

**Multi-Frequency Synthesis with
MERLIN:
A New Technique in Aperture
Synthesis Imaging.**

A thesis submitted to the
University Of Manchester
for the degree of
Doctor of Philosophy
in the
Faculty of Science
by
John E. Conway

June 1988

Abstract

The fidelity of radio images of complex objects produced by the aperture synthesis technique depends on the quality of the aperture coverage (or uv coverage) of the array used. The application of non-linear deconvolution algorithms to images made from uv data which are insufficiently sampled will result in so called 'reconstruction errors'.

If the aperture coverage can be improved these reconstruction errors can be reduced. This is best achieved by adding more antennae to an array, in practice however this method of improvement is limited by economic factors. An alternative is to note that the dimensions of the uv plane are measured in wavelengths and therefore by varying the wavelength of observation the uv coverage can be filled in. To this idea (by no means new) I give the name 'Multi-Frequency Synthesis' (MFS).

The obvious problem with MFS is that the source structure will in general vary with frequency. Fortunately, from both observation and theory it seems that we can generally expect that this frequency dependence will have a relatively simple form. At each point on the source the intensity will usually be a power law with frequency with an exponent (the 'spectral index') which varies with position. Most of this thesis will be spent describing the effects that such spectral behaviour introduces into our images and then considering algorithms for removing these effects.

A mathematical description of the form of error introduced by spectral variations of this type is presented. These errors are then described both qualitatively and quantitatively for the specific case of MFS observations with the MERLIN array. This array is typical of many sparse arrays often found in VLBI. The effect on the level of error of many different observing parameters such as the range and number of frequencies, declination etc are discussed. Also considered is the influence of different source structures. It is found that for realistic MFS observations the size of the errors are relatively small and for many sources below the expected thermal noise.

For other sources, in order to successfully use the MFS technique, special

processing algorithms must be applied to remove the significant spectral effects. Various approaches are considered. The most promising algorithm is called 'Double Deconvolution' and is based on extending the principles behind the CLEAN algorithm. This algorithm is tested on simulated MERLIN MFS data and shown to perform well. The simulation consists of imaging a model source which should be typical of the more difficult targets envisaged for MERLIN MFS. Given this success it appears likely that the problem of spectral effects are essentially solved for the MERLIN MFS case.

An analysis is made of the practical problems in carrying out MFS. The technical details of a working MFS system are briefly considered as are the constraints imposed by RF interference. Methods of phase and amplitude calibration of MFS data are also discussed. The effects of non-power law spectra are evaluated and found not to be significant at cm wavelengths except in the cores of extragalactic radio sources. In this case simple additional processing should remove the spectral effects.

Finally, real MERLIN MFS observations of the quasar 3C179 are described. At present the range of frequencies accessible to MERLIN is restricted by its RF technology. This means that the improvement in aperture coverage possible using MFS is therefore limited. Despite this limitation a slight but definite improvement in fidelity is achieved upon images produced by conventional single-frequency techniques.

Previous Study

The author attended Downing College, Cambridge and obtained an Honours degree in Natural Sciences (Physics) from the University of Cambridge in 1984. He registered for the M.Sc. course in Radio Astronomy at Jodrell Bank in October of the same year. After being awarded a Diploma for Advanced Studies in Science in 1985 he registered as a Ph.D student at the University of Manchester in October 1985.

Acknowledgements

I wish to thank Professor Sir Francis Graham Smith for the use of the facilities at Jodrell Bank and the Science and Engineering Research Council for a Research Studentship. Thanks also to the technical staff at Jodrell, for their often unsung efforts to run MERLIN, and to the programmers both at Jodrell and in the AIPS group for the many man-years of software without which radio imaging would be impossible. Personal thanks to Ken Harrop for doing the photographs so quickly. Thanks also to Janet Eaton for help with the typing.

I wish to thank all my friends at Jodrell for just 'being there', for sharing laughter, late night conversation and good and bad times. I especially want to thank Harbie, Dave and Simon. They have had more than their fair share of having to live with me, and have (almost) got used to my bad habits. I wish to remember the MSc class of '84-'85 including Genevieve, Andy Wells, Dave Astbury, Jenanne and everybody else who made up the best MSc year ever. Thank you too to all the other students, too numerous to mention, who have helped provide a human as well as an academic home for me at Jodrell.

Thanks to my family for their support and encouragement, both in the last three years and throughout my academic career.

Thanks to Tim Cornwell for useful discussions. Finally, I must thank my supervisor Dr P.N. Wilkinson for proposing an interesting thesis topic, for sorting out priorities and most of all for his constant enthusiasm and interest in my work.

Declaration

Except where stated the work contained in this thesis is that of the author alone. Chapter 4 considers some algorithms proposed by other authors but the present author is solely responsible for the discussion and interpretation of these algorithms. In Chapter 6 the discussion of a future RF system for MERLIN MFS is based on discussions with J.A. Battilana and others. Software used for the analysis and processing of data throughout this thesis is generally the work of others unless otherwise stated.

The work presented in this thesis is that of the author except where stated otherwise. No portion of this work has been submitted in support of an application for another degree or qualification at this or any other university or seat of learning.

J.E. Conway

John E. Conway
Nuffield Radio Astronomy Laboratories,
Jodrell Bank,
Macclesfield,
Cheshire.
June 1988.

To Dad, Mum and Sue.

For we know in part and we prophesy in part.

I Corinthians, 13. v 9.

For now we see through a glass darkly; but then
face to face: now I know in part; but then shall
I know even as I am known.

I Corinthians, 13, v 12.

Conventions

The radio spectral index is defined according to $S \propto \nu^{-\alpha}$.

The abbreviation FWHM means the Full Width Half-Maximum of a gaussian.

The abbreviation FT denotes a Fourier Transform.

Contents

1	Introduction	1
1.1	General Introduction	1
1.2	The Development of Aperture Synthesis Imaging	1
1.2.1	Early Aperture Synthesis	1
1.2.2	Deconvolution	4
1.2.3	Aperture Synthesis with Sparse Arrays	5
1.3	The MFS Idea	7
1.4	History of the MFS Idea	10
1.4.1	Early MFS Ideas	10
1.4.2	Modern MFS	17
1.5	Implementing MFS :	21
1.6	MFS on MERLIN	22
1.6.1	Introduction	22
1.6.2	MERLIN UV Coverage and MFS	23
1.6.3	Imaging Specifications for MERLIN MFS	25
1.7	Astrophysical Impact of MFS	28
1.8	Outline of Thesis	29
2	MFS and Frequency Dependent Structure I : The Form of the Spectral Effects	30
2.1	Introduction	30
2.2	The Nature of Frequency Dependent Structure	30
2.2.1	Frequency Dependent Structure and MFS	30
2.2.2	Synchrotron Emission Spectra	31
2.2.3	The Origin of Power Law Spectra	34
2.3	Effects of Power Law Spectra on MFS	38
2.3.1	Introduction	38
2.3.2	Non-Spatially-Varying Spectral Index Distributions	39
2.3.3	SVPL Spectra and the Linear Expansion	41

2.4	Some Features of the Spectral Dirty Beams	45
2.4.1	The SDB's and their UV Weights	45
2.4.2	The Centres of the SDB's	45
2.4.3	The Sidelobes of the Beams	48
2.4.4	The Derivative Relationship Between Beams	49
2.4.5	Summary	53
2.5	The Size of Linear Spectral Errors	53
2.5.1	The Coefficient Images	53
2.5.2	Linear Spectral Errors from a Point Dominated Source .	55
2.6	The Logarithmic Expansion	58
2.6.1	Description and Properties of the Logarithmic Expansion	58
2.6.2	Logarithmic Spectral Errors from a Point Dominated Source	62
3	MFS and Frequency Dependent Structure II: Spectral Effects in Realistic Cases	65
3.1	The D_1 Beam as a Function of Observing Parameters	65
3.1.1	Effects of Array Geometry	65
3.1.2	The Fractional Bandspread	71
3.1.3	The Frequency Coverage	72
3.1.4	Declination Effects	75
3.1.5	Summary	75
3.2	Spectral Effects from Extended Sources	77
3.2.1	First Order Spectral Effects	77
3.2.2	Simulations of First Order Spectral Effects in Realistic Sources	82
3.2.3	Simulations of Second Order Spectral Errors	86
3.3	Effects of Deconvolving Images Containing Spectral Errors . . .	91
3.3.1	Requirements For The Convergence of CLEAN	91
3.3.2	Simulations	92
4	Algorithms for Removing Spectral Effects 1	96

4.1	Classifying the MFS Problem	96
4.1.1	Introduction	96
4.1.2	MFS and the Size of Spectral Errors	96
4.1.3	MFS and Aperture Sampling	99
4.1.4	The Oversampled Limit	101
4.2	Map and Combine	103
4.2.1	The Algorithm	104
4.2.2	Effects of Single Frequency Reconstruction Errors	105
4.2.3	Correlation Between Single Frequency Reconstruction Errors	106
4.2.4	Simulations of Single Frequency Reconstruction Errors	106
4.2.5	Testing Map and Combine	107
4.3	Map and Stack	111
4.3.1	The Algorithm	111
4.3.2	Effects of Single Frequency Reconstruction Errors	112
4.3.3	Effectiveness of the Algorithm	113
4.4	Data Weighting Methods	116
4.4.1	Double-Sampled Data	116
4.4.2	The Cornwell Method	118
4.4.3	The Gridded Cornwell Algorithm	120
4.4.4	Simulations of the Gridded Weighting Algorithm	121
4.4.5	Other Approaches to Weighting	124
4.5	MEM Based Methods	127
4.6	Summary	130
5	Algorithms for Removing Spectral Effects II: Double Decon- volution	132
5.1	The Double Deconvolution Philosophy	132
5.2	Implementing the Double Deconvolution Philosophy	133
5.2.1	A CLEAN Based Double Deconvolution	133
5.2.2	An Illustration of A CLEAN based Double Deconvolution.	136
5.2.3	Double Deconvolution as a Generalised CLEAN	138

5.2.4	The Mechanism of Convergence	138
5.3	Convergence Conditions for Double Deconvolution	145
5.3.1	Introduction	145
5.3.2	Convergence of CLEAN	146
5.3.3	Convergence of DD to Fit MFS Data	147
5.3.4	Separating the I and $I\alpha'$ Distributions	149
5.3.5	Summary	152
5.4	Simulations of Double Deconvolution	153
5.4.1	Software Implementation	153
5.4.2	The Simulation Target	154
5.4.3	Control of the DD Simulation	156
5.4.4	Adding VLA Data	157
5.4.5	The Performance of the DD Simulation	159
5.4.6	Reducing the Effects of Undercleaning	164
5.4.7	Testing VLA Assisted Double Deconvolution	165
5.5	Future Developments	168
5.5.1	Improved Software Implementations of DD	168
5.5.2	Other Implementations of the DD Philosophy	171
6	Practical Aspects of MFS	173
6.1	Introduction	173
6.2	Interference and MFS	174
6.2.1	Introduction	174
6.2.2	Effects of Uncorrelated Interference	176
6.2.3	Effects of Correlated Interference	177
6.2.4	Results of Interference Surveys in L-band	179
6.2.5	Implications of L-Band Interference for MFS	185
6.2.6	Results of Interference Surveys Around 408 MHz	188
6.3	RF Technology and the Implementation of MFS	189
6.3.1	The Present L-Band MERLIN RF System	189
6.3.2	LO Generation and RF Filters for MFS	191
6.4	Phase Self-Calibration and MFS	194

6.4.1	Introduction	194
6.4.2	Self-Calibration and Reconstruction Errors	195
6.4.3	Parallel Self-Calibration	197
6.4.4	Linking Phase Errors	198
6.4.5	Results of Frequency-Switching Observations	202
6.4.6	Possible Methods of Phase Linking	205
6.4.7	Phase Referencing	208
6.4.8	Summary	209
6.5	Amplitude Calibration	210
6.5.1	Introduction	210
6.5.2	External Calibration	210
6.5.3	Self-Calibration	211
6.6	Non-Power Law Spectra	213
6.6.1	Introduction	213
6.6.2	Spectral Ageing	214
6.6.3	Synchrotron Self-Absorbtion	219
6.6.4	Free-Free Absorbtion	225
6.6.5	Superposition of Power Law Spectra	226
6.6.6	Summary	226
7	MERLIN MFS Observations of 3C179	228
7.1	Introduction	228
7.1.1	Previous MFS Observations	228
7.1.2	MFS and the Present MERLIN RF System	229
7.1.3	Selecting a Target Source	230
7.1.4	3C179 as an Astronomical Target	231
7.1.5	Expected Spectral Effects	231
7.2	Observations and Initial Processing	237
7.2.1	Observations	237
7.2.2	Calibration	240
7.2.3	Processing the MFS Data	243
7.2.4	The Single-Frequency Images	244

7.3	MFS Imaging	246
7.3.1	Combining the Frequencies	246
7.3.2	Complementary VLA Observations of 3C179	247
7.3.3	Combining MERLIN and VLA Data	252
7.3.4	Comparison of MFS and Conventional Images of 3C179	255
7.4	Conclusions	260
8	Conclusions and Suggestions for Future Work	261
8.1	Summary	261
8.2	Priorities for Future Work	267
8.2.1	Investigation of the Effects of Phase Errors on MFS	267
8.2.2	Amplitude Errors and MFS	269
8.2.3	A User-Friendly Implementation of the DD Algorithm	269
8.2.4	Understanding and Developing DD	270
8.2.5	Developing ‘Phase Linking’	270
8.2.6	Further Interference Surveys	271
8.2.7	Core Subtraction Simulations	271
8.2.8	Extensions of MFS	271
8.3	The Impact of MFS on MERLIN and VLBI	272
8.3.1	Implementing MFS on Long Baseline Arrays	272
8.3.2	Astronomical Targets for MFS with MERLIN and VLBI	275
8.3.3	Observations of Jets	276
8.3.4	Radio Spectral Studies	277
8.3.5	Radio/Optical Comparisons	278
8.3.6	Proper Motions	279
8.4	The Impact of MFS on Other Arrays	280
8.4.1	MFS on the AT	280
8.4.2	MFS on the VLA	282
8.4.3	MFS on the GMRT	282
8.4.4	MFS on the 5km Telescope	283
A	The Derivative Relationship for Linear SDB’s	286

B The Logarithmic Expansion	288
C Proof that the MFS Dirty Image is In Range of D_0	293
D The Gridded Cornwell Algorithm	297
E The Convergence of Double Deconvolution	303
F Separating the Effects of the I and $I\alpha'$ Distributions, with Double Deconvolution	307
Bibliography	311

Chapter 1

Introduction

1.1 General Introduction

History shows that progress in the subject of radio astronomy has largely been determined by the progress in the techniques of observation (Edge and Mulkay 1976). From the very earliest days the development of new instruments and observational techniques has occupied a large fraction of the effort of radio astronomers and this attention to technical advance has led to important astrophysical discoveries. Even today when radio astronomy is a much more mature subject considerable effort is still rightly spent on improving the techniques and technology of observation.

This thesis will present a new imaging technique, called ‘Multi-Frequency Synthesis’ (MFS). This technique should significantly improve the reliability of continuum images made by aperture synthesis methods. In the next section a brief history of aperture synthesis is presented to set the scene and demonstrate the need for methods to improve the fidelity of imaging. Some familiarity with aperture synthesis methods is assumed, for a simple introduction to the technique the reader is referred to Fomalont and Wright (1974) and for a more rigorous description to Swenson and Mathur (1968).

1.2 The Development of Aperture Synthesis Imaging

1.2.1 Early Aperture Synthesis

The realisation that the correlation between signals from two separated elements of an interferometer equaled a Fourier component of the sky brightness

distribution was present from the earliest days of radio astronomy (see e.g. Bracewell 1984). This principle was used by several groups in various ways to image the radio sky or set limits on source structure (Hanbury Brown et al 1952), but it was the group at Cambridge (see e.g. Scheuer 1984) that most successfully incorporated the above principle into a coherent method of imaging. The principle of synthesizing a large aperture from a series of interferometer observations was first fully described by Ryle and Hewish (1960) and the idea of using the Earth's rotation to achieve a well filled synthetic aperture though effectively used by O'Brien (1953) was first made explicit in the design of the One-Mile telescope (Ryle 1962) which was the prototype of all later East-West interferometer arrays.

The corresponding theory for the processing of data from synthesis arrays had earlier been dealt with by Bracewell (1958) who showed that the source brightness distribution could be determined uniquely (at a given resolution and within the noise) provided that the sampling of the synthetic aperture was sufficiently dense and regular. Essentially he showed that aperture synthesis imaging had a sampling requirement analogous to a two dimensional Shannon sampling theorem, a criterion well known in the analysis of time domain signals in signal processing. The design of the One Mile Telescope (Ryle 1962) utilised one continuously movable antenna on a railway track so that the aperture could be sampled as densely as was necessary to achieve a unique reconstruction.

Except for later modifications to deal with baselines which were other than East-West (Rowson 1963), Ryle's concept of an array that used movable elements and the Earth's rotation to fill a synthetic aperture together with the sampling theory of Bracewell (1958) gave a complete technique of radio imaging. Expanded versions of the One Mile Telescope were built at Cambridge (i.e. The 5-km Telescope (Ryle 1972)) and in the Netherlands (The Westerbork Synthesis Radio Telescope or WSRT (Baars et al 1973)) and proved very successful.

The work of Bracewell showed that on Fourier inverting the uv data the so

called ‘principal solution’ (also often referred to as the ‘dirty map’) would be equal to the true image convolved with a synthesised beam (the ‘dirty beam’) which is the Fourier transform of the uv sampling function. Even if the uv plane is completely filled out to the longest baseline this coverage is uneven because of the sharp transition from the sampled region to the unsampled region, this sharp transition gives rise to ‘sidelobes’ in the synthesised beam response. One way of reducing this effect is to multiply the samples before inversion by parameters(‘weights’) which decrease with increasing baseline length so reducing the sharpness of the edge to the aperture coverage (i.e. effectively forcing the sky plane image to be band-limited so ensuring that the Bracewell sampling theorem strictly holds). Such a technique, which is sometimes referred to as ‘grading’ or ‘apodization’ was applied to data both from the One Mile telescope and the WSRT array (Somerén Greve 1973). It can be argued that although successful such a technique is very wasteful of the data on the long baselines, since this data is weighted by very small values and therefore is effectively discarded.

Another type of artifact that occurs on directly inverting data from the above East-West arrays is that of ringlobes or ‘grating responses’. The E-W baselines are generally arranged to have baseline lengths which are multiples of some unit separation, such an array is termed a ‘grating array’. It gives rise on the sky to a grating response which has large sidelobes occurring at radii which are multiples of a radius which is inversely proportional to the size of the unit separation. By arranging the unit separation to be small enough the diameter of the ringlobe can be made much larger than the field of interest. Very strong sources a long way from the source of interest can however produce significant ringlobes that lie over the source. To ensure that such ringlobes are never important requires very dense sampling of the uv plane so that the whole primary beam of the antennae can be reliably imaged. Even when the antennae can be relocated it is sometimes impractical to achieve sufficient sampling because of the constraints of observing time and possible problems with bad data due to equipment failure etc.

In order to achieve the removal of sidelobes and ringlobes in cases of irregular or imperfect sampling the CLEAN algorithm was devised by Högbom (1974). The FT of the uv data (the ‘Dirty Map’) consists of the true image convolved with a dirty beam containing the ringlobe and sidelobe responses. The aim of the CLEAN method is to remove the effects of this dirty beam, it is therefore an example of a so-called ‘deconvolution’ method.

The CLEAN algorithm was first employed successfully on the WSRT but it was increasingly applied to data from arrays of small numbers of randomly placed antennae (‘sparse arrays’). The application of CLEAN and other deconvolution methods have today become indispensable to the processing of data from these arrays.

1.2.2 Deconvolution

The CLEAN algorithm works by searching for dirty beam-like artifacts in the dirty map and then iteratively removing the effects of these beam-like artifacts. Although proposed as a semi-empirical method for improving images later authors (including Schwarz 1978) have shown mathematically that under certain conditions the algorithm will converge. In removing sidelobes the algorithm implicitly extrapolates the model uv data (the Fourier transform of the image estimate) into the gaps between the ellipses in the uv plane on which the visibility is sampled. The CLEAN algorithm also reduces the contribution of the near-in sidelobes caused by the sharp edge to the uv coverage by effectively extrapolating slightly beyond the maximum baseline length. The two properties of extrapolation and interpolation applied to the observed data rely on the *a priori* knowledge contained within the processing algorithm.

The CLEAN algorithm contains an in-built *a priori* preference for point-like structure and since the source model is built up point by point the image that emerges is usually one of the most compact images that is consistent with the data. Incomplete uv coverage ensures that there are an infinite number of different images that can fit the observed data. CLEAN chooses one of these based on the above *a priori* bias.

One shortcoming of CLEAN compared to other deconvolution methods is that the *a priori* information contained within the CLEAN algorithm is implicit rather than explicit because the algorithm is defined solely as a set of procedures (Tan 1986). Other processing algorithms such as the Maximum Entropy Method (MEM) (e.g. see Narayan and Nityananda 1986 for a review) have more explicit definitions of the *a priori* assumptions used and find solutions to problems explicitly posed as selecting an image that fits the data whilst maximizing some function of the image (Titterton 1985).

Both CLEAN and MEM are non-linear algorithms, as must be any algorithm which is required to interpolate into unmeasured regions of the uv plane (Nityananda and Narayan 1982). In a similar manner to the way non-linear filters are required in electronics to generate new temporal frequencies from a set of existing frequencies a non-linear filter is required here to generate unmeasured spatial frequencies from the measured spatial frequencies. The non-linearity of the processing algorithms means that modest improvements achieved in the uv coverage can give larger improvements in image quality. On the negative side if the quality of the uv coverage becomes very inadequate for the complexity of the source the non-linearity can give rise to significant artifacts or ‘reconstruction errors’. For the CLEAN algorithm these artifacts often take the form of ‘ripples’ (Cornwell 1983) (see section(4.2.2)).

1.2.3 Aperture Synthesis with Sparse Arrays

With the advent of synthesis arrays which had only a small number of elements with no regular distribution as in the VLBI arrays, the problems of sidelobes became much greater than with regular arrays. Rather than being concentrated into ringlobes of large diameter the sidelobes often take the form of large subsidiary maxima close to the main lobe, which can have levels of ten or twenty percent of the peak. With such arrays there is usually no possibility of moving the antennae to fill the uv plane and so reduce sidelobe levels.

To obtain reasonable images from these arrays the use of non-linear processing algorithms becomes essential, since it is necessary to interpolate the

image transform across large and irregular gaps in the uv plane. Reconstruction is made even more difficult by the fact that observations from these long baseline arrays are also seriously contaminated by the effects of ionospheric or tropospheric propagation effects which contaminate the recorded phase of the visibilities. So called ‘self-calibration’ algorithms have been developed (e.g. see Pearson and Readhead 1984 for a review) which can largely correct for these phase errors, however the application of these algorithms at the same time as deconvolution further complicates the process of imaging and gives still greater potential for reconstruction errors to be introduced.

The great demands on image reconstruction algorithms when complex images are made from sparse uv coverage has led to increasing reliance on *a priori* knowledge as against observational knowledge. One commonly used variant of the CLEAN algorithm for instance allows the interactive placing of ‘windows’ in regions in which the user believes real flux exists; the search for flux is then confined to regions within such windows. The placing of the windows can be made based on reliable information about the source brightness distribution at a different resolution or frequency in which case the process can be justified. In other cases the placing of windows is done in an exploratory manner using the judgement and experience of the user, in which case the justification can only be on the basis of results.

Whilst some forms of *a priori* knowledge are rigorously based, such as the requirement that the source is positive to within the noise, other forms of constraints such as the window constraints described above are often more provisional. If the assumptions made are actually wrong then the resulting images will also be wrong. In the worst possible cases when the data are sparse and the source complex the use of interactive *a priori* prejudice might allow a user to produce an image biased to include the features he or she desires to see.

Experience with interactive non-linear processing of data from the MERLIN array (Davies, Anderson and Morison 1980) has shown that there can be significant problems with image fidelity (Junor 1987, Stephens 1987). It is

however possible with care to make good images far exceeding the quality expected when the array was first planned. These images provide good science but require for their production some experience and a large amount of processing time. Even so, occasionally images are produced which have astrophysically significant errors. The small but finite possibility of producing such images is an inhibition to the astrophysical interpretation of MERLIN images.

With the improvements planned for MERLIN (MERLIN Phase 2 proposal 1987) much of the expected science will lie in detecting detailed radio structures or in comparing MERLIN images with results from other instruments such as the VLA or the HST. It is likely that in the future ever higher fidelity images will be required as the astrophysical questions asked become more detailed and quantitative. For sparse arrays such as MERLIN and the VLBI arrays there must come a limit at which inadequate uv coverage cannot be reliably compensated for by increased use of deconvolution algorithms.

1.3 The MFS Idea

It appears that to make images with increasing reliability it will be necessary to improve substantially the uv coverage of our arrays. The obvious way in which to do this is to build more telescopes, however antennae with useful collecting area are very expensive to build. Although building new antennae will have some impact on the uv coverage of MERLIN and the VLBI arrays a really dramatic improvement in uv coverage is precluded by economic factors.

To improve the uv coverage of an array of fixed antennae without building new antennae might at first seem impossible. If, however, we note that the dimensions of the uv plane are measured in units of wavelengths then it can be seen that by observing at several frequencies it should be possible to build up a much improved uv coverage. This idea, to which the name Multi-Frequency Synthesis (MFS) is given, is by no means new (see section(1.4)), this thesis however seeks to be the first detailed study of the possibilities of MFS.

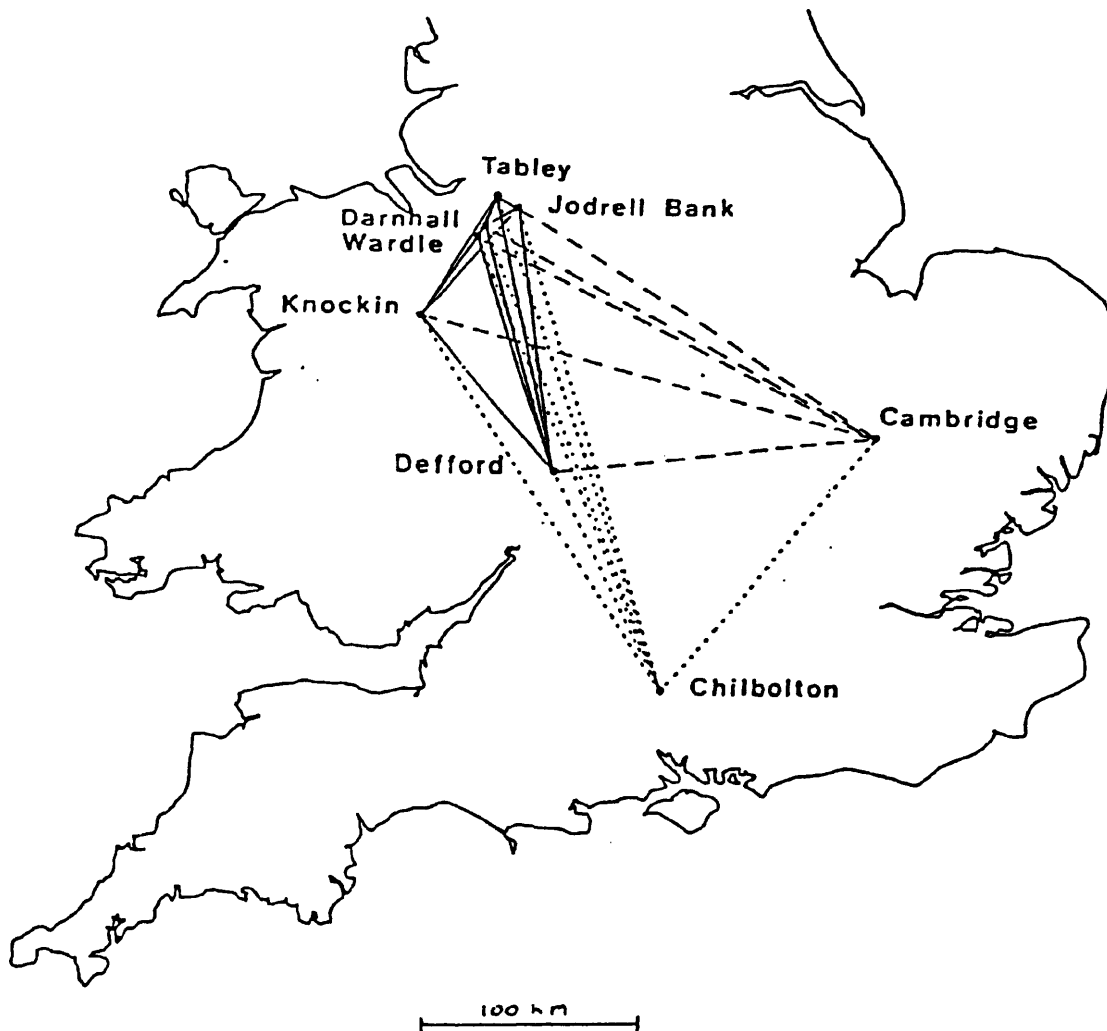


Figure 1.1: Locations of MERLIN antennae and baselines. Cambridge has recently been added to the array whilst the Chilbolton site has been proposed for inclusion into an extended MERLIN. The baselines to these two antennae are shown by broken lines.




Figure 1.2: MERLIN 8-station UV coverages at declination 50° .

a) Conventional single frequency observations at 1660 MHz.

b) MFS observations, 5 frequencies spread over 25% from 1350 MHz to 1750 MHz.


In fig(1.2a) is shown the conventional uv coverage expected from the 8-station extended MERLIN (see fig(1.1)) for a source at a declination of 50° , whilst in fig(1.2b) is shown the corresponding uv coverage if observations are made at 5 separate frequencies spread over a frequency range of 25%. There is clearly a very dramatic improvement, the MFS uv coverage is in fact equivalent to that which would be obtained from conventional, single-frequency observations, with an array of 17 telescopes. The sort of improvement in image quality that can be produced by using this enhanced uv coverage can be seen by comparing fig(1.3a) and fig(1.3b) where the first image shows the reconstruction of a model source by CLEAN given the conventional single frequency uv coverage and the second image shows the reconstruction assuming the MFS uv coverage. The model image used is shown in grayscale in fig(1.4a) and in a contour representation in fig(1.5a) (see sect(1.6.3)). These simulations simply demonstrate the improvements possible due to enhancing the uv coverage, they ignore the one obvious problem of MFS; that is that the source does not necessarily look the same at all the frequencies observed. Much of this thesis will be concerned with the effects of such frequency dependent structure (see Chapters 2 and 3) and with possible ways of correcting for these effects (see Chapters 4 and 5). It will be found that in most radio sources of interest the structure will vary with frequency in a smooth manner and that the frequency dependence at any point on such sources is generally quite simple, these facts will ensure that the ‘spectral errors’ are quite modest and that algorithms can be devised to remove them.

1.4 History of the MFS Idea

1.4.1 Early MFS Ideas

The realisation that the the dimensions of the uv plane are measured in wavelengths and therefore that the effective uv coverage of an array can be altered by switching frequency no doubt goes back to the earliest days of aperture synthesis. The first explicit mention in the literature however appears to be

a)



b)





Figure 1.3: CLEAN Reconstructions of a model image with data sampled by

a) The single frequency uv coverage of Figure 1.2a

b) The MFS uv coverage of Figure 1.2b.

Both images shown in linear greyscale plots from -1 mJy.beam^{-1} to 1 mJy.beam^{-1} .

a)



b)



Figure 1.4: Model source used in imaging simulations throughout this thesis.

- a) The model intensity image at the reference frequency of 1540 MHz. Linear plot from -1 mJy.beam^{-1} to 1 mJy.beam^{-1} .
- b) Model spectral index distribution $\alpha(x,y)$. Linear plot from $\alpha = 0.5$ (lightest) to $\alpha = 2.0$ (darkest).

