Precise starshade stationkeeping and pointing with a Zernike wavefront sensor

Michael Bottom, Stefan Martin, Carl Seubert, Eric Cady, Shannon Zareh, Stuart Shaklan

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109 USA

ABSTRACT

Starshades, large occulters positioned tens of thousands of kilometers in front of space telescopes, offer one of the few paths to imaging and characterizing Earth-like extrasolar planets. However, for the starshade to generate a sufficiently dark shadow on the telescope, the two must be coaligned to just 1 meter laterally, even at these large separations. The principal challenge to achieving this level of control is in determining the position of the starshade with respect to the space telescope.

In this paper, we present numerical simulations and laboratory results demonstrating that a Zernike wavefront sensor coupled to a WFIRST-type telescope is able to deliver the stationkeeping precision required, by measuring light outside of the science wavelengths. The sensor can determine the starshade lateral position to centimeter level in seconds of open shutter time for stars brighter than eighth magnitude, with a capture range of 10 meters. We discuss the potential for fast (ms) tip/tilt pointing control at the milli-arcsecond level by illuminating the sensor with a laser mounted on the starshade. Finally, we present early laboratory results.

Keywords: Starshades, exoplanets, high contrast imaging, formation flying

1. INTRODUCTION

The two challenges of imaging extrasolar planets around other stars are the problem of brightness and angular separation. A typical star outshines a planet by anywhere from 10^4 to 10^{10}, depending on the wavelength, age, and size; and the separation on sky of a planet around a nearby star is on the order of tens of milliarcseconds, less than the width of a human hair viewed 100 meters away. The only way forward solution is to artificially remove the starlight. One way to proceed is to suppress the starlight inside the telescope using a device called a coronagraph, and requires exquisite wavefront control to correct for optical imperfections and aberrations. Another way is to prevent the starlight from entering the optical system in the first place, by creating an artificial stellar eclipse on the telescope using a free-floating occulter called a starshade.[1,2]

In order to block the star while still seeing the planet, the starshade's angular size (as seen from the telescope) must be on order of the desired inner working angle,[3] which should be smaller than the angular separation of the planet and star. Combined with manufacturing tolerances and the requirement that the shadow cast on the telescope must be larger than the aperture, this typically leads to starshade designs tens of meters in size at separations of tens of thousands of kilometers from the telescope. Furthermore, the starshade must preserve its lateral position with respect to the telescope over the course of an observation to about 1 meter laterally, before starlight leaking around the edge becomes a problem.[4]

This control problem has two parts. First, the starshade's position must be sensed to very high accuracy by the telescope; second, the position must be dynamically corrected. The problem of correction is actually the easier one–the typical gravity gradients experienced by a telescope-starshade combination at Lagrange point 2, for example, are quite similar to those experienced in low-Earth orbit, where centimeter level control is routinely achieved.[5] The problem of sensing the relative position is more challenging. At 50,000 kilometers, motion of 5 cm corresponds to a nanoradian of angular separation.

In this work, we consider the use of a Zernike wavefront sensor as a pupil-plane position and tilt sensor. We perform numerical simulations using a realistic WFIRST-like telescope and detector, gravitational trajectories, and errors to show the amount of time needed to detect the position of the starshade to better than 5 cm.

Send correspondence to mbottom@jpl.nasa.gov

(The WFIRST project has been directed to study compatibility with a starshade, but as of now, NASA has not decided to launch a starshade to accompany WFIRST. This decision is pending the 2020 Decadal Survey.) We show that for stars brighter than eighth magnitude, a Zernike sensor can determine the position to better than 5 cm in one second. Additionally, the capture range is sufficient to "pull in" the starshade from a distance few meters away, avoiding the "dead zone" in position sensing when transitioning from ten meters to one meter away. We also perform calculations showing how the use of a laser beacon on the starshade can be used to separate tip/tilt errors of the telescope from lateral errors, and provide laser power requirements. Finally, we present preliminary laboratory work on starshade guidance.

2. STARSHADE NUMERICAL SIMULATIONS

2.1 Overview

Starshades have separation and wavelength-dependent shadow depths. To get good performance at a different wavelength, the starshade may be moved in the axial direction by thousands of kilometers. This may be understood as a consequence of the starshade being designed for a particular Fresnel number $F \propto D/(\lambda z)$, where D is the size of the starshade, z the separation, and λ the wavelength. While the peak suppression at a given wavelength and starshade/telescope separation may be better than 10^{10}, an octave from this design wavelength, the shadow brightness may be higher by a factor of 10^5 or more. It is this "leaked" light that is particularly useful for sensing the starshade's position.[6] While science observations can proceed at the design wavelength, the out-of-band light may be used in a separate wavefront sensing arm of the optical system to determine the starshade position and determine whether a trajectory correction is necessary to keep it aligned.

For WFIRST, there are three different science filters that are of interest, the "red" science band from 600-1000nm, the "green" science band for wavelengths of 550-800, and the "blue" science band for wavelengths of 400-650 nm. Figure 1 shows the different bands of interest, with starshade performance at shown as a function of wavelength. When performing science in the red band, blue light (separated by a dichroic) are used to sense the starshade position. For blue science, the opposite is true. For green science, either red or blue light may be used.

