

# Computational geometry and the design of aperture-synthesis telescopes

## I. A measure of the quality of the $uv$ -plane coverage

A. S. Webster

Institute for Astronomy, School of Physics, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh EH9 3HJ, UK  
e-mail: asw@roe.ac.uk

Received 13 November 2003 / Accepted 21 July 2004

**Abstract.** When designing an aperture-synthesis telescope for use in astronomy it is essential to find a layout of the antennas on the site that provides good coverage of the  $uv$ -plane. Large holes in the distribution of baseline points in that plane are particularly undesirable, and so an algorithm was devised for locating all the holes in any given design and assessing whether they are unacceptably large. It finds, in each region of the  $uv$ -plane, the largest circle that does not contain any of the baseline points and then takes the properties of that empty circle, such as the radius and the position of the centre, as measures of the properties of the local hole. The algorithm is based on the Delaunay triangulation of the set of baseline points and makes use of the result that the circumscribed circle of any Delaunay triangle is always empty. The largest empty circles in the  $uv$ -plane are readily found in this way, and the algorithm selects from these circles a non-overlapping set that is a good match to the intuitive concept of the largest holes in the distribution of baseline points. Modified algorithms are also presented that rank each circle not by its absolute radius but by the radius normalized to the mean radius of the circles in the same neighbourhood of the  $uv$ -plane; these versions make proper allowance for the large-scale variation of the density of baselines caused by the taper and are therefore of greater use to the designers of a telescope. The algorithms are efficient, requiring  $O(N_B \log N_B)$  operations where  $N_B$  is the number of baselines, and may be used in several ways. One is to take the radius of the largest empty circle as a figure of demerit so that, of a set of trial configurations, the one with the smallest value may be regarded as having the best  $uv$ -plane coverage. Another is to take the baselines on the circumference of the largest circle to indicate which antennas are the most suitable for repositioning in order to improve the quality of the  $uv$ -plane coverage of a trial configuration.

**Key words.** telescopes – instrumentation: interferometers

### 1. Introduction

Aperture-synthesis interferometry is an indirect technique. The desired result of an observation with a telescope of this type is an image of a particular field of the sky but the output of the interferometer is a sampling of the correlation function  $V(u, v)$  of the brightness of that field (for a full definition of this function see e.g., Perley et al. 1989; Taylor et al. 1999; Thompson et al. 2001). Although the correlation function is defined over a two-dimensional space conventionally termed the  $uv$ -plane (defined in Sect. 2 below), it is not an image in any sense and has to be converted into one before it is of much use. The underlying relationship between the correlation function and the image is the two-dimensional Fourier transform and is therefore well understood in principle, but the act of conversion is not straightforward in practice because the knowledge of the correlation function provided by the interferometer is always incomplete. The 2-dimensional Fourier transform relates two functions whose support is continuous and doubly infinite and so its use requires

a knowledge of the correlation function at every point in the  $uv$ -plane over the continuous range  $-\infty \leq u \leq \infty$  and  $-\infty \leq v \leq \infty$  (Bracewell 1965), whereas at any instant the telescope only samples the function at a finite set of points lying within a restricted region of that plane. This limitation on the sampling has two causes, the first being the finiteness of the extent of the telescope, which limits the accessible part of the  $uv$ -plane to a finite region centred on the origin, while the second is the finiteness of the number of antennas, which limits the sampling of that region to a finite number of points. The limitation means not only that the Fourier transform cannot be applied directly in order to derive the image but also that the image itself is not completely determined. Various techniques have been devised to circumvent these difficulties: the restriction of the sampling to a region near the origin is generally dealt with by attempting only to produce an image with a finite resolution, corresponding typically to the brightness distribution across the field of view after convolution with a broadening function of finite width, while the restriction of the sampling to a finite number

of points may be eased by methods that amount to interpolating the function to neighbouring points (Perley et al. 1989; Taylor et al. 1999; Thompson et al. 2001).

The incompleteness of the sampling has consequences that reach well beyond the choice of data-reduction algorithms, however, and the concern of this paper is with its influence on the design of telescopes, and in particular on the siting of the antennas. The incompleteness is relevant to this subject because, although the data-reduction algorithms may determine how the correlation measurements are interpolated in the  $uv$ -plane between the sampled points, the positions of the points are determined by where the antennas were sited in the first place. It follows that the designers of an aperture-synthesis telescope must consider whether any proposed siting of the antennas is ideal to support the best possible reduction of the astronomical data or whether a better siting can be found.

The distribution of the sampled points in the  $uv$ -plane may be considered separately on large and small scales. On large scales the distribution is commonly referred to as the taper and is determined by the large-scale distribution of antennas over the site. On small scales the topic of interest is the detailed distribution of the sampled points and the spaces between them. The present study is entirely concerned with the small-scale structure, and especially with the important matter of ensuring that there are no large holes between the sampled points. Such holes are undesirable because within them the correlation function is poorly sampled and, although the function could still be estimated there by interpolation from the sampled points surrounding the hole, interpolation is less reliable and less accurate between points that are widely spaced than between those that are close together (e.g., Walker 1984; Hjellming 1989).

The specific question addressed here, then, is how best to go about assessing the sizes of the holes in the distribution of baseline points in the  $uv$ -plane when designing a new telescope. Existing algorithms for assessing the quality of the coverage of the  $uv$ -plane are concerned either with the large-scale distribution or with a blend of the large- and small-scale distributions. In a typical example the  $uv$ -plane is divided up into square bins and the number of nodes falling within each bin is counted (Walker 1984); these counts are then used to form a statistical criterion of uniformity, but the large and small scales are thoroughly mixed together in the process and the method does not involve the identification of individual holes.

A technique that has been widely employed in the past to assess holes and other features of the  $uv$ -plane distribution is visual inspection, making use of the well-known ability of the human brain to recognize patterns (Walker 1984), but it is not really suitable for modern interferometers. The issue of how best to assess the holes in the coverage of the  $uv$ -plane arose in a particularly acute form while carrying out design studies for the Atacama Large Millimetre Array (Shaver 1996; Wootten 2001; ESO 2003; NRAO 2003), which will consist of 64 antennas of 12 m diameter that may be sited on any of about 250 antenna stations in order to provide many useful configurations. It is planned that adjacent configurations should differ in the siting of only a few antennas in order that the width of the synthesized beam may be finely controlled, and this results in a large number of configurations. These configurations must,

moreover, be designed simultaneously because they have antenna stations in common. The large number of configurations, the need to design them simultaneously and the large number of baseline points for each configuration (4032) make designing the telescope a major problem, for which assessing the quality of the  $uv$ -plane coverage by visual inspection is too subjective to be reliable and too laborious to be feasible. A quantitative, rational method of finding and assessing the holes in the coverage was needed, in which the process of assessment could be left entirely to the computer, but no such method could be found in the literature.

