What Fourier Plane Coverage is Right for the MMA?

M.A. Holdaway
National Radio Astronomy Observatory
Tucson, AZ 85721

July 9, 1996

Abstract

Simulations of various abstract Fourier plane distributions indicates that a moderately centrally condensed (u, v) coverage has much better noise characteristics upon tapering than a uniform (u, v) coverage, but still maintains a very nearly Gaussian beam. To achieve the same untapered resolution, an array with centrally condensed (u, v) coverage will require maximum baselines which are about 25-35% longer than the maximum baselines in an array with uniform (u, v) coverage. Due to the larger sidelobes of its synthesized beam, images produced from a uniform Fourier plane distribution may have lower dynamic range than images produced from an intrinsically tapered Fourier plane distribution, though preliminary results are conflicting. The largest three out of four MMA configurations are currently designed to be some sort of ring, either an ellipse or a Reuleaux triangle, which will result in approximately uniform Fourier plane coverage. I argue that we should consider a moderately centrally condensed Fourier plane distribution for all of the MMA configurations, with the possible exception of the largest array.

1 Introduction

Early work in designing the optimal array configurations for the MMA (Cornwell, 1984) and for the SMA (Keto, 1992) assumed that the (u, v) points should be uniformly distributed across the part of the Fourier plane which was to be sampled. This assumption led to ring-like arrays. This is intuitive since the autocorrelation of a ring is a uniform disk with a delta function at the origin. However, the arguments for a uniform Fourier plane distribution have not been on ground as solid as the methods which have been created to give us arrays resulting in uniform coverage.

Radio astronomers, especially in the VLBI community, have struggled for years to get something for next to nothing, reconstructing images from poorly sampled Fourier plane data. Given an arbitrarily complex object, we must have complete (u, v) coverage, sampling the Fourier plane at the Nyquist rate. The MMA's most compact configuration has essentially complete Fourier plane coverage (out to its maximum baseline) in a snapshot, and its a good thing too, since each field of the mosaics it makes will often be filled with complicated structure on all scales. When you don't completely sample the Fourier plane for an arbitrarily complex object, you are trying to solve for more independent resolution elements than you have independent data measurements. Fortunately, most astrophysical sources are not "arbitrarily complex", but rather vary in some reasonable manner. Algorithms such as MEM are therefore able to construct reasonably good images even in ill-determined problems due to MEM's bias towards simple images, and image quality degrades gracefully as the imaging problem becomes less well-determined.

On the other hand, there are many objects which are simple enough to permit high quality imaging with relatively poor sampling of the Fourier plane. For example, even without a support constraint explicitly forbidding emission from being reconstructed outside some region of the image, a point source can be imaged well with only a few measured visibilities. Objects of intermediate complexity can be imaged with high quality with moderate Fourier plane coverage, especially if a strong support constraint holds, limiting the emission to a small part of the primary beam.

The desire to design an interferometric array which produces nearly uniform Fourier plane coverage derives from the quest to obtain complete Fourier plane coverage, or as nearly complete Fourier plane coverage as is allowed by the number of antennas in the array. However, this goal does not properly address the kinds of objects which the array will be imaging or the reconstruction algorithms which are currently in use. It may be a dangerous activity to hypothesize how an array will be used based on our current knowledge of the universe, but it may also gain us a great deal if our estimates of the array use are correct. For example:

- If the MMA were only to observe very bright and complicated sources such as the planets, the sun, and nearby continuum sources, at the full natural resolution of each configuration, striving for uniform Fourier plane coverage might be a good strategy.

- If the MMA were to observe mainly sources which were marginally detectable (and indeed, astronomer's optimism will result in many sources not being detected at the full resolution of the configuration in which they were observed), then we would not be so acutely concerned with the ability to achieve extremely high quality images, but with the ability to optimally pull images out of our noisy data.

These two situations place very different demands on the Fourier plane distributions, and because the MMA will often find itself squarely in each of these situations, some compromise between these competing demands is warranted.

2 Tapering the MMA

We will want to taper the MMA's Fourier coverage in at least two common situations:

- we simply don't have the SNR to see what we want to, and tapering will increase the surface brightness sensitivity. (For uniform coverage, doubling the linear beam size increases the beam area by a factor of 4, reduces the number of baselines by a factor of 4, and increases noise by a factor of 2. Hence, T_b noise goes down by 1/2.) Since we have a factor of 4 in size between arrays, this will be done often as people try to observe in the highest resolution configuration just barely possible, and find out that the longest baselines are just a bit too long. Tapering the velocity resolution is an alternative to tapering the spatial resolution, at least for spectral line observations. Tapering the spatial resolution is more effective at increasing the brightness sensitivity, but the tradeoffs between tapering the velocity resolution and the spatial resolution depend upon the scientific objectives of the observations.

- many studies will compare different molecular transitions at different frequencies, and hence, at different resolutions. An accurate comparison requires the same resolution, so N-1 out of N compared images will be tapered to some extent. Since the configuration sizes are quantized in factors of about 4, significant tapering will often be required. Quantities derived from molecular transition comparisons will also be limited by the thermal noise of the input images, so optimal noise behavior upon tapering will buy us lower noise in these derived quantities.

Since the images will often be limited by thermal noise, designing antenna configurations which give good imaging and low noise both at full resolution and when tapered will increase the scientific throughput of the MMA.

3 Noise Behavior of Different Fourier Plane Distributions When Tapered

Figure 1 shows the radial Fourier plane distributions which are obtained with a ring array and a filled array. The central peak in the case of the ring array is inevitable for a large number of antennas. The radial Fourier plane distribution of the ring array is basically uniform, as the central peak does not cover very much of the 2-D (u,v) area. Hence, we justify approximating the ring array as a uniform Fourier plane coverage. The filled array gives a centrally condensed Fourier plane coverage which we approximate as a function which decreases linearly with (u, v) distance.

