

Current laboratory performance of starlight suppression systems and potential pathways to desired Habitable Worlds Observatory exoplanet science capabilities

Bertrand Mennesson^{a,*}, Ruslan Belikov^b, Emiel Por^c, Eugene Serabyn^a,
Garreth Ruane^a, A. J. Eldorado Riggs^a, Dan Sirbu^b, Laurent Pueyo^c,
Remi Soummer^c, Jeremy Kasdin^d, Stuart Shaklan^a, Byoung-Joon Seo^a,
Christopher Stark^e, Eric Cady^a, Pin Chen^a, Brendan Crill^a, Kevin Fogarty^b,
Alexandra Greenbaum^f, Olivier Guyon^g, Roser Juanola-Parramon^e, Brian Kern^a,
John Krist^a, Bruce Macintosh^h, David Marx^a, Dimitri Mawet^{a,i}, Camilo Mejia Prada^a,
Rhonda Morgan^a, Bijan Nemati^j, Leonid Pogorelyuk^k, Susan Redmond^l,
Sara Seager^l, Nicholas Siegler^a, Karl Stapelfeldt^a, Sarah Steiger^c, John Trauger^a,
James K. Wallace^a, Marie Ygouf^a and Neil Zimmerman^e

^aCalifornia Institute of Technology, Jet Propulsion Laboratory, Pasadena, California, United States

^bNASA Ames Research Center, Moffett Field, California, United States

^cSpace Telescope Science Institute, Baltimore, Maryland, United States

^dPrinceton University, Department of Mechanical and Aerospace Engineering, Princeton, New Jersey, United States

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

^fIPAC, California Institute of Technology, Pasadena, California, United States

^gNational Institutes of Natural Sciences, Subaru Telescope, National Astronomical Observatory of Japan, Hilo, Hawaii, United States

^hUniversity of California Observatories, Santa Cruz, California, United States

ⁱCaltech, Department of Astronomy, Pasadena, California, United States

^jTellus 1 Scientific, LLC, Huntsville, Alabama, United States

^kRensselaer Polytechnic Institute, Troy, New York, United States

^lMassachusetts Institute of Technology, Department of Earth, Atmospheric and Planetary Sciences, Cambridge, Massachusetts, United States

ABSTRACT. We summarize the current best polychromatic (~10% to 20% bandwidth) contrast performance demonstrated in the laboratory by different starlight suppression approaches and systems designed to directly characterize exoplanets around nearby stars. We present results obtained by internal coronagraph and external starshade experimental testbeds using entrance apertures equivalent to off-axis or on-axis telescopes, either monolithic or segmented. For a given angular separation and spectral bandwidth, the performance of each starlight suppression system is characterized by the values of “raw” contrast (before image processing), off-axis (exoplanet) core throughput, and post-calibration contrast (the final 1-sigma detection limit of off-axis point sources, after image processing). Together, the first two parameters set the minimum exposure time required for observations of exoplanets at a given signal-to-noise, i.e., assuming perfect subtraction of background residuals down to the photon noise limit. In practice, residual starlight speckle fluctuations during the exposure will not be perfectly estimated nor subtracted, resulting in a finite post-calibrated contrast and exoplanet detection limit whatever the exposure time. To place the current laboratory results in the perspective of the future Habitable Worlds Observatory (HWO) mission, we simulate visible observations of a fiducial Earth/Sun twin system at 12 pc, assuming a 6 m (inscribed diameter) collecting aperture and a realistic end-to-end optical throughput. The exposure times required for

*Address all correspondence to Bertrand Mennesson, Bertrand.mennesson@jpl.nasa.gov

broadband exo-Earth detection (20% bandwidth around $\lambda = 0.55 \mu\text{m}$) and visible spectroscopic observations ($R = 70$) are then computed assuming various levels of starlight suppression performance, including the values currently demonstrated in the laboratory. Using spectroscopic exposure time as a simple metric, our results point to key starlight suppression system design performance improvements and trades to be conducted in support of HWO's exoplanet science capabilities. These trades may be explored via numerical studies, lab experiments, and high-contrast space-based observations and demonstrations.

© The Authors. Published by SPIE under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JATIS.10.3.035004](https://doi.org/10.1117/1.JATIS.10.3.035004)]

Keywords: exoplanets; coronagraph; starshade; starlight suppression

Paper 24029G received Mar. 28, 2024; revised Aug. 21, 2024; accepted Aug. 23, 2024; published Sep. 20, 2024.

1 Introduction

To enable a broad range of exciting potential missions across the electromagnetic spectrum, the 2021 US Decadal Survey in Astronomy and Astrophysics¹ recently recommended the creation of a Great Observatories mission and technology maturation program. It further recommended that the first mission to be advanced under this program should be a large (~ 6 m) infrared/optical/ultraviolet space telescope that will “search for biosignatures from a robust number of about ~ 25 habitable zone planets and be a transformative facility for general astrophysics.” This prospective mission has since been called the Habitable Worlds Observatory (HWO).

To detect and spectrally characterize small exoplanets orbiting in the habitable zones of nearby Sun-like stars, HWO must include a powerful starlight suppression system able to cancel on-axis light very efficiently—typically by 10 orders of magnitude—while simultaneously passing a significant fraction of the nearby exoplanet light. In addition, various interesting molecules and spectral features dictate that the observations should be conducted over a spectral bandwidth as broad as possible. For example, the Habitable Exoplanet Observatory (HabEx) and large ultraviolet optical infrared surveyor (LUVOIR) pre-decadal large mission concept studies^{2,3} baselined a wavelength range of 0.2 to 1.8 μm for exoplanet direct observations, enabling searches for a variety of potential bio-signatures in rocky planets' atmospheres and placing them into context. Obtaining near-ultraviolet (UV) to near-infrared spectra of all targets would indeed be ideal, especially for exo-Earth characterization (e.g., Refs. 4 and 5). However, observing in the near UV and the infrared presents well-known additional challenges in terms of sensitivity and spatial resolution, respectively. A more limited wavelength range of ~ 0.45 to $\sim 1 \mu\text{m}$ has also been previously considered^{6–8} for the spectral characterization of directly imaged terrestrial—and gas giant—planets in reflected light. This “visible” range would still provide access to strong spectral features for a limited collection of gases (CO_2 , H_2O , CH_4 , O_2 , and H_2), as well as the opportunity to identify atmospheric Rayleigh scattering. The detailed exoplanet science objectives of HWO are yet to be specified in detail. But, even if the narrower 0.45 to 1 μm wavelength range was adopted as a minimum threshold for direct exoplanet observations with HWO, it would already represent a spectral bandwidth of 76% (defined as $\Delta\lambda/\lambda_{\text{center}}$). This would call for a starlight suppression system that can either operate over the whole spectral range at once or observe over a limited set of “reasonably broad” (here taken to mean at least 10%) individual spectral channels within that range.

Over the last 10 to 20 years, many laboratory testbeds have been set up worldwide⁹ to take on the challenge of broadband high-contrast (see Sec. 2.1 for rigorous definitions of the contrast terms used in this paper) starlight suppression to directly image and characterize Earth-like planets. They have used different approaches, operating conditions (e.g., in air versus vacuum), and aperture types (clear versus obscured, monolithic versus segmented).

Starlight can be suppressed between the telescope's primary mirror and the detectors by “internal coronagraphs” or between the star and the telescope by “external coronagraphs.” We will follow the standard practice and refer to internal coronagraphs as simply “coronagraphs” and external coronagraphs as “starshades.”

For coronagraphs, testbeds have been developed in both air and vacuum. The current in-air testbeds have tackled many different—and in most cases complementary—aspects of the overall technological challenge. For instance, the “Tres Haute Dynamique” (French for “very high contrast”) testbed^{10–13} of the Paris Observatory concentrated on ultra-broadband starlight suppression and a specific focal plane wavefront-sensing approach using a self-coherent camera. The NASA Ames Coronagraph Experiment testbed focused on multi-star wavefront control for high-contrast imaging around binary stars,^{14,15} as well as phase-induced apodization techniques to provide high-throughput observations at small separations^{16,17} in collaboration with the Subaru SCExAO team. The high-contrast high-resolution spectroscopy for the segmented telescopes testbed (HCST) located at the California Institute of Technology concentrated on apodization techniques for small inner working angle (IWA) vector vortex coronagraphs (VVC4)^{18,19} and the use of (single-mode) fiber-fed high-resolution spectrographs to distinguish planetary signals from starlight residuals.^{20,21} The High-Contrast Imager for Complex Aperture Telescope (HiCAT testbed), located at the Space Telescope Science Institute, specialized in wavefront sensing and control (WFS&C) as well as amplitude and phase apodization techniques for high-contrast coronagraphy on segmented or heavily obscured apertures.^{22–27}

To our knowledge—and recognizing that the information provided hereafter represents a snapshot in time as of March 2024—vacuum coronagraphic results have only been obtained at three facilities. The most prolific one is NASA’s High Contrast Imaging Testbed (HCIT) located at the Jet Propulsion Laboratory (JPL). It is currently the only vacuum coronagraphic facility that has demonstrated coronagraphic dark holes over spectral bandwidths of 10% or more with contrast levels deeper than 10^{-9} at 3 to $4\lambda/D$ separations. The HCIT is a set of vacuum coronagraphic testbeds that has been made available for coronagraphic experiments conducted by the US community under the NASA Strategic Astrophysics Technology (SAT) program over the last ~15 years. It has been used for testing many different coronagraphic approaches, including some of those mentioned above. The HCIT was also used for key milestone demonstrations in support of the Roman Coronagraph flight instrument development between 2014 and 2020. It currently hosts two decadal survey testbeds (DSTs)^{28,29} specifically designed to push coronagraphic contrast down to the levels required by HWO. The DSTs are the state of the art in vacuum testbeds based on the accumulated experience of the HCIT facility. The other two facilities are located at the Lockheed Martin Advanced Technology Center, where vacuum coronagraphic demonstrations occurred in 2014–2015 in support of the EXCEDE study,³⁰ and at the University of Arizona, where a new vacuum chamber and space coronagraph optical bench (SCoOB) were recently commissioned. (During the referring process, we became aware of 10% and 15% bandwidth vacuum coronagraphic results recently obtained with SCoOB,³¹ demonstrating an average contrast performance in the few 10^{-8} range over a single-sided dark hole extending from 3 to $10\lambda/D$.)

Recent laboratory results with starshades have been obtained in air with the Princeton starshade testbed.^{32,33} As with internal coronagraphs, we present here only broadband results, which were obtained over a spectral bandwidth of 12%.

In this paper, we collect in one place the recent starlight-suppression results most relevant to HWO to assess (1) the status of the field and (2) the work still needed to reach the performance notionally required by HWO.

In Sec. 2, we review the instrumental contrast performance currently demonstrated in the laboratory as of March 2024 (also see the regularly updated compilation of testbed performance demonstrations maintained by Brendan Crill on behalf of the Exoplanet Exploration program located at Ref. 34), covering various telescope types: on- and off-axis monoliths and on- and off-axis segmented apertures. In line with the high-level science objectives of HWO and the anticipated needs of HWO’s starlight suppression system, we only consider here approaches that have demonstrated “deep” (defined as better than 10^{-7}) contrast in the lab over instantaneous bandwidths of ~10% or more and at angular separations $\leq 4\lambda/D$. Polychromatic performance is itself often a good indicator of coronagraph (and testbed) maturity as achromatizing coronagraphic masks and simultaneously controlling both amplitude and phase across large bandwidths with deformable mirrors (DM) remain challenging.

In Sec. 3, we introduce the other two key performance parameters (KPPs) of a starlight suppression system beyond instrument contrast: the off-axis throughput for close-in exoplanet

light detection, and the post-calibration contrast obtained after removing residual starlight from science images. Given the extreme levels of starlight suppression required, laboratory experiments have mainly concentrated so far on demonstrating deep contrast at small angular separations. However, off-axis planet signal transmission and post-calibration capabilities become the crucial performance parameters once a threshold raw contrast performance has been reached.

In Sec. 4, we start exploring the relationship among the three starlight suppression system KPPs and HWO’s exoplanet science capabilities. As an initial simple step in that exploration, we compare the signals expected from exo-Earth targets, (exo)-zodiacal dust structures, and starlight (attenuated by variable factors) as a function of stellar distance and wavelength. We then adopt as a representative fiducial target the case of an Earth twin seen at quadrature around a Sun twin at 12 pc and compute the exposure times required for photometric detection and visible spectroscopy of that target under various KPP levels, including those currently demonstrated in the lab. This provides an initial mapping of the high-level starlight suppression KPPs values—or combinations thereof—required for viable scientific observations and helps identify possible trades among individual KPPs. For instance, it provides some insight into the relative science value of improving raw contrast versus increasing off-axis throughput or improving raw speckle calibration.

The actual system implementation trades, including the detailed telescope design and coronagraph mask selection, and the higher system level trades of, e.g., telescope wavefront stability (or knowledge) versus coronagraph resilience to aberrations are beyond the scope of this paper. At the 10^{-10} level, the contrast performance degradation caused by even tens of picometers of uncorrected wavefront changes will generally be noticeable by coronagraphs, and it will differ among them. The effect will be much less on starshades, which can tolerate nanometers of drifts before seeing any contrast degradation. The contrast sensitivity of different coronagraphs to telescope pointing jitter (and hence, stellar diameter) and aberrations changes caused by vibrations or slow thermal drifts is a crucial characteristic to include in future implementation trades. This aspect is being investigated separately^{35–37} and will be further studied in the context of future HWO technology maturation efforts.

We end in Sec. 5 with a list of suggested key improvements, high-priority technical trades, and immediate laboratory demonstrations needed to develop starlight suppression systems toward the notional exoplanet science needs of future HWO missions.

2 Starlight Suppression Laboratory Experiments and Raw Contrast Results

Once the exoplanet science mission objectives are clearly identified with quantitative figures of merit clearly defined, a close-to-optimum mission operations concept can be derived to maximize science yield.^{38–41} A usual example is the number of exo-Earths spectra that can be obtained over a specified wavelength range at a desired signal-to-noise ratio (SNR) and spectral resolution within some total observing time. Assuming that the telescope diameter and list of nearby Sun-like target stars to draw from are fixed, the exo-Earth science yield is essentially set by three starlight suppression system characteristics (KPPs), which vary as a function of spectral bandwidth and location around the star:⁴²

- “Raw contrast” is defined as the instrumental contrast measured in the dark hole region of the science images before any differential imaging and other post-calibration techniques are applied. The raw contrast is field-dependent and, for a given stellar diameter, is a property solely of the optical system. At any given position in the dark hole, a planet with a planet-to-star flux ratio equal to the local raw contrast value would produce the same number of counts as observed at that location from the partially suppressed star. To minimize the exposure time required for exo-Earth observations—especially spectroscopy—the raw contrast value must be low enough that the irreducible photon noise associated with other astrophysical background sources, such as zodi or exozodi dust, dominates over the stellar shot noise. In practice, that threshold raw contrast level varies with stellar distance, wavelength, and exozodi brightness and ranges from $\sim 10^{-10}$ to $\sim 10^{-8}$ (Sec. 4.1).
- Post-calibration contrast, also field-dependent, is defined as the minimum planet-to-star flux ratio detectable with an SNR of 1 after calibration and processing of the raw science

images (Secs. 3.2 and 4.2). For example, for visible observations of planets with a 10^{-10} flux ratio (same as an Earth twin seen at quadrature around a Sun twin) at an SNR of 20, the post-calibration contrast root mean square (rms) error must be $<5 \times 10^{-12}$.

- Off-axis (core) throughput is defined as the fraction of off-axis (planet) light transmitted through the starlight suppression system into a specified off-axis point spread function (PSF) core area (Sec. 3.1). The off-axis core-throughput two-dimensional map hence captures both the effects of PSF sharpness and finite transmission close to the star. This metric is more informative than the starlight suppression system IWA, which is defined as the angular separation at which the system's off-axis transmission reaches 50% of its maximum off-axis value. Although the IWA is a useful performance parameter, it can also be deceptive when used alone. For example, a system with a larger (= worse) IWA but a higher maximum off-axis transmission may still outperform a system with a smaller (= better) IWA but a lower off-axis transmission.

As stated above, we only consider here approaches that, as of March 2024, have demonstrated better than 10^{-7} contrast in the laboratory over instantaneous bandwidths of $\sim 10\%$ or more and at angular separations $\leq 4\lambda/D$. We acknowledge that this choice of threshold for the inclusion of detailed results is somewhat arbitrary, but it has the advantage of capturing the best results obtained to date on each of the four possible aperture types (segmented versus monolith and on- versus off-axis). All lab demonstrations discussed hereafter have been carried out by investigators selected under the auspices of the NASA SAT program and were conducted at visible wavelengths. Except for the HiCAT results, which were obtained in air, all other coronagraphic performance results reported in this section were obtained in vacuum in the HCIT. The HCIT uses passive mechanical vibration isolation to reduce pointing drifts and jitter (down to $\sim 0.005\lambda/D$ rms at $0.55\mu\text{m}^{28}$) and active correction of residual wavefront aberrations, providing a space-like environment that is turbulence-free and thermally stable at a level of 10 mK per day or less. Although there are still some small dynamical effects, e.g., fast residual pointing jitter at the milliarcsecond level, the vacuum lab performance results reported hereafter can then essentially be seen as “static” results. Finally, for all coronagraphic results presented hereafter, two types of DMs were used: 48×48 actuators DMs from Adaptive Optics Associates (Xinetics AOX, Devens, Massachusetts, United States) or ~ 1000 to ~ 2000 actuator DMs from the Boston Micromachines Corporation (BMC, Cambridge, Massachusetts, United States).

2.1 Starlight Suppression System Contrast Definitions

Before describing the results achieved in the laboratory, it is worth noting that the contrast level of a starlight suppression system is generally expressed in terms of either normalized intensity or raw contrast, which correspond to two different quantities.

For any pixel in the dark hole, the normalized intensity (NI) is defined as the ratio of the number of residual starlight counts measured in that pixel when the star is on-axis, with all starlight suppression optics in the beam path, to the maximum number of counts per pixel measured under the exact same conditions, but with the coronagraphic focal plane mask (FPM) (or the starshade mask) removed to let starlight through. Note that the denominator in this ratio is the same for all pixels, meaning that the NI map is a scaled version of the observed 2D coronagraphic image. NI is the quantity of choice when assessing the degree of starlight suppression. But, it says nothing about the system's ability to transmit the signal of off-axis sources (e.g., exoplanets or extended disk structures). All testbed contrast results presented in this section are expressed in terms of NI, unless specified otherwise.

On the other hand, the raw contrast (RC) computed at any pixel in the coronagraphic dark hole accounts for spatial variations in the off-axis transmission of the FPM. It is defined as the ratio of the number of residual starlight counts measured in a pixel (or photometric aperture PA centered around it) when the star is on-axis with all starlight suppression optics in, to the number of counts measured in that same pixel (or photometric aperture) under the same conditions, but after shifting the input star to that pixel. For a sky position \mathbf{r}_0 , the instrumental raw contrast is hence given as⁴³

$$\text{RC}(\mathbf{r}_0) = \frac{\int_{\text{PA}(\mathbf{r}_0)} \text{PSF}_{\text{coro}}(\mathbf{r}, 0) \cdot d\mathbf{r}}{\int_{\text{PA}(\mathbf{r}_0)} \text{PSF}_{\text{coro}}(\mathbf{r}, \mathbf{r}_0) \cdot d\mathbf{r}}, \quad (1)$$

where $\text{PSF}_{\text{coro}}(\mathbf{r}, \mathbf{r}')$ is the (field-dependent) coronagraphic PSF measured at image plane position vector \mathbf{r} for a source at sky position \mathbf{r}' , with all of the coronagraphic masks in place and DMs set to minimize starlight in the desired dark hole region of the image plane. Physically, at a given on-sky location, the RC value corresponds to the planet flux relative to the star that would produce the same signal through the system as the residual starlight. The RC is then a more relevant quantity than NI when assessing the detectability of exoplanets or extended circumstellar disks.

The RC per pixel is equal to the NI per pixel divided by the FPM (2D) off-axis transmission response. This means that it is always worse than NI and that it becomes infinite on-axis. The two figures of merit are only equal (approximately) for sources far enough off-axis to be unaffected by the FPM, yet close enough to the axis to be unaffected by the edges of any downstream field stop. A comparison of NI and RC values as a function of angular separation is given in Ref. 42 for various coronagraphs. For the starlight-suppression systems considered here, the NI and RC values only differ by a few tens of percent or less for separations greater than $3\lambda/D$. We also note that the RC in a multi-pixel photometric aperture is generally worse than the RC per pixel because the off-axis PSF maximum pixel level is higher than the mean off-axis PSF level in the photometric aperture.

2.2 Coronagraph Demonstrations for Off-Axis Monolithic Telescopes

In terms of fundamental physics limits, coronagraph RC performance is not strongly dependent on the aperture shape, obstructions, or segmentations.⁴⁴ However, all else being equal, a plain circular aperture without any obscuration, as given by an off-axis telescope, is more straightforward to work with than one with obscurations, and the best broadband coronagraphic laboratory results to date have been achieved for this type of aperture. They were obtained at JPL's HCIT, and fairly similar levels of RC performance were reached with a “classical” Lyot coronagraph (CLC), a hybrid Lyot coronagraph (HLC), and a VVC4. For completeness, we also present the broadband results obtained with a phase-induced amplitude apodization (PIAA) coronagraph, which demonstrated worse contrast but accessed smaller angular separations.

2.2.1 Classical Lyot coronagraph

The conceptually simplest coronagraph design is the CLC (Fig. 1) invented by Bernard Lyot for observations of the solar corona nearly a century ago. On-axis starlight cancellation is achieved via the combination of an opaque circular mask located in an intermediate focal plane and an aperture stop located in a downstream pupil plane conjugate. The downstream aperture stop, called a Lyot stop, is used to reduce the amount of starlight diffracted inside the pupil by the sharp edges of the FPM. As the FPM gets smaller, a larger fraction of the aperture needs to be blocked by the Lyot stop to reduce the FPM diffractive effects. This means that, for a CLC, a compromise must be made between the ability to detect sources at small separations and coronagraph throughput. Modern coronagraphs, such as the ones described in in Secs. 2.2.2–2.2.4,




Fig. 1 Schematic optical layout of a CLC, with a hard-edge FPM inserted in an intermediary focal plane and a Lyot stop located in a pupil plane conjugate. For the DST CLC experiment considered here,⁴⁵ the entrance pupil consisted of an unobscured circular aperture, and the FPM size was set to $2.7\lambda/D$ in radius. To control starlight diffraction over the 10% spectral bandwidth, the Lyot stop blocked any light located outside a $0.2D$ to $0.675D$ annulus, resulting in significant transmission losses.

use pupil apodization or phase mask techniques to overcome that limitation and theoretically provide higher sensitivity at smaller separations than a CLC.

However, given its relative simplicity to set up and model, a CLC was tested with an unobscured, circular pupil using the DST. The DST has two 48×48 DMs from Adaptive Optics Associates (Xinetics = AOX) and is equipped with a low-order WFS&C subsystem to sense and correct residual dynamic wavefront disturbances in the vacuum chamber (see Sec. 3). Conversely, to the hybrid Lyot mask coronagraph presented in the next subsection, the FPM is a simple binary mask in amplitude, with no spatially varying phase effects purposely introduced by the mask itself. For this reason, we refer to it as a CLC, although the two DMs are controlled to achieve a complex-valued apodization of the pupil.

For this demonstration, the Lyot stop was quite lossy: it blocked any light coming from outside 67.5% of the original beam radius and from inside the inner 20%. Blocking both the inner and outer regions of the beam was required to achieve broadband wavefront control with the two DMs and deep broadband starlight cancellation (Garreth Ruane, JPL, private communication). Using this setup, an average starlight suppression level of $\sim 3.8 \times 10^{-10}$ (NI) was demonstrated over a 360 deg dark hole region covering 3 to $8 \lambda/D$ in separation, using linearly polarized broadband light (10% bandwidth centered at $0.55 \mu\text{m}$) and a linear polarization analyzer.⁴⁵ The resulting high-contrast image is shown in Fig. 2(a). Further tests and analysis suggest that $\sim 80\%$ of the observed contrast is limited in nearly equal parts by four different effects: DM actuator finite stroke resolution (due to 16 bits electronics at the time), chromatic control residual, focal plane occulter ghost, and residual line-of-sight jitter. The origin of the last 20% contribution ($\sim 5 \times 10^{-11}$) to the observed contrast floor is unknown, implying that other processes must be considered.

