

Exoplanet imaging scheduling optimization for an orbiting starshade working with Extremely Large Telescopes

Eliad Peretz^{Ⓜ, a,*} John C. Mather^{Ⓜ, a} Kevin Hall,^{a,b} Lucas Pabarcus^{Ⓜ, a,c}
Clarissa M. Canzoniero,^{a,d} Kelsey Gilchrist^{Ⓜ, a,e} Matan Lieber-Kotz^{Ⓜ, a,f}
Richard Slonaker,^a Wayne H. Yu,^a Steven Hughes,^a Sun Hur-Diaz,^a
Adam Koenig^{Ⓜ, g} and Simone D'Amico^{Ⓜ, g}

^aNASA Goddard Space Flight Center, Greenbelt, Maryland, United States

^bUniversity of Maryland, Department of Astronomy, College Park, Maryland, United States

^cWesleyan University, Middletown, Connecticut, United States

^dPensacola State College, Pensacola, Florida, United States

^eUniversity of Chicago, Chicago, Illinois, United States

^fRice University, Houston, Texas, United States

^gStanford University, Aeronautics and Astronautics Department, Stanford, California,
United States

Abstract. We present optimized observation schedules for a distributed configuration of the Remote Occulter Mission. Accounting for refueling rounds, we show that an Earth-orbiting Remote Occulter could enable up to 158 ground-based observations of 80 exoplanetary targets in a mission lifetime. We develop two target lists, provide exposure time estimates for each potential target star, present an analytic approach for determining target observability, and estimate the cost of station-keeping and retargeting maneuvers required to maintain such a mission. We optimize the mission observation schedule over these cost and science delivery estimates using deterministic and metaheuristic optimization methods with varying degrees of operator intervention and conclude by assessing mission profile sensitivity to both isolated and accumulated cost and design perturbations. © The Authors. Published by SPIE under a Creative Commons Attribution 4.0 Unported License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JATIS.7.2.021213](https://doi.org/10.1117/1.JATIS.7.2.021213)]

Keywords: starshade; imaging; exoplanet; schedule optimization; remote occulter.

Paper 20123SS received Aug. 12, 2020; accepted for publication Jan. 12, 2021; published online Feb. 12, 2021.

1 Introduction

The National Academies of Science delivered a consensus study report for the exoplanet science community in 2018 that concluded that a starshade-based or coronagraph-based direct imaging mission is the best path to measure the reflected light spectra of temperate planets around Sun-like stars.¹ This finding is further supported by prior reports.² The Remote Occulter Mission (R-O), a proposed Earth-orbiting starshade used in conjunction with the Extremely Large Telescope (ELT), has been shown to enable high-quality direct images and spectra of temperate planets around Sun-like stars,³ and is one promising avenue toward this goal. Recent studies of the R-O mission have further defined its operational concept, including determining its time-dependent observable sky regions, evaluating potential targets, and assessing available exposure times.^{4,5,6}

Starshade operational regimes are fuel-limited because they are required to maneuver between targets (here termed retargeting) and maintain line-of-sight (LOS) alignment with their paired telescope during observations (here termed station-keeping). Previous work⁷ in optimizing starshade science deliverables has focused on space-telescope coupled configurations in L2 halo orbits, requiring sophisticated techniques to define stars with maximal completeness curves leading to estimating overall exoplanet yield.^{8–10}

*Address all correspondence to Eliad Peretz, eliad.peretz@nasa.gov

We focus here on developing adaptable mission planning tools, which will inform subsequent developments to the R-O’s mission architecture. This paper develops algorithmic processes to find and present optimized observation schedules for the R-O mission over various operational conditions. Two realistic target lists are formulated from established star catalogs (Sec. 2.1). High-fidelity exposure times are estimated for each candidate target using the Starshade Imaging Simulation Toolkit for Exoplanet Reconnaissance (SISTER) software package (Sec. 2.2).¹¹ Incorporating a dynamic time-bound observable sky region confirms the dates and times that each target can be viewed, allowing us to estimate the available exposure times given observational constraints (Sec. 2.3). We then calculate the associated required station-keeping cost for each target star as a function of declination (δ) and the relative right ascension (α) to local zenith during observation (Sec. 2.4). Finally, we estimate the retargeting cost associated with maneuvering between targets (Sec. 2.5).

With the operational constraints defined (Sec. 3.1), we introduce a high-level overview of and optimized schedules produced by both the “augmented deterministic” (Sec. 3.3) and “hybrid metaheuristic” optimization methods (Sec. 3.2). For each, we display the sky trajectory and accumulated delta-V costs for both station-keeping and retargeting maneuvers. We conduct a holistic analysis of both methods’ sensitivity to key parameters: station-keeping delta-V, retargeting delta-V, and retargeting time (Sec. 4.1) in addition to mission lifetime and target list choice (Sec. 4.2), followed by a sensitivity analysis of these parameters’ accumulated effects (Sec. 4.3). Finally, we provide conclusions and set the ground for future work (Sec. 5).

2 Methodology

2.1 Establishing a Target List

We select target stars from the main sequence that lie within 30 pc of Earth with conservative temperatures ranging between 3000 and 6500 K. We collect these stars from three individual catalogs: The 2nd Data Release from the Global Astrometric Interferometer for Astrophysics (Gaia DR2), Tycho,¹² and Hipparcos.¹³

All stars must pass specific astrometric and photometric requirements. These requirements follow a conservative approach by removing erroneous stars as well as binary star systems, which possess poor astrometric solutions¹⁴ and finalize by filtering any nonmain sequence stars and any duplicates.

To ensure candidate exoplanets lie sufficiently outside the starshade inner working angle (IWA), we examine the IWA for the Remote Occulter, which is typically defined for starshades and coronagraphs to be the angular distance from the center of the imaging plane where 50% throughput is achieved.¹⁵ We perform a parallel analysis by finding the SNR of Earth with different angular separations from the Sun, as can be seen in Fig. 1. The resulting curve reveals that for the Remote Occulter, an angular separation of 35 milliarcseconds (mas) yields a normalized SNR value of 50%.

Previous work on solar habitable zones has estimated that the habitable region expands from around 0.95 AU¹⁶ inner radii and extend to outer radii of 1.7¹⁷ or 2.4¹⁸ AU for a Sun-like star. We then calculate the angular separation of a comparably scaled (1.7 AU) habitable zone for each star in the potential target list using the following relation:

$$\theta_{\text{HZ}} = \tan^{-1} \left(\frac{\sqrt{\frac{L_*}{L_\odot}}}{r_{\odot-*}} \right), \quad (1)$$

where θ_{HZ} is the inner habitable zone angle and $r_{\odot-*}$ is the distance of the target exoplanetary system from the Sun, given that the habitable zone for a given star can be established by scaling to its luminosity using $\sqrt{\frac{L_*}{L_\odot}}$, where L_\odot is the luminosity of the Sun and L_* is the luminosity of a target star.¹⁹ To ensure each target HZ is nearly fully covered, we select target stars with habitable zones θ_{HZ} larger than the Remote Occulter IWA θ_{IWA} or 35 mas.

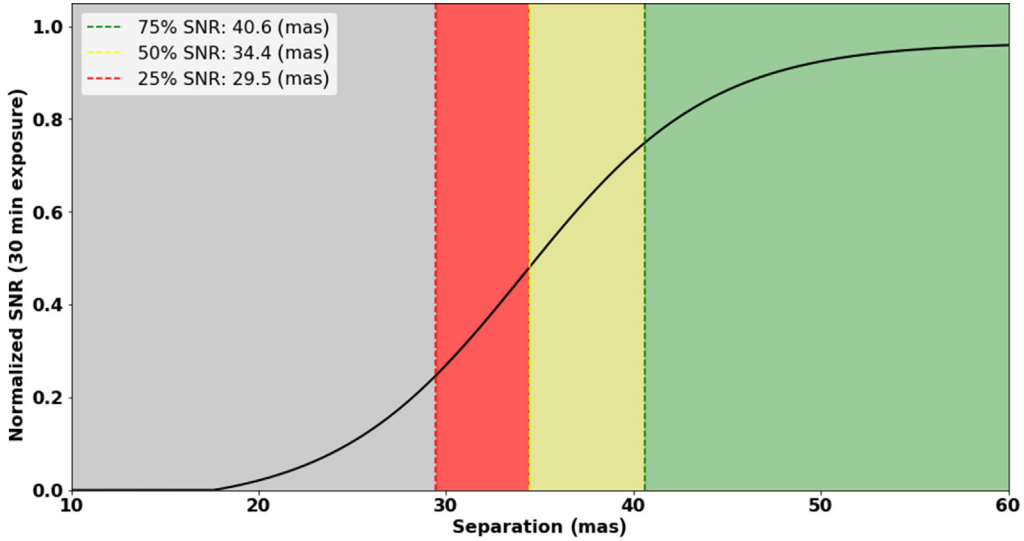


Fig. 1 The imaging performance of the Remote Occulter is examined by plotting the normalized SNR against the angular separation of the planet and host star. The SNR is for Earth at different angular separations from the Sun at a distance of 10 parsecs and normalized against the SNR achieved at 58 mas, which has a value of 156. The allotted exposure time is 30 min for all data points. There are four regions identified that correspond to angular separations that yield different normalized SNR values. The green region corresponds to angular separations that yield normalized SNR values larger than 75%. The yellow region yields values between 50% and 75%, and the red region yields 25% and 50%. Notice that the Remote Occulter IWA at 50% is 35 mas. The gray region corresponds to angular separations that are too close to the host star to attain strong SNR values.

