Extraction of extrasolar planet spectra from realistically-simulated, wavefront-corrected, coronagraph fields

John E. Krist, Stuart A. Shaklan, Marie B. Levine
Jet Propulsion Laboratory
4800 Oak Grove Drive, Pasadena, CA, USA 91109

ABSTRACT

The observation of an extrasolar planet in reflected light using a coronagraphic telescope requires a several-orders-of-magnitude reduction of scattered light around a star by controlling instrumental wavefront errors with deformable mirrors (DMs). This creates a "dark hole" in the background in which diminished speckles of scattered light are at a level similar to or less than that of a planet. There are a number of methods for detecting the planet in this dark field, including image subtraction from another telescope orientation. However, extracting the spectrum of the planet from such a field may be more difficult because the chromatic variation of the speckles can hide spectral features, and the need to disperse the light on a detector increases noise. To investigate the conditions in which spectral extractions are feasible, a physical optics propagation simulation was created of a band-limited Lyot coronagraphic telescope with realistic phase and amplitude errors on each optic and dual sequential DMs to correct the resulting wavefront errors. The Electric Field Conjugation method was used to set the DMs to create a broadband (25%) dark hole from λ = 625 - 878 nm. Simulated planet spectra were then extracted from within this hole using a variety of methods including roll subtraction and spectral filtering. The results of these experiments are presented.

Keywords: Coronagraph, extrasolar planets, Terrestrial Planet Finder

1. INTRODUCTION

The search for planets around stars other than our Sun is a popular topic in modern astronomy. Over 200 planets, most larger than Jupiter, have been indirectly detected by variations in the integrated light of the star+planet systems, either due to radial velocity, transit, or microlensing effects. However, there has yet been no bona-fide observation of light from a planet spatially resolved from its star. This is due to the observational difficulties caused by the large planet-to-star brightness contrasts (10^{-9} for Jupiter and 10^{-10} for Earth in visible light) and small apparent separations between the stars and planets (0.25 arcseconds for an Earth-sized orbit at maximum elongation around a star at 4 parsecs).

Diffraction and scattering from even the best telescope optics create a halo of light around the image of a star against which it is difficult to find a planet. Technology has only recently progressed to the point that it is feasible to significantly reduce these instrumental effects. Coronagraphs can be used to suppress the diffraction pattern of the star created by the edges of the optics or obscurations (*e.g*. the edge of the primary mirror), and many different configurations have been developed. Deformable mirrors (DMs) can correct wavefront aberrations that scatter light caused by optical imperfections, such as figuring, polishing, and coating errors. Ground-based telescopes are now using coronagraphs and DMs to search for the thermal emissions from large, young planets at infrared wavelengths. Observations of large planets in visible light and of any terrestrial (Earth-like) planets, however, require space telescopes that are free from the distorting effects of the atmosphere and large thermal and mechanical variations that can destabilize optics. It is possible to use a combination of a coronagraph and DMs to create a deep "dark hole" around the image of a star[1]. Inside the hole the instrumental background is reduced by several orders of magnitude, to a level equal or similar to that of a planet. There has been considerable progress in the development of algorithms for determining the optimal settings of the DMs that create these holes[2], and laboratory experiments have demonstrated them in space-like environments[3]. With just a coronagraph but no DMs to create a dark hole, the Poisson noise in the background would completely swamp the signal from a planet.

Space Telescopes and Instrumentation 2008: Optical, Infrared, and Millimeter,
edited by Jacobus M. Oschmann, Jr., Mattheus W. M. de Graauw, Howard A. MacEwen,
Proc. of SPIE Vol. 7010, 701044, (2008) · 0277-786X/08/$18 · doi: 10.1117/12.789788

To further improve contrast, image processing is used to reduce the remaining dark hole background. The background can be subtracted using a similar observation of another star or the same star taken at a different orientation of the telescope (*roll subtraction*). The sky, including the planet, will appear to change position on the detector while the instrumental pattern (the dark hole background) remains fixed. This works if the instrumental background, and thus the optics, remains stable. The point source image by itself provides limited information on the planet, so a spectrum is needed to more fully characterize it. Extracting a useful spectrum from the dark hole background is more difficult than just extracting a point source image. The contrast of a significant terrestrial-planet, visible-light spectral line relative to the unsuppressed stellar flux might only be 2×10^{-11}. The background will also vary with wavelength and will do so differently over the field. The flux must be dispersed as well, reducing the signal-to-noise ratio per pixel compared to an imaging observation, where most of the signal over a broad bandpass is concentrated in a compact core. Thus, spectral characterization requires much longer exposure times than detection.

To understand the limits of extracting planet spectra from observations, it is necessary to realistically model an observation, including the effects of optical errors, wavefront control, spectral dispersion, and limited exposure time. Presented here are simulations of observations of Earth as obtained with a hypothetical 4 meter space telescope that used a band-limited coronagraph with deformable mirrors over a wavelength range of $\lambda = 683 - 878$ nm. The spectrum of the planet was measured with an integral field spectrograph and extracted from the background's using roll subtraction and filtering techniques.

2. OPTICAL SYSTEM MODEL

An accurate representation of an extrasolar planet observation requires a realistic model of the telescope, wavefront control system (DM), coronagraph, and detector. At the contrast levels of terrestrial planets, both phase and amplitude wavefront aberrations scatter light, so reasonable estimates of the surface and coating errors on each optic must be assumed. Also, a realistic model of the deformable mirror surface with actuator influence functions is needed.

