

Flight designs and pupil error mitigation for the bowtie shaped pupil coronagraph on the Nancy Grace Roman Space Telescope

Jessica Gersh-Range,^{a,*} A J Eldorado Riggs^{©,b}, and N. Jeremy Kasdin^c

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

^bJet Propulsion Laboratory, California Institute of Technology, Pasadena, California, United States

^cUniversity of San Francisco, San Francisco, California, United States

Abstract. The Nancy Grace Roman Space Telescope will carry a coronagraph instrument (CGI) that will serve as a demonstrator for technologies needed for future high-contrast imaging missions in space, including deformable mirrors (DMs) to correct high-order wavefront errors that would otherwise limit the achievable contrast. The CGI has three baselined interchangeable observing configurations, one of which is a bowtie shaped pupil coronagraph for high-contrast spectroscopy. We present the flight designs for two closely related mask configurations of the bowtie shaped pupil coronagraph: a baseline 0-deg mask configuration for the technology demonstration and a 60-deg mask configuration contributed by the NASA Exoplanet Exploration Program. The shaped pupil mask and Lyot stop for each mask configuration result from an iterative process that maximizes the core throughput subject to constraints on other performance metrics, such as the contrast: a linear program optimizes the shaped pupil mask for a given Lyot stop, and the optimization repeats for various Lyot stops until the highest-throughput combination is identifiable. The flight designs for the baseline and rotated mask configurations have core throughputs of 4.50% and 3.89%, respectively, at $4 \frac{\lambda}{D}$ and are robust to conservative estimates of potential pupil errors such as misalignments and manufacturing errors. If these estimates are exceeded in flight, the DMs can be used to mitigate the effects of the excess error. © *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.8.2.025003](https://doi.org/10.1117/1.JATIS.8.2.025003)]

Keywords: bowtie shaped pupil coronagraph; Nancy Grace Roman Space Telescope; high-contrast spectroscopy; pupil padding; deformable mirror.

Paper 21146G received Nov. 10, 2021; accepted for publication Apr. 14, 2022; published online May 31, 2022.

1 Introduction

Since time immemorial, people have looked to the stars and wondered if there might be life elsewhere in the universe. This question has appeared in some form in every decadal survey for astronomy since 1972,¹⁻⁶ and over the decades, multiple techniques have been developed to detect exoplanets, including those in the habitable zone of their parent star.⁷⁻¹³ Two of these techniques, transits and direct imaging, can be used to spectroscopically analyze the light from the exoplanet, thereby enabling a search for biomarkers that suggest the possibility of life. Currently, a space-based direct imaging mission provides one of the best paths to the ultimate goal of searching for life on an exoplanet similar to Earth, an exoplanet of Earth-like size in the habitable zone of a Sun-like star.¹³ Large space-based direct imaging mission concepts, such as the Habitable Exoplanet Observatory (HabEx)¹⁴ and the Large UV/Optical/IR Surveyor (LUVOIR),¹⁵ are being developed with this goal in mind, and a 6-m class infrared/optical/ultraviolet telescope for high-contrast imaging and spectroscopy is listed in the most recent

*Address all correspondence to Jessica Gersh-Range, jg36@princeton.edu

decadal survey as the highest-priority recommendation for space frontier missions.⁶ However, the challenge of directly imaging and spectroscopically characterizing an Earth-like exoplanet is not insignificant; there are currently numerous technical gaps that will need to be addressed.^{16,17}

To advance our general understanding of high-contrast imaging in space, NASA's Nancy Grace Roman Space Telescope will carry a coronagraph instrument (CGI) that will serve as a demonstrator for some of the technology that will be needed for future missions. Notably, the CGI will be the first space-based coronagraph to use deformable mirrors (DMs) to correct high-order wavefront errors that would otherwise limit the achievable contrast, to incorporate low-order wavefront sensing and control, and to use precision-fabricated masks specifically designed to maximize the throughput.^{18,19} As a technology demonstrator, the CGI has three baseline interchangeable observing configurations, one of which is for high-contrast spectroscopy. (The other two observing configurations are for imaging: a narrow-field-of-view hybrid Lyot coronagraph designed for a 10% spectral band centered at 575 nm, and a wide-field-of-view shaped pupil coronagraph designed for a 10% spectral band centered at 825 nm; imaging over other spectral bands is possible but unsupported.²⁰⁻²²) The spectroscopy configuration uses a shaped pupil coronagraph to achieve a contrast of 3×10^{-9} within a two-sided bowtie-shaped region spanning 3.0 to $9.0 \frac{\lambda}{D}$ for wavelengths across a 15% spectral band centered at either 660 or 730 nm, and the coronagraph masks are followed by a slit-prism spectrograph with a spectral resolution of $R = 50$ at the center of the selected spectral band.^{20,21,23} Due to the shape of the high-contrast region (also called the dark hole), this coronagraph is referred to as the bowtie shaped pupil coronagraph.

The bowtie shaped pupil coronagraph has two closely related mask configurations: a baseline mask configuration that will be used for the technology demonstration, and a mask configuration that is rotated 60 deg relative to the baseline mask configuration. (To accommodate two different 15% spectral bands, there are two variations of each of these mask configurations. However, because the two variations of a given mask configuration differ only in their physical scaling, we will not discuss them individually in this paper. A full list of every designed configuration for the CGI may be found in a paper by Riggs et al.²¹) The rotated mask configuration is not part of the technology demonstration but contributed by the NASA Exoplanet Exploration Program (ExEP) so that it can be available for use if scheduling and funding constraints allow, and it is important to note that the rotated mask configuration is being accommodated on the grounds that it does no harm to the baseline design and causes no schedule delay.²⁴ The rationale for contributing the rotated mask configuration is that it improves the annular field of view coverage and increases the probability of detecting our canonical target exoplanet, 47 Ursae Majoris c (47 UMa c); with the two bowties and a ± 13 -deg roll, the detection probability during the nominal technology demonstration and beyond is nearly as high as it would be if the coronagraph had a 360-deg field of view due to the combination of the two bowtie openings, the bowties precessing, and orbital motion.²⁵⁻²⁷ (However, we emphasize that no observing time is currently allocated to the rotated mask configuration.)

The two mask configurations each produce an appropriately rotated dark hole by using a set of three masks: a shaped pupil mask for the first pupil plane and bowtie-shaped masks for the first focal plane and the Lyot plane (the focal plane mask and Lyot stop, respectively), as shown in Fig. 1. The focal plane mask design is determined by a set of requirements on the field of view and rotated to match the specified orientation; it is identical for the two mask configurations except for the orientation. The shaped pupil mask and Lyot stop, by comparison, are products of a design survey that searches for the combination with the highest throughput. In this survey, a shaped pupil mask is optimized for a particular Lyot stop (Sec. 2), and the optimization is repeated for a range of Lyot stops (Sec. 3). The final designs for the baseline and rotated mask configurations have core throughputs of 4.50% and 3.89%, respectively, for an off-axis source at $4 \frac{\lambda}{D}$. These designs are robust to conservative (three-sigma) estimates of potential pupil errors (such as misalignments and manufacturing errors), but in the event that these estimates are exceeded, the DMs that correct for high-order wavefront errors may be used to mitigate the effects of the pupil errors as well (Sec. 4).

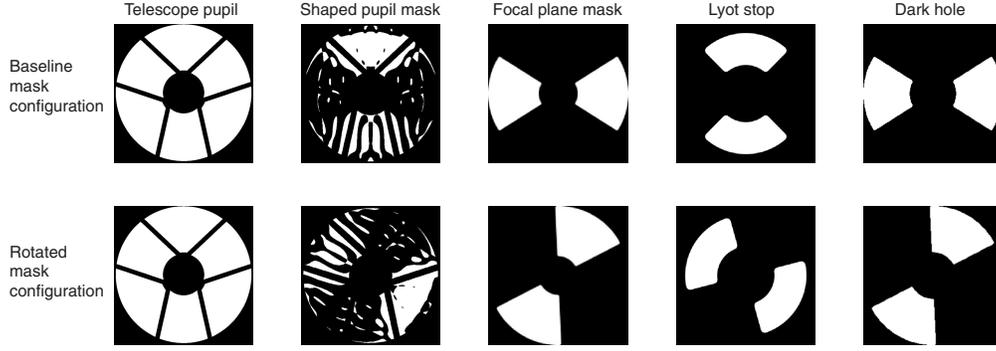


Fig. 1 Comparison of the masks in each plane for the baseline and rotated mask configurations of the bowtie shaped pupil coronagraph. The bowtie shaped pupil coronagraph produces a dark hole by using a shaped pupil mask in the first pupil plane, a bowtie-shaped mask in the first focal plane (the focal plane mask), and a bowtie-shaped mask in the second pupil plane (the Lyot stop). The Lyot stop is rotated 90 deg with respect to the focal plane mask, and the two bowties are not necessarily identical. For the baseline mask configuration, the dark hole is symmetric about both the horizontal (x) and vertical (y) axes, with the bowtie lobes along the horizontal axis. For the rotated mask configuration contributed by ExEP, the shaped pupil mask, focal plane mask, and Lyot stop are designed to produce a dark hole that is rotated 60 deg from the baseline. For reference, the nominal telescope pupil is also shown.

2 Shaped Pupil Mask Design Process

Due to computational constraints, the shaped pupil mask and the Lyot stop are designed using an iterative optimization process that, subject to constraints on the contrast and other performance metrics, seeks to identify the combination with the highest core throughput T , which is defined in the official CGI project requirements²⁸ as the ratio of the power passing through the half-max core of the point spread function (PSF) in the final image plane, P_{halfmax} , to the power passing through the pupil plane in the absence of a shaped pupil mask, $P_{\text{pupil,total}}$:

$$T = \frac{P_{\text{halfmax}}}{P_{\text{pupil,total}}}. \quad (1)$$

Because the PSF for the bowtie shaped pupil coronagraph has multiple bright lobes, it is important to note that the half-max core of the PSF is defined as the set of all points in the image plane for which the intensity is at least half of the maximum value, not just the points in the main lobe (Fig. 2).

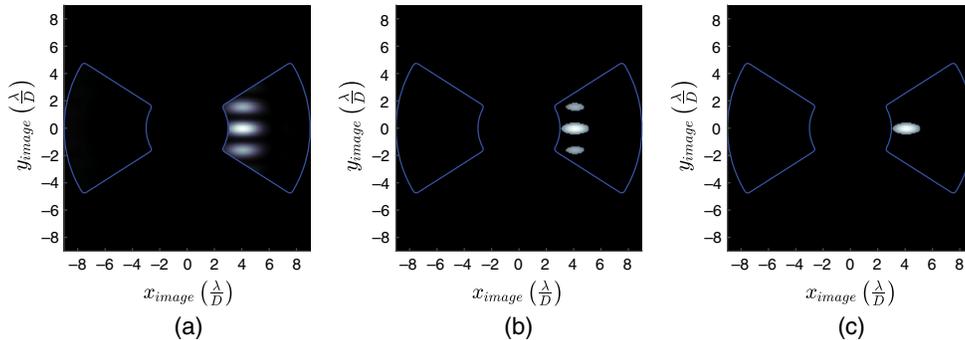


Fig. 2 Comparison of definitions of the half-max core of the PSF. The PSF for the bowtie shaped pupil coronagraph (a) consists of three lobes, so two definitions of the half-max core exist. Following the conventions set by the CGI Project, we define the half-max core as the set of all points for which the intensity is at least half of the maximum value (b). A more restrictive definition that is sometimes used includes only the points in the main lobe that have an intensity of at least half of the maximum value (c). The three plots are each shown with a linear scale, with the dark hole outlined in blue.

The throughput value is dependent on the location of the point source used to calculate the PSF, and we have selected a point source at $4 \frac{\lambda}{D}$. This value was chosen because it corresponds to the location of 47 UMa c, our canonical target exoplanet that was chosen because it is a confirmed exoplanet around a star with a V magnitude of almost 5 and a small angular diameter.²⁹

The emphasis on maximizing the core throughput was driven by the fact that it can be one of the primary factors limiting the ability to take spectroscopic measurements; however, in scenarios dominated by speckle instability, the effects of adjusting the throughput are more muted. This metric changed from a more signal-to-noise ratio (SNR) metric (that accounted for PSF sharpness) when the CGI switched from having an integral field spectrograph (IFS) to a slit spectrograph. For the IFS, the area of the PSF was a concern because observations were dark-current limited. With the switch to the slit spectrograph, the planet PSF is no longer broken up across several lenslets, thereby reducing the PSF area on the detector by a factor of a few. As a result, the challenge is to increase the throughput to reduce the photon noise of the planet itself (which is assumed to be faint, with a contrast on the order of 1×10^{-9}). Throughput is also an easy-to-use proxy for sharpness, while sharpness makes optimizations intractable.

The first step of the design process is to construct an appropriate starting point for the shaped pupil mask optimization by creating a representation of the telescope pupil with the correct orientation (Sec. 2.2) and adding padding to compensate for potential pupil errors, such as clocking, rotation, translation, and magnification (Sec. 2.3). The resulting padded telescope pupil serves as the shaped pupil mask optimization starting point throughout an entire design survey, but the padding may be adjusted for different surveys. To complete a design survey, an optimization routine determines the highest-throughput shaped pupil mask for a given Lyot stop (Sec. 2.1), and the optimization is repeated for a range of potential Lyot stops until the highest-throughput combination can be identified (Sec. 3).

2.1 Optimizing a Shaped Pupil Mask

The bowtie shaped pupil masks are optimized using a linear program that maximizes the peak pixel value for an off-axis planet (a proxy for the throughput) while also constraining the contrast and the sensitivity to tip and tilt. This linear program (written in AMPL³⁰) treats all of the equations for propagating light through the coronagraph as constraints, and it incorporates tip and tilt with the addition of separate propagations for each type of aberrated wavefront. A thorough description of how to model the propagation of light through the coronagraph and the specific techniques and equations used by the linear program may be found in papers by Zimmerman and Riggs and are not repeated here.^{31–33}

Directly constraining the sensitivity to tip and tilt leads to shaped pupil masks that are more robust to residual pointing jitter and resolved stellar diameters, but it also reduces the throughput. It is therefore important to balance the benefits of reducing the tip and tilt sensitivity against the benefits of increasing the throughput. The linear program includes a coefficient in the aberrated wavefront propagations that allows the amount of tip and tilt in an optimization to be adjusted, with higher amounts producing shaped pupil masks with lower sensitivity. For the final flight shaped pupil masks, the value of this coefficient was selected based on the results of a limited design survey conducted by the CGI Integrated Modeling Team that evaluated the overall performance of masks optimized with various amounts of tip and tilt included.³⁴

Because the linear program simulates propagating light through the coronagraph to evaluate each potential shaped pupil mask design, it requires knowledge of the coronagraph masks in each plane as well as the desired shape of the dark hole. As a result, two-dimensional arrays representing the shaped pupil mask starting point, the focal plane mask, and the Lyot stop are required inputs to the linear program, and the shape of the dark hole is specified by a set of constraints. To keep the optimization problem tractable, we require y-axis symmetry for each of the coronagraph masks and the dark hole and exploit additional symmetry wherever possible.³⁵

The y-axis symmetry requirement significantly reduces the size of the optimization problem, but it also means that care must be taken when defining the arrays for the coronagraph masks. For the baseline mask configuration, the dark hole and almost all of the coronagraph masks are symmetric about the y axis due to the choice of orientation; the bowties for the focal plane mask, Lyot stop, and dark hole each have two axes of symmetry, and these axes align with the x and y

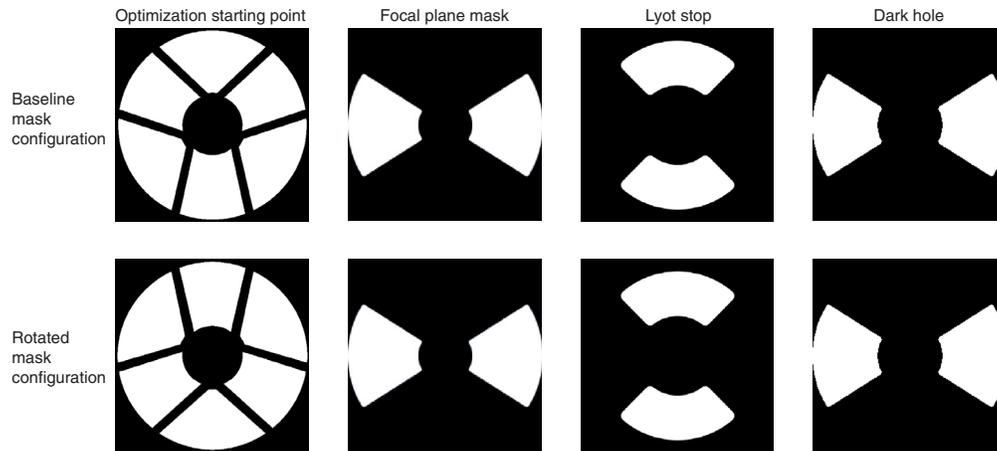


Fig. 3 Sample shaped pupil mask optimization starting points, focal plane masks, Lyot stops, and dark holes used for the baseline and rotated mask configurations. Although the full-plane masks are shown, the optimization process assumes y -axis symmetry and uses only the right half (or less when additional symmetry can be exploited). For the baseline (0-deg) mask configuration, y -axis symmetry is enforced using a symmetrized version of the telescope pupil as the starting point for the shaped pupil mask optimization; the remaining masks and the dark hole require no modification. To enforce symmetry for the rotated (60-deg) mask configuration, the telescope pupil is rotated instead of the focal plane mask, Lyot stop, and dark hole; the optimization starting point is a symmetrized version of the rotated telescope pupil.

axes for the 0-deg orientation. The shaped pupil mask, by comparison, is symmetric only because the linear program assumes it to be and outputs only one half of the mask. Because the telescope pupil is slightly asymmetric, we use a symmetrized version as the shaped pupil mask optimization starting point to ensure that the shaped pupil mask completely covers the telescope pupil as assumed (Fig. 3). For the rotated (60-deg) mask configuration, neither the dark hole nor any of the coronagraph masks are symmetric about the y axis (as illustrated in Fig. 1). To satisfy the symmetry requirement, we optimize a shaped pupil mask for the rotated mask configuration by rotating the telescope pupil 60 deg counterclockwise and keeping the baseline orientation for the remaining planes; the shaped pupil optimization starting point is therefore a symmetrized version of the rotated telescope pupil (Fig. 3). It is also worth noting that although we can optimize shaped pupil masks for a 90-deg mask configuration because the axes of symmetry for all three bowties again align with the x and y axes, the throughput is significantly lower due to the less-advantageous orientation of the dark hole with respect to the struts supporting the secondary mirror (Fig. 4).

