AT/25.1.1/011

SELFCAL and the AT

Ray Norris and Mike Kesteven

25 November 1984

1.0 SUMMARY

Recent tests of the SELFCAL error correction algorithm on simulated AT data revealed an unexpected problem: in some circumstances the maps were symmetrised, so that some components of the source model would appear reflected into the other half of the map. Although we have not yet reached a full understanding of this problem, its causes, and cures, we have achieved an empirical understanding, and suggest possible techniques which may solve the problem. This document is a progress report in which we take stock of what we now know, and invite comments which may guide our subsequent exploration.

1.1 Introduction

When a radio interferometer observes a source, the data are subject to errors which are predominantly telescope-based (i.e. depend only on the telescope, and not on the baseline). Such errors may typically be caused by atmosphere, ionosphere, or receivers, and may appear as errors in either phase or amplitude. Generally speaking, phase errors have a greater effect than amplitude errors on the quality of the final map, and it is these phase errors with which this report is primarily concerned. Telescope-based phase errors may be corrected by:

1. Phase referencing to a nearby unresolved source
2. Use of redundancy techniques.
3. Various algorithms (closure phase, hybrid mapping, etc.) which are known collectively as SELFCAL.

SELFCAL techniques have been widely used, with much success, on many arrays (e.g. VLA, MERLIN, VLBI) for some time, and are generally regarded as reliable and robust. The implementation used for the tests described here is the AIPS program ASCAL (or its derivative BSCAL).

ASCAL is used as follows. It is first given a model of the source. This model may be as simple as a point source, or as complex as a CLEANed map of the source. Then, for each integration period, ASCAL tries to minimise the difference bewteen the actual visiblity data and the predicted data from the model, by introducing complex telescope gain errors. These gain errors are then used to correct the data, which are then transformed and CLEANed. Any subsequent iteration will generally use this CLEAN map (or a portion of it) as a model for ASCAL.

A common strategy is to use a point source for the initial model. ASCAL will then try to incorporate any phase variations in the observed data into telescope errors, to leave 'corrected' data with zero phase. This will tend to symmetrise the map, but on the next iteration, when the CLEAN map (or, usually, only the CLEAN components down to the first negative component) is used as a model, the symmetry will be reduced, and will disappear in subsequent iterations.

2.0 THE INITIAL PROBLEM

We first became aware of a problem when simulated multi-day AT observations of the test source SPIRAL (see Figure 1) produced a stable symmetrical solution after a point source had been used for the initial model. Subsequent iterations failed to remove this symmetry. To investigate this effect, a series of tests were run using 1-day data in order to reduce CPU time. Figure 2 shows the effect of CLEANing the raw data, without applying SELFCAL. Figure 3 shows the effect of several iterations of SELFCAL to the same data, to which no telescope errors have been added. The SELFCALled solution is clearly symmetrised, and it retains this symmetry through subsequent iterations.

Further tests appeared to show that, if a model consisting of the components of the test source down to a flux limit F was used, then all flux below F was symmetrised (see Figure 4). However, we are less confident of this result because of the low dynamic range of the maps. Furthermore, this effect seemed sensitive to the complexity of the model and the quality of the uv data.

Subsequent tests were run on the simpler test source WONKY (Figures 5 and 6). In this case, the map would always converge to the correct solution, but it would often do so very slowly after having first gone to the symmetrical solution. However, the tests showed that this problem seemed to be exacerbated by the linear AT array: the problem was reduced in its severity if natural (rather than uniform) weighting was used, if the integration interval for ASCAL was increased, if the number of antennas was increased, or if the antennas were redistributed in a two-dimensional array.

A hand-waving explanation of this phenomenon might run along the following lines. For a one-dimensional array, each gain solution (corresponding to one integration period) would sample only one direction in the uv plane. The difference between two possible solutions (e.g. the model described in the next section) might appear only at one position angle, and thus affect only a small proportion of the independant gain solutions.

3.0 A SIMPLER EXAMPLE

To illustrate the problem, a very simple example is chosen. Consider a source consiting of a 1Jy point source at the phase centre, plus a 1Jy point source offset a distance Δ (Figure 7a). This model will be called solution a. The observed visiblity on a baseline $\underline{s}$ will have a phase θ, and an amplitude of $(1+\cos\theta)$, where $\theta=2\pi\Delta.\underline{s}$. If a point source model is given to ASCAL, ASCAL will find an exact solution (solution b) in which the observed phase is identically zero, and the gain of telescope at position $\underline{s}_1$ is $2\pi\Delta.\underline{s}_1$. Note that this is equivalent to saying that the closure phase of either solution is identically zero. When mapped, this will give the image shown in Figure 7b, corresponding to a visibility amplitude of $\sin(\theta)/\sin(\theta/2)$. The amplitudes for the two models are plotted in Figure 8. Note that, since the closure phase is zero in either case, the models can be distinguished only by their amplitudes, which differ significantly only in the region of $\theta=\pi$. In subsequent iterations, ASCAL will push towards solution b, and CLEAN will push towards solution a. The winner depends on the signal to noise ratio, uv coverage, etc. Clearly, there will be circumstances in which solution b will be a stable solution.

Examples of two trials using different signal-to-noise ratios are given in Figures 9 and 10.

4.0 POSSIBLE CURES

The example above can easily be made to converge on the correct solution by limiting the telescope errors in any one iteration of ASCAL. An elegant way of implementing this would be to do a constrained minimisation when solving for gains. An inelegant way, and the way tried in the tests, is simply to replace any phase errors $>10^{o}$ by a 10^{o} error. However, it is not clear whether this will work for more complex source distributions. It may also fail in the presence of real large telescope errors (such as may be encountered at Q band, etc.)

Alternatively, preliminary tests indicate that use of an uncalibrated CLEAN map, down to the first negative component, is unlikely to converge on the symmetrical solution however bad that initial map is. Clearly, however, this method will also fail if the data are so bad as to produce an uncalibrated map with only one component above the negative component level.

5.0 CONCLUSIONS

Our tentative conclusions are:

1. In most circumstances, SELFCAL will work on the AT, but
2. Care will be required in the choice of the initial model
3. This suggests that fairly frequent observations of calibration sources will be needed.
4. The computing requirements are likely to be larger than originally anticipated, since several iterations are generally needed.
5. On the wider astronomical front, there is some cause for worry as to whether maps produced on sparsely filled arrays such as MERLIN or VLBI, for which an initial point source model has been used for SELFCAL, may have been symmetrised.

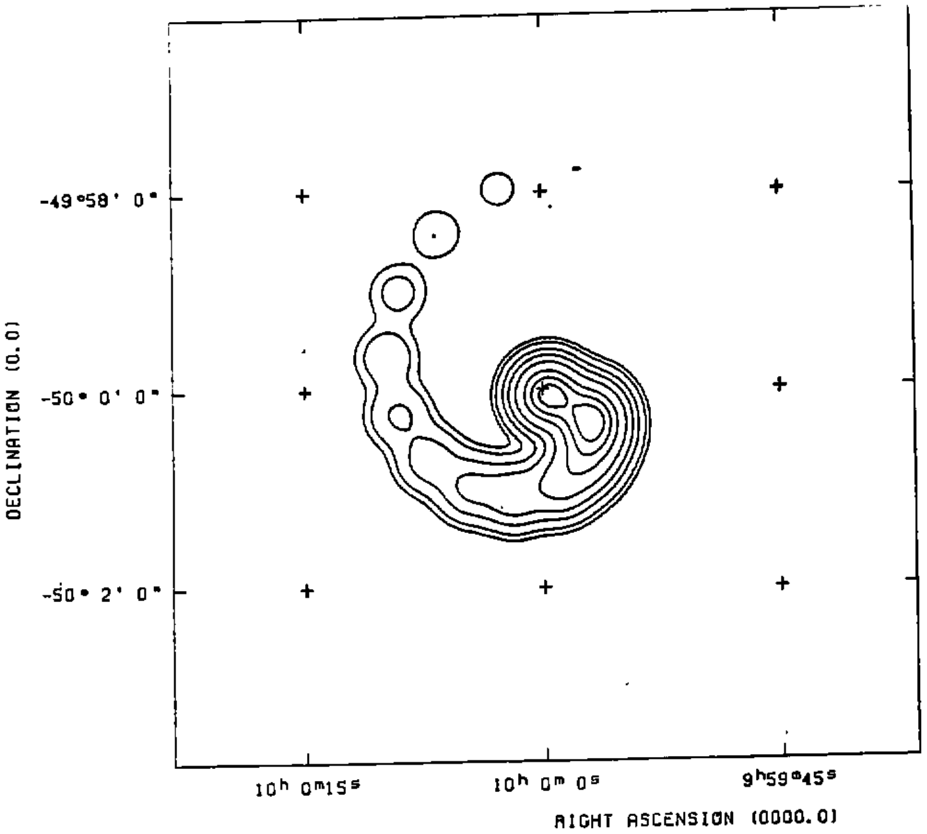

Figure 1 Model distribution of the test source SPIRAL. In this and all other maps (except where stated otherwise), contours are plotted at intervals of -2, -1, 1, 2, 5, 10, 20, 40, 60, 80% of the peak flux.

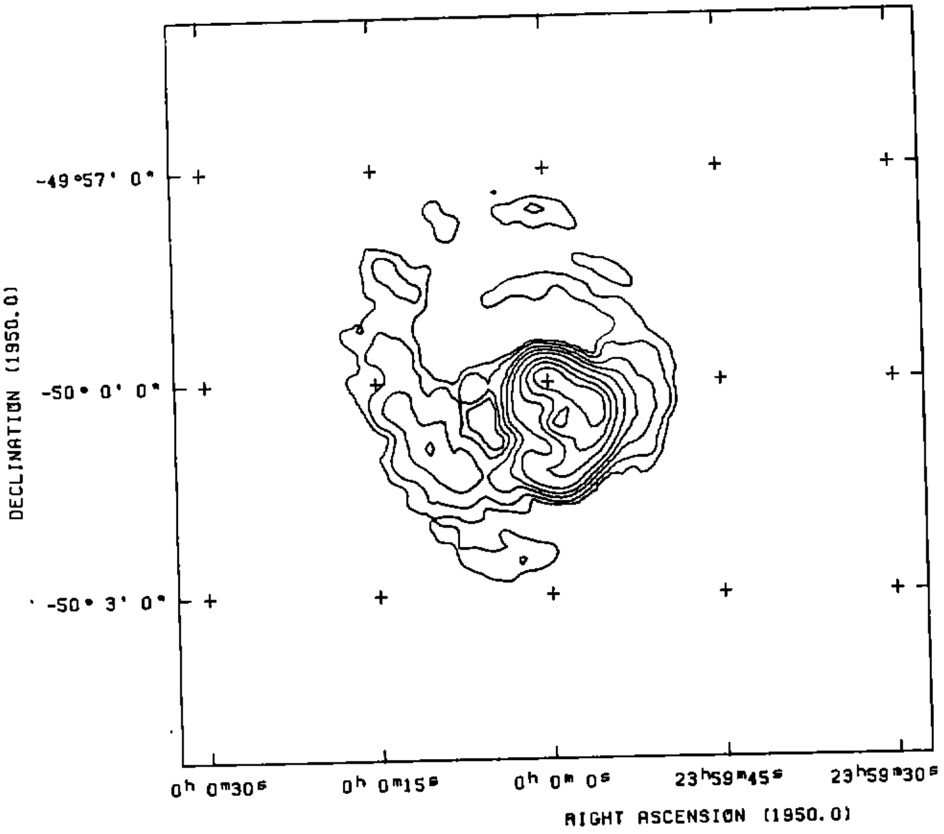

Figure 2 The CLEANed raw data from a simulated one-day observation of SPIRAL.