2.1 Wavefront propagation

The high contrast required for extrasolar planet imaging requires precise control of the wavefront to sub-Angstrom levels, and the simulations of such must be of similar accuracy. This requires that the wavefront propagation algorithms properly handle the changes that occur as the beam travels from optic to optic, including the permutation of phase errors into amplitude errors and vice-versa due to the Talbot effect[4]. Simple far-field diffraction (Fraunhoffer) calculations that can only compute the wavefront at pupil or image planes are insufficient for this task, so near-field diffraction (Fresnel) algorithms are needed. The simulations presented here were generated using the PROPER[5] optical propagation library for IDL (Interactive Data Language). PROPER has a variety of routines for adding phase and amplitude errors, creating complex apertures, using a deformable mirror, and propagating the wavefront using both near- and far-field algorithms. PROPER is free and publicly available from *www.openchannelsoftware.com*. PROPER has previously been used to model a similar telescope+coronagraph system[6]. The new simulations presented here used more realistic surface and amplitude errors, along with a different wavefront optimization method, and cover a 25% wavelength bandpass (polychromatic images were created by computing many monochromatic ones and then averaging them).

2.2 Telescope optics

The simulations assumed an unobstructed (off-axis) telescope with a circular, 4 meter diameter primary. The simplified layout is shown in Figure 1. Each optic was represented by surfaces with two-dimensional phase and amplitude error maps. Both types of maps were realizations of power spectral density (PSD) curves that specified the amount of error at each spatial frequency. The phase (polishing) errors were set according to the Terrestrial Planet Finder Coronagraph (TPF-C) error budget[4]. They are within the current state-the-art for surface finishing.

Amplitude errors caused by sub-percent reflectivity variations have not been thoroughly characterized for large optics, primarily because there are few applications sensitive to them. We have obtained high precision reflectivity measurements of small witness samples that were placed along a line spanning the radius of a large coating chamber at ITT Corporation. From these a one-dimensional reflectivity PSD curve was derived that was converted to a 2-D PSD[7] suitable for defining the amplitude error maps. The derived reflectivity had a root-mean-square variation of 0.04% over all measurable spatial scales and a mean of 96.9%. The 2-D PSD had a distribution that can be reasonably approximated by an f^{-3} power law, where f is the spatial frequency. The amplitude maps in the simulations assumed this PSD profile. The phase error PSDs also followed the same power law but with a different normalization for each optic. The total

wavefront phase error seen by the occulter was 8.6 nm RMS (because the occulter blocks most of the light from the star, scatter from optical errors after it are not very significant). For comparison, the Hubble Space Telescope, which was not designed for high contrast imaging, has an 18 nm RMS wavefront error[8].

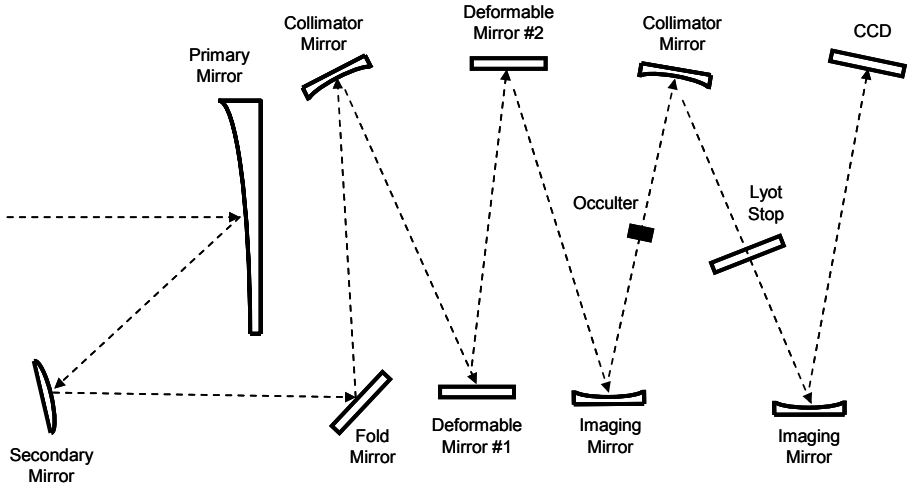

Figure 1. Schematic layout of the simulated optical system in imaging mode (not to scale). An integral field spectrograph would replace the CCD in spectroscopic mode. Deformable mirror #1 and the Lyot stop are located at pupil planes. The occulter and CCD are located at image planes.

2.3 Deformable mirrors

Deformable mirrors are used to correct for wavefront errors and modify the electromagnetic field at the final image plane to produce a dark region around the star so that planets are more easily detected. Two deformable mirrors, one placed at a pupil and one not, can be used to correct for both phase and amplitude errors in a system over a symmetric region centered on the star[4]. The maximum size of a dark hole depends on the number of DM actuators, N, across the pupil; the greater the number of actuators, the larger hole can be (maximum outer radius = $N/2 \times \lambda/D$ radians, where λ is the wavelength and D is the telescope diameter). To allow for greater control of the wavefront and improve contrast, a hole smaller than the maximum allowed diameter is usually created. In the simulations the DM was a square grid of uniformly-spaced actuators with N = 30, and the dark hole outer radius was 13 λ/D radians at λ = 683 nm (0.46 arcseconds). Larger DMs are available, but the computation time required for the greater number of elements during the wavefront reconstruction and correction phase was considered prohibitive for this study. Actuators are modeled in PROPER using measured influence functions[9] from a real DM that describe how the mirror facesheet deforms with actuator pistons.

