PART II

APERTURE SYNTHESIS WITH LIMITED OR NO PHASE INFORMATION

https://doi.org/10.1017/S0252921100074753 Published online by Cambridge University Press

FUNDAMENTAL ASPECTS OF APERTURE SYNTHESIS WITH LIMITED OR NO PHASE INFORMATION

(Invited paper)

J.E. Baldwin and P.J. Warner
Cavendish Laboratory, Cambridge, England.

Much ingenuity has been devoted to the determination of mutual coherence functions at radio wavelengths. Measurements of the amplitude and phase of the interference fringes with uncertainties of only 1 percent are now attainable. Unfortunately the atmosphere intervenes between the measured quantities and the mutual coherence function itself and the phase uncertainties introduced by the ionospheric and tropospheric irregularities are often the factor limiting the quality of radioastronomical maps.

In this paper we describe methods for making maps from data with poorly determined or even unknown phases. We discuss in turn the effects of errors on the maps, methods of phase determination using reference sources, the properties of hybrid maps and their application to mapping of fields containing point sources, extended sources or a mixture of the two. We assume in all our discussions that the data are adequately sampled for the problem in hand and will leave the problems associated with incomplete sampling of the mutual coherence function to the cleaners and those able to maximise entropy.

1. THE EFFECTS OF ERRORS ON THE MAPS

We enquire first under what circumstances the phases limit the quality of a map. Suppose the 2-dimensional mutual coherence function comprises a set of amplitudes and phases $A(x,y),\phi(x,y)$ having uncertainties ΔA, $\Delta\phi$. Since the fourier transform is a linear process, the resulting images can be considered as the sum of the ideal map and those obtained by fourier transformation of $(\Delta A,\phi)$ and $(iA\Delta\phi,\phi)$. These two error maps have similar properties; the only outstanding difference arising from the fact that in most work $A(-x,-y)$ and $\phi(-x,-y)$ are not measured but derived assuming the mutual coherence function to be conjugate symmetric. Thus the error map associated with ΔA is symmetric whilst that associated with $\Delta\phi$ is antisymmetric. The rms sidelobe levels are proportional to ΔA and $\Delta\phi$. There is little published data on the amplitude stability of aperture synthesis observations but there can be little doubt that the variations are mainly

C. van Schooneveld (ed.), Image Formation from Coherence Functions in Astronomy, 67–82.

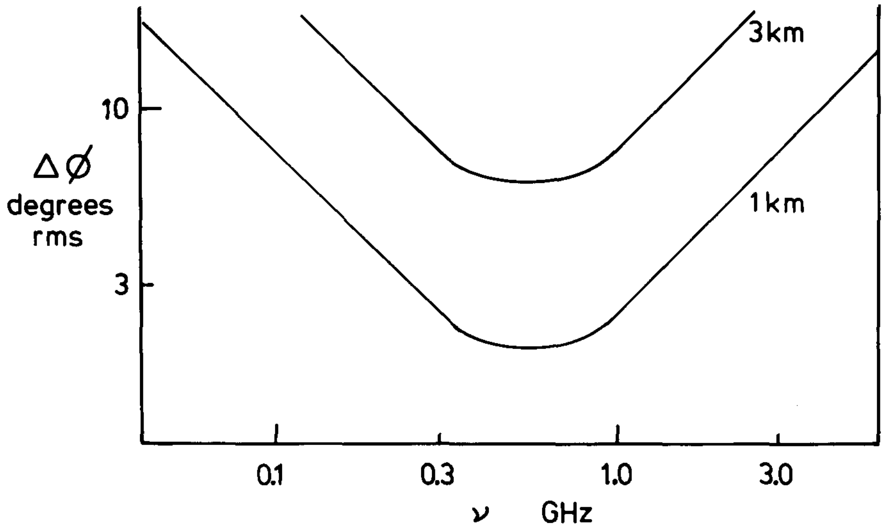

Fig 1. Typical rms phase variations due to ionospheric and tropospheric irregularities as a function of frequency and baseline.

instrumental in origin. Experience at Cambridge suggests that $\Delta A/A$ may be as small as 1 per cent under good conditions. The vagueness about ΔA merely reflects the fact that it is usually small compared with $\Delta\phi$. No authors have claimed phase uncertainties as small as 0.5 deg (i.e. 1 per cent of a radian). Even if instrumental sources of error are excluded the atmospheric contribution to ϕ is important for most synthesis telescopes. Fig 1 shows the rms phase deviations expected under average conditions as a function of frequency and baseline. Variations in atmospheric conditions may cause changes of at least a factor of 3 about these mean values. It is clear that for all existing aperture synthesis instruments the effects of phase deviations are likely to outweigh those of amplitude.

There is a wide range in the time scales and linear scales of the phase fluctuations. The most important ones are the long period large scale features, which produce an rms sidelobe level of roughly $0.2\ t^{\frac{1}{2}}(\mathrm{hrs})\Delta\phi_{max}(\mathrm{rad})$ over an area of sky of $\sim 16\ t^{-2}(\mathrm{hrs})$ beam areas. As an example the 5 km telescope operating at 15 GHz is affected most by fluctuations with time scales of roughly 30 mins giving rms phase fluctuations $\Delta\phi_{max}$ (at the largest baseline) of 30 deg. The resulting sidelobes are 6 per cent over an area of 64 beams.

2. PHASE DETERMINATION FROM REFERENCE SOURCES

Astronomers have devised many calibration techniques to overcome the desperation caused by poor phases. All rely on the idea of a reference source. In what follows we shall assume that everything possible by way of phase calibration using sources external to the field of view has been done, but that the resulting maps are still limited by unknown atmospheric or instrumental phase fluctuations.

