



JPL Document D- 74426

TECHNOLOGY DEVELOPMENT FOR EXOPLANET MISSIONS (TDEM)

**Assessing the performance limits of internal
coronagraphs through end-to-end modeling**

Technology Milestone #2 Results Report

**John Krist, PI
Ruslan Belikov
Dimitri Mawet
Dwight Moody
Laurent Pueyo
Stuart Shaklan
John Trauger**

7 December 2012, Revised 1 July 2013

Approved for public distribution

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

© 2012, 2013 All rights reserved.

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

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

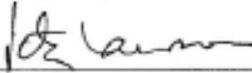
Released by



John Krist
Principal Investigator

26 Oct 2012
Date

Approved by



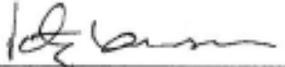
Peter Lawson
Exoplanet Exploration Program Chief Technologist, JPL

31 Oct. 2012
Date



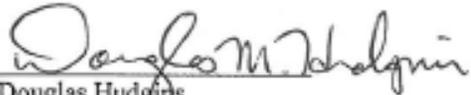
Marie Levine
Exoplanet Exploration Program Technology Manager, JPL

10/31/12
Date



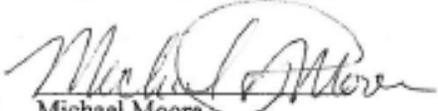
Peter Lawson
Exoplanet Exploration Program Manager (Acting), JPL

31 Oct. 2012
Date



Douglas Hudgins
Exoplanet Exploration Program Scientist, NASA HQ

8 Nov. 2012
Date



Michael Moore
Exoplanet Exploration Program Executive, NASA HQ

9/11/12
Date

Assessing the performance limits of internal coronagraphs

Table of Contents

1. Objectives.....	1
2. Introduction.....	1
2.1. Coronagraphic wavefront control	1
2.2. Goals of this study	2
2.3. Application to future NASA missions	3
2.4. Caveats	3
3. Milestone #2 Description	5
3.1. Milestone Prerequisites.....	5
3.1.1 Contrast definition	5
3.1.2 Coronagraph contrast and image plane field dimensions	5
3.1.3 Optical system layout	5
3.1.4 Wavefront control.....	6
3.2. Milestone Requirements	6
3.3. Milestone Metrics	10
4. Success Criteria.....	11
5. Milestone #2 Results	12
6. Milestone Certification Data Package.....	14
7. References.....	15
Appendix A: Layouts.....	17
Appendix B: Wavefront control	18
Appendix C: Results for the Hybrid BandLimited Coronagraph (HBLC).....	23
Appendix D: Results for the Vector Vortex Coronagraph (VVC).....	28
Appendix E: Results for the Phase-Induced Amplitude Apodization (PIAA) coronagraph.....	34
Appendix F. Acronyms.....	49

TDEM Milestone #2 Results Report: Assessing the Performance Limits of Internal Coronagraphs through End-to-End Modeling

1. Objectives

In support of NASA's Exoplanet Exploration Program and the ROSES Technology Development for Exoplanet Missions (TDEM), this report documents the achievement of Milestone #2 for our TDEM study, which is an assessment of the theoretical performance limits of selected coronagraphs as derived through numerical end-to-end modeling of a system with realistic optical aberrations. This milestone used the propagation algorithms developed in Milestone #1 (Krist et al. 2010, 2011) to characterize the wavefront control behavior of each coronagraph as predicted by numerical simulations (rather than hardware testing) and to identify the limiting factors for achieving 10^{-10} contrast over a broad bandpass. These results will help define the instrumental, testbed, and space mission configurations necessary to demonstrate and operate at this level using realistic technologies.

This report details only those aspects relevant to Milestone #2. The reader should review the Milestone #1 report for details on the context for the overall investigation, the coronagraphic technologies under study (hybrid bandlimited coronagraph (HBLC), vector vortex coronagraph (VVC), and phase-induced amplitude apodization (PIAA)), and the associated wavefront propagation algorithms.

2. Introduction

The technology milestone described here serves to gauge the developmental progress of optical modeling for a space-based coronagraphic mission such as ACCESS (Trauger et al. 2008) or the Terrestrial Planet Finder Coronagraph (TPF-C; Traub et al. 2006) that would detect and characterize exoplanets.

2.1. Coronagraphic wavefront control

Having the means to simulate propagation of a wavefront through a system, as developed in Milestone #1, is not sufficient to predict the contrast limit of a given telescope and coronagraph. Each coronagraph responds differently to wavefront aberrations and sets particular limits on the ability to control those errors with deformable mirrors, especially over a broad wavelength range (Shaklan & Green 2006). Therefore, the models must be executed within a wavefront control framework like that used in real systems that senses the simulated electric field at the image plane and then determines the deformable mirror (DM) actuator settings necessary to minimize the scattered light around the star (Give'on et al. 2007; Krist, Trauger, & Moody 2006; Krist et al. 2009).

There are advantages to simulating the various coronagraphs using the same basic framework. The performance of each coronagraph can be compared to the others on even terms, given that the same front-end aberrations and wavefront control methods are used. Piecing together the results from different studies with varying layouts and modeling methods will not provide the consistency necessary to identify the true capabilities of each coronagraph. Unforeseen and perhaps subtle causes for the poor performance of a coronagraph might be more readily identified by comparing its behavior to the others, especially how each responds to changes in the wavefront caused by the deformable mirrors. Any problems can be diagnosed as being system-level (if all coronagraphs have poor performance) or particular to just one coronagraph (or its modeling technique).

Using simulations to determine how a coronagraph responds to wavefront control is critical to demonstrating the technological readiness of these systems. Prior to building and installing a coronagraph on a testbed, modeling can be used to decide what configuration may be required to provide the required performance (e.g., the location and surface quality of critical optics, the number and positioning of DMs, etc.). It can also highlight differences between testbed and proposed mission layouts. For example, PIAA testbed experiments so far (Belikov et al. 2009) have used a single DM *after* the PIAA optics (after the beam has been remapped and apodized), but proposed mission concepts use DMs located *before* PIAA, prior to wavefront remapping, to provide the maximum outer working angle. This study is the first to accurately simulate the performance of PIAA in such a flight-like layout.

2.2. Goals of this study

Our study was specifically aimed at determining if there are fundamental, wavefront-modifying properties of the various coronagraph designs that, when used in a realistically aberrated system with wavefront control, would prevent attaining 10^{-10} contrast, the commonly accepted level for Earth-twin visible-light, imaging missions. A coronagraph only suppresses the diffraction pattern produced by the telescope. If the optical system is perfect and all the light is concentrated in the diffraction pattern, then any of the three designs used here (HBLC, VVC, PIAA) would be able to suppress the starlight to below 10^{-10} contrast, by design. However, any real system has aberrations that create speckles of scattered light that must be suppressed using wavefront control (deformable mirrors). A coronagraph, even one whose components are perfect, may alter the aberrations in a manner that prevents the wavefront control system from reducing the errors below the desired level. This may be due to wavefront remapping, conversion of phase errors to amplitude errors, etc. Such behavior may not be readily apparent until the coronagraph is implemented in an aberrated system with wavefront control. Due to the expense of hardware and testbeds, it is prudent to first predict the coronagraphic performance using end-to-end modeling in a simulated, realistically-aberrated system.

Using the propagation algorithms established in Milestone 1, we conducted end-to-end numerical modeling of each coronagraph in a realistically-aberrated optical system with wavefront control to determine its performance over a $\lambda = 500 - 600$ nm bandpass. The

mean contrasts for the modeled systems were evaluated in an imaging field of dimension $r = 2.5 - 18 \lambda_c/D$ radians ($\lambda_c = 550$ nm) around the central source. From these trials, characteristics of each coronagraph that affected the contrast limits with wavefront control were noted and will be described.

We note that this study was purely concerned with numerical modeling of the coronagraphs. It did not involve any actual hardware implementations of the systems, such as testbeds like HCIT. The goal was to understand the behavior of the coronagraphs before dedicating resources to placing them in testbeds or telescopes.

2.3. Application to future NASA missions

Any future mission that uses one of these coronagraphs will require the algorithms and parameters derived in this study to

- Determine the performance of the coronagraph in real-world conditions
- Plan testbeds used to evaluate prototype and flight coronagraphs that properly replicate flight layouts and properties
- Define system layouts and optical parameters that provide sufficient performance
- Generate the DM response matrix that is used on-orbit for determining the DM settings that produce a dark hole in the image plane around the star that allows for high contrast imaging

2.4. Caveats

Scalar versus vector propagation

The modeling undertaken in this study assumed scalar propagation of the wavefront. Vector propagation, which includes the physical effects of electric field interactions with conductive and non-conductive materials at small scales, was not used. Vector propagation becomes important when small apertures may act as waveguides and the electrical properties of the aperture substrate are significant (Lieber et al. 2005). In this study the impact of any vectorial effects would be mainly in the PIAA binary post-apodizer, the small occulting spot at the center of the VVC mask, and the amplitude-modifying structure of the HBLC. Based on previous studies for the Terrestrial Planet Finder Coronagraph, we expect such effects to be small as these structures are thin (i.e. we are not using thick apertures such as those used for early shaped pupil experiments). The realm of vector propagation was well beyond the time and financial limits of this study. We note that although the models do not use vector propagations, the HCIT laboratory results to date are consistent with predictions using scalar models to contrasts of $<10^{-9}$ broadband and 2×10^{-10} narrowband.

Polarization

We also did not account for the effects of polarization. It is known that polarization-induced aberrations can limit contrast at the levels we are concerned with here ($\sim 10^{-10}$) in all of the coronagraphs being considered (Elias et al. 2004; Balasubramanian et al. 2011).

These can be minimized to some degree with judicious choices for coatings and system layouts, but they cannot be completely negated. We therefore assumed that our simulations represented one polarization channel. It is known that the VVC requires a single polarization to provide broadband contrast at the 10^{-10} level.

Wavefront sensing

This study also did not attempt to replicate a realistic on-orbit wavefront sensing and control sequence in which the telescope pointing and thermal effects in the optics cause wavefront changes over time. The goal was to understand the fundamental limits of a static system with the understanding that it represents the best case scenario for a dynamic one.

Time-dependent wavefront variations

This study did not include time-dependent wavefront variations due to stresses from the thermal and dynamic environment as might be experienced in flight. Whereas these factors are important for the relative comparison of coronagraph methods, they were beyond the scope of this study. The only sources of contrast degradation that were modeled were those internal to the optical system. Other sources of contrast loss, external to the instrument, would need to be considered as part of a more comprehensive error budget in future modeling studies.

Coronagraph realism

The representations of the coronagraphs evaluated in this study were not all at the same level of reality, and so the performance results cannot be directly compared among the designs. The HBLC design was fairly realistic, with wavelength-dependent amplitude and phase variations determined from thin-film calculations based on the known properties of the metal and dielectric layers. This design could actually be fabricated and used in a real instrument with reasonable expectations that its performance would match that predicted by the models over a broad bandpass. Note that deviations in the mask design caused by fabrication errors were not included.

The initial PIAA coronagraph representation was also fairly realistic. The binary post-apodizer mask was designed to be within the capabilities of current lithographic/etching methods by placing a limit on allowed width of the thinnest ring. For a real system, a revised mask would be devised incorporating the effects of diffraction to provide better performance at the inner working angle (the mask used in the study was designed for a non-diffractive beam). Measured surface error maps of actual PIAA optics were used in the end-to-end modeling. These errors were shown to be too large to provide the required contrast, and so they were iteratively reduced until they did. However, the final adjusted maps are not within current fabrication capabilities, primarily at the edge of the PIAA M1 optic. Correction with a DM would require a large number of actuators (>140 across the pupil diameter) and strokes of ± 60 nm.

The VVC mask representation did not include any errors or wavelength-dependent, material-induced effects. The results represent what would be expected from a perfectly fabricated five-layer design using known materials and coating thicknesses. However, interference effects were not included. The VVC was therefore the least realistically represented system in this study.

3. Milestone #2 Description

Using the algorithms established in Milestone 1, we will assess the relative performance of HBLC, VVC, and PIAA coronagraphs via end-to-end modeling in a realistic optical system with wavefront control to achieve a numerically-predicted mean contrast of 10^{-10} within a specified annulus centered on the star integrated over a ~20% bandpass. Representative parameters of the optical system and wavefront control subsystem required to meet this contrast requirement will be derived.

3.1. Milestone Prerequisites

3.1.1 Contrast definition

Contrast is defined in this study as the ratio of the peak pixel value of the unocculted stellar point spread function (PSF) to the mean per-pixel surface brightness measured within a specified field around the star, corrected for spatial variations in transmission. A field contrast of 10^{-10} would indicate that a field point source (planet) 10^{10} times fainter than the star would have a peak pixel value equal to the mean per-pixel field brightness. The image fields in our study were sampled at $0.4 \lambda/D$ radians/pixel at $\lambda = 500$ nm ($1.2\times$ better than Nyquist). To correct for transmission variations near the inner working angle due to the occulter, the PSF peak pixel value at each radius was determined by stepping the source from the center and across the occulter.

3.1.2 Coronagraph contrast and image plane field dimensions

The imaging field of concern was an annulus centered on the star extending between $r = 2.5 \lambda_c/D - 18 \lambda_c/D$ radians in the image plane measured across a $\lambda = 500 - 600$ nm bandpass ($\lambda_c = 550$ nm, $D =$ telescope diameter). The inner radius was set by the 50% transmission point of the occulter. The outer radius was set by the number of deformable mirror actuators across the pupil (46 in this study) and the shortest passband wavelength ($18 \lambda_c/D \approx 20 \lambda/D$ at $\lambda = 500$ nm). All of the coronagraphic designs evaluated in this study were tailored to provide *in an aberration-free system* a mean contrast of $<10^{-10}$ within this field.

3.1.3 Optical system layout

The same optical system layouts used for the efficiency tests in Milestone #1 were used for Milestone #2. There were two layouts; one common to VVC and HBLC and one for PIAA. Both systems included two deformable mirrors in series for wavefront control. The systems were identical up to and including the 2nd DM. The system was represented as an unfolded (linear) layout and implemented using the PROPER optical propagation software (Krist 2007) with the custom routines developed in Milestone #1 specific for the coronagraphs.

3.1.4 Wavefront control

The wavefront was controlled using two deformable mirrors with 46 actuators across the pupil. The DM was modeled in PROPER using measured actuator surface influence functions of the Xinetics DM used in the HCIT. In these simulations, the DM actuators pistoned exactly by the commanded amount, unlike those in real DMs which typically have 10% piston errors (these errors eventually iterate out and a precision of 0.3 Angstroms can be achieved). The wavefront control algorithm used was Electric Field Conjugation (EFC; Give'on et al. 2007). Rather than using DM probing to sense the complex field at the final focus from intensity images, as is done in real systems, the computed field was used directly.

3.2. Milestone Requirements

***Milestone #2 Requirement:** Numerical models of each coronagraph will attempt to predict, after wavefront correction with deformable mirrors, a contrast of $\leq 10^{-10}$ in a realistically-aberrated optical system, quantified as the predicted mean level within a $2.5 - 18 \lambda/D$ annulus centered on the star over a $\lambda = 500 - 600$ nm bandpass. For all three coronagraphs, the optical system will be the same, including all simulated surface errors, from the primary mirror up to the second deformable mirror.*

Rationale:

Contrast

The brightness contrast relative to the star of an Earth-twin is $\sim 10^{-10}$ at visible wavelengths. The signal from the planet must be distinguished from the instrumentally-produced speckles inside the dark hole field around the star. We assume that this can be adequately accomplished in noisy images using post-processing (e.g. roll subtraction or reference star subtraction) if the peak pixel of the planet's point spread function is equal to the mean per-pixel speckle brightness.

Bandpass

In a real system, the contrast must be achievable over a broad ($\sim 20\%$) bandpass to either allow for deep integration in a wide bandpass filter or measurement using a spectrograph.

Field annulus

The inner radius was set where the occulter transmission was 50%. The outer radius was limited by the number of actuators on the DM.

Aberrations

The HBLC and vector vortex coronagraphs utilized the same optical system, excluding different focal plane and Lyot masks. The PIAA coronagraph required a different optical system between the second deformable mirror and final focusing optic. Each surface had realistic phase errors (from figuring and polishing) and amplitude errors (from coating non-uniformities). Synthetic two-dimensional error maps were generated from power spectral density (PSD) curves derived from actual optics and with error levels within current fabrication capabilities (off-axis aspheric mirrors for extreme UV lithography

with ~20 cm diameters have been fabricated with surface errors as low as 0.15 nm RMS). The primary mirror PSD was derived from actual large optics (>1.5 m) measurements, while the PSDs for the smaller elements were derived from measured extreme ultraviolet lithography optics. Interferometrically-measured surface error maps of the current best set of PIAA M1 and M2 optics were used for those surfaces, as they have aberration patterns unique to PIAA. They were modified as necessary to reach the required contrast.

The amplitude errors caused by coating non-uniformities were generally estimated, as accurate reflectivity maps are not common for such optics. The amplitude error distribution for the primary was derived from a set of measurement samples taken in a coating chamber at ITT. The chamber was designed to provide a very uniform coating. It was assumed that multiple smaller optics could be coated in such a chamber to provide even more uniform reflectivities (scaled to the beam size on the optics).

Table 1 lists the phase and amplitude RMS error levels for each optic. The synthetic error maps for the primary and 1st collimator mirrors are shown in Figure 1, with corresponding phase error PSDs plotted in Figure 2. We show in Appendix D and Figure 21 that the phase errors on the non-PIAA optics could be substantially increased and the target contrast would still be met.

Table 1. Optical wavefront errors

Optic	Phase error RMS	Amplitude error RMS $\times 10^{-5}$
<i>Common optics:</i>		
Primary	8.3 nm	22.0
Secondary	2.7 nm	4.2
Fold #1	1.5 nm	5.5
Collimator	1.9 nm	6.5
DM #1	5.0 nm	5.5
DM #2	5.0 nm	5.5
Final focusing lens	1.7 nm	2.9
Final fold	1.5 nm	5.5
<i>HBLC & VVC only:</i>		
Occulter focuser	2.0 nm	5.0
Pupil imager	2.3 nm	5.4
Fold #2	1.5 nm	5.5
<i>PIAA only:</i>		
PIAA collimator #1	1.8 nm	5.9
PIAA collimator fold	2.0 nm	5.5
PIAA collimator #2	1.7 nm	4.9
PIAA M1	varies	5.1
PIAA M2	varies	2.9
PIAA post-apodizer collimator	1.5 nm	1.0
PIAA occulter focuser	1.1 nm	2.5
Reverse PIAA M2	3.2 nm	3.5
Reverse PIAA M1	3.7 nm	3.0

The maps were filtered to limit errors to spatial frequencies of ≤ 48 cycles/diameter. This was done to prevent wrap-around of higher spatial frequency aberrations caused by the Fourier transforms used in the propagations. It also prevented aliasing of aberrations during propagation through PIAA, where low spatial frequencies are compressed to higher ones during wavefront remapping (the change in phase between two adjacent samples in a wavefront must be π radians or less to prevent numerical errors). The limit of 48 cycles provided enough high spatial frequency aberrations to introduce realistic amounts of frequency folded speckles, which are a primary limit on achievable broadband contrast; having aberrations at only those spatial frequencies that directly create speckles inside the dark hole (< 20 cycles/D) would not reproduce this important effect. On optics located between the forward PIAA M2 and reverse PIAA M2 mirrors, the errors were limited to ≤ 85 cycles/D. This is because higher spatial frequencies in PIAA-remapped coordinate space are remapped to lower ones by the reverse PIAA optics, as seen at the final image plane. The measured PIAA M1 and M2 maps (Figure 3) were filtered to ≤ 48 cycles/D.

