

Reconsidering the advantages of the three-dimensional representation of the interferometric transform for imaging with non-coplanar baselines and wide fields of view

D. M. P. Smith, A. Young, and D. B. Davidson

Department of Electrical and Electronic Engineering, Stellenbosch University, 7600 Stellenbosch, South Africa
e-mail: davidson@sun.ac.za

Received 24 June 2015 / Accepted 25 April 2017

ABSTRACT

Radio telescopes with baselines that span thousands of kilometres and with fields of view that span tens of degrees have been recently deployed, such as the Low Frequency Array, and are currently being developed, such as the Square Kilometre Array. Additionally, there are proposals for space-based instruments with all-sky imaging capabilities, such as the Orbiting Low Frequency Array. Such telescopes produce observations with three-dimensional visibility distributions and curved image domains. In most work to date, the visibility distribution has been converted to a planar form to compute the brightness map using a two-dimensional Fourier transform. The celestial sphere is faceted in order to counter pixel distortion at wide angles, with each such facet requiring a unique planar form of the visibility distribution. Under the above conditions, the computational and storage complexities of this approach can become excessive. On the other hand, when using the direct Fourier transform approach, which maintains the three-dimensional shapes of the visibility distribution and celestial sphere, the non-coplanar visibility component requires no special attention. Furthermore, as the celestial samples are placed directly on the curved surface of the celestial sphere, pixel distortion at wide angles is avoided. In this paper, a number of examples illustrate that under these conditions (very long baselines and very wide fields of view) the costs of the direct Fourier transform may be comparable to (or even lower than) methods that utilise the two-dimensional fast Fourier transform.

Key words. methods: analytical – methods: numerical – techniques: image processing – techniques: interferometric

1. Introduction

The two-dimensional (2D) representation of the interferometric transform was developed for radio telescopes that have nearly coplanar antennas and small field of views (FoVs). In this case, the visibility and celestial distributions are approximated as parallel 2D planes, so that the 2D fast Fourier transform (FFT) can be applied directly. The results for instruments that satisfy these conditions, from the earliest synthesis arrays such as OMT (Ryle et al. 1965) to contemporary instruments such as KAT (Carignan et al. 2013), validate this approach. Now, a range of telescopes which have been recently deployed or are currently being developed have large non-coplanar baselines, such as SKA (Dewdney et al. 2009); or subtend wide FoVs, such as PAPER (Jacobs et al. 2011); or both, such as LOFAR (van Haarlem et al. 2013).

Under such conditions encountered with new instruments, the three-dimensional (3D) shape of the visibility distribution and the curved surface of the celestial sphere deviate sufficiently from being planar that the usual approximation is no longer valid. In order to exploit the reduction in computational complexity offered by the 2D FFT, the visibility distribution can be converted into a set of approximately equivalent 2D forms (Cornwell et al. 2008; Offringa et al. 2014) and the celestial sky can be faceted into a set of 2D planes (Cornwell 1988). For very long baselines and for very wide FoVs, the computational and the storage costs for these methods can become high compared to the alternative, namely using the direct Fourier transform (FT)

to compute the brightness map on the curved surface of the celestial sphere directly from the sampled visibility distribution.

In this paper, such an approach is investigated that utilises two concepts: firstly, maintaining the 3D shapes of the visibility distribution and celestial sphere; and secondly, using the direct FT to compute the brightness map on the curved surface of the celestial sphere. Neither of these concepts is novel. For example, in Perley (1999) the 3D shape of the visibility distribution is maintained by faceting the celestial sphere as a set of 2D planes; this set of planes increases in size as the extent of the baselines or the FoVs expand. It should be noted that in defining the celestial sphere with only two of its components as independent, the FoVs are restricted to that of a hemisphere. Given that proposed space-based radio telescopes would have FoVs of a sphere and that terrestrial instruments currently being deployed and developed could have baselines thousands of kilometres long, the 3D representation of the interferometric transform is reviewed in this paper as an alternate approach under these conditions.

