

S5: Starshade technology to TRL5

Milestone 4 Final Report:

Lateral formation sensing and control

Thibault Flinois¹, Michael Bottom¹, Stefan Martin¹, Daniel Scharf¹, Megan C. Davis¹, Stuart Shaklan¹

¹Jet Propulsion Laboratory, California Institute of Technology

14 Nov 2018

This document is cleared for public release (JPL CL 19-0503).

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

© 2018 California Institute of Technology. Government sponsorship acknowledged.

Table of Contents

1	<i>Introduction</i>	3
1.1	Starshades and formation flying	3
1.2	Statement of technology gap	3
1.3	Approach to TRL5	4
2	<i>Summary of results</i>	4
3	<i>Lateral sensing</i>	6
3.1	Lateral sensing approach: overview	6
3.2	Radiometric calculations	7
3.2.1	Target star brightness	7
3.2.2	Starshade optical design	7
3.2.3	Telescope optical efficiency	8
3.3	Analytic simulations	9
3.4	Numerical simulations	10
3.4.1	Computational requirements	12
3.4.2	Range extension and initial acquisition	13
3.5	Laboratory demonstration	14
3.5.1	Overview	14
3.5.2	Testbed overview	14
3.5.3	Experimental design.....	16
3.5.4	Results.....	18

3.6	Results and discussion	19
4	Formation flying	21
4.1	Formation flying approach: overview	21
4.2	Simulation setup	22
4.2.1	Simulation environment overview	22
4.2.2	Initial conditions	23
4.2.3	Thruster models	24
4.2.4	Sensor models	25
4.3	Prototype control algorithms	26
4.3.1	Estimation	27
4.3.2	Control	27
4.3.3	Thrust allocation	29
4.4	Monte Carlo simulations	30
4.5	Results and discussion	31
4.5.1	Typical formation flying behavior	31
4.5.2	Overall Monte Carlo simulation results	33
4.5.3	Relevance of results to HabEx	35
5	Conclusions	37
6	Appendix	38
6.1	List of acronyms	38
6.2	Table of formation flying simulation parameters	39
6.3	Matching algorithm and computational steps	42
6.4	Comparison between SLATE testbed and flight expectation	42
7	Bibliography	43

1 Introduction

1.1 Starshades and formation flying

Starshades are one of the most promising technologies that may enable direct observations and characterization of Earth-like exoplanets in the next decade. A Starshade is a precisely-shaped screen that is designed to be carefully positioned at a specific distance along the line of sight of a telescope, which is pointed towards a target star. The Starshade shape is optimized to ensure that the starlight diffracting around it does not enter the telescope aperture. This creates a very dark “star shadow” and allows direct imaging of much dimmer exoplanets orbiting the target star.

The NASA Astrophysics Division has commissioned a Starshade Technology Development Activity (S5) whose purpose is to advance Starshade technology to Technology Readiness Level 5 (TRL5). This report focuses specifically on the advancement of Starshade lateral formation sensing and control technology as detailed in Section 2.2 of the S5 Technology Development plan [1]. As stated in that plan, the work needed to bring the lateral position sensing to TRL5 involves three main tasks:

“First, the optical model that predicts the out-of-band suppression pattern will be verified by comparison to images collected in the SLATE. Next, the algorithm that infers lateral offset distance by comparison to an offset image library will be tested by applying it to images collected at the SLATE with known offset and the flight SNR. Finally, a MATLAB model that simulates the lateral control servo using the lateral sensing algorithm will be run to demonstrate that lateral position control.”

The aim of this report is to document the work done as part of these activities and show that the results achieved demonstrate that the sensing of the Starshade lateral offset with respect to the telescope-to-star axis is at TRL5.

1.2 Statement of technology gap

The purpose of formation flying and lateral control is to ensure that starlight does not leak into the telescope aperture and overwhelm the faint exoplanet signal. The shadow cast by the Starshade has a finite extent, so the telescope and the Starshade must fly in formation such that their relative position is maintained to the accuracy required to perform exoplanet science. This must be done in the presence of differential gravitational and solar radiation pressure accelerations experienced by the two spacecraft. For missions like WFIRST and HabEx, the Starshade shadow is ~2 m wider than the telescope pupil, which leads to a ~1 m requirement on lateral control. The following technology gap, with corresponding milestone and “Key Performance Parameter” have thus been formulated:

S5 technology gap: *“the ability to sense and control the lateral offset between the Starshade and the telescope maintaining the desired contrast long enough for full science integration. The*

technology required to close this gap is a validated technique for sensing lateral displacements of the Starshade from the line of sight between the telescope and exoplanet host star to the necessary precision and accuracy.” [1] [2]

S5 Milestone: *“Starshade Lateral Alignment Testbed validates the sensor model by demonstrating lateral offset position accuracy to a flight equivalent of ± 30 cm. Control system simulation using validated sensor model demonstrates on-orbit lateral position control to within ± 1 m.”*

S5 Key Performance Parameter: *“verify lateral position sensor accuracy and that it supports ± 1 m control via simulation”.* The lateral sensing accuracy allocation is of 30 cm (3σ), with a further goal of sensing to 10 cm (3σ), assuming a measurement rate of 1Hz. [1]. Note, this 1 Hz rate is an allocation that ensures the Starshade receives a sufficient number of measurements as it crosses the lateral control region.

1.3 Approach to TRL5

Bringing this technology to TRL5 revolves around two main activities, which form the basis of the two main sections of this report:

- **Section 3:** development of a lateral sensing scheme appropriate for target flight environment; demonstration of its end-to-end performance with analytic, numerical and experimental data; delivery of a sensor model of expected flight capability in operational environment.
- **Section 4:** demonstration of formation flying performance using this sensing scheme with extremely conservative performance assumptions, in a high-fidelity end-to-end simulation of the space environment, corresponding to stressing operational conditions.

Although the context of TRL 5 assumes the Starshade Rendezvous mission with WFIRST in establishing its KPPs, article fidelities, and relevant environments, the technology is not tightly constrained to the current baseline WFIRST design. The formation flying sensing approach and control algorithms would be generally applicable to any optical space telescope of ~ 2.4 m aperture operating in an Earth-Sun L2 orbit, so long as it includes a pupil plane sensor consistent with the lateral position sensing concept.

2 Summary of results

We developed a sensor scheme based on imaging of the residual starlight diffracted by the Starshade in a pupil imager. The sensor uses an “image library” matching algorithm to determine the lateral displacement of the Starshade. To analyze the expected performance of the sensor, a thorough radiometric (throughput) budget was first created. Second, a simple analytic model was developed to get an initial assessment of the sensor accuracy. A detailed numerical model was then constructed using optical propagation from the star to the detector. This model was used along with the image matching algorithm to predict the flight sensor performance, finding good agreement with the simple analytic model. Finally, the image matching approach was validated in a laboratory experiment, finding good agreement with numerical simulation. **The analytic and**

numerical simulations predicted flight sensor performance better than the requirement of sensing the position to 30 centimeters in 1 second of exposure time, achieving a precision of less than 10 centimeters for stars *at least ten times fainter* than any target stars. Lab validation of the sensor model reproduced numerical simulation results to within a factor of ~2 in lateral precision, still well within the requirement.

We then developed a high-fidelity formation flying simulation environment to demonstrate that the proposed lateral offset sensor is adequate for controlling the lateral position of a Starshade with respect to a telescope within a 1m-radius. This was done by designing flight-like estimation and control algorithms that command Starshade thrusters to counteract the relative lateral acceleration experienced by the Starshade and the telescope. The estimation of the relative lateral position is based on the proposed sensing scheme, using extremely conservative assumptions for the sensor performance. As thruster firings interrupt scientific measurements, the control algorithms were designed to maximize the drift time between thruster firings. **Through Monte Carlo simulations, it was demonstrated that the shear sensing approach allows successful control of the relative lateral position of the Starshade to within 1 meter. This was achieved across all varied parameters, in stressing operational conditions, and even with the extremely conservative sensor performance assumed. It was also shown that thruster firings are only required every 10-15 minutes, thus providing high observational efficiency for Starshade science.**

3 Lateral sensing

3.1 Lateral sensing approach: overview

S5 requires a sensing scheme that allows for determining the displacement of the Starshade with respect to the telescope line of sight to the star to an accuracy of 30 cm (3σ) in 1 second. The sensor scheme we developed to meet this challenge is based on imaging the residual starlight diffracted by the Starshade in a pupil imager inside the WFIRST telescope, intended as an internal wavefront sensor for the coronagraph.¹ The concept of sensing starshade displacement via leaked starlight has long been discussed in the literature. [3] [4] [5] [6] The coronagraph instrument uses a pupil imager for a high-speed pointing control loop, which is not a technology gap and is not used in Starshade operations.

A Starshade converts the flat incident wavefront into a complex pattern at the pupil of the telescope, which consists of a central Arago spot tens of centimeters wide and radial structures due to the petals (Figure 3-1). This pattern is extremely faint ($<10^{-10}$) at the science wavelength, but brightens substantially when viewed outside the designed science band (10^{-3} - 10^{-4}); see Figure 3-2. The Arago spot moves with the Starshade position, and a pupil imager can be used to infer the relative displacement by determining the position of the spot using light outside the science band.

The algorithmic approach we developed uses “image library” matching to determine the displacement. [7] A grid of precalculated optical models of the Starshade optical field at different offsets, propagated through the telescope, is compared to the data from the pupil imager. A least-squares minimization algorithm determines the closest image match and corresponding offset in x and y .

The sensor model depends on the amount of starlight incident on the Starshade, the effect of the Starshade and telescope optics, and the sensing algorithm. To analyze these effects, we first developed a radiometric (throughput) budget consisting of the target stars, Starshade, and internal telescope optics. Second, we developed a simple analytic model of the sensor performance using the radiometric budget and scaling arguments. We then constructed a detailed numerical model using optical propagation between the star, Starshade, telescope, and detector. With this model, we used the image matching approach to predict the flight sensor performance, finding good agreement with the simple analytic model. Finally, we validated the image matching approach in a laboratory experiment, finding good agreement with numerical simulation.

The analytic and numerical simulations predicted flight sensor performance exceeding the requirement of sensing the position to 30 centimeters in 1 second of exposure time, achieving better than 10 centimeters of precision for stars *at least ten times fainter* than any target stars. Lab validation of the sensor model reproduced numerical simulation results to within a factor of ~ 2 , still well below the requirement.

¹ In the WFIRST coronagraph, the pupil imager is used in conjunction with a Zernike wavefront sensor, an elegant optical device that converts the flat intensity of the pupil image into an interference pattern encoding wavefront errors.

3.2 Radiometric calculations

The performance of the lateral sensing approach depends primarily on the amount of starlight available at the detector. This depends on

- the brightness of the target star (Section 3.2.1)
- the optical design of the Starshade and chosen wavelength range for sensing (Section 3.2.2)
- the internal optics of the telescope, including the detector efficiency (Section 3.2.3)

The radiometric calculations in this section are used as inputs to the analytic and numerical simulations (Sections 3.3, 3.4) of the sensor performance, and for signal-to-noise settings in the experimental results (Section 3.5).

3.2.1 Target star brightness

Starshade targets are typically bright, nearby main-sequence stars of spectral types FGK. While the target list has not been finalized, the latest version [8] consists of fewer than 30 stars ranging in magnitude from -1.4 (Sirius A) to 5.3 (Gliese 783A). The range in brightness of these stars is nearly a factor of 1000.

Since the range of brightness is mostly due to this magnitude range, and not the spectral differences (for FGK stars, photon flux rates vary by a maximum of 50% from the G-type stellar spectrum), we elected to use a G-type stellar atmosphere spectrum for the simulations. Specifically, we used a G5V ($T_{\text{eff}} = 5750\text{K}$, $\log g = 0.0$, $[\text{Fe}/\text{H}] = 0$) spectrum from the ATLAS9 stellar atmosphere models, a standard library used widely in astronomy [9]. The radiometric propagation of the G5V spectrum was validated against in-orbit measurements of the ASTM extraterrestrial solar reference spectrum [10], at a distance of 1 AU, with agreement at the 2% level. Further checks of fluence (photons/m²/s) against photometric zero points of standard astronomical filters were consistent at the ~5% level, with the discrepancy likely due to filter transmission function approximations. The input stellar fluxes are thus a negligible source of error or uncertainty for the sensor model.

3.2.2 Starshade optical design

The NI2-24 Starshade (designed by Eric Cady), is 26 meters tip-to-tip, with 24 petals (see Figure 3-1). Each petal is 8 meters long and 2 meters wide, and the geometric inner working angle (angular distance from center to tip) is 72 mas at the nominal operating distance of 37,242 km. It is designed to provide a deep starlight suppression simultaneously across a 425-552 nm band for a 2.4 m telescope located anywhere within the 4.4m-diameter shadow region.

Starshade optical models predict the wavelength-dependent transmission as a function of distance from the occulter. Significant progress has been made on numerical modeling of Starshade diffraction patterns, and modeling results now match experiments at the $<10^{-8}$ level [11].




Figure 3-1: (Left) NI2-24 Starshade design (E. Cady). The Starshade is ~26 meters tip to tip. (Right) Light pattern at the location of the WFIRST telescope, stretched to show detail, with the 2.4 meter pupil shown for scale.

