

MeerKAT time smearing simulation: draft 1

Laura Richter

August 25, 2014

1 Introduction

The MeerKAT correlator will produce a stream of output visibilities with a period of 0.5 seconds. This data stream will be further time-averaged in the ingestor. Longer averaging times in the ingestor reduce the data rate and hence reduce hardware costs. However, longer integration times can also impact the science data products through time smearing. This document describes a series of simulations performed to inform the choice of optimal averaging time.

Time smearing arises because averaging data in time corresponds to averaging data in uv-space. Distance of a source from the field centre translates into a phase slope in uv-space, with larger slope for larger distance from the centre. If uv-space averaging is performed over a substantial fraction of the phase slope of a source before imaging, then the source amplitude in the image will reduce and the source structure will be distorted.

The precise nature of the amplitude reduction and distortion is a function of the uv-coverage, the position of the source in the field, and the shape of the source. Analytical approximations have been performed, which support an integration time of 2 seconds for MeerKAT (Schwardt, 2014). These simulations provide an additional means to verify the analytical result.

2 Method

The effect of time smearing on MeerKAT images was investigated by simulating time smearing for multiple fields, using a simple sky model, and averaging the visibility data over successively larger intervals before imaging.

The sky model consisted of a pattern of point sources extending to approximately twice the primary beam radius. No primary beam effects were included in the model, so that reduction in amplitude with averaging period can be attributed solely to time smearing. It was found in initial attempts of the simulation that wide-field effects (which also increase with distance from the field centre) could not be neglected. This was subsequently accounted for by imaging with casapy wide-field gridding mode, using multiple w-projection planes, as described below.

The simulation was performed for three fields, chosen from the MeerKAT Large Survey Project fields. The fields used are given in Table 1, and labelled F1 to F3. For each field the imaging was performed with both natural and uniform weighting.

The sky modelling, averaging and imaging were all performed in casapy. For each of the simulations the following steps were followed:

2.1 Template measurement set creation

A MeerKAT measurement set for each field was created using the LOFAR software suite task makems. Each measurement set consisted of a 12 hour observation with 1s integration time, for a single channel of width 26.1 kHz, at start frequency 1.38 GHz, with arbitrary start time 15:00 01/01/2014.

2.2 Simulated sky

The measurement set was filled with a simulated \times -shaped distribution of 1 Jy point sources, using casapy.

Simulation label	Survey field	Right Ascension	Declination
F1	Laduma / Mightee2	03:32:00	-27:00:00
F2	Mightee4	10:00:00	02:00:00
F3	Fornax	03:30:00	-35:00:00

Table 1: List of simulated field names and positions.

The MeerKAT dish diameter is 13.5m, giving a primary beam FWHM of ~ 55 arcmin at 1.38GHz. The cross distribution was designed to cover approximately twice this extent, as we expect observers to require accurate images out to at least the primary beam FWHM. The source positions were incremented 5 arcmin \times 5 arcmin in Right Ascension and Declination from the centre, so that the radial distances of the sources from the centre were $\{0, 7.07, 14.14, 21.21, 28.28, 35.36, 42.43, 49.50, 56.57\}$.

2.3 Time averaging and imaging

The measurement set was time averaged in one second increments up to 40 seconds. The resulting data sets, from $\Delta t = 1$ s integration time to $\Delta t = 40$ s integration time, were then inverted, without CLEANing. The image size was 8000 by 8000 pixels of 0.75 arcseconds each in Right Ascension and Declination. The cell size was chosen to well sample the uniform imaging restoring beam, which is $\sim 7 \times 7$ arcseconds (Table 2). The images were not CLEANed to avoid introducing additional effects that may vary between averaged data sets due to source amplitude differences (and hence different CLEAN implementations).

Wide-field grid mode was used, with multiple w-projection planes. The number of planes necessary is a matter of active research (“We are working on formula”, Myers, 2014). The number of planes necessary for this simulation was therefore determined empirically, using the F1 Laduma/Mightee2 field, imaged with uniform weighting, and testing the number of planes required to reduce wide-field effects on the most exterior sources to an acceptable level. This is described in Appendix A. The final number of w-projection planes was chosen to be 72.

2.4 Source amplitude measurement

The amplitude of each of the simulated sources across all of the time-averaged images was extracted using an automated method: The peak position of each source in the original 1s image was determined by finding islands above a threshold of 0.85 Jy/beam in the 1s map, and then finding the peak pixel position of that island. The amplitudes in all of the maps were then taken as the peak in a 20×20 arcsecond box centred on the 1s peak position. The peaks for each Δt were found in a box, rather than simply at the 1s peak position, to account for possible variation in peak position due to smearing. The box size was chosen to be approximately the natural weighting restoring beam FWHM.

The quantity of interest in these simulations is the decrease in source amplitude as a function of increasing averaging time. Two systematic effects on the amplitude must be accounted for, before the amplitude decrease can be attributed to time smearing alone: sidelobes and the w-term. The use of 72 w-projection planes reduces the amplitude variation across the field to a $\sim 3\%$, dependent on the distance of the source from the phase centre (Appendix A). Sidelobe noise is a factor because the images are unCLEANed, and the source maxima may vary from source to source depending on the level of sidelobe noise beneath the source. This amplitude effect is not directly dependent on distance from the centre: it depends on the placing of the sources relative to the sidelobes of other nearby sources

To remove the effect of residual w-projection amplitude variation and sidelobe noise amplitude variation, we consider the percentage relative amplitude reduction, rather than absolute amplitude value. For each source the 1s amplitude was subtracted from the amplitudes at subsequent averaging times. This method is discussed in more detail in Section 4.

Simulation label	Survey field	Natural	Uniform
F1	Laduma / Mightee2	14.95×12.59	7.39×6.61
F2	Mightee4	17.52×12.64	8.91×7.45
F3	Fornax	14.69×12.48	7.17×6.48

Table 2: Restoring beams for the natural and uniform weighted images of each field. The restoring beam dimensions are listed as Right Ascension \times Declination in units or arcseconds in each case.