As the underlying theory is already well established, only the key equations related to the 3D representation of the interferometric transform are repeated here. Instead, this paper is dedicated to a series of case studies developed to highlight the advantages of making use of this approach for these types of cases, and this is the main contribution of this paper. The examples compare the computational and storage costs of this method with two methods based on the 2D representation of the interferometric transform: w -projection (Cornwell et al. 2008) and w -stacking (Offringa et al. 2014), with both 2D approaches utilising faceting (Cornwell 1988), if necessary. It will be shown that

under certain conditions, the cost of the 3D approach is equivalent to, or in some cases better than that of these 2D approaches.

2. Interferometric transform

A radio telescope using synthesis imaging consists of antennas (or stations comprising an array) that track the phase centre vector \mathbf{s}_0 and sample the emissions around this position in the celestial sky. Owing to the rotation of the Earth over the duration of the observation, T_{obs} , the spacing between the antennas (the baselines) shifts in relation to \mathbf{s}_0 . The result is a visibility distribution comprising elliptical loci that fill a cylindrical volume with limits set by the maximum baselines, $\pm u_{\text{max}}, \pm v_{\text{max}}, \pm w_{\text{max}}$. Due to the astronomical distances involved, the location of the celestial sources are defined solely by their position in the sky, thereby reducing the celestial sky to a sphere of unit radius.

When only one polarisation of the emissions is considered at a time, the interferometric equation is traditionally used, which is derived from the van Cittert-Zernike theorem (Thompson et al. 2004, Ch. 14). This is the approach followed in this paper. Methods aimed at full polarimetric imaging are better developed using the measurement equation (Hamaker et al. 1996), as for example in Carozzi & Woan (2009) and Smith et al. (2015). Based on the detailed development in Thompson (1999), the key terms of the interferometric equation are reproduced here using the same notation. The visibility at baseline vector \mathbf{b} , $V(\mathbf{b})$, based on the radio brightness at celestial position vector $\boldsymbol{\sigma}$, $I(\boldsymbol{\sigma})$, is

$$V(\mathbf{b}) = \int_S A(\boldsymbol{\sigma}) I(\boldsymbol{\sigma}) \exp(-j2\pi \frac{\mathbf{b} \cdot \boldsymbol{\sigma}}{\lambda}) d\Omega \quad (1)$$

where $A(\boldsymbol{\sigma})$ is the primary beam power pattern at $\boldsymbol{\sigma}$, λ is the wavelength of observation, S is the celestial surface, \cdot is the 3D vector dot product, and $d\Omega$ is the infinitesimal solid angle. This formulation assumes identical antennas, but can be extended.

Based on the extension of Thompson (1999) in Perley (1999), the generalised form of (1) can be rewritten in the form of a 3D representation of the interferometric equation, as in

$$V(u, v, w) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} A(\ell, m, n) I(\ell, m, n) \times \exp[-j2\pi(u\ell + vm + w[n-1])] d\ell dm dn \quad (2)$$

which is a 3D Fourier transform with inverse

$$I(\ell, m, n) = A(\ell, m, n)^{-1} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} V(u, v, w) \times \exp[+j2\pi(u\ell + vm + w[n-1])] dudvdw \quad (3)$$

where the brightness is *only* non-zero on the unit sphere of the celestial sky, when $\sqrt{\ell^2 + m^2 + n^2} = 1$. The baseline vector normalised to wavelength, \mathbf{b}/λ , has components (u, v, w) that point towards east, towards north and towards \mathbf{s}_0 , respectively; and $\boldsymbol{\sigma}$ has components (ℓ, m, n) that are direction cosines that are measured with respect to the u, v, w axes, respectively. The integration over the celestial surface could be defined by two spherical components; however three are required for (2) and (3) to form a 3D Fourier pair. In this paper, as the FoV is not constrained to a hemisphere, the n component is not defined by the ℓ and m components, altering the form of (3) from that in Perley (1999).

The infinite bounds of the 3D Fourier pair of (2) and (3) reduce to finite bounds determined by the maximum baselines,

Table 1. Costs of 3D approach.