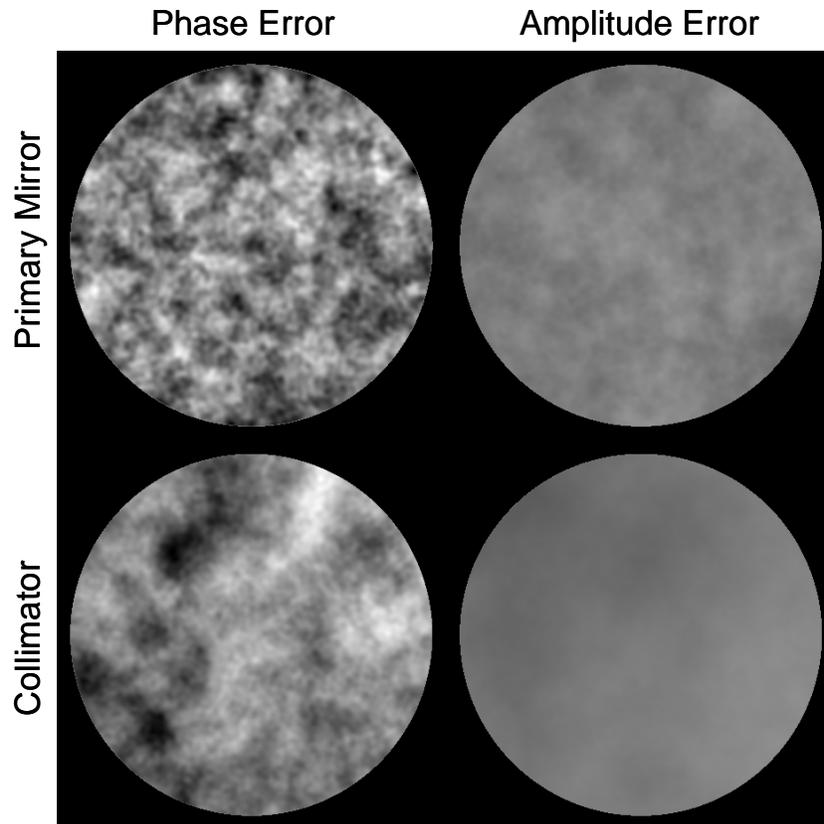


Figure 1 - Synthetic phase and amplitude error maps for optics used in the simulations. The errors of each type are shown at the same brightness scale.

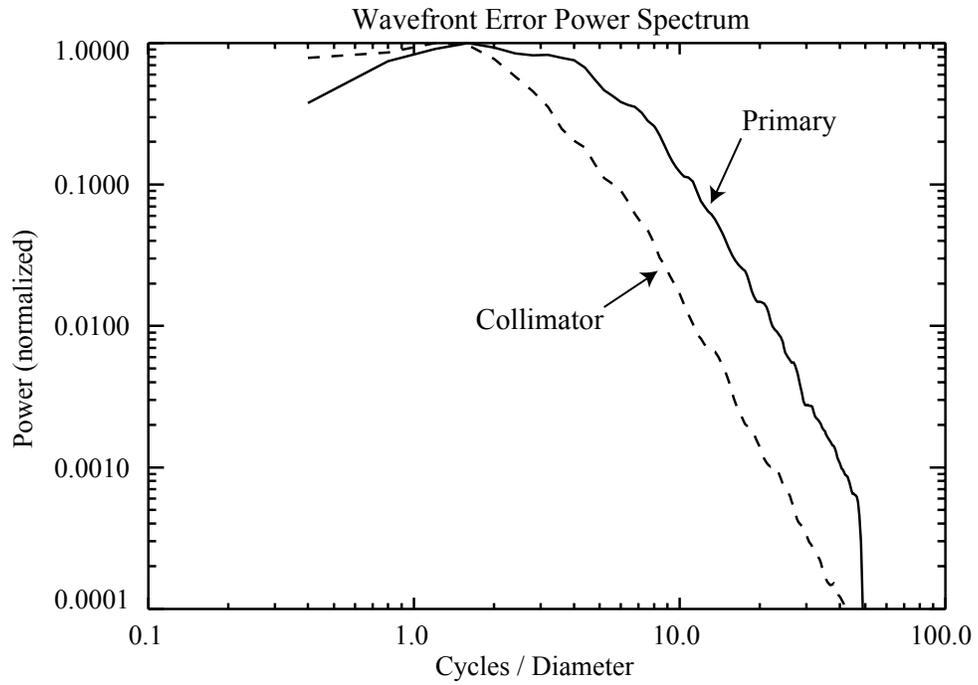


Figure 2 - Power spectral density curves of two synthetic optics used in the simulations. Note that the errors go to zero outside of 48 cycles/diameter.

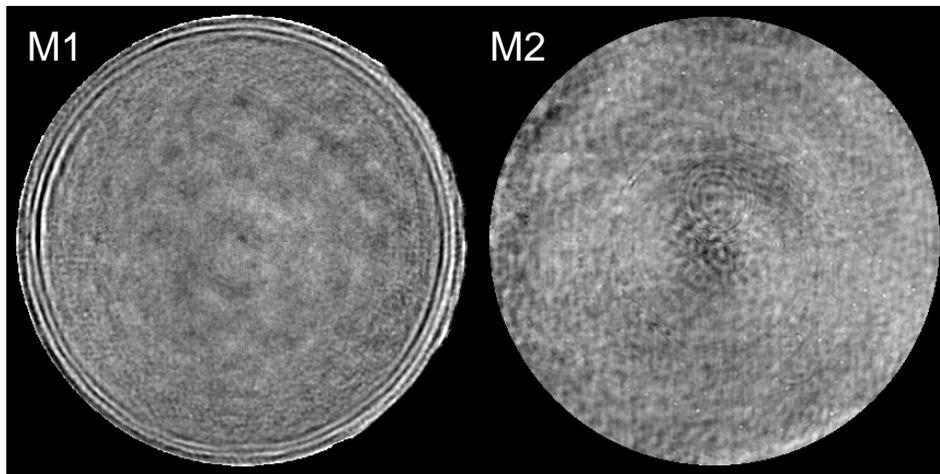


Figure 3 - Measured surface error maps for the PIAA 2 optics. M1 is shown between ± 30 nm of surface error and M2 between ± 15 nm.

3.3. Milestone Metrics

Milestone #2 Metric:

The mean contrast will be predicted within the specified annular region integrated over a $\lambda = 500 - 600 \text{ nm}$ passband.

Rationale:

The first step in simulating wavefront control in the model system was to set the 1st DM (the one conjugate to the primary mirror) to correct for phase errors in the wavefront at the occulter focal plane. In a real system, this would be done by using phase retrieval to derive the phase aberrations by defocusing the telescope and iteratively fitting the measured intensity image in the science detector using known constraints (pupil size, amount of defocus). In the simulations, the computed field at the occulter was extracted and Fourier transformed to the corresponding exit pupil, and then the DM surface was fit to the phase term. The opposite pattern was set on the DM, “flattening” the phase aberrations. This wavefront flattening reduced the speckle field intensity by about $10\times - 100\times$ and provided a starting point for EFC wavefront control.

To create the high-contrast dark hole, the entering wavefront was propagated through the system to the final focus where the complex-valued field was measured (sensed) within the specified annulus around the star. This was done separately at five monochromatic wavelengths that evenly spanned the broad passband. The measured fields, along with the DM response matrix (see Milestone #1), was used by EFC to determine the DM actuator settings that reduced the light within the annulus. The sense-control-propagate process was repeated until convergence was achieved (usually within 25-50 iterations, given the perfect functionality of the DMs).

Contrast at the end of each iteration was measured by:

- Converting each of the five complex-valued monochromatic fields (E_λ) to intensity: $I_\lambda = |E_\lambda|^2$
- Adding the monochromatic images together to create a broadband image
- Dividing the broadband image by the peak pixel value of the unocculted stellar PSF (accounting for the occulter radial transmission profile) to convert it to units of contrast
- Computing the mean of the pixels within the annulus

If, after convergence, the mean contrast was $>10^{-10}$, adjustments were made to the system layout, optical errors, and/or EFC settings (e.g., regularization), and the process repeated.

4. Success Criteria

The following items summarize the requirements and metrics detailed in Sections 2 and 3.

4.1 The optical systems modeled included the elements required to implement realistic coronagraphs and a telescope, including two deformable mirrors in series for wavefront control (46 actuators across the beam). These systems were identical up to and including the second deformable mirror.

4.2 The optics had realistic surface (phase) and coating (amplitude) errors derived from actual optics.

4.3 The wavefronts were propagated from surface to surface using the PROPER library for IDL along with the coronagraphic propagation and representation codes developed and verified in Milestone #1.

4.4 The fields of interest at the final focus were each an annulus of $2.5 \lambda_c/D \leq r \leq 18.0 \lambda_c/D$ radians centered on the star ($\lambda_c = 550$ nm, $D =$ diameter of the primary).

4.5 The wavelength bandpass was $\lambda = 500 - 600$ nm. It was evenly sampled at 5 monochromatic wavelengths to both sense the field and produce the polychromatic image.

4.6 Electric Field Conjugation was used to determine deformable mirror settings that produced a high-contrast field in each coronagraphic system integrated over a $\lambda = 500 - 600$ nm bandpass (uniform weighting across λ).

4.7 The mean contrast achieved within each dark hole field is reported, where contrast is the per-pixel intensity divided by the peak of the unocculted stellar point spread function.

4.8 The optical system parameters were modified as necessary to achieve a simulated system performance of 10^{-10} or better contrast within the dark hole. See the appendices in the Milestone #1 whitepaper for descriptions of the layouts, coronagraphs, and propagators.

5. Milestone #2 Results

The algorithms developed and verified in Milestone #1 were used to perform end-to-end modeling of the PIAA, HBLC, and VVC in a realistically-aberrated system. The final results are described in detail for each coronagraph in Appendices C, D, and E.

Some modifications were required to the initial systems. It was necessary to modify the layout of the PIAA system to include a 3rd DM between the forward PIAA optics and the occulter because the wavefront remapping compressed aberrations on surfaces after PIAA to spatial frequencies higher than those correctable by the DMs located prior to PIAA. It was also necessary to significantly reduce the errors on the PIAA optics (the original errors were from measured PIAA optics). The HBLC/VVC layout was unaltered. We note that neither the HBLC nor VVC included fabrication errors in their masks.

The final contrast maps are shown for the HBLC (Figure 4), VVC (Figure 5), and PIAA (Figure 6). These maps demonstrate that the Milestone #2 requirement of 10^{-10} mean contrast within the dark hole region over a $\lambda = 500 - 600$ nm bandpass was achieved in each case.

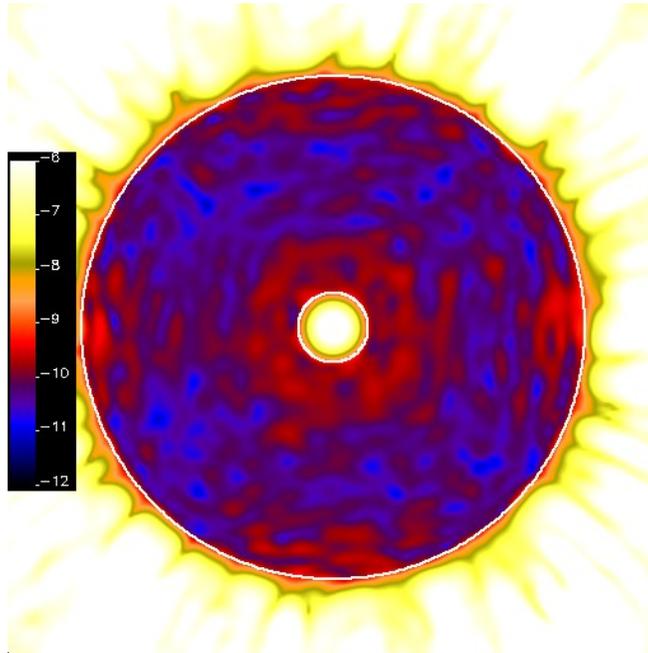


Figure 4 - Final broadband contrast over the field achieved using the HBLC. $\text{Log}_{10}(\text{contrast})$ is shown on the color scale. The inscribed circles mark the $r = 2.5$ and $18 \lambda/D$ boundaries of the contrast measurement region. The mean contrast in this region is 9.3×10^{-11} . The residual structure is largely due to the wavelength-dependent coating behavior included in the HBLC mask representation. See Appendix C for more details on the HBLC results and a radial contrast plot of this field (Figure 14).

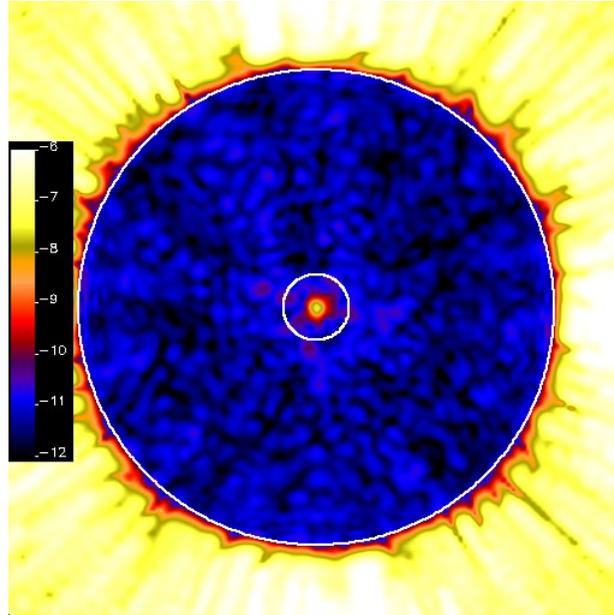


Figure 5 - Final broadband contrast over the field achieved using the VVC. $\text{Log}_{10}(\text{contrast})$ is shown on the color scale. The mean contrast in the dark hole region is 6.0×10^{-12} . It is possible that if coating-dependent thin-film effects, including interference, were included, there would be a greater, non-uniform level of residuals, like those seen for the HBLC or PIAA. See Appendix D for more details on the VVC results and a radial contrast plot of this field (Figure 21).

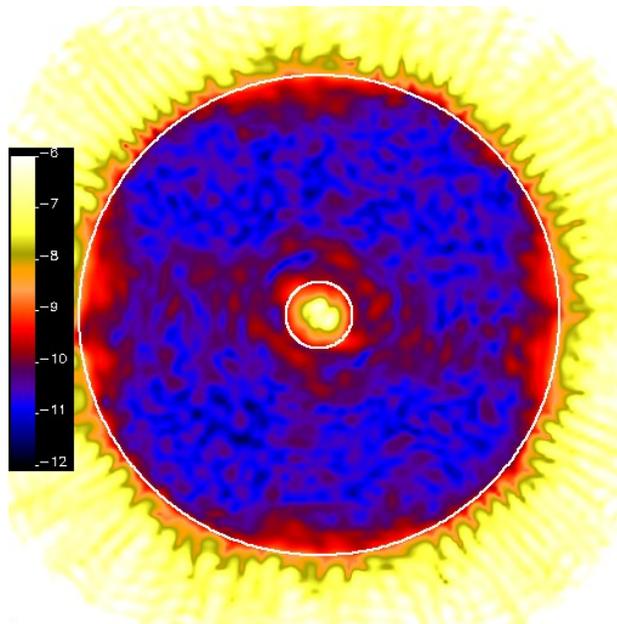


Figure 6 - Final broadband contrast over the field achieved using PIAA (M1 and M2 modified error maps, 3 DMs). $\text{Log}_{10}(\text{contrast})$ is shown on the color scale. The mean contrast in this region is 8.6×10^{-11} . See Appendix E for more details on the PIAA results and a radial contrast plot of this field (Figure 32).

6. Milestone Certification Data Package

The results of this study are being reported to and reviewed by the Exoplanet Exploration Program and NASA Headquarters. The documentation and data products that provide evidence that the requirements of this milestone have been met are:

- a. Documentation detailing the modeled optical system layouts, optical surface characteristics, and wavefront control algorithm parameters (see appendices). This includes the PROPER optical prescriptions (text files) and the error maps for each optic (FITS files) (these are being distributed in a separate data set and are not included in this document). The unique wavefront control behaviors of each coronagraph are discussed. Recommendations for future work, testbed and flight configurations, and computer resources (for wavefront control) are provided.
- b. The fields at the final image planes, in units of contrast, which demonstrate that the milestone contrast requirement has been met. These are being distributed as FITS files and shown as color-coded contrast maps in this document.

7. References

- Balasubramanian, K., Shaklan, S., Give'on, A., Cady, E., and Marchen, L., "Deep UV to NIR space telescopes and exoplanets coronagraphs: a trade study on throughput, polarization, mirror coating options and requirements", *Proc. SPIE*, 8151, 81511G (2011).
- Belikov, R., Pluzhnik, E., Connelley, M. S., Lynch, D. H., Witteborn, F. C., Cahoy, K., L., Guyon, O., Greene, T. P., McKelvey, M. E., "First results on a new PIAA coronagraph testbed at NASA Ames," *Proc SPIE* 7440, 74400J (2009).
- Bordé, P., and Traub, W., "High-Contrast Imaging from Space: Speckle Nulling in a Low-Aberration Regime", *Astrophysical Journal*, 638, 488 (2006).
- Elias, N., Bates, R., and Turner-Valle, J., "Polarization analysis for Terrestrial Planet Finder coronagraph designs", *Proc. SPIE*, 5555, 248 (2004).
- Give'on, A., Kern, B., Shaklan, S., Moody, D., Pueyo, L., "Broadband wavefront correction for high-contrast imaging systems," *Proc. SPIE*, 6691, 66910A (2007).
- Give'on, A., Kern, B., Shaklan, S., "Pair-wise, deformable mirror, image plane-based diversity electric field estimation for high contrast coronagraphy", *Proc. SPIE*, 8141, 815110 (2011).
- Groff, T., Carlotti, A., Kasdin, N., "Progress on broadband control and deformable mirror tolerances in a 2-DM system", *Proc. SPIE*, 8151, 81510Z (2011).
- Krist, J., Trauger, J., Moody, D., "Studying a simple TPF-C", *Proc. SPIE*, 6265 (2006).
- Krist, J., Moody, D., Mawet, D., Trauger, J., Belikov, R., Shaklan, S., Guyon, O., Vanderbei, R., "End-to-end simulations of different coronagraphic techniques," *Proc. SPIE*, 7440 (2009).
- Krist, J., Belikov, R., Mawet, D., Moody, D., Pueyo, L., Shaklan, S., Trauger, J., "Assessing the performance limits of internal coronagraphs through end-to-end modeling: Technology Milestone #1 Whitepaper", JPL Document D-66100 (2010).
- Krist, J., Belikov, R., Pueyo, L., Mawet, D., Moody, D., Trauger, J., Shaklan, S., "Assessing the performance limits of internal coronagraphs through end-to-end modeling: a NASA TDEM study," *Proc. SPIE*, 8151 (2011).
- Krist, J., "PROPER: An optical modeling program for IDL," *Proc. SPIE*, 6675, 66750P (2007).
- Krist, J., Trauger, J., Moody, D., "Studying a simple TPF-C", *Proc. SPIE*, 6265 (2006).
- Lieber, M., Neureuther, A., Ceperley, D., Kasdin, J., Hoppe, D., Eisenman, A., "Evaluating the end-to-end performance of TPF-C with vector propagation models: Part I – pupil mask effects," *Proc. SPIE*, 5905, 59050K (2005).
- Pueyo, L., and Kasdin, N., "Polychromatic Compensation of Propagated Aberrations for High-Contrast Imaging", *Astrophysical Journal*, 666, 609 (2007).

