

Fast computation of Lyot-style coronagraph propagation

R. Soummer^{*1,2}, L. Pueyo³,
A. Sivaramakrishnan^{1,2,4}, and R. J. Vanderbei⁵

¹Department of Astrophysics, American Museum of Natural History, New York, NY 10024

²Center for Adaptive Optics, University of California, Santa Cruz, CA 95064

³Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544

⁴Department of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794

⁵Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544

corresponding author:

rsoummer@amnh.org

Abstract: We present a new method for numerical propagations through Lyot-style coronagraphs using finite size occulting masks. Standard methods for coronagraphic simulations involve Fast Fourier Transforms (FFT) of very large arrays, and computing power is an issue for the design and tolerancing of coronagraphs on Extremely Large Telescopes (ELT) in order to handle both the speed and memory requirements. Our method combines a semi-analytical approach with non-FFT based Fourier transform algorithms. It enables both fast and memory-efficient computations without introducing any additional approximations. Typical speed improvements based on computation costs are of twenty to fifty for propagations from pupil to Lyot plane, with thirty to sixty times less memory needed. Our method makes it possible to perform numerical coronagraphic studies even in the case of ELTs using a contemporary commercial laptop computer, or any standard commercial workstation computer.

© 2007 Optical Society of America

OCIS codes: (350.1260) Astronomical optics; (110.1080) Active or adaptive optics; (070.2025) Discrete optical signal processing; (070.2575) Fractional Fourier transforms; (070.7145) Ultrafast processing

References and links

1. B. Macintosh, J. Graham, D. Palmer, R. Doyon, D. Gavel, J. Larkin, B. Oppenheimer, L. Saddlemyer, J. K. Wallace, B. Bauman, J. Evans, D. Erikson, K. Morzinski, D. Phillion, L. Poyneer, A. Sivaramakrishnan, R. Soummer, S. Thibault, and J.-P. Veran, "The Gemini Planet Imager," in *Advances in Adaptive Optics II. Edited by Ellerbroek, Brent L.; Bonaccini Calia, Domenico. Proceedings of the SPIE, Volume 6272, pp. (2006).* (2006).
2. J. L. Beuzit, D. Mouillet, C. Moutou, K. Dohlen, T. Fusco, P. Puget, S. Udry, R. Gratton, H. M. Schmid, M. Feldt, M. Kasper, and The Vlt-Pf Consortium, "A "Planet Finder" instrument for the VLT," in *Tenth Anniversary of 51 Peg-b: Status of and prospects for hot Jupiter studies*, L. Arnold, F. Bouchy, and C. Moutou, eds., pp. 353–355 (2006).
3. M. Tamura, K. Hodapp, H. Takami, L. Abe, H. Suto, O. Guyon, S. Jacobson, R. Kandori, J.-I. Morino, N. Murakami, V. Stahlberger, R. Suzuki, A. Tavrov, H. Yamada, J. Nishikawa, N. Ukita, J. Hashimoto, H. Izumiura, M. Hayashi, T. Nakajima, and T. Nishimura, "Concept and science of HiCIAO: high contrast instrument for the Subaru next generation adaptive optics," in *Ground-based and Airborne Instrumentation for Astronomy. Edited by McLean, Ian S.; Iye, Masanori. Proceedings of the SPIE, Volume 6269, pp. 62690V (2006).*, vol. 6269 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference* (2006).

4. B. Macintosh, M. Troy, R. Doyon, J. Graham, K. Baker, B. Bauman, C. Marois, D. Palmer, D. Phillion, L. Poyneer, I. Crossfield, P. Dumont, B. M. Levine, M. Shao, G. Serabyn, C. Shelton, G. Vasisht, J. K. Wallace, J.-F. Lavigne, P. Valee, N. Rowlands, K. Tam, and D. Hackett, "Extreme adaptive optics for the Thirty Meter Telescope," in *Advances in Adaptive Optics II. Edited by Ellerbroek, Brent L.; Bonaccini Calia, Domenico. Proceedings of the SPIE, Volume 6272, pp. 62720N (2006).*, vol. 6272 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference (2006)*.
5. L. Close, "Extrasolar Planet Imaging with the Giant Magellan Telescope," in *Proceedings of the conference In the Spirit of Bernard Lyot: The Direct Detection of Planets and Circumstellar Disks in the 21st Century. June 04 - 08, 2007. University of California, Berkeley, CA, USA. Edited by Paul Kalas., P. Kalas, ed. (2007)*.
6. C. Verinaud, M. Kasper, J.-L. Beuzit, N. Yaitskova, V. Korkiakoski, K. Dohlen, P. Baudoz, T. Fusco, L. Mugnier, and N. Thatte, "EPICS Performance Evaluation through Analytical and Numerical Modeling," in *Proceedings of the conference In the Spirit of Bernard Lyot: The Direct Detection of Planets and Circumstellar Disks in the 21st Century. June 04 - 08, 2007. University of California, Berkeley, CA, USA. Edited by Paul Kalas., P. Kalas, ed. (2007)*.
7. O. Guyon, J. R. P. Angel, C. Bowers, J. Burge, A. Burrows, J. Codona, T. Greene, M. Iye, J. Kasting, H. Martin, D. W. McCarthy, Jr., V. Meadows, M. Meyer, E. A. Pluzhnik, N. Sleep, T. Spears, M. Tamura, D. Tenerelli, R. Vanderbei, B. Woodgate, R. A. Woodruff, and N. J. Woolf, "Telescope to observe planetary systems (TOPS): a high throughput 1.2-m visible telescope with a small inner working angle," in *Space Telescopes and Instrumentation I: Optical, Infrared, and Millimeter. Edited by Mather, John C.; MacEwen, Howard A.; de Graauw, Matheus W. M.. Proceedings of the SPIE, Volume 6265, pp. 62651R (2006).*, vol. 6265 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference (2006)*.
8. J. T. Trauger and W. A. Traub, "A laboratory demonstration of the capability to image an Earth-like extrasolar planet," *Nature* **446**, 771–773 (2007).
9. C. Aime and R. Soummer, "The Usefulness and Limits of Coronagraphy in the Presence of Pinned Speckles," *Astrophys. J.* **612**, L85–L88 (2004).
10. C. Cavarroc, A. Boccaletti, P. Baudoz, T. Fusco, and D. Rouan, "Fundamental limitations on Earth-like planet detection with extremely large telescopes," *A&A* **447**, 397–403 (2006).
11. O. Guyon, E. A. Pluzhnik, M. J. Kuchner, B. Collins, and S. T. Ridgway, "Theoretical Limits on Extrasolar Terrestrial Planet Detection with Coronagraphs," *Astrophys. J.* **167**, 81–99 (2006). [arXiv:astro-ph/0608506](https://arxiv.org/abs/astro-ph/0608506).
12. R. Soummer, A. Ferrari, C. Aime, and L. Jolissaint, "Speckle noise and dynamic range in coronagraphic images," *ApJ*, accepted (2007).
13. B. Lyot, "Étude de la couronne solaire en dehors des éclipses. Avec 16 figures dans le texte." *Zeitschrift fur Astrophysics* **5**, 73–+ (1932).
14. C. Aime, R. Soummer, and A. Ferrari, "Total coronagraphic extinction of rectangular apertures using linear prolate apodizations," *A&A* **389**, 334–344 (2002).
15. R. Soummer, C. Aime, and P. E. Falloon, "Stellar coronagraphy with prolate apodized circular apertures," *A&A* **397**, 1161–1172 (2003).
16. R. Soummer, "Apodized Pupil Lyot Coronagraphs for Arbitrary Telescope Apertures," *Astrophys. J.* **618**, L161–L164 (2005).
17. R. Soummer, L. Pueyo, A. Ferrari, C. Aime, A. Sivaramkrishnan, and N. Yaitskova, "Apodized Pupil Lyot Coronagraphs for Arbitrary Telescope Apertures. II. Application to Extremely Large Telescopes," submitted to *ApJ* (2007).
18. F. Roddier and C. Roddier, "Stellar Coronagraph with Phase Mask," *PASP* **109**, 815–820 (1997).
19. R. Soummer, K. Dohlen, and C. Aime, "Achromatic dual-zone phase mask stellar coronagraph," *A&A* **403**, 369–381 (2003).
20. D. Rouan, P. Riaud, A. Boccaletti, Y. Clénet, and A. Labeyrie, "The Four-Quadrant Phase-Mask Coronagraph. I. Principle," *PASP* **112**, 1479–1486 (2000).
21. L. Abe, F. Vakili, and A. Boccaletti, "The achromatic phase knife coronagraph," *A&A* **374**, 1161–1168 (2001).
22. M. J. Kuchner and W. A. Traub, "A Coronagraph with a Band-limited Mask for Finding Terrestrial Planets," *ApJ* **570**, 900–908 (2002).
23. G. Foo, D. M. Palacios, and G. A. Swartzlander, Jr., "Optical vortex coronagraph," *Optics Letters* **30**, 3308–3310 (2005).
24. J. Goodman, *Introduction to Fourier Optics* (Mac Graw Hill, 1996).
25. R. N. Bracewell, *The Fourier Transform & Its Applications*, McGraw-Hill series in electrical and computer engineering., 3rd ed. (McGraw-Hill Science/Engineering/Math, Boston, 1999).
26. A. Ferrari, R. Soummer, and C. Aime, "Introduction to stellar coronagraphy," *C. R. Acad. Sci Paris* **in press** (2007).
27. D. H. Bailey and P. N. Swartztrauber, "The Fractional Fourier Transform and Applications," *SIAM Review* **33**(3), 389–404 (1991).
28. L. R. Rabiner, R. W. Schafer, and C. M. Rader, "The Chirp z-Transform Algorithm and its Application," *Bell Sys. Tech. J.* **48**(5), 1249–1291 (1969).
29. J. O. Smith, *Mathematics of the Discrete Fourier Transform (DFT)* (W3K Publishing,