Function	Computational	Storage
Observe		N_{vis}
Transform	$N_{\text{vis}}N_{\text{img}}$	N_{img}

B_{max} , and the angular resolution of the antennas, θ_{ant} , respectively. Furthermore, this Fourier pair reduces to a discrete form due to the sampling of the visibilities at discrete points determined by the sampling interval of the observation, ΔT_{obs} . The radio brightness is computed from these sampled visibilities at discrete positions in the sky. In the following section, the computational and storage costs, O_{cmp} and O_{str} , of the discretised form of (3) will be determined for the 3D approach proposed in this paper, as well as for the 2D approach traditionally used.

3. Synthesis imaging

When the 3D approach is used, the sampled visibility distribution is directly transformed onto the curved surface of the celestial sphere, thereby avoiding pixel distortion at wide angles. Furthermore, as the 3D shape of the visibility distribution is maintained, the non-coplanar visibility component requires no special treatment. The costs of this approach are listed in Table 1, where,

$$N_{\text{vis}} = N_{\text{frq}}N_{\text{ant}}(N_{\text{ant}} - 1)(T_{\text{obs}}/\Delta T_{\text{obs}}) \quad (4)$$

is the number of sampled visibilities, N_{frq} is the number of frequency bins, and N_{ant} is the number of antennas; and

$$N_{\text{img}} = \pi [1 - \cos(\theta_{\text{ant}}/2)] / \sin^2(\theta_{\text{int}}/2) \quad (5)$$

is the number of samples on the curved surface of the celestial sky and is twice as fine as the angular resolution of the interferometer, θ_{int} . The costs of the 3D approach are unaffected by the extent of the visibility distribution or the extent of the FoV, and are only dependent on the number of samples in each domain.

Under certain conditions, the 3D Fourier pair of (2) and (3) reduce to the 2D representation of the interferometric equation (Perley 1999). The visibility and celestial distributions are approximated by planar surfaces, with the visibility distribution convolved onto a regular grid perpendicular to the celestial distribution. The Nyquist sampling intervals in the gridded celestial sphere, $\Delta\ell$, Δm , and gridded visibility distribution, Δu , Δv , are

$$\begin{aligned} \Delta\ell &\leq 1/(2u_{\text{max}}), & \Delta m &\leq 1/(2v_{\text{max}}) \\ \Delta u &\leq 1/(2\ell_{\text{max}}), & \Delta v &\leq 1/(2m_{\text{max}}) \end{aligned} \quad (6)$$

where $\ell_{\text{max}}, m_{\text{max}} = \sin(\theta_{\text{ant}}/2)$ are the maximum extents in ℓ, m . The number of pixels along one dimension in both domains is

$$N_{\text{pix}} = \max(2\ell_{\text{max}}/\Delta\ell, 2m_{\text{max}}/\Delta m, 2u_{\text{max}}/\Delta u, 2v_{\text{max}}/\Delta v) \quad (7)$$

which are equal due to the N -to- N mapping of the FFT.

When using the 2D approach, the brightness map is computed from the gridded visibilities, with this transform having a much lower O_{cmp} than the 3D approach as traditionally $N_{\text{pix}}^2 \approx N_{\text{img}} \ll N_{\text{vis}}$. However, in order to make use of an FFT, there are additional O_{cmp} and O_{str} to convolve the sampled visibility distribution onto a grid. The costs of the 2D FFT are listed in Table 2, where N_{ker} is the one-dimensional (1D) size of the convolution kernel, which is typically set to 7 (Cornwell et al. 2008). In the next section, the additional costs of the functions required to correct for the non-coplanar visibility component and to avoid pixel distortion at wide angles will be investigated.

Table 2. Costs of 2D approach for small FoV and short baselines.

Function	Computational	Storage
Observe		N_{vis}
Grid	$N_{\text{vis}}N_{\text{ker}}^2$	N_{pix}^2
Transform	$N_{\text{pix}}^2 \log N_{\text{pix}}^2$	N_{pix}^2

Table 3. Costs of 2D approach using w -projection and faceting.

