The VLA Upgrade Project Memo Series

Number 7

Investigating Antenna Locations for the VLA A+ Array

M. A. Holdaway
Rick Perley

10 June 1996

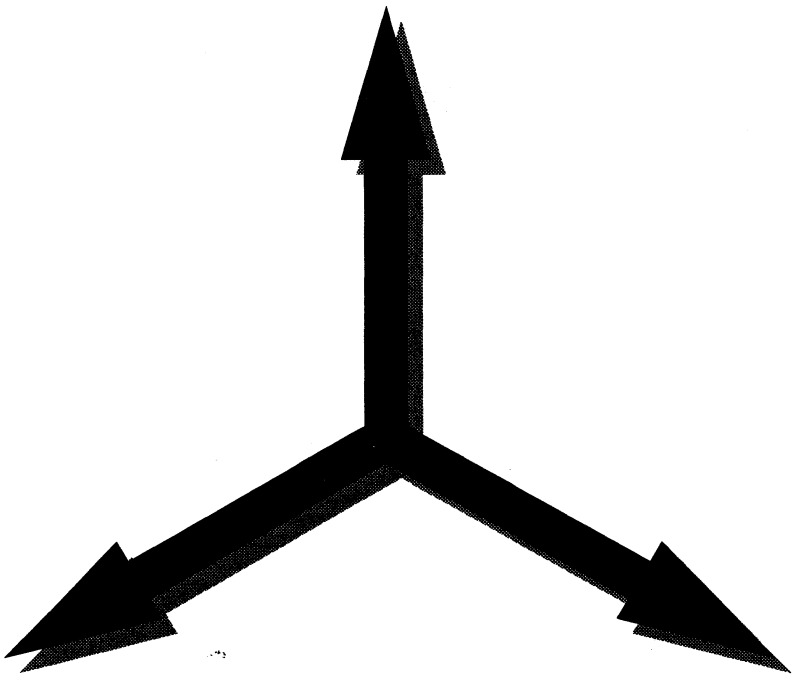

Returning the Instrument to the State of the Art

National Radio Astronomy Observatory

Investigating Antenna Locations for the VLA A^+ Array

M.A. Holdaway
National Radio Astronomy Observatory
949 N. Cherry Ave.
Tucson, AZ 85721-0655
mholdawa@nrao.edu

Rick Perley
National Radio Astronomy Observatory
1003 Lopezville Rd
Socorro, NM 87801
rperley@nrao.edu

June 10, 1996

Abstract

We explore a few possible algorithms for finding four to six antenna positions for the VLA Upgrade's A^+ configuration, which will bridge the gap between the VLA's A configuration and the inner VLBA antennas. A modified neural networks approach (Keto, 1992) does not seem to work well for this problem in which most antennas are fixed and a centrally condensed (u, v) coverage is desired. We develop an algorithm based on random assignment of antenna positions and antenna position wiggles of decreasing size, with configurations differentiated by the size of holes in the (u, v) plane and the rms sidelobe level of the point spread function. We used this algorithm to generate 300 km configurations with visually appealing Fourier plane coverage and full track point spread functions with RMS sidelobes below 1%. The best configurations made with four upgrade antennas are significantly better than Walker's (VLBA Project Book, 1988) four antenna configurations (admitedly an unfair test since Walker's design criteria was to produce a better VLBA). While adding more antennas to the A^+ array steadily improves the configuration quality, the improvements are incremental. We also find that a 100 km A^+ array has only marginally better brightness sensitivity than a 300 km A^+ array tapered to the same resolution.

1 Introduction

The arrangement of antennas in an interferometric array is very important for the quality of the images which the array will produce, but as yet, researchers cannot give an unambiguous

answer as to what the arrangement should be. The VLA's antenna configurations were based on misconceptions about how the imaging would be performed and on how images at different frequencies would be compared. The VLBA's antenna configuration was chosen based on the desire to fill the Fourier plane reasonably well, with emphasis on the shorter baselines which were conspicuously absent from *ad hoc* VLBI arrays. (Holes in the inner part of the Fourier plane can affect image quality more than holes in the outer part of the Fourier plane.) Both of these arrays produce highly centrally condensed Fourier plane distributions, which result in low resolution naturally weighted images with optimal noise, or high resolution uniformly weighted images with noise which is much higher than optimal.

Cornwell (1986) and Keto (1992) argued for uniform Fourier plane coverage, which led them to ring-like array configurations. Such arrays result in (u, v) coverage which attempts to sample as many regions of the Fourier plane as possible. The naturally weighted and uniformly weighted beams of these uniform Fourier plane distributions are very similar, so no noise or resolution is lost either way. However, another way to overcome the tradeoff between resolution and noise for centrally condensed Fourier plane coverages is through robust weighting (Briggs, 1995).

As there is an unfortunately large gap between the VLA and the inner VLBA, the VLA upgrade plans to add a few antennas to provide continuous Fourier plane coverage from the shortest VLA spacings to the longest VLBI spacings, permitting constant resolution imaging capabilities for a wide range of frequencies and resolutions, which is essential for many scientific problems. However, choosing optimal locations for these new antennas is a very tricky problem, as we must clarify how the antennas will be used with the VLA and the VLBA.

The inner VLBA antennas and the new upgrade antennas will probably be connected to the VLA correlator via fiber optic cables, and their signals will be correlated in real time, without observing with the other VLBA antennas. We call this the A^+ array. Hence, the first requirement for the upgrade configuration is that *the VLA A^+ array must provide good imaging characteristics in a stand-alone mode.* We arbitrarily choose to optimize the A^+ array for baselines out to 300 km (about 8.5 times higher resolution than the VLA's A array).

The new upgrade antennas and the VLA (in phased array mode or up to four individual VLA antennas) will also be used with the entire VLBA, so the second requirement for the upgrade configuration is that *it must substantially enhance the imaging capability of the VLBA.* We do not address this issue here.

The centrally condensed Fourier plane distribution of both the VLA and the VLBA make the implicit statement that holes in the inner part of the Fourier plane are more damaging than holes in the outer part of the Fourier plane. The visibility function often changes more quickly in the inner Fourier plane, so holes in the inner Fourier plane result in more damage to images than holes in the outer part of the Fourier plane. In considering the placement of the upgrade antennas, we should try to stick with this philosophy rather than attempt to uniformly fill the gaps in the Fourier plane.

Finally, we are not attempting to generate the final placement of the VLA upgrade antennas. Rather, we are trying to explore how the problem may be solved. There are a number of parameters, such as the size of the A^+ array or the manner in which we weight the holes as a

function of position in the (u, v) plane, which need to be refined in future work. Also, there are places where it is clearly not acceptable to place antennas, such as in the middle of the Gila Wilderness, and the algorithms used have not attempted to deal with this problem. In any event, this work may inspire someone to invent a better algorithm for suggesting antenna placement.

2 Keto's Neural Network Algorithm

Keto (1992) has written a neural network code which iteratively adjusts the antenna locations until a uniform Fourier plane coverage is achieved. The algorithm picks a point in the Fourier plane randomly within some maximum (u, v) limit and pulls the closest (u, v) sample towards that point, adjusting the coordinates of the antennas correspondingly. This process is repeated thousands of times, and with each iteration, the antennas become stiffer and move smaller distances towards the random points. The end result is a Fourier plane coverage in which any point in the Fourier plane is close to a (u, v) sample.