2.2 Constructing an Appropriately Rotated Unpadded Telescope Pupil

The first step to generating the starting point for the shaped pupil mask optimization is to construct a representation of the telescope pupil that has been rotated appropriately for the desired mask configuration. This representation is based on a map of the nominal CGI entrance pupil that was generated using the FRED raytracing software program.^{36–38} With 7980 and 8054 illuminated pixels along the x and y axes, respectively, this pupil map has a much higher resolution than the shaped pupil mask optimization process needs (1000 pixels along an axis is sufficient), and because this map is provided as a single array for the entire pupil, it does not enable the manipulation of individual elements, such as the struts, directly. However, the ability to arbitrarily modify any pupil element facilitates adding specialized pupil padding tailored to provide the desired degree of robustness to individual pupil errors (Sec. 2.3); although pupil padding can be added for global errors (translation, magnification, and clocking) by manipulating the provided array as a whole, adding specific pupil padding for component-level errors (misalignments and errors in the sizes of the struts, central obscuration, and tabs) requires working with the individual components directly. The ability to manipulate individual pupil elements also facilitates

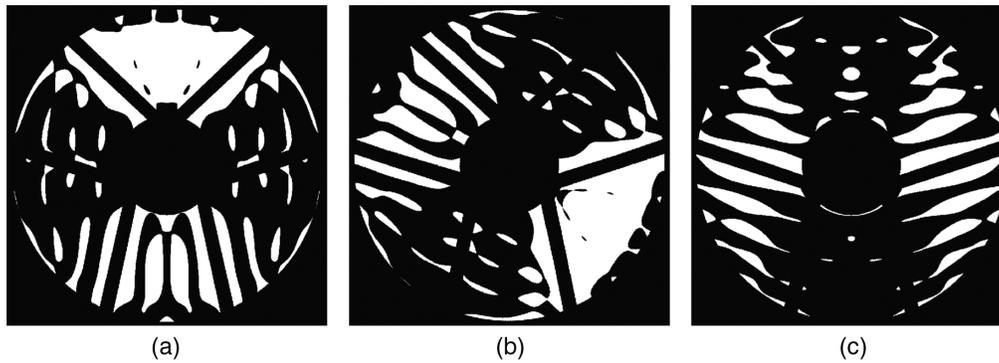


Fig. 4 Sample shaped pupil masks optimized for 0-, 60-, and 90-deg mask configurations. Although the 0-, 60-, and 90-deg mask configurations can all be set up to have the y -axis symmetry required by the optimization routine, the resulting shaped pupil masks are not all equivalent. The shaped pupil masks for the 0-deg and 60-deg mask configurations (a and b, respectively) have comparable throughputs due to the nearly identical positioning of the dark hole relative to the struts. The shaped pupil mask for the 90-deg mask configuration (c), however, has approximately half the throughput, making it a less desirable option for the rotated mask configuration.

studying the impacts of both the pupil padding and any pupil errors (Sec. 4). It is therefore preferable to work with a modifiable, lower-resolution representation generated by a telescope pupil model that simplifies the true telescope pupil into a collection of well-defined lines and shapes that can be individually manipulated as desired.

In the telescope pupil model, the pupil is divided into multiple components that are each modeled individually: the primary mirror, the secondary mirror, the six struts that support the secondary, and three tabs (which create small obscurations next to the secondary). The primary and secondary mirrors are each represented by an ellipse, and the struts are each represented (at least initially) by a rectangle. The tabs are then represented by sections of ellipses that fall between the appropriate struts. The sets of parameters that define the individual ellipses and rectangles are determined independently with one exception: the two tabs that are not located on the y axis for a pupil orientation of 0 deg are assumed to share the same ellipse. A list of the parameter values and an explanation of the methods used to infer them are provided in the [Appendix](#). Once defined, these component models are combined to form the full telescope pupil model (Fig. 5).

To provide important cross-checks and additional insights into the behavior of the telescope pupil as it is rotated (for optimizing a 60-deg shaped pupil mask) and padded (modified by blocking additional regions of the pupil plane to enforce symmetry and account for potential pupil errors), we use two independent methods to construct the specific telescope pupil models for the baseline and rotated mask configurations. (The two mask configurations were also designed semi-independently, with the baseline mask configuration primarily designed at the Jet Propulsion Laboratory, JPL, and the rotated mask configuration primarily designed at Princeton, as described in Sec. 3). For the baseline mask configuration, the telescope pupil model is generated with PROPER,³⁹ a library of optical propagation routines developed at JPL, as the superposition of elliptical and rectangular obscurations. For the rotated mask configuration, the telescope pupil model consists of a set of algebraic equations (explained below) that allow the individual edges of the obscurations to be manipulated directly, and the equations are written in a way that allows the pupil to be rotated as needed. Although not identical, the telescope pupil representations resulting from these two models agree very closely when generated for the same orientation, with only negligible differences (Fig. 6).

In the algebraic model, we generate a binary representation of the telescope pupil by first defining sets delineated by ellipses and lines that describe the clear area of the unobscured pupil, S_{unobsc} , and the individual obscurations due to the secondary, struts, and tabs, and then applying these set definitions to a grid of points. By binary, we mean that a point in the pupil plane is either fully clear or fully obscured, with clear points assigned a value of 1 and obscured points (as well as points in the pupil plane that fall outside S_{unobsc}) assigned a value of 0. The total obscuration

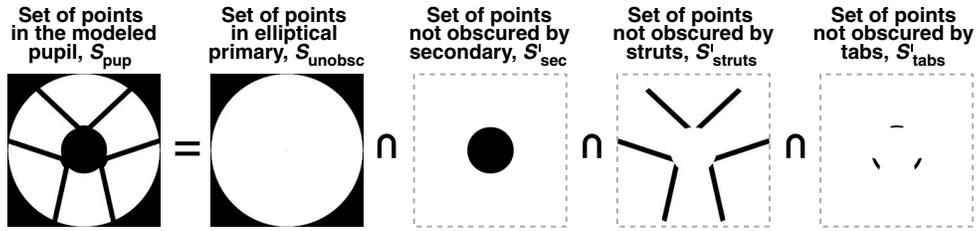


Fig. 5 Telescope pupil model. When optimizing the baseline and rotated mask configurations, it is useful to work with a lower-resolution telescope pupil representation generated by a model that divides the pupil into a set of individually modeled components: an elliptical primary mirror, an elliptical secondary mirror, six initially rectangular struts, and three tabs that are sections of ellipses between pairs of struts. For the baseline mask configuration, the telescope pupil model superposes elliptical and rectangular obscurations in PROPER. For the rotated mask configuration, the telescope pupil model consists of a set of algebraic equations that define the set S_{unobsc} of points in the unobscured primary, the set S'_{sec} of points not obscured by the secondary, the set S'_{struts} of points not obscured by the struts, and the set S'_{tabs} of points not obscured by the tabs. The set of points S_{pup} in the clear area of the modeled telescope pupil is then defined as $S_{pup} = S_{unobsc} \cap S'_{sec} \cap S'_{struts} \cap S'_{tabs}$. When these set definitions are applied to a grid of points, with a value of 1 assigned to points within S_{pup} and a value of 0 assigned to points within $S'_{pup} = S'_{unobsc} \cup S_{sec} \cup S_{struts} \cup S_{tabs}$, the resulting binary array is the telescope pupil representation that, in the absence of any pupil padding, becomes the starting point for the shaped pupil mask optimization. (When pupil padding is added to enforce symmetry and account for potential pupil errors, the set definitions are modified as described in Sec. 2.3 before generating the array that represents the telescope pupil during the shaped pupil mask optimization.) For illustrative purposes, the modeled telescope pupil is shown here in the 0-deg orientation; for the rotated mask configuration, the modeled telescope pupil is rotated 60 deg counterclockwise, as shown in Fig. 3.

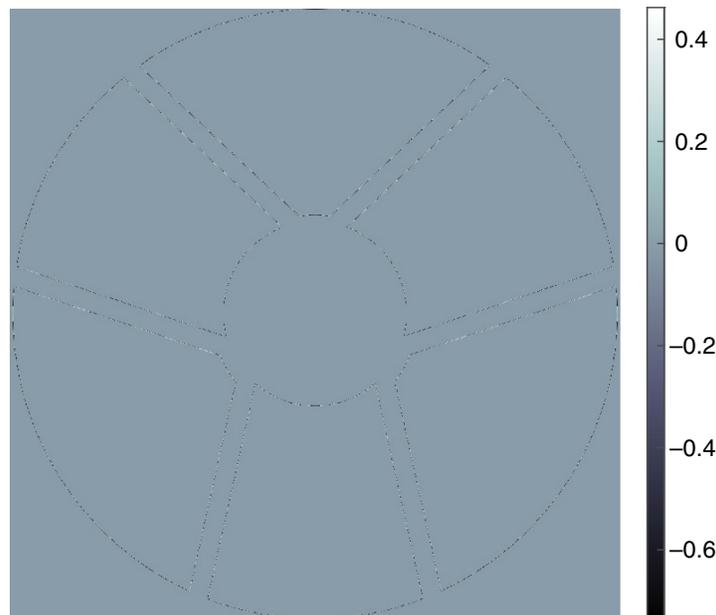


Fig. 6 Difference between the algebraic and PROPER-generated representations of the telescope pupil. The two representations agree very closely, with only minor variations along the edges of the obscurations. These minor variations are due to slight differences in the way that the two approaches interpret what it means to model the pupil: the algebraic representation uses binary edges and covers every pixel of the original pupil map, while the PROPER-generated representation uses gray edges and does not require every pixel of the original pupil map to be covered. If the two representations had been identical, this figure would be uniformly gray.

S_{obsc} is the union of all of the individual obscurations, and the clear area of the pupil is the set of points that fall within S_{unobsc} but outside S_{obsc} , as illustrated in Fig. 5.

To allow the pupil to be rotated arbitrarily, the perimeters of the primary and secondary are described by ellipses of the form

$$\frac{[(x - x_c) \cos \theta + (y - y_c) \sin \theta]^2}{a^2} + \frac{[(x - x_c) \sin \theta + (y - y_c) \cos \theta]^2}{b^2} = 1, \quad (2)$$

where, in the absence of any pupil errors or padding, a is the semiminor axis length, b is the semimajor axis length, θ is the orientation of the desired mask configuration (0 deg for the baseline mask configuration and 60 deg for the rotated mask configuration), and x_c and y_c are the x and y coordinates, respectively, that specify the location of the ellipse center. Because any collected light must land on the primary mirror, S_{unobsc} consists of the set of points that fall within the primary's perimeter:

$$S_{\text{unobsc}} = \left\{ (x, y) : \frac{[(x - x_{c_p}) \cos \theta + (y - y_{c_p}) \sin \theta]^2}{a_p^2} + \frac{[(x - x_{c_p}) \sin \theta + (y - y_{c_p}) \cos \theta]^2}{b_p^2} \leq 1 \right\}, \quad (3)$$

where the subscript p indicates parameter values specific to the primary. For the secondary, the inequality follows the same pattern, but the meaning is reversed; the set names are chosen such that, for any component i of the modeled telescope pupil, S_i is the set of points within (or obscured by) i . The set S_{sec} , therefore, describes the obscuration due to the secondary instead of the clear area outside the secondary:

$$S_{\text{sec}} = \left\{ (x, y) : \frac{[(x - x_{c_{\text{sec}}}) \cos \theta + (y - y_{c_{\text{sec}}}) \sin \theta]^2}{a_{\text{sec}}^2} + \frac{[(x - x_{c_{\text{sec}}}) \sin \theta + (y - y_{c_{\text{sec}}}) \cos \theta]^2}{b_{\text{sec}}^2} \leq 1 \right\}, \quad (4)$$

where the subscript sec indicates values specific to the secondary. The obscurations due to the tabs are defined similarly, but additional logic is required to restrict each of these obscurations to the subset of points within the ellipse that fall between the appropriate struts and outside S_{sec} .

To determine the obscuration due to the struts, it is useful to start by modeling a generic strut as a rectangle of length l and half width w_b centered on the origin. A particular strut can then be modeled by rotating this generic rectangle by an angle $\theta_{\text{strut}} + \theta$, where θ_{strut} is the nominal orientation of the specified strut, and translating the result so that its center is located at $\begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} x_{c_{\text{strut}}} \\ y_{c_{\text{strut}}} \end{bmatrix}$, where $x_{c_{\text{strut}}}$ and $y_{c_{\text{strut}}}$ are the x and y coordinates, respectively, that specify the location of the strut's center in the absence of any pupil rotation. This rectangular strut model is useful for determining the locations of the strut's four corners in the absence of any pupil padding and misalignments. Because the addition of pupil padding can modify the shape of the strut's obscuration so that the long edges are no longer parallel, it is ultimately beneficial to model the strut by using the corners to determine the lines that define the strut's longer sides and specifying the obscuration as the set of points that fall between these lines (truncated by the perimeter of the primary on one side and the secondary on the other). The specific definition for this set varies depending on the particular strut and the overall pupil orientation because these factors affect which corners define the upper of the two lines as well as how to bound the set to prevent it from poking through the opposite side of the secondary.

2.3 Padding the Pupil

Once the telescope pupil model has been constructed, the next step in generating the starting point for the shaped pupil optimization is to add the pupil padding. By padding, we mean that additional regions of the pupil plane are blocked to enforce symmetry (symmetry padding) or to ensure that the shaped pupil mask completely covers the telescope pupil even if errors are present (error padding). The symmetry padding is required to satisfy an assumption made during the

optimization process, and the amount that is needed depends on the pupil orientation (Sec. 2.3.2). The error padding, however, is more discretionary; although we can apply enough error padding to compensate for the maximum expected pupil error, we can also use less if there are other means of compensating for at least part of the potential pupil error (Sec. 2.3.1). Both types of pupil padding may be applied, but the application process varies depending on the type of padding and the method used to generate the telescope pupil model.

