

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

The MKID Camera for PICTURE-C and Photon Counting Stochastic Speckle Discrimination

Permalink

<https://escholarship.org/uc/item/87w8q68w>

Author

Bockstiegel, Clinton

Publication Date

2019

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

The MKID Camera for PICTURE-C and Photon Counting Stochastic Speckle Discrimination

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

Doctor of Philosophy
in
Physics

by

Clinton David Bockstiegel

Committee in charge:

Professor Benjamin A. Mazin, Chair
Professor Timothy D. Brandt
Professor Joseph F. Hennawi

September 2019

The Dissertation of Clinton David Bockstiegel is approved.

Professor Timothy D. Brandt

Professor Joseph F. Hennawi

Professor Benjamin A. Mazin, Committee Chair

August 2019

The MKID Camera for PICTURE-C and Photon Counting Stochastic Speckle
Discrimination

Copyright © 2019

by

Clinton David Bockstiegel

For Giulia and Adele.

Acknowledgements

This PhD would not have been possible without the support of many people, both in and out of the academic setting. I would like to thank everyone who helped me, no matter how much. In particular I want to make the following acknowledgements.

Professional

First I want to thank my thesis advisor, Ben Mazin, without whom none of this MKID work would have been possible. I am very grateful for the opportunities that he has created to simultaneously work in low temperature detectors and in high contrast astronomy. Both fields are cool in their own right, but working on both at the same time has taught me a combination of skills I don't think I could have found anywhere else.

To all of my labmates, past and present, thank you. I've enjoyed your company, and the conversations and trips I've had with you made working the the lab fun.

Thanks to Tim Brandt. It was his idea to use inter-photon arrival times to improve on existing SSD work. His guidance was absolutely crucial to the materialization of the project. He also showed me that statistics can be fun and not just scary. Thanks also to Alex Walter; we worked very closely on the analysis, code development, and understanding of speckles in general. I enjoyed working with both of you tremendously.

Thanks to Jiansong Gao. The tunable coupler was his idea, and he helped enormously with the analysis. Thanks also to Yiwen Wang, Mike Vissers, Li Fan Wei, Saptarshi Chaudhuri and Hannes Hubmayr for supporting the project with measurements, device fabrication, modeling and general discussion.

The staff in the Physics Department do a lot of work behind the scenes in helping our projects come together; thank you all.

Thanks to the Palomar day crew, support astronomers, and Monastary staff. You

made us feel welcome there, always ready to lend helping hands, anecdotes, and friendly smiles as we struggled to get light into DARKNESS.

My salary for the PICTURE-C work came from NASA grant NNX15AG23G.

Personal

My family, immediate and extended, has supported me, unwaveringly, for all my life. I am where I am today because of them. Thank you all, and especially Mom, Pop, and Evan; I hope I've made you proud.

Finally, to my wife, Giulia. Thank you for being my best friend and fiercest advocate. I can't wait to start our next chapter together.

Curriculum Vitæ

Clinton David Bockstiegel

Education

2019 (expected)	Ph.D. in Physics, University of California, Santa Barbara.
2017	M.A. in Physics, University of California, Santa Barbara.
2011	B.A. in Physics, University of Colorado, Boulder, CO.

Selected Publications

- A. B. Walter, C. Bockstiegel, T. Brandt, B. A. Mazin, “Stochastic Speckle Discrimination with Time-Tagged Photon Lists: Digging Below the Speckle Noise Floor”, Publications of the Astronomical Society of the Pacific (2019) **Submitted**, [[hrefhttps://arxiv.org/abs/1906.03354](https://arxiv.org/abs/1906.03354)arXiv:1906.03354]
- C. Bockstiegel, Y. Wang, M. R. Vissers, L. F. Wei, S. Chaudhuri, J. Hubmayr, J. Gao, “A Tunable Coupler for Superconducting Microwave Resonators Using a Non-linear Kinetic Inductance Transmission Line”, Applied Physics Letters 108, 222604 (2016)
- C. Bockstiegel, J. Gao, M. R. Vissers, M. Sandberg, S. Chaudhuri, A. Sanders, L. R. Vale, K. D. Irwin, D. P. Pappas, “Development of a Broadband NbTiN Traveling Wave Parametric Amplifier for MKID Readout”, Journal of Low Temperature Physics 176, 476–482 (2014)

Abstract

The MKID Camera for PICTURE-C and Photon Counting Stochastic Speckle
Discrimination

by

Clinton David Bockstiegel

High contrast imaging is a powerful technique for studying exoplanets and circumstellar disks. Ground based observations rely on adaptive optics to correct wavefront aberrations and starlight suppression techniques such as coronagraphy to reach contrast levels suitable for detecting substellar companions. Sensitivity is limited by coherent interference patterns, called speckles, in the science focal plane resulting from imperfect wavefront correction. As adaptive optics systems and methods for mitigating speckles from non-common path errors have matured in recent years, the contrast ratios remain limited by atmospheric speckles with lifetimes < 1 sec.

Microwave kinetic inductance detectors (MKIDs) are fast photon counting detectors that can simultaneously measure a photon's arrival time, its energy, and its x-y location within the focal plane. Armed with these capabilities, MKID science cameras promise to substantially improve the signal to noise of exoplanet and circumstellar disk detections by sampling speckle intensities quickly with photon limited noise.

In this thesis I present the development of an MKID-based science camera for a NASA funded balloon-borne observatory called PICTURE-C (the Planetary Imaging Concept Testbed Using a Recoverable Experiment-Coronagraph). The balloon will fly at an altitude of ≈ 40 km, placing it above most of the atmosphere and thereby removing atmospheric aberrations, with the goals of characterizing nearby debris disks and demonstrating the operation of MKIDs in a near space environment.

I will then discuss photon counting stochastic speckle discrimination, an analysis technique in which the variability of inter-photon arrival times is used to separate residual stellar speckles from companion light. This technique can improve the signal to noise of planet detections over that of long exposures where many speckle realizations are averaged over.

Finally I present a tunable coupler that can be used to adjust the effective coupling strength between a superconducting microresonator and its microwave feedline in situ for low temperature measurements. This scheme may find applications in astronomy for optimizing the readout of MKIDs or in quantum information science for optimizing qubit lifetimes.

Contents

Curriculum Vitae	vii
Abstract	viii
1 Introduction	1
1.1 Overview	1
1.2 Detecting Exoplanets	3
1.3 Scientific Motivation for Direct Imaging	6
1.4 Tools for Direct Imaging	12
1.5 Speckles	22
1.6 MKIDs	27
2 PICTURE-C	41
2.1 Project Overview	41
2.2 The Science of PICTURE-C: Imaging Nearby Debris Disks	43
2.3 Work at UMass Lowell: Gondola and Optics	45
2.4 Design of the PICTURE-C MKID Camera	47
2.5 ADR	70
2.6 Testing the PICTURE-C Cryostat	73
2.7 PICTURE-C Cryostat Electronics	75
2.8 Readout: Generation 3	76
3 Stochastic Speckle Discrimination	80
3.1 Speckles and the Modified Rician Distribution	80
3.2 Poisson Limited Modified Rician with Binned Lightcurves	86
3.3 Likelihood Function of Poisson Limited Modified Rician with No Companion	88
3.4 Adding a Companion to the Starlight in Binned Lightcurves	88
3.5 Bin-Free SSD	91
3.6 Finding the Maximum Likelihood	99
3.7 To Marginalize or Not to Marginalize?	100
3.8 Signal Detection Theory Overview	102

3.9	Mock Photon Lists	103
3.10	Binned vs. Bin-free SSD	108
3.11	Signal to Noise Ratio	113
3.12	Effect of Coronagraph on SSD	119
3.13	When to Use SSD	120
4	Tunable Coupler	123
4.1	Resonator Coupling	123
4.2	Derivation of 1-port Resonator Scattering Parameters with Tunable Coupler	124
4.3	Nonlinear Transmission Line	137
4.4	Conclusions	145
	Bibliography	148

Chapter 1

Introduction

1.1 Overview

This thesis is concerned with applications of Microwave Kinetic Inductance Detectors (MKIDs). As we will see, MKIDs have potential to be a very powerful tool for high contrast astronomy at optical and near-IR wavelengths, which is especially interesting for exoplanet science. I spent most of my time at UCSB working on three instruments, DARKNESS, MEC, and PICTURE-C, cameras that use MKIDs to simultaneously record images and discriminate the colors of light within their passband.

In this introductory chapter, I discuss the science motivation behind direct imaging, give a brief description of the optical tools required for it, and describe how our optical MKIDs work. This gives context for the work that I have done both on instrumentation and data analysis techniques, which is described in detail in the three subsequent chapters.

Chapter 2 presents a detailed description of the design and construction of the PICTURE-C cryostat. PICTURE-C is an MKID camera that will fly on a balloon, whose science goal is to image debris disks. Planet formation is a hot topic in astron-

omy, and we can gather a great deal of information about planetary systems and their histories by looking at disks, as they are considered to be signposts of exoplanets. The PICTURE-C cryostat builds on DARKNESS [1], a 10,000 pixel MKID camera. Starting in 2014, I was deeply involved in the construction work of the DARKNESS cryostat that led to its commissioning in 2016 at Palomar, and I led the setup of the instrument for two later observing runs. The experience that I gained over those three years was valuable for informing the design of PICTURE-C, making the cryostat easier to use leading up to the balloon flight.

In Chapter 3 I describe an analysis technique called stochastic speckle discrimination (SSD). Together with Tim Brandt and Alex Walter, we improved upon previous SSD work to take full advantage of the photon counting capabilities of MKIDs [2]. Simulations show clear improvements in signal to noise of planet detections and improvements in the contrast floor, and the technique is complementary to standard post-processing algorithms.

In Chapter four, I present a project carried out with collaborators at NIST. We designed, built, and tested a tunable coupler for optimizing MKID readout [3]. The ideal condition for maximizing signal to noise with MKIDs is to match their internal and coupling quality factors. This is made difficult by small errors during device fabrication. Our tunable coupler demonstrates a way to vary the coupling quality factor and achieve a matched condition after the device fabrication. This scheme may be able to improve MKID sensitivity, and quantum information science could also benefit by increasing the fidelity of superconducting qubits.




Figure 1.1: The accessible regions of the mass vs. period parameter space are highly dependent upon the detection method. For example, directly imaged planets are most easily accessed at large separations and when they are bright, a consequence of being young and hot. Image downloaded from exoplanets.ipac.caltech.edu.

1.2 Detecting Exoplanets

The field of exoplanet science is one of the fastest growing areas of astronomy today. The advent of various technologies, observatories, and data reduction techniques has produced an explosion in the number of exoplanet detections over the past two decades and characterizing attributes such as mass and host star separation has become routine. With a catalog of over 4,000 confirmed exoplanets [4], the community is well into an era of exploring the physics of exoplanets by examining statistical trends.

Astronomers can hunt for exoplanets in several ways, and while a number of wonderful textbooks have been written, including [5] and [6], we will attempt to summarize here in a few words the rough idea of how the most common methods work.

- **Transits:** This is the most prolific method for finding exoplanets, largely because of the Kepler Space Telescope, logging over 3000 discoveries. The brightness of a star is monitored for extended periods of time to obtain a light curve. Periodic dips in the light curve can indicate the presence of substellar companions. Transits are only possible if the system has a nearly edge-on inclination to our line of sight, and shorter period planets are easier to detect because they have more transits per unit time. In addition to the orbital period, the planet radius can be measured from the transit depth, as the amount of starlight blocked by the planet depends on its radius.
- **Radial Velocity (RV):** As planets move in their orbits, their stars move as well. A high resolution spectrograph records periodic Doppler shifts in stellar absorption or emission lines. Unless the orbital inclination i is known, RV only provides $M \sin i$ and orbital parameters like the period.
- **Microlensing:** When a foreground star moves across the same line of sight as some background object it can gravitationally lens the light from the more distant object, producing a light curve with a characteristic shape. If the foreground star has a planet, the planet can affect the shape of the light curve. Similar to transits, watching for microlensing signatures requires long observations of an object's intensity. Microlensing can measure a planet's mass and sometimes its orbit.
- **Direct Imaging:** It is possible to essentially take a photograph of a planet, although the technical challenges make this method difficult. Ground based observations

require adaptive optics to correct the wavefront, and careful removal of contamination from the host star is necessary for detecting companions with much lower intensities than the star. This is the mode of observation that we wish to employ with our MKID cameras. The inner working angles of high contrast systems place lower bounds on the angular separation of detectable planets, restricting the parameter space to planets far away from their star, and contrast limitations place lower bounds on the mass of detectable planets.

- **Timing Variations:** The presence of a planet can affect the orbital periods of other companions. Small, periodic changes in a companion's orbital period can be explained by otherwise unseen planets. Additionally, pulsar timing variations can imply slight radial velocity variations, and this was actually the first method used to detect an exoplanet in 1992 [7].
- **Orbital Brightness Modulation:** A planet orbiting close to its star can exert large tidal effects on the host. Changes in the shape of the star may be measurable to an observer because the visible surface area would be modulated by the planet's orbital position. Tidal forces scale with the inverse distance cubed, so this method is highly selective to close-in planets.
- **Astrometry:** The movements of a star when a substellar companion orbits can show up as wobbles in the star's position. Detecting wobbles with very precise position measurements can indicate the presence of a companion. Astrometry can measure a planet's mass and its orbital parameters.

This list hints at some of the technical challenges associated with observing exoplanets. The challenges are different for each method, leading to the clustering of planets of the same detection method in Figure 1.1. This is known as a selection effect, and

careful attention must be given to accounting for biases when drawing conclusions about exoplanet demographics. Calibrating for selection effects is made easier when different detection methods overlap in parameter space. Furthermore, different methods are complementary, so planets that might not be picked up by some method could be discovered by a different method. In the next sections we will discuss the goals of direct imaging and some of the tools with which smaller, closer-in planets are being pursued.

1.3 Scientific Motivation for Direct Imaging

1.3.1 Detecting Nearby Exoplanets in the Habitable Zone

Exoplanet science is growing rapidly today precisely because of the large gaps that remain in our knowledge. The Kepler mission in particular has made great strides in quantifying properties about the demographics and occurrence rates of planets, simply because of the sheer number of detections made possible by the observatory. However, as described in the previous section, transits have a low detection rate because of the inclination and orbital period constraints, meaning that only statistical conclusions can be drawn about overall occurrence rates of close-in planets. A star that does not have a transit signal in its light curve may still host planets.

RV measurements have yielded several hundred exoplanet discoveries. The detection method produces a somewhat more complete census of stars hosting planets than the transit method because a much wider range of orbital inclinations can lead to detections. But the parameter space is also constrained, this time to mostly planets of higher mass and longer orbital periods. For the most part, the populations of known RV and transiting exoplanets don't overlap in the mass vs. period parameter space.

Direct imaging has begun to fill in the region of parameter space encompassing mas-

sive planets at large separations, and improvements in technology coupled with upcoming thirty meter class telescopes will dramatically expand the reach of direct imaging to lower mass and smaller separations, allowing us to probe systems in the stellar neighborhood and see into their habitable zones. Further, because of the nature of direct imaging, null detections are much more likely to be the result of an actual absence of planets orbiting a star. Indeed, this helped lead to the revelation that massive planets (5-13 M_{Jup}) at large separations (30-300 AU) are rare ($\sim 0.006/\text{star}$) [8]. Direct imaging of nearby stars will be crucial to approaching a complete census of exoplanets in our neighborhood in the coming years.

One of the goals of direct imaging (and exoplanet science in general) is to find rocky planets in the habitable zone. As we search for life beyond our own planet, these seem like an obvious place to check because of the similarities to Earth. Liquid water in particular is a strong biomarker, and that coupled with molecular species out of chemical equilibrium would be a very strong indication that life may be present. Observing an Earth twin is a tall order because of the extreme contrast ratios ($\sim 10^{-10}$) and tiny angular separations ($\sim 0.1''$ at 10 pc). It may be possible to detect planets in the habitable zone of nearby M-dwarfs from the ground because the contrast and IWA requirements are more favorable, but it is still unclear whether this is the case and space-based observatories may be required. HabEx¹ and LUVOIR² are examples of proposed NASA missions that would find nearby habitable exoplanets. However, these missions, should they materialize, are far off in the future. In the meantime, ground-based systems are necessary to develop technologies providing higher contrast, smaller inner working angles, and better exoplanet characterization.

¹<https://www.jpl.nasa.gov/habex/>

²<https://asd.gsfc.nasa.gov/luvoir/>

1.3.2 Planet Formation

In addition to taking a census of exoplanets that are nearby and habitable, it is also crucial to study their atmospheres and compositions to inform planet formation models. Molecular compositions of atmospheres can be inferred from spectra obtained at every point in an image with integral field spectrographs, such that broad features of an exoplanet can easily be isolated from other objects in the same image, such as the host star.

Transit spectroscopy, whereby the transit depth is measured as a function of wavelength and used to infer spectral features of a planet's atmosphere, is the only way to probe its composition other than direct imaging. However, because it relies on starlight passing through the planet's atmosphere, stellar photon noise constrains the data quality. Direct imaging spectroscopy avoids this confusion because the planet light is isolated from that of the star when an image is formed in the science camera, yielding higher signal to noise [9]. Measuring broad spectral features of exoplanets from 0.7 - 1.4 μm is one of the main science goals driving the development of our MKID instruments.

More broadly, near-IR observations can track thermal emissions from young, massive planets. As planets form, gravitational energy is converted into thermal energy implying that young massive planets are warm and emit strongly at IR wavelengths. Capturing images and spectra of planets shortly after their formation provides clues for how planets form, like whether core accretion or gravitational instability is the preferred pathway.

1.3.3 Debris Disks

High contrast imaging systems are also well suited to observe circumstellar disks, which can be comprised of both gas and dust. The compositions and distributions of disk material change dramatically as the system ages, starting with gas-rich protoplanetary

disks. Material in such a disk is largely blown out of the system by stellar radiation in the first few Myr, but some also accretes into planets, planetissimals, and smaller dusty grains.

Older disks, called debris disks, are made up largely of residual dusty material left over from the protoplanetary progenitor and contain only trace amounts of gas. While small grains are still removed by Poynting-Robertson drag, stellar wind, and radiation pressure, they are continually replenished by a collisional cascade in which larger particles break up into smaller particles [10]. Hence, debris disks can be relatively long lived, with some as old as several Gyr [11]. Our own Kuiper and asteroid belts are an example.

Disks are interesting to exoplanet observers for a few reasons. Their presence has been shown to correlate with the presence of giant planets [12], suggesting that they are a kind of signpost for potentially undiscovered planets. Furthermore, disks can interact gravitationally with planets. This might lead to enhanced dust production through stirring, as well as interesting morphologies like warps. These interactions, together with coupled histories can inform models of planet formation and evolution, as well as give insights into the compositions of planetary systems.

Finally, debris disks often appear with two separate components. They are typically identified by two peaks in thermal IR emission as shown in Figure 1.2, implying that the two components have different temperatures and therefore separations from their stars. Long term secular perturbations or material accretion by planets may provide mechanisms for clearing gaps in a disk leaving more than one belt of material [10]. Some famous examples of two-belt systems exhibiting connections between debris disks and planets are HR8799, β Pictoris, HD 95086, and 51 Eridani, presented in Figure 1.3. We should note that debris disk gaps do not necessarily indicate the presence of planets, and indeed deep direct imaging searches for planets around Vega have yielded null results [13].

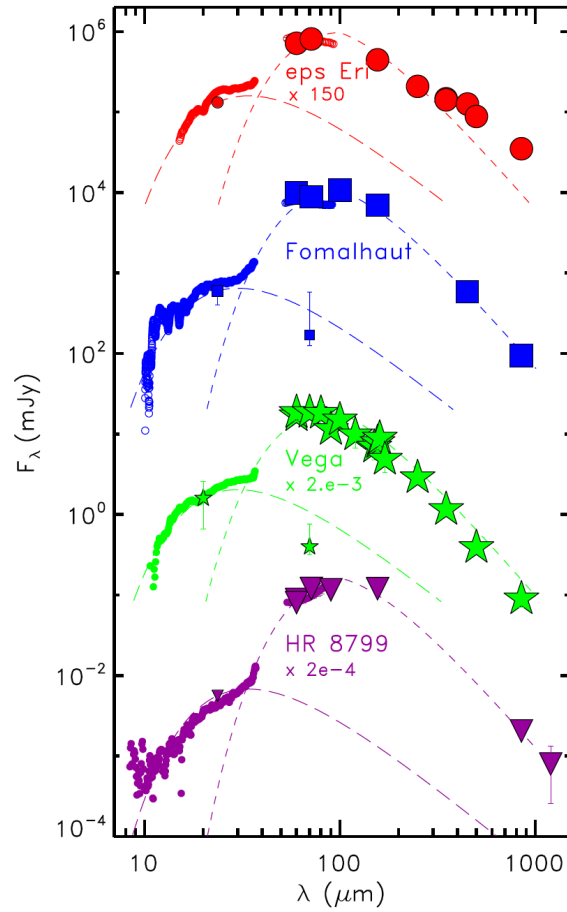


Figure 1.2: Excess spectral energy distributions from ϵ Eridani, Fomalhaut, Vega, and HR8799 indicating the presence of debris disks with two components of different temperatures. The warmer components have temperatures of 150-170 K while the cooler ones have temperature of 45-50 K. Planet formation or secular perturbations can sometimes account for the gaps. Reproduced from [14].




Figure 1.3: Images of systems with imaged planets and debris disks. HR8799 has four planets $5\text{--}10 M_{\text{Jup}}$ in mass, β Pictoris's disk is nearly edge-on and has a $13 M_{\text{Jup}}$ planet whose presence was predicted from the warp in the disk, HD95086 has a $5 M_{\text{Jup}}$ planet, and 51 Eridani has a $2 M_{\text{Jup}}$ planet. Reproduced from [8], with original images taken from [15], [16], [17], and [18].

1.4 Tools for Direct Imaging

Directly imaging exoplanets is challenging because of the large contrast ratios and the small angular separations between planets and their stars. Even relatively easy targets, such as the Jovian planets of HR8799, which have angular separations ranging from about 0.5 to 1.7 arcsec and contrast ratios of $\sim 10^{-5}$ in H band (1630 nm) [19] illustrate the extreme requirements of high contrast imaging systems. The first step to observing targets like these from the ground is correcting for atmospheric turbulence with adaptive optics. Next comes suppressing light from the host star, which can be done in a few ways. For systems like HR8799 it is sufficient to subtract an image of a reference star of the same stellar type from the image of the target, but the technique’s usefulness is limited because of noise in the images. A more aggressive approach is to use a coronagraph, which actually removes starlight before it even reaches the detector, and together with adaptive optics it is a key ingredient to imaging fainter exoplanets. Finally, the last key piece of hardware is a sensitive camera with good noise properties, preferably with wavelength sensitivity. MKIDs are an example of such a camera.

Here we will go over the basics of adaptive optics and coronagraphy, and then describe the detectors our lab builds for imaging exoplanets.

1.4.1 Adaptive Optics

In the absence of optical aberrations, telescopes and their cameras should operate at the diffraction limit, where the telescope diameter D and the operating wavelength λ set the angular resolution, defined as the “width” of the point spread function (PSF), λ/D . As light propagates through the atmosphere, the wavefronts become distorted because of inhomogeneities in temperature and to a lesser extent pressure. The Fried parameter, r_0 , is the aperture size over which the wavefront phase varies by ≈ 1 radian. r_0 is highly

dependent on geographic location, with larger values corresponding to better observatory sites with less turbulence and higher altitudes.

If D is smaller than r_0 , the optical system will be diffraction limited with a PSF defined by the shape of the telescope pupil, while if $D > r_0$ it will be seeing limited, where the PSF evolves along with atmospheric turbulence and long exposures exhibit a seeing halo. Seeing limited observations consequently have angular resolutions of λ/r_0 . For context, $r_0 \approx 10$ cm at Palomar for visible wavelengths, meaning that even with a mirror diameter of 5.1 meters (200 inches), the angular resolution without wavefront correction is no better than that of a 10 cm (4 inch) amateur telescope. In order to beat atmospheric seeing and approach diffraction limited resolution, we must introduce an adaptive optics (AO) system that corrects the aberrated wavefront.

The basic ingredients of an AO system are: a deformable mirror (DM) whose shape can be controlled in real time, a wavefront sensor (WFS) that measures aberrations in the wavefront, and a control system that crunches the data from the WFS to determine what corrections should be applied to the DM. Figure 1.4 illustrates this basic AO system layout where the corrected wavefront is directed to a science camera. Designs can of course be more complicated, for example the PALM-3000 AO system at Palomar has two deformable mirrors, one for low order wavefront correction and one for high order wavefront correction [20].

DMs are typically fabricated as a thin piece of glass attached to hundreds or thousands of micro-electromechanical actuators that can independently push or pull on the glass with maximum strokes on the order of a micron. The DM control system can apply arbitrary shapes to the mirror, with the aim of compensating for phase aberrations in the wavefront. “Extreme AO” is a phrase often employed to denote a system with thousands of DM actuators.

Several schemes for constructing WFSs exist, but here we will highlight the ones used




Figure 1.4: Adaptive optics systems correct wavefront distortions introduced by the atmosphere in real time. The wavefront sensor (WFS) measures the shape of the incoming wavefront and the control system uses that information to apply feedback to the deformable mirror (DM). In this classic AO layout, the light path to the science camera (labelled High-resolution Camera) is not identical to that of the WFS, which can lead to “non-common path errors” that create quasistatic speckles in the science image. Reproduced from lyot.org

by our MKID instruments. First, a Shack-Hartmann WFS that uses an array of tiny lenslets in the pupil plane to form spots of light on a detector. If the wavefront is tilted at the location of a particular lenslet, the spot formed there will be displaced from the optical axis by an amount proportional to the gradient of the wavefront. The combined measurements of the wavefront gradient across the entire pupil plane define the shape of the wavefront. PALM-3000 has a Shack-Hartmann WFS that we use while observing with our MKID instrument DARKNESS, and the PICTURE-C mission will also employ this type of WFS.




Figure 1.5: Shack-Hartmann wavefront sensors use microlens arrays in the pupil plane to image spots onto a detector. The displacements of the spots from the optical axis are functions of the local wavefront gradient, and taken altogether can be used to reconstruct the shape of the wavefront. Reproduced from <http://www.ctio.noao.edu/~atokovin/tutorial/part3/wfs.html>.

Pyramid WFSs use a glass pyramid with an extremely finely ground point located in the focal plane. Four images of the pupil appear behind the pyramid to be imaged by a detector, as shown in Figure 1.6. Aberrations in the wavefront are measured by comparing the intensity patterns in each of the pupils. For example, tip/tilt aberrations

move the PSF laterally away from the pyramid tip, which either increases or decreases the intensity of the pupils in each quadrant. The Subaru Coronagraphic Extreme AO (SCExAO) system, which we use for observations with MEC, uses a pyramid WFS.




Figure 1.6: The tip of a glass pyramid is placed in the focal plane, such that some fraction of the light is directed each of four locations on a detector. The detector is placed in the pupil plane, and the wavefront shape can be measured by comparing the four images of the pupil. Reproduced from <http://www.ctio.noao.edu/~atokovin/tutorial/part3/wfs.html>.

AO systems need to perform corrections continuously because wind and turbulence in the atmosphere constantly change the wavefront distortions. The temporal evolution is well approximated by the frozen flow model, in which one or more layers of fixed aberration maps move across the telescope beam, each with its own amplitude, wind speed V_0 , and direction [21]. For a fixed location in the beam, the wavefront phase will vary by 1 radian RMS with a coherence time $\tau_{\text{wavefront}} = 0.31V_0/r_0$. This timescale sets the rate at which AO systems must update the wavefront corrections.

Wavefront aberrations are also wavelength dependent. This can be understood intuitively by considering that smaller wavelengths will be sensitive to smaller turbulent perturbations in the atmosphere. Fried first defined r_0 in 1965 [22], showing that it scales with $\lambda^{6/5}$. Wavefront correction is easier at longer wavelengths, where there are fewer independent turbulent cells within the telescope beam.

The performance of an AO system is often described by the Strehl Ratio (SR), the

ratio of peak intensity of the aberrated image to the maximum intensity achievable if the system were diffraction limited. $SR = 1$ corresponds to perfect AO correction, with performance degrading as SR approaches zero. The Marechal approximation relates SR to the standard deviation σ of the wavefront phase across the pupil,

$$SR \approx e^{-\sigma^2}. \quad (1.1)$$

Visible wavelength correction is significantly more challenging than in the near-IR. The number of actuators n needed to achieve Nyquist sampling and a given level of AO correction scales with the number of r_0 cells within the telescope beam, so that $n \propto r_0^{-2} \propto \lambda^{-12/5}$. Further, the correction rate must also increase at shorter wavelengths to keep up with the temporal evolution as $\lambda^{-6/5}$. Modern Extreme AO systems routinely achieve $SR \sim 85\%$ in H-band (centered at $\lambda = 1630$ nm) but only $SR \sim 27\%$ in the visible ($\lambda = 500$ nm) [9]. However, advances in WFS and DM technology will allow AO systems to tackle visible wavelengths in the coming years. For example, MagAO-X promises to achieve $SR > 70\%$ at 656 nm upon deployment [23].

As shown in Figure 1.7, an AO corrected focal plane image consists of the diffraction limited PSF and a time varying speckle field from wavefront correction residuals. Speckles set the noise floor and limit the contrast throughout an image. Speckle noise suppression is an active area of research with a number of promising avenues for improvement. In Chapter 3 we will discuss one such method for post-processing data with time-tagged photons that takes advantage of the known functional form of speckles' temporal intensity distribution.

