

Some Synthesis Telescope imaging algorithms to remove nonisoplanatic and other nasty artifacts

A.G. Willis

National Research Council of Canada, Herzberg Institute of Astrophysics, Dominion Radio Astrophysical Observatory, P.O. Box 248, Penticton, BC, Canada V2A 6K3
e-mail: Tony.Willis@hia.nrc.ca

Received August 11, 1998; accepted February 18, 1999

Abstract. The small diameter antennas of the DRAO Synthesis Telescope give this instrument a very wide field of view (107 arcmin FWHM at 21 cm wavelength and 330 arcmin at 74 cm wavelength). Consequently the telescope has recently begun a systematic survey of a large region of the galactic plane. However, due to the wide field of view, the varying quality of the components used to build the telescope, and the long time required to synthesize a given field, artifacts may appear in ST images which are not easily removable by standard self-calibration techniques. This paper describes the specialized algorithms which have been developed at DRAO to remove such image defects.

Key words: instrumentation: interferometers — methods: miscellaneous — techniques: image processing — techniques: interferometric — telescopes

1. Introduction

The DRAO Synthesis Telescope (ST) consists of seven 9 m antennas having equatorial mounts that are situated along a 620 m east-west baseline. It operates at radio wavelengths of 21 and 74 cm. A total of 12 half-days of observing is required to completely sample the UV plane for a particular piece of sky. The small size of the antennas means that the array has a very wide field of view; the primary beam HPBW is 107 and 330 arcmin at 21 and 74 cm respectively. This wide field of view, coupled with a shortest spacing of 13 metres, makes the telescope an exceptionally useful instrument for studies of our own galaxy; the telescope is currently dedicated to a systematic survey of that part of the galactic plane above declination 30 degrees that is visible from the DRAO site. This project is called the Canadian Galactic Plane Survey (CGPS).

Send offprint requests to: A.G. Willis

The telescope has been gradually built over a 25-year time period using components acquired from an eclectic variety of sources; some antennas were even obtained from a San Francisco junk yard (Galt, private communication)! In contrast, most well-known aperture synthesis radio telescopes, such as the Very Large Array (VLA) or Westerbork Synthesis Radio Telescope (WSRT), were designed and manufactured to some pre-defined performance criteria – e.g. all dishes should have the same primary beam shape, all dishes should have a well-specified RMS surface accuracy allowing observations to a given wavelength, etc.

This combination of eclectic mechanical and electronic components is coupled with a very large primary beam. Consequently the images produced by the DRAO ST often contain artifacts visible above the noise that cannot be easily removed by standard self-calibration algorithms developed for telescopes such as the VLA or WSRT.

In this Paper I describe the special image processing techniques that have been developed to remove these instrumental effects from DRAO images. The end result is a set of high quality images for the Canadian Galactic Plane Survey (CGPS). These techniques may have applications to other synthesis arrays, especially when these arrays are pushed to the limits of performance at, for example, very high frequencies.

2. Problems

The primary method of obtaining high quality images from synthesis arrays is the technique of self-calibration (see Thompson et al. 1986, Chapter 11, Sect. 2, and references therein). In standard self calibration one takes observed visibility data D and model visibility data M for the baselines in the array and, over some time period where one assumes the antenna gains to be uniform over the field of view and constant, one does a least-squares fit to the visibility data to find a complex gain G_{ant} for each

antenna in the array. For each baseline in the array a gain relationship G between data and model is then given by

$$D = GM \quad (1)$$

where,

$$G = G_i G_j^* \quad (2)$$

and G_i and G_j are the individual antenna gains of the two antennas i and j whose correlated signal gives the visibility on this baseline. By then performing the operation of dividing the data by the gain (D/G) one effectively “adjusts” the data toward the model and hopefully improves the data and the corresponding image.

However there are situations where the assumption of a uniform complex gain is incorrect. At low frequencies, the observed data can be affected by ionospheric irregularities that introduce phase errors which are dependent on position in the field of view (Cotton 1989). The synthesis array itself can also introduce position dependent distortions and therefore position dependent gain errors. Cotton (1989) describes some of the common distortions that affect synthesis data. For example non-coplanar terms must be correctly accounted for in order to produce undistorted images from a non-coplanar array (see also Perley 1989).

The DRAO ST has position dependent distortions because the seven antennas in the array have two different sizes and thus different beam shapes. The two antennas at each end of the array, and thus the pair forming the longest baseline are each 9.15 metres in diameter with a HPBW of 101.8 arcmin at 21 cm wavelength; the remaining five antennas are all 8.54 metres in diameter, with a beam size of 108.8 arcmin. Since each antenna can be correlated with every other antenna, we obtain a total of 21 baselines divided up as follows: 1 baseline of $B \times B$, 10 $L \times L$, and 10 $B \times L$ ($B = \text{big}$, $L = \text{little antenna}$). An analysis of overlapping fields from the Canadian Galactic Plane Survey (CGPS) shows that we end up with a weighted mean beam having a HPBW of 107.2 arcmin (Gibson, private communication). Corresponding differences exist at 74 cm wavelength.

The differing sizes of the antennas means that the ST has a point spread function (PSF) that varies radially from the field centre, in contrast to the uniform PSF based on UV coverage that is normally assumed in synthesis imaging techniques. In Sect. 3.1 I shall present an algorithm for adjusting the data associated with that part of an image which lies well inside the main beams of the antennas where the different primary beam responses are well understood. It enables the observer to avoid computing a separate PSF for every direction in the field, and standard self-calibration solutions can then be applied to improve the image.

Another error occasionally affects DRAO ST images if a time variable source is situated in the field of view. If the incoming radio signal is constant one can produce a correct image of the sky from observations made over a variety of time periods. Unfortunately point sources whose

signal varies significantly during an observation produce an observed response that does not agree with the PSF calculated under the assumption of an statistically invariant sky. The end result is that one cannot properly clean the effects of such sources from a field with the standard clean algorithm. A number of sources with short period variability have been detected in the CGPS survey. In Sect. 3.2 I describe how the effects due to time variability may be removed by an inversion of the self-calibration procedure.

However, not all instrumental distortions can be removed by an adaption of the self-calibration algorithm. The wide field of view of the ST is both its strength and its weakness. The telescope can easily make wide field images of large scale emission complexes in our galaxy, but at the same time interfering sources such as Cas A, Cyg A or the Sun are likely to be detected in either the main primary beam or in a sidelobe. Grating rings from these objects often run through the area of interest in the field and it is critical to remove the effects of these sources.

If a source is located at a large distance from the field centre its computed model visibility may differ significantly from that observed. This may happen, for example, because the source is located in an antenna sidelobe where the instrumental response has not been properly measured and deviates from that computed on the basis of purely geometric considerations. (The appendix of this paper gives a detailed description of the assumptions and procedures that go into computing model visibilities for the DRAO ST.) It is also possible for instrumental effects to have a baseline dependent origin. If baseline dependent faults, such as bandpass mismatches, are significant, the fundamental assumption of self-calibration, that gain errors can be described by purely antenna based errors, breaks down. Cornwell & Fomalont (1989) list additional causes of of baseline-based errors.

In Sect. 4.1 through Sect. 4.5 I describe an algorithm which has been quite successful at removing sources affected by distortions due to either unknown or baseline-based errors.