2.3.1 Padding for pupil errors

Because the shaped pupil mask is optimized for a specific representation of the telescope pupil, problems can arise when there is a mismatch between the modeled and true telescope pupils. In particular, the contrast and throughput are no longer guaranteed to reach their design values if part of the true telescope pupil is not covered by the shaped pupil mask. This scenario can be caused by manufacturing uncertainties or misalignments, so a conservative approach to preventing the performance degradation is to add padding to the shaped pupil optimization starting point to account for these potential errors up to some specified amount. Although this error padding improves the robustness to pupil errors, it comes at the cost of some of the core throughput. The list of errors and their tolerances therefore plays an important role in the performance of the final design.

For both the baseline and rotated mask configurations, we pad for a total of seven errors. Three of these are due to global misalignments (translations, clocking, and magnification), and three are due to component-level manufacturing uncertainties (for the struts; the secondary, also called the central obscuration; and the tabs). The final error, the rolloff, accounts for the fact that light reflected from the very outer edge of the primary is unusable due to the process by which the mirror was ground and polished. The conservative tolerance values assigned to each of these errors are identical for the baseline and rotated mask configurations with one exception: the baseline mask configuration uses a slightly higher value for the translation tolerance based on a more conservative interpretation of a lateral pupil misalignment requirement (Table 1). This choice was made to slightly increase the robustness of the baseline mask configuration and reduce its technical risk. For the rotated mask configuration, which is contributed by ExEP so that it can be available for use if scheduling and funding constraints allow and is not part of the technology demonstration, the smaller value was selected to slightly increase the core throughput (by two-tenths of a percentage point).

The error padding is applied to the shaped pupil optimization starting point in one of two ways, both of which account for worst-case scenarios in which all of the errors are present simultaneously at the most extreme values specified by the tolerances. For the baseline mask

Table 1 Tolerances used for error padding. The diameter D is the full aperture diameter including the rolloff, not the diameter of the smaller aperture once the rolloff is masked.

Error type	Tolerance	
	More conservative (baseline)	Conservative (rotated)
Rolloff	$-0.5\%D$, increased based on max expected translation, clocking, and magnification	$-0.5\%D$, increased based on max expected translation, clocking, and magnification
Translation	$\pm 0.1\%D$	$\pm 0.062\%D$
Clocking	± 6 mrad	± 6 mrad
Magnification	$[+0.75 \ -1.05]\%D$	$[+0.75 \ -1.05]\%D$
Struts	$0.129\%D$	$0.129\%D$
Central obscuration	$0.141\%D$	$0.141\%D$
Tabs	$0.141\%D$	$0.141\%D$

configuration, the error padding is applied by compiling multiple versions of the array that represents the modeled telescope pupil. First, the component-level error padding and the rolloff padding are applied by modifying the telescope pupil model; the component-level error padding is applied by increasing the widths of the struts and the sizes of the secondary and tabs by the amounts specified in the tolerances, and the rolloff padding is applied by decreasing the radius of the primary. The global error padding is then applied by using PROPER to generate versions of the modified telescope pupil for a range of possible translation, clocking, and magnification combinations and taking the union of the padding cases.

For the rotated mask configuration, the error padding is applied by modifying the sets of equations used to generate the individual elements of the telescope pupil model. Following the conventions established for the baseline mask configuration, we first account for the component-level errors and rolloff and then pad for magnification, clocking, and translation. With all of the types of error padding incorporated, the telescope pupil model is no longer a simple collection of ellipses and rectangles, but it is still relatively straightforward.

The first modifications are made to the definition of $S_{\text{unobs}}c$ with the goal of padding the primary's perimeter so that the unusable area is always completely blocked by the shaped pupil mask even in the presence of pupil errors. If the primary had been modeled as the set of points within a circle centered at the origin, the modifications would have been as simple as reducing the circle's radius to a value that accounts for the rolloff tolerance as well as the magnification and translation tolerances. Conservatively, this value is

$$r_p = (1 + m_-)(r_{p_0} - \iota) - t\sqrt{2}, \quad (5)$$

where r_{p_0} is the radius of the unpadded circle, m_- is the negative magnification tolerance, t is the absolute value of the translation tolerance, and ι is the rolloff tolerance. The negative magnification tolerance is used because the unusable portion of the primary comes into view as the telescope pupil is demagnified, and the translation tolerance is multiplied by the square root of two to reflect a conservative interpretation that the telescope pupil can translate by the tolerance amount in x and y simultaneously.

However, because the primary is modeled as the set of points within a general ellipse, we must also account for clocking and the projection effects that modify the shape of the unusable region. We pad for clocking errors by considering the two boundary-case ellipses in addition to the single ellipse that defined the primary. These boundary-case ellipses are rotated by the maximum and minimum values specified by the clocking tolerance, and the sets of points $S_{p_{\pm}}$ within these ellipses are defined as

$$S_{p_{\pm}} = \left\{ (x, y) : \frac{[(x - x_{c_{p_{\pm}}}) \cos(\theta \pm \theta_c) + (y - y_{c_{p_{\pm}}}) \sin(\theta \pm \theta_c)]^2}{\left[(1 + m_-) \left(a_p - \frac{a_p}{r_{p_0} \iota} \right) - t\sqrt{2} \right]^2} + \frac{[(x - x_{c_{p_{\pm}}}) \sin(\theta \pm \theta_c) + (y - y_{c_{p_{\pm}}}) \cos(\theta \pm \theta_c)]^2}{\left[(1 + m_-) \left(b_p - \frac{b_p}{r_{p_0} \iota} \right) - t\sqrt{2} \right]^2} \leq 1 \right\}, \quad (6)$$

where the signs in the subscripts of S_p , x_{c_p} , and y_{c_p} reflect the sign of the clocking value used to define the boundary ellipse; θ_c is the absolute value of the clocking tolerance; the factors of $\frac{a_p}{r_{p_0}}$ and $\frac{b_p}{r_{p_0}}$ scale the rolloff tolerance to the appropriate amount for each axis of the boundary ellipse; and $x_{c_{p_{\pm}}}$ and $y_{c_{p_{\pm}}}$ are the coordinates locating the center of the boundary ellipse:

$$\begin{bmatrix} x_{c_{p_{\pm}}} \\ y_{c_{p_{\pm}}} \end{bmatrix} = (1 + m_-) \begin{bmatrix} \cos(\theta \pm \theta_c) & -\sin(\theta \pm \theta_c) \\ \sin(\theta \pm \theta_c) & \cos(\theta \pm \theta_c) \end{bmatrix} \begin{bmatrix} x_{c_p} \\ y_{c_p} \end{bmatrix}. \quad (7)$$

The clear area of the unobscured primary is then the intersection of the sets of points contained within each of the boundary ellipses and the original ellipse [Fig. 7(a)]:

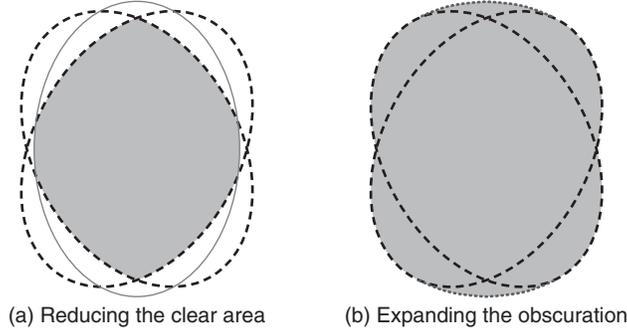


Fig. 7 Padding for regions defined by ellipses. Because the primary and secondary are both modeled as ellipses instead of circles, their padding process must account for clocking. This is done in part by considering two boundary ellipses (drawn with dashed lines): one that has the maximum expected counterclockwise rotation and one that has the maximum expected clockwise rotation. For the primary, the padding reduces the clear area to the intersection of the sets of points within the two boundary ellipses and the original, unrotated ellipse; the result, S_{unobsc} , is illustrated by the shaded region in (a). For the secondary, the padding increases the size of the obscuration to cover all of the possible orientations spanned by the boundary ellipses. The result, S_{sec} , shown by the shaded region in (b), consists of the union of the sets of points within the two boundary ellipses as well as the circular sectors swept out by the major axis as it is rotated from one extreme to the other. (The arcs traced by the endpoints of the major axis are drawn with dotted lines.) Note that, for illustrative purposes, the rotations and eccentricity in (a) and (b) are greatly exaggerated.

$$S_{\text{unobsc}} = S_{\text{unobsc}_+} \cap S_{\text{unobsc}_-} \cap S_{\text{unobsc}_0}, \quad (8)$$

where S_{unobsc_0} represents the set of points within the original, unrotated ellipse [Eq. (3)].

Although the secondary is also modeled as the set of points within a general ellipse, the modifications to S_{sec} are the opposite of the modifications to S_{unobsc} because S_{sec} is an obscuration instead of a clear area. Instead of reducing the size of the boundary ellipses and taking the intersection of the sets of points that fall within each ellipse, we now increase the size of the boundary ellipses and take the union of the sets of points that fall within these ellipses and all of the clocking cases that they bound [Fig. 7(b)]. Because the endpoints of the major axis of the secondary's ellipse trace arcs as the ellipse is rotated from one boundary case to the other, we can represent the union of the sets of points that fall within all clocking cases as the union of the sets of points S_{sec_\pm} that fall within the boundary ellipses and the sets of points S_{sec_1} and S_{sec_2} that fall within the circular sectors swept out by the major axis:

$$S_{\text{sec}} = S_{\text{sec}_+} \cup S_{\text{sec}_-} \cup S_{\text{sec}_1} \cup S_{\text{sec}_2}. \quad (9)$$

The sets S_{sec_\pm} are defined as

$$S_{\text{sec}_\pm} = \left\{ (x, y) : \frac{[(x - x_{c_{\text{sec}_\pm}}) \cos(\theta \pm \theta_c) + (y - y_{c_{\text{sec}_\pm}}) \sin(\theta \pm \theta_c)]^2}{[(1 + m_+)(a_{\text{sec}} + \delta r_{\text{sec}}) + t\sqrt{2}]^2} + \frac{[(x - x_{c_{\text{sec}_\pm}}) \sin(\theta \pm \theta_c) + (y - y_{c_{\text{sec}_\pm}}) \cos(\theta \pm \theta_c)]^2}{[(1 + m_+)(b_{\text{sec}} + \delta r_{\text{sec}}) + t\sqrt{2}]^2} \leq 1 \right\}, \quad (10)$$

where, as before, the signs of S_{sec} , $x_{c_{\text{sec}}}$, and $y_{c_{\text{sec}}}$ reflect the sign of the clocking value used to define the boundary ellipse; m_+ is the positive magnification tolerance; δr_{sec} is the absolute value of the alignment tolerance for the secondary; and $x_{c_{\text{sec}_\pm}}$ and $y_{c_{\text{sec}_\pm}}$ are the coordinates locating the center of the boundary ellipse:

$$\begin{bmatrix} x_{c_{\text{sec}_\pm}} \\ y_{c_{\text{sec}_\pm}} \end{bmatrix} = (1 + m_+) \begin{bmatrix} \cos(\theta \pm \theta_c) & -\sin(\theta \pm \theta_c) \\ \sin(\theta \pm \theta_c) & \cos(\theta \pm \theta_c) \end{bmatrix} \begin{bmatrix} x_{c_{\text{sec}}} \\ y_{c_{\text{sec}}} \end{bmatrix}. \quad (11)$$

To define S_{sec_1} and S_{sec_2} , we first define a boundary line between the two circular sectors swept out by the major axis:

$$y_{\text{bound}} = \tan \theta (x - \cos \theta) + \sin \theta. \quad (12)$$

Then, we consider each circular sector in turn. Although the specific equations and inequalities that define the circular sectors depend on θ , the overall method for defining S_{sec_1} and S_{sec_2} does not. For S_{sec_1} , we locate the endpoints $V_{\text{sec}_{1\pm}}$ of the traced arc, define the lines $\overleftrightarrow{V_{\text{sec}_{1+}}O}$ and $\overleftrightarrow{V_{\text{sec}_{1-}}O}$ that connect these endpoints to the origin O , and identify the circle of which arc $V_{\text{sec}_{1+}}V_{\text{sec}_{1-}}$ is a part. Then, we define S_{sec_1} as the set of points within this circle that fall between $\overleftrightarrow{V_{\text{sec}_{1+}}O}$ and $\overleftrightarrow{V_{\text{sec}_{1-}}O}$ and are on the correct side of y_{bound} . The set S_{sec_2} is defined similarly.

The three sets S_{tabs_1} , S_{tabs_2} , and S_{tabs_3} that represent the obscurations due to the tabs are modified in the same manner as S_{sec} . The additional logic that bounds S_{tabs_1} , S_{tabs_2} , and S_{tabs_3} between the appropriate struts is unaffected by the padding, although it does require the boundaries of the padded struts to be defined first. The specific inequalities and equations are orientation-dependent and vary by tab.

The six sets that represent the obscurations due to the individual struts are the most complicated to define because there are multiple details that depend on the specific strut and the pupil orientation. The first step is to follow the procedure outlined at the end of Sec. 2.2, with the strut half width increased by δw_b to account for the strut tolerance. For simplicity, let $\overleftrightarrow{V_{\text{strut}_1}V_{\text{strut}_2}}$ and $\overleftrightarrow{V_{\text{strut}_3}V_{\text{strut}_4}}$ be the boundary lines that define the longer sides of strut i , and let S_{strut_i} be the set of points that fall between $\overleftrightarrow{V_{\text{strut}_1}V_{\text{strut}_2}}$ and $\overleftrightarrow{V_{\text{strut}_3}V_{\text{strut}_4}}$, bounded by the perimeter of the primary on one end and the perimeter of the secondary on the other (Fig. 8). Then, we pad for each of the global errors in turn by identifying the boundary cases that move $\overleftrightarrow{V_{\text{strut}_1}V_{\text{strut}_2}}$ and $\overleftrightarrow{V_{\text{strut}_3}V_{\text{strut}_4}}$ the farthest outside S_{strut_i} and updating the definitions of the lines and S_{strut_i} accordingly. For example, to pad for magnification errors, we magnify one line by m_+ and the other by m_- , choosing the case that increases the strut width. Then, to pad for clocking errors, we rotate one of the new lines by θ_c and the other by $-\theta_c$, again choosing the case that increases the size of the obscuration. Finally, we pad for translation errors by shifting the resulting lines by $\pm t$ in x and $\pm t$ in y , with the sign for each shift chosen based on which moves the specified boundary line away from S_{strut_i} . The signs chosen for one line will be the opposite of those for the other.

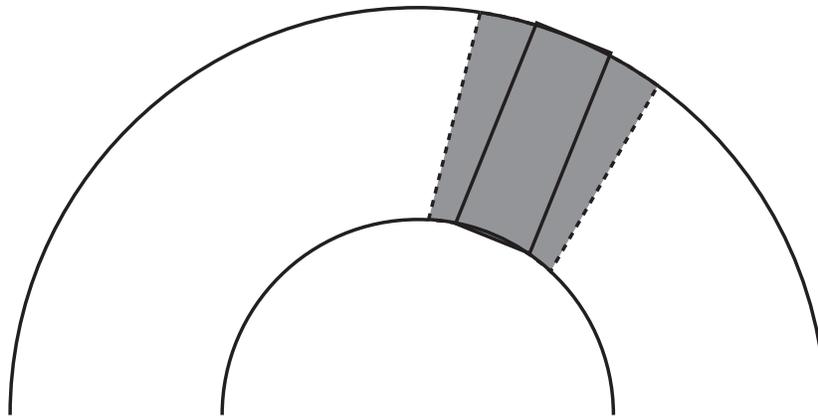


Fig. 8 Padding the struts. For simplicity, a strut is initially modeled as a rectangle placed between the primary and secondary, as shown for an example strut by the solid lines. To define the obscuration due to a padded strut, it helps to replace this rectangle with a pair of lines (dashed) that define the longer sides; the obscuration is then defined as the set of points (shaded) between these lines and between the perimeters of the primary and secondary, with an additional restriction to ensure that the obscuration is limited to the correct side of the pupil. As each type of padding is applied, the (dashed) lines move farther apart, and their angles also change once the clocking padding is applied. Note that the strut width and padding are greatly exaggerated for illustrative purposes.

