THE ASTROPHYSICAL JOURNAL, 223:25–36, 1978 July 1

THE MAPPING OF COMPACT RADIO SOURCES FROM VLBI DATA

A. C. S. READHEAD AND P. N. WILKINSON
Owens Valley Radio Observatory, California Institute of Technology
Received 1977 November 23; accepted 1978 January 16

ABSTRACT

We describe a method for recovering the brightness distributions of compact radio sources from VLBI observations of closure phase and fringe amplitude. Our approach is to deduce the visibility phase from these data. We show that reliable "hybrid" maps of complexity comparable with that of early synthesis results on extended sources can be obtained with networks of four or more existing telescopes.

Subject headings: instruments — interferometry — radio sources: general

I. INTRODUCTION

The determination of the structure of extragalactic radio sources on a scale 10^{-4} to 10^{-1} arcsec is crucial to our understanding of the physical processes which produce and sustain these objects. Such resolution is obtainable only in very long baseline interferometry (VLBI), and this technique has provided valuable clues to the nature and behavior of the energy sources in the nuclei of galaxies. However, the astrophysical knowledge which has been gained thus far has been limited, not only because of the logistic difficulty of making and analyzing the observations, but also by uncertainties in the interpretation of interferometric data when only a small range of baselines are used with no phase information. If VLBI is to achieve its full potential, it is necessary to demonstrate the feasibility of making reliable maps of compact radio sources and to delineate those areas in which uncertainties arise. These are the purposes of this paper. We describe a new technique for reconstructing brightness distributions which makes use of closure phase (see § II) and the extent of the (u, v) coverage in addition to the fringe amplitudes.

Our approach, described in § III, is to infer the visibility phase from the observed fringe amplitude and closure phase, and we demonstrate in § IV that this can be done provided there is sufficient (u, v) coverage. We can never obtain all of the complex visibility data or completely sample the (u, v)-plane, and the complex visibility function can vary in an unsampled region of the (u, v)-plane in a manner which cannot be deduced from the data. That this is not the case in all unsampled regions results (i) from the requirement that the source brightness be greater than or equal to zero at all points on the sky, and (ii) from a knowledge of the total extent of the source. Multistation VLBI observations are now made routinely, in many cases with (u, v) coverage which would be sufficient for mapping the sources if the interferometers were phase-stable. We therefore assume throughout this paper that the (u, v) coverage is adequate for recovering the brightness distribution from the amplitude and phase data which would be obtained with phase-stable interferometers. We use the CLEAN restoration technique (Högbom 1974; Schwarz 1977*a*, *b*) and in § V we discuss the limitation which this imposes on our method.

We have investigated the reliability of our procedure by performing a series of blind tests on simulated VLBI data. As a result of this study we have become convinced that one can deduce radio structure accurately even if the individual interferometers in an array are not phase-stable. The closure phase disposes of the 180° position-angle ambiguity which is present when amplitude data alone are used, but a basic limitation is that the position of the source is still not recoverable. The claim that we are actually "mapping" sources needs some justification because we do not measure the visibility phase directly but only the closure phase. Our justification is that, although we do indeed begin with models of the sources, the closure phases of the final maps agree, to within the noise, with the observations. This method has proved foolproof in all tests thus far. A more detailed discussion of this point will be given elsewhere. Following Baldwin and Warner (1977), we will refer to maps constructed from incomplete phase data as "hybrid" maps.

II. THE CLOSURE PHASE AND ITS USE

The technique of deriving visibility phase information from interferometry by summing the observed phases around closed loops of interferometer baselines was first suggested by Jennison (1958). Consider two telescopes denoted by i and j. The visibility phase, ψ_{ij}, on this baseline is related to the observed phase, ϕ_{ij}, by the equation

$$\phi_{ij} = \psi_{ij} + \theta_{ij}, \quad (1)$$

where θ_{ij} is the sum of phase perturbations due to propagation effects along the line of sight, oscillator drifts, and uncertainties in the source position and the baseline. When ϕ_{ij} is summed around a closed loop of three telescopes i, j, and k, thus forming the closure phase C_{ijk}, the θ cancel almost exactly; thus

$$C_{ijk} = \phi_{ij} + \phi_{jk} - \phi_{ik} \approx \psi_{ij} + \psi_{jk} - \psi_{ik}\,. \quad (2)$$

The cancellation of the θ_{ij} has been discussed by Jennison (1958), and its application to VLBI by Rogers *et al.* (1974). Rogers *et al.* express equation (2) as an equality. We note, in passing, that this is not exact if the rate errors of the clocks at each station are nonzero. However, since the rate errors are typically less than ~ 1 part in 10^{11}, the phase error introduced by assuming an equality in equation (2) may be neglected. When equation (2) is used the closure phase is formed at a given time, and allowance must be made for the delay between the arrival of the wave front at the different antennas. The explicit form of this correction depends on how the delay is implemented in a given VLBI processor (e.g., Clark 1973).

Visibility phase information can be obtained via the closure phase for any source which produces detectable fringes on an interferometer array of at least three telescopes. Each closure phase contains as much information about the shape of the source as each visibility amplitude and should therefore be included when attempting to reconstruct the brightness distribution. This can be done in a variety of ways; the simplest is to incorporate the closure phase into standard model-fitting procedures. The success of Purcell's model of 3C 147 (Wilkinson *et al.* 1977), derived from amplitudes alone, leaves us in little doubt that this would work reliably even for relatively complex sources. However, for such sources the model-fitting process can be long and tedious, and therefore quicker, more automated methods of arriving at a final answer are desirable. Wittels *et al.* (1976) have used truncated two-dimensional Fourier series, constrained to fit the amplitude and closure phase data, to approximate the source brightness distribution. The method works satisfactorily for simple sources, but no examples of its use on complex sources have yet been presented. Fort and Yee (1976) use the observed amplitudes with phases calculated from an initial source model. The closure phase is used to correct the estimated visibility phase, and a hybrid map is produced. All negative regions are set to zero, and this map provides a new model. The process is repeated until a stable solution is reached. In the relatively simple examples which they show, the maps converge to give a stable solution which is a good representation of the true brightness distribution.

Our approach is similar to this. We also attempt to derive the visibility phases on all baselines in order to produce a hybrid map by the Fourier transformation of complex visibility data. Where we differ from Fort and Yee is in a more direct use of the observed closure phases at each iteration and in our use of the known position of sidelobes in successive maps. Results have already been published (Wilkinson *et al.* 1977) for 3C 147. The method works particularly well for this source because it contains a barely resolved component which can be used as a phase reference. We shall show in § IV that more general brightness distributions can also be reconstructed by using our method.

III. RECOVERING THE VISIBILITY PHASE FROM THE CLOSURE PHASE

Consider interferometer observations with N telescopes. There are $N(N-1)/2$ combinations of telescope pairs and at most $(N-1)(N-2)/2$ independent closure phase relations. There are therefore $N-1$ too few linear combinations of the closure phases to solve for all the visibility phases, and it is necessary to obtain an independent estimate of the phase on $N-1$ baselines. If there were sufficient redundancy and if the visibility phase on at least one baseline were known, all the visibility phases could be derived directly (Jennison 1958). However, this is never the case in present-day VLBI observations.

