

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

DARKNESS: The First Microwave Kinetic Inductance Detector Integral Field Spectrograph for Exoplanet Imaging

Permalink

<https://escholarship.org/uc/item/1hj1q7s1>

Author

Meeker, Seth Richard

Publication Date

2017

Peer reviewed|Thesis/dissertation

UNIVERSITY of CALIFORNIA
Santa Barbara

**DARKNESS: The First Microwave Kinetic Inductance Detector Integral
Field Spectrograph for Exoplanet Imaging**

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

Doctor of Philosophy

in

Physics

by

Seth Richard Meeker

Committee in charge:

Professor Benjamin Mazin, Chair

Professor Philip Lubin

Dr. Rachel Street

Professor Lars Bildsten

September 2017

The dissertation of Seth Richard Meeker is approved:

Professor Lars Bildsten

Dr. Rachel Street

Professor Philip Lubin

Professor Benjamin Mazin, Chair

September 2017

Copyright © 2017
by Seth Richard Meeker

Dedicated to my parents.
Thank you for a lifetime of support.

Acknowledgements

First, some technical acknowledgements and attributions:

DARKNESS’s development has been a massive endeavor on behalf of a large collaboration and all members of the Mazin lab, and I will attempt to give proper attribution of major components here. The readout in particular, described in Chapter 3, is an extremely complicated subsystem. The readout hardware was designed by Ben Mazin and Paschal (né Matt) Strader with our collaborators at Fermilab. The readout firmware and control software were developed by Sean McHugh, Paschal Strader, Neelay Fruitwala, and Alex Walter. The readout crate was assembled by Giulia Collura. The FPGA chips and ADCs were generously donated by Xilinx and Analog Devices, respectively.

The DARKNESS devices described in Chapters 3 & 4 were fabricated by Bruce Bumble at JPL and Paul Szypryt, Gerhard Ulbricht, and Grégoire Coiffard at UCSB.

The DARKNESS cryostat was manufactured by Precision Cryogenic, and the ADR unit by High Precision Devices Inc. Clint Bockstiegel sacrificed many hours in the lab with me debugging the cryostat.

I owe tremendous gratitude to our collaborators at JPL and Caltech, including Mike Bottom, Rick Burruss, Dimitri Mawet, Gene Serabyn, Chris Shelton, and Gautam Vasisht, who all provided invaluable high-contrast instrumentation expertise throughout DARKNESS’s development. Without Mike’s expert hand running the SDC we would have undoubtedly lost many hours of observing time. A similar thank you goes to Rick for his expertise and on-call assistance with P3K. Observations at Palomar were made possible through Dimitri’s and Gautam’s time. I thank them for their mentorship, and for believing in MKIDs enough to share their time allocation for this project.

Special thanks goes to the Palomar observatory staff who are too numerous to mention here, because we tend to rely on every single person there at one point or another. Developing instruments for Palomar Observatory has been a pleasure thanks to the talented and helpful people. Hopefully the unique and exciting challenges we provide with our superconducting detectors make up for the tripped breakers and deactivated WiFi.

A large portion of the design and simulation work presented in Chapter 3 was originally published as “Design and Development Status of MKID Integral Field Spectrographs for High Contrast Imaging”, in Proc. of AO4ELT4, 2015.

Finally, I very gratefully thank our funding agencies. DARKNESS was funded through an NSF ATI grant, AST-1308556. Much of the UVOIR MKID development that enabled DARKNESS was funded through various NASA grants, and I was funded through most

of my Ph.D. by a NASA Office of the Chief Technologist Space Technology Research Fellowship (NSTRF), which afforded me many unique and exciting opportunities I would otherwise have not experienced.

Now some personal acknowledgements:

There are countless people to thank for their love, support, and/or guidance along this doctoral journey. I especially would not be here today without my wise Ph.D. advisor Ben, my labmates and classmates, my friends and family, and my wonderful wife Erin.

Ben, thank you for your many years of guidance, and for passing on a tremendous amount of your knowledge. I can't say you taught me everything I know, but its certainly a large fraction. This project would not have been possible without your vision, and we would not have made such rapid progress with anyone else helming the ship. Despite the sometimes grueling observing trips, occasional occupational hazards, and subsequent jokes about said hazards (my eye is fine, thanks) it has truly been a pleasure working with you. Also thanks for all the Palomar snacks and for having a sense of humor.

I must thank all of my labmates, past and present, for all of their hard work and for making the lab an all-around fun and enriching place to be (as much as it can be when you're opening the DARKNESS cryostat at midnight for the third time in two weeks). Thank you all for teaching me so much, and giving me opportunities to share my knowledge in return.

I want to thank my family for supporting me not only during my Ph.D., but throughout my education. I hope I've made you all proud. Please stop telling people I'm a rocket scientist though.

To my friends: I want to thank you all for enduring seven years of hearing, "Sorry I can't make it, I'll be in the lab :(" and still being my friends. I think we still managed to have some pretty good times.

Finally, to my wife Erin. It's amazing to think that I met you on week 1 of grad school and we now enter the next phase of life together. Thank you for being my partner in life, for supporting me in my goals and inspiring me with your ambition since the day we met.

Curriculum Vitæ

Seth Richard Meeker

Education

- 2017 (Expected) Ph.D., Physics with an Astrophysics Emphasis, University of California, Santa Barbara
- 2014 M.A., Physics with an Astrophysics Emphasis, University of California, Santa Barbara
- 2009 B.S., Astrophysics, University of California, Los Angeles

Honors & Awards

- 2017 NASA Postdoctoral Program Fellowship (declined)
- 2011-2015 NASA Space Technology Research Fellowship (NSTRF)
- 2011 NSF Graduate Research Fellowship Honorable Mention

Selected Publications

“DARKNESS: a Microwave Kinetic Inductance Detector integral field spectrograph for high-contrast astronomy”, **S. R. Meeker**, B. A. Mazin, A. B. Walter, et al. *PASP* (submitted)

“Design and Development Status of MKID Integral Field Spectrographs for High Contrast Imaging”, **S. R. Meeker**, B. A. Mazin, R. Jensen-Clem, et al. *Proc. of AO4ELT4*, 2015

“The ARCONS Pipeline: Data Reduction for MKID Arrays”, J. C. van Eyken, M. J. Strader, A. B. Walter, **S. R. Meeker**, et al. *ApJS*, 219, 14, 2015

“ARCONS: A 2024 Pixel Optical through Near-IR Cryogenic Imaging Spectrophotometer”, B. A. Mazin, **S. R. Meeker**, M. J. Strader, et al. *PASP*, 125, 1348-1361, 2013

“A superconducting focal plane array for ultraviolet, optical, and near-infrared astrophysics”, B. A. Mazin, B. Bumble, **S. R. Meeker**, et al. *Optics Express*, 20, 1503, 2012

Abstract

DARKNESS: The First Microwave Kinetic Inductance Detector Integral Field Spectrograph for Exoplanet Imaging

by

Seth Richard Meeker

High-contrast imaging is a powerful technique for the study of exoplanets. Combining extreme adaptive optics to correct for atmospheric turbulence, a coronagraph to suppress diffraction from the telescope aperture, and an integral field spectrograph to obtain a spectrum at every spatial element in the final image, ground-based high contrast instruments can effectively remove on-axis star light to characterize nearby faint companions and disks. Current state-of-the-art high-contrast imagers operating at near-infrared wavelengths regularly achieve contrast ratios $< 10^{-6}$ at $0.5''$ separations. For young systems ($\lesssim 10$ Myr) at 10 pc, this roughly translates to detectability of Jupiter mass planets in 5 AU orbits. Tighter separations may be achieved with larger telescope apertures, but deeper contrasts are limited from the ground by residual atmospheric aberrations. Unsensed and uncorrected wavefront aberrations lead to a pattern of coherent speckles in the final image that evolve on a range of timescales from a few milliseconds to tens of minutes. The most problematic speckle population, referred to as *atmospheric* speckles, have lifetimes of roughly 1 s causing them to average slowly in long exposures. After

subtraction of the long lived quasi-static speckles in post-processing, atmospheric speckle noise sets the ultimate contrast limits.

In this thesis we present DARKNESS (the DARK-speckle Near-infrared Energy-resolving Superconducting Spectrophotometer), the first demonstration platform to utilize optical/near-infrared Microwave Kinetic Inductance Detectors (MKIDs) for high-contrast imaging. The photon counting and simultaneous low-resolution spectroscopy provided by MKIDs enable real-time speckle control techniques and post-processing speckle suppression at framerates capable of resolving the atmospheric speckles. We describe the motivation, design, and characterization of DARKNESS, its deployment behind the PALM-3000 extreme adaptive optics system and the Stellar Double Coronagraph at Palomar Observatory, early speckle characterization results at \sim ms timescales, and future prospects for implementing this data in useful speckle removal schemes.

Contents

1	Introduction Part I: Exoplanet Imaging	1
1.1	The Current High-Contrast Haul	5
1.2	Instrumentation	9
1.2.1	Extreme Adaptive Optics	13
1.2.2	Coronagraphy	17
1.3	Speckles	23
1.3.1	Speckle Statistics	27
1.3.2	Speckle Suppression Techniques	29
1.4	High Contrast Instrumentation at Palomar Observatory	34
1.4.1	PALM-3000 Extreme Adaptive Optics	35
1.4.2	Project 1640	36
1.4.3	The Stellar Double Coronagraph	36
1.5	Organization of this Thesis	37
2	Introduction Part II: Microwave Kinetic Inductance Detectors	39
2.1	Operating Principle	40
2.2	Tuning MKID sensitivity	45
2.3	Recent MKID Advancements	50
2.4	ARCONS	51
3	DARKNESS Design	52
3.1	MKID Array	54
3.1.1	Pixel Design	55
3.1.2	Array Layout	58
3.1.3	Fabrication	62
3.2	Cryostat	64
3.2.1	Mounting to the AO bench	66
3.2.2	Wiring	67
3.3	Optical Design	69
3.3.1	Modifications to the SDC	74
3.4	Readout	77
3.4.1	Electronics Rack	80
3.5	Performance Simulations	80

3.5.1	Constraining HR 8799	85
4	DARKNESS Characterization and Commissioning	87
4.1	In-lab Verifications	88
4.1.1	MKID Quality and Yield	88
4.1.2	D-3 Sensitivity and Energy Resolution	90
4.1.3	DARKNESS Throughput	91
4.2	On-sky Commissioning	93
4.3	Reduction and Analysis Pipeline	95
4.3.1	Optical Checkout	96
4.3.2	On-sky Contrast	99
4.3.3	Pupil Viewer	101
5	Studying Speckle Lifetimes and Implications for Statistical Speckle Discrimination	103
5.1	Speckle Lifetime Review	104
5.2	Observations	107
5.3	Autocorrelation Analysis	108
5.3.1	A Seconds-Long Decorrelation of Possible Instrumental Origin . .	110
5.3.2	A Strong Anti-correlation at 12 ms Spacing	112
5.4	Implications for Statistical Speckle Discrimination	115
5.4.1	Leveraging Autocorrelation Analysis for SSD	115
5.4.2	First Attempt at SSD with DARKNESS	116
6	Future Work & Conclusions	120
A	QE Testbed	123
A.1	QE Testbed Design	124
A.2	QE Testbed Verification	127
A.3	Could we use a smaller $f/\#$?	130
	Bibliography	133

List of Figures

1.1	All confirmed exoplanets plotted with mass vs. orbital separation.	3
1.2	Discovery images of HR8799e.	7
1.3	L-T transition color-magnitude diagram	8
1.4	Meta-analysis of first-generation exoplanet imaging survey results.	10
1.5	Black body flux of several Solar System bodies and a putative hot Jupiter, as seen from 10 pc away.	12
1.6	Contrast of an Earth-like planet compared against the diffraction pattern from a 4-m telescope.	13
1.7	Schematic of a basic AO system.	15
1.8	Contrast vs. rms WFE	16
1.9	Layout of a basic Lyot coronagraph.	18
1.10	Cross sections from 1-D coronagraph	19
1.11	Pupil apodizer from P1640.	22
1.12	Operating principle of a vector vortex coronagraph.	24
1.13	Simulations of sinusoidal phase aberrations over a telescope pupil resulting in image plane speckles.	26
1.14	Recent state-of-the-art contrast curves from SPHERE and GPI.	27
1.15	Modified Rician PDFs with varying I_c/I_s	28
1.16	Example of spectral differential imaging for speckle suppression.	31
1.17	Example of stochastic speckle discrimination for speckle suppression.	33
1.18	Previous speckle nulling results from Palomar.	35
2.1	MKID operating principle.	42
2.2	Example of two MKID resonator geometries.	43
2.3	Frequency domain multiplexing circuit diagram and frequency comb.	44
2.4	Typical phase pulse from a UVOIR MKID.	45
2.5	Examples of two simulated resonators with varying Q_c	48
3.1	DARKNESS with the SDC mounted in the Hale Telescope Cassegrain cage.	53
3.2	D-1 base pixel design schematic.	56
3.3	D-1 sensitivity verification.	59
3.4	D-1 frequency map showing how resonators are placed spatially in the array.	61
3.5	30 MHz frequency sweep of a D-1 feedline.	62

3.6	Comparison of individual pixel design and orientation from D-1 compared to SCI-4 from ARCONS.	63
3.7	DARKNESS cryostat schematic.	65
3.8	DARKNESS cryostat photos.	67
3.9	Transmission curve of DARKNESS's cryogenic IR-blocking/bandpass filter.	70
3.10	CAD design of DARKNESS relay optics with Zemax ray trace overlay.	72
3.11	Zemax spot diagrams at DARKNESS MLA.	73
3.12	Photographs of Lyot coronagraph optics installed in the SDC for DARKNESS.	76
3.13	Block diagram of 2nd generation UVOIR MKID readout.	79
3.14	Photograph of DARKNESS electronics rack.	81
3.15	PROPER simulated speckle patterns of P1640 coronagraph at DARKNESS's observing wavelengths.	83
3.16	PROPER simulated contrast curves assuming P1640 optics and Dark-speckle post processing.	84
4.1	D-3 array characterization.	89
4.2	Typical D-3 808 nm pulses.	91
4.3	Histogram of D-3 energy resolutions.	92
4.4	DARKNESS measured/theoretical throughput vs. wavelength.	94
4.5	DARKNESS commissioning image gallery.	97
4.6	Mean J-band contrast vs. angular separation.	100
4.7	Pupil image from DARKNESS relay optics finder camera/pupil imager.	102
5.1	P3K internal white light (WL) source vs. π Her speckle correlation comparison	109
5.2	Seconds decorrelation	111
5.3	Milliseconds decorrelation	112
5.4	40 Hz Speckle modulation.	113
5.5	Shapiro-Wilk test for normality as a function of speckle exposure time.	116
5.6	SAO65921 speckle vs. companion intensity distributions.	117
5.7	SAO65921 speckle vs. companion Shapiro-Wilk test.	119
A.1	QE testbed design.	125
A.2	QE testbed Zemax spot diagram at photo-diode pupil plane.	128
A.3	Zemax simulation of QE testbed image plane uniformity.	129
A.4	Measured QE testbed image plane uniformity using variable aperture.	130
A.5	DARKNESS lightcurve of a QE measurement.	131
A.6	Zemax simulation of enclosed energy at MLA focus for varying $f/\#$	132

List of Tables

3.1	DARKNESS Overview	54
3.2	D-1 Design Parameters	56
3.3	Selectable optics parameters	77

Glossary of Acronyms

ACF — Autocorrelation Function
ADC — Analog to Digital Converter
ADI — Angular Differential Imaging
ADR — Adiabatic Demagnetization Refrigerator
AO — Adaptive Optics
APLC — Apodized Pupil Lyot Coronagraph
AR — Anti-reflection
ARCONS — the ARray Camera for Optical to Near-infrared Spectrophotometry
BD — Brown Dwarf
CCD — Charge-coupled Device
CPW — Co-planar Waveguide
DAC — Digital to Analog Converter
DARKNESS — DARK-speckle Near infrared Energy Resolving Superconducting Spectrophotometer
DM — Deformable Mirror
DS — Dark Speckle
FOV — Field of View
FPM — Focal Plane Mask
GPI — the Gemini Planet Imager
HEMT — High Electron Mobility Transistor
IDC — Interdigitated Capacitor
IFS — Integral Field Spectrograph
IS — Integrating Sphere
IWA — Inner Working Angle
MKID — Microwave Kinetic Inductance Detector
MLA — Microlens Array
MR — Modified Rician
NCPA — Non-Common Path Aberrations
OAP — Off-axis Parabola mirror
P3K — PALM-3000 adaptive optics system
P1640 — Project 1640
PD — Photodiode
PDF — Probability Density Function
PHARO — the Palomar High Angular Resolution Observer
PSD — Power Spectral Density
PSF — Point Spread Function
PtSi — Platinum Silicide
QE — Quantum Efficiency
ROACH — Reconfigurable Open-Architecture Computing Hardware
SCEAO — the Subaru Coronagraphic Extreme Adaptive Optics system

SDC — the Stellar Double Coronagraph
SDI — Spectral Differential Imaging
SPHERE — the Spectro-Polarimetric High-contrast Exoplanet REsearch instrument
SR — Strehl Ratio
SSD — Stochastic Speckle Discrimination
TiN — Titanium Nitride
UVOIR — Ultraviolet, Optical, and Near-infrared
VIP — Vortex Image Processing pipeline
VVC — Vector Vortex Coronagraph
WFE — Wavefront Error
WFS — Wavefront Sensor
WL — White Light
XAO — Extreme Adaptive Optics

Chapter 1

Introduction Part I: Exoplanet

Imaging

The study of exoplanets has undergone a veritable explosion since the first discovery of two *bona fide* planetary mass companions around a millisecond pulsar in 1992 (Wolszczan & Frail, 1992). In the subsequent 25 years, nearly 3000 exoplanets have been confirmed¹ using a variety of discovery techniques. The (relatively brief) history of the field has been a tale of constant surprise from our solar-centric perspective. For example:

- The first exoplanets, PSR B1257+12 a and b mentioned above, were discovered around a millisecond pulsar formed by the merging of two white dwarfs, and were ultimately two of only four such "pulsar planets" discovered to date.

¹2950 to be exact, according to exoplanets.org as of this writing in July 2017

- The first exoplanet discovered around a main sequence star, 51 Pegasi b (Mayor & Queloz, 1995) , was a Jupiter-mass companion orbiting its solar-type host every 4 days — or roughly one-tenth the planet-star separation of Mercury from our sun — and represented the first of an unexpected population of “hot Jupiters.”
- The first exoplanet to be imaged directly, 2M1207b, orbits its late-type (M8) host very far out (at ~ 40 AU) yet has an estimated mass of $\sim 5 \pm 2 M_{Jup}$ and can be spectrally classified as a brown dwarf (Chauvin et al., 2004; Ricci et al., 2017).
- The most successful planet-finding mission to date, the Kepler Space Telescope, set out with the lofty goal of constraining the prevalence of Earth-size exoplanets in our galaxy, and in the process discovered that the *actual* most prevalent class of planets has sizes between those of Earth and Neptune² — a class for which we had no precedent in our Solar System (Batalha et al., 2013; Fressin et al., 2013; Howard et al., 2012).

This is but a small sample of the unexpected discoveries that upended much of our understanding of planet formation, migration, and evolution, which was previously informed entirely by the composition and organization of our own Solar System.

To provide an exhaustive history of the field is clearly beyond the scope of this thesis (see Seager (2010) or Perryman (2011) for comprehensive, though slightly outdated overviews of the field), but we want to highlight that the milestones listed above all

²...initially believed to be a single broad population and referred to as either “Super Earths” or “sub-Neptunes” depending on how much you like to receive funding, though recent work shows that there may indeed be two distinct classes of planets with radii between 1 to 6 R_{\oplus} (Fulton et al., 2017).

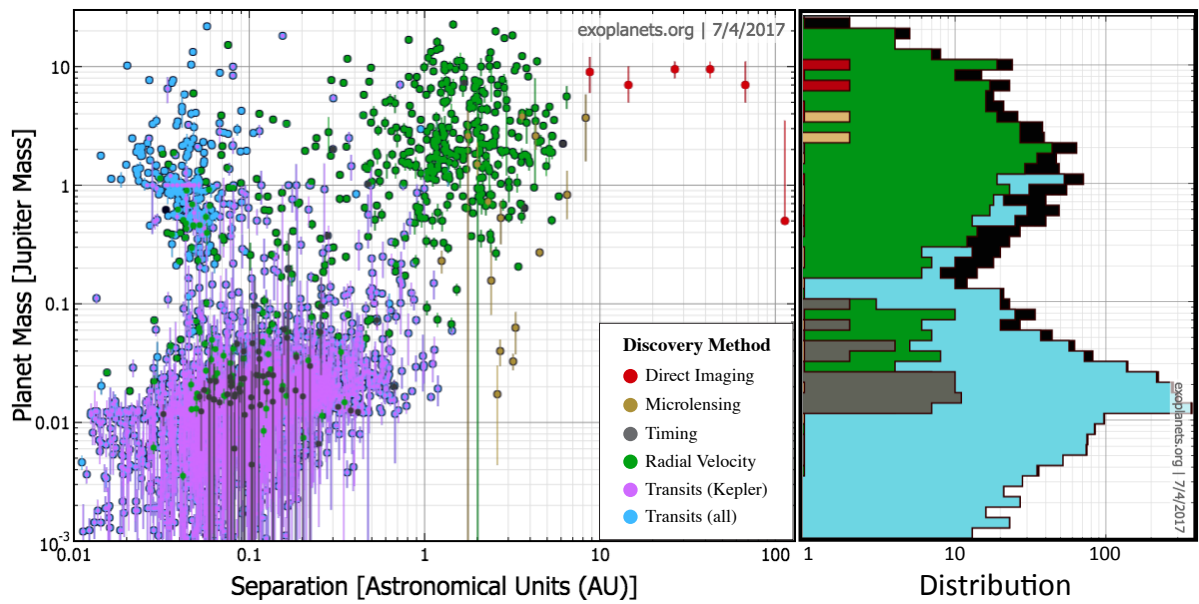


Figure 1.1: All confirmed exoplanets (according to exoplanets.org as of July 4, 2017). (Left) Scatter plot showing the mass vs. orbital distance from the host star for all exoplanets. (Right) The mass distribution of all known exoplanets, presented as a projection of the y-axis from the left plot. The cumulative distribution is given in black; otherwise both plots share the same legend.

utilized a different observing technique. Figure 1.1 gives the standard summary plot of the confirmed exoplanets to date presented as a function of their orbital distance and occurrence rate vs. mass. When distinguished by discovery technique it's clear that different methods are sensitive to different regions of mass/separation parameter space, and have varying degrees of success as "discovery" tools. In this thesis we will focus specifically on direct imaging, a technique that evolves in an "outside-in" manner with ongoing development, compared to the more "inside-out" evolution of transit and radial velocity surveys that must gain sensitivity to exoplanets at larger separation by waiting out their longer orbital periods. While direct imaging has not proven to be as exciting as a discovery tool, it is yet to reach its full potential in this regard. As we'll discuss more in the next section, it seems the universe simply does not produce many exoplanets in the mass/separation space where direct imaging is currently sensitive. These limits, however, are not fundamental. One of the merits of direct imaging is that once the technology pushes these limits to fainter planets at smaller separations, direct imaging will be able to detect exoplanets immediately and in single observations, without requiring a fortuitous orbital inclination relative to our line of sight (like transits), a years long time baseline to observe enough orbits to build a detectable signal (like transits and radial velocity), or constant monitoring to capture very rapid or singular events (like transits and microlensing).

Another significant merit to high-contrast imaging is, as the only observing technique that allows the direct observation of a planet's light rather than the planet's effect on a

host (or lensed) star’s light, it provides a straightforward way to spectroscopically characterize exoplanet atmospheres.³ Its true utility is likely not as a discovery technique, but as a characterization tool. In the near future, when the detection limits of radial velocity and direct imaging surveys overlap, the combined information from the two will yield an unprecedented capability for constraining exoplanet masses, orbits, bulk composition, and atmospheric chemistry simultaneously. However, we are not there yet, and making this future a reality is the primary goal of DARKNESS. This instrument serves as a pathfinder for high-contrast imaging with a new type of detector technology — Microwave Kinetic Inductance Detectors (MKIDs) — a superconducting detector technology that provides read-noise free photon counting and dispersionless low-resolution spectroscopy at optical and near-IR wavelengths. In this chapter we will delve deeper into the current state-of-the-art in high-contrast imaging, and illustrate places where MKIDs may meet the technological needs to advance the field.

1.1 The Current High-Contrast Haul

Even within the already fast-paced field of exoplanet research, high-contrast imaging is still evolving rapidly and is especially instrument-driven. Early ground-based surveys in the mid to late- 2000’s doubled as testbeds for high-contrast instrumentation and algorithm development while putting strong constraints on the frequency of young, giant

³We acknowledge that transit spectroscopy is another popular technique, however, it can only be applied to close-in, transiting exoplanets with highly inflated atmospheres, making it of limited utility beyond hot Jupiters and hot Super Earths.

exoplanets at large separations (~ 10 s to 100 s of AU) from their host stars (see Bowler (2016) for a comprehensive review of these first-generation surveys). Despite surveying hundreds of stars across spectral types B through M with high sensitivity to planets $\gtrsim 5M_{Jup}$ beyond ~ 30 AU separation, these surveys discovered only a handful of planetary mass companions (see Table 1 in Bowler (2016)).

One of the earliest and most well known systems, HR8799 (see Figure 1.2), provides a perfect case study for the science potential of direct imaging (see Currie (2016) for a comprehensive review of the HR8799 literature). Upon discovery, the four planetary mass companions immediately raised questions about how we draw the line between sub-stellar objects (brown dwarfs or BDs) and exoplanets. The system architecture resembles a planetary system — as there are no known systems where multiple BDs independently orbit a single star — and although they spectroscopically resemble BDs, their infrared colors clearly place them off the typical L-T transition expected from field BDs (see Figure 1.3). On the other hand, estimates for their mass (admittedly using bolometric corrections and cooling curves established from observations of field BDs) originally placed them near the Deuterium burning limit ($\sim 13M_{Jup}$) and in-situ formation of such large objects at large separations challenged prevailing core-accretion theories for planet formation (Marois et al., 2008b, 2010). The question becomes: can you form planetary mass companions via gravitational instability? Or *vis-a-versa*: can you form sub-stellar objects via core-accretion? Should these objects be defined by their formation mechanism or by the Deuterium burning mass limit?




Figure 1.2: Discovery images of HR8799e from Marois et al. (2010). Planets b, c, and d were previously discovered in Marois et al. (2008b). See text for details about this system.

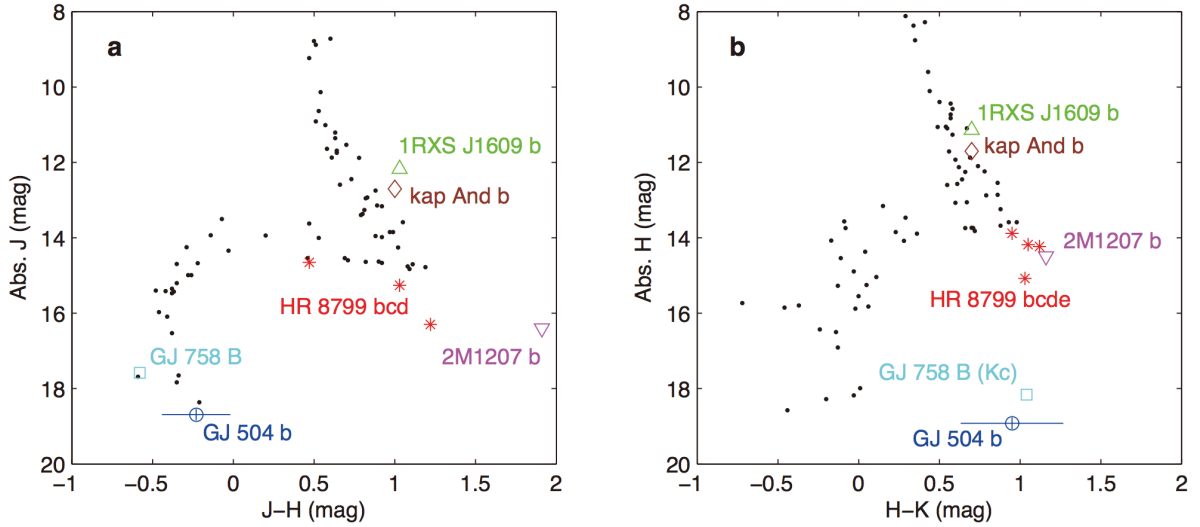


Figure 1.3: Color-magnitude diagrams for several directly imaged planetary mass companions, including HR 8799 bcde, with field LT-type brown dwarfs shown as black dots. From Kuzuhara et al. (2013).

These questions are certainly not unique to HR8799bcde, nor have they been fully resolved. Near-infrared spectroscopy of these planets has informed extensive studies of giant exoplanet atmospheres constraining non-equilibrium chemistry, cloud coverage, and surface gravity (Madhusudhan et al., 2011; Marley et al., 2012), and for the case of HR8799b and c provides tantalizing evidence of enhanced C/O ratio, implying in-situ formation by core-accretion (Barman et al., 2015; Konopacky et al., 2013; Öberg et al., 2011). A clear picture may be forming for this one system’s history, but a complete understanding of planet formation in general requires more data points with this level of characterization. Since these enigmatic planetary-mass companions have proven to be exceedingly rare at large separation, the hope is for the next generation of high-contrast surveys to find success at smaller separations, probing closer to where giant exoplanets

are expected to form.

