A direct comparison of exoEarth yields for starshades and coronagraphs

Christopher C. Stark[a], Eric Cady[b], Mark Clampin[c], Shawn Domagal-Goldman[c], Doug Lisman[b], Avi M. Mandell[c], Michael W. McElwain[c], Aki Roberge[c], Tyler D. Robinson[d], Dmitry Savransky[e], Stuart B. Shaklan[b], and Karl R. Stapelfeldt[c]

[a]Space Telescope Science Institute, 3700 San Martin Dr, Baltimore, MD 21218, USA; cstark@stsci.edu
[b]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA
[c]NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA
[d]Department of Astronomy and Astrophysics, University of California, Santa Cruz, CA 95064, USA
[e]Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853, USA

ABSTRACT

The scale and design of a future mission capable of directly imaging extrasolar planets will be influenced by the detectable number (yield) of potentially Earth-like planets. Currently, coronagraphs and starshades are being considered as instruments for such a mission. We will use a novel code to estimate and compare the yields for starshade- and coronagraph-based missions. We will show yield scaling relationships for each instrument and discuss the impact of astrophysical and instrumental noise on yields. Although the absolute yields are dependent on several yet-unknown parameters, we will present several limiting cases allowing us to bound the yield comparison.

Keywords: telescopes — methods:numerical — planetary systems

1. INTRODUCTION

A key science driver for future space telescopes is the direct detection and characterization of extrasolar planets that are potentially Earth-like (exoEarth candidates). To directly image such faint planets, the glare from their host stars must be reduced by a factor of ~ 10 billion.

Currently, two proposed technologies are being investigated to achieve this factor of 10^{-10} contrast: coronagraphs, which remove starlight inside of the telescope, and starshades, which fly in formation many millions of meters from the telescope and occult starlight externally. These instruments have differing capabilities and limitations. For example, coronagraphs will have lower throughput than a starshade, while starshades will be fuel-constrained and less nimble than a coronagraph-based mission.

The decision as to which instrument to use for the detection and characterization of exoEarth candidates will be informed by their technical feasibility, type and range of science that they can produce, and the quantity of that science. In this paper, we address the latter two of these. We attempt an apples-to-apples comparison of starshade- and coronagraph-based missions, calculating exoEarth candidate yield for both. We use a novel yield calculation code designed to maximize the yield for both mission concepts. We examine how this yield responds to astrophysical uncertainties as well as differences in the scientific output and approach to characterizing exoEarth candidates.

Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave, edited by Howard A. MacEwen, Giovanni G. Fazio, Makenzie Lystrup, Proc. of SPIE Vol. 9904, 99041U · © 2016 SPIE · CCC code: 0277-786X/16/$18 · doi: 10.1117/12.2233201

2. METHODS

To estimate exoEarth candidate yield for coronagraph- and starshade-based missions, we use the altruistic yield optimization (AYO) code initially developed by Ref. 1. For a detailed description on how AYO is implemented for a coronagraph, see Refs. 1 and 2. For starshade yields, see Ref. 3. Refs. 1–3 also detail exposure time calculations, assumptions about noise sources, etc. Here we provide a brief, high level overview of how yields are calculated and maximized for both instrument concepts.

In this manuscript, we limit our investigation to missions that use either a coronagraph or a starshade. We do not consider missions with multiple starshades, nor do we consider hybrid missions that use both coronagraphs and starshades. For both instrument concepts, we assume that we do not know anything about the planetary systems in advance and the mission is responsible for the initial search for exoEarth candidates as well as follow-up spectral characterization.

Because we assume no knowledge of which stars host exoEarth candidates, we cannot determine a yield based on known planets. Thus, yield estimation is a probabilistic calculation. Instead of determining whether a planet is detected, we calculate the probability of detecting a planet in a given observation, if such a planet exists. This quantity, called the completeness of the observation, is defined in Ref. 4.

To calculate completeness, we distribute a large number of synthetic Earth twins around each star in the Hipparcos catalog within 50 pc, resulting in a "cloud" of exoEarths around each star. We distribute planets over all orbits, orientations, and phases consistent with the planet residing in an assumed habitable zone. We illuminate each synthetic exoEarth with starlight and calculate the planet brightness. For an assumed set of mission parameters, we then calculate the exposure time of every synthetic planet in each planet cloud. For a given detection metric, we can then determine the fraction of the cloud that is detectable within a given exposure time—this quantity is the completeness. The yield for a single observation is simply the probability of detecting a planet if it exists (completeness) times the probability that the planet exists ($\eta_\oplus$).