Function	Computational	Storage
Observe		N_{vis}
w -Correction	$N_{\text{fac}}N_{\text{vis}}N_{\text{prj}}^2$	$N_{\text{fac}}N_{\text{vis}}$
Grid	$N_{\text{fac}}N_{\text{vis}}N_{\text{ker}}^2$	$N_{\text{fac}}N_{\text{pix}}^2$
Transform	$N_{\text{fac}}N_{\text{pix}}^2 \log N_{\text{pix}}^2$	$N_{\text{fac}}N_{\text{pix}}^2$
Stitch	$N_{\text{fac}}N_{\text{pix}}^2N_{\text{ker}}^2$	N_{pix}^2

3.1. Approaches to dealing with the 3D problem

In the general case, the discrete visibilities are defined by three orthogonal coordinates. However, in certain cases, it is inefficient to convolve these samples onto an uniform 3D grid and then use a 3D FFT to compute the brightness map for the following reasons. For arrays characterised by long baselines on a curved surface like Low Frequency Array (LOFAR), the discrete visibilities fill the visibility volume very sparsely, and only a small number of gridded visibilities would be non-zero. Also, as the celestial sky is constrained to a spherical surface, only a very small number of celestial samples on the 3D grid would satisfy $\sqrt{\ell^2 + m^2 + n^2} = 1$. Therefore, to compute N_{pix}^2 celestial samples with true emissions, a much larger number of samples would need to be calculated. As such, the methods adopted to date rather correct for the non-coplanar component, and these have been successful for instruments with moderate baselines, limited deviation from non-coplanarity and restricted FoVs.

The w -projection algorithm moves the sampled visibilities onto a plane that is perpendicular to the tracking position by convolving out the non-coplanar component of the baselines (Cornwell et al. 2008). If faceting is required, the sampled visibilities are moved to N_{fac} planes, each parallel to a different facet. The costs related to this approach are listed in Table 3, where

$$N_{\text{prj}} = \lambda_{\text{max}}^2 W_{\text{max}}^2 D_{\text{max}}^{-4} \quad (8)$$

is the 1D size of the w -projection kernel, λ_{max} is the maximum wavelength, W_{max} is the maximum non-coplanar baseline, and D_{max} is the maximum antenna dimension.

Alternatively, the w -stacking algorithm splits the visibility volume into a number of planes, each of which is perpendicular to the tracking position. Each plane is FT to the image domain, with the non-coplanar components corrected for in the imaging domain (Humphreys & Cornwell 2011). If faceting is required, the sampled visibilities are split into N_{stk} planes for each facet. The computational costs and storage costs related to this approach are listed in Table 4, where (Offringa et al. 2014)

$$N_{\text{stk}} \gg 2\pi(w_{\text{max}} - w_{\text{min}}) \max_{\ell, m} \left(1 - \sqrt{1 - \ell^2 - m^2}\right) \quad (9)$$

is the number of planes required to avoid aliasing, and w_{max} and w_{min} are the extents of the non-coplanar baseline components.

Table 4. Costs of 2D approach using w -stacking and faceting.

Function	Computational	Storage
Observe		N_{vis}
w -Correction	$N_{\text{fac}}N_{\text{stk}}N_{\text{pix}}^2$	$N_{\text{fac}}N_{\text{stk}}N_{\text{pix}}^2$
Grid	$N_{\text{fac}}N_{\text{vis}}N_{\text{ker}}^2$	$N_{\text{fac}}N_{\text{stk}}N_{\text{pix}}^2$
Transform	$N_{\text{fac}}N_{\text{stk}}N_{\text{pix}}^2 \log N_{\text{pix}}^2$	$N_{\text{fac}}N_{\text{stk}}N_{\text{pix}}^2$
Stitch	$N_{\text{fac}}N_{\text{pix}}^2N_{\text{ker}}^2$	N_{pix}^2

In approximating the visibility distribution as a flat plane, the 2D approach introduces a computation error of (Cotton 1999)

$$\text{error} \approx \pi w \left(\sqrt{\ell^2 + m^2} \right). \quad (10)$$

While w -correction techniques reduce this error, practicable implementation of these methods discretise the w range of the visibilities (Offringa et al. 2014), thereby introducing some error. No such approximation is necessary when using the 3D approach. When the costs of the 2D approach become comparable to that of the 3D approach, the 3D approach is the preferred option, due to its simpler implementation and because it avoids introducing any error due to the non-coplanarity of the visibility distribution.

