Swimming Against the Current: DFT Imaging Revisited

W. D. Cotton May 31, 2018

***Abstract*—Use of the long ignored, but simple, DFT imaging technique is reexamined and compared in computational efficiency and image quality to the more popular, but more complex, grid & FFT method. In spite of impressive advances in computer technology, the DFT technique is still orders of magnitude slower than grid & FFT for a very small EVLA test case and the difference is expected to increase dramatically with problem size for modern radio interferometer arrays. This is not the case for optical/IR interferometers where the data and image volumes are small enough that the computational cost of imaging (although not deconvolution) are insignificant. The images produced by the DFT and grid & FFT methods are essential indistinguishable except at the very edge of images where the anti-alias filtering in grid & FFT is inadequate to completely suppress aliasing.**

***Index Terms*—interferometry, direct Fourier transform**

I. Introduction

In days of old IBM mainframes ruled Big Iron and radio interferometers produced tiny amounts of data. This was especially true of VLBI when 3, 4 or 5 antenna arrays were commonly used and images were a couple dozen pixels on a side at best. During this period the simple "DFT" ("direct Fourier transform") method of forming dirty images from visibility data was viable. One such technique is described in [1]. The cost of this technique is of order $n \times m$ cosine evaluations where n is the number of visibility estimates and m is the number of pixel elements. In recent decades both n and m have grown dramatically and the DFT imaging has been abandoned in favor of the more complex but more efficient "grid and FFT'. However, optical/IR interferometry is still in the small data-small image regime. As computation power has grown, the ability to compute massive numbers of cosines has also grown so a comparison of the DFT and grid & FFT methods is interesting.

A potential issue with grid & FFT imaging is the implicit assumption that the visibility samples are randomly sampled on the grid. This may not be a good assumption when small amounts of data are involved. Furthermore, the FFT algorithm assumes that the sky being imaged is strictly periodic; the failure of this presumption allows aliasing of structure out of the field of view into the field of view. Sidelobes of structure in the field of view which would appear outside the field of view will also be aliased. The gridding process must then include an anti-aliasing filter and a correction made to the derived image for the effects of this filter. A comparison of the image results is useful as well as of the computational speeds. This memo discusses a comparison of DFT and grid & FFT methods using the Obit package [2] [1].

National Radio Astronomy Observatory, 520 Edgemont Rd., Charlottesville, VA, 22903 USA email: bcotton@nrao.edu

II. DFT Imaging

An early VLBI DFT imaging and self calibration method is described in [1]. The basic technique is to make a cosine transform of the sampled visibility data. The cosine transform is appropriate in the case of a real sky image giving rise to visibility data with Hermitian symmetry. In the case of a limited field of view where the "w" component of the baseline can be ignored, the dirty image is

$$M_{i,j} = \sum_{k=0}^{k=n} w_k a_k \cos(-\phi_k - 2\pi(u_k x_i - v_k y_j)) / \sum_{k=0}^{k=n} w_k$$

and the dirty Beam is given by

$$B_{i,j} = \sum_{k=0}^{k=n} w_k \cos(-2\pi(u_k x_i - v_k y_j)) / \sum_{k=0}^{k=n} w_k$$

where a_k is the visibility amplitude, ϕ_k the visibility phase, w_k is the visibility weight and u_k and v_k are the uv coordinates for visibility k. x_i is the "X" (Right Ascension) coordinate value of the i cells and y_i is the"Y" (Declination) coordinate of the j cells of the image and beam grids.

With modern interferometers a "visibility" is really a visibility spectrum involving samples at an array of sky frequencies. This adds an extra summation over frequency around the summations in the equations above.

A. Fast Cosine Evaluation

Large numbers of cosine evaluations are a common problem and fast libraries are available to make use of the vector hardware in modern CPUs. The vector instructions use operands the width of the memory bus which is either 128 bits (4 floats) or 256 bits (8 floats) on current processors. These allow 4 or 8 floating point operations at the cost of 1. Each of these has several levels of implementation, SSE2 (128 bit) is available on all cpus of interest; AVX is the basic 256 bit version and AVX2 an enhanced version. Fast cosine implementations are available as header files with inline routines implemented as vector intrinsics[2]. The SSE version, sse_mathfun.h is available at https:github.com/juj/MathGeoLib/blob/master/src/Math/sse_mathfun.h and the AVX/AVX2 version, avx_mathfun.h at https://github.com/reyoung/avx_mathfun.

[1] http://www.cv.nrao.edu/∼bcotton/Obit.html

[2] Look like function calls but get compiled into simple vector instructions

TABLE I
EVLA TIMING TEST

machine	imager	nThread	real sec	CPU sec	ratio
panther	Imager	2	0.84	0.91	
panther	DFTIm	2	791	1565	1884
smeagle	Imager	16	1.09	0.91	
smeagle	DFTIm	16	176	2474	323

Notes: ratio is twice the DFTIm real time divided by the Imager real time as DFTIm only makes an image or beam and Imager makes both,

B. Threading

The principle data dependency in the equations in Section II is the summation. This dependency is broken and multiple threads are feasible if each thread has its private copy of summation arrays for $M_{i,j}$ or $B_{i,j}$. At the end the summation, arrays are combined and normalized (by $1/\sum_{k=0}^{k=n} w_k$). Multi-threading can thus increase the performance of the DFT algorithm.

