



JPL Document D-81778

TECHNOLOGY DEVELOPMENT FOR EXOPLANET MISSIONS

Technology Milestone Report

Phase-Induced Amplitude Apodization (PIAA) Technology Development, Milestone 1

Monochromatic Contrast Demonstration

Olivier Guyon (PI)

University of Arizona

Brian Kern (Co-I), Jet Propulsion Laboratory

Andreas Kuhnert, Albert Niessner, K. Balasubramanian

20 Feb 2014

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

© 2014. All rights reserved

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

APPROVALS

Released by



Olivier Guyon, Principal Investigator
University of Arizona

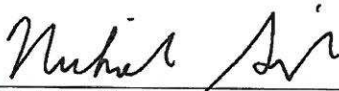
FEB 25, 2014

Approved by



Peter R. Lawson
Exoplanet Exploration Program Chief Technologist, JPL

2/25/2014



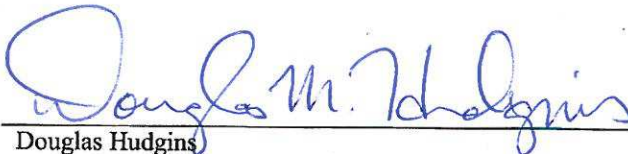
Nicholas Siegler
Exoplanet Exploration Program Technology Manager, JPL

2/25/14



Gary Blackwood
Exoplanet Exploration Program Manager, JPL

2/26/2014



Douglas Hudgins
Exoplanet Exploration Program Scientist, NASA HQ

3/6/2014



John Gagosian
Exoplanet Exploration Program Executive, NASA HQ

3/10/2014

LIST OF ACRONYMS

APLC	apodized pupil Lyot coronagraph
CCD	charge-coupled device
CLOWFS	coronagraphic low-order wavefront Sensor
DM	deformable mirror
HCIT	High Contrast Imaging Testbed
IWA	inner working angle
JPL	Jet Propulsion Laboratory
OAP	off-axis parabola
PIAA	phase-induced amplitude apodization
PSF	point spread function
SAT	Strategic Astrophysics Technology
TDEM	Technology Development for Exoplanet Missions
TPF-C	Terrestrial Planet Finder Coronagraph
TRL	technology readiness level

TABLE OF CONTENTS

1. ABSTRACT.....	1
2. INTRODUCTION	1
2.1. Definition and significance of milestone	1
2.2. Technical approach	2
3. EXPERIMENT CONFIGURATION AND TECHNIQUE	3
3.1. Optical layout.....	3
3.2. Off-axis distances.....	6
3.3. Image reduction and photometric normalization	8
3.4. Performance considerations	11
3.4.1. Throughput and inner working angle.....	11
3.4.2. Image spread	11
3.5 Wavefront control	13
3.5.1. Pupil-plane phase correction.....	13
3.5.2. Image-plane wavefront control	14
3.5.3. Dark hole creation.....	14
4. MILESTONE SPECIFICATION	15
4.1. Definitions.....	15
4.2. Success Criteria.....	18
4.3. Certification process	19
5. MILESTONE RESULTS.....	20
5.1. Description of milestone runs	20
5.2. Statistics of the milestone runs.....	21
5.3. Further analysis of intensity data: Histograms.....	21
5.4. Further analysis of intensity data: Radial scatter plots.....	23
5.5. Further analysis of intensity data: Coherent / incoherent light	24
5.6. Further analysis of intensity data: Effect of dark hole shape and size	26
5.7. Description of data analysis	27
6. CONCLUSIONS.....	27
Appendix 1. Occulter description and CLOWFS operation	27
Appendix 2. Calculation of λ/D_{sky}	28
Appendix 3. Model of testbed for wavefront control.....	29
Appendix 4. Applicability of Gaussian statistics	30
Appendix 5. Data reduction procedures.....	30
REFERENCES	31

1. ABSTRACT

This report describes the completion of Milestone 1 of the Phase-Induced Amplitude Apodization (PIAA) Technology Development experiment, funded under the NASA 2009 Technology Development for Exoplanet Missions grants. PIAA Milestone 1 is a demonstration of monochromatic coronagraphic starlight suppression at small inner working angles, at levels relevant for a space-based exoplanet imaging mission. The milestone is a requirement to “Demonstrate using Phase-Induced Amplitude Apodization a baseline contrast averaging 10^{-9} between a $2 \lambda/D$ inner working angle and a $4 \lambda/D$ outer working angle, in monochromatic light at a wavelength in the range of $400 \text{ nm} \leq \lambda \leq 900 \text{ nm}$.”

This document describes the scientific basis for targeting 10^{-9} contrasts in a coronagraph and the relevance of small inner working angles, the hardware comprising this coronagraph, the calibration and data analysis procedures, and the results of three Milestone 1 data runs satisfying the requirements. The intensity levels achieved in the three milestone data runs averaged between 5 and 8×10^{-10} , each of which is inconsistent with (*i.e.*, is better than) a 10^{-9} contrast with more than 90% confidence. A sample image is shown in Fig. 1.

2. INTRODUCTION

2.1. Definition and significance of milestone

In support of NASA’s Exoplanet Exploration Program and under the Strategic Astrophysics Technology (SAT), Technology Development for Exoplanet Missions (TDEM) funding program, our team has been advancing the Technology Readiness Level (TRL) for the PIAA starlight suppression technique. The purpose of the PIAA Milestone 1 (monochromatic contrast demonstration at $2 \lambda/D$), the methodology for computing the milestone metric, and the success criteria against which the milestone are evaluated, are described in detail in the PIAA Milestone 1 white paper [1].

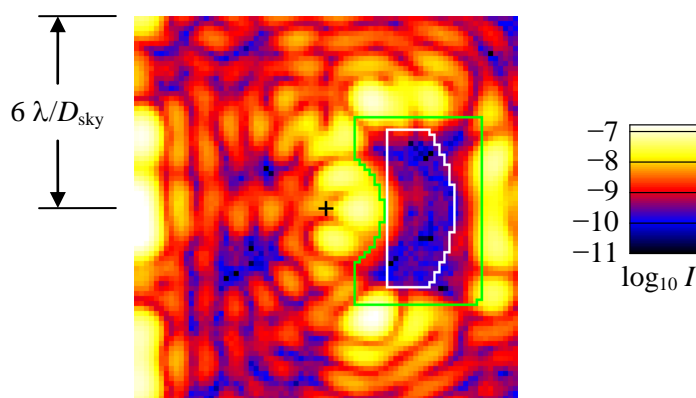


Fig. 1. Measured coronagraphic image showing dark hole with mean intensity $< 10^{-9}$. The field of view is $\pm 6 \lambda/D_{\text{sky}}$ on each side. Intensities are normalized to the occulter-out peak. The source center is marked with a small black cross, and the “scored” region border with inner edge at $x = 2 \lambda/D_{\text{sky}}$, outer radius $4 \lambda/D_{\text{sky}}$, is shown in white. The green line marks the edge of the occulter, the boundary between being fully opaque and fully transmitting. The mean scored intensity in this image is 5×10^{-10} .

The PIAA Milestone 1 definition, as stated in the white paper, is:

“Demonstrate using Phase-Induced Amplitude Apodization a baseline contrast averaging 10^{-9} between a $2 \lambda/D$ inner working angle and a $4 \lambda/D$ outer working angle, in monochromatic light at a wavelength in the range of $400 \text{ nm} \leq \lambda \leq 900 \text{ nm}$.”

This milestone addresses narrowband starlight suppression at small inner working angles with a PIAA coronagraph, a high-efficiency coronagraphy technique enabling high-contrast imaging at a small inner working angle ($\sim 2 \lambda/D$ for the configuration to be tested for this milestone) [2-11]. This milestone demonstration was performed on an optical table including key elements of any future high contrast imaging instrument for a future space mission (deformable mirror, science camera, coronagraph optics and active wavefront control), and is therefore of high relevance for direct imaging of exoplanets and disks from space.

Thanks to its high throughput, good angular resolution and small inner working angle (IWA), PIAA allows detection and characterization of habitable planets with a smaller telescope than was originally thought to be required [12-13]. While the theoretical potential of the PIAA has been widely recognized and estimated in mission concept studies, the realization of PIAA in future missions awaits proof of technical readiness traceable to space-based exoplanet mission science requirements and implementation constraints. PIAA Milestone 1 is an important step in this direction, demonstrating key advantages of the approach (small IWA, high contrast, with a high-throughput, high angular resolution architecture). A numerical comparison of a “generalized” PIAA architecture to other general architectures is presented in Ref. 9, while a comparison of coronagraph laboratory test results (including those presented here) appears in Ref. 14. These comparisons highlight the advantages, and motivation, for developing PIAA.

2.2. Technical approach

PIAA is a lossless beam apodization technique. Beam apodization is very useful in coronagraphy: an apodized pupil produces a high contrast image free of Airy rings. The conventional method to apodize the pupil is to introduce in the beam a mask which is fully transmissive in the center and becomes opaque at the edge of the pupil. With PIAA, the same apodized pupil is created by geometric redistribution of the light rather than selective absorption. This is achieved by aspheric optics (mirrors or lenses), as illustrated in Fig. 2.

The Airy rings produced with a conventional imaging telescope are due to diffraction originating from the sharp edges of the pupil. Pupil apodization, by creating a soft-edged pupil, can therefore greatly reduce these rings, and can be used either by itself or in combination with other coronagraph techniques (for example, the apodized pupil Lyot coronagraph which combines pupil apodization with Lyot coronagraphy). The conventional technique used to apodize the telescope beam is with an amplitude mask (continuous or binary) with variable transmission from the center to the edge of the pupil. This leads to a difficult compromise between reaching high contrast (which requires a strong apodization) and achieving high coronagraph throughput and good angular resolution (which both favor a weak apodization). In a conventional apodized pupil coronagraph, for example, reaching 10^{-10} contrast requires a strong apodization with a 10% throughput which degrades the telescope's angular resolution to approximately $2.5 \lambda/D$ (instead of $1 \lambda/D$) and does not allow high contrast imaging within approximately $4 \lambda/D$.




Fig. 2. PIAA uses aspheric optics to apodize a beam.

With PIAA optics, strong apodizations can be achieved with no loss in throughput or angular resolution, enabling high contrast imaging at small angular separation from the optical axis with almost no loss in efficiency. The PIAA's inner working angle at high contrast ranges from $0.64 \lambda/D$ for an aggressive PIAA design to $2 \lambda/D$ for a more conventional PIAA design (design choice depends on the goal contrast, manufacturing capabilities, ability to mitigate chromatic issues and angular size of the central source). PIAA does not absorb light, and it therefore preserves the sensitivity and angular resolution of the telescope. The introduction of a Lyot stop, as explained below, results in a small decrease in throughput. When implemented with mirrors, PIAA can be made to operate at high contrast over a wide spectral band.

The performance gain offered by PIAA for detection and characterization of exoplanets over other coronagraphs is quantified in Ref. 9. Compared to the more conventional coronagraphs which were considered for the Terrestrial Planet Finder Coronagraph (TPF-C), adopting the PIAA is equivalent to a $2\times$ to $3\times$ gain in telescope diameter. Some PIAA-based concepts such as PIAACMC allow even higher performance, with high contrast detection of exoplanets closer in than $2 \lambda/D$ [11].

3. EXPERIMENT CONFIGURATION AND TECHNIQUE

3.1. Optical layout

The JPL High Contrast Imaging Testbed (HCIT) PIAA testbed was assembled in 2009 and has been conducting high-contrast imaging experiments since [15-17]. It uses the first generation of PIAA mirrors fabricated by Axsys [8]. A schematic representation of the optical train is shown in Fig. 3. The optics are mounted on a $5' \times 8'$ table that is placed in a vacuum chamber. The layout of the optics on the table is shown in Fig. 4. The vacuum chamber typically reaches pressures around 10^{-6} Torr.

The optical train begins outside of the vacuum chamber, with a 5 mW, 807.5 nm fiber-pigtailed diode laser, frequency-stabilized with a distributed Bragg fiber grating (nominally a single longitudinal mode with linewidth 14 MHz). This light is carried in a single-mode fiber through a vacuum feedthrough, and illuminates a pinhole at the source (at the left of Fig. 3). The light then passes through the PIAA mirrors, creating an apodized pupil at the output of PIAA M2, with a nominally spherical wavefront. The PIAA




Fig. 3. Unfolded optical train to the science camera, not to scale. The marginal rays defined by the pupil stop (P1) are shown in blue, the marginal rays defined by the Lyot stop (P3) are shown in green. The planes conjugate to the source are easily identified as the locations where marginal rays cross; the occulter lies in a plane conjugate to the source. The planes conjugate to the pupil stop (P1) are labeled P0 – P4, and include PIAA M2 (P0), the DM (P2), and the Lyot stop (P3). The science camera is mounted on a translation stage that can reach a plane conjugate to the source (F4) and, by translating downstream, a plane conjugate to the pupil (P4), a range shown here in orange. The source and occulter can each be moved in 3 dimensions, and the pupil stop and Lyot stop can be removed. The CLOWFS system is not shown in this figure.