The shear (Starshade lateral offset) sensing operates in a very different regime from 10^{-8} , using diffracted light outside the science band with contrast at the 10^{-3} - 10^{-4} level. As such, we have high confidence in the optical Starshade models, as they have been validated at the fourth or fifth decimal place for our relevant contrast levels. For this work, we used these well-validated codes to compute Starshade contrasts at the relevant wavelengths, using the Starshade design baselined for the WFIRST-Rendezvous mission.

3.2.3 Telescope optical efficiency

We based the telescope efficiency measurements on the latest design of the WFIRST coronagraph internal optical chain. [12] The optical efficiency was derived from spectral transmission curves of the ~20 optical elements and their coatings, including 8 coatings of HRC Silver, 9 coatings of Denton Silver, 2 uncoated aluminum deformable mirrors, and 3 coated lenses in the low-order wavefront sensor path. An additional 10% loss was added to simulate the Starshade dichroic [13]. (See Figure 3-2 for the efficiency curve). The detector efficiency was taken from the EMCCD-201 datasheet, the chip baselined for the low-order wavefront sensor camera, with the flight-qualified “midband” coating.




Figure 3-2: Throughput parameters for the numerical simulations. In the green science band, guiding is accomplished with a small portion of the leaked light around the Starshade. The yellow line shows a representative spectrum of a G5V star, the dashed line shows the total system efficiency (detector QE + optical efficiency), and the green line shows the Starshade transmission.

While it is possible that the total telescope efficiency may change in later designs, even extreme changes to CGI optics (e.g., 50% efficiency loss) would not significantly affect the sensor model, as shown in the following sections.

3.3 Analytic simulations

It is possible to do a rough calculation of the expected performance of the shear sensor without reference to the particular details of the electric field propagation through the telescope, using only basic Fresnel scaling relationships, knowledge of the optical system efficiency and detector parameters. While this does not substitute for detailed modeling, it is a useful sanity check on the results from numerical simulations and laboratory demonstrations, and provides insight into the scaling of the precision with target star magnitude.

The analytic approach relies on the following steps. First, the Arago spot size (in physical units) can be predicted without knowledge of the detailed shape of the Starshade petals (near the optical axis, the intensity pattern of the Arago spot closely resembles a Bessel function). Under these assumptions, $r \sim \text{FWHM} \sim D / (4\pi F)$, where r is the Arago spot radius, FWHM is the spot full-width at half-maximum, D is the Starshade diameter, and F is the Fresnel number of the system. The absolute intensity of this spot can be predicted from the system efficiency as in the previous section. The signal-to-noise of this spot can be determined from the plate scale of the detector and its noise via the “CCD equation,” [14] as $\text{SNR} \sim N_{\text{spot}} / \sqrt{N_{\text{spot}} + n_{\text{ap}} \cdot \text{RN}^2}$, where SNR is the signal-to-noise ratio, N_{spot} is the number of photons in the spot, n_{ap} is the number of pixels taken up by the spot, and RN is the readout noise (per pixel) of the detector. Finally, the centroid precision can be determined from a well-known scaling of centroid precision, spot full-width half-maximum, and signal-to-noise ratio commonly used in astrometry [15]; $\sigma = \text{FWHM} / (c \cdot \text{SNR})$,

where σ is the centroid precision, FWHM is the spot full-width at half-maximum, SNR is the spot signal-to-noise ratio, and c is a constant which depends on the spot shape, but is usually about 2.

There are a few things missing in this approximation. First, while the central Arago spot size can be predicted fairly well, the region surrounding the spot is not well approximated by a Bessel function, and these regions will contribute to the centroid approximation and the matched-filter approach. Details of the off-axis optical intensity are totally ignored and captured by one number, the “Starshade contrast.” Similarly, effects due to pupil and spider obscuration of the telescope are ignored. As stated before, these effects will make the analytic approximation somewhat inaccurate, but it is still useful as an order-of-magnitude check.

Taking values for the efficiency in the previous section, Starshade diameters of 26 meters, and detector sampling of 32 pixels across the pupil with a read noise of $3e^-/\text{pixel}/\text{frame}$, results of the analytic model are presented in Figure 3-3. Note that for stars brighter than 6th magnitude, which comprise all the target stars, the analytic model predicts 3σ precision better than 2 cm in all science bands.

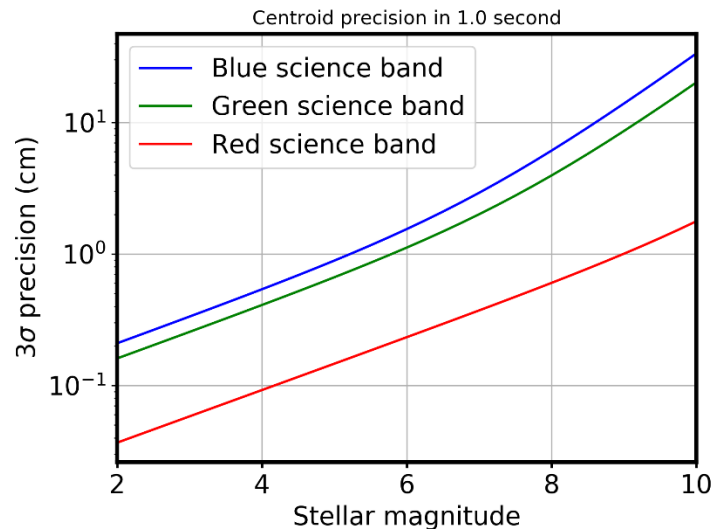


Figure 3-3 Analytic model of Starshade shear precision, for the red, green, and blue science bands, in one second of integration time, calculated at wavelengths 496 nm, 424 nm, and 937 nm. Note all target stars have magnitudes less than 6.

3.4 Numerical simulations

In addition to the analytic simulations, we constructed a full numerical model of the Starshade and telescope optical path. These simulations used electric field propagation past the Starshade, and through the telescope and wavefront sensor to determine the light pattern and intensity on the wavefront sensor. First, the electric field of the Starshade was calculated using the boundary integral method of Cady et al. [16]. The electric field was then propagated through the telescope model as follows: 1) shifted to account for the input offset positions, 2) multiplied by the telescope pupil aperture function to add the central obscuration, outer diameter and spiders, 3) propagated to the Zernike sensor plane, 4) multiplied by the Zernike phase function, 5) propagated to the low-order wavefront sensor camera, and 6) converted to an intensity. We used Fourier transforms to propagate between focal and pupil planes. This procedure was repeated for a variety of offset positions to build up the image library.

The Zernike dimple is only useful for low-order wavefront sensing; for shear sensing, it serves no function. Its effect on the shear signal is to create mild intensity gradients over the image, and actually scatter a good deal (10-20%) of light outside the pupil imager, similar to a coronagraph. The reason we included it is that it reduces the flux, making it a more conservative assumption; it is expected to exist in the system in a baseline configuration; and we did some preliminary tests on simultaneous shear and tip/tilt sensing, which are not part of this milestone but could be useful in the future.

The image matching algorithm is a least-squares procedure. Each input image SI is converted into a vector of n pixels (for example, n is 1024 for a 32x32 image), is normalized by the sum of the image intensities yielding mSI , and the scalar $e_m = (L_m - mSI)^2$ is calculated for each image in the library, L_m . The library image with the lowest e_m is selected, and its position is taken to be the Starshade position.

The matching algorithm is used with a sufficiently sampled image library such that the grid spacing is much smaller than the required accuracy. Spacings of 2 centimeters over regions of 1.25-2 meters were used in assessing computational requirements and statistical performance over the control region. Better resolution could be obtained by interpolating between the best fit grid neighborhood, but this is ignored due to the negligible effect on the formation flying performance (see Section 4).

To obtain statistical assessments of the sensor precision, individual pupil images were taken from the image library, reduced in intensity by scaling to account for optical efficiencies and star brightness (see Section 3.2) and then degraded with Poisson noise and read noise at the expected level ($3e^-/\text{pixel}$).² The noisy image is then matched to the image library using the least-squares minimization algorithm (see Figure 3-4), and the best-fit offset positions are saved. (Details of this algorithm may be found in the appendix.) Repeating this procedure 300 times with random noise inputs allows Monte Carlo statistical assessments of the offset precision at different offsets, as shown in Figure 3-4.

² Additionally, we examined the effects of jitter by convolving the electric field with a jitter kernel of 15 mas (1σ), the expected level in WFIRST, which had an almost unnoticeable effect.




Figure 3-4: **Left:** A least-squares algorithm matches the noisy camera image to the closest image in the image library, in this case to the image with shear (0.0, 0.6). Running this repeatedly on noisy images allows for statistical determination of the error in shear measurements. **Right:** Output of numerical simulations of guiding in the blue science band ($\lambda_c = 937$ nm) with an 8th magnitude star. The ovals show the 1, 2, and 3 σ error contours, with the requirement as a dashed circle. The grey circle shows the 1 meter control region.

This procedure was repeated for the red, green and blue science bands. **The worst-case precision predicted by the simulations in the control region was 3.0, 5.1, and 11.7 cm (3 σ) in the red, green, and blue science bands, for stars of 10th, 8th, and 8th magnitude, respectively.** Table 3-1 shows more data on the simulation results. This guiding precision is predicted to easily be sufficient for all target stars. Furthermore, we note reasonably good agreement with the analytic predictions, within 10% for the red and green bands, and 40% for the blue band.

Science band	Star magnitude	Guiding wavelengths (nm)	Weighted wavelength (nm)	Median 3 σ error (cm)	Analytic 3 σ error (cm)
Red	10.0	400-540	496	1.6	1.6
Green	8.0	400-435	424	3.6	3.9
Blue	8.0	870-1000	937	9.7	6.1

Table 3-1: Results of numerical simulations for pointing precision. The 30 cm 3- σ error requirement is exceeded in all cases.

We note two limitations of our simulations. First, we used a single, intensity weighted mean wavelength in the model, but did not compute the simulation over a broad band. This is appropriate, because the sharp Starshade transmission function out of band (proportional to λ^x where x is up to 15, see Figure 3-2) means that only a narrow range of wavelengths will contribute meaningfully to the flux, so a monochromatic simulation is sufficient. Furthermore, we did not consider motion blur in the simulations. This effect is also minor, as the maximum Starshade velocity in the control region is expected to be less than 1 cm/second, and required camera exposure times are expected to be well less than one second for most target stars.

3.4.1 Computational requirements

The image library approach is a brute-force solution to the position determination problem, but it does not appear to pose undue memory demands. For a 2m x 2m square control region (more than four times larger than the 1 meter radial deadband) calculated every 1 cm, storing each of the

40,000 32x32 pixel image at 16 bits/pixel comes to 82 Megabytes per channel, for a **total storage requirement of 246 Megabytes**. This number assumes no compression and an unreasonably high bit depth, as 16 bits are excessive for representing the dynamic range of the Starshade shadow.

Is the WFIRST LEON4 computer able to effectively run the image matching algorithm? A coronagraph Tiger team estimated an “effective” number of floating point operations per second (FLOPs) at $7.6e7$. The matching should be performed in less than one second. The number of floating point operations required to preprocess and compare the input image to the library is $3nm+2n+1$ (see algorithm in the appendix), where m is the number of images in the library and n is the number of pixels in the image, taken here to be 1024 ($=32*32$). Allowing for a 40% margin on the computational speed and time, we exceed the allocation if comparing the input image to each of the 40,000 reference images. However, the Starshade does not move very far from frame to frame, as it has a maximum speed of <1 cm/s. As such, the search area may be reasonably limited to a few cm in the 1 second interval between computations. In this case, **the computational demands are easily met using less than 3% of the allocation**, as seen in Table 3-2 below:

Algorithm	FLOPS	Time	% of total
Tip/tilt loop	1.00e5	1.3 ms	0.2 %
Centroiding and relative bearing	1.10e5	1.4 ms	0.2 %
Shear sensing $m=40,000$	4.92e8	6474 ms	1078 %
Shear sensing, 3 cm search area	1.19e5	1.6 ms	0.3 %
Offset sensing, 5 cm search area	3.15e5	4.1 ms	0.7 %
Offset sensing, 9 cm search area	1.00e6	13.2 ms	2.2 %
Sum (w/9 cm search area)	1.21e6	15.9 ms	2.6 %
<i>Total available (40% margin)</i>	<i>4.56e7</i>	<i>600 ms</i>	<i>100.0 %</i>

Table 3-2: Computing budget for image matching algorithm

We note that image matching is not a particularly unique or clever algorithm, and other faster and more sophisticated ones may be devised moving towards flight.³ The reason for using the image library approach here is that the least squares algorithm is functionally equivalent to a matched filter, which minimizes the additional noise added by the reduction algorithm.

3.4.2 Range extension and initial acquisition

The current Starshade acquisition sequence has a transition point between far (>10 meters) and science (<1 meter) shear separations. In the far zone, the Starshade beacon blinks and the modulating signal is measured with respect to the target star to determine the Starshade’s position and velocity. This measurement is done using the wide-field instrument. This works up until the point where the Starshade begins to substantially occult the star, which occurs around a ~ 10 meter shear separation. To acquire after this point, the current plan is to perform a “putting” maneuver where the calculated position and velocity are used to inform a thruster fire with a high probability of intersecting the 1 meter deadband.

