MMA Memo 216
Self-Similar Spiral Geometries for the LSA/MMA

John Conway
Onsala Space Observatory
S-43992, Sweden
email: jconway@oso.chalmers.se

June 18, 1998

Abstract

In order to reduce the construction and operations costs of the LSA/MMA a large degree of sharing of pads and other infrastructure between different sized configurations is highly desirable. Motivated by this argument logarithmic spiral geometries for the LSA/MMA have been investigated. It is found that certain classes of such arrays with high pitch angle spiral arms, and some random noise in the positioning of antennas along each spiral arm, have uv coverage distributions and naturally weighted beams which are very close to Gaussian. It is argued that for the intermediate sized configurations between 0.2km and 3km such uv distributions and beams are close to optimum. The use of a spiral geometry allows the LSA/MMA to have continuously variable resolution, and allows for the possibility of dispensing with set configurations entirely. If however set configurations are used it is possible to use many configurations with small ratios in size between them (i.e. factors of 2); while only needing to reconfigure about one third of the total number of antennas between configurations. The flexibility of the design means that the operational mode can be decided after construction in response to scientific requirements.

1 Introduction

Published simulations of LSA/MMA array configurations to date have concentrated on either ringlike/closed-figure distributions or pseudo-random antenna distributions (see the summary of Helfer and Holdaway 1998, MMA Memo 198). While these configurations have been optimised for a given array size there has not yet been much consideration of how to share antenna pads or other infrastructure between all the different array sizes that are required (although see Webster, MMA Memo 214). A large degree of sharing of pads, cable trenches, and the rough bulldozed 'roads' along which antennas will be moved is obviously highly desirable in order to keep construction costs down. In addition a high degree of pad sharing between configurations would greatly reduce the amount of antenna moves that are required; reducing the transporter and manpower requirements and/or allowing more frequent moves with smaller ratios between array sizes.

In order to minimise pad and other infrastructure costs it seems worthwhile to try to investigate configurations based on some self-similar geometry. At the first LSA Array Configuration Working Group meeting in Paris at the end of April I suggested array configurations based on logarithmic spirals to achieve these aims. Here I present results of some preliminary simulations of this concept. I note that these simulations have benefited significantly from the comments and suggestions of other members of the LSA Array Configuration Working Group at the Paris meeting.

2 Target uv Coverages

Although the investigation of spiral geometries was initially motivated by the desire to share pads and infrastructure between configurations (see Section 1) it turns out (See Sect 3) that certain tightly wound spiral geometries naturally give excellent snapshot uv coverages.

At the Paris meeting there was a discussion of what the target uv coverage for the LSA/MMA should be, a subject that has also been discussed in several MMA memos. It is clear that the largest and smallest arrays have their own special requirements and constraints. For the smallest array, complete coverage of the aperture plane is required (sampling of all uv grid cells), and furthermore the array should also be optimised for mosaicing. Some form of very closely packed 'crystalline array' has therefore been suggested. The special constraints on this array include antenna shadowing and practical issues of access by the antenna transporter. For the largest array (10km-12km diameter) the site places strong constraints. Both this and the fact that this largest configuration will be often be used for scientific problems where we desire the highest possible resolution suggests a 'closed-figure' configuration, giving almost uniform coverage of the uv plane (see Holdaway,Foster,Morita 1996, MMA Memo 153). The large near-in sidelobes which result from the abrupt termination of the aperture coverage at its outer boundary can be dealt with by standard deconvolution techniques. Since for this array images will probably be sensitivity limited, absolutely optimum imaging performance is not required, and deconvolution algorithms will be entirely adequate to give sufficiently high dynamic range (and possibly superesolved) images which reach the thermal noise.

In contrast to the smallest and largest arrays, for the intermediate arrays which fill the 3km diameter plain, there is considerable choice available for the target uv coverage. Several members of the LSA Array Configuration Working Group (not myself initially) argued strongly for a Gaussian distribution of uv points for these intermediate arrays. Obviously such a distribution gives a natural weighted (maximum sensitivity) beam which is also approximately a Gaussian. A final image convolved by a Gaussian is obviously what astronomers seem to want since CLEAN images and even MEM images (which naturally give a variable resolution, best estimate of brightness per pixel) are usually in a final step convolved with a Gaussian 'restoring' beam. Such a process can be thought of as an attempt to give a final estimate of what the sky would look like if we first convolved it with a Gaussian before imaging it, since the Fourier transform of the restored image equals the visibilities that would be observed in that hypothetical case. Such final images with definite resolution are what are generally desired for both qualitative and quantitative analysis. Since Gaussian beams are what are wanted in the end it makes obvious sense to aim for a dirty beam which is as close to a Gaussian as possible, before applying deconvolution algorithms, and finally restoring. It goes without saying that such Gaussian dirty beams will also have minimum (i.e. zero) sidelobes. The one disadvantage of a Gaussian uv coverage is that for a given maximum baseline length the resolution is low because the coverage is so centrally condensed. However this is not a problem in a multi-configuration instrument, to get higher resolution one can just go to the next largest Gaussian configuration, augmented by a ring-like geometry for the largest array.