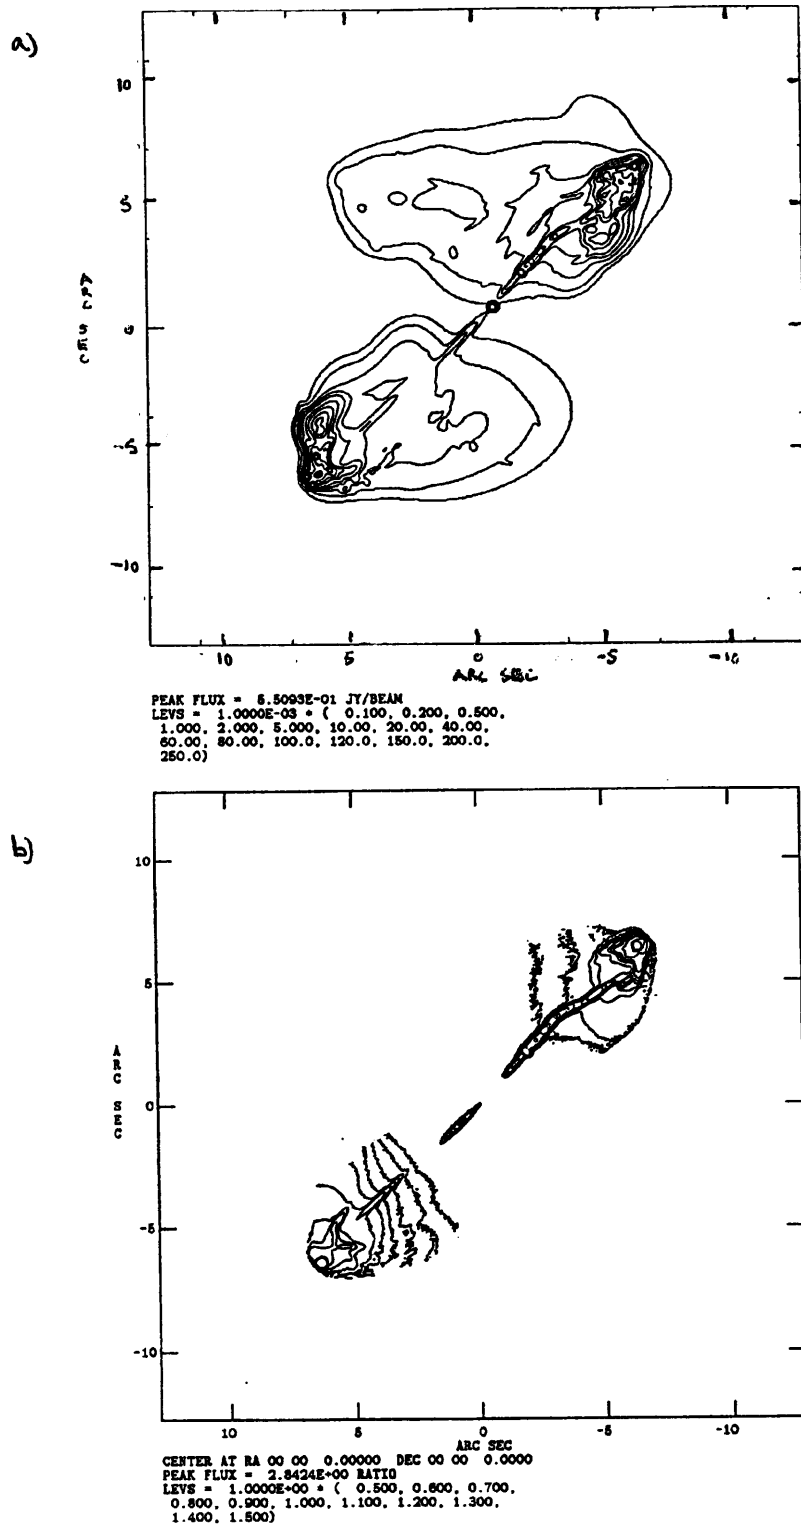


Figure 1.5: The model source. a) The intensity distribution at the reference frequency (1530 MHz). b) The spectral index distribution. Beam has a FWHM of 0.2".

due to Kock and Stone (1958). The MFS idea is not explicitly stated in modern terms of improving the uv coverage, but instead as a method to increase the resolution and decrease the sidelobes of a ‘pencil beam’ response of an array, i.e. the instantaneous response of an array in which the signals from the antennae are combined physically together. These ‘pencil beams’ in modern terms are equal to the Fourier transform of the instantaneous uv coverage of the array. The authors noted that either a single very large (i.e. octave) continuous bandpass or a series of discrete frequencies could be used to reduce the sidelobe response of the pencil beams.

The two above situations were shown to be equivalent respectively to monochromatic observations with a physically extended antenna and monochromatic observations with a series of physical baselines. This idea of ‘space-frequency equivalence’ follows directly from the realisation that the dimensions of the uv plane are measured in wavelengths so that the same aperture coverage can be achieved from different combinations of physical antenna and frequency coverage. A mathematical demonstration of the above ‘space-frequency equivalence’ appears in a later paper by Swenson and Mathur(1969) whilst McPhie(1979) discusses some of the limitations of the principle.

Kock and Stone(1958) mentioned two possible implementations of ‘space-frequency’ equivalence, i.e. either the use of one very wide bandpass or the use of several frequencies. Whilst the latter method is the main subject of this thesis the former method is also of interest and has been used. The Texas interferometer (Douglas et al 1973) used a single wide bandpass of 8% to improve the beam response in the direction parallel to the baselines whilst using extended cylindrical elements to control the beam shape in the perpendicular direction. This interferometer was successfully used as a survey instrument to find positions of unresolved sources.

Although very useful in the above mode such wide bandwidth systems cannot be used directly for the mapping of any but the simplest sources because of the problem of ‘bandwidth smearing’ (Cotton 1985), which is akin to ‘chromatic aberration’ in optics. The wide bandwidth system effectively samples

the visibility over a range of uv radii yet only returns a single estimate of the visibility, this visibility will be some mean of the visibilities over the range of uv radii sampled by the bandpass. The above process inevitably degrades the data and causes a radial smearing in any images produced which increases with distance from the centre of the image.

One way of separating the visibility contributions at each frequency within the bandpass is to correlate the signals with different time delays and hence different phase gradients across the bandpass, on Fourier transforming the resulting function of visibility against delay an estimate of the visibility against frequency can be found (D'Addario 1985). The visibility corresponding to particular parts of the uv plane can then be assigned. The above system gives results equivalent to having many separate sub-channels within the wide bandpass and therefore mimics a multi-frequency system. Such a wide bandwidth system was proposed for the planned VLA (and VLBI arrays) by MacPhie(1967). The author discusses in great detail how such a wide bandwidth system could produce great savings in the numbers of antenna required to achieve a full filling of the uv plane. Although the author explains how the problem of bandwidth smearing is avoided by using multi-delay correlation there is no discussion of the separate problem that occurs because the source has frequency dependent structure.

The alternative to observing with one wide bandpass is to observe at several different frequencies as described in section(1.3). The improved filling of the uv plane that results should allow higher fidelity imaging of complex sources. Multi-frequency observations can also be used however in a slightly different way to find accurate positions in an analogous manner to the wide bandwidth Texas interferometer.

A simple interferometer operating with a wide bandpass has a cosine sky plane response modulated by an envelope whose width depends on the bandwidth of the system. The width of this modulating beam or 'delay response beam' is inversely proportional to the bandwidth. The decrease in amplitude away from the centre of the delay response beam occurs because signals from

such positions are correlated with some residual delay. The phase corresponding to this delay is proportional to frequency, so that if there are a range of frequencies observed the net effect is to average the visibility over the bandpass with some phase rate. Away from the centre of the delay beam where there are large residual delays, the phase gradient with frequency will be large. Averaging over the frequencies within the bandpass will therefore give a significant decorrelation and so the amplitude of the response will decrease away from the centre of the delay beam.

One effective way in which to detect accurate positions of point sources is to measure the above phase gradient directly by making multi-frequency observations. From the phase gradient the residual delay can be then be found, knowing the delay as a function of hour angle will then allow an accurate source position to be found.

The above method is used in astrometry and geodesy and is termed 'Bandwidth Synthesis' (Thompson, Moran and Swenson 1986, Rogers 1970) and has proved very successful at finding absolute source positions to very great accuracy (see e.g. Fanselow et al 1984). The name originates from the original aim of creating from discrete frequencies a position finding instrument that is equivalent in performance to one wide bandwidth interferometer.

Whilst the above method of 'bandwidth synthesis'(BS) is similar to MFS in that it uses several frequencies the term as generally used refers to a very specific technique used with astrometric or geodetic observations.

Bandwidth Synthesis differs from MFS in that the target source is assumed to be almost point-like and the aim is to find the absolute position of this source. In contrast in MFS imaging the aim is to fill the uv plane and so constrain the structure of complex sources. The major difference between the techniques is that in BS the only quantity of interest is the gradient of phase with frequency rather than the total amplitude and phase as in MFS.

Although the term 'Bandwidth Synthesis' has been used loosely to describe the MFS method (Cornwell 1984) in this thesis the term will be used only

to refer to the specific method of delay astrometry for which the term was originally coined.

1.4.2 Modern MFS


All of the above discussion has concerned relatives of the MFS idea as defined in section(1.3), relatives that use the general idea of ‘Space-Frequency’ equivalence in different ways. What is the history of the MFS technique as originally defined?

Although the MFS idea was no doubt known to the earliest practitioners of aperture synthesis the earliest description of MFS aperture synthesis observations, indeed the first explicit mention of the idea in modern terms appears to be due to Conway and Stannard (1975). These authors combined together a number of single baseline observations of 3C273 made at wavelengths of 18,21, 31,49 and 74cm to form an image of the source, the resulting total uv coverage is shown in fig(1.6a).


The success of this imaging is remarkable especially given the fact that in addition to the problems caused because the data were observed at several frequencies the phase errors introduced by the atmosphere/ionosphere into the single baseline observations made the recorded phase unusable. Phases were assigned to the data by using the unresolved core component as a phase reference, by noting the maxima and minima of amplitude as the responses from the core and resolved components beat against each other the astronomical phase could be estimated.

Having corrected the phases, the effects of the core component were subtracted from the data at each frequency. The amplitudes at each frequency were then normalised according to estimates made of the spectrum of the resolved component taken from lunar occultation observations. Finally the above data was Fourier inverted and deconvolved using the CLEAN algorithm, the resulting image is presented in fig(1.6).

Remarkably the main features of this image were shown to be correct when later high fidelity images were made by conventional single frequency aperture



a) UV coverage built up from multi-frequency observations. Observations A and B from Jodrell-Defford and Wardle-Defford baselines at 74cm. Observations C from Jodrell to a mobile antenna at 74cm. Observations D to H made with the Jodrell-Wardle baseline at 18, 21, 31, 49 and 74cm respectively.



b) The image resulting from the MFS observations (left) compared to a modern VLA image of the same area at 1665 MHz (right) (after Perley). Both images have a resolution of 1''. The lines in the VLA image indicate degree and direction of linear polarisation.

Figure 1.6: The pioneering MFS observations of Conway and Stannard (1975) of the jet of 3C 273.

synthesis (Davis, Muxlow and Conway 1985). The success of these MFS observations of 3C273 over an octave in frequency make modern plans for MFS look modest. These observations showed that improved images were available from MFS despite all the problems of calibration and frequency dependent structure, albeit only at quite small dynamic ranges (approximately 20:1).

Despite the success of the above observations MFS methods were not pursued or discussed in the literature again till Cole(1979), apparently unaware of the work of Conway and Stannard, explicitly described a system of MFS as a way of improving uv coverages. This author also pointed out the fact that frequency dependent structure will cause artifacts in MFS images, though no attempt was made to estimate the size or nature of these artifacts.

The MFS idea was also suggested by Andrejanov et al (1984) to be used in connection with VLBI from orbiting antennae. If such an antenna is placed in an orbit of large dimensions (several Earth radii) so as to get high resolution there will be large gaps in uv coverage. The above authors suggested that these gaps could be partially filled by using many frequencies spread over a wide range of frequencies, an octave range of frequency being suggested.

Cornwell (1984) has made the most significant contribution to date in assessing the possibilities of MFS. The author considered the case for MFS, in particular he suggested it for the proposed Very Long Baseline Array (VLBA) (Romney 1987). A mathematical analysis of the form of the artifacts that are introduced by frequency dependent structure is presented although no quantitative estimates of the size of these errors is made. A couple of suggestions were made for algorithms which could remove the effects of frequency dependent structure. These algorithms will be considered in chapter 4.

Independently of the above work Conway (1985) analysed the effects of spectral effects in a very similar manner to that of Cornwell. Also presented were quantitative estimates of the size of spectral effects and the results of crude MFS MERLIN observations of the source 3C249.1. The analysis of spectral effects which will be presented in chapters 2 and 3 is based on extending the analysis contained in this initial work.

Manchester (1985) suggested the use of MFS on the proposed compact element of the Australia Array (Norris 1987). Simulations were carried out which showed that MFS can significantly improve the properties of the array's dirty beam. No analysis of the effects of frequency-dependent structure were carried out however. In section (8.4.1) the prospects for MFS on the AT will be discussed further.

Several authors including Braun, Gull and Perley(1987) have successfully carried out MFS observations with the the VLA in recent years. The VLA (Very Large Array) (Napier 1983 et al) is a very versatile instrument with the capability to observe at a very wide range of frequencies in any particular observational band. With such a capability MFS observations are easily carried out. Even in normal operation the array observes simultaneously in two 50MHz bandpasses whose centres are separated by 50MHz, at 1420 MHz this range of frequency corresponds to a two frequency MFS system with a fractional bandspread of 3%. Images made even at the highest dynamic range do not show any obvious ill effects due to frequency dependent structure, a situation which is encouraging for the prospects for wide bandspread MFS.

The total 100MHz bandwidth is split into two separate channels in order to reduce the effects of bandwidth smearing rather than deliberately to improve the uv coverage. Indeed the improvement in uv coverage with only a 3% bandspread is slight. The versatile VLA observing system does however allow the observing frequencies to be placed much further apart in order to improve significantly the uv coverage. One example of this mode of observing is from Murphy(1988) where observations were made simultaneously at the frequencies of 1465 MHz and 1640 MHz giving a fractional bandspread of 11%. The sources observed were a sample of core-dominated sources and special procedures were required to correct for the effects of their strong cores which had very different spectral behaviours from their extended emission.

Another more spectacular example of MFS on the VLA is that of Braun, Gull and Perley (1987) in which observations at 4 frequencies spread over up to 15% in frequency were combined when making observations of the supernova

remnant Cassiopeia A at both L and C bands. This particular source is so complex that even long tracks with the VLA do not produce sufficiently dense uv coverage to fully ensure the fidelity of reconstruction. The improvements in uv coverage possible using MFS are therefore highly desirable. Supernova remnants are also ideal targets for MFS imaging because there is evidence that their spectral behaviour is particularly simple (Siebler et al 1981) and so spectral effects are easily removed from the MFS data (see sect(2.3.2)).

Although the VLA design is such that MFS observations are relatively easy to make, in many ways the improvement in image quality that can be achieved by MFS on the VLA is limited because the conventional uv coverage is so good. Only on the most complex sources such as SNR's is it likely to give very significant improvements. Whilst it is true that on short 'snapshot' observations the uv coverage is much poorer than that achieved by long tracks, MFS does not dramatically improve the uv coverage of such observations. The effect of MFS on uv coverage is to spread out the visibility observations in uv radius, after applying MFS, gaps in azimuth will remain and these will limit the improvement in uv coverage possible. It appears that generally MFS is likely to make less impact on the VLA than on sparse arrays.

1.5 Implementing MFS

One way of implementing MFS on an array is to observe at all the frequencies simultaneously. Such an implementation has the disadvantage that separate receivers are required at each frequency. On the other hand such an implementation can have the advantage of increasing the total bandwidth with a consequent decrease in the noise. For MERLIN, however, which uses a microwave link to transfer the data from the antenna to the correlator, the total instantaneous bandwidth that can be used is limited by the bandwidth of the link. Similarly in VLBI the recording systems have limited capacities which restricts the instantaneous bandwidth that can be used. In such systems the total available bandwidth must be apportioned between the different

observing frequencies and so there will be no improvement in sensitivity from simultaneous observing.

A second implementation consists of making separate long observations at each frequency, spending a whole day at one frequency and then observing at another frequency. This method has the disadvantage that if an N frequency MFS observation is required N days of observations are needed. Although the large amount of observing time will produce an image with lower thermal noise than that from a single-day single-frequency observation there are only a limited number of sources which can justify the large amounts of telescope time needed for this implementation. Technically the above implementation is the easiest to arrange and so this mode of observing will no doubt be used on MERLIN at least in the first instance.

The last way of implementing MFS is to switch frequency rapidly. By switching on timescales of the order of minutes the uv measurements can be widely distributed over the uv plane yet a full MFS observation can be carried out within a single day. The frequency switching method also has another possible advantage in addition to operational flexibility, the phase errors introduced at the different frequencies by propagation effects in the ionosphere or troposphere will not be independent and this fact can be used as an extra constraint in the solution for these phase errors. This possibility is discussed in section(6.4.4). For MERLIN, of the three possibilities, the frequency switching implementation probably provides the best compromise between technical practicality and operational flexibility.

1.6 MFS on MERLIN

1.6.1 Introduction

Although the MFS technique is of general applicability on all arrays this thesis will concentrate on MFS as applied to MERLIN. MERLIN has a geometry typical of many sparse arrays of antennae (such as are found in VLBI arrays). Hopefully MFS can be successfully tested on MERLIN before being applied

more generally to these other arrays. With the continuing enhancement of MERLIN the capability exists to build in a hardware capability to make MFS observations, so creating a great improvement in the performance of the array in its own right and demonstrating the feasibility of MFS generally.

1.6.2 MERLIN UV Coverage and MFS

How many frequencies and what fractional range of frequency ('Bandspread') is required for a useful MFS system on MERLIN, simply from the point of view of uv coverage?

The maximum number of frequencies will probably be limited by the practical problems of finding frequencies sufficiently free of radio interference. This subject is discussed in detail in Chapter 6 where it is concluded that over a 30% range in frequency at L-band (1350-1750 MHz) seven frequencies is the practical maximum. Using less than the above number of frequencies will simply increase the size of residual holes in the uv plane proportionately. For the most part MFS using five frequencies will be assumed throughout this thesis, five frequencies simply being a useful compromise between improved uv coverage and practicality. Where appropriate the effects of varying the number of frequencies will be noted.

Assuming five frequencies we next consider the uv coverage of MFS MERLIN as a function of fractional bandspread. In fig(1.7) are shown uv coverages for the 8 antennae of the extended MERLIN at a declination of 50° for different fractional bandspreads. It can be seen that in order to fill the uv plane completely a fractional bandspread of at least 25% is required. Reasonable filling is however achieved with a 15% bandspread. Note that despite these large bandspreads MFS will always leave a hole in the very centre of the uv plane. For sources which have a large amount of extended flux this poor central sampling will have a much more serious effect on fidelity than any other gaps. For the highest fidelity images it seems likely that it will be beneficial to combine VLA 'A'-Array data with MERLIN MFS data.

Assuming a 25% fractional bandspread and 5 frequencies the MERLIN only




Figure 1.7: MFS uv coverages of 8-station MERLIN at a declination of 50° for various fractional bandwidths and 5 frequencies.

uv coverage is shown at various declinations in fig(1.8). From this plot it can be seen that full filling of the uv plane is only achieved at high declinations. Even so there is quite significant enhancement in uv coverage over that achieved by single frequency observations (fig(1.9)) even at declinations close to the celestial equator.

1.6.3 Imaging Specifications for MERLIN MFS

It will be demonstrated in the coming chapters that the effects of frequency dependent structure will cause artifacts in MFS images. Later on methods will be presented which attempt to remove these artifacts. In order to decide whether success has been achieved in making MFS applicable to the full range of sources accessible to MERLIN some target problem must be defined which provides some 'worst case' against which to test MFS.

The source model used for most of the simulations in this thesis is presented in fig(1.4) and fig(1.5). These figures give both the intensity distribution and the model spectral index distribution. The structure of this model is based on VLA observations of extended FRII sources such as Cygnus A (Perley, Dreher and Cowan 1984). This model should provide an approximation to the above 'worst case' source. The image is assumed to be at the most probable frequency band for MFS, i.e. L-band. Each point on the source is assumed to have a power law spectra (see sect(2.3.3) for details) with an exponent (the 'spectral index' $\alpha(x,y)$) which varies with position on the source. This spectral index varies from approximately 0.5 at the bright 'hotspots' to greater than 2.0 in the extended lobes (see fig(1.4b)). The spectral index on the very brightest parts of the source at the three hotspots varies over a range of 0.2. The hotspots have fluxes of around 500 mJy.beam⁻¹ and the total source flux is 10.5Jy.

With the FWHM resolution of extended MERLIN being 0.15" at L-band the total source extends over a total area of greater than 6×10^4 beam areas. This is an area much larger than is likely to be reliably imaged by conventional methods with an 8 antenna array according to the criteria of Wilkinson(1983). It can therefore potentially provide a useful demonstration of the advantages




Figure 1.8: MFS uv coverages at various declinations using the 8-station MERLIN with a 25% bandwidth and 5 frequencies.




Figure 1.9: Single frequency uv coverages for the 8-station MERLIN at various declinations.

of MFS for image reconstruction. The model also deliberately contains a wide range of structures with many different spatial frequencies to test out MFS reconstruction methods fully.

With the increased sensitivity expected from improved MERLIN a typical observation at L-band lasting 24 hours will produce images with rms noises of order $25\mu\text{Jy.beam}^{-1}$ (R.J. Davis, private communication). This means that, excluding the core component (whose effects can be removed at each frequency) the maximum ‘dynamic range’ required between the hotspots and the thermal noise will be less than $\approx 20,000:1$. If it can be shown that spectral effects due to the hotspots and any other parts of the source can be corrected to better than this level then MFS will be applicable to virtually all of the complex targets accessible to extended MERLIN.

1.7 Astrophysical Impact of MFS

As described in sect(1.2.2) reliance on non-linear processing algorithms for the processing of MERLIN or VLBI data has meant that the resulting images although astrophysically useful are not always perfectly reliable. MFS can be usefully applied to the extended MERLIN and to VLBI arrays to enhance uv coverage and hence reliability. This should allow imaging of the gross features of sources with greater confidence, it should also allow these instruments to partake in new, more quantitative science in which the highest fidelity is required.

Examples of such science includes optical/radio comparisons at HST resolutions, possibly tracing interactions between radio jets and line emitting clouds (e.g. Heckman et al 1982). Joint radio spectral studies with MERLIN and the VLA are another area of interest, these should allow us to determine the spectral age of synchrotron emitting electrons (Myers and Spangler 1985) with higher accuracy than at present. The accurate determination of proper motions or time evolution of radio knots or hotspots is an area in which

the highest fidelity will be required and in which MFS MERLIN in combination with the EVN can make an impact. An understanding of the dynamics of radio jets will probably require highly accurate images to compare against the results of numerical simulations (Norman, Smarr and Winkler 1982, Wilson and Falle 1985, Lind 1986) and here too MFS may be vital.

In many areas improved fidelity through MFS will make contributions to astrophysics, for more details see section(8.3.2) and section(8.4).

1.8 Outline of Thesis

In the next chapter the nature of the effects that are introduced into MFS images by frequency dependent structure are described. In Chapter 3 we consider the level of these spectral errors in realistic cases of extended sources. In Chapter 4 some simple algorithms for removing these effects are investigated and shown not to be generally practical in the case of MFS with sparse arrays such as MERLIN. In Chapter 5, however, a more successful algorithm ‘Double Deconvolution’ will be described. Simulations with this algorithm suggest that it should allow MFS to be practical on the whole range of sources accessible to future MERLIN. In Chapter 6 some of the practical problems of MFS are described including the nature of the interference environment, hardware implementations of MFS, amplitude and phase self-calibration of MFS data and the effects of complex spectra including radio cores. Chapter 7 describes the results of real narrow bandspread MERLIN MFS observations of the radio source 3C179. Finally in Chapter 8 conclusions are drawn as to the prospects for MFS, particularly as applied to MERLIN, suggestions for future work made and the astronomical impact of MFS assessed.

Chapter 2

MFS and Frequency Dependent Structure I : The Form of the Spectral Effects

2.1 Introduction

We have seen in the previous chapter that the MFS technique can dramatically improve the uv coverage of an array and so potentially greatly improve its imaging capability. There is however an obvious problem with this technique, this is that the source brightness distribution will not be the same at all the observing frequencies. This fact will introduce unpleasant effects into our images, the nature, size and removal of these effects are the main subjects of this thesis. In this chapter and the next the form and magnitude of these ‘spectral effects’ will be outlined. Later on in Chapters 4 and 5 algorithms will be developed which attempt to remove these effects and so allow an estimate of the source structure at some mean frequency to be produced.

2.2 The Nature of Frequency Dependent Structure

2.2.1 Frequency Dependent Structure and MFS

If the source brightness distribution is completely different at each observing frequency then it is obvious that the MFS technique is not useful. For instance, if MFS observations are made of a region dominated by radiation from narrow spectral lines then using information from one frequency cannot help in imaging the sky at another frequency. To be useful MFS requires that the

source brightness distributions at each observing frequency are correlated. In this section the nature of frequency dependent structure in continuum images from synchrotron sources is considered; it is found that emission is indeed highly correlated and has a simple frequency dependence.

For successful MFS the changes between frequencies must either be so small that they cause no sky plane effects or, if they are not insignificant, the brightness at each pixel on the source must be a simple function of frequency. In this second case only a small number of parameters will be needed to characterise the spectrum at each pixel. Processing methods can be envisaged which can estimate all the significant pixel parameters from the MFS data. One of these parameters will be the flux density at some reference frequency, the estimation of which is the prime objective of MFS. If two parameters at each pixel are required to fit the spectrum then the number of unknowns that any imaging algorithm needs to solve for are doubled. If there are more than two observing frequencies then the ratio of observables to unknowns is still increased over conventional imaging, thus making it plausible that MFS reconstructions will be more reliable than conventional single-frequency images.

2.2.2 Synchrotron Emission Spectra

We are fortunate that the mechanisms that produce continuum radio emission usually generate simple spectra over the ranges of frequency that we are considering for MFS, they can therefore be described to sufficient accuracy by only a few parameters. We will concentrate in this thesis on the synchrotron emission mechanism which dominates the continuum emission from complex sources, such as radio galaxies, quasars and supernova remnants. Although there are other emission mechanisms of importance in radio astronomy such as the free-free emission mechanism (Longair 1981) these are not important in the above objects which are the main targets for MERLIN and VLBI imaging. We shall see that under most conditions the spectra of radiation produced by the synchrotron mechanism is conveniently described by only two parameters

at each point on the source so that MFS should be easily adapted to deal with significant spectral effects.

The origin of synchrotron emission lies in the acceleration of ultra-relativistic electrons in magnetic fields (Ginzburg and Syrovatskii 1969). Synchrotron emission was first observed in cyclotron particle-acceleration devices; the first suggestions that radio sources might be emitting radiation by this mechanism were made by Alfvén and Herlofson (1950) among others. At this time technical difficulties prevented immediate observational confirmation in the radio that synchrotron emission was the dominant emission mechanism of the compact radio sources. These early suggestions were largely forgotten till Shklovskii (1953) made the proposal that both the optical and radio emission from the Crab nebula were due to synchrotron radiation. The tell-tale presence of high polarisation was confirmed first in the optical and later in the radio (Mayer et al. 1957) showing that synchrotron was indeed the emission mechanism for this object. Similar polarisation observations of discrete sources and the galactic background confirmed that these also were synchrotron sources.

In the late 1950's and early 1960's multifrequency observations of synchrotron sources were made in order to determine the spectra of these objects (e.g. Whitfield 1957, Conway et al 1963). For a large fraction of sources it was found that the spectrum was well fitted by a power law, with some spectral index α such that

$$I(\nu) = I(\nu_o)(\nu/\nu_o)^{-\alpha} \quad (2.1)$$

where $I(\nu)$ and $I(\nu_o)$ are the intensities at the general frequency ν and at the 'reference frequency' ν_o respectively. In many objects the above power spectrum was found to extend over a wide range of frequencies. In the Crab Nebula for instance a power law spectrum extends over a range of greater than 10^4 in frequency at radio wavelengths (Baldwin 1971).

Some of the total integrated flux spectra of sources which were not power law could be explained as the superposition of power law spectra presumably from different components of the radio source. High resolution aperture synthesis maps of extragalactic sources confirmed this explanation for most

extended sources (e.g. MacDonald et al 1968). In these sources the intensity at sky coordinates x, y and frequency ν can be written as

$$I(x, y, \nu) = I(x, y, \nu_o)(\nu/\nu_o)^{-\alpha(x, y)} \quad (2.2)$$

At each position on the source the spectrum is a power law to good approximation but the spectral index of this power law varies with position. The empirical discovery that the spectral behaviour of synchrotron emission at any position on a source x, y can be described over a wide range of frequency by only two parameters $I(x, y, \nu_o)$ and $\alpha(x, y)$ is important for MFS. As described earlier it should then be possible to form an algorithm to estimate all these parameters from the MFS data and still create final images better than those from single frequency observations.

The work of Ginzburg and Syrovatskii (1965) and of Pacholczyk (1970) showed how a power law emission spectrum results from an energy distribution of ultra-relativistic electrons which is also power law, i.e. a distribution such that the number dN of electrons with energies between E and $E+dE$ follows

$$dN = N_o(E/E_o)^{-\gamma(x, y)}dE \quad (2.3)$$

where $\gamma(x, y)$ is the electron energy distribution index at some position on the source and N_o is the electron energy density at energy E_o . These authors also showed that the spectral index α and the electron energy distribution index γ are related by

$$\alpha = (\gamma - 1)/2 \quad (2.4)$$

The above relationship between the underlying electron energy spectrum and the emission spectrum was derived by summing contributions from the spectra of single electrons (see fig 2.1) over the electron energy population distribution. Note that the width of this single electron spectrum is wide compared to the frequency of peak emission, this means that even if for some reason the underlying electron energy distribution were not smooth but instead consisted only of say monoenergetic electrons the emission spectra can be expected to be smooth over at least an octave in frequency. Even in the worst case therefore

there will be some degree of correlation between the emission at frequencies spread over 30% or less.

2.2.3 The Origin of Power Law Spectra

The fact that the observed emission spectra are usually very accurately power law means that the underlying electron spectra must be power law. Given the attractiveness of such power law emission spectra for the practical application of MFS it is important to ask about the theoretical explanation for the underlying power law electron energy spectra. The basis must be sought in the mechanism by which synchrotron emitting electrons are accelerated to ultra-relativistic velocities. A clue to this mechanism comes from the probable site of electron acceleration. The dominant structure of quasars and radio galaxies is of a compact 'core' straddled by diffuse 'lobes' of emission. In the more powerful of these sources the lobes often contain compact and very bright regions of emission known as 'hotspots'. The model of Blandford and Rees (1974) explains these regions as places where jets of material interact with the intergalactic or intercluster media and in many sources these jets have now been observed (e.g. Bridle and Perley 1984). These jets of material will often be travelling with a supersonic velocity in the intergalactic medium, the interaction of jet and external medium will therefore consist of a shock, or complex of shocks and it is thought that these shocks are the primary sites of electron acceleration for the electrons that emit in the radio lobes. Evidence for hotspots as sites for electron acceleration comes from observations of the the rapid decrease in flux and the increasing spectral index (due to 'spectral ageing' and adiabatic expansion (Kardashev 1962)) as we move away from hotspots into radio lobes (Myers and Spangler 1985). The low frequency power law spectra of the lobe emission can be explained if the original electron population created at the hotspots is power law.

How are electrons accelerated to the very high energies required to explain the observed synchrotron emission and can these mechanisms explain the power law energy spectra that have been inferred for these electrons?




Figure 2.1: The spectral emissivity of an ultrarelativistic electron due to synchrotron radiation. The maximum of the emission spectrum is at frequency $\nu = 0.29\nu_c$ where ν_c , the critical frequency, is $\frac{3}{2}\gamma^2\nu_g \sin \theta$.

$$\nu_g = \frac{eB}{2\pi m_e}$$

$\theta \approx$ pitch angle of electron

The most favoured mechanism for electron acceleration is known as diffusive shock acceleration (DSA) or as the 'First Order Fermi Process' and was first proposed in the late 1970's by several authors including Bell (1978). The mechanism consists of high energy electrons gaining small amounts of energy in passing through a strong shock and being scattered by irregularities in the magnetic field on the other side. The electrons gain energy because electrons with isotropic velocities in the frame of the upstream or downstream flows which then cross the shock will meet irregularities with systematic velocities toward the shock. Electrons reflected off these irregularities will gain energy. A small fraction of electrons scattered on the downstream side will be advected away and leave the region of the shock but the remainder at some point will diffuse back across the shock front where they can start the cycle over again, gaining more energy every time they cross the shock front. Electrons may undertake a very large number of such cycles before escaping and can thereby reach very large energies.

A useful analogy for this mechanism is that of a rubber ball bouncing between two plates which approach each other at some velocity, one of the plates having a small fraction of its surface covered with holes, in this analogy the mechanism by which the ball can gain energy can readily be visualised. The ball represents an electron whilst the two plates represent the scattering centres on either side of the shock, the holes represent the constant probability that an electron can escape downstream away from the accelerating region. If the shock is strong in a medium of monatomic gas then the approach velocity can be shown to be $V_s/4$ where V_s is the shock velocity in the upstream medium.

The DSA mechanism is thoroughly reviewed by Drury (1983) whilst an illuminating description of the way in which this mechanism produces a power law spectra is given by Longair (1981). Power law electron energy spectra are shown to occur provided the fractional increase in energy and probability of escape per cycle are both independent of energy (Webb et al 1985). Both of the above requirements should arise easily in radio sources in which the

electrons are relativistic and have velocities almost independent of energy. It therefore appears that the ubiquity of power law spectra can be understood theoretically.

Note that the simple model of the relativistic electrons as test particles only affected by the conditions in the thermal gas does not strictly hold in most realistic cases. The pressure due to relativistic electrons will usually be comparable to the pressure due to thermal gas and in this case the theory must be elaborated (Drury 1983, Coleman and Bicknell 1988). Although analysis suggests that in the special situation of the pressure being *dominated* by a relativistic electron population a power law will still result, no theory has demonstrated that such a spectrum occurs in the general case. It has been conjectured however that the spectrum would still follow a power law (Coleman and Bicknell 1988).