³ For example, given the circular symmetry of the problem, it is not even clear that a library is needed—an image gradient based approach would likely perform well. This problem is also a natural choice for a neural-network.

We performed numerical simulations to determine whether the same pupil sensing approach could be used to eliminate this blind spot. The answer is that there does not appear to be any challenge to extending this technique to 10 meters, eliminating the blind spot. The only caveat is the computational storage of the library. However, this can be reduced by using larger grid spacings of 5-10 cm, which are acceptable spacings for a $\sim 100 \text{ m}^2$ area. There is also substantially more light in the outer parts of the diffraction pattern (see Figure 3-1). Understanding whether this should be implemented as a sensing mode for initial acquisition is a topic for future research.

3.5 Laboratory demonstration

3.5.1 Overview

Starshades are designed to operate in the optical Fresnel regime, where the dimensionless “Fresnel number” $F \sim r^2/(\lambda Z)$ is in an intermediate regime of < 20 . Here r is the Starshade radius, λ the wavelength, and Z the separation, and optical propagation physics is preserved when the Fresnel number is the same. Due to the large propagation distances of $> 10000 \text{ km}$, (note the radius of the Earth is $\sim 6400 \text{ km}$), it is not possible to optically validate a full-scale Starshade. However, aspects of a flight-like setup may be tested by quadratically decreasing distance Z with Starshade radius r . Thus, scaling down the Starshade by a factor of one thousand will require a propagation distance one million times shorter, allowing the possibility of lab tests with reasonable propagation distances of 1-100 meters.

We built the Starshade Lateral Alignment Testbed (SLATE) to validate the lateral position sensing algorithms in the lab. The testbed consists of a fiber laser, collimating lens system and Starshade mounted on a motorized x-y stage, a fold mirror to increase path length, and a camera to image the Arago pattern (see Figure 3-5). This testbed operates at Arago spot contrasts of $\sim 10^{-3}$, within the range of the flight expectations of $10^{-3} - 10^{-4}$. **We note that similarly to the simulations, the tests were conducted with simulated star magnitudes at least ten times fainter than any star on the target list.**

It is important to note that SLATE was **not** intended to validate Starshade contrast at levels needed for exoplanet science. These tasks each have their own challenges and associated milestones. The purpose of SLATE was to demonstrate that numerical predictions of the performance of the image matching algorithm do not grossly over-predict sensing accuracy. Additionally, experimental tests could reveal unexpected failure cases or modes.

3.5.2 Testbed overview

SLATE is a beam launcher, a fold mirror, and a camera. The beam launcher consists of an optical fiber, a 100 mm doublet collimating lens, and the Starshade mask mounted in a tube. The tube is small enough to fit on a two-axis stage, creating a movable beam launcher to simulate shear offsets corresponding to Starshade motion. The fold mirror is fixed, and its only purpose is to increase the propagation length on the limited optical table. The camera is also fixed, and takes the place of the low-order wavefront sensor camera.

The camera sees “pupil” images, but we do not have a telescope simulator in the beam to create either the pupil of the WFIRST telescope, the Zernike phase plate, or any of the assorted recollimating and refocusing optics. This is by choice, as every optical surface adds noise and complication. To simulate the telescope pupil, we just mask out the pixels on the camera corresponding to the effective pupil obscuration. These pixels are not expected to be used in the image matching algorithm in flight either.

SLATE can create optical sensing signals similar to those expected in space, but deviates in many ways from a “perfect” formation flying lab setup. Below, we list some differences between an ideal situation and our lab setup. A summary of the differences between the test setup and flight expectation is presented in Table 6-4.




Figure 3-5: **Left:** Schematic of SLATE; movable fiber beam launcher, and fixed fold mirror and camera. **Right:** Image of testbed, partially uncovered. The camera may be translated along the rail to access different Fresnel numbers.

3.5.2.1 Optical considerations

The incident beam on the Starshade in flight will be a flat wave of starlight with effectively no aberration. Our testbed uses a single mode fiber laser collimated by an off-the-shelf doublet lens, producing a Gaussian beam incident on the Starshade⁴ is made of chrome deposited on an optical flat. From a spectral perspective, the use of a laser is reasonable, as the sensing wavebands are not particularly broad; at ~10% optical bandwidth, and one side of the band will be much brighter than the other in most cases. The Gaussian beam shape is also not what would be expected from a star, but this only marginally changes the Arago spot contrast. (For this lab setup, the apodization of the beam is important for preventing diffraction from the edges of the optics and their mounts.)

Any imperfections in the optical system affect contrast. Dust is a major concern, but is fairly easy to remove. Oclusions (scratches or nicks) on the optics are not possible to remove, and require replacing the damaged part. More generally, surface roughness in the optics and cannot be removed or mitigated. Optical surface roughness can be described by a power spectral density (PSD), a curve describing the amount of deviation of the surface as a function of spatial frequency

⁴ Starshade masks were provided by Opto-line

of the deviation. Optical surfaces generally follow a PSD that is flat up to a certain spatial frequency (say, 200 cycles per meter), then rolls off as spatial frequency to the \sim third power. The PSD curve can be described by the RMS amplitude (in nm, dominated by lower frequencies), the roll-off frequency, and the slope of the falloff.

Our analysis of optical tolerances indicated that very stringent scattering requirements would be necessary to operate at Fresnel numbers of ~ 7 , with RMS surface amplitude error of ~ 5 nm required. This is very challenging to achieve with bulk, static optics; for example, even interferometer reference flats are specified to $1/20^{\text{th}}$ of a wave or ~ 30 nm. One way to proceed would be to create an adaptive optics system using a deformable mirror and wavefront sensor. Another approach would be to avoid the use of any optics in the system and use a pinhole and etched Starshade; this is the approach taken by the ~ 70 meter Princeton testbed [11]. Neither of these approaches was particularly appealing due to cost and complexity.

However, the tolerance analysis indicated that operation at Fresnel numbers of 4-5 would be achievable with bulk optics, assuming dust and occlusions were avoided. We opted to operate in this regime, and capture the properties of the higher Fresnel numbers by scaling the relative pupil size with respect to the spot to reproduce the essential morphological properties of the Arago spot.

3.5.2.2 Camera parameters

The performance of the flight EMCCD detector baselined for the LOWFS exceeds the lab camera by factors ranging from 20 (read noise) to 10,000 (dark current). As such, attempts to match exposure times and fluence levels in the testbed to flight levels would result in a very different signal-to-noise ratio than that delivered by the EMCCD. Instead, we adjusted the exposure times and laser power to match the empirically measured signal-to-noise ratio of the spot, which ranged from 3-8 depending on the wavelength.

Another unknown at this time is the pixel scale of the LOWFS camera, which is expected to be between 16 and 55 pixels across the pupil diameter. While the pixel scale will be sufficient for Nyquist sampling the Arago spot, too fine of a sampling can degrade performance as the readout noise per photon will increase. Inside these limits, the pixel scale does not meaningfully affect the performance. We chose a pixel scale of 32×32 pixels, and binned the camera data to this value using linear interpolation.

3.5.3 Experimental design

The experimental setup used signal-to-noise ratios and spot sizes (as fractions of the pupil diameter) determined from the expected flight-like values. These are detailed in Table 6-4. The relative scaling between the testbed motion and flight motion (in units of millimeters per cm) were calculated analytically and verified experimentally. We note that each camera frame had some minor preprocessing applied to it that would not necessarily be duplicated a space-like environment, such as bias/background subtraction and intensity equalization. This is in part due to features of the test camera, which included a fixed bias offset, and confounding factors such as background illumination from neighboring labs. The flight camera controller and background light levels are not expected to be limiting factors in space. Experimentation with different preprocessing showed modest changes in performance.

3.5.3.1 2.4.1 Image library computation

In order to build the sensor model, we first needed to create an image library for the lab. We computed the miniature Starshade diffraction pattern on the sensor with models of the lab optics, which are simpler than the CGI, involving the fiber output beam, collimating lens, and Starshade (they do not include a Zernike phase plate). Obviously, sizes and distances are different as well, at ~ 3 meter separations and a ~ 6 mm Starshade rather than $\sim 40,000$ km separations and a 26 meter Starshade. We took this optical diffraction pattern and computed it at the expected LOWFS camera pixel scale of 32×32 pixels. An example of the computed and measured diffraction pattern is shown in Figure 3-6. While our contrast measurements were consistent to 25% of expectation, note that in the matching algorithm absolute values of contrast are not relevant because the images are normalized.




Figure 3-6: **Left:** SLATE optical propagation pattern in linear and log space. **Right:** Computed optical model corresponding to the testbed setup. Both scales are identical in their stretch.

3.5.3.2 Signal strength determination

We determined the correct laser diode voltage to use by empirically measuring the signal-to-noise ratio of the pixels (after being binned to the LOWFs plate scale) as approximated by $\text{SNR} = \text{mean}/(\text{standard deviation})$. We matched the empirical signal-to-noise ratio to that expected in space from stars fainter than 8th magnitude at a typical spot size, which ranges from 3 to 8. The exact nature of the noise will change between the flight detector and SLATE; the former will be almost purely Poissonian, while the latter includes Poisson, readout, and dark noise. It would be in principle possible to independently characterize the different SLATE detector noise sources, and their combined distributions, but we opted to use the empirical SNR instead. Some of the variation between our results may be due to this statistical effect.

3.5.3.3 Data acquisition and analysis

While the actuator encoder positions could be used for open loop positioning, there was a slight tilt to the actuators with respect to the optical axis, as well as motor backlash. Rather than try to calibrate these actuator imperfections, we determined the position of the starshade directly from

the camera images and ran an acquisition loop to go to the predetermined measurement positions, spaced apart by 30 cm (effective). To avoid errors in position due to optical and detector noise, during the position acquisition, the laser was turned to a very bright level. Then, the laser was turned down to the previously determined “science intensity” and 100 frames were taken at the science flux levels. After this, the beam launcher moved to the next grid point. See Figure 3-7.




Figure 3-7: Left: An image from the lab camera. Right: Matched image from the image library. The signal-to-noise ratio of the spot is approximately 7.5.

For each grid point, the 100 images were matched to the library, and the best match positions were collected and used to generate error ellipses shown in the next section.

3.5.4 Results

The formation flying results from the lab were fairly consistent with numerical expectations, with a factor of ~50% discrepancy between the simulations and lab results. This factor, obtained at much lower signal level than expected in flight, is still well within the error allowed by the milestone, and significantly below the errors used in the sensor model discussed in the next section. The SLATE results, compared to numerical expectations, are presented in Table 3-3. From the empirical data covariance matrix at each position, we generate error ellipses showing, 1, 2, and 3 σ contours. Since each data point has a different error, we present the median result (the median of all the ellipse axes) and worst result (the worst individual ellipse axis) in the Table. Plots of the results are presented in Figure 3-8.

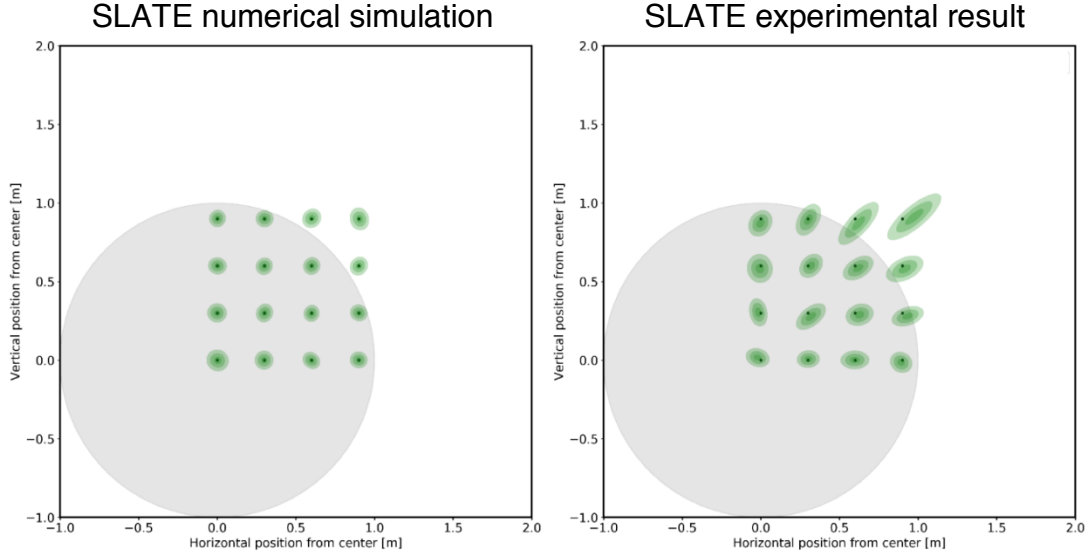


Figure 3-8: Comparison between accuracy of lab-generated and simulation-based models

Sim 3σ (worst)	Sim 3σ (median)	SLATE 3σ (worst)	SLATE 3σ (median)	Discrepancy (worst)	Discrepancy (median)
6.7 cm	4.0 cm	10.2 cm	6.2 cm	55%	52%

Table 3-3: Comparison between accuracy of lab-generated and simulation-based models

There are a few reasons for the worse lab performance than predicted. The primary reason is optical noise; that is, blobs of bright light forming structures in the shadow that would not be there in the CGI flight optical system (see 2.3.1.1). This noise is created by light scattered by the imperfect optics. It is the dominant source of noise, and the extra power leads to both statistical and systematic errors. The statistical errors are due to the combination of Poisson, dark, and readout noise; the systematic errors are due to the matching algorithm biasing towards the scattered light structures rather than only to the central Arago spot. (These systematic errors are visible in Figure 3-8 as slight shifts in the midpoint of the error ellipses compared to the setpoints).

