

Comparison of methods for exploiting symmetry in calculating the Fourier spectra of circularly symmetric functions

Joseph W. Goodman*

Stanford University, Department of Electrical Engineering, Palo Alto, California, United States

Abstract. Four different methods for calculating the radial profile of the Fourier transform of a circularly symmetric function are discussed and evaluated. Applications to optics, in which the circularly symmetric function represents the distribution of field across a circular aperture, are many. The four approaches include: (1) numerical integration of the zero-order Hankel transform equation (using the Gauss–Kronrod method); (2) extraction of the radial profile of a two-dimensional Fourier transform; (3) the projection-transform method that makes use of the projection-slice theorem of Fourier analysis; and (4) the quasifast Hankel transform introduced by Siegman. The results indicate that the projection-transform method is the fastest method and that it has accuracy that is second only to numerical integration. © 2020 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.OE.59.8.083105](https://doi.org/10.1117/1.OE.59.8.083105)]

Keywords: Fourier transform; Fourier optics; circular symmetry; diffraction; comparison of methods.

Paper 20200715 received Jun. 17, 2020; accepted for publication Aug. 5, 2020; published online Aug. 24, 2020.

1 Introduction

The purpose of this paper is to explore methods for fast calculation of Fourier transforms of two-dimensional (2-D) functions with circular symmetry. It comes as no surprise that taking advantage of symmetries should potentially speed up such calculations. While existing programs such as MATLAB and *Mathematica* can calculate Fourier transforms rather quickly, speedups could be helpful in many programs that use iterative methods involving a Fourier transform on each pass. Examples include resonator calculations, phase retrieval calculations, computer-generated phase-only hologram design, as well as others. We explore and compare the calculation speeds of several different methods that exploit circular symmetry.

Throughout we will be using Fourier and inverse Fourier transforms defined as

$$\begin{aligned} G(\nu_x, \nu_y) &= \iint_{-\infty}^{\infty} g(x, y) \exp[-j2\pi(\nu_x x + \nu_y y)] dx dy \\ g(x, y) &= \iint_{-\infty}^{\infty} G(\nu_x, \nu_y) \exp[+j2\pi(\nu_x x + \nu_y y)] d\nu_x d\nu_y. \end{aligned} \quad (1)$$

One-dimensional (1-D) versions of these transforms will also be used.

For circularly symmetric functions, the Fourier transform is also circularly symmetric. This fact allows the transform to be characterized by what will be called a “radial profile,” which is a central slice through the spectrum extending from the origin to some predetermined upper frequency limit. For example, the zero-order Hankel transform, discussed below, calculates a radial profile of the spectrum, valid from the origin to, in principle, an infinite limit of frequency. All versions of the Fourier transform will be compared on the basis of their radial profiles, but to some finite upper limit of frequency beyond which the values of the spectrum are not of interest.

The novelty of this study lies not in introducing new techniques, but rather in comparing various existing techniques. We have not considered in this study methods that require

*Address all correspondence to Joseph W. Goodman, E-mail: goodman@ee.stanford.edu

prestorage of values of Bessel functions of various orders and various arguments. Such methods are more user-implementation dependent and should be evaluated for both speed and accuracy.

2 Circularly Symmetric Functions in Optics

We consider only circularly symmetric functions describable by a complex-valued function of radius $g(r)$, where $r = \sqrt{x^2 + y^2}$. In optics, the cases in which circularly symmetric functions are encountered are plentiful. Circular symmetry is often found in sources, lenses, and exit pupils. Note that the circularly symmetric function can contain both circularly symmetric attenuations and circularly symmetric phases. Calculation of focal plane distributions for circular-symmetric systems with spherical aberration and focusing errors satisfies the circularly symmetric constraint. Calculation of Fraunhofer and Fresnel diffraction patterns of circularly symmetric apertures requires Fourier transforms of circularly symmetric functions. In what follows, we examine four different approaches to calculating the Fourier transforms of such functions, and ultimately compare their computational complexities, error properties, and speed in a specific example.

3 Background on the Methods

3.1 Integration of the Zero-Order Hankel Transform Equation

The traditional method for finding the radial profile of the Fourier spectrum of a circularly symmetric function is by means of the zero-order Hankel transform. Let $g(r)$ again represent the radial profile of the aperture function and $G(\rho)$ the radial profile of its circularly symmetric Fourier spectrum. Then

$$G(\rho) = 2\pi \int_0^{\infty} r g(r) J_0(2\pi r \rho) dr, \quad (2)$$

where $J_0(x)$ is the zero-order Bessel function of the first kind. In some cases, this integral can be directly evaluated by a symbolic integration program, yielding an analytic solution. In others, numerical integration can be used to obtain a highly accurate result. Both of these methods are relatively slow when compared with methods based on FFTs. However, we take these results as the gold standard against which the errors associated with other methods are measured.

3.2 Direct 2-D Fourier Transform

It is always possible to construct a rectangular array of samples of $g(\sqrt{x^2 + y^2})$ and perform a 2-D FFT on the array, eventually displaying a radial profile of the result by selecting half of a proper column or row of the transform. Such an approach is straightforward. However, there is a question to be answered that requires some experimentation. Namely, how much is it necessary to pad the samples of a circular aperture with zeros in order to get adequate sampling in the frequency domain and thereby get a reasonably accurate solution. Let M represent the number of samples along a diameter of the circular aperture on the vertical or horizontal axes, and let N represent the total number of samples along the same direction for the padded array. The answer to this question will be deferred to the later section on errors. The computation time required to find the Fourier transform of the circular aperture is the time required for an $N \times N$ FFT. Such a computation requires $2N^2 \log_2 N$ complex multiplies and $2N^2$ complex additions.

