

DARKNESS: A MICROWAVE KINETIC INDUCTANCE DETECTOR INTEGRAL FIELD SPECTROGRAPH FOR HIGH-CONTRAST ASTRONOMY

SETH R. MEEKER,^{1,2} BENJAMIN A. MAZIN,¹ ALEX B. WALTER,¹ PASCHAL STRADER,³ NEELAY FRUITWALA,¹
CLINT BOCKSTIEGEL,¹ PAUL SZYPRYT,¹ GERHARD ULBRICHT,⁴ GRÉGOIRE COIFFARD,¹ BRUCE BUMBLE,² GIULIA COLLURA,¹
RUPERT DODKINS,⁵ ISABEL LIPARTITO,¹ NICHOLAS ZOBRIST,¹ MICHAEL BOTTOM,² J. CHRIS SHELTON,² DIMITRI MAWET,⁶
JULIAN C. VAN EYKEN,⁷ GAUTAM VASISHT,² AND EUGENE SERABYN²

¹*Department of Physics, University of California Santa Barbara, Santa Barbara, California 93106, USA*

²*Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, California 91109, USA*

³*Dominican School of Philosophy and Theology, Berkeley, California 94708, USA*

⁴*School of Cosmic Physics, Dublin Institute of Advanced Studies, Dublin 4, Ireland*

⁵*Department of Physics, University of Oxford, Oxford OX1 3RH, UK*

⁶*Department of Astronomy, California Institute of Technology, 1200 E. California Blvd., Pasadena, California, 91125, USA*

⁷*IPAC, Mail Code 100-22, California Institute of Technology, 1200 E. California Blvd., Pasadena, California, 91125, USA*

ABSTRACT

We present DARKNESS (the DARK-speckle Near-infrared Energy-resolving Superconducting Spectrophotometer), the first of several planned integral field spectrographs to use optical/near-infrared Microwave Kinetic Inductance Detectors (MKIDs) for high-contrast imaging. The photon counting and simultaneous low-resolution spectroscopy provided by MKIDs will enable real-time speckle control techniques and post-processing speckle suppression at framerates capable of resolving the atmospheric speckles that currently limit high-contrast imaging from the ground. DARKNESS is now operational behind the PALM-3000 extreme adaptive optics system and the Stellar Double Coronagraph at Palomar Observatory. Here we describe the motivation, design, and characterization of the instrument, early on-sky results, and future prospects.

Keywords: instrumentation: detectors, adaptive optics, coronagraphs, spectrographs — planets and satellites: detection

1. INTRODUCTION & MOTIVATION

The field of exoplanet imaging is young and advancing rapidly. Early ground-based surveys doubled as testbeds for high-contrast instrumentation and algorithm development while putting strong constraints on the frequency of giant 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). More recently, a new class of dedicated near-IR high-contrast instruments have 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.¹ However, these instruments continue to find fewer giant exoplanets than predicted ([Macintosh et al. 2006](#)), suggesting a possible discrepancy between the planet mass function extrapolated from radial velocity surveys and the true giant exoplanet mass function ([Bowler 2016](#)). To fully understand the mass distribution of giant exoplanets and their formation mechanisms we must achieve deeper contrasts at smaller angular separations.

These specialized instruments have similar anatomies. First, an extreme Adaptive Optics (XAO) system corrects the atmosphere’s distortion of the incoming light.² After XAO the on-axis star light is removed with a coronagraph, leaving any off-axis light from a planet or disk mostly unaffected. Finally, the remaining light is sent to the science instrument backend for imaging, spectroscopy, and/or polarimetry depending on the particular system’s science instrument suite. While the exact capabilities vary from one system to another, each of the systems listed above features an integral field spectrograph (IFS) among their science instruments that returns a spectrum at each spatial element in the final image ([Hinkley et al. 2008; Larkin et al. 2014; Claudi et al. 2008; Groff et al. 2014](#)). With this configuration the majority of the overwhelming starlight can be suppressed, revealing nearby faint companions or disk systems with simultaneous imaging and spectroscopy ([Oppenheimer et al. 2013](#)).

Unfortunately, AO correction is never perfect — residual atmospheric aberrations and non-common path aberrations (NCPAs) between the wavefront sensor (WFS) and science arm result in some starlight escaping coronagraphic rejec-

tion, creating a pattern of coherent *speckles* in the final image that resemble faint companions. Quasi-static speckles — resulting from static aberrations such as from the instrument optics — have long decorrelation times on the order of minutes to hours ([Hinkley et al. 2007](#)). Several mitigation strategies exist for these static/quasi-static aberrations. NCPAs can be partially calibrated out before going on-sky with phase retrieval or interferometric calibration ([Burruss et al. 2010; Cady et al. 2013](#)), and long-lived speckles are typically subtracted in post-processing using a variety of differential imaging techniques like angular differential imaging (ADI; [Marois et al. 2006](#)) or spectral differential imaging (SDI or spectral deconvolution; [Sparks & Ford 2002; Crepp et al. 2011](#)). These differential imaging strategies utilize some form of image diversity to generate a reference point spread function (PSF) as an estimate for the quasi-static speckle pattern. However, speckles resulting from residual atmospheric aberrations have decorrelation times on the order of milliseconds to seconds ([Macintosh et al. 2005](#)), too fast to be resolved with conventional IFS exposures, and must simply be averaged into a smooth halo during long integrations.

The speckle noise that remains after post-processing imposes the state-of-the-art planet-star contrast limits across all ground-based high-contrast instruments. Overcoming this speckle barrier requires a wavefront correction scheme that employs a focal plane wavefront sensor (FPWFS) to capture NCPAs, operating at kHz frame rates to track even the atmospheric speckles ([Guyon 2005](#)). This application necessitates the development of fast, low-noise, near-infrared focal plane detectors. Here we introduce DARKNESS, 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.