I have simulated (u, v) data with abstract Fourier plane distributions, ie, Fourier coverages which are not associated with any physical distribution of antennas, and investigated some of the properties of these Fourier plane distributions. The radial profiles of uniform and centrally condensed (linearly decreasing) Fourier plane distributions are shown in Figure 2. To make the comparison more fair to the uniform Fourier distribution case, we stretched the centrally condensed distribution by 30% so the naturally weighted beam would be the same size as for the uniform distribution. (If we make the maximum baseline of the two distributions the same, rather than the resolution, the relative increase in SNR of the centrally condensed Fourier plane distribution upon tapering is exaggerated.)

Other moderately centrally condensed Fourier plane distributions have been studied and give similar noise curves after normalizing the width of the distribution to give the same natural resolution. I made simulations of 28,000 random (u, v) points which were consistent with these distributions. On a 128 by 64 Fourier plane grid, essentially every cell within the maximum (u, v) distance had one or more (u, v) samples for each of these distributions. These simulated (u, v) data sets were gridded with natural weighting, tapered to the desired resolution, and Fourier transformed to yield the synthesized beam. Untapered robust and uniform weightings were also investigated. I explored tapers as large as 4 times the full resolution beam. (Larger tapers won't be used often, as they provide the same resolution as the next smaller array configuration, but much less efficiently.)

The two features of the synthesized beam which we study here are the amount of sensitivity lost in the tapering, and the level to which the main lobe of the synthesized beam conforms to a Gaussian. The extent to which the beam mimics a Gaussian is important to imaging and analysis algorithms but is not fundamental. Since almost every radio astronomical researcher convolves their clean components or maximum entropy images with a Gaussian beam, we will assume that a Gaussian beam is desirable. If the synthesized beam falls off faster than a Gaussian, the sensitivity on the longer baselines is wasted as we convolve with a Gaussian with wider wings. On the other hand, a *highly* centrally condensed Fourier plane distribution, such as is obtained from the VLA, causes a synthesized beam with very broad wings which results in imaging problems since most of the array sensitivity is on very short spacings.

Figure 3 illustrates the point source sensitivity loss as a function of taper for the uniform and centrally condensed distributions. Consistent with intuition, the centrally condensed distribution does not lose sensitivity with taper as quickly as the uniform distribution. At 4 arcsec taper (almost a factor of 2 lower resolution than naturally weighted), the uniform distribution has 16% higher noise than the centrally condensed case; at 6 arcsec taper, the uniform distribution has 21% higher noise than the centrally condensed case. The upturned spurs at the high resolution end of Figure 3 are the robust and uniformly weighted results. While the centrally condensed Fourier plane distribution suffers much in sensitivity when robust or uniform weighting are used, it is able to achieve much higher resolution than its naturally weighted case. With uniform Fourier plane coverage almost no increase in resolution and little loss of sensitivity occur.

Figure 4 shows the fractional difference between the integrals of the synthesized beam and its best fit Gaussian, down to the 0.10 level of the Gaussian. Both uniform and centrally condensed distributions show significant departures from Gaussian beams at the highest resolution, but the uniform distribution produces much higher deviations. When the Fourier distribution is highly tapered, both beams are very similar to a Gaussian, the uniform distribution being a few tenths of a percent away from Gaussian and the centrally condensed distribution being about 1% away from Gaussian. The nearly Gaussian beams which result upon tapering indicate that the relative noise improvements of the centrally condensed Fourier plane coverage over the uniform Fourier plane coverage are not at the expense of imaging quality.

3.1 Interpretation of Taper Results

What is the optimal Fourier plane distribution? Assuming a Gaussian beam in the final image is desired to aid in the interpretation and analysis of the image, it is clear that a uniform Fourier distribution is NOT optimal. Figure 5 shows cuts through the uniform distribution's beam (very nearly a J1 Bessel function) and its "best fit" Gaussian. The Gaussian has much broader wings than the uniform distribution's beam, or the uniform distribution has too many long baselines to yield a nice beam. If the uniform coverage is tapered to produce a beam which is more Gaussian, about 30% of the sensitivity is lost. While I can't say at this point which Fourier plane distribution is optimal, I can say that the uniform distribution is sub-optimal with respect to resulting beam shape and noise behavior upon tapering.

3.2 Caveats

All of our reasoning has been based on the assumption that we have so many antennas and (u, v) samples that every cell which needs to be sampled can be sampled, and we are then dealing with the issue of where the extra (u, v) samples should go. The conclusions we draw for the 40 element MMA, for which this will often be true, will be quite different from the conclusions that might be drawn from a 6 element instrument, which is really stretching to fill as many cells as possible. A typical MMA field observed with the D array will be filled with complicated structure and requires complete (u, v) coverage. The centrally condensed Fourier plane distribution which results from the requirement of maximum surface brightness sensitivity in the D array also meets the complete (u, v) coverage requirement. In C array, either uniform or the linear distributions will usually fill most cells, permitting good imaging of complex objects. If typical fields observed by the MMA's larger configurations are not filled with complicated structure, but are splattered with regions of complicated structure, we may relax the goal of putting a (u, v) sample in absolutely every cell, and we may instead ask where we should put the (u, v) samples to do the most good.

