

TDEM Final Report: Formation Flying for External Occulters

Award #: NNX15AD31G

PI: N. Jeremy Kasdin

Anthony Harness, Leonel Palacios, N. Jeremy Kasdin
Princeton University

September 30, 2021

Signature Page



N. Jeremy Kasdin (Principal Investigator)
Princeton University

Date

Nick Siegler (ExEP Program Chief Technologist)
Exoplanet Exploration Program, NASA / JPL - California Institute of Technology

Date

Brendan Crill (Deputy ExEP Program Chief Technologist)
Exoplanet Exploration Program, NASA / JPL - California Institute of Technology

Date

Doug Hudgins (ExEP Program Scientist)
NASA - HQ

Date

Contents

Executive Summary	1
1 Introduction	2
2 TDEM Milestones	3
3 Methods	4
3.1 Formation Flying Scheme	4
3.2 Control + Estimation Algorithms	7
3.3 Formation Sensing	9
3.4 Simulations	12
3.5 Testbed Hardware	13
4 Results	16
4.1 Milestone 1: Simulation results showing $3\text{-}\sigma$ lateral alignment < 1 meter	17
4.2 Milestone 2: Monte Carlo analysis of reliable sensor transition	19
4.3 Milestone 3: Lab results showing centimeter-level sensing from pupil images	20
4.4 Milestone 4: Lab results showing closed-loop control, with simultaneous high contrast measurements	23
5 Conclusion	25
5.1 Future Work: Towards new control/estimation schemes	25
Acknowledgements	26

Executive Summary

This report details the methods and results of the TDEM activity exploring formation flying of starshade external occulters. A combination of analytical, simulation, and experimental studies are used to validate our approach to formation sensing and control between starshade and telescope. We present results supporting the successful completion of four milestones: 1) closed-loop simulations demonstrating $3\text{-}\sigma$ formation flying accuracy that meets the lateral alignment requirement of 1 meter; 2) Monte Carlo analyses demonstrating a reliable capability to transition from retargeting to observation modes; 3) experimental demonstration of centimeter-level accuracy in position sensing; and 4) experimental verification of closed-loop control with simultaneous high-contrast measurements.

Formation flying is achieved through a linear quadratic controller with integral action and position estimation is done with an unscented Kalman filter. High precision position measurements are made by imaging the out-of-band diffraction pattern incident on the telescope's aperture and using a non-linear least squares method to solve an analytic approximation of the diffraction pattern. Performance of these methods are verified with a hardware demonstration of closed-loop control with simultaneous high-contrast observations performed in the Princeton Starshade Testbed. We demonstrate a $3\text{-}\sigma$ lateral alignment of 87 cm achieved via a $3\text{-}\sigma$ position sensing accuracy of 7.5 cm.

Our results presented here demonstrate that we have met all four of the TDEM milestones for this study.

1 Introduction

A primary goal of the exoplanet research community—and humanity itself—is the detection and characterization of Earth-like planets. This goal is a central component of the NASA strategic plan ([14], Goal 2) and one of the top three science objectives in the ASTRO 2010 decadal survey [3]. More poetically, it strikes to the heart of the ancient, fundamental questions Where did we come from? and Are we alone? [13]

Many approaches for imaging these planets have been investigated over the decades, including nulling interferometers, both structurally connected and free-flying, pupil interferometers, various types of coronagraphs, and external occulters or starshades. Each has advantages and disadvantages, and all are at various levels of technological readiness. In this final report we focus on starshades; more specifically, we describe our successful TDEM studying the feasibility of the hardware and software approach to formation flying a starshade and telescope to the required accuracy. A starshade is a promising approach to creating the needed high-contrast for imaging planets as small as the Earth and in the habitable zone of their parent star. One particularly exciting concept is the Starshade Rendezvous mission (SRM), where a modest size starshade is flown to rendezvous with the Nancy Grace Roman Space Telescope and operate with the Coronagraph Instrument. [16] We used that mission as a guide for selecting the various parameters in the simulations and experiments described in this report.

Because of the potential of starshades, NASA has embarked on a series of studies to raise the Technology Readiness Level (TRL) of essential starshade technology to TRL 5 (known as Starshade to TRL 5 or S5) [21]. Key among them is precision formation flying to keep the telescope in the deep shadow cast by the starshade. Our TDEM award supported analytical, simulation, and experimental studies to confirm our approach for stationkeeping the starshade and telescope.

In the S5 work of Ref. 4, the authors present a formation sensing scheme based on an image library matching algorithm to extract the lateral displacement of the starshade from pupil images of the residual starlight [2]. The end-to-end performance of this scheme was demonstrated with analytic, numerical, and experimental data, achieving a position precision of better than 10 cm ($3\text{-}\sigma$). They also present a formation control scheme [5] and used high-fidelity Monte Carlo simulations to demonstrate robust control of the starshade to within 1 meter of the line-of-sight. The results of Ref. 4 demonstrate that the technology gap of starshade formation sensing is at TRL 5.

The work presented in this report builds upon that previous work to further demonstrate the viability of this approach to formation sensing. In particular, we study sensing and control in the acquisition phase of retargeting when the starshade is still 10's of meters from alignment and we demonstrate a reliable sensor transition to the observation phase. Additionally, our work provides a hardware demonstration of closed-loop control with simultaneous high-contrast observations. (Note we also use a different sensing and control scheme, demonstrating the robustness of starshade formation flying.)

This final report details the methods and results of this TDEM activity and communicates the successful completion of the four TDEM milestones introduced in Section 2. In

Section 3 we summarize the methods used in demonstrating closed-loop formation flying in simulation and hardware. In Section 4 we present the results from the four milestone activities and use the results to argue the successful completion of each milestone; this represents the bulk of the work and of this report. Finally, we conclude in Section 5 and outline directions for future work. Additional details of the control and estimation algorithms are presented in Ref. [15], our peer-reviewed publication in *Acta Astronautica* that represents the dissemination of this work to the broader community. Additional details of the testbed are presented in Ref. [7].

2 TDEM Milestones

We briefly state the quantitative milestones that determine the successful completion of this TDEM effort; discussion of the completion of these milestones is saved for Section 4.

TDEM Milestones:

1. Verify, via a closed-loop simulation employing realistic sensor and actuator models informed by experiments, that the $3\text{-}\sigma$ formation flying accuracy meets the lateral alignment requirement of 1 m.
2. Demonstrate a reliable capability to transition from retargeting to observation modes via Monte Carlo analyses starting from post-retargeting formation dispersions through precision formation sensor acquisition to steady-state 1 m alignment.
3. Demonstrate, with a breadboard precision formation alignment sensor, positioning accuracy of $1/8\times$ telescope diameter, scalable to flight requirements in a hardware testbed.
4. Verify positioning accuracy and control with hardware-in-the-loop testing using the formation sensor in the scaled occulter test facility and validate performance using optical contrast measurements at visible wavelengths.

Note: The originally proposed Milestone 3 aimed to demonstrate sub-pixel centroiding accuracy scalable to flight requirements. However, as understanding of the formation sensing problem has evolved, we have switched from an alignment sensor operating in the focal plane, to one that operates in the pupil plane of the telescope. As such, sub-pixel centroiding is no longer relevant to the accuracy achieved by the formation sensor. The pupil plane sensor extracts the lateral alignment from images of the diffraction pattern incident on the telescope’s aperture and the position accuracy is measured relative to the size of the telescope aperture. We use the S5 [21] requirement of 30 cm for the Starshade Rendezvous mission with a 2.4 m diameter telescope to set the Milestone 3 accuracy requirement to $1/8\times$ the telescope diameter.

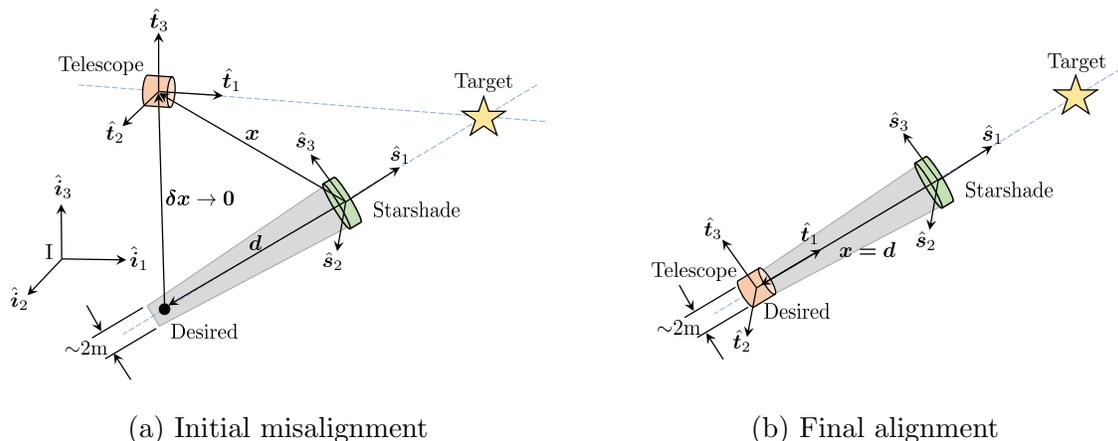


Figure 1: The formation flying problem

3 Methods

This section details the design of a formation flying controller for starshade alignment from the acquisition to observation phases. Trajectory tracking is achieved through a linear quadratic controller with integral action and estimation is done with an unscented Kalman filter. Later, this controller is used to perform hardware-in-the-loop (HWIL) simulations in the existing Princeton Starshade Testbed to validate formation sensing and control algorithms while maintaining high contrast with a flight-like starshade. More details and background on the algorithms can be found in Ref. [15].

3.1 Formation Flying Scheme

The formation flying problem at hand is that of formation acquisition and station-keeping as described in Fig. 1. Consider two spacecraft orbiting around the Earth-Sun L2 point with one designated as the leader and the other as the follower. According to the configuration of our laboratory, and without loss of generality, in this work the starshade is selected as the leader and the telescope as the follower. Both the starshade and the telescope are assumed to be continuously pointing toward the target star by means of an attitude control system (to be presented in future work). The telescope acquires formation with respect to the starshade by relocating itself towards the desired position while the alignment error $\delta \mathbf{x} \rightarrow 0$. Then, once located at the desired position, the telescope performs station-keeping for a determined period of time. Figure 2 provides a flowchart summarizing the sensing and control steps, which depend on the current sensing region, detailed in Section 3.3.