There are a number of methods of improving maps which use reference sources *within* the area of the map. Any point source on the map may be used to give the interferometer phase for each spacing by 12 hour averaging, but signal-to-noise and confusion will limit accuracy. It may even be possible to use shorter term phase averages to correct the phase. At first sight this seems an ideal way of overcoming atmospheric phase problems. The source is exactly where it is needed. But the apparent fluctuations in the phase of the reference are partly due to the atmosphere and partly real, i.e. due to the interference between the reference and other sources in the field. Removal of the fluctuations by some type of running mean distorts the map. In general the map is accurate in the outer parts and progressively less so as the distance from the reference source is decreased.

The limit of this procedure is reached when the time scale for correcting the phase is extremely short. All variations of the phase are attributed to the atmosphere and the apparent phase of the reference source is set equal to zero. We have discussed this approximation elsewhere (Baldwin & Warner, 1976).

The reference source need not be a single object. Any distribution of sources whose relative positions and flux densities are known can be used to calculate reference amplitudes and phases as a function of HA and spacing which can then be used for calibrating the observation. An example of this technique applied to observations of the field of view of the 5C7 survey (Pearson & Kus, 1978) is shown in Fig 2. There are several sources having flux densities in the range 0.5 to 1 Jy and the comparison between reference phase and the observation is so good that the interferometer phase can be measured to an accuracy of 2 deg.

These techniques of internal calibration require the isolation of an observed phase and its correction to some reference level. Two separate effects can combine to make this impossible: the signal-to-noise may be inadequate to correct variations on short time scales and it may not be possible to construct a good enough phase reference. An obvious criticism of this method is that it implies prior knowledge of the distribution of sources in the field of view. How this difficulty is overcome is described below.

3. HYBRID MAPS AND THEIR APPLICATION

We have recently described an iterative technique derived from crystallographic procedures to make maps from amplitude data alone (Baldwin & Warner, 1978). The starting point is a map made using the square of the observed amplitudes, which is independent of the observed phases. It is the autocorrelation of the true distribution: we call it a *Map of* A^2. It is possible to analyse this map to obtain an

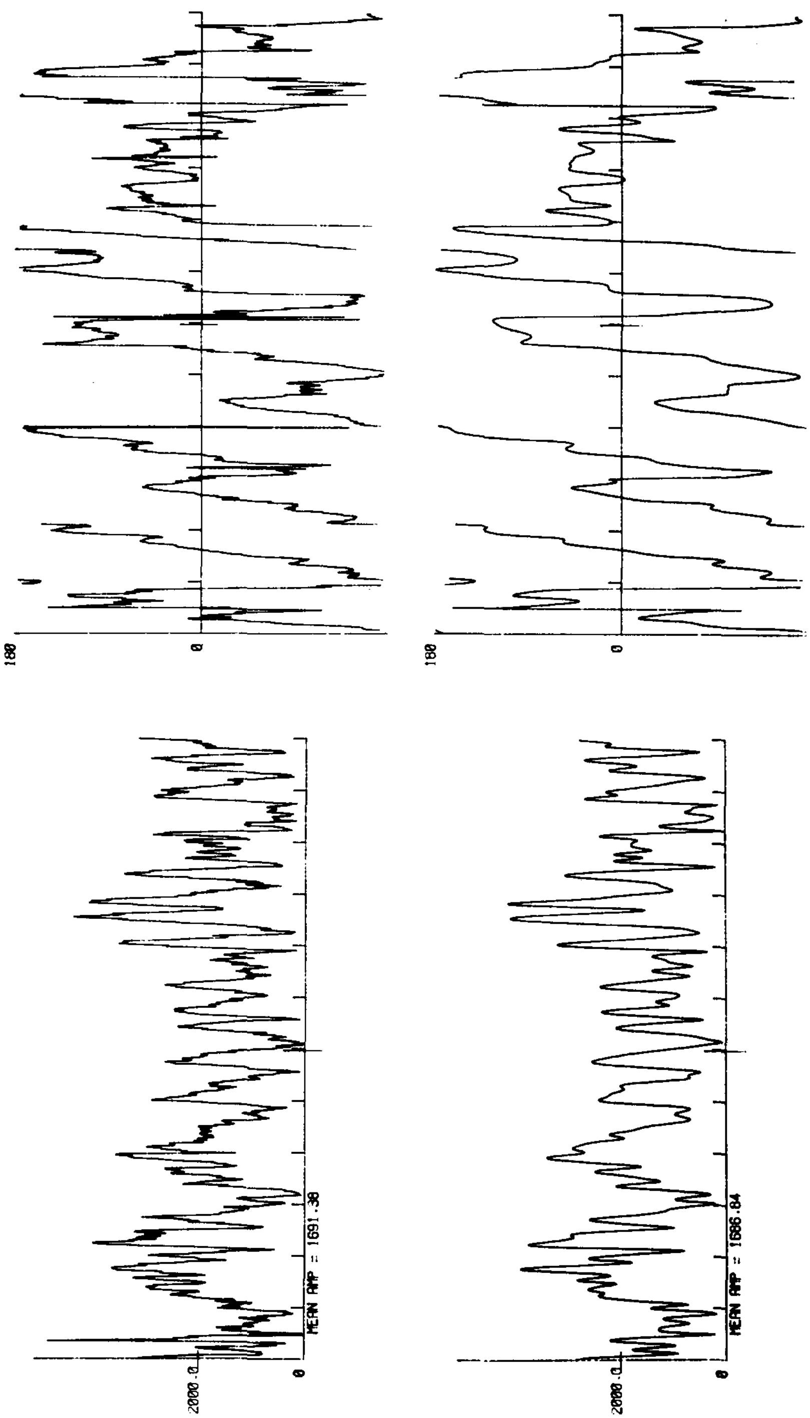

Fig 2. Calibration using field sources as an internal reference. Amplitudes and phases as a function of Hour Angle at a baseline of 718λ for the 5C7 survey. Upper curves; observations. Lower curves; 186 source model.

https://doi.org/10.1017/S0252921100074753 Published online by Cambridge University Press

approximate model of the true brightness distribution. However, if observed phases are available, the map based on them can provide the starting model. This is used to calculate phases which are combined with the observed amplitudes in a *Hybrid map*. Strictly, hybrid maps are made using amplitudes which are appropriate to one brightness distribution with phases which are appropriate to a different one. They have been used by several authors (e.g. Fort & Yee, 1976; Baldwin & Warner, 1978; Readhead & Wilkinson, 1978). An understanding of their properties is important as a basis for assessing the progress of the iteration, its convergence and speed of convergence.