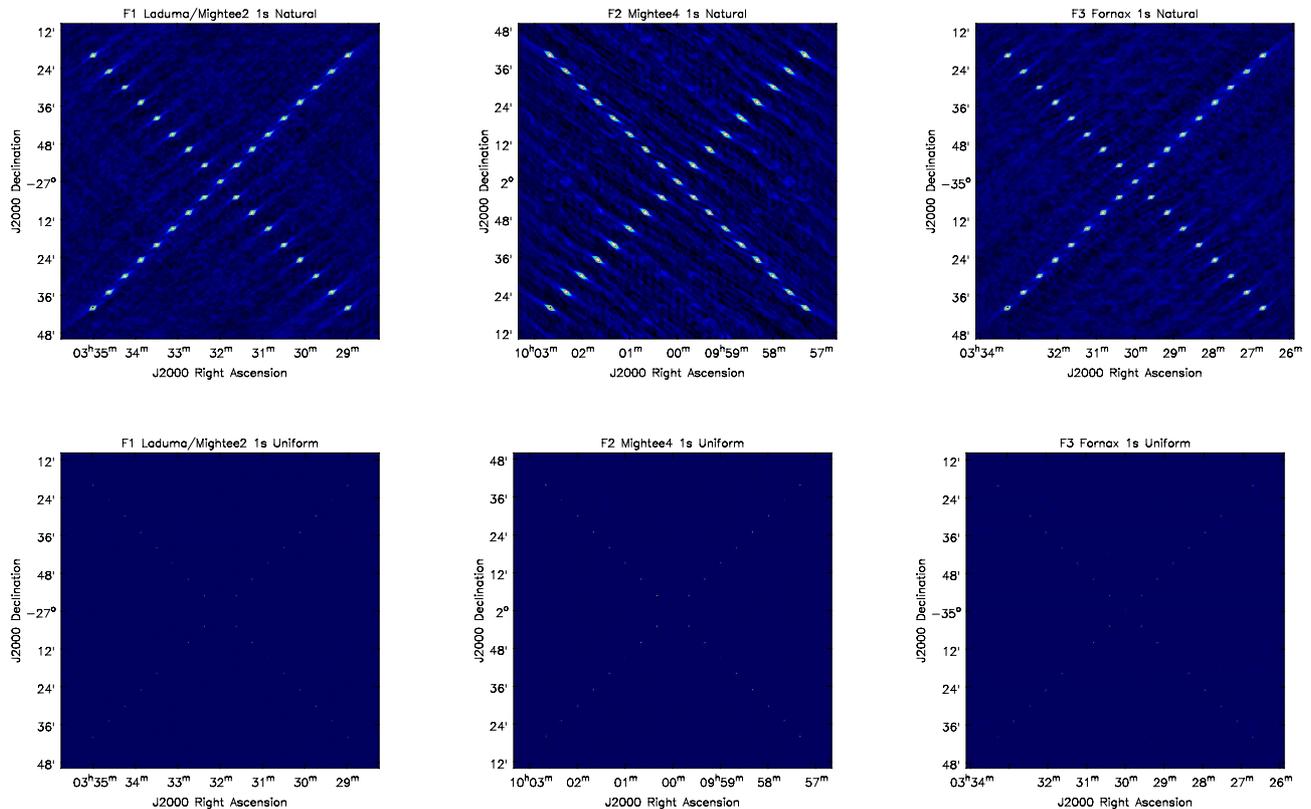


Figure 1: Inverted images of the $\Delta t = 1$ s data sets for the three fields, for natural weighting (top) and uniform weighting (bottom). The colour scale ranges are $[-0.03, 0.3]$ for all of the images.

3 Results

The inverted maps of the $\Delta t = 1$ s data sets are shown in Figure 1 for the three fields, showing natural and uniform weighting. Table 2 lists the restoring beam of each $\Delta t = 1$ s image. Figure 2 illustrates the radial distance dependent reduction in amplitude between the $\Delta t = 1$ s and 40s images for the F1 Laduma/Mightee2 simulation.

Figure 3 shows the uv-coverage for each of the fields, as well as a uv-coverage histogram of each to illustrate the density of uv-coverage. The uv-coverage histogram is in units of projected radial baseline length $B = \text{uvdistance} \times \lambda$, for comparison with later figures.

4 Analysis

4.1 Systematics

As discussed in Section 2.4, although the sources were added to the visibility data with amplitudes of 1 Jy, their amplitudes in the uncleaned maps will be less than this due to residual w-term effects, and due to a portion of their flux being located in sidelobes, and due to the variables sidelobe noise across the map.

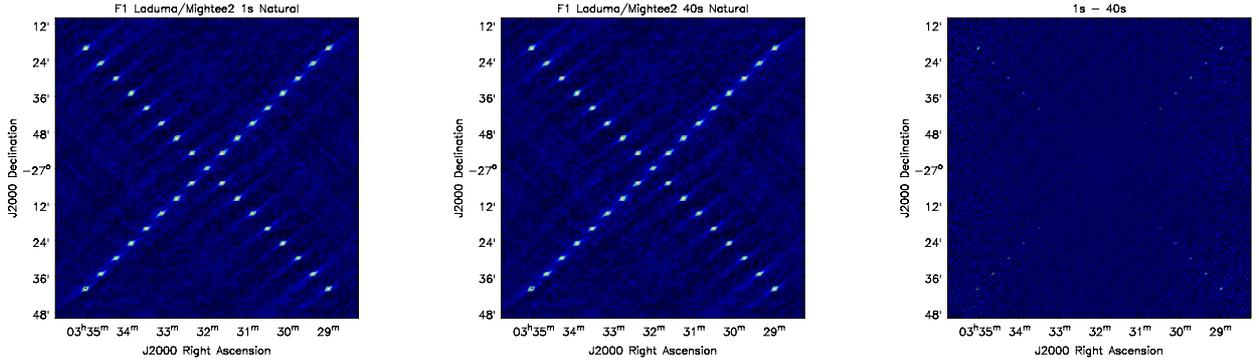


Figure 2: From left to right: inverted image of the F1 $\Delta t = 1\text{s}$ data set, inverted image of the F1 $\Delta t = 40\text{s}$ data set, both imaged with natural weighting, and an image of the difference between the two. The colour scale ranges are $[-0.03, 0.3]$ for the first two images, and $[-0.001, 0.01]$ for the difference image.

Figure 4 (left) illustrates this for the F1 Laduma/Mightee2 simulation. For each source, the peak amplitudes are plotted versus time averaging, up to $\Delta t = 20\text{s}$ to zoom in on the amplitude offsets. Sources at the same distance from the centre are grouped by linestyle. There are four sources at each distance, due to the simulated source configuration, except for the centre source. There are clear amplitude offsets between curves at a the same distance. This is most evident for the natural weighting images.

Figure 4 (right) illustrates the sidelobe noise variation across the image. The mean and standard deviation of the map were calculated in blocks of 100×400 pixels in Right Ascension and Declination in a strip across the centre of the image. Measurements over the inner two blocks are missing because they fall over the central source. The measured noise is purely sidelobe noise, as thermal noise was not included in the simulation. The increase in sidelobe noise towards the centre of the image, where the sources are closer together, is evident.