It was therefore decided to attempt to devise an original method, the main problem being to find an objective, geometrical construct that is a good match to the subjective idea of a hole in a distribution of points. The construct could then be taken as a proxy for the hole, with the properties of any particular instance, such as the position and size, being taken as representing the properties of the corresponding hole. Given that the construct required is geometrical in nature, it seemed logical to look to the field of computational geometry for ideas (Preparata & Shamos 1985; O'Rourke 1998; Okabe et al. 2000); this, as the name implies, is concerned with the application of computers to all aspects of geometry and although many of its concepts do not appear to be relevant to the present problem, those relating to sets of points distributed in the plane certainly do. They include constructs such as the Voronoi tessellation and the Delaunay triangulation, which are reviewed in Sect. 2 in order to find out whether any of them, or any construct derived from them, matches the intuitive concept of a hole in the distribution of points. One promising candidate is found and developed as the basis of an empirical technique for assessing the quality of any  $uv$ -plane distribution.

## 2. Computational geometry

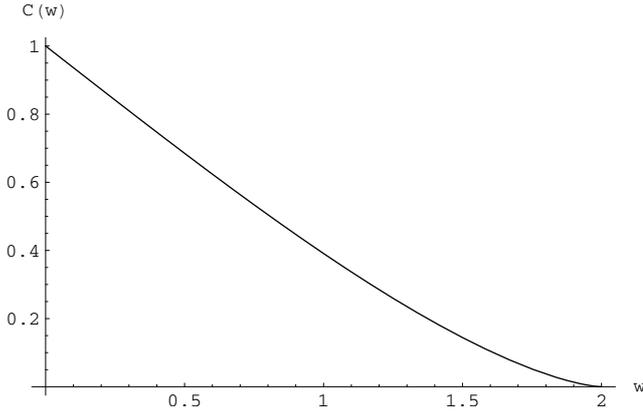
In any given configuration, the antennas are taken to be sited in a horizontal plane in which the Cartesian coordinates of the  $i$ th antenna are  $(x_i, y_i)$ . For each pair of antennas  $i \neq j$  there is a point in the  $uv$ -plane with coordinates

$$\begin{aligned} u_{ij} &= x_i - x_j \\ v_{ij} &= y_i - y_j \end{aligned} \quad (1)$$

and an inverse point with  $i$  and  $j$  reversed. These points are termed here baseline nodes, in order to distinguish them from general points lying anywhere in the  $uv$ -plane.

In order to define a problem that is well posed, the distribution of nodes is studied here in the simple form given in Eq. (1) with all possible complications set aside. Thus the effects of projection that would be introduced when studying a source in an arbitrary direction in the sky are ignored, which restricts the development to observations made at the zenith, while the baseline nodes are taken to be points rather than elongated tracks, which restricts the development to observations made in snapshot mode in which the effects of the rotation of the Earth are negligible. The consequences of making these assumptions differently are discussed in Sect. 3.3.

In order to illustrate the problem with an example, two specific instances of the same general design are studied.



**Fig. 1.** The autocorrelation function of a uniform disc of unit radius, from Eq. (2).

The design is based on a distribution of antennas on the site that is uniform within a disc of unit radius and zero outside. The electromagnetic properties of this case are well understood: the expected  $uv$ -plane distribution corresponds to the autocorrelation function of the uniform disc and is

$$C(w) = \frac{2}{\pi} \left[ \cos^{-1} \left( \frac{w}{2} \right) - \frac{w \sqrt{4 - w^2}}{4} \right] \text{ for } 0 \leq w \leq 2 \quad (2)$$

where  $w = \sqrt{u^2 + v^2}$  is the radial coordinate in the  $uv$ -plane (see Fig. 1), and the synthesized beam corresponds to the Airy pattern.

It is straightforward to use a computer to generate particular instances of arrays based on this general design by the use of a random-number subroutine. The position of each antenna in the  $xy$ -plane may be derived from its polar coordinates  $(r, \theta)$

$$\left. \begin{aligned} r_i &= U(0, 1)^{\frac{1}{2}} \\ \theta_i &= U(0, 2\pi) \end{aligned} \right\} \text{ for } i = 1, 2, \dots, N_A \quad (3)$$

where  $N_A$  is the number of antennas and  $U(a, b)$  represents the output of a pseudo-random number generator that returns independent values uniformly distributed in the interval  $a \leq U < b$ . Two instances of configurations generated in this way are given in Figs. 2a and c for  $N_A = 20$ . While the corresponding  $uv$ -plane diagrams (Figs. 2b and d) show baseline nodes well scattered over the plane, with the density of nodes dropping toward the edge of the plot in accordance with Fig. 1, the two distributions are not equally acceptable because the nodes in Fig. 2d have a substantial hole near  $(u, v) = (0.5, 0)$ , and of course an identical hole near the inverse point  $(-0.5, 0)$ , while those in Fig. 2b lack such a feature. Visual inspection of the two diagrams is enough to establish a clear preference for the upper configuration, and so the lower configuration is taken as a test case to be subjected to various algorithms in a search for one that will reliably identify its principal defect, the large hole among the baseline nodes.

### 2.1. The Voronoi tessellation

The first construct considered is the Voronoi tessellation (or tiling), which is shown in Fig. 3a for the test case. The  $uv$ -plane

is seen to be covered by a set of polygonal tiles, each of which contains precisely one node and consists of the set of all points in the plane that are closer to that node than to any other (Okabe et al. 2000). These Voronoi tiles are potentially useful in finding the hole in the distribution because their size is sensitive to the local density of baseline nodes: if the density of nodes is  $\sigma$ , then the average area of the tiles is proportional to  $\sigma^{-1}$  and the average diameter to  $\sigma^{-\frac{1}{2}}$ , and it might in principle be possible to locate the hole by finding a tile that is anomalously large. The tiles in the figure do seem to show this general dependence on density: those near the centre are smaller on average than those near the periphery of the diagram, as appropriate to a density of nodes following the predicted taper (Fig. 1), and the tiles within the hole are larger than those immediately surrounding it.

