

An Exact, Non-iterative Mojette Inversion Technique Utilising Ghosts

Shekhar Chandra¹, Imants Svalbe¹, and Jean-Pierre Guédon²

¹ School of Physics, Monash University, Australia

² IRCCyN-IVC, École polytechnique de l'Université de Nantes, France

Abstract. Mojette projections of discrete pixel arrays form good approximations to experimental parallel-beam x-ray intensity absorption profiles. They are discrete sums taken at angles defined by rational fractions. Mojette-like projections form a “half-way house” between a conventional sinogram and fully digital projection data. A new direct and exact image reconstruction technique is proposed here to invert arbitrary but sufficient sets of Mojette data. This new method does not require iterative, statistical solution methods, nor does it use the efficient but noise-sensitive “corner-based” inversion method. It instead exploits the exact invertibility of the prime-sized array Finite Radon Transform (FRT), and the fact that all Mojette projections can be mapped directly into FRT projections. The algorithm uses redundant or “calibrated” areas of an image to expand any asymmetric Mojette set into the smallest symmetric FRT set that contains all of the Mojette data without any re-binning. FRT data will be missing at all angles where Mojette data is not provided, but can be recovered exactly from the “ghost projections” that are generated by back-projecting all the known data across the calibrated regions of the reconstructed image space. Algorithms are presented to enable efficient image reconstruction from any exact Mojette projection set, with a view to extending this approach to invert real x-ray data.

1 Introduction

Tomographic reconstruction of discrete images is typically based on a sinogram of projection data obtained using finite aperture detectors oriented at some set of angles in continuous space. The mix of continuous and discrete sampling causes problems with image artefacts. To ensure the uniqueness of the reconstruction, some interpolation or smoothing is required within the Filtered Back-Projection or Fourier Inversion methods [1]. Any discrete object can, however, be reconstructed exactly from a sufficient set of discrete projections. Significant advances in discrete inversion methods have appeared in [2, 3, 4]. The discrete projections of the Mojette Transform (MT) of [5] resemble closely the form of real data projections, but the projections are restricted to angles whose tangents are rational fractions and the number of rays per projection varies with the projection angle. These MT projections are matched to the grid on which the object is reconstructed, removing the conflict between discrete and continuous sampling and

provide the promise of a more faithful reconstruction process and, potentially, less dose.

A disadvantage of the very general MT approach is the lack of a direct algorithm that avoids iteration [6] or the use of noise-sensitive “corner-based” inversion approaches [5]. The discrete Finite Radon Transform (FRT) of [7] has a very simple, direct and exact inversion algorithm that is a consequence of its prime number array-size and because of its fixed length, periodically wrapped projection structure. The FRT is a particular case of the MT. Mojette projections can be mapped into FRT form. The FRT inverse could then be used to reconstruct the image, but the number of symmetric FRT projections is usually larger than the asymmetric MT set. Providing the MT set contains sufficient information to reconstruct the data exactly, we propose a method to recover the full FRT projection set and thus provide a direct inversion method for the MT via the FRT formalism. We begin with a brief overview of the MT and the FRT.

1.1 The Mojette Transform (MT)

The MT is a discrete linogram [8] transform for objects of arbitrary shape where only a minimum projection set needs to be defined [5]. The projection angles are confined to the Farey Sequence q_i/p_i , which is a set of irreducible rational fractions where $\text{gcd}(p_i, q_i) = 1$ with $i = 1, \dots, \mu$ and μ is the total number of projections [9]. The number of bins in a given MT projection i , \mathcal{M}_i , depends on the angle $\theta_i = \tan^{-1}(q_i/p_i)$ through

$$\mathcal{M}_i = |p_i|(Q - 1) + |q_i|(P - 1) + 1, \tag{1}$$

for a $P \times Q$ object. An example of a MT for a 4×4 image is given figure 1. The minimum number of projections μ_{\min} required for exact reconstruction is




