Maximized exoEarth candidate yields for starshades

Christopher C. Stark
Stuart Shaklan
Doug Lisman
Eric Cady
Dmitry Savransky
Aki Roberge
Avi M. Mandell

Maximized exoEarth candidate yields for starshades

Christopher C. Stark,[a,*] **Stuart Shaklan,**[b] **Doug Lisman,**[b] **Eric Cady,**[b] **Dmitry Savransky,**[c]
Aki Roberge,[d] **and Avi M. Mandell**[d]
[a]Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, Maryland 21218, United States
[b]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109, United States
[c]Cornell University, Sibley School of Mechanical and Aerospace Engineering, Ithaca, New York 14853, United States
[d]NASA Goddard Space Flight Center, Greenbelt, Maryland 20771, United States

Abstract. The design and scale of a future mission to directly image and characterize potentially Earth-like planets will be impacted, to some degree, by the expected yield of such planets. Recent efforts to increase the estimated yields, by creating observation plans optimized for the detection and characterization of Earth-twins, have focused solely on coronagraphic instruments; starshade-based missions could benefit from a similar analysis. Here we explore how to prioritize observations for a starshade given the limiting resources of both fuel and time, present analytic expressions to estimate fuel use, and provide efficient numerical techniques for maximizing the yield of starshades. We implemented these techniques to create an approximate design reference mission code for starshades and used this code to investigate how exoEarth candidate yield responds to changes in mission, instrument, and astrophysical parameters for missions with a single starshade. We find that a starshade mission operates most efficiently somewhere between the fuel- and exposuretime-limited regimes and, as a result, is less sensitive to photometric noise sources as well as parameters controlling the photon collection rate in comparison to a coronagraph. We produced optimistic yield curves for starshades, assuming our optimized observation plans are schedulable and future starshades are not thrust-limited. Given these yield curves, detecting and characterizing several dozen exoEarth candidates requires either multiple starshades or an $\eta_{\text{Earth}} \gtrsim 0.3$. © *2016 Society of Photo-Optical Instrumentation Engineers (SPIE)* [DOI: 10.1117/1.JATIS.2.4.041204]

Keywords: telescopes; numerical methods; planetary systems.

Paper 16009SS received Feb. 11, 2016; accepted for publication May 13, 2016; published online Jun. 14, 2016.

1 Introduction

The design of future large-aperture space missions will be influenced by the desire to detect and characterize potentially Earth-like planets around other stars (exoEarth candidates). Recent estimates of the yield of exoEarth candidates have focused primarily on missions that use an internal coronagraph to suppress starlight.[1–4] However, coronagraphs can have relatively low throughput, potentially limited bandwidth requiring many observations to obtain a full spectrum, and require extremely stable mirrors. External occulters (starshades), on the other hand, having relatively high throughput, operate over a relatively wide bandpass, and, because the starlight is suppressed prior to diffraction by the primary mirror, do not require similar levels of telescope mirror stability.[5,6]

Yield estimates for coronagraph-based mission concepts have evolved recently from noise floor limited searches to Earth-twin optimized searches. These optimized searches select individual exposure times to maximize overall mission yield, resulting in roughly a tripling of the expected yields for coronagraphs.[3,4] Scaling relationships from these yield-maximized estimates have informed coronagraph-based mission design, e.g., yield is more sensitive to telescope aperture than total mission lifetime.[3] While studies of starshade-based mission yields have explored the impact of astrophysical parameters and the starshade's flight path for specific mission concepts,[7–9] yield models have not yet investigated in detail the balance between exposure time and fuel use necessary to maximize the yield of a starshade-based mission.

Here we apply the optimization techniques developed by Stark et al.[3,4] to starshade-based missions. We first solve for the optimal observation plan for the mission and then estimate the mission-long fuel requirements of such a plan using scaling relationships. In Sec. 2, we present our baseline mission assumptions. We then discuss updates to the optimization code for the case of starshades in Sec. 3. We present and discuss our results in Secs. 4 and 5.

2 Assumptions and Caveats

2.1 Scope and Intent of Models

ExoEarth candidate yield calculations are performed with what is commonly referred to as a design reference mission (DRM) code. Yield estimates published in the literature have been calculated using different DRM codes that vary significantly in scope, complexity, and intent. Some DRM codes are simply back-of-the-envelope calculations based on the exposure times of Earth-twins at quadrature around nearby stars. Other DRM codes explicitly calculate the visibility of every star given a well-defined field of regard for the mission, create a realistic observation schedule, and perform a dynamic traveling salesman-type calculation. For starshade-based missions, most yield estimates have adopted the latter, which we refer to as "mission execution simulators" because they simulate a mission from beginning to end, making decisions on the fly and calculating starshade fuel used when slewing between specific targets.[9]

Our DRM code, like the code of Refs. 3 and 4, is neither a back-of-the-envelope calculation nor a mission execution simulator. Our code creates an optimized static observation plan, which, when executed, will produce higher yields than any other static observation plan. By static, we mean that the code does not update the exposure times of the plan based on information learned over the course of the mission. However, our code inherently assumes that spectral characterization time is only applied to true exoEarth candidates, which can only be true if the decision to characterize is made on the fly. Our code handles the bookkeeping for this spectral characterization time by weighting each observation's possible spectral characterization time by the probability that a planet is detected during that observation. Thus, it may be more accurate to call our optimized observation plans semistatic.

Our code does not determine the absolute times or order of observations and, therefore, does not determine an observation schedule. Rather, we assume our observation plan is schedulable and that an efficient path is possible for the starshade. Previous studies have shown that it is possible to schedule an efficient starshade path that meets fuel constraints for a predetermined set of targets and observations.[8,9]

The fundamentals of our exoEarth candidate yield calculations are based on the completeness methods pioneered by Ref. 1. We start by distributing a large number of synthetic exoEarths around known main sequence stars, drawn from the Hipparcos catalog and vetted for binarity.[3] We distribute a cloud of synthetic planets around each star, sampling all possible habitable zone orbits, inclinations, and phases. We illuminate each synthetic planet with starlight and calculate the planet's exposure time. For each star, we can then determine the fraction of the synthetic cloud that is observable (i.e., completeness) as a function of exposure time. The completeness is equivalent to the probability of detecting the planet, if that planet exists.

Based on the completeness of each star as a function of time, we then prioritize stars by a benefit-to-cost ratio and down-select to those observations that fit within the mission lifetime. Our code can calculate multiple visits to each star, enabling the detection of planets that were unobservable during the previous visit. Our code also optimally distributes exposure time, such that it maximizes yield.[4]

Here we extend this code to starshades. To facilitate a direct comparison with the coronagraph results of Ref. 4, we make identical astrophysical assumptions where possible. For a detailed description of the code and justification of assumptions, we refer the reader to Refs. 3 and 4.

2.2 Baseline Astrophysical Assumptions

With exception to the local zodiacal light, we made the same astrophysical assumptions as Stark et al.[4] We adopted the same habitable zone (HZ) given by Ref. 10 for a Sun-like star and scaling with stellar luminosity. We assumed all synthetic exoEarths are on circular orbits, spaced logarithmically in semimajor axis, and have a V band geometric albedo of 0.2. We also adopted the same exozodi definition and baseline exozodi level of 3 zodis, the details of which are described in Stark et al.[3,4] We assumed the same baseline value of $\eta_\oplus = 0.1$ (except for Sec. 4.1, in which we temporarily adopted $\eta_\oplus = 0.5$ to compare our yield with previous results). We also used the same input target catalog, vetted in an identical

fashion. Table 1 summarizes all of our baseline astrophysical assumptions.

Stark et al.[3,4] assumed targets could be observed at solar longitudes ~135 deg, where the zodiacal light is minimized and varies slowly with solar longitude. The zodiacal brightness was calculated for each star using this solar longitude and the star's ecliptic latitude (see Appendix B in Ref. 3). As a result, the mean local zodiacal light brightness $z = 23$ mag arc sec^{-2}.

However, to avoid reflecting sunlight into the telescope, starshades observe at solar elongations 30 deg $\lesssim \alpha \lesssim$ 90 deg, where the local zodi can be as much as 14 times brighter. To calculate the local zodiacal light surface brightness for each star, we transformed the stellar equatorial coordinates to ecliptic coordinates. Because our code does not schedule each observation or calculate detailed starshade/telescope pointing, we assumed each target is observed at the median solar elongation $\alpha = 60$ deg. Combined with the ecliptic coordinates of each target star, this assumed elongation determines the solar longitude and ecliptic latitude of observation. We then interpolated Table 17 in Ref. 11 to calculate the local zodiacal light surface brightness for each of these pointings. For our baseline starshade mission, we found a mean local zodi brightness of $z = 22.1$ mag arc sec^{-2}, as listed in Table 1. In practice, we used the local zodi brightness calculated individually for each target.

Table 1 Baseline astrophysical parameters.

Parameter	Value	Description
$\eta_\oplus$	0.1	Fraction of Sun-like stars with an exoEarth candidate
R_p	$1 R_\oplus$	Planet radius
a	[0.75, 1.77] AU[a]	Semimajor axis (uniform in log a)
e	0	Eccentricity (circular orbits)
$\cos i$	[−1,1]	Cosine of inclination (uniform distribution)
ω	[0,2π)	Argument of pericenter (uniform distribution)
M	[0,2π)	Mean anomaly (uniform distribution)
Φ	Lambertian	Phase function
A_G	0.2	Geometric albedo of planet at 0.55 and 1 μm
z	22.1 mag arc sec^{-2}[b]	Average V band surface brightness of zodiacal light
x	22 mag arc sec^{-2}[c]	V band surface brightness of 1 zodi of exozodiacal dust
n	3	Number of zodis for all stars

[a]a given for a solar twin. The habitable zone is scaled by $\sqrt{L_\star/L_\odot}$ after calculating projected separation s_p.
[b]Zodiacal light is brighter for a starshade than a coronagraph—see Sec. 2.2—and varies with ecliptic latitude—see Ref. 11.
[c]For solar twin; varies with spectral type—see Appendix C in Ref. 3.

2.3 Baseline Mission Assumptions

Table 2 lists the baseline mission parameters. For comparison with Stark et al.,[4] we duplicated as many of the mission parameters as possible. ExoEarth candidates were assumed to be detected at V band and characterized at a wavelength of 1 μm to search for water. Spectral characterization time counted against the total exposure time budget. Only exoEarth candidates simultaneously observable at 0.55 and 1.0 μm counted toward the yield. We assumed exoEarth candidates and other planets or background objects can be distinguished by color, less demanding spectral information, or a simple orbital analysis, such that negligible time is used taking spectra of non-exoEarth candidates. We required the same characterization signal-to-noise ratio (S/N) and spectral resolving power requirements as Stark et al.[4]

Like Stark et al.,[4] we allowed stars to be observed multiple times and included revisit completeness. However, because many targets will have limited visibility windows for a starshade-based mission and the time between visits to a star will likely be dictated by schedulability, which we do not model, we only allowed up to five visits per star and did not allow the time between visits to be optimized. For each target star, we set the time between visits equal to the median orbital period of synthetic HZ planets around that star, allowing the synthetic planets to mostly lose phase coherence.[2] These assumptions are ultimately unimportant and impacted the yield by ≲10%, as optimization typically selects just one or two visits to each star.

Like Stark et al.,[4] we ignored the details of overheads. Unlike a coronagraph, starshades will not require an extended wavefront control time. Therefore, whereas Stark et al.[4] assumed one year of exposure time and one year of wavefront control/overheads, we could have justifiably adopted two years of

Table 2 Baseline mission parameters.