The procedure we use is outlined in Figure 1. We begin by examining the visibility amplitudes and the closure phases and marking the positions of maximum and minimum amplitudes in the (u, v)-plane. We then attempt to construct a simple model which reproduces the principal maxima and minima over a limited range in the (u, v)-plane. In practice a crude fit to the data on a few of the baselines is usually adequate, and more detailed modelling procedures were not used in the tests discussed in § IV. Note here that we are assuming that if the amplitudes predicted by the model fit the observed amplitudes well, then, apart from the usual ambiguity in orientation of 180°, the phases will also be a reasonable fit to the observed phases. Our experience with Purcell's model of 3C 147 (Wilkinson *et al.* 1977) and the tests described below support this assumption.

To dispose of the 180° ambiguity in the orientation of the source, we calculate the closure phases of the model for one orientation and compare this with the observations. The alternative orientation simply produces closure phases with the opposite sign; it is therefore easy to see which of the two possibilities is correct, and this is chosen as the initial model. We set the visibility phase equal to the model phase on $N-1$ baselines and use the $(N-1)(N-2)/2$ closure phase relations to deduce phases on the remaining baselines. We then have enough amplitudes and phases to make the complex Fourier transformation, as follows.

The *observed* fringe amplitudes and the *derived* visibility phases are convolved onto a regular grid in (u, v) space by using a Gaussian convolution function. The point at the origin is assigned the fringe amplitude appropriate to the total flux of the source and convolved by the same function. Zero values are assigned to grid points not filled by this procedure. This array is then transformed into the sky plane, resulting in the so-called dirty map of the source which is disturbed by the errors in the derived phases and by sidelobes due to the incomplete (u, v) coverage. The next step takes account of our knowledge of the (u, v) coverage by using the CLEAN procedure to reduce the amplitude of these sidelobes. We choose the area of the dirty map to be cleaned on the basis of our knowledge of

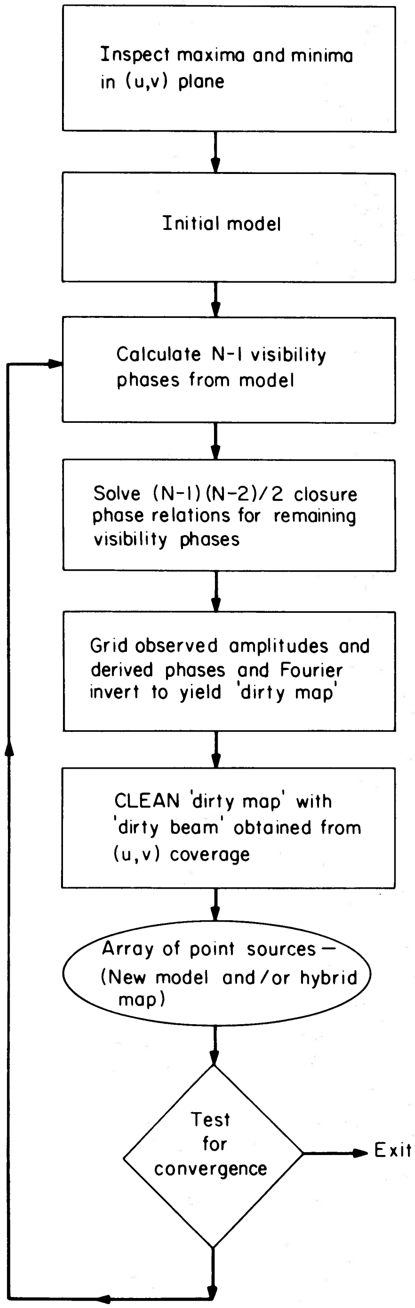

FIG. 1.—Flow diagram of the image reconstruction procedure.

the angular extent of the source from other observations and from direct inspection of the visibility amplitudes. In our experience any gross errors in the original model, and hence in the calculated phases, become apparent at this stage through the appearance of large negative regions in the CLEAN map; if this happens it is usually better to make a new start with a different model. Conversely, we find that if the negative areas are small, the model is a fair approximation to the true brightness distribution, and we return to the first step and replace the original model with the array of point sources produced by CLEAN. We then repeat the whole procedure and continue to iterate until a stable or near-stable solution is reached. For simple sources we find that the convergence is rapid, and only three or four iterations are needed to reach a stable map. For more complex sources, or for poorer (u, v) coverage, we find that, after an initial rapid convergence in the first three or four iterations, the map does not stabilize completely but continues to vary slightly between one iteration and the next. Thus it is possible to produce a number of very similar maps which have low sidelobe levels and fit the data reasonably well, but which differ from each other in small details.

Why does convergence occur at all? In the early stages the brightness distributions used to calculate the phases at each iteration do not fit the data perfectly. These can therefore be regarded as the sum of the true brightness distribution, which is everywhere positive, and an error distribution, which can have positive and negative regions. The error distribution produces errors in the derived phases. When the derived phases are coupled with the observed amplitudes and Fourier transformed, the phase errors tend to scatter brightness components over a larger area of sky than is occupied by the real source (see Fig. 3*a*). A similar effect has been noted by Walther (1963). If we cleaned over the total area occupied by both real and spurious components, we would produce an array of point sources which, when Fourier transformed, would fit the observed amplitudes and the derived phases very closely. This would be so regardless of what the phases were, since the point source array is produced by a linear process. In this case the hybrid map would not change between iterations. However, in accordance with the original basis of CLEAN, we suspect that most of the sky is empty and therefore restrict our search to a small area, or "window," which we believe includes the true source. We also stop the source subtraction in intermediate iterations when the flux density of the last point source is still well above the noise of the residual map. As a further filter against nonphysical brightness distributions, we do not return any negative point source if its amplitude is at least 10% of the strongest point source. Note that small negative point sources should be admitted in CLEAN if the "loop gain" (Högbom 1974) is not infinitesimal. Thus, since we return only those plausible components of the brightness distribution which lie within the chosen window at each iteration, we systematically erase the error distribution. The calculated phases therefore tend toward the correct ones (e.g., Fig. 3*b*) and the map converges toward the true distribution. In general, of course, some parts of the error distribution always lie within the window and may not be rejected by this process. Whether or not these errors are completely erased depends on the (u, v) coverage, and we have found that the speed of convergence depends on the number of independent picture elements in the window compared with the number of independent observed data points. For simple sources the error distribution is reduced to an insignificant level, and convergence is fast. For large complex sources the amount of rejection per iteration is smaller, convergence is slower, and the final hybrid maps may still contain significant errors.

Since the most subjective step in this procedure is the choice of the initial source model, it is advisable to begin with more than one model and to check that they converge to similar hybrid maps. As we show in § IV, the starting model is not critical for simple sources, but for more complex sources different starting models can lead to hybrid maps which differ in small details. It is always possible to reduce these discrepancies by increasing the amount of (u, v) coverage, but, although this is the most desirable way of resolving minor uncertainties, in practice it may not be easy to do. To make the technique more useful, we have therefore examined three ways of discriminating among the different hybrid maps and thereby picking the best solution.

a) Compare the rms noise left on the residual dirty maps after the subtraction of the point sources; in the initial rapid convergence it decreases by about a factor of 10 to a value close to that estimated from the observed scatter on the fringe amplitudes and closure phases. We have found that the better maps have the least noise.

b) Compare the data with the amplitudes and closure phases predicted by the final array of point sources.

c) Use the array of point sources to calculate phases on a *different* set of $N - 1$ baselines, and hence make a new hybrid map. Comparing this with the old map provides a simple test, in the map domain, of whether all the phases predicted by the array of point sources are consistent with the observed closure phases.