During the initial visit to a system, some of the synthetic planets may be too faint or too close to their host stars to meet our detection criteria. One can increase the completeness of a given system, and therefore the yield of a mission, by performing multiple visits to each star. By revisiting a star at a later time, planets that were previously undetectable (crescent phase) may evolve in their orbits to more favorable phases (gibbous). Our code includes these revisits and determines the optimum delay time between visits based on the remaining completeness possible.[2]

To calculate the total yield of a mission, we simply add up the yield of each observation until the limiting resource of the mission has been exhausted. For coronagraphs and starshades, these limiting resources are very different. Here we describe these limiting resources and how they are optimally distributed to maximize yield.

2.1 Coronagraph-based Missions

We assume a coronagraph-based mission is designed given some nominal mission lifetime and that consumables like propellent are not a limiting resource. As such, the only limiting resource for a coronagraph-based mission is time. Under this assumption, the maximization of yield is a relatively straight-forward process of optimizing the exposure times of each observation.

For each observation, we calculate the completeness C as a function of exposure time τ. The derivative of this curve, $dC/d\tau$, tells us the return on our investment that we'll receive if we invest a small amount of time in the observation. If we hold an "auction" and award exposure time in small packets to the observations that have the largest $dC/d\tau$, the end result will be that all observations will have the same slope.[5] We can therefore guess the optimal completeness curve slope ahead of time and evaluate the exposure time of each observation at this slope. For each optimal slope guess, we prioritize targets based on their benefit to cost ratio, C/τ, and down-select to those observations that fit within an assumed exposure time budget. By varying our guess of the optimal completeness curve slope, we can find the peak yield. In practice, we must also optimize the delay time between revisits during this process (for details, see Ref. 2).

The result of this optimization is a set of stars and exposure times that maximize the yield for the mission. Each star is assigned a number of visits, with exposure times and delay times given for each visit. The resulting

observation plan is static; if the astrophysical assumptions are correct and the plan is executed exactly, the maximum yield will be achieved. In reality, we don't know many astrophysical parameters prior to launching the mission. Realistic observation plans will "learn" as they are executed, adapting to astrophysical conditions. Previous efforts have shown a modest increase in yield if one can learn about the distribution of exozodi levels.[2] In this paper, we limit ourselves to static observation plans and bound the yield by examining a broad range of astrophysical assumptions.

2.2 Starshade-based Missions

In addition to a limited mission lifetime, starshade-based missions will have another primary limiting resource: fuel to keep the starshade and telescope aligned (stationkeeping) and fuel to move the starshade to the next target (slewing). Thus to optimize the observation plan for a starshade-based mission, we must optimally distribute both exposure time and fuel, and optimally balance these two resources. This process is complicated by the fact that the fuel use for a starshade is a traveling salesman-type problem, making on-the-fly decision making more critical.

Ref. 3 proposed a work-around, allowing rapid maximized yield calculations for a starshade. The methods proposed required several significant assumptions. First, the fuel required to visit many targets over a full mission can be approximated by the total number of slews, each evaluated using the median slew distance and median slew time. Second, we can ignore schedulability and assume that the optimized observation plan can be scheduled without a significant impact to the required fuel. Third, the difference in solar radiation pressure felt by the starshade and telescope can be mitigated without using significant fuel. Under these assumptions, Ref. 3 presented analytic scaling relationships relating a starshade's fuel use to the number of slews and stars observed, total exposure and mission time, and mission parameters describing the starshade's size, mass, optical performance, and propulsion system.

Using these scaling relationships, Ref. 3 showed that for a given set of astrophysical and mission parameters, yield is controlled by a 5-dimensional parameter space that can be explored to maximize the yield. Briefly, in addition to the completeness curve slope, one must optimize the number of unique stars visited, the number of slews, the slew efficiency, and the balance between the time and fuel resources. For a detailed description of how this is performed, see Ref. 3.

3. SCIENTIFIC GOALS

In this paper, our goal is to use the code of Refs. 2 and 3 to perform an apples-to-apples comparison of the yields for coronagraph- and starshade-based missions. The yields of these missions will depend on the scientific goals for the mission. Here, we adopt the same "minimally-compelling" science goals adopted by Refs. 2 and 3: we require each detected exoEarth to be detected at roughly visible wavelengths and spectrally characterized at 1 μm to search for evidence of a water vapor absorption. Because the stars that will be selected for observation are typically brightest in V band, we explicitly assume detections must include observations at 0.55 μm.

Because spectral characterization time will be tens of times longer than initial detection, spectra are relatively costly. Currently it is unclear how we would decide whether to spectrally characterize a detected point source. Ideally we'd only characterize the exoEarth candidates, but the chances of confusion with other planet types and background objects are likely high. Because the coronagraph and starshade are very different missions with different capabilities, the method by which we distinguish exoEarth candidates from other objects could greatly impact the yield. Thus, we will examine 4 different observational approaches to achieving these scientific goals. These approaches are as follows:

Method 1. Detection only, assuming follow-up spectral characterization occurs after the nominal mission.