In just the last 6 years, this new class of dedicated high-contrast imaging instruments has come online, including Project 1640 at Palomar (P1640; Hinkley et al., 2011), The Gemini Planet Imager at Gemini South (GPI; Macintosh et al., 2006, 2014), the Spectro-Polarimetric High-contrast Exoplanet REsearch instrument at VLT (SPHERE; Dohlen et al., 2006; Zurlo et al., 2014), and the Subaru Coronagraphic Extreme Adaptive Optics system at Subaru (SCEXAO; Jovanovic et al., 2015), with the goal of imaging exoplanetary systems at separations below 10 AU.⁴ However, these instruments continue to find fewer giant exoplanets than predicted,⁵ suggesting that the discrepancy between the planet mass function extrapolated from radial velocity surveys and the true giant exoplanet mass function (see Figure 1.4) remains unresolved. To fully understand the mass distribution of giant exoplanets and their formation mechanisms, and continue pushing inward toward reflected-light imaging of terrestrial planets, we must achieve deeper contrasts at smaller angular separations. To inform a discussion of the current limitations to these observations, let's first define the technologies required to perform them.

1.2 Instrumentation

The zeroth order challenge to performing direct imaging and spectroscopy of exoplanets is overcoming the overwhelming light from their nearby host stars. Figure 1.5, taken from

⁴e.g. GPI's inner working angle (IWA) in H-band is ~ 1 AU for a star at 10 pc, with the deepest contrasts achieved by ~ 5 or 6 AU.

⁵There are two published discoveries so far between them, with only one of those actually pushing the mass/separation boundaries set by previous surveys (Macintosh et al., 2015; Chauvin et al., 2017)

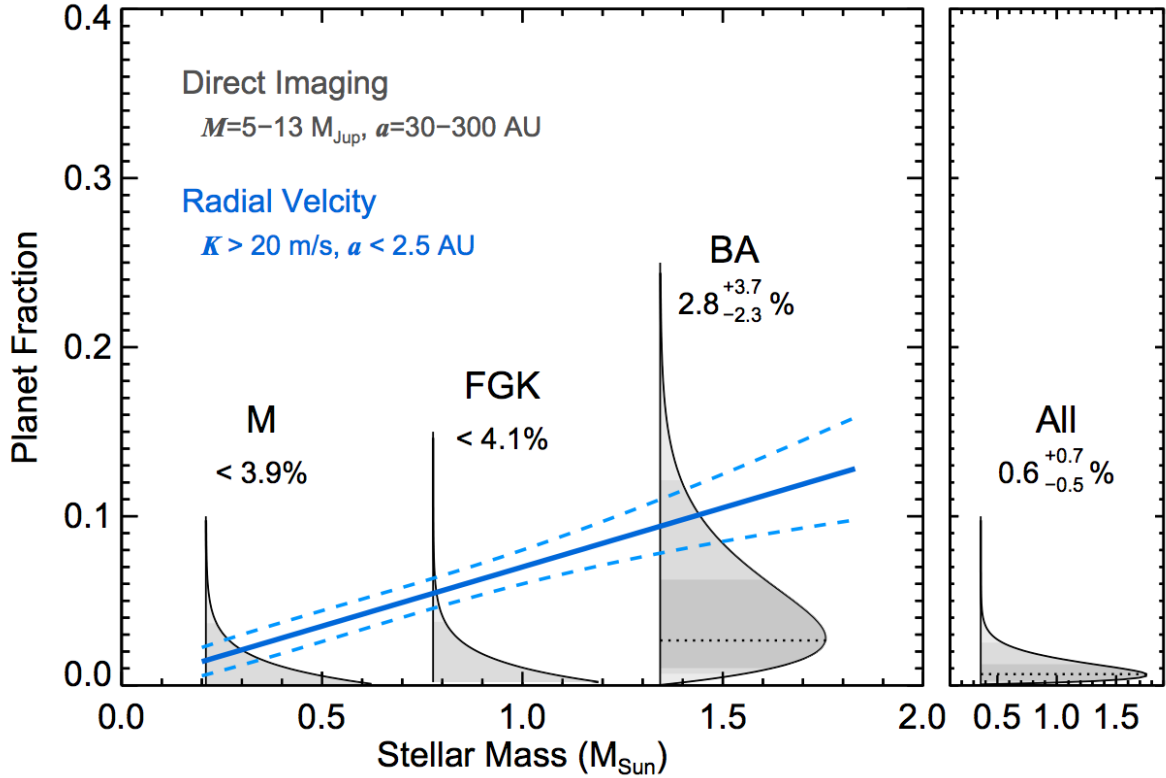


Figure 1.4: Probability distributions for occurrence rates of giant planets from meta-analysis of direct imaging surveys beyond 30 AU compared against the giant exoplanet frequency at small separations, determined by radial velocity results below 2.5 AU. This figure taken from Bowler (2016). When this style of extrapolation was originally used to infer planet yield for the GPI survey, it suggested that nearly ~ 100 planets could be discovered (Macintosh et al., 2006). In the first 2.5 years of the GPI survey, one discovery has been published (Macintosh et al., 2015).

Seager & Deming (2010), partially illustrates this challenge by comparing the blackbody fluxes of several solar system planets (and a putative hot Jupiter) with a sunlike star, as they would be observed from 10 pc. An Earth-like planet is roughly 10-billion times fainter than the host star at visible wavelengths. When we compare this contrast, at the requisite separation, to the expected diffraction in a real imaging system, as shown in Figure 1.6, the challenge becomes clear. Even with a perfect wavefront, the diffracted starlight will overwhelm faint companions at close separation, necessitating some strategy to control the diffraction and remove as much starlight as possible. Of course in reality the wavefront will not be perfect. When performing high-contrast imaging from the ground the wavefront is severely distorted by the atmosphere, but even from space aberrations in the optics will scatter unwanted starlight.

To meet these challenges, dedicated high-contrast instruments typically have a common anatomy. First, an extreme Adaptive Optics (XAO) system corrects the atmosphere’s distortion of the incoming light.⁶ After XAO the on-axis star light is removed with a coronagraph, leaving any off-axis light from a planet or disk unaffected. Finally, the remaining light is sent to an integral field spectrograph (IFS) that returns a spectrum at each spatial element in the final image. With this configuration the majority of the overwhelming starlight can be suppressed, revealing nearby faint companions or disk systems with simultaneous imaging and spectroscopy. Though we limit this discussion to current ground-based systems, it should be noted that any such future mission (including

⁶Here “extreme” means thousands of actuators performing the wavefront correction, compared to only hundreds in the AO systems that early surveys used




Figure 1.5: From Seager & Deming (2010). Black body flux of several Solar System bodies and a putative hot Jupiter, as seen from 10 pc away. The planet spectra show two peaks: one at long wavelength from their thermal emission and one at visible wavelengths from reflected starlight. From this figure we can see the typical 10^{-10} contrast quoted for exo-Earth imaging, which assumes imaging in the visible from a space based telescope to achieve the necessary angular resolution despite more favorable contrasts in the mid-IR.

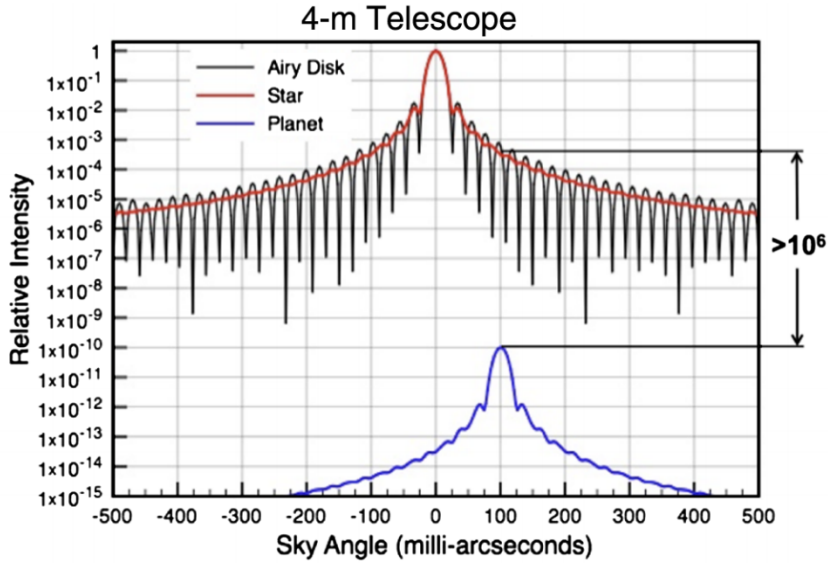


Figure 1.6: A cross section of the diffraction pattern from a perfect 4-m circular aperture telescope, measured at $\lambda = 500$ nm. A planet with Earth-like contrast (10^{-10}) and separation (100 mas at 10 pc) relative to a sunlike star is shown for comparison. This figure is taken from Lyon & Clampin (2012).

WFIRST-AFTA) will almost certainly contain the key elements listed here.

1.2.1 Extreme Adaptive Optics

The first, and in many ways most crucial subsystem in a high-contrast instrument is the XAO. Figure 1.7 shows the most basic AO schematic: an aberrated wavefront enters the telescope, bounces off a deformable mirror (DM) placed at a conjugate pupil plane, and is split by a beamsplitter to allow some science light to pass, while redirecting a portion to a wavefront sensor (WFS), also at a conjugate pupil plane, that measures the phase distortions in the wavefront. This measurement is then used to command the DM shape such that the wavefront reflected from the DM is nearly flat. The system operates in

a closed-loop, where the WFS constantly measures the residual wavefront error (WFE) and commands DM correction, typically at \sim kHz rates. With good enough correction, the diffraction limited performance of the telescope can be recovered. The metric for this performance is known as the Strehl ratio (SR), effectively the percentage of light that is gathered from the seeing halo back into the diffraction limited point spread function (PSF) core, which is well approximated by the Maréchal approximation

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

where $\sigma = \frac{2\pi\delta}{\lambda}$ and δ is the rms WFE in nm. A recent derivation of this approximation can be found in Ross (2009). Uncorrected starlight remains in the seeing halo, which has a characteristic width of λ/r_0 where r_0 is the Fried parameter, a measure of the atmospheric correlation length. From the Fried parameter and the mean wind speed, we also know the atmospheric correlation timescale, or Greenwood time delay:

$$\tau_0 = 0.314r_0/\bar{v} \quad (1.2)$$

which effectively sets the rate that an AO system must operate.

Also pertinent to high-contrast imaging, besides the quality of correction, is the area over which the wavefront can be corrected. The maximum control radius in the final image, given as θ_{max} , is directly related to the number of actuators available in the DM as:

$$\theta_{max} = \frac{N_{act}}{2} \frac{\lambda}{D} \quad (1.3)$$

where N_{act} is the number of DM actuators across the pupil. Relating the above values and scaling laws to final contrast as a function of image position is not straightforward,




Figure 1.7: The classic schematic showing the basic feedback loop of an AO system. An aberrated wavefront enters the telescope, bounces off a deformable mirror, and is then split to send some light to a wavefront sensor that measures the phase distortion of the wavefront. This information is then used to command the deformable mirror to take a shape such that the reflected wavefront is flattened. This figure is from lyot.org, courtesy of Claire Max.

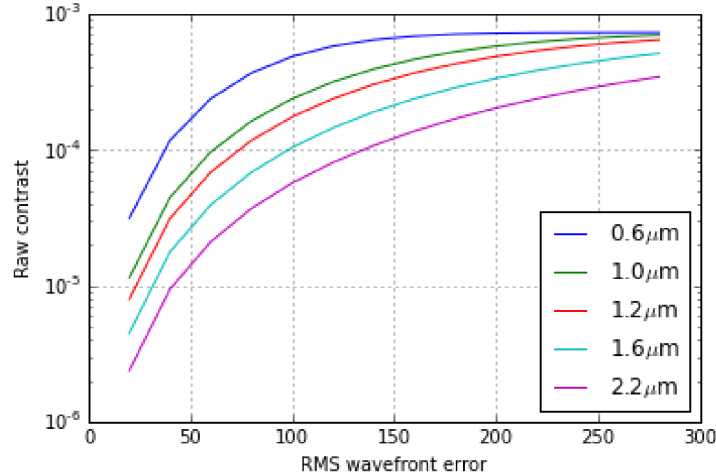


Figure 1.8: Plot of Equation 1.4 showing average contrast in an AO corrected image, as a function of rms WFE delivered by the AO system and observing wavelength. This figure taken from Milli et al. (2016).

but a rule of thumb follows from the above relations to demonstrate how vital the initial AO correction is (Serabyn et al., 2007). Since the uncorrected starlight is $(1 - SR)\%$ of total light, and it is scattered over roughly N_{act}^2 spatial elements in the final image, the mean contrast in a $1\lambda/D$ aperture can be approximated as:

$$C \approx \frac{1 - SR}{N_{act}^2} \quad (1.4)$$

Figure 1.8 plots this as a function of wavelength and rms WFE. Typical rms WFE for XAO systems is 100 nm, limited by the delay time between WFS measurement and DM correction. For a full discussion of an XAO system’s demonstrated error budget the interested reader is directed toward Poyneer et al. (2016), and for a general understanding of the vast complexity of AO systems Hardy (1998) is a landmark text and continues to be a valuable resource.

1.2.2 Coronagraphy

As we saw in Figure 1.6, even with a perfectly flat incident wavefront the diffracted starlight will still be many orders of magnitude brighter than a faint companion. A coronagraph must be used to reduce the diffracted starlight. Figure 1.9 shows a 1-D schematic of the most basic modern coronagraph, known as a Lyot Coronagraph. Starting with a perfect, uniform wavefront, an image of the star is formed and the core of the PSF is blocked with a small occulting mask in the center of the image. Any off-axis light from an exoplanet will pass unaffected. The focal plane mask has the added effect of diffracting residual starlight to the outskirts of the subsequent pupil plane. There a second mask, slightly undersized relative to the pupil and interchangeably referred to as a Lyot mask or pupil mask hereafter, removes the majority of remaining starlight while mostly ignoring off-axis planet light which has not been preferentially redirected to the outskirts. Off-axis planet light will just experience a loss in throughput determined by the % undersizing of the stop.

Sivaramakrishnan et al. (2001) provides a simple yet informative exposition of this one-dimensional design, summarized in Figure 1.10. For this discussion we first note a couple things. The coronagraph is actually acting on the complex electric field of the wavefront, denoted as $E(x)$ or $E(\theta)$ for pupil plane or focal plane coordinates, respectively, though we are accustomed to observing the wavefront as intensity, which is simply $|E|^2$. Throughout this discussion, and later in this thesis⁷ we will be making the assump-

⁷...and actually earlier when we were discussing the control radius in an XAO image being determined by the sampling of the pupil plane

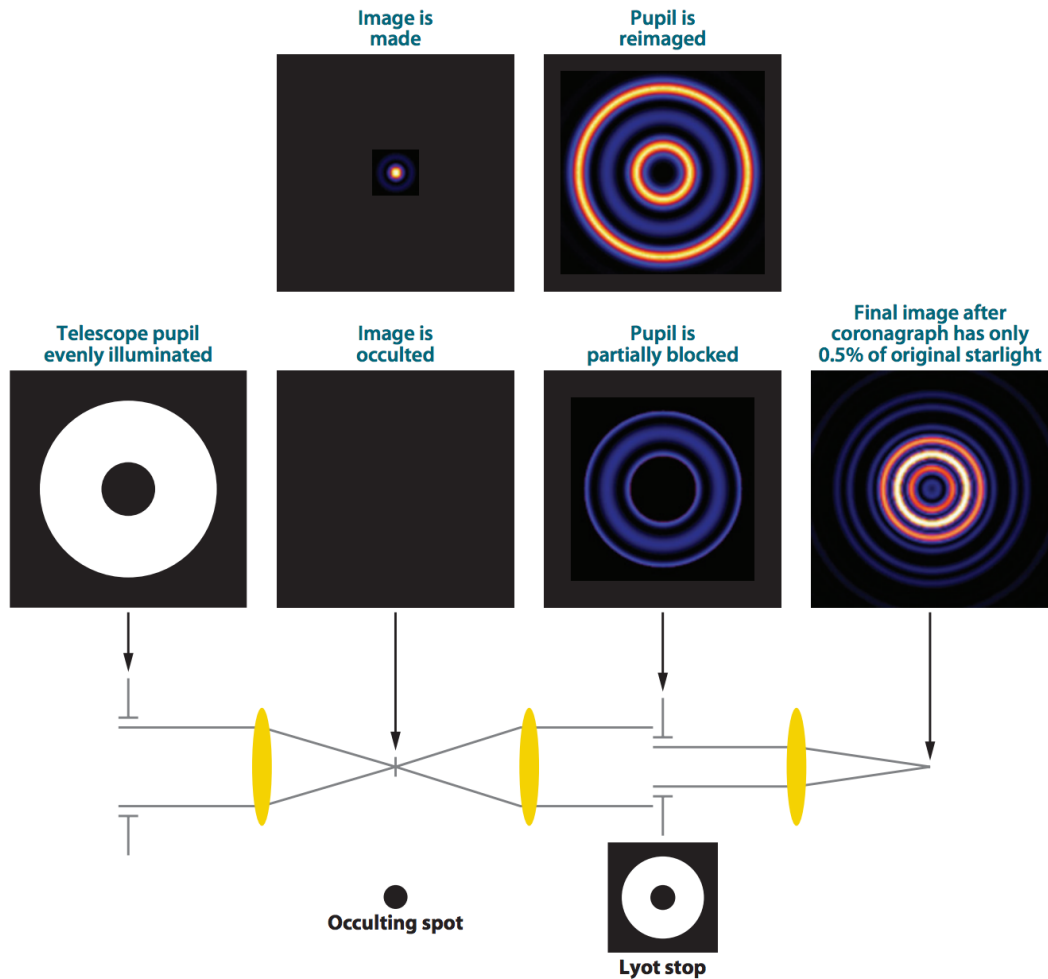


Figure 1.9: Basic Lyot coronagraph. The star is brought to a focus and blocked with an occulting mask, then the residual diffracted light is blocked at the outskirts of a subsequent pupil. Any off-axis planet light, assuming it has far enough angular separation to avoid the focal plane mask, is largely unaffected. This figure taken from Oppenheimer & Hinkley (2009)




Figure 1.10: 1-D coronagraph from Sivaramakrishnan et al. (2001). See text for details.

tion that the pupil plane and focal plane can be directly related by a Fourier transform, also known as the *Fraunhofer approximation*, which will provide valuable physical intuition for moving quickly between spatial modes in the pupil plane and PSF location in the image in Section 1.3. A handy summary of these physical optics approximations, particularly as utilized in coronagraphy, can be found in Traub & Oppenheimer (2010).

We begin with a plane wave at the telescope pupil defined as:

$$E = E_0 \operatorname{Re}(e^{i\phi}) \quad (1.5)$$

The telescope pupil is defined as a rectangular function:

$$\Pi(x) = \begin{cases} 1 & \text{for } |x| < \frac{1}{2} \\ 0 & \text{otherwise} \end{cases} \quad (1.6)$$

Such that

$$E_a = E\Pi(x\lambda/D) \quad (1.7)$$

where D is the telescope pupil diameter. Propagating from the pupil plane (a) to image plane (b) is equivalent to taking a Fourier Transform of this rectangularly masked field resulting in a sinc function (dropping proportionality constants), $E_b \propto \text{sinc}(D\theta/\lambda)$. Here they deviate from the typical hard edged focal plane mask and multiply by a transmission profile $1 - w(\theta D/\lambda s)$, with $w(\theta) = \exp(-\theta^2/2)$ defined such that if $w(\theta)$ has width of order unity, the stop will cover roughly s resolution elements (in Figure 1.10 a Gaussian mask is used with $s = 5$, equivalent to width = $5\lambda/D$). This mask blocks most of the power in the core of the PSF. Again, taking the Fourier Transform to propagate from image plane (d) to pupil plane (e), the sinc function turns back to the rectangular function, but now convolved with the Fourier Transform of the transmission profile. Carrying through the convolution gives:

$$E_e \propto \Pi\frac{x\lambda}{D} - \frac{s\lambda}{D}\Pi\frac{x\lambda}{D} * W(sx) \quad (1.8)$$

The residual diffracted light is concentrated at $\pm D/2$ in pupil plane (e), with wings related to the transform of the focal plane mask profile. The more the masked transmission in focal plane (d) resembles a pure sinusoid, by the Fourier Transform, the more the light is concentrated at a discrete location in the pupil plane, and the more effectively it

is masked with a Lyot stop.

The coronagraph discussed here is one of a large family of coronagraph designs that can operate on either the field amplitude (such as the opaque masks used above) or on the phase, preferentially moving light outside of regions of interest without sacrificing throughput or spatial coverage in the image plane — the so called inner working angle (IWA) of the coronagraph is set by the finite size of an opaque mask, but is theoretically the diffraction limit for a phase mask. A similar formalism to the one above can be found in Traub & Oppenheimer (2010) covering several coronagraph designs, including the Gaussian focal plane mask and the classic hard-edged focal plane mask from Figure 1.9. A more recent review of coronagraphs with small IWA can be found in Mawet et al. (2012). For now we will content ourselves with a quick description of two particular designs that are relevant to the work in this thesis.

Apodized Pupil Lyot Coronagraph

The apodized pupil lyot coronagraph (APLC) is a relatively straightforward extension of the Lyot coronagraph described at length above. The main idea is that the diffraction from a hard edged pupil can be minimized by “softening” the pupil edge. This can be achieved with a partially transmissive pupil mask before the standard Lyot coronagraph masks, at the expense of throughput. The ideal apodization function would gather 100% of the starlight into the PSF core, where it could be blocked with a reasonably sized, even hard edged, focal plane mask (FPM). Unfortunately this is impossible with a finite aperture, but *prolate spheroidal* functions have been found to be an adequate functional




Figure 1.11: Transmission profile of the P1640 apodizer, with detailed views of the microdot fabrication technique used to make it. This figure taken from Hinkley et al. (2011)

form for the apodization (Soummer et al., 2003), and define the apodization pattern in the P1640, GPI, and SPHERE APLCs.

Vector Vortex Coronagraph

A vector vortex coronagraph (VVC) is a style of focal plane mask that operates on the phase of the complex field, rather than the amplitude. An optical vortex is a helical phase profile in an optical field of the form $e^{i\theta}$ with θ being the azimuthal angle about the center of the optic,⁸ which creates a singularity of zero intensity at the center by total destructive interference. This null spot is then conserved and expands as a field propagates along the optical axis. Mawet et al. (2005) first proposed the use of optical vortices as a focal plane mask in a coronagraph, whereby the star's PSF is placed on the vortex, ejecting all light to the outskirts of a subsequent pupil, which is then masked by a Lyot stop. Their particular implementation is now referred to as the vector vortex coronagraph, since it creates the

⁸Apologies for the confusing multiple definitions of θ . This is the only paragraph where we use it this way, then it's back to being the spatial coordinate in the image plane for the remainder of the Chapter.

phase ramp using a half-wave plate (HWP) with spatially varying birefringence to rotate the polarization vectors of incident light as a function of azimuthal angle (see Figure 1.13). The first on-sky implementation in near-IR wavelengths utilized liquid crystal polymers (Mawet et al., 2009, 2010b). In the simplest design, with a single vortex at a focal plane and a matching pupil stop, obscurations due to an on-axis secondary mirror and its support spiders are not taken into account and will degrade the light-rejection of the vortex (Mawet et al., 2010a). To mitigate this problem, Mawet et al. (2011) propose a dual-stage configuration, in which a second vortex is added to the final image plane of the single-stage layout, again with remaining light blocked by an appropriately placed pupil stop. The principle advantages of vortex-based coronagraphs include small IWA (down to $1\lambda/D$), high throughput, a simple optical layout, and broadband performance. However, the trade-off is that a VVC requires very precise alignment making it especially susceptible to low-order aberrations that take light out of the PSF core and would still be well masked by a classic Lyot coronagraph.

1.3 Speckles

We now turn to the dominant noise source in high-contrast imaging, often referred to as “speckles.” Even after all the work described above to create a perfect wavefront and mitigate diffraction with a coronagraph, imperfect AO correction and unsensed aberrations downstream of the AO (so-called non-common path aberrations or NCPA) result in starlight escaping coronagraphic rejection, creating a pattern of coherent speckles in

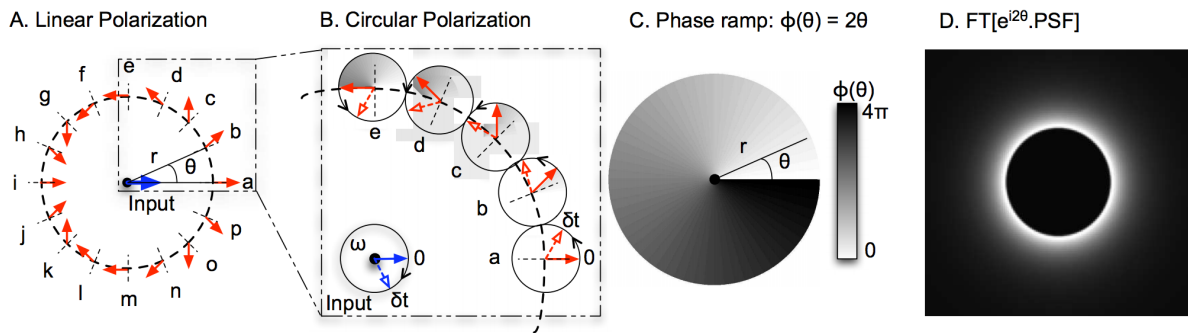


Figure 1.12: The operating principle of a vector vortex, taken from Mawet et al. (2010b). (a) A rotationally symmetric half-wave plate (HWP) with optical axis (short dashed lines) that rotates about the center of the optic. The effect on linearly polarized incident light (blue arrow) is that the polarization direction is rotated about the center by twice the azimuthal angle θ (red arrows). (b) By the definition of circular polarization being simply linear polarization rotating at some angular frequency, a rotation of the polarization vector is strictly equivalent to a phase delay. (c) Incident light is thus imprinted with a helical phase ramp. (d) Multiplying the PSF by this phase ramp generates perfect rejection at the center, which, when propagated to the subsequent pupil plane (the Fourier Transform of the PSF times this phase ramp), extends the rejection to the original pupil diameter and ejects all light outside this. See Mawet et al. (2005) Appendix C for the complete mathematical derivation of this rejection.

the final image. These speckles resemble faint companions — i.e. they have size $\sim \lambda/D$, same as a planet PSF, and comparable contrast. Thus, every modern high-contrast survey is designed with many levels of speckle mitigation as a fundamental consideration alongside the AO and coronagraph systems (we'll get back to mitigation strategies in Section 1.3.2).

