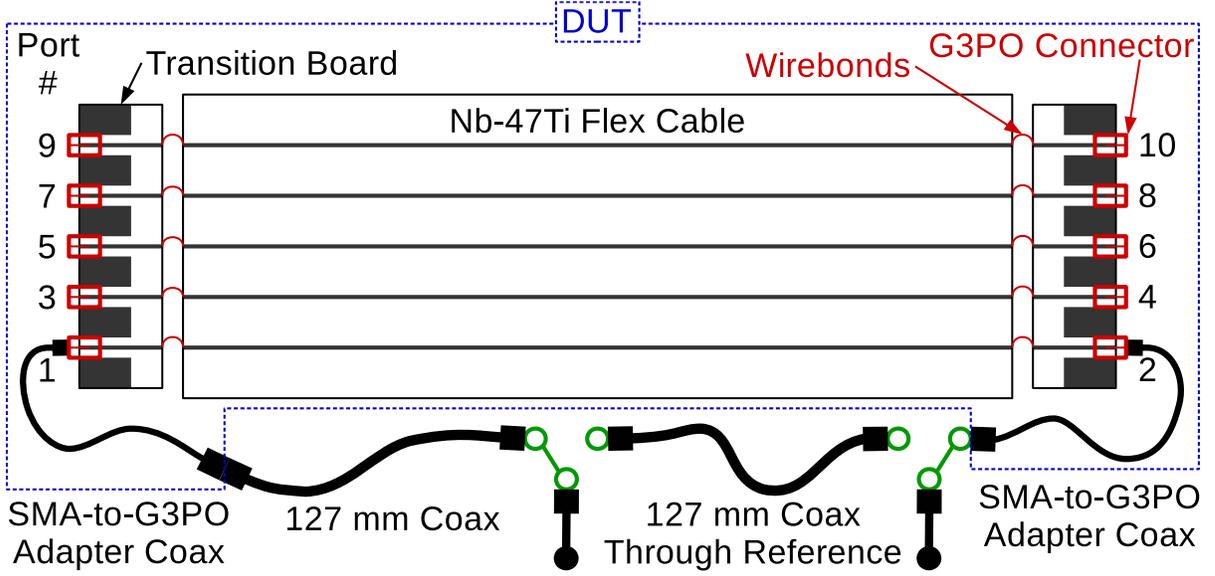


Figure 3.5: Diagram of the S_{21} microwave measurement setup at the mixing chamber of the dilution refrigerator. The Device Under Test (DUT) consists of 203 mm Koaxis SMA-to-G3PO coaxes and the flex cable assembly including the microstrip flex circuit, wirebonds, transition boards, and G3PO edge mount connectors. Disconnected traces on the flex cable were terminated with 50Ω loads. The SMA switches allow for an in-situ through calibration without interrupting the dilution cooling cycle. S_{41} is measured during a separate cool down.

We can further substantiate this hypothesis with a 2-port analytical model where the wirebonds are represented as lumped element inductors and the superconducting flex cable is modeled by a perfectly conducting transmission line with dielectric losses (See Figure 3.7). By writing the ABCD matrix for each element in the model we can derive a cascade ABCD matrix for the entire network (Pozar 2011). The complex scattering parameter corresponding to transmission as a function of angular frequency ω is:

$$\begin{aligned}
 S_{21}(\omega) = & [(Z_f/(2Z_0) + Z_0/(2Z_f) \\
 & + Z_w^2/(2Z_0Z_f) + Z_w/Z_f) \sinh \gamma l \\
 & + (Z_w/Z_0 + 1) \cosh \gamma l]^{-1}
 \end{aligned} \tag{3.1}$$

with $Z_0 = 50 \Omega$ ports, the purely imaginary wirebond impedance $Z_w = i\omega L$, and

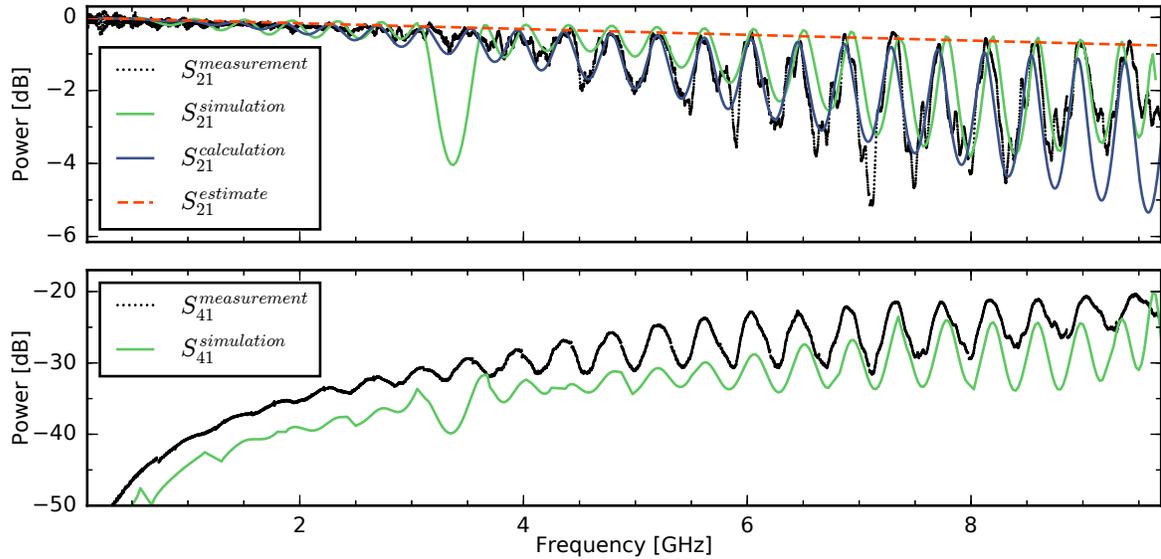


Figure 3.6: Typical S_{21} (top) and S_{41} (bottom) measurements with frequency are shown of the DUT which includes the assembled flex cable and short coax cables. Blue shows the analytical model calculation of the flex cable as a transmission line with dielectric losses terminated by wirebonds as lumped element inductors. In green, we use the wirebond inductance and dielectric loss tangent estimated from the analytical calculation and simulate the response with a planar EM model in Sonnet. In red, we estimate the insertion loss without wirebonds. The calculation and simulation agree with S_{21} measurements while the simulation of S_{41} is systematically 2–4 dB lower than the measured values.

the complex propagation constant $\gamma = \omega\sqrt{\epsilon_{eff}}(\tan\delta + 2i)/(2c)$. From time domain reflectometry measurements we determine the flex cable’s impedance Z_f is very close to $50\ \Omega$. We used a least-squares fit to match the DUT’s transmission data with two free parameters, wirebond inductance L , and the flex cable’s average dielectric loss tangent $\tan\delta$. The best fit plotted in Figure 3.6 uses $L = 1.0\ \text{nH}$ and $\tan\delta = 0.0015$. Removing the inductive impedance mismatch in the model gives an insertion loss of 0.6 dB at 8 GHz.

We are unable to independently measure the resistive losses from the copper adapter coaxes which may be significant and calculations from first principles are insufficiently accurate due to the uncertainty in the adapter coax’s copper purity. Thus, the flex cable’s



Figure 3.7: Schematic of a single trace of the flex cable assembly as a two port equivalent circuit. The flex cable is modeled as a perfectly conducting transmission line with dielectric losses and the wirebonds are modeled as lumped element inductors with inductance L . We ignore additional components of the DUT including the copper transition board, G3PO connectors, and SMA-to-G3PO adapter coaxes.

dielectric loss tangent may actually be much smaller than 0.0015 since the two port analytical toy model ignores effects from the copper adapter coax, G3PO connector, copper transition board, and cross coupling to neighboring traces. This upper bound appears small considering the room temperature loss tangents of Kapton $\tan \delta = 0.003$ and the bondplys $\tan \delta = 0.03$. Tuckerman et al. (2016) found in a similar measurement the room temperature dissipation factor for polyimides PI-2611 and HD-4100 was suppressed by a factor of a hundred at 20 mK.

We use Sonnet EM software tools to model the transmission and estimate the crosstalk from the more complicated multi-port circuit. The simulation applied the $\tan \delta = 0.0015$ and lumped element 1.0 nH inductors to match the results of the analytical calculation. While the transmission magnitude and oscillations agreed, we found the nearest neighbor crosstalk measured was 2–4 dB worse than the estimate from the Sonnet simulation. This may be because the transition board for the G3PO coaxial interface was ignored in the simulation. Sonnet is a planar simulator so accurately modeling the transition circuit becomes computationally expensive due to its inherent 3D nature.

3.2.4 Heat Flow Calculations

The thermal conductance of the flex cable is calculated from literature values of the thermal conductivity of Nb-47Ti (Daal et al. 2019; Olson 1993) and Nikaflex (Kellaris

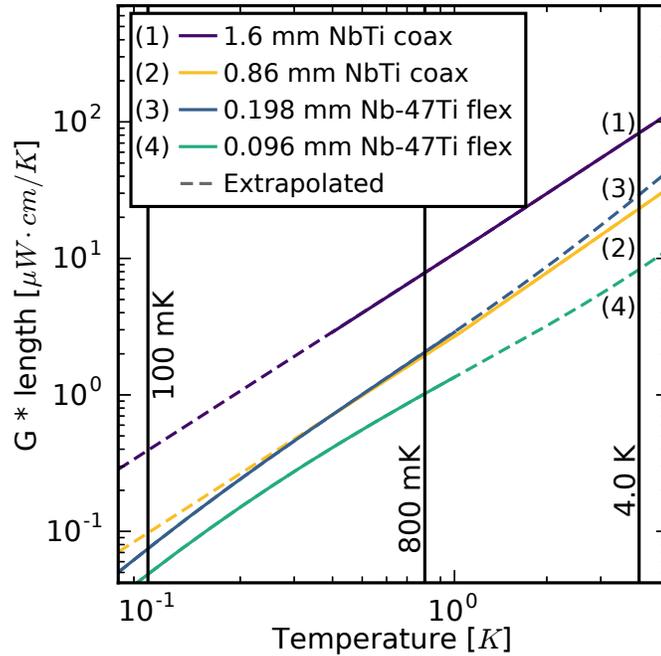


Figure 3.8: We integrate the thermal conductivity over the cross sectional area as a proxy for thermal conductance, G , independent of cable length. This work calculates values per trace of the Nb-47Ti flex cables from literature values of constituent materials (Daal et al. 2019; Olson 1993; Kellaris et al. 2014). This can be compared to the directly measured conductance of the smallest commercially available superconducting Nb-47Ti coax cables from Coax Co., LTD. We show here their \varnothing 1.6 mm Nb-47Ti Coax (SC-160/50-NbTi-NbTi)(Kushino et al. 2008) and their \varnothing 0.86 mm diameter version (SC-086/50-NbTi-NbTi)(Kushino and Kasai 2016). The dashed portion of the lines indicate data extrapolation to the temperature range of interest.

et al. 2014) and their respective cross sections. We assume that the Hanwha HGB adhesive has the same thermal conductivity as Nikaflex since it is similar to the epoxy that dominates the thermal conductivity in Nikaflex. In Figure 3.8 we show the calculated thermal properties per trace of our flex cable compared to the measured commercially fabricated superconducting cables from Coax Co., LTD.

The heat load can be calculated from values in Figure 3.8 by integrating over temperature and dividing by cable length. For our 5 trace, 22 cm long, 1.8 cm wide, 0.198 mm-thick flex cable we estimate a thermal load of 60 nW on the 0.1 K ADR stage if the cable

Table 3.1: Summary of Thermal, Mechanical, and Microwave Properties of the Nb-47Ti Flexible Microstrip Cables Presented in This Section Compared to the Best Commercially Available Nb-47Ti Coax Cables from Coax Co., LTD. The Size Indicates the Total Thickness of the Flex Cable or Outer Diameter of the Coax Cable. The Thermal Load with Cable Length per Trace is Calculated for Cables Spanning 4 K to 0.8 K and Cables From 0.8 K to 0.1 K. Measurements of the Insertion Loss and Crosstalk are Described in Section 3.2.3. The Estimated Insertion Loss for the 0.198 mm Thick Nb-47Ti Flex Cable Gives our Model Estimate When the Wirebond Impedance Mismatch in the Assembled Flex Cable is Removed. We Have Not Measured the Microwave Properties of the Thinner Nb-47Ti Flex Cables.

	Nb-47Ti Flex		Nb-47Ti Coax ^a	
	0.198	0.096	∅1.6	∅0.86
Size [mm]	0.198	0.096	∅1.6	∅0.86
Trace Pitch [mm]	3.556	3.556	>13	>13
Min Inside Bend Radius [mm]	6.4	3.2	6.4	3.2
Thermal Load per Trace [μW cm]	8.4	2.8	26.	7.1
	0.8 K Stage			
	0.1 K Stage	0.13	0.07	0.51
Insertion Loss at 8 GHz [dB/m]	<11.4	–	<0.4	<0.5
	Measured			
	Estimated	<2.7	NA	NA
Crosstalk at 8 GHz [db]	–25.	–	NA	NA

^aData from Coax Co catalog: http://www.coax.co.jp/en/wcaxp/wp-content/themes/coax/pdf/cryogenic_cable_catalogue.pdf.

spans 11 cm from the 0.8 K stage. This is roughly equivalent to five of the 0.86 mm Nb-47Ti coaxes from Coax Co. The 0.096 mm thick flex cable is half this. Table 3.1 compares the thermal load from the flex cables and Coax Co.’s superconducting coax cables.

3.2.5 Conclusions on the Microwave Wiring

We fabricated scalable superconducting flexible microwave push-on cables for 10,000+ sensor MKID arrays being developed at UCSB. They could be useful in a wide variety of sub-kelvin applications with highly sensitive devices that require a large number of and tightly packed microwave feedthroughs capable of large and overlapping bandwidth. While for these prototypes we use a panel manufacturing process, the fabrication tech-

nique of laminating Nb-47Ti foil to a polyimide substrate and wet etching is amenable to a roll-to-roll manufacturing process for very long flexible cables. The thermal conductance below 4 K for the 0.198 mm thick Nb-47Ti flex cable is roughly equivalent per trace to the smallest (\varnothing 0.86 mm) Nb-47Ti coax cables from Coax Co.. The second iteration of the flex cable is only 0.096 mm thick and will have half the thermal load. In either event, an ADR's thermal budget can be accommodated by increasing the flex cable's length.

We found an average insertion loss of <2.7 dB/m at 8 GHz for the 22 cm long, 0.198 mm-thick flex cable assuming the impedance mismatch caused by the wirebonds can be removed. This loss is more than five times the measured insertion loss for commercially available superconducting Nb-47Ti coax cables from Coax Co. which cannot be explained by the different cable geometry. The loss is likely do to an unknown combination of factors including the uncorrected loss from the copper adapter coaxes in the DUT, and a difference in the dielectric loss properties of Kapton and Teflon below 1K which are unexplored in literature.