Alternative arguments for uniform coverage of the aperture plane (i.e. Keto 1997, ApJ 475,843) are based on the argument that given the requirements of maximum sensitivity and resolution (and given a limited circle within which antennas can be placed) it is best to sample as many uv cells as possible out to the maximum baseline, reducing the number that must be estimated via deconvolution algorithms. This criteria seems reasonable; however in order to get the desired final beam that is without near-in sidelobes (without resorting to strong tapering which loses sensitivity) deconvolution algorithms are required to extrapolate well beyond the edge of the sampled uv coverage. Hence the value of the visibility in many unsampled uv cells must still be estimated, but now they are in the outer edges and require extrapolation as compared to the interior where the data must instead be interpolated.

In the context of the above uv grid sampling arguments the alternative Gaussian distributed uv coverage can perhaps be thought of as one which has nearly uniform coverage in the centre plus outliers which can constrain the necessary extrapolation beyond this well sampled region, and therefore makes optimum use of a limited number of uv points.

As we discuss in the next section it turns out that certain classes of spiral arrays give snapshot uv coverages which are very close to Gaussian. In addition therefore to the infrastructure issue such spiral arrays seem to give snapshot coverages, which if the above arguments are accepted, are close to the optimum desired for the intermediate arrays.

3 Spiral Array Simulations - General

It is obviously an advantage in minimising pad resources etc in a multi-resolution design to use some form of self-similar pattern, however other patterns exist, including nested circles or nested triangles, why therefore chose logarithmic spirals? One clear advantage of such spirals (in which along each arm the radius increases by the same fraction for each unit change in azimuth) is that they are continuously self-similar, this means that if we move all the antennas out along the spiral arms by any amount, then except for a rotation, we get scaled versions of the array and uv coverage (see Sect 5). Nested circles or triangles with fixed ratios between the scales of the nested elements don't have this property, its is only possible to change the scale by discrete values to get self similar uv coverages. Another advantage of spirals is that they naturally avoid the tiered or 'wedding cake' uv coverage which comes from nested circles or triangles.

There are clearly a very large number of potential spiral patterns which can be investigated. The fundamental variables to consider include the number of arms of the spiral. Another parameter is the ratio of the radius out to the last filled pad on each arm to the radius out to the first filled pad (which we call RRAD in the simulations). Finally we have the angle in azimuth the arms turn through in going from the first filled pad to the last filled pad (DTHETA). Together RRAD and DTHETA determine the 'pitch angle' of the resulting spiral arms. On a perfectly regular logarithmic spiral the azimuth angles between adjacent pads on a spiral arm are constant and equal DTHETA/(NTEL-1), where NTEL is the number of antennas per arm. Adjacent pads also have a constant ratio of radius from the origin which equals $RRAD^{1/(NTEL-1)}$. In practice it is found desirable (see Sect 4) to constrain the antennas to lie on spiral arms but add some random noise to the azimuth values chosen for the pads; giving us yet more parameters to vary.

Because of the large number of parameters a complete investigation of all possible spiral arrays has not yet been carried out. Preliminary simulations of 1, 2 and 3 arm spirals however suggest that the 3 arm versions are the best so far tested; the resulting 6 fold symmetries in the uv plane give almost circular synthesised beams for zenith snapshot observations. Simulations of spirals with more than 3 arms have not yet been made, these should also be investigated. Obtained uv coverages appear to be best for fairly tightly wound spirals. For small values of DTHETA the three arm pattern reduces to the dreaded VLA 'Y' shape with all its undesirable properties for snapshots. Good uv coverages are only obtained for large values of RRAD when DTHETA exceeds about one turn. The resulting array configurations come close then to the 'donut' arrays discussed by Kogan 1998, (MMA Memo 212), although more centrally condensed.

Figs 1 2d illustrate the uv coverage and natural weighted beam obtained for one of the best spiral arrays tested so far, the snapshot uv coverage of which is very close to Gaussian. Given the limited number of spirals so far tested there is no guarantee that this array is close to optimum; however in the rest of this memo we will concentrate on this specific design in order to illustrate some of the properties of spiral arrays in general.

