Research Article Vol. 37, No. 4 / April 2020 / Journal of the Optical Society of America A 629

Journal of the Optical Society of America A
OPTICS, IMAGE SCIENCE, AND VISION

Field stop diffraction and sampling effects in apodized pupil Lyot coronagraphs

Scott D. Will* and James R. Fienup

The Institute of Optics, University of Rochester, Rochester, New York 14627, USA
**Corresponding author: scott.will@rochester.edu*

Received 21 October 2019; revised 14 February 2020; accepted 14 February 2020; posted 14 February 2020 (Doc. ID 379609); published 20 March 2020

We study the predicted performance of two apodized pupil Lyot coronagraph designs in the presence of an occulter-plane field stop. We discuss techniques for capturing diffraction effects when the radius of the stop is larger than the field of view of an ordinary numerical diffraction model, including mask upsampling and analytical focal-plane envelope functions. We simulate a closed-loop coronagraphic wavefront control to assess the extent to which such diffraction effects can be compensated using deformable mirrors. We show that for the designs considered, field stop diffraction effects are significant at diameters considerably larger than the instrument field of view, suggesting the need to explicitly include a focal-plane stop in the design process.

https://doi.org/10.1364/JOSAA.379609

1. INTRODUCTION

Over the past three decades, direct imaging of extrasolar planets and circumstellar dust clouds has become an area of great interest in observational astronomy, and is a primary science goal of the WFIRST and proposed LUVOIR space observatories [1,2]. A number of techniques for this purpose have been developed, broadly categorized as Lyot coronagraphy, external occulting with star-shade masks, phase-induced amplitude apodization, and nulling interferometry. Within the family of Lyot-type coronagraphs, there exist numerous varieties including phase-masking coronagraphs [3,4], vortex coronagraphs [5–7], hybrid Lyot coronagraphs [8,9], apodized pupil Lyot coronagraphs (APLCs) [10–12], and shaped pupil Lyot coronagraphs [13–17].

Lyot-type coronagraphs employ a series of diffractive masks designed to engineer the on-axis stellar image formed by the telescope to minimize diffracted starlight in an off-axis region of interest known as the dark zone, while transmitting the faint signal from an off-axis orbiting exoplanet so that it can be detected and characterized. The simplest form of a Lyot-type coronagraph is a 4-F afocal relay followed by an additional Fourier transforming stage, shown in Fig. 1. In this configuration, a reimaged telescope pupil at the entrance pupil of the coronagraph (plane A), possibly coinciding with an apodizing or beam-shaping mask, is focused onto an occulting mask (plane B), which blocks the core of an on-axis stellar image while allowing light from the planet to pass. The beam is then Fourier transformed to a second pupil plane with a Lyot stop (plane C), which suppresses the vast majority of the remaining starlight for spatial frequencies inside the dark zone. A final Fourier transform forms a real image on an optical detector (plane D). For APLCs, the occulting mask is generally a small, opaque circular focal spot, while shaped-pupil Lyot coronagraphs tend to employ a focal-plane mask transmissive only over part of the plane, such as an annular or bowtie-shaped mask [18].

The intensity ratio of starlight to planetary light collected by the telescope is estimated to be at least 10^{10} for Earth-like exoplanets orbiting Sun-like stars at visible wavelengths, with extremely small angular separations on the order of 0.1 arcsec or less [19]. Due to the extremely tight tolerances required to realize instruments capable of such high levels of starlight suppression, accurate numerical diffraction modeling is a vital tool for analyzing the expected performance of proposed coronagraph designs under ideal and flight-like conditions, and has been the topic of considerable effort in recent years [20–28].

Early APLC designs, proposed by Aime *et al.* [10,11,29,30], were created to operate with simple unobscured, monolithic entrance pupils, and employed families of continuous-valued apodizing functions such as the prolate spheroidal wave functions. These functions have the property of maximizing the energy concentrated in the mainlobe of the far-field diffraction pattern, enabling high levels of starlight suppression with small occulting masks [11]. With the rise of complicated pupils with segments, support struts, and/or spiders, more recent APLC designs have borrowed pupil-shaping methods developed for shaped pupil Lyot coronagraphs and employ entrance pupil masks consisting of a grid of unit cells whose transmission pattern is numerically optimized, for a given Lyot stop and occulter, to maximize the open area while achieving a fixed desired dark

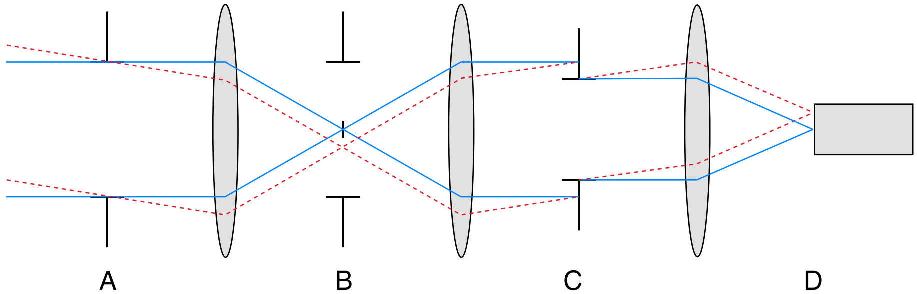

Fig. 1. Basic unfolded layout of a Lyot-type coronagraph. In an APLC, planes A, B, C, and D correspond to the apodizer, occulter, Lyot stop, and detector planes, respectively. Blue rays indicate on-axis starlight that is suppressed while dashed red rays indicate off-axis light from the orbiting exoplanet.

hole contrast ratio over a fixed region of the detector plane [12,13].

A number of techniques have been proposed and studied for fabricating binary apodizing masks. One recent procedure developed by Balasubramaniam *et al.* [31,32] consists of etching a set of highly absorbing regions onto a reflective aluminum-coated silicon substrate. In these fabrication modes, each sample of the numerical apodizing mask solution is realized as a lithographic resolution element with some finite characteristic width.

In coronagraph diffraction simulations, including the forward model utilized in mask design algorithms [33], the field immediately in front of the occulter (which is proportional to the Fourier transform of the apodizer in the paraxial regime) is computed by performing a discrete Fourier transform (DFT) of the field transmitted by the apodizer. In this case, the spatial extent of the computable occulter-plane field is limited by classical sampling theory to the Nyquist limit imposed by the sample spacing of the apodizing mask. This implies that it is not possible with this approach to characterize the effects of diffraction from focal-plane stops in Plane B of Fig. 1 much larger than the Nyquist limit. The focal-plane stop may be a field stop explicitly inserted into the coronagraph to limit the field of view of the instrument, or may represent the boundaries of the filter wheel or stage used to support the occulting mask itself. We demonstrate numerically that field stop diffraction is potentially significant for current coronagraph designs not explicitly optimized to include a focal-plane field stop.

Additionally, for masks with relatively coarse sample spacing, we have observed that after upsampling the pupil-plane masks, the predicted on-axis stellar image departs from the image computed using the original native-resolution arrays by a non-negligible amount. We analyze this effect in the context of numerical Fourier optics, and discuss implications for coronagraph design and numerical modeling.

This article is structured as follows. In Section 2, we review basic analytical and numerical forward modeling strategies for Lyot-type coronagraphs. In Section 3, we discuss the properties of coronagraph mask upsampling, derive an analytical envelope function that enables the paraxial occulter-plane field to be computed over an arbitrarily large field of view in the paraxial regime, and show that this field represents the limit as pupil-plane mask sampling becomes infinitely fine. In Section 4, we utilize this framework to present and analyze the results of numerical models of two APLC designs incorporating a focal-plane stop of varying radius and under multiple oversampling conditions, both with and without coronagraphic wavefront control for active diffraction suppression.

A. Notation

In the discussion that follows, continuous functions are represented using parenthetical arguments, such as $g(x, y)$. Discrete functions are denoted using integer index variables inside square brackets, for example, a function $h[m, n]$ indexed by the integer values m and n.

The continuous coordinate pairs (x, y) and (θ_x, θ_y) are used to represent transverse pupil-plane coordinates and paraxial focal-plane angles, respectively, while their discrete counterparts are the index pairs $[m, n]$ and $[p, q]$.

In order to be consistent with the notation used elsewhere, in this article, we have adopted the convention that the indices $[m, n]$ correspond directly with the continuous variables (x, y), so that x_m and y_n denote the mth x location and nth y location, respectively. This is distinct from the notation used by many programming languages such as MATLAB and Python, in which $[m, n]$ indexes an array in [row, column] format.

The sampled representation of a continuous function is represented explicitly by appending an index variable after the continuous coordinate arguments, such as $g(x, y)[m, n]$, which denotes a two-dimensional discrete function, indexed by the sequences $[m, n]$, which is obtained by sampling the continuous function $g(x, y)$. However, in many cases, for the sake of notational brevity, sampled functions are denoted directly using integer indices, i.e. $g[m, n]$. In these cases, the notation $g(x, y)$ and $g[m, n]$ are implicitly understood as the continuous function g and its discrete, sampled representation, respectively.

2. BASIC FORWARD MODELING

APLCs consist of three sequential diffractive masks, as illustrated in Fig. 1: an apodizing mask $A(x, y)$ placed at a reimaged telescope pupil $P(x, y)$ (plane A), a small on-axis focal-plane spot called the occulter (plane B) represented by the transmittance $1 - M(\theta_x, \theta_y)$, and a Lyot stop $L(x, y)$ placed in a second reimaged pupil plane (plane C), which suppresses stellar sidelobe energy in the dark zone not directly affected by the occulter.

The field in the Lyot stop plane is then Fourier transformed onto a detector, where the final image is formed.

