Performance of TPF's High-Contrast Imaging Testbed: Modeling and simulations

Erkin Sidick*, Fang Shi, Scott Basinger, Dwight Moody, Andrew E. Lowman, Andreas C. Kuhnert, and John T. Trauger
Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, USA 91109

ABSTRACT

The performance of the high-contrast imaging testbed (HCIT) at JPL is investigated through optical modeling and simulations. The analytical tool is an optical simulation algorithm developed by combining the HCIT's optical model with a speckle-nulling algorithm that operates directly on coronagraphic images, an algorithm identical to the one currently being used on the HCIT to actively suppress scattered light via a deformable mirror. It is capable of performing full three-dimensional end-to-end near-field diffraction analysis on the HCIT's optical system. By conducting extensive speckle-nulling optimization, we clarify the HCIT's capability and limitations in terms of its contrast performance under various realistic conditions. Considered cases include non-ideal occulting masks, such as a mask with parasitic phase-delay errors (i.e., a not band-limited occulting mask) and one with damped ripples in its transmittance profiles, as well as the phase errors of all optics. Most of the information gathered on the HCIT's optical components through measurement and characterization over the last several years at JPL has been used in this analysis to make the predictions as accurate as possible. Our simulations predict that the contrast values obtainable on the HCIT with narrow-band (monochromatic) illumination at 785nm wavelength are $C_m=1.58\times10^{-11}$ (mean) and $C_4=5.11\times10^{-11}$(at $4\lambda/D$), in contrast to the measured results of $C_m\sim6\times10^{-10}$ and $C_4\sim8\times10^{-10}$, respectively. In this paper we report our findings about the monochromatic light performance of the HCIT. We will describe the results of our investigation about the HCIT's broad-band performance in an upcoming paper.

Keywords: TPF, integrated modeling, speckle-nulling, high-contrast coronagraph

1. INTRODUCTION

High-contrast imaging testbed (HCIT) at JPL is Terrestrial Planet Finder (TPF) Coronagraph's primary platform for experimentation [1-2]. It is used to provide laboratory validation of key technologies as well as demonstration of a flight-traceable approach to implementation. In order to validate the HCIT's performance through modeling and error budget analysis, we have implemented an optical simulation algorithm by combining the HCIT's optical model [3] with a speckle-nulling algorithm. This simulation algorithm uses MACOS (Modeling and Analysis for Controlled Optical Systems) [4], a versatile optical modeling tool developed at JPL and used in many flight projects. Hence it is capable of performing full near-field diffraction analysis on the HCIT's optical system. Such a diffraction analysis is required to evaluate the HCIT's performance in terms of design tolerances and based on its ultimate metric—the contrast ratio. MACOS also allows seamless interfacing with structural and thermal models, thus eventually enabling us to create fully integrated optical/structural/thermal models to evaluate the dynamic effects of disturbances on the overall performance of the HCIT.

We have used our HCIT simulation algorithm to investigate the performance of the HCIT's optical system. The structural design of the optical system as well as the parameters of various optical elements used in the analysis are almost identical to those of the implemented HCIT, and the simulation takes into account the measured surface errors of the off-axis parabolic mirrors and the flat mirrors as well as the parasitic phase errors of the occulting mask that have been independently measured. We have also examined the effects of smoothed ripples in the optical density (OD) profiles of the occulting masks fabricated at JPL on the HCIT's performance. The initial results of our modeling predict a monochromatic mean contrast of 2×10^{-11}, and a contrast of 5×10^{-11} at an angular separation of $4\lambda/D$ within a "half dark

---

*Erkin.Sidick@jpl.nasa.gov; Phone 1 818 393-7585; Fax 1 818 393-9471; www.jpl.nasa.gov

hole". These results are better than the experimental results obtained to date on the HCIT by more than an order of magnitude. This paper is about the HCIT's performance under the monochromatic beam illumination. We will report on the HCIT's broad-band performance in our future publication.

2. OPTICAL SYSTEM AND COMPONENTS

2.1 The HCIT Optical System

The MACOS optical model realistically mimics the HCIT. The centerpiece of the HCIT is a 5x7 foot optical table, on which are installed the coronagraph elements as schematically shown in Figure 1. The whole optical table sits in a thermally controlled vacuum tank evacuated with a scroll pump to 10 milli-Torr or better. The tank provides the desired environment, including vibration isolation, atmospheric turbulence isolation and sub-Kelvin thermal stability. An artificial starlight is created by a 5μm pinhole illuminated by an optical fiber. Various light sources can be fed into the other end of the fiber outside the vacuum tank, such as a narrow band laser beam or a broadband light source with 20% or more bandwidth. An off-axis parabolic mirror (OAP1) collimates the light from the pinhole and directs it to a high-density deformable mirror (DM), which performs the wavefront control. The aperture mask on the DM defines the system pupil of the HCIT. After the DM, the collimated light is re-imaged onto the focal plane of the occulting mask by OAP2 and a flat mirror (FM1). The occulting mask attenuates the starlight, and almost has no effect on the light of a planet if present. The "back-end" of the system, from the occulting mask to the back focus plane, supports experimentation with diverse coronagraph configurations and apodizations. This optical setup corresponds to a classical Lyot coronagraph: an occulting mask at the highly-corrected focal plane plus a Lyot mask at the subsequent pupil plane. A flat mirror (FM2) and OAP3 re-collimate the light passing through the occulter mask and form a same-size sharp image of the DM pupil at the Lyot plane. A Lyot stop blocks the ring-like residual light diffracted off the occulter mask