4 Spiral Arrays - An Example Array

The array configuration shown in Fig 1a, is for a homogenous 63 element LSA/MMA which has three arms (hence 21 antennas per arm) with DTHETA=1.5 turns and RRAD of 8 (a similar simulation for a 36 element MMA is shown in Fig 3). It follows that for each half a turn along each spiral arm the radius from the centre increases by a factor of two. Random perturbations on antenna azimuths of r.m.s. 20% of the mean azimuth angle between antennas have been added to break the regularity. The importance of this for removing spiral sidelobes is illustrated in Fig 2.

The zenith snapshot uv coverage which results from the above array is shown in Fig 1b, while Fig 1c shows the density of uv points per unit area within annuli of different radius in the uv plane. The solid line in the same figure shows a best fitting Gaussian against uv distance. Except at very small radii the fall-off of uv point density with uv distance is remarkably close to Gaussian, even showing the required long Gaussian tail beyond uv distance 1500m (see Fig 1c). Part of the reason for this excellent uv coverage may be that that the outer perimeter of the antenna configuration (see Fig 1a) is close to being a Reuleaux Triangle with scaled nested Reuleaux Triangles within this outer perimeter.

The departures from a Gaussian distribution at small uv radius shown in Fig 1c are influenced by the exact choice of the random noise to be added to the antenna azimuths (i.e the departures are different for different runs of the Monte-Carlo simulation). The ratio between the longest and shortest baselines (which in this simulation is 75) also depends critically on the exact set of perturbations in pad positions used. Rather than simply drawing these perturbations from a completely random Gaussian distribution an optimised pseudo-random distribution of perturbation factors could probably be chosen to optimise the short baseline uv coverage, and optimise the beam shape (Kogan 1997, MMA Memo 171). In addition the short baseline density could be controlled somewhat by giving these short baselines non-unity weights; since relatively few points are effected the loss in sensitivity would not be significant (about 5% for the array shown in Figure 1).

The dirty beam obtained with pure natural weighting, shown as a slice in Fig1d and as a grey scale in Fig2d, is of good quality. As expected from the obtained uv coverage the main beam is close to Gaussian in shape. The addition of perturbations on the azimuth have completely removed any trace of the spiral sidelobes which are seen in simulations of completely regular spiral patterns (see Figs 2a,b). The peak amplitude sidelobe is only 2%; this could be reduced by optimum choice of azimuth perturbation factors and uv weighting to make the radial uv density even closer to Gaussian.

Simulations with fewer antennas, i.e. for a 36 element MMA only array, are shown in Fig 3. In this case the pure natural weighted beam peak sidelobe is only a factor of two larger than in the 63 element case. It appears that the positive properties of spiral arrays are general and do not depend on having a very large number of elements.

5 Scaled Array Properties

The scale-free properties of spiral arrays are illustrated in Figs 4 5. The dots here show the pad positions (of which there are 42 per arm in this example) while the circles indicate the position of the 21 antennas on each arm. Fig 4a shows the starting array which fills the 3km diameter plain, in which on each spiral arm pads 22 to 42 are occupied. We can reduce the array size by moving antennas inwards from the pads at the ends of the spiral arms to unpopulated pads toward the centre. So the first move would be to take the antenna on pad 42 of each arm and move it to pad 21 (see Fig 4b). The resulting antenna distributions after moving 3 and 7 antennas per arm are also shown in Fig 4c,d respectively. It is obvious that as the antennas are moved inward the antenna distribution rotates and shrinks giving a 'zoom lens' effect. The final plot in Fig 4d corresponds to having moved one third of all the antennas on each arm. Given the properties of this spiral in which the radius increases by a factor 8 between the lowest and highest populated pads while the angle changes by 1.5 turns, moving one third of the antennas corresponds to a change in scale of the whole array by a factor of 2 and rotation of half a turn. The resulting beam will therefore be the same as the original except reduced in size by a factor of two.

There are a number of possible ways to utilise the scale free property of spiral arrays in operating the LSA/MMA.

Continuously Moving Array. We need not have any set intermediate configurations at all, we could always be moving at least one antenna (Viallefond 1998, LSA/MMA Feasibility Study, 1998). Since the resolution of the array would be self-similar but continuously variable, observations could be scheduled when the resolution at the desired frequency was exactly as required by the astronomer. Using this facility it would be possible to make images in different molecules or transitions at exactly the same resolution which is obviously very important for line ratio studies.

One problem that was pointed out at the Paris meeting with this scheme would be the calibration of antenna positions. However it was realised that this need not be to much of problem when we consider that the exact antenna positions need only be known when the data is *analysed* not when *observed*. We could still therefore make do with calibration runs every few days.