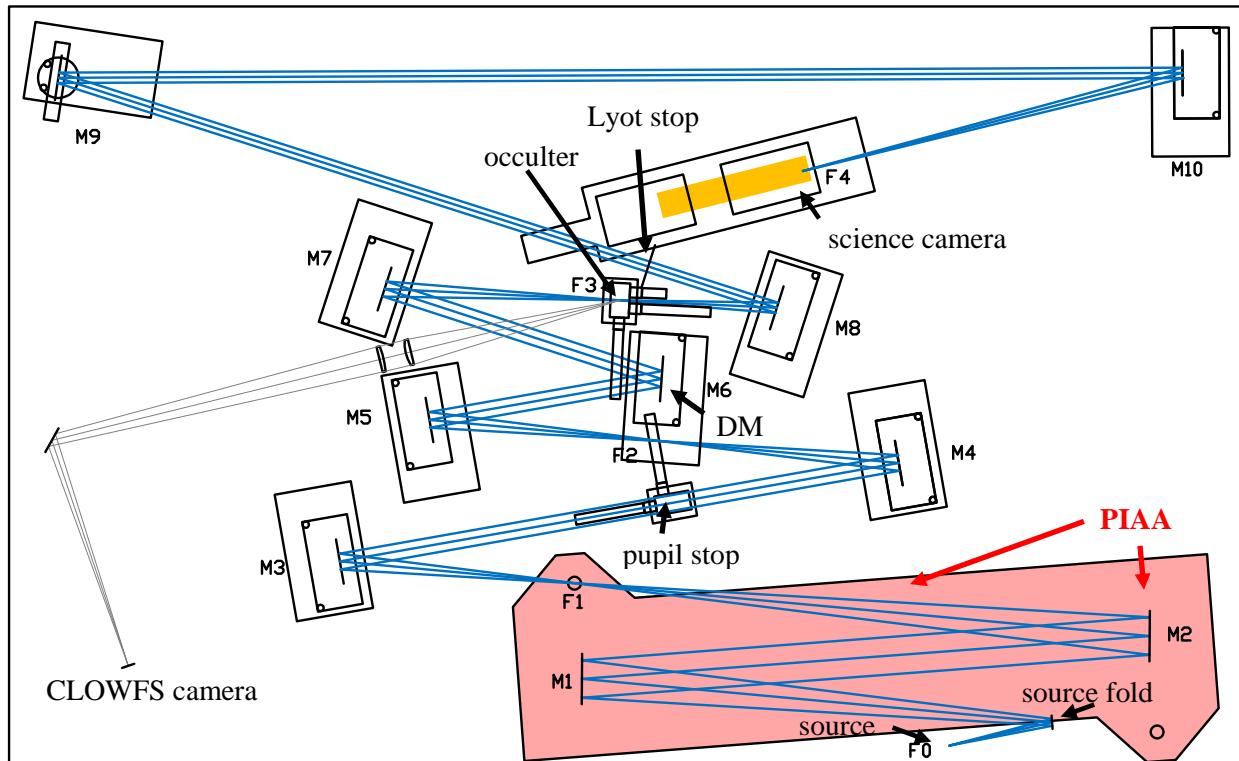


Fig. 4. Layout of optics on table, to scale. The source is at the bottom-right (F0). Light reflected off the occulter (gray rays) is relayed to the CLOWFS camera (bottom left). M9 is a flat mirror. The source fold mirror and DM have piezo actuators for tip-tilt (used by CLOWFS).

M2 is relayed to a pupil stop (a circular aperture), and then relayed again to the deformable mirror (DM), a 32x32-actuator Xinetics electrostrictive DM, with a 1 mm actuator pitch and a gold-coated continuous face sheet. The source is then imaged onto the occulter, recollimated and passed through a Lyot stop (another circular aperture), and reimaged onto the camera. The camera can translate farther downstream to a location conjugated to the pupil stop (and to PIAA M2, the DM, and the Lyot stop). A linear polarizer is in place at the camera, so all images represent a single linear polarization state.

At the occulter, much of the source light is reflected into the coronagraphic low-order wavefront sensor (CLOWFS) system, described in detail in PIAA TDEM Milestone 2 [18]. The occulter used here is not the same as the simple circular occulter described in Ref. 18, but is designed to transmit nothing in most of the regions of the point-spread function (PSF) that are not being controlled. A transmission image of this occulter (the light that passes through to the science camera) is shown in Fig. 5. The occulter is freestanding, in the sense that there is no substrate in the region where transmission occurs; the transmission of 1 is through a region of vacuum. The inner radius of the transmitting region is $1.7 \lambda/D_{\text{sky}}$, where λ/D_{sky} is defined in Section 3.2 below. The occulter, and its operation with the CLOWFS system, is discussed in more detail in Appendix 1.

While the PIAA mirror surfaces, which cause remapping of the entrance pupil and its lossless apodization, are nonconventional, the optical train downstream of the PIAA output is quite conventional. The illumination of the PIAA entrance pupil (the far-field of the source pinhole) is a nominally spherical wavefront with near-uniform amplitude across the aperture. Tracking the PIAA output apodization (*i.e.*, the wavefront at PIAA M2, identified in figures as P0) through the system, as shown in Fig. 6, with the occulter out it simply encounters two successive concentric circular aperture stops – at the pupil stop (P1), and the Lyot stop (P3). The pupil stop is included because the edges of the PIAA mirrors (the outer 10% of the mirror diameter) were not specified to produce a spherical wavefront (for ease of manufacturing), *i.e.*, the specified clear aperture is smaller than the mirror diameter, and the pupil stop removes the region of the pupil with poor wavefront quality. The Lyot stop is an integral part of the coronagraph architecture, acting in conjunction with the occulter (when it is in place) and the wavefront control.




Fig. 5. Occulter transmission. This transmission is measured at the science camera, ranging from 0 to 1, seen with the Lyot stop removed. The occulter is positioned so that the source image is centered on the red cross at the center of this image. The full image is $\pm 10 \lambda/D_{\text{sky}}$ on each side (wider than in Fig. 1). The dimensions of the mask parameters are shown in λ/D_{sky} units, defined in Sec. 3.2. The inner radius is $1.7 \lambda/D_{\text{sky}}$, while the farthest corners are located at $5.5 \lambda/D_{\text{sky}}$ (quadrature sum of $x = 4.7$ and $y = 2.9$).

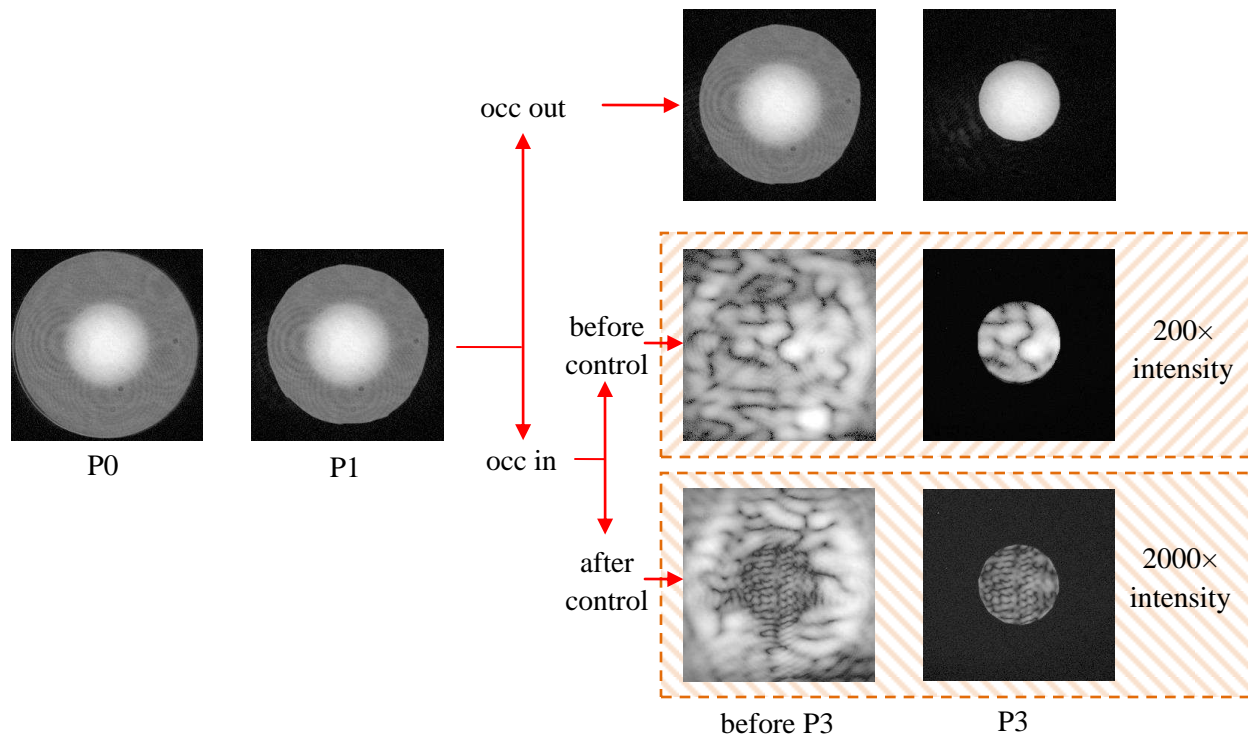


Fig. 6. Pupil images at PIAA M2 (P0), pupil stop (P1), upstream side of Lyot stop (before P3) and the Lyot stop (P3). Each image is log-scaled (with 2^{14} logarithmic range) intensity; the gray “uniform” intensity levels near the outside edge of P0 and P1 are approximately 1% of the center. The top row is observed without the occulter in place, so the P1 image is the same as before P3. With the occulter in, before wavefront control, enough light has been blocked that the images are shown here multiplied by 200 \times , and after wavefront control, by 2000 \times .

Fig. 7 shows a sequence of source images through the system, following the same progression through the optical train as in Fig. 6. What is readily apparent from the bottom rows of Figs. 6 and 7, is that the wavefront control does not attempt to reduce all of the light passing through the occulter, but allows a significant amount of light to pass the occulter (see bottom row, “occulter” column of Fig. 7), which is then eliminated by the Lyot stop (see bottom row of Fig. 6, bottom row, “F4” column of Fig. 7). In a categorical sense, after the lossless apodization created by the PIAA optics, the rest of the system is an apodized pupil Lyot coronagraph (APLC), although with an occulter that is not a simple circular occulter and a Lyot stop that is far smaller than simple diffraction near the edge would dictate. This is not an APLC designed for full on-axis extinction, however.