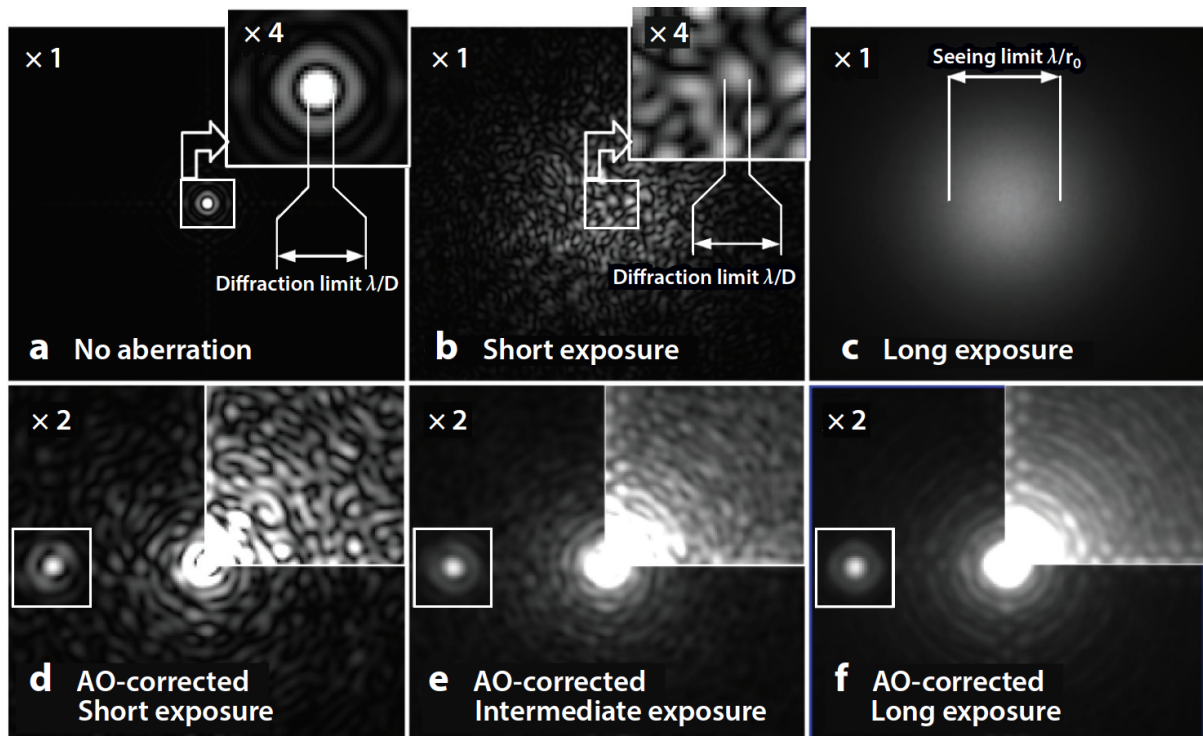


Figure 1.7: Simulated PSFs showing the result of AO correction. (a) Diffraction limited PSF in the case of no wavefront aberrations with size λ/D . (b) A short exposure PSF with aberrations and no AO correction. The granular spots are referred to as speckles, each with a size λ/D . (c) Long exposure, seeing-limited halo, whose width is λ/r_0 . (d) AO corrected short exposure, where speckles are still visible but most of the light is concentrated in the diffraction limited PSF. (e) and (f) Intermediate and long exposure AO corrected images where the intensity has been integrated over many speckle lifetimes. Reproduced from [9].

1.4.2 Coronagraphs

While AO systems dramatically improve images by taking starlight away from the seeing halo where planets and circumstellar disks might be and putting it back into the point spread function, there is often still too much stellar light obscuring companions. As such, other tools must sometimes be used to improve the contrast in an image. Coronagraphs are a piece of optical hardware that suppresses starlight while observations are ongoing³.

To understand how coronagraphs work, it is useful to first consider the Fourier optics approximation, where pupil planes are the Fourier transforms of image planes. Lenses or curved mirrors determine the locations of these planes, and one can in principle go back and forth between focal and pupil planes many times. The primary mirror of a telescope is a pupil plane, and its curvature brings the light it collects to an image in the focal plane. An idealized, uniformly illuminated circular telescope mirror with constant phase has a Fourier transform equal to an Airy pattern, with a central core and diffraction rings as shown in Figure 1.8.

A coronagraph is a combination of one or more filters in one or more focal and pupil planes. The earliest successful coronagraph was developed by Bernard Lyot, an opaque occulting spot placed at the center of the Airy core with a radius equal to or larger than the core [25]. Diffraction from the spot scatters light away from the optical axis, so a circular Lyot stop often follows in the next pupil plane to block light at large radii and pass light at small radii. See Figure 1.9 for an illustration. After re-imaging to another focal plane, the starlight is suppressed to about 1% its original intensity. Off-axis light coming from other objects in the field, say from an exoplanet, is largely unaffected in

³PSF subtraction is another way to remove starlight from an image, where a reference PSF is built up and then subtracted from the science data. This was the first method used to detect the planets orbiting HR8799 [24].

this process.

Coronagraph designs have come a long way since the Lyot coronagraph, with more complicated designs using softening of the edges, known as apodization, and phase manipulation [26]. In the future we hope to employ vector vortex coronagraphs (VVCs) [27] in observations with MKID instruments.




Figure 1.8: (left) A uniformly illuminated circular pupil has a fourier transform equal to an Airy pattern (right). The width of the Airy pattern is defined by the pupil diameter D and the wavelength λ . Reproduced from wikipedia.org.




Figure 1.9: Illustration of the Lyot coronagraph and Lyot stop. Black indicates areas where light is blocked. At left is the focal plane mask, a spot blocking all light within a radius on the order of a few λ/D . At right is the Lyot stop, the pupil plane mask, which blocks light diffracted to large radii by the focal plane mask. Also shown are a central obscuration and straight lines, which mimic the area obscured by a secondary mirror and its support struts, referred to as spiders. This is the kind of coronagraph we have employed with DARKNESS at Palomar.

1.4.3 Integral Field Spectrographs: DARKNESS, MEC, and PICTURE-C

Integral field spectrographs (IFSs) are a class of instrument that combine traditional imaging with wavelength sensitivity. They are constructed such that each spatial resolution element, or pixel, can measure the spectrum of the light hitting the pixel. IFS have become popular for high contrast imaging because their “data cubes” with two spatial dimensions and one wavelength dimension naturally enable measurements of a substellar companion’s spectrum and a differential imaging technique exploiting the chromatic behavior of speckles, discussed below. At UCSB we have received funding to build three IFSs using MKIDs as the detectors. In later sections will describe how we construct MKID arrays and the physics behind their wavelength sensitivity.

The first instrument, the DARK speckle Near-IR Energy resolving Superconducting Spectrophotometer [1] (DARKNESS), was commissioned in 2016 at Palomar, and served as a testbed for the development of a number of hardware, software, and device fabrication techniques. It has a 10,000 pixel MKID array in a liquid helium cooled cryostat with an adiabatic demagnetization refrigerator cooling the detectors to a temperature of 100 mK.

The MKID Exoplanet Camera (MEC) is the second instrument, and was deployed to the Subaru Telescope in 2018 with a 20,000 pixel detector. Using a closed-cycle helium refrigerator instead of liquid cryogenics, it borrowed many component designs from DARKNESS and included some improvements so that it could be operated remotely. With a larger telescope mirror and a better location, MEC promises to outperform DARKNESS in imaging young self-luminous Jovian planets. The wavelength range of DARKNESS and MEC is 800-1400 nm.

The MKID camera for the Planetary Imaging Concept Testbed Using a Recoverable

Experiment-Coronagraph (PICTURE-C) is our third instrument, and will be flown on a balloon observatory to image debris disks and possibly exoplanets. The high altitude platform circumvents many of the challenges of high contrast astronomy coming from atmospheric effects, but with a small primary mirror its angular resolution will be lower than either DARKNESS or MEC. It will demonstrate the operation of MKIDs in a space-like environment, and provide valuable feedback for using MKIDs on future space missions. Its 10,000 pixel array will cover 540-660 nm.

Photographs of all three instruments are shown in Figure 1.10.




Figure 1.10: The three MKID instruments constructed at UCSB. Left to right, DARKNESS, MEC, and PICTURE-C.

1.5 Speckles

We turn now to the noise that limits high contrast observations. Even after correcting for atmospheric turbulence with the AO system, aberrations persist giving rise to a grainy interference pattern in the focal plane of the science camera. The grains, or “speckles” as they are known, can masquerade as substellar companions because they have the

same characteristic size, λ/D . Residuals from the AO system leave us with atmospheric speckles that change with timescales that depend on wind speed V_0 , telescope diameter D , and the Fried parameter r_0 , typically on the order of tens or hundreds of milliseconds. As mentioned in Section 1.4.1, atmospheric turbulence can be described by a frozen flow model, meaning that speckles are both spatially and temporally correlated. Speckles can also be produced by aberrations in the instrument’s optics that go undetected by the WFS, referred to as non-common path errors with variability timescales ranging from minutes to hours. Known as quasistatic speckles, their temporal evolution is typically ascribed to telescope movements or changes in temperature leading to flexure. Residuals in the shape of the DM can also produce speckles, but with a timescale set by the update rate of the DM.

Various approaches have been developed for mitigating speckles. Here we will briefly describe some of the approaches that we hope to employ with our MKID instruments.

1.5.1 Speckle Statistics

Astronomers realized in the 1990’s that with fast detectors like avalanche photodiodes and adaptive optics that the temporal variability of speckles in astronomical imaging can be exploited to improve the signal to noise of planet detections. The initial approach was rather simple, in which fast imaging would attempt to capture “dark speckles”, short-lived speckles with zero intensity arising from a lucky combination of atmospheric turbulence. Given enough time, dark speckles will appear throughout an image. Locations that never achieve a truly zero intensity could be identified as a planet [28, 29].

The statistics governing laser speckle intensities have been studied extensively over the course of many decades. J. W. Goodman published work in 1975 [30] describing this behavior, and in particular derived the laser speckle intensity distributions. He showed

that the presence of a deterministic diffraction pattern combined with a random one results in a distribution known as a modified Rician (MR).

V. F. Canales and M. P. Cagigal published a series of papers in 1999 in which they proposed that laser speckle statistics could be applied to high contrast astronomy [31]. They showed that the deterministic and random components of the MR distribution can be calculated from the mean and variance, and also how very low light levels affect the MR distribution [32, 33]. In 2001 they made an in-lab demonstration with a laser and a controllable phase screen to confirm that MR statistics correctly describe the distributions recorded with a set of images with different phase screens [34].

M. P. Fitzgerald et. al. used on-sky observations to show that MR statistics describe real astronomical data in 2006 using millisecond imaging. They showed that speckles can have multiple characteristic decorrelation times ranging from tens of milliseconds to about one second. They also used the MR fits to separate the raw measured intensities into deterministic and random components for every pixel in a set of images and produce maps of those components [35].

From 2008 to 2010 several papers were published by S. Gladysz and collaborators in which they coined the phrase “stochastic speckle discrimination” (SSD) as an approach to mitigating quasistatic speckles and demonstrate it with on-sky data. The idea is to fit intensity distributions with MR and use the ratio of the deterministic to random components as a discriminator for identifying companions in the field. The ratio determines the skewness of the MR distribution, and they show that skewness performs better than other metrics for improving the contrast [36, 37, 38].

Chapter 3 goes in to more detail about speckles and presents an improved SSD approach that takes advantage of the photon counting nature of MKIDs. By avoiding using traditional exposures and instead examining interphoton arrival times, it is possible to significantly improve the signal to noise of a planet embedded in a speckle field. We

anticipate that photon-counting SSD will be most effective for removing atmospheric speckles, while longer lived quasistatic speckles can still be removed by other means.

1.5.2 Speckle Nulling

Speckle nulling is an iterative approach to actively remove quasispeckles from the science camera's focal plane before and during observations [39]. The technique relies on using the DM to create artificial probe speckles in the focal plane. A sinusoidal pattern is applied to the DM with a handful of different amplitudes and phases to determine what combination is best for nulling a particular speckle. The process is repeated for all speckles within a user defined control region, and demonstrations with DARKNESS at Palomar have shown this to be a promising way to suppress contaminating starlight, as illustrated in Figure 1.11. Efforts are ongoing with a similar demonstration using MEC and SCExAO (Fruitwala et.al. in prep).




Figure 1.11: Demonstration of speckle nulling with the internal white light source of PALM-3000. The two images on the left show a before and after comparison of the intensity in DARKNESS' focal plane, and the plot shows the total intensity within the red control region decreasing over several minutes. Image courtesy of Neelay Fruitwala.

1.5.3 Spectral Differential Imaging

Spectral differential imaging uses image cubes from integral field spectrographs to exploit the chromatic behavior of speckles [40]. The position of speckles in the focal plane is wavelength dependent, such that the angular separation of a speckle from the center of the image scales linearly with λ . A speckle field at one wavelength, scaled appropriately, looks the same as a speckle field at a different wavelength. Using this fact, speckles can be subtracted from an image cube. The locations of companions on the other hand are not wavelength dependent, so the subtraction process will leave companion signals largely unscathed. An illustration of this process is shown in Figure 1.12.




Figure 1.12: Illustration showing spectral differential imaging technique. An integral field spectrograph produces image cubes with two spatial dimensions and one wavelength dimension. The chromatic behavior of speckles can be exploited to do self-subtraction, whereby the image from one wavelength slice is appropriately scaled and then subtracted from a different wavelength slice. Panel b reproduces panel a with a stretch, so that the speckles appear as they would at $1.58\mu\text{m}$. Panel c shows the image slice at $1.58\mu\text{m}$. Panel d is the difference between b and c, and shows the companion in the bottom right as a combination of complimentary spots. Reproduced from [41].

1.5.4 Angular Differential Imaging

Angular differential imaging relies on field rotation to separate quasistatic speckles from true astrophysical objects [42]. On an alt-az telescope, the field of view rotates with time, but quasistatic aberrations remain fixed in the focal plane. If the rotation rate is

fast enough the quasistatic speckles won't change much over the course of an observation. With a known modulation of the objects location and fixed contaminating speckles, it is possible to improve the signal to noise of the companion.

1.6 MKIDs

1.6.1 Principles and Design Strategy

At low temperatures, certain materials exhibit superconducting behavior in which electrons can pair up to form bosons called Cooper Pairs. The energy required to break a Cooper pair is known as the superconducting gap energy 2Δ , and scales with the critical transition temperature T_c ,

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

where k_B is the Boltzmann constant. The value of T_c is material dependent.

Currents in superconductors are comprised almost entirely of the movement of Cooper pairs. Because they are bosons, they move as a superfluid with zero DC resistance. However, under the influence of AC fields superconductors can have large values of reactance. This phenomenon can be understood by considering that Cooper pairs have mass, and therefore momentum when they move. They respond to a changing electric field slowly, gaining or losing momentum some amount of time after the field changes. Such a lag can be modelled as an inductance, whose value for a thin film⁴ at zero temperature is

$$L_s = \frac{\hbar}{\pi\Delta} \frac{\rho_n}{t}, \quad (1.3)$$

where ρ_n is the normal state resistivity of the material and t is the film thickness [43].

⁴We consider only thin film properties because our MKIDs are made from thin films. "Thin" simply means that the thickness t of the film is much smaller than the penetration depth of the fields.

We note that normal conductors also have kinetic inductance for the same reason: their charge carriers, electrons, have momentum when moving. However, the kinetic inductance in normal conductors is small compared to geometrical inductances, and is typically ignored at frequencies $\lesssim 100$ GHz.

For a superconductor with $T_c = 1$ K⁵, the energy required to break a Cooper pair is $2\Delta = 3 \times 10^{-4}$ eV, much smaller than the energy of a $\lambda = 1$ μm photon, about 1.2 eV. So when a photon of this energy is absorbed in the superconducting film, potentially thousands of Cooper pairs are broken, changing the surface impedance of the film by an amount that can be observed with a proper detection method. Microwave kinetic inductance detectors are small resonant circuits whose job is to measure these variations in surface impedance.

MKIDs were conceived by J. Zmuidzinas and H. G. Leduc in 1999, and early designs and demonstrations were published in 2003 [44]. Originally intended for submillimeter astronomy in which they would measure incident power and be coupled to a feedhorn or antenna, it was quickly realized that direct absorption of higher energy photons (optical, UV, and X-ray) in the superconductor could also be observed as isolated events. Extensive work has subsequently expanded upon the original idea, with more than a dozen proposed or commissioned instruments at various wavelengths using MKIDs as the detectors. A comprehensive discussion of the physics of MKIDs is beyond the scope of this thesis (many theses have in fact been written on the topic). Instead we will give a brief discussion of how they operate and refer the reader to a review article by J. Zmuidzinas [43], B. A. Mazin's dissertation [45], and J. Gao's dissertation [46], the resources I have found most useful for understanding how MKIDs work.

MKIDs are typically designed as shunt-coupled microresonators fabricated with one or more layers of thin superconducting film. Perhaps the simplest design is shown in

⁵ T_c of our platinum silicide MKIDs is typically around 0.93 K.

Figure 1.13, consisting of a shorted coplanar waveguide (CPW). Because the photon detection scheme relies on variations in the inductance of the film, the resonators are modelled as LC tank circuits as shown in Figure 1.14. In practice, they are characterized by three parameters (ignoring parasitic effects of the measurement system): the resonant frequency f_r , the internal quality factor Q_i , and the coupling quality factor Q_c . f_r is related to the inductance L and capacitance C of the model tank circuit by $f_r = 1/2\pi\sqrt{LC}$, while Q_i and Q_c are related to the total quality factor $Q = f_r/\Delta f$ (resonant frequency divided by the full width half max of the resonance dip) by

$$\frac{1}{Q} = \frac{1}{Q_c} + \frac{1}{Q_i}. \quad (1.4)$$




Figure 1.13: Quarter wave coplanar waveguide resonator capacitively coupled to a feedline. The length, width, and gap of the resonator define its resonant frequency. The size of the coupling capacitor and its proximity to the transmission line set the coupling quality factor Q_c , the internal quality factor Q_i depends largely upon the material and film deposition process, and together they define the total quality factor $Q = (1/Q_i + 1/Q_c)^{-1}$. Reproduced and modified from [43].




Figure 1.14: A frequency multiplexed array of MKID detectors can be modelled as many LC tank circuits in parallel. Each has a unique combination of L and C and therefore a unique resonant frequency $f_r = 1/2\pi\sqrt{LC}$. The MKIDs are driven simultaneously by microwave probe tones propagating from port 1 to port 2, and room temperature electronics demodulate the signals to measure changes in each resonator’s inductance. Reproduced from [43].

The transmission response, given as the forward scattering amplitude⁶, is well described by [46]

$$S_{21} = 1 - \frac{Q_i/Q_c}{1 + 2j Q_i \frac{f-f_r}{f_r}}. \quad (1.5)$$

$|S_{21}|^2$ is a Lorentzian profile, the characteristic “resonator dip” shown in the upper left panels of Figure 1.15. In the complex plane S_{21} is a circle with diameter $d = Q/Q_c$, as shown in the lower left panels of the Figure. It is useful for measuring signals from MKIDs to consider the phase of S_{21} with respect to the center of the circle, shown in the upper right panels, which is described by

$$\theta = -\arctan \left[2Q \frac{f - f_r}{f_r} \right]. \quad (1.6)$$

Any variations of the kinetic inductance from absorbed photons will change f_r , moving the phase curve downward in frequency as illustrated by the time evolution in Figure 1.15.

The forward transmission many line-widths away from the resonant frequency such that $|f - f_r/Q| \gg 1$ is $S_{21} = 1$, meaning that signals far away from the resonator in frequency space will pass by it unperturbed. This suggests a natural way to multiplex

⁶The forward scattering amplitude is the ratio of the voltage wave leaving port 2 to the voltage wave incident at port 1: $S_{21} = V_2^-/V_1^+$.

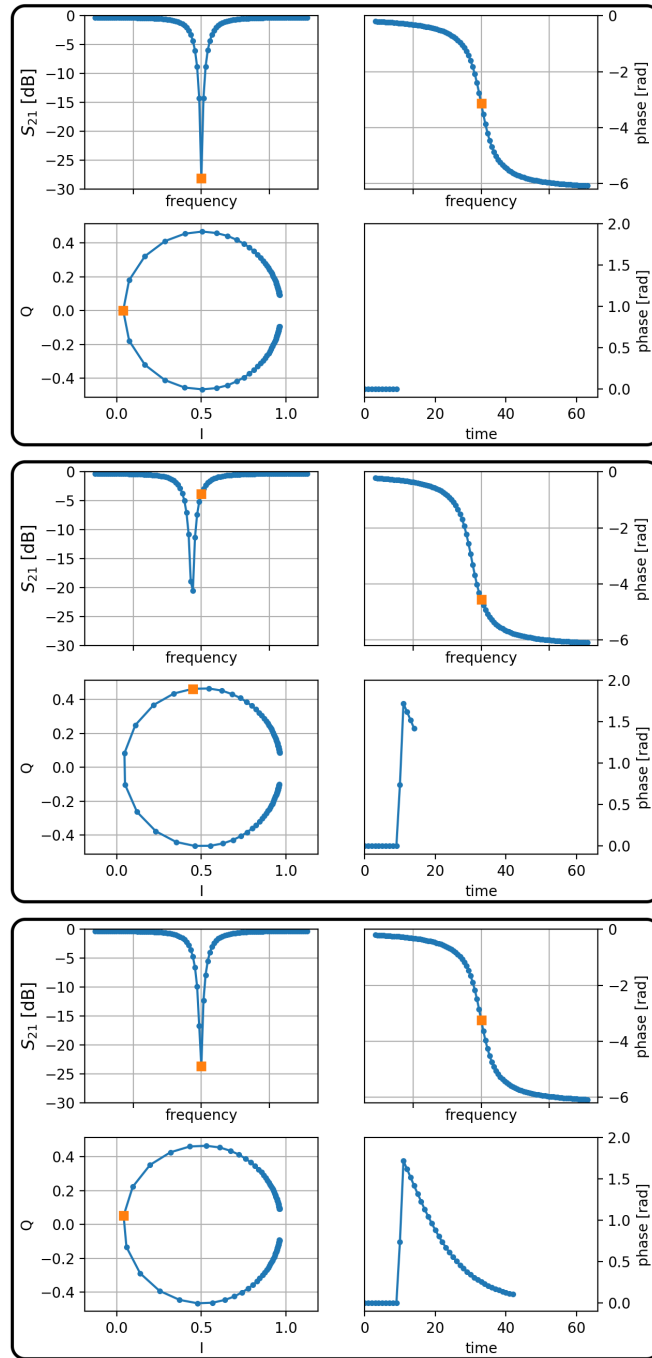


Figure 1.15: Time evolution of simulated MKID signal of a photon absorption. In each group of four panels, the upper left plot is the magnitude of the transmission, the lower left is same data in the complex plane, the upper right is the phase measured from the center of the circle vs. frequency, and the lower right is the (negative) relative phase of the orange square as a function of time. The first group of plots shows the resonator at equilibrium, the second group is immediately after a photon is absorbed in the superconductor, and the last group is when the resonator has nearly returned to equilibrium.

them in the frequency domain. Thousands of detectors can be read out in parallel with one transmission line, requiring only that each MKID resonator is assigned a unique frequency during lithography, as indicated in Figure 1.14.

Reading out thousands of MKIDs simultaneously requires high speed, high resolution digital-to-analog and analog-to-digital converters to generate a comb of microwave tones such that each MKID is excited by a single tone whose frequency matches its own f_r . The phase of each tone is carefully monitored, and photon detections are registered as pulses in the phase timestream, as shown in the lower right panel of Figure 1.15.

The size of a pulse depends on the fractional change in kinetic inductance and therefore the fractional change in the number of Cooper pairs. Photons with more energy break more Cooper pairs on average, meaning that the pulse height is a direct measurement of photon energy. However, there are limits to the range of resolvable energies. Photons with too low an energy produce pulses that cannot be distinguished from amplifier and other types of noise, while photons with too high an energy saturate the phase curve. As such, careful attention must be given to the design of the resonators to accommodate the observation wavelengths.

At optical and near-IR wavelengths, photons couple to MKIDs through an absorption event in which a shower of electron excitations in the superconductor gradually transitions into a localized bubble of phonon excitations which later return to thermal equilibrium [47]. This is in contrast to longer wavelength astronomy, where photons couple into superconducting films through antennas or transducers fed by feedhorns [48, 49]. To precisely control the sensitivity of our MKIDs, we employ a lumped-element design where a meandered inductor and interdigitated capacitor (IDC) form discrete elements of the tank circuit as shown in Figure 1.16. After choosing the pixel pitch, which is determined by the desired plate scale⁷ (the field of view of one pixel) and lithographic

⁷DARKNESS's pixel pitch is $150 \mu\text{m}$ and plate scale is 22 mas/pixel .

precision, designing the MKIDs for a specific energy range entails finding an optimal combination of four parameters.




Figure 1.16: Microscope image of a single MKID pixel. The microresonator consists of an interdigitated capacitor and meandered inductor etched into a thin superconducting film. The geometry and surface inductance of the film define the resonance frequency. The resonator is capacitively coupled to a transmission line carrying signals that excite the resonator. Photons absorbed in the inductor change the surface impedance, and therefore the resonance frequency, which can be monitored as fluctuations in the phase of a probe tone.

The first parameter is the material. Different superconductors each have their own critical temperature T_c , and even this can be changed somewhat through fabrication techniques. Equation 1.2 says that the gap energy Δ scales linearly with T_c , suggesting that photons absorbed in materials with lower T_c will break more Cooper pairs resulting in larger signals, making it desirable to use materials with the lowest possible T_c . On the other hand, the cryostat at one's disposal determines how cold a device can get. Limited budgets constrain our instruments to using adiabatic demagnetization refrigerators (ADR), capable of cooling devices to about 100 mK. It has been found empirically that

Q_i improves with decreasing temperature until $\sim T_c/8$. Hence, $0.8 < T_c < 1$ K is a good range for our MKIDs. Our material of choice is platinum silicide, which is relatively easy to fabricate and has $T_c = 900$ mK [50].

The second design parameter is the surface inductance L_s of the material, with thicker films possessing smaller inductance. The minimum thickness is constrained by the opacity of the material; films that are too thin start to become transparent, impacting the ability to absorb photons. Our PtSi films are 55 nm thick with a surface inductance of 10.5 pH/ \square and absorption ranging from 30 to 40% across the 700 to 1500 nm band. The third parameter is the inductor volume which, together with L_s , sets the number of Cooper pairs within the inductor and the total inductance. The volume, combined with Δ , sets the fractional change in the number of Cooper pairs for a photon of a given energy. Finally, Q determines the phase shift as a function of frequency shift. The slope of Equation 1.6 at $f = f_0$ is $d\theta/df = -2Q$; for small fluctuations in inductance, resonators with higher Q will have larger phase pulses. Q_c is much more reproducible across different device fabrication runs than Q_i , so we design the resonators such that $Q_c \ll Q_i$, giving us $Q \approx Q_c$. The result is that the phase response is controlled by Q_c .

1.6.2 Energy Resolution

While the pulse height of a photon with a given wavelength λ can be estimated from the quantities described above, other effects can lead to very different results in practice. Because of this, it becomes necessary to calibrate the MKID response to photons of different wavelengths. We do a “wavecal” by illuminating the detectors with several monochromatic lasers with known wavelengths and fitting the peaks of the pulse height distributions to a polynomial to get λ as a function of pulse height, which we later use for science observations. Figure 1.17 shows an example of a single pulse, and the inset

shows the pulse height distribution from simultaneous illumination with three lasers. The apparent line profiles are much wider than the true widths, leading us to conclude that the detector has a limited energy resolution, given by

$$R = \frac{\lambda}{\Delta\lambda}, \quad (1.7)$$

where for each laser, λ is the center of the peak and $\Delta\lambda$ is the Full Width Half Max (FWHM) of the peak.

There are several ways that R can degrade in our detectors. First, two level systems (TLS) at the superconductor-substrate interface provide a mechanism for macroscopic fluctuations of the dielectric constant, resulting in variations of the value of the MKID's capacitance [51]. TLS noise can be mitigated through careful sample preparation or by altering the resonator design. Second, the cryogenic HEMT amplifiers we use have effective noise temperatures of ~ 2 K, above the physical temperature of the MKIDs and thus contributing extra noise downstream from the detectors. This can be improved by using components with lower noise temperatures, such as traveling wave parametric amplifiers [52, 53, 54]. Third, geometric effects can produce inhomogeneities in the current distribution within the superconducting film such that photons absorbed in different locations on the MKID inductor may have different spatial distributions of broken Cooper pairs. This can also be mitigated by changing the resonator design. Fourth, the energy down-conversion process after a photon is absorbed results in an athermal phonon distribution. Hot phonons may escape the superconductor during down-conversion, carrying away energy that may have otherwise gone into breaking additional Cooper pairs. In-lab measurements yield $R \sim 6-7$, while MEC is limited by thermal noise coupling at $R \sim 3-4$. Recent measurements with parametric amplifiers demonstrated $R \sim 10$ [54].




Figure 1.17: A 671 nm photon absorbed in a platinum silicide film. The quasiparticle recombination time, indicated by the decay time of the pulse, is $36 \pm 2 \mu\text{s}$. Inset: The spectrum of three lasers of wavelengths 406.6, 671.0, and 982.1 nm. The smooth line is a fit of three Gaussians indicating nearly uniform energy resolution of $R = \lambda/\Delta\lambda = 8$ across the 400-1100 nm range. Reproduced from [50].

1.6.3 Timing Resolution

We design our MKIDs to have resonance frequencies covering $f_0 = 4 - 8$ GHz, with each feedline coupled to 2000 resonators. The frequency spacing between resonators should be 2 MHz, but nonuniformities in the superconducting film and other defects cause some resonators to have the wrong frequencies. This often results in two or more resonators being very close to each other and sometimes colliding. The channelization process implemented in the readout electronics firmware uses low pass filters (LPF) to isolate the signals from each resonator from its nearest neighbors [55]. If we choose a large cutoff frequency for the LPF, we can sample the phase of resonators more quickly, resulting in higher precision for photon timestamps. The tradeoff is that with bigger frequency bins (whose sizes depend on the LPF cutoff), more resonators may fall into the same bin which effectively kills both pixels. We have found that a cutoff of 100

kHz strikes a good balance, resulting in a timing resolution falling between 1 and 10 μ s depending on the type of filter used. This means that every photon recorded with our MKIDs has a timestamp associated with it whose precision is a few μ s. As we will see in Chapter 3, this turns out to be a powerful tool for measuring the intensity distribution of speckles as they evolve with time and separating contaminating starlight from planet light.

1.6.4 Quantum Efficiency

As alluded to earlier, the choice of superconducting material and its thickness determine its quantum efficiency (QE), the fraction of photons absorbed in the film. Our MKIDs' QE is relatively low, ranging from about 30 to 40% from 700 to 1500 nm. In a regime where we are starved for photons (our goal is to observe dim planets many lightyears away after all), the QE can limit the science we are able to do.