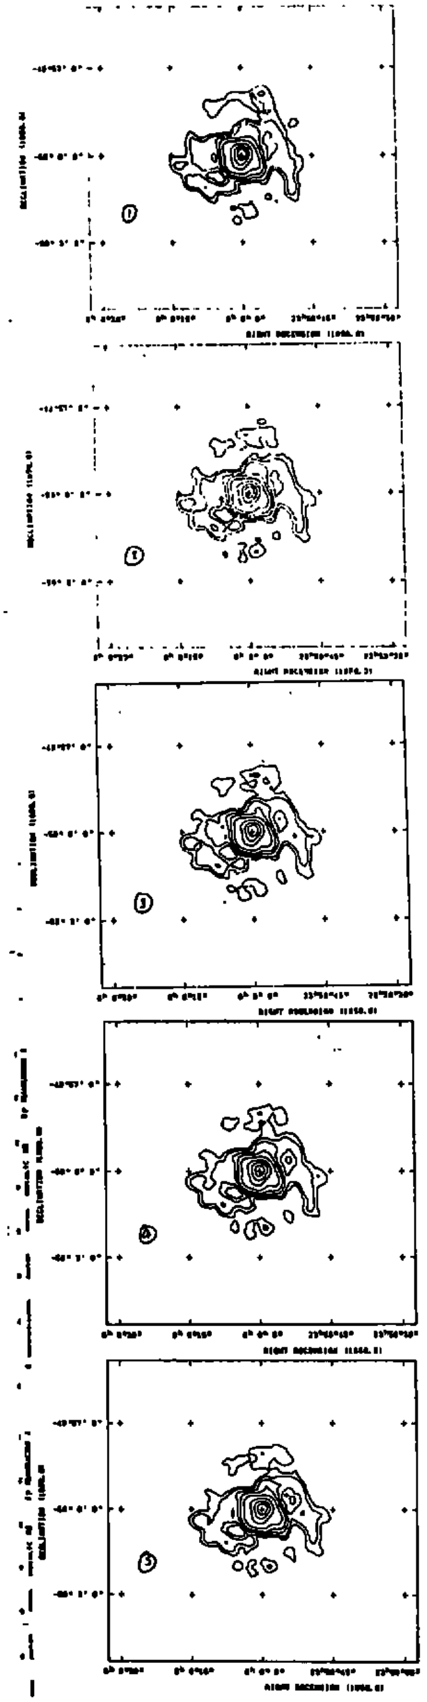

Figure 3 The ASCALled maps from the same data as for Figure 2. In this and all other maps, the array DR3D4 has been used with system noise but no telescope errors added, and ASCAL has been used on phase only. The first iteration of ASCAL (top map) used a point source model, and subsequent iterations each used the previous CLEANed map down to the first negative contour. Note how the maps are all roughly symmetrical.

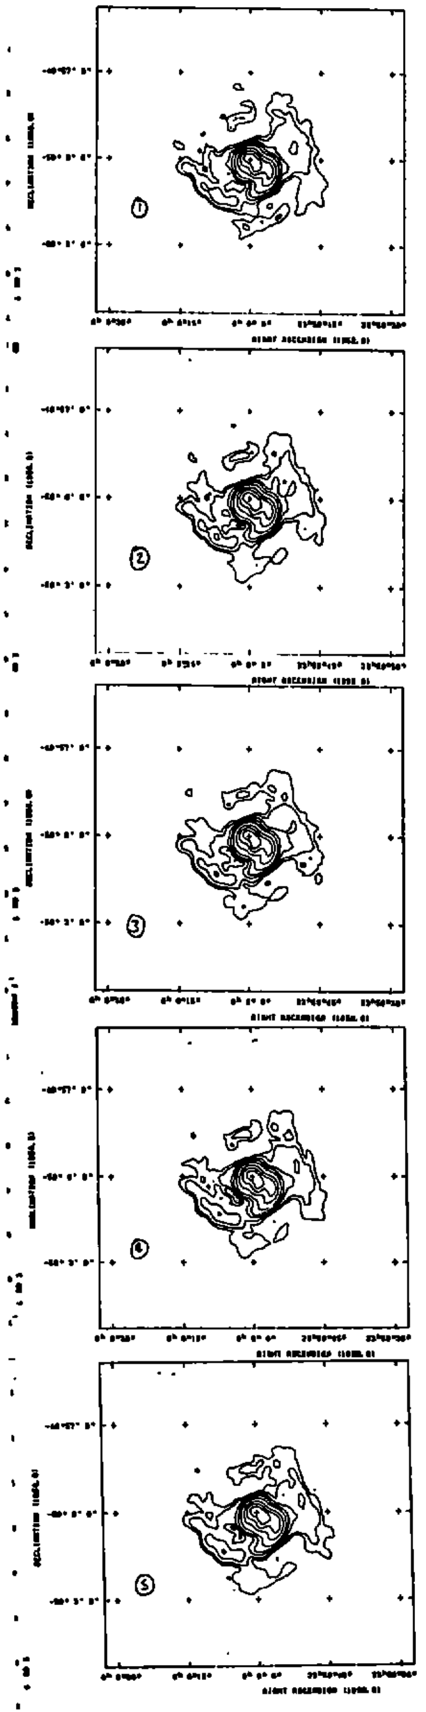

ITERATION
#1

#2

#3

#4

#5

Figure 4 The ASCALled maps from the same data as for Figures 2 and 3. In this case, however, the model used for the initial iteration consisted of the first 50 components of Figure 2. Note that although the symmetrical components near the centre have been eradicated, much of the weaker symmetrical structure appears to remain.

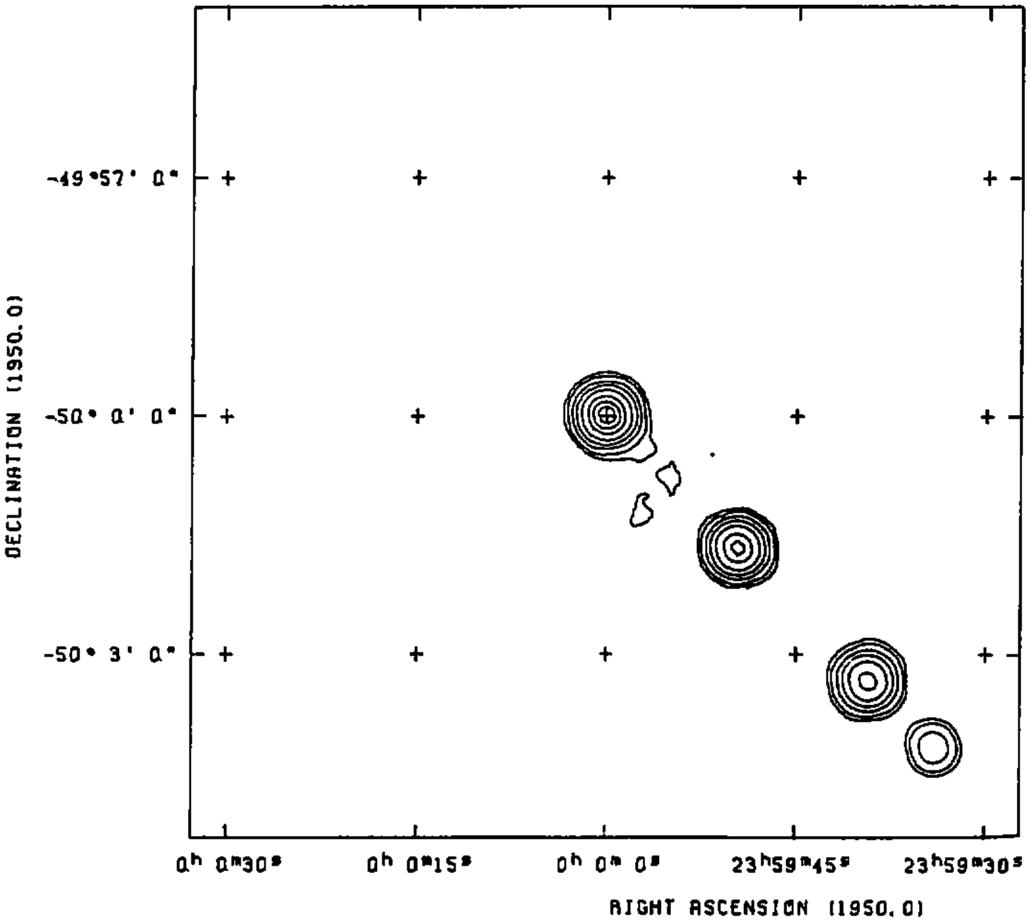

Figure 5 The CLEANed raw data from a simulated one-day observation of the source WONKY.