3.2. Off-axis distances

The PIAA coronagraph architecture presents an unusual complication in describing locations in the coronagraph field. In the Gaussian optics description (first-order, paraxial optics), the Lagrange invariant (or optical invariant) ensures that when comparing two conjugate planes, the lateral magnification corresponds to the change in focal ratios (geometrically defined by marginal rays). This makes it so that lateral distances in conjugate planes can be normalized by the local F/#. This is equally true if normalized to $(f\lambda/D)$, the diffraction limit (this is simply multiplying the focal ratio f/D by λ , which does not vary). In an infinite conjugate system (i.e., with a telescope front-end), it would be typical to describe distances

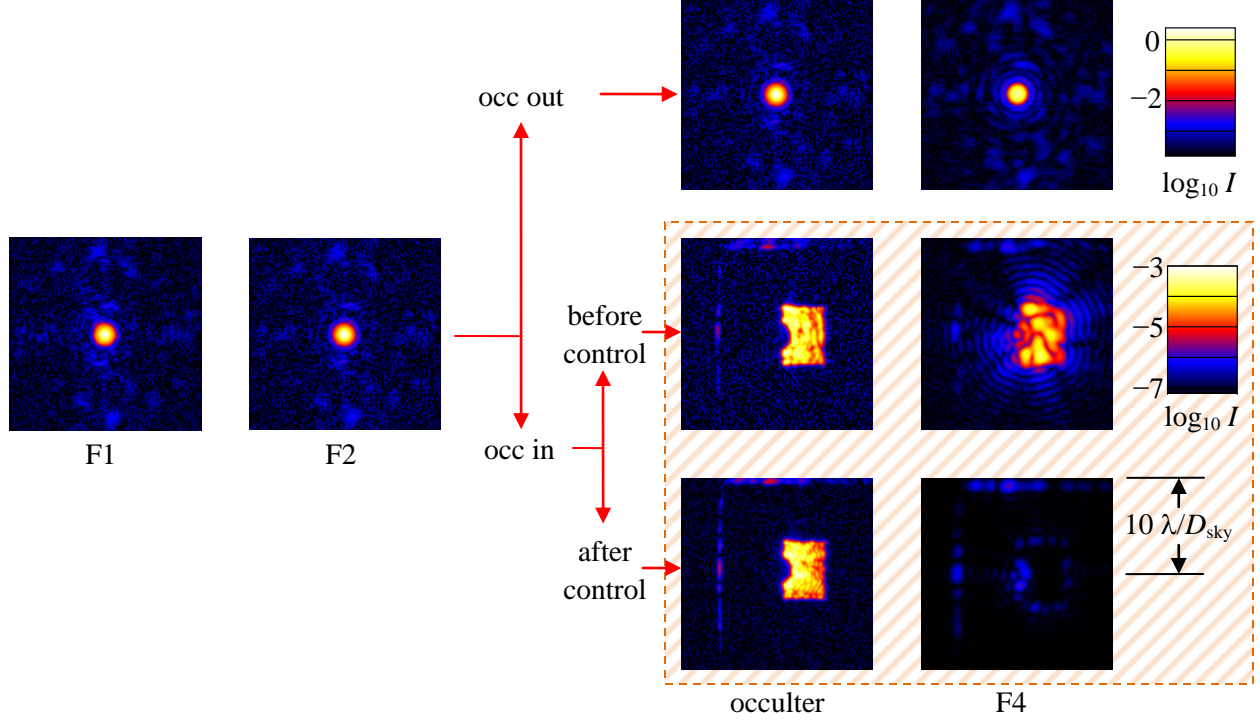


Fig. 7. Source images at F1 (after PIAA), F2 (after pupil stop), occulter, and F4 (at the camera, after the Lyot stop). Each image is log-scaled (with 2^{14} logarithmic range) intensity; two different colorbars are shown (lower light levels are seen with the occulter in, and correspond to the colorbar inside the hatched box). Intensities are normalized to the peak of the occulter-out, Lyot-in image (top right), so with the Lyot stop and the occulter out (see F1 and F2), the intensities rise above 1 (up to 2.5). The top row is observed without the occulter in place.

normalized to (λ/D) , an angular measure, with the implicit assumption that every image plane is normalized to $(f\lambda/D)$ with (f/D) defined by the marginal ray angles at that plane.

PIAA optics are not first-order optics, however, and off-axis distances in conjugate planes experience an additional magnification that is unrelated to changes in F/#. This is essentially due to the fact that the marginal rays pass through PIAA optics differently than the rays carrying the bulk of the intensity. Conjugate points are no longer located at the same multiples of (λ/D) or $(f\lambda/D)$.

To scale image planes in a manner relevant to the on-sky angles, all off-axis distances in this study will be quoted in units of λ/D_{sky} , which is defined by:

$$\lambda/D_{\text{sky}} = \begin{cases} f\lambda/D, & \text{image planes upstream of PIAA,} \\ Mf\lambda/D, & \text{image planes downstream of PIAA,} \end{cases}$$

where (f/D) is assumed to be $2/\tan \theta_m$, with θ_m the angle between marginal and chief rays, and the wavelength $\lambda = 807.5$ nm. The M introduced here is the remapping magnification, described below. Using these expressions, λ/D_{sky} represents positions on the sky as would be seen by a telescope feeding the coronagraph, appropriate to the locations where planets would be imaged.




Fig. 8. Evolution of image morphology for off-axis sources. The left-hand column is log-scaled, the right-hand column is linearly scaled. The three rows have the source offset by 0, 2, and $4 \lambda/D_{\text{sky}}$. The red cross marks the centroid of the light (to the nearest pixel). Each image is approximately $20 \lambda/D_{\text{sky}} \times 10 \lambda/D_{\text{sky}}$. All images were taken with the occulter out, and the Lyot stop in.

In addition to the remapping magnification experienced by an off-axis source, off-axis PSFs are aberrated by the PIAA optics (differently than an on-axis source). An example is shown in Fig. 8. Because the point-spread function is not shift-invariant, some metric must be chosen to represent the “position” of an aberrated image. The remapping magnification M is calculated using the centroid of the image of an off-axis source after passing through the PIAA optics. Used this way, the relationship between source location and image centroid is linear to $\ll 1\%$ over the distances relevant to this study.

The measurements that went into the calculation of λ/D_{sky} , measured in pixels, are detailed in Appendix 2, giving an answer of $1 \lambda/D_{\text{sky}} = 6.34 \text{ pix}$.

3.3. Image reduction and photometric normalization

As images are read off of the science camera, they are first bias- and dark-subtracted. The bias level is read from a region of charge-coupled device (CCD) overscan, in each image. The dark field calibration involves taking images in which the shutter is not activated, to produce a dark reference frame, which is then subtracted from the bias-subtracted CCD images. The mean dark level relevant to dark hole intensity measurements typically amounts to 10^{-10} , and the dark reference frames are measured to vary over month timescales by less than 10^{-11} , so errors in the dark subtraction are estimated to be less than 10^{-11} .

All intensity measurements relevant to the coronagraph are normalized to the occulter-out, Lyot-in intensity peak. The process of referencing coronagraphic images with intensities well below 10^{-9} , to occulter-out images with intensity at 1.0, is done in three stages, with each stage spanning approximately 10^3 in dynamic range. In each case, a 10^3 dynamic range in an individual scene is used to relate intensities to the next stage, with the exposure times increasing by 10^3 between stages. Fig. 9 gives a schematic representation of the 3-stage photometric normalization chain.




Fig. 9. Three stages of photometric normalization. (LEFT) Stage 1, remove occulter, add 10^{-3} speckle using DM, measure ratio of star (1.0) to speckle ($\sim 10^{-3}$). (CENTER) Stage 2, replace occulter, measure ratio of speckle ($\sim 10^{-3}$) to reference region (in green, $\sim 10^{-6}$). (RIGHT) Stage 3, remove speckle, scale entire image so that reference region brightness matches stage 2. Each stage involves intra-scene ratios with dynamic range $\sim 10^3$, with exposure times increasing by 10^3 between stages. Each panel is $\pm 10 \lambda/D_{\text{sky}}$ on each side, log-scaled. The Lyot stop is in for all stages. The exposure times listed are representative of the first two milestone runs.

The first two stages, taken together, establish the normalized intensity of a pre-defined photometric reference region (shown in green in the center and right panels of Fig. 9), near 10^{-6} . This sequence is repeated infrequently (before and after each milestone run, typically). With the reference region calibrated, every successive long-exposure coronagraph image (the right-hand panel of Fig. 9) is scaled in intensity so that the reference region matches its calibrated intensity. In this way, every long-exposure coronagraph image is “locally” normalized, making the final answers insensitive to source throughput variations or to calibration of exposure time linearity.

The reference region itself is chosen to capture light that “leaks” through a defect in the occulter. This is a low-transmission sliver reaching through the outer PSF (located $\sim 9 \lambda/D_{\text{sky}}$ above the source image). The source PSF landing in this location (see the right panel of Fig. 9) is affected very little by DM changes (the Nyquist condition on the DM corresponds to $6 \lambda/D_{\text{sky}}$). However, before correction (see center-right panel of Fig. 7), light passing through the main opening of the occulter is diffracted by the Lyot stop to overlap the reference region; after this light is removed by wavefront control (compare to bottom-right panel of Fig. 7), the light falling in the reference region decreases by $\sim 10\%$, and the reference region needs to be re-calibrated.

In practice, the reference region is calibrated before wavefront control to provide an initial approximate normalization for real-time use, then after wavefront control has completed, the reference region is re-calibrated using the final wavefront, and all previously acquired images are re-normalized to take on their final values. This post-processing normalization accommodates both changes in the PSF landing on the occulter leak itself, and changes in the light diffracted from the main opening of the occulter by the Lyot stop.

The relationship between calibration stages 1 and 2 is established by introducing a 10^{-3} speckle, done by adding a sinusoid to the DM. The calibration sequence relies on this speckle being the same in stage 1 as in stage 2. This is accomplished by taking a series of images for each stage, one image before changing the DM, one image with a sinusoid added, and a third image with the same sinusoid subtracted. Analyzing the three-image sequence allows the speckle intensity to be measured “AC-coupled,” i.e., with no influence from the underlying E -field (which is different with and without the occulter present).

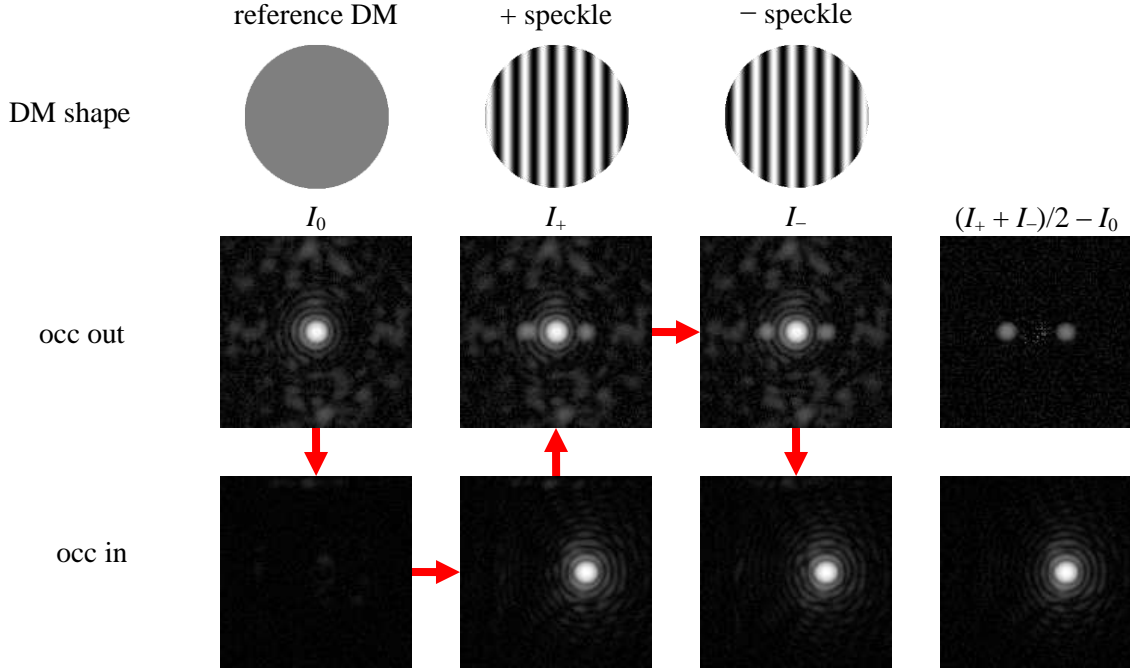


Fig. 10. Speckle ratio measurement sequence. Images are taken in the order directed by the red arrows, first by alternating occulter out / in (with exposure times different by $\sim \times 10^3$), then by modulating the DM to add a positive and negative speckle. The calculation in the right-hand column represents only the “AC-coupled” term, i.e., the speckles alone. The right-hand speckle of the “occ out” row is the same normalized brightness as that of the “occ in” row.

As is shown in Fig. 10, by measuring the I_0 , I_+ , and I_- images, corresponding to the addition of a “+” and a “-” speckle, the speckles can be isolated from I_0 by calculating $(I_+ + I_-)/2 - I_0$. The brightness of the occ out speckle is the same as the brightness of the occ in speckle (with $\ll 1\%$ effect from propagation through the occulter and Lyot stop), allowing stage 1 and stage 2 to be calibrated to each other (see Fig. 9).

In practice, a good test of the robustness of this photometric technique is to track a second reference region, normalized by the first, and see how its brightness changes. One obvious candidate for this second region is the “vertical” stripe to the left of the source, seen easily in the right panel of Fig. 9. This vertical stripe is located closer to the source center than the original reference region (at about $x = -5 \lambda/D_{\text{sky}}$), which makes it more susceptible to variability due to DM changes during wavefront control. The observed variations in the normalized intensity of this second reference region are typically less than 0.5% rms, once the large wavefront control DM changes are complete (i.e., once the contrast is $< 10^{-8}$). It can then be inferred that the variations in the primary photometric reference region are smaller than that, and are negligible.

