

TECHNOLOGY DEVELOPMENT FOR EXOPLANET MISSIONS

Final Technology Milestone Report
NASA Grant No. NNX12AI51G

JPL Document D-93928

INTEGRATED CORONAGRAPH DESIGN AND WAVEFRONT AMPLITUDE CONTROL USING TWO DEFORMABLE MIRRORS

PROF. N. JEREMY KASDIN
PRINCIPAL INVESTIGATOR
PRINCETON UNIVERSITY
MECHANICAL AND AEROSPACE ENGINEERING

August, 2014

Coinvestigators and Collaborators:

Stuart Shaklan, Robert Vanderbei, Tyler D. Groff, Alexis Carlotti, Laurent Pueyo, Michael Carr,
A. J. Eldorado Riggs

Approvals

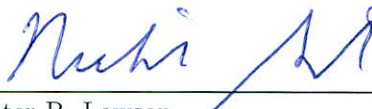
Released by:



11/3/14

N. Jeremy Kasdin
Principal Investigator, Princeton University

Approved by:

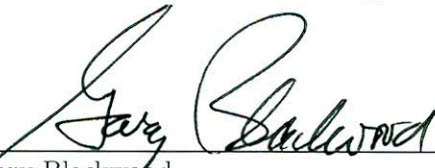
 in behalf of chief technologist 11/12/14

Peter R. Lawson
Exoplanet Exploration Program Chief Technologist, JPL



11/12/14

Nicholas Siegler
Exoplanet Exploration Technology Manager, JPL



11/20/14

Gary Blackwood
Exoplanet Exploration Program Manager, JPL



12/10/14

Douglas Hudgins
Exoplanet Exploration Program Scientist, NASA HQ



12/10/14

John Gagosian
Exoplanet Exploration Program Executive, NASA HQ

CONTENTS

1. Executive Summary	4
2. Certification	4
3. Success Criteria	5
4. Shaped Pupil Design and Manufacture	5
5. Testbed Layout	6
6. Two Deformable Mirrors in the HCIT	8
7. Wavefront Estimation	9
7.1. Algorithm Development	9
7.2. Implementation at the HCIT	9
7.3. Comparison	10
8. Wavefront Control	10
8.1. EFC	11
8.2. Stroke Minimization	11
8.3. Comparison of Lab Results	12
9. Results	14
9.1. Experimental Results	14
9.2. Limitations	14
10. Conclusions	16
11. Acknowledgement	17
Appendix A. Success Criteria from White Paper	18
Appendix B. Table of Acronyms	19
References	20

1. EXECUTIVE SUMMARY

In this final report, we document the results of our ROSES Technology Development for Exoplanet Missions (TDEM) two-year research program to advance the technology associated with using two deformable mirrors to perform amplitude and phase control in a coronagraph. We described the goals of the program in our Technology Milestone Whitepaper, JPL Document D-81164 dated October, 2013. We will refer to that often in this report.

While we did not meet our milestone, largely due to a reprioritization of facilities at JPL after the start of the WFIRST/AFTA coronagraph development, we did verify the effectiveness of using two deformable mirrors to create dark holes on both sides of the image. We achieved a monochromatic contrast of 3.6×10^{-9} in the High Contrast Imaging Testbed (HCIT) using one 32×32 DM (with 23×23 actuators illuminated) and one 64×64 DM (with 48×48 actuators illuminated) and a shaped pupil coronagraph, just shy of our primary milestone requirement of 1×10^{-9} .

As a reference, our milestones from the whitepaper were:

TDEM Primary Milestone:

- Demonstrate with 90% confidence that the system can achieve symmetric dark holes in the image plane in monochromatic light with an expected average contrast plus 3-sigma $\leq 1 \times 10^{-9}$ from 5-9 λ/D using two deformable mirrors in series.

TDEM Secondary Milestone:

- Demonstrate with 90% confidence that the system can achieve symmetric dark holes in the image plane in a 10% band about the central wavelength with an expected average contrast plus 3-sigma $\leq 5 \times 10^{-9}$ from 5-9 λ/D using two deformable mirrors in series.

In the remainder of this report we describe the design of the shaped pupil coronagraph, the algorithm development, and the experiments we performed.

2. CERTIFICATION

In this section, we reference the list of items for the certification data package from Section 6 of the milestone whitepaper and identify where they can be found in this report. The certification item description from the whitepaper is given in italics followed by the reference information in roman type.