Finally, we establish a second target list of stars identified by near-future missions including the Nancy Grace Roman Space Telescope (NGRST),²⁰ the Large Ultraviolet Optical Infrared Surveyor (LUVOIR),²¹ and the Habitable Exoplanet Imaging Mission (HabEx).¹⁵ We then impose the same constraints used for generating the Gaia-sourced target list. This results in 525 and 275 target stars in the Gaia-sourced and mission-sourced target lists, respectively. These two target lists will represent high-density and low-density star distributions for use in this study's subsequent analyses.

2.2 Exposure Time Calculations

We calculate exposure times using SISTER, which allows us to perform high-fidelity simulations of exoplanet system imaging for a wide variety of telescope, detector, and exoplanet system parameters. We use a submodule developed to support the R-O mission to model observation performance with each target star.

Each simulation performed by SISTER yields multiple sets of output data for various parameters including: total electron counts from the planet S_p , simulation-generated noise, and background from local zodiacal dust, star leakage, and exo-zodiacal dust S_b , number of pixels N_p , and detector read noise R_N . We further define Δt as the total exposure time and Δt_f as the exposure time for each frame. These outputs enable calculations of the signal-to-noise ratio (SNR) for a given simulated planet. This method is described in further detail by Hildebrandt et al.¹¹ and Peretz et al.² We define SNR as

$$\text{SNR} = \frac{S_p}{\sqrt{S_p + S_b + N_p R_N^2 \frac{\Delta t}{\Delta t_f}}}. \quad (2)$$

For each star, we place an earth-like planet within the habitable zone defined in Sec. 2.1 and require for each target star an SNR > 10 . For this study, we impose a minimum 30-min

Table 1 The distribution of different spectral types and their range of distances. The exposure time τ for imaging and SNR correspond to the median distance, and the stellar luminosity L_* used for each computation corresponds to the median.

Spectral type	Number of stars	Min (pc)	Median (pc)	Max (pc)	τ (min)	L_*	SNR
F	144	3.5	22.0	30.0	30	3	39.1
G	324	1.35	22.5	30.0	30	1	19.1
K	52	1.20	13.3	24.8	30	0.4	28.3
M	5	3.28	5.8	19.11	30	0.1	82.5

observation time for targets that result in shorter time exposures (which is nearly any target on the list) by that imposing a conservative limit. Table 1 displays parameters for the Gaia target list including, for each spectral type: the number of targets; the minimum, median, and maximum ranges; and, for the median range, the resultant exposure time, and SNR.

2.3 Time-Bound Observable Sky Region

Per Peretz et al. (2019),⁴ there are four primary constraints that drive the available observation window for each position on the celestial sphere over a given night. First, the telescope zenith-target star angle θ_{z-*} must be under 60 deg to enable R-O observations. Second, we constrain the Sun altitude to be at least 18 deg below the horizon, a night threshold that limits diffuse skylight at low-target altitudes and defines the initial and final zenith hour angles. This corresponds to a minimum zenith-Sun angle $\theta_{z-\odot}$ of 108 deg. Third, the Sun must be behind the Earth-facing starshade plane by at least 1 deg to minimize solar glare, meaning the starshade-Sun angle $\theta_{s-\odot}$ must be smaller than 89 deg. This produces an unobservable area centered on the telescope zenith hour angle at each sidereal midnight, which may be reduced by tilting the starshade normal away from the telescope-target star LOS by an angle θ_{s-*} up to 30 deg.

Each celestial coordinate's (RA, DEC) nightly available exposure times are determined by the constraints above and are calculated throughout the mission period. Python-based scripts have previously been successfully developed to iteratively and numerically compute the observable sky window throughout the year,⁴ but they require run times on the order of days and computational expenses in GBS of data. To rapidly compute observation availability in our optimization algorithms, we derive analytical functions for the geometry-driven observable bounds that fully define the overall observable contour, before verifying their equivalence to prior iterative solutions.

Each target's accumulated observation time is the difference between the initial and final times that it crosses the telescope zenith-centered local observable cone. By treating the projected zenith right ascension's angular speed, ω_z as constant, rotating the same rate as the earth, ω_{\oplus} , we make the following mathematical approximation, where the Earth's axial tilt ϵ_{\oplus} is sufficiently small (~ 24 deg):

$$\omega_z = \frac{d\alpha_z}{dt} = \frac{d}{dt} \arctan[\cos(\epsilon_{\oplus}) \tan(\omega_{\oplus} t)] \approx \frac{2\pi}{\text{day}}. \quad (3)$$

Available observation time is then directly proportional to the angular distance between a desired target's right ascension and the outer initial/final zenith boundaries, marked by the right and leftmost red and blue curve halves in Fig. 2. Indeed, the numerically calculated contour in Fig. 2 cutaway A exhibits linearly dependent observable times as a function of RA (α) for a fixed DEC (δ). Therefore, the overall observable contour at a given date may be directly derived from the observation window's boundaries. The regions in which the red and blue bounded observable domains overlap may be observed for the entire night as defined here, and available observation times for targets linearly decrease to zero as a function of the angular distance from the inner boundary.

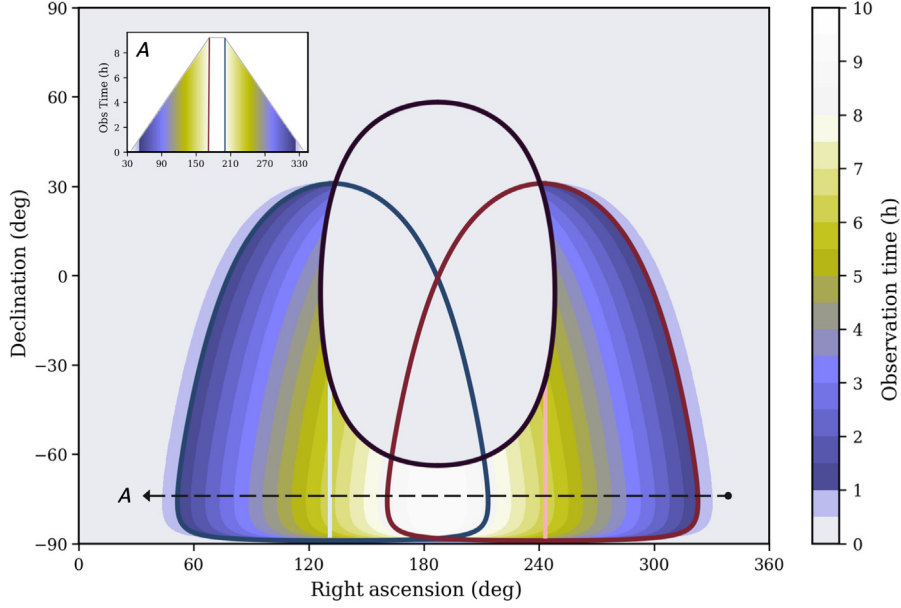


Fig. 2 Ecliptic-projected peak numerical observation duration and matching analytic bounds. The red/pink and blue/light blue curves, respectively, reflect the projected observable region bounds and local zenith right ascension at the start and end of night, and are offset to indicate regions that may be observed for at least 1 h. The central black ellipse indicates the sky regions that cannot be observed due to starshade–Sun reflectance. Inset A shows a cross section of available observation times at a declination of -75 deg.

By way of reference frame and coordinate system conversions, these boundaries are analytically derived in ecliptic coordinates. Due to the intrinsic symmetry produced by assuming a negligible impact of earth axial tilt on projected angular speed in our region of interest, they are functionalized in a halved form (around midnight’s zenith) where bounding RA is a function of DEC:

$$t_{\text{window}}[t] = \frac{1}{\omega_z} \left(2\pi - \arccos \left\{ \frac{\cos(\theta_{z-*-\min}) - \sin(\lambda_T) \sin[\delta_{\odot}(t)]}{\cos(\lambda_T) \cos[\delta_{\odot}(t)]} \right\} \right), \quad (4)$$

$$\alpha_{s-\odot-\max}[\delta, t] = \alpha_{\odot}(t) + \arccos \left\{ \frac{\cos(\theta_{s-\odot-\max} + \theta_{s-*-\max}) - \cos\left[\frac{\pi}{2} - \delta_{\odot}(t)\right] \cos\left(\frac{\pi}{2} - \delta\right)}{\sin\left[\frac{\pi}{2} - \delta_{\odot}(t)\right] \sin\left(\frac{\pi}{2} - \delta\right)} \right\}, \quad (5)$$

$$\alpha_{z-\max}[\Delta t, \delta, t] = \alpha_{\odot} + 2\pi + \omega_z(\Delta t - t_{\text{window}}) - \arccos \left[\frac{\cos(\theta_{z-*-\max}) - \cos\left(\frac{\pi}{2} - \lambda_T\right) \cos\left(\frac{\pi}{2} - \delta\right)}{\sin\left(\frac{\pi}{2} - \lambda_T\right) \sin\left(\frac{\pi}{2} - \delta\right)} \right], \quad (6)$$

where $(\alpha_{\odot}, \delta_{\odot})$ is the position of the Sun in celestial coordinates, t is the epoch, λ_T is the telescope latitude, and Δt is the exposure time. By overlaying our previously developed numerical solutions’ contour and these observable curves in Fig. 2, we see that both are in agreement. Precomputing this window for each target allows us to more rapidly predict observable regions in submillisecond time, reducing the computational constraints of optimizing the observational schedule. The center of the initial and final observable regions (marked in pink and light blue) may also be used to determine best case off-zenith observation angles for station-keeping cost calculations. Targets inside of these bounds are passed by the zenith over a night, whereas targets outside have best case observation off-zenith angles in proportion to their right ascension

bound's relative angular distance. A detailed derivation of these analytical equations is presented in [Appendix C](#).