Are these tests sufficient to ensure that the derived phases are the same, to within the noise level, as the true visibility phases? We will first address the problem, well known in physics, of phase recovery from the intensity distribution alone. In the one-dimensional case some progress has been made toward deriving the phases from the amplitudes (e.g., Walther 1963; Bates 1969). In general there is not a unique solution, but for Fourier transforms of band-limited functions, such as we are considering here, there are a finite number of solutions that need be considered (Burge *et al.* 1976). Many of these solutions can be rejected because they give rise to negative brightness regions. No complete theory has been developed for the two-dimensional case, but the requirement that the brightness distribution be positive appears to be more restrictive in two dimensions than in one. In a series of tests Napier and Bates (1974) found that for two-dimensional brightness distributions only one solution gave a wholly positive brightness distribution in each case. Similarly, recent tests on real data by Baldwin and Warner (1977) support the assumption, implicit in their work, that the correct visibility phases (barring the usual 180° ambiguity) can be deduced from amplitude data alone.

In the present case we are not limited to amplitude data since we also use the closure phases. These place very strong constraints on the visibility phases, and our experience with the blind tests, described in § IV, suggests that the information contained in the visibility amplitudes and closure phases is sufficient to determine correctly the visibility phases. Since we assume that there is sufficient (u, v) coverage for CLEAN to produce a reliable map (see § V), we believe that the hybrid map will be the same as the true map if the above tests are satisfied. This problem will be discussed in more detail elsewhere.

To summarize, the chosen hybrid map satisfies the following criteria: (i) the brightness distribution is everywhere positive or zero within the noise level; (ii) the visibility amplitudes derived from the map fit the observed data to within the noise level; (iii) the closure phases derived from the map fit the observed closure phases to within the noise level.

IV. TESTS OF THE IMAGE-RECONSTRUCTION PROCEDURE

A rigorous theory of what constitutes sufficient (u, v) coverage for this method to work must incorporate the signal-to-noise ratio at all sampled points in the (u, v)-plane and also allow for the unequal spatial frequency sampling in two dimensions which always obtains in VLBI observations. Rather than attempt to develop such a theory, we have tested the method by making a series of reconstructions from simulated VLBI data. The (u, v) coverage assumed for tests 1 through 4, shown in Figure 2*a*, is that available for a source at declination 50° on the baselines formed by telescopes at Owens Valley, California; Fort Davis, Texas; Green Bank, West Virginia; and Jodrell Bank, UK. For the assumed observing frequency of 610 MHz ($\lambda \approx 49$ cm) this gives a dirty beam whose principal response is nearly circular and has a FWHM of 10 milli-arcsec; we therefore used a sidelobe-free restoring beam which was a circular Gaussian of the same dimensions. In this case the grid spacing was set at 4 milli-arcsec. For the fifth test the source was taken to be at declination 5°, and the Haystack, Massachusetts, telescope was assumed rather than Jodrell Bank. The (u, v) coverage is shown in Figure 2*b*. Here the maximum resolution is about 18 milli-arcsec, and the grid spacing was 8 milli-arcsec. A computer program calculated the expected fringe amplitudes and closure phases for test sources consisting of Gaussian components. Random noise (5% rms fringe amplitude and 10° rms closure phase) was added to these simulated data; the amplitudes and phase noises were uncorrelated. This is a relatively high signal-to-noise ratio to assume, but such data have been obtained in 3C 147 at 610 MHz (Wilkinson *et al.* 1977).

It is important to stress that the test sources, which were devised by M. H. Cohen, were not revealed to us until we had decided upon a final interpretation of the simulated data. Thus, as in the real situation, at the start of each reconstruction we knew only the calculated fringe amplitudes, the closure phases, and the

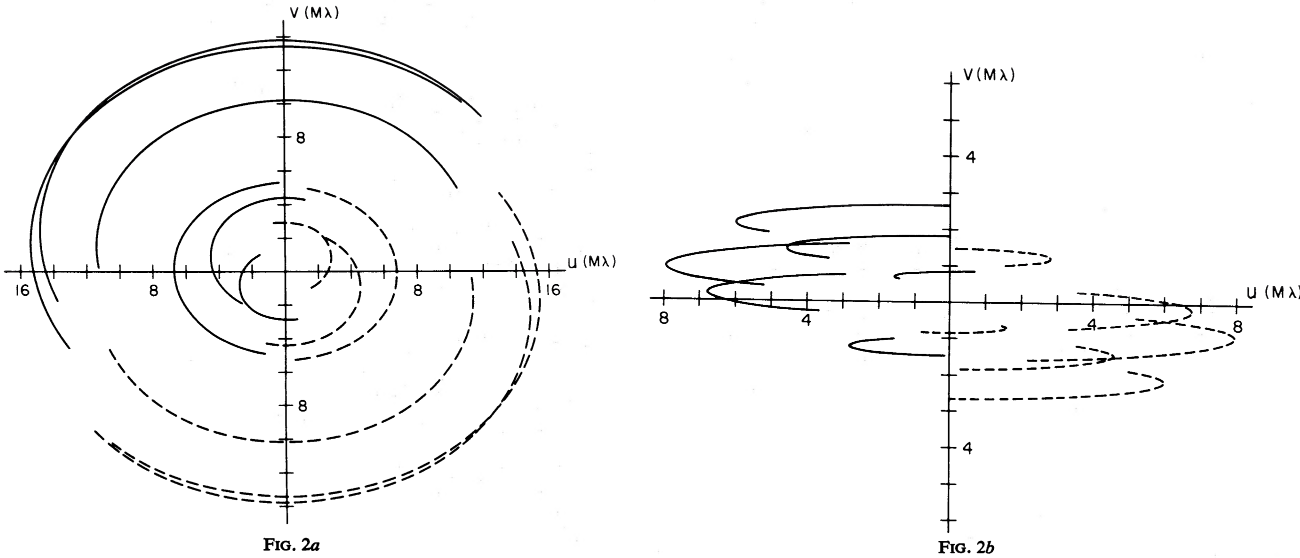

FIG. 2.—(*a*) The (*u*, *v*) coverage obtained at 610 MHz for a source at $\delta = 50°$ using telescopes at Owens Valley, California; Fort Davis, Texas; Green Bank, West Virginia; and Jodrell Bank, UK. (*b*) The (*u*, *v*) coverage at 610 MHz for a source at $\delta = 5°$ using the same telescopes as in (*a*) except that Haystack, Massachusetts replaces Jodrell Bank.