Parameter	Value	Description
D	4 m	Telescope diameter
I_{sk}	300 s	Specific impulse of stationkeeping propellant (chemical)
I_{slew}	3000 s	Specific impulse of slew propellant (electric)
ε_{sk}	0.8	Efficiency of stationkeeping fuel use
ε_{slew}	0.91[a,b]	Efficiency of slew fuel use
D_{ss}	61.7 m[b,c]	Diameter of starshade
m_{dry}	2742 kg[b,d]	Dry mass of starshade spacecraft including contingency
m_{tot}	9800 kg	Total initial mass of starshade spacecraft including fuel
$\mathcal{T}$	10 N[e]	Thrust
T	0.65	End-to-end facility throughput, excluding photometric aperture factor
X	0.7	Photometric aperture radius in λ/D

Table 2 (Continued).

Parameter	Value	Description
Υ	0.69	Fraction of Airy pattern contained within photometric aperture
Ω	$b\pi(X\lambda/D)^2$ radians[b]	Solid angle subtended by photometric aperture
b	1.0	Areal broadening of the planet's point spread function (PSF)
IWA	60 mas	Inner working angle
OWA	∞	Outer working angle
Δmag_{floor}	2.5–2.5 log ζ_d[b]	Systematic noise floor (i.e., dimmest point source detectable at S/N)
λ_d	0.55 μm	Central wavelength for detection (V band)
$\Delta\lambda_d$	0.22 μm	Bandwidth for detection
S/N_d	7	Signal-to-noise ratio required for broadband detection of a planet
ζ_d	10^{-10}	Raw contrast level in detection region, relative to theoretical Airy pattern peak
$CR_{b,detector,d}$	0	Detector noise count rate for detection
λ_c	1.0 μm	Central wavelength for spectral characterization
R_c	50	Spectral resolving power required for spectral characterization
S/N_c	5	Signal-to-noise ratio per spectral resolution element required for spectral characterization
ζ_c	10^{-10}	Raw contrast level for spectral characterization, relative to theoretical Airy pattern peak
$CR_{b,detector,c}$	0	Detector noise count rate for spectral characterization

[a]Optimized value chosen by code given high thrust assumption.
[b]Dependent quantity.
[c]Optimized for a given bandpass, IWA, contrast, and telescope diameter; see Sec. 2.5.
[d]See Sec. 2.5.
[e]Value chosen to ensure starshade operates far from thrust-limited regime.

total exposure time for a starshade mission. As we will show later, this time is better spent slewing than observing. Thus, we adopted a total mission lifetime of five years and allowed our code to optimize the balance between slew and exposure time. We assumed telescope repointing and fine alignment of the starshade make up a negligibly small fraction of the total mission lifetime.

Compared to Stark et al.,[4] we changed a number of instrument parameters to better reflect expected starshade performance. First,

we required the starshade to operate over the full bandpass of 0.5 to 1.0 μm. Therefore, we adopted a contrast of 10^{-10} for both detection and characterization. Because the identification of exoEarth candidates may require color information, we adopted half the starshade's bandpass as the bandwidth, or 0.22 μm, which is twice as large as was assumed for the coronagraph in Ref. 4. Finally, we increased the total system throughput from 0.2 to 0.65, consistent with the optical throughput and detector quantum efficiency assumed by the Exo-S instrument design[12] (the usable fraction of the planet's light, Υ, contained within the assumed photometric aperture is an additional factor that further decreases the effective throughput).

Unlike the coronagraph's inner working angle (IWA), expressed in units of λ/D by Stark et al.,[4] we adopted a baseline 60 mas IWA for our starshade that is independent of aperture size and adopted an infinite outer working angle. We reason that the IWA of a starshade is determined by the distance between the telescope and starshade as well as the diameter of the starshade and will likely be limited by fuel usage and manufacturability. We therefore should not expect the same scaling relationship between yield and aperture size as found by Stark et al.[4] for a coronagraph.

Together, the contrast, bandwidth, IWA, and telescope aperture size will determine the starshade's diameter D_{ss}, distance z, and mass m_{ss}, which we discuss in Sec. 2.5. All starshade diameters discussed in this paper are tip-to-tip diameters and IWA $= D_{ss}/2z$. We note that the baseline mission parameters discussed here correspond to a 61.7 m tip-to-tip diameter starshade at 106 Mm with a dry mass of 2742 kg.

We also implemented fuel constraints in our simulations. We adopted a high-efficiency electric propulsion system for slewing with a specific impulse of 3000 s and a chemical propulsion system for stationkeeping with a specific impulse of 300 s. For realistic propulsion systems, thrust will likely be limited to <1 N. However, in an effort to investigate upper limits on the yields for starshade missions, we chose to work far from the thrust-limited regime and adopted a thrust of 10 N, an assumption we address in Sec. 4.2. Finally, we chose a total mass limit of 9800 kg, the Delta IV Heavy payload mass limit to the Sun-Earth L2 point.[13] We investigate other mass limits in Sec. 4.2.

We reduced the baseline telescope diameter from the 10 m aperture assumed by Stark et al.[4] to 4 m. We do not intend the baseline starshade mission's yield calculated here to be similar to the baseline coronagraph's yield calculated by Stark et al.[4] The assumed baseline values simply represent a reasonable starting point comparable to previous works. Later, we will determine what mission parameters are required to achieve a given yield.

We note that not all of the baseline parameters are independent quantities. For example, if we were to deviate from the baseline telescope diameter of 4 m while keeping the contrast and IWA fixed, the starshade diameter would have to change. We indicate dependent quantities with footnote "b" in Table 2.

2.4 Uncertainty in Yield

We note that all yield estimates are probabilistic quantities. We do not know ahead of time which stars will actually host the planets, so it is possible that certain instances of nature may be more or less productive than other instances. Thus, all yield estimates should have an intrinsic uncertainty based on

the possible arrangement of planetary systems. This uncertainty has been included by some previous studies.[7]

However, this source of uncertainty is significantly less than other sources, including astrophysical quantities (e.g., η_{Earth} and median exozodi level), unmodeled instrument effects (e.g., contrast as a function of stellar diameter), and the as-yet unstudied impact of observational techniques and planet discrimination methods (e.g., characterization by orbit constraint versus characterization by spectroscopy). Because of this, we choose not to plot the uncertainties associated with planetary arrangement alone. However, we emphasize that the yield estimates presented in this work will likely change as these parameters and effects are better understood and that each yield is accompanied by an intrinsic uncertainty.

2.5 Starshade Design

We adopted numerically optimized starshades designed to produce the adopted IWA and contrast over the full wavelength band from 0.5 to 1.0 μm while minimizing both the starshade diameter and telescope–starshade separation.[6,14] To numerically design a starshade, we approximated the electric field near the center of the starshade shadow as the field following an azimuthally symmetric apodization and used a linear optimization to produce a radial profile that meets any desired input constraints. With a sufficient number of petals, this approximation to a starshade with a shaped edge becomes arbitrarily good.

While this optimization is usually done with physical parameters—lengths, wavelengths, angular distances—these parameters have some degeneracies. For example, identical starshade shapes may be used at different distances from the telescope if a different IWA and bandpass are assumed. We rewrote the optimization in terms of four nondimensional, nondegenerate parameters, enabling efficient coverage of the parameter space of possible optimized starshade designs.[15] We then interpolated over this grid to estimate the diameter and distance of a starshade defined by a given IWA, contrast, and telescope diameter.

We note that starshade contrast can vary with stellar separation while our DRM code adopts a singular uniform contrast exterior to the IWA. To simplify the problem, we adopted a singular contrast metric for each starshade design, evaluated by averaging the contrast over a 1 λ/D wide annulus centered at 1 λ/D exterior to the IWA.

For each starshade, we calculated the dry mass as

$$m_{dry} = (1 + \mathcal{C}) \times (m_{ss} + m_{bus}),\qquad(1)$$

where $\mathcal{C} = 0.3$ is the contingency factor, m_{ss} is the starshade payload, and m_{bus} is the spacecraft bus. For the starshade payload mass, we used a linear fit to recent starshade designs,[12] such that

$$m_{ss} = 35 \times \left(\frac{D_{ss}}{1\ m}\right) - 620\ kg.\qquad(2)$$

For the bus mass, we adopted the following formula, based on the analysis of Ref. 12, which scales bus mass with starshade and fuel mass:

$$m_{bus} = 615 + 0.09\left(\frac{m_{ss}}{1\ kg} - 560\right)$$
$$+ 0.03\left(\frac{m_{tot}}{1\ kg} - \frac{m_{dry}}{1\ kg}\right) kg. \qquad (3)$$

Substitution of Eqs. (2) and (3) into Eq. (1) gives the final dry mass equation we used.

$$m_{dry} \approx 52 \times \left(\frac{D_{ss}}{1\ m}\right) + 0.04 \times \left(\frac{m_{tot}}{1\ kg}\right) - 154\ kg. \qquad (4)$$

3 Methods

To maximize the yield of a coronagraph-based mission, Stark et al.[4] optimized the use of the single limiting resource, exposure time. To maximize the yield of a starshade, we must recognize that starshades have two limiting resources: exposure time and fuel. For a fixed-length mission, these two resources are not independent; fuel use depends on slew time, which is the fraction of the mission not used for exoplanet science exposures. For example, one could reduce fuel consumption per target by sacrificing exposure time and awarding it to slew time, allowing one to visit more targets for the same total fuel mass.

To better understand this trade-off, we find it useful to first proceed in Sec. 3.1 with a simplistic expression for fuel use and optimize yield under the incorrect assumption that time and fuel are independent resources. We refer to this scenario as the "time-budgeted" optimization scenario. We will then present an approximate analytic expression for fuel use in Sec. 3.2, relating fuel mass to exposure time, and show that the time-budgeted optimization does not ideally balance fuel use with exposure time. Finally, in Sec. 3.3, we will detail how to arrive at an ideally optimized starshade yield for a given fuel mass by using our analytic fuel use expression to balance among slew time, exposure time, number of slews, and number of targets. We refer to this latter method as the "fuel-time balanced" optimization method.

3.1 Time-Budgeted Optimization

Stark et al.[3] discuss an algorithm to maximize the completeness, C, and thus yield, by optimizing the observation plan. This algorithm consists of the following steps:

1. Award exposure time $d\tau$ to the i'th observation with the largest completeness curve slope $dC_i/d\tau$, where τ is the exposure time. As a result of this process, all observations will achieve the same final completeness curve slope.

2. Prioritize observations by the benefit-to-cost ratio C_i/τ_i and then down-select to those observations that fit within the total exposure time budget.

The prioritization in the second step assumes that the cost of an observation is simply the exposure time, which is valid for a coronagraph-based mission where time is the limiting resource. Starshades, however, have an additional limiting resource: fuel. The cost metric for a starshade observation is therefore some combination of fuel and exposure time. In practice, calculating the fuel associated with an observation using a starshade is nontrivial and depends on many factors, including the propulsion assumed, stationkeeping during observation, number of

slews, distance between each slew, and the time allowed for each slew. Because our code does not calculate an observation schedule, we do not know the precise distance of any single slew of the starshade, and it is therefore technically impossible to prioritize the observations by the actual fuel usage.

In light of this, we can simplify the problem and approximate the fractional fuel used by a single observation as $1/n_{slews}$, where n_{slews} is the number of starshade slews performed. Under the assumption that each slew consumes fractional fuel equivalent to $1/n_{slews}$, one choice for the cost metric of the i'th observation could be $1/n_{slews} + \tau_i/\Sigma\tau$, where $\Sigma\tau$ is the total exposure time budget. However, this cost fails to capture the slew-limited regime: if n_{slews} is sufficiently small, the mission may not use up all of the exposure time, in which case exposure time should not be considered a cost. For example, imagine a scenario in which only 10 slews of the starshade are allowed for the baseline mission. In this case, the baseline mission could completely search the observable portion of 10 HZs to the noise floor in far less than the exposure time allowed. In the slew-limited regime (i.e., fuel-limited regime), the exposure time of any individual target is irrelevant and the cost metric should be the fuel consumption only.