The simplest case is that mentioned above, the *phase zero map*. The amplitude is correct for the particular distribution of emission whilst the model for the phase is a point source at the centre of the map. Where there is a bright source in the field of view the map made from this combination has the following properties:

(1) The bright reference source is at the map centre.
(2) Every other source occurs at its correct relative position in the map but with half its true intensity.
(3) There is an identical spurious image for every source situated with the reference source as a centre of symmetry.
(4) There are fainter spurious peaks which correspond to the harmonics of, and cross modulation terms between, the various sources in the map. They occur at positions in simple geometrical relationships to the positions of the true sources.

(We shall show an example of the more complicated case where no source dominates later).

The analysis of such a map is straightforward only if the reference source is situated at the edge of the field of view so that the correct and spurious images can be distinguished easily. A case where this technique was applicable has been described by Riley & Pooley (1978) who used it for studies of 3C 123 at 15 GHz with the 5 km telescope. The structure of 3C 123 shows a large range in surface brightness and a bright, only slightly resolved feature, occurs at one side of the source. The original map using the measured phases is shown in Fig 3a whilst that obtained by setting the phases to zero is illustrated in Fig 3b. The improvement over the uncorrected map for the discrimination of faint features is a factor of about 3.

A more commonly occurring case is that of a hybrid map containing many point sources. If the relative positions and flux densities of some of the sources in the field of view are known then a hybrid map can be made with the observed amplitudes and the phases of a trial distribution composed of these point sources. The properties of this kind of hybrid map are:

(1) The sources in the trial occur in their correct relative positions and with nearly their correct intensities.
(2) Features in the true map but not in the trial appear in their correct positions but with half their true intensities.
(3) A large number of fainter spurious responses, both positive

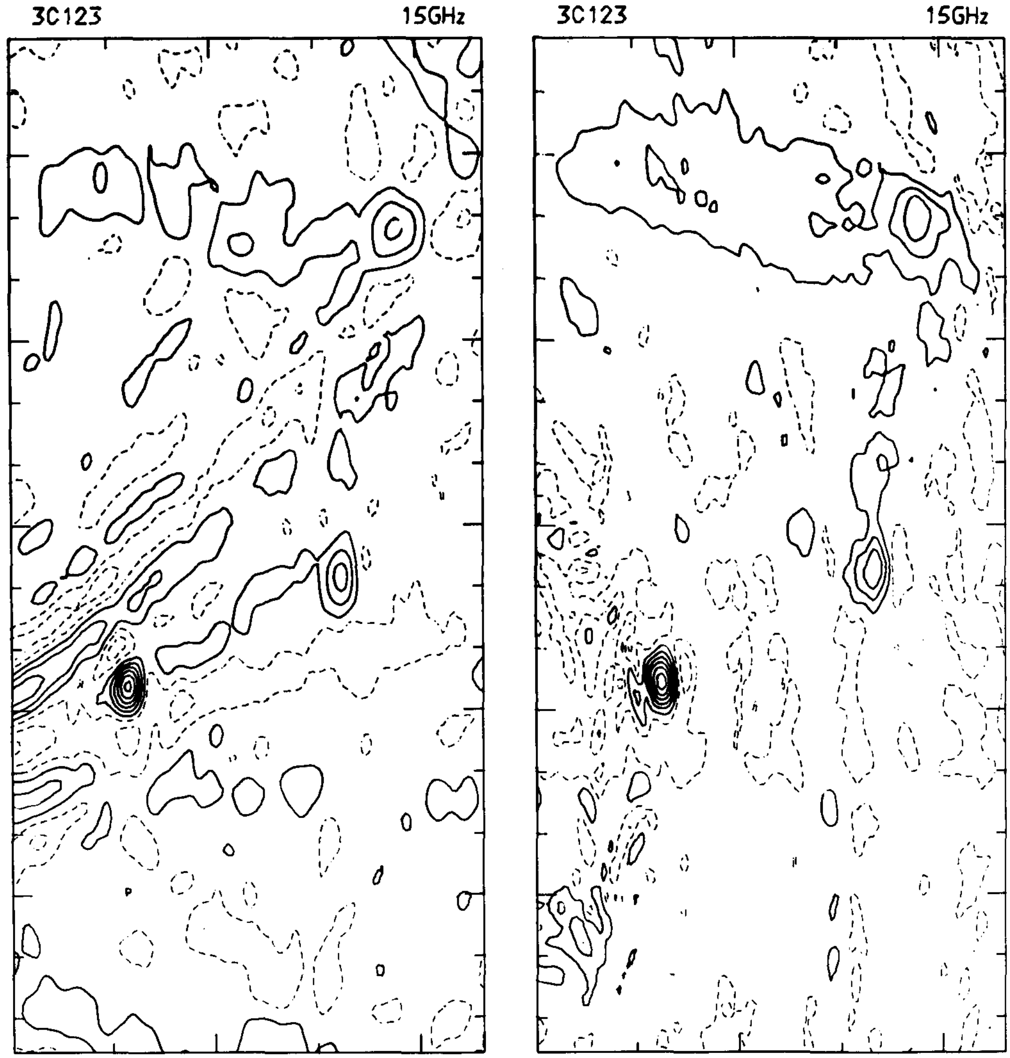

Fig 3. Improvement of sidelobe levels by using phase zero maps. Observations of 3C 123 at 15 GHz with a resolution of 0.65 arcsec. Left hand map; observed amplitudes and phases. Right hand map; observed amplitudes and phase zero. The bright reference source is the SE component which lies off the lower left hand edge of each picture. It would have 70 contours in Fig 3a and 140 in Fig 3b.

and negative occur at positions in a simple geometrical relation to the sources in the true map.