- Pueyo, L., Kay, J., Kasdin, N., Groff, T., McElain, M., Give'on, A., Belikov, R.,
“Optimal dark hole generation via two deformable mirrors with stroke
minimization”, *Applied Optics*, 48, 6296 (2009).
- Shaklan, S., Green, J., “Reflectivity and optical surface height requirements in a
broadband coronagraph. 1. Contrast floor due to controllable spatial frequencies,”
Applied Optics, 45, 5143 (2006).
- Traub, W., et al., “TPF-C: status and recent progress,” *Proc. SPIE*, 6268, 62680T (2006).
- Trauger, J., Stapelfeldt, K., Traub, W., Henry, C., Krist, J., et al. “ACCESS: a NASA
mission concept study of an actively corrected coronagraph for exoplanet system
studies,” *Proc. SPIE*, 7010, 701029 (2008).

Appendix A: Layouts

The schematic layouts of the two coronagraphic systems are shown here.

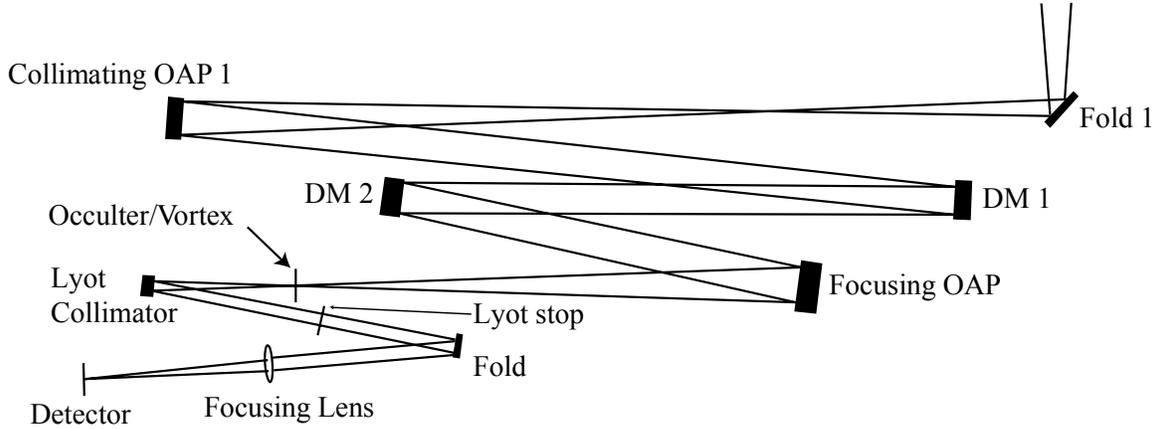


Figure 7 - Schematic optical layout for the HBLC/VVC. Not shown are the telescope primary and secondary mirrors that feed Fold 1 in the upper right.

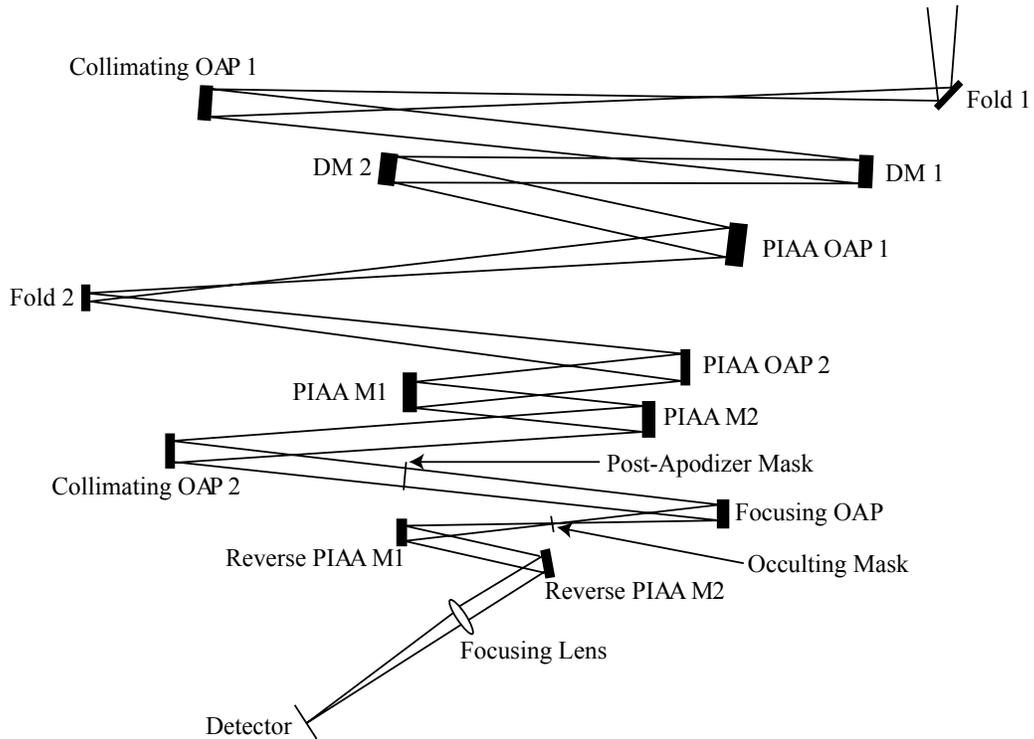


Figure 8 - Schematic optical layout for the PIAA coronagraph. Not shown are the telescope primary and secondary mirrors that feed Fold 1 in the upper right.

Appendix B: Wavefront control

Wavefront sensing

The first step in wavefront control is wavefront sensing, which entails measuring both the phase and amplitude. In a real system, the wavefront must be sensed by some means, either by using an interferometer or by deriving the field by perturbing it in some known manner and measuring the intensity change. In most current coronagraph testbeds like HCIT, a known pattern is placed on the DM and the resulting change in intensity in the final image plane is observed. Using a series of intensity measurements made with different DM settings, the wavefront can be derived (Give'on et al. 2011). The intensity measurements are typically made at a number of wavelengths using medium or narrow band filters to characterize the wavelength-dependent aberrations. Work is underway on trying to reduce the number of wavelengths required (Groff et al. 2011).

In our simulations, we avoided wavefront sensing altogether and simply used the complex field computed by the programs.

Wavefront control using Electric Field Conjugation

We used the iterative Electric Field Conjugation (EFC) method (Give'on et al. 2007) to determine the DM settings that reduced the intensity of the speckles within the targeted field around the central source. EFC is actually a fairly simple means of solving for the DM solution, but it requires significant computational resources. It assumes that the wavefront control process is linear by the equation:

$$Ax = B$$

where A is the DM response matrix (two-dimensional) that describes how the image field changes for a piston of each actuator, x is the vector of DM actuator settings, and B is the complex field at the image plane. Because this is a linear approximation and propagation is actually nonlinear, EFC is an iterative process.

The heart of EFC is the DM response matrix. It describes how each actuator changes the complex field at each pixel in the dark hole region. Whether it is used in correcting a real system or a simulation, this matrix is computed using a model. Each actuator on each DM is pistoned by some amount and then the wavefront is propagated through the entire system. Each coronagraph alters the wavefront change in a different manner (Figure 9).

The complex field values within the dark hole field are stored at each wavelength used to sense the wavefront. The 2-D array has the dimensions $[N_{act} \times N_{dm}, N_{pix} \times N_{\lambda}]$, where N_{act} is the number of active actuators on a single DM, N_{dm} is the number of DMs in the system, N_{pix} is the number of pixels within the dark hole field, and N_{λ} is the number of wavelengths at which the wavefront is sensed (or in our case, computed).

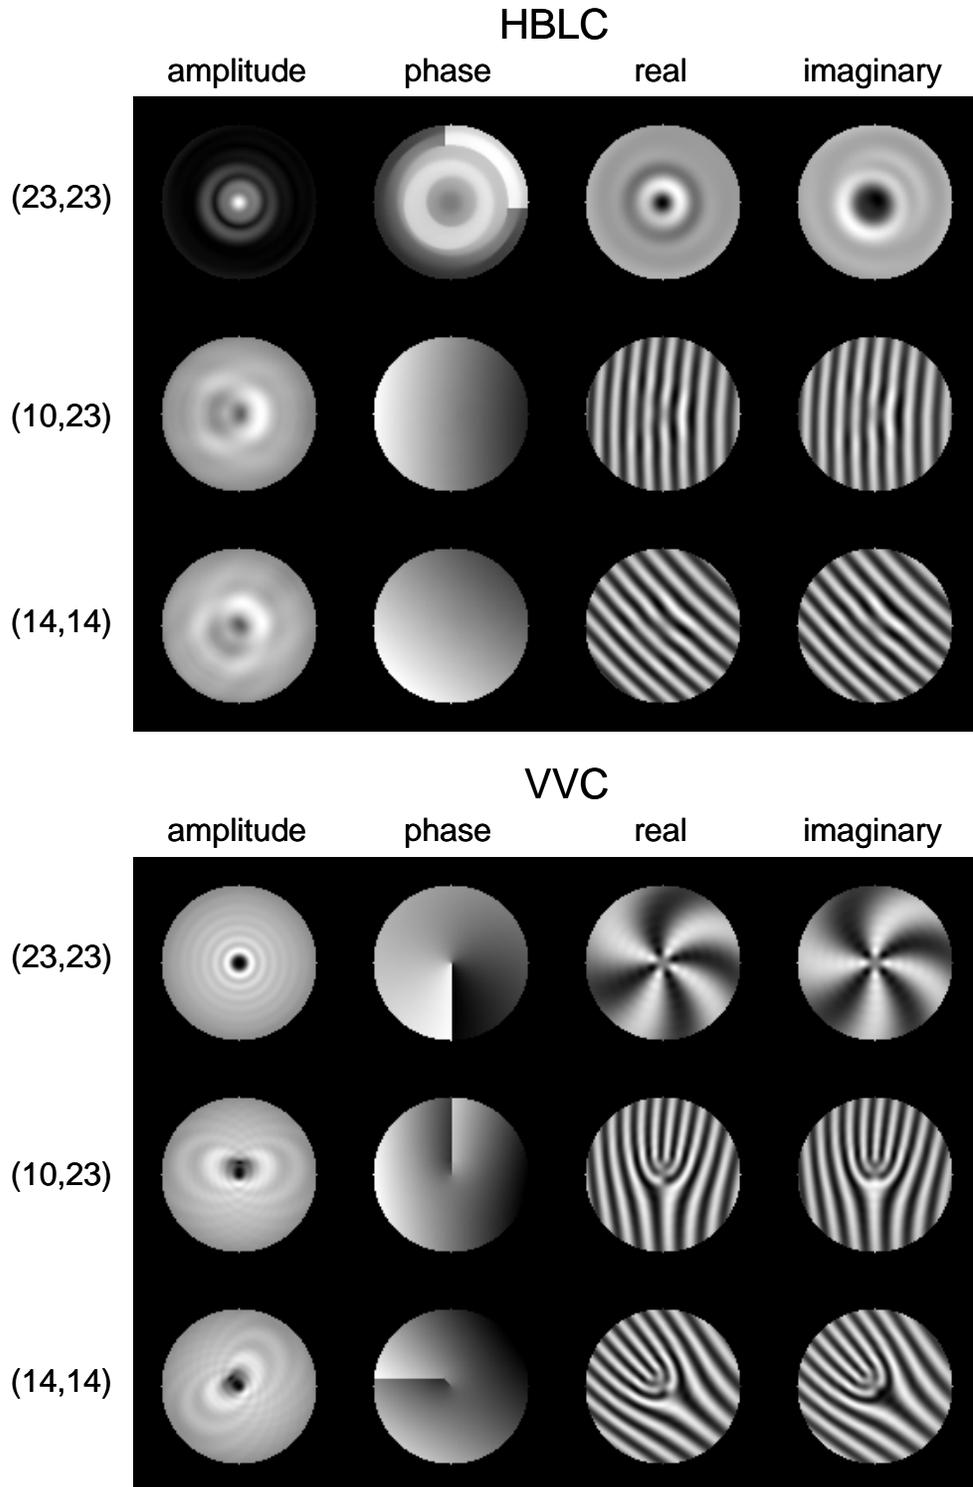


Figure 9 - Changes in the electric field distribution ($\lambda = 550$ nm) in the final image plane when the specified actuator on DM #1 is pistoned. Only the regions within the $18 \lambda/D$ outer radius of the dark hole are shown. The actuator at (23,23) in the HBLC is largely masked by the central spot in the Lyot stop, causing the unusual appearance in the field (it also has remnant errors from phase unwrapping). The wavefront is centered between DM actuators 23 & 24. (Continued on next page)

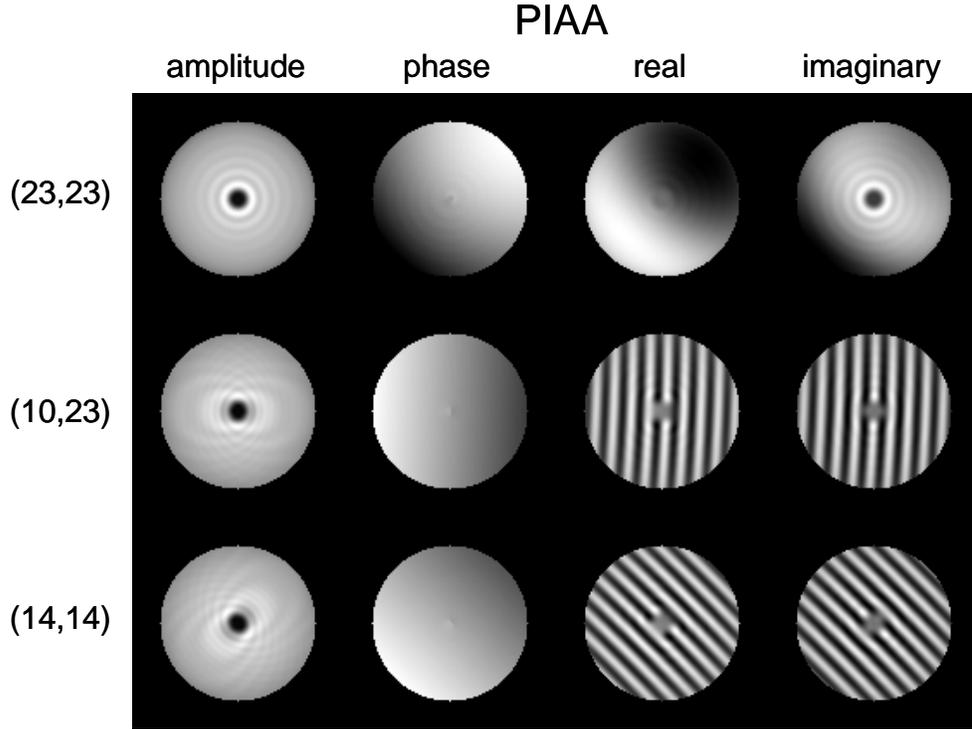


Figure 9 (continued).

Note that the number of pixels is dependent on the sampling. In our trials, we used a sampling of $0.4 \lambda/D$ at the shortest wavelength of the passband (500 nm).

Each image plane pixel can be given a weight, if one wishes to “pound” on specific parts of the image. This is done by simply multiplying the corresponding elements in the A and B matrices. One could, for instance, place greater emphasis on pixels near the inner working angle by giving those pixels a weight of 1.0 and a lower one to others.

The matrix in this form is still not yet suitable for practical use. It would allow large actuator strokes that would cause the solution to diverge. To limit the strokes to more stable conditions (damping), regularization is applied to the matrix. The first step is to reformat each element in the matrix from a complex value to two reals, doubling the number of elements but keeping the same size in bytes: $[N_{act} \times N_{dm}, 2 \times N_{pix} \times N_{\lambda}]$. Next, at the end of each column, a column vector of size $N_{act} \times N_{dm}$ is added on. The n -th element in this vector corresponding to the n -th column (actuator) is set to the regularization value (which is much less than 1.0, and the damping increases with the regularization value). So, the size of the DM response matrix is

$$[N_{act} \times N_{dm}, 2 \times N_{pix} \times N_{\lambda} + N_{act} \times N_{dm}].$$

A zero vector of $N_{act} \times N_{dm}$ elements is added on to the actuator stroke vector x to keep array sizes compatible.

We note that besides EFC there are other algorithms for determining the optimal DM settings to create a dark hole, such as energy minimization (Bordé & Traub 2006) and stroke minimization (Pueyo et al. 2009).

Inverting the response matrix

The goal of EFC wavefront control is to determine the actuator settings (x) that minimize the energy in the dark hole (B) given the known change in the wavefront caused by the DM actuators (A). The previous equation is thus reordered and negated:

$$x_{correction} = -A^{-1}B$$

Typically, A is not a square matrix and thus is not directly invertible. A pseudo-inversion can be done using singular value decomposition (SVD) that provides a solution in the least-squares sense.

The DM response matrices used in this study were large, over 2 GB in size, and SVD is computationally demanding. We initially used the SVD routines provided by IDL (Interactive Data Language), our primary computational environment. However, these routines took over 40 minutes to “invert” a matrix. We eventually used the SVD routine in the optimized math library provided by Advanced Micro Devices (AMD); Intel also offers a similar library, though it is not free like AMD’s. We generated a response matrix in IDL, wrote it to disk, then ran a separate C program linked to the optimized library that read it in, ran SVD, and wrote out the result. The inversion took only a few minutes this way.

Flattening the wavefront

With the inverted response matrix in hand, EFC can be run. However, it is advisable to first “flatten” the wavefront. This is done by deriving the phase error in the wavefront incident on the focal plane mask and setting the pupil-plane conjugate DM to negate (flatten) the errors. In a real system the phase errors could be measured using phase retrieval or some other wavefront sensing technique. In our simulations, we simply propagated a wavefront through the aberrated system (with flat DM settings) up to, but not including, the focal plane mask. The beam was then propagated to a pupil where the phase error map was extracted (Figure 10). The DM surface was then fitted to this map and negated. Flattening the wavefront removes the largest errors in the system and provides a relatively nominal starting point from which to use EFC (Figure 11). This keeps the system in a relatively linear regime of a non-linear process.