Most troublesome, speckles are spatially and temporally correlated, varying on several characteristic timescales. Speckles caused by static aberrations — such as from the instrument optics — have long decorrelation times governed by changing instrumental parameters like temperature and gravity vector, and can be subtracted using a variety of differential imaging techniques that employ some form of image diversity to generate a reference PSF of the quasi-static speckle field. However, speckles resulting from residual atmospheric aberrations have decorrelation times $\lesssim 1$ second, and are typically averaged into a smooth halo during long exposures (see discussion in Section 8.1 of Hinkley et al. (2011) or Section 3.1.1 of Jovanovic et al. (2015)). Early work made different assumptions about the lifetimes of these speckles, some assuming they randomized with every AO update, and others assuming a lifetime proportional to the telescope aperture clearing time, D_{tel}/v_{wind} . Macintosh et al. (2005) find that both timescales are valid depending on error source, where speckles caused by AO measurement error will decorrelate with lifetime related to the AO update rate, but speckles caused by atmospheric effects (i.e. fitting error, bandwidth error) have lifetimes proportional to D/v regardless of r_0 . The takeaway is that the total variance of atmospheric speckles is dominated by these longer

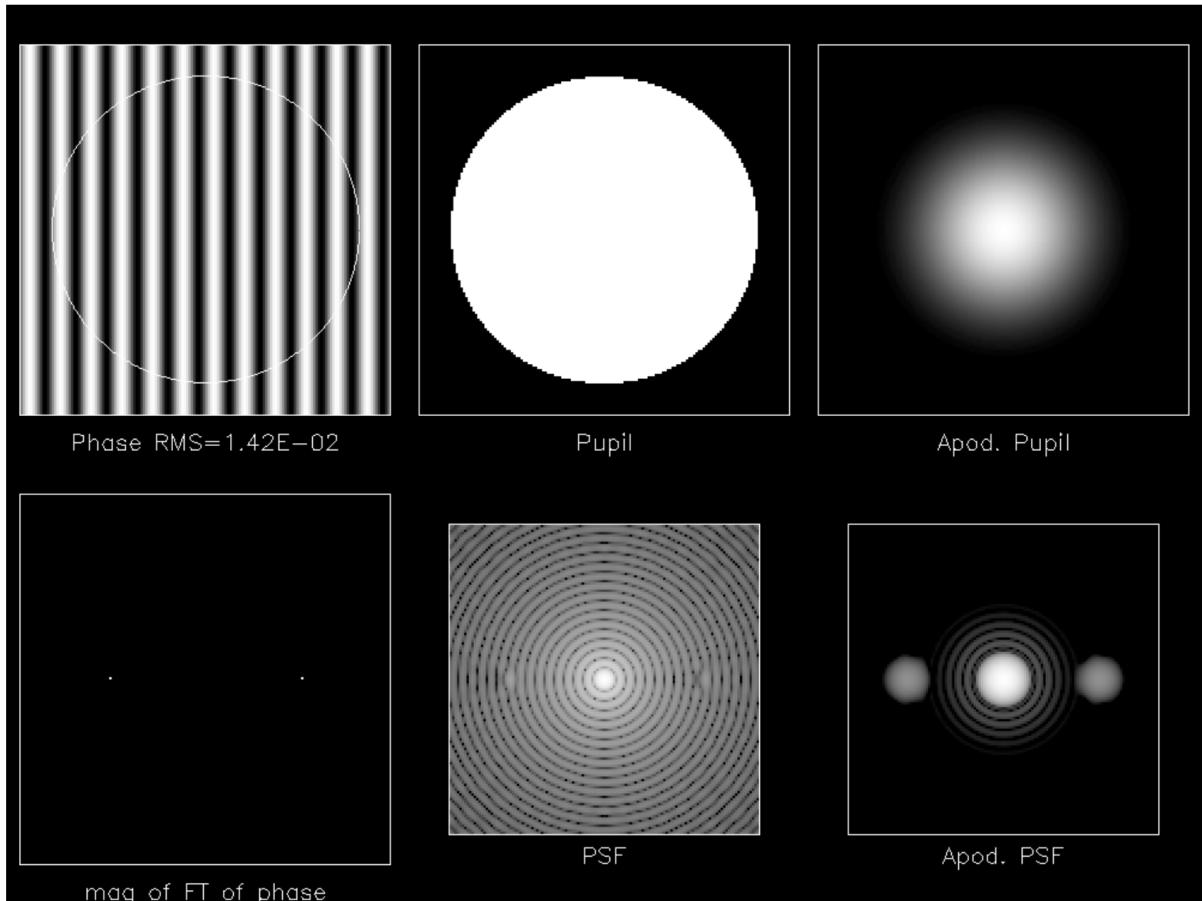


Figure 1.13: Simulation of a sinusoidal phase aberration over a telescope pupil, taken from Macintosh et al. (2005), demonstrating how a phase aberration in the pupil plane with spatial frequency $2\pi/x_0$, $\phi(x) \propto \cos(\frac{2\pi}{x_0}x)$, becomes a symmetric pair of speckles at $\pm\lambda/x_0$ in the focal plane. For a mode with N oscillations across the pupil diameter D , $x_0 = D/N$ and speckles appear N beam widths from the center of the image: $\theta = \pm N\frac{\lambda}{D}$. Larger amplitude oscillation will generate brighter speckles. In the presence of diffraction the speckles are “pinned” to the Airy pattern (Perrin et al., 2003; Aime & Soummer, 2004). If diffraction is removed (say, by a coronagraph or pupil apodization) the speckle halo becomes the dominant noise source.

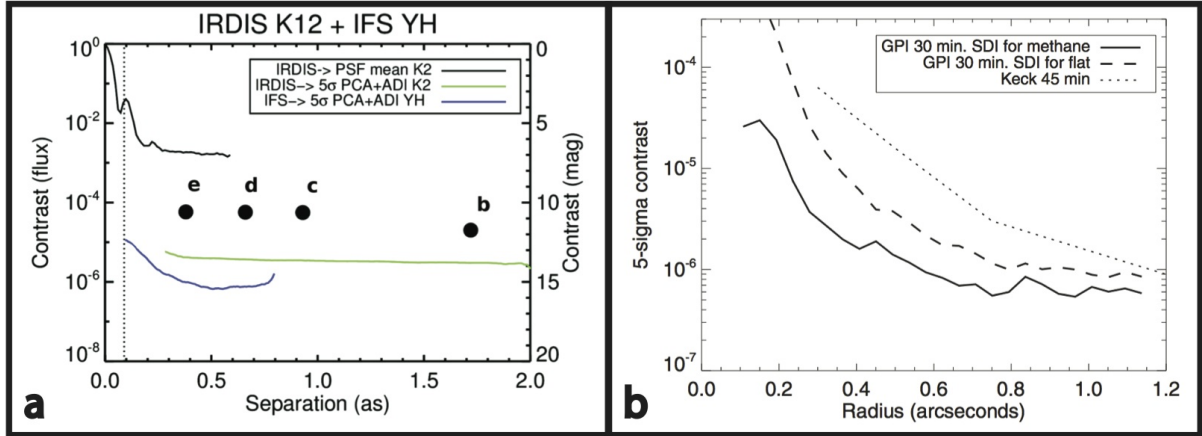


Figure 1.14: Recently published contrasts curves from (a) SPHERE (Zurlo et al., 2016) and (b) GPI (Macintosh et al., 2014), the two instruments currently capable of delivering the highest contrasts. Both plots show 5- σ contrast as a function of separation from the host star after a variety of point-spread function (PSF) removal techniques to mitigate residual speckles (the SPHERE curve also includes the K2-band photometry of HR8799’s four known companions). Both achieve roughly 10^{-6} planet-star contrasts by ~ 0.5 arcsecond separations on bright stars under ideal conditions.

lifetimes, averaging down slowly in long exposures and imposing the current state-of-the-art planet-star contrast limits across all high-contrast instruments: roughly 10^{-6} , corresponding to detectable exoplanets with masses larger than Jupiter (see Figure 1.14).

1.3.1 Speckle Statistics

A significant body of work has been dedicated to the underlying probability density function (PDF) from which a speckle’s intensities are drawn, which is known to be a modified Rician (MR):

$$p_{MR}(I) = \frac{1}{I_s} \exp\left(-\frac{I + I_c}{I_s}\right) I_0\left(\frac{2\sqrt{II_c}}{I_s}\right) \quad (1.9)$$

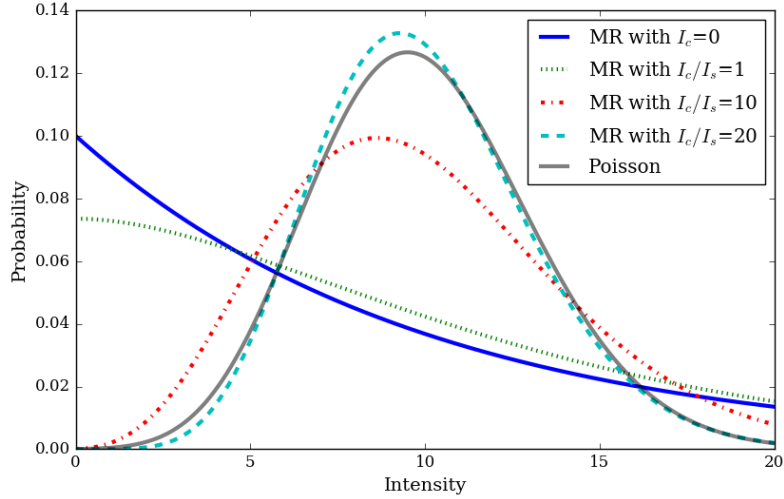


Figure 1.15: Several modified Rician (MR) PDFs with various values of I_c/I_s , compared against a Poisson PDF, all with the same expectation value of 10.

characterized by the static PSF contribution, I_c , and random speckle intensities, I_s , where $I_0(x)$ is the zero-order modified Bessel function of the first kind (see Goodman (2005); Perrin et al. (2003); Aime & Soummer (2004); Fitzgerald & Graham (2006); Soummer et al. (2007) to name a few). The mean and variance of I are:

$$\mu_I = I_c + I_s \quad (1.10)$$

$$\sigma_I^2 = I_s^2 + 2I_c I_s \quad (1.11)$$

The ratio I_c/I_s essentially characterizes the skewness. Figure 1.15 shows a few MR PDFs for reference. In the limit $I_c \rightarrow 0$ we get the exponential statistics of “pure speckle” which we will use later.

When a large number of independent and identically distributed (i.i.d.) speckles are co-added, their statistics will become Gaussian by the central limit theorem, which is what most studies assume when quoting “ $5\text{-}\sigma$ ” contrast curves. Most high-contrast observations rely on whitening of the statistics, whereby PSF subtraction (described in the next section) and other post-processing removes the correlated noise from (quasi-) static speckles, leaving the atmospheric speckles to average together. Assuming atmospheric speckles decorrelate quickly relative to the exposure time, they can be considered i.i.d. and Gaussian statistics can be applied (see discussion in Section 1.1 of Mawet et al. (2014)). However, care must be taken when making this assumption, as improper treatment of the speckles’ statistics (i.e. assuming Gaussianity when speckle noise is still correlated and thus the i.i.d. criterion is not true) can lead to severely overestimated confidence levels (Marois et al., 2008a).

1.3.2 Speckle Suppression Techniques

Speckle suppression, either optically with some form of on-sky wavefront calibration or in post-processing, is as critical to achieving high-contrast as the technologies described earlier, and informs both the instrument design and observing strategy. For most post-processing techniques, the trick is to generate a library of reference PSFs that capture the quasi-static speckle pattern. The most straightforward way to achieve this is via Reference Differential Imaging (RDI) where a nearby reference star is observed, ideally with similar apparent magnitudes in the chosen WFS and science bands, such that the

optical configuration and AO performance are as close as possible to those of the science target, resulting in similar quasi-static speckle patterns. The most popular observing strategy is Angular Differential Imaging (ADI; Marois et al., 2006) wherein the instrument rotator is turned off (for an alt/az telescope) allowing the field of view to rotate around the image center while optical aberrations (i.e. the speckle pattern) remains fixed. Principal Component Analysis (PCA) or some other similar algorithm (Lafrenière et al., 2007; Soummer et al., 2012; Meshkat et al., 2014) can then be applied to this observation sequence to remove the dominant, constant speckle pattern from each frame, and the frames are then de-rotated by the known amount due to sidereal motion and stacked to bring out any faint companions. Of course, for this technique to be successful a significant amount of image rotation must be present for any planets to move noticeably relative to the speckles, making it insensitive to speckles that change on timescales shorter than 10s of minutes. Below we will highlight a few techniques that do not inherently rely on long timescales, and would especially benefit from the existence of a read-noise free IFS capable of kHz frame rates.

Spectral Differential Imaging

Spectral Differential Imaging (SDI, sometimes referred to as Spectral Deconvolution; Sparks & Ford, 2002) takes advantage of the chromatic nature of the speckles to distinguish them from faint companions. Figure 1.16 shows an example of the basic concept. A speckle's position in the focal plane changes radially proportional to λ , while any faint companions will be fixed. By scaling the speckle pattern in a shorter wavelength frame




Figure 1.16: Series of SDI images, taken from Crepp et al. (2011). An image from a single spectral channel in an IFS data cube (a) is stretched to match an image at a longer wavelength (c). The speckle structure is clearly correlated in the middle two images, making the stretched image (b) a valid reference PSF for subtraction. In the difference image (d) the faint companion is clearly visible since its position is independent of wavelength.

from an image cube to match that at a longer wavelength, the scaled speckle pattern can be used as a reference and subtracted. This speckle suppression strategy encourages the use of an IFS, which is already desirable for spectral characterization of any observed planets, and is thus standard on all current generation high-contrast instruments (Hinkley et al., 2008; Larkin et al., 2014; Claudi et al., 2008; Groff et al., 2014).⁹

Statistical Speckle Discrimination

One of the oldest suggested speckle discrimination techniques was the Dark-Speckle technique (DS; Labeyrie, 1995; Boccaletti et al., 1998, 2001). The theory states that if you take a short exposure, regions of highly destructive interference in the speckle pattern will leave completely black spots with zero photon counts. However, if this destructive interference occurs at the position of a planetary companion, the planet will contribute some intensity and a totally dark spot will not occur. By mapping the positions of these

⁹an IFS is even slated to fly with the WFIRST-AFTA coronagraph (McElwain et al., 2016).

dark speckles in subsequent short exposures as the pattern changes with time, one can generate a dark map of the region with companions appearing as bright spots. The ability to register a zero-photon event is strongly dependent on the camera’s sampling across a speckle. The deeper a dark speckle is resolved, the more easily a statistical deviation from zero can be detected. Starting from the assumption of exponential “pure speckle” statistics mentioned earlier,¹⁰ Boccaletti et al. (1998) give the following equation for companion signal to noise ratio (SNR) achievable for a given set of observation parameters:

$$SNR = \frac{N_{\star}}{R} \sqrt{\frac{tT}{\frac{j+tN_{\star}}{G}}} \quad (1.12)$$

where R is the ratio of star brightness to planet brightness, G is the AO “gain” (the ratio of star brightness to halo brightness where halo refers to the starlight remaining outside of the Airy peak), t is the individual exposure time, T is the total integration time, N_{\star} is the total counts received from the host star per second, and j is the number of pixels within a speckle area. Previous attempts to utilize the Dark Speckle method in the near-IR, where young Exoplanet contrasts are optimal for detection, have also been hampered by significant read-noise from the available IR detectors (Boccaletti et al., 2001).

Stochastic Speckle Discrimination (SSD; Gladysz & Christou, 2008, 2009) is similar to DS in that it starts from the statistics of speckle intensity, but is more general because it doesn’t care about the absolute level of speckle intensities or whether $I_c = 0$, just

¹⁰This method pre-dates our full understanding of the static speckle contribution and subsequent MR speckle PDF.

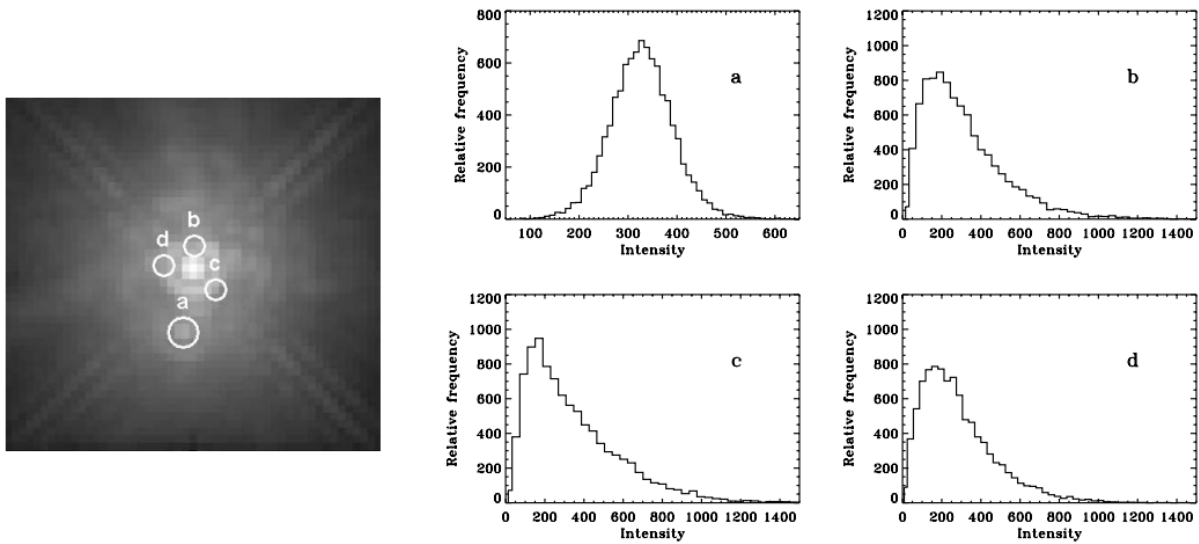


Figure 1.17: Demonstration of Stochastic Speckle Discrimination (SSD) taken from Gladysz & Christou (2008). Long exposure image is shown on the left, with histograms of intensity taken from 4 locations on the right. Location (a) contains a faint companion.

that speckle and exoplanet intensities follow different distribution functions (MR and Poissonian, respectively). By making a histogram of a pixel's intensity from thousands of successive short exposures a distribution containing only speckle contribution can be distinguished from one containing speckle+exoplanet (see Figure 1.17). To optimally implement either of these statistically based techniques requires low-noise near-IR detectors capable of \sim ms exposures, a technology that was previously unavailable (Boccaletti et al., 2001; Gladysz & Christou, 2008).

Speckle Nulling

While the previous techniques described in this section accomplished speckle discrimination in post-processing, speckle nulling (Bordé & Traub, 2006; Martinache et al., 2014) is a method for performing active speckle removal during observation. A feedback loop

is implemented parallel to the main XAO loop in which the science camera is used as a focal plane wavefront sensor (FPWFS) commanding DM corrections to correct NCPAs if the loop is slow, and possibly atmospheric speckle if the science camera can provide kHz frame rates to sustain a $\gtrsim 100$ Hz feedback loop. Since the pupil plane (where the DM resides) and the image plane are related by a Fourier Transform, a speckle’s position and intensity in the image plane gives the spatial frequency, angle, and amplitude of a corresponding sine wave on the DM. The only unknown is the sine wave’s phase offset, which can be determined with probe patterns on the DM, stepping the expected waveform through several phase values and watching the target speckle dim or brighten. With the phase measured, all the necessary information is available to apply the inverse waveform on the DM that will destructively interfere with the speckle. This process can be implemented in a closed-loop to continuously probe and null a region of speckles to dig a “dark-hole” in the image (see Figure 1.18).

1.4 High Contrast Instrumentation at Palomar Observatory

DARKNESS’s baseline design assumes integration with previous and existing high-contrast instrumentation at Palomar Observatory, and the commissioning and early science observations discussed in Chapters 4 and 5 were conducted from there. We will briefly introduce the relevant Palomar instrumentation here that will be referenced often in




Figure 1.18: Results from the Stellar Double Coronagraph’s in-lab speckle nulling loop at Palomar, operating on static aberrations with the PHARO imager, modified slightly from Bottom et al. (2016b). (a) The initial result of PSF correction still leaves many residual speckles in the focal plane. (b) Four iterations of speckle nulling remove most of the residual wavefront errors (c) Nine iterations get to within a factor of two of the detector read noise from 5 to 25 λ/D . The white polygon demarcates the control region, where contrast is improved by factors of 3-6.

Chapter 3 when discussing DARKNESS’s design.

1.4.1 PALM-3000 Extreme Adaptive Optics

PALM-3000 (P3K) is the second-generation adaptive optics (AO) facility at Palomar, and was engineered with high-contrast imaging as a primary science motivator. It is the first of the new generation of extreme AO systems with 3000 actuators in the active correction area, or roughly 64 actuators across the telescope pupil, and is designed to achieve 105 nm RMS residual wavefront error. By Equation 1.1 this translates to Strehl ratios of 58%, 76% and 88% at I, J, and H-bands, respectively. A full accounting of the predicted error budget and early on sky results can be found in Dekany et al. (2013).

1.4.2 Project 1640

Project 1640 (P1640, Hinkley et al., 2011), is the first dedicated high-contrast imaging instrument to combine the key ingredients listed above: XAO, a coronagraph, and an IFS. P3K provides the wavefront correction prior to P1640. The starlight is then suppressed with an APLC optimized for J and H band operation (Soummer et al., 2011) with an IWA of $\sim 3\lambda/D$. Integrated within the coronagraph is a precision wavefront calibration subsystem (CAL, Cady et al., 2013) that corrects for NCPA downstream of P3K using rejected light from the coronagraph focal and pupil plane masks. CAL is capable of further correcting wavefront phase errors over the entire field of view, or correcting phase and amplitude errors over half of the field. Science data is taken with a lenslet based IFS with 200×200 spaxels and spectral resolution of 30-100 (Hinkley et al., 2008). This IFS is the piece that DARKNESS was originally intended to replace when P1640 completed its nominal survey. As of this writing P1640 has conducted its final observations, but is being converted into a precision radial velocity instrument. It will still serve as our point of reference for performance simulations (see Chapter 3) as it had well understood optical properties and on-sky results as of the early stages of DARKNESS’s design.

1.4.3 The Stellar Double Coronagraph

To take advantage of the full suite of coronagraphs developed for Palomar’s Hale Telescope, DARKNESS was also slated to work with the Stellar Double Coronagraph (SDC, Bottom et al., 2016b), which was under assembly during DARKNESS’s design phase,

and typically utilizes the Palomar High Angular Resolution Observer (PHARO, Hayward et al., 2001) for its backend imager. The SDC is a flexible coronagraphic testbed with two internal focal planes and two pupil planes, allowing for a number of coronagraphic configurations. Its baseline design is for the dual-stage vortex and the apodized vortex (Mawet et al., 2011; Serabyn et al., 2011; Mawet et al., 2013), but its reconfigurable design makes it an attractive demonstration platform for new technologies and techniques (Bottom et al., 2016a, 2017), especially DARKNESS. As of 2016, the SDC became DARKNESS’s anticipated first-light front-end, requiring a re-working of the intermediate optics that are presented in Chapter 3.

1.5 Organization of this Thesis

To image exoplanet populations at smaller separations and lower masses from the ground requires a wavefront correction scheme that employs a focal plane wavefront sensor (FPWFS) scheme to sense NCPAs, operating at kHz frame rates to track the atmospheric speckles. This application necessitates the development of fast, low-noise, near-infrared focal plane detectors, and DARKNESS is a critical testbed for one such technology: Microwave Kinetic Inductance Detectors (MKIDs). With the unique capabilities of MKIDs, DARKNESS (and its successors) can simultaneously serve as the FPWFS and low-resolution IFS for science data.

In Chapter 2 we will properly introduce MKIDs, focusing on general aspects related to the design and sensitivity tuning of large format optical/near-IR MKID arrays and recent

advancements. In Chapter 3 we will present the detailed design of DARKNESS, including performance simulations, specifics of the MKID array design, mechanical and optical design, and readout. In Chapter 4 we will present DARKNESS's in-lab characterization, including performance of the MKID array and instrument total throughput, as well as first-light results and commissioning data verifying DARKNESS's on-sky performance with the SDC. In Chapter 5 we will discuss early science results, particularly an analysis of speckle decorrelation times on ms timescales with implications for performing SSD. In Chapter 6 we will conclude with near-future planned work and further prospects for DARKNESS.

Chapter 2

Introduction Part II: Microwave Kinetic Inductance Detectors

Microwave Kinetic Inductance Detectors (MKIDs; Day et al., 2003) are a superconducting detector technology that feature inherently simple geometric design and employ microwave multiplexing techniques to enable relatively low cost kilopixel, and potentially megapixel, arrays. Originally conceived for sub-millimeter astronomy applications, MKIDs have recently been developed for UV, optical, and near-IR (UVOIR) wavelengths (Mazin et al., 2012; Marsden et al., 2012). At these wavelengths MKIDs detect individual photons with time resolution of a few microseconds, are capable of measuring individual photon energies to within a few percent, and have no analogue for the read-noise or dark current present in conventional semiconductor-based detectors. These features make them very promising detector candidates for integration into a photon-counting

IFS that can tackle the speckle suppression challenges outlined in Chapter 1. For comprehensive discussions of the electrodynamics of superconducting micro-resonators the interested reader is directed toward (Mazin, 2004; Gao, 2008; Zmuidzinas, 2011). In this chapter we will delve a bit deeper into the operating principle behind MKIDs, but particularly focus on scaling relations, design considerations, and recent optical/near-IR MKID development pertinent to the arrays used in DARKNESS.

2.1 Operating Principle

We begin with the necessary basics of superconductivity leading in to key principles of MKID design. When cooled below a critical temperature, T_c , a superconductor’s DC resistance drops to zero. The “supercurrent” is composed of pairs of bound electrons known as Cooper pairs, held together with a binding energy

$$2\Delta \approx 3.5k_B T_c \tag{2.1}$$

that are able to travel indefinitely without scattering. Here Δ is the gap energy of the superconductor, which is close to the zero-temperature value $\Delta_0 = 1.76k_B T_c$ for temperatures well below T_c , and k_B is the Boltzmann constant. However, superconductors have non-zero AC impedance owing to the inertia of the Cooper pairs, which introduces a time lag in the conductivity. The AC conductivity is expressed by the Drude model as

$$\sigma(\omega) = \frac{\sigma_{DC}}{1 + i\omega\tau} \tag{2.2}$$

where τ is the scattering time for an electron (or Cooper pair) in the metal, σ_{DC} is the DC conductivity, $i = \sqrt{-1}$, and ω is angular frequency. For normal metals τ is incredibly small such that $\sigma(\omega) \approx \sigma_{DC}$. For a superconductor, both σ_{DC} and $\omega\tau \rightarrow \infty$, and their ratio is finite. This time lag is effectively a stored energy (an inductance) in the surface impedance of the superconductor:

$$Z_s = R_s + i\omega L_s \quad (2.3)$$

The surface resistance R_s is non-zero due to thermally excited quasiparticles,¹ but is very small such that $\omega L_s \gg R_s$. In the limit that the superconducting metal film is thin compared to the field penetration depth (which is always the regime we will be dealing with in our MKIDs) the surface inductance can be expressed as

$$L_s = \frac{\hbar}{\pi\Delta_0} \frac{\rho_n}{t} \quad (2.4)$$

where t is the film thickness, ρ_n is the normal state resistivity, and \hbar is the reduced Planck constant. L_s is typically quoted in units of pH/sq, so a strip of superconductor with length l and width w will have total kinetic inductance $L_{ki} \approx L_s(l/w)$ and volume $V = lwt$ (Zmuidzinas, 2011). This will be a fundamental design parameter later for our MKIDs, and is especially handy since it depends entirely on variables that can be measured in the lab.²

Kinetic inductance is the phenomenon exploited by MKIDs to make useful photon detectors — hence the name “kinetic inductance” detectors. The general operating prin-

¹We define quasiparticles as: electrons that are not bound as Cooper pairs

²We note that Gao (2008) provides an alternative expression in the thin film limit: $L_s = \frac{\lambda_L^2 \omega \mu_0}{t}$ where λ_L is the London penetration depth of the superconductor. This expression gives essentially the same value as Equation 2.4 for a given film.

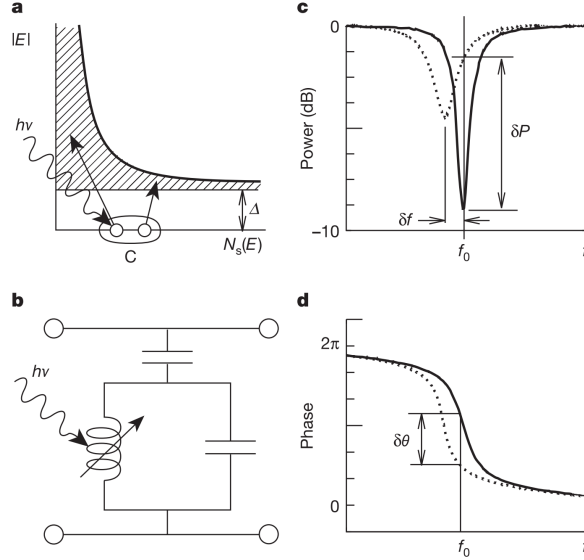


Figure 2.1: MKID operating principle, taken from Day et al. (2003). Incident photons break Cooper pairs in a superconductor. If the superconductor is used as the inductor in a LC oscillator this change in Cooper pair density results in a shift in resonance frequency and amplitude. See text for more details.

ciple is summarized in Figure 2.1. Panel (a) shows the quasiparticle density of states in grey, with Cooper pairs C separated by energy gap Δ . When a photon with energy $h\nu$ strikes a superconductor it will break a number of Cooper pairs given by:

$$dN_{cp} \approx \eta h\nu / 2\Delta \quad (2.5)$$

where $\eta \approx 0.57$ is the efficiency for converting energy to quasiparticles while the rest is lost as phonons (Kozorezov et al., 2000). This change in Cooper pair density briefly changes the surface impedance of the superconductor, which then returns more slowly to its steady state value as the quasiparticle system cools back down. This decay time is referred to later as the quasiparticle lifetime. Fundamentally, we wish to measure the number of broken Cooper pairs. This is achieved by fabricating the superconductor

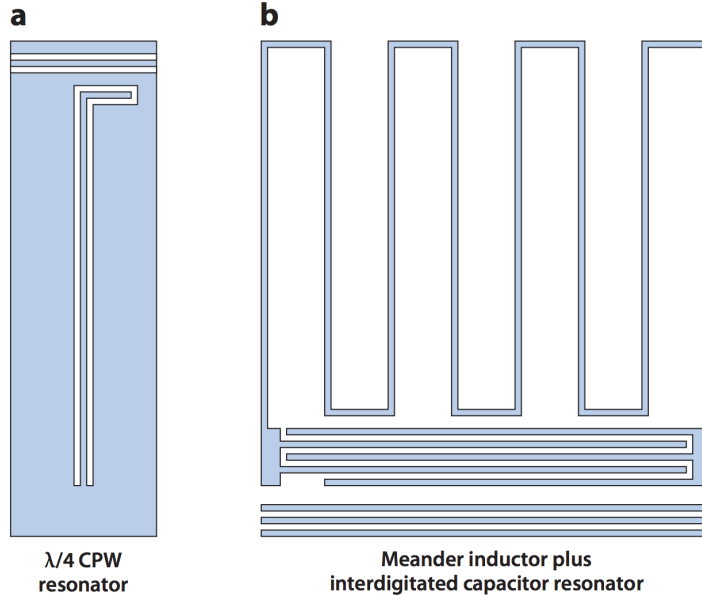


Figure 2.2: Examples of two different resonator geometries mentioned in the text, taken from Zmuidzinas (2011). For both designs, blue represents superconducting film and white is areas of exposed substrate. (a) A $\lambda/4$ coplanar-waveguide (CPW) shunt-coupled transmission line resonator. (b) A lumped-element resonator with meandered inductor and interdigitated capacitor. Both are coupled to a CPW transmission line with similar center strip and gaps, but finite ground strips in the lumped element example.

into a high frequency resonator — typically \sim GHz, where we get the “microwave” in MKID — with frequency $\omega_0 \approx 1/\sqrt{LC}$, with equivalent circuit shown in panel (b). A change in Cooper pair density results in a large change of the kinetic inductance, which in turn causes a measurable shift in the resonator frequency proportional to the number of broken Cooper pairs, shown in panel (c). This frequency shift imparts a change in phase and amplitude of a microwave probe signal transmitted past the resonant circuit, shown in panel (d).