Systematic errors in the camera also contribute. The dark level in the camera drifts continuously, and while we always took a background frame before a science frame, the noise on top of these frames can appear as a changing noise gradient from one side of the detector to the other. Another issue was flat-field correction: we did not solve for a flat-field on the camera and thus differences in per-pixel gain can create a spatially dependent systematic error signal. These errors are not expected to be present in the flight detector, which will be very stable and well-characterized.

3.6 Results and discussion

In this section, we evaluated the applicability of a pupil imager as a sensor for formation flying the Starshade with the WFIRST telescope. First, we determined the radiometric budget for the sensor using realistic target star and optical efficiency measurements. We used this budget to evaluate an analytic model of sensor performance. Furthermore, we implemented a custom optical propagation code to simulate the pupil images including propagation through the various optical

surfaces. We determined the accuracy to which the position of the Starshade could be recovered using an image matching approach, where the image from the pupil camera is matched against a library of precalculated model images of known offsets. We also built a lab demonstration that created flight-like optical conditions, matching Arago spot size, signal-to-noise ratio, and detector pixel scale at similar Fresnel number.

The simulations and lab experiments both used incident fluxes significantly fainter than those expected from target stars, yet both demonstrated that this sensing approach is well-matched to the task of determining the position of the Starshade to ≤ 30 cm (3σ) in one second. The analytic calculations and numerical optical propagation experiments beat the 30cm requirement by a factor of ~ 3 , with stars at least 10 times fainter than the faintest target star. The lab demonstrations were consistent with the numerical simulations to within a factor of ~ 2 .

The results here are fairly robust to changes to the CGI system efficiency or target list. Factors of a few decrease in incident fluxes would mainly degrade the blue science band performance, and it is trivial to increase the exposure time on the sensor to compensate, as the 600-800 seconds per trajectory are oversampled by the 1 second measurement cadence and modest < 1 cm/s speeds of the Starshade in the deadband (see next section). The main risk is a significant modification of the Starshade transmission function. The sensing signal depends on the out-of-band light, provided at the wings of the transmission function. Without taking formation flying into consideration, it would be possible to design a Starshade with adequate science performance but poor formation flying performance.

4 Formation flying

4.1 Formation flying approach: **overview**

The purpose of this section is to demonstrate that the shear sensing approach presented in the previous section allows controlling the lateral position of a Starshade in representative and stressing conditions within a 1m-radius disk normal to and centered on the telescope-to-target-star line. Stressing conditions are understood as conditions that are at least as conservative as the worst-case conditions expected for a WFIRST Rendezvous mission concept. The formation flying framework demonstrated here also applies directly to HabEx, as discussed briefly in Section 4.5.3.

The formation flying control performance requirement is met by commanding Starshade thrusters to counteract the relative lateral acceleration experienced by the Starshade and the telescope due to differential gravity acceleration and solar radiation pressure, thus ensuring the Starshade lateral offset remains inside the 1-m control region. The onboard estimation of the lateral offset is based on the sensing approach described in Section 3.

The simulations presented in this section include the orbital dynamics of both a telescope spacecraft and a Starshade spacecraft, as well as the (prescribed) attitude for the spinning Starshade spacecraft. Both spacecraft are in a halo orbit around the Sun-Earth L_2 Lagrange point and are subject to the gravitational acceleration due to the Sun, Earth, Moon, and solar system planets. The Starshade is in addition subject to solar radiation pressure.

As the shear sensor model from Section 3 vastly outperforms the requirement of 30cm accuracy (even for stars 10x fainter than the faintest star on the target list), the formation flying simulations are based on an even more conservative model, whereby the standard deviation of the shear position measurement error is no less than 10cm for all shear positions (i.e. 30cm, 3σ). See Section 4.2.4 for further details.

This simulation environment was used to test a set of proposed formation flying algorithms, also mostly developed specifically for S5. These algorithms include a Kalman filter that uses range and shear measurements to estimate the relative acceleration, velocity, and position between the two spacecraft. As thruster firings interrupt science observations, the control algorithms were also developed to maximize the drift time between burns.

Through Monte Carlo simulations, it was demonstrated that the shear sensing approach allows successful control of the relative lateral position of the Starshade to within 1 meter. This was achieved across all varied parameters, in stressing operational conditions, and even with the extremely conservative sensor performance assumed. It was also shown that thruster firings are only required every 10-15 minutes, thus providing high observational efficiency for Starshade science.

4.2 Simulation setup

4.2.1 Simulation environment overview

The MATLAB-based environment presented here propagates the position of a telescope and a Starshade spacecraft from initial conditions in three-dimensional space, under the gravitational influence of Earth, Sun, Moon, solar system planets, and solar radiation pressure.

4.2.1.1 *Starshade spacecraft*

The Starshade is composed of a 10-meter-diameter disk with 24 petals of 10.66 m² area each, distributed around the disk circumference. The Starshade spacecraft (including the bus and the Starshade) dry mass is assumed to be 1200 kg, and only 100 kg of propellant is assumed to be nominally on board (i.e. close to the dry mass for conservatism, as the mass uncertainty will represent a larger percentage of the total mass). A 2% propellant mass knowledge uncertainty is standard for typical spacecraft. Based on an assumption of 2500 kg maximum propellant mass, 50 kg (3σ) mass uncertainty was assumed on board (see Section 6 for a summary of parameters and uncertainties considered).

4.2.1.2 *Orbital dynamics*

The orbital propagation uses a constant time-step 5th order Runge-Kutta ODE solver. The ephemerides of the bodies were obtained from JPL's SPICE [17] library and validated against MONTE simulations (JP's high-fidelity mission design tool, validated operationally on several NASA deep space missions, including Cassini, Mars Science Laboratory, Juno [18]). For starting conditions similar to the ones considered in this report, the magnitude of the relative acceleration error over 24 hours was found to be of the order of 10^{-10} m/s², corresponding to about 0.001% of the relative acceleration vector magnitude. This negligible error is expected to be due to minor differences in the integrators or gravitational parameters.

4.2.1.3 *Solar radiation pressure*

A standard solar radiation pressure (SRP) model was implemented, whereby the total SRP acceleration is decomposed into absorptive, specular reflective, and diffusive reflective components. The area of the Sun-facing side of the Starshade is assumed to be covered in black Kapton, and the surface was therefore assumed to be 92% absorptive, 4% specularly reflective and 4% diffusively reflective. SRP acceleration is considered to be negligible for WFIRST due to the comparatively negligible surface area.

4.2.1.4 *Attitude dynamics*

As attitude control for a spinning spacecraft is not considered new technology, it is only included here to the extent that it affects the formation position control: a worst-case prescribed attitude is imposed on the Starshade, whereby the Starshade has a constant spin rate of 2°/s (1/3 rpm) around the Starshade axis, with this spin *axis* freely precessing at a constant rate of about 4°/s around the Starshade-to-telescope axis with a constant offset angle of 1° (no nutation). The offset in the spin axis corresponds to the current pointing requirement of $\pm 1^\circ$ (3σ) for the Starshade [13]. Modelling the Starshade spin in this manner demonstrates that the control algorithms – in particular the

thruster on-time allocator – are robust to relatively fast changes in attitude in the presence of measurement and thruster delays.

This prescribed attitude is generated by propagating the torque-free attitude dynamics of the Starshade throughout the simulation, with initial conditions chosen to provide the desired rotation and precession behavior. The attitude propagation is based on the same constant time-step 5th order Runge-Kutta ODE solver as the orbit dynamics and was validated to machine precision.

4.2.2 Initial conditions

The driving disturbance for formation flying is the relative lateral acceleration experienced by the two spacecraft, due to gravity and solar radiation pressure. For the simulations presented here, the initial positions of WFIRST and Starshade were chosen to maximize this disturbance, given the expected WFIRST and Starshade operational constraints.

An analysis was performed to characterize how the relative lateral acceleration is affected by WFIRST's initial positions on its L₂ trajectory (as provided by WFIRST in the form of preliminary ephemeris data), as well as the Starshade's position relative to WFIRST.

The worst-case relative lateral acceleration was found in portions of the WFIRST trajectory that are the closest to Earth. The worst-case Starshade position was found to correspond to large WFIRST-Starshade distances (the maximum range for a WFIRST rendezvous mission is 38.8Mm [13]) and Starshade-telescope-Earth angles (referred to simply as “Earth angle” here) around 40°-45°.

The identified worst-case initial conditions are illustrated in Figure 4-1, where the range is 38.8Mm (largest WFIRST rendezvous mission range), the Earth angle is 43.5°. This initial condition corresponds to an initial relative lateral acceleration of 15.2 μm/s², or 1.55 μg. This corresponds to a Starshade-telescope-Sun angle of 46.5°, which is lower than the minimum acceptable angle of 54° for a WFIRST rendezvous mission. As this corresponds to a larger formation flying disturbance however, it was chosen for added conservatism.

This analysis also showed that for conditions where the relative lateral acceleration is large, the lateral acceleration's rate of change (both in magnitude and direction) is negligible: the relative lateral acceleration magnitude rate of change is of the order of 7×10⁻¹⁴ m/s³ and the angular rate of change of the acceleration direction is of the order of 0.1 arcseconds/s, i.e. about 3×10⁻⁵ °/s.

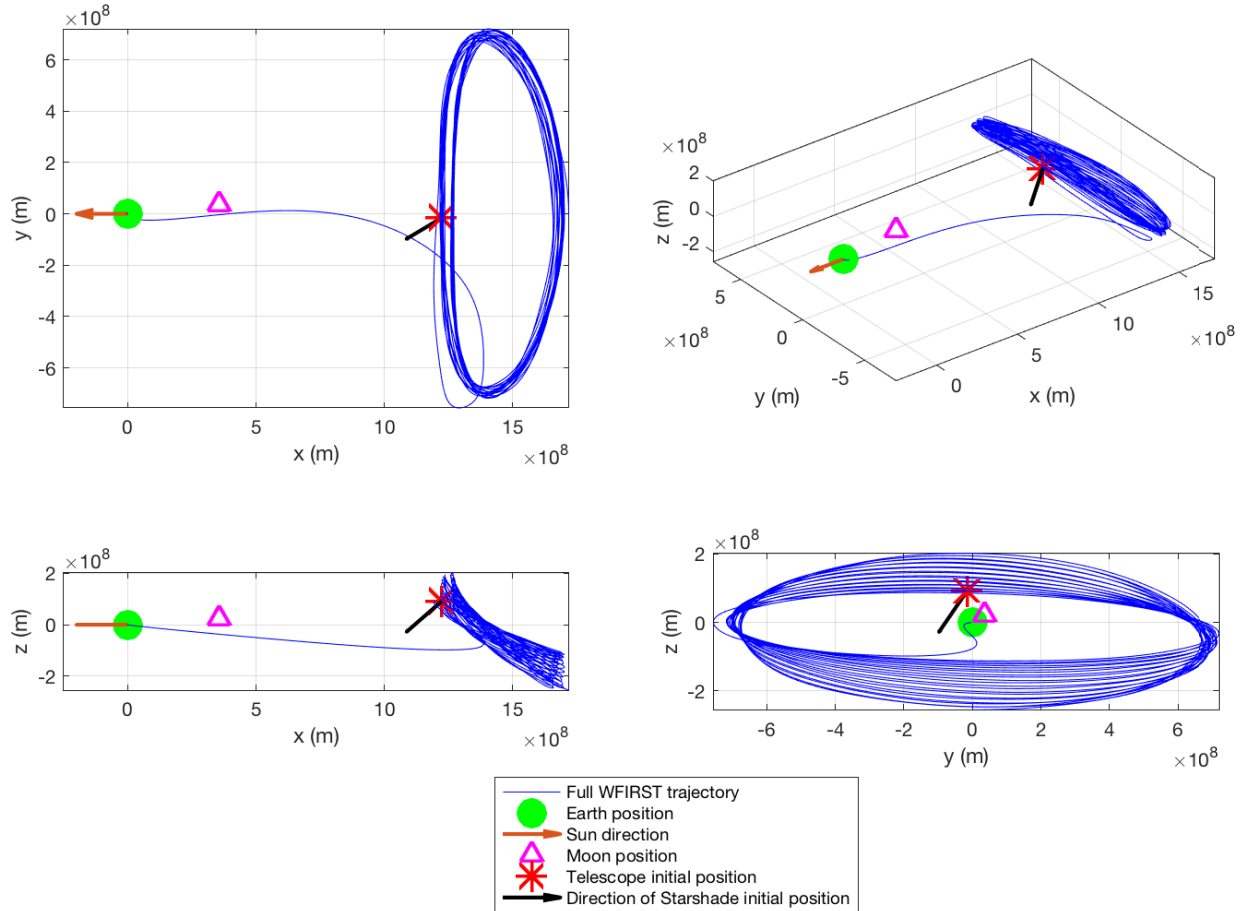


Figure 4-1: Reference initial positions of the telescope and Starshade for the formation flying simulations

4.2.3 Thruster models

A 16-thruster configuration was chosen for the Starshade, as shown in Figure 4-2. The thrusters are all nominally oriented at a 45° angle to the Starshade spin axis (x axis) and their projection onto the plane normal to the spin axis (y-z plane) is either oriented in the $\pm y$ or $\pm z$ direction. The thrusters assumed to be used on board are flight qualified Aerojet Rocketdyne R-6D bipropellant thrusters [19], with nominal thrust of 22 N, allowing short burns relative to the Starshade spin rate⁵.

