Ann. Rev. Astron. Astrophys. 1984. 22: 97–130

IMAGE FORMATION BY SELF-CALIBRATION IN RADIO ASTRONOMY

T. J. Pearson and A. C. S. Readhead

Owens Valley Radio Observatory, California Institute of Technology, Pasadena, California 91125

1. INTRODUCTION

Since its first application to astronomy in the late 1940s (67, 95, 98), radio interferometry has developed into an imaging tool of unparalleled power. The complex visibility measured by a radio interferometer is a single Fourier component of the sky brightness distribution, so a map of the sky can be made by using Fourier synthesis to combine measurements obtained with different interferometer baselines. This technique, called *aperture synthesis* (97), was the basis of a number of instruments designed in the 1950s, in which several fixed antennas or two or more movable antennas were used to provide the necessary range of baselines. In the 1960s and 1970s, several telescopes were built that made use of the Earth's rotation to change the interferometer baselines, notably the One-Mile and Five-Kilometre telescopes at Cambridge in England (96), and the Westerbork Synthesis Radio Telescope (WSRT) in the Netherlands (3, 56). The most powerful of these is the Very Large Array (VLA) of the National Radio Astronomy Observatory of the United States (52, 54, 112). The VLA, with a maximum baseline of 35 km, routinely makes maps that surpass the resolution of the best optical telescopes. Also during the 1960s, baselines were increased to hundreds of kilometers by using radio links to transmit the signal from one antenna to the other (35, 80), and eventually to thousands of kilometers by the technique of Very Long Baseline Interferometry (VLBI), in which the signals are recorded on magnetic tape at each antenna and later played back for cross-correlation (8, 16, 23, 24).

A common problem in aperture synthesis is that the phase of the complex visibility is badly corrupted by instrumental and propagation effects. Over

0066–4146/84/0915–0097$02.00

short baselines, these problems can generally be overcome by careful calibration of the instrument using unresolved sources. For a variety of reasons this solution cannot be applied to VLBI, and thus for some years it was thought that VLBI could never become a true imaging technique. It had been known since the early 1950s that part of the phase information could be recovered from badly corrupted phases when three or more antennas were used simultaneously (59), but methods of applying this information in the construction of images were not explored until the 1970s. The breakthrough came with the realization that this partial phase information could be used, together with the constraint that the source brightness distribution must be positive, to recover the full phase information and hence to make reliable images (86). This knowledge has stimulated the development of methods of mapping that can be used in the presence of large systematic errors in both the amplitude and the phase of the measured visibilities. These methods, now generally known as *self-calibration* or *hybrid mapping*, make use of the fact that most of the errors are associated with individual antennas, and that it is possible to form combinations of the measured visibilities, called *closure phases* and *closure amplitudes*, that are independent of such errors. Although the methods were originally developed for VLBI, where the errors are so large as to prevent conventional mapping, they are now also used in connected-element interferometers such as the VLA and the WSRT, and in the Jodrell Bank MERLIN array in England (33), to correct the much smaller errors that occur in such instruments; their application to VLBI has shown that it is possible to make reliable maps with submilliarcsecond resolution.

This review describes these self-calibration methods and their application to both conventional and very long baseline arrays and includes examples of the high-quality images that they produce. After first defining closure phase and closure amplitude in Section 2, we review conventional aperture synthesis methods in Sections 3 and 4. The new iterative self-calibration algorithms are the subject of Section 5. In Section 6 we consider the capabilities and limitations of these algorithms. Finally, in Section 7 we mention some applications in optical and infrared astronomy and in new radio instruments presently planned or under construction.

2. CLOSURE QUANTITIES

2.1 *Interferometry*

The theory of a radio interferometer has been described in a number of papers (e.g. 17, 39, 40, 89, 111, 114). In order to clarify the problems involved in calibrating measurements made with an interferometer, we first briefly summarize this theory.

Let $v_m(t)$ and $v_n(t)$ be the analytic signal representation of the voltages received by two antennas m and n, expressed as functions of time t. The interferometer measures the *mutual coherence function* (11) of the voltages received by the two antennas,

$$\Gamma_{mn}(\tau) = \langle v_m(t) v_n^*(t+\tau) \rangle, \qquad 1.$$

at some time delay τ (which may be adjusted); the angular brackets represent a time average. In continuum observing, the delay τ is adjusted so that $|\tau| \ll 1/\Delta\nu$, where $\Delta\nu$ is the bandwidth of the antenna and associated electronics. In this case the interferometer measures the *mutual intensity* or *complex visibility* $\Gamma_{mn}(0)$. The van Cittert–Zernike theorem (11, 111) shows that $\Gamma_{mn}(0)$ is proportional to a Fourier component of the sky brightness distribution:

$$\Gamma_{mn}(0) \propto \int\int I(\mathbf{s}) \exp(-2\pi i \mathbf{s} \cdot \mathbf{B}_{mn}/\lambda)\, d\Omega, \qquad 2.$$

where $I(\mathbf{s})$ represents the strength of the signal received from direction $\mathbf{s}$ (a unit vector), $\mathbf{B}_{mn}$ is the vector separation of the two antennas (baseline vector), λ is the wavelength at the center of the received band, and the integration ($d\Omega$) is over the celestial sphere. In spectral line observing, the delay dependence of $\Gamma_{mn}(\tau)$ is measured, and then the Fourier transform of $\Gamma_{mn}(\tau)$ with respect to τ gives the frequency dependence of $I(\mathbf{s})$. In the technique of aperture synthesis (which we discuss further in Section 3), one seeks to measure Γ_{mn} at a sufficient number of different antenna separations $\mathbf{B}_{mn}$ to form an estimate of $I(\mathbf{s})$ by inversion of Equation 2.

The quantity actually measured in a real interferometer is the complex visibility Γ_{mn} multiplied by a complex gain factor G_{mn}:

$$V_{mn} = G_{mn}\Gamma_{mn} + \text{noise}. \qquad 3.$$

The gain factor G represents amplitude and phase errors introduced by the interferometer itself and during the propagation of the signals through space and through the Earth's ionosphere and troposphere. If the interferometer electronics are sufficiently stable and well understood, it is possible in principle to calculate their contribution to G; but as there are always additional contributions that cannot be calculated, it is the usual practice to determine the G_{mn} factors from observations of sources of known structure, usually point sources. These calibration observations must be made sufficiently frequently that any time dependence of the G factors can be determined, and the calibration sources must be sufficiently close to the object under study that the propagation effects are indistinguishable.

The phase errors due to refraction in the troposphere become significant for frequencies higher than 2 GHz and baselines longer than 1 km, and they

vary on time scales between minutes and hours, owing to changes in water-vapor content of the atmosphere (e.g. 2, 50, 53, 66). As it is not practicable to make calibration observations more frequently than once every few minutes, conventional calibration can be almost impossible for baselines longer than a few kilometers at frequencies higher than 10 GHz. At frequencies below 1 GHz, the effects of tropospheric refraction are small compared with ionospheric refraction, which also makes calibration difficult on long baselines. In VLBI, there is the added problem of phase fluctuations in the local oscillators, which can reduce the coherence time to a minute or less. Amplitude errors are generally less important, and have longer time scales, than the phase errors, but again at high frequencies (>5 GHz) variable absorption by the atmosphere and mispointing of the antennas can introduce large amplitude errors with a time scale of a few minutes that cannot be calibrated.

If the complex gain errors G are varying faster than they can be measured, it would appear that it is impossible to recover the complex visibility Γ from the measurements. Progress can be made, however, when more than two antennas are used simultaneously. This is because the number of unknown gain factors is then less than the total number of measurements. In many cases, the gain factors can be associated with individual antennas:

$$G_{mn} = g_m g_n^*, \tag{4.}$$

where g_m is a complex gain factor for the individual antenna m. Most of the errors that occur in real interferometers can be factorized in this way; we discuss the origin and effects of errors that cannot be factorized in Section 6. In an array of N antennas, up to $N(N-1)/2$ simultaneous interferometer baselines can be formed. If all the antenna pairs are cross-correlated, this factorization reduces the number of unknown (complex) gain factors that must be determined by calibration from $N(N-1)/2$ to N. When N is less than $N(N-1)/2$, it is thus possible in principle to determine both the gain factors themselves and some parameters of the complex visibility of the source from the observations. This is the basis of the self-calibration methods that are the subject of this review.

In Sections 2.2 and 2.3, we enumerate explicitly two sets of parameters, the *closure phases* and *closure amplitudes*, that can be derived from observations with an N-antenna array even if the antenna gain factors g_m are unknown. It is convenient for this purpose to express the terms in Equation 3 as amplitudes and phases:

$$\begin{aligned} V_{mn} &= |V_{mn}| \exp(i\phi_{mn}), \\ G_{mn} &= |g_m|\,|g_n| \exp(i\theta_m) \exp(-i\theta_n), \\ \Gamma_{mn} &= |\Gamma_{mn}| \exp(i\psi_{mn}). \end{aligned} \tag{5.}$$

Equation 5 defines the observed interferometer phase ϕ_{mn}, the true visibility phase ψ_{mn}, and the phase errors θ_m and θ_n associated with each antenna.

2.2 *Closure Phase*

From Equations 3 and 5, the visibility phase ψ_{mn} on baseline mn is related to the observed phase ϕ_{mn} by

$$\phi_{mn} = \psi_{mn} + \theta_m - \theta_n + \varepsilon_{mn}, \tag{6.}$$

where ε_{mn} represents the noise in the measurement. The *closure phase* Ψ_{lmn} is formed by summing the observed phases around a triangle of baselines lm, mn, and nl:

$$\begin{aligned}\Psi_{lmn} &= \phi_{lm} + \phi_{mn} + \phi_{nl} \\ &= \psi_{lm} + \psi_{mn} + \psi_{nl} + \varepsilon_{lm} + \varepsilon_{mn} + \varepsilon_{nl}.\end{aligned} \tag{7.}$$

It can be seen that the θ_m terms cancel completely. This is true for any closed loop of baselines, not just a triangle. Thus the observed closure phase is uncorrupted by phase errors that are associated with individual antennas. These errors include those due to propagation effects along the line of sight, oscillator drifts, and uncertainties in the source position and baseline vectors. Certain residual errors that do not cancel remain, but these "closure errors" are usually small; we discuss them in Section 6. Figure 1 shows some measurements of phase on three baselines of the VLA, together with the calculated closure phase: the cancellation of antenna-based errors is dramatic.

The closure phase was first used by Jennison in triple-interferometer experiments at Jodrell Bank in the early 1950s (59, 60). As later interferometers became more phase-stable, the technique was no longer needed, until it was revived by Rogers and coworkers (92) for use in VLBI, where propagation effects and oscillator errors make the direct measurement of visibility phases impossible. The term "closure phase" is due to Rogers.