With the celestial distribution approximated as a plane, the outer pixels of the brightness map are distorted when the FoV exceeds $\sqrt{\lambda/B_{\text{max}}}$ (Perley 1999). The FoV is divided into facets, N_{fac} , to maintain the 2D approximation (Cornwell 1988), where

$$N_{\text{fac}} \geq [1 - \cos(\theta_{\text{ant}}/2)] / [1 - \cos(\sqrt{\theta_{\text{int}}}/2)]. \quad (11)$$

For each facet of the FoV, a unique gridded visibility distribution is required that is parallel to its own facet. To avoid aliasing, the size of each of these gridded visibility distributions must satisfy the Nyquist criterion of the full FoV, even though only the fraction of undistorted pixels, those that fall within the $\sqrt{\lambda/B_{\text{max}}}$ limit, are stitched together into the resultant brightness map.

For the 2D approach, each FT is performed on the dataset as a whole, while the functions that deal with w -correction and pixel distortion allow for computations to be performed sequentially or in parallel. For example, each facet can be processed sequentially to reduce the storage cost (by reusing data structures that store intermediate results) or all of the facets can be processed in parallel to speed up the computational time. Similar procedures can be used to deal with the w -corrections planes. For the 3D approach, each celestial sample can be computed independently, allowing for up to N_{img} sequential or parallel operations.

While the shapes of the visibility and celestial distributions are that of a cylinder and a spherical cap, respectively, gridding imposes rectangular configurations. The pixels that lie outside of these shapes have no effect on the brightness map, but increase the costs of the 2D approach. Furthermore, as planar facets only approximate a sphere, overlap between the facets is unavoidable, resulting in N_{fac} being larger than that required by (11). Conversely, with the 3D approach, there is no direct link between the number of samples in the two domains. Therefore, no sample is computed that lies outside of the bounds of either distribution.

One reason why the 2D approach is faster than the 3D approach for most cases is because the gridding function reduces the number of visibilities the 2D approach has to FT. As the values of these gridded visibilities are not equal to the values of the visibility at these grid points (Jackson et al. 1991), the FT of the

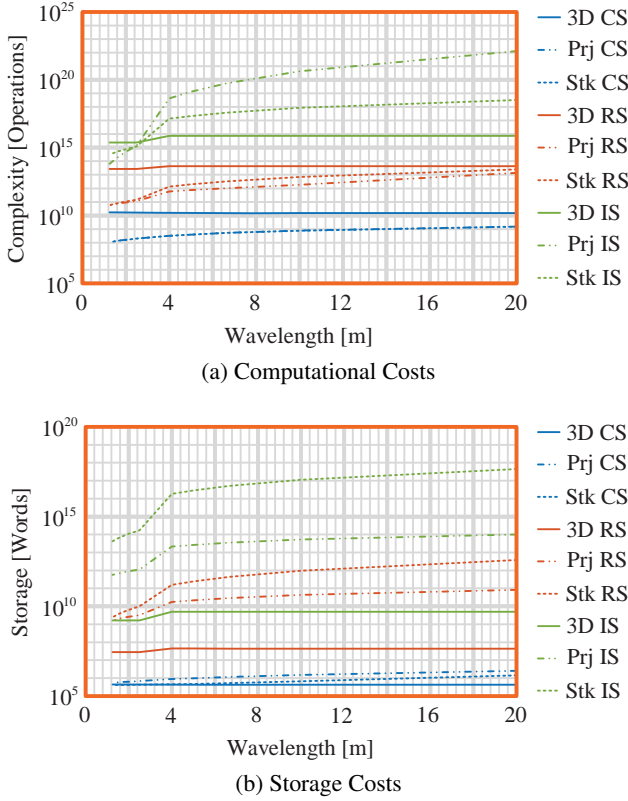


Fig. 1. Costs when $s_0 = 0^\circ$ with inner low-band antenna (LBA) stations used.

gridded visibility would not be equal to that of the true visibility. Thus, to reduce the costs of the 3D approach, a subset of sampled visibilities could be transformed. The brightness map of this subset of sampled visibilities would be more accurate than that of the gridded visibilities, as actual visibilities are transformed.