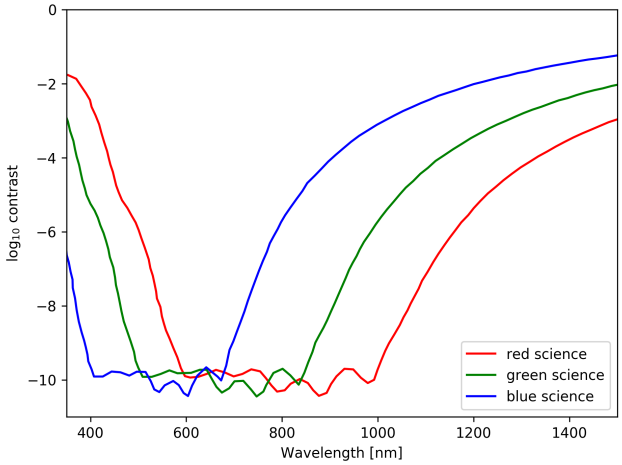

Figure 1: Starshade performance at in the different science bands, as a function of wavelength. Note the rapid decline in performance outside of the specified science band. For example, in the blue science band the suppression (or contrast) is about 10^{-10} at 500 nm but only about 10^{-3} at 900 nm.

The question of starshade position sensing is thus one of radiometry: for a star of magnitude M_V, how long does it take to sense the position to an accuracy of better than 5 cm? To answer this, one needs to know how much out-of-band light is present, and how many photons are needed to sense the starshade position to a certain accuracy. The answer depends on a few factors, such as the spectrum of the target star, the transmission

profile of the starshade, throughput of the telescope optics, wavefront sensing scheme, and detector performance. Additionally, realistic spacecraft dynamics, such as starshade trajectory and telescope vibrations should be considered.

We use the following strategy to determine the answer to this question, which we summarize here and explain further in the following sections. First, we generate 2-d post-starshade maps of the stellar electric field at the position of the telescope. We propagate these electric fields through the telescope optics and onto the wavefront sensor detector plane. We repeat this for a grid of telescope separations, which we combine into a (noise free) image library. To determine the sensing accuracy, we repeat the exercise except include the effects of nonidealities, including Poisson shot noise (which depends on the star's brightness), detector noise, and telescope jitter. Essentially, this gives us a "noisy image," with noise comparable to that expected in reality. We then run a matched filter algorithm between the clean image library and the noisy image, storing the derived position. Performing the experiment with thousands of noisy images at a particular grid point allows us to determine the scatter in recovered positions, and hence the intrinsic sensing accuracy.

A similar approach is used to explore two further avenues of interest. First, we consider the problem of sensing the starshade position at an intermediate offset from the optimal position, up to 10 meters away, as a method of initial alignment. We show that there is sufficient light to guide the starshade into the optimal position from this distance. Second, we consider the problem of correcting for not only telescope position (shear) errors, but tip/tilt jitter. While no stars are bright enough to generate enough tip/tilt signal, we show a modestly powered laser mounted on the starshade can add enough signal to simultaneously allow for tip/tilt and shear sensing.

2.2 Electric field propagation

For this simulation, we use a 34 meter starshade model operated at a design Fresnel number of $F \approx 12$. We propagate a plane wave from the starshade position to the telescope, and plot the electric field on a 20x20m grid of 2 cm/pixel. The electric field is normalized such that the square of the amplitude corresponds to the flux ratio at that position. For example, if the electric field amplitude at the telescope input plane is 0.1 at a particular pixel, the starshade is suppressing the intensity by a factor of 100 as compared to what the electric field would be with no starshade present. In order to efficiently and accurately perform the Fresnel propagation calculations, we use the method introduced in Cady et al. 2012,[7] where Babinet's principle and a boundary value integral are used.

To capture the input electric field, the starshade electric field is multiplied by the telescope input pupil, interpolated to the same pixel scale. To simulate starshade/telescope shear misalignment, the telescope pupil is shifted with respect to the starshade before multiplication. To simulate tip/tilt error, the starshade electric field is multiplied by a phasor $e^{2\pi i\theta/\lambda}$, where θ is the tip/tilt angle and λ is the wavelength.

Once past the telescope pupil, we use Fraunhofer propagation to transform it to the first focal plane of the Zernike wavefront sensor. We simulate a "pure" Zernike wavefront sensor, which is a phase plate that multiplies the electric field by 180 degrees within a 0.5 λ/D radius on-axis area and leaves it unchanged outside that area. The performance of a Zernike sensor depends weakly on this size; tradeoffs between linearity and sensitivity may be had by making the dislocation larger or smaller.[8] After the Zernike sensor, a similar propagation is performed to the plane of the detector, which is at a pupil. At this point, the image is resampled to the expected 16x16 pixels on the detector.

This procedure is repeated for a grid of shear positions at 2 cm spacings, so that the expected light intensity on the detector may be calculated for these many points. The set of positions and images (an [N x 16 x 16] matrix, where N is the number of grid points) is used to form a reference image library.

2.3 Photon budget and detector noise

In reality, while the reference image library describes the intensity pattern of the light on the detector, the actual amount of photons measured will depend on the brightness of the star, integration time per frame of the camera, and efficiency losses through the optics. Also, the number of photons measured per frame will vary due to shot noise; with the number of photons k being measured described by a probability distribution $p(k) = e^{-\lambda}\lambda^k/k!$,

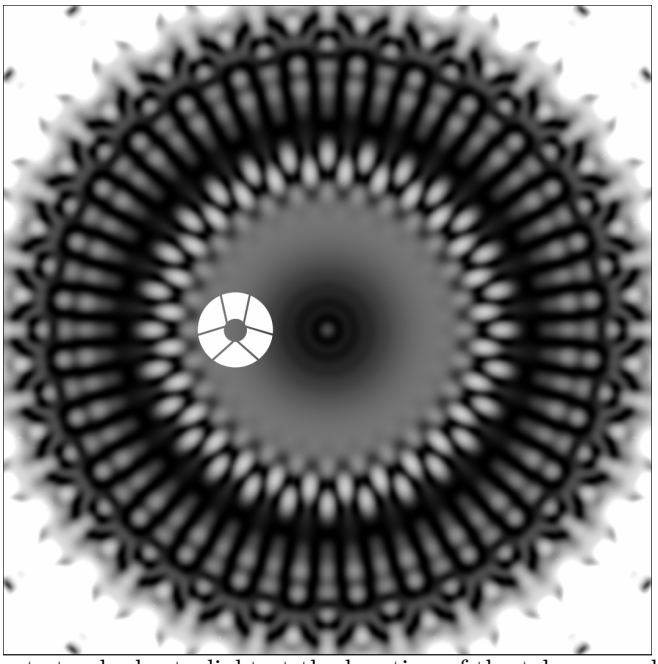

Figure 2: The intensity of post-starshade starlight at the location of the telescope, shown on a nonlinear stretch to show features. The relative size of the WFIRST telescope pupil is overlaid. This image corresponds to a square 20m on a side.

where λ is the mean intensity, that is, the mean number of photons per frame at that particular pixel. On top of this, detectors will add their own noise.