- (1) *A narrative report, including a discussion of how each element of the milestone was met, and a narrative summary of the overall milestone achievement.*
This narrative report responds to item (1) of the certification data package.
- (2) *A complete description of the HCIT layout and optical system used with the significant characteristics.*
See Section 5.
- (3) *The sets of images used for calibration of the reference star.*
Calibration was performed using the standard HCIT procedure by the HCIT personnel. The procedure they follow is to: 1) calibrate the contrast of the quilting order to the main PSF, 2) put the focal plane mask in place, and 3) use the quilting order peak brightness from then on to calibrate the contrast in the dark hole (without ever moving the mask again). We unfortunately don't have the images used for this process.
- (4) *Microscope images of the shaped pupil and focal plane masks with simulated PSFs.*
See Section 4.

- (5) *Calibrated final images of the coronagraph contrast field for each experiment run, for both the monochromatic and broadband experiments.*

See Section 9 for the final calibrated images of the single monochromatic experiment. There were no broadband experiments as we ran out of time in the facility.

- (6) *Calibrated curves showing the contrast convergence of the control algorithm for all runs.*

See Section 7.3.

- (7) *A histogram of the brightness distribution of pixels in the dark hole for each of the final images in the data set and for the combined data.*

Since there was only a single experimental run with one final image of the dark hole, we were not able to produce histograms across multiple images. However, Fig. 8 shows a histogram of the brightness distribution of pixels (in units of contrast) across the dark hole for that final image.

3. SUCCESS CRITERIA

Appendix A reproduces verbatim the success criteria from the milestone whitepaper. Here we reference those criteria and indicate where in this report evidence of success or a description of the process can be found.

Primary Milestone:

- (1) The HCIT was successfully brought to vacuum with all alignments completed.
- (2) The calibration was successfully completed.
- (3) The new Kalman filter was implemented and successfully obtained estimates faster than the previous method. This was used with stroke minimization to successfully create two symmetric dark holes.
- (4) The mean contrast over the dark hole was computed to be 3×10^{-9} .
- (5) Unfortunately, there was not time to repeat the experiment.
- (6) Because there was no time to repeat the experiment, no statistical analysis was completed.

Secondary Milestone:

The Secondary Milestone was not attempted due to lack of access to the HCIT facility.

4. SHAPED PUPIL DESIGN AND MANUFACTURE

One shaped pupil mask was used for this entire experimental run. The shaped pupil is a ripple mask similar to that used in Princeton's 2007 run at the High Contrast Imaging Testbed (HCIT) at Jet Propulsion Laboratory (JPL).⁽¹⁾ The mask was designed in the ideal case to achieve 1.0×10^{-10} contrast or better at all points in the discovery space, the region of high contrast in a perfect optical system, which in this case extends from 4.5 to $56 \lambda/D$ in symmetric 90° degree wedges. This difficult contrast constraint limited the IWA to $4.5 \lambda/D$. The outer working angle was chosen to be $56 \lambda/D$ to prevent the PSF copies from DM1 quilting orders from leaking into the discovery space. The discovery space stretched across two 90° arcs on opposite sides of the image plane. The high design contrast also limited the transmission of the mask to 15% for an unobstructed circular aperture. The corresponding focal plane mask was slightly oversized to allow for slight misalignments in the testbed. It blocked up to an IWA of $4.75 \lambda/D$ (at 808nm) and left openings 80° -across for dark hole correction in the 90° discovery region. Because the shaped pupil is a binary mask, it is fundamentally achromatic and its PSF scales with wavelength. We note that the *correction region* of the shape pupil is slightly smaller. This is the sub-region of the discovery space in which the DMs recover high contrast from the aberrated PSF. The correction region ranged from a radius of $4.8 \lambda/D$ to an x-distance of $9.5 \lambda/D$ over symmetric 80° -degree wedges. To avoid accidentally




FIGURE 1. (Left) Two-dimensional image of the 10^{-10} , $56 \lambda/D$ shaped pupil design. White areas are transmissive. The red circle illustrates the pupil diameter. (Right) Schematic of the first wafer layout, including several shaped pupils and the corresponding image plane masks.

measuring pixels obscured by the focal plane mask, in the results quoted below pixels were counted only from an IWA of $5.0 \lambda/D$ over 75-degree openings.