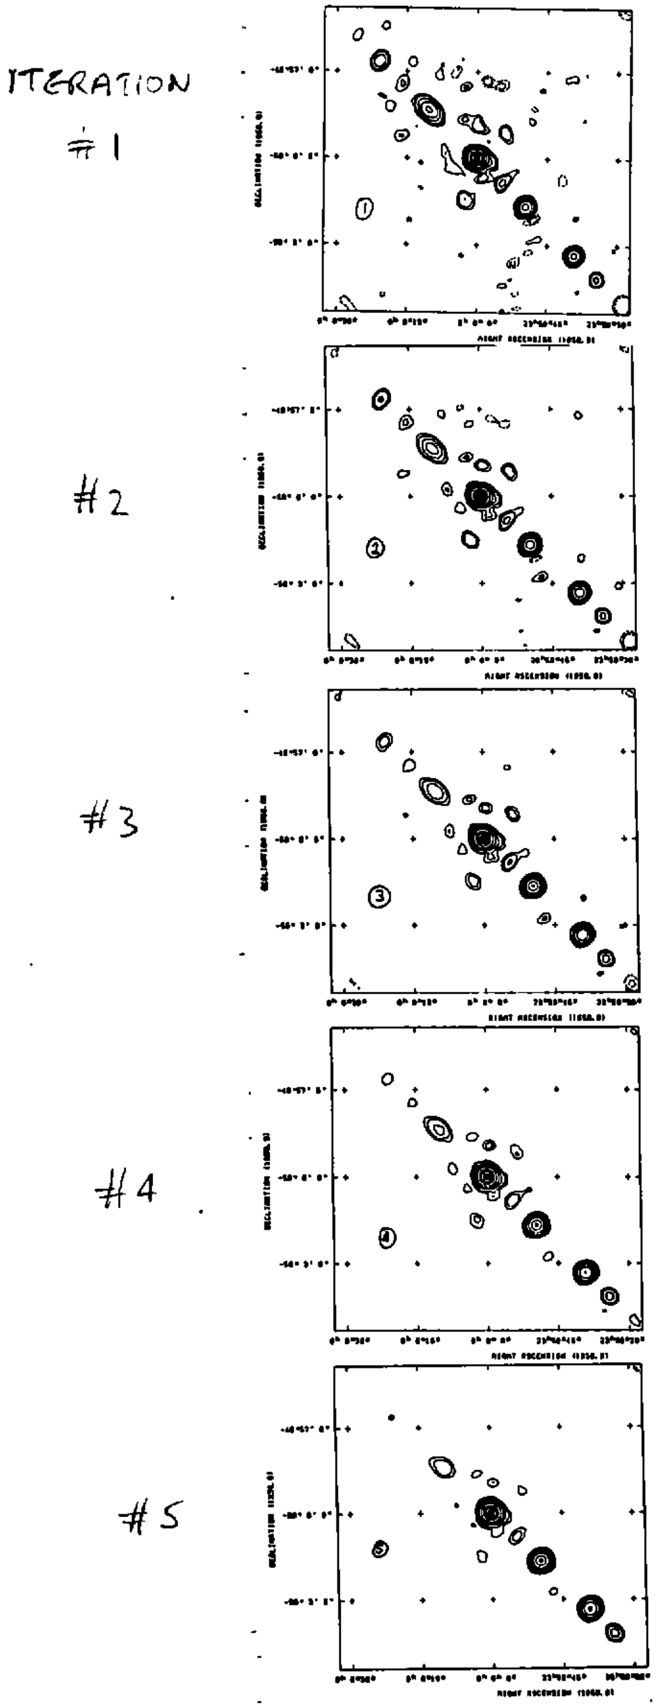

Figure 6 The ASCALled maps from the same data as for Figure 5. The first iteration of ASCAL (top map) used a point source model, and subsequent iterations each used the previous CLEANed map down to the first negative contour. In the case of this simpler model, the process does seem to be converging (although slowly) on the correct solution.

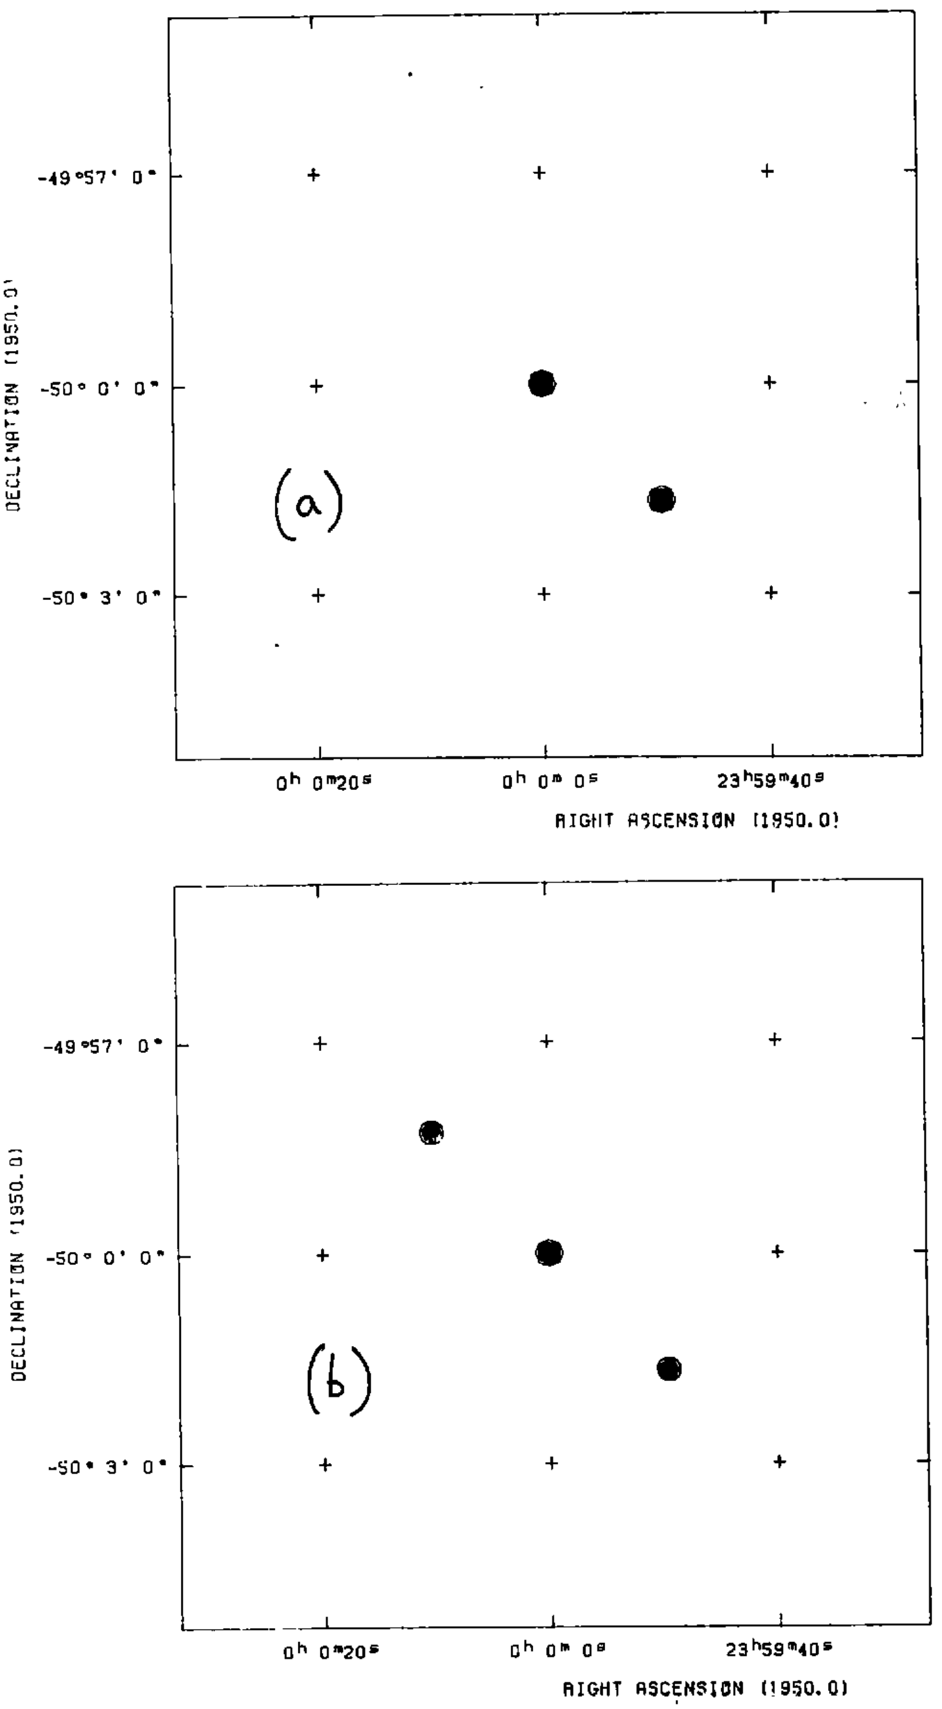

<u>Figure</u> <u>7</u> (a) Model sky for the double source discussed in the text (Model a). (b) Model sky for the symmetrical source discused in the text (Model b).

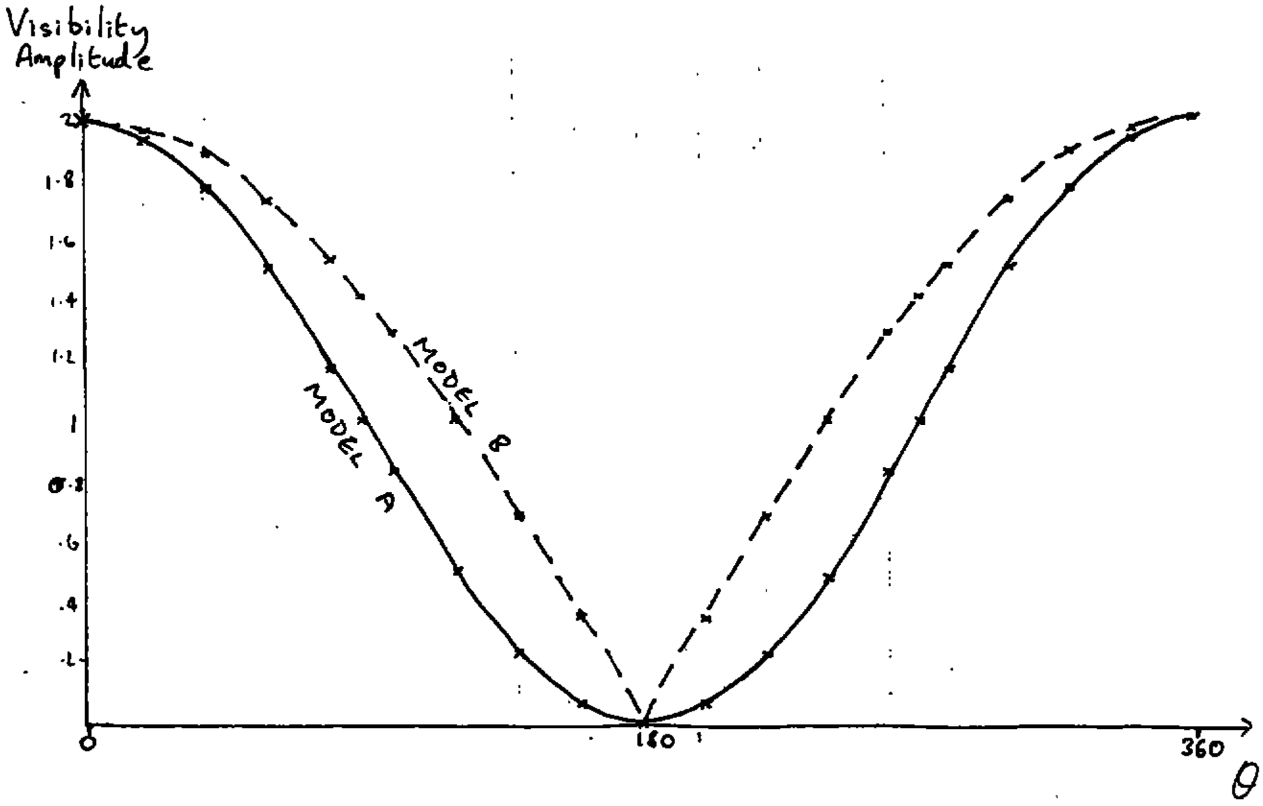

Figure 8 Plots of visibility amplitude which would be observed by an interferometer for the two models shown in Figure 7. The closure phases would be identical for the two models.