Efforts are underway to improve the fraction of photons absorbed by the detectors by applying an antireflection (AR) coating to the inductor. Preliminary results using a $\text{SiO}_2/\text{Ta}_2\text{O}_5$ bilayer have demonstrated an increase of quantum efficiency of $\sim 20\%$ with respect to the values measured without AR coating across the science band, although we suspect that optical alignment issues may have been limiting this result and subsequent measurements of the QE may be even higher.

1.6.5 Data Format

The format of the data coming out of the MKID instruments is very different from that of traditional imaging cameras. It is constructed such that individual photons are recorded by 64 bit words, or “photon packets”, containing four pieces of information. Each packet includes the peak phase of the pulse (remember that phase maps to wave-

length), the timestamp at which the photon was detected, and the (x,y) location within the MKID array.

1.6.6 Reaching the Photon Noise Limit

In semiconductor detectors, such as CCDs, photons are detected by measuring the number of excited electrons in specific parts of the device after an exposure. CCD performance is characterized by, among other things, “dark current” and read noise. Dark current comes from a population of excited electrons that were not produced by photons. Instead, they come from thermal fluctuations in various regions of the detector, and the result is an extra noise term in the intensity. Cooling the chip can reduce dark current, leaving read noise to dominate the detector sensitivity. Read noise comes from the conversion of photo-electrons to voltages.

MKIDs have no dark current or read noise because of the different detection method. As described above, MKIDs are sensitive to changes in the population of Cooper pairs. The superconductor gap energy Δ is thousands of times smaller than the semiconductor bandgap, and the change in the population of Cooper pairs is consequently much larger than the population of photo-electrons in the semiconductor. MKIDs do have phase noise, as shown on the pulse in Figure 1.17, but triggering thresholds in the readout system can make the rate of false photon counts arbitrarily small. The result is that MKIDs are photon-noise limited. This is profound, because it means that measurements of photon statistics are limited by fundamental physics and not by detector engineering.

1.6.7 The Power of MKIDs

As the exoplanet community seeks better performance from integral field spectrographs in order to detect and characterize Earth-like planets in the habitable zone through

direct imaging, MKIDs answer many of the requirements needed to explore this new area of parameter space. MKIDs combine the ability to simultaneously produce images, count individual photons with high timing precision, and measure their wavelength. These abilities make MKIDs a unique and very powerful tool.




Figure 1.18: The 20,000 pixel MKID array for the MKID Exoplanet Camera (MEC).

1.6.8 MKID Pipeline

The MKID pipeline, based on the original ARCONS pipeline [56], is an open-source, python-based set of software for processing MKID data, and can be found at <https://github.com/MazinLab/MKIDPipeline>. With every photon saved as a 64-bit word, if we assume that all the pixels in a 10,000 pixel DARKNESS array are registering 1000 counts per second, the data rate is 80 Megabytes/sec. Because the rate is fairly large, raw data is recorded to disk in one-second long binary files. The first step in the data pipeline is to unpack these binary files and convert them to a more accessible format.

We use a hierarchical data format known as HDF5 or simply “h5”.

After conversion to h5 files, calibrations are performed in which coefficients from the wavecal convert raw phase to wavelength for every pixel and a flatcal correction for differences in quantum efficiency. Other higher order effects are also calibrated, including spectral weights measured from the flatcal at different wavelengths to correct for wavelength dependent QE, a phase noise weight to correct for photons with phases near the cutoff threshold, and a correction for count rate linearity.

The pipeline is still under development, but significant progress has been made toward streamlining and integrating the various modules to minimize reliance on the user. Stay tuned.

Chapter 2

PICTURE-C

2.1 Project Overview

The Planetary Imaging Concept Testbed Using a Recoverable Experiment-Coronagraph (PICTURE-C) [57] is a system designed specifically to image debris disks at 600 nm around a handful of nearby stars. It is a NASA balloon mission with two proposed flights hosting a 0.6 meter primary mirror, adaptive optics (AO) system, vector vortex coronagraph (VVC), and MKID camera. It is designed to achieve a contrast floor of at least 10^{-7} in a half plane ranging from $1.7 \lambda/D$ to $10 \lambda/D$ (see Figure 2.1) using the MKIDs as a high order wavefront sensor. Our lab at UCSB is building the MKID science camera, a near clone of DARKNESS [1], while our collaborators at UMass Lowell are building the pointing structure and optics.

The first flight, using an EMCCD as the science camera, possibly as early as September 2019, will fly from Fort Sumner, New Mexico for one night of observations. The gondola, AO system, and supporting subsystems will be flight-demonstrated then. The second flight will use the MKID camera and the VVC, again for one night of observations.

Observing debris disks at visible wavelengths is difficult because, as discussed in

Section 1.4.1, current visible AO systems can provide only low Strehl ratios from the ground. This severely limits one’s ability to resolve the spatial extent and features of a disk. However, r_0 increases with altitude, meaning that AO correction becomes easier the higher one goes. A short duration helium balloon will raise PICTURE-C to an altitude of 35-45 km, where $r_0 \approx 40$ m [58]. With $D \ll r_0$, the atmosphere is no longer the dominant source of wavefront distortions. Flexure from swinging of the gondola and aberrations from imperfect optics remain, still requiring AO in order to achieve the 10^{-7} contrast specification.




Figure 2.1: The PICTURE-C AO and coronagraph are designed to achieve better than 10^{-7} contrast in a dark hole covering a half-plane ranging from $1.7 \lambda/D$ to $10 \lambda/D$. Reproduced from [59].

PICTURE-C’s science goal to directly image debris disks from a high altitude balloon is challenging because of the system’s complexity. If successful, it will also serve to demonstrate and raise the technology readiness level, a metric used by NASA to gauge how close technologies are to being spaceflight worthy, of both MKIDs and VVCs to TRL 6. This is important as future space missions may seek to use these technologies.

Establishing that they can work in a balloon environment will show that they are nearly ready for deployment in space.

2.2 The Science of PICTURE-C: Imaging Nearby Debris Disks

Debris disks are a type of low mass circumstellar disk orbiting main sequence stars. They are optically thin and gas poor, and sustained by a steady state collisional cascade. Large particles formed early on in a system’s evolution collide with each other and fragment, forming smaller particles. The particles range from micron-sized dust grains to planetesimals hundreds of kilometers across. The so-called “blowout size” sets the lower bound of the grain size distribution. Grains smaller than this are more strongly affected by radiation pressure from their star than its gravity, and henceforth are blown away from the star. The radiation and drag forces on larger, gravitationally bound grains and planetesimals influence these particles to an extent that depends on their mass and cross section, affecting their orbits differently [10]. Planets can also influence the spatial distribution of dust in the disk, making debris disk structure an indicator for their presence [60].

Millimeter and larger sized grains are more efficient at emitting long wavelength radiation than smaller grains, allowing one to image the parts of a disk made up of them in the far-IR or sub/millimeter. Micron sized grains are poor emitters, but are efficient at scattering wavelengths comparable to their size and can be detected with visible and near-IR imaging. Imaging a disk from the visible out to the millimeter allows for mapping of the spatial distributions of all sizes of dust, and produces rich details about the circumstellar environment. Images from different epochs can also demonstrate

interesting temporal behavior, like the movement of features in the AU Mic debris disk on timescales of only a few years [61], as shown in Figure 2.2.

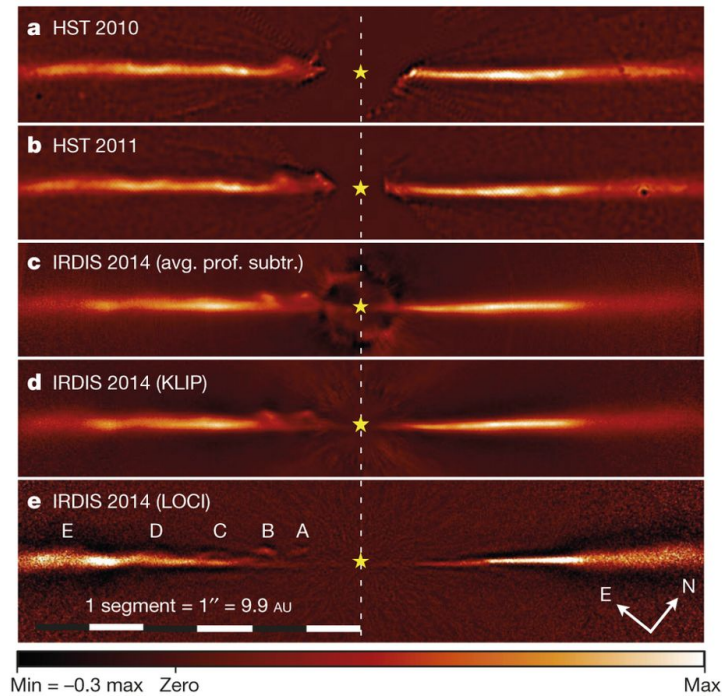


Figure 2.2: Direct images of AU Mic from three epochs clearly show the movement of features in the debris disk. Reproduced from [61].

About a hundred debris disks have been spatially resolved, mostly in thermal emission with the Herschel Space Observatory, JCMT, or ALMA, while only about 40 have been resolved in scattered light [62]. Considering their close relationship to planets, a more complete understanding of debris disks can inform planet formation models and their role as signposts for as-yet undiscovered exoplanets.

2.3 Work at UMass Lowell: Gondola and Optics

2.3.1 Gondola

Pointing of the PICTURE-C telescope will be done by the Wallops Arc Second Pointer (WASP), a system able to actively control pitch and yaw to provide stable pointing for several hours with sub-arcsecond accuracy [63]. The WASP was developed by NASA engineers to provide a flexible, reusable, general purpose platform to enable science projects with restrictive pointing requirements. Balloons typically spin when aloft, so the WASP works in conjunction with a standard gondola rotator, an azimuth control device located above the gondola providing rough azimuth positioning with an accuracy of 1 to 3 degrees.

The WASP will house the primary mirror, the adaptive optics and coronagraph, and the MKID science camera. The gondola structure must also accommodate all of the electronic control systems, batteries, and radiator panels for dissipating waste heat, though these pieces will not be part of the pointed section to minimize the weight whose orientation must be controlled.

2.3.2 Optics

The primary mirror (M1) is a 60 cm (24 inch) off-axis paraboloid (OAP) accompanied by a off-axis elliptical secondary mirror (M2) that directs the beam into the coronagraphic optics enclosure. A pickoff mirror with a hole in the center, located at the second focus of M2, directs part of the beam to an acquisition camera with a 10 arcminute FOV. The hole defines a 20 arcsecond FOV. The tertiary mirror (M3) collimates the beam and reimages the entrance pupil (M1) onto the low order deformable mirror (LODM), an ALPAO DM97-15 with 97 actuators and tip-tilt control. The LODM corrects low order




Figure 2.3: CAD model of the PICTURE-C gondola. At the top is the gondola rotator, and below is the rectangular frame supporting the optics and MKID cryostat (purple).

aberrations at 200-500 Hz using one of two low order wavefront sensors (LOWFS). One LOWFS, a Shack-Hartmann WFS, uses out-of-band blue light at 430-540 nm to provide low precision, high dynamic range measurements of the wavefront. The second LOWFS is a low dynamic range, high precision reflective Lyot stop sensor that uses light rejected by the coronagraph for common-path, in-band measurements of the wavefront [64, 65]. Two OAPs (OAP1 and OAP2) reimage the beam onto the high order deformable mirror (HODM), a 952 actuator Kilo-1.5 DM manufactured by Boston Micromachines Corporation. A second pair of OAPs (OAP3 and OAP4) reimage the beam onto the Lyot stop, with the VVC at the focus between them. A set of polarizers and quarterwave plates filter chromatic leakages, and finally two active flat mirrors (FM1 and FM2) direct the f/662 beam to the MKID science camera. See Figure 2.4 for a ray trace diagram illustrating the optical path.

Electric field conjugation (EFC) [67] will be used to reach the contrast floor in the MKID focal plane. Briefly, EFC takes knowledge about the complex electric field and the DM's influence matrix to determine the corrections the DM should apply. The electric field is measured with a calibration prior to turning on the control loop and then combined with a detailed model of the optical system (including the DM influence matrix) to dig the dark hole. Once the dark hole is established, an integrator control loop will provide corrections needed to maintain it.

2.4 Design of the PICTURE-C MKID Camera

The PICTURE-C MKID cryostat, shown in Figures 2.5, 2.6 and 2.7, is cooled to 4 K with liquid helium and an intermediate liquid nitrogen stage maintaining 77 K. Cooling with the evaporation of these liquids is standard practice for balloon experiments because




Figure 2.4: Top: Side view of a ray trace diagram showing the design of the PICTURE-C optics, developed by our collaborators at UMass Lowell. Bottom: Top view of the optical bench hosting the AO system and coronagraph. Reproduced from [66].

they require no external power¹. The “cold stage” holding the MKIDs at 100 mK is cooled by an Adiabatic Demagnetization Refrigerator (ADR) that provides ~ 12 hours at base temperature per cycle. An anti-reflection coated window and optical filters at 77 K and 4 K define the optical bandpass of 540-660 nm while rejecting out of band light. The MKID array has 10,000 pixels whose drive tones are brought to the cold stage with five parallel transmission lines, each a combination of commercial coax and custom flexible cables. Temperatures are monitored at each stage with either silicon diode or ruthenium oxide thermometers.

In many ways PICTURE-C is a clone of DARKNESS. The overall cryostat layout borrowed heavily from its predecessor, with the same type of tanks, ADR, and MKID device support structure. There were, however, a number of improvements over the DARKNESS design that we wished to implement. The most important change was an optional pulse tube cooler that could cool the fridge without cryogenics for in-lab testing, making the fridge drastically easier and cheaper to operate. Another important change was the size of the cryogen tanks, because we wished to increase the hold time to about 50 hours from DARKNESS’s 40. The microwave transmission line “flex” cables in DARKNESS had problems with crosstalk, whereby signals from one feedline would leak into those adjacent to it and disrupt the readout of pixels that happened to have very similar frequencies.

There are improvements outside the cryostat as well. First, the thermometry and ADR control electronics had to be reduced in scale and simplified in operation. Second, the balloon payload will carry batteries, so the power supplying all of the electronics will be DC, unlike for DARKNESS which uses standard 110 VAC power. Finally, the Gen2 readout electronics for DARKNESS is a power-hungry piece of hardware, consuming about 1.4 kW. Development of Gen3 is ongoing, using state-of-the-art FPGAs with

¹Cryocooler vibrations are also a concern for certain experiments, but are not problematic for MKIDs.

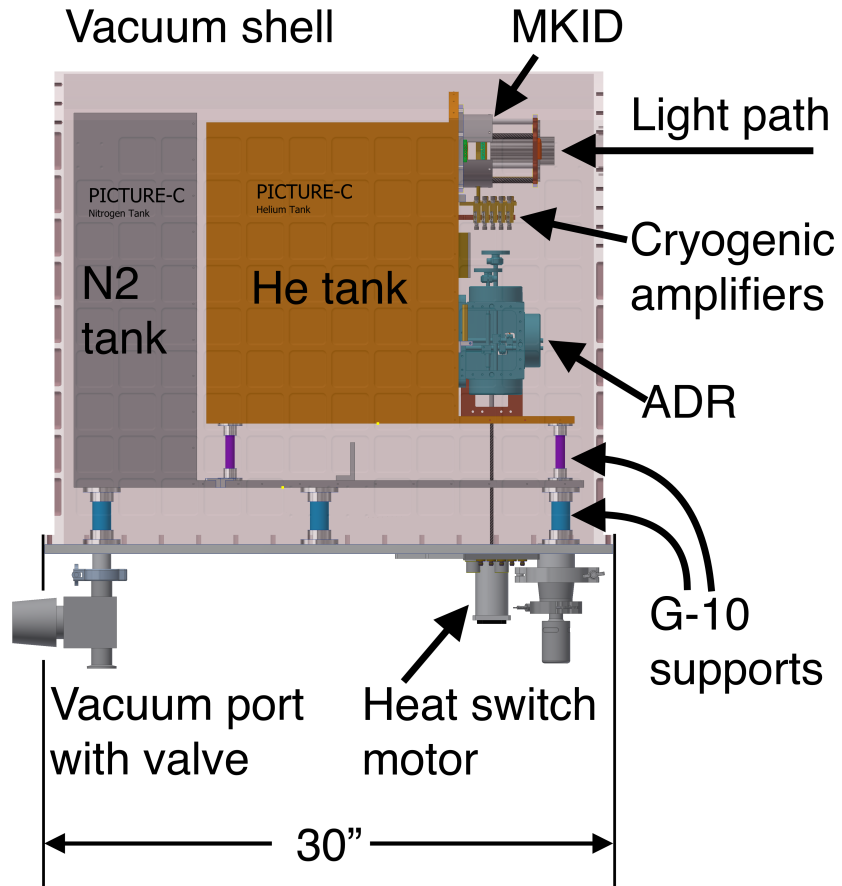


Figure 2.5: Side view drawing of the PICTURE-C cryostat in the upright flight orientation. Liquid helium maintains a bath temperature of ~ 4.2 K with an intermediate liquid nitrogen stage. An adiabatic demagnetization refrigerator cools the MKID devices to 100 mK. Cryogenic HEMT amplifiers provide ~ 40 dB of gain immediately after the MKIDs.



Figure 2.6: The PICTURE-C cryostat mounted on its lab frame with the pulse tube installed. The cryostat is upside down compared to the flight configuration.




Figure 2.7: A composite image showing the device stage and ADR of the PICTURE-C cryostat.

integrated ADC/DACs that can do the same job with about 10% the power.

As the first high contrast imaging camera of its kind, DARKNESS served as a remarkable testbed for a number of hardware and software developments. From it we learned a great deal about what worked, what didn't work, and what could have been better. Here we will detail some of the important features of PICTURE-C, pointing out the lessons learned from DARKNESS and the improvements made in response to them.

2.4.1 Cryogen Hold Time

The desired hold time (the time between topping off and running out) for PICTURE-C was 50 hours for both the liquid nitrogen and liquid helium. Hold times depend on the volume of cryogen and the heat load seen by each: a greater volume will take longer to boil off, and a larger heat load will evaporate the liquids more quickly. With the hold time given as a specification, calculating the required volume (i.e. designing the cryostat) necessitated an estimate of the heat load. Estimating heat loads, especially on cryostats whose final designs are not yet known, is challenging. We used DARKNESS as a benchmark to calculate the effective heat loads and inform decisions on the PICTURE-C design. The intent with DARKNESS was to have a helium hold time of about 40 hours and a nitrogen hold time of about 80. The hold times were about 40 hours for both liquids. Here we will discuss the discrepancy at 77 K and the validation at 4 K for DARKNESS.

Measuring the heat loads of the cryogen was straightforward: use the hold time (40 hrs), tank volume (~ 9.8 L for N_2 , ~ 22 L for He), heats of vaporization (5.56 kJ/mol for N_2 , 0.0829 kJ/mol for He), densities (.807 g/mL for LN_2 , .125 g/mL for LHe), molecular

Table 2.1: Estimated heat loads for DARKNESS nitrogen tank during design phase. The estimate was significantly lower than what we measured in the lab.

Item	Heat Load (Watts)
G-10 Supports	2.2
Radiation from 300 K shell	1.4
2x 50 cm length 16 gauge copper wire for magnet	0.5
10x SS-085 Coax	0.1
Thermometry and HEMT Wiring	0.07
Total Estimated 77 K Heat Load	4.27
Measured 77 K Heat Load	10.9

weights (28 g/mol for N₂, 4 g/mol for He) and some dimensional analysis to obtain

$$\frac{9.8 \text{ L}}{40 \text{ hr}} \times \frac{10^3 \text{ mL}}{\text{L}} \times \frac{.807 \text{ g}}{\text{mL}} \times \frac{1 \text{ mol}}{28 \text{ g}} \times \frac{5,560 \text{ J}}{\text{mol}} \times \frac{1 \text{ hr}}{3600 \text{ s}} = \mathbf{10.9 \text{ W}} \quad (2.1)$$

for the nitrogen stage and

$$\frac{22 \text{ L}}{40 \text{ hr}} \times \frac{10^3 \text{ mL}}{\text{L}} \times \frac{.125 \text{ g}}{\text{mL}} \times \frac{1 \text{ mol}}{4 \text{ g}} \times \frac{829 \text{ J}}{\text{mol}} \times \frac{1 \text{ hr}}{3600 \text{ s}} = \mathbf{0.4 \text{ W}} \quad (2.2)$$

for the helium stage.

The early estimates listed in Table 2.1 for conductive components agree with follow up estimates based on more careful measurements of lengths and thicknesses. However, there is one large, conductive heat load not included in the table: the stainless steel fill and vent lines providing access to the cryogen tanks.

We can estimate one dimensional conductive heat loads by numerically integrating Fourier's law of heat conduction,

$$P = \frac{A}{l} \int_{T_{\text{low}}}^{T_{\text{high}}} \kappa(T) dT, \quad (2.3)$$

where P is the power transferred from one end of a wire (or any piece of hardware), A

Table 2.2: Estimated heat loads for DARKNESS helium tank during design phase.

Item	Heat Load (Watts)
G-10 Supports	0.14
Radiation from 77 K shell	0.13
77K Radiation on 4K window	0.0025
High T_c Superconductor Leads	0.0063
10x SS-085 Coax	0.06
Thermometry and HEMT Wiring	0.045
Heat Switch Rod	0.01
Total Estimated 4 K Heat Load	0.39
Measured 4 K Heat Load	0.4

and l are the cross-sectional area and length of the object of interest, and $\kappa(T)$ is the thermal conductivity as a function of temperature T . Measurements of $\kappa(T)$ for various materials can be downloaded from NIST's website².

Each tank has a fill line with an outer diameter (OD) of .25" and inner diameter (ID) of .21", as well as a vent line with OD = .375" and ID = .335", with a section that extends into the tank to increase the effective length and reduce the heat load at the cold end as shown in Figure 2.8. While the extensions limits the heat load from the vacuum shell to the nitrogen tank to a fraction of a Watt, a heat sink at 77 K for the helium tubes 1.625" away from the 300 K plate results in a heat load of about 1.5 W, bringing the total estimated heat load in Table 2.1 to ~ 5.8 W, still less than Equation 2.1.

Investigating the radiation load revealed that the original estimate was low. The Stefan-Boltzmann law, valid for blackbody radiation, gives the power emitted by an object with surface area A , emissivity ϵ , and temperature T as

$$P_{\text{emitted}} = \epsilon A \sigma T^4, \quad (2.4)$$

²<https://trc.nist.gov/cryogenics/materials/materialproperties.htm>

where σ is the Stefan-Boltzmann constant. Conversely, the power absorbed by an object bathed in a blackbody radiation field from all directions is $P_{\text{absorbed}} = \epsilon A \sigma T_{\text{field}}^4$. The sum of these is the total radiation load on the object, which in this case is the nitrogen stage of the DARKNESS cryostat,

$$P_{\text{rad}} = \epsilon A \sigma (T_{\text{field}}^4 - T^4). \quad (2.5)$$

where $T_{\text{field}} \approx 300$ K and $T \approx 77$ K. DARKNESS is constructed mainly of 6061 Aluminum, which has an emissivity of $\epsilon = 0.09$ [68], giving $P_{\text{rad}} \approx 45$ W³. This is much larger than Equation 2.1, but DARKNESS also has insulation made of tens of layers of aluminized mylar which can account for the discrepancy. If we include the effect of the insulation as an effective emissivity so that Equation 2.5 matches Equation 2.1, we obtain $\epsilon_{\text{eff}} \approx 0.02$.

The total shown in Table 2.2 is very close to the measured heat load on the helium stage of DARKNESS shown in Equation 2.2. Similar calculations to those presented for the nitrogen stage validate the individual heat loads with the exception of the missing stainless steel tubes, contributing an additional 36 mW.

After the heat loads were understood in DARKNESS, design of the PICTURE-C cryogen tanks commenced. Accounting for the missing heat loads in the nitrogen tank and increasing the hold times of each tank individually resulted in a helium and nitrogen tank volumes of 26.8 L and 18.8 L respectively (compared to DARKNESS's 22 L and 9.8 L), with hold times of 49 and 76 hours.

³It is hard to overstate that 300 Kelvin radiation matters a lot when designing a cryogenic testbed!




Figure 2.8: The stainless steel vent and fill lines extend into each tank to increase the length and reduce the heat load. Image courtesy Dick Gummer at Precision Cryogenic Systems.

2.4.2 Removable Pulse Tube Cooler

The steps involved in cooling a cryostat with liquid helium are as follows. First, air must be evacuated from the vacuum chamber such that the internal pressure is below about 10^{-3} mbar. Second, both cryogen tanks are filled with liquid nitrogen and allowed to equilibrate at 77 K. Following this, the nitrogen is removed from the helium tank, which is then filled with liquid helium. While this method is actually fairly quick, taking roughly 18 hours to go from room temperature to 4 K, it is very expensive and requires a great deal of care because of the hazards involved with the cryogens. Furthermore, the tanks must be refilled at regular intervals.

Closed cycle refrigerators, called Pulse Tube Coolers [69], using He as the working substance can readily cool testbeds to helium temperatures. They are commercially available, easy to use, and very reliable. After installation, operating them is as simple as turning on a compressor unit and waiting for the testbed to reach 4 K. The flight of PICTURE-C requires the fridge to be cold for less than 24 hours with liquid cryogens, but all laboratory tests leading up to the flight can use a pulse tube.

A simple cost-benefit analysis based on helium consumption of DARKNESS and the price of a new pulse tube cooler estimated that we would break even after about 20 cooldowns. When helium shortages that disrupt deliveries, the labor required to manage helium cooldowns, and the fact that our lab already possessed an extra helium compressor⁴ (a critical part of a pulse tube system), we decided to purchase a pulse tube cooler. Based on the heat loads from DARKNESS detailed in Section 2.4.1, we determined that a two-stage Cryomech model PT405 cryocooler would satisfy the requirements of PICTURE-C. The PT405 is certified to lift 0.5 W at 4.2 K and 25 W at 65 K.

In order to minimize the modifications to the DARKNESS design, we opted to insert

⁴ARCONS, the first optical MKID instrument, used an air-cooled compressor during observing runs at Palomar [70]. It was decommissioned after its final observing run in October 2014.

the Cryomech PT405 pulse tube from below. The design of the PT405 requires that it be operated vertically, and without liquids in the tanks the cryostat has no preference for its orientation. To this end, a scheme for flipping the cryostat upside down was devised in which one person could do it alone. A drawing of the inverted cryostat is shown in Figure 2.10.




Figure 2.9: The Cryomech PT405, two-stage pulse tube head is rated to lift 0.5 W at the 4.2 Kelvin stage and 25 W at 65 Kelvin stage, sufficient for PICTURE-C. Reproduced from <https://www.cryomech.com/>.

To facilitate a relatively easy installation and removal process for the pulse tube we increased the space between consecutive stages to 3" (compared to 2") in DARKNESS. The extra space allows one to easily reach between the plates and operate the tools used to install and remove fasteners. A single allen key is the only tool required.




Figure 2.10: Side view of the PICTURE-C cryostat with the vacuum shell and radiation shields removed. The cryostat is upside-down, the orientation used in the lab for cooling with the pulse tube cooler.

2.4.3 Copper Rope Heat Straps

Thermal expansion is a concern when mechanically joining parts made of different materials. We wished to avoid a situation where the pulse tube head might be pushing or pulling on the nitrogen and helium tanks, as well as provide flexibility to the relative locations of the tanks and pulse tube. Commercial cryostats avoid this by using high purity copper ropes terminated in blocks of high purity copper to make flexible thermal connections. The copper rope heat sinks for PICTURE-C needed to have a custom design, but quotes from cryostat manufacturers to build the heat sinks prompted the decision to have them made by the UCSB physics machine shop.

Connections between the 1/4" copper rope and the copper blocks were made with a tungsten inert gas (TIG) welder, using excess rope as the filler material. While the process probably introduced some contamination into the material and thereby reduced the thermal conductivity, the heat sinks provided sufficient conductivity in cryogenic tests.

To quantify the quality of thermal connection between the pulse tube head and the helium tank, two silicon diode thermometers were placed at opposite ends of the heat sink. One measured the temperature of the pulse tube cold head, while the other measured the helium tank. The values recorded were 3.45 and 3.56 K respectively. A third thermometer recorded the temperature of the ADR unit (located farther away from the pulse tube but still on the helium tank) as 4.48 K. These three measurements show that the copper heat sinks do not limit the transfer of heat away from the helium tank. Rather, it is either limited by the thermal conductivity of the tank itself, constructed of 6061 aluminum alloy, or by the interface between the heat sink and the tank. Figure 2.11 shows the 4 Kelvin heat sink installed in the cryostat.

After installing the ADR components (described in the next section), cooling the




Figure 2.11: Side view of the the PICTURE-C cryostat showing the copper rope heat sink connecting the pulse tube to the helium tank.

cryostat from room temperature to 4 K consistently takes about 48 hours. While this is slow compared to cooling with nitrogen and helium, cooling with the pulse tube is far less complicated and costly.

2.4.4 Removable Device Stage

Careful consideration was given to the PICTURE-C design so that users could quickly remove the entire structure by disconnecting the ADR heat sinks, a thermometer wire, and the microwave transmission lines. The goal was to make MKID devices easier and straightforward to install and remove. Difficulties with installing devices and wiring in DARKNESS informed the newer design.

The structure itself is secured with screws to the 4 K plate such that alignment can be maintained with less than a millimeter of translational error after reinstalling it. Figure 2.12 shows the structure installed in the cryostat without microwave transmission

lines.

A set of two copper rope heat straps was assembled to form the thermal connections between the ADR and the device stage. Each one uses three parallel 1/8" high purity oxygen free high conductivity copper (OFHC) ropes with machined OFHC blocks at the ends. The ropes were inserted into slightly oversized holes in the blocks, and a hydraulic press was used to compress the blocks making a strong thermal and mechanical connection. The assemblies can be seen in Figure 2.12.

2.4.5 Flex and Flax Cables

Low temperature experiments involving microwave devices usually rely on small coaxial cables to transmit signals between 4K and the device stage. Depending on the requirements, they might be made of a normal metal such as cupronickel (CuNi) or a superconductor such as niobium titanium (NbTi). Normal metal coax is cheaper and easier to manufacture, but can have significant attenuation and thermal conductivity. NbTi is more difficult to work with, but is nearly lossless at microwave frequencies with a much lower thermal conductivity. PICTURE-C requires 5 feedlines, as did DARKNESS. DARKNESS served to demonstrate a novel superconducting technology for low temperature microwave transmission lines, and we outline below the lessons learned that informed the new PICTURE-C design.