The thruster performance assumed in the simulations is as follows: *known* thrust direction and magnitude biases of 1° (3σ) and 1% (3σ) were assumed respectively for each thruster. Additionally, further *unknown* biases of 0.75° (3σ) and 2% (3σ) were added. Finally, for each

⁵ In the technology development plan [1], MR-103M thrusters, i.e. 1N-class hydrazine monopropellant thrusters are mentioned. It was found that the larger thrusters considered here were preferable as they allowed for a simpler thruster allocation design.

individual thruster burn, *unknown* errors of 0.1° (3σ) and 1% (3σ) were applied. Each thruster was assumed to have a minimum on time of 5ms, corresponding to a nominal minimum impulse per thruster of $84.6\mu\text{m/s}$ (for a 1300kg spacecraft mass). The quantization level of the thrust on time was set to 0.5ms. A 1 second delay was applied between any command and the start of the corresponding burn.

The thruster performance was chosen based on a combination of available thruster data, Europa Clipper or Mars Science Laboratory performance, and rules of thumb from SMEs at JPL.




Figure 4-2: Three schematic views of the thruster configuration, shown by thruster plume directions

4.2.4 Sensor models

The inputs to the estimation and control algorithms described in Section 4.3 are range, attitude, and shear measurements, generated using the sensor models described in this section.

The range sensor is based on a standard S-band system with an accuracy of ± 500 m (3σ), as specified in the Starshade interface requirements document [13].

Starshade attitude sensing was implemented as a truth + noise model, with a 20 arcsecond magnitude (3σ) attitude error and a $0.01^\circ/\text{s}$ magnitude (3σ) angular rate error. This is consistent with standard performance from flight-proven attitude sensors.

The shear sensor is based on the sensing scheme presented in Section 3. The performance of the sensor depends on the shear position of the Starshade with respect to the telescope-star axis. As a result, for any shear position, the sensor performance is expressed as a covariance matrix.

As the sensor model from Section 3 was found to outperform its requirement of measuring the shear position to 30 cm (3σ) and in order to conservatively demonstrate successful formation flying within the defined control region to this requirement, the original allocation of 30 cm shear measurement accuracy was maintained by *scaling* the error of the sensor model from Section 3

such that for all shear positions, no measurement has a standard deviation smaller than 10 cm (i.e. 30 cm, 3σ) in any direction. This corresponds to multiplying the standard deviations for all shear positions in the sensor model from Section 3 for the worst-performing band (the “blue” band) and for a star of magnitude 8 by a factor of 3.67. This approach thus maintains the spatial variation of the sensor accuracy, but with extremely conservative performance. This is illustrated in Figure 4-3.

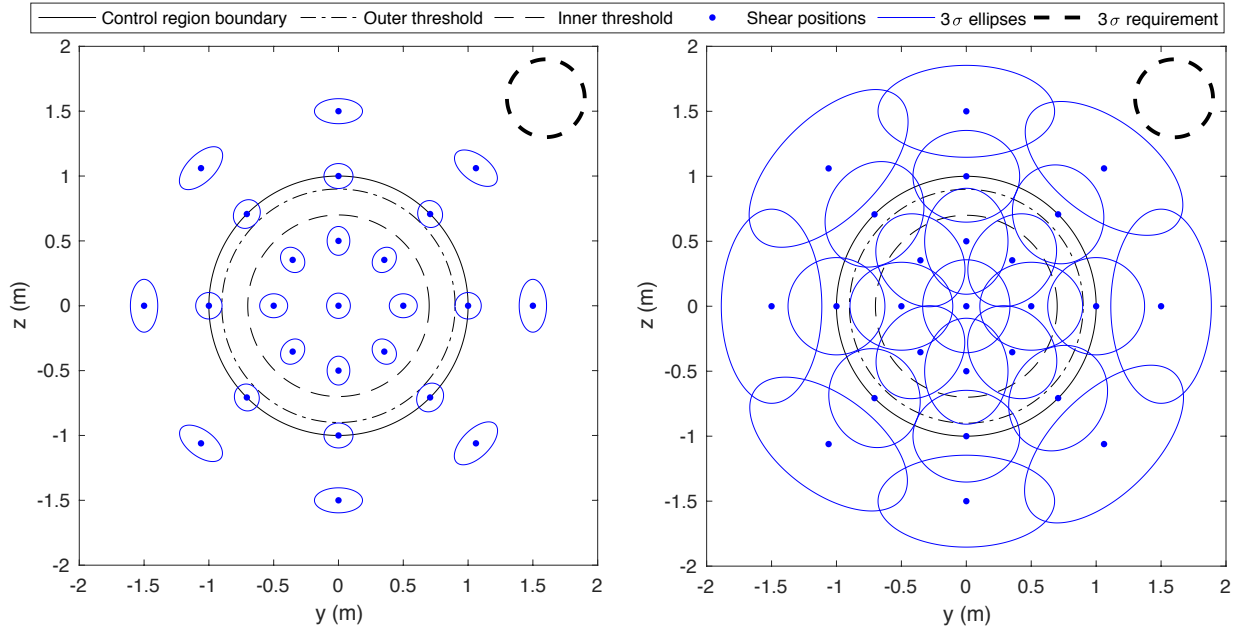


Figure 4-3: Shear sensor model: **Left:** original model for 8th magnitude stars in the “blue” band. **Right:** Scaled, extremely conservative model used in the formation flying simulations.

For both range and shear sensors, further errors were added to the sensor models: while measurements are expected to be provided every second, the average measurement rate is assumed have an error of 10ms (3σ) to account for clock drift. This drift is assumed to be unknown on board and thus results in a constant time-tagging error bias due to constant error in the integration time compensation on board). A 3 ms (3σ) variation in the actual measurement time variation was added for each measurement and a further 100 ms (3σ) time-tagging error was assumed, as only ground-based clock synchronization between the two spacecraft is currently guaranteed (as opposed to more accurate clock correlation through the S-band). All measurements are assumed to be delayed by half a measurement period (about 0.5 seconds) to account for measurement integration time, and a further 1 second delay was added to account for other delays in the system.

4.3 Prototype control algorithms

The purpose of the control algorithms presented here is to process raw measurements generated using the sensor models described above and to trigger thruster firings in order to meet the control requirement.

4.3.1 Estimation

A Kalman filter was implemented to estimate the three-dimensional relative position, velocity, and acceleration between the two spacecraft. All three variables are required for the control algorithms (see Section 4.3.2). The Kalman filter is only fed with delayed and noisy sensor measurements and their associated time-tags, and is designed based on a three-dimensional constant acceleration model. This simple model choice is justified since it was found that the rate of change of the relative lateral acceleration is negligible at the time-scales considered in the present simulations (as mentioned in Section 4.2.2).

Kalman filter design is centered around the tuning of a process noise matrix and a measurement noise matrix. The process noise matrix is tuned differently depending on whether thrusters are firing (larger process noise is assumed when thrusters are firing). This avoids unwanted transients in the velocity and acceleration estimation. For conservatism, no measurements are used during thruster firings. The measurement covariance matrix tuning is dependent on the current relative lateral position of the Starshade and is based on the accuracy of the longitudinal sensor model and the (extremely conservative) lateral sensor model.

The attitude estimation algorithm propagates the on-board attitude estimate linearly to the desired time, based on the unfiltered attitude state measurement. This is required to ensure the attitude state at the required time step is fed to the thruster on-time allocator.

4.3.2 Control

Although the identified technology gap only focuses on controlling lateral position to within a 1m-radius disk, the task undertaken in the present formation flying simulations is more challenging. First, although this is not considered a technology challenge, the relative longitudinal position must also be controlled to within a specified control region to maintain the telescope in the Starshade’s “dark shadow”. A large control region of ± 250 km is considered in this case. Second, thruster firings are expected to generate plumes that can scatter sunlight and be bright enough to degrade scientific observations. Therefore, science must be interrupted when thruster fire so the controller should be designed such that thruster firings take place as rarely as possible and ideally in a predictable manner.

4.3.2.1 Lateral control

To maximize drift time between thruster firings, a new optimal deadbanding algorithm was developed. Typically, this type of control problem is solved using single-axis deadbanding. However, here the deadband is not only two-dimensional, but also disk-shaped. As a result, highly suboptimal results can be expected if two uncoordinated single-axis deadbanding controllers are used. This motivated the development of the new optimal disk-deadbanding algorithm proposed here [20]. With this algorithm thrusters are only fired when the relative lateral position reaches the edge of a disk-shaped region. The optimal Δv leads to a trajectory that maximizes the drift time until the next burn is commanded, while maintaining the trajectory within the lateral control region. After the first two burns and in the absence of any disturbances, the optimal trajectory bounces

from the very “bottom” to the very “top” of the diameter of the disk that is aligned with the relative acceleration vector. This provides the longest possible drift time between thruster firings (see left part of Figure 4-4).

In practice, estimation and thruster errors often lead to slightly suboptimal yet acceptable trajectories. To ensure that thruster firings are not triggered unnecessarily but that the trajectory always remains within the 1 m-radius disk, an inner and an outer threshold were defined as follows:

- 1) **Outer threshold:** An optimal thruster firing is commanded if:
 - a) the estimated relative lateral position is outside the outer circle, and
 - b) the estimated relative lateral velocity is directed outwards.

Here, the radius of the outer circle was tuned to 90% of the actual deadband radius to ensure that a burn is always triggered before the trajectory reaches the boundary of the 1m-radius control region.

- 2) **Inner threshold:** an optimal thruster firing is commanded if:
 - a) the estimated relative lateral position is outside the inner circle, and
 - b) the estimated relative lateral velocity is directed outwards, and
 - c) the radial component of the estimated relative lateral acceleration is directed outwards.

The purpose of this double threshold approach is to allow for small discrepancies and overshoots (larger than the radius of the first threshold but smaller than the radius of the second threshold) in the lateral trajectory without triggering a thruster firing. Here, the radius of the inner circle was tuned to 70% of the actual deadband radius (see right part of Figure 4-4, where it is shown that thruster firings are only triggered for large overshoots that may lead to a control region boundary crossing if not corrected).

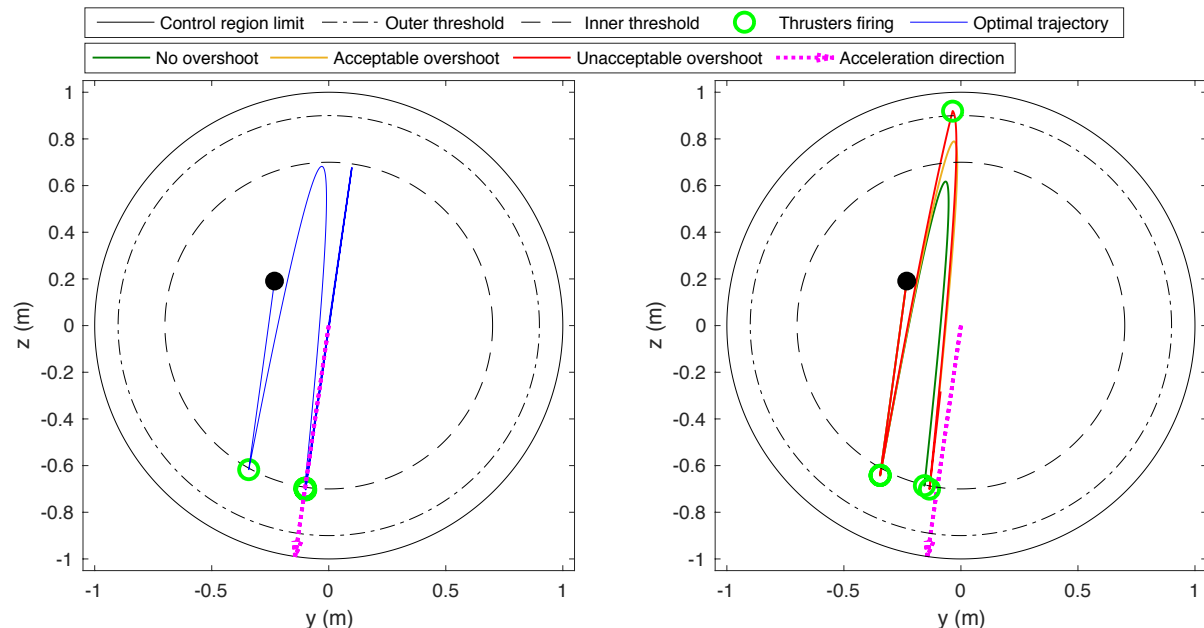


Figure 4-4: Illustration of the lateral control approach. **Left:** Optimal disk deadbanding with no errors or disturbances. **Right:** Illustration of the double threshold control approach, allowing small overshoots but ensuring the trajectory does not cross the control region boundary.