In order to apply Keto's algorithm to the VLA upgrade problem, we had to make a few changes:

- the positions of the existing VLBA and VLA antennas are fixed. They contribute to the (u, v) coverage, but do not participate in the iterative antenna dance.
- because of the very sparse (u, v) coverage in VLBI, we considered long ($\sim$12 hour) tracks through the Fourier plane (Foster, 1996). (Keto's original algorithm optimized for a snapshot at the zenith.)
- because the VLA/VLBA system has a highly centrally condensed Fourier plane coverage, we do not seek uniform coverage but a Gaussian coverage. Actually, the VLA/VLBA system (with only four VLA antennas) has a Fourier plane density which varies radially approximately as $r^{-1.6}$, but we use a Gaussian of half width 160 km because it was simple to program. We attempted to recover a Gaussian (u, v) distribution by distributing the "random" points picked in the Fourier plane with a Gaussian distribution (ie, a Gaussian "picking function").

After experimenting with the modified Keto algorithm, it became clear that there are implicit forces operating which bias the solutions towards placing antennas on some sort of ring. The upgrade antennas usually were placed in a partial, distorted ring about 150-200 km away from the VLA, each separated by about 50 km. We obtained similar results for uniform picking functions and Gaussian picking functions of a variety of widths. These configurations produce (u, v) coverages which had some large patches of very uniform coverage adjacent to large patches at similar radial distances which were very sparsely covered. As we will see below, the Keto configurations are quantitatively much worse than configurations obtained through other less elegant methods.

While our experience with the Keto algorithm was unsatisfactory, there are probably further modification that could be made to the algorithm which could improve its performance on this problem.

3 Cornwell's Simulated Annealing Algorithm

Cornwell (1986) wrote a simulated annealing algorithm which moves antennas about, calculates the "energy" of the resulting Fourier coverage (defined as the potential energy of the system of charged, repelling particles at the antenna locations, each with 1/r potentials), subject to the constraint that the antennas lie within a circular region. Lower energy configurations were always accepted, and higher energy configurations were accepted with a probability of $e^{-E/kT}$, where E is the configuration energy and T is a fictitious temperature. As the iterations proceed, the antennas are moved by smaller amounts and the temperature slowly drops, hopefully resulting in a nearly minimum energy configuration. The possibility of accepting higher energy configurations permits escape from local minima in the energy.

We did not even try the Cornwell algorithm. Like the Keto algorithm, it would have placed the upgrade antennas in a ring specified by the circular region within which we were willing to place antennas.

Again, some modification of the Cornwell algorithm, such as weighting the holes in the inner Fourier plane more highly than the holes in the outer Fourier plane, would probably improve its performance in this problem.

4 The "Modified Random" Algorithm

Since being clever doesn't seem to help, perhaps its time to try being dumb. Simply searching all possible locations for the four or so upgrade antennas is prohibitively time consuming. However, algorithms such as simulated annealing don't aim at getting the very best solution, but at getting one of many nearly optimal solutions. Typically, the minima of the functions we are trying to optimize are broad and shallow, so we don't need to find the very best solution to get comparable performance.

We selected random locations within 250 km of the VLA for the upgrade antennas. In order to evaluate the success of these random configurations, we investigate both the quality of the uniformly weighted point spread function (PSF) generated from the (u, v) data with a 300 km taper, and the Fourier plane coverage. While there is a general trend between the quality of the PSF and the quality of the Fourier plane coverage (smaller Fourier holes generally result in smaller PSF sidelobes), some low-hole configurations (such as a completely uniform coverage with a sharp maximum baseline cutoff) will produce large sidelobes.

The quality of the PSF is gauged by the rms sidelobe level outside of a mask which covers the main beam. The quality of the Fourier plane coverage is gauged by the sum of the size of the holes in the Fourier plane on baselines between 30 and 300 km on a 10 km grid, weighted by the radial (u, v) distance to the -1.5 power. In order to be computationally fast, we approximate

the size of a hole at a given (u, v) cell as the number of empty cells in the u direction times the number of empty cells in the v direction. Since a similar number will be repeated for each cell which lies in a given hole, we are actually weighting the holes as the square of the hole area, which strongly discourages large holes. As we mentioned above, a hole in the center of the Fourier plane has a more sever affect on image reconstruction than a hole of the same size in the outer Fourier plane. To account for this, we have weighted the holes as the radial (u, v) distance raised to the -1.5 power. We need to weight the holes by a power less than -1.0, since this power just compensates for the bias towards long baselines which, through earth rotation synthesis, span more Fourier cells and hence reduce more holes than a short baseline would. To arrive at a single number, we integrate the weighted hole image between the range of interest. Finally, we weighted this (u, v) hole measure such that it would be about twice as large as the RMS sidelobe level, so the sidelobe level basically distinguishes between different configurations of similar Fourier plane coverage quality.

It is simple to include different source declinations in the optimization. We took $\delta = -15°$ to be representative of the range $-30° < \delta < 0°$, with a weight of 0.37, $\delta = 15°$ to be representative of the range $0° < \delta < 30°$, with a weight of 0.37, and $\delta = 45°$ to be representative of the range $30° < \delta < 60°$, with a weight of 0.26. We ignored the north polar cap since it has a small fraction of the visible sky. The weights are proportional to the fraction of sky which falls into each declination range.

As we increase the number of "random" iterations, we sample the possible phase space of antenna locations better. In order to sample the antenna plane of size r_{max} to a resolution of d with N_{ant} movable antennas, the number of iterations N_{iter} must exceed

$$N_{iter} > ((r_{max}/d)^2)!/\left(((r_{max}/d)^2 - N_{ant})!N_{ant}!\right). \quad (1)$$

With $N_{ant} = 4$, this expression leads to about 2000 iterations to sample to a resolution of $r_{max}/4$. Obviously, this will take forever to get any sort of resolution. However, if we take a few of the best configurations which the "random" iterations have found and wiggle the antennas randomly by order d, we can proceed much more quickly. After wiggling the antennas for many iterations, we again select a few of the best configurations and cut the size of the wiggle. We name these iterations the "wiggling" stage of the algorithm. How many "wiggling" iterations of wiggle amplitude d will approximately sample the antenna subspace such that we can decrease the wiggle by a factor of δ? Approximately

$$N_{iter} > (\delta^2)^{N_{ant}}, \quad (2)$$

or about 256 iterations for $\delta = 2$ and $N_{ant} = 4$. It is most efficient to decrease the wiggles by small δ factors and take more wiggle stages. Also, several wiggle stages are more effective at sampling the space of antenna locations than the random stage, assuming that the wiggle stages do not get trapped in a local minimum.

Being trapped in a local minimum is a real concern. If we had simply made a uniform grid of some coarse resolution and tried every possible combination of antenna locations, then reduced the resolution and performed systematic (ie, non-random) antenna "wiggles", we would

certainly have gotten stuck in a local minimum. It is our hope that by randomizing the antenna locations, running more iterations than would seem required, and by exploring several of the best configurations at each stage, we will jump out of many of the local minima. However, the bottom line is that the "best" configuration is not all that much better than a "good" configuration, as is demonstrated by our results below. Indeed, most of the configurations with truly poor (u, v) coverage and PSF sidelobes are the result of placing one or more of the upgrade antennas right next to another antenna. So, to first order, the rule of thumb "Don't do anything dumb" is good enough to prevent you from making a bad configuration.