3.3 Projection-Transform Method

The projection-slice theorem of Fourier analysis states that a projection through a 2-D function $g(x, y)$ onto an axis at angle $+\theta$ with respect to the x axis has a 1-D Fourier transform that is a slice through the 2-D spectrum G of g along an central axis at angle $+\theta$ to the ν_x axis in the frequency domain. Note that for a circularly symmetric function, projections in all directions are

the same, so any one projection direction contains all the needed information. A final 1-D Fourier transform of the padded projection then yields a slice through the 2-D spectrum, from which the radial profile of the circularly symmetric spectrum is known.

The origin of the projection-transform method for finding Fourier transforms of circularly symmetric functions appears to lie with the work of Bracewell in 1956,¹ but the method was introduced to the signal processing community by Oppenheim et al., in 1978² (see also Refs. 3 and 4). The latter works apply to evaluation of Hankel transforms of any order, but here only the zero-order transform is of interest due to the assumption of circular symmetry.

Suppose that a circularly symmetric function containing variations of amplitude and/or phase is sampled in an $M \times M$ rectilinear grid. If M again represents the number of samples along a diameter of the circular aperture on the horizontal or vertical axes, projections onto such axes require summations of M samples down each of M columns (or across each of M rows), for a total of M^2 complex additions. It is possible to perform column summations for only those sample points that lie within the aperture, rather than in a full $M \times M$ rectangular array, but this only reduces the number of samples by a factor of about $\pi/4$, the ratio of the area of a circle of diameter b to the area of a square of side b . The 1-D Fourier transform is carried out by an FFT, but padding of the projection sequence with zeros is necessary to sample the transform closely enough to yield accurate results. Exactly how much padding is needed will be discussed in the later section on errors. The computation requires $N \log_2 N$ complex multiples and $N + M^2$ complex additions, where the term M^2 arises from the projection operation.

The number of additions can be reduced further by projecting only one quarter of the circular aperture onto the horizontal (or vertical) projection axis. This reduces the number additions to $(M/2)^2$ from M^2 . After computing the $(M/2)^2$ projection samples, they can be doubled in value to complete half of the projection samples, and then that result can be reflected about the origin to obtain the symmetrical results, thus completing the full projection sequence. Note that if this method is used to reduce the number of additions, M should be an even number so that there are no samples on the horizontal or vertical axes that will be counted twice. When comparing computation times, the time required to fill the full projection sequence from the partial result should be included. In later sections, we ignore this potential speedup and perform projections through the entire square array of size $M \times M$.

3.4 Quasifast Hankel Transform

Siegman⁵ devised a method he called the quasifast Hankel Transform that can be applied here. Siegman credited this method to Gardner et al.,⁶ although Agrawal and Lax⁷ pointed out a much earlier use of a similar transformation. This method makes exponential substitutions for r and ρ in Eq. (2), converting the Hankel transform to a cross-correlation integral. This method can be applied to a Hankel transform of any order, but we focus on the zero-order transform here because we have assumed circular symmetry. The substitutions are

$$r = r_0 e^{\alpha x}, \quad \rho = \rho_0 e^{\alpha y}, \tag{3}$$

yielding a correlation integral

$$\hat{G}(y) = \int_{-\infty}^{\infty} \hat{g}(x) \hat{j}(x + y) dx, \tag{4}$$

where $\hat{g}(x) = r_0 e^{\alpha x} g(r_0 e^{\alpha x})$, $\hat{G}(y) = \rho_0 e^{\alpha y} G(\rho_0 e^{\alpha y})$, and the Bessel kernel becomes

$$\hat{j}(x) = 2\pi \alpha r_0 \rho_0 e^{\alpha x} J_0(2\pi r_0 \rho_0 e^{\alpha x}). \tag{5}$$

We must now discretize $\hat{g}(x)$ and $\hat{j}(x)$ in order to create a discrete convolution that can be carried out by a series of fast Fourier transforms. One problem is evident immediately. Since the lower limit 0 of the Hankel transform has mapped to $-\infty$ in the correlation integral, any finite discrete approximation to the correlation integral must have a region of zeros at and around the origin, a region we call the “donut hole.” Agrawal and Lax⁷ addressed this problem by simply adding to

the result a constant in the donut hole, the constant being the value of the Hankel transform at zero frequency.

Uniformly spaced sampling is done in the x and y domains, necessary for FFTs to be applied to calculating the correlation integral. Such sampling results in nonuniformly spaced samples in the r and ρ domains. The correlation integral is calculated via a series of FFTs according to

$$\hat{G}(y) = \text{FFT}^{-1}(\text{FFT}[\hat{g}(y)] \times \{\text{FFT}[\hat{j}(y)]\}^*), \quad (6)$$

where FFT is the forward fast Fourier transform, FFT^{-1} is the inverse fast Fourier transform, and $*$ indicates the conjugation.