Let ψ_A, $\tilde{\psi}_B$, ψ_C, and $\tilde{\psi}_D$ denote the fields in the plane of the apodizer, occulter, Lyot stop, and detector, respectively, and let $\psi(x, y)$ denote the field illuminating the coronagraph entrance pupil. The tildes in this notation are employed to remind us that these symbols represent optical fields in image (focal) planes. Assuming that the coronagraph optics are positioned so that the pupil and focal-planes are related by simple Fourier transforms, the simplest continuous-domain forward model used to calculate an on-axis stellar image is

$$\psi_A(x, y) = P(x, y)A(x, y)\psi(x, y), \tag{1}$$

$$\tilde{\psi}_B(\theta_x, \theta_y) = \left[1 - M(\theta_x, \theta_y)\right] \mathcal{F}\left\{\psi_A(x, y)\right\}(\theta_x, \theta_y), \tag{2}$$

$$\psi_C(x, y) = L(x, y)\mathcal{F}\left\{\tilde{\psi}_B(\theta_x, \theta_y)\right\}(x, y), \tag{3}$$

$$\tilde{\psi}_D(\theta_x, \theta_y) = \mathcal{F}\left\{\psi_C(x, y)\right\}(\theta_x, \theta_y), \tag{4}$$

where each of the fields is given after interacting with the mask in the plane in which they are defined, where the Fourier transform $\mathcal{F}\{\cdot\}$ is defined as

$$\mathcal{F}\{g(x, y)\}(f_x, f_y) = \iint_{-\infty}^{\infty} g(x, y) \times \exp\left\{-i2\pi(x f_x + y f_y)\right\} \mathrm{d}x\mathrm{d}y, \tag{5}$$

and recalling that the spatial frequency variables (f_x, f_y) are proportional to the paraxial angles (θ_x, θ_y) via

$$(f_x, f_y) = \left(\frac{\theta_x}{\lambda_0}, \frac{\theta_y}{\lambda_0}\right), \tag{6}$$

where λ_0 is the center wavelength of interest, assuming narrowband light for any given propagation. In APLCs, it is usually the case that the support of the apodizing mask is a subset of the support of the pupil, and thus $A(x, y)P(x, y) = A(x, y)$. Additionally, when determining the as-designed (aberration-free) stellar image, we generally let $\psi(x, y) = 1$ (a unit-amplitude on-axis plane wave), so that $\psi_A(x, y) = A(x, y)$. Note, however, that this assumption is not valid for coronagraph designs that utilize deformable mirrors to shape the input field prior to entering the coronagraph instrument.

To derive a discrete model suitable for numerical diffraction calculations, we sample the input and output variables with N_A and N_B datapoints, respectively, along each dimension, defining the discrete coordinates $x_m \triangleq m\Delta x$, $y_n \triangleq n\Delta x$, $f_{x,p} \triangleq p\Delta f_x$, $f_{y,q} \triangleq q\Delta f_x$, with $0 \le m, n < N_A$ and $0 \le p, q < N_B$, and assuming for simplicity that the sample spacing is identical along both directions. Approximating Eq. (5) as a Riemann sum, we obtain the discrete Fourier transformation

$$\mathrm{DFT}\left\{g[m, n]\right\}[p, q] = (\Delta x)^2 \sum_{m=0}^{N_A-1} \sum_{n=0}^{N_A-1} g[m, n] \times \exp\left\{-i2\pi \Delta x \Delta f_x (mp + nq)\right\}. \tag{7}$$

The Nyquist limit, which determines the highest-frequency spatial sinusoid representable without frequency aliasing, is given by

$$f_N = \frac{1}{2\Delta x}, \tag{8}$$

which, from Eq. (6), implies that, for a given pupil-plane sample spacing, the sub-Nyquist region of the focal-plane has a half-width (half-bandwidth) of

$$\theta_{\max} = \frac{\lambda_0}{2\Delta x}. \tag{9}$$

For a pupil of diameter D sampled on an $N_A \times N_A$ input array with no zero-padding, we have $\Delta x = D/N_A$, and, hence,

$$\theta_{\max} = \frac{\lambda_0 N_A}{2D} = \frac{N_A}{2}\frac{\lambda_0}{D}. \tag{10}$$

If we fix the number of datapoints so that $N_B = QN_A \triangleq N$, where Q is an integer describing the embedding of the $N_A \times N_A$ pupil-domain array in a larger $QN_A \times QN_A$ array of zeros, and additionally constrain the input and output sample spacings to satisfy $\Delta x \Delta f_x = 1/N$, then Eq. (7) reduces to

$$(\Delta x)^2 \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} g[m, n] \exp\left\{-i2\pi(mp + nq)/N\right\}, \tag{11}$$

which is implemented with high efficiency by the family of fast Fourier transform (FFT) algorithms. In this article, we draw an explicit distinction between the general discrete Fourier transformation as defined in Eq. (7), and the FFT implementation based on Eq. (11). Alternative implementations of the DFT, such as the matrix triple product (MTP), also called the matrix Fourier transform (MFT) [34,35], enable the DFT to be computed with more flexibility, such as arbitrary output sample spacing and field of view, than with the rigid sampling parameters imposed by the FFT, possibly at the cost of reduced speed and increased memory requirements.

Using the above results, we can formulate a basic approach to numerically computing the on-axis stellar image resulting from an APLC as follows:

$$\psi_A[m, n] = P[m, n]A[m, n]\psi[m, n], \tag{12}$$

$$\tilde{\psi}_B[p, q] = (1 - M[p, q])\, \mathrm{DFT}\left\{\psi_A[m, n]\right\}[p, q], \tag{13}$$

$$\psi_C[m, n] = L[m, n]\mathrm{DFT}\left\{\tilde{\psi}_B[p, q]\right\}[m, n], \tag{14}$$

$$\tilde{\psi}_D[p, q] = \mathrm{DFT}\left\{\psi_C[m, n]\right\}[p, q]. \tag{15}$$

The semianalytical algorithm developed by Soummer *et al.* [35] provides a highly efficient implementation of this forward model in APLCs by utilizing Babinet's principle. Expanding Eq. (14) using the definition of $\tilde{\psi}_B$ in Eq. (13), and temporarily ignoring array indices for brevity, we have

$$\psi_C = L\mathrm{DFT}\left\{\mathrm{DFT}\left\{\psi_A\right\} - M\mathrm{DFT}\left\{\psi_A\right\}\right\}. \tag{16}$$

Defining the quantity $\psi'_B \triangleq M\mathrm{DFT}\{\psi_A\}$ as the field impinging on the opaque part of the occulter, and recognizing that $\mathrm{DFT}\{\mathrm{DFT}\{A[m,n]\}\} = A[-m,-n]$, we obtain

$$\psi_C[m,n] = L(\psi_A[-m,-n] - \mathrm{DFT}\{\psi'_B\}). \quad \textbf{(17)}$$

Hence, the field immediately before the Lyot stop may be computed by transforming the occulter-plane field evaluated within the support of the occulter only, and subtracting the result from the coordinate-reversed apodizing mask. Because the occulter is generally small ($<5\lambda_0/D$) in comparison to $\theta_{\max}$, this insight, coupled with the flexibility of the MTP algorithm, enables a dramatic reduction in required computational resources.

Up to this point, we have considered only the simplest four-plane coronagraph model. The Gaussian pilot beam algorithm utilized in the PROPER and POPPY optical modeling packages [36,37] for full end-to-end Fresnel propagation is outside the scope of this article; however, the results derived in the following sections remain valid.

A. Including a Focal-Plane Field Stop

If a field stop with transmittance $F(\theta_x, \theta_y)$ with radius R_F is inserted into the occulter-plane, then Eq. (2) becomes

$$\tilde{\psi}_B(\theta_x,\theta_y) = [F(\theta_x,\theta_y) - M(\theta_x,\theta_y)]\mathcal{F}\{\psi_A(x,y)\}(\theta_x,\theta_y), \quad \textbf{(18)}$$

and its discrete counterpart in Eq. (13) becomes

$$\tilde{\psi}_B[p,q] = (F[p,q] - M[p,q])\,\mathrm{DFT}\{\psi_A[m,n]\}[p,q]. \quad \textbf{(19)}$$

Here, we are presented with a problem: if the radius of the focal-plane stop is larger than $\sqrt{2}\theta_{\max}$, the length from the center to the corner of the computational array, then effectively $F[p,q] = 1$, and diffraction from its edges is not included in the model. If $\theta_{\max} \le R_F \le \sqrt{2}\theta_{\max}$, then the field will be partially clipped by the edge of $F[p,q]$, and the effects of the field stop will only be partially (and incorrectly) modeled. Fortunately, the unique geometry of the pupil-plane APLC masks enables us to increase the focal-plane angular bandwidth simply and efficiently to overcome this problem, as we will show in the next section.

The expression for the Lyot-plane field in Eq. (17) no longer holds in the presence of a field stop because $\mathrm{DFT}\{F\mathrm{DFT}\{A\}\} \ne A[-m,-n]$, and so the semianalytical algorithm in most basic form is unable to model the field stop, regardless of radius. To circumvent this, one can calculate the direct field $\mathrm{DFT}\{F\mathrm{DFT}\{A\}\}$ with much coarser sample spacing in the focal-plane than the occulter field ψ'_B, which is slower than the basic semianalytical algorithm but retains a computational advantage over the "direct" model described in Eqs. (12)–(15).

3. SAMPLING AND DIFFRACTION CALCULATIONS WITH BINARY PUPIL MASKS

Consider a pupil-plane apodizing mask $A(x,y)$ that can be described as a superposition of a regular grid of non-overlapping unit cell functions $E(x,y)$, such that

$$A(x,y) = \sum_m \sum_n A[m,n] E\left(\frac{x - m\Delta x}{\Delta x}, \frac{y - n\Delta x}{\Delta x}\right). \quad \textbf{(20)}$$

Here, the integer pair $[m,n]$ with $0 \le m \le N_C - 1$ and $0 \le n \le N_C - 1$ is used to index the discrete set of unit cells, weighted by the coefficients $A[m,n]$, which are simply the samples of the numerical apodizing mask solution. For the masks considered here, which are fabricated using a photo-lithographic manufacturing process, we choose $E(x,y) = \mathrm{rect}(x/\Delta x)\mathrm{rect}(y/\Delta y)$.

A. Mask Upsampling

The simplest approach to increase the computed focal-plane area is to upsample the apodizer so that each unit cell is sampled by multiple array elements. Denoting the apodizer upsampling operation by integer factor K by $\mathbb{S}\{A[m,n];K\}[m',n']$, where m' and n' represent the indices of the upsampled pupil-plane coordinates with step size $\Delta x/K$, the propagation approach becomes

$$\psi_A[m',n'] = \mathbb{S}\{A[m,n];K\}[m',n'], \quad \textbf{(21)}$$

$$\tilde{\psi}_B[p,q] = (1 - M[p,q])\,\mathrm{DFT}\left\{\psi_A[m',n']\right\}[p,q], \quad \textbf{(22)}$$

$$\psi_C[m',n'] = \mathbb{S}\{L[m,n];K\}[m',n']\,\mathrm{DFT} \times \left\{\tilde{\psi}_B[p,q]\right\}[m',n'], \quad \textbf{(23)}$$

$$\tilde{\psi}_D[p,q] = \mathrm{DFT}\left\{\psi_C[m',n']\right\}. \quad \textbf{(24)}$$

Decreasing sample spacing by K in a pupil plane causes the Nyquist frequency f_N in Eq. (8), and hence the angular bandwidth $\theta_{\max}$ in Eq. (10), to increase commensurately by K. Note in Eq. (23) that the Lyot stop is upsampled in addition to the apodizer; this is necessary in order for the components at the edge of the enlarged occulter-plane to be adequately sampled in the Lyot plane.