In order to determine the mean intensity, we consider the stellar intensity and efficiency and propagation losses through the optics.

To approximate a stellar spectrum, we use the ATLAS9 synthetic spectral atlas.[9] We use a model of a G-type star of effective temperature 5750K, and solar metallicity and surface gravity. The flux density is converted to photons/m^2/s after accounting for stellar magnitude (or stellar distance and radius) and filter bandpass. The stellar model was checked against two separate sources, the actual solar radiance measured at the top of the atmosphere*, and the zero-point magnitude of different filter systems,[10] showing excellent agreement (at the few percent level) in spectrum and number of photons measured for both cases. The advantage of a synthetic atlas is that different spectral types can be easily investigated without changes to the code architecture. The intensity of this spectrum is scaled by the stellar apparent magnitude via the distance modulus formula.

Table 1: default

Item	Model	Value	Notes
Detector QE	EMCCD39	N/A	Manufacturer QE curve
Silver Mirrors	FSS99-600	N/A	Nineteen mirrors
Lenses	N/A	0.99^4	Four lenses
Dichroic	N/A	0.95	Single dichroic
Red science filter	N/A	486-614 nm, T=100%	Square filter
Blue science filter	N/A	860-1000 nm, T=100%	Square filter

*ASTM G173-03(2012), Standard Tables for Reference Solar Spectral Irradiances: Direct Normal and Hemispherical on 37 Tilted Surface, ASTM International, West Conshohocken, PA, 2012

Inside the telescope, further losses are due to reflections (nineteen), four lenses, and one dichroic needed for the total path to the wavefront sensor detector. Silver reflection data was approximated assuming a FSS99-600 coating[†], and the lenses were assumed to be 99% efficient each, with the dichroic being 95% efficient. Furthermore, the detector quantum efficiency must be incorporated (detector noise is not a radiometric loss, but a separate source of noise), and for this we use the manufacturer response curve. Individual filter transmission losses are not considered, but are assumed to be insignificant compared to the net losses from optical components. Figure 3 shows the associated losses from the pupil to the detector. Table 1 gives information on the individual subcomponents.

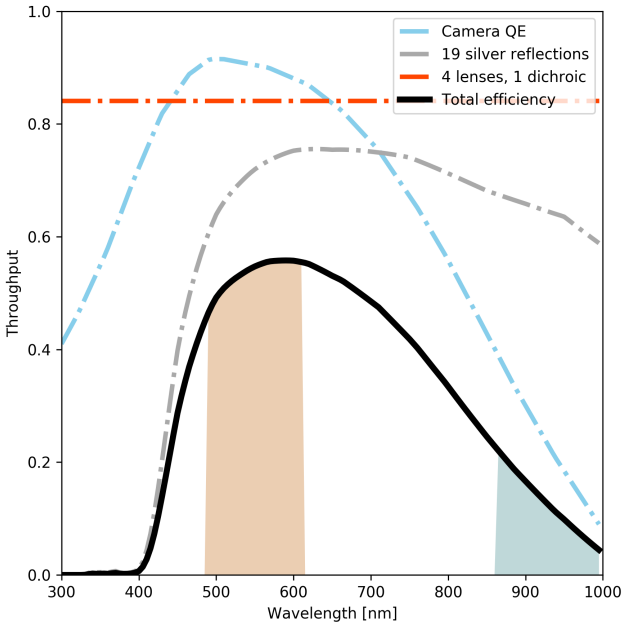

Figure 3: Losses from lenses, mirror reflections, and dichroic, with the associated filter bandpasses shown for red science (500-600 nm wavefront sensing) and blue science (860-1000nm wavefront sensing).

Once the throughput-corrected electric field on the detector is known, this must be converted into a mean intensity of photons/pixel/frame, which is done by multiplying the mean photons/second by a selected exposure time. To convert to a single realistic detector readout, the photons/pixel/frame are modified according to Poisson noise, and Gaussian detector read noise is added at a level of 3 to 5 electrons/pixel/frame. For a particular shear offset between the telescope and starshade, this exercise is performed 1000 times, so that there are 1000 independent realizations of image plane data. For each of these realizations, the image data is used in a matched filter against the template of the image library. In noisy data, the matched filter will sometimes mismatch, and with the 1000 samples, it is possible to build up statistics demonstrating the accuracy and precision of the wavefront sensor by comparing the recovered positions with the true positions. By adjusting the exposure time for different stellar magnitudes, it is possible to determine how much exposure time is needed to limit positional errors to an acceptable level.