1.1. Microwave Kinetic Inductance Detectors

MKIDs ([Day et al. 2003](#)) are an emerging low temperature detector (LTD) technology that, through microwave multiplexing techniques and inherently simple geometric design, enables 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; Szypryt et al. 2017](#)). 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.

MKIDs operate on the principle that incident photons will briefly change the surface impedance of a superconductor through the kinetic inductance effect, caused by the finite

¹ e.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.

² Here “extreme” means thousands of actuators performing the wavefront correction, compared to only hundreds in the AO systems that early surveys used

mass of the Cooper pairs that compose the supercurrent. To reverse the supercurrent direction in an AC field requires the extraction of stored Cooper pair kinetic energy, which ultimately manifests as an additional surface inductance that dominates the overall surface impedance. When an incident photon breaks Cooper pairs the superconductor's surface impedance changes, and this change can be measured by using the superconductor as a variable inductor in a resonant circuit. The change in surface impedance will shift the frequency of the resonance, which, in turn, causes a shift in the amplitude and phase of a microwave probe signal that is coupled to the circuit. By fabricating individual resonators with different resonant frequencies a single microwave transmission line can read out thousands of resonators simultaneously through frequency domain multiplexing (McHugh et al. 2012).

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 rapidly moves off-resonance (in ~ 1 μ s) upon absorbing the photon, then decays back to its steady state more slowly (in ~ 20 μ s) according to the quasiparticle lifetime. 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 numbers of Cooper pairs, creating pulses in phase with height proportional to photon energy, which provides the intrinsic energy resolution of MKIDs.

The capabilities of UVOIR MKIDs have been 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 Coudé focus at Palomar and Lick Observatories, featuring a 2024 pixel MKID array optimized for a 0.4 μ m to 1.1 μ 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; Collura et al. 2017). DARKNESS inherits significantly from its predecessor, ARCONS, especially in the MKID design (Section 2.2), readout electronics (Section 2.4), and analysis software (Section 4.1).

1.2. High-speed Speckle Correction Techniques

Here we mention a few speckle suppression techniques that will especially benefit from the advent of MKIDs for high-contrast applications.

1.2.1. Statistical Speckle Discrimination

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 mod-

ified Rician (MR) form given by

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

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 (Perrin et al. 2003; Goodman 2005; Aime & Soummer 2004; Fitzgerald & Graham 2006; Soummer et al. 2007). The ratio I_c/I_s essentially characterizes the skewness, and the mean and variance of I are:

$$\mu_I = I_c + I_s \quad (2)$$

$$\sigma_I^2 = I_s^2 + 2I_c I_s \quad (3)$$

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- σ " contrast curves. 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 underestimated false-alarm probabilities (Marois et al. 2008).

Previously, studies have relied on whitening of the statistics, whereby PSF subtraction and other post-processing removes the correlated noise from (quasi-) static speckles. This leaves the atmospheric speckles to average together, and since they decorrelate quickly 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, averaging of atmospheric speckles is exactly what we hope to avoid.

When imaging faster than the atmospheric speckle decorrelation time their MR statistics can be preserved, and this information can be used in post-processing to distinguish them from astrophysical sources. Methods like Dark-Speckle Imaging (DS; Labeyrie 1995; Boccaletti et al. 1998, 2001) and Stochastic Speckle Discrimination (SSD; Gladysz & Christou 2008, 2009) rely on the fact that speckle intensities and companion intensities are drawn from 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. To implement these techniques optimally requires low-noise near-IR detectors capable of \sim ms exposures, a technology that was previously unavailable (Boccaletti et al. 2001; Gladysz & Christou 2008), making them promising techniques to pursue with MKIDs.

1.2.2. Speckle Nulling

Speckle nulling (Bordé & Traub 2006; Martinache et al. 2014; Bottom et al. 2016a) is one of several techniques for

Table 1. DARKNESS Overview

Parameters	Values
MKID Array Materials	PtSi on Sapphire substrate w/Nb ground plane
Array Format	80×125 pixels
Pixel Pitch	150 μm
Plate Scale	19.8 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 cm)	66 × 33 × 60
Cryostat Weight	110 kg

performing focal plane wavefront control, but is most attractive for our purposes because it lends itself most easily to a fast 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 speckle’s phase, which can be determined with probe patterns on the DM, stepping the expected waveform through just a few 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. With a science camera capable of $\sim\text{kHz}$ frame rates this technique can even cancel out atmospheric speckles in real time.

2. DARKNESS OVERVIEW

DARKNESS is the first of several planned IFs (see also MEC (Meeker et al. 2015) and PICTURE-C (Cook et al. 2015)) built to demonstrate the potential of MKIDs for high-contrast astronomy. Its baseline design is for operation at Palomar Observatory with the PALM-3000 (P3K; Dekany et al. 2013) XAO system and the Stellar Double Coronagraph (SDC; Bottom et al. 2016b). Here we provide an overview of the instrument as it currently operates with the SDC, while highlighting the flexible aspects of its design that enable its future travel to other observatories (see Section 5). A summary of key instrument parameters is provided in Table 1.

2.1. Cryostat




Figure 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.

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 66 cm long × 33 cm wide × 60 cm 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 110 kg. A detailed accounting of the cryostat internals can be found in Meeker et al. (2015). We briefly summarize key features and measured performance here.

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 thermal radiation shielding at 77 K for ~ 40 hour hold time on a single fill. Nested inside the 77 K shield is a 24 liter LHe tank and another layer of thermal 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 per fill.

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.

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 hollow Vespel SCP-5050 rods (Figure 2 (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.