4 A Global View of the MMA's Multiple Configurations

Many astronomers will want images with the highest angular resolution which will still give adequate brightness sensitivity in the amount of time the scheduling committee has granted them to observe their source. (There are exceptions to this statement, researchers looking for microwave background fluctuations of a certain scale, for example.) There is a natural tradeoff between resolution and brightness sensitivity. This tradeoff can be made continuously by tapering a single array configuration, or more efficient from an operational standpoint but less efficient from a scientific standpoint, discretely by switching among arrays of different size.[1] The combined effect of tapering and switching among different array configurations

[1]Another means of making this tradeoff is by observing for different amounts of time in different arrays. In general, all the MMA's configurations will have some very short spacings, permitting good single configuration imaging, but some experiments will require multi-configuration observations.

is shown schematically in Figure 6 for the cases of uniform and centrally condensed Fourier plane distributions. The solid straight line represents the brightness sensitivity as a function of resolution for some constant amount of observing time which would result if the MMA antennas could be continuously reconfigured to any resolution. The letters A, B, C, and D indicate where on this "optimal" line the resolution and brightness sensitivity of the actual arrays lie (assuming the resolution of each is separated by a factor of 4.0). The dashed, discontinuous lines indicate the result of tapering each array, ie, trading resolution for brightness sensitivity. The top dashed lines illustrate how a uniform Fourier plane distribution responds to tapering, and the lower dashed lines illustrate how a centrally condensed Fourier plane distribution respond to tapering. We do not consider tapering to a resolution lower than that of the next smaller array. As each smaller array enters in, we get a large improvement in surface brightness sensitivity as no tapering is used at full resolution. These lines represent the limits of detection, and sources with T_B below the lines will not be detected. The heavy curves represent the surface brightness as a function of resolution for three Gaussian sources of the same flux but of sizes 0.2, 0.5, and 1.0 arcseconds. To the right, each source is unresolved. As each source is observed at higher resolution, it eventually becomes undetectable, or "resolved out", by the arrays.

We assume that there is no *a priori* source size which is more likely or more important than other source sizes. Under this assumption, we would desire that our discrete array configurations be able to come as close to the solid straight line as possible. We can do this by maximizing the number of configurations which are financially and operationally feasible, and by optimizing the way in which the array sensitivity degrades as we taper to lower resolution.

We now compare the brightness sensitivity of the uniform and centrally condensed Fourier plane distributions for the same amount of observing time. Observing in a "C" configuration (resolution of about 1.1 arcsec), all three sample sources are firmly detected at full resolution and there is little difference between the two Fourier plane distributions. When the 1 arcsec Gaussian source is observed with the uniform "B" configuration (resolution of about 0.28 arcsec), it is basically undetected at all tapers. However, the 1 arcsec source is detected by the centrally condensed "B" array at beams larger than 0.6 arcsec. Either "B" configuration will detect the 0.5 arcsec source at full resolution. Both the uniform and condensed "A" configurations require some tapering to detect the 0.2 arcsec source, but the centrally condensed configuration requires less tapering, and can therefore image the 0.2" source at somewhat higher resolution than the uniform configuration.

5 Dynamic Range of Images

Above, we have addressed the issue of noise performance with respect to tapering the beam, which is pertinent for imaging weak sources. Another consideration in choosing a Fourier plane distribution for the MMA is the image dynamic range which can be achieved for very bright sources. Interferometric images of very bright sources are not limited by thermal noise, but by other systematic errors. VLA images of sources dominated by a single bright unresolved source are dynamic range limited at the level of a few 100,000:1 by unknown systematic effects. VLA images of more complicated sources dominated by extended emission are typically dynamic range limited at the level of 10,000:1, presumably due to deconvolution errors. It seems likely that we can design the MMA configurations to give high dynamic range deconvolution. The Fourier plane distribution of an array determines the character of the sidelobes in the point spread function, and a point spread function with small sidelobes encourages a high dynamic range deconvolution. Such reasoning would indicate that the uniform Fourier plane coverage would result in poorer images than the tapered coverage.

Simulations using the abstract Fourier plane coverages mentioned above without adding thermal noise have given conflicting results. I have simulated observations of our standard MMA source model, a planet model, and a random collection of point sources. The simulated data were *naturally weighted*, gridded, Fourier transformed, and deconvolved using both CLEAN and MEM. Uniform Fourier plane distribution results in images with dynamic range at full resolution which is 10-50% than images made from a centrally condensed Fourier plane distribution. As the visibilities are tapered, the uniform distribution images become comparable to the centrally condensed images after the resolution is degraded by 50-100%. However, Morita (private communication) has performed simulations using abstract uniform and moderately centrally condensed Fourier plane distributions which have been imaged using *uniform weighting*. Morita finds that the uniform Fourier plane coverage results in higher dynamic range images. These conflicting results are not necessarily in conflict since the uniform weighting will up-weight the sparsely sampled long baseline (u, v) points of the tapered Fourier distribution, possibly resulting in larger image reconstruction errors.

6 Antenna Based Errors

We have not investigated the difference of the effects of antenna based errors on images formed from uniform and centrally condensed Fourier plane distributions. It has been argued that since there is more flux, and since the visibility function is changing faster, on short baselines, an error on a short baseline will have a larger effect on the resulting image's quality than a similar fractional error incurred on a long baseline. By this argument, we might achieve better images if we give a measure of redundancy on the shorter baselines to reduce the effects of these errors by averaging. However, since there are more short spacings than in the uniform case, any given antenna based gain error will affect more short baselines than in the case of a uniform Fourier plane distribution. If the centrally condensed Fourier plane coverage is taken seriously, we should perform simulations to determine which distribution is superior with respect to antenna based gain errors.