2.3.2 Padding for symmetry

Because the shaped pupil mask optimization algorithm requires that the pupil be symmetric about the y axis (as explained in Sec. 2.1), the naturally asymmetric telescope pupil model is forced to be symmetric through the addition of pupil padding. This padding is added by manipulating the array that represents the modeled telescope pupil, so, for the analytically modeled telescope pupil, the symmetry padding is applied after any error padding. The set of pixels that correspond to the pupil asymmetries, S_{asym} , is identified by subtracting a mirror image of the left half of the pupil from the right half. Padding is then added to the right half of the pupil by setting the pixel value equal to zero wherever $S_{\text{asym}} > 0$. Similarly, padding is added to the left half of the pupil by setting the pixel value equal to zero wherever the mirror image of S_{asym} is negative. The two padded halves are then recombined to form the symmetrized version of the telescope pupil. Although it is possible to achieve the same result by simply padding one side of the pupil and combining the result with its mirror image, we track the amount of padding added to each side of the pupil in order to transparently explain where each pixel of padding originates.

Unlike the error padding, the symmetry padding is affected by the pupil orientation due to the lack of perfect rotational symmetry. (As noted in Sec. 2.1, we optimize shaped pupil masks with the pupil rotated by the specified amount when needed to make the problem tractable; in reality, the focal plane mask, Lyot stop, and dark hole are rotated.) The differences are noticeable even between the 0-deg, 60-deg, and -60 -deg orientations, which are all nearly symmetric about the y axis (Fig. 9). The ± 60 -deg orientations have significantly more symmetry padding than the 0-deg orientation for three reasons, all of which are due to the CGI being mounted slightly off-axis: the ellipses that describe the outer diameters of the primary, the secondary, and the tabs are no longer symmetric about the y axis; the two tabs that do not fall on the y axis for the ± 60 -deg orientations are no longer fit to the same ellipse; and the strut pairs are less symmetric. Ultimately, these changes lead to a throughput loss of about half a percentage point compared with the throughput of the equivalent 0-deg mask configuration. For the rotated mask configuration, the $+60$ -deg orientation was selected over the -60 -deg orientation due to its slightly higher symmetry and therefore slightly better throughput.

3 Flight Design Selection

While the shaped pupil mask optimization determines the best shaped pupil mask for a given Lyot stop, the design survey identifies the Lyot stop that leads to the highest-throughput design. The Lyot stop for the bowtie shaped pupil coronagraph is a bowtie-shaped mask rotated 90 deg relative to the focal plane mask; this orientation aligns the Lyot stop with the reflective regions of the shaped pupil mask, and the bowtie shape is well-matched to the regions illuminated by off-axis sources³² (Fig. 10). A Lyot stop is specified by three main parameters: the inner and outer diameters of the opening and the opening angle [Fig. 11(a)]. These parameters determine the basic perimeter of the bowtie, which consists of four line segments and four arcs. To improve manufacturability and reduce the risk of damage during launch, the basic bowtie shape is modified slightly: the eight sharp corners where a line segment and an arc intersect are replaced by fillets [Fig. 11(b)]. Each fillet is a section of a circle of radius dr that is tangent to the line segment at one end and tangent to the arc at the other, allowing the perimeter to smoothly transition from the line segment to the arc. Although dr is adjustable, it is held constant unless being optimized explicitly.

For the design surveys, we considered Lyot stops with inner diameters of $20\%D$ to $43\%D$, outer diameters of $76\%D$ to $96\%D$, and opening angles of 80 deg to 100 deg. (The diameter D is the full aperture diameter, not the smaller diameter once the rolloff is masked.) Due to time constraints, this parameter space was sampled coarsely at first and then successively finely until a maximum-throughput design could be identified in a parameter cube of resolution $1\%D$ in the inner and outer diameters and 1 deg in the opening angle; at this resolution, the throughput changes very slowly around the maximum point, and additional refinement is unnecessary. Although the design surveys for the baseline and rotated mask configurations (Secs. 3.1 and 3.2, respectively) all followed this basic procedure, the specific process varied because the two mask configurations were designed semi-independently (with the baseline mask configuration

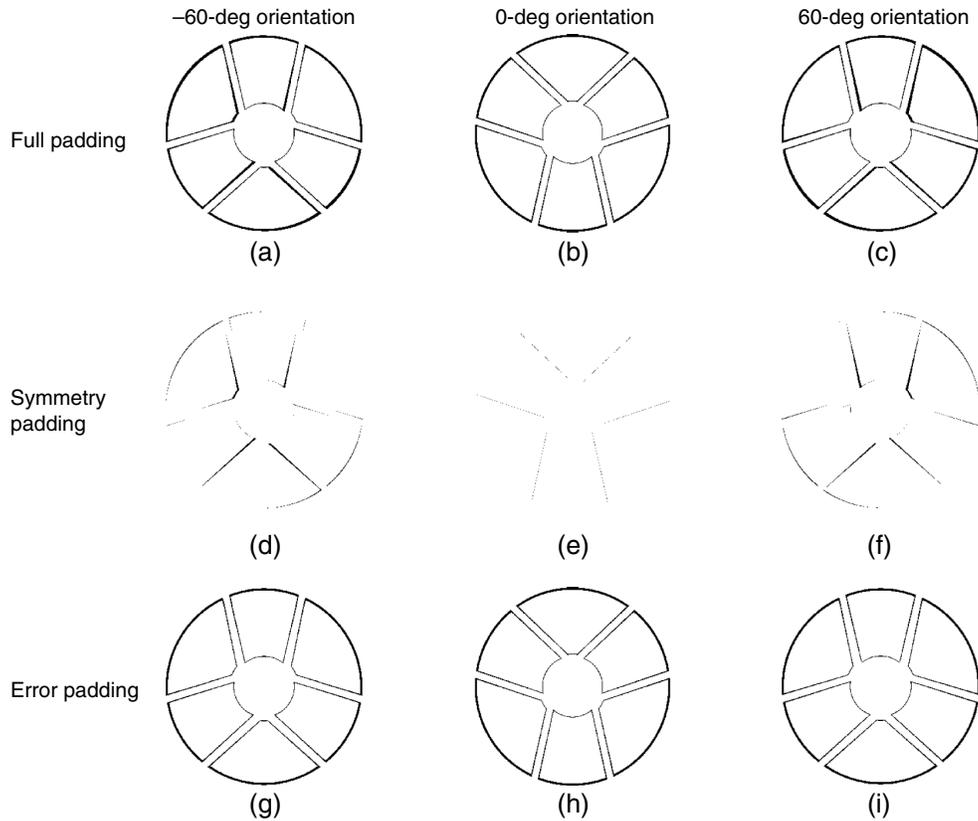


Fig. 9 Padding variations due to orientation-dependent asymmetries. There are two types of pupil padding: padding to enforce y -axis symmetry and padding to account for possible pupil errors. The total pupil padding depends on the choice of mask configuration orientation because the pupil is rotated for the shaped pupil optimization instead of the other masks. Although the amount of error padding is consistent for 0-, 60-, and -60 -deg orientations of the Roman pupil, the symmetry padding varies noticeably. The ± 60 -deg orientations require more symmetry padding than the 0-deg orientation, with the -60 -deg orientation needing the most, and the additional padding leads to a throughput loss of about half a percentage point relative to the throughput of the 0-deg mask configuration.

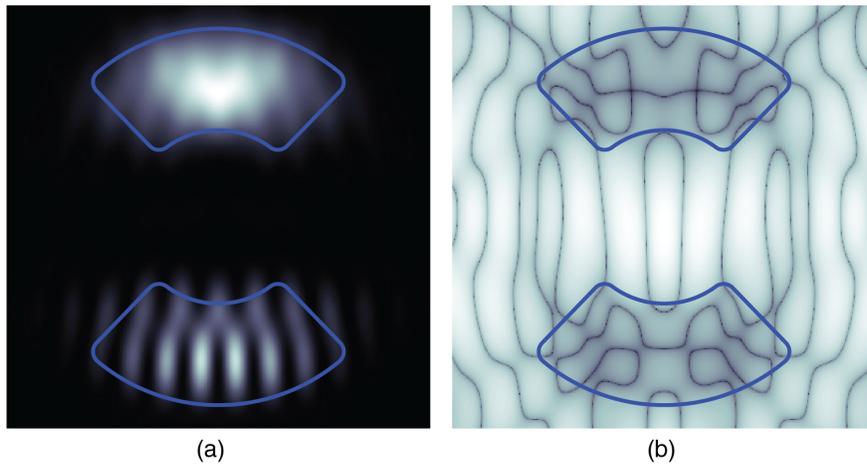


Fig. 10 Rationale behind the Lyot stop shape. A bowtie-shaped Lyot stop captures most of the light from a planet (a), while blocking most of the starlight (b). To emphasize the regions illuminated by the off-axis light as well as the rejection of the on-axis light, (a) is shown with a linear scale, while (b) is shown with a logarithmic scale.

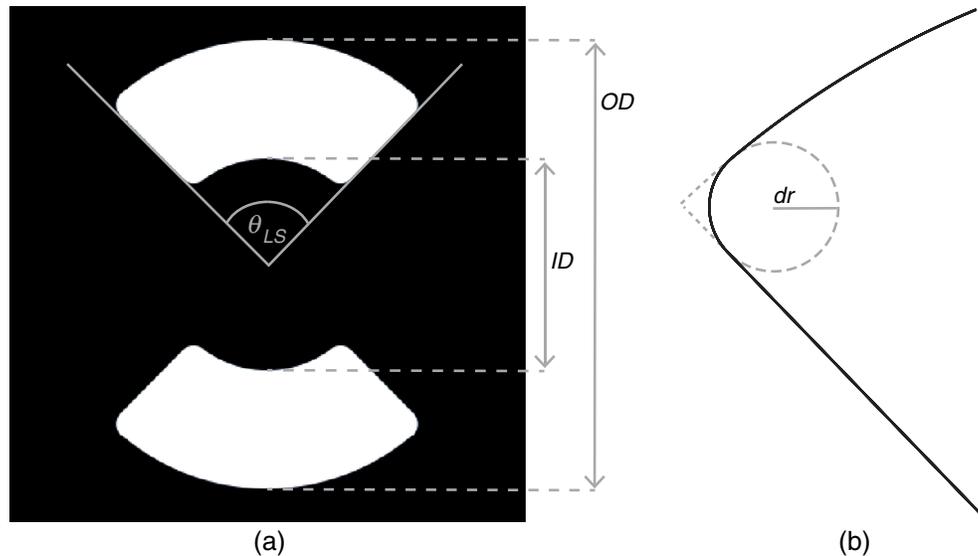


Fig. 11 Lyot stop geometry. The Lyot stop is a bowtie-shaped mask rotated 90 deg relative to the focal plane mask and specified by the inner and outer diameters (ID and OD , respectively) of the open areas and the opening angle θ_{LS} , as shown in (a). The sharp corners of the bowtie are replaced with fillets; instead of transitioning abruptly from a line segment to an arc or vice versa, the perimeter traces a small arc that is tangent to the line segment at one end and tangent to the large arc at the other (b). This small arc is a section of a circle of radius dr , and this radius is called the radius of curvature for the fillets.

primarily designed at JPL and the rotated mask configuration primarily designed at Princeton) and under fluctuating time constraints. Despite the differences in the design methodologies, the flight designs for the baseline and rotated mask configurations (Sec. 3.3) are strikingly similar, as would be expected if the telescope pupil were perfectly rotationally symmetric.

3.1 Design Surveys for the Baseline Mask Configuration

To quickly narrow the parameter space for the baseline mask configuration, the design process began with a preliminary design survey that coarsely sampled the set of possible Lyot stops while generating shaped pupil masks with a lower resolution than the full manufacturing resolution. Although the lower-resolution (250 pixel \times 250 pixel) shaped pupil masks are not purely binary as required, they offer comparable performance to full-resolution (1000 pixel \times 1000 pixel) shaped pupil masks in simulations while also being significantly faster to optimize. (The shaped pupil mask optimization routine does not guarantee that the resulting shaped pupil mask will be perfectly binary, but when the shaped pupil mask is optimized at the full manufacturing resolution, the contrast does not degrade when the pixel values are rounded to produce a purely binary mask.) The results of this survey suggest that the highest-throughput design corresponds to a Lyot stop with an inner diameter of $41\%D$ to $42\%D$, an outer diameter of $89\%D$, and an opening angle of 87 deg to 88 deg (Figs. 12 and 13). In the interest of time, the parameter space for subsequent design surveys was therefore narrowed to inner diameters of $39\%D$ to $43\%D$, outer diameters of $87\%D$ to $92\%D$, and an opening angle of 88 deg; the opening angle was fixed due to time constraints and the minimal benefit of varying it near the maximum-throughput point.

In response to concerns about launch stresses, we conducted an additional low-resolution design survey to determine the radius of curvature for the Lyot stop fillets. The results of this survey indicate that although a radius of $2\%D$ leads to the highest throughput, the radius can be increased to $3\%D$ without a significant throughput loss (Fig. 14). We therefore selected a radius of $3\%D$ to reduce the risk of damage to the Lyot stop during launch. For this radius and an opening angle of 88 deg, a Lyot stop with an inner diameter of $41\%D$ and an outer diameter of $89\%D$ leads to the highest throughput for the baseline mask configuration, as confirmed by a final high-resolution design survey (Fig. 15).

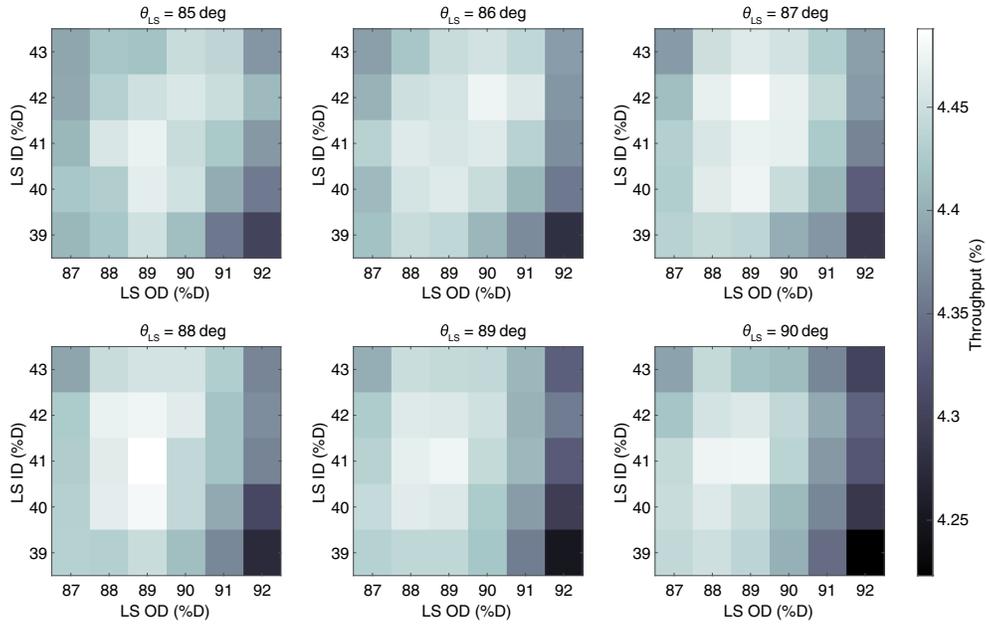


Fig. 12 Preliminary low-resolution design survey results for the baseline mask configuration. The results of this survey, which are used to narrow the parameter space for subsequent surveys, suggest that the highest-throughput design corresponds to a Lyot stop with an inner diameter of 41% D or 42% D , an outer diameter of 89% D , and an opening angle of 87 deg to 88 deg.