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

2.1.2. 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 heatsinking the outer conductor only 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 wt% Niobium/47 wt% 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 2 (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 available in Walter et al. (2018).

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 4–8 GHz Low Noise Factory High Electron Mobility Transistor (HEMT) amplifier with a noise temperature of 2 K. 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.

2.2. MKID Array

DARKNESS houses a 10,000 pixel (80x125) MKID array with 150 μm pixel pitch, optimized for a 0.8 μm to 1.4 μm bandwidth. DARKNESS pixels follow a similar lumped-element design to those used in ARCONS (Marsden et al. 2012; Mazin et al. 2013), and early array designs presented in Meeker et al. (2015) were made out of titanium nitride (TiN) on Silicon (Si) substrates and also followed a similar layout and fabrication process to ARCONS arrays. (Leduc et al. 2010; Mazin et al. 2012). The most recent DARKNESS arrays (those that produced the on-sky results presented in Section 4) follow similar design principles, but take advantage of recent work exploring platinum silicide (PtSi) as an alternative to TiN (Szypryt et al. 2016). Figure 3 shows the PtSi on Sapphire array (known as D-3) used on DARKNESS’s most recent observing run, mounted in its microwave package with accompanying microscope photos revealing details of the individual pixels. A full reporting of the design, fabrication, and performance of these PtSi arrays is presented in Szypryt et al. (2017). While resonator internal quality factor, Q_i , is high in single-layer test chips — comparable to the values regularly achieved with TiN — work is ongoing to preserve high quality through the full array fabrication process. Nonetheless, our current D-3 arrays demonstrate sensitivity and energy resolution on-par with or better than that seen in DARKNESS’s early TiN arrays (see Section 3), with dramatically improved fabrication uniformity (Szypryt et al. 2016) and photometric stability (hot-pixel behavior in our TiN ar-




Figure 2. (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. (Right) Close-up of the device mounting structure with magnetic shield and 1 K baffle removed.




Figure 3. (Left) A 10,000 MKID PtSi on Sapphire D-3 array mounted in its microwave package. (Center) Detail image of the array showing the CPW transmission line with bondpad for one feedline. (Right) Further detail of several MKID pixels. The densely meandered patches at the top of each pixel are the photosensitive inductors, and the large sparse sections are the interdigitated capacitors used to tune each MKID to a unique resonant frequency.

rays is quantified in [van Eyken et al. \(2015\)](#), and elimination of this behavior with PtSi will be investigated fully in a forthcoming manuscript).

2.3. Optics

2.3.1. Cryostat Optics

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 (MLA) that concentrates the light from the final image plane onto the photo-sensitive inductor of each pixel.

The 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 entrance window is 12.5 mm thick AR-coated fused Silica. The 77 K and 4 K windows are both N-BK7 glass, 10 mm and 20 mm thick, respectively. These cold windows are coated with a custom IR-blocking filter from Custom Scientific, and define the cutoffs of our observing band. Since these are reflective coatings, both windows are mounted at 3° relative to the incident beam to reduce ghosting. The transmission curve of a single filter is shown in Figure 4. A first iteration of these blocking filters was not as stringent, and




Figure 4. Theoretical transmission curve of one IR-blocking/bandpass filter, courtesy of Custom Scientific. The transmission of the thick N-BK7 substrate is not included. 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.

we were sensitive to a small, but noticeable population of 300 K blackbody photons where the filter transmission increased slightly around 2.7 μm . With the final filter curves presented here, when combined with the IR blocking ability of the N-BK7 substrates above 2.8 μm , we expect < 1 photons per pixel per second of 300 K blackbody leak within DARKNESS's bandpass and < 15 photons per pixel per second below 4 μm . The level of IR blocking provided by the filters obviates the need for a cold pupil stop.

2.3.2. Relay Optics

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 full optical train, including the Hale Telescope, P3K, the SDC, and DARKNESS, targeting Nyquist sampling of the diffraction beamwidth (λ/D) across our observing band. Resulting spot diagrams are shown in Figure 6.

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 5. 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"x4" 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 mir-

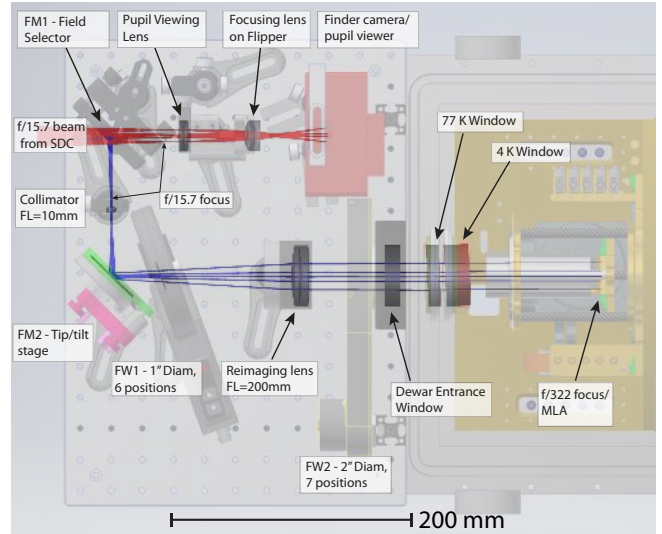


Figure 5. DARKNESS foreoptics layout when operating with the SDC.

ror 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 6-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 2.

This entire fore-optics enclosure is easily removable to 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.