An important result which follows naturally from the DSA acceleration process is that it predicts (Bell 1978) a spectral index which depends on Mach number. For strong non-relativistic shocks (i.e. for large Mach number shocks) in monatomic media the theory suggests that an electron energy index of $\gamma = 2.0$ should be produced. This in turn will give rise to an emission spectral index at the site of acceleration of $\alpha = 0.5$. At low frequencies hotspot spectral indices apparently have $\alpha = 0.50$ to 0.65 (see Stephens 1987), so that to first order the theory seems correct. Spectral indices larger than 0.5 for strong shocks can be explained by considering the effects of synchrotron losses after electrons have been accelerated at the shock (Heavens and Meisenheimer 1987).

Radio emission from the jets which supply radio lobes is thought to originate from electrons accelerated by shocks either internal to the jet or in a turbulent boundary layer around the jet. The same mechanisms of electron acceleration will operate in jets as in hotspots (Achterberg 1986) and power law spectra with a small range of spectral index should occur here too. Observations of jets give $\alpha = 0.60 \pm 0.15$ (Bridle 1982).

The convenient situation arises that the brightest parts of many sources

not only obey a power-law spectra but also have spectral indices which show only small variations. This will allow us to carry out preliminary processing of the MFS data in which the data is corrected as if the source has a constant index, so substantially reducing spectral effects (see sect(2.3.2)).

So far we have assumed that all parts of all sources have power law spectra. However on applying synchrotron theory (Pacholczyk 1970) to the calculation of theoretical spectra it is found that especially at high and low frequencies large departures from power law spectra are predicted due to such mechanisms as ‘self-absorption’ or ‘electron ageing’. Such departures are often observed in radio ‘cores’ which show complex spectra and in other parts of radio sources at high and low frequencies. These departures from power law spectra and their effects on MFS will be considered in Chapter 6 where it will be shown that, except for radio cores which are point-like and can be dealt with special ways, the assumption of power law spectra holds very well at the moderate frequencies (i.e. L-band) at which MFS is intended to operate.

In this chapter we proceed to develop the theory of MFS assuming that the effects of any radio cores with complex spectra have been removed from the data so that the residual data can be assumed to correspond to a source with a perfect spatially varying power law (SVPL) spectral distribution.

2.3 Effects of Power Law Spectra on MFS

2.3.1 Introduction

We have seen (sect(2.2)) that in many cases a SVPL synchrotron spectra described by equation(2.2) is both theoretically expected and observationally observed.

In MFS, observations are made of the complex visibility at several distinct frequencies ν_i . At each such frequency the observed visibility function $V(u,v,\nu_i)$ will be non-zero only on parts of the uv plane sampled by the various baselines at frequency ν_i . The observed visibility at u,v at frequency ν_i is

therefore

$$V(u, v, \nu_i) = S(u, v, \nu_i) \overline{I(x, y, \nu_i)} \quad (2.5)$$

where the overline symbol indicates a Fourier transform. $S(u, v, \nu_i)$ is the uv sampling function consisting of δ functions at observed points in the uv plane (whose areas depend on the ‘weights’ assigned to each observation). The above equation utilises the Zenicke-VanCittert theorem (e.g. Swenson and Mathur 1968) which states that the visibility at frequency ν_i is equal to the Fourier transform of $I(x, y, \nu_i)$, the brightness distribution of the source at frequency ν_i .

If the observations at the different frequencies cover the same hour angle range and use the same antennae then the frequency dependence of the uv coverage can be described by

$$S(u, v, \nu) = S_o(u \frac{\nu}{\nu_o}, v \frac{\nu}{\nu_o}, \nu_o) \quad (2.6)$$

where $S_o(u, v, \nu_o)$ is the effective uv coverage function of the array at the reference frequency ν_o . The Fourier transform of the uv coverage at frequency ν_i , $S(u, v, \nu_i)$ will correspond to the ‘dirty beam’, $D(x, y, \nu_i)$ at frequency ν_i , so that

$$D(x, y, \nu_i) = \overline{S(u, v, \nu_i)} \quad (2.7)$$

These dirty beams will be identical in form except for a change in radial scale with frequency.

2.3.2 Non-Spatially-Varying Spectral Index Distributions

Consider the situation in which the spectral index is constant across the source so that $\alpha(x, y) = \alpha_o$. In this case

$$I(x, y, \nu_i) = I_o(x, y) \left(\frac{\nu_i}{\nu_o} \right)^{-\alpha_o} \quad (2.8)$$

where

$$I_o(x, y) = I(x, y, \nu_o) \quad (2.9)$$

i.e.

$$\overline{I(x, y, \nu_i)} = \overline{I_o(x, y)} \left(\frac{\nu_i}{\nu_o} \right)^{-\alpha_o} \quad (2.10)$$

Thus

$$V(u, v, \nu_i) = \overline{I_o(x, y)} S(u, v, \nu_i) \left(\frac{\nu_i}{\nu_o} \right)^{-\alpha_o} \quad (2.11)$$

If the observed visibilities $V(u, v, \nu_i)$ at frequency ν_i are multiplied by the factor $(\nu_i/\nu_o)^{\alpha_o}$ to form $V_c(u, v, \nu_i)$ the result at a particular point in the uv plane is equal to the visibility that would have been observed if the observing frequency had been the reference frequency ν_o . Thus

$$V_c(u, v, \nu_i) = (\nu_i/\nu_o)^{\alpha_o} V(u, v, \nu_i) \quad (2.12)$$

$$= S(u, v, \nu_i) \overline{I_o(x, y)} \quad (2.13)$$

On combining the corrected uv data at each frequency we obtain the total uv data set V_T .

$$V_T(u, v) = \sum_{i=1}^N V_c(u, v, \nu_i) \quad (2.14)$$

The combined data set now consists of the Fourier transform of the intensity at ν_o sampled by the combined multifrequency uv coverage. On Fourier inverting to produce a ‘dirty map’ we obtain

$$\overline{V_T} = I_o(x, y) * D_o \quad (2.15)$$

where

$$D_o = \overline{\sum_{i=1}^N S(u, v, \nu_i)} \quad (2.16)$$

is the multifrequency dirty beam, i.e. the Fourier transform of the total uv coverage, and the symbol $*$ denotes convolution. This dirty map can then be deconvolved using standard algorithms such as CLEAN or MEM with D_o chosen as the ‘dirty beam’. The resulting images should have higher fidelity than conventional images because of the greatly improved effective uv coverage $\sum_{i=1}^N S(u, v, \nu_i)$.

The obvious effect of a non-spatially varying spectral index is simply to scale the flux at each observed frequency ν_i , by $(\nu_i/\nu_o)^{-\alpha_o}$. This multiplicative

effect can be easily removed from the observed uv data at each frequency before combining the different frequency data sets together. For a SVPL spectra however the effects of spectral index on each frequency channel are not so straightforward and more complicated procedures will be required to remove spectral effects.

2.3.3 SVPL Spectra and the Linear Expansion

In this section the effects that arise if there is a SVPL spectral distribution are considered. We first note that, although the spectral index will always vary across real sources, it is often found that the spectral index of the brightest parts of radio sources does not vary very much. The spectral indices of hotspots for instance are predicted to be around $\alpha = 0.5$ (Stephens 1987)) whilst jets have $\alpha = 0.60 \pm 0.15$ (Bridle 1982). In such situations it is clearly beneficial to first of all ‘correct’ the observed multi-frequency data at each frequency for an ‘overall’ spectral index α_o , as described in the previous section. The result of this processing should be to reduce gross spectral effects, but to leave residual effects due to the remaining spatially varying part of the spectral index. If we originally have the observed visibility at frequency ν_i , $V(u, v, \nu_i)$ such that

$$V(u, v, \nu_i) = \overline{(\nu_i/\nu_o)^{-\alpha(x,y)} I_o(x,y)} S(u, v, \nu_i) \quad (2.17)$$

and this sampled visibility is corrected for an ‘overall’ spectral index α_o as in the above section then we obtain

$$V_c(u, v, \nu_i) = \overline{(\nu_i/\nu_o)^{\alpha_o} (\nu_i/\nu_o)^{-\alpha(x,y)} I_o(x,y)} S(u, v, \nu_i) \quad (2.18)$$

$$= \overline{(\nu_i/\nu_o)^{(-\alpha(x,y)+\alpha_o)} I_o(x,y)} S(u, v, \nu_i) \quad (2.19)$$

The corrected uv data will therefore be equivalent to that due to a source with an effective spatially varying spectral index of $\alpha'(x, y) = (\alpha(x, y) - \alpha_o)$ instead of $\alpha(x, y)$.

In nearly all cases it should be advantageous to remove some overall spectral index α_o from the data. The optimum value of α_o will be close to the

spectral index of the brightest points of the source rather than the spectral index of the integrated source. We will assume in what follows that such an ‘overall’ correction has been applied so that the data are only affected by the residual effective spectral index α' . The corrected uv data from all the observing frequencies are as before combined to form a total uv data set V_T , i.e. we form

$$V_T(u, v) = \overline{\sum_{i=1}^N (\nu_i/\nu_o)^{-\alpha'(x,y)} I_o(x, y) S(u, v, \nu_i)} \quad (2.20)$$

To proceed we substitute $\nu_i/\nu_o = (1 - \Delta\nu_i/\nu_o)$ where $\Delta\nu_i = \nu_o - \nu_i$ and then expand the intensity at frequency ν_i as a Taylor series in $\Delta\nu_i$.

$$I_o(\nu_i/\nu_o)^{-\alpha'(x,y)} = I_o(1 - (\Delta\nu_i/\nu_o))^{-\alpha'(x,y)} \quad (2.21)$$

$$= I_o + (\Delta\nu_i/\nu_o) I_o \alpha'(x, y) + (\Delta\nu_i/\nu_o)^2 I_o \alpha'(\alpha' + 1)/2 \dots \dots \quad (2.22)$$

$$= I_o + \sum_{n=1}^{\infty} (\Delta\nu_i/\nu_o)^n (1/n!) I_o \prod_{m=1}^n (\alpha' - m + n) \quad (2.23)$$

substituting this expansion into our expression for V_T in equation (2.20) we obtain

$$V_T = \sum_{i=1}^N \overline{[1 + (\Delta\nu_i/\nu_o) \alpha'(x, y) + 1/2 (\Delta\nu_i/\nu_o)^2 \alpha'(\alpha' + 1) \dots] I_o(x, y) S(u, v, \nu_i)} \quad (2.24)$$

$$= \overline{I_o} \sum_{i=1}^N S(u, v, \nu_i) + \overline{I_o \alpha'} \sum_{i=1}^N (\Delta\nu_i/\nu_o) S(u, v, \nu_i) + \frac{1}{2} \overline{I_o \alpha'(\alpha' + 1)} \sum_{i=1}^N (\Delta\nu_i/\nu_o)^2 S(u, v, \nu_i) \dots \dots \quad (2.25)$$

On Fourier transforming the combined MFS data V_T to form a dirty map we will obtain

$$\overline{V_T} = (I_o * D_o) + (I_o \alpha' * D_1) + \frac{1}{2} (I_o \alpha'(\alpha' + 1) * D_2) \dots \quad (2.26)$$

where the symbol $*$ denotes convolution and the ‘beams’ D_o, D_1, D_2 etc are defined by

$$D_n = \overline{\sum_{i=1}^N (\Delta\nu_i/\nu_o)^n S(u, v, \nu_i)} \quad (2.27)$$

so that D_n (which will be referred to as the n th order, spectral dirty beam or SDB) is the Fourier transform of the multifrequency uv coverage function

with each observation at frequency ν_i weighted by the value $(\Delta\nu_i/\nu_o)^n$. The weighting schemes that correspond to the different orders of spectral dirty beam are illustrated in fig(2.2) for a simple 3 element East-West array. The above method of representing spectral effects in terms of a Taylor expansion in $(\Delta\nu_i/\nu_o)$ was independently developed by Cornwell (1984) and Conway (1985) and will be referred to as the ‘linear expansion’ analysis. Note that the spectral dirty beams referred to in the above derivation can also be defined as linear combinations of the single frequency dirty beams, from equ(2.27) it follows that

$$D_n = (1/N) \sum_{i=1}^N (\Delta\nu_i/\nu_o)^n D(x, y, \nu_i) \quad (2.28)$$

where N is the number of frequencies and the factor $1/N$ must be included because the Fourier transform operation includes a normalisation operation. Equation(2.28) follows from the fact that the FT of $S(u, v, \nu_i)$ is the dirty beam at frequency ν_i . We can see that the n th order spectral dirty beam is simply some weighted sum of the dirty beams corresponding to each separate frequency. In the simplest case of having only two observing frequencies with a separation of $\Delta\Omega$ in frequency and choosing ν_o as the average of the two frequencies then the resulting D_1 beam will simply be $\Delta\Omega/4\nu_o$ times the difference between the two single frequency dirty beams.

The effects of the expected power law spectrum of most radio sources on MFS can be described by the so-called ‘linear expansion’ as shown above. In a useful expansion the contributions of the terms should decrease as rapidly as possible with increasing order. To what extent does this happen and how many terms are required to characterise all the spectral effects to any given accuracy? To answer these questions the properties of both the SDB’s, and the quantities with which they are convolved (the ‘coefficient images’), must be followed as the order is increased. In the next section the nature and size of the SDB’s are described whilst in section(2.5.1) the ‘coefficient images’ are considered.

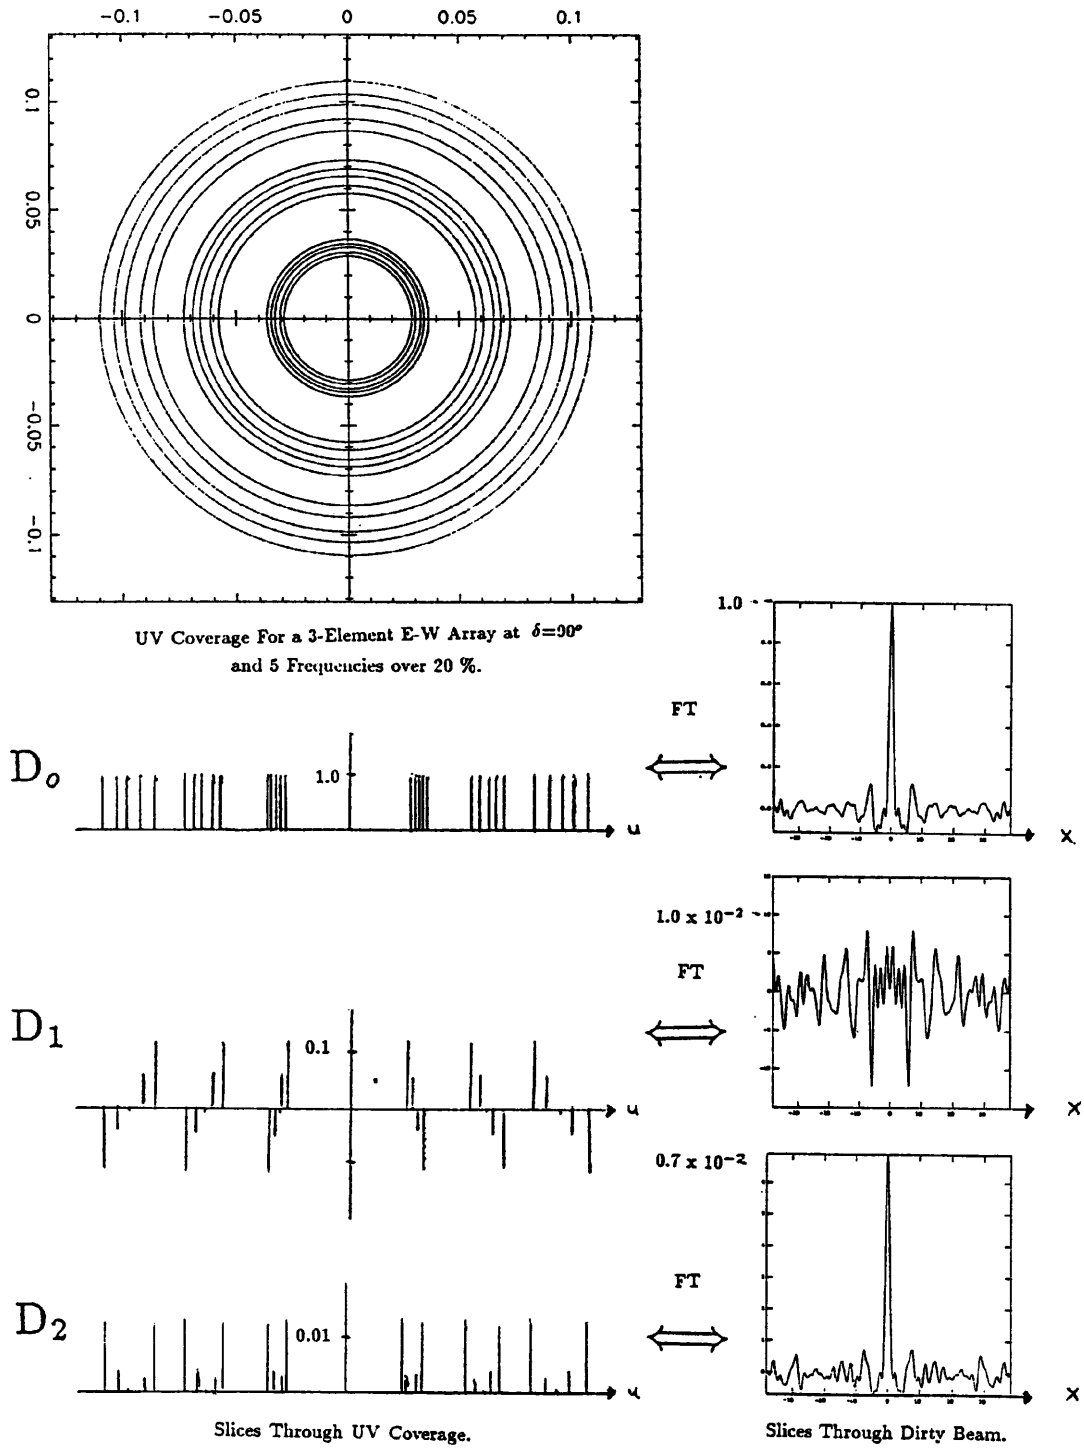


Figure 2.2: UV Weightings and Spectral Dirty Beams for a Simple 3 Element Interferometer

2.4 Some Features of the Spectral Dirty Beams

2.4.1 The SDB's and their UV Weights

The weightings of the uv coverage that correspond to the Fourier transform of the n th order spectral dirty beam, D_n , are $(\Delta\nu_i/\nu_o)^n$. These weightings are illustrated for $n = 0, 1, 2$ in fig(2.2). For a range of frequency or 'bandspread' of $\pm 10\%$ the quantity $(\Delta\nu_i/\nu_o) < 0.1$, where ν_o is the reference frequency which is selected to be in the centre of the band. For this bandspread the weighting $(\Delta\nu_i/\nu_o)^n$ decreases very rapidly as n increases, in this particular case the weighting of the frequencies at the top and bottom of the band decrease by a factor of 10 for each increase in order. With each increasing order the weighting of some intermediate frequency will decrease by some greater factor.

The overall size of the weights and hence the corresponding size of the spectral dirty beams must obviously decrease rather rapidly as n increases. What, however, about the nature of the weighting and its effects on the form of the resulting spectral dirty beam?

We can distinguish sharply here between the weightings when n is even (even spectral dirty beams) and those where n is odd (odd spectral dirty beams). In the case of n being even the quantity $(\Delta\nu_i/\nu_o)^n$ will be positive for all ν_i , whether ν_i is greater or less than ν_o ; however for n which is odd the sign of the weighting will depend on the sign of $(\nu_o - \nu_i)$. The differences between the nature of the two weighting functions are illustrated for a simple uv coverage in fig(2.2). As we shall see below this difference between the signs of the weights causes the odd spectral dirty beams to have no central peak.

2.4.2 The Centres of the SDB's

The D_0 and D_1 spectral dirty beams that result from a realistic MERLIN aperture plane coverage, and five frequencies spread over 25% in frequency, are illustrated in fig(2.3) and fig(2.4). The most notable difference that we see between the beams in these two figures is the absence of a central peak in

the case of the $n=1$ beam.

We recall the well known result (Bracewell 1965) that the value at $x=0$ $y=0$ of any function of x and y , $f(x,y)$ is equal to the integral over all space of its Fourier transform. In the case of the spectral dirty beams, their Fourier transforms consist of a set of delta functions multiplied by values which depend on the observational frequency. Since the areas of the delta functions in uv space are unity the integral of the area of the sampling function will equal the summation of the weighting values over all the uv measurements. The result of this summation (normalised by the number of uv points) will equal the value at the centre of the appropriate order spectral dirty beam.

For even beams the weights are all positive, the corresponding beams are very similar in form to the normal dirty beams with their large central lobes and subsidiary sidelobes. It follows from the positivity of the uv weights that the central lobe of such beams will also be their peak. This peak value is equal to the mean weight given to the corresponding uv coverage, a value which is easily calculated.

For an odd spectral dirty beam, whose uv weights consist of both positive and negative values, the sum over the uv plane will be reduced compared to adding up values that are all positive; hence the small central value of the D_1 beam. Because of this effect the peak of the D_1 does not occur at its centre but rather at some 'sidelobe'. The level of this peak is not as easily calculated as in the even case.

An important consequence of having weights of different sizes for $n=1$ is that we may choose the free parameter ν_o , the reference frequency, so that the central value of the $n=1$ spectral dirty beam is exactly zero. Using the relationship connecting the central value of a function to the integral over its Fourier transform, (i.e. the sum of its uv weights) we note that after normalisation,

$$D_n(x = 0, y = 0) = \frac{\sum_{i=1}^N M_i (\Delta \nu_i / \nu_o)^n W_i}{\sum_{i=1}^N M_i W_i} \quad (2.29)$$




Figure 2.3: Greyscale plot of the D₀ beam. Linear plot from -5×10^{-2} to 5×10^{-2} . Figure is 25.6" on a side. **Central peak value is unity.**




Figure 2.4: Greyscale plot of the D₁ beam. Linear plot from -5×10^{-3} to 5×10^{-3} . Figure is 25.6" on a side.

where M_i is the number of uv points at the frequency ν_i and W_i is the mean weight attached to the uv points of frequency ν_i due to the various conventional weighting schemes such as ‘natural’ and ‘uniform’ weighting (Sramek and Schwab 1985). If we require that $D_1(0,0) = 0$ then it follows that,

$$\nu_o = \frac{\sum M_i W_i \nu_i}{\sum M_i W_i} \quad (2.30)$$

if we choose ‘natural’ weighting so $W_i = 1$, then the value of ν_o that gives $D_1(0)=0$ will be the arithmetic mean of the observing frequencies averaged over all the uv points. If in addition there are the same numbers of uv points at each frequency then ν_o equals the mean of all the frequencies. In the more usual case of uniform weighting the exact calculation of the optimum reference frequency will depend on the values of W_i in equation(2.29) which in turn depends on the uniform weights assigned by the mapping program.

If ν_o is chosen so that $D_1(0,0)=0$ this will mean that at the position of any point source the dirty map will be unbiased by the effects of 1st order terms due to that point itself. This means that if we deconvolve a MFS dirty map with the D_o beam then any isolated point source will have its D_o sidelobes successfully removed by CLEAN. To a first approximation the presence of the first order spectral effects will not affect the deconvolution.

2.4.3 The Sidelobes of the Beams

Besides the absence of a central peak in the case of odd spectral dirty beams, how else do the even and odd beams differ? From fig(2.3) and fig(2.4) we see that away from the very centre the general features of the D_o and D_1 beams are similar. From fig(2.3) and fig(2.4) we see that in both cases the near-in sidelobes consist of roughly elliptical arcs whilst further out the beam structure is almost noise-like. The behaviour at large radius can be understood as being due to summing of the contributions from each uv point with almost random phase.

If the relative sizes of the sidelobes of the two beams are compared it is found that the level of the D_1 beam in the centre appears to be nearer

to 10^{-2} of D_0 than the 10^{-1} expected simply from considering the relative sizes of the weightings, only at large radii do the relative size of the sidelobes approach the expected value. The reason for the lower levels is again that the D_1 beam, being odd, has weights of different sign. Toward the centre of the beam where the effects of the weightings are added coherently there will be a significant reduction in the case of odd beams. At large radii where the effects are added incoherently and the sidelobes appear random this mechanism will not be effective and the relative sizes of the sidelobes will simply depend on the amplitudes of the weights.

The reduction in magnitude of the near-in sidelobes by the above effect is important, it means that the peak levels of the odd beams do not depend on the amplitudes of the uv weights in the simple way of the even beams. The peak levels are instead much lower than would be predicted from a knowledge of the mean amplitudes of the weights alone. As an example take the first order beam shown in fig(2.4) where the bandwidth is 25%. Despite the fact that the typical amplitudes of its uv weights are around 10^{-1} we find that the peak level of the D_1 beam is the much smaller value of 5×10^{-3} . This peak value means that for a point source of flux I and residual spectral index α' the peak first order sidelobe level will be about $I\alpha' \times 5 \times 10^{-3}$. This small level of spectral sidelobe means that spectral effects can often be ignored entirely because the expected effects are less than the noise.

2.4.4 The Derivative Relationship Between Beams

Away from the centre the general structure of the sidelobes of even and odd types of beam are very similar, consisting of arc like sidelobes (see fig(2.3) and fig(2.4)). Close inspection of these sidelobes in D_0 and D_1 shows that the maxima and minima of the two beams do not coincide, in fact they appear to be almost orthogonal (fig 2.5). This property is a consequence of the definition of the beams.

In section(2.3.3) it was described how D_n can be considered to be made up of a linear combination of the normal dirty beams from each individual




Figure 2.5: Comparison of D_0 and D_1 Beams Showing their Approximate Orthogonality. North-South slices taken through fig 2.3. and fig 2.4

observing frequency. If we use the same antennae and hour angle ranges at each frequency then the shapes of these beams will be identical except for a change in scale of their radial dimensions. The D_n beam can be considered as the linear sum of the individual frequency dirty beams $D(\nu_i)$ multiplied by the factor $(\Delta\nu/\nu_o)^n$. The value of D_1 in polar coordinates, at radius θ and azimuth ϕ will therefore be

$$D_1(\theta, \phi) = (1/N) \sum_{i=1}^N D(\theta, \phi, \nu_i) (\Delta\nu_i/\nu_o) \quad (2.31)$$

where $D(\theta, \phi, \nu_i)$ refers to the value of the single frequency dirty beam at frequency ν_i at polar coordinates θ, ϕ . We note that the single frequency D beams scale with frequency so that

$$D(\theta, \phi, \nu_i) = D((\nu_i/\nu_o)\theta, \phi, \nu_o) \quad (2.32)$$

i.e. the single frequency beam at ν_i is a version of a single frequency beam at the reference frequency scaled in radius by a factor depending on ν_i/ν_o . Inserting the above relation into our summation we see that for $\nu_i > \nu_o$ the contribution to D_1 at a particular θ_o will be a positive number times $D(\theta', \nu_o, \phi)$ where $\theta' > \theta_o$, whilst for $\nu_i < \nu_o$ the contribution will be a negative number times $D(\theta', \nu_o, \phi)$ where $\theta' < \theta_o$. The value of D_1 at θ_o , will therefore be related to the change of $D(\theta', \nu_o, \phi)$ with respect to θ at θ_o . A given fractional range in frequency means that we sample $D(\nu_o)$ over the same fractional range in θ . For a small enough θ_o the $D(\nu_o)$ beam will change almost linearly with θ over this range and hence the value of the D_1 beam will depend on the local gradient of $D(\nu_o)$. This simple argument gives us a physical picture of the nature of the D_1 beam as being related to the derivative of the single frequency dirty beam with radius. In Appendix A we consider the exact relationships between the derivatives of $D(\nu_o)$ and any D_n . We can, however, note here that the above argument explains our observation that the sidelobes of D_o and D_1 are roughly orthogonal.

If the sidelobes of $D(\nu_o)$ along a given azimuth are roughly cosine-like then the derivative of $D(\nu_o)$ with respect to θ and hence the D_1 beam will be roughly sine-like and so orthogonal to $D(\nu_o)$. Near to the centre the sidelobes

of D_o will be approximately equal to those of $D(\nu_o)$ and so D_o and D_1 are roughly orthogonal.

An important quantitative result of the more detailed analysis given in Appendix A is that the size of D_1 at θ_o from the centre is equal to

$$D_1(\theta_o, \phi) \approx (1/N)C_{1,1} \frac{dD(\nu_o)}{d\theta} \Big|_{\theta=\theta_o} \cdot \theta_o \quad (2.33)$$

where

$$C_{n,m} = \sum (\Delta\nu_i/\nu_o)^{n+m} (-1)^m \quad (2.34)$$

For the case of 5 frequencies spread equally over 25% in frequency then $C_{1,1} = 3.9 \times 10^{-2}$. The above formula will be a good approximation provided θ_o is not too large. If the peak level of D_1 occurs in the region where the above equation is valid, as it normally does, then the maximum sidelobe level of D_1 will occur at $\theta = \theta_{max}$ such that

$$\frac{dD(\nu_o, \theta, \phi)}{d\theta} \cdot \theta \quad (2.35)$$

is a maximum and will be equal to

$$C_{1,1} \frac{1}{N} \theta_{max} \frac{dD(\nu_o)}{d\theta} \Big|_{\theta=\theta_{max}} \quad (2.36)$$

Given a knowledge of the the conventional single frequency dirty beam it is therefore possible to estimate the peak level of the D_1 beam. For a set of almost randomly distributed antennae it is a general observation that the basic features of $D(\nu_o)$ depend only weakly on the exact uv coverage. For arrays of less than say a dozen randomly distributed antennae the $D(\nu_o)$ beam will consist of a large central value with the addition of sidelobes of level 5%-20% of the peak. In such arrays therefore the behaviour of the quantity in equ(2.35), and hence D_1 will be only weakly dependent on uv coverage. For most single frequency dirty beams the peak of the function in equ(2.35) will occur near radius $\theta = b$ where b is the HPBW of the central main lobe. The value of the gradient of $D(\nu_o)$ at this point is

$$dD/d\theta \approx 1/2b \quad (2.37)$$

because the value of $D(\nu_o)$ falls from unity to zero in going from the centre to $\theta = 2b$. The peak value of D_1 can therefore be expected to be about $C_{1,1}/2N$ under a wide range of conditions i.e. at a level of 3.9×10^{-3} for the example of five frequencies over 25%. We shall see in section(3.1.2) that the above argument predicts the actual behaviour of the peak of the D_1 beam reasonably well.

2.4.5 Summary

We have seen that because of the nature of the uv weighting schemes we can expect the spectral dirty beams to decrease in level rapidly when the fractional bandspread is significantly less than unity. We have also seen that the odd numbered beams have considerably smaller peak values than would be expected from the above argument. This is because whilst even beams are similar in structure to conventional dirty beams odd beams differ in having no central peak and having small central values. This is due to the fact that the D_1 beam and all other odd beams are produced from weightings that are both positive and negative in the uv plane.

A useful physical picture of the D_1 beam is to consider its relationship to the derivative of the single frequency beam at frequency ν_o . Pursuing this idea quantitatively gives us a means to estimate the peak level of D_1 and other odd beams by a method that is relatively independent of the exact details of uv coverage. For a realistic example of a 25% bandspread and 5 frequencies this method allows us to calculate the peak level of the D_1 beam to be less than $1/200$. The peak size of even beams is easily calculated by finding the mean weight attached to the corresponding uv coverage. In the above case a peak level smaller than $1/100$ can be predicted for the D_2 beam.

2.5 The Size of Linear Spectral Errors

2.5.1 The Coefficient Images

Having discussed the dirty beams we now consider the quantities with which these spectral dirty beams are convolved in equ(2.26), quantities which we will term ‘coefficient images’. Once we know the behaviour of both the beams and the coefficient images the expected level of spectral errors can then be estimated.

We recall from the above linear expansion equation that the the n th order term is given by

$$D_n * I\alpha'(\alpha' + 1)\dots(\alpha' + n - 1)/n! \quad (2.38)$$

where the n th order coefficient image is the quantity to the right of the convolution operation.

It is useful to consider approximations for the coefficient images based on our knowledge of how $I(x,y)$ and $\alpha'(x,y)$ behave in realistic cases. In ‘interesting’ images the range of I represented is often greater than 10^3 . In contrast the range of α' is rarely as great as -1.0 to $+2.0$. Because of the large range in I compared to the comparatively narrow range of α' the n th order coefficient is usually largest where I is largest and smallest where I is smallest.

The regions of bright emission (hotspots, knots, jets etc) which will dominate the coefficient images often have only narrow ranges of spectral index. This means that after removing a suitable overall spectral index α_o the residual spectral index α' will only span a small range in these bright regions. To take the case of the ‘target source’ described in sect(1.6.3) although α varies over the range 0.5 to 2.0 , the α in the dominating bright regions lies in the range 0.5 to 0.7 , so that after removing $\alpha_o = 0.6$ the range in α' will be only 0.1 . In the regions of high I which dominate the coefficient images it is therefore usually the case that $|\alpha'| \ll 1$. If we evaluate the n th order coefficient image using the above assumption we obtain

$$\frac{I\alpha'(\alpha' + 1)\dots(\alpha' + n - 1)}{n!} \approx \frac{I\alpha'(n - 1)!}{n!} = \frac{I\alpha'}{n} \quad (2.39)$$

This result shows that the size of the largest parts of the n th order coefficient image decreases only as $(1/n)$, i.e. only slowly with increasing order. This slow decrease is a disappointment as it may mean that several terms of the linear expansion are required to describe spectral effects adequately.