- <http://www.w3k.org/books/>, 2007).
30. C. Aime, "Principle of an Achromatic Prolate Apodized Lyot Coronagraph," *PASP***117**, 1012–+ (2005).
 31. URL <http://www.fft.w.org/speed/>.
 32. J. P. Lloyd and A. Sivaramakrishnan, "Tip-Tilt Error in Lyot Coronagraphs," *Astrophys. J.* **621**, 1153–1158 (2005). [arXiv:astro-ph/0503661](https://arxiv.org/abs/astro-ph/0503661).
 33. A. Sivaramakrishnan and J. P. Lloyd, "Spiders in Lyot Coronagraphs," *Astrophys. J.* **633**, 528–533 (2005). [arXiv:astro-ph/0506564](https://arxiv.org/abs/astro-ph/0506564).
 34. A. Sivaramakrishnan, C. D. Koresko, R. B. Makidon, T. Berkefeld, and M. J. Kuchner, "Ground-based Coronagraphy with High-order Adaptive Optics," *Astrophys. J.* **552**, 397–408 (2001). [astro-ph/0103012](https://arxiv.org/abs/astro-ph/0103012).
 35. L. Bluestein, "A linear filtering approach to the computation of discrete Fourier transform," *IEEE Transactions on Audio and Electroacoustics* **18**, 451–455 (1970).
 36. URL <http://www.primatelabs.ca/geekbench/>.
-

1. Introduction

The field of high contrast imaging has expanded rapidly in the past few years, driven by the prospect of science enabled by the direct detection and characterization of extrasolar planets and faint circumstellar disks. Several observatories have launched studies and development of such projects, e.g. Gemini Planet Imager (GPI) [1], SPHERE [2], or HiCIAO [3]. In the future, Extremely Large Telescopes (ELT) [4, 5, 6] will provide the higher angular resolution necessary to study more distant, planet forming regions, than today's eight meter class telescopes, and also faint old giant planets. The study of Earth-like planets will probably have to wait for space-based instruments [7, 8]. Ground-based imaging requires the association of coronagraphy and extreme adaptive optics (ExAO), which can be studied both theoretically and numerically [9, 10, 11, 12].

In this paper we study the numerical modeling of Lyot type coronagraphs [13], including Apodized Pupil Lyot Coronagraphs (APLC) [14, 15, 16, 17], or phase masks [18, 19]. These coronagraphs consists of the succession of binary filters (pupil, occulting mask, Lyot stop). The method we describe does not provide any significant improvement in the case of infinite-size focal plane masks [20, 21, 22, 23], and we will not discuss these coronagraphs here. Under the classical approximations of Fourier Optics [24], a Fourier Transform (FT) relationship exists between the field amplitude at two successive planes of the coronagraph. The pupil, with its finite support, is not band-limited and so cannot be perfectly sampled according to the Shannon-Nyquist sampling theorem [25]. The same problem pertains to the occulting mask in the focal plane. A remedy is to impose a very high sampling both in the pupil and in the focal plane, in order to represent small features (e.g. segmentation or secondary mirror support structures, occulting spot). Note that Nyquist sampling of the focal plane with two pixels per resolution element (λ/D), is not sufficient here. Indeed, the small occulting mask has a typical size of about $5\lambda/D$, and would not be well represented by a disk of ten pixels in diameter. This is even more problematic for phase masks of size $\sim \lambda/D$ [18, 19, 26]. A considerable oversampling of the PSF is therefore required for credible coronagraphic modeling.

This two-fold sampling requirement in two Fourier-conjugate planes is not only a severe practical problem, but raises a fundamental problem associated to the uncertainty principle. In order to preserve complete information in the FT, a high sampling in one domain means that the corresponding FT must be calculated to a very high frequency in the reciprocal domain. This is achieved optically when forming successively pupil and focal plane images. Numerically, this can be reproduced using Fast Fourier Transforms (FFT), but this can lead to overwhelming demands on memory and computing power. In Sec.2 we show that this fundamental limitation makes classical FFTs poorly suited for Lyot-type coronagraphic computations. In Sec.4 we analyze the analytical formalism of coronagraphic propagation, and conclude that a partial FT algorithm providing arbitrary sampling in a limited area is more appropriate for coronagraphy. Several methods can be used for that purpose [27, 28], and we detail a matrix-based FT, [29]

that can be implemented easily in any high-level language, with the advantage of being easily (or automatically) parallelized on multi-cpu computers.