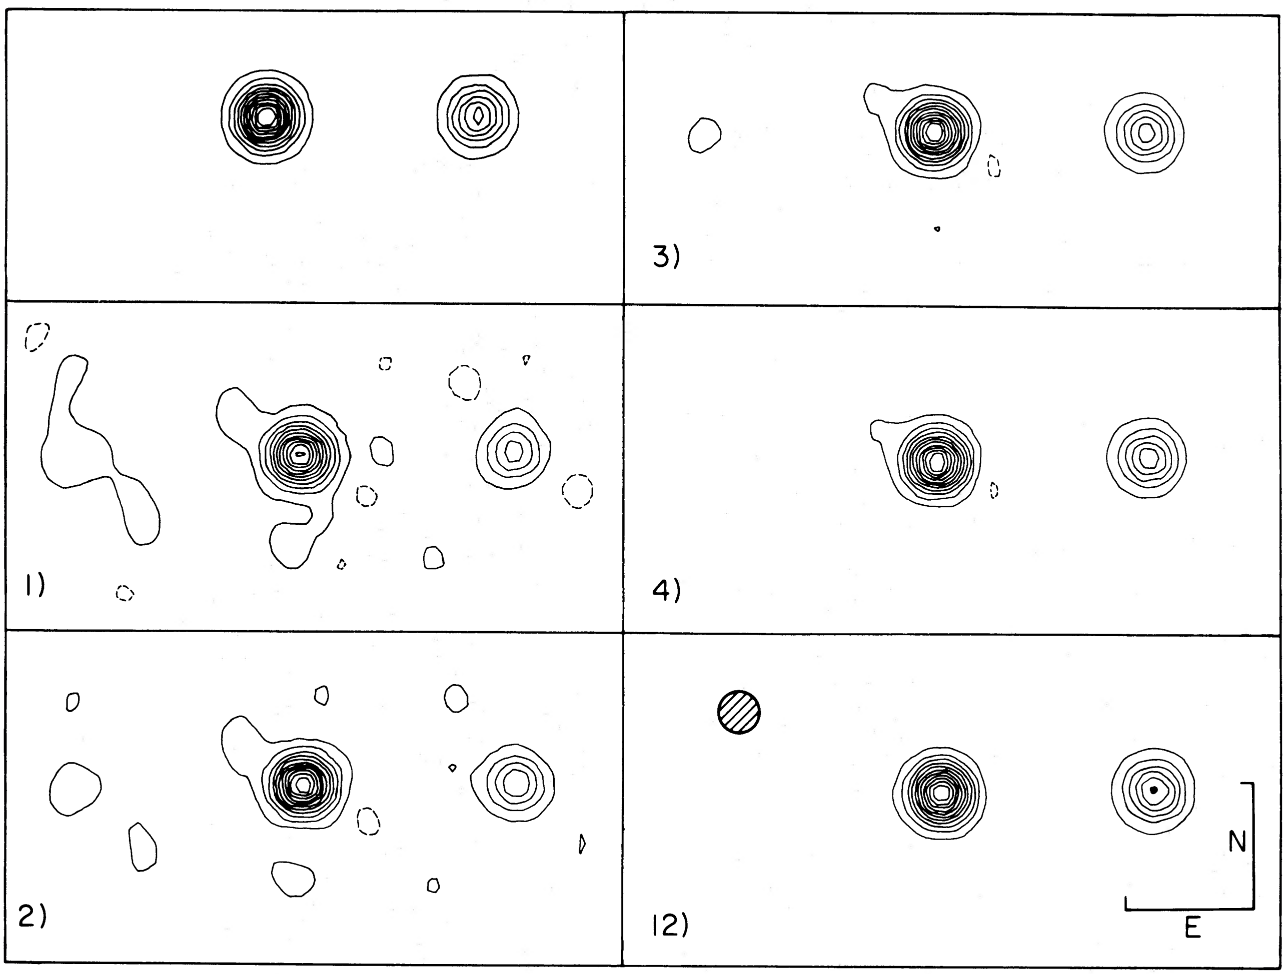

FIG. 3*a*.—Test 1 [(*u*, *v*) coverage Fig. 2*a*]. *Top left*, CLEAN solution assuming full phase data. *Others*, various iterations using closure phase. (Contours −5%, 5%, 10%, 20%, . . . , 90% of the peak brightness). In tests 1–5 the arms of the L each represent 30 milli-arcsec.

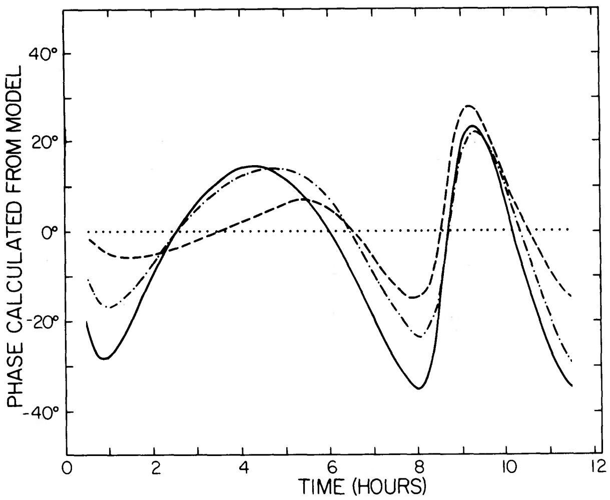

FIG. 3*b*.—A comparison of the phase calculated from the model with the true visibility phase (*solid line*) on the NRAO–Fort Davis baseline. *Dots*, first iteration; *dashes*, 2d iteration; *dash-dot curve*, 4th iteration. The 12th iteration is indistinguishable from the solid line.

total flux density of the source. We did ask that the test sources fall within a square 150 milli-arcsec on a side (i.e., ~15 times the FWHM diameter of the synthesized main beam), but as the calculated amplitudes and closure phases produced by several of these models varied in a more complex manner than anything actually observed to date, we do not regard this as a major restriction.

a) *Test 1. 2:1 Point Double*

For this example only, we had prior knowledge of the test source, and no noise was added to the data. The source was chosen to be very simple so that it would be easy to make a detailed comparison of the maps and derived phases at each iteration and to illustrate the convergence to the true distribution.

At the top left of Figure 3*a* we show, not the test source, but the "CLEAN solution" which comes from assuming full visibility phase and amplitude data, throughout the interval in which the source can be observed at all four stations, and therefore represents the best result we could have obtained given the (u, v) coverage and the effectiveness of the CLEAN process. In the rest of Figure 3*a* we show successive iterations during the reconstruction from amplitude and closure phase data (the particular iteration is indicated at the bottom left-hand corner of each map); note that we cleaned over the whole area of each map. The first iteration is the result of setting all the visibility phases on three of the baselines to zero, equivalent to starting with a point source at the center of the map. This produces a fairly good map, in spite of the fact that the derived phases differ by up to $\pm 60°$ from the true ones. By the fourth iteration the reconstructed map is good, while excellent agreement with the CLEAN solution comes around the seventh iteration. By the twelfth the map is virtually indistinguishable from the CLEAN solution. In Figure 3*b* are shown the phases derived from the model, on the Green Bank–Fort Davis baseline, at three stages of the procedure.