The normalization determined by the photometric calibration routine is measured to vary by 1.6% rms over repeated calibrations. Taking the 1.6% initial calibration uncertainty and the 0.5% upper-limit on variability during the high-contrast portion of a run as independent of one another, and adding them in quadrature, an overall calibration uncertainty of 1.7% is used for the statistical confidence error analysis in Sections 4.1.8 and 5.1.

The important thing to note about this construction is that there is no reliance on stability of the light source, or on repeatability or linearity of the exposure times. During normal operation, every frame contains a measurement of the reference region, so that this normalization implicitly accommodates any combination of source and exposure time / throughput variability. The specific calibration sequence, taken in 3 steps (see Fig. 9), involves the addition of a dynamic speckle to transition from stage 1 to stage 2, rather than trying to span a 10^6 dynamic range in one transition (without any common intra-scene reference).

3.4. Performance considerations

Any starlight suppression system must be examined under several criteria, which should include residual intensity level (*i.e.*, contrast), optical bandwidth, inner working angle, throughput, and off-axis image quality. For this monochromatic experiment, the bandwidth is very narrow (1 MHz linewidth typical according to laser spec sheet, but not measured for this device— fractional bandwidth at 807.5 nm wavelength is $\sim 3 \times 10^{-9}$). The residual intensity level is the primary focus of this milestone, and is analyzed in great detail below. The throughput, inner working angle, and off-axis image quality are described further in this section.

3.4.1. Throughput and inner working angle

The occulter inner working angle is typically defined as the planet position at which the coronagraph throughput (total integrated light) is 0.5. This location can be different than the boundary of the dark hole region in which the wavefront is controlled, and different again than the edge beyond which intensities are reported for “scoring” purposes. Before quoting the occulter inner working angle, a convention must be specified for the normalization of throughput.

There are three normalizations of throughput that are readily generalized. These are:

1. Coronagraph transmission, normalized to peak occulter-in, Lyot-in transmission
2. Coronagraph transmission, normalized to unocculted, Lyot-in transmission
3. Coronagraph transmission, normalized to unocculted, no-Lyot transmission

All three are plotted in Fig. 11. Note that none of these three conventions considers reflectivities of mirrors or detector QE, or any other “total” throughput metric, as these numbers can change greatly unrelated to the coronagraph architecture – the PIAA remapping, occulting mask, and Lyot stop are the architecture choices under test, so the other factors are implicitly removed by these normalization conventions. Of these three normalization conventions, the first appears most commonly in other coronagraph studies. The third, on the other hand, measures how much light is absorbed the two “stops,” *i.e.*, the occulter and the Lyot stop, and is the only definition of coronagraph throughput that can meaningfully be compared across different architectures (neglecting reflective losses). All cases are for a single polarization state. Adopting these three conventions separately, the occulter inner working angle is 1.70, 1.82, and 1.88 λ/D_{sky} , respectively (the occulter inner edge is at 1.7 λ/D_{sky} from Section 3.1).

3.4.2. Image spread

To measure off-axis image spread, we present two metrics as a function of off-axis source position: the peak unocculted intensity, and the peak intensity through the occulter, both with the Lyot stop in. A plot




Fig. 11. Coronagraph throughput for off-axis source, with three different normalizations: (LEFT) normalized to peak throughput, (CENTER) normalized to unocculted, Lyot-in transmission, (RIGHT) normalized to unocculted, no-Lyot transmission. In each panel, the solid black line is a model cut along the x axis, the dashed black line is a model cut along a line to the corner of the mask open area, and the blue points are measured along x . The throughput crosses $T=0.5$ at 1.70 , 1.82 , and $1.88 \lambda/D_{\text{sky}}$.

of the modeled and measured values of these is in Fig. 12. The distortion caused by the PIAA remapping decreases the peak image intensity for off-axis sources, as the image is spread out over a larger area (see Fig. 8, related to the blue points in Fig. 12). This effect can be essentially eliminated by introducing an inverse PIAA set of optics at the end of the optical train, a technique that has been convincingly demonstrated on other PIAA testbeds [19].

The normalization used in Fig. 12 does not itself allow comparison of image quality to other coronagraph architectures. On-axis, the PIAA remapping greatly concentrates the PSF of the source. This can be




Fig. 12. Peak image brightness for off-axis source positions, normalized to on-axis peak unocculted brightness. The blue lines and points are unocculted peak brightnesses, the green lines and points are occulted peak brightnesses; the lines are modeled and the points are measured. The Lyot stop is in for all cases.

quantified by using an alternate metric, that of the intensity-weighted PSF area – the area a PSF occupies should be inversely related to the peak intensity. The intensity-weighted area of an Airy pattern (PSF of an unocculted, uniformly illuminated circular pupil, i.e., no PIAA optics) is $1.3 (\lambda/D)^2$, while the PIAA used here produces an unocculted PSF with area $0.5 (\lambda/D_{\text{sky}})^2$ with no Lyot stop, or $1.1 (\lambda/D_{\text{sky}})^2$ with the testbed Lyot stop in – even with the PIAA Lyot stop in place, the on-axis PIAA PSF is more concentrated than an Airy pattern. Compare this to typical band-limited Lyot coronagraphs, which may use a linear Lyot stops with $\varepsilon = 0.36$ for example [20], raising the unocculted PSF area to $2.3 (\lambda/D)^2$. This means that even with the off-axis image degradation that comes from not using inverse PIAA optics (by a factor of $I_{\text{max}}/I_{0,\text{max}} \sim 0.6$, for $2 < r/(\lambda/D_{\text{sky}}) < 4$ from Fig. 12), the image quality is far better than that of laboratory-tested band-limited Lyot PSFs – a band-limited Lyot coronagraph would need to use linear $\varepsilon = 0.25$ to have a similar image quality to the no-inverse-PIAA image quality demonstrated here at $r = 2 - 4 \lambda/D_{\text{sky}}$. The starting PSF used to normalize Fig. 12 is much sharper than standard Airy patterns, independently from throughput considerations (as in Section 3.4.1). Ref. 9 establishes a thorough comparison of different coronagraph architectures along similar lines, although the PIAA configuration examined in that reference is different than what is used here, as are the band-limited Lyot designs.

3.5 Wavefront control

3.5.1. Pupil-plane phase correction

With no a priori knowledge, when the DM is turned on the actuators are typically commanded to a uniform mid-range voltage. The exit pupil wavefront phase is then a sum (through the appropriate propagation) of errors arising from the PIAA optics, the off-axis parabolas (OAPs) and flats and their misalignment, and the initial DM surface shape. With the occulter and Lyot stop out, we perform a phase retrieval technique which gathers data at 10 different camera locations (in the yellow area of Figs. 3-4), then performs a combination of modified Gerchberg-Saxton iterative solutions and nonlinear optimization of parametrized pupil-plane phase modes. At the end of this procedure, we obtain a complex-valued estimate of wavefront phase in each of these planes, one of which is conjugate to the pupil.

With a technique for obtaining pupil-plane phase estimates, we iteratively change the DM shape to flatten the wavefront phase, and obtain a new phase estimate, until the wavefront is acceptably flat (at the spatial frequencies under control). Sample pupil-plane amplitude and phase maps are shown in Fig. 13, after flattening the phase. This sequence typically completes after 2 iterations from a uniform voltage.

It is appropriate to designate these steps a form of pupil-plane wavefront phase correction, which we call flattening. This flattening technique is repeated after every milestone run, so that the starting point before




Fig. 13. (LEFT) E -field amplitude and (RIGHT) phase after pupil-plane phase correction. Amplitude is shown in log-scale, where the outside region is about 10% the peak amplitude. Phase is shown with a 2.3 rad full-scale stretch. The flattening removes low-order phase errors; residual phase errors are above the actuator Nyquist frequency, which contribute light only outside the dark hole.

image-plane wavefront control is a reasonably well-corrected pupil-plane phase.

3.5.2. Image-plane wavefront control

With the pupil-plane phase flattened, the coronagraph field-of-view still contains light in the dark hole at the 5×10^{-5} level, verified by measurement on the testbed, modeling with the measured phase, and modeling with the flattest phase the DM could hypothetically produce (all give the same number to $< 10\%$). Through modeling, the effect of successive corrections on the dark hole average intensity could be derived:

1. 2×10^{-5} if the pupil plane phase could be made exactly uniform (i.e, flattened by a DM with an infinite number of infinitely small actuators)
2. 2×10^{-6} if the pupil plane amplitude errors could be removed. The existing pupil plane amplitude errors are likely due to surface errors propagating to the pupil plane, reflectivity variations, contamination of optics, etc., which cannot be eliminated using a phase-only correction in the pupil plane.

Note that even with perfect “correction,” the APLC-like configuration of the PIAA coronagraph does not provide starlight suppression beyond the 10^{-6} level. In this context, what is needed for high contrast is the ability to actively cancel light in the image plane, i.e., to control the PSF and force it to take on a shape that is different than the notional “perfect” wavefront. This is emphasized by describing the process as wavefront control rather than wavefront correction. The distinction here is that no constraints are put on the pupil-plane wavefront, only the image-plane intensity in the dark hole. The DM and wavefront control may drive the pupil-plane wavefront to be any arbitrary shape, as long as those changes improve the dark hole contrast. This process is explained in Ref. 21.

Beyond the initial phase flattening, the wavefront control is all performed with the camera at the source image. Controlling only the DM actuators, the relationship between differential actuator motions and the intensities observed in the dark hole is both nontrivial and ambiguous; it is nontrivial because of the propagation shown in Fig. 3, and is ambiguous because an intensity measurement does not describe the complex nature of the image-plane E -field necessary to determine how to create destructive interference.

The image-plane wavefront estimation and control used in this experiment is very similar to that described in Ref. 21. An estimate of the image-plane complex E -field is obtained, and a linearization of the effect of DM actuator motions on image-plane complex E -field is used to determine what DM setting will best eliminate the image-plane E -field. This estimation – control sequence is iterated as needed. The model describing the optical configuration and the effect of actuators on the image plane is described in Appendix 3.

3.5.3. Dark hole creation

The choice to include regions of the coronagraph image in the image-plane wavefront control or not defines a dark hole boundary. The milestone specification describes a region with $x > 2 \lambda/D$, $r < 4 \lambda/D$, which defines a “D” shape, in which the intensities are measured and reported. The wavefront control is typically specified to cover a larger region, because the edges of the controlled region are often not as dark as the central areas.




Fig. 14. Relative positions of occulter edge (green), scored region (white), and milestone “D” (black). The black line underlies the white line everywhere except near the top and bottom. Intensity image is from the same run as in Fig. 1.

The occulting mask used suffered an error during fabrication, in which it did not have as large a y-extent as desired (see Fig. 5). As such, the full “D” shape from the milestone specification was unavailable. The scored region used for this demonstration was as much of the “D” shape that was transmitted by the occulter, as shown in Fig 14, with a y-extent limited to $\pm 2.5 \lambda/D_{\text{sky}}$, whereas the full “D” would extend to $\pm 3.5 \lambda/D_{\text{sky}}$. A detailed description of the effect of this reduced dark hole size appears in Section 5.6. The conclusion of the detailed analysis is that if the correct occulter had been used, the change in the dark hole size from the full “D” to the white region of Fig. 14 would have a negligible impact on the performance. This analysis motivated the decision not to fabricate and re-test a new occulter, as opposed to proceeding to testing broadband contrast performance (the follow-on to this milestone).

The specification of the dark hole wavefront control region (as opposed to the scored region) was varied between milestone runs, with some attempt at determining its influence on the contrast performance. Those variations are not analyzed in this report, as all of the choices of wavefront control regions resulted in performance that exceeded the milestone requirements.

4. MILESTONE SPECIFICATION

The milestone definitions, success criteria, and certification process are described in the white paper, and are duplicated here in abridged form, with minor wording and notation changes. Any significant differences with respect to the original definitions are identified in italics, and separate notes are set in brackets. Note that the section numbers are referenced to this document, and so will not match individual section numbers in the white paper.

4.1. Definitions

4.1.1. “Raw” Image and “Calibrated” Image. A “raw” image is the 2D array of pixel values image obtained by reading out a CCD camera. A “calibrated” image is a raw image that has had background bias subtracted *and reference dark images subtracted*.

[The original wording included flat fielding, which was not performed. The average effect of flat fielding errors over the relevant areas should be negligible.]