To capture the transition between the slew- and exposure time-limited regimes, we can adopt a linear combination of fractional fuel usage and fractional exposure time usage, such that the benefit-to-cost ratio for the i'th observation is given by

$$\frac{Benefit}{Cost} = \frac{C_i}{\kappa/n_{slews} + (1-\kappa)\tau_i/\Sigma\tau}, \qquad (5)$$

where $0 < \kappa < 1$. We can allow the code to choose the value of κ that maximizes the yield by calculating the yield for 1000 different values of κ. In practice, this negligibly increases the run time of the yield estimation code, as the majority of the run time is spent on other optimizations like the first step listed above. We note that we tried other cost metrics, including raising the fuel and exposure time usage to optimized powers, but found the linear combination above to achieve the highest yields in the transition between the fuel- and exposure time-limited regimes.

By adopting $1/n_{slews}$ as the fractional fuel used per slew, we are setting the fuel consumption of each slew equal to the average fuel per slew, a reasonable approximation for many slews. A more detailed simulation would take into account the path of the starshade and optimize the fuel usage. However, keep in mind that our fuel cost metric does not say anything about the absolute amount of fuel used by the mission, a calculation we will address later.

We implemented the above method and calculated the optimized yield as a function of the two limiting resources n_{slews} and $\Sigma\tau$. Figure 1 shows yield contours for the baseline mission parameters in black. In this plot, the y-axis corresponds to the number of slews possible—the number of slews chosen by the code that maximizes the yield could be smaller.

Each yield contour has horizontal and vertical asymptotes. The horizontal asymptotes correspond to the slew-limited regime, in which additional photons do not produce higher yields because $\sim n_{slews}$ HZs have been observed as completely as possible. The vertical asymptotes correspond to the exposure time-limited regime, in which additional slews do not produce higher yields because the mission does not have enough exposure time to use them. The elbow at the lower-left of each curve represents the transition between the fuel- and exposure time-limited regimes, where a starshade-based mission can efficiently
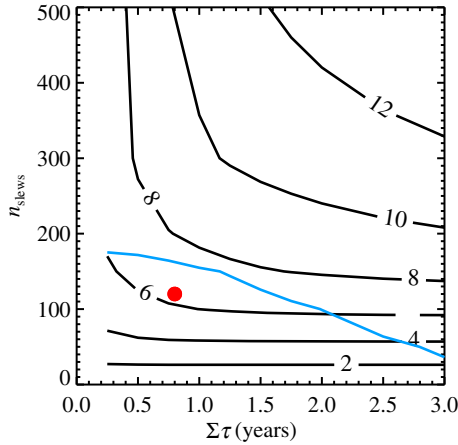
Fig. 1 ExoEarth candidate yield contours for the baseline mission (black) as functions of total exposure time and number of starshade slews using the nonideal time-budgeted optimization method. The red dot marks the location of the baseline mission. A hypothetical contour of constant fuel mass (blue) illustrates that one can trade off exposure time for slews to increase yield; time-budgeted optimization does not necessarily produce the ideal solution.

use all of the allowed slews and all of the exposure time simultaneously.

Of course for a given mission, we are not allowed any combination of n_{slews} and $\Sigma\tau$ that we desire—this choice is governed by the available fuel mass and usage. Figure 1 plots a single hypothetical fuel mass contour in blue, roughly based on the equations presented in Sec. 3.2. The blue curve illustrates that the available fuel mass should determine the precise combination of n_{slews} and $\Sigma\tau$ that maximizes yield, an additional optimization that our time-budgeted method does not perform.

3.2 Analytic Expressions for Fuel Use

As shown in Sec. 3.1, to maximize the yield of a starshade mission, we need to optimally balance exposure time and the number of slews for a given mission. Because the total slew time is the remainder of the mission not used for exposures, this balancing act requires relating the number of slews to the total slew time, i.e., an expression for fuel use. Here we provide a method to calculate fuel use based on simple scaling relationships.

Basic rocket physics tells us

$$\ln\frac{m + dm}{m} = \frac{\Delta v}{v_{ex}}, \qquad (6)$$

where m is the initial mass, dm is the change in mass (fuel expended), Δv is the change in velocity of the starshade, and v_{ex} is the exhaust velocity of the propellant. Due to concerns about exhaust from an ion propulsion system interfering with observations, recent starshade designs have assumed separate propulsion systems for slewing and stationkeeping, adopting an efficient ion propulsion for the former and chemical propulsion for the latter.[12,7] Separating these two factors, we can rewrite the above expression in terms of the specific impulse of each propulsion system as

$$\ln\frac{m + dm}{m} = \frac{\Delta v_{slew}}{gI_{slew}} + \frac{\Delta v_{sk}}{gI_{sk}}, \qquad (7)$$

where $g = 9.8\ \mathrm{m\,s^{-2}}$. Below we present approximate analytic expressions for fuel used by slewing and stationkeeping for each observation. We include only the terms that are commonly included in the literature, neglecting radiation pressure on the starshade and the contribution of the orbital dynamics to the slew time. We then explain how to use these expressions to estimate mission-long fuel use.

3.2.1 Stationkeeping

Figure 2 shows the approximate geometry of the starshade and telescope orbiting about L2 at a distance r_{orbit} (exaggerated for illustration), as seen from the north ecliptic pole at a time when they are located in the ecliptic plane. For a more precise illustration showing the starshade and telescope orbits, see Fig. 2 in Ref. 16. Given that $r_{orbit} \ll 1$ AU, solar elongation $\approx\alpha$.

The starshade and telescope must be kept aligned to within roughly a meter or so in the lateral direction to ensure that the telescope stays within the starshade's shadow.[12] While the tolerance is much looser in the axial direction ($\hat{z}$), the distance offset must be controlled at some point, either as part of stationkeeping or during the slew of the telescope. For simplicity, we decide to include the axial disturbances in the stationkeeping fuel use budget.

In a noninertial rotating reference frame, such as that rotating with L2 about the Sun–Earth barycenter, we can express the acceleration felt by any object as
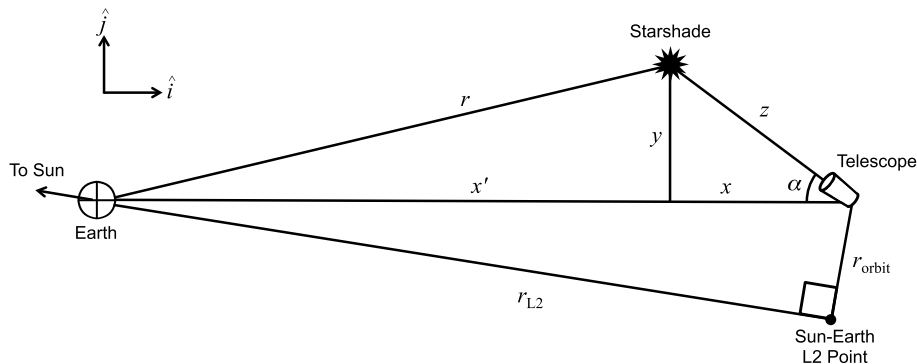

Fig. 2 Diagram of telescope and starshade orientation with respect to Earth observing at an angle α. Because $r_{orbit} \ll 1$ AU, α is approximately the solar elongation.

$$\vec{a}' = \vec{a} - \vec{\Omega} \times (\vec{\Omega} \times \vec{r}) - 2\vec{\Omega} \times \vec{v}', \tag{8}$$

where the first term $\vec{a}$ is the acceleration felt in the inertial frame due to real forces, the second term is the centrifugal pseudo-force term, the third term is the Coriolis pseudo-force term, $\vec{\Omega}$ is the angular rotation rate vector of the rotating frame, $\vec{r}$ is the position of the object in the inertial frame, and $\vec{v}'$ is the velocity of the object in the rotating frame.

For the moment, we will ignore the fact that the telescope and starshade are orbiting about L2 and adopt the orientation illustrated by Fig. 2. For both the starshade and the telescope, which are assumed to be orbiting the Sun at the same rate as Earth, the Sun's gravitational pull should approximately cancel out with each object's centrifugal term. Thus, we are left with terms corresponding to the Earth's gravitational pull and the Coriolis force. We start with the Coriolis force.

Both the telescope and starshade are moving on an orbit about L2, such that they both have velocities in the reference frame rotating with Earth about the Sun and therefore both have associated Coriolis terms. Because we are ultimately interested in the difference in acceleration between the two spacecrafts and to first order their velocities are the same, the majority of the Coriolis effect cancels out. However, the starshade and telescope will have a non-negligible difference in velocity to maintain alignment with a given star. During an observation, the telescope must change its pointing and the starshade must move relative to the telescope to track the motion of the stars. Given the orientation shown in Fig. 2, this means the starshade must have a velocity relative to the telescope in the rotating reference frame, which is roughly directed along $\hat{i}$. Thus, we can roughly express the Coriolis effect on the starshade as

$$\vec{a}_{\text{SS,coriolis}} \approx 2\Omega v_{\text{SS}}\hat{k}, \tag{9}$$

where v_{SS} is the speed of the starshade relative to the telescope in the $\hat{i}$ direction and $\hat{k}$ is directed out of the page. To track the sky when both spacecrafts are in the ecliptic, $v_{\text{SS}} \approx \Omega z$, such that

$$\vec{a}_{\text{SS,coriolis}} \approx 2\Omega^2 z\hat{k}. \tag{10}$$

Given that $\Omega = 2\pi$ year^{-1}, or $\approx 2 \times 10^{-7}$ s^{-1}, the difference in acceleration in the frame rotating with Earth between the starshade and telescope due to the Coriolis effect is $\approx 1.6 \times 10^{-7}(z/1\ \text{Mm})$ m s^{-2} for targets in the ecliptic. When looking at targets at the ecliptic poles, $\vec{\Omega}$ and $\vec{v}_{\text{SS}}$ are antiparallel, such that there is no difference in the Coriolis terms. Assuming targets are uniformly distributed over the sky, such that the median target is at an absolute declination of 30 deg, the median difference in acceleration due to the Coriolis effect should be roughly a factor of $\cos 30$ deg smaller, such that $a_{\text{SS,coriolis}} \approx 1.4 \times 10^{-7}(z/1\ \text{Mm})$ m s^{-2}.

Now we address the first two terms of Eq. (8). Because we are interested in the difference in the forces acting on the starshade and telescope, and because the centrifugal terms cancel with the Sun's gravitational force, the difference in the remaining forces can be well approximated by the difference in Earth's gravitational influence on the two spacecrafts. These terms can therefore be approximated as

$$\vec{a}_T = -\frac{GM_\oplus}{r'^2}\hat{i}, \tag{11}$$

$$\vec{a}_{\text{SS}} = -\frac{GM_\oplus}{r^2}\hat{r}. \tag{12}$$

Taking the difference of these two vectors expressed in terms of $\hat{i}$ and $\hat{j}$, and substituting $\mathcal{Z} = z/r'$, where $r' = \sqrt{r_{\text{L2}}^2 + r_{\text{orbit}}^2}$, gives

$$\Delta\vec{a} \approx \frac{GM_\oplus}{(r_{\text{L2}}^2 + r_{\text{orbit}}^2)}\left[1 - \frac{1 - \mathcal{Z}\cos\alpha}{(1 - 2\mathcal{Z}\cos\alpha)^{3/2}}\hat{i} \right.$$
$$\left. - \frac{\mathcal{Z}\sin\alpha}{(1 - 2\mathcal{Z}\cos\alpha)^{3/2}}\hat{j}\right]. \tag{13}$$

Expanding to first order in $\mathcal{Z}$ and taking the magnitude gives

$$\Delta a \approx \frac{GM_\oplus z}{(r_{\text{L2}}^2 + r_{\text{orbit}}^2)^{3/2}}(1 + 3\cos^2\alpha)^{1/2}. \tag{14}$$

Assuming an average solar elongation of $\alpha = 60$ deg and $r_{\text{orbit}} = 0.5r_{\text{L2}}$, we find $\Delta a \approx 1.1 \times 10^{-7}(z/1\ \text{Mm})$ m s^{-2} when the starshade and telescope are located in the ecliptic plane.