3. Modifications to the self-calibration algorithm

3.1. Different size antennas

I now turn to the problem of processing images generated from an array having antennas with different HPBWs (101.8 and 108.8 arcmin at 21 cm in the case of the DRAO ST). Measurements of the separate antenna patterns made by Dougherty (unpublished DRAO technical memos) show that at 21 cm the maximum difference in response is about 5% at a distance of 64 arcmin from the field centre (the primary beam response of the 9 m antennas is 32.0% of the peak and that of the 8 m antennas is 37.1% of the peak at this distance; the weighted mean attenuation is 36.0%). If one happens to detect a

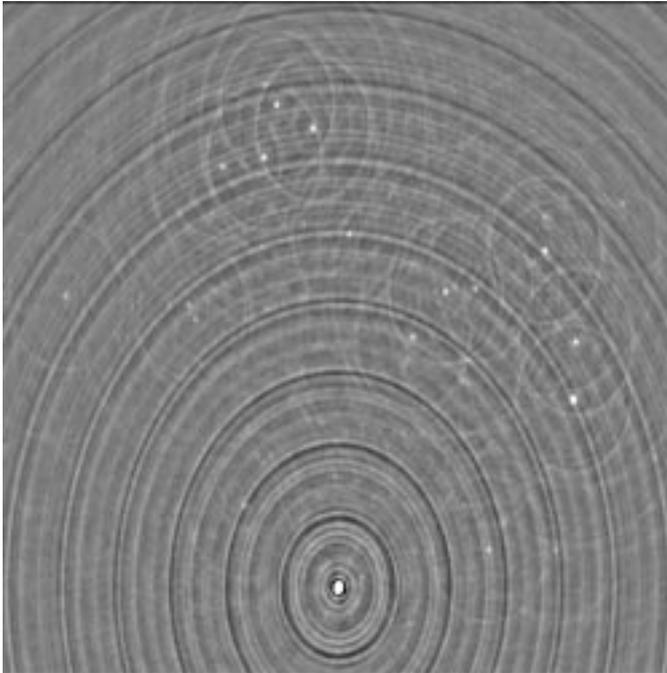


Fig. 1. A DRAO Synthesis Telescope 21 cm (1420 MHz) image of an area containing the strong source 3C 147 at a distance of 64 arcmin to the south of the field centre. The image has a width of 2.8 degrees. It shows the results of doing a standard phase-only self-calibration and clean of the field. Residual rings having amplitudes in the range 6 to 30 mJy (0.09 to 0.44% of the peak signal) surround 3C 147, which has an attenuated flux density of 6.7 Jy

strong source at a fairly large distance from the field centre, this discrepancy will limit both the dynamic range attainable through standard self-calibration procedures and how well one can clean the associated image. The synthesized antenna pattern (PSF) used to compute the response to clean components is derived solely from the UV locations sampled by the array as it tracks the field. In reality, however, each source in a DRAO field will have a slightly different PSF that depends on its distance from the field centre because of the different sensitivities of the individual antenna pair combinations that sample the UV plane. Also, if one has a strong source at a large distance from the field centre and weak sources elsewhere, the contribution from the strong source will dominate the gain solutions derived from self-calibration causing the gain solutions to be incorrect for sources elsewhere in the field.

This effect is quite noticeable if one observes a strong source situated at a distance from the field centre where the beams are significantly discrepant. Figure 1 shows the results of a 12 hour test observation at 21 cm wavelength in which the ST field centre was placed at a distance of exactly 64 arcmin from the strong source 3C 147. There should be a maximum discrepancy between the primary beams of the large and small antennas at this distance. The image shows the results of doing standard cleaning

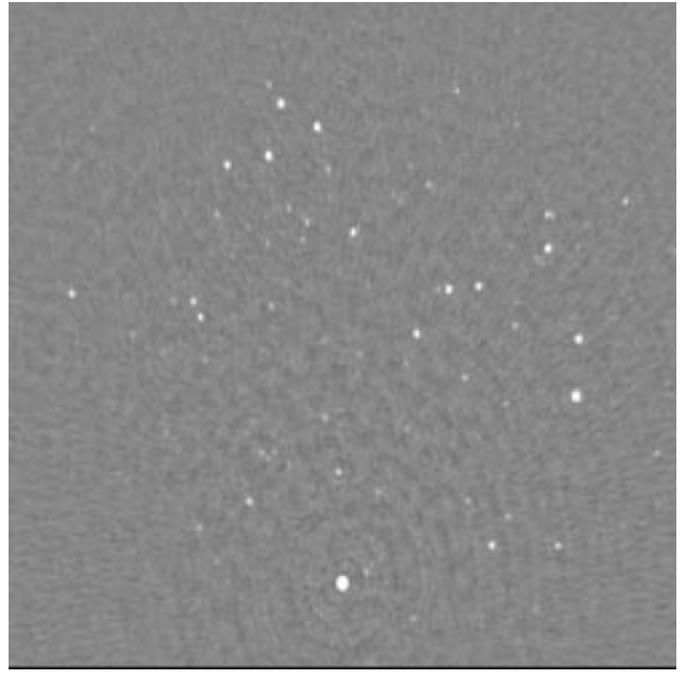


Fig. 2. The 3C 147 field after first subtracting off model visibilities corresponding to clean components in the UV plane and then adding these clean components back into the image when the clean restoration is done. The rms noise in the vicinity of 3C 147 is reduced to 1.3 mJy, or 0.02% of the peak signal. The image half-tone scale is similar to that of Fig. 1

and phase-only self-calibration under the assumption that the antennas all have the same HPBW. Residual grating rings are present because of the discrepancy between the actual and assumed PSF response.

How does one proceed to solve this problem? When doing self-calibration one normally obtains antenna-based amplitude and phase corrections by feeding the gain solver a collection of normalized baseline data, D/M , where D is the observed visibility and M is the model visibility. In the present case one must compute 3 separate sets of model visibility data corresponding to three separate model skys for antenna pair combinations $L \times L$, $L \times B$ and $B \times B$ starting from an image which represents a weighted mean of these combinations. (Here, L = little antenna, B = big antenna).

In order to obtain the model visibilities one does the following steps:

1) Clean the combined image using the Cotton-Schwab algorithm (see Cotton 1989, and Cornwell & Braun 1989). Then for each clean component with derived signal S from the combined image determine its distance from the field centre.

2) Calculate its value on the sky before attenuation by the mean synthesized beam (i.e. correct the clean component for the mean primary beam attenuation A)

$$S' = S/A \quad (3)$$

where A is the mean primary beam attenuation.

3) Then compute the signal S'' which would be seen by a particular antenna pair combination

$$S'' = S' \sqrt{A_1 A_2} \quad (4)$$

where A_1 is the attenuation factor for the first antenna of the baseline antenna pair and A_2 is the attenuation factor for the second member of the pair.

4) Compute the model visibilities M for the baseline associated with the antenna pair by taking the Fourier Transform of the ‘image’ of all the S'' clean components and divide the observed visibilities, D , by M to get the normalized signal D/M for the baseline.

5) Repeat this procedure for each baseline in the array, and feed the normalized visibilities into the self-calibration solution solver, to get antenna-based amplitude and phase corrections that should be independent of the positions of the clean components in the field.

6) Apply the resulting antenna based amplitude and phase corrections to the observed visibilities in the usual way to self-calibrate the data.

Then one would usually make a new image from the corrected visibility data and clean the resulting image. If the previous six steps have gone according to plan the new image should have an improved dynamic range, and one may go further with another round of self-calibration by repeating steps 1) to 6) again. This cycle may take place several times until little or no improvement is seen in the final image.

The resulting restored image should have a higher dynamic range and indeed it does for the 3C 147 field as can be seen in Fig. 2.