2.4 Coronagraph

There are a large number of coronagraph designs, most of which remain theoretical. To date, the only one that has approached the contrast levels required for terrestrial extrasolar planet detection in laboratory experiments is the classical Lyot coronagraph[10]. That demonstration used a band-limited[11] occulting mask at an intermediate image plane with a subsequent pupil (Lyot) stop. For the simulations presented here, a circular, band-limited, 4th-order (1-sinc^2 amplitude transmission) occulter was used. In reality, such graded transmission occulters introduce phase errors proportional to both the transmission and wavelength, creating wavefront errors that limit the ability to create a dark hole over a broad bandpass[12]. Recent advances in material selection and the application of dielectric coatings are reducing these effects[13]. Given the rapid rate of progress on this issue, a perfect occulter with no transmission-induced phase errors was assumed for the simulations. The occulter was sized to have 50% intensity transmission at a radius of 3.5 λ/D radians at λ = 878 nm (0.16 arcseconds). The clear area of the corresponding Lyot stop in the subsequent pupil plane was 66% of the diameter of the pupil, reducing throughput by 56%.

2.5 Detector (camera/spectrograph)

After the coronagraph, the image was fed to a integral field spectrograph (IFS)[14]. The image was integrated onto a two-dimensional array of square lenslets that covered the area of the dark hole. The lenslets Nyquist-sampled the field PSF at the shortest wavelength of 683 nm (27 milliarcsec/lenslet). In reality, this division of the field can be done using lenslets or a segmented mirror (*i.e.* an image slicer). In the simulations, the lenslet represented an area over which the image was integrated without any optical propagation being calculated. In an actual IFS, the beam from each lenslet would be fed to a diffraction grating. The simulated dispersed spectrum was integrated onto CCD detector pixels. The sampling of the spectrum was set to be two pixels per spectral resolution unit λ/R. In the final simulations presented here, $R = 70$ and the spectra are 36 pixels long.

At visible wavelengths CCDs are the current detectors of choice. There has been considerable progress in making CCDs that have essentially zero read noise[15], and they will certainly have advanced further by the time a coronagraphic telescope like the one presented here is launch. Therefore, while the simulated spectra include Poisson noise, they do not have read noise.

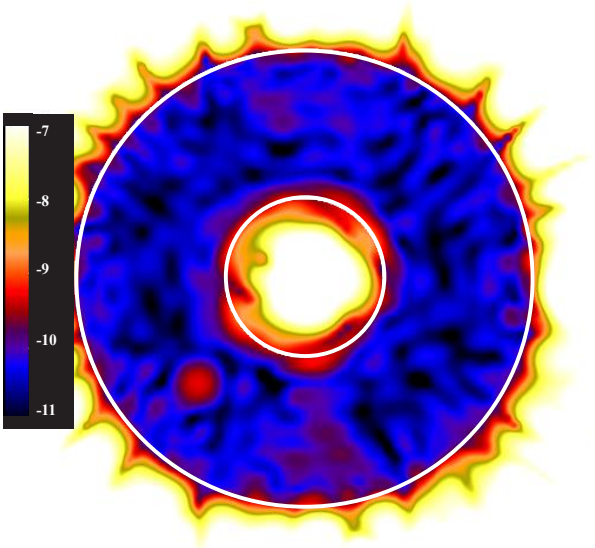

Figure 2. Contrast map of a simulated dark hole around a star over a λ=683 – 878 nm bandpass. The throughput pattern of the occulter has been removed. The color bar indicates $\log_{10}$(contrast). A planet with 5×10^{-10} contrast has been added in the lower left quadrant. The overplotted circles show the inner (3.5λ/D @ λ=878 nm, 0.16 arcsec) and outer (13 λ/D @ λ=683 nm, 0.46 arcsec) radii. Contrast is computed by dividing the flux in a pixel by the peak pixel flux of the star.

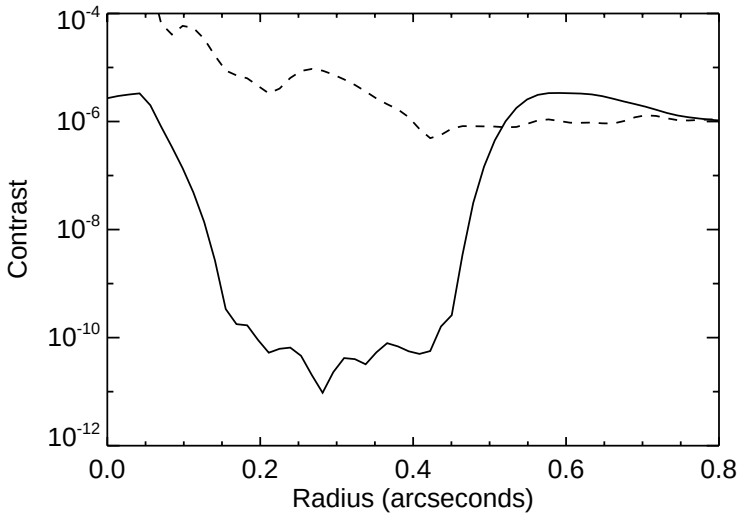

Figure 3. Plot along one radius of the dark hole shown in Figure 2. The dashed line is the contrast before wavefront control and the solid line is after. Note that the background actually increases outside the dark hole after correction.

3. OBSERVATION SIMULATIONS

3.1 Wavefront control and creating dark holes

The previous study[6] used a technique called "speckle nulling" to create the dark hole around the star. That method identifies a number of isolated speckles and alters the deformable mirrors to probe for the type (phase or amplitude) of error that created each speckle and determines the DM patterns necessary to null it. The procedure is very slow, operating on a limited number of speckles at a time. A new and more versatile method[2], electric field conjugation (EFC), has since been developed that offers improved performance in both time and achievable contrast by operating on the entire dark hole field. It makes the approximation that the effect of small ($<<\lambda$) pistons of the DM actuators on the complex-valued electric field (EF) at the final image plane can be described as a linear process. Using assumptions for these functions and the measured EF, a least-squares solution can be obtained for the DM patterns that creates (or at least approaches) a zero-valued dark hole.