2.5.2 Linear Spectral Errors from a Point Dominated Source

To determine the size of the spectral sidelobes we must consider the contribution from the different terms in the linear expansion. We can find these contributions by convolving the appropriate n th order beam with the corresponding n th order coefficient. The exact result of this convolution will naturally depend on the form of the intensity and spectral index distributions. A simple and important case however is that of a source dominated by a single point component of flux I and residual spectral index α' . In this case each coefficient image will also be point-like, the resulting n th order contribution will therefore equal the n th order SDB multiplied by a quantity that depends only on I and α' .

In sect(2.4.5) order of magnitude estimates of the sidelobe levels of the different SDB's were made. These estimates were based only on a knowledge of the fractional bandspread and the number of frequencies. We shall see in Chapter 3 that these estimates are fairly robust and are not greatly affected by the details of uv coverage or declination etc. These estimates can therefore be used to give a idea of the contributions of different orders of spectral sidelobes in the case of a point source. Because the assumed levels of the beams are only weakly dependent on the observing parameters the resulting estimates will be of general utility in estimating (to within say a factor of two) the spectral errors in most observations.

As in sect(2.4.5) we will assume the standard 5 frequencies over a 25% bandspread. The resulting D_0 beam which has a central value of unity will have sidelobes of less than about 0.1. The reference frequency can be chosen so that the D_1 beam can be assumed to have a zero central value. The peak

Order	On Source Errors	Near-in Sidelobes
0	I	I/10
1	0	$I\alpha'/200$
2	$I\alpha'/200$	$I\alpha'/2000$

Table 2.1: Contributions from the Different Orders of the Linear Expansion for a Point Source

of this beam will be approximately $1/200$. Finally the D_2 beam can be taken to have a central value of about $1/100$ and sidelobes which are $1/10$ of this peak at a level of $1/1000$. Note that for the particular case of the 8-station MERLIN at a moderate declination (say 50°) these levels for the SDB's are if anything slight overestimates.

In Table(2.1) are listed the contributions from the different orders of the linear expansion assuming the above sizes of beam and assuming a point source of flux I and residual spectral index α' . The spectral errors are quantified at two different positions; 'on-source' i.e. at the position of the point component and 'near-in' which refers to the maximum level of sidelobe within a few beam diameters of the point component. It is assumed in Table(2.1) that α' is small so that the approximation of equ(2.39) holds. The figures quoted in Table(2.1) are approximate but they suffice to give a general impression of the contribution from each order of the expansion. In the next chapter simulations will be used to find the contributions of the different orders in the case of non-point components and realistic distributions of $I(x,y)$ and $\alpha'(x,y)$. It will be found that these results are similar to the those found for the point source case.

The results in Table(2.1) show that because D_0 and D_2 have similar peak values and because of the slow $(I\alpha'/n)$ decrease in the size of the coefficient images, the peaks of the first and second order spectral sidelobes are similar. The contributions of the first and second order spectral sidelobes together dominate the spectral contribution to the MFS dirty map. The peak level of each will be of order $I\alpha'x(D_1)_{max}$ which, since $(D_1)_{max}=1/200$, gives $I\alpha'/200$. If the thermal noise is greater than this value, that is to say that the noise limited dynamic range is less than $200/\alpha'$ then we can ignore first, second

and higher order spectral terms. We can then proceed by ignoring all spectral effects entirely and simply process the MFS data as we would conventional uv data by attempting to ‘deconvolve’ D_0 . If the noise is less than $I\alpha'/200$ then a full description strictly requires the contributions of at least the zeroth, first and second order spectral terms.

In Chapters 4 and 5 we will examine algorithms which attempt to separate out the sky plane effects of various orders that contribute to the MFS image. Once the effects of all orders other than the zeroth are removed the corresponding uv data can be mapped conventionally and the benefit of improved uv coverage gained without the corrupting effects of spectral errors. The most successful of these algorithms will assume that the sky plane effects of the MFS data are well described by only two terms of the expansion, i.e. the zeroth and the first, the effects of higher terms will cause problems for these algorithms.

For the point source case the peak effect of the first and second order terms in the linear expansion are comparable. However, we would much rather have a second order contribution which is less than the first order contribution so that the correction algorithms could be readily applied.

The fact that the maximum sky effect of second order errors occur ‘on source’, because of the large central peak of the D_2 beam, to some extent ensures that the errors that arise from a significant second order contribution are relatively benign and will not interact with the ‘two order’ processing algorithms. Although such errors might be significantly larger than the noise they will only be larger than the noise on regions of strong emission. The largest fractional error due to these ‘on source’ second order errors will therefore be limited. In our example case the fractional error will be $\alpha'/200$, the small bias this error introduces will never effect the interpretation of gross features and even in quantitative work it is unlikely to be significant since it will probably be less than the error in the map’s overall flux scale.

The peak contribution of the second order spectral errors described above

will define a ‘Strong Limit’ such that, if the noise is less than the peak second order contribution, the second order terms cannot be strictly ignored. If the on-source bias errors are considered acceptable however then the critical quantity is the level of error caused by the sidelobes rather than the peak of the D_2 beam. These errors will be about a factor of 10 less and will define a ‘weak limit.’ Provided a source is within the weak limit it will be found that the ‘two order’ processing algorithms will continue to work well. There will be some interaction with the on-source ‘bias’ errors however. Because these will be assumed to be part of the source structure in any ‘deconvolution’ step a contribution equal to the bias convolved with D_o will be removed. This interaction will not be serious; because the D_0 and D_2 beams are very similar in shape it will if anything reduce the effective contribution to the image due to D_2 sidelobes.

Despite this distinction between ‘Strong’ and ‘Weak’ limits, which will extend the range of sources over which the two order approximation holds, it clearly would be best if the contributions of the various terms decreased more rapidly with order. In the next section we will consider a modified expansion which will provide a more compact description of spectral effects so that the two order approximation will be valid for a larger range of sources.

2.6 The Logarithmic Expansion

2.6.1 Description and Properties of the Logarithmic Expansion

The linear expansion that we have discussed in the previous section is only one expansion out of an infinity of possibilities. Instead of expanding the function $I(x,y,\nu_i)$ in terms of powers of $(\Delta\nu_i/\nu_o)$ it is possible instead to expand it as a power series in any function of $\nu, f(\nu)$. It is shown in Appendix B that choosing the function $f(\nu_i)$ to be $\ln(\nu_o/\nu_i)$ when the source has a SVPL spectral distribution has advantages over the linear expansion. This new ‘logarithmic expansion’ is superior because the contributions of the higher terms are less

than for the linear expansion. In this section we will not attempt to derive the logarithmic expansion, this is done in Appendix B, instead the form of the expansion will be briefly presented and its main features described.

In Appendix B it is shown that if an ‘overall’ spectral index α_o is first removed from the MFS data as described in section(2.3.2) so leaving a space varying spectral index $\alpha'(x, y)$ and then the resulting uv data is inverted to form a dirty map then this dirty map can be expressed as

$$V_T = (I * D'_o) + (I\alpha' * D'_1) + (I\alpha'^2/2 * D'_2) + \dots \quad (2.40)$$

$$= \sum_{n=0}^{\infty} (I\alpha'^n * D'_n) / n! \quad (2.41)$$

This differs from the linear expansion in that the coefficient images are changed to

$$(I\alpha'^n / n!) \quad (2.42)$$

as a consequence of this change the definition of the spectral dirty beam of order n, D'_n is changed so that

$$D'_n = \overline{\sum_{i=1}^N [\ln(\nu_o / \nu_i)]^n \cdot S(u, v, \nu_i)} \quad (2.43)$$

The logarithmic spectral dirty beams are therefore exactly analogous to the linear spectral dirty beams except that instead of having a uv point observed at ν_i weighted by $(\Delta\nu_i / \nu_o)^n$ it is instead weighted by $[\ln(\nu_o / \nu_i)]^n$.

For both the linear and logarithmic cases the uv weightings when n is even are all positive. The sign of the weightings do not depend on whether $\nu_i > \nu_o$ or $\nu_i < \nu_o$. In contrast, under either expansion, when n is odd the sign of the weightings will depend on the sign of $\nu_o - \nu_i$. The different nature of the weightings for even and odd logarithmic beams will cause differences in form between the different types of beam just as in the case of the linear expansion (see section(2.4)). We can therefore expect the general properties of the logarithmic SDB's to be similiar, i.e. we can expect odd SDB's to have small centres and relatively small peak values.

The differences between the values of the weightings under the logarithmic expansion and under the linear expansion are quite small when we have

	Frequency (MHz)				
	1730	1615	1540	1415	1350
$\Delta\nu/\nu_o$	-0.131	-0.056	-0.006	0.075	0.118
$\ln(\nu_o/\nu)$	-0.123	-0.054	-0.006	0.078	0.132

Table 2.2: Weightings of the Linear and Logarithmic Expansions

a moderate bandwidth. In Table(2.2) are compared the values of the logarithmic and the linear weighting at five frequencies for a total bandwidth of approximately 25%. The same ‘reference’ frequency of 1530 MHz is assumed for both types of expansion. The reason for the close similarity between the weightings is that for small values of $\nu_o - \nu_i$ compared with ν_o it can be shown that $\ln(\nu_o/\nu_i) \approx \Delta\nu_i/\nu_o$ (see Appendix B) and hence in the limit of small fractional bandwidths the two weighting schemes become identical. For the example of the 25% bandwidth the differences as shown in Table(2.2) amount to only ten percent at the highest and the lowest frequencies. Because the two weightings are so similar in size and because the weightings exhibit the same differences in form between odd and even orders for any given set of frequencies the logarithmic and linear beams of order n are very similar. In fig(2.6) are shown slices through the D_1 and D'_1 beams, in this case the reference frequencies are taken to be the arithmetic and geometric mean frequencies. The two types of beams are very similar.

In section(2.4.2) it was shown that the central value of the linear D_1 beam could under certain conditions be made zero by choosing the reference frequency to be the *arithmetic* mean of the observed frequencies. If we wish to make the central value of the logarithmic beam D'_1 zero, we should instead take the expansion about a reference frequency ν_o which is the *geometric* mean of the frequencies over all uv points. To see this consider the condition that $D_1(\theta = 0, \phi) = 0$.

$$\frac{\sum M_i W_i [\ln(\nu_o/\nu_i)]^n}{\sum M_i W_i} = 0 \quad (2.44)$$

where M_i is the number of uv points at frequency ν_i and W_i is the mean weight at frequency ν_i due to natural or uniform weighting schemes. All summations are taken over the frequency index i . From the above equation




Figure 2.6: Comparison of Slices Through Logarithmic and Linear First-Order Spectral Dirty Beams.

we obtain

$$\sum M_i W_i \ln(\nu_i) = \sum M_i W_i \ln(\nu_o) \quad (2.45)$$

hence

$$\nu_o = \left[\prod_{i=1}^N \nu_i^{M_i W_i} \right]^{(1/\sum M_i W_i)} \quad (2.46)$$

Which if W_i and M_i are constant for all i is the geometric mean of the observing frequencies.

2.6.2 Logarithmic Spectral Errors from a Point Dominated Source

We have seen that the linear and logarithmic beams are very similar in both size and overall morphology. The sizes of for instance, the first order spectral dirty beams in the two cases seem comparable and so from this point of view there does not seem to be any benefit in using the logarithmic expansion. The advantage of using the logarithmic expansion rather than the linear lies instead in the form of the coefficient images with which the various orders of beam are convolved to form the MFS dirty image. In the linear case the n th order coefficient image is

$$\frac{I\alpha'(\alpha' + 1) \dots (\alpha' + n - 1)}{n} \approx \frac{I\alpha'}{n} \quad (2.47)$$

where the approximation holds if $|\alpha'| \ll 1$, whilst for the logarithmic expansion the coefficient image is given in equ(2.42). If $\alpha' \ll 1$ then the logarithmic coefficient will be smaller than the linear expansion by the factor

$$\frac{\alpha'^{n-1}}{(n-1)!} \quad (2.48)$$

which for $n=2$ amounts to α' .

We have already seen how we can expect the values of α' to be quite small in bright regions of a source. Going from a linear to logarithmic expansion can therefore have a significant effect in reducing the size of the coefficient images at these bright regions and hence the spectral sidelobes which surround them. These spectral sidelobes will often dominate the spectral error contribution and so their reduction will significantly reduce the overall spectral errors.

Order	On Source Errors	Near-in Sidelobes
0	I	I/10
1	0	$I\alpha'/200$
2	$I\alpha'^2/200$	$I\alpha'^2/2000$

Table 2.3: Contributions from the Different Orders of the Logarithmic Expansion for a Point Source.

The MFS algorithms that we will consider in Chapters 3 and 4 will require that the dirty map can be well described within the noise by only two orders, the zeroth and first. Unfortunately under the linear expansion the peak contribution of the second order term will often be comparable to the contribution of the first order term. The use of the logarithmic expansion will help reduce this second order contribution so making the two order approximation more applicable.

To compare the logarithmic and linear expansions we can again consider the situation of a point source of flux I and residual spectral index α' (see section(2.5.2). In Table(2.3) are listed the contributions of the different orders of spectral sidelobe, from these figures the corresponding weak and strong limits on dynamic range at which the two order approximation holds can be defined.

The most important change which is evident in comparing Table(2.3) to Table(2.1) is that the strong dynamic range limit for the two order approximation changes from $200/\alpha'$ to $200/\alpha'^2$ in going from the linear to the logarithmic expansion. For bright parts of a radio source it can usually be assumed that $|\alpha'| < 0.2$ and often $|\alpha'| < 0.1$ so that the improvement in dynamic range if such bright regions are dominant is between 5 and 10.

This increase in the dynamic range limit should allow us to use the two order approximation for sources as ‘difficult’ as our target source. For this target source we recall that $\alpha' = 0.1$ at the hotspots. If these hotspots have flux of approximately 500mJy then the logarithmic expansion gives a peak second order contribution (strong limit) at 0.5×10^{-4} of I (a dynamic range of 2×10^4)

or about $25 \mu \text{ Jy.beam}^{-1}$ which is near the expected thermal noise. The corresponding weak limit is $2.5 \mu \text{ Jy.beam}^{-1}$ which is well below the noise. This calculation is only approximate and assumes that point-source hotspots dominate the second order contribution. For a more realistic case see section(3.2.3) where the second order effects are simulated for the model target source for both the linear and logarithmic expansions. These simulations show a marked superiority for the logarithmic expansion over the linear.

It has been shown in this section that the logarithmic expansion is better than the linear because of its faster convergence. Because there are no competing disadvantages associated with it we shall employ the logarithmic expansion rather than the linear expansion in the rest of the thesis. In any future reference to beams or expansions it should be assumed that the logarithmic variety is meant unless otherwise stated. The unprimed symbol for the n th order SDB, viz D_n will be taken to refer to the logarithmic version of the dirty beam.

In the present chapter we have considered the mathematical form of the spectral contributions, in the next chapter the spectral errors will be considered as functions of practical observational parameters such as declination and fractional bandspread for the specific case of MERLIN. In the second half of Chapter 3 the assessment of the contributions from SDB's will be extended to include non-point dominated images.

Chapter 3

MFS and Frequency Dependent Structure II: Spectral Effects in Realistic Cases

3.1 The D_1 Beam as a Function of Observing Parameters

In this section we extend our discussion of the size and nature of the D_1 beam to see how the beam changes as we change observing parameters such as bandspread and declination. If we assume a two order expansion then study of the D_1 beam is vital since it is solely through this beam that the effects of spectral index will influence the dirty map. Knowledge of its properties will then allow us to define the sources in which spectral effects can be ignored entirely and those in which processing will be required to remove the effects of the D_1 beam.

3.1.1 Effects of Array Geometry

How dependent is the size and form of D_1 on the geometry of the observing array? For the most part simulations in this thesis will concentrate on MERLIN, but it is as well to briefly consider other array geometries. In fig(3.1) are shown slices through the D_1 beam for three arrays, namely MERLIN (which we assume is typical of many current VLBI arrays); an example of a typical East-West array, namely the 8 antenna 5km array (Ryle 1975); and a VLA snapshot observation. In all cases the same five frequencies spread over 25% in frequency and a source declination of 50° are assumed.

The results for long-period observations with MERLIN and the 5km array show similarities despite the different numbers of antennae and their different distribution. The spectral sidelobes for the VLA in ‘snapshot’ mode using the ‘standard frequencies’ are however somewhat larger. The fact that the D_1 sidelobes for the two long period observations are similar can be understood in terms of the discussion in section(2.4.4) about the relationship between the D_1 beam and the derivatives of the single frequency beam $D(\nu_o)$. For both MERLIN and the East-West array the single frequency beams are rather similar, consisting of a main lobe surrounded by ring-like structures which vary smoothly with radius.

In contrast the single frequency sidelobes for the VLA are very different. We have the characteristic ‘star’ shape beam and along the arms of the star the level of the beam varies in an irregular fashion. This is due at root to the concentration of uv samples along intra-arm baselines. The effect of the irregular single-frequency beam is to increase the level of the D_1 beam because this depends on the derivatives of the single frequency dirty beam. Overall it seems that the behaviour of the D_1 beam can be deduced from examining the properties of the single frequency beams alone, and hence from the single-frequency uv coverage.

This principle seems secure for the case of random arrays of telescopes or grating arrays; however, for arrays which have a power law scaling of baseline length it will sometimes break down when the frequency coverage has a special relationship to the uv coverage. In such arrays it is possible to choose a set of frequencies so that different physical baselines at different frequencies exactly, or very nearly, sample the same uv points, a situation which we will term ‘overlap’ or ‘double sampling’. Although partial overlap is possible for forms of array geometry other than power law arrays such behaviour is dependent on coincidences between particular physical baselines and frequencies. Only in the case of scaled arrays is it possible to arrange that most baselines will be involved in such overlap.

To take an extreme example consider a power law or scaled array such that




Figure 3.1: Slices through D₁ beams from different arrays assuming five frequencies over 25%. The slices through the 5km and MERLIN beams are in the N-S direction whilst the VLA slice is at an azimuth at which the sidelobes are a maximum.

along any radius in the uv plane the uv tracks obey $r = r_o f^n$ where r is the distance from the centre of the uv plane; r_o is a constant, n an integer and f some scale factor. If observations are made at two frequencies $f^{1/2}\nu_o$ and $f^{-1/2}\nu_o$ then we will get exact overlap between the uv points due to the lower frequencies and those due to the higher frequencies (see fig(3.2)). Under the logarithmic expansion the uv weights given to the FT of the D_1 beam at the two frequencies will be $\ln(f^{-1/2})$ and $\ln(f^{1/2})$ respectively. These weights will be equal in magnitude but opposite in sign. Because points due to the two frequencies will be coincident in the uv plane their effects on D_1 will cancel out exactly. Only the longest physical baseline at the higher frequency and the shortest physical baseline at the lower frequency will be unpaired and so contribute to the D_1 beam. If these two baselines are deleted the result will be a D_1 beam which is zero. For this array of special geometry the choice of a particular set of frequencies results in a D_1 beam whose size cannot be directly forecast from the properties of the single frequency dirty beam or the single frequency uv coverage.

A obvious drawback in having overlap is that, because uv points are paired, the total effective uv coverage is worse than it would be if there were no special relationship between the baseline lengths and the frequencies. If all uv points are paired the effective uv coverage is as it would be if we had $N/2$ separate frequencies rather than N . A two frequency MFS system is therefore pointless for improving the fidelity of the intensity image. The special relationship ensures that D_1 is zero but the uv coverage is no better than in conventional imaging. If observations are made at a third frequency, ν_o , then we do get a useful improvement; the extra frequency being at the reference frequency adds nothing to D_1 or any other SDB but yields extra uv coverage (see fig(3.2)). The whole system of three frequencies will have double the uv coverage obtained in conventional methods yet there will be no first order spectral contributions because D_1 is zero; indeed there will be no contributions from any odd order because all odd order SDB's will be zero.

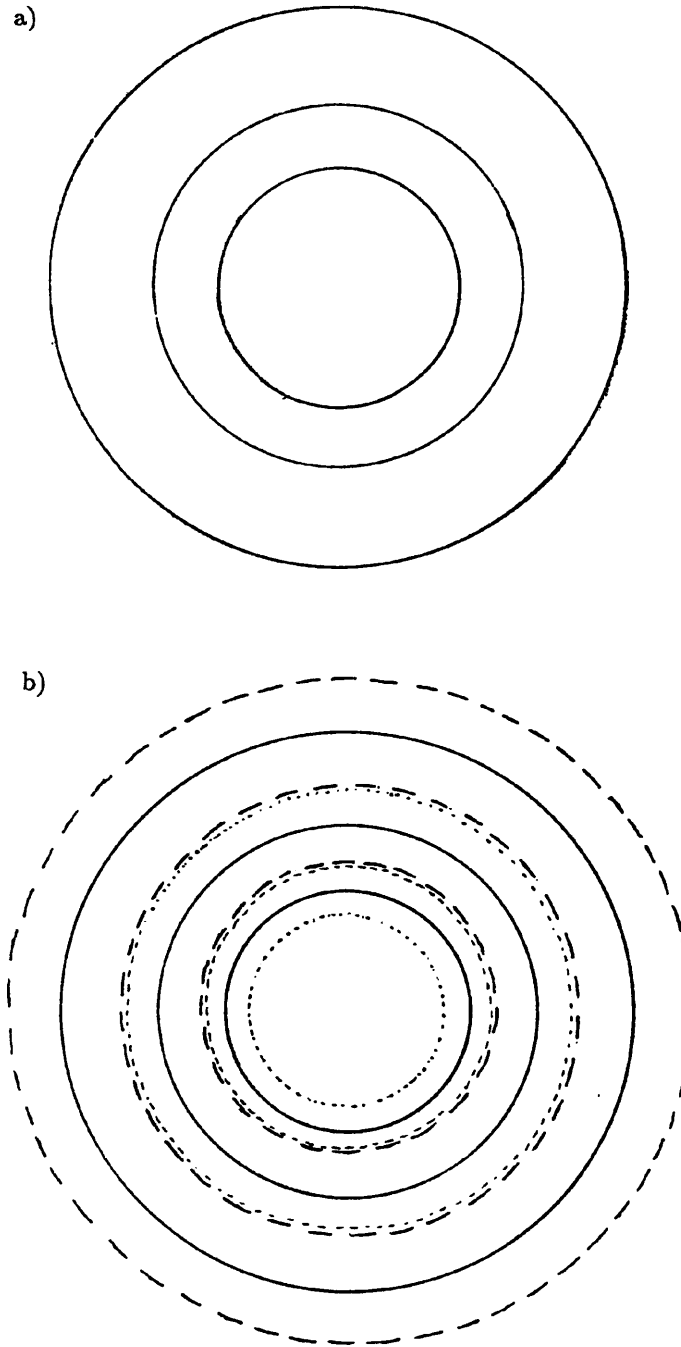


Figure 3.2: Examples Demonstrating Scaled-Array MFS. a) Single frequency uv coverage. b) A three frequency system with double the uv coverage of a single frequency observation. The dotted lines represent observations at the lower frequency, the dashed lines represent the higher frequency and the solid lines the geometric mean frequency.

For systems of more than three frequencies with a scaled array the optimum choice of frequencies to reduce spectral effects becomes complicated. It should be possible to arrange things so that nearly all uv points are paired; however, unlike in the three frequency case it will usually not be possible to pair together uv points with D_1 weightings which are opposite in sign but equal in magnitude. This means that spectral effects will not be automatically reduced to zero. In such systems one can however arrange to apply extra weightings to each uv measurement so that first order spectral effects do cancel out. In chapter 4 such methods will be described in more detail.

With special array geometries there are clearly effective but specialised ways to reduce spectral effects. In this thesis, however, I will not pursue scaled-array methods, mainly because at present no such arrays actually exist. Although the antennae along the arms of the VLA are often described as fitting a ‘power-law’ distribution the nature of this distribution is quite different from that discussed above. For the VLA the distance of the n th antenna along an arm $r(n)$ fits $r(n) = r_o(mn)^f$ with $f = 1.716$ and the integer m being between 1 and 4 depending on the ‘array’ chosen (Napier et al 1983). In contrast in a frequency scaled array $r(n) = r_o f^n$. Although the design of the VLA array geometry does allow ‘scaled array’ observations between widely separated frequencies to be made (i.e. between 20cm and 6cm, a ratio of 3.285 in frequency) it also makes general frequency overlap between frequencies spread over a typical MFS range of 25% impossible. Because the radial distance of one antenna and the next along an arm most definitely *does not* have a simple ratio regular overlap cannot be achieved.

A second reason for not pursuing the case of arrays with special geometries is that in order to get overlap, and so get the large reductions in spectral effects, the special geometries demand special sets of frequencies. It is likely that in practice the choice of frequencies will be strongly limited by the presence of RF interference (see Chapter 6). This will, in turn, limit the use of special-geometry-based MFS, in densely populated industrialised areas at least. It appears however that the present design of the GMRT (Swarup 1987) (see

sect(8.4.3)) consists of a ‘Y’ with the same ‘power law’ distribution as the VLA. Use of special geometry methods therefore seems precluded at a site where the interference environment might well allow the use of special frequencies.

Whether having frequency scaled arrays is a good idea is an interesting question. The situation is in many ways analogous to the question of redundancy in arrays used for self-calibration (Cornwell 1985). Having redundancy allows special techniques to be used yet it reduces the quality of the uv coverage achievable with a given number of antennae. If general purpose methods can be devised which do not require redundancy then it is probably better to have no redundancy and have a better uv coverage. For the MFS case, if general methods for removing spectral effects are available it may again be best to dispense with special array geometries and instead have better uv coverage.

In contrast to special geometry arrays, for sparse arrays with no special geometry the only demand on the observing frequencies is that for the purposes of having optimum uv coverage the frequencies should be roughly evenly spread across the bandspread : a much looser constraint. This type of MFS does however demand more sophisticated subsequent processing of the data to account for spectral effects. The nature of this processing is the main subject of this thesis and will be pursued in the coming chapters.

3.1.2 The Fractional Bandspread

Alongside array geometry the most important factor affecting the size of the D_1 beam is the size of the fractional bandspread. In section(2.4.4) we discussed how we would expect the size of the D_1 beam to change with this parameter; how well do these predictions fit what we find in practice? In fig(3.3) we show north-south slices through the D_1 beam for the case of MERLIN observing a source at a declination of 50° with 5 frequencies distributed over different bandspreads ranging from 5% to 30%. The results broadly confirm our expectations. The outer sidelobes increase almost linearly with increasing fractional bandspread whereas the near-in sidelobes increase faster with fractional bandspread and so become more prominent relative to the far sidelobes. The fact

that the D_1 beam is related to the derivative of the single frequency dirty beam as $D_1 \propto \theta dD(\nu_o)/d\theta$, allows us to predict that the value of D_1 close to the centre of the beam scales as the fractional bandwidth squared (see Appendix A). This derivation though originally devised for the linear expansion, will hold to high accuracy for the logarithmic expansion because of the close similarity of the two types of beam over the moderate bandwidths used. The expected trend of D_1 with fractional bandwidth seems to be roughly confirmed by the results in fig(3.3).

The beam derivative model also predicts an absolute numerical value as well as a trend with frequency. How well does this predicted level compare with the results of simulations? In section(2.4.4) we estimated the peak contribution of a D_1 beam with a bandwidth of 25% and 5 equally spaced frequencies to be $C_2/2N$, where $C_2 = 3.9 \times 10^{-2}$ and N is the number of frequencies. This gives us an estimate for the largest D_1 sidelobe of 3.9×10^{-3} compared with the actual value of 5.0×10^{-3} ; this is in reasonable agreement given the assumptions involved. For the smaller bandwidths notice that the peak sidelobe is no longer coincident with the central sidelobes; for the bandwidths of 5%, 10% and 15% the peak sidelobes along the slices occur 2-3 arc seconds from the centre. This means that for these bandwidths our rough calculation of the peak D_1 sidelobe, based on the assumption that the peak effect occurs near the centre, will only be a rough guide. Even so it provides an estimate correct to within a factor of 2 of the correct value.

3.1.3 The Frequency Coverage

As well as the total fractional bandwidth the number of observing frequencies within the bandwidth will also affect the size of D_1 . If we start with two frequencies at the top and bottom of the band then adding extra frequencies will always reduce the size of D_1 . The uv points at the new frequencies will have weightings that are smaller than those of the frequency extrema. When we form the D_1 beam or any other beam we always normalise by the total number of uv points and so the mean amplitude of the weighting will be




Figure 3.3: D_1 Beams With Different Fractional Bandspreads. Standard MERLIN observing parameters are assumed. The slices shown are N-S and pass through the centre of the beam.

reduced, thus the level of the D_1 beam will be reduced as well. From the beam derivative model we can expect that, in the region of the beam where first derivatives dominate, the size of the contribution to the D_1 beam will scale as $C_2/N \approx \sum (\nu_i/\nu_o)^2/N$ where we sum over i from 1 to the number of frequencies N (see Appendix A). Strictly this formula applies to the linear beams but it should apply to reasonable accuracy to logarithmic beams too. If the total bandwidth is kept constant and the number of frequencies is varied from 2 frequencies to 5 then we get a reduction of a factor of about 2 in the size of the D_1 sidelobes. Increasing the number of frequencies still further produces only modest improvements in the peak D_1 level, for 9 frequencies for instance the D_1 level is 5/6 of its value for 5 frequencies. This saturation effect should be expected since as we let N become large the value of C_2/N approaches the integral of $(\nu_i/\nu_o)^2$ over the bandwidth i.e. it becomes independent of the exact number of frequencies. *It is important to note however that although increasing the number of frequencies to large values has only a modest effect in reducing the D_1 beam it is always of benefit in improving the total effective uv coverage.*

As well as the number of frequencies N within the bandwidth, another aspect of the frequency coverage is the distribution of these N frequencies. We have already seen that the choice of frequencies is especially significant for arrays of special geometry; what, however, about the case of random arrays such as MERLIN ? As far as the total effective uv coverage is concerned the best choice is to have the frequencies equally spaced in frequency. As far as size of D_1 sidelobes is concerned it is of advantage to have one frequency at or near the reference frequency, because this gives uv coverage without contributing to any of the spectral dirty beams. The combination of having equally spaced frequencies and observing at the reference frequency suggests an odd number of frequencies. Calculating the estimates of D_1 based on the formula C_2/N for both odd and even numbers of frequencies does show that an odd number gives slightly lower values for D_1 sidelobes than the general trend. Finally, note that choosing the frequencies to be symmetrical in $\ln \nu$

about the reference frequency ν_o will (in the absence of additional weighting and assuming the same numbers of observations at each frequency) ensure that the central values of all odd numbered logarithmic spectral dirty beams are zero.

It must be emphasised that the effects of the exact distribution of frequencies on the level of D_1 are quite small. In practice the choice of frequencies will probably be made on the basis of the evenness of the resulting MFS uv coverage and, most importantly, the absence of RF interference (see Chapter 6).

3.1.4 Declination Effects

We next consider how the MERLIN D_1 beam changes as the declination is varied. It is found that the D_1 sidelobe levels change with declination in a similar way to the sidelobes of the D_o beam. In fig(3.4) are shown North-South slices through the D_1 and D_o beams for several declinations. The correspondence between the sidelobe levels of the D_o and D_1 beams seen in this figure readily fits into the ‘beam derivative’ picture of the D_1 beam. Note that the two dimensional representations of the D_o and D_1 beams (not shown) indicates that at low declinations the form of D_1 beam is similar to that of the D_o beam in that it has large ‘North-South’ sidelobes, the North-South slices shown therefore give an idea of the worst levels of these sidelobes. These results are comforting in that they show that nothing peculiar happens to the D_1 beam at low declinations. In Chapter 1 it was shown that at low declinations the MERLIN overall uv coverage does not gain dramatically from MFS observations. Low declination MFS observations are likely to have a lower priority mainly because of the more modest improvement in uv coverage rather than the increased size of the D_1 beam.

3.1.5 Summary

In this section we have discussed how the sidelobe levels of the D_1 beam changes according to the various observing parameters. It has been found




Figure 3.4: North-South Slices Through the D_0 and D_1 Beams at Different Declinations. Standard MERLIN observing parameters assumed.

that only two parameters have a strong effect on the size of the sidelobes of the D_1 beam, namely, the array geometry and the fractional bandwidth. Other parameters are of lesser importance, perhaps causing variations in the level of the D_1 beam of order of a factor of 2.

Whilst analysing the effects of the various parameters it has been found that the beam derivative model first described in section(2.4.4) is very useful. This model, used in conjunction with the value of the fractional bandwidth, is sufficient to give us a reasonable estimate of the size of the D_1 beam in the case of a random array.

