

Fast and Accurate Polar Fourier Transform

A. Averbuch* R.R. Coifman† D.L. Donoho‡ M. Elad§ M. Israeli¶

May 22nd, 2005 – Revised Version

Abstract

In a wide range of applied problems of 2-D and 3-D imaging a continuous formulation of the problem places great emphasis on obtaining and manipulating the Fourier transform in Polar coordinates. However, the translation of continuum ideas into practical work with data sampled on a Cartesian grid is problematic. In this article we develop a fast high accuracy Polar FFT. For a given two-dimensional signal of size $N \times N$, the proposed algorithm's complexity is $O(N^2 \log N)$, just like in a Cartesian 2D-FFT. A special feature of our approach is that it involves only 1-D equispaced FFT's and 1-D interpolations. A central tool in our method is the Pseudo-Polar FFT, an FFT where the evaluation frequencies lie in an oversampled set of non-angularly equispaced points. We describe the concept of Pseudo-Polar domain, including fast forward and inverse transforms. For those interested primarily in Polar FFT's, the Pseudo-Polar FFT plays the role of a halfway point – a nearly-Polar system from which conversion to Polar Coordinates uses processes relying purely on 1-D FFT's and interpolation operations. We describe the conversion process, and give an error analysis of it. We compare accuracy results obtained by a Cartesian-based unequally-sampled FFT method to ours, both algorithms using a small-support interpolation and no pre-compensating, and show marked advantage to the use of the Pseudo-Polar initial grid.

Keywords: Polar coordinates, Cartesian coordinates, Pseudo-Polar coordinates, fast Fourier transform, unequally-sampled FFT, interpolation, linogram.

*Department of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel

†Department of mathematics, Yale University, New Haven CT 06520-8283 USA

‡Department of Statistics, Stanford University, Stanford 94305-9025 CA. USA.

§Department of Computer Science, The Technion, Haifa 32000, Israel

¶Department of Computer Science, The Technion, Haifa 32000, Israel.

1 Introduction

1.1 Polar Fourier Transform

Fourier analysis is a fundamental tool in mathematics and mathematical physics, and also in signal and image processing. The discovery, popularization, and digital realization of fast algorithms for Fourier analysis – so called FFT – has had far reaching implications in science and technology in recent decades. The scientific computing community regards the FFT as one of the leading algorithmic achievements of the 20th century [1]. In fact, even ordinary consumer-level applications now involve FFT's – think of web browser decoding JPEG images – so that development of new tools for Fourier analysis of digital data may be of potentially major significance. In this paper we develop tools associated with Fourier analysis in which the set of frequencies is equispaced when viewed in Polar coordinates.

Let $f(x) = f(x_1, x_2)$ be a function on the plane $x = (x_1, x_2) \in \mathbf{R}^2$. Let

$$\hat{f}(\phi) = \int f(x) \exp(-i\phi'x) dx$$

be the usual continuum Fourier transform of f . Writing the frequency $\phi = \{r\cos(\theta), r\sin(\theta)\}$ in Polar coordinates, we let

$$\tilde{f}(r, \theta) = \hat{f}(\phi(r, \theta)).$$

In this paper, the term *Polar Fourier Transform* will always refer to the operation

$$\tilde{f}(r, \theta) = \mathcal{PF}\{f(x)\},$$

namely, getting $f(x)$ in Cartesian variables and computing $\tilde{f}(r, \theta)$ defined with Polar variables.

While changes of variables are, of course, banal *per se*, their significance lies in the change of viewpoint they provide. As it turns out, several fundamental procedures for manipulation of such continuum functions f can best be defined in terms of Polar variables. The Polar FT can be a powerful tool for organizing our understanding of operators and functions on the two dimensional continuum. Much the same can be said of higher dimensions.

1.2 Digital Problematics

The relatively simple and obvious continuum concepts, posing little or no intellectual challenge, correspond to concrete processes operating on sampled digital data, which are important and widely used. Such is the case in image rotation, image registration, tomography, and more. It seems natural to ask if the conceptual simplicity and clarity of the continuum Polar FT can be used in some way to assist, improve, or simplify practical procedures for digital processing, which appear as digitally sampled realizations of rotation, registration, radon transformation, and so on.

This leads naturally to the question of whether there could be a Polar FT for discrete data, particularly a fast algorithm, or *Polar FFT*. Such a hypothetical entity should have many of the properties

of the continuum Polar FT, including relations to rotation, registration, Radon transform, and so on, and yet would be definable and rapidly computed for digital data in the now ubiquitous equispaced Cartesian format.

The prevailing belief seems to be that there is no such algorithm. For example, in Briggs' treatise *The FFT: an Owner's Manual for the Discrete Fourier Transform* [2], which is widely considered comprehensive and authoritative, the index contains the entry "Polar FFT", continuing with "no FFT for, 284"! Indeed, several difficulties stand in our way to construct such a fast Polar FT for discrete data: (i) the need to define the appropriate Polar grid in the frequency domain; (ii) finding a fast way to evaluate (or approximate) the FT on this grid points; and (iii) enabling a stable transform inversion.

1.3 This Paper's Contribution

In this paper we propose a notion of a Polar FT which is well suited for digital data – a procedure which is faithful to the continuum Polar FT concept, highly accurate¹, fast, and generally applicable. As in Figure 1, we define the Polar grid of frequencies $\xi_{p,q} = \{\xi_x[p, q], \xi_y[p, q]\}$ in the circle inscribed in the fundamental region $\xi \in [-\pi, \pi)^2$, and, given digital Cartesian data $f[i_1, i_2]$ we define the Polar FT to be the collection of samples $\{F(\xi_{p,q})\}$, where $F(\xi_{p,q})$ is the trigonometric polynomial

$$F(\xi_{p,q}) = \sum_{i_1=0}^{N-1} \sum_{i_2=0}^{N-1} f[i_1, i_2] \exp(-i(i_1 \xi_x[p, q] - i_2 \xi_y[p, q])).$$

The sampled frequency points are given by

$$\left\{ \left[\xi_x[p, q] = \frac{\pi p}{N} \cos(\pi q/2N), \quad \xi_y[p, q] = \frac{\pi p}{N} \sin(\pi q/2N) \right] \text{ for } -N \leq p \leq N-1, 0 \leq q \leq 2N-1 \right\}.$$

Thinking of the Polar Discrete Fourier Transform (PDFT) mapping $\mathcal{PDFT} : f[i_1, i_2] \rightarrow F(\xi_{p,q})$ as a linear operator, we also consider a generalized inverse procedure of it, going back from discrete Polar Fourier data to Cartesian spatial data.

In this paper we carefully define the Polar FT concept for digital data, the associated fast algorithms, and discuss its features such as accuracy and computational complexity. We will make special effort to describe five advantages of the proposed PFFT, *speed*, *accuracy*, *stability*, *vectorizability* and *non-expansivity*. These terms will be given meaning as we deepen our description. A brief and partial version of this work appeared in [3], and here we give an expanded and more detailed description of this work.

1.4 Relation to State of the Art

Two existing bodies of literature contain ideas relevant to the above definition of Polar FT – one corresponds to methods to compute the inverse Radon transform, and the other treating the general problem of evaluating the Fourier transform on non-equally spaced frequency points.

¹This implies that the proposed algorithm gives an approximated evaluation of the transform with high accuracy.

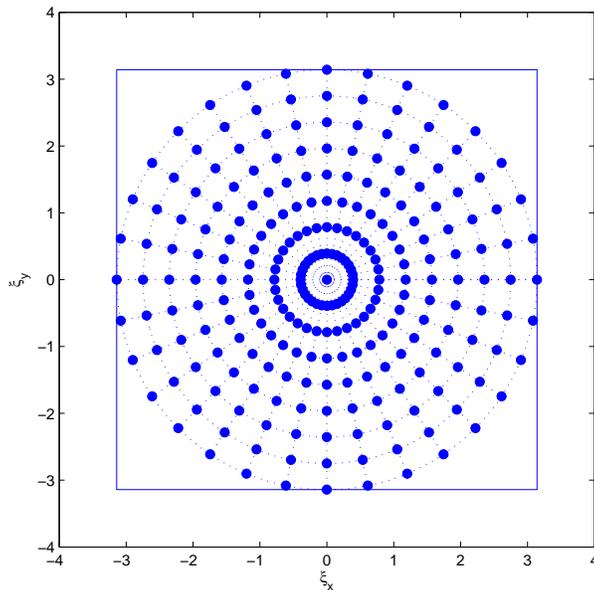


Figure 1: Polar grid - intersection of 8 concentric circles and 16 angularly equispaced rays.

The literature of tomography concerns reconstruction of an image from a collection of its one-dimensional integrals. An important approach to this problem is the direct Fourier method [4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16]. The direct Fourier method uses the projection-slice theorem, which suggests that one can convert a collection of projection data $(Rf)(\cdot, \theta)$ into a two-dimensional Fourier data $\hat{f}(\xi)$ on a Polar grid, and then reconstruct by Fourier inversion.

In the second body of literature [17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29] one has data at an equally spaced Cartesian grid, but wishes to evaluate its discrete Fourier transform as a non-equispaced non Cartesian set of frequencies. This problem is known as either the USFFT (short for Unequally Spaced FFT), or the NUFFT (Non-Uniform FFT) problem.

It is not hard to see the direct relevance of the above literature to the definition we have given. The Polar sampling set defined above is certainly a non-equispaced, non-Cartesian set. Thus, the problem of computing $F(\xi_{p,q})$ from digital data $f[i_1, i_2]$ is explicitly a problem of evaluating the Fourier transform of f at unequally spaced frequencies. In other words, the problem of computing the PDFT is a special case of the USFFT problem. Second, note that, if we take the projection data $(Rf)(\cdot, \theta)$ and perform a discrete Fourier transform in the (discretely sampled) t -variable, we effectively have $\hat{f}(\xi_{p,q})$. Hence, the problem of Radon inversion is closely connected to the problem of going *back* from discrete Polar Fourier data $\hat{f}(\xi_{p,q})$ to Cartesian spatial data $f[i_1, i_2]$.

The existing state-of-the-art expressed by the above literature can be adapted to produce an approximated Polar FFT (PFFT) and its inverse. Adapting ideas from the USFFT literature, the forward Polar-FFT can be accomplished by the following stages: (i) pre-multiplication of $f[i_1, i_2]$ by pre-computed weights $s[i_1, i_2]$; (ii) computation of the 2D-FFT on an oversampled Cartesian grid; and (iii) interpolation of the desired values on the Polar grid based on those given on the Cartesian one. Reversing

these operations can provide the adjoint operator. The inversion of the transform can be accomplished, for example, by Least-Squares fitting, as proposed in [28]. This approach is similar to ideas available in the direct Fourier reconstruction literature mentioned above.

The approach we propose here for PFFT is in fact similar to the above, using a special type of USFFT method, where a different launching grid is used, rather than the Cartesian one. Our method factors the problem into two steps: first, a *Pseudo-Polar* FFT is applied, in which a Pseudo-Polar sampling set is used, and second, a conversion from Pseudo-Polar to Polar FT is performed.