Several important points are illustrated by these figures. Firstly, as seen in the left-hand plots, the most extreme offsets from 1 Jy are for the sources at and closest to the centre, where the interaction between the larger amplitude inner sidelobes is most extreme, due to the proximity of the sources. This is shown by the higher standard deviation in the centre of the maps, visible in the right-hand plot.

Secondly, the standard deviation plots show that the sidelobe noise across the map is considerably larger for the natural weighted images, about double as high in the centre of the image. The effects of this are seen in the amplitude versus time averaging plots on the left, where the amplitude decrease between different source distances is not a function of distance (as it is for w-term effects) and is notably low for the central source (an effect not seen for the uniform map, where the sidelobes are more compact).

Thirdly, there is a clear difference between the amplitude offset of each of the four same-distance sources in the natural images. The reasons for this are clear from Figure 1: depending on the angle of the sidelobes, each four same-distance sources will fall in different locations relative to the sidelobes of adjacent sources. Some of the sources may fall on positive sidelobe regions, and consequently have increased amplitude, and some not.

Fourthly, although there is this amplitude offset between the different sources at each distance, the form of the amplitude fall-off with time averaging time appears very similar across same-distance sources. This implies that the amplitude offset due to sidelobe noise (and the w-term) is approximately the same across the time averaged data sets.

Would not expect it to be completely identical, as the source amplitudes change with time-averaging, so the sidelobe levels will vary accordingly. However, this variation will be small: of the order of the difference in amplitude (less than a couple of percent up to a maximum of of about 20%) times the sidelobe level of an adjacent source at the location of the source in question ($< 1\%$).

In order to remove the effect of the variable sidelobe noise for each source, as well as residual w-term effects, and to distill the effects of time smearing, we consider the amplitude reduction for each source