The Cotton-Schwab algorithm works because the model visibilities that are computed in steps 3) and 4) above contain adjustments for the separate antenna attenuations and thus agree with the true instrumental response to the incoming signal. Since these visibilities are subtracted off in the UV plane, which is then inverted back to the image plane for the next clean cycle, one avoids large discrepancies between the actual and assumed PSF response.

The current DRAO software package does not have a direct equivalent of the Cotton-Schwab variant of clean, which is implemented in the AIPS task MX. However, a strategy of cleaning an image to a level to somewhat above about 2% of the initial peak signal (and thus avoiding the error levels shown in Fig. 1), subtracting off the corresponding visibilities in the UV plane, generating a new image from the residual UV data, and then cleaning this residual image, etc., allows one to obtain a reasonable approximation to the Cotton-Schwab algorithm. Before generating the final restored image all clean components obtained in the various steps are added together.

The strategy of doing self-cal with big and little antennas together is necessitated by the fact that there is a total of only 21 correlated baselines in the DRAO ST

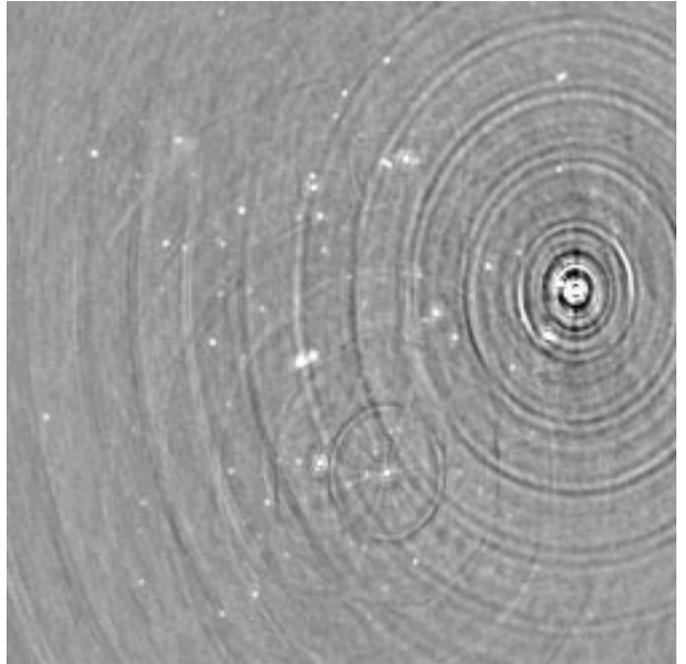


Fig. 3. A DRAO Synthesis Telescope 21 cm (1420 MHz) image of an area containing PSR1930+20. The image, 2.8 degrees across, shows the field as seen before any processing. The pulsar itself has a mean flux of about 189 mJy. The brightest source in the field has a signal of 2 Jy

array. If one had an array with large numbers of both big and little antennas then one could investigate algorithms where data from sub arrays of big antennas was processed separately from data obtained with the small antennas.

3.2. Variable source removal

As I mentioned earlier in Sect. 2, a source whose signal varies significantly over the time period of a synthesis observation violates the assumption that the sky is invariant and therefore the standard clean algorithm will fail to remove the effects of such a source. Figure 3 shows a field surrounding pulsar PSR 1939+20 before any image processing has been done. Figure 4 shows the same area after standard cleaning and self-calibration has been done.

If the variable source is sufficiently strong one can remove its effects by using a variant of the standard self-calibration procedure. In standard self calibration (see the beginning of Sect. 2) by performing the operation of dividing the data by the calculated gain (D/G) one effectively ‘adjusts’ the data toward the model and hopefully improves the data.

However, there is nothing to stop the astronomer from taking an opposite approach with a variable source. Use an imperfect model, the time averaged signal (represented by the clean components found in the 12 day averaged image, or the flux density measured in the averaged image), and move this imperfect model toward the visibility

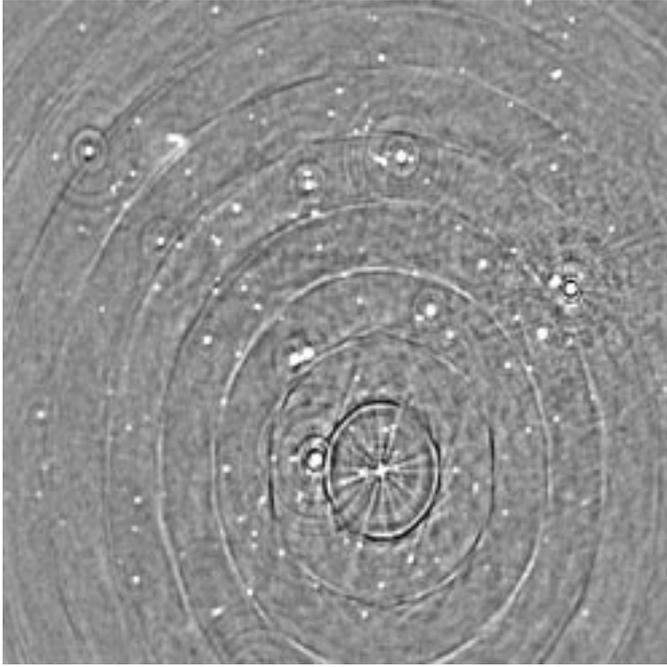


Fig. 4. The pulsar field after standard cleaning and self-calibration has been done. Because of time variations in the pulsar signal, residual grating rings and spokes having amplitudes of up to 3% and 2% respectively of the mean flux surround the pulsar

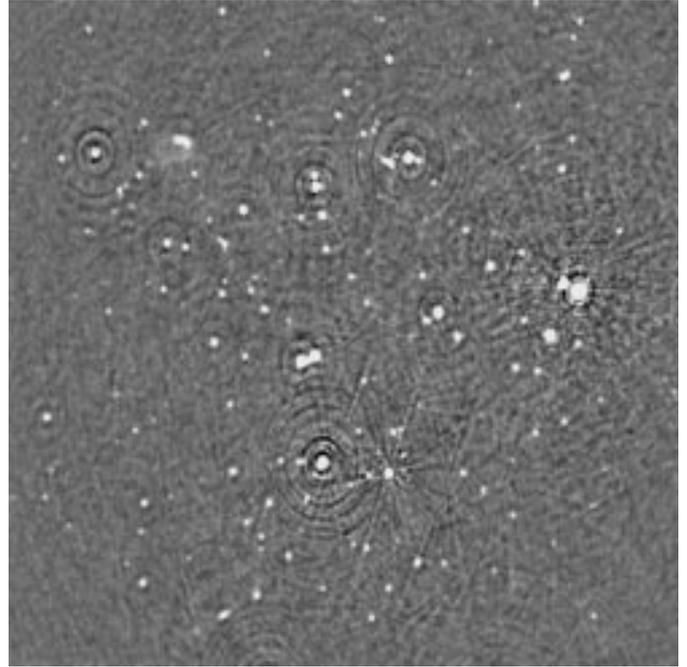


Fig. 5. This image shows the PSR1930+20 field after application of the inverted self-calibration procedure. The remaining spokes due to the pulsar have a level of about 0.5% (1 mJy) of the pulsar's mean signal. Some residual rings with amplitudes 1 to 2% of the source peaks are visible around several sources in the field

data (which one can consider as perfect since the visibility data contain the time varying information). One can then obtain the time varying gain corrections which should be applied to the model to “degrade” the model toward the data.