Comparisons between classical propagation with FFTs and our semi-analytical method show a considerable gain both in speed and memory requirements. This therefore opens new possibilities for coronagraphic studies hitherto impossible without large computer cluster facilities. Among them, the study of extremely fine sampling of the pupil plane in the case of ELTs, of extremely fine sampling of the occulting plane mask, or of end-to-end simulations of an ExAO coronagraph to simulate a long exposure, are achievable with standard desktop equipment.

2. Lyot-type Coronagraphs

In this section we recall the general formalism of Lyot-type coronagraphs, following the notation of Aime *et al.* and Soummer *et al.* [14, 15, 17]. The general layout is given in Fig.1. The setup consists of an ensemble of apodizers, masks and stops in four successive planes A, B, C, D , respectively, where A is the entrance aperture, B is the focal plane with the occulting mask, C is an image of the entrance aperture where a pupil mask called Lyot stop is placed and D is the final image plane. We will consider the usual approximations of paraxial optics [24], and that the optical layout is properly designed to cancel the quadratic phase terms associated with Fresnel propagation, so that a FT relationship exists between two successive planes. The tele-




Fig. 1. Illustration of the four coronagraphic planes: the pupil corresponds to Plane A (possibly apodized). A focal masks (hard-edged, or phase mask) is placed in the focal plane B, and a Lyot stop (possibly undersized) in plane C.

scope aperture function with the position vector $\mathbf{r} = (x, y)$ is denoted by $P(\mathbf{r})$ (index function equal to 1 inside the aperture \mathcal{P}). This aperture can be apodized by a function $\Phi(\mathbf{r})$. Note that these functions do not have to be radial. A mask of transmission $1 - \varepsilon M(\mathbf{r})$ is placed in the focal plane. M is the index function that describes the mask shape \mathcal{M} , equal to 1 inside the coronagraphic mask and 0 outside. L is the index function of the Lyot stop. We recall that both the Lyot coronagraph (opaque mask) and the Roddier coronagraph (π phase mask) can be described by this common formalism, with $\varepsilon = 1$ for Lyot and $\varepsilon = 2$ for Roddier using that $e^{i\pi} = -1$, as detailed in [14, 26]. The field amplitude in the four successive planes are:

$$\Psi_A(\mathbf{r}) = P(\mathbf{r})\Phi(\mathbf{r}) \quad (1)$$

$$\Psi_B(\mathbf{r}) = \widehat{\Psi}_A(\mathbf{r})(1 - \varepsilon M(\mathbf{r})) \quad (2)$$

$$\Psi_C(\mathbf{r}) = (\Psi_A(\mathbf{r}) - \varepsilon \Psi_A(\mathbf{r}) * \widehat{M}(\mathbf{r})) L(\mathbf{r}) \quad (3)$$

$$\Psi_D(\mathbf{r}) = (\widehat{\Psi}_A(\mathbf{r}) - \varepsilon \widehat{\Psi}_A(\mathbf{r}) M(\mathbf{r})) * \widehat{L}(\mathbf{r}) \quad (4)$$

In these equations, $\widehat{\cdot}$ denotes the Fourier Transform, $*$ the convolution product, and we assume that the focal lengths of the successive optical systems are identical (if not, an appropriate change of variables leads to the same result). Also, we re-orientate the axis in the opposite

direction and omit the -1 proportionality factor for better legibility here. Also, we assumed a monochromatic propagation in these equations. The effect of the wavelength can be added easily to this formalism: in the focal plane, the size of the point spread function is proportional to the wavelength, but the mask (hard-edged or phase) has a fixed size. It can be readily shown formally [14, 15, 30] that changing the wavelength is formally equivalent to changing the mask size.

3. Classical numerical propagation and FFTs

In order to compute the coronagraphic propagation, we are interested in the numerical evaluation of the Fourier integral between each plane. Without loss of generality we consider the case of a one-dimensional signal $f(x)$ with $x \in [-\gamma D/2, \gamma D/2]$ and $f(x) = 0$ for $|x| > D/2$, where γ is a padding coefficient that will be later related to the resolution of the FT in the image plane. We illustrate the relationship between plane A and plane B. We first recall classical results of Fourier analysis [27]: the sampled Continuous Fourier Transform (CFT) can be written as a Riemann sum:

$$\widehat{F}(u_k) = \int_{-\gamma D/2}^{\gamma D/2} f(x) e^{-i2\pi x u_k} dx \quad (5)$$

$$\simeq \delta x \sum_{n=0}^{N_A-1} f(x_n) e^{-i2\pi x_n u_k} \quad (6)$$

where $x_n = (n - \gamma N_A/2) \delta x$, $n \in [0, \gamma N_A - 1]$, corresponding to the sampling points of $f(x)$, and $u_k = (k - N_B/2) \delta u$, $k \in [0, N_B - 1]$ are the sampled values of the Fourier transform. N_A is the number of pixels along the pupil diameter D , and N_B is the number of pixels in the focal plane. Note that under the Riemann sum approximation, independent sampling grids can be chosen in both domains. However, when one chooses the same size $N = \gamma N_A = N_B$ for both arrays, and the integration steps as:

$$\delta x \delta u = \frac{1}{N}, \quad (7)$$

then Eq. 6 greatly simplifies to:

$$\widehat{F}(u_k) = \frac{\gamma D}{N} (-1)^{N/2-k} \sum_{n=0}^{N-1} (-1)^n f(x_n) e^{-i2\pi kn/N}, \quad (8)$$

where the sum is now a Discrete Fourier Transform (DFT), which can be computed very efficiently using FFT algorithms. The number of floating point operations (flops) in a radix-2 Cooley-Tukey FFT of size N is $5N \log_2 N$. It is generally assumed that this is also an approximation for the number of operations in any complex FFT (see for example the documentation of the FFTW package [31]).

It is critical to realize that this possibility of using FFTs is obtained at the expense of a fixed relationship between the focal plane sampling and the padding factor, forced by Eq. 7: $\delta u = \lambda / (\gamma D)$, as illustrated on Fig. 2. Under this condition, the occulting mask of m resolution elements is therefore sampled with γm pixels. Because stellar coronagraphs use very small masks (typically $4 \sim 5\lambda/D$ for an APLC, and λ/D for a PM or DZPM), decent sampling of these masks (at the very least a few tens of pixels) imposes large zero padding factors γ . A general consensus for coronagraphic calculation is to use $\gamma = 6$ or $\gamma = 8$. The memory and computational speed problems associated with classical coronagraphic algorithms stem from the requirement of finely sampling the image plane mask. This can lead to prohibitively large padding factors for some applications which can benefit for particularly high sampling, such as

tip-tilt tolerancing [32], study of atmospheric differential refraction effects, mask ellipticity or roughness studies.