For upsampling general unit cell geometries, an appropriate numerical algorithm for approximating the desired shape on a discrete grid must be employed, which may involve computations of nontrivial complexity. However, for square unit cells, upsampling can be performed efficiently and exactly using a *tile-upsampling* operation $\mathbb{S}_T\{A[m,n];K\}$, in which each element of $A[m,n]$ is replaced by a $K \times K$ square of identical samples to produce a new array K times larger along each dimension.

For any finite value of K, upsampling the apodizer produces an occulter-plane field that is an approximation to the true (continuous variable) field, and due to the introduction of larger arrays, the computational complexity of the propagation increases. However, upsampling is simple to incorporate into propagation models, and as the upsampling factor K tends to infinity, the result of Eq. (21) converges to the true field in the occulter-plane to within the paraxial approximation, because $\mathrm{DFT}\{E[m',n']\}$, which is a Riemann sum approximation to $\mathcal{F}\{E(x,y)\}$, converges asymptotically to $\mathcal{F}\{E(x,y)\}[m',n']$ as the pupil-plane sample spacing decreases.

B. Limiting Behavior: The Analytical Envelope

We can alternatively adopt a more efficient analytical approach. Taking the continuous Fourier transform of $A(x, y)$ given in Eq. (20) and using the linearity property, the Fourier similarity (scaling) theorem, and the Fourier shift theorem, we obtain

$$\begin{aligned}\mathcal{F}\{A(x,y)\} &= \sum_m \sum_n A[m,n]\mathcal{F}\left\{E\left(\frac{x-m\Delta x}{\Delta x}, \frac{y-n\Delta x}{\Delta x}\right)\right\} \\ &= (\Delta x)^2 \tilde{E}(\Delta x f_x, \Delta x f_y)\sum_m \sum_n A[m,n] \\ &\quad \times \exp\left\{-i2\pi(x_m f_x + y_n f_y)\right\},\end{aligned} \tag{25}$$

where the envelope function $\tilde{E}(f_x, f_y)$ is the continuous Fourier transform of the unit cell function $E(x, y)$, and we have defined $x_m \triangleq m\Delta x$ and $y_n \triangleq n\Delta x$ as discrete spatial coordinates.

The summation in Eq. (25) is a two-dimensional analogue of the discrete-time Fourier transform (DTFT) [38] with discrete spatial coordinates and continuous spatial frequency coordinates. Sampling this "discrete-space" Fourier transform in spatial frequency by defining $f_{x,p} \triangleq p\Delta f_x$ and $f_{y,q} \triangleq q\Delta f_y$, we obtain a discrete, frequency-domain function indexed by the integer pair $[p, q]$,

$$\begin{aligned}\mathcal{F}\{A(x,y)\}[p,q] &= (\Delta x)^2 \tilde{E}(\Delta x f_x, \Delta x f_y)[p,q] \\ &\quad \times \sum_m \sum_n A[m,n]\exp\left\{-i2\pi(x_m f_{x,q} + y_n f_{y,p})\right\},\end{aligned} \tag{26}$$

which is exactly equal to

$$\begin{aligned}\mathcal{F}\{A(x,y)\}[p,q] &= \tilde{E}(\Delta x f_x, \Delta x f_y)[p,q] \\ &\quad \times \mathrm{DFT}\{A[m,n]\}[p,q].\end{aligned} \tag{27}$$

In other words, in the paraxial regime, the *exact* field diffracted by the apodizing mask may be obtained numerically by computing the DFT of the numerical apodizer solution (with no upsampling), and multiplying the result by a discretely sampled envelope function, which has an analytical expression for certain choices of unit cell geometry. For instance, for the square unit cell case considered here, the paraxial focal-plane field is given by

$$\begin{aligned}\mathcal{F}\{A(x,y)\}[p,q] &= \mathrm{sinc}(\Delta x f_x)\mathrm{sinc}(\Delta x f_y)[p,q] \\ &\quad \times \mathrm{DFT}\{A[m,n]\}[p,q].\end{aligned} \tag{28}$$

Note that we have made no assumptions about the values of the weights $A[m, n]$. Thus, the above formulation is true for any array $A[m, n]$, and is not limited to binary masks.

As a simple demonstration of this principle, consider a one-dimensional rectangle function $g(x)$ composed of L subrectangles with unit cell width Δx, i.e.,

$$g(x) = \mathrm{rect}\left(\frac{x}{L\Delta x}\right) = \sum_{k=0}^{L-1}\mathrm{rect}\left(\frac{x-k\Delta x}{\Delta x}\right). \tag{29}$$

Also consider the discrete representation $g[n]$ of $g(x)$ comprised of L nonzero samples,

$$g[n] = \mathrm{rect}\left(\frac{n\Delta x}{L\Delta x}\right) = \mathrm{rect}\left(\frac{n}{L}\right). \tag{30}$$

The length-N DFT of $g[n]$ from x to frequency variable f, indexed by p, is the Dirichlet kernel with period N,

$$\mathrm{DFT}\{g[n]\}[p] = \frac{\sin(\pi Lp/N)}{\sin(\pi p/N)}. \tag{31}$$

The unit cell considered here has the form $\mathrm{rect}(x/\Delta x)$, and its continuous Fourier transform is

$$\mathcal{F}\left\{\mathrm{rect}\left(\frac{x}{\Delta x}\right)\right\} = \Delta x\frac{\sin(\pi\Delta x f)}{\pi\Delta x f} = \Delta x\,\mathrm{sinc}(\Delta x f). \tag{32}$$

Sampling using $f[p] = p\Delta f$ with $\Delta f\Delta x = 1/N$, and multiplying by the DFT of $g[n]$,

$$\begin{aligned}\mathcal{F}\left\{\mathrm{rect}\left(\frac{x}{\Delta x}\right)\right\}\mathrm{DFT}\{g[n]\}[p] &= \Delta x\frac{\sin(\pi\Delta x\Delta fp)}{\pi\Delta x\Delta fp}\frac{\sin(\pi Lp/N)}{\sin(\pi p/N)} \\ &= \Delta x\frac{\sin(\pi\Delta x\Delta fp)}{\pi\Delta x\Delta fp}\frac{\sin(\pi L\Delta x\Delta fp)}{\sin(\pi\Delta x\Delta fp)} \\ &= L\Delta x\frac{\sin(\pi L\Delta x\Delta fp)}{\pi L\Delta x\Delta fp} \\ &= (L\Delta x)\mathrm{sinc}((L\Delta x)(p\Delta f)),\end{aligned} \tag{33}$$

which is the continuous Fourier transform of $g(x)$, sampled onto a frequency axis with $f[p] = p\Delta f$.

This example is illustrated in Fig. 2. We observe that the analytical envelope function has a value of zero precisely at the center locations of all DFT periods except for the fundamental period, which is a consequence of the fact that for functions of the form in Eq. (20), the sample spacing of the array is equal to the width of the unit cells. After multiplication by the envelope function, the resulting field is identical to the analytical Fourier transform everywhere, even beyond the fundamental period of the DFT.

Similarly, by adopting a purely discrete model in the pupil plane, it is possible to obtain a discrete analog to the continuous-domain envelope derived above. For square unit cells, this envelope is given by the Dirichlet kernel (see Appendix A), which is the discrete-space Fourier transform of a rectangle function with K nonzero samples along each direction,

$$\tilde{E}_D(f_x, f_y) = \frac{1}{K^2}\frac{\sin(\pi Kf_x)}{\sin(\pi f_x)}\frac{\sin(\pi Kf_y)}{\sin(\pi f_y)}. \tag{34}$$

Denoting the DFT of the tile-upsampled apodizer by

$$\tilde{A}'[p,q] \triangleq \mathrm{DFT}\left\{\mathbb{S}_T\{A[m,n];K\}[m',n']\right\}[p,q], \tag{35}$$

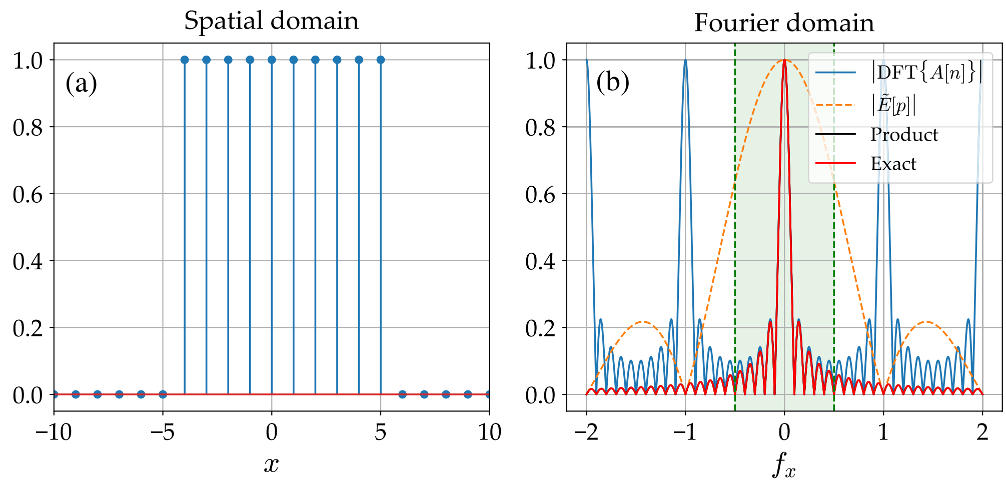

Fig. 2. A demonstration of the principle of exact Fourier transform recovery using the discrete Fourier transform and analytical envelope multiplication. (a) A 10-sample binary rectangle function, representing samples of a continuous-domain rectangle function. (b) (blue line) The magnitude of the DFT of the binary rectangle, extended to multiple periods (the shaded area denoting the fundamental period), along with the unit-amplitude correction term (orange dashed line), and their product (black), which is identical to the continuous Fourier transform of the original, continuous rectangle function (red solid line, behind which the black line is hidden).

we find that

$$\tilde{A}'[p,q] = \tilde{E}_D(\Delta x' f_x, \Delta x' f_y)[p,q] \mathrm{DFT}\{A[m,n]\}[p,q], \tag{36}$$

where $\Delta x' = \Delta x / K$ is the pupil-plane sample spacing after upsampling, and the product is performed elementwise. In other words, multiplying the DFT of the native-resolution apodizer by the Dirichlet envelope yields exactly the DFT of the tile-upsampled apodizer. As the upsampling factor K tends toward infinity and the pupil-plane sample spacing $\Delta x'$ tends to zero, we have (in one dimension)

$$\begin{aligned} \lim_{K\to\infty} \tilde{E}_D(\Delta x' f_x) &= \lim_{K\to\infty} \frac{\sin\left(\pi K \Delta x' f_x\right)}{K \sin\left(\pi \Delta x' f_x\right)} \\ &= \lim_{K\to\infty} \frac{\sin\left(\pi \Delta x f_x\right)}{K \sin\left(\pi (\Delta x / K) f_x\right)} \\ &= \frac{\sin\left(\pi \Delta x f_x\right)}{\pi \Delta x f_x} \\ &= \tilde{E}(\Delta x f_x). \end{aligned} \tag{37}$$

As shown above, since the continuous-domain sinc envelope $\tilde{E}(\Delta x f_x)$ provides the exact paraxial field when multiplied by the DFT of the apodizer, this provides a more rigorous argument that the numerical approximation of the paraxial focal-plane field approaches the true field as the pupil-plane upsampling factor increases.