Several avenues exist to avoid the impedance mismatch including a better transition board design or plating of the Nb-47Ti with a solderable metal like copper or aluminum. The plating technique is especially desirable because it would remove the need for a transition board and enable the fabrication of via structures and therefore more complicated multilayer designs such as stripline. The greater isolation enjoyed by the stripline structure over microstrip could be used to improve the cross talk between traces, although the measured -25 dB nearest neighbor cross talk is sufficiently low for our current MKID applications. Additionally, the 0.096 mm thick flex cable fabricated will have improved cross talk performance due to the thinner traces and their closer proximity to the ground plane.

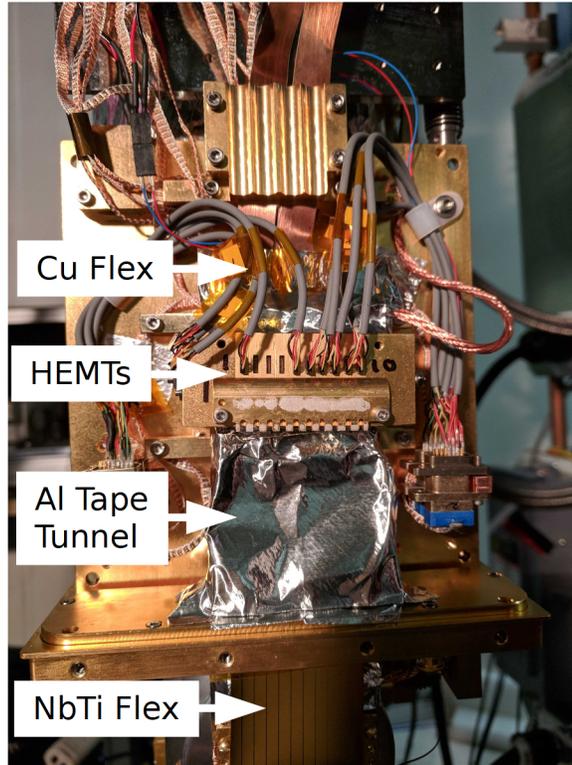


Figure 3.9: The output Nb-47Ti Flex cables connect to the HEMTs under the 60 K shield. By surrounding the exposed microstrip cables in aluminum tape the noise is reduced.

3.2.6 Comments on Performance in MEC

We found that the microstrip cables in MEC caused noise on the MKID phase response equivalent to about 50 K thermal noise, much larger than the expected 1.8 K HEMT amplifier noise. The HEMTs are bolted to a 4 K copper plate but that section of the cryostat is surrounded by the 60 K (nominally 57.5 K) heat shield. We speculate that the exposed microstrip flex cable was picking up 60 K black body radiation which was amplified as thermal noise by the HEMT. By surrounding the flex cable with an aluminum tape tunnel as shown in Figure 3.9 this noise is almost completely mitigated.

When the MEC readout is run with all feedlines at full power we found that one or more HEMTs were oscillating. We speculate that this is caused by bad isolation in the

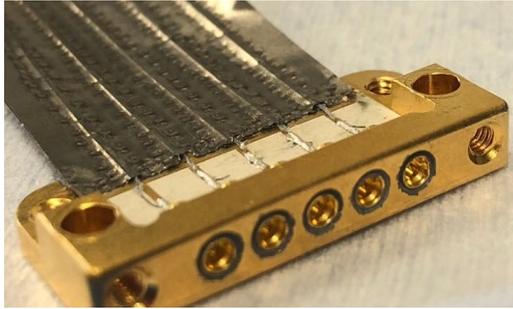


Figure 3.10: A new design for a superconducting microwave cable. It crimp welds NbTi foil around commercially available NbTi wires for improved isolation and mechanical flexibility compared to the old flex cable design. This is preliminary work.

Nb-47Ti flex cable and/or Cu flex cable that causes a feedback loop in the amplifier. We were able to mitigate this by slightly adjusting the gains in the HEMT amplifiers and completely removing power from the HEMTs in feedlines 2, 3, and 4 (bad feedlines) and 5, 6 (good feedlines). However, the overall noise performance at Subaru is not as good as when a single feedline is powered in the lab at UCSB. Future flex upgrades will fix this issue.

Finally, the Nb-47Ti flex cables used in MEC turned out more delicate than expected. The current flex cable traces were made with Nb-47Ti foil that was tempered at 2 mil thickness then rolled to 0.010 mm thick. We suspect that this process greatly reduced the metal's ductility because it was not re-annealed afterwards.

3.2.7 Preliminary Work for MEC Upgrade

Preliminary work on a new superconducting flex cable has begun to solve these problems. The new design stretches commercially available 3 mil diameter NbTi wire (with dielectric coating) and crimp welds the NbTi foil around it as a shield. This is shown in Figure 3.10. The new cable has improved isolation both in terms of crosstalk and pickup and by using commercially available NbTi wires for the signal trace the ductility is improved. Testing is on-going.

Chapter 4



Instrument Control, Setup, and Readout Software

4.1 Temperature Control of Cryostat

We developed a Labview GUI (see Figure 4.1) for the automated control of the cryostat temperature. The cryostat needs to regulate the temperature of the MKID array at 0.090 ± 0.001 K for stable operation of the instrument. We have confirmed operation of this software at the summit with the following specifications:

- 17+ continuous hours hold time at 90 mK
- 3 hour recycle time at 4 K (MKIDs cannot be read out during this time)
- Continuous logging of temperature and electronics rack status

Conveniently, the software allows the automated cycling of the ADR to begin at a specified time. This software runs on the Intel NUC PC located in the MEC electronics rack on the IR Nasmyth platform.

The He compressor is controlled remotely and its status logged through a Labview executable provided by Cryomech.

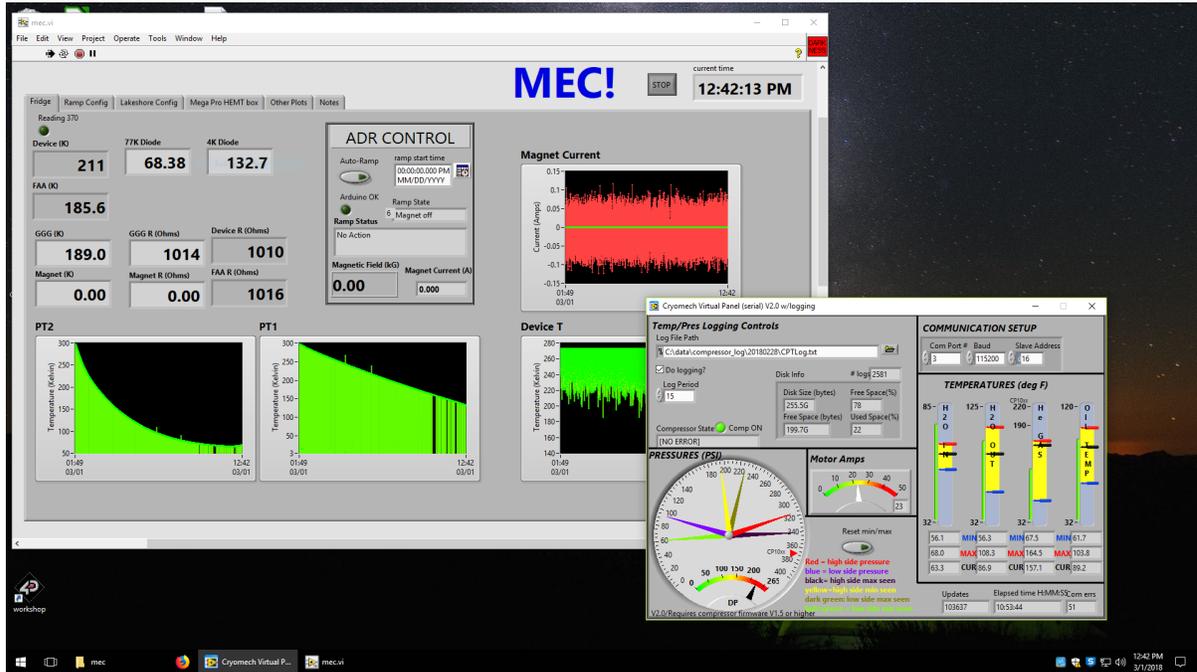


Figure 4.1: Screenshot of MEC temperature control GUI and remote compressor control GUI.

4.2 Pixel Calibration and Pre-Setup

Before operating a new MKID array every pixel must be individually calibrated. This is done by probing the MKID array using the readout electronics (see Section 2.6) and processing the optimal operating conditions for each pixel with automated scripts¹. The steps are described in the following subsections.

4.2.1 Resonator Frequency Identification and Power Optimization

While each MKID pixel has a specific design frequency, its true resonant frequency is scattered due to inhomogeneities in the superconducting film and nanoprocessing. Furthermore, when temperature cycling to room temperature and back there are often

¹Software publicly available on github: <https://github.com/MazinLab/MKIDReadout>

semi-global shifts in the resonator frequencies possibly due to a shift in the magnetic flux traps around the feedline or resonator circuit. This means that the MKID resonant frequency is not known a priori and must be measured by examining the transmission of a frequency sweep.

Using the digital readout we set up a random frequency comb on the DAC and sweep the local oscillator frequency in order to measure the transmission across the entire 4-8 GHz readout band. We repeat this sweep with various levels of power attenuation to attain a power sweep. Resonators are identified in the frequency sweep as loops in the IQ plane or as sharp dips in the transmission. The power sweep data is analyzed by a deep neural net using Tensor Flow in Python to determine the optimal resonator power as described in Dodkins et al. (2018).

4.2.2 Choosing the LO Frequency

Each readout board can handle resonators within 2 GHz of bandwidth as determined by the bandwidth of the DAC and enforced by an anti-aliasing filter. We choose a local oscillator frequency in the middle of the band; around 5 GHz for the low frequency half or around 7 GHz for the high frequency half. We can not read out resonators that are more than ± 1 GHz away from the LO or within 250 KHz of each other. Furthermore, we want to avoid tones in the sideband aliasing across the LO on top of a different tone in the other sideband. We use a monte carlo simulation to optimize the choice of LO to read out the largest number of pixels.

4.2.3 Calculating Pixel Optimal Filters

We use the digital readout to record the phase timestream of a resonator sampled at $1 \mu s$ while illuminating the MKID array with white light. We average photon events into

a template and measure the phase noise in order to build an optimal filter that is unique for each pixel. These optimal filters are loaded into the ROACH2 firmware to improve the signal-to-noise of the on-board photon triggering.

4.2.4 Determining Pixel Location in the Array (Beammapping)

Since a pixel's resonant frequency is different from its design frequency it is unknown where that pixel is physically located in the MKID array. We call the procedure of mapping each pixel's frequency to spatial location "beammapping." To beammap, we image the screen of a Samsung Galaxy S6 cell phone directly onto the MKID array as described in Bottom et al. (2018). With a custom Android app, a bar of light is swept across the screen in the 'x' and then 'y' direction. By determining the time at which an MKID pixel observes light, its relative location in the array can be inferred.

4.3 Instrument Setup, Control, and Observing Strategy

After the pixels have been calibrated as described in Section 4.2 the instrument is ready to be set up for an on-sky observation night. For this we follow the following procedures:

4.3.1 Readout Board Initialization

Every time the readout boards are powered they must be initialized (hardware described in Section 2.6). This includes the following steps:

1. Connect

The readout computer queries the readout boards over the private network and confirms that it has booted up.

2. Program Vertex 6 FPGA

The firmware is uploaded and programmed into the Vertex 6 FPGA on the ROACH2 board.

3. Initialize Vertex 7 FPGA

The UART communication on the ADC/DAC board is initialized for the Vertex 7 FPGA. The ADC, DAC, LO chip, and attenuators are configured.

4. Calibrate ZDOK

The ZDOK high speed connection between the Vertex 7 and Vertex 6 FPGAs is calibrated.

5. Calibrate QDR

The QDR static ram on the ROACH2 board is formatted so that the DDS LUT can be loaded later.

These steps need only be performed once, however, if there are issues with the readout or the pixels seem especially noisy then it is recommended to re-initialize or power cycle the board.

4.3.2 High Templar: Readout Setup

We set up the pixels for readout using a Python GUI called "High Templar." This must be done every time the cryostat is cycled from 4 K down to the 90 mk operating temperature (once per observing night). A screenshot of the initialization GUI and High Templar GUI running on the MEC readout computer is shown in Figure 4.2. The setup steps are listed below:

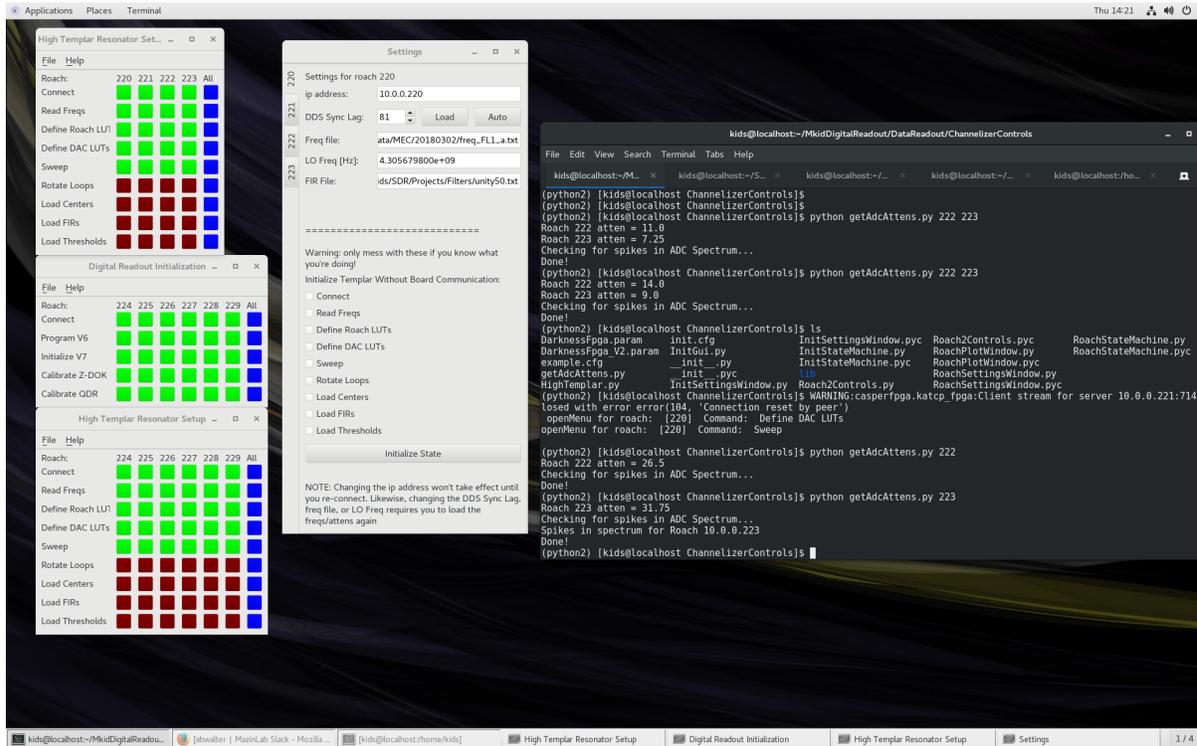


Figure 4.2: Screenshot taken March 29, 2018 of the High Templar GUI and Init GUI setting up MEC’s pixel readout at Subaru Telescope.