Represent again the number of samples across a horizontal or vertical diameter of the circular function by M . The question remains as to the number of samples N required in the computation. In Siegman's procedure, if $M/2$ samples of \hat{G} are desired, \hat{g} is padded with $M/2$ zeros to make a length $N = M$ sequence, and \hat{j} is extended to an M length sequence. These length M sequences are Fourier transformed, and only the first $M/2$ samples of the correlation are retained, thus deleting negative-frequency components that are not part of the radial profile. The total number of complex multiplications required is therefore $3(M \log_2 M) + M$ and the number of complex additions is $3M$. These measures of complexity do not include the computations required to transform $rg(r)$ into $\hat{g}(x)$, $J_0(2\pi\rho)$ into $\hat{j}(x)$, and $\hat{G}(y)$ into $G(\rho)$. These coordinate transformations can take a significant fraction of the total computation time for this method, as illustrated in a later section.

Siegman⁵ derives several relationships between the parameters used in this problem, based on sampling requirements. If $\rho = \beta$ is the highest frequency of interest in the ρ (frequency) domain, then the lower end truncation point r_0 in the r domain, where the sample spacing is finest, should not be larger than $1/K_1$ cycles of the highest frequency β . The constant K_1 should be greater than or equal to 2 but is usually chosen to be 4 to be safe. Similarly, at the upper truncation point, $r = b$, the sample spacing should not be greater than $1/K_2$ cycles of β , where again K_2 is usually chosen to be 4. Analogous arguments can be applied in the ρ domain, leading to a set of equations that should be satisfied in choosing the parameters for the computation:

$$M/2 = K_2\beta b \log_2(K_1\beta b), \quad (7)$$

$$\alpha \exp(\alpha M/2) = K_1/K_2, \quad (8)$$

$$r_0\rho_0 = (K_2/K_1^2)\alpha. \quad (9)$$

Given a choice for the size M of the calculation and choices for K_1 and K_2 , Eq. (7) represents a nonlinear equation that can be solved for βb , which is the space bandwidth product of the calculation. Likewise, given choices of the same parameters, Eq. (8) defines a nonlinear equation that can be solved for α . Finally, from Eq. (9), the product $r_0\rho_0$ can be calculated once the value of α is known.

Figure 1 shows plots of the radial profiles obtained for three different choices of total number of samples M using the quasifast Hankel transform. As can be seen from these results, in all three cases, the donut hole has a radius of about $\Delta\rho \approx 0.1$ to the right of the origin, but grows slightly as M is reduced. The primary differences between these results are twofold. First, the gap between the value of the result at the edge of the donut hole and the exact result at the origin grows larger as M is reduced. Second, the maximum value of ρ for which the profile can be plotted also grows smaller as M is reduced.

The severity of the donut-hole problem depends on how one chooses to divide up the space-bandwidth product between the radius b of the circle in the space domain and the radius β in the frequency plane over which one desires accurate results. Figure 2 shows the consequences of different divisions of a constant space-bandwidth product βb between b and β . In the left column are the "exact" results obtained from the analytic solution of the Hankel transform equation. On the right are the results obtained by the quasifast algorithm for different choices of b and β , their product being held constant. All other parameters of the algorithm are held constant, including the total number of samples ($M = 256$) in the calculation. In accord with the requirements of the

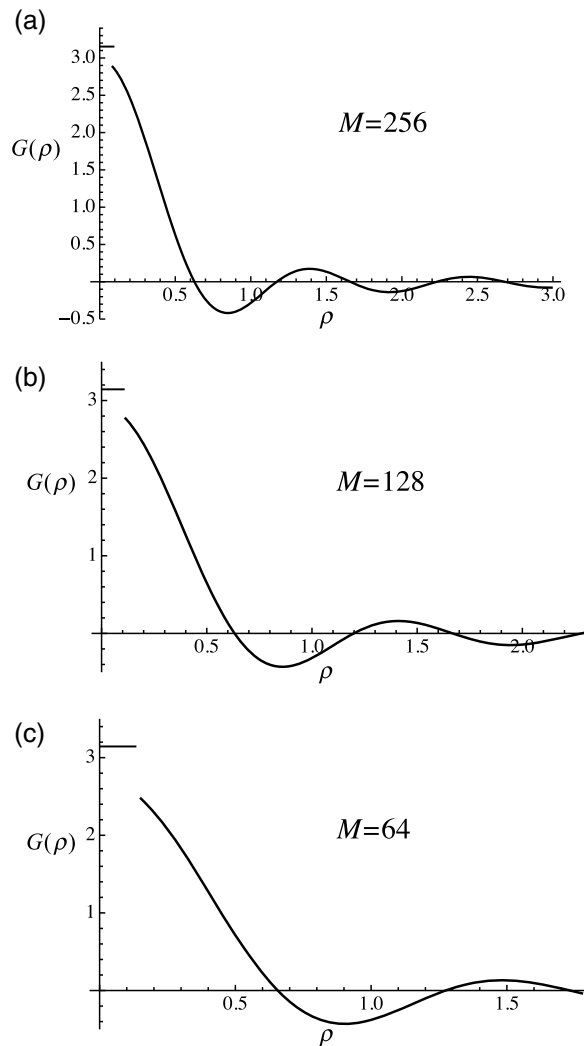


Fig. 1 Results of calculations of the radial profiles of Fourier transforms of a uniform circular aperture using the quasifast Hankel transform algorithm, with the following parameter values: (a) $M = 256$, $\alpha = 0.0279$, $r_0 = \sqrt{\alpha}/2$, $\rho_0 = \sqrt{\alpha}/2$, $b = 1$, $\beta = 8.945$, $K_1 = K_2 = 4$; (b) $M = 128$, $\alpha = 0.047$, $r_0 = \sqrt{\alpha}/2$, $\rho_0 = \sqrt{\alpha}/2$, $b = 1$, $\beta = 5.254$, $K_1 = K_2 = 4$; and (c) $M = 64$, $\alpha = 0.079$, $r_0 = \sqrt{\alpha}/2$, $\rho_0 = \sqrt{\alpha}/2$, $b = 1$, $\beta = 3.155$, $K_1 = K_2 = 4$. The horizontal lines at 3.1416 represent the value of the analytic profile at the origin.

algorithm, the last 128 of the calculated samples are discarded. The first 128 samples are displayed with second-order interpolation.

