

EFFECTIVENESS OF THE CORRELATOR FIELD OF VIEW WEIGHTING TECHNIQUE IN SOURCE ATTENUATION

Dylan R. Nelson
University of California Berkeley, Berkeley, CA

Shep S. Doeleman, Divya Oberoi, Colin J. Lonsdale and Roger J. Cappallo
Massachusetts Institute of Technology, Haystack Observatory, Westford, MA

(Submitted 18 August, 2006)

Abstract. The science requirements of next-generation radio telescope arrays present a new set of challenges to traditional imaging and data processing techniques. Instruments such as the MWA, SKA, and other arrays implementing large numbers of small diameter dishes have a naturally large field-of-view due to the small diameter of individual telescopes. In order to achieve high image fidelity and dynamic ranges, noise contributions from off-center sources must be reduced – a task traditionally requiring imaging of the full field-of-view. However, implementation of this subtraction requires unreasonable computational resources while also generating unmanageable volumes of data. One approach to this problem involves dynamic control over the field-of-view, implemented in software as a weighting function internal to the integration routine of the correlator. Using this technique one can effectively reduce the noise contribution levels from sources outside the region of interest, and in some cases dramatically reduce the volume of data exiting the correlator for post-processing. In this paper we focus on verifying the effectiveness of this technique, implemented through the MIT Array Performance Simulator (MAPS) using simulated data sets. Additionally, several dimensions of possible parameter space are explored in order to test limitations and determine design requirements of this approach, including the impact of variable levels of radio frequency interference (RFI) excision on image fidelity and off-center source signal rejection.

1. Introduction

Next-generation radio telescope arrays are currently being designed to facilitate many scientific goals, which promise unparalleled views into a wide variety of current problems in physics and astrophysics. These goals are effectively driving the development of future instruments – as they present a number of technical challenges which previously have not been of great concern. The development of the Square Kilometer Array (SKA), in particular, has focused current efforts towards investigating and quantifying the technical requirements of future science. As a result, it is known that improvements in image fidelity and dynamic range are needed by several orders of magnitude over current radio telescopes in order to achieve desired sensitivity levels.

The problem arises because the sky is relatively full of radio sources, especially when observing with an instrument of high sensitivity. In order to achieve high angular resolution with a correspondingly good sensitivity, a radio telescope must handle the inevitable noise contributions from off-center sources which add sidelobe noise and confusion to the central region of interest. Conventional solutions require imaging the full field of view (FOV) in order to identify and subtract out contributions from these sources. However, potential designs such as the USSKA Proposal emphasize the need for a “large-number of small-diameter” (LNSD)

elements, which, due to the small diameter of individual telescopes, are naturally sensitive to a large field of view on the sky. Combined with the extremely high angular resolution specifications, imaging of the full field of view is projected to create unmanageable data volumes. Numbers referenced (Lonsdale, 2003) have been of order petabytes per second. Even with Moore's Law type extrapolations, to the amounts of potentially available computing power when the SKA will become operational, it would not be feasible to handle such data volumes in a cost effective manner.

As a result, novel techniques for reducing the amount of data output from the correlator are required. One such technique is correlator field of view weighting – where the effective field of view is precisely constrained within the correlator through the use of weighting functions applied during the integration stage. This technique can theoretically reduce contributions from sources outside the central region of interest, thereby enabling higher sensitivities without the cost of full FOV imaging. In order to accomplish this without loss of image fidelity, the robustness of the FOV weighting technique must be explored under various regimes of potential use. We first explore the implementation of the technique within the MIT Array Performance Simulator (MAPS), and then utilize the simulation architecture to verify the effectiveness of the approach.

2. Field of View Weighting

The technique we introduce here, field of view weighting, exploits the Fourier transform relationship between the (u,v) plane and the image, or sky, plane. Essentially, convolving in the (u,v) plane by a windowing function is equivalent to multiplying the image plane by the Fourier transform of that windowing function. In this way, the field of view of the data is essentially restricted before it ever exits the correlator hardware. The advantage of this approach is that FOV shaping occurs before any post-processing, and thus reduces data volumes generated by LNSD concept instruments.

2.1 Convolution Functions

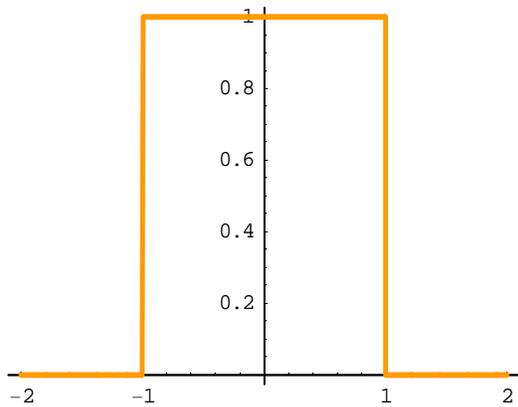
There are several different families of convolution functions which could be used to implement effective FOV weighting. Factors to consider include the resolution required for the convolution function in the (u,v) plane, the amplitude attenuation rate of the Fourier transform in the image plane, and any sort of extended sinusoidal or ripple structure which exists on large distance scales from the function maximum.

For the remainder of this paper we focus on a simple top-hat windowing function – effectively a unit step function. The top-hat is a demonstrator class convolution function, in that it is easy to implement within the simulation software, and has a Fourier transform amenable to easy measurement in the image plane. Thus, in order to gauge the effectiveness of correlator FOV weighting, we can apply a top-hat windowing function uniformly in the (u,v) plane and measure the sky plane to determine if the attenuation as a function of radius is as expected. The relation between the top-hat and its Fourier transform is given by:

$$w\left(\frac{r}{2a}\right) = \begin{cases} 1, & r < a \\ 0, & r > a \end{cases} \Leftrightarrow \mathcal{W}(q) = \frac{J_1(2\pi qa)}{\pi qa} \quad (1)$$

Here r is the radial distance on the (u,v) plane, and q is the radial distance on the sky plane. The function on the left represents the top-hat windowing function, which is implemented as weighting all values outside of a certain radius to zero, while weighting all those within that same radius to one. Its Fourier transform is a Jinc function, where J_1 is a Bessel function of the first kind. Visual representation allows a more direct approach to understanding the characteristics of both the convolution function and its Fourier transform:

1D Step Function



1D Sinc Function

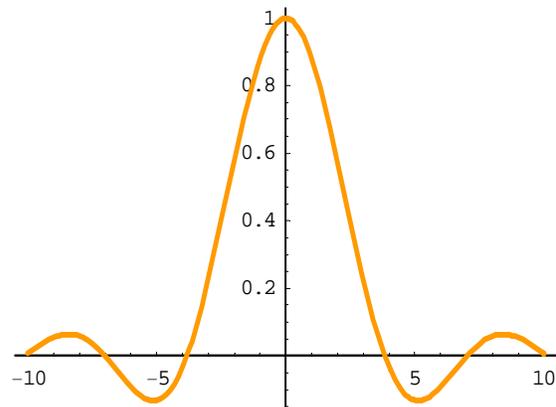
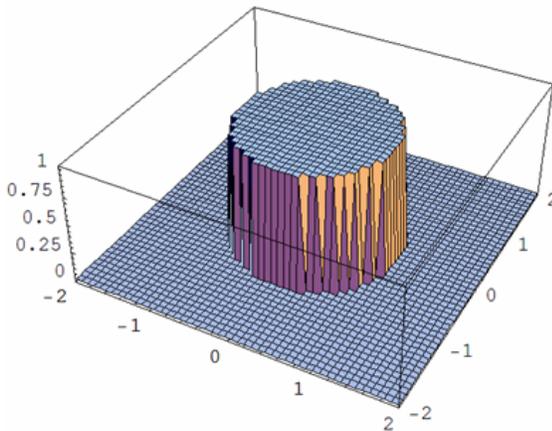


Figure 1. One dimensional representation of the top-hat convolution function and its Fourier transform, the Sinc function.