Figure 4 shows an example of how the simulation proceeds.[‡] Here, the telescope pupil is decentered from the nominal position by about 0.75 meters. The leftmost panel shows an overlay of the telescope pupil on the out-of-band diffracted intensity. The right panel shows a spatial plot of a trajectory correction maneuver, where the blue circle shows the 1m radius control region. The parabolic arc demarcates the trajectory correction maneuver, where a thruster fire occurs at the bottom right (4 o' clock), causing a ∼11 minute "freefall" trajectory in the control region, after which another thruster fire is required. The upper inset shows a "perfect" image of the

[†]Quantum Coating Incorporated

[‡](Animations of the trajectory simulation may be found at *mb2448.github.io/ss_traj.zip*)

Star magnitude (solar spectrum)	Red science req. (sec) for 1σ uncertainty <5cm	Blue science req. (sec) for 1σ uncertainty <5cm
8	1 (**2 cm**)	6.25
7	0.4	2.5
6	0.16	1 (**3 cm**)
5	0.064	0.4
4	0.0256	0.16
3	0.01	0.064

Table 2: The required exposure time, in seconds, to get a positional uncertainty of less than 5cm using out-of-band starlight to guide the starshade. The left column corresponds to the stellar apparent magnitude, the middle shows the exposure time requirement for the red science band (using blue light to guide) and the right column shows the exposure time requirement for the blue science band (using red light to guide). The bold text shows the exact uncertainty in the numerical simulation result, which is scaled for other magnitudes. Band specifications and radiometry assumptions are shown in Table 1.

out-of-band light on the wavefront sensor 100 seconds into the maneuver; the lower inset shows the same after Poisson and detector read noise are added corresponding to 1 second of exposure time. The noisy image is matched to a library of clean images to determine the position. The lower panel shows the results of 1000 such noise realizations at five different trajectory points, with the error ovals showing the standard deviation of the positional determination, about 2 cm in this case.

We investigated whether telescope jitter would contribute meaningfully to the error budget. It doesn't. The dominant sources of jitter in the system comes from the reaction wheels and the cryogenic cooler, which both contribute to about 15 mas of tip tilt jitter with major modes at 45 and 95 Hz, respectively. To simulate jitter, we performed the same propagations through the system but with an additional injection of a random tip and tilt term to the input wavefront. The tip and tilt terms were selected by a 2d Gaussian random variable with a radial 1-σ of 15 mas. The end result was a modest increase in positional error of about 15% in the brightest data (eg. 2.3 cm RMS vs 2 cm RMS with no jitter), which is negligible with respect to the radiometric error budget. Another way of saying this is jitter is undetectable for fainter stars, and minor for bright stars.

Table 2 shows the results of these numerical simulations.

2.4 Jitter control

Unlike coronagraphs, which require a very high level of wavefront control and are extremely sensitive to tip/tilt jitter, starshades are somewhat less so. Despite this, it may be desirable to control jitter, as it has the effect of blurring out the point-spread function. In the case of spectroscopy, this corresponds to a throughput loss, as less of the light makes it down the slit or lenslet and dispersed into the spectrum of interest. A 15 mas rms jitter term corresponds to a reduction in peak PSF amplitude of 22% and 9% for 500 nm and 850 nm operating wavelengths, respectively. Reducing jitter to 5 mas rms causes just a 4% and 1% reduction, respectively, so a goal of less than 5 mas of jitter is reasonable.

We note here that the previous simulations did not actually benefit from the design of the Zernike wavefront sensor, but could have been performed with any pupil imaging system. However, to control jitter, the tip/tilt must be sensed at high speed. A pupil imaging sensor is blind to jitter and any low-order aberrations; this is where the use of a Zernike sensor is required. The Zernike sensor is the low-order aberration sensor baselined for WFIRST due to its high photon efficiency; that is, its ability to measure aberrations accurately even in low-flux conditions.

The WFIRST telescope is expected to experience jitter from two main sources, the cryogenic cooler and the reaction wheels, which both contribute to about 15 mas of tip/tilt jitter with dominant modes at 45 and 95 Hz, respectively. In order to effectively correct these modes, the jitter must be sensed about ten times faster than the fastest frequency. This leads to the conclusion that an approximately 1 kHz rate of sensing is needed to control the jitter.

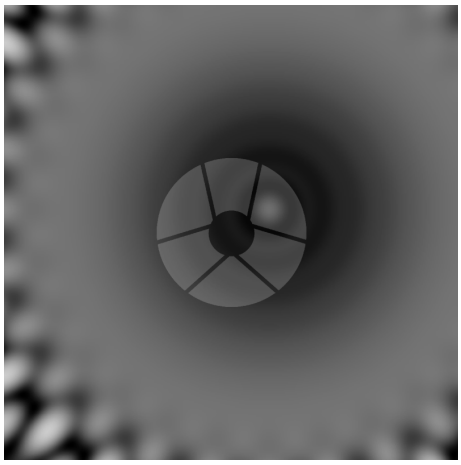

(a)

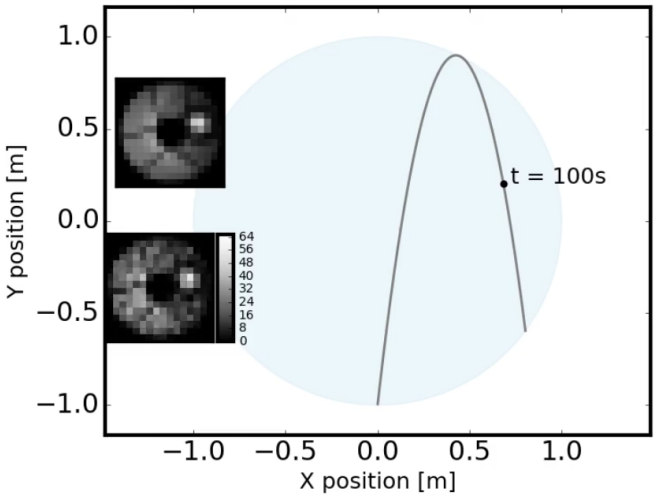

(b)