In order to test our ideas about the convergence mechanism we repeated this test, restricting the window to include only the true source. In this case, with the area of the window about 5 times smaller than before, convergence was roughly twice as fast. We also tested the effect of having more visibility data by assuming an extra telescope (Haystack) in the array, thus giving 10 baselines and six closure phases. In this case, again with the smaller window, convergence was roughly 3 times as fast as in Figure 3*a*. In practice, of course, one would never start from a point-source model because the double nature of the source is obvious from a cursory glance at the visibility data. However, this example does show that for simple sources no initial interpretation of the data is required to obtain the correct solution.

We now discuss the results of four "blind" tests. In Figures 4–7, (*a*) shows the test source, and (*b*) shows the CLEAN solution, i.e., the solution which is obtained by assuming full visibility phase and amplitude data with the (u, v) coverage shown in either Figure 2*a* or Figure 2*b*. Subsequent hybrid maps were obtained by using closure phase data. In each case the true distribution (*a*) has not been convolved with the restoring beam used in CLEAN.

b) *Test 2*

Figure 4*c* shows the preferred hybrid map which is a fairly accurate representation of the test source given that its maximum extent is ~15 beam diameters. There are some significant errors, but this was the first source to be tried, and we had therefore not developed all the "tricks of the trade" which came with subsequent experience. In particular we did not use the test of changing the baselines on which the phases were calculated (test [*c*]). We are confident that a result much closer to that of Figure 4*b* would have been achieved had we tackled this source later in the series.

c) *Test 3*

Figure 5*c* shows our peferred solution. The only significant error is that the "leading edge" of the southern component is somewhat broadened. This was a feature of all the maps we made on this source and demonstrates that our method will not converge absolutely when the source covers a large number of picture elements. However, test (*c*) did show that Figure 5*c* still contained significant errors.

d) *Test 4*

Figure 6*c* shows the preferred hybrid map assuming data from six baselines. The main features are well

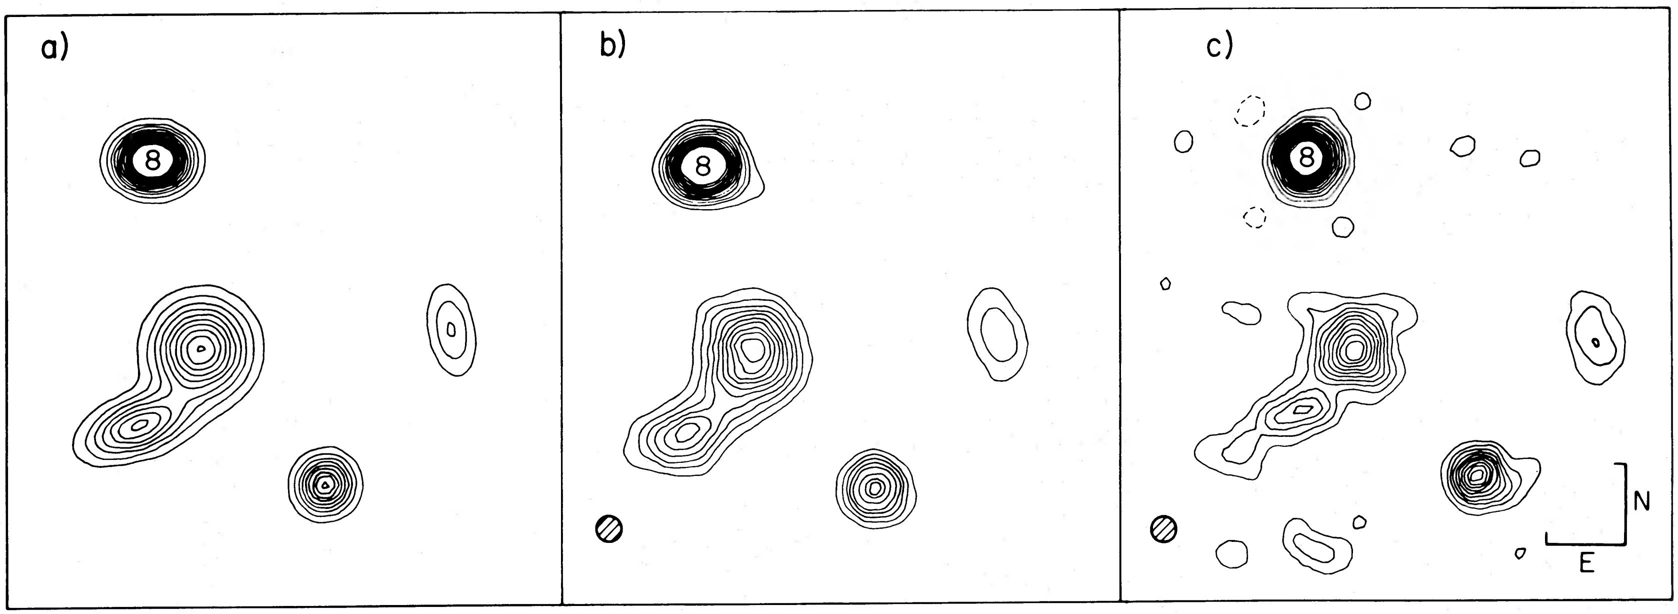

FIG. 4.—Test 2 [(*u*, *v*) coverage, Fig. 2*a*]. (*a*) True distribution. (*b*) CLEAN solution. (*c*) Hybrid map using closure phase (contours -5%, 5%, 10%, . . ., 60% of the peak brightness). Note that in all these tests the true distribution has not been convolved with the restoring beam used in CLEAN. Negative contours are dashed.