1. Connect

The readout computer queries the readout boards and confirms that their firmware has been initialized.

2. Load Frequencies

The optimized frequencies and powers for the pixel probe tones are read from a text file.

3. Define ROACH2 Look Up Table (LUT)

The FFT bin that each pixel’s probe tone falls in is calculated and uploaded into the firmware’s channel selection block. The Direct Digital Synthesis (DDS) tone for each pixel is calculated and loaded into the DDS LUT (for digital down conversion from RF to DC).

4. Define DAC LUT

The Local Oscillator (LO) frequency is loaded into the LO chip on the IF board. The frequency comb to probe the pixels is calculated and loaded into the DDR3 memory on the ADC/DAC board. The global DAC attenuators are set to maximize the dynamic range used on the DAC while keeping the resonator probe tones at the correct power. The global ADC attenuator is optimized so that the incoming signal on the ADC uses the maximum amount of dynamic range without overflow.

5. Frequency Sweep

The LO tone is swept in steps of 10 kHz for 500 kHz (this is customizable) to measure the pixel's resonator IQ loops.

6. Rotate IQ Loops

The phase of each pixel's probe tone is calculated and the negative of that phase is used to recalculate the DDS tones and load them into the DDS LUT. This makes each resonator's IQ loop have a nominal phase of zero after the DDS.

7. Translate IQ Loops

A frequency sweep is performed again to measure the IQ loop. The center of each resonator's IQ loop is calculated and loaded into the ROACH2 firmware so that the phase can be calculated in the photon triggering block.

8. Load Finite Impulse Response (FIR) Filter

The optimal filter coefficients for each pixel are loaded into firmware.

9. Load Thresholds

The phase response for each pixel is sampled in the dark and the noise measured. A (nominal) 4 sigma noise threshold is calculated and loaded into the

firmware triggering block. Only when the phase overcomes this threshold will the firmware trigger a photon event.

4.3.3 Dashboard: Instrument Control

The real-time monitoring and instrument control Python GUI is called "Dashboard." Dashboard automatically starts the c-program "packetmaster" which performs the following operations:

- Listen over the 10 Gb Ethernet port and collect photon packets sent by the readout boards.
- Collate the photon packets into one second chunks and write them directly to disk as raw binary files.
- Histogram the photons in images and save those images in shared memory. The effective exposure time is variable and a wavelength cut optional.

Dashboard displays the real-time image created by the packetmaster to the user. It also sends instructions to the filter wheel server and dithering mirror server that are running on the MEC PC in the instrument rack. Finally, Dashboard periodically logs the instrument state including telescope TCS information. A screenshot of Dashboard is shown in Figure 4.3.

4.3.4 Wavelength, Flat Field, and Pointing Calibration Data

Part of the observing procedure is to collect wavelength, flat, and pointing calibration data that is required for later analysis. Wavelength calibration data needs to be collected at least once per night of observing (and every time the resonators go through

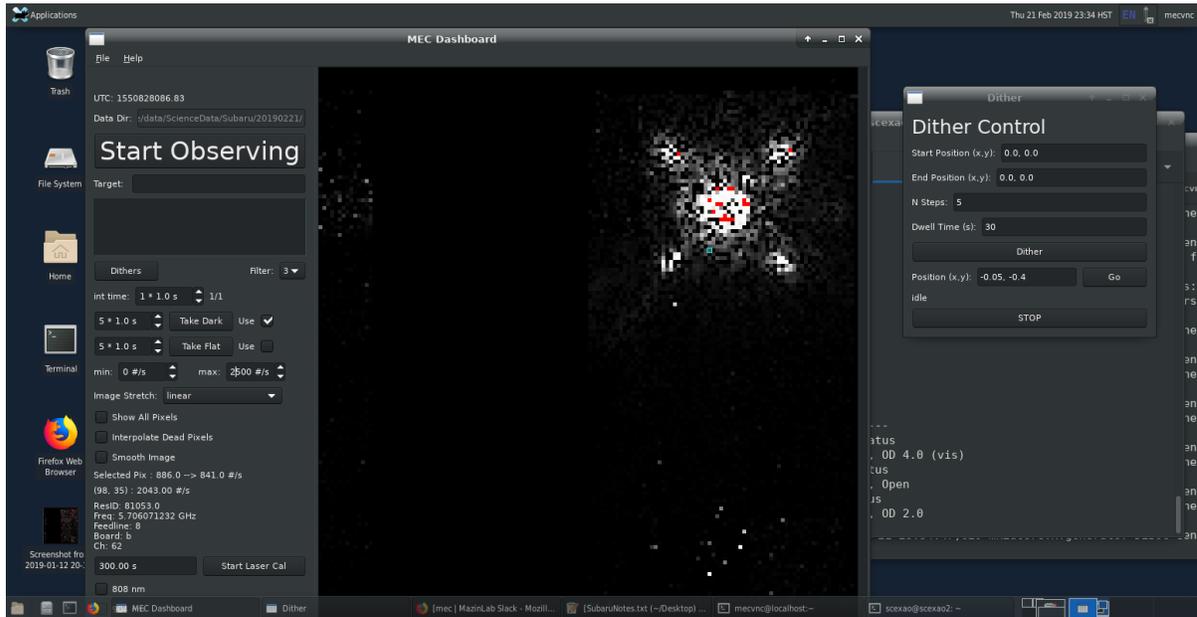


Figure 4.3: Screenshot taken February 21, 2019 of the Dashboard GUI with real-time display. The SCEXAO white light source is imaged on the MKID array along with four artificial satellite speckles.

the rotate/translate/threshold setup). Flat fields and Pointing calibration can be done once every few nights.

To collect data for the wavelength calibration we use NKT Photonics' supercontinuum source (SuperK) with a laser line tunable filter (LLTF) and an integrating sphere to image a monochromatic flat field. The wavelength can be freely chosen (with FWHM of 2-5 nm) with the tunable filter and the brightness adjusted by the laser source and the ND filters in MEC's filter wheel. Typically we aim for about 1000 photons/second/pixel for 60-100 seconds for five wavelengths between 800-1500 nm.

For the flat field calibration we use a bright white light internal to Subaru Telescope's AO188 which is imaged as a flat field on the MKID array. We typically aim for a few minutes of data with count rates around 500-700 photons/second/pixel.

The pointing calibration is a calibration of the (u, v) angle of the dithering mirror in the MEC fore-optics to the (x, y) pixel location of an on-axis point source. This is

done by imaging the internal white light source from SCEExAO. Additionally, if the plate scale was known previously, we can check for small changes in the plate scale due to realignments in the SCEExAO optics by dithering.

Chapter 5



MEC On-Sky Observations and Data Processing

To analyze MEC data we use the MKID Data Pipeline available on Github: <https://github.com/MazinLab/MKIDPipeline>.

5.1 Calibration of Raw Data

The MKID pipeline begins by transforming the raw binary data into calibrated photon lists in a fast access database. This process is outlined in Figure 5.1. Many of the steps were implemented for the ARCONS pipeline (Eyken et al. 2015) and were copied or slightly modified for MEC.

1. Processing Binary Data into a HDF5 Database

The raw binary files are interpreted and saved into a HDF5 table format. This involves parsing the photon packets into integers (timestamp, pixel location) and floats (phase height). Two weights are associated with each photon. The photon list is rearranged so that it is time ordered.

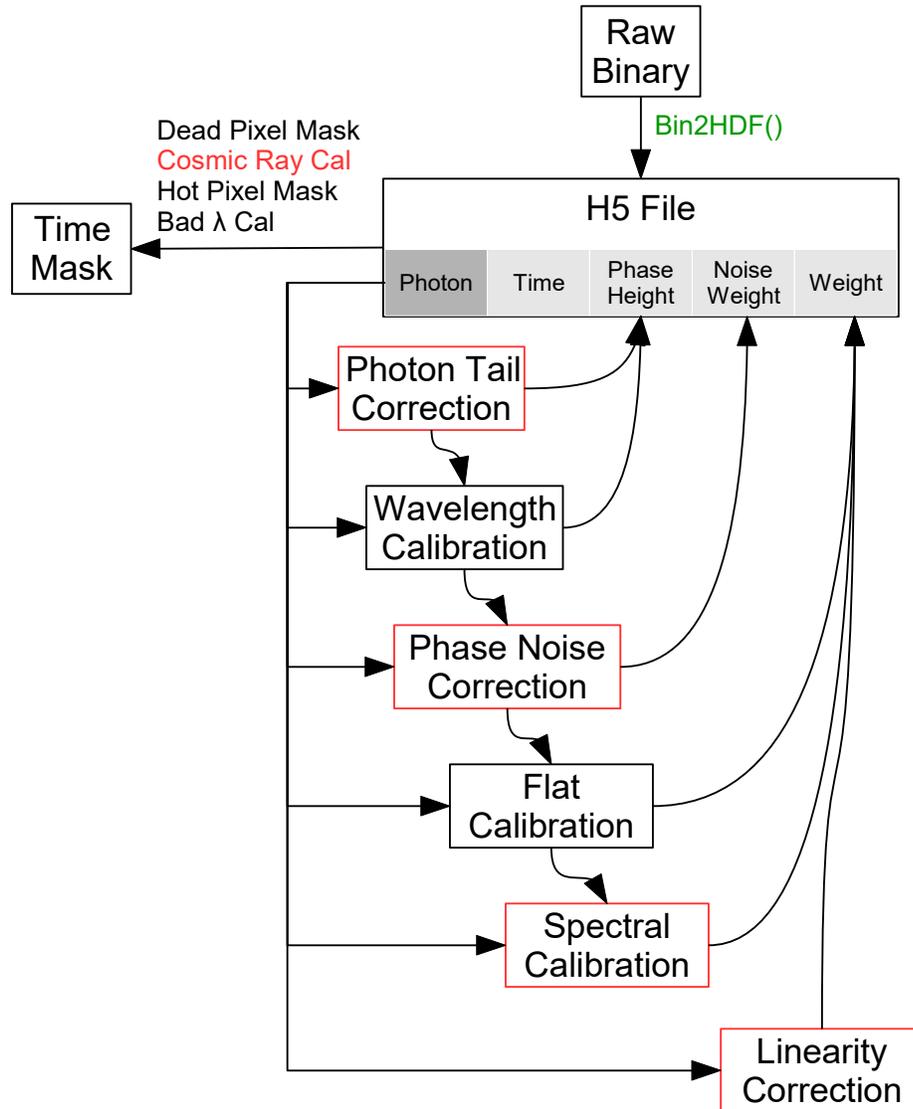


Figure 5.1: Flowchart of MKID pipeline for calibrating raw MKID data. Red indicates a procedure not yet implemented.

2. Linearity Correction

This calibration accounts for missing photons due to the detector dead time. This dead time is introduced in the firmware to avoid photon pile-up and limit the total count rate. It is applied to the weight column for each photon depending on the local count rate. It is not yet implemented because it is a small correction for low count rates.

3. Photon Tail Correction

This calibration step is not implemented because it is a small effect and it is not clear that it is needed. Essentially, if two photons arrive too close together in time, the second photon's phase peak will ride on the tail of the first photon's. This can be subtracted off since we understand the shape of the phase peak from the optimal filter template. This calibration would correct the wavelength of some photons.

4. Wavelength Calibration

The wavelength calibration is essential. Monochromatic calibration data is used to determine a quadratic, $a^2\phi + b\phi + c = hc/\lambda$, that converts phase, ϕ , into wavelength, λ . To save space, the phase height float column is overwritten with the calibrated wavelength for each photon in a pixel with a valid wavelength solution.

5. Phase Noise Correction

This calibration step is not implemented because it is difficult. For photons at the low energy end of the band (approaching 1500 nm) the phase height is small and the population may be contaminated by noise events that act as false positives. This can systematically skew the flat weight at the red end. If this calibration step is implemented correctly, it will assign a weight to each photon describing how likely it is to be a real photon as a function of that photon's wavelength. Bluer photons are very likely to be real while redder photons are less likely. This is a difficult procedure because every pixel has a different energy resolution, a different threshold for phase noise, and the phase noise may increase with higher count rate or change based on the incident spectrum.

6. Flat Field Calibration

This is a standard flat field calibration that determines the inter-pixel relative QE. This is done as a function of wavelength. The resulting flat weight is applied to the weight column for each photon.

7. Spectral Calibration

This calibration accounts for the total QE for the detector and system as a function of wavelength. It multiplies into the weight column of the HDF5 table to merge with the flat weight. It is not yet implemented in the MKID pipeline.

8. Hot Pixel Masking

Hot pixels (or cold or otherwise badly behaving pixels) are prevalent in our data because we try to include as many pixels as possible. These pixels must be removed from the final data product.

5.2 Image Processing

After the photon lists have been calibrated we can begin the analysis of scientific data. The easiest to understand procedure is to bin the photon list in time and wavelength to create images after which a classical astronomy pipeline can be used.

Figure 5.2 shows the five star system Theta¹ Orionis B taken in y-band by MEC at Subaru Telescope on Jan. 12, 2019. B1, behind the coronagraph, is a spectroscopic eclipsing binary with nominal magnitude of $V=7.96$. B2 and B3 are resolved by Subaru Telescope's 8.2 m mirror in conjunction with the AO correction provided by SCExAO. The four astro speckles artificially created by the DM are about 5.5 magnitudes fainter than B1. B4 is 4.98 magnitudes fainter than B1 in H-band (Close et al. 2013).