The algorithm contains the following steps:

1) Obtain a visibility data set D' in which one has deleted the signal E from all other sources in the field as best one can

$$D' = D - E. \quad (5)$$

Since D' is used to compute gain solutions for the variable source, it is critical that the signal E contains a very good estimate of the flux from all other sources in the field.

2) Take visibilities M computed from the imperfect (time averaged) model of the variable source and, using D' and M , calculate the least-squares antenna-based gain corrections as a function of time using a standard self-cal procedure; then obtain the corresponding baseline-based gain corrections G from Eq. (1) and Eq. (2).

3) After obtaining the gain corrections G , instead of dividing the visibility data D' by G as was done in Sect. 2, multiply the model visibilities by the gain corrections and subtract the resulting visibilities from D'

$$D'' = D' - GM. \quad (6)$$

One should end up with a visibility set D'' which is essentially the noise during the observation.

4) Finally, to get a set of visibilities which represent the proper time-averaged field add visibilities for E and M back into the data.

$$D''' = D'' + E + M. \quad (7)$$

Images made from visibilities D''' should give a very good time-averaged image for the field. Figure 5 shows the pulsar field after use of the above algorithm. Most of the effects due to the pulsar have been removed. (The ability of this method to remove very short time fluctuations will obviously depend on the intrinsic signal-to-noise ratio of the source. The weaker the source, the longer the time interval over which we must average the data in order to get acceptable solutions for the gains (Cornwell & Fomalont 1989).)

As an alternative to step 2) above, one can directly start with a baseline-based approach in which one computes an average gain correction G for each baseline

$$G = D'/M \quad (8)$$

over some time interval. Figure 6 shows the results of using this method. The baseline-based approach (Fig. 6) seems to give slightly better results as far as subtracting off a time variable source is concerned. However some distortion of the noise in the area close to the pulsar occurs because the gain correction, G , is derived from simple baseline averaging, rather than from an antenna-based least-squares fit.

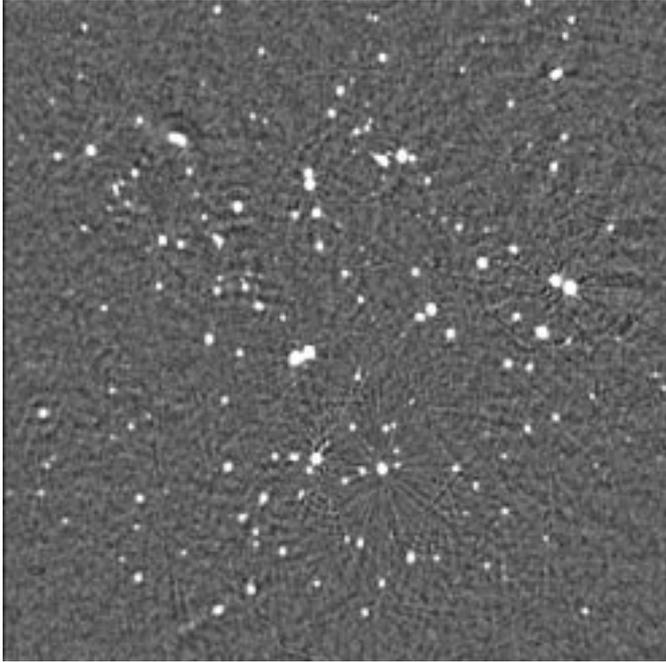


Fig. 6. This image shows the PSR1930+20 field after applying a baseline-based averaging procedure to the individual baseline data. The remaining spokes due to the pulsar have a level of about 0.2% of the pulsar's mean signal. Residual rings around other sources in the field were removed by the MODCAL algorithm described in Sects. 4.1 and 4.5. The RMS noise in this field is about 0.3 mJy

Note that residual grating rings visible around some of the sources in Fig. 5 are no longer visible in Fig. 6. They were removed by the MODCAL algorithm described in Sects. 4.1 and 4.5.

4. The MODCAL procedure

4.1. Theory

Once the primary beam attenuates a signal to about 5% of its intrinsic value, the attenuation factor usually has quite a large error and consequently a calculation of the actual value becomes meaningless. Other instrumental effects such as bandwidth smearing and time smearing also distort the image at large distances. Consequently, for example, CGPS 21 cm continuum mosaics cut off the field of view of an individual field being joined to the mosaic at a radius of 100 arcmin. This distance corresponds to a primary beam attenuation of 6.6% of the intrinsic value of the signal.

I mentioned above that because the DRAO ST has a wide field of view and a narrow bandwidth there is a good chance that any single field is contaminated by responses to one of Cyg A, Cas A, or Tau A. A typical field in which Cyg A is detected is shown in Fig. 7. Other contaminating signals come from man-made interference and solar emission. In all these cases, if one assumes a purely geometric

relationship between the source position and the telescope baseline (see Eqs. (A2) and (A3) of the Appendix), computation of model visibilities from clean components, followed by subtraction of the model visibilities from those observed, usually does not lead to successful removal of the contaminating source. However I shall now describe a method which has been very successful at removing the effects of such sources from DRAO fields, even when the exact instrumental response (such as a sidelobe gain) to the incoming signal is unknown.

Usually, the initial step is the derivation of model visibilities for the contaminating source. Then, for each baseline in the array one compares the model visibilities with those observed and calculates time dependent complex gain factors. These gain factors are used to degrade the model visibilities toward those actually observed, and the degraded model visibilities are then subtracted from the observed visibilities. The algorithm to determine these gain factors is called MODCAL (after MODel CALibration). It is derived from a procedure originally developed by Lloyd Higgs to improve the dynamic range of the four antenna version of the ST that was in service before the current seven antenna system came on-line in 1991. An initial description of this algorithm was given in Willis & Higgs (1996).

In order to remove a contaminating source by using the MODCAL algorithm one proceeds in detail as follows:

- 1) Compute model visibilities M for this contaminating source and model visibilities E for the remaining sources in the field.
- 2) Delete the model visibilities E corresponding to all other sources in the field from the observed visibilities D to obtain visibilities D' .

$$D' = D - E. \quad (9)$$

It is desirable that the signal E contains a very good representation of the signal from all other sources in the field. This may not always be possible in practice, and some iteration is usually required (see Sects. 4.2).

D' should now represent the observed response to the contaminating source. At some particular instant in time the relationship between the observed visibility, D' , and the contaminating source model visibility, M , can be described by a simple equation of the form

$$D' = GM + C \quad (10)$$

where G is a gain factor and C is a quantity which represents the contribution to the visibility from the radio emission in the rest of the field. If E accurately represented the contribution to the observed visibilities from other sources in the field, C should have a minimal value. All quantities are complex.

- 3) If G and C are assumed to be approximately constant over some time interval, then one can now solve for G and C in the time interval by simple least-squares fitting techniques. If the combined real and imaginary

residuals are to be minimized one obtains the two complex equations

$$\sum_{i=1}^N D'_i M_i^* = G \sum_{i=1}^N M_i M_i^* + C \sum_{i=1}^N M_i^* \quad (11)$$

and

$$\sum_{i=1}^N D'_i = G \sum_{i=1}^N M_i + C \sum_{i=1}^N 1 \quad (12)$$

where M^* is the complex conjugate of the model visibility and N is the total number of visibility samples in the time interval. One can easily solve for the complex G and C terms by standard matrix inversion techniques.

At DRAO the above least squares analysis is usually done on one hour's worth of visibility data. Since a standard observation with this telescope requires 12 hours to completely sample the UV domain, one ends up with 12 values of G and C for each baseline. To obtain the MODCAL solutions for G and C one must assume that one is solving for a constant value of C over each time interval. This assumption is usually incorrect, and may affect the solution for G , especially for any baselines where C represents a significant contribution to the total visibility. As I will show in the next subsection, one can get around this problem by an iterative procedure.