The tiles do not, however, indicate the presence of the main hole in a simple way because no single one of them corresponds to the hole: instead of the hole being filled by only one tile, it contains about ten of them. Another difficulty is that every tile contains a node and is therefore not empty, and this again prevents any tile from being regarded as a good representation of a hole. An empty polygon would be better in this regard.

### 2.2. The Delaunay triangulation

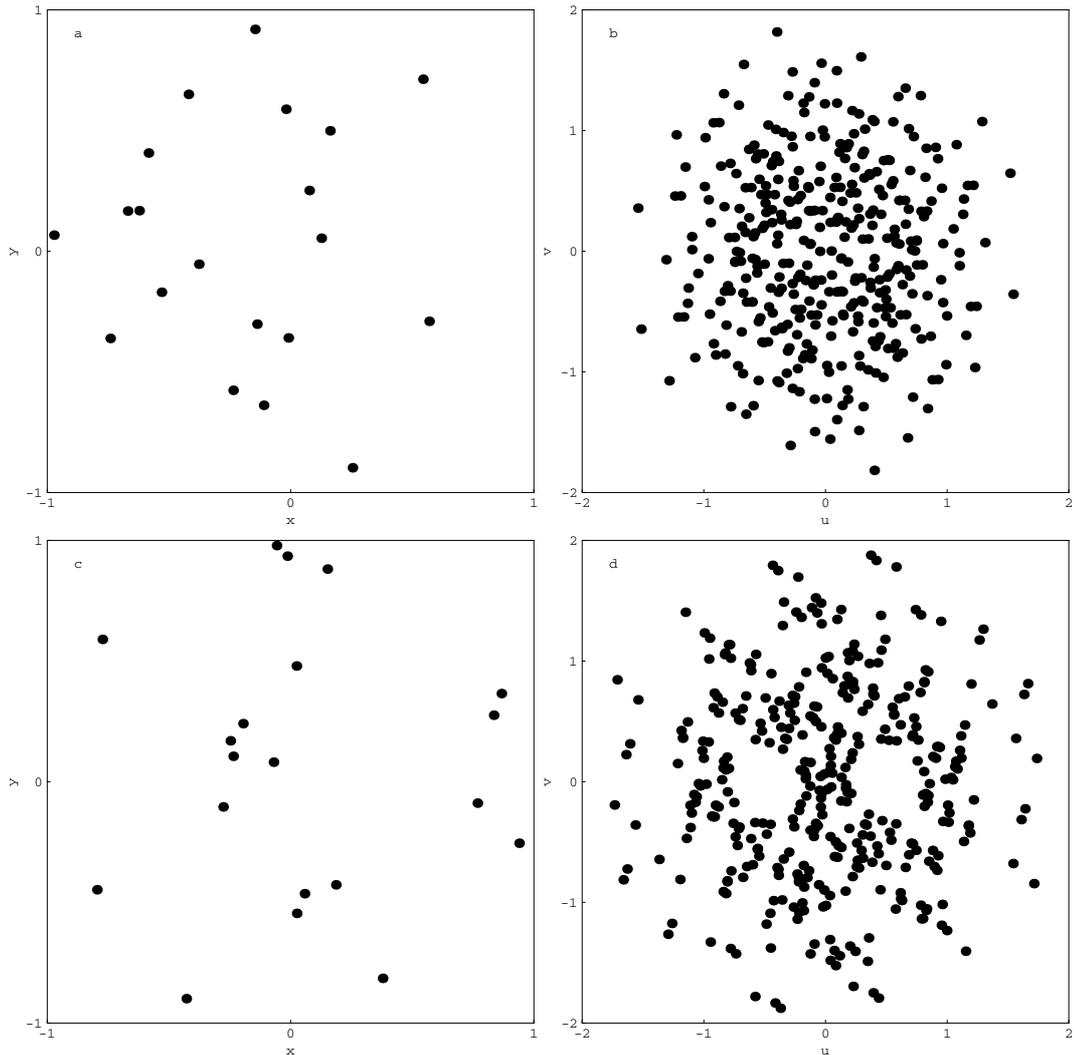
The Delaunay triangulation may be defined in several ways, one of which is as the dual of the Voronoi tessellation (Okabe et al. 2000): two baseline nodes are said to be natural neighbours if and only if their Voronoi tiles share an edge, and the Delaunay triangulation is obtained by joining every pair of natural neighbours by a segment of a straight line (Fig. 3b). There is an enormous number of different triangulations of a set of points as large as the set of baseline nodes under consideration here, but the Delaunay triangulation is usually unique (see Sect. 3.1) and always has a number of special properties, two of which will be exploited below.

The interiors of the Delaunay triangles are all empty of baseline nodes, and so the triangles are more suitable than the Voronoi tiles as potential representations of holes. In contrast to the Voronoi tiles, the triangles do not tessellate the entire  $uv$ -plane but are confined to a finite polygon termed the convex hull of the nodes whose edges form the outer boundary in Fig. 3b; a polygon is convex if none of its internal angles is greater than  $180^\circ$ , and the hull is the smallest convex polygon that contains all the nodes.

The areas and diameters of the Delaunay triangles have the desirable property of being sensitive to the local density of baseline nodes  $\sigma$  in the same way as those of the Voronoi tiles. The triangulation also shares a defect with the tessellation because no one triangle corresponds properly to the hole. There are about six large triangles inside the hole in Fig. 3b, and while this is fewer than the Voronoi tiles inside the hole it is still too many and just one would have been preferable.