DARKNESS needed five independent transmission lines to read out a full MKID array, for a total of ten connections between 4 K and the device stage (input/output). Heat load estimates showed that commercial coax that could be properly connectorized for DARKNESS would present too high of a heat load to the MKID device stage. A "flex cable" solution was proposed in which a superconducting NbTi foil, an insulating Kapton layer, and five parallel NbTi microstrip feedlines as shown in Figure 2.13 could be




Figure 2.12: The device stage can be removed from PICTURE-C by quickly disconnecting the copper rope heat sinks for the FAA and GGG (shown), the MKID thermometer cable (shown), and the microwave transmission lines (not shown, but will include coax and flax cables).

laminated to meet the thermal conductivity requirement [71] and simultaneously provide structural stability. MEC, with ten feedlines, and PICTURE-C, with five feedlines, would also require flex cables.




Figure 2.13: Cross section of the microstrip flex cable. A NbTi foil formed the ground plane, glue and kapton formed the insulator, and an etched NbTi foil produced the microstrip lines. While the transmission of this design was good, crosstalk between adjacent feedlines was high enough to disrupt the readout of pixels on different feedlines with the same f_r .

During initial prototyping of the flex cable assembly, which required connectorizing the ends to mate with commercially available components and heat sinking the middle of the cable, we discovered that the NbTi microstrips were extremely brittle. They would often crack when the cable was curved with the feedlines on the outside of the bend radius. Bend radii smaller than $\sim .25''$ would stretch the NbTi strips, resulting in small fractures that would break the microstrip continuity. A single fracture could block the readout of an entire bank of 2000 pixels, making it critical that we avoid damaging the feedlines. Handling and installing the cables in DARKNESS without damaging them proved nearly impossible because of multiple sharp bends.

We promptly made adjustments to the assembly to prevent cracking. Principally, we machined curved clamps without sharp edges for the connectors and a “half-pipe” heat sink in the middle, such that the tightest bend kept the microstrips on the inner radius and under compression. The final assembly used in DARKNESS is shown in Figure 2.14, with all five feedlines on both the input and output cables working.

Cross-talk measurements of the flex cables showed that at the upper end of the MKID band, around 8 GHz, the isolation between adjacent feedlines was about -25 dB. This

was higher than we had anticipated, and indeed resulted in readout interference for some MKID resonators. If two adjacent feedlines have resonators with resonant frequencies less than 100 kHz apart (the cutoff frequency of a low pass filter in the firmware), signals from one feedline could leak into the other with enough power such that beating would trigger the readout system to register photons that did not exist. Because the fields and signal strengths from each feedline are typically very similar, the beating effect was symmetric, making *both* pixels unusable.

The shortcomings of the flex cables have been more profound in MEC. It uses input and output flex cables with ten feedlines each, so it has been especially troublesome to read out the full MKID array due to microstrip fractures. Furthermore, the exposed microstrips behave like antennas to ambient electric fields. This is problematic for two reasons. First, the cable is bathed in ~ 60 K blackbody radiation which can couple to the feedlines, degrading signal to noise. Second, the ~ 40 dB cryogenic amplifiers can develop feedback. The microstrips can radiate energy with the same efficiency as they can absorb it, so exposed cables on the output side of the amplifiers can transmit signals to the input side. A temporary workaround has been to reduce the gain of the amplifiers, but a true solution would entail better isolation of the input and output ports of the amplifiers. DARKNESS had avoided issues with blackbody pickup because the flex cables are only exposed to 4 K radiation, which is just above the effective noise temperature of the amplifiers. Amplifier feedback was also avoided because the transmission line after the amps consisted entirely of coaxial cable, which completely contains the fields.

While the DARKNESS flex cables would suffice for reading out the PICTURE-C MKIDs and we have the parts required to assemble them, the anticipated deployment date (June of 2020) permits enough time to build and test an upgraded version that we refer to as a “flax” cable. The upgrade, shown in Figure 2.15, will consist of NbTi wire with a PTFE insulator sandwiched between two NbTi foil ground planes. Circular




Figure 2.14: Top: A CAD model showing the five conductor flex cable with connectors and an intermediate heat sink. Bottom: Photograph of the flex cable installed in DARKNESS. Differential thermal contraction would curl the cable in the direction of the feedlines, resulting in touches and degraded thermal performance. Several clips made of G-10 fiberglass were used to constrain this movement as the cable cooled, and can be seen here.

grooves stamped into the foil will receive the wire, forming coaxial transmission lines. Spot welds are used to join the ground planes on either side of each wire. These cables are currently under development, but measurements of initial prototypes indicate better isolation than the flex cables. Both PICTURE-C and MEC will use flex cables, and DARKNESS will likely be upgraded as well. The final assemblies for PICTURE-C will be very similar to that shown in Figure 2.14, with modifications to accommodate the cross-sectional profile of the flax.




Figure 2.15: A proposed upgrade, dubbed a “flax” (flex-coax) cable, sandwiches a fine NbTi wire with a teflon insulator between two NbTi foil ground planes. Shown in cross section here, the ground planes are both stamped to create semi-circle grooves to receive the insulated wire. This scheme should provide better isolation between feedlines and from ambient fields than the older flex cables. Preliminary measurements of a flax prototype indicate that adjacent feedlines have 10-20 dB better isolation than in the flex cable. PICTURE-C will likely use flax cables.

2.4.6 Tank Baffles

During its flight, PICTURE-C will experience lots of accelerations. One concern about this is sloshing of the liquids in the cryostat, because the center of mass can change very quickly and disrupt telescope pointing. Baffles were installed inside both the nitrogen and helium tanks to minimize the movement of liquids without limiting their access to all parts of the tanks. They are constructed of aluminum sheet metal, and add less than 1 kg of mass. Figure 2.16 shows the design of the helium tank baffle.




Figure 2.16: Left: CAD model showing the design of the baffle for the helium tank, which should minimize sloshing of the liquid without forming independent cavities. A similar baffle will reduce sloshing in the nitrogen tank. Right: Perforated aluminum sheet metal actually used in the cryostat. Image by Dick Gummer at Precision Cryogenics.

2.4.7 Fill and Vent Port Fittings

The PICTURE-C design reused much of DARKNESS's, including filling the cryogenics from below. Unlike DARKNESS, it will operate in a rarified atmosphere with an ambient pressure of $P \sim 1$ mbar, about 10^3 times smaller than sea level. The boil-off rate of the cryogenics will be controlled with absolute pressure valves, requiring fittings more robust to cryogenic temperatures than rubber O-rings. ConFlat[®] is a type of vacuum fitting using deformable, single-use copper gaskets between two stainless steel knife-edges. They are extremely good at protecting against leaks, and can endure a wide range of temperatures. The PICTURE-C fill and vent ports were built with ConFlat[®] fittings to prevent leaks and provide a simple design using off-the-shelf components.

2.5 ADR

An adiabatic demagnetization refrigerator (ADR) shown in Figure 2.17 was procured from High Precision Devices for cooling the MKID devices from 4 K to 100 mK. It uses a combination of two paramagnetic salt pills made of ferric ammonium alum (FAA) and gadolinium gallium garnet (GGG) in the bore of a 4 Tesla superconducting magnet and a mechanical heat switch. It takes advantage of the fact that paramagnetic materials have a lower entropy in a magnetic field because as the field strength increases, the spins in the lattice become more ordered.

The steps to use the ADR after cooling the cryostat are as follows.

1. A mechanical heat switch is closed, connecting the salt pills to the 4 K thermal bath (the helium tank). This should be done at the beginning of the cooldown, so that the pills reach 4 K as soon as possible.
2. Current through the superconducting magnet is ramped up at a constant rate such




Figure 2.17: The two-stage adiabatic demagnetization refrigerator uses a superconducting magnet and paramagnetic salt pills to provide over 100 mJ of cooling from 4 K. This is sufficient to bring the temperature of the MKIDs and their support structure to 100 mK for 12 hours.

that the back-emf doesn't exceed 250 mV, a value specified by the manufacturer to prevent damage. 5 mA/s reaches the maximum current of 9.4 A in about 30 minutes and stays below the limit. Decreasing the entropy in this step is an exothermic process, but the heat switch connection transfers the excess heat into the 4 K bath.

3. The pills are "soaked" in the magnetic field for about 30 minutes, allowing them to completely thermalize.
4. The heat switch connection is opened, thermally isolating the salt pills.
5. The current is ramped back down to zero at the same rate as before. As the field decreases the spins become more disordered, increasing the entropy and consequently lowering the temperature.

The principle of operation is the same for both salt pills. The difference between them is their base temperatures and heat capacities. The cold stage FAA routinely achieves temperatures around 80 mK in PICTURE-C, but has a very small heat capacity ($\lesssim 100$ mJ). The warmer GGG typically reaches 800 mK during a magnet cycle, but is able to absorb $\gtrsim 1$ J⁵. The purpose of GGG is to provide an intermediate buffer between 4 K and 100 mK by routing all the wiring to the FAA stage through a heat sink mounted on the GGG, reducing the heat load on the FAA and thereby extending its hold time.

Maintaining a constant operating temperature is important because there is a population of thermally excited quasiparticles (broken Cooper pairs) in the MKIDs. Increasing the temperature has the effect of lowering the internal quality factor Q_i in the resonators. As such, we use closed-loop feedback to keep their temperature at $100 \pm \sim 0.3$ mK. The temperature is measured with a thermistor mounted to the MKID device box and a proportional-integral (PI) controller sends an analog voltage to the magnet power supply, which then adjusts a small current (typically below ~ 140 mA) to correct for any

⁵These values come from simulations by High Precision Devices.

difference between the measured temperature and the 100 mK setpoint. Figure 2.18 shows how the regulation current changes over time. For more details on how ADRs work, see [72, 73].




Figure 2.18: The 100 mK temperature of the MKIDs is maintained by continuously adjusting the current through the superconducting magnet. Shown here is a demonstration of the PICTURE-C ADR before (orange) and after (green) fixing an internal radiation leak compared to a typical curve from DARKNESS (blue). The point at which the current reaches zero is what we refer to as the ADR hold time.

2.6 Testing the PICTURE-C Cryostat

Verification of the PICTURE-C cryostat has involved several cooldowns, with various problems being diagnosed and resolved as we neared a complete assembly. First, the pulse tube was shown to regularly maintain a base temperature of ≈ 4.5 K as measured by a thermometer in the ADR unit, and to bring the nitrogen stage to ≈ 61 K. According to the capacity curve published by the manufacturer shown in Figure 2.19, these temperatures

corresponded to heat loads of about 0.7 W and 20 W for the helium and nitrogen tanks respectively. The helium stage heat load was worrisome because it was significantly higher than predicted.

Possible causes for this discrepancy were investigated, including light leaks allowing 300 K radiation through the nitrogen stage radiation shields, incomplete coverage with multilayer insulation at 4 K, and residual helium gas within the helium tank transferring heat from the SS fill/vent lines to the innards of the cryostat. Light leaks and the insulation coverage proved to have a negligible effect, while residual helium gas did appear to be an issue with one cooldown where the temperature was abnormally high at ≈ 5.5 K. The helium tank temperature remained at or above 4.5 K despite multiple attempts to reduce it.

After considering the locations of the thermometers compared to the pulse tube, we realized that the heat load may actually be smaller than previously thought due to a temperature gradient. Indeed, a thermometer located at the pulse tube head later measured ≈ 3.5 K (similar to the measurement of the temperature gradient from the top to the bottom of the copper rope heat straps described earlier). Again using the capacity curve in Figure 2.19, the heat load appears to be ≈ 0.4 W on the helium stage as indicated by the red circle.

The difference in heat loads between the DARKNESS nitrogen stage in Equation 2.1 and the 25 W estimate from the capacity curve for PICTURE-C can likely be attributed to incomplete multilayer insulation on the nitrogen stage. While this extra heat presents no difference to the operation of the pulse tube, care should be taken when using liquid cryogens to cover up as much of the nitrogen stage as possible with insulation to minimize heating from 300 K blackbody radiation.




Figure 2.19: Capacity curves of the PT405 pulse tube. Measurements of both pulse tube stages give an indication of their respective heat loads. We estimate the helium stage is absorbing ≈ 0.4 W while the nitrogen stage absorbs ≈ 25 W. Reproduced and modified from <https://www.cryomech.com/>.

2.7 PICTURE-C Cryostat Electronics

The thermometry and ADR wiring replicated those of DARKNESS, such that the DARKNESS electronics rack, shown in Figure 2.20 could be used while the PICTURE-C hardware was still being procured and assembled. Weighing 256 kg (564 pounds), taking up ≈ 0.8 m³, and consuming ≈ 1700 W altogether, the rack design was not optimal for a balloon observatory. Scaling down the PICTURE-C electronics in weight, size, and power requires scaling down the individual components wherever possible. For the most part, DARKNESS uses very versatile instruments with many features that are unnecessary for its operation. We intend to replace most of them with instruments having a smaller set of capabilities.

Replacing the temperature monitors is straightforward. DARKNESS uses two units from Lakeshore Cryotronics for all the temperature readout and ADR control. Smaller

and more power efficient commercial equipment already exists, as shown in Figure 2.21. Replacing the magnet power supply also should not present many difficulties. DARKNESS uses a general purpose power supply with a lot of mass consisting mostly of a large transformer for converting 110 VAC wall power to small DC voltages. PICTURE-C will borrow a flight tested design from the PIPER collaboration [74], shown in Figure 2.21, to replace the Kepco. The HEMT power supply will largely replicate the design from DARKNESS, using the same components. Table 2.3 compares the DARKNESS electronics to those of PICTURE-C. Replacements for the 10 MHz rubidium clock for timing synchronization, a network switch for transferring data from the readout electronics, and control computers have not been selected yet, but will use commercial products.

2.8 Readout: Generation 3

The MKID readout electronics for DARKNESS and MEC are described at length in [55]. Briefly, they are comprised of 10 (20 for MEC) identical units, each consisting of a ROACH2 board designed by the CASPER collaboration, an ADC/DAC board designed by Fermilab, and an RF/IF board also designed by Fermilab. A single unit covers 2 GHz of bandwidth, such that two units together can cover the 4-8 GHz band of DARKNESS/MEC MKID array.

While this “Gen2”⁶ readout is stable and reliable, it consumes approximately 1.4 kW when reading a full five-feedline MKID array, a large load for a balloon-borne observatory. This is problematic for two reasons. First, it would require a large number of batteries during the flight, and second, waste heat is removed from the balloon gondola via thermal radiation and not with air-cooling. A blackbody panel maintaining a temperature of 300 K and emitting into space (the panel has to face upward) would have a surface area of

⁶The ARCONS readout was Gen1. See [70] for details.




Figure 2.20: The DARKNESS rack used mostly commercially available equipment for the ADR power supply and temperature monitors, networking, and precision timing reference. The HEMT power supply and MKID readout crate were assembled in house. PICTURE-C is able to use this rack while assembly of its own electronics is ongoing. Once finished, it will be battery powered, significantly smaller, and consume less power.

Table 2.3: Summary of the sizes, masses, and power consumptions of the electronics supporting the operation of DARKNESS and PICTURE-C. All dimensions given in cm, masses in kg, and powers in W. The computers for the readout and fridge control are not included here. [†]Estimates. [‡]Measured while regulating magnet current, but not while ramping, which would add an estimated 10W. *Maximum power estimate during ramp is battery voltage time magnet current, 28 V×10 A, lasting for ~ 1 hr. When not ramping, power consumption will be much lower, estimated at 10 W. [§]These do not include the chassis and cooling hardware.

DARKNESS						PICTURE-C					
Item	W	L	H	Mass	Power	Item	W	L	H	Mass	Power [†]
Kepeco ADR power supply	32	44	12	23	} 312 [‡]	PIPER ADR power supply [†]	6	12	4	0.5	280/10*
LS372 temp monitor	44	37	9	6.8		SRS 921 temp monitor	8	9	18	1	10
LS336 temp monitor	44	37	9	7.8		SRS 960 PID controller	8	9	18	1	10
10 MHz ref	44	36	4	3.3		LS240 temp monitor	2	10	12	.12	2.4
HEMT power supply	44	37	9	3.5		HEMT power supply [†]	20	20	30	3	5
Network switch	36	48	4	5.4		Network & 10 MHz	-	-	-	-	-
Gen2 readout crate	48	69	46	71	1430	Gen3 readout crate [†]	30	15	32	7	250
Total	120.8[§]				1742	12.6[§]				560/290	



Figure 2.21: Left: A Lakeshore 240-2P will monitor the nitrogen and helium stage temperatures. Center: A SRS SIM921 AC resistance bridge will measure the temperature of the device stage, while an SRS SIM960 (not pictured, but looks similar to the 921) provides PI feedback to the ADR power supply. Right: The PIPER power supply will source current for cycling and regulating the ADR.

about 3 m^2 , large enough to potentially interfere with other systems during the flight.

Because of this and future MKID projects demanding more pixels, we started working toward a “Gen3” readout using more advanced microwave components to decrease the power load, physical footprint, and cost per pixel. This new device combines an FPGA with multiple analog to digital converters and digital to analog converters, streamlining the signal flow and dramatically reducing the power requirement from 1740 W to $\approx 290 \text{ W}$. Preliminary testing with Xilinx ZCU111 evaluation boards⁷ has shown promising results, and we remain confident that a fully functional Gen3 readout will be ready in time for the PICTURE-C flight in 2020.

⁷See <https://www.xilinx.com/products/boards-and-kits/zcu111.html#overview>.

Chapter 3

Stochastic Speckle Discrimination

3.1 Speckles and the Modified Rician Distribution

Coherent light passing through a distorting medium or scattering from a rough surface produces granular interference patterns where the individual grains are referred to as speckles. Speckles set the noise floor of high contrast observations, because as we saw in Chapter 1, starlight is much more intense than light from companions like planets, and atmospheric and optical aberrations distort the stellar wavefront which can obscure planets.

With the ubiquity of consumer grade laser pointers, it is an easy experiment to observe speckle patterns that arise from coherent, monochromatic light as shown in Figure 3.1. Interestingly, these patterns change depending on the relative position of the laser source, the scattering surface, and the observer's location. Researchers first began investigating speckle phenomena in the 1960's when continuous wave lasers became commercially available, and a large body of work has since been dedicated to understanding their causes and implications.

Speckles arise from the interference of many electric field phasors with random rela-

tive phases. A laser striking a piece of paper is an excellent example of such a situation, because an incident, coherent beam hits a material with a surface roughness comparable to or greater than the wavelength of the laser. The small variations behave as tiny scatterers, and because the laser is *temporally* coherent, the electric fields from the scatterers add with each other to produce regions of constructive or destructive interference.




Figure 3.1: A speckle pattern is produced when coherent light scatters off of rough surfaces. Reproduced from [30].




Figure 3.2: The modified rician distribution describes the statistics of the intensity of a constant phasor plus a sum of a random walk. I_s is the mean of the magnitude of the sum of random walk phasors. Adapted from [75].

Light can also be *spatially* coherent, as is the case with wavefronts from distant stars before they get corrupted by Earth's atmosphere. To understand why starlight originating from a thermal source becomes coherent we turn to the Cittert Zernike theorem,

which states that the spatial coherence area is given by

$$A_c = \frac{L^2 \lambda^2}{\pi d^2}$$

where d is the diameter of the source, L is the distance from the telescope to the source, and λ is the observation wavelength. Wavefronts smooth out and flatten as the distance from the source increases. To give an example, a 10 meter diameter telescope observing at $1 \mu\text{m}$ sets a threshold of $L/d \gtrsim 1.6 \times 10^7$. This condition is easily satisfied for large A-type stars like Vega ($L/d \approx 7.2 \times 10^7$) and HR8799 ($L/d \approx 5.1 \times 10^8$), meaning that wavefronts arriving from those stars will be flat across the aperture of a 10 meter (or even 30 meter) diameter telescope.

Turbulence and other nonuniformities in Earth's atmosphere distort stellar wavefronts so that they are no longer flat and the phase becomes position dependent. As described in Section 1.4.1, the length scale of the distortions depends primarily upon r_0 which is much smaller than the mirror diameters of both Subaru and Palomar at $1 \mu\text{m}$. Even under the best seeing conditions, there will be hundreds of r_0 diameter sub-apertures across the primary mirrors as shown in Figure 3.3.

Each sub-aperture on the primary mirror contributes an electric field phasor with a random phase, and their sum in the focal plane creates a speckle pattern. The random phasors also interfere with the deterministic diffraction pattern defined by the shape of the pupil (the primary mirror). Goodman showed in 1975 [30] that the probability of obtaining an intensity I when a deterministic phasor is added to the sum of a large number of independent random phasors is described by the ‘‘modified rician’’ (MR) distribution,

$$p_{\text{MR}} [I|I_c, I_s] = \frac{1}{I_s} \exp \left[-\frac{I + I_c}{I_s} \right] I_0 \left[\frac{2\sqrt{II_c}}{I_s} \right], \quad (3.1)$$




Figure 3.3: The telescope primary mirror with diameter D has many small subapertures the size of r_0 . The phase of the wavefront is relatively constant for each subaperture, but their relative phases are uniformly distributed from 0 to 2π .

where I_c is the intensity of the deterministic (“constant”) phasor, I_s is the average intensity of the random (“speckle”) phasor sum alone, and I_0 is the zero-order modified Bessel function of the first kind. Examples of MR distributions with different values of I_c are shown in Figure 3.4. The mean intensity and variance are

$$\langle I \rangle = I_c + I_s \quad (3.2)$$

$$\sigma_I^2 = I_s^2 + 2 I_c I_s. \quad (3.3)$$




Figure 3.4: Modified Rician distributions with different values of I_c with $I_s = 0.1$ held constant in arbitrary units. I is a continuous parameter. Reproduced from [76]

With imaging optics, both I_c and I_s become functions of position in the focal plane. I_c includes the PSF, while I_s is the seeing halo. When $I_c \neq 0$, such as on top of an Airy ring as indicated in Figure 3.5, the speckles are said to be “partially developed” because there is a deterministic phasor, whereas when $I_c = 0$ at Airy pattern zeros, the intensity distribution follows a negative exponential and the speckles are “fully developed” because

there is no deterministic component,

$$p[I|I_s] = \frac{1}{I_s} \exp\left[-\frac{I}{I_s}\right]. \quad (3.4)$$

In this case the mean and variance are $\langle I \rangle = I_s$ and $\sigma_I^2 = I_s^2$.




Figure 3.5: Cross section of the surface intensity profiles of an Airy pattern and seeing halo. With partial AO correction I_c will be dominated by the Airy profile, such that I_c is zero at Airy minima. I_s is the seeing halo shown by the blue curve.

So far we've only considered the intensity distribution at a fixed point in time. As we move forward in time, the random phasors at the primary mirror will evolve with timescales set by r_0 , D and wind speed. After a decorrelation time τ , a new realization of speckles will form in every part of the focal plane image. If we assume that I_c and I_s are constant in time (which is not actually a perfect assumption, as we'll discuss later), then each new speckle realization has an intensity drawn from Equation 3.1. In other words, Goodman's derivation was originally for the spatial distribution of speckle intensity, but we can exploit a symmetry argument and posit that MR statistics also

govern the temporal distribution of speckle intensity.

Gladysz and Christou proposed in 2008 [36] to use fast exposures to measure this temporal variability and fit the intensity distributions for every location in the focal plane to extract I_c and I_s .

3.2 Poisson Limited Modified Rician with Binned Lightcurves

In a low light regime where shot noise becomes significant we must adjust our model of the intensity distribution accordingly. To do this, we apply a Poisson-Mandel transformation to Equation 3.1 [76, 32]. The probability of receiving n photon counts in a bin given I_c and I_s

$$p_{\star}[n|I_c, I_s] = \int_0^{\infty} \frac{I^n}{n!} \exp(-I) p_{\text{MR}}(I) dI \quad (3.5)$$

$$p_{\star}[n|I_c, I_s] = \frac{1}{I_s + 1} \left(1 + \frac{1}{I_s}\right)^{-n} \exp\left(-\frac{I_c}{I_s}\right) L_n\left(\frac{-I_c}{I_s^2 + I_s}\right) \exp\left(\frac{I_c}{I_s^2 + I_s}\right),$$

where L_n is the n^{th} Laguerre polynomial¹. Examples of this distribution are shown in Figure 3.6. The mean intensity is unchanged, but the variance now includes an extra term from Poisson statistics $\sigma_p^2 = I_c + I_s$ such that Equation 3.3 becomes

$$\sigma_I^2 = I_s^2 + 2 I_c I_s + I_c + I_s. \quad (3.6)$$

¹[76] gives this result with the Kummer confluent hypergeometric function instead of the Laguerre polynomial. The relationship between these functions is ${}_1F_1(n+1, 1, x) = L_n(-x) \exp(x)$. It was found early on when we began examining binned vs bin-free statistics that the `scipy` implementation of ${}_1F_1$ has a much slower execution time than $L_n(-x) \exp(x)$. For reference, wikipedia's Laguerre polynomial definition is $L_n(x) = \frac{e^x}{n!} \frac{d^n}{dx^n} (e^{-x} x^n) = \frac{1}{n!} \left(\frac{d}{dx} - 1\right)^n x^n$.




Figure 3.6: The modified rician distribution in the low light, photon counting regime suffers a Poisson-Mandel transformation. n is number of photons detected in a light curve bin. I_s has the same value for each curve, and four representative values of I_c show how the distribution changes. $I_c = 0$ is the situation where the speckle pattern is fully developed, and the intensity has no constant component. The units are counts/bin.

3.3 Likelihood Function of Poisson Limited Modified Rician with No Companion

If we assume that only starlight is contributing to the focal plane intensity, the likelihood of getting n_i counts in the i th bin of the light curve, given I_c and I_s is $\mathcal{L}_i = p_\star[n_i|I_c, I_s]$. We can write the likelihood of obtaining the entire light curve as the product of all the individual likelihoods,

$$\mathcal{L} = \prod_i^N \mathcal{L}_i$$

$$\ln \mathcal{L} = \sum_i^N \ln \mathcal{L}_i$$

$$\ln \mathcal{L} = N \left(\ln \left[\frac{1}{I_s + 1} \right] - \frac{I_c}{I_s} \right) + \sum_i^N \left(n_i \ln \left[\frac{I_s}{1 + I_s} \right] + \ln \left[L_{n_i} \left(-\frac{I_c}{I_s^2 + I_s} \right) \exp \left(\frac{I_c}{I_s^2 + I_s} \right) \right] \right)$$

$$\ln \mathcal{L} = N \left(\ln \left[\frac{1}{I_s + 1} \right] - \frac{I_c}{I_s} \right) + N \left(\frac{I_c}{I_s^2 + I_s} \right) + \sum_i^N \left(n_i \ln \left[\frac{I_s}{1 + I_s} \right] + \ln \left[L_{n_i} \left(-\frac{I_c}{I_s^2 + I_s} \right) \right] \right) \quad (3.7)$$

3.4 Adding a Companion to the Starlight in Binned Lightcurves

So far, all of this discussion has concerned only light coming from the host star. I_c is the constant part of the star's PSF, and I_s is the speckle part. If our optical system were perfect and the atmosphere introduced no aberrations, we would only see I_c in the image plane and I_s would be zero. A sub-stellar companion's light is not included in I_c

or I_s , so now let's try to figure out how it fits in to our binned speckle analysis.

We'll call the intensity of the planet (sub-stellar companion) I_p^2 . I_p is a coherent, off-axis PSF that is *incoherent* with the starlight. We are still in the photon counting regime, so the photon arrival times from the planet should follow a Poisson process. At the planet's location, we should expect to see the sum of the planet light and starlight. The probability distribution of the sum of two or more independent distributions is the *convolution* of the independent distributions. This means that we should expect that the photons at the planet's location will follow an intensity distribution that is the convolution of a poisson limited modified rician with a poisson distribution.

The planet light should follow

$$p_p[m|I_p] = e^{-I_p} \frac{I_p^m}{m!}$$

where I_p is the mean count rate and m is the number of counts per bin. For convenience, we'll rewrite Equation 3.5 here

$$p_\star[n|I_c, I_s] = \frac{1}{I_s + 1} \left(1 + \frac{1}{I_s}\right)^{-n} \exp\left(-\frac{I_c}{I_s}\right) L_n\left(\frac{-I_c}{I_s^2 + I_s}\right) \exp\left(\frac{I_c}{I_s^2 + I_s}\right)$$

Now we calculate the distribution that we would actually measure. The probability of getting k counts will depend on the number of counts n received from p_\star , as well as the number of counts m received from p_p . $k = m + n$, so the result is

$$p[k|I_p, I_c, I_s] = \sum_{m=0}^k p_p[m|I_p] p_\star[k - m|I_c, I_s] \quad (3.8)$$

²If “planet” seems too presumptuous, “Poisson” also works.