4) Once the least-squares values of G have been obtained, one fits sequences of orthogonal polynomials to their amplitude and phase equivalents using the method of Peterson (1979). These best fitting polynomial sequences are then used to calculate a complex gain adjustment, G_{adj} , at each original sample point of a given baseline in the UV plane. Multiplication of the model visibility at this location, M , by G_{adj} gives one a best-fit contaminating visibility. Subtraction of this contaminating visibility from the original visibility D' yields a data visibility, D'' , which should not have any signal due to the interfering source.

$$D'' = D' - G_{\text{adj}} M. \quad (13)$$

5) Finally, to get a set of visibilities which represent the observed field without the contaminating source, add visibilities E back into the data.

$$D''' = D'' + E. \quad (14)$$

Images made from visibilities D''' should give a high quality view of the field.

One can do least-squares fitting on one-hour bins of DRAO data followed by further fitting of orthogonal polynomials because the DRAO antennas have equatorial mounts. Consequently most unknown instrumental responses to an incoming signal are expected to change slowly, if at all, as a function of time. This will not necessarily be the case for arrays of telescopes with altazimuth mounts, such as the VLA, and it is not clear that the MODCAL algorithm has any applicability to such arrays.

4.2. Removal of strong sources

If one wishes to remove the effects of sources such as Cyg A, Cas A and Tau A from a field, Cambridge or VLA images of these sources are normally used to compute model visibilities by the procedure described in the Appendix. These images contain Fourier components well beyond the maximum required by the DRAO Synthesis Telescope's longest baseline of approximately 600 metres.

The accuracy with which model visibilities can be computed for the DRAO ST from images obtained with other telescopes depends on how well one can emulate the instrumental responses of the ST. For example, the effect of bandwidth smearing on model visibilities is currently calculated under the assumption that the bandpass function is square (Eq. (A1)). However this theoretical bandpass deviates in shape from what is actually found. Recent measurements by Thorsley (unpublished DRAO technical memo) show that the 21 cm passbands are skewed and asymmetrical in shape, and differ from baseline to baseline. Consequently the model visibilities will be in error. However, this error should be compensated to some extent, by a counteracting adjustment to the gain factors G that will occur during the least-squares fit for the determination of G . Therefore the complex visibility, GM , that one wants to subtract from the original visibility, D' , should be relatively insensitive to errors in the modelling procedure.

(The asymmetric bandpasses are certainly one of the limiting factors in our ability to obtain high dynamic range in DRAO images by standard self-calibration procedures. Future enhancements to the modelling software will take this property into account.)

To get the best MODCAL results one proceeds in an iterative manner. Do an initial run of MODCAL on the original observed visibility data D to get a preliminary subtraction of the interfering source. This subtraction will not be perfect because the visibilities being fitted still contain the contribution E from the other sources in the field, and images produced from the MODCAL output visibilities may contain some minor artifacts near the field centre. These artifacts occur because of our intermingling of the G and C terms in the least squares fit. As one might expect, the artifacts tend to have low spatial frequencies - indicative of poor solutions on the short spacings, where the contribution of C will be dominant.

Once one has a preliminary removal of the interfering source, one can usually get a reasonable model of the emission at the field centre by a combination of cleaning and self-calibration. One can now proceed to improve the MODCAL solution by inverting this model into the UV plane and subtracting its visibilities (E) from the original observed visibilities D . The majority of the remaining signal D' will be due to the contaminating source; the C component of the observed visibilities should be minimal.

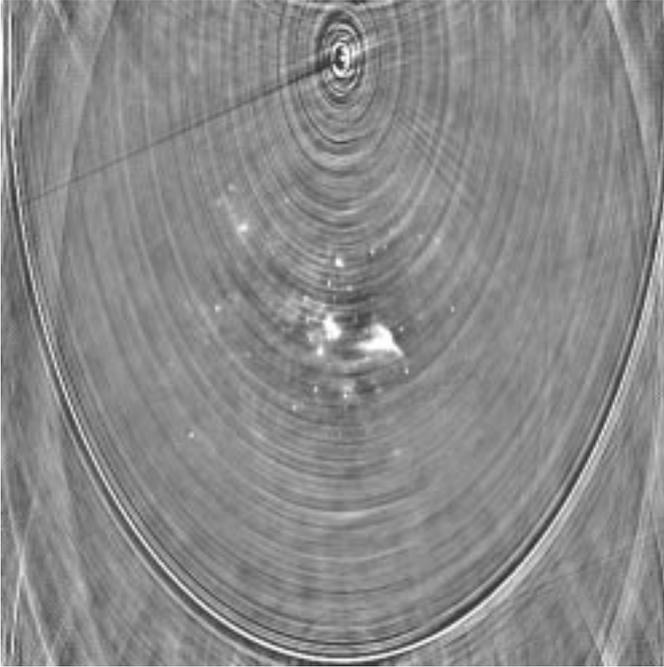


Fig. 7. A DRAO Synthesis Telescope 74 cm (408 MHz) image centred on the CTB 80 region. The total field of view is about 19.7 degrees. Cyg A is detected with a peak flux density of approximately 36 Jy at a distance of 8.2 degrees to the north of the field centre. Low level grating artifacts having amplitudes of 50 to 100 mJy prevent our studying the emission regions at the field centre. These grating artifacts cannot be successfully removed by conventional clean or self-calibration procedures

Then proceed to do a second MODCAL on these altered visibilities as described in Sect. 4.1. Now one will get a very good least squares estimate of the gain G regardless of baseline, and consequently, a very good subtraction of the contaminating source. Finally, add back to the second set of modified visibilities D'' the visibilities E generated by the clean component model of the emission at the field centre.

One now has a set of visibilities D''' from which one can generate images of the field undisturbed by the contaminating source. An example of an image corrected by the procedure is shown in Fig. 8.

4.3. Removal of solar interference

At 74 cm, solar interference can be a problem as the Sun is often detected in antenna sidelobes; it then produces significant fringes on shorter baselines. One does not want to flag short spacing data contaminated by solar emission unnecessarily, since short spacings contain information about the broad structure of a field. A simple solution which allows the astronomer to remove most of the solar interference is to proceed by approximating the solar image as a uniform disk with a radius of 17.5 arcmin.

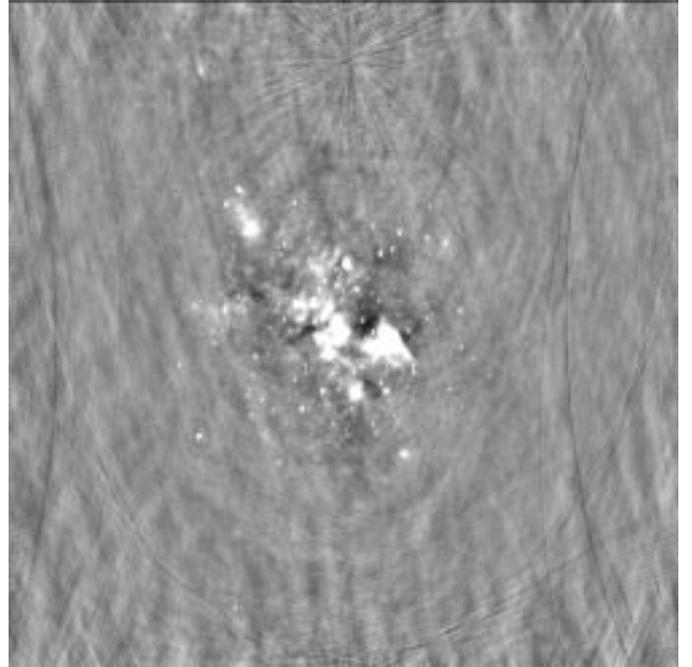


Fig. 8. The CTB 80 field after removal of Cyg A with the MODCAL procedure. All that is left of Cyg A are a few faint spokes and stripes centred at its position with amplitudes in the range 10 to 50 mJy. The astronomer can now study the radio emitting regions at the field centre to a level below 25 mJy per beam

The Fourier transform of a uniform disk is a first order Bessel function (Bracewell 1965); so one can easily calculate the visibility amplitude at a given distance from the UV origin as

$$A = RJ_1(2\pi RD)/D \quad (15)$$

where A is the visibility amplitude, R the solar radius in radians, D the distance from the UV origin in wavelengths, and J_1 is the first order Bessel function.