While our algorithm really is dumb, and does take a good deal of cpu time to run, it was quite simple to program. It consists of a unix shell script which calls SDE tasks, and took about 20 minutes to write, once we decided what we wanted to do. Trading computer time for astronomer time isn't so dumb after all.

5 Results of the Modified Random Algorithm

If you are confused about what exactly the Modified Random Algorithm does or how it works, ignore the last section and look at some of the results. For the four antenna case, we ran 2000 iterations in the random stage, and 500 iterations per starting configuration in each of the 80 km, 40 km and 20 km wiggle stages. The three best configurations from each stage were used as starting configurations for the subsequent stages. Figure 1 shows the best 10 configurations made with four upgrade antennas at the end of the random stage, the 80 km wiggle, the 40 km wiggle, and the 20 km wiggle stages. Existing VLA and VLBA stations are marked with a filled square while upgrade stations are marked with an empty square. The 10 best configurations from each stage are laid on top of each other to show the range of possibilities at each stage. Initially, the upgrade antennas are widely distributed, but as the wiggle amplitude becomes smaller, the upgrade stations settle down into variations on a single state, even though multiple starting configurations from the previous stages were carried through to each stage.

Another way of looking at the results of our algorithm is to plot the rms sidelobe level against the net weighted (u, v) hole size for each configuration. Since we have created thousands of configurations, such a plot is difficult to look at and to print, so we have made a schematic figure which indicates the boundaries of (u, v) hole-rms sidelobe space within which the configurations in each stage are found. Figure 2 shows that the four upgrade antenna configurations produced by the random stage span a large range in quality. The very bad configurations tend to be redundant, placing upgrade antennas next to each other or next to existing antennas. Each stage of wiggle iterations produces a small improvement over the last stage, and the level of improvement decreases as the wiggle amplitude decreases.

So, while we cannot be certain that we have found the "best" configuration (in fact, it is clear that we haven't), we can state that we have found a good configuration which is considerably better than many others. For example, our best configuration is considerably better than the result of the modified neural networks configuration and Craig Walker's configuration

(Figure 2). Craig Walker was not trying to solve the same problem as we are trying to solve, so it is not at all surprising that his configuration is not as good as our best configurations. Also, Walker's principle motivation was to obtain good, uniform (u, v) coverage, which explains the location of his configuration on our plot, indicating good (u, v) coverage but not so good PSF sidelobes.

Figure 3 is a schematic of the entire range covered by the very best through the very worst configurations from all stages, displayed for the four, five, and six upgrade antenna cases. The bottom left hand tips of the regions dip to lower and lower rms sidelobe and integrated (u, v) hole levels as we increase the number of upgrade antennas, but the improvements are incremental. Hence, you would need at least seven antennas to reduce the rms sidelobe and (u, v) hole level to be half of the level for the best configuration produced with four upgrade antennas. On the other hand, the best configuration with four upgrade antennas is considerably better than many six upgrade antenna configurations, so wise placement of antennas is worth many antennas.

The ten best configurations obtained from the final stage (20 km wiggles) for the four, five, and six upgrade antenna cases are shown in Figures 4. Table 1 lists the antenna longitude, latitude, and site name for the best configurations obtained for the four, five, and six upgrade antenna cases. Each upgrade antenna can be moved 10-20 km for better access to roads and power without significantly affecting the Fourier plane coverage or properties of the point spread function. A few of the antenna sites, such as Elk Mountain in the Gila wilderness, the Gila Cliff Dwellings, and the Black River site in Arizona, are clearly not possible. It would not be difficult to create an antenna plane mask which prohibited placing antennas in wilderness areas or on Indian reservations. Since many configurations have very similar beam and (u, v) hole characteristics, we can find excellent configurations which are environmentally acceptable.

The Fourier plane coverages and beams for -15, 15, and 45 degree declinations for the best configuration with four, five, and six upgrade antennas are shown in Figures 5 through 7.

6 Brightness Sensitivity

It is somewhat arbitrary that we have chosen to optimize for a 300 km array. Such an array size is a compromise between extending the capabilities of the VLA and extending the capabilities of the VLBA. However, it can be argued that a $\sim$100 km configuration which extends the VLA's resolution by a factor of about 2.85 (the expansion factor between adjacent VLA configurations) is the right way to spend extra antennas. Such an array would have superior brightness sensitivity, Fourier coverage, and imaging capabilities for both strong and weak sources, and should be considered.

We present an example 100 km configuration made by adding five upgrade antennas to Pie Town and the VLA A Array (Figures 8 and 9). We have not gone through any elaborate optimization process for this configuration, just laid the antennas down by eye. However, the Fourier pane coverage is very good, with most gaps being the same size as the distance between A array antennas at the arm ends. The few larger gaps could undoubtedly be fixed

Four Upgrade Antennas		
Longitude	Latitude	Site
-109.06571	33.77203	Alpine, AZ
-105.91119	32.95187	High Rolls, NM
-106.15098	33.44207	Three Rivers, NM
-109.10673	33.31532	Glenwood, NM
Five Upgrade Antennas		
Longitude	Latitude	Site
-105.32571	33.32017	San Patricio, NM
-109.94344	33.52937	Black River, AZ
-110.00505	34.14544	Pine Top, AZ
-107.21213	32.18726	Akela, NM
-108.51912	33.57188	Elk Mountain, NM
Six Upgrade Antennas		
Longitude	Latitude	Site
-105.41932	33.06540	Elk River, NM
-110.48065	33.37835	San Carlos, AZ
-109.63974	34.15905	Green's Peak, AZ
-107.44289	32.02733	Akela, NM
-108.11215	33.28230	Gila Cliff Dwellings, NM
-108.78598	33.15079	Buckhorn, NM

Table 1: Longitude and Latitude of antenna sites for four, five, and six upgrade antenna configurations. These antenna locations could be moved by 10-20 km without significantly affecting the quality of the configurations.

via optimization. Figure 10 shows the Fourier plane coverages and beams for comparison with the 300 km configurations.

The surface brightness noise is proportional to the point source noise divided by the beam area. Since the point source noise is constant among arrays of different sizes, a small configuration with a larger beam will have lower brightness noise (ie, the 100 km array has lower brightness noise than the 300 km array). However, to compare these two arrays on an equal footing, we must compare the brightness noise at the same resolution. As we taper in the Fourier plane to obtain a larger beam, we throw away some long baseline visibilities and our point source noise increases. However, our beam area increases faster than the points source noise increases, and tapering results in higher brightness sensitivity, but not as high as if no tapering were required.

The A^+ array is in an interesting situation with regards to brightness sensitivity. Most of the sensitivity of the array is in the center of the Fourier plane (ie, the VLA in its A array),

and very little sensitivity is in the region beyond 35 km. In order to achieve good resolution, we must use some sort of uniform weighting. In this case pure uniform weighting drastically down weights the inner Fourier plane and throws away much of the array's sensitivity. Robust weighting (Briggs, 1995) is a hybrid weighting scheme which preserves much of the resolution of uniform weighting and much of the sensitivity of natural weighting, and something like robust weighting is precisely what the A^+ array requires. Robust weighting still needs to down weight the central part of the Fourier plane to get high resolution. As we taper to get lower resolution, robust weighting needs to down weight less, and the point source noise can actually *decrease*, therefore making the brightness sensitivity improve more than expected. Unlike tapering naturally weighted data, the brightness sensitivity improvement which results from tapering robustly weighted data is not a smooth function of resolution. Figure 11 shows the point source noise as a function of resolution for the 300 km and 100 km arrays with five upgrade antennas. These point source sensitivities are generated for 12 hour tracks using 1 GHz bandwidth in two polarizations, using current 8 GHz T_{sys} and 27 VLA antennas, 5 upgrade antennas, and the PT, LA, FD, and KP VLBA stations, tapered accordingly to achieve the plotted resolution. Placing the antennas into a more compact (100 km) A^+ array results in surface brightness sensitivity which is only marginally better than the 300 km A^+ arrays. The image fidelity of a 100 km A^+ array will probably be better than that of a 300 km array at the same resolution.