By fabricating individual resonators with different resonant frequencies a single microwave transmission line can read out thousands of resonators simultaneously through




Figure 2.3: Frequency domain multiplexing (FDM) concept. Thousands of resonators can be fabricated with unique frequencies and coupled in parallel to a single microwave transmission line. This figure from Mazin et al. (2013)

frequency domain multiplexing (FDM; McHugh et al., 2012), monitoring each pixel with a unique probe tone. This frequency comb is generated with room temperature electronics, simply transmitted through the MKID transmission line, amplified at low temperature with a single amplifier per feedline, then analyzed again in room temperature electronics. This scheme minimizes the wiring complexity at the device, requiring only two wires to read out thousands of pixels.

Since the gap energy of a superconductor is very small (e.g. $\sim 10^{-4}$ eV for Aluminum), a high energy (~ 1 eV) optical photon will break many thousands of Cooper pairs, creating a “pulse” in the probe signal’s phase as the resonator moves off-resonance rapidly (in $\sim 1 \mu\text{s}$) upon absorbing the photon, then decays back to its steady state more slowly (in




Figure 2.4: A typical phase pulse in a UVOIR MKID. This figure from Mazin et al. (2012)

$\sim 20\text{-}50 \mu\text{s}$) depending on the quasiparticle lifetime (see Figure 2.4). The sharpness of the pulse rise-time, combined with the continuous readout of every resonator simultaneously, provides the high time resolution of MKIDs. Photons of different energies will break different quantities of Cooper pairs, creating pulses in phase with height proportional to photon energy, which provides the intrinsic energy resolution of MKIDs.

2.2 Tuning MKID sensitivity

We define MKID *sensitivity* as the desired probe tone phase shift per number of quasiparticles released by a photon absorption: $d\theta/dN_{qp}$. To inform a discussion of how we can design for a particular sensitivity, we take a quick detour to define more explicitly some physical parameters of our resonators. Transmission past a parallel transmission-

line resonator such as shown in Figure 2.1 is properly referred to as the forward scattering element S_{21} of the two-port circuit scattering matrix, and is actually a complex value. Figure 2.1 (c) shows total transmitted power, or $|S_{21}|^2$. When S_{21} is decomposed into real (I for in-phase) and imaginary (Q for quadrature) parts it traces a circle in the complex plane (referred to later as the IQ loop) for an ideal resonator. Resonator quality factor, $Q = \omega_0/\Delta\omega(FWHM)$, is a measure of how much energy is stored in the resonator divided by energy lost per cycle (high Q implies low rate of energy loss). Total Q is defined as:

$$\frac{1}{Q} = \frac{1}{Q_i} + \frac{1}{Q_c} \quad (2.6)$$

where Q_i is the quality factor due to internal or radiative losses in the resonator, and Q_c is due to energy lost in coupling to the transmission line. For typical UVOIR MKIDs we can fabricate resonators with Q_i in excess of 500,000, such that Q is dominated by the Q_c , which we control in fabrication. Mazin (2004) provides analytical expressions for S_{21} and S_{21}^{min} depending only on the resonator parameters that we care about, namely Q_i and Q_c :

$$S_{21} = \frac{S_{21}^{min} + 2iQ\delta x}{1 + 2iQ\delta x} \quad (2.7)$$

$$S_{21}^{min} = \frac{Q_c}{Q_i + Q_c} \quad (2.8)$$

where δx is the fractional frequency shift $(f - f_0)/f_0$ and $f_0 = \omega_0/2\pi$. By fitting an IQ loop with these general expressions we can extract Q_i and Q_c from real resonator data.

We use them to more clearly illustrate the relation of frequency and phase in Figure 2.5, but in design work we generally skip the toy models in favor of full 2.5D E&M simulations with SONNET.³

Now back to the discussion of sensitivity. As given in Equation 2.5, we expect the number of broken Cooper pairs to be proportional to the binding energy divided by the energy of the absorbed photon (dN_{qp} is then just $2dN_{cp}$). When designing MKIDs for desired sensitivity we can tune a handful of parameters in the pixel design, while many constraints are imposed without wiggle room. A reasonable plate scale for the desired astronomical application (typically wants pixels to be smaller) balanced with the achievable lithographic precision in our fabrication and readout frequencies available in our FDM scheme (both want pixels to be larger) loosely imposes an overall pixel pitch on the order of 100s of μm .

It has been shown empirically that resonator quality improves with lower operating temperature, down to $\sim T_c/8$, until it deviates from Mattis-Bardeen theory and lower temperatures no longer yield improvements (Zmuidzinas, 2011). Current cryogenic technology enables flexible compact instrumentation that can comfortably hold a 100 mK base temperature for a night of observing, thus imposing a T_c constraint of ~ 0.8 to 1 K. This T_c constraint trickles down to a constraint on the superconducting material we can use, and also factors into that material's L_s .

The primary knobs we can turn in the pixel design to find the desired sensitivity are then restricted to: film thickness, which is a loose handle on L_s , but still heavily

³www.sonnetsoftware.com/

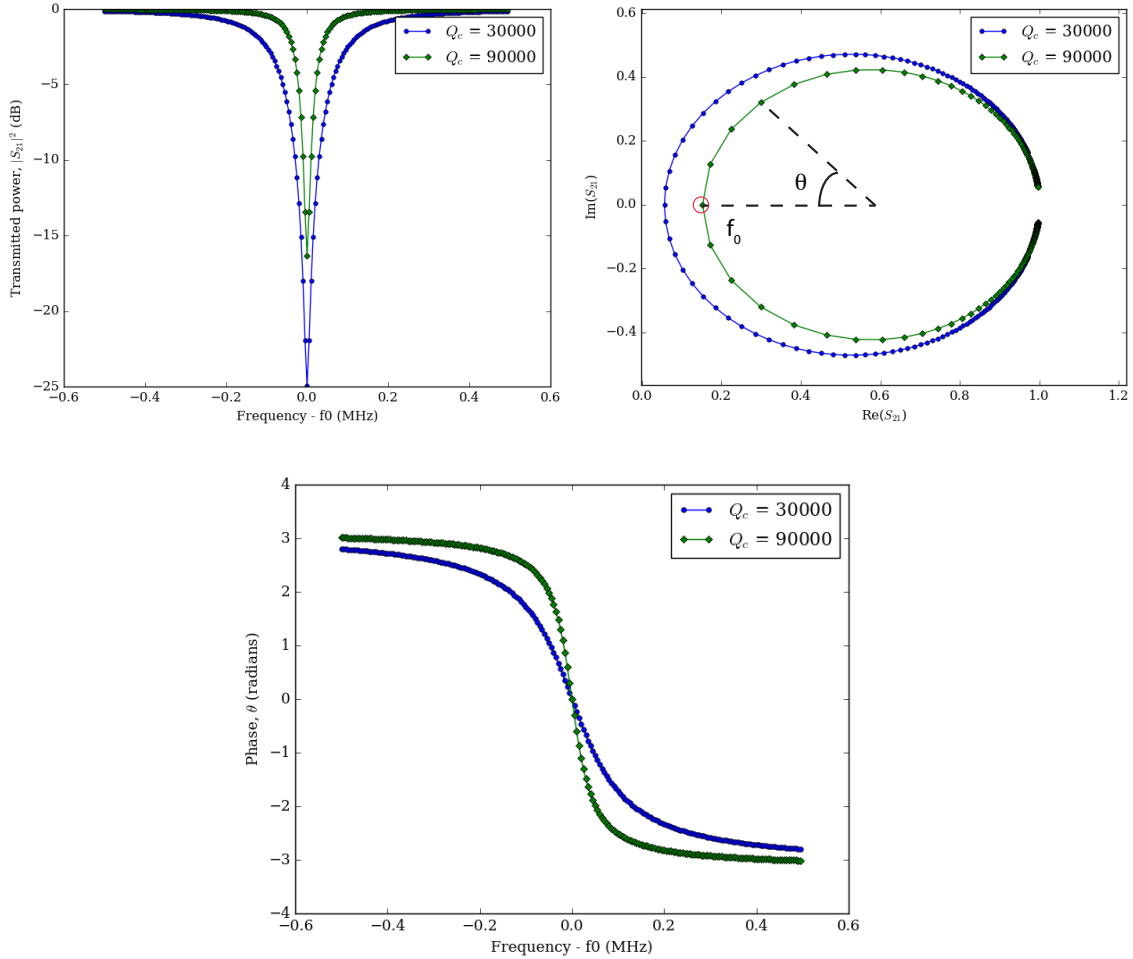


Figure 2.5: Examples of two simulated resonators with $Q_i = 500\text{k}$ and $f_0 = 5$ GHz, but different Q_c . The top left panel shows the way we typically observe our resonators, as the dip in transmitted power ($|S_{21}|^2$) through a two-port microwave transmission line. Feedline transmission away from resonance has been normalized to 1 here, and both resonators are plotted at constant frequency interval of 5 kHz. The top right panel shows both resonators decomposed into I and Q values, where the resonance frequency f_0 is circled in red and phase θ is measured as the angle from f_0 relative to the loop center. The lower Q resonator traces a larger loop, such that the same shift in frequency (same number of frequency points) spans a smaller θ . The lower panel shows this explicitly.

constrained by transparency issues if the film is too thin; inductor volume, which controls the total N_{cp} for a given L_s ; and Q_c which determines the phase shift for a given frequency shift. From Figure 2.5 it is clear that a higher Q resonator will see a much larger phase shift for a given change in frequency, and qualitatively you would expect the fractional change dN_{cp}/N_{cp} to have less impact for a larger inductor volume and larger N_{cp} .

Mazin (2004) derives the following (somewhat messy) expression for estimating resonator sensitivity that explicitly displays these Q and V dependencies:

$$\frac{d\theta}{dN_{qp}} = 3\sqrt{\frac{2}{\pi}} \cdot \frac{\hbar\omega I_0(\zeta) - (2\Delta + \hbar\omega)I_1(\zeta)}{\sqrt{\Delta k_B T}(2\Delta + k_B T)} \cdot \frac{\alpha Q}{N_0 V} \quad (2.9)$$

where $I_{0,1}$ are modified Bessel functions of the first kind, $\zeta = \hbar\omega/2k_B T$, α is the kinetic inductance fraction L_{ki}/L_{tot} (≈ 1 for our designs assuming the thin film limit and $L_{ki} \gg L_m$, the standard magnetic inductance), and N_0 is the quasiparticle single-spin density of states. For titanium nitride (TiN), our material of choice for most of DARKNESS's early life, a theoretical N_0 value was quoted in Leduc et al. (2010) as $8.7 \times 10^9 \mu\text{m}^{-3} \text{ eV}^{-1}$, however, Gao et al. (2012) present a measured value roughly 4x higher, which we will use here: $N_0 = 3.9 \times 10^{10} \mu\text{m}^{-3} \text{ eV}^{-1}$. A couple other caveats: the derivation in Mazin (2004) assumes a thick film limit, whereas we are in the thin film limit which simply adds the factor of 3 out front (Zmuidzinas, 2011); assumptions are made that the fractional frequency deviation $(f_0 - f)/f_0$ is very small, $Q \approx Q_c$, $T \ll T_c$, and dR_s is very small relative to dL_s .

Applying Equation 2.9 to the resonator design in Section 3.1.1 we find a 1 eV photon should create a $\sim 120^\circ$ phase pulse, a decent approximation to what we measure in reality.

While this is a useful tool for getting us in the ballpark with a design, honing in on the desired sensitivity ultimately requires quite a bit of fabrication/testing iteration.

2.3 Recent MKID Advancements

Since their invention, MKIDs have undergone significant evolution to not only improve their performance as detectors, but to facilitate the production of large arrays. In particular, the lumped element KID (LEKID) design presented by Doyle et al. (2008) offered a more convenient pixel geometry that simplified fabrication and provided a straightforward path toward CCD-style focal plane arrays. The use of interdigitated capacitors (IDCs) in this lumped element design has the added benefit of reducing two-level system (TLS) noise⁴ that limits individual MKID noise performance (Noroozian et al., 2009). A subsequent geometry improvement was the use of a double-meandered inductor to reduce inter-pixel crosstalk (Noroozian et al., 2012). Previous work focused on mostly aluminum and niobium based resonator designs, but Leduc et al. (2010) identified TiN as an especially attractive material for use in MKIDs due to its tunable T_c , large kinetic inductance fraction, and robust material properties. While a significant fraction of this development effort has been in the service of sub-millimeter MKID cameras such as MUSIC (Schlaerth et al., 2010; Maloney et al., 2010), NIKA (Monfardini et al., 2010), and MAKO (Swenson et al., 2012), the select improvements summarized here were the foundation of the pixel

⁴TLS are tunneling states in amorphous solids, and appear at the substrate-superconductor interface in MKIDs even when depositing the superconducting film directly on a clean substrate. An in-depth investigation of this noise source can be found in Gao (2008).

architecture utilized in DARKNESS (shown in detail in Section 3.1.1), as well as the first UVOIR MKID demonstrator, ARCONS.

2.4 ARCONS

The full capabilities of UVOIR MKIDs were first demonstrated on-sky with the ARray Camera for Optical to Near-infrared Spectrophotometry (ARCONS; Mazin et al., 2013), a seeing-limited IFS designed for the Coud  focus at Palomar and Lick Observatories, featuring a 2024 pixel TiN MKID array (Mazin et al., 2012; Marsden et al., 2012) optimized for a $0.4 \mu\text{m}$ to $1.1 \mu\text{m}$ bandwidth. ARCONS was the first MKID camera at any wavelength to produce published astronomical science results (Strader et al., 2013; Szypryt et al., 2014; Strader et al., 2016). DARKNESS inherits significantly from ARCONS, especially in the MKID pixel and array design (Section 3.1), readout electronics (Section 3.4), and analysis software (Section 4.3).

Chapter 3

DARKNESS Design

DARKNESS is the first of several planned IFSs (see also Chapter 6) built to demonstrate the potential of MKIDs for high-contrast astronomy. It was originally designed for operation at Palomar Observatory with the PALM-3000 (P3K; Dekany et al., 2013) XAO system and either P1640 or the Stellar Double Coronagraph (SDC; Bottom et al., 2016b), which had not been commissioned as of DARKNESS’s design phase. With the recent completion of P1640’s science survey, and its impending transformation into a precision radial velocity instrument, DARKNESS’s baseline design is now for integration with only the SDC. Here we provide an overview of the instrument as it currently operates at Palomar Observatory with the SDC, while highlighting the flexible aspects of its design that enable its future travel to other observatories (see Chapter 6). A summary of key instrument parameters is provided in Table 3.1.




Figure 3.1: DARKNESS (on the right) and the SDC (on the left) attached to the P3K bench and installed in the Hale Telescope Cassegrain cage. DARKNESS is shown here with its foreoptics box attached, but before connecting external cabling. The SDC electronics board typically hangs next to DARKNESS as well, but is removed here for clarity.

Table 3.1: DARKNESS Overview

Parameters	Values
D-1,2 MKID Array Materials	TiN on Silicon Substrate w/Nb ground plane
D-3 MKID Array Materials	PtSi on Sapphire w/Nb ground plane
Array Format	80×125 pixels
Pixel Pitch	150 μm
Plate Scale	22 mas/pix
Wavelength Coverage	0.8 to 1.4 μm
Spectral Resolution ($\lambda/\Delta\lambda$)	7 to 5
Operating Temperature	100 mK
Cryostat 100 mK Hold Time	13 hours
Cryostat 4 K & 77 K Hold Times	40 hours
Cryostat Dimensions (L×W×H inches)	26 × 13 × 23.5
Cryostat Weight	110 kg

3.1 MKID Array

The 10,000-pixel MKID array for DARKNESS is an evolution of the 2,024-pixel design (known as SCI-4) implemented in ARCONS. In the first year of operation DARKNESS has undergone two detector revisions. Here we present the design for the first generation DARKNESS array, D-1. Subsequent designs, D-2 and D-3, add incremental improvements to tweak sensitivity, improve fabrication yield, and transition to a new superconductor material (PtSi was selected over TiN as our MKID material moving forward because we can fabricate much more uniform films with little effort (Szypryt et al., 2015, 2016)), but the design philosophy and approximate dimensions from D-1 hold true. A manuscript dedicated to large format PtSi MKID arrays is currently in preparation to expand upon the design improvements since D-1 (Szypryt et al. (2017) in prep.).

3.1.1 Pixel Design

The D-1 base pixel design inherits many traits from our lessons learned with SCI-4, as evidenced by the qualitatively similar resonator geometry, while making significant changes to optimize for a new wavelength regime and to ease the transition to a much larger array. Figure 3.2 shows a schematic of the base D-1 pixel design with key design parameters and material properties listed in Table 3.2. Resonators are capacitively coupled to the feedline with a coupling bar that runs along the bottom of the pixel's capacitor and serves as an extension of the coplanar waveguide (CPW) feedline's center strip. The length of this coupling bar controls the coupling quality factor, Q_c , of the resonator. For the resonators in D-1, Q_i is typically $> 10^5$ so Q_c dominates. D-1 pixels are designed for $Q \approx Q_c=30,000$.

The overall MKID pixel size is determined by several competing factors including readout frequency, plate scale, and sensitivity, as described in Section 2.2. The fractional change in Cooper pair density within the inductor, $\delta N_{CP}/N_{CP}$, determines the MKID's frequency shift, which, when combined with the width of the resonance set by Q , determines the phase shift measured by the probe signal. While Q has been included as a free parameter in our design discussion, in practice it is also tied to the density of resonators packed into our readout bandwidth. To avoid crosstalk, resonators need to be placed several linewidths apart in frequency space. We compromise with ~ 2 MHz resonator spacing and $Q_c \approx 30,000$.

Now, for a given Q , the ratio $\delta N_{CP}/N_{CP}$ must be optimized to give the desired

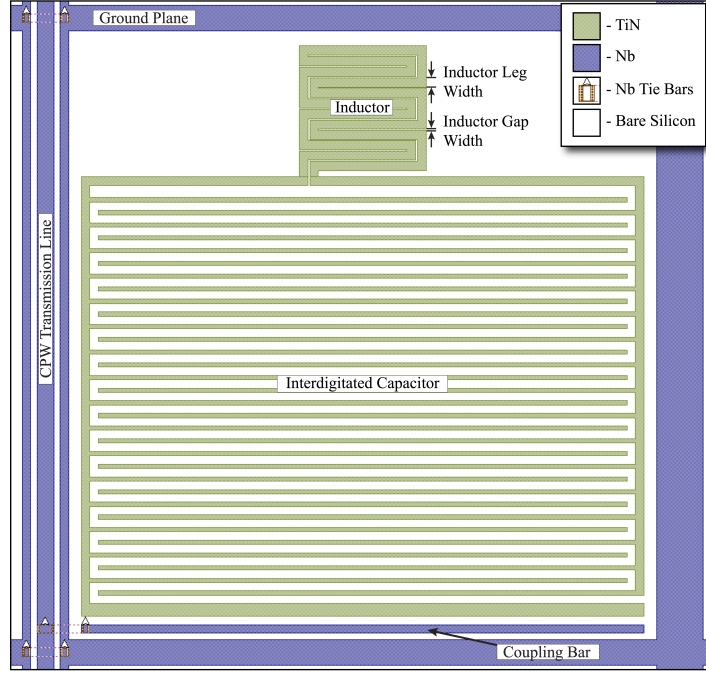


Figure 3.2: Schematic of the D-1 base MKID design with key parameters listed in Table 3.2.

Table 3.2: D-1 Design Parameters

Resonator Material	Titanium Nitride (TiN)
Superconducting Transition Temperature (T_c)	1 K
Resistivity (ρ_n)	$110 \mu\Omega \text{ cm}$
Film Thickness (t)	60 nm
Surface Inductance (L_s) [†]	24 pH/sq
Inductor Width \times Height	$30 \times 30 \mu\text{m}$
Average Inductor Leg Width	$2.5 \mu\text{m}$
Inductor Gap Width	$0.3 \mu\text{m}$
Interdigitated Capacitor (IDC) Width \times Height	$133 \times 104 \mu\text{m}$
IDC Leg Width	$1 \mu\text{m}$
IDC Gap Width	$2 \mu\text{m}$
Total Kinetic Inductance (L_k)	5018 pH
Total Magnetic Inductance	182 pH
Total Capacitance	0.25 pF
Resonant frequency	4.41 GHz
Quality Factor ($Q = f/\Delta f_{FWHM}$)	30,000

[†] L_s is calculated using Eqn. 2.4.

maximum phase shift from the highest energy photon the instrument will detect. The highest energy photon (and therefore δN_{CP}) is set by astronomical application, so the MKID sensitivity can only be tuned by adjusting the inductor volume to control N_{CP} . As demonstrated earlier, Equation 2.9 gives us a decent approximation of the necessary volume after setting all other design parameters. However, in practice, we have a well characterized starting point with SCI-4 that used the same t , Q , T_c , and L_s as D-1, meaning we must only scale V accordingly. SCI-4 featured a $40 \times 40 \mu\text{m}$ inductor with 60 nm thick TiN film, which gave the desired maximum phase-shift for $0.4 \mu\text{m}$ photons. To re-calibrate this sensitivity to $0.8 \mu\text{m}$ photons for DARKNESS, while keeping the same material and film thickness, the inductor area was shrunk by nearly a factor of two to $30 \times 30 \mu\text{m}$. Note that adjusting the volume by using a thinner film is not an option as TiN absorption drops off quickly for films thinner than $\sim 50 \text{ nm}$ (Roquiny et al., 1999).

With inductor geometry firmly set by this sensitivity argument, the capacitor geometry is then tailored to achieve the desired resonant frequency ($\sim 4.4 \text{ GHz}$ for the D-1 base design) within a small enough area to allow for 10,000 pixels to fit on a reasonably sized chip. The exact D-1 pixel size was chosen to achieve a $150 \mu\text{m}$ pixel pitch to match a lenslet pitch that was readily available in commercial microlens arrays (see Section 3.3). This resizing had the added benefit of allowing a straightforward magnification in the camera optics using off-the-shelf components to maintain the same plate-scale as P1640's $75 \mu\text{m}$ lenslet pitch, though this was not a driving factor in the design.

To verify the sensitivity of the D-1 pixel design we briefly present data here from

an engineering grade D-1 chip as tested in the lab.¹ We illuminate the array with two monochromatic light sources: a HgAr lamp with 694 nm filter and a 982 nm laser. Pulses from these sources are shown in Figure 3.3 (Left) for a $Q_m=54,000$ pixel.² The mean pulse height from 694 nm photons is $\sim 100^\circ$, however, it is worth noting that this is a higher Q_m pixel, so the phase offset for a given photon energy is greater than that expected for pixels with median Q_m . For pixels near the median Q_m we measure 694 nm pulse heights of $\sim 80^\circ$, matching the response of typical ARCONS MKIDs to 254 nm photons (Mazin et al., 2012) and demonstrating that our sensitivity is tuned roughly as desired.

3.1.2 Array Layout

Instead of being one monolithic array, D-1 is actually five identical 2,000-pixel (80×25) segments with one feedline per segment. During fabrication (Section 3.1.3) each lithography step is performed by repeating the 2,000-pixel mask one next to the other with a stepper to pattern a full D-1 chip. This strategy reduces the need for multiple photolithography masks for the various layers of the fabrication process (all the necessary D-1 layers fit on a single mask), and could be expanded to scale arbitrarily in one dimension until we fill a 100 or 150 mm wafer.

The resonators in the 2,000-pixel mask are designed to cover a ~ 4 GHz bandwidth with 1.8 MHz spacing between resonators and a 200 MHz gap between low and high

¹In Chapter 4 we will present a full characterization of DARKNESS's state-of-the-art D-3 array.

²When referring to total *measured* Q we will use Q_m .




Figure 3.3: (Left) A phase time stream from a $Q_m=54,000$ D-1 pixel while being illuminated with 694 nm and 982 nm photons. Pulse height is defined as the maximum phase excursion, in degrees, measured by the microwave probe signal during a photon event, with more energetic photons resulting in greater pulse heights. (Right) Histogram of thousands of pulses from 694 nm and 982 nm light. By measuring the FWHM of the pulse height distribution in response to monochromatic light we determine the energy resolution ($R=E/\Delta E$) of the MKID at that wavelength. For a typical engineering sample D-1 pixel we find $R\approx 5$ at 694 nm and $R\approx 3$ at 982 nm.

frequency halves. Frequency tuning is accomplished by drawing back pairs of interdigitated capacitor legs. At higher resonant frequencies the coupling bar must be shortened to maintain $Q_c = 30,000$ for all pixels. Otherwise, high frequency resonators will couple more effectively to the microwave transmission line and Q_c will decrease systematically with increasing frequency. These altered capacitors and correspondingly shorter coupling bars are evident in Figure 3.6. Rigorously designing and simulating 2,000 individual pixels would not be feasible. Instead we simulate several (10-15) resonators in SONNET at frequencies across our bandwidth to build empirical scalings for capacitor and coupling bar shrinkage vs. frequency, then use these scaling relations to programmatically

generate the 2,000 pixels. During this simulation phase we also adjust the tapering of the inductor leg widths as a function of frequency to maintain a relatively uniform current density distribution in the inductor across all frequencies (more details about this tapering process can be found in Mazin et al. (2012)).

During array layout pixels are assigned to positions on the chip according to an algorithm that tries to randomize their position, but simultaneously maximizes the physical distance between pixels with similar resonant frequencies to avoid cross-talk (nearest neighbors have $\Delta f > 50$ MHz). This algorithm is also sensitive to known non-uniformities in the device fabrication. Sputtered TiN can show large T_c gradients across a wafer (Visers et al., 2013), which causes resonators to shift away from their design frequencies since $f \propto L^{-1/2}$ and $L \propto T_c^{-1}$. A frequency “collision” is defined as when two or more resonators land within 500 kHz of each other, such that we can no longer ensure that each unique readout tone is coupling to only a single pixel. These collisions are the dominant source of unusable pixels in our current arrays. Figure 3.4 (Left) shows the frequency map used for the 1,000 pixels covering 4.4-6.4 GHz. This pattern simply repeats in rows 40-80 for the 6.4-8.4 GHz resonators, with a 200 MHz gap between the halves. Figure 3.4 (Right) shows the percent of unusable pixels due to collisions expected from this mapping for various orientations and intensities of frequency gradient across the chip. The frequency map is optimized such that we always lose a predictable number of pixels due to collisions ($\sim 12 \pm 1\%$) regardless of randomly oriented T_c gradients.

Increasing feedline yield is a high priority in D-1. With SCI-4 we encountered several




Figure 3.4: (Left) Frequency map of 1,000 pixels covering 4.4-6.4 GHz demonstrating the randomized placement of pixels with constraints to maximize physical distance between resonators with similar frequencies (nearest neighbors have $\Delta f > 50$ MHz) while also making the array insensitive to uncontrollable gradients in T_c . (Right) The percent of unusable pixels due to frequency collisions expected from the map on (Left) for various orientations and intensities of frequency gradient across the chip. We expect to lose $12 \pm 1\%$ of pixels to collisions.

fabrication defects that reduced our yield, most commonly caused by metallization steps that were not adequately etched, leaving enough metal behind to short the feedline center strip to the ground plane upon its deposition. Even a single small defect or scratch can render an entire feedline useless, and moving from a 2,024 pixel/2 feedline array to a 10,000 pixel/5 feedline array will exacerbate this problem. In anticipation of this issue we re-oriented the pixels in D-1 and reflected them over either side of the feedline, effectively cutting each feedline length in half and reducing the odds of a defect breaking it. Figure 3.6 shows a side-by-side comparison of individual SCI-4 and D-1 pixels demonstrating this re-orientation. This symmetry breaking required minor changes to our coupling




Figure 3.5: A 30 MHz frequency sweep of a D-1 feedline taken with a vector network analyzer showing several resonators, including an instance of a “collision” where two or more resonators are closer than 500 kHz in frequency, rendering all but one of them unusable. The two resonances on the far left narrowly avoid our 500 kHz spacing criterion.

geometry and layout software, but will prove to be a necessary step as we move to larger arrays.

3.1.3 Fabrication

D-1 is fabricated in the Microdevices Lab at NASA’s Jet Propulsion Laboratory (JPL). The fabrication procedure for D-1 closely follows that described in Mazin et al. (2012) for achieving high quality TiN resonators on a high resistivity Si substrate with a niobium (Nb) ground plane and CPW feedline, so we will only summarize recent improvements to the process here.