A common method to mitigate these mask sampling problems is to use a gray-pixel approximation, where gray pixels are used at the edge of the mask: their value is defined as the ratio of the area covered by the mask to the total pixel area. This technique corresponds to a slight numerical apodization of the focal plane mask, akin to the use of an anti-aliasing filter. The same approach can be used for small features in the pupil plane, such as secondary mirror support structures. However, there are cases where the gray approximation cannot be used. For example, the tolerancing of the effects of the small pupil features (segmentation, spiders [33]), and their mitigation by the Lyot Stop cannot be done with gray pixel approximation. Prolate apodizers for APLCs need to be defined without gray pixels in the pupil [17]. In the case of phase masks (Roddier [18] or Dual Zone [19]), the gray approximation is meaningless for the phase mask, and large padding coefficients have to be used since these masks are very small (typically one resolution element).

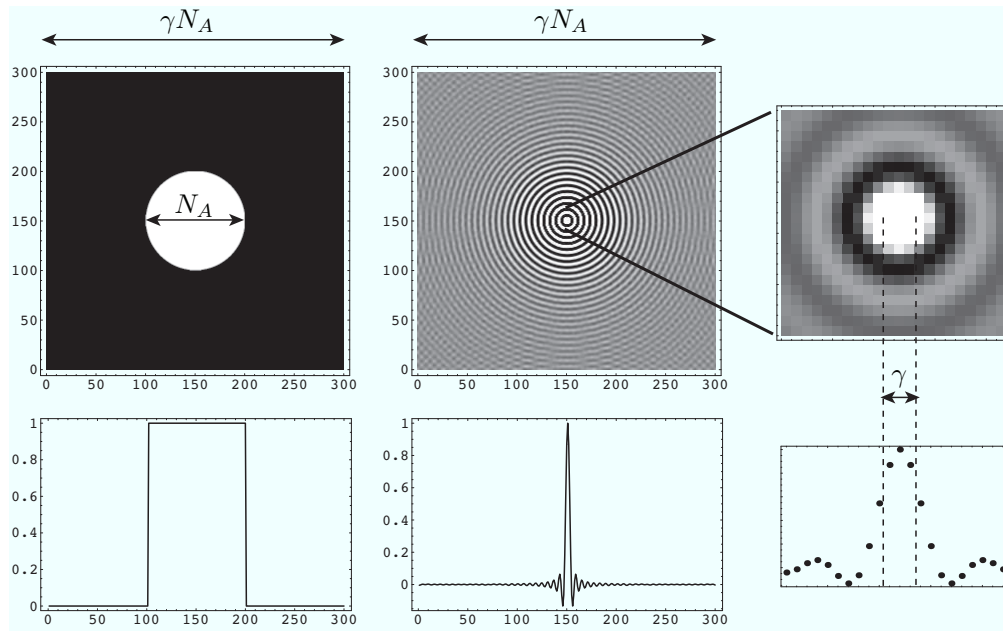


Fig. 2. Sampling relationship between pupil (left column) and image plane (center), using FFTs. The padded pupil plane array is γ times larger than the actual pupil (here $\gamma = 3$), and its Fourier transform features three pixels per unit of angular resolution, as shown in the zoomed image of the core (right). In this case a $4 \sim 5$ resolution element mask would not be sufficiently sampled. A general consensus for coronagraphic calculations is to use $\gamma = 6$ or $\gamma = 8$.

4. Semi-analytical coronagraphic propagations

4.1. Principle

As shown in Fig. 2, FFT methods address the image plane sampling requirement by zero padding the pupil, and mimicking the optical propagation, preserving the complete information between planes. A much better approach can be derived from the analytical expression of

the field in the Lyot plane, simply re-writing Eq. 3 as:

$$\Psi_C(\mathbf{r}) = (\Psi_A(\mathbf{r}) - \varepsilon \mathcal{F}[\mathcal{F}[\Psi_A(\mathbf{r})]M(\mathbf{r})]) L(\mathbf{r}), \quad (9)$$

where we can readily identify that the first FT of the pupil field amplitude $\mathcal{F}[\Psi_A(\mathbf{r})]$ is truncated by the occulting spot $M(\mathbf{r})$, and that the second FT is truncated by the Lyot Stop $L(\mathbf{r})$. This means that we are only interested in the knowledge of the FTs *inside* limited areas, viz., the limited occulting mask area, and the limited Lyot stop area. We can thus completely circumvent the sampling problem by restricting the information of the FTs to these two areas: the semi-analytical approach consists of computing these limited-area FTs numerically, and subtracting the result from the pupil complex amplitude, according Eq. 9.

In the case of one-dimensional problems (rectangular apertures [14], or perfect circular apertures [15, 19]), the semi-analytical approach can be applied by calculating directly these two limited-area FTs (or Hankel Transforms for circular apertures) using a one-dimensional numerical integration algorithm. This approach presents some advantages for phase masks calculations [19], but such direct integrations cannot be used efficiently in the two-dimensional case for computation reasons.

In the general two-dimensional case, the two limited-area FTs can be calculated using partial FT methods. Several methods exist to calculate such partial FTs, such as the Fractional Fourier Transform or chirp z-transform, [27, 28], and we describe a very simple matrix-based method in the next section [29]. These partial FT methods are intrinsically slower than a FFT if exactly the same computation is performed. However, since we are only interested in a small number of points within limited areas (occulting mask *and* Lyot stop), the effect is to replace fast calculations with very large arrays by slow calculations with very small arrays. Note that for $\gamma = 8$, only 40×40 values of the FT need to be calculated in the first focal plane for a $5\lambda/D$ mask, independently of the size of the pupil n pixels.

The semi-analytical propagation from the pupil to the Lyot plane is performed as follows:

- Define the pupil field amplitude *without* zero padding as a $N_A \times N_A$ array, where N_A is the diameter of the pupil in pixels.
- Calculate the FT to an arbitrary high sampling inside the area *limited* to the occulting mask as a $N_B \times N_B$ array, with e.g. $N_B = 40$ for a $5\lambda/D$ mask with $\gamma = 8$.
- Calculate the FT of this previous field amplitude inside an area *limited* to the pupil ($N_A \times N_A$ array).
- Reverse the spatial axes because two successive FTs restore the function while changing the sign of the variable, and subtract the result from the pupil field amplitude (An inverse FFT can also be used between plane B and plane C to avoid reversing the axis).

Note that for convenience all masks (pupil and occulor) are defined with their centers located on the pixel $N_A/2 + 1$ (or $N_B/2 + 1$), and that the focal plane mask can be defined using the gray pixel approximation, as for the classical FFT method. These computation steps are illustrated in Fig.3 and compared to the classical FFT approach. The figure shows the actual arrays that are calculated for both method: note that the semi-analytical method does not use zero-padding in the pupil plane and calculates the focal plane amplitude only *inside* the occulting mask, instead of *outside* in the case of the classical FFT method.