Using this equation for the amplitude and an ephemeris giving the Sun's position as a function of time, one computes model visibilities for those hour angles where the Sun was visible above the local DRAO horizon. The solar fringes can then be removed by application of the MODCAL algorithm.

An excellent example of solar removal by means of this method is seen in a CGPS survey field centred just to the southwest of the W3, W4, W5 complex. Figure 9 shows the image made from the original data. In addition to the objects of interest, a significant ripple is present in the field.

After application of the MODCAL removal procedure, one obtains the image shown in Fig. 10. The low level solar waves have completely disappeared, and the astronomer can proceed with cleaning, self-calibration etc.

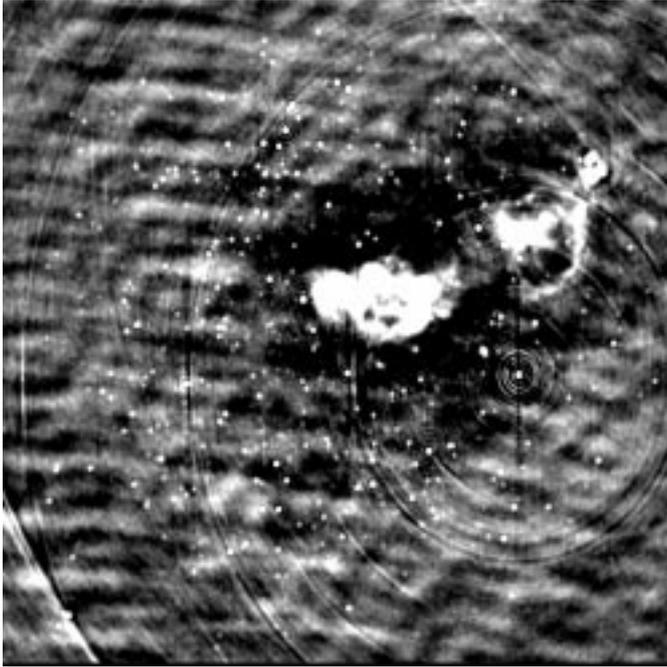


Fig. 9. A DRAO Synthesis Telescope 74 cm (408 MHz) image of an area to the southwest of the W3, W4, W5 complex. The field is 9.9 degrees across. The low level ripples with an amplitude of approximately 30 mJy that are present throughout the image are due to the Sun

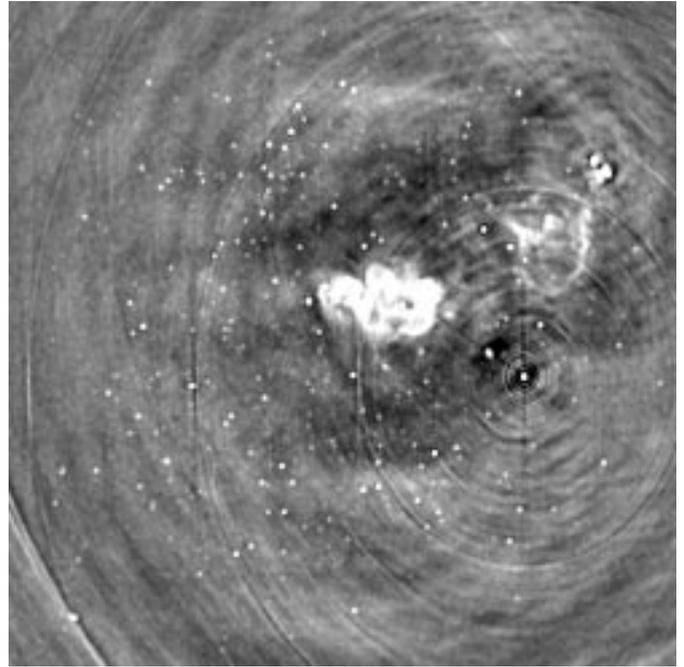


Fig. 10. The image of the W5 area after processing the visibilities with the solar removal procedure outlined in the article. Some low level systematic offsets with an amplitude of approximately 8 to 25 mJy remain. These irregularities are caused by a lack of short UV spacings in the synthesized image

4.4. Removal of system and man-made interference

The previous two subsections described how to remove interfering objects by starting with some predefined model, either a high resolution image, clean components, or an abstract mathematical definition. However the MODCAL removal procedure usually works very well to remove features that do not have either a predefined structure or known origin. Figure 11 shows another 74 cm field near Cyg A after initial removal of Cyg A by MODCAL. A faint grating ring is visible crossing through the lower middle part of the image. Some detective work revealed that the cause of this artifact is a “ghost” Cyg A (see Fig. 12) located 12 h away in right ascension from Cyg A but at the same declination! The exact cause of this artifact is not known but it is probably caused by some error in the 408 MHz fringe stopping or delay system. I infer this because the phase of an object varies according to a $\sin(t)$ term (Eq. (A3)) where t is the hour angle. Since $\sin(t + \pi)$ equals $-\sin(t)$, I suspect that incorrect fringe stopping or delay setting is somehow injecting a small percentage of the Cyg A signal into a negative sin term.

In any case one can remove the effects of this ghost signal from the region of interest quite easily. Simply use the image shown in Fig. 12 as the model, invert it into the UV plane, and apply MODCAL in the way described previously. The end result is shown in Fig. 13. The unwanted grating ring has disappeared.

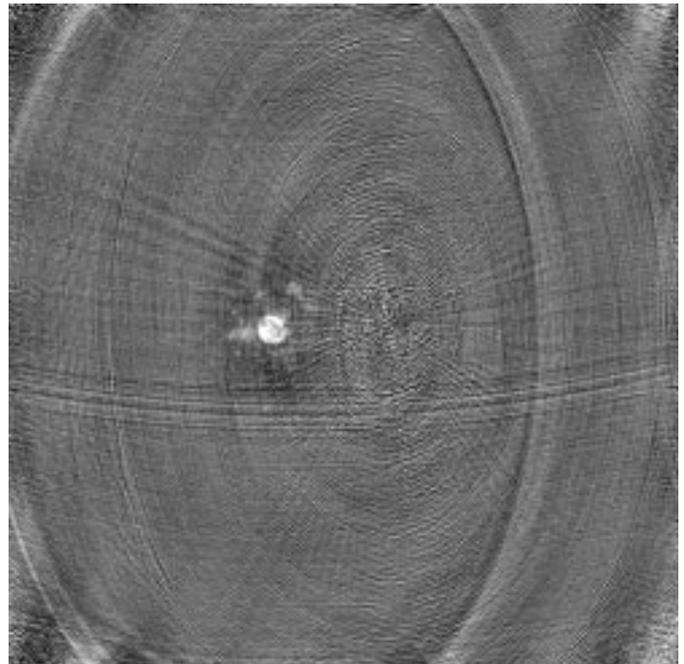


Fig. 11. 74 cm wide field image (23.9 degrees on one side) of a field after initial removal of Cyg A by the MODCAL algorithm. Cyg A has been removed reasonably well, but a faint low-level ring-like feature, having a systematic amplitude offset of 100 to 300 mJy, remains, stretching from the lower left to the lower right part of the image. In the original image Cyg A had an attenuated peak flux density of 2300 Jy. Diffuse emission from other objects in the field that is as faint as about 200 mJy can be seen