The motion of the starshade spacecraft is given by the differential equation of motion

$$\ddot{\mathbf{r}}_S = - \sum_{i=1}^N \mu_i \left(\frac{\mathbf{r}_{Si}}{\|\mathbf{r}_{Si}\|^3} \right) + \mathbf{u}_S \quad \text{for } i = 1, 2, \dots, N, \quad (1)$$

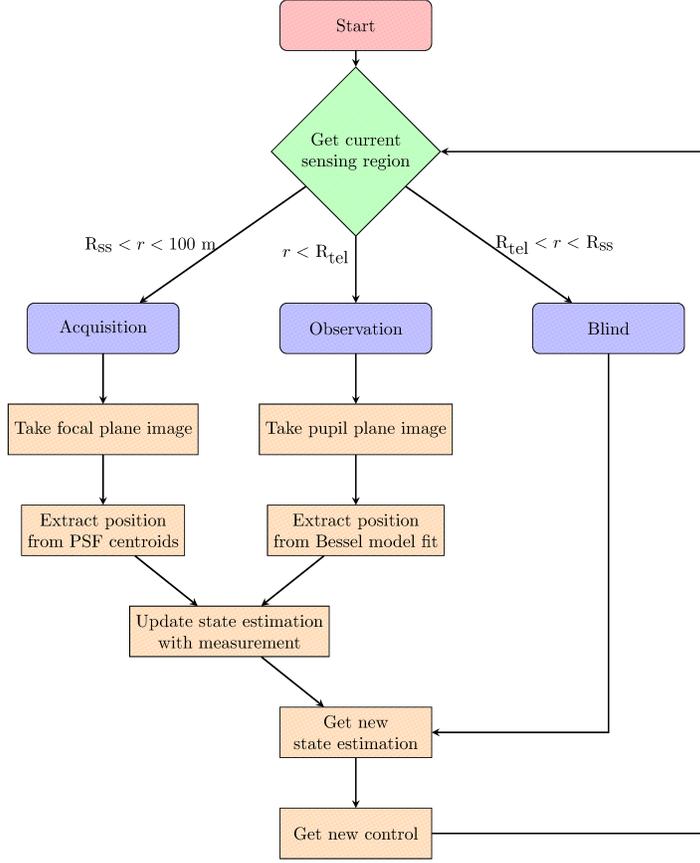


Figure 2: Flowchart summarizing the sensing and control steps depending on the current sensing region. The current sensing region (explained in Section 3.3) depends on the current offset (r), relative to the starshade radius (R_{SS}) and telescope radius (R_{tel}).

where \mathbf{r}_{S_i} is the position vector of the starshade with respect to the i -th celestial body, μ_i is its gravitational parameter, N is the total number of celestial bodies considered in the summation, and the term \mathbf{u}_S corresponds to control inputs. Likewise, the motion of the telescope spacecraft is given by

$$\ddot{\mathbf{r}}_T = - \sum_{i=1}^N \mu_i \left(\frac{\mathbf{r}_{T_i}}{\|\mathbf{r}_{T_i}\|^3} \right) + \mathbf{u}_T \quad \text{for } i = 1, 2, \dots, N. \quad (2)$$

The relative position of the telescope with respect to the starshade is defined as $\Delta \mathbf{r} = \mathbf{r}_T - \mathbf{r}_S$ which yields the equations of relative motion

$$\Delta \ddot{\mathbf{r}} = - \sum_{i=1}^N \mu_i \left(\frac{\mathbf{r}_{T_i}}{\|\mathbf{r}_{T_i}\|^3} - \frac{\mathbf{r}_{S_i}}{\|\mathbf{r}_{S_i}\|^3} \right) + (\mathbf{u}_T - \mathbf{u}_S) + \mathbf{w} \quad \text{for } i = 1, 2, \dots, N, \quad (3)$$

where the term $\mathbf{w} \sim \mathcal{N}(\mu_{SRP}, \sigma)$ adds uncertainty to the relative motion using a normal random distribution around its mean, μ_{SRP} , defined as the average of the differential acceleration caused by the solar radiation pressure on the relative motion. In order to employ the control and estimation techniques described in later sections, it is necessary to linearize Eq. (3). The resulting linear, time-varying system is defined as a deterministic state-space system with the state vector $\mathbf{x} = [\Delta\mathbf{r} \ \Delta\dot{\mathbf{r}}]^T$ and has the usual structure $\dot{\mathbf{x}}(t) = \mathbf{A}(t)\mathbf{x} + \mathbf{B}\mathbf{u}(t)$, which, for the state variables as defined, becomes

$$\begin{bmatrix} \Delta\dot{\mathbf{r}} \\ \Delta\ddot{\mathbf{r}} \end{bmatrix} = \begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{I}_{3 \times 3} \\ \mathbf{\Theta} & \mathbf{0}_{3 \times 3} \end{bmatrix} \begin{bmatrix} \Delta\mathbf{r} \\ \Delta\dot{\mathbf{r}} \end{bmatrix} + \begin{bmatrix} \mathbf{0}_{3 \times 3} \\ \mathbf{I}_{3 \times 3} \end{bmatrix} (\mathbf{u}_T - \mathbf{u}_S), \quad (4)$$

where $\mathbf{\Theta}$ is defined as in Ref. [12]

$$\mathbf{\Theta} = - \sum_{i=1}^N \frac{\mu_i}{\|\mathbf{r}_{Si}\|^3} \{ \mathbf{I}_{3 \times 3} + 3 [\mathbf{e}_{Si} \ \mathbf{e}_{Si}^T] \}, \quad (5)$$

and \mathbf{e}_{Si} is the unit vector along \mathbf{r}_{Si} . It also is necessary to define how to obtain reference trajectories for the relative motion of the telescope. Once the telescope converges toward these trajectories by means of the actions of the controller (formation acquisition), the telescope/starshade system remains aligned (formation-keeping) toward a target star during science mode. For this purpose, the unit vector of the relative position between the starshade and the target star, \mathbf{e}_{S^*} , is required. If a target star is selected with right ascension γ (rad), declination δ (rad), and distance from the Earth D (km), then its position with respect to Earth \mathbf{r}_{*E} is obtained in the inertial frame as

$$\mathbf{r}_{*E} = D \begin{bmatrix} \cos(\delta) \cos(\gamma) \\ \cos(\delta) \sin(\gamma) \\ \sin(\delta) \end{bmatrix}. \quad (6)$$

Given the position vector of the starshade with respect to the Earth, \mathbf{r}_{SE} , then the unit vector of the position of the target star with respect to the starshade, \mathbf{e}_{*S} , is defined as:

$$\mathbf{e}_{*S} = \frac{\mathbf{r}_{*E} - \mathbf{r}_{SE}}{\|\mathbf{r}_{*E} - \mathbf{r}_{SE}\|} \quad (7)$$

and the desired reference trajectory of the telescope is

$$\mathbf{r}_R = -\bar{D}\mathbf{e}_{*S}, \quad (8)$$

where \bar{D} (defined here between 20,000 - 50,000 km) is the desired distance between the telescope and the starshade. The desired initial velocity is simply approximated from a first order difference from the desired position. The controller design is presented next.

3.2 Control + Estimation Algorithms

During the acquisition phase, the telescope must align with the starshade within the allowed tolerances of the mission. In this phase, velocities are small and maneuvering distances and times are short. Therefore, a discrete-time, linear quadratic regulator with integral action (LQRi) and an Unscented Kalman Filter (UKF) has been selected for this task. Additionally, in order to provide a realistic scenario, we assume the spacecraft is equipped with on-off thrusters (with nominal thrust of 22 N). Therefore, the control commands from the LQRi controller must be converted to on-off via pulse width modulation (PWM). A schematic representation of the controller is presented in Fig. 3.

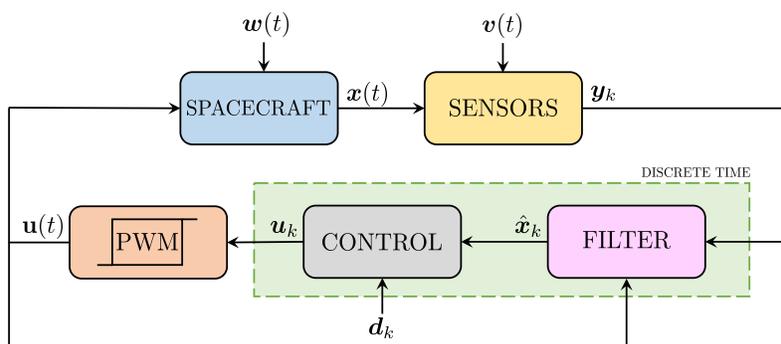


Figure 3: Schematic representation of the control/estimation scheme

The LQRi approach requires the linearized version of the actual relative motion as in Eq. (4) and is designed in such way that it tracks the reference trajectory presented in the previous section. The integral action acts directly on the relative position and adds a certain degree of robustness to the approach by compensating for the nonlinearities and uncertainty in the plant. The LQRi is defined in incremental quantities to require most of the computations to be done using small quantities only [17]. This approach, together with the integral action, adds an extra degree of robustness to the controller, especially since the problem consists of position vectors of the order of thousands of kilometers. Consider the discrete-time version of Eq. (4) as

$$\mathbf{x}_{k+1} = \mathbf{A}_k \mathbf{x}_k + \mathbf{B}_k \mathbf{u}_k \quad (9)$$

$$\mathbf{y}_k = \mathbf{C}_k \mathbf{x}_k \quad (10)$$

and in incremental form as

$$\Delta \mathbf{x}_{k+1} = \mathbf{A}_k \Delta \mathbf{x}_k + \mathbf{B}_k \Delta \mathbf{u}_k, \quad (11)$$

where $\Delta \mathbf{x}_k = \mathbf{x}_k - \mathbf{x}_{k-1}$. Also, the measurement equation is rewritten as

$$\mathbf{y}_k - \mathbf{r}_k = \mathbf{y}_{k-1} - \mathbf{r}_k + \mathbf{C}_k \Delta \mathbf{x}_k. \quad (12)$$

Putting these equation together with the integral action into an augmented system we have:

$$\begin{bmatrix} \Delta \mathbf{x}_{k+1} \\ \mathbf{y}_k - \mathbf{r}_k \end{bmatrix} = \begin{bmatrix} \mathbf{A}_k & \mathbf{0}_{n \times m} \\ \mathbf{C}_k & \mathbf{I}_{m \times m} \end{bmatrix} \begin{bmatrix} \Delta \mathbf{x}_k \\ \mathbf{y}_{k-1} - \mathbf{r}_k \end{bmatrix} + \begin{bmatrix} \mathbf{B}_k \\ \mathbf{0}_{m \times r} \end{bmatrix} \Delta \mathbf{u}_k \quad (13)$$

$$\mathbf{y}_k - \mathbf{r}_k = \begin{bmatrix} \mathbf{C}_k & \mathbf{I}_{m \times m} \end{bmatrix} \begin{bmatrix} \Delta \mathbf{x}_k \\ \mathbf{y}_{k-1} - \mathbf{r}_k \end{bmatrix}, \quad (14)$$

which can also be expressed in compact form as

$$\tilde{\mathbf{x}}_{k+1} = \tilde{\mathbf{A}}_k \tilde{\mathbf{x}}_k + \tilde{\mathbf{B}}_k \Delta \mathbf{u}_k \quad (15)$$

$$\tilde{\mathbf{y}}_k = \tilde{\mathbf{C}}_k \tilde{\mathbf{x}}_k. \quad (16)$$

Next, consider the following cost function

$$J_k = \frac{1}{2} \sum_{k=i}^{\tau-1} \left[\tilde{\mathbf{x}}_k^T \tilde{\mathbf{Q}}_k \tilde{\mathbf{x}}_k + \Delta \mathbf{u}_k^T \mathbf{R}_k \Delta \mathbf{u}_k \right], \quad (17)$$

where the weight matrix $\tilde{\mathbf{Q}}_k = \tilde{\mathbf{C}}_k^T \mathbf{Q}_k \tilde{\mathbf{C}}_k$. Thus, the required optimal control feedback signal required to track a reference trajectory is defined as

$$\Delta \mathbf{u}_k = \mathbf{K}_k \tilde{\mathbf{x}}_k, \quad (18)$$

where the control gain is obtained from the usual LQR algorithm with the equations

$$\mathbf{K}_k = - \left(\mathbf{R}_k + \tilde{\mathbf{B}}_k^T \mathbf{S}_k \tilde{\mathbf{B}}_k \right)^{-1} \tilde{\mathbf{B}}_k^T \mathbf{S}_k \tilde{\mathbf{A}}_k \quad (19)$$

and

$$\mathbf{S}_k = \tilde{\mathbf{Q}}_k + \mathbf{K}_k^T \mathbf{R}_k \mathbf{K}_k + \left(\tilde{\mathbf{A}}_k + \tilde{\mathbf{B}}_k \mathbf{K}_k \right)^T \mathbf{S}_k \left(\tilde{\mathbf{A}}_k + \tilde{\mathbf{B}}_k \mathbf{K}_k \right). \quad (20)$$

During science observations, the positioning requirement is simply that the telescope remains within the deepest part of the starshade's shadow. The size of this deep shadow is a variable in the mission design; most conceptual designs to date have assumed 1 meter (in radius) larger than the telescope's aperture; this determines the alignment tolerance. The aggressiveness of the control law is tuned such that the telescope always remains within the ± 1 m box, while not being overly aggressive. Using a tighter control law consumes more fuel and requires more thruster firings, which interrupt science observations.