The difference in direction between the Coriolis and gravitational acceleration terms changes with time and depends on the orbit about L2 and the coordinates of the target star. Thus, we cannot simply add these two terms in quadrature; improved estimates of the overall relative accelerations require detailed modeling of the orbital dynamics, which is beyond the scope of this paper. However, the above simple estimates illustrate that the dominant terms controlling the stationkeeping fuel budget are both proportional to starshade–telescope separation, z, and of order $\sim 1 \times 10^{-7}(z/1\ \text{Mm})$ m s^{-2}. This is in agreement with previous studies that have performed such detailed modeling. For example, Savransky et al.[7] adopted $\Delta a = 1.4 \times 10^{-7}(z/1\ \text{Mm})$ m s^{-2} using the orbit-averaged results of Ref. 16. For this study, we adopt the same value of Δa and express stationkeeping fuel use for the i'th observation as

$$\frac{\Delta v_{i,\text{sk}}}{gI_{\text{sk}}} \approx \frac{\Delta a\tau_i}{gI_{\text{sk}}\epsilon_{\text{sk}}}, \tag{15}$$

where ϵ_{sk} is the efficiency of stationkeeping fuel use, typically ~ 0.8 due to the cant of the thrusters.

In the above derivation, we neglected radiation pressure. At the Sun-Earth L2 point, the absolute acceleration due to radiation pressure on a 100-m-diameter starshade at a solar elongation of 60 deg, with a mass of 5000 kg, is $a_{\text{RP}} \approx 5 \times 10^{-6}$ m s^{-2}. This is on par with the acceleration difference due to Earth's gravity at a separation of $z = 100$ Mm. However, the relative radiation pressure force will be offset by the pressure on the telescope, which we assume can be modified to make the difference in acceleration due to radiation pressure negligible.

We also neglected the fuel demands of a spinning starshade. If the starshade must spin for dynamical stability or to mitigate the effects of petal defects, additional fuel will be required to either spin up/down the starshade or change the starshade angular momentum vector before and after slewing.

3.2.2 Slewing

While the starshade and telescope are orbiting near L2, the starshade must move relative to the telescope to point at a new target. Figure 3 illustrates two potential slews for a starshade by
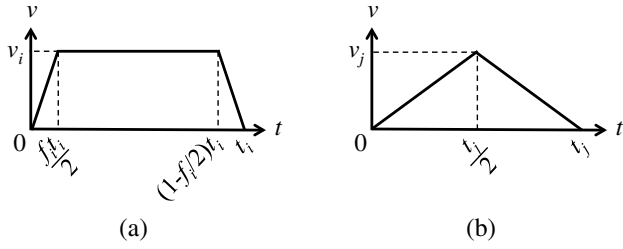
Fig. 3 (a) Example starshade slewing velocity profiles of an efficient slew and (b) an inefficient, thrust-limited slew.

plotting velocity as a function of time. We have assumed that the mass of the starshade changes negligibly during a given slew, such that the initial and final accelerations are equal and the velocity profiles are symmetric. On the left-hand plot, a small fraction f_i of the total slew time of the i'th slew is spent accelerating/decelerating, such that the majority of the slew is spent coasting; this velocity profile represents efficient use of starshade fuel. On the right, $f_j = 1$ and the starshade accelerates/decelerates for the entire duration of the slew; this velocity profile represents inefficient thrust-limited use of starshade fuel.

Assuming the starshade is not operating in the thrust-limited regime, such that $0 < f_i < 1$, we can express the distance traversed during the i'th slew as

$$s_i = \frac{1}{2} a_i \left(\frac{t_i f_i}{2} \right)^2 + v_i t_i (1 - f_i) + \frac{1}{2} a_i \left(\frac{t_i f_i}{2} \right)^2. \quad (16)$$

Given that $v_i = a_i t_i f_i / 2$, and $\langle v \rangle_i = s_i / t_i$, we can rewrite this as

$$\langle v \rangle_i = v_i \left(1 - \frac{f_i}{2} \right), \quad (17)$$

and note that the slew efficiency $\epsilon_{i,\text{slew}} = (1 - f_i/2)$, where $0 < f_i < 1$.

We can also rewrite Eq. (16) as

$$s_i = \frac{1}{2} a_i t_i^2 \left(f_i - \frac{f_i^2}{2} \right). \quad (18)$$

Substituting $a_i = \mathcal{T}/m_i$ for the starshade acceleration, where $\mathcal{T}$ is the thrust and m_i is the starshade mass during the i'th slew,

$$s_i = \frac{1}{2} \frac{\mathcal{T}}{m_i} t_i^2 \left(f_i - \frac{f_i^2}{2} \right). \quad (19)$$

We make a common assumption used by other studies to date: the slew efficiency is constant for all slews, such that $f_i = f$ and $\epsilon_{i,\text{slew}} = \epsilon_{\text{slew}}$. Given Eq. (19), this means that the slew time of the i'th slew is given by

$$t_i = \sqrt{\frac{2 m_i s_i}{\mathcal{T}(f - f^2/2)}}. \quad (20)$$

In reality, each s_i will be different and will depend on optimization of the starshade's path. However, we are interested in a scaling relationship that is approximately valid for the mission as a whole. Thus, we approximate $s_i \approx s$, where s is the typical slew distance. We approximate the typical slew distance by assuming that n_{targets} are distributed uniformly over the sky. Under this assumption, we can divide the sky up into n_{targets} patches, similar to what is shown by the left diagram in Fig. 4. The gray line in this diagram traces a maximally efficient path through these 16 targets, where each slew requires a slew distance of x. The diagram on the right shows that if we quadruple the number of points to an 8×8 grid, a maximally efficient path requires a slew distance of $x/2$. Thus, the slew distance is proportional to the inverse square root of the number of targets.

Assume the starshade travels along the chord between targets distributed uniformly over the surface of a sphere of radius z, where z is the separation between the starshade and telescope.

$$s = 2z \sin \sqrt{\frac{\pi}{n_{\text{targets}}}}. \quad (21)$$

Equation (17) allows us to express the Δv required for the i'th slew ($= 2v_i$) as

$$\frac{\Delta v_{i,\text{slew}}}{g I_{\text{slew}}} = \frac{2 s_i}{g I_{\text{slew}} \epsilon_{\text{slew}} t_i}. \quad (22)$$
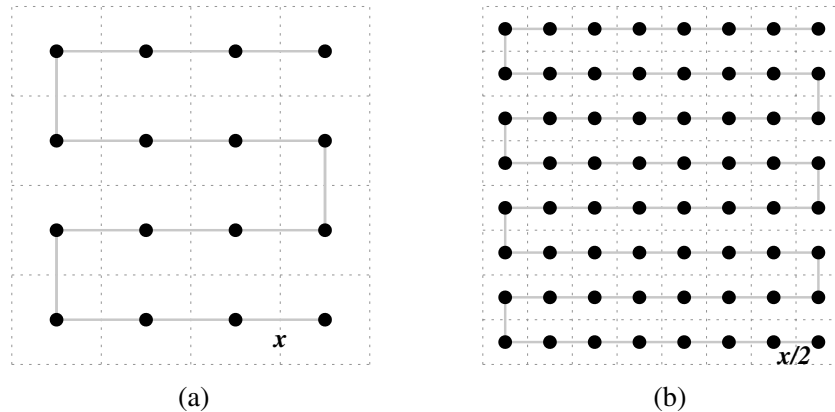

Fig. 4 (a) and (b) Diagrams illustrating how the maximally efficient path (gray line) changes as the number of uniformly distributed targets (dots) increases. The diagram on the right has four times as many targets, but the slew distance is only half that of the diagram on the left.

Equations (20) and (22), with $s_i = s$ given by Eq. (21), constitute our final slew fuel use expressions.

3.2.3 Caveats

In the above derivations, a number of assumptions and approximations were made. First, Eq. (21) assumed that targets were uniformly distributed over the sky and that a maximally efficient path would be chosen. In reality, the starshade may in fact backtrack or otherwise deviate from an optimal path due to pointing constraints and limited knowledge of the ideal path. This may lead one to believe that we will underestimate the amount of fuel required. However, we also note that detailed starshade DRMs, like those of Savransky et al.,[7] choose targets in part based upon their slew distance. As a result, targets may not be uniformly distributed over the sky and the typical slew distance may in fact be smaller than we estimate. Additionally, our slew estimates do not calculate the detailed quasi-halo orbital dynamics of the telescope and starshade.[16]

We also remind the reader that our calculations do not simulate the execution of a mission like those of Savransky et al.[7] We do not model the detailed pointing of the starshade/telescope nor do we schedule each observation given a defined field of regard. Instead, we simply adopt the expected median solar elongation of 60 deg for each target and assume all observations are schedulabe. To prevent yield gains due to unrealistic timing optimizations, we choose not to optimize the time between revisits.

3.2.4 Fuel use calculation method

To calculate the fuel mass for a set of (n_{slews}, $n_{targets}$, $\Sigma\tau$, and Σt_{slew}), we employ the following procedure, starting with the dry mass $m_{dry} = m_0$:

1. Calculate the starshade mass after the i'th observation's stationkeeping as

$$m_i' = m_i \exp\left(\frac{\Delta a_{transverse}\tau_i}{gI_{sk}\varepsilon_{sk}}\right). \tag{23}$$

2. Calculate the slew distance for the i'th slew $s_i = s$ given by Eq. (21).

3. Calculate the slew time for the i'th slew t_i using Eq. (20) with m_i' above for m_i.

4. Calculate the starshade mass after the i'th observation's slew as

$$m_{i+1} = m_i' \exp\left(\frac{2\,s}{g\,I_{slew}\,\varepsilon_{slew}\,t_i}\right). \tag{24}$$

The above procedure is calculated until $i = n_{slews}$, at which point the total mass is known, and the fuel mass can be determined by subtracting off the dry mass.

We compared our fuel mass calculations to 100 individual simulations of a single starshade design by Savransky et al.[7] When adopting the starshade parameters and lateral disturbance acceleration used by Savransky et al.,[7] our estimated fuel mass was 4% larger than that calculated by Savransky et al.[7] on average, with a standard deviation of 4%, and produced an average slew time that fit with the five-year budget adopted

by Savransky et al.[7] Given the exponential nature of the fuel mass expression, we consider this to be a good agreement.

3.3 Fuel-Time Balanced Optimization

Now that we have illustrated the importance of balancing exposure time with slew time to maximize yield, we discuss an alternative, preferred method for yield maximization that automates this balancing process. In Sec. 3.1, we modified the benefit-to-cost ratio to take into account fractional fuel use, treating each slew equally. While it may seem reasonable to consider fuel mass as our limiting fuel resource, this would require knowing the order of observations before planning them, as the fuel mass expended for the i'th observation depends on the mass of the starshade during the i'th observation. Additionally, because of the exponential decay of starshade mass, it is actually the mass ratio of each observation that matters. Therefore, we consider $\Delta v_i/v_{ex}$ [the right-hand side of Eq. (6)] our limiting resource.

Under this assumption, and demanding that we ignore the order of observations, only the stationkeeping term, which is linearly proportional to τ_i, is unique to each observation. As a result, our method of distributing exposure time $d\tau$ to each observation such that all observations have equal slopes $dC_i/d\tau_i$, detailed in step 1 of Sec. 3.1, is still valid and simultaneously optimizes both the fuel and exposure time resource distributions under the assumption that all slews are equal.