2D Step Function



2D Jinc Function

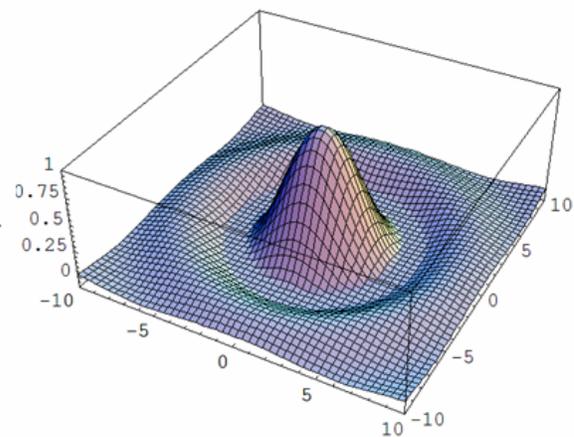


Figure 2. Two dimensional representations of the top-hat convolution function and its Fourier transform, the Jinc function.

In both the one and two dimensional cases we observe the sharp fall off of the Sinc or Jinc function with increasing radius. The ripple structure is also evident, which extends to infinity beyond the plotted region. The clear effect when multiplying the image plane by the two dimensional Jinc function is that sources near the region of interest are weighted more highly and thus emphasized, while sources further away are down-weighted and so effectively attenuated.

The degree to which we can control the precision of this attenuation is a function of the resolution in time and frequency within the integration step. This can be seen in Figure 2, where a finite number of divisions along both axes results in a shape which is only perfectly circular in the limit as the resolution in time and frequency approaches infinity.

Although the top-hat function is used exclusively for the remainder of this paper, it is important to note that there are several other potential candidates for convolution functions. A Gaussian convolution function has the following Fourier transform relationship:

$$W(\mathbf{r}) = e^{-\pi a^2 r^2} \Leftrightarrow \mathcal{W}(\mathbf{q}) = e^{-\pi q^2/a^2} \quad (2)$$

That is, the Fourier transform of a Gaussian is once again a Gaussian. This is beneficial since the function multiplying the sky plane monotonically asymptotes to zero, instead of containing sinusoidal structure. There are also more complex functions which have even better imaging characteristics, such as prolate spheroidal wave functions, whose Fourier transforms fall off more rapidly than either of the previous two examples.

2.2 Implementation

The implementation of the field of view weighting within MAPS is surprising simple, requiring no underlying changes in structure to the program. Since a single integration is preformed for each baseline in order to generate a single visibility point, we must only enforce a weighting scheme that corresponds to our desired windowing function. This is, effectively, the same as convolution of the (u,v) plane by the windowing function, and so the result will be a multiplication in the image plane by the Fourier transform of that windowing function.

The top-hat function implementation consists only of calculating the distance between the reference u,v point (at the center of the patch) and the current u,v cell being integrated over. If that distance exceeds our desired radius of the top-hat function, weighting for that point is set to zero, otherwise it is set to one. This weighted contribution must be then added into the integral, while the final visibility must be normalized for correlator weights.

2.3 Verification

In order to test the functionality of the FOV weighting technique and insure that source intensity in the image plane is being attenuated as expected, a series of simulations were run. A “plus-configuration” sky was generated by placing sources of known intensities at known locations. Specifically, 25 sources were placed in a 4-arm arrangement with a single source at the center. Spacing between sources was 75”, each source having an intensity of 1 Jansky. All other observing parameters are as described in the RFI excision tests, in order to facilitate a simplified implementation.

The goal of such a simulation is to monitor the falloff in source intensities with increasing distance from the map center. Plotting those sources intensities as a function of radius, and over-plotting the Jinc fall off function should reveal how closely the image plane follows the expected values. Image maps of the sky configuration without any FOV weighting as compared to the sky configuration with FOV weighting are shown, as well as the aforementioned plot of source intensities.

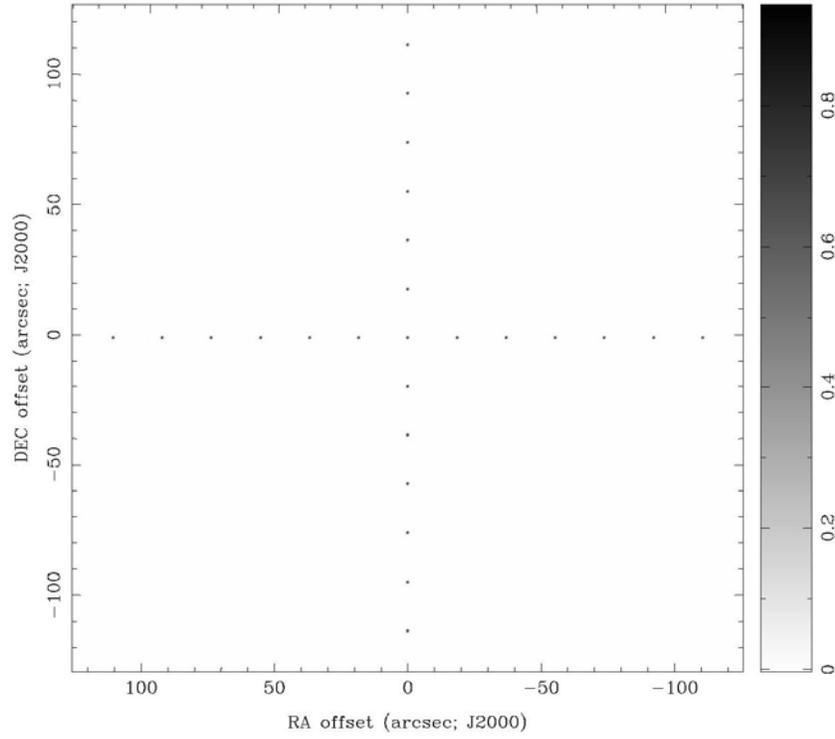


Figure 3. Image map produced with Miriad of the “plus-configuration” sky – normal simulation with no field of view weighting.