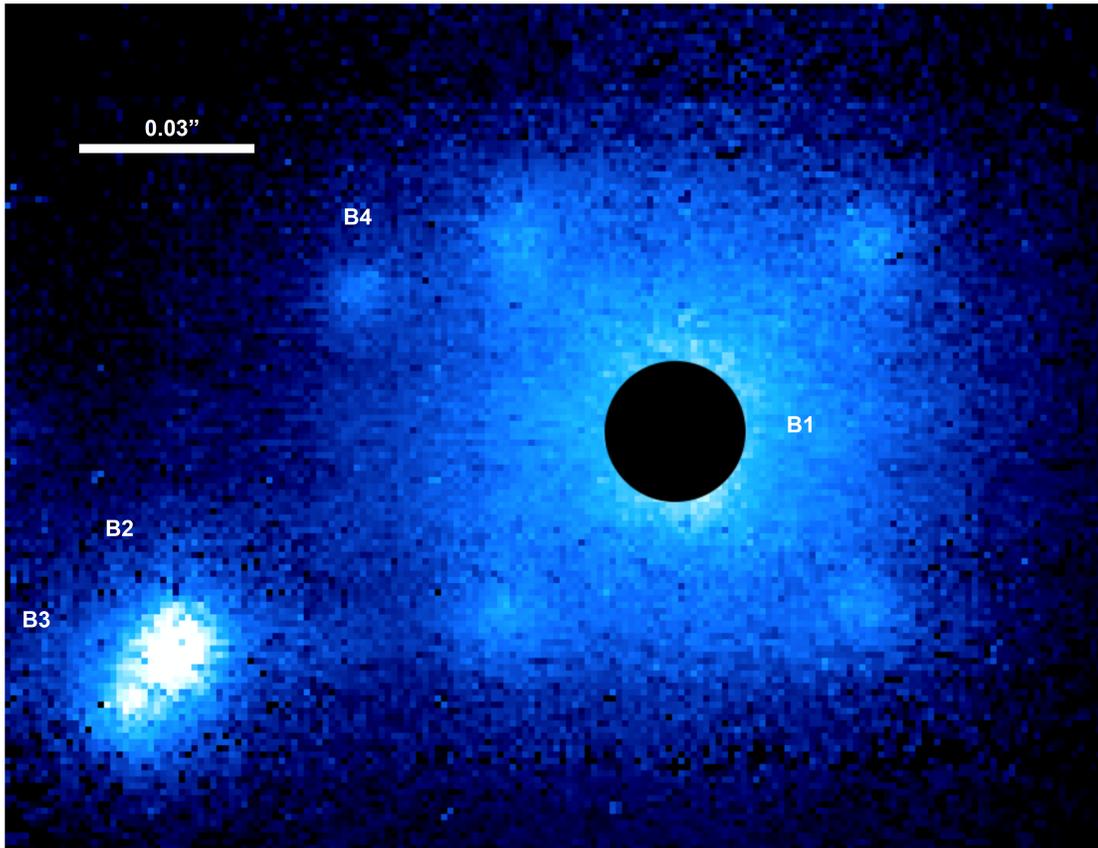


Figure 5.2: y -band image of Theta¹ Orionis B taken by MEC at the Subaru Telescope on Jan. 12, 2019. 25 30 s dithered images are stacked then smoothed by a Gaussian filter. The black circle indicates the location of the coronagraph on B1. B4 is 4.98 magnitudes fainter in H band than B1. The four artificial speckles are for astrometric and photometric calibration.

The image uses data from 25 30 s dithered positions on-sky. We use a drizzling code to place the list of observed photons in a virtual pixel grid that correctly takes into account the translation from the dither position as well as sky rotation as observed from the Nasmyth port.

Chapter 6



Stochastic Speckle Discrimination

The work presented in this chapter is published in Walter et al. (2019).

6.1 Introduction to the SSD Algorithm

Stochastic Speckle Discrimination (SSD) is a post-processing technique for direct planet detection designed to reduce the additional noise caused by speckle fluctuations by temporally resolving the intensity variations in the focal plane. It relies on the difference in shape between the on-axis and off-axis intensity distributions (Soummer and Ferrari 2007; Gladysz et al. 2008) or the difference between various moments or combinations thereof (Gladysz and Christou 2009; Gladysz, Yaitskova, and Christou 2010) to enhance the signal of faint companions hidden in stellar light. The first algorithm (Gladysz and Christou 2008) compared the empirical intensity distribution at the location of a suspected companion to a theoretical probability density function (PDF) of a speckle. The difference between distributions was quantified and translated to a confidence level associated with the detection decision. The moment-based approach compared only the parameters of measured distributions (Gladysz and Christou 2009; Gladysz, Yaitskova,

and Christou 2010) or their combinations. Expressions involving these parameters are chosen so that stellar light is suppressed while light from a companion is amplified. SSD-like approaches have been shown to reduce speckle noise from even fast atmospheric speckles and improve the contrast limit for the detection of substellar companions (Frazin 2016; Meeker et al. 2018; Stangalini et al. 2018).

In this chapter, we present an improved version of SSD to exploit noise-free photon-counting cameras like MEC, the Microwave Kinetic Inductance Detector (MKID) Exoplanet Camera (Walter et al. 2018; Meeker et al. 2018) on Subaru Telescope’s SCEAO instrument (Lozi et al. 2018b). Our approach statistically distinguishes a combination of constant and speckle intensity (both ultimately from the bright star) from incoherent light (from a planet or disk). At small separations, where ADI and SDI are least effective and the scientific questions are most pressing (Mawet et al. 2012), we expect SSD to offer strong improvements in the limiting detection contrast. This chapter describes the new photon counting SSD technique and demonstrates its performance on simulated data.

We organize the chapter as follows. In Section 6.2 we detail the simulation of photon lists to emulate the data expected from photon-counting cameras like MEC. Section 6.3 describes a formal extension of previous SSD analysis techniques with millisecond images using a maximum likelihood algorithm. We simulate the performance for various atmospheric conditions, planet brightnesses, and effective exposure times. Section 6.4 presents the new photon-counting SSD algorithm that estimates the incoherent light from a companion or disk directly from individual photon arrival timestamps. We demonstrate the algorithm on a simulated telescope image. We discuss the main results in Section 6.5, and conclude with Section 6.6.

6.2 Simulating Photon Arrival Times

6.2.1 Modeling the Stellar Speckle Intensity

The statistics governing the off-axis intensity distribution of coherent light with a partially developed speckle pattern have been studied at length. Originally derived by Goodman (1975) and verified experimentally by Cagigal and Canales (2001) and Fitzgerald and Graham (2006), the probability of getting an instantaneous intensity I given I_c and I_s is governed by the Modified Rician (MR) distribution, defined as:

$$\rho_{\text{MR}} [I|I_c, I_s] = \frac{1}{I_s} \exp \left[-\frac{I + I_c}{I_s} \right] I_0 \left[\frac{2\sqrt{II_c}}{I_s} \right], \quad (6.1)$$

where $I_0[x]$ denotes the zero-order modified Bessel function of the first kind. The parameter I_c represents the intensity of the "constant" part of the diffraction pattern, i.e. the PSF of a star without the atmosphere, while I_s is the intensity of the seeing halo which manifests as a "speckle" pattern (Canales and Cagigal 1999; Aime and Soummer 2004; Soummer et al. 2007).

Typically the AO system attempts to confine all the speckles into the constant diffraction pattern. A coronagraph can remove or transform this coherent portion out of the image plane. For all of the calculations in this chapter we assume that the Strehl ratio, and by extension I_c and I_s , remain constant.

Speckle intensity is correlated temporally. In the limit of Kolmogorov atmospheric turbulence and frozen flow, the speckle decorrelation time can be thought of as the wind crossing time across the telescope pupil. In reality, the turbulence is not Kolmogorovic, the atmospheric turbulence evolves, and there are often multiple turbulence layers with different wind speeds and directions (Roddier, Gilli, and Lund 1982; Macintosh et al. 2005). Other high speed processes such as dome seeing, telescope vibrations from wind

buffeting, or the AO loop itself can further complicate the speckle temporal power spectrum density (PSD) (Stangalini et al. 2016). This may result in a faster speckle decorrelation time, a temporal PSD described by multiple exponential timescales, or possibly a dependence on the position in the focal plane.

For the purposes of creating the simulated data we characterize the speckle PSD as a simple exponential decay with a characteristic speckle lifetime of $\tau_s = 0.1$ s to roughly match empirical data in the near infrared (Fitzgerald and Graham 2006; Meeker et al. 2018; Goebel et al. 2018). The speckle lifetime can change drastically depending on the atmospheric conditions but our results here can be qualitatively understood in those cases by scaling all parameters in time. A core condition for the SSD technique is for many photons to arrive in a single speckle decorrelation time: if this is satisfied, individual speckle fluctuations can be probed without relying on prior knowledge of the true speckle temporal PSD.

6.2.2 Modeling an Incoherent Source

For high levels of AO correction that minimize the variations in instantaneous Strehl the light from a companion can be approximated as a delta function with a Poisson intensity distribution (Gladysz et al. 2008). This works well for faint sources where only the core of the PSF is visible. By ignoring variations in the Strehl we simplify our model; incoherent Poisson sources are parameterized by an intensity I_p and can be readily injected into images with the relevant photons uniquely identified in simulation. Any incoherent Poisson sources will be represented by this term including binary companions, planetary companions, extended sources, dark current, and read noise. For the coherent stellar light, in the case that $I_s = 0$, the MR distribution reduces to a Poisson distribution with intensity I_c and will be indistinguishable from I_p . In the demonstrations that follow,

we identify I_p as an injected planetary companion.

The intensity fluctuations associated with speckles elevate the noise floor above the typical shot noise, i.e., it is harder to measure the planet's intensity when it is embedded in a boiling speckle field. If the speckle temporal information is marginalized over, then we show in Appendix A that the total variance of temporally correlated intensities obeying MR statistics combined with photon shot noise is

$$\sigma_{I,\text{tot}}^2 \approx \frac{2\tau_s (I_s^2 + 2I_c I_s) + I_c + I_s + I_p}{T_{\text{tot}}}. \quad (6.2)$$

The variance of the measured intensity is inversely proportional to the total integration time, T_{tot} , as expected. This is the long exposure ($t_{\text{exp}} \gg \tau_s$) photon noise limit faced by all PSF subtraction techniques like ADI and SDI. In the limit that $2\tau_s I_s \ll 1$, this variance reduces to pure shot noise on the number of photons. For high contrast imaging, typical parameters might be $\tau_s = 0.1$ s, $I_s = 50$ s⁻¹, and $2\tau_s I_s = 10$, in which case the noise from speckle fluctuations will dominate. In Section 6.4, we show that this noise can be overcome by temporally resolving individual fluctuations in the speckle background.

6.2.3 Generating Mock Photon Lists

We have developed code¹ for quickly generating mock photon lists with an optional injected planet, corresponding to the output of a single MKID-like pixel which is single photon counting with low noise. The photon lists obey the following rules:

1. The underlying intensities are MR distributed, but are correlated in time with $\langle I(t)I(t + \delta t) - \langle I \rangle^2 \rangle \propto \exp[-\delta t/\tau_s]$.
2. I_c , I_s , and I_p are independently specified by the user, such that I_c and I_s govern

¹Part of MKID Pipeline python package available at <https://github.com/MazinLab/MKIDPipeline>

the MR statistics of the stellar intensity and I_p is the mean count rate of a Poisson source. The total intensity should have an expectation value of $\langle I \rangle = I_c + I_s + I_p$.

3. Due to the intrinsic dead time in an MKID, photons are removed from the list if they arrive within τ_0 of the previous valid photon’s arrival time.

The procedure for generating photon timestamps begins with creating a correlated list of random numbers that follow a Gaussian distribution. The random numbers are transformed to a uniform distribution ranging from 0 to 1, and finally transformed again to a MR distribution. The correlated MR sequence defines the intensity in photon counts for a small ($\ll \tau_s$) time bin (we used $200 \mu\text{s}$), and a Poisson draw on that “instantaneous” intensity determines the number of photons that will finally be placed into that bin. The photons are distributed according to a uniform distribution in each bin. For simplicity, we assume neighboring pixels to have uncorrelated photon lists.

Since MKIDs do not have the same dark current or readout noise as conventional semiconductor detectors we do not add any additional noise. However, during high count rates photons can be lost due to a firmware triggering lockout that acts as a non-paralyzable dead time (Eyken et al. 2015). We use $\tau_0 = 10 \mu\text{s}$ to match the latest firmware implemented on MEC. This dead time formulation could be used with quasi-photon counting EMCCDs to account for photon pile-up.

6.3 SSD with Short Exposure Images

Past efforts have been successful in detecting the position of (real or synthetic) faint companions in high contrast images. Gladysz and Christou (2008) and Meeker et al. (2018) fit histograms of light curves (see top/middle of Figure 6.1) and discriminate planets by extracting parameters of hypothesized distributions using e.g. the method of

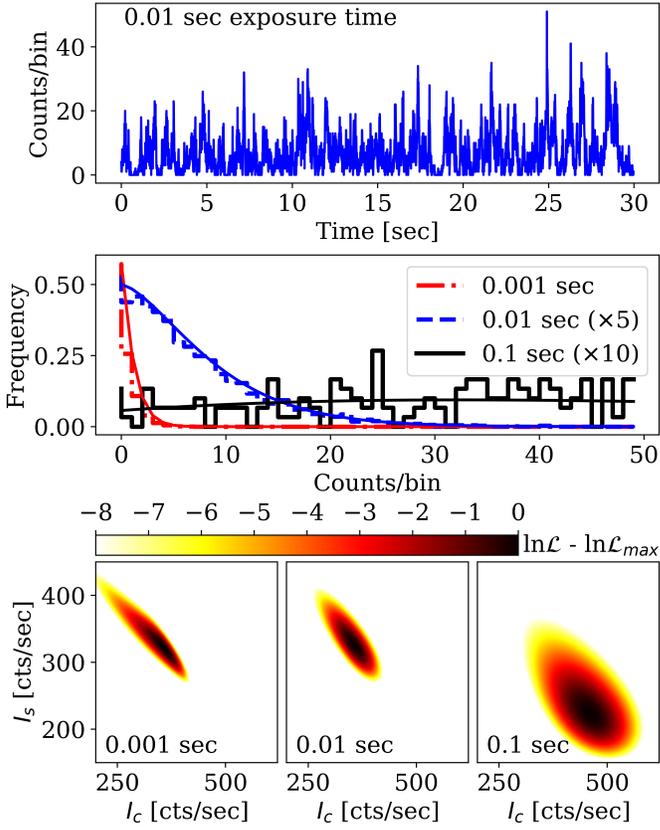


Figure 6.1: Top: a 30 s mock photon list simulated using the parameters $I_c = 300$, $I_s = 300$, $I_p = 0$ photons/s with $\tau_s = 0.1$ s, is binned into 0.01 s exposures to form a light curve. Middle: The same photon list is binned by three different exposure times and used to plot intensity histograms. The best-fit modified Rician functions are overplotted. The distribution changes shape as the bin size is varied; the corresponding fitted parameters also change. The histograms for 0.01 and 0.1 s exposure times have been scaled by factors of 5 and 10 respectively. Bottom: the likelihood is marginalized over I_p for the three different exposure times to illustrate how the best fit parameters evolve with bin size. The darkest area in the plot represents the maximum likelihood and determines the best fit values of I_c and I_s . The same photon list was used for all plots.

moments and comparing those parameters, or their combinations, to the rest of the field. For example, a large I_c/I_s ratio can be used as a merit function since faint companions can masquerade as static speckles. Stangalini et al. (2018) use a discriminator from a direct recurrence quantification analysis of the light curves to search for companions. Gladysz and Christou (2009) and Gladysz, Yaitzkova, and Christou (2010) also fit histograms of the light curves but use a deconvolution of the modeled probability distribution functions to further extract the companion’s photometry.

In this section we extend previous work by extracting the position and photometry of astrophysical sources (point or extended sources) using a maximum likelihood algorithm.