2.3 *Closure Amplitude*

With four antennas, k, l, m, n, it is possible to form ratios of the visibility amplitudes that are independent of the antenna gain factors:

$$A_{klmn} = \frac{|V_{kl}|\,|V_{mn}|}{|V_{km}|\,|V_{ln}|} = \frac{|\Gamma_{kl}|\,|\Gamma_{mn}|}{|\Gamma_{km}|\,|\Gamma_{ln}|}. \tag{8.}$$

Such ratios are called *closure amplitudes*. If all six interferometer pairs formed by the four antennas are correlated, three different closure amplitudes can be calculated (A_{klmn}, A_{klnm}, A_{knml}), but only two of these are independent. Ratios of this type were first measured with a four-antenna interferometer by Twiss et al. (116, 117). Their application to VLBI was first

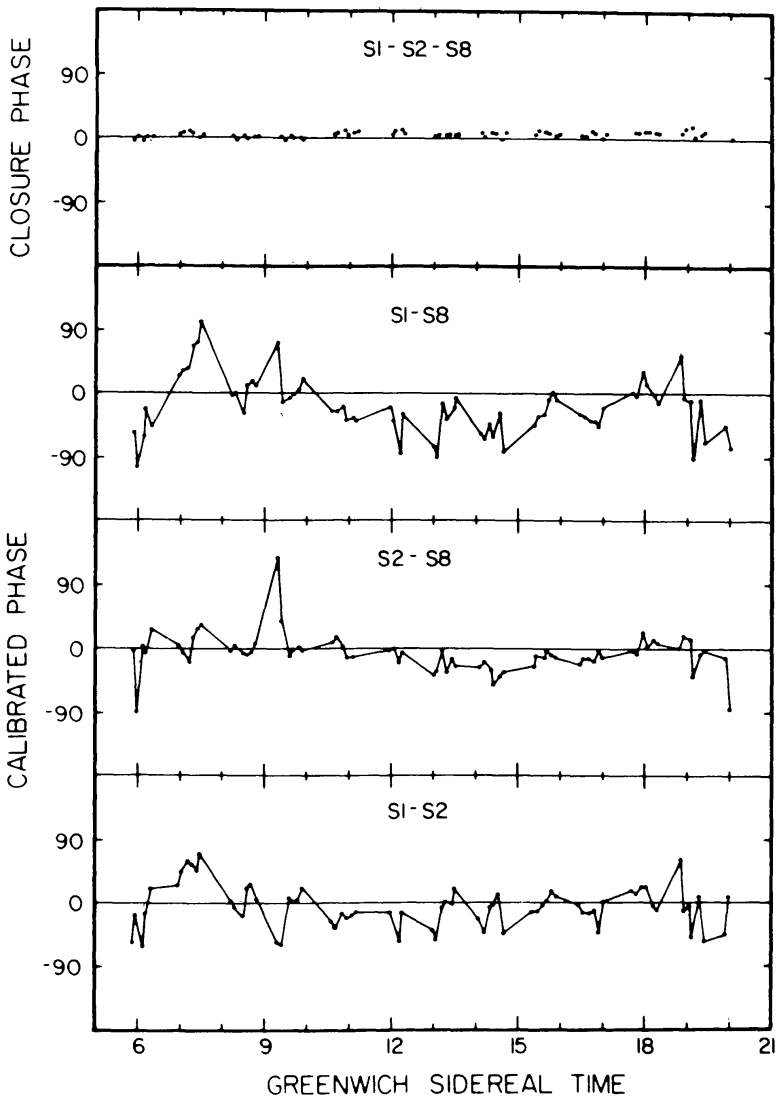

Figure 1 Illustration of the elimination of antenna-based phase errors by the calculation of closure phase. The lower three frames show the visibility phases observed on three baselines of the VLA in bad weather, after conventional calibration on a nearby point source. The upper frame shows the closure phase formed by summing the three visibility phases. The scatter is greatly reduced; the deviations from zero are due to structure in the source (the quasar 3C 147). Reproduced, with permission, from Readhead et al. (84).

discussed by Readhead et al. (85), who introduced the term "closure amplitude" by analogy with "closure phase." Note that at least four antennas are required to measure a quantity that is independent of all the antenna gain factors. With fewer antennas, though, it is still possible to eliminate some of the gain factors; for example, Smith (110), using three antennas and a single correlator, was able to eliminate the unknown gain of one of the antennas. Jennison (59) used three antennas and three correlators, and by monitoring the total power received by one antenna, he

could eliminate all the antenna gains; this method is equivalent to using the closure amplitude with one baseline of zero length.

2.4 *Interpretation of Closure Phase and Closure Amplitude*

In an array of N antennas, with $N(N-1)/2$ interferometer pairs, there are N unknown complex antenna gain factors g_m, that is, N unknown amplitude factors and N unknown phases. As only phase differences are important, however, one of the unknown phases can be set to zero, leaving a total of $2N-1$ unknowns. The $N(N-1)/2$ complex visibility measurements thus yield $N(N-1)-(2N-1) = N^2-3N+1$ parameters of the true source visibility. These may be divided into $(N-1)(N-2)/2$ independent closure phases and $N(N-3)/2$ independent closure amplitudes. In an N-antenna array, the fraction of the total visibility information available in the closure quantities is $(N-2)/N$ in the closure phases and $(N-3)/(N-1)$ in the closure amplitudes. These ratios show the rewards to be gained by increasing the number of antennas in the array; with only 4 antennas, 50% of the phase information and 33% of the amplitude information is available, while for 10 antennas (a typical VLBI configuration), these ratios increase to 80% and 78%, respectively.

The interpretation of measurements of closure phase and closure amplitude is difficult, as, unlike the visibilities, they are not related to the sky brightness distribution by a simple Fourier transform relationship. If only closure phases and closure amplitudes are measured, knowledge of both the absolute strength and the absolute position of the source is lost. This can be seen by (*a*) applying a scale factor to $I(\mathbf{s})$ in Equation 2, which merely introduces a scale factor in Γ that cancels when the closure amplitude is calculated; and by (*b*) shifting the function $I(\mathbf{s})$, which introduces phase factors in Γ that cancel when the closure phase is calculated.

In the early work with closure phase and closure amplitude (59, 116), the antennas could be arranged in a redundant configuration, so that all the interferometer baselines were parallel and multiples of a single minimum baseline. In this case it is possible to solve the closure relations to compute the visibilities on all the baselines in terms of that on the minimum baseline, so that the *only* information that is lost is the absolute strength and position of the source. Both the strength and the position of the source can be recovered if it is possible to calibrate the visibility measurement on a single baseline. Rogstad (93) has given an example of a redundant six-antenna array with which it is possible to compute the phases on all multiples of the minimum baseline up to ten, and Rhodes & Goodman (88) have indicated how this technique can be extended to two-dimensional arrays. In VLBI arrays, however, it is not usually possible to arrange the antennas in a redundant configuration, and hence other methods must be used to derive

the visibility phases. The methods used, nowadays called either *hybrid mapping* or *self-calibration*, are the subject of Section 5; but we must first digress to discuss the techniques of conventional aperture synthesis, of which the self-calibration methods are an adaptation.

3. APERTURE SYNTHESIS

3.1 *Basic Theory*

Let us rewrite Equation 2 in a coordinate system in which $\mathbf{B}_{mn}/\lambda = (u, v, w)$ and $\mathbf{s} = (x, y, (1-x^2-y^2)^{1/2})$:

$$\Gamma(u,v,w) = \int\int I'(\mathbf{s}) \exp[-2\pi i(ux+vy+w(1-x^2-y^2)^{1/2}] \, dx \, dy, \qquad 9.$$

where we have written $\Gamma(u, v, w)$ for $\Gamma_{mn}(0)$, and $I'(\mathbf{s}) = I(\mathbf{s})/(1-x^2-y^2)^{1/2}$. The sky-coordinates x and y are direction cosines relative to the center of the field of interest. In all the cases that we consider, it is possible to ignore w, and thus Equation 9 reduces to a two-dimensional Fourier transform:

$$\Gamma(u,v) = \int\int I'(\mathbf{s}) \exp[-2\pi i(ux+vy)] \, dx \, dy. \qquad 10.$$

Equation 10 is the fundamental equation of aperture synthesis, and the fundamental problem in aperture synthesis is, given a number of measurements of Γ at different points in the uv-plane, to invert Equation 10 in order to estimate the sky brightness I.

The treatment of w deserves a little discussion. In going from Equation 9 to Equation 10 we have removed a term $-2\pi w$ from the phase, which is usually applied as a phase correction to the measured visibilities, i.e.

$$\Gamma(u,v) = \Gamma(u,v,w) \exp[2\pi i w],$$

and we have ignored higher-order terms in w. When this is not a valid approximation, Equation 9 does not reduce to a Fourier transform, and its inversion is much more difficult (17, 114). Fortunately, w can be ignored in all cases in which all the interferometer baselines lie in a single plane (in which case w may be arranged to be zero by choice of coordinate system) and in cases in which the field of view occupied by the source is sufficiently small $[2\pi w(x^2+y^2) \ll 1]$, which includes almost all VLBI observations.

The conventional method of inverting Equation 10 is based on the inverse Fourier transform:

$$I'(x,y) = \int\int \Gamma(u,v) \exp[2\pi i(ux+vy)] \, du \, dv. \qquad 11.$$

The measurements are made at a set of only M points (u_j, v_j) in the uv-plane, and the Fourier transform must be approximated by the sum

$$I''(x, y) = \sum_{j=1}^{M} \Gamma(u_j, v_j) W_j \exp[2\pi i(u_j x + v_j y)], \quad 12.$$

where W_j is a "weight" associated with the jth point. This gives an estimate $I''(x, y)$ of the sky brightness, usually called the "dirty map," which is the convolution of the true sky-brightness $I'(x, y)$ with a point-spread function or "dirty beam":

$$I''(x, y) = I'(x, y) \otimes P(x, y), \quad 13.$$

where $\otimes$ denotes two-dimensional convolution. The point-spread function $P(x, y)$ is given by the right-hand side of Equation 12, with all the $\Gamma(u_j, v_j)$ terms replaced by 1. The weights may be chosen to adjust the shape of the dirty beam.

In practice, a Fast Fourier Transform (FFT) algorithm is usually used to compute Equation 12, which requires that the data be sampled on a regular rectangular grid. The choice of weights W_j, the techniques for interpolating the measured data onto a rectangular grid, and the distortions that this procedure introduces into the map are described elsewhere (17, 32, 114).

3.2 *Image Restoration*

The estimate of the sky brightness distribution provided by Equation 12 is a *linear* combination of the measured quantities. It is a convolution of the sky brightness distribution with a beam shape determined by the sampling of the uv-plane and the weights assigned to the measured points. Although this map will be sufficient to show many features of the source in cases where the uv-plane coverage is good, it is not the best possible representation of the sky, and it contains artifacts, most notably the positive and unphysical negative sidelobes around bright peaks. If the sampling of the uv-plane is irregular and uneven, the dirty beam will have large sidelobes that will confuse and obscure the structure of interest in the map. These are due to the erroneous assumption that the visibility is zero at spatial frequencies where it was not measured. The choice of weights W_j is necessarily a compromise between high resolution, low sidelobe level, and good signal-to-noise ratio (78).

A considerable improvement is possible by using *nonlinear* methods of image restoration, and a number of such methods have been introduced into radio astronomy in the last decade. These all attempt to estimate, or "restore," the unmeasured Fourier components in order to produce a more physical map; some of them can actually estimate components at larger spatial frequencies than any that were measured, and thus achieve an

increase in resolution (superresolution). These methods also have the advantage over the linear methods that they can take into account the constraint that the map must be everywhere positive or zero (brightness is never negative), and they can also make use of a priori knowledge about the extent of the source and the statistics of the measurement process. Nonlinear methods can always do a better job than any linear technique (121), but the restoration is not unique, because in almost all cases the number N of parameters defining the map (e.g. the brightness at N pixels on a rectangular array) is greater than the number of constraints ($2M$ if complex visibility measurements are made at M points, but less if self-calibration is required). Thus the image reconstruction method must select one of the possible solutions, and the choice of reconstruction method is tantamount to a choice of criteria for selecting the "best" solution. A large number of nonlinear restoration methods have been suggested and experimented with; here, we confine our discussion to three of the most popular: the Högbom CLEAN method, the Gerchberg–Saxton method, and the maximum entropy methods. These methods have matured somewhat since they were last reviewed in this series (13).