In the next section we will continue the investigation of the effects of spectral index on the MFS dirty map by extending the discussion to realistic models of the $I\alpha'$ distribution. For definiteness the 'standard' MFS observing parameters will be assumed, i.e. 5 frequencies approximately evenly distributed over 25% in frequency and a source at declination 50 degrees. Using the results of the above section the corresponding effects for different values of the observing parameters can then be estimated if necessary.

3.2 Spectral Effects from Extended Sources

3.2.1 First Order Spectral Effects

In this section we will consider the spectral error contributions from extended components. Given the complex structure of many real sources some approximations will be necessary in order to get a quantitative idea of first order spectral effects. We will proceed by modelling complex distributions of emission in $I\alpha'$ as collections of gaussian components. The strategy will be to consider the effect of convolving single gaussians of different sizes with the D_1 beam and then to see how the spectral contributions from a set of gaussians add up to produce the total spectral error contribution.

On convolving gaussians with the D_1 beam it is found that the resulting functions, unsurprisingly, consist of images that only contain spatial frequencies less than the reciprocal of the characteristic size of the gaussian. A set of

slices showing the morphology of the first order spectral effects for different sized gaussians is presented in fig(3.5). Note that in contrast to the D_1 beam these first order contributions have central values of the same order of size as their sidelobes. The peak amplitudes of first order spectral errors from gaussians, as functions of their changing FWHM's are described by the graphs in fig(3.6). The lower graph shows the maximum contribution as we change the FWHM whilst keeping the total $I\alpha'$ flux of the gaussian fixed at 1 Jy. The upper graph shows the peak as we keep the maximum brightness a constant. The first graph shows that, for a given total flux, the first order effect is maximum when the gaussian is as small as possible. The reduction for extended components can be understood as being due to the positive and negative sidelobes of D_1 giving rise to cancellation when convolved with extended features. In the second graph we see that, as the size of a gaussian of given peak flux increases, the above effect is sufficient almost to counteract the increase due to the increased flux in the component. Comparing the two graphs shows that knowledge of the peak brightness of a component is much more informative in predicting the expected peak first order effects than is the total flux, since the effect of the latter is much more dependent on the scale size of the gaussian. This is intuitively what one would expect.

From the lower graph in fig(3.6), and the observation that the maximum spectral effect of a given total flux in $I\alpha'$ occurs if the flux is concentrated in a point, we can immediately set an upper limit to the peak first order error of

$$D_{1,max} \int \int I\alpha'(x,y) dx dy \quad (3.1)$$

where $D_{1,max}$ is the peak D_1 sidelobe level and the integral is the total $I\alpha'$ flux in the source. If we do not know this quantity we can estimate an upper limit from the total flux and the maximum expected magnitude of α' . For most extended sources the formula will give an estimate for spectral effects which is wildly pessimistic, for point dominated sources it will however give a good estimate.

For the extended sources estimates of spectral effects based on peak brightness give a more accurate means of estimating the first order spectral effects




Figure 3.5: North-South Slices Through D_1 Convolved With Gaussians of Different FWHM. Standard observing parameters are assumed. Arbitrary vertical scale; see fig(3.6) for numerical results.





Figure 3.6: Peak Spectral Sidelobe From Gaussians of Different Diameter.

than using total flux. If there is a component of peak brightness $I\alpha'$ and characteristic size θ then we can say this contributes a peak effect of

$$D_{1,max} I\alpha' f(\theta) \quad (3.2)$$

where the scale size factor $f(\theta)$ is defined as the ratio of the peak first order effect of a gaussian component of size θ to the largest D_1 sidelobe from a point source of the same peak brightness. Note that from the definition that $f(\theta=0)=1$. In fig(3.6) $f(\theta)$ increases from unity to about 4.5 in going from $\theta = 0''$ to $\theta = 8''$.

To use the brightness approach successfully we must define how the radio source is built up from gaussian components. One useful model for radio lobes is to consider them as superpositions of gaussians of different peak $I\alpha'$ brightness and diameters θ with each gaussian being concentric to within the value of its FWHM. The resulting peak first order spectral sidelobes can then be estimated to be less than the sum of the peak spectral sidelobes from the individual gaussians.

To see this, first consider combining the effects from gaussians of quite different diameter. As we have seen from fig(3.5) the first order spectral contribution of a gaussian of size θ tends to consist of a central lobe of width comparable to θ and peak sidelobes which occur within a few times θ from the centre. If we combine the effects of two gaussians whose dimensions are quite different and which are almost concentric then the peak sidelobe due to the more compact gaussian will almost certainly lie on top of the main lobe of the larger gaussian. The main lobe of the larger gaussian will usually contain its peak sidelobe level. The total peak sidelobe will therefore be approximately the sum of the peak spectral sidelobes of each gaussian.

In the case of two gaussians of comparable size which are concentric then in the limit as they approach each other in size we again get a total peak effect equal to the the sum of the peak contribution of each. From the above two cases it is plausible to suggest

$$D_{1,max} \sum_{i=1,N} f(\theta_i)(I\alpha')_i \quad (3.3)$$

as an estimate of the peak total spectral contribution, where the sum is taken over the N component gaussians. Note that even if the adding assumption breaks down the above estimate should still be a good upper limit to the total spectral contribution. The above expression is in turn less than

$$D_{1,max} f(\theta_{max}) \sum (I\alpha')_i \quad (3.4)$$

where $f(\theta_{max})$ is the scale size factor of the largest gaussian. The above limit applies because $f(\theta)$ increases monotonically with θ . If the components are approximately concentric then $\sum (I\alpha')_i = (I\alpha')_{peak}$, the peak brightness in $I\alpha'$ and so our estimate becomes

$$D_{1,max} f(\theta_{max}) (I\alpha')_{max} \quad (3.5)$$

This is a formula based on parameters that are easily estimated.

Some semi-heuristic rules have been presented for estimating the peak spectral contributions of realistic $I\alpha'$ distributions modelled as collections of gaussian components. It is always possible to find complex distributions with which these rules fail, but they will be found to be useful in most situations. It will usually be the case that spectral contributions will be dominated by one or two components and this will make the estimation of peak spectral effects relatively easy. The significant component will often fall into one of the classes of structure that have been discussed i.e. point components, single gaussians and concentric gaussians. Using the above rules, approximate estimates of the overall spectral error contribution can then be made. In the next section estimates of spectral errors will be compared with the results of numerical simulations for a realistic source and will be found to be reasonably accurate.

3.2.2 Simulations of First Order Spectral Effects in Realistic Sources

The first order (logarithmic) spectral errors of the standard source model (see sect(1.6.3)) were simulated assuming first that the effects of the central core had been removed (see sect(6.6.3)) and then the data had been corrected for an 'overall' spectral index α_o of 0.5 (see sect(2.3.5)).

The observations were assumed to be made in the 'standard' mode of 5 frequencies spread over approximately 25% in frequency with the 8 station extended MERLIN, the declination was taken as 50 degrees. In fig(3.7) are shown slices from the bottom left hand corner to the top-right hand corner of the model intensity (diagram a) and the residual spectral index distribution α' (diagram b). In fig(1.4) of Chapter 1 are shown greyscale representations of the same intensity and spectral index distributions. The first order contribution $I\alpha' * D_1$ that results from this model is shown in greyscale in (a) of figure(3.8). Inspection of this greyscale image shows D_1 -like features centered on points in both the north and south lobes. The spacing between these ring-like features indicates that they are largely due either to unresolved or very compact components in the $I\alpha'$ map. Examples of such compact components can be seen in the I and α' distributions, e.g. the two hotspots with $\alpha' = \pm 0.1$ and I flux of 300-400mJy give compact components in the $I\alpha'$ map at the 30-40mJy level. As well as high spatial frequency features there are also low level ripples around the two lobes with characteristic lobe spacings of 2-3 arc seconds. These features will be caused by features in the $I\alpha'$ image with characteristic size 2 to 3 arcseconds.

How well can the errors found in the above simulations be estimated from simple rules? In figure(3.9) is shown a contour map of the true $I\alpha'$ distribution, whilst alongside is a sketch representation of the dominant components of the $I\alpha'$ model. Note that because the α' distribution has a much smaller range than the I distribution these $I\alpha'$ components correspond to distinct components of the I map. It is likely that, in most cases, one can identify such distinct components solely from the intensity map. With some estimates or guesses about the α' distribution it should then be possible to at least set upper limits to the $I\alpha'$ contribution of each component. In table(3.1) are listed the peak $I\alpha'$ brightness and total flux of the most prominent components identified in fig(3.9). Taking the average of the major and minor axis dimensions of each elliptical component as θ , estimates for the structure factor $f(\theta)$ can be made using the graph in fig(3.6). Assuming a peak value for $D_{1,max}$ of 5×10^{-3} an






Figure 3.7: Slices Through Model Intensity and Spectral Index Distributions.
a) Diagonal slice from bottom left to top right of the I model in fig(1.4a). b) Similar slice through the α model in fig(1.4b) after correction for an 'overall' spectral index $\alpha'=0.5$.

a)



b)




Figure 3.8: Greyscale Images of First and Second Order Spectral Errors

- a) First order errors, linear plot from $-0.5 \text{ mJy} \cdot \text{Beam}^{-1}$ to $0.5 \text{ mJy} \cdot \text{Beam}^{-1}$.
- b) Second order errors, linear plot from $-10 \text{ microJy} \cdot \text{Beam}^{-1}$ to $60 \text{ microJy} \cdot \text{Beam}^{-1}$.

estimate for the peak error contribution of each component of $D_{1,max} f(\theta) I\alpha'$ can be made, where $I\alpha'$ is the peak brightness of each component. From table(2.1) we see that the most significant spectral errors are due to the point sources such as the hotspots (components A,G and H) and the various compact components scattered through the radio lobes. These components explain the high spatial frequency structure found in fig(3.8). Also of significance are the contributions from components B and I which give rise to the lower spatial frequency structure. The amplitudes of these low spatial frequency ripples are predicted to be $110\mu\text{Jybeam}^{-1}$ (Table 3.1), which is comparable to the amplitude actually seen in figure(3.8).

The figures in Table(3.1) can be used to estimate the peak total effect due to first order spectral sidelobes. The spectral sidelobes around G and H extend to a diameter of less than 3 arcseconds (see fig(3.8)), component I on the other hand will have a diameter of at least 3 arcseconds. It is likely that the peak sidelobes of G and H will lie on top of the 'main lobe' due to component I. The total peak will therefore be of order the peak due to I plus the peak due to G and H together. Because the two compact components are close together their large sidelobes will overlap and their total peak contribution will be the sum of the peak effects due to G and H separately. The predicted peak can therefore be expected to be approximately the sum of the peak effects of G,H and I, which comes to $450\mu\text{Jybeam}^{-1}$. This compares well to the peak observed value of $442\mu\text{Jybeam}^{-1}$.

3.2.3 Simulations of Second Order Spectral Errors

The methods for removing spectral effects presented in Chapter 4 and 5 will assume that the second order effects are smaller than the thermal noise. It is therefore important to confirm that this is the case in practice. In diagram (a) of fig(3.10) is shown the second order contribution along a slice assuming the linear expansion, i.e. a second order contribution $I\alpha'(\alpha'+1)/2 * D_2$, whilst in (b) is shown an equivalent slice through the second order term of the logarithmic expansion, $I\alpha'^2/2 * D_2$. Comparison between diagrams (a) and (b)

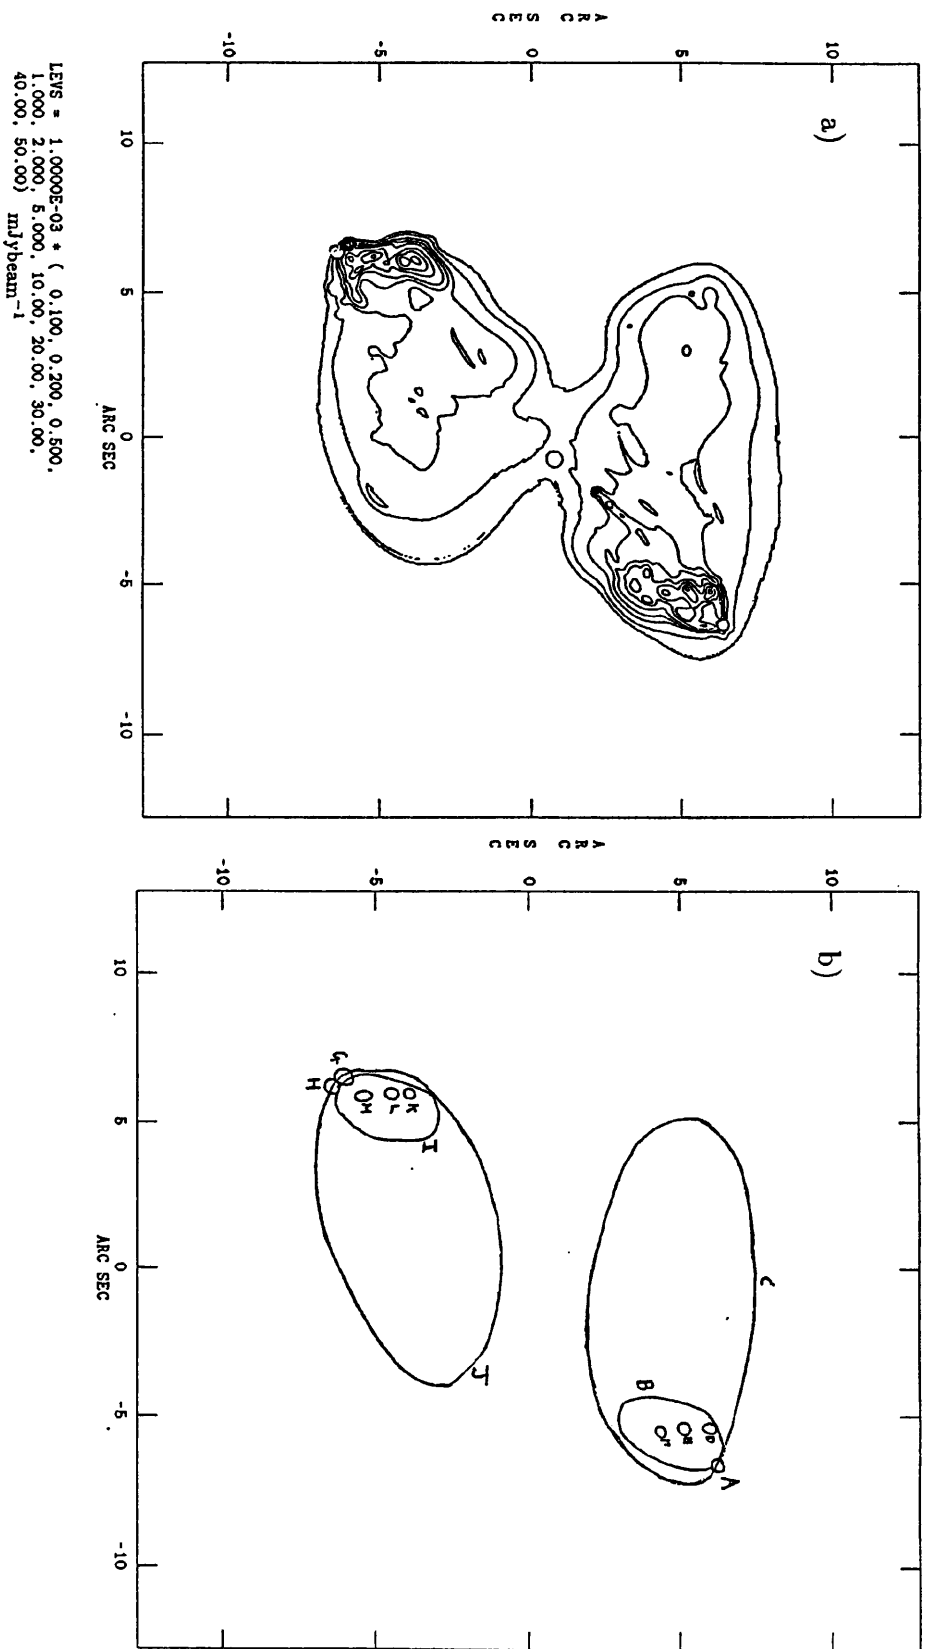


Figure 3.9: Distinct Components Within the Model $I\alpha'$ Image. a) Contour map of the $I\alpha'$ image. b) Identification of distinct components

Component	Dimensions (Beam FWHM's)	Peak Flux (mJy.Beam ⁻¹)	Mean Dimension θ (arcseconds)	$f(\theta)$	$D_{1,max} I\alpha' f(\theta)$ (mJy.Beam ⁻¹)
A	0x0	-29	0	1	0.145
B	18x9	9.5	2.0	2.1	0.110
C	42x24	0.7	5.5	3.9	0.013
D	2x2	20	0.3	1	0.100
E	2x2	20	0.3	1	0.100
F	2x2	20	0.3	1	0.100
G	0x0	34	0	1	0.170
H	0x0	-34	0	1	0.170
I	16x10	5	2.0	2.1	0.110
J	42x24	0.7	5.5	3.9	0.013
K	2x2	15	0.3	1	0.075
L	2x2	15	0.3	1	0.075
M	2x2	15	0.3	1	0.075

Table 3.1: Parameters of the components of the $I\alpha'$ image identified in fig(3.9) and their expected spectral effects.

shows the advantages of the logarithmic expansion over the linear in reducing the size of the second order contribution, particularly in those regions such as the hotspots which have relatively small α' .

The peak contribution for the logarithmic expansion along the slice shown in (b) is $50\mu\text{Jybeam}^{-1}$. At the point in the slice where this peak occurs the intensity is of order 50mJybeam^{-1} and $\alpha'=0.6$, giving $I\alpha'^2/2=9\text{mJy}$. The peak level of D_2 for this observation is approximately $1/140$ and this gives us an expected second order contribution of approximately $64\mu\text{Jybeam}^{-1}$ quite close to the observed value of $50\mu\text{Jybeam}^{-1}$. A peak effect of this size will be of order 3σ for a one day observation with extended MERLIN. Second order errors will therefore be of significance with respect to thermal noise at points near the maximum. Note that the second order effects will only be above the noise within a small region in this realistic case. Even if the noise were lower (say because we carry out MFS by a 5 day observation, with a different frequency observed each day) the second order contribution would not necessarily affect the interpretation of the source structure. In fig(3.8b) is shown a greyscale representation of the logarithmic second order contribution. The largest second order effects will lie 'on source' and so will only cause a slight bias as described in section(2.5.2); they will not significantly affect the apparent structure. As an example the $50\mu\text{Jybeam}^{-1}$ second order feature lies on top of a region of intensity 50mJybeam^{-1} and so will only cause a small fractional error of 1 part in 1000.

In addition to the errors 'on source' there are sidelobes surrounding the regions of high $I\alpha'^2$, due to the convolution of the coefficient image with the D_2 beam. These sidelobes can be clearly seen in fig(3.8b). Because of the fact that the D_2 beam is rather like a conventional dirty beam with a large main lobe and sidelobes of order 10% of the peak this image looks rather like a conventional dirty image.

The size of the above sidelobes are important in determining the so-called 'weak limit' of the second order contribution (see sect(2.5.2)). Since the D_2 sidelobe level is of order 10% and the peak contribution to the second order




Figure 3.10: Slices Comparing the Second Order Spectral Effects Under the Logarithmic and Linear Expansions. a) Slice through the second order spectral errors under the linear expansion i.e. $I\alpha'(\alpha' + 1)/2*D_2$. b) Slice through the second order errors under the logarithmic expansion i.e. through $I\alpha'^2/2*D'_2$.

errors is at $50 \mu\text{Jybeam}^{-1}$ this gives sidelobes at a level of about $5\mu\text{Jybeam}^{-1}$. For extended MERLIN this ‘weak limit’ will be well below the noise in virtually all sources of interest.

3.3 Effects of Deconvolving Images Containing Spectral Errors

3.3.1 Requirements For The Convergence of CLEAN

The processing of all aperture synthesis data from sparse arrays requires that D_o sidelobes are removed by deconvolution algorithms, such as CLEAN (Högbom 1974) or MEM (e.g. Narayan and Nityananda 1986). Such algorithms implicitly assume a dirty image of the form $I\alpha'*D_o$. In the case of MFS data we have seen that the dirty image will instead be of the form $(I*D_o)+(I\alpha'*D_1)$. What happens if we apply CLEAN or another deconvolution algorithm to this data? Is it possible that the extra term will interact with the deconvolution algorithm in such a way as to cause serious errors in the CLEAN map? This question is of interest when considering numerical estimates of the size of spectral contributions to the MFS dirty map. Any interaction of the deconvolution algorithm with the dirty map might multiply the effect of the spectral errors in the dirty map (see section(2.6.2)). In this case the stated dynamic range limits at which the contributions from first or higher order terms can be safely ignored would no longer be accurate. In this section it will be shown by simulation that no such serious interaction occurs, this lack of interaction can also be demonstrated theoretically.

Schwarz (1978) has considered the situations under which CLEAN is guaranteed to converge to a solution. One vital requirement is that the dirty image be ‘in range’ of the dirty beam, that is to say that there must be at least one image which when convolved with the dirty beam equals the dirty map. The problem we face in deconvolving the MFS dirty image with D_o in the presence of errors due to D_1 is that this requirement will not obviously hold in the case

of finite sized images and beams. Tan (1986) has shown that in nearly all cases the absence of the above condition results in a divergence in the sense that the CLEAN image will accumulate increasing amplitudes of the ‘invisible distributions’ of the dirty beam D_o as the algorithm progresses. In the alternative vector-matrix language of the deconvolution problem (as used by Tan(1986)), we will get increasing flux in the zero eigenvectors of the D_o beam matrix.

In considering whether a dirty image is ‘in range’ of the MFS dirty map the vital term in the expression for the dirty map is the second which contains a convolution with D_1 . To see whether the second term is in range the sets of eigenvalues of the eigenvectors of the two beam matrices must be compared. If there exists any non-zero eigenvector of the D_1 beam which is a zero eigenvector of D_o then the MFS dirty map will not be in range. In the case that we allow cleaning over an infinite ‘window’, so that the image I and the beam matrices D_o and D_1 are infinite, then the zero eigenvectors of both matrices will be sines and cosines with spatial frequencies corresponding to uv points not sampled by either the D_o or D_1 uv coverages. In this case all the non-zero eigenvectors of D_o are non-zero eigenvectors of D_1 and vice-versa, the dirty image is therefore in range and no divergence occurs.

In the case of finite-sized images the fact that the dirty map is still in range of the dirty beam is not as obvious as in the infinite image case. In Appendix C however it is proved mathematically that here too there should be no problems with deconvolution in the presence of spectral errors.

3.3.2 Simulations

In order to demonstrate that there are no special problems caused by cleaning first order spectral effects two sets of simulations have been carried out. In the first set, a D_1 beam is cleaned deeply with a D_o beam, in the second a MFS dirty image of the target source containing spectral errors was deeply cleaned.

The results of the first test are shown in figure(3.11). Diagram(a) shows

a slice through the D_1 beam taken diagonally through the centre, whilst diagram(b) shows a similar slice through the CLEAN map that results from deep cleaning within a window of 2 arcseconds by 2 arcseconds centered on the centre of the beam. A total of 12,000 clean components were selected with a loop gain of 0.1, which reduced the peak residual from 5mJy to 0.1mJy (see diagram(c)). As can be seen from comparing (a) and (b) no dramatic errors have been introduced by the cleaning process, the dirty and cleaned images appearing remarkably similar. Inspection of the 'error' slice (d), (i.e (b)-(a)) only confirms this conclusion. Despite the deep cleaning there appear to be no obvious effects due to large amplitude zero eigenvectors.

It might be objected that this test is unrealistic because in real situations we will be cleaning a mixture of structures caused by D_0 and D_1 sidelobes and possibly there might be some form of interaction between these different sidelobes which causes undesirable effects. To test this, model MFS data for the standard model radio source was cleaned for 30,000 iterations down to a residual of $0.25 \text{ mJybeam}^{-1}$. Slices parallel to the southern jet from the resulting image are shown in fig(3.12). In (a) is shown the intensity distribution along this slice whilst in (b) is shown the expected contribution along the slice due to $I\alpha'D_1$. The result of taking the difference between the CLEAN image and the true I model (the 'error' map) is shown in (c). This difference can be almost entirely accounted for (except for the overall negative level caused by lack of short spacings) as being due to the clean map copying the effects of spectral sidelobes into the image model. The size of the error of the clean map is reasonably well described by the size of the spectral sidelobe contribution to the dirty map, and no obvious divergence appears to have occurred.

Note that the fact that spectral sidelobes when cleaned with the D_0 beam are simply transferred largely unmodified into the CLEAN map will be useful in sect(5.2.4) when analysing a particular algorithm for removing spectral effects.




Figure 3.11: CLEANing the D_1 Beam with the D_0 Beam.

- a) N-S slice through the D_1 beam.
- b) Slice through CLEAN image, i.e. CLEAN components plus residuals.
- c) Slice through the residual image showing the area and depth of CLEANing.
- d) Difference image between a) and b)




Figure 3.12: CLEANing a Realistic Dirty Image Containing Spectral Errors.
a) Slice through intensity image parallel to southern jet.
b) First order spectral errors along the same slice.
c) Difference between a CLEAN map and a) showing effects of spectral errors shown in b).

Chapter 4

Algorithms for Removing Spectral Effects 1

4.1 Classifying the MFS Problem

4.1.1 Introduction

Chapters 2 and 3 discussed the form of the map errors caused by the fact that most sources have frequency dependent structure. In this chapter and the next algorithms will be described that are designed to remove the spectral errors caused by this frequency dependence.

Fig(4.1) shows the various algorithms that will be described in this chapter and the next and their interrelationship. This diagram will provide a guide to our discussion by classifying the algorithms in terms of their key features. Note that the diagram has two main levels of branching. The first level divides the problem according to the size of the spectral effects compared with the noise. What matters is how many orders of the spectral expansion are required to describe the sky plane to within the thermal noise. In the second level of branching the algorithms are roughly divided according to the quality of uv coverage required for their successful operation.

4.1.2 MFS and the Size of Spectral Errors

The question of the number of orders of the spectral expansions needed to describe the MFS dirty map to within the noise occupied a large fraction of Chapters 2 and 3. It was concluded that by using the ‘logarithmic expansion’ two order expansions are sufficient for virtually all complex MERLIN targets. It was also shown that there is an important subset of sources in which a

single order is sufficient, i.e. in which all spectral effects are below the noise. In these sources spectral effects can be ignored entirely and the dirty image deconvolved in a conventional manner. The analysis in Appendix C shows that we can be assured that deconvolution will not increase the effects of spectral errors. If spectral effects are negligible in the dirty image they will also be negligible in the final image.

In this chapter we will limit our discussion to the two order problem, because this form of the MFS problem is most relevant to the imaging of the largest fraction of sources. The single order problem can be considered as a special case of the two order problem where the contributions of the higher order term is small. We will not discuss in detail the MFS problem when more than two orders dominate. Only a small fraction of sources will fall into this class and even for these the second order spectral errors are likely to produce astrophysically benign map errors (sect(2.5.2)).

However should it become important in the future two approaches are possible in dealing with the three (or higher order) problem. One approach is to produce three order analogues of certain algorithms designed to solve the two order problem. Another approach is to apply two order algorithms to the three order problem to get estimates of I and $I\alpha'$ that are approximately correct. From these two images the expected contributions from higher terms can be calculated and their effects removed from the data, because all the higher order coefficient images can be generated from the estimates of the first two coefficient images. After such a correction the two order approximation will be more accurate and the two order algorithm can again be applied to the corrected data to get a new I and $I\alpha'$ estimate. This procedure can obviously be applied iteratively. Both of the above approaches to solving the higher order problem rest on applying or modifying algorithms developed to deal with the two order problem and therefore rest on an understanding of these two order algorithms.

The two order problem that must be solved is to find $I_o(x,y)$ (and at least implicitly $I\alpha'(x,y)$) given our dirty map $\overline{V}_T(x,y)$ produced from inverting the

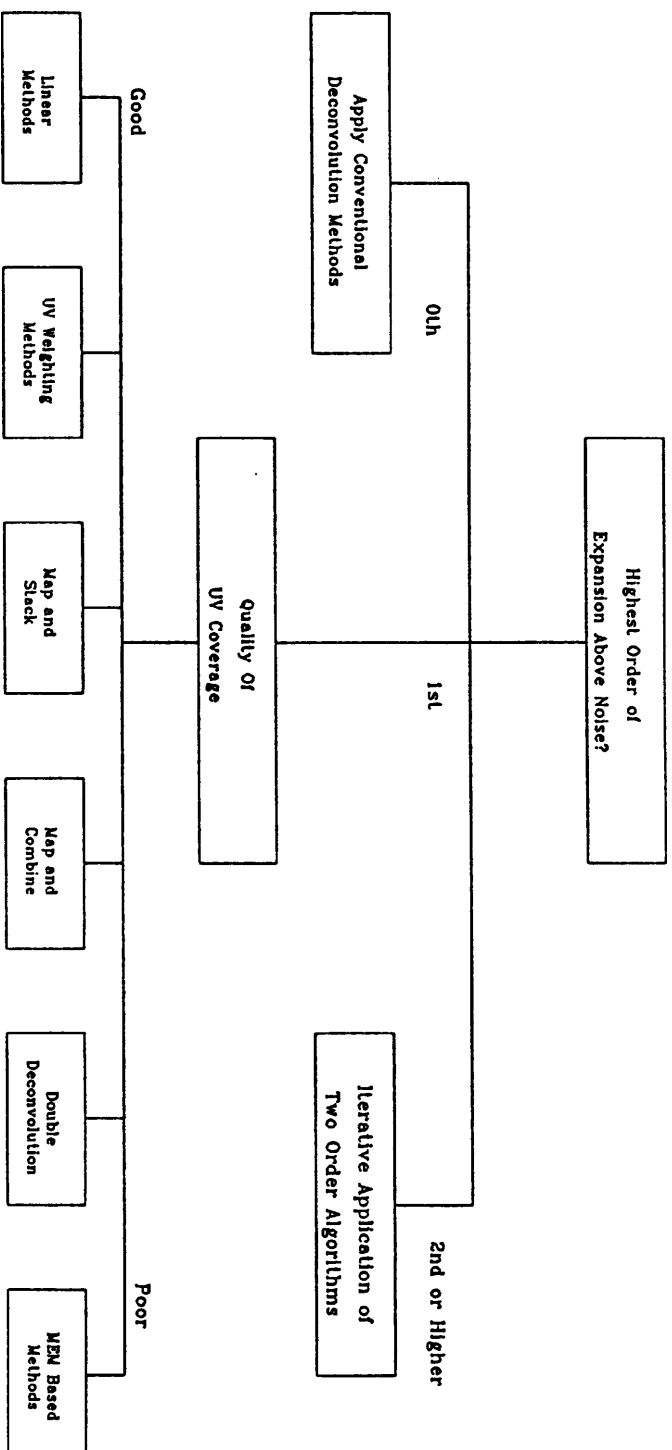


Figure 4.1: Relationships between MFS algorithms.

multifrequency data. We know from section(2.3.3) of Chapter 2 that this $\overline{V_T}$ is related to I_o and $I\alpha'$ by

$$\overline{V_T}(x, y) = (I_o * D_o) + (I\alpha' * D_1) \quad (4.1)$$

where D_o and D_1 are the logarithmic spectral dirty beams of order 0 and 1.

4.1.3 MFS and Aperture Sampling

The second level of branching in fig(4.1) depends on the degree and nature of the aperture sampling. The appropriate choice of algorithm for solving for I and $I\alpha'$ will depend on these quantities.

The nature of the MFS problem will obviously depend on the number of independent samples of the uv plane. At one extreme when the source is simple and the uv coverage dense the MFS problem is easier than when the source is complicated and the uv coverage poor. In these two circumstances it is not surprising that different algorithms should be required for the solution of the MFS problem. In the first case when the source is 'oversampled' there should, in the absence of noise be a unique solution which can be arrived at in a direct manner. The simplest possible example of such oversampling in the MFS case would be to estimate the I and $I\alpha'$ flux of a source known *a priori* to be a point. Only two observations made at different frequencies are required to uniquely determine both of these quantities. In the oversampled case the problem can be seen as one of solving a set of simultaneous equations which link the measured visibilities to the unknowns. Only direct *linear* methods are required to solve for these overdetermined systems of equations and give unique solutions. If the source contains N filled pixel areas then there will be $2N$ unknowns (i.e. I and $I\alpha'$ at each pixel). In the absence of noise only $2N$ uv plane samples are mathematically required to define a unique solution. Strictly speaking the distribution of the uv points does not matter in the absence of noise. When noise is present however the observations must be evenly distributed if the solution is not to be badly conditioned.

In the opposite limit, when the data is undersampled, different methods

must be employed. The problem posed in linear algebra will have no unique solution. As described in Chapter 1 in this situation the processing algorithm must interpolate into unsampled regions of the uv plane, this necessarily implies the use of algorithms that are *non-linear*.