6.3.1 Maximum Likelihood Model for Discrete Light Curves

We model the intensity in the focal plane with three parameters I_c , I_s , and I_p and calculate their most likely values with an algorithm operating on the light curve. This allows for a direct detection of the non-stellar intensity, I_p , in the form of a point source like a planet or an extended source like a protoplanetary disk. Furthermore, we can naturally introduce photon noise from a low-intensity photon counting regime. By allowing the MR distribution in Equation (6.1) to suffer a Poisson-Mandel transformation (Cagigal and Canales 1999; Aime and Soummer 2004), the discrete stellar intensity distribution becomes

$$\begin{aligned}
 p_\star[n|I_c, I_s] &= \int_0^\infty \frac{I^n}{n!} \exp[-I] \rho_{\text{MR}}[I] dI \\
 &= \frac{1}{I_s + 1} \left(1 + \frac{1}{I_s}\right)^{-n} \exp\left[-\frac{I_c}{I_s}\right] \\
 &\quad \times L_n\left[\frac{-I_c}{I_s^2 + I_s}\right] \exp\left[\frac{I_c}{I_s^2 + I_s}\right], \tag{6.3}
 \end{aligned}$$

where L_n is the n^{th} Laguerre polynomial. The units of intensity for this section are number of photons per exposure time, t_{exp} . If a planet (or some other source that is incoherent with the star) exists in the field, the intensity distribution at that location will be the discrete convolution of p_\star with the planet's independent probability distribution p_p . To simplify, we assume the incoherent source has a Poisson probability distribution, $p_p[m|I_p] = \exp[-I_p] I_p^m / m!$ with average intensity I_p . The likelihood of the i th bin of a light curve containing k photons given I_c , I_s , and I_p is

$$\mathcal{L}_i = \sum_{m=0}^{k_i} p_p[m|I_p] p_\star[k_i - m|I_c, I_s], \tag{6.4}$$

which becomes

$$\mathcal{L}_i = \frac{1}{I_s + 1} \exp \left[-I_p - \frac{I_c I_s}{I_s^2 + I_s} \right] \times \sum_{m=0}^{k_i} \frac{I_p^m}{m!} \left(1 + \frac{1}{I_s} \right)^{-(k_i - m)} L_{k_i - m} \left[\frac{-I_c}{I_s^2 + I_s} \right]. \quad (6.5)$$

The likelihood of the entire light curve is $\mathcal{L} = \prod_i \mathcal{L}_i$. We use a L-BFGS-B optimizer to find the most likely values of I_c , I_s , and I_p .

The maximum likelihood estimate of the intensity distribution is overplotted onto the light curve histograms for three different exposure times in the middle panel of Figure 6.1. The bottom panel shows the likelihood functions marginalized over I_p for these three exposure times.

6.3.2 Performance of Millisecond Imaging SSD

To understand the performance of this millisecond imaging SSD algorithm and the extent to which it can improve exoplanet detections, we produce receiver operator characteristic (ROC) curves (Tanner Jr. and Swets 1954; DeLong, DeLong, and Clarke-Pearson 1988; Krzanowski and Hand 2009; Jensen-Clem et al. 2017). We achieve this by generating an ensemble of mock photon lists with the same nominal values of I_c , I_s , and $I_p = 0$ and calculating maximum likelihood estimates to build up a distribution of I_p corresponding to the signal-absent hypothesis. Next we inject a planet with $I_p > 0$ and calculate maximum likelihood estimates to build a distribution for I_p corresponding to the signal-present hypothesis. These distributions are shown in the left column of Figure 6.2. By choosing a detection threshold, we can calculate a true positive rate and false positive rate from the signal-present distribution and signal-absent distribution respectively. The ROC curves in the right column of Figure 6.2 are generated by varying the detection thresholds.

The right column of Figure 6.2 demonstrates the performance of the millisecond

imaging SSD algorithm for different exposure times. The shape of the discrete intensity distribution can change depending on exposure time, which systematically affects the resulting maximum likelihood estimates (MLE) for I_c , I_s , and I_p along with their uncertainties. In some cases there is a significant probability for the most likely value of I_p to be equal to zero. In the case that $I_c \gg I_s$, the MLE probability can be multimodal with a significant peak at MLE $I_p \approx \text{True } I_c + \text{True } I_p$. Both these behaviors appear for the same reason: when there is little modulation by I_s of photons associated with I_c , then I_c becomes difficult to distinguish from I_p . However, the total flux is still accurately recovered.

We find that there exists an optimal camera frame rate that leads to the most precise probability distribution for I_c , I_s , and I_p . For exposure times that are too short ($t_{\text{exp}} \ll \delta t$, the photon inter-arrival time) the correlation between subsequent photon arrivals is lost because frames are interpreted as an unordered set containing only 0 or 1 photons. For exposure times that are too long ($t_{\text{exp}} \gg \tau_s$) the speckle temporal information is averaged over. This tends to artificially decrease the I_s component of the speckle intensity directly in exchange for increasing the I_c component (see Fig 6.1 bottom). For those that use I_c/I_s as a merit function this means the part of the image without a planet will be polluted making it harder to discriminate the location of a companion.

6.4 SSD in the Photon Counting Regime

6.4.1 Maximum Likelihood Model for Photon Arrival Times

In the previous section the most precise extraction of the model parameters I_c , I_s , and I_p required the use of an optimal exposure time. Since this optimal exposure time is dependent on the speckle intensity as well as the speckle decorrelation time, the optimal

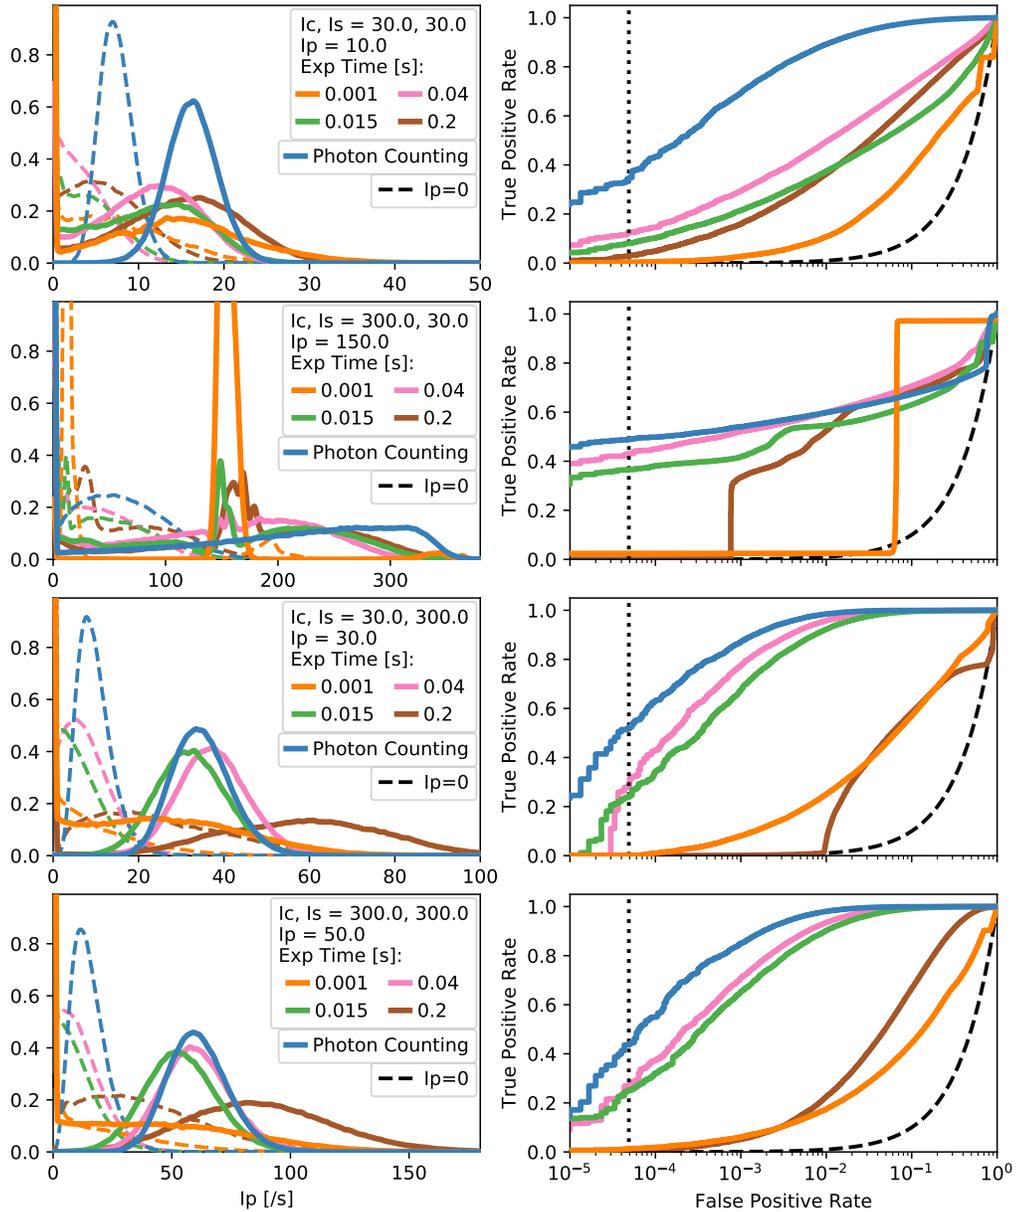


Figure 6.2: Performance of our millisecond imaging SSD algorithm (Section 6.3) compared to our photon counting SSD algorithm (Section 6.4). Left panels: histograms of the maximum likelihood estimates of I_p , computed using $3 \cdot 10^5$ 30 s mock photon lists. Maximum likelihood estimates for I_p are calculated for various effective exposure times and for the case with (solid line) and without ($I_p = 0$, dashed line) an injected planet. The y -axis in the left column is arbitrary. The probability distributions are used to calculate the true/false positive rates for the receiver operator characteristic (ROC) curve (right). The vertical dotted line at $1/20000$ (for the 20000 pixels in MEC) roughly indicates the maximum tolerable false positive rate. The full photon-counting SSD algorithm (blue lines) described in Section 6.4 outperforms the cases with nonzero exposure times.

exposure time will vary across the image. Conversely, if a single exposure time is chosen there may be a systematic loss of precision across some portions of the image. In order to avoid this pitfall, we develop the posterior probability for I_c , I_s , and I_p directly from the set of photon inter-arrival times.

We start by considering the (normalized) probability density for the next inter-photon arrival interval, δt , given a fixed intensity I :

$$p[\delta t|I] = Ie^{-I\delta t}. \quad (6.6)$$

The speckle field intensity is not fixed but varies in time with the modified Rician probability density described in Equation (6.1). At a fixed point in time, the probability density of the next photon arrival time given I_c and I_s becomes

$$p[\delta t|I_c, I_s] = \int_0^\infty p[\delta t|I]\rho_{\text{MR}}[I|I_c, I_s]dI, \quad (6.7)$$

where we have integrated the Poisson probability density from Equation (6.6) over all possible instantaneous stellar intensities, I . Intensities are considered to be in units of photons per second. We assume here that the speckle lifetime is much longer than the time it takes a photon to arrive, $\tau_s \gg \delta t$.

With a set of photon inter-arrival times, $\{\delta t_i\}$, we want the relative probability that a δt is realized. At moments that happen to have higher instantaneous intensities the number of short δt_i 's will be increased. Thus, the relative probability of a photon inter-arrival time being in our data becomes

$$p[\delta t|I_c, I_s] \propto \int_0^\infty p[\delta t|I]\rho_{\text{MR}}[I|I_c, I_s]IdI. \quad (6.8)$$

Finally, we consider the case that light incoherent with the star, such as from a planet,

is in the field. We consider only the simplest case in which this source has constant intensity I_p and is governed by Poisson statistics. In this case, the relative probability of δt given I_c, I_s, I_p is

$$p[\delta t|I_c, I_s, I_p] \propto \int_0^\infty p[\delta t|I + I_p]\rho_{\text{MR}}[I|I_c, I_s](I + I_p)dI, \quad (6.9)$$

which can be evaluated analytically. We then find the normalization constant by setting

$$\int_{\tau_0}^\infty p[\tau|I_c, I_s, I_p]d\tau = 1. \quad (6.10)$$

Equation (6.10) accounts for the non-paralyzable detector dead time, τ_0 , intrinsic to MKIDs by replacing the lower limit of integration with τ_0 . In Equation (6.9), we set the likelihood equal to zero for $\delta t < \tau_0$. It is convenient to use the change of variables,

$$u_i = \frac{1}{1 + I_s\delta t_i}, \quad u_{\text{max}} = \frac{1}{1 + I_s\tau_0}. \quad (6.11)$$

The log-likelihood, $\log(\mathcal{L})[I_c, I_s, I_p]$, then becomes

$$\begin{aligned} \log(\mathcal{L}) &= \sum_{i=1}^N p[\delta t_i|I_c, I_s, I_p] \\ &= \sum_{i=1}^N (u_i - 1) \left(\frac{I_p}{I_s u_i} + \frac{I_c}{I_s} \right) \\ &\quad + \sum_{i=1}^N \log [I_c^2 u_i^5 + 4I_c I_s u_i^4 + (2I_s^2 + 2I_p I_c) u_i^3 \\ &\quad \quad \quad + 2I_p I_s u_i^2 + I_p^2 u_i] \\ &\quad - N \frac{(u_{\text{max}} - 1)(I_p + I_c u_{\text{max}})}{I_s u_{\text{max}}} \\ &\quad - N \log [I_c u_{\text{max}}^3 + I_s u_{\text{max}}^2 + I_p u_{\text{max}}]. \end{aligned} \quad (6.12)$$

As in Section 6.3, we assume that AO performance remains stable so that I_c and I_s remain constant over the course of observations.

Equation (6.12) for the photon counting SSD algorithm replaces Equation (6.5) from Section 6.3. We use a Newton conjugate-gradient search to find the maximum of the log-likelihood space and recover the best estimates for I_c , I_s , and I_p .

The photon-counting SSD algorithm consistently outperforms the millisecond imaging SSD algorithm from Section 6.3, which marginalizes over temporal information via the exposure time (see Figure 6.2). In Figure 6.3 we show the performance of the photon counting SSD algorithm under conditions given by various combinations of I_c and I_s , and with various planet brightnesses I_p . The algorithm performs well with a high true positive detection rate even for the case when I_c and I_s are both large. However, the performance suffers in the case of $I_c \gg I_s$. As in millisecond imaging SSD, I_p and I_c become indistinguishable in this limit.