Since we are assuming the state measurement is incomplete and noisy, an estimation strategy must be included in the design. Given its capacity to handle nonlinear systems and its derivative-free design, in this project the Unscented Kalman filter (UKF) is selected as the estimation strategy. Table 1 summarizes the UKF. Since this is a standard estimation tool, the interested reader may consult references [20, 10], for additional details.

Table 1: Unscented Kalman filter equations

Time update:

$$\begin{aligned}\hat{\mathbf{x}}_k^{(i)} &= \mathbf{f}\left(\hat{\mathbf{x}}_{k-1}^{(i)}, \mathbf{u}_k, t_k\right) \\ \hat{\mathbf{x}}_k^- &= \sum_{i=0}^{2n} W_i^m \hat{\mathbf{x}}_k^{(i)} \\ \mathbf{P}_k^- &= \sum_{i=0}^{2n} W_i^c \left(\hat{\mathbf{x}}_k^{(i)} - \hat{\mathbf{x}}_k^-\right) \left(\hat{\mathbf{x}}_k^{(i)} - \hat{\mathbf{x}}_k^-\right)^T + \mathbf{Q}'_{k-1} \\ \hat{\mathbf{y}}_k^{(i)} &= \mathbf{h}\left(\hat{\mathbf{x}}_k^{(i)}, t_k\right)\end{aligned}$$

Measurement update:

$$\begin{aligned}\mathbf{P}_y &= \sum_{i=0}^{2n} W_i^c \left(\hat{\mathbf{y}}_k^{(i)} - \hat{\mathbf{y}}_k\right) \left(\hat{\mathbf{y}}_k^{(i)} - \hat{\mathbf{y}}_k\right)^T + \mathbf{R}'_k \\ \mathbf{P}_{xy} &= \sum_{i=0}^{2n} W_i^c \left(\hat{\mathbf{x}}_k^{(i)} - \hat{\mathbf{x}}_k^-\right) \left(\hat{\mathbf{y}}_k^{(i)} - \hat{\mathbf{y}}_k\right)^T \\ \mathbf{K}' &= \mathbf{P}_{xy} \mathbf{P}_y^{-1} \\ \hat{\mathbf{x}}_k^+ &= \hat{\mathbf{x}}_k^- + \mathbf{K}' \left(\mathbf{y}_k - \hat{\mathbf{y}}_k\right) \\ \mathbf{P}_k^+ &= \mathbf{P}_k^- - \mathbf{K}' \mathbf{P}_y \mathbf{K}'^T\end{aligned}$$

3.3 Formation Sensing

The starshade operates by casting a deep shadow over the aperture of the telescope that blocks the starlight and allows only the exoplanet’s light into the telescope. During science operations, the telescope must stay within the deepest part of the shadow, lest the image be ruined by light diffracting around the starshade. The depth of the shadow is a steep function of radius and the intensity increases by 10 orders of magnitudes across the size of the starshade, which is typically 10’s of meters across. The width of the deep shadow can be made wider than the telescope aperture to loosen alignment tolerances, but that requires a larger starshade operating at a larger separation, which decreases the number of targets accessible for a given mission lifetime. Thus, better formation-keeping performance helps to maximize the scientific yield.

There are three regions of formation sensing we consider in this study, each of which use a different approach to obtaining the relative offset between the starshade and the line of sight between the telescope and target star. For each case, the axial distance between starshade and telescope is provided by RF ranging, with an accuracy of ~ 5 meters. The three formation sensing regions are: the **acquisition** region ($R_{\text{starshade}} < \text{offset} < 100$ m) where both the starshade and target star are in the telescope’s field of view and a beacon on the starshade provides a signal for a bearing measurement; the **blind** region ($R_{\text{telescope}} < \text{offset} < R_{\text{starshade}}$) where the star disappears behind the starshade and no longer provides

an accurate bearing measurement; and the **observation** region (offset $< R_{\text{telescope}}$) during which science observations are made and the starshade’s position is extracted from the diffraction pattern incident on the telescope. The three regions are discussed in further detail below.

3.3.1 Acquisition Region

As the starshade completes its retargeting phase and makes its approach to the line of sight to the target star, it will enter the field of view of the telescope and, with both the starshade and target star visible, we can exploit the resolution provided by the large aperture and obtain a bearing measurement between the two. A laser beacon on the starshade provides a strong signal from which to extract the starshade’s position. By synchronizing the observations with toggling of the beacon, a difference image easily differentiates on the focal plane the beacon’s signal from that of the target star.

The position accuracy in this region depends on the distance to the starshade and the accuracy in centroiding the beacon’s and star’s point spread function on the detector, which scales linearly with the size of the primary mirror. For SRM, with a 2.4 m primary mirror, the starshade at 26,000 km distance, and conservatively assuming we can centroid to only 1/5 the width of the $\lambda = 500$ nm PSF, we estimate a position accuracy of ~ 1.1 meters.

3.3.2 Blind Region

Once the starshade begins to occult the target star, we start to lose the star’s signal and that signal becomes distorted as the light diffracts around the starshade, leading to a disconnect between the position of the star and the center of light in the image [18]. Observing outside of the starshade’s design bandpass can recover some of the loss in flux, but the distortion problem still remains. The power received to estimate the offset could be used [18], but a study of that approach is out of scope for this work. Instead, we decided to be “blind” while in this region and rely on our UKF estimator to provide a state estimation without updating it with measurements. As will be shown, this is sufficient to reliably reach the observation region where the fine alignment sensor takes over.

3.3.3 Observation Region

For most starshade designs being considered, the shadow is typically made larger than the aperture by ~ 1 meter in radius. Controlling to within this region is made possible by having knowledge of the lateral position by 10’s of cm of accuracy. Over 10,000’s of kilometers of intersatellite separation, this position knowledge translates to a sub-milliarcsecond angle measurement, a feat difficult even for dedicated astrometric missions. Instead, we directly sample the diffraction pattern produced by the starshade, which tracks the position of the starshade one-to-one [4], thus allowing us to forgo making an angular measurement while obtaining the high precision position measurements over large separations.

Outside of the starshade’s operating wavelength bandpass, the starlight suppression quickly deteriorates and a strong signal of light diffracting around the starshade re-emerges

and forms a diffraction peak called Poisson’s (or Arago’s) spot (see Fig. 4). This spot is continuously formed along the central axis of the starshade such that there is a one-to-one mapping from the lateral position of the starshade to that of the spot. A sensor operating at an out-of-band wavelength can image the entrance pupil of the telescope, locate the peak in the diffraction incident on the aperture, and extract the lateral position of the starshade. Operating outside of the starshade’s bandpass means the signal is only attenuated by a factor of $10^2 - 10^3$ (rather than 10^{10} in the science band), meaning that a position can be determined with only seconds-long exposure times. In the case of the Roman telescope, the coronagraph’s wavefront sensor is conjugate to the aperture and can provide the pupil images for alignment sensing.

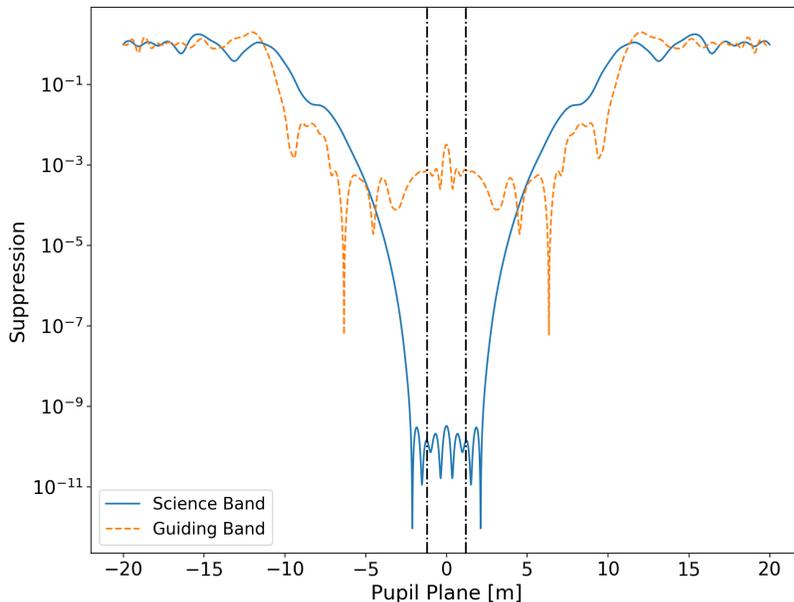


Figure 4: Suppression (relative extinction of starlight) at the telescope aperture for the science bandpass (solid line) and guiding bandpass (dashed line). The vertical lines denote the radius of the telescope. The deep shadow is 1 m larger than the telescope in radius.

To locate the Poisson spot and extract the starshade’s position from the pupil image, we use a simple model based on the approximate solution to the diffraction equation. Other methods, such as centroiding on the intensity or a look-up table of pre-computed images [4, 2], could also be used. The diffraction pattern of the starshade in out-of-band light is similar to that of a circular disk and can be approximated by a Bessel function [8]. For a starshade of radius R , separated by distance z , and operating at wavelength λ , the intensity at position (x, y) at the telescope’s pupil is approximated by,

$$I(x, y) \approx J_0^2 \left(\frac{2\pi R \sqrt{(x - x_s)^2 + (y - y_s)^2}}{\lambda z} \right), \quad (21)$$

where the unknown starshade’s lateral position is given by (x_s, y_s) and can be solved for via non-linear least squares fit to the pupil image. Pixels that are in the shadow of the secondary mirror are known in advance and are not included in the least squares fit. Fig. 5 shows an experimental image of the diffraction pattern of the starshade in the lab, which represents a 26 m diameter starshade separated by 26,000 km. The Bessel function nature of the diffraction is clearly visible. The results of Section 4.3 show < 5 cm accuracy is achievable by the pupil sensor for a range of signal-to-noise ratios (SNR).