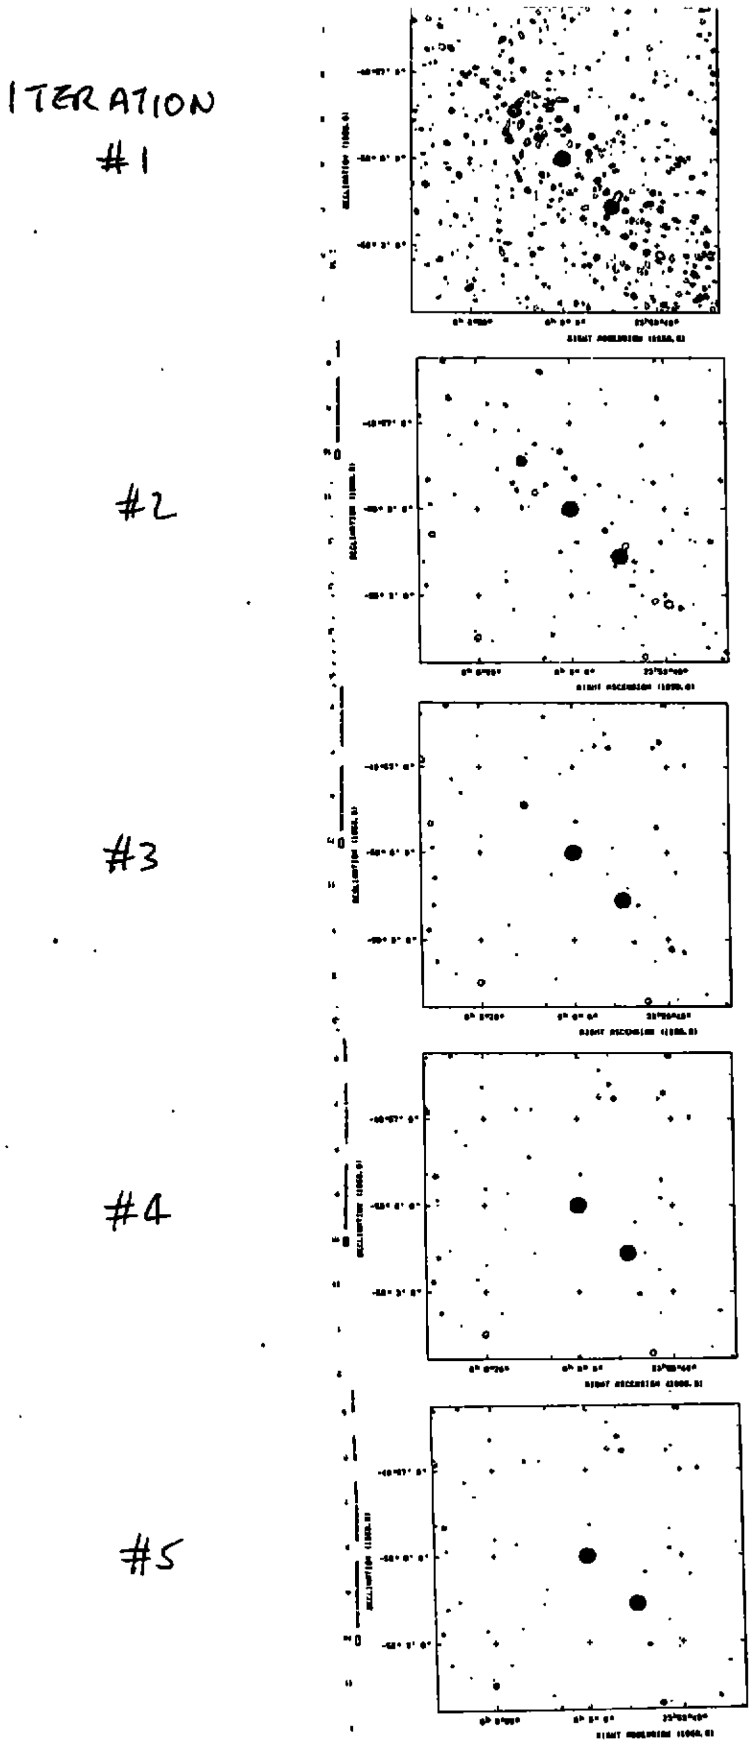

Figure 9 Result of using ASCAL on simulated observations of the double with good signal-to-noise. The initial symmetry has finally disappeared, and the map has converged on the correct solution (model a).

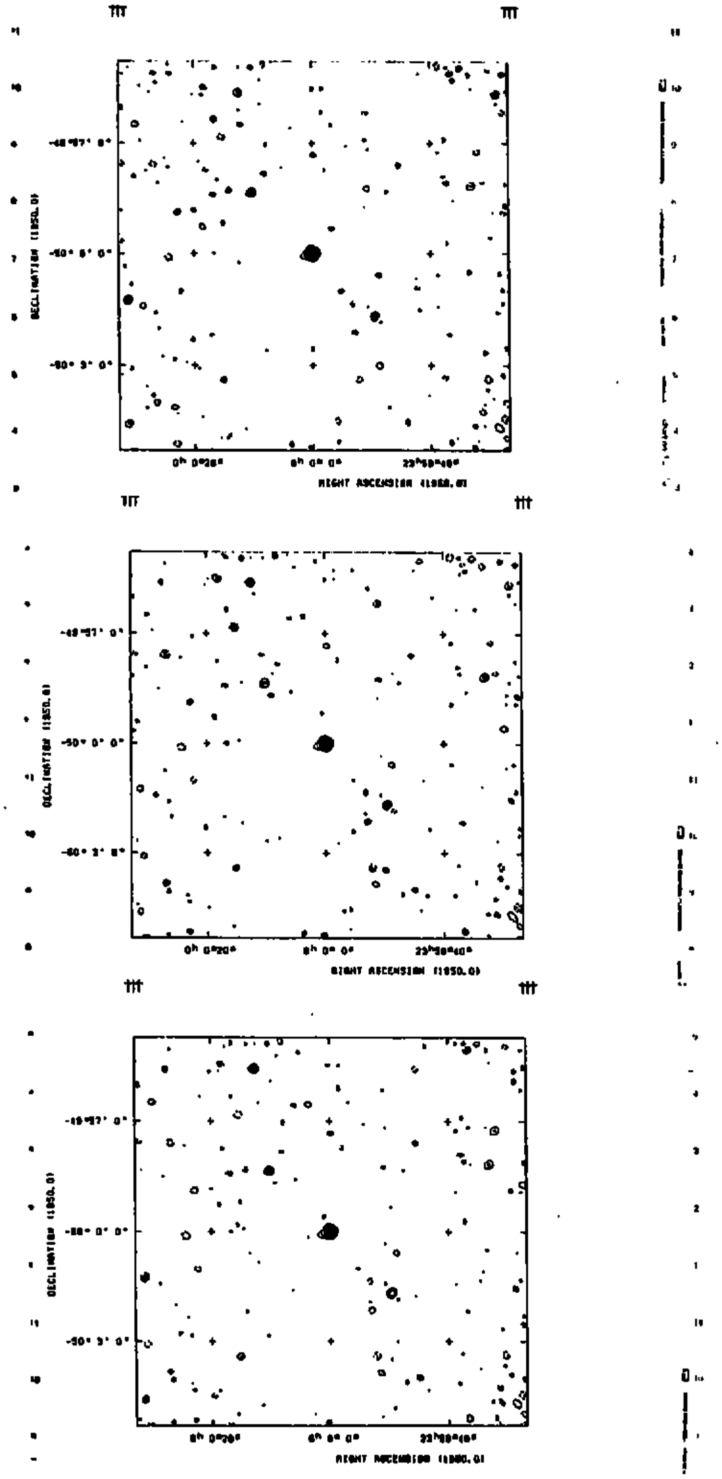

Figure 10 Result of using ASCAL on simulated observations of the double with poor signal-to-noise. The map is dominated by a single, central component with weak symmetrical sources either side (something like model b), although these still barely exceed the noise. Contour levels are at -10, -5, 5, 10, 20, 30, 40, 50, 60, 70, 80, 90 % of the peak. In subsequent iterations, the southwest component slowly started to increase in intensity, indicating a slow migration to solution (a).

FILE COPY
28 NOV 1984
AT

AT/25.1.1/011

SELFCAL and the AT

Ray Norris and Mike Kesteven

25 November 1984

1.0 SUMMARY

Recent tests of the SELFCAL error correction algorithm on simulated AT data revealed an unexpected problem: in some circumstances the maps were symmetrised, so that some components of the source model would appear reflected into the other half of the map. Although we have not yet reached a full understanding of this problem, its causes, and cures, we have achieved an empirical understanding, and suggest possible techniques which may solve the problem. This document is a progress report in which we take stock of what we now know, and invite comments which may guide our subsequent exploration.

1.1 Introduction

When a radio interferometer observes a source, the data are subject to errors which are predominantly telescope-based (i.e. depend only on the telescope, and not on the baseline). Such errors may typically be caused by atmosphere, ionosphere, or receivers, and may appear as errors in either phase or amplitude. Generally speaking, phase errors have a greater effect than amplitude errors on the quality of the final map, and it is these phase errors with which this report is primarily concerned. Telescope-based phase errors may be corrected by:

1. Phase referencing to a nearby unresolved source
2. Use of redundancy techniques.
3. Various algorithms (closure phase, hybrid mapping, etc.) which are known collectively as SELFCAL.

SELFCAL techniques have been widely used, with much success, on many arrays (e.g. VLA, MERLIN, VLBI) for some time, and are generally regarded as reliable and robust. The implementation used for the tests described here is the AIPS program ASCAL (or its derivative BSCAL).

ASCAL is used as follows. It is first given a model of the source. This model may be as simple as a point source, or as complex as a CLEANed map of the source. Then, for each integration period, ASCAL tries to minimise the difference bewteen the actual visiblity data and the predicted data from the model, by introducing complex telescope gain errors. These gain errors are then used to correct the data, which are then transformed and CLEANed. Any subsequent iteration will generally use this CLEAN map (or a portion of it) as a model for ASCAL.