4.1.2. “Scratch” is a DM setting *which delivers a flat pupil phase [see Section 3.5.1]. [This is not the same definition that was used for TPF-C HCIT milestones, which used a uniform voltage.]*

4.1.3. “Star”. The “star” is *an unresolved pinhole illuminated by an external laser [the original wording used a bare fiber tip, which was changed because the pinhole provided better performance].*

4.1.4. “Wavefront control iteration”. A “wavefront control iteration” is a measurement of the complex E -field in the dark zone followed by a DM correction intended to remove light in the dark zone.

4.1.5. The “*normalized intensity*” is a calibrated image normalized by the peak value of the unocculted source if it were measured under identical conditions (same illumination, exposure time, etc., *see Section 3.3*).

[*The original wording used the word “contrast”, which is defined differently for different coronagraph experiments; normalized intensity is specific and unambiguous.*]

4.1.6. The “*mean intensity*” is the average value of the normalized intensity over the dark zone adopted for the experiment. [*As in 4.1.5, this is normalized intensities instead of various definitions of contrast.*]

4.1.7. “*Data set*” is a contiguous set of at least four iterations and their associated mean intensity measurements. [*This definition was added for clarity in its use in other contexts. While the white paper mentions a possible duration for the data set, i.e., about an hour, only the number of iterations is used here. The elapsed time is detailed in Table 5.2.*]

4.1.8. “Statistical Confidence”. The *mean intensity* values (defined in 4.1.6) have associated noise estimates, with contributions from photon shot noise (Poisson statistics), detector read noise, *changes to the wavefront, environmental variations, and normalization uncertainty*. The goal of the statistical analysis is to determine the confidence with which the data demonstrate that the expected mean intensity is below 10^{-9} .

The notation used in this description begins with the terms contributing to an individual measurement,

$$m_i = \mu (1 + \epsilon) + \delta m_i,$$

where m_i is the mean intensity measured in iteration i , μ is the expected mean intensity, ϵ is the photometric normalization error, and δm_i is the iteration-specific variation for iteration i . Each of these quantities and their distributions are described below. The goal of this analysis is to test whether μ is below 10^{-9} , by measuring several m_i .

The determination of confidence intervals requires assumptions of probability distributions that underlie the measured mean intensity values. In this analysis, two distributions are assumed; the first is a distribution of systematic errors that take on a single value through an entire data set, and the second is the distribution of variations from one iteration to the next. Both distributions are assumed to be Gaussian distributions, but the estimation of variances of the two distributions differs.

The only significant contributor to the systematic errors is assumed to be the calibration of the photometric normalization (bias and dark subtraction, also applied uniformly through an entire data set, contribute negligible errors compared to photometric normalization), described in Section 3.3. This produces measurements that experience an additional multiplicative factor $(1 + \epsilon)$, where ϵ is normally distributed with mean 0 and variance σ_ϵ^2 . A single random variate ϵ is applied to all measurements in an individual data set (it appears without a subscript i). σ_ϵ^2 is estimated from repeated photometric calibration cycles (not from the measured mean intensities), described in Section 3.3.

Each iteration experiences variation in the measured mean intensity due to all changes between iterations, δm_i . These variations comprise photon shot noise, detector read noise, deliberate wavefront changes due to ongoing wavefront control, and environmental temporal changes to the testbed. The distribution of δm_i is assumed to be Gaussian (ignoring details of Poisson noise, for example), with mean 0 and variance σ_{each}^2 . The variance is estimated from the m_i samples themselves, so the sample mean of δm_i values over a data set comprising many iterations will follow a Student's- t distribution.

The sample mean of the mean intensities over a data set with n iterations, called here the data set mean, is denoted \bar{m} , defined by

$$\bar{m} = \sum_{i=1}^n m_i / n = \mu (1 + \varepsilon) + \sum_{i=1}^n \delta m_i / n$$

Note the degeneracy in the use of the word “mean,” which is discriminated by the “mean intensity” in a single iteration as compared to the “data set mean” which is the average of the “mean intensity” over several iterations. It is helpful here to define a data set mean variate,

$$\Delta m = \sum_{i=1}^n \delta m_i / n,$$

which then implies, by construction,

$$\bar{m} = \mu (1 + \varepsilon) + \Delta m.$$

The sample variance measured on the individual m_i measurements, σ_{each}^2 , is defined as,

$$\sigma_{\text{each}}^2 = \sum_{i=1}^n (m_i - \bar{m})^2 / (n - 1)$$

The estimated variance of the data set mean, σ_m^2 (the overbar on the m subscript is dropped for notational simplicity), is,

$$\sigma_m^2 = \sigma_{\text{each}}^2 / n$$

From these assumptions, the data set mean \bar{m} is a random variate with a Student's- t distribution for $(n - 1)$ degrees of freedom, about an expected value $\mu(1+\varepsilon)$ with variance σ_m^2 , or equivalently, Δm is distributed about expected an value 0 with variance σ_m^2 .

The total variance of the data set mean values about the expected mean μ can be calculated as the simple sum of the iteration-independent and the systematic variances (scaled appropriately),

$$\sigma^2 = \sigma_m^2 + \mu^2 \sigma_\varepsilon^2,$$

and while this is accurate, the confidence intervals about μ should be treated in a more careful manner to respect the uncertainty in σ_m^2 , namely by using the Student's- t distribution explicitly.

The probability distribution of measuring a data set mean \bar{m} , given the expected value μ and σ_ε^2 and σ_m^2 (estimated from individual m_i measurements), is a marginal projection of a joint probability distribution, which is the product of a Gaussian distribution in ε and a Student's- t distribution of Δm , and the 90% one-sided confidence interval has a lower limit C_{90} defined by,

$$0.90 = \int_{C_{90}}^{\infty} d\bar{m} \int_{-\infty}^{\infty} d\varepsilon' p_{\varepsilon}(\varepsilon') p_{\Delta m}(\bar{m} - \mu(1 + \varepsilon')),$$

where p_{ε} is normal, $N(0, \sigma_{\varepsilon}^2)$, and $p_{\Delta m}$ is the Student's- t probability density for $(n - 1)$ d.o.f. and variance σ_m^2 .

The formal definition of the milestone confidence requirement is that the hypothesis that the measured data set mean \bar{m} was drawn from the population with expected mean μ and the distributions described above, can be rejected with 90% confidence. Equivalently, the measured data set mean is below the lower limit of the 90% one-sided confidence interval for an expected mean intensity $\mu = 10^{-9}$. To simplify the expression for C_{90} , when σ_{ε} dominates σ_m/μ , then $\sigma \sim \mu\sigma_{\varepsilon}$, and

$$C_{90} \sim \mu - \sigma\sqrt{2} \operatorname{erf}^{-1}(2(0.90 - 0.5)) = 10^{-9} - 1.28\sigma.$$

When σ_m dominates $\mu\sigma_{\varepsilon}$, $\sigma \sim \sigma_m$, and for $n = 4$ (the minimum number of iterations in a data set),

$$C_{90} \sim 10^{-9} - 1.65\sigma.$$

For intermediate cases of $\sigma_m \sim \mu\sigma_{\varepsilon}$, the full integration must be performed, and the coefficient multiplying σ will vary between 1.28 and 1.65. Whatever value C_{90} comes to, the milestone statistical confidence requirement is,

$$\bar{m} < C_{90}$$

[This is a restatement of the description in the white paper, with careful attention to a precise hypothesis test under a frequentist (not Bayesian) interpretation, and its rejection with 90% confidence. The treatment of Student's- t distribution is also more precise; the final success criterion is mathematically identical to the statement in the white paper for the case of large n , but expressed in a frequentist context. A discussion of the applicability of Gaussian statistics appears in Appendix 4.]

4.2. Success Criteria

4.2.1. Illumination is monochromatic light in single or dual polarization at a wavelength in the range of $400 \text{ nm} < \lambda < 900 \text{ nm}$.

Rationale: This milestone is a monochromatic experiment to demonstrate feasibility of the approach at a wavelength in the science band of planned (or considered) NASA missions.

[As described in Sec. 3.1, the illumination was a monochromatic 807.5 nm laser, observed through a linear polarization analyzer at the science camera.]

4.2.2. A mean intensity metric of 10^{-9} or smaller shall be achieved in a 2 to 4 λ/D dark zone, as defined in Sec. 3.5.3.

Rationale: This provides evidence that the field is sufficiently dark (10^{-9} expected exozodi level) to be useful for searching planets, and test whether there is a fundamental limitation at the inner working angle.

[See Section 5.1, Table 5.1 for intensity, Section 5.6 for analysis of effect of dark hole shape]

4.2.3. Criterion 4.2.2, averaged over the data set, shall be met with a confidence of 90% or better, as defined in Sec. 4.1.8. Sufficient data must be taken to justify this statistical confidence.

Rationale: Assuming the intensity values have a Gaussian distribution about the mean intensity, this demonstrates a statistical confidence of 90% that the mean intensity goal has been reached.

[See Section 5.1, Table 5.1.]

4.2.4. Elements 4.2.1 – 4.2.3 must be satisfied on three separate occasions with a reset of the wavefront control system software (DM set to scratch) between each demonstration.

Rationale: This provides evidence of the repeatability of the contrast demonstration. The wavefront control system software reset between data sets ensures that the three data sets can be considered as independent and do not represent an unusually good configuration that cannot be reproduced. There is no time requirement for the demonstrations, other than the time required to meet the statistics stipulated in the success criteria. [*Shortened from the original.*]

[See Section 5.1, Table 5.1.]

4.3. Certification process

The milestone certification data package will contain the following explanations, charts, and data products.

4.3.1. A narrative report, including a discussion of how each element of the milestone was met, an explanation of each image or group of images, appropriate tables and summary charts, and a narrative summary of the overall milestone achievement.

4.3.2. Calibrated images of the coronagraph transmittance profile.

4.3.3. Calibrated images of the 3 sets of data, with appropriate numerical or color-coded or grayscale coded intensity values indicated, and with coordinate scales indicated in units of Airy distance (λ/D), all in demonstration of achieving the milestone elements.

4.3.4. A histogram of the intensity distribution of pixels in the dark field for each of the high contrast images in the data set, and for the combined data acquired in each data set.

4.3.5. A set of intensity measurement values for each of the 3 data sets.

4.3.6. A description of the residual components of the residual light in the dark zone: static coherent light, dynamic coherent light (due to time-variable pointing errors and wavefront changes too rapid to be fully corrected by the wavefront control loop) and incoherent light (ghosts, polarization leaks).

4.3.7. A step by step description of all data processing and analysis performed, along with source code and algorithm description. This will be provided in sufficient detail so an independent analysis of the raw data can be applied outside our team.

5. MILESTONE RESULTS

5.1. Description of milestone runs

The milestone results were split into three runs, with four iterations in each data set (see Sec. 4.1.7). Images from the three runs are presented in Fig. 15, and a statistical summary of the mean contrasts is shown in Table 1.

Each of the three runs was operated under different wavefront control parameters or with optical differences. The wavefront control is free to choose to control arbitrary regions in the image plane, and leave the rest uncontrolled. This controlled region does not necessarily align with the occulter transmitting region or the scored region (green and white regions in Fig. 15). For runs R1 and R3, the regions controlled nearly matched the transmitting region of the occulter. For run R2, the region controlled was much tighter to the scored region, i.e., the controlled region had smaller area than for runs




Fig. 15. Individual images from milestone runs R1 – R3. Each runs shows 4 images, as well as the average of the 4 images. Each image extends from -1 to $+6 \lambda/D_{\text{sky}}$ in x , $\pm 4 \lambda/D_{\text{sky}}$ in y . The black cross is the center of the source image, the green box is the edge of the occulter, the white box is the scored region of the dark hole.

Run	$m_1 \times 10^{10}$	$m_2 \times 10^{10}$	$m_3 \times 10^{10}$	$m_4 \times 10^{10}$	$\bar{m} \times 10^{10}$	$\sigma_{\text{each}} \times 10^{10}$	$\sigma_m \times 10^{10}$	$\sigma \times 10^{10}$	$C_{90} \times 10^{10}$
R1	8.19	8.87	9.43	9.62	9.03	0.64	0.32	0.36	9.43
R2	6.26	6.14	6.66	6.01	6.27	0.28	0.14	0.22	9.67
R3	5.22	5.86	5.06	6.75	5.72	0.77	0.39	0.43	9.32

Table 1. Statistics of individual iterations and averages for runs R1-R3. The two green columns represent the milestone requirement, $\bar{m} < C_{90}$, which can be quickly verified in each case.