Fig. 1. An example of a Mojette Transform for a discrete image of size 4×4 . The bold lines within the right-hand grid shows a possible reconstruction path using a corner-based inversion method [5].

dictated by the Katz criterion

$$N \leq 1 + \max \left(\sum_{i=1}^{\mu_{\min}} |p_i|, \sum_{i=1}^{\mu_{\min}} |q_i| \right) \tag{2}$$

for an $N \times N$ object. Note that one can easily over-specify the projection set (i.e., $\mu > \mu_{\min}$) [10]. When the Katz criterion is not satisfied, ghosts or phantoms arise in the reconstructed image [11, 5]. These ghosts were first studied by Katz (1977, [12]) and are artefacts formed within an image that sum to zero at certain projection angles. Examples of ghosts are shown in figure 2. The goal of this paper is to enable direct the use of the FRT formalism to reconstruct exact images from arbitrary Mojette projection data.




Fig. 2. Depiction of simple ghosts in an image. (a) shows image values that lead to ghosts in the four projection directions shown. (b) shows a 61×61 greyscale image of ghosts (black denotes negative and white denotes positive greyscale values). Here the image will “disappear” when viewed at 15 of the 62 FRT projection directions.

1.2 The Finite Radon Transform (FRT)

The FRT $R(t, m)$ of image $I(x, y)$ is a discrete prime-periodic transform where the image space is considered to be a torus of size $\mathbf{p} \times \mathbf{p}$ pixels, where \mathbf{p} is prime [7]. An example of a projection in the FRT and how these projections relate to the MT [13, 14] is given in figure 3.

The FRT projections $R(t, m)$ sum all pixels lying along the lines

$$\begin{aligned} y &\equiv mx + t \pmod{\mathbf{p}}, & \text{for } 0 \leq m < \mathbf{p} \\ x &\equiv t, & \text{for } m = \mathbf{p} \end{aligned} \tag{3}$$

where $x, y, m, t \in \mathbb{Z}$, with lines at each translate $t = 0, \dots, \mathbf{p} - 1$ having slope m . Each projection consists of \mathbf{p} translates which wrap around the image modulo \mathbf{p} . Due to the primality of \mathbf{p} , each translate of a projection sums precisely \mathbf{p} pixels, ensuring that each bin has \mathbf{p} terms within it. This in turn ensures that every pixel is sampled exactly once for each projection. When all $\mathbf{p} + 1$ projections are taken, the FRT can then be inverted exactly as shown in figure 4(a). An algebraic example of FRT projection and inversion is shown in figure 5. When




Fig. 3. (a) shows how projections are taken in $I(x, y)$ to form $R(t, m)$. Here the $m = 2$ projection has translate $t = 0$ on a $p = 5$ grid. (b) shows how FRT projections can be mapped to Mojette projections along the direction of nearest neighbour lines [13].

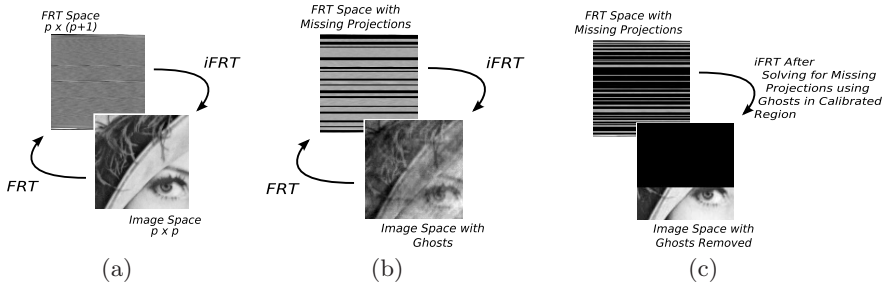


Fig. 4. (a) shows an example of an FRT space and its reconstructed image. The transform is exact and one-to-one. (b) shows the ghosts that form when projections are missing from the FRT (shown as black regions). (c) shows how a sufficient calibrated image region (black) enables exact reconstruction of a sub-region.

any FRT projections are missing, the reconstruction is only partially correct and has ghost artefacts at each pixel corresponding to a superposition of a (negative) contribution from each missing projection. The result of FRT inversion when some projections are missing is shown in figure 4(b). However, given a calibrated region of sufficient size, the missing projections can be recovered from the ghosts (see figure 4(c)). In the next two sections we discuss how ghosts are structured, how they can be used to recover missing projections and what constitutes a sufficient calibrated region.