2.3.3. 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 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

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

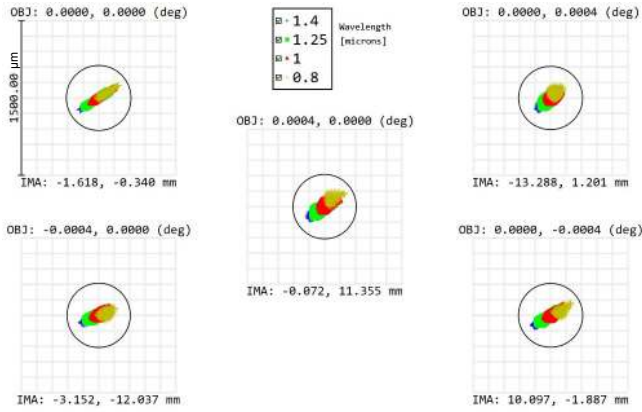


Figure 6. Spot diagrams from several fields roughly covering the extent of the DARKNESS FOV. Airy radius is shown as a black circle, calculated at $0.8 \mu\text{m}$. These diagrams show that even at the edges of the array the optical performance of the system will be diffraction limited (or rather, governed 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 expect Nyquist sampling (or better) of the diffraction limited PSF across our band with our $150 \mu\text{m}$ pixel pitch.

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).

2.3.4. Changes 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 (Bottom et al. 2016b; 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




Figure 7. Lyot Coronagraph optics in the SDC. (Left) Focal plane masks (Right) Lyot mask

use of a VVC at these wavelengths (though this may change very soon - see discussion in Section 5).

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, that is used to focus DARKNESS relative to the coronagraph focal plane. The FPM assembly attaches to the linear stage via magnetic mounts that are accurate to a few microns, allowing for easily reproducible alignment when swapping between our Lyot configuration and the SDC’s standard VVCs (Bottom et al. 2016b). 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 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 7 with selectable parameters listed in Table 2, 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 InGaAs quad-cell detector near its first focal plane that maintains very precise alignment of the target star on the coronagraph FPM. In SDC’s standard configuration — optimized for K-band observations — a dichroic sends all J-band light to this IR tracker. When operating with DARKNESS we replace the dichroic with a reflective 1450 nm shortpass filter. A similar swap is also made

Table 2. 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 _{ss} , Mauna Kea J, Diffuser (high grit), Diffuser (low grit), Open

in P3K where the typical 960 nm dichroic that splits WFS and science light would take part of DARKNESS's bandwidth at the blue end for the WFS. For DARKNESS observations we replace this with a 750 nm dichroic.

2.4. Readout Electronics

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, one 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 development 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 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 sum-

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

marize 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 8 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.

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 offset 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 image which it writes to disk every second.

2.4.1. Electronics Rack

DARKNESS's readout electronics, ADR magnet power supply, HEMT power supply, thermometry control, and GPS

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.

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

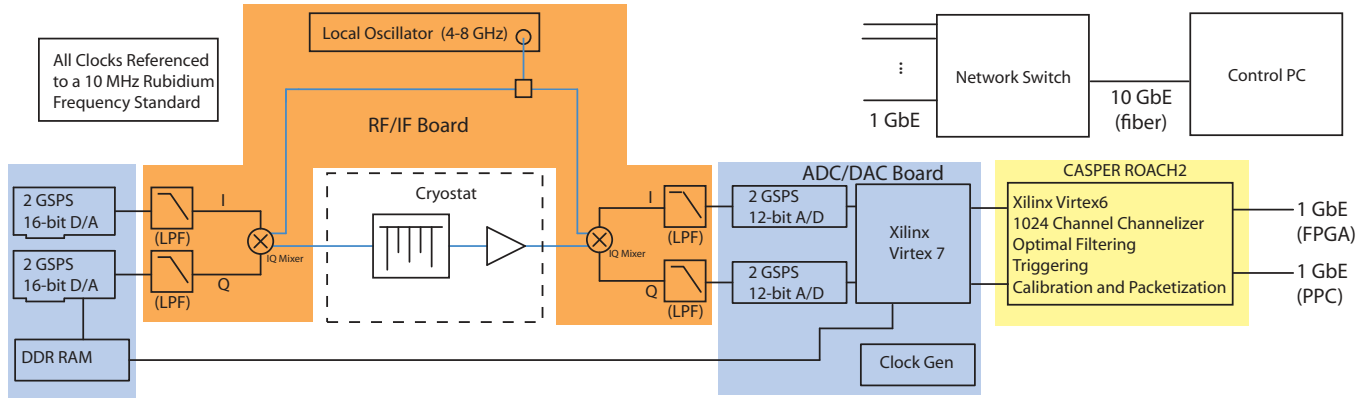


Figure 8. Block diagram of the second generation UVOIR MKID readout. The blue, orange, and yellow shaded functions take place on the ADC/DAC, IF, and ROACH2 boards, respectively.

reference time source are installed in an electronics rack that attaches to the outside of the Cassegrain cage during observations (see Figure 9 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. LABORATORY VERIFICATION

3.1. MKID Sensitivity and Energy Resolution

To confirm the sensitivity of the detectors we illuminate the array with a laser source in the lab while reading out a handful of resonators with a small-scale version of our readout. Figure 10 shows a couple events from an 808 nm laser. The phase offsets are near the desired value of $\sim 100^\circ$ for the blue end of our bandpass.

To calibrate the spectral resolution across the entire array we illuminate with a series of narrow band lasers at different wavelengths. For a given wavelength, λ , we estimate a pixel’s spectral resolution, $R = \lambda/\Delta\lambda$, by the full-width-half-maximum (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 11 shows the distribution of measured R from our state-of-the-art D-3 array for pixels that received full phase-to-wavelength solutions. As shown in [Szypryt et al. \(2017\)](#), the current PtSi D-3 arrays exhibit large fluctuations in Q_i with resonator frequency, resulting in a significant number of pixels having too low energy resolution to be calibrated in the manner described in Section 2.3.3.