JPL's Microdevices Laboratory (MDL) fabricated this shaped pupil and others, along with the corresponding focal plane mask (FPM), using Deep Reactive Ion Etching (DRIE) of a Silicon-On-Insulator (SOI) wafer. The original wafer was 400 microns thick and was etched down to 40 microns near the openings for sharper edges. To prevent transmission of light through the mask, the front, un-etched side of the mask was coated with a 200 nm gold layer. As with the SP manufactured in 2007, the dashing method was used to prevent the thinning tails of the ripples from becoming waveguides. In the dashing technique, 25-micron diameter circles were placed where the ripple tails would be less than 12-microns wide, and the circles are spaced as to conserve the amount of material that would have been etched out by the ripple tails. The 2007 paper and the milestone whitepaper have a more detailed analysis and figures that show dashing.(1)

Despite the innovation of 2-D optimized masks to account for obstructed apertures, (2) we utilized a nearly identical shaped pupil for this series of tests as we used in 2007 because we were restricted to a transmissive design. The Ripple3 SP used in 2007 had a maximum contrast of 3.0×10^{-10} in the discovery space, had edges 50-microns thick, and was aluminum plated.(1) The main purpose of this HCIT run was to test 2-DM estimation and control methods, so we chose to use the proven SP manufacturing technique from before. Fig. 1 shows the design of the pupil and the wafer layout for etching. Figure 2 shows a microscope image of the final manufactured shaped pupil for the HCIT tests.

5. TESTBED LAYOUT

In the vacuum chamber of the HCIT as diagrammed in Fig. 3, the beam expands to an off-axis parabola (OAP1). A 64×64 Xinetics DM with 1-mm pitch is located at a pupil in the collimated beam and has a 62mm circular stop. The beam reflects off OAP2 and then off DM2 as it converges.

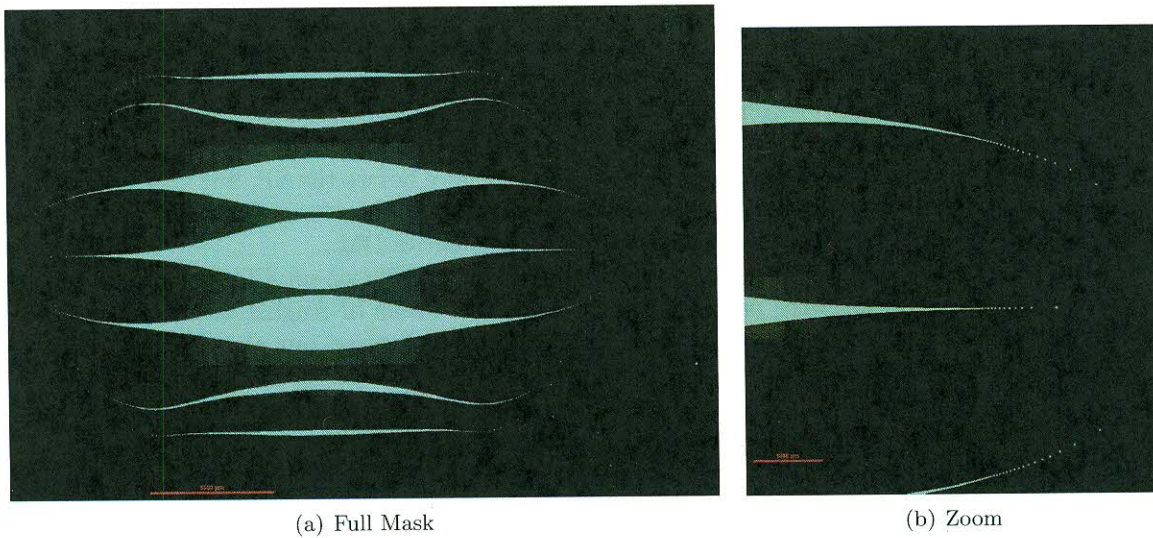


FIGURE 2. A microscope image of the manufactured shaped pupil form the microdevices laboratory. (Left) Full mask image. (Right) Zoom in on the dashing.




FIGURE 3. Optical layout of the JPL HCIT. The collimated beam propagates from DM1 at the first pupil to OAP2. The converging beam hits DM2 before being refocused and passing through a shaped pupil. At the next focus the core of the PSF is blocked with an image plane mask, and the 80° of the 90° search areas are re-imaged onto the final science camera.

The Lyot occulter plane is empty because the shaped pupil by itself requires neither a Lyot occulter nor a Lyot stop. After collimating off OAP3, the collimated beam passes through the shaped pupil. OAP4 then focuses the beam onto a bowtie-shaped focal plane mask whose sole purpose is to block

the core of the PSF for better dynamic range on the science camera. Two more OAPs then re-image the PSF onto the science camera. The camera is on a linear stage to allow for phase retrieval and DM actuator registration before control and estimation procedures begin.

The HCIT could already accommodate shaped pupil coronagraphs, so the only major re-configuration was the inclusion of a second DM for quasi-static speckle suppression. We would have preferred to have DM2 in the collimated beam somewhere before the SP, but that would have necessitated moving all the optics either before or after DM1. Therefore, the fold mirror in the converging beam after OAP2 was replaced with a 32×32 Xinetics DM.