Factor of Two Scaling. If we did still want set array configurations, because of the high degree of pad sharing we could separate the sizes of adjacent arrays by much smaller ratios than the factors of 3 or 4 so far discussed in the MMA design documents. A factor of two between array sizes seems like a natural ratio to choose, and is probably the smallest ratio over which there is likely to a significant change in imaging properties. Fig 5 shows how we could cover the range of scales between the smallest close-packed crystalline array and a spiral array which fills the 3km plain, with only 4 configurations each differing in size by a factor of 2 from the one before. For a 63 element LSA/MMA array with 21 antennas and 42 pads per arm, the largest 3km array spiral would populate pads 22 to 42, the next smallest pads 15 to 35, the next pads 8 to 28 and finally the most compact spiral would use pads 1 to 21 (see Figs 5a,b,c,d). In moving from one of these set arrays to another we need only move one third of the antennas. While the move was occurring we could still use the array since it would still have approximately 4/9 of its baselines available.

Flexibility Perhaps the most important point about scale-free spiral arrays is that they are very flexible, hence *we don't need to know how we will operate the array in detail before the array is constructed.* We could decide once it is built between using a continuously moving array, factor of 2 set arrays or some other option. Since we don't know what the scientific requirements will be in 10 - 20 years time this inherent flexibility is one of the most important arguments for this type of design.

6 Spiral Arrays within the overall LSA/MMA Design

As we discussed in Section 1, the most compact array and the 10km-12km array have their own special design requirements, and we only propose spiral arrays for the intermediate scale configurations between these extremes.

If we use the example spiral array discussed in Sections 4 and 5 then by using only 2 pads per antenna we can cover continuously the range of baselines between the most compact 'crystalline' array and a spiral which just fits within the 3km plain. If we operated set configurations we could then span the above range with only four configurations (see Fig 5). For the most compact spiral array the longest baseline would be about 300m, with most antennas being within a dense pack of around 100m in diameter. This smallest spiral therefore naturally connects with whatever design is chosen for the smallest 'crystalline' or 'parking lot' array and could share some pads.

The total pad requirements in this scheme would be 1 pad per antenna for the crystalline array (possibly this could be reduced by pad sharing with the spiral), 2 pads per antenna in the spiral and 1 pad per antenna on the 10km/12km diameter ring, giving a total requirement of 4 pads per antenna. One modification that should be considered is to have an additional ring-like, high resolution 3km array. The 3km spiral is quite condensed, without something intermediate between this array and the highest resolution 12km array there will be a serious gap in resolution. A high resolution 3km array could be achieved by bending the outer spiral arms in Fig 1a into a ring or Reuleaux triangle, and placing most of the antennas on this perimeter. Other alternatives are donut shaped arrays (Kogan 1998, MMA memo 212); perhaps half of the pad positions could be shared with the 3km spiral giving a total requirement of 4.5 pads per antenna for the whole LSA/MMA.

7 Infrastructure

One advantage of a self-similar spiral pattern compared to completely randomised arrays is the reduced construction costs due to minimisation of conduits, cables and roads.

Conduits for optic fibres, power cables and the 'road' along which the transporters would move antennas would follow the spiral pattern; and in addition there would would be roads and conduits running North-South and East-West along the x and y axes of Fig 1a. Since the antennas share the use of conduits and roads this self-similar design minimises the amount of construction required. In addition the optic fibres from each antenna to the correlator located at the spiral centre, will first follow along the spiral path and then one of the N-S or E-W access trenches, hence the mean length of fibre required to reach any given antenna is only about 40% larger than the absolute minimum set by the line-of-sight distance from antenna to correlator building at the spiral centre.

8 Conclusions and Future Work

In summary 'grand design' spiral arrays seem to offer some advantages in the aspects of uv coverage, construction and overall array operations and therefore seem well worth investigating in addition to the ring-like or randomised arrays so far considered.

A number of issues need to be addressed in future work

- Carry out a full exploration of all the parameter space of possible spirals to find the optimum.

- Include terrain restrictions, which will be most important for the 3km spiral. A quick look at the available terrain (see Fig 1 in Kogan (1998), MMA Memo 202), suggests that a rough spiral pattern can be fitted in. Fortunately some non-regularity in antenna positioning is required anyway to obtain the best beam.

- Optimisation of antenna positions along the arms to obtain near Gaussian uv sampling and the lowest possible sidelobes in the synthesised beams.

