MeerKAT core vs. core+spur: a comparison of image quality.

I M Stewart

November 3, 2010

Abstract

The effect on continuum imaging quality of an additional 7 'spur' antennas for MeerKAT is examined. Dirty images corresponding to 3 different sky scenarios are created for a variety of different array memberships and visibility weighting schemes. It is shown that, for the full core+spur array, even omitting core-only baselines (which make up 82% of the total) yields a much superior image (in terms of both resolution *and* sensitivity) than any which can be made just with the core. A new weighting scheme is also proposed which makes the best of the uneven distribution of spur baselines.

1 Introduction

The original design for MeerKAT included, in addition to the 80-dish core array within a roughly 8 km diameter circle, an additional 7 'spur' dishes arranged along the approximately 50 km length of the roughly SE-aligned site access road. The current proposal at the time of writing replaces the 80 core dishes with 64 dishes of slightly larger diameter; it has also recently been proposed to abolish the extra 7 spur dishes. The present study was performed with the aim of exploring the cost to continuum science of losing the spur dishes.

The study is a simulation of the radio continuum sky at L band (the radio band which includes the rest-frame HI emission line at 1420 MHz) at the microjansky level. Since no-one really knows how the sky will look at this depth, several input scenarios were used. A variety of combinations of array layout and weighting were then applied to this data.

There are several issues which one would like to obtain knowledge about: e.g. sensitivity, resolution, confusion. Sensitivity, that is, how deeply one can detect sources, is a function primarily of the instrumental noise. The amount of noise present in the final image varies in a straightforward way with bandwidth and observation time, but in a less straightforward way on the weighting of visibilities. Resolution and confusion are related issues, but one can make the distinction between a situation in which the structure of individual sources is usually unresolved, but in which the density of sources at that depth is low enough for confusion between sources to be rare, and the situation in which separate sources are frequently confused.

In this report the input data to the simulation is described in section 2.1, and issues connected with processing this data to yield a set of visibilities are discussed in section 2.2. Gridding, weighting and imaging issues are discussed in section 2.3. The results are described in section 3.

2 Description of the simulations

2.1 The simulated sky

Three simulated input skies were used: two containing just point sources, the third being a more sophisticated simulated sky obtained via the SKADS database (Wilman et al [2008]).

There seems to be almost no knowlege of the progress of $n(S)$, the sky density of sources as a function of flux density S, below approximately 50 μJy. Published reports found to be useful in extrapolating this function were Condon [1984], Huynh et al [2005], Fomalont et al [2006], Kellermann et al [2008] and Gervasi et al [2008].

The shape of the $n(S)$ function over a short range of S values is most conveniently described as a power law. As one progresses towards fainter flux density values, a knee in the curve is seen at approximately 1 Jy. Below this point, the power of the increase stays fairly constant at a level significantly below the Euclidean value of 2.5 until 0.1 mJy, at which point it appears to steepen again. Condon [1984] has speculated that this shape reflects two separate populations of continuum sources, namely AGNs at the bright end and SFGs at the faint end.

Since there are few measurements below the observed flattening at 0.1 mJy, the progress to yet fainter values of S must be held to be poorly constrained. That is the reason that two different scenarios were employed for the point-source input skies in the present study.

2.1.1 'Sparse' point sources

The 'sparse' sky comprised a list of 7620 sources with random positions and flux densities. The positions were chosen from a distribution which was uniform within a square of 205 arcsec on a side. The flux densities followed a power-law of the form

$$n(S) = S^{-(2.5+\epsilon_2)}/A_2 \ \mathrm{sr}^{-1}\mathrm{Jy}^{-1},$$

which is the faint-end limit of equation (2) of Gervasi et al [2008]. The constants were set to $\epsilon_2 = -0.862$ and $A_2 = 1.116 \times 10^{-4}$, which are the values fitted to the 1400 MHz data by Gervasi et al and given in their tables 1 and 2.

The flux density distribution of the simulated sources was truncated at 1×10^{-9} Jy and 1×10^{-4} Jy. The faint-end cutoff was essentially arbitrary but certainly at this level one would expect that the contribution of individual sources to the confusion within the image would be on an average rather than an individual basis, and that fainter sources could only raise the average sky brightness slightly. I note that, if this power law is assumed to continue to $S = 0$, the total flux density per steradian is a finite quantity, thus avoiding Olbers' paradox, although the total number density of sources is not.

The bright-end cutoff is possibly more controversial, but one may argue that if any sources brighter than 1×10^{-4} Jy remain uncleaned from the data then it is probably pointless to concern ourselves with the details of antenna layouts, since we are not going to see anything really faint whatever we do.

2.1.2 'Dense' point sources

For the 'dense' sources I simply laid a ruler on figure 8 of Kellermann et al to obtain the shallowest reasonable slope on that graph consistent with the observed densities (in particular those reported from Fomalont et al [2006]) at around 1×10^{-4} Jy. The equivalent values of ϵ_2 and A_2 were -0.1 and 0.42 respectively. The sparse and faint curves cross at about 1.3×10^{-4} Jy.

The bright-end cutoff was the same as for the sparse sources but the far higher numbers at low S values forced a retreat of the faint-end cutoff to 1×10^{-8} Jy. Despite this precaution there were 266,589 sources in this input catalog.

2.1.3 The SKADS sources

The SKADS image (see top right panel in figure 3) exhibits a density of sources which appears about half-way between the sparse and dense point-source catalogs. It shows quite a number of sources which are resolved on a scale of a few arcseconds.

2.2 Simulating the visibilities

2.2.1 Approximations

For an array with as many antennas and spectral channels as MeerKAT will have, the task of simulating realistic data is a heavy one, because there is so much data to simulate. Because of the limited computing power available, it was necessary to cut some corners. To begin with, it was assumed that the calibration was perfect; that all sources in the sky had flat spectra, allowing the uncomplicated use of wide-band frequency synthesis; and that all sources were unpolarised. Distortions due to noncoplanarity were neglected. Far-field sources and all other effects of the primary beam were neglected.