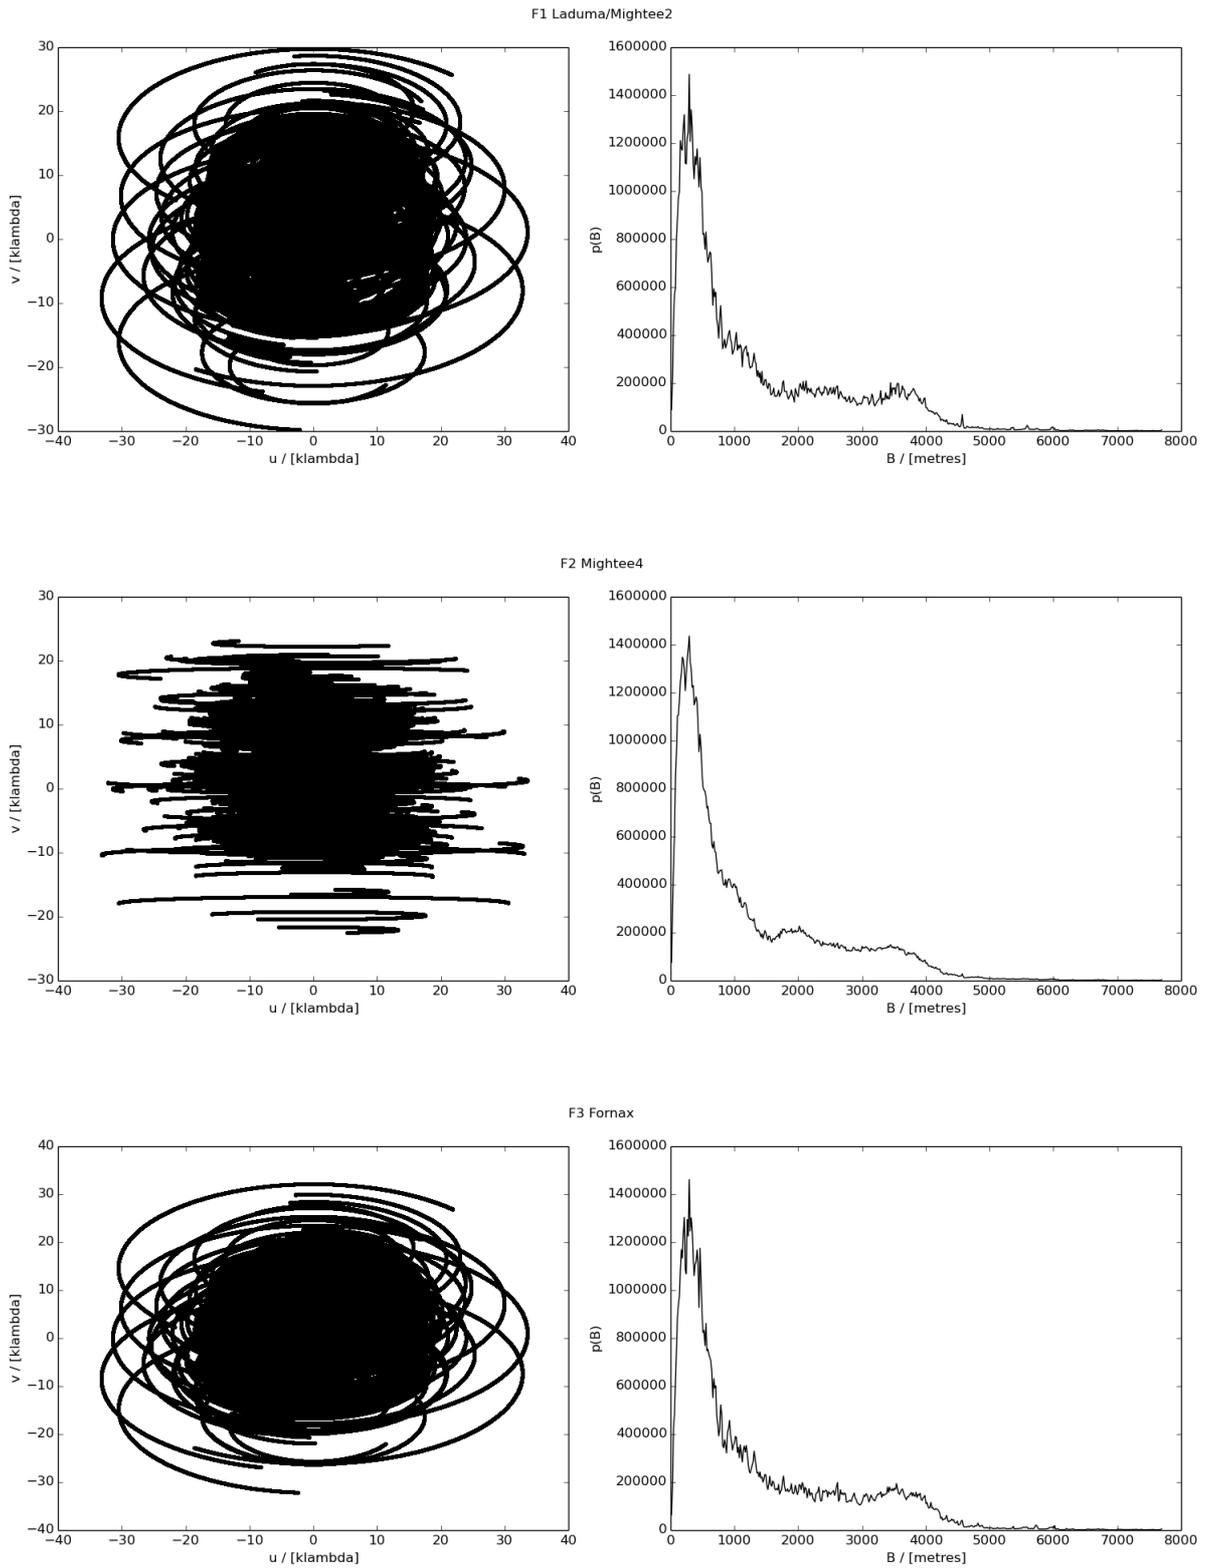


Figure 3: For each of the simulated fields, plots of uv-coverage in units of wavelength (left), and histograms of the uv-coverage density $p(B)$ as a function of baseline length B (right).

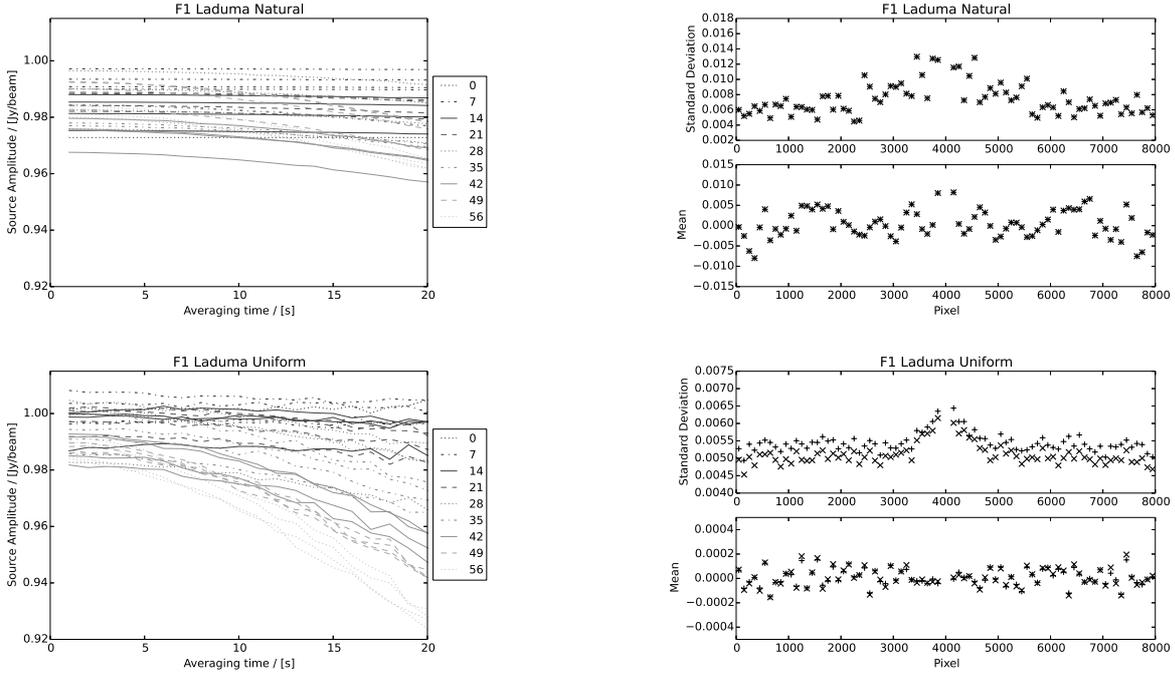


Figure 4: Plots of the F1 Laduma/Mightee2 field, imaged using natural weighting (top) and uniform weighting (bottom). (Left) Plot of amplitude versus time for each of the sources. The linestyles are grouped by distance from the centre, as shown in the legend on the right, which is in units of arcminutes. (Right) Plot of mean and standard deviation in boxes of pixel dimension 100×400 pixels across the centre of the image, as described in the text. The means and standard deviations are plotted for the $\Delta t = 1$ s image (marker '+') and the $\Delta t = 40$ s image (marker 'x').

relative to its $\Delta t = 1$ s amplitude. Figures 5 and 6 show the percentage decrease with averaging time for each field and source distance, where the percentage decrease is calculated as

$$\text{Amplitude reduction} = 100 * (A_{\Delta t=t_i} - A_{\Delta t=1})/A_{\Delta t=1} \quad (1)$$

for each amplitude A at time average $\Delta t = t_i$.

In order to quantify possible variation in sidelobe noise amplitude offsets between different time-averaged images, the figures show the average percentage decrease between the same-distance sources, as well as a grey window around each average curve, which shows the the maximum to minimum percentage decrease over that set of sources. This grey window provides an estimate of two errors: i) pixellisation variations in measurements of the peak amplitude (which are larger for the uniform maps with smaller restoring beam), as well as ii) changes in the sidelobe noise amplitude offset with averaging time. The four sources sample difference sidelobe noise regions, based on the relative location of the source to nearby sidelobes. So the effect of changes in the sidelobe noise amplitude offset will vary between the sources, and the range of the change provides an indication of the scale of the effect.

4.2 Analytic models

Figures 5 and 6 also show curves for one of the theoretical time smearing relations presented in Schwarzt (2014):

$$\text{Average reduction} = 1 - 0.884 \times 10^{-9} (B/D)^2 (\theta_p \Delta t)^2 \quad (2)$$

for dish diameter D , baseline length B , and θ_p the distance of the source from the centre, in units of primary beam FWHM (Schwarzt, 2014). This relation was formulated for circular uv-coverage and Gaussian tapering of the uv-data. The simulated data are not well described by Equation 2, as seen in the figures. The discrepancy is discussed in the next section.