At the heart of the method proposed here for the PFFT we use the Pseudo-Polar FFT – an FFT where the evaluation frequencies lie in an oversampled set of non-angularly equispaced points (see Figure 2). The Pseudo-Polar FFT offers us a near-Polar frequency coordinate system for which exact rapid FT evaluation is possible [30]. Whereas the Polar grid points sit at the intersection between linearly growing concentric circles and angularly equispaced rays, the Pseudo-Polar points sit at the intersection between linearly growing concentric squares and a specific choice of angularly non-equispaced rays.

The Pseudo-Polar transform plays the role of a halfway point – a nearly-Polar system from which conversion to Polar Coordinates uses processes relying purely on 1-D FFT’s and interpolation operations. We present this conversion process, along with an error analysis of it, showing far improved performance compared to the Cartesian-based USFFT-based alternative approach. The reason for the expected improvement in accuracy is the ability of the Pseudo-Polar grid to provide a space-varying sampling of the frequency domain, which is close in density to the Polar destination one. More specifically, the Pseudo-Polar grid gives a denser sampling near the origin, enabling better interpolation performance. In comparison, starting with a Cartesian grid, getting similar density (and thus accuracy) near the origin requires an increased oversampling all-over the frequency domain. Note also that for most signals, their energy resides near the origin, and thus, better treatment there implies higher overall accuracy.

As mentioned before, the above described approach enjoys the following important advantages:

- *Complexity*: For a given two-dimensional signal of size $N \times N$, the complexity of both the proposed Polar FFT and its inverse is of order $N^2 \log N$, just like in a Cartesian 2D-FFT operating on the same input array.
- *Accuracy*: Since polynomial interpolations are involved in the algorithm’s path, exact results cannot be claimed. However, as will be shown later, the applied interpolations are expected to yield highly accurate results due to specific features of the frequency domain function slices. The accuracy obtained is substantially better than the one expected with Cartesian-based USFFT methods. We should note that this claim is demonstrated in our work for short-support interpolation schemes that provide an overall simple algorithm. Recent work by Fessler and Sutton [28] have clearly shown a marked improvement in the obtained accuracy with Cartesian USFFT when using large support Kaiser-Bessel interpolation, and pre-compensation. Further work is required to use those options with initial Pseudo-Polar grid and explore ways to further improve our method.

- *Stability*: This property refers mostly to our ability to invert the transform, and claim almost one-to-one mapping between spatial array and its inverse PFFT result. Two ingredients are responsible for such a behavior - adequate conditioning of the signal, prior to its transform, and a proper set of Polar frequency samples that would be considered representative.
- *Vectorizability*: The process we propose is organized as a series of purely one-dimensional signal manipulation processes, in which data form a single row or column of an array are subjected to a series of Fourier transforms, 1D interpolations, and ancillary operations. Owing to the use of vectorized operations, the overall algorithm is "cache-friendly"².
- *Non-expansivity*: There is no drastic oversampling of the underlying array. As we show later on, a similar accuracy by a Cartesian-based USFFT method requires much higher oversampling.

1.5 Content

This paper is organized as follows: Section 2 describes the Pseudo-Polar Fourier grid, and the forward/inverse transforms over this grid. In this section we closely follow the work in [30]. Section 3 then discusses the conversion from Pseudo-Polar – to –Polar for the forward transform, and from Polar – to – Pseudo-Polar in the inverse one. We show how these conversions amount to one-dimensional operations, and discuss reasons for the high interpolation accuracy obtained. Section 4 analyzes the proposed algorithm from several aspects. Accuracy is studied by experimenting on images, and by finding the worst-case scenarios maximizing the approximation error. Section 5 gives a brief description of a freely available software performing the proposed PFFT transform. This software also includes the code to reproduce this paper’s results. Section 6 concludes this paper, with discussion on future work and open questions.

2 Pseudo-Polar Fourier Transform

The Pseudo-Polar Fourier transform is based on a definition of a Polar-like 2D grid that enables fast Fourier computation. This grid has been explored by many since the 1970-s. The pioneers in this field are Mersereau and Oppenheim [31] who proposed the concentric squares grid as an alternative to the Polar grid. Work by Pasciak [32], Edholm and Herman [33], and Lawton [34] showed that fast exact evaluation on such grids is possible. Later work by Munson and others [8, 10] have shown how these ideas can be extended and used for tomography. Recently, the Pseudo-Polar grid was proposed as the base for a stable forward and inverse Radon transform called *Fast Slant-Stack* [30, 35]. It has also been used for image registration [36].

²We should note that Cartesian-based USFFT methods may also be made cache-efficient in various ways [25]. However, this property is far more natural when the data is structured as in the Pseudo-Polar coordinates.

As we shall see next, in this work we strongly build on this notion of near-Polar grid, and exploit the fast and stable forward and inverse transforms with the frequency domain sampled with such coordinate system. In this section we cover the basics of the Pseudo-Polar grid and its use for forward and inverse transform.

Before we start we remind the reader of the Cartesian grid 2-D DFT to set notations. Letting the horizontal and vertical frequencies

$$\left\{ \xi_x = \frac{2\pi k_1}{N}, \xi_y = \frac{2\pi k_2}{N} \right\}_{k_1, k_2=0}^{N-1},$$

we get the familiar DFT in 2 dimensions by

$$\hat{f}(\xi_x, \xi_y) = \sum_{i_1=0}^{N-1} \sum_{i_2=0}^{N-1} f[i_1, i_2] \cdot \exp(-i(i_1 \xi_x + i_2 \xi_y)). \quad (1)$$

2.1 The Pseudo-Polar Coordinate System

We start by defining the Pseudo-Polar grid points in the frequency domain. These points are separated into two groups – the Basically Vertical (BV) and the Basically Horizontal (BH) subsets, given by

$$BV = \left\{ \xi_y = \frac{\pi \ell}{N} \text{ for } -N \leq \ell < N, \xi_x = \xi_y \cdot \frac{2m}{N} \text{ for } -\frac{N}{2} \leq m < \frac{N}{2} \right\} \quad (2)$$

and

$$BH = \left\{ \xi_x = \frac{\pi \ell}{N} \text{ for } -N \leq \ell < N, \xi_y = \xi_x \cdot \frac{2m}{N} \text{ for } -\frac{N}{2} < m \leq \frac{N}{2} \right\} \quad (3)$$

Figure 2 depicts this grid; BV points are marked with the filled disks and BH ones are marked as circles. Several properties about this grid should be mentioned:

- The grid points are at intersections of linearly growing concentric squares with angularly non-equispaced rays. The squares' sides are of size $\pi k/N$, $k = 0, 1, \dots, N$. The BH rays have equispaced slope: $2k/N$, $k = -N/2 + 1, -N/2 + 2, \dots, N/2$. The BV rays are similar but with clockwise rotation of 90° .
- The suggested grid is Polar-like. The main two differences are the concentric squares replacing concentric circles, and equispaced rays in slope replaced by equispaced rays in angle. This resemblance will be exploited towards the development of a true Polar-FFT algorithm.
- Referring to the BV (and similarly to the BH) points, we can see that they are organized on lines with angles ranging from $-\pi$ to π . Along each such line, the points are spread uniformly, though in an angle-dependent way. These two properties are the driving force for the ability to compute the Fourier transform in a fast way for this grid.
- When looking at Figure 2 one might get the impression that some points on the outer-most square are missing. However, since $\hat{f}(\xi_x, \xi_y)$ is periodic with a period of 2π in both axes, these points are actually redundant.

- In computing and storing the Fourier transform for these points, our data structure is given by two simple 2D arrays – one for the BV and the other for the BH sample sets. Referring to the BV array, the vertical axis corresponds to the index ℓ and the horizontal to m (see (2) and (3)). Thus, our array has $2N$ rows and N columns. Overall we have $4N^2$ frequency sample points, originating from an image of size $N \times N$. The factor 4 oversampling is helpful for numerical stability [30]. This data structure implies that when we draw a row or column from these arrays, we do not necessarily refer to horizontal or vertical set of frequency points, but rather refer to points along one ray or points along one of the concentric squares.
- In our treatment, the processing of the BV and the BH data are completely parallel. Our description will refer to the BV points only.

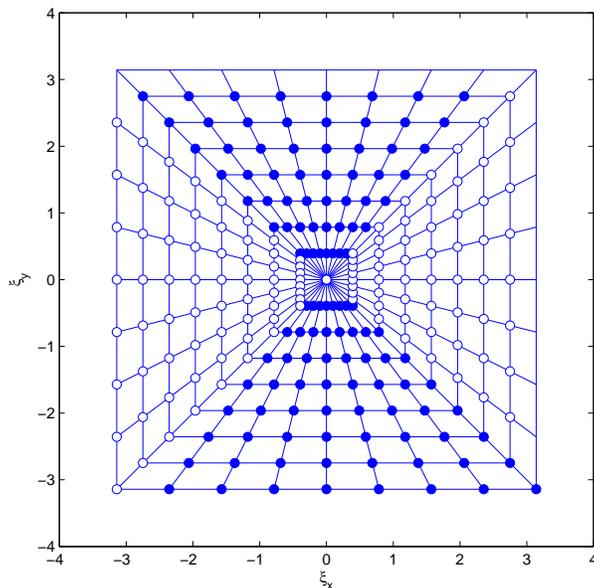


Figure 2: The Pseudo-Polar grid and its separation into BV (filled circles) and BH (empty circles) coordinates ($N = 8$) - intersection of 8 concentric squares and 16 slope-equispaced rays.

2.2 A Fast Forward Transform

Given the Pseudo-Polar BV grid points in the frequency domain, we are interested in computing the Fourier transform values. In what follows we shall show that simple 1D-FFT operations can be used to achieve this goal.

Theorem 1 *Given a 2D signal $f[i_1, i_2]$, $0 \leq i_1, i_2 < N$, evaluation of the FT on the Pseudo-Polar grid as defined in (2) and (3) can be done by 1D operations only, and with complexity of $140N^2 \log N$ flops.*

Proof: This result can be considered known, applying ideas behind the work of Pasciak [32], Edholm-Herman [33], and Lawton [34] to the Pseudo-Polar grid, which, however, is not precisely the grid they

considered. Essentially, the result as stated has been obtained in [30], and is spelled-out here for completeness.