Princeton operated the HCIT remotely by exchanging FITS files via an FTP server at JPL. All of Princeton's two-DM tests were performed in the last two weeks of August 2013, as the earlier part of the summer was dedicated to setup in the HCIT and adapting our algorithms and model for the HCIT optical configuration.

6. TWO DEFORMABLE MIRRORS IN THE HCIT

The two deformable mirrors are by far the two worst optical surfaces in the system. Fig. 4(a) shows the shaped pupil and Fig. 4(b) shows the resulting PSF for a flat wavefront. In Fig. 4(c) we show the image of the shaped pupil with the actual amplitude aberrations (after the phase has been flattened). The main structure of the amplitude errors matches the quilting pattern of the DMs. In Fig. 4(d) we can see a simulated image of the aberrated PSF with an average contrast of 9.5×10^{-6} from 4 to $9.5 \lambda/D$.

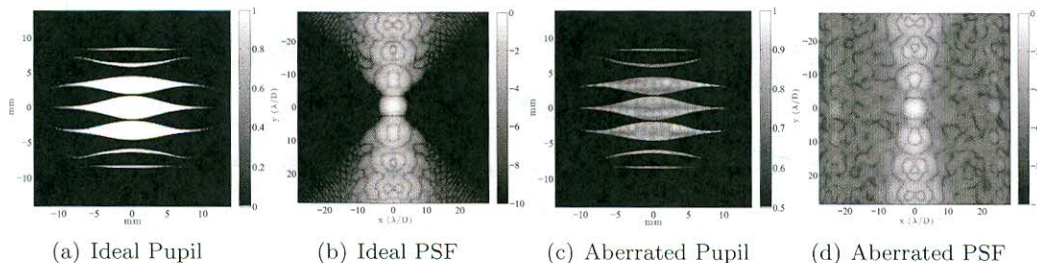


FIGURE 4. The ideal pupil (a) produces two 90° search areas at better than 10^{-10} contrast (b). With the measured amplitude aberrations provided by the phase flattening and retrieval procedures at the HCIT (c), the simulated PSF from the SP has an expected starting contrast of 9.5×10^{-6} (d), which agrees well with the actual measured value of 7.8×10^{-6} . The PSFs are on a log scale shown by the gray-scale bars.

Before testing estimation and control algorithms, we ran established HCIT routines to flatten the exit pupil phase and register the location of the DM actuators. The science camera is on a linear stage, which allowed phase retrieval of the exit pupil electric field. By poking sets of non-neighboring actuators and performing phase retrieval each time, the locations and gains of the actuators for each DM were obtained. The phase at the exit pupil was then flattened and the resulting electric field was re-obtained via phase retrieval. Flattening the pupil phase decreased the starting contrast in the correction region from about 10^{-4} to about 10^{-5} contrast. We also used phase retrieval with and without the SP rotated into the pupil to determine the alignment of the SP.

To construct the control effect matrix (also called the Jacobian or G matrix) in the final image plane for each DM, we utilized the flattened electric field from the phase retrieval as our starting point. Because both DM1 and the SP were in pupils, we simply modeled the SP as being at

DM1 and then performed a Fourier transform to the image plane for each actuator. DM2 was somewhat harder to model. From Zhou and Burge,(3) we see that we can still utilize the Talbot effect for amplitude and phase control across the whole image plane with DM2 in a converging beam. Using geometric optics, we find that we can adequately model DM2 as though it were still in the collimated beam $z_{eq} = 3.88m$ downstream of DM1 based on

$$z_{eq} = f_{OAP2}^2/R, \quad (1)$$

where f_{OAP2} is the focal length of OAP2 and R is the distance of DM2 to the following focus. The advantage of modeling DM2 as being in a collimated beam is that it lets us directly use, without any scaling, the actuator positions obtained via phase retrieval of the exit pupil. The approximation of DM2 being in a collimated beam is slightly inaccurate because the field in the converging beam is supposed to be curved, not flat. Because OAP2 has such a long focal length of $1.54m$ and only about 23 mm of DM2 is illuminated, however, the small curvature at DM2 is negligible. The error that this assumption introduces is on par with the uncertainty in the influence function shape and can thus be ignored.