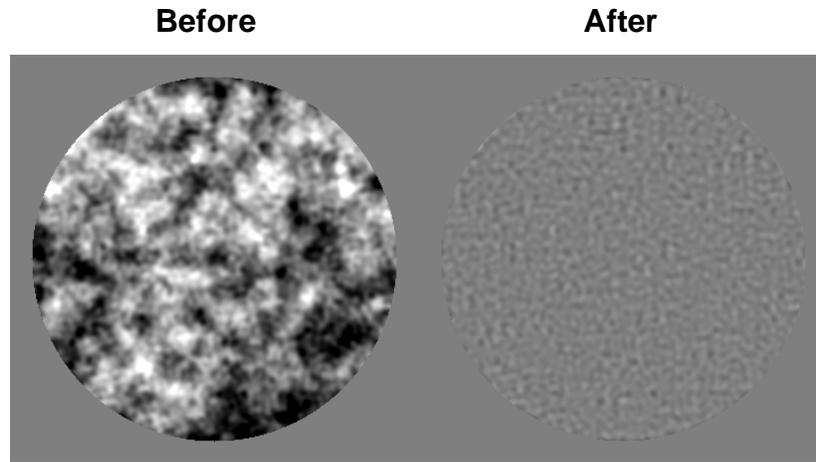


Figure 10 - (Left) Wavefront phase error in the pupil conjugate to the focal plane (9.4 nm RMS). (Right) The wavefront after flattening with the DM (1.05 nm RMS). Both are shown at the same scale (± 20 nm wavefront error).

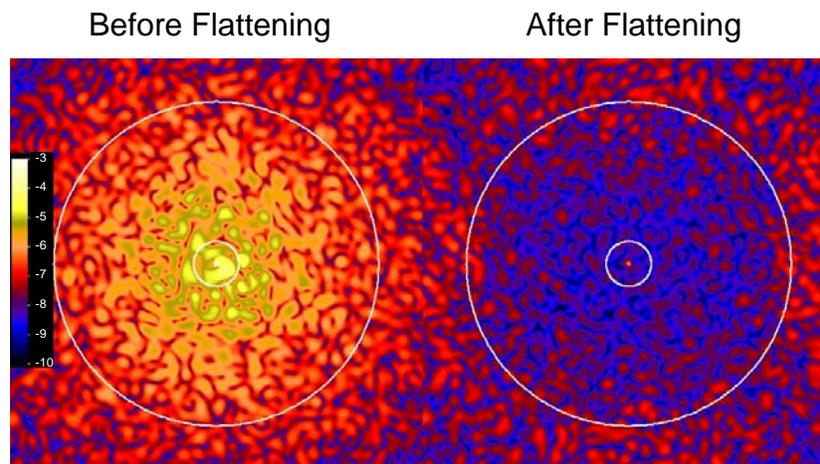


Figure 11 - Monochromatic (550 nm) contrast in the VVC final image plane before (left) and after (right) flattening the wavefront. The scale is shown in $\log_{10}(\text{contrast})$. Note that the intensity stretch here is different from the other contrast maps in this document due to the higher uncorrected speckle level.

Appendix C: Results for the Hybrid BandLimited Coronagraph (HBLC)

Coronagraph description

The HBLC is described in detail in the Milestone #1 report, so we briefly summarize it here. The HBLC uses a combination of three components to implement diffraction pattern suppression:

- A focal plane mask with circularly-patterned gradient amplitude transmission (metallic) and phase modifying (dielectric) coatings
- A Lyot stop at a pupil image subsequent to the focal plane mask that blocks the remaining diffracted light concentrated around the edge of the pupil (and in this version, at the center as well using a small opaque spot)
- Two deformable mirrors in sequence to provide phase and amplitude wavefront control

All three components were iteratively optimized in an aberration-free system to create a diffraction-suppressed dark hole with the specified inner and outer radii and bandpass. The wavelength-dependent amplitude and phase dependent properties of the metal (nickel) and dielectric (cryolite) coatings were included as determined by thin-film calculations.

Wavefront control parameters and generating the DM response matrix

The DM response matrix used by EFC was generated by propagating individual actuator pokes at each wavelength sampling the bandpass and recording the complex field computed at the final image plane. The HBLC masks are circularly symmetric, and the DM patterns are 4-fold symmetric, allowing the 1-for-8 sector shortcut to be used when generating the response matrix.

The DM actuators were pistoned by 0.5 nm. As a test, after convergence of an EFC run using this value, a new response matrix was generated using 0.1 nm pistons and the control process was continued. However, this provided no further improvement.

We experimented with different image weighting schemes. First, we used the same scheme applied to the vector vortex coronagraph control (described later), where a uniform weighting of 1.0 was applied over $r = 0.9 - 18 \lambda_c/D$ and 0.2 for $r < 0.9 \lambda_c/D$. This produced, after 50 iterations with a regularization of 0.5×10^{-5} , a mean contrast of 9.8×10^{-11} over $r = 2.5 - 18 \lambda_c/D$ and 2.2×10^{-10} over $r = 2.5 - 3.5 \lambda_c/D$. An alternative scheme set the weight to 1.0 within a $r = 2.2 - 18 \lambda_c/D$ annulus and 0.0 elsewhere, resulting in a mean contrast of 9.4×10^{-11} from $2.5 - 18 \lambda_c/D$ and 2.4×10^{-10} from $2.5 - 3.5 \lambda_c/D$. We investigated more localized weighting as well, including adjusting the weights to vary linearly with inverse radius, so that pixels near the IWA were given greater emphasis. This did not provide a better result than using uniform weighting. We also tried adjusting the weighting after having converged using the uniform weights, so that pixels above 10^{-10} contrast were given 4 – 8 times greater weighting than the others.

This provided a few percent improvement at the IWA but with a corresponding reduction in the contrast over the entire field. Because it produced the best contrast over the full field, we chose to use uniform weighting over a $2.2 - 18 \lambda_c/D$ annulus (the contrast was still evaluated over $2.5 - 18 \lambda_c/D$).

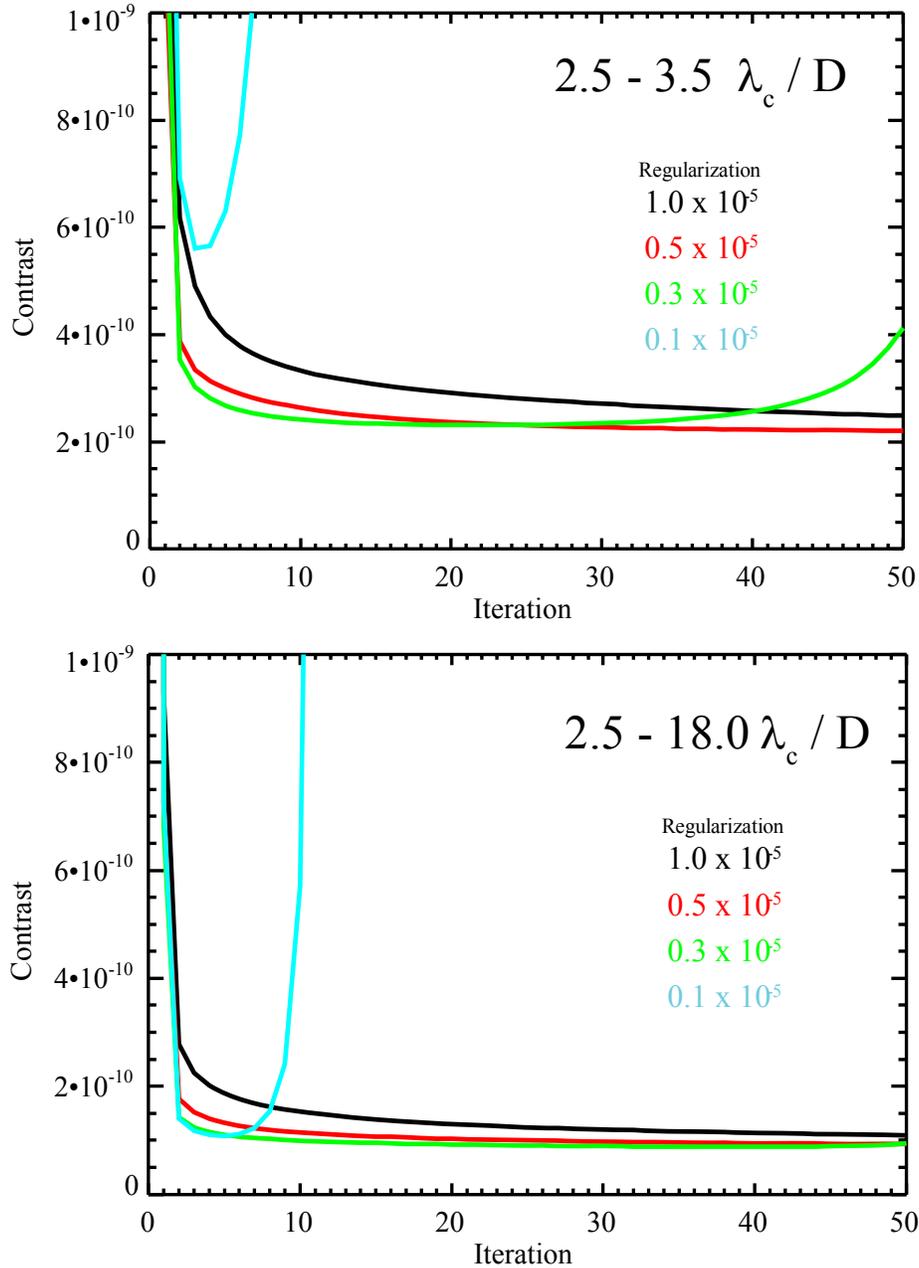


Figure 12 - Plots of mean broadband contrast with the HBLC versus EFC iteration, measured within an annulus of (top) $r = 2.5 - 3.5 \lambda_c/D$ and (bottom) $2.5 - 18.0 \lambda_c/D$ radians ($\lambda_c = 550$ nm). Plots are shown for different values of regularization used in the corresponding DM response matrices. Iteration 0 is the result using the initial “flattening” settings on the 1st DM, before running EFC.

Trial EFC runs used different regularization values in the DM response matrix. As shown in the plots in Figure 12, there appeared to be a “sweet spot” using a regularization of 0.5×10^{-5} . Lower regularization values (less actuator damping) provided rapid improvements in the first few iterations but the solution would eventually diverge. A higher value (more damping) slowed the convergence rate. After 150 iterations, a regularization value of 1.0×10^{-5} reached mean contrasts of 2.4×10^{-10} and 9.7×10^{-11} (for $r = 2.5\text{-}3.5 \lambda/D$, $2.5\text{-}18 \lambda/D$, respectively). After only 50 iterations, a regularization value of 0.5×10^{-5} provided contrasts of 2.4×10^{-10} and 9.4×10^{-11} for the same annuli, which we adopted here.

The target dark hole contained 7596 pixels and there were 1804 active actuators on each DM. The DM response matrix for 5 wavelengths was 2.1 GB in size.

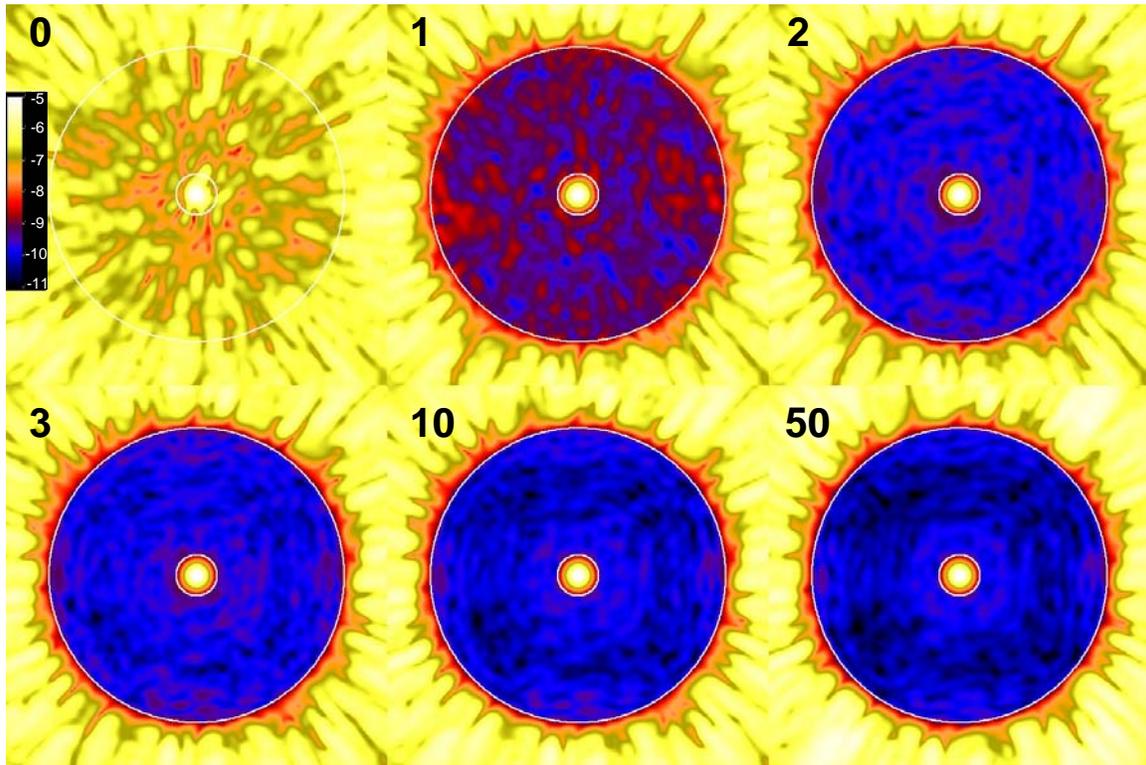


Figure 13 - Broadband contrast maps versus EFC iteration for the HBLC (regularization = 0.5×10^{-5}). The color scale shows $\log_{10}(\text{contrast})$. Iteration 0 is the result using the initial “flattening” settings on the 1st DM, before running EFC. The inscribed circles denote radii of 2.5 and $18.0 \lambda_c/D$ radians.

Final results

The final result used a DM response matrix with a regularization of 0.5×10^{-5} and uniform image weighting from $2.2 - 18 \lambda_c/D$. EFC was used for 50 iterations (Figure 13 & Figure 14), at the end of which the mean contrast was 9.3×10^{-11} over $r = 2.5 - 18 \lambda_c/D$ (meeting the milestone requirement of 10^{-10}) and 2.2×10^{-10} over $r = 2.5 - 3.5 \lambda_c/D$. This was achieved without needing to modify the system from the initial conditions.

As shown in Figure 15, there is considerable variation in contrast with wavelength. With an ideal coronagraph, one might expect that the contrast would be best near the central wavelength and degrade away from it. However, the combination of the wavelength-dependent amplitude-and-phase-modifying mask and DM settings in the HBLC results in a different behavior.

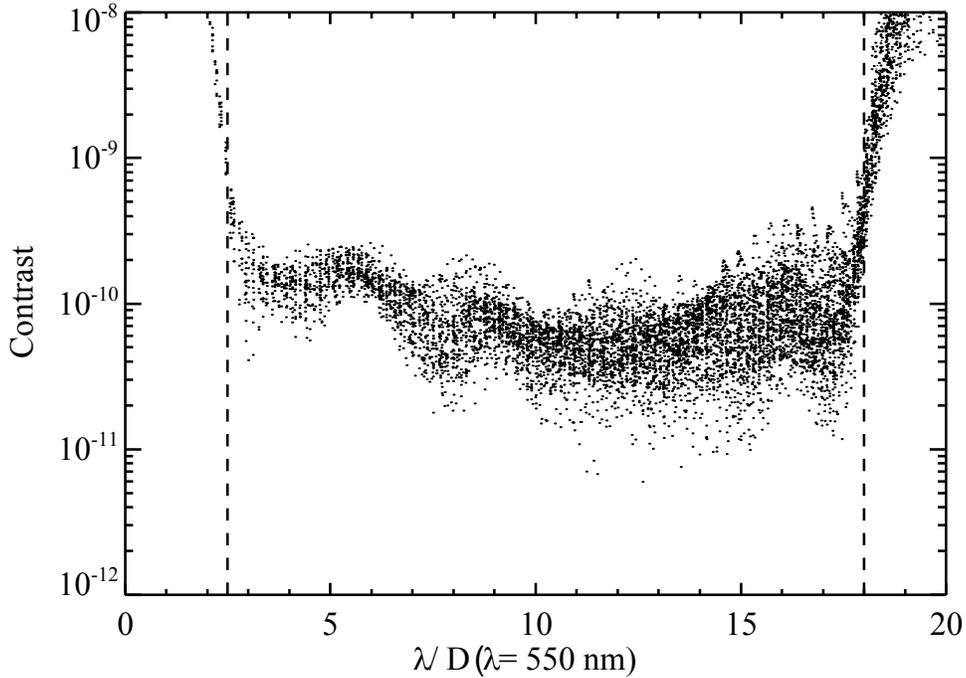


Figure 14 - Contrast versus radius for each pixel in the dark hole region of the final HBLC EFC solution. The inner and outer working angles of 2.5 and 18 λ/D are indicated. The mean contrast in the dark hole was 9.3×10^{-11} .

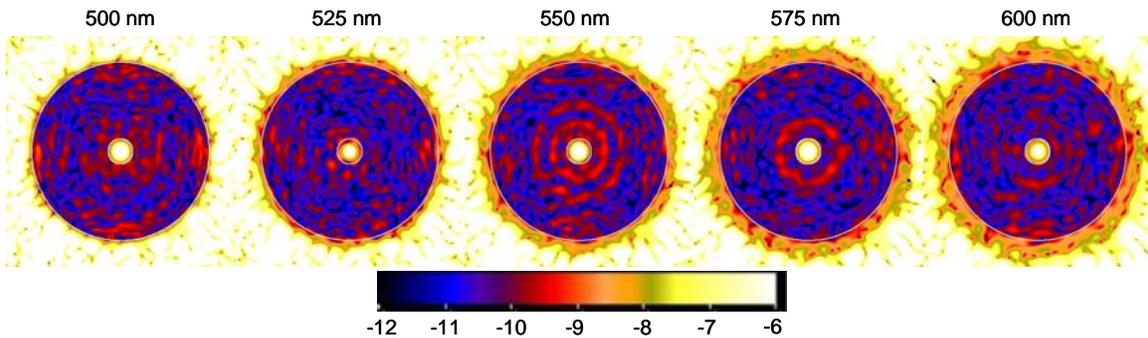


Figure 15 - Final HBLC broadband contrast decomposed into monochromatic contrasts over the field. $\text{Log}_{10}(\text{contrast})$ is shown on the color scale (note that the contrast range is different here than for the other contrast maps in order to emphasize the variations with wavelength). The inscribed circles mark the $r = 2.5$ and 18 λ/D boundaries of the contrast measurement region. At the central wavelength (550 nm), the mean contrast is 9.4×10^{-11} over the dark hole.