In this study, the image plane complex-valued response function for each actuator was numerically computed using a PROPER model of the system that included a representation of the DM actuator surface influence function. Actuators were individually pistoned and the resulting DM surface was propagated through an aberration-free version of the system. All of the response functions were combined into a two-dimensional array, and regularization was used to damp unstable solutions. This matrix was then inverted using singular value decomposition, and the DM actuator settings, DM_{dark}, that null the measured EF within the dark hole were computed.

The EF can be measured via probing using the DM[2], though for simplicity and speed the field here was simply obtained from the computed propagation through the system. To create a dark hole over the given bandpass, the EF and the actuator image plane response functions were computed at multiple wavelengths and included in the EF and DM response matrices. In this study they were computed at five monochromatic wavelengths that evenly sampled the λ = 683 – 878 nm bandpass. In a real system, the sensing would likely use wider bandpass (>5%) filters and would have to deal with image noise.

Because the assumption of linearity is not completely accurate, the EFC process requires some iteration to converge to a good solution. After 25 iterations, the dark hole median contrast in the simulations was 7×10^{-11} (Figures 2 & 3). In a real system, this is about the lowest contrast level one could expect given the effects of noise and exposure time limits during the wavefront correction process. The simulated observations presented here assumed that the dark hole was created by running the EFC procedure using coronagraphic images of a very bright calibration star. The majority of the extrasolar planet candidate stars are fainter than V=4, so it would take too much time to create the dark hole using them.

To allow investigation of the background chromatic variations, a sequence of monochromatic dark hole images was generated using the same DM_{dark} settings at 70 wavelengths that evenly spanned the bandpass. These formed a dark hole datacube, DH_0, with dimensions [x, y, λ]. To simulate how changes in the wavefront can alter the background, separate dark hole datacubes, DH_1 & DH_2, were created by adding more phase aberrations on the primary mirror without any correction by the DMs. These took the form of an additional phase map created with the same PSD specification as the primary mirror but with a low level of error (same phase map with DH_1=0.1 Å and DH_2=0.03 Å RMS). Because of the slope in the PSD, these errors are primarily low-spatial-frequency aberrations.

The plots in Figure 4 (shown before processing by the IFS) demonstrate that the background spectra are dominated by low-order variations of 50-100% of the mean, with lower-amplitude, higher-frequency modulations. The variations are not consistent with wavelength – the contrast is worse at long wavelengths for some regions and better for others. The scattered light near the inner edge of the dark hole is more difficult to control than the other regions, and it has a contrast of $\sim 10^{-9}$ at longer wavelengths. Because the higher-frequency variations are low, the spectra suggest that filtering to extract spectral line features is feasible.

3.2 Noise sources

Each observation comprises a sequence of many separate exposures. The maximum exposure time was about 2000 seconds to prevent too many pixels from being contaminated by cosmic rays. In practice the images would be summed together with rejection of those pixels affected by rays, but none have been added in these simulations.

Given that a zero-read-noise CCD is assumed for the simulations, the major noise source was Poisson noise in the planet's signal and from the background. Additional noise sources (e.g. flat field and dark frame instabilities) were not

considered. The background consisted of the remaining starlight scattered by the optics and the sky (Zodiacal light). Light from a zodiacal-like dust disk surrounding a star may be significant as well but was not considered (the density, dust distribution, grain properties, and inclination factor into the apparent brightness of the disk). Many days of total integration time are needed to achieve a sufficient signal-to-noise ratio for terrestrial planets, meaning that one target would require hundreds of similar, individual exposures.

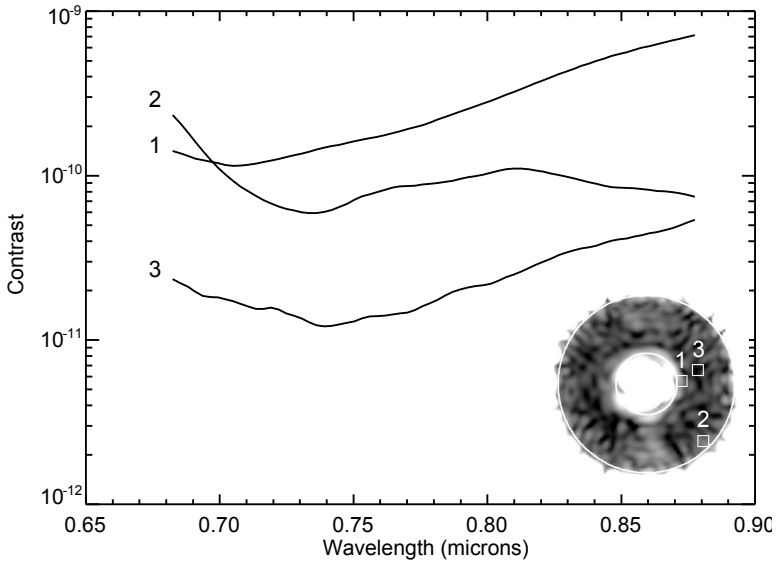

Figure 4. Plots of the mean contrast versus wavelength within the noted selected regions of the dark hole (effective spectral resolution R = 245 – 315).