However, we still must optimize the balance between total exposure time and total slew time. To do so, we implement the following procedure, which is a modification to the "equal-slope" method of Ref. 17:

1. Given a set of mission parameters, calculate the completeness curves of every observation, $dC_i/d\tau$.

2. Guess the optimal slope of all completeness curves, $(dC/d\tau)_o$.

3. Find the exposure time τ_i for each observation at which $dC_i/d\tau = (dC/d\tau)_o$.

4. Assume $\epsilon_{slew} = 0.5$.

5. Assume a valid combination of $n_{targets}$ and n_{slews}, such that $n_{targets} \leq n_{slews}$, starting with values smaller than the optimal values.

6. Prioritize observations by the benefit-to-cost ratio at the optimized exposure time, $C_i(\tau_i)/X_i$.

7. Select all top-priority observations that fulfill the $n_{targets}$ and n_{slews} constraints. Calculate the total observation plus slew time, as well as the total starshade mass. If both are within the limits of your mission parameters, determine yield $\eta_{planet}\Sigma C_i$ and store this valid solution.

8. Repeat steps 6 and 7 while varying κ over a fine grid of values spanning $0 \leq \kappa \leq 1$.

9. Repeat steps 5 through 8 while varying $n_{targets}$ and n_{slews}.

10. Repeat steps 4 through 9, covering all possible values of ϵ_{slew}.

11. Repeat steps 3 through 10 to optimize the completeness curve slope.

12. Select the combination that produces the highest yield.

The above algorithm essentially describes nested for loops that are designed to scan the five-dimensional parameter space controlling yield. Figure 5 illustrates this multiparameter space for our baseline mission parameters. The plot on the left shows the yield as a function of $(dC/d\tau)_o$. At each point in this plot, the code scanned the four-dimensional $(\kappa, \epsilon_{slew}, n_{targets}, n_{slews})$ subspace and returned the maximum from that subspace. The red dot represents the final yield, maximized over all five parameters.

Because we cannot easily illustrate the four-dimensional subspace, we must show projections of the yield. The middle panel shows a contour projection of the yield in the $(\kappa, \epsilon_{slew})$ subspace, where each yield value represents the maximum from the remaining two parameters, $n_{targets}$ and n_{slews}. The middle panel corresponds to the complete parameter space scanned for just the red dot in the left plot.

Finally, the right contour plot shows the yield of the $(n_{targets}, n_{slews})$ subspace. All yields in this contour plot correspond to the single point at the maximum of the middle panel. The region to the bottom-right, with $n_{slews} < n_{targets}$, is excluded, as you must have at least as many observations as targets. The region to the top-left is excluded by our limit of five visits. The region at the top is excluded by fuel limitations. All yield surfaces are well behaved and an obvious global maximum exists. Our code indeed scans all relevant parameter space and achieves yield maximization.

The above optimization algorithm is written primarily for clarity. If one implemented this algorithm explicitly as written, it would require four nested loops to cover all parameter space with a sorting routine performed on the innermost loop (for the prioritization metric), producing prohibitively long run times. In practice, the sorting/prioritization loop can be moved to just interior to the slope optimization loop as long as κ is finely sampled. We find that our choice of 1000 values for κ is adequate for all tests performed.

Adding status checks for intelligent loop-breaking and improving the initial value guesses can significantly reduce the run time of the above algorithm. Doing so, we are able to run the above algorithm with our baseline mission parameters in 3 min on a single 2.9 GHz CPU. These starshade yield calculations are on par with the coronagraph calculations performed by Stark et al.;[4] the increase in run time due to the above optimization is mostly offset by the limited number of revisits to each star and lack of revisit optimization that was employed by Stark et al.[4]

4 Results

We have implemented the above algorithm in our DRM code to efficiently calculate optimized yields for starshade missions using the baseline parameters and input target list described in Sec. 2. Here we compare our results with previous work and calculate the yield for starshade missions as a function of different mission parameters.

4.1 Comparison with Previous Work

To check our calculations, we compared our results with those of Savransky et al.[7] We chose to compare our yield to that calculated by Savransky et al.[7] for the case of $\eta_\oplus = 0.5$. Thus, for this section only, we will deviate from our baseline assumption of $\eta_\oplus = 0.1$.

Savransky et al.[7] calculated the yield of unique Earth-twin detections using an external occulter combined with a 4 m telescope. Savransky et al.[7] simulated the operation of a starshade mission from beginning to end, determining a flight path consistent with targets' visibility windows and realistically calculating fuel use for slews and stationkeeping. Savransky et al.[7] randomly distributed planets around each star, such that the yield of a given simulation depended on the random arrangement of planets. To estimate the expected (mean) yield as well as the yield uncertainty, Savransky et al.[7] performed 100 such simulations. Our yield simulations explicitly calculate
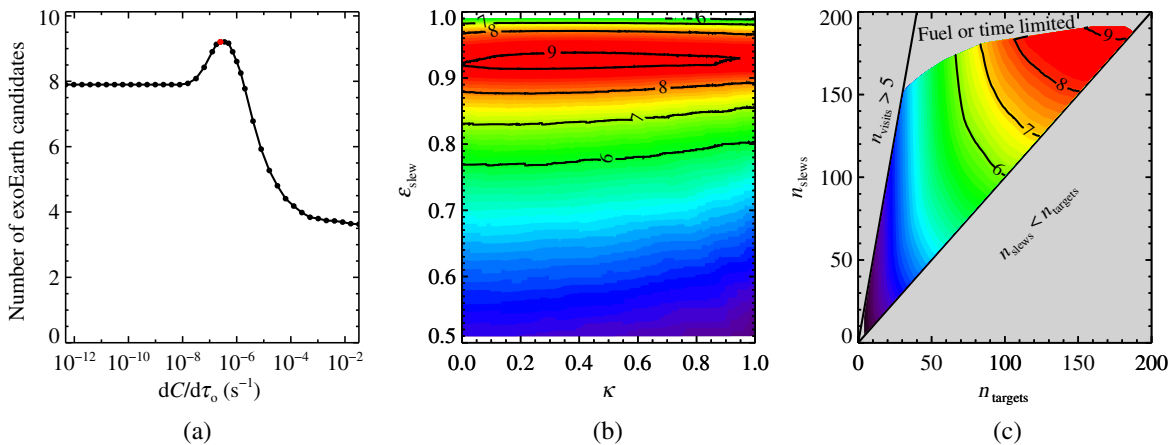

Fig. 5 The five-dimensional space over which our code maximizes yield for the baseline mission. (a) Yield as a function of completeness curve slope. The red dot marks the global yield maximum after scanning the remaining four parameters. (b) Yield contours for two of the remaining parameters. All points in this plot correspond to the maximum achieved after scanning the remaining two parameters and represent the expanded dimensionality of the red point in the left plot. (c) Yield contours of the remaining two parameters, representing the expanded dimensionality of a single point in the middle plot.

the expected mean yield—therefore, we will compare our yield estimates with the average, or expected yield value, reported by Savransky et al.[7]

We modified our code to adopt the same assumptions as Savransky et al.[7] We modified our baseline assumptions of an Earth-twin to that used by Savransky et al.,[7] as well as the HZ definition, distribution of orbits within the HZ, and exozodi definition. We used a 4 m telescope aperture and adopted identical starshade performance parameters: a 51.2-m-diameter starshade operating at 70.4 Mm to provide a 59 mas IWA (evaluated at the 50% throughput point), 10^{-10} contrast, throughput of 0.52, and 20% bandpass. We assumed an identical propulsion system ($I_{slew} = 4160$ s, $I_{sk} = 220$ s, $\epsilon_{sk} = 0.75$, $\mathcal{T} = 0.45$ N), identical dry mass ($m_{dry} = 4210$ kg), and identical initial mass (6000 kg). We adopted identical spectral characterization requirements ($R_c = 70$, $S/N_c = 11$), PSF scale (factor of 1.5 greater in solid angle than Airy pattern) and sampling (27 pixels inside of FWHM), detector noise parameters [read noise (RN) = 3 pix^{-1} read^{-1}, dark current dark current = 0.001 pix^{-1} s^{-1}, $t_{read} = 1000$ s], and approximately the same detection S/N of 4.

From the top-right panel of Fig. 9 in Ref. 7, for $\eta = 0.5$, Savransky et al.[7] estimated $\approx 16 \pm 3$ exoEarth candidates. This data point corresponds to an average of 118 observations, 89 unique stars, and an average total exposure time of ~0.8 years. When adopting identical mission parameters and letting our code optimize the observation plan, we calculated an yield of ~22 exoEarth candidates, ~35% higher than the expected yield of Savransky et al.[7] Our optimized observation plan includes 113 observations and 85 unique stars, fewer total observations than Savransky et al.,[7] and ~0.8 years of total exposure time.

To determine the source of this yield discrepancy, we compared the observations calculated by our code to that of Savransky et al.[7] We determined that the observation plans of Savransky et al.[7] visited a relatively small number of stars many times, as can be seen by the green and red points in Fig. 3 of Ref. 7. Roughly 20 total slews of the starshade were devoted to third visits or greater, which contribute little to the total yield. In contrast, our code chose not to visit any star more than twice. If one were to take all of the observations spent on visits beyond the second visit and devote those to new stars, the yield estimate of Savransky et al.[7] would have

increased by ~20%. We therefore conclude that the majority of our increase in yield is due to a difference in how revisits are distributed.

The remaining discrepancy in yield can therefore be attributed to time optimization or target selection. To address the first, we turned off exposure time optimization in our code and forced our code to always observe to the noise floor. The resulting yield decreased to 21 exoEarth candidates, implying that exposure time optimization can only explain a 7% change in yield.

A majority of the yield discrepancy has been explained by differences in revisit selection and exposure time optimization. The remaining ~10% yield discrepancy is likely a result of target selection differences. To check, we recalculated our yield while limiting the code to the exact set of stars chosen by Savransky et al.[7] Our yield was reduced to 19 exoEarth candidates when adhering to the target stars chosen by Savransky et al.,[7] suggesting that target selection differences can explain the remaining yield differences.

To determine the source of the target selection differences, we plot the targets selected by our code as well as those selected by all 100 simulations of Savransky et al.[7] in Fig. 6. Savransky et al.[7] explicitly removed late K and M stars due to the short orbital periods of HZ planets around those stars. However, we recalculated our yield while banning stars with $L_\star < 0.2$ and found that our yield decreased by a negligible 3%; low-mass stars are not the source of the discrepancy.

Some of the targets selected by Savransky et al.[7] were identified as giant stars and close-separation binaries by our code. However, if we were to include these stars in our target list, our yield would only increase. Giant stars are therefore not the source of the discrepancy.

Thus, the difference in target selection is primarily a small number of nearby F, G, and K stars that have very short exposure times and were not selected by Savransky et al.,[7] e.g. α Cen B. It may be surprising that a small handful of stars can change the yield by ~10%. However, in this case, η_{Earth} was assumed to be 50%, such that only ~4 nearby stars can add ~2 exoEarth candidates to a yield of ~20.