Figure 9. DARKNESS electronics rack in lab at UCSB. This rack holds 5 readout cartridges with 10 sets of the readout boards described in Section 2.4, 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.

From this data we see that the median spectral resolution across our 0.8 to 1.4 μm band goes from ~ 7 to 5, roughly equivalent in width to the standard near-IR photometric filters at these wavelengths (i.e. Y, J). The current energy resolution is consistent with the measured phase noise in our resonators, dominated by HEMT noise ([Mazin et al. 2012](#))




Figure 10. Typical D-3 808 nm pulses.




Figure 11. Histogram of individual D-3 pixel energy resolutions.

and two-level system (TLS) noise.⁷ Further development is required to improve energy resolution through higher quality resonator fabrication (see [Szypryt et al. \(2017\)](#)) and lower noise-temperature amplifiers.

3.2. 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. 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

⁷ TLSs 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 et al. \(2008\)](#); [Gao \(2008\)](#).

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.

Figure 12 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 throughput 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 ([Szypryt et al. 2016](#)), and D-3 inductor fill factor. The biggest unknown in this theoretical curve is the spot size at the MLA focus. From Zemax diffraction simulations we assume 80% ensquared energy at the inductor, but this still does not account for all the lost flux. We attribute the remaining lost flux to MLA misalignment (roughly 20 μm of lateral MLA misalignment relative to the MKID inductor centers could account for a factor of 2 loss in QE) and are investigating strategies to improve our MLA mounting precision.

The above throughput calculation is for DARKNESS only. To assess the total capabilities of the instrument we estimate the throughputs of P3K and the SDC to be roughly 50%⁸ and 80% respectively ([M. Bottom 2018](#), private communication).

4. COMMISSIONING

DARKNESS traveled to Palomar Observatory for commissioning in July 2016, then again for ongoing commissioning and science verification in November 2016, April 2017, and October 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 a UVOIR MKID on sky, and April 2017 marked the first deployment of a PtSi MKID array on sky.

4.1. Data Reduction

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 ARCONS data is saved directly

⁸ [Dekany et al. \(2013\)](#) estimate photodetection efficiency at the P3K WFS as 24%, and this value was measured as 21% in 2015 ([R. Burruss](#), private communication). Assuming a conservative 92% average efficiency per optic, this corresponds to $\sim 50\%$ final transmission of the science beam at the dichroic.

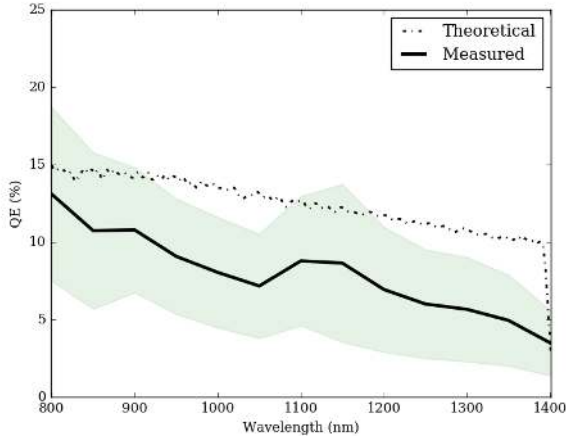


Figure 12. DARKNESS measured throughput vs. wavelength compared against theoretical prediction. The solid curve and shaded region shows the median and 1σ spread in throughput measured by all pixels. For the theoretical curves we’ve assumed flat 93% transmission through the 300 K window, applied the manufacturer 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.

to H5 files as a list of photon packets organized by pixel for a user defined total exposure time. The DARKNESS pipeline retains the use of the H5 format for its intermediate files, however, DARKNESS data is communicated to the control computer continuously every $500 \mu\text{s}$ and recorded to raw binary files every second as a timestream of photon packets. The first step in the reduction pipeline is to collect the desired 1-second binary files for a given target and format them into an indexed H5 file that can be funneled into the subsequent modules. After that, calibrations are carried out following the procedure in van Eyken et al. (2015) and we perform additional 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 deviation from the calibration procedures in van Eyken et al. (2015) is how we perform hot pixel masking. As mentioned in Section 2.2 PtSi on Sapphire resonators do not show the same random “switching” behavior seen in our TiN on Silicon arrays. Hot pixels are now mostly related to non-ideal readout parameters, manifesting as either con-




Figure 13. Median 1-second J-band image of 10 Uma, a spectroscopic binary with separation of $0.42''$ 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.

stantly 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 time to the desired short exposure time, and these frames are then dark subtracted, flat corrected, aligned, and smoothed with a Gaussian kernel having λ/D FWHM.

4.2. Plate Scale Verification

In November 2016 we observed 10 Ursae Majoris (10 Uma), a spectroscopic binary with separation of $0.42''$ at the time of observation, $V_{prim}=3.96$, and $\Delta V \approx 2$. Figure 13 shows a J-band image of the system with the coronagraph FPM installed, and inset shows FPM removed to reveal the primary. We fit a centroid to both objects, and using the known separation we calculate an on-sky plate scale of $19.8 \text{ mas per pixel}$. The diffraction limit λ/D for a 5.1 m telescope at $1.25 \mu\text{m}$ is roughly 50 mas , so we achieve Nyquist sampling at these wavelengths, but are slightly sub-Nyquist at $0.8 \mu\text{m}$.