The values at the origin of the analytic solution change with the value of b because the areas of the circular apertures are changing. The Fourier transforms also become more concentrated toward the origin as the radius in the space domain increases, as expected. The donut hole is smallest for small b and large β but increases with increasing b and decreasing β . In the case at the right-hand bottom of the figure, for which $b = 8.95$ and $\beta = 1$, the donut hole encompasses both the main lobe and part of the first sidelobe off the radial profile, yielding a very poor result.

3.5 Comparison of Computational Complexities of the Above Methods

Table 1 summarizes the operation counts for three of the approaches.

From a purely theoretical point-of-view, based on the computational complexities, we see that the 2-D FFT has the greatest computational complexity of the three FFT-based methods by far. The projection-transform method has fewer complex multiplies and additions than the 2-D

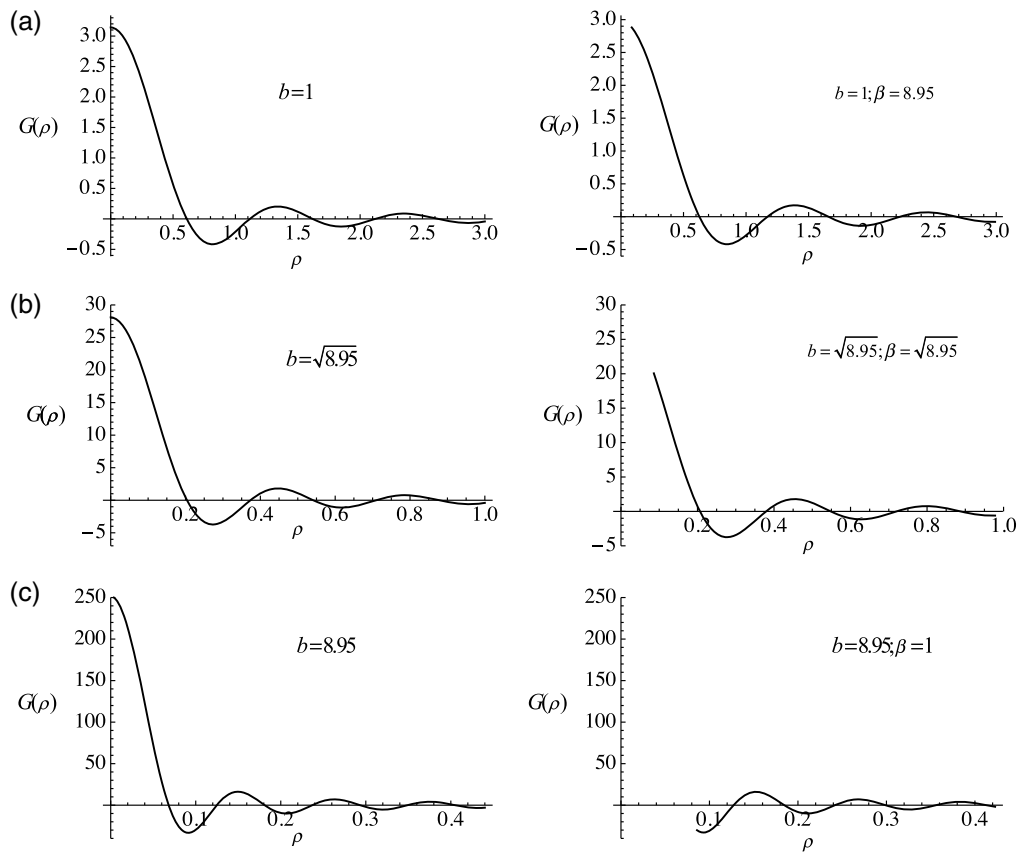


Fig. 2 Results of calculations of the radial profiles of Fourier transforms of a uniform circular aperture of radius b for constant space-bandwidth product βb but a changing division between the space factor b and the frequency factor β . The left column contains results from analytic solutions of the Hankel transform equation for a uniform circular aperture. The right column contains the results of the quasifast algorithm corresponding to each such result. The parameters b and β are divided as follows: (a) $b = 1, \beta = 8.95$; (b) $b = \beta = \sqrt{8.95}$; and (c) $b = 8.95, \beta = 1$. The values of the other parameters used are in all cases: $\alpha = 0.0279484, r_0 = \sqrt{\alpha}/2, \rho_0 = \sqrt{\alpha}/2, K_1 = K_2 = 4,$ and $M = 256$.

Table 1 Numbers of complex multiplies and complex additions for three methods for computing a frequency-domain radial profile. An array size of $N \times N = 4M \times 4M$ has been assumed for the 2-D FFT case, an array length of $4M$ for the 1-D FFT case, and only M samples in the quasifast Hankel transform approach.