6.4.2 Maximum A Posteriori Estimation

Prior knowledge of a parameter can improve the estimates of I_c , I_s , and I_p . In our case, we commonly have information on the I_c parameter which corresponds to the static or quasi-static speckle point spread function either from a telescope model or measured on a reference star. For such situations, the log likelihood function in Equation (6.12) can be modified with a Gaussian prior as

$$\log(\mathcal{L}) \rightarrow \log(\mathcal{L}) - \frac{1}{2} \left(\left(I_c - \tilde{I}_c \right) / \sigma[\tilde{I}_c] \right)^2 \quad (6.13)$$

where $\tilde{I}_c \pm \sigma[\tilde{I}_c]$ is the prior on I_c . The new estimates of I_c , I_s , and I_p become maximum a posteriori (MAP) estimates.

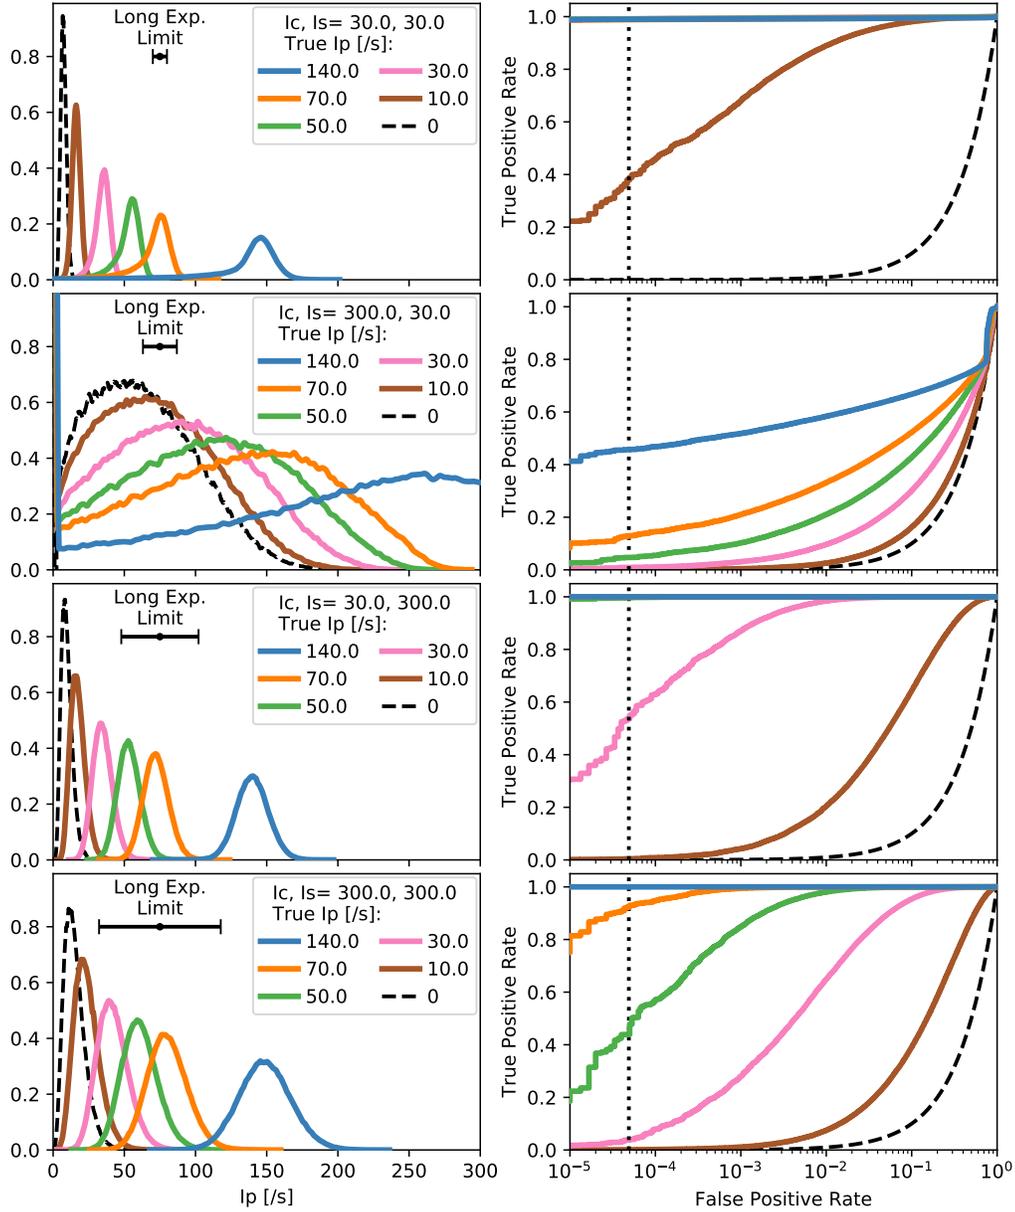


Figure 6.3: Performance of our photon-counting SSD algorithm. Left panels: histograms of the maximum likelihood estimates of I_p , computed using $3 \cdot 10^5$ 30 s mock photon lists for each set of parameters. The $\pm\sigma$ for the long exposure photon noise limit (Equation (6.2)) is shown with an error bar. The y -axis in the left column is arbitrary. The probability distributions are used to calculate the true/false positive rates for the receiver operator characteristic (ROC) curve (right panels). The vertical dotted line at $1/20000$ (for the 20000 pixels in MEC) roughly indicates the maximum tolerable false positive rate. The algorithm performs well with a high true positive detection rate even for the case when I_c, I_s are both large. However, the performance suffers in the case of $I_c \gg I_s$ (row 2).

6.4.3 Performance on Simulated Telescope Image

We evaluate the performance of the photon-counting SSD algorithm on a simulated 30 s telescope image without a coronagraph. We identify I_c as the Airy-ring pattern of a diffraction limited telescope with a circular unobstructed aperture, I_s as the seeing halo from atmospheric speckles with a Strehl ratio of 0.7, and I_p as a series of injected planets at various brightnesses and separations. The faintest planet’s total intensity is 40 photons/s, for a contrast of $5 \cdot 10^{-5}$ with the host star, and is separated by $3.5 \lambda/D$. Assuming a 5% end-to-end throughput on an 8.2 m telescope, the stellar magnitude in the near infrared is approximately $J = 10$. We assume a speckle decorrelation time of $\tau_s = 0.1$ s and detector dead time of $\tau_0 = 10 \mu\text{s}$.

Each pixel has an independent 30 s photon list generated from the “True” I_c , I_s , and I_p shown in panels (c), (f), and (i) respectively of Figure 6.4. Spatial correlations in the photon lists are ignored for simplicity. The average intensity realized for each pixel is shown in Figure 6.4a. Figure 6.4b is the average intensity minus the expected light from the star, (True $I_c + \text{True } I_s$), which illustrates the best possible long exposure PSF subtraction (compare the background variance to Equation (6.2)). The MLE I_c , I_s , and I_p are shown in Figure 6.4d, g, and j respectively. In Figure 6.4e, h, and k, we calculate the MAP estimates. We used the True $I_c \pm 3 \cdot \sqrt{\text{True } I_c}$ as a Gaussian prior on I_c ; in practice one could use a telescope model or a reference PSF. The central black dot with radius $1.22 \lambda/D$ is not a coronagraph but simply obscures the on-axis light for convenience.

Comparing Figure 6.4d to Figure 6.4b shows that the MLE I_p from the photon counting SSD algorithm recovers the injected planets better than a perfect stellar PSF subtraction (i.e. subtraction of the true I_c and I_s). Figure 6.4d and e show that the MAP estimate for I_p is not significantly better than the MLE I_p (although the MAP estimate

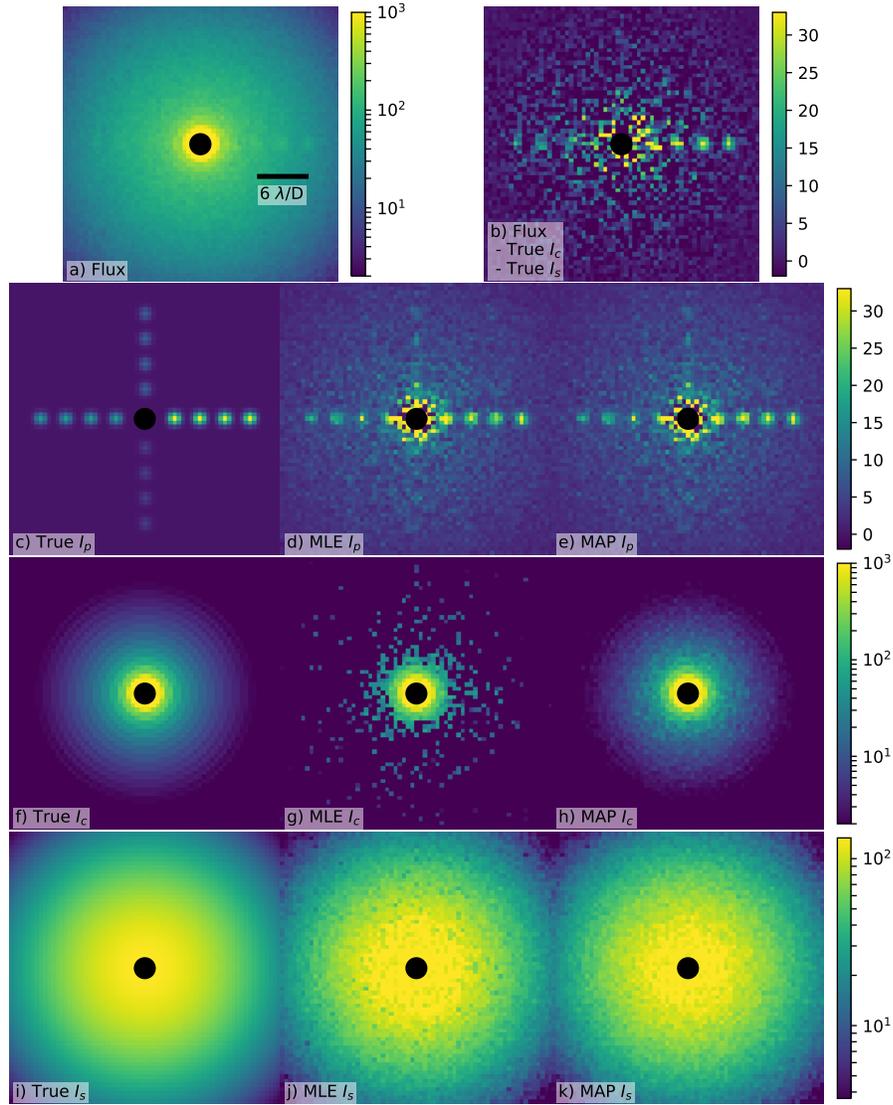


Figure 6.4: Performance of our photon-counting SSD algorithm on simulated telescope images. Panel (a) shows the average intensity of simulated photon lists in each pixel; panels (c), (f), and (i) show the parameters used. Panel (b) subtracts stellar flux from (f) and (i) from the average intensity, (a), illustrating the speckle variance in a long exposure image. Panel (b) represents the theoretical limit of perfect PSF subtraction in a single exposure subject to MR intensity fluctuations. The photon-counting SSD algorithm results in the maximum likelihood estimates shown in (d), (g), and (j). Using a priori knowledge of the I_c parameter (we used a Gaussian prior of $\text{True } I_c \pm 3 \cdot \sqrt{\text{True } I_c}$) we generate the maximum a posteriori (MAP) estimates shown in (e), (h), and (k). The planet signals extracted from the MAP I_p estimate in (e) are not significantly improved compared to the MLE I_p in (d). However, the SSD results in (d) and (e) both extract the planet better than the perfect PSF subtraction shown in (b). All images represent 30 seconds of data on a magnitude $J = 10$ star with an 8.2 m telescope; all units are photons/second.

Table 6.1: Companion SSD Signal-to-Noise Ratio

Contrast	Separation = $3.5 \lambda/D$			Separation = $6.5 \lambda/D$			Separation = $9.5 \lambda/D$			Separation = $12.5 \lambda/D$		
	MLE	MAP	Limit ^a	MLE	MAP	Limit ^a	MLE	MAP	Limit ^a	MLE	MAP	Limit ^a
$4 \cdot 10^{-4}$	4.8	5.3	2.3	6.8	7.0	3.8	7.6	8.3	5.7	7.0	9.4	8.8
$2 \cdot 10^{-4}$	2.8	3.0	1.2	4.0	4.0	1.9	4.9	4.9	2.9	5.6	5.9	4.5
$1 \cdot 10^{-4}$	1.6	1.7	0.6	2.3	2.3	1.0	2.8	2.8	1.4	3.5	3.5	2.3
$5 \cdot 10^{-5}$	0.9	0.9	0.3	1.3	1.2	0.5	1.6	1.5	0.7	2.0	2.0	1.1

^aDetection limit for a 30 s long exposure with perfect PSF subtraction of the stellar light.

for I_c is more precise). This is surprising because the MAP estimate includes a prior on I_c that should help discriminate between I_c and I_p . That this is not the case indicates we can take full advantage of the photon counting SSD algorithm without prior knowledge of the telescope PSF.

For the central pixel of each planet, we used 10^5 independent photon lists to calculate the signal-to-noise ratio $S/N = (\langle I_p \rangle - \langle \text{Background} \rangle) / (\text{std. dev.} \langle I_p \rangle)$ where the $\langle \text{Background} \rangle$ is estimated by not injecting a planet. These are recorded in Table 6.1. The long exposure photon noise limit is also recorded in Table 6.1 where the estimated I_p is equal to the total flux minus the light from the star, ($\text{True } I_c + \text{True } I_s$). While the results are from the simulated ensemble, they match the results from Equation (6.2). Table 6.1 is representative of 30 seconds of data, but the S/N will scale with $\sqrt{T_{\text{tot}}}$. For a 2 minute exposure, all values in Table 6.1 should be scaled up by a factor of 2. The S/N was calculated using only the central pixel for convenience but would be larger if the surrounding pixels were considered.

6.5 Discussion of SSD Simulations

While it remains impossible to beat the photon shot noise \sqrt{N} , Figure 6.4 shows that the photon counting SSD algorithm can beat the long exposure ($t_{\text{exp}} \gg \tau_s$) photon noise limit described by Equation (6.2). Table 6.1 quantifies the improvement in the S/N as a

factor of 3 in the case of faintest planet ($5 \cdot 10^{-5}$) at the nearest separation ($3.5 \lambda/D$). This is possible because the speckle fluctuations are temporally resolved ($t_{\text{exp}} \ll \tau_s$) and individual speckles are probed by multiple photons ($\delta t \ll \tau_s$). In the case that the fluctuations from stellar speckles dominate the variance of the total intensity ($2I_s\tau_s \gg 1$), fast, noiseless detectors like MKIDs or EMCCDs are needed to dig beneath the noise. This can occur in high contrast imaging at separations between $1 \lambda/D > r < \text{seeing radius}$ but is especially important at $\lesssim 5 \lambda/D$ where ADI and SDI start to lose their effectiveness (depending on spectral coverage, sky rotation, and AO performance).