The continuous-domain analytical envelope allows the occulter-plane field to be computed over an arbitrarily large area, in the paraxial regime, without the requirement of additional pupil-plane upsampling. However, because this envelope has theoretically infinite extent and is nonperiodic, calculating the corresponding exact field in the Lyot stop requires large values of K and extremely fine sample spacing to avoid Gibbs ringing artifacts in the Lyot stop field that corrupt the final stellar image. However, as we have shown, the Dirichlet envelope converges to the analytical case as K becomes large, and does not introduce Gibbs ringing in the Lyot stop plane, which makes the Dirichlet envelope a good choice in practice.

1. *Halftoning*

In general, when numerically optimizing apodizing masks for coronagraphs with complex pupil structure, mask solutions are mostly zero-one valued, but contain a small number of gray-level (partially transmissive) pixels. These gray-level solutions are typically used in first-order evaluations of inner and outer working angle (OWA) and dark-zone contrast. In our simulations, we have consequently assumed that the samples of the gray-level solution correspond one-to-one with the finite lithographic fabrication elements of a physically realized mask (see Section 1).

In reality, current fabrication processes can only represent binary (zero-one) reflectivity values at each pixel of a mask solution. Therefore, prior to fabrication, solutions are tile-upsampled and halftoned, commonly using the Floyd–Steinberg error diffusion algorithm [39], with the goal of producing a binary mask with similar diffraction properties to the gray-level solution within some acceptable degree of accuracy [13].

Because the halftoned masks also comprise regular arrays of square unit cells, the results derived in this work are equally applicable to this case. Furthermore, as mentioned in Section 1, the number of gray-level pixels is generally a small proportion of the total, and consequently the difference in the treatment and behavior of gray-level and halftoned masks is expected to be small.

The relationship between gray-level and halftoned masks, along with the behavior of halftoned masks when unit cell extent is accounted for, will be examined more closely in future work.

4. RESULTS

A. Effect of Field Stop on Coronagraph Contrast

We introduced a circular focal-plane field stop with the radius varying from $32\lambda_0/D$ to $1024\lambda_0/D$ into numerical models of two separate APLC designs for the LUVOIR telescope architecture "A" [1,40], denoted as LUVOIR-A 20170828 and LUVOIR-A 20180119. The 20180119 design has 512×512 sample pupil-plane masks and an annular dark zone from 3.5 to $12\lambda_0/D$, while the 20170828 design has 256×256 sample pupil-plane masks and an annular dark zone from 4 to $10\lambda_0/D$. Table 1 summarizes these parameters, while Fig. 3 shows the pupil-plane masks for each design. Figure 4 shows aberration-free stellar images without a field stop or pupil-plane upsampling. Shown is the contrast $I(\theta_x, \theta_y)/I_{00}$, where $I(\theta_x, \theta_y)$ is the intensity of the stellar image at each location in the detector plane and I_{00} is the on-axis intensity of the unocculted stellar image.

For each focal-plane field stop radius R_F, the average dark-zone contrast for each design was computed as the average contrast within an annulus with an inner working angle (IWA) and OWA specified by each design, shown in Table 1. To probe the approximation error inherent in finite sample spacing, the mean contrast as a function of R_F was reevaluated for multiple values of the oversampling parameter K. All analysis was performed in the monochromatic, aberration-free regime with $0.25\lambda_0/D$ sample spacing in the final image plane.

Stellar images for each coronagraph were computed using the direct four-plane diffraction model outlined in Eqs. (12)–(15). Because of the large number of datapoints resulting in each plane when pupil-plane masks are upsampled, several speed optimizations were implemented to reduce computational overhead. First, the DFT was implemented using the MTP algorithm described in Section 2, which, unlike the FFT, does not require pupil-plane arrays to be zero-padded. Second, the flexible output field of view of the MTP algorithm was used to evaluate the occulter-plane field, up to the smaller of $\theta_{\max}$ or R_F along each axis, in order to reduce the number of datapoints in the occulter-plane for field stop radii smaller than the available focal-plane angular bandwidth. Similarly, the detector-plane field was evaluated only up to the OWA for each coronagraph design, which is typically much smaller than $\theta_{\max}$. Finally, the Dirichlet envelope, Eq. (34), was utilized to evaluate the DFT of the tile-upsampled apodizer $\tilde{A}'[p, q]$ using the DFT of the native-resolution apodizer given in Eqs. (35)-(36), further reducing the number of datapoints in each propagation.

As a point of reference, the most computationally expensive optical propagation in the study was the case for $K = 16$ oversampling of the LUVOIR-A 20180119 pupil-plane masks. In this case, the pupil-plane arrays were 8192×8192 samples in size, and the occulter-plane array was 8200×8200 samples in size, representing a half-width of $1025\lambda_0/D$ sampled at $0.25\lambda_0/D$ per pixel, which ensured that the $1024\lambda_0/D$ field stop radius was entirely included in the array. The final image-plane array was 104×104 samples in size, representing a half-width of $13\lambda_0/D$ at $0.25\lambda_0/D$ per pixel. Performing the same analysis using the FFT algorithm, which would require the pupil-plane arrays to be zero-padded by a factor of four to achieve the equivalent image-plane sample spacing, would have resulted in arrays of size $32{,}678 \times 32{,}678$ samples at every plane, representing a dramatic increase in required computational resources.

Figure 5 shows the results of this analysis. In general, as one would expect, the effects of field stop diffraction diminish as the radius of the field stop increases. However, because of the reasons outlined in Section 2.A, as the field stop becomes larger than the available bandwidth for any particular value of K, its effects are underestimated, or lost entirely, because the edge of the field stop is no longer explicitly represented, either wholly or in part, in the computational array. We can see this by observing that for a fixed field stop radius (horizontal axis), upsampling the pupil-plane masks to include the edge of the field stop in the occulter-plane results in a degradation of the predicted coronagraph contrast, in some cases by multiple orders of magnitude. This pattern was true for every field stop radius considered in this study.

Notice in Fig. 5 that at the largest field stop radius considered, $1024\lambda_0/D$, the contrast target of 10^{-10} is missed by approximately an order of magnitude in both designs. At this radius, the field stop edge only falls entirely within the available focal-plane angular bandwidth for $K \geq 8$ for the LUVOIR-A 20170828 design, which has $N_A = 256$ samples across the entrance pupil, or for $K \geq 4$ for the LUVOIR-A 20180119 design, which has $N_A = 512$, and therefore these datasets are the only reliable indicators of coronagraph performance with this particular field stop size.

Table 1. Number of Samples Across Pupil-Plane Masks N_A, Occulter Radius R_M, Inner Working Angle (IWA), Outer Working Angle (OWA), and Contrast Target for the Two Coronagraph Designs Analyzed in This Paper[a]

	N_A	$R_M[\lambda_0/D]$	IWA $[\lambda_0/D]$	OWA $[\lambda_0/D]$	Contrast Target
LUVOIR-A 20170828	256	4	4	10	10^{-10}
LUVOIR-A 20180119	512	3.5	3.5	12	10^{-10}

[a]Focal-plane quantities are given in diffraction widths λ_0/D, where λ_0 is the central propagation wavelength and D is the diameter of the coronagraph entrance pupil.

B. Mitigation Using Closed-Loop Wavefront Control

In many scenarios where aberrations and imperfections in the optical train degrade coronagraph contrast, closed-loop wavefront control is able to recover the as-designed contrast target. To understand the extent to which closed-loop wavefront control is able to compensate for the diffraction from the focal-plane field stop, we introduced a pair of deformable mirrors as shown in Fig. 6, one in the entrance pupil of the coronagraph and one located at Fresnel number $N_F = D_{\mathrm{DM}}^2/(4\lambda_0 z) = 1562.5$ away from the entrance pupil, where D_{DM} is the beam diameter at the pupil-plane deformable mirror, λ_0 is the center wavelength, and z is the inter-DM distance along the optical axis. Using two deformable mirrors in this fashion is a standard approach to coronagraphic wavefront control [41]. For an assumed

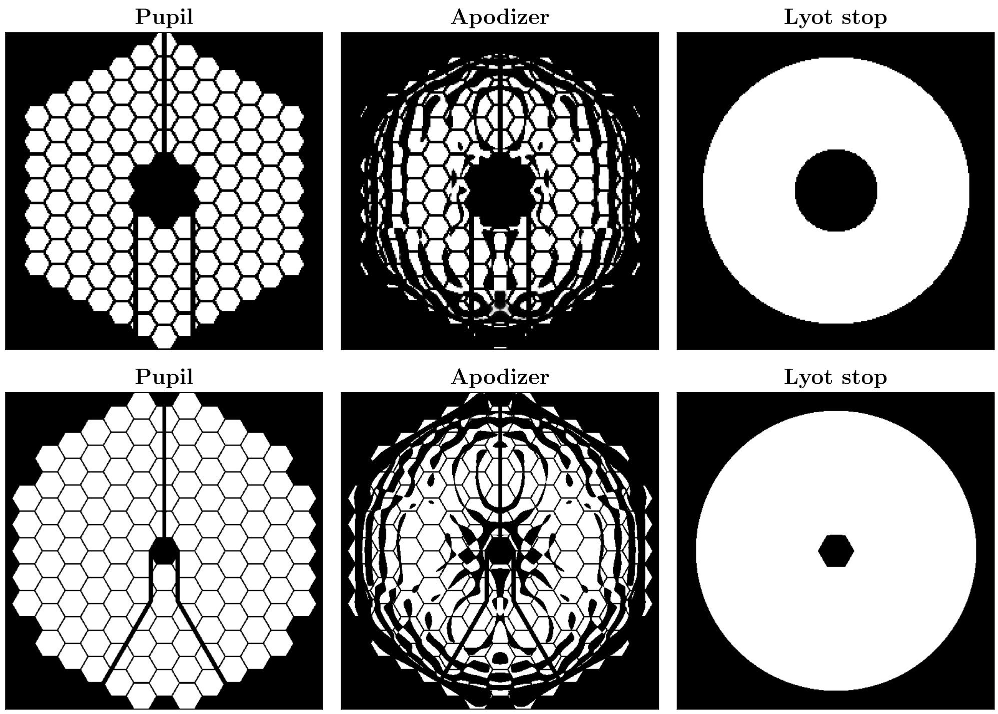

Fig. 3. Pupil-plane masks for the LUVOIR-A 20170828 (top) and 20180119 (bottom) APLC designs.