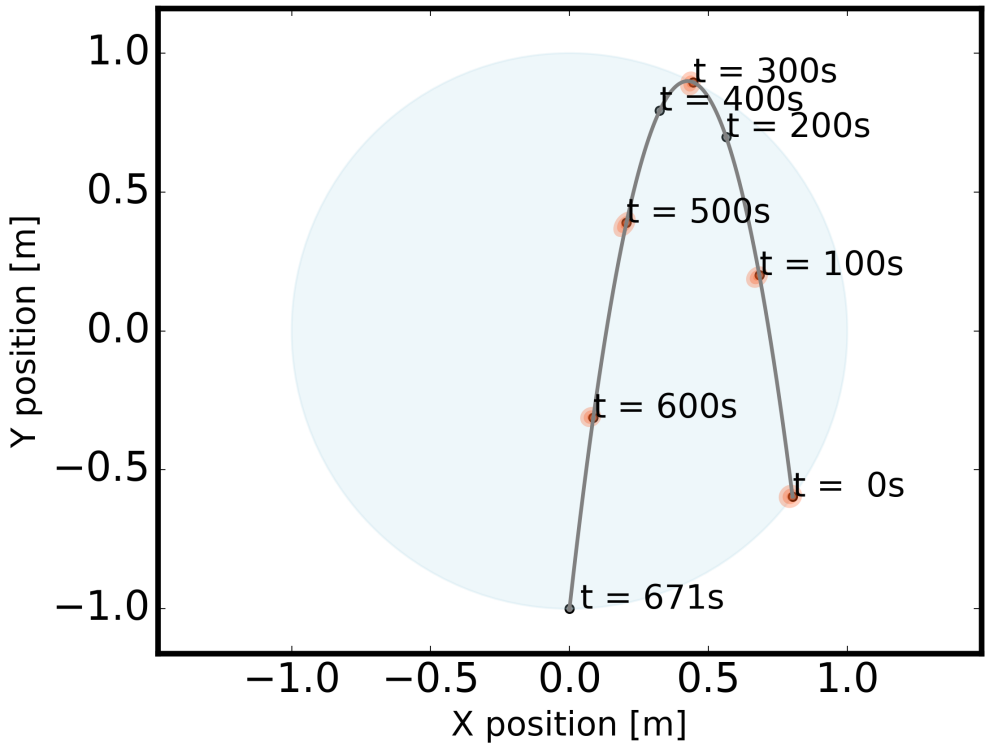

(c)

Figure 4: Pointing precision example for the red science band for a 8th magnitude G-type star and exposure time of 1 second. (Animations of the trajectory simulation may be found at *mb2448.github.io/ss_traj.zip*) Here the starshade is performing an eleven minute trajectory correction maneuver to stay in the dark hole, firing its thrusters at time $t = 0$ and having a quasi parabolic trajectory in the L2 gravity gradient. (a) shows an overlay of the telescope pupil on the intensity (compare to Figure 2) when decentered by about 0.75 meters, about 100 seconds after the thruster fire. The upper inset of (b) shows what the detector would see in the absence of noise, the lower shows the same when Poisson and detector read noise are added. One thousand realizations of noise are used to determine the sensing precision by comparing to a library of clean images at different positions. The blue circle corresponds to the 1 meter radius dark hole, and the parabolic shape shows the dynamic trajectory after the thruster fire. (c) shows the position sensing precision (orange ellipsoids) when calculated at five different points along the trajectory. The 1-σ error in this case is about 2 cm.

It is clear that there is not enough starlight available to sense the jitter at 1 kHz. For example, in Table 2, it can be seen that just determining the position requires about 1 second of time for stars of 6th to 8th magnitude, so in 0.001 seconds (ie 1kHz), there will simply not be enough light there. (Sensing 1000 times faster corresponds to a radiometric loss of 7.5 magnitudes.) Sensing very slow drifts is possible, on the other hand. By creating a library of both tip/tilt and positional offsets and simultaneously solving for both, we were able to determine the tip/tilt to a few milli-arcseconds RMS with a few seconds of exposure time. This is probably sufficient for pointing drift control, but will do nothing to control jitter. A solution to the jitter issue can be found by mounting a laser on the starshade and using the laser beam to illuminate the telescope optics. In this way, the laser beam provides a "reference" light source, and can be made bright enough to enable measurement of the fast reaction wheel and cryocooler oscillations. To get a sense of the light levels involved, a 20 mW[§] laser beam at 550 nm and with a 1.2 degree opening half angle provides 69000 photons/m^2/s about the same number of red science band photons as a post-starshade 4th magnitude star.

We first define two operational modes of the starshade tip/tilt sensing. In the first mode, we assume that we know essentially where the starshade is in shear (to $\pm$7cm), so that we may subtract off the shear signal from the camera data. In this case, the wavefront sensor will operate continuously and constantly update shear and tip/tilt value. In the second mode, we assume we don't know where it is at all, so we need to overpower the star signal with the laser to extract the tip/tilt value. Here, the laser will turn off periodically to measure the shear position of the starshade.

To determine how powerful the laser needs to be, we simulate camera images including starlight, laser light, and appropriate Poisson and read noise at known tip/tilt and shear offsets. We then recover the tip/tilt offsets using the same matched-filter technique as discussed before. (For the first sensing mode, we first subtract the approximate shear signal of the starshade from the camera data.) Running this simulation hundreds of times allows for a measure of the precision of the recovery. Increasing the laser power can bring the error in position recovery down to acceptable levels. We show two such simulation results in Figure 5. Different lasers will be needed for the different science wavelengths, and we chose laser wavelengths that are common and can be used at high power settings (550 and 950 nm, corresponding to HeNe and Nd:YAG).

It turns out that a moderately powerful laser is needed even in the case where the starshade position is known. In the red science case, a laser of power 0.5 Watts is sufficient, while for the blue science case, a laser of power 2 Watts is sufficient. For the case where the strategy is to overpower the stellar signal with the laser, 2 Watts and 5 Watts respectively for the red and blue science case is sufficient.