7. WAVEFRONT ESTIMATION

7.1. Algorithm Development. DM Diversity is the most commonly used technique for wavefront estimation in high contrast imaging experiments. Other sources contain in-depth derivations,(4; 5) so we will just summarize here. DM Diversity is a type of focal plane wavefront estimation in which the wavefront sensor is the science camera to avoid non-common path errors. In DM Diversity a measurement is the image from a DM probe shape $+\phi$ minus the image from the negative of that probe shape $-\phi$; this is called a probe image pair. Two image pairs are needed to invert the measurement and yield the image plane electric field estimate, and three or more image pairs are necessary for a least-squares estimate. Because of the stability of the HCIT and its high-quality camera, only two image pairs are typically used there for each iteration. The main drawback of DM Diversity is that it is a batch process estimator. That is, after each iteration the electric field estimate is thrown away. Because exposures at high contrast in a space observatory will be hours long, it would be beneficial to utilize the previous wavefront estimate data to reduce the number of necessary exposures. This motivated the Kalman filter developed in Groff et al. 2013 that is used as the standard estimation method at the Princeton High Contrast Imaging Laboratory (HCIL).(5) The Kalman filter initializes with DM Diversity for the first control iteration and then requires as few as 1 image pair for each wavefront estimate update.

7.2. Implementation at the HCIT. Before we could compare DM Diversity and the Kalman filter at the HCIT, we had to update our implementation of DM Diversity. In our initial one-sided dark hole tests, our DM Diversity estimator could not replicate the actual image below about 3×10^{-7} contrast, whereas the implementation previously used at the HCIT still matched at contrasts on the order of 10^{-9} . The issue was that the Princeton estimator relied on the linearized system model to predict the change in the image plane's electric field from each probe. The HCIT's method used the intensity data from the actual probe images to compute the amplitude change and the nonlinear model of the system to compute the phase change. Once the Princeton DM Diversity code was changed to follow the HCIT method, our estimator easily allowed us to control down to nearly 10^{-9} contrast. The Kalman filter also requires computing the change in electric field at the image plane from the probes, so we similarly corrected the code for that estimator as well.

The Kalman filter for wavefront estimation needed tuning for the HCIT testbed. We fixed the value of the sensor noise since the readout noise from the camera was known. The process noise value was then tuned until the contrast converged the fastest, which occurred at a value corresponding to about a 3-nm uncertainty in the actuation height. As found by Groff, (5) when using single image pair updates with the Kalman filter, a good probe is necessary for contrast

improvement at each iteration of estimation and control. Using the same probe every iteration or cycling through two or three probes in a set yielded poor contrast convergence at the HCIT. Cycling through four probe shapes gave good results, so we used four probe pairs for all the Kalman filter results presented.

7.3. Comparison. After correcting our algorithms and tuning the Kalman filter, we were able to compare the performance of the two estimators directly. Fig. 5 plots contrast versus number of images for DM Diversity and the Kalman filter, both with stroke minimization as the controller. We do not plot the iteration number because DM Diversity requires five images per iteration (1 image plus 2 probe pairs: $I_0, I_1^+, I_1^-, I_2^+, I_2^-$) whereas the Kalman filter requires 5 images for the first iteration and 3 for each update (1 image plus 1 probe pair: I_0, I_1^+, I_1^-). We see that the Kalman filter starts off slightly slower but then reaches 1×10^{-8} in half the number of images required by DM Diversity. Thus, the Kalman filter is more practical for a space mission than a batch process estimator because it requires far fewer images for estimation.

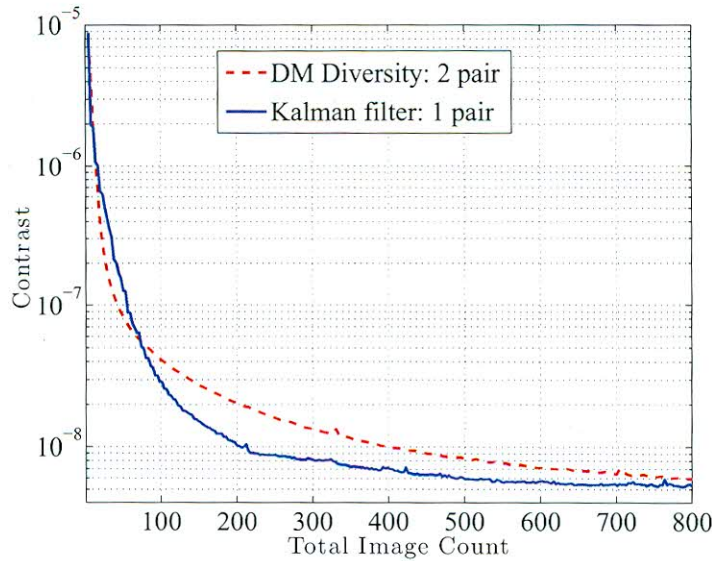


FIGURE 5. Plot showing contrast versus number of images used for two different estimators in the experimental run at the HCIT. The Kalman filter starts worse but eventually requires only about half the images that DM Diversity does to reach a given contrast. Both estimators used stroke minimization for control.