An intrinsic flaw in the simulations presented here is that the reported amplitude reduction measurements are made relative to the $\Delta t = 1$ amplitude. They are not absolute measurements of amplitude

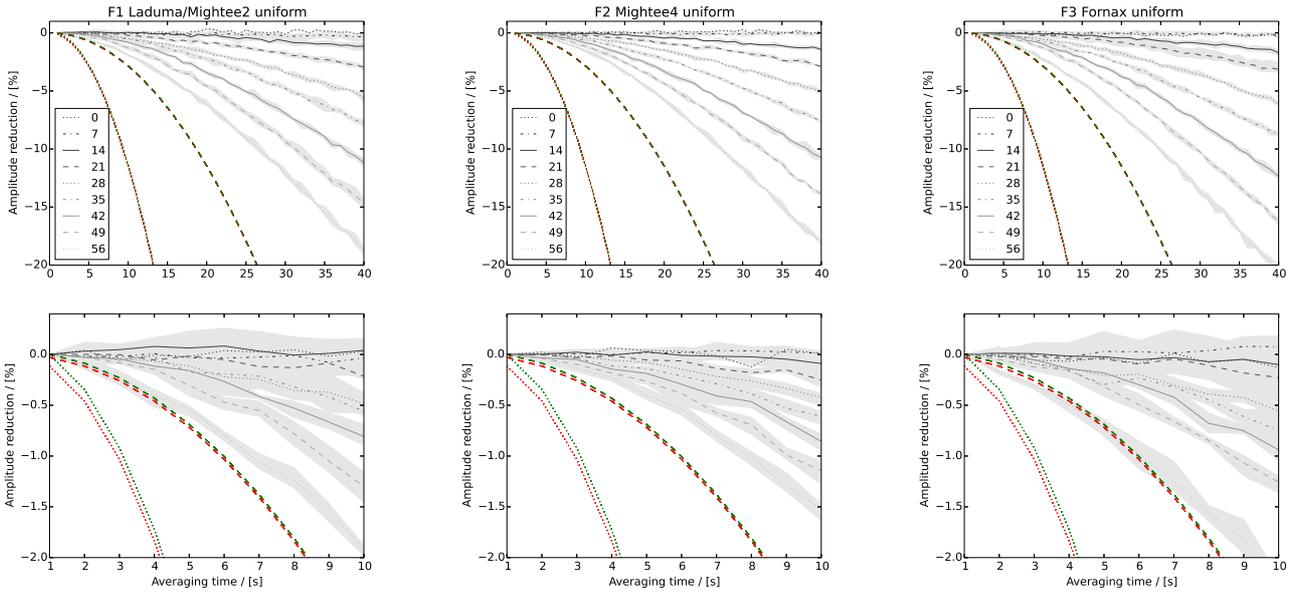


Figure 5: Plots of amplitude reduction versus time averaging for each of the fields, imaged using uniform weighting. The curves show average relative offset across all sources at a particular distance, as labelled in the legend. The grey region around each curve shows the range of values over all of the sources at each distance. The lower plots zoom in on the averaging times up to 10s. The red and green curves show theoretical percentage amplitude reduction, using a baseline of 8000 km (dotted) and 4000 km (dashed), as described in the text.

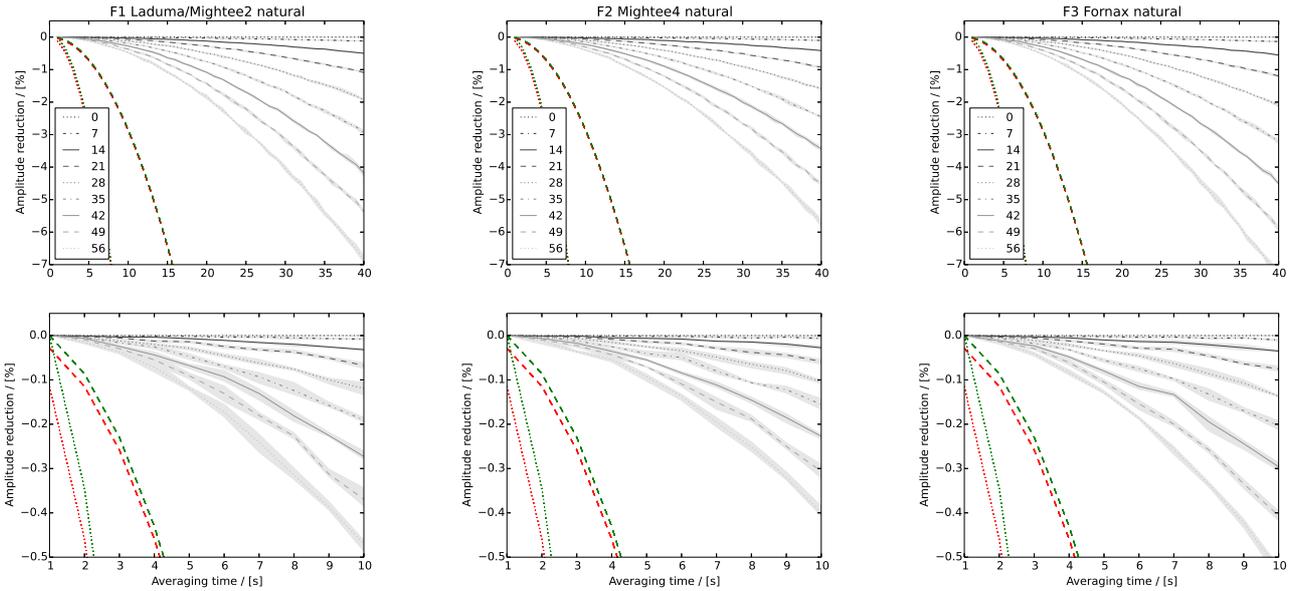


Figure 6: Plots of amplitude reduction versus time averaging for each of the fields, imaged using natural weighting, as for Figure 5.

Distance [arcmin]	F1		F2		F3	
	uniform	natural	uniform	natural	uniform	natural
7		1434		1288		1496
14		1409		1306		1504
21	2517	1387	2496	1287	2483	1488
28	2719	1400	2533	1288	2883	1467
35	2398	1404	2498	1288	2739	1478
42	2431	1403	2452	1279	2645	1474
49	2547	1391	2447	1262	2620	1466
56	2675	1388	2467	1253	2893	1453
Average B_{ave}	2548	1402	2482	1282	2710	1478
B_{max}		7696		7696		7696
$B_{max}/3$		2565		2565		2565
B_{mean}		1309		1221		1329

Table 3: Fitted average baseline lengths B_{ave} , for each field and weighting, and the average over all source distances. For each field, the maximum baseline distance B_{max} , one third of the maximum, and the mean B_{mean} are listed, as calculated from the uv-coverage.

loss due to time smearing. However, because the $\Delta t = 1$ s amplitude loss is small (a fraction of the loss observed between $\Delta t = 1$ s and $\Delta t = 2$ s in this simulation), we can consider the relative amplitude reductions presented here to be approximately the same as the absolute amplitude reduction. We confirm that this is reasonable, to within the uncertainty of the simulation, by plotting each theoretical curve on Figures 5 and 6, and subsequent figures, in red for the absolute amplitude loss and in green as amplitude loss relative to $\Delta t = 1$ s, as defined in Equation 1. As can be seen in the figures, the difference between these two curves is smaller than the range of the grey window representing the uncertainties.

