Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Computational Physics 206 (2005) 1–5
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

Short Note

The type 3 nonuniform FFT and its applications

June-Yub Lee [a,*], Leslie Greengard [b]

[a] Department of Mathematics, Ewha Womans University, 11-1 Daehyundong, Seodaemoongu, Seoul, 120-750, Korea
[b] Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA

Received 16 August 2004; received in revised form 11 October 2004; accepted 10 December 2004
Available online 22 January 2005

Abstract

The nonequispaced or nonuniform fast Fourier transform (NUFFT) arises in a variety of application areas, including imaging processing and the numerical solution of partial differential equations. In its most general form, it takes as input an irregular sampling of a function and seeks to compute its Fourier transform at a nonuniform sampling of frequency locations. This is sometimes referred to as the NUFFT of type 3. Like the fast Fourier transform, the amount of work required is of the order $O(N \log N)$, where N denotes the number of sampling points in both the physical and spectral domains. In this short note, we present the essential ideas underlying the algorithm in simple terms. We also illustrate its utility with application to problems in magnetic resonance imagin and heat flow.

Keywords: Nonuniform fast Fourier transform; Fourier integral; Heat equation; Magnetic resonance imaging

1. Introduction

The nonuniform or unequispaced fast Fourier transform (NUFFT) arises in a variety of application areas, including imaging processing and the numerical solution of partial differential equations. In its simplest form, one is given an irregular sampling of a function at N data points and seeks to compute the coefficients of the corresponding Fourier series. Following Dutt and Rokhlin [2], we refer to this as a type 1 transform. The adjoint of this procedure is that of evaluating a given Fourier series at a set of nonuniform target points, which we refer to as a type 2 transform. Like the fast Fourier transform (FFT) for equispaced data, the NUFFT reduces the cost of the computation from $O(N^2)$ operations to $O(N \log N)$ operations. This short note is not intended as a review article, and we refer the reader to a sampling of the relevant

* Corresponding author. Tel.: +82 2 3277 3451; fax: +82 2 3277 2289.
E-mail addresses: jylee@math.ewha.ac.kr, jylee@cims.nyu.edu (J.-Y. Lee).

0021-9991/$ - see front matter
doi:10.1016/j.jcp.2004.12.004

literature in the papers [1,3,4,7,8]. Additional citations can be found in our earlier paper [5], where we describe simple and efficient implementations of these NUFFTs.

Here, we consider the most general such transform, where both the "physical" and "Fourier" space sampling are nonuniform and discuss some of its applications. More precisely, in d dimensions, we consider the computation of

$$F_k = \sum_{j=0}^{N-1} f_j e^{\pm i\mathbf{s}_k\cdot\mathbf{x}_j}, \tag{1}$$

at N locations $\mathbf{s}_k$. This is sometimes referred to as a type 3 version of the NUFFT. The algorithm is not new. Descriptions can be found, for example, in [1,2,7,8]. Nevertheless, it is worth describing the essential features of the method in simple terms. Moreover, the type 3 NUFFT has interesting applications which do not appear to be widely appreciated.

2. The nonuniform FFT of type 3

We can think of (1) as a discretization of the continuous Fourier transform,

$$F(\mathbf{s}) = \frac{1}{(2\pi)^d}\int_{-\infty}^{\infty}\cdots\int_{-\infty}^{\infty} f(\mathbf{x})e^{\pm i\mathbf{s}\cdot\mathbf{x}}\,d\mathbf{x} \tag{2}$$

using nonuniformly sampled discretization points and evaluated at nonuniformly sampled frequencies.

All existing nonuniform FFTs are based, in essence, on combining a local interpolation scheme with the standard FFT. The type 3 NUFFT is no exception. For the sake of simplicity, however, we limit the detailed discussion to the one dimensional case.