A more recent, higher throughput, CLC experiment achieved similar contrast performance over a 20% spectral bandwidth with a similar DST laboratory setup but over a one-sided dark hole [Fig. 2(b)] extending from 5 to $13.5 \lambda/D$, i.e., located at larger separations.⁴⁶ As the spectral bandwidth increases, coronagraphic starlight cancellation requires more degrees of freedom (i.e., DM actuators) to control phase and amplitude or may only be achieved over a smaller dark-hole region. Different dark-hole extents may be considered for initial exoplanet blind search detections (full dark hole) versus follow-up spectroscopic observations (e.g., half dark hole).

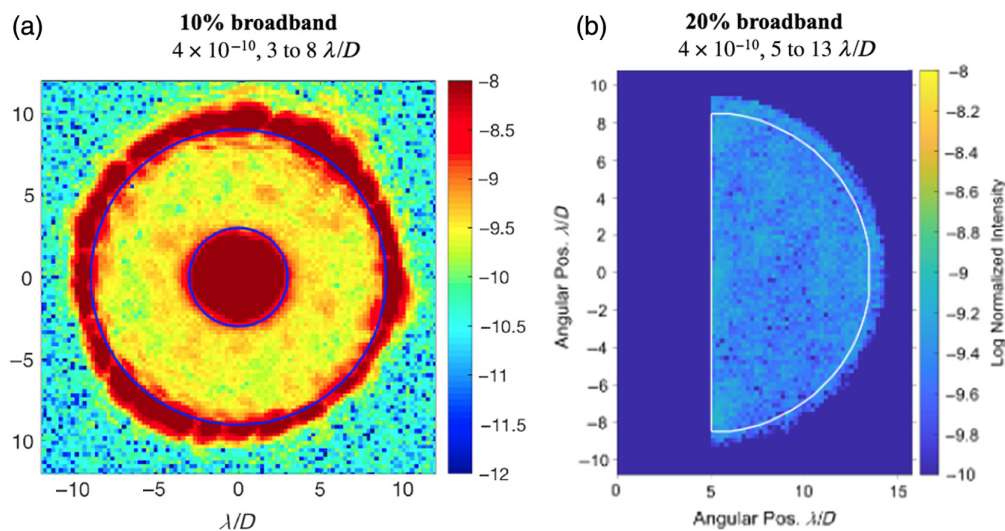


Fig. 2 (a) CLC results obtained in the HCIT DST testbed using linearly polarized broadband light (10% bandwidth centered at $0.55 \mu\text{m}$). The average contrast (NI) measured over the 360 deg dark hole for separations ranging from 3 to $8 \lambda/D$ was 3.82×10^{-10} . (Adapted from Ref. 45). (b) DST CLC results obtained with the same setup but over a 20% bandwidth centered at $0.56 \mu\text{m}$ on a one-sided dark-hole extending from 5 to $13.5 \lambda/D$. The average contrast measured over that one-sided dark hole was 3.97×10^{-10} . (Adapted from Ref. 46.)

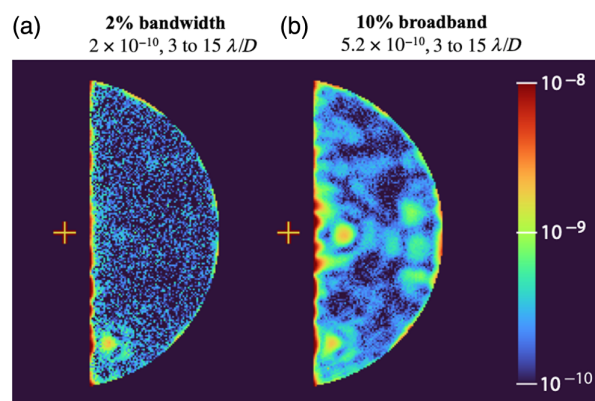


Fig. 3 HLC results obtained in the HCIT testbed using unpolarized light [(a) 2% bandwidth and (b) 10% bandwidth] centered at $0.80 \mu\text{m}$. The average contrast (NI) measured over the one-sided dark hole, which extends from 3 to $15 \lambda/D$, was 2.0×10^{-10} in the 2% bandwidth and 5.2×10^{-10} in the 10% bandwidth. The average contrast measured in the inner 3 to $4 \lambda/D$ region was 3.2×10^{-10} in the 2% bandwidth and 6.0×10^{-10} in the 10% bandwidth.

2.2.2 Hybrid Lyot coronagraph

To improve the coronagraph off-axis throughput at small separations, the FPM spatial transmission profile should be apodized, rather than constant in the CLC case. For an HLC, the optical layout is the same as in Fig. 1. But, the FPM amplitude transmittance, as well as its induced phase shift, is now varying spatially. The FPM amplitude and phase profiles are jointly optimized with the DM solution to minimize contrast over a specified dark hole region and spectral range.

The best HLC laboratory results reported to date (Fig. 3) used a linear occulting mask.^{47,48} The average contrast (NI) measured over the one-sided dark hole, which extends from 3 to $15 \lambda/D$, was 2.0×10^{-10} in the 2% bandwidth and 5.2×10^{-10} in the 10% bandwidth. The FPM complex apodization was provided by thickness-profiled metallic and dielectric films superimposed on a glass substrate. The 10% bandwidth results are fairly similar to those achieved with the CLC but obtained under different operating conditions (one-sided versus full dark hole, dual polarization versus single polarization, and single DM versus two DMs), making an absolute comparison difficult.

Compared with the 10% spectral bandwidth case, contrast results obtained with the same HLC mask were $\sim 2\times$ better over the 2% bandwidth and $\sim 3\times$ worse with the 20% bandwidth. The observed degradation of contrast performance with increased bandwidth, together with a model analysis of the observed speckles morphology and inspection of fabrication hardware, all pointed to a calibration error for the FPM dielectric thickness, an issue that has since been remedied.

2.2.3 Vector vortex charge 4 coronagraph

There are two kinds of known optical vortices: the scalar optical vortex,^{49,50} implemented by a structural (i.e., variable thickness) helix providing a scalar optical phase delay, which applies to the two orthogonal polarization components of natural light identically, and the VVC4,^{51–54} implemented by a rotationally symmetric half-wave plate, providing a “geometrical” phase shift that applies opposite phase screws to the two orthogonal circular polarization states.

In contrast to the Lyot coronagraphs, the vortex FPM is a transparent optic that imparts a spiral phase shift of the form $\exp(il\varphi)$ on the incident field, where l is an even nonzero integer known as the “charge” and φ is the azimuth angle in the focal plane. Light from an on-axis source (i.e., the star) that passes through a circular unobscured entrance pupil of radius a and the transparent vortex FPM is completely diffracted outside the downstream Lyot stop of radius b , with $b < a$, providing a high-throughput solution with perfect starlight suppression in the ideal case. Figure 4 shows the schematic of a vortex coronagraph layout, which is seen to be identical to the CLC/HLC layout except for its phase-based FPM.

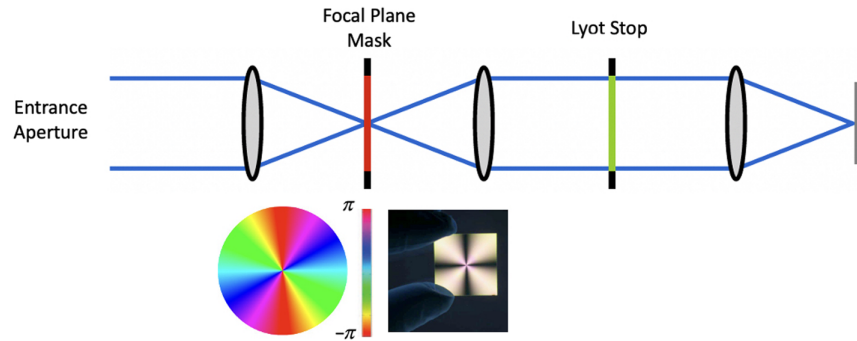


Fig. 4 Schematic layout of a vortex coronagraph, including its focal plane phase mask of complex transmittance $\exp(il\varphi)$ (illustrated with $l = 2$) and circular Lyot stop. For the results described in this section, the entrance pupil consisted of an unobscured circular aperture. A vector vortex of charge $l = 4$ was used, meaning that the focal plane phase rotated twice as fast as represented in the figure, and the Lyot stop opening diameter was set to 80% of the entrance pupil diameter.

In addition, because the focal plane phase mask is transparent instead of opaque, the vortex coronagraph provides high throughput for off-axis point-like sources, i.e., planets located at small angular separations from the star. In practice, as its charge increases, an optical vortex coronagraph becomes less sensitive to low-order aberrations,⁵⁴ but its throughput at small angular separations also decreases. As shown by previous studies,^{2,3} vortex charges of four or six provide a reasonable compromise between the desired resilience to low-order aberrations and high sensitivity to close-in, e.g., habitable zone, exoplanets. A charge 6 VVC was ultimately preferred in the HabEx case as it exhibited far less sensitivity to tip-tilt and astigmatism, the main low-order aberrations induced by polarization cross-talk effects in the HabEx telescope beam train.⁵⁵ It is also worth noting that, although the theoretical performance of vortex coronagraphs does not fundamentally change for segmented primary mirrors, assuming that the segment co-phasing and pupil apodization requirements can be met, their throughput and robustness to wavefront errors are currently found to degrade significantly on centrally obscured (on-axis) telescopes unless hybridized with the PIAA coronagraph.⁵⁶

The deepest rejection has been obtained with an unobscured circular aperture with a vectorial vortex coronagraph of charge four (VVC4) in the DST using the same two 48×48 AOX DMs as in the CLC case. In this case, the vortex geometrical phase ramp was implemented using birefringent liquid crystal polymer waveplates, with their fast axis orientation rotating as a function of azimuth angle. The contrast (NI) achieved by the system on linearly polarized light was 1.6×10^{-9} and 5.9×10^{-9} in 10% (Fig. 5) and 20% optical bandwidths, respectively, averaged over 3 to $10 \lambda/D$ separations on one side of the pseudo-star.⁵⁷ The residual starlight intensity was

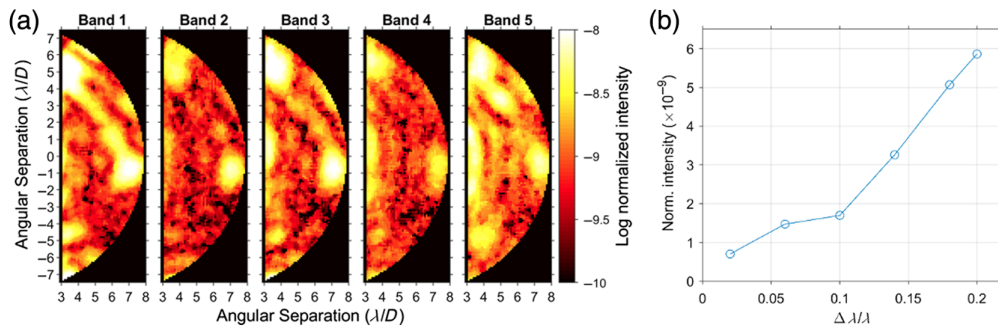


Fig. 5 (a) Raw NI images obtained in five 2% sub-bands with a VVC4 operating on an unobscured circular aperture. All five images are obtained with the same DM settings, determined to minimize the total starlight residuals over the synthetic 10% bandwidth centered around $0.65 \mu\text{m}$. The resulting NI spatial average over the one-sided 3 to $10 \lambda/D$ dark hole is 1.6×10^{-9} . (b) Spatial average of NI measured over the dark hole with the same VVC4 setup but optimizing the DM settings for spectral bandwidths ranging from 2% to 20%. Panels a and b are adapted from Ref. 57.

shown to be coherent (i.e., correlating with changes in DM surface) and likely dominated by spatial imperfections in the vortex FPM. As a result of these imperfections, the NI scales roughly with the square of the spectral bandwidth for bandwidths $>10\%$. Future work aims to achieve 5×10^{-10} contrast in a 10% bandwidth using a new generation of vector vortex masks with tighter tolerances on their imperfections and defects, improved suppression of the polarization leakage, and a new set of higher-resolution DM controller electronics to reduce quantization errors.

In addition to this dual DM VVC4 experiment, tests were also conducted with a single BMC DM, also using polarized light over a 10% bandwidth.⁵⁸ The RC degraded by a factor of 4 to 5 in the inner 3 to $5 \lambda/D$ region and by a factor of ~ 2 over the full one-sided 3 to $10 \lambda/D$ dark hole, with an average contrast measured at 3.2×10^{-9} . Although the VVC4 clear aperture results are worse with a single DM than two, they provide the relevant point of comparison for the VVC4 segmented mask results presented in Sec. 2.3, which were obtained with the same single DM.

2.2.4 Phase-induced amplitude apodization coronagraph

Phase-induced amplitude apodization⁵⁹ (PIAA) is an alternative to classical pupil apodization techniques that use a lossy amplitude pupil mask. In the case of PIAA, an achromatic apodized pupil can be obtained by reflection of an unapodized wavefront on two aspheric mirrors (the PIAA optics) carefully shaped to redistribute the light in the pupil. The surface quality of the PIAA optics and their alignment are critical for optimum performance. Provided those can be achieved or mitigated (see below), the PIAA approach theoretically provides a smaller IWA than the previous coronagraph types and—because the phase-induced apodization is nearly lossless—a high coronagraphic throughput with almost no loss in PSF sharpness close to the optical axis. It is worth noting that, because of the intervening PIAA optics (with only the first PIAA mirror located in a pupil plane), the PSF is no longer translation invariant and becomes less concentrated for off-axis sources. That image distortion becomes significant for separations larger than a few λ/D but can be remedied by adding a set of “inverse PIAA optics” downstream of the coronagraph masks, essential to reform the telescope entrance pupil. Because starlight is already canceled at that point, the inverse PIAA optics surface requirements are not as stringent as for the upstream set of PIAA optics. But, they still add two extra reflections and some complexity to the system (Fig. 6).

The surface quality requirements of PIAA depend on the details of the system. In particular, a classic PIAA system [as opposed to, e.g., the PIAA complex-valued mask coronagraph (PIAACMC) system described in Sec. 2.5] typically has more aggressive PIAA apodization and more stringent wavefront error requirements. Classic PIAA mirrors manufactured in 2010 (“generation 2” by L3 Tinsley), designed for 10^{-9} contrast in broadband, had surface rms errors of 4.65 and 3.8 nm. With a pair of 32×32 DMs, the contrast limit due to surface errors predicted




Fig. 6 Schematic layout of a PIAA coronagraph, including the two PIAA optics (aspheric mirrors) generating the beam apodization, and the inverse set of PIAA optics used to reform the entrance pupil downstream of the coronagraph to correct for off-axis PSF distortion. For the HCIT experiment, the incoming beam into the PIAA optics was diverging, and the FPM limited the dark hole extent to 2 to $4 \lambda/D$ (transmissive region of inserted SEM image) to concentrate on the performance achieved at small separations. No inverse PIAA optics were used or necessary to correct distortions that close to the optical axis.

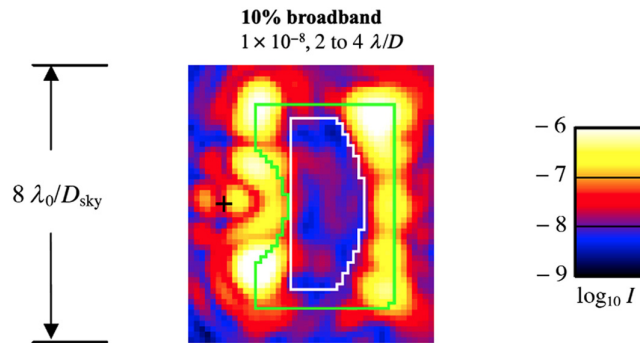


Fig. 7 PIAA starlight suppression results obtained in the HCIT over a 10% spectral bandwidth centered around $\lambda_0 = 0.8 \mu\text{m}$. The contrast image extends from -1 to $+6\lambda_0/D$ in x , $\pm 4\lambda_0/D$ in y . The source center (“star” position) is marked with a small black cross, and the “scored” region border with an inner edge at $x = 2\lambda_0/D$ and outer radius $4\lambda_0/D$ is shown in white. The green line marks the edge of the occulter, the boundary between being fully opaque and fully transmitting. The mean contrast measured over the dark hole region is 10^{-8} , to be compared with 5×10^{-10} in monochromatic light. (Adapted from Ref. 62.)

by models was 4×10^{-10} in a 10% wide band around 800 nm.⁶⁰ On the other hand, for hybrid PIAA designs such as PIAACMC or PIAA-Vortex,⁵⁶ the PIAA apodization is milder, leading to more relaxed requirements on surface quality. For example, lower-cost PIAA mirrors manufactured for a PIAACMC project had an empirically measured wavefront rms error of 65 nm, yet models predicted wavefront-error-limited contrast of 3×10^{-10} (which also included errors from the CMC mask), using a single 952-actuator DM (~ 26 actuators per pupil diameter) in a 10% band around 650 nm.⁶¹ DMs with more actuators (available today) and a second DM are expected to improve this theoretical contrast limit if needed in the HWO context, even without the need for improvements in surface quality (which is also possible).

A classic PIAA coronagraph was tested in vacuum in the HCIT in 2013 on a circular monolithic aperture using a single DM for wavefront control.^{62,63} The testbed reached 10^{-8} mean broadband (10% bandwidth) contrast averaged over a one-sided dark hole extending from 2 to $4\lambda/D$ (Fig. 7) with a linear polarizer inserted immediately before the science camera. Interestingly, the monochromatic contrast demonstrated at $0.8 \mu\text{m}$ with the same setup was 5×10^{-10} , and the significant degradation in performance with spectral bandwidth is believed to be due to surface errors on the first-generation PIAA mirrors. As mentioned above, this was addressed by a second-generation set with improved surface quality, as well as the emergence of hybrid PIAA concepts with more relaxed wavefront error requirements (see Sec. 2.5).

2.3 Coronagraph Demonstrations for Off-Axis Segmented Telescopes

After off-axis monoliths, the next most favorable aperture for coronagraphy, from the standpoint of current coronagraph designs and testbed results, is a segmented off-axis aperture. As in the monolithic off-axis case, this entrance pupil has no central obscuration, but it is discontinuous and slightly obscured by segment gaps. The best broadband vacuum results to date with this aperture type have been obtained with a VVC4 using a segmented mask to simulate the amplitude effect of segment gaps on the entrance aperture (Sec. 2.3.1). No HLC tests have been conducted to date in vacuum on segmented apertures. The only segmented aperture Lyot results have instead been obtained in air, with a phase-apodized-pupil Lyot coronagraph (PAPLC) (see Sec. 2.3.2 for definitions). The PAPLC results used a segmented DM as an actual discontinuous wavefront, allowing for the simulation of the effects of both amplitude and phase discontinuities between telescope segments. The VVC4 and PAPLC results are described in Secs. 2.3.1–2.3.2.

2.3.1 Vector vortex charge 4 coronagraph

The above-described VVC4 (Sec. 2.2.3) was also tested in the HCIT DST⁵⁸ with a hexagonally segmented pupil mask located at the entrance of the coronagraph [Fig. 8(a)], representing a—perfectly phased—off-axis segmented entrance pupil with five segments across. The width of

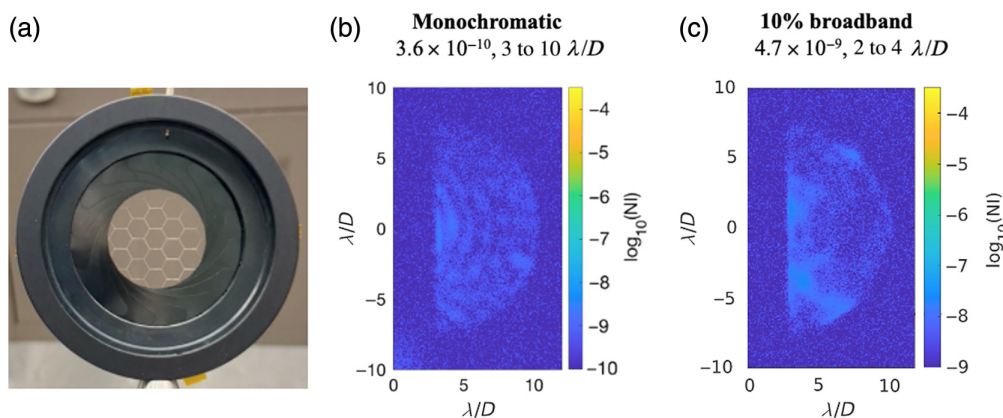


Fig. 8 (a) Entrance aperture used in the segmented VVC4 HCIT/DST lab experiments. (b) Raw NI images obtained using a $0.637 \mu\text{m}$ laser diode. The NI spatial average measured over the one-sided 3 to $10 \lambda/D$ dark hole is 3.6×10^{-10} . (c) Raw NI images obtained over a 10% spectral bandwidth centered at $0.650 \mu\text{m}$. The normalized-intensity spatial average measured over the one-sided 3 to $10 \lambda/D$ dark hole is 4.7×10^{-9} . (Adapted from Ref. 58.)

the segment “amplitude gaps” was $\sim 60 \mu\text{m}$ or 0.2% of the diameter. This is larger than the notional 0.1% relative gap size currently considered for HWO but is in the family. In this case, a single BMC DM with ~ 2000 actuators was used for wavefront correction, and the Lyot stop inner opening diameter was increased to $0.87D$.

The effect of segmentation was initially tested at a single wavelength using a $0.637 \mu\text{m}$ diode laser. Interestingly, the monochromatic contrast performance was found to be the same with or without segmentation of the entrance aperture, which is the reason that we mention it here. The monochromatic contrast (NI) demonstrated with this setup was 3.6×10^{-10} averaged over a 3 to $10 \lambda/D$ one-sided dark hole. Starlight residuals were dominated by slight Airy rings centered on the star [Fig. 8(b)] that did not react to DM actuation. They were attributed to polarization leakage caused by imperfect retardance of the vector vortex mask or the finite extinction ratio of the circular polarizer-analyzer pair used to isolate a single circular polarization through the system.

Using the segmented pupil in broadband light (10% bandwidth centered at $0.650 \mu\text{m}$), the spatially averaged NI over the dark hole increased ~ 10 -fold, to 4.7×10^{-9} . The observed degradation with bandwidth is thought to be primarily caused by chromatic errors in the vector vortex mask itself.⁵⁸ However, the broadband performance observed in the segmented case is also 1.5×10^{-9} worse than the 3.2×10^{-9} achieved with the exact same single-DM broadband VVC4 setup in the monolithic pupil case. So, although the demonstrated VVC4 contrast performance always degrades as the spectral bandwidth increases, that degradation appears slightly more pronounced in the segmented case than in the monolithic one. More tests will be required to disentangle the impact of mask defects from the intrinsic effects of segmentation on broadband VVC performance.

For this HCIT VVC4 demonstration, the diffractive effects of segment gaps were suppressed using phase apodization provided by the DM. An alternative approach used an amplitude-apodized vector vortex (AVC), tested in air in Caltech’s HCST. In this case, an efficient apodizer was manufactured using carbon nanotubes acting as microdots to encode a grayscale pattern. The in-air AVC setup demonstrated a spatial mean contrast of 8.4×10^{-9} over a one-sided dark hole extending from 5 to $10 \lambda/D$ using polarized light over a 10% spectral bandwidth.¹⁸ There are ongoing developments to repeat such AVC demonstrations in vacuum and extend them to smaller separations consistent with HWO needs.

2.3.2 Phase-apodized-pupil Lyot coronagraph

The PAPLC inherently produces a one-sided dark hole.⁶⁴ Similar to other one-sided dark holes presented above, it is better suited to exoplanet spectral characterization than initial blind

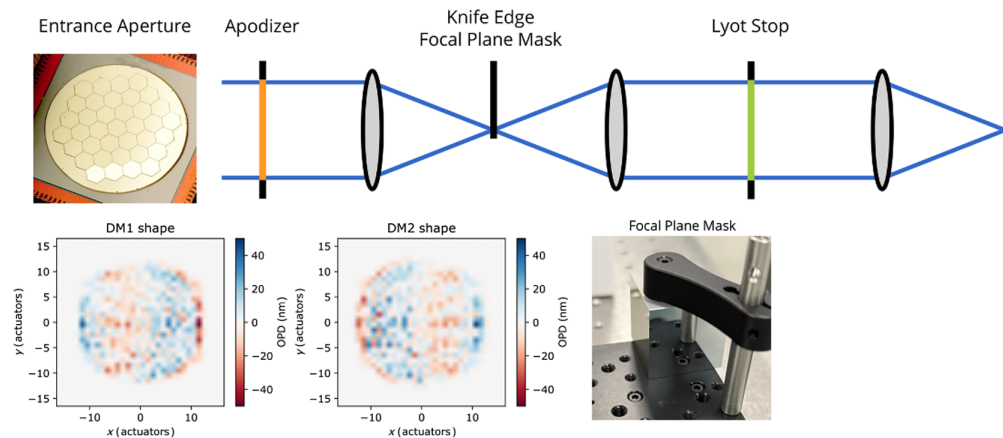


Fig. 9 Conceptual layout of the PAPLC design tested in air on the HiCAT testbed. The entrance aperture, produced by the IrisAO DM, contains 37 individually movable hexagonal segments. Two Boston Micromachines Corporation Kilo-DMs are used to create the phase-induced apodization (DM maps shown) and to perform high-order wavefront control. Note that, instead of reflection off the phase knife, this figure schematically shows the beam propagating to the Lyot stop in transmission.

searches and broadband detection of exoplanets. It uses pupil-plane phase apodization to apodize the PSF falling on an (off-set) knife edge FPM to create a one-sided dark zone with IWAs theoretically as small as $1.5 \lambda/D$. The apodizer can be implemented using a phase-only mask or using a pair of DMs. The PSF is not aligned with the edges of the FPM but is offset by 1 to $2 \lambda/D$. With the phase-only apodizer, this offset is applied using the apodizer phase map itself, producing a PSF with a position that is invariant with respect to the knife edge and creating an inherently achromatic coronagraph. With other apodizers, the offset needs to be applied by moving the PSF on the FPM. Naïve designs already provide robustness against aberrations along the knife edge direction. With more advanced designs of the apodizer, this can theoretically be extended to all low-order aberrations up to the specified order.