This is the discrete convolution of p_p and p_* . Plugging in p_p and p_* .

$$p[k|I_p, I_c, I_s] = \sum_{m=0}^k \exp(-I_p) \frac{I_p^m}{m!} \frac{1}{I_s + 1} \left(1 + \frac{1}{I_s}\right)^{-(k-m)} \times \exp\left(-\frac{I_c}{I_s}\right) L_{k-m}\left(\frac{-I_c}{I_s^2 + I_s}\right) \exp\left(\frac{I_c}{I_s^2 + I_s}\right) \quad (3.9)$$

We can pull some things outside the sum to get

$$p[k|I_p, I_c, I_s] = \frac{1}{I_s + 1} \exp\left(-I_p - \frac{I_c}{I_s} + \frac{I_c}{I_s^2 + I_s}\right) \sum_{m=0}^k \frac{I_p^m}{m!} \left(1 + \frac{1}{I_s}\right)^{-(k-m)} L_{k-m}\left(\frac{-I_c}{I_s^2 + I_s}\right)$$

$$p[k|I_p, I_c, I_s] = \frac{1}{I_s + 1} \exp\left(-I_p - \frac{I_c I_s}{I_s^2 + I_s}\right) \sum_{m=0}^k \frac{I_p^m}{m!} \left(1 + \frac{1}{I_s}\right)^{-(k-m)} L_{k-m}\left(\frac{-I_c}{I_s^2 + I_s}\right) \quad (3.10)$$

The likelihood of the i th bin of a light curve is




Figure 3.7: $p[k|I_p, I_c, I_s]$ for different values of I_p with I_c and I_s held constant.

$$\mathcal{L}_i = \sum_{m=0}^{k_i} p_p[m|I_p] p_\star[k_i - m|I_c, I_s], \quad (3.11)$$

$$\mathcal{L}_i = \frac{1}{I_s + 1} \exp\left(-I_p - \frac{I_c I_s}{I_s^2 + I_s}\right) \sum_{m=0}^{k_i} \frac{I_p^m}{m!} \left(1 + \frac{1}{I_s}\right)^{-(k_i-m)} L_{k_i-m}\left(\frac{-I_c}{I_s^2 + I_s}\right). \quad (3.12)$$

Recalling that the likelihood of the entire light curve is the product of all the individual likelihoods and taking the logarithm,

$$\ln \mathcal{L} = \sum_i \ln \mathcal{L}_i$$

$$\ln \mathcal{L} = \sum_i \ln (p[k_i|I_p, I_c, I_s]) \quad (3.13)$$

Because $p[k_i|I_p, I_c, I_s]$ is a sum, there is little we can do to simplify Equation 3.13.

$$\ln \mathcal{L} = \ln \left[\frac{1}{I_s + 1} \right] + \left(-I_p - \frac{I_c I_s}{I_s^2 + I_s} \right) \sum_i \ln \left[\sum_{m=0}^{k_i} \frac{I_p^m}{m!} \left(1 + \frac{1}{I_s}\right)^{-(k_i-m)} L_{k_i-m}\left(\frac{-I_c}{I_s^2 + I_s}\right) \right] \quad (3.14)$$

3.5 Bin-Free SSD

The goal of SSD is to exploit the temporal variability of speckles to remove them and improve the signal to noise of a substellar companion detection. The discussion thus far has only considered intensities that have been binned in time, as this occurs with traditional cameras that have to integrate light in order to build up enough signal to measure intensity accurately. This necessarily leads to some amount of information loss for the temporal behavior, even when one is able to sample the intensity faster than the

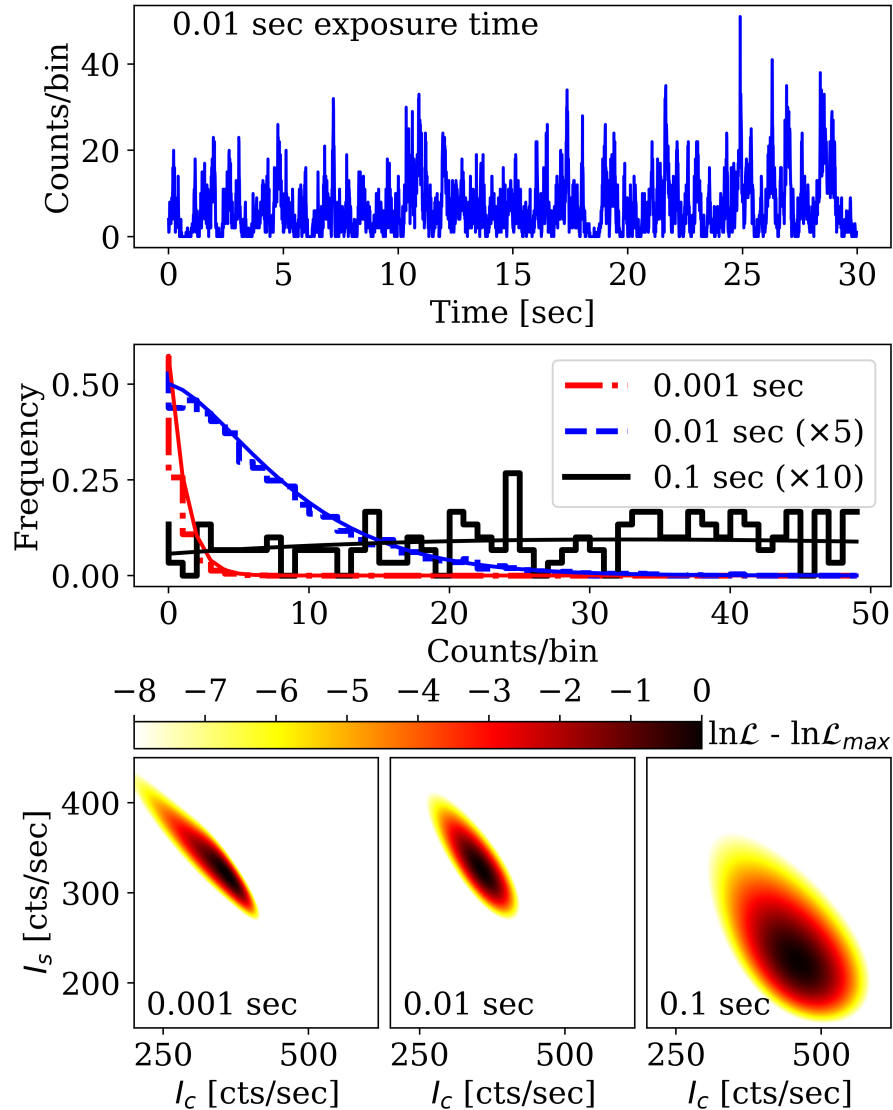


Figure 3.8: Top: a 30 s mock photon list simulated using the parameters $I_c = 300$, $I_s = 300$, $I_p = 0$ photons/s with $\tau = 0.1$ s, is binned into 0.01 s exposures to form a light curve. Middle: The same photon list is binned by three different exposure times and used to plot intensity histograms. The best-fit modified Rician functions are overlotted. 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.

changes in intensity.

Conventional imaging cameras in astronomy typically have to integrate over an exposure time because of technological limitations, but it is possible to count individual photons. Indeed, researchers have used photomultiplier tubes to do so since the 1930's, but improvements in fabrication techniques have only recently made it possible to build compact photon counting imagers, including MKIDs, with good signal to noise. Photon counting detectors can potentially capture all of the temporal variability and thus provide as much information as possible to an SSD algorithm. Specifically, we can use the differences in photon arrival times to calculate the likelihood of a photon list given a set of model parameters, and then adjusting the parameters until the likelihood is maximized. As before, the parameters in the model we will use are I_c , I_s , and I_p . However, the likelihood function that we wish to use gives the probability of measuring a set of inter-photon arrival times, δt_i , instead of a light curve containing the number of photons per exposure time, k_i .

3.5.1 Bin-Free Likelihood Function

When using conventional imaging cameras for high contrast astronomy, the exposure times are set by the constraints of the camera and the requirements of the science case, and cannot be made shorter after the data is collected. Combining frames could make exposures longer. This can significantly affect the results of SSD. If one is indeed constrained to using exposures, care should be taken to ensure that they are not too short or too long. If the exposures are too short there won't be enough light to provide a good measurement of the intensity, and if they are too long the changes in speckle intensity will be averaged out. The timescale constraints for doing SSD with exposures are

$$\langle \delta t \rangle \ll t_{\text{exp}} \ll \tau, \quad (3.15)$$

where $\langle \delta t \rangle$ is the mean inter-photon arrival time, t_{exp} is the camera exposure time, and τ is the speckle decorrelation time.

When using a photon counting instrument such as one equipped with MKIDs, it behooves us to avoid using traditional exposures entirely and remove one of the timescale constraints so that we only need to ensure that

$$\langle \delta t \rangle \ll \tau. \quad (3.16)$$

In other words, we want to make sure that many photons are detected before a speckle's intensity changes very much.

Here we highlight some of the important points of the bin-free log-likelihood derivation from an unpublished internal note by Tim Brandt [77]. First, we begin with the probability of an inter-photon arrival time δt given a mean photon count rate I

$$p[\delta t|I] = I e^{-I\delta t}. \quad (3.17)$$

The probability getting δt at a fixed point in time is

$$\int_0^\infty p[\delta t|I] p[I] dI, \quad (3.18)$$

but we want to know the relative probability of getting δt in a photon list. As I increases, there are more chances to measure δt , meaning that the integrand should be proportional to I ,

$$\int_0^\infty p[\delta t|I] p[I] I dI. \quad (3.19)$$

If I follows a Poisson distribution, $p[I]$ would be a delta function and the integral would evaluate to Equation 3.17. However, the intensity varies according to Equation 3.1, which

we will restate here for convenience,

$$p_{MR}[I|I_c, I_s] = \frac{1}{I_s} \exp\left[-\frac{I + I_c}{I_s}\right] I_0\left[\frac{2\sqrt{II_c}}{I_s}\right],$$

where the zero-order modified Bessel function of the first kind is

$$I_0[x] = \sum_{k=0}^{\infty} \frac{(\frac{1}{4}x^2)^k}{(k!)^2}.$$

Evaluating the integral and simplifying yields

$$p[\delta t|I_c, I_s] = \frac{1}{(I_c + I_s)(1 + I_s\delta t)^5} \times \left\{ (I_c^2 + 4I_c(I_s + I_s^2\delta t) + 2(I_s + I_s^2\delta t)^2) \exp\left[\frac{-I_c\delta t}{1 + I_s\delta t}\right] \right\}. \quad (3.20)$$

Photon counting detectors typically have some finite deadtime after registering a photon event, and MKIDs are no exception. Accounting for this entails changing the lower limit on the integral in Equation 3.19 to the deadtime t_0 . After evaluating we obtain

$$p[\delta t|I_c, I_s] = \frac{(1 + I_s t_0)^3}{(I_c + I_s + I_s^2 t_0)(1 + I_s\delta t)^5} \times \left\{ (I_c^2 + 4I_c(I_s + I_s^2\delta t) + 2(I_s + I_s^2\delta t)^2) \exp\left[\frac{-I_c\delta t}{1 + I_s\delta t}\right] \right\} \times \exp\left[\frac{I_c t_0}{1 + I_s t_0}\right]. \quad (3.21)$$

Going forward we will only show the expressions that include deadtime, but the simpler expressions can be found by setting $t_0 = 0$.

$p[\delta t|I_c, I_s]$ is the likelihood \mathcal{L}_i of obtaining a single value of δt , and so the likelihood \mathcal{L} of obtaining an entire set of δt 's is the product of their individual likelihoods. For computational reasons, it is convenient to calculate $\ln \mathcal{L}$, which is the sum of the individual $\ln \mathcal{L}_i$'s.

So far, our calculations have only concerned the statistics of starlight intensity. We must also include the planet (or other companion), I_p . Suppose we have two sources with intensities I_A and I_B where I_A is starlight following a modified rician distribution and I_B is planet light with a constant intensity. If we observe a photon, the probability of obtaining a second photon that came from I_A after a time δt is

$$p[\delta t|I_A] = I_A \exp[-(I_A + I_B)\delta t]. \quad (3.22)$$

A similar expression gives the probability of detecting a second photon coming from I_B . Normalizing these probabilities yields identical expressions, meaning that the distribution of inter-photon arrival times is the same whether the second photon is coming from I_A or I_B . The integral we need to evaluate becomes

$$p[\delta t|I_c, I_s, I_p] = \int_0^\infty p[\delta t|I + I_p] p[I] (I + I_p) dI. \quad (3.23)$$

Defining

$$u_i = \frac{1}{1 + I_s \delta t_i}$$

and

$$u_{\max} = \frac{1}{1 + I_s \delta t_0},$$

we obtain the following expression for the normalized likelihood of the i th δt ,

$$\mathcal{L}_i = \frac{\exp \left[(u_i - 1) \left(\frac{I_p}{I_s u_i} + \frac{I_c}{I_s} \right) \right] \{ I_c^2 u_i^5 + 4I_c I_s u_i^4 + (2I_s^2 + 2I_p I_c) u_i^3 + 2I_p I_s u_i^2 + I_p^2 u_i \}}{(I_c u_{\max}^3 + I_s u_{\max}^2 + I_p u_{\max}) \exp \left[\frac{(u_{\max} - 1)(I_p + I_c u_{\max})}{I_s u_{\max}} \right]}, \quad (3.24)$$

$\mathcal{L} = \prod_i \mathcal{L}_i$, so the log-likelihood for a set of N δt 's is

$$\begin{aligned} \ln \mathcal{L} = & \sum_i (u_i - 1) \left(\frac{I_p}{I_s u_i} + \frac{I_c}{I_s} \right) \\ & + \sum_i \ln [I_c^2 u_i^5 + 4I_c I_s u_i^4 + (2I_s^2 + 2I_p I_c) u_i^3 + 2I_p I_s u_i^2 + I_p^2 u_i] \\ & - N \frac{(u_{\max} - 1)(I_p + I_c u_{\max})}{I_s u_{\max}} - N \ln [I_c u_{\max}^3 + I_s u_{\max}^2 + I_p u_{\max}]. \end{aligned} \quad (3.25)$$

Now that we have the expression for $\ln \mathcal{L}$, we can find the values of I_c , I_s , and I_p that maximize it for a given dataset of δt 's. $\ln \mathcal{L}$ is a well behaved function that has only one maximum, so finding the location of the peak in the I_c , I_s , I_p parameter space is fairly straightforward.

While of little use when analyzing data, it may be helpful to visualize the shape of the likelihood function when I_c , I_s , and I_p are held constant and δt is varied. \mathcal{L} decreases monotonically with δt , with a shape loosely resembling an exponential, as shown in Figure 3.9. It shows the probability of detecting zero photons for a time δt when an initial photon is detected at time $t = 0$. At small δt the likelihood is large, meaning that it's very likely that a second photon will *not* be detected at time $t < \delta t$. For larger δt it is less likely that a second photon will not be detected, implying that detecting a second photon at $t < \delta t$ is very likely.

3.5.2 Including a Prior on I_c

We can include prior knowledge of the stellar PSF and quasistatic speckles if we know the focal plane intensity without atmospheric speckles or a companion. The shape of the PSF is determined by the telescope pupil and subsequent optics (including any coronagraph), so one can in principle use a “perfect” PSF as part of the prior for I_c . More difficult is the task of including quasistatic speckles, which cannot be determined




Figure 3.9: Normalized \mathcal{L} vs. δt . \mathcal{L} is the probability of a photon arriving $> \delta t$ after the previous one. An exponential, corresponding to a Poisson process with a mean equal to $1/(I_c + I_s + I_p)$ is shown for comparison.

from a model of the optics in the system but instead would need to be measured.

Such a measurement could be done in much the same way that Reference star Differential Imaging (RDI, [78]) is used to subtract a reference PSF from direct images to reveal substellar companions. However, instead of subtracting the reference PSF from the raw data, we would instead use it as a prior for I_c to improve the estimate of I_p .

To understand how to include a prior, we can start with Bayes' theorem,

$$p(M|D) = \frac{p(D|M)p(M)}{p(D)}, \quad (3.26)$$

where $p(M|D)$ is the probability of the model given the data, also known as the posterior distribution, $p(D|M)$ is the probability of the data given the model, also known as the likelihood, $p(M)$ is the probability of the model, also known as the prior, and $p(D)$ is

the difficult-to-calculate probability of the data, which we will treat as a normalization constant and will have no effect on the result. If we assume that our measurement of the prior is normally distributed, the logarithm of 3.26 becomes

$$\ln(\text{posterior probability}) \propto \ln \mathcal{L} + \frac{1}{2} \left(\frac{I_c - I_{c,\text{prior}}}{\sigma_{I_{c,\text{prior}}}} \right)^2, \quad (3.27)$$

where $I_{c,\text{prior}}$ is the best estimate for the value of I_c and $\sigma_{I_{c,\text{prior}}}$ is the uncertainty. Our goal then is to find the parameters that maximize $\ln(\text{posterior probability})$.

While this appears to be a promising avenue for improving signal to noise on measurements of I_p , we will see in the following sections, particularly in Table 3.1, that the improvements depend on the relative values of I_c and I_s and should not be applied universally. Given this, as well as the difficulty associated with obtaining $I_{c,\text{prior}}$, the case for including a prior is somewhat weak.

3.6 Finding the Maximum Likelihood

The MKID data pipeline is written in python and is freely available to the public³. As a group we have written extensive code for calibrating and reducing the data from the various MKID instruments, including a module for SSD. Broadly, a list of photon timestamps from each functional pixel is passed to the function and it returns the best fit parameters of I_c , I_s , and I_p .

We implemented this with the help of the `scipy.optimize.minimize` module, where we actually minimize $-\ln \mathcal{L}$. Depending on the method used to find the peak, one can also specify the Jacobian and Hessian matrices to decrease the convergence time. The Newton Conjugate-Gradient method performs quite well for the bin-free likelihood,

³<https://github.com/MazinLab/MKIDPipeline>

usually finding the optimal parameters of a 30 second dataset with a mean count rate of $\langle I \rangle \sim 360$ cps within 50 ms as shown in Figure 3.10. The number of floating point operations used to calculate $\ln \mathcal{L}$ scales linearly with the number of photons in a dataset, so photon lists with higher count rates take longer to process.




Figure 3.10: The Newton Conjugate Gradient minimization routine usually finds the optimal values of I_c , I_s , and I_p within 50 ms for $\langle I \rangle$.

The same approach can be used for finding parameters that maximize the posterior distribution (Equation 3.27). It is necessary only to include the extra term with the log-likelihood, and the changes to the derivatives are simple.

3.7 To Marginalize or Not to Marginalize?

Maximum likelihood searches for optimal parameters often involve marginalizing, or integrating, out the parameters in the likelihood function that aren't needed. In our case we are after the “best” value of I_p while I_c and I_s are to be discarded. In the previous

section we assumed that the “best” I_p was that which corresponded to the peak of the likelihood function in Equation 3.25. It’s conceivable that marginalizing the likelihood over I_c and I_s may produce a different result for I_p than simply locating the peak in the phase space spanned by I_c , I_s , and I_p .

Testing this is straightforward, but computationally intensive: compare the distributions of I_p corresponding to the maximum likelihood and the maximum marginalized likelihood. The result, shown in Figure 3.11, is that any difference between the distributions is small. This is fortunate because the time required to map out the likelihood function in three dimensions, even when using algorithms to avoid doing calculations in regions of parameter space yielding negligibly small values of the likelihood, will be orders of magnitude longer.




Figure 3.11: Example showing the distribution of I_p taken from the peak of the likelihood function (black) is the same as that of I_p taken from the peak of the likelihood marginalized over I_c and I_s (blue). Because mapping out the likelihood function in the I_c , I_s , and I_p phase space was extremely computationally expensive compared to using a minimization algorithm, only 100 photon lists were used to produce the blue curve. Shown here are two combinations of I_c and I_s used to produce the photon lists.

3.8 Signal Detection Theory Overview

To characterize the performance of this SSD technique, we turn to signal detection theory, which provides a way to compare and quantify rates of false positive vs true positive detections. As we will see later, this is useful for directly comparing binned vs. bin-free analyses, as well as informing the completeness of exoplanet observations.

If we assume that a signal is present (H_1 hypothesis) and succeed in detecting it, we have a True Positive result, while if we fail in detecting it we have a False Negative result. Alternatively, if we assume that no signal is present (the H_0 hypothesis) and achieve a detection we have a False Positive result, and if we detect nothing we have a True Negative result. We define the True Positive Fraction (TPF) and False Positive Fraction (FPF) as

$$\text{TPF} = \int_{\xi}^{\infty} p(x|H_1)dx$$

$$\text{FPF} = \int_{\xi}^{\infty} p(x|H_0)dx,$$

where ξ is the detection threshold and $p(x|H)$ is the normalized distribution of intensity measurements. We assume here that p is normalized such that the total area under each curve is unity. These definitions assume that we have distributions of many measurements both with and without an actual signal. Our goal is to approach $\text{TPF} = 1$ and $\text{FPF} = 0$ so that as many positive detections as possible are attributable to true signals.

Figure 3.12 illustrates hypothetical distributions obtained under signal absent and signal present hypotheses. The FPF is the area under the curve to the right of the threshold in the left panel, and similarly the TPF is the area under the curve to the right of the threshold in the right panel. The noise is the width of the H_0 distribution and the signal is the separation between the means of the H_1 and H_0 distributions. We should note that while the widths of the distributions under both hypotheses is the same

in this example, this is not generally true but doesn't change how the FPF and TPF are calculated.




Figure 3.12: Left: A mean-subtracted gaussian noise distribution. The detection threshold was set to 3σ ($x = 3$), so that the noise exceeds the threshold a fraction .001 of the time. Equivalently, the false positive fraction is .001. Right: Intensity distribution of planet with mean at $x = 3$ and the same noise. The same detection threshold leaves half of the area under the curve to the right of the threshold, meaning the true positive fraction is 0.5. This means that a detection threshold of 3σ will only detect half of the planets with this intensity. Reproduced from [79].

The TPF/FPF trade space is often displayed graphically with Receiver Operator Characteristic (ROC) curves as shown in Figure 3.13. A ROC curve plots the TPF against the FPF for a range of thresholds while the signal and the noise are kept constant. The shape of the curve depends on the signal strength and noise, with better SNR suggested by curves that approach $\text{FPF} = 0$ and $\text{TPF} = 1$. ROC curves are generally produced from simulated data and are used to understand the completeness of real data.

3.9 Mock Photon Lists

With the goal of understanding and characterizing the performance of our maximum likelihood approach to separating companion light from stellar light, we also developed code to produce mock photon lists to simulate single pixel data that one of our MKID instruments might produce. The mock photon lists obey the following rules:




Figure 3.13: Example ROC curves generated with mock photon lists. As the signal to noise ratio (SNR) increases, the curve approaches the corner where the TPF = 1 and the FPF = 0.

- The overall intensity is MR distributed, but varies in time continuously with a decorrelation time τ such that $\langle I(t)I(t + \delta t) - \langle I \rangle^2 \rangle \propto \exp[-\delta t/\tau]$. τ is specified by the user.
- I_c , I_s , and I_p are independently specified by the user, such that I_c and I_s govern the stellar intensity and I_p defines the mean count rate of the Poisson source. The total intensity should have an expectation value of $\langle I \rangle = I_c + I_s + I_p$.
- There is a deadtime t_0 after each photon to match an intrinsic deadtime of MKIDs, in which no new photons will be “detected”. t_0 is specified by the user.

The procedure for generating photon timestamps begins with creating a list of random numbers drawn from a Gaussian distribution, which is then passed to an algorithm that rearranges gaussian distributed numbers to be exponentially correlated with their




Figure 3.14: Comparison of the maximum likelihood value distributions for $I_p = 0$ (no signal) and $I_p = 20$ cps for $I_c = 30$ cps and $I_s = 300$ cps. The threshold indicated by the black line is at 18 cps, and the shaded area under each curve divided by the total area under each curve represent the FPF (red) and TPF (blue) for that threshold.

neighbors with a scale set by τ . This serves to make the final intensities correlated in time. The correlated, gaussian distributed numbers are transformed to a uniform distribution ranging from 0 to 1, and finally transformed to a MR distribution. The intermediate step serves to avoid a single complicated transformation in favor of two simpler ones.

The list now defines the mean intensity as a function of time, with timesteps of 200 μs . We have another intermediate step, where a Poisson draw at each timestep for that “instantaneous” intensity determines the number of photons to be placed in each bin. At this point we have a binned light curve, similar to previous sections. Finally, random numbers are drawn from a uniform distribution to determine the placement of each photon within its 200 μs bin (chosen to be $\ll \tau$). At the end, the function returns a sequential list of photon timestamps.

Since MKIDs do not have the same dark current or readout noise as conventional semiconductor detectors we do not add any additional noise. However, to prevent exceedingly high count rates from hot pixels and photon pulse pileup we implement a dead time in the readout firmware so that after a photon is recorded to disk, no subsequent photons may be registered until a predefined time has passed. We emulate this deadtime when generating mock photon lists, with a default value of $t_0 = 10 \mu s$ to match the latest firmware implemented on MEC. This dead time formulation could also be used with quasi-photon counting EMCCDs to account for photon pile-up, so its use is not limited to simulating MKID data.

We have implicitly assumed that I_c , I_s , and I_p are constant for the duration of a mock photon list, and that there is only one speckle decorrelation time. Real high contrast observations will violate all of these assumptions, but the maximum likelihood formalism discussed above can be generalized to include variations in the intensity parameters.

Soummer et. al. [80] and Gladysz et. al. [38] proposed that the intensity of adaptively corrected images is not MR distributed at the core of the PSF. Rather, the peak core

intensity, proportional to the theoretical strelh ratio SR , follows

$$p[S|SR, m] = \frac{m}{2} \frac{\left[\frac{m \ln(S)}{2 \ln(SR)}\right]^{(m-3)/2} \exp\left[-\frac{m \ln(S)}{2 \ln(SR)}\right]}{\Gamma\left(\frac{m-1}{2}\right) |\ln(SR)| S}, \quad (3.28)$$

where S is the instantaneous strelh, m is the number of independent cells in the wavefront SR is related to the statistical variance according to the Maréchal approximation [81], $SR = \exp(-\sigma_\phi^2)$, where σ_ϕ^2 is the pupil phase variance. This is to say that SR , and therefore I_c , is not constant with time. I_p will also follow Equation 3.28 because the quality of the AO correction is the same for companions as for the central star, and I_s will follow a distribution for $1 - SR$ because starlight that doesn't go into the PSF instead goes into the seeing halo. While future high contrast imaging may benefit from including variations in SR , this is beyond the scope of this thesis.

As noted in Section 1.4.1, speckles can have more than one decorrelation time. Atmospheric and deformable mirror (DM) residual speckles could perhaps be modeled by two exponentials with different timescales. The decorrelation time is not a parameter in the bin-free SSD analysis, and so should not affect the result unless $\tau I_s \ll 1$, which is true for DM residuals. This is a harder problem to solve, as the number of photons available depends primarily upon the telescope diameter and the optical throughput. Because our MKID instruments cannot tolerate count rates sufficiently high to sample DM error speckles, including another decorretion time in the code for generating photon lists will not inform this SSD analysis technique. It may be worthwhile if higher count rate detectors are available.

In the end, our goal is to measure the intensity of a planet, I_p , and discard contaminating starlight. Now that we have a way to generate mock photon lists quickly, we can start exploring some of the behaviors of the maximum likelihood approach. Figure 3.15 shows a series of simulations, with the four panels representing various combinations of

I_c , I_s . For each panel, the same four values of I_p were used as the parameters to generate 10^5 mock photon lists, including $I_p = 0$, and then obtain the likelihood estimates for I_p from each photon list.

ROC curves produced from these distributions are shown in Figure 3.16. Unsurprisingly, the signal to noise ratio is best when both I_c and I_s are small. The lower right panels in both Figures illustrate that the maximum likelihood estimates perform poorly when I_c/I_s is large. This can be understood by considering the fundamental assumption of the bin-free SSD technique: there is a deterministic intensity (I_c) being modulated by a random intensity (I_s). When I_c/I_s is large, the modulation is small and the overall intensity distribution approaches Poisson statistics, making it hard to discriminate between the three parameters. We will explore this further below in a discussion of signal to noise.

3.10 Binned vs. Bin-free SSD

The great advantage of using the photon inter-arrival times to calculate a likelihood instead of the number of photons in a given bin is that one doesn't need to choose a bin size. Usually the point is moot because a bin size is chosen for the observer because of the science camera. But when photon counting is an option, fitting the components of the intensity will not require the bin size as an extra parameter.

We can test the hypothesis that bin-free SSD is better at discriminating companions from background by generating ensembles of photon lists as described in the previous section, binning the photon timestamps according to an effective exposure time t_{exp} to form a light curve, finding the combination of I_c , I_s , and I_p that maximizes 3.13, and then repeating these steps with a new t_{exp} as illustrated in Figure 3.8.

ROC curves of I_p distributions for a range of t_{exp} as presented in Figure 3.17 indeed




Figure 3.15: The distributions of the best fit I_p for 10^5 photon lists for each combination of I_c and I_s . With the exception of the lower right panel, the distributions show that I_p can be estimated with good signal to noise. As expected, the distribution moves to the right as the true value for I_p increases.

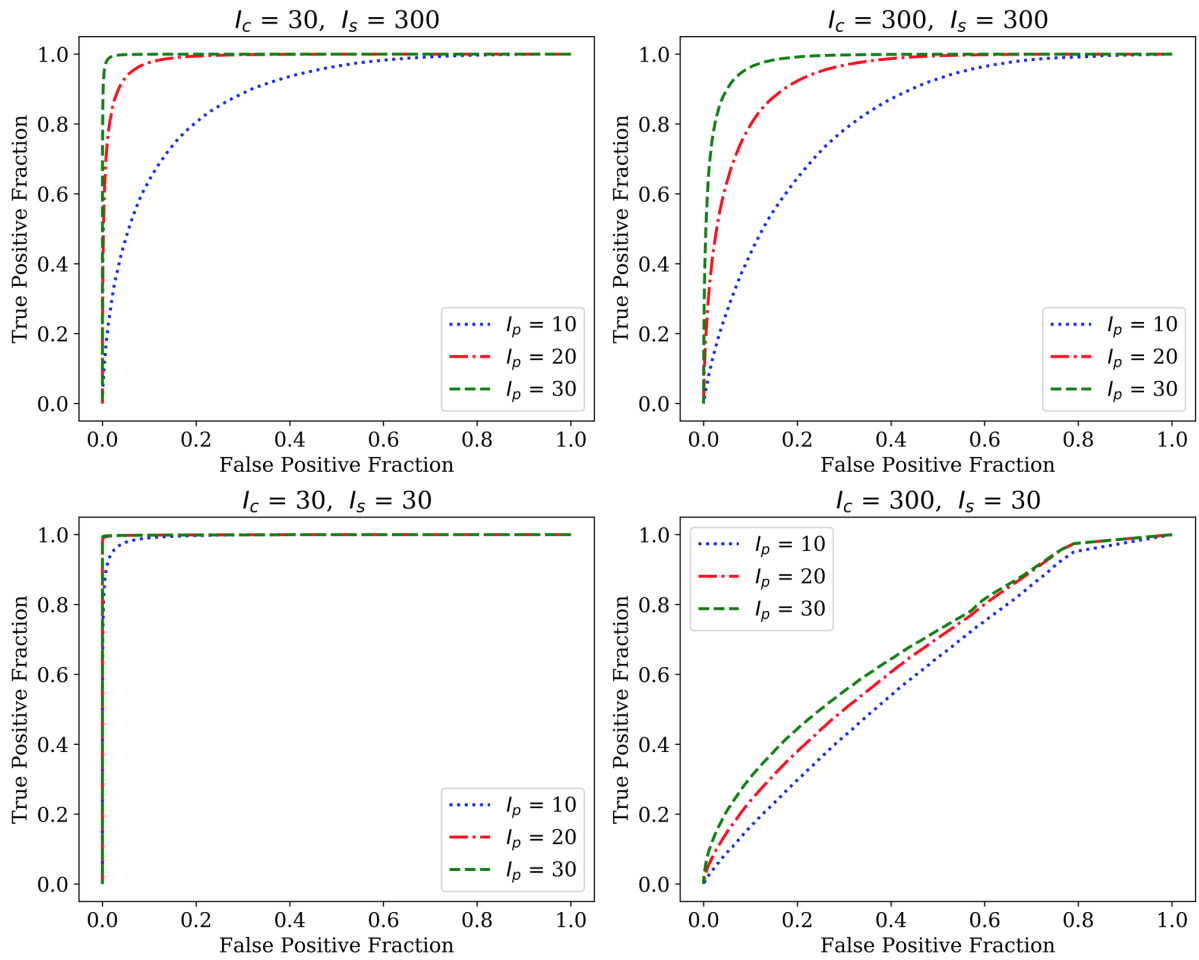


Figure 3.16: The SNR of the photon counting SSD results for I_p suffer noticeably when I_c dominates I_s .

show that the SNR of binned data never exceeds bin-free data. A crude metric for measuring the relative SNR of the groups of ROC curves is the Area Under the ROC Curve, or AUC. While AUC doesn't capture all the details in the changing shape of ROC curves, it can provide useful information about overall trends. The AUC for Figure 3.17 are shown in Figure 3.18. There is a clear peak in the AUC for each combination of I_c and I_s as a function of t_{exp} . This "optimal" t_{exp} is consistently smaller than the decorrelation time of $\tau = 0.1$ sec used to generate the light curves, and also seems to depend on the values of I_c and I_s .