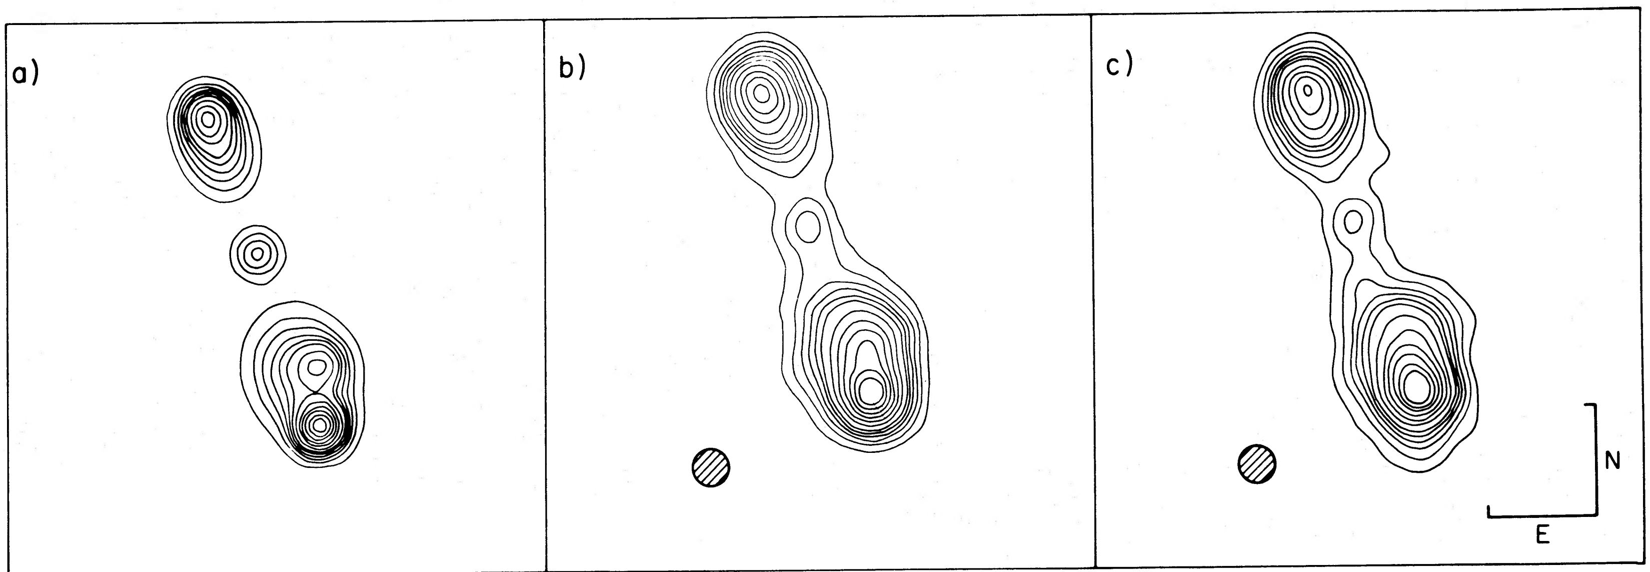

FIG. 5.—Test 3 [(u, v) coverage, Fig. 2*a*]. (*a*) True distribution. (*b*) CLEAN solutions. (*c*) Hybrid map using closure phase. (Contours as in Fig. 3*a*.)

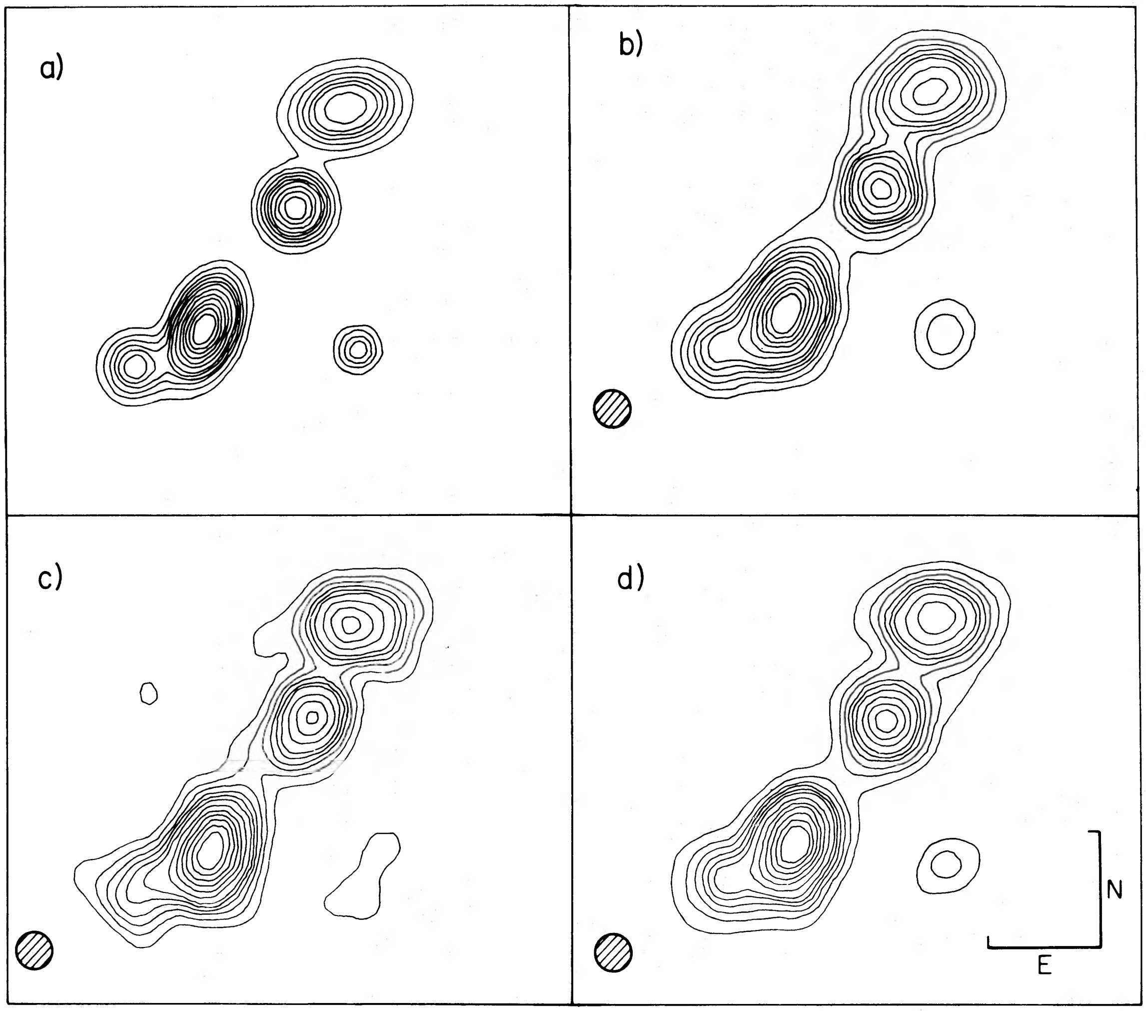

FIG. 6.—Test 4 [(u, v) coverage, Fig. 2*a*]. (*a*) True distribution. (*b*) CLEAN solution. (*c*) Hybrid map using closure phase. (*d*) Hybrid map assuming a further station (Haystack) in the array. (Contours as in Fig. 3*a*.)

reproduced, but the presence of the weak component in the SW was not clearly established. On this source we tried having the grid spacing (to 2 milli-arcsec), but the improvements in the map were barely significant. We again tested the effect of having more visibility data by including the Haystack telescope in the array. Figure 6*d* shows the result of three further iterations; the marked improvement in the quality of the map is clear.

e) Test 5. A Source at a Low Declination

The clean beam is elongated here because of the low declination, which restricts the (u, v) coverage as shown in Figure 1*b*. The preferred map, shown in Figure 7*c*, is satisfyingly similar to the CLEAN solution. Figure 7*d* was obtained from the same point source array as Figure 7*c* but here we have used a circular restoring beam (FWHM = 20 milli-arcsec).

V. LIMITS TO THE PROCESS

We now try to give some practical rules of thumb for the complexity of sources that can reliably be recovered from VLBI observations with four or more telescopes. The fundamental limits to the process we have outlined are clearly set by the performance of CLEAN with full phase data. Since it is difficult to obtain a feeling for these limits from the literature, we have made additional tests assuming the (u, v) coverage of Figure 1*a* and the same signal-to-noise ratio as in the tests discussed in § IV.