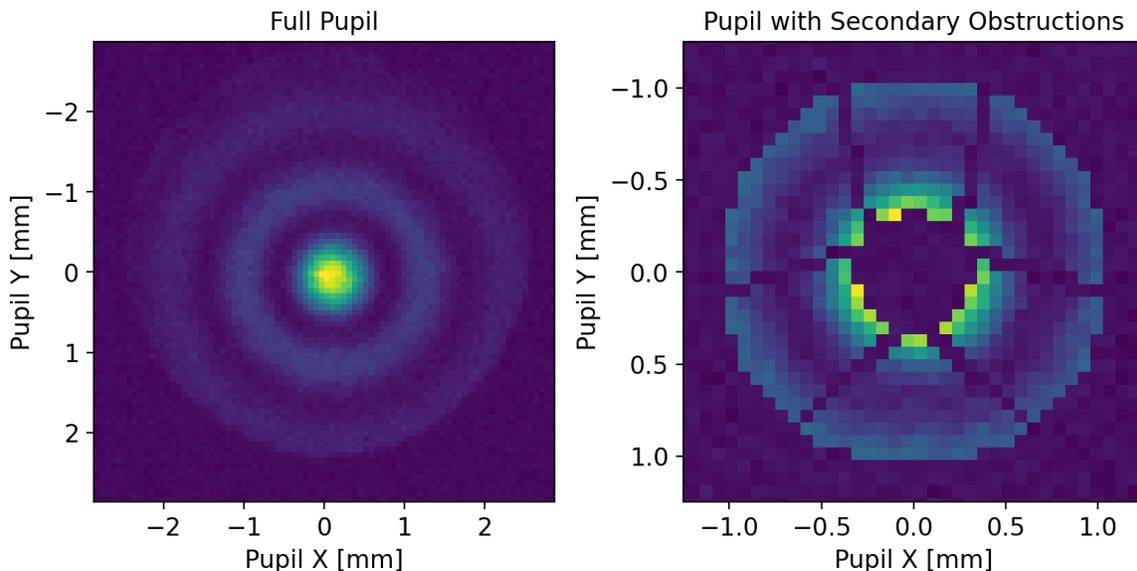


Figure 5: Experimental pupil image showing diffraction pattern of starshade in the lab, which represents a 26 m diameter starshade separated by 26,000 km. Left: full pupil image. Right: pupil image with secondary mirror obstructions artificially imposed.

3.4 Simulations

The HWIL experiments were designed and built to be interchangeable with a simulation-based testbed emulator. The emulator allowed us to develop and debug the control+estimation and position extraction algorithms and quickly implement them in hardware. It also enabled additional robustness studies through Monte Carlo simulations. Each piece of hardware, including communication between hardware components, has a software equivalent. Focal and pupil plane images used by the alignment sensor were simulated with diffraction calculations using optical models that have been validated to high contrast levels [7, 6]. A detector model consistent with the measured noise properties of the lab detector added simulated noise to the images.

3.5 Testbed Hardware

3.5.1 Overview

The experimental work for this activity was done in the Princeton Starshade Testbed, which was designed to perform sub-scale tests [7] of starshades to experimentally validate scalar optical models at a flight-like Fresnel number of ~ 15 (the Fresnel number being the dimensionless similarity parameter governing optical physics). The testbed includes a laser simulating the target star, a 25 mm mask corresponding to the starshade, and an optical system with detector representing the telescope. A schematic and image of the testbed are shown in Fig. 6; a detailed description of the testbed can be found in Ref. [7].

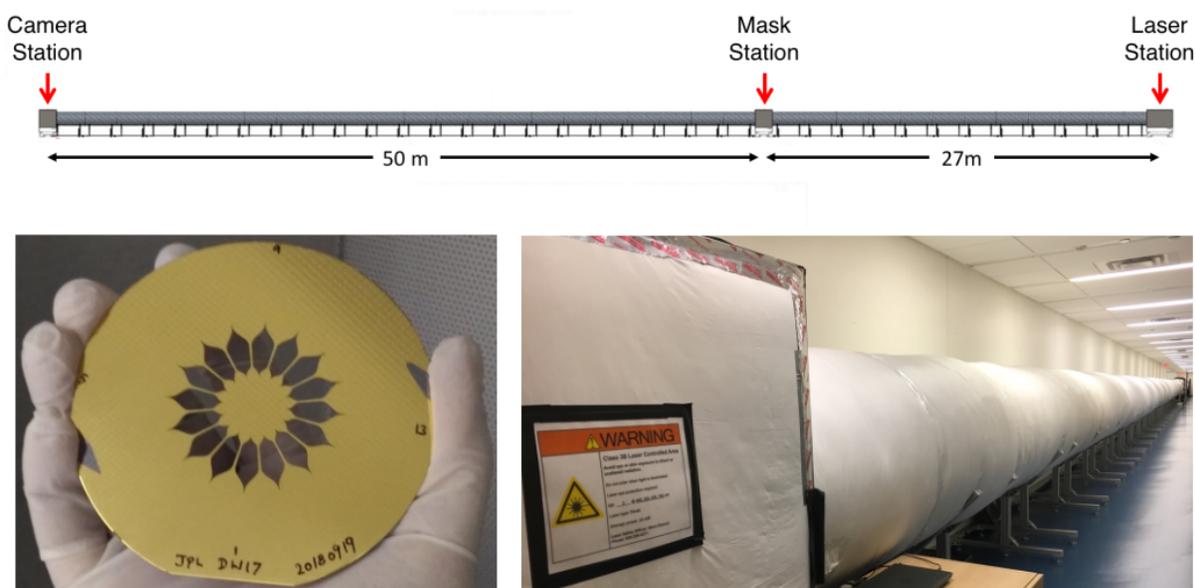


Figure 6: Top: schematic of testbed. Bottom Left: image of starshade mask. Bottom Right: image of testbed.

Since the testbed was built to validate optical diffraction models, the starshade size and separation are scaled to mimic the diffraction expected in flight. This is done by operating at a similar Fresnel number: $N = R^2/(\lambda Z_{\text{eff}})$, where R is the starshade radius, λ is the wavelength of light, and Z_{eff} is the effective starshade-telescope separation, which accounts for a light source at a finite distance, $Z_{\text{eff}} = Z_{\text{tel}}Z_{\text{source}}/(Z_{\text{tel}} + Z_{\text{source}})$. This results in the sizes of the starshade and telescope being scaled by the factor $\gamma = \sqrt{Z_{\text{space,tel}}/Z_{\text{lab,eff}}} \approx 1200$. Details on the laboratory and flight configurations are provided in Table 2.

	Laboratory	SRM
Telescope diameter (D)	1.98 mm	2.4 m
Starshade diameter ($2R$)	25.06 mm	26 m
Telescope - starshade sep. (Z_{tel})	50.0 m	26,000 km
Source - starshade sep. (Z_{source})	27.45 m	> 3 parsec
Guiding bandpass	405 nm	425 - 552 nm
Science bandpass	641 - 725 nm	615 - 800 nm
Fresnel number (at $\lambda = 405$ nm)	22	16
Pupil image resolution	62 μm /pixel	75 mm/pixel

Table 2: Physical parameters for the laboratory experiment and the Starshade Rendezvous Mission [19] architecture.

3.5.2 Light Source

The light source serving as the artificial star is a fiber-fed, monochromatic laser diode operating at 405 nm (for guiding observations) and 641 nm (for science observations). Fibers from the two wavelength channels are joined with a fused fiber coupler such that switching between wavelengths is done by toggling each channel on/off. The fiber into the testbed is single-mode for the science bandpass, but is multi-mode for the guiding light, although that should have a minimal impact on the diffraction pattern. The fiber terminates with a collimator and the output Gaussian beam is focused by an objective lens through a pinhole to spatially filter high-order aberrations.

3.5.3 Starshade

The starshade mask, shown in Fig. 6, is lithographically etched in a silicon wafer and coated with a thin metallic layer; details on the manufacturing process can be found in Ref. [1]. The starshade mask consists of an inner starshade representative of a free floating occulter that is supported in a silicon wafer via radial struts. The outer ring of the support wafer is also apodized to minimize diffraction. This starshade design has been demonstrated to achieve 10^{-10} contrast across its designed bandpass of 630 – 730 nm [7]. The starshade remains stationary throughout the experiment, but can be removed from the beam line in order to take images of the unocculted source for flux calibrations.

We imitate the laser beacon on the starshade with a simple LED mounted next to the starshade that can be toggled remotely. Since the laboratory starshade masks are too fragile to mount anything to them, the LED is mounted to the wall a few centimeters away and this distance is appropriately accounted for in the offset calculations.

3.5.4 Telescope

The optical system in the testbed that represents the space telescope is contained on a breadboard mounted to a motorized XY stage. In contrast to the space mission, the telescope in the lab performs the alignment motions, while the starshade remains stationary. At each time step of the experiment, the motorized stage moves to account for the change in relative position between starshade and telescope due to the accelerations from dynamics and controlled thrusts computed via the simulation model of starshade motion.

3.5.5 Pupil Imaging Sensor

The optical system can operate in a far field mode, which is focused at the light source and represents the science imaging mode, or in a pupil imaging mode, which is focused on the entrance aperture and is used as the alignment sensor. A motorized stage toggles a lens in/out of the beam line to switch between the modes. The detector on the backend is an Andor iXon Ultra 888 EMCCD, with 13 μm pixel size, liquid-cooled to -90°C , and operated with its conventional amplifier for better noise performance. Properties of the detector are summarized in Table 3.

Parameter	Value
Pixel size	13 μm
Inverse Gain	0.79 e^-/count
Read out noise	4.8 $\text{e}^-/\text{pixel}/\text{frame}$
Dark current	$7 \times 10^{-4} \text{e}^-/\text{pixel}/\text{s}$
CIC noise	0.0025 $\text{e}^-/\text{pixel}/\text{frame}$

Table 3: Noise properties of the pupil imaging sensor detector.

The aperture size is determined by the scaling factor set by the chosen value of $Z_{\text{space, tel}}$. The number of pixels that span the aperture in the pupil image has not yet been determined for a flight design, so we conservatively chose 32×32 pixels; more pixels would increase the accuracy of the position determination.

The optics in the lab are unobstructed lenses, while a space telescope would have an obstruction from the secondary mirror. To capture the loss in information of the diffraction pattern due to the obstruction, the obscuration pattern for the Roman Space Telescope is artificially imposed on the pupil images (see Fig. 5).

3.5.6 HWIL Operations

Figure 7 provides a flowchart summarizing the operational steps in conducting an HWIL experiment.

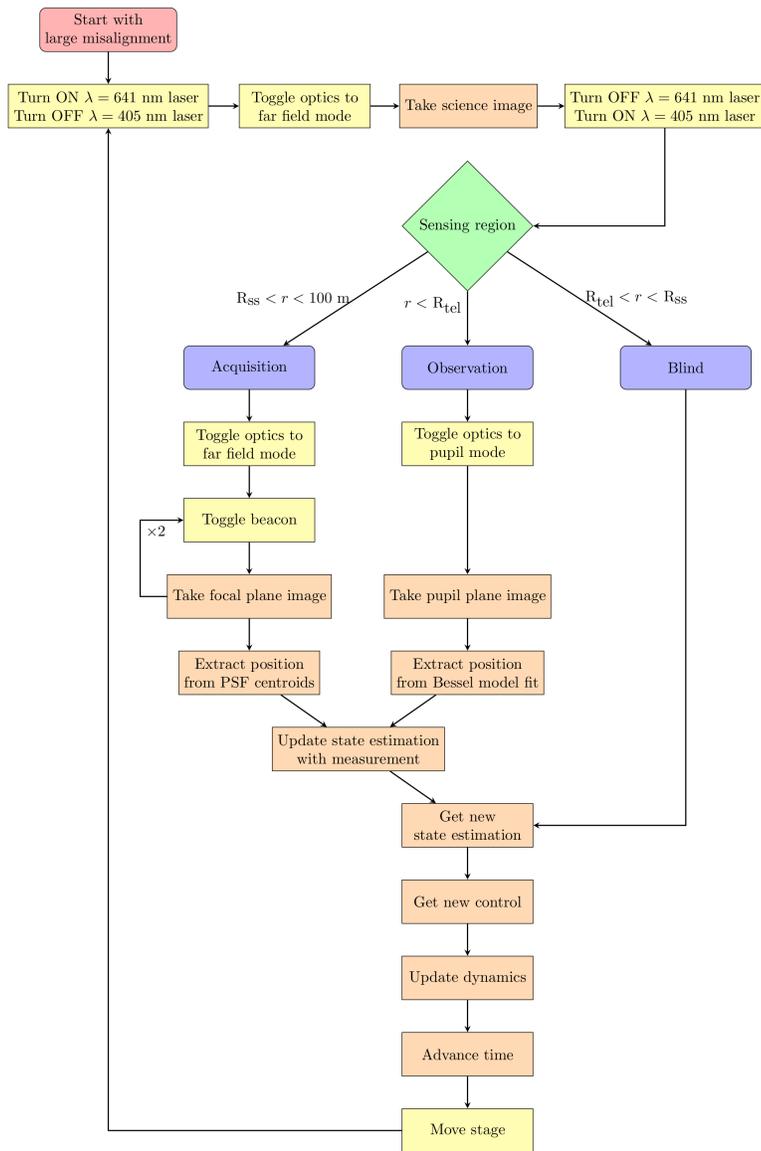


Figure 7: Flowchart for HWIL experiments.