Method 2. 3-band detection and immediate spectral characterization of only the exoEarth candidates, assuming we can differentiate perfectly between all planet types and background objects via planet color, photometry, projected stellar separation, and debris disk information.

Method 3. Detection and immediate spectral characterization, assuming we cannot differentiate between anything (spectral characterization required after every observation).

Method 4. Detection, followed by orbit determination (6 total visits to system required), followed by spectral characterization of only the exoEarth candidates.

For Method 1, we assume the full bandwidth of the instrument(s) can contribute to a single broadband detection. As a result, the mission may not gain any color information on any of the objects. Follow-up spectral characterization after the nominal mission, or by a future mission, may therefore be somewhat uncertain. However, the yield associated with this method should represent a maximum yield and serves as an important comparison to understand the impact of Methods 2–4.

For Method 2, we assume perfect differentiation between exoEarth candidates and all other point sources prior to selecting objects for spectral characterization. To enable this, we assume that 2 bands of color information must be obtained, a reasonable minimum requirement. Thus we divide the full bandwidth of our instruments(s) accordingly.

For Method 3, we assume we cannot differentiate anything and must obtain a spectrum after every observation. Thus we allow the initial detection to use the full bandwidth of the instrument(s).

Finally, for Method 4, we assume each star must be observed multiple times to confirm proper motion with the star and constrain the planet's orbit before spectral characterization occurs. We allow the full bandwidth of the instrument(s) to contribute to the detection, noting that the multiple visits should provide sufficient photons for color information prior to spectral characterization. The minimum number of visits required to adequately constrain the orbit is unknown—we make the reasonable assumption that at least 6 visits are required to each star.

4. ASSUMPTIONS

Our astrophysical assumptions are identical to those made by Refs. 2 and 3 and are presented in Table 1. We assume both missions are attempting to maximize the yield of Earth-twins located in the habitable zone (HZ) of their host stars. We adopt the habitable zone given by Ref. 6 for a Sun-like star and scale the boundaries with the square root of the stellar luminosity. Target stars are drawn from the Hipparcos catalog and vetted for binarity.[1]

Our detection requirements are also identical to those made by Refs. 2 and 3. We require a S/N of 7 for detection and calculate the required exposure time based on the photon noise limit, as described in Refs. 1 and 2. Spectral characterizations are calculated at a wavelength of $\lambda = 1.0$ μm with a spectral resolution $R = 50$ and a S/N of 5 per spectral channel to search for a water absorption feature. Table 2 summarizes these assumptions for both mission concepts.

For our baseline coronagraph, we make similar, though not identical, assumptions to Ref. 2. We assume 3 coronagraphs, each with 20% bandwidth, roughly spanning the full range from 0.55–1.0 μm. We adopt 3 coronagraphs, instead of 2 adopted by Ref. 2, to provide 2-color information on each detected point source. These three coronagraphs all have identical inner working angles (IWAs) measured on the sky, equal to $3.6\lambda/D$ at $\lambda = 0.55$ μm and $2\lambda/D$ at $\lambda = 1.0$ μm, such that any planets detectable at the shortest wavelengths are also detectable at the longest. We assume identical end-to-end throughput of 20% for each of these coronagraphs. The contrast of the two short-wavelength coronagraphs is assumed to be 10^{-10} while the contrast of the longest-wavelength coronagraph is assumed to be 5×10^{-10}, to compensate for the fact that it must have a relatively small IWA in units of λ/D. We assume all 3 coronagraphs can be operated in parallel.

Like Ref. 2, we assume a 5 year mission with 2 years devoted to exoplanet science, of which 1 year is untracked overheads, resulting in only 1 year of total exposure time for exoEarth candidate search/characterization. We assume zero detector noise. We assume a larger outer working angle (OWA) than Ref. 2, adopting $75\lambda/D$ instead of $15\lambda/D$. This larger OWA negligibly impacts the yield and is ultimately unimportant. Table 3 provides a list of our coronagraph mission parameters.