Method	Multiplications	Additions
2-D FFT	$2N^2 \log_2 N$	$2N^2$
Projection transform	$N \log_2 N$	$N + M^2$
Quasifast Hankel	$3(M \log_2 M) + M$	$3M$

FFT. Complex multiplies take more time than complex additions, so the number of multiplies is particularly important. Finally, the quasifast Hankel transform has the least complex multiplies and the least complex additions. However, as mentioned previously, the computational complexity cited does not include the time taken to perform the required coordinate transformations, which can be a significant fraction of the total time required for this method. When comparing these methods in a real example, these times will be included. Finally, the times required to determine the parameters for this transform have been ignored, since once $M, K_1,$ and K_2 are

chosen, many transforms can be performed with these same parameters. The computation times realized in practice will depend on the computer and the software used.

Computational complexity alone does not include information about the errors or the actual computation times associated with the various methods. We turn attention to the error properties of these algorithms next.

4 Comparison of Absolute Errors at the Sampling Points

The total absolute errors of each of the three methods (other than numerical integration) are assessed in this section. There are three potential sources of error. First, a major source of error is associated with the discrete representation of a continuous circle. Note that as we increase the number of samples representing the circle, the smaller the errors due to this source will be. Note also that the quasifast Hankel transform does not suffer from this source of inaccuracy. A second source of error is undersampling of the transform, with the result that the accuracy of the representation of the transform is decreased and interpolation errors occur. In this case, it is the padding of the discrete space-domain function that determines the spacing of the samples in the transform domain, and therefore padding the array with a large number of zeros decreases these errors. A third source is aliasing in the transform domain, which is potentially associated with any discrete Fourier transform. The level of aliasing is determined by the spacing of the discrete samples in the space domain. Thus, increasing the density of the samples in the space domain simultaneously decreases two types of error: the error representing the discrete representation of the circular function and the aliasing error.

In what follows, it is convenient to assess errors associated with the absolute value of the transform, since the absolute values are independent of exactly where in the array of samples the circular function is placed. The choice of a uniform circle as the space-domain function is convenient because the exact transform is analytically known:

$$G(\rho) = bJ_1(2\pi b\rho)/\rho, \quad (10)$$

where b is the radius of the circular space-domain function.

4.1 2-D Fourier Transform Method

Simulations have shown that the errors associated with the 2-D Fourier transform method and the projection-transform method are identical. This is perhaps not surprising in view of the projection-slice theorem. Selection of a radial profile from the 2-D transform is in effect selecting half of a slice through that transform. Nonetheless, there is at least one topic worth discussing separately, which we defer to the projection-slice error section.

The number of samples of the circle in a diameter along either the x or the y axes is again represented by M and the total padded square array of samples is size N on a side. Consider Fig. 3 in which are shown (a) absolute value of samples of the transform, as calculated by either the 2-D Fourier method or the projection-transform method (the results are the same), (b) the corresponding transform obtained from the exact analytical solution, and (c) a superposition of interpolated samples of the exact analytical solution and the errors represented by the absolute differences between the calculated and exact solutions. Linear interpolation has been used between sample points in the first two cases. The vertical axes have logarithmic scales, and the transforms have been normalized to unity at the origin. Because of this normalization, the errors shown can be viewed as fractional errors with respect to the maximum value of the transforms, which occur at the origins. As indicated in the caption, these results are for the specific choices of $M = 128$ and $N = 512$ and are presented as an example of the relevant results, from which maximum errors are determined.

Figure 4 shows a plot of the maximum value of the errors as defined above, for different values of M and N . In all cases, $N = 4M$; that is, the $M \times M$ samples of the circle are padded by $3M/2$ zeros in each direction.

What conclusions can we draw from these results? The primary conclusion is that as M is increased while holding N as a constant multiplier of M , the density in the sampling of the

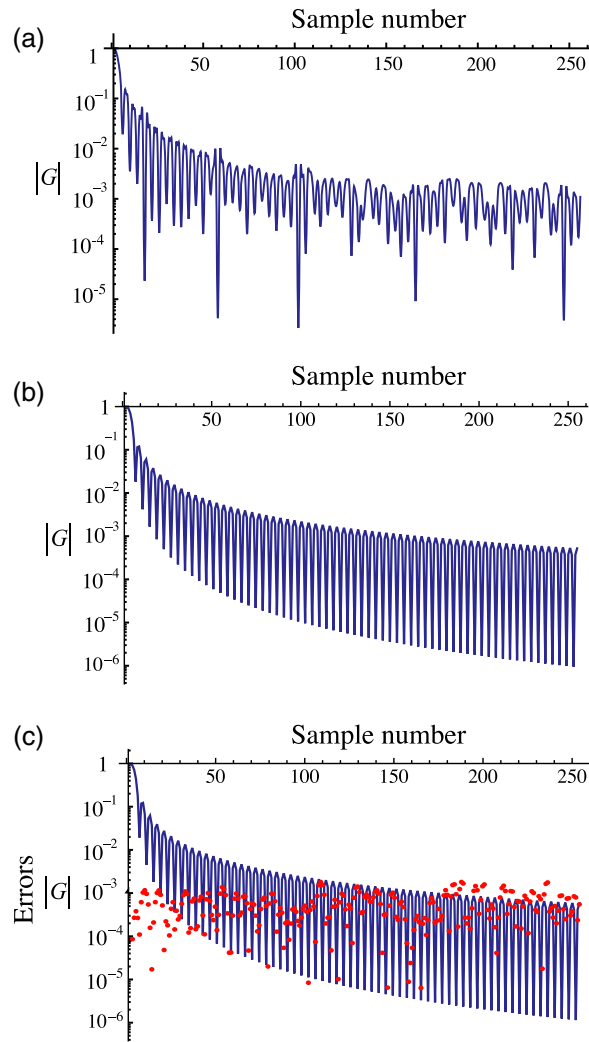


Fig. 3 Results obtained with the 2-D Fourier transform method and the projection-transform method. (a) Radial profile obtained by discrete calculation; (b) radial profile obtained from samples of the exact analytical solution; and (c) errors in the absolute value of the transform at each sample location overlaid with the same curve as (b). For these results, $M = 128$ and $N = 512$. Results are only shown for the first $N/2$ samples, which represent the radial profile. Linear interpolation has been used for the continuous curves.