After depositing the TiN we etch it into a “resonator outline” pattern, leaving square




Figure 3.6: A comparison of individual pixels from the edges of the ARCONS SCI-4 array design and the DARKNESS D-1 array design. Both images are set to the same scale. The 90° reorientation allows D-1 pixels to be reflected over the feedline while keeping the inductors on the correct gridding for the microlens array (Section 3.3). The result is that the feedline winds between every two columns of pixels for DARKNESS, rather than every column as for ARCONS, effectively halving the total feedline length. This reorientation comes with an accompanying re-positioning of the coupling bars. Also evident in this comparison is the shrinking inductor design required to tune DARKNESS’s sensitivity to lower energy photons as discussed in Section 3.1.1. Color differences between the two microscope images are not significant as both arrays are fabricated from the same materials.

TiN patches in the locations where resonators will be, but not yet etching the resonator slots. Fabrication then proceeds with the Nb transmission line, followed by SiO_2 cross-overs, and ground plane. An additional step is added to place gold bond pads around the border of the ground plane for improved heat sinking when wire-bonding the chip into its box. The final step is then to etch the resonator patterns into the TiN patches in a “cookie cutter” fashion. By breaking the resonator patterning into two steps this way, the TiN is still deposited on clean Si (a necessary precaution to achieve high quality resonators), and the fine resonator features are not created until the end of the process,

ensuring they cannot be damaged or dirtied by the lift-off steps required for the Nb and gold.

D-3 Fabrication

The most recent generation of DARKNESS arrays, D-3, utilize PtSi as the superconducting material. The general fabrication process is still very similar to the TiN process outlined above, but clearly some changes were required for the process to play well with the new material, especially in the primary deposition (PtSi with $T_c=1$ K is made in a very different process than TiN) and various protection layers. Szypryt et al. (2016) and Szypryt et al. (2017) in prep. give complete accounts of how we make high quality PtSi resonators and maintain that quality when expanding to full D-3 arrays.

3.2 Cryostat

DARKNESS's cryostat is a liquid cryogen pre-cooled Adiabatic Demagnetization Refrigerator (ADR) capable of reaching temperatures below 100 mK. The custom dewar, built by Precision Cryogenics, is manufactured mostly from 6061-T6 aluminum and is designed as a drop-in replacement for the SDC's usual backend imager, PHARO (Hayward et al., 2001). The ADR unit from High Precision Devices was integrated into the dewar at UCSB. The complete cryostat measures roughly 26 inches long \times 13 inches wide \times 23.5 inches tall, filling a similar envelope on the P3K bench as PHARO, but with a few extra inches of height to accommodate extra cryogen volume. The weight (when filled) is roughly




Figure 3.7: Schematic of the DARKNESS cryostat with several pertinent dimensions provided (all in inches). Detailed labeling of the internal components can be found in Figure 3.8.

110 kg.

A liquid cryogen design was selected due to space constraints in the Hale Telescope Cassegrain cage, preventing the use of pulse tube cooling which would require heavy and rigid drag-lines. Internal to the 300 K vacuum shell is an 8 liter LN₂ tank that maintains a layer of radiation shielding at 77 K for ~ 40 hour hold time. Nested inside the 77 K shield is a 24 liter LHe tank and another layer of radiation shielding surrounding the 4 K experimental volume where the ADR unit and detector package reside. The 4 K stage has a similar ~ 40 hour hold time.

The ADR acts as a single-shot magnetic cooler that brings the MKID array down to

100 mK where the temperature is stabilized with a feedback loop to the ADR magnet power supply. We achieve a 100 mK hold time of 13 hours on-sky, more than sufficient for a night of calibrations and observations. Special care has been taken to isolate and shield the MKID array from the ADR’s magnetic field. The detector package and ADR are mounted to the 4 K stage far apart, with the 100 mK ADR cold finger attached to the MKID array via a flexible copper strap, and the detector package is enclosed in an Amumetal magnetic shield, visible in Figure 3.8 (Left).

The detector package is comprised of three stages: a base plate at 4 K that also holds the magnetic shield, an intermediate 1 K ring that stands off from the 4 K base on Carbon fiber supports, and the 100 mK stage suspended from the 1 K ring by Vespel SCP-5050 rods (Figure 3.8 (Right)). This design thermally isolates the 100 mK stage, places the MKID array far from the magnetic shield opening where field leakage will be strongest, and also accommodates a 1 K baffle to block scattered light and 4 K blackbody radiation that could increase the phase-noise floor of the detector.

3.2.1 Mounting to the AO bench

DARKNESS is held in a custom mounting cradle by three pins near the top of the vacuum shell: two on the sides near the front and one on the rear face. This cradle closely follows the design of the P1640 IFS mounting bracket (Hinkley et al., 2011), including ± 10 degrees of pitch adjustment using a screw-jack at the rear of the instrument and ~ 1 inch of focus adjustment by hanging the instrument from three Bosch-Rexroth ball




Figure 3.8: (Left) Photograph of DARKNESS cryostat with 300 K vacuum shell and 77 K and 4 K radiation shields removed to show 4 K experiment volume with ADR and device mounting stage. The device stage has the top half of its magnetic shield removed for clarity. (Right) Close-up of the device mounting structure with magnetic shield top and 1 K baffle removed.

rails. These runners are then attached to an Aluminum 7071 mounting plate that screws directly to the P3K bench.

3.2.2 Wiring

Microwave Signal Path: DARKNESS requires five feedlines to read out an entire array. The signal paths begin with hermetic SMA bulkhead connectors, bringing the signals in through the bottom face of the cryostat. Five laser welded stainless steel 0.087 inch semi-rigid coax cables bring the signal from the inside of the 300 K shell to the 4 K stage, with a heat sinking clamp at 77 K along the way. From the feedthroughs at the bottom of the 4 K stage the signals pass through 20 dB attenuators for reducing

room temperature Johnson noise, then hand-formable SMA-to-SMA cables bring the signals to the MKID mounting structure. From here, SMA-to-G3PO cables connect to a superconducting 53% Niobium/47% Titanium (NbTi) flex cable.

We have fabricated custom 0.096 mm thick microstrip NbTi/polyimide/NbTi flex cables to allow for a high density of feedlines while minimizing heat load from 4 K to 100 mK as compared to five individual NbTi coax. An example of this flex cable is visible in Figure 3.8 (Right), connecting from the 4 K base of the device mounting structure to the 100 mK stage with an intermediate heat sink at the 1 K stage. A full report of their design and performance is in preparation (Walter et al. (2017) submitted).

These flex cables connect to the MKID box through five small G3PO connectors allowing for a much more compact detector package than standard SMA connectors. The box-mounted G3PO connectors are solder connected to gold-plated copper on duroid co-planar waveguide (CPW) transition boards, which are then aluminum wire-bonded to the MKID chip. After passing through the MKID array the five signals are brought out through the same series of CPW board, G3PO connectors, custom NbTi microstrip flexcable, and G3PO-to-SMA wires. At 4 K each line is amplified by a Low Noise Factory High Electron Mobility Transistor (HEMT) amplifier. Hand-formable SMA-to-SMA cables again bring the signal to the bottom of the 4 K plate, and stainless steel coax take the signals from there to 300 K.

DC wiring: DARKNESS has two 24-pin DC wire bundles going to 4 K to provide HEMT biasing and thermometry. Current is supplied to the ADR magnet by two DC

leads using copper wire from 300 K to 77 K, high- T_c superconductor from 77 K to 4 K, and superconducting (NbTi) wire from 4 K to the magnet.

3.3 Optical Design

DARKNESS's cryostat optics are very simple, requiring only an entrance window, a pair of cold, IR-blocking filters at the 77 K and 4 K stages, and a microlens array that concentrates the light from the final image plane onto the photo-sensitive inductor of each pixel.

The entrance window is 12.5 mm thick AR-coated fused Silica. The 77 K and 4 K windows are both BK7 glass, 10 mm and 20 mm thick, respectively. These windows are coated with a custom IR-blocking filter from Custom Scientific, and define the cutoffs of our observing band. The transmission curve of a single filter is shown in Figure 3.9. Since these are reflective coatings, both windows are mounted at a 3° angle relative to the incident beam to reduce ghosting. The microlens array (MLA) from Advanced Microoptic Systems is composed of roughly 100 x 145 lenslets with 150 μm pitch, and made from 1 mm thick STIH53 glass.

The majority of DARKNESS's optical complexity is in the warm re-imaging optics that convert the SDC's $f/15.7$ output beam to the $f/322$ required for DARKNESS. The main constraints to this optical design are the diffraction limit of the telescope at DARKNESS's operating wavelengths and the need for telecentricity to ensure proper functioning of DARKNESS's microlens array. We performed Zemax simulations of the




Figure 3.9: Transmission curve of one IR-blocking/bandpass filter. Blocking is specified as $T_{Avg} \leq 0.4\%$ from 1.5 to 2.8 μm . Our filter stack combines two such filters, providing roughly OD4 blocking in the specified range.

full optical train, including the Hale Telescope, P3K, the SDC, and DARKNESS to verify diffraction limited performance and Nyquist sampling of the diffraction beamwidth (λ/D) across our observing band. Resulting spot diagrams are shown in Figure 3.11.

The layout of the fore-optics enclosure, including a ray trace of the re-imaging optics and a finder camera/pupil imaging arm, is shown in Figure 3.10. The $f/15.7$ beam is first folded by a "field selector" comprised of a reflective aluminum rectangle deposited on the center of a BK7 window. This optic sends the central $3'' \times 4''$ of the FOV to the $f/322$ re-imaging optics and DARKNESS, while passing the surrounding full FOV to an SBIG STF-8300M CCD camera. The SBIG arm can switch between imaging (for target acquisition) and pupil viewing mode by flipping in an optional lens. The science beam

is collimated and then folded again toward DARKNESS's entrance window. This fold mirror is on a remote controlled 3-axis Picomotor stage, allowing for fine adjustment of the FOV on the MKID array and automated dithering routines to fill in dead pixels. After the Picomotor mirror, the collimated beam passes through a 5-position 1" filter wheel with a selection of neutral density filters and is then re-focused to an $f/322$ beam. A second, 7-position 2" filter wheel is placed just before the DARKNESS front window, providing a selection of color filters and also serving as the instrument's "shutter." A summary of the selectable filters can be found in Table 3.3.

This entire fore-optics enclosure is easily removable so DARKNESS can accommodate a variety of observing configurations³ without disrupting the optical configuration used with SDC or making any changes to cold optics. Focus placement after removing/replacing the fore-optics is repeatable to within the day-to-day focus drift experienced over a several day observing run.

Wavelength Calibration

As part of our calibration procedure, we must map phase offset to photon energy for each pixel using the measured response from a series of known laser sources. We employ a similar setup to the one described in Mazin et al. (2013) for ARCONS using a custom package that holds several laser diodes controlled remotely via Arduino. The lasers used for DARKNESS operate at 808 nm, 920 nm, 980 nm, 1120 nm, and 1310 nm. The

³...including: directly accepting the P3K beam, operating in a seeing-limited mode at the Cassegrain focus, or traveling to other observatories

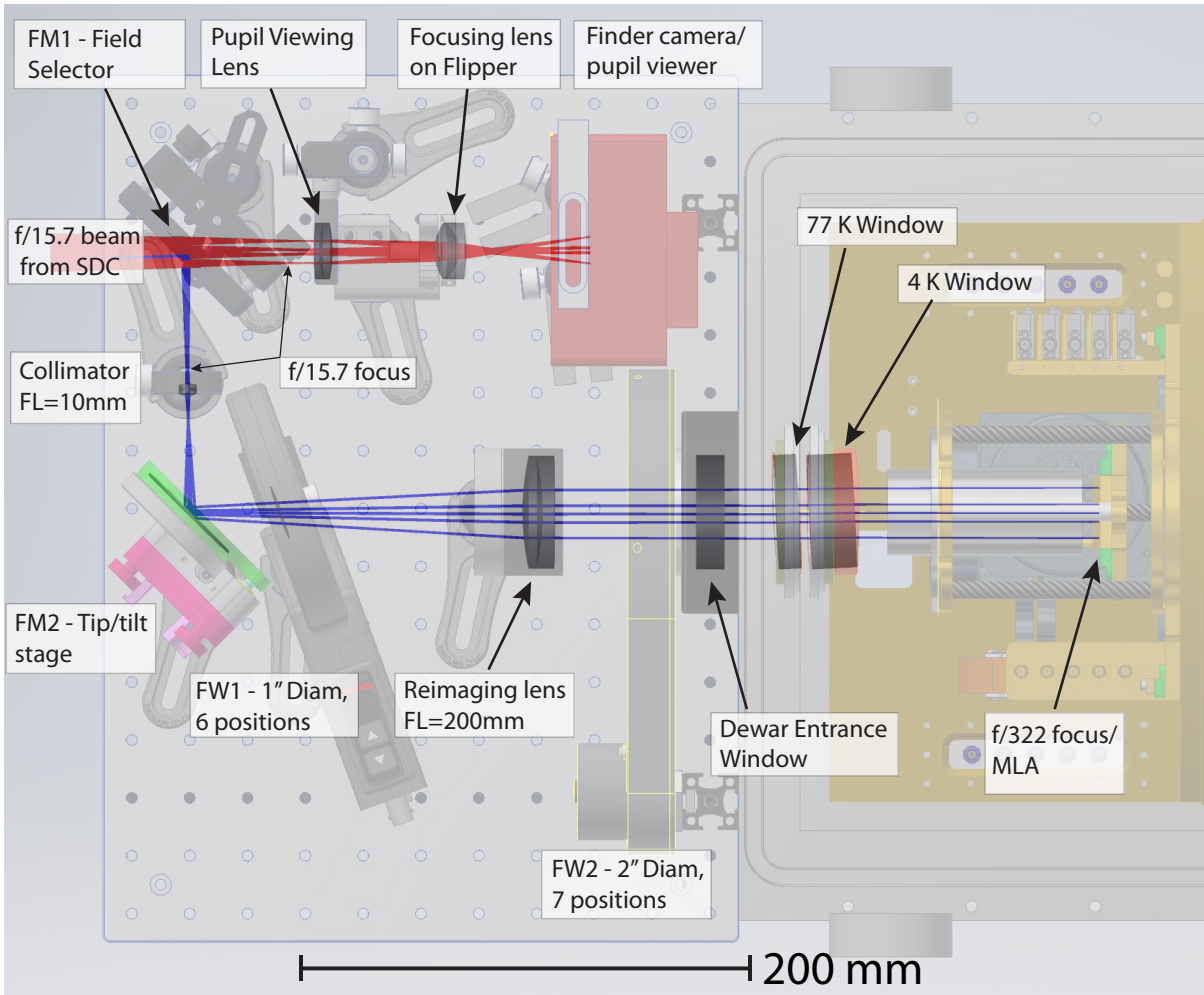


Figure 3.10: DARKNESS foreoptics layout when operating with the SDC.

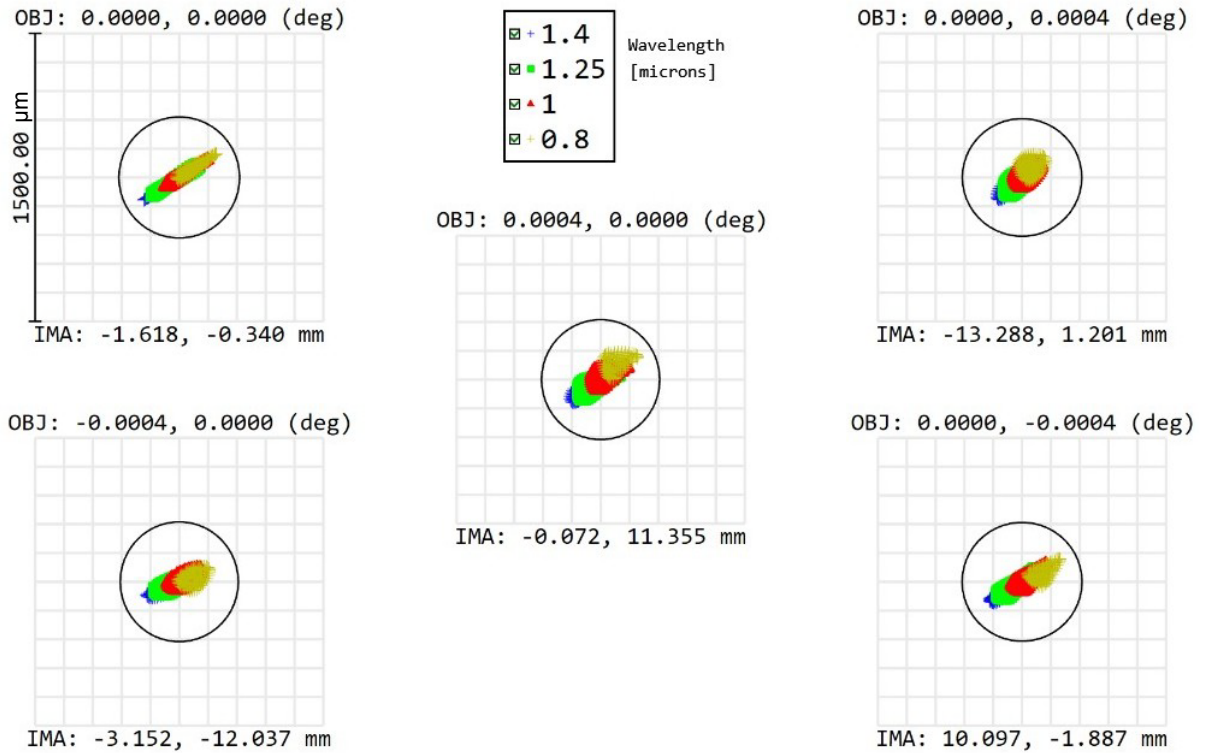


Figure 3.11: Spot diagrams from several fields roughly covering the extent of the DARK-NESS FOV. Airy radius is shown as black circle, calculated at 0.8 μm . These diagrams show that even at the edges of the array the optical performance of the system will be diffraction limited (or rather, governed by the wavefront quality achieved with P3K) and chromatic effects are negligible. With an Airy radius of $\sim 300 \mu\text{m}$ at our shortest wavelength we are Nyquist sampling (or better) the diffraction limited PSF across our band with our 150 μm pixel pitch.

diodes shine into an integrating sphere with fiber output, allowing us to mount the laser box assembly wherever is convenient (typically on the SDC electronics board alongside DARKNESS) while using a fiber to bring the light into our fore-optics box.

In the fore-optics this fiber is installed next to the re-imaging lens, then directed at the DARKNESS front window to simply flood illuminate the detector with the help of a diffuser in FW2. Here uniformity is not a priority as much as decent count rate on every pixel. During observations, typically while tuning AO on a new target, we perform a wavelength calibration by closing FW1, moving FW2 to the diffuser position, then cycling through the lasers taking ~ 1 minute of data from each. In processing, each laser peak is fit by a Gaussian to locate the mean phase offset from that wavelength for each pixel. The five peak locations are subsequently fit with a second-order polynomial to provide a complete mapping of phase-offset back to wavelength, which is applied as the first calibration step in our processing pipeline (this procedure is essentially unchanged from van Eyken et al. (2015)).

3.3.1 Modifications to the SDC

The SDC is a flexible coronagraph platform that features two internal focal planes and two pupil planes for deploying a variety of coronagraphic configurations, including a dual vector vortex coronagraph (VVC) designed to overcome diffraction from the secondary and spider obscurations in the pupil (Mawet et al., 2005, 2011). To optimally utilize a VVC requires superb correction from the XAO system, and very high Strehl ratio. Currently,

P3K does not provide adequate Strehl ratio below J-band to justify the use of a VVC at these wavelengths (though this may change very soon — see discussion in Chapter 6). With this consideration, and to minimize time-consuming coronagraph alignment while debugging DARKNESS, we fabricated and installed a simple Lyot coronagraph for use in SDC. The focal plane mask (FPM) is actually a set of three reflective aluminum spots of various diameters, sputter deposited on a single fused silica substrate, and located at the SDC’s first focal plane. Similar to the standard SDC configuration, this FPM is installed on a linear translation stage along with a fiber ferrule, co-focused with the FPM, allowing fiber injection directly into SDC and DARKNESS which is used to focus DARKNESS relative to the coronagraph focal plane. A Lyot mask, fabricated by deep reactive ion etching through a Silicon substrate, is installed at the SDC’s first pupil plane. This mask can easily be made reflective to use the rejected starlight for low-order wavefront sensing (LOWFS) in future planned upgrades to the SDC. The secondary focal and pupil planes in the SDC are not used. With a selection of FPM diameters and Lyot mask sizes we can choose the desired configuration based on observing conditions, and this flexibility also allows for easy re-configuration when optimizing for different observing wavelengths. The masks used during commissioning, shown in Figure 3.12 with selectable parameters listed in Table 3.3, were optimized for J-band operation since this is where we expected the best correction from P3K and focused the majority of commissioning observations.

The SDC includes an internal IR quad-cell detector near its first focal plane that maintains very precise alignment of the target star on the coronagraph FPM. In the




Figure 3.12: Lyot Coronagraph optics in the SDC. (Left) Focal plane masks (Right) Lyot mask

standard SDC configuration — optimized for K-band observations — a dichroic sends all J-band light to this IR tracker. When operating with DARKNESS we replace this dichroic with a 55/45 pellicle to simply share the J-band light between DARKNESS and the IR tracker. A similar compromise is also made with the P3K dichroic where the typical optic would take part of DARKNESS’s bandwidth at the blue end for the WFS, so we install a 50/50 beamsplitter to avoid this. In the future we could recover some throughput by acquiring an appropriate dichroic for P3K, but a proper solution for the SDC tracker would require a new quad cell and is likely too invasive to be feasible.

Table 3.3: Selectable optics parameters

Parameters	Values
<i>Lyot Coronagraph</i>	
FPM Diameters (λ/D at 1.25 μm)	5, 6.6, 8.2
Undersized Lyot Stop Factor	10%, 15%, 20%
<i>Foreoptics</i>	
FW1 Neutral Density Filters	Closed, OD 3.0, OD 2.5, OD 1.5, OD 0.5, Open
FW2 Diffusers & Color Filters	Closed, Y, z_s , Mauna Kea J, Diffuser (high grit), Diffuser (low grit), Open

3.4 Readout

DARKNESS’s readout hardware and photon detection firmware are an evolution of those used in ARCONS, and follow the same strategy as that outlined in McHugh et al. (2012). In general, these instruments use a heterodyne mixing scheme, where a set of probe tones is generated for each MKID resonator. These tones are then passed through the device where the effects of illumination on the MKID array are imprinted on the probes, and this altered signal is compared against the original to detect the individual photon strikes. In the DARKNESS readout scheme, frequency comb generation/conversion and photon detection are handled with a combination of three boards: a ROACH2 board, a combination ADC/DAC board, and an intermediate frequency (IF) board. ROACH2 is the second generation of the CASPER Reconfigurable Open-Architecture Computing Hardware (ROACH), a platform originally intended for radio astronomy, and selected for our purposes for the flexibility and rapid development it enables.⁴ The additional resources provided by ROACH2 and the advances in ADC/DAC technology since ARCONS devel-

⁴More information on ROACH2 can be found on the CASPER wiki: casper.berkeley.edu/wiki/ROACH2

opment has enabled a substantial increase in readout density. One set of DARKNESS boards is capable of reading out 1024 MKIDs in 2 GHz bandwidth, so each DARKNESS feedline thus requires two sets of boards for resonators covering 4 GHz of bandwidth, for a total of 10 sets for 10,000 pixels. For comparison ARCONS required eight ROACH boards for 2024 pixels. We summarize DARKNESS's board specifications, readout signal chain, and photon detection here, but encourage interested readers to consult Strader (2016) for significantly more detailed descriptions.

Figure 3.13 provides a block diagram of the readout chain. Definition of the tone frequencies and backend signal processing are handled on the Virtex 6 field programmable gate array (FPGA) on the ROACH2 board. The probe tones for each resonator are created as complex waveforms at low frequency, generated as separate real (I or in-phase) and imaginary (Q or quadrature) components using dual 2 GSPS (giga-samples per second) 16-bit DACs. These I and Q components are then combined on the IF board and also mixed with a local oscillator (LO) up to our MKID frequencies (4 to 8.5 GHz). The summed waveforms (representing a "comb" in frequency space) are then sent to DARKNESS where they pass through the MKID, are amplified at 4 K with HEMT amplifiers, then brought back to the readout electronics. After another round of amplification on the IF board, the signals are mixed down to base-band and are broken back into I and Q components, then sent to the ADC/DAC board for digital conversion with dual 2 GSPS 12-bit ADCs. Finally this I and Q data is sent to the ROACH2 for channelization, filtering, and photon detection.

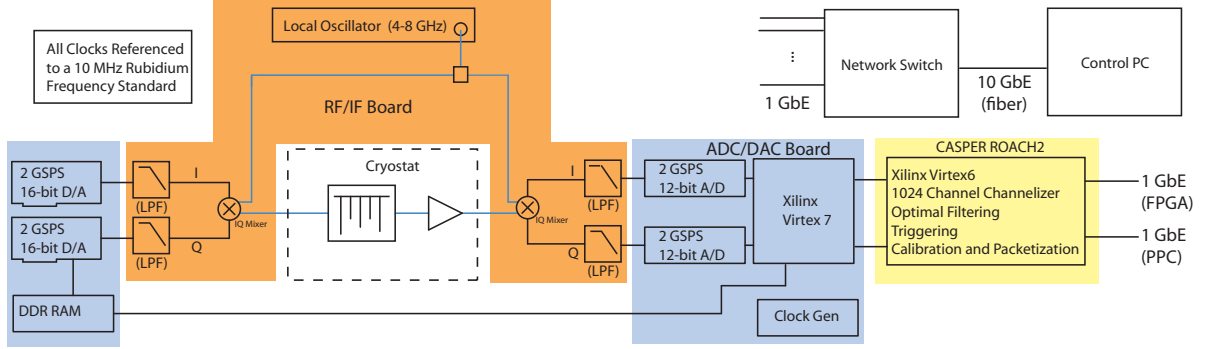


Figure 3.13: Block diagram of the second generation UVOIR MKID readout. The blue, orange, and yellow shaded regions take place on the ADC/DAC, IF, and ROACH2 boards, respectively.

In the ROACH2 firmware, individual tones are separated out and downconverted to 0 Hz with a two stage channelization process. The I and Q data for each channel is low pass filtered, then converted to phase⁵ and filtered with a finite impulse response (FIR) filter with coefficients customized to each channel’s unique pulse shape. This shape is determined using the formalism of Optimal (Weiner) Filtering. After filtering, phase excursions that pass some threshold⁶ are flagged as photon events, and the photon is stored in a buffer as a 64-bit word that includes the arrival timestamp, phase height of the event, and resonator ID for the pixel where the photon was detected. Every 0.5 ms or 100 photons (whichever comes first) the photons in every ROACH buffer are sent over 1 GbE connection to a HP Procurve switch. The switch is connected to the data acquisition (DAQ) computer with a 10 GbE fiber link. The computer collects the photons packets and writes them to a 80 TB RAID6 array continuously. It also computes a quicklook

⁵ $\phi = \arctan\left(\frac{I - I_{center}}{Q - Q_{center}}\right)$ where (I_{center}, Q_{center}) is the center of the resonator loop in the I/Q plane.

⁶Typically 4σ below zero phase, where σ is the standard deviation in optimally filtered phase noise for each pixel.

image which it writes to disk every second.

3.4.1 Electronics Rack

DARKNESS's readout electronics, ADR magnet power supply, HEMT power supply, thermometry control, and GPS reference timesource are installed in an electronics rack that attaches to the outside of the Cassegrain cage during observations (see Figure 3.14 for a view of the rack in the AO lab). The readout electronics are installed as server blades into a crate that supplies power, 1 PPS, and 10 MHz signals to each set of boards (2 board sets per blade). Total power consumption is dominated by this readout crate that contributes 1.4 kW of the 1.7 kW total power budget. Cooling of the rack is achieved with two fan tray/heat exchanger pairs, one dedicated to the crate and one to the rest of the rack components, integrated with the P3K glycol cooling system through a copper manifold in the back of the rack. Total weight is roughly 100 kg.

3.5 Performance Simulations

To estimate the potential contrast gain of DARKNESS across our target wavelength range we conducted simulations of the P1640 coronagraph using the IDL wavefront propagation software, PROPER (Krist, 2007). At the time of this design work we used the P1640 coronagraph as our baseline since its optical properties were well understood, however a version of P1640's APLC masks could easily be installed in the SDC. We assumed no changes to P1640, which was optimized for H and J band observations, and aimed to



Figure 3.14: DARKNESS electronics rack in lab at UCSB. This rack holds 5 readout cartridges with 10 sets of the readout boards described in Section 3.4, 2 per cartridge, installed in a custom readout crate at the bottom of the rack that distributes power, 1 PPS, and 10 MHz signals to the boards. Above the crate is a network switch, time source, thermometry control, HEMT power supply, and ADR magnet power supply.

answer two questions: what contrast will P1640 deliver at $0.8 \mu\text{m}$ where its performance was not previously characterized, but DARKNESS's observing band begins? How much contrast improvement is possible on top of this raw contrast using the Dark Speckle technique and assuming DARKNESS's instrument parameters?