For our baseline starshade mission, we make similar, though not identical, assumptions to Ref. 3. These assumptions are summarized in Table 4. Like our coronagraph mission, we assume a 5 year mission lifetime. Unlike the coronagraph mission, we allow the starshade to use as much of this time for exoEarth candidate

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*	Semi-major axis (uniform in $\log a$)
e	0	Eccentricity (circular orbits)
$\cos i$	$[-1, 1]$	Cosine of inclination (uniform distribution)
ω	$[0, 2\pi)$	Argument of pericenter (uniform distribution)
M	$[0, 2\pi)$	Mean anomaly (uniform distribution)
Φ	Lambertian	Phase function
A_G	0.2	Geometric albedo of planet from 0.55–1 μm
z	23 mag arcsec^{-2}†	Average V band surface brightness of zodiacal light
x	22 mag arcsec^{-2}‡	V band surface brightness of 1 zodi of exozodiacal dust
n	3	Number of zodis for all stars

*a given for a solar twin. The habitable zone is scaled by $\sqrt{L_\star/L_\odot}$.

†Due to the pointing constraints, for a starshade we calculate the zodiacal brightness for each target star at a solar elongation of 60°, resulting in a median value of 22 mag arcsec^{-2}.

‡For Solar twin. Varies with spectral type—see Appendix C in Ref. 1.

Table 2. Baseline Detection Requirements

Parameter	Value	Description
$\mathrm{S/N_d}$	7	Signal to noise ratio required for broadband detection of a planet
λ_c	1.0 μm	Wavelength for spectral characterization
R_c	50	Spectral resolving power required for spectral characterization
$\mathrm{S/N_c}$	5	Signal to noise ratio per spectral channel required for spectral characterization

Table 3. Baseline Coronagraph Mission Parameters

Parameter	Value	Description
$\Sigma\tau$	1 yr	Total exposure time of the mission
X	0.7	Photometry aperture radius in λ/D
Υ	0.69	Fraction of Airy pattern contained within photometry aperture
Ω	$b\pi(X\lambda/D)^2$ radians	Solid angle subtended by photometry aperture
b	1.0	Areal broadening of the planet's PSF
T	0.2	End-to-end facility throughput for all coronagraphs
$\Delta\text{mag}_{\text{floor}}$	27.5	Systematic noise floor (faintest detectable point source)
$\text{CR}_{\text{b,detector}}$	0	Detector noise count rate for all coronagraphs
		Coronagraph 1
λ_1	0.55 μm	Central wavelength
$\Delta\lambda_1$	0.11 μm	Bandwidth for detection
ζ_1	10^{-10}	Uniform raw contrast level, relative to theoretical Airy pattern peak
IWA	$3.6\lambda_1/D$	Inner working angle
OWA	$75\lambda_1/D$	Outer working angle
		Coronagraph 2
λ_2	0.725 μm	Central wavelength
$\Delta\lambda_2$	0.145 μm	Bandwidth for detection
ζ_2	10^{-10}	Uniform raw contrast level, relative to theoretical Airy pattern peak
IWA	$2.8\lambda_2/D$	Inner working angle
OWA	$75\lambda_2/D$	Outer working angle
		Coronagraph 3
λ_3	1.0 μm	Central wavelength
$\Delta\lambda_3$	0.2 μm	Bandwidth for detection
ζ_3	5×10^{-10}	Uniform raw contrast level, relative to theoretical Airy pattern peak
IWA	$2.0\lambda_3/D$	Inner working angle
OWA	$75\lambda_3/D$	Outer working angle

exposures as is desired. The total exposure time for exoEarth candidate exposures is ultimately determined by an optimized balance between slew time and exposure time. Although up to 5 years are allowed, typically only 1–1.5 years is selected. We ignore the time for overheads associated with telescope repointing, initial starshade alignment, settling, etc., assuming they will be negligibly small.

We assume a single, broad bandpass for the starshade spanning $\lambda = 0.5$ μm to 1.0 μm, which we will later artificially break into three separate bandpasses when adopting observation Method 2, as this method assumes 2-color information on each detected point source. We assume a single uniform raw contrast of 10^{-10} across the whole bandpass, a relatively high end-to-end throughput of 65%, and a fixed IWA of 60 mas. This combination of bandpass, contrast, and IWA, along with an assumed telescope aperture, define the starshade diameter. We adopt minimized starshade diameters calculated for Ref. 3. We adopt the same starshade mass scaling relationship as Ref. 3.

We assume the starshade is launched in its own rocket to L2, with an initial fuel mass determined by the difference between the rocket payload mass limit and the starshade dry mass. We adopt the same scaling relationship between starshade diameter and dry mass as given in Ref. 3. We note that the optimum IWA for a starshade varies a bit with the payload mass limit of the rocket in which it is launched. This variation is relatively small and we focus mainly on a launch in a Delta IV Heavy, for which the 60 mas IWA is most valid.