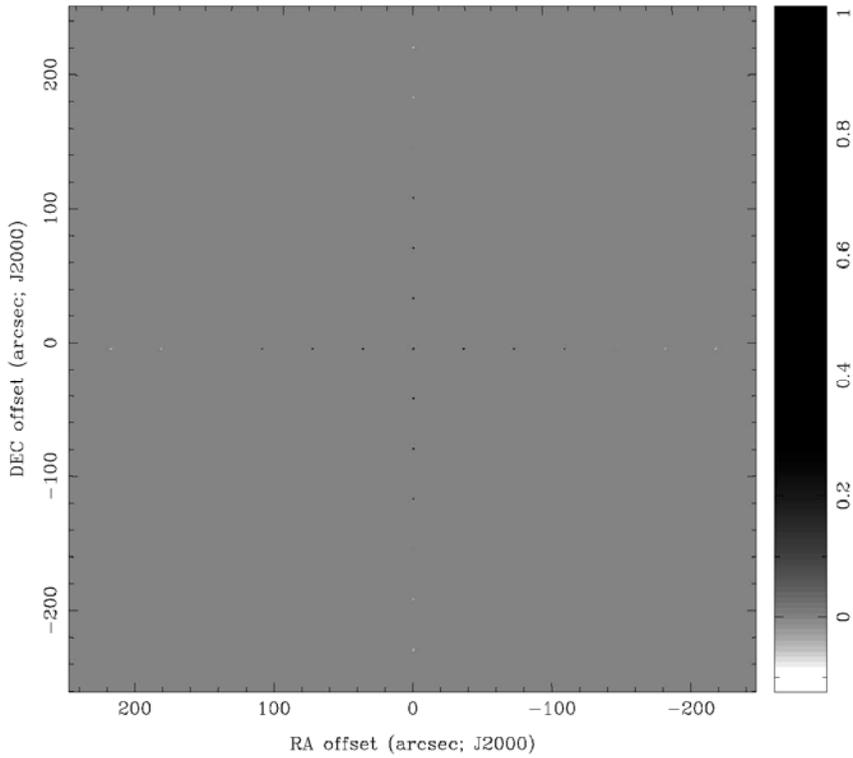


Figure 4. Image map produced with Miriad of the “plus-configuration” sky – simulation with a top-hat windowing function of radius 800 wavelengths enforced.

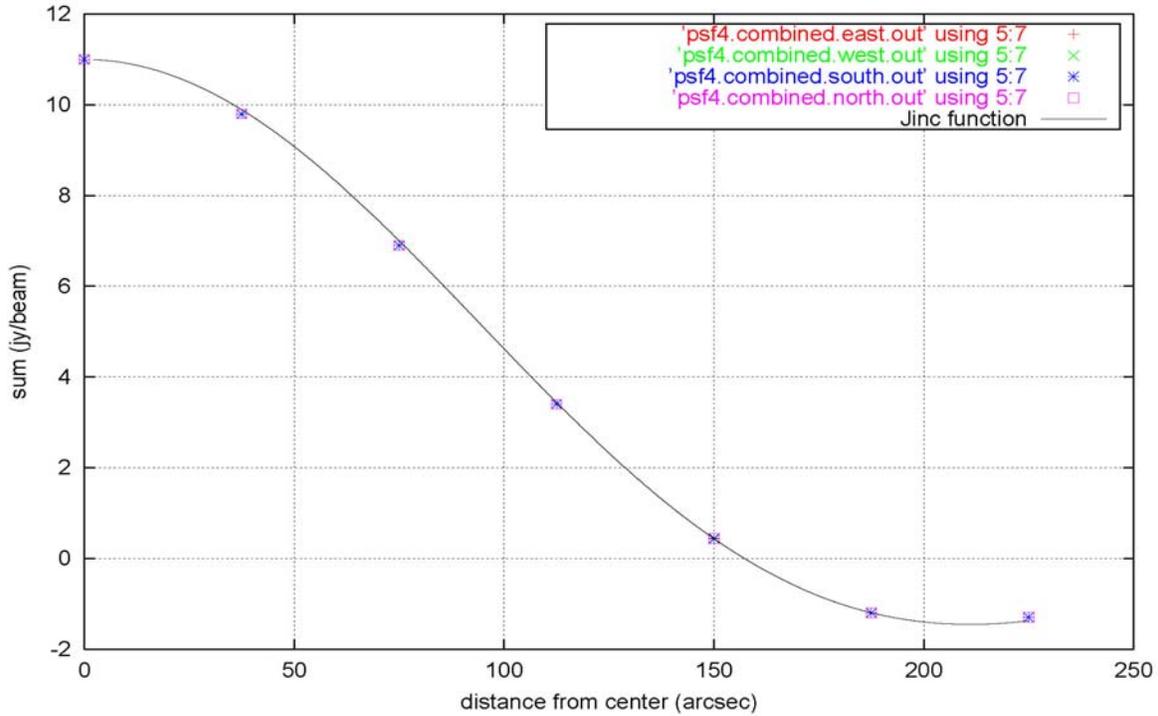


Figure 5. Source intensity versus expected fall off rate. North, south, east and west represent measurements at each distance along different arms. The Jinc function is scaled vertically to match the magnitude of the center point.

We can see that source intensity attenuation follows our expectation extremely closely. This agreement between measured and predicted fall off clearly demonstrates the ability of the FOV weighting technique to affect source suppression away from the phase center. This suppression can be tailored as desired by manipulating the windowing function applied in the (u,v) plane.

3. RFI Excision

Although, in the ideal case, the effectiveness of the field of view weighting technique has been established, there are a large number of confounding factors which will impact its performance in real-world applications. In this section we focus on a preliminary analysis of the effects of missing time/frequency “(t,f)” data on the effectiveness and accuracy of the FOV weighting technique in achieving off-center source attenuation.

3.1 Concept

Data gaps in the (t,f) domain are mainly expected to be due to automated Radio Frequency Interference (RFI) excision conducted prior to correlation. In modeling RFI excision there are two distinct types which occur – narrow band and wide band. Narrow band occurs in well-defined frequency bands which cover relatively small fractions of the total bandwidth, and may occur in discrete time intervals or across the entire time span of an observation. This type of RFI originates from transmission sources which produce interference at specific frequencies, corrupting those frequency bands while leaving others intact. In contrast, wide band RFI results when a transmission source swamps all frequency bands for relatively brief time periods, and may be intermittent, periodic or single occurrence. In the interest of time for these demonstrator

tests, the type of excision was limited to narrow band RFI occupying specific frequency bands throughout the bandwidth for the entirety of the observation.

The need to quantify the effects of RFI arise due to the differing representation of missing data in the (t,f) plane versus the (u,v) plane. The (t,f) extent over which an integration is performed to generate a visibility is uniform for all baselines. However, the (u,v) patch geometry differs on a baseline per baseline basis. As a result, the same (t,f) patch is mapped on to (u,v) patches of different geometry as well as size. Excised time and frequency bands are thus mapped to different data gaps in the (u,v) patch for each baseline. If a single convolution function is then applied consistently across those patches it will experience bands of excision to different extents for each baseline. As a result, the convolution function applied in the (u,v) plane is *no longer uniform* among baselines.