III. IMPLEMENTATION

DFT imaging is implemented in Obit Task DFTIm which uses the SSE/AVX math libraries[3] implemented in the ObitSinCos and ObitVecFunc modules. Data sets are "strip mined" with large blocks of visibilities being read at a time and divided among the threads being used. Threading uses gthreads thread pools to reduce the overhead of starting threads and the same pool of threads is used to process the entire data-set. Each thread has a double precision copy of the accumulation array to reduce the numerical noise of the potentially large number of adds into the array. After the full data-set has been processed, the thread accumulation arrays are combined and normalized. Either a beam or dirty image is created in a given run of the program. Only "natural" weighting is implemented.

Deconvolution can be done using Obit python class CleanImage. AIPS task APCLN can be used but images must be a power of 2 in each dimension.

IV. TESTING

Testing used two machines: panther, a laptop with two 2.6 GHz cores with AVX2, 8 GByte of RAM and an SSD disk; and smeagle, a workstation with sixteen 3.1 GHz cores with AVX, 256 GByte of RAM, SSD and raid 5 disks.

A. EVLA

This test data-set was a double snapshot image with the EVLA in D configuration. This included 600 visibility spectra with 896 channels each. Stokes I was imaged onto an 800×800 cell image. The grid & FFT used Obit task Imager with equivalent parameters to the DFTIm imaging. Timing results are given in Table I. The difference in run times for the panther and smeagle tests is about that expected given the differences in CPU speeds, number of threads and AVX/AVX2.

[3]I couldn't get the simple cosine functions to work and used the sincos function which is almost as fast.

The images produced by DFTIm and Imager are displayed in Figure 1 using the same stretch. The only areas of significant difference are the very edge pixels where aliasing effects are significant in the grid & FFT image.

B. VLTI

Optical/IR interferometers are characterized by small number of baselines and small images. The differences between DFT and grid & FFT imaging quality might be significant. For this test, the uv coverage, 5 visibility samples with 11 frequency channels is representative of a snapshot with the ESO VLTI interferometer on Cerro Paranal, Chile. The data and images sizes are so small that timing tests will not produce meaningful results and only the image comparison is done. A single point source with no noise is used as the sky model. A comparison of the results is shown in Figure 2. The only areas of significant difference are the very edge pixels where aliasing effects are significant in the grid & FFT image.

V. DISCUSSION

The tests described above strongly suggest that DFT imaging is not practical in the near term for radio interferometers. The test in Section IV-A is tiny by current standards yet the run-time ratio was 2000 for the two core laptop and 300 for the sixteen core workstation. The run time of DFT imaging is expected to scale linearly with cpu speed, number of threads, data volume and image size. This scaling of the cost of DFT imaging is such that moderate to large radio array projects are prohibitively expensive for no significant gain. This is not the case for optical/IR interferometers as the data-sets are so small that the computation cost of imaging is insignificant.

The image comparisons of the DFT and grid & FFT imaging showed significant differences only at the very edge of the images where the anti-aliasing filter in the grid & FFT method is failing to completely suppress aliased responses. Thus, the grid & FFT imaging is adequate if a sufficient guard-band is allowed along the edge of the image which is not used in the analysis. Aliasing is a result of forcing the data onto a grid and using the FFT algorithm which assumes an infinitely repeating sky. DFT imaging is closer to a true Fourier transform so lacks the aliasing problem. This has been understood for a very long time.