In practice it is often useful to plot the *difference map* i.e. hybrid map minus trial map.

It is clear that the best way of improving the trial is to add to it those features in (2) but at twice their apparent intensity. It is this property of hybrid maps which is the basis for the iterative improvement of the trial map. The problem is how to distinguish the real features from the spurious peaks. The brightest feature in the hybrid map will almost always correspond to a real source; the spurious peaks are fainter by a factor $a_m a_n/\Sigma a^2$ where a_m is the flux density of a source in the trial map and the sum is over all sources in the trial map. If the number of sources of comparable intensity in the trial map is large then the factor may be << 1, many new peaks can be included correctly at each iteration and convergence to the true map is extremely rapid.

We have described a test of this procedure using data from a 5C survey. The 5C7 survey (Pearson & Kus, 1978) at 408 MHz found 5 sources between 0.4 and 1 Jy and about a further 250 above the limiting flux density of 12 mJy. The noise level on such maps should be 0.5 mJy but Pearson & Kus found fluctuation with an rms of 2 mJy; they attributed this discrepancy to the effects of confusion and to calibration errors.

Using observed amplitudes and model phases the noise level is expected to be slightly more than $\sqrt{2}$ times worse than the ideal value, i.e. 0.8 mJy. On a difference map the effects of confusion are greatly reduced so the phaseless method offers the possibility of reaching fainter sources. Recent analysis shows that our earlier failure to reach the expected limiting sensitivity was mainly due to a lack of perseverence. Now after 6 iterations, 186 sources have been incorporated in the trial map down to a flux density of 14 mJy. The difference map shows about a further 100 sources down to 10 mJy and the rms noise level, excluding these sources, is 1.8 mJy or about 2.2 times that expected theoretically. The position and flux densities are in good agreement with those of Pearson & Kus. We do not know the reason for the remaining discrepancy in the noise level. A necessary precaution in reaching a result as good as this was the omission from the analysis of those samples in the aperture plane in which the amplitude fell below 2σ. Inclusion of such points both raises the apparent noise level on the difference map and was found to increase the flux densities of the sources in the trial map by about 4 per cent above their true values.

We conclude that this technique is successful for fields composed of point sources and that it may provide a useful extension of the dynamic range in maps which are limited by phase fluctuations in the data.

An important question concerning the method is how good must the starting model be? A simple estimate is that at least half the amplitude in any sample must be due to sources in the trial map. Under these circumstances the true phase must always lie within $^{\pi}/_{2}$ of the trial phase and the synthesis, though of poor quality, is certainly good enough to identify correctly the brightest peaks in the difference map. A series of model examples is shown in Fig 4. The true map contains 10 sources (in fact lying at the positions of the 10 brightest sources in the 5C7 survey). Difference maps corresponding to trial maps containing 1, 2, 3, and 4 sources respectively are illustrated in Figs 4a, 4b, 4c, 4d. The positions of the 10 sources have been marked by dots. All the other peaks are spurious images. In Fig 4d it is clear that all of the remaining true sources are brighter than the spurious images. In Fig 4b the brightest true source is only 1.7 times brighter than the largest spurious image and at most 2 sources could be added to the trial map at the next iteration. It is not disastrous to the method if a few spurious peaks are incorrectly assumed to be true sources, it only causes some oscillation in the convergence. Our experience in analysing the 5C7 data was that, at each iteration, the flux density of the faintest sources incorporated in the trial map could be reduced by a factor of about 2.

The quality of the initial trial map would be of no great importance if the analysis could be shown to converge from any starting point and that the result was unique. This last point is an important one. For any given distribution of amplitudes in the aperture plane there may be a large number of distributions of phase which would give acceptable maps. Provided that $\phi(x,y) = -\phi(-x,-y)$ the map will be real and the only criterion for choice then is that the intensity must be positive and that there may be some knowledge of the extent of the source. However, alternative solutions satisfying both criteria can sometimes be found in simple one dimensional cases.

Even if the positivity constraint is not satisfied strictly the region of negative intensity may be shallow and the map appear to be of acceptable quality. To illustrate how misleading an incorrect trial can be, a hybrid map made from a uniform amplitude A with trial phases appropriate to the same 10 sources used in Figs 4 is shown in Fig 5. Strong peaks occur at the positions of the 10 sources with flux densities in the ratio used in the trial. There are many negative sources on the map but the most intense are only about 10 per cent of the intensity of the main sources. This result bears no relationship to what all astronomers would agree is the correct solution to the uniform A problem, namely a single point source, and it illustrates how dangerous wrong initial models can be.

This ambiguity in the solution contrasts sharply with the apparent uniqueness of the result in the analysis of the 5C7 data. Could other solutions have been found for 5C7? Not if the basic assumption that the map is composed of point sources is adhered to. Each unresolved peak in the map of A^2 arises from the coincidence of two point sources in the autocorrelation process rather than from some ingenious distribution of positive and negative amplitudes which gives a value of the

https://doi.org/10.1017/S0252921100074753 Published online by Cambridge University Press