4 Results

The objective of the simulation tests is to align the telescope with the starshade within the allowed tolerances of the mission for science operations. In agreement with the design of our laboratory, it is assumed that the starshade is orbiting freely around the L2 point and the telescope regulates its relative position with respect to the starshade by its control systems. Additionally, the required positions of celestial bodies in Eq. (3) are obtained from JPL's ephemeris DE432 and include information about the Sun, Mercury, Venus, Earth, Earth's

moon, Mars, Jupiter, Saturn, Uranus and Neptune (meaning $N = 10$). The controller and filter weighting matrices were tuned empirically and they are presented in Table 4 along with the remaining features of the experiment.

Table 4: Features of the HWIL experiments.

Feature	Value
Intersatellite separation	26,000 km
Target star	Epsilon Eridani [†]
Maximum thrust firing	22 N
Simulation time	2 hours
Exposure time	0.5 seconds
SNR of Spot of Arago	5
μ_{SRP}	3.6×10^{-10} km/s ²
σ_{SRP}	3.6×10^{-11} km/s ²
Time step	1 second

[†] $\alpha = 53.2292^\circ$, $\delta = -9.4581^\circ$ and $D = 9.90538 \times 10^{13}$ km

4.1 Milestone 1: Simulation results showing 3- σ lateral alignment < 1 meter

In this section we present simulation results demonstrating lateral alignment within the 1 meter requirement. The simulations were performed using the values in Table 4, starting in the observation region with a 0.8 meter misalignment. We first present results from a representative simulation and then present results from Monte Carlo simulations showing that we reliably meet the requirement of 3- σ lateral alignment < 1 meter.

The performance of the relative alignment error in the line-of-sight reference frame of the telescope is shown in Fig. 8. Only the two components of the pupil plane (y and z) are shown because these are the ones most relevant to science operations. The telescope sustains the alignment with the starshade within the prescribed error tolerance of ± 1 m (shown in red dotted lines) for the entire simulation. Figure 9 shows the instances of thruster firings throughout the simulation; the performance is summarized in Table 5. If there were no control applied to maintain alignment, the starshade would drift away to 25 m within an hour.

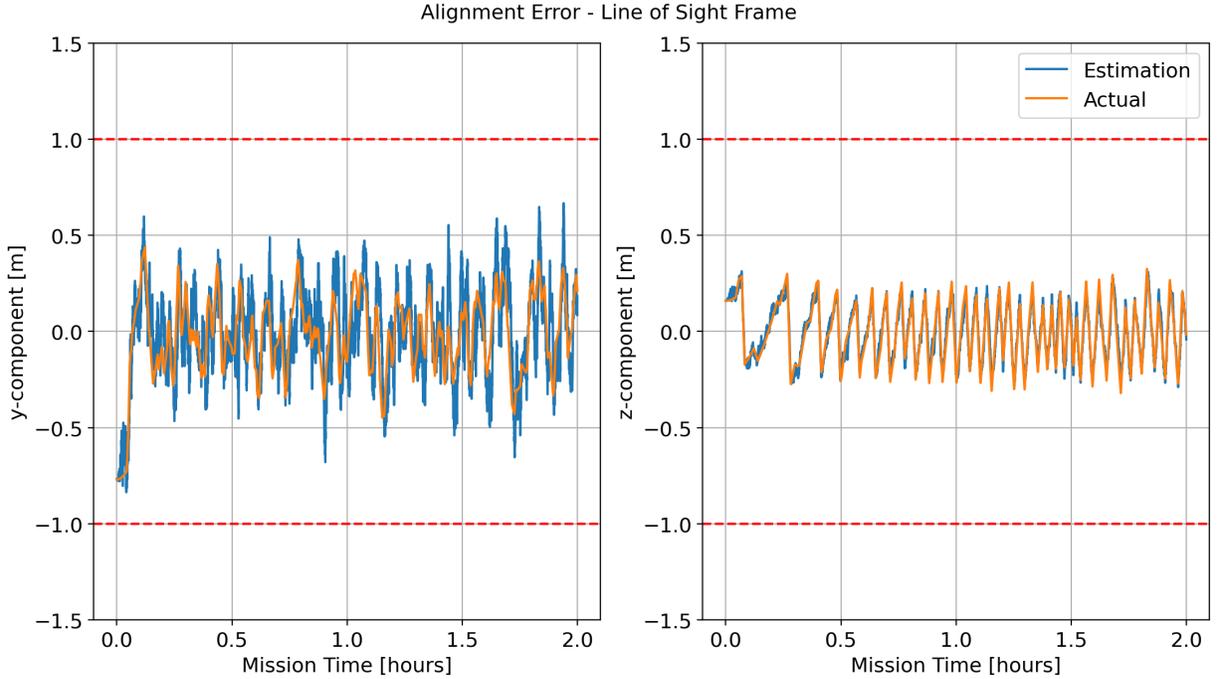


Figure 8: **Simulation** starting with a 0.8 m misalignment. The alignment error, estimated and actual, is shown in the two relevant coordinates of the telescope’s line of sight frame. The red dashed lines mark the ± 1 meter requirement.

Table 5: Results from a representative simulation starting with 0.8 m misalignment.

	Simulation of 2 hour mission
3- σ alignment error	50.7 cm
Time with alignment error < 1 m	100%
Time spent thrusting	4.3%
ΔV	3.4 m/s

To evaluate the reliability of the performance, we ran 1,000 Monte Carlo simulations, starting with a random initial misalignment $\in [0, 0.95]$ meters and SNR $\in [3, 10]$, and running for 1 hour mission time. The results of these simulations are shown in Fig. 10, where we show the distributions of the alignment errors at each time step and the percentage of time spent thrusting in each simulation. Out of the 1,000 simulations, only 1 simulation leaves the ± 1 m bounds (and does so for only 30 seconds). Once settled (after ~ 5 minutes), the 3- σ alignment error for all timesteps is 58.7 cm, the average time spent thrusting is 4.2%, and the average ΔV is 1.7 m/s. The goal in this work was to demonstrate a control scheme that provided a solution to the formation flying problem; as such, we did not focus

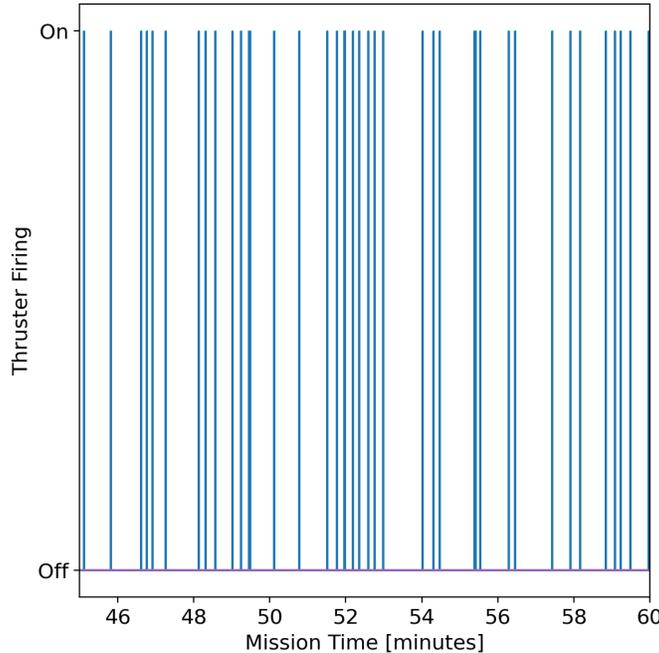


Figure 9: **Simulation** starting with a 0.8 m misalignment. The thruster firing instances are shown as a function of mission time for a fifteen minute window.

on optimizing the duty cycle of observations and left this work for future studies. Allowing a looser control while in the deep shadow could help reduce the number of thruster firings. Nevertheless, 4.2% is very small and leaves the vast majority of the time available for science.

As stated, using the results from these Monte Carlo simulations, we calculate the 3- σ lateral alignment to be 58.7 cm. This is the primary result demonstrating that Milestone 1 has been satisfied.

4.2 Milestone 2: Monte Carlo analysis of reliable sensor transition

A critical stage in the retargeting process is the transition from the acquisition region, where the starshade enters the primary telescope’s field of view, to the observation region, where the pupil sensor takes over and science observations begin. In this transition region, the telescope is ‘blind’ to position measurements and we must rely solely on the UKF estimator without measurement updates. To demonstrate a reliable transition, 1,000 simulations were run starting at random positions with > 100 meters separation and various SNR’s.

The initial conditions of the Monte Carlo analysis are drawn from uniform distributions: the initial misalignment is drawn from a distribution of radii $\in [100 \text{ m}, 150 \text{ m}]$ and a distribution of angles $\in [0, 2\pi)$; the SNR is drawn from a distribution $\in [3, 10]$. Each simulation is run for 1 hour mission time. For 1,000 simulations, 100% successfully tran-

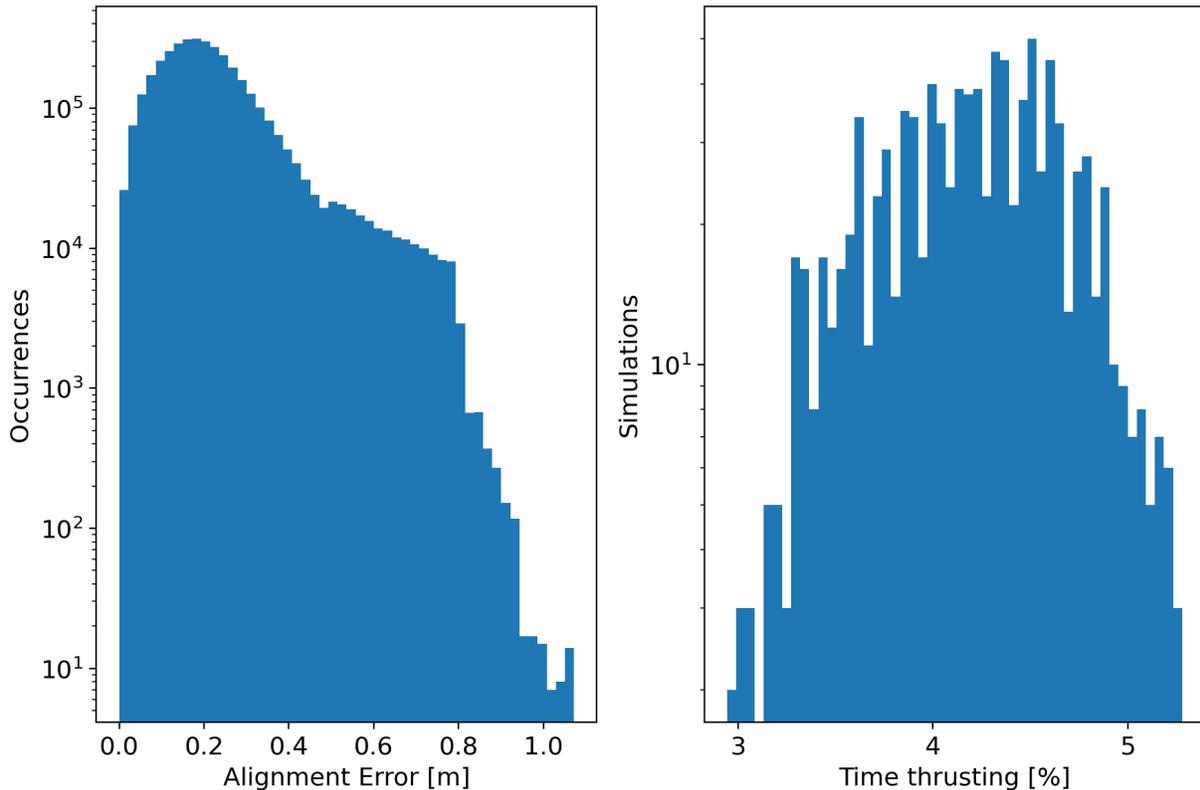


Figure 10: Results of 1,000 Monte Carlo simulations starting in observation region. Left: distribution of alignment errors for every timestep. Right: distribution of the time spent thrusting for each simulation.