4.3.2.2 Longitudinal control

Controlling relative longitudinal position to within ± 250 km with a ± 500 m-accurate sensor can be done in several ways and is not considered a technology challenge. The purpose of the longitudinal control algorithm proposed here is therefore to show that longitudinal control can be performed without degrading the performance of the lateral control, while ensuring that the telescope-to-Starshade range remains within the longitudinal control region.

The longitudinal controller was designed as follows: first, longitudinal burns are only commanded at the same time as a lateral burn to avoid shortening drift time and thus lowering observational efficiency. Second, large longitudinal burns would have residuals that could significantly disturb the lateral trajectory. As a result, the magnitude of these “synchronized” longitudinal burns is capped at 50% of the lateral burn magnitude. Third, longitudinal burns are only commanded if the relative longitudinal velocity estimate is above a certain threshold (0.1m/s). This is required because the range sensor is only accurate to ± 500 m, so the relative longitudinal velocity estimate can be in the wrong direction for small relative velocity magnitudes.

Note that this “velocity damping” approach was not designed to robustly maintain the relative longitudinal position in a tight control region. It is only intended to significantly increase the amount time it takes for the edge of the ± 250 km longitudinal control region to be reached. Note also that in most scenarios, due to the very large ± 250 km longitudinal control region size, no longitudinal control would actually be required for days depending on the relative longitudinal position, velocity, and acceleration at the start of the observation period.

4.3.2.3 Attitude control

Closed-loop attitude control was not included in the simulations. As the spacecraft is spin-stabilized, only small and infrequent Δv s would be necessary to maintain the attitude of the Starshade within $\pm 1^\circ$ (3σ). It is therefore expected that it would be possible to synchronize these with the lateral station-keeping burns. Even if an additional attitude adjustment burn is occasionally required between lateral station-keeping burns, this would only affect the drift time between burns and not the stated TRL5 objective, i.e. maintaining the lateral position offset under 1-m.

4.3.3 Thrust allocation

The thrust allocator translates Δv commands into thruster on-times that realize as well as possible the desired torque-free impulse. In the present simulations, the QPCAP (Quadratic Programming algorithm for the Control Allocation Problem) thrust allocation algorithm was used. This algorithm [21] was developed internally at JPL and flew on the Mars Science Laboratory (MSL) descent stage and has been used extensively in formation flying studies and demonstrations (e.g. [22]).

Although QPCAP assumes the attitude is constant throughout the burn execution, the Starshade is spinning and the thruster firings are executed with a significant delay. As a result, the attitude fed

to QPCAP must be a propagated attitude, corresponding to when the burn is expected to actually begin. A linear attitude propagation algorithm was therefore implemented for this purpose.

4.4 Monte Carlo simulations

A set of Monte Carlo simulations were run using the setup described in the sections above in order to show that in realistic and stressing conditions, the lateral sensing approach described in this document can be used to robustly control lateral position of a Starshade to within a 1m-radius disk, while providing high observational efficiency for science. Assuming the Starshade position is successfully maintained within the control region, the main parameter of interest is therefore drift time between individual thruster firings.

Each Monte Carlo run is propagated for 6 hours from an initial relative position and velocity during which the on-board algorithms autonomously control the formation position. The parameters varied between runs are summarized in the tables in Section 6, including mass uncertainty, thruster and sensor errors and biases, initial position with respect to the center of the control region, initial velocity, and initial estimation error. Several parameters are also distributed randomly *within* each Monte Carlo run – e.g. individual sensor measurements, associated measurement time and time-tag, burn-dependent magnitude and direction errors.

For a given set of initial conditions and biases, a finite simulation time is required in order to obtain a converged estimate of the average drift time between burns. With a simulation time of 6 hours, the mean drift time was found to converge to within approximately 5% of the long-term average. Since all biases are fixed within a single simulation, it is not surprising that the statistics of the drift time do not require a very long simulation time to converge.

The total number of Monte Carlo runs required to obtain overall converged statistics was also investigated. It was found that 60 Monte Carlo runs (of 6 hours of formation flying each) provided sufficient confidence in the results and allowed identifying potential problematic scenarios and corner cases. This is illustrated in Figure 4-5, which shows that the average of the per-simulation mean drift time is converged to well below 1% with 60 simulations.

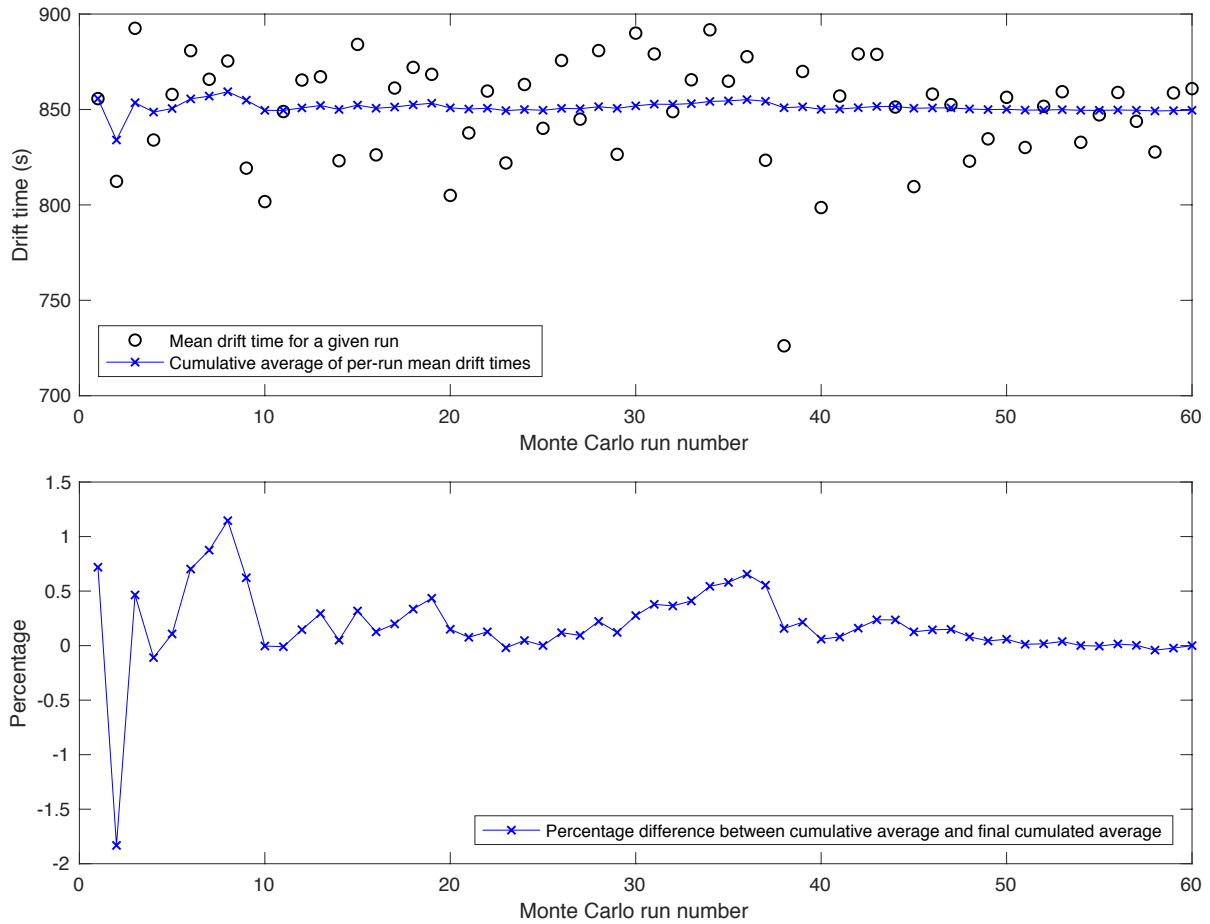


Figure 4-5: Convergence of mean drift time over the 60 Monte Carlo simulations of 6h each

4.5 Results and discussion

4.5.1 Typical formation flying behavior

Figure 4-6 shows a typical deadbanding trajectory over one hour. As expected, the lateral trajectory remains inside the inner threshold, and from the second burn onwards, the trajectory remains close to the optimal trajectory, i.e. aligned with the diameter of the control region that is aligned with the relative lateral acceleration.

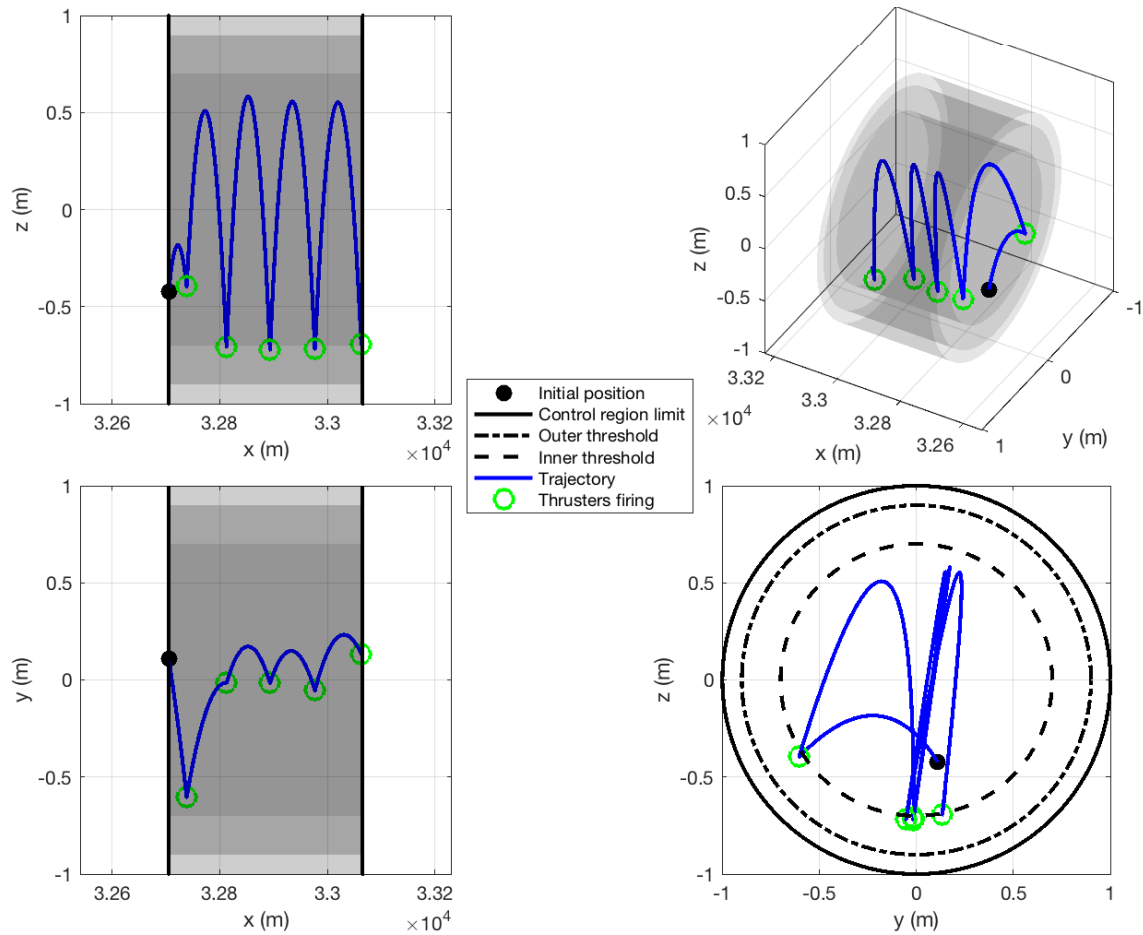


Figure 4-6: Example of a typical three-dimensional trajectory over 1h, showing the two thresholds and the control region boundary.

Figure 4-7 shows the behavior of the longitudinal control algorithm. In this case, longitudinal burns are triggered at two instances where the estimated longitudinal velocity is above the threshold value of 0.1 m/s at the time when a lateral burn is triggered. Around $t=2755$ s, the real value of the longitudinal velocity is above the threshold but the estimated value is just below, so no burn is triggered.




Figure 4-7: Example behavior of the longitudinal estimation and control over 1h

Figure 4-8 illustrates the behavior of the Kalman filter. After a thruster firing, the position and velocity estimation converge in about 200 seconds with the current Kalman filter tuning. As required, the acceleration estimation is not disturbed by the thruster firings.