Further to this, given the software and computing resources at hand, it proved impossible to adopt a spectral channel width and time sampling frequency which were fine enough to avoid significant smearing over the full field of view: my simulation program ran out of memory when I tried to do this. When dealing with the 64-dish core, the smallest practical values were found to be $\Delta\nu = 20$ MHz and $\Delta t = 120$ s. These are still far too large to avoid bandwidth smearing over the full $\sim$1 degree field of MeerKAT; however, matters were arranged so that this would not matter. Firstly, only a small $\sim$200 arcsecond square region of the primary beam was imaged in the present studies. This is about 6% of the diameter of the primary beam. A real observation of this area would still suffer bandwidth smearing with the specified sampling but not nearly so badly as would be the case if one imaged the full primary beam.

Secondly, the only thing which was simulated via the 'standard process' of first visibilities then gridding then back-transforming, was noise values. In contrast to this, the contribution made to the dirty images by sources was calculated by convolving an appropriate 'source image' by the dirty beam. Since with this method the visibilities from sources are never calculated, one does not have to worry about smearing. The disadvantages inherent in this method are twofold. Firstly, sources not located in the centre of image pixels are not treated correctly: such sources in reality give rise to 'sinc-like' far-field ripples (see Briggs [1995]). Secondly, the vignetting due to the Fourier inverse of the gridding kernel is absent.

Correct convolution of off-pixel-centre sources could be achieved in cases where some tapering is applied to the visibilities: namely by applying the tapering a priori as it were via a matching convolution of the source image. In fact the source images were constructed with the size of the taper Gaussian specified as an input, so the correct pixel values for off-centre sources could be calculated and applied to begin with. It is impossible to achieve a correct source image if one starts from an image which does not have this pre-convolution: each source in this case contributes flux to only a single pixel, thus any subsequent convolution will generate a source profile which is symmetric about the centre of that pixel, rather than about the true source position. Such 'pre-tapering' was naturally not possible in the case of naturally-weighted images: but here the dirty beams are so diffuse that ripples due to off-centre sources would be so smeared out they would make a negligible contribution in any case.

Neglecting the gridding-kernel vignetting probably makes the signal-to-noise slightly larger towards the edges of the image than it would be in reality. This was not thought to be a severe effect, but a correction would be easy to simulate if necessary.

2.2.2 The simulated observations

Three sets of 'noise only' visibilities were constructed:

1. 64-dish core.
2. Only baselines involving the extra 7 dishes.
3. The full 64+7 array.

The last was not created separately but by adding the gridded visibilities from the first two schemes. This is trivial in the case of natural weighting but has to be done carefully for any

other weighting scheme (see section 2.3.1 for details). The antenna positions are shown on figure 1. (In passing I wonder whether the positions for the core antennas are correct: the core appears flattened in the north-south direction.)

The full details of the simulated observation were as follows:

- RA and dec of the Chandra Deep Field South (RA = 53.11667 deg, dec = -27.80833 deg).[1]
- 42×20 MHz spectral channels spanning 890 to 1730 MHz.
- 359×120 s time samples spanning an hour-angle range from -6 hr to +6 hr.

The uncertainty in the real or imaginary part of a single cross-correlation (visibility) is given by the following formula, which is the weak-source, identical-antenna limit of equation 9-10 of Wroble & Walker [1999]:

$$\Delta S = 1.0 \times 10^{26} \times \frac{kT\sqrt{2}}{\eta A \sqrt{\Delta\nu \Delta t}} \text{ Jy}, \tag{1}$$

where k is Boltzmann's constant, T is the system temperature in kelvin, η the combined efficiency of the antenna and the post-antenna signal chain, A is the antenna collecting area in m^2, $\Delta\nu$ is the channel width in Hz and Δt the integration time in seconds. The following values were used in the calculation of the noise figures in the present study:

- $T = 40$ K
- $\eta = 0.7$
- $A = 2\pi r^2$ with $r = 7.5$ m
- $\Delta\nu = 2 \times 10^7$ Hz
- $\Delta t = 120$ s

The resulting value is 12.9 mJy per visibility component.

The final data (including all the figures presented in section 3) were designed to correspond to a total observed time on source of 1000 hours. This was simulated by dividing the noise dirty images by $\sqrt{1000/12}$ before adding them to the source component. (The slight difference here between siderial and terrestrial duration was neglected.)

2.3 Gridding and weighting choices

It was decided to create images of size 2048×2048 pixels, each pixel being square with a side length of 0.1 arcseconds. This translates to a grid-cell size in the UV plane of about 1007 waves (about 211 m at 21 cm). The simplest possible gridding scheme was employed, namely simply to add together the values located in each UV pixel. The simplification is acceptible because there are no external sources which can be aliased back into the field; and it is desirable because it allows one to treat the weights on a pixel-by-pixel basis. Gridding produces two separate 'UV-plane' images: one comprising the data values (just noise in the present case), a second for which the value at a given pixel is equal to the number of visibilities which fall within that pixel. The first of these yields, when inverse Fourier transformed, the dirty image; the second yields the dirty beam. Both are divided, after back-transforming, by the value of the [0,0] pixel[2] of the raw dirty beam, so that the correct 'janskys per beam' units of the dirty image are preserved.

Three weighting schemes were employed in general, although not all weighting schemes were applied to all three layout combinations. The weight schemes were natural, tapered

[1] This comes close to the zenith for MeerKAT but zenith avoidance and parallactic angle rotation were neglected in the present study.