Using the definition of the Fourier transform for discrete functions in Equation (1), and plugging the coordinates from Equation (2) and (3) we obtain

$$\begin{aligned}
\hat{f}(\xi_x, \xi_y) = \hat{f}[m, \ell] &= \sum_{i_1=0}^{N-1} \sum_{i_2=0}^{N-1} f[i_1, i_2] \cdot \exp(-i(i_1\xi_x + i_2\xi_y)) \\
&= \sum_{i_1=0}^{N-1} \sum_{i_2=0}^{N-1} f[i_1, i_2] \cdot \exp\left(-i\left(\frac{2\pi i_1 \ell m}{N^2} + \frac{\pi i_2 \ell}{N}\right)\right) \\
&= \sum_{i_1=0}^{N-1} \exp\left(-\frac{i2\pi i_1 \ell m}{N^2}\right) \sum_{i_2=0}^{N-1} f[i_1, i_2] \cdot \exp\left(-\frac{i\pi i_2 \ell}{N}\right).
\end{aligned} \tag{4}$$

Concentrating on the inner summation part, we define

$$\hat{f}_1[i_1, \ell] = \sum_{i_2=0}^{N-1} f[i_1, i_2] \cdot \exp\left(-\frac{i\pi i_2 \ell}{N}\right). \tag{5}$$

This expression stands for a 1D-FFT on the columns of the zero padded array $f[i_1, i_2]$. In order to show this, assume that $f[i_1, i_2]$ is zero padded to yield

$$f_Z[i_1, i_2] = \begin{cases} f[i_1, i_2] & 0 \leq i_2 < N \\ 0 & N \leq i_2 < 2N \end{cases}$$

Thus,

$$\hat{f}_1[i_1, \ell] = \sum_{i_2=0}^{N-1} f[i_1, i_2] \cdot \exp\left(-\frac{i\pi i_2 \ell}{N}\right) = \sum_{i_2=0}^{2N-1} f_Z[i_1, i_2] \cdot \exp\left(-\frac{i2\pi i_2 \ell}{2N}\right). \tag{6}$$

This expression stands for 1D-FFT of an array of $2N$ samples. However, another complicating factor is the range of ℓ ($-N \leq \ell \leq N-1$), as opposed to the regular FFT range that starts at 0. Thus, defining $n = \ell + N$ we get

$$\begin{aligned}
\hat{f}_1[i_1, n] &= \sum_{i_2=0}^{2N-1} f_Z[i_1, i_2] \cdot \exp\left(-\frac{i2\pi i_2(n-N)}{2N}\right) \\
&= \sum_{i_2=0}^{2N-1} f_Z[i_1, i_2] \cdot (-1)^{i_2} \cdot \exp\left(-\frac{i2\pi i_2 n}{2N}\right).
\end{aligned} \tag{7}$$

To summarize this part, computation of the inner summation (w.r.t. i_2) amounts to 1D-FFT of length $2N$ on the columns of the zero padded and pre-multiplied³ array $f[i_1, i_2] \cdot (-1)^{i_2}$. The amount of operations needed for this stage are $5 \cdot 2N \log(2N)$ per each column, and $10N^2 \log(2N)$ operations for all the N columns.

³We have to multiply our array by $(-1)^{i_2}$, but this could be replaced by later shift, owing to the periodic nature of the transform.

We return now to Equation (4), and assume that the above process has been completed. Thus we hold the $N \times 2N$ array $\hat{f}_1[i_1, \ell]$, and we should proceed by computing the second summation

$$\hat{f}[m, \ell] = \sum_{i_1=0}^{N-1} \hat{f}_1[i_1, \ell] \exp\left(-\frac{i2\pi i_1 m \ell}{N^2}\right) = \sum_{i_1=0}^{N-1} \hat{f}_1[i_1, \ell] \exp\left(-\frac{i2\pi i_1 m}{N} \cdot \frac{\ell}{N}\right). \quad (8)$$

Removal of the factor $\alpha = \ell/N$ in the exponent turns this expression into a regular 1D-FFT, this time applied on the rows of the array \hat{f}_1 . With this factor α in the above summation, the required operation is known as the Chirp-Z [37], or the Fractional Fourier Transform (FRFT) [38]. In Appendix A we show how this transform can be computed efficiently for any α , with $30N \log N$ operations, based on 1D-FFT use. As before, due to a shifted range of the index m , one has either to shift after the transform, or modulate the array, prior to the transform.

To summarize, we have $2N$ rows that are to be put through a Chirp-Z transform. Thus we need $60N^2 \log N$ operations for this stage. Adding to the previous stage complexity, we get an overall of $70N^2 \log N$ operations for the computation of the transform for the BV part of the grid. To cover both the BV and the BH grid points we need $140N^2 \log N$. \square

In this paper the Pseudo-Polar grid is a stepping stone towards the Polar coordinates system. Since interpolations are to be used to convert from Pseudo-Polar to Polar coordinates, we should consider computing the Pseudo-Polar FFT on a more densely-spaced grid. We state here without proof the following result: Given a 2D signal $f[i_1, i_2]$, $0 \leq i_1, i_2 < N$, the evaluation of the FT on the oversampled Pseudo-Polar grid with NS concentric squares and $2NP$ slope-equispaced rays can be obtained by 1D vector operations only, and with complexity of $120N^2 PS \log(NS)$ flops. In fact, this result could be shown to follow immediately from Theorem 1 with zero-padding of the original image.

2.3 Fast Inverse Transform and Quasi-Parseval Relationship

The Pseudo-Polar FFT can be inverted by the method of Least-Squares (LS) (see [28, 30] for the same concept). To see this, consider a matrix-vector formulation of our processes. For an image $f[i_1, i_2]$ of size $N \times N$, we define the column vector \underline{f} of length N^2 , containing the image pixel values in column-stack ordering. This vector is multiplied by a matrix $\mathbf{T}_{PP} \in \mathcal{C}^{4N^2 \times N^2}$, representing the Pseudo-Polar FFT. The outcome of this multiplication is a vector $\underline{\hat{f}}$ of length $4N^2$, containing the samples of the Fourier transform on the Pseudo-Polar grid. Given $\underline{\hat{f}}$, inversion of the Pseudo-Polar Fourier transform is achieved by solving

$$\underline{f} = \text{Arg min}_{\underline{x}} \left\| \mathbf{T}_{PP} \underline{x} - \underline{\hat{f}} \right\|_2^2 = (\mathbf{T}_{PP}^H \mathbf{T}_{PP})^{-1} \mathbf{T}_{PP}^H \underline{\hat{f}} = \mathbf{T}_{PP}^+ \underline{\hat{f}}. \quad (9)$$

where \mathbf{T}_{PP}^+ denotes generalized inverse.

Of course, the matrix-based solution is useless as a computational approach for reasonable sizes of input arrays; a Pseudo direct inversion of the matrix \mathbf{T}_{PP} is computationally prohibitive, and definitely beyond the desired $N^2 \log(N)$ complexity. Instead, we approach the optimization problem iteratively

by the iteration relation

$$\underline{f}_{k+1} = \underline{f}_k - \mathbf{D} \mathbf{T}_{PP}^H \left(\mathbf{T}_{PP} \underline{f}_k - \hat{\underline{f}} \right). \quad (10)$$

The expression $\mathbf{T}_{PP}^H \left(\mathbf{T}_{PP} \underline{x} - \hat{\underline{f}} \right)$ is the function's gradient, and the multiplication by a positive definite matrix \mathbf{D} guarantees descent in the LS error. If chosen properly, \mathbf{D} could speed-up convergence of this algorithm to the true solution, as posed in (9). In [30] a specific choice of diagonal matrix \mathbf{D} is proposed that down-weights near-origin points in order to equalize the condition number. It is shown that with few (2 – 6) iterations this iterative process achieves high accuracy solutions. As to the initialization, it could be chosen as zeros for simplicity. This way we obtain a fast inverse transform of the same complexity as the forward one, as every iteration requires the application of both the forward transform (multiplying with \mathbf{T}_{PP}) and its adjoint (multiplying with \mathbf{T}_{PP}^H). It still remains to be seen that the adjoint is computable with the same complexity as the forward transform, which is the claim of the next Theorem.

Theorem 2 *Given a 2D array $\hat{f}[m, \ell]$ of size $2N \times 2N$ representing a Pseudo-Polar grid sampling in the frequency domain, the evaluation of the adjoint Pseudo-Polar FFT to produce an $N \times N$ image can be done by 1D operations only, and with complexity of $O\{N^2 \log(N)\}$ operations.*

Proof: Posed in a matrix notation, multiplication by \mathbf{T}_{PP}^H requires taking the conjugate of each element in the matrix \mathbf{T}_{PP} , and summing with respect to columns, rather than rows. Thus, using the relation posed in (4) we have that the regular (forward) transform is applied by

$$\hat{f}[m, \ell] = \sum_{i_1=0}^{N-1} \sum_{i_2=0}^{N-1} f[i_1, i_2] \cdot \exp \left(-i \left(\frac{2\pi i_1 \ell m}{N^2} + \frac{\pi i_2 \ell}{N} \right) \right).$$

Similarly, referring to the basically vertical (BV) values only, the adjoint is achieved by

$$\begin{aligned} \tilde{f}[i_1, i_2] &= \sum_{m=-N/2}^{N/2-1} \sum_{\ell=-N}^{N-1} \hat{f}[m, \ell] \cdot \exp \left(i \left(\frac{2\pi i_1 \ell m}{N^2} + \frac{\pi i_2 \ell}{N} \right) \right) \\ &= \sum_{\ell=-N}^{N-1} \exp \left(i \frac{\pi i_2 \ell}{N} \right) \sum_{m=-N/2}^{N/2-1} \hat{f}[m, \ell] \cdot \exp \left(i \frac{2\pi i_1 \ell m}{N^2} \right). \end{aligned} \quad (11)$$

The inner summation could be written as

$$\tilde{f}[i_1, \ell] = \sum_{m=-N/2}^{N/2-1} \hat{f}[m, \ell] \cdot \exp \left(i \frac{2\pi i_1 m}{N} \cdot \frac{\ell}{N} \right),$$

and this is a Fractional FFT, just as we obtained with the forward transform. We have seen that this operation could be done in $O\{N \log(N)\}$ per every ℓ , producing N values. Thus, for $-N \leq \ell \leq N - 1$ we need to perform $O\{N^2 \log(N)\}$ operations. Once performed, we then obtain the expression

$$\tilde{f}[i_1, i_2] = \sum_{\ell=-N}^{N-1} \exp \left(i \frac{\pi i_2 \ell}{N} \right) \tilde{f}[i_1, \ell], \quad (12)$$

and this is a regular Inverse-FFT, which requires $O\{N \log(N)\}$ per every $0 \leq i_1 \leq N - 1$. Thus, again we get that $O\{N^2 \log(N)\}$ operations are required to conclude this part of the computations.

As a last point in this proof, we should refer similarly to the BH rays. The complete adjoint operation is obtained by performing the same process as described above, and adding the two resulting arrays. \square

3 From Pseudo-Polar to Polar

Similar to the Cartesian-based USFFT approach, we suggest to compute the Polar-FT values based on a different grid for which a fast algorithm exists, and then go to the Polar coordinates via an interpolation stage. However, instead of using the Cartesian grid in the first stage, we use the Pseudo-Polar grid of the previous section. Since this grid is closer to the Polar destination coordinates, there is a reason to believe that this approach will lead to better accuracy and thus lower oversampling requirements. However, as we shall see next, beyond the proximity of the Pseudo-Polar coordinates to the Polar ones, another very important benefit is the ability to perform the necessary interpolations via pure 1D operations without losing accuracy. This property is vital in understanding the superiority of the proposed scheme over traditional Cartesian-based USFFT methods.