### 2.3. The Delaunay circumcircles

In order to investigate further the problem of finding a construct that is a good match to the concept of a hole in a distribution of



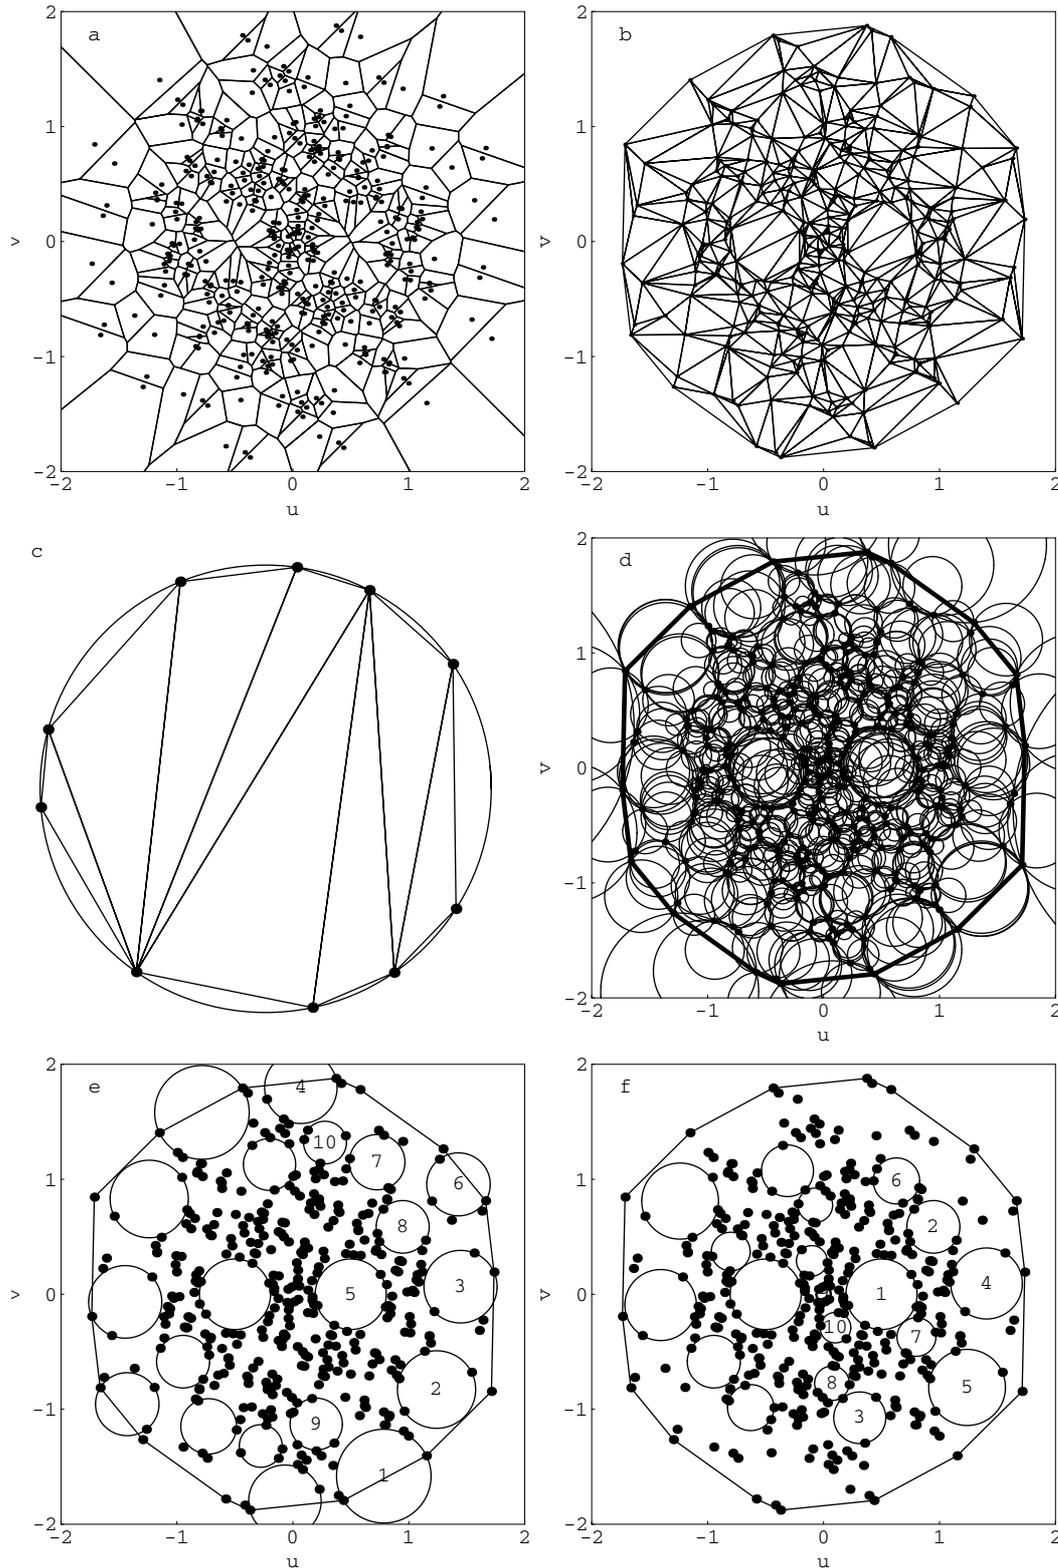
**Fig. 2.** Two antenna configurations (a) and (c) and their  $uv$ -plane diagrams (b) and (d)). The distribution of the baselines in (b) is broadly acceptable but that in (d) has an undesirably large hole near  $(u, v) = (0.5, 0)$ .

nodes, a simple model of such a hole was devised by randomly siting ten nodes on the circumference of a circle; the circle, nodes and Delaunay triangulation are shown in Fig. 3c. The nodes do not represent the baselines of any feasible array of antennas, and the model is intended simply to represent a hole in a set of nodes that is isolated and is not surrounded by a large number of other nodes. The triangulation consists of eight triangles, none of which is particularly large, so none of them on its own is suitable to represent the hole.

There are geometrical entities in that diagram that are suitable, however, and one of them is the union of all the triangles and their boundaries. This union is the largest empty convex polygon whose edges are Delaunay edges (i.e. edges of Delaunay triangles) and is the convex hull of the ten nodes. The convex hull is easily calculated and seems at first to be an interesting candidate for an approximation to the hole. Difficulties arise, however, when considering the larger problem of approximating the hole in the distribution of baseline nodes given in Fig. 2d. The convex hull of the nodes in that case encloses the entire figure and bears little relationship to the hole, which is

a relatively small feature embedded within it. In this context, therefore, some other method is required in order to generate the desired polygon. The polygon is an empty, convex polygon whose vertices are a subset of the nodes and is the largest of that type that can be formed from nodes in the vicinity of the hole, but this definition does not appear to correspond to any of the standard constructs of computational geometry and devising a general algorithm for it does not seem to be straightforward.