sitioned to the observation region and achieved a steady state. The left panel of Fig. 11 shows the distribution of UKF prediction errors at each time step in the 1,000 simulations, separated into the three sensing regions. Also shown in Fig. 11 are the distributions of the time in the shadow (middle panel) and time spent thrusting (right panel). Once reaching a steady state, 96% of the simulations stay within the deep shadow for > 99.7% of the time. The average time spent thrusting is 4.4% and the average ΔV is 2.3 m/s. These Monte Carlo simulations show a reliable transition from retargeting to observation modes and we conclude that Milestone 2 is successfully met.

4.3 Milestone 3: Lab results showing centimeter-level sensing from pupil images

The hardware-in-the-loop test (HILT) is carried out by moving the camera to the desired alignment with the starshade in one-to-one motion with the telescope actual dynamics and the command signal of the control/estimation system. This motion is captured in a set of pupil plane images (shown in Fig. 15) at different instants of time. These images directly

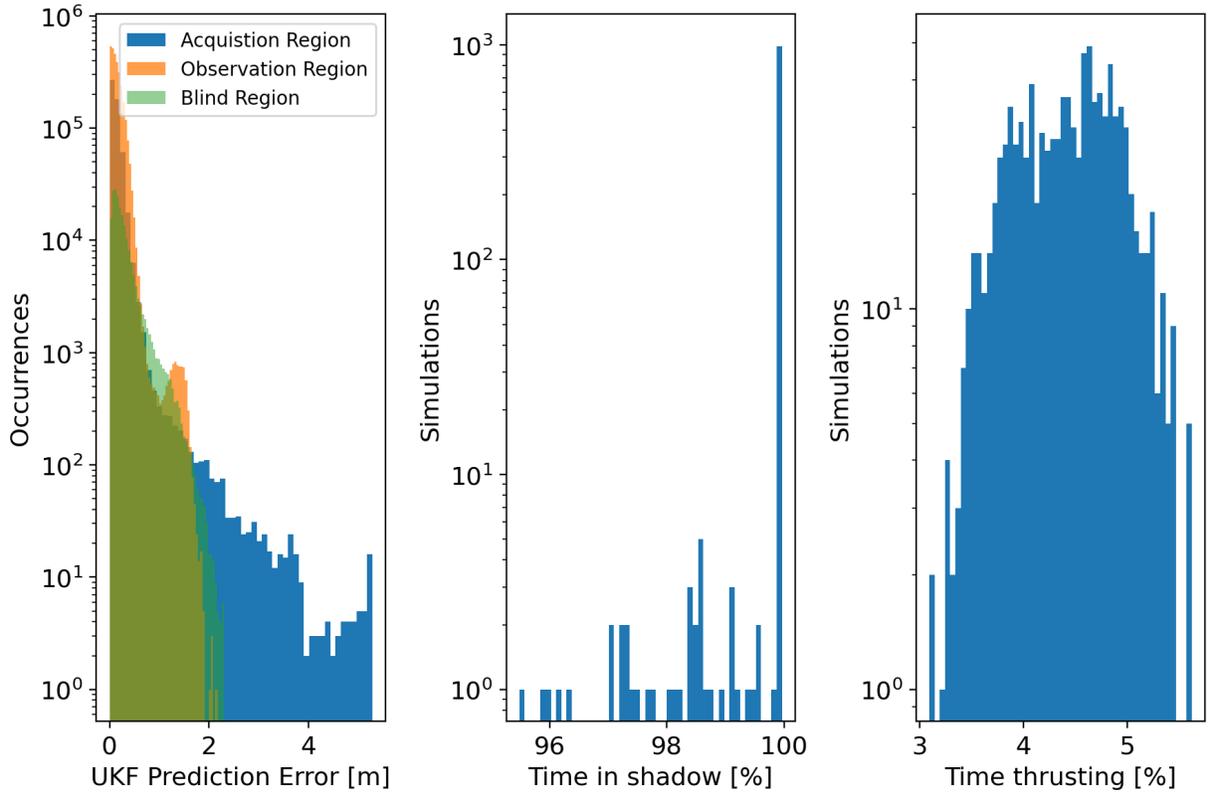


Figure 11: Results of 1,000 Monte Carlo simulations starting in acquisition region. Left: distribution of UKF prediction errors for every timestep; errors are separated by their sensing region. Middle: distribution of the time spent within the ± 1 m deep shadow for each simulation. Right: distribution of the time spent thrusting for each simulation.

sample the diffraction pattern incident on the telescope aperture; obstructions from the telescope’s secondary mirror are artificially imposed on the images (see Fig. 5).

Using a non-linear least squares fit of Eq. (21) to the pupil images, an accuracy of better than 5 cm is achievable for a range of SNR’s. Figure 12 shows the results from extracting positions from 2,000 simulated images. The left panel shows a histogram of the position error (distance between sensed position and true position) for different SNR’s. The SNR is estimated from the signal in the FWHM of the spot of Arago in the center of the diffraction pattern and is quoted per pixel across the spot (FWHM = 10 pixels for a 32×32 pixel image). The quoted SNR is calculated under the assumption that the spot of Arago is not blocked by the secondary mirror; when the central spot goes behind the secondary, the position is extracted from the diffraction rings, which are fainter and thus have a lower SNR. As such, the quoted SNR should simply be used as a metric to compare runs with different exposure times. The right panel shows the median error (and $\pm 1\sigma$ bounds) as a function SNR; even for a SNR as low as 3, the pupil sensor performs well and is able to achieve centimeter level accuracy. The least squares fit is achieved via the Levenberg-Marquardt

algorithm, which converges within 8 function evaluations, taking < 1 milli-second on a standard laptop. It was beyond the scope of this work to map requirements to a flight computer, but we do not believe computation time is a limiting factor.

The least squares algorithm requires an initial guess of the solution, which is usually provided by the UKF state estimation. For the image simulations shown in Fig. 12 (which have no guess provided by the UKF), the initial guess is randomly drawn from a normal distribution centered about the true position with a conservative standard deviation of 20 cm. If no initial guess (or an initial guess of 0 misalignment) is provided, the least squares algorithm will occasionally fail (about 25% of the time) to converge to an accurate solution. However, this does not pose an operational problem, as there will always be an initial guess provided by the UKF. During the alignment phase, images of the starshade’s beacon provides the measurement update to the UKF. Additionally, using the brightest pixel in the image as an initial guess is sufficient to reliably converge to an accurate solution.

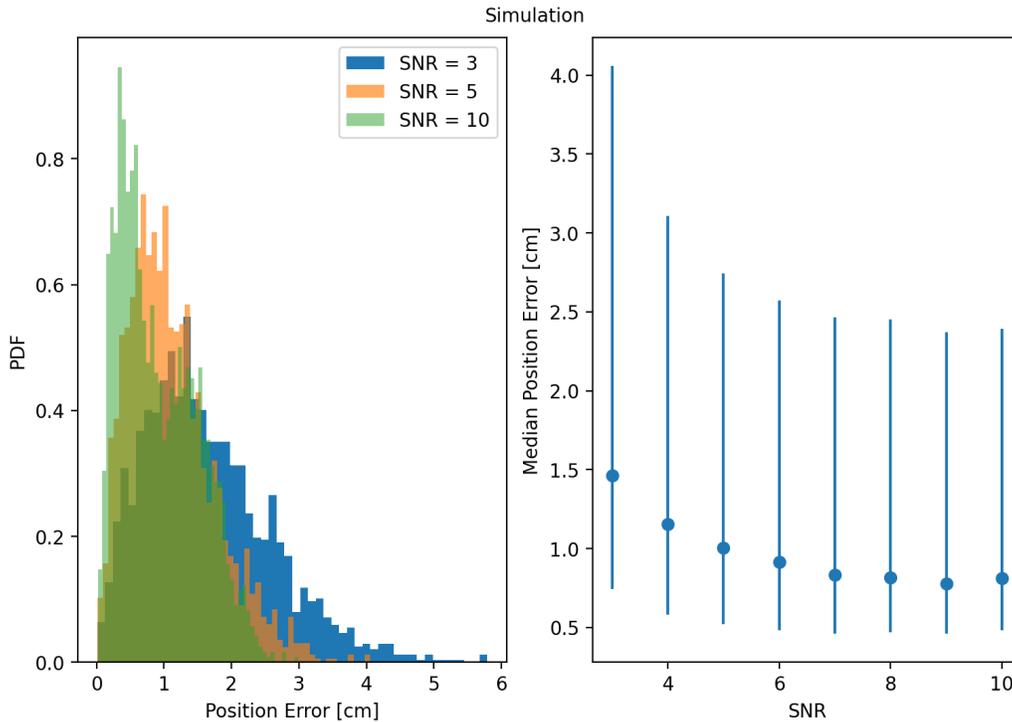


Figure 12: **Simulation** results of position sensing. Left: histogram of position error for different SNR. Right: median error as a function of SNR; lower and upper error bars capture $\pm 1\sigma$ about the median.

Figure 13 shows experimental results for extracting the position from pupil plane images. The left panel shows the distribution of position errors for various SNR’s and number of pixels across the pupil. The right panel shows the median error (and $\pm 1\sigma$ bounds) as a function SNR and number of pixels. The performance is better with more pixels across the pupil as there is more information provided to the least squares algorithm. Across all

SNR’s the $3\text{-}\sigma$ position sensing accuracy is 7.5 cm, well within the criteria set for Milestone 3. The results of Fig. 13 demonstrate we have successfully met Milestone 3.