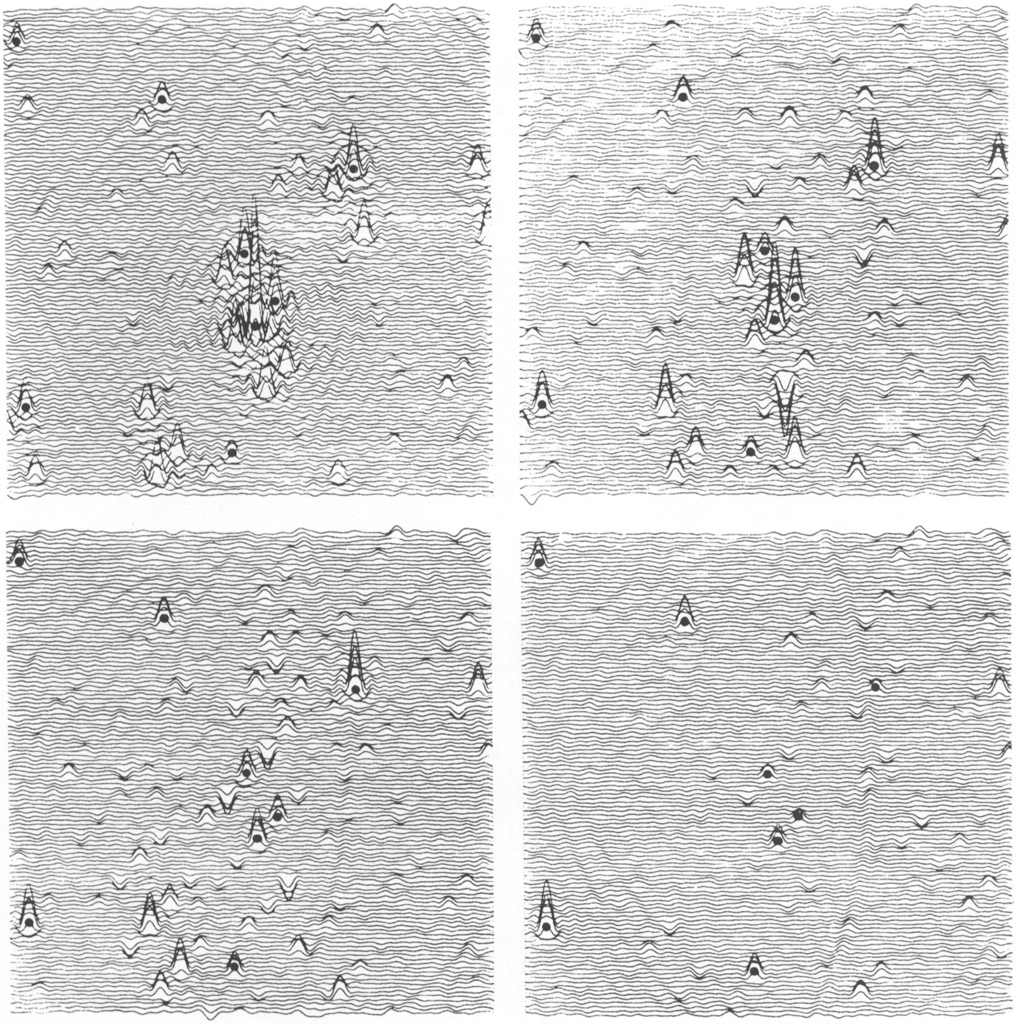

Fig 4. The effect of different trial distributions on the interpretation of hybrid maps. Difference maps made for a distribution of 10 point sources and trial distributions comprising successively 1, 2, 3 and 4 of the sources. Positions of 8 of the sources are marked by • , the remaining 2 lie outside the plotted area. The four sources of the trials are the central three and the upper right hand one. Notice that residual peaks at their positions gradually become smaller.

autocorrelation function at that point and zero at all surrounding points. The N(N-1) vectors in the A^2 map then uniquely define the arrangement of the N sources. Even if some of the faint vectors are undetected, there is usually no possibility for ambiguity in the arrangement of the bright sources.

https://doi.org/10.1017/S0252921100074753 Published online by Cambridge University Press

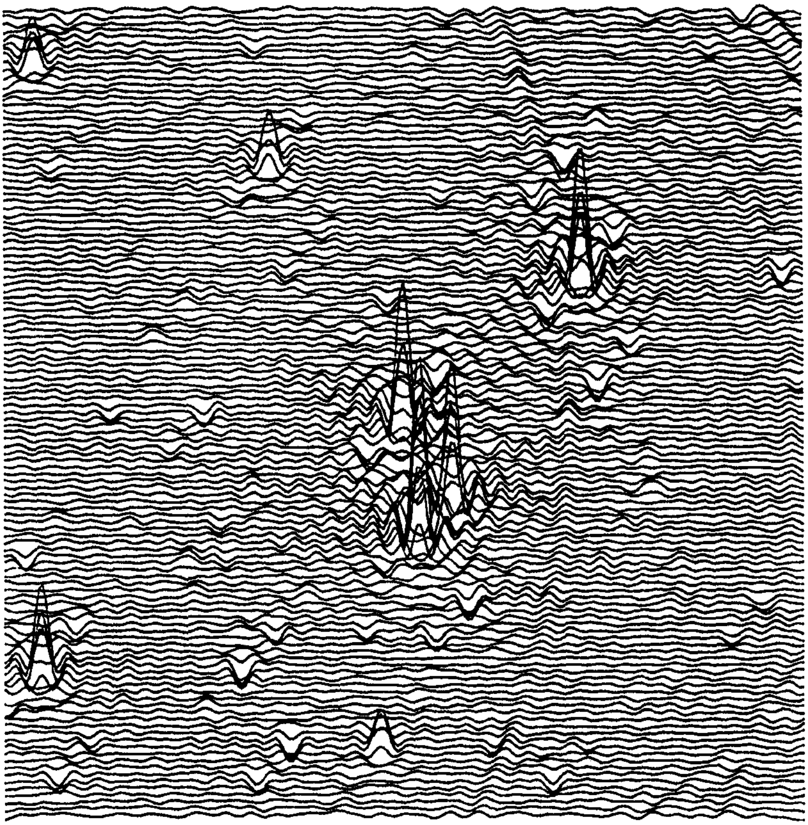

Fig 5. Hybrid map made from constant amplitude and phases appropriate to the distribution of the 10 sources in Fig 4.