As a sanity check for our prescription of the P1640 optics we first simulated an unaberrated version of the full telescope, AO, and coronagraph optical train with a uniform input beam. We verified the accuracy of this model using Lyot plane images and contrast profiles provided in the P1640 design documentation.⁷ Wavefront errors were then added, taking into consideration an imperfect input beam and aberrated optics. Mirror surface aberrations were approximated with power spectral density (PSD) error maps, using RMS values supplied in the P1640 documentation. A PSD phase screen was also applied to the coronagraph entrance beam with a 105 nm RMS amplitude to represent the residual WFE delivered by P3K under median seeing (Dekany et al., 2013). We extracted a simulated Strehl ratio of 85% at 1650 nm, matching the expectation from Equation 1.1. Figure 3.15 shows the resulting image plane showing point spread functions (PSFs) with realistically simulated speckle patterns and contrast ratios matching on-sky results (Oppenheimer et al., 2012, 2013).

With this framework in place we extend our simulations down through $0.8 \mu\text{m}$, covering the entire DARKNESS bandwidth. Figure 3.16 shows contrast profiles with raw coronagraphic suppression, coronagraph+SDI, and coronagraph+Dark Speckle. SDI performance is approximated as an order of magnitude improvement in contrast, following

⁷P1640 Design Document: www.astro.caltech.edu/~shinkley/PROJECT_1640_files/P1640_DandO_compress.pdf




Figure 3.15: Simulated P1640 I and J-band PSFs with 105 nm RMS post-AO wavefront error and aberrated coronagraph optics. Both images are roughly 4 arcseconds per side.

Crepp et al. (2011)’s results. Dark Speckle performance is estimated using Equation 1.12 for a SNR of 5 from a 3 hour total integration with 20 ms exposures, and should be considered a lower limit where static speckles have been perfectly accounted for. Here G is given by the raw coronagraph contrast, N_{\star} is calculated from known HR 8799 magnitudes with a 6% estimated system throughput, and j is given by the designed DARK-NESS plate scale. We expect that when operating in unison SDI and Dark Speckle will provide higher contrasts than either one is capable of alone since SDI will be very useful for removing static speckles that would otherwise limit the utility of the Dark Speckle technique (Boccaletti et al., 2000).




Figure 3.16: Simulated P1640 I and J-band contrast curves with 105 nm post-AO RMS wavefront error and aberrated coronagraph optics. HR 8799 planets are overplotted using available known photometry in J-band and photometry derived from model spectra in I-band (and J-band for planet e) (Madhusudhan et al., 2011). Spectral differential imaging (SDI) performance is estimated as an order of magnitude increase in contrast ratio, following Crepp et al. (2011)’s analysis. $5\text{-}\sigma$ Dark Speckle (DS) suppression is estimated using Equation 1.12 and should be considered a lower limit where static speckles have been perfectly accounted for. We expect that when operating in unison SDI and Dark Speckle will certainly provide higher contrasts than either one is capable of alone since SDI will be very useful for removing static speckles that would otherwise limit the utility of the Dark Speckle technique (Boccaletti et al., 2000). We find all planets should be detectable in both bands with DARKNESS, allowing us to place constraints on current atmosphere models should they be detected at the expected contrasts.

3.5.1 Constraining HR 8799

To demonstrate the science potential of DARKNESS we have overplotted HR 8799b, c, d, and e on our simulated contrast curves (Figure 3.16). H and J-band photometry for planets b, c, and d (as well as HR 8799 itself) are taken from the original discovery paper (Marois et al., 2008b). H-band photometry for HR 8799e is provided by Skemer et al. (2012). To simulate the J-band photometry for HR 8799e and the I-band photometry for all four planets we apply the atmosphere models of Madhusudhan et al. (2011). Following their suggestions for the best-fitting models (and Skemer et al. (2012)’s suggestion for HR 8799e) we apply simulated I and J-band filters and extract absolute magnitudes in each band. To estimate HR 8799’s I-band magnitude we applied color tables found in Zombeck (2007) scaled to the previously measured colors for HR 8799 (all roughly 2x those expected from Zombeck (2007) for a main sequence A5V star). The simulated J-band contrast for HR 8799d and e places them just outside P1640’s detectable regime, matching results from P1640 at this wavelength (Oppenheimer et al., 2013). With the added suppression provided by DARKNESS, J-band spectral data can be obtained from Palomar for these objects.

The I-band contrast yields intriguing results. None of the planets are accessible with P1640 and standard SDI speckle suppression, even if the current IFS was sensitive to those wavelengths. However, all planets do fall into a regime that is theoretically accessible with Dark Speckle imaging, given the current models’ I-band magnitudes. Thus HR8799 provides an excellent test of DARKNESS’s true capability, with the potential to place

new spectral constraints on giant exoplanet atmosphere models.

Chapter 4

DARKNESS Characterization and Commissioning

DARKNESS was commissioned in July 2016, and subsequently returned to Palomar in November 2016 and April 2017 (hereafter we'll refer to these observing trips as 2016a, 2016b, 2017a). As we've previously alluded to, UVOIR MKID development is an ongoing process. With the technology still maturing at a rapid rate, every DARKNESS observing run provides new opportunities to incorporate the improvements we've made to the detector.¹ As such, the detector used during the 2017a observing trip (a PtSi D-3) is not the same as that used for the first two trips (a TiN D-2), and will likely be replaced in time for the next trip. Rather than comprehensively compare every iteration of the detector, we will outline the battery of validation tests we perform in the lab

¹The readout is maturing rapidly in parallel with the detectors, and also contributes significantly to our ability to field a well-behaved instrument. Improvements on this front are mostly related to generating clean probe tones and are discussed in full in Strader (2016).

for each candidate array, presenting results only from the most recent D-3 device. This can be seen as a snapshot of DARKNESS’s current state, with commentary on areas of expected improvement for D-4. Section 4.2 presents verification of DARKNESS’s optical performance and will include a mix of results from the three completed observing trips.

4.1 In-lab Verifications

4.1.1 MKID Quality and Yield

The first step in characterizing a new array fresh out of the fabrication is to get a sense of the resonator yield and typical Q_m . We use a vector network analyzer (VNA) to sweep each feedline’s transmission as a function of frequency — a process we refer to as the *widesweep* — to locate all resonator dips. The individual resonators are identified (mostly automatically, with some manual intervention) and frequency collisions are flagged and removed. The 2017a D-3 showed roughly 1725 of 2000 (86%) resonators per feedline in raw yield, and 1500 of 2000 (75%) after removing collisions. The widesweep data for each resonator is then fit using Equations 2.7 and 2.8 to extract the average quality factors (Q_m and Q_i). Figure 4.1 shows the results of this widesweep and resonator fitting analysis for a D-3 feedline.

From these plots we see that the resonator placement is excellent — D-3 aims for resonator bandwidth covering 4 to 8.2 GHz with a 200 MHz gap between low and high frequency halves — however, Q_m is lower than desired due to Q_i being quite low. From

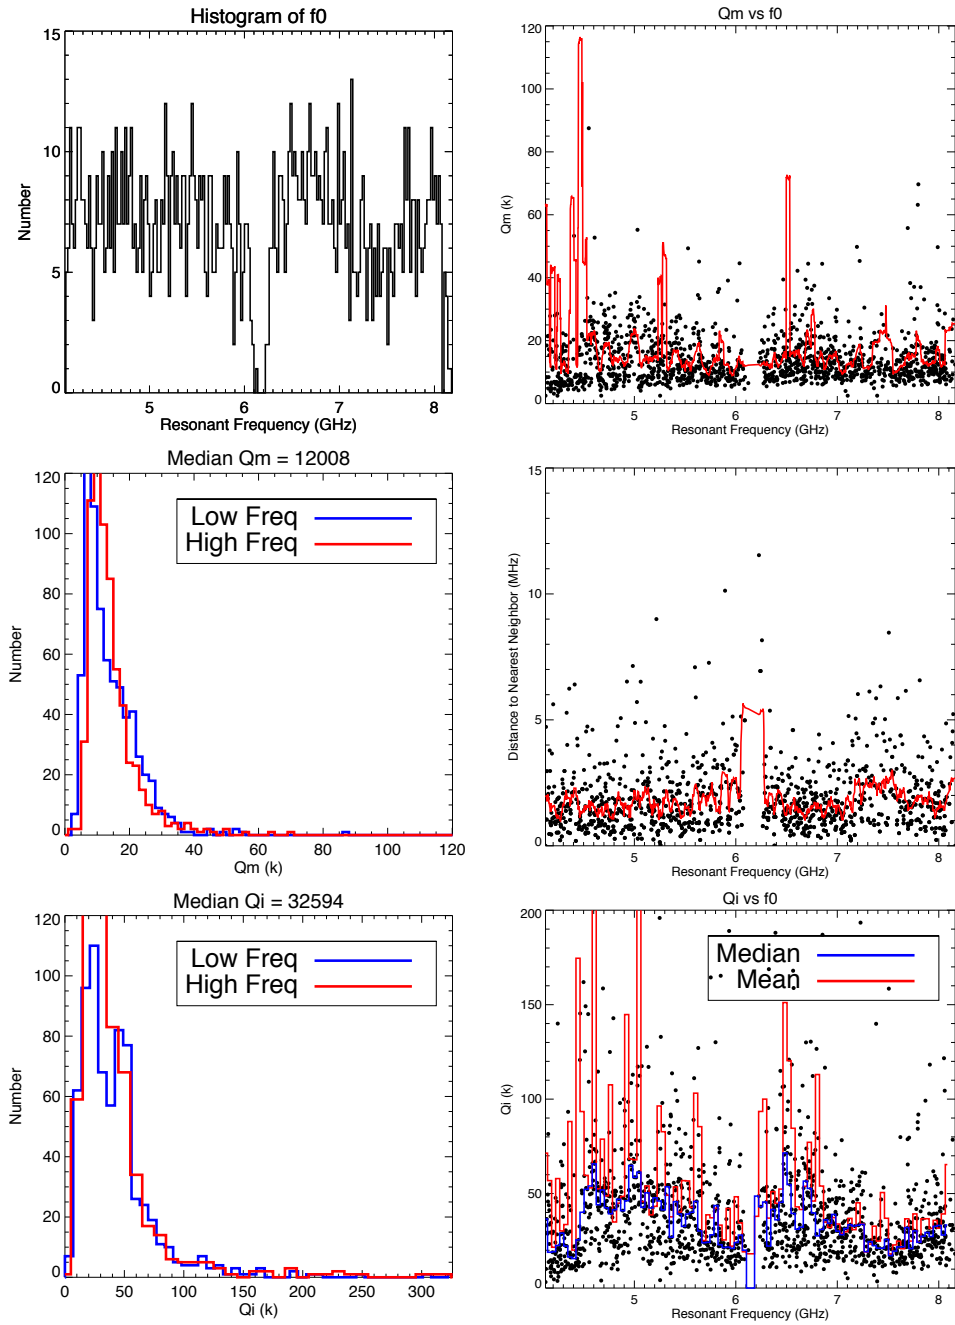


Figure 4.1: Output of resonator fits for a typical D-3 feedline from DARKNESS's 2017a run.

Equation 2.6 we can see that median Q_c is actually close to its design goal of 30,000. From our experience with the first TiN arrays in ARCONS, expanding from single layer test chips to full kilopixel arrays invariably degrades the resonators because of the extra processing required. D-3's low Q_i is likely due to such fabrication issues, and will increase as we explore better techniques for preserving the resonators throughout fabrication. More peculiar is the trend in Q_i seen as a function of frequency. This implies that resonators near some frequencies are losing energy in an additional unexpected manner, such as coupling to a box resonance mode or standing wave on the feedline. We suspect this is being caused by the periodicity of our ground plane cross over placement, and have eliminated this behavior in subsequent devices with randomized cross overs. Despite the lower than usual Q_i , D-3's resonators are quite sensitive and show decent energy resolution.

4.1.2 D-3 Sensitivity and Energy Resolution

To confirm the sensitivity of the detectors, we proceed with the same measurement described in Section 3.1.1 where we illuminate the array with a light source in the lab while reading out a handful of resonators with a small-scale version of our readout. Figure 4.2 shows a couple D-3 pulses from an 808 nm laser. The pulse heights are near the desired value (100°) for the blue end of our bandpass, but considering that Q_m is nearly $1/2$ the designed value this means the resonator volume will need to be tweaked in future iterations to correct the sensitivity once we achieve high Q_i and the expected




Figure 4.2: Typical D-3 808 nm pulses.

Q_m .

To calibrate the response across the entire array we illuminate with a series of narrow band lasers. For a given laser line, we estimate a pixel’s energy resolution R at that wavelength by the FWHM of a Gaussian fit to the pixel’s measured spectrum. Using our full digital readout we determine R as a function of λ for every pixel. Figure 4.3 shows the distribution of measured R s from our 2017a D-3 array at two different wavelengths. From this data we see that the D-3’s median energy resolution across our 0.8 to 1.4 μm band goes from ~ 7 to 5.

4.1.3 DARKNESS Throughput

To measure the absolute throughput of DARKNESS including the 300 K, 77 K, and 4 K windows, MLA, and MKID QE, we’ve constructed a “quantum efficiency” (QE) testbed.




Figure 4.3: Histogram of D-3 energy resolutions.

This testbed uses a monochromator and integrating sphere to generate a uniformly illuminated object plane of given λ . We then re-image this object plane to a $\sim f/300$ beam and use a rotating fold mirror to send the light out to DARKNESS, or to a calibrated photo-diode inside the testbed enclosure. By dividing the flux measured at DARKNESS by the flux measured on the photo-diode we obtain an absolute measurement of the instrument throughput at each MKID pixel, and with the monochromator we can perform this measurement at discrete wavelengths across our bandpass. See Appendix A for more detail on the QE testbed design.

Figure 4.4 shows the median throughput measured across the array in 50 nm steps from 0.8 to 1.4 μm (solid black curve). The shaded region is the $1\text{-}\sigma$ standard deviation in throughputs measured by the individual pixels. This measured curve is compared against the throughput expected from the known transmission of our windows + filters,

PtSi measured absorptivity, and D-3 inductor fill factor. The biggest unknown in this theoretical curve is the spot size at the MLA focus. The dotted curve assumes 100% of the light is focused onto the inductor by the MLA, but from Zemax simulations we know this is not the case — i.e. ensquared energy at the extent of the inductor is not 100% even assuming perfect MLA focus and alignment. The dashed-dotted curve assumes 80% ensquared energy at the inductor, our best guess from Zemax simulations, but still does not account for all the lost flux. We attribute the remaining lost flux to MLA defocus and/or misalignment and are investigating strategies to improve our MLA mounting procedure.

4.2 On-sky Commissioning

DARKNESS travelled to Palomar Observatory for commissioning in July 2016, then again for ongoing commissioning and science verification in November 2016 and April 2017. First-light was achieved on July 26, 2016, marking the first demonstration of J-band imaging with an MKID array on sky and the first diffraction limited images obtained with an MKID on sky, and April 2017 marked the first deployment of a PtSi MKID array on sky.




Figure 4.4: DARKNESS measured throughput vs. wavelength compared against theoretical prediction. The solid curve and shaded region shows the median and 1σ spread in throughput measured by all pixels. For the theoretical curves we've assumed flat 93% transmission through the 300 K window, applied the manufacturer transmission curves for the filters at 77 K and 4 K (roughly 97% across our band), assumed MLA fill factor of 93% and transmission of 98% (from manufacturer), inductor fill factor of 90% due to gaps between the meandered line, and PtSi measured absorptivity from Szypryt et al. (2016). These parameters are all measured or otherwise well constrained, leaving the MLA spot size as the final factor. From Zemax simulation we expect ensquared energy at the inductor to be 80%, but this is very sensitive to MLA focus and alignment.

4.3 Reduction and Analysis Pipeline

DARKNESS’s data reduction is largely based on the ARCONS pipeline (van Eyken et al., 2015), however the fully integrated pipeline to produce final calibrated photon lists is not yet available for DARKNESS data. Development effort so far has been focused on speeding up the initial accessing of the binary data. The ARCONS pipeline is written around the HDF5 file format (H5), and observation data is saved directly to H5 files as a list of photon packets organized by pixel for a user defined total exposure time. To facilitate future high-speed feedback from DARKNESS to P3K, DARKNESS data is communicated to the control computer every $500 \mu\text{s}$ and recorded to binary file every second as a timestream of photon packets, not sorted by pixel. The first step in the reduction pipeline is to collect the desired 1-second files for a given target and format them into an indexed H5 file that can be funneled into the existing ARCONS pipeline. After that, calibrations are carried out with the mostly unmodified ARCONS-pipeline modules and we perform subsequent processing and analysis (image registration, photometry, etc.) with a combination of custom *Python* scripts (mostly *NumPy* and *SciPy*) and the Vortex Image Processing pipeline (VIP; Gomez Gonzalez et al., 2017). Ultimately, these integrated packages will be streamlined and serve as the foundation for our statistical speckle suppression pipeline, to be presented in future work.

The only major revision to the ARCONS pipeline is how we perform hot pixel masking. As mentioned in van Eyken et al. (2015), TiN on Si MKIDs exhibited random “switching” behavior that caused pixels to appear hot at random intervals and for ran-

dom durations, and required a sophisticated hot pixel cleaning script. PtSi on Sapphire resonators do not show this behavior. Hot pixels are now mostly related to non-ideal readout parameters, manifesting as either constantly high count rates (bad phase threshold) or “beating” in intensity (two readout tones too close together). We simply flag these bad pixels using periodic dark exposures.

To facilitate quicker verification of our on-sky imaging performance, we performed most commissioning observations through a J-band filter. For the results presented in the remainder of this section we forego the spectral calibrations and treat this data as conventional imaging. The photon time streams are binned in software to frames with 1-second effective integrations that are then dark subtracted, flat corrected, and aligned.

4.3.1 Optical Checkout

Figure 4.5 shows a gallery of DARKNESS on-sky images, one from each observing trip, demonstrating the qualitative optical performance. In addition to the reduction above, the images shown here are composed of several median-combined frames from multiple dither positions to fill in missing pixels, which are then smoothed with a Gaussian kernel with radius λ/D . The top left frame, SAO 65485, is the official first-light image from July 2016. We see a diffraction-limited, though still highly aberrated PSF. Downstream from P3K there are several off-axis parabola (OAP) mirrors in SDC that are expected to introduce significant astigmatism, so the low order aberrations we see here are not surprising. This astigmatism can be corrected by two methods. Some can be taken out

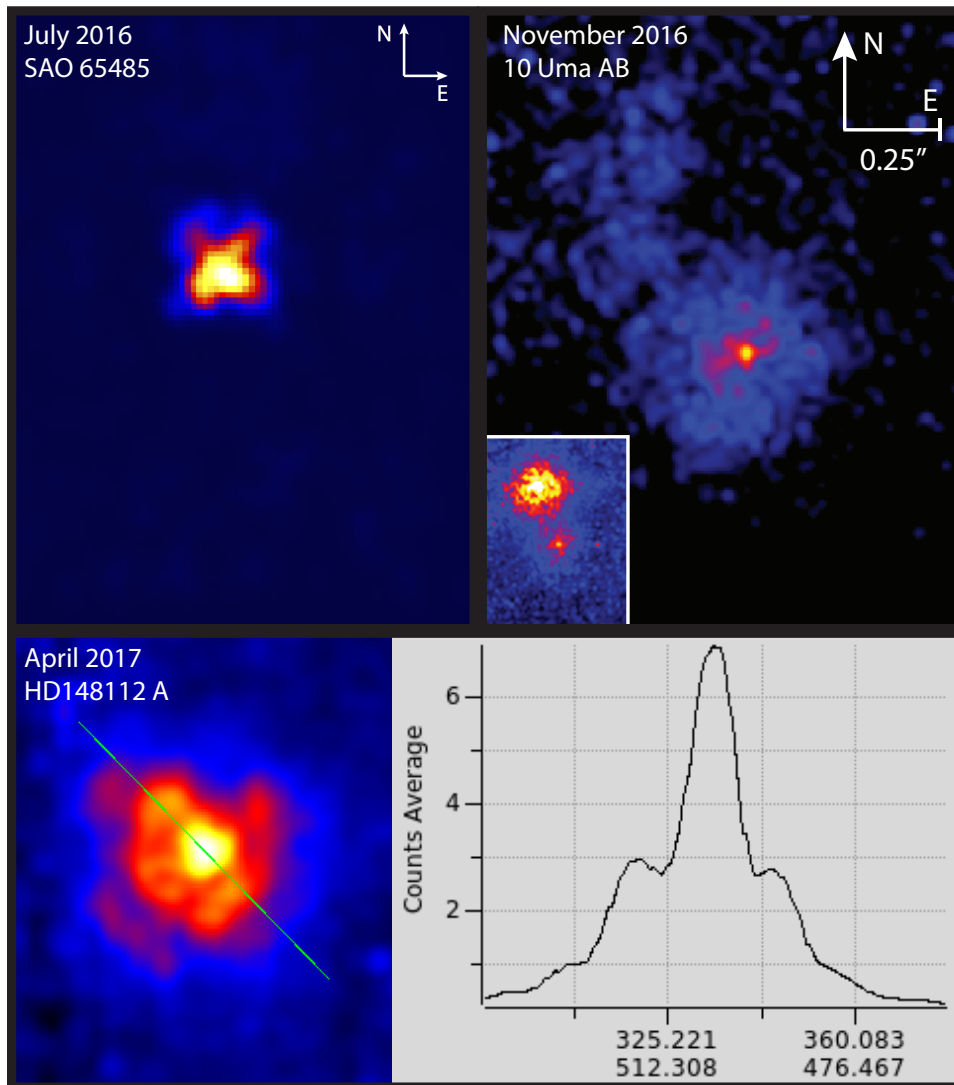


Figure 4.5: Collection of on-sky images from DARKNESS's first three observing runs, qualitatively showing the verification of PSF quality, coronagraph rejection, and plate scale. See text for details.

by adjusting the light-path through the SDC such that astigmatism from one OAP is partially cancelled by another. Low order aberrations can also be removed by *Zernike tuning*, wherein P3K’s low-order DM setpoint is manually adjusted by putting power into certain low-order modes.² We prefer the first method if possible because it assumes that the PSF is “perfect” at the SDC FPM where it matters most for coronagraph performance and aberrations are introduced downstream. Tuning Zernikes using DARKNESS images will make the final PSF look nice, but actually aberrate the PSF at the FPM.

The top right frame from November 2016 is a binary, 10 Uma AB, with separation of 0.47” at the time of observation, $V_{prim}=3.96$, and $\Delta V \approx 2$. The large frame shows the system with the coronagraph FPM installed, and inset shows FPM removed to reveal the primary. For this observation we’ve cleaned the PSF using only Zernike tuning, but considering the size of our FPM, contrast degradation is minimal. From this observation we can fit a centroid to both objects, and using the known separation we calculate an on-sky plate scale of 22 mas per pixel. The diffraction limit λ/D for a 5.1 m telescope at 1.25 μm is roughly 50 mas. So we achieve Nyquist sampling at these wavelengths, but are slightly sub-Nyquist at 0.8 μm .

The bottom panel is mostly a “glamour shot” from the April 2017 run after astigmatism off-loading with SDC and Zernike tuning. P3K reported 0.7” seeing, and in this image of HD148112 A and the accompanying projection we can see the first, and possibly second, airy rings of the diffraction pattern with a little help from some pinned speckles.

²P3K does its correction by decomposing the WFE into Zernike polynomials, which are related to familiar optical aberrations. The first few Zernike modes are 0: piston, 1: Tip, 2: Tilt, 3: Defocus, 4: Astigmatism, 5: Astigmatism, 6: Coma, 7: Coma, 8: Spherical Aberration...

4.3.2 On-sky Contrast

To estimate the raw contrast achieved with the coronagraph we observed π Herculis (J=0.79) on April 9th, 2017. We first observed the unocculted PSF using a 1% neutral density (ND) filter to ensure we could perform photometry on an unsaturated PSF core, and also with the coronagraph Lyot stop installed to ensure proper normalization of the throughput. We then moved the coronagraph FPM in and removed the ND filter to observe the surrounding speckle pattern. Examples of the unocculted and occulted long-exposure images are shown in Figure 4.6 (Left). We proceed using tools from the Vortex Image Processing pipeline (VIP; Gomez Gonzalez et al., 2017) to perform basic aperture photometry on the unocculted PSF and estimate contrast as a function of angular separation. In the standard procedure, average intensity at a given separation is estimated with a ring of λ/D sized apertures. Figure 4.6 (Left) shows an example of the apertures used at $4\lambda/D$ separation for both occulted and unocculted images. These intensities are then normalized by the intensity of the unocculted, unsaturated PSF core measured with the same aperture. Figure 4.6 (Right) shows the resulting contrast curves of the occulted and unocculted PSF, plotting mean raw contrast (i.e. no post-processing to remove the static speckle pattern) measured in $1\lambda/D$ annuli with 1σ error bars. We see that the coronagraph is providing at least an order of magnitude raw rejection at all separations. We also note that the 1σ error estimates do not include the now-standard small number statistics correction from Mawet et al. (2014).

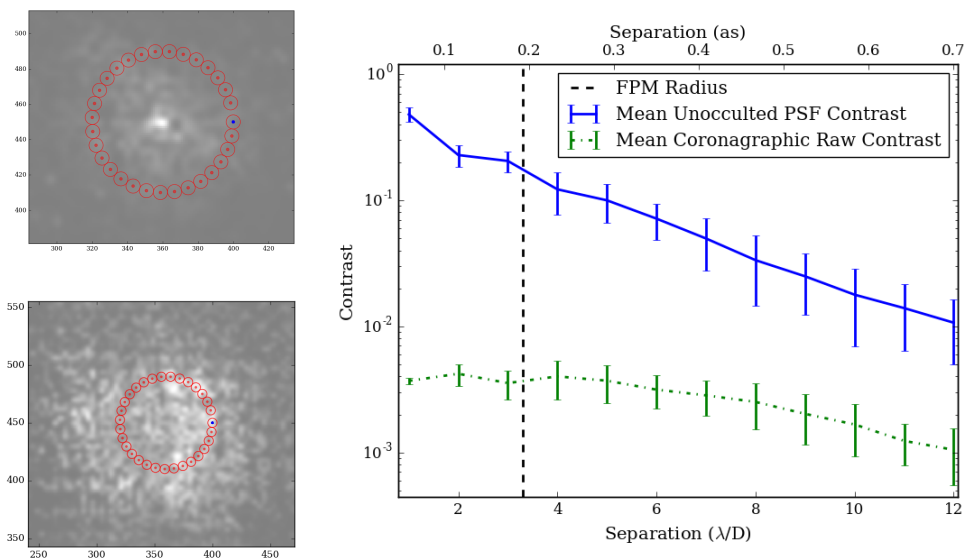


Figure 4.6: Mean J-band contrast as a function of angular separation with and without coronagraph FPM installed. (Left) Long exposure images of the unocculted (Top) and occulted (Bottom) PSF provided by the processing pipeline VIP showing the ring of apertures used to estimate contrast at $4\lambda/D$ separation. Both images are *log* scaled, but with different intensity minima and maxima, and set to slightly different spatial scales. (Right) the resulting contrast vs. separation curve after taking the mean of intensities measured by similar aperture annuli at λ/D spacing and normalizing by the unocculted, unsaturated PSF core.

4.3.3 Pupil Viewer

The DARKNESS fore-optics include a second optical arm for target finding and pupil viewing. The large FOV facilitates easy target acquisition on-sky and the pupil viewing mode is especially critical to check Lyot mask alignment relative to the telescope pupil (demonstrated in Figure 4.7). This mode is also quite useful for coronagraph FPM alignment. A coarse alignment is most easily achieved by moving the star relative to the FPM position using small telescope offsets and observing the dimming in the final pupil, then viewing the occulted PSF on the DARKNESS array to finely adjust the alignment. In Figure 4.7 we see that we achieve a very clean pupil image in this calibration arm.




Figure 4.7: Image of the Hale 200" Telescope pupil as seen by DARKNESS's fore-optics pupil viewing camera. This mode is used to align the coronagraph pupil mask with the telescope pupil (the mask is still visibly clocked relative to the telescope's spiders in this image) and can help with coronagraph focal plane mask alignment by searching for the mask position where the residual light in this pupil image is minimized.

Chapter 5

Studying Speckle Lifetimes and Implications for Statistical Speckle Discrimination

With DARKNESS successfully commissioned and basic operations demonstrated on sky, we return to the issue of atmospheric speckles discussed in Chapter 1. In this Chapter we show a preliminary investigation of speckle lifetimes at very short timescales and implications for statistical speckle discrimination to demonstrate the unique power of MKIDs.

5.1 Speckle Lifetime Review

In Chapter 1 we made mention of a few characteristic lifetimes that we may expect to see speckles evolve over, which we will expand upon here. We will skip the discussion of quasi-static speckles since an enormous body of work has been dedicated to overcoming that noise floor already.

Macintosh et al. (2005) finds that speckles resulting from imperfect AO correction can show two characteristic timescales. The shorter of these two is related to AO measurement error. Since these are essentially random errors (i.e. due to noise in the WFS measurement) we can naively expect them to refresh with every AO update. For an AO system running at 1 kHz this translates to 1 ms decorrelation times, though Macintosh et al. (2005) point out that the proportionality will depend on the AO controller and expect $\tau_{meas} \propto \frac{1}{10}\Delta t_{AO}$ where Δt_{AO} is the AO update rate. Resolving this timescale will be challenging, since it would require count rates on the order of 10k photons per second per speckle to achieve a decent signal in 1 ms effective exposure (DARKNESS has a soft limit in firmware of 2500 photons per second per pixel to avoid photon pile-up and readout buffer overflows). We may resolve this in the brightest speckles, but will be limited in our ability to probe this timescale over the entire AO correction region.