Like Ref. 3, we adopt 2 separate propulsion systems for stationkeeping and slewing, due to the concern of bright ion plumes interfering with exposures. We also ignore the effects of solar radiation pressure on stationkeeping, for which mitigation systems are not yet understood. Whereas Ref. 3 adopted an unrealistically large thrust of 10 N, here we adopt a more realistic, yet optimistic, thrust of 2 N.

5. RESULTS & DISCUSSION

Figure 1 shows the yield of a coronagraph-based mission as a function of aperture size for each of the 4 observation methods described in Section 3. The solid line shows the yield for Method 1, in which all 3 20% bandpass coronagraphs contribute simultaneously to the broadband detection of exoEarth candidates, no color information is obtained, and no spectral characterizations are performed. All spectral characterizations are assumed to happen after completion of the nominal mission and are not included in the mission exposure time budget. This method represents an upper bound for the yield.

The dotted line shows the yield for Method 2, in which the planet must be detected in each bandpass independently at a S/N= 7, providing a 2-color discriminator. With Method 2, we assume that this color information, along with any astrometric, photometric, and debris disk information obtained, can perfectly discriminate between exoEarth candidates and all other point sources. Spectral characterizations are performed to search for water absorption at 1 μm, and are limited to only the detected exoEarth candidates. The time required for spectral characterization counts against the total exposure time budget. The reduced bandpass used for detection, along with the required spectral characterization time, reduces the yield by a factor of $\sim$0.7 compared to Method 1.

The dashed line shows the yield for Method 3, in which we assume that the discriminators described in Method 2 won't work and we commit to performing a spectral characterization *after every observation*. Because we perform a spectral characterization after every observation, color information will not be necessary, so we allow all 3 20% bandpass coronagraphs to contribute to the initial detection. Compared to Method 2, Method 3 reduces the yield by a factor of $\sim$0.65, or a factor of $\sim$ 0.45 compared to Method 1.

Finally, the dot-dashed line shows the yield for Method 4, in which we require at least 6 visits to each system prior to attempting a spectral characterization. Because multiple visits are performed, we allow all 3 20% bandpass coronagraphs to contribute to detections in each of these visits and assume that the 6 visits, when co-added, will provide 2-color information as well. This method would provide all of the discriminators of Method 2, plus orbital constraints that can be used to determine whether a planet is in the HZ, and possibly mass constraints if astrometric measurements can be made simultaneously. Therefore, we assume this method provides perfect differentiation between all planet types and background objects, and we only devote spectral characterization time to the detected exoEarth candidates. The time required for spectral characterization counts against the total exposure time budget.

Table 4. Baseline Starshade Mission Parameters

Parameter	Value	Description
$\Sigma\tau$	< 5 yr[a]	Total exposure time of the mission
X	0.7	Photometric aperture radius in λ/D
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 PSF
T	0.65	End-to-end facility throughput, excluding photometric aperture factor
$\Delta\text{mag}_{\text{floor}}$	27.5	Systematic noise floor (faintest detectable point source)
$\text{CR}_{\text{b,detector}}$	0	Detector noise count rate
		Starshade Parameters
λ	0.75 μm	Central wavelength
$\Delta\lambda$	0.5 μm	Bandwidth
ζ	10^{-10}	Uniform raw contrast level, relative to theoretical Airy pattern peak
IWA	60 mas	Inner working angle
OWA	∞	Outer working angle
D_{ss}	61.7 m[b]	Diameter of starshade
m_{dry}	2742 kg[b]	Dry mass of starshade spacecraft including contingency
		Propulsion Parameters
I_{sk}	300 s	Specific impulse of station keeping propellant (chemical)
I_{slew}	3000 s	Specific impulse of slew propellant (electric)
ϵ_{sk}	0.8	Efficiency of station keeping fuel use
$\mathcal{T}$	2 N	Thrust

[a]Code optimizes total exposure time by balancing fuel and time use. Up to 5 yrs allowed; ∼1 yr typical.

[b]Dependent quantity that varies with telescope aperture; values listed corresponds to $D = 4$ m.

Interestingly, Method 4, which represents our most thorough approach to vetting planets prior to investing spectral characterization time in them, produces one of the highest yields for a coronagraph. This is because coronagraphs are nimble and prefer to observe systems multiple times anyway;[2] the added requirement of at least 6 visits to a system does not actually alter the observation plan significantly. Additionally, because we assumed that the planet detections can be co-added to obtain high enough S/N for color information, the multi-coronagraph detection bandpass provides significantly faster detections than Method 2.

Figure 2 shows the same yields curves for the starshade-based mission. Black lines correspond to a 9800 kg launch mass, approximately that of a Delta IV Heavy rocket to L2. Gray lines correspond to a 3600 kg launch mass, approximately that of a Falcon 9 rocket to L2. We do not show yield curves for an SLS Block 1 launch, as they are nearly identical to that of the Delta IV Heavy; there is little benefit to launching our baseline starshade in rockets more capable than the Delta IV Heavy.