7 Discussion

We will use the section to round up all the loose ends we mentioned in the text. While we cannot cut them off yet, at least we can put them in one place:

- What should be the maximum (u, v) baseline over which the A^+ array is optimized? This work focused on 300 km arrays, but also looked at one 100 km array.
- While a 100 km array has a lot more 35-100 km baselines than the 300 km array, the brightness sensitivities of the two arrays at the same resolution are actually not very different.
- How should we evaluate the combination of the VLBA plus the VLA A^+ array?
- How should we weight the importance of (u, v) holes as a function of radial (u, v) distance? We weighted our holes as the radial (u, v) distance raised to the -1.5 power.
- How many upgrade antennas do we need? While adding antennas incrementally produces incremental improvements in the objective configuration quality measures, the (u, v) coverages produced with six upgrade antennas subjectively look a lot better than the four antenna (u, v) coverages.

- We need to consider where we can and cannot realistically place antennas during the optimization, rather than perform an unconstrained optimization and then try to fit the resulting array to realistic antenna locations.
- What other criteria should be used to gauge the effectiveness of a configuration? Our "algorithm" is flexible enough that we could incorporate anything which was not too cpu intensive.
- Should we devise a more intelligent algorithm?
- Any proposed configurations should be put to the test by simulating realistic data and comparing the imaged simulation data with the true brightness distribution from which the data were generated.

References

Briggs, D.S., 1995, Thesis, "High Fidelity Deconvolution of Moderately Resolved Sources", New Mexico Institute of Mining and Technology, Socorro, NM.
Cornwell, T.J., 1986, MMA Memo 38, "Crystalline Antenna Arrays".
Foster, Scott M., 1996, MMA Memo in preparation.
Keto, Eric, 1992, SMA Memo, "Cybernetic Design for Cross-Correlation Interferometers".
Very Long Baseline Array Project Book, Version 7 1988, NRAO, Charlottesville VA, pp 1-1 to 1-20.

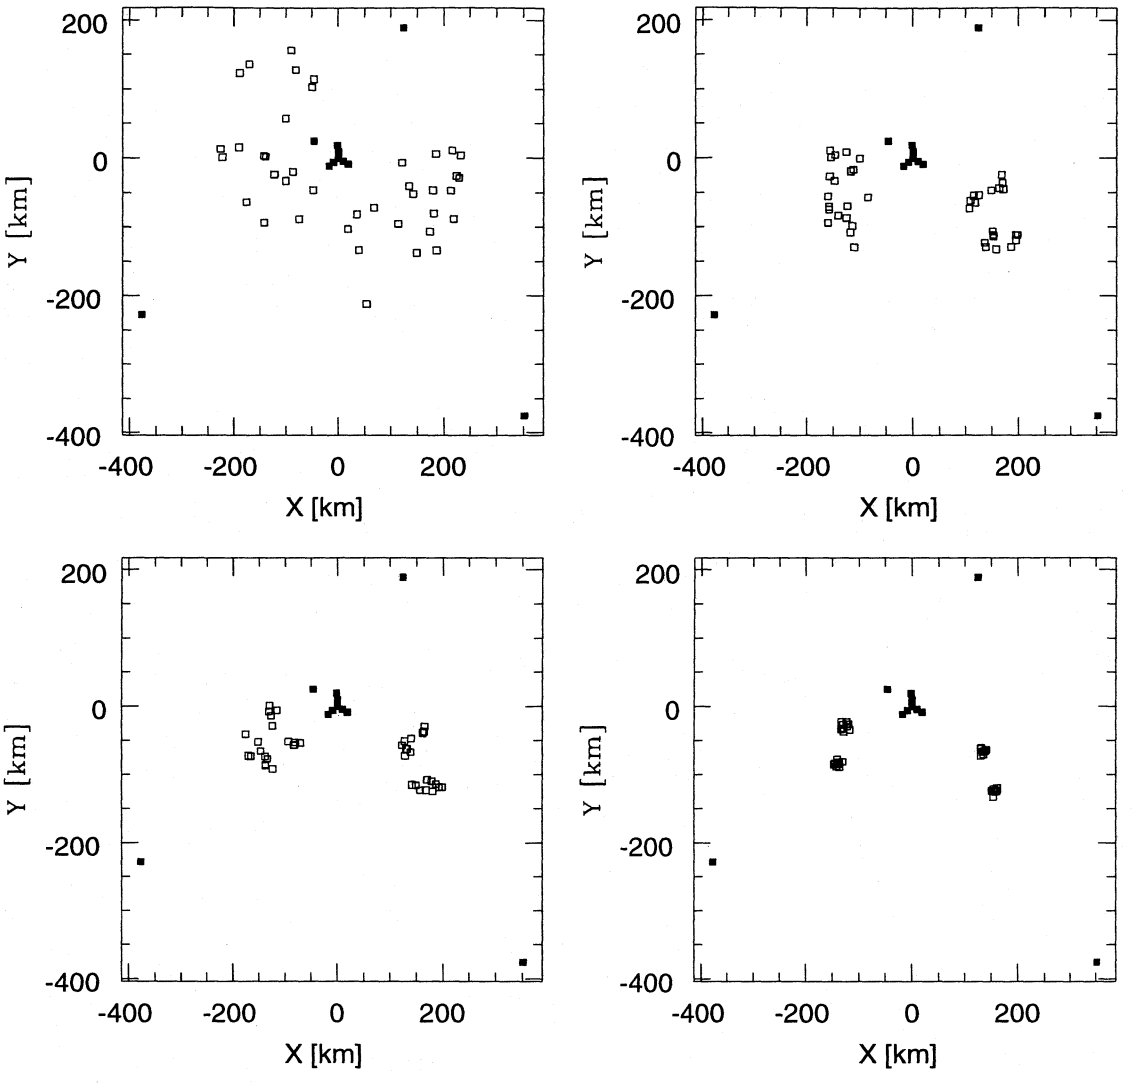

Figure 1: The stages of our antenna placement algorithm are illustrated by taking the best 10 four antenna configurations from each stage and displaying them atop each other. The filled boxes are existing antennas and the open boxes are possible locations for the four upgrade antennas. Upper left: 10 configurations after the random stage; upper right: 10 configurations after wiggling by 80 km; lower left: 10 configurations after wiggling by 40 km; lower right: 10 configurations after wiggling by 20 km.