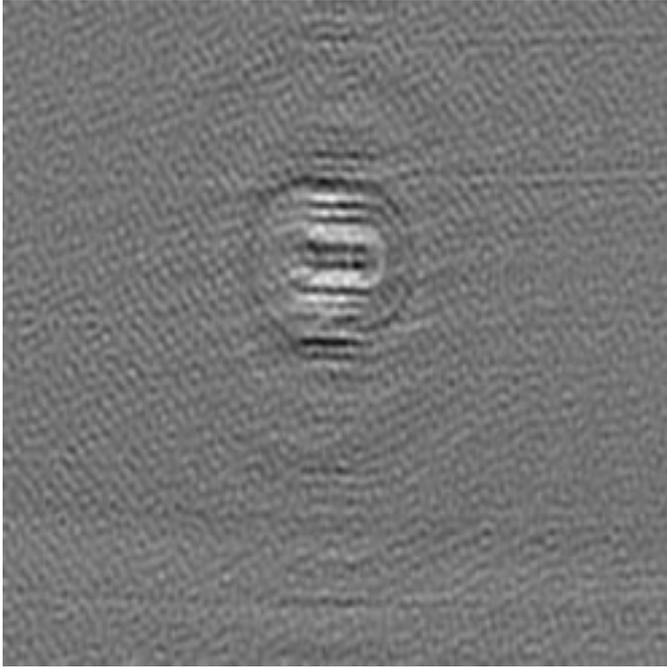


Fig. 12. 74 cm image (6 degrees across) of a field centred at J2000 α 7 h 59 m 27.92 s, δ 40 deg 44' 0.7", 12 hours away in right ascension from Cyg A but at the same declination. The "object" present in the image is responsible for the ring feature seen in Fig. 11. It has a maximum signal of 2 Jy. The RMS noise in the field is approximately 145 mJy

The astronomer can remove most ground-based man-made interference in a similar manner. Since a ground-based source is not moving, an interferometer data collection system usually interprets its signal as coming from a fixed point on the sky, namely the celestial pole. So all that needs to be done is to make an image of the region surrounding the pole, and use this image as the MODCAL model.

4.5. Removal of residual grating rings

It is not uncommon for low-level residual rings to remain in DRAO images after self-calibration, even when one adjusts for different antenna sizes. Such low level residual rings may be seen in Fig. 5. There are several potential causes of these residual effects. Firstly, holography measurements by Gray (in preparation) indicate that some dishes have significant surface irregularities. Secondly, as was discussed in Sect. 4.2 the correlator bandpasses have significant departures from the ideal rectangular shape. Both these effects are likely to introduce effects that either vary across the field (dish surface irregularities), or are baseline dependent (correlator bandpasses). In theory one could attempt to model these effects. In practice, it is easier to remove these effects using the MODCAL procedure.

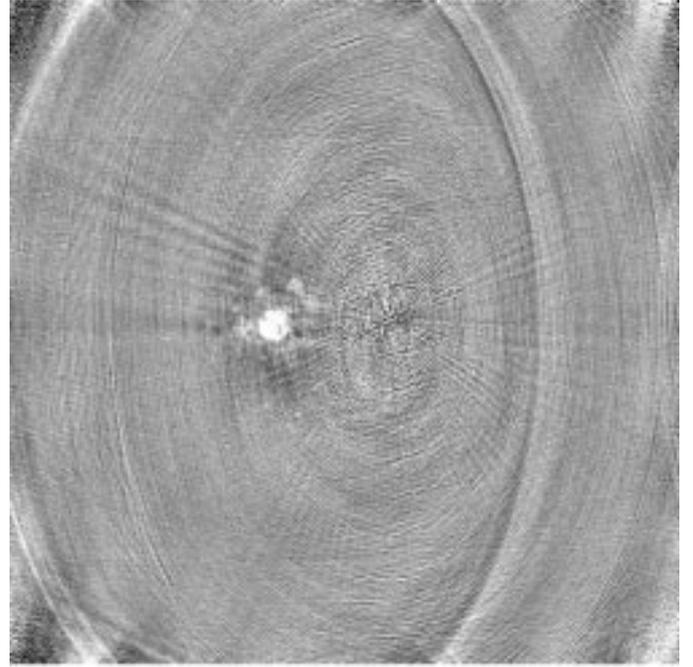


Fig. 13. A 74 cm image of the Cyg A area after use of MODCAL to remove artifacts due to the object seen in Fig. 12

All that is necessary is for each source that has residual rings:

- 1) Subtract out model visibilities corresponding to the other sources in the field.
- 2) Create model visibilities for the offending source and run the MODCAL algorithm to remove the baseline dependent effects of the offending source.
- 3) Add the model visibilities created in step 2 to the adjusted raw data visibilities and continue on to the next source. In Fig. 6 the background sources have had their residual rings removed by this method.

5. Conclusions

The algorithms described in this paper have proved very successful at producing high quality images for the Canadian Galactic Plane Survey. The MODCAL method has been especially useful in removing the effects of strong sources such as Cyg A from many fields. Further, the results presented in Sect. 4.4 and in Sect. 4.5 showed that it is effective at removing many artifacts of unknown origin from ST images. It is not, however, a panacea for all system problems. In both Fig. 8 and Fig. 13 faint residual spokes remain centred on the position of Cyg A. Although there is a suspicion that these spokes are due to antenna pointing errors, efforts to find an effective procedure to remove them have been unsuccessful to date.

Continued improvement in the dynamic range of the telescope should be possible by further incorporation of

measured instrumental characteristics, such as the bandpass shapes mentioned in Sect. 4.2, into the procedures for computing model visibilities and correction factors.

Acknowledgements. I thank Lloyd Higgs for suggesting the adoption of the MODCAL algorithm for use with seven-antenna ST data. S. Dougherty functioned as chief tester of many of the routines described here while he reduced CGPS data. The referee, J. Hamaker, proposed numerous improvements to the presentation of the paper. The DRAO Synthesis Telescope is operated as a national facility by the National Research Council of Canada.

Appendix A: Computation of model visibilities

Most of the procedures described in this paper require the computation of model visibilities, which are then compared with actual data. Accurate computation of these model visibilities is therefore important.

Most standard self-calibration techniques derive the model visibilities from Fourier inversion of one or more clean components (see e.g. Cornwell & Fomalont 1989). Clean components have the advantage that they represent a deconvolution of an observed image into a set of point-components. As such they usually, but not always, reflect exactly what was actually observed; if the observed visibility data are distorted by instrumental effects, the distribution of clean components will generally reflect this distortion. Consequently model visibilities generated from the clean components should appear similar to those observed.