7 Recommendations for MMA Configurations

The argument which drives the D configuration antenna layout is to maximize the surface brightness sensitivity, which requires a centrally condensed Fourier plane distribution, achieved by placing the antennas as close together as possible. No change in the current design is suggested for the D configuration. Until now, the antenna configuration for the larger arrays has been driven by the desire to provide uniform Fourier plane coverage, a desire which has not been supported by any analysis. I argue that the larger configurations should have moderately centrally condensed Fourier plane coverages to provide better shaped naturally weighted beams and better noise performance with respect to tapering. One possible exception is the largest array, which might be a ring-like array to get the most sensitivity at the resolution of the longest baselines.

The argument which drives the D configuration antenna layout is to maximize the surface brightness sensitivity, which requires a centrally condensed Fourier plane distribution, achieved by placing the antennas as close together as possible. No change in the current design is suggested for the D configuration. Until now, the antenna configuration for the larger arrays has been driven by the desire to provide uniform Fourier plane coverage, a desire which has not been supported by any analysis. I argue that the larger configurations should have moderately centrally condensed Fourier plane coverages to provide better shaped naturally weighted beams and better noise performance with respect to tapering. One possible exception is the largest array, which might be a ring-like array to get the most sensitivity at the resolution of the longest baselines.

Does the MMA still need four configurations? Absolutely! Even for the centrally condensed (linearly decreasing) Fourier plane distribution studied here, the C array will have a factor of 3 better sensitivity than the B array tapered to the C array's full resolution. An even more highly centrally condensed coverage would not lose sensitivity so quickly, but will result in a very bad synthesized beam with a wide plateau at its base, indicating that there is not enough sensitivity on the long baselines, and is undesirable.

Antennas randomly distributed over a circular region with a uniform deviate will give (u, v) coverage which approximates the centrally condensed distribution which we've been using. Some of the inner antenna pads for a given array may be common with the outer pads of the next smaller array. Even with tricks like this, the configurations with centrally condensed (u, v) coverage will require much more road and cables than the ring-like arrays currently under consideration. Tim Cornwell once suggested that the MMA arrays be built out of concentric rings of antenna stations, with two adjacent rings being populated each with half the antennas to produce a tapered (u, v) coverage. To obtain a monotonically decreasing Fourier plane density as a function of baseline, each array would have to be 3.0 times larger than its adjacent smaller array, which would require 5 different arrays to span from the smallest to the largest array. Such an array falls right in between the uniform and the centrally condensed linearly decreasing distribution for sensitivity loss as a function of taper, but due to sharp corners in it's three tiered "wedding cake" Fourier distribution, the image fidelity suffers a great deal in simulations. To what extent can the sharp edges be rounded off by a careful layout of the antennas on the two rings? We need to study this further.

For the future, we need to debate the merits of uniform and centrally condensed Fourier plane coverage for the MMA, clarify the shape of the optimal beam, clarify the use of multiple configurations in the MMA's imaging, further explore the various options for centrally condensed Fourier plane distributions, and design antenna configurations which result in the desired centrally condensed Fourier plane distribution and permit economical road, power, and communications layouts. Furthermore, we need to investigate the possibility of flexible layout of antenna pads which would permit arrays with several Fourier plane distributions.