The PAPLC laboratory tests were conducted in air at STScI’s HiCAT testbed, where full 360-deg broad-band (9% bandwidth) dark holes were also demonstrated with a “regular” (amplitude-based) apodized pupil Lyot coronagraph at separations $>4.4 \lambda D$.²⁷ The HiCAT light source illuminates the pupil mask (Lenox Laser, Glen Arm, Maryland, United States), the IrisAO segmented DM, the reflective apodizer (a fold mirror for the PAPLC), and two continuous Boston Micromachine Corporation Kilo-DM DMs, each with 952 actuators. These DMs are used for both the apodizer of the PAPLC and for high-order wavefront control. The light is then focused onto the edge of a right-angle prism (Thorlabs, Newton, New Jersey, United States) forming the FPM, of which the reflected part is collimated onto a circular Lyot stop (Lenox Laser). Finally, the light is focused on the science camera. HiCAT has not seen polarization differences at the contrast levels accessible in air so far and operates in unpolarized light. A 37-segment unobstructed hexagonally segmented input aperture was used for these tests (Fig. 9).

The mean contrast observed with the PAPLC setup in 9% broadband light (centered at $0.66 \mu\text{m}$) was 4.2×10^{-8} from 2 to $13 \lambda/D$, and for 25% broadband light, was 9.5×10^{-8} from 2 to $12 \lambda/D$ (Fig. 10 and Ref. 27 and Por et al. in prep). The current performance appears to be limited by the chromaticity effects of the apodizer, which should be correctable according to wavefront control simulations.

2.4 Coronagraph Demonstrations for On-Axis Monolithic Telescopes

Broadband vacuum coronagraphic demonstrations for on-axis monolithic apertures have been obtained in the case of the Roman $D = 2.4$ m telescope, which features a heavily obscured aperture with a 30% central obscuration and six secondary struts that are $0.03 D$ wide. Roman’s complex entrance pupil provides a challenge for coronagraph architects and is likely not a representative of the HWO aperture. However, we describe it here to illustrate the trade between coronagraphic RC and off-axis core throughput, two of the main starlight suppression parameters

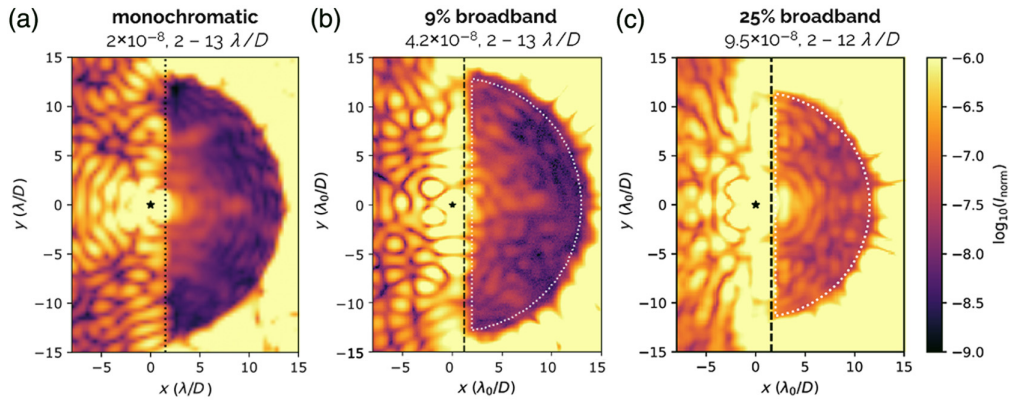


Fig. 10 PAPLC in-air laboratory results. The vertical dashed line indicates the knife edge location at $2 \lambda_0/D$. (a) Monochromatic one-sided dark hole obtained with a laser diode at $0.64 \mu\text{m}$. (b) 9% bandwidth one-sided dark hole obtained around $\lambda_0 = 0.66 \mu\text{m}$. (c) 25% bandwidth one-sided dark hole obtained around $\lambda_0 = 0.66 \mu\text{m}$. The mean contrasts observed are given in the text. (Adapted from Ref. 27 and Por et al. in prep.)

defined above. HLC and shaped pupil coronagraph (SPC) masks were developed and optimized for operation with Roman at different visible wavelengths and over various dark hole extents.⁶⁵ Both types of masks were tested in the HCIT on a Roman-like pupil using a pair of AOX 48×48 DMs for wavefront control and demonstrated deep broadband contrast through early laboratory testing on representative Roman technology testbeds,^{66,67} as summarized hereafter. However, high-contrast performance and reasonable resilience to wavefront changes or misalignments were only achieved at the expense of reducing the off-axis core (planet) throughput.⁶⁸

We note for completion that a back-up PIAA-based architecture (called PIAACMC, see Sec. 2.5) was also developed for the Roman aperture with the goal of providing high contrast down to separations of only $1.3 \lambda/D$. However, due to limited resources and schedule and the more complex nature of PIAACMC, it did not undergo sufficient lab testing to determine its ultimate contrast performance,⁶⁹ and only the HLC and SPC modes will fly with Roman.

2.4.1 Roman HLC coronagraph

The layout of the Roman HLC coronagraph is shown in Fig. 11. As in the off-axis monolith HLC case (Sec. 2.2.2), the FPM amplitude and phase profiles are jointly optimized with the DM solution to minimize contrast over a specified dark hole region and spectral range. To reach deep, broadband (10% to 15% bandwidth depending on coronagraphic mode) suppression of the starlight diffracted from the heavily obscured Roman telescope pupil, significant wavefront distortion is injected by the DMs, which degrades the off-axis PSF by scattering light from the core into extended wings. In addition, a significantly undersized Lyot stop is used to control broadband diffraction and sensitivity to wavefront drifts.⁶⁸ Both effects contribute to a low core throughput, the price to pay to reach and maintain high contrast with the Roman entrance aperture.




Fig. 11 Optical layout of the Roman HLC coronagraph, showing the Roman telescope entrance pupil, a microscope image of the HLC FPM, and the HLC Lyot stop.

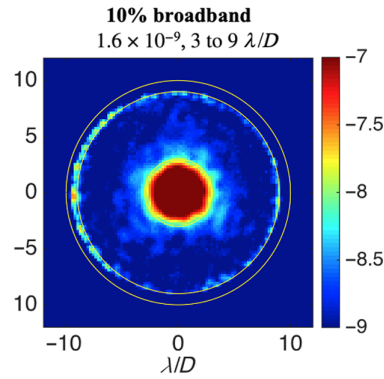


Fig. 12 Best contrast measurements obtained in the lab with the Roman HLC coronagraph mask over 10% bandwidth (2019 technology milestone results⁶⁶).

The Roman HLC coronagraph demonstrated 1.6×10^{-9} contrast averaged from 3 to $9 \lambda/D$ over a 360 deg dark hole using unpolarized light over a 10% bandwidth centered at $0.55 \mu\text{m}$ ⁶⁶ (Fig. 12).

2.4.2 Roman SPC coronagraph

The SPC principle^{70,71} is to use optimum binary amplitude apodization of the entrance pupil to minimize diffraction effects in the subsequent focal plane. In the Roman case, the binary-shaped pupil spatial distribution, FPM (including a field stop), and Lyot stop designs were jointly optimized numerically.⁶⁵ As in the HLC case above, many performance metrics were considered; these include spectral bandwidth, field of view, contrast, core throughput, encircled energy, DM surface height, and low-order aberration sensitivity. The layout of the “SPC-spec” coronagraph used for slit-spectroscopy of point sources with Roman is shown in Fig. 13. There are three such combinations of pupil apodizer, FPM, and Lyot stop, rotated by 60 deg from each other to cover a full 360 deg dark hole.

The Roman SPC coronagraph demonstrated 4.1×10^{-9} contrast averaged from 3 to $9 \lambda/D$ over a dark hole covering $2 \times 65^\circ$ deg in azimuth, using unpolarized light over a 10% bandwidth centered at $0.55 \mu\text{m}$ ⁶⁷ (Fig. 14). The same setup reached 1.1×10^{-8} contrast over 18% bandwidth. The final filter bandwidth for spectroscopy with the flight coronagraph instrument was set at 15%.

Because the SPC relies on amplitude apodization on the pupil, it is intrinsically less transmissive than coronagraphs using phase-only amplitude apodization or FPMs. Also, additional padding of the SPC mask along the spiders was necessary due to uncertainties in pupil alignment, distortion, and magnification. Both effects contributed to a low core throughput for the Roman SPC coronagraphs, similar to what was observed for the Roman HLC.




Fig. 13 Optical layout of the Roman SPC dedicated to spectroscopic observations (SPC-spec), showing the Roman telescope entrance pupil, the binary-shaped pupil used to apodize the pupil, the opaque (amplitude only) bow-tie shaped FPM, and the HLC Lyot stop.




Fig. 14 Best contrast measurements obtained in the lab with the Roman SPC-spec coronagraph over 10% bandwidth (2019 technology milestone results⁶⁷).

2.5 Coronagraph Demonstrations for On-Axis Segmented Telescopes

2.5.1 Phase-induced amplitude apodization complex mask coronagraph

Although the ultimate theoretical performance of coronagraphs is not strongly dependent on aperture type,⁴⁴ for on-axis telescopes, the secondary obscuration degrades the observed contrast performance of current coronagraphs. One way to mitigate this is with a variant of the classical PIAA coronagraph (Sec. 2.2.4) called the PIAA complex-valued mask coronagraph (PIAACMC). In theory, a PIAACMC setup can extend PIAA’s intrinsic advantages in terms of IWA and throughput to any aperture shape,^{16,72} including on-axis segmented apertures. In addition to the two PIAA aspheric mirrors (PIAA optics), PIAACMC also uses a reflective phase shifting FPM, i.e., a special “complex valued mask” in place of the simple black spot occulter of the classical PIAA architecture. In practice, such an “ideal” mask can be implemented by a fully reflective holographic mask, which emulates the far-field response of the ideal CMC mask. In theory, the PIAACMC coronagraph with an ideal CMC mask and no manufacturing errors is capable of suppressing light from an on-axis point source for any arbitrary aperture in broadband light. Models have also shown that, with the empirically measured errors on PIAA mirrors and CMC mask, 3×10^{-10} contrasts are theoretically achievable in a 10% band with a 952-actuator DM (26 actuators per pupil diameter).⁶¹ The tolerance on zone heights for a CMC mask to achieve 10^{-10} contrast is $\sim 10\%$ height error rms.

The PIAACMC laboratory tests were conducted under vacuum at the HCIT facility, producing the best contrast results so far for a centrally obstructed segmented aperture.^{73,74} Key components of the layout are as follows (Fig. 15). The input light source is relayed to a pupil mask representing the LUVOIR A pupil (manufactured by Lambda Consulting using carbon nanotube technology) and then relayed to a PIAA system (by NuTek, Brea, California, United States), consisting of a tube with two PIAA mirrors with a hole in the middle corresponding to the pupil obstruction. A Boston Micromachines Corporation Kilo-DM with 952 actuators is conjugated to the downstream PIAA mirror (PIAA2) and performs wavefront control. Finally, a CMC mask (by NASA JPL Micro Devices Lab, Los Angeles, California, United States), Lyot Stop (by Lambda Consulting, Harwood, Maryland, United States), and field stop complete the coronagraphic setup. A polarizer (analyzer) was also placed in front of the camera to isolate polarization effects from other limiting factors. The testbed design was optimized for a one-sided (180 deg) dark zone using one DM. For the sake of simplicity and as the lab experiment aimed at demonstrating starlight suppression rather than off-axis point source detection, no “inverse PIAA optics” were used downstream of the Lyot stop to correct for the image distortion created off-axis




Fig. 15 Layout of the PIAACMC design tested under vacuum in the HCIT. The entrance aperture, seen in reflection, emulates the on-axis segmented telescope proposed as part of the LUVOIR A mission study concept. A Boston Micromachines Corporation Kilo-DM (not shown) is conjugated to the second PIAA optic and performs wavefront control.

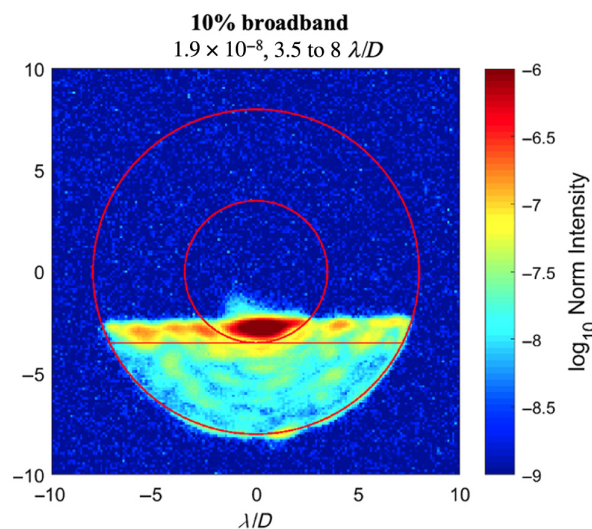


Fig. 16 PIAACMC laboratory results: 10% bandwidth one-sided dark hole obtained around $0.65 \mu\text{m}$. The mean contrast observed between 3.5 and $8 \lambda/D$ is 1.9×10^{-8} . (Adapted from Ref. 74.)

by the PIAA optics. Such optics would need to be added for a PIAA-based coronagraph on HWO, adding a bit of complexity and extra losses.

The mean contrast performance obtained by the HCIT PIAACMC setup between 3.5 and $8 \lambda/D$ in a 10% band was 1.9×10^{-8} (Fig. 16), which is comparable (within roughly a factor of 2) to the monochromatic and 2% cases.⁷³ The 10% bandwidth case appeared to be limited by coherent chromatic effects.

2.6 Starshade Demonstrations

In contrast to an (internal) coronagraph, where starlight suppression occurs by means of an optical system located after the telescope, a starshade, located tens of thousands of kilometers in front of the telescope along the line of sight to the target star, blocks starlight before it enters the telescope^{2,8,75,76} (Fig. 17). As a result, the contrast performance is independent of the type of telescope aperture used: on-axis or off-axis and monolithic or segmented.

The starshade acts as a finite diffraction apodizer, and its shape is numerically optimized to minimize the amount of starlight diffracted across the telescope aperture over some user-specified wavelength range, potentially very broad, with designs theoretically providing better than 10^{-10} contrast over a bandwidth of 100% or more. A key starshade design parameter is its Fresnel number (F), defined at a given observing wavelength λ as $F = D_s^2 / 4\lambda z$, where D_s is the starshade diameter and z is its distance from the telescope. Sources located at angular separations

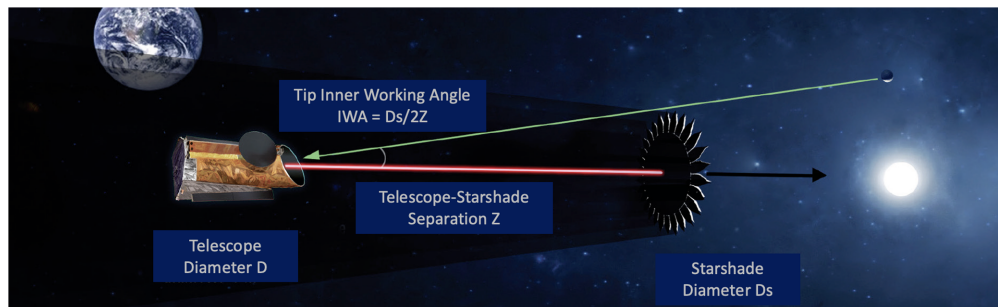


Fig. 17 Schematic exoplanet observations using an external starshade flying tens of thousands of kilometers in front of the telescope, along the line of sight to the target star. (Adapted from Refs. 2 and 8.)

greater than $D_s/2z$ are unaffected by the starshade, meaning that the starshade IWA is of order $D_s/2z$ (although sources inside that angle are still partially transmitted). Adopting this definition of the starshade IWA, often referred to as the “tip IWA,” the starshade diameter required to reach a fixed IWA at a given wavelength is then $D_s = 2\lambda.F/\text{IWA}$. This means that the starshade diameter providing access to sources at a fixed physical angular separation is, to the first order, independent of telescope diameter (D). Conversely, if the IWA is constrained to scale as $1/D$, e.g., if the desired IWA is $2\lambda/D$, the starshade diameter is simply $F * D$. For the flight applications considered here, the Fresnel numbers of order 10 are considered a reasonable compromise between starshade manufacturability (e.g., the contrast sensitivity to shape errors, which increases at small F numbers) and starshade size.

A miniature starshade mask (25.06 mm in diameter) was tested at the Princeton Starshade Testbed (S5 Milestone M1b³²), which is 80 m long and capable of testing 1/1000th scale starshades at a flight-like Fresnel number. As shown in the schematic of Fig. 18, the beam from a spatially filtered laser propagates 27.5 m before it is diffracted by the starshade mask, after which it propagates another 50 m to a camera with an aperture diameter of 5 mm, sitting in the

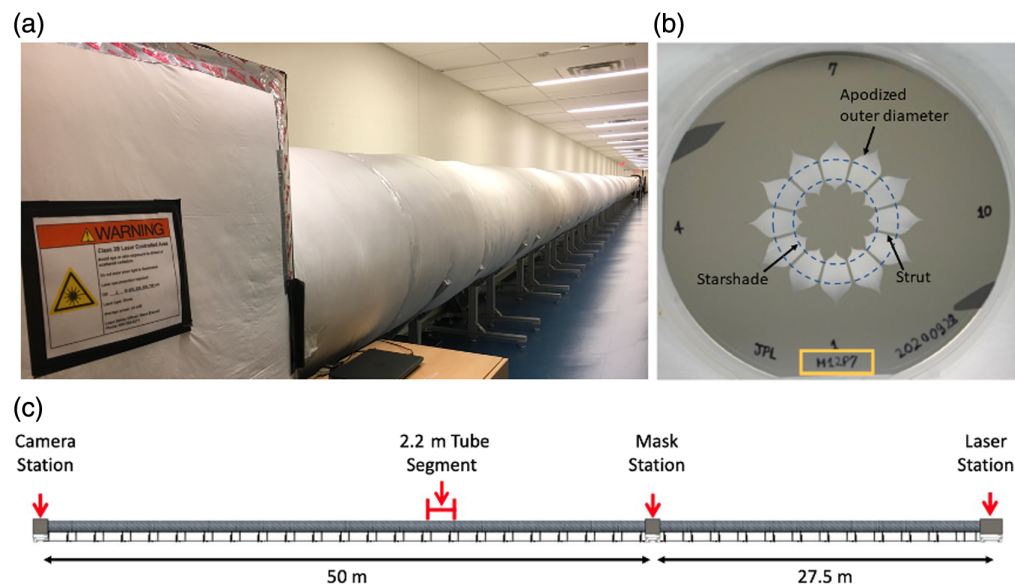


Fig. 18 (a) Standing at the camera station and looking down the length of the testbed toward the starshade mask and laser stations inside the Princeton Frick building. The beam propagates in air. The white insulation minimizes temperature variations and therefore turbulence in the tube. (b) Starshade pattern etched into a silicon-on-insulator wafer, manufactured at the Microdevices Lab at JPL. Interior to the inner blue circle is the inner starshade representing a free-floating occulter. The inner starshade is supported in the wafer by radial struts. The outer blue circle marks the start of the outer apodization function. (c) The layout of the testbed showing distances between the laser, starshade, and camera. (Adapted from Ref. 32.)

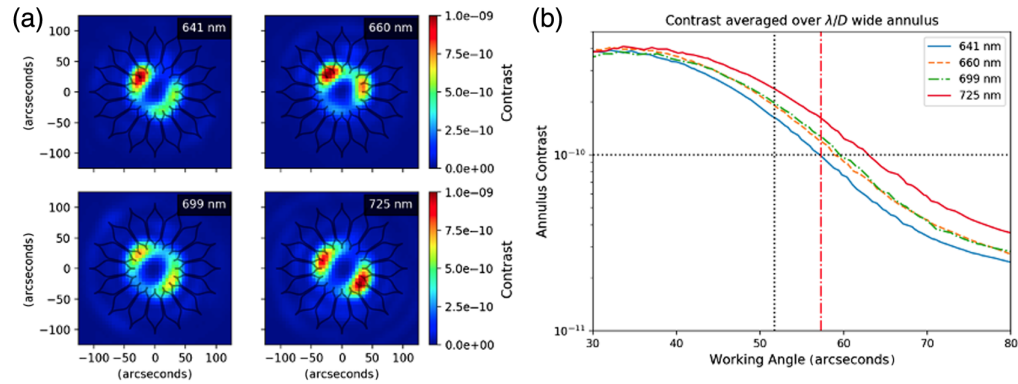


Fig. 19 (a) 2D contrast (NI) maps measured on the Princeton testbed with the same starshade mask at four different wavelengths spanning 12% in bandwidth. The localized bright $\sim 10^{-9}$ lobes observed close to the starshade tip IWA are thick screen effects caused by interactions between the polarized source and the edges of the small mask. (b) Azimuthal mean contrast versus separation. The vertical black dotted line indicates the starshade mask IWA of $51.7''$, corresponding to $1.7 \lambda/D$ at $0.725 \mu\text{m}$ for the 5 mm diameter aperture used. The separation where 10^{-10} contrast is reached corresponds to $\sim 2.1 \lambda/D$ for all wavelengths considered. (Adapted from Ref. 32.)

starshade's shadow. The beam line is enclosed in a 1 m diameter air-filled tube to seal the testbed from stray light and dust and to help stabilize the air inside. To emulate a broadband source, contrast measurements were obtained with a strongly (96%) linearly polarized laser source successively operating at four different wavelengths ranging from 0.641 to $0.725 \mu\text{m}$, spanning an overall spectral bandwidth of 12%. The experiment Fresnel number ranged from ~ 12.2 at $0.725 \mu\text{m}$ to ~ 13.8 at $0.641 \mu\text{m}$, commensurate with flight values. Given that the laser source is not at infinity, the effective starshade-telescope separation to be used for Fresnel number calculations³² is 17.72 m. As illustrated in Fig. 19(b), the experiment demonstrated a broadband azimuthal mean contrast of $\sim 2 \times 10^{-10}$ at a tip IWA of $1.7 \lambda/D$ at $0.725 \mu\text{m}$ ($1.9 \lambda/D$ at $0.641 \mu\text{m}$) and reached 10^{-10} contrast at $\sim 2.1 \lambda/D$ for all wavelengths. The contrast performance further improved at larger angles and eventually was limited by Rayleigh scattering by air molecules.^{32,33} Detailed modeling of non-scalar diffraction effects³³ indicates that the contrast performance and brighter lobes observed close to the IWA [Fig. 19(a)] were limited by non-scalar diffraction (thick screen) effects, where the polarized light interacts with the edges of the mask. Non-scalar diffraction predicts that such effects will be completely negligible (>1000 lower) on a much larger (>10 m) flight-size starshade. In the case of a much larger HWO-compatible starshade, which would have to be tens of meters in diameter, the main contrast limitations are instead expected to come from formation flying, deployment, petal shape, and positioning accuracy, as well as solar glint.²

2.7 Lab Results Summary

Figure 20 summarizes the contrast performance demonstrated in the lab as a function of angular separation for all starlight suppression systems considered above.

Before attempting to make any comparison, it is important to acknowledge that the lab experiments used different setups, as shown in Table 1, which summarizes the performance reached by each system and its main operating conditions: aperture type; DM type and format size; range of angular separations and azimuths over which the mean contrast is computed; central wavelength and bandwidth; polarization state; facility used for the demonstration; and vacuum versus in-air operation. The testbeds used even differ in a given facility, as e.g., the HCIT DST and its carbon fiber composite optical bench only became available in 2019.