2 Ghosts in the FRT

Consider a single missing projection or ghost, labelled as $\tilde{a} = \{a_0, \dots, a_{p-1}\}$, located at row m_a in $R(t, m)$. In the FRT inversion, each projection from $R(t, m)$

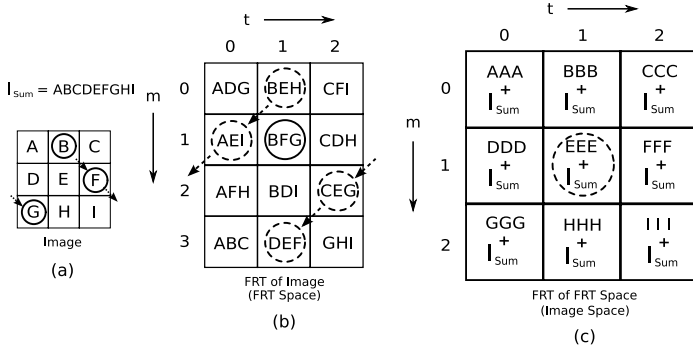


Fig. 5. Projections of the image (a) are taken using equation (3) and placed into FRT space (b) (see example pixels in closed circles). To invert the FRT, projections of the FRT space are taken in the same way as before, but with $m' = p - m$ for projections $0 < m < p$ (see example pixels in broken circles). Then every pixel in the result (c) is normalised by subtracting I_{Sum} and dividing by p to recover the exact image (a).

is mapped, as a whole, onto each row in image space, translated (and wrapped periodically) by m_x pixels at row x . The ghost \tilde{a} also shifts periodically by $-m_a$ on each subsequent row of the image. This structure is shown in insets A and B of figure 6. The periodic shift structure of the rows can also be viewed along the column direction, as shown in inset C of figure 6. In the column data, the ghost bin indices increment at multiples of m_j . Conversion between whole rows and columns of ghost data is possible using this property. The following section presents the algorithm to exploit this periodic structure in solving for missing projections.

3 The Mojette Ghost-Recovery Algorithm (MGR)

The ghost structure represented in figure 6 forms the basis for recovering the extra FRT projections needed to invert MT data. An image containing a calibrated region (i.e., where the image values are known) has the ghosts structured as depicted in inset A of figure 6. It is possible to unscramble these ghost superpositions (or signals) exactly in this region to solve for the missing FRT projections. These signals essentially form a set of coupled linear equations that require $\mathcal{N} \times p$ calibrated row or column ghost pixels to solve for the $\mathcal{N} \times p$ missing bins of projection values. In [6] a conjugate gradient approach was used to unpack Mojette bins into consistent image data. Here the simple arithmetic nature of the FRT is exploited to provide a direct algebraic reconstruction. The idea of using redundant image regions has been used in [10] and [15] for encryption of data as projections. The MGR algorithm is demonstrated in figure 7.