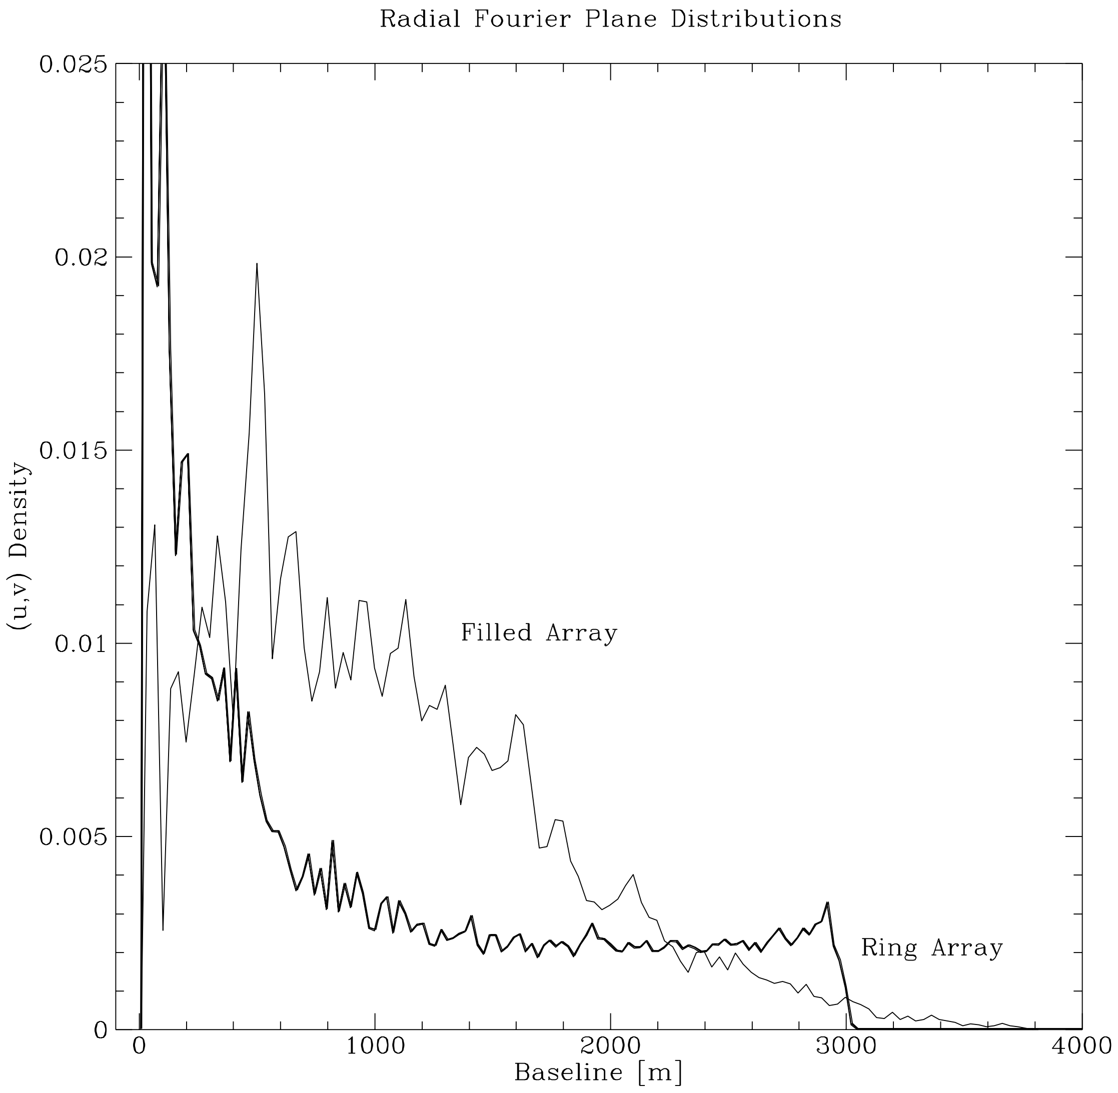

Figure 1: Radial Fourier plane distributions from actual ring and filled circle antenna configurations

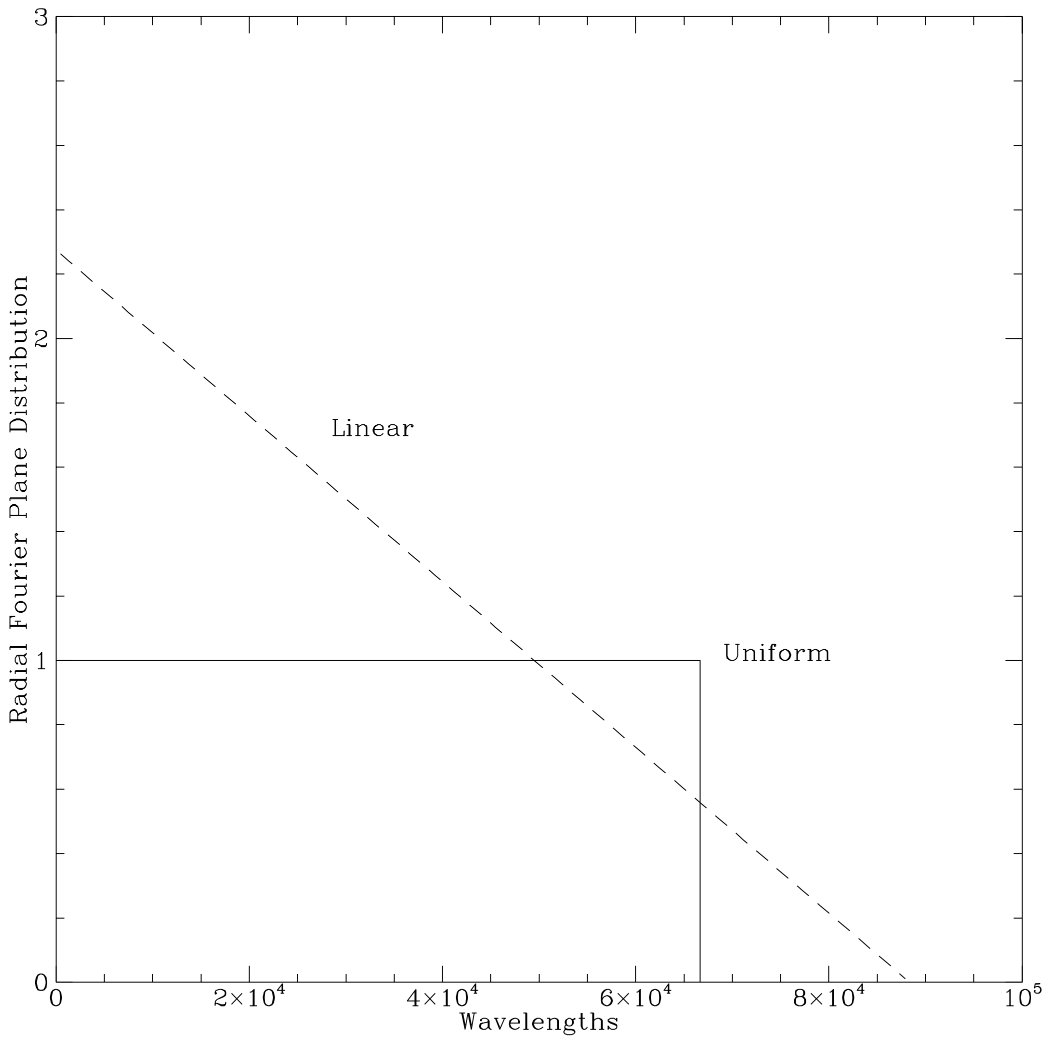

Figure 2: The uniform and centrally condensed (linearly decreasing) Fourier plane distributions used in this study. The radial cuts have been plotted such that the integral over the 2-D Fourier plane is the same for both distributions. Just arguing from geometry, trading a little bit of baseline density in the expansive outer part of the plane buys us a lot of baseline density in the cozy inner plane.