In place C, the Lyot stop can be equal to the pupil size in the case of an APLC, or more generally undersized to optimize the coronagraphic efficiency according to various possible criteria [34, 17]. If the Lyot stop size has been previously chosen, the calculated area can be limited to the Lyot stop itself to accelerate further the computation, but in practice we usually calculate

the Lyot plane amplitude over the size of the entire pupil. If necessary, it is straightforward to oversize the calculated region, at the expense of computing efficiency, for example to see the light diffracted outside the pupil.

Avoiding the typical six-to-eight fold zero-padding enables the use of much smaller arrays, than FFT calculations require. For example, in the case of GPI ($D = 8m$), the secondary mirror support structures are about 1cm wide. Understanding the effects of spiders on coronagraphic performance is important since they appear bright in the Lyot plane and must be masked out by the Lyot stop [33]. For example, with 4 pixels per spider, $N_A = 3200$ pixels are needed in the pupil. Standard $\gamma = 8$ padding would require 25600×25600 FFTs. The semi analytical method enables this calculation in a few seconds on a standard desktop computer (Table 2).

It is important to note that the calculations of the partial FTs inside the mask and the Lyot stop does not correspond to an additional approximation, as it may seem that we discard the information *outside* the focal plane mask. We do not lose any information because the method is based on the analytical formulation of Eq. 3 (or Eq. 9). In particular, it can be easily verified that both method, give the *exact* same results to the numerical precision, when the same sampling is used.

4.2. Matrix direct Fourier transform

In this section we describe a simple matrix Fourier transform (MFT), which can be used for the semi-analytical method. In order to restrict the computation of the FT of the pupil to the image plane size m expressed in resolution elements units (λ/D), we choose the sampling step in plane B such that: $du = m/N_B$. We compute the Riemann sum directly using a matrix formulation of Eq.6:

$$\begin{pmatrix} \widehat{F}(u_0) \\ \dots \\ \widehat{F}(u_k) \\ \dots \\ \widehat{F}(u_{N_B-1}) \end{pmatrix} = \begin{pmatrix} e^{-2i\pi x_0 u_0} & \dots & e^{-2i\pi x_k u_0} & \dots & e^{-2i\pi x_{N_A-1} u_0} \\ \dots & \dots & \dots & \dots & \dots \\ e^{-2i\pi x_0 u_k} & \dots & e^{-2i\pi x_k u_k} & \dots & e^{-2i\pi x_{N_A-1} u_k} \\ \dots & \dots & \dots & \dots & \dots \\ e^{-2i\pi x_0 u_{N_B-1}} & \dots & e^{-2i\pi x_k u_{N_B-1}} & \dots & e^{-2i\pi x_{N_A-1} u_{N_B-1}} \end{pmatrix} \begin{pmatrix} f(x_0) \\ \dots \\ f(x_k) \\ \dots \\ f(x_{N_A-1}) \end{pmatrix}, \quad (10)$$

which can be rewritten as:

$$\widehat{F}(\mathbf{U}) = e^{-2i\pi \mathbf{U} \mathbf{X}^T} \cdot f(\mathbf{X}), \quad (11)$$

where $\mathbf{U} = (u_0, \dots, u_{N_B-1})^T$, $\mathbf{X} = (x_0, \dots, x_{N_A-1})^T$, and $\exp(\cdot)$ is the element-wise exponential of a matrix. Two-dimensional FTs can be implemented straightforwardly as follows:

- Define the four vectors $\mathbf{U} = (u_0, \dots, u_{N_B-1})^T$, $\mathbf{X} = (x_0, \dots, x_{N_A-1})^T$, $\mathbf{V} = (v_0, \dots, v_{N_B-1})^T$, $\mathbf{Y} = (y_0, \dots, y_{N_A-1})^T$.
- The vector elements are $x_k = y_k = (k - N_A/2) \times 1/N_A$ and $u_l = v_l = (l - N_B/2) \times m/N_A$, for $k = [0, \dots, N_A - 1]$ and $l = [0, \dots, N_B - 1]$.
- The two-dimension FT is obtained by computing the two matrix products:

$$\widehat{F}(\mathbf{U}, \mathbf{H}) = \frac{m}{N_A N_B} e^{-2i\pi \mathbf{U} \mathbf{X}^T} \cdot f(\mathbf{X}, \mathbf{Y}) \cdot e^{-2i\pi \mathbf{Y} \mathbf{V}^T}, \quad (12)$$

where the normalization coefficient $m/(N_A N_B)$ imposes the conservation of energy according to the Parseval theorem [25]: the energy in the limited-area FT is a fraction of the total energy of the FT, corresponding to the limited area, which was calculated.

- With this definition, the Fourier transform is centered in a similar fashion to the FFT, with the zero frequency at the pixel $N_B/2 + 1$.




Fig. 3. Illustration of each computing steps for the classical method using FFTs (left) and the semi-analytical method (SAM) for a possible ELT geometry and an APLC. With the classical method, a six-fold zero padding is used in the pupil (6 pixel per resolution element in the focal plane). The insert shows the application of the five resolution element opaque mask. These computation involves two 4500×4500 FFTs. In the case of the semi-analytical method we do not use zero padding of the pupil. In the focal plane, we only need to calculate the points of the FT that are *inside* of the occulting mask (note the difference with the FFT method, where the points *outside* the mask are needed). The Lyot plane is obtained in the area limited to the pupil geometry by analytically forming the interference between the pupil and the wave diffracted by the mask. In this example, we obtain a speed improvement of about 15 using 36 times less memory.

Note that when the conditions of Eq.7 are verified, both the FFT and the matrix FT give identical results. This should be used as a sanity check when implementing the technique. This is not surprising since the matrix method is a direct implementation of the Riemann sum, which is also what is computed by a FFT very efficiently taking advantage of simplifications for particular sampling conditions (Eq.8). With the MFT, the arbitrary sampling du in the focal plane is obtained at the expense of partial information on the FT. Note that in the case of the MFT, γ still corresponds to the sampling of the FT (number of pixels per resolution element), but not to the zero-padding coefficient. The use of the semi-analytical method enables to break the sampling constraints of the FFT, and to provide increased flexibility in the computation, since the number of pixel in the pupil and in the focal plane are totally independent in this case.