Figure 7(a) shows three rows from the calibrated region where a total of three ghosts (\tilde{a} , \tilde{b} , \tilde{c} and $\mathcal{N} = 3$) are superimposed. We can first eliminate ghost \tilde{a} by cyclically shifting (or rotating) row 1 and row 2 by $-m_a$ and $-2m_a$ (i.e.,

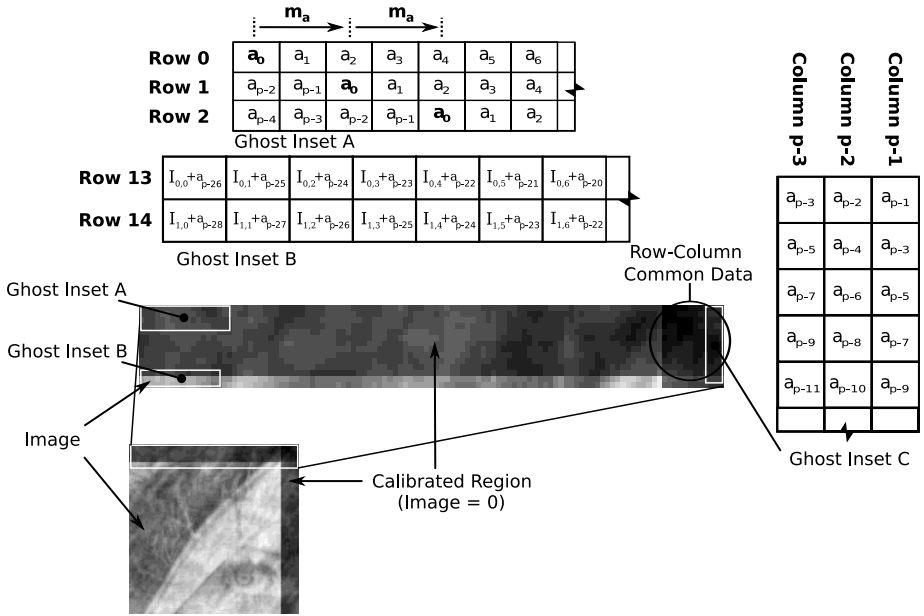


Fig. 6. Ghost structure in image space. The image (Lena) is of 100×100 pixels embedded (by padding) in a prime image space of 113×113 . If we assume there is only one missing projection $\tilde{\mathbf{a}}$ in FRT space ($\mathcal{N} = 1$) at $m_a = 2$, then inset A shows how the ghost will be structured inside the calibrated region. Inset B shows how artefacts $\tilde{\mathbf{a}}$ affects the image data. All ghosts will be structured in the same m -dependent pattern. If multiple ghosts are present, the resulting pattern is a superposition of each ghost pattern.

to the left) respectively and differencing. The result is shown in figure 7(b) which contains a mixture of shifted ghosts $\tilde{\mathbf{b}}$ and $\tilde{\mathbf{c}}$ as differences. However, the differences have the same ghost order but are shifted with respect to each other. Therefore, we may remove the $\tilde{\mathbf{b}}$ differences by aligning them by $-(m_b - m_a)$ and differencing again. The result is shown in figure 7(c), where only differences containing ghost $\tilde{\mathbf{c}}$ remain. The differencing process leaves a systematic pattern (marked by the arrows) caused by the relative shifts with respect to m_c used to remove the other ghosts. A consecutive sum of the $\tilde{\mathbf{c}}$ differences (where the previous result is carried on to the next) at steps $m_c - m_b$ (\mathfrak{p} times) and then at steps $m_c - m_a$ (\mathfrak{p} times), cancels the negative terms in the pattern, leaving only ghost $\tilde{\mathbf{c}}$. Note that this is only made possible by the properties of prime congruence, as attempting to integrate the differences having lengths other than \mathfrak{p} will result in an incomplete ghost. This sum (a discrete integration) leaves the result with a constant offset that can be removed after each integration (see figure 7(d)). The result $\tilde{\mathbf{c}}$ can then be back substituted and the process repeated to solve for the remaining ghosts. If the redundant region is sufficient (i.e., there are at least \mathcal{N} full rows or columns in the calibrated region), all