2.4 Station-Keeping Cost: Chemical Delta-V

Current designs of the Remote Occulter mission concept place the spacecraft in a highly elliptical orbit whose path crosses a target of interest as viewed from the ground telescope near its apogee of $\sim 200,000$ km. To achieve proper contrast, we require the R-O to remain within ± 6 m of the LOS from the ground telescope to the target star, a relaxed requirement in comparison to other starshade-based missions.^{21,22} To estimate the delta-V cost associated with matching the starshade velocity with a desired target, here termed station-keeping maneuvers, we use the analytical approximation of delta-V costs for a short observation developed by Koenig and D'Amico including their defined reference frames.²³ To do so, we first examine the differential angular accelerations between the ground-based observatory and the R-O in the defined L-frame:

$$a_{s-*} = \omega_{\oplus}^2 r_{\oplus-T} \cos(\lambda_T) \begin{bmatrix} -\cos(\delta_*) \cos(\omega_{\oplus} t) \\ -\sin(\omega_{\oplus} t) \\ \sin(\delta_*) \cos(\omega_{\oplus} t) \end{bmatrix}, \quad (7)$$

where λ_T is the telescope latitude, δ_* is the declination of the astronomical target, and ω_{\oplus} is the Earth's angular velocity. The acceleration of the starshade a_s is given by the following equation:

$$a_s = -\mu \frac{r_{\oplus-s}}{r_{\oplus-s}^3}, \quad (8)$$

where μ is the standard gravitational parameter and $r_{\oplus-s}$ is the Earth-centered starshade vector. We then calculate station-keeping delta-V costs for a period Δt in seconds of observation using the following equation:

$$\Delta v = \omega^2 r_{\oplus-T} \cos \lambda_T \int_{t_0}^{t_0+\Delta t} \sqrt{\sin^2(\omega_{\oplus} t) + \sin^2(\delta_*) \cos^2(\omega_{\oplus} t)} dt, \quad (9)$$

where δ_* is the declination of the target star and t_0 is the time at the start of the observation measured from the time, in which the target is closest to zenith. This equation can be resolved using single-step Euler integration to approximate short observations less than an hour:

$$\Delta v = \omega_{\oplus}^2 r_{\oplus-T} \cos(\lambda_T) \Delta t \sqrt{\sin^2(\omega_{\oplus} t_c) + \sin^2(\delta_*) \cos^2(\omega_{\oplus} t_c)}, \quad (10)$$

where t_c is the epoch in the middle of observation such that $t_c = t_0 + \frac{\Delta t}{2}$. We find that observations generally require more starshade acceleration or deceleration for targets that are away from zero declination and/or are low on the horizon along the ecliptic, termed "relative right ascension." This trend is invariant with respect to telescope latitude, whose cosine simply scales the overall delta-V costs. Figure 3 presents these delta-V costs as a function of target declination and relative right-ascension for the ELT, located at Cerro Armazones.

2.5 Retargeting Cost: Electrical Delta-V

The starshade observations described in this paper occur near orbital apogee to minimize LOS station-keeping during observation, and therefore, require orbital reconfiguration to change targets. Preliminary Collocation Stand-Alone Library and Toolkit (CSALT)²⁴ optimization simulations assessing the operational cost of these low thrust transfers have found a close linear correlation between intertarget angular distance and transfer cost for targets within 30 deg.³

We then calculate transfer costs during optimization as a CSALT-informed scaling of intertarget angular distance θ , calculated as follows:




Fig. 3 Contour plot of station-keeping delta-V costs (m/s) for an hour-long observation as a function of target declination and relative right ascension (target star RA relative to the telescope zenith RA at the time of the observation). Green lines denote 15 m/s increments.

$$\cos(\theta) = \cos(\alpha_2 - \alpha_1) \cos(\delta_1) \cos(\delta_2) + \sin(\delta_1) \sin(\delta_2), \quad (11)$$

where initial and posttransfer targets are located at (α_1, δ_1) and (α_2, δ_2) . Based on previous system studies, we conservatively approximate $\frac{\Delta V}{\Delta \theta} = 30$ m/s/deg for optimization,²³ and further set target transfer operations to take a number of days equal to the degree distance between targets, or 5 days for targets closer than 5 deg. We review transfer cost and time assumption accuracy and associated observation delivery sensitivity in Sec. 4.

3 Schedule Optimization

3.1 Defining Frame and Goals

We set the science orbit delivery date to January 2035 and baseline a maximum of four subsequent refueling operations over 7 years. Based on the previous analyses by Koenig and D’Amico (2019), we use a “distributed” 19 Ton system architecture composed of a 8000 kg Starshade and 5000 kg servicer, with a combined 6000 kg of propellant budget. The Starshade system independently maintains line of site during observation using a conservative 280 s Isp chemical propulsion system and docks with the servicer for orbital reconfiguration using its 2800 s Isp electric propulsion system. The retargeting servicer also refuels the starshade, whose tank volume is set to contain 1000 kg of chemical propellant. Since the starshade is expected to have a dedicated launch, no “operational” delta-V is expended to place the starshade in its initial science orbit.

Within the allocated 6000 kg of spacecraft propellant, various relative chemical and electrical fuel masses may be selected, as calculated in Appendix C. In Fig. 4, we explore the effect of varying relative propulsion system masses on the delivered mission delta-V, from which we optimize the selected system architecture for delivered number of observations. A range of electrical delta-V values is provided to account for variation in relative starshade/servicer propellant consumption throughout a mission.

We find that optimized star sequences select targets averaging between 6 deg and 8 deg for our primary target list depending on location in the sky. Assuming a uniformly distributed set of targets and a total of 2 h of observations at ± 10 deg declination and 15 deg off zenith—corresponding to half of observations occurring in the outer half of each zenith domain, we can approximate the typical chemical–electrical cost ratio, $\frac{\overline{\Delta V}_C}{\overline{\Delta V}_E}$ as ideally between 3 and 3.5, or 2 to 8

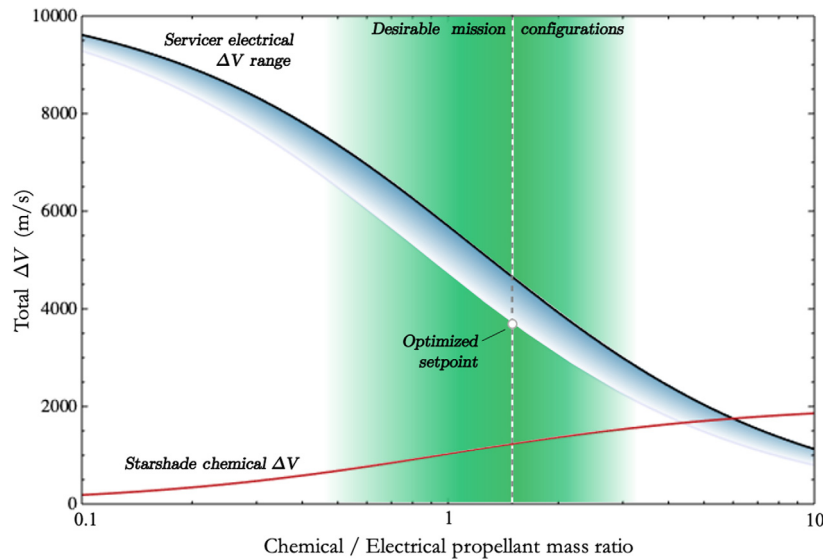


Fig. 4 Total available interrefuel chemical and electrical ΔV as a function of relative spacecraft propellant mass distribution. The range of expected mass ratios and specific optimized value used in this paper is indicated by the green contour and dashed gray line, respectively.

in extremis. These ranges are bounded by the green set of mass ratios in Fig. 4. Following an optimization of this ratio in our target sequencing efforts, we conservatively set initial and refuel delta-V budgets to 3685 electrical m/s for retargeting and 1325 chemical m/s for station-keeping.

After the starshade independently maintains LOS for an observation, both servicer and starshade must reconfigure their orbits and dock. We reflect this operational cost by doubling the starshade station-keeping costs defined in Fig. 10. In the case of repeated observations of the same target, we estimate that an additional 10 m/s of electrical propellant is used by the servicer to counteract orbital perturbations. We expect refuelling operations to occur over up to three orbit periods. While this could hypothetically occur during a target transfer, we set refuels to occur during idle time in this paper.

Both methods—(1) hybrid metaheuristic and (2) augmented deterministic—utilize the tools developed in Sec. 2 as input parameters, establishing a baseline for performance analysis. Using each method, we produce a chronological list of the observed targets along with their associated station-keeping and retargeting delta-V costs for their optimized schedule.