Sensitivity to outer radius extent

A coronagraph works best when the PSF is as close to perfect as possible. For the HBLC and VVC, that requires that the PSF is largely intact and extends to a large outer radius. Truncation of the PSF using a field aperture will both introduce additional diffraction (the PSF wings diffract into the pupil interior) while also blocking out spatial frequencies necessary to reconstruct a clean image of the pupil at the Lyot stop (similarly, it is necessary to have a clean pupil image at the PIAA M1 optic).

As an additional experiment to demonstrate this, we used our final HBLC dark hole solution and introduced a circular field aperture at the intermediate focus located between the first fold mirror and the DM collimation optic. The wavefront solution was not updated after introducing the aperture. The contrast maps are shown in Figure 16 for aperture masks with unlimited, 48, 96, and 144 λ_c/D outer radii. The mean contrasts in the dark hole are, respectively, 9.3×10^{-11} , 4.0×10^{-10} , 1.0×10^{-10} , and 9.3×10^{-11} . It is likely that the VVC and PIAA have similar sensitivities to outer radius truncation. This indicates that an outer radius of at least $\sim 130 \lambda/D$ is required for 10^{-10} contrast.

We note that it is possible to gain back some of the loss by using wavefront control to suppress some of the diffracted light introduced by the aperture. However, there is a limit to how well this can be done due to chromatic effects.

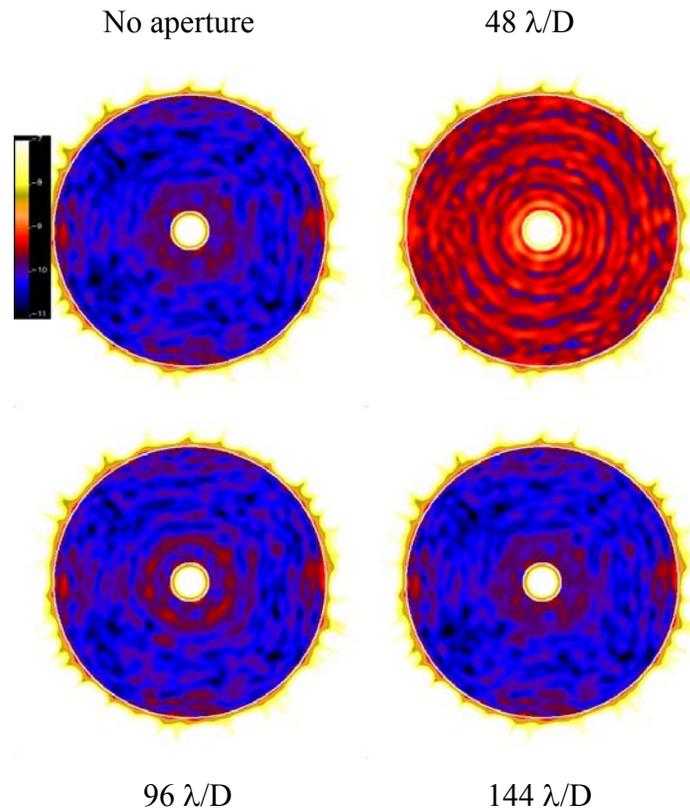


Figure 16. Comparison of fields after the introduction of an aperture stop at an intermediate focus before the HBLC occulter (aperture outer radius is given in units of λ_c/D). Contrast is displayed from 10^{-11} to 10^{-7} .

Appendix D: Results for the Vector Vortex Coronagraph (VVC)

Coronagraph description

As detailed in the Milestone #1 report, the VVC focal plane mask used in these simulations had a charge of 4 (the resulting phase spiral rotates by $4 \times 2\pi$ radians), resulting in 4th order aberration rejection. An opaque circular spot of $r = 0.25 \lambda_c/D$ at the center of the mask covered the “confusion zone”, the region where, in a real mask, there would be large fabrication errors near the central singularity. A simple, circular Lyot stop was used at the reimaged pupil plane after the VVC mask with a clear aperture of 90% the diameter of the beam.

As discussed in the Milestone #1 report, the broadband performance of a mask with less than 3 layers would be poor ($\sim 10^{-4}$ - 10^{-8} contrast) due to “chromatic leakage”. A real mask would require 5 layers of polymer and operate in a single polarization channel (as we assumed in all of our simulations). Such a design would ideally create a contrast floor of $< 10^{-12}$ (this does not include reflections/scattering from interference effects, which have not been computed). We assumed that we were using a 5-layer mask, and because the predicted leakage term was far below our contrast requirement, we did not explicitly include it in our simulations.

Wavefront control parameters and generating the DM response matrix

The VVC DM response matrix was generated in the same manner as for the HBLC. Because of the charge=4 phase ramp at the VVC mask, the system has 4-fold symmetry, so a 1-for-4 quadrant shortcut was used (actuators were poked only in one 90° sector of the DM). The DM actuators were pistoned by 0.5 nm. We used a uniform weighting of 1.0 over $r = 0.9 - 18 \lambda_c/D$ and 0.2 for $r < 0.9 \lambda_c/D$. The 5-wavelength response matrix was 2.2 GB.

We ran EFC separately using response matrices with the regularization set to 1×10^{-5} and 0.5×10^{-5} . As shown in Figure 17, the latter value resulted in faster convergence, especially near the inner working angle, and a slightly better contrast (6.0×10^{-12} over the full dark hole field compared to 8.6×10^{-12}).

Final results

After 25 iterations of EFC (Figure 18, Figure 19) we reached mean contrasts of 6.0×10^{-12} over $r = 2.5 - 18 \lambda_c/D$ (easily meeting the Milestone #2 requirement) and 1.5×10^{-11} over $r = 2.5 - 3.5 \lambda_c/D$.

Unlike the HBLC, the contrast is best at the central wavelength (Figure 20). This difference is likely due to the representation of the mask being wavelength independent (excluding the fact that the central spot does not scale with wavelength). If a more realistic representation of the VVC mask were to be used (i.e., one including interference effects between coating layers), then the wavelength dependence would likely differ.

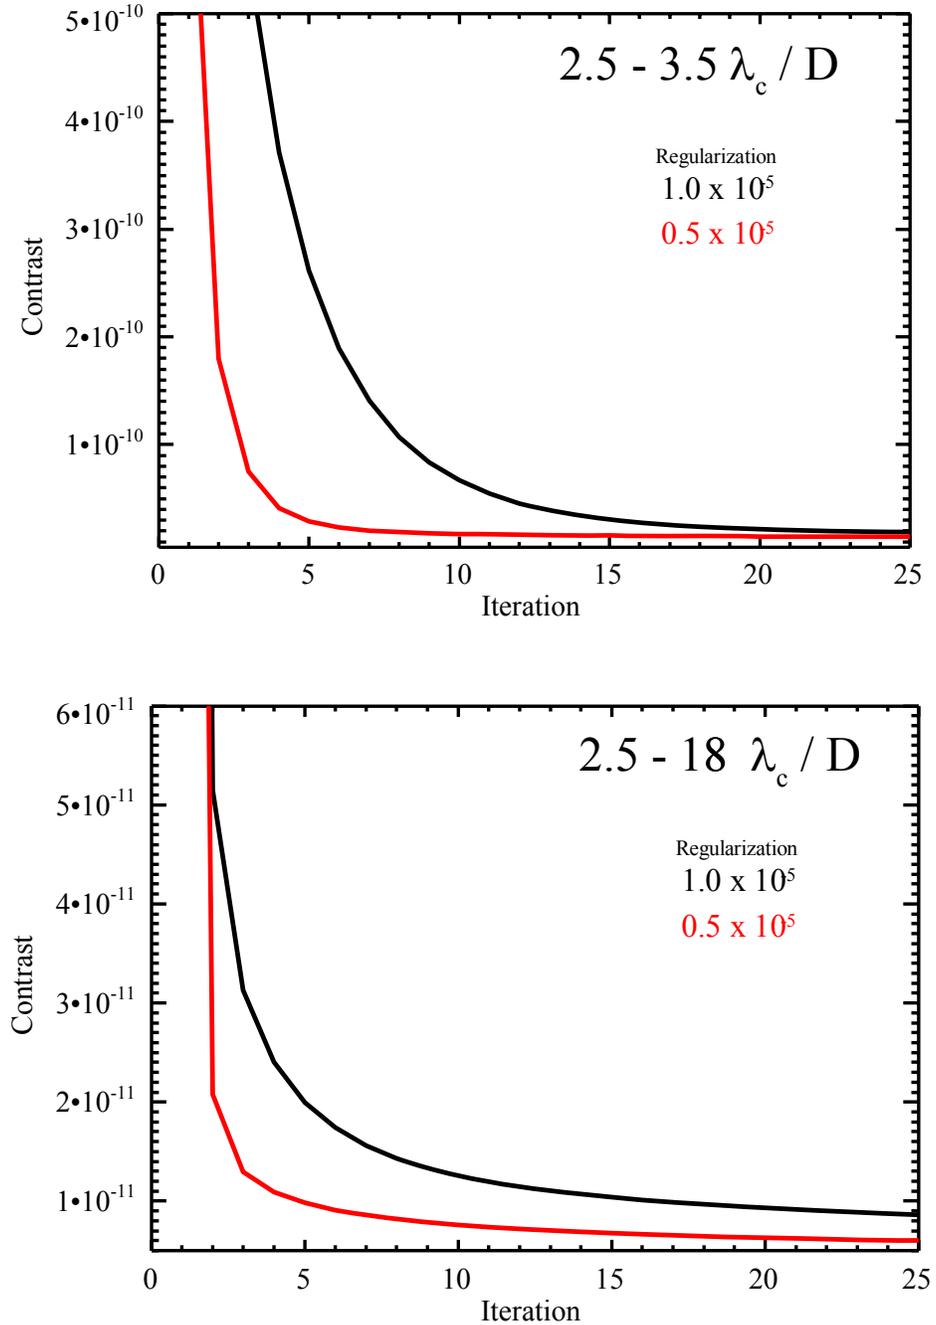


Figure 17 - Plots of mean broadband contrast with the VVC versus EFC iteration, measured within an annulus of (top) $r = 2.5 - 3.5 \lambda_c / D$ and (bottom) $2.5 - 18.0 \lambda_c / D$ radians ($\lambda_c = 550 \text{ nm}$). Plots are shown for different values of regularization used in the corresponding DM response matrices. Iteration 0 is the result using the initial “flattening” settings on the 1st DM, before running EFC.

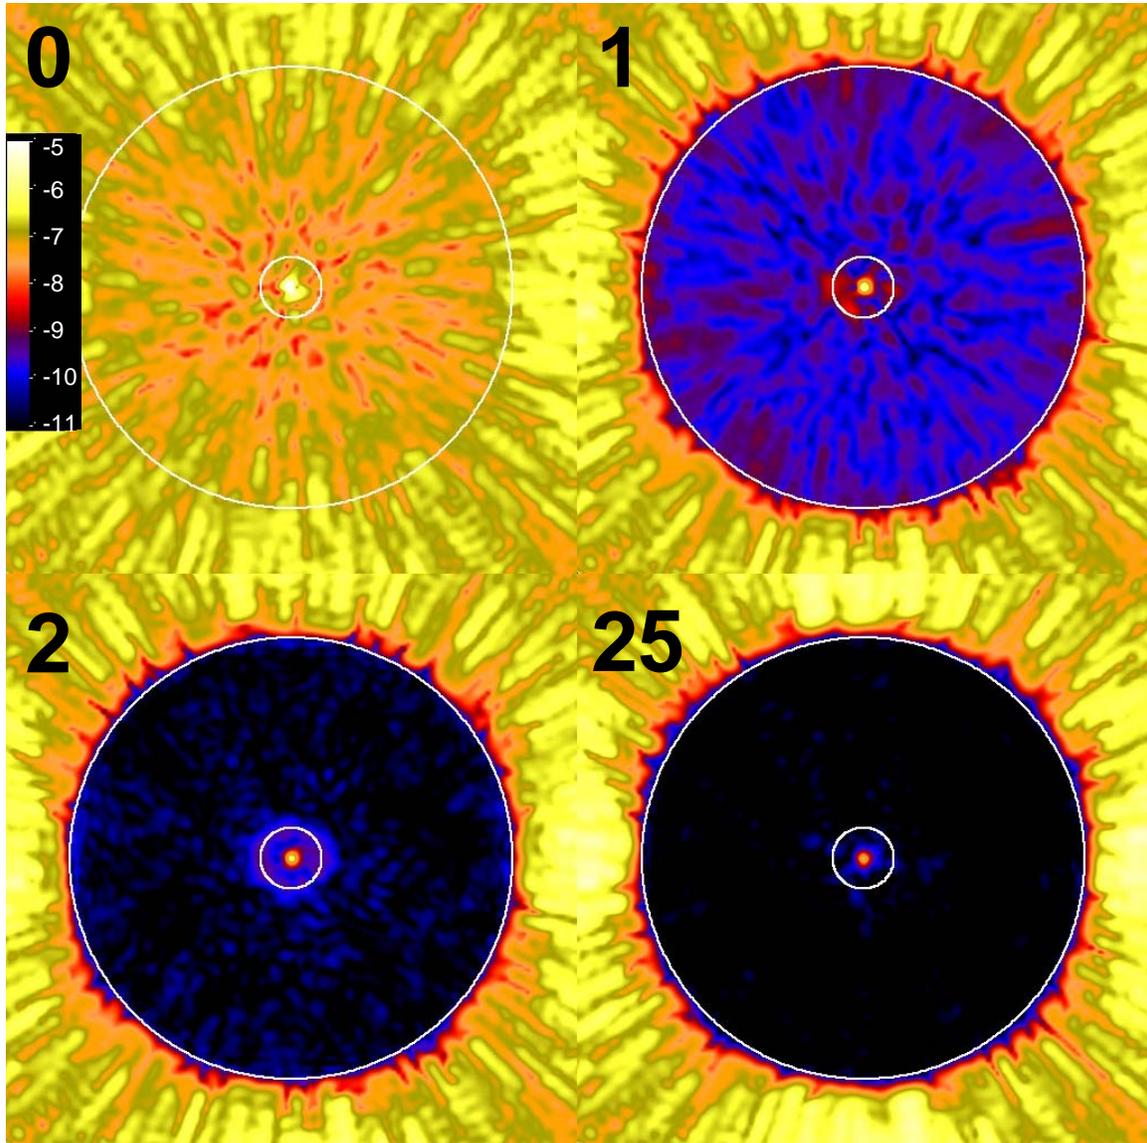


Figure 18 - Broadband contrast maps versus EFC iteration for the VVC (regularization = 0.5×10^{-5}). The color scale shows $\log_{10}(\text{contrast})$. Iteration 0 is the result using the initial “flattening” settings on the 1st DM, before running EFC. The inscribed circles denote radii of 2.5 and 18.0 λ_c/D radians.

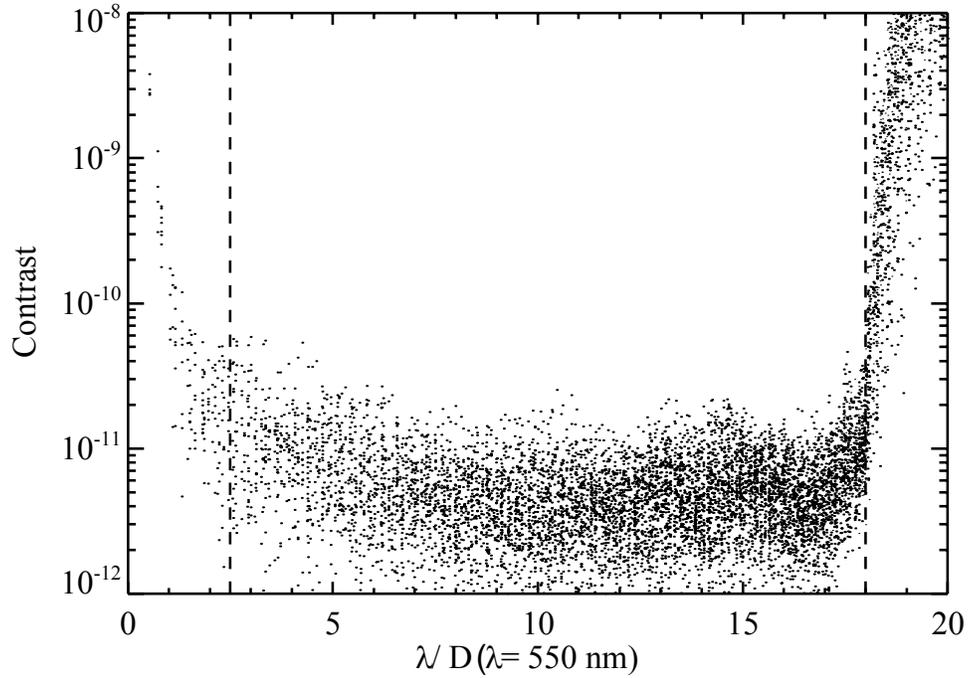


Figure 19 - Contrast versus radius for each pixel in the dark hole region of the final VVC EFC solution. The inner and outer working angles of 2.5 and 18 λ/D are indicated.

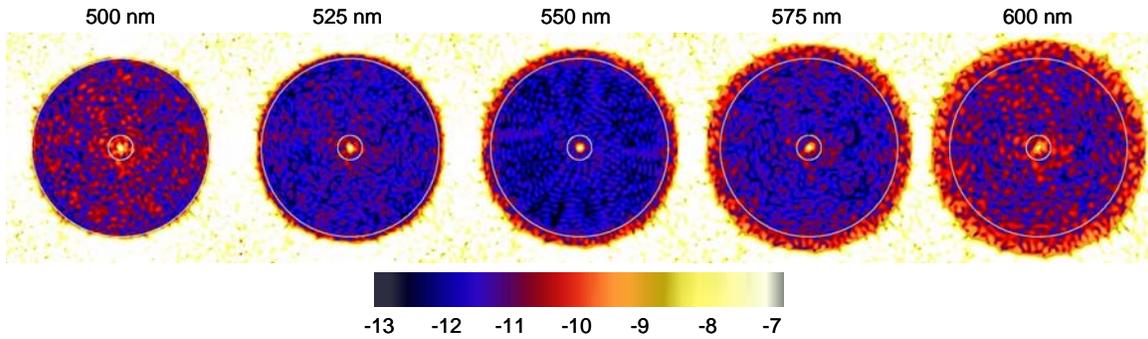


Figure 20 - Final VVC broadband contrast decomposed into monochromatic contrasts over the field. $\text{Log}_{10}(\text{contrast})$ is shown on the color scale (note that the contrast range is different here than for the other contrast maps in order to emphasize the variations with wavelength). The inscribed circles mark the $r = 2.5$ and 18 λ/D boundaries of the contrast measurement region. At the central wavelength (550 nm), the mean contrast is 1.8×10^{-12} over the dark hole.

Performance limitations

The VVC used in these simulations easily met the milestone mean contrast requirement over the full dark hole, and its performance at the inner working angle was also excellent.

The VVC in an aberration-free system had lower values (5.3×10^{-13} over $r = 2.5 - 18 \lambda/D$, 5.4×10^{-12} over $r = 2.5 - 3.5 \lambda/D$).