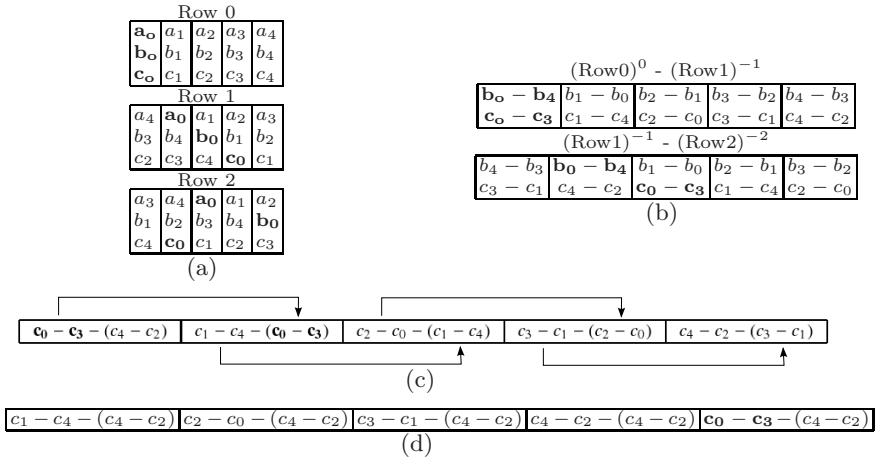


Fig. 7. An example of the MGR algorithm for $p = 5$ and $\mathcal{N} = 3$. (a) shows ghosts $\tilde{\mathbf{a}}$, $\tilde{\mathbf{b}}$ and $\tilde{\mathbf{c}}$ for $m = 1, 2, 3$ superimposed in Row 0,1 and 2. The addition signs are removed for clarity. (b) shows ghost $\tilde{\mathbf{a}}$ eliminated after alignment and pair differences of rows 0, 1 & 2. The superscripts on row labels denote direction and size of cyclic shifts. (c) shows ghost $\tilde{\mathbf{a}}$ and $\tilde{\mathbf{b}}$ eliminated, leaving a summation of four shifted copies of ghost $\tilde{\mathbf{c}}$. At each step the remaining ghosts are compounded by the differencing. Integration (at steps shown by arrows) simplifies terms to leave the ghost $\tilde{\mathbf{c}}$ with constant offset. (d) shows the first integration result at step $m_c - m_b$. Offset $-(c_4 - c_2)$ is removed and the next integration done at step $m_c - m_a$.

missing projections can be recovered exactly. In summary, the algorithm for solving missing projections is:

1. Align the first ghost (say $\tilde{\mathbf{a}}$) by shifting the row or column data cyclically using $-m_a$ and differencing $\mathcal{N} - 1$ pairs of signals. This removes ghost $\tilde{\mathbf{a}}$ from all those signals. Shifting by $-m_a$ also shifts all other ghosts by $-m_a$.
2. Align the next ghost (say $\tilde{\mathbf{b}}$) by the new relative shift $-(m_b - m_a)$ and difference. This removes ghost $\tilde{\mathbf{b}}$ from the signals.
3. Repeat until only the \mathcal{N}^{th} ghost is left. Note that the solve order is completely arbitrary (i.e., in figure 7 we could have solved for $\tilde{\mathbf{b}}$, followed by $\tilde{\mathbf{a}}$ to get $\tilde{\mathbf{c}}$).
4. Integrate the last ghost (removing the integration constant after each integration) to get the missing projection. Back substitute and repeat the algorithm to determine the other ghosts.

The advantageous features of this algorithm are:

1. The algorithm processes just \mathcal{N} 1D cyclic blocks of data, each of length p .
2. The algorithm is exact and non-iterative in nature.
3. The solve order is arbitrary. When processing noisy projection data, this may be an advantage relative to “corner-based” methods.
4. The algorithm is easily and highly parallelised, because each missing projection can be solved independently.

MGR will also help to understand the nature of artefacts in CT reconstructions arising from non-uniqueness as discussed by [11], [12] and [16]. In the next section, the FRT ghosts are utilised to demonstrate a simple inversion mechanism for the Mojette transform.