[2] All Fourier operations are assumed to involve images which have been cyclically shifted such that the most significant pixel is at position [0,0], rather than at the centre of the image, as makes for more comfortable viewing.

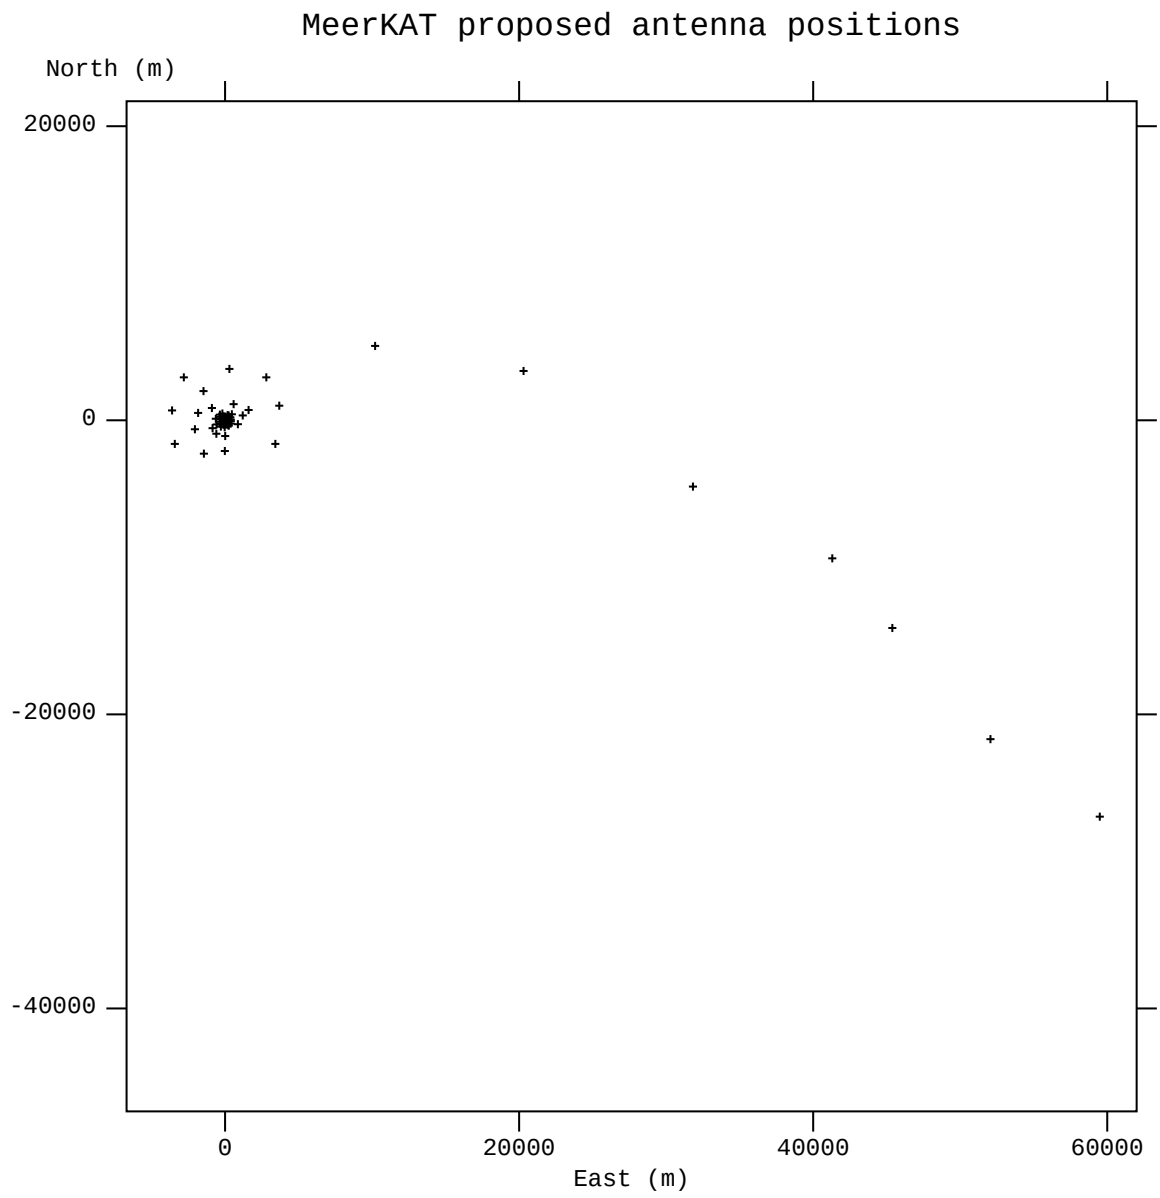


Figure 1: A schematic showing the currently proposed positions of the antennas for the extended (i.e. core + spur) MeerKAT array.

uniform, and a third which I have named Feathered Tapered Uniform (FTU - see section 2.3.2).

All weight schemes involved the creation of a weight image. The raw gridded UV data and UV beam were multiplied by this weight image before back-transforming. For natural weighting, the weight image was everywhere set to unity. For any variant of uniform weighting, the starting weight image is the reciprocal of the number of visibilities at each pixel, except where there are none, at which the weight was set to zero. Additional tapering or feathering was then applied to the raw weight image.

Gaussian tapers were used exclusively. The values used are given in table 1. The relation between Gaussian widths in the UV and sky planes is

$$\sigma_{\rm sky} = \frac{N}{2\pi\sigma_{\rm UV}}$$

where both $\sigma_{\rm sky}$ and $\sigma_{\rm UV}$ are in image pixels, and N is the size of the image in that direction.

Table 1: Gaussian taper parameters.

Array	$\sigma_{\rm sky}$ (arcsec)	$\sigma_{\rm UV}$ (kλ)
64	1.5	21.9
7	0.5	65.7
71 X[a]	0.18	182
71 Y[a]	0.32	103

[a] The Gaussian was then rotated clockwise by 35 degrees.