CLEANING The most widely used technique of image reconstruction in radio interferometry is the iterative beam-removing method called CLEAN; as this is a major part of current self-calibration algorithms, we shall discuss it in some detail. CLEAN was developed by Högbom in the early 1970s (55, 94, 107). It starts with a dirty map made by the usual linear Fourier inversion procedure and attempts to decompose this into a number of components, each of which has the shape of the dirty beam. The algorithm searches the dirty map for the pixel with largest absolute value $|I_{\max}|$ and subtracts a dirty beam pattern centered on this point with amplitude $\gamma I_{\max}$ (the factor γ is called the *loop gain*). Then the residual map is searched for the next largest pixel, and a second beam shape is subtracted, and so on. The iteration is terminated when the maximum residual is consistent with the expected noise level. The result of the iteration is a residual map containing noise and low-level source contributions, plus a list of the amplitudes and positions of the components removed. The list of components, considered as an array of delta functions, is a model of the sky brightness distribution that will reproduce the observed visibilities (within the noise). As this is not convenient for display, the delta functions are usually convolved with a "clean beam" and added back to the residuals to produce a "clean map." The clean beam can be chosen arbitrarily: usually a truncated elliptical Gaussian of about the same size as the central lobe of the dirty beam is used, thus producing a map with the same resolution as the original dirty map but uncontaminated by sidelobes. The Fourier

transform of the clean map does *not* reproduce the measured visibilities, but rather the measured visibilities multiplied by a taper function that is the Fourier transform of the clean beam.

It has been shown (105, 106) that CLEAN is equivalent to a least-squares fit of sine functions to the visibility data. There is now considerable experience with CLEAN in practice (e.g. 26). A fast algorithm based on the FFT has been introduced by Clark (21), and a modification of CLEAN with improved behavior has been suggested by Cornwell (28). Perhaps the least-understood aspect of CLEAN is the convolution of the delta-function components with the clean beam: there are no theoretical criteria for choosing the clean beam. Although the normal practice is to use a clean beam comparable to the central lobe of the dirty beam, it is sometimes possible to achieve superresolution by using a smaller clean beam. Experiments have shown, however, that such superresolution is untrustworthy, unless the signal-to-noise ratio is very high and the *uv*-plane coverage is very good. In no case has superresolution by a factor >2 been demonstrated.

THE GERCHBERG–SAXTON ALGORITHM The Gerchberg–Saxton or Fienup algorithm was originally introduced for solving the phase problem in electron microscopy (45) and has been adapted for reconstructing an image from measurements of the visibility amplitudes only (37, 38). Although it has not gained a large following in radio astronomy, it is closely related to the iterative self-calibration algorithms that are the subject of Section 5. It is an iterative algorithm involving repeated Fourier transformation between the map (or sky) plane and the *uv*-plane, imposing the known constraints in each domain. In radio astronomy, the constraints are that the map should nowhere be negative and that the *uv*-plane data should agree with the measurements where available; the map can also be constrained to be zero outside one or more selected "windows." Thus, starting with an estimate of the visibilities in the *uv*-plane, as in the standard linear Fourier inversion, we transform to the sky plane, set all negative regions and any regions outside the window to zero, and transform back to the *uv*-plane. To start the next iteration, we replace the visibilities thus derived by the measured visibilities where they are known. The procedure, therefore, allows arbitrary values for the visibility in unsampled regions of the *uv*-plane. Gerchberg (44) has shown that this method is capable of superresolution. Högbom (55) has discussed the use of this method in radio astronomy and pointed out its main drawback (which it shares with the maximum entropy methods)—that the resolution achieved is data dependent and varies across the map. Methods of this type may be regarded as special cases of more general "constrained iterative restoration algorithms" (100).

MAXIMUM ENTROPY AND RELATED METHODS The limited sampling of the *uv*-plane means that there will be an infinite number of possible maps consistent with the data within the accuracy of the measurements. Not all of these possibilities will be everywhere positive. The maximum entropy methods, which were introduced to radio astronomers by Ables (1), attempt to choose the map that is in some sense the most "likely" or most "reasonable" a priori, and that satisfies the positivity constraint. If we represent the map by the vector $\mathbf{b} = (b_1, b_2, \ldots, b_N)$ formed from the brightnesses of its N pixels and define an "entropy" function $Q(\mathbf{b})$, the methods involve choosing $\mathbf{b}$ to maximize $Q(\mathbf{b})$ subject to the constraint

$$\sum_{j=1}^{M} |\Gamma(u_j, v_j) - V(u_j, v_j)|^2/\sigma_j^2 \leq M, \qquad 14.$$

where $\Gamma(u_j, v_j)$ is the complex visibility predicted by the map $\mathbf{b}$, $V(u_j, v_j)$ are the visibility measurements, the sum is over the M sampled (u, v) points, and σ_j^2 is the variance of the jth point. The left-hand side of Equation 14 is a χ^2 measure of the disagreement between the visibilities predicted by the map and the measured visibilities, and the constraint ensures that the resulting map will be consistent with the data, taking into account its noise statistics, which are assumed to be Gaussian. Other models of the noise statistics, or other constraints (19), can be incorporated in the maximum entropy method, but this is the most usual. The solution of the constrained maximization of $Q(\mathbf{b})$ is not a trivial matter, but Wernecke (122) and Gull & Daniell (48) have developed iterative numerical techniques that can solve the problem in a reasonable amount of computer time by use of the FFT algorithm.

For the entropy $Q(\mathbf{b})$, two functions have been suggested:

$$Q_1 = \sum_n \log b_n, \qquad 15.$$

$$Q_2 = -\sum_n b_n \log b_n. \qquad 16.$$

There is still some controversy, which we do not enter into here, as to which of these functions should be used and which, if either, deserves to be called entropy. The Q_1 function, favored by Wernecke & D'Addario (122, 123), is a measure of the entropy of the electromagnetic field sampled by the antenna, while Q_2, favored by Frieden (42) and Gull & Daniell (48, 49), is a measure of the configurational entropy or information content of the map itself. In practice, it appears that both functions give similar maps when there are sufficient data. Other "entropy" functions with similar behavior have been suggested (74), corresponding to different a priori assumptions about the sky brightness. A radically different approach is the "maximum sharpness"

criterion (4, 71):

$$Q_3 = -\sum_n b_n^2. \qquad 17.$$

The maximum entropy methods have not yet gained a strong following among radio astronomers, probably because the behavior of these methods is not yet fully understood. For example, they can superresolve, in some cases by up to a factor of two, but the reliability of this superresolution has not yet been demonstrated. Nor are all astronomers convinced that the maximum entropy map is the best possible; for example, the maximum of the entropy function may be very broad, and in such cases it is not really appropriate to suggest that any single map is a good estimate of the sky brightness. The maximum entropy methods do, however, have the advantage that they can identify such cases.

4. INTERNAL CALIBRATION METHODS

The calibration problems described above can be overcome if there is a suitable reference source in the same isoplanatic patch as the object. Ideally, such a reference source should be unresolved, and it must be possible to separate its contribution to the visibility measurements from that of the source under study. Such a separation is possible if the reference source is sufficiently far removed from the object, or if the reference source and the object appear in different frequency channels. If the displacement between the reference source and the object is large enough, the interferometer delay (τ in Equation 1) and fringe rate (the rate of change of the visibility phase) will be different for reference source and object, which allows their contributions to the visibility to be separated (82). The power of this technique is illustrated by the VLBI detection of a third compact radio component in the gravitational lens 0957 + 561 (47). In one particular class of object, the cosmic maser sources, it is possible to separate an internal reference source by frequency. The cosmic masers (both OH and H_2O) generally consist of a number of well-separated components, some of which are unresolved or have very simple structure; the maser lines are narrow, and the different velocities of the components cause them to appear in different frequency channels. In some cases, one frequency channel is found to contain a single unresolved source, which may be used as a reference for calibrating the other channels. Such cases are rare, but if it is possible to obtain the brightness distribution in a particular frequency channel without using phase (e.g. by model-fitting using the amplitude measurements), or if a single channel can be mapped by one of the hybrid-mapping methods of Section 5 (77), the data in that channel can be used as a phase reference. Two

methods, "fringe rate mapping" and "phase mapping," have been used to map maser sources.

In fringe rate mapping (69, 70, 119), narrowband spectra of the object are examined, and a reasonably strong unblended spectral feature is chosen as the reference source. The most sensitive interferometer baseline is selected for comparing the fringe rates of all the features. The interferometer phases at each point in the spectrum are adjusted so that the reference source has zero phase at all times, and hence zero fringe rate. The residual fringe rates that remain for the other sources are due to position offsets from the reference source, and the relative positions of the different spectral features can thus be determined from the relative fringe rates. In principle, accuracy of about 10^{-4} arcsec can be achieved by this method (43).

In the phase-mapping (or aperture synthesis) method (87, 120), a reference source in one channel is used to calibrate the phases in all the other channels, which can then be mapped by the usual aperture synthesis methods. The drawback of this method for VLBI observations of masers is that the maser components are spread over a field of view that may be thousands of times larger than the beamwidth, so that the computation involved in aperture synthesis is prohibitive. In some cases it is useful to begin by making a fringe rate map of the entire masing region using the most sensitive interferometer pair, followed by low-resolution phase maps, and finally by phase maps at the full resolution of selected areas encompassing the most interesting groups of masers. The reference features are sometimes spatially resolved, in which case they must be modeled before being used as a reference. The relative phases of the other features depend on baseline vectors, station clocks, and propagation effects, as well as on the positions and the frequencies of these features. The errors are usually dominated by clock and baseline errors, and the positional accuracy achieved is usually about 5×10^{-5} arcsec.

The internal calibration methods are restricted in scope because of the requirement that there be a calibration source close to the object under study. The hybrid-mapping methods, which are the subject of the next section, circumvent this difficulty by using the object itself as the calibration source.

5. HYBRID-MAPPING METHODS

When insufficient data are available to make a dirty map, the standard mapping and image-restoration procedures cannot be used. The missing data may be, for example, the visibility phases (in which case we want to derive a map from the amplitudes only, and thus implicitly make an estimate of the phases) or the complex antenna gains (in which case we want

to estimate the antenna gains as well as the sky brightness distribution from the available data). We might even be faced with the problem of mapping with phase information but no amplitudes.

A number of iterative procedures have been introduced for tackling problems of this type. Although they differ in detail, they all have a similar structure, involving repeated Fourier transformation between the map plane and the *uv*-plane, with adjustments in each domain. They could be regarded as variants of the Gerchberg–Saxton algorithm; but in radio astronomy they are generally termed "hybrid-mapping" methods, following Baldwin & Warner (6) and Readhead & Wilkinson (86).

A single iteration consists of the following steps (Figure 2): (*a*) Starting with a *model* of the sky-brightness distribution, called a *trial map*, obtain *trial visibilities* by Fourier transformation. (*b*) Combine the (incomplete) visibility measurements with sufficient information from the trial visibilities to complete the sampling of the *uv*-plane. (*c*) Compute a *hybrid map* by Fourier inversion of these estimated visibilities. Baldwin & Warner (6) defined a hybrid map to be a map made using amplitudes that are appropriate to one sky brightness distribution and phases that are appropriate to a different one. It seems reasonable to generalize the definition to include maps made from any mixture of measurements and guesses from the trial map. (*d*) Examine the hybrid map and use it to derive a better trial map for the next iteration.