In terms of computation cost, the matrix FT involves 2 complex-matrix products of the form:

$$F = E_1 \cdot f \cdot E_2, \quad (13)$$

where the respective sizes of the three matrix E_1, f, E_2 are: $N_B \times N_A, N_A \times N_A$, and $N_B \times N_A$. We consider the case of a complex FT for generality. Each element of the result in a matrix product is the result of N_A multiplications and $N_A - 1$ additions. While complex addition requires 2 flops (floating point operations), complex multiplications require 6 flops. Since the two matrices E_1 and E_2 can be calculated beforehand, in a similar fashion that FFT plans are generated, the total number of operations for the product $f \cdot E_2$ is therefore: $N_A N_B (8N_A - 2) \simeq 8N_A^2 N_B$, assuming N_A large enough. The number of flops involved in the first complex MFT (pupil to focal) is therefore:

$$n(MFT) = 8(N_A^2 N_B + N_A N_B^2) - 2N_A N_B - 2N_B^2. \quad (14)$$

In the particular case where both pupil and image plane have the same array size N , it is interesting to note that the number of operations for the matrix method is proportional to N^3 and not N^4 as for a direct calculation of the DFT. This is because the two successive matrix products take advantage of some redundancy in intermediate calculations. The FFT is much more efficient in this case, and the relative number of operations is:

$$\frac{n(FFT)}{n(MFT)} = \frac{5 \log_2(N)}{8N}. \quad (15)$$

For example a 2000×2000 FT requires 36 times more operations with a MFT than with a FFT. This calculation, which takes about 0.5s with a FFT on a on a 2GHz Apple G5 (1 Gflop/s), would take 18s with a MFT. The performance comparison between the FFT propagation method and the semi-analytical method is given in Sec. 5.

4.3. Fast fractional Fourier transform

The computation of the FT in a limited area can also be made using a Fractional Fourier Transform. This method is thoroughly presented by Bailey and Swartztrauber [27] and we refer to these authors for a more detailed description of this algorithm. Consider the Riemann sum in Eq. 6, and assume that there is no zero padding $x_n = n\delta x - (N_A \delta x)/2$, $n \in [0, N_A - 1]$ with $\delta x = D/(N_A)$ and $u_k = k\delta u - (N_B \delta u)/2$, $k \in [0, N_B - 1]$ with $\delta u = 1/(\gamma D)$ and $N_B = m\gamma$, where m is the size of the image plane mask in resolution element units (λ/D). Note that here γ does not correspond to an oversize of the pupil but directly to the number of pixels per unit of angular resolution in the image plane. The Riemann sum can be re-written as:

$$\hat{F}(u_k) = \frac{D}{N} e^{i(\pi/\gamma)(k-N_B/2)} \sum_{n=0}^{N_A} e^{i\pi n N_B / (\gamma N_A)} f(x_n) e^{i2\pi k n / (\gamma N_A)} \quad (16)$$

which corresponds to a Partial Fourier Transform of the series $g_n = \exp(i\pi n N_B / (\gamma N_A)) f(x_n)$. This sum can be evaluated using Bluestein's technique [35], which is related to the chirp z-transform. By noticing that $2kn = n^2 + k^2 - (k - n)^2$ the Partial Fourier Transform above can be written as a convolution:

$$\widehat{F}(u_k) = \frac{D}{N} e^{i(\pi/\gamma)(k - N_B/2)} \sum_{n=0}^{N_A} e^{i\pi n N_B / (\gamma N_A)} f(x_n) e^{i\pi k^2 / (\gamma N_A)} e^{-i\pi(n-k)^2 / (\gamma N_A)} \quad (17)$$

Consequently, the computation of the field in the image plane to an arbitrary sampling can be computed using a discrete convolution of sequences of length N_A . We direct the reader to the aforementioned references for details on these computation schemes. In order to use fast circular convolution algorithms, the incoming sequence needs to be zero-padded by factor of two, and then the computational cost is three FFTs of size $2N_A$. When the field in the image plane is computed up to the Nyquist limit, the computational cost of the method is $20N_A \log_2(N_A)$. When the size of the final array is $N_B \neq N_A$, partial convolution algorithms can be used, and they reduce the computational cost to $20N_A \log_2(N_B)$ [27].

5. Performance comparison and practical applications

5.1. Computation costs for coronagraphic propagation

We compare the computation costs of coronagraphic propagation from pupil to Lyot plane in the case of the FFT and semi-analytical methods. As described before, coronagraphic propagation with FFTs requires pupil zero-padding by a coefficient γ . The array sizes are therefore $(\gamma N_A)^2$, where N_A is the number of pixel across the pupil diameter and the number of flops involved in each FFT is $10(\gamma N_A)^2 \log_2(\gamma N_A)$. The total number of operations for a propagation from pupil to Lyot plane is:

$$n(FFT) = 20(\gamma N_A)^2 \log_2(\gamma N_A) + 6(\gamma m)^2, \quad (18)$$

where we also include the multiplication by the focal plane mask ($6(\gamma m)^2$ flops), which can be neglected in most cases of interest.

With the semi-analytical method, the size of the first FT is limited to the mask area, $N_B = \gamma m$ where m is the mask size in resolution elements (we remind here that typically $m = 5$ for a Lyot coronagraph and $m = 1$ for a Roddier or DZPM). The number of operations for the first MFT (Eq. 14). The second MFT transforms the $N_B \times N_B$ array into a $N_A \times N_A$ array and its cost is $n(MFT) = 8(N_A^2 N_B + N_A N_B^2) - 2N_A N_B - 2N_A^2$. The propagation from pupil to Lyot plane includes the two MFTs, the multiplication by the focal plane mask ($6N_B^2$ flops) and the subtraction of the result from the pupil ($2N_A^2$ flops). The total cost of a semi-analytical propagation from pupil to Lyot plane is therefore:

$$n(SAM) = 16(N_A^2 \gamma m + N_A (\gamma m)^2) + 4(\gamma m)^2 - 4N_A N_B, \quad (19)$$

where the last terms can be neglected in most cases of interest. The relative computation cost between the classical FFT propagation and the semi-analytical method is therefore approximately:

$$\frac{n(FFT)}{n(SAM)} = \frac{5 \gamma N_A \log_2(\gamma N_A)}{4 m N_A + \gamma m^2}, \quad (20)$$

neglecting the multiplication by the focal plane mask. This expression shows that there are always more operations with the FFT than with the semi-analytical method, and that the relative cost of the FFT increases with the padding factor γ and the pupil size N_A .

In Fig.4, Fig.5, and Fig.6, we illustrate the gain between the semi-analytical and FFT methods, using the ratio of their computation costs given in Eq.18 and Eq.20. We first give some

results for a typical Lyot coronagraph simulation on an 8-meter class telescope, and two cases where the computation gain can be even larger: an ELT simulation with fine pupil sampling, or a phase mask coronagraph where very fine sampling is required in the focal plane.