However, if one wishes to compute model visibilities from objects other than clean components, then instrumental effects must usually be taken into consideration. Two common instrumental effects that cause distortion of images obtained by a radio interferometer are bandwidth smearing and time-averaging of raw visibility data. These effects cause the image of a point source to become increasingly distorted as its distance from the field centre increases.

An excellent introduction to the effects of bandwidth smearing and time averaging of visibility data on the resulting image is given in Sects. 6.3 and 6.4 of Thompson et al. (1986). In summary, bandwidth smearing will distort an image parallel to the direction to the field centre while time averaging distorts the image perpendicular to the field centre direction. These effects become increasingly important the larger the distance from the field centre becomes. When one deconvolves an image into clean components the distribution of the clean components reflects this distortion and their corresponding model visibilities will suffer the same attenuation as those observed (see e.g. Cotton 1989).

However, there are a number of cases where use of clean components is not the best way to derive model visibilities. Because the DRAO ST has a wide field of view and

is predominantly used to observe that part of the galactic plane visible in the northern hemisphere, there is a good chance that any single field is contaminated by responses to one of Cyg A, Cas A, or Tau A.

These sources, which have angular sizes in the range 2 to 5 arcmin are usually only modestly resolved by the DRAO synthesized beam, which has an east-west size of 3 and 1 arcmin at 74 and 21 cm wavelength respectively. Decomposition of these sources into clean components and calculating visibilities based on clean components does not seem to work particularly well. There are several reasons for this; firstly clean components do not form a particularly good model for a slightly resolved source (Briggs & Cornwell 1993). Secondly, these sources are often detected at very large distances from the field centre, in the side-lobes of an antenna with unknown phase and amplitude response. Currently, visibility data on all ST baselines is averaged to a 90 second integration period. If a source is very far away from the field centre, not only does simple time smearing distort the sampled data, but aliasing due to under sampling of the fringes may occur. In such cases, clean components may represent a ‘‘Global Average’’ of the data, but produce relatively poor agreement between model visibilities and actual data on any given baseline.

Better model visibilities for these bright sources are generally obtained by taking a high resolution image of the source obtained by another array that includes spatial information out to wavenumbers much higher than can be sampled by the DRAO interferometer, and inverting this image back into the UV plane. Since the strong sources the ST detects are usually at quite large distances from the field centre, one must then explicitly correct the model visibilities produced by this method for both time and bandwidth smearing effects.

Assuming that the bandpass function is square, one can describe the bandwidth attenuation, A , at a given UV location D , for a source at a distance ξ from the field centre, by a sinc function ($\sin(x)/x$) where, for bandwidth smearing,

$$x = \pi D \xi B / c. \quad (\text{A1})$$

Here B is the bandwidth and c is the velocity of light (see Thompson et al. 1986). (In Sect. 4.2 of this paper, I discuss the limitations of this assumption in the case of the DRAO ST.)

The attenuation due to time smearing is also a function of increasing distance from the field centre and of increasing baseline length (Eq. (6.67) of Thompson et al. 1986). One can model the time smearing effect by computing model visibilities in some sub increment of the actual time integration period used for the observations, and then averaging the computed model visibilities together over each actual integration period. At present the DRAO software package uses a sub increment length of 10 seconds, or 1/9 of the final integration period.

If one wishes to simply derive model visibilities for one source then the fastest way to generate them is to FFT

the high resolution image into the Fourier domain and then interpolate from the UV grid to derive the visibility function at an observed UV location. Visibilities derived by this method can then be corrected for bandwidth and time smearing by application of the equations and methods outlined above.

This procedure already introduces one source of potential error: we must use some sort of interpolation function to obtain visibilities at observed UV locations from values computed at UV grid locations.

However, if an FFT inversion of a model image which contains components at widely differing distances and directions from the field centre is done and one wishes to properly correct the FFTd visibilities for bandwidth and time smearing, there is an additional problem: the degridded model visibilities already represent a summation of the visibilities due to the different sources of emission. At best, one can only compute some sort of weighted mean correction for time and bandwidth smearing.

More accurate model visibilities can be obtained if one starts from the actual definition of the phase term in the Fourier relationship between a celestial source and an observed visibility as seen by a radio interferometer

$$\text{phase} = 2\pi(Ul - Vm) \quad (\text{A2})$$

where U and V are the UV locations of the sampled data point, and l and m are the direction cosines of the source on the celestial sphere. For an east west interferometer, such as the ST, this equation reduces to the rather simple result

$$\text{phase} = 2\pi R(\cos \delta \sin t - \cos \delta_o \sin t_o) \quad (\text{A3})$$

by substitution of the appropriate definitions for U , V , l and m (see e.g. Brouw 1971, p. 25), here R is the actual baseline length in wavelengths, δ and t are the declination and hour angle of the source (or of each pixel in the image), and δ_o and t_o are the declination and hour angle of the phase centre of the array.

If one uses this direct Fourier transform relationship for each sky pixel one can compute correct model visibilities at each observed UV location and make a proper correction for bandwidth and time smearing. This approach also allows one to correct for other instrumental distortions that are a function of direction in the sky. The final model visibility is obtained by summing the visibilities corresponding to each image pixel. Although use of this method implies a considerable increase in computer time over that required for FFT methods, computation

of model visibilities is an excellent example of what computer scientists call an ‘‘embarrassingly parallel’’ application. The work can be easily spread out over a network of workstations.

Note that an image (as opposed to clean components) represents a convolution of emission from the sky with the point spread function (PSF) of the telescope that made the image. Therefore, when an image is inverted into the Fourier domain, the intrinsic visibilities will be multiplied by the Fourier transform of the originating instrument’s PSF. Since the images inverted by this procedure normally have a considerably higher resolution than does the ST, in practice this means that there is an arbitrary amplitude gain ratio between model visibilities generated by this method and actual visibilities observed by the ST. The easiest way to determine the value of this ratio is to make an image from the model visibilities, and compare the maximum signal in this model image with that in an image derived from the actual observations. The model visibilities can then be scaled by the value found for this ratio.

References

- Bracewell R., 1965, *The Fourier Transform and Its Applications*. New York: McGraw-Hill, p. 248
- Briggs D.S., Cornwell T.J., 1993, in *IAU Symposium 158, Very High Angular Resolution Imaging*, Robertson J.G., Tango W.J. (eds.). Dordrecht: Kluwer
- Brouw W.N., 1971, Ph.D. Thesis. Leiden University
- Cornwell T., Braun R., 1989, in *Synthesis Imaging in Radio Astronomy*, Perley R.A., Schwab F.R., Bridle A.H. (eds.) ASP Conf. Ser. 6, 167
- Cornwell T., Fomalont E.B., 1989, in *Synthesis Imaging in Radio Astronomy*, Perley R.A., Schwab F.R., Bridle A.H. (eds.) ASP Conf. Ser. 6, 185
- Cotton W.D., 1989, in *Synthesis Imaging in Radio Astronomy*, Perley R.A., Schwab F.R., Bridle A.H. (eds.) ASP Conf. Ser. 6, 233
- Perley R.A., 1989, in *Synthesis Imaging in Radio Astronomy*, Perley R.A., Schwab F.R., Bridle A.H. (eds.) ASP Conf. Ser. 6, 259
- Peterson D.M., 1979, *PASP* 91, 546
- Thompson A.R., Moran J.M., Swenson, G.W., 1986, *Interferometry and Synthesis in Radio Astronomy*. New York: Wiley
- Willis A.G., Higgs L.A., 1996, in *Astronomical Data Analysis, Software and Systems V*, Jacoby G.H., Barnes J. (eds.) ASP Conf. Ser. 101, 211