3.1 Pseudo-Polar – Polar: Grid Conversion

We define the Polar coordinate system based on the Pseudo-Polar one, with manipulations that lay out the necessary interpolation stages discussed later on. Starting from the basically-vertical frequency sampling points in the Pseudo-Polar grid as given in (2)

$$BV = \left\{ \xi_y = \frac{\pi \ell}{N} \text{ for } -N \leq \ell < N, \xi_x = \frac{2\pi m \ell}{N^2} \text{ for } -\frac{N}{2} \leq m < \frac{N}{2} \right\},$$

the Polar ones are obtained by two operations:

1. **Rotate the Rays:** In order to obtain an angularly-uniform ray sampling as in the Polar coordinate system, the rays must be rotated. This is done by replacing the term $2m/N$ in ξ_x above with $\tan(\pi m/2N)$, leading to the grid points

$$BV_{new} = \left\{ \xi_y = \frac{\pi \ell}{N} \text{ for } -N \leq \ell < N, \xi_x = \frac{\pi \ell}{N} \cdot \tan\left(\frac{\pi m}{2N}\right) \text{ for } -\frac{N}{2} \leq m < \frac{N}{2} \right\}.$$

The result is a set of points organized on concentric squares as before, but the rays are now equispaced in angle rather than slope. Figure 3 depicts this step as an interpolation stage. Rotating the rays amounts to a 1D operation along horizontal lines (for the BV points). A set of N equispaced points along this line are replaced by a new set of N points along the same line in different locations (marked as small squares) implementing equispaced angular sampling.

An interesting property of this interpolation stage is the fact that the underlying function to be interpolated assumes a simple form, which implies that interpolation accuracy is expected to be

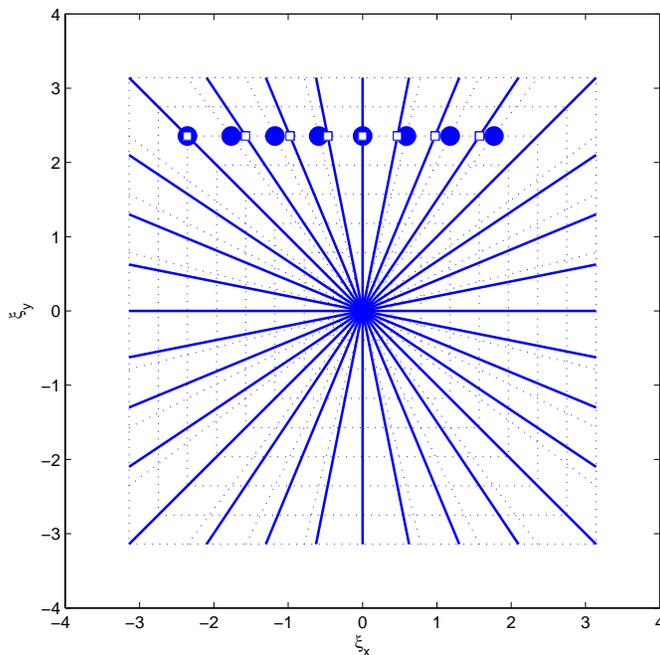


Figure 3: The first interpolation stage for rotating the rays. Circles represent the known Pseudo-Polar grid points, and squares are the desired equi-angularly spaced rays.

high, even for low oversampling factors. Furthermore, the given points are densely arranged around the desired locations, thus ensuring a better interpolation accuracy (compared to the same number of points spread from end-to-end, as in the Cartesian USFFT).

Returning to Equation (4) and referring to a specific row with fixed ξ_y , we have

$$\begin{aligned}
 \hat{f}(m, \xi_y) &= \sum_{i_1=0}^{N-1} \sum_{i_2=0}^{N-1} f[i_1, i_2] \cdot \exp\left(-i \left[\frac{2mi_1}{N} + i_2 \right] \xi_y\right) \\
 &= \sum_{i_1=0}^{N-1} \left[\sum_{i_2=0}^{N-1} \exp(-i\xi_y i_2) f[i_1, i_2] \right] \cdot \exp\left(-i \frac{2mi_1}{N} \xi_y\right) \\
 &= \sum_{i_1=0}^{N-1} C[i_1] \cdot \left[\exp\left(-i \frac{2i_1}{N} \xi_y\right) \right]^m
 \end{aligned} \tag{13}$$

and this is a complex trigonometric polynomial of order N .

One implication of this observation is that given N samples of this function, all information about this function is given. Theoretically, this implies that a perfect interpolation is possible if for every destination points all the N given samples are used. From a more practical point of view, the above observation implies that this function is relatively smooth and with a moderate oversampling, a near-perfect interpolation is expected for a small neighborhood operation.

2. **Circle the Squares:** In order to obtain concentric circles as required in the Polar coordinate system, we need to ‘circle the squares’. This is done by dividing both ξ_x and ξ_y by a constant

along each ray, based on its angle, and therefore a function of the parameter m , being

$$R[m] = \sqrt{1 + \tan^2\left(\frac{\pi m}{2N}\right)} \quad (14)$$

The resulting grid is given by

$$\left\{ \begin{array}{l} \xi_y = \frac{\pi \ell}{NR[m]} \quad \text{for } -N \leq \ell < N \\ \xi_x = \frac{\pi \ell}{NR[m]} \cdot \tan\left(\frac{\pi m}{2N}\right) \quad \text{for } -\frac{N}{2} \leq m < \frac{N}{2} \end{array} \right\}.$$

Figure 4 depicts this step as an interpolation stage. Circling the squares amounts to a 1D operation along rays. A set of $2N$ equispaced points is replaced by a new set of $2N$ points along the same line in different locations (marked as small squares). This time the destination points are equispaced, but with a different spacing.

We have seen above that the first interpolation stage is applied on a trigonometric polynomial, which explains the expected accuracy obtained. In this later stage, the function over which we interpolate is not of the same simple form, but nevertheless smooth enough. Referring to one ray of slope α we have that the 1D slice to work on is given by

$$\begin{aligned} \hat{f}(\xi_x, \xi_y) &= \sum_{i_1=0}^{N-1} \sum_{i_2=0}^{N-1} f[i_1, i_2] \cdot \exp(-i(i_1 \xi_x + i_2 \xi_y)) \\ &= \sum_{i_1=0}^{N-1} \sum_{i_2=0}^{N-1} f[i_1, i_2] \cdot \exp(-i(i_1 + \alpha i_2) \xi_x) = \tilde{f}(\xi_x) \end{aligned}$$

Consider the function $\tilde{f}(z)$ described above as a 1D function and take its Fourier transform

$$\begin{aligned} \mathcal{F}\{\tilde{f}(z)\} &= \int_{-\infty}^{\infty} \tilde{f}(z) \exp(-i\omega z) dz \\ &= \int_{-\infty}^{\infty} \sum_{i_1=0}^{N-1} \sum_{i_2=0}^{N-1} f[i_1, i_2] \cdot \exp(-i(i_1 + \alpha i_2)z) \exp(-i\omega z) dz \\ &= \sum_{i_1=0}^{N-1} \sum_{i_2=0}^{N-1} f[i_1, i_2] \cdot \delta(i_1 + \alpha i_2 + \omega). \end{aligned}$$

This expression assumes its maximal frequency at $i_1 = i_2 = N - 1$ and its minimal frequency at $i_1 = i_2 = 0$. The frequency support of this function is the interval $[-(N - 1)(1 + \alpha), 0]$. Thus, for $\alpha = 1$ we obtain the widest bandwidth of $2N$, and then the Nyquist-sampling rate for this function is $T_{max} = \pi/N$.

This means that the Fourier transform \hat{f} restricted to the ray is a band-limited function with required maximal sampling period of π/N . If this function is sampled in this rate along the entire ray (from $-\infty$ to ∞) we have a complete representation of it that enables knowledge of its values at any location. In our case, for a limited interval $z \in [-\pi, \pi]$ we have $2N$ samples, which is the critical sampling rate exactly. Thus, with a moderate oversampling we may expect to represent

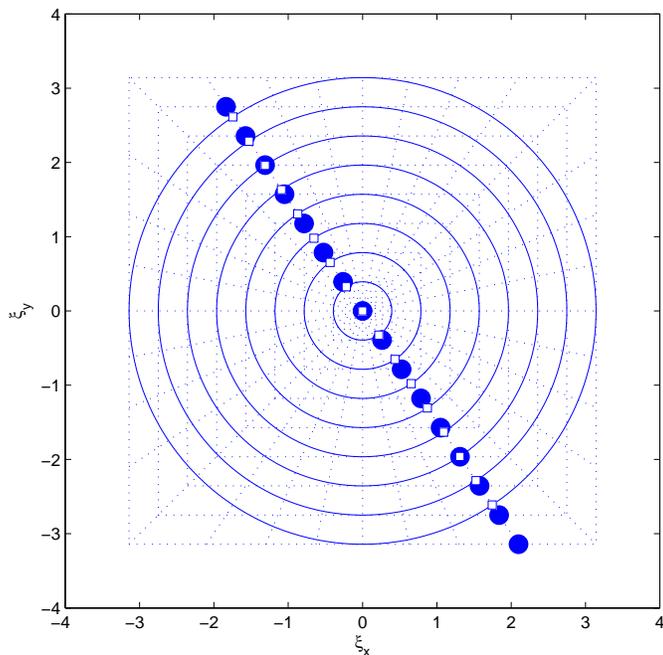


Figure 4: Second Interpolation stage for circling the squares. Small circular points represent the known grid points after the previous interpolation stage, and the square points are the desired final interpolated values.

this function very accurately. This implies that this function is relatively smooth as well and lends itself to high-accuracy interpolation.

In performing the Polar-FFT, we start with computation of the Fourier transform over a Pseudo-Polar grid, followed by the two interpolation stages presented above. As we have seen, in order to obtain high accuracy, the Pseudo-Polar grid should be oversampled, both radially and angularly. One approach to oversampling is the use of the over-sampled Pseudo-Polar grid as presented in the previous section (e.g., by zero padding). Alternatively, as the Pseudo-Polar fast transform can be broken into two parts, this separation could be used for better efficiency. For the BV coordinates, the vertical FFT can be done with oversampling by zero padding as before, and then, as we go to the second phase of Fractional FFT per each row, we can apply this stage one row at a time with the interpolation that leads directly to the rotated rays. The memory savings are substantial (directly proportional to the oversampling factor P). Note that in performing the 1D interpolations, complete rows or columns from the result array are brought to the cache, and processed to produce the desired values. Thus, memory management becomes extremely effective, as opposed to the regular USFFT methods that require general memory access that may lead to many cache misses.