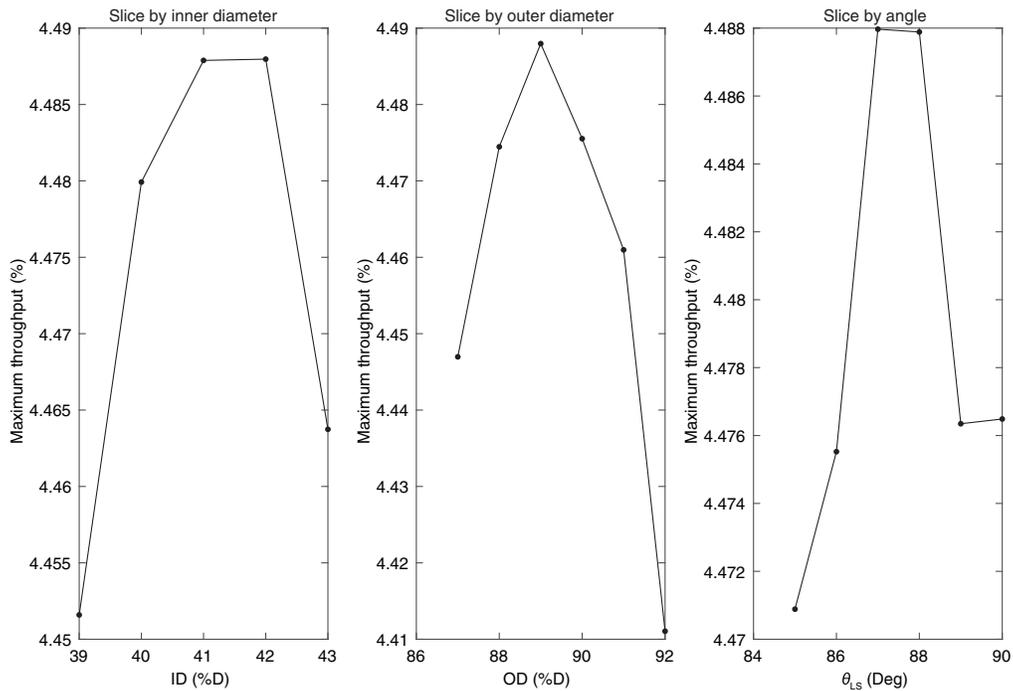


Fig. 13 Identification of the maximum-throughput design point for the preliminary low-resolution design survey for the baseline mask configuration. Because multiple design points can have a similar throughput, it often helps to look at a set of three plots: the maximum throughput as a function of inner diameter, the maximum throughput as a function of outer diameter, and the maximum throughput as a function of angle. For the preliminary low-resolution design survey for the baseline mask configuration, an inner diameter of 42% D and an opening angle of 87 deg lead to the maximum throughput, but the maximum throughput changes very slowly with variations in these two parameters. An opening angle of 88 deg was selected for subsequent baseline design surveys.

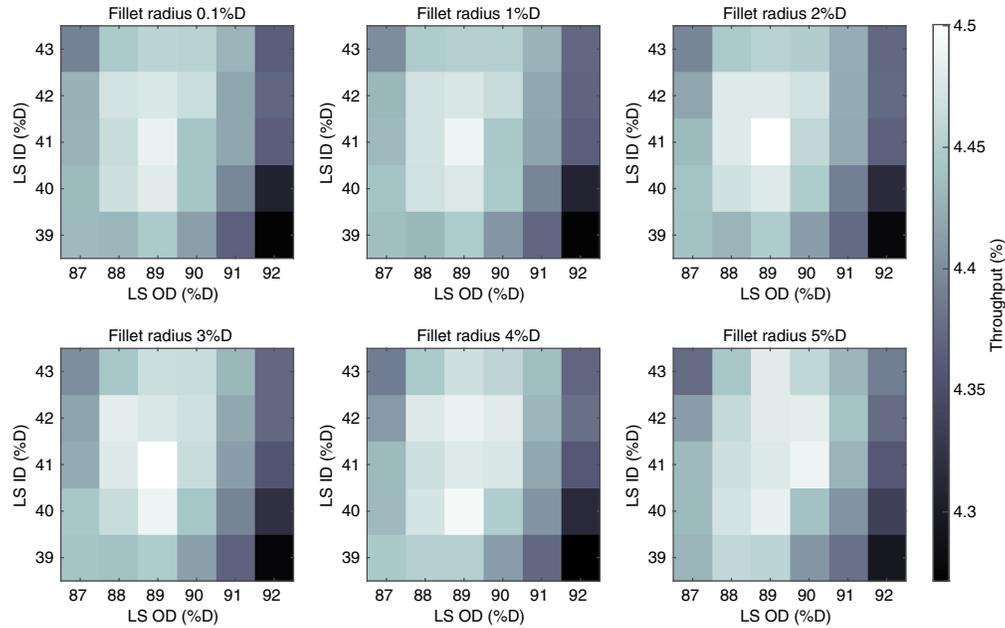


Fig. 14 Results from the low-resolution Lyot stop fillet radius of curvature survey. For the radius of curvature survey, Lyot stops with inner diameters of $39%D$ to $43%D$, outer diameters of $87%D$ to $92%D$, and radii of curvature of $0.1%D$ to $5%D$ were considered. The opening angle was restricted to 88 deg due to time constraints and the minimal impact of varying the angle near the maximum-throughput point of the parameter space. A Lyot stop with an inner diameter of $41%D$, an outer diameter of $89%D$, and a fillet radius of $2%D$ led to the highest throughput (4.501%), but increasing the fillet radius to $3%D$ negligibly affected the throughput, lowering it slightly to 4.499% . The larger radius of curvature was therefore selected due to concerns about launch stresses, and it was used for both the baseline and rotated mask configurations.

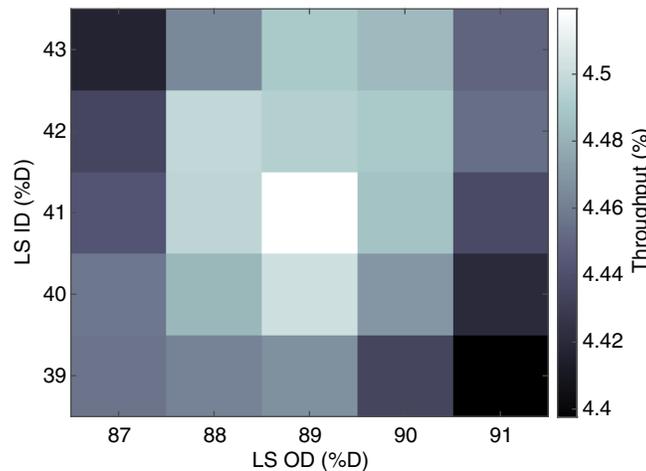


Fig. 15 Results of the final, full-resolution design survey for the baseline mask configuration. For a fixed opening angle of 88 deg and a fillet radius of $3%D$, the highest-throughput design used a Lyot stop with an inner diameter of $41%D$ and an outer diameter of $89%D$. This design was therefore selected for the baseline mask configuration.

3.2 Design Surveys for the Rotated Mask Configuration

The design surveys for the rotated mask configuration were performed concurrently with the design surveys for the baseline mask configuration and followed a similar process, incorporating information from the baseline design survey results as needed. For example, once the final value for the radius of curvature for the Lyot stop fillets, dr , was determined from the results of a

design survey for the baseline mask configuration (Fig. 14), the value of dr used for the rotated mask configuration was updated to match. All of the design surveys for the rotated mask configuration generated full-resolution shaped pupil masks, and the differences from one survey to the next consisted of updates to the telescope pupil, tolerances, and other design parameters. Due to time constraints, the parameter space (consisting of possible values for the inner and outer diameters of the Lyot stop and the opening angle) was explored iteratively with successively targeted mini surveys until the maximum-throughput point could be identified, and the final design survey was initially centered on the maximum-throughput point of the penultimate survey (Figs. 16 and 17), which differed only by using a slightly more conservative translation tolerance. The final survey produced two designs with equivalent throughputs (Figs. 18 and 19), but the degeneracy was broken once dr was updated to its final value. Time constraints prohibited running a separate design survey to incorporate this last update, so the shaped pupil masks for the two highest-throughput point designs were reoptimized with updated Lyot stops; the resulting design with the higher throughput was selected for the rotated mask configuration.

3.3 Final Designs for the Baseline and Rotated Mask Configurations

The flight designs for the baseline and rotated mask configurations are very similar because the orientation of the dark hole relative to the struts is nearly equivalent for the two cases. The shaped pupil masks look almost like rotated versions of each other (Fig. 20), and the Lyot stops for the two mask configurations use very similar parameters: an inner diameter, outer diameter, and opening angle of $41\%D$, $89\%D$, and 88 deg, respectively, for the baseline mask configuration and $41\%D$, $87\%D$, and 89 deg, respectively, for the rotated mask configuration. However, while the two designs are very similar, they have different throughputs; the baseline mask configuration has a core throughput of 4.50% at $4 \frac{\lambda}{D}$, whereas the rotated mask configuration has a core throughput of 3.89% at $4 \frac{\lambda}{D}$. [For reference, the throughput for the CGI input pupil alone is 34.4% when calculated using Eq. (1); note that this definition neglects all reflectivities and transmittivities and accounts only for the diffraction of the pupil features.] The difference in throughput is due to differences in the amount of pupil padding: the starting point for the shaped pupil optimization for the rotated mask configuration required more symmetry padding

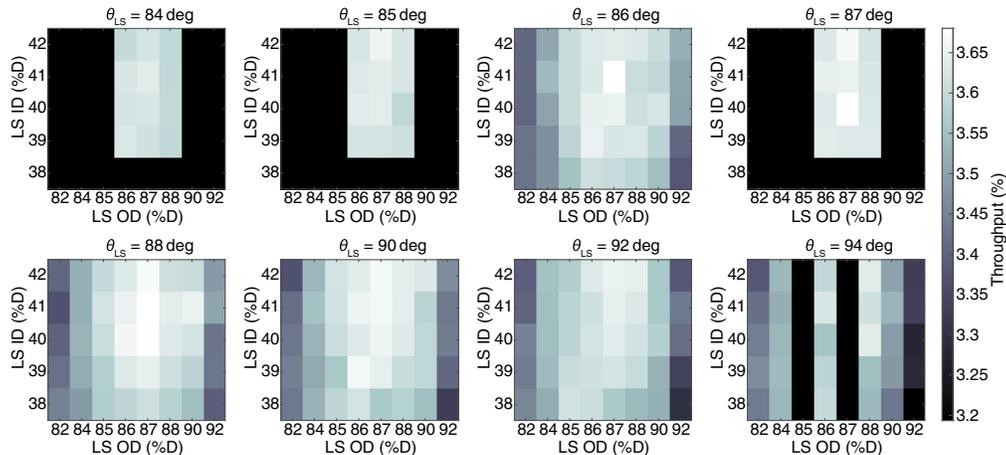


Fig. 16 Results of the penultimate design survey for the rotated mask configuration. Due to time constraints, the parameter space was explored iteratively with successively targeted mini surveys; the regions shown in black were not evaluated. The survey concluded when the maximum-throughput point could be identified in the center of a parameter space cube with a resolution of $1\%D$ in the inner and outer diameters and 1 deg in the opening angle. This point corresponded to a Lyot stop with an inner diameter of $40\%D$, an outer diameter of $87\%D$, and an opening angle of 87 deg. However, the throughput was only negligibly smaller for a Lyot stop with an inner diameter of $41\%D$, an outer diameter of $87\%D$, and an opening angle of 86 deg. This survey used a more conservative translation tolerance than the final survey, and the fillet radius was updated after both surveys concluded.

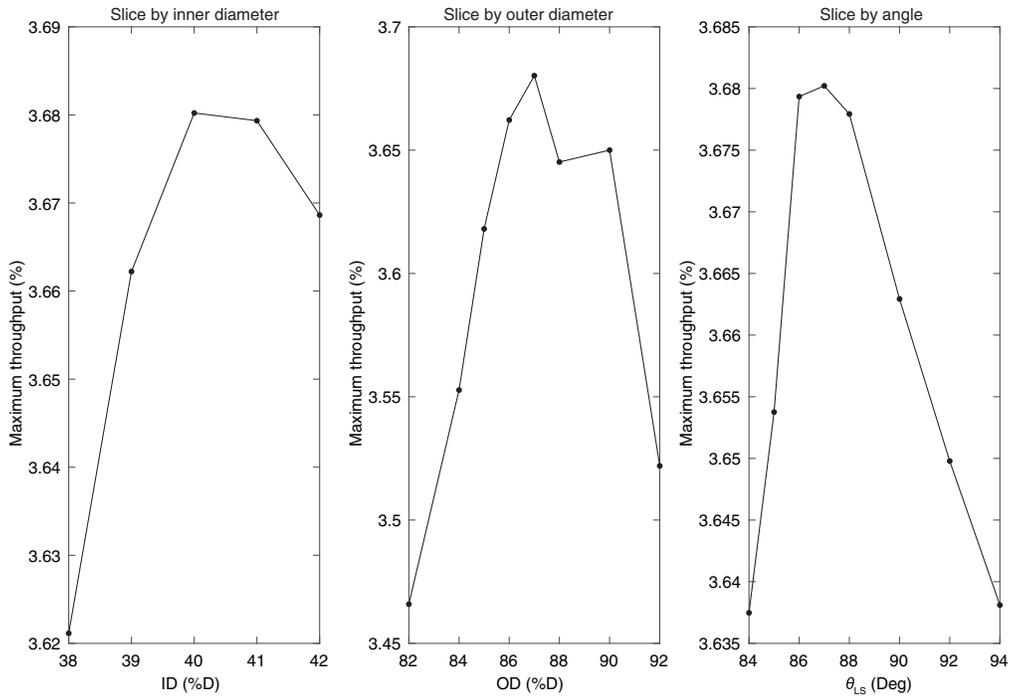


Fig. 17 Identification of the maximum-throughput design point for the penultimate design survey for the rotated mask configuration. Although a Lyot stop with an inner diameter of $40\%D$, an outer diameter of $87\%D$, and an opening angle of 87 deg led to the highest throughput, the maximum throughput changes very slowly for small variations in these three parameters. The changes are particularly slow for variations in the opening angle.

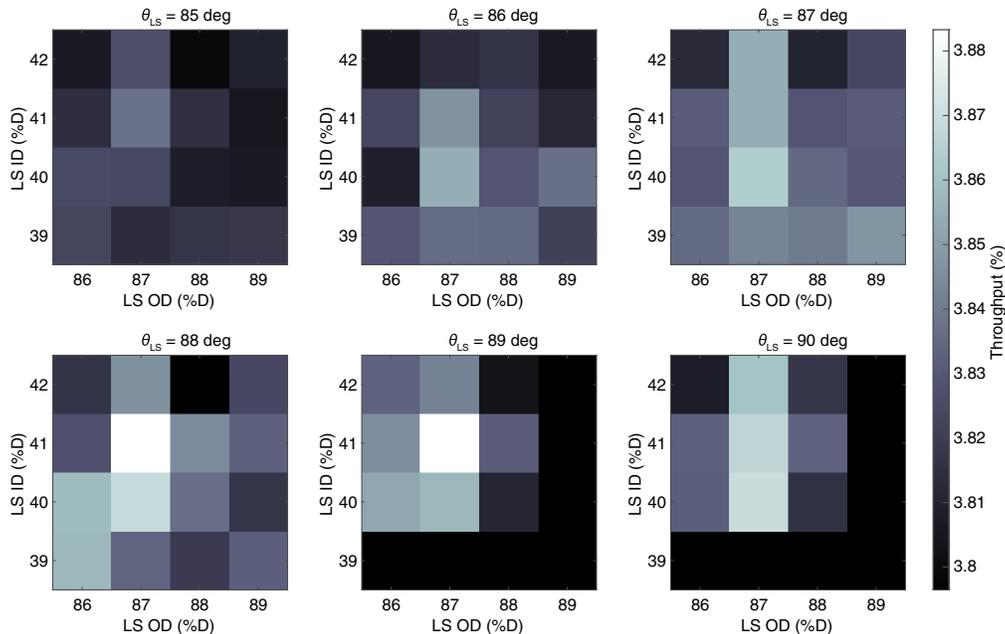


Fig. 18 Results of the final design survey for the rotated mask configuration. Due to time constraints, the final design survey was initially centered around the highest-throughput point of the penultimate design survey, which differed only by using a more conservative translation tolerance. The surveyed portion of the parameter space was then modified with successively targeted mini surveys; the regions shown in black were not evaluated. Two Lyot stops differing only in the opening angle produced the maximum throughput. However, this degeneracy broke in favor of the larger opening angle once the fillet radius was updated to its final value for the two point designs.