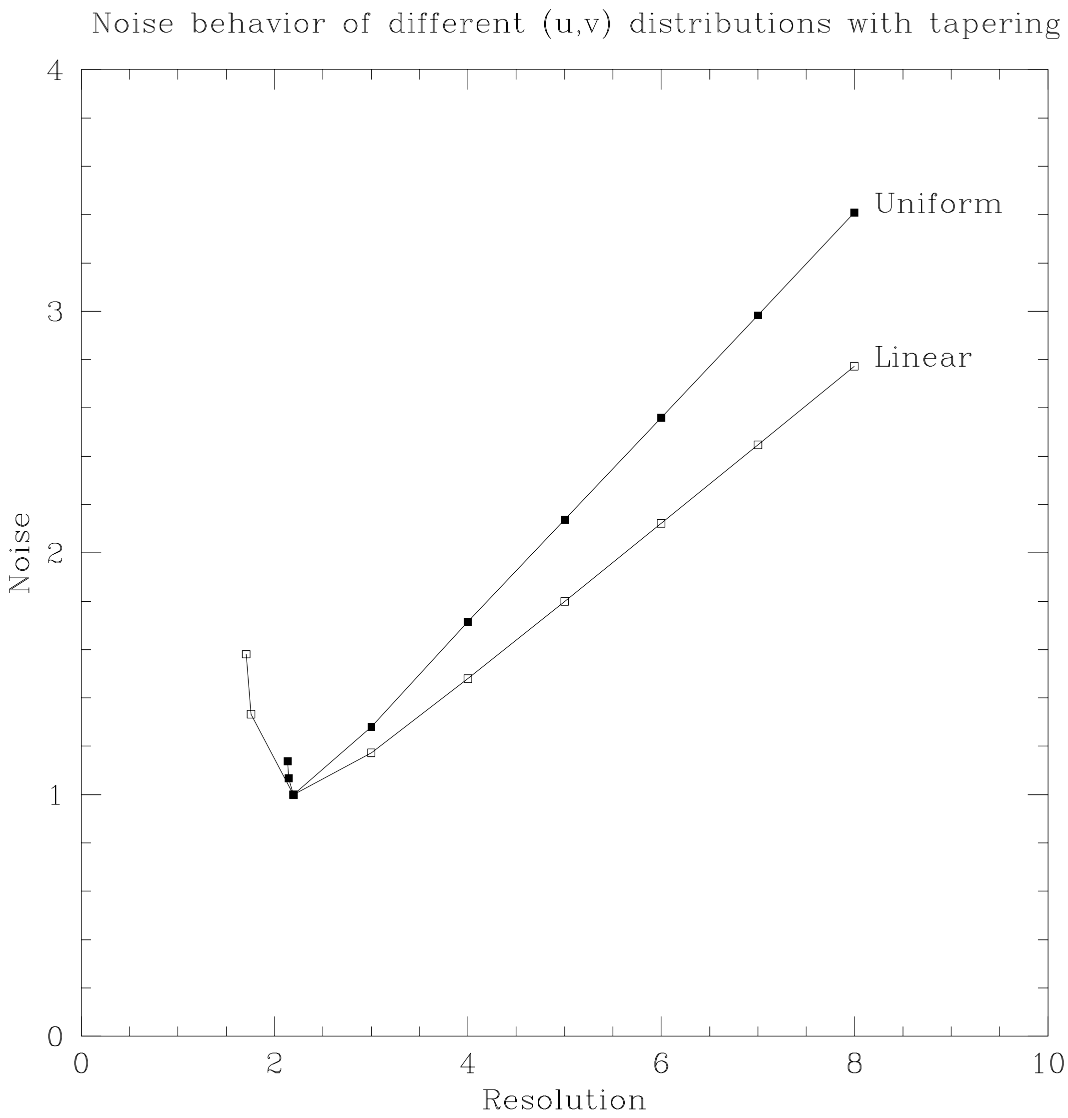

Figure 3: How does point source sensitivity degrade as we taper the Fourier plane coverage? We plot the point source sensitivity, normalized to full array sensitivity, as a function of resolution for uniform and linearly decreasing Fourier plane distributions. The upraised spurs at the highest resolutions are for robust and uniform weighting.