3. INITIAL ACQUISITION OF THE STARSHADE

Starshades must slew long distances between to re-target, and these slews can take weeks, spanning thousands of kilometers. Sensing the position of the starshade during this maneuver requires a combination of coarse and medium range sensing. The coarse range sensing is provided by ground-based facilities (such as the Deep Space Network), and can deliver 100 km of precision. For medium sensing, a small star camera mounted on the space telescope simultaneously observes background stars and a blinking optical beacon on the starshade. The medium sensor can acquire the starshade to a shear offset of approximately 30 meters,[5] at which point the stellar and starshade signal start to blend.

To get from the 30 meter medium-bearing range to the 1 meter range necessary for the fine bearing sensor to take over (and the science operations), a proposed strategy has been to perform a "putt" into the dark hole, combining knowledge of the starshade position from the medium bearing sensor with on-board velocity knowledge. The positional uncertainty before the "putt" maneuver is about 0.4 meters, and once the maneuver is initiated, the medium bearing star camera does not provide any further information, as the partially blocked starlight and beacon light are not separable. Therefore, the thruster fire towards the dark zone necessitates traversing a "blind spot" in the sensing where no direct positional information is available, and estimation combining the onboard accelerometers and a Kalman filter must be used to acquire the dark hole with a single thruster fire. If the thruster fire is too slow, the position uncertainty can grow unacceptably fast, but if the maneuver is too fast, then the spacecraft will lack sufficient thrust to stop once it gets into the dark hole. While

[§]Note that "laser power" refers to the optical power leaving the device, not the input power supply.

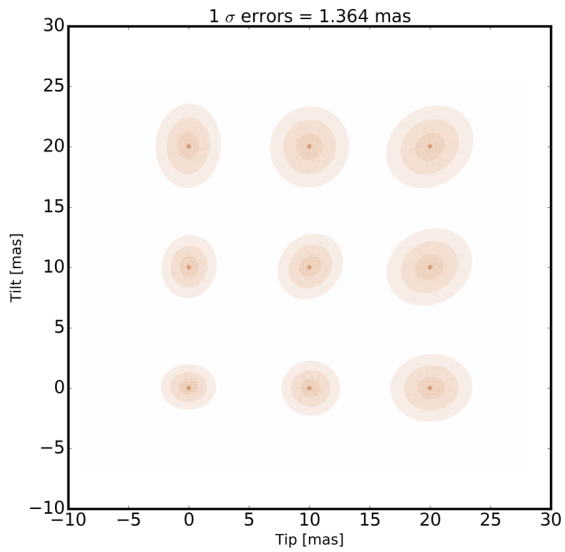

(a) Red science, 550 nm laser, 0.5 Watts

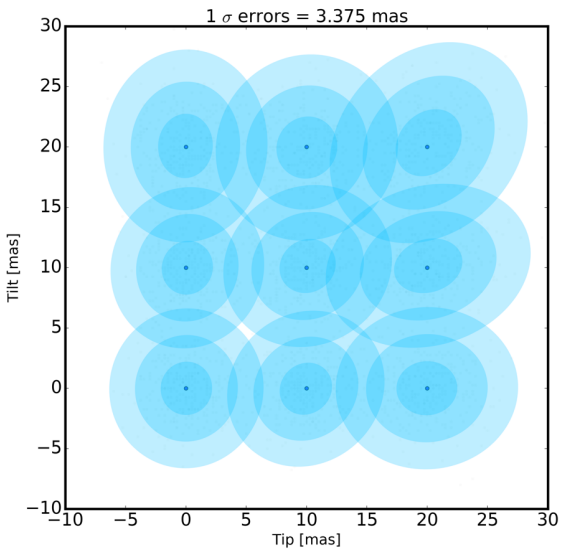

(b) Blue science laser, 950 nm, 2 Watts

Figure 5: Simulations of tip/tilt pointing precision. Tip/tilt combinations are injected at 10 mas intervals, along with stellar light from a 1st magnitude star and laser. Noise from Poisson and readout noise is added corresponding to a 1 ms exposure. After subtracting an estimate of the shear position (carried by the starlight), the tip/tilt is calculated using a matched filter. The spread in recovered tip/tilt offsets shows the pointing precision. (a) a 0.5 Watt laser at 550 nm is sufficient to recover the tip/tilt offsets to better than 2 mas, for the "red science" case. (b) a 2 Watt 950 nm laser is sufficient to recover the tip/tilt offsets to better than 3.5 mas, for the "blue science" case.

appropriate choices of velocity have shown the maneuver to succeed with high probability, here we show that it is possible to avoid the uncertainty altogether, as the pupil plane sensor can easily provide cm-level precision in a 10 meter radius away from perfect alignment. This reduces the risk from the putt maneuver and ensures there are no significant "blind spots" in the acquisition process.

There are a few reasons why sensing the position using the pupil sensor after the medium-bearing sensor is advantageous. First, outside of the dark hole the brightness climbs rapidly, so there are relatively more photons to work with. Second, there is plenty of diffracted structure with spatial scales smaller than the telescope pupil (see Figure 2); this makes feature recognition relatively simple with a matched filter. Finally, even in the absence of these features and any image library for a matched filter, the decline in flux towards the center allows for a simple "gradient descent" algorithm to effectively acquire the dark hole.

To determine the effectiveness of the positional acquisition, we injected the telescope pupil at various radial positions separated by 0.5 meters in a 20x20 meter square surrounding the dark hole, and simulated different camera images of 1 second length targeting an 8th-magnitude star. Tip/tilt jitter and camera noise were injected at the 15 mas and 3e/pixel/frame level. We only used the higher-throughput red science band in this case, as the initial bearing acquisition can be done at any wavength. We recovered the position using a matched filter against a library of known (noiseless) camera images of known location in this 20x20 square. Repeating this experiment hundreds of times allows for determination of the accuracy and precision of this recovery procedure.