3.3 Integral field spectrograph

The wavelength range of λ = 683 – 878 nm was chosen because it contains relatively strong O_2 (λ=0.76 μm) and H_2O (λ=0.73 μm & 0.82 μm) absorption lines in Earth's spectrum that can be identified at a resolution of R = 70. At lower resolutions our experiments have shown that these lines cannot be reliably distinguished from the background variations. More prominent lines exist at longer wavelengths, but those would require a larger telescope to obtain the resolution necessary to separate the image of the planet from the star. A well-sampled synthetic spectrum of the modern Earth[16], including clouds, was filtered to a R = 70 spectral resolution. The dark hole datacubes were weighted by a G2V spectrum and also smoothed to R = 70. The fluxes were computed using realistic reflectivities and a mean CCD efficiency. To include planets, PSFs appropriate for each wavelength bin and with the same sampling as the lenslets were added to the datacube at certain *x,y* locations, normalized to the planet's spectrum.

4. EXTRACTING PLANET SPECTRA

The dark hole in our simulations had a median per-pixel surface brightness similar to that of the peak pixel of an Earth-twin planet's PSF. Because of spatial and spectral variations in the background, it is not possible to robustly identify the planet either as a point source or by its spectrum at this contrast level. The background variations must be removed. Two methods have been investigated here: roll subtraction and filtering the spectra.

4.1 Roll subtraction

Roll subtraction is the most straightforward method for removing the background signal. As the telescope is rolled about the line of sight to the star, the image of the sky will appear to rotate on the detector, but the instrumental artifacts (dark hole background speckles) will remain stationary. Subtracting two images (or in this case, two datacubes) taken at different orientations removes the instrumental background and leaves positive and negative images of the sky (planets) located at different positions on the detector. Using an image of the same star rather than another as the reference avoids changes to the background caused by stellar color differences. Such mismatches would create subtraction artifacts. Knowing the roll angles used, more elaborate algorithms can solve for the static and dynamic components using images from two or more orientations[10].

The subtraction quality is critically dependent on the wavefront stability over time and roll angle. A spectral datacube from one orientation represents days of exposure time, and at least two such datacubes are required at different orientations. If a spectral line with a contrast of 2×10^{-11} relative to the star is to be identified, then the background contrast over the combined duration of the two observations can vary by only a fraction of that (it is less stringent for simple point source detection where the contrast is 10^{-10} over the full bandpass). The wavefront must thus remain constant to within a fraction of an Angstrom over weeks and a range of roll angles. Wavefront variations may be induced by thermal and mechanical settling after a slew to the new target or rotation about the star. It is assumed that the telescope is stationed in an L2 orbit, away from the heating from the Earth that can create wavefront variations, such as the orbital focus changes that occur in the Hubble Space Telescope[17]. Continual wavefront monitoring and control may not be practical for maintaining the wavefront to the required stability. Determining the wavefront variations that could be expected from thermal and mechanical effects in our model telescope is beyond the scope of this study. However, the sensitivity of spectral extraction quality to an assumed level of wavefront variation can be investigated.

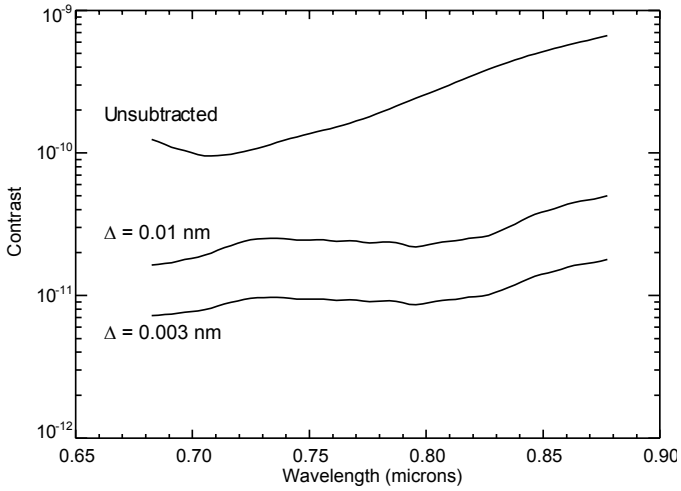

Figure 5. Contrast relative to the stellar flux of background region #1 (see Figure 4) before and after subtraction of a reference image. The reference is a similar observation in which the RMS wavefront phase has changed by the indicated amount without additional correction by the DMs. This approximates a change in the optical system between two observations taken at different roll orientations of the telescope.

To simulate roll subtraction, DH_1 and DH_2, representing the dark hole after rolling the telescope, were subtracted from DH_0, representing the pre-roll dark hole. Figure 5 plots the background spectra within region #1 of Figure 4 (near the inner hole radius) before and after the subtractions without any spectral downsampling. The plots show that the background level was reduced by factors of 4 – 14 using a reference with an 0.1 Å RMS wavefront difference and 10 – 35 with 0.03 Å. The contrast variation versus wavelength decreased from 6×10^{-10} before subtraction to 3.5×10^{-11} (0.1 Å) and 1.3×10^{-11} (0.03 Å) afterwards. The largest variations were of low order, and those with scales similar to the widths of spectral lines are at the 10^{-12} contrast level.

It is assumed that the rotation angle must be sufficient to separate the two planet images by at least the PSF's Airy disk diameter to ensure that the planet does not significantly subtract out its own signal. At the longest wavelength the disk diameter for our coronagraphic system is 0.16 arcsec (note that the effective diameter of the telescope is reduced by the coronagraph's Lyot stop, broadening the planet's PSF). A roll of 60º would be required at the inner dark hole radius of 0.16". It is unlikely that such a large roll angle would be practical in reality as the resulting thermal and mechanical changes may well introduce intolerable wavefront instabilities, and it may not be possible to position the solar arrays properly at both angles during the same visit. The roll angle would be a more feasible 10º at 0.9" from the star, but that is twice the dark hole radius in these simulations. Alternative methods may be needed to reduce the background because of this limitation.