As for the complexity of the overall algorithm, we have seen before that an amount of $120N^2SP \log_2(NS)$ operations are required for the computation of the oversampled Pseudo-Polar FFT. Those operations are followed by the first interpolation requiring $O(N^2S)$ operations (every row uses NP points to compute N

values requires $O(N)$ operations, and there are NS of those). The second interpolation requires $O(N^2)$ operations (every ray with $2NS$ values is used to produce $2N$ new values, and there are N rays). Thus, the overall operation count is dominated by the $120N^2SP \log_2(NS)$ operations of the Pseudo-Polar FFT.

In terms of memory requirements, the overcomplete Pseudo-Polar FFT requires $4N^2SP$ float-values, and once those are allocated, all other operations can be done within this array. Alternatively, using row-wise interpolation, only $4N^2S$ values are required. Also, a factor 2 saving can be obtained if a proper separation between the basically vertical and basically horizontal parts is applied. Further savings could be obtained by breaking each of these groups (BV and BH) to sub-parts with small overlap.

All the above applies to a general image of size $N \times N$, resulting with an output Polar FT array of size $2N \times 2N$. If higher oversampling is desired in the destination grid, all the above is still applicable. One simple way to get such oversampling is to zero-pad the input array (this way modifying the initial N).

The inversion of the Polar-FFT requires an implementation of the inverse Pseudo-Polar FFT as described in the previous section. Starting with a set of Polar coordinates in the frequency domain, we first apply two interpolation stages - *squaring the circles*, and then *rotating the rays* - just as described above but in reversed order. For such operation to succeed, we assume that the Polar grid is given in an over-complete manner (i.e., we start with an array of $2NP[rays] \times 2NS[circles]$ and interpolate to a Pseudo-Polar grid with $2N[rays] \times 2N[squares]$). Once those coordinates are filled with values, inverse Pseudo-Polar FFT is applied as described previously, using the adjoint operator. The same arguments posed above explain the suitability of the 1D interpolation and their expected accuracy in the above reverse process.

3.2 Disk Band-Limited Support As Regularization

A major difficulty in using the Polar coordinate system with digitally sampled data arises from the square shape of the fundamental frequency domain $[-\pi, +\pi]^2$. In our Polar FFT the corners are not represented, and this may lead to some loss of information. From an algebraic point of view, one may say that the matrix representing the Polar FFT we have defined is not invertible, or badly conditioned.

A way around this problem is to assume that there is no content in those non-sampled corners. This means that the image should be supported in the frequency domain inside a disk of radius π , leaving the corners empty. If the given image is sampled at $\sqrt{2}$ times the Nyquist rate on both Cartesian axes then its frequency domain support is the square $[-\pi/\sqrt{2}, +\pi/\sqrt{2}]^2$, and this square is contained in the disk required. Thus, given any image, we can easily verify that this condition holds: zero-padding the image by factor $\sqrt{2}$ in both axes, and applying 2D-FFT followed by 2D-IFFT, the result image is twice as big (in pixel count) and its frequency domain support is as required. All these operations add $O(N^2 \log_2(N))$ flops to the Polar FFT that follows.

In the inversion transform from the frequency Polar-coordinates to the spatial domain, the knowledge about the corners having zero content is valuable, and could be used to further stabilize the inverse transform. This is done in the interpolation to the Pseudo-Polar grid, by assigning zero values to all grid-points outside the π -radius circle⁴.

Figure 5 presents the effect of the corner-nulling process proposed here on the condition number of the transform matrix. For an input array of size $N \times N$ (with N assuming the values 4, 6, 8, ..., 30) we compose the transform matrix for the $4N^2$ Polar coordinates and compute its condition number. As can be seen from the graph, this value grows exponentially. As an example, for a moderate size of 30×30 input array, the condition number is 2×10^9 , implying that inversion by Least-Squares will require numerous iterations, and will be highly sensitive to numerical errors.

The second curve refers to the same transform matrix with an augmented part referring to the corners. The augmentation part is obtained by generating a regular $4N \times 4N$ Cartesian grid transform matrix, and choosing the rows that refer to the corners (outside the π -radius disk). Roughly speaking, we have doubled the number of rows, and the condition-number this time is well controlled. For an input array of size 30×30 the condition number is 6.6, meaning that standard iterative techniques of linear algebra achieve a high accuracy inversion.

This numerical experiment comes to show that zero-forcing at the corners stabilizes the inversion process. However, in practice, there is no need for the augmentation as done in the prescribed experiment - all that is needed is nulling values outside the frequency support disk of radius π , and the improved condition-number presented here holds true nevertheless.

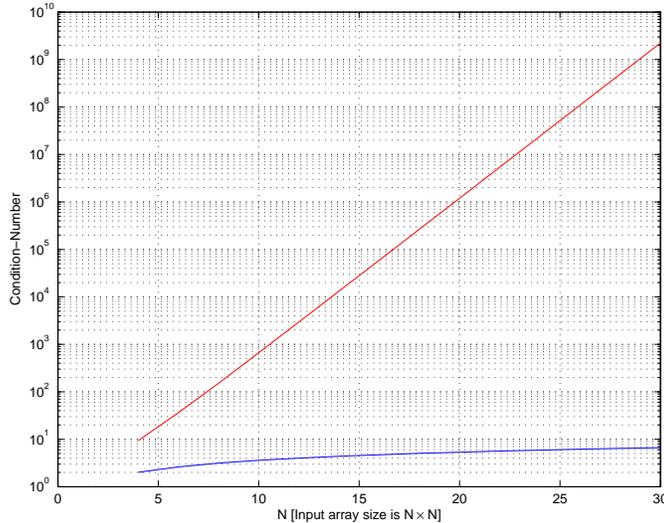


Figure 5: The condition-number for the transform matrix for a direct Polar-coordinates transform, and for an augmented one, that nulls the corners.

⁴Note that this discussion is not to be confused with the need to apply preconditioning in the inversion of the Pseudo-Polar transform.

4 Analysis of the Polar FFT Algorithm

4.1 Experimental Evidence - USFFT versus Polar-FFT

We apply the USFFT and the Polar-FFT tools to the image *Lena* (scaled down to size 64×64 pixels, in order to speedup this and later experiments), and compare error behaviors. In particular, we are interested in studying the effect of the oversampling. Figure 6 presents the results of the first experiment.

For the Cartesian-based USFFT method, we use a 2D Hermite interpolation on a 2×2 support. It exploits known Fourier values on the Cartesian grid, along with their horizontal and vertical derivatives (altogether, 12 values), in order to fit a polynomial. Those directional derivatives are obtained by pre-multiplying the image by linear terms and applying the 2D-FFT three times [21]. For the Pseudo-Polar based method, the first interpolation stage is also based on Hermite interpolation with a 2-samples support (as for the USFFT, but applied in 1D, and exploiting 4 known values on this support to fit a cubic). The second stage interpolation uses a piecewise cubic spline using the Fourier values alone, and forcing alignment of the first derivatives⁵. Thus, the interpolation schemes deployed on our Polar-FFT are inferior to the one practiced for the Cartesian-USFFT. Those schemes remain fixed throughout the reported experiments. As already mentioned in the introduction, better interpolation methods can be used for both algorithms. Here we chose to concentrate on short-support and relatively simple methods that are easy to understand and implement, and ones that eventually lead to an acceptable level of accuracy, nevertheless.

Figure 6 presents both ℓ^1 and ℓ^2 measures of error. The error presented is a relative error, normalized by the appropriate norm of the ideal result. The horizontal axis in the plots is the overall oversampling ratio. For the Cartesian-based USFFT, for an oversampling factor of S in each axis, the overall oversampling ratio is S^2 . Similarly, for our method, the overall oversampling is PS , where S and P are the oversampling ratios for the number of squares and number of rays. Several choices of P and S are tested.

We see that for the proper choice of P and S ($S \approx 5P$) the Polar-FFT method is far more accurate than the USFFT, under either error measure. For a fixed P there is a saturation effect in the approximation error as $S \rightarrow \infty$ because errors caused by the first interpolation stage dominate and cannot be compensated. The higher needed oversampling along the rays can be explained in two ways: (i) the radial interpolation stage of the PFFT uses a lower-order procedure, based on splines instead of Hermite interpolation; and (ii) the radial interpolation task is more difficult as the signal to be interpolated is less smooth (see previous Section). Extensive experiments with different signals, their sizes, and metrics of error evaluations, all confirm that the above conclusions comparing the two methods are typical.

⁵Evaluating the directional derivatives for this stage requires a replication of the first stage interpolation for those as well, and thus was avoided, settling for a weaker scheme.

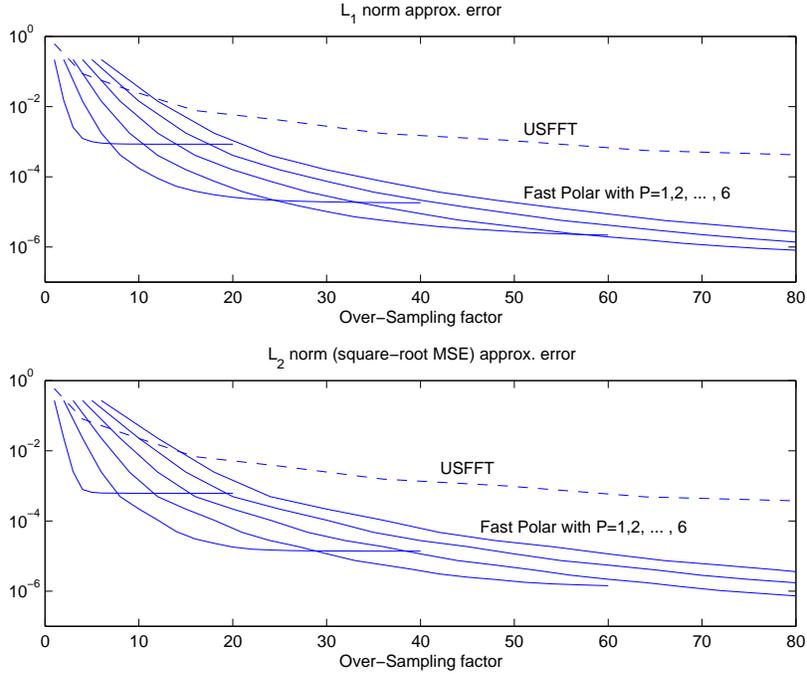


Figure 6: USFFT and the Polar-FFT error as a function of the over-sampling. The USFFT oversampling is the same along the two axes, while for the PFFT we fix the angular oversampling ($P = 1, 2, \dots, 6$, and thus 6 curves) and vary the radial one $S = 1, 2, \dots, 20$.

4.2 Frequency Domain Spread of the Error

Given the specific image and fixing the oversampling ratio, we study how the errors of the two methods behave in the frequency domain. Intuitively, we know that Cartesian-based USFFT method should cause higher errors near the origin where the interpolation starts from an inherently coarser grid. In contrast, we expect the Polar-FFT to perform equally well in most frequency locations.