4 MGR-Based Inversion

The essential difficulty of inverting the MT is the arbitrary number of projections and the variable number of bins in each projection. The fixed properties of the FRT, together with the ability to recover ghosts, provide an explicit inversion scheme for the MT. Mojette projections can be mapped into an FRT space with only a minor re-ordering of the translates t without any interpolation (as was required for reconstructing the real sinogram data in [13]). The net effect of placing MT projections without interpolation in FRT space is to embed the sought image within a larger image space. The size of the larger image space is dictated by the prime \mathfrak{p} chosen so that the longest MT projection (the largest \mathcal{M}_i as given by equation (1)) fits into FRT space. Merging of bins is avoided in order to preserve the information in the Mojette set under the Katz Criterion. Since MT projections tend to be longer and smaller in number than for FRT projections, the resulting FRT space largely consists of missing projections. After recovering the ghosts from the calibrated region in image space using the MGR algorithm, the Mojette projections may be inverted using the very simple inverse FRT algorithm. This final inversion of the FRT can also be done efficiently by selecting the appropriate translates from $R(t, m)$ space so only the embedded image is reconstructed (ignoring the calibrated area). Hence, the MT inversion scheme is:

1. Reorder and place MT projections into a sufficiently large FRT space.
2. Recover the missing FRT projections. This is done by creating ghosts and using the MGR algorithm to untangle each of the missing projections.
3. Invert the FRT data using only the translates in FRT space that correspond to the embedded image.

The scheme has the favourable property of not relying on the corners of the image (as required by most other exact inversion methods) and should be more noise tolerant. In the following section, initial results of the MGR inversion scheme are presented and its current limitations are discussed.

5 Results

The MGR scheme was applied to the MT inversion of a 100×100 image. The MT of the image is shown in figure 8(a) and the subsequent exact reconstruction in figure 8(b). In order to solve the system, the total number of missing projections could not be greater than the size of the calibrated region. As the FRT has $\mathfrak{p} + 1$ projections, a total of 101 MT projections had to be taken. This meant solving




Fig. 8. Illustration of the exact and non-iterative image reconstruction using MGR inversion. (a) shows the MT for the Lena image (100×100 pixels) using 101 projections. (b) shows the exact inversion result of the MGR algorithm using rows. The larger image space is $\mathbf{p} = 1289$ and required solving for 1188 missing projections. The reconstructed image, mostly calibrated space, has been truncated for display.

for $1289 - 101 = 1188$ missing projections, since the longest Mojette projection was $\mathcal{M}_i(\max) = 1288$ (hence $\mathbf{p} = 1289$) corresponding to $p = -7$ and $q = 6$.

An interesting phenomenon also occurs during the differencing process. When taking differences while shifting periodically, the magnitude of the differences compound as shown in figure 7(c). When \mathcal{N} is large, the pattern of resultant differences of the \mathcal{N}^{th} ghost are almost purely sinusoidal. This sinusoidal effect is probably due to the differences being zero mean and periodic. The differences are always zero mean because all FRT projections must sum to the same constant I_{Sum} . The process was replicated using a computer algebra package and solved analytically. Here the ghost bin addresses were represented by variables, as in figure 7. The weight each ghost bin contributes to the ghost differencing process is shown in figure 9. This shows that the sinusoidal behaviour is independent of image data and dependent on geometry of the problem. For large \mathcal{N} , the same sinusoidal behaviour leads to rapid numerical growth of the difference values (reaching 10^{40} in the case of figure 8(b) for $\mathcal{N} = 1188$). An example of the rapid numerical growth is given in figure 10. This problem so far has been overcome by using arbitrary-precision integers for the differencing and integration processes in the MGR algorithm. Controlling the numerical growth is an area of future work which is discussed in the next section.