One factor that limits the performance of any of the coronagraphs is the level of optical system aberrations. The optical surface errors we assumed in the simulations were representative of very high quality, custom-made components (though even better optics have been fabricated). But how would the contrast degrade if we had chosen lesser-quality (cheaper, though definitely not cheap) optics? To investigate this, we ran EFC a couple of times on the same VVC system but with the phase errors on the non-primary-mirror optics increased by $3\times$, without and with the primary mirror phase errors increased by $2\times$ (the default phase error levels are given in Table 1). As before, we used 0.5×10^{-5} regularization, 25 iterations, uniform weighting from $0.9 - 18 \lambda/D$, and a new initial DM setting was derived for each case. The resulting contrast maps (Figure 21) show that even in the latter case we readily achieved the milestone requirement. In the 1st case, the mean contrasts were 3.1×10^{-11} ($2.5 - 18 \lambda/D$) and 6.4×10^{-11} ($2.5 - 3.5 \lambda/D$). In the 2nd case, the contrasts were 8.3×10^{-11} ($2.5 - 18 \lambda/D$) and 1.6×10^{-10} ($2.5 - 3.5 \lambda/D$); most of this degradation was likely due to an increase in folding-frequency speckle intensity caused by errors on the primary with spatial frequencies above the DM control limit.

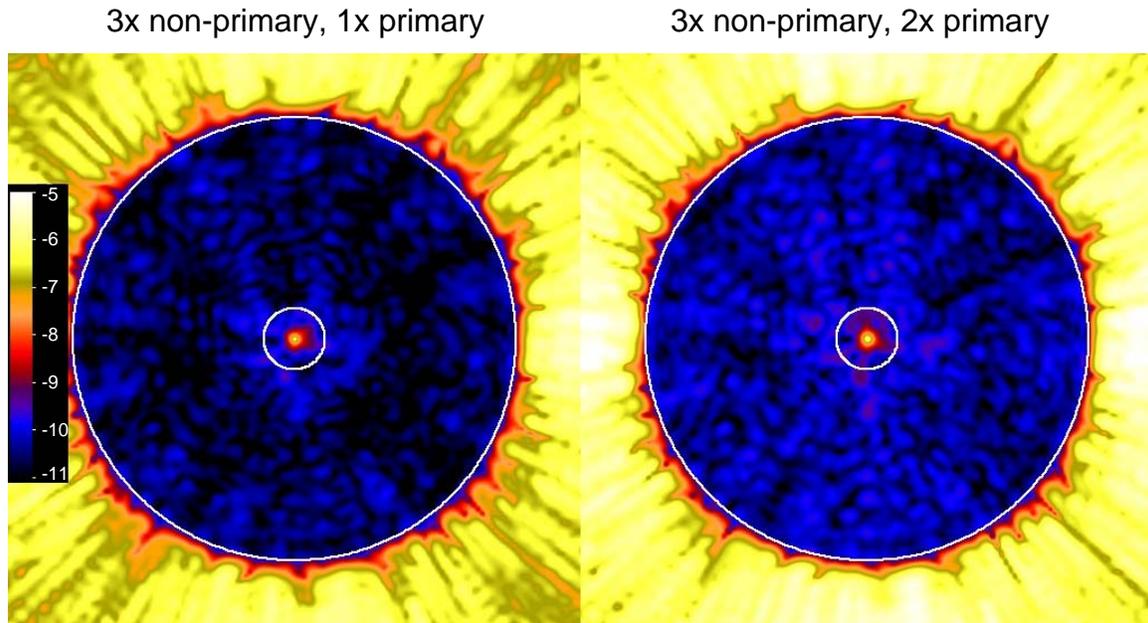


Figure 21 - Broadband contrast after 25 iterations of EFC for a VVC system with (LEFT) the default non-primary mirror phase errors increased by $3\times$ and (RIGHT) the primary mirror phase errors also increased by $2\times$. Log_{10} contrast is shown on the color scale.

Folding frequency aberrations have a λ^{-2} dependence that cannot be fully corrected over a broad bandpass by a dual sequential DM system, which is able to deal with those errors that vary proportionally to λ and λ^{-1} . A dual DM system in a Michelson interferometer

configuration is, however, capable of correcting the λ^{-2} dependent folding frequency errors, but not the others. Pueyo & Kasdin (2007) proposed a 3 DM system composed of one DM followed by a 2-DM Michelson system to control all of these errors. The net result of our study for the three coronagraphs, however, is that with current optical fabrication technology, the contrast floor caused by folding frequency effects are not a limiting factor for 10^{-10} contrast imaging, but rather mask and DM fabrication errors.

We note that the VVC EFC runs converged more rapidly than those for the HBLC. This may be due to the inclusion of wavelength-dependent coating effects in the simulations. EFC may have to work harder to determine a good solution. With a more realistic representation of coating effects, it might take more iterations for VVC to converge as well.

Future work

Our results here show that there are no show-stoppers in the fundamental VVC design that would prevent reaching 10^{-10} contrast. However, the system used here was essentially perfect, with no errors in the vortex mask and no wavelength-dependent phase and amplitude dispersions introduced by the coatings. More work is needed to implement such errors, though they would likely be heavily dependent on particular fabrication methods. So far, VVC has not been developed as much as HBLC or PIAA, and it requires more time to mature.

During our experiments, we concluded that for the particular VVC propagation technique we used (specifically, propagating the effect of the occulting spot) to be accurate, it was necessary to not have any optics between the focal plane mask and the subsequent collimating mirror. This prevents light near the center of the beam from being scattered by surfaces before the vortex has had a chance to form and clear out the center.

Appendix E: Results for the Phase-Induced Amplitude Apodization (PIAA) coronagraph

Coronagraph description

The PIAA optics were described in detail in the Milestone #1 document. We use a system that incorporates a set of forward PIAA optics to apodize the beam by geometrically distorting it. A binary post-apodizer mask consisting of a series of concentric rings is located at a pupil after the forward PIAA optics; it offloads some of the apodization duty from the PIAA optics to reduce fabrication requirements and edge diffraction. A small occulting spot is located at the focus of the PIAA optics. After the occulter the beam passes through a reverse set of PIAA optics that restore the beam to its original mapping.

Generating the DM response matrix

Because of the need to use separate algorithms for propagating between the PIAA optics and the time involved in doing those calculations, the DM response matrix generation process was different for PIAA than for the other coronagraphs. As with the others, each DM actuator was pistoned and the corresponding change in the field at the final image plane was recorded. However, the order of the propagations was optimized especially for PIAA.

We started the sequence by propagating at each wavelength an unaberrated wavefront from the primary mirror to the DM (#1 or #2, depending on which DM the actuators were being poked) using PROPER and saving the field to disk (this forward part of the system did not change if the DMs were altered). Next, a DM actuator was poked and the perturbed field was propagated to PIAA M1 using PROPER. A Fourier transform was then used to compute the X,Y spatial frequency spectrum of this field, which was stored to disk with the field spectra for all the other pokes. The spatial frequencies in the spectra were limited to 60 cycles/diameter in the forward PIAA and 22.4 in the reverse (the field stop at the occulter actually limits the spatial frequencies to 20 cycles/D).

After all of the field spectra at M1 for all of the actuators were computed, the PASP forward-PIAA propagation matrix (see Milestone #1) for each wavelength was generated for each X,Y Fourier component (ripple). This involved analytically propagating and remapping the rippled wavefront corresponding to each component from M1 to M2 and storing the field in the matrix. We again took advantage of perfect symmetry in our model system to reduce the computational load by a factor of 8 by propagating only the ripples corresponding to a 45° sector of the frequency spectrum. We could then use those fields, with appropriate rotates and/or transposes, for the fields corresponding to the other sectors of the spectrum. There were over 70,000 values in the spatial frequency spectrum (which is limited to <60 cycles/D with a sampling of 0.4 cycles/D/element), and each corresponding wavefront (ripple pattern) was 460 × 460 pixels using double precision complex numbers (16 bytes/pixel). Even using the factor of 8 reduction, the propagation matrix was 30 gigabytes.

The M1 field spectrum corresponding to each actuator poke was read in and multiplied by the propagation matrix, producing the wavefront at M2, to which the S-Huygens-generated edge diffraction component of M1 was added. The result was then stored to disk. Using a propagation matrix and processing all of the DM poke wavefronts in one stage avoided the need to recompute the Fourier term propagations for each poke. The M2 wavefronts were then serially read in and propagated using PROPER from M2 to the filtered version of the post-apodizer, and then on to the occulter and the entrance of the reverse PIAA. There, the spatial frequency spectrum was computed and saved to disk. Another propagation matrix was computed for the reverse PIAA, all the DM poke wavefronts propagated through it using PASP, and then PROPER was used to the final focus. The total computation time for generating the response matrix was 33.8 hours on a dual quad-core Xeon workstation with 48 GB of memory, which meets the milestone requirement of <48 hours. It required the combined use of PROPER and custom, multithreaded C code.

PIAA optical aberrations

Unlike the VVC and HBLC, we had surface error maps of actual PIAA optics that we could incorporate into our simulations. The interferometer-measured maps (Figure 3) were of the PIAA 2 optics fabricated by Tinsley, which produce the same apodization profile as our model PIAA does. These are currently the best set of PIAA optics available. Use of these maps allowed us to incorporate PIAA-unique aberrations that would not have been reproduced by synthetic maps that have isotropic surfaces that lack spatially-correlated structure. The most obvious example of this is the region of narrow rings along the edge of M1, which have surface departures of more than ± 30 nm. The curvature is very high here, making it difficult to polish out errors even with ion beam techniques. Prior to use, these maps were filtered to <48 cycles/D and then interpolated using a damped sinc method to match the sampling of the wavefront in the simulations.

Generating the dark hole

Because we propagated an entire wavefront at each wavelength in sequence when running EFC, a more direct progression through the system was used than when the response matrix was computed. As before, the wavefront was propagated from the primary mirror to the DM and then on to PIAA M1 using PROPER. At M1, the Fourier transform of the wavefront was taken and each corresponding Fourier component (ripple) was propagated through the system. This is different from before, where a propagation matrix was constructed for all of the components. We found that with a single wavefront at sequential wavelengths it was quicker to compute the propagation of each component separately than to read in the entire propagation matrix from disk (which is faster when propagating multiple wavefronts at a single wavelength). The same happens at the reverse PIAA. The total time to propagate once through the system was 5 minutes on the previously detailed workstation (1.7 minutes just from M1 to M2).

Contrast with a perfect wavefront

The use of a binary post-apodizer creates a contrast floor, even in the absence of any aberrations. Figure 22 (left panel) shows the contrast map of the unaberrated system. The mean contrast in the dark hole is 2.0×10^{-11} and the median is 1.0×10^{-11} . At the inner working angle the contrast is 6×10^{-10} , and it falls rapidly to $< 10^{-11}$ towards the outer working angle. A couple of iterations of EFC is sufficient to reduce the mean contrast to 3.5×10^{-12} and the contrast at the IWA to 2×10^{-10} . Thus, our implementation of PIAA was not a limiting factor for contrast, except that it would be desirable for the IWA contrast to be $< 10^{-10}$. The bulk of the residuals are at the longer wavelengths, indicating that making the occulter larger might reduce some of the light.

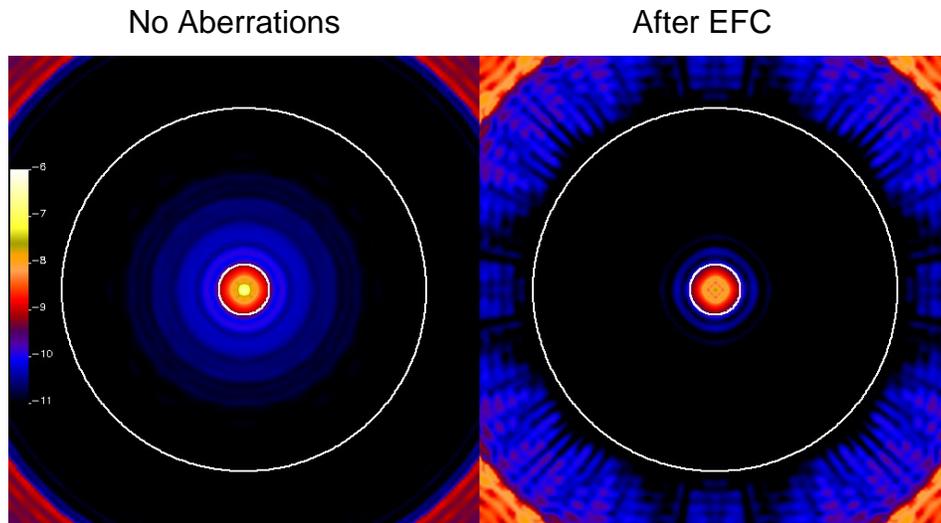


Figure 22 - Contrast ($\lambda=500 - 600$ nm) for the unaberrated PIAA coronagraphic system with binary post-apodizer, before (left) and after (right) two iterations of wavefront control with EFC (2 DMs). The circles are 2.5 & $18 \lambda/D$ ($\lambda = 550$ nm). $\text{Log}_{10}(\text{contrast})$ is shown in the color bar (contrast range = 10^{-11} to 10^{-6}).

Results with only M1 aberrations

The M1 mirror provides the bulk of the apodization in a PIAA system. The most critical region of the optic is the edge, where the curvature rapidly varies. The measured M1 map shows a series of concentric rings along the edge, remnants of the figuring and polishing process. In the unfiltered map these have spatial frequencies centered around ~ 66 cycles/diameter (but including lower frequency terms), which would require a DM with at least $66 \times 2 = 133$ actuators across the pupil to correct. In the map filtered to < 48 cycles/D used in our simulations, the remaining spatial frequencies appear centered at ~ 37 cycles/D, requiring a DM with 74 actuators/D, still far above our DM's density of 46 actuators/D (the filtering also reduces their amplitude by about 40%). They therefore could only be weakly corrected by the DM in the simulations. The remaining errors scatter light over the entire remapped pupil and over the entire dark hole. The filtered

map has an RMS wavefront error ($2\times$ surface) of 6.2 nm with peak amplitudes of ± 40 nm.

Figure 23 shows the dark hole region after running EFC in the two DM system where only M1 was aberrated. The mean contrast in the dark hole was 1.7×10^{-9} and the median was 7.1×10^{-10} . To verify that the edge errors were the dominant cause of the high level of light in the dark hole, we zeroed the errors within the interior of the map, leaving errors only in the outer 10% of the radius. Running EFC again produced a dark hole almost identical to that produced by the complete map. A pronounced, nearly-horizontal streak running across the dark hole is oriented perpendicular to the highest ridges along the edge of the map with a contrast of $>10^{-8}$.

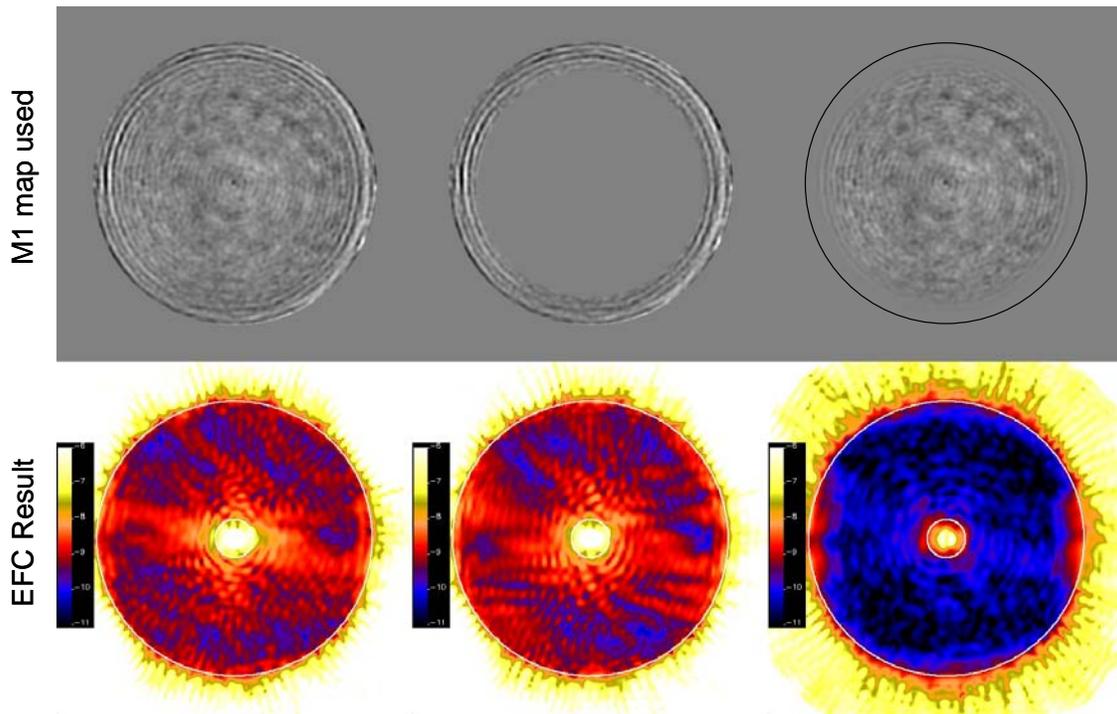


Figure 23 - Results of running EFC in a system in which only the PIAA M1 mirror is aberrated, using the measured error map. The map used is shown on top with the corresponding EFC-generated dark hole on the bottom. (Left) The full M1 error map is used. (Middle) Errors in the interior of the M1 optic were set to zero. The resulting dark hole is very similar to that for the full error map, indicating that the errors at the edge of the optic are the primary limitation. (Right) Errors at the edge of the optic are set to zero, resulting in a good dark hole. The contrast scale ranges from 10^{-11} to 10^{-6} .

To determine how much lower the errors needed to be in order to reach the required contrast, we adjusted the M1 error map (Figure 24). The edge errors were decreased by multiplying the map with a cosine taper on the outer 30% radius (3.6 nm RMS wavefront error, ± 18 nm peak errors, the highest outer edge errors reduced by $\sim 3\times$) and running EFC again. This significantly improved the contrast in the dark hole, but there were still high values in the streak and around the inner radius (mean = 2.8×10^{-10} , median = 9.5×10^{-11}). After that, a cosine-squared taper was used (same full-map statistics as with the

cosine taper, but the outer errors reduced by $\sim 10\times$ relative to the non-tapered map), reducing the mean after EFC to 1.5×10^{-10} and the median to 4.1×10^{-11} . There were still small regions of about 10^{-9} contrast near the inner radius and along the outer radius in the direction of the streak. Given that these results remained too large, the entire map (after the cosine-squared tapered) was reduced in amplitude by a factor of two. The end result had a mean contrast of 3.3×10^{-11} and a median of 8.0×10^{-12} . This indicates that the errors in the center of the M1 optic need to be reduce by $2\times$ and those along the outer edge by $\sim 20\times$. While this may be feasible for the interior, the large reduction in the outer errors may be beyond current polishing techniques.