Equation 2 was derived from a more general relation from Bridle & Schwab (1989), given by

$$\overline{R_{\Delta t}} \approx 1 - (\pi^2 w_e^2 / 12) \Theta^2 \Delta t^2 B_{ave}^2 / \lambda^2 \quad (3)$$

where w_e is the Earth's rotational velocity, $w_e = 7.27 \times 10^{-5}$ radians.sec⁻¹, Θ is the distance of the source from the phase centre, λ is the wavelength of the observation and $B_{ave}^2 = \overline{B_x^2 + B_y^2}$ is the average of the squared lengths of the projected baseline vectors.

We fit Equation 3 to each of the curves presented in Figures 5 and 6, fitting for the average baseline length B_{ave} . Bridle & Schwab (1989) note that Equation 3 is only valid for small intensity losses, so we limit the fit to times $\Delta t = 1$ to 15s. The fit was necessarily performed to the relative percentage data, so Equation 3 was normalised relative to the $\Delta t = 1$ s value as described in Equation 1 for the fit. The results of these fits are plotted over the original curves in Figure 7, where red and green denote absolute percentage amplitude loss (the fitted B_{ave} applied to Equation 3 directly) and relative percentage loss (the fitted B_{ave} applied to Equation 3 then normalised following Equation 1).

The fitted values of B_{ave} are listed in Table 3, along with several baseline length quantities measured from the uv-coverage plotted in Figure 3: the maximum value of the radial projected baseline distance, one third of the maximum, and the mean. Fits of the 7 and 14 arcmin sources are absent from the uniform image weighting simulation as the amplitude reductions in the $\Delta t = 1$ to 15s region were too small, relative to the noise, for these two cases, resulting in poor fits.

5 Discussion

We note several features of the results presented above:

i) The time smearing effect is less extreme for natural weighting images than uniform weighting images. This is expected, as the natural images weight the inner uv-data higher, minimising the effect of longer baselines where the time averaging effect on amplitude is most extreme.