Of course, it is important to note that our code does not explicitly calculate whether a given set of observations is schedulable. It is therefore possible that some of our yield gain is due to observations that are simply unschedulable. However, given
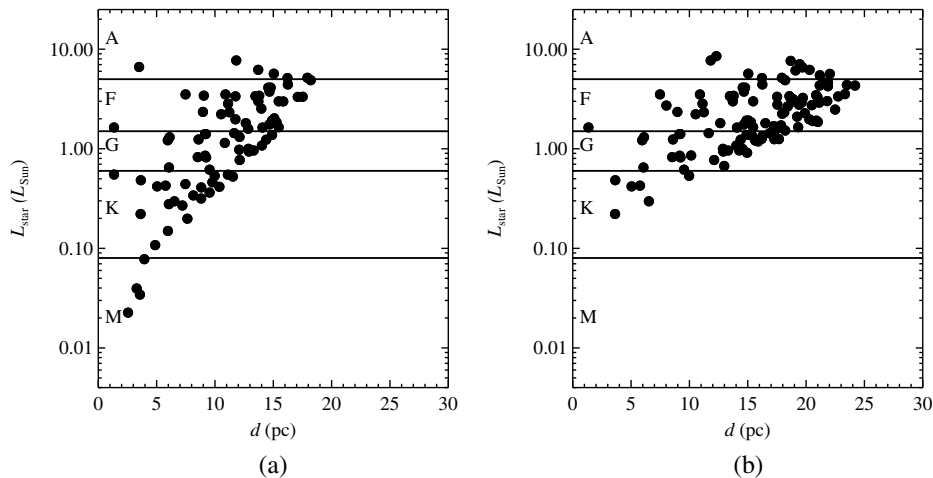

Fig. 6 Target list comparison with Savransky et al.[7] (a) Targets selected by the DRM described in this paper and (b) targets selected at least once in 100 DRM simulations by Savransky et al.[7]

that our selected targets are more-or-less uniformly distributed over the sky, and that ecliptic latitudes > ~45 deg or −45 deg are continuously observable, we find this unlikely.

4.2 Yields for Starshade Missions

We estimated the exoEarth candidate yield for single starshade mission for three different launch mass limits: 3600, 9800, and 25,000 kg, corresponding to the Sun-Earth L2 payload mass limits for Falcon 9, Delta IV Heavy, and SLS Block 1 rockets, respectively.[13] We do not budget for the telescope mass and allow the starshade to consume the full launch mass limit. We calculated the yield as a function of telescope aperture size, IWA, and contrast. We investigated telescope diameters ranging from 2 to 10 m in steps of 1 m, IWAs ranging from 30 to 90 mas in steps of 10 mas, and logarithmically spaced contrasts ranging from 10^{-11} to 10^{-9} for each launch mass limit.

For each combination of D, IWA, and ζ, we interpolated our grid of optimized starshade diameters to determine D_{ss} and z. All other starshade and mission parameters were kept equal to the baseline assumptions listed in Table 2. We note that we did not vary the bandpass of the starshade (0.5 to 1.0 μm), which was motivated scientifically. However, by changing the separation distance of a fixed-size starshade, one can reduce the bandpass to improve the IWA, potentially increasing the yield by a modest amount, albeit at a different set of wavelengths.[18]

For every calculation, we adopted our large baseline thrust of 10 N. This unrealistically large thrust ensured that we were working far from the thrust-limited regime and our yield estimates may therefore be optimistic. We address this assumption in more detail below.

The left panel of Fig. 7 shows the exoEarth candidate yield as a function of telescope aperture size, where the dotted, solid, and dashed lines correspond to launch mass limits of 3600, 9800, and 25,000 kg, respectively. The color of each plotted point corresponds to the starshade diameter as shown by the color bar on the right. Next to each data point, we also list the optimum number of stars and slews chosen by the code, given in the format $n_{slews}/n_{targets}$. Note that the number of slews is equivalent to the number of observations. Observations with zero slews correspond to dry masses that exceeded the payload limits.

In each plot, we report the sensitivity of the yield to changes in the relevant mission parameters, ϕ. For each launch mass limit, we provide ϕ at the baseline values, with each value of ϕ evaluated at the corresponding symbol on the curve. For example, the left panel shows that at the location of the downward pointed triangle, the sensitivity of the yield to changes in D is 0.11. As discussed in Ref. 3, the yield sensitivity is calculated as the fractional change in yield per fractional change in mission parameter, i.e.,

$$\phi(x_0) = \frac{\Delta N_{EC}}{\Delta x} \frac{x_0}{N_{EC}}, \qquad (25)$$

where x is the independent mission parameter and N_{EC} is the exoEarth candidate yield at x_0. Note that $\phi(x_0)$ is equivalent to the exponent of the power law relationship between N_{EC} and x at $x = x_0$, i.e., near $x = x_0$, $N_{EC} \propto x^{\phi(x_0)}$. We note that while the sensitivity metric is equivalent to the exponent of the power law relationship at the location of the red dot, the entirety of the yield curves are clearly not well described by power law relationships.

For launch mass limits of 3600 and 9800 kg, the yield has a moderate dependence on telescope aperture size. For comparison, Stark et al.[4] showed the yield for coronagraph scales roughly as D^2. This moderate dependence results from two facts. First, the telescope aperture impacts the photon collection rate and the PSF size, but it does not affect the IWA as it does for a coronagraph. More importantly, the starshade does not operate entirely in the exposure time-limited regime as a coronagraph mission does, so shorter exposure times have less impact on the total yield. Finally, as the color of the dots indicate, larger apertures require larger starshades at greater distances, which increases fuel use. This latter factor can actually degrade the yield in some cases, as shown by the dotted line in Fig. 7.

Figure 7(a) also tells us about the ideal ratio of $n_{slews}/n_{targets}$. For small D, where the mission is operating closer to the exposure time-limited regime, $n_{slews} > n_{targets}$, indicating that revisits are valuable. As the aperture size increases, exposure times decrease and slews become more expensive, such that the mission is operating closer to the fuel-limited regime. In the fuel-limited regime, $n_{slews} \sim n_{targets}$, indicating that (potentially longer) single visits are more productive.
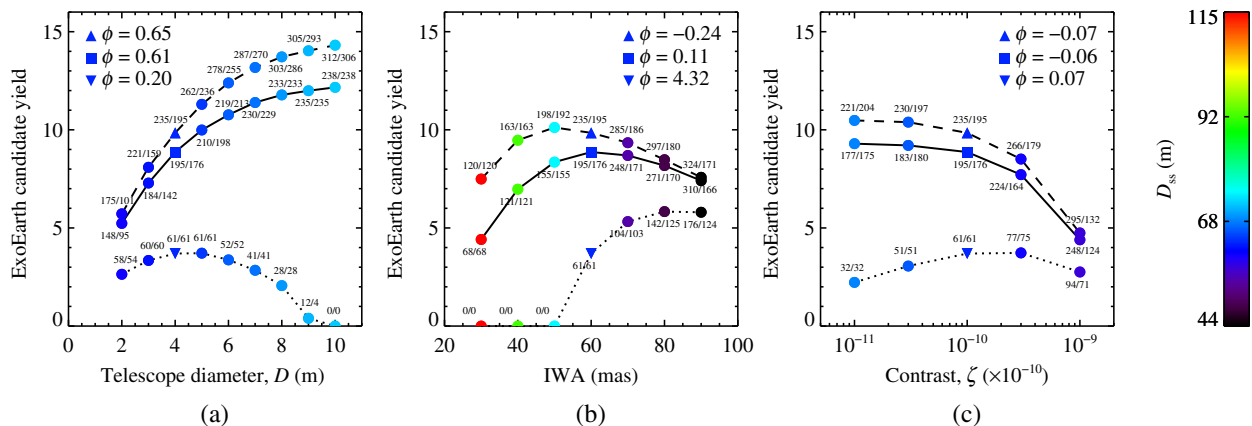

Fig. 7 ExoEarth candidate yield as a function of (a) telescope diameter, (b) IWA, and (c) contrast for the baseline mission assuming launch mass limits of 3600 (dotted line), 9800 (solid line), and 25,000 kg (dashed line). Color indicates starshade diameter and $n_{slews}/n_{targets}$ is listed next to each data point. The sensitivity of yield, ϕ, is given at the baseline parameter values of each curve.

Figure 7(b) shows the yield as a function of IWA for the baseline mission. An optimum IWA is clearly evident. The yield decreases at small IWA because large, massive starshades are required with large separation distances, greatly impacting the number of slews that can be made. On the other hand, the yield is reduced at large IWA because fewer HZs are accessible. As previously discussed, in the fuel-limited regime of small IWA, $n_{slews} \approx n_{targets}$. In the large IWA scenario, $n_{slews} \approx 2n_{targets}$, indicating that revisits are very valuable, a result of operating closer to the exposure time-limited regime and having fewer accessible HZs. We note that the optimum IWA appears to vary with the launch mass limit while all other parameters are kept fixed.

As Fig. 7(c) illustrates, the yield is generally a weak function of contrast. As long as the starshade mass does not dominate the launch mass limit, yield generally decreases as contrast is degraded. This is in spite of the fact that as contrast is degraded, starshade diameter and separation are reduced (as indicated by the color of the data points).

The sensitivity of the yield to changes in mission parameters varies significantly over parameter space. For example, Fig. 8 shows the same sensitivity curves as Fig. 7 when adopting baseline values of $D = 6$ m, IWA = 50 mas, and $\zeta = 3 \times 10^{-10}$. However, the qualitative shape of the yield curves remains similar.

Figure 9 shows the exoEarth candidate yield over four dimensions. We plot two-dimensional contours of yield as a function of IWA and contrast for four different telescope apertures (rows) and three different launch mass limits (columns). The lower-left corner of each plot shows reduced yield—this region corresponds to high contrast and small IWA such that the starshade payload mass can dominate the launch mass limit, leaving little room for fuel. An optimal IWA clearly exists, and it does not appear to be a strong function of D; the launch mass limit appears to predominantly determine the optimal starshade IWA.

Figure 10 shows exoEarth candidate yield as a function of launch mass limit for two different IWAs and three different telescope apertures. The purple dotted, black solid, and green dashed lines correspond to telescope apertures of 2, 4, and 6 m, respectively. Larger launch mass limits correspond to larger initial fuel masses. In all cases, the yield ramps up quickly at low mass limits and then flattens out such that yield becomes a weak function of initial fuel mass. This flattening of the yield curves simply reflects the exponential nature of the rocket mass equation; to achieve an additional slew, we must add fuel, but that additional fuel adds mass that we must move.

Figure 11 shows exoEarth candidate yields as a function of total mission length for the baseline mission, where launch mass limits of 3600, 9800, and 25,000 kg are indicated by dotted, solid, and dashed lines, respectively. Although yield is a moderately weak function of total mission length, it is on par with a coronagraph's relationship between yield and total exposure time.[4] This result may be somewhat surprising, as starshades do not operate entirely in the exposure time-limited regime and yield curves shown above have been relatively insensitive to the photon collection rate. However, there is a difference between exposure time (photon collection rate) and total mission time, and Fig. 11 shows the latter, which can be used to lengthen slew time and reduce fuel consumption per slew.

Because the yield of a starshade depends on future propulsion technology, the sensitivity of the yield to propulsion parameters is an important check on starshade performance. The left and right panels of Fig. 12 show exoEarth candidate yield as a function of I_{slew} and I_{sk}, respectively, for each of the assumed launch mass limits. The yield for our baseline starshade mission is weakly dependent on the specific impulse for both slewing and stationkeeping; a factor of 2 change in the specific impulse (or any other factor in the slew fuel use expression) impacts the baseline yield by ~15%. We note that these scaling relationships are not the result of operating far from the thrust-limited regime (we find similar sensitivity with $\mathcal{T} = 1$ N), but will change depending on the location in parameter space. For example, a smaller aperture paired with a more massive starshade payload would lead to a greater dependence on specific impulse.

All of the calculations in this paper assume a thrust of 10 N such that the starshade is operating far from the thrust-limited regime. To illustrate this regime, Fig. 13 shows the yield as a function of thrust. For a launch mass limit of 3600 kg, the yield is flat for $\mathcal{T} > 0.5$ N. For launch mass limits of 9800 and 25,000 kg, the yield is thrust-limited for $\mathcal{T} \lesssim 2$ N and $\mathcal{T} \lesssim 8$ N, respectively; our adopted thrust of 10 N ensures that all simulations do not operate in the thrust-limited regime. More realistically, lower thrust values of ~1 N could reduce the baseline mission yields by ~15%.