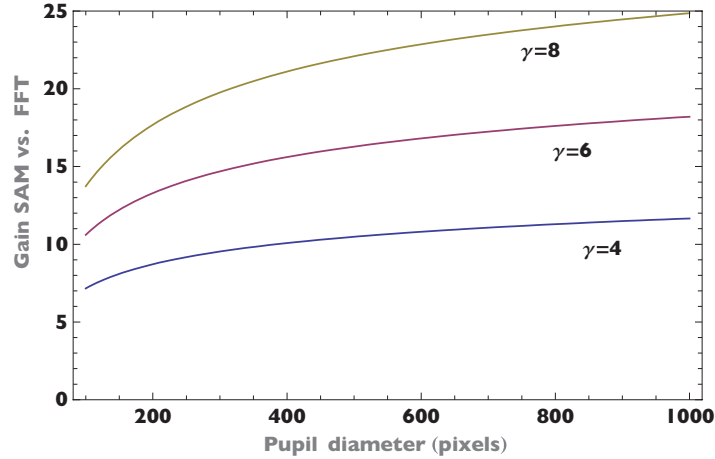


Fig. 4. Gain between SAM and FFT for a propagation from pupil to Lyot plane, as a function of the number of pixels across the pupil. The gain is based on the evaluation of the computation costs for both methods. We use common zero-padding coefficients ($\gamma = 4, 6, 8$) and a mask size $m = 5\lambda/D$. This corresponds to the regime of simulations of Lyot coronagraphs for 8-meter class telescopes.




Fig. 5. Same as Fig.4, but in the case of an ELT where a very large number of pixel is used in the pupil.

5.2. Case of direct imaging

It is also interesting to discuss briefly the case of direct imaging as a potential application of the MFT. Although in most cases the MFT requires more computations than a FFT, it is possible to identify a few cases where it can be more efficient. This happens when the calculated area is limited and/or when a fine sampling is needed. We illustrate the case of a point spread

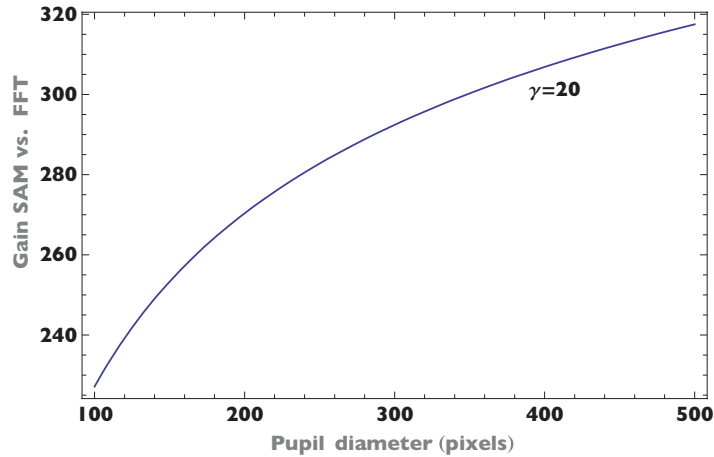


Fig. 6. Same as Fig.4, but for a Roddier or Dual Zone coronagraph with a mask size of $m = \lambda/D$ and an oversizing $\gamma = 20$. This corresponds to a maximum FFT size of 10000×10000 for the 500 pixel pupil.

function (PSF) calculation with a field of view (FOV) of 50 resolution elements in Fig.7. In this case, both techniques have similar costs with a slight advantage for one or the other depending on pupil size and sampling parameter γ . If the PSF is only calculated to a FOV of 20 resolution elements, the MFT provides a gain by a factor of several over the FFT (Fig.8). In some applications where only the center of the PSF is needed, gains comparable to that of the coronagraphic cases can be obtained.

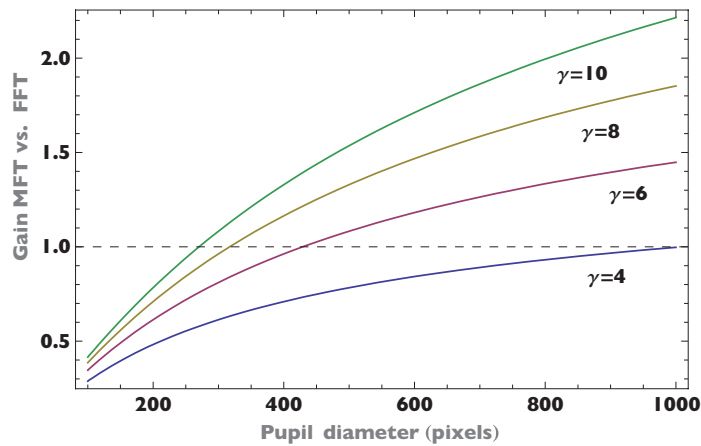


Fig. 7. Non-coronagraphic, direct imaging with FFT and MFT where the PSF is calculated to a field of view of 50 resolution element (25 in radius). Both MFT and FFT have similar costs with a slight advantage to MFTs when higher sampling is used.




Fig. 8. Non-coronagraphic, direct imaging with FFT and MFT where the PSF is calculated to a field of view of 20 resolution element (10 in radius). For a limited field of view, the MFT can be significantly faster than the FFT. This can be interesting for coronagraphic and extreme adaptive optics simulations.

5.3. Practical applications and implementation

In this section we discuss the actual performance of the semi-analytical method for a specific hardware, a 2003 Apple G5 dual processor 2Ghz computer with 4Gb RAM. We measured the performance of this machine using the package FFTW (version 3) [31] and report the results in Fig.9. We considered the case of double precision complex 2D arrays, which is appropriate for coronagraphic simulations. These results are consistent with those given on the FFTW website, but we extend them up to 8192×8192 FFT to verify that the speed does not depend strongly on the array size. The calculations have been made with the option `FFTW_PATIENT`, which optimizes the FFT plan prior to the calculation of the FFT itself. We used single-threaded FFTs although the FFTW packages provides the possibility of multi-threaded calculations for further speed improvement on multi-core machines.

In order to evaluate the machine's performance for the MFT, we considered the geekbench benchmark software [36] which reports 1.4 Gflop/s for single-threaded dot products and 2.6 Gflop/s for multi-threaded ones. It seems that dot products are slightly more efficient than the FFTW on this hardware (1.4 Gflop/s vs. 1 Gflop/s). This may be due to the easier implementation of the dot-product which takes direct advantage of the processor's vector engine. An interesting aspect of the MFT is that it can take advantage of recent multi-core computers, with easy and efficient multi-threading of linear algebra calculations.

High level languages such as Mathematica or Matlab are well optimized for vector computations. For example, we obtain 2.5 Gflop/s with Mathematica for multi-threaded dot-products on double precision complex arrays of sizes between 1000×1000 and 10000×10000 . This result is consistent with the benchmark results this machine (2.6 Gflop/s). In the case of a propagation from pupil to Lyot plane, the speed goes down to 1Gflop/s with Mathematica because the other operations involved in the semi-analytical method are not fully optimized with this language.

In Table 1, we compare the two methods for three pupil sizes (256×256 , 512×512 , 1024×1024). N_A and N_B are the number of pixels used in the calculation, in the pupil and focal planes. We use the same number of pixels in the Lyot plane as in the pupil. We show the effect of the sampling improvement from $\lambda/4D$ to $\lambda/8D$, obtained by zero-padding the pupil with FFTs.