- Consider more complex schemes for populating the pad positions, the present simulations assume that all adjacent pads are occupied, i.e. in the largest configuration the 21 pads between pad number 22 an 42 are all occupied and none others; more complex schemes can of course be utilised.

- Include a slight 10% or 20% stretching in the North-South direction to give an almost a circular snapshot beam at a wide range of declinations (Foster 1994, MMA Memo 119).

- Consider hybrid arrays for observations of sources in the far North or South. This can probably be achieved when expanding the array by moving antennas from the inner to outer parts of the spiral first on whichever spiral arm is northernmost at the time.

- Consider non-snapshot i.e. long track observations. The synthesised beam in this case will be a linear combination of the snapshot beams, which are in turn rotated and stretched versions of the zenith snapshot beam. Since each of these snapshot beams will be approximately Gaussian we can expect very low sidelobes in the synthesised beam.

- Consider in detail, the interface between the smallest spiral (Fig 5d) and the close packed 'crystalline' arrays. Furthermore consider the interface between the largest spiral and the 3km and 10km high resolution 'ring' arrays. We discussed a high resolution 3km array in Sect 6. In addition the 10km ring could be formed by extending two of the three arms of the 3km spiral to wrap around the Chascon mountain (see Fig 3.1.2. in the LSA/MMA Feasibility study 1998) forming a closed loop.

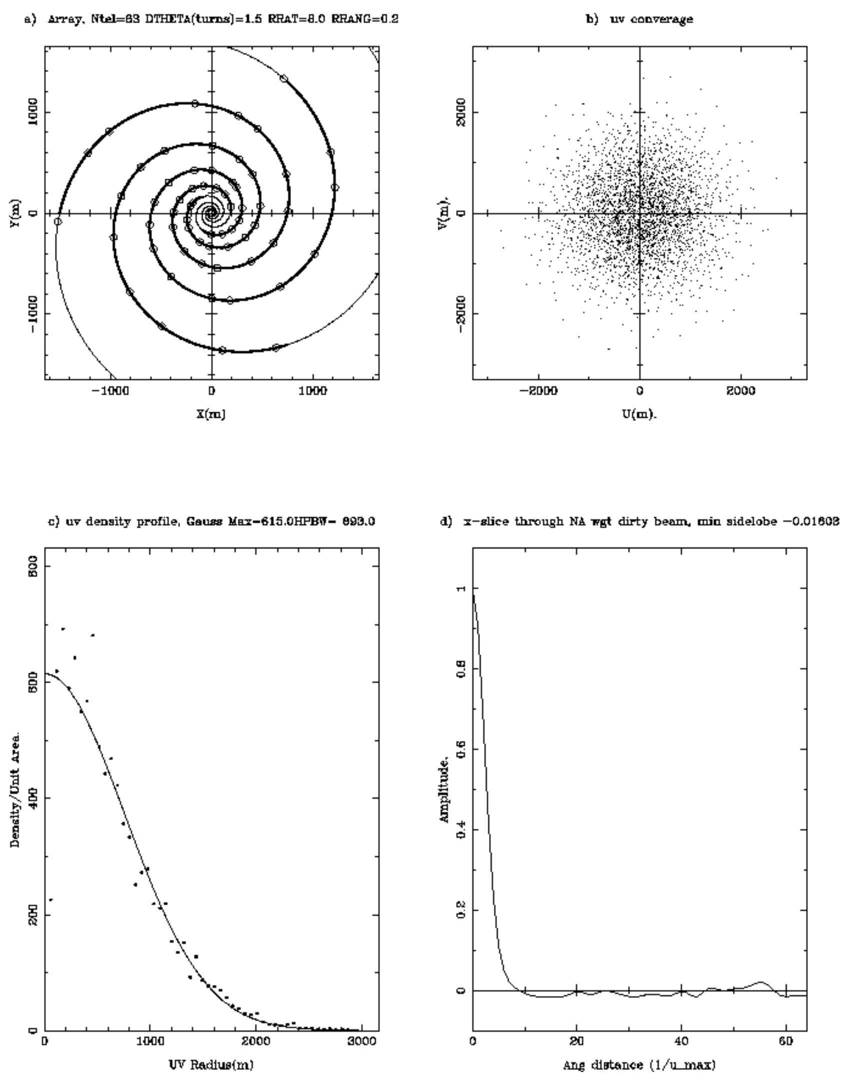

Figure 1: Simulation of a homogenous 63 element LSA/MMA spiral array with DTHETA=1.5 turns per arm over a ratio of radius RRAD = 8 between the first antenna and last antenna on each arm. a) Antenna layout. Nb. Circles indicate antenna positions, spiral lines indicate the paths taken by the 'road' and conduits for optical fibres, power cables etc. b) uv coverage. c) UV point density versus uv radius. d) Slice through the natural weight dirty beam. This array will decrease by a factor of two in size if one third of the antennas are moved from the ends of each spiral arm to the inner parts of that arm (see Figs 4 5).