Using the same test image from the previous experiment (scaled down *Lena*), Figure 7 shows the relative errors obtained by the USFFT ($S = 9$ in both axes), and the Polar-FFT ($S = 20$, $P = 4$). The plots show log (with base 10, being crucial for a good visualization) of the absolute value of the transform errors, divided by the mean absolute value of the true transform, so as to get a relative error evaluation. The frequency domain is presented on Cartesian axes in order to give more intuitive frequency-content description. Since the frequency domain is sampled in Polar coordinates in the transforms, we use a Voronoi diagram to slice the plane into pieces per each Polar frequency sample. This is a matter of visual display only.

Our expectations about the distribution of errors in the frequency domain are validated - the Cartesian-USFFT concentrates the error at the origin, while the Polar-FFT errors are spread around. Moreover, the errors obtained by the Polar-FFT are substantially smaller.

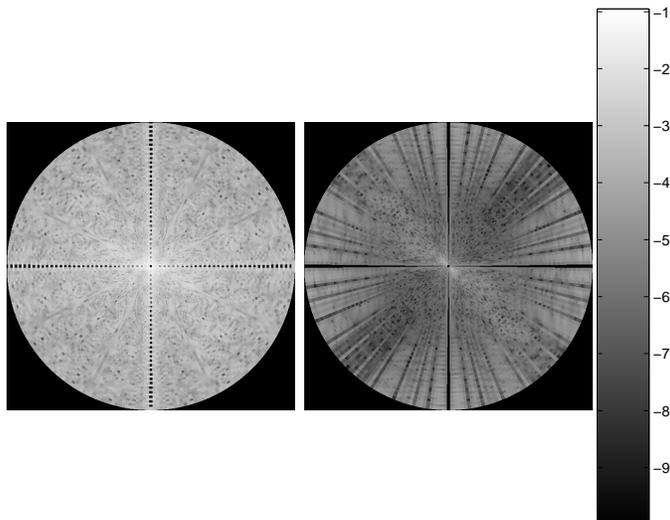


Figure 7: The relative error as a function of frequency location using a specific image. This Figure shows the error obtained by the USFFT method (left), and the PFFT algorithm errors (right). The error scale is relative and logarithmic.

4.3 Frequency Domain Spread of the Error - Worst-Case Analysis

A limitation of the above experiment is that it relies on a specific choice of a signal. We now address the same question from a worst-case point of view: for a specific frequency location in the destination Polar grid, what is the worst possible signal, maximizing the error at this point, and how big is this maximum error? Answering these questions per each location, we can draw a worst-case error map in the frequency domain and compare for both the USFFT and the Polar-FFT methods.

Let x of size $N \times N$ denote the signal to be transformed. The transform result has $2N$ rays containing $2N$ samples each. As all the involved operations are linear, we adopt a matrix-vector representation for the exact (\mathbf{T}_e), the USFFT (\mathbf{T}_u) and the Polar-FFT (\mathbf{T}_p). All these matrices are of size $4N^2 \times N^2$. For the given signal \underline{x} of length N^2 (due to its lexicographic ordering), the transform error for the USFFT is $(\mathbf{T}_e - \mathbf{T}_u)\underline{x}$, with a similar expression for the Polar-FFT method. To maximize this error is to solve

$$\max_{\underline{x}} \frac{\|\mathbf{W}(\mathbf{T}_e - \mathbf{T}_u)\underline{x}\|_2^2}{\|\underline{x}\|_2^2}. \quad (15)$$

\mathbf{W} is a diagonal matrix with all entries being zero apart from the one element chosen in the frequency domain. Thus, essentially, only one row of the matrix $\mathbf{T}_e - \mathbf{T}_u$ is used - denote this as \underline{e}^T . Clearly, by Cauchy's inequality, the maximum is obtained for $\underline{x} = \underline{e}$, and the error obtained is $\|\underline{e}\|_2$.

Figure 8 presents these worst-case errors⁶ as a function of frequency for both methods when $N = 16$.

⁶In this and later experiments based on the matrix-vector representation of the transforms, we use low value for N because of the induced matrix sizes. One could use much higher N values if instead of explicitly forming the matrices, the transform and its adjoint operator are applied.

The Polar-FFT performs far better in all locations. Note that our relatively low choice of $N = 16$ in this experiment causes special sampling effects.

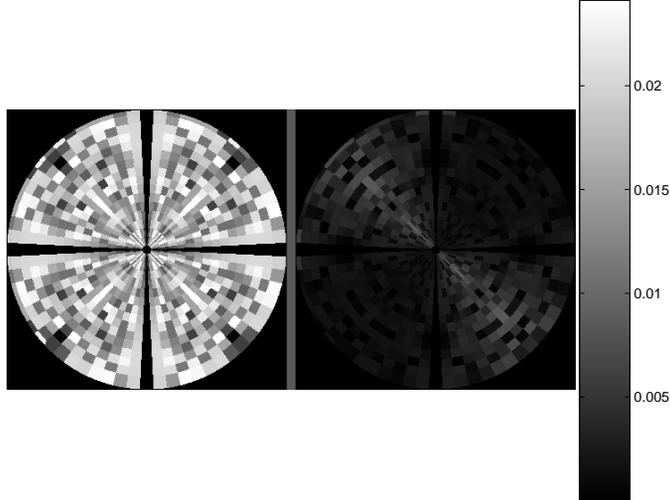


Figure 8: The error as a function of frequency location: using the worst case signal per location.

4.4 Worst Case Error via Eigenspace Analysis

4.4.1 Direct Approach

we desire A plot OF the worst case ℓ^2 error of each method under a fixed oversampling. Using matrix-vector notation, for a given signal \underline{x} , the transform error of the USFFT is $(\mathbf{T}_e - \mathbf{T}_u)\underline{x}$. Now we solve the optimization problem

$$\max_{\underline{x}} \frac{\|(\mathbf{T}_e - \mathbf{T}_u)\underline{x}\|_2^2}{\|\underline{x}\|_2^2}, \quad (16)$$

seeking the worst-possible signal \underline{x} maximizing error, subject to unit ℓ^2 -norm. The answer is of-course the first right singular vector of $(\mathbf{T}_e - \mathbf{T}_u)$ [39], and the value of (16) is the square of the first singular value (a square-root is required to align with previous error measures).

Figure 9 presents the real and imaginary parts of these worst-case signals for the USFFT ($S = 9$) and the Polar-FFT ($S = 20, P = 4$), and the absolute frequency description of this signal. The USFFT's maximal square-root MSE error is 3.6×10^{-3} while the Polar-FFT's is 1.9×10^{-4} . Again, the USFFT method is weaker, and its worst signal is concentrated near the frequency origin where the method is weakest. Note that the worst signal is modulated (notice the shift from the center in the spatial domain) to result in a very non-smooth frequency behavior.

4.4.2 Relative Approach

A major problem with the above analysis is the difficulty in understanding the meaning of the error found, being a ratio between energies in the frequency and the spatial domains. An interesting alternative

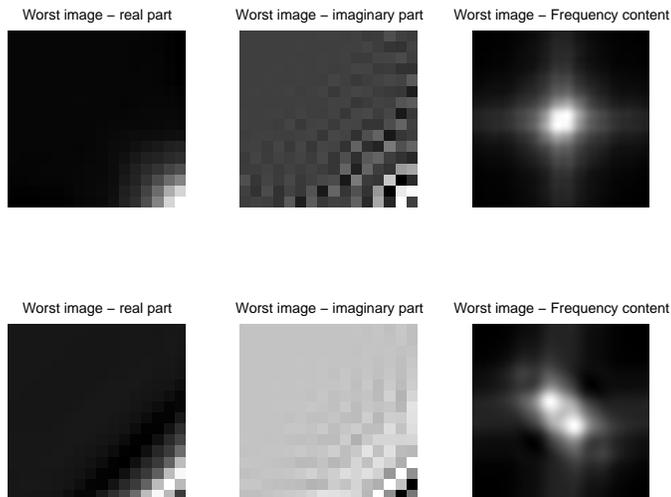


Figure 9: Worst case signal - Direct eigenvalue approach. The top row shows the worst signal for the USFFT - the real spatial part (left), the imaginary spatial part (middle), and its spectrum (right). The bottom row shows the same for the PFFT.

is the definition of worst signals by

$$\max_{\underline{x}} \frac{\|(\mathbf{T}_e - \mathbf{T}_u)\underline{x}\|_2^2}{\|\mathbf{T}_e\underline{x}\|_2^2}. \quad (17)$$

Put in words, we seek the worst error relative to transform size in Polar frequency coordinates. This problem amounts to a generalized eigenvalue problem [39]. Figure 10 presents the results for Cartesian USFFT method (maximal error is 3.6×10^{-4}) and Polar-FFT method (worst error is 4.5×10^{-5}). Both methods lead to similar worst-case signals, with energy mostly falling outside the circle of radius π , so that the denominator in the above definition is nearly zero. Nevertheless, our Polar-FFT has an advantages over Cartesian coordinates.

4.4.3 Relative Approach with Support Constraint

In order to avoid signals of the sort obtained earlier, we seek the worst case signal for (17) with the side-constraint of no energy in the frequency domain outside the circle of radius π . Thus we solve

$$\max_{\underline{x}} \frac{\|(\mathbf{T}_e - \mathbf{T}_u)\underline{x}\|_2^2}{\|\mathbf{T}_e\underline{x}\|_2^2} \quad \text{subject to} \quad \mathbf{F}_1\underline{x} = 0, \quad (18)$$

where \mathbf{F}_1 represents the regular Cartesian FFT in a predetermined density, restricted to frequency points outside the circle. Alternatively, we reformulate this using

$$\max_{\{\underline{x} \mid \mathbf{F}_1\underline{x}=0\}} \frac{\|(\mathbf{T}_e - \mathbf{T}_u)\underline{x}\|_2^2}{\|\mathbf{T}_e\underline{x}\|_2^2} \approx \max_{\underline{x}} \frac{\|(\mathbf{T}_e - \mathbf{T}_u)\underline{x}\|_2^2}{\|\mathbf{T}_e\underline{x}\|_2^2 + \lambda \|\mathbf{F}_1\underline{x}\|_2^2}. \quad (19)$$

With this formulation we again have a generalized eigenvalue problem. For large λ , (19) yields an approximate solution of the original problem (18). We chose $\lambda = 1000$ and find that higher values does

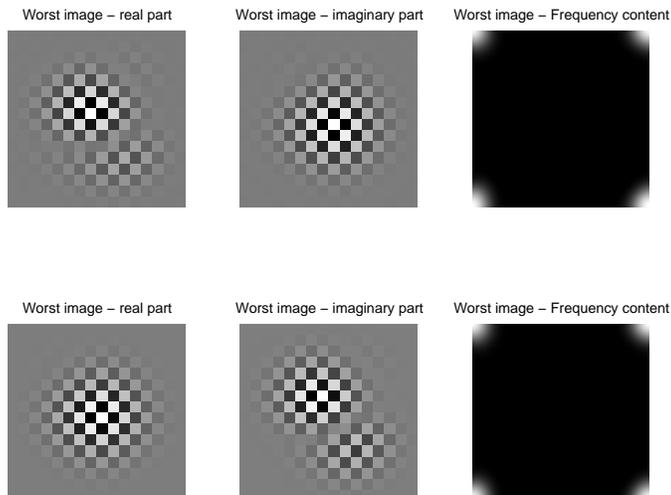


Figure 10: Worst case signal - relative eigenvalue approach. The top row shows the worst signal for the USFFT - the real spatial part (left), the imaginary spatial part (middle), and its spectrum (right). The bottom row shows the same for the PFFT.

not change the results perceptibly. Those results are shown in Figure 11. The USFFT worst error is 6.5×10^{-5} and with the Polar-FFT method the error we obtain is 4.2×10^{-6} .