Fig. 4 Error maximum absolute values calculated by taking the maximum of the absolute values of the differences between the absolute values of the 2-D FFT result and the absolute value of the sampled exact analytical result, both for $N/2$ points in the two spectra. M takes the values 32, 64, 128, 256, 1024, and 2048, and in all cases $N = 4M$.

space-domain function increases, thus increasing the accuracy of the digitized circular function. (See Ref. 8 for an improved representation of a discrete circle.) At the same time, the density of samples within the transform remains constant, a consequence of keeping N as a constant multiplier of M . The reduction of errors as M is increased is primarily due to the reduction in errors in representing the uniform circular function by a discrete array with a finite number of samples. Figure 3 does show some increase in the level of errors when the transform coefficient index goes beyond about $N/4$, which may be a consequence of some aliasing, but the effect on sample values with index smaller than $N/4$ appears to be negligible.

4.2 Projection-Transform Method

As has been said, the projection-transform method yields exactly the same results as the 2-D Fourier transform method, as far as the radial profile of the spectrum and the errors are concerned.

There is one result concerning the projection operation that is interesting and will be presented here, because it sheds some light on the errors associated with a discrete approximation to a circular function. It is easy to calculate an analytical expression for the projection of a uniform circle with value unity and radius b onto any central axis. The result is

$$p(x) = \int_{-b}^b g(x, y) dy = 2\sqrt{b^2 - x^2}, \tag{11}$$

where b is again the radius of the circle represented by $g(x, y)$. This exact expression for the projection can be compared with the discrete projection obtained by integrating down each column (or across each row) of the rectangular array representing the digitized circle. By gradually increasing M , the transition between a bad representation of the circle and a good representation of the circle will be seen.

Figure 5 shows plots of the projection of a uniform discrete circle for various numbers of $M \times M$ samples in the representation of the circle. The differences between the exact analytic projection of the circle and a discrete projection through a digitized circle are the main source of errors for the projection-transform method. But the errors associated with the projections arise from the digitized circle itself, just as they do for the 2-D FFT method, and ultimately they are the same for these two methods.




Fig. 5 Projections through an $M \times M$ digitized circle for various numbers of samples.

4.3 Quasifast Hankel Transform Method

Evaluating the errors in the quasifast Hankel transform method is a bit more difficult than for the other methods. The reason is that there are a number of different parameters to choose, and different choices yield different advantages. Note that Eqs. (7) and (9) only specify the products βb and $r_0\rho_0$, but not the values of the individual variables in these products. As a consequence, experimentation is required to find the best assignments for these variables. In all cases, the space-domain function chosen is a uniform circle of value 1 and radius b , since the analytic solution for the transform is known.

With respect to the product βb , it has already been mentioned that the donut hole has been found to be present but minimized when we choose $b = 1$ and choose β as the total product βb . With respect to the product $r_0\rho_0$, we adopt Siegman's choice of $r_0 = \sqrt{\alpha}/2$ and $\rho_0 = \sqrt{\alpha}/2$ which yields well-shaped sidelobes.

Figure 6 shows a plot of the absolute value of the analytic solution for the main lobe and several sidelobes when $M = 512$ samples are taken inside the exact analytical solution. The samples are not equidistant, due to the nonlinear stretching of the ρ axis characteristic of the quasifast Hankel transform method. Shown on the same plot as discrete points are the error levels at the sampling points, where the error is between the samples of the transform computed by the quasifast Hankel transform and the exact values taken from the analytical solution. Note that the maximum errors occur where the slope of the transform is highest.

Finally, in Fig. 7 is shown the maximum error as a function of the number M of samples in the radius of the circle on a loglog plot. All the error values for the quasifast Hankel transform are larger than the corresponding error values for the same values of M with the 2-D Fourier transform or the projection-transform approaches. This is true even though the errors associated with the donut hole have not been included. Padding the two sequences \hat{f} and \hat{j} with additional zeros has not been found to improve the maximum errors appreciably. The estimates of error have been made with certain choices of the parameters relevant for the quasifast Hankel transform. Our choice of parameters followed Siegman's choices. It is always possible that there are other choices among the myriad of possibilities that might yield lower errors than reported here, but we have been unable to find them.

4.4 Numerical Integration of the Hankel Transform Equation

Since the exact analytic solution of the zero-order Hankel transform of a uniform radial profile of a circle is known, it is a simple matter to find the absolute errors between the exact result and the result of numerical integration of the Hankel transform equation. The method used for numerical integration in the Gauss–Kronrod method, which is the default method *Mathematica* chooses for this problem. Calculations for various values of M show that these absolute errors are at




Fig. 6 Calculated absolute spectrum (normalized) and error values for the quasifast Hankel transform of a uniform circle. The solid curve represents the result of the calculation, with linear interpolation between sample points. The red dots represent errors at the sample points. Both the solid curve and the error values are normalized by the value the transform would have at the origin. The donut hole has been preserved in the solution, but no errors have been represented in that region. The number of samples in the diameter of the circle is $M = 512$. Parameter values $\alpha = 0.0161231$, $r_0 = \rho_0 = \sqrt{\alpha}/2$, $b = 1$, and $\beta = 15.5057$ have been used.