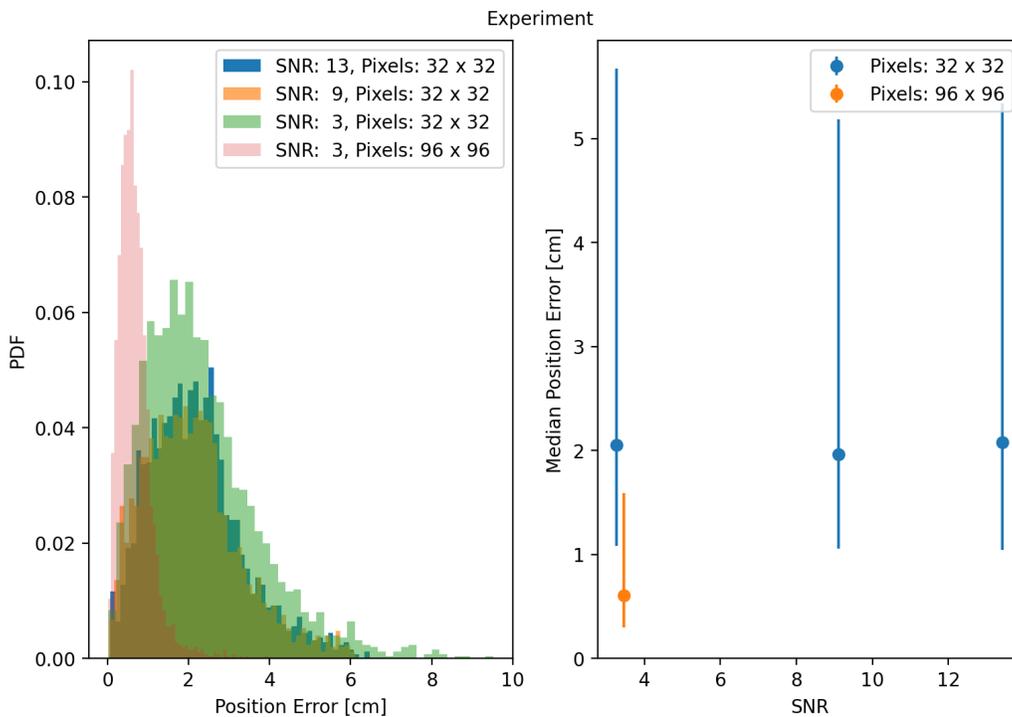


Figure 13: **Experimental** results of position sensing. Left: histogram of position error for different SNR and image size. Right: median error as a function of SNR for different image sizes; lower and upper error bars capture $\pm 1\sigma$ about the median.

4.4 Milestone 4: Lab results showing closed-loop control, with simultaneous high contrast measurements

The starshade testbed was designed to achieve high contrast observations with sub-scale starshades [7]. Figure 14 shows one such high contrast observation from Ref. 7. These observations are taken at $\lambda = 641$ nm, which is in the bandpass for which the starshade is designed to achieve high contrast. In Fig. 14, the outline of the starshade is overlaid, along with a dashed circle which represents the starshade’s inner working angle (the smallest separation to search for exoplanets). Due to non-scalar diffraction effects that arise from the narrow gaps between the starshade petals, 7×10^{-10} contrast lobes aligned with the polarization vector appear inside the inner working, which does not affect the science performance. The peak contrast at the inner working angle is 1.1×10^{-10} and quickly improves with angular distance.

We can validate the formation flying performance by simultaneously achieving high contrast at a wavelength in the starshade’s designed bandpass. Figure 15 shows a progression

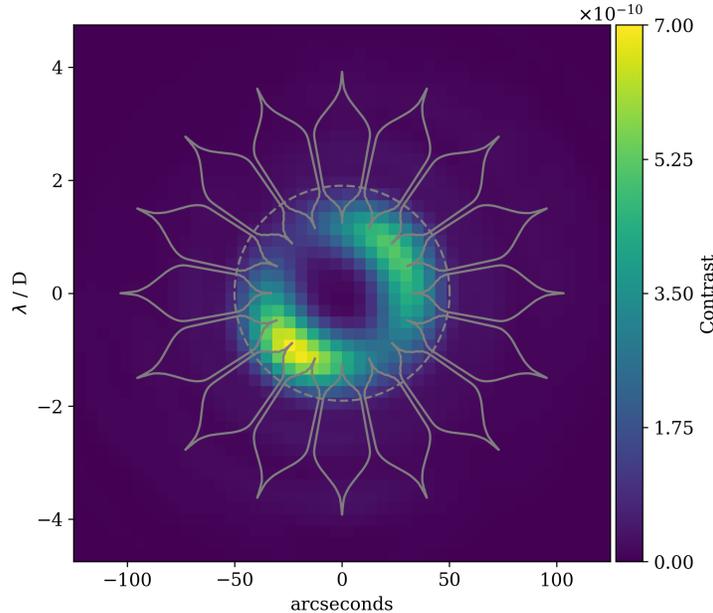


Figure 14: Experimental high contrast observation (focused at infinity; $\lambda = 641$ nm) from Ref. 7. The outline of the starshade is overlaid; the dashed circle represents the starshade’s inner working angle. The contrast at the inner working angle is 1.1×10^{-10} , limited by non-scalar diffraction effects due to small gaps between the starshade petals.

of images as the telescope moves into alignment with the starshade. The left column of Fig. 15 shows pupil plane images at $\lambda = 405$ nm in the guiding bandpass; the right column shows focal plane images at $\lambda = 641$ nm in the science bandpass. When the telescope is far off-axis, the contrast is $1000\times$ higher than the nominal contrast. As the telescope moves into alignment within the deep shadow, the contrast is reduced until it reaches its nominal contrast level. Also note that as time goes by, the telescope center exhibits a transient behavior, oscillating around the center of the diffraction pattern of the starshade, until it finally reaches a steady-state alignment within the prescribed tolerance. The bright lobes emanating from the edge of the starshade (when the telescope is in the deep shadow) are due to non-scalar diffraction arising from the narrow gaps between petals. These lobes are aligned horizontally with the input polarization direction (these observations differ from that of Fig. 14 in that they were conducted with a polarized analyzer in front of the camera, which leads to a shift in the bright lobes and a higher contrast).

Figure 16 shows experimental results of the closed-loop control. This experiment was run with an initial misalignment of 62 m and a SNR of 5. The telescope reaches a stable state in the ± 1 m region within 20 minutes and remains there for the rest of the 2 hour observation time. Figure 17 shows the full range of the offset distance. Once in a stable

state, the $3\text{-}\sigma$ alignment error is 86.6 cm and 4.2% of the simulation is spent thrusting, using a total ΔV of 3.4 m/s. The $3\text{-}\sigma$ alignment error from the experiment (86.6 cm) is higher than that of the simulations (58.7 cm), though both meet the 1 meter requirement. We attribute this difference to additional disturbances in the lab that are not captured in our dynamics model, the largest of which is motion in the diffraction pattern due to atmospheric motion in the testbed. While the controller is able to compensate for the additional disturbances, their presence changes the effect of the tuned gains and leads to a larger (though still sufficient) alignment error. The experimental result of 86.6 cm $3\text{-}\sigma$ alignment error, combined with the simultaneous high contrast measurements shown in Figure 15, demonstrates successful completion of Milestone 4.

5 Conclusion

In this project, we presented the design and implementation of a hardware demonstration on the Princeton Starshade Testbed to validate formation flying sensing and control algorithms while maintaining high-contrast with a starshade. These algorithms include a discrete-time linear quadratic regulator with integral action, an unscented Kalman filter, and high-precision position sensing using a pupil image sensor of the starshade’s diffraction pattern. Simulated and experimental results confirm favorable performance. This includes position error convergence within required tolerances, a robust steady-state to allow science mode activities, and low values of total ΔV during steady-state operation. We have demonstrated a closed-loop lateral alignment of 87 cm ($3\text{-}\sigma$) and position sensing of 7.5 cm ($3\text{-}\sigma$) with a medium fidelity hardware testbed. These results suggest the success of the selected strategy and have successfully demonstrated the completion of the four TDEM Milestones.

5.1 Future Work: Towards new control/estimation schemes

While our studies have been successful, there is still room for improvement in the performance of the control and estimation schemes. The UKF was showing considerable sensitivity to changes in tuning parameters, sometimes leading to unstable relative motion. This problem was mainly due to the performance of the controller, which was influenced by the restrictions imposed by the pulse width modulation on its range of action, and its linear nature. This is not an uncommon problem in linear control theory where the designer assumes the structure of the dynamics of the controlled system is known and can be linearized. However, such systems are often difficult to identify (and sometimes it is impossible), especially if there are nonlinear impositions such as the pulse width modulator or random signals.

We suggest three ideas that may help to alleviate these concerns. First, a more numerically accurate estimator, such as the Cubature Kalman filter (CKF) [9], could be selected. The CKF is derivative free and is considered a theoretical optimum nonlinear filter, approximating the Bayesian filter while preserving second order information. Second, an online neural approach could be embedded in the estimation in order to learn the nonlinear

process in real-time. The estimation algorithm (e.g. CKF) would be in charge of the estimation of the state and, at the same time, the weights of the neural network [22]. Third, instead of the two previous ideas, we could design a controller based on neural networks and reinforcement learning (RL) [11]. Under this approach, an algorithm such as RL is used to update the network weights in real-time in order to minimize a cost usually defined as an error function. The controller has now the ability to adapt/learn online exclusively from the reading of the sensors.

Acknowledgements

This work was performed with the support of NASA Technology Development for Exoplanet Missions (TDEM) award # AWD1004331. The authors would like to thank Phil Willems for support in this work. The authors would also like to thank Douglas Chin and Michael Hauge for help in writing software for hardware control and Michael Galvin for help in construction of the testbed.

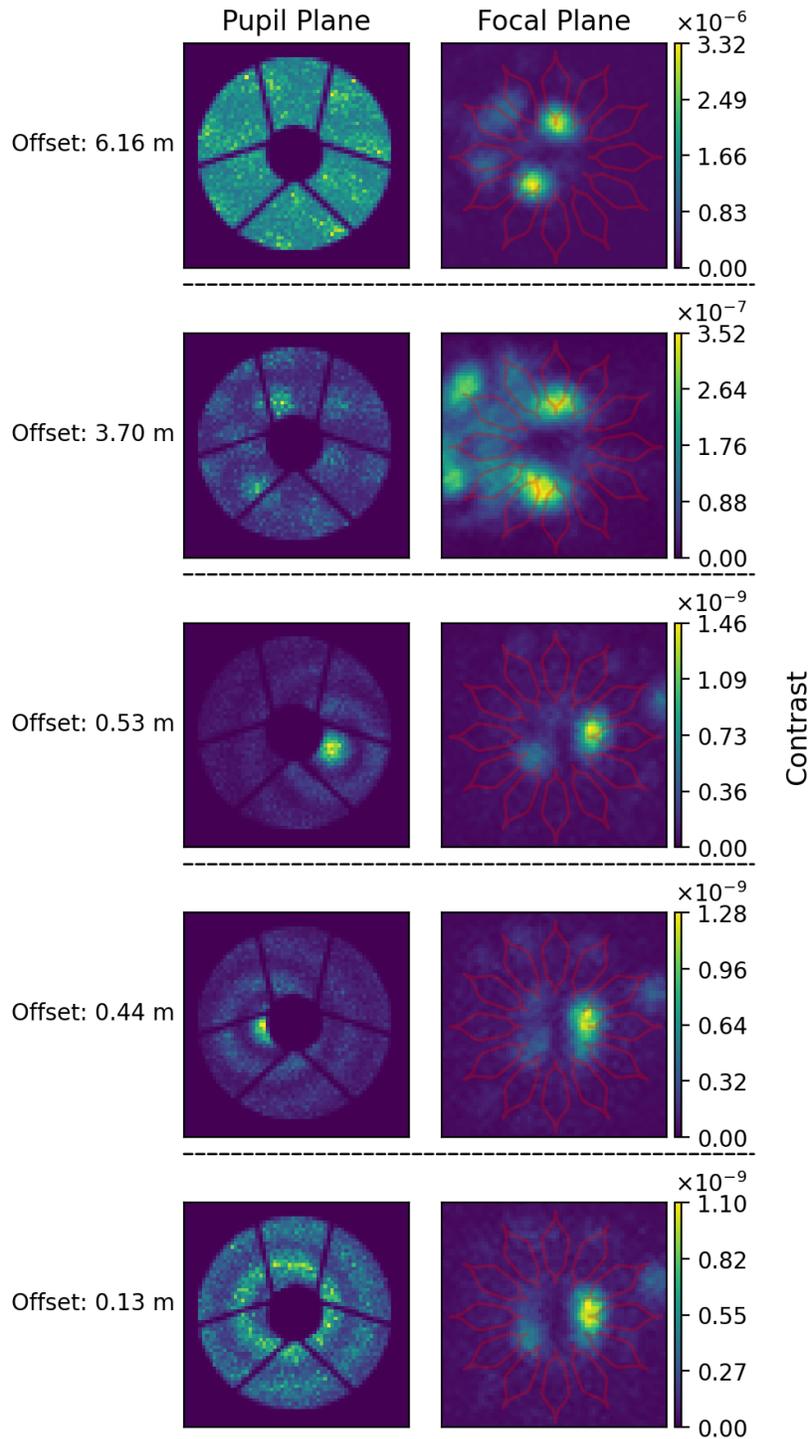


Figure 15: **Experiment:** pupil plane (left column) images at $\lambda = 405$ nm and focal plane (right column) images at $\lambda = 641$ nm for different lateral offsets. The colorbar applies to the right column only. **Note** the change in colorbar scale and much higher contrast level when the starshade is misaligned. These data were taken earlier than those used for the analyses of this report and thus the secondary structure shadow assumes a larger aperture than data shown earlier in the report.