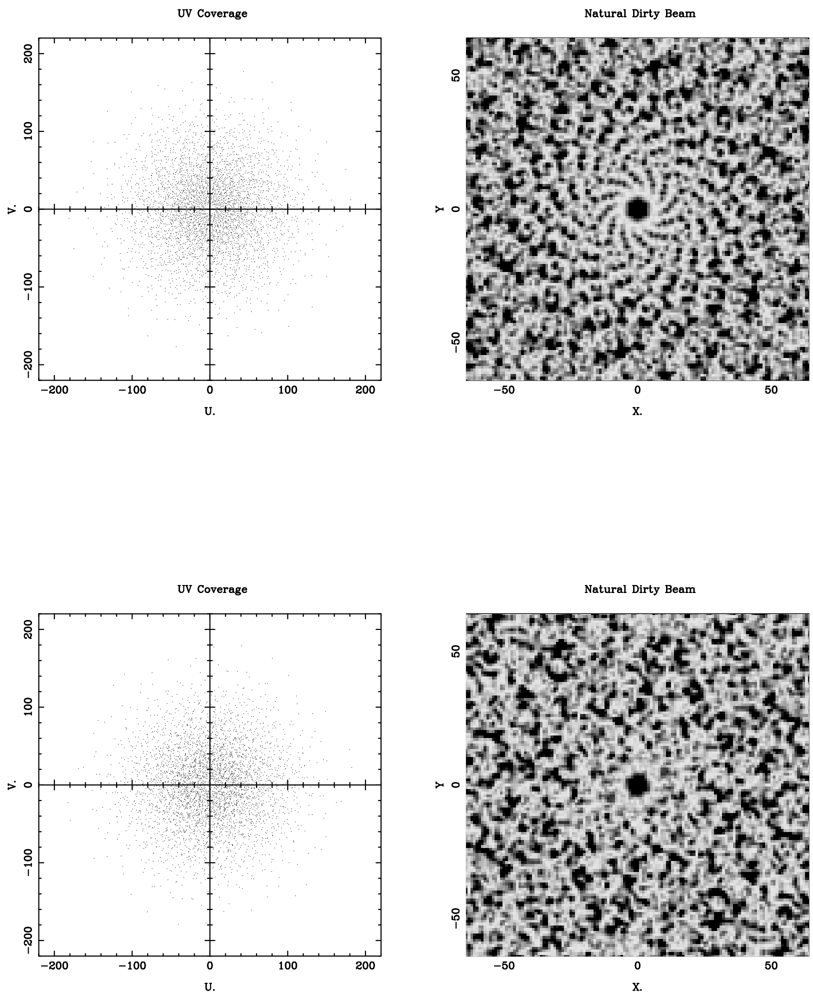

Figure 2: (a, Top Left) uv coverage for a perfectly regular spiral array with DTHETA=1.5, RRAD=8 and 21 antennas per arm. (b, Top Right) The resulting dirty beam pattern, white is -0.02, black is 0.02. Note the spiral sidelobes. (c, Bottom Left) uv coverage after including random perturbations in the pad azimuths of r.m.s. 20%, i.e. for the antenna configuration shown in Fig 1a. (d, Bottom Right) the resulting natural beam pattern, white is -0.02, black is 0.02. The spiral sidelobes present in the case of a perfectly regular spiral array have been removed.