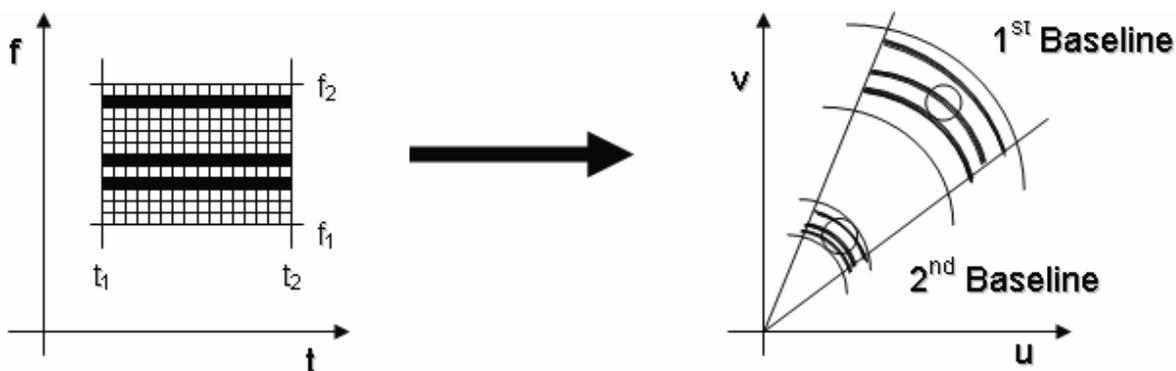


Figure 6. Relation between the time/frequency plane and the u,v plane. In this example three excised frequency bands are represented as black in the (t,f) plane. These three bands represent the same frequency ranges for all baselines, however once mapped into the (u,v) plane the data lost differs on a baseline per baseline basis. If the same size convolution function is the applied to each baseline (shown), the resultant shape of the convolution function is no longer uniform among all baselines.

The result is a modification of the corresponding Fourier transform of the convolution function, which is now variable for different visibilities. Since these visibilities now correspond to different fields of view on the sky, the effective FOV imposed in the image plane will correspond to an *averaged* FOV. This will compromise the efficiency with which this technique can attenuate source contribution away from the region of interest. The impact of RFI excision on FOV shaping is difficult to calculate analytically, however, the effects can be simulated and potentially modeled, giving a feel for the type of errors incurred.

3.2 Subdivision Algorithm Simulation Modifications

A problem arises in the implementation of RFI excision within the existing architecture of MAPS. The existing machinery was suitable for demonstrating a proof of concept for the FOV shaping technique, but not for simulating the effects of RFI excision.

Specifically, if the simulator intends to generate a single visibility point for each baseline, as is currently done in actual correlators, it places a single convolution function centered in the (u,v) patch for each baseline. However, if frequency bands are then excised, the behavior on small vs. large baselines varies dramatically. On a shorter baseline the convolution function fills the majority of the (u,v) patch and is, as a result, sensitive to the majority of any excision. With

the same excision parameters, however, a longer baseline could be completely unaffected if the convolution function occupied only a small portion of the (u,v) patch. If this was the case, the baseline would only be sensitive to a narrow range of frequency and time excision.

This behavior is undesirable since next-generation correlator architectures utilizing FOV weighting schemes would undoubtedly implement an algorithm to subdivide (u,v) patches based on their ability to fit multiple convolution function sized sub-patches. In order to accurately simulate the effects of RFI excision, then, the simulator architecture was modified to accommodate this new behavior.

To simplify the implementation certain assumptions were made. First, the observing array was placed at the North Pole (Latitude 88:00:00.0, East Longitude: 00:00:00.0), and the observed region of the sky was set to be directly overhead (Dec: 88:00:00.0, RA: 00:00:00.0). In both cases the exact North Pole is avoided due to certain degenerate behavior associated with the motion of the Earth's rotation at that point. Nevertheless, our approximation is close enough that the paths traced by individual baselines in the (u,v) plane are essentially circular. This insures that there is no time variability of the radius at which a given frequency for a given baseline falls. As a result the (u,v) patches over a given snapshot are bounded in frequency by circular arcs and in time by radial vectors. We can then avoid the computational complexity arising due to changes in geometric projection of the baselines – when viewed from the source – as a function of time. It should be noted that this limitation is not of direct relevance to our current goal, and does not limit the validity of these results in any way. This simplifies the computation of the subdivision algorithm by allowing it to assume a general correspondence between the (t,f) and (u,v) planes.

The initial implementation was also designed to accommodate only snapshot observations, where the observing time was equal to the correlator integration time, and the observing bandwidth was equal to the correlator channel bandwidth. This insures that a single point is a priori generated for each baseline. The needed functionality to allow the subdivision algorithm to handle multiple time and frequency slices exists, but was not used for this test.

The logic of the implementation proceeds within the baseline loop of the simulator by first pre-computing the extent in frequency/time space needed to fit a single convolution function (the size of which is specified in wavelengths on the u,v plane). This size is variable among baselines. The time extent required to contain a convolution function is given by:

$$\Delta t_{\text{req}} = (2 R \sigma) \cdot \frac{c}{\nu} \cdot \frac{1}{D} \cdot \frac{1}{\omega_e} \quad (3)$$

The frequency extent required to contain a convolution function is similarly:

$$\Delta \nu_{\text{req}} = (2 R) \cdot c \cdot \frac{1}{D} \quad (4)$$

Where in both cases R is the radius of the convolution function, c is the speed of light, ν is the central observation frequency, D is the u,v distance for that baseline, ω_e is the angular rotation speed of the Earth (0.000072921159 rad/s), and σ is a padding factor inversely proportional to the time resolution of the integration function.

The number of times that convolution function can fit within the f,t extent of the observation is then calculated, and truncated to an integer value. Next, the f,t space is divided

equally by that number of times. The central time and frequency values for each subdivision, as well as the upper and lower bounds, are then fed into the subdivision visibility point generation logic. This consists of two loops running over the time and frequency subdivisions on a baseline per baseline basis, passing the modified f, t values into the integration stage of the correlator.

The loop logic had to work around existing architecture within the visgen program, which was not designed to handle more than a single visibility point for a given time and frequency slice. Several different approaches were attempted and discarded. All these placed the additional visibilities as the results of *effective* correlator frequency channel visibilities, in order to work within the output file structure of visgen as much as possible. This file structure includes references to frequency tables and visibility counts, which were previously independent of baseline, but are now a dynamic function of each baseline. Modifying the header structure to properly account for these changes is infeasible, as the downstream conversion programs (specifically, vis2ms and glish) expect certain output from visgen.

Several attempts to circumvent this problem were made, including writing the visibility frequency structure inside the frequency subdivision loop using a modified indices structure. This failed, as downstream handling requires a constant number of correlator frequency channels for each observing frequency slice. This is no longer the case, since the frequency integration extent is baseline dependent.