For the off-axis monolith case, Table 1 includes the CLC, HLC, and VVC4 results but not the PIAA ones, which were significantly worse and obtained over a small separation range. Also, the on-axis monolithic HLC and SPC Roman coronagraph results were presented above for completeness but omitted in the summary table because the heavily obscured Roman entrance pupil does not represent a viable telescope aperture for HWO. Indeed, the HLC and SPC designed for




Fig. 20 Best azimuthal mean contrast (NI) demonstrated to date by different starlight suppression approaches and laboratory experiments over a $\sim 10\%$ spectral bandwidth. The x-axis shows the angular separation in units of λ/D , where D is the entrance pupil inscribed diameter, λ is the central wavelength of the bandpass. Coronagraphic results were obtained with either one or two DMs; for different aperture types: off-axis monolith (plain curves); off-axis segmented (dashed curves); on-axis monolith (dashed-dotted curve); and on-axis segmented (dotted curve). The sub-scale starshade results (orange curve) are independent of the aperture type considered. A summary of the experimental conditions used by the main starlight suppression demonstrations considered is given in Table 1.

the Roman aperture have off-axis (planet) throughput values below 5% (Fig. 21), too low to spectrally characterize Earth-like exoplanets within reasonable exposure times (Sec. 4.4).

For coronagraphs, the current best 10% bandwidth performance has been achieved on monolithic off-axis (clear) apertures. Thus far, the only coronagraph architecture to have achieved better than 10^{-9} mean contrast at a separation smaller than $4 \lambda/D$ over a full 360 deg dark hole is the CLC coronagraph operating behind a clear circular aperture. However, the CLC laboratory setup used for that demonstration relied on a poorly transmissive Lyot stop and had low ($\sim 8\%$ to 13%) off-axis core throughput in the central 3 to $4.5 \lambda/D$ region, likely too low for getting visible spectra of exo-Earths with a 6 m HWO mission (Sec. 4). Next in terms of contrast performance is the linear HLC, which also provided better than 10^{-9} mean contrast over $\sim 10\%$ bandwidth but only over a one-sided dark hole. The linear HLC lab setup offers slightly better core throughput ($\sim 10\%$ to 20%) than the CLC coronagraph in the central 3 to $4.5 \lambda/D$ region. The next best contrast performance, also achieved over a clear aperture and a one-sided dark hole, was obtained with the VVC4 coronagraph, which reached a contrast of a few 10^{-9} together with a significantly higher core throughput ($\sim 38\%$) at $3 \lambda/D$.

In terms of segmented off-axis results, the only direct comparison comes from the broadband contrast performance obtained by the VVC4 on a segmented aperture with a single DM (blue dashed curve of Fig. 20). It is on average $1.5\times$ worse (1.6×10^{-9} higher) than in the monolithic 1DM case (upper solid blue curve).

This provides a first order of magnitude of the segmentation effect, but only at contrast levels of order a few 10^{-9} , in a regime in which the ultimate VVC4 broadband performance is still limited by residual spatial defects in the vortex mask rather than by the aperture type. To better assess the impact of segmentation itself, a more systematic comparison of vacuum results must be conducted with the same coronagraph setups, changing only the aperture from monolithic to segmented. This should include comparative tests of the CLC and HLC coronagraphs (only used so far on monolithic apertures), the PAPLC (only tested in air), and improved VVC masks recently manufactured.

Finally, the on-axis segmented aperture results (dashed green curve obtained with a PIAACMC) are currently significantly worse than the results obtained with the VVC4 on an

Table 1 Summary of current broadband starlight suppression lab results and operating conditions (as of March 2024). The results are expressed in units of NI. CLC, classical Lyot coronagraph; HLC, hybrid Lyot coronagraph; VVC4, vector vortex charge 4 coronagraph; PAPLC, phase-apodized pupil Lyot coronagraph; PIAACMC, phase-induced amplitude apodization complex mask coronagraph. A 64×64 DM was used for the HLC test⁽¹⁾, but the aperture diameter only extended over 48 actuators. In the CLC, VVC4 monolith, and PAPLC cases, the results are shown over $\sim 10\%$ and $\sim 20\%$ to 25% spectral bandwidths. The latter results are indicated in italics. All results were obtained in vacuum, except for the PAPLC and starshade mask. In the starshade case, starlight is canceled before the collecting aperture: the results are independent of the telescope aperture type and no active wavefront control is used (no DMs).

Lab setup	CLC	HLC	VVC4	VVC4	PAPLC	PIAACMC	Starshade mask
Aperture type	Off-axis monolith			Off-axis segmented		On-axis segmented	Any
DMS	2 AOX 48×48	1 AOX $64 \times 64^*$	2 AOX 48×48	1 BMC 2k	2 BMC 1k	1 BMC 1k	None
Central wavelength (μm)	0.550	0.800	0.635	0.635	0.660	0.650	0.680
Spectral bandwidth	10% 20%	10% 20%	10% 20%	10%	9% 25%	10%	12%
Number of polarizations	1	2	1	1	2	1	1
Dark hole separation range	3 to $8\lambda/D$ 5 to $13\lambda/D$	3 to $15\lambda/D$	3 to $10\lambda/D$	3 to $10\lambda/D$	2 to $13\lambda/D$ 2 to $12\lambda/D$	3.5 to $8\lambda/D$	1.7 to $7\lambda/D$
Dark hole extent	Full <i>One-side</i>	One-side	One-side	One-side	One-side	One-side	Full
Mean NI over dark hole	4×10^{-10}	5.2×10^{-10}	1.6×10^{-9}	4.7×10^{-9}	4.2×10^{-8}	1.8×10^{-8}	2×10^{-11}
Mean NI at larger bandwidth	4×10^{-10}	7.8×10^{-9}	5.9×10^{-9}	1.1×10^{-8}	9.5×10^{-8}	$\sim 7 \times 10^{-8}$	2.1×10^{-11}
NI at $3\lambda/D$	1.6×10^{-9}	6.0×10^{-10} 2.3×10^{-9}	2.4×10^{-9} 8.9×10^{-9}	1.1×10^{-8}	2.4×10^{-7}	$\sim 7 \times 10^{-8}$	2.1×10^{-11}
Testbed core throughput at $3\lambda/D$ and $4.4\lambda/D$	0.08/0.13	0.10/0.19	0.38/0.46	0.38/0.46	0.51/0.53	0.60/0.61	0.68/0.68
Facility and testbed	HCIT-2 DST	HCIT	HCIT-2 DST	HCIT-2 DST	HiCAT	HCIT-2	Princeton Frick
Vacuum operation	Y	Y	Y	Y	N	Y	N
Main reference	Seo et al. ⁴⁵	Trauger et al. ⁴⁸	Ruane et al. ⁵⁷	Riggs et al. ⁵⁸	Por et al. ⁶⁴	Marx et al. ⁷⁴	Harness et al. ³²

Note: bold characters are used to highlight key differentiators between the various set-ups.




Fig. 21 Core throughput of the laboratory starlight-suppression setups presented in Sec. 2. Core throughputs are given as a function of angular separation in units of λ/D , where λ is the central wavelength and D is the inscribed diameter of the entrance aperture. At a given separation, the core throughput is computed within a circular aperture of radius $0.7 \lambda/D$ for all systems, except for the Roman HLC and SPC-spec coronagraphs, which have highly spatially extended PSFs and for which the PSF FWHM region is used instead. In the starshade case, the angular separation is computed at $0.683 \mu\text{m}$, the central wavelength of the Princeton 12% bandwidth lab setup, and the core throughput is nearly achromatic versus physical angular separation in arcsec. The PIAACMC curve assumes that inverse PIAA optics are used to correct off-axis PSF distortion (although they were not part of the original lab setup).

off-axis segmented aperture and a single DM as well (blue dashed curve). However, it must be recognized that, with only one vacuum coronagraphic experiment conducted so far, the testbed time devoted to on-axis segmented aperture demonstrations has also been significantly less than for off-axis apertures.

In terms of laboratory broadband starlight suppression results, the best contrast performance to date has been achieved with the starshade Princeton testbed. The mean contrast demonstrated by a $\sim 1/1000$ th scale starshade mask operating at a flight-like Fresnel number was 2×10^{-11} over a 360 deg dark hole extending from 1.7 to $7 \lambda/D$, with a core throughput of 68%. At a separation of $2 \lambda/D$, the contrast was 2×10^{-10} , again with a core throughput of 68%.

3 Beyond Raw Contrast

Given the extreme level of starlight suppression required for direct exo-Earth observations, laboratory experiments have mainly concentrated so far on demonstrating deep broadband ($> \sim 10\%$) RC at small angular separations, as reported in Sec. 2. However, once the contrast performance is good enough that residual starlight levels fall significantly below the signal from irreducible astrophysical background sources (zodi and exozodi dust), the exposure time required to characterize exo-Earth planets becomes independent of RC. As a result, the predicted yield increases marginally with further RC improvements.³⁹ Depending on the stellar distance and wavelength, that transition occurs at a threshold RC level that ranges from $\sim 10^{-10}$ to $\sim 10^{-8}$ (Sec. 4.1).

Once that RC threshold performance is reached, the number of exo-Earths characterized depends primarily on two starlight suppression KPPs: the overall optical system off-axis core throughput, especially its value at small angular separations, and the post-calibrated contrast uncertainty achieved after estimating and subtracting residual starlight speckles in raw science images. These KPPs are discussed in Secs. 3.1 and 3.2 below.

3.1 Core Throughput

The core throughput is the fraction of light from a point source at a given off-axis location that passes through the coronagraph or starshade masks and falls within a specified photometric

aperture centered at that location. The photometric aperture can be chosen to optimize the detection of off-axis point sources, e.g., using a matched filter. The matched filter will depend on the telescope aperture type, the field-dependent PSF shape at the planet location, and the background structure. It may provide significant sensitivity (SNR) gains over simple photometric apertures commonly used, such as a circular aperture with a radius equal to the off-axis PSF full width at half maximum (FWHM) or to a fixed value that maximizes the SNR of an Airy pattern against a spatially uniform background. (In the case of a perfectly circular beam of diameter d , the photometric aperture radius value that maximizes the signal-to-noise of an Airy pattern against a spatially uniform background is $0.665 \lambda/d$. Assuming a Lyot stop whose diameter is 95% that of the telescope pupil of diameter D , this corresponds to $0.7 \lambda/D$.) With two exceptions, all throughput values reported in this section use the latter approach and a constant circular photometric aperture of radius $0.7 \lambda/D$, where D is the inscribed circular diameter of the telescope. The exceptions are the values for the Roman HLC and SPC-spec coronagraphs, which have bright PSF wings. In these two cases, the FWHM definition is used instead. The resulting photometric aperture radii are $0.57 \lambda/D$ for HLC and $\sim 0.88 \lambda/D$ (averaged between the X and Y axes) for SPC-spec.⁶⁸

In the coronagraph case, the core throughput value captures transmission losses and off-axis PSF broadening due to the FPM, Lyot stop, pupil amplitude apodization (if any), and the DM settings used to generate the dark hole. In the starshade case, the core throughput is approximately given by the product of the off-axis transmission of the laboratory mask by the fraction of the telescope PSF falling within the $0.7 \lambda/D$ radius photometric aperture (0.68 for a perfect Airy pattern).

Figure 21 shows the theoretical core throughput versus angular separation of the tested starlight suppression systems presented in Sec. 2, under the experimental conditions and aperture types used in the laboratory. The Roman HLC and SPC coronagraphs exhibit significantly lower core throughputs than the other coronagraphs not because these coronagraph types have intrinsically low core throughput but because of the heavily obscured Roman telescope entrance pupil.⁶⁸ Their core throughput would be significantly higher on the less obscured aperture expected for HWO. The CLC, with a full 360-deg dark hole and 10% bandwidth, used a Lyot stop blocking a substantial fraction of the aperture and is lowest in terms of core throughput performance for the HWO aperture. Depending on separation, it is three to four times lower than for the other coronagraph experiments considered, which are aimed at preserving high throughputs at small separations. The PAPLC, PIAACMC (+ inverse PIAA inserted before science focal plane), and VVC4 all have $>20\%$ off-axis core throughput values down to $2 \lambda/D$ or less. The starshade core throughput curve is quite steep and reaches its maximum value when the off-axis point source separation exceeds the angle sustained by the starshade when seen from the telescope, known as the “tip IWA,” which corresponds to $\sim 1.8 \lambda/D$ at the mean operating wavelength of $0.683 \mu\text{m}$.

Comparing Figs. 20 and 21, the coronagraphs with the best contrast performance currently demonstrated in the lab (e.g., CLC and linear HLC) tend to have a lower off-axis throughput, whereas those with the best sensitivity to sources located at small separations (e.g., PIAACMC) tend to have worse demonstrated contrast. Given that the HCIT essentially provides a static wavefront, this observed trend does not reflect the higher sensitivity to aberration drifts expected for coronagraphs with small IWA, which should only show up in dynamic testing. Yet, this small statistics trend is still observed.

At all separations larger than $\sim 1.5 \lambda/D$, the theoretical core throughput is the highest in the starshade case as there are neither occulter nor Lyot-plane losses inside the telescope. It reaches its maximum value of 68% (a perfect Airy pattern) at the tip of IWA.

3.2 Residual Contrast after Wavefront Control and Image Processing

Dark hole contrasts will degrade in the presence of drifts. Active wavefront control for RC stabilization and image calibration techniques have received more attention in recent years. We highlight hereafter some of these developments. As detailed in Sec. 4, the detection and spectral characterization of faint exoplanets require a combination of optical starlight suppression and image calibration (called “post-processing” in the coronagraph community) techniques. Depending on the stellar distance, observing wavelength, and RC performance, planets of interest might be significantly dimmer than the residual starlight speckles and other background

sources, such as the solar zodiacal light signal (spatially uniform over the dark hole) or the spatially varying exozodiacal signal in the target planetary system. Of specific interest is the accurate estimation and calibration of starlight residuals, which have the added complexity of varying over short timescales. As shown in Ref. 42, the relevant quantity is the spatial standard deviation of the starlight speckles after image calibration. Any residual error in the estimation of the speckles field in the science images will result in an irreducible systematic noise floor and a finite planet-to-star flux ratio detection limit, whatever the exposure time (Sec. 4.2). This error and the resulting post-calibration contrast map residuals can be minimized in two complementary ways: (1) RC stabilization using WFS&C during the observations and (2) accurate knowledge and subtraction of uncorrected residual speckles present in science images. The first method corrects for wavefront drifts in real time. It aims both at improving RC over the science exposure and reducing the post-calibrated starlight residuals induced by RC fluctuations. The second method is “after the fact,” meaning that even if some speckle fluctuations cannot be perfectly corrected in real time (e.g., due to time lag or imperfect DM calibration), knowing what they were may further improve post-calibration residuals and get closer to the photon noise limit. We give hereafter some examples of speckle stabilization and speckle estimation techniques. We also acknowledge that a given WFS&C approach may both stabilize contrast and provide estimates of any uncorrected contrast fluctuations.

3.2.1 Estimating and correcting the electric field

Wavefront sensing and control are commonly used for ground-based observations to reduce the effect of turbulence and for laboratory deep starlight suppression demonstrations using coronagraphs. As very-high-contrast coronagraphic observations with HWO are concerned, an important goal is to limit non-common path errors between the wavefront sensor and the science beam path. This can be accomplished via direct measurements in the science focal plane or by picking up light reflected by the FPM. To reach deep contrast, both low-order and high-order WFS&C are required, likely over different timescales. There have been numerous high-contrast WFS&C experiments worldwide, and we only highlight a few here to illustrate some approaches relevant to HWO, starting with low-order WFS&C.

One notable laboratory experiment is the active sensing and control of low-order aberrations ($Z_2 - Z_{11}$: tip-tilt, defocus, astigmatism, coma, trefoil, and spherical) that were demonstrated at flight such as photon flux in one of the HCIT vacuum chambers for the Roman coronagraph.⁷⁷ In this case, the low-order WFS&C systems use a low-order Zernike wavefront sensor (ZWFS), which picks up starlight reflecting off the metallic occulting spot of the FPM, which is $\sim 6\lambda/D$ in diameter. The system was able to maintain broadband (11% bandwidth centered at $0.55\ \mu\text{m}$) coronagraph contrast below 10^{-8} with stability better than 10^{-9} over periods of up to an hour in the presence of a purposely injected line of sight and low-order wavefront changes commensurate with the expected Roman space telescope on orbit jitter and thermal drift levels. Perturbations were injected at one DM and then corrected at the other (located in a pupil plane conjugate) based on the ZWFS signal. This performance was demonstrated in both SPC and HLC modes, down to the photon flux expected from a $V = 2$ star. At the lower flux expected from a $V = 5$ star (closer to the magnitude of a typical HWO target), the low-order WFS&C system still suppressed the Roman-like line-of-sight disturbances ($Z_2 - Z_3$) down to a post-correction jitter level of 0.35 mas rms.⁷⁷ Additional $V = 5$ tests demonstrated the ability of the low-order WFS&C system to keep the 10-min mean values of each low-order aberration ($Z_4 - Z_{11}$) within 10 pm over a period of 10 h (Brandon Dube, JPL, personal correspondence). Wavefront stabilization over such timescales is in line with the Roman coronagraph operations concept, which baselines a series of 2-h-long individual target star observations at alternating telescope rolls, together with higher-order touch-ups on a bright reference star every 10 h or so. Low-order wavefront control of the first 11 Zernike modes using an FPM-filtered ZWFS was also demonstrated in air on the HiCAT testbed, using instead a CLC and a segmented aperture.⁷⁸ In that experiment, the in-air mean dark hole level between 7.6 and $12.8\lambda/D$ was stabilized at a contrast level of 7×10^{-8} with a standard deviation of 7×10^{-9} over a period of 15 min, in the presence of random low-order ($Z_2 - Z_{11}$) wavefront perturbations purposely introduced every 2 s.

An inherent limitation of these experiments is that, because the reflected starlight beam is spatially filtered by the small FPM occulter, only low spatial frequencies (up to three or four cycles across the pupil) can be sensed by the ZWFS. Sensing residual phasing errors on a segmented telescope will require a WFS that is sensitive to higher spatial frequencies than that of the Roman coronagraph's WFS. HWO could have five or more mirror segments across the primary. Segment-to-segment errors in the form of rigid body motions will require at least a few wavefront samples per segmented mirror area and at least two across to measure relative tilts. As a result, the FPM reflective region needs to be more than $20\lambda/D$ in diameter to detect and correct such errors. If the WFS needs to measure DM instabilities at the actuator level, many more samples (potentially 48 to 128 actuators across) and a much larger reflecting region (respectively, >24 to $64\lambda/D$) will be required.

A first possible solution to sense these higher spatial frequencies at high sensitivity is to use out-of-band wavefront sensing as is commonly done by ground-based adaptive optics systems. It was also proposed in the context of the LUVOIR study⁷⁹ in which one of three parallel coronagraph channels, operating at different wavelengths, would be used for sensing the segment motions of the primary and the other two would be used for science integrations. In that case, the ZWFS would be fed by a dichroic located upstream of the science coronagraphs, resulting in a fair amount of non-common path errors between the WFS and science channels. Other implementations, including a dual-purpose Lyot coronagraph (DPLC) mask for simultaneous high-contrast imaging and high spatial resolution wavefront sensing,⁸⁰ have been suggested. In this proposed approach, a reflective FPM is used to feed light to the WFS and minimize non-common path errors. The design uses a tiered metallic focal plane occulter to suppress the starlight transmitted through the mask at the science wavelength and a dichroic-coated substrate that reflects out-of-band light to a ZWFS.

A second approach to sense and correct high-order wavefront drifts in real time is the so-called “dark zone maintenance” (DZM) algorithm.⁸¹ In that case, the wavefront drift is estimated via dark hole intensity measurement while dithering the DM actuators (similar in some sense to lock-in amplification). In addition, a predictive scheme based on an extended Kalman filter is used for maintaining the dark hole. This DZM approach has been validated in air in the HiCAT laboratory⁸² at low star-equivalent flux levels with monochromatic $0.638\ \mu\text{m}$ light. Over an annular dark zone extending from 5.8 to $9.8\ \lambda/D$, the closed-loop contrast was maintained at 5.3×10^{-8} with a standard deviation of 6.4×10^{-9} over a period of ~ 35 h, in the presence of a DM random walk drift of 20 pm rms per iteration (Fig. 22).

A third approach to enable fast monitoring of wavefront changes at all spatial frequencies is linear dark field control (LDFC).^{83,84} In that case, there is no DM probing or dithering: the contrast in the dark hole is locked by monitoring the temporal evolution of bright speckles “outside”

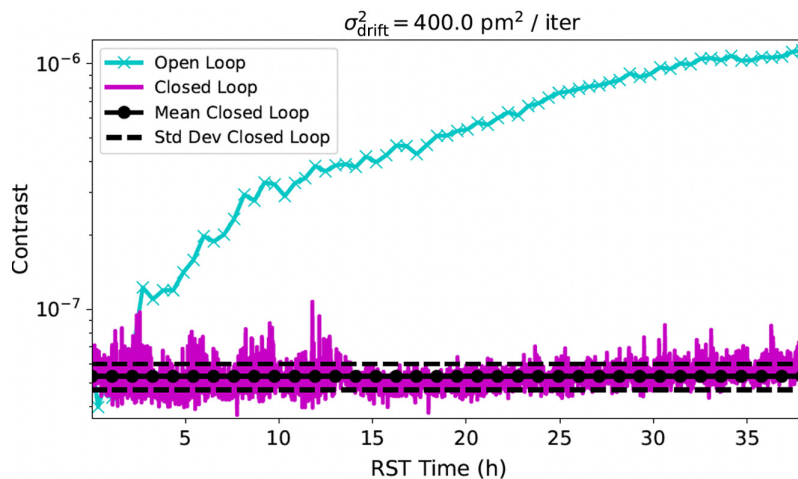


Fig. 22 Mean dark-zone contrast measured versus time for the HiCAT low SNR DZM experiment. In the presence of a DM random walk drift of 20 pm rms per iteration, the open-loop contrast diverges to 1.1×10^{-6} after 35 h. Conversely, the closed-loop contrast is maintained at 5.3×10^{-8} with a standard deviation of 6.4×10^{-9} . The iteration time is 39 s. (From Ref. 82.)

the dark hole, either spatially or spectrally (as in the DPLC). Changes in these bright field (BF) regions are highly correlated to the same wavefront changes that spoil the deep halo suppression in the dark field (DF). As the BF images are significantly brighter than the DF images, they can be acquired at a higher cadence, and no starlight needs to be directed to the DF during science exposures. By calibrating or computing the linear changes in the BF against wavefront changes, an LDFC servo can in principle maintain high contrast in the DF during science exposures. An initial demonstration was conducted in air with the Ames coronagraphic experiment testbed^{85,86} using spatial LDFC (i.e., bright uncorrected regions) to stabilize a 5×10^{-7} dark hole extending from 1.5 to $5.2 \lambda/D$ down to an rms stability level of 6.5×10^{-8} in the presence of injected perturbations.

An interesting upcoming trade for WFS&C in the context of HWO's coronagraphs is then between the LDFC approach using bright light outside the dark hole, the out-of-band high-order DPLC-like ZWFS approach, and a more conventional (Roman coronagraph-like) low-order in-band ZWFS working together with a (DZM-like) DM-dithering scheme with sensing in the science dark hole. Testing of all three approaches using a common testbed and coronagraph mask would be most informative.

3.2.2 Estimating and subtracting residual speckles

Whatever WFS scheme is used to minimize the instrumental RC and its fluctuations, a residual starlight speckle field will remain at some level in the science images. For optimal exoplanet detection, the RC should be deep enough that residual starlight is fainter than the solar zodi plus exozodi background. The speckle field must then be estimated and subtracted accurately enough to get back to that background shot noise limit. For a given science image, the speckle field may be estimated using a library of coronagraphic images obtained at different times on the same target (but at a different roll angle of the telescope) or on reference targets. This is the principle of, e.g., the successful locally optimized combination of images (LOCI)⁸⁷ and Karhunen-Loeve image projection (KLIP)⁸⁸ high-contrast post-processing approaches commonly applied to reference (RDI) and angular differential imaging (ADI) of coronagraphic observations down to detection limits of $\sim 10^{-6}$ to 10^{-5} , from the ground and with HST⁸⁹/JWST.^{90,91} However, these are “blind” correction methods, in which the speckle field at the time of science observations is estimated through observations obtained at a different time and possibly on a different star (RDI). Accordingly, such methods place strong requirements on wavefront and speckle temporal stability during telescope slews (RDI) or rolls (ADI). They may also suffer from variations in astrophysical source properties, e.g., stellar diameter and exozodi level variations between the two RDI stars or target source asymmetries in the case of ADI.