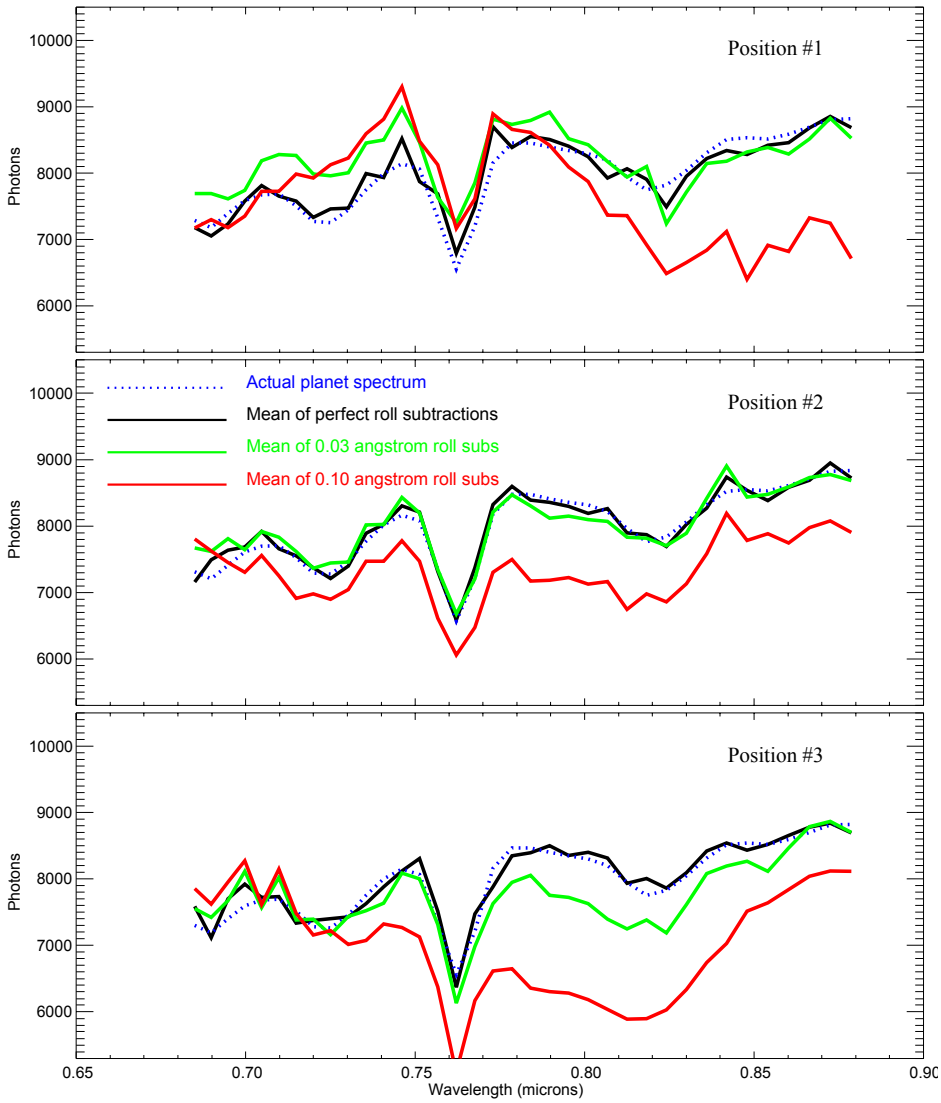

Figure 6. Spectra of Earth around a V=4 star extracted from the simulated dark hole using roll subtraction to remove the background signal. The total integration time is 20 days per roll. The spectra represent the flux per spectral bin integrated over 3×3 lenslet arrays centered at the locations shown in Figure 4. The dashed blue line is the noiseless model spectrum of the Earth. The black line is the mean spectrum using perfect roll subtraction (no wavelength change between rolls). The green line is the mean with a 0.1 Å RMS wavefront change between rolls, and the red line is the mean with a 0.03 Å RMS change.

IFS datacubes were generated for a G2V star of magnitude V=4, which is at the bright end of the TPF sample of nearby stars[18]. Simulated Earth PSFs, with the appropriate spectral weighting, were added to the optimized (pre-roll) datacube at the locations of the three regions shown in Figure 4 (the PSFs had the same contrast, 10^{-10}, regardless of location). The same PSFs were added to the post-roll datacubes at locations offset 60º in azimuth about the star from those three positions. Poisson noise was added and the post-roll datacubes were subtracted from the pre-roll one. The total signal within a 3×3 lenslet region centered on each planet at each wavelength was computed to generate the spectrum. At the post-roll locations where the planet signals were negative in the subtractions, the spectra were negated and then averaged with the pre-roll planet spectra.

The results are shown in Figure 6 for observations of 20 days (864 exposures) per orientation. The simulations with perfect subtractions (the background did not change except for different realizations of Poisson noise) reliably revealed all three major lines present in the planet spectrum and the larger-scale variation of the overall flux. With a 0.03 Å RMS wavefront difference between rolls, all the lines could be identified, though the larger-scale variations were not reliably seen at all positions within the dark hole. With the larger 0.1 Å RMS difference, only the O_2 line at λ=0.76 μm was identifiable at all positions – the large scale variations were lost and the weaker lines could not be reliably identified among the other small-scale variations.