When using the 2D approach, to increase the celestial sampling density beyond that set by (7), the extent of the visibility distribution would need to be zero padded. This increases the costs of the 2D approach by multiplying N_{pix} by N_{pad} , where N_{pad} is the factor by which the sampling density is increased along one dimension. While these zero points are necessary to make use of the FFT, they have no effect on the brightness map. Conversely, when using the 3D approach, the sampling density in a sector of the FoV can be increased without artificially increasing the number of discrete visibilities through zero-padding.

4. Case studies comparing 2D and 3D approaches

To illustrate the effect of extending the baselines and FoVs on the costs of the 2D approach, examples using LOFAR will be presented in this section. In the first case, only the $N_{\text{ant}} = 24$ core stations (CS) are used. In the second case, the baselines are extended by using the CS and the $N_{\text{ant}} = 14$ remote stations (RS). In the third case, the baselines are further extended by using the CS, the RS and the $N_{\text{ant}} = 8$ international stations (IS). In each case, the FoV is extended by reducing the frequency of observation from the upper limit of the high-band antenna (HBA) to the lower limit of the LBA. Also, the non-coplanar baseline is extended by increasing the angle of the phase centre vector away from zenith. Furthermore, the FoV is extended by switching from the outer to the inner LBA stations.

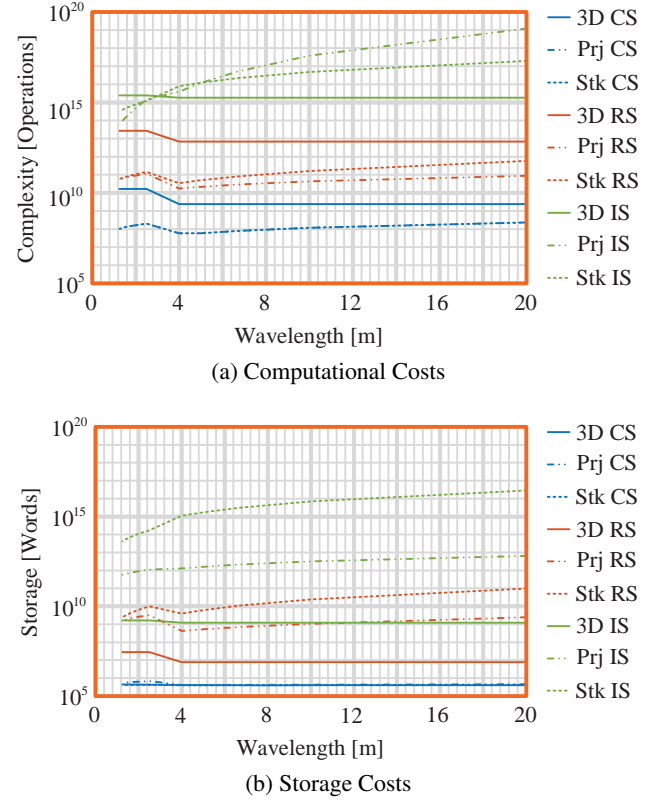


Fig. 2. Costs when $s_0 = 0^\circ$ with outer LBA stations used.

The angular resolution and location of the LOFAR stations are taken from van Haarlem et al. (2013), where