The procedure should be iterated until the hybrid map is indis-

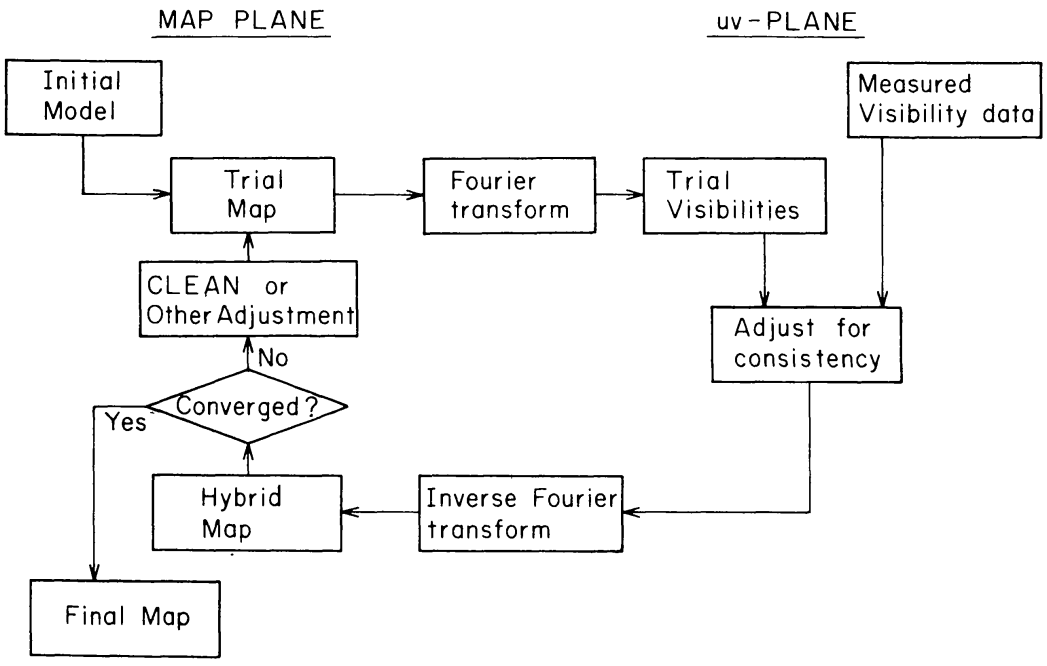

Figure 2 The iterative hybrid-mapping algorithm.

tinguishable from the trial map. There is no guarantee that the procedure will converge—this depends principally on step (*d*)—but if it does, the final hybrid map is the best estimate of the correct map.

The various hybrid-mapping algorithms in the literature differ in how they combine measured data with trial-map data [step (*b*)] and in the procedure for deriving the trial map [step (*d*)]. They all require an initial trial map to start the iteration, which may be obtained in a number of ways. In Sections 5.1, 5.2, and 5.3 we describe first the method of Baldwin & Warner for phaseless aperture synthesis and then discuss two methods for mapping in the presence of large antenna gain and phase errors: the original method of Readhead & Wilkinson, and the more general self-calibration method of Schwab and Cornwell & Wilkinson. Following this discussion, we examine the adaptation of maximum entropy for self-calibration (Section 5.4), as well as two related iterative calibration procedures: the use of redundancy (Section 5.5) and global fringe-fitting in VLBI (Section 5.6).

5.1 *Method of Baldwin & Warner*

Baldwin & Warner (5–7) were concerned with the problem of phaseless aperture synthesis. The hybrid map in this case is made from the measured amplitudes and trial phases. They considered sky brightness distributions consisting of a number of point sources in an otherwise empty field, and they showed that it is fairly easy in this case to obtain an initial trial map from the map of amplitude squared (the autocorrelation function of the true map) and to identify features that are missing from the trial map by inspection of the hybrid map.

5.2 *Method of Readhead & Wilkinson*

If we have measurements of amplitude and closure phase only in an array of N antennas, there are $(N-1)$ unknown phase errors that must be estimated by hybrid mapping. Readhead & Wilkinson (86, 125), inspired by the work of Fort & Yee (41), chose to use phases computed from the trial map on $(N-1)$ of the $N(N-1)/2$ baselines, and to use the measured closure phases to deduce phases for the remaining baselines. This requires the observer to specify which visibility phases should be obtained from the trial map; this choice affects the speed of convergence of the algorithm, but it should not change the ultimate solution. Cotton (31) suggested that it would be preferable to choose the $(N-1)$ unknown phase errors to minimize the differences between the hybrid and trial visibilities. Both Readhead & Wilkinson and Cotton used CLEAN on the hybrid map to generate a delta-function model to be used as the next trial map. This has to be done with care, since if CLEAN is allowed to converge, the delta-function trial map

would give identical visibilities to the hybrid map, and the iteration would not advance. The solution to this is to not clean too deeply, to omit large negative CLEAN components from the trial map, and to restrict CLEAN to a small "window" (30).

Readhead & Wilkinson made "blind tests" of their method that showed that the procedure usually does converge quite quickly to the correct solution. It has been used widely in VLBI since its first application to observations of the quasar 3C 147 (126) and has been applied to VLA observations made in conditions of poor phase stability (84): Figure 3 shows the great improvement in the map obtained for 3C 147. Readhead et al. (85) extended the method to observations with poor amplitude calibration by utilizing the closure amplitudes in a similar way to the closure phases, using amplitudes from the trial map on some baselines and deducing the remainder from the closure amplitudes. This method has successfully corrected amplitude errors of a factor of 20 (83).

Fort & Yee (41) and Rogers (90) suggested variants of the Readhead–Wilkinson method in which instead of using CLEAN, they obtained a trial map by setting to zero (*a*) regions of low or negative brightness in the hybrid map, and (*b*) regions outside an a priori window.

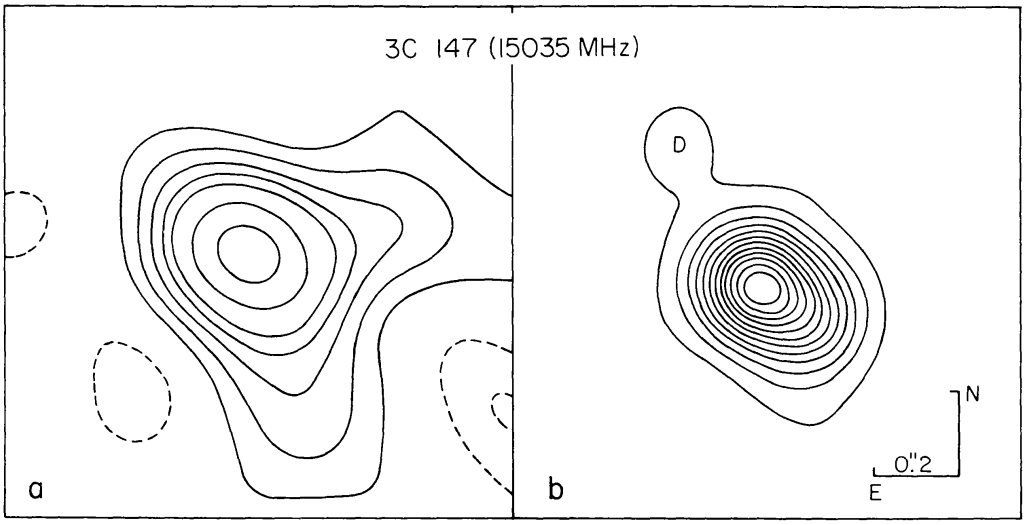

Figure 3 Comparison of maps made with and without self-calibration. (*a*) VLA map of 3C 147 made with conventionally calibrated phases, some of which are shown in Figure 1. The calibrated phases are severely corrupted owing to bad weather during the observations. All of the extended features in this map are spurious; note the large negative sidelobes. Contour levels: −20, −10, 10, 20, 30,...% of the peak value. (*b*) Hybrid map made from the same observations, using the closure phases. The extensions seen here have been confirmed by subsequent observations. Contour levels: 4, 12, 20,...% of the peak value. Reproduced, with permission, from Readhead et al. (84).

5.3 *Method of Schwab and Cornwell & Wilkinson*

In the Readhead–Wilkinson method, the closure phases and closure amplitudes are computed explicitly, and the closure relations are solved on a baseline-by-baseline basis. Both Schwab (102) and Cornwell & Wilkinson (29) showed that it was preferable to solve for the unknown antenna gains directly, without explicit calculation of the closure quantities, as this allows a correct treatment of noise in the measurements and is more economical of computer resources. This procedure involves adjusting the estimates of the N complex antenna gain errors to minimize the mean square difference between the measured visibilities and the true visibilities:

$$\sum_{m<n} w_{mn}|V_{mn}-g_m g_n^* \Gamma_{mn}|^2, \qquad 18.$$

where w_{mn} are weighting factors that can be chosen on the basis of the observed signal-to-noise ratio. This is a nonlinear least-squares problem that must be solved iteratively. When used in hybrid mapping, an estimate of Γ_{mn} is obtained from the trial map, and the estimates of the gain-factors g_m are used with the measured visibilities V_{mn} to form a hybrid map for the next iteration. Both Schwab and Cornwell & Wilkinson used CLEAN to derive a new trial map from the hybrid map. Cornwell & Wilkinson improved on Schwab's method by showing how to take account of prior information about the signal-to-noise ratio of the data and the expected variances and time scales of the gain errors. Least-squares (Equation 18) is not necessarily the best way of choosing the g-factors. An alternative that is less sensitive to occasional large, non-Gaussian errors in the measurements is to minimize the sum of the absolute values of the differences (the L_1 norm) instead of the sum of their squares (the L_2 norm) (14, 103).

Maps made from VLA data (by Schwab's method) and from MERLIN data (by the Cornwell–Wilkinson method) are shown in Figures 4 and 5. Other examples can be found in the article by Bridle & Perley in this volume (15).

5.4 *Maximum Entropy Methods*

The maximum entropy methods can be adapted to self-calibration by changing the constraints on the maximization to ensure that the maximum entropy map is consistent not with the visibilities but with the measured closure quantities. The algorithms that have been used for solving the maximum entropy problem all involve repeated transformation between the sky and the *uv*-plane, so it is a fairly simple matter to incorporate the closure constraints in the *uv*-plane in the same way as is done in the Readhead–Wilkinson or Schwab–Cornwell–Wilkinson methods. We have

Figure 4 Map of the radio galaxy Her A (3C 348), made using the VLA in three configurations at 5 GHz. The data were self-calibrated using Schwab's method. The resolution is 0".5, and the radio brightness ranges from 0.05 to 22 mJy per beam area. Reproduced, with permission, from Dreher & Feigelson (34). Copyright © 1984 by Macmillan Journals Limited.