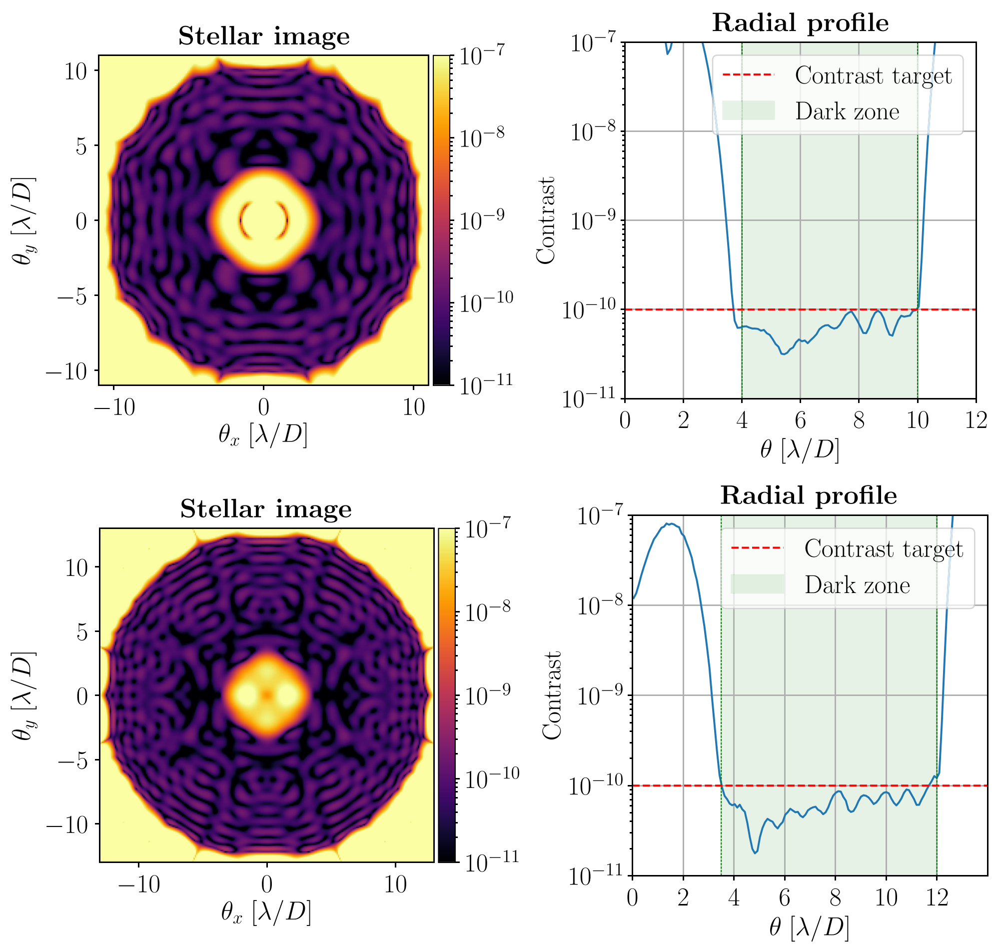

Fig. 4. On-axis, aberration-free stellar image with no pupil-plane upsampling, from the LUVOIR-A 20170828 (top) and 20180119 (bottom) APLC designs, whose pupil-plane masks are shown in Fig. 3. Azimuthal averages of each image are shown on the right.

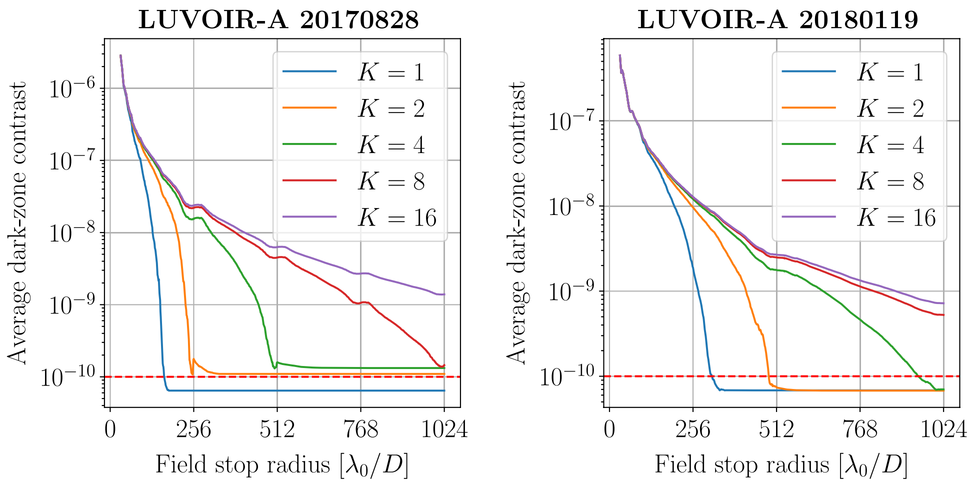

Fig. 5. Average dark-zone contrast as a function of field stop radius, for two different APLC designs and for multiple values of upsampling factor K. The contrast target for each design is marked by a red dashed line.

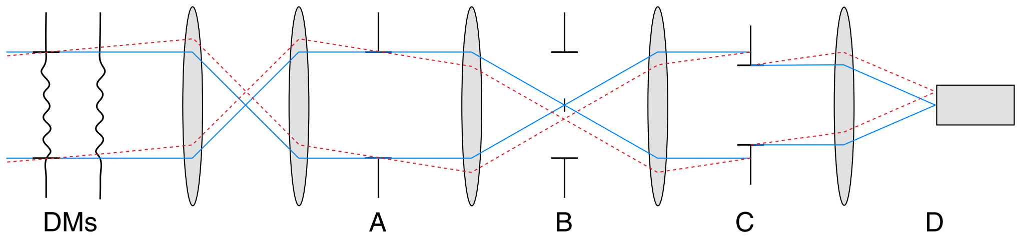

Fig. 6. Lyot coronagraph layout with deformable mirrors (DMs) for active wavefront control. The first deformable mirror is coincident with the entrance pupil of the coronagraph, and is conjugate to the apodizer plane (plane A). The second deformable mirror is in an intermediate plane at a distance corresponding to Fresnel number $N_F = 1562.5$. The two deformable mirrors together enable simultaneous compensation of amplitude and phase aberrations over a 360° dark zone. As in Figure 1, solid blue rays and dashed red rays represent stellar and planetary light, respectively.

deformable mirror diameter of 50 mm and a 500 nm center wavelength, this Fresnel number corresponds to a distance of 800 mm, which is consistent with the optical design of the LUVOIR-A coronagraph [42]. The deformable mirrors were each modeled as having a 48×48 grid of actuators over the coronagraph entrance pupil, and utilized the influence function model used by the PROPER optical modeling library [36].

We simulated monochromatic wavefront control loops at $\lambda_0 = 500$ nm for both coronagraph designs shown in Fig. 3 with field stops of varying radius. Moreover, we simulated each choice of field stop radius both with no pupil-plane upsampling ($K = 1$), and with $K = 16$, the largest value considered in the forward-modeling simulations described in the previous section. The control loop utilized the stroke minimization algorithm [41,43] to iteratively compute the optimal deformable mirror commands, and assumed perfect knowledge of the electric field in the coronagraph dark zone. No noise was introduced into the model. For each combination of coronagraph, field stop radius, and upsampling factor, we simulated 80 control iterations, which we chose so as to enable all simulated loops to converge to the minimum-achievable contrast. The control matrix (Jacobian matrix) used by the control algorithm was generated by a model with no field stop or pupil-plane upsampling, so that the results represent the achievable contrast without *a priori* knowledge of either effect. In each control iteration, a multiplicative decrease in total integrated dark-zone intensity of 0.8 over the previous iteration was requested. Because the Dirichlet envelope given in Eq. (34) and utilized in the forward model in Section 4.A is not strictly valid in the presence of non-plane wave illumination, we tile-upsampled the pupil-plane masks for the $K = 16$ case.

Figures 7 and 8 show the results of this analysis. Figure 7 shows a plot of mean dark-zone contrast versus field stop radius for both designs, with upsampling factors $K = 1$ and $K = 16$. For the LUVOIR-A 20170828 design, the smallest field stop radius for which the average dark-zone contrast after wavefront control is 10^{-10} or smaller is $350\lambda_0/D$, approximately 90 times the radius of the focal-plane mask. The corresponding value for the LUVOIR-A 20180119 design is $300\lambda_0/D$, approximately 85 times the radius of the focal-plane mask. As a point of reference, the as-designed field of view of the LUVOIR-A coronagraph is $64\lambda_0/D$ [42], so assuming that a field stop with a radius similar to this would be inserted into the coronagraph for stray light control is not unreasonable.

Figure 8 shows the root mean square (RMS) and peak-to-valley (PTV) deformable mirror actuator stroke corresponding to each case. Though the RMS actuator stroke is relatively small for all cases, on the order of single nanometers, the PTV stroke is considerably larger, due to the actuator stroke being strongly concentrated around a small number of actuators. The solutions for the LUVOIR-A 20170828 design with $R_F = 350\lambda_0/D$

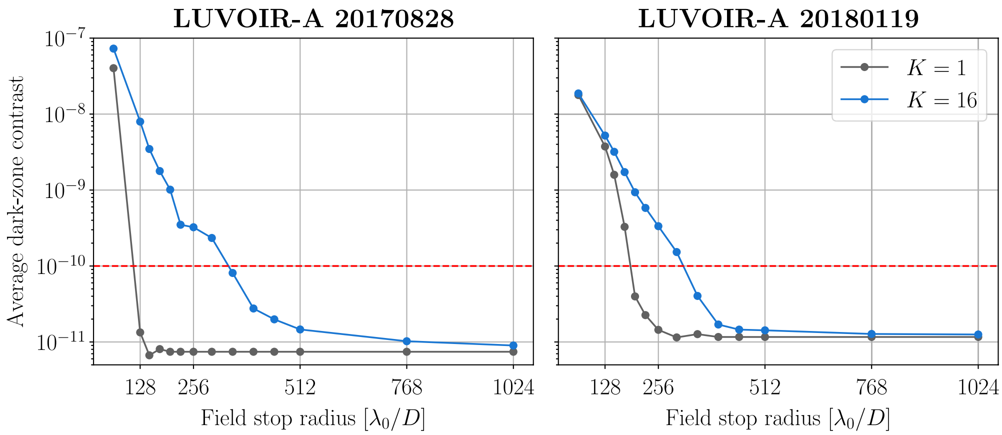

Fig. 7. Average dark-zone contrast as a function of field stop radius, for two different APLC designs and for two values of upsampling factor K, after closed-loop wavefront control. The contrast target for each design is marked by a red dashed line. The results for both coronagraphs are shown on the same vertical scale. For the most accurate model considered ($K = 16$), the smallest field stop for which 10^{-10} contrast can be recovered without *a priori* knowledge in the wavefront control forward model is approximately $350\lambda_0/D$ for the 20170828 design and $300\lambda_0/D$ for the 20180119 design. The less accurate model without pupil-plane upsampling ($K = 1$) fails to predict the full contrast degradation associated with finite field stop radius, with the discrepancy most significant near the nominal Nyquist limit, which is $128\lambda/D$ for the 20170828 design and $256\lambda/D$ for the 20180119 design.