Note first that Eq. (1) can be interpreted as the continuous Fourier transform of the function

$$f(x) = \sqrt{2\pi}\sum_{j=0}^{N-1} f_j\delta(x-x_j) \tag{3}$$

evaluated at the point $s = s_k$. We obviously cannot sample the function $f(x)$ on a uniform mesh, so we begin by convolving $f(x)$ with the (one-dimensional) Gaussian $g_\tau(x) = e^{-x^2/4\tau}$. Thus, we define $f_\tau(x)$ by

$$f_\tau(x) = f * g_\tau(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} f(y)g_\tau(x-y)\,dy. \tag{4}$$

Since f_τ is now a smooth, infinitely differentiable function, it can be well-resolved by a uniform mesh in x whose spacing is determined by the parameter τ. For this, we assume that $|x_j| \leqslant X$. We define the discrete equispaced samples of f_τ by

$$f_\tau(n\Delta_x) = \sum_{j=0}^{N-1} f_j g_\tau(n\Delta_x - x_j). \tag{5}$$

Let us now carry out a less intuitive anti-diffusion step, whose purpose will become clear shortly. For this, we define the (continuous) Fourier transform of $g_\tau(x)$ by $G_\tau(s)$. A straightforward calculation shows that $G_\tau(s) = \sqrt{2\tau}e^{-s^2\tau}$. We let

$$f_\tau^{-\sigma}(x) = \frac{f_\tau(x)}{G_\sigma(x)} = \frac{1}{\sqrt{2\sigma}}e^{\sigma x^2}f_\tau(x). \tag{6}$$

The Fourier transform of $f_\tau^{-\sigma}$, namely

$$F_\tau^{-\sigma}(s) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} f_\tau^{-\sigma}(x)e^{-isx}\,dx,$$

can be computed with high accuracy on a uniform grid in s with spacing Δ_s using the standard FFT on a sufficiently fine grid. That is,

$$F_\tau^{-\sigma}(m\Delta_s) \approx \frac{\Delta_x}{\sqrt{2\pi}} \sum_n f_\tau^{-\sigma}(n\Delta_x) \mathrm{e}^{-\mathrm{i}mn\Delta_x\Delta_s}. \tag{7}$$

Remark 1. Here, we assume that $|s_k| \leqslant S$. Because of the rapid decay of the Gaussian in (5), we can ignore contributions to $f_\tau\ (n\Delta_x)$ from points x_j more than a certain distance away with an exponentially small error. We will denote by m_{sp} the number of grid points to which we extend the influence of a point source in each direction. Given m_{sp}, we choose $\Delta_x \leq \frac{\pi}{S}\frac{1}{R}$, $\Delta_s \leq \frac{\pi}{X+m_{sp}\Delta_x}\frac{1}{R}$, and $M_r = \frac{2\pi}{\Delta_x\Delta_s}$, where the oversampling parameter $R > 1$. The actual values of R and m_{sp} can be optimized once the accuracy requirements are known.

The next step is to recover the values $F_\tau(s_k)$ by convolving $F_\tau^{-\sigma}(s)$ with $g_\sigma(s)$: i.e.,

$$F_\tau(s_k) = \left(F_\tau^{-\sigma} * g_\sigma\right)(s_k) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} F_\tau^{-\sigma}(s) g_\sigma(s_k - s)\,\mathrm{d}s \approx \frac{\Delta_s}{\sqrt{2\pi}} \sum_m F_\tau^{-\sigma}(m\Delta_s)\, g_\sigma(s_k - m\Delta_s). \tag{8}$$

This follows from the convolution theorem and the fact that we already carried out the deconvolution step in (6).

Once the values $F_\tau(s_k)$ are known, it remains only to correct for the initial smoothing to obtain the desired values $F(s_k)$. An elementary calculation shows that

$$F(s_k) = \frac{1}{\sqrt{2\tau}} \mathrm{e}^{s_k^2\tau} F_\tau(s_k). \tag{9}$$

This follows again from the convolution theorem. The actual implementation requires a complete specification of all details. We do not repeat the analysis of [2] here, but summarize the relevant results as follows: if we let

$$\Delta_x \leq \frac{\pi}{S}\frac{1}{\sqrt{2}}, \quad \Delta_s \leq \frac{\pi}{X + m_{sp}\Delta_x}\frac{1}{\sqrt{2}}, \quad \tau \cong \frac{\Delta_x}{2\pi}\frac{m_{sp}}{\sqrt{2}(\sqrt{2}-1)}, \quad \sigma \cong \frac{\Delta_s}{2\pi}\frac{m_{sp}}{\sqrt{2}(\sqrt{2}-1)}, \quad M_r = \frac{2\pi}{\Delta_x\Delta_s}, \tag{10}$$

then carrying out the convolutions in (5) and (8) with $m_{sp} = 9$ yields about six digits of accuracy. Carrying out the convolutions with $m_{sp} = 18$ yields about twelve digits of accuracy. Efficient implementation of these steps can be carried out using the fast Gaussian gridding algorithm of [5]. The higher dimensional versions involve more notation but are obvious extensions of the one-dimensional scheme. Appropriate values of M_r, Δ_x and Δ_s can be chosen for each dimension separately.