2.3.1 Heterogeneous image construction and its discontents

In the case of tapered uniform weighting, because the noise image and the noise-free source image are created separately, some care has to be taken to ensure that they share the same amplitude scale. The starting points are the noise dirty image and the dirty beam, both created via pure, untapered uniform weighting. The algorithm employed to arrive at the correctly tapered noise image, source image and dirty beam is as follows.

1. Convolve the raw dirty beam by the Fourier inversion of the desired tapering function. Note that this 'taper beam' must first be scaled such that its value at the [0,0] pixel is unity. Call the result B.
2. Measure the value of the [0,0] pixel of the result. Call this value α.
3. Construct a source image by painting on the appropriate amplitude of the taper beam at each source location.
4. Convolve this with the raw dirty beam. Call the result $I_{\rm src}$.
5. Convolve the raw noise dirty image by the taper beam. Call the result $I_{\rm noise}$.
6. Divide B, $I_{\rm src}$ and $I_{\rm noise}$ by α.
7. Finally, add the now properly scaled $I_{\rm src}$ and $I_{\rm noise}$ to give the final dirty image.

2.3.2 Feathered Tapered Uniform (FTU) weighting

The distribution of baselines resulting purely from the core antennas is remarkably uniform and symmetric; sadly the same cannot be said of the spur baselines. Since these result from antennas which lie broadly speaking along a single line, directed somewhat south of east from the core, for all sky locations except in the extreme south these generate visibilities which lie within an elliptical locus. Matters are even worse if a 24 hour observation is impossible since in this case the ellipsoidal regions are truncated. A fair sample is shown (naturally

weighted) in the top left panel of figure 2. The upshot of such an asymmetrical coverage of the UV plane is a strangely-shaped, highly non-Gaussian dirty beam. (The apparently highly symmetric nature of the naturally-weighted beam in the bottom left panel of figure 2 seems to belie this. However this is because the angular scale chosen for this panel is relatively coarse. At this scale the dominant contribution to the beam is from the core baselines, which are highly symmetric. A better idea of the shape of the spur beam can be gained from row 1 of figure 5.)

At this point it is of value to make a distinction in the taxonomy of beam problems between the shape of the inner beam and the amplitude and extent of beam far sidelobes. Previously the tendency has been to consider these together, since the same remedy has been applied to amend both: that is, tapering. Provided the distribution of baselines is reasonably uniform, without a large hole in the centre, which is certainly the case for wide-band frequency-synthesis MeerKAT observations, then severe enough tapering can provide almost any degree of beam improvement of both varieties. The downside is the cost in resolution (due to the shrinkage in equivalent width of the weighted baseline distribution) and possibly also sensitivity (if a significant portion of the baselines are down-weighted).

As pointed out by Briggs [1995], however, the far sidelobes of the dirty beam are mostly generated by sudden steps or jumps in the density of baselines in the UV plane. Uniform weighting eradicates internal steps in the density function, at the cost usually of inflating the relative height of the final step where the density falls to zero. Tapering can lower this again but at the costs already described. Briggs suggested an alternate scheme (subsequently implemented for example via the UVBOX and UVBXFN parameters in the AIPS task IMAGR) which involves a crude blurring or smoothing of the density function. In 2007 I described (slide 26 of Stewart [2007]) an alternative scheme for feathering out steps in the density of baselines. This scheme involves an iterative minimization of the far sidelobes directly by transforming them to the UV plane and removing some fraction of the result from the weight image. The method can fail to converge if care is not taken with gain and masking, but it is capable of producing weight images which are beautifully smooth, yet which do not sacrifice many of the baselines. Such an example is shown in the top right panel of figure 2. The beams corresponding to the respective weight images are shown in the lower two panels. The FTU beam can be seen to be very compact (because the weights don't begin to fall off until near the natural edge). Although its inner beam shape is somewhat complicated, it is dominated by a Gaussian centre because of the moderate degree of taper applied after the feathering was calculated. For faint sources, all one sees is this centre. And distant sidelobes are very tenuous indeed.

2.3.3 RMS noise in the dirty images

The calculation of noise amplitude in the dirty image falls naturally into two parts: the calculation for natural weighting (which can be shown to provide the lowest possible image noise RMS), and the increase in noise RMS due to any departure from natural weighting. The former is just given by equation 9-21 of Wroble & Walker [1999]:

$$\Delta I = \frac{\Delta S}{\sqrt{N_{\mathrm{baselines}} N_{\mathrm{channels}} N_{\mathrm{timesamples}}}},$$

where ΔS is as given in equation 1. Obviously this only holds if ΔS is independent of baseline. This is a reasonable assumption for the homogeneous dishes which will comprise MeerKAT.

The increase f due to non-natural weighting is easy to calculate in the simple gridding scheme adopted here. It is given by