4.3. Raw Contrast

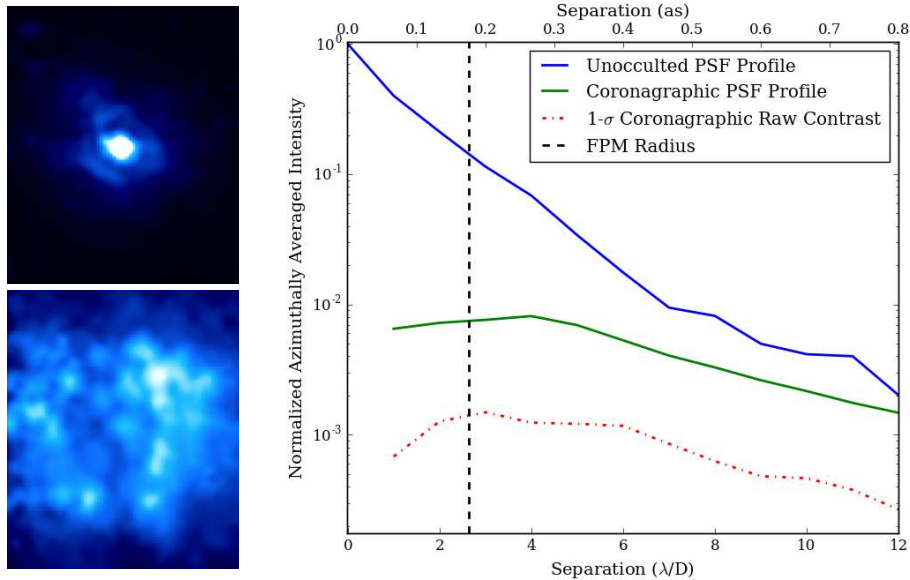


Figure 14. Raw J-band contrast (no post-processing speckle subtraction) as a function of angular separation with and without coronagraph FPM installed, measured on-sky with π Herculis ($J=0.79$). (Top Left) Median of 124 one-second frames of the unocculted, unsaturated PSF. (Bottom Left) Median of 1350 one-second frames of the coronagraphic PSF. Both images are *linear* scaled, but with different intensity minima and maxima. (Right) Azimuthally averaged PSF profiles and $1-\sigma$ contrast vs. separation curve measured in annuli at λ/D spacing and normalized by the unocculted, unsaturated PSF core. *Note:* This data was taken with the 20% undersized Lyot stop, which corresponds to $\lambda/D = 63$ mas in J-band.

To estimate the raw contrast achieved with the coronagraph we use observations of π Herculis ($J=0.79$) from April 2017, shown in Figure 14. We first observed the unocculted PSF using our OD2.5 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. We proceed using tools from VIP to perform basic aperture photometry on the unocculted PSF and estimate contrast as a function of angular separation. In the standard procedure, standard deviation at a given separation is estimated with a ring of λ/D sized apertures. These intensities are then normalized by the intensity of the unocculted, unsaturated PSF core measured with the same aperture. Figure 14 (Left) shows the un-occulted and coronagraphic PSFs. Figure 14 (Right) shows the resulting raw contrast curves (i.e. no post-processing to remove the static speckle pattern), plotting azimuthally averaged PSF profiles and $1-\sigma$ standard deviation measured in $1\lambda/D$ annuli.

We see the coronagraph provides little suppression as there is not significant diffraction to suppress in this low Strehl regime, but we are optimistic that this will improve with P3K’s impending upgrade.

4.4. Preliminary Speckle Statistics Results

Here we preview an investigation of speckle vs. companion statistics at very short timescales to demonstrate the promise of MKIDs. On October 4, 2017 we observed the multiple star system 32 Pegasi, of which the closest companion (32 Peg Ab, a stellar companion despite the lower-case “b” in the name) has 0.4” separation and $\Delta V=4.03$ (Mason et al. 2001).⁹ Figure 15 shows the median J-band image from a set of 1-second effective exposures with the coronagraph FPM installed. 32 Peg Ab is marked at location A, and by adding two sine-waves to the DM shape we have placed satellite speckles (copies of the stellar PSF) of similar intensity to 32 Peg Ab at locations B and C.

With this dataset we show proof-of-concept of the SSD technique. 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. With the power of MKIDs we can dynamically se-

⁹ Unfortunately the data available on Vizier for this object only provides the known visual magnitudes of the two components of the system, and only the spectral type of the primary, precluding an extrapolation of ΔJ .




Figure 15. Median 1-second J-band image of 32 Peg Ab (at location A) and satellite speckles generated by placing two sine wave shapes on the P3K DM (locations B, C). The white box demarcates the region shown in Figure 17.

lect this exposure time in post-processing. Macintosh et al. (2005) find that atmospheric speckles evolve on timescales related to the aperture clearing time, $\tau_0 = 0.6 D/\bar{v}$, where D is the telescope aperture diameter, and \bar{v} is the mean wind speed. With $D = 5.1$ m for Palomar’s Hale telescope, and a prevailing wind speed of 9 m/s at the time of observation, this lifetime is approximately 340 ms. Using a 40 second subset of the full 32 Peg dataset we bin the photon timestream to a conservative 20 ms effective exposures, and plot the resulting histograms of measured intensity at locations A, B, and C, shown in Figure 16.