Any reconstruction process is basically limited by the size of the largest "holes" in the (u, v)-plane; conventionally the maximum extent of sky which can be synthesized is $1/\Delta u$ by $1/\Delta v$, where Δu and Δv define the size of the hole. When the baselines are irregularly spaced the holes are of varying sizes, and it is difficult to characterize Δu and Δv. However, for an N station network, the largest holes have dimensions

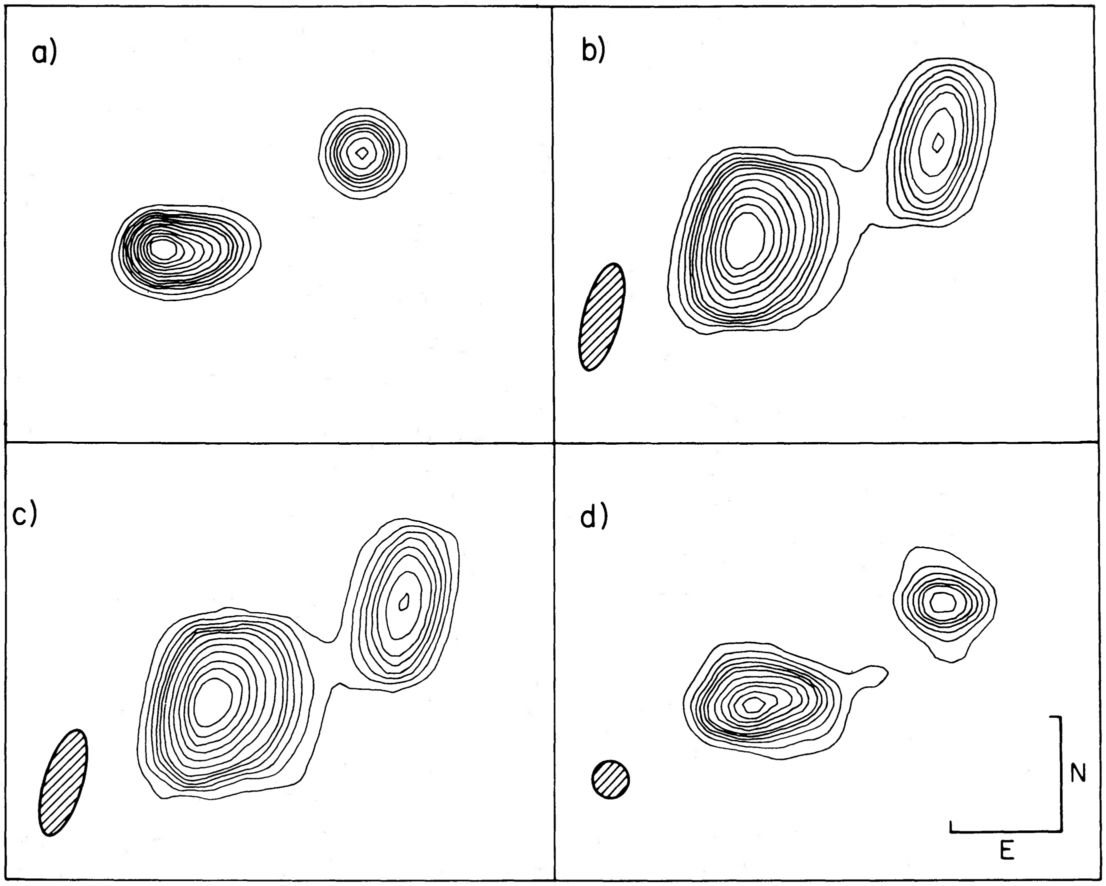

FIG. 7.—Test 5 [(u, v) coverage, Fig. 2*b*]. (*a*) True distribution. (*b*) CLEAN solution. (*c*) Hybrid map using closure phase. (*d*) same as (*c*) except for the shape of the clean beam. (Contours as in Fig. 3*a*.)

$\gtrsim 2/N(N-1)$ of the maximum excursions in the (u, v)-plane. Thus we might expect to be limited to a map with $\lesssim [N(N-1)/2]^2$ picture elements, where a picture element has area (FWHM of beam)^2.

In order to test this limit we constructed the source shown in Figure 8*a*, which covers ~100 picture elements, and attempted to recover it from the full amplitude and phase data assuming the (u, v) coverage shown in Figure 2*a*. The result for this four-station network is shown in Figure 8*c*. It is clear that the coverage is not sufficient, but adding one more station (Haystack) to the network yields the improved map shown in Figure 8*b*. Such tests support the results of Rogstad (private communication) that *arbitrary* brightness distributions of area equal to the conventional limit, discussed in the previous paragraph, can reliably be recovered by CLEAN if the signal-to-noise ratio is high (~20:1). If instead the source consists of a small number of regions of high surface brightness, as in the tests discussed in the last section, considerably larger regions of sky can be cleaned; for example, the CLEAN map in test 2 covers ~260 beam areas. Many extragalactic sources are also considerably more extended in one dimension than the other, and in this case sources resembling 3C 147 (Wilkinson *et al.* 1977) can be reconstructed out to ~25 beam diameters (i.e., ~75 beam areas) with full amplitude and phase data from only six baselines. With more baselines the limiting areas increase in proportion to the amount of visibility data available.

A practical limitation of any image reconstruction process is the time required to reach an acceptable solution. In the simulations of § IV we allowed ourselves only ~2 days, and ~20 jobs, on a batch processing system. Thus our method is considerably faster than conventional model fitting. However, there are probably some compact radio sources which are difficult to model on $N-1$ baselines, but which could, nevertheless, be mapped with full amplitude and phase information on $N(N-1)/2$ baselines (e.g., Fig. 8*a*). In such cases more time is required to derive an adequate starting model, but it should be emphasized that this map is recoverable from amplitude and closure phase data, since the number of unknown phase data is finite, and incorrect hybrid maps can be eliminated by means of the tests discussed in § III.

We have considered mainly the case of four telescopes, the minimum for our process to be useful.