3.2 Hybrid Metaheuristic Method

The hybrid metaheuristic approach combines problem-informed local search and solution refinement methods with a globally tuning evolutionary algorithm. The evolutionary algorithm's evolved state vector guides the target selection process by differential weighting of various cost and heuristic considerations to globally propagate sequences from all initial targets until propellant is consumed before selecting the top 10% of candidate sequences.²⁵ These sequences are then refined with a k-opt algorithm that exploits transfer cost graph symmetries, and are perturbed by the incorporation of time-dependent observable windows and off-angle station-keeping costs.²⁶ A high-level overview is shown in Fig. 5. This method ultimately produces a globally optimized solution whose priority is total observations. It, therefore, maximizes the use of all available mission resources. For instance, if additional time is provided, the algorithm will use the time to select lower-cost stars by waiting to observe when there are lower relative right ascension costs.

Using the analytical observability bounds presented in Sec. 2.3, we predetermine the times of year that every target may be observed, shown left for an example star in Fig. 6. This and the evolved target priorities inform each subsequent target selection, with the additional ability to wait for a more optimal time within the window. A visual of these features is provided in Sec. 6,

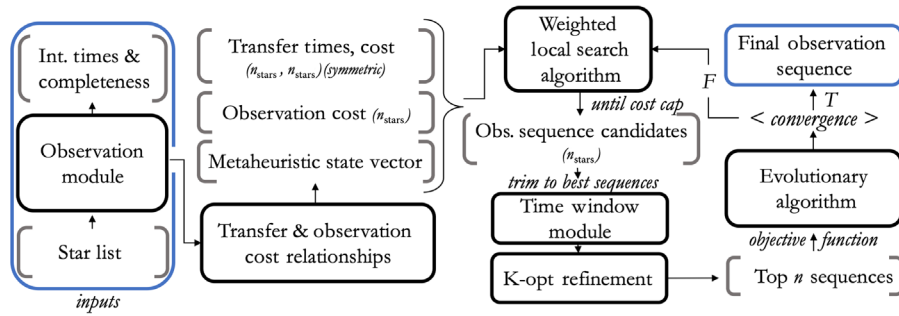


Fig. 5 Schedule optimization procedure used for the hybrid metaheuristic method. Note that the star target list, analytical bounds of the observable sky window, and integration times are precalculated.




Fig. 6 Time-dependent observable window bounds used for stars at -30 deg declination (L). Example target transfer priorities for a specific initial target star overlaid on network priority weights, marked in blue and gray, respectively (R).

whereas a full discussion of the target prioritization, evolution, and timing logic is presented in [Appendix 6.4](#).

The results of the hybrid metaheuristic method for the parameters described in [Sec. 3.2](#) are displayed below. Over a mission lifetime of almost 7 years from December 3, 2035 to July 1, 2042, 60 unique targets are observed. [Figure 7\(a\)](#) displays the cumulative retargeting and station-keeping costs as a percentage of total mission delta-V, whereas [Fig. 7\(b\)](#) displays the cumulative number of targets, all as a function of mission lifetime.

[Figure 8](#) shows the optimized observation schedule displayed on the relevant portion of the sky. For the full list of targets displayed in this figure, along with the associated delta-V costs and the date of observation, see [Fig. 17](#) in [Appendix A](#).

3.3 Augmented Deterministic Method

The augmented deterministic approach employs an intuitive global search followed by a triply branching greedy selection to choose targets and determine observation times.²⁷ We find the observable sky window’s bounds numerically for a full calendar year instead of using analytic heuristics described in [Eqs. \(4\)–\(6\)](#).

We intuitively select a target based on the principles discussed in [Sec. 2](#). The station-keeping delta-V cost is calculated for all stars within a set arc-length radius of the initial target. The retargeting delta-V costs are also calculated between the initial target and all identified stars. We forecast the cost of future maneuvers for these targets four more and select the first target in the minimal-cost sequence. Two observation dates and delta-V costs are recorded within the first valid 30 days. The final observation schedule is set when either the available fuel after refueling is consumed, or the set mission timeline of 7 years is met. [Figure 9](#) provides a

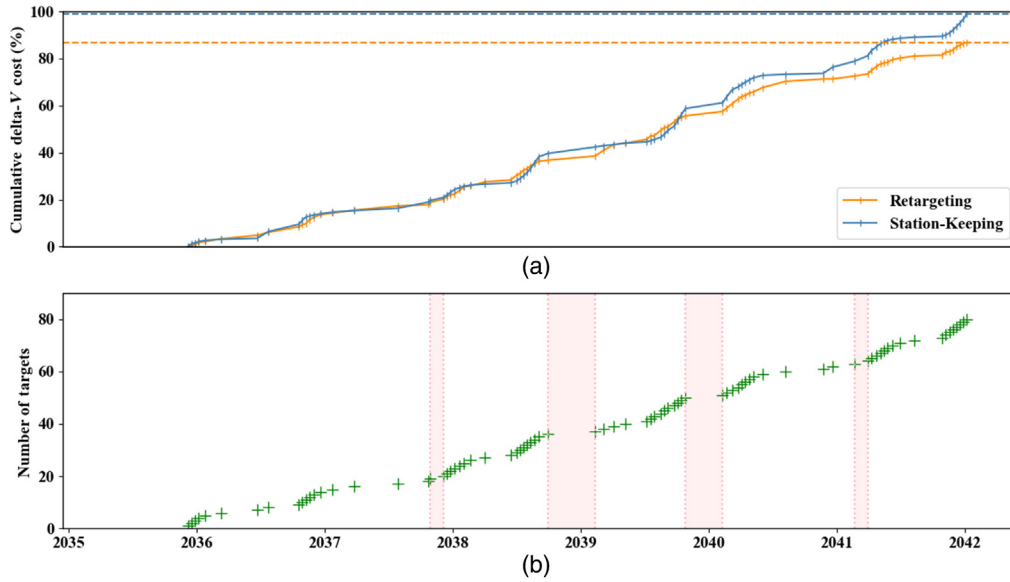


Fig. 7 (a) The cumulative retargeting (orange) and station-keeping (blue) delta-V costs for the hybrid metaheuristic method over time as percents of total mission capacity. Each tick represents an individual target and all costs associated with it. (b) Target count (marked as +) over the mission duration and indicates planned refueling windows in light red.




Fig. 8 Results from the metaheuristic target sequence calculations are mapped onto the sky. Targets are connected with black lines to form a sequence and are colored chronologically by their observation date, spanning the mission timeline from 0 to 7 years, as indicated on the colorbar. Sequenced targets are sized based on the predicted station-keeping costs for these dates. A green circle around a sequence target indicates that a refueling occurs before that target is observed. Unobserved targets are visible in gray around the sequence. Unused targets are sized according to predicted available station-keeping costs.

high-level overview of this method, whereas Fig. 10 shows a visualization of the selection process. This method ultimately produces a regionally unconstrained local solution and is optimized for mission lifetime. This solution diverges from the optimal sequence as the number of targets increases.²⁸ If additional time is provided, the greedy nature of the selection process in the algorithm will not use this time to add lower-cost stars to the end of the schedule. It will instead prefer to maintain the lower lifetime of the mission.

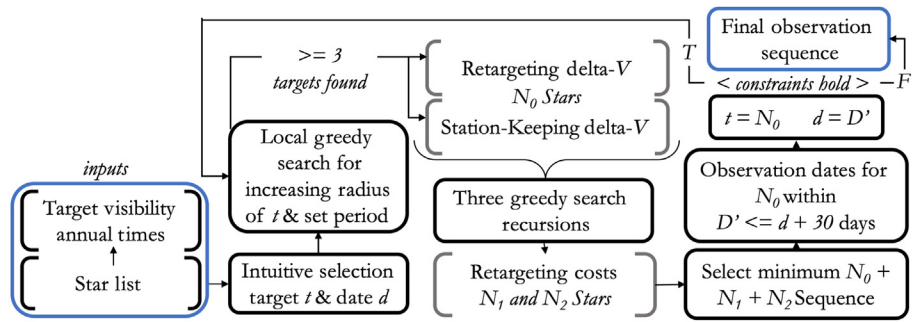


Fig. 9 Schedule optimization procedure used for the augmented deterministic method. Note that the star target list and the numerically determined values for the annual visibility times are precalculated.

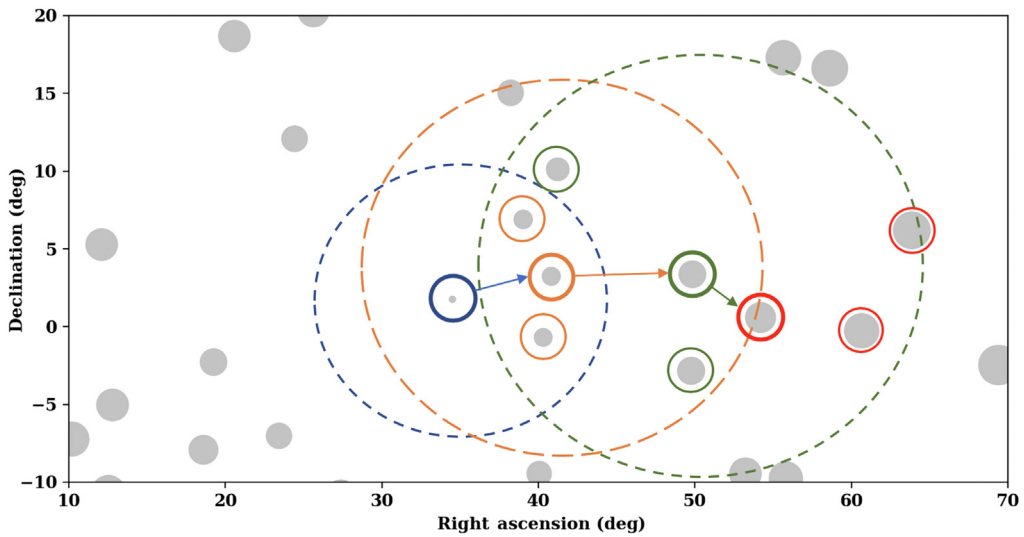


Fig. 10 A visualization of the deterministic method, the size of each target reflects the sum of the minimal station-keeping costs and retargeting costs from the blue target star t . A triply branching greedy selection provides an estimation of the ideal observation sequence by selecting orange target N_0 with minimal future costs. Note that the search radius increases for stars later in the predictive sequence as the neighborhood star density decreases because three potential targets must be identified at minimum.