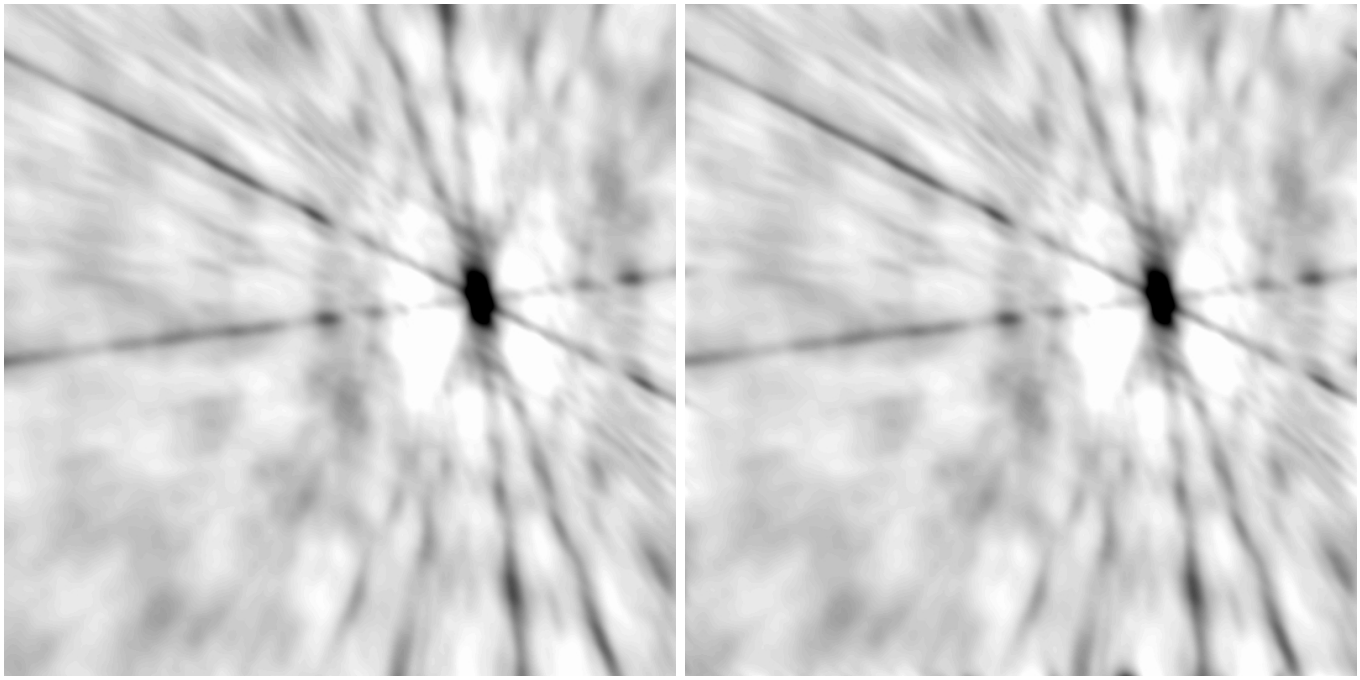

Fig. 1. Dirty images of EVLA test using DFTIm (left) and Imager (right). Shown as reverse gray-scale from -0.02 to 0.1 Jy. Some differences appear along the very edges and are due to inadequate alias suppression in the grid & FFT image.

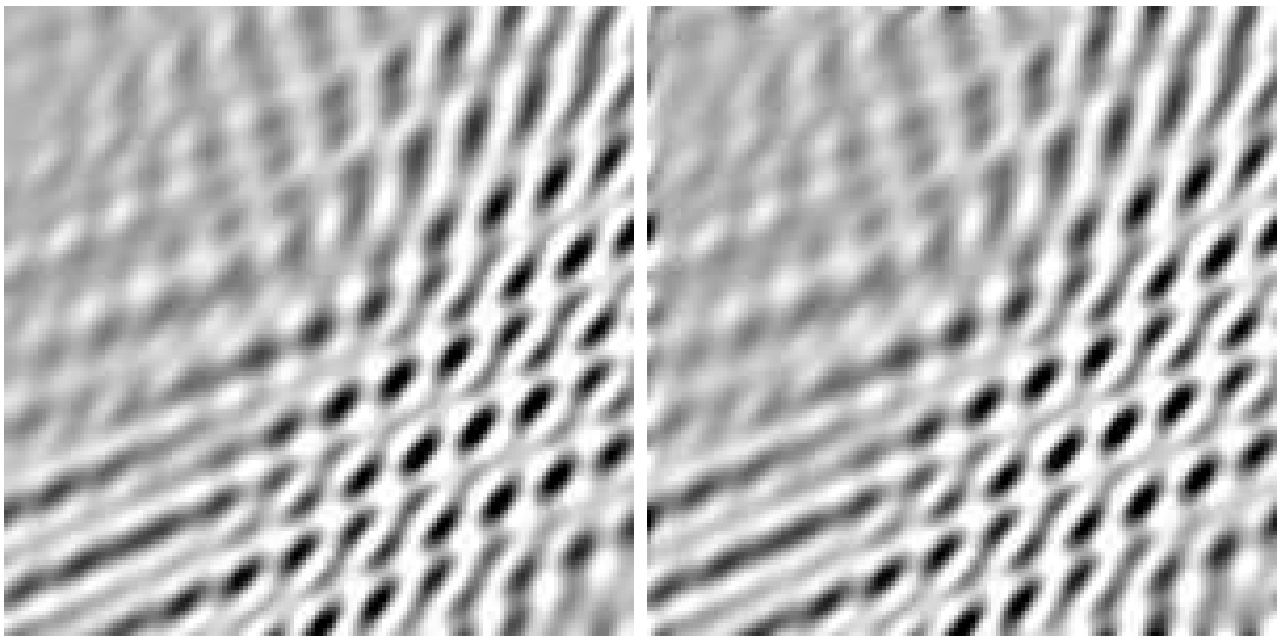

Fig. 2. Dirty images of VLTI test using DFTIm (left) and Imager (right). Shown as reverse gray-scale from -0.2 to 0.5 Jy. Some differences appear along the very edges and are due to inadequate alias suppression in the grid & FFT image.

REFERENCES

[1] W. D. Cotton, “A method of mapping compact structure in radio sources using vlbi observations,” *Astron. J.*, vol. 84, p. 1122, 1979.
[2] W. D. Cotton, “Obit: A Development Environment for Astronomical Algorithms,” *PASP*, vol. 120, pp. 439–448, 2008.