Table 1. Roll subtraction measured goodness-of-fit and O_2 (λ=0.76 μm) spectral line fitting errors

(statistics for combined observations taken at two roll angles)

	Full Spectrum χ^2	O_2 Line Fit Error Relative to Model		
		Depth (% Error)	Center λ (% Error)	Width (% Error)
Position #1				
5 days/roll				
Perfect	0.83	2.5	0.0	12.1
0.03 Å	1.06	0.0	0.1	15.3
0.10 Å	4.30	2.2	0.1	28.9
10 days/roll				
Perfect	0.40	0.6	0.0	8.7
0.03 Å	0.61	1.7	0.1	2.6
0.10 Å	3.97	5.2	0.1	5.6
20 days/roll				
Perfect	0.18	1.7	0.0	1.3
0.03 Å	0.56	4.2	0.1	11.3
0.10 Å	3.98	3.3	0.1	4.6
Position #2				
5 days/roll				
Perfect	0.29	0.6	0.0	20.2
0.03 Å	0.39	7.4	0.0	2.4
0.10 Å	1.08	6.9	0.1	10.3
10 days/roll				
Perfect	0.14	4.8	0.0	10.8
0.03 Å	0.20	1.2	0.0	6.3
0.10 Å	0.86	3.4	0.1	5.6
20 days/roll				
Perfect	0.07	0.6	0.0	4.9
0.03 Å	0.15	1.4	0.0	4.7
0.10 Å	0.89	2.7	0.1	0.0

To quantify the quality of the extracted spectra, the goodness-of-fit was characterized by computing a simple chi-square* (χ^2). In addition, Gaussian profiles were fit to the 0.76 μm O_2 line and compared to the fit to the line in the input spectrum. At the signal levels of the simulations, the variations in a given spectral pixel due to Poisson noise can be high enough to obscure any general trends in quality versus exposure time. To reduce this effect on the statistics, 20 separate

* $\chi^2 = \sum_{\lambda} (y_{\exp} - y_{obs})^2 / y_{\exp}$, where y_{exp} and y_{obs} are the expected and observed spectral values, respectively

subtractions were generated for each position, each with separate realizations of Poisson noise, and the median of the χ^2 values for each was computed. To allow for comparison between exposure times, the spectra used for computing χ^2 were divided by the number of 2000 sec exposures that comprised them. Also, in order to compare the results with those from filtering presented in the next section, the offsets between the extracted and model spectra were adjusted to minimize χ^2. The line depths, center wavelength, and full widths at half depth were obtained. The results are shown in Table 1 for positions 1 and 2 identified in Figure 4 and for different integration times.

The measurements show a significant reduction in spectrum quality when the wavefront difference between rolls is greater than 0.03 Å RMS. With a 0.1 Å RMS difference, the extracted spectrum quality is limited by background wavelength variations rather than detector noise, so additional exposure time does not provide much gain.

4.2 Spectral filtering

As discussed in the prior section, roll subtraction may not be practical for observing planets at small angles from the star due to the large rotations required to change the PSF position enough to prevent self-subtraction. The technique also requires great stability between rolls and over the combined duration of the observations. Ideally, the planet's spectrum would be differentiated from the background's without resorting to rolling the telescope. One way would be to search for spectral features that differ between the planet and the background. If the background spectrum changes in a manner similar to that of the planet's, then this would not be possible, but the simulations show that the background varies slowly.

The differences in the scale of the variations suggest that high pass filtering can be used to remove the slowly varying background. The simplest way to do this is to subtract a smoothed version of the spectrum from itself. The choice of the shape and width of the smoothing kernel can affect the results. We have experimented with Gaussian and top-hat smoothing kernels with widths ½ to 6× that of the O_2 λ=0.76 μm line width. As with the roll subtractions, the χ^2 was computed and the O_2 line fitted (Table 2). Because filtering out the low-order components of the spectrum subtracts the overall offset, the χ^2 was chosen that minimizes the difference between the extracted and model spectra. The best results were obtained with a top-hat kernel with a width equal to 7 spectral pixels (R=λ/3.5).

The extracted spectra are shown in Figure 7 for a 20 day observation of an Earth around a V=4 star. The O_2 line can be reliably identified, as can the 0.73 μm H_2O line. The other H_2O line is not identifiable in the position #1 spectrum (which has the higher background level), but is in the position #2 one, though there is noise that mimics lines of similar depth.

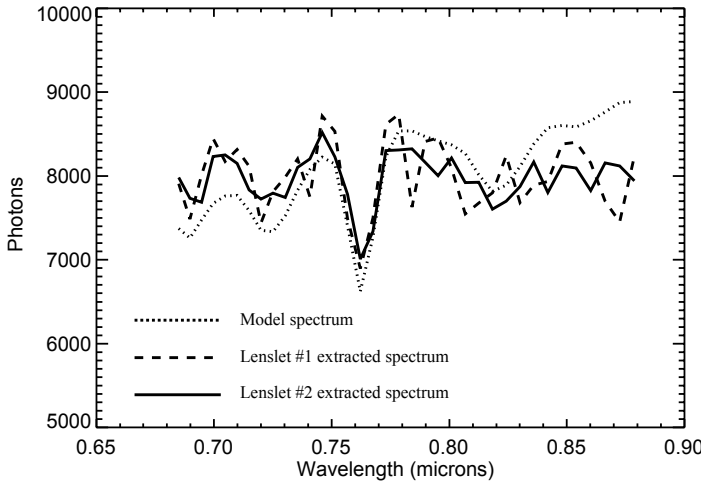

Figure 7. Spectrum of the Earth extracted using filtering from a simulated observation of a V=4 star with 20 days of total exposure time and with Poisson noise and background speckle variations. The extracted spectra are shown for two locations identified in Figure 4.