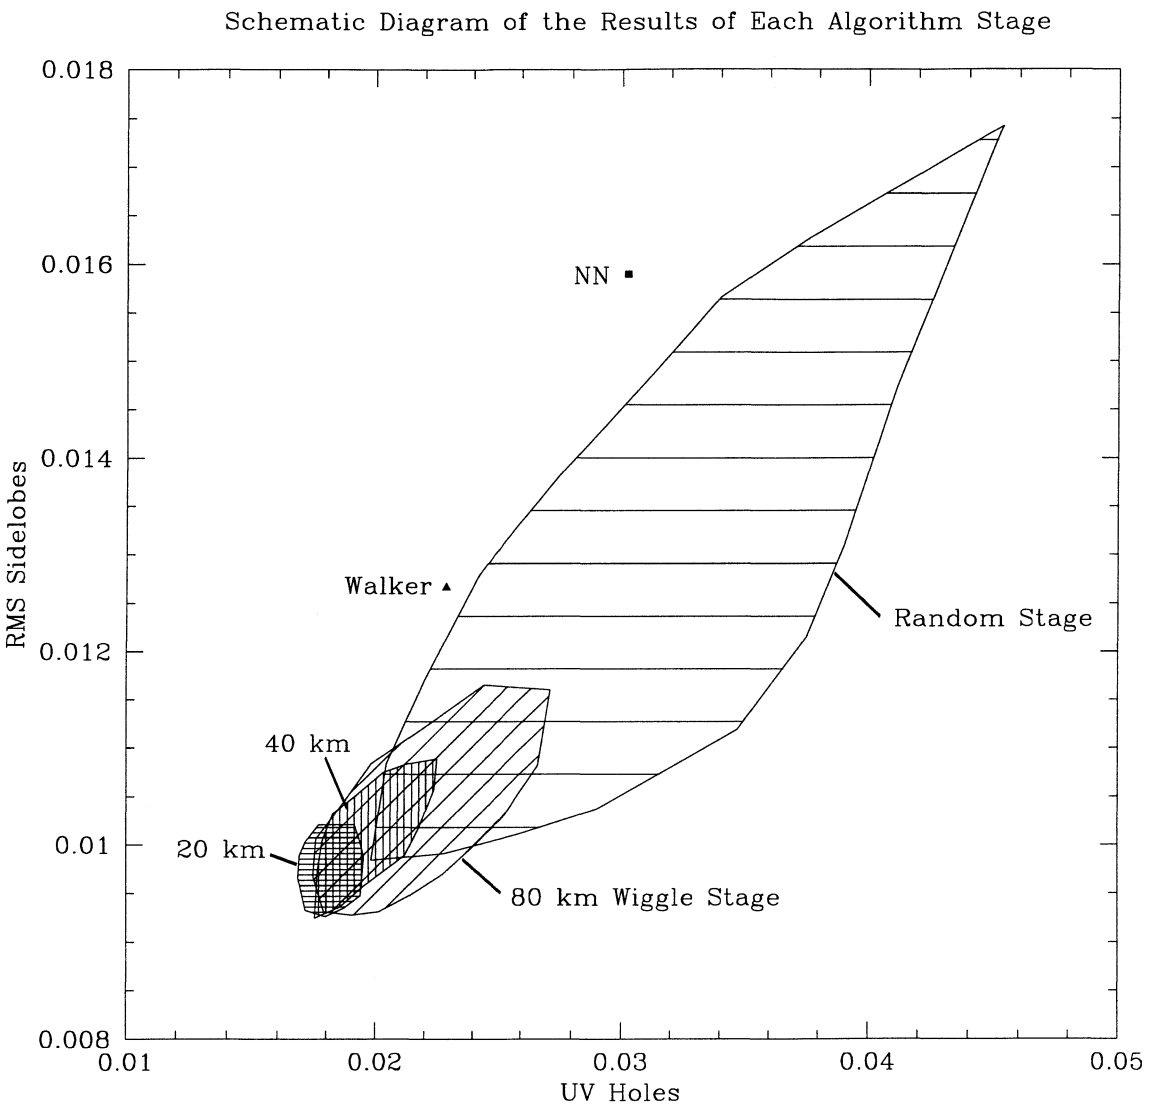

Figure 2: The range in quality of the Fourier coverage and point spread functions for all generated four antenna configurations is schematically displayed for the four stages of our antenna placement algorithm. Also listed are Walker's configuration and the neural network (NN) configuration discussed in the text.

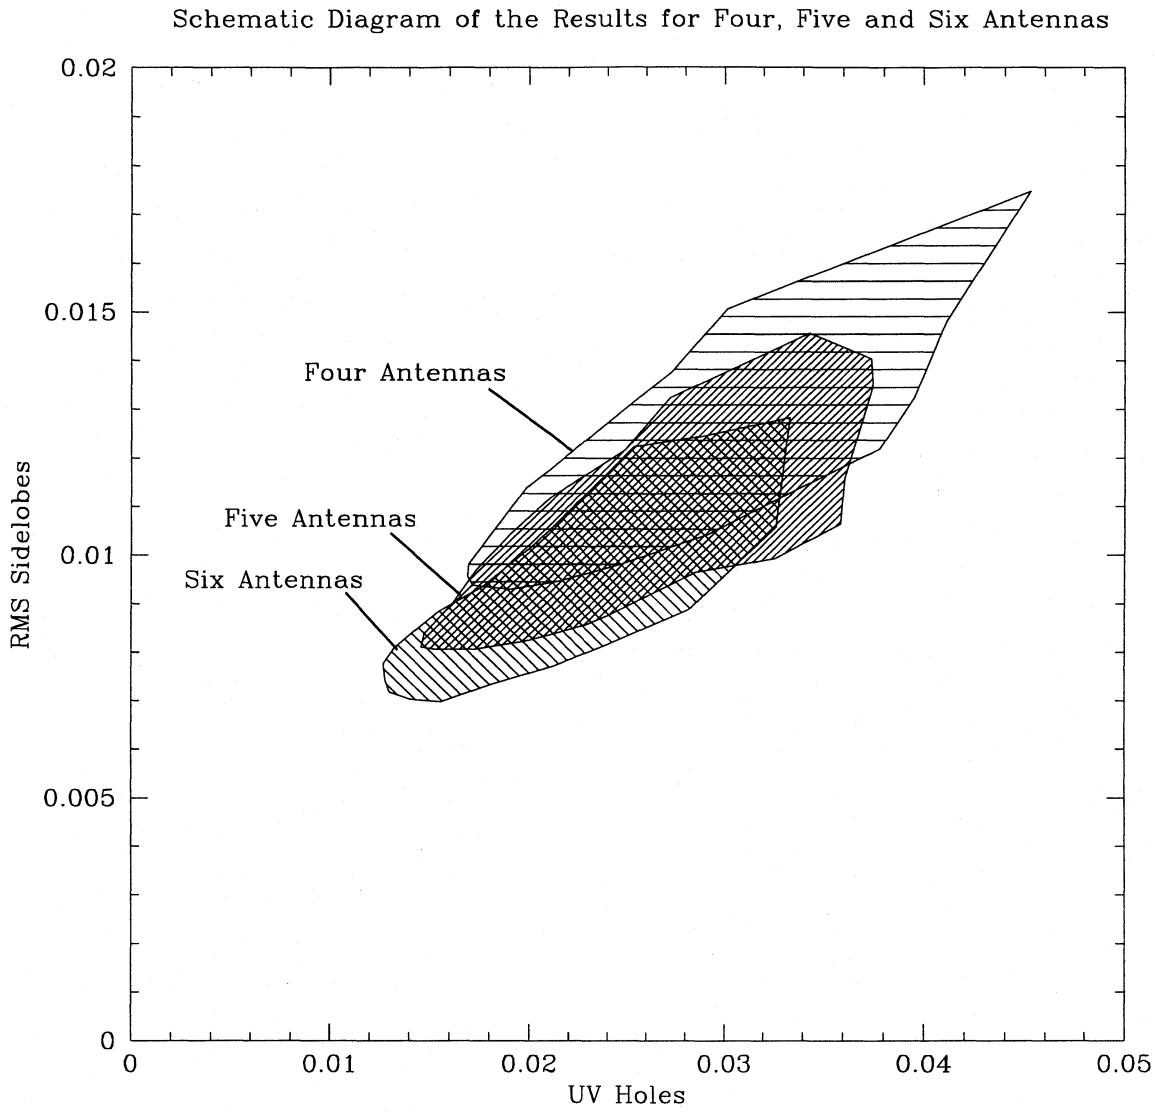

Figure 3: The range in quality of the Fourier coverage and point spread functions over all stages for four, five, and six upgrade antennas.