Results are shown in Figure 6. The positional determination is precise at no worse than an 8 cm level out to ten meters radially. The worst accuracy is just outside the dark hole zone (1-3 meters out), where there are fewer features and the flux is somewhat decreased compared to further out.

One drawback of this simulation is that we did not include motion blur. The "putting" paradigm calls for starshade velocities in the 5-15 cm/s range for dark hole acquisition, which would increase the error the positional retrieval if not accommodated, as this is similar to the error due to shot noise and camera noise in the 1-second

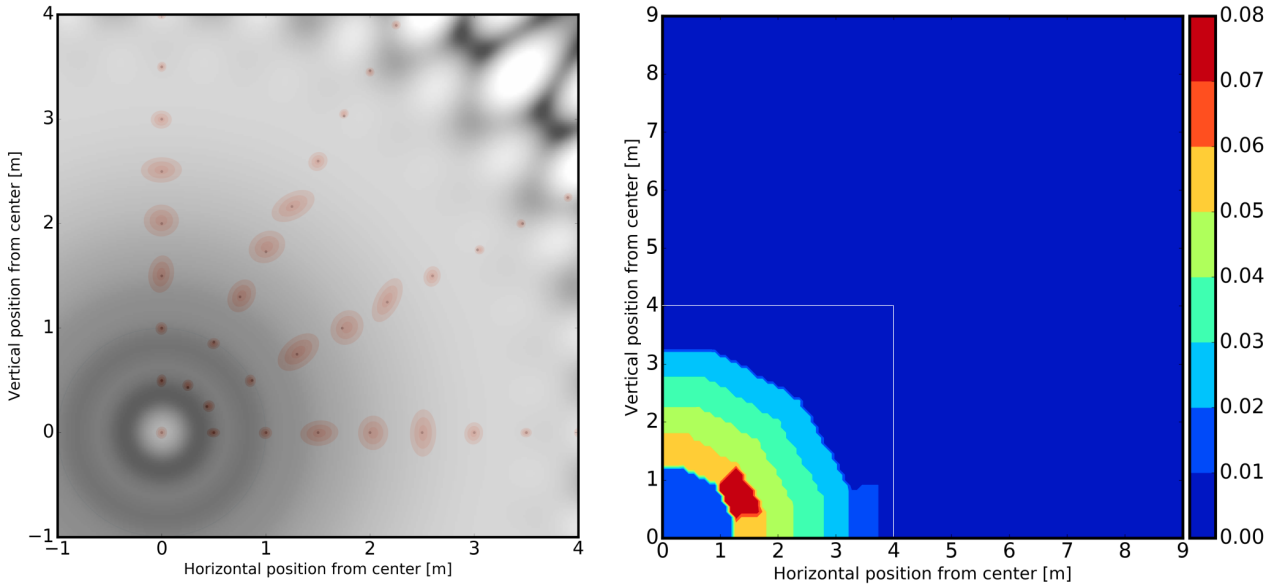

Figure 6: Results of the initial acquisition simulation. (left) the 1-σ error ellipses corresponding to a 1-second exposure when behind an 8th magnitude star overlaid against the (scaled) intensity distribution, showing an inset of a 4 meter radius from the center of the dark hole. (right) the positional accuracy out to 9 meters, with values interpolated. The right inset square is indicated on the image. For the entire field, the errors in position are below 10 cm, showing that initial acquisition can be easily accomplished with a pupil-plane sensor.

exposures around the faintest stars in the target list. However, either slowing the starshade down or updating the image library to include simulated motion blur should gain back much of this around these stars. It should be noted that the 5 cm/s speed was determined because of the lack of positional information; in this case we would be able to measure the position fairly easily. For brighter stars this effect could probably be ignored as the exposure times could simply be reduced–for example, for a 3rd magnitude star the exposure time for the same accuracy would be about 0.01 s, during which the starshade would have moved less than 2 mm. This is well below the photometric error of the position measurement.

4. LAB DEMONSTRATION

We are building a laboratory experiment to demonstrate the guidance and control aspect of the starshade at realistic Fresnel numbers, similar in spirit to Kim et al.,[11] but operating with more relaxed tolerances, as the guiding signal is present at 10^{-3} to 10^{-4} contrast rather than the 10^{-9} required to image planets. The setup of the lab experiment uses a red diode fiber laser to simulate starlight, collimated by an off-axis parabola to a $\sim$ 30mm beam incident on a small (6mm) starshade. The starshade is a metal mask mounted on a glass substrate, which is coated with anti-reflective material to help with stray light mitigation. The entire starshade mount is connected to movable stages to simulate the local gravity gradient and corrective thruster fires. An imaging sensor about 2 meters away from the starshade records the intensity pattern of the diffracted light.

Initial results with the testbed used a different setup, where a fiber off-axis parabola collimator was used to create the flat beam incident on the starshade. The collimator turned out to be of insufficient optical quality, leading to too much scattered light, and also had a somewhat smaller beam diameter than was desirable. The second attempt used doublet lenses to collimate the diverging fiber output beam. These had reasonably good optical quality, but diffraction from the edges of the mounts caused too much diffraction to be useful. These imperfections and their effects on contrast were confirmed theoretically using a physical optics propagator package (POPPY[12]) to model our lab setup.