5. CONCLUSIONS

The results presented here represent the most complete end-to-end simulations of coronagraphic terrestrial exoplanet characterization observations to date. They included a model of a 4 meter diameter space telescope having realistic phase and amplitude wavefront aberrations, used deformable mirrors with an optimization algorithm to alter the wavefront to

create a dark hole over a λ=683-878 nm bandpass, and applied different methods to distinguish Earth spectra from noisy background variations as measured with an integral field spectrograph with a spectral resolution of $R = 70$.

Roll subtraction and high-pass filtering were used to extract the planetary spectra. Comparisons of the fit statistics for the two methods show that for equal total integration times roll subtraction will provide better results if temporal wavefront variations are below 0.1 Å RMS for the assumed error distribution. The quality of the extracted spectrum rapidly degraded between 0.03 Å and 0.1 Å RMS of change. Both methods produced similar λ=0.76 μm O_2 line fit statistics for the same integration times. However, the weaker H_2O lines were more apparent in the roll subtractions.

Roll subtraction appears to be necessary to reliably measure spectral features weaker than the O_2 line and to detect low-order variations in the spectrum. The cost, however, is a stringent wavefront stability requirement so that the change in the background is kept to less than 10^{-11} in terms of contrast relative to the star (less for simple detection of the planet in a broadband image). The need to roll the telescope enough to move the planet fully out of the background region may place unrealizable constraints on the optical stability.

Spectral filtering offers hope that for measuring the most prominent spectral lines only observations from one orientation of the telescope are needed. This significantly reduces the stability requirements. Filtering, however, can remove low-order spectral variations that may be important.

Table 2. Spectrum filtering measured goodness-of-fit and O_2 (λ=0.76 μm) spectral line fitting errors

	Full Spectrum χ^2	O_2 Line Fit Error Relative to Model		
		Depth (% Error)	Center (% Error)	Width (% Error)
Position #1				
10 days	1.68	14.3	0.1	20.8
20 days	1.37	4.7	0.1	15.8
40 days	1.14	2.9	0.1	6.8
Position #2				
10 days	1.14	1.5	0.1	14.9
20 days	1.02	3.9	0.1	5.4
40 days	1.00	1.5	0.1	8.4

ACKNOWLEDGEMENTS

The authors thank Lisa Kaltenegger for providing the simulated Earth spectra. The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

[1] Malbet, F., Yu, J.W., and Shao, M., "High-Dynamic Range Imaging Using a Deformable Mirror for Space Coronagraphy," Pub. of the Astronomical Society of the Pacific, 107, 386 (1995).

[2] Give'on, A., Kern, B., Shaklan, S., Moody, D., and Pueyo, L., "Broadband wavefront correction algorithm for high-contrast imaging systems," Proc. SPIE 6691, 66910A (2007).

[3] Trauger, J., Give'on, A., Gordon, B., Kern, B., Kuhnert, A., Moody, D., Niessner, A., Shi, F., Wilson, D., and Burrows, C., "Laboratory demonstrations of high-contrast imaging for space coronagraphy," Proc. SPIE 6693, 66930X (2007).

[4] Shaklan, S., and 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).

[5] Krist, J., "PROPER: an optical propagation library for IDL," Proc. SPIE 6675, 66750P (2007).
[6] Krist, J., Trauger, J., and Moody, D., "Studying a simple TPF-C," Proc. SPIE 6265, 62653O (2006).
[7] Church, E., and Takacs, P., "Optimal estimation of finish parameters," Proc. SPIE 1530, 71 (1991).
[8] Krist, J., and Burrows, C., "Phase retrieval analysis of pre- and post-repair Hubble Space Telescope images," Applied Optics, 34, 4951 (1995).
[9] Ealey, M., and Trauger, J., "High-density deformable mirrors to enable coronagraph planet detection," Proc. SPIE 5166, 172 (2004).
[10] Trauger, J., and Traub, W., "A laboratory demonstration of the capability to image an Earth-like extrasolar planet," Nature, 446, 771 (2007).
[11] Kuchner, M., and Traub, W., "A Coronagraph with a Band-Limited Mask for Finding Terrestrial Planets," Astrophysical Journal, 570, 900 (2002).
[12] Halverson, P., Ftaclas, M., Balasubramanian, K., Hoppe, D., and Wilson, D., "Measurement of wavefront phase delay and optical density in apodized coronagraphic mask materials," Proc. SPIE 5905, 59051I (2005).
[13] Moody, D., and Trauger, J., "Hybrid Lyot coronagraph masks and wavefront control for improved spectral bandwidth and throughput," Proc. SPIE 6693, 66931I (2007).
[14] Allington-Smith, J., "Basic principles of integral field spectroscopy," New Astronomy Reviews, 50, 244 (2006).
[15] Wen, Y., Rauscher, B., Baker, R., Clampin, M., Fochie, P., Heap, S., Hilton, G., Jorden, P., Linder, D., Mott, B., Pool, P., Waczynski, A., and Woodgate, B., "Individual photon counting using e2v L3 CCDs for low background astronomical spectroscopy," Proc. SPIE 6276, 6276H (2006).
[16] Kaltenegger, L., Traub, W., and Jucks, K., "Spectral Evolution of an Earth-like Planet," Astrophysical Journal, 658, 598 (2007).
[17] Lallo, M., Makidon, R., Casertano, S., and Krist, J., "Temporal optical behavior of HST: focus, coma, and astigmatism history," Proc. SPIE 6270, 62701N (2006).
[18] Brown, R., "Single-Visit Photometric and Obscurational Completeness," Astrophysical Journal, 624, 1010 (2005).