A common strategy is to use a point source for the initial model. ASCAL will then try to incorporate any phase variations in the observed data into telescope errors, to leave 'corrected' data with zero phase. This will tend to symmetrise the map, but on the next iteration, when the CLEAN map (or, usually, only the CLEAN components down to the first negative component) is used as a model, the symmetry will be reduced, and will disappear in subsequent iterations.

2.0 THE INITIAL PROBLEM

We first became aware of a problem when simulated multi-day AT observations of the test source SPIRAL (see Figure 1) produced a stable symmetrical solution after a point source had been used for the initial model. Subsequent iterations failed to remove this symmetry. To investigate this effect, a series of tests were run using 1-day data in order to reduce CPU time. Figure 2 shows the effect of CLEANing the raw data, without applying SELFCAL. Figure 3 shows the effect of several iterations of SELFCAL to the same data, to which no telescope errors have been added. The SELFCALled solution is clearly symmetrised, and it retains this symmetry through subsequent iterations.

Further tests appeared to show that, if a model consisting of the components of the test source down to a flux limit F was used, then all flux below F was symmetrised (see Figure 4). However, we are less confident of this result because of the low dynamic range of the maps. Furthermore, this effect seemed sensitive to the complexity of the model and the quality of the uv data.

Subsequent tests were run on the simpler test source WONKY (Figures 5 and 6). In this case, the map would always converge to the correct solution, but it would often do so very slowly after having first gone to the symmetrical solution. However, the tests showed that this problem seemed to be exacerbated by the linear AT array: the problem was reduced in its severity if natural (rather than uniform) weighting was used, if the integration interval for ASCAL was increased, if the number of antennas was increased, or if the antennas were redistributed in a two-dimensional array.

A hand-waving explanation of this phenomenon might run along the following lines. For a one-dimensional array, each gain solution (corresponding to one integration period) would sample only one direction in the uv plane. The difference between two possible solutions (e.g. the model described in the next section) might appear only at one position angle, and thus affect only a small proportion of the independant gain solutions.

3.0 A SIMPLER EXAMPLE

To illustrate the problem, a very simple example is chosen. Consider a source consiting of a 1Jy point source at the phase centre, plus a 1Jy point source offset a distance Δ (Figure 7a). This model will be called solution a. The observed visiblity on a baseline $\underline{s}$ will have a phase θ, and an amplitude of $(1+\cos\theta)$, where $\theta=2\pi\Delta.\underline{s}$. If a point source model is given to ASCAL, ASCAL will find an exact solution (solution b) in which the observed phase is identically zero, and the gain of telescope at position $\underline{s}_i$ is $2\pi\Delta.\underline{s}_i$. Note that this is equivalent to saying that the closure phase of either solution is identically zero. When mapped, this will give the image shown in Figure 7b, corresponding to a visibility amplitude of $\sin(\theta)/\sin(\theta/2)$. The amplitudes for the two models are plotted in Figure 8. Note that, since the closure phase is zero in either case, the models can be distinguished only by their amplitudes, which differ significantly only in the region of $\theta=\pi$. In subsequent iterations, ASCAL will push towards solution b, and CLEAN will push towards solution a. The winner depends on the signal to noise ratio, uv coverage, etc. Clearly, there will be circumstances in which solution b will be a stable solution.

Examples of two trials using different signal-to-noise ratios are given in Figures 9 and 10.

4.0 POSSIBLE CURES

The example above can easily be made to converge on the correct solution by limiting the telescope errors in any one iteration of ASCAL. An elegant way of implementing this would be to do a constrained minimisation when solving for gains. An inelegant way, and the way tried in the tests, is simply to replace any phase errors $>10^{\circ}$ by a 10° error. However, it is not clear whether this will work for more complex source distributions. It may also fail in the presence of real large telescope errors (such as may be encountered at Q band, etc.)

Alternatively, preliminary tests indicate that use of an uncalibrated CLEAN map, down to the first negative component, is unlikely to converge on the symmetrical solution however bad that initial map is. Clearly, however, this method will also fail if the data are so bad as to produce an uncalibrated map with only one component above the negative component level.

5.0 CONCLUSIONS

Our tentative conclusions are:

1. In most circumstances, SELFCAL will work on the AT, but
2. Care will be required in the choice of the initial model
3. This suggests that fairly frequent observations of calibration sources will be needed.
4. The computing requirements are likely to be larger than originally anticipated, since several iterations are generally needed.
5. On the wider astronomical front, there is some cause for worry as to whether maps produced on sparsely filled arrays such as MERLIN or VLBI, for which an initial point source model has been used for SELFCAL, may have been symmetrised.

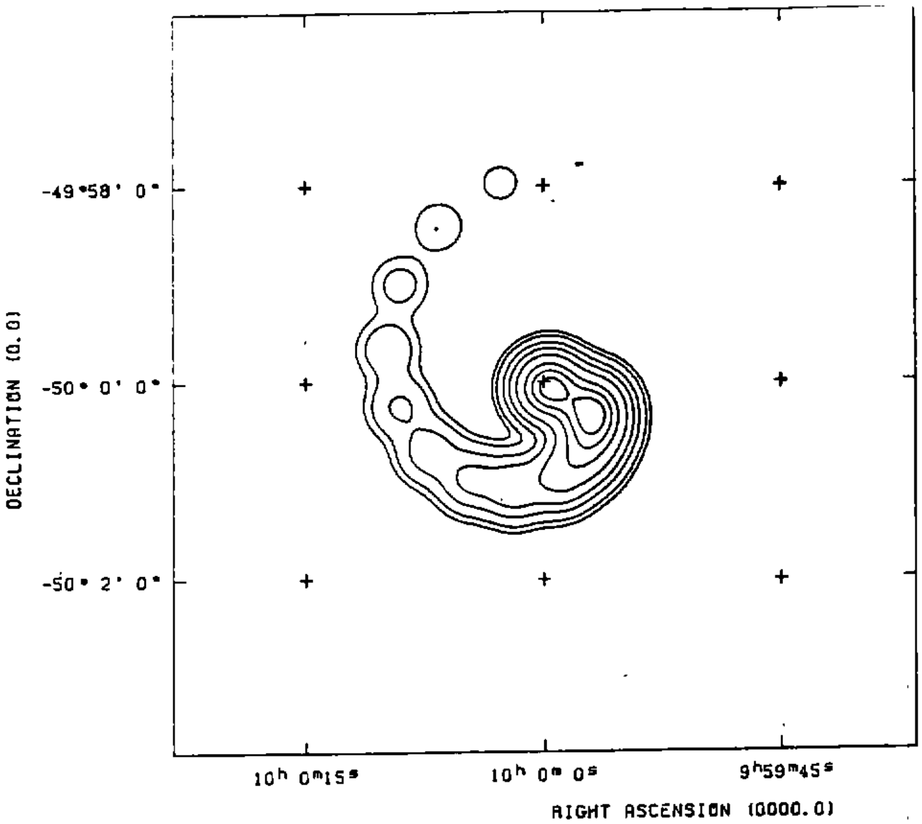

Figure 1 Model distribution of the test source SPIRAL. In this and all other maps (except where stated otherwise), contours are plotted at intervals of -2, -1, 1, 2, 5, 10, 20, 40, 60, 80% of the peak flux.

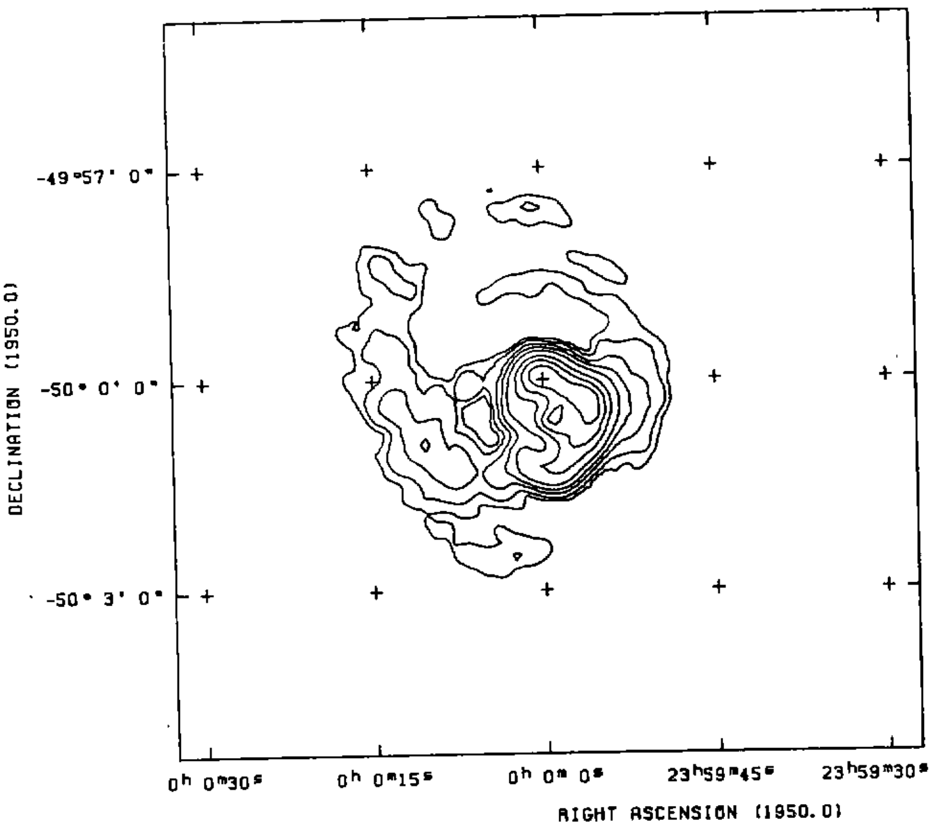

Figure 2 The CLEANed raw data from a simulated one-day observation of SPIRAL.