Fig. 9. Normalized speed expressed in Mflop/s using the FFTW3 package [31] on a 2GHz dual-processor 2003 Apple G5 with 4Gb RAM. The normalized speed is defined as the number of operations divided by the execution time: $10N^2 \text{Log}_2(N)/T$ where T is the time for one FFT (single-threaded, double precision complex arrays). The red dots correspond to power of twos arrays. This hardware delivers approximately 1Gflop/s over the range of array sizes, which translates for example into 2.3s for 4096×4096 transforms, and 10s for 8192×8192 transforms.

Note that for the semi-analytical method, the number of pixels N_A corresponds to the pupil size itself, and that N_B corresponds to the focal plane mask size. For the FFT, N_A corresponds to the zero-padded pupil. The timings given in the table correspond to the time for 2 FFTs calculated with FFTW3, neglecting the application of the focal plane mask and zero-padding time. For the semi-analytical method, we give timings obtained with Mathematica as example.

Coronagraphic simulations are usually performed with a $\lambda/6D$ to $\lambda/8D$ sampling, in order to keep the FFTs manageable. With the semi-analytical method, we can increase the precision significantly, and we typically use a $\lambda/10D$ or $\lambda/20D$ sampling for the focal plane mask, including gray pixel approximation. This corresponds to a $40 \sim 50$ or $80 \sim 100$ pixel diameter mask for a $4 \sim 5$ resolution element mask.

The semi-analytical method offers a total flexibility in terms of sampling in both planes, and there is no particular limitation to the type of calculations. We present a few examples in Table 2, which would hardly be achieved with FFT propagations. For example, one can choose to increase the sampling in the pupil, with a modest $\lambda/8D$ sampling in the focal plane, in order to study the effect of segmentation of an ELT. We show an example of a 6000×6000 pupil array. This corresponds to calculations that are typically made on very large computer clusters using the classical method. This would correspond to 36000×36000 FFTs with $\gamma = 6$. It is also possible to oversample dramatically the occulting mask. This can be used to study the effect of the actual mask shape, such as ellipticity, edge roughness, fine alignment, or effects of atmospheric differential refraction. We give an extreme example with 1000 pixels in the pupil and a resolution of $\lambda/200D$. This calculation would correspond to 200000×200000 FFTs.

Another application is the simulation of Roddier or Dual Zone phase masks (DZPM), where the masks have a typical size of λ/D , making the sampling problem even more critical for FFTs. The optimization of DZPM [19], which was obtained using the semi-analytical approach by integrating directly the Hankel transforms, can be studied by direct numerical simulations with our method.

Table 1. Performance comparison of FFT-based and semi-analytical methods (SAM). N_A is the number of pixels across the pupil diameter, and N_B is the size of the array in the focal plane. The FFT timings correspond to 2 FFTs calculated with FFTW3, and the SAM timings correspond to an actual propagation from pupil to Lyot plane, with Mathematica.

Method	Pupil Plane		Image Plane		Time
	N_A	Array size	Sampling	Array size (N_B)	
FFT	256	1024	$\lambda/4D$	1024	0.22 s
FFT	256	2048	$\lambda/8D$	2048	1.0 s
SAM	256	256	$\lambda/8D$	40	0.06 s
FFT	512	2048	$\lambda/4D$	2048	1.0 s
FFT	512	4096	$\lambda/8D$	4096	4.7 s
SAM	512	512	$\lambda/8D$	40	0.19 s
FFT	1024	4096	$\lambda/4D$	4096	4.7 s
FFT	1024	8192	$\lambda/8D$	8192	20.2 s
SAM	1024	1024	$\lambda/8D$	40	0.70 s

Table 2. Performance of semi-analytical simulations for double precision complex arrays, with their possible applications. These calculations cannot be performed using the FFT method on commercial workstations, as they would correspond respectively to 200000×200000 , 25600×25600 , and 48000×48000 FFTs.

Method	Pupil Plane		Image Plane		Time	Application
	N_A	Array size	Sampling	N_B		
SAM	1000	1000	$\lambda/200D$	1000	13.5 s	Mask Tolerancing
SAM	3200	3200	$\lambda/8D$	40	6.1 s	4 pixel/spider on Gemini
SAM	6000	6000	$\lambda/8D$	40	22 s	5 mm/pixel for 30 m ELT

6. Conclusion

In this paper we introduced a semi-analytical method to calculate the numerical propagation through a Lyot-style coronagraph, without any additional approximation. We showed that FFTs are not appropriate for Lyot coronagraphy because of fundamental sampling limitation, and should be avoided for these calculations. The semi-analytical method is derived straightforwardly from the analysis of the wave propagation, and requires the use of a Fourier transform method producing arbitrary sampling in a limited area. Several methods exist to perform such transforms, and we suggest a matrix-based Fourier propagator that can be implemented straightforwardly and efficiently in any language.

For coronagraphic studies for ExAO on ELTs and eight meter class telescopes, the typical speed improvement based on computation costs is of at least a factor twenty to fifty compared to the classical FFT coronagraphic propagation method. In addition to speed, our semi-analytical method is particularly efficient in terms of memory, as it does not involve any zero padding, and typical memory improvement is of a factor fifty. In the case of a full propagation from pupil to final focal plane, if the final FT is performed using a FFT, the semi-analytical method is approximately three times faster than a FFT-based propagation, since the propagation time is entirely dominated by the final FT. If a reduced field of view is calculated in the final image,

which is often the case in extreme adaptive optics coronagraphy, an MFT can be used in lieu of an FFT. For example a factor of 5 can be gained on the last FT for $N_A = 400$ and $\gamma = 10$, and the overall gain with the semi-analytical method is a factor of 15.

The semi-analytical method brings a total flexibility in terms of sampling, since the sampling in the pupil and focal plane are totally independent. This enables calculations hitherto impossible, or requiring access to large and expensive computers clusters. For example the study of fine occulting mask features and alignment, or fine pupil structures like segmentation, is now achievable on a good laptop or on any commercial workstation.

The method may find interesting applications where a FT has to be calculated in a limited area. Applications may be found in adaptive optics for anti-aliasing spatial filters, pyramid wavefront sensing, etc.

Our method also enables fast computation of optimal apodizer for APLCs, which involves an iterative algorithm and makes computing efficiency even more critical. APLCs are one of the most promising concepts for high contrast imaging instruments on ELTs, and their study requires the optimization of the mask size according to various criteria and for broadband light. This means that a theoretical apodizer has to be calculated for each mask size in the optimization study [17]. A complete exploration of the parameter space, including effects of segmentation and spiders on the coronagraph is now made possible with acceptable computing times. Our method could be utilized in several current efforts focussed on ground and spaced based coronagraphy.

Acknowledgements

RS acknowledges partial support by a Michelson Postdoctoral Fellowship, under contract to the Jet Propulsion Laboratory (JPL) funded by NASA. JPL is managed for NASA by the California Institute of Technology and by a AMNH Kalbfleisch Fellowship. This work has also been partially supported by the National Science Foundation Science and Technology Center for Adaptive Optics, managed by the University of California at Santa Cruz under cooperative agreement AST 98-76783. The authors thank Lisa Poyneer, Jim Fienup for interesting discussions and Russell Makidon for helpful comments.