Attention was therefore directed to a second feature in Fig. 3c, namely the circle, which is another good approximation to the concept of the hole. It too has the desirable property of being empty (nodes on the circumference not being regarded as lying within) and, unlike the empty convex polygon, the circle is readily found because it is the circumcircle of any of the Delaunay triangles. Furthermore, the idea of using circles to represent holes can be made to carry over successfully from the model to the actual interferometric  $uv$ -plane distribution provided the circles are defined as the circumcircles of the Delaunay triangulation of the nodes. One of the special properties of the Delaunay triangulation is that it is the only



**Fig. 3.** **a)** the Voronoi tessellation of the baseline nodes; **b)** the Delaunay triangulation of the nodes; **c)** the Delaunay triangulation of ten points randomly placed on the circumference of a circle; **d)** the Delaunay circumcircles of the baseline nodes (the polygon in this and subsequent panels is the convex hull of the baseline nodes); **e)** the ten largest empty circles, the numbers indicating their ranking according to radius, together with their unlabelled inverses; **f)** as for the previous panel but with the circles ranked by their scaled radii.

triangulation whose circumcircles are all empty (Okabe et al. 2000) and, as was pointed out above, emptiness is a desirable property of any geometrical figure taken to represent a hole.

A second important property of the circumcircles is that they contain the largest empty circles that it is possible to draw amongst the nodes. This may be shown by defining  $R(u, v)$  to

be the radius of the largest empty circle centred on the arbitrary point  $(u, v)$  and then studying the local maxima of this function as  $u$  and  $v$  are varied.  $R(u, v)$  is evidently identical to the distance from  $(u, v)$  to the nearest node, and is a maximum wherever two conditions are met: firstly, that the point  $(u, v)$  is equidistant from three such nodes (i.e. lies at a Delaunay circumcentre or, what is identical, a Voronoi vertex); and secondly that  $(u, v)$  lies within the (Delaunay) triangle formed by those three nodes. An equivalent statement is that the local maxima of  $R$  all lie at the circumcentres of those Delaunay triangles that are acute, with none of their internal angles greater than a right angle. It follows that the largest empty circles that it is possible to draw in the various parts of the  $(u, v)$ -plane are all to be found amongst the Delaunay circumcircles.

The circumcircles of the Delaunay triangles of the baseline nodes are given in Fig. 3d, where they may all be seen to be indeed empty. The circumcircles appear to be more promising than any of the constructs tried so far because several of them lie within the main hole in the nodes and are more or less as large as the hole is.

There are, however, two matters that need to be addressed before the technique is suitable to assist the designers of a telescope, the first of which is the existence of circles well outside the region of the  $uv$ -plane occupied by the baseline nodes. These exterior circles cannot be said to lie among the nodes and therefore do not represent holes among the latter. They can, moreover, be very large and are in consequence likely to be undesirably prominent when sorting the circles by size to find the largest holes in the distribution. For both of these reasons it is necessary to eliminate them from further consideration, and some rational definition of an exterior circle must be adopted in order to do so. It is not difficult to devise definitions of exterior circles: they might, for example, be taken to be circles whose centres lie further from the origin of the  $uv$ -plane than the longest baseline. The particular definition adopted here is based on computational geometry for consistency with the general theme of this paper and a circle is defined to be exterior if its centre lies outside the convex hull of the baseline nodes. The hull is the smallest convex polygon that contains all the baseline nodes and so its boundary marks a natural division between the interior and the exterior.

The second matter is that, even after eliminating the exterior circles, many circles remain, and it would be preferable to reduce the number under consideration substantially. Many of the circles overlap each other, and as a result there can be several different approximations to the same hole whereas only one is required. This is particularly clear for the main hole in Fig. 3d, which has many large circles lying within it, one for each of the Delaunay triangles. A decision has to be made as to which of the overlapping circumcircles to select as representing the hole, and it seems most natural to opt for the largest of them.

It is therefore necessary to eliminate all the circles that overlap larger ones, and so a method was devised that makes use of two lists containing data records for the circles, each record containing the radius and the coordinates of the centre. The first, List A, contains the records of all the interior circles ranked in decreasing order of radius  $\rho_i$  while the second, List B,

is initially empty. The record of the largest circle is copied from List A to List B, and then that for the second-largest circle in List A is inspected and appended to the tail of List B if and only if the corresponding circle does not overlap the circle already in that list. The procedure is continued to progressively smaller entries in List A, appending each record to List B if and only if the corresponding circle does not overlap any of the circles already present in that list. At the end of this procedure, List B contains the desired set of non-overlapping circles, greatly reduced in number and ranked from largest to smallest.

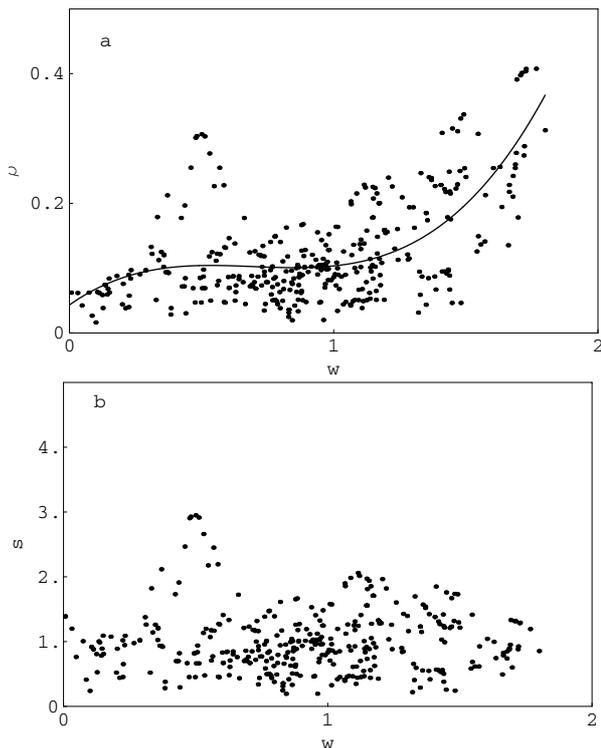
The ten largest circles for our example are plotted in Fig. 3e together with the baseline nodes: each circle appears twice because of the inversion symmetry, and those with  $u > 0$  are labelled with their ranking by radius, with the largest ranked first. It may be seen that the diagram is now much less cluttered than Fig. 3d. The circles do indeed appear to be good matches to the largest holes in the distribution of nodes and it does not appear possible to make any of the circles larger without causing it to engulf at least one node. The main hole near  $(u, v) = (0.5, 0)$  has been picked out clearly by a single circle, number 5.