A first alternative is to leverage auxiliary wavefront sensing data to estimate the starlight electric field at the time of the science observations.⁹² In particular, if BFs recorded at other wavelengths (e.g., DPLC WFS) or in different regions (LDFC WFS) of the science focal plane can indeed be used to reliably estimate the starlight residuals in the dark hole, then high-accuracy estimates of the speckle field can be made over much shorter timescales than when using the DF science images themselves. If such speckle “self-calibration” estimates⁹³ can be made faster than the speckle field changes, wavefront fluctuations will be accurately known. Even if the DMs cannot perfectly correct for these measured fluctuations in real time due to finite spatio-temporal response or chromatic effects, such fluctuations may be corrected after the fact in science images, potentially resulting in a significant relaxation of wavefront stability requirements. In this case, RC stability needs only to be maintained at a level below the (exo)-zodi background and no longer at a fraction of the planetary signal. As long as this condition is met, this essentially becomes a wavefront knowledge problem. Importantly, this scheme requires that the relationship between the WFS BF data and the DF is constant over time or at least that any time variable offset is slow enough to be measured. Machine learning algorithms trained to estimate the electric field in the DF based on BF measurements are being investigated.⁹⁴ Post-processing schemes using WFS information directly gathered in the dark hole at the science wavelength would likely suffer less from that limitation and have been suggested as well, for instance, using DM dither and reduced order estimation of the speckle electric field.^{95,96}

A second alternative is the use of medium to high spectral resolution ($R > 1000$) observations to distinguish residual starlight and (exo)-zodi background signals, essentially featureless or showing slow spectral variations, from planetary (or low-mass companion) signals that show distinctive molecular absorption features at specific wavelengths and exhibit line forests over the science bandpass. In this approach, major atmospheric species can be searched for one by one in the companion's atmosphere by fitting the spectra obtained with model templates of individual molecules (the “molecular mapping” technique). Initial demonstrations were obtained at a spectral resolution high enough to leverage the planet differential radial velocity relative to the star, resulting in the direct detections of young extrasolar giant planets with integral field spectrographs at both VLT^{97–99} and Keck^{20,100} and providing direct measurements of planet molecular content, as well as temperature, surface gravity, and radial and spin velocities. Whether this approach might be applicable to the significantly fainter (10^{-10} versus 10^{-5}) planets to be spectrally characterized by HWO remains to be investigated. Previous studies of such high dispersion coronagraphy in the case of the HabEx and LUVOIR designs^{101,102} suggested that a spectral resolution of ~ 1000 might be optimum. However, this value strongly depends on wavelength, detector properties, and overall system throughput and should be reassessed in the HWO context. We also note that this method recently demonstrated the detection of an $\sim 10^{-5}$ close-in brown dwarf companion with JWST using medium resolution ($R \sim 2700$) NIRSpec IFU data¹⁰³ and is also being developed at Keck in conjunction with a coronagraph for the first time, using a single-mode fiber to feed a high-spectral-resolution ($R \sim 35,000$) point source spectrograph.^{104,105}

Overall, the current developments in WFS&C, starlight speckle estimation, and spectral disambiguation point to two major upcoming trades in the context of HWO coronagraph design and requirements. The first trade is between speckle stability and speckle knowledge, assessing how far stability requirements might be realistically relaxed. The second one is about spectral resolution and detector noise, assessing how an increase in spectral resolution (with respect to the $R \sim 100$ value commonly assumed) might be possible and what benefits it might bring.

4 Illustrative Cases

In Sec. 2, we presented the current state of the art in terms of RC demonstrated in the lab. In Sec. 3, we introduced the other two main characteristics of a starlight-suppression system: off-axis (core) throughput and post-calibration contrast.

This conceptual framework can now be applied to start assessing how the currently demonstrated lab performance compares with the notional needs of HWO. HWO's specific exoplanet science figures of merit, characterization requirements, and candidate architectures are not defined yet, making complete design reference mission simulations premature. Instead, we concentrate here on a fiducial exo-Earth twin case and estimate the exposure time required to detect and spectrally characterize it for different combinations of starlight suppression system key performance KPPs. This approach using exposure time as a basic cost function already enables a rough exploration of the starlight suppression performance needed by HWO and the identification of several important design trades.

4.1 Astrophysical Signals

In this section, we aim to compare the signals coming from the solar zodiacal light, the exozodiacal light, the planet light, and the residual starlight. Following Stark et al.,³⁸ we adopt a mean surface brightness of the solar zodi in the V band of 23 mag/arcsec^2 and a surface brightness of 22 mag/arcsec^2 for one zodi worth of exozodi dust located in the habitable zone of a Sun-like star. Indeed, a solar zodi cloud analog would appear brighter when observed from afar as we would typically observe dust both above and below the mid-plane, and the $1/r^2$ illumination factor will always bias the background flux in any photometric aperture to higher values.

Figure 23 shows the “incident” astrophysical fluxes detected at 1 AU from a Sun-like (G2V) star, in a photometric region of radius $0.7 \lambda/D$, as a function of stellar distance, at wavelengths of 0.25, 0.55, 1.0, and 1.8 μm . All cases assume a 6 m (inscribed) diameter telescope with a starlight suppression system with unit end-to-end transmission (including detector) efficiency and an off-axis PSF described by a perfect Airy pattern i.e., with a core throughput of 0.68 within the prescribed photometric region. If the off-axis PSF is more extended than for a perfect Airy pattern,

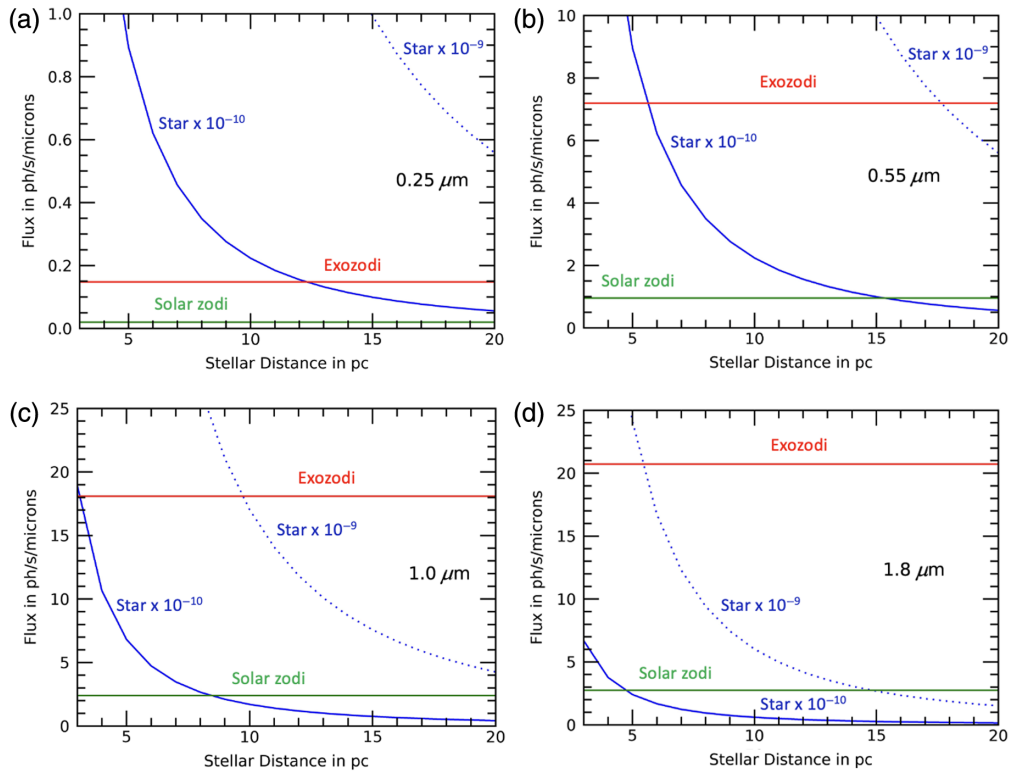


Fig. 23 Astrophysical fluxes detected at 1 AU from a Sun-like (G2V) star, in a photometric region of radius $0.7 \lambda/D$, as a function of stellar distance. The four panels correspond to wavelengths of (a) 0.25, (b) 0.55, (c) 1.0, and (d) 1.8 μm (note the different y-axis ranges). All cases assume a $D = 6$ m telescope with a starlight suppression system with unit end-to-end transmission (including detector) efficiency and an off-axis PSF described by a perfect Airy pattern (i.e., with a core throughput of 0.68 within the photometric region). The starlight suppression RC at 1AU is set to either 10^{-10} (plain blue curve) or 10^{-9} (dotted blue curve). The blue curves also correspond to the flux from an exoplanet with a planet-to-star flux ratio of 10^{-10} (an exo-Earth at quadrature) or 10^{-9} . The exozodi level is set to three zodi in all cases.

the stellar and planetary fluxes per photometric aperture will decrease accordingly. Conversely, the solar and exozodi fluxes per photometric aperture will remain essentially constant. The starlight-suppression RC at 1AU is set to either 10^{-10} (plain blue curve) or 10^{-9} (dotted blue curve). Because of the way that RC is defined, the plain blue curve also corresponds to the flux from an exoplanet with a planet-to-star flux ratio of 10^{-10} (an exo-Earth at quadrature), and the dotted blue curve corresponds to a planet flux ratio of 10^{-9} . The exozodi level is set to three zodi in all cases, the most likely median exozodi level derived from the LBTI HOSTS survey.¹⁰⁶ As a result, the assumed exozodi V-band surface brightness in the habitable zone is 20.8 mag/arcsec². Because the solar zodi and exozodi sources are spatially extended, their corresponding signal within the $(0.7 \lambda/D)$ photometric aperture is constant versus stellar distance but increases with wavelength. Conversely, the residual starlight signal decreases with stellar distance, meaning that observations will become dominated by “background” (exo)-zodi signals for distant-enough targets and/or deep-enough starlight cancellation.

At V band [around 0.55 μm , Fig. 23(b)], assuming an instrumental RC level of 10^{-10} , the (three zodi) exozodi signal dominates over starlight residuals, or equivalently the signal from a 10^{-10} Earth-like planet, for all stars further away than 6 pc, i.e., most of the targets. This says that at V band, unless most stars have significantly less than three zodi worth of dust, an RC of 10^{-10} or better will guarantee that most observations are limited by the exozodi background rather than starlight residuals, at least if the latter is constant or perfectly calibrated. The irreducible solar system zodi signal only dominates for stars further than ~ 15 pc. On the other hand, an RC of 10^{-9} at V-band results in starlight residuals being the dominant source of photon noise for all Sun-like stars within ~ 17 pc, likely most of the sample. However, if most stars have significantly

more than three zodis worth of dust, a high exozodi scenario still compatible with the LBTI exozodi survey data that derived a 95% confidence upper limit of 27 zodis on the median exozodi level of Sun-like stars, exozodiacal light will still dominate. Depending on star distance, an RC between 10^{-10} and 10^{-9} will hence be required to guarantee that background (exo-)zodi signals dominate over stellar shot noise.

At shorter wavelengths, e.g., in the near UV around $0.25 \mu\text{m}$ [Fig. 23(a)], the solar and exozodi backgrounds are both reduced due to the smaller beam and photometric region size compared with V band. An RC of 10^{-10} or lower is required to remain exozodi-background-limited for near UV observations of stars closer than ~ 12 pc. The RC requirements are hence most stringent at near UV wavelengths.

At longer wavelengths of $1 \mu\text{m}$ (close to the strong $0.94 \mu\text{m}$ water band) and $1.8 \mu\text{m}$, we see the opposite effect. Most observations would still be limited by the solar and exozodi backgrounds, even for RC as bad as 10^{-9} or even close to 10^{-8} [Figs. 23(c) and 23(d)]. The RC requirements are in principle least stringent at the longer infrared wavelengths. However, we caution that operating at an RC of 10^{-9} or higher will make the calibration of residual speckles below the planet level more challenging as it will require speckle calibration at even better relative precision to maintain the same noise floor. Exo-Earth observations at near-infrared wavelengths would also put a premium on the precise estimation of the exozodi signal at the planet's location. Indeed, at $1.8 \mu\text{m}$, three zodis worth of exozodi dust would create a signal $\sim 50\times$ brighter than an exo-Earth seen at quadrature around a star located at 12 pc.

It is worth highlighting that the level of exozodiacal light around Sun-like stars has been best estimated in the mid-infrared, down to an uncertainty of ~ 50 zodis (1σ) per individual star, assuming a solar-like grain size and density profile. It is much less constrained at HWO's near-UV to near-IR wavelengths. As shown in Fig. 23, even at only $3\times$ the solar system level, exozodi dust represents the dominant source of background shot noise for most exo-Earth observations at visible to near-infrared wavelengths. At 10 to $20\times$ the solar density level, exozodi dust clouds may also create confounding resonant structures.¹⁰⁷ The exozodi surface brightness—including its possible time variability^{108,109}—has then a large impact on the observability of exo-Earths around individual targets. Precursor observations at or close to HWO's wavelengths would thus be very valuable and should be conducted. Among them, an exozodi survey conducted by the Roman coronagraph at visible wavelengths would be particularly impactful.¹¹⁰⁻¹¹²

4.2 Signal-to-Noise Ratio and Time to SNR

For signal-to-noise calculations, we adopt the definitions and formalism described in previous work.⁴² A summary of all variable notations and definitions is given in Table 2. Over exposure time δt , for a given spectral bandpass, and assuming no systematic noise (photon noise only), the planet observation SNR is given as

$$\text{SNR} = \frac{r_{pl} \cdot \delta t}{\sqrt{r_n \cdot \delta t}}, \quad (2)$$

where r_{pl} is the planet photon rate detected in a circular photometric aperture centered around the planet location and r_n is the total photon rate detected from all astrophysical sources in the aperture: planet (r_{pl}), residual starlight (speckle rate r_{sp}), solar zodiacal light (r_{sz}), and exozodiacal light (r_{xz}). Strictly speaking, the planet photon noise contribution should only be included for a characterization (i.e., a flux measurement) SNR and not for a detection SNR. By definition

$$r_n = r_{pl} + r_{sp} + r_{sz} + r_{xz}. \quad (3)$$

Following the detailed calculations presented in Ref. 43, we have

$$r_{pl} = \eta_p \cdot \epsilon \cdot N_s, \quad (4)$$

and

$$r_{sp} = \eta_p \cdot \text{RC} \cdot N_s, \quad (5)$$

where η_p is the field-dependent point source core throughput computed at the planet location (x, y) over the specified photometric aperture, as defined in Sec. 3.1. Following Ref. 42, it is

Table 2 Summary and definition of variables used to estimate the SNR of exoplanet observations.

Variable	Definition
Dist	Stellar distance
D	Telescope diameter
Tr	Overall transmission and reflectivity of all non-coronagraphic optics in optical system
R	Spectral resolution
r_{pl}	Planet photon rate detected in photometric aperture centered at planet location, in photo-electrons per second
δt	Exposure time in seconds
r_n	Total photon rate from all, astrophysical sources detected in photometric aperture centered at planet location, in photo-electrons per second
r_{sp}	Residual starlight photon rate detected in photometric aperture centered at planet location, in photo-electrons per second
r_{sz}	Solar zodi photon rate detected in photometric aperture centered at planet location, in photo-electrons per second
r_{xz}	Exo-zodi photon rate detected in photometric aperture centered at planet location, in photo-electrons per second
ϵ	Mean astrophysical (intrinsic) planet-to-star flux ratio over the bandpass
N_s	Stellar photon rate detected over the whole focal plane in the absence of any coronagraph or starshade masks, in photo-electrons per second
RC	RC at planet location (see Sec. 2.1)
η_p	Core throughput at planet location, within photometric aperture (see Sec. 3.1)
η_{occ}	Occluder transmission resulting from all coronagraphic masks or the external starshade
η_{PSF}	Fraction of light in the planet PSF contained within photometric aperture
$f_{\Delta I} (<1)$	Starlight speckle subtraction effectiveness at planet location, after data calibration
$f_{\Delta Isz} (<1)$	Solar zodi subtraction effectiveness at planet location, after data calibration
$f_{\Delta Ixz} (<1)$	Exozodi subtraction effectiveness at planet location, after data calibration
$\sigma_{\Delta C} = f_{\Delta I} * RC$	1σ point source detection limit at planet location, after data calibration
SNR_0	Specified SNR required for detection or characterization
t_{SNR0}	Exposure time required to reach specified SNR
ϵ_{min}	Minimum planet-to-star flux ratio detectable at specified SNR

defined as the product of the overall occulter transmission $\eta_{occ}(x, y)$ (resulting from all coronagraphic masks or the external starshade) and the fraction of light in the planet PSF that ends up in the photometric aperture $\eta_{PSF}(x, y)$, so

$$\eta_p(x, y) = \eta_{occ}(x, y) \cdot \eta_{PSF}(x, y). \quad (6)$$

In the starshade case, $\eta_{occ}(x, y)$ is the off-axis transmission profile of the external occulter, reaching ~ 1 at the starshade petal tip. In the coronagraph case, the transmission of the FPM $\eta_{FPM}(x, y)$ is multiplied by the constant transmission of the Lyot stop (η_{LS}) and of any pupil plane mask apodizer (η_{PPM}), so

$$\eta_{occ}(x, y) = \eta_{PPM} \cdot \eta_{LS} \cdot \eta_{FPM}(x, y). \quad (7)$$

For uniformly extended sources that are spatially uniform over larger scales than the specified photometric aperture, such as zodiacal light and exozodi to a high degree, there are no PSF losses, and only the η_{occ} factor applies, so the equivalent throughput factor accounting for (exo)-zodi transmission through the starlight suppression system is given to the first order by

$$\eta_{sz}(x, y) = \eta_{xz}(x, y) = \eta_{\text{occ}}(x, y) = \eta_p(x, y) / \eta_{\text{PSF}}(x, y), \quad (8)$$

where

- ε is the mean astrophysical (intrinsic) planet-to-star flux ratio over the bandpass.
- RC is the starlight suppression system “raw contrast” at the planet location, defined in Sec. 2.1.
- N_s is the total number of stellar photons that would be detected per second through the optical system over some bandpass $\Delta\lambda$ in the absence of any coronagraphic (pupil and/or focal plane) masks, Lyot stop, or starshade mask, integrating starlight received over the whole focal plane, that is,⁴³

$$N_s = \int_{\Delta\lambda} \phi_s(\lambda) A \cdot q(\lambda) \cdot Tr(\lambda) \cdot d\lambda, \quad (9)$$

where $\phi_s(\lambda)$ is the incoming stellar flux in photons per unit area per unit time per unit wavelength at the primary mirror; A is the collecting area of the telescope; $q(\lambda)$ is the detector finite quantum efficiency; and $Tr(\lambda)$ is the overall transmission and reflectivity of all non-coronagraphic optics in the system that affects the star, planet, and (exo)-zodi light equally, including any spectral filters and polarizers.

As mentioned, the planet signal will be accompanied by a background signal made up of three components: residual starlight, solar zodi, and exozodi, some of which could be significantly brighter than the detected signal from an exo-Earth depending on stellar distance and observing wavelength (Sec. 4.1). Some form of differential imaging based on advanced post-processing techniques such as reference, angular, spectral, polarization, coherent differential imaging, or wavefront sensor-based estimates will be required to properly estimate the background and subtract it. The background subtraction will generally be imperfect, meaning that the planet SNR will include a systematic post-calibration error term and is rewritten as

$$\text{SNR} = \frac{r_{pl} \cdot \delta t}{\sqrt{r_n \cdot \delta t + r_{\Delta I}^2 \cdot \delta t^2}}. \quad (10)$$

The systematic noise term growing as $r_{\Delta I} \cdot \delta t$ captures the finite ability to calibrate the total background signal at the planet location after using post-processing. Following Ref. 42 and extending their approach to all three background sources, $r_{\Delta I}$ can be thought of as the residual total background rate after post-processing and is written as

$$r_{\Delta I} = f_{\Delta I} \cdot r_{sp} + f_{\Delta Isz} \cdot r_{sz} + f_{\Delta Ixz} \cdot r_{xz}, \quad (11)$$

where $f_{\Delta I} (<1)$ is the effectiveness of the starlight differential imaging suppression. Similarly, the $f_{\Delta Isz}$ and $f_{\Delta Ixz}$ coefficients measure the ability to reduce the effective solar and exozodi flux after image calibration. The solar zodi signal is spatially uniform in the dark hole and constant over time to the first order, making it relatively straightforward to calibrate precisely. Conversely, the exozodi signal depends on the target star itself. It will be estimated from the science images through modeling of the exozodi brightness distribution and its intensity at the planet location. In the remainder of this paper, we assume that the zodi and exozodi signals are perfectly calibrated, i.e., $f_{\Delta Isz} = f_{\Delta Ixz} = 0$. This is clearly an optimistic assumption for exozodi, especially at longer wavelengths, where exozodi may get significantly brighter than planet light (Sec 4.1). Regardless, we assume hereafter that the only systematic noise term remaining comes from the imperfect subtraction of starlight speckle residuals and their temporal fluctuations during science exposures, so

$$r_{\Delta I} = f_{\Delta I} \cdot r_{sp}. \quad (12)$$

In the case in which the speckle background mean value can be perfectly removed, $f_{\Delta I} = 0$ and the observations are purely shot noise limited. As a best-case scenario, we further assume zero detector noise in all subsequent SNR calculations.

Based on Eqs. (2)–(12) and assuming that differential imaging doubles the shot noise term (e.g., as in the case of ADI), the planet detection SNR is rewritten as

$$\text{SNR} = \frac{\eta_p \cdot \varepsilon \cdot N_s \cdot \delta t}{\sqrt{2 \cdot (\eta_p \cdot \varepsilon \cdot N_s + \eta_p \cdot RC \cdot N_s + r_{sz} + r_{xz}) \cdot \delta t + (f_{\Delta I} \cdot \eta_p \cdot RC \cdot N_s \cdot \delta t)^2}}. \quad (13)$$

If differential imaging is based on the observations of a brighter reference star or simultaneous observations of the target star (Sec. 3.2), the shot noise penalty will be lower, so Eq. (13) describes the worst-case impact of post-processing on SNR and exposure time required.

The post-processing factor $f_{DI} (< 1)$ corresponds physically to the ratio of the post-calibrated contrast spatial rms at the planet location, noted $\sigma_{\Delta C}$ hereafter, to the RC at the planet location. (It should not be confused with the post-processing gain used to compare the post-calibrated contrast obtained using two different post-processing techniques, e.g., KLIP vs. RDI¹¹³). Replacing in Eq. (13), the SNR is rewritten as

$$\text{SNR} = \frac{\eta_p \cdot \varepsilon \cdot N_s \cdot \delta t}{\sqrt{2 \cdot (\eta_p \cdot \varepsilon \cdot N_s + \eta_p \cdot RC \cdot N_s + r_{sz} + r_{xz}) \cdot \delta t + (\eta_p \sigma_{\Delta C} \cdot N_s \cdot \delta t)^2}}. \quad (14)$$

For a given astrophysical scene and non-coronagraphic end-to-end optical transmission, the achievable SNR hence depends on three starlight suppression parameters computed in the photometric aperture centered at the planet location: the off-axis core throughput (η_p), the RC, and the post calibrated contrast $\sigma_{\Delta C}$. For exposure times long enough that systematic effects prevail over photon noise, the achievable SNR approaches an asymptotic maximum value given as

$$\text{SNR} = \frac{\varepsilon}{f_{\Delta I} \cdot RC} = \frac{\varepsilon}{\sigma_{\Delta C}}. \quad (15)$$

This means that the minimum planet-to-star flux ratio ε detectable with an SNR of 1 after calibration of the raw images using differential imaging post-processing techniques is equal to $\sigma_{\Delta C}$ (Sec. 3.2). Whatever the exposure time, if the characterization SNR threshold is set to a value SNR_0 , no planet can be successfully observed with a planet-to-star flux ratio lower than

$$\varepsilon_{\min} = \text{SNR}_0 \cdot f_{\Delta I} \cdot RC = \text{SNR}_0 \cdot \sigma_{\Delta C}. \quad (16)$$

As an illustration, to spectrally characterize a planet with a flux ratio of $\varepsilon = 10^{-10}$ with an SNR of 10, the post-calibrated contrast rms $\sigma_{\Delta C}$ must be lower than 10^{-11} . For all planets brighter than ε_{\min} , the exposure time required to reach the threshold SNR is

$$t_{\text{SNR}_0} = \frac{2 \cdot \text{SNR}_0^2 \cdot (\eta_p \cdot \varepsilon \cdot N_s + \eta_p \cdot RC \cdot N_s + r_{sz} + r_{xz})}{(\eta_p \cdot \varepsilon \cdot N_s)^2 - \text{SNR}_0^2 \cdot (\eta_p \cdot \sigma_{\Delta C} \cdot N_s)^2}. \quad (17)$$

Although the time to reach a given SNR always decreases with improved RC, all other parameters being fixed, the minimum planet flux ratio accessible [Eq. (16)] is set by the post-calibration contrast rms $\sigma_{\Delta C}$, defined as the product of RC and speckle post-processing efficiency ($f_{\Delta I}$). This means that both terms must be considered and minimized through the instrument design and operations concept. It also provides the opportunity to trade one term for the other, at least within the range of RC values that preserve a reasonably short exposure time (Secs. 4.3 and 4.4).