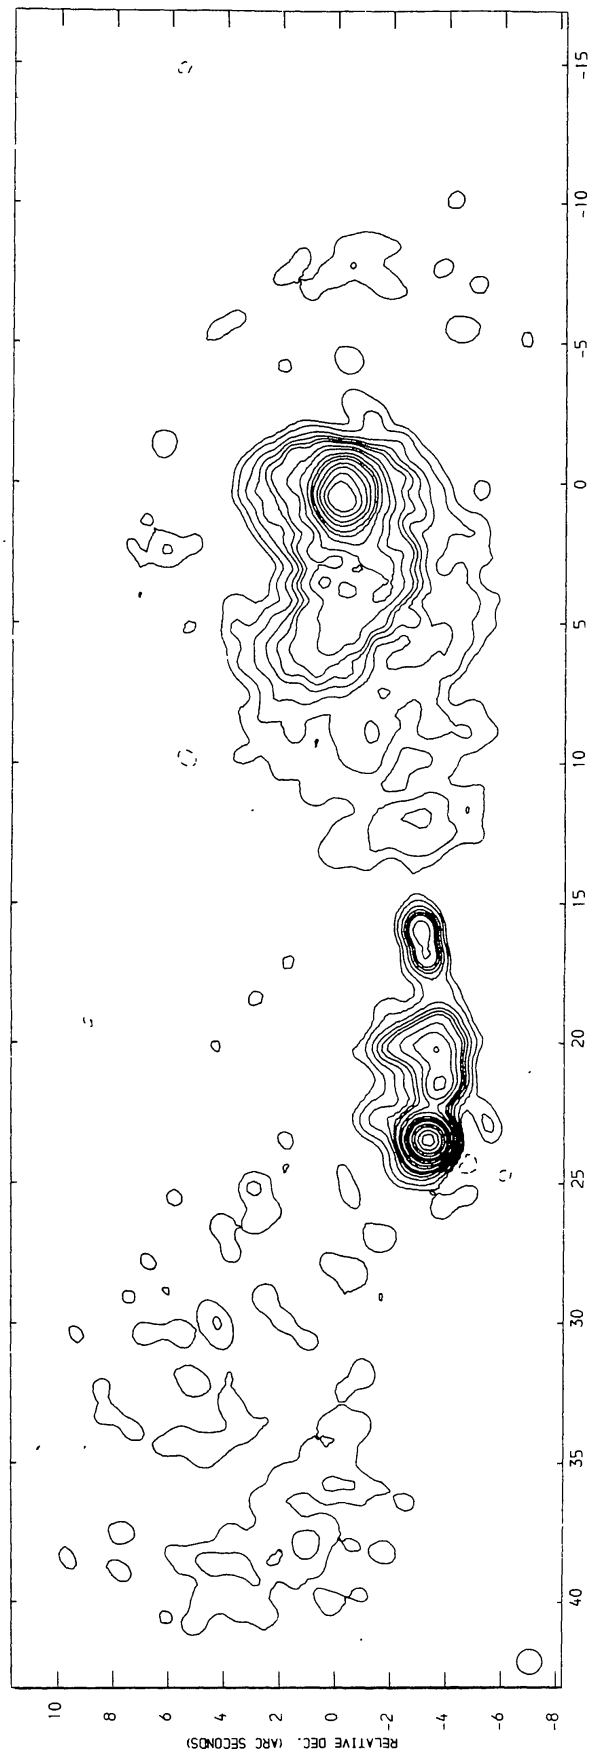

Figure 5 Map of the quasar 3C 249.1 at 408 MHz, made with the six-antenna MERLIN array using the Cornwell–Wilkinson self-calibration method. The lowest contour level is 0.5% of the peak brightness. Reproduced, with permission, from Lonsdale & Morison (65). Copyright © 1983 by the Royal Astronomical Society.

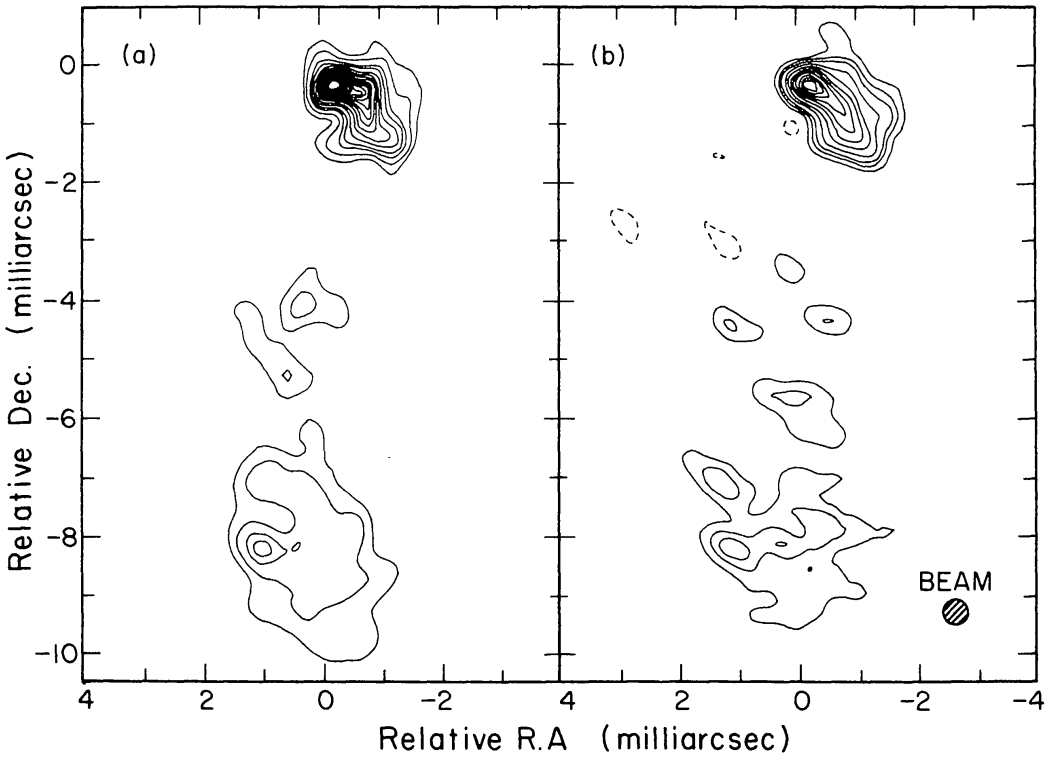

Figure 6 Maps of the nucleus of NGC 1275 (3C 84) at 22 GHz, obtained with a VLBI array of six antennas. (*a*) Map made by the maximum entropy method; (*b*) map made by the Readhead–Wilkinson method. The contour levels are the same in both maps, and the lowest contour is 4% of the peak in (*b*). The maximum entropy map has a higher peak owing to its higher resolution.

had considerable experience with a program of this sort written by S. F. Gull (private communication), and Sanroma & Estalella (99) have described some tests of a similar program. An example of a map made from VLBI data by the maximum entropy method is shown in Figure 6. The only problem with these methods is that the final map is more difficult to interpret than one with the same resolution at all points. It is clear, however, that as maximum entropy methods become better understood in conventional aperture synthesis, they will also be used for self-calibration.

5.5 *Redundancy*

As we mentioned in Section 2.4, if there is sufficient "redundancy" in the interferometer array, it is possible to recover almost all the visibility data from measurements of the closure quantities. A redundant array is one in which there are baselines of the same length and orientation between different pairs of antennas. If there is enough redundancy, the number of independent errors may be reduced to one gain and one phase error, and only the absolute intensity and position of the source are lost. For arrays with few redundant baselines the errors in the relative gains accumulate and

make the full relative gain solution unreliable, but these errors can be reduced in heavily redundant arrays if all of the redundant information is used. (This is not, however, a good argument for building a redundant array: with a fixed number of antennas, it is usually best to arrange them to optimize the sampling of the *uv*-plane, rather than to maximize redundancy.)

Noordam & de Bruyn (75, 76) have used this method on the WSRT, which has many redundant spacings. The WSRT is a linear east-west array that, at any instant, has a one-dimensional fan beam. The beam rotates with time, and the brightness distribution is built up from successive scans, each of duration $\ll 12$ hr. For each scan, all spacings that carry redundant information are used to obtain a weighted least-squares solution of the relative antenna gains. This fixes the relative gains of all antennas within one scan. Ionospheric and tropospheric wedges cause phase gradients that vary between scans, and there are a variety of effects that cause long-term overall gain variations. Successive scans must therefore be corrected for both phase gradients and overall gain variations. This is done by comparing the one-dimensional brightness distribution produced from the scan with a source model, and then correcting the level and position of the brightness distribution. This is nearly the same as the self-calibration procedures described above; the only difference is that the correction is done on a number of one-dimensional scans sequentially instead of on the whole source simultaneously.

When this method is used, the WSRT can achieve a very high dynamic range. The dynamic range is the ratio of the strongest feature to the weakest believable feature in the map. In their map of the active galaxy NGC 1275, Noordam & de Bruyn obtained a dynamic range of 10,000: 1, which is the highest that has yet been obtained. NGC 1275 is a very strong source; with weaker sources, the dynamic range will be limited by the noise in the measurements. An example of a high dynamic range map produced by this method is displayed in Figure 7, which shows the faint extended structure in the quasar 3C 345. The lowest contour is a factor of 8000 below the peak brightness of the core.

A similar method has been used to correct phase errors in an interferometer at Toyokawa Observatory (58).

5.6 *Global Fringe-Fitting in VLBI*

In VLBI, as in other types of interferometry, one measures the amplitude and phase of the complex visibility. Owing to uncertainties in the propagation delay, the baseline parameters, and the oscillator phase, the delay (τ in Equation 1) and the fringe rate (the rate of change of the measured visibility phase) cannot be predicted accurately, so the procedure

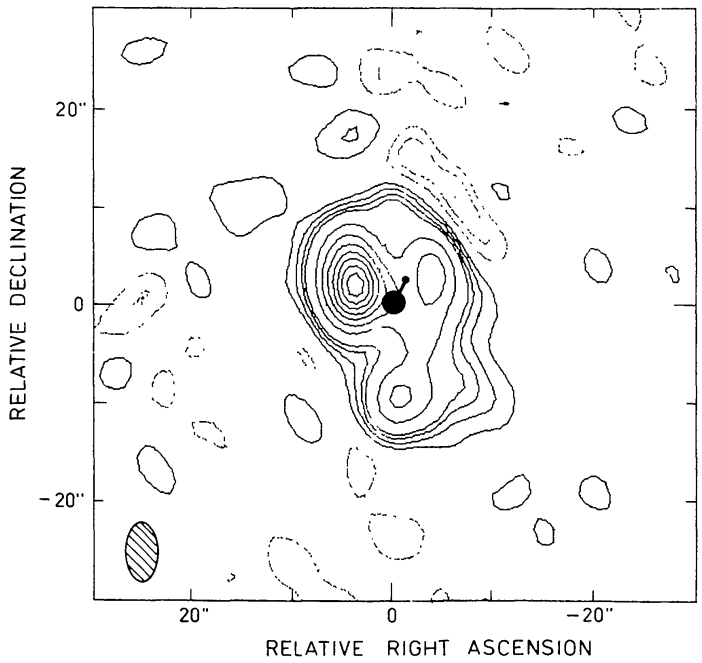

Figure 7 Map of the quasar 3C 345 at 5 GHz, made using the redundancy method on the WSRT. A point source of 10.15 Jy and a jet of 0.35 Jy have been subtracted at the positions indicated by black dots. The lowest contour level is 1/8000 of the peak brightness. Reproduced, with permission, from Schilizzi & de Bruyn (101). Copyright © 1983 by Macmillan Journals Limited.

adopted (70) is to cross-correlate the signals at several different trial delays and to integrate with several trial fringe rates, and then to search for the maximum amplitude as a function of delay and fringe rate. In the past, this search was carried out independently for each baseline. This requires a high signal-to-noise ratio on every baseline in order to obtain reliable detections. A large part of the uncertainty in delay and fringe rates is due to errors that cancel around triangles, like the closure phases. The number of unknown delays or fringe rates is N (the number of antennas), rather than $N(N-1)/2$ (the number of baselines). This constraint can be used to advantage if the data for all $N(N-1)/2$ baselines are used simultaneously in a *global* search instead of a baseline-by-baseline one. For example, if the source is unresolved, the "closure delays" and "closure fringe rates" are exactly zero, so that if fringes can be detected with good signal-to-noise ratio on two baselines, the delay and fringe rate can be predicted exactly for the third baseline of the triangle, which allows the third visibility to be measured even if its signal-to-noise ratio is low.