Fig. 7 Maximum absolute errors for values of M equal to 32, 64, 128, 256, 512, and 1024 using the quasifast Hankel transform method. The Fourier transforms all would have value unity at the origin if it were not for the donut hole, so the maximum errors can be regarded as fractional values of the value the transform would have at the origin. We have neglected any errors caused by the donut hole. For this curve, we have adopted Siegman's choice of $r_0 = \sqrt{\alpha}/2$ and $\rho_0 = \sqrt{\alpha}/2$. In addition, $b = 1$ and $\beta = \beta b$.

maximum a few times 10^{-15} . These errors are at least 10 orders of magnitude smaller than the best of those found for the various other numerical methods examined here. As a consequence, it is reasonable to consider both analytic solutions and numerical-integration solutions as the "exact" solutions with which other solutions are compared.

5 Comparison of Shape Errors Due to Inadequate Spectral Sampling Density

Errors that occur at the sampling points are one form of error. However, there is another form that also requires attention, namely, errors in form or shape of the spectrum due to inadequate spectral sampling density. It is possible that errors at the sampling points are acceptably low, but that nonetheless the continuous shape of the spectrum displayed is distorted. This type of error is most important for spectra with a great deal of detailed structure, such as the spectrum of a uniform circle, but is much less noticeable in less complicated spectra, such as the spectrum of a truncated Gaussian function. The appearance of the displayed spectrum is also often dependent of the type of interpolation used by the plotting program. This type of error is more difficult to quantify; in many respects, the acceptability of such errors depends on the observer's criterion for the required accuracy of the shape of the spectrum. This section addresses such errors.

5.1 2-D FFT and Projection-Transform Methods

Since these two methods produce exactly the same calculated transform and exactly the same errors at the sampling points, they can be discussed together. Referring the reader to Fig. 3, the displayed transforms by both methods will be identical. Zeros are not represented well in the calculated transform, due to the fact that the sampling points only occasionally fall close to a zero. However, a more serious problem is observed when the spectrum of Fig. 3(a) is displayed in an expanded view. Figure 8 shows the details of a low-frequency part of the spectrum in such a view. Part (a) of this figure is a plot calculated from a very densely sampled continuous analytic solution for the spectrum. Part (b) shows the interpolated discretely calculated spectrum displayed by the plotting program with linear interpolation between samples. Part (c) shows the same set of sample points this time interpolated with second-order interpolation. Part (d) shows the spectrum when the plotting program uses third-order spline interpolation. The distortions in (b) and (c) are clear. The spline fit is not perfect, but it is much closer to the analytic solution in terms of the shape of the sidelobes. Part (e) shows the magnitude of the spectrum, with error values, when the padding of the aperture is increased to yield $N = 1024$ while M remains 128. Linear interpolation has been used in part (e). The maximum error has not changed appreciably

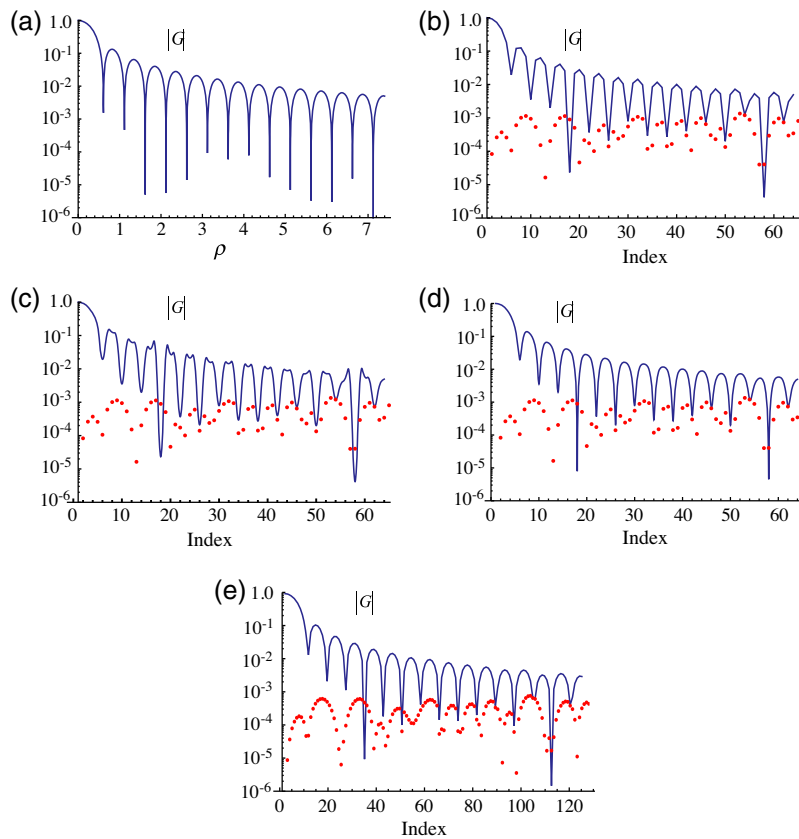


Fig. 8 Interpolated low-frequency portions of spectra. (a) Analytic solution, (b) sampled calculated spectrum with linear interpolation, (c) sampled calculated spectrum with quadratic interpolation, (d) sampled calculated spectrum with third-order spline interpolation. In parts (b)–(d), $M = 128$ and $N = 512$. In part (e), the padding has been increased so that the total number of samples is 1024 while M remains 128, and linear interpolation is used.

between case (b) and case (e), in spite of the more dense sampling of the spectrum, suggesting that the most important source of the errors is the discretization of the circle, rather than aliasing.