Whenever the uv coverage is not oversampled (which we expect to be the norm for MERLIN MFS) the processing algorithm must contain some non-linear component. The various methods in fig(4.1) can be distinguished by how this non-linear component is implemented. Some algorithms will apply non-linear deconvolution separately to the data at each frequency, other algorithms will attempt to process the data *in toto* as much as possible. From the very definition of non-linearity this second approach would seem far superior. The only problem will be finding a way of applying such overall processing within the structure of the MFS problem. The simplest method of processing for instance consists of combining the data directly in the uv plane and then applying deconvolution to this total data set. This treats the data as a whole but is affected by spectral errors. The argument behind some of the methods which deal with spectral effects via separate processing of the data at each frequency will be that this is the most ‘natural’ way to find and deal with spectral effects given the structure of the problem.

A second aspect of the uv coverage other than just the overall density that will affect one particular algorithm, namely the ‘uv weighting methods’, is the presence or otherwise of so called ‘double sampling’ of the uv plane first described in Chapter 3. This is the close proximity in the uv plane of samples made at different frequencies on different physical baselines.

Given the different degrees of non-linearity and methods of applying it which occur within the different algorithms there will be different requirements on the quality of uv coverage needed before each can be usefully used. On the left of fig(4.1) are algorithms that are most applicable when the single-frequency or multi-frequency uv coverages are almost oversampled. On the extreme left the so-called ‘Linear methods’ can only be applied when the uv

plane is formally oversampled. In principle a direct solution to the corresponding set of linear equations could be attempted, in practice a variety of linear methods could be used or adapted (Andrews and Hunt 1977). Next we have ‘uv weighting’ (see sect(4.4)) methods which use entirely linear processing to remove spectral effects and only after this attempt non-linear deconvolution of the corrected data. These methods require as a prerequisite exact or approximately ‘double sampled’ data. The ‘Map and Combine’ algorithm (see sect(4.2)) uses non-linearity in forming separate images at the different frequencies but then removes frequency-dependent structure by linear manipulation of the resulting images. The ‘Map and Stack’ (see sect(4.3)) algorithm is similar in that it estimates the $I\alpha'$ distribution from the separate deconvolutions at each frequency, it differs in that the subsequent correction is carried out in the uv plane and one final non-linear deconvolution of the whole corrected data is carried out. We will see that the combination of non-linear and linear processing used in the above two algorithms will not generally be very effective with the type of MFS uv coverage achievable with MERLIN. Finally on the far right of fig(4.1) we have ‘Double Deconvolution’(see Chapter 5) and ‘MEM based methods’ (see sect(4.5)) which use non-linear deconvolution algorithms such as CLEAN or MEM applied to the data ‘in toto’ to estimate both the I and $I\alpha'$ distributions.

It will be found in the course of this chapter and the next that the best algorithms for MERLIN MFS will be those that lie at the right hand side of fig(4.1) and in particular the so-called ‘Double Deconvolution’ algorithm. Before describing this method in Chapter 5, it will be necessary to consider other simpler algorithms. Much of this chapter will deal with describing these other methods. Hopefully the ideas introduced whilst considering these algorithms will make the nature of the MFS problem clearer.

4.1.4 The Oversampled Limit

Can we derive an order of magnitude estimate of the condition for oversampling? Wilkinson(1983) gave order of magnitude estimates for the maximum

size of source that can be reliably mapped by conventional methods with a 6 antennae MERLIN. These limits have been found to be justified empirically by experience with MERLIN mapping. The limits are based on the size of holes in the uv plane. Such an approach can be extended to the MFS case.

Say that we have n physical baselines and f frequencies, this will give us (nf) observations along any radius vector in the uv plane. If we make the simplifying assumption that these uv observations are equally spaced in radius out to the maximum uv distance u_{max} we get a radial separation in uv distance between adjacent uv tracks of (u_{max}/nf) . We now assume that we sample along each baseline at an interval of $(1/\theta_{max})$. If we attempt to reconstruct the image by fitting data which sample each uv track much more often than this we find that the result will be very ill-conditioned with respect to noise. Assuming the above sampling interval the number of uv points will be equal to the total length of the uv tracks divided by $(1/\theta_{max})$, the total length being approximately the total area divided by the area per unit length of uv track.

$$\frac{\pi u_{max}^2}{u_{max}/nf} = \pi u_{max} n f \quad (4.2)$$

Dividing this by $(1/\theta_{max})$ gives the total number of observations. If we include conjugate points the number of uv points therefore amounts to $2\pi u_{max} n f \theta_{max}$. If the image is b beams by b beams then $\theta_{max} = b(1/u_{max})$. If we have two pixels per beam and two unknowns per pixel (I and $I\alpha'$) then the requirement for oversampling will be approximately.

$$2(2b)^2 \leq 2\pi n f b \quad (4.3)$$

or

$$b \leq (\pi/4) n f \quad (4.4)$$

For an 8-station extended MERLIN with $n = 28$ and $f=5$ this gives $b=110$ which, if we have 2 pixels per beam, amounts to a region 220 pixels square. This analysis is only approximate and ignores the fact that we must self-calibrate phase and amplitude. The telescope-based phase and amplitude errors are additional unknowns which the data must constrain, their necessary

solution will result in a smaller size of field for which oversampling will apply. For an 8 telescope array these extra overheads will not however greatly reduce the maximum critical field size.

Most of the complex sources which extended MERLIN will attempt to image will require image sizes between 256 and 512 pixels square. It therefore appears that such sources will be moderately undersampled by the MFS uv coverage. This implies that some non-linear processing will be required. The undersampling due to single frequency uv coverages will be significantly greater. This means that methods based on the separate deconvolution of the data at each frequency are likely to introduce substantial reconstruction errors. We shall see that the large size of these single frequency reconstruction errors makes these methods unattractive.

4.2 Map and Combine

The simplest methods that attempt to deal with the effects of frequency dependent structure in MFS are based on imaging the data separately at each of the frequencies.

Any algorithm that removes spectral effects must either estimate the spectral index distribution and then remove the effects of this distribution from all the data or find some process that manipulates the data so as to cancel out the spectral effects. It can be argued that the most natural domain in which to do either of the above is in the sky plane domain, from separate images made at the different frequencies.

We shall briefly consider the sky-plane-based algorithms in this section and the next. We will see that although they are intuitively simple they do not generally give dramatic improvements in image quality. The reason for this is that the non-linear part of the processing occurs in the separate deconvolution of the data at the different frequencies. As explained earlier this piecemeal deconvolution is unlikely to be the best way of processing the data. In creating the individual channel maps single-frequency reconstruction errors occur due

to the poor uv coverage, any algorithm that relies on these images to solve the spectral error problem must to some extent incorporate the single frequency reconstruction errors into the final MFS image. Algorithms which use the whole data to estimate the spectral effects rather than going through the single frequency reconstructions will give superior performance as we shall see in later sections and in chapter 5.

4.2.1 The Algorithm

The simplest procedure for producing an improved composite image from a MFS data set would be to image and deconvolve the data at each frequency separately and then add all the resulting images together. The resulting image should then be better than any of the individual channel maps since we could expect the size of the reconstruction errors to be reduced by a statistical factor that depends on the degree of correlation between reconstruction errors at the different frequencies.

Because the spectrum at each point on the radio source is a power law simple adding is not strictly appropriate. The correct way to combine the channel maps is instead to take the *geometric mean* of all the maps. This simple algorithmn of separate deconvolutions and combination of images in the sky plane we dignify with the name ‘Map and Combine’.

Consider making images at each of the frequencies ν_i , $I(\nu_i)$ so that

$$I(\nu_i) = I(\nu_o)(\nu_i/\nu_o)^{-\alpha'(x,y)} + E(x, y, \nu_i) \quad (4.5)$$

where $E(x, y, \nu_i)$ is the single-frequency reconstruction error at frequency ν_i . A composite image can be formed by taking these images and forming a geometric mean I_{geom} . Ignoring the reconstruction errors the resulting image will be

$$I_{geom} = \left(\prod_{i=1}^N I(\nu_i) \right)^{1/N} = I(\nu_o) \left[\prod_{i=1}^N (\nu_i/\nu_o)^{-\alpha'(x,y)} \right]^{1/N} \quad (4.6)$$

and so $I_{geom} = I(\nu_o)$ where ν_o is such that

$$\prod_{i=1}^N (\nu_i/\nu_o) = 1 \quad (4.7)$$

i.e. where ν_o is the geometric mean of all the observing frequencies.

4.2.2 Effects of Single Frequency Reconstruction Errors

Next we follow the propagation of the effects of single-frequency reconstruction errors $E(x,y,\nu_i)$ into the final image. Take the case of combining images made at two frequencies ν_1 and ν_2 , $I(\nu_1)$ and $I(\nu_2)$, each with its error $E(\nu_i)$. Substituting equ(4.5) into equa(4.6) we obtain

$$I_{geom} = (I(\nu_1) I(\nu_2))^{1/2} \quad (4.8)$$

$$= I_o[1 + (E_1/I_o)(\nu_1/\nu_o)^{\alpha'(x,y)} + (E_2/I_o)(\nu_2/\nu_o)^{\alpha'(x,y)} + E_1 E_2 / I_o^2]^{1/2} \quad (4.9)$$

When the sizes of E are small compared to the intensity image I i.e. when the fractional errors are small, the product term $E_1 E_2$ can be neglected and the quantity raised to the power one half expanded as a power series. The final error is approximately additive and equal to

$$(E_1/2)(\nu_1/\nu_o)^{\alpha'} + (E_2/2)(\nu_2/\nu_o)^{\alpha'} \quad (4.10)$$

In the general case of N different images then it can be shown that the approximate error will be

$$E_{geom} = (1/N) \sum E_i (\nu_i/\nu_o)^{\alpha'} \quad (4.11)$$

this is reasonable since, in the limit as the fractional bandspread becomes small the geometrical mean gives the same result as arithmetic averaging. The total error should therefore be approximately the mean of the errors in each of the channel maps.

The extent to which sky plane combination reduces the size of reconstruction errors will depend on the degree of correlation of the reconstruction errors. If they are highly correlated then there will be little reduction. If they are uncorrelated there will be a statistical reduction so that that the level of errors in a N frequency observation are approximately reduced by the factor $N^{-1/2}$. Only if the errors are anti-correlated is there a large improvement in fidelity.

4.2.3 Correlation Between Single Frequency Reconstruction Errors

The single frequency reconstruction errors are caused by CLEAN badly interpolating the model visibility into gaps in the uv plane. Because the uv coverage is quite different at the different frequencies (or else there is no point in MFS) we can expect that the reconstruction errors will not be highly correlated. If there is a significant bandwidth so that the gaps in the total uv coverage are well filled in then we can expect that the uv coverages at the extremum frequencies will be almost orthogonal, i.e. the gaps in one coverage will correspond to well sampled tracks at the other frequency.

Although some work has been done on the nature of CLEAN reconstruction errors, it is difficult to say much in detail about the nature and degree of correlation of these errors in any particular set of single-frequency reconstructions.

It can be expected that they will depend in some complex way on the source structure as well as on the uv coverage. Various authors have attempted to explain the mechanisms that cause CLEAN errors which in general appear as ripple artifacts in the images (see fig(4.3a)). Schwarz(1984) has argued that the iterative removal of dirty beams with negative sidelobes produces the ripples. Cornwell(1983) and Tan(1986) have shown however that the ripples can and do occur even if the dirty beam has no negative sidelobes, the crucial properties of the dirty beam instead lie in its eigenvalue spectrum. Tan(1986) also shows to some extent how the model is interpolated into gaps in the uv plane and how bad interpolation causes the ripples; the ripples are generated both as transitory phenomena and in the final solution when the uv coverage is poor.

4.2.4 Simulations of Single Frequency Reconstruction Errors

Despite theoretical work the only way to determine the nature and correlation of the single-frequency reconstruction errors is by simulation. Figure(4.3a) and

figure(4.4) show the results of such simulations of single frequency deconvolutions presented as a series of ‘error’ maps, i.e. as the difference between the reconstructions and the true source at each of the frequencies. The model source used was an earlier version of the target model first described in Chapter 1. The model used here is essentially the same as that described in Chapter 1 except for a reduction in the amount of extended emission (see fig(4.2)).

In fig(4.3a) can be seen typical examples of different types of CLEAN artifact obtained on deconvolving simulated data from this model source. First there is a low amplitude, high spatial frequency, ripple that runs perpendicular to the radio source axis; this error can be seen most readily ‘off-source’. In the uv plane this error would correspond to a point-like error of small amplitude within a gap in the uv coverage. Secondly, in the diffuse lobes the structure appears ‘mottled.’ This could be thought of as being the superposition of many ripple-like errors of different orientations, all being non-zero only within the area of the radio lobes. Lastly there appear to be low spatial frequency ripples that are confined to the very brightest parts of the source close to the two sets of hotspots.

To see the degree to which reconstruction errors at the different frequencies are correlated examine the error maps presented in figs(4.3a) and fig(4.4). The amplitudes of the reconstruction errors are found to be different at the different frequencies; this is most easily seen in the case of the low spatial frequency ripples. Close inspection of the wavelengths of the different ripples confirms, however, that to a first approximation they do scale inversely with observing frequency as would be expected if they are caused by errors in particular gaps in the uv plane.

4.2.5 Testing Map and Combine

Figure(4.3b) shows the error map that corresponds to the image formed by taking the geometric mean of the five single frequency images. The resulting spectral errors are reduced by a factor of 2 or 3 over the typical level of reconstruction error seen in the single-frequency images. This level of improvement

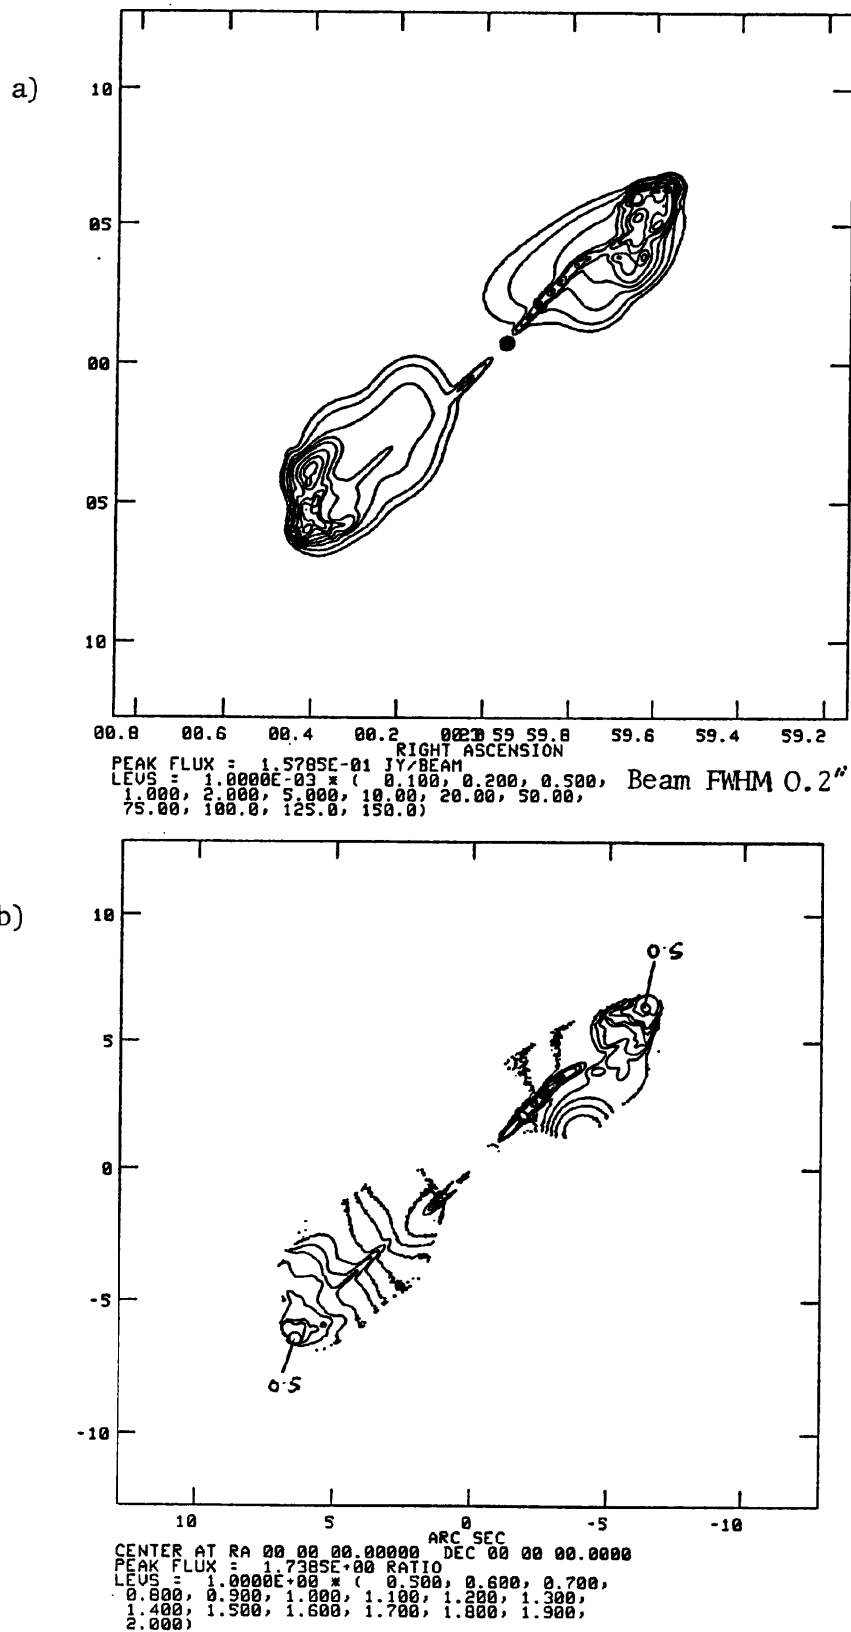



Figure 4.2: The model source used for simulations in this chapter a) The intensity distribution at the reference frequency. b) The spectral index distribution.

a)



b)




Figure 4.3a: Error between a single-frequency reconstruction at 1540 MHz and the true image. Linear plot from $-10 \text{ mJy.Beam}^{-1}$ to 10 mJy.Beam^{-1} . Restoring beam is $0.15''$

b: Error map for the result of the "Map and Combine" algorithm. Scale and beam as above.

(NB. The central dark spot is an artifact)




Figure 4.4: Error between single-frequency reconstructions at 1350, 1415, 1615 and 1730 MHz and the true image. Linear plots from $-10 \text{ mJy.Beam}^{-1}$ to 10 mJy.Beam^{-1} . Restoring beam $0.15''$.

is in line with the statistical reduction expected if the errors are uncorrelated. The modest improvement in fidelity achieved will mean that this MFS processing method is only of limited applicability. Its poor performance is to be expected because the method does not properly utilise the non-linearity of deconvolution.

It must be possible to do much better than ‘Map and Combine’. Imagine that the MFS uv coverage was sufficiently dense that the uv plane were ‘over-sampled’ but the single frequency uv coverage was badly undersampled. In this situation the I and α distributions could be uniquely derived in the absence of noise. ‘Map and Combine’ in contrast would still only give a statistical reduction upon the large single frequency reconstruction errors.

4.3 Map and Stack

4.3.1 The Algorithm

The ‘Map and Combine’ method avoided the problems of ‘spectral sidelobes’ by deconvolving each frequency channel separately rather than first combining the uv data and then deconvolving. The resulting images still contained the effects of frequency dependent structure however. Allbeit in a simple way, the intensity at each pixel simply following a power-law spectrum. By forming the geometrical mean of the separte channel maps these simple frequency dependent effects could therefore be easily removed. The ‘Map and Combine’ algorithm was however found to be inadequate because in the end the MFS intensity image relied *only* on the separate deconvolutions, which have poor uv coverage and hence low fidelity.

It might be wondered whether an algorithm could be found that estimates the spectral index distribution in the sky plane from the single frequency images in the obvious way yet allows the appropriately corrected data to participate in a single deconvolution step. The uv coverage of the final deconvolution to create the final I image should then be excellent and so the resulting map fidelity high. Such a method could hopefully combine the spectral estimation

advantages of sky-plane based methods with the advantages of *in toto* deconvolution of uv plane combined data. In this section an algorithm of this type is described called ‘Map and Stack’.

‘Map and Stack’ works by taking the single channel images, notionally ‘stacking’ these images into a data cube and then interpolating a best fit power law spectrum at each pixel. From the estimated spectral index distribution $\alpha(x,y)$ and intensity distribution at the reference frequency $I_o(x,y)$ corrections can be calculated to apply to the MFS data so as to remove the spectral effects. The quantities $I\alpha'$, $I\alpha'^2$ etc could be calculated and their effects removed from the MFS data according to the logarithmic or linear expansions of sections(2.6.1) and (2.3.3). The corrected data could then be inverted and deconvolved with the multi-frequency dirty beam.

Different levels of accuracy can be considered for the spectral interpolation, estimates of the gradient of the function $\ln[I(\nu_i)]$ against $\ln(\nu_i)$ could be made, and this would give a spectral index at each pixel. Often however, it is sufficient to assume that the spectrum is linear across the bandspread. The gradient of intensity against frequency measured would then directly equal $I\alpha'$, the only significant coefficient image that must be removed from the data if the two order approximation is assumed. Even in the case that the more accurate model is used the resulting estimate for $I\alpha'$ that emerges should be close to that which comes from the linear model and so in the following analysis this approximation is assumed.

4.3.2 Effects of Single Frequency Reconstruction Errors

The effect of using single frequency reconstructions, which contain errors, to correct out spectral effects must be calculated. Clearly some part of these errors will be fed back into the MFS data and these will finally limit the fidelity that can be achieved by the algorithm.

It is assumed that the uv data sets at each frequency are first rescaled to correct for an overall spectral index (see section(2.3.2)). The data at each of the frequencies are then separately deconvolved. From the logarithmic

expansion the estimate of the intensity map at frequency ν_i will be equal to

$$I'(\nu_i) = I_o + I\alpha' \ln(\nu_o/\nu_i) \dots + E(\nu_i) \quad (4.12)$$

where E is the reconstruction error. Assuming that the higher terms of the expansion can be neglected then an estimate of the $I\alpha'$ distribution can be made using the channel maps thus

$$I\alpha'_{ms} = \frac{\sum \ln(\nu_o/\nu_i) I'(\nu_i)}{\sum (\ln(\nu_o/\nu_i))^2} \quad (4.13)$$

where the summations are over i , so that

$$I\alpha'_{ms} = I\alpha' + \frac{\sum \ln(\nu_o/\nu_i) E(\nu_i)}{\sum (\ln(\nu_o/\nu_i))^2} \quad (4.14)$$

On taking this estimate of $I\alpha'$, removing its effect from the MFS data and then inverting the corrected data we obtain for the dirty map

$$I_o * D_o + \left[\frac{\sum \ln(\nu_o/\nu_i) \cdot E(\nu_i)}{\sum (\ln(\nu_o/\nu_i))^2} \right] * D_1 \quad (4.15)$$

The second term will contain the limiting error due to feeding back reconstruction errors.

4.3.3 Effectiveness of the Algorithm

In this section we consider the effectiveness of the algorithm in reducing spectral effects in two simple cases, first when an image is dominated by extended emission containing extensive ripple errors and secondly when the image is point dominated.

In the first case we will assume for simplicity that all the single frequency images are perfect (i.e. $E(\nu_k)=0$) except one at ν_i which contains a simple ripple with peak amplitude F Jy.beam $^{-1}$ extending over an area θ by θ . In 'natural' units the amplitude of the ripple amounts to of order $F u_{max}^2$ Jy.Sr $^{-1}$ where u_{max} is approximately the longest baseline in the array measured in wavelengths. The FT of this reconstruction error is shown in fig(4.5), it can be seen that the error lies in a gap in the uv coverage. The integral over area of the Fourier transform of this error will equal the size of the amplitude of

the sky plane ripple. The area of the error in the uv plane will be θ^{-2} which means that the peak in the uv plane will be

$$F(u_{max}\theta)^2 = F(\theta/\theta_{res})^2 \quad (4.16)$$

where θ_{res} is the resolution of the array. It can be seen that the size of the peak depends on the number of beam areas over which the ripple error extends.

We next examine the effect that this reconstruction error at frequency ν_i introduces into the dirty map via equ(4.15). We will assume that the reconstruction error is sampled by the uv track due to observations at frequency ν_j as shown by the dotted track in fig(4.5). The resulting sky plane error will be ripple-like and from equ(4.15) will have an amplitude

$$(N/N_o)F(\theta/\theta_{res})^2 \cdot \frac{\ln(\nu_o/\nu_i)}{\sum(\ln(\nu_o/\nu_i))^2} \ln(\nu_o/\nu_j) \quad (4.17)$$

Where N_o is the total number of visibilities and N the number of uv points at which the reconstruction error is sampled. The ratio (N/N_o) will equal the fraction of the total length of uv track which samples the reconstruction error. Since the error extends over $(1/\theta)$ then using the expression for the total length of uv track given in equ(4.2) we obtain

$$(N/N_o) = (\theta_{res}/\theta) \frac{1}{\pi n f} \quad (4.18)$$

so that the amplitude of the ripples after processing is

$$F \frac{(\theta/\theta_{res}) \ln(\nu_o/\nu_i) \ln(\nu_o/\nu_j)}{\pi n f \sum(\ln(\nu_o/\nu_i))^2} \quad (4.19)$$

For extended MERLIN with $n=28$, $f=5$, a bandspread of 20% and the frequency ν_i at one end of this band and ν_j at the other end of the band then it can be shown that provided $\theta \leq 870\theta_o$ the amplitude of the ripples will be decreased by the 'Map and Stack' processing.

The above analysis is specific to one particular set of circumstances, it considers the response of 'Map and Stack' to extended ripple errors. It shows that the level of ripple error can be reduced, the critical problem with the algorithm is however likely to be that such single frequency errors

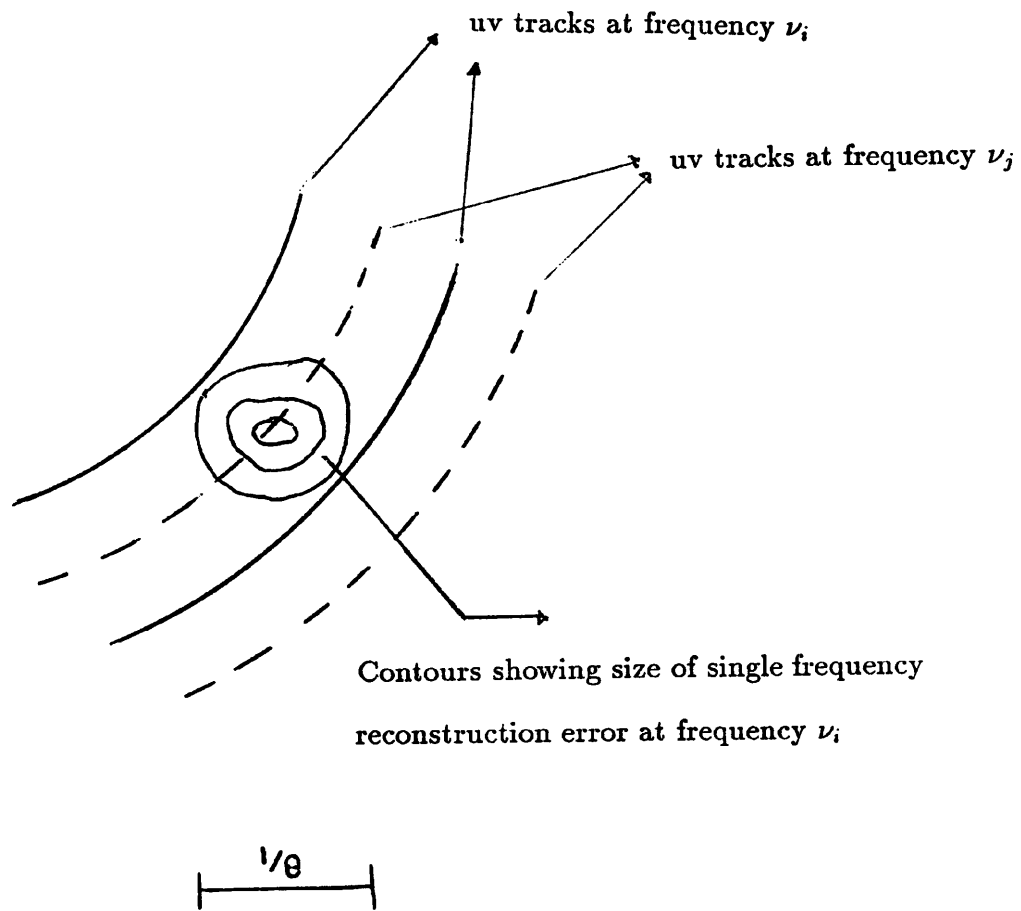


Figure 4.5: Diagram showing the uv sampling of single frequency reconstruction errors in the 'Map and Stack' algorithm

are likely to be converted to D_1 like errors. This effect will be most obvious when we must represent the $I\alpha'$ image over compact regions.

Consider a hotspot with a given $I\alpha'$, and consider estimating this unknown quantity from two single frequency images separated in frequency by $\Delta\Omega$. If the two channel maps contain reconstruction errors of size E then the resulting estimate will be in error by of order

$$(\nu_o/\Delta\Omega)(E/I).I \quad (4.20)$$

In order to be useful we will require that the effective error in the estimated spectral index $(\nu_o/\Delta\Omega)(E/I) \leq \alpha'$ i.e. it must be less than the residual spectral error. Because we can often estimate and remove ‘overall’ spectral indices at hotspots to better than 0.1 then we can assume that the residual spectral index is $\alpha' = 0.1$. For a bandspread of 20%, to improve upon this estimate using ‘Map and Stack’ would require the fractional errors at the hotspots in the single frequency maps to be accurate to better than 2%. For the case of more than two frequencies the criterion would be less severe, but not dramatically so. The figure of a 2% fractional error is therefore a reasonable first estimate of the required single frequency fidelity.

The size of single frequency fractional reconstruction errors will depend in some complex way on the relative complexity of the source. Predicting the size of these errors is not generally possible. From the simulations in section(4.2.2), however, it appears that for a source as complex as the ‘target source’ errors of order 1% can occur at the hotspots. If this is typical, then for the complex sources which should be the targets for MFS imaging, ‘Map and Stack’ at best appears of only marginal utility, even if these sources are point-dominated.

4.4 Data Weighting Methods

4.4.1 Double-Sampled Data

In Chapter 3 (sect(3.1.1)) we briefly considered the subject of special array geometries designed so that the uv tracks from different physical baselines

at different frequencies overlapped. With this geometry there would be two measurements at different frequencies at each sampled point in the uv plane i.e. ‘double-sampled data’. Given two observations it will then be possible to solve for the two unknowns at each uv point viz $\overline{I_o}$ and $\overline{I\alpha'}$. With this geometry we can directly identify and eliminate the uv plane effects of spectral index by removing the effects of $I\alpha'$ at each uv point.

If at a given uv point we have two measurements, one made at the frequency ν_1 , $V(\nu_1) = \overline{I(\nu_1)}$ the other at ν_2 , $V(\nu_2) = \overline{I(\nu_2)}$ then using the two order approximation we have

$$V(\nu_1) = \overline{I_o} + \ln(\nu_o/\nu_1)\overline{I\alpha'} \quad (4.21)$$

$$V(\nu_2) = \overline{I_o} + \ln(\nu_o/\nu_2)\overline{I\alpha'} \quad (4.22)$$

By weighting the two observations appropriately and then adding them we obtain V_{comb} .

$$V_{comb} = V(\nu_1)\ln(\nu_o/\nu_2) + V(\nu_2)\ln(\nu_o/\nu_1) \quad (4.23)$$

$$= [\ln(\nu_o/\nu_1) + \ln(\nu_o/\nu_2)]\overline{I_o} \quad (4.24)$$

we can therefore recover the true reference frequency visibility from V_{comb} and so eliminate the effects of spectral index.

Methods of this type were first discussed in Cornwell (1984), who analysed the effect of spectral index in terms of a power law expansion (based on the linear model rather than the logarithmic) and then proposed a method based implicitly on the above ‘double sampling’ to reduce the two order problem to a conventional mapping problem by eliminating spectral effects. The Cornwell algorithm is, in effect, a generalization of the ‘double sampling’ methods in that it can be applied to data having any uv coverage, but we will see that in its purest form it only gives useful results for those arrays in which a significant number of points are exactly ‘double sampled’, since other isolated uv points will contribute nothing to the solution. It is however useful to discuss it here because it is easily adapted into an algorithm which can deal with data which is only approximately ‘double sampled’; i.e. in which uv samples at different

frequencies lie within $(1/\theta)$, where θ is the characteristic angular size of the source.