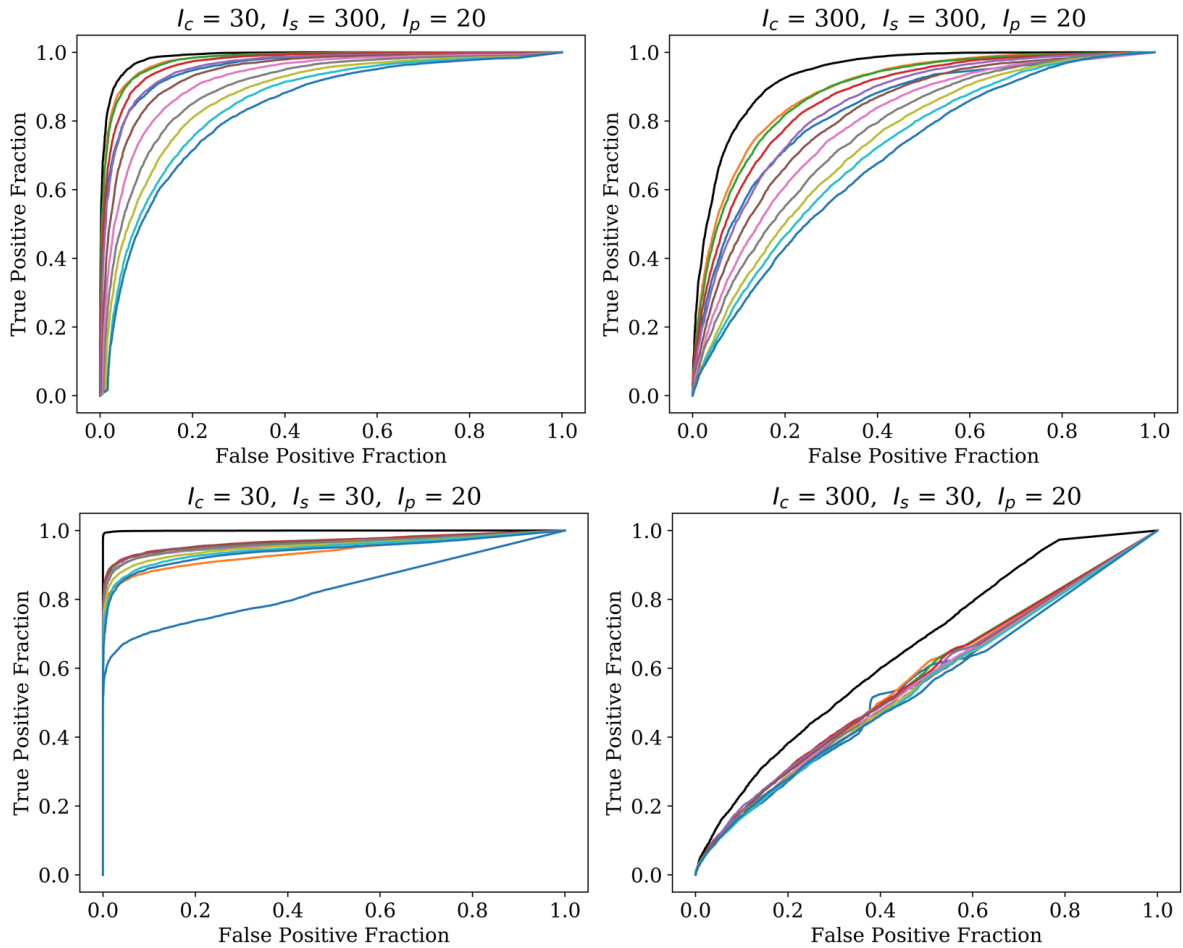


Figure 3.17: The shape of the ROC curves depends on the effective exposure time when using binned SSD. The colored curves indicate different exposure times, while black indicates bin-free SSD.




Figure 3.18: Area Under the ROC Curve gives an indication of relative SNR vs. effective exposure time. The optimal effective exposure times are smaller than the true decorrelation time $\tau = 0.1$ sec used to generate the photon lists, but still depend on the combination of I_c and I_s . This suggests that even if the decorrelation time is known, choosing the optimal bin size may be difficult, and in any case the AUC never exceeds that of bin-free SSD, indicated by the black lines.

3.11 Signal to Noise Ratio

Perhaps the most useful quantity for determining the effectiveness of the SSD analysis described here is the signal to noise ratio, or SNR. For the discussion here, we will consider two SNRs, one for the SSD results for I_p and one for the total intensity. First, we define the signal to noise for the bin-free SSD results as

$$\text{SNR}_{\text{SSD}} = \frac{\langle I_p \rangle - \langle \text{Background} \rangle}{\sigma_{I_p}}, \quad (3.29)$$

where $\langle I_p \rangle$ is the mean SSD estimate of I_p from an ensemble of photon lists with $I_{p,\text{true}} > 0$, $\langle \text{Background} \rangle$ is the mean SSD estimate for I_p from an ensemble with $I_{p,\text{true}} = 0$, and σ_{I_p} is the standard deviation of the SSD estimates for I_p . Equation 3.29 also applies for maximum posterior estimates.

A graphical representation of this is presented in Figures 3.19 and 3.20. The blue curves show the distributions of maximum likelihood estimates for I_p , with faint red lines indicating the means. The difference between the red lines is the numerator of Equation 3.29, and the standard deviation of the solid blue curve ($I_p = 30$) is the denominator.

Having already established that bin-free SSD performs better at extracting I_p from a photon list than SSD with binned data in the previous section, we now wish to quantify the improvement that bin-free SSD may provide over traditional, long exposure imaging. To simulate long exposures we calculate the mean intensities of the 30 second photon lists used for the bin-free analysis, and then examine the distributions of those. This averages over all of the temporal variability within a photon list and replaces it with a single number, which we refer to as the “long exposure intensity”. The signal to noise of

the long exposure intensities is

$$\text{SNR}_{\text{long exposure}} = \frac{\langle I_{\text{tot}} \rangle - I_c - I_s}{\sigma_{I_{\text{tot}}}}, \quad (3.30)$$

where $\langle I_{\text{tot}} \rangle$ is the mean of the distribution of long exposure intensities, I_c and I_s are the true values used to create the photon lists, and $\sigma_{I_{\text{tot}}}$ is the standard deviation of the long exposure intensities. This is the signal to noise ratio we should expect from traditional cameras with integrated exposures that don't have photon counting abilities. The long exposure intensity distributions are shown in black in Figures 3.19 and 3.20. As before, the faint red lines indicate the means, whose difference is the numerator in Equation 3.30. As expected, $I_c + I_s$ is the mean of the semi-transparent black distributions. The denominator of the SNR is the standard deviation of the solid black curves.

In the case where $I_s = 10I_c$ (Figure 3.19), $\text{SNR}_{\text{SSD}} = 3.63$ and $\text{SNR}_{\text{long exposure}} = 1.17$, and we can clearly see that the black distributions overlap more than the blue distributions. However, when $I_c = 10I_s$ we see the signal to noise of SSD is actually worse than with long exposures, $\text{SNR}_{\text{SSD}} = 0.59$ and $\text{SNR}_{\text{long exposure}} = 2.57$. SNR values for various combinations of I_c , I_s , and I_p are compiled in Table 3.1.

We can conclude from these two examples that bin-free SSD doesn't always give better measurements of I_p . The reason for the poor SNR when $I_c > I_s$ is the same as was discussed in Section 3.9: SSD relies on the intensity of a constant phasor being modulated by a random phasor. If this modulation is small compared to the magnitude of the deterministic phasor, then the overall intensity approaches a Poisson distribution and becomes hard to separate into I_c , I_s , and I_p .

In addition to producing large ensembles of photon lists all with the same values of I_c , I_s , and I_p , it is also possible to make simulated high contrast images where the intensities vary with position and perform SSD analysis on those. Figure 3.21 shows




Figure 3.19: The blue histograms (left) show the distributions for I_p obtained from SSD, while the black histograms (right) show the distributions for the total mean intensities $\langle I \rangle$ of the same photon lists. The semi-transparent curves come from photon lists where $I_{p,true} = 0$, while the dark curves are from $I_{p,true} = 30$. The means of each distribution are indicated by the vertical red lines. The SNR for each type of signal (I_p vs. $\langle I \rangle$) can be thought of as the separation of the means divided by the widths of the $I_{p,true} = 30$ distributions. Here, the SNR of the SSD result is 3.63, while the SNR of the total intensity is 1.17. The length of both datasets was 30 seconds. All count rates given in counts per second.

Table 3.1: The signal to noise ratios for maximum likelihood estimates (MLE), maximum posterior estimates (MAP), and for long exposures suggest that SSD estimates for I_p are better than that of long exposures when $I_s \gtrsim I_c$, but are worse when $I_c > I_s$. 10^5 photon lists were generated for each combination of I_c , I_s , and I_p whose values are indicated in the table; the decorrelation time $\tau = 0.1$ second, and each photon list represented 30 seconds of data.

	$I_c = 30, I_s = 30$			$I_c = 30, I_s = 300$			$I_c = 300, I_s = 30$			$I_c = 300, I_s = 300$		
I_p	MLE	MAP	long	MLE	MAP	long	MLE	MAP	long	MLE	MAP	long
10	4.38	4.20	2.28	1.46	1.36	0.39	0.32	0.46	0.85	1.25	0.64	0.24
20	5.37	6.69	4.52	2.62	2.41	0.78	0.49	1.30	1.72	2.04	1.27	0.48
30	5.62	8.83	6.72	3.63	3.33	1.17	0.59	2.20	2.57	2.71	1.84	0.73




Figure 3.20: When $I_c \gg I_s$ the SNR of the SSD result is worse (0.59) than that of the total intensity (2.57) because the I_p distributions have a much larger width than the $\langle I \rangle$ distributions. This kind of situation may be present when the Strehl ratio is very high ($\gtrsim 0.95$) and a coronagraph is not present. Care should be taken to avoid using SSD to analyze data falling into this category. All count rates given in counts per second.

such a simulation with $SR = 0.7$. Here we start with “true” images, where I_c is an Airy pattern, I_s is a Moffat profile with seeing width of $W_{\text{halo}} \sim 0.5''$ given by [82]

$$f_{\text{halo}}(\theta) = \frac{0.488}{W_{\text{halo}}^2} \left[1 + \frac{11}{6} \left(\frac{\theta}{W_{\text{halo}}} \right)^2 \right]^{-11/6}, \quad (3.31)$$

and I_p is a sum of Airy patterns with intensities specified by contrast ratios of 4×10^{-4} , 2×10^{-4} , 1×10^{-4} , and 5×10^{-5} . Maximum likelihood estimates were found for each of the 65×65 pixels, where the PSF, seeing halo, and planet intensities were used to generate independent photon lists. Comparing the image of I_p to that of perfectly subtracting the PSF and halo from the total intensity, i.e. $I - I_c - I_s$, we see that SSD performs better for fainter planets at smaller separations (see panels b and h in Figure 3.21.). We should note that $I_s > I_c$ for every pixel outside of the PSF core, making this simulation favorable for SSD. The values of each at the planet locations are listed in Table 3.2.




Figure 3.21: Previous page: Simulation comparing images of I_p and perfect subtraction of I_c and I_s from the mean intensity. Panel (a) represents a long exposure of the total intensity, i.e. what would be recorded by a traditional camera, by generating individual photon lists for each pixel using the “true” maps for the diffraction limited PSF (I_c), seeing halo (I_s), and planets (I_p) in panels (c), (e), and (g), respectively. Panel (b) shows perfect subtraction of I_c and I_s from the total intensity. Panels (d), (f), and (h) indicate the maximum likelihood estimates (MLE) from SSD analysis. Comparison of panels (b) and (h) show that SSD gives better SNR than perfect reference subtraction. The planet separations are 3.5, 6.5, 9.5, and 12.5 λ/D with peak intensities of 40 (upper right), 20 (lower left), 10 (upper left), and 5 (lower right) counts/sec corresponding to contrast ratios of 4×10^{-4} , 2×10^{-4} , 1×10^{-4} , 5×10^{-5} . The values of I_c and I_s at the planet locations are listed in Table 3.2. The simulation is 30 seconds of data of a $J = 10$ magnitude star observed with a 8.2 meter telescope with Strehl ratio of 0.7 and $\tau = 0.1$ second. The seeing halo is represented by a Moffat profile with FWHM = 0.5 arcsecond. The obscuration at the center of each frame has a radius of 1.22 λ/D to hide the core of the stellar PSF to aid the color scale of the image, but does not suggest that a coronagraph has been used. The width of each frame is 26.9 λ/D .

Table 3.2: I_c and I_s are functions of separation in the simulated images shown in Figure 3.21. Listed here are their values at the planet locations.

Separation (λ/D)	I_c (counts/sec)	I_s (counts/sec)
3.5	104.5	147.5
6.5	12.0	109.6
9.5	2.5	74.1
12.5	0.70	48.3

3.12 Effect of Coronagraph on SSD

In [80], Soummer et. al. explored how I_c and I_s are affected by the Strehl ratio and the addition of a coronagraph. Strehl is the fraction of total stellar light contained in the PSF, so that $SR = 1$ signifies perfect AO correction. The majority of the light in I_c is the PSF, so roughly speaking $I_c \propto SR$. The fraction of stellar light not contained in the PSF goes into quasistatic speckles and the seeing halo, so that $I_s \propto 1 - SR$. See Figure 3.22 for an illustration.

Recall that the job of a coronagraph is to suppress the PSF of a star and effectively




Figure 3.22: Numerical simulations illustrating the decomposition of the mean intensity PSF into I_c and I_s for a direct, noncoronagraphic image. The speckle halo has a lower intensity when the SR is larger. I_c changes very little when $SR \gtrsim 90\%$. Reproduced from [80].

reduce its intensity. While coronagraphs of various complexities can be extremely effective [26] at removing a PSF, they won't help with removing speckles of either the quasistatic or atmospheric variety. A coronagraph's design depends on knowing the deterministic intensity pattern in the focal plane, and speckles are by definition non-deterministic. In other words, a coronagraph will reduce the part of I_c coming from the stellar PSF but not quasistatic speckles produced by optical imperfections, and will not affect I_s at all, as shown in Figure 3.23.

3.13 When to Use SSD

Stochastic speckle discrimination is a powerful tool that has the potential improve the signal to noise of substellar companion detections. On a pixel by pixel level it leverages a known speckle intensity distribution to separate variable intensity I_s from static intensity, which can be further decomposed into static starlight I_c and static companion light I_p .




Figure 3.23: Depending on the type and quality of a coronagraph, the stellar PSF intensity will be reduced. Illustrated here is the effect of an apodized pupil Lyot coronagraph (APLC) with the AO system providing 95% SR. A long exposure image without a coronagraph has contributions from both I_c and I_s , whereas with the coronagraph it is completely dominated by the speckle halo. Reproduced from [80].

Unlike angular differential imaging (ADI) and spectra differential imaging (SDI), SSD performance does not suffer at small angular separations because those algorithms rely on modulation that improves with larger separations. In ADI, a companion traces an arc in the focal plane with length proportional to its separation, while SDI relies on the chromatic behavior of speckles, which move more at larger separations. SSD, in contrast, will have the same performance at different separations given the same combination of I_c , I_s , and I_p .

However, SSD should not be used indiscriminately. In Section 3.11 we saw that it produces results with better SNR than long exposures only when $I_s > I_c$. When that condition isn't satisfied, the SSD results are markedly *worse* than the long exposure case, even when using a prior for I_c . This suggests that noncoronagraphic observations with excellent strehl ratios that yield images where speckle noise is not dominating may be served better by analyses that don't use SSD, such as reference PSF subtraction.

On the other hand, if a coronagraph *is* being used and the stellar PSF's intensity is below that of the speckle noise floor, measures should be taken while observations are ongoing to mitigate quasistatic speckles, as they can still dominate atmospheric speckles. Approaches like speckle nulling and electric field conjugation can push I_c below I_s in at least some parts of the focal plane.

Chapter 4

Tunable Coupler

4.1 Resonator Coupling

Superconducting microwave resonators, such as MKIDs or qubits, are usually defined lithographically from thin superconducting films. The coupling strength between resonators and their feedline, characterized by the coupling quality factor Q_c , is a critical design parameter which sets the lifetime in qubit experiments and affects the detector sensitivity of MKIDs. Usually Q_c is fixed by the geometry of the coupling section of the resonator. In certain applications, however, actively tuning the coupling strength and Q_c is desired. For example, a tunable coupler is useful for qubit readout, control [83, 84], and inter-qubit communication [85, 86]. A superconducting qubit's fidelity and bandwidth depend on, among other factors, how strongly the qubit couples to the measurement system. MKIDs would also benefit from a tunable coupler scheme to achieve the critical coupling condition $Q_c = Q_i$, where Q_i is the resonator internal quality factor, to compensate for changes or uncertainty in the optical loading conditions during lab tests and telescope observations. MKID signal to noise is maximized when Q_c matches Q_i , so maintaining this equality improves detector performance.




Figure 4.1: Example of a $\lambda/4$ resonator. Blue represents the superconducting film, and white represents gaps. The microwave feedline is at the top, the coupling capacitor is the elbow near the feedline, the short to the ground plane is near the bottom, and the length of the resonator defines the resonance frequency. The coupling quality factor Q_c is fixed by the geometry of the feedline and resonator. Reproduced from [43].

Previously, a tunable coupler has been demonstrated using a dc-SQUID as the tunable element, which requires Josephson junction (JJ) fabrication [87]. However, Josephson junctions are complicated to fabricate and require great care during the multistep process. Here, we present a design and experiment for a tunable coupler using a superconducting transmission line with nonlinear kinetic inductance. This tunable coupler is easy to construct and can tune Q_c of a generic resonator by more than a factor of 40 in situ.

4.2 Derivation of 1-port Resonator Scattering Parameters with Tunable Coupler

Lithographically defined superconducting microwave resonators are photon cavities that store microwave photons. The analogy between a one-sided optical cavity and a 1-




Figure 4.2: (a) Cartoon showing an optical cavity, (b) a resonator capacitively coupled to a feedline with one port, and (c) a resonator capacitively coupled to a feedline in the proposed tunable coupler scheme.

port capacitively-coupled microwave resonator is illustrated in Figure 4.2(a) and (b). In both cases, optical (microwave) photons leak in and out of the cavity (resonator) through the partially transparent wall (coupling capacitor) with a transmission coefficient t and associated Q_c given by

$$Q_c = 2\pi \frac{\text{energy stored in cavity}}{\text{energy leaked out per cycle}} = \frac{4m\pi}{|t|^2}, \quad (4.1)$$

where m is the optical (electrical) length of the cavity (resonator) in number of wavelengths λ .

In order to tune Q_c , the magnitude of the transmission coefficient $|t|$ has to be varied. Previously, this has been implemented in an inductive coupling scheme, where the mutual inductance controlling the coupling strength is tuned by a flux-biased SQUID circuit [84, 85]. Here, we describe a tunable capacitive coupling scheme without SQUIDs. In our setup, we shunt the coupling capacitor to ground through a lossless nonlinear transmission line whose phase length is variable, as shown in Figure 4.2(c). As the phase length βl is varied, the shunt impedance of the transmission line varies periodically from an open ($\beta l = \pi/2 + n\pi$) to a short ($\beta l = n\pi$), effectively turning the coupling on and off. Here




Figure 4.3: Two-port capacitor coupler network. The input line impedance is Z_1 and the output line impedance is Z_2 .

β is the propagation constant, l is the physical length of the transmission line, and n is an integer. From microwave network analysis, we find that the wave reflected from the transmission line branch interferes with the direct input wave giving rise to a $1 + \gamma_B$ modulation factor on t as shown in Figure 4.2c. The combined coupler block, indicated by the dashed box in Figure 4.2(c), has an effective transmission coefficient t' and coupling quality factor Q'_c . Let us go through the derivation of these parameters now.

4.2.1 Network Analysis of a 1-port Resonator with Fixed Coupling Capacitor and the Derivation of Q_c

Capacitor

First, consider the 2-port capacitor shown in Figure 4.3 with line impedance Z_1 on port 1 and Z_2 on port 2. We want to calculate the scattering parameters of this 2-port network. The input impedance looking into port 1 is

$$Z_{in} = \frac{1}{j\omega C} + Z_2. \quad (4.2)$$

Using notation from [88], the reflection coefficient

$$S_{11} = \frac{V_1^-}{V_1^+} = \frac{Z_{in} - Z_1}{Z_{in} + Z_1}$$

$$S_{11} = \frac{1 + j\omega CZ_2 - j\omega CZ_1}{1 + j\omega CZ_2 + j\omega CZ_1}, \quad (4.3)$$

where V_1^+ and V_1^- are the voltage amplitudes of the incident and reflected waves at port

1. Defining $\delta_1 = \omega CZ_1$ and $\delta_2 = \omega CZ_2$, S_{11} becomes

$$S_{11} = \frac{1 + j\delta_2 - j\delta_1}{1 + j\delta_2 + j\delta_1}. \quad (4.4)$$

Taylor expanding,

$$S_{11} = (1 + j(\delta_2 - \delta_1)) (1 - j(\delta_2 + \delta_1) - (\delta_2 + \delta_1)^2 + \mathcal{O}(\delta^3)) \quad (4.5)$$

If we assume that the capacitor is the dominant impedance in this network, $\delta_1 \ll 1$ and $\delta_2 \ll 1$ and we can make the approximation

$$S_{11} = 1 - j2\delta_1 - 2\delta_1\delta_2. \quad (4.6)$$

A similar calculation for the reflection from port 2 gives

$$S_{22} = 1 - j2\delta_2 - 2\delta_1\delta_2. \quad (4.7)$$

Let us turn to S_{21} and S_{12} . Again using a voltage divider we get

$$V_2 = \frac{Z_2}{Z_2 + \frac{1}{j\omega C}} V_1,$$

where V_1 is the voltage at port 1, which is the sum of the incident and reflected waves

$$V_1 = V_1^+ + V_1^-$$

$$V_1 = V_1^+(1 + S_{11}) = V_1^+(2 - j2\delta_1).$$

There is no wave incident on port 2, so $V_2^+ = 0$, and

$$S_{21} = \frac{V_2^-}{V_1^+}$$

$$S_{21} = \frac{j\omega CZ_2}{j\omega CZ_2 + 1}(2 - j2\delta_1) \quad (4.8)$$

Making the same approximation as before gives

$$S_{21} = j\delta_2(1 - j\delta_2)(2 - j2\delta_1),$$

and keeping the leading terms and simplifying leaves

$$S_{21} = j2\delta_2. \quad (4.9)$$

A similar derivation can be done for the other port, so that $S_{12} = j2\delta_1$. Scattering matrices for passive components must be symmetrical, so that the transmission coefficient is the same for both directions, i.e. $S_{21} = S_{12}$. Our transmission coefficients are not the same because of the different characteristic impedances on each port of the capacitor. To correct this we need to normalize the scattering parameters to remove their dependence on the characteristic impedances.

$$S_{21} \Rightarrow S_{21} \sqrt{\frac{Z_1}{Z_2}}$$

$$S_{21} = j2\sqrt{\delta_1\delta_2}, \quad (4.10)$$

$$S_{12} \Rightarrow S_{12} \sqrt{\frac{Z_2}{Z_1}}$$

$$S_{12} = j2\sqrt{\delta_1\delta_2}. \quad (4.11)$$

The scattering matrix for the capacitor is

$$\mathbf{S} = \begin{bmatrix} 1 - j2\delta_1 - 2\delta_1\delta_2 & j2\sqrt{\delta_1\delta_2} \\ j2\sqrt{\delta_1\delta_2} & 1 - j2\delta_2 - 2\delta_1\delta_2 \end{bmatrix} \quad (4.12)$$

1-port Resonator Readout

Now we'll derive the reflection coefficient from a 1-port resonator. The resonator can be modeled as a $\lambda/4$ shorted transmission line with length l and propagation constant γ . The reflection coefficient of the shorted transmission is

$$\Gamma = -e^{-2\gamma l}. \quad (4.13)$$

The minus sign in front comes from the short at the end of the line, and the propagation constant can be decomposed into a real and imaginary part (because the line is lossy), $\gamma = \alpha + j\beta$. Substituting, the reflection coefficient becomes

$$\Gamma = -e^{-2\alpha l - 2j\beta l}. \quad (4.14)$$

We can rewrite the lossy part in terms of frequency and the internal quality factor Q_i , where Q_i can be thought of as the number of times the resonator will ring before the

stored energy is reduced by $1/e$. We define it as

$$\alpha l = \frac{\pi}{4} \frac{1}{Q_i}. \quad (4.15)$$

The $\pi/4$ denotes the electrical length of the resonator. Γ becomes

$$\Gamma = -e^{-\pi/2Q_i} e^{-2j\beta l}.$$

Rewriting βl in terms of phase velocity and frequency,

$$\beta l = \frac{2\pi f l}{v_p}, \quad (4.16)$$

and defining $f_{\lambda/4} = v_p/4l$,

$$\beta l = \frac{\pi f}{2f_{\lambda/4}}. \quad (4.17)$$

Reading out a $\lambda/4$ resonator requires $f \approx f_{\lambda/4}$, implying that $e^{-j\pi f/f_{\lambda/4}}$ is very close to -1. In other words, its imaginary part is very small. Plugging this into Γ and rewriting, we get

$$\Gamma = -e^{-\pi/2Q_i} e^{-j\pi f/f_{\lambda/4}}$$

$$\Gamma = e^{-\pi/2Q_i} e^{-j\pi f/f_{\lambda/4}} e^{-j\pi}$$

$$\Gamma = e^{-\pi/2Q_i - j\pi \frac{f - f_{\lambda/4}}{f_{\lambda/4}}}. \quad (4.18)$$

Q_i is large in good superconducting films ($\gtrsim 10^6$ for good platinum silicide and niobium titanium nitride), so we can approximate

$$\Gamma \approx 1 - \frac{\pi}{2Q_i} - j\pi \frac{f - f_{\lambda/4}}{f_{\lambda/4}}, \quad (4.19)$$

and

$$\frac{1}{\Gamma} \approx 1 + \frac{\pi}{2Q_i} + j\pi \frac{f - f_{\lambda/4}}{f_{\lambda/4}}. \quad (4.20)$$

1-port Readout of Resonator with Fixed Coupling Capacitor

For convenience, we'll re-write the capacitor scattering matrix derived earlier with “0” and “r” subscripts to denote the feedline and resonator respectively.

$$\mathbf{S} = \begin{bmatrix} 1 - j2\delta_0 - 2\delta_0\delta_r & j2\sqrt{\delta_0\delta_r} \\ j2\sqrt{\delta_0\delta_r} & 1 - j2\delta_r - 2\delta_0\delta_r \end{bmatrix}$$

As mentioned earlier, superconducting microwave resonators couple to their feedlines with a capacitance as shown in Figure 4.2b ($m = 1/4$) with coupling strength characterized by the parameter Q_c . The reflection coefficient R_{11} from port 1 of the capacitor/resonator network is

$$R_{11} = S_{11} + \frac{S_{12}\Gamma S_{21}}{1 - S_{22}\Gamma}. \quad (4.21)$$

Let's go through each piece to understand where this comes from. First, the S_{11} term is the first reflection of the wave off of the capacitor. In the second term, S_{21} is the transmission of the wave through the capacitor, Γ is the resonator reflection coefficient described in the previous section, S_{12} is the transmission of the wave back through the capacitor, and the denominator accounts for multiple internal reflections in the resonator (it's ringing after all).

The coupling quality factor Q_c is

$$Q_c = 2\pi \frac{\text{energy stored in resonator}}{\text{energy leak from port 2 to port 1 per cycle}} = \frac{2\pi}{2|S_{21}|^2}.$$

The factor of 2 in the denominator comes from port 2 reflecting the wave 2 times per cycle.

$$Q_c = \frac{\pi}{|S_{21}|^2} = \frac{\pi}{4\delta_0\delta_r}. \quad (4.22)$$

Plugging in Q_c , assuming that $S_{11} \approx 1$, and using $S_{21} = S_{12}$, R_{11} becomes

$$R_{11} \approx 1 + \frac{S_{21}^2}{\frac{1}{\Gamma} - S_{22}} \quad (4.23)$$

$$R_{11} = 1 - \frac{\frac{\pi}{Q_c}}{\frac{\pi}{2Q_i} + \frac{\pi}{2Q_c} + j\pi \left(\frac{f-f_{\lambda/4}}{f_{\lambda/4}} + \frac{2\delta_0}{\pi} \right)} \quad (4.24)$$

We make another definition and approximation where

$$f_r = f_{\lambda/4} \left(1 - \frac{2\delta_0}{\pi} \right) \quad (4.25)$$

$$R_{11} \approx 1 - \frac{\frac{\pi}{Q_c}}{\frac{\pi}{2Q_i} + \frac{\pi}{2Q_c} + j\pi \left(\frac{f-f_r}{f_r} \right)}. \quad (4.26)$$

The total quality factor of the resonator is typically defined as $1/Q = 1/Q_i + 1/Q_c$,

$$R_{11} = 1 - \frac{\frac{2}{Q_c}}{\frac{1}{Q_i} + \frac{1}{Q_c} + j2 \left(\frac{f-f_r}{f_r} \right)} \quad (4.27)$$

$$R_{11} = 1 - \frac{2\frac{Q}{Q_c}}{1 + j2Q \left(\frac{f-f_r}{f_r} \right)}. \quad (4.28)$$

When $f = f_r$, $R_{11} = 1 - 2Q/Q_c$, implying a resonance circle diameter of $D = 2Q/Q_c$, as illustrated in Figure 4.4. The condition for critical coupling is $Q_i = Q_c$, giving $R_{11} = 0$ when $f = f_r$ meaning that 100% of the power is dissipated in the resonator when it is




Figure 4.4: The diameter of the resonance loop is $D = 2Q/Q_c$. If the resonator is overcoupled $D < 1$, if it's under coupled $D > 1$, and if it's critically coupled $D = 1$.

driven at its resonance frequency.