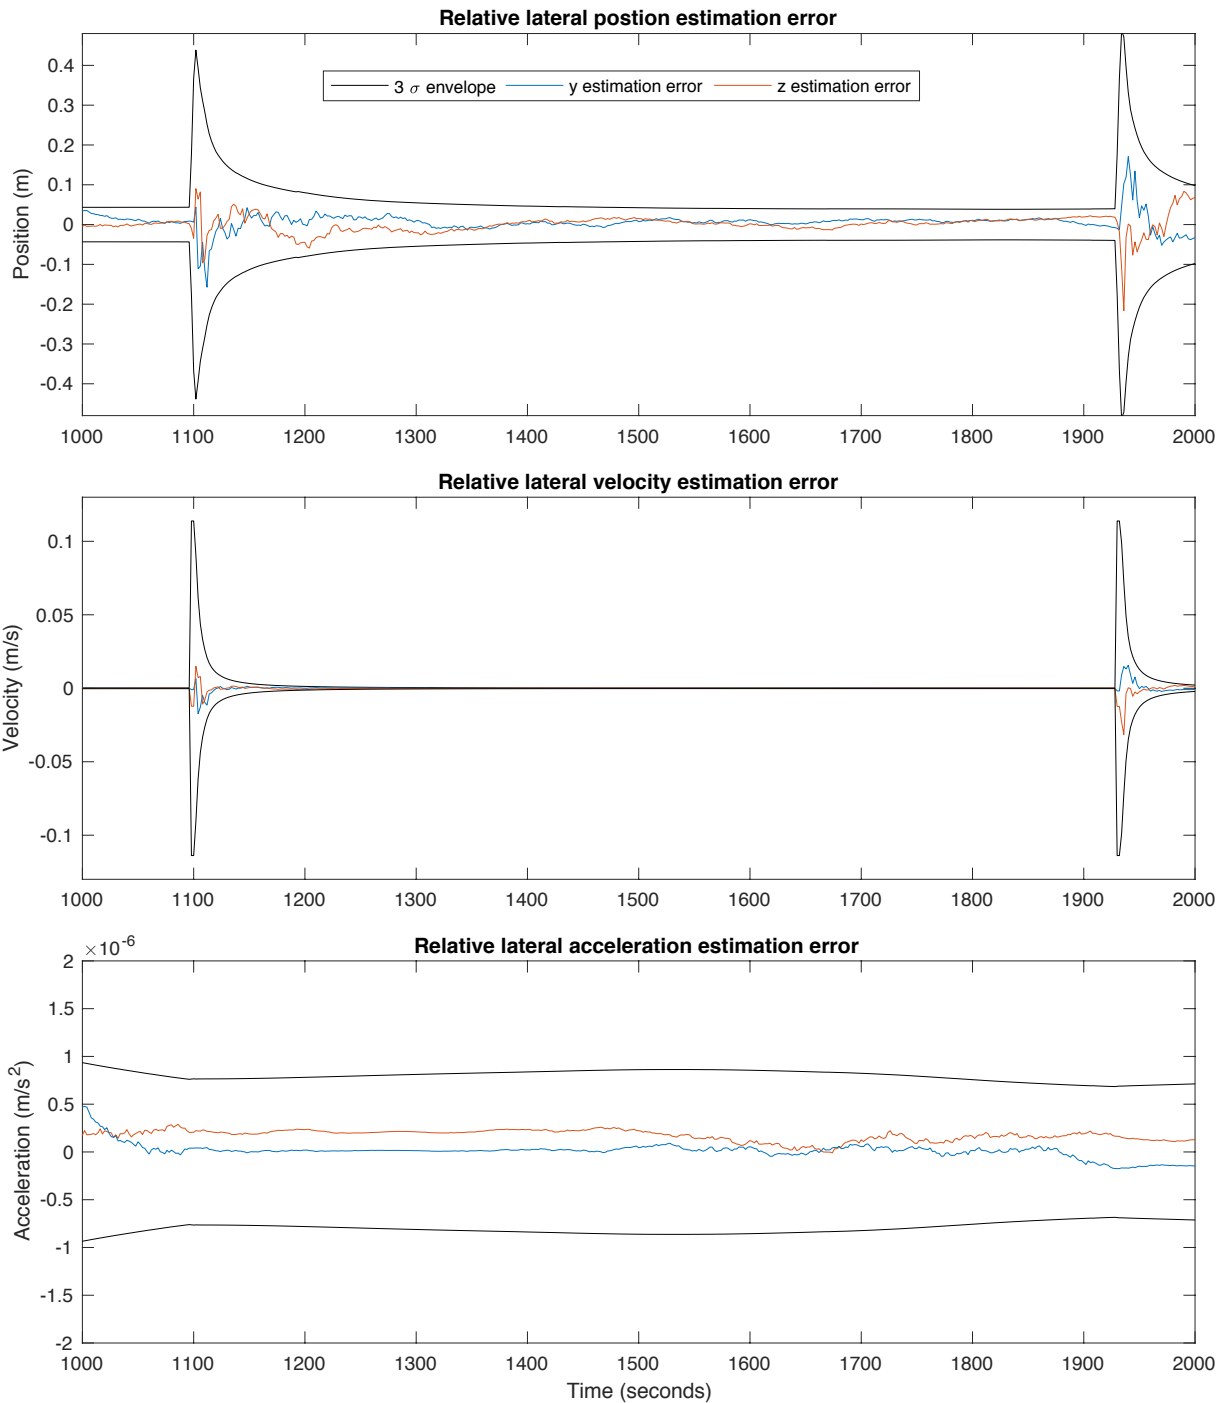


Figure 4-8: Example behavior of the lateral components of the estimator between two burns.

4.5.2 Overall Monte Carlo simulation results

Figure 4-9 shows the histogram and the cumulative distribution function of drift time, compared to ideal drift times, expected for optimal deadbanding along the deadband diameter that is aligned with the relative lateral acceleration vector (with no disturbances). The first drift periods when the trajectory is converging to a steady-state are ignored in the statistics. In all cases considered, the lateral offset of the Starshade was found to remain within both the lateral and longitudinal control regions. The mean drift time across all simulations was found to be 850 seconds. This value is marginally lower than the ideal drift time for a perfect deadbanding within the 2nd threshold given a relative acceleration of $15.2 \mu\text{m/s}^2$, i.e. 858 seconds. This is likely due to the fact that although small thrust *magnitude* errors compared to the ideal burn are equally likely to increase and decrease the drift time (and thus are not expected to introduce a shift in the mean drift time), all thrust *direction* errors lead to a reduction in the drift time. Similarly, all trajectories *starting* away from the optimal location along the inner threshold boundary will also lead to reduced drift times.

Figure 4-9 also shows that having an inner threshold at 70% of the actual deadband limit (i.e. a 0.7 m instead of 1 m nominal deadbanding radius) reduces the ideal drift time from 1026 seconds (17 minutes, 6 seconds) to 858 seconds (14 minutes, 18 seconds). It also shows that in very rare cases (<1% of the burns), the initial burn error was large enough for the outer threshold to be reached thus triggering a corrective burn and approximately halving the drift time. This behavior was found to take place when the on-board estimate of the mass is much higher than the actual mass, leading to systematically excessively large Δv commands. In the present set of simulations, this occurred (several times) for one of the simulations which had a mass estimation error (bias) of 41 kg (i.e. 2.5σ). This indicates that a trade-off must therefore be performed: increasing the threshold radii will lead to an increased nominal drift time but a higher probability that corrective burns will be occasionally required, each time compromising one or two science frames. Note that once in orbit, if it is found that correction burns are often required, then the on-board mass estimate could simply be adjusted.

Figure 4-10 is the Q-Q plot of the drift times over the entire Monte Carlo campaign. It shows that from about -2σ to $+1.5\sigma$, the drift time distribution over the entire Monte Carlo campaign is approximately normally distributed. The rare cases where correction burns halve the drift time appear distinctly in the 300 to 500 seconds and -3.5σ to -2.5σ range. Furthermore, the suboptimal initial positions at the start of drift periods and non-ideal burn directions that lead to a small reduction in the average drift time also add some skewness to the distribution, as indicated by the slightly lower slope above $+1.5\sigma$.

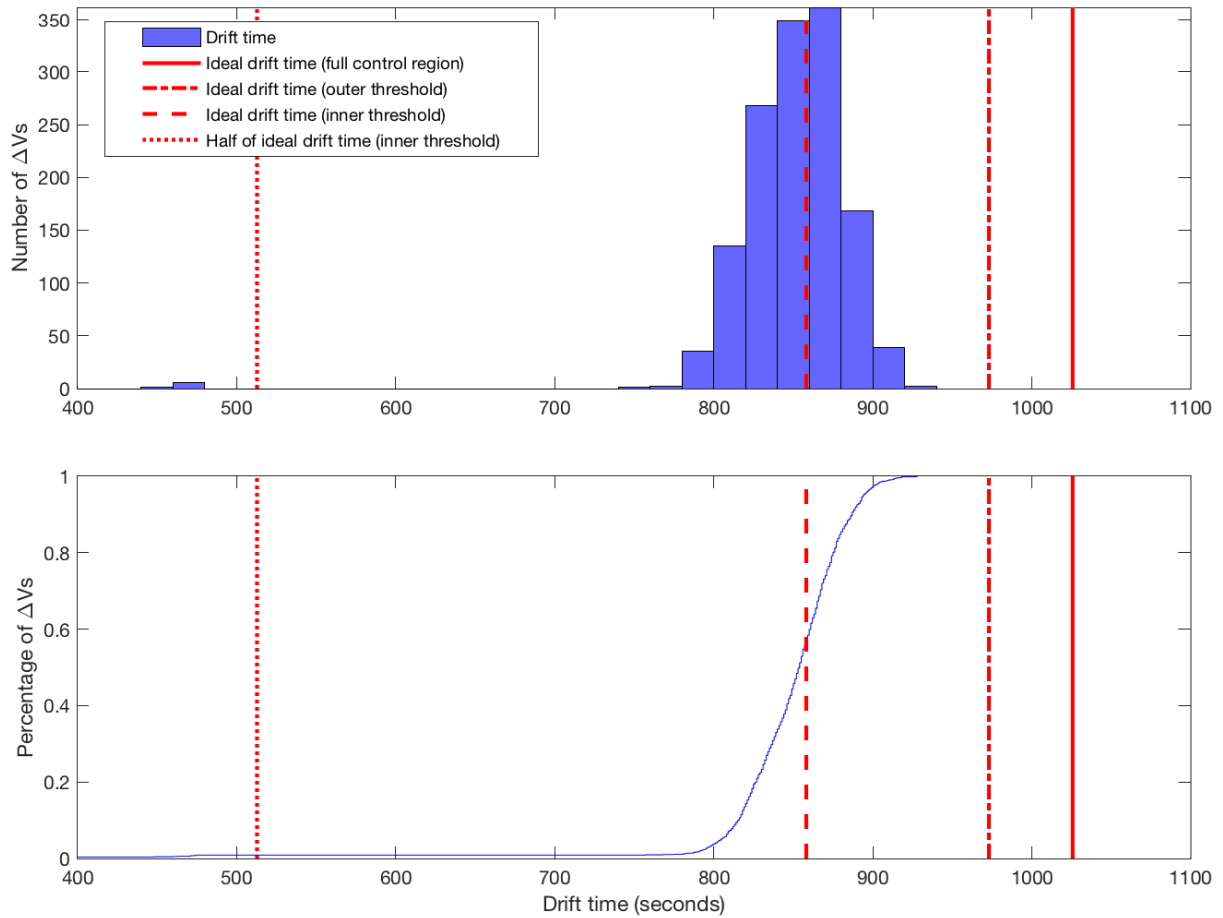


Figure 4-9: Summary of results of the Monte Carlo campaign. **Top:** Histogram of drift times. **Bottom:** Cumulative distribution function of drift times




Figure 4-10: Q-Q plot of drift times over entire Monte Carlo campaign

4.5.3 Relevance of results to HabEx

Although the design of HabEx is not finalized, a preliminary analysis was performed to test the applicability of the formation flying framework proposed here to HabEx by varying the key parameters that are expected to affect the formation flying performance.

HabEx is also expected to be in a halo orbit at the Earth-Sun L_2 point, so the WFIRST trajectory was kept here, and the telescope-to-Starshade distance was increased 76.6 Mm. Based on the HabEx Interim Report [23], a dry mass of 6394 kg was considered and the 2% propellant mass uncertainty assumption was maintained (based in this case on a maximum propellant mass of 7007 kg). The Starshade surface area was increased by a factor of 4 compared to the WFIRST Starshade. The rest of the parameters were not changed.

With this setup, the worst-case initial conditions analysis was repeated. This led to qualitatively similar results: the worst-case telescope initial position is close to the Sun-Earth line, it minimizes the distance to the Earth and Sun, and the Earth-angle is around 42.5° . The worst-case relative lateral acceleration value was found to be $31.5 \mu\text{m/s}^2$, or $3.21 \mu\text{g}$. The ideal drift time for perfect deadbanding is thus correspondingly smaller: 710 seconds (11 minutes, 50 seconds) for the full deadband and 595 seconds (9 minutes, 55 seconds) with a 70% inner threshold.

The full 60-run Monte Carlo campaign was then repeated with the HabEx parameters. The drift time between burns was found consistent with the relative acceleration and the overall behavior of the formation was found to be qualitatively similar. This confirms that the deadbanding approach proposed here is likely to extend to HabEx's larger relative lateral acceleration values.