#### 2.4. The rescaled circumcircles

While the algorithm presented in the previous section is largely successful, it does have one remaining defect: it ranks the circle representing the main hole fifth rather than first. The reason it fails to rank the main hole first is the existence of the large-scale gradient in the density of nodes caused by the taper, which causes the outer nodes to be spaced more widely on average than the inner. The result is that the largest circumcircles are mostly to be found around the periphery of the cloud of nodes. This behaviour of the algorithm is undesirable because the taper is a planned feature of the distribution of nodes and is not an unplanned defect that requires the attention of the designers of the telescope. The main hole is a different matter: it is considerably larger than expected from the taper and is very definitely an unplanned defect that does require their attention.

The algorithm therefore requires modification to remove the effect of the taper. In order to ensure that unexpectedly large holes are ranked highly wherever they lie, it must scale the size of each circumcircle relative to some estimate of the expected size given the taper and then take that relative size rather than the absolute size as the criterion for ranking. There are several ways of defining the expected size, one of which may be applied when the taper is known as it is in the present case in the form of the autocorrelation function  $C(w)$  in Eq. (2). The average radius of a circumcircle whose centre is a distance  $w$  from the origin of the  $uv$ -plane is proportional to  $C(w)^{-\frac{1}{2}}$  so that, if the  $i$ th interior circumcircle has radius  $\rho_i$  and is centred at  $(u_i, v_i)$ , then the scaled radius  $\rho_i C(w_i)^{\frac{1}{2}}$  (where  $w_i = [u_i^2 + v_i^2]^{\frac{1}{2}}$ ) has the desired property of indicating the size of the circumcircle relative to the local mean.

A second way, which may be applied whether the taper is known explicitly or not, is to define the average local size empirically by fitting a smooth function to a plot of the radii of all the interior circumcircles as a function of the positions of their centres in the  $uv$ -plane. This is particularly



**Fig. 4.** **a)** The radii  $\rho_i$  of the circumcircles plotted against the distance of their centres from the origin of the  $uw$ -plane,  $w_i$ . The curve is  $\bar{\rho}(w)$ , a cubic fitted by least-squares to the points; **b)** the scaled radii of the circumcircles  $s_i$  as a function of  $w_i$ .

straightforward when, as here, the design has an underlying circular symmetry because  $\rho_i$  may be plotted against  $w_i$  for all the circles and a smooth curve fitted. This plot is given for the present example in Fig. 4a where, despite the scatter, the general rise in  $\rho(w)$  with increasing  $w$  caused by the taper may be seen. A polynomial is a suitable choice of fitting function and  $\bar{\rho}(w)$ , the curve in Fig. 4a, is a cubic that was fitted by an unweighted least-squares method. The scaled radius is now defined as

$$s_i = \frac{\rho_i}{\bar{\rho}(w_i)} \quad (4)$$

and Fig. 4b gives a plot of  $s_i$  against  $w_i$ . It may be seen that the rising trend caused by the taper is no longer present, the points being scattered instead about the mean value  $s = 1$  at all values of  $w$ , with the result that the points related to the hole near  $w = 0.5$  have the largest values of  $s$ .

The largest scaled circles may be found by using the two-list method to eliminate overlapping circles as before, with the sole difference that the entries in List A must be ranked at the outset in decreasing order of  $s_i$  rather than  $\rho_i$ . The ten highest-ranked circles in List B are plotted in Fig. 3f where it may be seen that the circles are scattered more uniformly across the diagram than in Fig. 3e and are no longer concentrated to the periphery, demonstrating the successful removal of the effects of the taper. Most importantly, the main hole is picked out by the first-ranked circle, as desired. The three largest scaled radii are  $s_1 = 2.95$ ,  $s_2 = 2.05$  and  $s_3 = 2.01$  (where the subscripts

give the ranking in List B), so the main hole is considerably the most prominent in terms of the scaled size.

## 2.5. A variation

A third method of identifying which circumcircles are of most interest is a hybrid of the previous two. The algorithm for this variation proceeds as for the first case, in Sect. 2.3 above, setting up List A in decreasing order of  $\rho_i$  and generating List B of non-overlapping circles from it. Instead of stopping at that, however, the algorithm then continues by ranking the contents of List B in decreasing order of  $s_i$ . Each hole is thus represented by the circumcircle of largest absolute size that fits within it, but is then ranked by its size relative to the local mean. The resulting plot would show the same circumcircles as for the unscaled algorithm (as in Fig. 3e) but numbered differently.

The results of this and the previous algorithm that also makes use of  $s_i$  usually resemble each other closely, the main difference being that they sometimes select a different circle to represent any given hole. This is clear by comparing Figs. 3e and f, where although circles 5 in the former and 1 in the latter are identical, circles 3 and 4 are not and neither are circles 2 and 5. Nevertheless, the latter pairs differ very little and so it hardly matters which algorithm is selected.

The difference between the circles indicating a particular hole can occasionally be larger, however, especially when the hole is not approximately circular. An example is circle 7 in Fig. 3e, which lies in the same hole as circle 6 in Fig. 3f but differs appreciably in location and in size. The reason is that the hole in which they lie is pear-shaped, and the first algorithm fits the large bulge of the pear while the second fits the narrower extension because it lies nearer the origin of the  $uw$ -plane where the density of nodes is higher. The choice of algorithm is only significant if the designers of the telescope have a preference as to which of these alternative circles they wish to have drawn to their attention, but there is little to choose otherwise.

## 3. Discussion

### 3.1. Mathematical issues

There are two mathematical issues to discuss, the first of which is the complexity of the algorithms described here. The complexity is determined by  $N_B$ , the total number of baseline nodes, which is equal to twice the number of antenna pairs

$$N_B = N_A(N_A - 1). \quad (5)$$

The principal algorithm, the construction of the Delaunay triangulation, requires  $O(N_B \log N_B)$  operations (Fortune 1987; Okabe et al. 2000). The number of triangles is not determined exactly by  $N_B$  but is approximately  $2N_B$  (Okabe et al. 2000), and so the construction of the circumcircles, one per triangle, requires  $O(N_B)$  operations. List A has one entry per circle and so requires  $O(N_B)$  operations to set up; it must then be sorted, which can be done in  $O(N_B \log N_B)$  operations (Press et al. 1997). Finally, the entries in List A are considered one by one for possible transfer to List B, which is  $O(N_B)$ . The most complex procedures in the computation determine the complexity

of the whole, and so the algorithm requires  $O(N_B \log N_B)$  operations overall, or  $O(N_A^2 \log N_A)$  operations from Eq. (5).