8. WAVEFRONT CONTROL

Assuming a small, complex electric field aberration Ag to the nominal field A and a small DM surface deformation in phase as ϕ , our expression for the electric field at the pupil plane is

$$\begin{aligned} E_{DM}(\xi, \eta) &= A(\xi, \eta)(1 + g(\xi, \eta))e^{i\phi(\xi, \eta)} \\ &\approx A(\xi, \eta)(1 + g(\xi, \eta)) + iA(\xi, \eta)\phi(\xi, \eta) \end{aligned} \quad (2)$$

The propagation to the image plane is encompassed in a linear operator $\mathcal{C}\{\cdot\}$ such that the corrected field at the final image is

$$E_{cor}(x, y) = \mathcal{C}\{A(\xi, \eta)(1 + g(\xi, \eta))\} + i\mathcal{C}\{A(\xi, \eta)\phi(\xi, \eta)\}. \quad (3)$$

Here we will drop the coordinates for succinctness. For small DM actuation, we can approximate the second term as a linear control effect matrix (or Jacobian) G times the vector of actuator commands u ,

$$\mathcal{C}\{A\phi\} = Gu, \quad (4)$$

which lets us write the corrected image plane electric field as

$$E_{cor} = E_{ab} + iGu, \quad (5)$$

where $E_{ab} = \mathcal{C}\{A(1 + g)\}$ is the aberrated electric field without DM actuation.

8.1. EFC. This derivation of Electric Field Conjugation (EFC) follows that of Give'on except that this notation does not split complex terms into real and imaginary parts.(4) In EFC we wish to cancel the electric field to create a dark hole in a certain region. By setting $E_{cor} = 0$, we find from Eq. 5 that the optimal (real-valued) command to minimize the energy in the dark hole is

$$\begin{aligned} u_{opt} &= \mathcal{R}\{iG^{-1}E_{ab}\} \\ &= -\mathcal{I}\{G^{-1}E_{ab}\}, \end{aligned} \quad (6)$$

where $\mathcal{R}\{\cdot\}$ and $\mathcal{I}\{\cdot\}$ yield the real and imaginary parts, respectively. The matrix G is usually rectangular and over-determined, so we must take the left pseudoinverse, which is equivalent to a least-squares error optimization, instead of the inverse

$$\begin{aligned} u_{opt} &= -\mathcal{I}\{G^L E_{ab}\} \\ &= -(G^*G)^{-1}\mathcal{I}\{G^*E_{ab}\}. \end{aligned} \quad (7)$$

To damp the stroke on the actuators, Give'on augmented the matrices G and E_{ab} into

$$\tilde{G} = \begin{bmatrix} G \\ \alpha\mathbb{I} \end{bmatrix} \quad (8)$$

and

$$\tilde{E}_{ab} = \begin{bmatrix} E_{ab} \\ \mathbf{0} \end{bmatrix}. \quad (9)$$

Substituting \tilde{G} for G and \tilde{E}_{ab} for E_{ab} in Eq. 7, we find that the optimal command to conjugate the electric field while limiting DM stroke is

$$u_{opt} = -(G^*G + \alpha^2\mathbb{I})^{-1}\mathcal{I}\{G^*E_{ab}\}, \quad (10)$$

which we recognize as a Tikhonov regularization with α^2 as our regularization parameter.

8.2. Stroke Minimization. Another approach, developed by Pueyo et al.,(6) is to minimize the actuators strokes directly while requiring a contrast goal at each iteration to be met. We can write this as

$$\begin{aligned} \text{minimize } u^T u &= \sum_{k=1}^{N_{act}} a_k^2 \\ \text{subject to } C_{DH} &\leq C_{target}, \end{aligned} \quad (11)$$

where a_k is the command to the k th actuator, C_{DH} is the contrast in the dark hole, and C_{target} is the target contrast.