5 Conclusions

The Starshade mission requires formation flying the telescope and Starshade to within a 1-meter control radius, leading to a sensing accuracy requirement of 30 cm in 1 second. This work developed a sensor model and control framework that can achieve this level of performance with very conservative assumptions about target brightness, spacecraft performance, and space environment.

The sensor model is based on the Starshade shadow, which consists of a circular “Arago” spot in the center and radial patterns surrounding it. This shadow is very dark at the science wavelengths, but brightens by many orders of magnitude at wavelengths outside the science band. At these wavelengths, a pupil camera can detect this relatively bright signal, and the position of the Arago spot can be used to determine the relative shear of the Starshade with respect to the telescope.

Radiometric calculations of the amount of light at the camera using target star fluxes, Starshade suppression curves, and telescope optical efficiency were combined with analytic calculations and detailed numerical simulations to predict the sensor sensitivity. The analytic and numerical simulations were in good agreement and predicted performance easily exceeding the S5 requirement for stars at least ten times fainter than any target star. Laboratory experiments at similar signal levels also showed good agreement with predictions from numerical simulations, despite the presence of systematic errors not expected in the space environment.

This sensor scheme was then used in a dedicated high-fidelity formation-flying simulation environment, whereby it was shown that, even with extremely conservative sensor model performance assumptions and starting from worst-case initial conditions, it is possible to control the shear position of the Starshade spacecraft with respect to the telescope-to-star line to within a 1m-radius disk. This was achieved using a set of flight-like formation flying algorithms that command thruster firings to correct the drift induced by the relative acceleration experienced by the formation. These algorithms were developed to maximize the drift time between thruster firings, and were found to provide an approximately periodic drift time of about 14 minutes on average for worst-case acceleration conditions corresponding to a WFIRST rendezvous mission. A preliminary analysis indicated that this approach is likely to also be applicable to HabEx, where the larger relative acceleration leads to lower drift times of the order of 10 minutes. In both cases, the control approach provides high observational efficiency for exoplanet science.

This confirms that the Key Performance parameter from Section 1.2 has been demonstrated, thus reaching the corresponding S5 Milestone associated with the formation flying technology gap, and in turn confirming that the shear sensing approach considered in this report has reached TRL5.

6 Appendix

6.1 List of acronyms

ASTM	American Society for Testing and Materials
CCD	Charge-coupled device
CGI	Coronagraph Instrument
FLOPS	Floating-point operations per second
FWHM	Full-width at half-maximum
HCIT	High contrast imaging testbed
LOWFS	Low-order wavefront sensor
MONTE	Mission Analysis, Operations, and Navigation Toolkit Environment
MSL	Mars Science Laboratory
ODE	Ordinary differential equation
PSD	Power spectral density
QE	Quantum efficiency
QPCAP	Quadratic programming thrust allocation algorithm
RMS	Root-mean-square
RN	Readout noise
SLATE	Starshade Lateral Alignment Testbed
SME	Subject Matter Expert
SNR	Signal-to-noise ratio
SRP	Solar radiation pressure
TRL	Technology Readiness Level

6.2 Table of formation flying simulation parameters

Parameter	Known/ Unknown on board	Value	Units
Time			
Single Monte Carlo run final time	N/A	6	h
SPICE ephemeris time step	N/A	360	s
Attitude dynamics time step	N/A	0.5	s
Attitude control time step	Known	0.5	s
Position dynamics time step	N/A	1	s
Position control time step	Known	1	s
Spacecraft physical properties			
Starshade spacecraft dry mass	Known	1200	kg
Starshade propellant mass	Known	100	kg
Propellant mass uncertainty (3σ)	Unknown	50	kg
Starshade hub diameter	Known	10	m
Starshade petal area (24 petals)	Known	10.66	m ²
Starshade surface absorptivity	Known	92	%
Starshade surface specular reflectivity	Known	4	%
Starshade surface diffusive reflectivity	Known	4	%
X principal moment of inertia	Known	42390	kg.m ²
Y principal moment of inertia	Known	21289	kg.m ²
Z principal moment of inertia	Known	21289	kg.m ²

Table 6-1: Parameters used in formation flying simulations Part 1/3. Items highlighted in blue are kept constant within a given Monte Carlo run but varied from run to run. Items highlighted in yellow vary during each simulation.

Parameter	Known/ Unknown on board	Value	Units
Actuator models			
Nominal angle of thruster to Starshade spin axis	Known	45	°
Nominal individual thruster force	Known	22	N
Individual thruster thrust magnitude bias (3σ)	Known	1	%
Individual thruster thrust angle bias (3σ)	Known	1	°
Individual thruster thrust magnitude bias (3σ)	Unknown	2	%
Individual thruster thrust angle bias (3σ)	Unknown	0.75	°
Individual thruster thrust magnitude variation (each burn) (3σ)	Unknown	1	%
Individual thruster thrust angle variation (each burn) (3σ)	Unknown	0.1	°
Thruster minimum on time real value	Unknown	5	ms
Thruster minimum on time assumed on board	Known	5.25	ms
Thruster on-time quantization	Known	0.5	ms
Delay between command and start of thrust assumed on board	Known	1	s
Variation in delay between command and start of thrust (3σ)	Unknown	50	ms
Sensor models			
Best shear measurement accuracy across all shear positions (3σ)	Known	0.3	m
Range measurement accuracy (3σ)	Known	± 500	m
Attitude measurement accuracy (3σ)	Known	20	arcsec
Angular rate measurement accuracy (3σ)	Known	0.01	°/s
Nominal shear/range measurement sample time	Known	1	s
Shear/range measurement average sample time uncertainty (3σ)	Unknown	10	ms
Shear/range measurement sample time noise (3σ)	Unknown	3	ms
Shear/range measurement time-tagging uncertainty (3σ)	Unknown	100	ms
Shear/range measurement systematic delay	Unknown	1	s

Table 6-2: Parameters used in formation flying simulations Part 2/3. Items highlighted in blue are kept constant within a given Monte Carlo run but varied from run to run. Items highlighted in yellow vary during each simulation.

Parameter	Known/ Unknown on board	Value	Units
Position control			
Longitudinal Δv magnitude cap	Known	50	% of lateral Δv
Longitudinal Δv deadzone	Known	[-0.1 , 0.1]	m/s
Full longitudinal control region	Known	± 250	km
Lateral control region radius	Known	1	m
Lateral control region threshold radii	Known	[0.7 , 0.9]	m
Estimation			
Initial acceleration uncertainty per axis (3σ)	Known	3	$\mu\text{m/s}^2$
Initial longitudinal velocity estimation uncertainty (3σ)	Known	0.2	m/s
Initial lateral velocity estimation uncertainty per axis (3σ)	Known	0.3	mm/s
Initial longitudinal position estimation uncertainty (3σ)	Known	500	m
Initial lateral position estimation uncertainty per axis (3σ)	Known	0.3	m
Position dynamics			
Telescope - Starshade nominal distance	Unknown	38.8	Mm
Earth - Telescope - Starshade nominal angle	Unknown	43.5	$^\circ$
Initial longitudinal velocity uncertainty (3σ)	Unknown	0.2	m/s
Initial lateral velocity uncertainty per axis (3σ)	Unknown	7.5	mm/s
Initial longitudinal position uncertainty (3σ)	Unknown	150	km
Initial lateral position uncertainty per axis (3σ)	Unknown	0.7	m
Attitude dynamics			
Starshade spin rate	Unknown	2	$^\circ/\text{s}$
Starshade spin axis offset from Starshade-telescope axis	Unknown	1	$^\circ$
Starshade axis precession rate	Unknown	~ 4	$^\circ/\text{s}$

Table 6-3: Parameters used in formation flying simulations Part 3/3. Items highlighted in blue are kept constant within a given Monte Carlo run but varied from run to run. Items highlighted in yellow vary during each simulation.

6.3 Matching algorithm and computational steps

Assuming there are m images in the library, and n pixels per image:

- Sum raw image intensities SI n-1
- Calculate $mSi = n/SI$ to get inverse of mean intensity 1
- Multiply raw image intensities by mSi n
- Subtract each library image from normalized image n * m
- Square result for each pixel for each library image n * m
- Sum all squared differences for each library image (n-1) * m
- Find minimum error among all images m-1 with binary search

The total cost is $3nm+2n+1$

6.4 Comparison between SLATE testbed and flight expectation

Parameter	Flight expectation	SLATE testbed
Fresnel number	5-7	4.5
Light type	broadband starlight (50-100 nm filtered)	632 nm laser
Wavefront quality	~14nm wavefront error	>500 nm wavefront error
Beam apodization	None	Gaussian
Camera chip	e2v CCD201	SBIG KAF402-me
Camera read noise	2 electrons	40 electrons
Camera dark current	1.5e-4 electrons/pixel/sec	2 electrons/pixel/sec
Camera clock-induced charge	0.02 electrons	<1 electron
Camera flat field calibration	excellent	none
Arago spot FWHM	10 pixels /32x32 detector	10 pixels/ 32x32 pixels
Arago spot SNR	5/pixel in FWHM	5/pixel in FWHM

Table 6-4: Summary of the differences between the test setup and flight expectation. “Flight expectation” refers to current measurements that have been measured by the WFIRST CGI Camera team, and are consistent with the EMCCD201 datasheet.

7 Bibliography

- [1] P. Willems, "Starshade to TRL5 (S5) Technology Development Plan".
- [2] B. Crill and N. Siegler, "Exoplanet Exploration Program, 2018 Technology Plan Appendix, JPL Document No: D-101271".
- [3] M. C. Noecker, "Alignment of a terrestrial planet finder starshade at 20-100 megameters.," *Techniques and Instrumentation for Detection of Exoplanets III.*, vol. 6693, 2007.
- [4] W. Cash, "The New Worlds Observer: the astrophysics strategic mission concept study," *EPJ Web of Conferences*, vol. 16, 2011.
- [5] M. C. Noecker, "Alignment control for an external-occulting terrestrial planet finder mission," *Advances in the Astronautical Sciences*, vol. 137, no. 19, 2010.
- [6] A. Harness and W. Cash, "Harness, Anthony, and Webster Cash. "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] M. Bottom, "Precise starshade stationkeeping and pointing with a Zernike wavefront sensor," *Techniques and Instrumentation for Detection of Exoplanets VIII.*, vol. 10400, 2017.
- [8] A. Romero-Wolf, *Priv. Communication*.
- [9] F. Castelli and R. Kurucz, "New grids of ATLAS9 model atmospheres," *arXiv preprint astro-ph/0405087*, 2004.
- [10] C. Fröhlich and J. Lean, "Total solar irradiance variations: The construction of a composite and its comparison with models," *Symposium-International Astronomical Union*, vol. 185, 1998.
- [11] A. Harness and et al, "Modeling and performance predictions for the Princeton Starshade Testbed," *Techniques and Instrumentation for Detection of Exoplanets VIII, International Society for Optics and Photonics*, vol. 10400, 2017.
- [12] H. Tang, *Priv. communication*.
- [13] O. Hsu, P. D. Lisman and M. Vess, "Starshade Accommodations Interface Requirements Document".
- [14] S. B. Howell, *Handbook of CCD astronomy, Vol 5*, vol. 5, Cambridge University Press, 2006.
- [15] I. R. King, "Accuracy of measurement of star images on pixel array," *Publications of the Astronomical Society of the Pacific*, vol. 95.564, no. 163, 1983.
- [16] E. Cady, "Boundary diffraction wave integrals for diffraction modeling of external occulters," *Optics Express*, vol. 20.14, pp. 15196-15208, 2012.
- [17] "SPICE Homepage," [Online]. Available: <https://naif.jpl.nasa.gov/naif/>.
- [18] "MONTE Brochure," [Online]. Available: <https://montepy.jpl.nasa.gov/static/document/brochure.pdf>.
- [19] A. Rocketdyne, "Aerojet Rocketdyne Bipropellant Thrusters Data Sheet," [Online]. Available: <https://www.rocket.com/files/aerojet/documents/Capabilities/PDFs/Bipropellant%20Data%20Sheets.pdf>.

- [20] T. L. B. Flinois, C. Seubert and D. P. Scharf, "Fine formation flying for a Starshade rendezvous mission," *To be published*, 2019.
- [21] G. Singh, *QPCAP: A Quadratic Programming Algorithm for the Control Allocation Problem*, Pasadena, CA: California Institute of Technology, JPL Engineering Memorandum 3457-03-013, 91125, Tech. Rep, 2003.
- [22] D. P. Scharf, J. A. Keim and F. Y. Hadaegh, "Flight-Like Ground Demonstrations of Precision Maneuvers for Spacecraft Formations — Part I," *IEEE Systems Journal*, vol. 4, no. 1, pp. 84-95, 2010.
- [23] "Habex Interim Report," [Online]. Available: https://www.jpl.nasa.gov/habex/pdf/HabEx_Interim_Report.pdf.
- [24] M. Bottom, L. S. Neat, L. K. Harding, P. Morrissey, S. R. Meeker and R. T. Demers, Smartphone scene generator for efficient characterization of visible imaging detectors, Proceedings Volume 10709, High Energy, Optical, and Infrared Detectors for Astronomy VIII; 107092R, 2018.