In summary, we made three strong assumptions in the above SNR calculations. Two are optimistic, and one is conservative. On the one hand, we assumed negligible detector noise, which may be overly optimistic depending on the spectral resolution used, as well as “perfect” subtraction (down to the photon noise limit) of the (exo-)zodi background at the planet location. Both are best-case scenarios. On the other hand, we assumed an ADI-like post-processing approach that doubles the background shot noise and observing time, which is a worst-case scenario.

In the case in which there is no systematic stellar noise ($\sigma_{\Delta C} = 0$) and the exozodi background flux dominates over the residual stellar flux, exoplanet flux, and solar zodi flux, we have

$$t_{\text{SNR}_0} \propto \frac{\text{SNR}_0^2 \cdot r_{\text{XZ}}}{(\eta_p \cdot \epsilon \cdot N_s)^2}. \quad (18)$$

Because the exozodi background is spatially extended, its detected flux does not depend on stellar distance or telescope diameter and scales as

$$r_{\text{XZ}} \propto \frac{\eta_{\text{occ}} \cdot Tr}{R}. \quad (19)$$

Using Eqs. (6) and (9), we get

$$\eta_p \cdot N_s \propto \frac{\eta_{\text{occ}} \cdot \eta_{\text{PSF}} \cdot Tr \cdot D^2}{R \cdot \text{Dist}^2}, \quad (20)$$

where R is the spectral resolution, D is the telescope diameter, and Dist is the stellar distance. Combining the previous three equations, we finally get that the exposure time required to reach a specified SNR_0 scales as

$$t_{\text{SNR}_0} \propto \frac{R \cdot \text{SNR}_0^2 \cdot \text{Dist}^4}{\epsilon^2 \cdot \eta_{\text{occ}} \cdot Tr \cdot \eta_{\text{PSF}}^2 \cdot D^4} = \frac{R \cdot \text{SNR}_0^2 \cdot \text{Dist}^4}{\epsilon^2 \cdot \eta_p \cdot Tr \cdot \eta_{\text{PSF}} \cdot D^4}. \quad (21)$$

This shows that, in the exozodi-limited regime expected for observations of the most distant and time-consuming targets at visible to near-infrared wavelengths, a gain in the off-axis coronagraphic point source PSF concentration (η_{PSF}) reduces exposure time faster than the overall optical transmission of the system ($\eta_{\text{occ}} \cdot Tr$) at the planet location. It also shows that doubling the spectral resolution is less hurtful than doubling the required SNR (at least within the negligible detector noise assumption adopted) and that the exposure time required scales as $(\text{Dist}/D)^4$, as expected for point source observations limited by a spatially extended background source. In addition, the core throughput (η_p) at a given planet separation will also improve as D increases, providing a further reduction in exposure time.

4.3 Fiducial Exo-Earth Detection

The list of nearby solar-type stars amenable to exo-Earth searches and spectral characterization at near-UV to near-infrared wavelengths is reasonably well known from previous studies,^{2,3} and an improved provisional HWO target list was recently compiled by the Exoplanet Exploration Program.¹¹⁴ That preliminary list suggests that about half of the anticipated number of needed HWO targets (47 ‘‘Tier A’’ stars) could be accessed with an IWA of 83 mas, which corresponds to an exo-Earth seen at quadrature around a Sun-like star at 12 pc. Although a significant fraction of potential HWO targets will be farther, we then adopt hereafter as a representative fiducial case a Sun/Earth system at 12 pc with an exozodiacal dust density level three times higher than in the Solar System.¹⁰⁶

To assess the influence of the starlight suppression KPPs (raw contrast, post-calibrated rms contrast, and off-axis core throughput), we computed the exposure times [Eq. (17)] required for the ($\text{SNR} = 7$) broadband ($R = 5$) visible detection of a fiducial 12 pc exo-Earth for five combinations of raw and post-calibrated contrasts and four different core throughput values at the planet location. For the latter, we used four illustrative curves of core throughput versus separation, meant to represent three illustrative coronagraph cases and one starshade case (Fig. 24).

Given that HWO’s entrance pupil and overall design requirements have not yet been specified, optimized coronagraph masks have not been defined for HWO, and their core throughput versus separation is even less. To deal with that uncertainty and assess the relative merit of different options, we adopted three coronagraph types or ‘‘bins’’ that cover the full range of possible core throughput values and maturity levels:

- Type 1 coronagraphs represent systems with the lowest core throughput and highest technical maturity as evidenced by the quantity and quality of available high contrast ($<10^{-9}$) lab data (Figs. 20 and 21). This type includes the CLC, HLC, and SPC coronagraphs that will fly aboard the Roman Space Telescope.¹¹⁵ All suffer from low core transmission due to

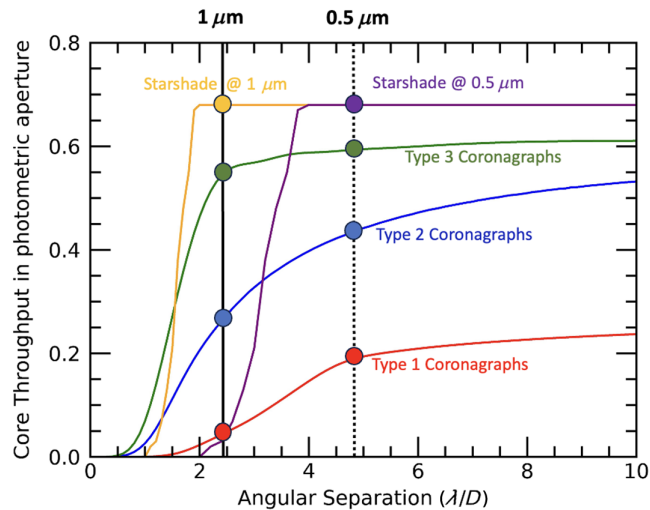


Fig. 24 Core throughput versus separation curves adopted for the four starlight suppression systems considered. We consider three families (types) of coronagraphs and corresponding representative core throughput curves for each. Type 1 (red curve) represents high-maturity coronagraphs with low off-axis core throughput. Type 2 (blue curve) represents medium-maturity coronagraphs with higher off-axis core throughput. Type 3 (green curve) represents the theoretical limit of coronagraphs and corresponds to the lowest maturity/highest core throughput case. For the notional broadband 60 m starshade case considered, the core throughput only depends on physical separation, so the throughput curve versus separation expressed in λ/D units depends on the wavelength (orange: $1 \mu\text{m}$, purple: $0.5 \mu\text{m}$). The vertical black lines indicate the fiducial planet (83 mas) separation in λ/D units at $0.5 \mu\text{m}$ (dotted) and $1 \mu\text{m}$ (plain). The colored circles indicate the planet core throughput values for each starlight suppression system at 0.5 and $1 \mu\text{m}$.

pupil amplitude apodization (SPC), extended Lyot masks, and redistribution of light from the planet PSF's core into the wings because of the masks or due to the wavefront modulation imposed on the DMs as part of the dark hole solution (HLC). Such effects are especially impactful if the aperture is heavily obscured, as in the case of Roman. The core throughput versus separation curve adopted for type 1 coronagraphs is the average of several APLC, SPC, and HLC coronagraph designs proposed³⁶ for the segmented aperture defined by the Ultrastable Observatory Roadmap Team (USORT). Although the HWO aperture will likely differ, the USORT-defined aperture³⁵ can be deemed representative enough for the broad coronagraph types considered here.

- Type 2 coronagraphs represent systems with medium-high core throughput and technical maturity. An example is the family of VVC coronagraphs, which have significantly higher throughput at small separations than type 1 coronagraphs (e.g., Fig. 21) but are intrinsically working on a single polarization. No VVC will fly on Roman, but a VVC4 has been extensively tested in the lab on both clear and segmented apertures with measured contrast levels of a few 10^{-9} (Fig. 20). The core throughput versus separation curve adopted for type 2 coronagraphs is the average of the VVC4 and VVC6 core throughput curves, implicitly assuming that comparable core throughput and contrast performance might be achieved on a segmented aperture with a dual polarization coronagraph.¹¹⁶
- Type 3 coronagraphs represent systems with the highest possible core throughput and the lowest technical maturity. This type includes optimal N 'th order coronagraphs that could theoretically be built using photonic solutions,^{44,117,118} PAPLC designs not yet tested in vacuum, and small-inner-working-angle PIAA (CMC) designs so far limited to $\sim 10^{-8}$ contrast levels in the lab. The core throughput versus separation curve adopted for type 3 coronagraphs is the average of the fourth-, sixth-, and eighth-order optimal designs⁴⁴ (selected to limit sensitivity to stellar finite diameter) created for the USORT segmented aperture.³⁶ In addition to providing smaller IWAs, types 2 and 3 coronagraphs have higher throughput than type 1 for two reasons: (i) they use masks with intrinsically higher transmission close to the optical axis; (ii) their field PSFs are very close in morphology and

photometry to the non-coronagraphic ones, meaning that their PSF core fraction is close to that of a perfect Airy pattern (0.68) except near and within the IWA defined by the FPM.

- In the starshade case, the theoretical core throughput is the highest as no pupil mask, Lyot mask, or DM-induced PSF distortion is necessary to cancel starlight. For point sources located further out than the angle sustained by the starshade petal tip, the core throughput is set by the telescope PSF. The core versus separation curve adopted in the starshade case is computed for a 60 m-diameter starshade^{119,120} flying 95 Mm in front of a 6 m telescope. It is designed to provide a constant 65 mas (tip) IWA between ~ 0.45 and $\sim 1 \mu\text{m}$. Although the laboratory starshade mask testbed reached the deepest starlight cancellation level to date, no starshade has ever flown, and the technical maturity of a large space-based starshade can be considered low.

In addition to the core throughput term, the exposure time calculation includes the end-to-end optical throughput (Tr) of the system, which represents the overall transmission and reflectivity of all other (non-coronagraphic) intervening optics. In line with previous studies,² we assumed $\text{Tr} = 0.3$ for broadband imaging observations with the three coronagraph systems and $\text{Tr} = 0.5$ with the starshade. In the starshade case, because starlight is canceled before entering the telescope, there are no DMs, and no re-imaging optics are required to form additional focal and pupil planes, resulting in a simplified beam train with significantly fewer optical elements. All instrumental and astronomical parameter assumptions used in the exposure time calculations are summarized in Table 3. Note that dual-polarization observations are assumed over a 20% spectral bandwidth and that all detector noise terms are assumed to be negligible compared with the total photon noise from astrophysical background sources. In line with the SNR calculations presented in Sec 4.2, the solar and exozodi backgrounds are assumed to contribute zero mean photon noise, and ADI-like image processing is assumed, doubling the effective background noise.

For the combination of RC and post-calibration rms contrast ($\sigma_{\Delta C}$), we assumed that five different levels of performance could be achieved at the fiducial planet separation ($\sim 4.4 \lambda/D$): ($\text{RC} = 10^{-10}$, $\sigma_{\Delta C} = 0$), ($\text{RC} = 10^{-10}$, $\sigma_{\Delta C} = 5 \times 10^{-12}$), ($\text{RC} = 10^{-9}$, $\sigma_{\Delta C} = 0$), ($\text{RC} = 10^{-9}$, $\sigma_{\Delta C} = 5 \times 10^{-12}$), and ($\text{RC} = 10^{-8}$, $\sigma_{\Delta C} = 0$). We note that, for the detection of a fiducial exo-Earth with a planet-to-star flux ratio of 10^{-10} at an SNR of 7, a post-calibration rms contrast of $\sim 1.4 \times 10^{-11}$ or better is required. This common set of contrast performance values was assumed for all starlight suppression systems, regardless of the contrast levels currently demonstrated in the lab. For comparison (Sec. 2), the best lab contrasts obtained to date at $\sim 4 \lambda/D$ separation range from $\sim 10^{-10}$ with a starshade mask to $\sim 10^{-9}$ for coronagraphs of types 1 and 2 all obtained over $\sim 10\%$ to 12% bandwidth rather than the 20% assumed here, either over a half dark hole or at a single polarization.

The resulting broadband detection times for each starlight suppression type are shown in Fig. 25 for the five contrast performance levels assumed (x -axis). Absolute exposure times strongly depend on the assumptions listed in Table 3 and should be regarded as ballpark estimates. The relative exposure times among the different cases are more informative because they show the relative impact of RC, post-calibration efficiency, and system-specific core throughput.

For a given starlight suppression system, the time to achieve a given SNR increases with RC and shoots up when the contrast goes from 10^{-9} to 10^{-8} . That is even in the case of perfect speckle calibration ($\sigma_{\Delta C} = 0$) and due to prohibitive stellar shot noise. Broadband visible detection remains theoretically possible within a few hours to a day for RC values up to 10^{-9} , as long as speckles can be post-calibrated to a high accuracy (i.e., $\sigma_{\Delta C} = 5 \times 10^{-12}$ or less). However, stabilizing the RC at a given level becomes more difficult as RC worsens. Indeed, for small fluctuations, the contrast degradation caused by a given wavefront drift goes as the square root of the initial RC level,⁴² analogous to the well-known “pinned speckles” effect.

At a fixed combination of RC and post-calibration efficiency (fixed x -coordinate), the exposure time required depends only on the core off-axis throughput of each system. A significant penalty is observed for type 1 coronagraphs compared with other systems, and the starshade system provides the shortest exposure time. However, a 60 m starshade flying 95 Mm away will take weeks to travel between targets, and this time-to-SNR advantage will only hold if general astrophysics observations are conducted in the meantime. If exo-Earth blind searches and orbit determinations are required via broadband observations and revisits of many target stars,

Table 3 List of astrophysical and instrument parameters used for broadband Earth twin detection. Parameters being varied among the different starlight suppression systems are listed in bold. C1, C2, or C3 designates coronagraphs of types 1, 2, or 3, respectively.

Parameter	Value
Stellar type	Solar twin
Stellar distance	12 pc
Planet type	Earth twin
Planet orbit semi-major axis	1 AU
Planet illumination phase and flux ratio	Quadrature; resulting in 10^{-10} planet-to-star flux ratio
Solar zodiacal light surface brightness at planet location	23 mag/arcsec ² at V band
Exozodiacal light surface brightness at planet location	22 mag/arcsec ² at V band for a 1 zodi solar analog
Exozodi level	3 zodis
Core throughput at planet location within the photometric aperture (η_p)	Varies with starlight suppression system (see Fig. 24)
Core PSF fraction at planet location (η_{PSF})	0.4 for C1; 0.65 for C2 and C3; 0.68 for starshade
(Exo)-zodi throughput at planet location [Eq. (8)] $\eta_{\text{sz}} = \eta_{\text{xz}} \sim \eta_{\text{occ}} = \eta_p / \eta_{\text{PSF}}$	$\eta_p / \eta_{\text{PSF}}$
Telescope (inscribed) diameter	6 m
Central obscuration	None
Central wavelength	0.55 μm
Spectral resolution	5 (i.e., 20% bandwidth)
End-to-end optical throughput (Tr), excluding all starlight suppression masks and detector quantum efficiency	0.3 for all coronagraphs 0.5 for starshade
Radius of photometric aperture	$0.7 \lambda / D^a$ (centered at planet location)
RC at planet location (RC)	10^{-10}, 10^{-9} or 10^{-8}
Post calibration contrast spatial rms at the planet location ($\sigma_{\Delta C}$)	0 or 5×10^{-12}
Detector quantum efficiency (QE)	0.9
Detector noise	Negligible
Number of polarizations	2
Detection signal-to-noise	7

^aAlthough it does not necessarily provide the optimum SNR for all starlight suppression systems and planet locations, we assume hereafter that the photometric aperture radius is $0.7 \lambda / D$. This value maximizes SNR in the case of a point source Airy pattern to be detected against a spatially uniform background.

coronagraphic observations may still be preferred overall for broadband blind detection and orbit determination of exo-Earths.

4.4 Fiducial Case for Exo-Earth Spectroscopy

For spectroscopic observations, we adopt the additional instrument and astrophysical assumptions listed in Table 4. We defined four 20% spectral bands covering the full 0.45 to 1 μm spectral




Fig. 25 Exposure times required to detect a fiducial exo-Earth at 12 pc with an SNR of 7 over a 20% bandwidth centered at $0.55 \mu\text{m}$, if various levels of raw/post-calibration contrast performance (X-axis) could be reached by four different starlight suppression systems: coronagraphs of types 1, 2, and 3 (C1, C2, and C3) and starshade. For reference, the approximate RC performance level currently achieved in the lab by each system type over the 10% to 20% bandwidth is indicated by a filled circle (assuming perfect post-calibration of raw starlight speckles, i.e., $\sigma_{\Delta C} = 0$). If all systems were able to achieve the same contrast performance, systems with the highest exoplanet transmission require the shortest exposures. The assumed exozodi level is three zodis in all cases.

region and computed the exposure time required to reach an SNR of 10 per $R = 70$ spectral bin at the central wavelength of each band, assuming a continuum planet flux. We assumed that an integrated field spectrograph is used, resulting in lower end-to-end throughput than for broadband detection. In line with the previous analysis,² we adopted an end-to-end coronagraphic IFS throughput of 0.20 and a starshade IFS throughput of 0.33, both constant with wavelength. The detector noise is still assumed to be negligible; given the lower astrophysical flux per spectral bin and the larger number of pixels used in an IFS, this is a significantly more driving assumption than in the broadband imaging case and one that will impact spectrometer selection and resolution.

4.4.1 Exposure times required for O_2 search

Exposure times required for spectroscopy across band 3 were computed in the middle of the band, at $0.747 \mu\text{m}$, in the continuum near the critically important O_2 A band at $0.76 \mu\text{m}$. The same five combinations of instrumental RC and post-processing performance are assumed as in Sec. 4.3, but with reduced end-to-end throughputs for IFS spectroscopy, as defined in Table 4.

As in the case of broadband detection, for the same contrast performance, the system with the highest exoplanet transmission requires the shortest exposures. In past studies,^{2,3} the exposure time limit for exo-Earth spectroscopy was set to 60 days. Although somewhat arbitrary, that threshold is probably in the right ballpark given that changes in the apparent planet separation, planet illumination phase, and corresponding spectrum may already be significant over such a timescale, as well as possible changes in exozodi dust brightness due to resonant dust structures in nearby orbit or other time variable astrophysical phenomena. In addition, multiplying that 60-day upper limit by the number of exo-Earths spectrally characterized would already result in several years of observing time. Adopting this notional cutoff limit here for simplicity, spectroscopic ($R = 70$) measurements of a fiducial exo-Earth at 12 pc at an SNR of 10 per spectral bin around $0.75 \mu\text{m}$ (Fig. 26) would be within the capabilities of a starshade (in 1 to 2 weeks of exposure), as well as type 2 (5 to 9 weeks) and type 3 coronagraphs (2 to 4 weeks), that is, as long as they can operate with an RC of 10^{-9} or better and with a residual post-calibrated contrast floor ($\sigma_{\Delta C}$) of 5×10^{-12} or better, all at a separation of $\sim 3 \lambda/D$ (83 mas). For the highest maturity (type 1) coronagraphs, the core throughput is too low to enable reasonable exposure

Table 4 List of astrophysical and instrument parameters used for Earth twin visible spectroscopy (only parameters differing from Table 3 values are listed).

Parameter	Value
Planet type	Earth Twin with constant albedo (=0.2)
Solar zodiacal light surface brightness at planet location	23 mag/arcsec ² at V band. Gray opacity assumed at other wavelengths
Exozodiacal light surface brightness at planet location	22 mag/arcsec ² for a 1 zodi solar analog at V band. Gray opacity assumed at other wavelengths
Exozodi level	3 zodis
Spectral bands	Band 1: 0.45 to 0.55 μm Band 2: 0.55 to 0.672 μm Band 3: 0.672 to 0.822 μm Band 4: 0.822 to 1.0 μm
Spectral resolution	70
End-to-end optical throughput (Tr), excluding all starlight suppression masks and detector quantum efficiency	0.20 for IFS-based coronagraphic observations 0.33 for IFS-based starshade observations
Raw contrast at planet location (RC)	10^{-10} , 10^{-9} or 10^{-8}
Post calibration contrast spatial rms at the planet location ($\sigma_{\Delta C}$)	0 or 5×10^{-12}
Signal-to-noise per spectral bin	10 computed at each band central wavelength, assuming continuum flux

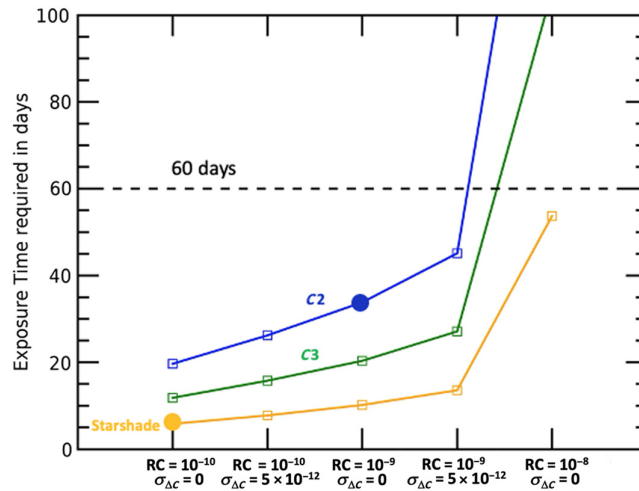


Fig. 26 Exposure time required for $R = 70$ spectroscopy of a fiducial exo-Earth at 12 pc with an SNR of 10 per spectral bin, computed at $0.747 \mu\text{m}$, for various levels of raw/post-calibration contrast performance (X-axis) achieved by different starlight suppression systems. For reference, the approximate raw-contrast performance level currently achieved in the lab by each system is indicated by a filled circle (assuming perfect post-calibration of raw starlight speckles, i.e., $\sigma_{\Delta C} = 0$). If all systems were able to achieve the same contrast performance, the starshade and type 3 coronagraph (C3) systems would require the shortest exposures due to higher planet throughput. The assumed exozodi level is three zodis in all cases. Exposure times for type 1 coronagraphs (not shown) are above 200 days for all contrast performance cases considered. The horizontal dashed line indicates a notional exposure time threshold of 60 days.

times for spectroscopy of the fiducial target, regardless of the contrast performance achieved (exposure times $> \sim 200$ days and off the chart). Spectroscopic observations at 10^{-8} RC become impractically long for all coronagraphs but may still be feasible for a starshade, assuming perfect calibration of residual starlight.

4.4.2 Exposure times required for spectroscopy across the full 0.45 to 1 μm range

We now compute the exposure times required for spectroscopy as a function of visible wavelength, assuming two different RC levels at all wavelengths: 10^{-10} or 10^{-9} . As above, the required exposure times are calculated for $\text{SNR} = 10$ spectroscopy per $R = 70$ spectral bin, but now at the central wavelengths of four 20% spectral bands covering the full 0.45 to 1 μm range (Table 4).

The change of exposure time with wavelength differs among the starlight suppression systems considered and depends on the RC performance (Fig. 27). For coronagraphs, the exposure time generally increases with visible wavelength because of two compounding effects: (i) the PSF core region area increases as λ^2 , capturing more of the spatially extended (exo-)zodi background; (ii) the planet apparent separation decreases in λ/D units, and the core throughput decreases accordingly (Fig. 24). Both effects are more marked at deeper contrast, where (exo-)zodi background dominates, and for lower core throughput and PSF core fraction. This means that the required exposure times increase faster with the wavelength for the 10^{-10} contrast than 10^{-9} and for type 2 coronagraphs than type 3. For the starshade case, the core throughput is by definition (Fig. 24) the same at all wavelengths in range, and only the effect of the increased (exo-)zodi beam etendue can be seen. It results in a modest increase of exposure time with wavelength in the 10^{-10} contrast case and almost a flat spectral response at 10^{-9} contrast.