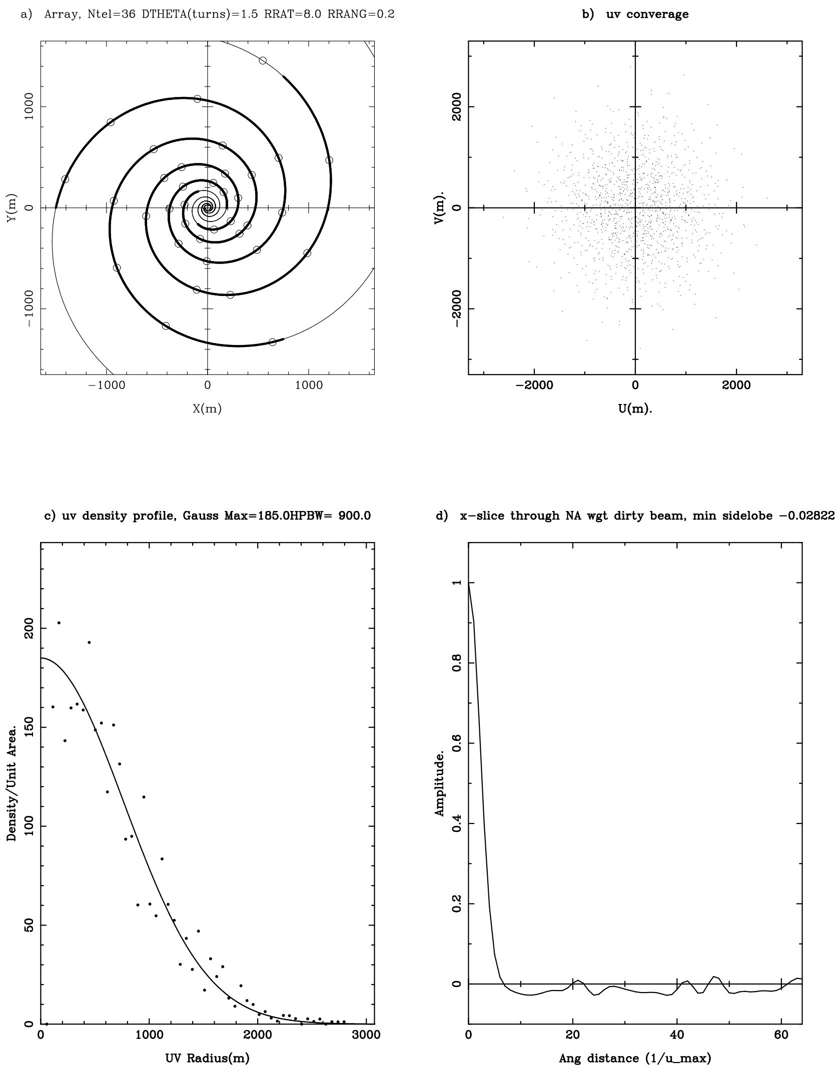

Figure 3: Simulation of a 36 element MMA-only spiral array with DTHETA=1.5 turns per arm over a ratio of radius RRAD = 8 between the first antenna and last antenna. a) Antenna layout. Circles indicate antenna positions, spiral lines indicate the paths taken by the 'road' and conduits for optical fibres, power cables etc. b) uv coverage. c) UV point density versus uv distance. d) Slice through the natural weighted dirty beam.