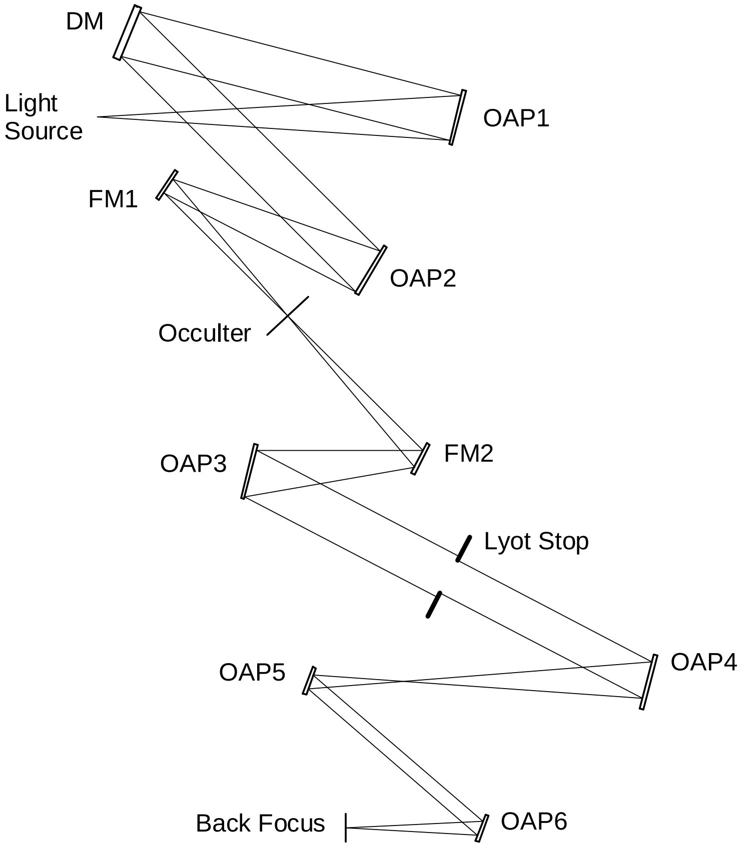

Figure 1. Schematic diagram of the High Contrast Imaging Testbed layout. The light source ("starlight") is a 5μm pinhole illuminated by an optical fiber, and a CCD science camera is located at the back focal plane for detecting the image of the "starlight".

while letting most of the planet light through. After OAP4 forms an image from the remaining stellar and planet lights, it is then magnified (M ≈3) by the OPA5-OAP6 pair for proper sampling on the CCD science camera located at the back focal plane. On the testbed, the wavefront control software takes the starlight image captured by the CCD camera as an input, finds a new set of DM control commands from it, then updates the DM settings with the new set of control commands.

In the following several sub-sections, we describe the characteristics and the physical parameters of various optical components used on the testbed. The various parameters to be described are the ones used in our modeling and simulations, and they are matched to the designed and/or measured parameters of the actual optical components as close as possible. We use an aperture mask diameter D=30mm throughout our analysis.

2.2 Occulting Mask

One of the occulting masks implemented on the HCIT is a gray-scale pattern written on high-energy e-beam sensitive (HEBS) glass with the 100 keV e-beam lithography facility at JPL [5]. It has a linear sinc^2 intensity profile with a one-dimensional transmittance profile $T(x)$ given by [6]:

$$T(x) = \left[1 - \left(\frac{\sin(\pi x / w)}{(\pi x / w)} \right)^2 \right]^2 . \tag{1}$$

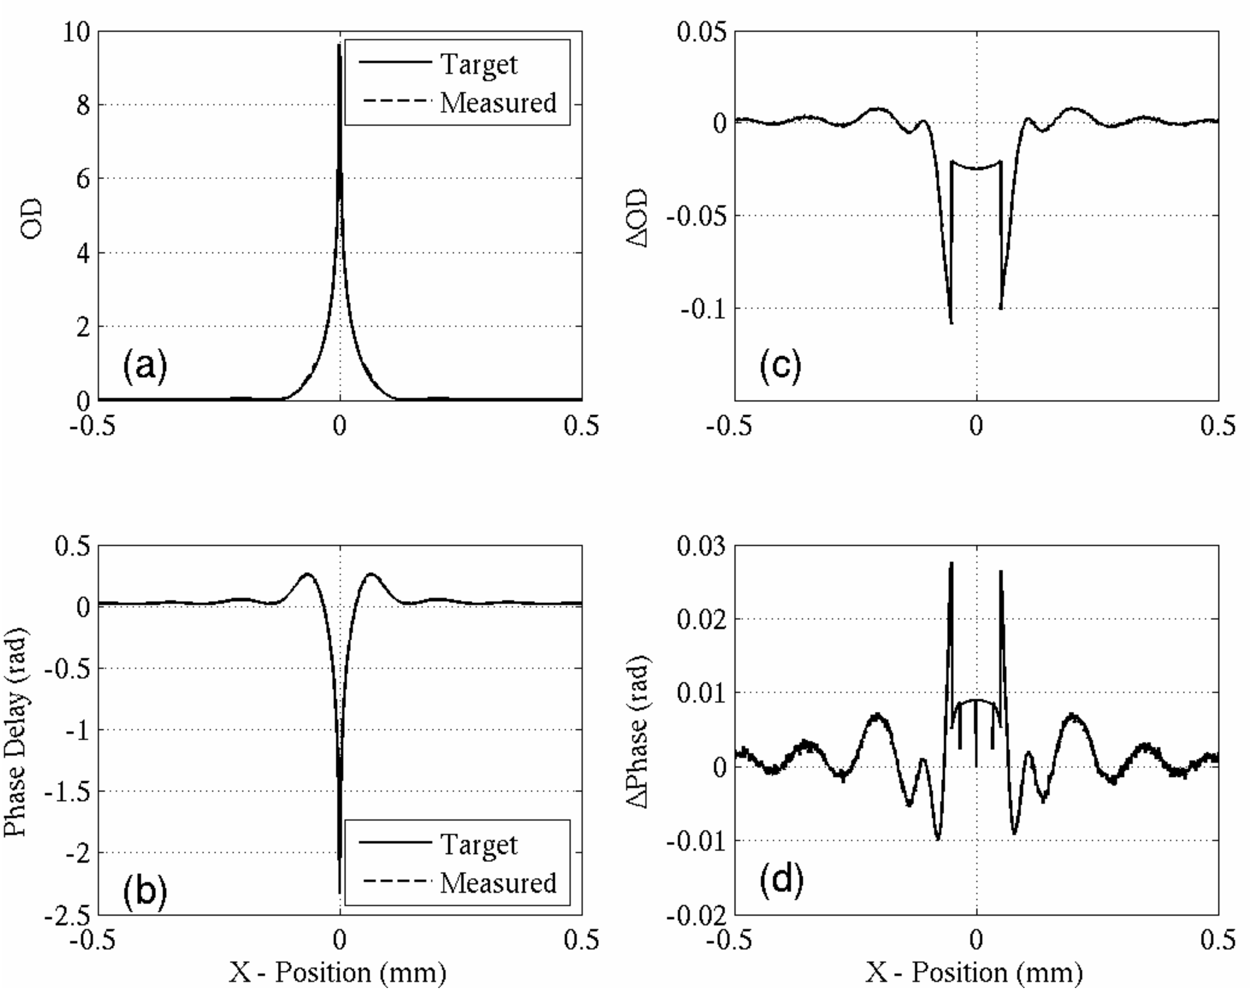

Figure 2. Target and measured profiles of a linear sinc^2 occulting mask OD and phase delay. (a) OD profiles; (b) Phase-delay profiles; (c) Difference in OD (target-measured); and (d) Difference in phase-delay (target-measured).

The occulting mask used in the current analysis has a design width of w=141.9μm. A group at JPL has also characterized the fundamental optical properties of the gray-scale occulting masks written on HEBS glass through measurement of the optical density (OD), the phase delay and the optical constants of the mask material and/or direct measurement of fabricated occulting masks [5, 7]. Figure 2 shows the target (or designed) and the measured OD and phase profiles as well as their differences for a linear sinc^2 occulting mask (Eqn. 1). The targeted width of this mask is w=141.9μm.