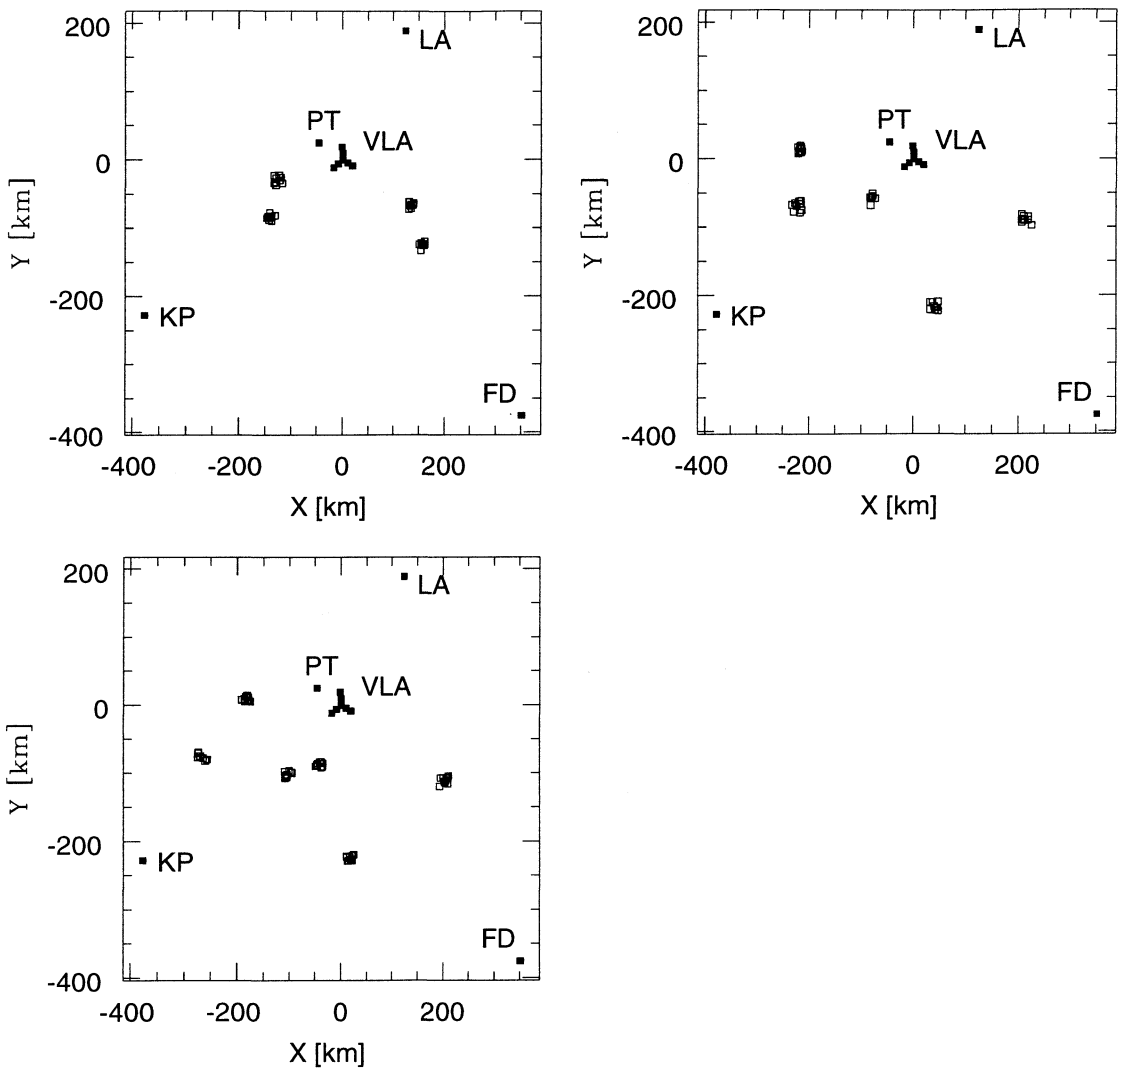

Figure 4: Antenna locations for the best 10 configurations with four, five, and six upgrade antenna after the 20 km wiggle stage. Filled boxes are existing VLBA and VLA stations while empty boxes are possible upgrade sites.

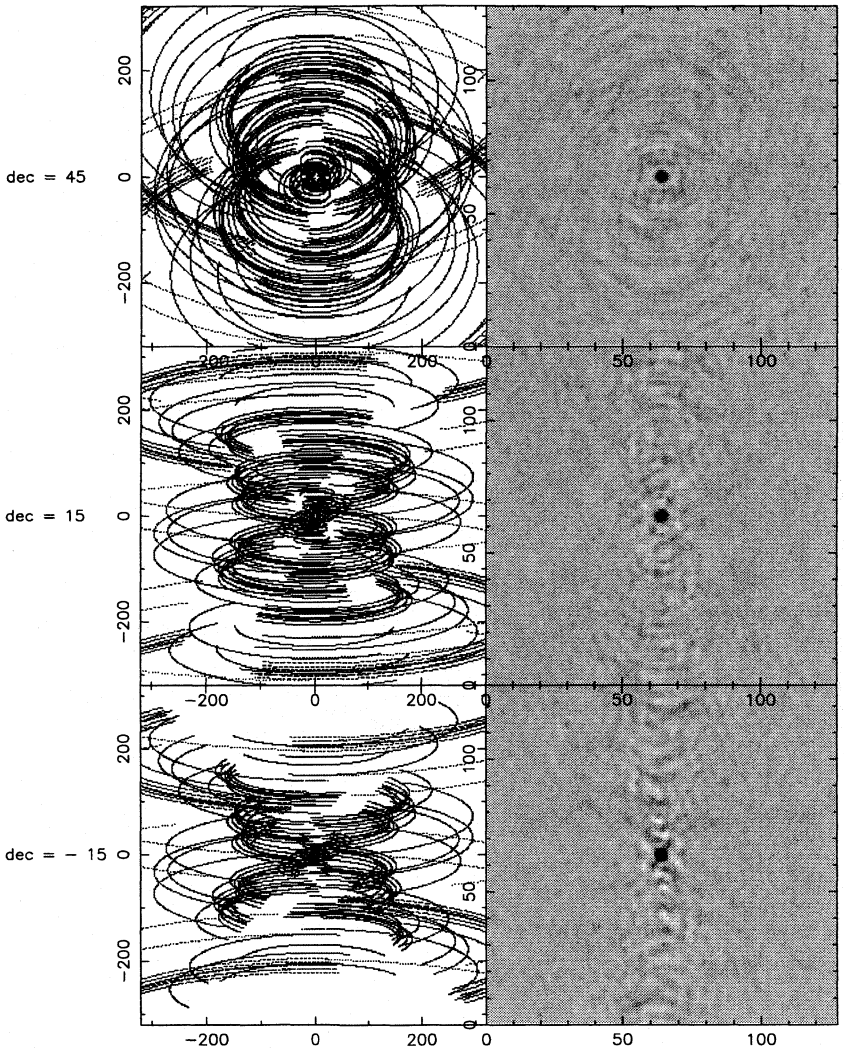

Figure 5: Fourier plane coverages over the inner 320 km and uniformly weighted point spread functions (greyscale range: -0.1 to 0.3) for the best four antenna configuration for declinations of 45, 15, and -15 degrees.