$$f = \frac{\sqrt{\sum \beta \sum w^2 \beta}}{\sum w \beta}$$

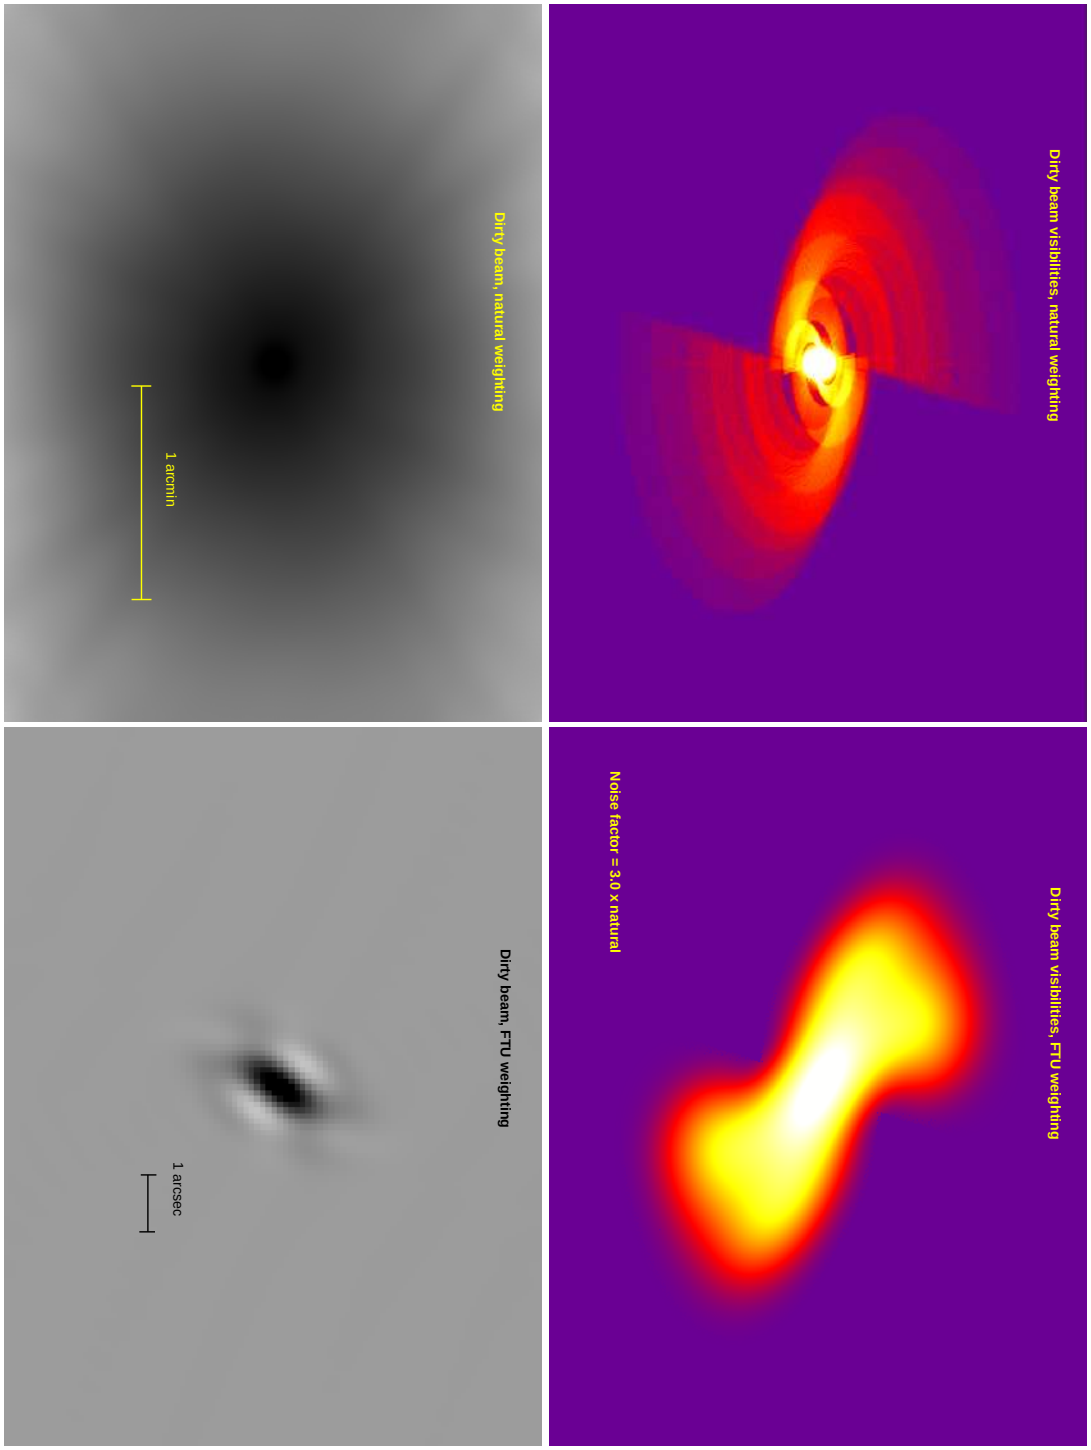


Figure 2: The upper row (with colourful colourmap) shows weighted 'beam' visibilities, that is, those visibilities which would result from a source at the phase centre, after weighting. The lower reverse-grey-scale row shows the beams which result. The left column is natural weighting; the right shows the improvements obtainable via FTU weighting, as described in section 2.3.2. Note that the angular scale of the 'natural' beam is much coarser than that for the 'FTU' beam.

where w is the UV-plane weight image and β the UV-plane beam image (which just records the number of visibilities per pixel).

In table 2 are listed the RMS noise values in μJy calculated for a 1000 hour observation. The RMS values measured directly from the images vary from these calculated values by no more than a few percent.

Table 2: Calculated image noise RMS values for 1000-hour observations (μJy).

Array	Natural	Tapered Uniform	FTU
64	0.26	2.24	-
7	0.53	-	-
64+7	0.23	1.13	0.69

3 Results and discussion

Figures 3, 4 and 5 give some indication of how the L-band radio continuum sky would look in a 1000-hour observation. The three columns of these images correspond to the three simulated skies, as described in section 2. The first row of figure 3 shows these input skies. The colour scale for this row has been chosen such that any source with more than a touch of purple at its centre has a flux density of greater than about 1 μJy. The remaining rows of these three figures show the effects of various array layouts and weighting schemes.

The 64-dish naturally-weighted view (Figure 3 row 2) has the lowest noise RMS, but this sensitivity is mostly at large spatial scales. These images are heavily confused, and one must conclude that it is probably pointless to employ natural weighting for continuum observations with MeerKAT at this depth.