There are two aspects of the above data that need to be explained. (i) The OD profile measurement was carried out on an occulting mask made with a Canyon Materials' glass HEBSSPT10. The experimental setup for measuring such OD profiles is described in reference [5]. The particular OD data shown here were taken at a wavelength of λ=785 nm. This setup has a noise floor of OD~3. It was found that the measured OD profile starts to deviate from the target at OD > 1, and such deviation becomes significant at OD~3 and above. Therefore, the measured OD values in Fig. 2 corresponding to OD > 1 have been obtained by fitting a model of the form given by Eqn. 1 to the measured data at OD ≤ 1. Obtained occulting mask width of this fit is w=144.0μm. Therefore, in the region where OD > 1 in Figs. 2(a) and 2(c), the difference in OD profiles is actually the difference in theoretical profiles calculated with w=141.9μm and w=144.0μm, respectively. (ii) The phase measurement was carried out with HEBS glass N20G158, darkened using 100 keV electron beam, 20 nA, at selected wavelengths of 532, 635, 785 and 830 nm. Details of this measurement are given in reference [7]. We have developed an analytical model for the phase versus OD curve at any given wavelength based on these data, and used it to obtain the phase versus position curves shown in Figs. 2(b) and 2(c). That is, the phase curves in Fig. 2 are an estimate that is based on the measured data of a different type of HEBS materials. A linear sinc^2 occulting mask being used on the testbed was made of an HEBS glass that is different from both HEBSSPT10 and N20G158. But in our analysis we always used occulting mask phase profiles that were estimated with the above approach. That is, the occulting mask phase profiles used in our analysis are not identical to those of the actual devices, but we expect they are good estimates of what we really have on our testbed.

2.3 Lyot Stop

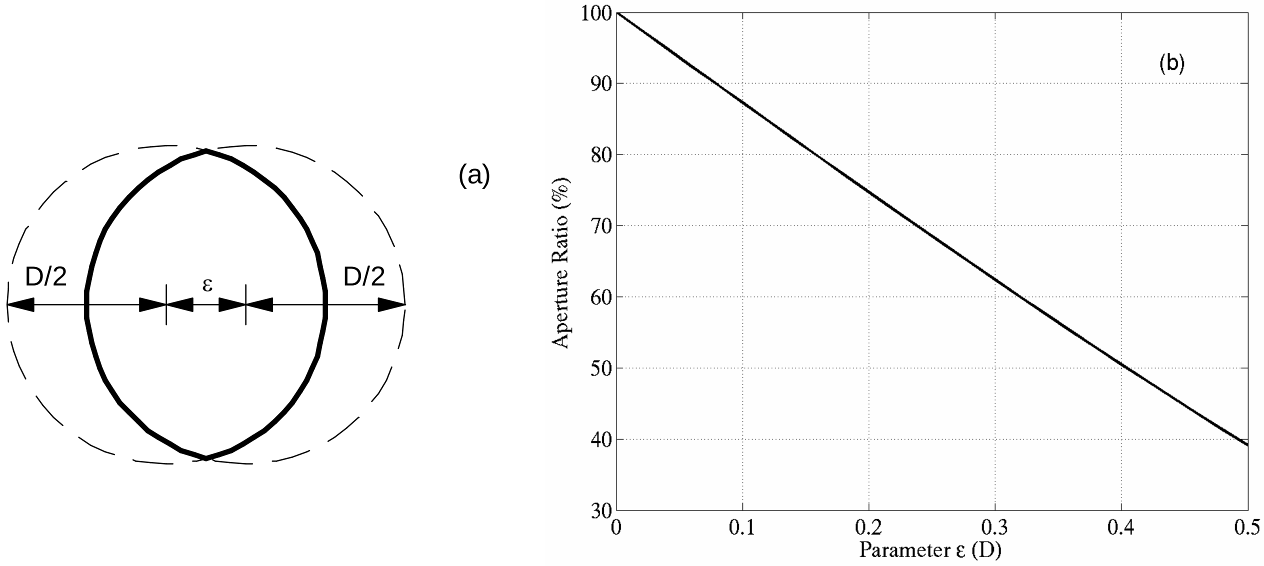

Figure 3. (a) The shape of the Lyot stop opening and the definition of Lyot stop parameter ε; (b) The ratio of Lyot stop aperture to the system aperture as a function of parameter ε.

Our Lyot stop is a simple blackened piece of sheet metal with a sharp edge. Its opening (Lyot stop aperture) has an eye-shape defined by two circles that are shifted with respect to each other in the horizontal direction by a distance of ε in units of *D*, the system aperture diameter, as shown in Fig. 3(a). The value of ε needs to be chosen based on the value of the occulting mask width parameter *w* [6], and these two parameters together determine the throughput characteristics of the coronagraph [6, 8]. Figure 3(b) shows the ratio of the Lyot stop opening to the system aperture as a function of ε. It can be considered as an estimate of this coronagraph's throughput. The Lyot stop used on our testbed with the *w*=141.9μm occulting mask has ε=0.4 and *D*=30mm.

2.4 OAPs and Flat Mirrors

In order to fully understand the optical characteristics of the HCIT, all OAP mirrors and flats have been carefully measured and characterized using a Zygo interferometer before they were put into the testbed [9]. This also allowed us to conveniently put the measurement data into our optical model. In our MACOS based optical analysis tool, we can include the measured phase error data of only one component at a time to study its effects on the coronagraph performance separately, or include all of them together to evaluate the overall system performance. Figure 4 shows the phase errors of the optics presently being used in the HCIT that are "seen" at the exit pupil of the optical system when no occulting mask and Lyot stop are present. Among them, Figs. 4(a)-(h) were obtained by "turning on" the measured phase errors of only one component at a time, and the one in Fig. 4(i) was obtained with all the phase error data "turned on". As can be seen, the HCIT optical system's combined phase error has an RMS value of ~36nm.