As shown by Ref. 3, the yield turns over at larger telescope apertures. This is because the starshade diameter and separation from the telescope increase with telescope aperture for a fixed IWA. At larger separations, the fuel demands for starshade slewing increase. At the same time, a larger starshade diameter equates to a greater starshade mass, such that it occupies more of the mass budget of the rocket. Thus, starshade slews become more demanding of fuel, and there is less fuel available.

For the starshade, Methods 2 and 3 are lower than Method 1. However, the impact of Methods 2 and 3 are not independent of telescope aperture size, which was roughly the case for coronagraph-based missions. For

the starshade, larger apertures are affected less because they are operating farther into the fuel-limited regime; smaller apertures are operating farther into the time-limited regime.

Unlike the coronagraph, Method 4 is substantially less than Methods 1–3. This is because Method 4 requires multiple visits to each system prior to obtaining a spectrum, and each visit requires slewing the starshade. Because the starshade is largely limited by fuel, this observational approach is very costly.

It is reasonable to conclude that if one desires to use a starshade to constrain the orbit of exoEarth candidates (e.g., to verify that they reside within the HZ), a wiser approach would be to first identify the promising exoEarth candidates via their spectra, then perform multiple visits to just those systems. We refer to this as Method 5. We do not investigate Method 5 in detail in this paper, but we do provide an approximation of this method. If 6 total visits are required for each exoEarth candidate, we can approximate the impact of Method 5 via $N'_{\mathrm{EC}} \approx N_{\mathrm{EC}}(1 - 5N_{\mathrm{EC}}/n_{\mathrm{slews}})$, where N_{EC} is the yield calculated via Methods 2 or 3 and n_{slews} is the number of slews performed for those methods. The thin, dashed-triple-dotted lines in Figure 2 shows this approximation applied to Method 3. Orbital characterization of just the exoEarth candidates reduces the yield of a starshade-based mission, most notably for larger aperture missions that operate well into the fuel-limited regime.

We can also provide an absolute upper limit on the yield of a starshade-based mission that observes via Method 5. We note that for the baseline starshade mission considered here, ~ 150 slews are performed. If six visits are required to each star, then $\sim 150/6$ stars can be visited the minimum number of times. Thus, even if $\eta_{\oplus}$ is unity and the multi-visit completeness of each system is 100%, the yield would be limited to 25 exoEarth candidates. In practice, the completeness will be $< 100\%$, and therefore the yield of a starshade will likely be < 25 exoEarth candidates if orbit determination is required.

Figure 3 shows a direct comparison of the coronagraph and starshade yields for each of the 4 methods, plus Method 5 for the starshade. Coronagraph yields are shown in black and starshade yields are shown in blue. For this figure, we ignore the Falcon 9 launch for the starshade and plot only on the Delta IV Heavy launch. The y-axis is logarithmic to highlight where the curves cross. The left and right panels show optimistic and pessimistic assumptions of the median exozodi level. In the optimistic 3 zodi scenario, the coronagraph and starshade yield curves cross at telescope apertures in the 4–7 m range, depending on the observational method used.

Because the starshade based mission is largely limited by fuel, and indirectly limited by exposure time, it is more robust to photometric sources of astrophysical noise.[3] As a result, for the pessimistic 60 zodi scenario, the starshade yields change much less than the coronagraph yields, resulting in the yield curves crossing at larger telescope apertures between 4.5–9 m.

The validity of these observational methods, and the assumptions behind them, have yet to be demonstrated with high-fidelity planet retrieval models. However, if they are approximately valid, and the assumed coronagraph and starshade performance parameters are realistic, starshades may lead to greater exoEarth candidate yields for missions with $D \lesssim 4$ m, while coronagraphs may have greater exoEarth candidate yields for $D \gtrsim 8$ m. If orbit determination is deemed necessary, coronagraphs may have greater exoEarth yield for $D \gtrsim 5$ m.

6. CONCLUSIONS

We provide an apples-to-apples comparison of exoEarth candidate yield as a function of telescope aperture size for starshade-based missions and corongraph-based missions. Starshade yields are relatively high for small aperture missions, but turn over as aperture size increases due to the increased fuel demands of larger, more distant starshades. Coronagraph yields start off relatively low, but continue to increase at larger apertures. These yields cross somewhere between 4 and 8 m, with starshades performing better at smaller apertures and coronagraphs performing better for larger apertures. If orbital constraints are necessary, coronagraph yields may be higher for apertures as small as 5 m. The observational approach to obtaining color, spectra, and orbital constraints can significantly impact the yield. To maximize the yield of a starshade-based mission, one should attempt to first distinguish between planet types with color, geometry, and photometry and then immediately spectrally characterize only the exoEarth candidates, if possible. Multi-visit orbit determination (for HZ verification) should occur after spectral characerization, if necessary. To maximize the yield of a coronagraph-based mission, one should perform multi-visit orbit determination first, then attempt to distinguish between exoEarth candidates and other planets, and then perform spectral characterizations.