Tapered uniform weighting of the 64-dish core views (Figure 4 row 1) improves the confusion markedly but at the cost of an almost 10-fold increase in 'thermal' noise because of the down-weighting of the bulk of the available baselines. Sources down to about 5 μJy appear to be visible. All but the largest sources in the SKADS image remain unresolved however.

To be fair, one could probably do better than tapered uniform weighting with the 64-dish core alone, both in reducing the noise and in obtaining a finer resolution. The advantage of the core baselines is that their distribution in the UV plane is very smooth and approximately circular over the whole range of declination angles. An FTU treatment such as was applied to the 64+7 array, as shown in row 2 of figure 5, would probably result in images which were more aesthetically pleasing at least.

Row 2 of figure 4 shows the improvement in resolution due to the long baselines available from the spur. Although the noise figure is half of that in the corresponding core-only image, visual inspection doesn't confirm any increase in sensitivity.

Row 1 of figure 5 shows a naturally-weighted image created using only those 469 baselines which include a spur dish. Compared to the core, the number of baselines is reduced by a factor of over 4; however this implies that the noise figure increases only as about root(4), i.e. 2. This is still only 0.5 μJy of noise though. Combined with the vastly better resolution of these wide-spread baselines, quite an acceptable picture of the sky is generated. Many of the extended objects in the SKADS field are resolved. Objects of flux density down to about 3 μJy or possibly even a bit lower are detected.

The final row of figure 5 represents the best image I have been able to make with all 71 dishes. The resolution is about 0.6 arcsec in the one o'clock direction and half that orthogonally. Sources of less than 3 μJy can be seen. In resolution and sensitivity it is about as good as the naturally-weighted spur baselines alone, but it is a much cleaner image and arguably of publication quality without further processing. Note that the resolution is so

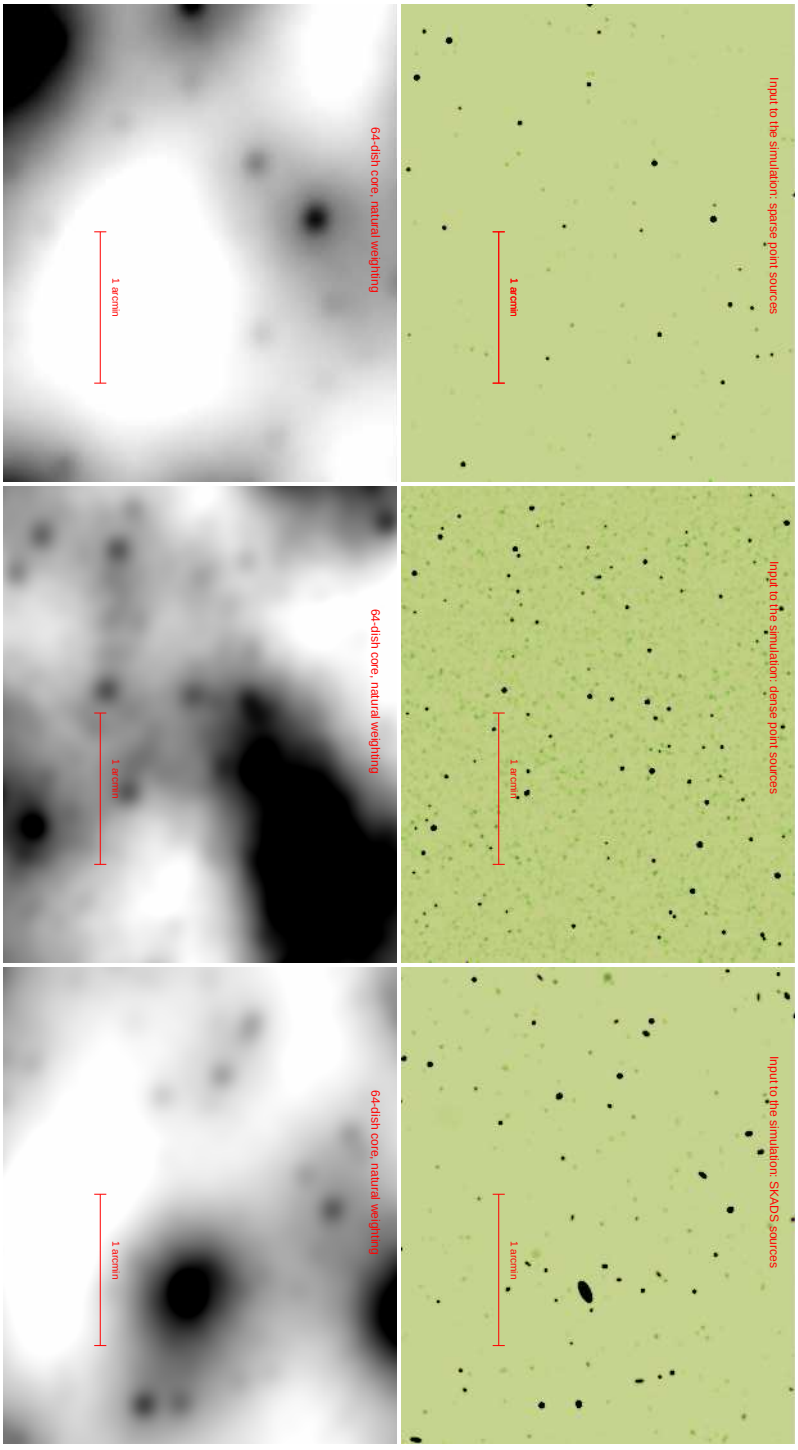


Figure 3: The top row (with greenish colourmap) shows the three simulated skies used in the present simulation. Any source with more than a touch of purple is $> 1\mu$Jy. All three images here have been convolved with an 0.5 arcsec Gaussian to enhance the visibility of compact sources. The lower, grey-scale row shows imaging of the respective skies with the core 64 dishes alone, naturally weighted.