As a second attempt we implemented an over sampling of correlator frequency channels at a constant number per observing frequency slice. The visibility point resulting from the convolution function subdivision is then placed in the appropriate over sampled channel. The accuracy of this technique is then a function of the over sampling resolution – where 100 channels were found moderately acceptable in placing of order 18 visibility points with sufficient frequency accuracy. The over sampling of correlator frequency channels fails in the imaging process, however, since each correlator frequency channel is expected to represent the frequency plane of an image cube in a real observation. This was found to be an intractable problem with no time efficient solution.

Our final implementation makes two improvements on both of these previous attempts. Firstly, it moves the output of the visibility block header as well as the visibility frequency structure inside the frequency subdivision loop. This has the result of producing more visibilities than predicted by the visgen time block header. We thus pad the number of expected visibilities to the maximum amount of subdivisions in the largest baseline times the number of baselines. Visibilities which for smaller baselines are not written out are identified and weighted to zero, in order to be filtered out in the inversion stage of imaging. Secondly, the implementation calculates the (u, v) coordinates for each visibility point in order to place it at the exact location corresponding to its central time and frequency. This is done once per visibility block header, and so now can be tailored for every visibility point generated by the subdivision algorithm.

An example output of the subdivision process is given below. UV coverage plots for a VLA-like array are given with no subdivisions (left), and with subdivisions occurring at most 18 times in both time and frequency for the given observing parameters (right).

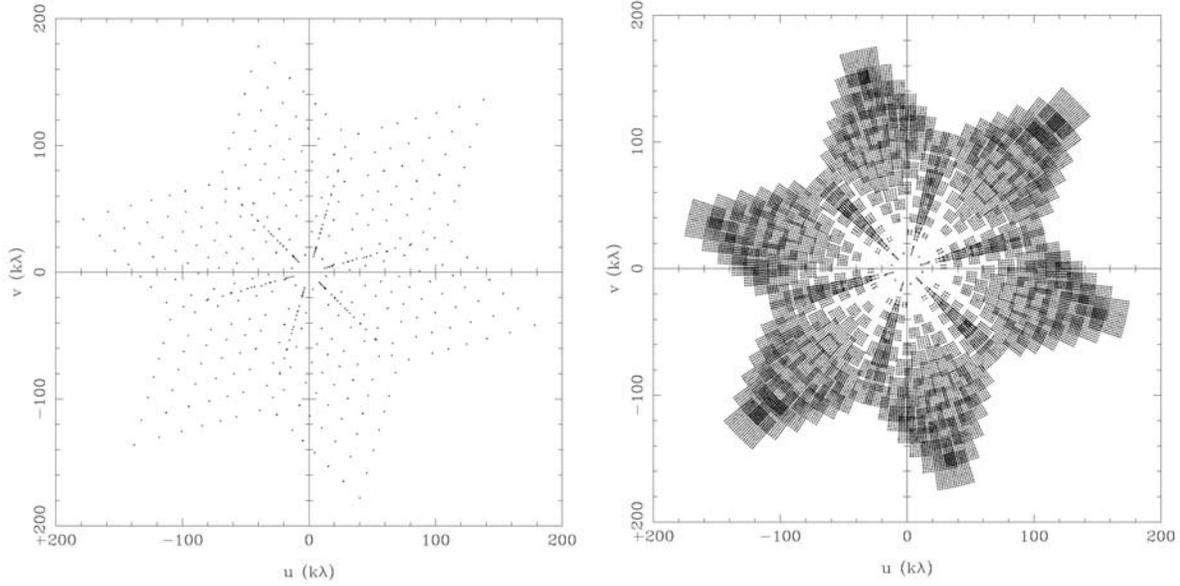


Figure 7. The effects of the subdivision algorithm on the u,v coverage plots of a VLA-like antenna configuration. The array consists of 21 elements, producing 210 visibility points without subdivision, as seen on the left. With these specific observing parameters, 14210 visibilities are produced after enabling u,v patch subdivisions.

Validating the visibility data produced with the (u,v) patch subdivision algorithm consisted of repeating the FOV weighting test – expecting to see source intensity attenuate as a function of distance from the phase center exactly corresponding to the structure of the Fourier transform of the convolution function. Image fidelity was found to be consistent with expectations.

3.4 RFI Excision Simulator Implementation

Having programmed the machinery to subdivide (u,v) patches into multiple convolution sized patches, we went for a quick approach to implementing narrow-band RFI excision. Within the integration stage there is a finite resolution in time and frequency, of order 64, 128, or 256 cells. In order to avoid issues whereby different frequency bands would be excised on different baselines due to the different coverage of these integration cells, we parameterized the smallest excision band size as the total observing bandwidth divided by the integrator frequency resolution.

For each simulation, then, a “ n parameter” is passed into visgen which represents the number of frequency bands, out of the total resolution bands, to excise. For example, an “ n parameter” of 32 when using 128 frequency subdivisions in the integrator corresponds to 25% RFI excision. After dividing the observing bandwidth into a number of bands equivalent to the integrator resolution, the requested number are randomly chosen and fixed in an array of excision band structures. This array is the passed into the integrator, which for each cell iterates over the requested excision bands and checks to see if the current cell being integrated over is within a band. If so, it weights that cell to zero, just as in the top-hat convolution function distance check. In order to parameterize the amount of RFI excision, then, we step through our “ n parameter,” which runs from 0 (no RFI excision) to 128 (100% RFI excision) in the case of the following tests.

3.3 Method

Having programmed the machinery to implement tests of RFI excision, we proceeded as follows. A configuration of 8 sources was randomly placed equidistant from the sky center (pictured below) at a distance equal to the null point of the jinc function. This was calculated to be approximately $157.165''$ when using a top-hat windowing convolution function of radius 800 wavelengths.