The intensity at any point in the image plane is the square of the magnitude of the linearized electric field

$$I_{im}(x, y) = |\mathcal{C}\{A(\xi, \eta)(1 + g(\xi, \eta))\} + i\mathcal{C}\{A(\xi, \eta)\phi(\xi, \eta)\}|^2. \quad (12)$$

The contrast is the sum of the pixel intensities in the dark hole divided by the number of pixels, N_{pix} , and the maximum intensity of the star, I_{00} . We then have an expression for contrast that is quadratic in u

$$\begin{aligned} C_{DH} &= \|I_{im}(x, y)\|_{DH}/N_{pix}I_{00} \\ &= (\langle E_{ab}, E_{ab} \rangle + 2\mathcal{R}\{\langle E_{ab}, iGu \rangle\} + \langle iGu, iGu \rangle)/N_{pix}I_{00} \\ &= (u^T G^* G u + 2u^T \mathcal{I}\{G^* E_{ab}\} + E_{ab}^* E_{ab})/N_{pix}I_{00}. \end{aligned} \quad (13)$$

We use Eq. 13 to write the cost function representing Eq. 11 as

$$\begin{aligned} J &= u^T u + \mu_0(C_{DH} - C_{target}) \\ &= \left(u^T (N_{pix}I_{00}\mathbb{I} + \mu_0 G^* G) u + 2\mu_0 u^T \mathcal{I}\{G^* E_{ab}\} + \mu_0 (E_{ab}^* E_{ab} - N_{pix}I_{00}C_{target}) \right) / N_{pix}I_{00}. \end{aligned} \quad (14)$$

To find the control that yields the minimum of this cost, we take the derivative with respect to u and set equal to zero

$$\begin{aligned} \frac{\partial J}{\partial u} = 0 &= (2/N_{pix}I_{00})((N_{pix}I_{00}\mathbb{I} + \mu_0 G^* G)u_{opt} + 2\mu_0 \mathcal{I}\{G^* E_{ab}\}) \\ 0 &= ((N_{pix}I_{00}/\mu_0)\mathbb{I} + G^* G)u_{opt} + \mathcal{I}\{G^* E_{ab}\}, \end{aligned} \quad (15)$$

before rearranging to obtain

$$u_{opt} = -(G^* G + (N_{pix}I_{00}/\mu_0)\mathbb{I})^{-1} \mathcal{I}\{G^* E_{ab}\}. \quad (16)$$

Comparing the optimal commands from Eqs. 10 and 16, we see that they are equivalent with the regularization parameter $N_{pix}I_{00}/\mu_0 = \alpha^2$. Thus, for a single DM in monochromatic light, the control command for stroke minimization and regularized EFC are exactly equivalent.

The method to extend EFC at the HCIT and stroke minimization at Princeton to two DMs was the same. The Jacobians for both DMs, G_1 and G_2 were combined into an augmented control effect matrix

$$G = [G_1 \ G_2]. \quad (17)$$

It is possible to choose separate optimization parameters for each DM, but both the HCIT and Princeton used a single variable for both DMs. Thus, for 2-DM control in monochromatic light, EFC and stroke minimization stay in the equivalent forms of Eqs. 10 and 16, respectively. Groff et al.(7) go into additional detail on the differences between EFC and stroke minimization.

8.3. Comparison of Lab Results. Despite having the same derivation for monochromatic light, the two control techniques differ in their lab implementation and thus yield different results. At the HCIT, the methodology with EFC is to keep the regularization parameter α^2 as small as possible so that the contrast is corrected faster at the expense of more actuator stroke. Actuators partially or fully obscured by the shaped pupil (or the Lyot stop for a Lyot coronagraph) are used only at medium to high contrast values to prevent large strokes for low-contrast correction. Typically, the lab operator allows more actuators to be used once the contrast stops converging at the previous setting. The best regularization parameter at each control iteration is determined experimentally with a regularization test, in which images are taken at several α^2 values and the image with the best contrast determines the next regularization value. This method is sufficient for lab purposes so long as hysteresis of the actuators is not a concern, but it would not be used in a space mission because of the increased number of images required.

Princeton implements stroke minimization with the opposite approach of EFC. The regularization parameter starts high (for more damping on the actuators) and is decreased until the contrast

target for the current iteration is met (as indicated in Eq. 11). Thus, all actuators can be used at each step and the linearized model is more likely to hold because of the small strokes used. In Princeton's HCIL, we can use the model of the optical system to determine the best regularization parameter at each iteration and have achieved our best contrast levels in only 30 control iterations. (5) In our limited time at the HCIT, we did not get the model-based line search on μ_0 working, so we instead used the regularization test method for choosing μ_0 .




FIGURE 6. (a) Corrected image after using EFC. Average contrast from 5 to 9 λ/D is 4.15×10^{-9} . (b) and (c) Voltage maps from EFC for DM1 and DM2, respectively. (d) Corrected image after using stroke minimization. Average contrast from 5 to 9.5 λ/D is 3.58×10^{-9} . (e) Voltage maps from stroke minimization for DM1 and DM2, respectively. Both (a) and (d) are shown in log scale. Note that the upper and lower scales are different, highlighting the much smaller stroke from stroke minimization.