If the source is resolved, however, the closure delays and closure fringe rates are not zero, and a more complicated procedure must be used. Schwab & Cotton (104) have devised and programmed such a method of performing the global solution for a resolved source. As in the hybrid-

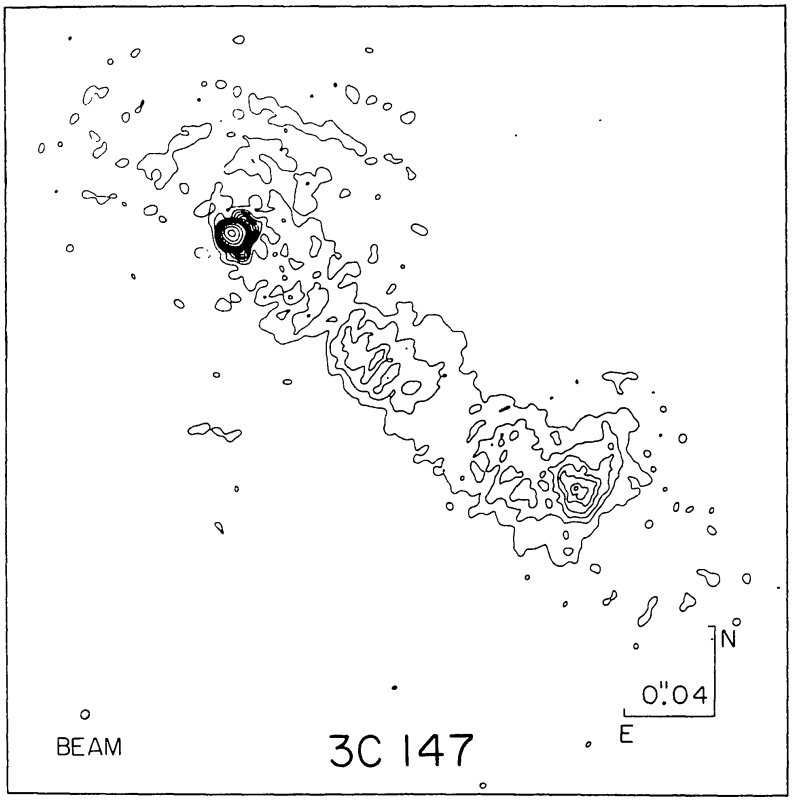

Figure 8 Ten-antenna VLBI map of the quasar 3C 147 at 1.7 GHz, made using the global fringe-fitting technique to obtain high signal-to-noise ratio and high dynamic range. This is a higher-resolution map of the object shown in Figure 3: the resolution is 0".004. The lowest contour level is 0.5% of the peak brightness. Reproduced, with permission, from Simon et al. (109).

mapping methods, a source model is used to predict the structure effects in delay and fringe rate, which are used in the global solution. The visibility measurements obtained are then used to make a map by self-calibration methods, and the map is in turn used as a model for another iteration of global fringe-fitting. Global fringe-fitting is thus a doubly iterative procedure, with an iterative self-calibration nested within each iteration of the fitting procedure. It is expensive in computer time, but the rewards it brings are well worth this investment, as it permits the use of data on baselines that are much too weak for single-baseline fitting and increases the overall signal-to-noise ratio. An example of a map made from VLBI data using this technique is shown in Figure 8.

6. CAPABILITIES AND LIMITATIONS

The various self-calibration algorithms attempt to estimate the sky brightness distribution from measurements of the closure phases and

closure amplitudes. If we knew the sky brightness distribution, we could calculate the closure quantities; thus self-calibration is a member of the general class of *inverse problems* that occur in many fields (e.g. 10, 81). Many techniques have been introduced for solving these problems, and some of them are closely related to the self-calibration algorithms. For any algorithm, we should consider the questions of *uniqueness*, *construction*, and *stability*. Uniqueness is a mathematical matter: given perfect observations, is there only one solution to the problem? If there is not, there certainly will not be a unique solution with incomplete and noisy data. The question of construction is, will the algorithm actually find the solution in a reasonable number of iterations? Finally, the algorithm is unstable if a small change in the measured quantities leads to a large change in the solution. None of these questions has yet been answered for self-calibration, but in this section we discuss the effects of incomplete *uv*-plane coverage and measurement errors on the convergence of the algorithms used for self-calibration.

6.1 *Uniqueness and uv-Plane Coverage*

One may naturally be worried that there might be several brightness distributions compatible with the closure quantities, and that the final self-calibrated map might not be unique, even if there were no errors in the measurements. We have already seen that if only the closure quantities are measured, we cannot determine the absolute strength and location of the source; that is, if $I(\mathbf{s})$ is a brightness distribution consistent with the closure quantities, then so is $cI(\mathbf{s}+\mathbf{s}')$ for any factor c and displacement $\mathbf{s}'$. Are these the *only* distributions consistent with the closure quantities? The question of uniqueness is tied up with that of the adequacy of the *uv*-plane sampling for mapping the source, and we should distinguish the problems of *uv*-plane coverage that are common to both conventional aperture synthesis and self-calibration from those that occur only with self-calibration.

There are no general rules available for determining whether a given sampling of the *uv*-plane is sufficient to map a source by aperture synthesis. To map a source that is confined within a rectangle of size θ radians, it is sufficient to sample the *uv*-plane on a rectangular grid with spacing $<1/(2\theta)$ wavelengths (the Nyquist sampling rate); but unless this sampling is extended indefinitely, there will be infinitely many brightness distributions consistent with the measured visibilities (12). If the *uv*-plane coverage extends to a maximum baseline of $u_{\max}$, the resolution of the map is $\sim 1/u_{\max}$, which amounts to saying that all the possible brightness distributions will look similar when convolved with a beam of this width. In practice, particularly with VLBI arrays, the sampling of the *uv*-plane can be sparse and irregular, and there are no clear-cut rules for what field of view can be mapped, or with what resolution; it depends very much on the

source structure and the signal-to-noise ratio. Because the total number of measurements is less, it is clear that the number of picture elements that can be mapped will be less when closure phases are measured in place of visibility phases, or closure amplitudes in place of visibility amplitudes.

A case that has received considerable attention is that of mapping with amplitudes but no phases. This "phase problem" occurs in X-ray diffraction, electron microscopy, speckle interferometry, and many other fields. In general, the solution is ambiguous. Complete knowledge of the visibility amplitudes is insufficient to define a unique brightness distribution, even with the constraints that the brightness distribution must be positive and confined to a finite solid angle; but it now appears that in *most* cases, the brightness distribution can be reconstructed uniquely (ignoring the "trivial" ambiguities: an arbitrary translation of the image or a rotation through 180° does not change the visibility amplitudes). There is still some controversy, however, as to the correct interpretation of the word "most" (e.g. 9, 18, 36, 57). It is likely that knowledge of the closure phases is sufficient to resolve the ambiguities, but this has yet to be proved.

The converse problem of mapping with phase measurements but no amplitude measurements has received less attention, as, in practice, phases are usually much harder to measure than amplitudes. An interesting review of phase-only mapping has been presented by Oppenheim & Lim (79), who point out that the phase of the visibility function carries more information about the brightness distribution than the amplitude does. They give examples of hybrid images formed from the visibility amplitudes of one image and the visibility phases of another. The hybrid image closely resembles the original image with the same phase. This effect was also noticed by Baldwin & Warner (6) in one of their tests. A hybrid map made with the correct phase but an amplitude of unity bears much more resemblance to the true map than one made with the correct amplitude and zero phase. As in the amplitude-only case, it appears that phase-only mapping can produce a unique solution (apart from a scale factor), or equivalently that the magnitude can be recovered from the phase, in many practical situations in which the source is confined to a finite area.

The mathematical question of uniqueness in self-calibration is unsolved. In real observations, the situation is confused by the incomplete *uv*-coverage and the errors in the measurements. But in practice it appears that the self-calibration algorithms do converge to a unique solution in cases where the *uv*-coverage is adequate. This is supported by the results of "blind tests" of the self-calibration methods (85, 86). As in normal aperture synthesis, it is not clear what constitutes "adequate" sampling of the *uv*-plane, but heuristically at least it appears that self-calibration requires more complete sampling than normal aperture synthesis.

6.2 *Noise*

In aperture synthesis with good *uv*-plane coverage, the fundamental limitation to the quality of the map is statistical noise in the measurements. Noise is introduced by the receivers, the microwave background, galactic background emission, atmospheric emission, pickup from the ground, and other causes. To a good approximation it may be regarded as white, Gaussian noise. When a linear image restoration procedure is used, it is straightforward to compute the effect on the map of Gaussian noise in the visibility measurements: it appears as Gaussian noise on the map with rms (in order of magnitude) $\sigma/\sqrt{M}$ per beam area, where σ is the rms noise in an individual visibility measurement, and M is the number of independent visibility measurements. The noise deflections on the map are correlated on a scale of the synthesized beamwidth. A more complete discussion can be found in the general reviews of aperture synthesis referred to earlier; a particularly good treatment is that of Napier (72). The noise behavior of the nonlinear image restoration methods is less well understood, and self-calibration is inherently nonlinear, so a rigorous understanding of the effect of noise on self-calibration is difficult. Self-calibration methods increase the noise level on the map because of errors in the determination of the complex antenna gain factors, which are themselves determined from the noisy data. Cornwell (25, 27) has shown that the rms noise level on the map is increased by a factor of $\sqrt{(N-1)/(N-2)}$ if only phases are corrected, or $\sqrt{(N-1)/(N-3)}$ if both phases and amplitudes are corrected. These factors are small when the number of antennas (N) is large. This, however, is only true if each individual visibility measurement can be made with a good signal-to-noise ratio in a time less than the time scale of variation of the antenna gain factors. If the signal-to-noise ratio is too small, it is not possible to make a reliable estimate of the antenna gain factors before they change, and self-calibration is impossible.

6.3 *Convergence*

A robust self-calibration algorithm should converge uniformly toward a unique solution and should be insensitive to the choice of starting model. In general, an arbitrary starting model can be used, though in practice convergence is greatly improved by starting with a sensible model; for example, a point source is a good starting model for mapping a source that is known to be barely resolved. When mapping with amplitudes only, the starting phases can be chosen to be zero or can be entirely random; when mapping with phases only, the amplitudes can be chosen to be unity or, even better, to decrease with increasing baseline in the manner of realistic sources. In a partially phase-stable array, like the VLA, an adequate

starting model can be obtained from a conventionally cleaned map; when the visibility phases are completely unknown, as in VLBI, a starting model must be obtained by least-squares model-fitting techniques. We have found that it is important to choose a model that fits the closure phases well; the fit to the amplitudes is less critical. This seems to be related to the observation mentioned earlier that phase carries more information about the source structure than does amplitude. It also helps to adjust only the phases in the first few iterations of self-calibration, and to only adjust the amplitudes when a good fit to the closure phases has been obtained. In a typical observation with the VLA, two iterations of phase adjustment, followed by two iterations of amplitude and phase adjustment, will be sufficient for the self-calibration process to converge. In a VLBI experiment, with larger initial errors, poorer *uv*-coverage, and a poorer starting model, 20 iterations may be needed for a complex source. Perhaps the most critical requirement for convergence is that the map-plane constraints be applied rigorously to exclude negative features, and features outside as small a window as possible, from the trial map (30). This is particularly important when the *uv*-coverage is poor. The window can be chosen in a variety of ways (for example, by inspection of maps made with other instruments or at other frequencies, or from the known field of view of the instrument). An objective method is to examine the autocorrelation function of the map, made by Fourier transformation of the squares of the visibility amplitudes (6, 38). Use of a small window in early iterations and a larger one in later iterations has sometimes been found to be helpful. These and other practical considerations have been discussed by Cornwell & Wilkinson (27, 30, 124).