ACKNOWLEDGMENTS

The results reported herein benefitted from collaborations and/or information exchange within NASA's Nexus for Exoplanet System Science (NExSS) research coordination network sponsored by NASA's Science Mission Directorate.

REFERENCES

[1] C. C. Stark, A. Roberge, A. Mandell, and T. D. Robinson, "Maximizing the ExoEarth Candidate Yield from a Future Direct Imaging Mission," *ApJ* **795**, p. 122, Nov. 2014.

[2] C. C. Stark, A. Roberge, A. Mandell, M. Clampin, S. D. Domagal-Goldman, M. W. McElwain, and K. R. Stapelfeldt, "Lower Limits on Aperture Size for an ExoEarth Detecting Coronagraphic Mission," *ApJ* **808**, p. 149, Aug. 2015.

[3] C. C. Stark, S. Shaklan, D. Lisman, E. Cady, D. Savransky, A. Roberge, and A. M. Mandell, "Maximized ExoEarth Candidate Yields for Starshades," *ArXiv e-prints* , May 2016.

[4] R. A. Brown, "Single-Visit Photometric and Obscurational Completeness," *ApJ* **624**, pp. 1010–1024, May 2005.

[5] 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," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **6693**, Sept. 2007.

[6] R. K. Kopparapu, R. Ramirez, J. F. Kasting, V. Eymet, T. D. Robinson, S. Mahadevan, R. C. Terrien, S. Domagal-Goldman, V. Meadows, and R. Deshpande, "Habitable Zones around Main-sequence Stars: New Estimates," *ApJ* **765**, p. 131, Mar. 2013.

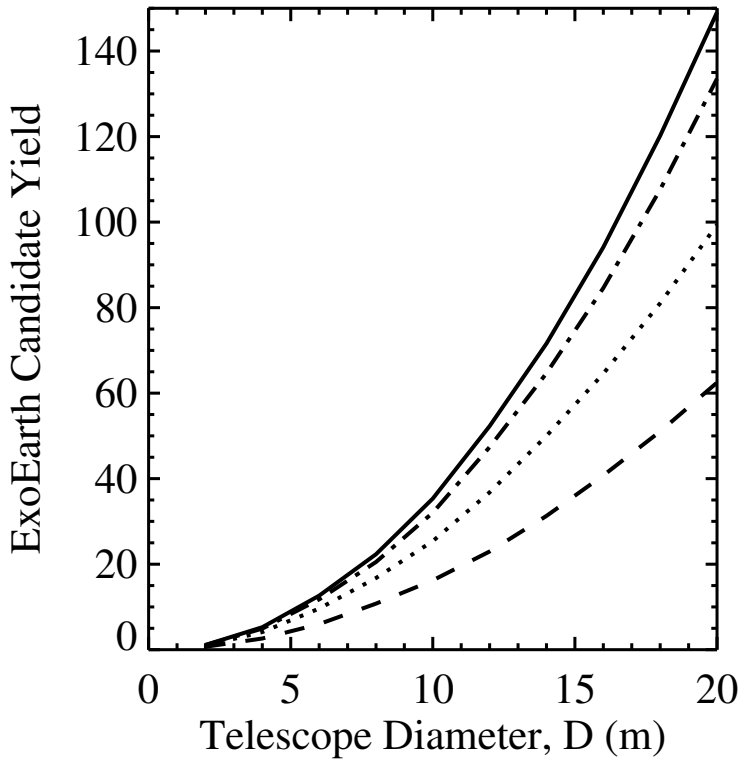

Figure 1 ExoEarth candidate yield for a coronagraph-based mission as a function of telescope aperture size. Observational Methods 1–4 are illustrated with solid, dotted, dashed, and dot-dashed lines, respectively.