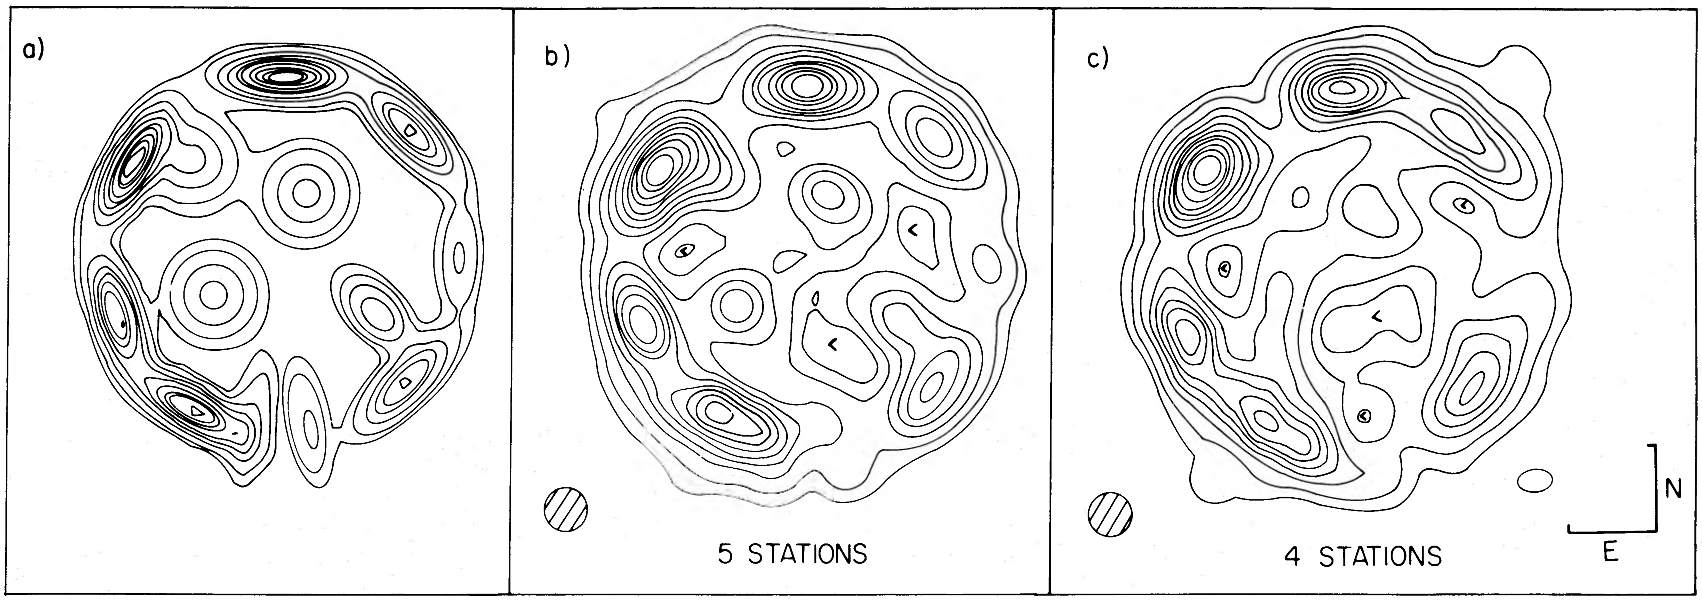

FIG. 8.—Test of capability of CLEAN on complex brightness distribution. (*a*) True distribution. (*b*) CLEAN solution with (u, v) coverage of Fig. 1*a* + Haystack. (*c*) CLEAN solution with (u, v) coverage of Fig. 1*a*. (Contours as in Fig. 3*a*, arms of the L represent 20 milli-arcsec.)

With more telescopes the ratio of the number of independent closure phase relations to the number of baselines increases. Thus the fraction of missing phase information decreases, and one may expect to see a corresponding increase in the effectiveness of the process.

How could our method be improved? First, there may be more effective ways of deriving phases from the amplitude data on $N - 1$ baselines (e.g., Napier and Bates 1974). Second, one might use a method of image reconstruction other than CLEAN. A fundamental criticism of CLEAN is that, by setting the visibility amplitude equal to zero for unsampled regions of the (u, v)-plane, one is supplying data which are known to be incorrect. The maximum entropy approach (Ables 1974; Ponsonby 1973; Gull and Daniell 1977) is more sensible in this respect, since it attempts to introduce the minimum amount of information and hence a minimum of false information in the unsampled areas. Tests have shown (Rogers 1976) that this method does indeed reproduce source structures more reliably than CLEAN but at the expense of considerably more computer time.

VI. CONCLUSIONS

We have shown that reliable hybrid maps can be made by using data which are easy to obtain in present day VLBI observations. The solution is derived in a more objective way, and more rapidly, than is the case with conventional model fitting. This is particularly so for sources of complexity similar to those simulated in § IV and for most, if not all, sources actually observed by VLBI thus far, since adequate starting models can easily be constructed on $N - 1$ baselines. VLBI arrays of four to six telescopes are now common, and our simulations show that brightness distributions of complexity comparable with many of the Cambridge One-Mile telescope maps (Macdonald, Kenderdine, and Neville 1968; Mackay 1969) can be obtained. With rms noise of 5% in amplitude and 10° in closure phase, a dynamic range of $\sim$20:1 can be achieved in the map.

At present, the only way of obtaining full phase information from VLBI is by using a nearby unresolved source as a phase reference. In the future such data will become increasingly available, and thereby relative positions can be determined with the full accuracy of VLBI. However, given the sensitivity of present-day VLBI systems, there are too few phase reference sources for this method to be generally useful. Even with improved sensitivity, phase referencing will be difficult at both low and high frequencies, due to irregularities in the ionosphere and atmosphere, respectively. Since the closure phase does not suffer from these disadvantages and is easy to obtain as long as fringe amplitudes can be measured on closed loops of baselines, it will be used increasingly in observations at these frequencies.

We have discussed the use of the closure phase in VLBI. However, this method is generally applicable to any interferometer array which is not phase-stable. It could therefore be used to extend the capabilities of existing aperture synthesis arrays to both lower and higher frequencies.

We thank M. H. Cohen for his continued encouragement and helpful criticisms, and for inventing the source models used in blind tests. We also thank D. H. Rogstad for many useful discussions regarding the use of CLEAN. A. C. S. R. gratefully acknowledges receipt of the Royal Society Weir Fellowship during part of this work. This work was supported in part by National Science Foundation grant MPS 73-04677.

REFERENCES

Ables, J. G. 1974, *Astr. Ap. Suppl.*, **15**, 383.
Baldwin, J. E., and Warner, P. J. 1977, *M.N.R.A.S.*, in press.
Bates, R. H. T. 1969, *M.N.R.A.S.*, **142**, 413–428.
Burge, R. E., Fiddy, M. A., Greenaway, A. H., and Ross, G. 1976, *Proc. Roy. Soc. London, A*, **350**, 191–212.
Clark, B. G. 1973, *Proc. IEEE*, **61**, 1242–1248.
Fort, D. N., and Yee, H. K. C. 1976, *Astr. Ap.*, **50**, 19.
Gull, S. F., and Daniell, G. J. 1977, *Nature*, in press.
Högbom, J. A. 1974, *Astr. Ap. Suppl.*, **15**, 417.
Jennison, R. C. 1958, *M.N.R.A.S.*, **118**, 276.
Macdonald, G. H., Kenderdine, S., and Neville, A. C. 1968, *M.N.R.A.S.*, **138**, 259.
Mackay, C. D. 1969, *M.N.R.A.S.*, **145**, 31.
Napier, P. J., and Bates, R. H. T. 1974, *Astr. Ap. Suppl.*, **15**, 427.
Ponsonby, J. E. B. 1973, *M.N.R.A.S.*, **163**, 369.
Rogers, A. E. E. 1976, in *Methods of Experimental Physics*, vol. **12C**, chap. 5.
Rogers, A. E. E., *et al.* 1974, *Ap. J.*, **193**, 293.
Schwarz, U. 1977*a*, *Astr. Ap.*, **54**, 863.
———. 1977*b*, *Astr. Ap.*, **61**, 91.
Walther, A. 1963, *Optica Acta*, **10**, 41–49.
Wilkinson, P. N., Readhead, A. C. S., Purcell, G. H., and Anderson, B. 1977, *Nature*, **269**, 764–768.
Wittels, J. J., *et al.* 1976, *A.J.*, **81**, 933.

A. C. S. Readhead: Mail Code 102-24, Owens Valley Radio Observatory, California Institution of Technology, Pasadena, CA 91125

P. N. Wilkinson: Nuffield Radio Astronomy Observatory, Jodrell Bank, Macclesfield, Cheshire SK11 9DL, England