The efficiency of the algorithm may be assessed by comparing the complexity with an obvious lower limit. Any algorithm for assessing the quality of the  $uv$ -plane coverage must at the very least read in the coordinates of all  $N_B$  nodes as input, which requires  $O(N_B)$  operations and sets a lower limit on the complexity. The algorithms presented here require  $O(N_B \log N_B)$  operations, only modestly greater than the limit, and so it is reasonable to conclude that the algorithms are efficient.

The second matter for discussion is the possibility that the Delaunay triangulation of any given set of baseline nodes is degenerate, i.e. not unique. The simplest case in which this happens is when four nodes are sited at the vertices of a square. In this case there are two Delaunay triangles, whose sides are made up of the sides of the square together with one of the diagonals, but either diagonal may be chosen and as a result there are two different Delaunay triangulations. This lack of uniqueness is of no concern here, however, because the Delaunay triangles themselves are not the object of attention but their circumcircles. Whichever Delaunay triangulation is selected from a degenerate set, the list of circumcircles is identical. This may be seen in the particular example of the square where the circumcircle of any of the triangles is the circumcircle of the square, whichever of the two triangulations is chosen.

The lack of uniqueness may nevertheless affect the choice of computer program to calculate the triangulation because, although some routines will arbitrarily choose one of the alternatives when the triangulation is degenerate, others will merely point out the ambiguity without resolving it (Okabe et al. 2000). The first kind of routine is required for the present purpose.

### 3.2. Applications

One application of the algorithms presented here is to make use of the values of  $s_i$  in order to compare different trial configurations of antennas. There are many ways of doing this, some of which are simpler than others, but they may all be considered methods of measuring the deviation of the distribution of radii from the ideal case in which all the circles are of the same size,  $s_i = \bar{s}$  where  $\bar{s}$  is the average value. An obvious measure is  $\sigma_s$ , the standard deviation of the  $s_i$  about  $\bar{s}$ , which is a figure of demerit, the smaller the better. Another general approach is to concentrate on the larger circles, those with  $s_i > \bar{s}$ , and to ignore the smaller ones. The rationale for this approach is that the small circles are not of primary importance; they are found where the baseline nodes are denser than the average, and the additional sampling of the correlation function they provide is not at all unwelcome in itself. The only problem associated with dense clumps of nodes is the indirect one, that if the nodes in these clumps had been distributed more evenly the low-density regions of the  $uv$ -plane might have been better sampled and so the largest circles would have been smaller. This may be addressed directly by considering only the very

large circles, and there are still many ways of doing this, one of which is to take the radius of the  $n$ th largest circle, where  $n$  might be chosen to be 3, say, or 5 or 10. Another is to calculate the average radius of the largest circles, say the 3 or 10 largest. The simplest, however, is to take  $s_1$ , the largest of the scaled radii, as a figure of demerit for any given configuration; the smaller the value, the better the configuration. This choice has the effect of making the comparison between configurations on the basis of the worst single defect in each.

Here we compare the two configurations given in Fig. 2 in this way. The configuration in Fig. 2a has  $s_1 = 1.65$ , considerably smaller than the  $s_1 = 2.95$  for Fig. 2c, and so this comparison provides rational grounds for inferring that the configuration in Fig. 2a has much the better coverage of the  $uv$ -plane, in agreement with the inference drawn in Sect. 2 above on subjective grounds.

A second application is to make use of the technique to help improve a given design. If a trial layout of the antennas gives a  $uv$ -plane distribution with one or more large holes, it may be improved by adjusting the positions of some of the antennas in such a way as to reduce the size of those holes. There are many alternative ways of doing this, but it is likely to be preferable to move as few antennas by as short a distance as possible in order to preserve the major features of the design. In reducing the size of any hole, the obvious baselines to adjust are the three on the circumference of the circumcircle, and the direction in which to move them is toward the centre of that circle. All the alternatives would involve longer moves, and so the three baselines on the circumcircle indicate which antennas should be moved and in which direction. This and other ways of optimizing designs will be discussed more fully in the second paper in this series.

The new algorithm was developed for application to the time-honoured technique of designing aperture-synthesis interferometers in which the general form of the distribution of antenna stations over the  $xy$ -plane is decided in advance; this immediately determines the large-scale structure of the distribution of nodes in the  $uv$ -plane and it only remains to check that the small-scale structure of that distribution is satisfactory for any specific siting of the antennas and to improve the siting of the antennas if it is not. This is not the only possible technique, however, and a different approach has recently been described in which the antennas are sited initially at arbitrary points in the  $xy$ -plane and then their positions are adjusted iteratively in order to optimize some function measuring the quality of the corresponding distribution in the  $uv$ -plane (Boone 2001). Since it is possible to define this function in order to control the quality of the small-scale distribution of nodes in the  $uv$ -plane as well as the large-, the new technique may be implemented in such a way that it is less likely to leave large holes in the coverage. The new algorithm may therefore not be as useful in this technique as in the older one, although it could still be employed as a final check that all has indeed gone well and to provide grounds for choosing between a number of solutions obtained by applying the same optimization to different starting positions for the antennas.

### 3.3. The assumptions

The development was based on a number of simplifying assumptions (see Sect. 2), but they are not fundamental to the technique and it can be adapted easily enough to take account of the more complicated situations that arise in practice.

The technique was described by confining attention to the simple case of snapshot observations made at the local zenith, but it is readily extended to snapshot observations made in other directions. For each direction, the  $uv$ -plane is defined as lying perpendicular to that direction, so the nodes are projected parallel to the observing direction into that plane and the largest empty circles within the projected distribution of nodes may be found as before. The properties of all the circles associated with all of the chosen directions may then be considered in order to evaluate any single configuration of antennas: a suitable figure of demerit, for example, would be the largest scaled radius of all these circles. In so far as this method is equivalent to finding ellipses in the zenith projection to represent the same holes that are already represented by circles, it does not usually add much to what can be found more economically from the zenith projection alone, although there is one important case when it does. This is when the antennas are not all sited in the same plane but stand on uneven ground and are therefore distributed above and below the best-fitting plane. It is then necessary to take the  $z$ -coordinates of the antennas into account when calculating the baselines and projecting them onto the  $uv$ -plane, and the differences in altitude need not be very large for the projected positions of the baseline nodes in the  $uv$ -plane to differ substantially from those obtained by applying a simple shear to the zenith projection. If the altitudes above and below the best-fitting plane are comparable in magnitude with the separation of neighbouring nodes in the  $uv$ -plane, the microstructure of the distribution of nodes will be significantly changed, and in this case it is likely to be well worth investigating the holes in the distributions corresponding to a large number of possible observing directions.