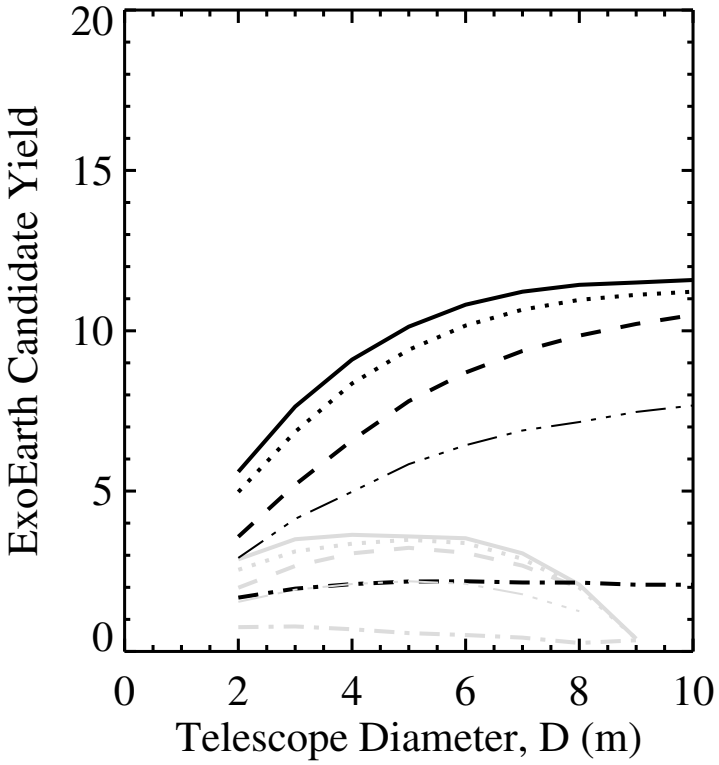

Figure 2 ExoEarth candidate yield for a starshade-based mission as a function of telescope aperture size. Observational Methods 1–4 are illustrated with solid, dotted, dashed, and dot-dashed lines, respectively. Black and gray lines correspond to Delta IV H and Falcon 9 launches, respectively. The thin dashed-triple-dotted line shows an approximate yield for the starshade assuming Method 3 is applied, and then 5 visits per detected exoEarth candidate are reserved for later orbit characterization.

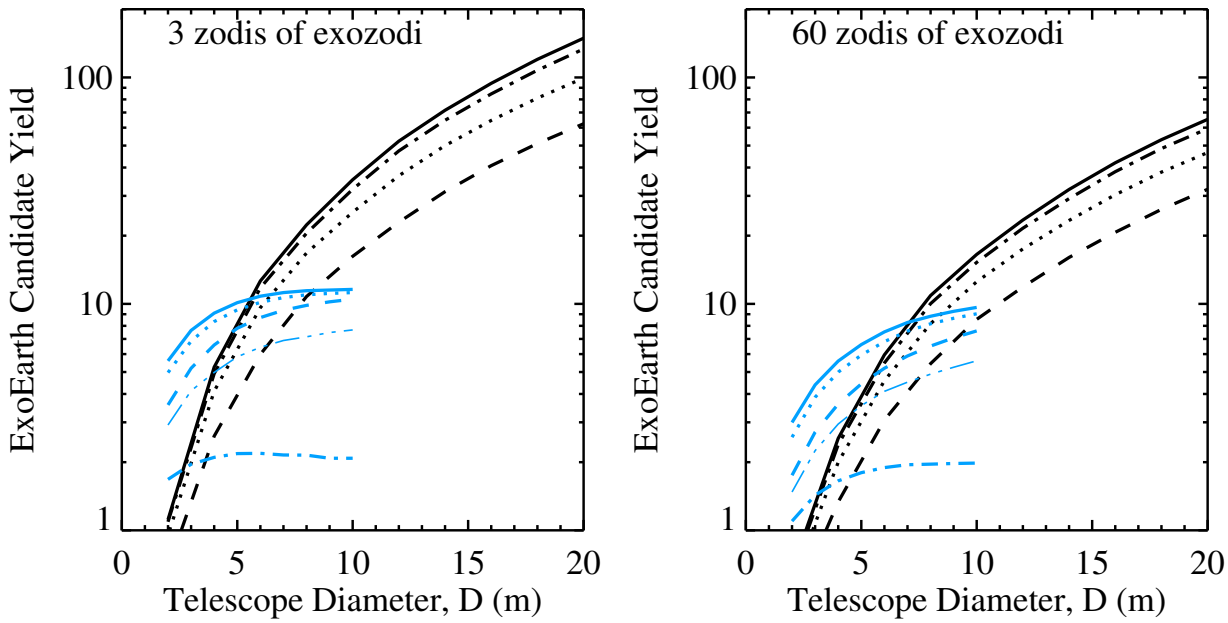

Figure 3 Direct comparison of coronagraph and starshade yields for Methods 1–4. The thin dashed-triple-dotted line presents an approximate yield for the alternate observation Method 5 for a starshade, in which orbit characterization is performed after spectral characterization.