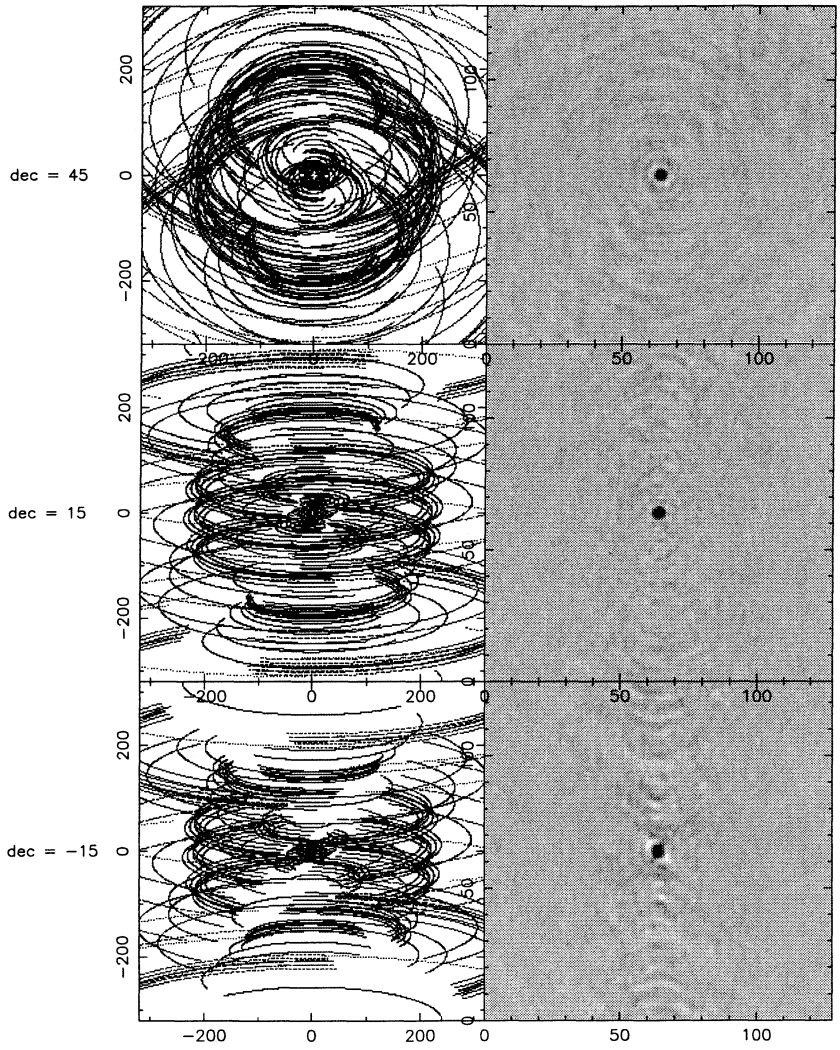

Figure 6: Fourier plane coverages over the inner 320 km and uniformly weighted point spread functions (greyscale range: -0.1 to 0.3) for the best five antenna configuration for declinations of 45, 15, and -15 degrees.

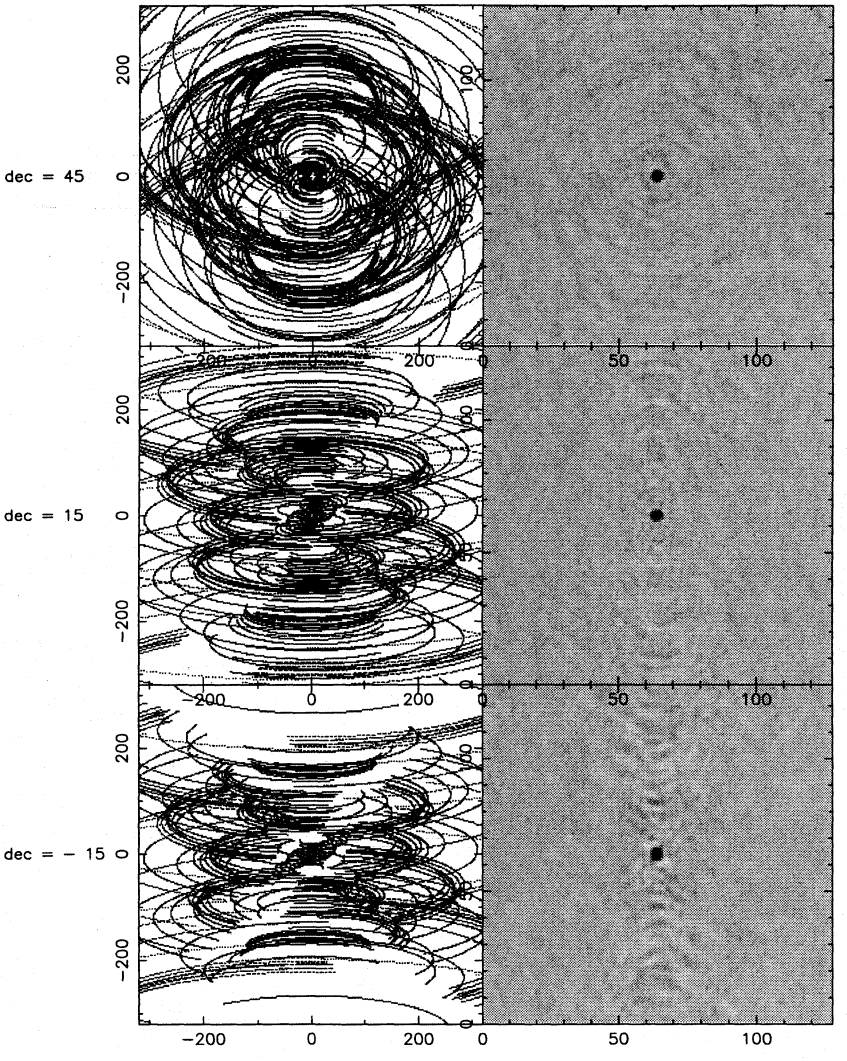

Figure 7: Fourier plane coverages over the inner 320 km and uniformly weighted point spread functions (greyscale range: -0.1 to 0.3) for the best six antenna configuration for declinations of 45, 15, and -15 degrees.