Finally, Fig. 14 shows the yield of our baseline mission as a function of the astrophysical assumptions η_{Earth}, average exoEarth candidate albedo, and median exozodi level for the three different launch mass limits. Comparing Fig. 14 to Fig. 14 of Stark et al.,[4] we see that our baseline starshade is
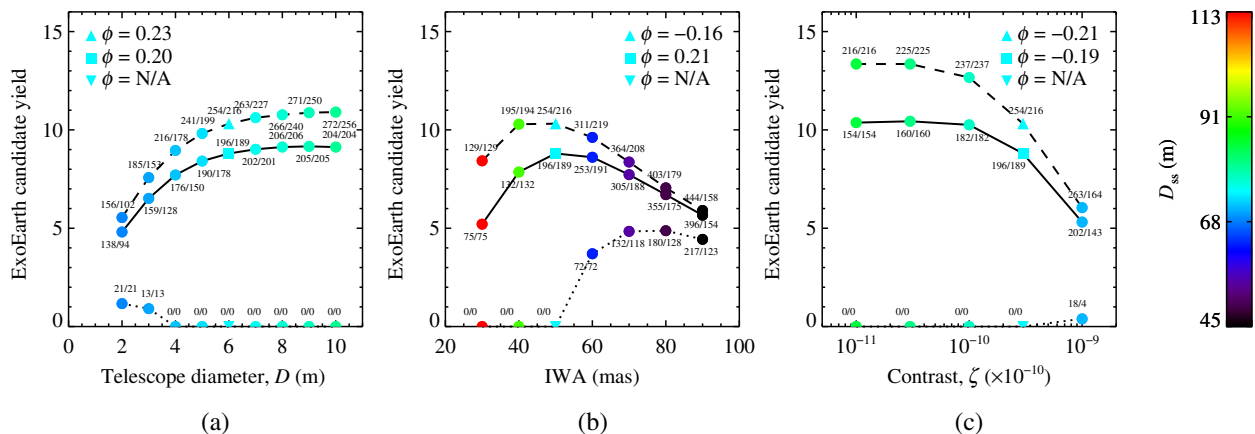

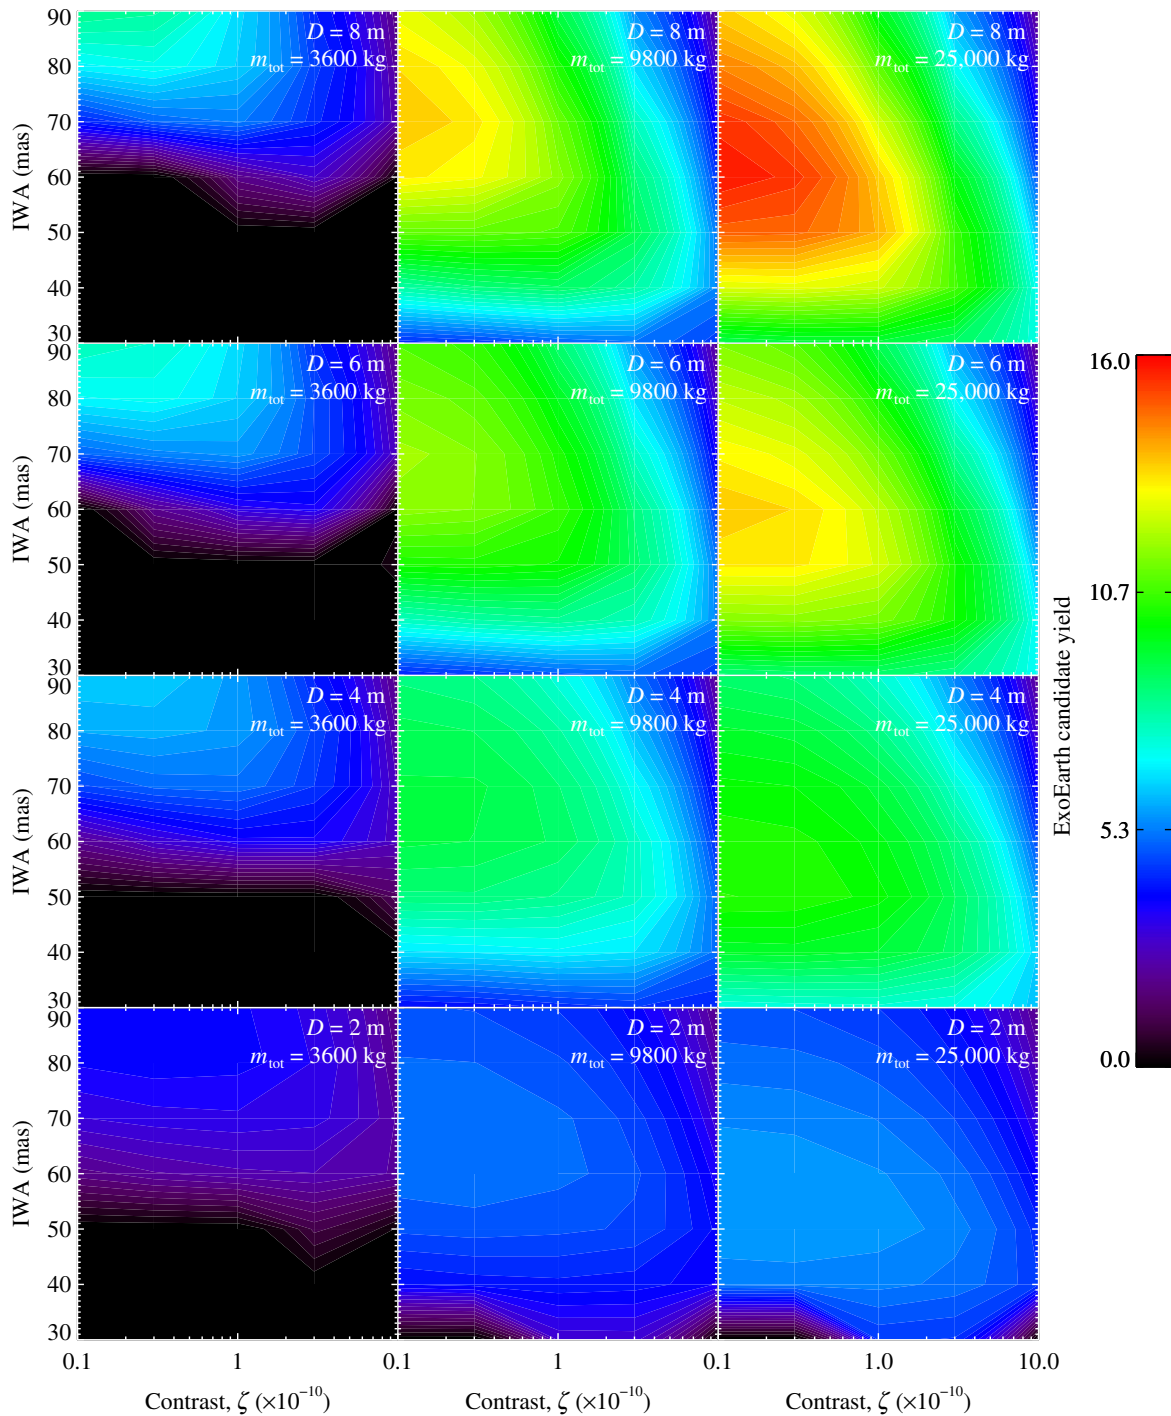Fig. 8 Same as Fig. 7, but for values of $D = 6$ m, IWA = 50 mas, and $\zeta = 3 \times 10^{-10}$.

Fig. 9 ExoEarth candidate yield contours as a function of IWA and contrast. Rows correspond to aperture sizes ranging from 2 to 8 m and columns correspond to total initial launch masses ranging from 3600 to 25,000 kg.

less sensitive to both the average geometric albedo of exoEarth candidates and exozodi level, consistent with the starshade not operating in the exposure time-limited regime.

5 Discussion

The scaling relationships for a starshade mission's exoEarth candidate yield are much more variable than those for a corona-graph and depend greatly upon the specific set of mission parameters being considered—specifically, the degree to which the starshade is operating in the fuel- or exposure time-limited regime. That being said, there are several obvious conclusions given the yield curves presented above.

First, starshades do not have any mission parameters that can impact yield to the same degree that aperture size impacts a coronagraph's yield. Although starshade yield can be extremely sensitive to IWA ($\propto$ IWA$^{4.3}$), this is limited to isolated regions of
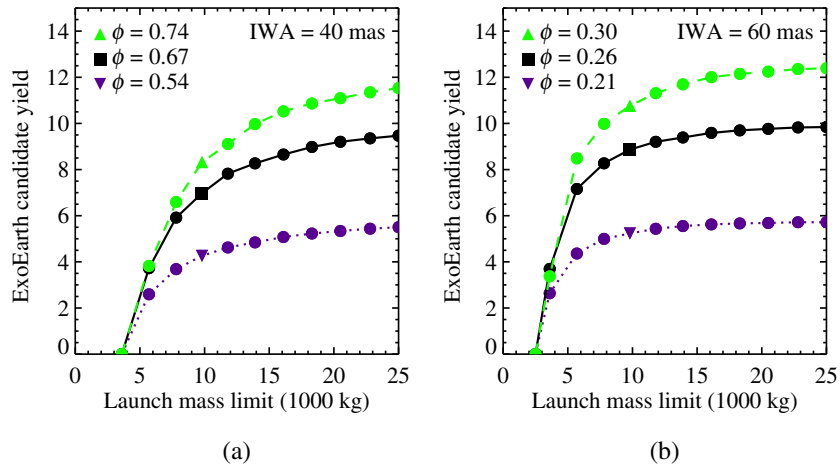
Fig. 10 ExoEarth candidate yield as a function of total launch mass. Lines correspond to telescope apertures of 2 (purple dotted), 4 (black solid), and 6 m (green dashed).
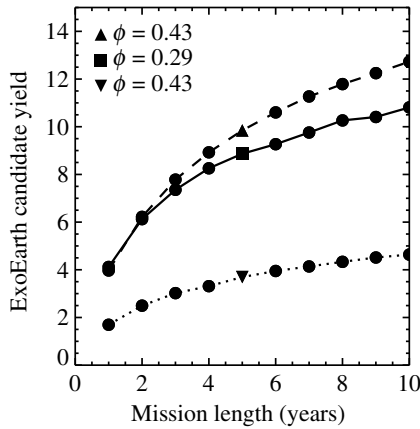

Fig. 11 ExoEarth candidate yield as a function of mission length for launch mass limits of 3600 (dotted line), 9800 (solid line), and 25,000 kg (dashed line).
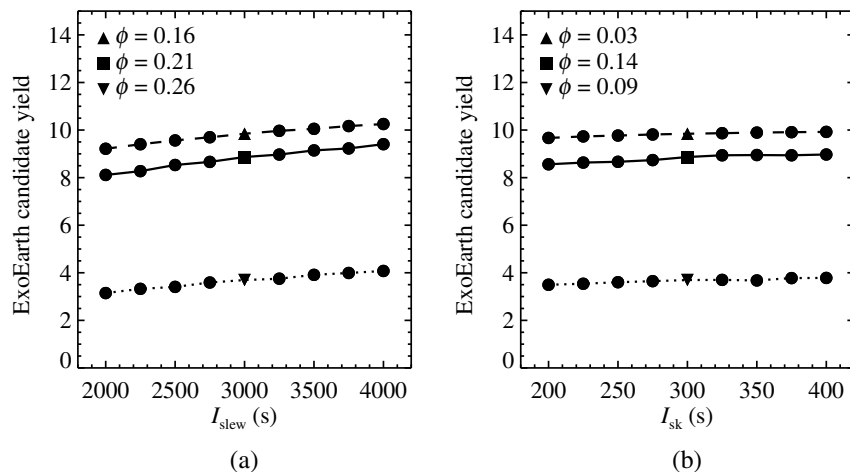
parameter space in which the starshade design is far from the optimal IWA. The optimal IWA limits the yield gain that can be realized from this relationship.