The results of the augmented deterministic method for the parameters described in Sec. 3.1 are displayed below. Over a mission lifetime of almost 7 years from January 1, 2035 to August 2, 2041, 69 unique targets are observed. Figure 11(a) displays the cumulative retargeting and station-keeping costs as a percentage of total mission delta-V, whereas Fig. 11(b) displays the cumulative number of targets, all as a function of mission lifetime.

Figure 12 shows the optimized observation schedule displayed on a portion of the sky. For the full list of targets displayed in this figure, along with the associated delta-V costs and the date of observation, see Fig. 18 in Appendix B.

4 Method Sensitivity Analysis

The results displayed in Secs. 3.2 and 3.3 are optimized based on previously outlined parameters for delta-V costs, mission fuel capacity, mission lifetime, transfer time between targets, and target list selection as defined in Secs. 2 and 3.1. However, the mission concept is likely to change during development due to the inevitable evolution of the science requirements and engineering capabilities.

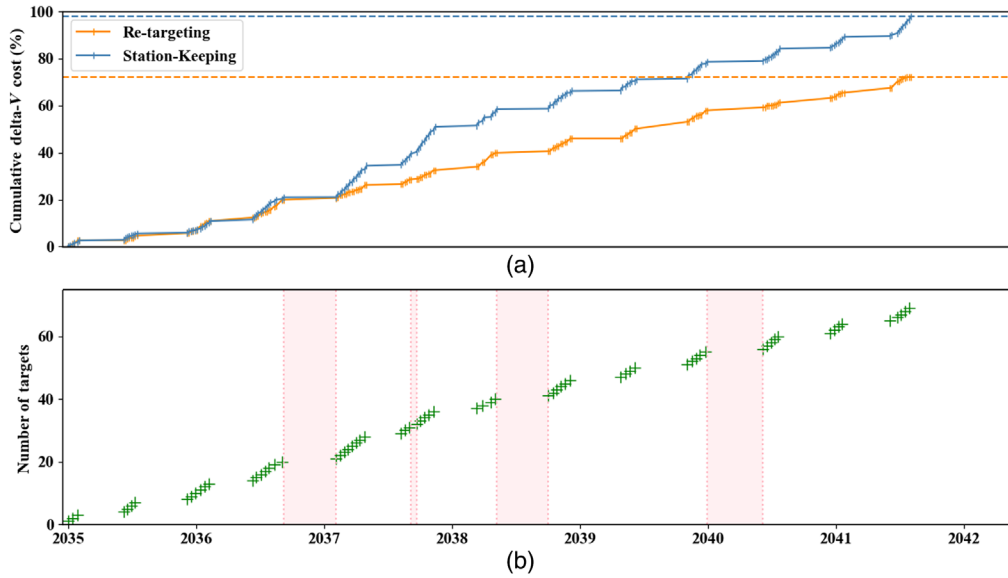


Fig. 11 (a) The cumulative retargeting (orange) and station-keeping (blue) delta-V costs for the augmented deterministic method over time as percents of total mission capacity. Each individual observation is marked with a tick. (b) Target count (marked as +) over the mission duration and indicates planned refueling windows in light red.




Fig. 12 Results from the augmented deterministic target sequence calculations are mapped onto the sky. Targets are connected with black lines to form a sequence and are colored chronologically by their observation date, spanning the mission timeline from 0 to 7 years, as indicated on the colorbar. Sequenced targets are sized based on the predicted station-keeping costs for these dates. A green circle around a sequence target indicates that a refueling occurs before that target is observed. Unobserved targets are visible in gray around the sequence. Unused targets are sized according to predicted available station-keeping costs.

To assess the impact of operational cost uncertainties, we evaluate the the individual and combined impact of varied retargeting costs, station-keeping costs, and transfer times on delivered targets in Sec. 4.1. We further evaluate changes to mission design parameters including target list and mission lifetime in Sec. 4.2, before reviewing net optimistic and pessimistic scenarios in Sec. 4.3.

4.1 Operational Sensitivity

To calculate sensitivity to delta-V cost parameters, we generate optimized target sequences under varying retargeting and station-keeping costs by multiplying their respective cost function results by factors between 0.5 and 2.0 before optimization. In addition to each-other (chemical and electrical), delta-V change sensitivities are also expected to be cross-correlated with the extent to which time is a mission constraint. We assess this impact by further varying coefficient factors on our current model of transfer time over the same range.

These uncertainties and their combined effects on target delivery correspond to a three-dimensional space where each point has an optimal target sequence. In Fig. 13, we sample the three base planes of two combined operational factors in this space using the metaheuristic approach, as well the individual effects on both optimization methods. We find that for isolated perturbing factors, target delivery is most sensitive to station-keeping cost, followed by retargeting cost and then retarget time, which only produces a slight impact on the order of 5 to 10 observed targets for doubled transfer duration. We find that scaling retargeting times has a relatively small impact on target delivery over the range of scaled retargeting costs [Fig. 13(b)], but a noticeable effect in combination with increased station-keeping costs when the time factor exceeds the station-keeping factor [Fig. 13(a)], which may indicate a “phase transition” of whether chemical fuel or mission duration is primary mission constraint. Excluding time factors, target delivery is almost entirely sensitive to station-keeping (chemical) costs, rather than retargeting (electric) costs [Fig. 13(c)].

We expect the globally optimal solution space to generally be smooth. The noise in optimized results (suboptimality) can then be attributed to the metaheuristic method’s selection priorities being optimized for all cost factors being unity and the augmented deterministic method utilizing a local search. Gaussian-filtered contours are presented to indicate overall trends and approximate this global space.

4.2 Target List and Mission Lifetime Sensitivity

The Gaia-sourced target list containing 525 targets offers a relatively high-density distribution of stars. We now compare the effectiveness of the methods when using a lower-density list of potential exoplanet host stars compiled from the NGRST, LUVOIR, and HabEx missions and display these results in Fig. 14.

For the augmented deterministic method, 69 stars are observed using the Gaia catalog, and 55 stars are observed using the ExoCat list.

The hybrid-metaheuristic method yields 80 stars using the Gaia list, and 53 stars with the ExoCat list. Where target density drops 48% between Gaia and ExoCat, we expect the distribution of intertarget distances to approximately shift by the root of density, a 28% drop. This is more closely proportional to the observed shift in target sequences, though there are obviously many other dominating factors at play, including characteristics of each target list, the observable window and “using up” desirable available stars.

To examine sensitivity to mission lifetime, we generate optimized schedules for mission lifetimes of 4, 7, and 10 years and then extract the number of targets. The 4-year mission has one less refueling available (3 total), whereas the 10-year mission has one more (5 total). Results are shown in Table 2.

As seen in Table 2, as mission lifetime increases, the number of targets observed also increases; however, these increases in mission lifetime do not scale in direct proportion with increases in the number of observed targets. These diminishing returns are due to optimization constraints, including distribution of target stars in the sky and the observable window.

4.3 Sensitivity to Accumulated Effects

We want to explore how the final number of targets is impacted by all parameter sensitivities when combined and therefore, draw conclusions beyond the sensitivity trends of individual parameters. We start with the nominal case corresponding to the results discussed in Sec. 3 and then apply $\pm 50\%$ factors in the following order: (1) station-keeping, (2) retargeting, (3) transfer