6 Further Work

The main remedy for the limitation in the number of projections required will be to fully utilise all parts of the redundant image area (i.e., a combination of

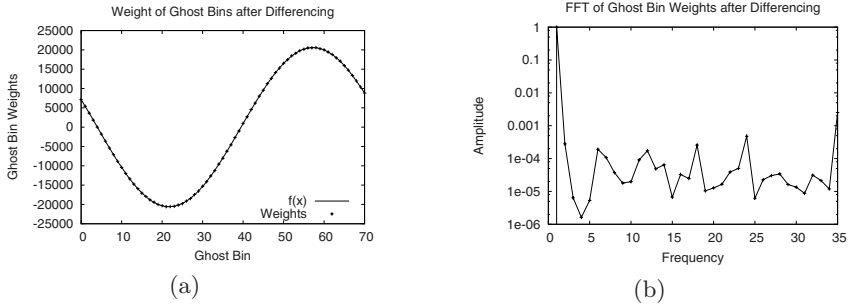


Fig. 9. The result of algebraically solving a ghost system with $p = 71$ and $\mathcal{N} = 58$ as described in section 3. (a) shows the distribution of weights of the missing projection bins at the end of the differencing process. Here the fit $f(x) = -a \sin(bx + c)$ where $a = 20591.2$, $b = 0.0885 \approx \frac{2\pi}{p}$ and $c = 12.212$. (b) shows Fourier transform coefficients of (a), normalised to peak absolute value and plotted on a logscale.




Fig. 10. The max and minimum positive values of the compounded ghosts $\mathcal{N} = 170$ during the differencing process. The max values have a fine-scale oscillatory imprint as shown in (b). Each period consists of one less point than the previous cycle, as each ghost is progressively eliminated.




Fig. 11. An example of solving $\mathcal{N} = 2$ using a row and a column. Here we denote ghosts as $\tilde{\mathbf{a}}$ and $\tilde{\mathbf{b}}$ having projections $m = 2, 3$ respectively. (a) shows the initial data showing row 1 and column $p - 1$. (b) shows the row put into column form and aligned so $\tilde{\mathbf{a}}$ can be removed. (c) shows $\tilde{\mathbf{a}}$ eliminated by differencing the two rows in table (b). The resulting $\tilde{\mathbf{b}}$ differences cannot be integrated because of the cancellation of like terms. This occurs regardless of the starting position of the integration. Also, the integration step required is no longer constant.

complete rows and columns of the calibrated region). Initial attempts to use both the rows and columns simultaneously have been frustrated by the interactions between row and column data (see figures 6 and 11) as p independent values per row or column are needed to untangle the ghosts and integrate the result. We are also considering an alternative algorithm that uses the calibrated area to retrieve the compounded ghost values from the reconstructed image area, pixel by pixel, rather than independently solving for each ghost in its entirety.

To remedy the problem of numerical growth, the compounding of the ghost differences must be understood. Any truncation or wrapping of the differences leads to incorrect results and any re-normalisation of the values will still require high precision floating point representation. Also note from the Fourier transform of figure 9(a) (shown in figure 9(b)) that the sinusoidal signature includes very small but essential contributions from other frequencies whose origins are unclear. Future work will also address the speed relative to other Mojette inversion algorithms as well as the robustness of MGR to noise in the projection data.

7 Conclusion

An MT projection set mapped into $R(t, m)$ space has missing FRT projections or ghosts that have a known, precise structure in the reconstructed image. This structure allows the construction of algorithms to recover the missing FRT projections. The algorithms rely on a calibrated region being present with the image within a larger image space. The MGR algorithm is exact, non-iterative and easily parallelised.

The algorithm was applied to invert the MT as shown in figure 8. Known MT projections are converted to $R(t, m)$ projections. The sought image then becomes a subset of a larger image space. Since the small number of MT projections leave many missing FRT projections, portions of the larger image space are used as a calibrated area for the MGR algorithm, in order to recover these projections. Once the FRT space is filled, an efficient subset inverse FRT can be used to invert the MT.