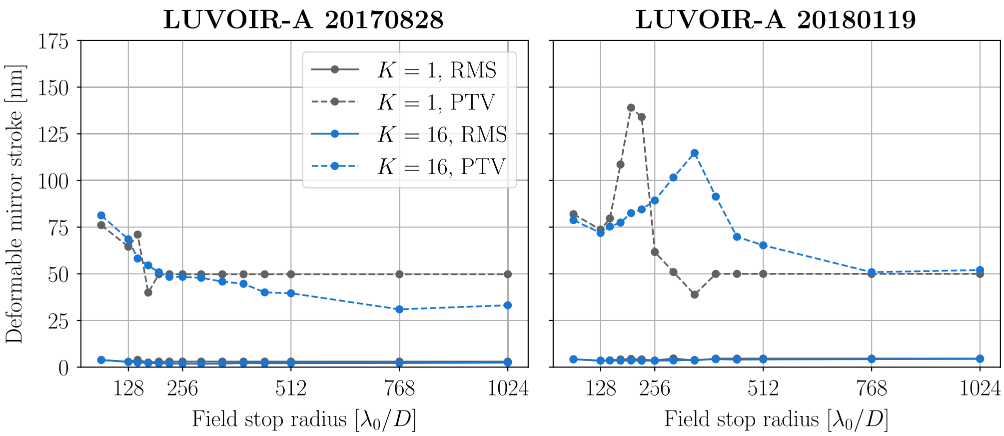

Fig. 8. Deformable mirror actuator stroke required to compensate for field stop diffraction effects, as a function of coronagraph, field stop radius, and pupil-plane upsampling factor. All commands are approximately zero mean. Though the root mean square (RMS) actuator stroke is on the order of nanometers, the commands display pronounced concentrated peaks, and the peak-to-valley (PTV) stroke is on the order of tens to hundreds of nanometers.

and $K = 16$, and for the LUVOIR-A 20180119 design with $R_F = 300\lambda_0/D$ and $K = 16$, are shown in Fig. 9. Because of the presence of pairs of adjacent actuators with strong commands of opposite sign, these solutions may be difficult to represent accurately on a continuous face-sheet deformable mirror in an experimental setting. The deformable mirror model used in this study approximates the total surface deformation as a linear combination of the individual influence functions of each actuator, and does not account for effects such as inter-actuator coupling and nonlinear deformation. Additionally, the concentration of actuator stroke near the edge of the pupil mask may make the solutions sensitive to misalignments between the deformable mirrors and the coronagraph entrance pupil. Therefore, the smallest-correctable radii reported above should be taken as a theoretical lower bound rather than a tight constraint.

C. Diffraction Effects from Tile-Upsampled Coronagraph Masks

Referring again to Fig. 5, we observe that for the LUVOIR-A 20170828 design, when the APLC masks are upsampled with $K > 1$, the mean dark-zone contrast does not return below 10^{-10} even when the field stop radius far exceeds the focal-plane bandwidth ($256\lambda_0/D$ for $K = 2$, or $512\lambda_0/D$ for $K = 4$). This indicates that some effect apart from field stop diffraction is introducing additional diffracted energy into the dark zone not captured by the nominal ($K = 1$) model. To study this more closely, we removed the field stop from the coronagraph model, tile-upsampled the apodizing mask and Lyot stop by $K = 8$, and computed the resultant stellar image, which is shown in Fig. 10(b). In the upsampled case, the stellar image contains additional high-energy rings within the dark zone, resulting in

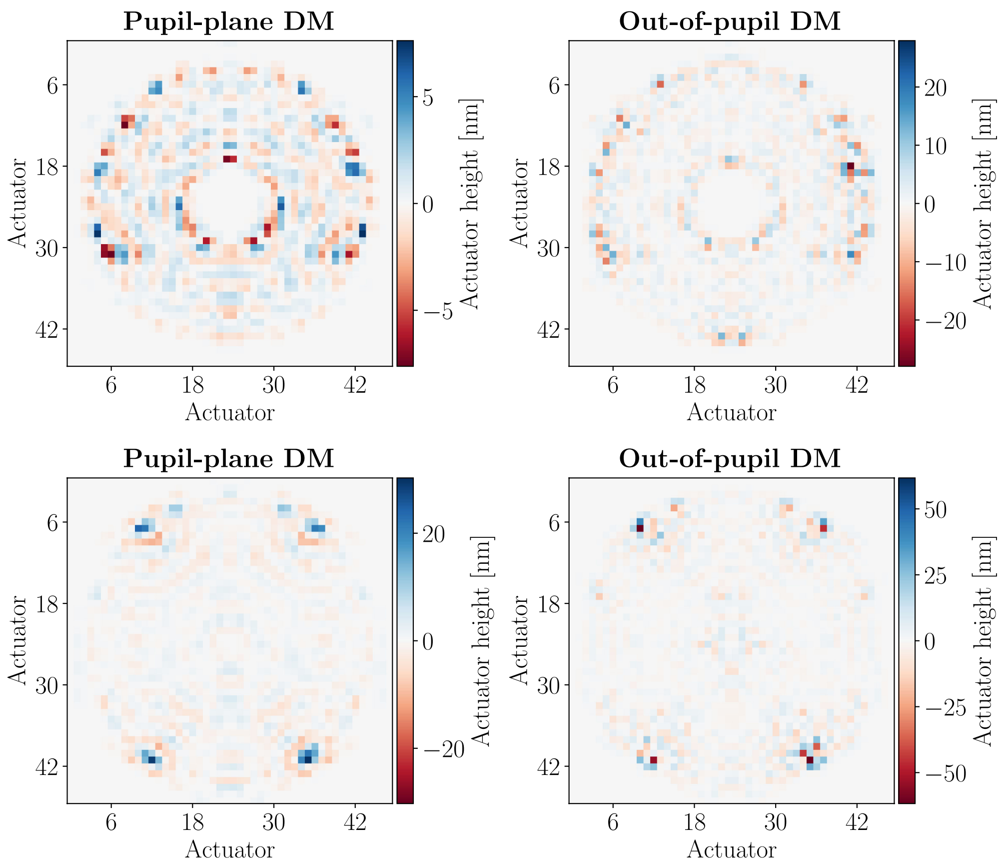

Fig. 9. Deformable mirror actuator commands for the LUVOIR-A 20170828 design with $R_F = 350\lambda_0/D$ and $K = 16$ (top) and the LUVOIR-A 20180119 design with $R_F = 300\lambda_0/D$ field stop and $K = 16$ (bottom). In both cases, the out-of-pupil deformable mirror command is highly concentrated along the outer edge of the pupil and contains strong commands with opposite sign on adjacent actuators.

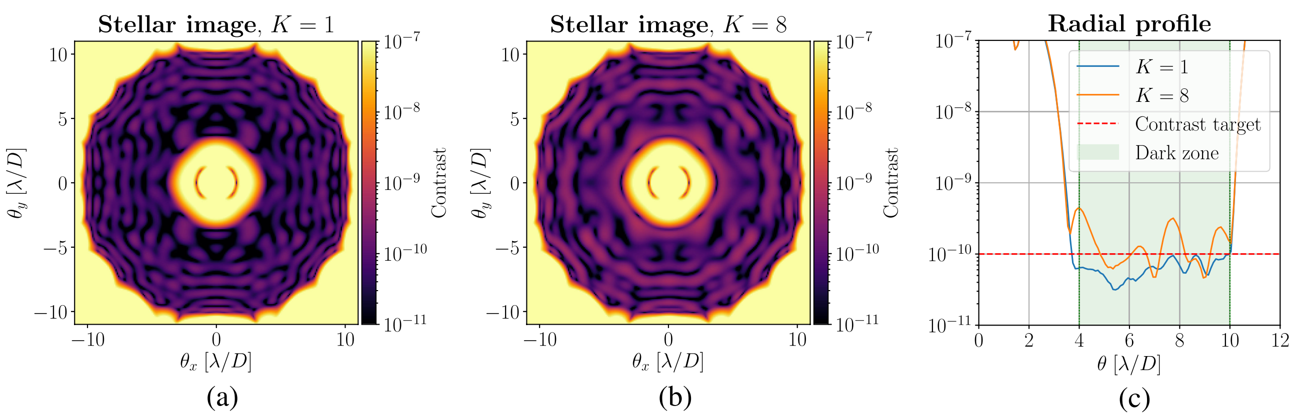

Fig. 10. Nominal stellar image from the LUVOIR-A 20170828 design (a) without and (b) with tile-upsampling of the pupil-plane masks by a factor of $K = 8$, and with no field stop. (c) Azimuthal averages of both images. Extra diffraction rings in the dark zone of the $K = 8$ image result in the coronagraph failing to achieve the desired 10^{-10} contrast target, indicated by a red dashed line.

the coronagraph failing to achieve the desired contrast target in several locations.

Recalling from Section 1 that the unit cells of the true, fabricated mask are well-modeled as squares, representing each cell as a $K \times K$ square of samples is more accurate in the spatial domain, and, indeed, we showed in Section 3.B that doing so causes the numerically calculated focal-plane field to become commensurately more accurate.

Another way to view this phenomenon is as follows. When each unit cell is represented with only a single sample, as in the $K = 1$ case, the contribution of each unit cell to the focal-plane field is a linear phase function only, because

$$\mathrm{DFT}\left\{\delta[n-k]\right\} = \exp\left\{-i2\pi(k\Delta x)(n\Delta f)\right\}, \tag{38}$$

where $\delta[n]$ is the Kronecker delta function. However, representing each unit cell by a $K \times K$ square of samples causes this unit-magnitude contribution to become modulated by the Dirichlet envelope in Eq. (34) and, hence, modifies the numerically calculated field in successive propagation steps.