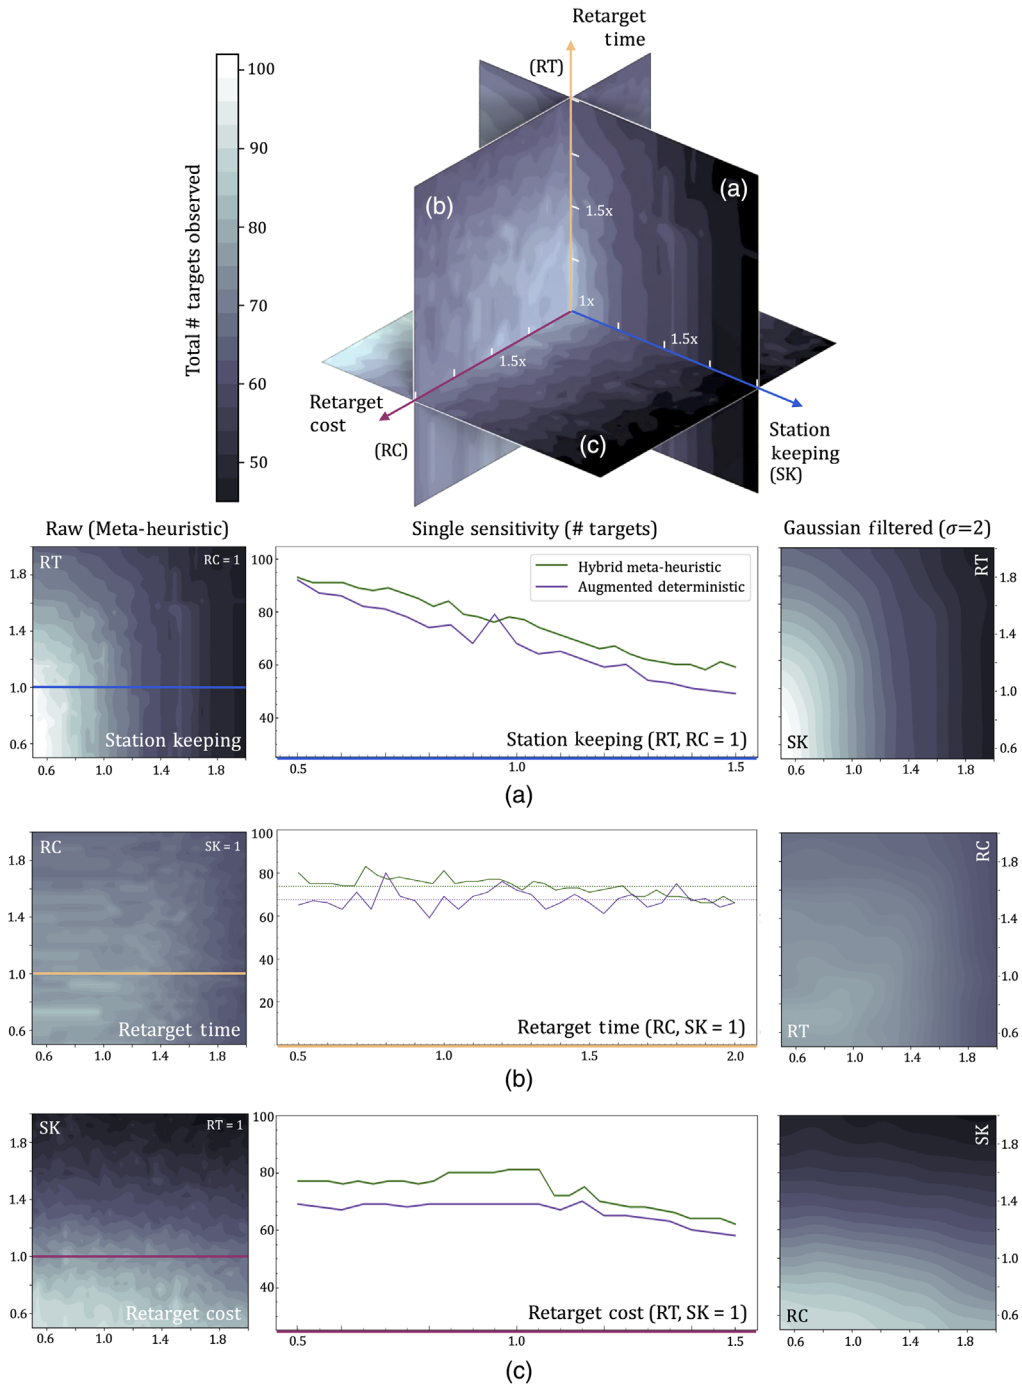


Fig. 13 Contour maps of the optimized target uncertainty space. Axes are operational cost factors, where retargeting time (orange/RT), retargeting costs (purple/RC), and station keeping costs (blue/SK) are uniformly scaled from their baseline values. Each row in (a)–(c) corresponds to the same sensitivity plane of two combined factors. The center column shows the target delivery sensitivity to isolated cost factors for both hybrid-metaheuristic and augmented deterministic methods.

time, and (4) mission lifetime. This exercise generates net optimistic and pessimistic scenarios where parameter costs are reduced and increased by 50%, respectively, and accumulated. We then generate an optimized schedule for each step taken and present the extracted target results in Fig. 15, starting from the left where the nominal solution is displayed, all the way to the right where all effects are accumulated.

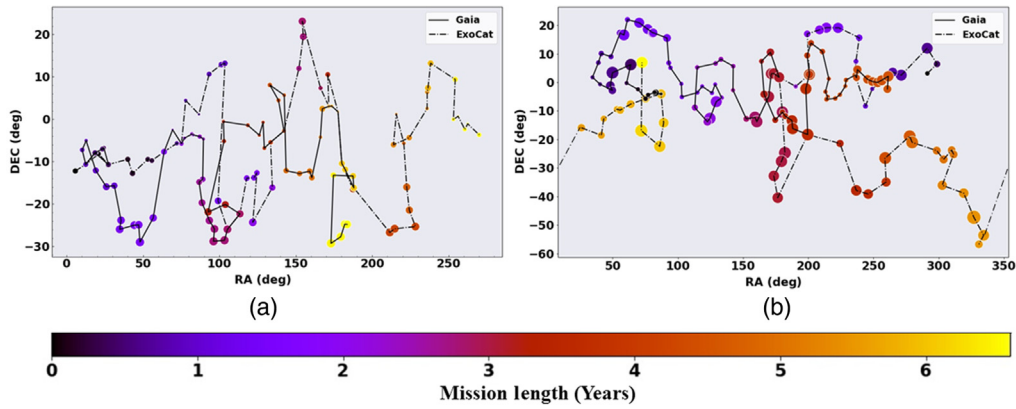


Fig. 14 The observation sequence built using both the Gaia and ExoCat. (a) The sequence corresponds to the augmented deterministic method and (b) the sequence corresponds to the hybrid-heuristic method. Each star is colored by the time and it is observed during the mission lifetime.

Table 2 The number of targets observed for each schedule optimization method. There are three numbers listed corresponding to 4, 7, and 10 years. The number of refuel missions, RM, is provided as well.

Method	4 Years	RM	7 Years	RM	10 Years	RM
Augmented deterministic	48	3	69	4	82	5
Hybrid metaheuristic	50	3	80	4	91	5

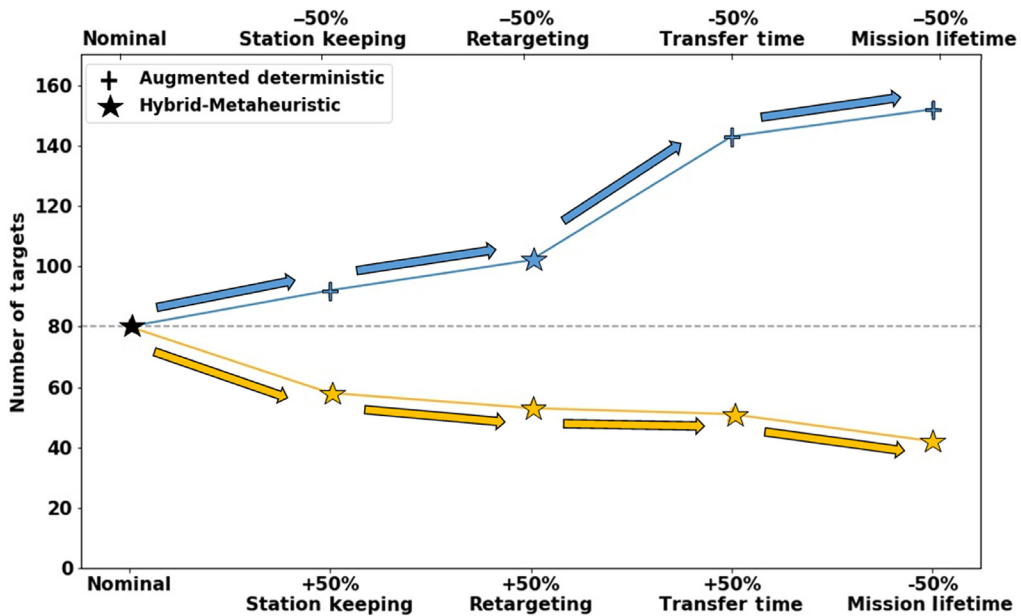


Fig. 15 The total number of observed targets for a given sensitivity parameter. As each plot moves along the horizontal axis, an additional optimization parameter has a factor of 1.5 or 0.5 applied. The plot on top corresponds to 0.5, and the bottom plot corresponds to 1.5. The mission lifetime parameters receive 0.5 and 1.5 on the top and bottom, respectively.

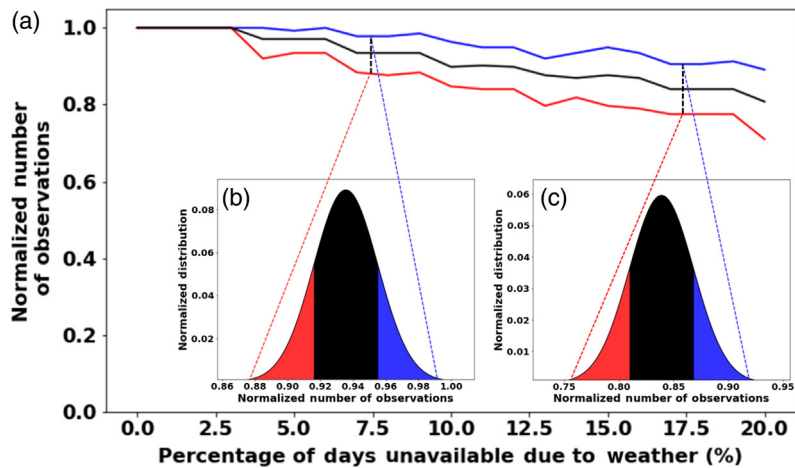


Fig. 16 Probability distributions for the number of observations possibility with a percentage of days unavailable due to weather and with no adaptation. (a) The maximum, minimum, and mean normalized number of observations is plotted as a function of days unavailable as blue, black, and red, respectively. The percentages of days unavailable range from 0% to 20%. We show the normalized probability distribution for two cases of percentages. The two cases B and C correspond to 7.5% and 17.5%, respectively. The minimum and maximum value for each distribution is connected to the plot above by dashed lines with matching colors.