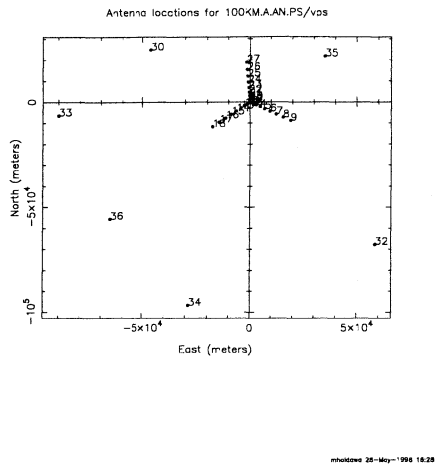

Figure 8: A 100 km antenna configuration with five upgrade antennas.

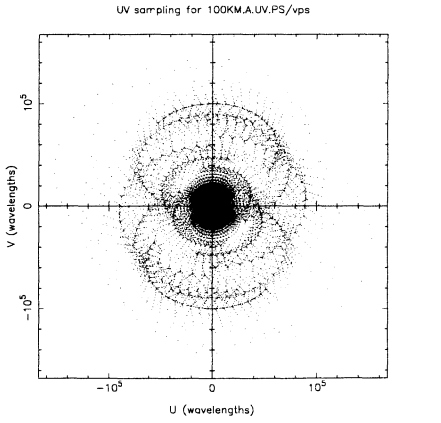

Figure 9: The 100 km antenna configuration results in excellent Fourier plane coverage and would produce excellent images with higher brightness sensitivity.

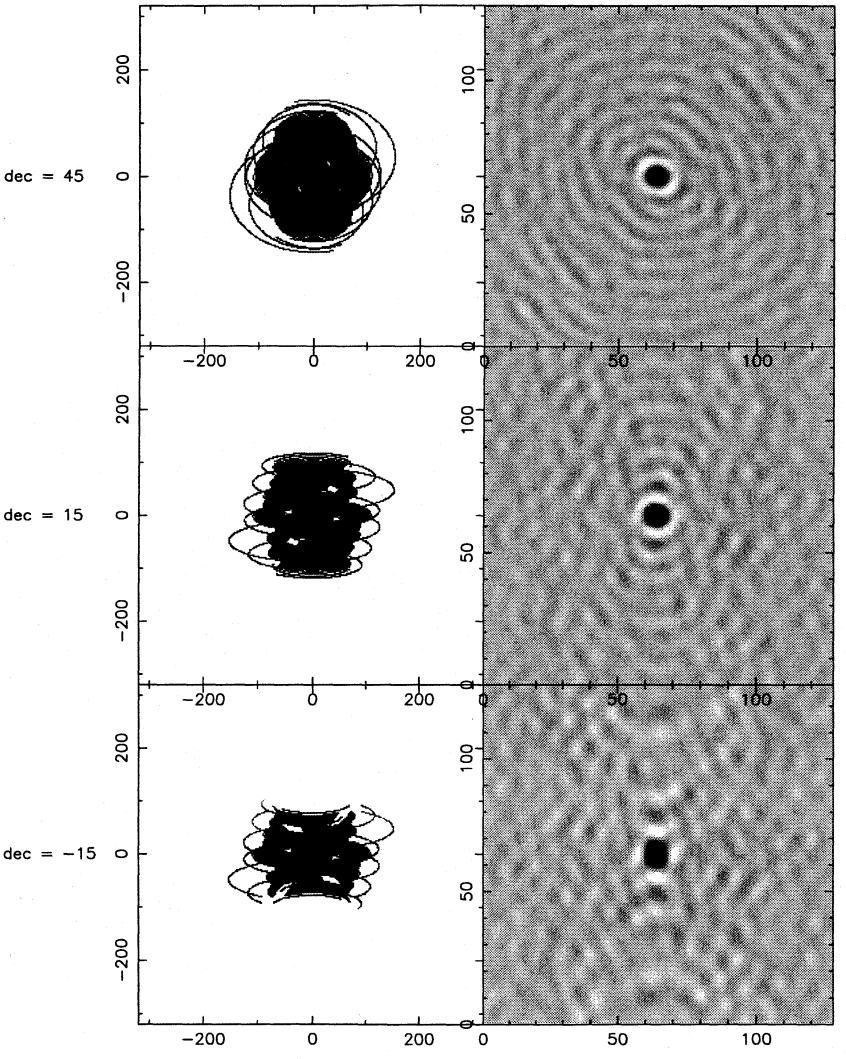

Figure 10: Fourier plane coverages over the inner 320 km and uniformly weighted point spread functions (greyscale range: -0.1 to 0.3) for the five antenna 100 km configuration for declinations of 45, 15, and -15 degrees.

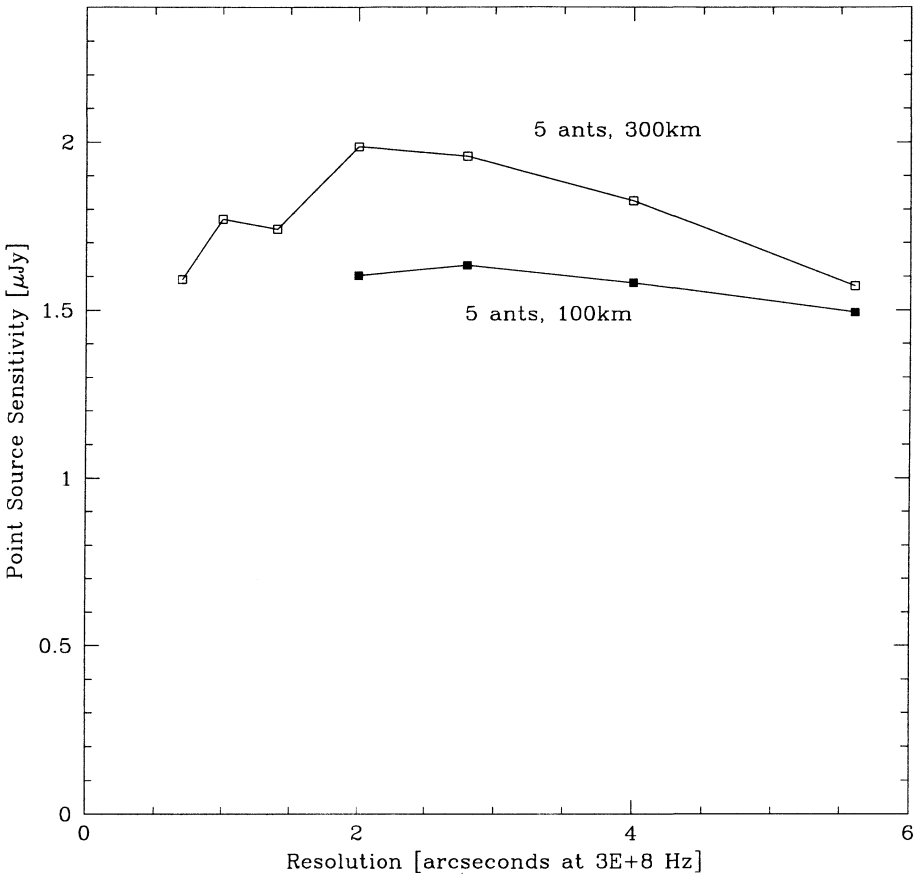

Figure 11: Point source noise as a function of resolution for robust weighted and tapered data from five upgrade antenna 300 km amd 100 km configurations. With robust weighting, the point source noise can actually decrease as we taper (throw away some long baseline data). The 300 km array's brightness sensitivity will be only moderately worse than a 100 km array.