The main implementation differences between the HCIT's EFC and Princeton's stroke minimization methods were that the HCIT limited the use of partially obscured actuators and they re-linearized about the new DM settings whenever they included more actuators in their control matrices. Because of the limited time in which Princeton's tests were conducted, we did not try to re-linearize about the previous DM settings once the contrast began to converge slowly. Re-linearization may indeed speed up the convergence of the control experiment at the expense of only about 15 minutes per recalculation of the Jacobians on an 8-core workstation running MATLAB. In Fig. 6, we compare the best results obtained using EFC and stroke minimization. Both methods achieved essentially the same final contrast, with EFC reaching 4.15×10^{-9} in Fig. 6(a) and stroke minimization reaching in 3.58×10^{-9} in Fig. 6(d). The strokes on the DMs were significantly different, however. EFC required a voltage actuation range on DM1 of 66.2 V versus 15.4 V for stroke minimization. These correspond to expected phase deformation ranges of 0.18λ and 0.07λ ,

respectively. The strokes for DM2 had ranges of 15.2 V (0.18λ) for EFC and 8.1 V (0.10λ) for stroke minimization. In the other 2-DM test using EFC, there were fewer manual adjustments of which actuators were used. The final contrast in that case was 6.56×10^{-9} and required DM stroke ranges about $6\times$ higher than those used with stroke minimization. The results clearly show that stroke minimization uses much smaller stroke than EFC, even when EFC is carefully adjusted. In a space mission it will be important to keep the actuator strokes as small as possible for a model-based calculation of the best regularization parameter at each step. Regularization tests are impractical for time-constrained space missions, so the larger actuations that EFC tends to make are less compatible with the DM superposition model currently used.

It is important to note that because of the limited time available a full exploration of all possible parameters in each method was not attempted. As a result, we are not drawing broad conclusions that stroke minimization is always better. Rather, it provided a new methodology for exploring a new part of the parameter space that demonstrated the same contrast is possible with much smaller stroke on the DMs. Future experimental work would benefit from a careful optimization and comparison of these two methods.

9. RESULTS

9.1. Experimental Results. Because we used most of the summer adapting and troubleshooting our algorithms for the HCIT, we had only two weeks to perform all our 2-DM tests. We therefore chose to utilize that full period for faster monochromatic tests of our control and estimation algorithms. The much lower throughput for broadband experiments would not have allowed us to tune the Kalman filter and perform several full control runs in the time allotted.

As described in previous sections, our best contrast achieved in symmetric dark holes was 3.6×10^{-9} in monochromatic light. Fig. 7(a) shows the full image before correction and Fig. 7(b) shows after correction. The correction region extended from the radius of the FPM ($r = 4.86\lambda/D$, since the 808 nm laser was swapped for a higher-power, more stable 790 nm laser) to $x = 9.5\lambda/D$ over the full 80° arcs passed by the FPM. The final contrast reported is in the region $r = 5.0\lambda/D$ to $x = 9.5\lambda/D$ over 75° arcs so that pixels covered by the FPM are not accidentally included in the contrast measurement. Fig. 8 shows a histogram of the contrast achieved by pixel. This provides a quantitative picture of the spread of contrast across the dark hole. It shows that the contrast is tightly clustered around the mean value and better with a relatively small number of pixels at contrasts greater than 5×10^{-9} .

9.2. Limitations. The dashing technique used to manufacture the ripple masks could be limiting the achievable contrast with SPs. Neither in this HCIT run nor the 2007 one did we reach contrasts below 10^{-9} , and both SPs used the same manufacturing method. Note that this is a concern only for this type of mask; future reflective shaped pupils will not require dashing of small features in their manufacture.

At the masked center of the PSF in Fig. 7(b), one can see light at 6×10^{-9} contrast. This spot appeared centered at the same location in images with single-sided correction, so we do not believe it is light diffracting around the FPM. The light is not a Poisson spot because the FPM is at a focal plane, not a pupil plane. We therefore believe that the spot is the brightest part of the PSF leaking through the mask, meaning that the attenuation is 6×10^{-9} where the silicon FPM is the full 400 microns thick. Near the edges of the FPM the silicon is only 40 microns thick, but the PSF is not as bright there and the 200 nm gold layer should be providing most of the attenuation. Thus, the transmitted light should still be far below the 10^{-9} levels we are reaching.

In Fig. 7(b) we can also see streaks of stray light where the FPM should be blocking the PSF. The vertical lines in several locations are clearly just saturated pixels bleeding onto their neighbors. However, the diagonal lines rippling away from the edges of the FPM indicate that light is diffracting around the FPM or that stray light is in the system. If there indeed is diffraction