Some results from these experiments are shown in Figure 7. While there is good agreement between the experiment and model contrast at the brighter parts, the core is brighter than required by more than a factor

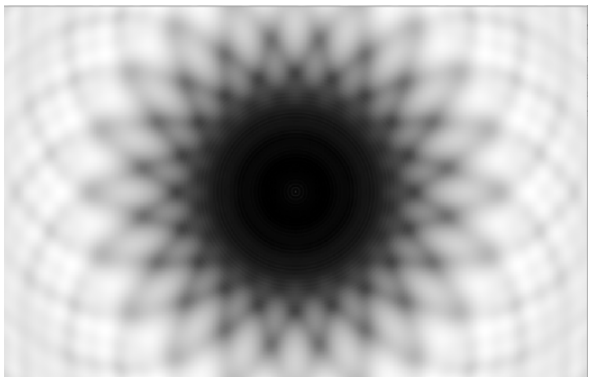

(a) Numerical simulation of lab experiment

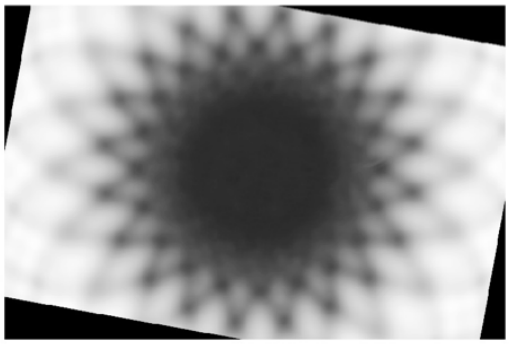

(b) Actual lab data

Figure 7: Comparison of lab data with numerical simulations. The data are shown on the same scale and stretch, with the lab data being rotated by about 10 degrees to match the orientation of the starshade mask. While the outer parts of the image show good agreement, the center is too bright in the lab data. This data was taken with our previous, suboptimal setup using lenses to collimate the beam, which suffered from edge diffraction effects.

of 10 in the lab setup. At the time of this submission, we were installing a higher quality off-axis parabola as our beam collimation source, with the expectation that this will provide the appropriate beam quality, size, and contrast to sense the guiding spot in the center of the field.

5. CONCLUSION

We have presented numerical simulations demonstrating the ability to perform precise stationkeeping with a WFIRST-type telescope and a starshade, for stellar brightnesses typical of exoplanet direct imaging mission targets. Sensing of the starshade position to <5 cm is achievable with modest exposure times for all target stars and in all wavelength ranges of interest. Furthermore, similar arguments show that sensing the position of the starshade from 10-20 meters away from perfect alignment is possible, negating the risk associated with a blind "putting" maneuver to acquire the starshade initially.

Sensing tip/tilt is more challenging. We showed that while slow drifts (over second timescales) can be sensed using the diffracted starlight and an internal Zernike wavefront sensor, fast jitter at the ∼100 Hz requires sensing at 1 kHz to control, well beyond what is possible when using just starlight. The solution is to augment the signal using an onboard laser at the starshade, but we found that relatively high powered lasers producing watts of optical power are required. We also argued that the modest gains in planet throughput expected from correcting tip/tilt may not be worth the additional complexity of laser pointing systems.

Finally, we presented some initial work on validating these simulations in a lab setting. While initial agreement is encouraging, more work is needed to match experimental results to our optical models of the testbed.

ACKNOWLEDGMENTS

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. (C) 2017 California Institute of Technology. Government sponsorship acknowledged.

REFERENCES

1. W. Cash, "Detection of Earth-like planets around nearby stars using a petal-shaped occulter," *Nature* **442**, pp. 51–53, July 2006.
2. M. C. Turnbull, T. Glassman, A. Roberge, W. Cash, C. Noecker, A. Lo, B. Mason, P. Oakley, and J. Bally, "The Search for Habitable Worlds. 1. The Viability of a Starshade Mission," *PASP* **124**, p. 418, May 2012.

3. T. Glassman, A. S. Lo, J. Arenberg, W. Cash, and C. Noecker, "Starshade scaling relations," in *Techniques and Instrumentation for Detection of Exoplanets IV*, *Proc. SPIE* **7440**, p. 744013, Aug. 2009.
4. M. C. Noecker, "Alignment of a terrestrial planet finder starshade at 20-100 megameters," in *Techniques and Instrumentation for Detection of Exoplanets III*, *Proc. SPIE* **6693**, p. 669306, Sept. 2007.
5. D. P. Scharf, S. R. Martin, C. C. Liebe, Z. H. Rahman, C. R. Seubert, M. C. Noecker, and G. H. Purcell, "Precision formation flying at megameter separations for exoplanet characterization," *Acta Astronautica* **123**, pp. 420–434, June 2016.
6. W. C. Anthony Harness, "Enabling formation flying of star shades for the search of earth-like exoplanets," in *8th International Workshop on Satellite Constellations and Formation Flying*, 2015.
7. E. Cady, "Boundary diffraction wave integrals for diffraction modeling of external occulters," *Opt. Express* **20**, pp. 15196–15208, Jul 2012.
8. M. N'Diaye, K. Dohlen, T. Fusco, and B. Paul, "Calibration of quasi-static aberrations in exoplanet direct-imaging instruments with a Zernike phase-mask sensor," *A&A* **555**, p. A94, July 2013.
9. F. Castelli and R. L. Kurucz, "New Grids of ATLAS9 Model Atmospheres," *ArXiv Astrophysics e-prints* , May 2004.
10. A. N. Cox, *Allen's astrophysical quantities*, 2000.
11. Y. Kim, D. Sirbu, M. Galvin, N. J. Kasdin, and R. J. Vanderbei, "Experimental study of starshade at flight Fresnel numbers in the laboratory," in *Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave*, *Proc. SPIE* **9904**, p. 99043G, July 2016.
12. M. Perrin, J. Long, E. Douglas, A. Sivaramakrishnan, and C. Slocum, "POPPY: Physical Optics Propagation in PYthon." Astrophysics Source Code Library, Feb. 2016.