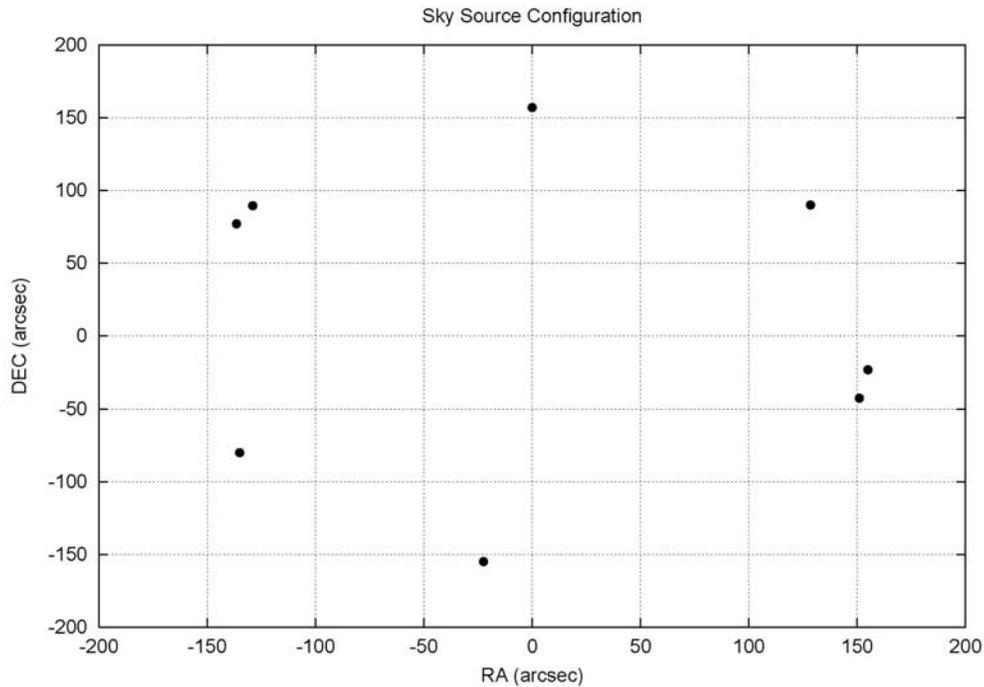


Figure 8. The configuration of sources used for the RFI excision tests. All sources have an intensity of 1 Jy, and were placed randomly on a circle of radius $157.165''$ centered at the origin.

The expectation was that, given a perfect agreement with fall off expectation, all eight sources should be attenuated to zero intensity in the case of no RFI excision and thus uniform convolution functions across all baselines. Requiring the placement of sources at integer pixel locations at this step introduced a small though measurable error, which is accounted for by taking the ratio of measurements with excision to the value of those same measurements with no excision.

This sky configuration was imaged with the `vla_array.pos2` array configuration, a VLA array design with baselines less than 2 km removed. The observing parameters used were: a field of view of $512.0''$, scan duration and integration time of 2400 seconds, scan bandwidth and correlator channel bandwidth of 240 MHz, with a central observing frequency of 1.42 GHz. The integration resolution in time and frequency were set equal at 128 subdivisions. Parameter space of amount of RFI excision is stepped through, keeping all other variables constant. For each “n parameter,” equivalent simulations were run 4 times in order to create an average result over the differing randomly selected frequency excision bands. In each simulation an analysis pipeline in Miriad was used to extract the intensity of a source at the null, the RMS at the map center, the RMS close to the edge of the map, and Gaussian fit profiles for a source at the null. The

Gaussian fits were intended to provide a measure of source shape deformation, however the *imfit* task in Miriad was unsuccessful, and those results are not included. Similarly, the outer map RMS was found to be nearly identical to the map center RMS, and so was not included.

3.4 Preliminary Results

After averaging over the 4 simulation runs per “n parameter” of RFI excision, where we took n=0 through n=102, the raw findings are given below:

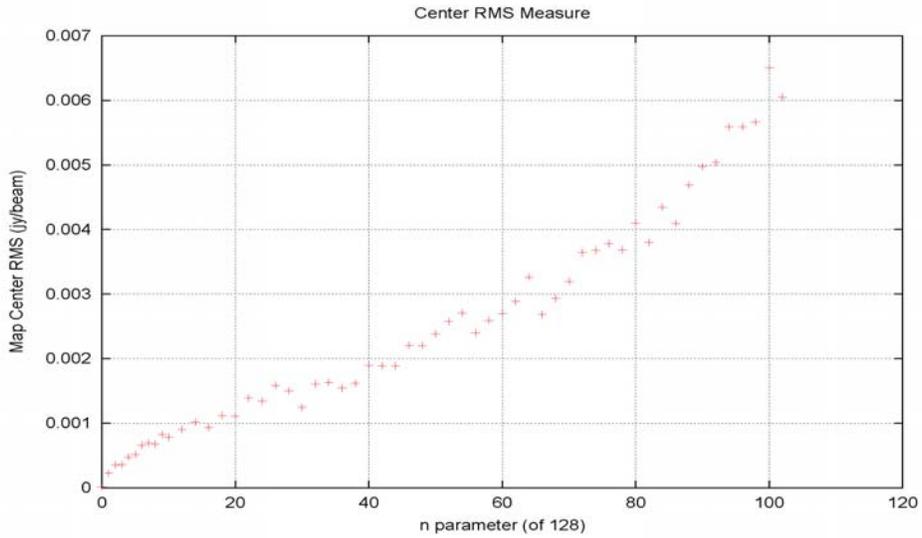


Figure 9. Deviation of the source in the jinc null location from “zero” intensity, measured as the increase in maximum pixel intensity at the source’s expected location.

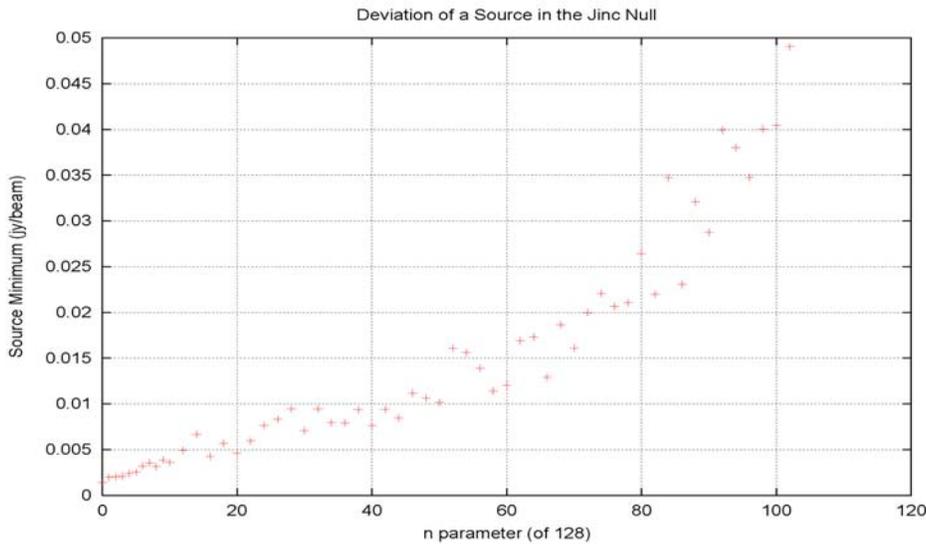


Figure 10. RMS measured at the phase center of the map.