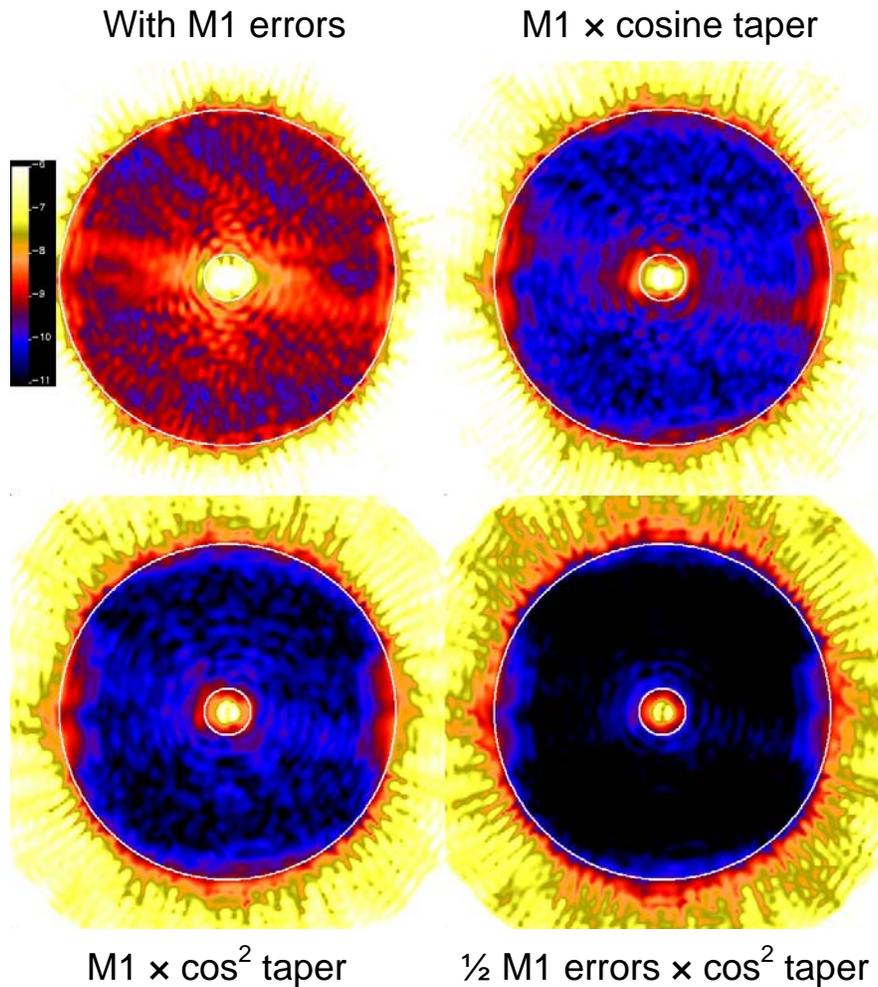


Figure 24 - EFC-generated (2 DMs) dark holes ($\lambda = 500\text{-}600\text{ nm}$) for a PIAA coronagraphic system with only M1 aberrated (using the measured map filtered to <48 cycles/diameter). (Top left) Result using filtered map. (Top right) Using filtered map with outer 30% radius multiplied by a cosine taper. (Bottom left) Using filtered map with outer 30% radius multiplied by a cosine-squared taper. (Bottom right) Same map as preceding result, but with all map errors reduced by one-half. $\text{Log}_{10}(\text{contrast})$ is shown in the color bar (contrast range = 10^{-11} to 10^{-6}).

Results with only M2 aberrations

As we did for M1, we ran EFC on a system in which only M2 was aberrated, using the measured error map. The <48 cycles/D filtered map had an RMS wavefront error of 6.8 nm with maximum excursions of ± 30 nm. With more gradual curvature than M1, the M2 errors appear more like those on a conventional high-quality optic. The M2 error levels, though, were still much higher than can routinely be achieved (<1 nm RMS) on custom-figured mirrors, indicating that some additional time (and, of course, money) would produce a smoother optic.

The EFC result (Figure 25) has a mean contrast of 1.7×10^{-10} and median of 4.2×10^{-11} . There was a prominent $\sim 10^{-9}$ contrast vertical streak crossing a portion of the dark hole. The DM patterns showed large strokes on a small number of actuators at the edge of DM #1, an attempt by the algorithm to introduce a compensating wavefront to null the streak.

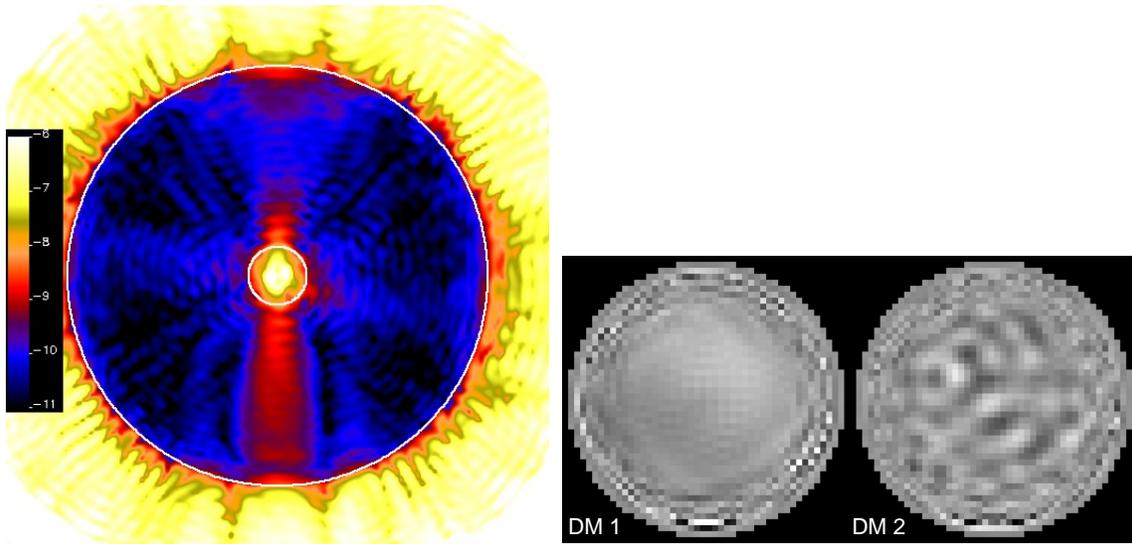


Figure 25 - (Left) EFC-generated (2 DMs) dark hole for a system with only M2 aberrated (using measured error map). The circles are at 2.5 and $18 \lambda/D$ ($\lambda = 550$ nm). $\text{Log}_{10}(\text{contrast})$ is shown in the color bar (contrast range = 10^{-11} to 10^{-6}). (Right) The DM actuator piston maps derived by EFC to create the dark hole. DM #1 is conjugate to M1. The piston amplitudes shown are ± 11 nm

Correcting post-PIAA optical errors with a DMs prior to PIAA

The preceding results highlight the limitations of using DMs upstream of the PIAA optics (i.e. in undistorted wavefront space) to correct wavefront errors from optics located after PIAA, specifically from M2 to the occulter (i.e. optics in distorted wavefront space). There are some subtle problems that complicate wavefront control in this configuration.

As illustrated in Figure 26, in the forward direction (from M1 to M2), the beam is compressed towards the center and stretched along the edge, but when backing through

the system (from M2 to M1 to the DMs), the opposite happens. A low spatial frequency error on M2 is seen by the DM as a higher frequency one along the outer edge. This can push the aberration beyond the spatial frequency control regime of the DMs so that only the central region of the wavefront can be well corrected (lower left image of Figure 26). The remaining errors will scatter light over the entire pupil and inside the dark hole region. A second problem is that when the pre-PIAA DMs are used to correct post-PIAA errors, they effectively introduce what amounts to a phase error at M1, which in turn causes a wavelength-dependent error in the PIAA-induced apodization that cannot be fully corrected.

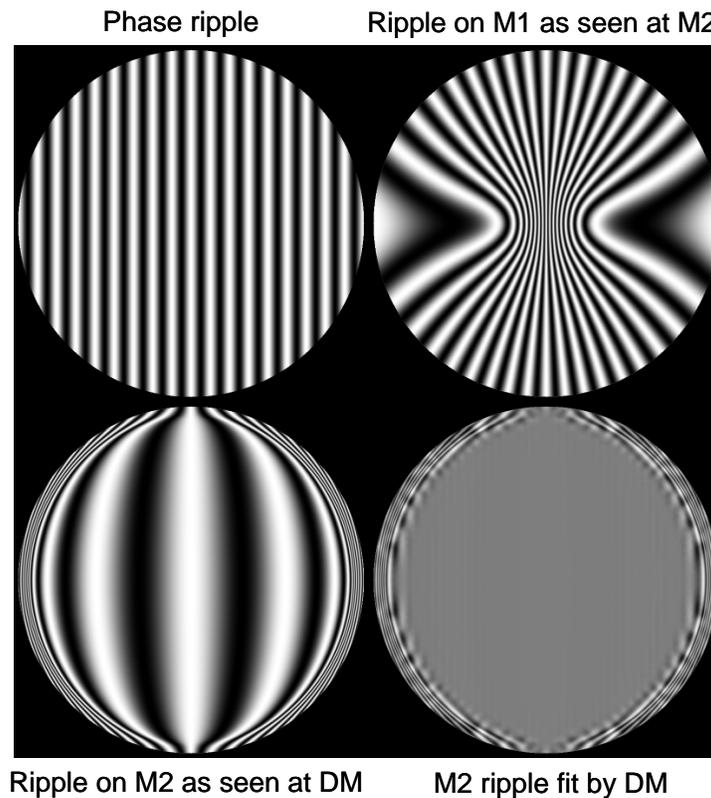


Figure 26 - A phase ripple on one side of PIAA as seen from the other side. (Top left) Example ripple of 18 cycles/diameter. (Top right) Ripple on M1, as seen from M2. (Bottom left) Ripple on M2 as seen from the deformable mirror located prior to PIAA and conjugate to M1. (Bottom right) The DM-corrected M2 ripple (46 actuators/D).

Adding a 3rd DM after PIAA

The previous examples demonstrated that there are inherent difficulties in controlling wavefront errors when the deformable mirrors are in the undistorted wavefront space but critical optics are in the distorted (PIAA apodized) space. A potential solution is to add a deformable mirror after the PIAA optics to directly correct aberrations from surfaces

between M2 and the occulter. This would also offload some of the control responsibility from the pre-PIAA DMs.

We added another DM at the plane of the post-apodizer (Figure 27), with 46 actuators across the opening of the mask (out to the last clear ring). The addition of these new actuators to the DM response function matrix required another stage of propagation. At each wavelength, an unaberrated wavefront was propagated from M1 to M2 with PASP, then to the 3rd DM using PROPER, and the result stored. For each actuator, this wavefront was read in, the actuator poked, and the resulting wavefront propagated to and through the reverse PIAA and on to the final image plane, using a similar procedure as for the prior DMs. The new response functions were added to the DM response matrix, which was then inverted. The 3rd DM increases the size of the response matrix to about 3.2 GB.

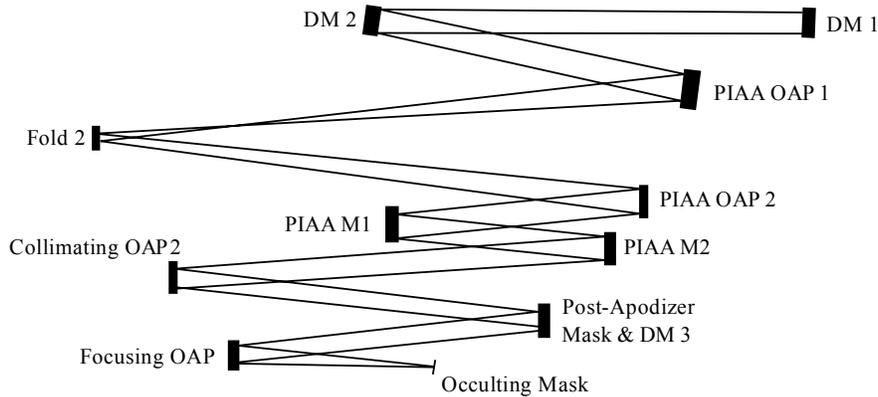


Figure 27. Schematic of a portion of the revised PIAA coronagraph layout, shown starting at DM 1 and through the PIAA focal plane occulter. The 3rd DM is added at the location of the post-apodizer mask.

We first ran EFC with three DMs in a system with only M1 aberrations and with and without the edge tapered downweighting of the M1 map errors. The resulting dark holes (Figure 28) showed that the 3rd DM helped to improve the overall contrast level compared to the 2 DM solution, but the streak from the large edge errors remained. With a tapered downweighting of the M1 edge aberrations, the three-DM results were only marginally better than those using two DMs. The 3rd DM had little control authority over errors from surfaces prior to PIAA. The beam falling on the 3rd DM is highly apodized, so only a small number of actuators actually have a strong effect.

When the 3-DM EFC is run on a system with only M2 aberrations, the improvement is significant (Figure 29). The mean contrast decreased to 1.1×10^{-11} , down from 1.7×10^{-10} with the 2-DM solution. The 10^{-9} streak was largely suppressed, and the only significant errors were near the inner radius, at the level of those seen in the unaberrated system limited by the post-apodizer.

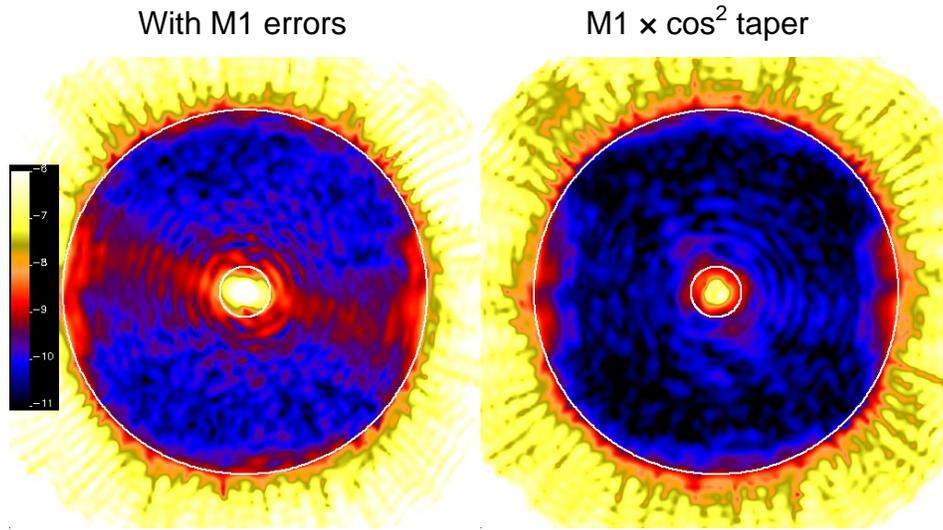


Figure 28 - EFC-generated dark holes ($\lambda = 500-600$ nm) for a PIAA system using 3 DMs and with only M1 aberrated (using the measured map filtered to <48 cycles/diameter). (Left) Result using filtered map. (Right) Using filtered map with outer 30% radius multiplied by a cosine-squared taper. These can be compared to the 2-DM solutions shown in Figure 24. $\text{Log}_{10}(\text{contrast})$ is shown in the color bar (contrast range = 10^{-11} to 10^{-6}).

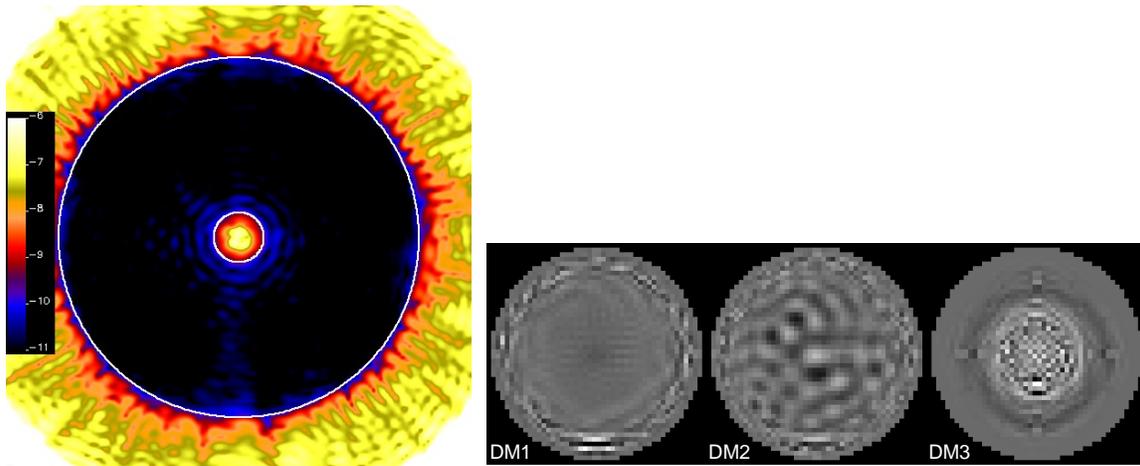


Figure 29 - (Left) EFC-generated (3 DMs) dark hole ($\lambda=500 - 600$ nm) for a system with only M2 aberrated (using measured error map). This can be compared to the 2-DM result in Figure 25. The circles are at 2.5 and $18 \lambda/D$ ($\lambda = 550$ nm). $\text{Log}_{10}(\text{contrast})$ is shown in the color bar (contrast range = 10^{-11} to 10^{-6}). (Right) The DM actuator piston maps derived by EFC to create the dark hole. DM #1 is conjugate to M1, DM #3 is

Final results

The final trials used 3 DMs with the $\frac{1}{2}$ amplitude, cosine-squared tapered, filtered M1 map and the filtered M2 map, plus the default errors on all the other optics. We experimented with different regularizations (Figure 30) and found that the 0.5×10^{-5} value that was optimal for the HBLC and VVC did not provide the best results for PIAA, causing early divergence, especially near the inner working angle. The addition of the 3rd DM may have made the inverted response matrix a bit less well conditioned. A value of 0.9×10^{-5} provided more damping and better results. A mean contrast of 8.6×10^{-11} over $2.5 - 18 \lambda_c/D$ was achieved after 41 iterations (Figure 31 and Figure 32), with a mean of 2.0×10^{-10} over $2.5 - 3.5 \lambda_c/D$ (the contrast in this inner region was down to 1.45×10^{-10} at iteration 14 before diverging). The Milestone 2 requirements were met. Whether this is realistic depends on how much improvement can be made in reducing the M1 errors. An even greater amount of damping was tried using a regularization value of 1.5×10^{-5} , which provided a slightly better contrast at the IWA ($\sim 1.28 \times 10^{-10}$ at iteration 45, after which it diverges) but did not converge after 50 iterations over the full field (9.5×10^{-11} contrast).

The contrast is optimal at the central wavelength, as demonstrated when the final solution is decomposed into separate wavelengths (Figure 33). In this case, the chromatic behavior is dependent on the post-apodizer (which was optimized for a broad bandpass) and the M1 surface errors, which introduce wavelength-dependent apodization of the beam (but have been suppressed to achieve the required contrast). Unlike HBLC, there are no wavelength-dependent coating effects to consider (assuming that the post-apodizer can be deposited on an anti-reflection coated substrate with benign wavelength dependent properties).