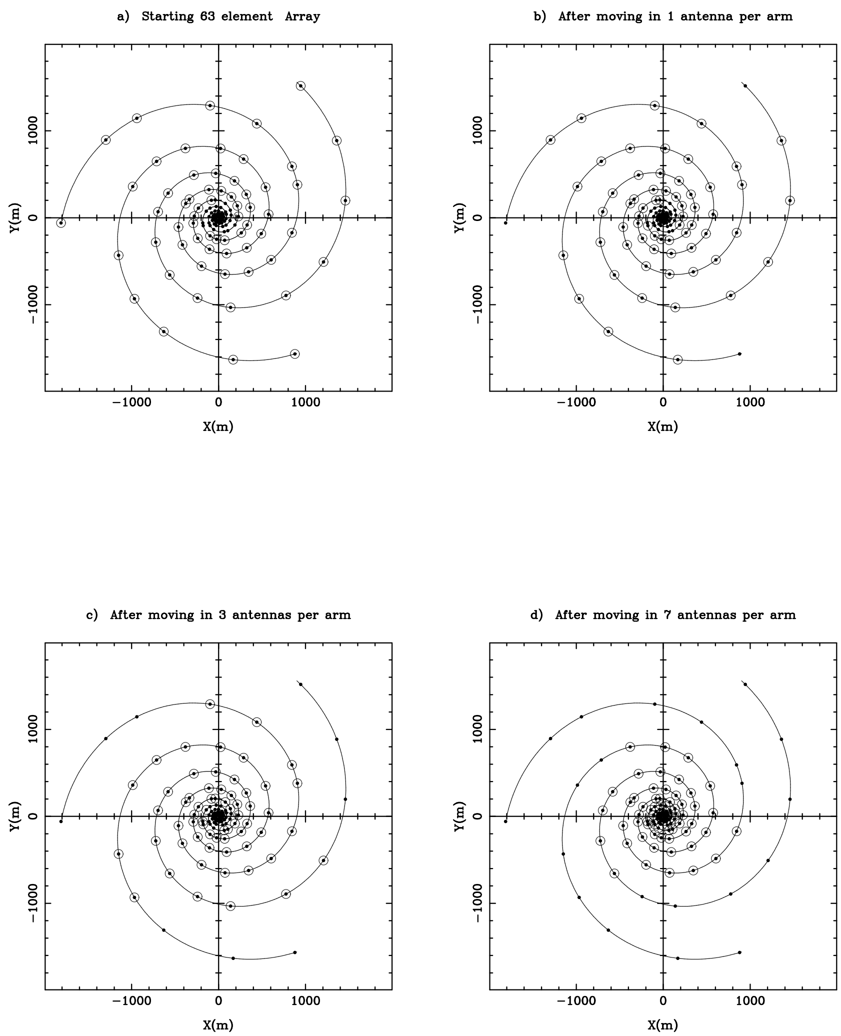

Figure 4: Illustration of the self-similarity of the array as antennas are moved inward from the ends of the spiral arms. Small dots indicate pad positions and circles the antennas. This 'zoom lens' effect allows the possibility of dispensing with fixed array configurations entirely and operating in a continuously moving mode (see text). In any case it allows great operational flexibility. As can be seen by from Fig 4d, after moving one third of the total number of antennas the antenna geometry is the same as the original except scaled by a factor of 2 in size and rotated by 180 degrees.

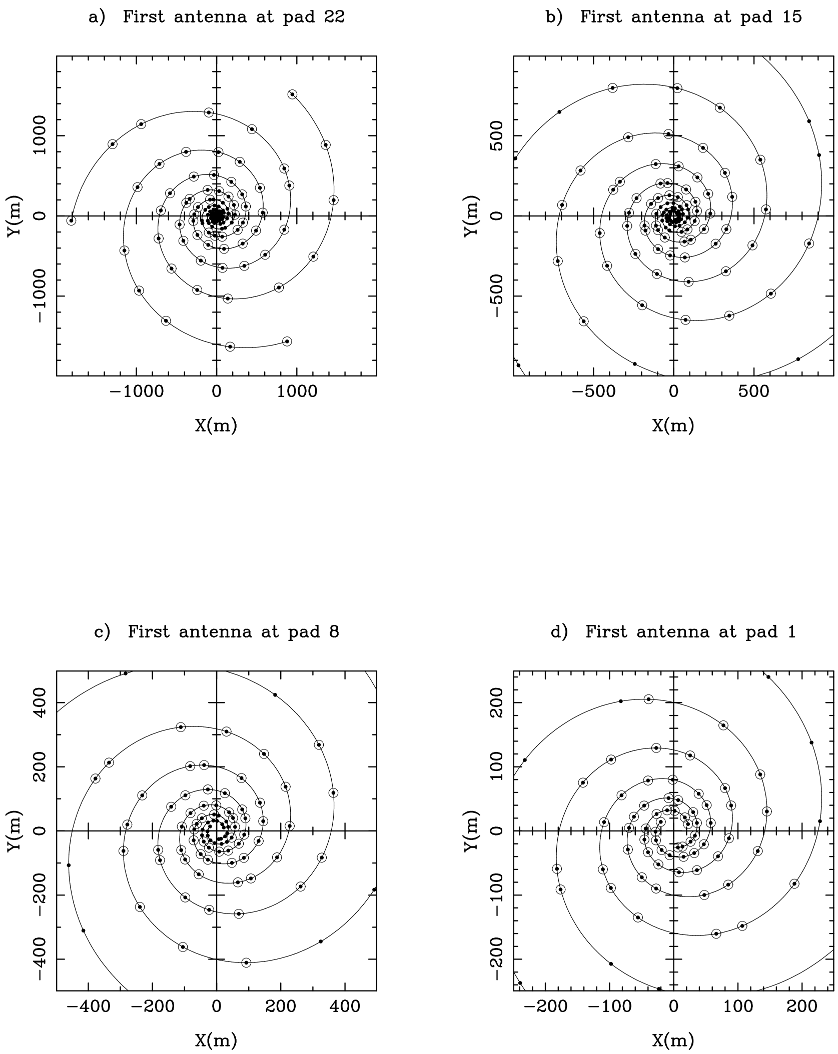

Figure 5: Illustration of how four configurations could be used to span the range of baseline distance between a spiral which fills the 3km plain and ultra-compact 'crystalline' array, in powers of 2, using a geometry with 21 antennas and 42 pads per arm. The scale of each panel is different from the previous one by a factor of 2. In moving from one configuration to the next smallest only 1/3 of the antennas need be moved. Dots are pad positions and circles denote antennas. Note that Fig 5b corresponds to Fig 4d plotted with a different scale. The smallest spiral array shown in Fig 5d might be modified to smoothly connect with the most compact 'crystalline' array.