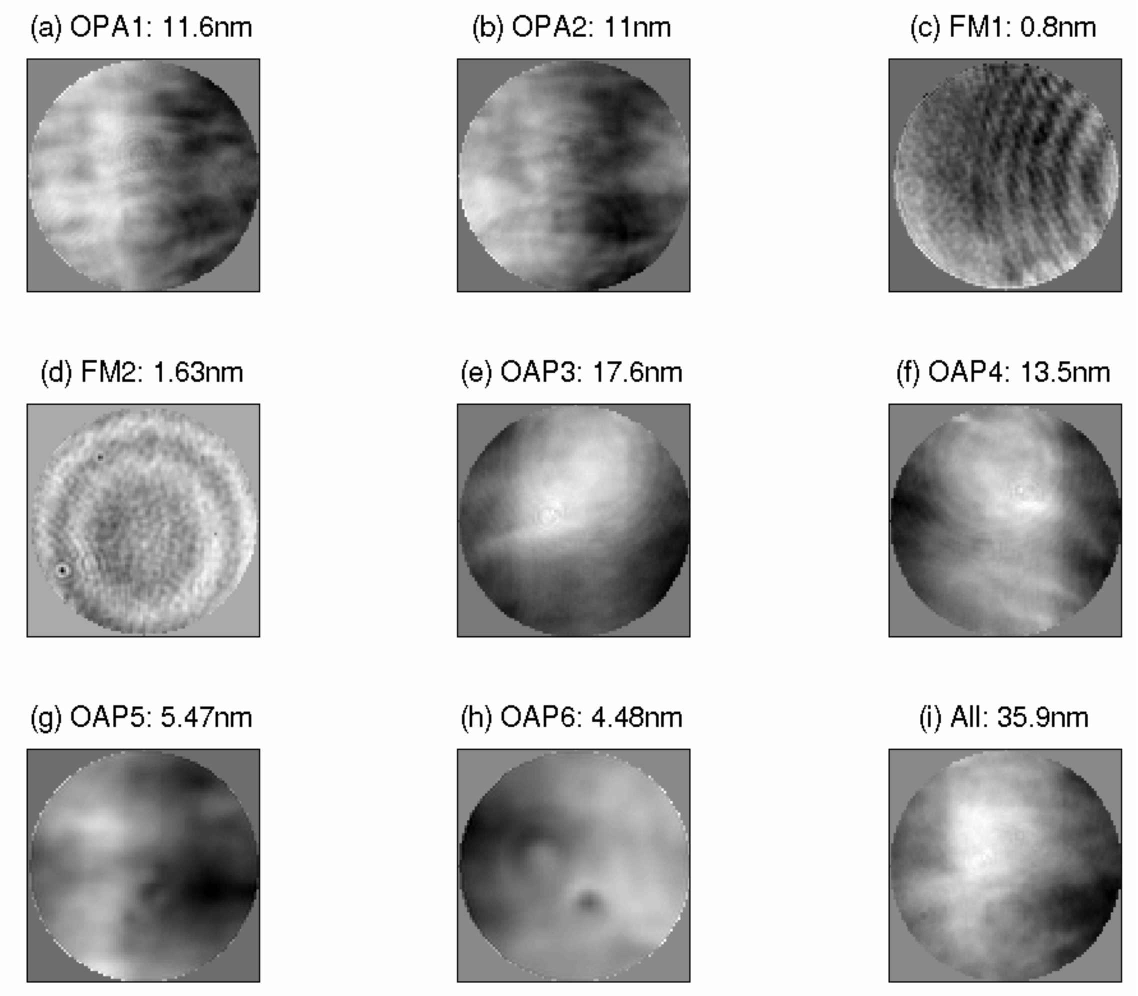

Figure 4. Wavefront error maps "seen" at the exit pupil of the HCIT optical system when no occulting mask and Lyot stop are present. Parts (a)-(h) were obtained by "turning on" the Zygo measurement data of one optic at a time in the MACOS optical model, and the last part was obtained by including the phase error data of all optics at once. Also shown in title of each figure is the RMS value of wavefront error map shown in the corresponding figure.

2.5 Light Source

The HCIT was equipped with several light sources to create "starlight" through a 5μm pinhole, including a white light source consisting of a supercontinuum combined with neutral density and spectral filters and fiber-coupled laser diodes having different center wavelengths. In this paper we consider only a 785nm laser source in our simulations. It has been found that the wavefront at the HCIT pupil develops some amplitude errors, including low-order droop, structural bands and high-order variations [9]. Among them, the low-order intensity variation of the pupil image comes mainly from the non-uniform illumination pattern of the artificial star source. Our MACOS based optical model has the capability to include such effects with varying degrees. Assume the HCIT pupil aperture $P(u,v)$ is defined by a non-uniform amplitude distribution $A(u,v)$ and a wavefront error distribution $\phi(u,v)$, that is

$$P(u,v) = A(u,v)\,\mathrm{e}^{j\phi(u,v)}. \qquad (2)$$

Figure 5 shows an example of the wavefront amplitude profile $A(u,0)$ of the HCIT. It was obtained with our MACOS optical model by setting the phase errors of all optics to zero. In this example, the amplitude drooping is 1.64% and we use this value throughout our analysis except stated otherwise. However, we assume that all of the optical components are free of wavefront amplitude errors.

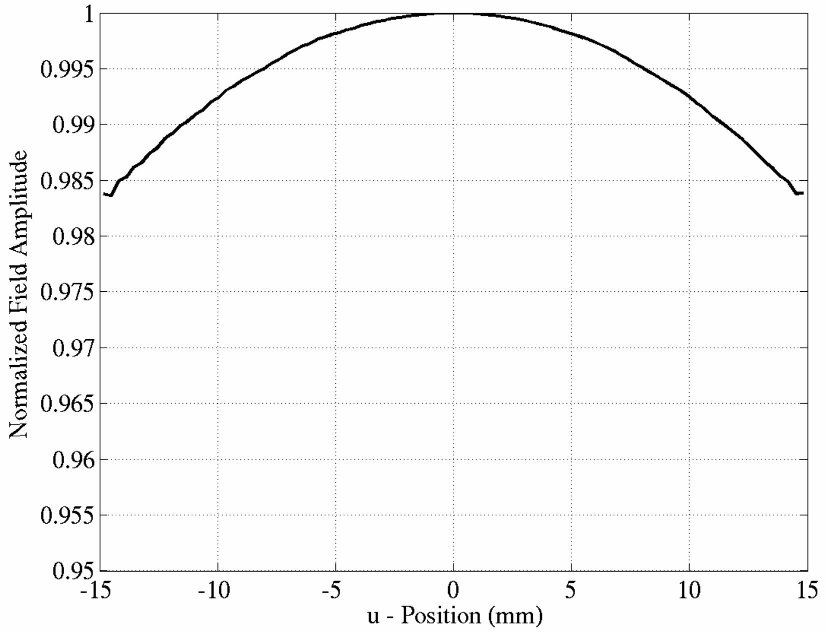

Figure 5. Drooping of the wavefront amplitude at the HCIT pupil due to the non-uniform illumination pattern of the artificial star source.