4. EXTENDED SOURCES

Can this method be applied to extended sources? Where the extended structure accounts for less than about half the total flux density there seem to be no problems in its application as the discussion of 3C 123 demonstrates. But in cases where no trial map can be made from point sources the problem is much more difficult. An essential feature of the analysis of an A^2 map to give a trial distribution is the recognition of three vectors which together make up a closed triangle in the true distribution, a process which is simple only when there are discrete peaks in the A^2 map. For extended sources the A^2 map has a continuous distribution of brightness and there is no way of picking out particular sets of three vectors. The trial distribution is likely to be a much poorer first guess. Furthermore there can be substantial overlapping of correct and spurious images in the hybrid map which makes interpretation and advancement of the iteration difficult.

In this case progress can be made on the assumption that large areas of the map are zero unless there is definite evidence to the contrary. The limits of the area can be obtained from a study of the A^2 map. The idea of restricting the extent of a source to a 'window' is related to the use of a 'window' in the CLEAN procedure (Schwarz 1978) though in that case the phases (or beam shape) are known. In our case it is the phases which are unknown. This type of window has been used in the iterative procedure of Readhead & Wilkinson (1978) in which they also incorporate phase closure measurements. Successive iterations differ because the decomposition of the map into point sources by CLEAN is restricted to the region of the window. It has not been clear to us exactly what function the window serves in connection with the cleaning procedure or how important it is compared with the closure phase measurements in guiding successive iterations towards the correct distribution of brightness.

At first sight it is surprising that setting a region surrounding the source to zero can exert any useful influence. A study of the measurements necessary to distinguish an asymmetric pair of point sources from a symmetric triple source using amplitude data alone, helps to clarify the situation and is of interest in its own right for sources such as 3C 345 whose structure is a matter of argument. A triple source, comprising a central object of amplitude A and two outer sources of amplitude a/2 at separations of θ and -θ from the central source, gives an amplitude as a function of spacing d(wavelengths) of

$$A + a\cos(2\pi d\theta)$$

whereas a double source with amplitudes A and a and separation θ gives an amplitude

$$(A^2+a^2)^{\frac{1}{2}}\left\{1 + \frac{aA}{(A^2+a^2)}\cos(2\pi d\theta) - \frac{1}{2}\frac{(aA)^2}{(A^2+a^2)^2}\cos^2(2\pi d\theta) + \ldots\right\}$$

which contains higher harmonics of the terms $\cos(2\pi d\theta)$. It is these terms which distinguish symmetric and asymmetric models. They give rise to the sharp features in the amplitude - spacing curves noticeable where the amplitude falls towards zero. The presence of these high harmonics in the amplitude implies that the sampling of the aperture plane necessary to specify the amplitude completely, for a source whose overall dimension is θ, is much finer than the normal information theory interval of $^1/_{2\theta}$. If the aperture plane is sampled at intervals of $^1/_{2\theta}$ it is impossible to distinguish the double and the triple models from amplitude measurements alone. If samples of A are taken at intervals of $^1/_{4\theta}$ then the fourier transform, taking an initial guess for the phase distribution of $\phi = 0$, correctly represents the triple source whereas the double source appears as a similar triple with extra responses at distances of 2θ, 3θ etc. from the central peak. It is at this point that knowledge of the maximum extent of the source becomes