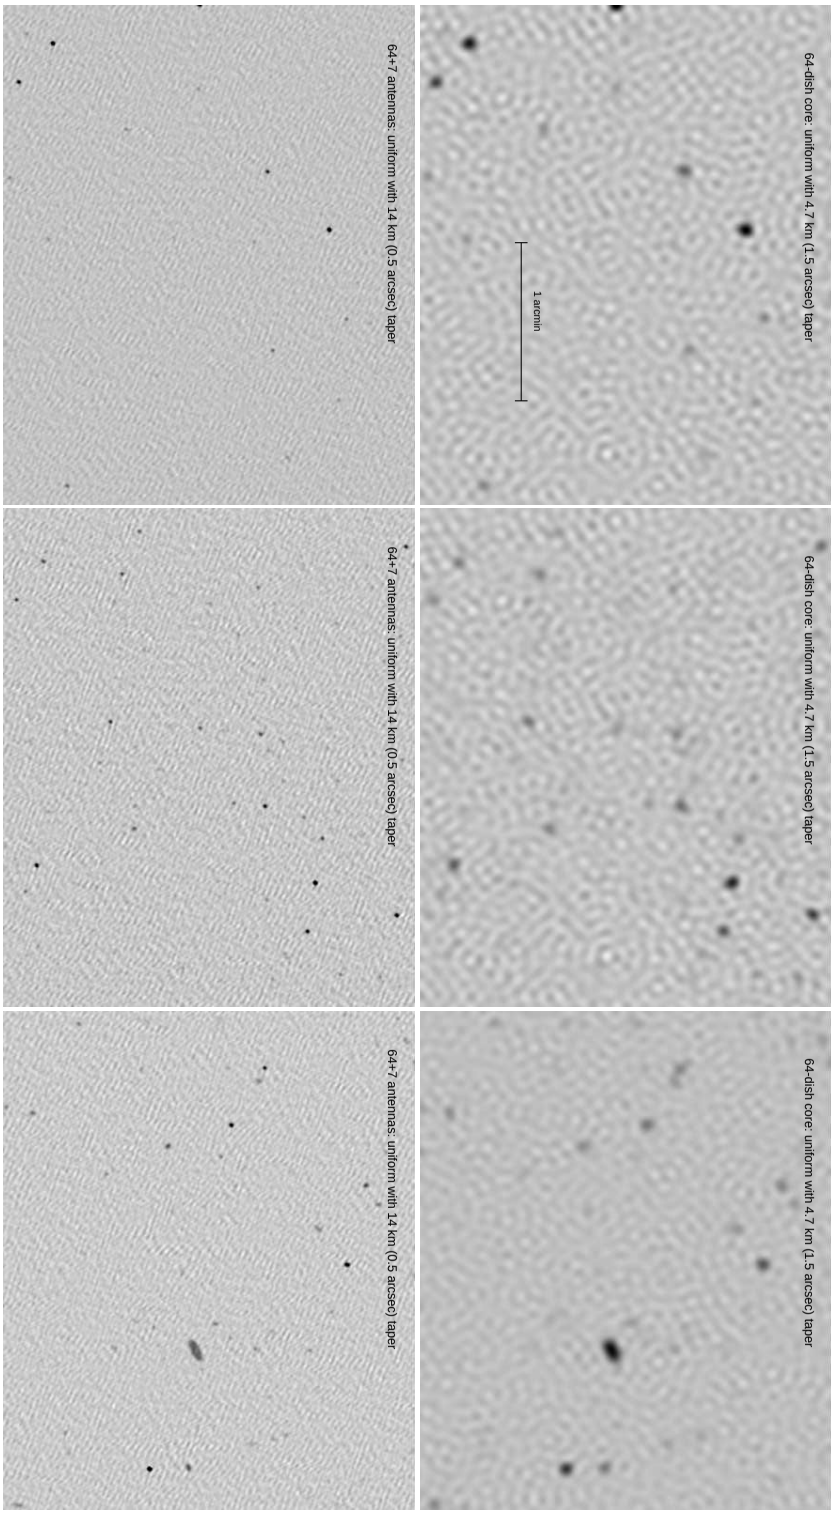


Figure 4: This figure is the second in a sequence of three showing the effect of various layouts and weighting schemes on dirty images of the three input simulated skies. The skies are shown in figure 3, row 1.