As the number of samples across the pupil increases and the unit cells consequently become smaller, the extra dark-zone energy introduced by tile-upsampling the pupil-plane masks becomes negligible. From Fig. 5, we see that this is indeed the case for the LUVOIR-A 20180119 design, whose pupil-plane masks are 512×512 arrays (compared to the 256×256 masks for the 20170828 design), and for which the mean dark-zone contrast is identical for $K = 1$, $K = 2$, and $K = 4$ after the field stop radius exceeds the Nyquist bandwidth and vanishes from the propagation. Therefore, this effect functionally serves to impose a lower limit on the number of samples with which apodizing masks can be optimized and while still approximating the continuous-domain within an acceptable degree of accuracy. This lower limit will be probed more finely in future work.

D. Impact on Other Lyot Coronagraph Architectures

The diffraction effects we have described here are not limited to the APLC. In principle, any coronagraph architecture that makes the implicit assumption of a focal-plane with infinite extent will be susceptible to unintended field stop diffraction; this includes coronagraphs with focal-plane phase masks such as the vortex coronagraph [5–7], hybrid Lyot coronagraph [8,9], phase-masking coronagraphs [3,4], and the classical Lyot coronagraph. These architectures will be considered in future work. Shaped pupil Lyot coronagraph designs with spatially restricted focal-plane masks such as bowtie masks [33] would be unaffected, because the diffraction from the outer edge of the focal-plane is already incorporated in the design. However, at coarse mask resolutions, the effects of finite unit cell size may become non-negligible.

5. CONCLUSION

In this article, we have examined the predicted behavior of two APLC designs in the presence of a limiting field stop in the plane of the occulting mask, which may be explicitly included in the instrument optical design, or may functionally be imposed by the boundary of the filter used to support the occulting mask itself. We have described, analyzed, and implemented several computational methods that utilize the unique geometrical structure of binary pupil-shaping masks to extend the available angular bandwidth in the coronagraph focal-plane, including upsampling and focal-plane envelope functions. We utilized these techniques to capture diffraction effects from large field stops whose radii exceed the angular bandwidth imposed by the Nyquist limit in classical sampling theory. For masks with rectangular lithographic unit cells, the proposed focal-plane envelope functions can be used to recover the exact Fourier transform of the apodizing mask in the paraxial regime.

Next, we simulated closed-loop coronagraphic wavefront control to assess whether deformable mirrors may be used to compensate. Our analysis indicates that without *a priori* knowledge of the focal-plane field stop in the wavefront control forward model, 10^{-10} contrast may be recovered for field stops with radii as small as approximately $300{-}350\lambda/D$. Future work will analyze the impact of focal-plane field stops for other varieties of the Lyot coronagraph.

Finally, we numerically demonstrated and analyzed artifacts in computed stellar images that arise when pupil-plane coronagraph masks are upsampled, and become non-negligible for low-resolution coronagraph designs, suggesting a lower limit on the number of datapoints that may be used in the design algorithm. The precise nature of these artifacts, as well as a refined estimate of the lower limit on design resolution, will be explored in future work.

APPENDIX A: DERIVATION OF DIRICHLET FOCAL-PLANE ENVELOPE

In this appendix, we derive a discrete analogue of the continuous-domain analytical envelope function described in Section 3.B. We begin with the discrete-time Fourier transformation (DTFT), an operation that maps discrete functions to a continuous, periodic frequency domain. For a thorough analysis of the DTFT and its properties, see Oppenheim and Schafer [38].

The DTFT of a discrete function $g[n]$ is defined as

$$\mathrm{DTFT}\{g[n]\} \triangleq \sum_{n=-\infty}^{\infty} g[n] \exp\{-i\omega n\}. \tag{A1}$$

The output function is periodic in the continuous angular frequency variable ω with period 2π; for this reason, we generally restrict ω to the range $[-\pi, \pi)$.

Writing in terms of a unitless linear frequency variable $\bar{f} = \omega/2\pi$, we have

$$\mathrm{DTFT}\{g[n]\} = \sum_{n=-\infty}^{\infty} g[n] \exp\left\{-i2\pi \bar{f} n\right\}, \tag{A2}$$

with $\bar{f} \in \left[-\frac{1}{2}, \frac{1}{2}\right)$. For temporal signals, the Nyquist frequency is $f_N = 1/(2\Delta t)$, where Δt is the time-domain sample spacing. Because it represents the largest frequency that any signal sampled with Δt can contain, we can draw a direct correspondence between the Nyquist frequency and the normalized frequency $\bar{f} = \frac{1}{2}$. Using this, we may relate linear frequency f, in units of s^{-1}, to $\bar{f}$ via $f \Delta t = \bar{f}$, and thus write down the DTFT in terms of f as

$$\begin{aligned}\mathrm{DTFT}\{g[n]\} &= \sum_{n=-\infty}^{\infty} g[n] \exp\{-i2\pi n \Delta t f\} \\ &= \sum_{n=-\infty}^{\infty} g[n] \exp\{-i2\pi t_n f\},\end{aligned} \tag{A3}$$

with $f \in [-f_N, f_N)$, and where $t_n \triangleq n\Delta t$ is the nth sampling time interval.

Now we consider the DTFT of a signal represented as a superposition of non-overlapping, length-K discrete rectangle functions $\mathrm{rect}_D(\cdot)$, weighted by $g[n]$,

$$\mathrm{rect}_D\left(\frac{n'}{K}\right) \triangleq \begin{cases} 1 & 0 \le n' \le K-1 \\ 0 & \text{otherwise,} \end{cases}$$

$$g'[n'] = \sum_n g[n]\mathrm{rect}_D\left(\frac{n'-nK}{K}\right),$$

$$\mathrm{DTFT}\{g'[n']\}_{n'\to f} = \sum_n g[n]\,\mathrm{DTFT}\left\{\mathrm{rect}_D\left(\frac{n'-nK}{K}\right)\right\}_{n'\to f}. \quad \textbf{(A4)}$$

As the notation above suggests, $g'[n']$ and $g[n]$ correspond directly to the two-dimensional tile-upsampled and native-resolution masks, respectively, described in Section 3.A. The samples of $g'[n']$ are spaced apart by interval $\Delta t/K$ in time, so the Nyquist frequency for g' is

$$f'_N = \frac{1}{2\Delta t/K} = K f_N. \quad \textbf{(A5)}$$

Using the shift theorem of the DTFT,

$$\mathrm{DTFT}\{g'[n'-n]\}_{n'\to f} = \exp\left\{-i2\pi n\frac{\Delta t}{K}f\right\} \times \mathrm{DTFT}\{g'[n']\}_{n'\to f}, \quad \textbf{(A6)}$$

this yields

$$\begin{aligned}&\sum_n g[n]\mathrm{DTFT}\left\{\mathrm{rect}_D\left(\frac{n'-nK}{K}\right)\right\}_{n'\to f}\\ &= \sum_n g[n]\exp\left\{-i2\pi n\frac{\Delta t}{K}Kf\right\}\mathrm{DTFT}\left\{\mathrm{rect}_D\left(\frac{n'}{K}\right)\right\}_{n'\to f}\\ &= \mathrm{DTFT}\left\{\mathrm{rect}_D\left(\frac{n'}{K}\right)\right\}_{n'\to f}\sum_n g[n]\exp\{-i2\pi n\Delta t f\}. \end{aligned} \quad \textbf{(A7)}$$

Using Eq. (A3),

$$\mathrm{DTFT}\{g'[n']\}_{n'\to f} = \mathrm{DTFT}\left\{\mathrm{rect}_D\left(\frac{n'}{K}\right)\right\}_{n'\to f} \times \mathrm{DTFT}\{g[n]\}_{n\to f}. \quad \textbf{(A8)}$$

Finally, sampling all terms onto a common frequency axis f with step size Δf reduces the DTFT to the familiar discrete Fourier transform,

$$\mathrm{DFT}\{g'[n']\} = \mathrm{DFT}\left\{\mathrm{rect}_D\left(\frac{n'}{K}\right)\right\}\mathrm{DFT}\{g[n]\}. \quad \textbf{(A9)}$$

Extending Eq. (A3) to two dimensions, we define the *discrete-space Fourier transform* (DSFT) as

$$\mathrm{DSFT}\{g[m,n]\} = \sum_m\sum_n g[m,n]\exp\{-i2\pi(x_m f_x + y_n f_y)\}. \quad \textbf{(A10)}$$

As before, in the above example, x_m and y_n describe a discrete set of spatial locations corresponding to the samples of the function $g[m,n]$, and f_x and f_y are continuous spatial frequency variables. Sampling f_x and f_y similarly yields the desired expression in terms of the two-dimensional DFT,

$$\mathrm{DFT}\{g'[m',n']\} = \mathrm{DFT}\left\{\mathrm{rect}_D\left(\frac{m'}{K}\right)\mathrm{rect}_D\left(\frac{n'}{K}\right)\right\} \times \mathrm{DFT}\{g[m,n]\}. \quad \textbf{(A11)}$$

As in the continuous-domain example derived in Section 3.B, the expression $\mathrm{DFT}\{\mathrm{rect}_D(\frac{m'}{K})\mathrm{rect}_D(\frac{n'}{K})\}$ can be evaluated analytically, since it has a closed-form expression given in terms of continuous frequency by

$$\mathrm{DSFT}\left\{\mathrm{rect}_D\left(\frac{m'}{K}\right)\mathrm{rect}_D\left(\frac{n'}{K}\right)\right\} = \frac{1}{K^2}\frac{\sin(\pi\Delta x f_x)}{\sin(\pi\Delta x f_x/K)}\frac{\sin(\pi\Delta x f_y)}{\sin(\pi\Delta x f_y/K)}. \quad \textbf{(A12)}$$

Funding. Goddard Space Flight Center (80NSSC18K0213).

Acknowledgment. The research reported here was performed at the Institute of Optics of the University of Rochester. The authors thank Matthew Bolcar, Roser Juanola-Parramon, Neil Zimmerman, Rémi Soummer, Laurent Pueyo, Christopher Stark, and Marshall Perrin for helpful discussions.

Disclosures. The authors declare no conflicts of interest.