R1 and R3. As a consequence, for example, the right-hand corners (top and bottom) of the transmitting region of the occulter are bright in R2 (but outside of the scored region).

For run R3, new optics were installed at the source head, improving the throughput from the output of the optical fiber to the pinhole that leads into the coronagraph (at “F0” in Fig. 4) by a factor of 500×. For runs R1 and R2, the exposure times were kept short enough to be time-efficient, but given the low throughput they did not produce deep exposures of the dark hole. For run R3, with the increased throughput, reasonable exposure times produced well-exposed dark holes. The reduced measurement noise in R3 can be readily seen in Fig. 15. Additionally, more care was taken in R3 to modulate the DM solutions to reduce quantization noise in the dark hole.

The differences in the data set means between the three runs may have a number of underlying causes. The difference between R1 and R2 may be due to the size of the controlled region (smaller region in R2, lower intensities). The difference between R1 and R3 may be due to the photon counts (more photons in R3, lower normalized intensities). These questions have not been explored.

5.2. Statistics of the milestone runs

Table 1 contains all of the statistical analyses that were required for the milestone. The means of the intensities in the dark hole are listed for each iteration, allowing determination of \bar{m} and σ_{each} , as defined in Section 4.1.8, which then determines σ_m . Knowing σ_m from the four intensity measurements and σ_s from photometric calibration (see 4.1.8) yields σ , and determines C_{90} . In each case, the milestone requirement is that $\bar{m} < C_{90}$, which is readily verified, with comfortable margins. Formally, this is the statement that the data reject with 90% confidence the hypothesis that the true mean intensity μ was 10^{-9} .

5.3. Further analysis of intensity data: Histograms

The intensity data is analyzed further in a handful of ways. The individual intensities per pixel are histogrammed and plotted in Fig. 16, for individual iterations and for the mean of all iterations in a given run. Because the full range of data span a large interval of intensities, the histogram bin edges are not spaced uniformly from the minimum to maximum values.

For milestone runs R1 and R2, with low flux levels, the measurement noise is high enough to generate a large number of low (even negative) values, which accumulate in the lowest intensity bin in Fig. 16. When the four iterations in each run are averaged, that noise decreases to a level where relatively few pixels have values below 2×10^{-10} . As a rough expectation, the flux levels in R1 and R2 are such that a pixel with intensity 6×10^{-10} would see read noise at 7×10^{-10} and Poisson noise at 3×10^{-10} , for a total measurement noise level per pixel of 8×10^{-10} , *i.e.*, measurement SNR per pixel less than 1 (see the white noise level in the dark hole in Fig. 15). Averaged over 349 pixels in the dark hole, the contribution of measurement noise to σ_{each} is expected to be 4×10^{-11} . The observed values of σ_{each} reported in Table 1 are estimated from 4 m_i values, which produce a 1-sigma uncertainty in the measurement noise contribution ranging from 2.2×10^{-11} to 5.4×10^{-11} . This compares well with the observed value $\sigma_{\text{each}} = 2.8 \times 10^{-11}$ for R2 (see Table 1), assuming that the contribution of wavefront control and dynamic environmental fluctuations were small. The use of Student’s-*t* distributions in the calculation of C_{90} incorporates this uncertainty in σ , by construction.

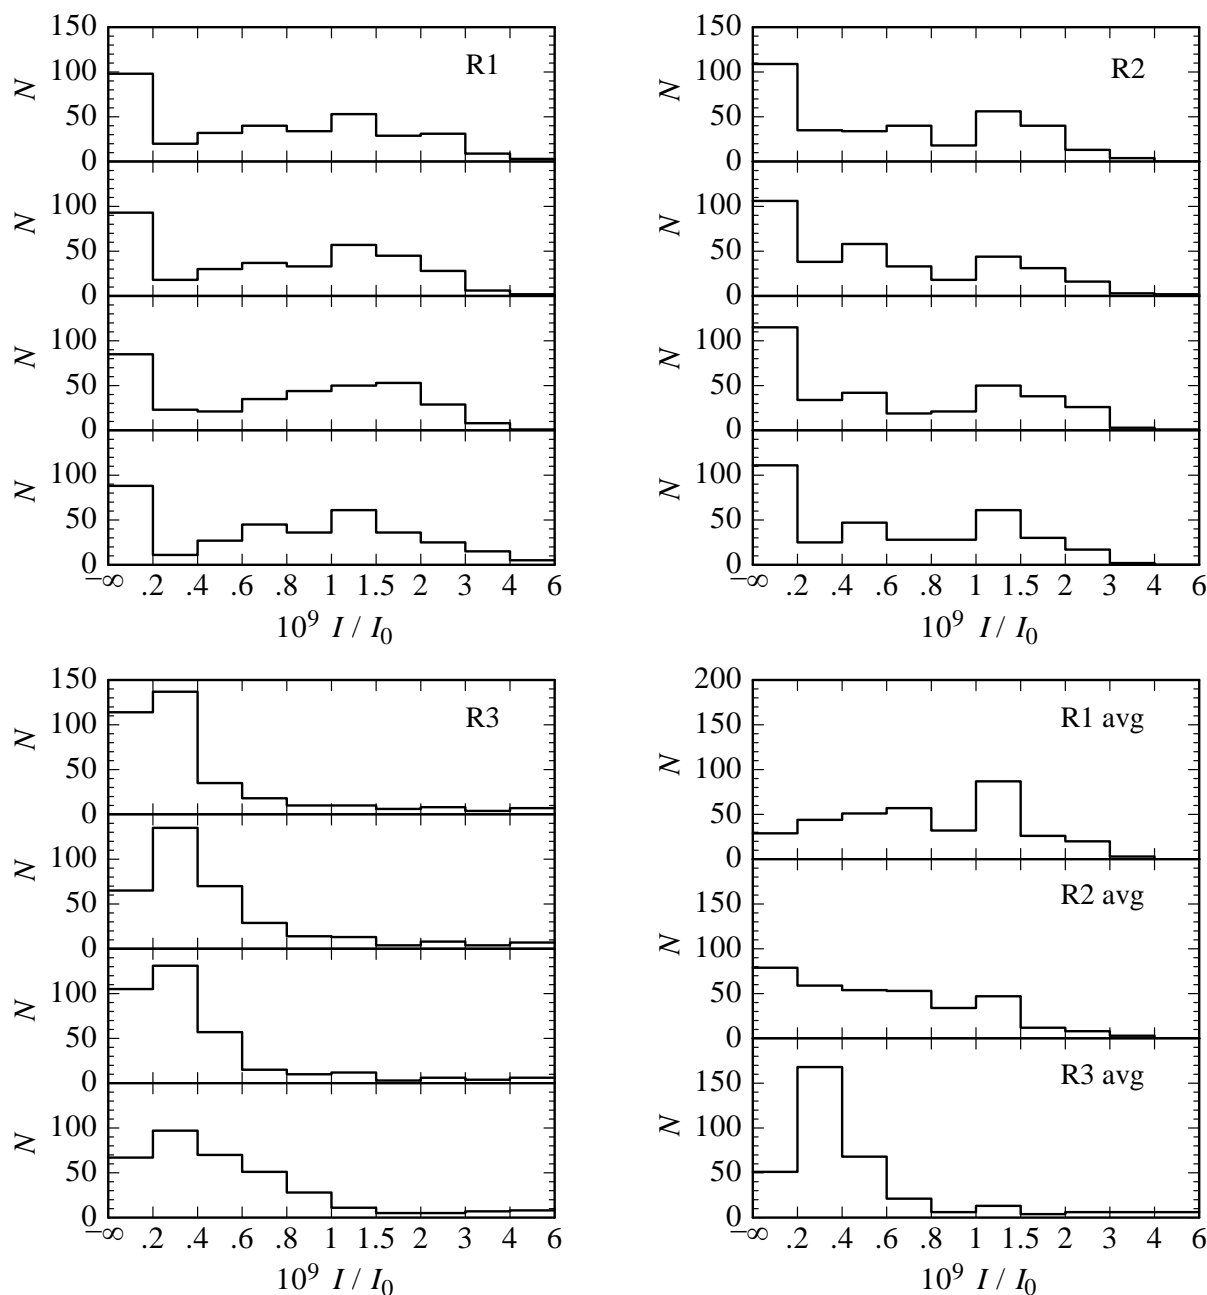


Fig. 16. Histograms of intensity values in dark hole. The top left, top right, and bottom left panels are 4 separate iterations from each of runs R1, R2, and R3. The bottom right panel has the averages of all four iterations in runs R1, R2, and R3. In every panel, there are 349 pixels in the dark hole, *i.e.*, the sum of all histogram bin N values is 349. Note that the bin edges (horizontal axis values) are not spaced uniformly.

Because the measurement noise levels are similar between R1 and R2, the larger observed σ_{each} value for R1 can be interpreted as containing a significant contribution from wavefront control changes to the dark hole during the four iterations, or from dynamic environmental conditions. This is a weak conclusion, as

$\sigma_{\text{each}} = 6.4 \times 10^{-11}$ is not quite a 2-sigma (one-tailed) disagreement from the measurement noise contribution of 4×10^{-11} .

The flux levels in the R3 iterations are more than $10\times$ larger, making the measurement noise negligible compared to the observed σ_{each} values. Here the conclusion is unambiguous that the wavefront control and / or dynamic environment dominate the observed σ_{each} values. This can be seen in Fig. 15, where for R3 the morphological changes from iteration to iteration are correlated over speckle sizes, rather than being dominated by a white (spatially-independent) noise pattern.

5.4. Further analysis of intensity data: Radial scatter plots

Another method to display the pixel-to-pixel variations is the scatter plots of intensity vs. radius in Fig. 17, which also displays the mean intensities binned over 1 pixel in radius, which corresponds to $(1/6.34) \lambda/D_{\text{sky}}$ -sized bins. What is apparent in Fig. 17 is the variety of profiles for the different runs. The dark-hole mean in R3 is dominated by the larger intensities at small radii, while the profile is much flatter




Fig. 17. Intensity vs. radius for milestone run 4-iteration averages. Each blue point is a single pixel, averaged over 4 iterations, from the right-hand column of Fig. 15. The solid black lines are binned averages, with bin width 1 pixel = $(1/6.34) \lambda/D_{\text{sky}}$. Triangles denote points below 10^{-10} .

for R1. The specific way in which wavefront control parameters, as well as measurement SNR, play into these issues has not been disentangled. The diversity of profiles in Fig. 17 is simply a case study in the range of results produced during normal operation, in the absence of any attempt to make them more uniform.

5.5. Further analysis of intensity data: Coherent / incoherent light

The wavefront estimation is a process of modulating the DM in a sequence of images, to produce changes in the intensities across the image plane, which can be analyzed along with the knowledge of the DM modulation to estimate the complex-valued image-plane E -field. In the course of this analysis, it is always true that the observed intensity modulation has an interpretation that includes an intensity component that does not participate in the coherent modulation due to the DM, according to the assumed wavefront propagation. In analogy with phase-shifting interferometers, this non-modulating component, which is assumed to be simply additive by intensity over the sequence of probed images, is labeled the “incoherent” intensity component. The complex E -field estimate is, by construction, the portion of the light that did modulate with the DM, so the intensity (*i.e.*, squared modulus of the E -field) is labeled the “coherent” intensity component. In its simplest interpretation, the incoherent component will not submit to wavefront control to be reduced further, as the DM was unable to modulate it, and so cannot cancel it. By construction, the incoherent intensity is defined as the difference between the observed “unprobed” intensity and the estimated coherent intensity.