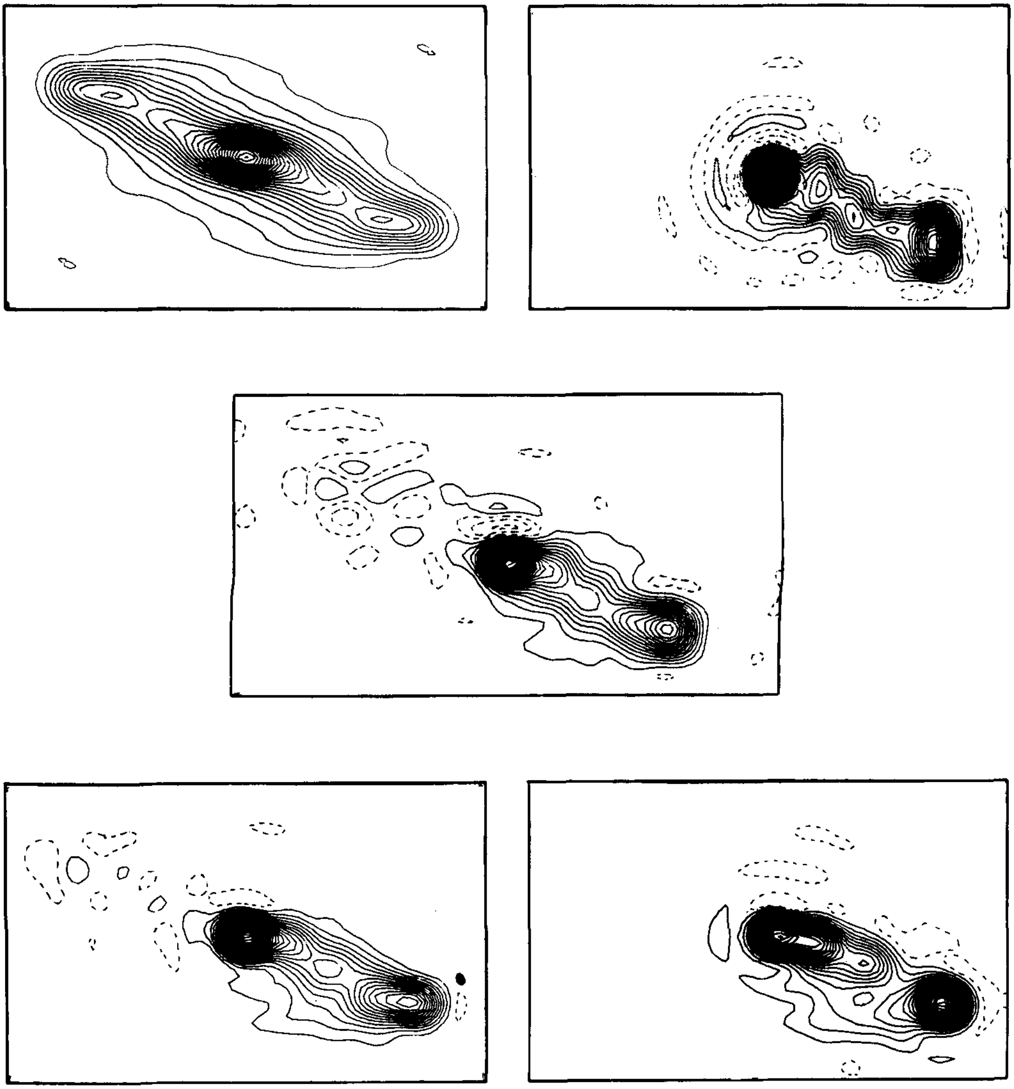

Fig 6. An attempt to derive the structure of 3C 382 from phaseless data using a window technique.

(a) Map of A^2, the autocorrelation of the true distribution.

(b) The initial trial map deduced from (a).

(c) The first hybrid map obtained from observed amplitudes and phases from (b).

(d) The hybrid map after 4 iterations. The window used to define the model for the phase calculation corresponds approximately to the first contour of the source.

(e) The map obtained from observed amplitudes and observed phases.

https://doi.org/10.1017/S0252921100074753 Published online by Cambridge University Press

valuable and enables one to reject the multiple interpretation.

We must now enquire whether this information is powerful enough to provide an iterative scheme which will converge to the correct distribution of intensity. As an example we consider observations of 3C 382 made with the One Mile Telescope at 1407 MHz (S.F. Burch private communication). Fig 6a shows the map of A^2: the extent of the contours is a direct measure of the maximum vector separations which exist between regions of emission in the true map. An accurate trial distribution is hard to deduce from the map of A^2 but it is clear that it must have a predominant axis. The rather uniform intensity away from the peak at the origin and along the axis indicates that the true distribution must have peaks at each end which are joined by some form of bridge. There is no evidence whether the peaks are symmetric or not. An arbitrary ratio of 2:1 in their flux densities was adopted. This initial trial distribution is shown in Fig 6b. The hybrid map, made from the observed amplitudes and the trial phases is shown in Fig 6c. Comparison with the true map (Fig 6e) demonstrates both the improvement over the initial trial and the existence of spurious regions well away from the source. For the next iteration the main body of the source was represented by a set of point sources, following the practice with CLEAN. After 4 cycles the resultant map (Fig 6d) was changing very slowly. In some respects it closely resembles the true map: the shape of the bridge, the faint extensions to the SE and the sharp cut off to the N are all reproduced. The intensity ratio of the two main peaks and their shapes are less well reproduced: the valley running across the bridge does not appear.

Can we distinguish whether the result is an alternative solution to the correct one or whether the speed of convergence is so slow that very many iterations are needed to reach it? A simpler case than 3C 382 is more easily examined. Consider a source distribution composed of four point sources. Suppose that the trial solution exists comprising the three outer sources in their correct positions. The overall extent of the distribution is easily obtained by examination of the map of A^2. The hybrid map (Fig 7a) shows 6 positive peaks inside that overall extent (the 3 trial sources and the peaks A, B, C). These are returned as the next trial distribution for calculating the phases of the second hybrid map. The convergence towards the correct solution is much slower than that discussed for 5C7 since spurious peaks are incorporated into the trial in successive iterations. The normal factor of 2 applied to features in the difference map cannot be applied since the procedure then diverges in this case. The intensities of the correct peaks and the 2 spurious ones are shown as a function of the number of iterations in Fig 7b. Roughly 20 iterations are required to reach a good solution. We do not know how much improved our bybrid map of 3C 382 would be after so many iterations.

5. CONCLUSIONS

Maps from aperture synthesis measurements are usually limited by the errors in the phases. We have investigated methods for improving