SSD will benefit ADI and SDI by reducing speckle noise from the data that is fed into those algorithms. ADI processing can be approached the same way as usual, but instead of using raw images one would use the I_p images produced with SSD. The modulation of the planet location will be unaffected by SSD. Similarly for SDI, the algorithm would be given I_p maps at various wavelengths. This approach would require wavelength information for each detected photon, which is an intrinsic feature of an MKID detector.

Our SSD algorithm does not perform well when $I_c \gg I_s$ as seen in Figure 6.3. In this regime there is little modulation of the static speckle intensity I_c by the atmospheric speckle field I_s , and as a result I_c starts to become indistinguishable from the Poisson distributed I_p . This can be greatly mitigated with a coronagraph and possibly active speckle nulling (Martinache et al. 2014) which directly reduce I_c in the image.

On a space-borne telescope atmospheric speckles are not a concern, implying that the image-plane intensity will have different temporal behavior from ground-based observatories. While our SSD algorithm relies on the intensity following a MR distribution, in general any distribution can be used so long as it is known. It may also be possible to modulate speckles in a controlled way using onboard AO.

6.6 SSD Algorithm Conclusions and Future Work

We can exploit photon arrival time statistics with a stochastic speckle discrimination (SSD) algorithm to distinguish planets from speckles. We first extend previous work with a formalized maximum likelihood algorithm operating on light curves with fixed, albeit fast, exposure times. We find that the likelihood space can sometimes result in bimodal behavior in the case of $I_c \gg I_s$. Additionally, the choice in exposure time can systematically skew the MLE of I_p and inflate its variance. More generally, with a fixed exposure time, the performance will change as a function of parameters I_c , and I_s . This is a problem because I_c and I_s can change with observing conditions as well as with separation from the host star.

To overcome these difficulties, we have developed a new photon-counting SSD algorithm that calculates the maximum likelihood for I_p directly from the individual photon arrival times. With this approach the likelihood space becomes smooth and unimodal and the precision is maximized. The planet detection performance can be better by a factor of 2 than perfect stellar PSF subtraction of a long exposure. This requires fast noiseless detectors like MKIDs.

We have made several simplifying assumptions in our analysis. We take the speckle temporal PSD of our simulated data to be described by the single exponential timescale τ_s , we assume I_c , I_s , and I_p remain constant (ignoring variations to the instantaneous Strehl), and we assume that the MR distribution accurately describes the off-axis stellar intensity. Finally, we ignore chromaticity. These assumptions represent avenues of exploration for future work. The speckle temporal PSD can be measured and used to more accurately simulate photon lists. Since the instantaneous Strehl can be measured, a future implementation of this algorithm might use that information in a more sophisticated model for the incoherent light like in Gladysz et al. (2008). This would require

an understanding of how Strehl variations effect I_c and I_s .

SSD algorithms for highly corrected AO images are ultimately constrained on two fronts. First, the photon arrival time δt must be much shorter than the speckle decorrelation time τ_s : we need many photons to characterize the properties of a materialized speckle. Second, the performance degrades when I_c is large but I_s is small, because the incoherent planet light masquerades as the static speckles described by I_c . Fortunately, this can be improved with a coronagraph.

In return, SSD algorithms are most useful when I_s is large, which is often inescapable at small inner working angles. Furthermore, the results are not directly dependent on separation (although they are dependent on I_c and I_s , which are larger at small separations). This makes SSD a powerful post-processing technique at small inner working angles where it can complement more established techniques like ADI and SDI. At high speckle intensities, SSD can even outperform the theoretical limits of a perfect implementation of ADI and SDI.

Appendix A

Deriving Speckle Noise in a Long Exposure

We derive here the noise of a long ($t_{\text{exp}} \gg \tau_s$) exposure subject to speckle statistics. This is not simply photon shot noise, but arises due to the fluctuations of the modified Rician itself. We consider a modified Rician parametrized by I_c and I_s with an exponential decorrelation time τ_s . The basic statistics of the distribution are

$$\langle I \rangle = I_c + I_s \tag{A1}$$

$$\sigma_I^2 = \langle I^2 \rangle - \langle I \rangle^2 = I_s^2 + 2I_c I_s \tag{A2}$$

$$\langle (I_i - \langle I \rangle)(I_j - \langle I \rangle) \rangle = \sigma_I^2 \exp \left[-\frac{|t_i - t_j|}{\tau_s} \right]. \tag{A3}$$

where brackets, $\langle \rangle$, denote an ensemble average.

We wish to compute the variance of the mean intensity measured over a finite time

interval T_{tot} , which we divide into N subintervals, each of length δt .

$$\bar{I} = \frac{1}{N} \sum_i I_i. \quad (\text{A4})$$

Here, the overbar denotes a temporal average. The variance is

$$\langle \bar{I}^2 \rangle - \langle \bar{I} \rangle^2 = \frac{1}{N^2} \left\langle \sum_i \sum_j I_i I_j \right\rangle - \frac{1}{N^2} \left\langle \sum_i I_i \right\rangle^2 \quad (\text{A5})$$

$$= \frac{1}{N^2} \left\langle \sum_i \sum_j (I_i - \langle I \rangle) (I_j - \langle I \rangle) \right\rangle \quad (\text{A6})$$

$$= \frac{\sigma_I^2}{N^2} \sum_i \sum_j \exp \left[-\frac{|t_i - t_j|}{\tau_s} \right]. \quad (\text{A7})$$

Now we will set $t_i = i\delta t$. For the third step below, we use $\delta t/\tau_s \ll 1$. More generally, we take $\delta t \rightarrow 0$, $N \rightarrow \infty$.

$$\begin{aligned} \frac{1}{N^2} \sum_i \sum_j \exp \left[-\frac{|t_i - t_j|}{\tau_s} \right] &= \\ &= \frac{1}{N^2} \sum_i \sum_j \exp \left[-\frac{|i - j|\delta t}{\tau_s} \right] \end{aligned} \quad (\text{A8})$$

$$= \frac{1}{N^2} \sum_{i=0}^{N-1} \left(\sum_{j=0}^{N-i-1} \exp \left[-\frac{j\delta t}{\tau_s} \right] + \sum_{j=0}^i \exp \left[-\frac{j\delta t}{\tau_s} \right] - 1 \right) \quad (\text{A9})$$

$$= \frac{1}{N^2} \sum_{i=0}^{N-1} \left(\frac{\tau_s}{\delta t} \left(2 - \exp \left[-\frac{(N-i)\delta t}{\tau_s} \right] - \exp \left[-\frac{(i+1)\delta t}{\tau_s} \right] \right) - 1 \right) \quad (\text{A10})$$

$$= \frac{2\tau_s}{N\delta t} - \frac{2\tau_s}{N\delta t} \frac{1}{N} \left(\sum_{i=1}^N \exp \left[-\frac{i\delta t}{\tau_s} \right] \right) - \frac{1}{N} \quad (\text{A11})$$

$$= \frac{2\tau_s}{N\delta t} - 2 \left(\frac{\tau_s}{N\delta t} \right)^2 \left(1 - \exp \left[-\frac{(N+1)\delta t}{\tau_s} \right] \right) + \frac{2\tau_s}{N^2\delta t} - \frac{1}{N}. \quad (\text{A12})$$

Taking $N \rightarrow \infty$, $\delta t \rightarrow 0$, and $N\delta t = T_{\text{tot}}$, we have

$$\langle \bar{I}^2 \rangle - \langle \bar{I} \rangle^2 = \left(\frac{2\tau_s}{T_{\text{tot}}} - 2 \left(\frac{\tau_s}{T_{\text{tot}}} \right)^2 \left(1 - \exp \left[-\frac{T_{\text{tot}}}{\tau_s} \right] \right) \right) \sigma_I^2 \quad (\text{A13})$$

$$= \left(\frac{2\tau_s}{T_{\text{tot}}} - 2 \left(\frac{\tau_s}{T_{\text{tot}}} \right)^2 \left(1 - \exp \left[-\frac{T_{\text{tot}}}{\tau_s} \right] \right) \right) (I_s^2 + 2I_c I_s). \quad (\text{A14})$$

In the limit of a very short integration, $T_{\text{tot}} \ll \tau_s$, the prefactor is unity (as expected). For an integration time much longer than the decorrelation time, $T_{\text{tot}} \gg \tau_s$ (as is more typical), Equation (A14) simplifies to

$$\langle \bar{I}^2 \rangle - \langle \bar{I} \rangle^2 \approx \left(\frac{2\tau_s}{T_{\text{tot}}} \right) (I_s^2 + 2I_c I_s). \quad (\text{A15})$$

Assuming $2I_s\tau_s \gg 1$ (the inter-photon arrival time from the speckle field is much shorter than the decorrelation time τ_s), Equation (A14) dominates over simple shot noise $\sigma^2 = (I_c + I_s)/T_{\text{tot}}$. Interestingly, setting the number of independent realizations of the modified Rician equal to T_{tot}/τ_s would miss the factor of two in Equation (A15). Adding shot noise back in, assuming $T_{\text{tot}} \gg \tau_s$, and including a component I_p incoherent with I_c and I_s , we have

$$\sigma_{I,\text{tot}}^2 \approx \frac{2\tau_s (I_s^2 + 2I_c I_s) + I_c + I_s + I_p}{T_{\text{tot}}}. \quad (\text{A16})$$

Bibliography

- Aime, C. and R. Soummer (2004). “Influence of speckle and Poisson noise on exoplanet detection with a coronagraph”. In: *2004 12th European Signal Processing Conference*, pp. 1071–1074.
- Allen, C. A., D. E. Franz, and S. H. Moseley (2006). “Compliant system of polyimide microwires for cryogenic detector applications”. In: *Journal of Vacuum Science & Technology A* 24.4, pp. 1552–1555.
- Barends, Rami et al. (2016). “Digitized Adiabatic Quantum Computing with a Superconducting Circuit”. In: *Nature* 534, pp. 222–226.
- Bellei, Francesco et al. (2016). “Free-space-coupled superconducting nanowire single-photon detectors for infrared optical communications”. In: *Opt. Express* 24.4, pp. 3248–3257.
- Bottom, Michael et al. (2016). “Speckle nulling wavefront control for Palomar and Keck”. In: *SPIE* 9909, pp. 9909 –9909 –16.
- Bottom, Michael et al. (2018). “Smartphone scene generator for efficient characterization of visible imaging detectors”. In: *Proc. SPIE*. Vol. 10709.
- Bowler, Brendan P. (2016). “Imaging Extrasolar Giant Planets”. In: *PASP* 128.968, p. 102001.
- Cagigal, Manuel P. and Vidal F. Canales (1999). “Photon statistics in adaptive optics”. In: *Optics in Atmospheric Propagation and Adaptive Systems III*. Ed. by Anton Kohnle and John D. Gonglewski. Vol. 3866. SPIE, pp. 165–175.
- Cagigal, Manuel Perez and Vidal F. Canales (2001). “Experimental checking of the Rician statistics in partially compensated wave fronts”. In: *Optical Engineering* 40.12, pp. 2690 –2697 –8.
- Calvo, M. et al. (2016). “The NIKA2 Instrument, A Dual-Band Kilopixel KID Array for Millimetric Astronomy”. In: *Journal of Low Temperature Physics* 184.3, pp. 816–823. ISSN: 1573-7357.
- Canales, Vidal F. and Manuel P. Cagigal (1999). “Rician distribution to describe speckle statistics in adaptive optics”. In: *Appl. Opt.* 38.5, pp. 766–771.
- Carbillet, Marcel et al. (2011). “Apodized Lyot coronagraph for SPHERE/VLT”. In: *Experimental Astronomy* 30.1, pp. 39–58. ISSN: 1572-9508. DOI: 10.1007/s10686-011-9219-4.

- Close, L. M. et al. (2013). “Diffraction-Limited Visible Light Images of Orion Trapezium Cluster with the Magellan Adaptive Secondary Adaptive Optics System (MagAO)”. In: *ApJ* 774.2, p. 94.
- Cook, Timothy et al. (2015). “Planetary Imaging Concept Testbed Using a Recoverable Experiment-Coronagraph (PICTURE C)”. In: *Journal of Astronomical Telescopes, Instruments, and Systems* 1.4, p. 044001.
- Crepp, Justin R. et al. (2011). “Speckle Suppression with the Project 1640 Integral Field Spectrograph”. In: *ApJ* 729.2, p. 132.
- Daal, M. et al. (2019). “Properties of selected structural and flat flexible cabling materials for low temperature applications”. In: *Cryogenics* 98, pp. 47–59. ISSN: 0011-2275.
- Daal, Miguel D. (2015). “Kinetic Inductance Detectors for Dark Matter Searches”. PhD thesis. University of California, Berkeley.
- Day, Peter et al. (2003). “A broadband superconducting detector suitable for use in large arrays”. In: *Nature* 425, pp. 817–21. DOI: 10.1038/nature02037.
- DeLong, ER, DM DeLong, and DL Clarke-Pearson (1988). “Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach”. In: *Biometrics* 44.3, 837–845. ISSN: 0006-341X.
- Dodkins, R. et al. (2018). “MKID digital readout tuning with deep learning”. In: *Astronomy and Computing* 23, pp. 60–71. ISSN: 2213-1337.
- Eyken, J. C. van et al. (2015). “The ARCONS Pipeline: Data Reduction for MKID Arrays”. In: *ApJS* 219.1, p. 14.
- Fitzgerald, M. P. and J. R. Graham (2006). “Speckle Statistics in Adaptively Corrected Images”. In: *ApJ* 637, pp. 541–547.
- Frazin, Richard A. (2016). “Empirical Green’s function approach for utilizing millisecond focal and pupil plane telemetry in exoplanet imaging”. In: *Adaptive Optics Systems V*. Vol. 9909. SPIE, p. 99093X.
- Gao, Jiansong (2008). “The Physics of Superconducting Microwave Resonators”. PhD thesis. California Institute of Technology.
- Gerard, Benjamin L. et al. (2019). “A Chromaticity Analysis and PSF Subtraction Techniques for SCEXAO/CHARIS Data”. In: *ApJ* 158.1, p. 36.
- Gladysz, S., N. Yaitskova, and J. C. Christou (2010). “Statistics of intensity in adaptive-optics images and their usefulness for detection and photometry of exoplanets”. In: *Journal of the Optical Society of America A* 27.27, A64–A75.
- Gladysz, Szymon and Julian C. Christou (2008). “Detection of Faint Companions through Stochastic Speckle Discrimination”. In: *ApJ* 684.2, pp. 1486–1495.
- (2009). “Reference-less Detection, Astrometry, and Photometry of Faint Companions with Adaptive Optics”. In: *ApJ* 698.1, pp. 28–42.
- Gladysz, Szymon et al. (2008). “Temporal Variability and Statistics of the Strehl Ratio in Adaptive-Optics Images”. In: *PASP* 120.872, pp. 1132–1143.
- Goebel, Sean B. et al. (2018). “Measurements of Speckle Lifetimes in Near-infrared Extreme Adaptive Optics Images for Optimizing Focal Plane Wavefront Control”. In: *PASP* 130.992, p. 104502.