6.4 *Closure Errors*

The self-calibration techniques discussed above are all based on the assumption that the gain and phase errors can be factorized by antenna. In practice there are instrumental errors that cannot be factorized. These "closure errors" may be regarded as systematic errors in the measurements of closure phase and closure amplitude. Unless the closure errors can somehow be removed, they are the limiting factor in the quality of the maps produced by self-calibration methods. A fractional error f in the estimated gains will introduce sidelobes on the map at a level of about $f/\sqrt{M}$, where M is the number of independent errors. For example, with constant 1% errors on 10 baselines, the dynamic range of the map would be limited to about 300:1. Thus, to achieve a dynamic range of 1000:1 or 10,000:1 (comparable to the best VLA maps), it is important to keep uncorrectable closure errors well below a level of 1% in gain or 0.5° in phase.

We should distinguish errors that remain constant in time from those

that vary. Constant errors can in principle be measured by observations of a point source and then eliminated in the data processing. Alternatively, if the additional number of unknowns is not large, they can be solved for in the hybrid mapping in the same way as the station gains (29). But if the errors are variable on a short time scale, they can be neither measured nor corrected, and thus they can only be tackled by improving the quality of the equipment.

A number of causes of closure errors in existing interferometer arrays have been identified (22). Some errors may arise in the hardware of the correlators themselves and in the subsequent electronics and software; but if these are identical for all the correlators in the array, they may be ignored. Thompson & D'Addario (113) have shown that the factorization of gains by antenna is, in general, not possible if the antenna transfer functions are different functions of frequency, and they have investigated in detail the effects of mismatched transfer functions. Usually, identical rectangular bandpasses are used at each antenna; as an example of the sort of error that can occur, a shift of the center frequency by 0.7% of the bandwidth can introduce a gain closure error of 1%. It is clear that when self-calibration methods are to be used, the tolerances on bandpass distortions are very stringent. In present VLBI systems, designed before the advent of self-calibration, large bandpass mismatches are not uncommon and are probably a limiting factor in image quality. Another serious problem in present VLBI systems has been pointed out by Rogers (91): a phase closure error results if the antennas differ in their polarization purity; for example, a 0.5% error in polarization purity can cause a 1° phase closure error.

If the antenna gains are varying on a short time scale, one must be careful not to average the data for too long a period, as this will introduce closure errors (124). This can be a problem in VLBI, where the coherence time of the atmosphere and oscillators is short and where there is a temptation to average for as long as possible to improve the signal-to-noise ratio.

We have implicitly assumed that the complex antenna gains are the same for all parts of the source being observed. In VLBI, the field of view is small, and this is usually a good assumption, but it breaks down when mapping large fields. For example, the atmosphere may not be homogeneous across the source. It is usually assumed that the atmosphere is "isoplanatic" over regions $\sim 1°$ at radio wavelengths, but this is not true in bad weather. Also, if the antenna pointing is not constant, gain errors are introduced that vary across the field of view, with the worst errors occurring where the slope of the antenna pattern is greatest. These are difficult to correct by self-calibration, though it may be possible to make some progress by self-calibrating different isoplanatic patches independently.

7. THE FUTURE

The development of self-calibration methods for radio astronomy has already proved extremely valuable. By eliminating sidelobes due to unavoidable amplitude and phase errors, it has enabled arrays such as the VLA, the WSRT, and MERLIN to far surpass initial expectations of their imaging quality; and it has been the single most important factor in transforming VLBI into a true imaging technique.

In principle the techniques of self-calibration are not confined to radio astronomy, and we can expect to see them applied at other frequencies in the future. Michelson (68) demonstrated the feasibility of optical astronomical interferometers in 1920, but the effects of atmospheric turbulence impose severe limitations on this technique. Problems arise because both the time scale of atmospheric fluctuations is short (10 ms) and the transverse coherence length of the wavefront is small (10 cm). If we assume an aperture of diameter 10 cm, an integration time of 10 ms, and a bandwidth of 10 nm, and if we allow for the imperfect photon efficiency of even the best CCD detectors, the limiting magnitude of a ground-based Michelson interferometer is 8. There are a number of projects in progress to construct ground-based Michelson interferometers (e.g. 63, 108), but so far these efforts have met with limited success due to the stringent rigidity requirements and the difficulty of tracking fringes. It would clearly be of great value to be able to use self-calibration methods with such instruments. Although methods have been suggested for measuring closure phase with optical arrays (46, 88, 93), they have not yet been applied in astronomy. The application of self-calibration to heterodyne infrared interferometers (61, 115) will be more straightforward, since the techniques involved are similar to those of radio interferometry.

In order to observe faint astronomical objects, apertures larger than 10 cm are needed; transverse coherence then no longer obtains across the aperture, and the image is spread out over an angle λ/d, where d is the transverse coherence length (about 10 cm). A short exposure yields a "speckle pattern" image, which is caused by interference of rays from different portions of the wavefront. Labeyrie (62) first pointed out that these speckle patterns can be used to derive source structure down to the resolution limit of the aperture. That this must be so can be seen from the fact that the speckle pattern is formed in part by the interference of rays from the extreme edges of the aperture, and these rays clearly contain information about the transverse coherence due to the source over these distances. If the instantaneous telescope response to a point source, the so-called instantaneous "point-spread function," were known, it would be

possible to obtain an image of the object by deconvolution. However, only the time-averaged point-spread function is known, which makes it difficult to extract phase information, although it is easy to construct the autocorrelation function of the image.

If two large apertures are operated as an interferometer, the speckles are modulated by interference fringes. Thus it might at first sight seem possible to use closure to extract phase information from these fringes. However, there is a random phase relationship between the fringes in different speckles. It might prove possible to remove this random component by the techniques of adaptive optics (e.g. 20, 51, 71). These techniques are closely related to the self-calibration methods used in radio astronomy (30): the adjustment of the phase of small elements of the telescope pupil is analogous to the adjustment of the complex antenna gains in a radio array. Unfortunately, the adjustment must be made in the 10-ms time scale of the atmospheric fluctuations, which limits the technique to bright objects. Radio astronomers have the advantage that they can record the visibility measurements and make the gain adjustments in a computer after all the measurements have been accumulated.

In the next decade we may expect to see new instruments developed to exploit the possibilities opened up by self-calibration in the radio and other wavebands. In the United States, the Astronomy Survey Committee of the National Research Council (73) has recommended the construction of a Very Long Baseline Array (VLBA) of radio telescopes designed to produce images with an angular resolution of 0.3 milliarcsec, and it is hoped that construction of this instrument will begin in 1984. Similar projects are under way in Canada (64) and Australia. The Astronomy Survey Committee has also recommended placing a VLBI antenna in space to increase the resolving power of ground-based arrays. We can also anticipate the development of large millimeter-wavelength arrays and infrared and optical interferometers on the ground or in space. All of these projects will rely on the self-calibration methods developed for radio astronomy to make high-quality images.

ACKNOWLEDGMENTS

We thank our colleagues who contributed to this review by discussion and comment, and by sending us preprints of their work. We are especially grateful to J. W. Dreher, D. H. Hough, T. W. B. Muxlow, R. T. Schilizzi, and R. S. Simon for providing the examples of self-calibration that appear in the figures. This work was supported in part by the US National Science Foundation (AST 8210259).

Literature Cited

1. Ables, J. G. 1974. *Astron. Astrophys. Suppl.* 15:383–93
2. Armstrong, J. W., Sramek, R. A. 1982. *Radio Sci.* 17:1579–86
3. Baars, J. W. M., Houghoudt, B. G. 1974. *Astron. Astrophys.* 31:323–31
4. Baker, P. L. 1981. *Astron. Astrophys.* 94:85–90
5. Baldwin, J. E., Warner, P. J. 1976. *MNRAS* 175:345–53
6. Baldwin, J. E., Warner, P. J. 1978. *MNRAS* 182:411–22
7. Baldwin, J. E., Warner, P. J. 1979. See Ref. 118, pp. 67–82
8. Bare, C., Clark, B. G., Kellermann, K. I., Cohen, M. H., Jauncey, D. L. 1967. *Science* 157:189–91
9. Bates, R. H. T. 1982. *Optik* 61:247–62
10. Boerner, W.-M., Jordan, A. K., Kay, I. W. 1981. *IEEE Trans. Antennas Propag.* AP-29:185–89
11. Born, M., Wolf, E. 1975. *Principles of Optics.* Oxford: Pergamon. 808 pp. 5th ed.
12. Bracewell, R. N. 1958. *Proc. IRE* 46: 97–105
13. Bracewell, R. N. 1979. *Ann. Rev. Astron. Astrophys.* 17:113–34
14. Branham, R. L. 1982. *Astron. J.* 87: 928–37
15. Bridle, A. H., Perley, R. A. 1984. *Ann. Rev. Astron. Astrophys.* 22:319–58
16. Broten, N. W., Legg, T. H., Locke, J. L., McLeish, C. W., Richards, R. S., et al. 1967. *Science* 156:1592–93; *Nature* 215:38
17. Brouw, W. N. 1975. In *Methods in Computational Physics*, ed. B. Alder, S. Fernbach, M. Rotenberg, 14:131–75. New York: Academic
18. Bruck, Yu. M., Sodin, L. G. 1979. *Opt. Commun.* 30:304–8
19. Bryan, R. K., Skilling, J. 1980. *MNRAS* 191:69–79
20. Buffington, A., Crawford, S., Pollaine, S. M., Orth, C. D., Muller, R. A. 1978. *Science* 200:489–94
21. Clark, B. G. 1980. *Astron. Astrophys.* 89:377–78
22. Clark, B. G. 1981. *VLA Sci. Memo. No. 137*, Natl. Radio Astron. Obs., Charlottesville, Va. 10 pp.
23. Cohen, M. H. 1969. *Ann. Rev. Astron. Astrophys.* 7:619–64
24. Cohen, M. H. 1973. *Proc. IEEE* 61: 1192–97
25. Cornwell, T. J. 1981. *VLA Sci. Memo. No. 135*, Natl. Radio Astron. Obs., Charlottesville, Va. 4 pp.
26. Cornwell, T. J. 1982. See Ref. 114, Lect. No. 9
27. Cornwell, T. J. 1982. See Ref. 114, Lect. No. 13
28. Cornwell, T. J. 1983. *Astron. Astrophys.* 121:281–85
29. Cornwell, T. J., Wilkinson, P. N. 1981. *MNRAS* 196:1067–86
30. Cornwell, T. J., Wilkinson, P. N. 1984. *Proc. URSI/IAU Symp. Meas. Process. Indirect Meas., Sydney, 1983*, ed. J. A. Roberts. Cambridge: Cambridge Univ. Press. In press
31. Cotton, W. D. 1979. *Astron. J.* 84: 1122–28
32. D'Addario, L. R. 1980. *Proc. Soc. Photo-Opt. Instrum. Eng.* 231:2–9
33. Davies, J. G., Anderson, B., Morison, I. 1980. *Nature* 288:64–66
34. Dreher, J. W., Feigelson, E. D. 1984. *Nature* 308:43–45
35. Elgaroy, O., Morris, D., Rowson, B. 1962. *MNRAS* 124:395–403
36. Fiddy, M. A., Brames, B. J., Dainty, J. C. 1983. *Opt. Lett.* 8:96–98
37. Fienup, J. R. 1978. *Opt. Lett.* 3:27–29
38. Fienup, J. R. 1982. *Appl. Opt.* 21:2758–69
39. Fomalont, E. B. 1973. *Proc. IEEE* 61:1211–18
40. Fomalont, E. B., Wright, M. C. H. 1974. In *Galactic and Extragalactic Radio Astronomy*, ed. G. L. Verschuur, K. I. Kellermann, pp. 256–90. New York: Springer-Verlag
41. Fort, D. N., Yee, H. K. C. 1976. *Astron. Astrophys.* 50:19–22
42. Frieden, B. R. 1972. *J. Opt. Soc. Am.* 62:511–18
43. Genzel, R., Reid, M. J., Moran, J. M., Downes, D. 1981. *Ap. J.* 244:884–902
44. Gerchberg, R. W. 1974. *Opt. Acta* 21:709–20
45. Gerchberg, R. W., Saxton, W. O. 1972. *Optik* 35:237–46
46. Goodman, J. W. 1970. In *Progress in Optics*, ed. E. Wolf, 8:1–50. Amsterdam: North-Holland
47. Gorenstein, M. V., Shapiro, I. I., Cohen, N. L., Corey, B. E., Falco, E. E., et al. 1983. *Science* 219:54–56
48. Gull, S. F., Daniell, G. J. 1978. *Nature* 272:686–90
49. Gull, S. F., Daniell, G. J. 1979. See Ref. 118, pp. 219–25
50. Hamaker, J. P. 1978. *Radio Sci.* 13: 873–91
51. Hardy, J. W. 1978. *Proc. IEEE* 66:651–97
52. Heeschen, D. S. 1981. In *Telescopes for the 1980s*, ed. G. Burbidge, A. Hewitt, pp. 1–61. Palo Alto, Calif: Ann. Rev. Inc.