The algorithm has two main drawbacks at the moment. The current use of either full rows or columns in the calibrated region leads to a limit on the number of MT projections able to be processed. The algorithm also leads to interesting ghost structures (see figure 9) and involves very large integer values as shown in figure 10. A further study of these ghost structures may reveal a means to control or suppress the numerical growth. The remaining area of concern is then to establish the robustness of the MGR algorithm to noise in the MT projection data. Initial results show that small levels of added noise are well tolerated, but more work is required to quantify these findings.

References

- [1] Kak, A.C., Slaney, M.: Principles of Computerized Tomographic Imaging. Society of Industrial and Applied Mathematics (2001)
- [2] Beylkin, G.: Discrete Radon transform. IEEE Trans. Acoust., Speech, Signal Processing 35(2), 162–172 (1987)

- [3] Averbuch, A., Donoho, D., Coifman, R., Israeli, M., Walden, J.: Fast slant stack: A notion of Radon transform for data on a cartesian grid which is rapidly computable, algebraically exact, geometrically faithful and invertible. Tech. Report, Stanford University 11 (2001)
- [4] Kingston, A., Svalbe, I.: Projective transforms on periodic discrete image arrays. *Advances in Imaging and Electron Physics* 139, 75–177 (2006)
- [5] Guédon, J.P., Normand, N.: The Mojette transform: the first ten years. In: Andrès, É., Damiand, G., Lienhardt, P. (eds.) *DGCI 2005. LNCS*, vol. 3429, pp. 79–91. Springer, Heidelberg (2005)
- [6] Servières, M., Idier, J., Normand, N., Guédon, J.P.: Conjugate gradient Mojette reconstruction. *Proceedings of the SPIE - The International Society for Optical Engineering* 5747(1), 2067–2074 (2005)
- [7] Matúš, F., Flusser, J.: Image representation via a finite Radon transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15(10), 996–1006 (1993)
- [8] Edholm, P., Herman, G.: Linograms in image reconstruction from projections: Part 1. back projecting using the Radon transform. In: *Proc. Twenty-First Annual Hawaii International Conference on System Sciences*, vol. 4, pp. 48–50 (1988)
- [9] Hardy, G., Wright, E.: *An Introduction to the Theory of Numbers*, 5th edn. Clarendon Press, Oxford (1979)
- [10] Normand, N., Guédon, J.P., Philippe, O., Barba, D.: Controlled redundancy for image coding and high-speed transmission. In: *Proceedings of the SPIE - The International Society for Optical Engineering*, vol. 2727, pp. 1070–1081 (1996)
- [11] Philippe, O., Guédon, J.P.: Correlation properties of the Mojette representation for non-exact image reconstruction. In: *ITG-Fachbericht Verlag, Proc. Picture Coding Symposium*, vol. 143, pp. 237–242 (1997)
- [12] Katz, M.: *Questions of Uniqueness and Resolution in Reconstruction from Projections*. Springer, Heidelberg (1977)
- [13] Svalbe, I., van der Spek, D.: Reconstruction of tomographic images using analog projections and the digital Radon transform. *Linear Algebra and Its Applications* 339, 125–145 (2001)
- [14] Servières, M., Guédon, J.P., Normand, N.: A discrete tomography approach to PET reconstruction. In: *Proc. 7th International Conf. on Fully 3D Reconstruction in Radiology and Nuclear Medicine* (2003)
- [15] Kingston, A., Svalbe, I.: Geometric effects in redundant keys used to encrypt data transformed by finite discrete Radon projections. In: *Proc. IEEE Digital Imaging Computing: techniques and applications*, Cairns, Australia (2005)
- [16] Boyd, J., Little, J.: Complementary data fusion for limited-angle tomography. In: *Proceedings 1994 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 288–294 (1994)