The next fastest timescale we may expect to resolve is the atmospheric coherence time, which we stated back in Equation 1.2 but will restate here for convenience:

$$\tau_0 = 0.314r_0/\bar{v} \tag{5.1}$$

where r_0 is the atmospheric coherence length (the Fried parameter), and \bar{v} is the mean

wind speed. Since $r_0 \propto \lambda^{6/5}$, this parameter depends on the wavelength of observation and prevailing wind speed at the time. At Palomar we assume $r_0(500 \text{ nm})=9.2 \text{ cm}$ (Dekany et al., 2013), which scales to J-band as $r_0(1250 \text{ nm})\approx 28 \text{ cm}$. For a wind speed of 9 m/s (average at Palomar is more like 5 or 6 m/s) we expect $\tau_0 \approx 10 \text{ ms}$, a typical value for speckle boiling time in the literature. This timescale is approaching the point where MKIDs may resolve the speckle evolution, but given the current count rate limitations we could not collect enough photons in the requisite time to do useful correction.

Macintosh et al. (2005) also showed that uncorrected atmospheric aberrations (i.e. from AO bandwidth error) can produce speckles with lifetimes related to the telescope aperture clearing time, with typical decorrelation times of:

$$\tau_{atmo} = 0.6D/\bar{v} \tag{5.2}$$

where D is the telescope aperture diameter. Again, using Palomar as our fiducial point and $\bar{v}=9 \text{ m/s}$ we expect this to be on the order of 0.3 s. For a larger telescope and better conditions this timescale can be greater than 1 s (e.g. $D=8 \text{ m}$ and $\bar{v}=3 \text{ m/s}$ gives $\tau_{atmo}=1.6\text{s}$).

For these random processes we expect the speckle intensity variance to scale according to the speckle decorrelation times: $\sigma^2 \propto t_{dec}/t_{int}$, and the total variance will be the sum of variances from the various speckle populations. For a conventional IFS that is not able to resolve these speckles in time, long integrations are taken to allow many realizations of the speckle field to average together into a smooth halo. From the above proportionality we can easily see why atmospheric speckles present the biggest challenge as their variance

will decrease the slowest of the random speckle populations. Understanding these various timescales has significant implications for applying a variety of speckle control and removal techniques. For statistical speckle suppression, we need to know the timescale over which the speckles decorrelate to ensure we aren't averaging over many realizations of the speckle pattern and blurring the statistics. For speckle nulling, if we want to track and remove the speckles at high frame rate we had better make sure they are behaving as we expect on those timescales.

In our analysis we largely follow the formalism layed out in Fitzgerald & Graham (2006) and Milli et al. (2016) for studying the temporal evolution of speckle fields via the autocorrelation function of a given speckle's lightcurve. Fitzgerald & Graham (2006) study adaptively corrected non-coronagraphic images from Lick Observatory using 5 ms exposures ($\Delta t=14.5$ ms), with the aim of resolving the sub-second decorrelation times and confirming the MR intensity PDF. Milli et al. (2016) present their work more in the context of probing quasi-static lifetimes for predicting performance of differential imaging, using coronagraphic XAO images recorded in a special windowed readout of the IRDIS subsystem in SPHERE to achieve 1.6 Hz frame rates. They seem to resolve a rapid decorrelation on the order of the aperture clearing timescale, but surprisingly find the same timescale with the internal calibration source in the absence of atmospheric aberrations.

With the unique capabilities of DARKNESS we conduct a similar analysis to the above mentioned works, and with the ability to set *ex post facto* integrations in the

software we can easily explore timescales from milliseconds to hours with a dynamic re-sampling of a single data set.

5.2 Observations

On the night of April 09, 2017 around UTC 12:30 we observed the bright star (magnitude $J=0.79$) π Herculis (hereafter π Her). This target was selected to maximize the photons in the residual speckle field even after coronagraphic rejection, and because it was near zenith at the time of observation (airmass ~ 1.01). The SDC was configured with the 330 mas diameter FPM ($6.6 \lambda/D$ in J band) and 15% undersized Lyot stop, and DARKNESS filter wheels were set to J-band and no neutral density. Seeing was reported as 1.0" by the Palomar 18-inch telescope seeing monitor, which is typically worse than the seeing actually experienced at the 200" dome.¹ Median wind speed at the time was 7.1 m/s and the P3K WFS was operating at 1.5 kHz. From these values we assume a ballpark $r_0(1250 \text{ nm}) \approx 30 \text{ cm}$ and $\tau_0(1250 \text{ nm}) \approx 12 \text{ ms}$, and $\tau_{atmo} \approx 0.4 \text{ s}$.

In total we collected ~ 1 hour of data on target while performing a variety of dither patterns to fill in missing pixels. To ease data reduction and avoid the need to concatenate light-curves from different pixels, we focus on a stable 2 minute chunk of data from the beginning of the observation during which the speckle pattern was stationary on the DARKNESS array. Figure 5.1 (Bottom Left) shows the median image of the long dither sequence after binning to 1 second effective exposures and aligning these 1 second frames.

¹Earlier in the night the seeing monitor reported 1.0" while P3K reported 0.7", but we did not query P3K while observing π Her.

The following afternoon we collected a companion data set on P3K’s internal calibration light source (referred to hereafter as white light or WL). DARKNESS and the SDC were setup with the same configuration (J-band filter, no ND, 330 mas FPM, 15% Lyot undersize) and P3K was operated at 1kHz. Like the on-sky data we took a long integration featuring several dither moves, with the median frame shown in Figure 5.1 (Top Left) following the same reduction as the π Her data.

5.3 Autocorrelation Analysis

To study the temporal evolution of the speckle fields we follow the autocorrelation formalism layed out in Fitzgerald & Graham (2006) and Milli et al. (2016). For a given pixel we take the raw photon timestream and bin it into a time series with effective short exposures chosen based on the speckle decorrelation times we wish to resolve. Assuming temporal stationarity during the observation (meaning the mean μ and variance σ^2 are time-independent) we can assume that the autocorrelation of the lightcurve depends only on lag time τ , such that:

$$R(\tau) = \frac{E[(I_t - \mu)(I_{t+\tau} - \mu)]}{\sigma^2} \quad (5.3)$$

where E is the expectation value operator. The normalization by σ^2 ensures $R(0)=1$ and subtraction of μ removes any DC signal, such that $R(\tau \gg \tau_{atmo})=0$ where τ_{atmo} is the longest expected speckle lifetime we’ll see (we are effectively assuming that any quasi-static speckle signal will be constant over the couple minutes we analyze here.) An example of this procedure is shown in Figure 5.1 with a comparison of the π Her and WL

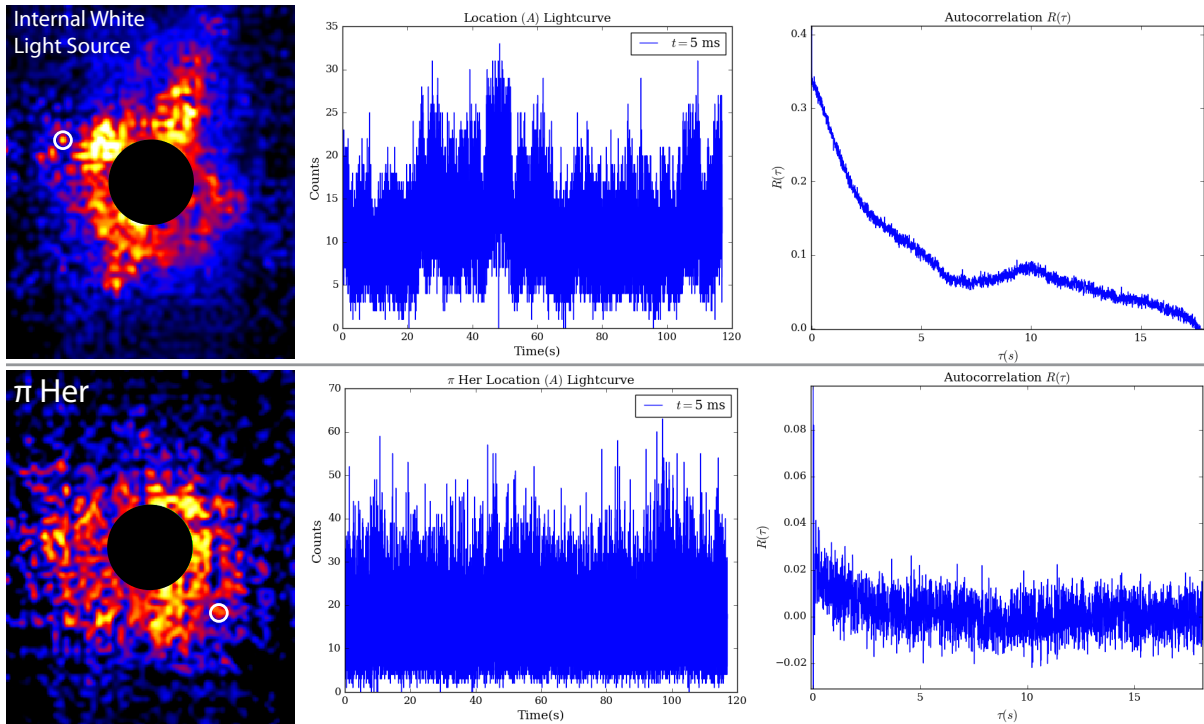


Figure 5.1: P3K internal white light (WL) source vs. π Her speckle correlation comparison. (Left) Long-exposure J-band images with the coronagraph FPM installed, with speckles highlighted in both at similar intensities and separations. (Center) lightcurves from just a 2-minute chunk of each dataset where we’ve taken the raw photon time streams from the selected speckles and binned in software to 5 ms effective exposures. (Right) Autocorrelation functions of these lightcurves. We see that the on-sky data is almost entirely uncorrelated after a fraction of a second, but shows some residual correlation for the first few seconds. In the absence of a turbulent atmosphere, the WL source shows much greater correlation for a longer time, however there is a similar linear decay to the on-sky data in the first few seconds. For both sets any quasi-static speckles with \gtrsim minute lifetimes would have been subtracted with the de-meaning done in the autocorrelation analysis.

data sets described above. The (Left) image in both registers shows the median J-band images with the coronagraph FPM installed, with speckles highlighted in both at similar intensities and separations. The (Center) panels in both registers show lightcurves from the 2-minute chunk of each dataset where we’ve taken the raw photon time streams from the selected speckles and binned in software to 5 ms effective exposures. The

(Right) panels show $R(\tau)$ for the selected speckles, from which we can extract various speckle decorrelation timescales. Ultimately we want to apply this information to find appropriate exposure times for performing SSD (Section 5.4), but first we investigate a couple unexplained features in the autocorrelation curve.

5.3.1 A Seconds-Long Decorrelation of Possible Instrumental Origin

Qualitatively, we see that the on-sky data is almost entirely decorrelated after a fraction of a second, but shows some residual decay for the first couple seconds. This is not surprising, as we would expect potentially three different rates of decorrelation based on the speckle timescales listed above, two of which may not be apparent at this scaling. Since we are most interested in τ_{atmo} we can rebin the photon timestream to highlight that particular range. With a 100 ms effective exposure time we expect to essentially average over features on $t < \tau_0 \approx 12$ ms timescales, but still resolve features on $t > \tau_{atmo} \approx 0.4$ s timescales. Indeed, this is what we see in Figure 5.2 (Left), showing the autocorrelation for the same π Her speckle, now using 100 ms exposures and zooming in to the shortest lag times. Fitting an exponential of the form $R(\tau) \propto \exp^{-\tau/\tau_c} + R_0$ to the autocorrelation, we extract a decay time $\tau_c = 2.1$ s. This is perhaps the right order of magnitude to be related to τ_{atmo} , but we see a similar timescale in the WL data as well, shown in Figure 5.2 (Right). The WL autocorrelation curve is not as easily interpreted, as it seems to re-correlate after a few seconds. Considering there is no wind or turbulent atmosphere to




Figure 5.2: (Left) ACF of π Her lightcurve. (Right) ACF of WL lightcurve. Both with 100 ms effective short exposures. Exponential fit gives 2 s decay time.

generate correlation features on short timescales, we may expect the speckle to remain correlated indefinitely. However, we still observe a steep decorrelation in the first few seconds, and if we attempt to fit an exponential to the region where the correlation is still trending downward, we extract a similar decay time $\tau_c=2.8$ s.

Even more puzzling is that a similar step decorrelation regime was reported by Milli et al. (2016) in SPHERE’s on-sky and internal calibration source data, also (perhaps coincidentally) on the order of a few $\times \tau_{atmo}$. These timescales are much too fast to be related to the flexure or long term temperature drift associated with quasi-static speckles, suggesting they may be caused by density fluctuations in the mostly stationary air within the instrument. Clearly more investigation is required, and high-cadence data from other such high-contrast instruments would be useful.




Figure 5.3: (Left) Autocorrelation of π Her lightcurve. (Right) Autocorrelation of WL lightcurve. Both with 2 ms effective short exposures. Exponential fits give 2 and 3 ms decay times, respectively. In the π Her figure we also find a noticeable anti-correlation at 12 ms lag time, corresponding to a periodic signal in the speckle intensity with roughly 40 Hz frequency.

5.3.2 A Strong Anti-correlation at 12 ms Spacing

With the speckles selected in Figure 5.1 we achieve the necessary count rate to explore even finer time sampling. Figure 5.3 (Left) shows the autocorrelation function for the selected π Her speckle using 2 ms effective exposures. We again fit an exponential decay to extract the timescale of this rapid decorrelation and find $\tau_c=2$ ms. Figure 5.3 shows the same analysis on WL, and returns a ~ 3 ms decay time. These two lifetimes are proportional to ($3\times$) the respective P3K frame rates during the two observations — 1.5 kHz on-sky, 1 kHz on WL leading to $\Delta t_{AO}=0.7$ and 1.0 ms, respectively — suggesting that we can indeed resolve this short decorrelation regime.

More striking though (and more important for practical purposes) is that the intensity becomes anti-correlated at 12 ms lag time, corresponding to a periodic signal in the




Figure 5.4: Lightcurve from selected π Her speckle showing an intensity modulation at ~ 40 Hz.

lightcurve with period around 24 ms (~ 40 Hz). Figure 5.4 shows a 200 ms span of the lightcurve where this “switching” behavior is apparent. This has severe implications if we wish to attempt speckle nulling at tens to hundreds of Hz to correct atmospheric speckles, as a periodic signal unrelated to the atmospheric speckle decorrelation time will confuse the speckle nulling algorithm — it will be impossible to distinguish this periodic dimming of the speckle from the dimming caused by a DM probe pattern that is out of phase with the true speckle, leading to an incorrect determination of the speckle phase and poor loop convergence.

The exact cause of this periodic signal is unclear, but we can identify likely candidates by considering other processes (physical and instrumental) with similar periods. Instrumentally, the only subsystem that operates near this rate is the tip-tilt mirror (TTM) of P3K, which has a known bump in the residual error PSD near 30 Hz, seen in both tip and

tilt directions, that is supposedly being injected by the system (Tesch, J. 2015, private communication). We don't see the anti-correlation in our WL data, however, this data was taken while the telescope was parked at zenith and there is precedent for telescope pointing drive vibrations to inject serious residuals into XAO systems (see for instance the surprise 60 Hz feature GPI discovered at Gemini South (Poyneer et al., 2014)). This potential culprit will need to be investigated fully with on-sky and WL data taken during telescope tracking, at a variety of pointings, and compared against a more careful logging of P3Ks TTM residuals.

Another likely culprit is the atmospheric coherence scale, considering the decorrelation happens at lag times $\approx \tau_0=12$ ms predicted at the beginning of this Chapter. Again we turn to Macintosh et al. (2005) where a similar modulation was found in simulations. They find that for a 2-layer atmosphere model composed of two phase screens, each featuring phase ripples with spatial frequency \mathbf{k}_0 and translated relative to each other at velocities \mathbf{v}_1 and \mathbf{v}_2 , speckles show an intensity modulation imprinted on top of their longer lifetime variations with period $\mathbf{k}_0 \cdot \Delta \mathbf{v}$. Ignoring wind directions and using a 40 Hz period with $2\pi/r_0$ to approximate k_0 yields $\Delta v=1.5$ m/s, which is not unreasonable considering the median wind speed overall was 7.1 m/s.² Again we cannot rule out this source for the modulation, but will explore it further as we analyze more on-sky data taken under a variety of wind conditions.

²The readily available Palomar weather records give median wind speed and direction, but no information about the vertical profile.

5.4 Implications for Statistical Speckle Discrimination

Our investigation of the various speckle decorrelation times at play provides valuable insight into systematic effects that could hinder a stable speckle nulling loop, but we would also like to leverage this information in post processing to perform statistical speckle discrimination (SSD).

5.4.1 Leveraging Autocorrelation Analysis for SSD

Gladysz & Christou (2008) point out that observers should avoid simply using the shortest possible exposure, as low count rates will tend toward Poissonian statistics. On the other hand, too long of an exposure could allow many realizations of the speckle pattern to average together, and by the central limit theorem the MR statistics will become Gaussian. Additionally they suggest that, contrary to conventional wisdom, exposures many times longer than the atmospheric coherence time could be used without penalty, but this is something we can test directly with our MKID data. We would expect that for sufficiently short time sampling, the intensity distribution of a speckle light curve will appear largely MR. As we re-sample the same data to longer exposures, averaging over greater numbers of decorrelation times, the statistics of each lightcurve should tend more and more toward Gaussian. Figure 5.5 shows the result of this experiment on our π Her test speckle where we resample the photon timestream at a series of short exposure




Figure 5.5: Shapiro-Wilk test for normality as a function of speckle exposure time. At very short exposure times the target π Her speckle shows clearly non-Gaussian intensity distribution, but as we increase our exposure times to average over a greater number of short decorrelation times we see the distribution tend toward Gaussian (higher W).

times, and in each instance check the normality of the lightcurve intensity distribution with a Shapiro-Wilk test. We see that for very short exposures the statistics are very clearly non-Gaussian (lower W) and around 30 ms exposures the test statistic W levels off. From this analysis we may select an optimal exposure time to maximize signal in each speckle while preserving the MR statistics.

5.4.2 First Attempt at SSD with DARKNESS

On April 9, 2017 we observed the bright double star system SAO 65921 which has $J_{prim}=4.6$, $\Delta V=3.9$, and separation $0.9''$ with position angle 143° at time of observation. This system was selected as an easy first test for performing SSD because the secondary would be bright enough to see above the speckle pattern in our real-time display, but

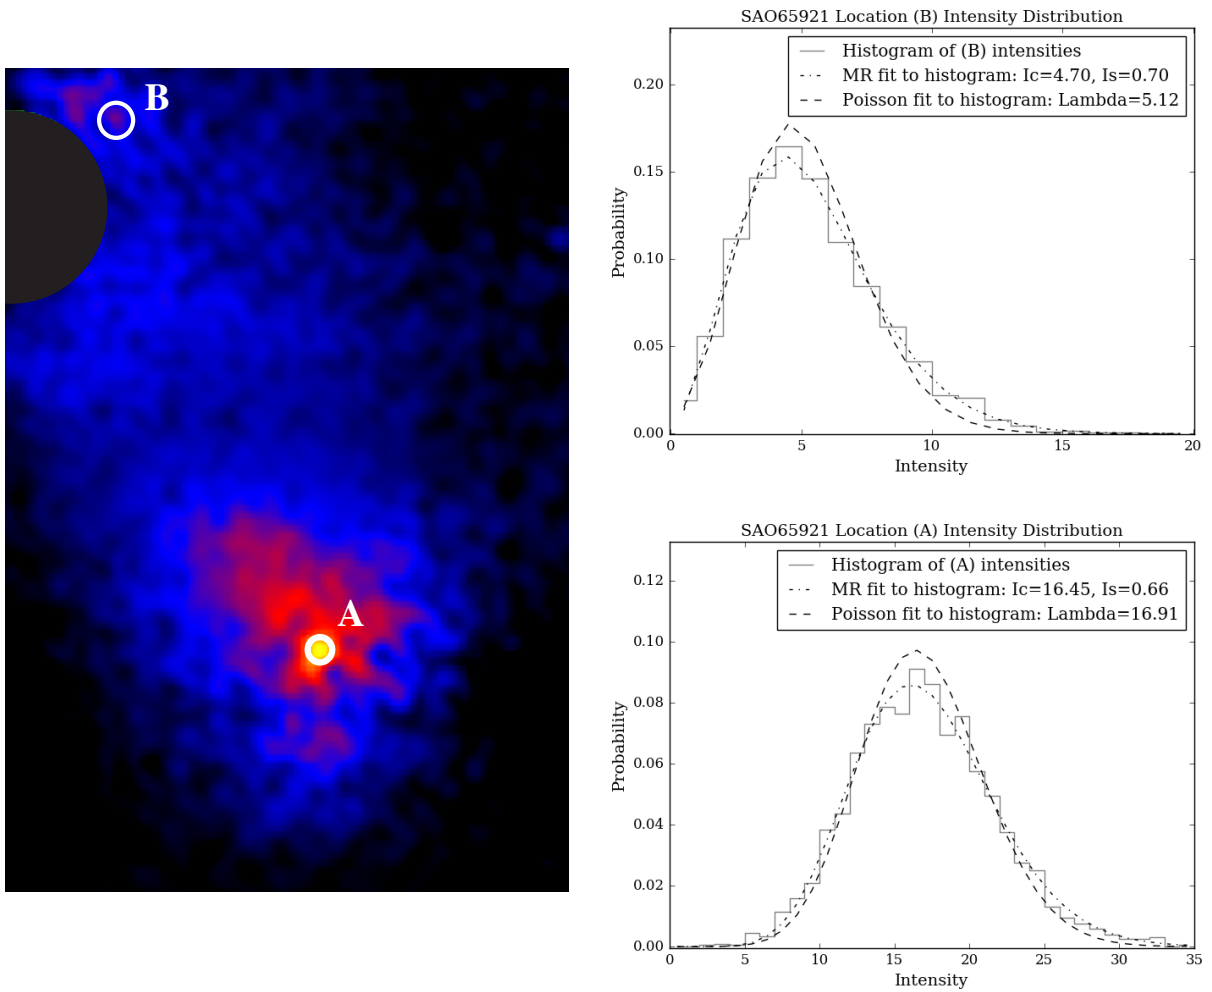


Figure 5.6: (Left) SAO 65921 median broadband frame following the same reduction as π Her described above. (Top Right) shows the resulting intensity distribution at the speckle location (B), and (Bottom Right) shows the same for the secondary location (A). By eye these distributions are qualitatively quite different, but still effectively fit by MR distributions, albeit with very different I_c/I_s ratios that characterize the MR skewness. For the speckle we see this ratio is 6.7, whereas for the companion it is nearly 25.

similar intensity overall to the low-order speckles near the FPM. Figure 5.6 (Left) shows the median broadband frame following the same reduction as π Her described above. In this case we observed the target broadband to increase the count rate in the speckles and also serve as a demonstration of spectral differential imaging (left for future work).

We acknowledge that the wavelength dependence of τ_0 will result in different decorrelation times for speckles at different wavelengths, which we must take into account. We perform the same Shapiro-Wilk analysis on the selected speckle (location B) as we did for π Her and find a similar roll-off around 30 ms. With median wind speed of 4.7 m/s during this observation we would expect the shortest atmospheric coherence time to be $\tau_0(800 \text{ nm})=11 \text{ ms}$. For the following analysis we safely select 5 ms effective exposures.

Figure 5.6 (Top Right) shows the resulting intensity distribution at the speckle location (B), and (Bottom Right) shows the same for the secondary location (A). By eye these distributions are qualitatively quite different, but we see that both are still effectively fit by a MR distribution (and both are poorly fit by Poissonians). Proceeding with a useful pipeline requires a quantitative way to distinguish these distributions. One lever may be the ratio of I_c/I_s that characterizes the MR skewness. For the speckle we see this ratio is 6.7, whereas for the companion it is nearly 25. A more powerful test may be a two-point statistical test such as a Kolmogorov-Smirnov test that simply checks if two populations are drawn from the same distribution, regardless of the functional form. However, a K-S test assumes some knowledge about the companion location, or else must be computed for every pair of λ/D elements at a given radius.

For the time being we will observe that location A's statistics appear quite a bit more Gaussian, and fall back on the Shapiro-Wilk test again. Figure 5.7 shows a comparison of W vs. exposure time for the secondary and speckle locations highlighted in 5.6, as well as the WL and π Her speckles from earlier. In all speckle cases the statistics remain




Figure 5.7: SAO 65921 speckle vs. companion Shapiro-Wilk test, with the same metric from WL and π Her speckles included for reference. In all speckle cases the statistics remain non-Gaussian out to long exposures while the secondary converges toward high W almost instantly.

non-Gaussian out to long exposures (in fact the WL speckle turns around at one point and seems to become more MR at larger exposure times) while the secondary converges toward high W very rapidly. The next step will be to apply this technique across the full image, making a map of W at a given short exposure time where the secondary should stand out with high SNR, but already this appears to be a promising avenue.

Chapter 6

Future Work & Conclusions

We have presented the design and initial on-sky performance of DARKNESS, a MKID based IFS for high contrast imaging. Including the successful commissioning in July 2016, DARKNESS has completed three observing trips to Palomar where it operates with the P3K XAO system and the SDC. These early trips marked several milestones for DARKNESS and MKID development in general, including the first deployment of a PtSi MKID array on-sky, the first diffraction limited observations with MKIDs, and the first J band observations with MKIDs. A robust high-contrast MKID data pipeline is in the works, but already the data is yielding unique insights into speckle behavior and high-contrast instrument systematics at high frame rates.

UVOIR MKID technology is still maturing rapidly — we have active research programs dedicated to improving the PtSi resonator Q_i which will improve raw pixel yield and energy resolution, as well as exploring direct deposition of anti-reflection coatings on

the MKID inductors to drastically improve detector QE. The results shown here can be considered as a snapshot of the current state of the instrument, which will be upgraded with a new MKID array when these improvements come online.

The SDC has demonstrated competitive contrasts with speckle nulling (Bottom et al., 2016b), however, the current bottleneck is the slow readout time of the science camera, PHARO (Hayward et al., 2001), resulting in overly long calibration times and limited speckle control due to their decorrelation timescales being comparable to the control loop delays. Because of this slow control bandwidth (~ 0.3 Hz), speckle nulling converges slowly and is mostly limited to correcting static wavefront aberrations. We are currently working with the P3K team to implement a faster communication scheme between the P3K and DARKNESS control computers, and aim to demonstrate speckle nulling on-sky at > 10 Hz rates in the coming months.

The WFS camera and real-time controller (RTC) for P3K will be upgraded in early 2018. With these upgrades in place we anticipate improved J-band Strehl ratio and have procured a J-band vector vortex FPM that will be installed in DARKNESS's open lanai slot. The RTC upgrade will also open the door for kHz feedback from DARKNESS to target atmospheric speckles.

DARKNESS represents the beginning of a multi-instrument effort to utilize MKIDs for exoplanet imaging. The development work invested in DARKNESS has simultaneously supported the MKID Exoplanet Camera (MEC), a 20 kilopixel MKID IFS for SCEXAO shipping to Subaru in late 2017, and PICTURE-C (Cook et al., 2015), a

balloon-borne high-contrast platform that will fly a DARKNESS clone in 2019. Additionally, thanks to its portable design, DARKNESS is slated to join MagAO-X (Males et al., 2016) in 2019 as the first high-contrast instrument with MKID IFS backend in the Southern Hemisphere.

Appendix A

QE Testbed

Integrating spheres (IS) are a useful tool for measuring detector quantum efficiency (QE) and uniformity. A light source is injected into the IS, which is a hollow cavity coated internally with a diffuse, highly reflective coating. When viewed from an output port, the interior of the IS appears as a *Lambertian surface*, meaning it has uniform, constant surface brightness regardless of viewing angle. In the typical setup for conventional detectors, uniform irradiance is the goal and the diverging output beam is directly coupled to the detectors.¹ However, we wish to use the IS for an absolute throughput measurement of DARKNESS and require a bit more care in re-imaging the light on to the DARKNESS array properly. We must especially take care in simulating the $f/\#$ and telecentricity of the beam DARKNESS receives at the telescope to ensure proper performance of the MLA, which we expect to be the biggest source of uncertainty in our total system throughput. In this Appendix we provide a brief overview of our QE testbed

¹A handy guide to IS theory and applications can be found at www.labsphere.com

design so the reader can understand the sources of uncertainty in the measurement presented in Chapter 4 and the tradeoffs made to reformat the testbed for DARKNESS and future instruments.

A.1 QE Testbed Design

Figure A.1 shows a schematic overlay of the mechanical and optical design of the QE testbed. The light source is a pair of Tungsten and Deuterium lamps injected into a monochromator that disperses the light and outputs a selected wavelength to the IS. The IS object plane is then reimaged to a telecentric $f/300$ beam with a simple aperture/lens pair, which can be steered to an optical photodiode (PD), IR PD, or out to DARKNESS with a rotating fold mirror (FM). We then compare the count rate from the uniform image plane illumination in each MKID pixel with the power measured on the calibrated PDs. Since the MKID pixels and PD sensors both have a well defined collecting area, assuming a perfectly uniform re-imaging of the IS object, we can easily convert both measurements to photons per second per unit area.