The decomposition of the observed intensity into coherent and incoherent intensity components, for the average of four iterations in each run, is shown in Fig. 18. While the interpretation of the incoherent intensity is nonnegative-definite, in the presence of read noise, negative values can be observed simply due

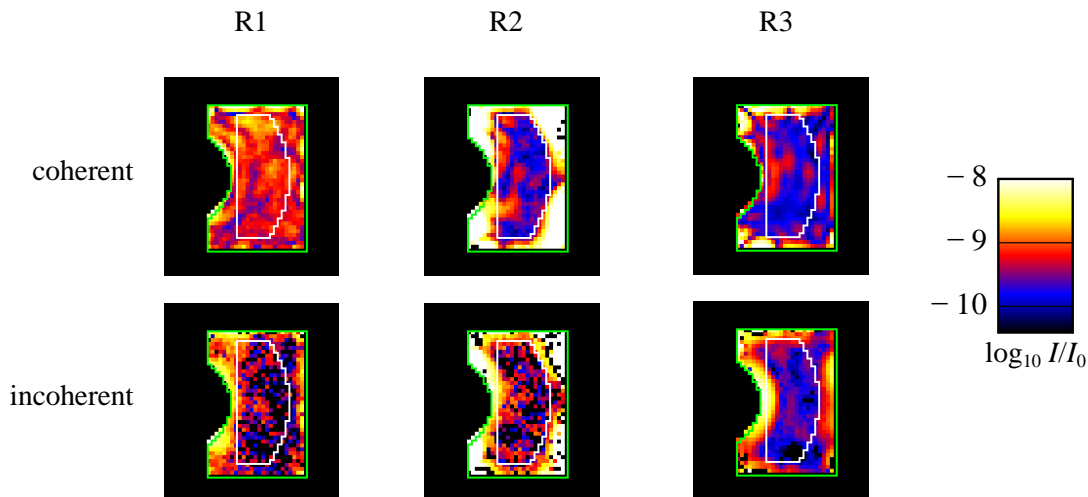


Fig. 18. Separation of light into coherent and incoherent components, averaged over all four iterations in each run. This calculation was done only over a region that was well probed; the edge of this region appears black in these figures (and nearly coincides with the occulter edge in green). Note that the choice of probes can reduce the effect of read noise on the coherent estimates, but for the incoherent estimates read noise dominates pixel-by-pixel for runs R1 and R2. Green and white regions are the same as in Fig. 14.

to measurement errors. In addition, since the magnitude of the coherent E -field estimate is a measurement of modulation amplitude, any uncorrelated noise will bias the coherent estimate to larger intensities, correspondingly biasing the incoherent intensity estimate to lower values. The presence of negative observed values of incoherent intensities in Fig. 18 is due to measurement noise.

The further decomposition of the coherent light into dynamic and static portions is handled somewhat simplistically in this treatment. The complex-valued nature of the coherent estimates allows the E -fields to be directly compared. The simplest statistic to apply to determine a dynamic variability is a variance, so at each pixel in the dark hole, a variance of the complex coherent E -field is measured across the four iterations in a run. The maps of these pixel-by-pixel variances are shown in Fig. 19. These variances, which have the dimensions of intensity, are not the same as the changes in intensity from iteration to iteration – for example, if for one pixel the E -field amplitude remained the same across four iterations but the phase changed through 2π , the variance of the E -field would be larger than the mean intensity, even though the coherent intensity was identical for the four iterations (i.e., zero variance in intensity). This E -field variance is the only metric of dynamic coherent component analyzed in this treatment.

The coherent E -field variances averaged over the dark hole, σ^2_{coh} , are listed in Table 2, along with the mean coherent intensity, \bar{m}_{coh} , and the mean incoherent intensity, \bar{m}_{inc} . The interpretation of the variance values is largely a question of the use of wavefront control between the four iterations. If the wavefront control was actively trying to change the E -field, the variance is large; if the wavefront control was making only small changes, then the variance is small. A more complicated treatment would further decompose this light based on the expected changes from wavefront control (i.e., what the DM changes should have produced), and the telemetry from the CLOWFS system. This treatment does not separate the wavefront control action from any environmental changes. It is true, for example, that the wavefront control was making larger changes to the dark hole during Run 1 than in the other runs, based on the DM motions. There was no attempt to make the wavefront control treatment uniform across the three runs.




Fig. 19. Coherent E -field variances, calculated over the four iterations in each run. Compare to the total coherent intensity panel of Fig. 17. Green and white regions are the same as in Fig. 14.

Run	$\bar{m} \times 10^{10}$	$\bar{m}_{\text{coh}} \times 10^{10}$	$\bar{m}_{\text{inc}} \times 10^{10}$	$\sigma^2_{\text{coh}} \times 10^{10}$
R1	9.03	7.24	1.79	5.58
R2	6.27	3.17	3.10	1.14
R3	5.72	2.20	3.52	1.00

Table 2. Mean coherent E -field variances in runs R1-R3. There is no milestone requirement on any of these numbers.

The static coherent component may be defined as the difference between the coherent intensity, \bar{m}_{coh} , and the coherent E -field variance, σ_{coh}^2 . This is not independently tabulated, but is readily available from Table 2.

5.6. Further analysis of intensity data: Effect of dark hole shape and size

The occulter used in this experiment was fabricated with an error, in that the opening was too small in its y -extent (see Fig. 5). The size of the mask did not allow the use of the full “D”-shape originally specified in the milestone definition (see Section 3.5.3 and Fig. 14). The difference in area between the full “D” shape and the milestone scored region is not large, a difference between 349 pixels and 403 pixels, equivalent to $8.7 (\lambda/D_{\text{sky}})^2$ and $10.0 (\lambda/D_{\text{sky}})^2$. A histogram of the radii at which these pixels lie is shown in the left panel of Fig. 20, where it is apparent that the missing pixels are located at radii between 3 and $4 \lambda/D_{\text{sky}}$. From Fig. 17, it is clear that the intensities at these radii are lower than in the $2\text{--}3 \lambda/D_{\text{sky}}$ range. From this argument, it would follow that if the correct occulter had been used in this experiment, and the full “D” shape had been measured, the mean intensities would have come out lower than what is reported in Tables 1 and 2.

A separate argument would be that larger areas are harder to control, due to a limited number of degrees of freedom allowed by the finite number of DM actuators. In this argument, the full “D” shape would have higher mean intensities than the scored region, because of the larger area. This argument can be tested, because the region subject to wavefront control was larger than the scored region, and a larger region can be evaluated. In this analysis, the “extended” region in Fig. 21 shown in cyan has the same area as that shown in black. The mean radii inside each region (the area-weighted radii) are nearly the same, $3.24 \lambda/D_{\text{sky}}$ for the full “D” and $3.23 \lambda/D_{\text{sky}}$ for the extended region, as a similar number of pixels were added at smaller and larger radii than the radii of those excluded. The right-hand panel in Fig. 20 shows the radii of the pixels in these regions. The mean intensity in the extended region was 5.69×10^{-10} for R3 (the run with the highest signal-to-noise ratio), as compared to 5.72×10^{-10} in the scored region. This difference is negligible in the context of the milestone statistics, and so the conclusion of this




Fig. 20. Number of pixels at given radii for (LEFT) scored region used for milestone report and (RIGHT) extended region for size analysis. The solid black line is the full “D” (in both panels), the dashed line is the scored region, and the cyan line is the extended region.




Fig. 21. Relative positions of occulter edge (green), milestone “D” (black), and extended region (cyan). The areas of the cyan and black regions are equal. The white region of previous images is the intersection of the black and cyan regions. Intensity image is the average from run R3.

analysis is that there would be no appreciable difference between the scored region and the full “D” region (if the occulter would have allowed it).

5.7. Description of data analysis

The detailed description of the data analysis, referenced in Section 4.3.7, appears as Appendix 5.

6. CONCLUSIONS

The data described in this report demonstrate the combination of PIAA optical elements and wavefront control producing a monochromatic dark hole with average intensity below the 10^{-9} level, at radii down to $2 \lambda/D_{\text{sky}}$. The details of the milestone requirements and the statistical analysis are presented in full, to validate the performance.

Appendix 1. Occulter description and CLOWFS operation

The design of the occulter used in this experiment is similar to that used in the PIAA TDEM Milestone 2, completed in 2012 [18]. For Milestone 2, the occulter (see left panel of Fig. 22) was patterned on glass and had three zones: a circular absorbing region, an annular reflective region, and a fully transmitting region (in order of increasing radii). The reflecting region sends light to the CLOWFS. For this experiment, the occulter still had three zones, but not in a circular geometry. The transmitting region (dark in the right panel of Fig. 22) was a void (etched out of a silicon substrate), the absorbing circle was “black silicon,” and the reflecting region (everywhere but the small circle or the void) was coated with aluminum to be reflective. The occulter used for this milestone was not the exact occulter imaged in the right panel of Fig. 22, but differed only in the specific dimensions of the open void (compare to Fig. 5), most notably the x dimension of Fig. 5.




Fig. 22. (LEFT) Optical microscope image of occulter used for PIAA TDEM Milestone 2, and (RIGHT) scanning electron microscope image of occulter similar to that used in this experiment.

There are two main differences between the two occulter shown in Fig. 22. The first is that the Milestone 2 occulter was patterned on glass, which contributed some unknown amount of scattering and ghost reflections. The second difference is that the reflecting region of the Milestone 1 occulter is much larger, and notably, it is not symmetric with respect to the absorbing circle. Because the PSF landing on the center of the occulter is relatively faint at the inner radius of the transmitting region, the fact that one side reflects that light while the other transmits makes little difference to the CLOWFS. Specifically, the CLOWFS signal is dominated by the light at the inner edge of the absorbing region, with very little influence from the inner edge of the transmitting region.

Appendix 2. Calculation of λ/D_{sky}

The remapping apodization of the PIAA mirrors creates a subtle difficulty in directly measuring λ/D_{sky} , in that the typical diffractive signatures from uniformly illuminated apertures (e.g., Airy rings) cannot be quantified as simply as in other systems. The simplest alternative is to directly measure the relevant parameters on the testbed. The most relevant parameter, the diameter of the entrance pupil, is not directly measurable because the pupil stop is located two optics downstream of the first mirror, with the PIAA remapping in between. For a variety of minor reasons, the full “clear aperture” of the PIAA mirror design was not used, in the sense that the postapodizer diameter is smaller than the PIAA clear aperture diameter observed in the same plane.

A variety of combinations of measured distances and design values was used to cross-check each other, and to establish an estimate of disagreement between them. The most straightforward of these combinations gave the value of λ/D_{sky} used throughout this report, $\lambda/D_{\text{sky}} = 6.34$ pix. This technique compiles the following measurements and design values:

1. Measure the ratio of the pupil stop diameter D_p to the full M2 diameter D_{M2} as imaged at the science camera (at plane P4 in Fig. 3)
2. Use the design remapping to determine the diameter on M1, D_1 , that remaps to D_p
3. Use the design focal length of M1, f_{M1} , to determine $(f/D)_{\text{src}} = f_{M1}/D_1$, the focal ratio at the source
4. Measure the wavelength λ to calculate $(f\lambda/D)_{\text{src}}$, the lateral distance at the source corresponding to λ/D_{sky}
5. Move the source, measure the image motion at the science camera in pixels versus the source motion measured in microns using the source encoders, to determine the lateral image magnification $M_{\text{src-cam}}$ [pix/ μm]
6. Multiply these to determine λ/D_{sky} measured in pixels

The values that enter into this calculation are $D_p/D_{M2} = 426.4 \text{ pix} / 496 \text{ pix} = 0.860$, $D_1 = 74.6 \text{ mm}$, $f_{M1} = 1141 \text{ mm}$, $(f/D)_{\text{src}} = 15.30$, $\lambda = 807.5 \text{ nm}$, $(f\lambda/D)_{\text{src}} = 12.35 \mu\text{m}$, $M_{\text{src-cam}} = 0.513 \text{ pix}/\mu\text{m}$, $\lambda/D_{\text{sky}} = 6.34 \text{ pix}$.

Using other measured quantities, $M_{\text{occ-cam}}$, $(f/D)_{\text{cam}}$, $(f\lambda/D)_{\text{cam}}$, determined independently from the above measures and from each other, a consistency check can be made. The worst-case disagreement is 2%, which is conservatively adopted as the uncertainty in λ/D_{sky} .

Appendix 3. Model of testbed for wavefront control

The model used to infer the relationship between actuator motions and dark hole intensities can be broken into three components: a DM model, an illumination model, and a propagation model. These models divide the optical train into three blocks, as shown in Fig. 23. For the current application, the illumination model is simply the application of the measured complex pupil-plane wavefront to the DM. Since this is measured adequately on the testbed, there is no need to rely on any assumptions about the source geometry, the PIAA optics, or any other concerns. This is then simply a measurement denoted

$$E_p(x_p, y_p), \quad \text{Illumination}$$

shown in Fig. 13.