In order to gain a better understanding of the change in source intensity and map center RMS as a function of RFI excision the same data is given reformatted as a ratio in relation to the n=0 values, which are 1.41E-03 Jy for source maximum and 1.21E-05 Jy/beam for the RMS at the center of the map:

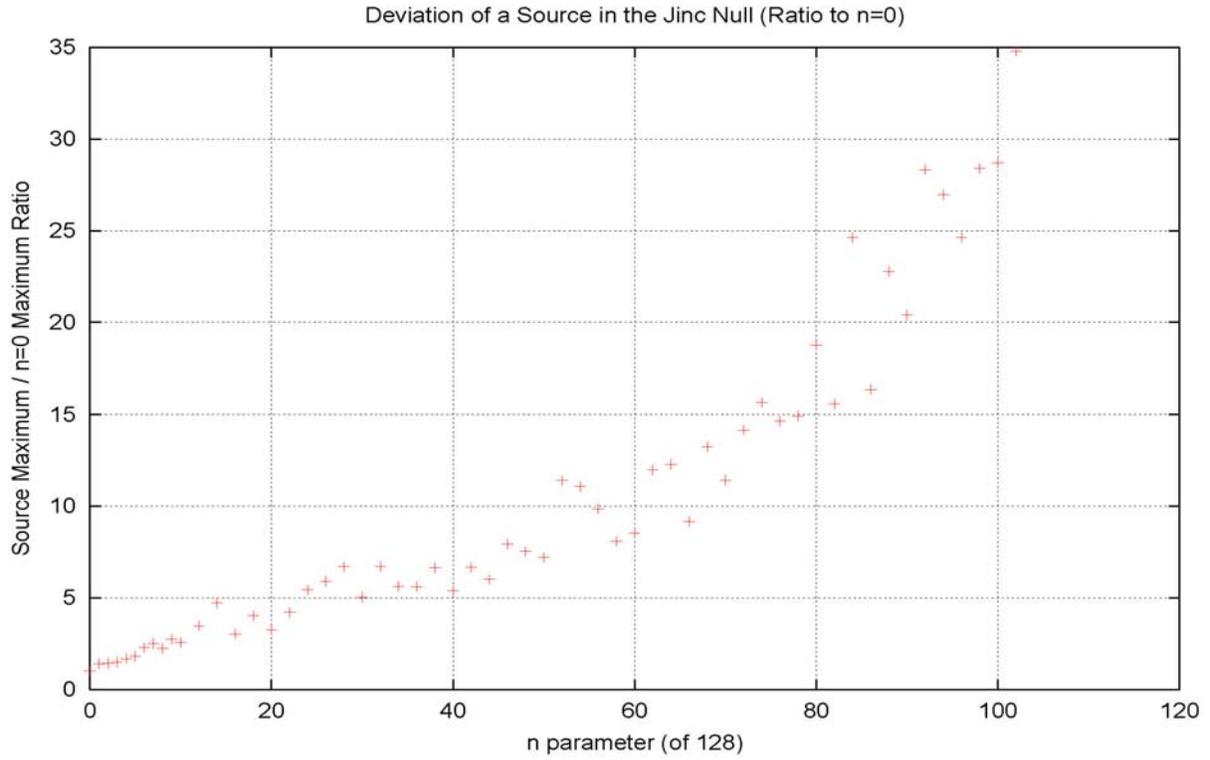


Figure 11. Variation in the ratio of the observed source maximum at the Jinc null, to that for n=0.

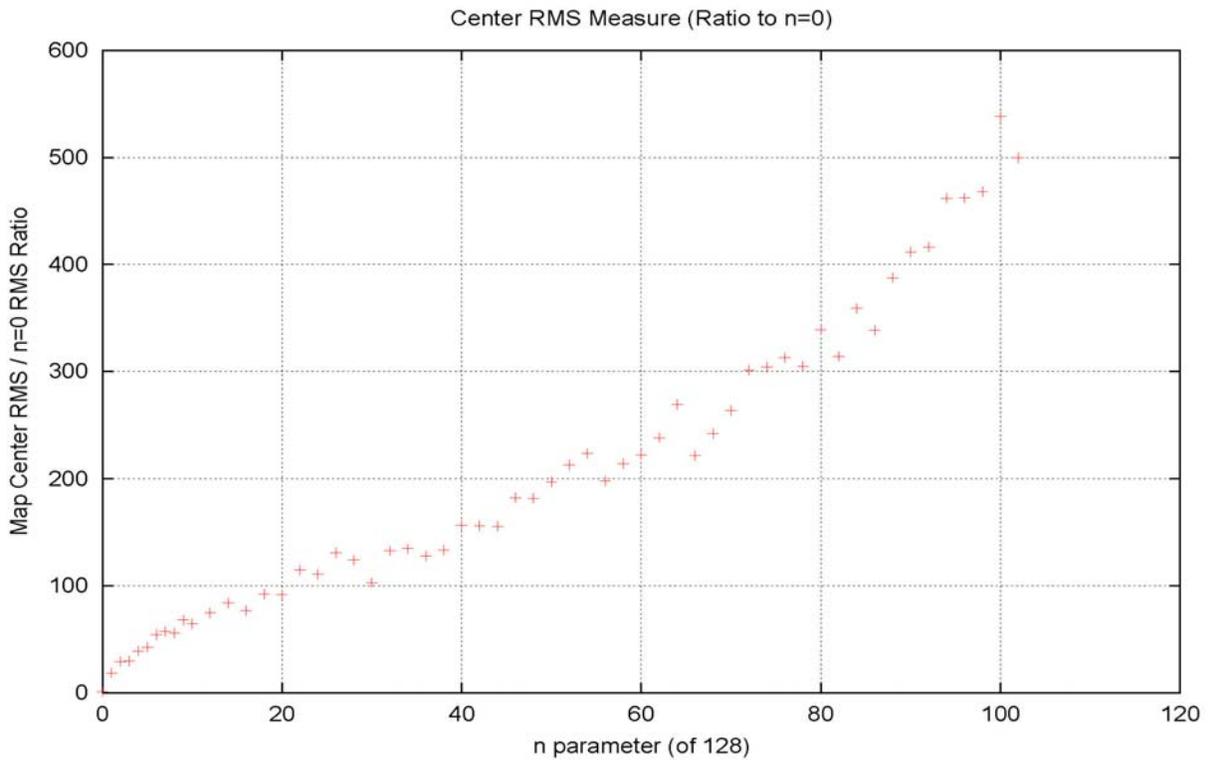


Figure 12. Variation in the ratio of the RMS at the center of the map, to that for n=0.

We can observe that, at reasonable levels of RFI excision, deviation from the expected behavior ($n=0$) appears linear. Although the rate of increase in map RMS is larger than that of the deviation of the fall off expectations of the Jinc function, they appear to share a linear relationship. To obtain a better quantitative measure of this, we can get an estimate for the residual dynamic range of the resultant images by taking the ratio of the source intensity peak to the image RMS.