4.4.2 The Cornwell Method

The two order approximation gives a dirty map \overline{V}_T

$$\overline{V}_T = d(o) = (I * D_o) + (I\alpha' * D_1). \quad (4.25)$$

The corresponding equation in the uv plane is

$$V_T = \overline{I} \cdot \overline{D}_o + \overline{I\alpha'} \cdot \overline{D}_1 = \overline{d(o)} \quad (4.26)$$

where \overline{D}_o and \overline{D}_1 are non-zero only on sampled points and the non-zero value of \overline{D}_1 depends on the observing frequency at which a given uv point was sampled. If the observing frequency is ν_i then this weight is equal to $\ln(\nu_o/\nu_i)$. In the limit that we have two uv measurements made at different frequencies ν_1 and ν_2 at exactly the same point in uv plane then the value of

$$\overline{D}_1 = \ln(\nu_o/\nu_1) + \ln(\nu_o/\nu_2) \quad (4.27)$$

at this uv point. Now we form a new dirty map $d(1)$ by first multiplying each observation at frequency ν_i by $\ln(\nu_o/\nu_i)$ before combining the weighted data in the uv plane and then inverting, then the FT of $d(1)$ becomes

$$\overline{d(1)} = \overline{I} \cdot \overline{D}_1 + \overline{I\alpha'} \cdot \overline{D}_2 \quad (4.28)$$

Where \overline{D}_1 and \overline{D}_2 have the normal definitions for singly sampled points but for overlapping points \overline{D}_1 has the form of equation(4.27) and \overline{D}_2 the form.

$$\overline{D}_2 = [\ln(\nu_o/\nu_1)]^2 + [\ln(\nu_o/\nu_2)]^2. \quad (4.29)$$

On inverting equ(4.28) we therefore obtain

$$d(1) = (I * D_1) + (I\alpha' * D_2) \quad (4.30)$$

where D_1 and D_2 are the Fourier transforms of \overline{D}_1 and \overline{D}_2 . We now try to find some combination of equ(4.25) and equ(4.30) that eliminates the effects of spectral index. Consider forming the quantity $[d(o) * D_2] - [d(1) * D_1] = d$

$$d = I * [(D_o * D_2) - (D_1 * D_1)] + I\alpha' * [(D_1 * D_2) - (D_2 * D_1)] \quad (4.31)$$

The convolution operation is commutative and so the quantity convolved with $I\alpha'$ is necessarily zero everywhere, therefore spectral effects are eliminated and we may write

$$d = I * D_{eff} \quad (4.32)$$

where

$$D_{eff} = (D_o * D_2) - (D_1 * D_1) \quad (4.33)$$

We can now proceed to reconstruct I by using a standard deconvolution algorithm using D_{eff} as the dirty beam.

The Fourier transform of D_{eff} is

$$\overline{D_{eff}} = \overline{D_o} \cdot \overline{D_2} + \overline{D_1} \cdot \overline{D_1} \quad (4.34)$$

For a point which is singly-sampled this quantity will be zero and so isolated points will contribute nothing to the dirty beam and hence nothing to the dirty map. If instead we consider points that are 'double sampled' then we obtain on substituting equ(4.27) and equ(4.29) into equ(4.33)

$$\overline{D_{eff}} = 2[\ln(\nu_o/\nu_1)]^2 + 2[\ln(\nu_o/\nu_2)]^2 - [\ln(\nu_o/\nu_1) + \ln(\nu_o/\nu_1)]^2 \quad (4.35)$$

$$= [\ln(\nu_o/\nu_1) - \ln(\nu_o/\nu_2)]^2 \quad (4.36)$$

which is non-zero only if $\nu_1 \neq \nu_2$. It appears that only uv points that are sampled twice at different frequencies will contribute to the MFS image.

Note that from equ(4.36) that the uv values which correspond to the effective dirty beam will depend on how close together in frequency the overlapping frequencies are, since this can vary over the uv plane we can get a very non-uniform distribution of sampling. To correct for this we can multiply the weight of the dirty beam at each 'double sampled' point by an appropriate amount \overline{w} , where

$$\overline{w} = \frac{1}{[\ln(\nu_o/\nu_1)]^2 - [\ln(\nu_o/\nu_2)]^2} \quad (4.37)$$

on 'double sampled' points and

$$\overline{w} = 0 \quad (4.38)$$

elsewhere.

If this additional weighting is applied then the weighted data will be equal to \bar{I} on the sampled uv points and so this form of the Cornwell algorithm gives an exactly equivalent result to the original ‘double sampling’ algorithm presented at the beginning of this section.

4.4.3 The Gridded Cornwell Algorithm

The major shortcoming of the above method appears to be that, in order to count in the deconvolution process, uv points must exactly overlap. This requirement seems overly stringent, because for sources of diameter θ we can expect that their Fourier transforms do not vary significantly over distances in the uv plane of less than $1/\theta$. If two sampling points at two different frequencies are separated by this distance or less they are effectively sampling values of the visibility that are highly correlated.

One way to introduce *a priori* knowledge about this correlation and so allow more uv points to contribute in weighting methods is to use the gridding operation that forms part of all the common imaging algorithms (Cornwell, private communication). When we form inversions of uv data we normally first convolve our sampled uv measurements onto a regular grid of spacing approximately $(1/\theta)$. This operation is known as ‘gridding’ and is performed so we can use the advantages of the Fast Fourier Transform (FFT) algorithm.

To use the gridding mechanism we take equ(4.25) and equ(4.30) and then convolve them with D_1 and D_2 respectively. Instead of doing an exact convolution we instead carry out what we will term a ‘gridded convolution’; that is we sample both quantities in the uv plane onto a regular grid, multiply these gridded values together, grid point by grid point, and then Fourier transform back to the sky plane to get our final result. If we apply the process to each of equations(4.25) and (4.30) and then subtract the results we obtain the new ‘dirty map’ d.

$$d = D_2 \bar{*} (D_o * I) - D_1 \bar{*} (D_1 * I) + D_2 \bar{*} (D_1 * I\alpha') - D_1 \bar{*} (D_2 * I\alpha') \quad (4.39)$$

where the symbol $\bar{*}$ denotes gridded convolution and $*$ denotes normal or exact

convolution. It can be shown (see Appendix D) that in general $D_2 \bar{*} (D_1 * I\alpha') \neq D_1 \bar{*} (D_2 * I\alpha')$ and so, unlike in its exact counterpart in this algorithm, the spectral effects are not eliminated entirely. However spectral effects are very substantially reduced within a region of diameter θ . An unfortunate side-effect of the gridded algorithm is that extra error terms are introduced because of the approximate nature of the convolution. Both the new error terms and the residual spectral effects will mean the dirty map does ^{not} exactly fit the form of a true image convolved with a dirty beam. This will mean by definition that the dirty map is not ‘in range’ of the dirty beam. In many cases however these error effects will be smaller than the thermal noise so that we can ignore them. The gridded algorithm will then be useful because it reduces spectral effects to below the noise whilst relaxing the exact ‘double sampling’ requirement. Because of the gridding operation two samples need only lie within the same grid cell. This means that many more points can contribute to the effective dirty beam D_{eff} , than would be the case for the exact algorithm.

4.4.4 Simulations of the Gridded Weighting Algorithm

Given the relaxation of the ‘double-sampling’ requirement that we obtain in going from the exact algorithm to the gridded algorithm we might ask whether the weighting method is now useful for our expected MFS system on MERLIN. In order to enter into the ‘effective uv coverage’ which is the FT of the effective dirty beam a uv point must lie within a distance of order $1/\theta$ in the uv plane to another uv measurement observed at a different frequency. If it does not then the point will be given a very small weighting and so effectively will not contribute to the effective dirty beam. For a simple source that lies within a small support, so that θ is small, many points will enter into the ‘effective uv coverage’. However for such simple sources conventional imaging would normally suffice. What we wish to know is what kind of effective uv coverage is produced when we try to make an image of a useful size, say of 200x200 beams. We obviously require in such cases that the effective uv coverage is substantially better than that of a single frequency observation or else the

whole point of MFS is lost. It is clearly fruitless to devise an algorithm that effectively eliminates the consequences of a space varying spectral index at the cost of reducing the effective MFS uv coverage to approximately the level of a single frequency observation or worse.

In order to assess the suitability of the ‘Gridded Cornwell’ method for MERLIN, the algorithm was tested assuming the standard situation of a 8 antenna MERLIN with 5 frequencies over a range of 25% and a source at declination 50° . Effective dirty beams and their corresponding effective uv coverage obtained after applying the algorithm are shown in fig(4.6).

The steps in producing these effective dirty beams follow from the definition of the effective dirty beam given in equ(4.39). The ‘gridded convolution’ operation denoted by the symbol $A \star B$ can be implemented using the standard convolution routines within existing software packages

Figure(4.6c) shows the effective dirty beam that results when we have an image size of 512×512 pixels and a cell spacing of $0.05''$, this gives a θ of $25.6''$ which corresponds to a grid cell spacing of approximately 4000 wavelengths. For comparison in the same figure are shown similar y axis slices through the dirty beams corresponding to single frequency uv coverage (a) and the multifrequency uv coverage (b). The most noticeable difference between the new effective beam and the other two is the decrease in resolution and the presence of significant ‘wings’ attached to the main lobe. Much clearer differences can be seen if we examine slices through the FFT’s of each beam. Examining such plots allows us to examine the ‘effective uv coverage’ corresponding to each beam. These slices are shown in fig(4.6). Note that the apparent non-uniformity of the amplitudes of the different samples is partly due to taking slices through uv sample points convolved with the effects of gridding and the effects of taking the FT of finite-sized sky images.

Figure(4.6a) shows the effective uv coverage for the single frequency data. Note the considerable gaps along the slice, these gaps are all virtually filled in in the second diagram (b) which shows the situation for multi-frequency uv coverage. In the final diagram we see the effect on this multi-frequency uv

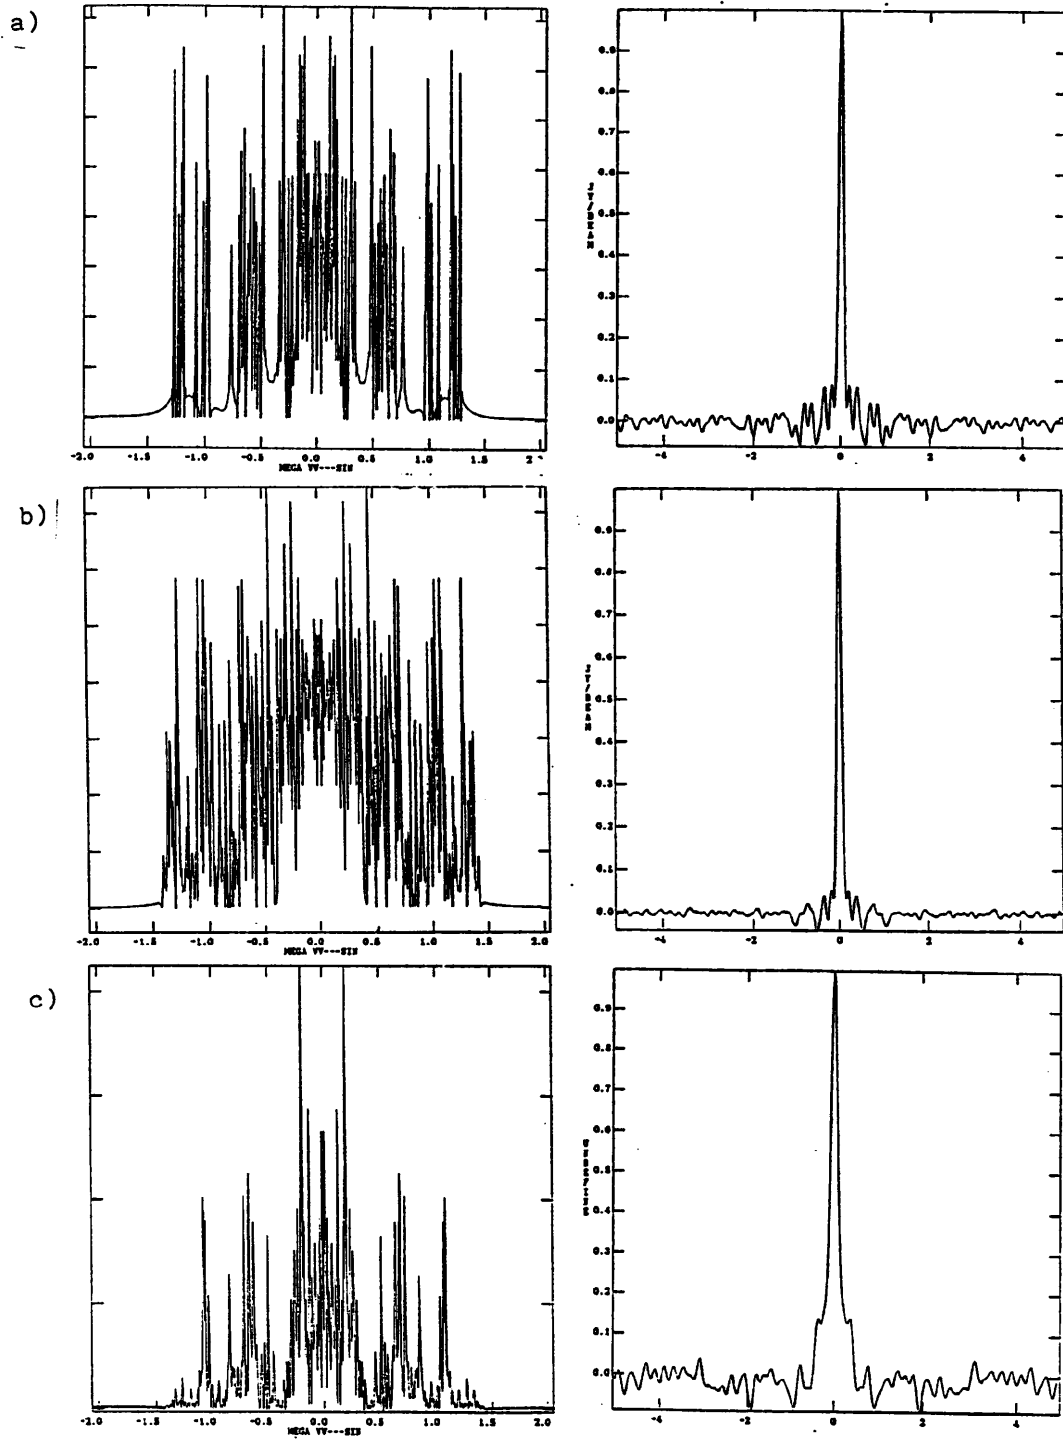


Figure 4.6: Slices through effective dirty beams (right) and effective uv coverages (left). a) Single-frequency observations. b) MFS observations. c) MFS observations after applying the gridded Cornwell method to reduce spectral effects.

coverage of applying the gridded Cornwell method to reduce spectral effects. In this effective uv coverage only uv points near the centre and at isolated points in the outer uv plane have significant weights as shown by the amplitude of the slice through the uv plane. These regions occur only where we get effective ‘double-sampling’. In the centre such ‘double-sampling’ is much more common simply because the gaps between uv tracks measured in wavelengths are much smaller and so the chances of uv tracks at different frequencies lying within the requisite gridding area are correspondingly higher. The increased dominance of shorter spacings in the effective uv coverage explains the lower resolution of the effective dirty beam. We can see that a reduction in effective resolution is likely to be an inevitable consequence of applying the Cornwell algorithm. This problem is also likely to get worse as the size of our image increases. From the examples we can see that, overall, the effective uv coverage after applying the algorithm in the case of a modest image size of 512x512 will be worse than the single frequency uv coverage. This means that the algorithm is pointless for useful sized images. The problem arises from the fact that the MERLIN MFS uv coverage is not sufficiently well ‘mixed’, in the sense of having uv tracks from different frequencies overlapping. To reduce spectral effects significantly we are forced to eliminate much of the extra data we obtain via MFS. The poor applicability of the method to sparse arrays was foreseen by Cornwell (1984). However it seems likely that for much better frequency mixed arrays, e.g. the VLA, the method may prove useful.

4.4.5 Other Approaches to Weighting

Cornwell’s original proposal was motivated by a desire to find a weighting value to give to every observation such that spectral effects would be eliminated over the whole sky. The exact Cornwell method that results is then exactly equivalent to the heuristically obvious method of processing exactly

‘double sampled’ data. Unfortunately in real life on real interferometers which do not possess special geometries exactly ‘double-sampled’ data do not exist. This forces us to consider an approximation to the method that utilizes

gridded data. The act of gridding introduces information about the image size, and the resulting effective dirty map has the biggest reductions in spectral contribution in the centre of the field with the size of the spectral contributions increasing toward the edge of the field. Because *we do not* require spectral effects to be *eliminated completely everywhere* we can effectively afford to relax our requirement as to how close two points must be to be ‘effectively doubled sampled’. In the gridded algorithm points need only be within the same grid cell to contribute to the image reconstruction. The resulting effective uv coverage is therefore much improved over the exact algorithm.

The above approach has clear advantages, it is not however necessarily the best use of the data. Instead of an algorithm that introduces information about the source size via the simple ‘first-order’ method of gridding, we could in principle derive an algorithm that is explicitly based on the reduction of spectral effects *below some threshold level over well defined regions*. Such an algorithm would have clear advantages, limits on spectral contributions in regions of interest could be set and the best effective uv coverage obtained consistent with reducing spectral effects to a required level.

In the exact Cornwell algorithm each uv measurement is multiplied by a weight depending on its position in the uv plane and its frequency of observation. If the uv data has no exact double-sampling then the weighting of each observation is equivalent to the multiplication in the uv plane by some weighting function $\bar{w}(u, v)$ i.e. by a function that depends only on position in the uv plane. This is in turn equivalent to convolving the dirty image by $W(x, y)$, the FT of $\bar{w}(u, v)$. In the two order approximation the dirty image before convolution is

$$d = (I_o * D_o) + (I_o \alpha' * D_1) \quad (4.40)$$

and after convolution/weighting it will be

$$(d * W) = (I_o * (D_o * W)) + (I_o * (D_1 * W)) \quad (4.41)$$

If we have $\bar{D}_o=1$ everywhere we have sampled data then $D_o * W = W$. The problem to be solved is to choose W so that the effective spectral dirty beam

$(D_1 * W)$ is smaller than some critical level over some defined region, whilst ensuring that the effective dirty beam $W(x,y)$ has 'good' properties. These 'good' properties include a compact main lobe (high resolution) and small sidelobes. In the uv plane these properties correspond to having evenly distributed effective uv coverage out to the maximum baseline length.

A useful region over which to reduce the size of the effective spectral dirty beam is the support of the autocorrelation of the source. This region can be defined as the region in which the autocorrelation is non-zero (taken to be the region over which the source is brighter than the noise). Ensuring that the effective dirty beam is exactly zero in this support means that spectral sidelobes around any point within the source do not overlay any other part of the source. If this condition is successfully achieved spectral effects could be entirely eliminated from the non-zero part of the image.

Given a sufficient number of uv points relative to the size of the source support we can normally achieve $(D_1 * W) = 0$ over the autocorrelation of support. Given a rectangular support of size $M \times N$ pixels and hence an autocorrelation of support of size $(2M \times 2N)$ we can almost certainly achieve our aim provided that we have more than $2MN$ uv measurements (reduced by two due to symmetry of the beam). Given our discussion in Chapter 2 of the size and nature of the D_1 beam where it was noted that its unweighted level is already small, we might hope that the additional weighting required to reduce the spectral sidelobes would need only to be modest.

An important constraint if we wish to use the deconvolution algorithm CLEAN to deal with the subsequent deconvolution problem $(W*d) = I * W$, is that the effective dirty beam W has a corresponding uv coverage with no negative values. When all of a beam's weights have positive sign the CLEAN algorithm is mathematically guaranteed to converge to produce an image that fits the uv data (Schwarz 1978). No similar reassurances can be given if there exist negative weights. With the above problems it can be appreciated that,

although the requirements of any weighting algorithm are easily stated, finding the optimum weighting to apply to the data is not trivial.

Because of the difficulties involved, algorithms that implement optimised weighting have not been pursued in this thesis. It must be doubted whether, despite the advantages expected in principle, the practical improvements in quality are generally worth the extra effort. In special cases, however, the differences might be more significant; a particular example might be when the source consists of two or more separated regions of non-zero emission (i.e. ‘disconnected support’) as is commonly found in sources consisting of two separated ‘lobes’. If each lobe lies within a square of size θ_o and the lobes are separated by the larger distance θ , then if we seek to image the full image with the gridded Cornwell method we will effectively exclude ‘double sampled’ uv measurements that are not within a distance of $1/\theta$. In the optimum weighting approach however we could choose to reduce the effective dirty beam only over the autocorrelation of the support, so we can expect the minimum required uv separation to be reduced to the potentially much larger value of $1/\theta_o$. This will increase the number of measurements which contribute to the effective dirty beam and hence the quality of the effective uv coverage.

4.5 MEM Based Methods

Any imaging problem in aperture synthesis can be stated as finding an image that ‘best’ fits the measured uv data, i.e. any problem can be viewed explicitly in terms of model fitting, in which some function which measures the goodness of fit is optimised. The simplest function that can be considered is the mean square fit between the model uv data and the observed uv data. When the uv plane is undersampled however the corresponding imaging problem is ill-posed, there will be many images that fit the data just as well. The corresponding matrix problem is one in which the convolution matrix has no inverse. In such situations the problem can be ‘regularised’ (Titterton 1985) by considering optimising a function which contains additional terms

other than the mean square error fit. These terms will depend only on the model image and not on the data and can embody some *a priori* knowledge about the source. The additional term for instance can be engineered to be small when the model image is smooth, if for some reason it is thought that this is likely. The resulting image from the optimisation will then be biased toward being smooth.

The Maximum Entropy Method (MEM) (see e.g. Narayan and Nityananda 1986) can be thought of as one member of a whole family of regularisation methods. The function that is usually minimised is of the form (Cornwell and Evans 1983)

$$J = H - c_1\chi^2 - c_2F \quad (4.42)$$

where the χ^2 measures the mean square fit between the model uv data and the observed uv data, F is the total flux of the model and c_1 and c_2 are constants. H is an entropy function of the form

$$H = - \sum_i b_i \ln(b_i/m_i) \quad (4.43)$$

Here b_i is the brightness of the model at the i th pixel, and m_i is the brightness of a 'default image' often chosen so that $m_i = \text{constant}$. The above minimisation can be viewed as finding the minimum of the function H , subject to the constraints $\chi^2 = \text{constant}$ and $F = \text{constant}$, by the method of Lagrange undetermined multipliers. The above constraints correspond to the model fitting the data within the noise and having a given total flux. The constants c_1 and c_2 can be adjusted so as to define the numerical values of the expected noise and total flux. The choice of the logarithmic term in the above function can be justified based on a theory of Bayesian statistics as described by Cornwell 1982, and Gull and Skilling 1984. The entropy function can however also be justified more pragmatically (Nityananda and Narayan 1982) in terms of its effects on the resulting image reconstructions.

The MEM method can be easily extended to deal with the MFS problem (T.J. Cornwell, private communication). Although this work is in no way associated with the author, it is useful to briefly describe it here for completeness.

The MFS problem can be posed in terms of optimising two separate images, the spectral index distribution and the intensity image at the reference frequency. The algorithm consists of two stages of optimisation, in the first stage the intensity image pixel values are adjusted whilst in the second stage the spectral index image is varied. The above process is repeated iteratively till the images converge.

For the $I(\nu_o)$ varying part of the algorithm the pixel values are changed so that the sum of the J values for the resulting images $I(\nu_i)$ and the data at each frequency ν_i are minimised. For the α varying half-cycle the $I(\nu_o)$ image is kept constant and $\alpha(x, y)$ adjusted so that the sum of the $J(\nu_i)$ is again minimised. Alternately some simpler optimisation function could be chosen, i.e. the sum of the mean square fits at each frequency (the sum of $\chi^2(\nu_i)$) for instance.

In reality a wide variety of optimisation functions could be used at each stage. The crucial distinctions of this approach to the MFS problem are that the problem is posed in terms of finding images that optimise such functions; that these functions contain non-linear entropy-like terms so as to bias the solutions toward having desirable properties; and that the intensity and spectral index images are varied separately and iteratively so as to optimise the functions. The detailed design of such MEM model fitting methods will to a large extent depend on the feasibility of finding numerical procedures that can optimise the penalty functions. Even with the simplest algorithm of this type however we can expect that the method will require very extensive computer resources.

Despite the above reservations the above approach has one very great virtue; the full MFS data is processed *in toto* at every stage of the algorithm. The algorithm can be non-linear in estimating both the intensity and spectral index distributions. This non-linearity should give useful improvements rather than generate large reconstruction errors because it is applied to as much data as possible at each stage so that the image is only ever mildly underdetermined by the data. Given the above properties the MEM approach is classed with

the ‘Double Deconvolution’ method on the right hand side of fig(4.1). Double Deconvolution will be the sole subject of the next chapter and we will consider in sect(5.5.2) the relationships between this algorithm and MEM methods.

4.6 Summary

We can consider the problem of removing spectral effects as one of estimating, by some means or other, $I\alpha'(x,y)$ and then removing its effects from the multi-frequency uv data. In Chapter 4 three general approaches have been considered for either implicitly or explicitly estimating the $I\alpha'$ distribution. In the first two approaches the $I\alpha'$ distribution can either be estimated by combining the MFS data at each frequency in the uv plane or in the sky plane. The final approach described in sect(4.5) saw the estimation of $I\alpha$ in terms of a explicit model fit.

As pointed out in section(4.2.2) and section(4.3.2) the problem with the two sky plane based methods, which use images made at the separate frequency channels is that each channel map will have reconstruction errors determined by the single frequency uv coverage and these errors will be carried through into the final MFS map. The basic problem with these algorithms is that the piecemeal application of deconvolution does not provide an effective way of applying the necessary non-linear processing when the uv data are undersampled.

Whilst the above approaches are sky plane based the weighting methods described in section(4.4) to remove spectral effects were uv plane based. These methods apply weightings to the uv data to remove exactly or approximately the effects of $I\alpha'$. In their purest forms these methods are wholly linear and use no *a priori* knowledge about the expected I or α' distributions except that they lie within the chosen dimensions of the image. The idea is to remove spectral effects by purely linear means. Once the correction for spectral effects is achieved the remaining MFS data are combined in the uv plane and non-linear deconvolution is applied to the data to give the final I image.

In practice the problem with linear methods such as these is that in order to successfully separate effects due to I and $I\alpha'$, one has to disregard a large fraction of our precious data by weighting much of it to zero. For example isolated uv points have to be zero weighted because there is no linear method which can separate the relative contributions due to I and $I\alpha'$. Zero weighting ensures that they do not contribute to the spectral errors but, as a consequence, they cannot contribute to the estimate of I either. When we finally come to apply non-linear deconvolution there is little data left and so the final results are poor. Because of this problem it appears that simple data weighting methods are not of much use in processing MFS data from sparse arrays like MERLIN (see section(4.4.4)).

In the next chapter we will consider using uv weighting in a slightly different way. This new approach will prove much more fruitful. We do not find a form of weighting which eliminates the effects of the $I\alpha'$ distribution. Instead the weighting will be chosen so that the effects of the unwanted distribution in the sky plane are in a form which can be easily distinguished by a deconvolution algorithm. Data weighting and the *a priori* knowledge within deconvolution algorithms will be used together to separate out the effects of the I and $I\alpha'$ distributions. The method will be non-linear in estimating both I and $I\alpha'$ and will always be applied to the data *in toto*, for these reasons we can expect its performance to be superior to the wholly sky and uv plane based methods described in this chapter.

Chapter 5

Algorithms for Removing Spectral Effects II: Double Deconvolution

5.1 The Double Deconvolution Philosophy

It might be wondered whether we can improve upon the data weighting methods described in Chapter 4 by introducing *a priori* information into the algorithms. Such *a priori* information might help suppress the spectral effects yet allow as much of the recorded information as possible to be used in estimating I. We might consider applying some linear process in the uv plane (i.e. data weighting) to emphasise the I or $I\alpha'$ distributions and then transform these data to the sky plane where non-linear deconvolution algorithms would utilise *a priori* information to separate the desired distribution from the undesired one.

The weighting schemes would be chosen to emphasise the brightness of one distribution relative to the other unwanted one whilst ensuring that as few uv data points as possible are weighted to small or zero values. The weights would also be chosen so that properties that would be recognised by the deconvolution algorithm would be emphasised in the wanted distribution and the same property suppressed in the contribution of the unwanted distribution. As an example, the weighting could ensure that the desired distribution produced compact structures in the residual map whilst the effects of the undesired distribution were scattered over the whole map. Any deconvolution algorithm which contained an *a priori* bias toward recognizing compact structures should then be largely successful in separating out the wanted distribution in the presence of errors caused by the unwanted distribution. The

deconvolution step would serve a similar role as it does when used in self-calibration algorithms where it helps in finding a phase solution which gives a compact image (Readhead and Wilkinson 1978).

The double deconvolution philosophy, unlike the simple uv weighting and ‘Map and Stack’ methods, appears both to utilise the whole of the uv data and to use *a priori* knowledge to separate the I and $I\alpha'$ distributions. For these reasons we can hope that it should be a superior approach.

In the next section we will concentrate on a more concrete description of a working algorithm which exploits the ‘double deconvolution’ philosophy.

5.2 Implementing the Double Deconvolution Philosophy

5.2.1 A CLEAN Based Double Deconvolution

Within the underlying philosophy given above very many varieties of actual algorithm are possible. The choices of the weighting functions to be employed, of the desirable properties to be emphasised and of the deconvolution algorithm to be used are wide, and all determine the final form of the algorithm. In this section we make simple choices for the above elements and describe one implementation of the double deconvolution philosophy.

We start with the familiar two order expression

$$d = (I * D_o) + (I\alpha' * D_1) \quad (5.1)$$

If we already possess some good estimate of the I distribution I_m then the effects of this can be removed from the data to produce a residual map

$$r_o = (I - I_m) * D_o + I\alpha' * D_1 \quad (5.2)$$

The next step in line with our philosophy is to find some uv weighting function (or equivalently some convolving function in the sky plane) which emphasizes the $I\alpha'$ contribution at the expense of the I contribution without weighting most of the uv data to zero. There are of course an infinite number of possible

weighting schemes to be considered. If, however, we consider the convolving function D_1 we will see it has some desirable properties. On applying this weighting to the residual uv data we obtain in the sky plane

$$r_o * D_1 = r_1 = (I - I_m) * D_1 + I\alpha' * D_2 \quad (5.3)$$

where we use the result that $D_n * D_m = D_{n+m}$ (this follows by considering the multiplication of the uv weights corresponding to each beam).*

From Chapter 2 we know that the properties of odd and even order spectral dirty beams are quite different. The odd beams have no large central peaks and their overall magnitudes are small, in contrast the even beams are quite similar in form to conventional dirty beams. The absence of a *central* peak in D_1 means that the effects of the residual $(I - I_m)$ distribution are scattered over the whole map whilst the effects due to $I\alpha'$ which is convolved with D_2 remain relatively compact.

The difference in form between odd and even beams together with the overall decrease in level with increasing order ensures that whilst the peak of the D_1 and D_2 beams are similar both are considerably smaller than the peak of D_o . For a bandspread of 20%-25% the D_o beam has a peak which is approximately 200 times greater than that of D_1 or D_2 .

The relative levels of the beams imply that after weighting the data the $I\alpha'$ distribution will be convolved with a beam which is of the same magnitude as before weighting. In contrast the I distribution will be convolved with a beam with much lower peak level than before. One overall effect of the weighting is therefore to emphasise the $I\alpha'$ contribution relative to the I distribution.

The next step according to our philosophy is to deconvolve the residual map r_1 with the dirty beam D_2 to provide an estimate of $I\alpha'_m$. Note in practice that it is usually convenient first to normalise the image r_1 by dividing by $D_2(0,0)$, the central value of the D_2 beam. The beam response that must be deconvolved out from the normalised residual map is then $D_2/D_2(0,0)$ which, like a normal dirty beam, has a central value of unity.

We choose as our deconvolution method the CLEAN algorithm which we

* This assumes as will the rest of our derivation that each SDB is formed from naturally weighted uv data.

weighting schemes to be considered. If, however, we consider the convolving function D_1 we will see it has some desirable properties. On applying this weighting to the residual uv data we obtain in the sky plane

$$r_o * D_1 = r_1 = (I - I_m) * D_1 + I\alpha' * D_2 \quad (5.3)$$

where we use the result that $D_n * D_m = D_{n+m}$ (this follows by considering the multiplication of the uv weights corresponding to each beam).*

From Chapter 2 we know that the properties of odd and even order spectral dirty beams are quite different. The odd beams have no large central peaks and their overall magnitudes are small, in contrast the even beams are quite similar in form to conventional dirty beams. The absence of a *central* peak in D_1 means that the effects of the residual $(I - I_m)$ distribution are scattered over the whole map whilst the effects due to $I\alpha'$ which is convolved with D_2 remain relatively compact.

The difference in form between odd and even beams together with the overall decrease in level with increasing order ensures that whilst the peak of the D_1 and D_2 beams are similar both are considerably smaller than the peak of D_o . For a bandspread of 20%-25% the D_o beam has a peak which is approximately 200 times greater than that of D_1 or D_2 .

The relative levels of the beams imply that after weighting the data the $I\alpha'$ distribution will be convolved with a beam which is of the same magnitude as before weighting. In contrast the I distribution will be convolved with a beam with much lower peak level than before. One overall effect of the weighting is therefore to emphasise the $I\alpha'$ contribution relative to the I distribution.

The next step according to our philosophy is to deconvolve the residual map r_1 with the dirty beam D_2 to provide an estimate of $I\alpha'_m$. Note in practice that it is usually convenient first to normalise the image r_1 by dividing by $D_2(0,0)$, the central value of the D_2 beam. The beam response that must be deconvolved out from the normalised residual map is then $D_2/D_2(0,0)$ which, like a normal dirty beam, has a central value of unity.