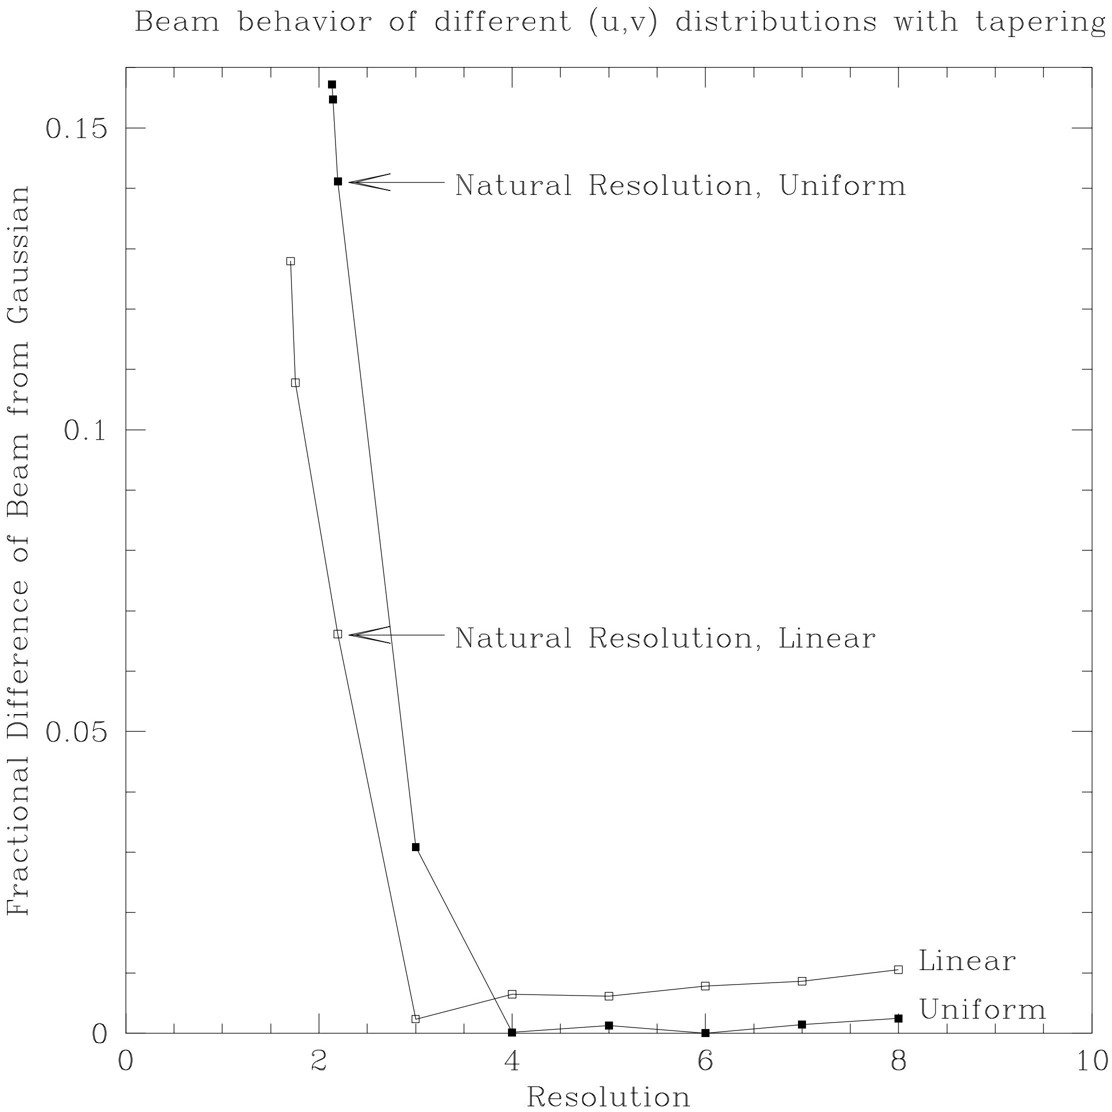

Figure 4: But will these Fourier plane distributions give me a nice beam? We plot the fractional difference of the beam integral and the integral of the fit Gaussian beam as a function of tapered beam size for the uniform and centrally condensed (linearly decreasing) Fourier plane distributions. Uniform Fourier coverage gives us a highly non-Gaussian beam at full natural resolution, due to its excess of long baselines. Uniformly weighted beams for both uniform and centrally condensed distributions are highly non-Gaussian.