For each accumulated sensitivity parameter applied, the results are displayed for the optimization method that yielded the largest number of targets. For the orange curve (pessimistic), the hybrid-heuristic method had the larger number of targets with 80, 58, 53, 51, and 42. The blue curve (optimistic) had two values come from the hybrid-heuristic method, and the final three come from the augmented deterministic method, at 80, 92, and 102, 143, and 152, respectively.

Weather effects could also play a role, yet it is worth noting that in the case of bad weather or seeing conditions, the R-O could maintain the same target orbit and observe again within an orbit period with a minimal ΔV (<10 m/s) cost. For the E-ELT case considered in this paper, Cerro Armazonas was chosen among many reasons due to its ideal observation conditions, with a level of over 89% clear nights (Mauna Kea 76%, in comparison) a year, with median seeing of 0.7 arc sec and coherence times of 4.5 ms.²⁹ In our optimization process, we add a 10 m/s cost and baseline two observation orbits regardless of having a bad/good night to be more conservative. For the optimized case, only 160 nights out of 7 years are required for observations, which provide the mission with ample opportunity to wait, retarget, or reoptimize the observation schedule and largely maintain its performance.

In Fig. 16, we evaluate how the number of potentially observed targets is impacted by bad observation conditions for the mission lifetime in a worst case scenario affording no adaptation to poor weather conditions. We conduct a statistical study, in which we evaluate mission observation profiles within a range of 0% to 20% and provide the profiles for two cases that have 7.5% and 17.5% nights that are unobservable.

5 Conclusions and Future Work

In this study, we explore schedule optimization for exoplanet imaging for a starshade working in conjunction with a ground-based telescope. We confirm that the Earth-orbiting R-O Mission can conduct over 160 observations of over 80 targets within the given mission constraints including: the observable window, retargeting, and station-keeping ΔV costs, transfer time between targets, mission lifetime, and fuel capacity. We find that both augmented deterministic and hybrid metaheuristic methods can provide reasonable schedules to be used.

We conducted sensitivity tests to examine the robustness of the proposed solutions to changes in the singular base parameters, followed by an analysis of the accumulated effects as they stack, and concluded that our schedules are minimally sensitive. These sensitivity tests can be used to inform changes to the mission architecture such as the electrical-to-chemical fuel ratio, and mass distribution between the starshade and the servicer. This, in turn, will inform the operation concept of the mission later since the target list will dynamically evolve as scientific discoveries occur and changes to models predicting exoplanet demographics will improve.

Future work will include further reliance on CSALT optimization for an accurate assessment of retargeting costs for each specific target star and associated orbit. Both the hybrid metaheuristic and augmented deterministic methods can be further refined to accommodate additional selection parameters. For example, the methods can be upgraded to ensure specific targets are included in a generated sequence or that a large number of observation could occur for a specific target. A follow up study to determine the operation concept ultimately establishing its scientific yield as well as a mechanical and thermal studies to further develop the R-O starshade design.

6 Appendices

6.1 Appendix A: Variable Table

The variables used in this study are described in Table 3.

Table 3 Description of variables used in this study.

Variable	Description	Subscript	Description
α	Right ascension	\odot	Sun
δ	Declination	s	Starshade
λ	Latitude	$*$	Target star
ϵ	Axial tilt	T	Telescope
Δt	Observation time	z	Telescope zenith
ω	Angular speed	\oplus	Earth
L	Luminosity	IWA	Inner working angle
θ	Relative angle	IHZ	Inner habitable zone

Prior work evaluating the geometric constraints that drive the Remote Occulter's observable window⁴ has made use of reference frame and coordinate conversions to derive the following relations of solar ecliptic position, telescope latitude, and target position that must be satisfied in order for a target to be observed at a given point in time:

$$\cos(\theta_{z-\odot-\max}) \geq \sin(\lambda_T) \sin[\delta_{\odot}(t)] + \cos(\lambda_T) \cos[\delta_{\odot}(t)] \cos[\alpha_{\odot}(t)], \quad (12)$$

$$\cos(\theta_{z-*-\max}) \leq \cos\left[\frac{\pi}{2} - \delta_z(t)\right] \cos\left(\frac{\pi}{2} - \delta\right) + \sin\left[\frac{\pi}{2} - \delta_z(t)\right] \sin\left(\frac{\pi}{2} - \delta\right) \cos[\alpha - \alpha_z(t)], \quad (13)$$

$$\begin{aligned} &\cos(\theta_{s-\odot-\max} + \theta_{s-*-\max}) \geq \\ &\cos\left[\frac{\pi}{2} - \delta_{\odot}(t)\right] \cos\left(\frac{\pi}{2} - \delta\right) + \sin\left[\frac{\pi}{2} - \delta_{\odot}(t)\right] \sin\left(\frac{\pi}{2} - \delta\right) \cos[\alpha - \alpha_{\odot}(t)], \end{aligned} \quad (14)$$

where the solar coordinates α_{\odot} and δ_{\odot} are determined by the equation of time for a given time of day. As Eq. (12) describes the relative zenith-Sun angle, the difference between the two times, in which the equality is satisfied on a given day is, therefore, the "length of night" defined in Eq. (4) as $t_{\text{window}}[t]$, where t is the day-rounded epoch.

A result may be obtained for the outermost starshade-Sun boundary in two dimensions on the ecliptic-projected sky through a similar process of forcing equality in Eq. (14) and solving for right ascension as a function of Declination and the combined tilted starshade-Sun constraint $\theta_{s-\odot} + \theta_{s-*}$. This produces one half of the unobservable cutoff region and is defined in Eq. (5). The full unobservable region, 15 may then be produced by mirroring around the local zenith's right ascension at midnight, which marks the observable window's line of symmetry:

$$\alpha_{\max-s-s}[\delta, t] \leq \alpha_{\text{unobservable}} \leq 2\alpha_s - 2\alpha_{\max-s-s}[\delta, t]. \quad (15)$$

When the zenith right ascension at midnight on a given day is offset by the Earth's rotation to the start of night, $\frac{t_{\text{window}}}{2}$ and applied to Eq. (13) at equality, the trailing edge (leftmost boundary) of the projected zenith observable cone at the start of night is found using Eq. (6). Offset by t_{window} , this curve represents the observable cone's trailing edge at the end of night. Both leading edges may then be produced by similarly mirroring around $\alpha_{\text{Sun}} + \frac{\pi}{2}$ as in Eq. (15), therefore, fully defining all observable boundaries. Available observation times and the best case off-angle observation costs are calculable from these boundaries using the processes described in Sec. 2.3.

6.4 Appendix D: Metaheuristic Algorithm Details

We utilize a combination of various heuristic parameters to inform target selection, where each works together into an overall $\frac{\text{value}}{\text{cost}}$ priority function, which is greedily maximized. Each consideration is individually weighted by coefficients of the evolved state vectors, (γ, β) where the length of γ is the number of candidate stars and β is variable dependent on which additional mission parameters are optimized, e.g., spacecraft mass ratio. To account for the value of subsequent observations in selecting a follow-on target, we incorporate a cost-weighted network summation of all possible transfers from each initial star. This is represented by Eq. (16), where each potential transfer from initial star i to candidate star j is scaled according to an inverse exponential of its value/cost ratio and summed. Additional evolved factors (γ_i) are added to tune the decay of network cost-relevance and prioritize prior evolutionary stages' known regions of interest. We generally define value (completeness) as Q , cost as C and priority as P :

$$P_{i_net} = \sum_{j=0}^{n_{\text{stars}}} \gamma_i e^{\frac{\beta_2 C_{j_obs} + \beta_3 C_{ij_trans}}{-\theta_k}}, \quad (16)$$

$$\hat{P}_{i_net} = \frac{P_{i_net}}{\max(P_{net})}, \quad \hat{C}_{obs} = \frac{C_{i_obs}}{\max(C_{obs})}, \quad \hat{C}_{transfer} = \frac{C_{i_transfer}}{\max(C_{ij_transfer})}, \quad (17)$$

$$P_{i_weighted} = \hat{p}_{i_net}^{\beta_4} + \frac{\beta_5 \hat{Q}_j^{\beta_6}}{\beta_7 \hat{C}_{obs}^{\beta_8} + \beta_9 \hat{C}_{transfer}^{\beta_{10}}} + \frac{\beta_{11}}{1 + e^{-\beta_{12}(\Delta RA - 360 \frac{\beta_{13}}{\rho_{max}})}}. \quad (18)$$

An evolved linear combination of maximum value-normalized (hatted) observation, transfer, and network costs [Eq. (17)] is then used to generate subsequent selection priorities, [Eq. (18)]. The top priority targets are selected until a mission cost cap is exceeded, in which point a refueling operation is initiated and network costs are recalculated. An example set of target priorities is shown right in Fig. 6.