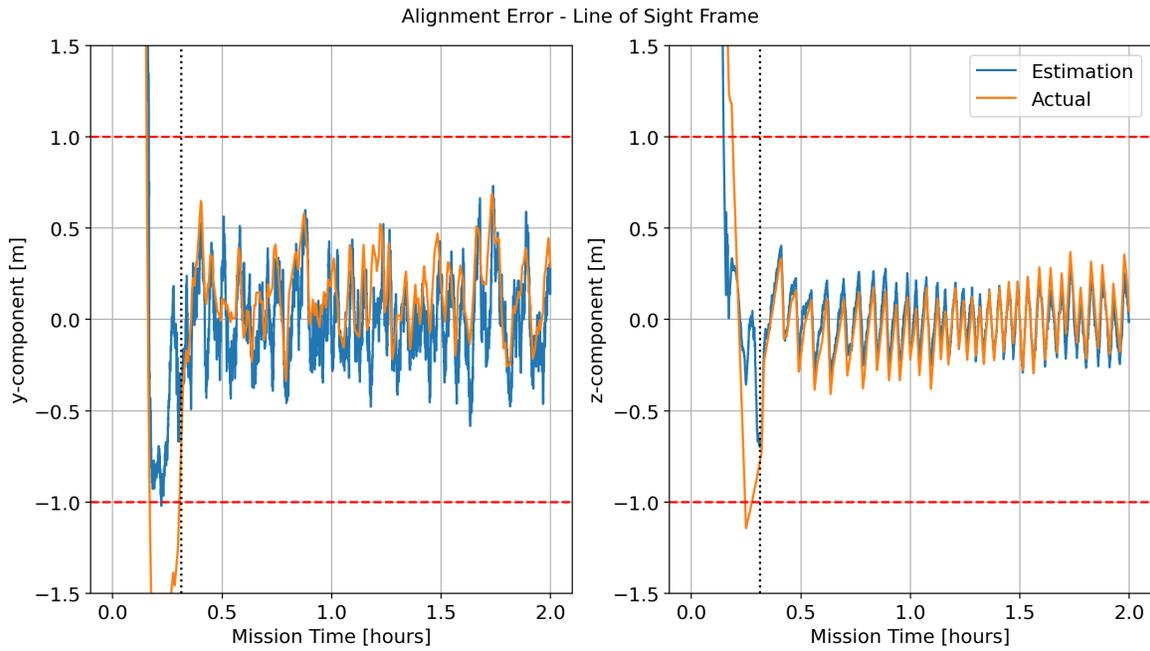


Figure 16: **Experiment** starting with a 62 m misalignment. The alignment error is shown in the two relevant coordinates of the telescope’s line of sight frame.

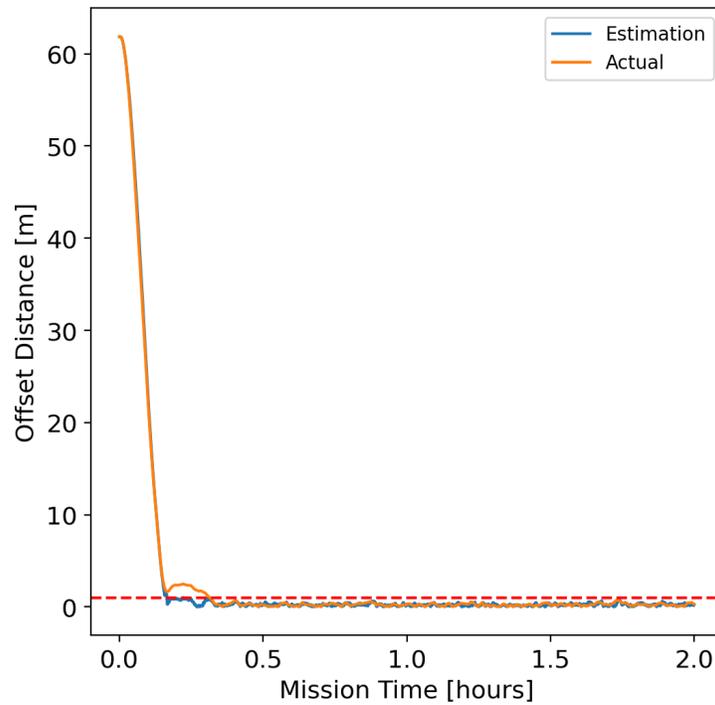


Figure 17: **Experiment** starting with a 62 m misalignment. The offset distance is shown as a function of mission time. The red dashed line marks 1 meter.

References

- [1] Kunjithapatham Balasubramanian, Stuart Shaklan, Victor White, Karl Yee, Richard Muller, Seneca Velling, Simon Vuong, Anthony Harness, and N Jeremy Kasdin. Recent progress in the fabrication of starshade masks for laboratory demonstration of concepts. In *Proc. SPIE*, volume 10698, pages 10698–224, 2018.
- [2] Michael Bottom, Stefan Martin, Eric Cady, Megan C. Davis, Thibault Flinois, Dan Scharf, Carl Seubert, Shannon K. Zareh, and Stuart Shaklan. Starshade formation flying I: optical sensing. *Journal of Astronomical Telescopes, Instruments, and Systems*, 6:015003, January 2020.
- [3] National Research Council. New worlds, new horizons in astronomy and astrophysics. Technical report, The National Academies Press, 2010.
- [4] T Flinois, M Bottom, S Martin, D Scharf, M Davis, and S Shaklan. S5: Starshade technology to TRL5 Milestone 4 Final Report: Lateral formation sensing and control. *Jet Propulsion Laboratory Publications*, 2018.
- [5] Thibault L. B. Flinois, Daniel P. Scharf, Carl R. Seubert, Michael Bottom, and Stefan R. Martin. Starshade formation flying II: formation control. *Journal of Astronomical Telescopes, Instruments, and Systems*, 6(2):1 – 28, 2020.
- [6] Anthony Harness, Stuart Shaklan, Phil Willems, N. Jeremy Kasdin, K. Balasubramanian, Victor White, Karl Yee, Philip Dumont, Rich Muller, Simon Vuong, and Michael Galvin. Optical experiments and model validation of perturbed starshade designs at the Princeton starshade testbed. In Stuart B. Shaklan and Garreth J. Ruane, editors, *Techniques and Instrumentation for Detection of Exoplanets X*, volume 11823, pages 159 – 167. International Society for Optics and Photonics, SPIE, 2021.
- [7] Anthony Harness, Stuart Shaklan, Phillip Willems, N. Jeremy Kasdin, Bala K. Balasubramanian, Philip Dumont, Victor White, Karl Yee, Rich Muller, and Michael Galvin. Optical verification experiments of sub-scale starshades. *Journal of Astronomical Telescopes, Instruments, and Systems*, 7(2):1 – 32, 2021.
- [8] Anthony D. Harness. *High Contrast Astronomy with Starshades*. PhD thesis, University of Colorado, 2016.
- [9] Simon Haykin. *Neural Networks and Learning Machines*. Pearson, 3rd edition, 2009.
- [10] Andrew D. Lewis. Is it worth learning differential geometric methods for modeling and control of mechanical systems? *Robotica*, 25(06):765–777, 2007.
- [11] Frank L. Lewis and Derong Liu. *Reinforcement Learning and Approximate Dynamic Programming for Feedback Control*. Wiley, 1st edition, 2012.

- [12] R. Luquette and R Sanner. Spacecraft Formation Control: Managing Line-of-Sight Drift Based on the Dynamics of Relative Motion. In *3rd International Symposium on Formation Flying, Missions and Technologies*, pages ESA Special Publication SP-654, 2008.
- [13] NASA. Exoplanet Exploration Program (ExEP) Program Elements. <http://exep.jpl.nasa.gov/programElements/>. Accessed March 2014.
- [14] NASA. 2011 NASA Strategic Plan. Technical report, 2011.
- [15] Leonel M. Palacios, Anthony Harness, and N. Jeremy Kasdin. Hardware-in-the-loop testing of formation flying control and sensing algorithms for starshade missions. *Acta Astronautica*, 171:97–105, 2020.
- [16] Andrew Romero-Wolf, Geoffrey Bryden, Sara Seager, N. Jeremy Kasdin, Jeff Booth, Matt Greenhouse, Doug Lisman, Bruce Macintosh, Stuart Shaklan, Melissa Vess, Steve Warwick, David Webb, John Ziemer, Andrew Gray, Michael Hughes, Greg Agnes, Jonathan W. Arenberg, S. Case Bradford, Michael Fong, Jennifer Gregory, Steve Matousek, Jason Rhodes, Phil Willems, Simone D’Amico, John Debes, Shawn Domagal-Goldman, Sergi Hildebrandt, Renyu Hu, Alina Kiessling, Nikole Lewis, Maxime Rizzo, Aki Roberge, Tyler Robinson, Leslie Rogers, Dmitry Savransky, and Chris Stark. Starshade rendezvous: exoplanet sensitivity and observing strategy. *Journal of Astronomical Telescopes, Instrumentation, and Systems*, 7(2), April-June 2021.
- [17] David Di Ruscio. Discrete LQ optimal control with integral action: A simple controller on incremental form for MIMO systems. *Modeling, Identification and Control*, 33(2):35–44, 2012.
- [18] Daniel P. Scharf, Stefan R. Martin, Carl Christian Liebe, Zahidul H. Rahman, Carl R. Seubert, Martin Charles Noecker, and George H. Purcell. Precision formation flying at megameter separations for exoplanet characterization. *Acta Astronautica*, 123:420–434, June 2016.
- [19] S. Seager, N. J. Kasdin, and Starshade Rendezvous Probe Team. Starshade Rendezvous Mission probe concept. In *American Astronomical Society Meeting Abstracts #231*, volume 231 of *American Astronomical Society Meeting Abstracts*, Jan. 2018. <https://smd-prod.s3.amazonaws.com/science-red/s3fs-public/atoms/files/Starshade2.pdf>.
- [20] E.a. a Wan and R. Van Der Merwe. The unscented Kalman filter for nonlinear estimation. In *Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC.*, pages 153–158, 2000.
- [21] Phil Willems. Starshade to TRL5 (S5) Technology Development Plan, 2018.
- [22] Miao Zhiyong, Zhang Yi, Zhao Kun, and Xun Fan. Nonlinear state estimation using neural-cubature kalman filter. *Automatika*, 58(3):347–353, 2017.