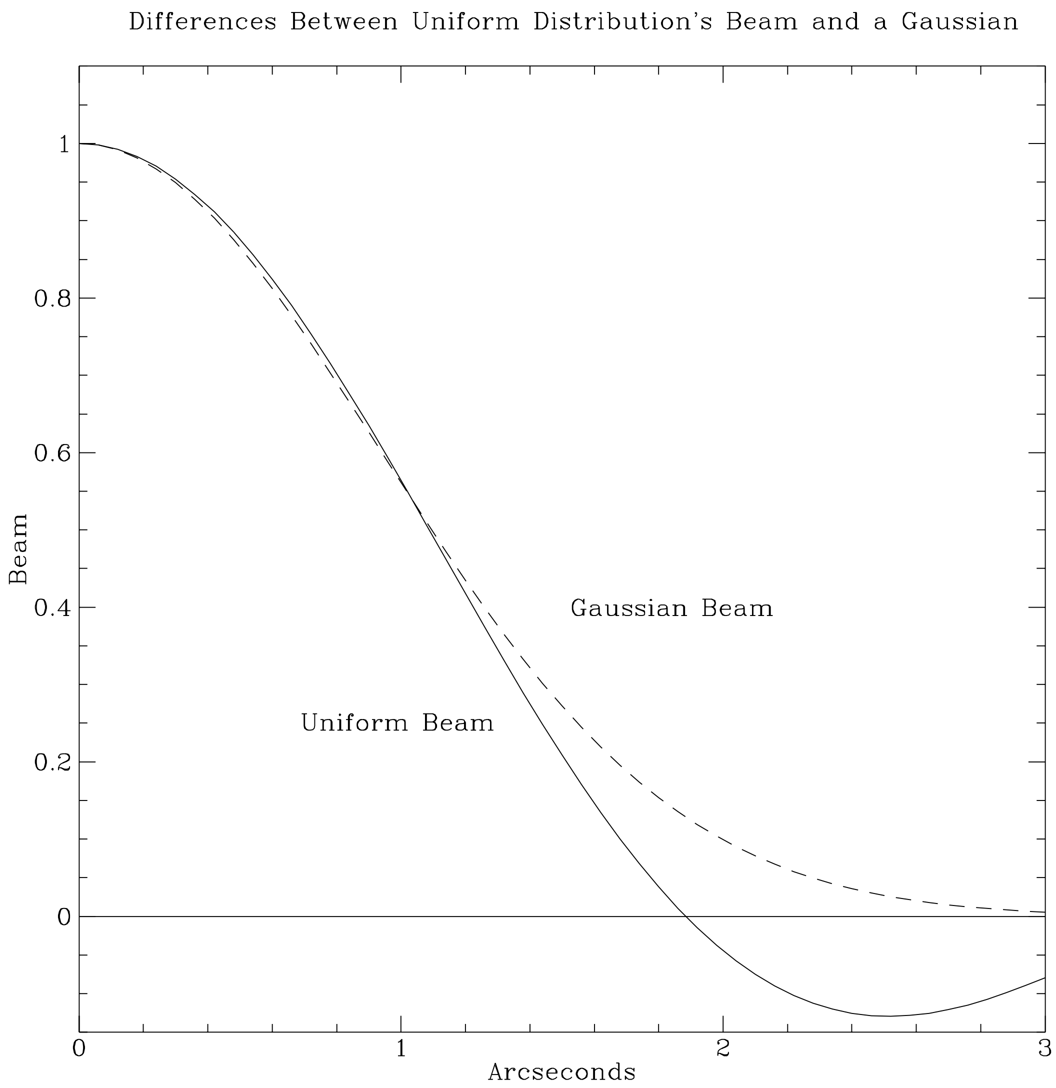

Figure 5: What's wrong with the uniform distribution's beam? Here are the radial profiles of the naturally weighted uniform Fourier distribution's beam and its best fit Gaussian. The uniform Fourier distribution has much too many long baselines to yield a good Gaussian beam. Conversely, in order to get a good Gaussian beam, the long baselines need to be tapered down, resulting in a loss of about 30% of the sensitivity.

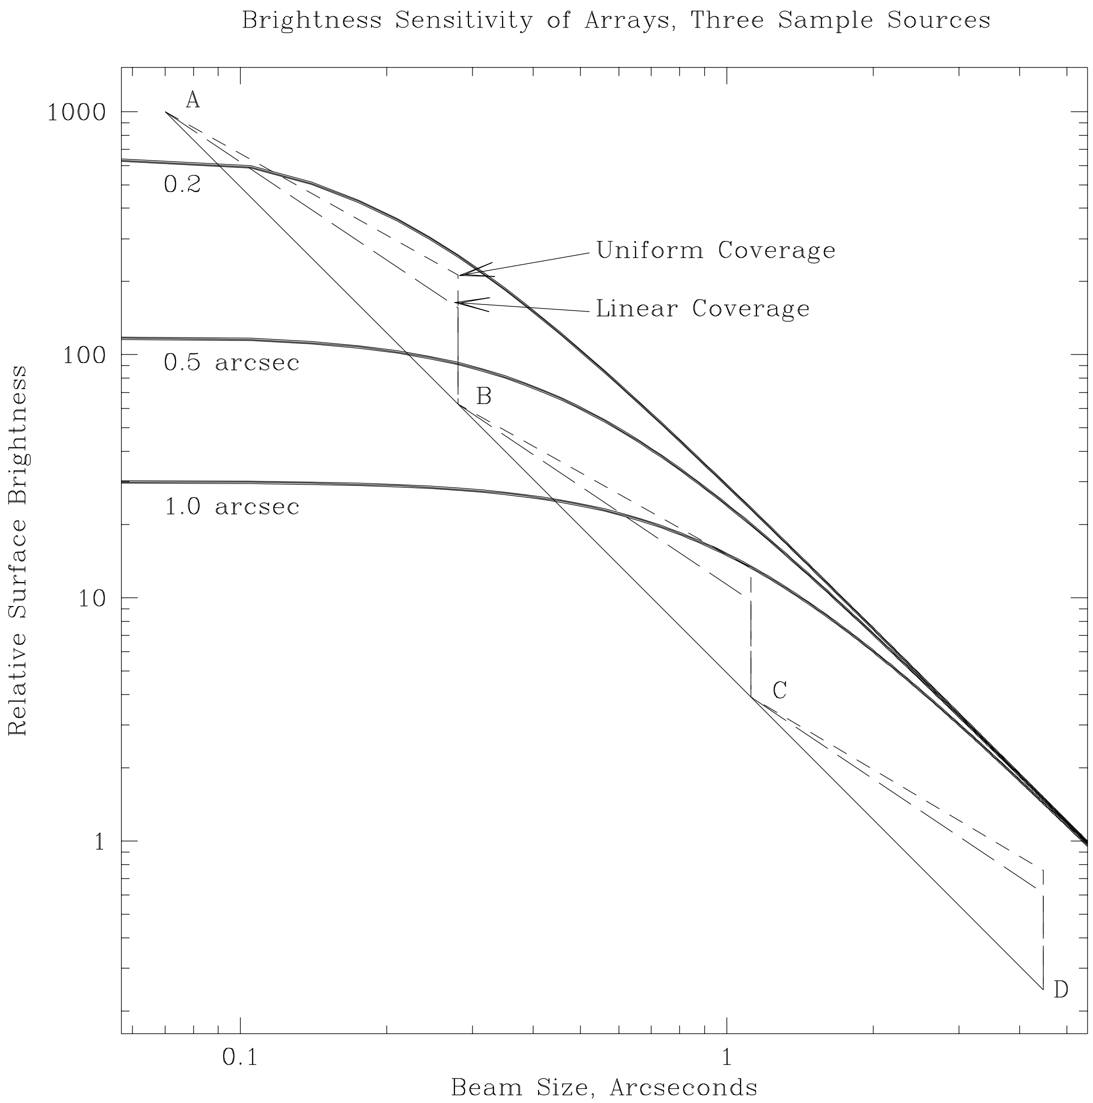

Figure 6: Schematic Log-Log plot of brightness sensitivity as a function of resolution for an infinitely reconfigurable array (solid straight line), uniform Fourier plane distribution versions of A, B, C, and D arrays subject to tapering (top dashed line), and centrally condensed Fourier plane distribution versions of A, B, C, abd D arrays (lower dashed line). The heavy curves represent surface brightness as a function of resolution for Gaussian sources of different sizes. The sources are detected when their curves lie above the array sensitivity lines.