Assuming a mission adopts an IWA close to the optimum IWA, we found yield could be as sensitive to aperture size as $\propto D^{0.6}$, but the exponent depends significantly on the mission parameters and is generally <0.6. However, this relationship also cannot increase yield without bound, and the yield curves eventually turn over at larger aperture sizes as starshade diameter, distance, and mass become too great. Because of this modest dependence on aperture size, the yield of a starshade mission with a 4 m aperture is only ~20 to 50% greater than that of a 2.4 m aperture, assuming a fixed IWA and contrast.

The yield of a starshade mission may be most sensitive to a parameter we did not consider: the number of starshades. Given the exponential nature of the rocket mass equation and the yield curves shown in Fig. 10, it is clearly more advantageous to launch two starshades with half the total launch mass than one starshade with the full launch mass limit, if possible. We leave detailed calculations of multiple starshades to future work, but speculate that doubling the number of starshades should produce slightly less than a doubling of yield as long
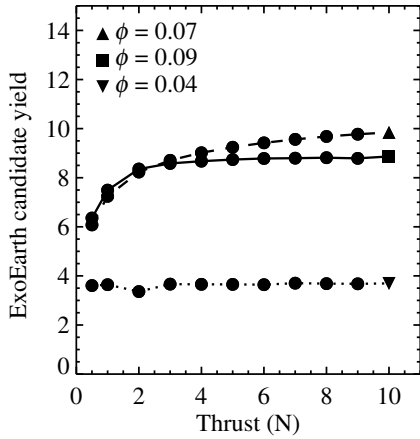

Fig. 12 ExoEarth candidate yield as a function of propulsion parameters (a) I_{slew} and (b) I_{sk}. Dotted, solid, and dashed lines correspond to launch mass limits of 3600, 9800, and 25,000 kg, respectively.

Fig. 13 ExoEarth candidate yield as a function of thrust. Dotted, solid, and dashed lines correspond to launch mass limits of 3600, 9800, and 25,000 kg, respectively.

as the two starshade payloads do not dominate the total launch mass limit. For example, whereas our baseline yield shown by the black square in the right panel of Fig. 10 achieved ∼9 exoEarth candidates, two such starshades may achieve ∼15 exoEarth candidates.

Under our baseline astrophysical assumptions, the largest yield estimate we calculated is ≈16 exoEarth candidates. We were unable to find an set of mission parameters that produced an yield of ∼30 to 60 exoEarth candidates, the yield goals suggested by Stark et al.[3,4] As shown in Fig. 14, the absolute yield of a starshade mission of course depends strongly on the value of η_{Earth}. If our best estimate of η_{Earth} increases by a factor of a few, a starshade mission could conceivably achieve such yields. We note that for consistency, we adopted $\eta_{Earth} = 0.1$, as was used in Ref. 4. Stark et al.[4] adopted the 1σ lower limit on η_{Earth} from recent planet occurrence rate estimates, explicitly accounting for uncertainty in this important parameter. The expected η_{Earth} value will likely evolve with time as our estimates improve, though it is unclear how or whether the 1σ lower limit on η_{Earth} will change.

Assuming the fuel mass scaling relationships presented in this paper hold true when performing detailed dynamical calculations for the starshade and telescope orbits, the yield curves

shown in this paper may be optimistic. We assumed a very large value of thrust (10 N), specifically chosen to operate far from the thrust-limited regime; real starshades may indeed be thrust-limited. We also did not budget for telescope mass and explicitly assumed the telescope and starshade launched separately.

Although a mission with a single starshade appears unable to achieve the same yield as a coronagraph mission under our baseline astrophysical assumptions, it is interesting to note that the starshade is less sensitive to photometric noise sources as well as parameters controlling the planet's photon collection rate. As shown by Fig. 14, the yield is a weak function of average exoEarth candidate albedo and a very weak function of median exozodi level. Starshades are more robust to some astrophysical sources of uncertainty.

Finally, we discuss two assumptions made by Stark et al.[4] that become important when comparing the performance of starshades to coronagraphs. First, Stark et al.[4] assumed, as we do in this paper as well, that the spectral characterization time required to achieve an $R = 50$, $S/N = 5$ spectrum is only devoted to planets that are actually exoEarth candidates (Earth-sized planets in the HZ). This implicitly assumes that we can differentiate between exoEarth candidates and other sources (background objects, crescent phase giant planets, and so on) via a combination of color information, debris disk/orbital inclination, absolute brightness, and so on, and that we can obtain this information without affecting the initial detection exposure time (this implies that we may require simultaneous multiband imaging). (This may sound like no time can be wasted on background objects or other types of planets, an optimistic assumption. However, because of the statistical nature of the detections and the most probable phase of finding an exoEarth, if one performs the detections with an energy-resolving detector, we will likely obtain some spectra for free, as shown by the distribution of realized S/N in Fig. 13 of Stark et al.[4]) This assumption may be more valid for a starshade given the continuous spectral coverage from 0.5 to 1 μm, and the negative impacts of such an assumption on yield would likely be less for the starshade than coronagraph since it does not operate in the exposure time-limited regime.

Second, Stark et al.[4] assumed, as we do in this paper, that no revisits are performed with the intent of establishing an orbit. If revisits are required to establish an HZ orbit for a given target,
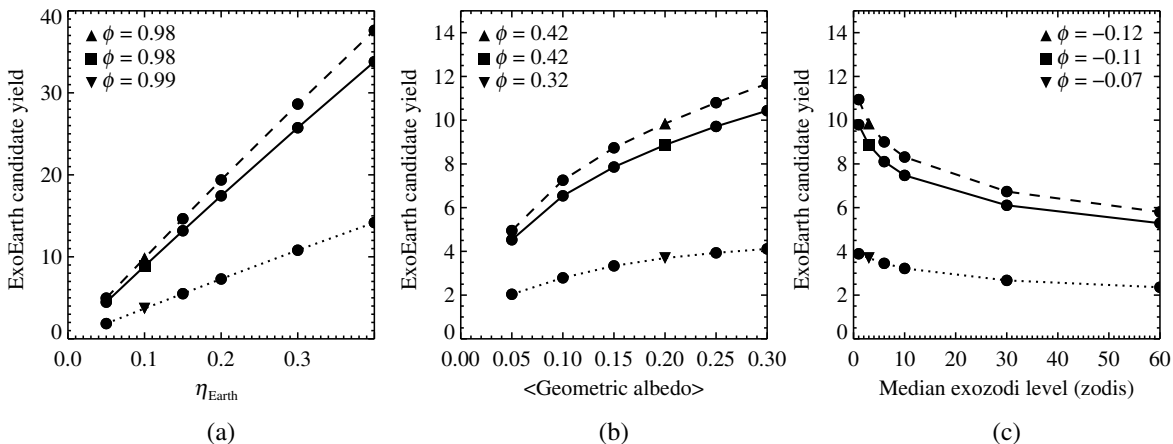

Fig. 14 ExoEarth candidate yield as a function of astrophysical parameters for launch mass limits of 3600 (dotted line), 9800 (solid line), and 25,000 kg (dashed line).

the yield of the fuel-limited starshade would be impacted more significantly than that of the coronagraph. Assuming a starshade mission can perform ~150 slews, if each exoEarth candidate requires approximately five observations for orbit determination, a substantial amount of fuel would be devoted to such follow-up. The coronagraph, on the other hand, would either devote some of its search time to such follow-up (which may not significantly impact yield as the yield of a coronagraph mission is a moderately weak function of total exposure time) or could simply perform such observations during an extended mission.

6 Conclusions

We presented analytic expressions for starshade fuel use and combined these with our yield optimization methods to estimate maximized exoEarth candidate yields for starshade-based missions, adopting the same astrophysical assumptions as Stark et al.[4] did for coronagraph-based missions. We find that starshade yields are maximized when the starshade operates somewhere between the time- and fuel-limited regimes. As a result, starshade yields are less sensitive to astrophysical photometric noise sources (e.g., exozodi and unsuppressed starlight) and average exoEarth albedo; starshades are more robust to some astrophysical uncertainties, with the notable exception of $\eta_\oplus$. However, this also results in starshade yields that are less sensitive to mission parameters controlling the planet photon collection rate (e.g., telescope aperture). Under the assumption $\eta_\oplus = 0.1$, we are unable to find a set of mission parameters that provides a yield of several dozen exoEarth candidates for a five-year mission with a single starshade. If $\eta_\oplus = 0.3$, a yield of several dozen becomes possible, but only with a telescope aperture $\gtrsim 7$ m, an ~70 m diameter starshade, a thrust of $\gtrsim 2$ N, and the launch mass limits of a Delta IV Heavy dedicated solely to the starshade. If $\eta_\oplus \gtrsim 0.4$, a yield of several dozen could be achieved with a telescope aperture of ~4 m. Flying multiple starshades, which we do not model in this paper, would relax these requirements.

References

1. R. A. Brown, "Single-visit photometric and obscurational completeness," *Astrophys. J.* **624**, 1010–1024 (2005).
2. R. A. Brown and R. Soummer, "New completeness methods for estimating exoplanet discoveries by direct detection," *Astrophys. J.* **715**, 122–131 (2010).
3. C. C. Stark et al., "Maximizing the exoEarth candidate yield from a future direct imaging mission," *Astrophys. J.* **795**, 122 (2014).
4. C. C. Stark et al., "Lower limits on aperture size for an exoEarth detecting coronagraphic mission," *Astrophys. J.* **808**, 149 (2015).
5. W. Cash, "Detection of Earth-like planets around nearby stars using a petal-shaped occulter," *Nature* **442**, 51–53 (2006).
6. R. J. Vanderbei, E. Cady, and N. J. Kasdin, "Optimal occulter design for finding extrasolar planets," *Astrophys. J.* **665**(1), 794–798 (2007).
7. D. Savransky, N. J. Kasdin, and E. Cady, "Analyzing the designs of planet-finding missions," *Publ. Astron. Soc. Pac.* **122**, 401–419 (2010).
8. M. C. Turnbull et al., "The search for habitable worlds. 1. The viability of a starshade mission," *Publ. Astron. Soc. Pac.* **124**, 418–447 (2012).
9. R. Trabert et al., "Design reference missions for the exoplanet starshade (exo-s) probe-class study," *Proc. SPIE* **9605**, 96050Y (2015).
10. R. K. Kopparapu et al., "Habitable zones around main-sequence stars: new estimates," *Astrophys. J.* **765**, 131 (2013).
11. C. Leinert et al., "The 1997 reference of diffuse night sky brightness," *Astron. Astrophys. Suppl. Ser.* **127**, 1–99 (1998).
12. S. Seager et al., "Exo-S: starshade probe-class exoplanet direct imaging mission concept final report," March 2015, http://exep.jpl.nasa.gov/stdt/Exo-S_Starshade_Probe_Class_Final_Report_150312_URS250118.pdf (23 May 2016).
13. N. Rioux, "A future large-aperture UVOIR space observatory: reference designs," *Proc. SPIE* **9602**, 960205 (2015).
14. E. Cady et al., "Progress on the occulter experiment at Princeton," *Proc. SPIE* **7440**, 744006 (2009).
15. E. Cady, "Nondimensional representations for occulter design and performance evaluation," *Proc. SPIE* **8151**, 815112 (2011).
16. E. Kolemen and N. J. Kasdin, "Optimal trajectory control of an occulter based planet finding telescope," 2007, https://www.princeton.edu/~ekolemen/publications/kolemen_gnc_feb_2007_breckenridge.pdf (23 May 2016).
17. S. L. Hunyadi, S. B. Shaklan, and R. A. Brown, "The lighter side of TPF-C: evaluating the scientific gain from a smaller mission concept," *Proc. SPIE* **6693**, 66930Q (2007).
18. S. Seager et al., "The exo-s probe class starshade mission," *Proc. SPIE* **9605**, 96050W (2015).

Biographies for the authors are not available.