3. RESULTS

A detailed and realistic optical model of the HCIT was set up by using its "as-built" prescription, and a speckle-nulling algorithm [1] was also implemented. This algorithm is identical in principle to what was implemented on the testbed, with some differences in the implementation detail. The fundamental principle of this algorithm was explained in reference [1] and will not be repeated here. In a typical speckle-nulling process on the testbed, the starlight image intensity is measured with the CCD science camera in the back focal plane, it is processed to obtain a new corrective phase map, the DM control setting is modified with this new information, and these iteration steps are repeated until the contrast reaches a desired value or the speckle-nulling process becomes ineffective. In our simulations, on the other hand, the final image plane intensity map is obtained by running MACOS with a given set of parameters. By carrying out end-to-end full diffractive propagation analysis with this optical model, we are not only able to evaluate the HCIT's performance under various different conditions, but also to simulate the speckle-nulling optimization process and make reasonable predictions on its capability in achieving our contrast target.

The DM used on the HCIT has 1024 (32x32) actuators arrayed on a 1mm pitch. It can produce a high-contrast half dark field in an area defined the by the controllable spatial frequency of the DM. The targeted high-contrast window in the HCIT coronagraph image is [xmin xmax ymin ymax] = [4 10 -10 10]*λ/D, with Mean Contrast and Contrast at 4λ/D as the ultimate metrics of the HCIT's performance. Here, the Mean Contrast is defined as the contrast value averaged over the window of [4 10 -10 10]*λ/D, and the Contrast at 4λ/D as the one averaged over a smaller window of [4 5 -0.5 0.5]*λ/D. And the quantity Contrast is calculated using Eqn. 12 in reference [8], but without any averaging.

In our simulations, the speckle-nulling is carried out with the following steps: 1) Adjust the de-space distance of the occulting mask from its nominal position so that the total energy at the final focal plane (or the scattered light) is minimized (usually 0 ~ 2mm). 2) Carry out speckle-nulling for 5 iterations in the full right half-plane, for 5 iterations in the full left half-plane, and repeat these two sub-steps one more time. This step reduces the amount of light leaking through the Lyot Stop. 3) Continue the speckle-nulling in the full right half-plane for 100 iterations. In this case, the size of the speckle-nulling window is determined by the controllable spatial-frequency bandwidth of the DM. 4) Continue the speckle-nulling inside a reduced window [3 11 -11 11]*λ/D for several hundred iterations. 5) Conduct the rest of speckle-nulling iterations in an even smaller window of [3 10 -10 10]*λ/D. There is an additional parameter except the speckle-nulling window that is needed in the speckle-nulling algorithm, namely, the "gain factor" γ. It is used to scale the correcting phase map (or DM corrective surface pattern) obtained by processing the coronagraph image intensity. The right value for γ depends on the implementation details of the algorithm, and should be determined through numerical experimentation. The goal is to make the speckle-nulling process as efficient as possible. It should be pointed out that the MACOS-based speckle-nulling simulation process is fairly time-consuming. In our case, for a grid size of 512x512 and with an Intel Xeon 3.2GHz dual-processor workstation, it took about two hours for 100 speckle-nulling iterations.

3.1 Ideal (band-limited) occulting mask

First, we consider several ideal cases of the occulting mask to understand the contrast noise floor and the ultimate performance capability of the HCIT. An ideal (band-limited) linear-sinc^2 occulting mask has a transmittance as defined by Eqn. 1 with its phase-delay the same at all OD values. In this part of our simulations, we set the occulting mask phase to zero, and evaluated the following three cases: 1) The light source does not exhibit drooping ("No Drooping") at the pupil, and all optics are free of phase errors ("No Zygo"). 2) The light source exhibits drooping as shown in Fig. 5, but all of the optics are free of phase errors. 3) The light source exhibits drooping, and all of the optics have phase errors as measured with the Zygo interferometer ("With Zygo") and shown in Fig. 4. Among these 3 cases, the speckle-nulling simulations were carried out for only the latter two cases. The detector noise is not taken into account in this paper.

Figure 6 displays the process of speckle-nulling in terms of Mean Contrast value. In this figure, the letter "a" indicates the point where the speckle-nulling window was changed from the full right half-plane to [3 11 -11 11]*λ/D rectangle, and the letter "b" indicates the point where this window was further reduced to [3 10 -10 10]*λ/D. Figure 7 shows the final results of the contrast map obtained for the above 3 cases. As we can see from these figures, the HCIT optical model has a contrast noise floor of C_m~1x10^{-14}. The wavefront amplitude drooping caused by non-uniform illumination raises this floor to C_m=1.75x10^{-11}, but it can be improved to C_m=2.56x10^{-13} by speckle-nulling for 700 iterations. That is, scattered light caused by the amplitude drooping of the source can be eliminated to certain degree through speckle-nulling. Here we stopped the speckle-nulling optimization after a total of 700 iterations because no much improvement is expected beyond that point. The intrinsic phase errors introduced by all of the optics into the optical system greatly degrade the contrast performance of the HCIT as expected. In this particular case, the mean contrast changes from C_m=1.75x10^{-11} to C_m~1.0x10^{-8} just because of the phase errors of all the optics. We could have increased this contrast value to C_m=2.75x10^{-11} after a total of 1000 speckle-nulling iterations, as shown in Figs. 6 and 7.

3.2 Realistic occulting mask

Next, we consider several cases of realistic occulting mask. In all of the analyses below, the amplitude drooping of the light source exhibited at the pupil is also taken into account. Modeling and simulations in this subsection help us understand what behavior we can expect from the HCIT under various practical conditions. Figures 8 and 9 show that

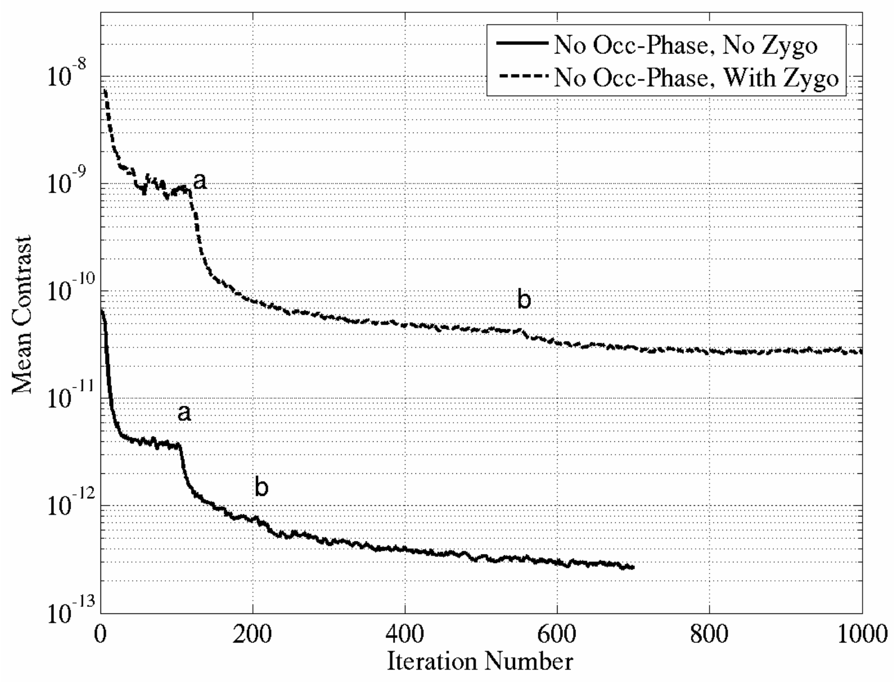

Figure 6. Mean Contrast during the speckle-nullling process. Letter "a" indicates the point where the speckle-nulling window was reduced from the full right-half plane to [3 11 -11 11]*λ/D, and the letter "b" indicates the point where it was further reduced to [3 10 -10 10]*λ/D.