The DM model itself can be considered to fall into two parts, one relating voltages to discretized actuator heights,

$$\vec{h} = H(\vec{v}), \quad \vec{h} = \{h_1, \dots, h_N\}, \quad [h_i] = \text{nm}, \quad \vec{v} = [v_1, \dots, v_N], \quad [v_i] = \text{V}, \quad \text{DM(a)}$$

defined by a function H for N actuators (1024 in total, although some are completely obscured), and another relating discretized actuator heights to a continuous DM surface shape s ,

$$s(x_p, y_p) = S(x_p, y_p; \vec{h}), \quad [s] = \text{nm}, \quad \text{DM(b)}$$

defined by a function S . Here, x_p and y_p are pupil-plane locations, which are applied to all planes conjugate to the pupil (i.e., the DM and Lyot planes). The effect of the DM surface shape on the E-field leaving the DM is

$$E_{\text{DM}}(x_p, y_p) = E_p(x_p, y_p) e^{i4\pi s/\lambda}$$

The propagation model is a sequence of three Fourier transforms \mathcal{F} and multiplication by the occulting mask M and Lyot stop L transmissions,

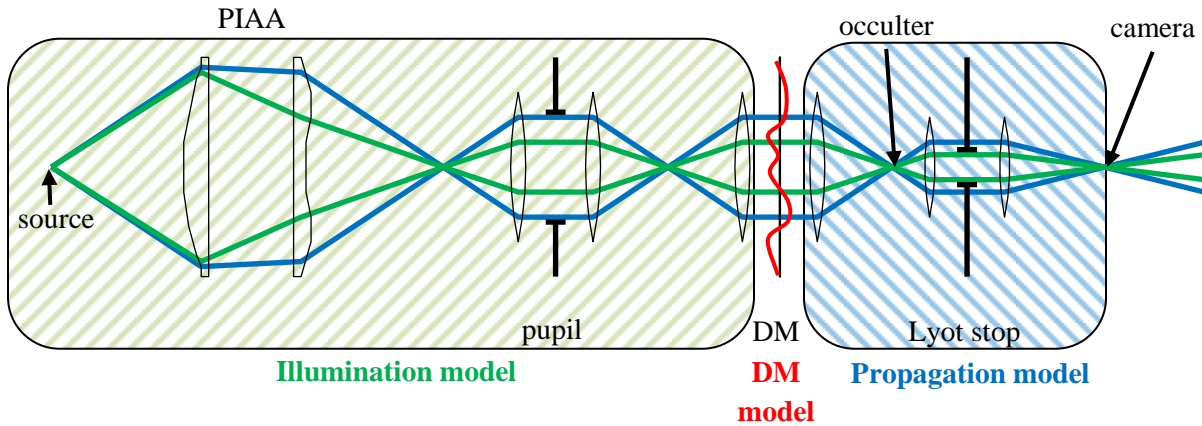


Fig. 23. Division of testbed model into three components, an illumination model, DM model, and propagation model.

$$E_{\text{occ}}(x_s, y_s) = \mathcal{F}_{\text{P-S}}\{E_{\text{DM}}(x_p, y_p)\} M(x_s, y_s) \quad \text{Propagation}$$

$$E_{\text{Ly}}(x_p, y_p) = \mathcal{F}_{\text{S-P}}\{E_{\text{occ}}(x_s, y_s)\} L(x_p, y_p)$$

$$E_{\text{cor}}(x_s, y_s) = \mathcal{F}_{\text{P-S}}\{E_{\text{Ly}}(x_p, y_p)\}$$

where E_{cor} is the E-field in the coronagraph image, with coordinates x_s, y_s . The mask transmission M is shown in Fig. 5, and effect of the Lyot transmission L is easily seen in Fig. 6.

E_{cor} is a linear function of E_{DM} , and E_{DM} is a product of a static complex field (w.r.t. actuator heights) and the phase effect of the DM surface. It is then straightforward to describe the Jacobian of E_{cor} as a function of actuator heights, $\partial E_{\text{cor}}/\partial h_i$, which is sampled at each spatial pixel in the dark hole from each actuator on the DM. The job of wavefront control is then to take an estimate of the complex E -field at every pixel in the dark hole, E_{cor} , and solve for a Δh_i set of actuator motions that will produce $\Delta E_{\text{cor}} = -E_{\text{cor}}$, reducing the intensity in the dark hole to zero. Because the relationship between E_{DM} and h_i is nonlinear, this solution is best treated as an iterative process, with each iteration further reducing the squared residuals remaining at the end of the previous iteration.

Appendix 4. Applicability of Gaussian statistics

The residual light in the dark hole, after wavefront control, has a mix of organized spatial structure (specifically radial structures, see Fig. 17) and speckle-sized variations. The intention of the requirement to analyze the applicability of Gaussian statistics was to isolate the radial variations and the disorganized residual structures. However, as is clear from Fig. 17, the radial structures are pronounced, and have different realizations between different runs. In addition (see Sec. 5.6), the area of the dark hole is limited to $8.7 (\lambda/D_{\text{sky}})^2$, giving only ~ 9 independent samples to examine spatial variations. With such a limited spatial sample set and the varied and pronounced differences in the organized structures in the dark holes, it is not reasonable to expect a discrimination between a normal distribution and the observed distribution of light decomposed into organized and disorganized structures. No further statistical analysis, beyond that reported above, was performed to determine the form of the distribution.

Appendix 5. Data reduction procedures

The data reduction performed consists of four steps: registration, bias subtraction, dark subtraction, and photometric normalization. Each step is described below.

The image registration is performed by regularly removing the occulter and observing the location of the unocculted peak brightness. This is implicitly part of the photometric calibration sequence described in Sec. 3.3. This location is then used as the coordinate zero-point for all further analysis.

The CCD frames include a bias region, which is an overscan of the serial register by approximately 150 pixels. To avoid low-level persistent readout issues, a margin of rows and columns from the first line, and the last active column are excluded from the region used to determine the bias levels. Specifically, the region from columns 1080 to 1200, and rows 200-600 are averaged to determine the amplifier bias counts in each image. This bias level is calculated and subtracted image by image.

The CCD readout takes approximately 6 s, during which time dark current is accumulating on each pixel. A zero-second exposure, tied to the 6 s readout, will have statistically identical dark current for every pixel in a given column. Therefore, the dark current accumulated during readout is determined by taking a number of zero-second exposures, and averaging over rows and over exposures (after bias subtraction of every exposure) to determine a 1-D (columnar) dark current profile specific to the readout time. This is called the readout dark current, typically including 100 exposures in the average.

After determining the readout dark current, a number of 60 s exposures are taken. These are bias subtracted and readout dark current subtracted (from each row), then averaged over the exposures to form a 2-D 60 s linear dark current map. All subsequent images will subtract this map, scaled by exposure time relative to 60 s.

The last step for coronagraphic images is the photometric normalization, described in principle in Sec. 3.3. Specifically, after registration of each image, the pixels in the “photometric reference region” (see Fig. 9) are summed, and multiplied by the pre-determined photometric calibration factor, to infer the unocculted source peak brightness. The image is divided by this peak brightness, to produce the normalized intensity map.

To quantify these procedures, after acquiring all the calibration images, we have the unocculted star center location (x_c, y_c) , the readout dark current $DC_0(x)$, the linear dark current $DC_L(x,y)$, the photometric calibration factor P , and the photometric region boundary corners (x_{p1}, y_{p1}) and (x_{p2}, y_{p2}) . Given a raw intensity $I_r(x,y)$, first the bias-and-dark-subtracted intensity $I_{bd}(x,y)$ is calculated,

$$I_{bd}(x,y) = I_r(x+x_c, y+y_c) - \langle I_r(1080:1200, 200:600) \rangle - DC_0(x+x_c) - DC_L(x+x_c, y+y_c) (t/60 \text{ s}),$$

And then the normalized intensity $I(x,y)$ is calculated as

$$I(x,y) = I_{bd}(x,y) / (P \sum I_{bd}(x_{p1}:x_{p2}, y_{p1}:y_{p2}))$$

For these runs, the photometric region boundary corners were $(-29,57)$ and $(26,66)$, and the photometric calibration factor P was 12065 for R1, 12065 for R2, and 12756 for R3.

REFERENCES

- [1] Guyon, O. et al. 2010, “Phase-Induced Amplitude Apodization Coronagraph: Monochromatic Contrast Demonstration,” <http://exep.jpl.nasa.gov/files/exep/Guyon%20Milestone%20WP%201%20Final%20signed%20CR%202011.pdf>
- [2] Guyon, O. 2003, “Phase-induced amplitude apodization of telescope pupils for extrasolar terrestrial planet imaging,” *A&A*, **404**, 379.
- [3] Traub, W.A. and Vanderbei, R.J. 2003, "Two-Mirror Apodization for High-Contrast Imaging," *ApJ* **599**, 695.
- [4] Vanderbei, R.J., Traub, W.A. 2005, "Pupil Mapping in Two Dimensions for High-Contrast Imaging," *ApJ*, 626, 1079.
- [5] Guyon, O., Pluzhnik, E.A., Galicher, R., Martinache, F., Ridgway, S.T., Woodruff, R.A. 2005, "Exoplanet Imaging with a Phase-induced Amplitude Apodization Coronagraph. I. Principle," *ApJ* **622**, 744.

- [6] Martinache, F., Guyon, O., Pluzhnik, E.A., Galicher, R., Ridgway, S.T. 2006, "Exoplanet Imaging with a Phase-Induced Amplitude Apodization Coronagraph. II. Performance," *ApJ* **639**, 1129.
- [7] Vanderbei, R.J. 2006, "Diffraction Analysis of Two-dimensional Pupil Mapping for High-Contrast Imaging," *ApJ* **636**, 528.
- [8] Pluzhnik, E.A., Guyon, O., Ridgway, S.T., Martinache, F., Woodruff, R.A., Blain, C., Galicher, R. 2006, "Exoplanet Imaging with a Phase-Induced Amplitude Apodization Coronagraph. III. Diffraction Effects and Coronagraph Design," *ApJ* **644**, 1246.
- [9] Guyon, O., Pluzhnik, E.A., Kuchner, M.J., Collins, B., Ridgway, S.T. 2006, "Theoretical Limits on Extrasolar Terrestrial Planet Detection with Coronagraphs", *ApJSS* **167**, 81.
- [10] Lozi, J., Martinache, F., Guyon, O. 2009, "Phase-Induced Amplitude Apodization on centrally obscured pupils: design and first laboratory demonstration for the Subaru Telescope pupil," *PASP* **121**, 1232.
- [11] Guyon, O., Martinache, F., Belikov, R., Soummer, R. 2010, "High Performance PIAA Coronagraphy with Complex Amplitude Focal Plane Masks", *ApJSS* **190**, 220.
- [12] PECO website: <http://zero.as.arizona.edu/CAAO/PECO/>
- [13] Levine, M. et al. 2009, "Levine, M., et al. 2009 Terrestrial Planet Finder – Coronagraph (TPF-C) Flight Baseline Mission Concept", white paper provided to the 2010 Decadal Survey.
- [14] Lawson, P. R. et al. 2013, "Survey of experimental results in high-contrast imaging for future exoplanet missions," *Proc. SPIE* **8864**, 88641F.
- [15] Kern, B., Give'on, A., Kuhnert, A., Levine-West, M.B., McMichael, I., Moody, D.C., Niessner, A.F., Pueyo, L., Shaklan, S.B., Traub, W.A., Trauger, J.T., Belikov, R., Guyon, O. 2009, "Phase-induced amplitude apodization (PIAA) coronagraph testing at the High Contrast Imaging Testbed," *Proc. SPIE* **7440**, 74400H.
- [16] Kern, B., Give'on, A., Kuhnert, A., Niessner, A., Guyon, O. 2011, "Laboratory testing of a Phase-Induced Amplitude Apodization (PIAA) coronagraph," *Proc. SPIE* **8151**, 815104.
- [17] Kern, B., Guyon, O., Kuhnert, A., Niessner, A., Martinache, F., Balasubramanian, K. 2013, "Laboratory demonstration of Phase Induced Amplitude Apodization (PIAA) coronagraph with better than 10^{-9} contrast," *Proc. SPIE* **8864**, 88640R.
- [18] Guyon, O., Kern, B., and Martinache, F., "Phase-Induced Amplitude Apodization (PIAA) Technology Development, Milestone #2: Instrument Tip-Tilt Control Demonstration at Sub-Milliarcsecond Levels," http://exep.jpl.nasa.gov/files/exep/PIAAmilestone2_final120821.pdf
- [19] Lozi, J., Martinache, F., and Guyon, O. 2009, "Phase-Induced Amplitude Apodization on Centrally Obscured Pupils: Design and First Laboratory Demonstration for the Subaru Telescope Pupil," *PASP* **121**, 1232.
- [20] Trauger, J. and Traub, W. 2007, "A laboratory demonstration of the capability to image an Earth-like extrasolar planet," *Nature* **446**, 771.
- [21] Give'on, A., Kern, B., Shaklan, S., Moody, D.C., Pueyo, L. 2007, "Broadband wavefront correction algorithm for high-contrast imaging systems," *Proc. SPIE* **6691**, 66910A.