Interestingly, as shown in Fig. 27, exposure times can be shorter at worse RC but higher throughput (e.g., a starshade operating at the 10^{-9} contrast versus a type 3 coronagraph at the 10^{-10} contrast, or a type 3 coronagraph operating at the 10^{-9} contrast versus a type 2 coronagraph at the 10^{-10} contrast, especially for the redder wavelengths). This observation illustrates the starlight suppression system trade that exists between core throughput and RC, as long as RC can be calibrated down to a constant level. For all starlight suppression systems considered, we further find that the spectroscopic ($R = 70$) exposure times increase sharply at longer near-infrared

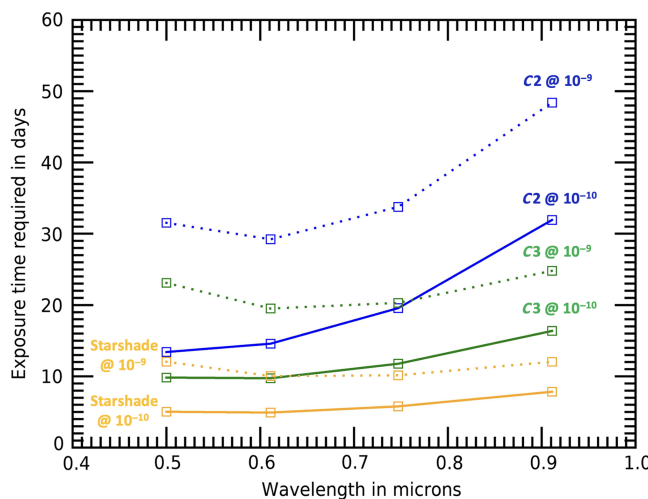


Fig. 27 Exposure times required for $R = 70$, $\text{SNR} = 10$ spectroscopy of a fiducial exo-Earth at 12 pc as a function of optical wavelength. For each starlight suppression system considered (starshade and coronagraphs of types 2 and 3), the RC level at the planet location is either set to 10^{-10} (plain curves) or 10^{-9} (dotted curves). All cases assume perfect speckle subtraction ($\sigma_{\Delta C} = 0$). Exposure times for type 1 coronagraphs (not shown) are around 100 days at 0.5 μm and increase sharply with wavelength.

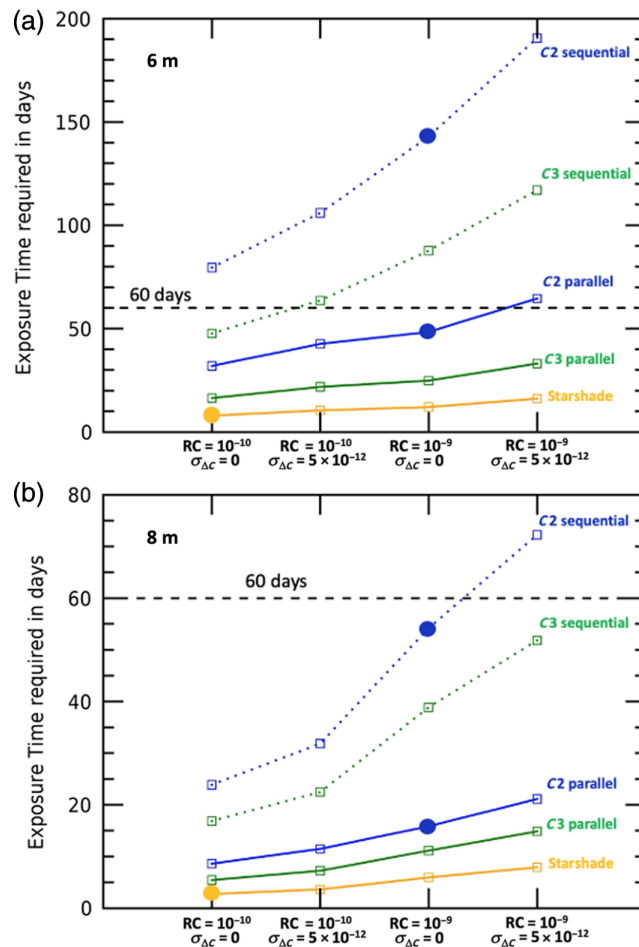


Fig. 28 Total exposure time (in days) required to measure the $R = 70$ SNR = 10 spectrum of a fiducial 12 pc exo-Earth from 0.45 to $1 \mu\text{m}$, with a notional 60-day threshold (horizontal dashed line). Five starlight suppression cases are considered: (i) an ultra-broadband (60 m) starshade system covering the full wavelength range at once; (ii) four medium throughput/medium maturity coronagraph channels (type C2) working in parallel over 20% bandwidth each; (iii) a single (type C2) coronagraph channel observing sequentially over four different 20% spectral bands, one at a time; (iv) four extremely high-throughput/low-maturity coronagraph channels (type C3) working in parallel over 20% bandwidth each; and (v) a single (type C3) coronagraph channel observing sequentially over four different 20% spectral bands, one at a time. For reference, the approximate RC performance currently achieved in the lab by each system is indicated by a filled circle (assuming perfect post-calibration of raw starlight speckles, i.e., $\sigma_{\Delta C} = 0$). (a) $D = 6$ m inscribed diameter telescope. (b) (Note the different y scales): $D = 8$ m inscribed diameter telescope, using a starshade with the same physical IWA (65 mas) as in the 6 m telescope case (i.e., a ~ 62 m starshade) but coronagraphs with IWAs improving as $1/D$.

wavelengths, reaching 100 days or more at $1.5 \mu\text{m}$ under the combined effects of lower planet core throughput and higher exozodi background per PSF core.

Next, we investigate the time required by each system to obtain a spectrum over the full 0.45 to $1 \mu\text{m}$ range, either at once, or sequentially, one spectral band at a time (Fig. 28), still using the astronomical and instrumental assumptions listed in Table 4. The starshade can in principle be designed for parallel spectroscopic observations over the full 0.45 to $1 \mu\text{m}$ wavelength range, without a change in operating conditions. Such “full spectral parallelization” is also possible for coronagraphs but would require several coronagraph spectral channels to run in parallel over smaller individual spectral band-passes using beam-splitters or dichroics with spectral phase and amplitude effects that are controllable. Each channel would have its own pair of DMs, wavelength-dependent DM actuator settings, and coronagraphic masks. Assuming that

dual-polarization coronagraphic measurements over a 20% bandwidth are achievable, four parallel coronagraphs would be needed to cover the 0.45 to 1 μm range simultaneously.

Looking at the fiducial 12 pc exo-Earth case, an $R = 70$ spectrum with an SNR of 10 per spectral bin could be obtained from 0.45 to 1 μm in 1 to 2 weeks by a starshade system, depending on the assumed contrast performance [Fig. 28(a), orange curve]. It would typically take $\sim 4\times$ longer, i.e., 4 to 8 weeks, depending on the contrast performance, for a fully spectrally multiplexed type 2 coronagraph system (blue curve), and $\sim 2\times$ longer, i.e., 2 to 4 weeks with a theoretically optimum coronagraph (type 3, green curve), also composed of four parallel 20% bandwidth channels. Still considering the 6 m telescope case and the fiducial exo-Earth, the total exposure time would be prohibitive (70 to 190 days) if spectra are instead obtained sequentially, 20% bandwidth at a time with a type 2 coronagraph (dotted blue curve), and still prohibitive (50 to 120 days) with type 3 coronagraphs under most contrast performance scenarios (dotted green curve). Thus, for internal coronagraphs, the throughput and instantaneous bandwidth, per channel or synthesized via parallel channels, accessible to observations are essential factors in reigning in spectroscopic integration times to acceptable levels. To get full 0.45 to 1 μm spectra of typical exo-Earths at the specified SNR (10) and spectral resolution ($R = 70$), and if coronagraph instruments can only access 20% bandwidth at a time, some degree of spectral “multiplexing” will be required, observing in parallel with several coronagraphic channels. An important trade for conducting spectroscopic measurements with coronagraphs is among the bandwidth of individual channels, the core throughput accessible per channel, and the number of channels desired for spectroscopic characterization, setting the degree of spectral parallelization and instrument complexity needed.

In the fiducial case adopted, parallel observations using a lower throughput coronagraph (e.g., C2 solid blue curve) can still be conducted faster than sequential observations with a higher throughput coronagraph (e.g., C3, dotted green curve). We also note that, as the contrast performance gets worse, parallel observations become increasingly favored over sequential ones. It can be understood as the following: as the contrast gets, e.g., from 10^{-10} to 10^{-9} , observations become less limited by the (exo-)zodi background and more limited by starlight residuals. As a result, the required observing times are more uniform across spectral channels (Fig. 27), which is when parallel observations become most advantageous over sequential ones (Fig. 28).

An additional knob to turn would obviously be to use a larger telescope, as illustrated by the 8 m (inscribed diameter) case shown in the Fig. 28(b). The relative exposure times as a function of the starlight suppression system throughput and level of spectral parallelization remain about the same as in the 6 m case. But, all exposure times go down by a factor of ~ 3 , as expected for observations limited by shot noise from the (exo-)zodi background (exposure time scaling as $\sim 1/D^4$).

One last trade illustrated in Fig. 28 (6 m panel, solid curves) is between RC and post-calibrated contrast. For instance, for the fully parallelized types 2 and 3 coronagraphs, the exposure time required assuming 10^{-9} RC with perfect speckle subtraction is close to that required at 10^{-10} RC but with a worse (5×10^{-12}) post-calibrated contrast. In other words, RC might be traded against speckle calibration to some degree, highlighting the importance of data processing and starlight residual removal.

5 Design Trades, Numerical Parametric Studies, and Future Experiments

HWO’s detailed science objectives and corresponding technical requirements remain to be defined. For direct exoplanet spectro-imaging, the basic observational parameters and quantitative figures of merit-driving mission design are not yet specified. This includes the number of objects to be detected and characterized for different planet types, as well as basic measurement parameters such as wavelength coverage, spectral resolution, and SNR. It is only after such high-level science figures of merit, physical parameters, and measurement requirements are established that realistic mission science yield simulations can be run to optimize the observations scheduling across the pool of available target stars^{38,40} and thoroughly assess the impact of individual performance parameters.

Pending such detailed mission yield simulations, we adopted in this paper a single representative exo-Earth target at 12 pc and used the observing time required for its detection and spectral characterization at visible wavelengths (loosely defined as 0.45 to 1 μm) as a simple figure of merit. The exposure time required for observations of a typical target certainly does not capture the complexity of science yield calculations over the full target sample. But, it still constitutes a useful metric to assess the feasibility of scientific observations and provides valuable insight into the performance requirements of the starlight suppression system, as well as on some of the major trades remaining to be explored, numerically or experimentally.

5.1 Numerical Science Yield Parametric Studies

Using realistic astronomical and instrument parameters (Tables 2 and 3), we find that the time required for spectroscopy of the adopted fiducial exo-Earth far exceeds that needed for broad-band detection and that spectral characterization drives the performance of the overall starlight suppression system. For a given spectral bandwidth, the required exposure times are set by three starlight suppression system KPPs: RC, post-calibrated contrast, and off-axis (planet) core throughput. The three parameters can be traded against each other to some degree, as, e.g., higher off-axis throughput can make up for worse RC, at least if speckles can be calibrated close to the photon noise level. Once science objectives are established, we suggest that a detailed multi-parameter trade study of science yield versus KPP should be conducted to map the combination of RC, post-calibrated contrast, and core throughput values that meet the science needs. This exploration can be done before implementation choices are made (e.g., trading coronagraphs with different throughputs and sensitivities to aberrations), providing a few possible winning options that can then be compared with existing lab performance so that technical gaps and further lab work can be properly identified and prioritized.

Another important trade highlighted in our spectroscopic exposure time calculations concerns the degree of spectral parallelization, i.e., the number of individual spectral channels needed to cover the desired wavelength range. Although a starshade can in principle be used for spectroscopy over the full 0.45 to 1 μm at the same time, internal coronagraphs, if limited to 20% bandwidth per individual channel, will require multiple parallel channels to cover the same range within a reasonable time. Covering the full visible range at once could, for instance, require four parallel coronagraph channels, and twice that many if individual channels can only measure a single polarization. For internal coronagraphs, the throughput and instantaneous spectral bandwidth accessible per channel or synthesized via parallel channels accessible to observations are hence essential factors to reigning in spectroscopic integration times to acceptable levels. A parametric study of the number of parallel coronagraph channels necessary to meet the full mission exoplanet spectroscopy yield objectives or sets of objectives reflecting different levels of ambition (e.g., spectral coverage^{4,5,121–123}) would then be most useful to drive requirements and technical solutions. Broader instantaneous coronagraph bandwidth may for instance be reachable at the expense of reducing the core throughput or dark-hole field of view.

Pending such detailed parameter studies, we find (Fig. 28) that full visible-range spectroscopy of an exo-Earth orbiting a Sun-like star at the typical distance of HWO targets (taken to be 12 pc) with a 6 m telescope is possible for starlight suppression systems providing an RC of 10^{-9} or better with a residual post-calibrated contrast floor ($\sigma_{\Delta C}$) of 5×10^{-12} or better and a core throughput greater than $\sim 30\%$, all at a separation of $< \sim 3\lambda/D$.

5.2 Technical Demonstrations: Coronagraphs

In comparison, the best coronagraphic lab results so far (Sec. 2, Fig. 20) have reached the required level of RC at 3 to $4\lambda/D$ separation but only over a 10% spectral bandwidth (HCIT CLC & HLC lab setups^{45–48}), with $< 15\%$ off-axis throughput and only on a monolithic unobscured aperture. The VVC4 10% bandwidth setup demonstrated RC values a few times worse at significantly higher core throughput (Fig. 21) but still with a monolithic aperture and using a single polarization.⁵⁷

For segmented off-axis apertures, the best polychromatic performance currently achieved⁵⁸ is $\sim 10\times$ worse, with $\sim 10^{-8}$ RC demonstrated at $3\lambda/D$. However, among the coronagraph types providing the best clear aperture results, only one (VVC4) was also tested in vacuum in that off-axis segmented configuration and using a single DM instead of two. This makes it currently

difficult to isolate the effect of segmentation on coronagraphic performance. In addition, we caution that the current vacuum “segmented” experiments used a segmented aperture mask in which only the amplitude shows discontinuities across the pupil, not the phase. Such experiments hence demonstrate the contrast limit accessible with perfectly co-phased segments but do not realistically probe the contrast sensitivity to individual segments tip-tilt and piston errors. Given that HWO is currently baselining a segmented telescope, this highlights the pressing need for more vacuum lab demonstrations of high-contrast coronagraphic imaging and spectroscopy on segmented apertures, especially with entrance pupils that exhibit both phase and amplitude discontinuities. Nonetheless, experiments on clear monolithic apertures should also still be pursued because they provide an adequate benchmark for assessing the intrinsic impact of segmentation. To provide the most direct comparison between the monolithic and segmented cases, it is desirable that the same setup be used, limiting changes to what is strictly required to deal with segmentation (e.g., mask adjustments but same operating conditions). Further vacuum tests conducted on monolithic apertures will also help to better investigate the systematics and current contrast limits set by the testbeds themselves (Sec. 2.2.1), as well as by manufactured mask residual defects. In most cases, it is inferred that mask defects are responsible for a significant part of the contrast performance degradation observed as spectral bandwidth increases. This highlights the need for improved mask manufacturing, as well as detailed mask inspection and characterization before coronagraphic testing. When incorporated in WF control algorithms, the precise knowledge of residual defects in manufactured masks may indeed result in faster dark hole convergence and improved contrast performance.

An additional degradation of broadband contrast performance is currently observed in the laboratory when switching from off-axis to on-axis segmented apertures, for which the best RC achieved⁷⁴ at $3\lambda/D$ is $\sim 7 \times 10^{-8}$, $\sim 100\times$ worse than in the off-axis monolithic case. However, that degradation is far larger than predicted by current models and again based on vacuum testing of a single coronagraph (PIAACMC). This calls for further laboratory testing in the next few years so that any firm conclusion on the gap currently observed in the lab between off- and on-axis coronagraphic results can be reached in a timely manner.

Given the extreme level of contrast (less than $\sim 10^{-9}$) required for direct exo-Earth spectroscopic observations, laboratory experiments have mainly concentrated so far on demonstrating deep broadband ($> \sim 10\%$) RC at small angular separations. However, once the contrast performance is good enough that residual starlight levels fall below the signal from irreducible astrophysical background sources (zodi and exozodi dust), the exposure time required to characterize exo-Earth planets becomes essentially independent of RC (at least if post-calibration of residual speckles can be achieved down to rms levels lower than the (exo-)zodi background shot noise). Consequently, the predicted total mission yield of exo-Earths spectrally characterized increases only marginally with further RC improvements.³⁹ Depending on the stellar distance and exact visible wavelength, that transition occurs at a threshold RC level ranging from $\sim 10^{-10}$ to $\sim 10^{-9}$ (Sec. 4.1). Once that RC threshold performance is reached, the ability to spectrally characterize exo-Earths primarily depends on the next two starlight suppression system KPPs: off-axis core throughput, especially at small angular separations, and post-calibrated contrast achieved after image processing. Both parameters have been fairly overlooked so far and deserve a lot more attention going forward.

As far as throughput is concerned, accessing high core throughput at separation as close as 2 to $3\lambda/D$ will be mandatory for successful spectroscopic observations of exo-Earths with HWO. There are three throughput components to maximize: the overall end-to-end optical transmission of the telescope + instrument system, excluding all starlight suppression optics; the cumulative transmission of all occulting masks; and the fraction of planet light captured in the off-axis PSF core. The first term applies equally to all sources in the field, whereas the latter two are field-dependent. Because observations of the most distant and time-consuming targets are expected to be limited by the spatially extended (exo-)zodi backgrounds, a special emphasis should be put on maximizing the PSF core fraction at the planet location rather than the other two terms (Sec. 4.2, Eq. 21). A trend currently observed is that the RC level demonstrated by coronagraph setups appears to degrade as their throughput at small separations increases (Figs. 20 and 21). Because coronagraphs with higher throughput at close-in separations also tend to exhibit higher sensitivity to low-order aberration drifts,^{36,37} this would be expected for dynamic contrast

performance but not in the lab under essentially static conditions. Whether this unpredicted trend is real or due to a small number of statistics remains to be explored.

The least experimentally constrained performance parameter, at least for contrast levels below 10^{-7} , is the post-calibrated contrast rms error ($\sigma_{\Delta C}$), achieved after estimation and subtraction of starlight speckles using differential imaging techniques. However, this parameter has a crucial impact on the feasibility of exo-Earth spectroscopic observations, which will require $\sigma_{\Delta C}$ values or $\sim 10^{-11}$ or lower. In addition, if post-calibration speckle residuals can be kept below the (exo)-zodi background shot noise level, the RC performance could be relaxed to the threshold value at which (exo)-zodi background dominates. Given the existing trade between RC and off-axis core throughput, this would open the possibility to operate at higher core throughput values than commonly thought, a very attractive prospect.

An even more crucial open question is whether the information provided contemporaneously with the science images (rather than through conventional asynchronous ADI/RDI), using either WFS data or data recorded outside the dark hole (spatially or spectrally), could yield precise estimates of starlight residuals and their uncorrected temporal fluctuations. The latter would represent a real paradigm shift, from a wavefront stability demand to a wavefront knowledge demand. Wavefront stability would still be required to maintain the average RC below the background-limited threshold level, but no longer within a fraction of the planetary signal. As in the RC case, tolerating higher levels of wavefront and RC fluctuations could have a tremendous impact on fundamental design trades such as telescope stability versus coronagraph resilience to aberrations (robustness) and coronagraph robustness versus core throughput. Diving into such trades is beyond the scope of this paper but is the object of past^{36,37} and on-going community work in the context of HWO's technology maturation. Given the importance of the post-calibrated contrast parameter for the design of HWO's starlight suppression system, substantial efforts should be spent to minimize it and estimate it, e.g., through dedicated coronagraphic lab experiments or using science + telemetry data from space coronagraph instruments on-board JWST and ultimately on Roman.⁹²

The ability to accurately subtract residual speckles may also depend heavily on the concepts of operations and observing scenarios. Exploring the overall system trade among observatory stability, the coronagraph WFS/C performance, and image processing contrast-improvement capabilities is a very high priority to allow for the definition of HWO's coronagraph requirements for the detection and spectroscopy of exoplanets, in particular for the driving science case of exo-Earth spectroscopic observations. The observatory and its coronagraph instruments (including their WFS/C system, operations concept, and post-processing strategy) should be jointly optimized as a single system, so the requirement allocations may be properly split between the telescope and the coronagraph systems. Making the necessary design trades will require a detailed error budget, informed by end-to-end integrated modeling of all of the different sub-system components as well as model validations of coronagraph testbed results and post-processing improvements, reaching all the way to projected multi-variate exoplanet science yields and molecular abundance spectral retrievals.

5.3 Technical Demonstrations: Starshades

The current baseline for exoplanet direct broadband detection with HWO is a coronagraphic system capable of conducting agile blind searches and orbital characterization of exoplanets via repeated observations at low spectral resolution ($R \sim 5$). Unless it can be refueled through in-orbit servicing, a starshade will be limited in its number of target slews and not so well suited to a blind exo-Earth search phase.¹²⁴ However, a starshade holds many promises for the spectral characterization of exoplanets and exo-Earths previously identified by HWO's coronagraph or other precursor observations: this includes deep contrast at small separations over large ($\sim 100\%$) instantaneous spectral bandwidths at high off-axis (planet) core throughput. Also, because it blocks starlight before entering the telescope, a starshade system works equally well with off-axis, on-axis, monolithic, or segmented apertures. As long as the telescope remains diffraction limited, no sub-nanometer wavefront stabilization system, additional starlight suppression masks, re-imaging optics, or DMs are required in the beam train. As a result, the overall system throughput is intrinsically high, the native telescope PSF is retained for off-axis sources, and the dark hole angular size is only limited by the detector size rather than by the number of

DM actuators. Some of these promises have been confirmed in the Princeton starshade lab, with a small-scale starshade mask experiment reaching 10^{-10} contrast over a 12% bandwidth down to a separation of only $2\lambda/D$,³² with starlight residuals falling off quickly at wider separations.

These impressive lab results demonstrate a level of starlight suppression commensurate with the detection and spectroscopy of exo-Earths. They were obtained with a much smaller starshade mask ($\sim 1/1000$ th scale) than required for a 6 m telescope. However, they were conducted at a flight-like Fresnel number, hence enabling testing and validation of vectorial finite diffraction theory in the appropriate physical regime. (In addition, thick screen effects and small-scale petal defects have a worse effect on contrast for a smaller starshade.) In fact, the challenges of operating a starshade in conjunction with a large space telescope lie elsewhere. Indeed, in the starshade case, the wavefront control and starlight suppression challenges are essentially shifted from the telescope and internal WFS/C systems to deploying a large external occulter with high shape accuracy and stability, flying in precise formation with the telescope to block the starlight and to mitigating scattered sunlight. NASA's starshade technology development (S5) activity^{125,126} has closed several milestones in these areas and in particular has closed two major technology gaps for formation flying sensing and starlight suppression/model validation.³³ Milestones pertaining to the remaining technology gaps, i.e., deployment accuracy, shape accuracy, and stability of prototype starshade petals of representative length (8 m), are expected to be closed by the end of FY24 and will concentrate on extending the currently demonstrated mechanical performance to higher-fidelity components.

Given the current lab results, as well as the potential for reaching deep broadband contrast at small separations and high throughput, it appears appropriate to keep the starshade approach in HWO's toolbox. When working in tandem with the HWO coronagraphic system, a starshade could significantly enhance spectroscopic performance in the visible and extend it to the near infrared for many targets, due to its small IWA. It may also offer a unique solution for photon-efficient exoplanet observations with HWO in the near UV, should the science community recommend them. A starshade could be launched with the HWO prime mission or a few years later after the coronagraph has identified which stars have exo-Earth candidates orbiting in their habitable zones.

If such a dual coronagraph plus starshade option were to be pursued, further starshade technology maturation would be required to test for manufacturing accuracy and thermal stability after full-scale petal development of an HWO-compatible starshade, which could be of order 60 m in diameter at visible/near-infrared wavelengths or ~ 35 m diameter for near-UV observations. In addition, the starshade approach cannot be fully tested at scale from the ground, and there is no technology demonstration mission such as the Roman coronagraph currently planned in the coming years. This raises the questions of whether a starshade space demonstration is opportune and necessary in the HWO context and what its minimum size should be to remain technically relevant.

5.4 Summary

The exo-Earth spectral characterization objectives of HWO will drive HWO's design and remain to be defined. This includes the range of insolation levels and sizes to be adopted for exo-Earths, the exact number to be characterized (or searched for at some completeness level), the wavelength range to be accessed, the SNR and spectral resolution desired, and the distribution of such observing parameters across the sample. As an illustrative scenario, we examined here the case of an Earth twin orbiting a Sun-like star at 12 pc with three times the level of solar zodiacal dust (our current best estimate) and required that a full spectrum extending from 0.45 to 1 μm be obtained at an SNR of 10 and spectral resolution of 70.