REFERENCES

1. L. Pueyo, N. Zimmerman, M. Bolcar, T. Groff, C. Stark, G. Ruane, J. Jewell, R. Soummer, K. S. Laurent, J. Wang, D. Redding, J. Mazoyer, K. Fograty, R. Juanola-Parramon, S. Domagal-Goldman, A. Roberge, O. Guyon, and A. Mandell, "The LUVOIR architecture 'A' coronagraph instrument," Proc. SPIE **10398**, 103980F (2017).
2. I. Poberezhskiy, F. Zhao, X. An, K. Balasubramanian, R. Belikov, E. Cady, R. Demers, R. Diaz, Q. Gong, B. Gordon, R. Goullioud, F. Greer, O. Guyon, M. Hoenk, J. N. Kasdin, B. Kern, J. Krist, A. Kuhnert, M. McElwain, B. Mennesson, D. Moody, R. Muller, B. Nemati, K. Patterson, A. Riggs, D. Ryan, B.-J. Seo, S. Shaklan, E. Sidick, F. Shi, N. Siegler, R. Soummer, H. Tang, J. Trauger, K. J. Wallace, X. Wang, V. White, D. Wilson, K. Yee, H. Zhou, and N. Zimmerman, "Technology development towards WFIRST-AFTA coronagraph," Proc. SPIE **9143**, 91430P (2014).
3. F. Roddier and C. Roddier, "Stellar coronograph with phase mask," Publ. Astron. Soc. Pac. **109**, 815–820 (1997).
4. D. Rouan, P. Riaud, A. Boccaletti, Y. Clénet, and A. Labeyrie, "The four-quadrant phase-mask coronagraph. I. Principle," Publ. Astron. Soc. Pac. **112**, 1479–1486 (2000).
5. G. Foo, D. M. Palacios, and G. A. Swartzlander, "Optical vortex coronagraph," Opt. Lett. **30**, 3308–3310 (2005).
6. D. Mawet, P. Riaud, O. Absil, and J. Surdej, "Annular groove phase mask coronagraph," Astrophys. J. **633**, 1191–1200 (2005).
7. G. Ruane, J. Jewell, D. Mawet, L. Pueyo, and S. Shaklan, "Apodized vortex coronagraph designs for segmented aperture telescopes," Proc. SPIE **9912**, 99122L (2016).
8. D. C. Moody and J. T. Trauger, "Hybrid Lyot coronagraph masks and wavefront control for improved spectral bandwidth and throughput," Proc. SPIE **6693**, 66931I (2007).

9. D. C. Moody, B. L. Gordon, and J. T. Trauger, "Design and demonstration of hybrid Lyot coronagraph masks for improved spectral bandwidth and throughput," Proc. SPIE **7010**, 70103P (2008).
10. C. Aime, R. Soummer, and A. Ferrari, "Total coronagraphic extinction of rectangular apertures using linear prolate apodizations," Astron. Astrophys. **389**, 334–344 (2002).
11. R. Soummer, C. Aime, and P. E. Falloon, "Stellar coronagraphy with prolate apodized circular apertures," Astron. Astrophys. **397**, 1161–1172 (2003).
12. M. N'Diaye, R. Soummer, L. Pueyo, A. Carlotti, C. C. Stark, and M. D. Perrin, "Apodized pupil Lyot coronagraphs for arbitrary apertures. V. Hybrid shaped pupil designs for imaging Earth-like planets with future space observatories," Astrophys. J. **818**, 163–171 (2016).
13. N. T. Zimmerman, M. N'Diaye, K. S. Laurent, R. Soummer, L. Pueyo, C. C. Stark, A. Sivaramakrishnan, M. Perrin, R. J. Vanderbei, N. J. Kasdin, S. Shaklan, and A. Carlotti, "Lyot coronagraph design study for large, segmented space telescope apertures," Proc. SPIE **9904**, 99041Y (2016).
14. A. J. E. Riggs, N. Zimmerman, A. Carlotti, N. J. Kasdin, and R. Vanderbei, "Shaped pupil design for future space telescopes," Proc. SPIE **9143**, 914325 (2014).
15. A. Carlotti, R. Vanderbei, and N. J. Kasdin, "Optimal pupil apodizations of arbitrary apertures for high-contrast imaging," Opt. Express **19**, 26796–26809 (2011).
16. R. J. Vanderbei, N. J. Kasdin, and D. N. Spergel, "Checkerboard-mask coronagraphs for high-contrast imaging," Astrophys. J. **615**, 555–561 (2004).
17. N. J. Kasdin, R. J. Vanderbei, D. N. Spergel, and M. G. Littman, "Extrasolar planet finding via optimal apodized-pupil and shaped-pupil coronagraphs," Astrophys. J. **582**, 1147–1161 (2003).
18. A. J. E. Riggs, N. T. Zimmerman, B. Nemati, and J. Krist, "Shaped pupil coronagraph design improvements for the WFIRST coronagraph instrument," Proc. SPIE **10400**, 104000O (2017).
19. M. Perryman, *The Exoplanet Handbook*, 2nd ed. (Cambridge University, 2018).
20. J. E. Krist, B. Nemati, and B. Mennesson, "Numerical modeling of the proposed WFIRST-AFTA coronagraphs and their predicted performances," J. Astron. Telesc. Instrum. Syst. **2**, 011003 (2016).
21. J. E. Krist, R. Belikov, L. Pueyo, D. P. Mawet, D. Moody, J. T. Trauger, and S. B. Shaklan, "Assessing the performance limits of internal coronagraphs through end-to-end modeling," Proc. SPIE **8864**, 88640P (2013).
22. J. E. Krist, "End-to-end numerical modeling of AFTA coronagraphs," Proc. SPIE **9143**, 91430V (2014).
23. J. Krist, A. J. Riggs, J. McGuire, H. Tang, N. Amiri, G. Gutt, L. Marchen, D. Marx, B. Nemati, N. Saini, E. Sidick, and H. Zhou, "WFIRST coronagraph optical modeling," Proc. SPIE **10400**, 1040004 (2017).
24. J. Krist, R. Effinger, B. Kern, M. Mandic, J. McGuire, D. Moody, P. Morrissey, I. Poberezhskiy, A. Riggs, N. Saini, E. Sidick, H. Tang, and J. Trauger, "WFIRST coronagraph flight performance modeling," Proc. SPIE **10698**, 106982K (2018).
25. B. Nemati, J. E. Krist, and B. Mennesson, "Sensitivity of the WFIRST coronagraph performance to key instrument parameters," Proc. SPIE **10400**, 1040007 (2017).
26. H. Zhou, J. Krist, and B. Nemati, "Diffraction modeling of finite subband EFC probing on dark hole contrast with WFIRST-CGI shaped pupil coronagraph," Proc. SPIE **9911**, 99111S (2016).
27. H. Zhou, B. Nemati, J. Krist, E. Cady, B. Kern, and I. Poberezhskiy, "Wavefront control performance modeling with WFIRST shaped pupil coronagraph testbed," Proc. SPIE **10400**, 1040005 (2017).
28. H. Zhou, J. Krist, E. Cady, and I. Poberezhskiy, "High accuracy coronagraph flight WFC model for WFIRST-CGI raw contrast sensitivity analysis," Proc. SPIE **10698**, 106982M (2018).
29. C. Aime, R. Soummer, and A. Ferrari, "Interferometric apodization of rectangular apertures—application to stellar coronagraphy," Astron. Astrophys. **379**, 697–707 (2001).
30. C. Aime, "Principle of an achromatic prolate apodized Lyot coronagraph," Publ. Astron. Soc. Pac. **117**, 1012–1019 (2005).
31. K. Balasubramanian, V. White, K. Yee, P. Echternach, R. Muller, M. Dickie, E. Cady, C. M. Prada, D. Ryan, I. Poberezhskiy, B. Kern, H. Zhou, J. Krist, B. Nemati, A. J. E. Riggs, N. T. Zimmerman, and N. J. Kasdin, "WFIRST-AFTA coronagraph shaped pupil masks: design, fabrication, and characterization," J. Astron. Telesc. Instrum. Syst. **2**, 011005 (2016).
32. K. Balasubramanian, A. J. E. Riggs, E. Cady, V. White, K. Yee, D. Wilson, P. Echternach, R. Muller, C. M. Prada, B.-J. Seo, F. Shi, D. Ryan, S. Fregoso, J. Metzman, and R. C. Wilson, "Systematic errors and defects in fabricated coronagraph masks and laboratory scale star-shade masks and their performance impact," Proc. SPIE **10400**, 104000C (2017).
33. N. T. Zimmerman, A. J. E. Riggs, N. J. Kasdin, A. Carlotti, and R. J. Vanderbei, "Shaped pupil Lyot coronagraphs: high-contrast solutions for restricted focal planes," J. Astron. Telesc. Instrum. Syst. **2**, 011012 (2016).
34. A. S. Jurling, M. D. Bergkoetter, and J. R. Fienup, "Techniques for arbitrary sampling in two-dimensional Fourier transforms," J. Opt. Soc. Am. **A35**, 1784–1796 (2018).
35. R. Soummer, L. Pueyo, A. Sivaramakrishnan, and R. J. Vanderbei, "Fast computation of Lyot-style coronagraph propagation," Opt. Express **15**, 15935–15951 (2007).
36. J. E. Krist, "PROPER: an optical propagation library for IDL," Proc. SPIE **6675**, 66750P (2007).
37. M. D. Perrin, R. Soummer, E. M. Elliot, M. D. Lallo, and A. Sivaramakrishnan, "Simulating point spread functions for the James Webb Space Telescope with WebbPSF," Proc. SPIE **84423**, 84423D (2012).
38. A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, 3rd ed. (Prentice Hall, 2009).
39. R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial grayscale," Proc. Soc. Inf. Disp. **17**, 75–77 (1976).
40. M. R. Bolcar, J. Crooke, J. E. Hylan, G. Bronke, C. Collins, J. Corsetti, J. Generie, Q. Gong, T. Groff, W. Hayden, A. Jones, B. Matonak, S. Park, L. Sacks, G. West, K. Yang, and N. Zimmerman, "The large UV/Optical/Infrared Surveyor (LUVOIR): decadal mission study update," Proc. SPIE **10698**, 106980O (2018).
41. T. D. Groff, A. J. E. Riggs, B. Kern, and N. J. Kasdin, "Methods and limitations of focal plane sensing, estimation, and control in high-contrast imaging," J. Astron. Telesc. Instrum. Syst. **2**, 011009 (2016).
42. Q. Gong, M. R. Bolcar, J. A. Corsetti, J. A. Crooke, J. Generie, T. D. Groff, J. Hylan, A. Jones, R. J. Parramon, A. M. Mandell, G. J. West, and N. T. Zimmerman, "Optical design of the Extreme Coronagraph for Living Planetary Systems instrument for the LUVOIR mission study," J. Astron. Telesc. Instrum. Syst. **5**, 025002 (2019).
43. L. Pueyo, J. Kay, N. J. Kasdin, T. Groff, M. McElwain, A. Give'on, and R. Belikov, "Optimal dark hole generation via two deformable mirrors with stroke minimization," Appl. Opt. **48**, 6296–6312 (2009).