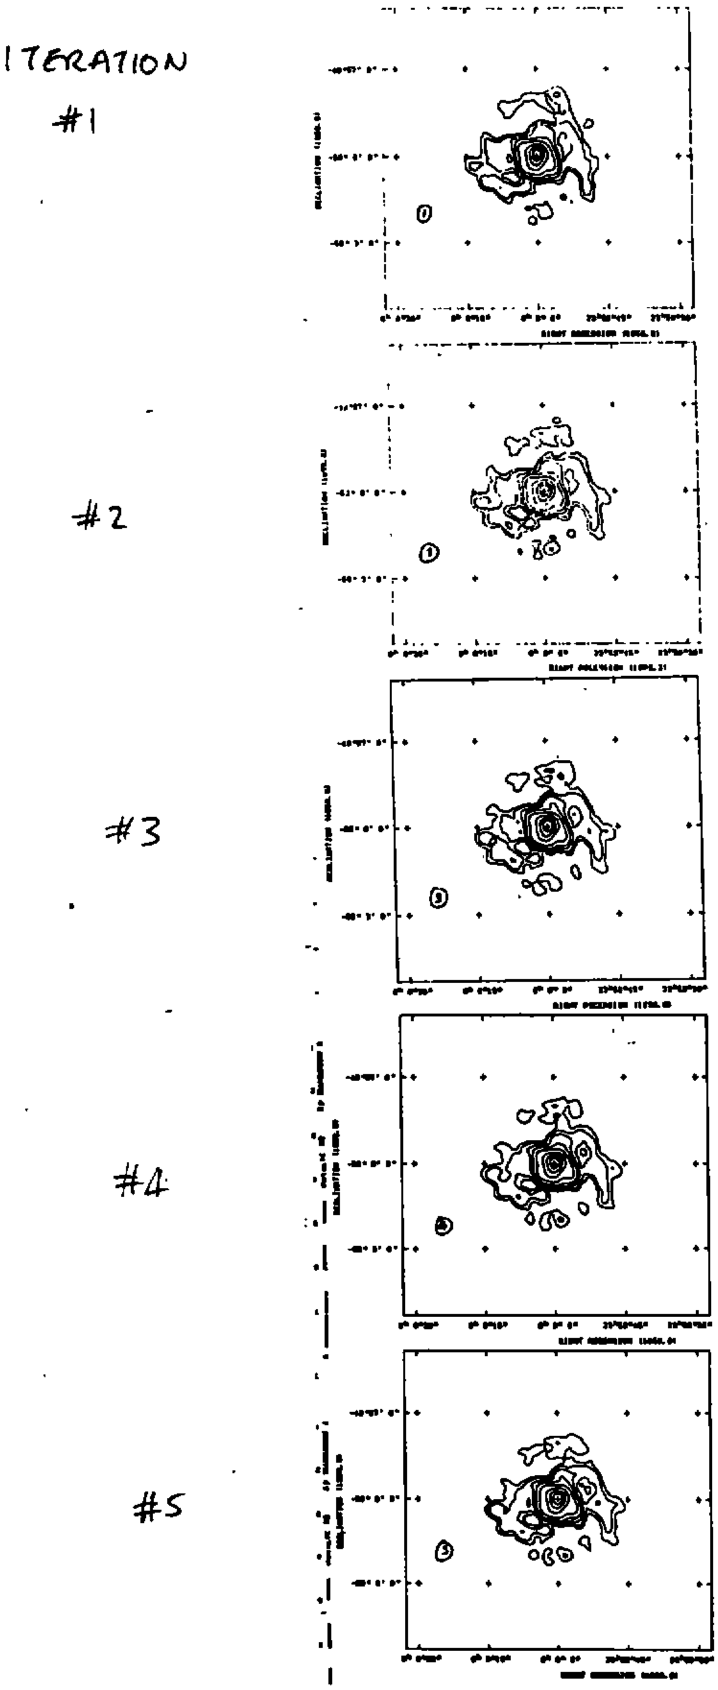

Figure 3 The ASCALled maps from the same data as for Figure 2. In this and all other maps, the array DR3D4 has been used with system noise but no telescope errors added, and ASCAL has been used on phase only. The first iteration of ASCAL (top map) used a point source model, and subsequent iterations each used the previous CLEANed map down to the first negative contour. Note how the maps are all roughly symmetrical.

ITERATION
#1

#2

#3

#4

#5

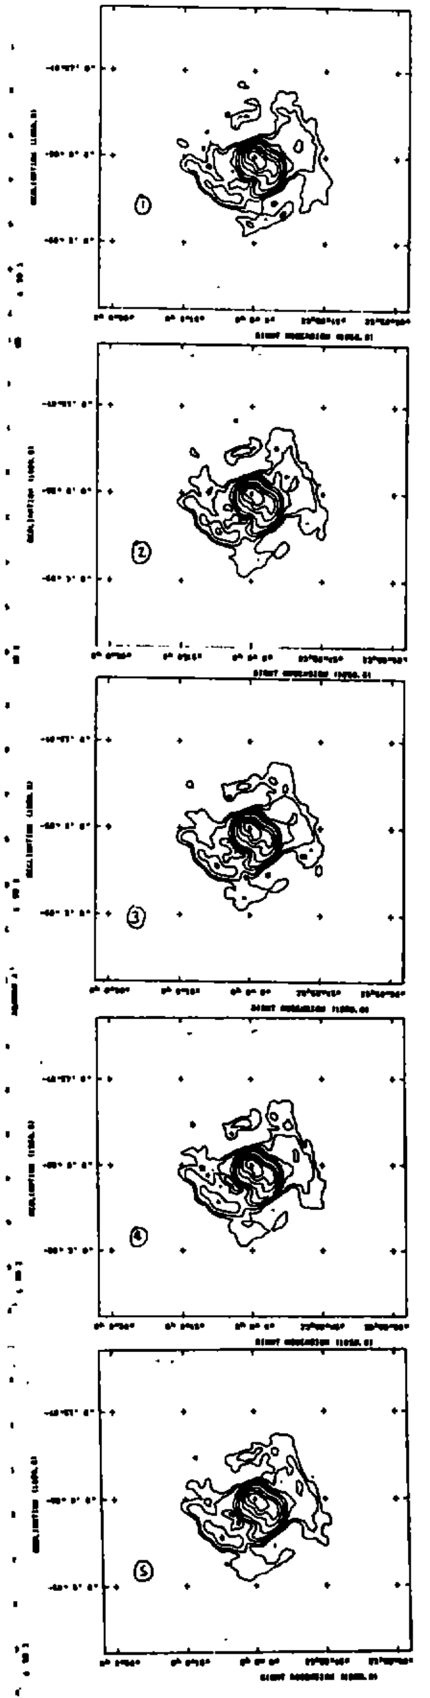

Figure 4 The ASCALled maps from the same data as for Figures 2 and 3. In this case, however, the model used for the initial iteration consisted of the first 50 components of Figure 2. Note that although the symmetrical components near the centre have been eradicated, much of the weaker symmetrical structure appears to remain.

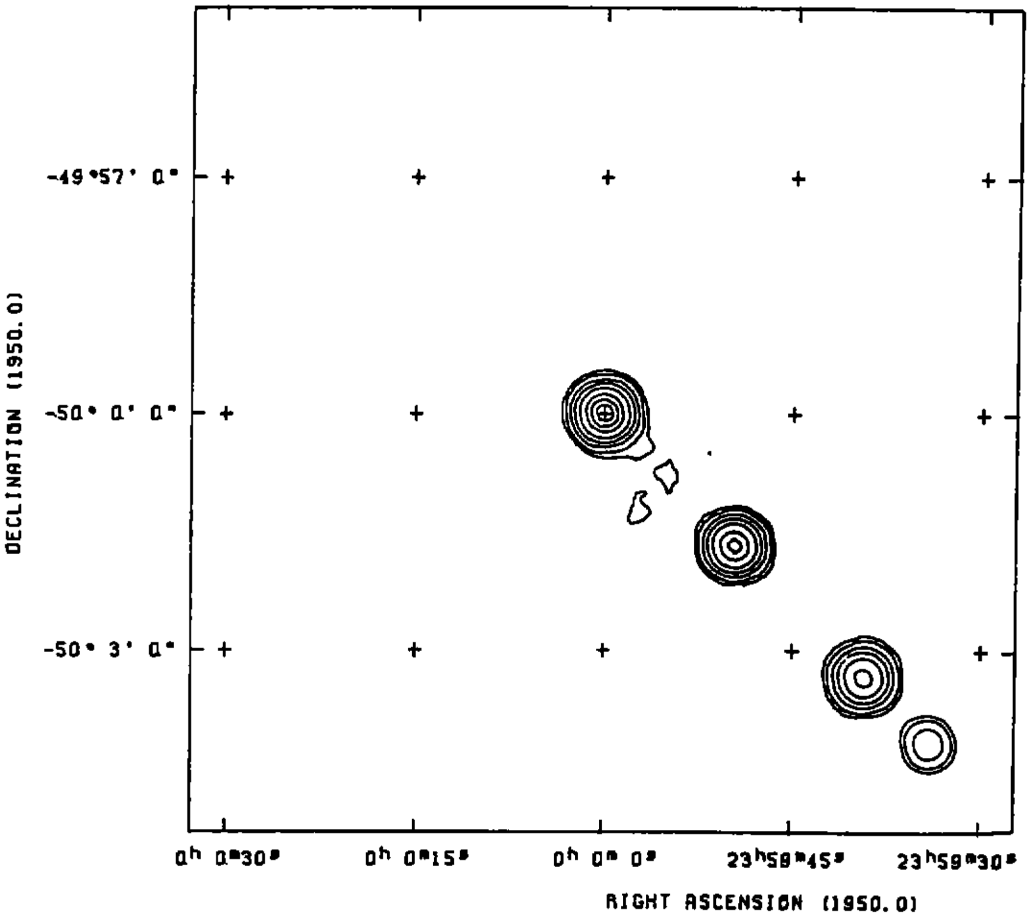

Figure 5 The CLEANed raw data from a simulated one-day observation of the source WONKY.