4.2.2 Network Analysis of a 1-port Resonator with Tunable Coupler and Derivation of Q'_c and t'

The derivation of the reflection coefficient shown above used a fixed coupling capacitor. Our goal with the tunable coupler is to vary the effective coupling strength between the feedline and the resonator. We can accomplish this with a resonator coupled to a transmission line with 2 ports, where one of those ports is connected to a long ($\beta l \gtrsim 10$) transmission line of variable electrical length with a short at the end. The other port brings the microwave probe to and from the resonator. A schematic of this model is shown in Figure 4.5. The resonator will be excited by standing waves on the transmission line, and tuning the length of the line will change the physical location of the nodes of the standing wave. When a voltage maximum excites the resonator we achieve maxi-

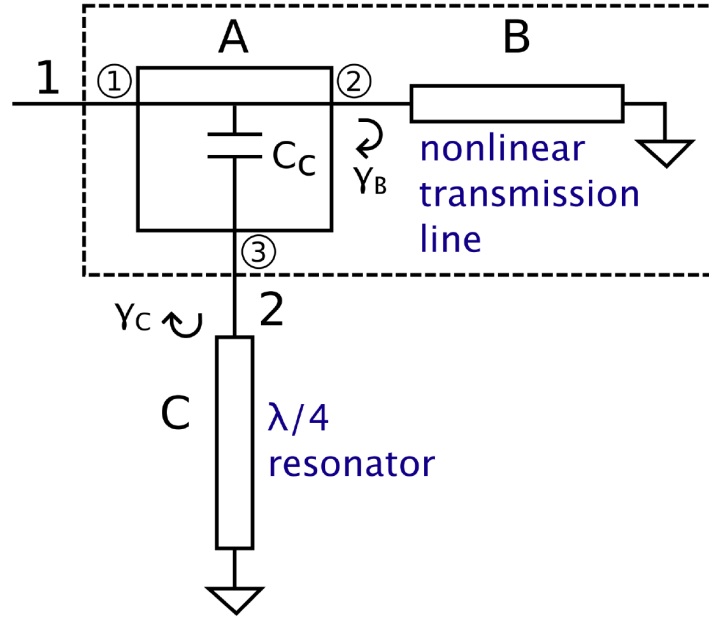


Figure 4.5: Network model of the tunable coupler setup. Reflection measurements are made at port 1. B represents the transmission line of variable electrical length, and the $\lambda/4$ resonator is shown at the bottom.

imum coupling (minimum effective Q_c), while nodes result in minimal coupling (maximum effective Q_c).

We will now derive the scattering parameters of the tunable coupler model network and match the relevant parameters to those of Figure 4.2b ($m = 1/4$). The scattering parameters for network A are derived in [46],

$$\mathbf{S}^A = \begin{bmatrix} -j\delta_0/2 & 1 - j\delta_0/2 & j\sqrt{\delta_0\delta_r} \\ 1 - j\delta_0/2 & -j\delta_0/2 & j\sqrt{\delta_0\delta_r} \\ j\sqrt{\delta_0\delta_r} & j\sqrt{\delta_0\delta_r} & 1 - j2\delta_r \end{bmatrix} \quad (4.29)$$

where $\delta_0 = \omega C_c Z_0$, $\delta_r = \omega C_c Z_r$, Z_0 is the impedance of ports 1 and 2, and Z_r is the impedance of port 3.

The scattering parameters of the full $A + B$ two port network are

$$S_{11}^{AB} = S_{11}^A + \frac{S_{12}^A \gamma_B S_{21}^A}{1 - S_{22}^A \gamma_B} \approx \gamma_B \quad (4.30)$$

$$S_{21}^{AB} = S_{31}^A + \frac{S_{21}^A \gamma_B S_{32}^A}{1 - S_{22}^A \gamma_B} \approx S_{31}^A (1 + \gamma_B) \quad (4.31)$$

$$S_{12}^{AB} = S_{13}^A + \frac{S_{23}^A \gamma_B S_{21}^A}{1 - S_{22}^A \gamma_B} \approx S_{13}^A (1 + \gamma_B) \quad (4.32)$$

$$S_{22}^{AB} = S_{33}^A + \frac{S_{23}^A \gamma_B S_{32}^A}{1 - S_{22}^A \gamma_B} \approx \sqrt{1 - |S_{21}^{AB}|^2} e^{j\phi} = \sqrt{1 - \frac{\pi}{Q_c}} (1 + j\phi) \approx 1 - \frac{1}{2} \frac{\pi}{Q_c} + j\phi \quad (4.33)$$

The approximations in the first three scattering parameters come about because δ_0 is small, while that of S_{22}^{AB} relies on energy conservation.

Finally, we include the resonator in the network model to write down the reflection coefficient from the entire network (blocks $A + B + C$). This is the quantity that we measure.

$$R_{11}^{ABC} = S_{11}^{AB} + \frac{S_{12}^{AB} \Gamma S_{21}^{AB}}{1 - S_{22}^{AB} \Gamma}, \quad (4.34)$$

where Γ is the reflection coefficient from the resonator, Equation 4.19. Plugging in the various quantities and simplifying gives us

$$R_{11}^{ABC} = \gamma_B - \frac{\delta_0 \delta_r (1 + \gamma_B)^2}{1/\Gamma - S_{22}^{AB}} \quad (4.35)$$

$$R_{11}^{ABC} = \gamma_B - \frac{\delta_0 \delta_r (1 + \gamma_B)^2}{\left(\frac{\pi}{2Q_c} + \frac{\pi}{2Q_i}\right) + j\pi \left(\frac{f - f_{\lambda/4}}{f_{\lambda/4}} - \frac{\phi}{\pi}\right)} \quad (4.36)$$

$$R_{11}^{\text{ABC}} = \gamma_B \left(1 - \frac{\delta_0 \delta_r (1 + \gamma_B)^2 / \gamma_B}{\left(\frac{\pi}{2Q_c} + \frac{\pi}{2Q_i} \right) + j \pi \left(\frac{f - f_{\lambda/4}}{f_{\lambda/4}} - \frac{\phi}{\pi} \right)} \right) \quad (4.37)$$

Inserting equation 4.22,

$$R_{11}^{\text{ABC}} = \gamma_B \left(1 - \frac{\frac{\pi}{4Q_c} (1 + \gamma_B)^2 / \gamma_B}{\frac{\pi}{2} \frac{1}{Q} + j \pi \left(\frac{f - f_{\lambda/4}}{f_{\lambda/4}} - \frac{\phi}{\pi} \right)} \right). \quad (4.38)$$

Defining the effective coupling quality factor

$$Q'_c = 4Q_c \frac{\gamma_B}{(1 + \gamma_B)^2}, \quad (4.39)$$

$$R_{11}^{\text{ABC}} = \gamma_B \left(1 - \frac{2Q/Q'_c}{1 + j 2Q \left(\frac{f - f_r}{f_r} \right)} \right). \quad (4.40)$$

This has the same form as the one-port resonator measurement shown in Equation 4.28

From Equation 4.2.1, the transmission coefficient for the fixed coupler

$$t = S_{21} = j 2 \sqrt{\delta_0 \delta_r} \quad (4.41)$$

and from Equations 4.29 and 4.31 we obtain the effective transmission coefficient for the the tunable coupler is

$$t' = S_{21}^{\text{AB}} = j \sqrt{\delta_0 \delta_r} (1 + \gamma_B), \quad (4.42)$$

$$t' = \frac{t}{2} (1 + \gamma_B). \quad (4.43)$$

We can use a trigonometric identity to express Q'_c in terms of the phase of γ_B ,

$$\begin{aligned}
Q'_c &= 4Q_c \frac{\gamma_B}{1 + 2\gamma_B + \gamma_B^2} \\
&= 4Q_c (\gamma_B^{-1} + 2 + \gamma_B)^{-1} \\
&= 4Q_c \left(2 + \frac{1}{|\gamma_B|} e^{-j\angle\gamma_B} + |\gamma_B| e^{j\angle\gamma_B} \right)^{-1} \\
&= \frac{4Q_c}{2 + 2\cos\angle\gamma_B} \\
Q'_c &= \frac{Q_c}{\cos^2(\angle\gamma_B/2)} = \frac{Q_c}{\cos^2(\beta l - \pi/2)}.
\end{aligned} \tag{4.44}$$

Similar to the treatment of the resonator described above, βl is the electrical length or equivalently the phase of the nonlinear transmission line. Equation 4.44 clearly suggests that the effective coupling quality factor Q'_c can be continuously varied from the original value of Q_c to infinity by varying the phase length of the nonlinear transmission line. The next step is to build such a transmission line.

4.3 Nonlinear Transmission Line

Certain superconductors, such as thin-film NbTiN, have large kinetic inductance fractions that also display nonlinear behavior. NbTiN has a current dependent kinetic inductance given by

$$L_k = L_{k,0} \left[1 + \left(\frac{I}{I_*} \right)^2 \right], \tag{4.45}$$

where I is the current flowing through the inductor and I_* is a characteristic current on the order of the superconducting critical current [43]. The nonlinear kinetic inductance of NbTiN has been found to exhibit extremely low dissipation and used in the construction of nonlinear superconducting devices, such as broadband traveling-wave parametric amplifiers [52, 53], frequency tunable resonators [89], and frequency combs [90]. In this

experiment, we adopt a lumped-element transmission line design that is modified from a conventional coplanar waveguide (CPW) transmission line.

The first version of the tunable coupler nonlinear transmission line was a 2.1 meter long coplanar waveguide wrapped into a spiral, as shown in Figure 4.6. The chips were recycled from a project in which we developed a traveling wave parametric amplifier described in [53]. The CPW center strip width and gaps were $2 \mu\text{m}$ with a characteristic impedance of $\sim 180 \Omega$. The center strip width and gap sizes were chosen such that the nonlinearity of the center strip was small enough to get significant gain but large enough for the resolution of the NIST cleanroom stepper.

Measurements of the spiral indicated that oscillations in the gain vs. frequency were the result of an impedance mismatch at the ports of the spiral. We implemented a Klopfenstein stair-stepped taper [88] to mitigate the effects of the change in impedance, which was somewhat effective. Oscillations of several dB in amplitude were still present. While the spiral geometry was sufficient to demonstrate the amplifier, a transmission line with lower characteristic impedance was desired.

Reducing the characteristic impedance of the transmission line to 50Ω while keeping the same center strip width required additional capacitance. We implemented this with an artificial transmission line geometry where the capacitive and inductive sections were essentially discrete elements.

We place $2 \mu\text{m}$ interdigitated capacitor (IDC) fingers directly on both sides of the $2 \mu\text{m}$ center strip, forming a “fishbone” pattern as shown in Figure 4.7. The added capacitance to ground compensates for the high kinetic inductance of the NbTiN center strip resulting in a lumped element transmission line with characteristic impedance of $Z_0 = 50 \Omega$, naturally matched to the measurement system. For the design parameters shown in Figure 4.7, we have verified $Z_0 = 48.5 \Omega$ with time-domain reflectometry (TDR) measurements in a LHe dunk test. As a comparison, a CPW made from 20 nm NbTiN




Figure 4.6: Initial design of spiral traveling wave parametric amplifier used as a long nonlinear transmission line for the tunable coupler setup. The chip was 2 cm wide and the spiral was 2.1 m long, made from 20 nm thick NbTiN. While the nonlinearity was sufficient to vary the coupling strength to a resonator (and also achieve > 10 dB of gain over several GHz), the difference in the characteristic impedance of the line and the rest of the measurement system resulted in significant reflections at the ports. The introduction of impedance tapers helped, but oscillations in the transmission with frequency were still clearly visible. Reproduced from [53].

with $2 \mu\text{m}$ center and gap would have an impedance of more than 200Ω . It should be pointed out that the tuning element is not limited to a 50Ω transmission line and in general any lossless component with a variable reactance may work as a tunable coupler. However, in other impedance mismatched cases, the phase of the reflection coefficient oscillates rapidly with frequency, limiting it to be useful only for narrow-band devices. The trade-off for the 50Ω fishbone tuner is that width is increased by about a factor of 3 from a meandered CPW implementation due to the added IDC fingers. On the other hand, the phase velocity $v_p = 1/\sqrt{LC}$ is lower, reducing the required length of the fishbone line by a factor of ~ 4 from the CPW counterpart.

Because the propagation constant is related to the nonlinear inductance by $\beta =$

$\omega\sqrt{L(I)C}$, the phase length may be modulated by a DC current (I_{dc})

$$\beta l = \beta_0 l \left[1 + \frac{\alpha}{2} \left(\frac{I_{dc}}{I_*} \right)^2 \right], \quad (4.46)$$

where $\beta_0 l$ corresponds to the phase length at zero current and α is the kinetic inductance fraction [43]. Therefore we have constructed a lossless transmission line whose phase length can be varied by injecting a DC current (neglecting the effect of a slight impedance change, which we will discuss later), as required by the tunable coupler scheme in Figure 4.2(c).




Figure 4.7: (a) Lumped element model of a nonlinear kinetic inductance transmission line where the inductance L depends on the current I . (b) One port of the “fishbone” nonlinear transmission line with IDC features. The design dimensions listed in the figure yield a transmission line close to 50Ω characteristic impedance. NbTiN is shown in blue and exposed substrate in white.

4.3.1 Reflection Coefficient of Nonlinear Transmission Line

I_{dc} affects the phase shift of the nonlinear transmission line in two ways. First, I_{dc} changes the phase velocity $v_p = 2\pi f/\beta$ as shown in Equation 4.46. Second, I_{dc} slightly alters the characteristic impedance of the line. The reflection coefficient of the shorted nonlinear transmission line is

$$\gamma_B = \frac{jz \tan \beta l - 1}{jz \tan \beta l + 1}, \quad (4.47)$$

where $z = Z_{\text{NTL}}/Z_0 \approx 1 + \frac{\alpha}{2} \left(\frac{I}{I_*}\right)^2$ is the ratio of the characteristic impedance of the line to the impedance of the measurement system. The phase of γ_B is

$$\angle \gamma_B = -2 \arctan \left[\left(1 + \frac{\alpha}{2} \left(\frac{I}{I_*}\right)^2 \right) \tan [\beta l] \right] + \pi. \quad (4.48)$$

4.3.2 Experimental Setup

Our Q_c tuning experiment was conducted using the setup shown in Figure 4.8 in an adiabatic demagnetization refrigerator (ADR) maintained at a bath temperature of 380 mK, well below the superconducting transition temperature of the rhenium resonators ($T_c \sim 1.7$ K) and the NbTiN fishbone transmission line ($T_c \sim 14$ K). The choice of rhenium for the resonators was based on the availability of chips and their internal quality factors ($Q_i \sim 3 \times 10^4$). A battery-powered, low noise power supply provided a DC current, which was injected into the feedline using a bias tee and a 1 k Ω series resistor. This DC current flows through the circulator, resonator feedline, fishbone transmission line, and finally through a short to ground. The output RF signal, the R_{11}^{ABC} reflection from the entire network including the resonator and the lumped transmission line, was amplified with a HEMT at 3 K. The center-to-ground short at one end of the fishbone transmission




Figure 4.8: (a) A diagram showing the tunable coupler setup. The DC current is injected through the bias tee and flows through the circulator, resonator feedline, fishbone transmission line, and finally through a short to ground. A larger DC current reduces the wavelength of the standing wave (illustrated by the red and blue sinusoidal waves) and varies the phase of the wave seen by the resonator. (b) Photograph of the resonator chip (left) and fishbone transmission line chip (right). One of the 8 resonators on the resonator chip and one of the 3 fishbone transmission lines were used in our experiment.

line and the interconnections between the resonator chip and the fishbone transmission line chip were made using aluminum wirebonds. While wirebonds indeed have a small series inductance, we assume that their effect was negligible.

We first measured the device reflection around its resonance frequency of 4.548 GHz using a vector network analyzer (VNA), while DC current injected into the fishbone transmission line was stepped from 0 to a maximum value of $I_{\max} = 3.32$ mA. For $I > I_{\max}$ part of the transmission line went normal, and the transmission was dramatically reduced. Figure 4.9(a) shows a waterfall plot of all the measured $|R_{11}^{ABC}|^2$ curves under different bias currents. One can see that with increasing DC current, the resonance dip first goes

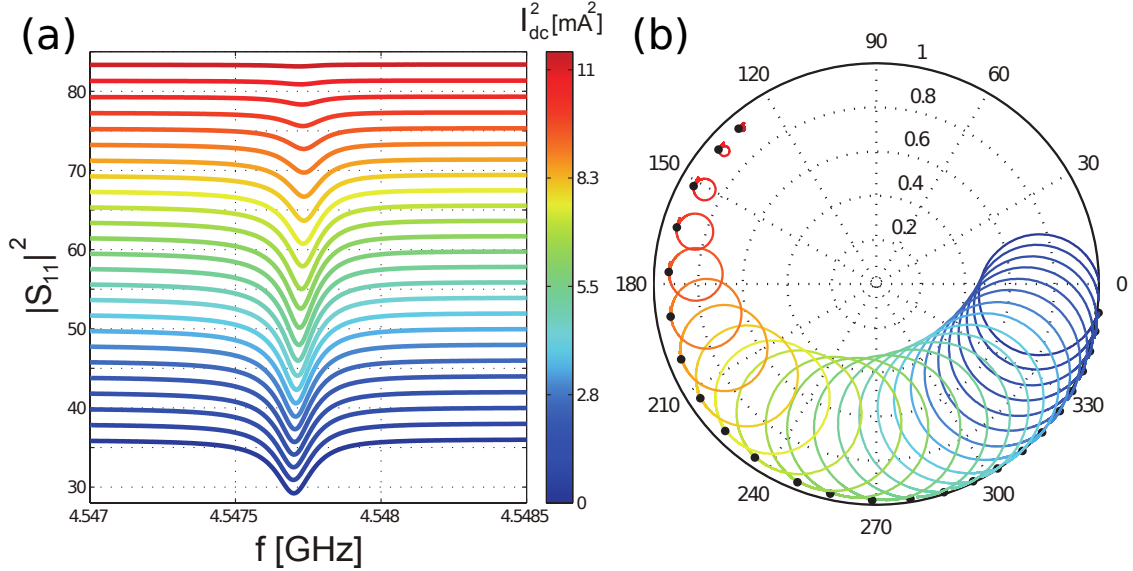


Figure 4.9: Network analyzer sweeps of the resonator-coupler circuit for different values of DC current. $|R_{11}^{ABC}|^2$ data are plotted in (a) with 2 dB offsets in the y-axis added for clarity. The resonance frequency is mostly unchanged but the depth of the resonance dip is modulated by varying Q_c . The normalized complex R_{11}^{ABC} data are plotted in (b). The diameter of the resonance circle is modulated by the effective coupling, Q_c . The first (lowest) frequency point in each sweep is indicated by the black dots.

deeper and then shallower, indicating the coupling quality factor Q_c of the resonator is effectively tuned by the DC currents.

To further analyze the reflection data, we normalized the complex R_{11}^{ABC} data at 380 mK for different DC currents to the R_{11}^{ABC} data with no injected DC current at 700 mK. At 700 mK the resonances die out due to thermal quasiparticles, leaving a clean background for reference. This normalization technique removes cable delays and losses from coaxial cables and was commonly used in resonator measurements in reference [46].

The normalized resonance loops for all the DC currents are shown in Figure 4.9(b). The diameter of each resonance loop is directly related to its effective coupling quality factor, Q'_c . With increasing DC current the centers of the resonance loops move clockwise and the diameters change, indicating both the phase of R_{11}^{ABC} and Q'_c change with current.

The black dots in Figure 4.9(b) indicate the first frequency point f_1 in each sweep. Measured at a fixed off-resonance frequency f_1 , $\gamma_B \approx R_{11}^{\text{ABC}}(f_1) = e^{i\theta}$ gives the reflection coefficient (normalized) from the nonlinear fishbone transmission line only. The phase of these black dots, θ , versus DC current squared, I_{dc}^2 , is plotted in Figure 4.10(a). According to Eqn. (4.46), $\theta \propto I_{\text{dc}}^2$. The θ vs. I_{dc}^2 curve roughly follows a linear trend but at higher currents a clear departure is visible. This effect has been previously observed in other nonlinear devices using NbTiN, where a quartic term I_{dc}^4 has to be included to fit the data [89, 91]. Such a quartic term may arise from higher order nonlinear response of kinetic inductance or from quasiparticle generation [89]. In spite of the detailed model of $\theta(I_{\text{dc}})$ and its underlying physics, Figure 4.10(a) shows that we have obtained a large nonlinear phase shift of over π , useful for Q_c tuning by injecting a DC current.

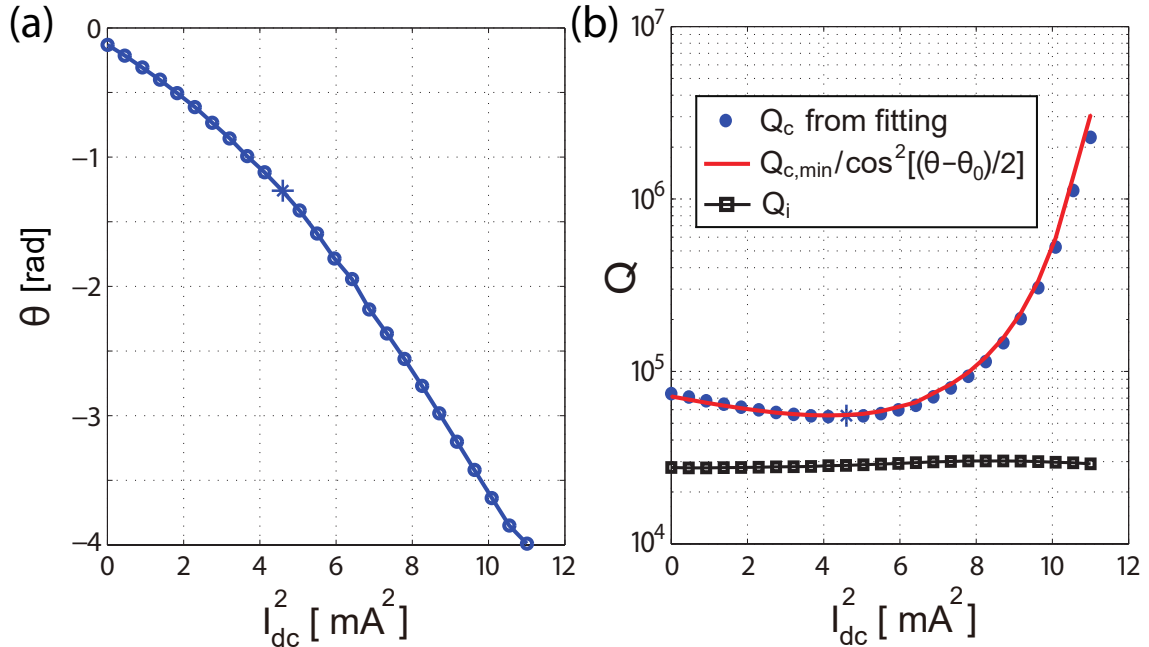


Figure 4.10: (a) Current-induced nonlinear phase shift θ and (b) quality factors as a function of DC current squared I_{dc}^2 . The minimum $Q_{c,\text{min}}$ and the correspondent nonlinear phase shift θ_0 are indicated by the stars.

To quantitatively assess the Q_c tunability, we retrieve the effective Q'_c for each value

of I_{dc} by fitting the complex R_{11}^{ABC} data to a theoretical model derived for our circuit configuration in Equation 4.40, which we will re-write here for convenience:

$$R_{11}^{\text{ABC}}(f) = \gamma_B \left[1 - \frac{2Q/Q'_c}{1 + j2Q \left(\frac{f-f_r}{f_r} \right)} \right]$$

using the resonator fitting procedure established in [46]. Here f_r and Q are the resonance frequency and total quality factor, respectively. Figure 4.10(b) shows the fitted Q_i (squares) and Q'_c (dotted) plotted against I_{dc}^2 . The internal loss Q_i remains nearly constant with increasing DC current while Q'_c ranges from a minimum $Q_{c,\text{min}} = 5.5 \times 10^4$ at $I_{\text{dc}} = 2.14$ mA to a maximum $Q_{c,\text{max}} = 2.3 \times 10^6$ at $I_{\text{dc}} = 3.32$ mA. This result clearly demonstrates that we have achieved tuning of Q_c by more than a factor of 40, while Q_i is unaffected. We further infer that when $Q_{c,\text{min}}$ is reached the DC current has induced an additional nonlinear phase $\theta_0 \sim 1.25$ rad in the transmission line, making the total phase length $\beta l = \pi/2 + n\pi$.

In order to check the agreement between our model and our data, we plot the theoretical model implied by Eqn. (4.44), $Q'_c = Q_{c,\text{min}} / \cos^2[(\theta - \theta_0)/2]$, where θ is the current induced nonlinear phase shift given by Figure 4.10(a), $Q_{c,\text{min}}$ is the lowest measured Q'_c and θ_0 is its corresponding nonlinear phase (indicated by the stars in Figure 4.10). As shown by the red curve in Figure 4.10(b), the theoretical model is in excellent agreement with the measured Q'_c .

4.4 Conclusions

With this experiment we demonstrated the tunable coupler scheme for superconducting microwave resonators proposed in Fig. 4.2(c) and achieved Q_c tuning over a factor of 40, from $Q_c \sim 5.5 \times 10^4$ to $Q_c \sim 2.3 \times 10^6$. Our scheme used the current-dependent

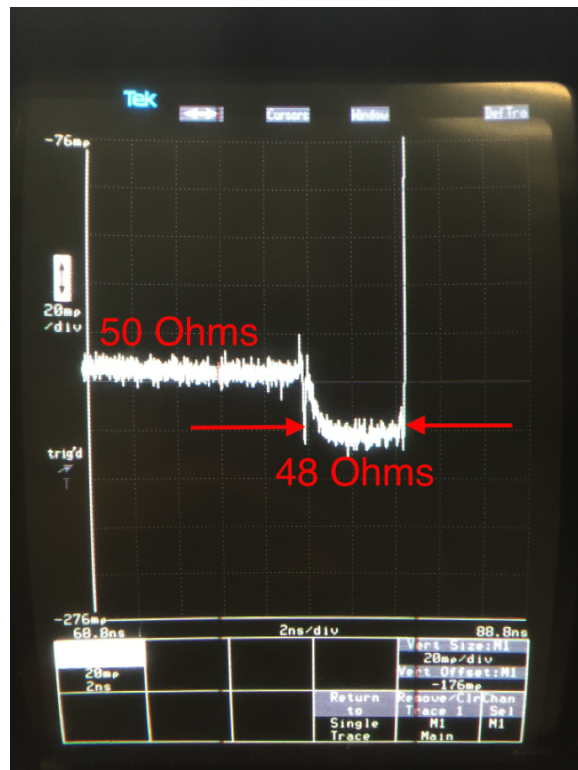


Figure 4.11: Photos of time domain reflectometry (TDR) measurements of the fishbone coupler with a measurement system impedance of 50Ω . The horizontal axis is time and the vertical axis is the reflection coefficient. The length of the fishbone is ≈ 4.4 ns and $\Delta\Gamma \approx -.02$, suggesting $Z_0 \approx 48 \Omega$.

nonlinear kinetic inductance in a superconducting NbTiN transmission line without any Josephson junctions. We introduced a lumped-element transmission line architecture with an interdigitated capacitor to compensate for the high kinetic inductance of NbTiN. Our “fishbone” design provides an elegant solution for making highly nonlinear transmission lines with $\sim 50 \Omega$ intrinsic characteristic impedance, and may be used in other applications such as kinetic inductance traveling-wave parametric amplifiers. This device is easy to construct and requires only single layer fabrication. Integrating the tunable coupler with resonators to be tuned onto the same chip should be uncomplicated. If a tight integration is needed, the footprint of the coupler circuit can be further reduced by shrinking the width of the spine and the IDC finger/gap of the fishbone. The demonstration described here was for a 1-port resonator; it is straightforward to extend the scheme to 2-port resonators with independent control of the coupling strength for the two ports.

Bibliography

- [1] S. R. Meeker, B. A. Mazin, A. B. Walter, P. Strader, N. Fruitwala, C. Bockstiegel, P. Szypryt, G. Ulbricht, G. Coiffard, B. Bumble, G. Cancelo, T. Zmuda, K. Treptow, N. Wilcer, G. Collura, R. Dodkins, I. Lipartito, N. Zobrist, M. Bottom, J. C. Shelton, D. Mawet, J. C. van Eyken, G. Vasisht, and E. Serabyn, *DARKNESS: A Microwave Kinetic Inductance Detector Integral Field Spectrograph for High-contrast Astronomy*, *Publications of the Astronomical Society of the Pacific* **130** (June, 2018) 065001, [arXiv:1803.1042].
- [2] A. B. Walter, C. Bockstiegel, T. D. Brandt, and B. A. Mazin, *Stochastic Speckle Discrimination with Time-Tagged Photon Lists: Digging Below the Speckle Noise Floor*, *arXiv e-prints* (Jun, 2019) arXiv:1906.03354, [arXiv:1906.0335].
- [3] C. Bockstiegel, Y. Wang, M. R. Vissers, L. F. Wei, S. Chaudhuri, J. Hubmayr, and J. Gao, *A tunable coupler for superconducting microwave resonators using a nonlinear kinetic inductance transmission line*, *Applied Physics Letters* **108** (2016), no. 22.
- [4] “The nasa exoplanet archive.” <https://exoplanetarchive.ipac.caltech.edu>, 2019. [Online; accessed 2-May-2019].
- [5] M. Perryman, *The Exoplanet Handbook*. 2018.
- [6] S. Seager, *Exoplanets*. 2010.
- [7] A. Wolszczan and D. A. Frail, *A planetary system around the millisecond pulsar PSR1257 + 12*, *Nature* **355** (Jan, 1992) 145–147.
- [8] B. P. Bowler, *Imaging Extrasolar Giant Planets*, arXiv:1605.0273.
- [9] O. Guyon, *Extreme Adaptive Optics*, *Annual Review of Astronomy and Astrophysics* **56** (Sep, 2018) 315–355.
- [10] M. C. Wyatt, *Evolution of debris disks.*, *Annual Review of Astron and Astrophys* **46** (Sep, 2008) 339–383.