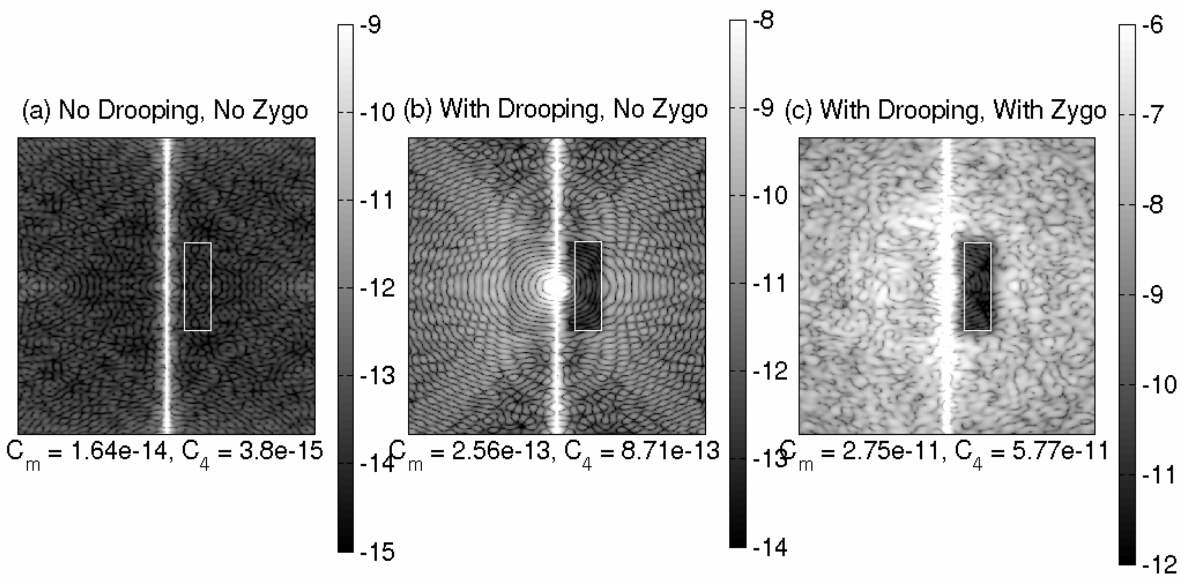

Figure 7. Log-scale contrast maps under three different conditions. The white-colored rectangular box in each graph indicates the field angle window corresponding to [4 10 -10 10]*λ/D. C_m shown as the x-label is the mean contrast inside the [4 10 -10 10]*λ/D box, and C_4 is the contrast at 4λ/D obtained by averaging the contrast values inside the [4 5 -0.5 0.5]*λ/D box. (a) No source-light amplitude drooping and no phase errors for all optics. (b) The source-light has amplitude drooping as shown in Fig. 5, but the optics have no phase errors. (c) Both the source-light amplitude drooping and the optics' phase errors are included.

when the occulting mask phase is present, the phase errors of the optics do not have much impact on the behavior of the HCiT. On the contrary, we found in most of the cases we have examined that the phase errors of the optics actually helpimprove the contrast faster during the early iterations of speckle-nulling, as seen from Fig. 8. This is one of the most significant results of this study. A comparison of Fig. 8 with Fig. 6 shows that either one of the occulting mask parasitic phase and the phase errors of the HCIT optics result in a great degradation in the HCIT's contrast performance; and among these two factors, the phase errors of the occulting mask have stronger adverse effects. That is, it seems that the occulting mask parasitic phase error is the main limiting factor in the HCIT's present contrast performance, not the phase errors of the HCIT optics. This is another significant finding of this study.

In order to understand the impact of the fabrication error in the occulting mask OD profile on the HCIT's performance, we put the measured OD profile and its corresponding phase data into our simulation code, and performed a new set of speckle-nulling simulations. The results of these simulations are also included in Figs. 8-9. Figure 10 shows the correcting phase maps used to obtain the contrast results presented in Fig. 9. The blown-up versions of the OD and phase profiles of the target and the measured occulting masks are shown in Fig. 11. To our surprise, the measured (and "distorted") OD profile worked much better as compared to the target OD profile. Our initial suspicion about the cause of this improved performance was on the difference of the occulting mask width parameter *w*: The targeted width was *w*=141.9μm, but the measured OD profile was best fit with *w*=144.0μm. In order to confirm our suspicion, we conducted another 1000 iterations of speckle-nulling simulations for a target occulting mask with *w*=144.0μm, but we obtained results very similar to those obtained for the target occulting mask with *w*=141.9μm in Figs. 8. Therefore, this set of results is not shown in Fig. 8. That is, it turned out that the fabrication error in the occulting mask width *w* is not responsible for the great improvement in the HCIT's contrast performance. Instead, it seems that the damped (or smoothed) ripples near the center of the occulting mask OD profile (Figure 10) result in better HCIT contrast performance, as is seen from Fig. 8. We are currently examining this phenomenon in more detail.

As we mentioned earlier, the typical best results obtained experimentally on the HCIT during a one-hour time period are C_m~6x10^{-10} and C_4~8x10^{-10}, respectively. The best results of our simulations conducted with the measured occulting mask OD profile are C_m=1.58x10^{-11} and C_4=5.11x10^{-11}, respectively. That is, the modeling results are better by more than an order of magnitude as compared the measured ones. We think this discrepancy in the contrast values between the HCIT experiment and our current simulations are caused by some unaccounted effects in the optical system. These effects may include, but are not limited to, the background scattered light on the testbed, uncertainty in the illumination beam amplitude non-uniformity, uncertainty on the phase errors of the actually used occulting mask, amplitude errors introduced by non-uniform reflectance of HCIT optics, polarization effects, detector noise, noise in DM electronics, variations in DM actuator gains and influence functions, and back-reflections of light from the occulting mask glass and from the CCD camera screen. Some of these effects are currently being investigated in more detail.