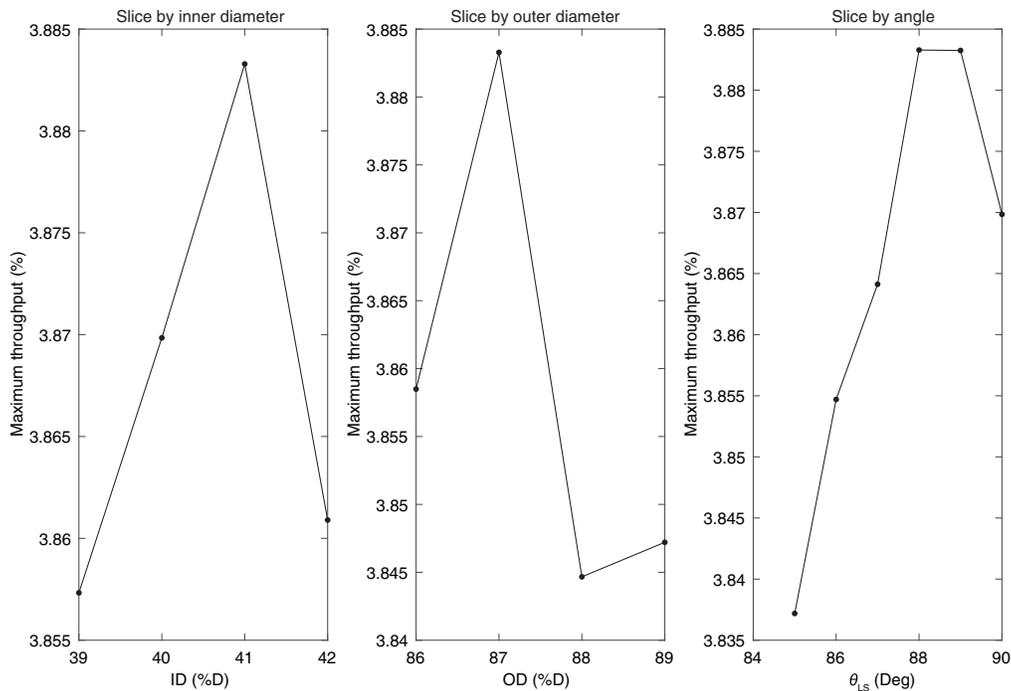


Fig. 19 Identification of the maximum-throughput design point for the final design survey for the rotated mask configuration. Although it is clear which inner diameter and outer diameter lead to the highest-throughput design, there are two equivalent choices for the opening angle. Updating the fillet radius to its final value breaks the degeneracy in favor of the larger opening angle, although the difference is very slight (0.03 percentage points of throughput).

(Sec. 2.3.2), and the aperture loss due to the additional symmetry padding was not fully offset by using a slightly less conservative translation tolerance. A full-plane optimization that did not require symmetry likely would have produced baseline and rotated designs with more comparable throughputs because there would have been no need for the symmetry padding, but this type of optimization was not tractable due to the large design space.

To briefly comment on the integration time required for $\text{SNR} \sim 5$ spectroscopy with the baseline mask configuration, we first note that specific integration time estimates depend strongly on how conservatively the CGI performance is modeled. The most favorable reflected light planets observable by the CGI are very challenging targets; they have separations near the inner working angle and flux ratios of a few 10^{-9} at best, and they are near the performance limit of the CGI, where contributions from random noise and speckle noise are both significant. There may be very few reflected light planets (such as HD 219134 h and 47 Uma c) for which $\text{SNR} \sim 5$ spectroscopy is possible, with integration time estimates ranging from about one day to several weeks. For the interested reader, the Roman Coronagraph Exposure Time Calculator⁴⁰ is a publicly available tool for estimating the integration times for different planetary system configurations.

4 Compensating for Excessive Errors

The flight designs for the baseline and rotated mask configurations are designed to mask the telescope pupil even with errors up to the amounts specified by the tolerances. However, these tolerances are three-sigma tolerances instead of not-to-exceed values, and it is possible (though not expected) for the pupil errors to be large enough that the telescope pupil is partially visible behind the shaped pupil mask, leading to a degradation in the contrast. Although using more pupil padding could have increased the robustness to pupil errors, it is preferable to compensate for any excessive errors by using the DMs that are part of the wavefront sensing and control system for the CGI.

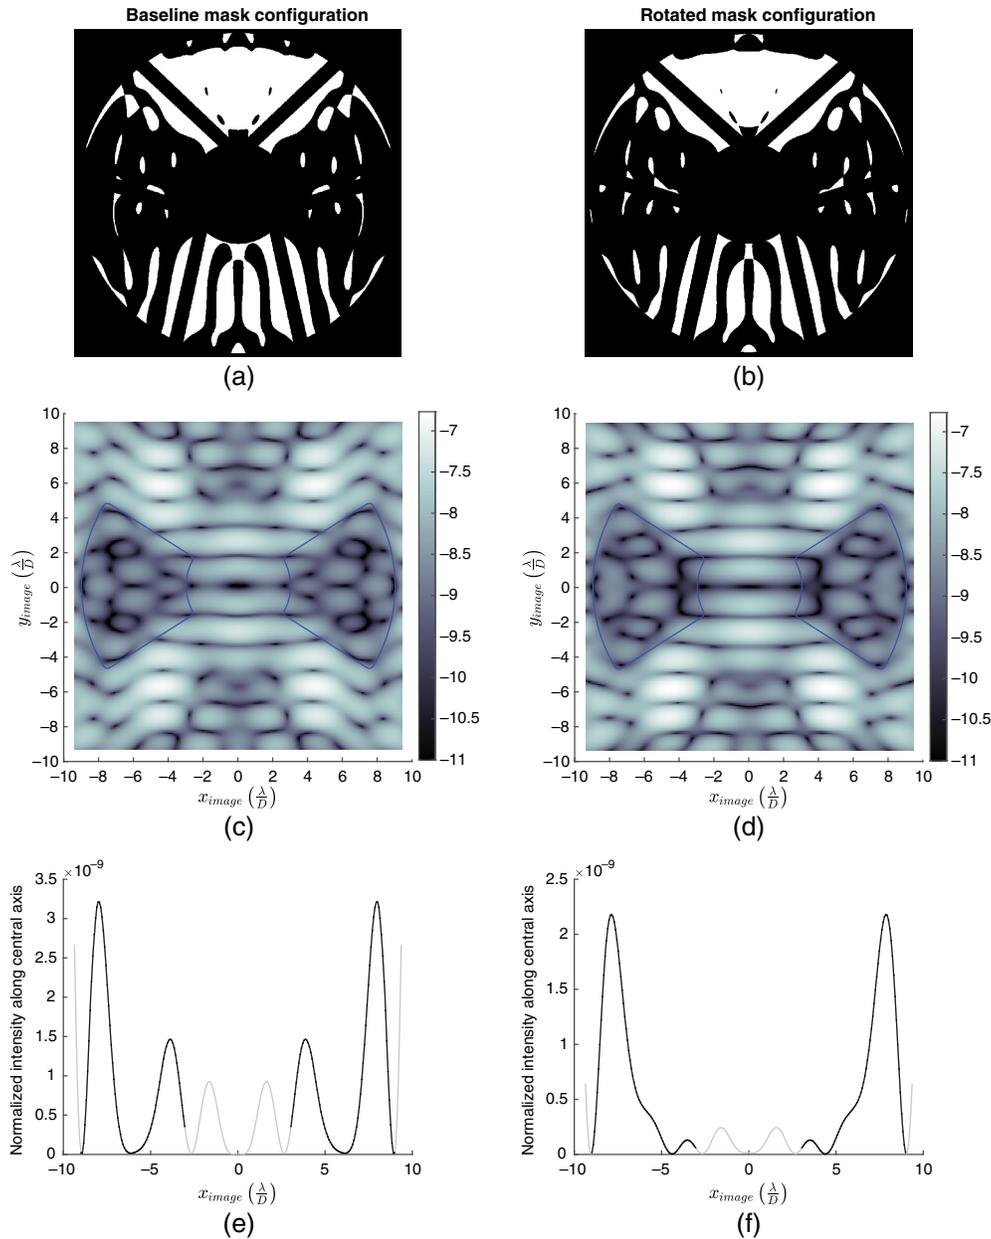


Fig. 20 Final shaped pupil masks for the baseline and rotated mask configurations, shown without the relative rotation. The two shaped pupil masks (a and b) are very similar, producing comparable dark holes (c and d) and contrasts (e and f). We would expect the designs to be identical for a perfectly rotationally symmetric telescope pupil; due to asymmetries in the true pupil, the rotated mask configuration requires more symmetry padding, leading to a lower throughput.

4.1 Costs of Additional Pupil Padding

Because the pupil padding is a permanent feature of the shaped pupil mask, the throughput losses due to the padding are present even if the pupil errors never approach the tolerance values. These losses can be significant, but they are not evenly distributed among the various types of pupil padding. To illustrate the general trends that describe how the different types of padding contribute to the total throughput loss, we generate two sets of telescope pupils for the rotated mask configuration with varying types and amounts of padding and calculate the amount of clear aperture for each pupil relative to that of the unpadded pupil (a quickly calculated proxy for the throughput loss). First, using the tolerance values, we pad for only one error type at a time to determine which types of padding most significantly affect the amount of clear aperture for the

rotated mask configuration's flight design. Then, we start with the fully padded telescope pupil that we used to optimize the flight design and increase one type of padding beyond its tolerance value to illustrate the cost associated with the additional padding.

As shown in Table 2, the rolloff padding is by far the dominant contributor to the aperture loss, followed by the padding for clocking, the struts, magnification, and translation. The padding for the secondary and the tabs, by comparison, minimally affects the aperture. If we increase the padding for any error beyond its tolerance, the amount of clear aperture decreases. The loss is minimal if we increase the padding for the tabs or the secondary, but it is noticeable for the other cases (Fig. 21). The loss is particularly dramatic if we increase the padding for magnification, and this is due to the corresponding change in the rolloff padding. [To prevent the unusable portion of the primary from coming into view, the rolloff padding is increased based on the padding for magnification, translation, and clocking. The clocking tolerance affects the rolloff padding because the modeled pupil is elliptical instead of circular, as explained in Eqs. (6)–(8)

Table 2 Impact of pupil padding on the clear aperture.

Padding type	% clear aperture (relative to the original telescope pupil)
Rolloff	94.6
Clocking	98.3
Struts	98.7
Magnification	98.9
Translation	98.9
Central obscuration	99.5
Tabs	99.6
Symmetry	99.6

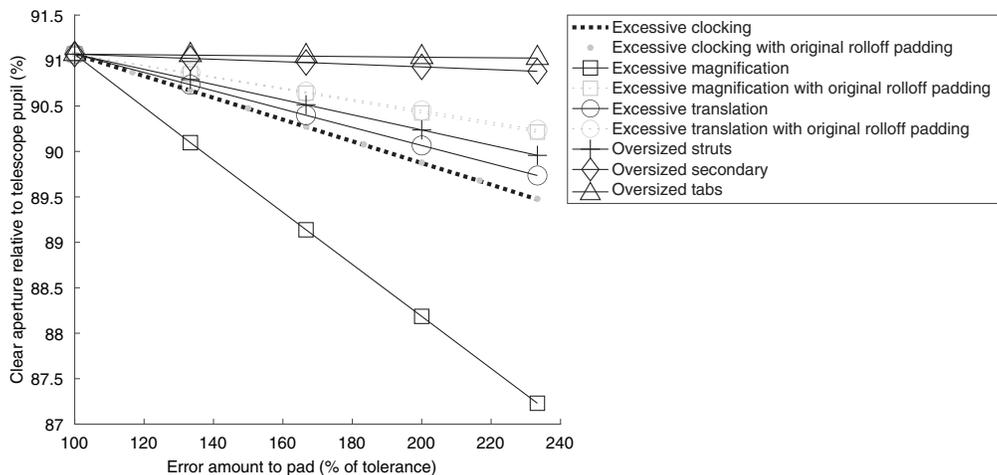


Fig. 21 Aperture costs of additional error padding. Although additional padding can improve the robustness to pupil errors, it costs some of the clear aperture. To investigate the costs associated with padding each type of error beyond its tolerance value, we start with a telescope pupil that has been padded for all errors at their tolerance levels and then continue to pad for increasing amounts of a given error. The most dramatic aperture loss occurs when padding is added for magnification errors that exceed the tolerance, but this is due to the corresponding increase in the rolloff padding, which depends on the limits used to pad for clocking, translation, and magnification errors. When the rolloff padding is held constant, clocking is the most costly error to pad beyond its tolerance.

and Fig. 7(a)]. With the rolloff padding held constant, additional padding for clocking errors becomes the most costly.

4.2 Using the DMs to Compensate for Excessive Errors

In addition to being potentially costly, padding beyond the three-sigma tolerances is unnecessary. To demonstrate this point, we evaluate how excessive errors affect the initial contrast (the contrast calculated without the DMs present and with no wavefront errors) and examine how well the DMs are able to compensate for the excessive error in some of the worst cases. For these tests, we start with a worst-case representation of the telescope pupil that has every error present at its tolerance value. In the case of magnification, which has an asymmetric tolerance, we choose the positive value, and it is also worth noting that we treat magnification as a resizing of the binary telescope pupil; we do not incorporate any edge blurring. Although this representation of the telescope pupil is still completely covered by the shaped pupil mask, the simultaneous presence of all error types at the levels used for padding reduces the extent to which the padding for one type of error can mask a different error. From this starting point, we increase a single error beyond its tolerance and calculate the initial contrast for the rotated mask configuration. Finally, we identify the worst of these error cases and evaluate how well the DMs are able to achieve a target contrast of 3.00×10^{-9} across a 15% band using an electric field conjugation (EFC) algorithm.⁴¹ This performance goal is related to the CGI requirements, and EFC is the baseline wavefront sensing and control algorithm for the CGI.

The most dramatic degradation in the initial contrast occurs as the clocking is increased beyond its tolerance. As the clocking increases 3 mrad past its tolerance (to 150% of its tolerance value), the initial contrast degrades by an order of magnitude (Fig. 22). Achieving the same amount of degradation for any of the other error cases that we consider, by comparison, requires

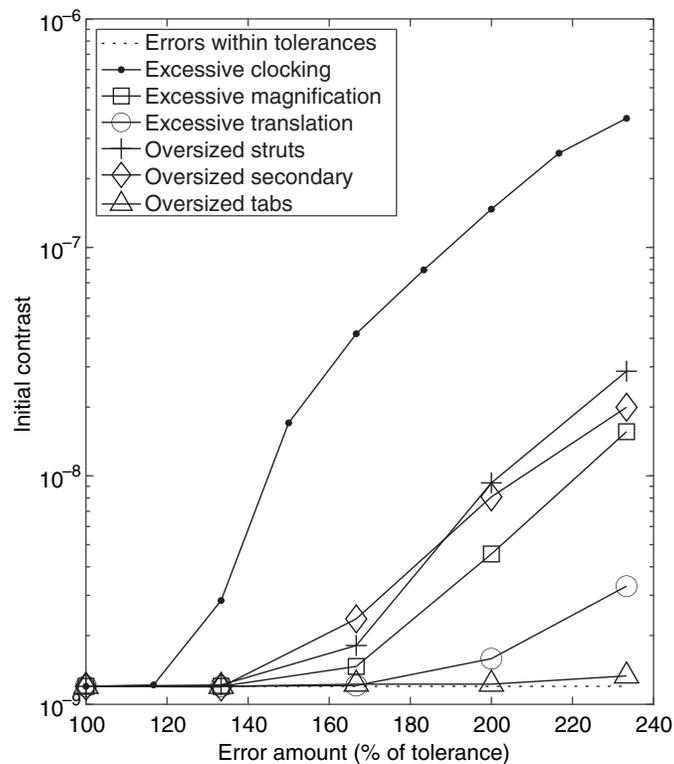


Fig. 22 Initial contrast degradation due to errors exceeding their tolerances. As an error exceeds the limits used for padding, the initial contrast degrades. The most rapid degradation is caused by excessive clocking errors, which are also among the costliest to pad. However, it is important to note that these results do not consider the extent to which the DMs can recover the contrast; they show the degradation in the starting point before any control is applied.

increasing the error to $>200\%$ of its tolerance. However, this contrast degradation can be mitigated at least partially by the DMs. In our example, even as much as 5 mrad of excessive clocking can be corrected fully by the DMs (Table 4 and Fig. 23), which is fortuitous because clocking is one of the costliest errors to pad.