We precalculate the times of year, in which each star in the target list crosses the outermost boundary of the observable window for its required observation time, the threshold in which an observation can cross zenith and incur no off-angle costs, and the starshade cut-off region. Observation times are calculated for the most promising target sequences by directly selecting the shortest relative time until the target star is next observable, if it is not within one orbit period. Additional evolved logic is used to decide whether it is more optimal to wait until relative right ascension-dependent acceleration (station-keeping) costs are lower at the expense of increasing time limitations and reducing the ability to keep up with the observable window. This decision is made when outside of the observable window by waiting until the current right ascension intersects an interpolated periodic function of target declination and time of year that is normalized in the time domain between the observable window's outer boundary and the point of no off-angle cost. If we exclude mission profiles with significant excess time, it is generally most optimal to observe whenever possible; a boundary scale that is everywhere at unity.

6.5 Appendix E: Delta-V Trades

We calculate the cumulative Starshade electrical and chemical delta-V at various relative system mass distributions with the rocket equation:

$$\Delta V = g_0 I_{sp} \ln \frac{m_0}{m_f}. \quad (19)$$

Using the relative specific impulse and system masses presented in Sec. 3.1, we bound the worst and the best-case servicer electrical delta-V by calculating for maneuvers made with a docked starshade and a full and empty chemical propellant tank, respectively. These extremes bound variance that could occur during mission operations due to variance in relative fractional electrical/chemical fuel consumption rates from unity.

We calculate the chemical delta-V delivered by the starshade by assuming the starshade is allocated an additional 100 kg of electrical propellant for precision maneuvering and docking, and setting a 1000-kg capacity chemical propellant tank that is refilled by the servicer between observations. This tank could be made smaller to improve total delivered chemical delta-V provided that the longest, most chemically expensive observation in the mission plan does not exceed its capacity.

Acknowledgments

We appreciate numerous discussions with the Goddard Space Flight Center Orbit Analysis Group, including David Folta, Rizwan Qureshi, Robert Pritchett, Donald Dichmann, Sun Hur-Diaz, and Steven Hughes. We thank the Team-X Laboratory at Jet Propulsion Laboratory for engineering analysis and cost estimation, and the Exoplanet Exploration Program (ExEP) Team for support and technical discussion, especially Gary Blackwood, Keith Warfield, Karl Stapelfeldt, and Erick Mamajek. As special thanks go to Stuart Shaklan and Sergi Hildebrandt for their guidance on operating the SISTER software package. Goddard Space Flight Center's Internal Research and Development funds (IRAD) supported work at Goddard, and NASA's ExEP supported the Team-X study at JPL and scientific and engineering workshops at Goddard.

References

1. National Academies of Science, Engineering, and Medicine, *Exoplanet Science Strategy*, The National Academies Press, Washington DC (2018).
2. National Research Council, *New Worlds, New Horizons in Astronomy and Astrophysics*, The National Academies Press, Washington DC (2010).
3. J. Mather et al., “Orbiting starshade: observing exoplanets at visible wavelengths with GMT, TMT, and ELT,” *Bull. Am. Astron. Soc.* **51**(7), 48 (2020).
4. E. Peretz et al., “Mapping the observable sky for a remote occulter working with ground-based telescopes,” *Proc. SPIE* **11117**, 11120S (2019).
5. E. Peretz et al., “Mapping the observable sky for a remote occulter working with ground-based telescopes,” *J. Astron. Telesc. Instrum. Syst.* **7**(2), 021212 (2021).
6. E. Peretz et al., “Exoplanet imaging performance envelopes for starshade based missions,” *J. Astron. Telesc. Instrum. Syst.* **7** (2021).
7. C. C. Stark et al., “Maximized exoearth candidate yields for starshades,” *J. Astron. Telesc. Instrum. Syst.* **2**(4), 041204 (2016).
8. D. R. Keithly et al., “Optimal scheduling of exoplanet direct imaging single-visit observations of a blind search survey,” *J. Astron. Telesc. Instrum. Syst.* **6**(2), 027001 (2020).
9. G. Soto et al., “Optimal starshade observation scheduling,” *Proc. SPIE* **10698**, 106984M (2018).
10. R. Trabert et al., “Design reference missions for the exoplanet starshade (Exo-S) probe-class study,” *Proc. SPIE* **9605**, 96050Y (2015).
11. S. R. Hildebrandt et al., “SISTER: Starshade Imaging Simulation Toolkit for Exoplanet Reconnaissance,” To be submitted to JATIS (2021).
12. ESA, “VizieR online data catalog: the Hipparcos and Tycho catalogues (ESA 1997),” VizieR Online Data Catalog, I/239 (1997).
13. E. Høg et al., “The Tycho-2 catalogue of the 2.5 million brightest stars,” *Astron. Astrophys.* **355**, L27–L30 (2000).
14. L. Lindgren et al., “Gaia data release 2—the astrometric solution,” *Astron. Astrophys.* **616**, A2 (2018).
15. B. S. Gaudi et al., “The Habitable Exoplanet Observatory (HabEx) mission concept study final report” (2020).
16. J. Leconte et al., “Increased insolation threshold for runaway greenhouse processes on Earth-like planets,” *Nature* **504**, 268–271 (2013).
17. R. M. Ramirez and L. Kaltenegger, “A volcanic hydrogen habitable zone,” *Astrophys. J. Lett.* **837**(1), L4 (2017).
18. R. K. Kopparapu et al., “Habitable zones around main-sequence stars: new estimates,” *Astrophys. J.* **765**, 131 (2013).
19. R. A. Brown, “Single-visit photometric and obscurational completeness,” *Astrophys. J.* **624**, 1010–1024 (2005).
20. J. Green et al., “Wide-Field Infrared Survey Telescope (WFIRST) final report,” arXiv:1208.4012 (2012).
21. The LUVOIR Team, “The LUVOIR mission concept study final report,” arXiv:1912.06219 (2019).
22. S. Seager et al., “Exo-S: starshade probe-class exoplanet direct imaging mission concept final report” (2015).
23. A. Koenig and S. D’Amico, “Orbit design and control for the earth-orbiting starshade mission,” Space Rendezvous Laboratory, Stanford University (2019).
24. S. Hughes et al., “Benchmarking the collocation stand-alone library and toolkit (CSALT),” **7**, NASA (2017).
25. K. A. D. Jong, *Evolutionary Computation: A Unified Approach*, MIT Press, Cambridge, MA (2006).
26. I. Khan, M. Maiti, and M. Maiti, “Coordinating particle swarm optimization, ant colony optimization and k-opt algorithm for traveling salesman problem,” in *Int. Conf. Math. Comput.*, pp. 103–119 (2017).
27. A. Sancetta, “Greedy algorithms for prediction,” *Bernoulli* **22**, 1227–1277 (2016).

28. E. Angelelli, R. Mansini, and M. Vindigni, "Exploring greedy criteria for the dynamic traveling purchaser problem," *Cent. Eur. J. Oper. Res.* **17**, 141–158 (2009).
29. M. Schöck et al., "Thirty Meter Telescope site testing I: overview," *Publ. Astron. Soc. Pac.* **121**, 384–395 (2009).

Eliad Peretz is a technology program manager and leads the small sat development lab at Goddard Space Flight Center. He is the ORCAS mission study PI, and the Remote Occulter Co-I, and has been positioned at GSFC since 2015. He received his BSc in aerospace engineering from the Israeli Institute of Technology and his PhD in mechanical aerospace engineering with a minor in applied physics from Cornell University.

John C. Mather is the senior project scientist on the James Webb Space Telescope. He is also a senior astrophysicist in the Observational Cosmology Laboratory located at NASA's Goddard Space Flight Center.

Kevin Hall is currently an astrophysics research assistant within the Department of Astronomy at the University of Maryland. He received his BS degree in astronomy and physics in 2019.

Lucas Pabarcus is a NASA Pathways student trainee in attitude control systems at Goddard Space Flight Center and a junior undergraduate student studying physics and astronomy at Wesleyan University.

Sun Hur-Diaz received her BS degree in aerospace engineering from The University of Texas at Austin and her MS and PhD degrees in aeronautics and astronautics from Stanford University. She joined NASA Goddard Space Flight Center as a member of the Navigation and Mission Design Branch. Her current projects at NASA include mission and navigation design of Lunar IceCube, leading the flight dynamics effort on DAVINCI+, and leading the research and development of flight software apps for autonomous onboard navigation, guidance, and control.

Adam Koenig received his BS degree in aerospace engineering from Wichita State University in 2012 and his MS and PhD degrees in aeronautics and astronautics from Stanford University in 2015, and 2019, respectively. Currently, he is a postdoctoral scholar in the Space Rendezvous Laboratory at Stanford University. His research interests include astrodynamics and advanced guidance, navigation, and control algorithms for distributed space systems.

Simone D'Amico received his BS and MS degrees from Politecnico di Milano in 2003 and the PhD from Delft University of Technology in 2010. From 2003 to 2014, he was research scientist and team leader at the German Aerospace Center (DLR). There, he gave key contributions to the design, development, and operations of spacecraft formation flying and rendezvous missions such as GRACE (United States/Germany), TanDEM-X (Germany), PRISMA (Sweden/Germany/France), and PROBA-3 (ESA). Since 2014, he has been assistant professor of Aeronautics and Astronautics at Stanford University, founding director of the Space Rendezvous Laboratory (SLAB), and Satellite Advisor of the Student Space Initiative (SSSI), Stanford's largest undergraduate organization.