- Gomez Gonzalez, C. A. et al. (2016). “Low-rank plus sparse decomposition for exoplanet detection in direct-imaging ADI sequences - The LLSG algorithm”. In: *A&A* 589, A54.
- Goodman, J. W. (1975). “Statistical Properties of Laser Speckle Patterns”. In: *Laser Speckle and Related Phenomena*. Ed. by J. C. Dainty. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 9–75. ISBN: 978-3-662-43205-1.
- Guyon, Olivier (2005). “Limits of Adaptive Optics for High-Contrast Imaging”. In: *ApJ* 629.1, pp. 592–614.
- Guyon, Olivier and Jared Males (2017). “Adaptive Optics Predictive Control with Empirical Orthogonal Functions (EOFs)”. In: *arXiv e-prints*. arXiv: 1707.00570 [astro-ph.IM].
- Guyon, Olivier et al. (2010). “High-Contrast Imaging and Wavefront Control with a PIAA Coronagraph: Laboratory System Validation”. In: *PASP* 122.887, pp. 71–84.
- Harris, A. I. et al. (2012). “Note: Cryogenic microstripline-on-Kapton microwave interconnects”. In: *Review of Scientific Instruments* 83.8, 086105.
- Holland, W. S. et al. (2013). “SCUBA-2: the 10 000 pixel bolometer camera on the James Clerk Maxwell Telescope”. In: *MNRAS* 430.4, pp. 2513–2533.
- Jensen-Clem, Rebecca et al. (2017). “A New Standard for Assessing the Performance of High Contrast Imaging Systems”. In: *AJ* 155.1, p. 19.
- Jovanovic, N. et al. (2015). “The Subaru Coronagraphic Extreme Adaptive Optics System: Enabling High-Contrast Imaging on Solar-System Scales”. In: *PASP* 127.955, p. 890.
- Kellaris, N. et al. (2014). “Sub-Kelvin Thermal Conductivity and Radioactivity of Some Useful Materials in Low Background Cryogenic Experiments”. In: *Journal of Low Temperature Physics* 176.3, pp. 201–208. ISSN: 1573-7357.
- Keppler, M. et al. (2018). “Discovery of a planetary-mass companion within the gap of the transition disk around PDS 70”. In: *A&A* 617, A44, A44.
- Krzanowski, Wojtek J. and David J. Hand (2009). *ROC Curves for Continuous Data*. 1st. Chapman & Hall/CRC. ISBN: 1439800219, 9781439800218.
- Kushino, A. and S. Kasai (2016). “Development of semi-rigid cables for low temperature superconducting detectors”. In: *APS March Meeting Abstracts*.
- Kushino, A. et al. (2008). “Development of Superconducting Coaxial Cables for Cryogenic Detectors”. In: *Journal of Low Temperature Physics* 151.3, pp. 650–654. ISSN: 1573-7357.
- Kuzuhara, M. et al. (2013). “Direct Imaging of a Cold Jovian Exoplanet in Orbit around the Sun-like Star GJ 504”. In: *ApJ* 774.1, 11, p. 11.
- Lafrenière, David et al. (2007). “A New Algorithm for Point-Spread Function Subtraction in High-Contrast Imaging: A Demonstration with Angular Differential Imaging”. In: *ApJ* 660.1, pp. 770–780.
- Lagrange, A. M. et al. (2010). “A Giant Planet Imaged in the Disk of the Young Star β Pictoris”. In: *Science* 329.5987, p. 57.

- Lovis, C. et al. (2017). “Atmospheric characterization of Proxima b by coupling the SPHERE high-contrast imager to the ESPRESSO spectrograph”. In: *Astronomy and Astrophysics* 599, A16, A16.
- Lozi, Julien et al. (2015). “SCEXAO: the first high contrast exoplanet imager on an ELT?” In: *Adaptive Optics for Extremely Large Telescopes 4 - Conference Proceedings*. Vol. 1. 1.
- Lozi, Julien et al. (2018a). “Characterizing vibrations at the Subaru Telescope for the Subaru coronagraphic extreme adaptive optics instrument”. In: *Journal of Astronomical Telescopes, Instruments, and Systems* 4.4, pp. 1–13–13.
- Lozi, Julien et al. (2018b). “SCEXAO, an instrument with a dual purpose: perform cutting-edge science and develop new technologies”. In: *Adaptive Optics Systems VI*. Vol. 10703. SPIE, p. 1070359.
- Macintosh, B. et al. (2015). “Discovery and spectroscopy of the young jovian planet 51 Eri b with the Gemini Planet Imager”. In: *Science* 350.6256, pp. 64–67.
- Macintosh, Bruce et al. (2005). “Speckle lifetimes in high-contrast adaptive optics”. In: *Astronomical Adaptive Optics Systems and Applications II*. Vol. 5903. SPIE, 59030J.
- Macintosh, Bruce et al. (2014). “First light of the Gemini Planet Imager”. In: *Proceedings of the National Academy of Sciences* 111.35, pp. 12661–12666. ISSN: 0027-8424. DOI: 10.1073/pnas.1304215111.
- Macintosh, Bruce A. et al. (2008). “The Gemini Planet Imager: from science to design to construction”. In: *SPIE* 7015, pp. 7015–7015–13.
- Males, Jared R. and Olivier Guyon (2018). “Ground-based adaptive optics coronagraphic performance under closed-loop predictive control”. In: *Journal of Astronomical Telescopes, Instruments, and Systems* 4.1, pp. 1–21–21.
- Males, Jared R. et al. (2018). “MagAO-X: project status and first laboratory results”. In: *SPIE* 10703.
- Marois, Christian et al. (2000). “Efficient Speckle Noise Attenuation in Faint Companion Imaging”. In: *PASP* 112.767, pp. 91–96.
- Marois, Christian et al. (2006). “Angular Differential Imaging: A Powerful High-Contrast Imaging Technique”. In: *ApJ* 641.1, pp. 556–564.
- Marois, Christian et al. (2008). “Direct Imaging of Multiple Planets Orbiting the Star HR 8799”. In: *Science* 322.5906, p. 1348.
- Marois, Christian et al. (2014). “GPI PSF subtraction with TLOCI: the next evolution in exoplanet/disk high-contrast imaging”. In: *Adaptive Optics Systems IV*. Vol. 9148. SPIE, 91480U.
- Martinache, Frantz et al. (2014). “On-Sky Speckle Nulling Demonstration at Small Angular Separation with SCEXAO”. In: *PASP* 126.940, pp. 565–572.
- Mawet, D. et al. (2014a). “Fundamental Limitations of High Contrast Imaging Set By Small Sample Statistics”. In: *ApJ* 792.2, p. 97.
- Mawet, D. et al. (2016). “Keck Planet Imager and Characterizer: concept and phased implementation”. In: *SPIE* 9909, pp. 9909–9909–7.

- Mawet, Dimitri et al. (2012). “Review of small-angle coronagraphic techniques in the wake of ground-based second-generation adaptive optics systems”. In: *Proc. SPIE*. Vol. 8442.
- Mawet, Dimitri et al. (2012). “Review of small-angle coronagraphic techniques in the wake of ground-based second-generation adaptive optics systems”. In: *Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave*. Vol. 8442. SPIE, p. 844204.
- Mawet, Dimitri et al. (2014b). “Demonstration of vortex coronagraph concepts for on-axis telescopes on the Palomar Stellar Double Coronagraph”. In: *SPIE* 9143, pp. 9143–9143–8.
- Mazin, Benjamin A. et al. (2012). “A superconducting focal plane array for ultraviolet, optical, and near-infrared astrophysics”. In: *Opt. Express* 20.2, pp. 1503–1511.
- Mazin, Benjamin A. et al. (2013). “ARCONS: A 2024 Pixel Optical through Near-IR Cryogenic Imaging Spectrophotometer”. In: *PASP* 125.933, pp. 1348–1361. ISSN: 00046280, 15383873.
- McGarey, Patrick et al. (2014). “A 16-channel flex circuit for cryogenic microwave signal transmission”. In: *Proc. SPIE*. Vol. 9153, 91532F–91532F–6.
- McHugh, Sean et al. (2012). “A readout for large arrays of microwave kinetic inductance detectors”. In: *Review of Scientific Instruments* 83.4, p. 044702. DOI: 10.1063/1.3700812.
- Meeker, Seth R. et al. (2018). “DARKNESS: A microwave kinetic inductance detector integral field spectrograph for high-contrast astronomy”. In: *PASP* 130.988. ISSN: 00046280.
- Miller, Kelsey, Olivier Guyon, and Jared Males (2017). “Spatial linear dark field control: stabilizing deep contrast for exoplanet imaging using bright speckles”. In: *Journal of Astronomical Telescopes, Instruments, and Systems* 3, pp. 3–3–11.
- Morley, Caroline V. et al. (2015). “Thermal Emission and Reflected Light Spectra of Super Earths with Flat Transmission Spectra”. In: *ApJ* 815.2, p. 110.
- Olson, J.R. (1993). “Thermal conductivity of some common cryostat materials between 0.05 and 2 K”. In: *Cryogenics* 33.7, pp. 729–731. ISSN: 0011-2275.
- Paik, K. W. and A. Mogro-Campero (1994). “Studies on the high-temperature superconductor (HTS)/metal/polymer dielectric interconnect structure for packaging applications”. In: *IEEE Transactions on Components, Packaging, and Manufacturing Technology: Part B* 17.3, pp. 435–441. ISSN: 1070-9894. DOI: 10.1109/96.311794.
- Pappas, C. G. et al. (2016). “High-Density Superconducting Cables for Advanced ACT-Pol”. In: *Journal of Low Temperature Physics* 184.1, pp. 473–479. ISSN: 1573-7357.
- Pozar, David M. (2011). “Microwave Engineering”. In: 4th ed. Wiley. Chap. Microwave Network Analysis. ISBN: 0470631554.
- Racine, René et al. (1999). “Speckle Noise and the Detection of Faint Companions”. In: *PASP* 111.759, pp. 587–594.
- Roddier, F, J M Gilli, and G Lund (1982). “On the origin of speckle boiling and its effects in stellar speckle interferometry”. In: *Journal of Optics* 13.5, pp. 263–271.

- Ruane, G. et al. (2018). “Review of high-contrast imaging systems for current and future ground- and space-based telescopes I: coronagraph design methods and optical performance metrics”. In: *Proc. SPIE*. Vol. 10698.
- Schlaerth, J. A. et al. (2012). “The Status of Music: A Multicolor Sub/millimeter MKID Instrument”. In: *Journal of Low Temperature Physics* 167.3, pp. 347–353. ISSN: 1573-7357.
- Seager, Sara and Drake Deming (2010). “Exoplanet Atmospheres”. In: *Annual Review of Astronomy and Astrophysics* 48.1, pp. 631–672.
- Skidmore, Warren et al. (2015). “Thirty Meter Telescope Detailed Science Case: 2015”. In: *Research in Astronomy and Astrophysics* 15.12, pp. 1945–2140.
- Soummer, Rémi and André Ferrari (2007). “The Strehl Ratio in Adaptive Optics Images: Statistics and Estimation”. In: *ApJ* 663.1, pp. L49–L52.
- Soummer, Rémi, Laurent Pueyo, and James Larkin (2012). “Detection and Characterization of Exoplanets and Disks Using Projections on Karhunen-Loève Eigenimages”. In: *ApJ* 755.2, L28, p. L28.
- Soummer, Remi et al. (2007). “Speckle Noise and Dynamic Range in Coronagraphic Images”. In: *ApJ* 669.1, pp. 642–656.
- Sparks, William B. and Holland C. Ford (2002). “Imaging Spectroscopy for Extrasolar Planet Detection”. In: *ApJ* 578.1, pp. 543–564.
- Stangalini, M. et al. (2018). “Recurrence Quantification Analysis as a Post-processing Technique in Adaptive Optics High-contrast Imaging”. In: *ApJ* 868.1, p. 6.
- Stangalini, Marco et al. (2016). “Speckle statistics in adaptive optics images at visible wavelengths”. In: *Adaptive Optics Systems V*. Vol. 9909. SPIE, 99097P.
- Strader, Matthew (2016). “Digital Readout for Microwave Kinetic Inductance Detectors and Applications in High Time Resolution Astronomy”. PhD thesis. University of California Santa Barbara.
- Szypryt, P. et al. (2017). “Large-format platinum silicide microwave kinetic inductance detectors for optical to near-IR astronomy”. In: *Opt. Express* 25.21, pp. 25894–25909.
- Tanner Jr., Wilson P. and John A. Swets (1954). “A decision-making theory of visual detection”. In: *American Psychological Association* 61 (6), pp. 401–409.
- Tuckerman, David B et al. (2016). “Flexible superconducting Nb transmission lines on thin film polyimide for quantum computing applications”. In: *Superconductor Science and Technology* 29.8, p. 084007.
- Ulbricht, Gerhard et al. (2015). “Highly Multiplexible Thermal Kinetic Inductance Detectors for X-Ray Imaging Spectroscopy”. In: *Appl. Phys. Lett.* 106, p. 251103. DOI: 10.1063/1.4923096.
- Verma, V. B. et al. (2014). “A four-pixel single-photon pulse-position array fabricated from WSi superconducting nanowire single-photon detectors”. In: *Applied Physics Letters* 104.5, 051115.
- Wadell, Brian C (1991). *Transmission Line Design Handbook*. Artech House Microwave Library. 685 Canton St, Norwood, MA 02062: Artech House Inc. ISBN: 0890064369.

- Walter, A. B. et al. (2018). “Laminated NbTi-on-Kapton Microstrip Cables for Flexible Sub-Kelvin RF Electronics”. In: *IEEE Transactions on Applied Superconductivity* 28.1, pp. 1–5. ISSN: 1051-8223.
- Walter, Alex B. et al. (2018). “MEC: the MKID exoplanet camera for high contrast astronomy at Subaru (Conference Presentation)”. In: *Ground-based and Airborne Instrumentation for Astronomy VII*. Vol. 10702. SPIE, p. 107020V.
- Walter, Alex B. et al. (2019). “Stochastic Speckle Discrimination with Time-Tagged Photon Lists: Digging Below the Speckle Noise Floor”. In: *PASP*, arXiv:1906.03354. arXiv: 1906.03354 [astro-ph.IM].
- Weers, H.J. van et al. (2013). “Niobium flex cable for low temperature high density interconnects”. In: *Cryogenics* 55-56, pp. 1–4. ISSN: 0011-2275.
- Woodcraft, Adam L. et al. (2010). “Thermal conductance at millikelvin temperatures of woven ribbon cable with phosphor-bronze clad superconducting wires”. In: *Cryogenics* 50.8, pp. 465–468. ISSN: 0011-2275.
- Yung, C. S. and B. H. Moeckly (2011). “Magnesium Diboride Flexible Flat Cables for Cryogenic Electronics”. In: *IEEE Transactions on Applied Superconductivity* 21.3, pp. 107–110. ISSN: 1051-8223. DOI: 10.1109/TASC.2010.2080655.