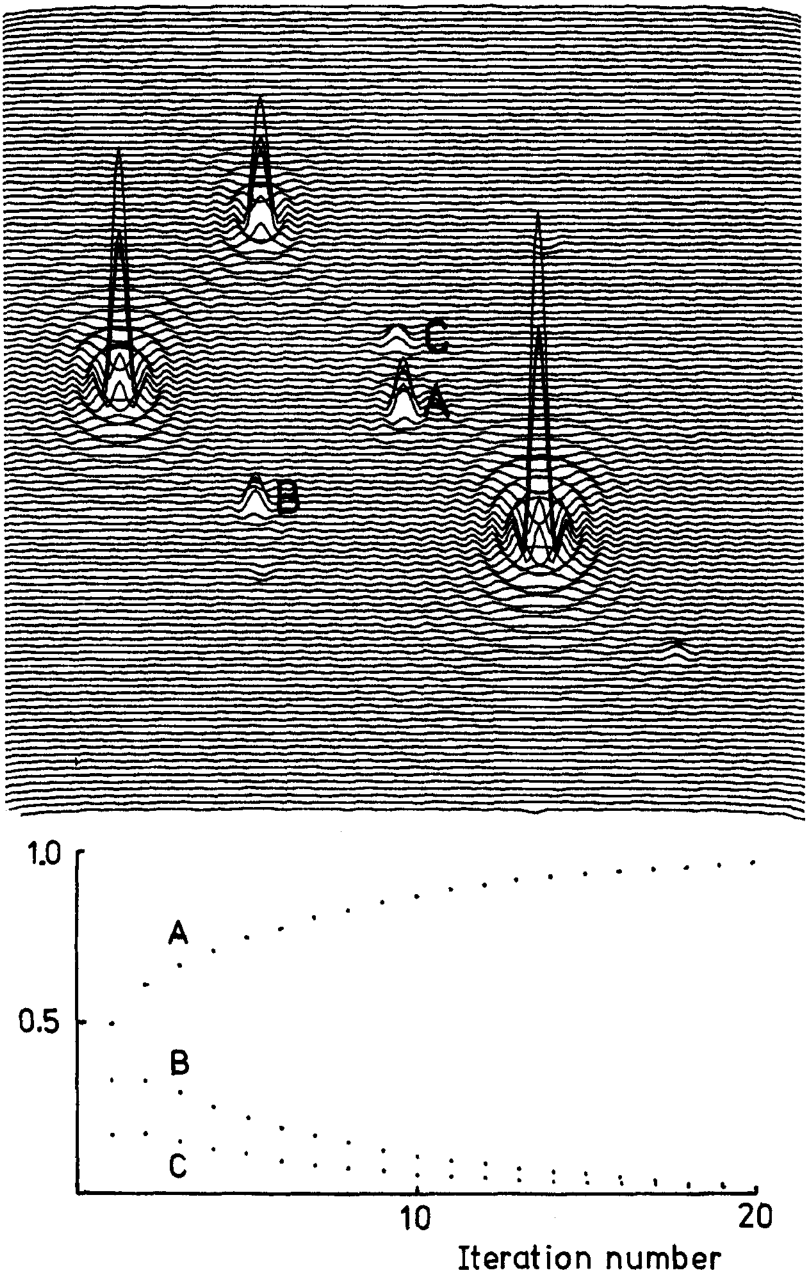

Fig 7. The use of a window technique on a field of 4 point sources.
(a) Hybrid map (1st iteration) made from amplitudes (3 bright sources and source A), together with phases (3 bright sources).
(b) Changes in the correct peak A and spurious peaks B and C in the hybrid map with successive iterations. The intensity scale is normalized to the correct value for peak A.

https://doi.org/10.1017/S0252921100074753 Published online by Cambridge University Press

their quality by using the observed phases, if any, only to provide an initial model of the intensity distribution and thereafter relying solely on the amplitude measurements. The use of hybrid maps provides an easily understood basis for a scheme of iterative improvement and analysis of their properties leads us to believe:

(1) The reconstruction of fields of point sources can be effected to a high degree of accuracy even in the absence of phase information.
(2) The main features of the structure of extended sources can be derived provided that the overall extent of the source is defined and that the amplitudes in the aperture plane have been sampled adequately.

For extended sources the following questions need to be answered:

(1) What characterises a sufficiently good initial model?
(2) How many acceptable alternative solutions exist and how similar are they?
(3) Can the 'window' iteration be improved on?

Although the problems tackled here have arisen in a purely radio-astronomical context, it is evident that the results can be carried over to other fields where the phases are poor or unmeasured such as in speckle interferometry.

REFERENCES

Baldwin, J.E. & Warner, P.J. 1976. Mon. Not. R. astr. Soc., 175, pp 345-354.
Baldwin, J.E. & Warner, P.J. 1978. Mon. Not. R. astr. Soc., 182, pp 411-422.
Fort, D.N. & Yee, H.K.C. 1976. Astron. Astrophys., 50, pp 19-22.
Pearson, T.J. & Kus, A.J. 1978. Mon. Not. R. astr. Soc., 182, pp 273-274.
Readhead, A.C.S. & Wilkinson, P.N. 1978. Astrophys. J., 223, pp 25-36.
Riley, J.M. & Pooley, G.G. 1978. Mon. Not. R. astr. Soc., 183, pp 245-256.
Schwarz, U.J. 1978. Astron. Astrophys., 65, pp 345-356.

DISCUSSION (see also discussion after paper by Wilkinson - Readhead.)

Comment R.H. HARTEN
If you choose a "wrong" component early in the iteration scheme, i.e. a wrong source position, will the method converge and is it self correcting or does it require a critical eye to spot it?
Reply J.E. BALDWIN
If an incorrect component is put into the trial at an early stage it

can cause serious problems, as was shown by the constant amplitude case I described. At a later stage an incorrect component merely causes some oscillations which slow down, but do not prevent, convergence.

Comment J.J. WITTELS
1) It is necessary to be very careful in choosing the box in which the source is considered to be confined, especially early in the modelling, and especially if an FFT algorithm is used.
2) Computer size constraints create, or can create, a difficulty in choosing the correct grid spacing and a large enough box size to adequately model both the long- and short-baseline data simultaneously.
3) Is it not preferable to use closure phase with a connected interferometer rather than to throw out the noisy phase data completely? Have you tried it?
Reply J.E. BALDWIN
1,2) A good, i.e. conservative, method of defining the windows is to look at the A^2 map. That map gives a correct measure of the maximum vector separations in the source and hence of the area which wholly contains the source.
3) The main phase problems in connected interferometers are associated with uniform phase gradients covering the whole telescope array. The closure phases provide no information about the presence or absence of such gradients. We tend to think of this situation as one in which there is one unknown, not in terms of phase closure relations. Otherwise I agree with what you say. We did not try it in the cases I have just described.

Comment T.W. COLE
Did the talk misrepresent the relevance of "information theory"? It gives one a direct measure of redundancies, the need for more measurements if one throws out phases, and how windowing reduces the amount of independent information in the spatial frequency plane.
Reply J.E. BALDWIN
I hope it didn't! I agree that more amplitude information is clearly needed if we have no phase information, but it is much more than a factor of 2. For instance an equal double source has an amplitude which drops to zero at some spacings and the cusps there require indefinitely fine sampling in the aperture plane to specify then completely.

Comment R. GORDON
You noted that the error in phase can be estimated as a function of wavelength and spacing of antennae. It seems that this *a priori* information could be useful in constructing hybrid maps. During each iteration each phase could be constrained to remain within the error range of its measurement.
Reply J.E. BALDWIN
My impression is that if the observed phases are sufficiently good to provide a map which is the basis of the initial trial, then subsequent interations are unlikely to depart significantly from those input phases. The point is that the amplitude information is a much less powerful constraint than the phases. But I agree that your suggestion might help if only to a small extent.