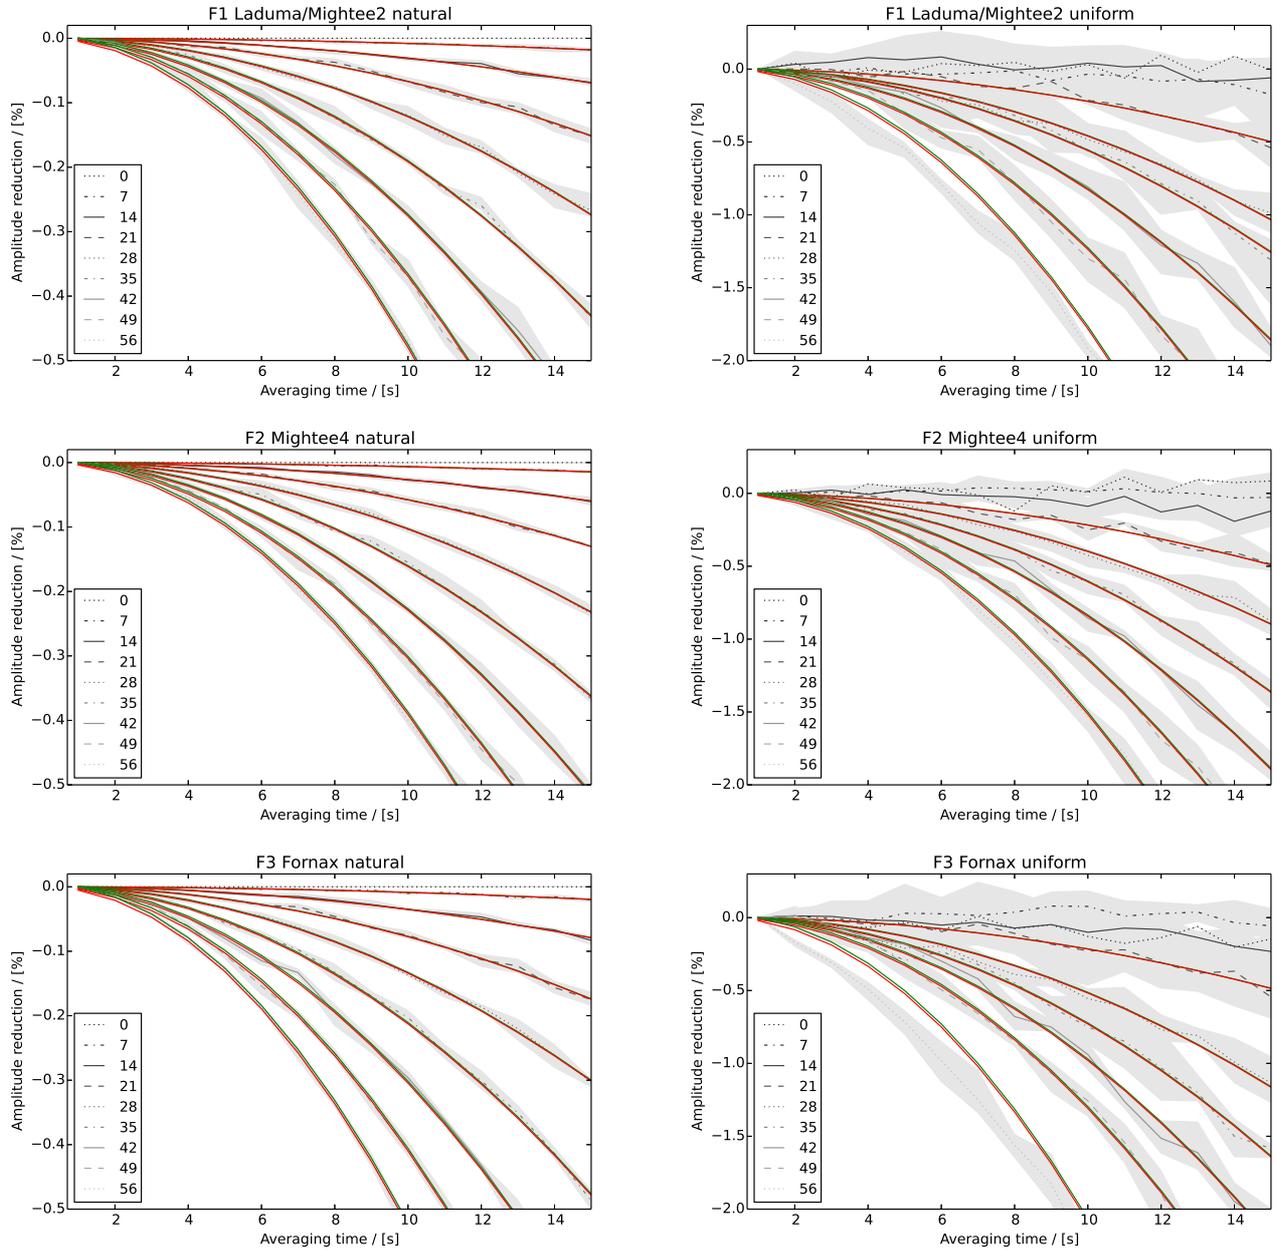


Figure 7: Plots of relative amplitude reduction versus time averaging for each of the fields, imaged using natural weighting (left) and uniform weighting (right). The grey and greyscale curves are repeated from Figures 5 and 6. The red and green curves are fits of the relative amplitude reduction versus time averaging data for average baseline, as described in the text.

ii) The amplitude reductions evident in Figures 5 and 6 are considerably less than those predicted by Equation 2, for a nominal maximum MeerKAT baseline of 8000 km. Taking a baseline length of 4000 km, provides a better approximation.

iii) The fits for average baseline length in Figure 7 approximate the results of the simulation reasonably well, especially for the natural weighting case. In the natural weighting case the fitted B_{ave} is close to the mean value B_{mean} . In the uniform weighting case the fitted value B_{ave} is close to $B_{max}/3$.

The baseline averaging in Equation 3 comes about because the observed amplitude reduction of a source will be due to the average effect of amplitude reduction over all baselines. So the baseline length in the equation must be averaged in some way over all baselines. This is not a simple average, however. Firstly, the amplitude reduction varies with location of the source in the sky, and hence the extent of averaging at different points in the uv-plane. This is not taken into account in a simple averaging of the radial baseline lengths. Secondly, the uv-coverage of a real interferometer is not a constant function that fills the uv-plane.

In the case of MeerKAT, the uv-coverage is heavily concentrated at lower baseline lengths (Figure Figure 3). We would consequently expect that for MeerKAT $B_{ave} < B_{max}/2$. The empirical B_{ave} values listed in Table 3 provide estimates of the averages. The use of B_{mean} for natural weighting and $B_{max}/3$ for uniform weighting may also provide a useful rule-of-thumb for the average baseline length to calculated time smearing for the MeerKAT uv-coverage.

The current intended MeerKAT ingest time averaging is 2s. The theoretical approximation of Equation 2 predicts an amplitude reduction of $\approx 0.5\%$ at a radius of ~ 60 arcmin for time averaging of 2s (Schwardt, 2014). These simulations show that for a maximum source radius of 56 arcmin, an $\approx 0.5\%$ amplitude reduction can be achieved for $\gtrsim 4$ s for uniform weighting images case, and considerably longer for natural weighting images (see Figures 5 to 7).

6 Conclusion

The primary conclusion of this simulation is that time smearing is a less significant effect than initial analytical approximations suggest, due to the concentration of MeerKAT antennas in a central core, and the consequent concentration of uv-coverage towards the centre of the distribution.

The specific choice of time averaging period may ultimately depend on the imaging parameters used by different surveys: surveys which require high resolution and image with more uniform weighting may require shorter averaging times. The acceptable amplitude reduction due to time smearing can be advised by the survey teams. The corresponding time averaging limits can be determined from simulations such as those performed here.

7 References

- Myers, S., 2014, Wide-Field Imaging I: Full-Beam Imaging & Surveys, Fourteenth Synthesis Imaging Workshop
 Perley, R., 1999, Imaging with Non-Coplanar Arrays, ASP Conference Series, 180, 391
 Bridle A.H. & Schwab F.R., 1989, Wide field imaging I: bandwidth and time average smearing, ASP Conference Series, 6, 247
 Schwardt, L., 29 May 2014, MeerKAT Time Smearing

A Appendix A: Wide-field effects

The number of w-projection planes required for the simulation described in this document was determined by repeating the simulation multiple times for the F1 Laduma/Mightee2 field, using a range of w-projection planes. The range of planes investigated was $\{12, 24, 36, 72, 108\}$, as well as a reference simulation with no wide-field correction. The full simulation described above from Section 2.1 to Section 2.4 was performed for each case, using uniform weighting to minimise the effect of sidelobe noise (discussed in Section 2.4). An island threshold of 0.5 was used in the automated amplitude extraction method, so that even sources severely effected by the w-term are measured.

The results of these simulations are shown in Figure 8. The simulation without widefield correction is shown in the top left panel of the figure. The absence of curves in this figure are due to the automated amplitude extraction method failing for these images, due to the w-term distortion in the images. A portion of the $\Delta t = 1\text{s}$ image is shown in Figure 9 to illustrate the distortion, which increases with distance from the centre.

The top panels in Figure 8 show the measured peak amplitude for each source, which drops with distance from the field centre, as expected for w-term effects. The results from the 72 and 108 plane simulations are identical: the convolution support reported by casapy is the same, except for repetitions of identical values. Presumably beyond a certain point adding extra planes has no effect, in the casapy implementation.

The lower panels show the percentage amplitude reduction for each source over time relative to the $\Delta t = 1\text{s}$ amplitude, as defined in Section 4. Note that for the simulations with number of w-projection planes 24 and above the percentage decrease is effectively the same for each simulation. This shows that the use of relative decrease effectively removes the overall residual w-term amplitude offsets.

The final number of planes to be used in the time smearing simulation was chosen to be 72, as 108 planes provide no additional benefit.