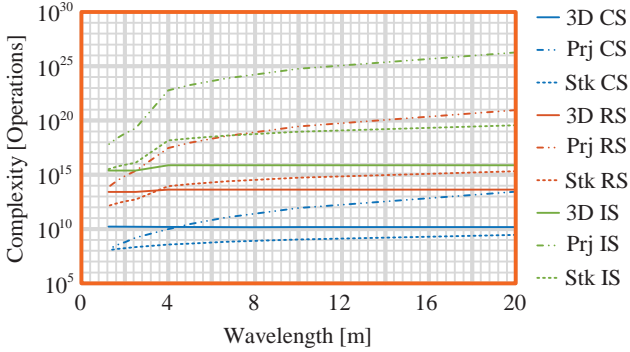
$$\begin{aligned} \theta_{\text{ant}} &= \alpha \lambda / D_{\text{max}} \\ \theta_{\text{int}} &= \alpha \lambda / B_{\text{max}} \end{aligned} \quad (12)$$

where α is the tapering of the antenna. For the cases in this paper, $T_{\text{obs}} = 12$ hr and $\Delta T_{\text{obs}} = 1$ min. When the LBA stations are used, $\alpha = 1.10$ and λ ranges from 4 to 20 m. When the HBA stations are used, $\alpha = 1.02$ and λ ranges from 1.25 to 2.5 m.


Using the equations listed in Tables 3 and 4, the results when the phase centre vector is pointing at zenith are depicted in Figs. 1 and 2. In the legend, “3D”, “Prj” and “Stk” denote using the 3D approach, using the 2D approach with w -projection, and using the 2D approach with w -stacking, respectively. Also, “CS”, “RS” and “IS” denote the three cases as defined above. The 3D approach is less storage intensive in all the cases, but is more computationally intensive in the first two cases. When the IS are used, the 2D approach is more computational intensive, with the 3D proposed for this case. Due to the wider FoV of the inner LBA stations, more facets are required, increasing the 2D approach costs, but with no effect on the 3D approach costs.

The results when the phase centre vector is pointing at 45° from zenith are depicted in Figs. 3 and 4. This misalignment of the phase centre vector increases the non-coplanarity of the visibility distribution, resulting in an increase in the costs of w -correction. Here, the 3D approach is less storage intensive in all the cases, and is only more computationally intensive in the first case. The 2D approach is more computational intensive for the RS case and the IS case, with the 3D proposed for these cases.


The effect of the extent of the baselines and the extent of the FoV on the value of N_{fac} , N_{prj} and N_{stk} for two cases are listed in Tables 5 and 6. In the first case, $\lambda = 240$ MHz and $s_0 = 0^\circ$



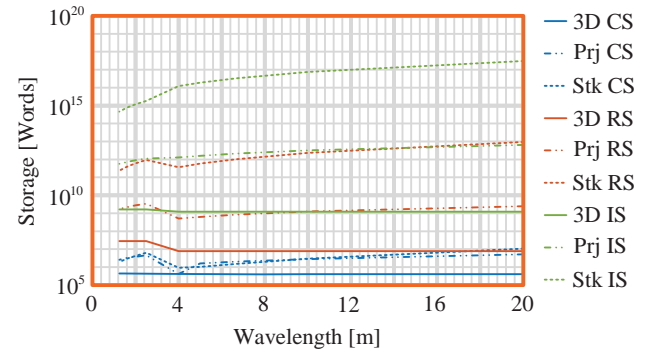
(a) Computational Costs



(a) Computational Costs



(b) Storage Costs



(b) Storage Costs

Fig. 3. Costs when $s_0 = 45^\circ$ with inner LBA stations used.

Fig. 4. Costs when $s_0 = 45^\circ$ with outer LBA stations used.

Table 5. Values for N_{fac} , N_{prj} and N_{stk} when $s_0 = 0^\circ$ at $\lambda = 240$ MHz.

Case	N_{fac}	N_{prj}	N_{stk}
Case I	5	1	1
Case II	91	1	2
Case III	517	239	51

Table 6. Values for N_{fac} , N_{prj} and N_{stk} when $s_0 = 45^\circ$ at $\lambda = 15$ MHz, with outer LBA stations used.

Case	N_{fac}	N_{prj}	N_{stk}
Case I	75	964	121
Case II	2517	597 671	3006
Case III	27 073	67 588 030	31 960

from zenith, resulting in almost planar distributions, with the 2D approach more appropriate. In the second case, $\lambda = 15$ MHz and $s_0 = 45^\circ$ from zenith, resulting in large non-coplanar baselines and a wide FoV, with the 3D approach more appropriate.