4.5 Signal-Space Ordering via Eigenspace Analysis

We return to the definition of error in Equation (17). In solving the maximization problem for the Polar-FFT we find the worst signal and the accompanying worst error. This worst signal represents only a 1-dimensional subspace of signals, and for signals orthogonal to it, all we can say is that the error on their transform computation is expected to be smaller, though we cannot say by how much.

Solving the same problem again while forcing the result to be orthonormal to the previous result, we obtain a second worst rank-1 sub-space. Repeating this process, we essentially find ranked orthonormal basis that represents the signal-space $\mathbf{C}^{N \times N}$, such that the first vector spans the worst signals, the second spans the “next” worst signals, and so on. The process described above is simply an eigenvalue problem for the matrix

$$(\mathbf{T}_e - \mathbf{T}_u)^H (\mathbf{T}_e - \mathbf{T}_u).$$

The eigenvectors are the ranked ortho-basis, and the eigenvalues are related squared errors. Denote the results of the above process on the Polar-FFT error by $\{\underline{u}_k\}_{k=1}^{k=N^2}$, such that \underline{u}_{N^2} is the worst signal. Then the sequence

$$\{\lambda_k\}_{k=1}^{N^2} = \left\{ \frac{\|(\mathbf{T}_e - \mathbf{T}_p) \underline{u}_k\|_2^2}{\|\underline{u}_k\|_2^2} \right\}_{k=1}^{N^2} \quad (20)$$

measures the errors induced by each subspace, from ranked best to the worst. Using the same ortho-basis,

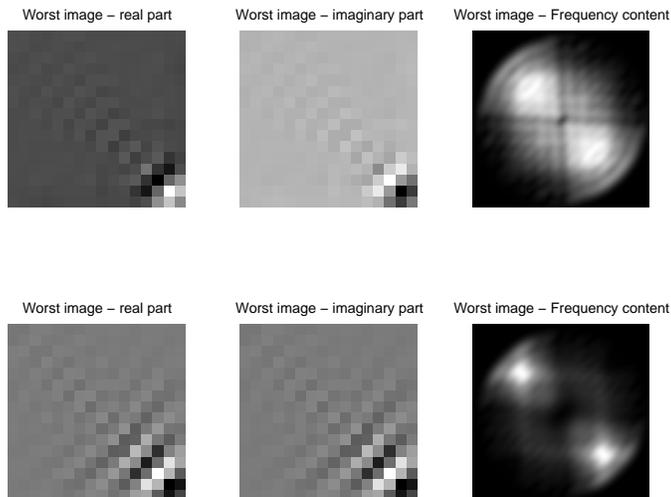


Figure 11: Worst case signal - relative eigenvalue approach with constraint. The top row shows the worst signal for the USFFT - the real spatial part (left), the imaginary spatial part (middle), and its spectrum (right). The bottom row shows the same for the PFFT.

we may compute

$$\{\delta_k\}_{k=1}^{N^2} = \left\{ \frac{\|(\mathbf{T}_e - \mathbf{T}_u)\mathbf{u}_k\|_2^2}{\|\mathbf{u}_k\|_2^2} \right\}_{k=1}^{N^2}, \quad (21)$$

the errors we will effectively obtain from the Cartesian USFFT method on the same subspaces. Plotting these two sequences in Figure 12, we compare the distribution of errors by subspace and obtain a complete picture about the transform errors of the two methods for all the signal space.

Figure 13 is similar but uses (18) to obtain the ortho-basis and the errors for its subspaces. In this case the eigenvalue problem becomes a generalized eigenvalue one. Similarly, Figure 14 presents results from the constrained generalized eigenvalue problem (19). All these figures show a consistent superiority (and even a growing one, as we turn from the worst-case direction to the remaining subspace) of our Polar-FFT over the USFFT method.

5 Available Software and Reproducible Results

One of the main objectives of this project has been development of a software tool that performs the computations given here. We intend to make this toolbox available to the public, in a fashion parallel to `Wavelab` and `BeamLab` libraries [40, 41]. Our implementation uses Matlab code to perform various tasks around the idea of Polar FFT. The PFFT toolbox is freely available in <http://www.cs.technion.ac.il/~elad/PolarFFT.zip>.

As part of the freely available software, we supply code reproducing every Figure in this paper that is based on computation. We believe that this will help researchers entering this area, and engineers

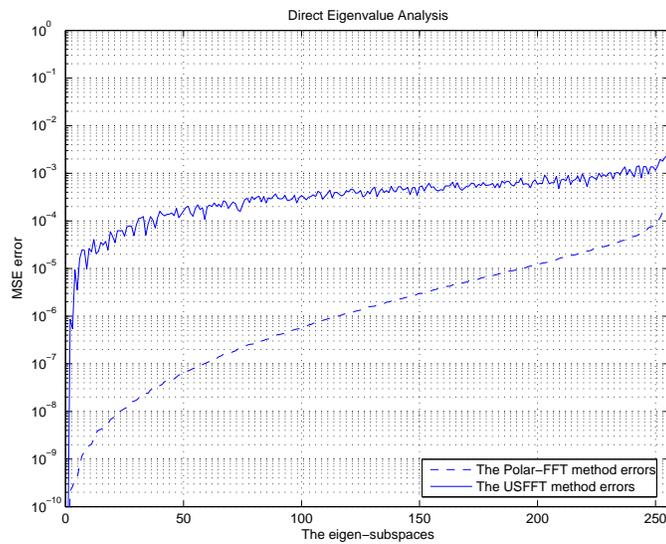


Figure 12: The ranked signal space - the Direct eigenvalue approach.

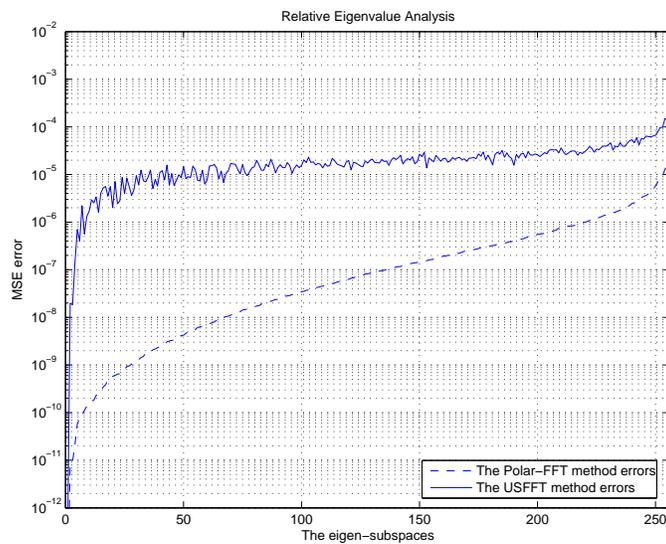


Figure 13: The ranked signal space - the relative eigenvalue approach.

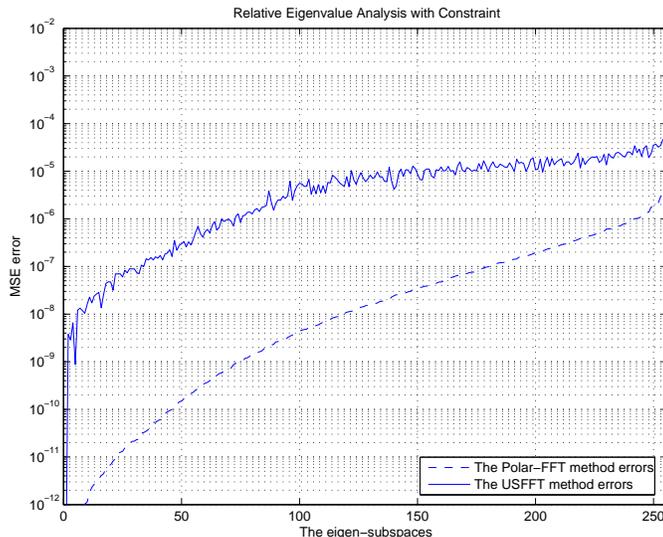


Figure 14: The ranked signal space - the relative eigenvalue approach with constraints.

interested in better understanding the details behind the verbal descriptions supplied here. This also agrees with the recent emerging standard of guaranteeing reproducible results in scientific publications [40].

6 Conclusions

In this article we have developed a rapid and accurate Polar FFT. For a given two-dimensional signal of size $N \times N$, the proposed algorithm produces a Polar FFT with $C \times N^2$ coordinate samples, where C is a chosen parameter (e.g., $C = 4$). The complexity of this algorithm is $O(CN^2 \log N)$, just like in a Cartesian 2D-FFT. Two special features of this algorithm are (i) it involves only 1-D equispaced FFT's and 1D interpolations, leading to a highly cache-aware algorithm; and (ii) it has very high accuracy for moderate oversampling factors.

The accuracy of the algorithm was studied both for specific images and from a worst-case standpoint. The new approach is shown to be far more accurate than Cartesian-based Unequally-Spaced Fast Fourier Transform (USFFT). Our analysis shows that the proposed scheme produces highly accurate transform values near the origin, where “typical” signals concentrate on. However, these comparative results are restricted to the use of simple and small-support interpolations and no pre-compensation. Further work is required in order to incorporate more advanced techniques from the USFFT literature.

In the present paper we emphasize the basic construction, leaving aside applications, improvements, and extensions. We believe that this paper serves the researchers/engineers interested in applying those tools in applications, offering a reasonable confidence in the accuracy and speed. The supplied software, along with this documentation, may make the adoption of these methods relatively painless.

Future work on this topic might consider various improvements to the interpolation stages imple-

mented in our algorithm, using recent results in the USFFT literature (e.g. [28, 29]). The proposed algorithm might have applications in stable and rapid inversion of Radon transform, fast image rotation and registration, and elsewhere. Extension to 3D will require definition of Polar-like⁷ coordinates.