To simulate the DM corrections, we used a suite of modular optical propagation and coronagraph modeling algorithms developed at Princeton. These algorithms implement Fresnel and Fraunhofer propagation using a matrix Fourier transform approach based on that of Soummer et al.,⁴² and the DM surfaces are modeled by superposing influence functions, which describe how a single actuator affects the DM surface when pistoned. For each actuator, we scale the influence function provided with PROPER⁴³ by the product of the gain (which we assume is 1 nm/V for all actuators) and the commanded voltage, and then we sum all of the individual influence functions (each centered over their respective actuators) to produce the modeled surface. To match the DMs used in the CGI, we represent the two DMs as 48×48 actuator arrays

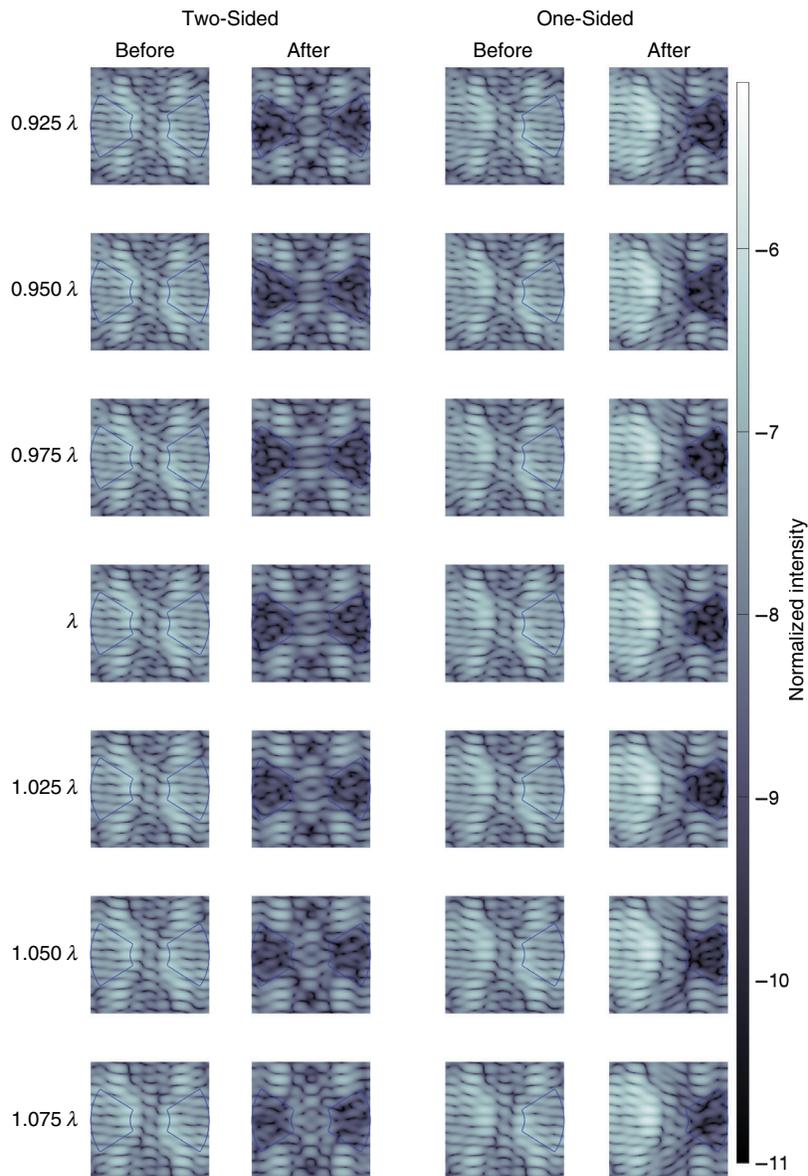


Fig. 23 Normalized intensity in the image plane before and after correction for 11 mrad of total clocking, shown for both two- and one-sided dark holes (outlined) at each of the seven controlled wavelengths. The DMs are able to completely correct the clocking error across the 15% band for either type of dark hole.

Table 3 Selected technical details for the DMs. This information (and more) can be found in a paper by Poberezhskiy et al.²⁰

Actuator array dimensions	48 × 48
Technology	Lead-magnesium-niobate (PMN) electrostrictive
Manufacturer	AOA Xinetics
Actuator spacing	~1 mm
Actuator stroke limit	~500 nm
Voltage range	0 to 100 V
Electronics resolution	16-bit

separated by 1 m;²⁰ selected technical information about the CGI DMs is provided for reference in Table 3.

The extent to which excessive errors can be corrected depends in part on the magnitude of the error and the choice of the dark hole. It is also worth noting that the ability to compensate for pupil errors is also dependent on the ratio of the square of the DM width to the DM separation, with smaller ratios preferred, but this ratio is fixed for the Roman Space Telescope. To illustrate the general patterns that describe how well the errors can be corrected, we consider two of our worst excessive error cases (11 and 14 mrad of total clocking) and use an EFC algorithm to improve the contrast over either a two-sided dark hole or a single-sided dark hole (Tables 4 and 5 and Figs. 23 and 24). Because the specific number of iterations required depends upon the size of the dark hole and the magnitude of the pupil error, the algorithm is set to run for 30 iterations, the maximum number of allowed wavefront sensing and control iterations planned for digging the initial dark holes on orbit for the baseline configurations of the hybrid Lyot coronagraph (HLC), the bowtie shaped pupil coronagraph, and the wide-field-of-view shaped pupil coronagraph. The postcorrection values shown in Tables 4 and 5 correspond to either the first iteration for which the contrast is beneath the target across the entire band or to the iteration for which the contrast has converged as close to the target as possible.

As we would expect, smaller errors require fewer control iterations and are more likely to be successfully corrected over the full bandwidth for a two-sided dark hole (Table 4 and Fig. 23). As the magnitude of the error increases, there is eventually a point at which there is simply too much error for the DMs to correct completely across the full band for a two-sided dark hole. At this point, the edges of the band fail to reach the target contrast. As the error continues to increase, the contrast at the outer edges of the band continues to deteriorate, and the fully corrected portion

Table 4 Control results for 11 mrad of total clocking (5 mrad beyond the tolerance).

Dark hole type	Iterations	Contrast achieved at selected wavelengths over 15% band						
		0.925 λ	0.950 λ	0.975 λ	λ	1.025 λ	1.050 λ	1.075 λ
Two-sided	0	6.15 $\times 10^{-8}$	6.77 $\times 10^{-8}$	7.18 $\times 10^{-8}$	7.44 $\times 10^{-8}$	7.57 $\times 10^{-8}$	7.98 $\times 10^{-8}$	8.74 $\times 10^{-8}$
	2	2.12 $\times 10^{-9}$	1.75 $\times 10^{-9}$	1.42 $\times 10^{-9}$	1.37 $\times 10^{-9}$	1.53 $\times 10^{-9}$	1.67 $\times 10^{-9}$	2.65 $\times 10^{-9}$
One-sided	0	6.11 $\times 10^{-8}$	6.69 $\times 10^{-8}$	7.09 $\times 10^{-8}$	7.31 $\times 10^{-8}$	7.39 $\times 10^{-8}$	7.78 $\times 10^{-8}$	8.53 $\times 10^{-8}$
	1	1.00 $\times 10^{-9}$	8.63 $\times 10^{-10}$	6.45 $\times 10^{-10}$	6.91 $\times 10^{-10}$	7.91 $\times 10^{-10}$	7.89 $\times 10^{-10}$	1.40 $\times 10^{-9}$

Table 5 Control results for 14 mrad of total clocking (8 mrad beyond the tolerance).

Dark hole type	Iterations	Contrast achieved at selected wavelengths over 15% band						
		0.925 λ	0.950 λ	0.975 λ	λ	1.025 λ	1.050 λ	1.075 λ
Two-sided	0	3.18 $\times 10^{-7}$	3.39 $\times 10^{-7}$	3.54 $\times 10^{-7}$	3.61 $\times 10^{-7}$	3.62 $\times 10^{-7}$	3.68 $\times 10^{-7}$	3.77 $\times 10^{-7}$
	4	8.25 $\times 10^{-9}$	5.95 $\times 10^{-9}$	4.39 $\times 10^{-9}$	3.50 $\times 10^{-9}$	3.26 $\times 10^{-9}$	3.54 $\times 10^{-9}$	5.23 $\times 10^{-9}$
One-sided	0	3.17 $\times 10^{-7}$	3.37 $\times 10^{-7}$	3.52 $\times 10^{-7}$	3.58 $\times 10^{-7}$	3.59 $\times 10^{-7}$	3.64 $\times 10^{-7}$	3.74 $\times 10^{-7}$
	1	2.59 $\times 10^{-9}$	1.82 $\times 10^{-9}$	1.32 $\times 10^{-9}$	1.22 $\times 10^{-9}$	1.36 $\times 10^{-9}$	1.48 $\times 10^{-9}$	2.20 $\times 10^{-9}$

of the band shrinks in toward the central wavelength. However, reducing the size of the dark hole by switching to a single-sided dark hole can enable the DMs to fully correct the error across the band once again (Table 5 and Fig. 24), and there are contributed field stops included in the CGI that allow for single-sided bowtie dark holes if that becomes necessary.

It is also worth noting that, for each of the example cases, the actuator heights and commanded voltages (shown in Table 6) are well within their limits (Table 3). For the cases in which the DMs successfully corrected the error, the largest absolute change in actuator height was 9.4 nm, which is about 2% of the stroke limit. For the example in which the DMs were unable to completely correct the error, the largest absolute change in actuator height was 22.4 nm. The root-mean-square (RMS) stroke for the nominal HLC design is about 30 nm RMS before accounting for any optical aberrations, whereas for the shaped pupil coronagraphs it is nominally zero because the apodizer does most of the contrast generation. Therefore, a ± 22.4 nm change in actuator height for the bowtie shaped pupil coronagraph in these cases with misalignments is well within the DM actuation capabilities. Overall, there is a range of excessive errors that can be completely corrected by the DMs across the full band, and beyond this range, the contrast degrades gracefully.

5 Summary

Currently, a space-based direct imaging mission offers one of the best paths to searching for life on an Earth-like exoplanet, and the level of interest in such a mission is underscored by the most recent decadal survey's recommendation of a 6-m class infrared/optical/ultraviolet telescope for high-contrast imaging and spectroscopy as the highest priority for space frontier missions. To help close some of the numerous technical gaps that need to be addressed, the CGI on the Nancy Grace Roman Space Telescope will serve as a demonstrator for high-contrast imaging technologies including DMs to correct high-order wavefront errors and precision-fabricated masks specifically designed to attenuate the relative intensity of the star off-axis while maximizing the core throughput. The CGI has three interchangeable observing configurations, one of which is a bowtie shaped pupil coronagraph for high-contrast spectroscopy.

The bowtie shaped pupil coronagraph has a baseline 0-deg mask configuration for the technology demonstration and a 60-deg mask configuration contributed by ExEP. Each mask configuration consists of a shaped pupil mask, a bowtie-shaped focal plane mask, and a bowtie-shaped Lyot stop. The designs of the shaped pupil mask and the Lyot stop are determined through an iterative process: a linear program optimizes the shaped pupil mask for a given Lyot stop, and the optimization is repeated for a range of potential Lyot stops until the highest-throughput combination is identifiable in a parameter space cube of resolution $1\%D$ in the inner and outer diameters of the Lyot stop and 1 deg in the opening angle. The resulting flight designs for the baseline and rotated mask configurations have core throughputs of 4.50% and 3.89%, respectively.

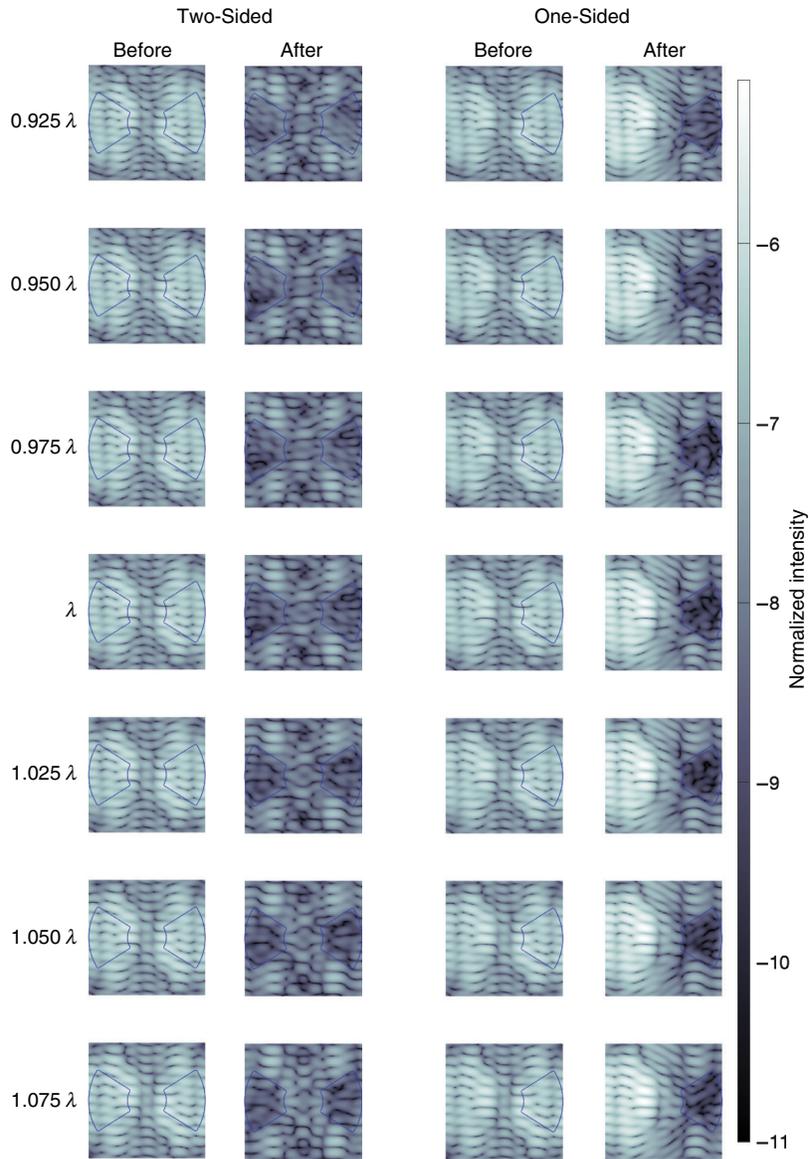


Fig. 24 Normalized intensity in the image plane before and after correction for 14 mrad of total clocking. Although the DMs cannot completely correct this much clocking error across the full 15% band for a two-sided dark hole, they can for a one-sided dark hole. Another alternative would be to reduce the controlled bandwidth because the contrast degrades at the edges of the band first.

Table 6 Maximum and minimum voltage commands and actuator heights for the clocking error cases.

Error case	Max. voltage (V)	Min. voltage (V)	Max. height (nm)	Min. height (nm)
11 mrad, two-sided	7.9	-10.0	8.3	-9.4
11 mrad, one-sided	1.3	-1.0	1.5	-1.1
14 mrad, two-sided	18.8	-23.9	19.8	-22.4
14 mrad, one-sided	2.6	-1.9	2.8	-2.2

To provide some robustness to errors, the shaped pupil mask optimization directly incorporates tip and tilt, and the optimization starting point is padded so that the telescope pupil is fully masked for conservative estimates of misalignments and manufacturing errors. This padding is specifically tailored to a list of error tolerances as opposed to being generated by rotating, translating, and scaling a provided high-resolution pupil array, and this new approach for generating robust shaped pupil designs is made tractable by a completely parametric model of the telescope pupil (which we also use to study the costs of adding various types of robustness and the impacts of pupil errors not masked by the padding). We have developed two versions of this model: one that generates a pupil by superposing elliptical and rectangular obscurations in PROPER, and one that consists of a set of algebraic equations that allow the individual edges of the obscurations to be manipulated directly. These equations are written in a way that allows the pupil to be rotated, facilitating designs for rotated mask configurations. There is also some padding to enforce y -axis symmetry, which is required to keep the optimization problem tractable, and the amount of symmetry padding depends on the mask configuration orientation.

However, the pupil padding costs throughput, and the throughput difference between the baseline and rotated mask configurations is largely due to the higher amount of symmetry padding required by the rotated mask configuration. The throughput loss due to pupil padding is also unevenly distributed among the various types of error padding, with the rolloff padding costing the most throughput and the padding for symmetry, the tabs, and the central obscuration costing the least.