5. Conclusion

In this paper, the 3D representation of the interferometric transform has been reviewed as a potential solution to dealing with radio synthesis telescopes that have antennas (or stations) that are non-coplanar and that have antenna beamwidths that subtend wide FoVs. This solution is proposed because for these cases the visibility distribution and celestial distribution fill 3D volumes.

This 3D representation offers a number of advantages when computing the brightness map that are not possible when using the 2D approach, as the direct FT does *not* require gridding. While the maximum extent of the visibilities restrict the resolution of the brightness map, there is no direct relation between the positions of the samples in the two domains. Thus, the celestial samples could be placed directly on the curved surface of the celestial sphere, thereby avoiding pixel distortions for wide FoVs, which could otherwise require an expensive faceting approach. Also, as the 3D shape of the visibility distribution is maintained, w -correction is avoided. The 3D approach is more costly when compared to the core functions of the 2D approach. For certain cases, the costs of the 2D approach exceed that of the 3D approach when dealing with wide FoVs and non-coplanar baselines, with the 3D approach preferred as it is more accurate. Additionally, selecting which discrete visibilities to use and which celestial positions to compute could be adaptively modified to alter the accuracy and speed of computing the brightness map.

This 3D representation makes no provision to deal with unequal antenna patterns between antenna pairs, unequal signal paths to such a pair, nor polarised emissions. With this in mind, an extension of the work in this paper has used a 3D representation of the measurement equation (Hamaker et al. 1996) to deal with non-coplanar arrays and wide FoVs (Smith et al. 2015).

Acknowledgements. This work was supported by SKA South Africa, the South African Research Chairs Initiative of the Department of Science and Technology, and the South African National Research Foundation.

References

- Carignan, C., Frank, B. S., Hess, K., et al. 2013, *AJ*, **146**, 48
Carozzi, T. D., & Woan, G. 2009, *MNRAS*, **395**, 1558

- Cornwell, T. J. 1988, [A&A](#), **2002**, 316
- Cornwell, T. J., Golap, K., & Bhatnagar, S. 2008, [IEEE J. Selected Topics in Signal Processing](#), **2**, 647
- Cotton, W. D. 1999, in *Synthesis Imaging in Radio Astronomy II*, eds. G. B. Taylor, C. L. Carilli, & R. A. Perley, ASP Conf. Ser., 180, 357
- Dewdney, P. E., Hall, P. J., Schilizzi, R. T., & Lazio, T. J. L. W. 2009, [Proc. IEEE](#), **97**, 1482
- Hamaker, J. P., Bregman, J. D., & Sault, R. J. 1996, [A&A](#), **117**, 137
- Humphreys, B., & Cornwell, T. 2011, *Analysis of Convolutional Resampling Algorithm Performance*, Memo 132, Square Kilometre Array
- Jackson, J. I., Meyer, C. H., Nishimura, D. G., & Macovski, A. 1991, [IEEE Transactions on Medical Imaging](#), **10**, 473
- Jacobs, D. C., Aguirre, J. E., Parsons, A. R., et al. 2011, [ApJ](#), **734**, L34
- Offringa, A. R., McKinley, B., Hurley-Walker, N., et al. 2014, [MNRAS](#), **444**, 606
- Perley, R. A. 1999, in *Synthesis Imaging in Radio Astronomy II*, eds. G. B. Taylor, C. L. Carilli, & R. A. Perley, ASP Conf. Ser., 180, 383
- Ryle, M., Elsmore, B., & Neville, A. C. 1965, [Nature](#), **205**, 1259
- Smith, D. M. P., Young, A., & Davidson, D. B. 2015, in *International Conference on Electromagnetics in Advanced Applications*, Turin
- Thompson, A. R. 1999, in *Synthesis Imaging in Radio Astronomy II*, eds. G. B. Taylor, C. L. Carilli, & R. A. Perley, ASP Conf. Ser., 180, 1
- Thompson, A. R., Moran, J. M., & Swenson, G. W. 2004, *Interferometry and Synthesis in Radio Astronomy*, 2nd edn. (Weinheim: Wiley), 383
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, [A&A](#), **556**, A2