- [11] L. A. Hillenbrand, J. M. Carpenter, J. S. Kim, M. R. Meyer, D. E. Backman, A. Moro-Martín, D. J. Hollenbach, D. C. Hines, I. Pascucci, and J. Bouwman, *The Complete Census of 70 μm -Bright Debris Disks within “The Formation and Evolution of Planetary Systems” Spitzer Legacy Survey of Sun-like Stars*, *Astrophysical Journal* **677** (Apr, 2008) 630–656, [arXiv:0801.0163].
- [12] T. Meshkat, D. Mawet, M. L. Bryan, S. Hinkley, B. P. Bowler, K. R. Stapelfeldt, K. Batygin, D. Padgett, F. Y. Morales, E. Serabyn, V. Christiaens, T. D. Brandt, and Z. Wahhaj, *A Direct Imaging Survey of Spitzer-detected Debris Disks: Occurrence of Giant Planets in Dusty Systems*, *Astronomical Journal* **154** (Dec, 2017) 245, [arXiv:1710.0418].
- [13] T. Meshkat, R. Nilsson, J. Aguilar, G. Vasisht, R. Oppenheimer, K. Y. L. Su, E. Cady, T. Lockhart, C. Matthews, R. Dekany, J. Leisenring, M. Ygouf, D. Mawet, L. Pueyo, and C. Beichman, *A Deep Search for Planets in the Inner 15 au around Vega*, *Astronomical Journal* **156** (Nov, 2018) 214, [arXiv:1809.0694].
- [14] K. Y. L. Su, G. H. Rieke, R. Malhotra, K. R. Stapelfeldt, A. M. Hughes, A. Bonsor, D. J. Wilner, Z. Balog, D. M. Watson, M. W. Werner, and K. A. Misselt, *Asteroid Belts in Debris Disk Twins: Vega and Fomalhaut*, *Astrophysical Journal* **763** (Feb, 2013) 118, [arXiv:1301.1331].
- [15] A. L. Maire, A. J. Skemer, P. M. Hinz, S. Desidera, S. Esposito, R. Gratton, F. Marzari, M. F. Skrutskie, B. A. Biller, D. Defrère, V. P. Bailey, J. M. Leisenring, D. Apai, M. Bonnefoy, W. Brandner, E. Buenzli, R. U. Claudi, L. M. Close, J. R. Crepp, R. J. De Rosa, J. A. Eisner, J. J. Fortney, T. Henning, K. H. Hofmann, T. G. Kopytova, J. R. Males, D. Mesa, K. M. Morzinski, A. Oza, J. Patience, E. Pinna, A. Rajan, D. Schertl, J. E. Schlieder, K. Y. L. Su, A. Vaz, K. Ward-Duong, G. Weigelt, and C. E. Woodward, *The LEECH Exoplanet Imaging Survey. Further constraints on the planet architecture of the HR 8799 system*, *Astronomy and Astrophysics* **576** (Apr, 2015) A133, [arXiv:1412.6989].
- [16] E. L. Nielsen, M. C. Liu, Z. Wahhaj, B. A. Biller, T. L. Hayward, J. R. Males, L. M. Close, K. M. Morzinski, A. J. Skemer, M. J. Kuchner, T. J. Rodigas, P. M. Hinz, M. Chun, C. Ftaclas, and D. W. Toomey, *The Gemini NICI Planet-Finding Campaign: The Orbit of the Young Exoplanet β Pictoris b*, *Astrophysical Journal* **794** (Oct, 2014) 158, [arXiv:1403.7195].
- [17] R. Galicher, J. Rameau, M. Bonnefoy, J. L. Baudino, T. Currie, A. Boccaletti, G. Chauvin, A. M. Lagrange, and C. Marois, *Near-infrared detection and characterization of the exoplanet HD 95086 b with the Gemini Planet Imager*, *Astronomy and Astrophysics* **565** (May, 2014) L4, [arXiv:1404.4635].
- [18] B. Macintosh, J. R. Graham, T. Barman, R. J. De Rosa, Q. Konopacky, M. S. Marley, C. Marois, E. L. Nielsen, L. Pueyo, A. Rajan, J. Rameau, D. Saumon, J. J.

- Wang, J. Patience, M. Ammons, P. Arriaga, E. Artigau, S. Beckwith, J. Brewster, S. Bruzzone, J. Bulger, B. Burningham, A. S. Burrows, C. Chen, E. Chiang, J. K. Chilcote, R. I. Dawson, R. Dong, R. Doyon, Z. H. Draper, G. Duchêne, T. M. Esposito, D. Fabrycky, M. P. Fitzgerald, K. B. Follette, J. J. Fortney, B. Gerard, S. Goodsell, A. Z. Greenbaum, P. Hibon, S. Hinkley, T. H. Cotten, L. W. Hung, P. Ingraham, M. Johnson-Groh, P. Kalas, D. Lafreniere, J. E. Larkin, J. Lee, M. Line, D. Long, J. Maire, F. Marchis, B. C. Matthews, C. E. Max, S. Metchev, M. A. Millar-Blanchaer, T. Mittal, C. V. Morley, K. M. Morzinski, R. Murray-Clay, R. Oppenheimer, D. W. Palmer, R. Patel, M. D. Perrin, L. A. Poyneer, R. R. Rafikov, F. T. Rantakyro, E. L. Rice, P. Rojo, A. R. Rudy, J. B. Ruffio, M. T. Ruiz, N. Sadakuni, L. Saddlemyer, M. Salama, D. Savransky, A. C. Schneider, A. Sivaramakrishnan, I. Song, R. Soummer, S. Thomas, G. Vasisht, J. K. Wallace, K. Ward-Duong, S. J. Wiktorowicz, S. G. Wolff, and B. Zuckerman, *Discovery and spectroscopy of the young jovian planet 51 Eri b with the Gemini Planet Imager*, *Science* **350** (Oct, 2015) 64–67, [arXiv:1508.0308].
- [19] A. Zurlo, A. Vigan, R. Galicher, A. L. Maire, D. Mesa, R. Gratton, G. Chauvin, M. Kasper, C. Moutou, M. Bonnefoy, S. Desidera, L. Abe, D. Apai, A. Baruffolo, P. Baudoz, J. Baudrand, J. L. Beuzit, P. Blancard, A. Boccaletti, F. Cantalloube, M. Carle, E. Cascone, J. Charton, R. U. Claudi, A. Costille, V. de Caprio, K. Dohlen, C. Dominik, D. Fantinel, P. Feautrier, M. Feldt, T. Fusco, P. Gigan, J. H. Girard, D. Gisler, L. Gluck, C. Gry, T. Henning, E. Hugot, M. Janson, M. Jaquet, A. M. Lagrange, M. Langlois, M. Llored, F. Madec, Y. Magnard, P. Martinez, D. Maurel, D. Mawet, M. R. Meyer, J. Milli, O. Moeller-Nilsson, D. Mouillet, A. Origné, A. Pavlov, C. Petit, P. Puget, S. P. Quanz, P. Rabou, J. Ramos, G. Rousset, A. Roux, B. Salasnich, G. Salter, J. F. Sauvage, H. M. Schmid, C. Soenke, E. Stadler, M. Suarez, M. Turatto, S. Udry, F. Vakili, Z. Wahhaj, F. Wildi, and J. Antichi, *First light of the VLT planet finder SPHERE. III. New spectrophotometry and astrometry of the HR 8799 exoplanetary system*, *Astronomy and Astrophysics* **587** (Mar, 2016) A57, [arXiv:1511.0408].
- [20] R. Dekany, J. Roberts, R. Burruss, A. Bouchez, T. Truong, C. Baranec, S. Guiwits, D. Hale, J. Angione, T. Trinh, J. Zolkower, J. C. Shelton, D. Palmer, J. Henning, E. Croner, M. Troy, D. McKenna, J. Tesch, S. Hildebrand t, and J. Milburn, *PALM-3000: Exoplanet Adaptive Optics for the 5 m Hale Telescope*, *Astrophysical Journal* **776** (Oct, 2013) 130, [arXiv:1309.1216].
- [21] L. Poyneer, M. van Dam, and J.-P. Véran, *Experimental verification of the frozen flow atmospheric turbulence assumption with use of astronomical adaptive optics telemetry*, *Journal of the Optical Society of America A* **26** (Mar, 2009) 833.
- [22] D. L. Fried, *Statistics of a geometric representation of wavefront distortion*, *J. Opt. Soc. Am.* **55** (Nov, 1965) 1427–1435.

- [23] J. R. Males, L. M. Close, K. Miller, L. Schatz, D. Doelman, J. Lumbres, F. Snik, A. Rodack, J. Knight, K. Van Gorkom, J. D. Long, A. Hedglen, M. Kautz, N. Jovanovic, K. Morzinski, O. Guyon, E. Douglas, K. B. Follette, J. Lozi, C. Bohlman, O. Durney, V. Gasho, P. Hinz, M. Ireland, M. Jean, C. Keller, M. Kenworthy, B. Mazin, J. Noenickx, D. Alfred, K. Perez, A. Sanchez, C. Sauve, A. Weinberger, and A. Conrad, *MagAO-X: project status and first laboratory results*, in *Proceedings of the SPIE*, vol. 10703 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 1070309, Jul, 2018. arXiv:1807.0431.
- [24] C. Marois, B. Macintosh, T. Barman, B. Zuckerman, I. Song, J. Patience, D. Lafrenière, and R. Doyon, *Direct Imaging of Multiple Planets Orbiting the Star HR 8799*, *Science* **322** (Nov, 2008) 1348, [arXiv:0811.2606].
- [25] B. Lyot, *The study of the solar corona and prominences without eclipses (George Darwin Lecture, 1939)*, *Monthly Notices of the RAS* **99** (May, 1939) 538.
- [26] G. Ruane, A. Riggs, J. Mazoyer, E. H. Por, M. N’Diaye, E. Huby, P. Baudoz, R. Galicher, E. Douglas, and J. Knight, *Review of high-contrast imaging systems for current and future ground- and space-based telescopes I: coronagraph design methods and optical performance metrics*, in *Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave*, vol. 10698 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 106982S, Aug, 2018. arXiv:1807.0704.
- [27] D. Mawet, E. Serabyn, K. Liewer, R. Burruss, J. Hickey, and D. Shemo, *The Vector Vortex Coronagraph: Laboratory Results and First Light at Palomar Observatory*, *Astrophysical Journal* **709** (Jan, 2010) 53–57, [arXiv:0912.2287].
- [28] A. Labeyrie, *Images of exo-planets obtainable from dark speckles in adaptive telescopes.*, *Astronomy and Astrophysics* **298** (Jun, 1995) 544.
- [29] A. Boccaletti, A. Labeyrie, and R. Ragazzoni, *Preliminary results of dark-speckle stellar coronagraphy*, *Astronomy and Astrophysics* **338** (Oct, 1998) 106–110, [astro-ph/9806144].
- [30] J. W. Goodman, *Statistical Properties of Laser Speckle Patterns*, pp. 9–75. Springer Berlin Heidelberg, Berlin, Heidelberg, 1975.
- [31] V. F. Canales and M. P. Cagigal, *Rician Distribution to Describe Speckle Statistics in Adaptive Optics*, *Applied Optics* **38** (Feb, 1999) 766–771.
- [32] M. P. Cagigal and V. F. Canales, *Photon statistics in adaptive optics*, in *Optics in Atmospheric Propagation and Adaptive Systems III* (A. Kohnle and J. D. Gonglewski, eds.), vol. 3866 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pp. 165–175, Dec, 1999.

- [33] M. P. Cagigal and V. F. Canales, *Photon statistics in adaptive optics*, in *Proceedings of the SPIE* (A. Kohnle and J. D. Gonglewski, eds.), vol. 3866 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pp. 165–175, Dec, 1999.
- [34] M. P. Cagigal and V. F. Canales, *Experimental checking of the Rician statistics in partially compensated wave fronts*, *Optical Engineering* **40** (Dec, 2001) 2690–2697.
- [35] M. P. Fitzgerald and J. R. Graham, *Speckle Statistics in Adaptively Corrected Images*, *Astrophysical Journal* **637** (Jan., 2006) 541–547, [astro-ph/0510503].
- [36] S. Gladysz and J. C. Christou, *Detection of faint companions through stochastic speckle discrimination*, *The Astrophysical Journal* **684** (sep, 2008) 1486–1495.
- [37] S. Gladysz and J. C. Christou, *Reference-Less Detection, Astrometry, and Photometry of Faint Companions with Adaptive Optics*, *Astrophysical Journal* **698** (Jun, 2009) 28–42, [arXiv:0805.1870].
- [38] S. Gladysz, N. Yaitskova, and J. C. Christou, *Statistics of intensity in adaptive-optics images and their usefulness for detection and photometry of exoplanets*, *Journal of the Optical Society of America A* **27** (Nov., 2010) A64–A75.
- [39] F. Martinache, O. Guyon, C. Clergeon, and C. Blain, *Speckle Control with a Remapped-Pupil PIAA Coronagraph*, *Publications of the Astronomical Society of the Pacific* **124** (Dec, 2012) 1288, [arXiv:1206.2996].
- [40] C. Marois, D. W. Phillion, and B. Macintosh, *Exoplanet detection with simultaneous spectral differential imaging: effects of out-of-pupil-plane optical aberrations*, in *Proceedings of the SPIE*, vol. 6269 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 62693M, Jun, 2006. astro-ph/0607002.
- [41] J. R. Crepp, L. Pueyo, D. Brenner, B. R. Oppenheimer, N. Zimmerman, S. Hinkley, I. Parry, D. King, G. Vasisht, C. Beichman, L. Hillenbrand, R. Dekany, M. Shao, R. Burruss, L. C. Roberts, A. Bouchez, J. Roberts, and R. Soummer, *Speckle Suppression with the Project 1640 Integral Field Spectrograph*, *Astrophysical Journal* **729** (Mar, 2011) 132, [arXiv:1012.4016].
- [42] C. Marois, D. Lafrenière, R. Doyon, B. Macintosh, and D. Nadeau, *Angular Differential Imaging: A Powerful High-Contrast Imaging Technique*, *Astrophysical Journal* **641** (Apr, 2006) 556–564, [astro-ph/0512335].
- [43] J. Zmuidzinas, *Superconducting microresonators: Physics and applications*, *Annual Review of Condensed Matter Physics* **3** (2012), no. 1 169–214.

- [44] P. K. Day, H. G. LeDuc, B. A. Mazin, A. Vayonakis, and J. Zmuidzinas, *A broadband superconducting detector suitable for use in large arrays*, *Nature* **425** (2003) 817.
- [45] B. A. Mazin, *Microwave kinetic inductance detectors*, .
- [46] J. Gao, *The Physics of Superconducting Microwave Resonators Thesis by*. PhD thesis, California Institute of Technology, 2008.
- [47] A. G. Kozorezov, J. K. Wigmore, D. Martin, P. Verhoeve, and A. Peacock, *Electron energy down-conversion in thin superconducting films*, *Physical Review B: Solid State* **75** (Mar, 2007) 094513.
- [48] B. Dober, J. A. Austermann, J. A. Beall, D. Becker, G. Che, H. M. Cho, M. Devlin, S. M. Duff, N. Galitzki, J. Gao, C. Groppi, G. C. Hilton, J. Hubmayr, K. D. Irwin, C. M. McKenney, D. Li, N. Lourie, P. Mauskopf, M. R. Vissers, and Y. Wang, *Optical Demonstration of THz, Dual-Polarization Sensitive Microwave Kinetic Inductance Detectors*, *Journal of Low Temperature Physics* **184** (Jul, 2016) 173–179, [arXiv:1603.0296].
- [49] J. Sayers, C. Bockstiegel, S. Brugger, N. G. Czakon, P. K. Day, T. P. Downes, R. P. Duan, J. Gao, A. K. Gill, J. Glenn, S. R. Golwala, M. I. Hollister, A. Lam, H. G. LeDuc, P. R. Maloney, B. A. Mazin, S. G. McHugh, D. A. Miller, A. K. Mroczkowski, O. Noroozian, H. T. Nguyen, J. A. Schlaerth, S. R. Siegel, A. Vayonakis, P. R. Wilson, and J. Zmuidzinas, *The status of MUSIC: the multiwavelength sub-millimeter inductance camera*, in *Proceedings of the SPIE*, vol. 9153 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 915304, Aug, 2014.
- [50] P. Szypryt, B. A. Mazin, G. Ulbricht, B. Bumble, S. R. Meeker, C. Bockstiegel, and A. B. Walter, *High quality factor platinum silicide microwave kinetic inductance detectors*, *Applied Physics Letters* **109** (2016), no. 15 151102.
- [51] J. Gao, M. Daal, A. Vayonakis, S. Kumar, J. Zmuidzinas, B. Sadoulet, B. A. Mazin, P. K. Day, and H. G. Leduc, *Experimental evidence for a surface distribution of two-level systems in superconducting lithographed microwave resonators*, *Applied Physics Letters* **92** (Apr, 2008) 152505, [arXiv:0802.4457].
- [52] B. Ho Eom, P. K. Day, H. G. LeDuc, and J. Zmuidzinas, *A wideband, low-noise superconducting amplifier with high dynamic range*, *Nature Physics* **8** (2012), no. 8 623–627, [arXiv:1201.2392].
- [53] C. Bockstiegel, J. Gao, M. R. Vissers, M. Sandberg, S. Chaudhuri, a. Sanders, L. R. Vale, K. D. Irwin, and D. P. Pappas, *Development of a broadband NbTiN traveling wave parametric amplifier for MKID readout*, *Journal of Low Temperature Physics* **176** (2014), no. 3-4 476–482.

- [54] N. Zobrist, B. H. Eom, P. Day, B. A. Mazin, S. R. Meeker, B. Bumble, H. G. LeDuc, G. Coiffard, P. Szypryt, N. Fruitwala, I. Lipartito, and C. Bockstiegel, *Wide-band parametric amplifier readout and resolution of optical microwave kinetic inductance detectors*, *Applied Physics Letters* **115** (2019), no. 4 042601, [<https://doi.org/10.1063/1.5098469>].
- [55] M. J. Strader, *Digital readout for microwave kinetic inductance detectors and applications in high time resolution astronomy*. PhD thesis, University of California, Santa Barbara, Jan, 2016.
- [56] J. C. van Eyken, M. J. Strader, A. B. Walter, S. R. Meeker, P. Szypryt, C. Stoughton, K. O'Brien, D. Marsden, N. K. Rice, Y. Lin, and B. A. Mazin, *The ARCONS Pipeline: Data Reduction for MKID Arrays*, *Astrophysical Journal, Supplement* **219** (Jul, 2015) 14, [arXiv:1507.05663].
- [57] T. Cook, K. Cahoy, S. Chakrabarti, E. Douglas, S. C. Finn, M. Kuchner, N. Lewis, A. Marinan, J. Martel, D. Mawet, B. Mazin, S. R. Meeker, C. Mendillo, G. Serabyn, D. Stuchlik, and M. Swain, *Planetary Imaging Concept Testbed Using a Recoverable Experiment-Coronagraph (PICTURE C)*, *Journal of Astronomical Telescopes, Instruments, and Systems* **1** (Oct, 2015) 044001.
- [58] W. A. Traub, P. Chen, B. Kern, and T. Matsuo, *Planetscope: an exoplanet coronagraph on a balloon platform*, in *Proceedings of the SPIE*, vol. 7010 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 70103S, Jul, 2008.
- [59] C. B. Mendillo, G. A. Howe, K. Hewawasam, J. Martel, S. C. Finn, T. A. Cook, and S. Chakrabarti, *Optical tolerances for the PICTURE-C mission: error budget for electric field conjugation, beam walk, surface scatter, and polarization aberration*, in *Proceedings of the SPIE*, vol. 10400 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 1040010, Sep, 2017.
- [60] A. V. Krivov, *Debris disks: seeing dust, thinking of planetesimals and planets*, *Research in Astronomy and Astrophysics* **10** (May, 2010) 383–414, [arXiv:1003.5229].
- [61] A. Boccaletti, C. Thalmann, A.-M. Lagrange, M. Janson, J.-C. Augereau, G. Schneider, J. Milli, C. Grady, J. Debes, M. Langlois, D. Mouillet, T. Henning, C. Dominik, A.-L. Maire, J.-L. Beuzit, J. Carson, K. Dohlen, N. Engler, M. Feldt, T. Fusco, C. Ginski, J. H. Girard, D. Hines, M. Kasper, D. Mawet, F. Ménard, M. R. Meyer, C. Moutou, J. Olofsson, T. Rodigas, J.-F. Sauvage, J. Schlieder, H. M. Schmid, M. Turatto, S. Udry, F. Vakili, A. Vigan, Z. Wahhaj, and J. Wisniewski, *Fast-moving features in the debris disk around AU Microscopii*, *Nature* **526** (Oct, 2015) 230–232.

- [62] É. Choquet, G. Bryden, M. D. Perrin, R. Soummer, J. C. Augereau, C. H. Chen, J. H. Debes, E. Gofas-Salas, J. B. Hagan, D. C. Hines, D. Mawet, F. Morales, L. Pueyo, A. Rajan, B. Ren, G. Schneider, C. C. Stark, and S. Wolff, *HD 104860 and HD 192758: Two Debris Disks Newly Imaged in Scattered Light with the Hubble Space Telescope*, *Astrophysical Journal* **854** (Feb, 2018) 53, [arXiv:1801.0542].
- [63] D. W. Stuchlik, *The Wallops Arc Second Pointer - A Balloon Borne Fine Pointing System*. <https://arc.aiaa.org/doi/pdf/10.2514/6.2015-3039>.
- [64] G. A. Howe, C. B. Mendillo, K. Hewawasam, J. Martel, S. C. Finn, T. A. Cook, and S. Chakrabarti, *The low-order wavefront control system for the PICTURE-C mission: preliminary testbed results from the Shack-Hartmann sensor*, in *Proceedings of the SPIE*, vol. 10400 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 104001Z, Sep, 2017.
- [65] K. Hewawasam, C. B. Mendillo, G. A. Howe, J. Martel, S. C. Finn, T. A. Cook, and S. Chakrabarti, *The low-order wavefront control system for the PICTURE-C mission: high-speed image acquisition and processing*, in *Proceedings of the SPIE*, vol. 10400 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 1040020, Sep, 2017.
- [66] C. B. Mendillo, G. A. Howe, K. Hewawasam, J. Martel, T. A. Cook, and S. Chakrabarti, *Polarization aberration analysis for the PICTURE-C exoplanetary coronagraph*, *Journal of Astronomical Telescopes, Instruments, and Systems* **5** (Apr, 2019) 025003.
- [67] A. Give'on, B. D. Kern, and S. Shaklan, *Pair-wise, deformable mirror, image plane-based diversity electric field estimation for high contrast coronagraphy*, in *Proceedings of the SPIE*, vol. 8151 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 815110, Oct, 2011.
- [68] *Emissivity Coefficients Materials*, 2003 (accessed July 21, 2019). https://www.engineeringtoolbox.com/emissivity-coefficients-d_447.html.
- [69] C. Wang, G. Thummes, and C. Heiden, *A two-stage pulse tube cooler operating below 4 K*, *Cryogenics* **37** (Jan, 1997) 159–164.
- [70] B. A. Mazin, S. R. Meeker, M. J. Strader, P. Szypryt, D. Marsden, J. C. van Eyken, G. E. Duggan, A. B. Walter, G. Ulbricht, M. Johnson, B. Bumble, K. O'Brien, and C. Stoughton, *ARCONS: A 2024 Pixel Optical through Near-IR Cryogenic Imaging Spectrophotometer*, *Publications of the ASP* **125** (Nov, 2013) 1348, [arXiv:1306.4674].

- [71] A. B. Walter, C. Bockstiegel, B. A. Mazin, and M. Daal, *Laminated NbTi-on-Kapton Microstrip Cables for Flexible Sub-Kelvin RF Electronics*, *IEEE Transactions on Applied Superconductivity* **28** (Jan, 2018) 1–5.
- [72] C. V. Heer, C. B. Barnes, and J. G. Daunt, *The Design and Operation of a Magnetic Refrigerator for Maintaining Temperatures below 1 K*, *Review of Scientific Instruments* **25** (Nov, 1954) 1088–1098.
- [73] R. D. Britt and P. L. Richards, *An adiabatic demagnetization refrigerator for infrared bolometers*, *International Journal of Infrared and Millimeter Waves* **2** (Nov, 1981) 1083–1096.
- [74] S. Pawlyk, P. A. R. Ade, D. Benford, C. L. Bennett, D. T. Chuss, R. Datta, J. L. Dotson, J. R. Eimer, D. J. Fixsen, N. N. Gand ilo, T. M. Essinger-Hileman, M. Halpern, G. Hilton, G. F. Hinshaw, K. Irwin, C. Jhabvala, M. Kimball, A. Kogut, L. Lowe, J. J. McMahon, T. M. Miller, P. Mirel, S. H. Moseley, S. Rodriguez, E. Sharp, P. Shirron, J. G. Staguhn, D. F. Sullivan, E. R. Switzer, P. Taraschi, C. E. Tucker, A. Walts, and E. J. Wollack, *The primordial inflation polarization explorer (PIPER): current status and performance of the first flight*, in *Proceedings of the SPIE*, vol. 10708 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 1070806, Jul, 2018.
- [75] J. W. Goodman, *Speckle Phenomena in Optics: Theory and Applications*. 2007.
- [76] C. Aime and R. Soummer, *Influence of speckle and poisson noise on exoplanet detection with a coronagraph*, in *2004 12th European Signal Processing Conference*, pp. 1071–1074, Sep., 2004.
- [77] T. D. Brandt, “Speckle statistics.” unpublished, July, 2018.
- [78] D. Lafrenière, C. Marois, R. Doyon, and T. Barman, *HST/NICMOS Detection of HR 8799 b in 1998*, *The Astrophysical Journal* **694** (Apr, 2009) L148–L152, [arXiv:0902.3247].
- [79] R. Jensen-Clem, D. Mawet, C. A. G. Gonzalez, O. Absil, R. Belikov, T. Currie, M. A. Kenworthy, C. Marois, J. Mazoyer, G. Ruane, A. Tanner, and F. Cantalloube, *A new standard for assessing the performance of high contrast imaging systems*, *The Astronomical Journal* **155** (dec, 2017) 19.
- [80] R. Soummer and A. Ferrari, *The Strehl Ratio in Adaptive Optics Images: Statistics and Estimation*, *Astrophysical Journal, Letters* **663** (Jul, 2007) L49–L52, [arXiv:0706.1718].
- [81] T. S. Ross, *Limitations and applicability of the maréchal approximation*, *Appl. Opt.* **48** (Apr, 2009) 1812–1818.

- [82] R. Racine, G. A. H. Walker, D. Nadeau, R. Doyon, and C. Marois, *Speckle Noise and the Detection of Faint Companions*, *Publications of the Astronomical Society of the Pacific* **111** (May, 1999) 587–594.
- [83] M. S. Allman, F. Altomare, J. D. Whittaker, K. Cicak, D. Li, A. Sirois, J. Strong, J. D. Teufel, and R. W. Simmonds, *Rf-squid-mediated coherent tunable coupling between a superconducting phase qubit and a lumped-element resonator*, *Phys. Rev. Lett.* **104** (Apr, 2010) 177004.
- [84] M. S. Allman, J. D. Whittaker, M. Castellanos-Beltran, K. Cicak, F. da Silva, M. P. DeFeo, F. Lecocq, A. Sirois, J. D. Teufel, J. Aumentado, and R. W. Simmonds, *Tunable resonant and nonresonant interactions between a phase qubit and lc resonator*, *Phys. Rev. Lett.* **112** (Mar, 2014) 123601.
- [85] Y. Yin, Y. Chen, D. Sank, P. J. J. O’Malley, T. C. White, R. Barends, J. Kelly, E. Lucero, M. Mariantoni, A. Megrant, C. Neill, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis, *Catch and release of microwave photon states*, *Phys. Rev. Lett.* **110** (Mar, 2013) 107001.
- [86] J. Wenner, Y. Yin, Y. Chen, R. Barends, B. Chiaro, E. Jeffrey, J. Kelly, A. Megrant, J. Y. Mutus, C. Neill, P. J. J. O’Malley, P. Roushan, D. Sank, A. Vainsencher, T. C. White, A. N. Korotkov, A. N. Cleland, and J. M. Martinis, *Catching time-reversed microwave coherent state photons with 99.4% absorption efficiency*, *Phys. Rev. Lett.* **112** (May, 2014) 210501.
- [87] R. Harris, A. J. Berkley, M. W. Johnson, P. Bunyk, S. Govorkov, M. C. Thom, S. Uchaikin, A. B. Wilson, J. Chung, E. Holtham, J. D. Biamonte, A. Y. Smirnov, M. H. S. Amin, and A. Maassen van den Brink, *Sign- and magnitude-tunable coupler for superconducting flux qubits*, *Phys. Rev. Lett.* **98** (Apr, 2007) 177001.
- [88] D. M. Pozar, *Microwave engineering*. Fourth edition. Hoboken, NJ : Wiley, [2012] ©2012.
- [89] M. R. Vissers, J. Hubmayr, M. O. Sandberg, S. Chaudhuri, C. Bockstiegel, and J. Gao, *Frequency-tunable Superconducting Resonators via Nonlinear Kinetic Inductance*, *Applied Physics Letters* **107** (2015), no. 062601.
- [90] R. P. Erickson, M. R. Vissers, M. Sandberg, S. R. Jefferts, and D. P. Pappas, *Frequency comb generation in superconducting resonators*, *Phys. Rev. Lett.* **113** (Oct, 2014) 187002.
- [91] A. Kher, P. K. Day, B. H. Eom, J. Zmuidzinas, and H. G. Leduc, *Kinetic inductance parametric up-converter*, *Journal of Low Temperature Physics* (2015) 1–6.