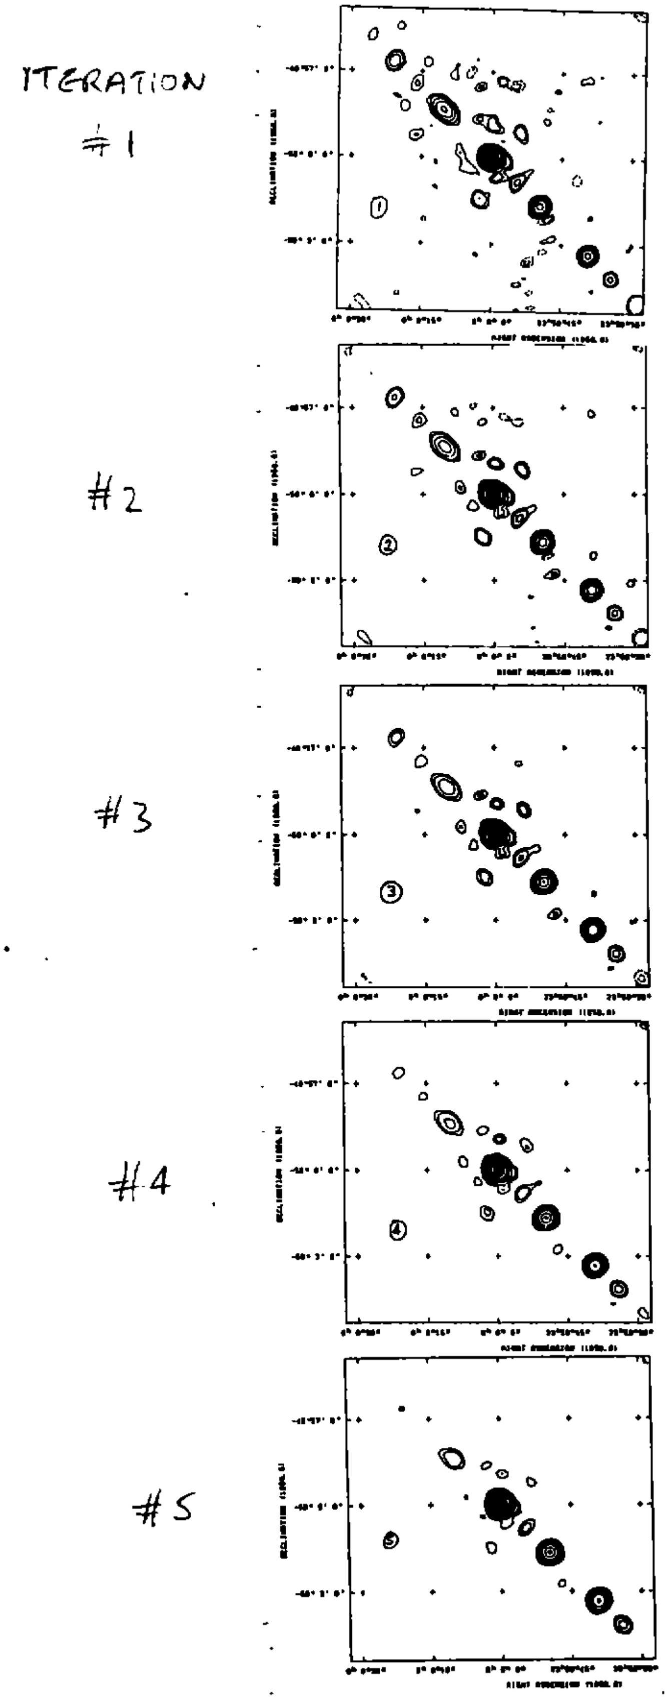

Figure 6 The ASCALled maps from the same data as for Figure 5. The first iteration of ASCAL (top map) used a point source model, and subsequent iterations each used the previous CLEANed map down to the first negative contour. In the case of this simpler model, the process does seem to be converging (although slowly) on the correct solution.

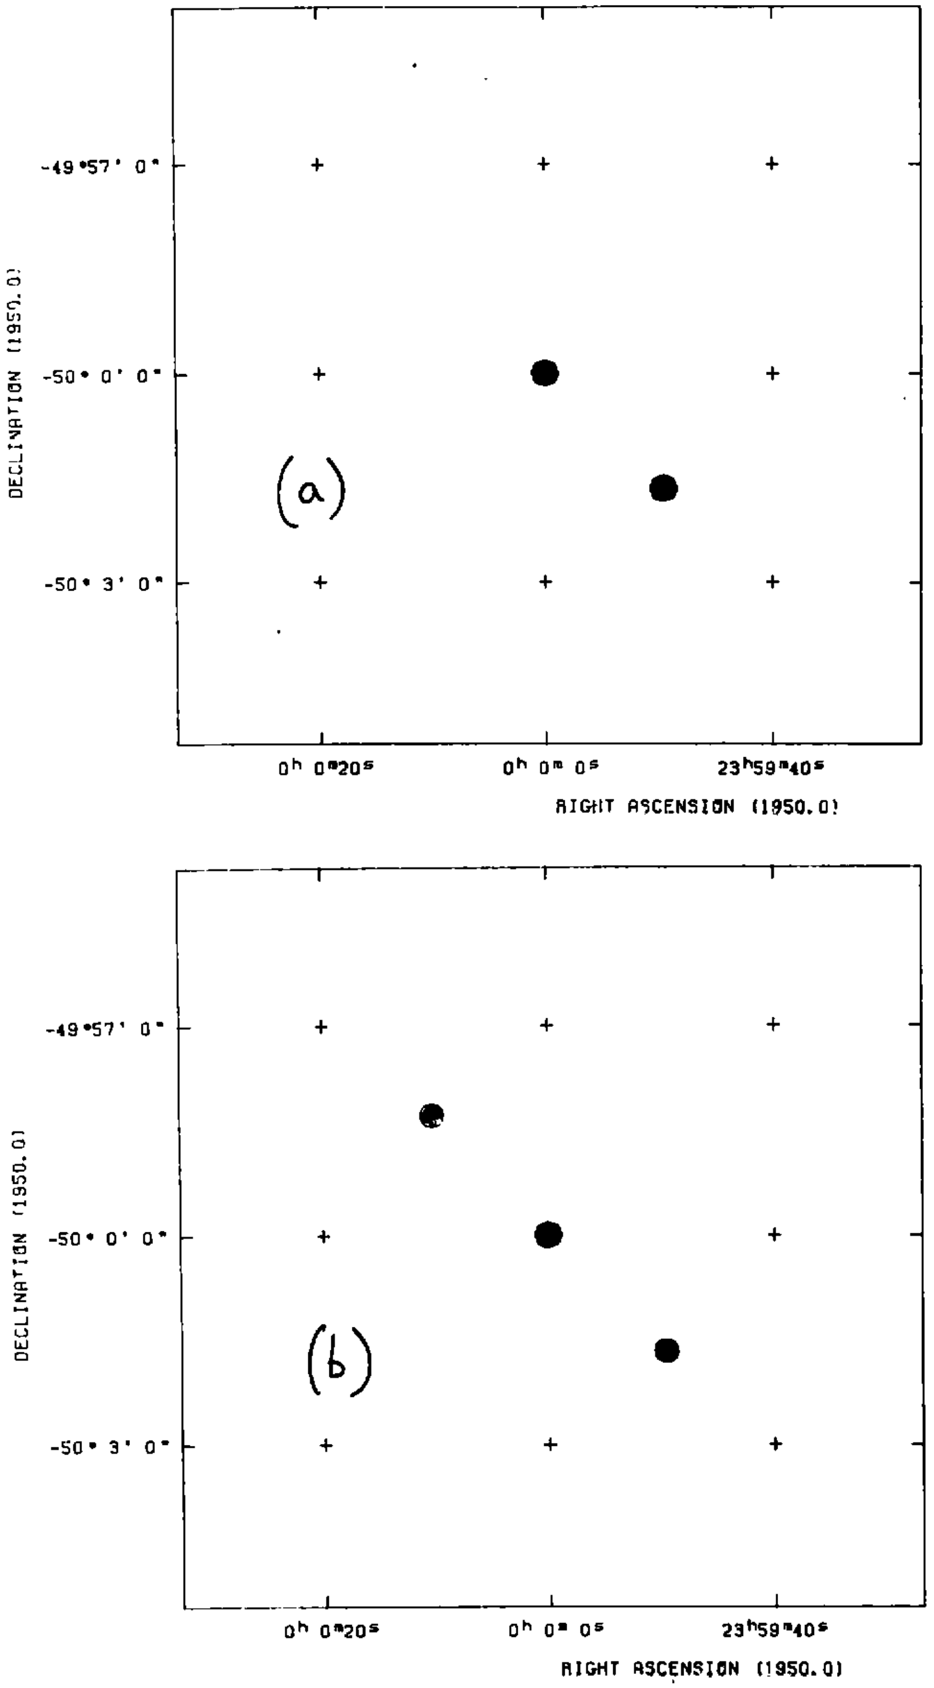

<u>Figure 7</u> (a) Model sky for the double source discussed in the text (Model a). (b) Model sky for the symmetrical source discused in the text (Model b).

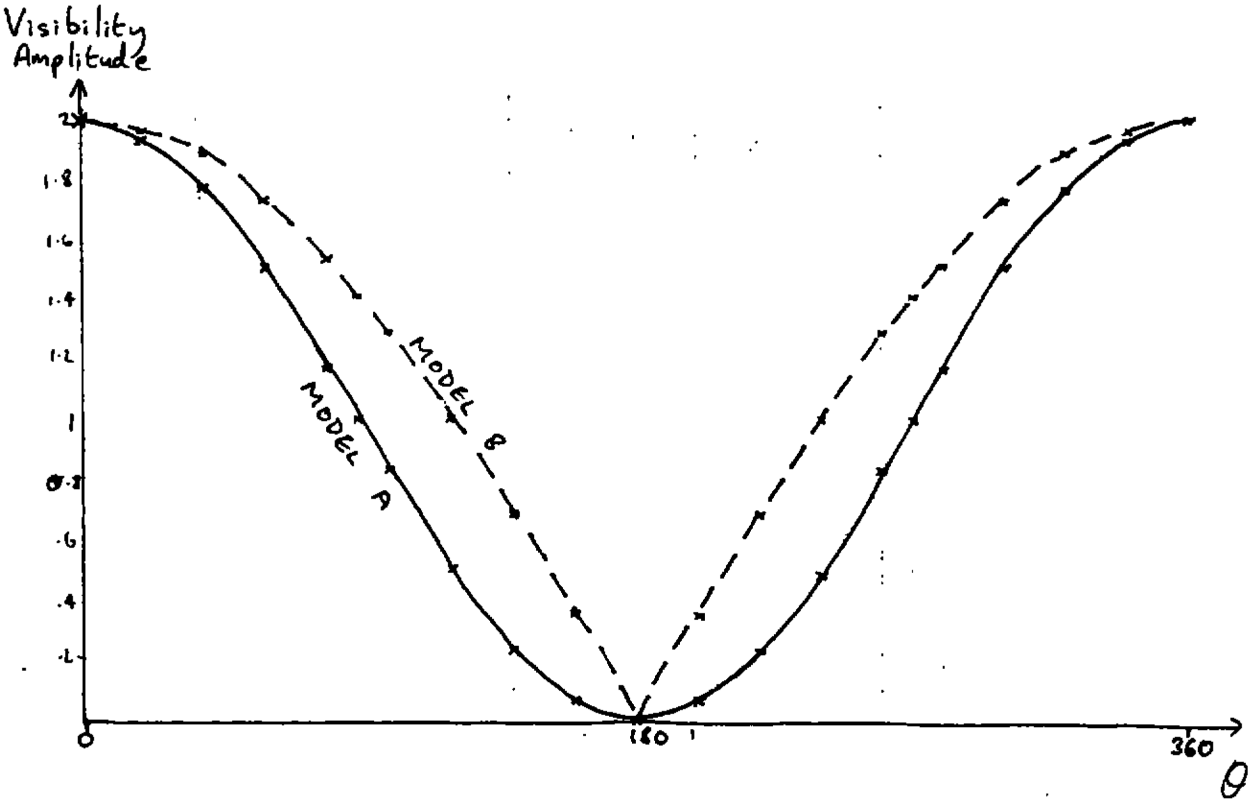

<u>Figure</u> <u>8</u> Plots of visibility amplitude which would be observed by an interferometer for the two models shown in Figure 7. The closure phases would be identical for the two models.