4. CONCLUSION

We have conducted a detailed investigation on the monochromatic wavelength behavior of JPL's high-contrast imaging testbed (HCIT) through optical modeling and simulations. By examining both the ideal and realistic cases of occulting masks as well as the HCIT's other optical components, we have found the following about the contrast performance of the HCIT: 1) The noise floor of the HCIT's present MACOS-based optical model in terms of the mean contrast is C_m~1x10^{-14}. 2) Under the realistic amplitude drooping condition of the illuminating light source, an ideal or band-limited occulting mask and error-free optics result in C_m=3.71x10^{-13} (mean contrast) and C_4=1.88x10^{-12} (contrast at 4λ/D), respectively. However, when the phase errors of the HCIT optics are included, these results degrade to C_m=2.75x10^{-11} and C_4=5.77x10^{-11}, respectively. 3) When a realistic occulting mask having parasitic phase-shift errors (or non-band-limited occulting mask) is used, the phase errors of the HCIT optics do not make much difference at the end of a half dark-hole speckle-nulling process. Both of these cases resulted in comparable results of C_m~9x10^{-11}at the end of 1000 speckle-nulling iterations. That is, when the occulting mask has a parasitic phase-shift error, completely eliminating the phase-errors of the HCIT's optics does not yield any advantage in the half dark-hole. 4) The linear-sinc^2 occulting masks fabricated at JPL have damped ripples near the center of their OD profiles as compared to targets. But one such occulting mask investigated in this paper gave much better mean contrast value as compared to the designed occulting mask, namely, C_m=1.58x10^{-11} (with OD ripple-damping) versus C_m≈9x10^{-11} (no OD ripple-damping). We are currently investigating this phenomenon in more detail and will report the results in a future publication.

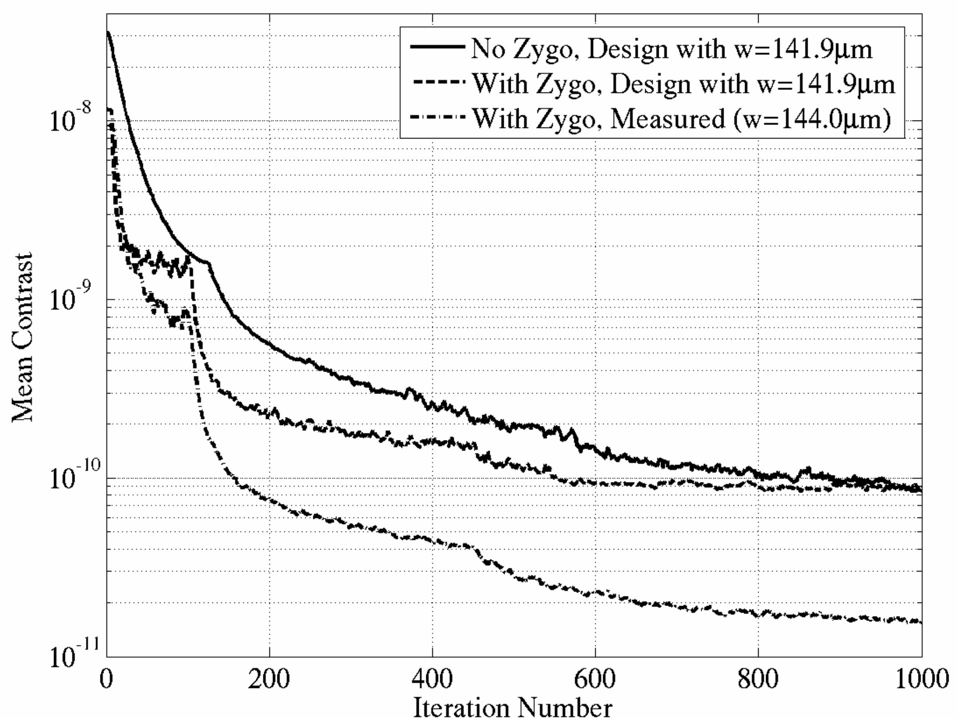

Figure 8. Mean Contrast during the speckle-nullling process. The fast improvement in Cm near iteration=100 takes place when the speckle-nulling window was reduced from the full right-half plane to [3 11 -11 11]*λ/D. The words "Design" and "Measured" in the figure legend indicate that either the designed (target) or the measured occulting mask profiles were used in the simulations.

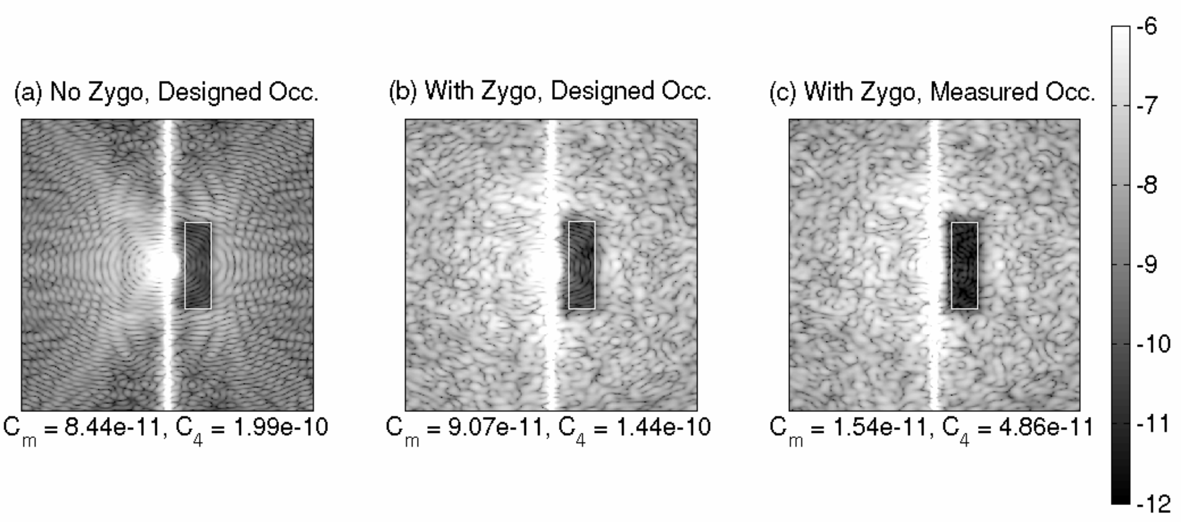

Figure 9. Log-scale contrast maps under three different conditions. The white-colored rectangular box in each graph indicates the field angle window corresponding to [4 10 -10 10]*λ/D. C_m shown as the x-label is the mean contrast inside the [4 10 -10 10]*λ/D box, and C_4 is the contrast at 4λ/D obtained by averaging the contrast values inside the [4 5 -0.5 0.5]*λ/D box. (a) No phase errors for all optics and with target occulting mask. (b) With phase errors for all optics and with target occulting mask. (c) With phase errors for all optics and with measured occulting mask.