Remark 2. Greater efficiency can be achieved by using convolution kernels other than Gaussians [1,3,7], but at an increased storage cost compared to the scheme of [5].

3. Numerical examples

The NUFFT of types 1, 2 and 3 have been implemented (in Fortran) with fast gridding in one, two and three dimensions. Detailed experiments were presented in [5] for the first two types in one and two dimensions. Here, we illustrate the performance of the type 3 transforms in the context of two applications.

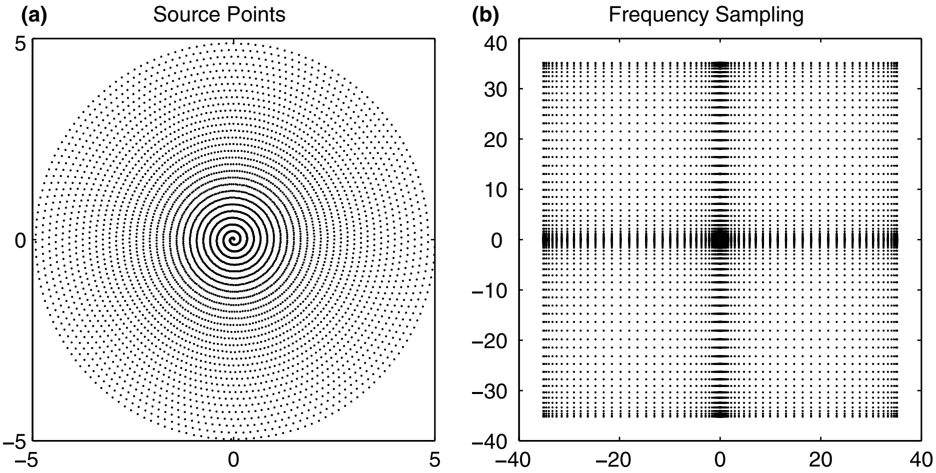

Fig. 1. The left-hand figure shows a set of sampling points for a function $f(\mathbf{x})$ and the right-hand figure shows a set of discretization points in Fourier space.

Example 1. An important application of the type 3 NUFFT is to heat flow in exterior domains. We have shown [6] that the free-space heat kernel is well-represented in the Fourier domain using a highly non-uniform sampling – one that clusters exponentially to the origin (Fig. 1(b)). In order to project a function onto these Fourier modes, one must evaluate sums of the form

$$F(\mathbf{k}_j) \approx \sum_{n=1}^{N} w_n f(\mathbf{x}_n) \mathrm{e}^{-\mathrm{i}\mathbf{k}_j \cdot \mathbf{x}_n},$$

where the points $\mathbf{x}_n$ are the sampling points for $f(x)$. (The tensor product discretization shown is easy to construct but not entirely optimal. More sophisticated quadratures following the general approach of [6] would achieve a modest reduction in the total number of nodes.) Suppose for example that $f(x)$ is a singular heat source concentrated on the curves depicted in Fig. 1(a). With N = 22,500 points and a spectral discretization using 150×150 nodes, the type 3 NUFFT with six digits of accuracy requires 0.4 s while the direct calculation requires about 110 s. If the data were uniformly spaced, the standard FFT would require only about 0.01 s.