Because the throughput losses due to pupil padding are fixed, it is important to balance the need for pupil padding against the throughput cost. The baseline and rotated mask configurations are padded to conservative three-sigma estimates of various pupil errors, which is expected to provide sufficient robustness. However, if the pupil errors exceed the estimates, the DMs can be used to compensate at least partially. In our simulations, excessive clocking errors lead to the most dramatic contrast degradation by far, but the DMs are able to completely compensate for the error across a 15% band in even the worst cases (11 mrad and 14 mrad of total clocking, with padding for only 6 mrad). The one caveat is that a single-sided dark hole is required for the 14 mrad case, while a two-sided dark hole can be used for the 11 mrad case. In general, the DMs are able to compensate for excessive pupil errors up to a point, after which the contrast degrades gracefully; the fully corrected portion of the band shrinks inward toward the central wavelength, and the corrected contrast at the ends of the band degrades as the error increases.

Future missions with obscured apertures may benefit from extending the use of the DMs a step further: selectively reducing the pupil padding to improve the throughput and relying on the DMs to compensate for the unpadded error as needed. The successful implementation of this approach and the extent to which the padding can be reduced depend strongly on the ratio of the square of the DM width to the DM separation, in addition to depending on details including the specific telescope pupil obscuration and the magnitudes of the expected errors. The throughput gains depend on the specific situation as well, but preliminary results suggest the gains can be as high as a percentage point or two in some cases. Continued investigations of the throughput gains achievable by reducing the pupil padding and relying on the DMs to compensate for pupil errors as needed, and the factors that affect the success and usefulness of this approach, are the subjects of future work.

6 Appendix: Pupil Model Parameters

The parameters for the two telescope pupil models were determined independently through separate trial-and-error approaches with different goals, and due to tight time constraints, much of the work was done by hand. For each model, we inferred a set of parameter values that reproduced the provided high-resolution array. Then, as a final step, we shifted the coordinates to center the modeled primary. (The pupil position in the provided array is arbitrary.)

For the PROPER-based model (used to develop the baseline mask configuration), the overall goal was to find a set of parameters that led to the best fit, which minimized the summed squared difference between the provided high-resolution pupil array and the PROPER-generated model. The first step was to determine the approximate lateral position and size (width or radius) of each

pupil component by eye. Then, this initial pupil model was refined through an iterative grid search in which one to three parameters were varied at a time for each individual obscuration.

For the algebraic model (used to develop the rotated mask configuration), the overall goal was to find a conservative set of parameters that led to a pupil with the smallest obscuration that still completely covered the obscuration in the provided pupil. The first step was to determine a set of parameters that produced a pupil with an obscuration that covered every pixel of the obscuration in the provided pupil; this was easily checked by subtracting the provided pupil from the modeled pupil. Then, using a plot of the difference of the modeled and provided pupils as a guide, the parameters were refined by hand for one pupil component at a time. Changes that reduced the number of pixels inside the modeled obscuration were kept so long as the provided obscuration was completely covered, and changes that increased the number of obscured pixels were discarded. The process ended when any further adjustments either increased the number of obscured pixels or uncovered any part of the obscuration in the provided pupil.

Although the parameters for the two pupil models were inferred through different methods with different goals, the final numerical values are very similar, as shown in Table 7. The differences in the angles of the last three struts reflect differences in the angle ranges used to define the nominal strut orientations. It is also worth noting that while the algebraic model briefly uses the strut length, the exact value is unimportant as long as it is positive.

Table 7 Pupil model parameters for the baseline and rotated mask configurations.

Pupil component	Parameter name	Model value	
		PROPER-based model (baseline mask config.)	Algebraic model (rotated mask config.)
Primary (ellipse)	Center loc. in x , x_{c_p}	0	0
	Center loc. in y , y_{c_p}	0	0
	Semiminor axis length (along x), a_p	0.4954	0.4952
	Semimajor axis length (along y), b_p	0.5000	0.4998
Secondary (ellipse)	Center loc. in x , $x_{c_{sec}}$	0	0
	Center loc. in y , $y_{c_{sec}}$	-3.6625×10^{-4}	-3.7249×10^{-4}
	Semiminor axis length (along x), a_{sec}	0.1502	0.1503
	Semimajor axis length (along y), b_{sec}	0.1515	0.1516
Strut 1 (rectangle)	Center loc. in x , $x_{c_{strut1}}$	0.1505	0.1520
	Center loc. in y , $y_{c_{strut1}}$	0.2742	0.2812
	Nominal orientation, θ_{strut1}	77.4858 deg	77.5060 deg
	Half width, w_{b_1}	0.0160	0.0161
	Length, l_1	0.55	0.55 (unimportant)
Strut 2 (rectangle)	Center loc. in x , $x_{c_{strut2}}$	0.2928	0.2931
	Center loc. in y , $y_{c_{strut2}}$	0.0043	0.0041
	Nominal orientation, θ_{strut2}	-18.5193 deg	-18.5199 deg
	Half width, w_{b_2}	0.0161	0.0162
	Length, l_2	0.55	0.55 (unimportant)

Table 7 (Continued).

Pupil component	Parameter name	Model value	
		PROPER-based model (baseline mask config.)	Algebraic model (rotated mask config.)
Strut 3 (rectangle)	Center loc. in x , $X_{c_{strut3}}$	0.1538	0.1715
	Center loc. in y , $Y_{c_{strut3}}$	-0.2598	-0.2765
	Nominal orientation, θ_{strut3}	-43.2889 deg	-43.3000 deg
	Half width, w_{b3}	0.0160	0.0161
	Length, l_3	0.55	0.55 (unimportant)
Strut 4 (rectangle)	Center loc. in x , $X_{c_{strut4}}$	-0.1661	-0.1675
	Center loc. in y , $Y_{c_{strut4}}$	-0.2702	-0.2716
	Nominal orientation, θ_{strut4}	-136.7115 deg	43.2900 deg
	Half width, w_{b4}	0.0160	0.0161
	Length, l_4	0.55	0.55 (unimportant)
Strut 5 (rectangle)	Center loc. in x , $X_{c_{strut5}}$	-0.2895	-0.2931
	Center loc. in y , $Y_{c_{strut5}}$	0.0063	0.0050
	Nominal orientation, θ_{strut5}	-161.4717 deg	18.5190 deg
	Half width, w_{b5}	0.0161	0.0162
	Length, l_5	0.55	0.55 (unimportant)
Strut 6 (rectangle)	Center loc. in x , $X_{c_{strut6}}$	-0.1483	-0.1511
	Center loc. in y , $Y_{c_{strut6}}$	0.2685	0.2812
	Nominal orientation, θ_{strut6}	102.5074 deg	-77.5000 deg
	Half width, w_{b6}	0.0160	0.0160
	Length, l_6	0.55	0.55 (unimportant)
Tab 1 (part of ellipse between struts 1 and 2)	Center loc. in x , $X_{c_{tab1}}$	0	0
	Center loc. in y , $Y_{c_{tab1}}$	0.0067	0.0102
	Semiminor or semimajor axis length (along x), a_{tab1}	0.1667	0.1660
	Semimajor or semiminor axis length (along y), b_{tab1}	0.1680	0.1635
Tab 2 (part of ellipse between struts 3 and 4)	Center loc. in x , $X_{c_{tab2}}$	0	0
	Center loc. in y , $Y_{c_{tab2}}$	0.0084	-0.0003
	Semiminor axis length (along x), a_{tab2}	0.1693	0.1612
	Semimajor axis length (along y), b_{tab2}	0.1706	0.1620

Table 7 (Continued).

Pupil component	Parameter name	Model value	
		PROPER-based model (baseline mask config.)	Algebraic model (rotated mask config.)
Tab 3 (part of ellipse between struts 5 and 6)	Center loc. in x , $x_{c_{tab_3}}$	0	0
	Center loc. in y , $y_{c_{tab_3}}$	0.0067	0.0102
	Semiminor or semimajor axis length (along x), a_{tab_3}	0.1667	0.1660
	Semimajor or semiminor axis length (along y), b_{tab_3}	0.1680	0.1635

Acknowledgments

Part of the work reported here was performed at Princeton University under NASA Contract NNG16PJ29C, and part was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

References

1. National Research Council, *Astronomy and Astrophysics for the 1970's, Volume 1: Report of the Astronomy Survey Committee*, National Academy Press, Washington, D.C. (1972).
2. National Research Council, *Astronomy and Astrophysics for the 1980's, Volume 1: Report of the Astronomy Survey Committee*, National Academy Press, Washington, D.C. (1982).
3. National Research Council, *The Decade of Discovery in Astronomy and Astrophysics*, National Academy Press, Washington, D.C. (1991).
4. National Research Council, *Astronomy and Astrophysics in the New Millennium*, National Academy Press, Washington, D.C. (2001).
5. National Research Council, *New Worlds, New Horizons in Astronomy and Astrophysics*, The National Academies Press, Washington, D.C. (2010).
6. National Research Council, *Pathways to Discovery in Astronomy and Astrophysics for the 2020s*, National Academies Press, Washington, D.C. (2021).
7. A. Wolszczan and D. A. Frail, "A planetary system around the millisecond pulsar PSR1257 + 12," *Nature* **355**, 145–147 (1992).
8. M. Mayor and D. Queloz, "A Jupiter-mass companion to a solar-type star," *Nature* **378**, 355–359 (1995).
9. I. A. Bond et al., "OGLE 2003-BLG-235/MOA 2003-BLG-53: a planetary microlensing event," *Astrophys. J.* **606**, L155–L158 (2004).
10. G. Chauvin et al., "A giant planet candidate near a young brown dwarf," *Astron. Astrophys.* **425**, L29–L32 (2004).
11. M. W. Muterspaugh et al., "The phases differential astrometry data archive. V. Candidate substellar companions to binary systems," *Astron. J.* **140**, 1657–1671 (2010).
12. W. J. Borucki, "KEPLER Mission: development and overview," *Rep. Prog. Phys.* **79**, 036901–50 (2016).
13. National Academies of Sciences, Engineering, and Medicine, *Exoplanet Science Strategy*, National Academies Press, Washington, D.C. (2018).
14. HabEx Study Team, "Exploring new worlds – understanding our universe," 2019, <https://www.jpl.nasa.gov/habex/pdf/HabEx-Final-Report-Public-Release-LINKED-0924.pdf>.
15. The LUVOIR Team, "The LUVOIR mission concept study final report," 2019, <https://asd.gsfc.nasa.gov/luvoir/reports/>.
16. B. Crill and N. Siegler, "Space technology for directly imaging and characterizing exo-Earths," *Proc. SPIE* **10398**, 103980H (2017).

17. B. Crill and N. Siegler, “Exoplanet Exploration Program 2019 technology plan appendix,” 2019, https://exoplanets.nasa.gov/internal_resources/1123/.
18. H. Tang et al., “The WFIRST coronagraph instrument phase B optical design,” *Proc. SPIE* **11117**, 111170C (2019).
19. N. J. Kasdin et al., “The Nancy Grace Roman Space Telescope Coronagraph Instrument (CGI) technology demonstration,” *Proc. SPIE* **11443**, 114431U (2020).
20. I. Y. Poberezhskiy et al., “Roman Space Telescope Coronagraph: engineering design and operating concept,” *Proc. SPIE* **11443**, 114431V (2021).
21. A. J. E. Riggs et al., “Flight mask designs of the Roman Space Telescope Coronagraph Instrument,” *Proc. SPIE* **11823**, 118231Y (2021).
22. J. Trauger et al., “Hybrid Lyot coronagraph for WFIRST-AFTA: coronagraph design and performance metrics,” *J. Astron. Telesc. Instrum. Syst.* **2**(1), 011013 (2016).
23. T. Groff et al., “Roman Space Telescope CGI: prism and polarizer characterization modes,” *Proc. SPIE* **11443**, 114433D (2021).
24. V. Bailey, “Contributed mask summary,” Presented at the Virtual “Stanford” meeting on August 5 (2020).
25. V. Bailey, “Possible mask options,” In the notes of the CGI coronagraph mask working group (2020).
26. E. Nielsen, “WFIRST: 47 UMa c and the Bowtie,” Presented to the CGI coronagraph mask working group (2020).
27. E. Nielsen, “Update of additional bowtie orientations 47 uma detectability,” Presented to the CGI coronagraph mask working group (2020).
28. R. Boucarut, RST Coronagraph Instrument Requirements Document (CGIRD), Goddard Space Flight Center, Greenbelt, Maryland (2021).
29. V. Bailey, Private communication to author, (2020).
30. R. Fourer, D. M. Gay, and B. W. Kernighan, *AMPL: A Modeling Language for Mathematical Programming*, 2nd ed., Duxbury Thomson (2003).
31. N. T. Zimmerman et al., “Shaped pupil Lyot coronagraphs: high-contrast solutions for restricted focal planes,” *J. Astron. Telesc. Instrum. Syst.* **2**, 011012 (2016).
32. A. Riggs et al., “Shaped pupil coronagraph design improvements for the WFIRST coronagraph instrument,” *Proc. SPIE* **10400**, 104000O (2017).
33. A. J. E. Riggs, G. Ruane, and B. D. Kern, “Directly constraining low-order aberration sensitivities in the WFIRST coronagraph design,” *Proc. SPIE* **11117**, 111170F (2019).
34. B. Nemati, Private communication to author, (2020).
35. R. J. Vanderbei, “Fast Fourier optimization,” *Math. Prog. Comput.* **4**, 53–69 (2012).
36. A. Riggs, “Summary of decisions on CGI pupil usage,” unpublished (2020).
37. “2020-05-13 CGI entrance pupil 8k as-built uncoated optics.fits,” 2020, This is the pupil file, http://github.com/ajeldorado/falco-matlab/blob/master/lib/masks/WFIRST/PhaseC/pupil_CGI-20200513_8k_binary_noF.png.
38. Photon Engineering, “FRED overview,” 2020, <https://photonengr.com/wp-content/uploads/2020/04/FREDOverview19.40.2.pdf>.
39. J. E. Krist, “PROPER: an optical propagation library for IDL,” *Proc. SPIE* **6675**, 66750P (2007).
40. S. R. Hildebrandt, “Roman coronagraph exposure time calculator,” <https://roman.ipac.caltech.edu/sims/ETC.html>.
41. T. D. Groff et al., “Methods and limitations of focal plane sensing, estimation, and control in high-contrast imaging,” *J. Astron. Telesc. Instrum. Syst.* **2**(1), 011009 (2016).
42. R. Soummer et al., “Fast computation of Lyot-style coronagraph propagation,” *Opt. Express* **15**(24), 15935–15951 (2007).
43. J. Krist et al., “PROPER: an optical propagation library for IDL, Python, and Matlab,” NASA Jet Propulsion Laboratory, California Institute of Technology, version 3.2.5 ed., 2021, <https://proper-library.sourceforge.net>.

Jessica Gersh-Range is an associate research scholar at Princeton’s High Contrast Imaging Laboratory, where she optimized shaped pupil Lyot coronagraph masks to support technology development efforts for the Nancy Grace Roman Space Telescope. She was the recipient of a

NASA Graduate Student Researchers Program fellowship at Marshall Space Flight Center, and she also worked at the Space Telescope Science Institute as a graduate student. She received her BA degree in physics from Swarthmore College and her PhD in mechanical engineering from Cornell University.

A J Eldorado Riggs is an optical engineer at the Jet Propulsion Laboratory. His research focuses on the high-contrast imaging of exoplanets, in particular, mask optimization and wavefront sensing and control for the Coronagraph Instrument on the Nancy Grace Roman Space Telescope. He received his BS degree in physics and mechanical engineering from Yale University and his PhD in mechanical and aerospace engineering from Princeton University in 2011 and 2016, respectively.

N. Jeremy Kasdin is the Eugene Higgins professor of mechanical and aerospace engineering, emeritus, at Princeton University and the former assistant dean for engineering at the University of San Francisco. He received his PhD from Stanford University in 1991. After being the chief systems engineer for NASA's Gravity Probe B spacecraft, he joined the Princeton faculty in 1999, where he researched high-contrast imaging technology for exoplanet imaging. From 2014 to 2016, he was the vice dean of the School of Engineering and Applied Science at Princeton. From 2014 to 2021, he was the adjutant scientist for the coronagraph instrument on NASA's Nancy Grace Roman Space Telescope.