Appendix A - Chirp-Z Transform (FRFT)

Definition 3 Given the discrete signal $x[n]$ over the support $0 \leq n < N$, and given an arbitrary scalar value α , the Chirp-Z Transform (also known as the Fractional Fourier Transform) is given by

$$\hat{x}[k] = \sum_{n=0}^{N-1} x[n] \exp\left(-\frac{i2\pi kn}{N} \cdot \alpha\right) \quad \text{for } 0 \leq k < N. \quad (\text{A-1})$$

As we have seen through the definition of the Pseudo-Polar Fourier transform, we need to perform this kind of transform as part of our process. The following theorem guarantees existence of a highly efficient algorithm for this transform. This theorem, and more importantly, the efficient process itself, are described in [37, 38].

Theorem 4 Evaluation of the Chirp-Z transform can be done with order of $30N \log N$ operations.

Proof: Our proof will be given by construction of such an efficient process. Our starting point is the trivial assignment $2kn = k^2 + n^2 - (k - n)^2$. Using this relationship into the definition above we obtain

$$\begin{aligned} \hat{x}[k] &= \sum_{n=0}^{N-1} x[n] \exp\left(-\frac{i\pi\alpha}{N} (k^2 + n^2 - (k - n)^2)\right) \\ &= \exp\left(-\frac{i\pi\alpha}{N} k^2\right) \sum_{n=0}^{N-1} x[n] \cdot \exp\left(-\frac{i\pi\alpha}{N} n^2\right) \cdot \exp\left(\frac{i\pi\alpha}{N} (k - n)^2\right). \end{aligned} \quad (\text{A-2})$$

If we define the sequence

$$s[n] = \exp\left(-\frac{i\pi\alpha}{N} n^2\right),$$

then the above equation can be rewritten as

$$\hat{x}[k] = s[k] \sum_{n=0}^{N-1} x[n] s[n] s[k - n]. \quad (\text{A-3})$$

This equation suggests a way to evaluate the transform, starting with pre-multiplication of $x[n]$ by $s[n]$, followed by a convolution with $s[n]$, and finally, post-multiplying with $s[n]$ again.

The main saving in computation comes from replacing convolution by a 1D-FFT based method. A 1D-FFT transform should be applied on the two sequences, multiply the results, and apply inverse transform on the resulting sequence. FFT should be computed for the two sequences with zero padding to length $2N$ in order to avoid cyclic effects.

Three 1D-FFT's are needed, costing each $10N \log N$ operations. The creation of $s[n]$ and the pre/post-multiplication by it require $O(N)$ additional operations. Thus, an overall of $30N \log N$ operations is needed for the entire process, as claimed by the theorem. \square

⁷There is no exact Polar sampling method for 3D, thus we must settle for Polar-like grids.

References

- [1] Dongarra, J. and Sullivan, F. (2000) Introduction the top 10 algorithms, Special issue of *Computing in Science and Engineering*, Vol. 2, No. 1.
- [2] Briggs, B. and Henson, E. (1995) *The FFT: an Owner's Manual for the Discrete Fourier Transform*, SIAM, Philadelphia.
- [3] Averbuch, A., Coifman, R.R., Donoho, D.L., Elad, M. and Israeli, M. (2003) Accurate and fast Polar ourier transform, The 37th Asilomar on Signals, Systems and Computers, Pacific Grove, CA. November.
- [4] Natterer, F. (1986) *The Mathematics of Computerized Tomography*. John Wiley & Sons, New York.
- [5] Stark, H., Woods, J.W., Paul, I. and Hingorani, R. (1981) Direct Fourier reconstruction in computer tomography, *IEEE Trans. on Acoustics, Speech and Signal Proc.* Vol. 29, No. 2, pp. 237-245.
- [6] Fan, H. and Sanz, J.L.C. (1985) Comments on 'Direct fourier reconstruction in computer tomography', *IEEE Trans. on Acoustics, Speech and Signal Proc.*, Vol 33, No. 2, pp. 446-449.
- [7] Matej, S. and Vajtersic, M. (1990) Parallel implementation of the direct Fourier reconstruction method in tomography, *Computers and Art. Intel.*, Vol. 9, No. 4, pp. 379-393.
- [8] Moraski, K. J. and Munson Jr., D.C. (1991) Fast tomographic reconstruction using chirp-z interpolation, in *Proceedings of 25th Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, pp. 1052-1056.
- [9] Delaney, A.H. and Bresler, Y. (1996) A fast and accurate Fourier algorithm for iterative parallel-beam tomography, *IEEE Trans. Image Processing*, pp. 740-753, May.
- [10] Choi, H. and Munson Jr., D.C., (1998) Direct-Fourier reconstruction in tomography and synthetic aperture radar, *Int. J. of Imaging Systems and Tech.*, Vol.9, No.1, pp. 1-13.
- [11] Walden, J. (2000) Analysis of the direct Fourier method for computer tomography, *IEEE Trans. on Medical Imaging*, Vol. 19, No.3, pp. 211-22.
- [12] Gottlieb, D., Gustafsson, B., and Forssen, P. (2000) On the direct Fourier method for computer tomography, *IEEE Transactions on Medical Imaging*, Vol. 19, No. 3, pp. 223-232.
- [13] Natterer, F. and Wübbeling, F (2000) *Mathematical Methods in Image Reconstruction*. SIAM, Philadelphia.
- [14] Basu, S. and Bresler, Y. (2000) An $O(n^2 \log n)$ filtered backprojection reconstruction algorithm for tomography, *IEEE Trans. Image Processing*, Vol. 9, No. 10, pp. 1760-1773.
- [15] Fourmont, K. (2003) Non equispaced fast Fourier transforms with applications to tomography, *J. Fourier Anal. Appl.*, Vol. 9, pp. 431-450.

- [16] Matej, S., Fessler, J.A. and Kazantsev, I.G. (2004) Iterative tomographic image reconstruction using Fourier-based forward and back-projectors, *IEEE Trans. on Medical Imaging*, Vol. 23, No. 4, pp. 401–412, April.
- [17] Suli, E. and Ware, A. (1991) A spectral method of characteristics for hyperbolic problems *SIAM J. Numer. Anal.*, Vol. 28, No. 2, pp. 423–445.
- [18] Boyd, J.P. (1992) Multipole expansions and Pseudospectral cardinal functions: A new generalization of the Fast Fourier Transform, *J. Comput. Phys.*, Vol. 103 No. 2, pp. 184–186.
- [19] Boyd, J.P. (1992) A fast algorithm for Chebyshev, Fourier, and sinc interpolation onto an irregular grid, *J. of Comp. Phys.*, Vol. 103, No. 2, pp.243–257.
- [20] Beylkin, G. (1995) On the fast Fourier transform of functions with singularities, *Applied Computational Harmonics*, Vol. 2, pp. 263–381.
- [21] Anderson, C. and Dahleh, M. D. (1998) Rapid computation of the discrete Fourier transform, *SIAM J. Sci. Comput.*, Vol. 17 pp. 913–919.
- [22] Dutt, A. and Rokhlin, V. (1993) Fast Fourier transforms for nonequispaced data, *SIAM J. Sci. Comput.*, Vol. 14, pp. 1368–1393.
- [23] Dutt, A. and Rokhlin, V. (1995) Fast Fourier transforms for nonequispaced data II, *Appl. Comput. Harmon. Anal.* Vol. 2, pp. 85–100.
- [24] Ware, A. F. (1998) Fast approximate Fourier transform for irregularly spaced data, *SIAM Review*, Vol. 40, No. 4, pp. 838–856.
- [25] Potts, D., Steidl, G. and Tasche, M. (2000) Fast Fourier transforms for nonequispaced data: A tutorial, in *Modern Sampling Theory: Mathematics and Application* by J. J. Benedetto and P. Ferreira, eds., ch. 12, pp. 253 - 274, Birkhauser.
- [26] Duijndam, A.j.W. and Schonewille, M.A. (1999), Nonuniform fast Fourier transform, *Geophysics*, Vol. 64, No. 2, pp. 539–551.
- [27] Nguyen, N. and Liu, Q.H. (1999) The regular Fourier matrices and nonuniform fast Fourier Transforms, *SIAM Journal of Scientific Computing*, Vol. 21, No. 1, pp. 283–293.
- [28] Fessler, J.A., and Sutton, B.P. (2003) Nonuniform Fast Fourier Transforms Using Min-Max Interpolation, *IEEE Trans. on Signal Processing*, Vol. 51, No. 2, pp. 560–574.
- [29] Greengard, L. and Lee, J-Y. (2004) Accelerating the Nonuniform Fast Fourier Transform, *SIAM Review*, Vol. 46, pp. 443–454.
- [30] Averbuch, A., Coifman, R., Donoho, D.L., and Israeli, M. (2003) Fast Slant Stack: A notion of Radon transform for data in a Cartesian grid which is rapidly computible, algebraically exact, geometrically faithful and invertible, to appear in *SIAM Scientific Computing*.

- [31] Mersereau, R. M. and Oppenheim, A. V., (1974) Digital reconstruction of multidimensional signals from their projections, *Proc. IEEE*, Vol. 62, No. 10, pp. 1319–1338.
- [32] Pasciak, J.E. (1981) A note on the Fourier algorithm for image reconstruction, Preprint AMD 896 Applied Mathematics Department, Brookhaven National Laboratory, Upton, New York 11973.
- [33] Edholm, P. and Herman, G.T., (1987) Linograms in image reconstruction from projections, *IEEE Trans. Medical Imaging*, MI-6, No. 4, pp. 301-307.
- [34] Lawton, W., (1988) A new Polar fourier-transform for computer-aided tomography and spotlight synthetic aperture radar, *Ieee Transactions On Acoustics Speech And Signal Processing*, Vol. 36, No. 6, pp. 931–933 June 1988.
- [35] Averbuch, A. and Shkolnisky, Y. (2003) 3D Fourier based discrete Radon transform, *Applied and Computational Harmonic Analysis* Vol. 15, No. 1, pp. 33–69, July.
- [36] Keller, Y., Averbuch, A., and Israeli, M. (2005) A PseudoPolar FFT technique for translation, rotation and scale-invariant image registration, *IEEE Trans. On Image Processing*, Vol. 14, No. 1, pp. 12–22.
- [37] Rabiner, L. R., Schafer, R. W. and Rader, C. M. (1969) The Chirp z-Transform algorithm and its application”, *Bell System Technical Journal*, Vol. 48, pp. 1249–1292.
- [38] Bailey, D. H. and Swartztrauber P. N. (1991) The fractional Fourier transform and applications, *SIAM Review*, Vol. 33, No. 3 pp. 389–404.
- [39] Golub, G.H. and Van Loan C.F, (1996) *Matrix Computations*, John Hopkins University Press, Third edition.
- [40] Buckheit, J. and Donoho, D.L. (1995) WaveLab and Reproducible Research, in *Wavelets and Statistics*, Springer-Verlag, Berlin.
- [41] Choi, C.S., Donoho, D.L., Flesia, A.G., Huo, X., Levi, O., and Shi, D. (2002) About BeamLab a Toolbox for New Multiscale Methodologies, Stanford University - Statistics Department Technical Report, <http://www-stat.stanford.edu/~beamlab>.