This QE bed was originally designed for making the same throughput measurement with ARCONS, but DARKNESS required a complete redesign to accommodate a much larger $f/\#$ ($f/19$ for ARCONS to $f/300$) without sacrificing too much light at a small aperture. From our experience with ARCONS, a 6.4 mm diameter aperture provided sufficient illumination at the MKID array using 0.6 mm slit widths in the monochromator. For DARKNESS the primary challenge was achieving the beam length necessary for

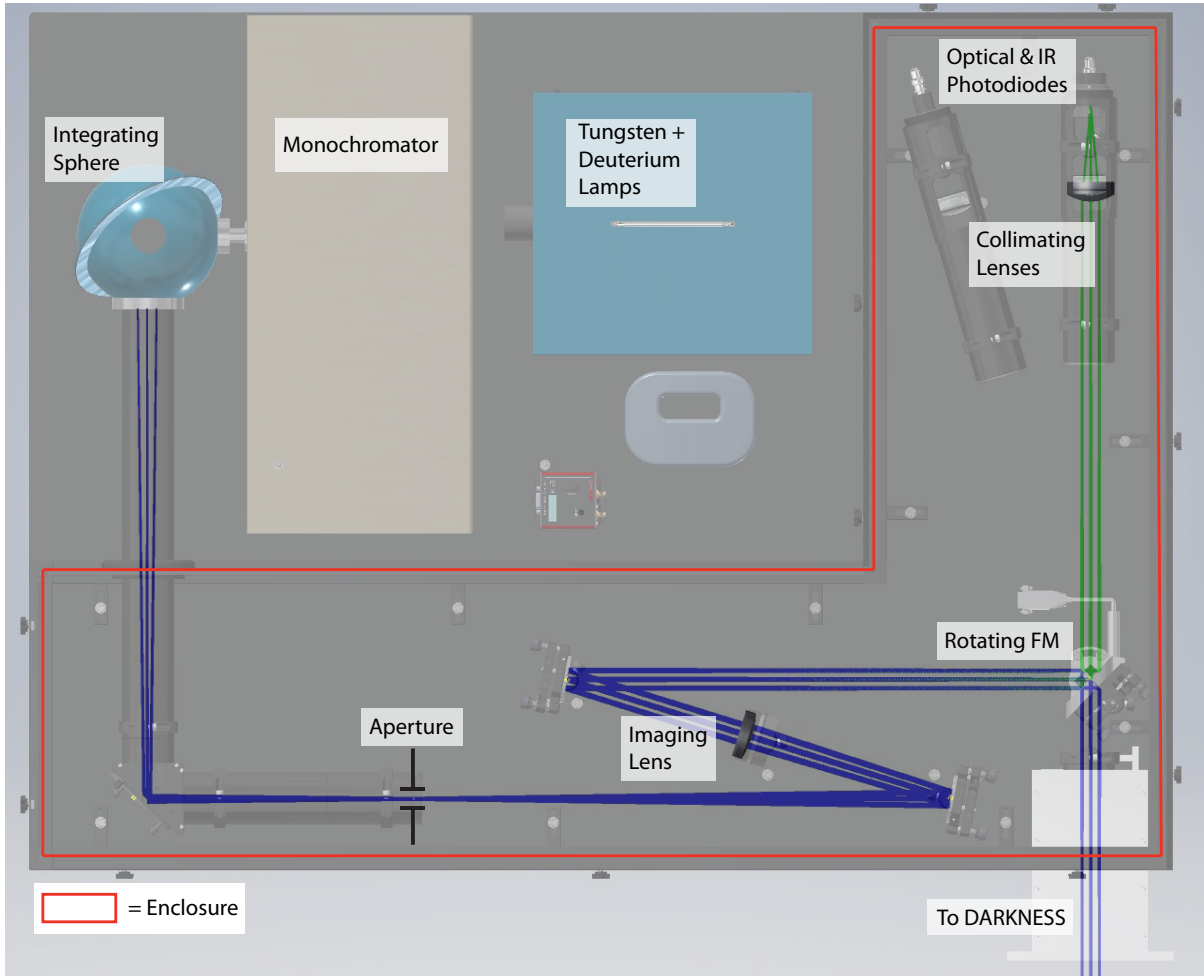


Figure A.1: Overlay of mechanical design and Zemax raytrace for QE testbed. See text for details. For reference, the optical bench is 48" on the long side.

$f/300$ with the largest possible aperture in the available bench space. With the folded design shown in Figure A.1 we use a lens with 750 mm focal length (Lens 1) and a 2.4 mm diameter aperture. The factor of ~ 9 lower throughput at the aperture can be partially compensated by opening the monochromator slits to increase source luminosity, at the expense of energy resolution. Resolution at the selected monochromator wavelength, $\lambda/\Delta\lambda$, can be determined with the inverse linear dispersion relation for the grating:

$$\frac{\Delta\lambda}{\Delta x} = \frac{10^6 \cos\beta}{mnf} \text{nm/mm} \quad (\text{A.1})$$

where Δx is the slit width in mm, β is the diffraction angle, m is the diffraction order, n is grating density in grooves/mm, and f is the exit focal length in mm. For our monochromator² $\Delta\lambda/\Delta x = 3.2$ nm/mm with the 1200 g/mm grating. The worst-case resolution, using the 600 g/mm grating and assuming fully open slits (3 mm) gives $R=52$ at $1 \mu\text{m}$ which is still far above our MKID energy resolution, meaning we can safely open the slits to maximum width. With the given aperture and fully open monochromator slits, we expect the count rate at the detector (assuming similar QE between ARCONS and DARKNESS) to be approximately 50% of that measured in ARCONS, which was typically 1000 cps. This is still an acceptable amount of illumination.

In the calibration arm with the PDs, we need to overcome a similar intensity issue: the IR photodiode did not receive enough light in the ARCONS configuration, but increasing the source luminosity would saturate the MKIDs. In the new testbed we overcome this issue by concentrating the uniformly distributed light from the image plane with an

²Spectral Products DK240 specs can be found at www.spectralproducts.com/dk240

additional collimating lens (Lens 2) just before the PDs. Effectively, this makes a final pupil plane with diameter given by the original aperture diameter in the re-imaging optics multiplied by the ratio of lens focal lengths, f :

$$D_{final} \approx D_{aper} \frac{f_{lens2}}{f_{lens1}} \quad (\text{A.2})$$

and we select a lens with appropriate f to make D_{final} much smaller than the PD sensor area. We select a $f=100$ mm lens to create a final pupil with $D_{final} < 1$ mm (the smaller IR PD sensor is 4×4 mm). With 100% of the light collected on the PD, we can back-out a power per unit area measurement by placing a well defined aperture mask just before the collimating lens. We now know exactly the collecting area in the image plane that is concentrated as a pupil on the PD and can determine the photons per second per unit area for comparison with the MKID data.

A.2 QE Testbed Verification

We performed extensive simulations in Zemax of the above design, checking for $f/\#$, telecentricity at the MLA (equivalent to an exit pupil at ∞), pupil diameter D_{final} at the PD, and uniformity of the final image plane's surface brightness.

Our first concern is whether our scheme for concentrating light on the PDs is valid. Figure A.2 shows the spot diagram and encircled energy calculation from the PD arm of the testbed. We see that the resulting pupil images are well aligned on the optical axis with spot sizes well within the PD sensor area. With proper optical alignment and focusing we can assume that 100% of the light is collected at the PD, but to safely turn

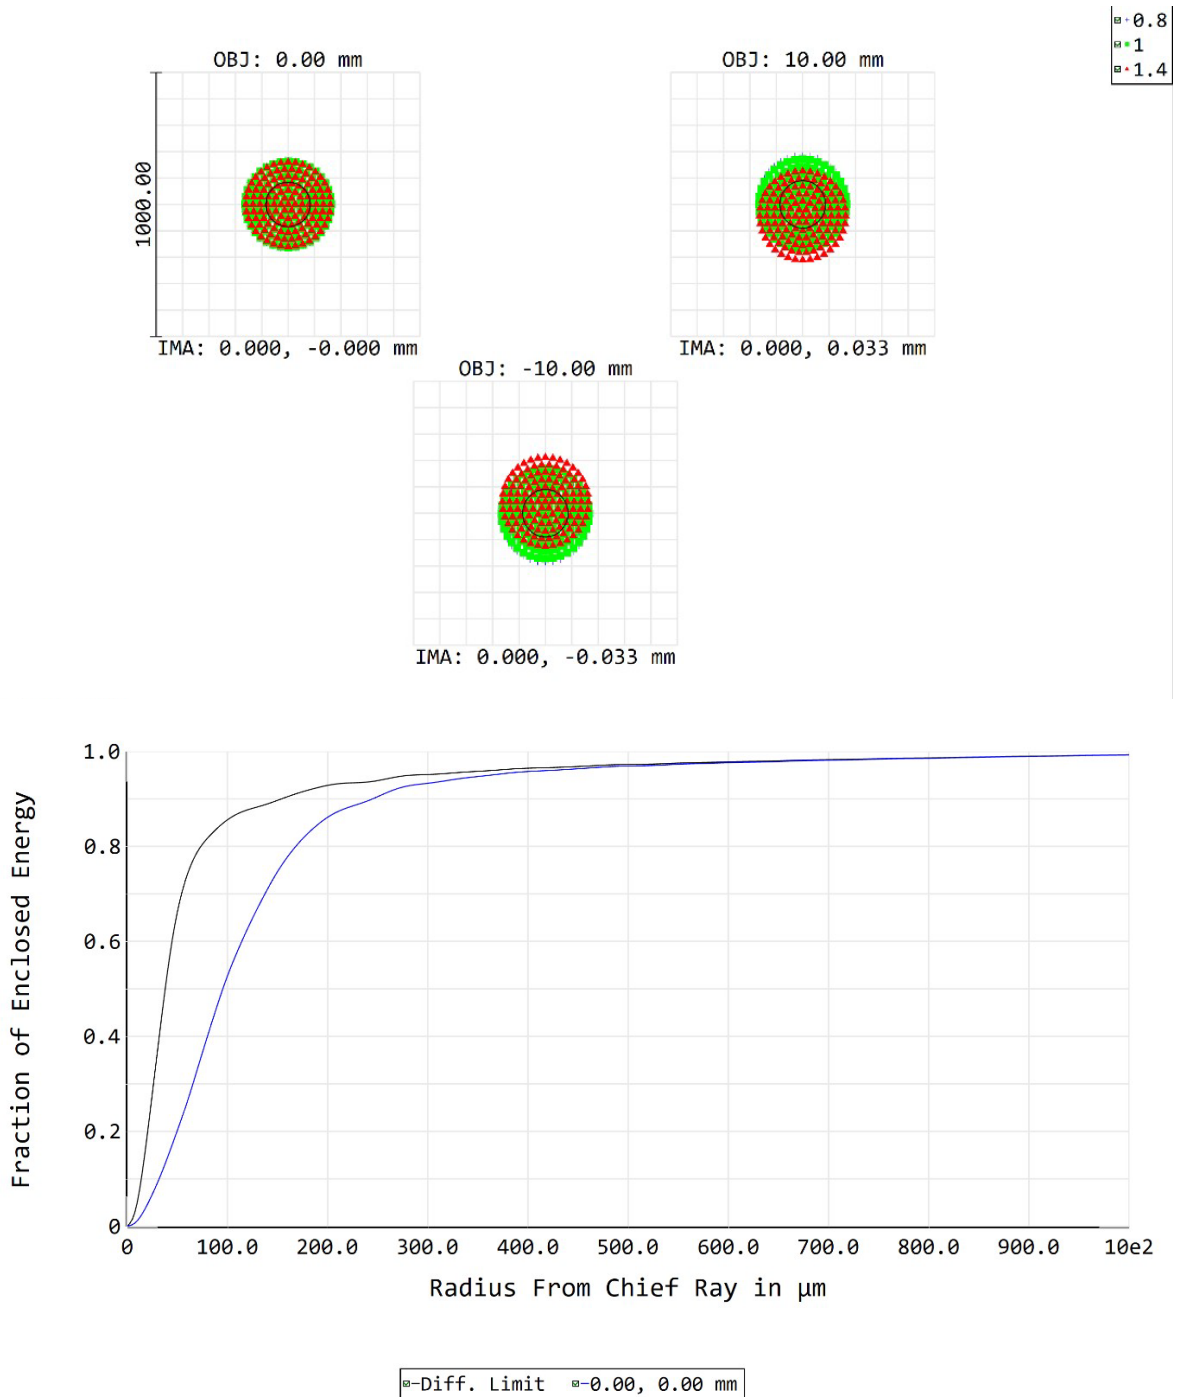


Figure A.2: Standard Zemax spot diagram of the simulated photo-diode arm and enquared energy plot, confirming that across DARKNESS's band we can expect the pupil images to be well within 100% enquared energy on the PD sensors.

the measured PD power into a counts per second per unit area value at the aperture just before the collimating lens we need to verify that the illumination is indeed uniform at that surface.

From simulation we expect the beam to be uniform to within $\sim 5\%$, assuming a perfectly uniform object plane (Figure A.3).

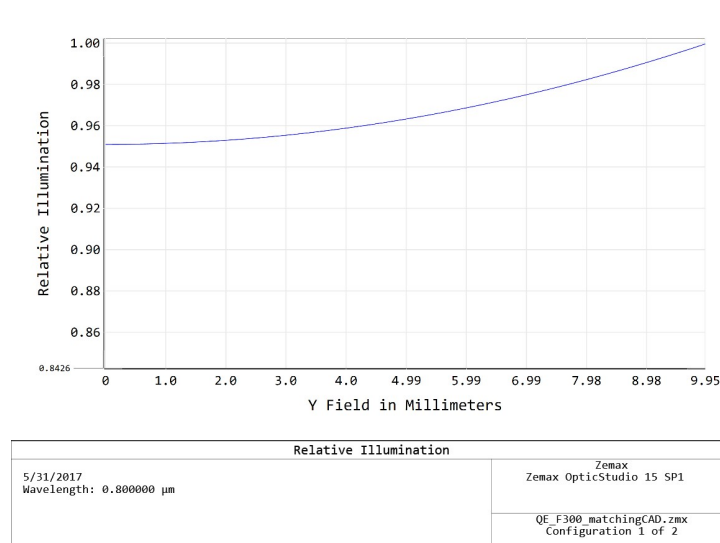


Figure A.3: Zemax simulation of relative intensity as a function of radius in the final image plane, showing expected uniformity to 5%.

We can experimentally confirm this simulation by stopping down the final image plane before the PD collimating lens using a variable aperture. If the image plane has uniform intensity we expect the measured PD power to scale linearly with collecting area. Figure A.4 shows the results of this measurement, and confirms that the intensity does seem to be constant across the image plane.

Finally, we must confirm that the measured count rate at DARKNESS’s MKID array is in a comfortable range (i.e. measurably above the dark count rate background but




Figure A.4: Measured uniformity of image plane using a variable aperture just before the PD collimating lens. Assuming the resulting pupil does indeed concentrate 100% of the collected light on the PD sensor, we see that the measured power does scale linearly with aperture area. Note: the plotted line is not a fit, just a line connecting each data point.

below our imposed ~ 2500 cps firmware buffer limit). Figure A.5 shows the measured lightcurve from a typical D-3 pixel during a recent DARKNESS QE measurement where we illuminate the array with a series of wavelengths across DARKNESS's band. The count rates at every wavelength fall nicely within our desired range. This confirms that our QE bed delivers the expected performance in both the PD and MKID arms.

A.3 Could we use a smaller $f/\#$?

The design constraints described above could be relaxed considerably if we didn't need a $f/300$ beam. This possibility is easily checked in Zemax by varying $f/\#$ and examining the ensquared energy at the MLA focus, as shown in Figure A.6. From this analysis it is




Figure A.5: Lightcurve showing the measured counts per second (cps) from a typical DARKNESS D-3 pixel during a full QE measurement. In this measurement we illuminate the array with a series of wavelengths from 0.8 to 1.4 μm in 50 nm steps. We see the lowest cps at 0.8 μm is still well above our background, and the highest cps at 1.2 μm is well below our saturation point.

clear (though probably not surprising) that different $f/\#$ will have an impact on the MLA performance at the 10% level. Considering that MLA misalignment is expected to be a considerable source of throughput loss, we decided it best to not complicate interpretation of the QE results with this additional model dependent uncertainty. Maintaining $f/\#$ as similar as possible to that at the telescope is key for ensuring a valid comparison to the count rate we measure on-sky.




Figure A.6: Comparison of ensquared energy of MLA focus for $f/100$ vs. $f/300$ designs. Even with perfect MLA focus and alignment we can expect the diffraction limited spot size to be comparable to the MKID inductor size, with nearly 10% variation in ensquared energy. This suggests that we must mimic the $f/\#$ expected at the telescope for the most accurate comparison to our on-sky measured throughput.

Bibliography

- Aime, C., & Soummer, R. 2004, *ApJL*, 612, L85
- Barman, T. S., Konopacky, Q. M., Macintosh, B., & Marois, C. 2015, *ApJ*, 804, 61
- Batalha, N. M., et al. 2013, *ApJS*, 204, 24
- Boccaletti, A., Labeyrie, A., & Ragazzoni, R. 1998, *A&AS*, 338, 106
- Boccaletti, A., Moutou, C., & Abe, L. 2000, *A&AS*, 141, 157
- Boccaletti, A., Moutou, C., Mouillet, D., Lagrange, A.-M., & Augereau, J.-C. 2001, *A&A*, 367, 371
- Bordé, P. J., & Traub, W. A. 2006, *ApJ*, 638, 488
- Bottom, M., Femenia, B., Huby, E., Mawet, D., Dekany, R., Milburn, J., & Serabyn, E. 2016a, in *Proc. SPIE*, Vol. 9909, *Adaptive Optics Systems V*, 990955
- Bottom, M., Wallace, J. K., Bartos, R. D., Shelton, J. C., & Serabyn, E. 2017, *MNRAS*, 464, 2937
- Bottom, M., et al. 2016b, *PASP*, 128, 075003
- Bowler, B. P. 2016, *PASP*, 128, 102001
- Cady, E., et al. 2013, in *Proc. SPIE*, Vol. 8864, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*
- Chauvin, G., Lagrange, A.-M., Dumas, C., Zuckerman, B., Mouillet, D., Song, I., Beuzit, J.-L., & Lowrance, P. 2004, *A&A*, 425, L29
- Chauvin, G., et al. 2017, *ArXiv e-prints*
- Claudi, R. U., et al. 2008, in *Proc. SPIE*, Vol. 7014, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 70143E

- Cook, T., et al. 2015, JATIS, 1, 044001
- Crepp, J. R., et al. 2011, ApJ, 729, 132
- Currie, T. 2016, ArXiv e-prints
- Day, P. K., LeDuc, H. G., Mazin, B. A., Vayonakis, A., & Zmuidzinas, J. 2003, Nature, 425, 817
- Dekany, R., et al. 2013, ApJ, 776, 130
- Dohlen, K., et al. 2006, in Proc. SPIE, Vol. 6269, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 62690Q
- Doyle, S., Mauskopf, P., Naylon, J., Porch, A., & Duncombe, C. 2008, Journal of Low Temperature Physics, 151, 530
- Fitzgerald, M. P., & Graham, J. R. 2006, ApJ, 637, 541
- Fressin, F., et al. 2013, ApJ, 766, 81
- Fulton, B. J., et al. 2017, ArXiv e-prints
- Gao, J. 2008, PhD thesis, California Institute of Technology
- Gao, J., et al. 2012, APL, 101, 142602
- Gladysz, S., & Christou, J. C. 2008, ApJ, 684, 1486
- . 2009, ApJ, 698, 28
- Gomez Gonzalez, C. A., et al. 2017, AJ, 154, 7
- Goodman, J. W. 2005, Introduction to Fourier optics
- Groff, T. D., et al. 2014, in Proc. SPIE, Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V, 91471W
- Hardy, J. W. 1998, Adaptive Optics for Astronomical Telescopes, 448
- Hayward, T. L., Brandl, B., Pirger, B., Blacken, C., Gull, G. E., Schoenwald, J., & Houck, J. R. 2001, PASP, 113, 105
- Hinkley, S., Oppenheimer, B. R., Brenner, D., Parry, I. R., Sivaramakrishnan, A., Soumerai, R., & King, D. 2008, in Proc. SPIE, Vol. 7015, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
- Hinkley, S., et al. 2011, PASP, 123, 74

- Howard, A. W., et al. 2012, *ApJS*, 201, 15
- Jovanovic, N., et al. 2015, *PASP*, 127, 890
- Konopacky, Q. M., Barman, T. S., Macintosh, B. A., & Marois, C. 2013, *Science*, 339, 1398
- Kozorezov, A. G., Volkov, A. F., Wigmore, J. K., Peacock, A., Poelaert, A., & den Hartog, R. 2000, *Phys. Rev. B*, 61, 11807
- Krist, J. E. 2007, in *Proc. SPIE*, Vol. 6675, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
- Kuzuhara, M., et al. 2013, *ApJ*, 774, 11
- Labeyrie, A. 1995, *A&A*, 298, 544
- Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, *ApJ*, 660, 770
- Larkin, J. E., et al. 2014, in *Proc. SPIE*, Vol. 9147, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 91471K
- Leduc, H. G., et al. 2010, *APL*, 97, 102509
- Lyon, R. G., & Clampin, M. 2012, *Optical Engineering*, 51, 011002
- Macintosh, B., Poyneer, L., Sivaramakrishnan, A., & Marois, C. 2005, in *Proc. SPIE*, Vol. 5903, *Astronomical Adaptive Optics Systems and Applications II*, ed. R. K. Tyson & M. Lloyd-Hart, 170–177
- Macintosh, B., et al. 2006, in *Proc. SPIE*, Vol. 6272, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
- Macintosh, B., et al. 2014, *PNAS*, 111, 12661
- . 2015, *Science*, 350, 64
- Madhusudhan, N., Burrows, A., & Currie, T. 2011, *ApJ*, 737, 34
- Males, J. R., et al. 2016, in *Proc. SPIE*, Vol. 9909, *Adaptive Optics Systems V*, 990952
- Maloney, P. R., et al. 2010, in *Proc. SPIE*, Vol. 7741, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 0
- Marley, M. S., Saumon, D., Cushing, M., Ackerman, A. S., Fortney, J. J., & Freedman, R. 2012, *ApJ*, 754, 135
- Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, 641, 556

- Marois, C., Lafrenière, D., Macintosh, B., & Doyon, R. 2008a, *ApJ*, 673, 647
- Marois, C., Macintosh, B., Barman, T., Zuckerman, B., Song, I., Patience, J., Lafrenière, D., & Doyon, R. 2008b, *Science*, 322, 1348
- Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., & Barman, T. 2010, *Nature*, 468, 1080
- Marsden, D., Mazin, B. A., Bumble, B., Meeker, S., O'Brien, K., McHugh, S., Strader, M., & Langman, E. 2012, in *Proc. SPIE*, Vol. 8453, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
- Martinache, F., et al. 2014, *PASP*, 126, 565
- Mawet, D., Pueyo, L., Carlotti, A., Mennesson, B., Serabyn, E., & Wallace, J. K. 2013, *ApJSal*, 209, 7
- Mawet, D., Pueyo, L., Moody, D., Krist, J., & Serabyn, E. 2010a, in *Proc. SPIE*, Vol. 7739, Modern Technologies in Space- and Ground-based Telescopes and Instrumentation, 773914
- Mawet, D., Riaud, P., Absil, O., & Surdej, J. 2005, *ApJ*, 633, 1191
- Mawet, D., Serabyn, E., Liewer, K., Burruss, R., Hickey, J., & Shemo, D. 2010b, *ApJ*, 709, 53
- Mawet, D., Serabyn, E., Liewer, K., Hanot, C., McEldowney, S., Shemo, D., & O'Brien, N. 2009, *Optics Express*, 17, 1902
- Mawet, D., Serabyn, E., Wallace, J. K., & Pueyo, L. 2011, *Optics Letters*, 36, 1506
- Mawet, D., et al. 2012, in *Proc. SPIE*, Vol. 8442, Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave, 844204
- Mawet, D., et al. 2014, *ApJ*, 792, 97
- Mayor, M., & Queloz, D. 1995, *Nature*, 378, 355
- Mazin, B. A. 2004, PhD thesis, California Institute of Technology
- Mazin, B. A., Bumble, B., Meeker, S. R., O'Brien, K., McHugh, S., & Langman, E. 2012, *Optics Express*, 20, 1503
- Mazin, B. A., et al. 2013, *PASP*, 125, 1348
- McElwain, M. W., et al. 2016, in *Proc. SPIE*, Vol. 9904, Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave, 99041A

- McHugh, S., Mazin, B. A., Serfass, B., Meeker, S., O'Brien, K., Duan, R., Raffanti, R., & Werthimer, D. 2012, *Review of Scientific Instruments*, 83, 044702
- Meshkat, T., Kenworthy, M. A., Quanz, S. P., & Amara, A. 2014, *ApJ*, 780, 17
- Milli, J., Mawet, D., Mouillet, D., Kasper, M., & Girard, J. H. 2016, *Astronomy at High Angular Resolution*, 439, 17
- Monfardini, A., et al. 2010, *A&A*, 521, A29
- Noroozian, O., Day, P. K., Eom, B. H., Leduc, H. G., & Zmuidzinas, J. 2012, *ArXiv e-prints*
- Noroozian, O., Gao, J., Zmuidzinas, J., LeDuc, H. G., & Mazin, B. A. 2009, in *American Institute of Physics Conference Series*, Vol. 1185, American Institute of Physics Conference Series, ed. B. Young, B. Cabrera, & A. Miller, 148–151
- Öberg, K. I., Murray-Clay, R., & Bergin, E. A. 2011, *ApJL*, 743, L16
- Oppenheimer, B. R., & Hinkley, S. 2009, *Annu. Rev. Astron. Astrophys.*, 47, 253
- Oppenheimer, B. R., et al. 2012, in *Proc. SPIE*, Vol. 8447, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
- Oppenheimer, B. R., et al. 2013, *ApJ*, 768, 24
- Perrin, M. D., Sivaramakrishnan, A., Makidon, R. B., Oppenheimer, B. R., & Graham, J. R. 2003, *ApJ*, 596, 702
- Perryman, M. 2011, *The Exoplanet Handbook*
- Poyneer, L. A., et al. 2014, in *Proc. SPIE*, Vol. 9148, Adaptive Optics Systems IV, 91480K
- Poyneer, L. A., et al. 2016, *Appl. Opt.*, 55, 323
- Ricci, L., et al. 2017, *AJ*, 154, 24
- Roquiny, P., Bodart, F., & Terwagne, G. 1999, *Surface and Coatings Technology*, 116, 278
- Ross, T. S. 2009, *Appl. Opt.*, 48, 1812
- Schlaerth, J. A., et al. 2010, in *Proc. SPIE*, Vol. 7741, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 9
- Seager, S. 2010, *Exoplanets*
- Seager, S., & Deming, D. 2010, *Annu. Rev. Astron. Astrophys.*, 48, 631

- Serabyn, E., Mawet, D., Wallace, J. K., Liewer, K., Trauger, J., Moody, D., & Kern, B. 2011, in Proc. SPIE, Vol. 8146, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
- Serabyn, E., Wallace, K., Troy, M., Mennesson, B., Haguenauer, P., Gappinger, R., & Burruss, R. 2007, ApJ, 658, 1386
- Sivaramakrishnan, A., Koresko, C. D., Makidon, R. B., Berkefeld, T., & Kuchner, M. J. 2001, ApJ, 552, 397
- Skemer, A. J., et al. 2012, ApJ, 753, 14
- Soummer, R., Aime, C., & Falloon, P. E. 2003, A&A, 397, 1161
- Soummer, R., Ferrari, A., Aime, C., & Jolissaint, L. 2007, ApJ, 669, 642
- Soummer, R., Pueyo, L., & Larkin, J. 2012, ApJL, 755, L28
- Soummer, R., Sivaramakrishnan, A., Pueyo, L., Macintosh, B., & Oppenheimer, B. R. 2011, ApJ, 729, 144
- Sparks, W. B., & Ford, H. C. 2002, ApJ, 578, 543
- Strader, M. J. 2016, PhD thesis, University of California Santa Barbara
- Strader, M. J., et al. 2013, ApJ, 779, L12
- . 2016, MNRAS, 459, 427
- Swenson, L. J., et al. 2012, in Proc. SPIE, Vol. 8452, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 0
- Szypryt, P., Mazin, B. A., Bumble, B., Leduc, H. G., & Baker, L. 2015, IEEE, 25
- Szypryt, P., Mazin, B. A., Ulbricht, G., Bumble, B., Meeker, S. R., Bockstiegel, C., & Walter, A. B. 2016, APL, 109, 151102
- Szypryt, P., et al. 2014, MNRAS, 439, 2765
- Traub, W. A., & Oppenheimer, B. R. 2010, Direct Imaging of Exoplanets, ed. S. Seager, 111–156
- van Eyken, J. C., et al. 2015, ApJS, 219, 14
- Vissers, M. R., Gao, J., Sandberg, M., Duff, S. M., Wisbey, D. S., Irwin, K. D., & Pappas, D. P. 2013, APL, 102, 232603
- Wolszczan, A., & Frail, D. A. 1992, Nature, 355, 145

Zmuidzinas, J. 2011, Annual Review of Condensed Matter Physics, 3, 169

Zombeck, M. 2007, Handbook of Space A&A: Third Edition (Cambridge University Press)

Zurlo, A., et al. 2014, A&A, 572, A85

—. 2016, A&A, 587, A57