53. Hinder, R., Ryle, M. 1971. *MNRAS* 154:229–53
54. Hjellming, R. M., Bignell, R. C. 1982. *Science* 216:1279–85
55. Högbom, J. A. 1974. *Astron. Astrophys. Suppl.* 15:417–26
56. Högbom, J. A., Brouw, W. N. 1974. *Astron. Astrophys.* 33:289–301
57. Huiser, A. M. J., van Toorn, P. 1980. *Opt. Lett.* 5:499–501
58. Ishiguro, M. 1974. *Astron. Astrophys. Suppl.* 15:431–43
59. Jennison, R. C. 1958. *MNRAS* 118: 276–84
60. Jennison, R. C. 1966. *Introduction to Radio Astronomy.* London: Newnes, New York: Philos. Libr., Inc. 160 pp.
61. Johnson, M. A., Betz, A. L., Townes, C. H. 1974. *Phys. Rev. Lett.* 33:1617–20
62. Labeyrie, A. 1970. *Astron. Astrophys.* 6:85–87
63. Labeyrie, A. 1978. *Ann. Rev. Astron. Astrophys.* 16:77–102
64. Legg, T. H. 1982. *Phys. Can.* (*La Physique au Canada*) 38:3–7
65. Lonsdale, C. J., Morison, I. 1983. *MNRAS* 203:833–51
66. Mathur, N. C., Grossi, M. D., Pearlman, M. R. 1970. *Radio Sci.* 5:1253–61
67. McReady, L. L., Pawsey, J. L., Payne-Scott, R. 1947. *Proc. R. Soc. London Ser. A* 190:357–75
68. Michelson, A. A. 1920. *Ap. J.* 51:257–62
69. Moran, J. M. 1976. In *Methods of Experimental Physics*, ed. M. L. Meeks, 12C:228–60. New York: Academic
70. Moran, J. M., Papadopoulos, G. D., Burke, B. F., Lo, K. Y., Schwartz, P. R., et al. 1973. *Ap. J.* 185:535–67
71. Muller, R. A., Buffington, A. 1974. *J. Opt. Soc. Am.* 64:1200–10
72. Napier, P. J. 1982. See Ref. 114, Lect. No. 3
73. National Research Council. 1982. *Astronomy and Astrophysics for the 1980s. Volume 1: Report of the Astronomy Survey Committee.* Washington DC: Natl. Acad. Press. 189 pp.
74. Nityananda, R., Narayan, R. 1982. *J. Astrophys. Astron.* 3:419–50
75. Noordam, J. E. 1981. *Proc. ESO Conf. Sci. Importance High Angular Resolut. Infrared Opt. Wavelengths, Garching*, ed. M. H. Ulrich, K. Kjär, pp. 257–61. Garching bei München, FRG: ESO
76. Noordam, J. E., de Bruyn, A. G. 1982. *Nature* 299:597–600
77. Norris, R. P. 1983. *Proc. Int. Conf. VLBI Tech., Toulouse, Fr., 1982*, pp. 341–44. Toulouse: Cent. Natl. Étud. Spat. 488 pp.
78. Oldenburg, D. W. 1976. *Geophys. J. R. Astron. Soc.* 44:413–31
79. Oppenheim, A. V., Lim, J. S. 1981. *Proc. IEEE* 69:529–41
80. Palmer, H. P., Rowson, B., Anderson, B., Donaldson, W., Miley, G. K., et al. 1967. *Nature* 213:789–90
81. Parker, R. L. 1977. *Ann. Rev. Earth Planet. Sci.* 5:35–64
82. Peckham, R. J. 1973. *MNRAS* 165:25–38
83. Readhead, A. C. S., Hough, D. H., Ewing, M. S., Walker, R. C., Romney, J. D. 1983. *Ap. J.* 265:107–31
84. Readhead, A. C. S., Napier, P. J., Bignell, R. C. 1980. *Ap. J. Lett.* 237: L55–60
85. Readhead, A. C. S., Walker, R. C., Pearson, T. J., Cohen, M. H. 1980. *Nature* 285:137–40
86. Readhead, A. C. S., Wilkinson, P. N. 1978. *Ap. J.* 223:25–36
87. Reid, M. J., Haschick, A. D., Burke, B. F., Moran, J. M., Johnston, K. J., Swenson, G. W. 1980. *Ap. J.* 239:89–111
88. Rhodes, W. T., Goodman, J. W. 1973. *J. Opt. Soc. Am.* 63:647–57
89. Rogers, A. E. E. 1976. In *Methods of Experimental Physics*, ed. M. L. Meeks, 12C:139–57. New York: Academic
90. Rogers, A. E. E. 1980. *Proc. Soc. Photo-Opt. Instrum. Eng.* 231:10–17
91. Rogers, A. E. E. 1983. *VLB Array Memo. No. 253*, Natl. Radio Astron. Obs., Charlottesville, Va. 2 pp.
92. Rogers, A. E. E., Hinteregger, H. F., Whitney, A. R., Counselman, C. C., Shapiro, I. I., et al. 1974. *Ap. J.* 193:293–301
93. Rogstad, D. H. 1968. *Appl. Opt.* 7:585–88
94. Rogstad, D. H., Shostak, G. S. 1971. *Astron. Astrophys.* 13:99–107
95. Ryle, M. 1952. *Proc. R. Soc. London Ser. A* 211:351–75
96. Ryle, M. 1975. Nobel Lecture delivered in Stockholm on 12 Dec. 1974. In *Les Prix Nobel en 1974*. Stockholm: P. A. Norstedt & Söner. Reprinted in *Science* 188:1071–79 and *Rev. Mod. Phys.* 47:557–66
97. Ryle, M., Hewish, A. 1960. *MNRAS* 120:220–30
98. Ryle, M., Vonberg, D. D. 1946. *Nature* 158:339–40
99. Sanroma, M., Estalella, R. 1983. *Proc. Int. Conf. VLBI Tech., Toulouse, Fr., 1982*, pp. 391–407. Toulouse: Cent. Natl. Étud. Spat. 488 pp.
100. Schafer, R. W., Mersereau, R. M., Richards, M. A. 1981. *Proc. IEEE* 69: 432–50

101. Schilizzi, R. T., de Bruyn, A. G. 1983. *Nature* 303:26–31
102. Schwab, F. R. 1980. *Proc. Soc. Photo-Opt. Instrum. Eng.* 231:18–25
103. Schwab, F. R. 1981. *VLA Sci. Memo. No. 136*, Natl. Radio Astron. Obs., Charlottesville, Va. 20 pp. and Erratum (Feb. 1982)
104. Schwab, F. R., Cotton, W. D. 1983. *Astron. J.* 88:688–94
105. Schwarz, U. J. 1978. *Astron. Astrophys.* 65:345–56
106. Schwarz, U. J. 1979. See Ref. 118, pp. 261–75
107. Schwarz, U. J., Cole, D. J., Morris, D. 1973. *Aust. J. Phys.* 26:661–73
108. Shao, M., Staelin, D. H. 1980. *Appl. Opt.* 19:1519–22
109. Simon, R. S., Readhead, A. C. S., Wilkinson, P. N. 1984. *Proc. IAU Symp. 110, VLBI and Compact Radio Sources.* Dordrecht: Reidel. In press
110. Smith, F. G. 1952. *Proc. Phys. Soc. London Sect. B* 65:971–80
111. Swenson, G. W., Mathur, N. C. 1968. *Proc. IEEE* 56:2114–30
112. Thompson, A. R., Clark, B. G., Wade, C. M., Napier, P. J. 1980. *Ap. J. Suppl.* 44:151–67
113. Thompson, A. R., D'Addario, L. R. 1982. *Radio Sci.* 17:357–69
114. Thompson, A. R., D'Addario, L. R., eds. 1982. *Synthesis Mapping, Proc. NRAO–VLA Workshop, Socorro, N. Mex.*, Green Bank, W. Va: Natl. Radio Astron. Obs.
115. Townes, C. H., Sutton, E. C. 1981. *Proc. ESO Conf. Sci. Importance High Angular Resolut. Infrared Opt. Wavelengths, Garching*, ed. M. H. Ulrich, K. Kjär, pp. 199–223. Garching bei München, FRG: ESO
116. Twiss, R. W., Carter, A. W. L., Little, A. G. 1960. *Observatory* 80:153–59
117. Twiss, R. W., Carter, A. W. L., Little, A. G. 1962. *Aust. J. Phys.* 15:378–86
118. Van Schooneveld, C., ed. 1979. *Proc. IAU Colloq. 49, Image Formation from Coherence Functions in Astronomy, Groningen, The Neth., 1978.* Dordrecht: Reidel. 340 pp.
119. Walker, R. C. 1981. *Astron. J.* 86:1323–31
120. Walker, R. C., Burke, B. F., Haschick, A. D., Crane, P. C., Moran, J. M., et al. 1978. *Ap. J.* 226:95–114
121. Wells, D. C. 1980. *Proc. Soc. Photo-Opt. Instrum. Eng.* 264:148–56
122. Wernecke, S. J. 1977. *Radio Sci.* 12: 831–44
123. Wernecke, S. J., D'Addario, L. R. 1977. *IEEE Trans. Comput.* C-26:351–64
124. Wilkinson, P. N. 1983. *Proc. Int. Conf. VLBI Tech., Toulouse, Fr., 1982*, pp. 375–89. Toulouse: Cent. Natl. Étud. Spat. 488 pp.
125. Wilkinson, P. N., Readhead, A. C. S. 1979. See Ref. 118, pp. 83–91
126. Wilkinson, P. N., Readhead, A. C. S., Purcell, G. H., Anderson, B. 1977. *Nature* 269:764–68