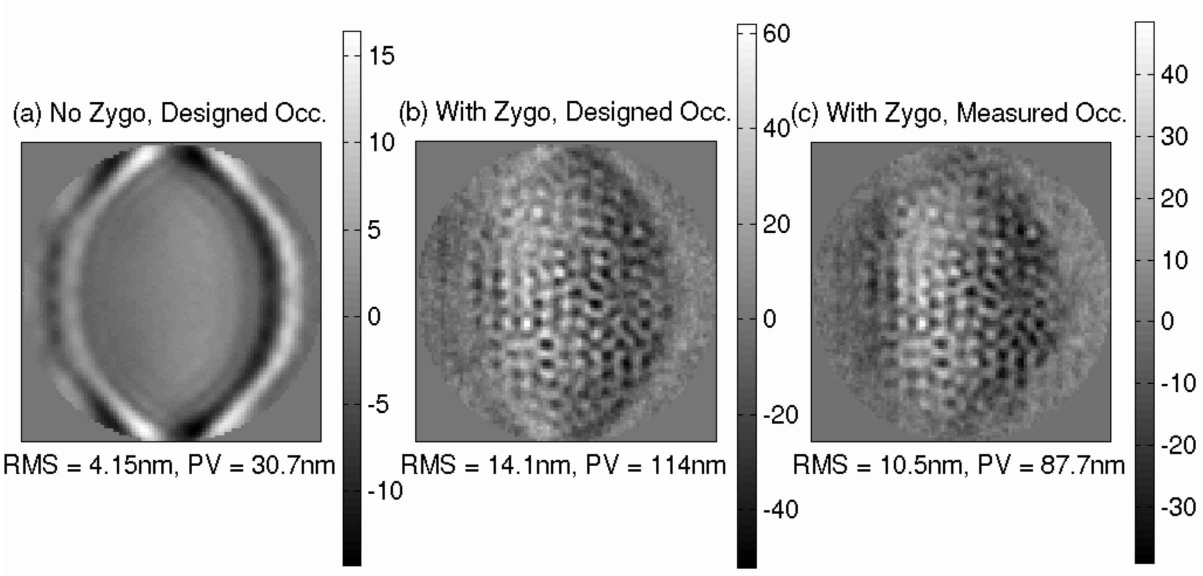

Figure 10. Corrective phase maps (or DM corrective surface patterns) that yield the contrast results given in Fig. 9. The unit of the colorbar is nm in all plots. Shown as the x-label are the RMS and peak-to-valley (PV) values of the corresponding phase map.

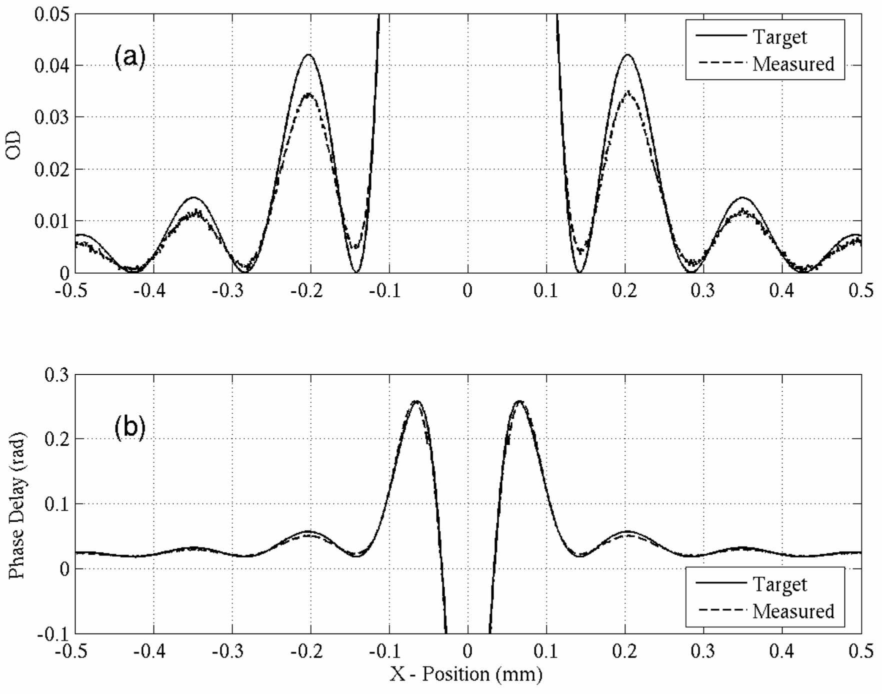

Figure 11. Exploded versions of Figures 2(a) and 2(b), respectively.

ACKNOWLEDGEMENTS

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

REFERENCES

1. John T. Trauger, Chris Burrows, Brian Gordon, Joseph J. Green, Andrew E. Lowman, Dwight Moody, Albert F. Niessner, Fang Shi, and Daniel Wilson, "Coronagraph contrast demonstrations with the high-contrast imaging testbed," Proc. SPIE, **5487**, 1330, 2004.
2. Andrew E. Lowman, John T. Trauger, Brian Gordon, Joseph J. Green, Dwight Moody, Albert F. Niessner, and Fang Shi, "High-contrast imaging testbed for the Terrestrial Planet Finder coronagraph," Proc. SPIE, **5487**, 1246, 2004.
3. Scott A. Basinger, and David C. Redding, "Terrestrial Planet Finder coronagraph optical modeling," Proc. SPIE, **5528**, 1, 2004.
4. *Modeling and Analysis for Controlled Optical Systems User's Manual*, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA.
5. Daniel W. Wilson, Paul D. Maker, John T. Trauger, and Tony B. Hull, "Eclipse apodization: realization of occulting spots and Lyot masks," Proc. SPIE, **4860**, 361, 2003.
6. M. Kuchner and W. Traub, "A coronagraph with a band-limited mask for finding terrestrial planets," *Astrophys. J.*, **570**, 900, 2002.
7. Peter G. Halverson, Micheal Z. Ftaclas, Kunjithapatham Balasubramanian, Daniel J. Hoppe, and Daniel W. Wilson, "Measurement of wavefront phase delay and optical density inapodized coronagraphic mask materials," Proc. SPIE, **5905**, 59051I, 2005.
8. Joseph J. Green and Stuart B. Shaklan, "Optimizing coronagraph designs to minimize their contrast sensitivity to low-order optical aberrations," Proc. SPIE, **5170**, 25, 2003.
9. Fang Shi, Andrew E. Lowman, Dwight C. Moody, Albert F. Niessner, and John T. Trauger, "Wavefront amplitude variation of TPF's high-contrast imaging testbed: modeling and experiment," Proc. SPIE, **5905**, 59051L, 2005.