Assuming such broad spectral characterization is desired, we identified three main options for conducting time-efficient spectroscopy of a significant number of exo-Earths at visible wavelengths with HWO: (i) improved broadband coronagraph systems, likely using several parallel spectral channels, that provide deep raw ($<10^{-9}$) and post-calibrated ($<5 \times 10^{-12}$) contrasts together with high off-axis (exoplanet) throughput at $< 3\lambda/D$ separation; (ii) use of a starshade for photon-efficient enhanced broadband spectroscopy of exo-Earths and other planets previously detected by HWO's coronagraph or by other precursor indirect observations (e.g., extreme precision radial velocimetry); and (iii) use of a larger (e.g., 8 m) telescope to significantly reduce

exposure time in the (exo-)zodi limited regime expected for many targets, especially the most distant and time-consuming ones. Any of these options would strongly benefit HWO's exo-Earths spectral characterization capabilities. They could also be exercised in conjunction for optimum performance.

Finally, extending HWO's exo-Earth spectroscopic measurements to the near UV, where planets are intrinsically faint in reflected light, or to the near infrared, where planets are less resolved from their parent star and the irreducible (exo-)zodi background signals are higher than at visible wavelengths, would make these options even more indispensable.

Our findings may provide useful input to the design trade space exploration to be conducted by the newly formed HWO technology maturation project office.

Code and Data Availability

NASA regulations govern the release of source code, including what can be released and how it is made available. Readers should contact the corresponding author if they would like copies of the software or data produced for this study.

Acknowledgments

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

References

1. National Academies of Sciences, Engineering, and Medicine, 2023, *Pathways to Discovery in Astronomy and Astrophysics for the 2020s*, The National Academies Press, Washington, DC (2021).
2. B. Gaudi et al., "The Habitable Exoplanet Observatory (HabEx) mission concept study final report," arXiv:2001.06683 (2020).
3. The LUVOIR Team, "The LUVOIR mission concept study final report," arXiv:1912.06219 (2019).
4. M. Damiano and R. Hu, "Reflected spectroscopy of small exoplanets. II. Characterization of terrestrial exoplanets," *Astron. J.* **163**, 299 (2022).
5. M. Damiano, R. Hu, and B. Mennesson, "Reflected spectroscopy of small exoplanets. III. Probing the UV band to measure biosignature gases," *Astron. J.* **166**, 157 (2023).
6. M. Levine et al., "Terrestrial Planet Finder Coronagraph (TPF-C) flight baseline concept," arXiv:0911.3200 (2009).
7. K. Stapelfeldt et al., "Exo-C: exoplanet direct imaging coronagraph probe mission study," 2015, <https://exoplanets.nasa.gov/exep/studies/probe-scale-stdt/>.
8. S. Seager et al., "Exo-S: starshade probe-class exoplanet direct imaging mission concept," 2015, https://exoplanets.nasa.gov/internal_resources/788/.
9. J. Mazoyer, "High-contrast testbeds for future space-based imaging exoplanet missions," arXiv:1907.09508 (2019).
10. P. Baudoz et al., "Optimization and performance of multi-deformable mirror correction on the THD2 bench," *Proc. SPIE* **10706**, 107062O (2018).
11. J. R. Delorme et al., "Laboratory validation of the dual-zone phase mask coronagraph in broadband light at the high-contrast imaging THD testbed," *Astron. Astrophys.* **592**, A119 (2016).
12. R. Galicher, "A family of phase masks for broadband coronagraphy example of the wrapped vortex phase mask theory and laboratory demonstration," *Astron. Astrophys.* **635**, A11 (2020).
13. A. Potier, "Comparing focal plane wavefront control techniques: numerical simulations and laboratory experiments," *Astron. Astrophys.* **635**, A192 (2020).
14. R. Belikov et al., "High-contrast imaging in multi-star systems: progress in technology development and lab results," *Proc. SPIE* **10400**, 104000W (2017).
15. D. Sirbu et al., "Initial laboratory demonstration of multi-star wavefront control at the occulting mask coronagraph testbed," *Proc. SPIE* **12680**, 126800Y (2022).
16. O. Guyon et al., "High performance Lyot and PIAA coronagraphy for arbitrarily shaped telescope apertures," *Astrophys. J.* **780**, 171 (2014).
17. R. Belikov et al., "Laboratory demonstration of high contrast with the PIAACMC coronagraph on an obstructed and segmented aperture," *Proc. SPIE* **12180**, 1218025 (2022).
18. A. Bertrou-Cantou et al., "High contrast spectroscopy testbed (CST): experimental results with an apodized vortex coronagraph," *Proc. SPIE* **12680**, 126801B (2023).
19. J. Llop-Sayson et al., "Broadband wavefront control with a single mode fiber: 10^{-8} raw contrast over 20% bandwidth," *Proc. SPIE* **12680**, 126801C (2023).

20. J. J. Wang et al., “Detection and bulk properties of the HR 8799 planets with high resolution spectroscopy,” *Astron. J.* **162**, 148 (2021).
21. D. Echeverri et al., “First light of the vortex fiber nulling mode on the Keck planet imager and characterizer,” *Proc. SPIE* **12680**, 126800M (2023).
22. L. Lebouleux et al., “High-contrast imager for Complex Aperture Telescopes (HiCAT). 4. Status and wavefront control development,” *Proc. SPIE* **9904**, 99043C (2016).
23. J. Mazoyer et al., “Correcting for the effects of pupil discontinuities with the ACA method,” *Proc. SPIE* **9904**, 99044T (2016).
24. M. N’Diaye et al., “Apodized Lyot coronagraphs for arbitrary apertures. V. Hybrid shaped pupil designs for imaging Earth-like planets with future space observatories,” *Astrophys. J.* **818**, 163 (2016).
25. R. Soummer et al., “High-contrast imager for complex aperture telescopes HiCAT: 5. First results with segmented aperture coronagraph and wavefront control,” *Proc. SPIE* **10698**, 106981O (2018).
26. R. Soummer et al., “High-contrast imager for complex aperture telescopes HiCAT: 8. Dark zone demonstration with simultaneous closed-loop low-order wavefront sensing and control,” *Proc. SPIE* **12180**, 1218026 (2022).
27. R. Soummer et al., “High-contrast imager for complex aperture telescopes HiCAT: 11. System-level static and dynamic demonstration of the Apodized Lyot coronagraph with a segmented aperture,” *Proc. SPIE* **13092**, 130921Z (2024).
28. K. D. Patterson, “Design description and commissioning performance of a stable coronagraph technology development testbed for direct imaging of Earth-like exoplanets,” *Proc. SPIE* **11117**, 111171U (2019).
29. S. Meeker, “The twin decadal survey testbeds in the high contrast imaging testbed facility at NASA’s Jet Propulsion Laboratory,” *Proc. SPIE* **11823**, 118230Y (2021).
30. D. Sirbu et al., “EXCEDE technology development IV: demonstration of polychromatic contrast in vacuum at $1.2 \lambda D$,” *Proc. SPIE* **9605**, 96050J (2015).
31. K. Van Gorkom et al., “The space coronagraph optical bench (SCoOB): 4. Vacuum performance of a high contrast imaging testbed,” arXiv:2406.18885 (2024).
32. A. Harness, N. J. Kasdin, and M. Galvin, “Starshade Technology Development Milestone Demonstrations, Milestone #1b report on broadband contrast performance,” 2019, <https://exoplanets.nasa.gov/exep/technology/starshade/>.
33. A. Harness et al., “Starshade Technology Development Milestone Demonstrations, Milestone #2 report on Optical Model Validation,” 2022, <https://exoplanets.nasa.gov/exep/technology/starshade/>.
34. B. P. Crill, “An appendix to the NASA exoplanet exploration program technology plan,” 2002, https://exoplanets.nasa.gov/internal_resources/2595.
35. R. Belikov et al., “Coronagraph design survey for future exoplanet direct imaging space missions: interim update,” *Proc. SPIE* **12860**, 126802G (2023).
36. R. Belikov et al., “Coronagraph design survey for future exoplanet direct imaging missions,” *Proc. SPIE* **13092** in press (2024).
37. R. Belikov et al., “Coronagraph design survey for future exoplanet direct imaging space missions: final report,” in prep (2024).
38. C. Stark et al., “Maximizing the exoearth candidate yield from a future direct imaging mission,” *Astrophys. J.* **795**, 122 (2014).
39. C. Stark et al., “Lower limits on aperture size for an exoearth detecting coronagraphic mission,” *Astrophys. J.* **808**, 149 (2015).
40. D. Savransky et al., “Quantifying the impact of schedulability on science yield of exoplanet missions,” *Proc. SPIE* **12680**, 126801K (2023).
41. D. Savransky, C. Delacroix, and D. Garrett, “Multi mission modeling for space-based exoplanet imagers,” *Proc. SPIE* **10400**, 104001L (2017).
42. B. Nemati et al., “Method for deriving optical telescope performance specifications for Earth-detecting coronagraphs,” *J. Astron. Telesc. Instrum. Syst.* **6**, 039002 (2020).
43. G. Ruane et al., “Review of high-contrast imaging systems for current and future ground- and space-based telescopes. I. Coronagraph design methods and optical performance metrics,” *Proc. SPIE* **10698**, 106982S (2018).
44. R. Belikov et al., “Theoretical performance limits for coronagraphs on obstructed and unobstructed apertures: how much can current designs be improved?,” *Proc. SPIE* **11823**, 118230W (2021).
45. B.-J. Seo et al., “Testbed demonstration of high-contrast coronagraph imaging in search for Earth-like exoplanets,” *Proc. SPIE* **11117**, 111171V (2019).
46. G. Allan et al., “Demonstration of coronagraph technology for high-contrast point spectroscopy of Exoearths,” *Proc. SPIE* **12680**, 1268019 (2023).
47. J. Trauger et al., “A hybrid Lyot coronagraph for the direct imaging and spectroscopy of exoplanet systems: recent results and prospects,” *Proc. SPIE* **8151**, 81510G (2011).

48. J. Trauger et al., “Complex apodization Lyot coronagraphy for the direct imaging of exoplanet systems: design, fabrication and laboratory demonstration,” *Proc. SPIE* **8442**, 84224Q (2012).
49. G. Foo et al., “Optical vortex coronagraph,” *Opt. Lett.* **30**, 3308 (2005).
50. G. Swartzlander, “Broadband nulling of a vortex phase mask,” *Opt. Lett.* **30**, 2876 (2005).
51. D. Mawet et al., “Annular groove phase mask coronagraph,” *Astrophys. J.* **633**, 1191 (2005).
52. D. Mawet et al., “Optical vectorial vortex coronagraphs using liquid crystal polymers: theory, manufacturing and laboratory demonstration,” *Opt. Express* **17**, 1902 (2009).
53. D. Mawet et al., “The vector vortex coronagraph: laboratory results and first light at Palomar Observatory,” *Astrophys. J.* **709**, 53 (2010).
54. G. Ruane et al., “Vortex coronagraphs for the Habitable exoplanet imaging mission concept: theoretical performance and telescope requirements,” *J. Astron. Telesc. Instrum. Syst.* **4**(1), 015004 (2018).
55. J. Krist et al., “Numerical modeling of the HabEx coronagraph,” *Proc. SPIE* **11117**, 1111705 (2019).
56. K. Fogarty et al., “The PIAA vortex coronagraph: a new coronagraph technology to maximize exo-Earth yields in the Astro2020 era,” *Proc. SPIE* **12180**, 1218028 (2022).
57. G. Ruane et al., “Broadband vector vortex coronagraph testing at NASA’s high contrast imaging testbed facility,” *Proc. SPIE* **12180**, 1218024 (2022).
58. A.J. Riggs et al., “High-contrast demonstration of a vector vortex coronagraph with a segmented, off-axis aperture,” *Proc. SPIE* **12180**, 121805I (2022).
59. O. Guyon, “Phase-induced amplitude apodization of telescope pupils for extrasolar terrestrial planet imaging,” *Astron. Astrophys.* **404**, 379 (2003).
60. R. Belikov, “Laboratory demonstration of high-contrast imaging at inner working angles $2 \lambda/D$ and better,” *Proc. SPIE* **8151**, 815102 (2011).
61. D. Sirbu, “Model validation of phase-induced amplitude apodization complex mask coronagraph for LUVOIR-A in vacuum,” *Proc. SPIE* **11823**, 118230R (2021).
62. B. Kern et al., “Laboratory demonstration of phase induced amplitude apodization (PIAA) coronagraph with better than 10^{-9} contrast,” *Proc. SPIE* **8864**, 88640R (2013).
63. O. Guyon et al., “Phase-induced amplitude apodization (PIAA) technology development milestone #3: 10% bandpass contrast demonstration,” 2014, <https://exoplanets.nasa.gov/exep/technology/TDEM-awards/>.
64. E. H. Por, “Phase-apodized-pupil Lyot coronagraphs for arbitrary telescope pupils,” *Astrophys. J.* **888**, 127 (2020).
65. A.J. Eldorado Riggs et al., “Flight mask designs of the Roman Space Telescope coronagraph instrument,” *Proc. SPIE* **11823**, 118231Y (2021).
66. B.-J. Seo, “Hybrid Lyot coronagraph for WFIRST: high-contrast broadband testbed demonstration,” *Proc. SPIE* **10400**, 104000F (2017).
67. E. Cady, “Shaped pupil coronagraphy with WFIRST: high-contrast broadband testbed demonstration,” *Proc. SPIE* **10400**, 104000E (2017).
68. J. E. Krist et al., “End-to-end numerical modeling of the Roman Space Telescope coronagraph,” *J. Astron. Telesc. Instrum. Syst.* **9**, 045002, (2023).
69. B. Kern et al., “Phase-induced amplitude apodization complex mask coronagraph fabrication, characterization, and modeling for WFIRST-AFTA,” *J. Astron. Telesc. Instrum. Syst.* **2**, 011014 (2016).
70. N. J. Kasdin et al., “Optimal one-dimensional apodizations and shaped pupils for planet finding coronagraphy,” *Appl. Opt.* **44**, 1117 (2005).
71. R. Belikov et al., “Demonstration of high contrast in 10% broadband light with the shaped pupil coronagraph,” *Proc. SPIE* **6693**, 66930Y (2007).
72. O. Guyon et al., “High performance PIAA coronagraphy with complex amplitude focal plane masks,” *Astrophys. J. Suppl. Ser.* **190**(2), 220 (2010).
73. R. Belikov et al., “Design and performance analysis of a PIAACMC coronagraph on a segmented aperture,” *Proc. SPIE* **10698**, 106981H (2018).
74. D. Marx et al., “Results from the laboratory demonstration of a PIAACMC coronagraph with a segmented aperture,” *Proc. SPIE* **11823**, 118230O (2021).
75. L. Spitzer, “The beginning and future of space astronomy,” *Am. Sci.* **50**, 473 (1962).
76. R.J. Vanderbei, E. Cady, and J. Kasdin, “Optimal occulter design for finding extrasolar planets,” *Astrophys. J.* **665**, 794 (2007).
77. F. Shi et al., “WFIRST low order wavefront sensing and control dynamic testbed performance under the flight-like photon flux,” *Proc. SPIE* **10698**, 106982O (2018).
78. R. Pourcelot et al., “Low-order wavefront control using a Zernike sensor through Lyot coronagraphs for exoplanet imaging. Blind stabilization of an image dark hole,” *Astron. Astrophys.* **663**, A49 (2022).
79. L. Pueyo et al., “The LUVOIR extreme coronagraph for living planetary systems (ECLIPS) I: searching and characterizing exoplanet gems,” *Proc. SPIE* **11117**, 1111703 (2019).
80. G. Ruane et al., “Dual purpose Lyot coronagraph masks for simultaneous high-contrast imaging and high resolution wavefront sensing,” *Proc. SPIE* **12680**, 1268018 (2023).

81. L. Pogorelyuk and J. Kasdin, “Dark hole maintenance and a posteriori intensity estimation in the presence of speckle drift in a high-contrast space coronagraph,” *Astrophys. J.* **873**, 95 (2019).
82. S. Redmond et al., “Dark hole maintenance for future coronagraphic space missions,” *Proc. SPIE* **12180**, 121802B (2022).
83. O. Guyon et al., “Spectral linear dark field control: stabilizing deep contrast for exoplanet imaging using out-of-band speckle field,” (2017).
84. P. K. Poon et al., “Experimental demonstration of spectral linear dark field control at NASA’s high contrast imaging testbeds,” *Proc. SPIE* **12680**, 126801G (2023).
85. O. Guyon et al., “Linear wavefront control high contrast imaging,” Milestone #1 report whitepaper, 2021, https://exoplanets.nasa.gov/internal_resources/1996/.
86. T. Currie et al., “Laboratory demonstration of spatial linear dark field control for imaging extrasolar planets in reflected light,” *Publ. Astron. Soc. Pac.* **132**, 104502 (2020).
87. D. Lafreniere et al., “A new algorithm for PSF subtraction in high-contrast imaging: a demonstration with angular differential imaging,” *Astrophys. J.* **660**, 770 (2007).
88. R. Soummer, L. Pueyo, and J. Larkin, “Detection and characterization of exoplanets and disks using projections on Karhunen-Loeve eigenimages,” *Astrophys. J.* **755**, L28 (2012).
89. J. Debes et al., “Pushing the limits of the coronagraphic occulters on the Hubble Space Telescope/Space Telescope Imaging Spectrograph,” *J. Astron. Telesc. Instrum. Syst.* **5**, 035003 (2019).
90. J. Kammerer et al., “Performance of near-infrared high-contrast imaging methods with JWST from commissioning,” *Proc. SPIE* **12180**, 121803N (2022).
91. A. Boccaletti et al., “JWST/MIRI coronagraphic performances as measured on-sky,” *Astron. Astrophys.* **667**, A165 (2022).
92. M. Ygouf, “Data processing for high-contrast imaging with the James Webb Space Telescope,” *Proc. SPIE* **11443**, 114433N (2020).
93. O. Guyon et al., “High contrast imaging at the photon noise limit with WFS-based PSF calibration,” *Proc. SPIE* **12185**, 121850E (2022).
94. O. Guyon et al., “Linear dark field control, Milestone #3 whitepaper: LDFC PSF calibration: 10x contrast gain to 10^{-8} contrast,” 2022, https://exoplanets.nasa.gov/internal_resources/2240/.
95. L. Pogorelyuk, J. Kasdin, and C. W. Rowley, “Reduced order estimation of the speckle electric field history for space-based coronagraphs,” *Astrophys. J.* **881**, 126 (2019).
96. L. Pogorelyuk, L. Pueyo, and J. Kasdin, “On the effects of pointing jitter, actuator drift, telescope rolls, and broadband detectors in dark hole maintenance and electric field order reduction,” *J. Astron. Telesc. Instrum. Syst.* **6**, 039001 (2020).
97. I. Snellen et al., “Spin of the young extrasolar planet β Pictoris b,” *Nature*, **509**, 63 (2014).
98. H. J. Hoeijmakers et al., “Medium-resolution integral-field spectroscopy for high contrast imaging. Molecule maps of the β Pictoris system with SINFONI,” *Astron. Astrophys.* **617**, A144 (2018).
99. S. Y. Haffert et al., “Two accreting protoplanets around the young star PDS 70,” *Nat. Astron.* **3**, 749 (2019).
100. Q.M. Konopacky et al., “Detection of carbon monoxide and water absorption lines in an exoplanet atmosphere,” *Science* **339**, 1398 (2013).
101. J. Wang et al., “Observing exoplanets with high dispersion coronagraphy. I. The scientific potential of current and next generation large ground and space telescopes,” *Astron. J.* **153**, 183 (2017).
102. J. Wang et al., “Baseline requirements for detecting biosignatures with the Habex and LUVOIR mission concepts,” *J. Astron. Telesc. Instrum. Syst.* **4**, 035001 (2018).
103. J-B. Ruffio et al., “JWST-TST high contrast: achieving direct spectroscopy of faint substellar companions next to bright stars with the NIRSpec IFU,” (2023).
104. D. Echeverri et al., “Vortex fiber nulling for exoplanet observations: implementation and first light,” *J. Astron. Telesc. Instrum. Syst.* **9**, 035002 (2023).
105. D. Echeverri et al., “Vortex fiber nulling for exoplanet observations: first direct detection of M dwarf companions around HIP 21543, HIP 94666, and HIP 50319,” *Astrophys. J. Lett.* **965**, L15 (2024).
106. S. Ertel et al., “The HOST survey for exozodiacal dust: observational results from the complete survey,” *Astron. J.* **159**, 177 (2020).
107. D. Defrere et al., “Direct imaging of exoearths embedded in clumpy debris disks,” *Proc. SPIE* **8442**, 84420M (2012).
108. S. Ertel et al., “A near-infrared interferometric survey of debris disks stars. V. Pionier search for variability,” *Astron. Astrophys.* **595**, A44 (2016).
109. P. Nunez et al., “A near-infrared interferometric survey of debris disks stars. VI. Extending the exozodiacal light survey with CHARA/JouFLU,” *Astron. Astrophys.* **608**, 113 (2017).
110. B. Mennesson et al., “Interplanetary dust around main sequence stars: origin, magnitude, and implications for exoplanet habitability searches,” *Bull. Am. Astron. Soc.* **51**, 3 (2019).
111. B. Mennesson et al., “The potential of exozodiacal disks observations with the WFIRST coronagraph instrument,” arXiv: 1909.02161 (2019).

112. E. Douglas et al., “Sensitivity of the Roman coronagraph instrument to exozodiacal dust,” *Publ. Astron. Soc. Pac.* **134**, 024402 (2022).
113. M. Ygouf et al., “Roman coronagraph instrument post-processing report—OS9 HLC,” 2021, https://roman.ipac.caltech.edu/docs/sims/20210402_Roman_CGI_post_processing_report_URS.pdf.
114. E. Mamajek and K. Stapelfeldt, “NASA ExEP mission star list for the Habitable Worlds Observatory 2023,” 2024, <https://arxiv.org/ftp/arxiv/papers/2402/2402.12414.pdf>.
115. B. Mennesson et al., “The Roman Space Telescope coronagraph technology demonstration: current status and relevance to future missions,” *Proc. SPIE* **12180**, 121801W (2022).
116. N. Desai et al., “Achromatizing scalar vortex coronagraphs with radial phase mask dimples,” *Proc. SPIE* **12680**, 126802B (2023).
117. N. Desai et al., “Integrated photonic-based coronagraphic systems for future space telescopes,” *Proc. SPIE* **12680**, 126801S (2023).
118. E. Serabyn, K. Liewer, and G. Ruane, “Geometric-phase-based phase-knife mask for stellar nulling and coronagraphy,” *Opt. Express* **32**, 19924 (2024).
119. S. Shaklan and D. Lisman, “Starshade attributes” presentation, towards starlight suppression for the Habitable Worlds Observatory workshop,” Caltech, https://exoplanets.nasa.gov/internal_resources/2837/Aug092023_1125_Shaklan_Lisman_Attributes-Towards-Starlight-Suppression-Workshop.pdf (accessed 9 August 2023).
120. R. Morgan et al., “Exoearth yield of a 6 m space telescope in the near-infrared,” *Proc. SPIE* **12680**, 126801L (2023).
121. C. Stark et al., “Optimized bandpasses for the Habitable Worlds Observatory’s exoearth survey,” *J. Astron. Telesc. Instrum. Syst.* **10**, 014005 (2024).
122. N. Latouf et al., “Bayesian analysis for remote biosignature identification on exoearths (BARBIE). I. Using grid-based nested sampling in coronagraphy observations simulation for H₂O,” *Astron. J.* **166**, 129 (2023).
123. N. Latouf et al., “Bayesian analysis for remote biosignature identification on exoearths (BARBIE). II. Using grid-based nested sampling in coronagraphy observations simulation for O₂ and O₃,” *Astron. J.* **167**, 27 (2024).
124. C. Stark et al., “A direct comparison of exoearth yields for starshades and coronagraphs,” *Proc. SPIE* **9904**, 99041U (2016).
125. P. Willems et al., “NASA’s starshade technology development activity,” *Proc. SPIE* **12180**, 121802N (2022).
126. “Starshade to TRL5 (S5) Technology Development Plan,” 2018, https://exoplanets.nasa.gov/internal_resources/1033/.

Bertrand Mennesson is a principal research scientist at the Jet Propulsion Laboratory, California Institute of Technology. He received his MS degree in physical engineering from the Ecole des Mines, France, in 1993, and his MS degree and PhD in astrophysics and space techniques from the University of Paris VII in 1994 and 1999, respectively. His expertise is in the design, assembly, and scientific exploitation of high-contrast high-resolution optical systems—coronagraphs, interferometers, and starshades—used for astronomical imaging. His main scientific focus is on the direct detection and spectroscopy of exoplanetary systems, ultimately leading to the search for life on Earth-like exoplanets through the use of new instruments and data reduction techniques. He is the author or co-author of over 300 publications in specialized journals and at professional conferences.

Biographies of the other authors are not available.