Qualitatively the distribution shape from location A is noticeably different from those of the speckles. Taking inspiration from Fitzgerald & Graham (2006) we fit the histograms with MR functions to infer values of I_c and I_s at each location. We then use the ratio I_c/I_s as a proxy for each histogram’s skewness, and as a means to quantify the differences. Though all three distributions can be fit by MR curves, $I_c/I_s \approx 30$ at location A and ≈ 10 at locations B and C. If we make a map of I_c/I_s at every pixel (shown in Figure 17) the effect is clear. With a maximum I_c/I_s of 30.7 at the center of the companion PSF compared to a mean across the map of 5.8 with SD = 3.8, the companion stands out in this parameter despite being indistinguishable from the speckles in average intensity. A full

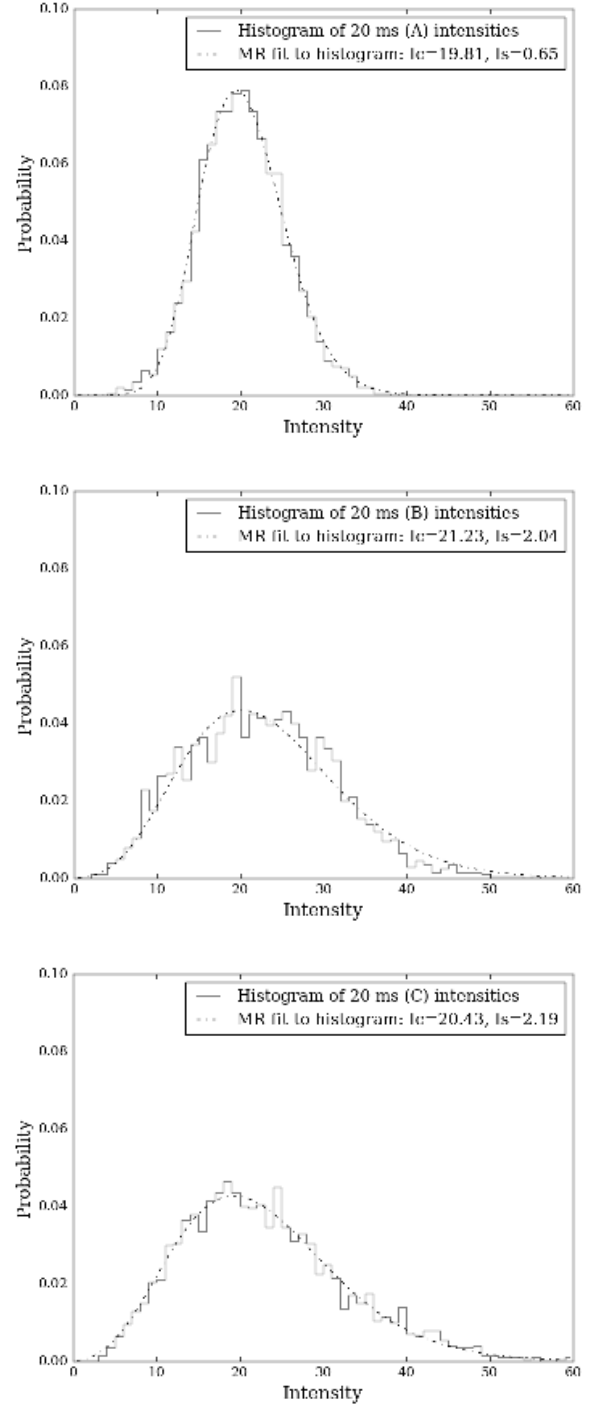


Figure 16. Intensity histograms from two-thousand 20 ms exposures for 32 Peg Ab (top panel) and satellite speckles (middle and bottom panels). Intensity units are photon counts in each aperture per 20 ms frame.




Figure 17. Map of I_c/I_s for each pixel.

analysis of the SNR gain from this test is beyond the scope of the current manuscript, but future work will investigate the performance of this and other statistical tests in maximizing companion SNR, with the ultimate goal of generalizing to use the photon arrival time statistics without the need for binning.

5. CONCLUSION AND ON-GOING WORK

We have presented the design and initial on-sky performance of DARKNESS, a MKID based IFS for high contrast imaging at Palomar. UVOIR MKID technology is still maturing rapidly — we have active research programs dedicated to improving the PtSi resonator quality factor, which will improve raw pixel yield and energy resolution, as well as exploring direct deposition of anti-reflection coating 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 long readout time of the science camera, PHARO (Hayward et al. 2001), resulting in slow conver-

gence and limited speckle control due to their decorrelation timescales being comparable to the control loop delays. 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 (Fruitwalla et al. in prep.).

The WFS camera and real-time controller (RTC) in 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 VVC for use with DARKNESS. The RTC upgrade will also open the door for kHz feedback from outside the main P3K control loop, enabling nulling of atmospheric speckles. With these upgrades in place we expect to overcome the current speckle noise barrier and achieve photon-noise limited contrasts.

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.

DARKNESS was funded by an NSF ATI grant AST-1308556. SRM and PS were supported throughout this work by NASA Office of the Chief Technologist’s Space Technology Research Fellowships (NSTRF). The research was carried out in part at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. We extend a special thanks to the amazing staff of Palomar Observatory. Their patience, knowledge, and support has been vital to the successful design and operation of DARKNESS.

Facilities: Hale (PALM-3000, Stellar Double Coronagraph)