Performance limitations

There are a few properties particular to PIAA that significantly impact its contrast performance. First, the large curvature at the outer edge of the M1 optic, necessary to apodize the beam, requires fine-scale polishing techniques. The surface errors in the current actual M1 optic (PIAA2 configuration) need to be reduced by factors of tens to produce a properly apodized beam. Alternative designs have been suggested that divide the apodization function among 3 or 4 optics, allowing easier fabrication. However, the additional number of optics and the corresponding sensitivities to alignments between them increase the complexity of the system, and this configuration is at yet unproven.

The next factor is the post-apodizer. Future designs of binary post-apodizers need to include the effects of diffraction, rather than assuming that the predicted geometric beam (no diffraction) can be used in the optimization process. Our results at the inner working angle are limited by our post-apodizer design, which was optimized using a geometric beam.

The remapping of the wavefront by PIAA means that aberrations on one side of the PIAA M1 optic have different spatial frequency distributions when projected to the other side. The main implication of this is that any DMs located prior to PIAA cannot fully correct lower-frequency aberrations on optics after M1 that have been, from the perspective of

these DMs, remapped to high spatial frequencies. This requires the addition of another DM after M1 and prior to the focal plane occulting mask. This DM can directly correct phase errors on those optics situated after M1 (such as M2 and fold/focusing optics).

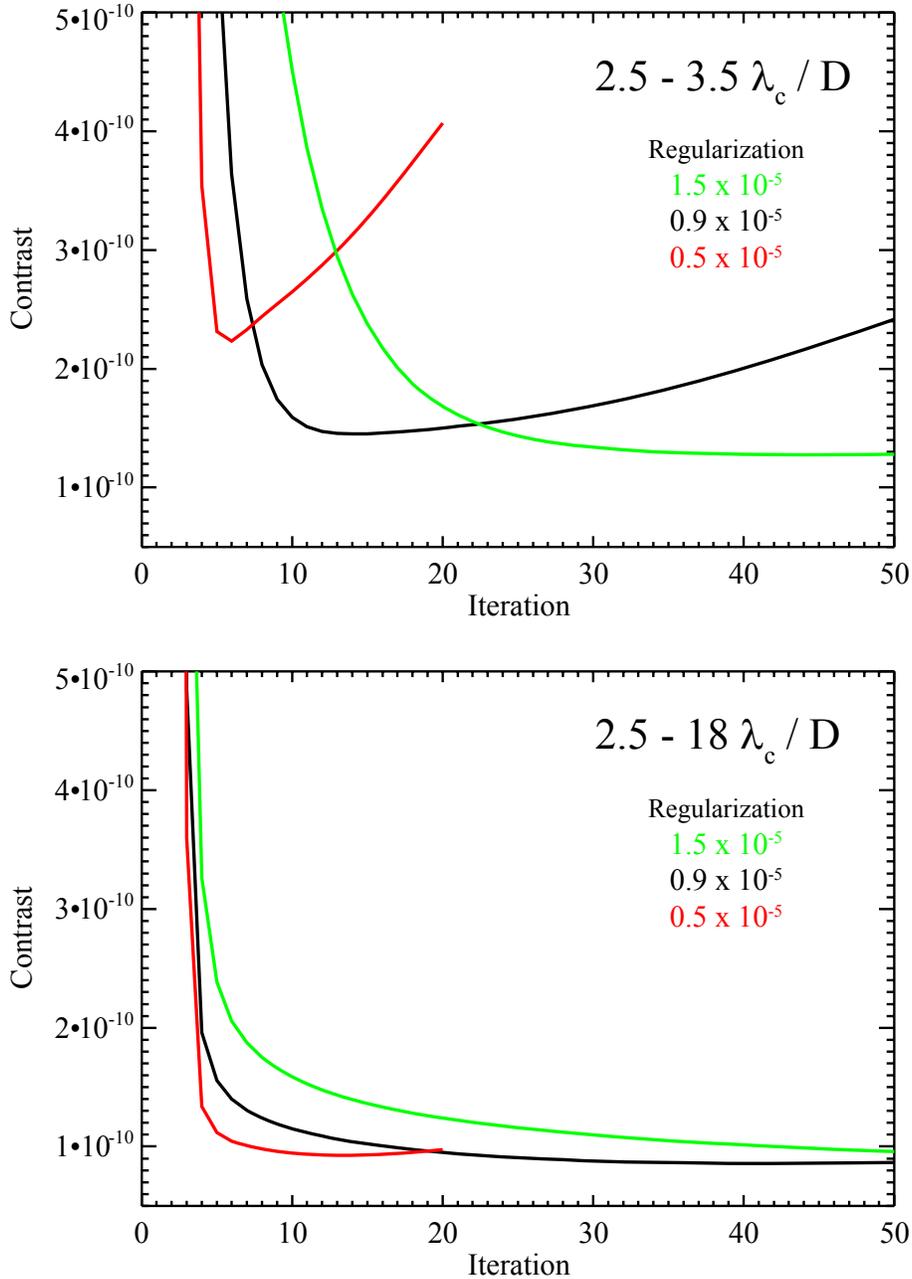


Figure 30 - Plots of mean broadband contrast with the PIAA versus EFC iteration (3 DMs, modified PIAA error maps), measured within an annulus of (top) $r = 2.5 - 3.5 \lambda_c / D$ and (bottom) $2.5 - 18.0 \lambda_c / D$ radians ($\lambda_c = 550$ nm). Plots are shown for different values of regularization used in the corresponding DM response matrices. Iteration 0 is the result using the initial “flattening” settings on the 1st DM, before running EFC.

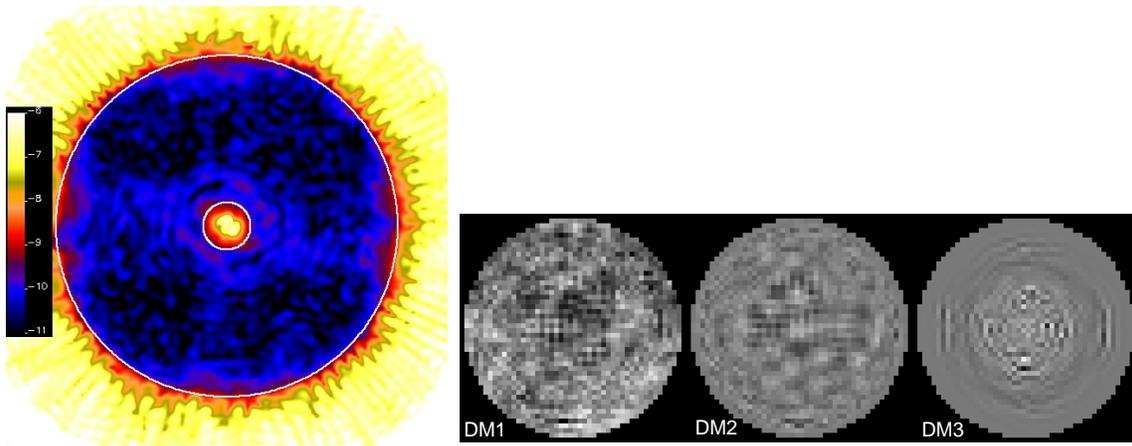


Figure 31 - (Left) EFC-generated dark hole ($\lambda = 500\text{-}600$ nm) for a PIAA system using 3 DMs and all surfaces aberrated. $\text{Log}_{10}(\text{contrast})$ is shown in the color bar (contrast range = 10^{-11} to 10^{-6}). (Right) The DM actuator piston maps derived by EFC to create the dark hole. DM #1 is conjugate to M1, DM #3 is located at the PIAA post-apodizer. The piston amplitudes shown are ± 15 nm.

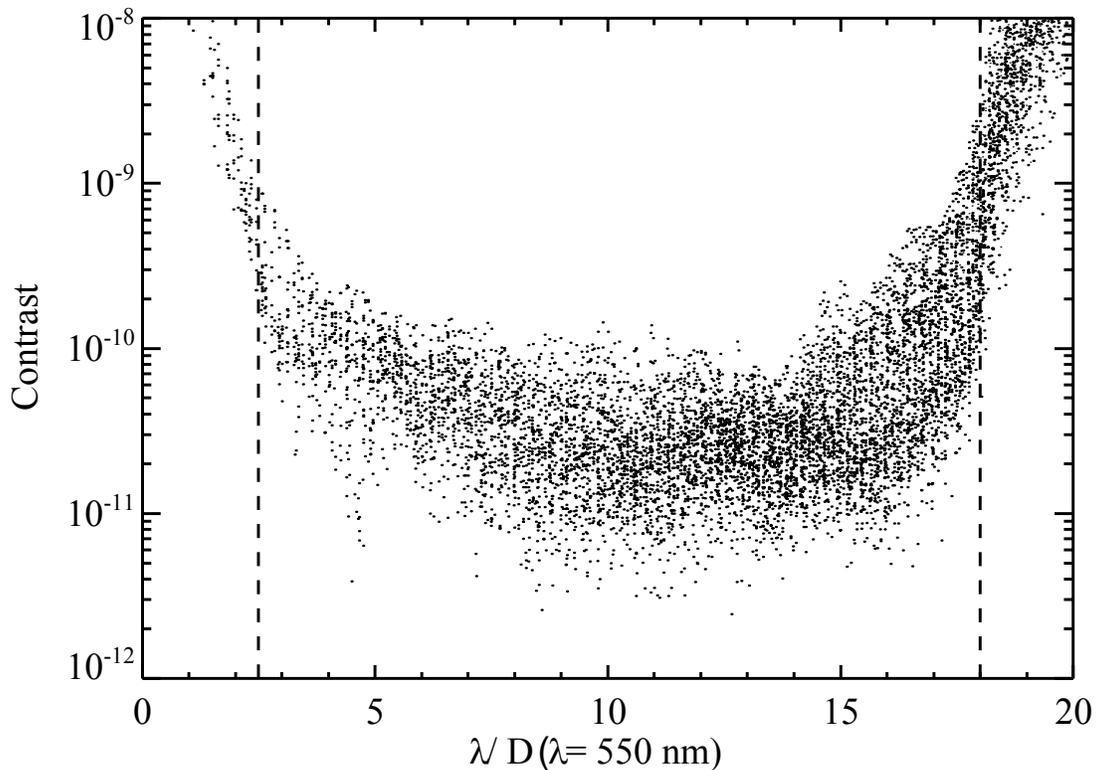


Figure 32 - Contrast versus radius for each pixel in the dark hole region of the final PIAA EFC solution. The inner and outer working angles of 2.5 and $18 \lambda/D$ are indicated.

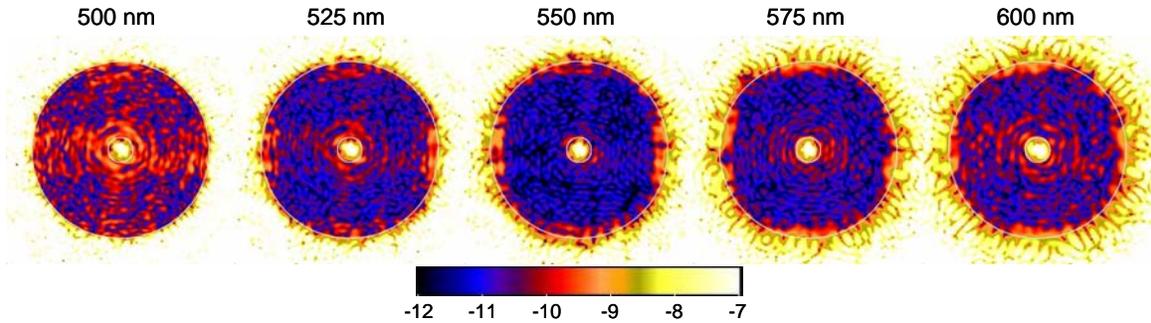


Figure 33 - Final PIAA broadband contrast decomposed into monochromatic contrasts over the field. $\text{Log}_{10}(\text{contrast})$ is shown on the color scale (note that the contrast range is different here than for the other contrast maps in order to emphasize the variations with wavelength). The inscribed circles mark the $r = 2.5$ and $18 \lambda/D$ boundaries of the contrast measurement region. At the central wavelength (550 nm), the mean contrast is 6.2×10^{-11} over the dark hole.

Building the DM response matrix using remapping rather than PASP

The PASP algorithm, while accurate and considerably faster than S-Huygens, is still computationally expensive. Generating the DM response matrix for a system like that studied here takes tens of hours unless significant computer resources are applied to parallelize the work. A “quick and dirty” alternative may use wavefront remapping via interpolation to model the effect of the PIAA optics rather than applying the explicit wavefront propagation provided by PASP.

The remapping technique was discussed in the Milestone 1 report. Briefly, it involves using interpolation to remap the wavefront at M1 onto M2 and applying the expected PIAA amplitude apodization function. This ignores conversions between phase and amplitude wavefront errors that actually occur between the M1 and M2 propagation, but it is much faster to execute than PASP. The accuracy of the remapping technique, as reported in the Milestone 1 report, was 7.3% for a 10^{-5} mean contrast field and 14.6% for a 10^{-10} contrast one, which would not meet the Milestone 1 requirement.

To assess if its reduced accuracy would adversely impact the ability to generate a suitably high contrast dark hole, the remapping method was used to generate the DM response matrix for the TDEM optical system. This matrix was then used with EFC to create the dark hole in the model system (including the reduced PIAA surface errors) in which the PIAA propagations were done using PASP. In essence, this tested the remapping DM matrix against a “real” system, with the assumption that PASP is sufficiently accurate to reflect reality.

The result of running EFC (Figure 34) shows that in this system the remapping was adequate to achieve a dark hole with 10^{-10} mean contrast. This was achieved with about twice the number of iterations than was required the PASP-generated matrix. If we assume that each remapping requires one second of CPU time to execute, then the total

time to create the response matrix by this method is reduced by nearly a factor of two compared to using PASP (from 34 hours down to 17 for 2080 actuators).

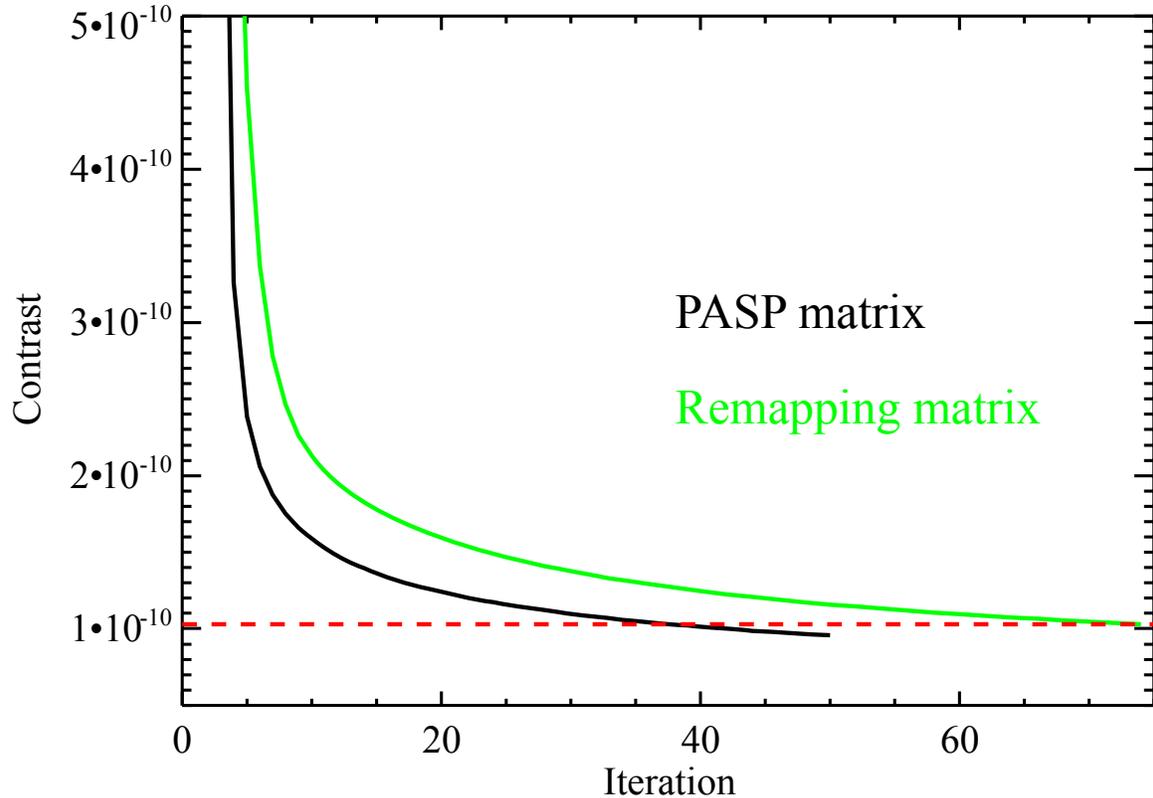


Figure 34. Mean contrast over the $2.5 - 18 \lambda/D$ field versus EFC iteration using a DM response matrix generated using PASP (black) and remapping via interpolation (green). The dashed red line demonstrates that it requires about twice as many iterations to achieve the same contrast using the remapping matrix. A regularization of 1.5×10^{-5} was used for both.

While remapping may be a valid method for generating the response matrix in this case, it is still important that future studies examine its suitability to a specific system, which is best achieved by incorporating PASP in the reference model. Any DM response matrix generated via modeling will have some mismatches compared to a real system - the differences between matrices generated via remapping and PASP may be smaller than between any model and reality.

Future work

PIAA testbed experiments so far have used, for simplicity, a DM located after M2 and prior to the occulter. This significantly limits the size of the dark hole field, so proposed space missions have favored configurations with DMs prior to PIAA. As we have shown, however, the wavefront control behavior of PIAA is such that the current testbed

experiments do not replicate the performance of the proposed space mission layouts. Future experiments need DMs both prior to and after the PIAA M1 optic to properly explore the proposed mission layouts.

The computational resources needed to propagate the wavefront between the M1 and M2 optics are fairly great. The PASP algorithm may be suitable for implementation on graphics processing units (GPUs), whose massively parallel architectures might provide substantial increases ($>10\times$) in speed when computing the DM response matrix. This is especially important because in a real system the DMs would not be geometrically perfect, as they are in the simulations, so one could not use the “trick” of using DM symmetry to reduce the response matrix generation time.

Appendix F. Acronyms

AMD	Advanced Micro Devices
D	Diameter of the entrance pupil (typically)
DM	Deformable mirror
EFC	Electric field conjugation
FITS	Flexible Image Transport Specification
GB	Gigabyte
HBLC	Hybrid band-limited coronagraph
HCIT	High Contrast Imaging Testbed
IDL	Interactive Data Language
IWA	Inner working angle
JPL	Jet Propulsion Laboratory
OWA	Outer working angle
PASP	PIAA angular spectrum propagation
PIAA	Phase induced amplitude apodization
PSD	Power spectral density
PSF	Point spread function
RAM	Random access memory
RMS	Root mean square
SVD	Singular value decomposition
TDEM	Technology development for exoplanets missions
TPF-C	Terrestrial Planet Finder - Coronagraph
VVC	Vector vortex coronagraph

----- END OF DOCUMENT -----