5.2 Quasifast Hankel Transform Method

Consider Fig. 6 which shows a plot of the spectrum magnitude obtained by this method when the number of spectral samples is $M = 512$ and linear interpolation is used. The parameters are as specified in Fig. 6. The red dots represent errors at the sample points normalized by the value the spectrum would have at zero frequency. While the errors at the sample points are relatively high, the shapes of the main lobe and the sidelines are good, especially at the low frequencies where the sampling is most dense. Unfortunately, a donut hole exists and the extent of the spectrum calculated by this method is small compared with other techniques. These latter drawbacks are not shared by the other methods.

6 Examples of the Computation Times for the Four Methods

To provide a specific example, simulations have been run for four methods using an 8-core iMac Pro (2017) running macOS Catalina version 10.15.2 and using *Mathematica*, version 12. In all cases but one, the total number of samples in the diameter of the circular function was $M = 256$, and the total number of samples after padding was $N = 1024$. In the case of the quasifast Hankel transform, two results are presented, one for $M = 256$ and one for $M = 1024$, the latter being included so that it is possible to compare the result with those that used padding to increase the number of samples in the spectrum to $N = 1024$. For the quasifast Hankel transform technique,

Table 2 Computation times in examples of four different methods for computing the radial profile of the Fourier transform of a uniform circular function. The 2-D Fourier method, the projection-transform method, and numerical integration method all computed 1024 sample values of the transform, of which 512 were selected for the radial profile. The quasifast Hankel transform method calculated two different cases: $M = 256$ and $M = 1024$. For the case $M = 256$, $\sim 1/4$ of this method's time was used in calculating \hat{j} , but this fraction decreases as M is increased.

Method	Average time (ms)	Standard deviation (ms)
2-D FFT	82.1	3.38
Projection transform	10.67	0.16
Quasifast Hankel ($M = 256$)	41.8	0.50
Quasifast Hankel ($M = 1024$)	64.8	0.43
Numerical integration	792	42

the time to create the arrays \hat{f} and \hat{j} in the x domain from the r domain, the time to perform the three 1-D Fourier transforms, and the time to change the results from the y correlation variable to the ρ frequency variable were included. The times to solve the nonlinear equations to determine the values of the parameters are not included, and in no case was plotting time for the result included. In the case of the projection-transform method, the time to compute the full sequence of projection samples and the time to apply a 1-D FFT to the padded projection sequence were included. Simulations of each method were run 10 times, and the average times and standard deviations were calculated. The time required for a numerical integration of the Hankel transform is also included. Table 2 contains the results.

7 Concluding Remarks

This paper has examined four different methods (including numerical integration) for computing the radial profile of the Fourier transform of a circular symmetric function. The projection-transform method is fastest in speed and has accuracy that is second only to numerical integration. For problems involving Fourier transforms of circularly symmetric functions of any kind, this method is the fastest.

The quasifast Hankel transform is faster than the full 2-D FFT method, but it does not achieve errors comparable with the 2-D FFT or with the projection-transform methods. Note that the times to perform the coordinate transformations have been included in the results of Table 2. This technique has the further disadvantage that there exists a donut hole in the results, which in general is not acceptable, except perhaps in the case of transforms that are zero at the origin. The only real advantage of this method is that it can be used for higher-order Hankel transforms than are needed here, but it should not be the method of choice for circularly symmetric functions.

Extracting a radial profile from a 2-D Fourier transform is third in speed, and numerical integration is the slowest but most accurate method.

Acknowledgments

Disclosures: The author has no relevant financial interests in the manuscript and no other potential conflicts of interest to disclose.

References

1. R. Bracewell, "Strip integration in radio astronomy," *Aust. J. Phys.* **9**, 198–217 (1956).
2. A. Oppenheim, G. Frisk, and D. Martinez, "An algorithm for numerical evaluation of the Hankel transform," *Proc. IEEE* **66**, 264–265 (1978).

3. A. Oppenheim, G. Frisk, and D. Martinez, "Computation of the Hankel transform using projections," *J. Acoust. Soc. Am.* **68**, 523–529 (1980).
4. E. Hansen, "Fast Hankel transform algorithm," *IEEE Trans. Acoust., Speech Signal Process.* **33**, 666–671 (1985).
5. A. Siegman, "Quasifast Hankel transform," *Opt. Lett.* **1**, 13–15 (1977).
6. D. Gardner et al., "Method of analysis of multicomponent exponential decay curves," *J. Chem. Phys.* **31**, 978–979 (1959).
7. G. Agrawal and M. Lax, "End correction to the quasi-fast Hankel transform for optical-propagation problems," *Opt. Lett.* **6**, 171–173 (1981).
8. S. Will and J. Fienup, "An algorithm for exact area-weighted antialiasing of discrete circular apertures," *J. Opt. Soc. Am. A* **37**, 688–696 (2020).

Joseph W. Goodman is the William Ayer Professor Emeritus in electrical engineering at Stanford University. He is a fellow of the SPIE and the OSA, and the author of several books on optics. His research interests cover a broad range of topics in Fourier optics and statistical optics.