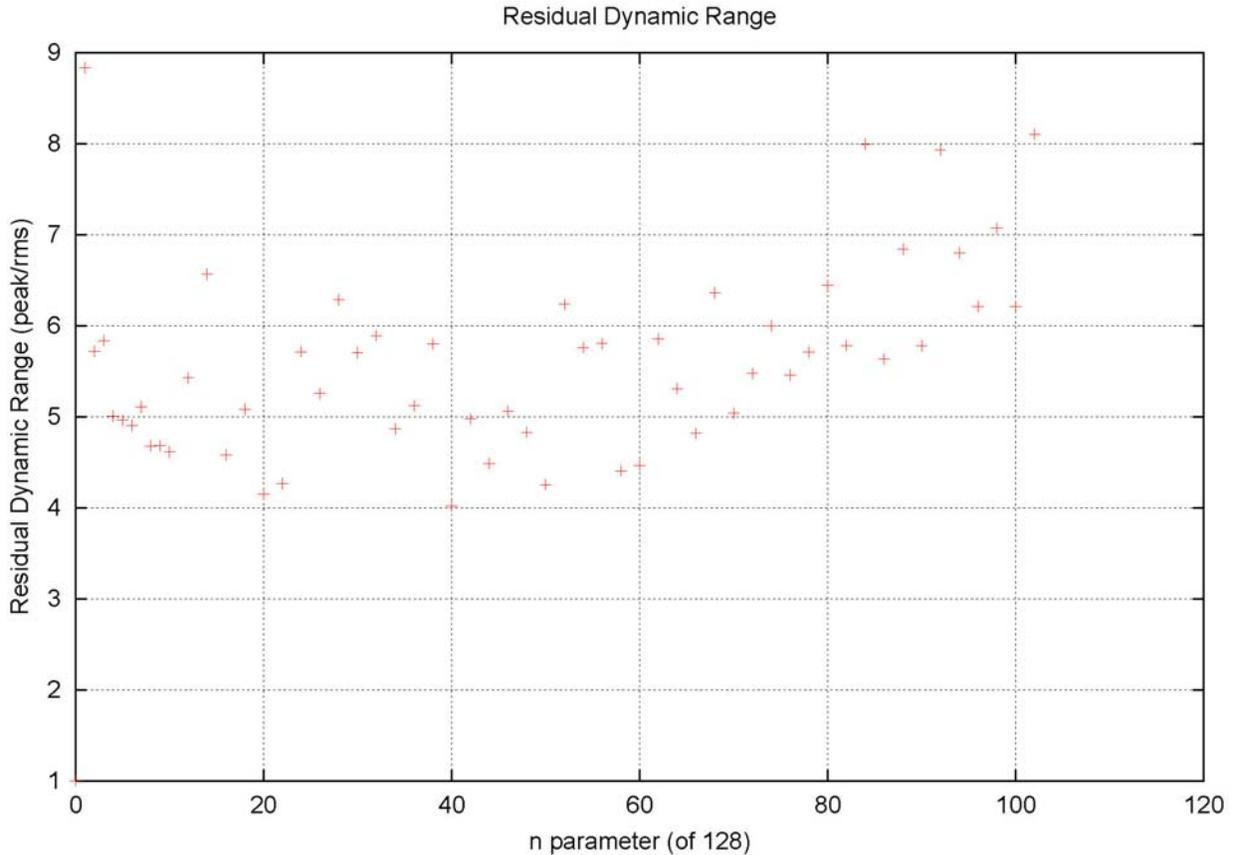


Figure 13. Residual dynamic range of the image across RFI excision space. Plotted as the ratio of the maximum pixel intensity of the source in the jinc null to the map RMS at the phase center.

To first order the dynamic range of the resultant images appears linear across the range of RFI excision. We expect a correspondence between the amount of intensity showing through at the jinc null location and the map RMS. The fact that we find a linear relationship is promising, as it indicates a lack of other confounding factors affecting map RMS in the face of RFI excision. It could have been the case that the non-uniformity of convolution functions among different baselines would have unexpected and unpredictable effects for image fidelity.

It appears, rather, that the decrease in image fidelity follows a pattern which can at least be predicted, and potentially modeled. Thus, although RFI excision has clear detrimental effects on image quality, the impact appears to be quantifiable. Additionally, it appears to lead to a gradual degradation of image quality in the expected regimes of RFI excision for the SKA. One caveat here that requires further exploration is the failure of *imfit* to successfully model any amount of source deformation occurring.

4. Conclusion

The effectiveness of the field of view weighting technique in attenuating contribution from sources away from the center region of interest is now established. What is less known is the ability of the technique to operate under confounding factors likely to be found in real world applications. The degradation in performance with increasing RFI, investigated here, seems to be both gradual and quantifiable in the regime of low RFI where the SKA is expected to operate. Additionally, the benefits of FOV weighting inside the correlator must be balanced against its original purpose. As the original problem was one of reducing data volume while maintaining image fidelity, we must ensure that the computational cost of FOV weighting does not outweigh its benefits. Here the issue of time and frequency resolution within the correlator comes into play, as the technique is only effective when high enough resolution is maintained inside the correlator up until the point where weighting is applied.

Analysis of this resolution requirement is only one area of the many that require exploration before any possible implementation studies, hardware or software based, can begin. These include possible limits to the dynamic range of baselines lengths, requirements in terms of calibration tolerance and drift over time, the level of time/frequency subdivision resolution required at the time of integration, and the effects of missing time/frequency data. Additionally, the amount of source suppression needed at various radii from the phase center is unknown – it is possible that at some finite distance an array such as the SKA can simply ignore insignificant source contribution. In order to test this requirement one should consider the use of unrealistic skies tailored for testing purposes, in addition to benchmark realistic skies or actual skies derived from sky maps, each of which have their own advantages and disadvantages. Finally, the computational feasibility for the actual implementation of a field of view weighting technique must be investigated.

It appears, though, that the creation of images with the sensitivity and dynamic range requirements of next-generation science objectives is a feasible goal even in the face of data volume issues facing LNSD array designs. As traditional techniques are not well suited to handling the demands of next-generation arrays, novel approaches to the problem are required. Field of view weighting within the correlator is one of these new techniques, and promises to be a strong candidate for addressing this problem, though a more critical examination is called for.

Acknowledgements

- ❖ Shep Doeleman – for taking the lead and spending countless hours massaging away the problems that arose.
- ❖ Divya Oberoi – for being a consistent resource, always willing to answer the many questions that came up.
- ❖ Colin Lonsdale & Roger Cappallo – for being there whenever one of those questions stumped the rest of us.
- ❖ Katherine de Kleer & Chris Beaumont – for knowing what was going on with AIPS.
- ❖ Richard Crowley – for making everything work.
- ❖ MIT Haystack – for hosting myself and all the other REUs.
- ❖ NSF – for funding the program.

References

- Bracewell, Ronald N.: 1986, *The Fourier Transform and Its Applications*. McGraw-Hill: Boston.
- Bunton, J.: 2005, *SKA Signal processing costs*, SKA Memo 65
- Cappallo, R. J.: 2002, *Visibility Integration*, MAPS Memo 6.
- Cornwell, T.J.: 2005, *SKA computing costs for a generic telescope model*, SKA Memo 64
- Lonsdale, C. J., Doeleman, S. S., and Oberoi, D.: 2005, *Efficient Imaging Strategies For Next-Generation Radio Arrays*, *Experimental Astronomy J.* **17**: 345-362.
- Lonsdale, C. J.: 2003, *Data rate and processing load considerations for the LNSD SKA design*, SKA Memo 32.
- Ed. Perley, Richard A.: 1985, *Synthesis Imaging: Course Notes from an NRAO Summer School*. NRAO: Green Bank.
- Thompson, A. R., Moran, J. M., and Swenson, G. W.: 1986, *Interferometry and Synthesis in Radio Astronomy*. John Wiley & Sons New York.