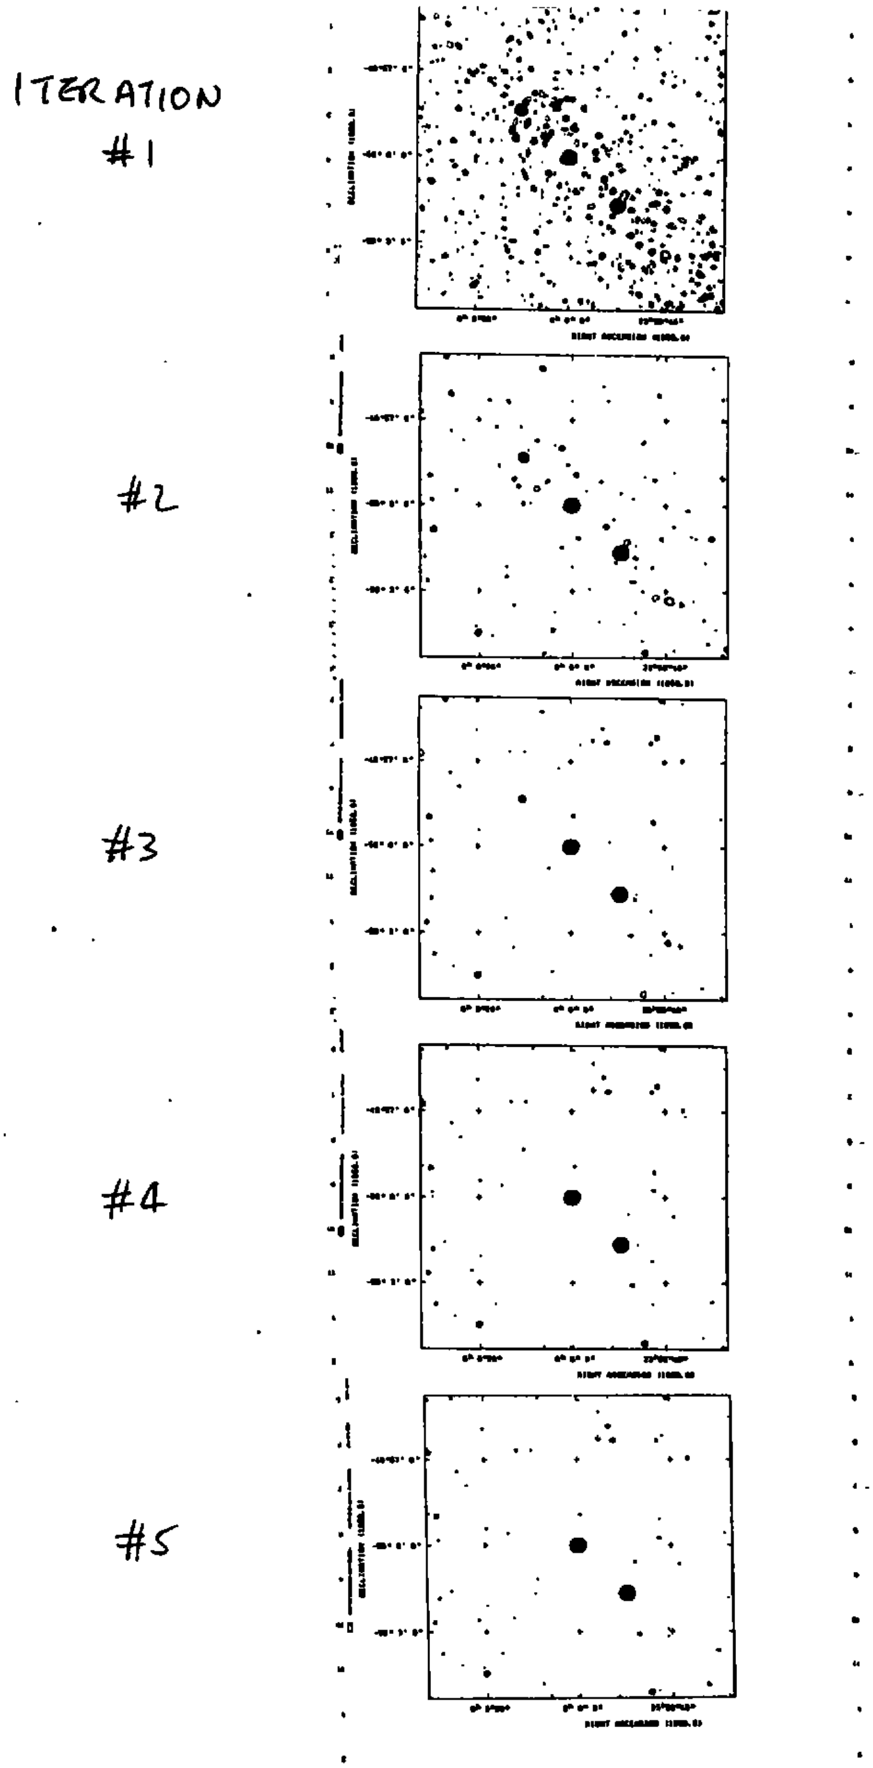

Figure 9 Result of using ASCAL on simulated observations of the double with good signal-to-noise. The initial symmetry has finally disappeared, and the map has converged on the correct solution (model a).

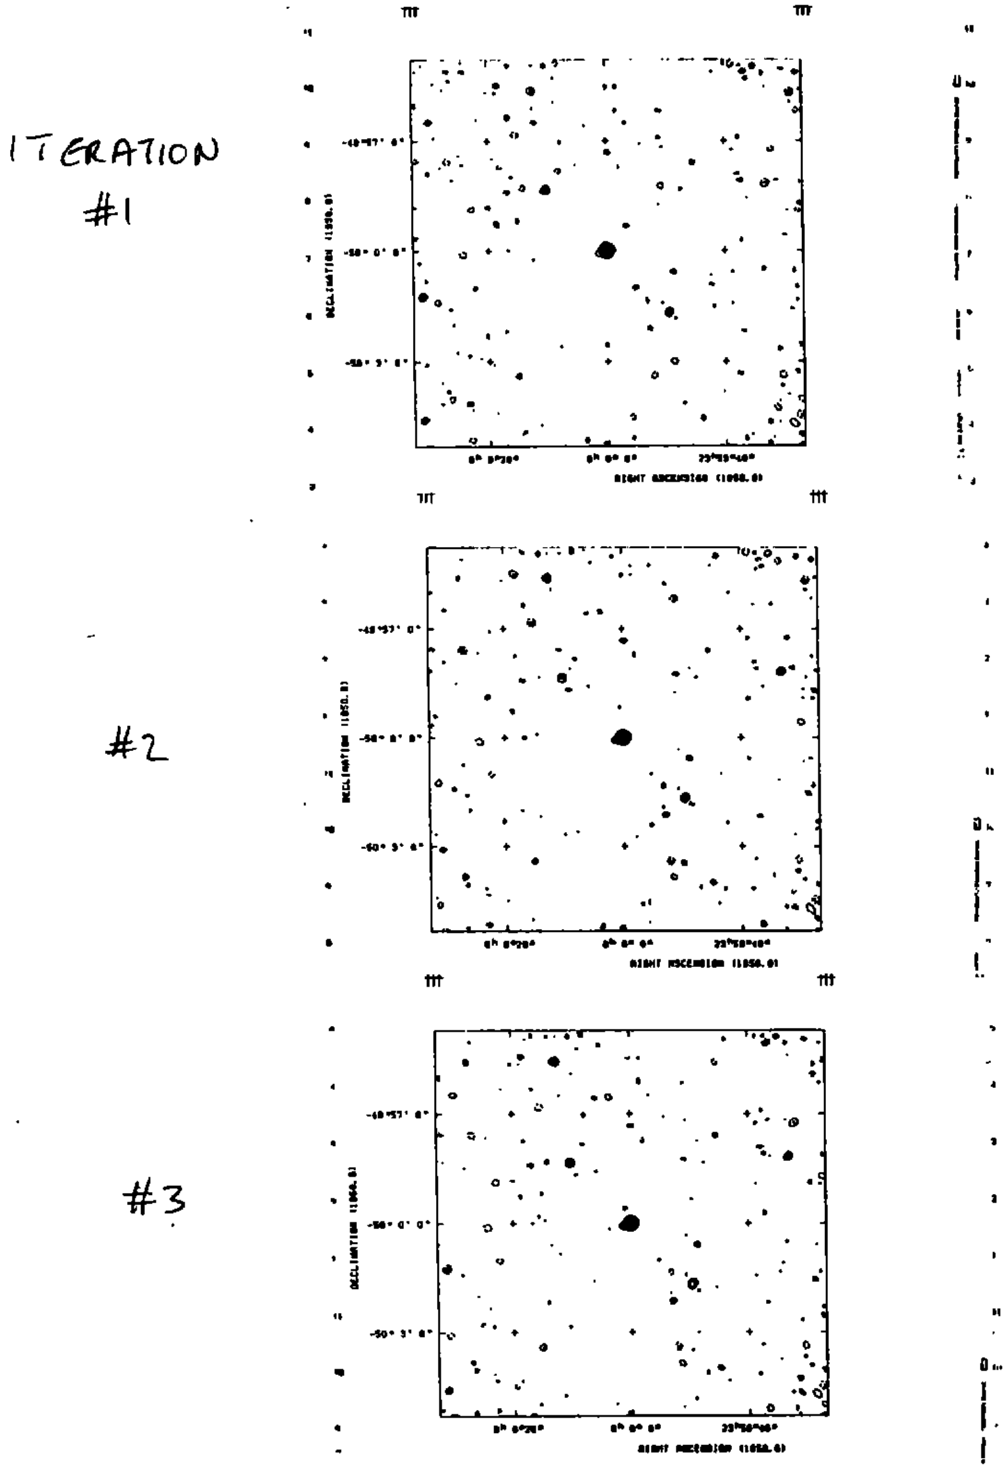

Figure 10 Result of using ASCAL on simulated observations of the double with poor signal-to-noise. The map is dominated by a single, central component with weak symmetrical sources either side (something like model b), although these still barely exceed the noise. Contour levels are at -10, -5, 5, 10, 20, 30, 40, 50, 60, 70, 80, 90 % of the peak. In subsequent iterations, the southwest component slowly started to increase in intensity, indicating a slow migration to solution (a).