Example 2. A second important application of the nonuniform FFT is to magnetic resonance imaging (MRI) [9]. Under the assumption of a perfectly homogeneous magnetic field, the MRI hardware is able to acquire the Fourier transform of a particular tissue property at selected points in the frequency domain. The signal produced during the "readout phase" at time t is given by

$$s(t) = \int \rho(\mathbf{x}) \mathrm{e}^{-\mathrm{i}2\pi\mathbf{k}(t)\cdot\mathbf{x}} \, \mathrm{d}\mathbf{x},$$

where $\mathbf{x} = (x^1, x^2)$ is a point in the two-dimensional image plane. In other words, $s(t)$ is precisely the value of the Fourier transform $\hat{\rho}$ at the location $\mathbf{k}(t) = (k_1(t), k_2(t))$. In the presence of a (known or unknown) field inhomogeneity given by $\phi(\mathbf{x})$, however, we have

$$s(t) = \int \rho(\mathbf{x}) \mathrm{e}^{-\mathrm{i}2\pi\mathbf{k}(t)\cdot\mathbf{x}} \mathrm{e}^{-\mathrm{i}\phi(\mathbf{x})t} \, \mathrm{d}\mathbf{x}. \tag{11}$$

Suppose now that one seeks to model the signal $s(t)$ from known functions of space $\rho(\mathbf{x})$ and $\phi(\mathbf{x})$. This requires the computation of

$$s(t_j) \approx \sum_{n=1}^{N} w_n \rho(\mathbf{x}_n) e^{-i2\pi \mathbf{k}(t_j)\cdot \mathbf{x}_n} e^{-i\phi(\mathbf{x}_n)t_j} = \sum_{n=1}^{N} w_n \rho(\mathbf{x}_n) e^{-i\mathbf{K}_j \cdot \mathbf{X}_n},$$

where $\mathbf{K}_j = (k_1(t_j), k_2(t_j), t_j)$ and $\mathbf{X}_n = (2\pi x_n^1, 2\pi x_n^2, \phi(\mathbf{x}_n))$. Thus, by embedding the data points in a higher dimensional space, one can carry out the transformation using a type 3 NUFFT. We cite one timing result, for a case where the (x^1, x^2) variables are discretized using a regular 128×128 mesh on the unit box, and k-space is traversed by a spiral trajectory similar to that shown in Fig. 1(a) up to a maximum frequency of 60. The time interval of the readout phase was scaled is to $[0, 1]$, and $\phi(\mathbf{x})$ was allowed to vary in the interval $[-5\pi, 5\pi]$. Under these conditions, direct evaluation required about 140 s, while the NUFFT required 3.8 s to obtain six digits of accuracy.

4. Conclusions

We have presented a simple version of the type 3 nonuniform FFT. It can be used to approximate the continuous Fourier transform when neither the spatial nor the Fourier domain spacing is regular. Further, it allows for the evaluation of more general integral operators such as (11) by embedding them in a higher dimensional space.

Acknowledgement

This work was supported by the Applied Mathematical Sciences Program of the US Department of Energy under Contract DEFGO200ER25053.

References

[1] G. Beylkin, On the fast Fourier transform of functions with singularities, Appl. Comput. Harmonic Anal. 2 (1995) 363–383.
[2] A. Dutt, V. Rokhlin, Fast Fourier tramsforms for nonequispaced data, SIAM J. Sci. Comput. 14 (1993) 1368–1383.
[3] J.A. Fessler, B.P. Sutton, Nonuniform Fast Fourier transforms using min–max interpolation, IEEE Trans. Signal Proc. 51 (2003) 560–574.
[4] K. Fourmont, Non-equispaced Fast Fourier transforms with applications to tomography, J. Fourier Anal. Appl. 9 (2003) 431–450.
[5] L. Greengard, J.-Y. Lee, Accelerating the nonuniform fast fourier transform, SIAM Rev. 46 (2004) 443–454.
[6] L. Greengard, P. Lin, Spectral approximation of the free-space heat kernel, Appl. Comput. Harmonic Anal. 9 (2000) 83–97.
[7] A. Nieslony, G. Steidl, Approximate factorizations of Fourier matrices with nonequispaced knots, Lin. Alg. Appl. 366 (2003) 337–351.
[8] D. Potts, G. Steidl, M. Tasche, Fast Fourier Transforms for Nonequispaced Data: A Tutorial, Modern Sampling Theory: Mathematics and Applications, in: J.J. Benedetto, P. Ferreira (Eds.), Birkhauser Boston, Applied and Numerical Harmonic Analysis Series, 2001, Chapter 12, pp. 249–274.
[9] B.P. Sutton, D.C. Noll, J.A. Fessler, Fast iterative reconstruction for MRI in the presence of field inhomogeneities, IEEE Trans. Med. Imag 22 (2003) 178–188.