The technique may also be extended to cover Earth-rotation observations. In the course of these observations, which are made over an extended period of time, the projected positions of the baseline nodes move steadily over the  $uv$ -plane in trajectories that are elliptical arcs. The holes in the coverage may be approximated in the same general way as before if the largest circles that fit amongst the arcs can be found. This appears to be a well-posed problem of computational geometry, but no exact solution is to be found in the standard reference (Okabe et al. 2000), not could we devise one. An approximate solution may readily be found, however, by breaking up each elliptical arc into a series of points and using the methods of this paper to find the largest circles that fit amongst the points. The principal consideration is to decide how large the interval between adjacent points should be: too large and the points would not approximate the arcs very well, but too small and the computational load would become heavy as a result of the large number of points (Sect. 3.1). The interval might be set so that a side of average length has several points along its length, perhaps 3–10.

Each circle in the Earth-rotation case lies typically within a “triangle” made up of three elliptical arcs. The shape of these triangles varies considerably from one to another, but in the outer parts of the coverage of the  $uv$ -plane, where the arcs tend to be almost parallel, the triangles are often elongated, and so the circles would all be relatively small, comparable in diameter to the smallest diameter of the triangle. The question then arises whether the size of such a circle is the appropriate choice to represent the size of the hole, given that the triangle has a small and a large diameter which might well differ considerably in size. The answer is that the circle is appropriate, because the visibility function is well-sampled along the arcs forming the side of the triangle and the interpolation of its value into the interior of the triangle is most accurately done across the narrowest dimension of the triangle, whose size is well approximated by the size of the circle.

## 4. Conclusions

A solution has been presented to the problem of devising an algorithm to find and quantify the holes in the distribution of baseline nodes in the  $uv$ -plane of a design for an aperture-synthesis interferometer. It makes use of the largest empty circle in each region of the plane as a proxy for the local hole. The algorithm first constructs the Delaunay triangulation of the nodes and then finds the circumcircles of all the triangles, which are always empty for this particular triangulation. Circles whose centres lie outside the convex hull of the baseline nodes are then ignored because they do not lie among the nodes and cannot be said to represent holes among them. Finally, the remaining circles are sorted into decreasing order of size and smaller circles are eliminated if they overlap larger ones, leaving a set of non-overlapping circles each of which corresponds to a hole in the  $uv$ -plane distribution of the baseline nodes. The largest of these circles corresponds to the largest hole in the  $uv$ -plane coverage.

Making use of the absolute size of each circle as the basis of the sorting has the unwanted effect that the largest circles lie preferentially in the part of the diagram where the density of nodes is lowest as a result of the desired taper, whereas the designers are usually more interested in finding holes that are unusually large for their locality. Modified algorithms were therefore presented which, by ranking the circles in decreasing order of size scaled relative to the local mean size, pick out egregiously large circles wherever they lie.

The properties of the circumcircles may be used in several ways. They may be used to decide which of a number of candidate designs for a given telescope has the best  $uv$ -plane coverage by taking the largest scaled radius of each design as a figure of demerit, choosing the design with the smallest value as having the best  $uv$ -plane coverage. They may also be used to help improve the  $uv$ -plane coverage of any given design because the three baselines on the circumference of the circumcircle with the largest scaled radius are those which, if moved towards the centre of the circle by adjusting the positions of some of the corresponding antennas, will provide the greatest improvement of the coverage while causing the least disturbance to the original design.

The largest empty circle that may be drawn among the nodes in any given region of the  $uv$ -plane is a mathematical concept that is a satisfactory match to the subjective concept of a hole among those nodes. It is precisely defined and readily computed, and may be used as the basis of a rational investigation of the quality of the coverage of the  $uv$ -plane by the baseline nodes.

## References

- Boone, F. 2001, *A&A*, 377, 368
- Bracewell, R. 1965, *The Fourier Transform and its Applications* (London: McGraw-Hill)
- ESO 2003, *The Atacama Large Millimetre Array*, <http://www.eso.org/projects/alma/>
- Fortune, S. J. 1987, *Algorithmica*, 2, 153
- Hjellming, R. M. 1989, in *Synthesis Imaging in Radio Astronomy*, ed. R. A. Perley, F. R. Schwab, & A. H. Bridle (San Francisco: Astronomical Society of the Pacific), 477
- NRAO 2003, *The Atacama Large Millimetre Array*, <http://www.alma.nrao.edu/>
- Okabe, A., Boots, B., Sugihara, K., & Chiu, S. N. 2000, *Spatial Tesselations: Concepts and Applications of Voronoi Diagrams*, 2nd ed. (Chichester: John Wiley & Sons)
- O'Rourke, J. 1998, *Computational Geometry in C*, 2nd ed. (Cambridge: Cambridge University Press)
- Perley, R. A., Schwab, F. R., & Bridle, A. H. 1989, *Synthesis Imaging in Radio Astronomy* (San Francisco: Astronomical Society of the Pacific)
- Preparata, F. P., & Shamos, M. I. 1985, *Computational Geometry: An Introduction* (Heidelberg: Springer)
- Press, W. H., Teukolsky, S. T., Vetterling, W. T., & Flannery, B. R. 1997, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed. (Cambridge: Cambridge University Press)
- Shaver, P. A. 1996, *Science with Large Millimetre Arrays*, (Heidelberg: Springer)
- Taylor, G. B., Carilli, C. L., & Perley, R. A. 1999, *Synthesis Imaging in Radio Astronomy II* (San Francisco: Astronomical Society of the Pacific)
- Thompson, A. R., Moran, J. M., & Swenson, G. W. 2001, *Interferometry and Synthesis in Radio Astronomy*, 2nd ed. (New York: Wiley)
- Walker, R. C. 1984, in *Indirect Imaging*, ed. J. A. Roberts (Cambridge: Cambridge UP), 53
- Wootten, A. 2001, *Science with the Atacama Large Millimeter Array (ALMA)* (San Francisco: Astronomical Society of the Pacific)