REFERENCES

- Aime, C., & Soummer, R. 2004, *ApJL*, **612**, L85
- Boccaletti, A., Labeyrie, A., & Ragazzoni, R. 1998, *A&AS*, **338**, 106
- 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., et al. 2016a, in *Proc. SPIE*, Vol. 9909, *Adaptive Optics Systems V*, 990955
- Bottom, M., Shelton, J. C., Wallace, J. K., et al. 2016b, *PASP*, **128**, 075003
- Bowler, B. P. 2016, *PASP*, **128**, 102001
- Burruss, R. S., Serabyn, E., Mawet, D. P., et al. 2010, in *Proc. SPIE*, Vol. 7736, *Adaptive Optics Systems II*, 77365X
- Cady, E., Baranec, C., Beichman, C., et al. 2013, in , [arXiv:1309.6357 \[astro-ph.IM\]](https://arxiv.org/abs/1309.6357)
- Claudi, R. U., Turatto, M., Gratton, R. G., et al. 2008, in , 70143E
- Collura, G., Strader, P., Meeker, S. R., et al. 2017, *ApJ*, **850**, 65
- Cook, T., Cahoy, K., Chakrabarti, S., et al. 2015, *JATIS*, **1**, 044001
- Crepp, J. R., Pueyo, L., Brenner, D., et al. 2011, *ApJ*, **729**, 132
- Day, P. K., LeDuc, H. G., Mazin, B. A., Vayonakis, A., & Zmuidzinas, J. 2003, *Nature*, **425**, 817
- Dekany, R., Roberts, J., Burruss, R., et al. 2013, *ApJ*, **776**, 130
- Dohlen, K., Beuzit, J.-L., Feldt, M., et al. 2006, in
- Fitzgerald, M. P., & Graham, J. R. 2006, *ApJ*, **637**, 541
- Gao, J. 2008, PhD thesis, California Institute of Technology
- Gao, J., Daal, M., Vayonakis, A., et al. 2008, *Applied Physics Letters*, **92**, 152505
- Gladysz, S., & Christou, J. C. 2008, *ApJ*, **684**, 1486
- . 2009, *ApJ*, **698**, 28
- Gomez Gonzalez, C. A., Wertz, O., Absil, O., et al. 2017, *AJ*, **154**, 7
- Goodman, J. W. 2005, Introduction to Fourier optics
- Groff, T. D., Kasdin, N. J., Limbach, M. A., et al. 2014, in *Proc. SPIE*, Vol. 9147, *Ground-based and Airborne Instrumentation for Astronomy V*, 91471W
- Guyon, O. 2005, *ApJ*, **629**, 592
- Hayward, T. L., Brandl, B., Pirger, B., et al. 2001, *PASP*, **113**, 105
- Hinkley, S., Oppenheimer, B. R., Brenner, D., et al. 2008, in
- Hinkley, S., Oppenheimer, B. R., Soummer, R., et al. 2007, *ApJ*, **654**, 633
- Hinkley, S., Oppenheimer, B. R., Zimmerman, N., et al. 2011, *PASP*, **123**, 74
- Jovanovic, N., Martinache, F., Guyon, O., et al. 2015, *PASP*, **127**, 890
- Labeyrie, A. 1995, *A&A*, **298**, 544
- Larkin, J. E., Chilcote, J. K., Aliado, T., et al. 2014, in , 91471K
- Leduc, H. G., Bumble, B., Day, P. K., et al. 2010, *APL*, **97**, 102509
- 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
- Macintosh, B., Graham, J., Palmer, D., et al. 2006, in
- Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, *PNAS*, **111**, 12661
- Males, J. R., Close, L. M., Guyon, O., et al. 2016, in *Proc. SPIE*, Vol. 9909, *Adaptive Optics Systems V*, 990952
- Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, **641**, 556
- Marois, C., Lafrenière, D., Macintosh, B., & Doyon, R. 2008, *ApJ*, **673**, 647
- Marsden, D., Mazin, B. A., Bumble, B., et al. 2012, in
- Martinache, F., Guyon, O., Jovanovic, N., et al. 2014, *PASP*, **126**, 565
- Mason, B. D., Wycoff, G. L., Hartkopf, W. I., Douglass, G. G., & Worley, C. E. 2001, *AJ*, **122**, 3466
- Mawet, D., Riaud, P., Absil, O., & Surdej, J. 2005, *ApJ*, **633**, 1191
- Mawet, D., Serabyn, E., Wallace, J. K., & Pueyo, L. 2011, *Optics Letters*, **36**, 1506
- Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, *ApJ*, **792**, 97
- Mazin, B. A., Bumble, B., Meeker, S. R., et al. 2012, *Optics Express*, **20**, 1503
- Mazin, B. A., Meeker, S. R., Strader, M. J., et al. 2013, *PASP*, **125**, 1348
- McHugh, S., Mazin, B. A., Serfass, B., et al. 2012, *Review of Scientific Instruments*, **83**, 044702
- Meeker, S. R., Mazin, B. A., Jensen-Clem, R., et al. 2015, in
- Oppenheimer, B. R., Baranec, C., Beichman, C., et al. 2013, *ApJ*, **768**, 24
- Perrin, M. D., Sivaramakrishnan, A., Makidon, R. B., Oppenheimer, B. R., & Graham, J. R. 2003, *ApJ*, **596**, 702
- Soummer, R., Ferrari, A., Aime, C., & Jolissaint, L. 2007, *ApJ*, **669**, 642
- Sparks, W. B., & Ford, H. C. 2002, *ApJ*, **578**, 543
- Strader, M. J. 2016, PhD thesis, University of California Santa Barbara
- Strader, M. J., Johnson, M. D., Mazin, B. A., et al. 2013, *ApJ*, **779**, L12
- Strader, M. J., Archibald, A. M., Meeker, S. R., et al. 2016, *MNRAS*, **459**, 427
- Szypryt, P., Mazin, B. A., Ulbricht, G., et al. 2016, *Applied Physics Letters*, **109**, 151102
- Szypryt, P., Duggan, G. E., Mazin, B. A., et al. 2014, *MNRAS*, **439**, 2765
- Szypryt, P., Meeker, S. R., Coiffard, G., et al. 2017, *Optics Express*, **25**, 25894

van Eyken, J. C., Strader, M. J., Walter, A. B., et al. 2015, [ApJS](#), 219, 14

Walter, A. B., Bockstiegel, C., Mazin, B. A., & Daal, M. 2018, [IEEE Transactions on Applied Superconductivity](#), 28, 1

Zurlo, A., Vigan, A., Mesa, D., et al. 2014, [A&A](#), 572, A85