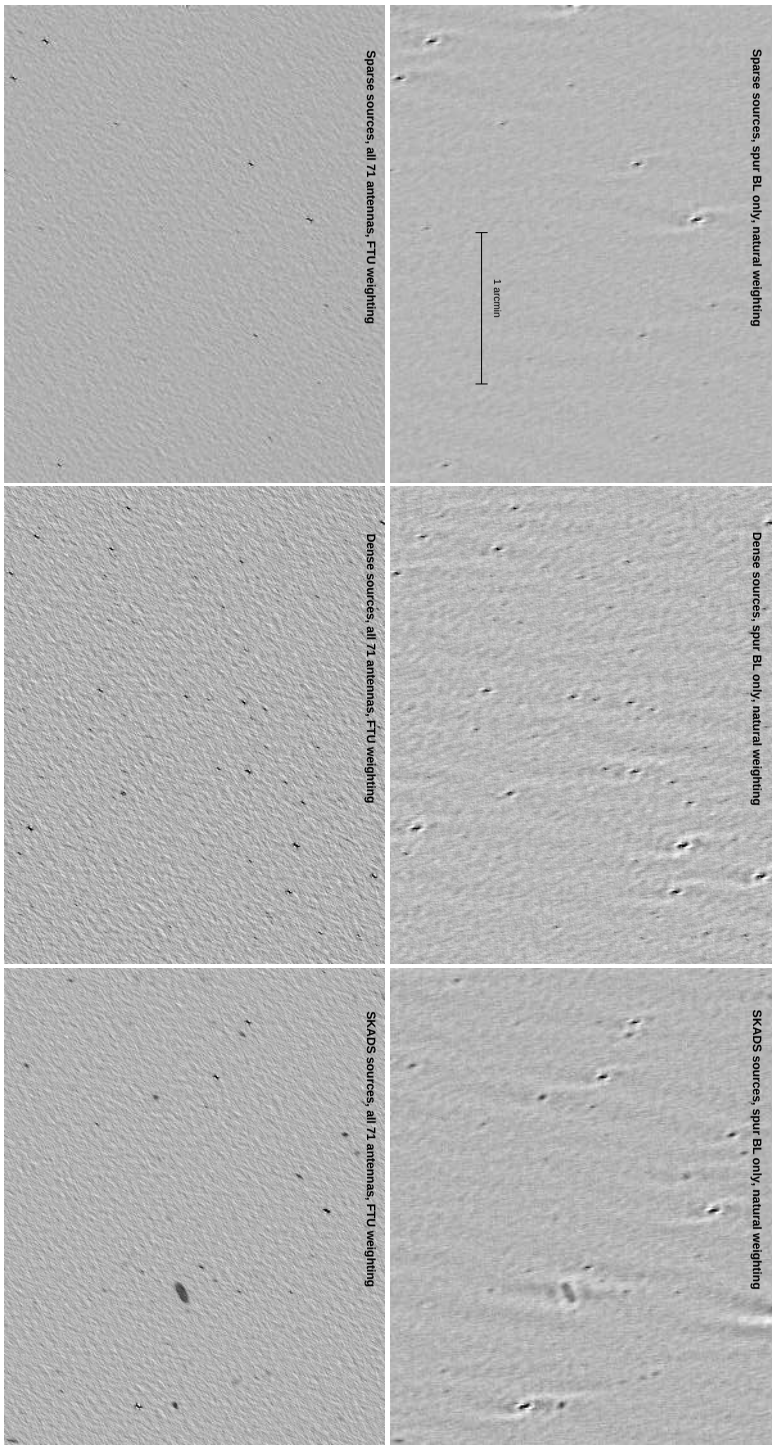


Figure 5: This figure is the final in a sequence of three showing the effect of various layouts and weighting schemes on dirty images of the three input simulated skies. The skies are shown in figure 3, row 1.

fine that compact sources, while present, do not 'jump at the eye' on this scale. Only when one zooms in to a smaller scale can the full sensitivity of this image be appreciated.

As a final comment I note that the proposed arrangement of spur dishes is probably not ideal. A greater concentration of dishes towards the distal point might well bring the naturally-weighted baseline density closer to the FTU output, which would result in a decreased noise figure. This is something which would need to be confirmed with some further simulations.

4 Conclusion

Comparison of simulations with various MeerKAT layouts shows that continuum images made using the core+spur array are significantly deeper and have 10 times the resolution of core-alone images. As shown by the SKADS image, current expectation is that continuum sources at μJy depth will have angular sizes on the order of a few arcseconds. This means quite simply that the spur can resolve such sources, but the core cannot. All knowledge of the morphology of these sources will thus be lost if only the core array is built: the continuum science which MeerKAT can do will then be reduced to something not much in advance of stamp collecting.

One objection to building the spur array has been that the poorly-shaped beam which will be the inevitable result of such a constrained set of baseline orientations will in effect render the project not worth pursuing. However as is shown in the present report, it is possible to greatly ameliorate such beams with judicious weighting, such that the output image is aesthetically pleasing and arguably already of publication quality. More accurately said, there will undoubtably be many more difficulties to overcome in producing such high-quality images from MeerKAT data, so actually what one should conclude here is that the spur beam shape can be pushed down to well below the level at which it is the worst of the problems we have to deal with.

Acknowledgements

Many thanks for help and advice to Brad Frank, Kurt van der Heyden and Matt Jarvis.

References

[1995] Briggs, D. S. 1995, "High Fidelity Deconvolution of Moderately Resolved Sources", PhD Thesis, New Mexico Institute of Mining and Technology. Available from http://www.aoc.nrao.edu/dissertations/dbriggs/

[1984] Condon J. J. 1984, Ap J 287, 461-474

[2006] Fomalont E. B., Kellermann K. I., Cowie L. L., Capak P., Barger A. J., Partridge R. B., Windhorst R. A. & Richards E. A. 2006, Ap J S 167, 103-160

[2008] Gervasi M., Tartari A., Zannoni M., Boella G. & Sironi G. 2008, Ap J 682, 223-230

[2005] Huynh M. T., Jackson C. A., Norris R. & Prandoni I., 2005, Astron. J 130, 1373-1388

[2008] Kellermann K. I., Fomalont E. B., Mainieri V., Padovani P., Rosati P., Shaver, P., Tozzi P. & Miller N., 2008, Ap J S 179, 71-94

[2007] Stewart I. M. 2007, "Wide-Band Imaging", European Radio Interferometry School, Sept 2007, http://www.mpifr-bonn.mpg.de/div/eris/talks/Stewart_MultiNu.pdf

[2008] Wilman R. J., Miller L., Jarvis M. J., Mauch T., Levrier F., Abdalla F. B., Rawlings S., Kloĉkner H.-R., Obreschkow D., Olteanu D. & Young S. 2008, MNRAS 388, 1335-1348

[1999] Wrobel J. M. & Walker R. C. 1999, "Synthesis Imaging in Radio Astronomy II", ASP Conf. Ser. 180, 171-186