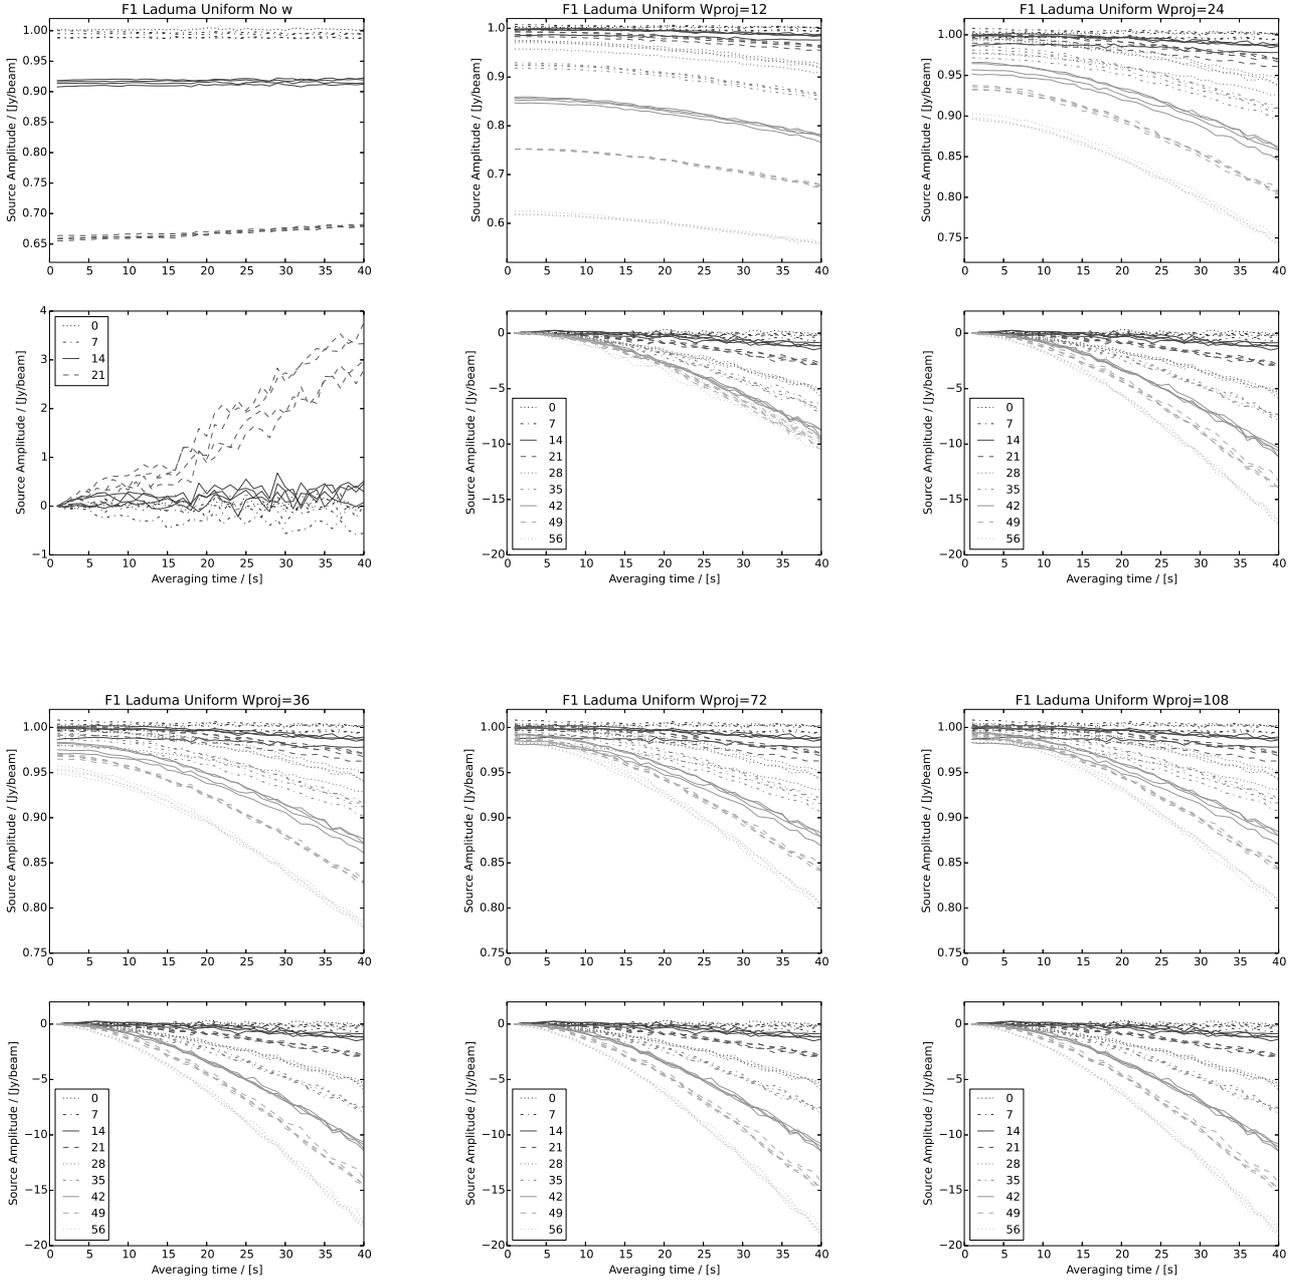


Figure 8: Plots of amplitude offset versus time averaging for field F1 Laduma/Mightee2, imaged using uniform weighting, for a range of wide-field imaging options (shown in the plot titles). For each wide-field option there are two panels. The upper panel shows the measured amplitudes for all of the sources. The lower panel the percentage amplitude reduction, as defined in Section 4. The linestyles are grouped by source distance from the centre, in units of arcminutes, as shown in the legend on each lower panel.

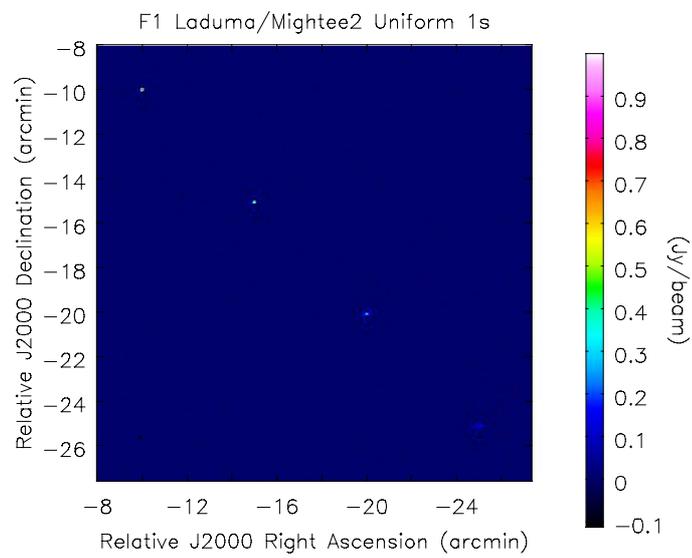


Figure 9: Inverted image of the F1 Fornax/Mightee2 uniform weighting $\Delta t = 1s$ data set, with no w -term correction. The image shows a portion of the bottom right section of the simulated source distribution, to illustrate the source shape distortion cause by the w -term, which increases with distance from the centre.