We choose as our deconvolution method the CLEAN algorithm which we

* This assumes as will the rest of our derivation that each SDB is formed from naturally weighted uv data.

know is especially good at recovering compact structure and which is convenient to use. Note that the choice of D_1 as our convolving function has an essential property if we choose to use CLEAN. It ensures that the resulting beam which we must deconvolve, viz D_2 , consists solely of positive weights in the uv plane, a situation which is essential if we wish that our CLEAN deconvolution does not diverge (Schwarz 1978).

Having made an estimate of $I\alpha'$ from our MFS data we can attempt to improve the estimate of the I distribution. We proceed in an analogous manner as in the $I\alpha'$ half-cycle by first removing the effects of our estimate of $I\alpha'$, $I\alpha'_m$, from the dirty map to produce the residual map

$$r_o = I * D_o + (I\alpha' - I\alpha'_m) * D_1 \quad (5.4)$$

At this point we could as in the $I\alpha'$ half cycle emphasise the new desired distribution I, relative to the undesired distribution $I\alpha'$, by multiplying the residual data by some weighting function in the uv plane. Such uv weighting is of course equivalent to convolution in the sky plane with some beam. Such methods of weighting were considered first in section(4.4). It was found that these methods tended to give small weight to points which were not paired with other points or 'double sampled' and because of this these methods were wasteful of data and hence unattractive. Rather than apply additional weighting it is probably best therefore to leave the residual data alone.

If we wish to consider that the residual uv data set has been convolved with a beam we can consider this as being D_o . Such a convolution of course leaves the residual map unchanged. If we recognise this notional convolution we see that the treatment of the residual data in the two half-cycles is analogous. In both cases the convolving beam is taken to be equal to the beam with which the wanted distribution is convolved in the two order expansion. We can consider these convolving beams to be chosen so as to 'recognise' the appropriate distribution in each half-cycle by finding the highest correlation between the convolving beam and the dirty map (as in the original justification behind CLEAN, see sect(5.2.3)).

If we choose no additional weighting as described above we proceed in the

I half-cycle by deconvolving the residual image r_o with D_o so as to form an updated estimate of I, I_m .

In the above method, therefore, we process the data in turn to produce estimates of $I\alpha'$ and I. Because these estimates will be inaccurate due to the presence of errors due to the effects of the undesired distribution, it seems sensible to apply the above processing for the I and $I\alpha'$ distributions iteratively. The full double deconvolution algorithm is therefore as shown in fig(5.1).

5.2.2 An Illustration of A CLEAN based Double Deconvolution.

The progress of a cycle of the DD algorithm is illustrated in fig(5.2) for a model data set derived from the ‘target’ double-lobed radio model source. We start in the I ‘half cycle’ with a residual map (a), ring-like features surrounding the hotspots are the effects of $(I\alpha' - I\alpha'_m) * D_1$, i.e. of uncleaned $I\alpha'$. A notional convolution with D_o leaves the resulting image (b) the same as (a). Given that we know the true $I\alpha'$ distribution exactly we can calculate the contribution to (a) and (b) due to uncleaned $I\alpha'$; this is shown in (c). This distribution, plus $I_o * D_o$ (not shown) should equal the total dirty map (a). We now apply CLEAN for a limited number of iterations to the total dirty map in an attempt to separate the effects of the two distributions by estimating I_o with I_m . The deconvolution is stopped before CLEAN components are heavily contaminated by the effects of the undesired distribution, $(I\alpha' - I\alpha'_m) * D_1$. Our estimate I_m (d) of I_o should have been improved by passing through the deconvolution ‘filter’.

Having made an estimate of I we are ready to start that half of the DD which estimates the $I\alpha'$ distribution. We first subtract the effects of I_m from the observed data to produce fig(5.2e) which is $(I - I_m) * D_o + I\alpha' * D_1$. We now apply our weighting function to the uv plane, or equivalently convolve in the sky with D_1 , specifically to enhance and recognise the $I\alpha'$ distribution. The resulting image $(I_o - I_m) * D_1 + I\alpha' * D_2$ is shown in fig(5.2f). It can be seen that the convolution step is very effective in enhancing the $I\alpha'$ distribution in

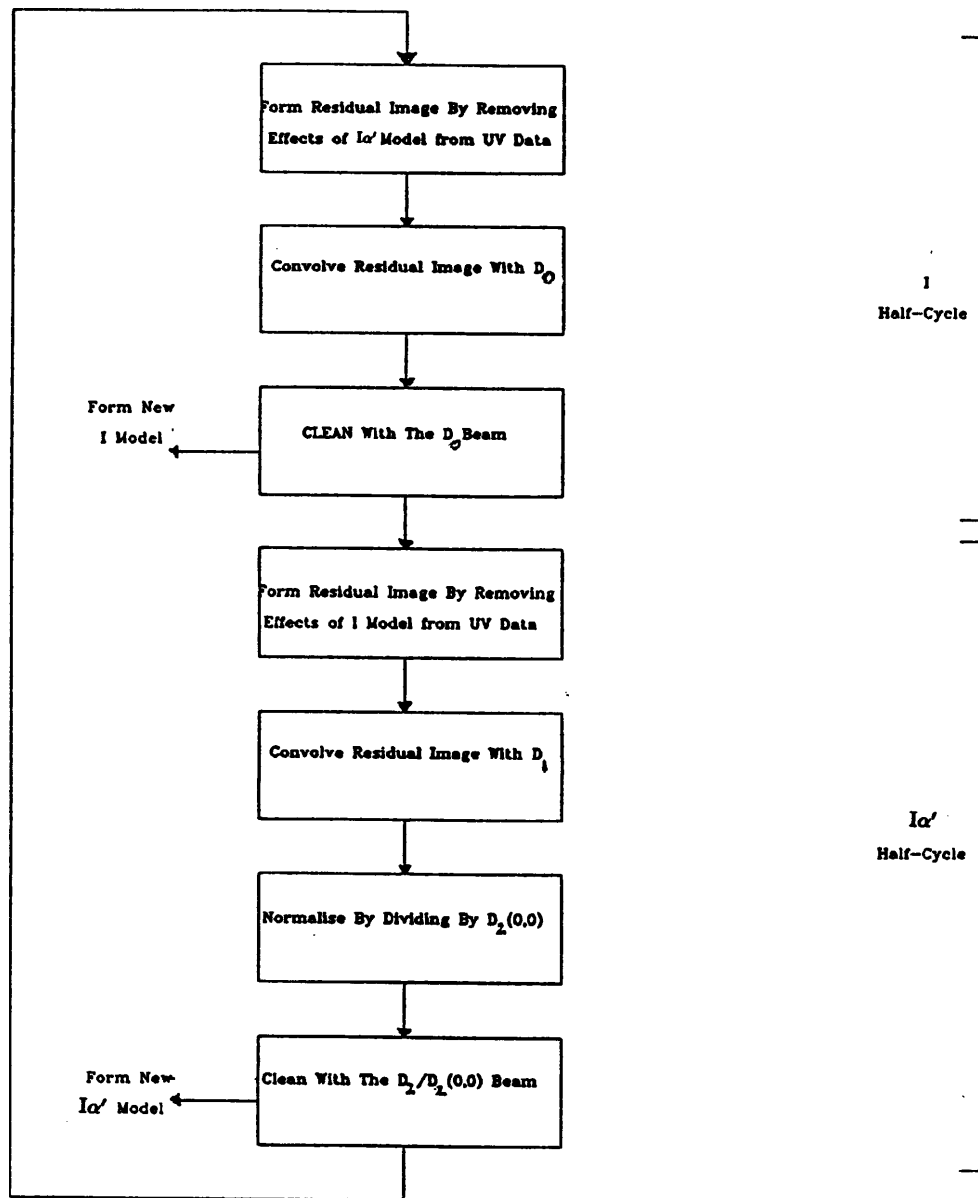


Figure 5.1: Flow diagram of the Double Deconvolution algorithm

the image. The image (fig(5.2f)) can be considered as the sum of contribution due to $I\alpha' * D_2$ (not shown) and an 'error' term due to uncleaned I , $(I - I_m) * D_1$ shown in fig(5.2g). In a similar fashion to the I deconvolution half of the cycle we now apply 'limited iteration' CLEAN with D_2 as the dirty beam. This will produce an estimate of $I\alpha'$, $I\alpha'_m$ (h).

5.2.3 Double Deconvolution as a Generalised CLEAN

The particular implementation of the DD algorithm described above can be thought of as a generalised form of CLEAN in which first we attempt to recognise and remove the effects of I in the image and then we recognise and remove the effects of $I\alpha'$.

The step in DD of convolving the residuals with D_1 so as to emphasise features in the dirty map due to the convolution of $I\alpha'$ with D_1 can be thought of as an attempt to recognise the features convolved with D_1 by finding those parts of the dirty map most correlated with it. Analogously, in the I half-cycle too, even though no explicit convolution is carried out, the dirty map can be considered as being convolved with D_o so as to recognise the sidelobes surrounding features in the I image.

Recognizing the effects of a dirty beam in an image and then removing its effects is exactly the principle behind the original proposal for the CLEAN algorithm (Högbom(1974)). The algorithm was defined as a method of finding the highest correlation between the dirty map and the dirty beam and then carrying out the iterative removal of the effects of this dirty beam. It appears that this implementation of DD can be considered as a generalised form of CLEAN which can be applied when the residual image is equal to the sum of two convolutions as is the case in the MFS problem.

5.2.4 The Mechanism of Convergence

Can we describe in some heuristic way how the DD algorithm actually separates the I and $I\alpha'$ distributions? Every time we attempt to form estimates of I or $I\alpha'$ we will be deconvolving in the presence of an error contribution

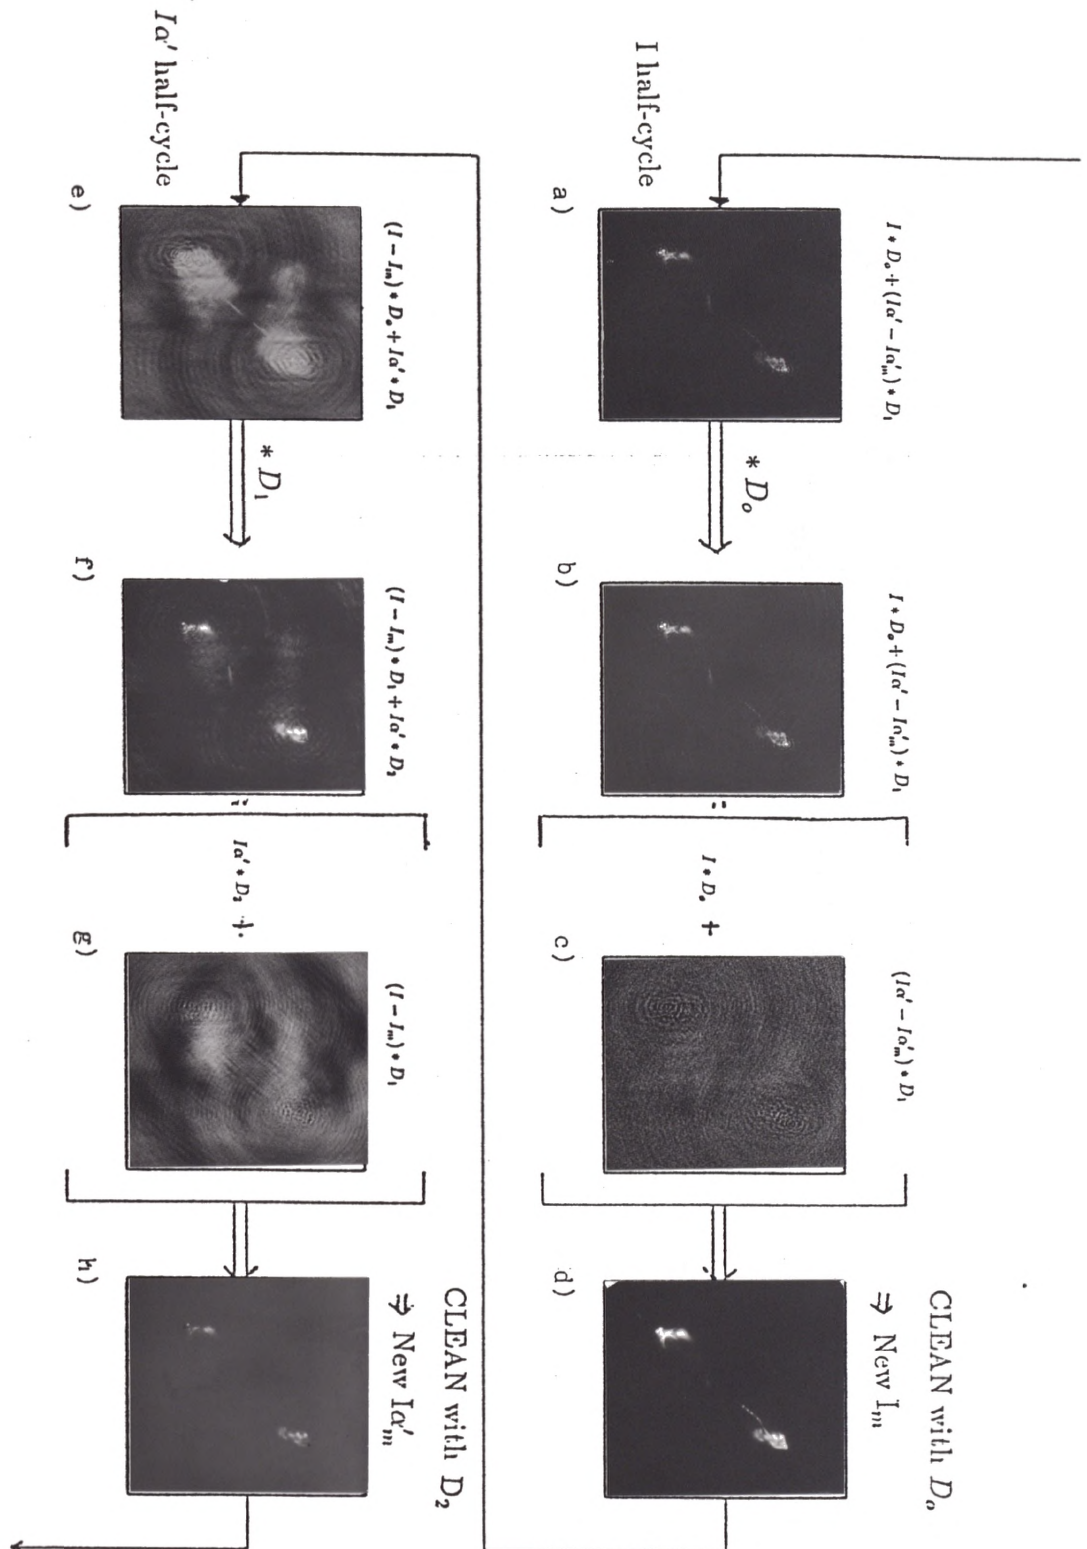


Figure 5.2: Pictorial illustration of the different stages of Double Deconvolution processing

in the dirty map. At least some of the features of this error distribution will contaminate our estimates of I or $I\alpha'$ no matter when CLEAN is terminated. When we subtract the effects of our estimates of I or $I\alpha'$ from the observed data at the beginning of a half cycle these errors will be fed back into the algorithm. Why should we believe that this does not cause the DD algorithm to diverge but rather still converge towards a successful separation between I and $I\alpha'$? Is it possible to identify the mechanism by which the algorithm recovers from 'overcleaning' in a particular half-cycle?

We shall see that it is possible to recover from 'overcleaning' because the uncleaned spectral sidelobes in the I domain will still contain enough information to allow us to form an improved model of $I\alpha'$. The spectral sidelobes away from the source will be almost unaffected by any overcleaning and will continue to contain information about the $I\alpha'$ distribution encoded within them. If we obtain an improved $I\alpha'$ model from this residual image using DD the spectral sidelobes in the next I half-cycle residual map can be reduced, this in turn eventually reduces the errors in the next $I\alpha'$ estimate. Over many cycles we can get a cumulative reduction of the contribution due to initially having an imperfect $I\alpha'$ model.

To make the above mechanism more explicit consider the initial dirty image in the I domain

$$(I * D_0) + (I\alpha' * D_1) \quad (5.5)$$

Consider CLEANing this dirty image with the D_0 beam for a limited number of iterations. The resulting I model will consist of a good model of the intensity distribution plus a contamination due to spectral errors. We know (from section(3.3)) that on CLEANing a dirty image which contains a contribution due to spectral sidelobes, the contamination in the resulting image will be approximately equal to the spectral error contribution within the area CLEANed. There is no special interaction between CLEAN and the spectral errors which increases their effects. The spectral errors are simply transferred from the dirty image into the CLEAN image. Given this the contamination

in our I model, I_m , can be modelled mathematically by

$$I_m = I + (I\alpha' * D_1)f \quad (5.6)$$

where f is a function of position which is non-zero only within F , the region over which CLEAN components are found. The function f will have its maximum value of unity at points within F which have been deeply CLEANed. If we proceed to the second half cycle, removing our model of I from the uv data and convolving with D_1 we obtain as our dirty map

$$r_1 = (I - I_m) * D_1 + I\alpha' * D_2 \quad (5.7)$$

using equ(5.6) we obtain

$$r_1 = -[f(I\alpha' * D_1)] * D_1 + (I\alpha' * D_2) \quad (5.8)$$

The value after normalisation will be

$$I\alpha' * D_2 / D_2(0,0) - (I\alpha' * f D_1) * D_1 / D_2(0,0) \quad (5.9)$$

the magnitude of the second error term will depend on f and in particular the area over which it is non-zero. If $f=1$ over all space then the second term will equal the first, if it is non-zero only over a small area then the error will be relatively small. In this second case when overcleaning occurs only over a small area, there will be sufficient information contained within the unaffected sidelobes so that even after only a single cycle of DD the $I\alpha'$ dirty map is almost free of errors.

If DD is applied for many half-cycles, the finite areas over which I and $I\alpha'$ models are CLEANed will continue to reduce the contribution from the initial error. As an example, after one half-cycle CLEANing in the I domain and one half-cycle CLEANing in the $I\alpha'$ domain the error on the I domain dirty image will be

$$\frac{f_2[f_1(D_1 * I\alpha') * D_1] * D_1}{D_2(0,0)} \quad (5.10)$$

where f_1 and f_2 are the functions describing the depth and area of CLEANing in the I and $I\alpha'$ half-cycles respectively. The above expression can easily be

generalised to any number of half-cycles. It can be seen that if the areas over which the f functions are non-zero are moderate in size we can expect a large decrease in the size of the error contribution after every half-cycle.

The separation between the two distributions clearly depends on the restriction of the area over which CLEANing takes place. If we allowed deep CLEANing over the whole sky then there would be no progress toward successful separation of the distributions. The restriction of the CLEANing so that it is largely only the wanted compact distribution that is CLEANed into the model rather than the widely scattered spectral errors lies at the heart of the DD algorithm. Given the importance of restricting the depth and extent of CLEANing we should therefore consider the nature of the CLEAN ‘termination conditions’.

Limiting CLEAN by stopping the process when the residuals reach a level equal to that of the expected contribution from the unwanted distribution is one possible method of control. It seems sensible to stop CLEANing when the residuals become dominated by the error distribution. For many sources this method will automatically ensure that a large fraction of the real structure is CLEANed, whilst ensuring that the error distribution is CLEANed only within a modest area. This method of terminating the DD algorithm can be termed the ‘residual limit’ method. Note that it only indirectly restricts the area of CLEANing to a ‘safe’ level. Methods of control using windows, would in contrast limit the area of CLEANing directly. Despite being indirect we shall see in the next section and Appendix F that under certain conditions CLEAN termination based on ‘residual limit’ control will be mathematically guaranteed to give a successful separation.

Despite the above, problems will often arise with ‘residual limit’ control when a point source lies embedded inside a diffuse lobe or halo. In this case both the I or $I\alpha'$ images consist of components of quite different scale sizes. Setting a low enough residual limit so that CLEANing occurs within a large enough area to represent the diffuse component will allow deep CLEANing of the spectral sidelobes around any compact component. Because the spectral

sidelobes from the compact component are superimposed on top of the diffuse component they can be CLEANed into the I model, even if the residual limit is larger than the expected level of spectral sidelobe contribution (see fig(5.3)). The CLEANing of the spectral sidelobes permitted in this case may be so extensive that it is not possible to subsequently estimate the $I\alpha'$ contribution of the compact component from the residual data. We will see this effect in our simulations of the progress of the double deconvolution algorithm which will be presented in section(5.4).

The above mechanism arises because setting a residual limit is only a good way of separating the two distributions when the source is compact. For a source containing many different scale sizes the residual limit method will instead be a poor way of separating between the two distributions. One way of improving things for such sources would be to set termination limits directly based on set 'CLEAN area limits'. A crude way to control the CLEANing in this way would be interactively to set CLEAN windows at each half-cycle. A more sophisticated method would operate automatically. A half-cycle would terminate when CLEAN components had been found on a critical number of pixels. The size of the critical area would be increased gradually and in a controlled manner every half cycle. Such methods of control can be termed 'CLEAN area limited'.

From the discussion of the mechanism of convergence it was concluded that the restriction of CLEAN area was the primary way in which convergence was achieved in DD. It therefore seems sensible to apply the area restriction directly rather than indirectly through residual limits. Although this method of controlling DD may prove superior to methods based on 'residual limits' as yet no mathematical criteria exists for deciding when a successful separation is assured.




Figure 5.3: Sketch diagrams illustrating the problems of applying DD to images containing structures of different scale size. It is assumed that the value of α' is constant so that the I and Ia' images are similar. In the case of the two component model the spectral sidelobes overlay the $I * D_o$ contribution and will be CLEANed into the I model before the critical residual level is reached. The resulting residual map (far right) may consequently contain too little information to allow for a successful solution for the Ia' distribution.

5.3 Convergence Conditions for Double Deconvolution

5.3.1 Introduction

In this section we will consider under which particular conditions the DD algorithm will converge to a solution and in which situations this solution will be a good approximation to the true sky distribution. We will specifically analyse the form of the algorithm that has been described above, i.e. we choose the convolving functions in the two half cycles to be respectively D_0 and D_1 and the deconvolution algorithm to be the method CLEAN.

In considering the behaviour of DD we will first consider under what conditions the final solutions for I and $I\alpha'$ give model uv data that fit the observed uv data. We must show there are no common circumstances in which the algorithm diverges. A second question is, given that the above convergence occurs, under what circumstances will the algorithm correctly divide the data at each uv point between the contribution due to I and that due to $I\alpha'$? This is the special problem which our MFS processing algorithms must attempt to solve.

Given that the observed uv data is correctly fitted and the contributions due to I and $I\alpha'$ are successfully separated there is still the problem of 'unseen' distributions. If there are insufficient observations to define a source structure uniquely then even if the I and $I\alpha'$ give correct data at the sampled uv points there are many possible ways for the I and $I\alpha'$ models to interpolate into the gaps between observations. How does the algorithm choose between the infinity of possible solutions that fit the data? Unfortunately it is a well known problem with the CLEAN algorithm, applied to undersampled data, that it is not easy to say which of the possible solutions it will choose even in conventional mapping (Tan 1986). Understanding how the algorithm behaves in this respect within the DD algorithm will therefore be very difficult indeed.

5.3.2 Convergence of CLEAN

In developing a mathematical description of the first two aspects of the convergence of the DD method, we will draw heavily on the work done by Schwarz (1978) on the convergence of the CLEAN algorithm. In Schwarz's paper the deconvolution problem that CLEAN seeks to solve in conventional mapping is presented in terms of solving the system of linear equations denoted by the matrix equation.

$$d = Bt_o \quad (5.11)$$

This form of describing the convolution problem is described more fully in Appendix C. In Schwarz's nomenclature, d is the dirty map vector of dimension N produced by 'lexicographically' ordering the pixel values (Andrews and Hunt 1977) and t_o is the vector representing the true source brightness distribution. The matrix B connecting d and t_o is an $N \times N$ matrix and it has a Toeplitz form (Andrews and Hunt 1977) with each row consisting of a shifted version of the dirty beam. The matrix B is known as the beam matrix and the operation that results from the matrix multiplication Bt_o is the convolution of t_o with the dirty beam.

Schwarz's (1978) proof shows that the estimate t of t_o that emerges from the CLEAN algorithm is such that the residual map r defined by

$$r = d - Bt = B(t_o - t) \quad (5.12)$$

and the norm

$$(t_o - t)^T B(t_o - t) \quad (5.13)$$

can both be reduced arbitrarily close to zero, or equivalently the model visibility calculated from t fits the observed data, provided that three conditions are met.

The first condition is that the dirty beam matrix B is symmetric, this condition always holds for data from an interferometer. The second requirement is that the dirty map is 'in the range' of the dirty beam (as first described in section(3.3.1)), i.e. there must exist a t' such that $d = Bt'$. The third and final requirement is that the beam matrix B be at least 'positive semi-definite',

that is any quadratic form involving B such as $v^T B v$ where v is an arbitrary vector must be greater or equal to zero for all possible vectors v . It is possible to show this will be the case if the weighting function in the uv plane corresponding to the dirty beam is everywhere positive or zero. If all three of the above conditions are met then Schwarz (1978) shows that convergence to fit the data is assured.

5.3.3 Convergence of DD to Fit MFS Data

For the case that we allow CLEANing in either half cycle of DD to occur over an infinite window the convergence to fit the MFS data follows directly from the conventional Schwarz proof. Consider deconvolving in the first half cycle to update our estimate of I . The dirty map will be $d = B_0 t_0 + B_1 s_0$, where the two matrices perform convolution with D_0 and D_1 respectively and where t_0 and s_0 denote the true I and $I\alpha'$ distributions (see Appendix C for details).

This dirty map will be in the range of B_0 because there exists at least one solution $t' = t_0 + B_1 s_0$ such that $d = B_0 t'$ where t' lies inside the window because it is infinite. If the matrix of the beam with which we deconvolve, B_0 is both symmetric and positive semi-definite it follows that the norm $(t - t_0)^T B_0 (t - t_0)$ continually decreases, can be made zero and hence the MFS uv data can be fitted exactly. As pointed out by Tan (1986) the proof of convergence of CLEAN does not depend on the sky plane distribution that causes the observed visibilities. Arbitrary values (i.e. even noise) of the visibilities at the observed uv points can be fitted by CLEAN provided the window we are allowed to CLEAN over is large enough (and provided there are no other constraints such as only allowing positive CLEAN components.)

Now consider deconvolving in our second half cycle to update $I\alpha'$. We can show that here too we can fit the uv data exactly by just finding a solution s and leaving $t=0$. If we allow infinite support for our solutions we see that the fit to the data always improves in both half-cycles of DD and we can eventually fit the data exactly. If infinite windows are used in both half-cycles

the separation of I and $I\alpha'$ chosen by DD will depend entirely on the choice of the points at which we switch between half-cycles of DD.

In realistic cases we must operate our DD algorithm with a finite-sized residual map and hence there is a finite 'window' in which our estimates of t and s are allowed to be non-zero. The finite window is both necessary, practically, and useful because restricting the solutions for t and s to particular areas will reduce the possible solutions that fit the data. The finite windows constitute *a priori* knowledge about the source support, which can be used to separate I and $I\alpha'$ effects.

In considering the convergence of double deconvolution in the case of finite windows, we must adapt the Schwarz (1978) proof of the convergence of CLEAN to this new algorithm. In Appendix E the analysis of Schwarz(1978) is applied separately to the CLEANing operation in the two different domains, i.e. in CLEANing for the I and $I\alpha'$ distributions. It is shown that the result is always to produce I and $I\alpha'$ models which continuously improve the fit between the total model uv data and the observed uv data. In no circumstances will this fit get worse provided the normal requirements for CLEAN hold. If the windows are 'correct', in that they encompass the non-zero regions of the I and $I\alpha'$ distributions, then the fit can in principle be made exact and the residuals can also be made zero.

DD will converge because the requirements necessary for convergence hold for the CLEAN operation in each of the two half-cycles. Both the beam matrices B_0 or B_2 will be symmetric because both beams are symmetric. The dirty map will be 'in the range' of the dirty beam matrix despite the presence of errors due to the other distribution as is shown in Appendix C. Lastly the beam matrices B_0 and B_2 will be at least positive semi-definite provided that the uv weights corresponding to the beams D_0 and D_2 are always positive. This is always true given the definitions of D_0 and D_2 .

The convergence to fit the data has been described in this section. The total uv model due to the I and $I\alpha'$ distributions will converge to fit the observed data at the sampled uv points. The next question that must be answered is

how well does DD separate out the effects of the I and $I\alpha'$ distributions at these sampled points?

5.3.4 Separating the I and $I\alpha'$ Distributions

Depending on the uv coverage and the size and shape of our windows the solution that emerges from DD can either be a unique solution or one of an infinity of possible solutions. Only in the former case can we be sure that our solution successfully separates the effects of I and $I\alpha'$. If the solutions are not unique it would be nice to know what role the cycle termination parameters of DD play in deciding which of the possible solutions is chosen by DD.

From Appendix F it is found that we get a unique, correct, solution from DD if the $N \times N$ pixel field size includes the regions where t_o and s_o are non-zero and the matrix C given by

$$C = \begin{pmatrix} B_o & B_1 \\ B_1 & B_2 \end{pmatrix} \quad (5.14)$$

is non-singular. The matrices B_o , B_1 and B_2 are the $N \times N$ dirty beam matrices corresponding to the different order spectral dirty beams. A necessary though not sufficient condition that C is non-singular is that the number of uv measurements S is greater than $2N$, where N is the total number of elements in the window. This seems a reasonable condition, given that there are M pixels in the image and there are two unknowns per pixel, i.e. I and $I\alpha'$, if we have more than $2N$ uv measurements we should be able to uniquely define all the unknowns. If the uv coverage consists of essentially randomly distributed uv points and the above inequality holds then C will normally be non-singular and the solution will be both unique and correct. Note this argument says nothing of the conditioning of the solution, to decide about any possible sensitivity to noise will require study of this conditioning and an understanding of how the non-linear DD algorithm deals with noisy data. Simulations do not show, however, that the treatment of noisy data is a serious problem (see sect(5.4)).

If we are not in the oversampled limit how should we steer the DD algorithm so that it finds a solution that successfully separates the contributions due to I and $I\alpha'$? As described in section(5.2.4) there are two possible ways of controlling the depth of CLEANing in each half cycle, one based on terminating CLEAN when a certain peak residual is reached and the other based on terminating when CLEAN components are found in a given area. According to the explanation of the mechanism of convergence given in section(5.2.4) this second method of control seems more natural. To date however only the residual limit method has been tested (see section(5.4)) and the only theoretical work which has been done has also been exclusively concerned with this termination method. The fact that our discussion below is confined to the residual limit method does not imply that this is the best method, only that it is the only one so far explored.

This question of separating the two component distributions with DD using residual limits is tackled in Appendix F where we consider the norms Q_0 and Q_1 . The norm Q_0 measures the fit of the model uv data corresponding to the I model, t , to the true data corresponding to the Fourier transform of the true sky, t_0 . The norm Q_2 measures a similar fit between the model $I\alpha'$ uv data and the true $I\alpha'$ data. Note that unlike Q_T (see Appendix E) which measures the fit of the total model to the data the norms Q_1 and Q_2 are inaccessible in a real deconvolution because we do not know the true I or $I\alpha'$ distributions. Despite their inaccessibility consideration of these norms is useful when we are considering the mathematical requirements for a successful separation. By finding the conditions under which the norms Q_0 and Q_2 continuously decrease we can find the conditions such that the DD solution converges toward a solution that both fits the overall data (as described in the previous section) but which also successfully separates between the I and $I\alpha'$ contributions.

Appendix F shows that provided we stop CLEANing in a given half cycle and go onto the next when the magnitude of the peak residual is still larger than a certain optimal value convergence is assured. When g is the loop gain

of the CLEAN this optimal value is $(2g/2g-g^2)$ times the size of the peak contribution of the error contribution. Note that the above factor is approximately unity if $g \ll 1$. For good separation between the two distributions we therefore have the reasonable condition that the CLEAN limit should be larger than the size of the error contributions.

The derivation in Appendix F is useful in showing us the form of the optimal ‘residual limit’ convergence condition. In practice however it does not give a complete solution to the question of how to run DD so as separate the two distributions. Because we do not know perfectly the distributions I and $I\alpha'$ (or else there would be no point in attempting to image the source) we cannot find exactly the expected spectral errors and therefore cannot calculate the size of the critical residuals.

What can be done, however, is to measure the statistics of the residual image and then use these measurements to set model-dependent limits on the size of the contribution due to the undesired distribution. From these estimates the CLEAN limits can then be set. As an example we could measure the rms of the residual image within a region distant from the radio source which is thought *a priori* to be empty. If we assume that this non-zero flux is entirely due to the far sidelobes of D_1 surrounding the unwanted distribution we can estimate an upper limit for the errors due to the unwanted distribution on-source and thence set a residual limit.

To convert the measured rms values to the critical spectral errors on-source we must use the semi-empirical rules of section(3.2.1) that relate the peak brightness on source to the level of spectral sidelobe. These relationships are estimated from a model of a radio source consisting of a collection of gaussian components and to this extent they will be model dependent. If the unwanted distribution contains features looking like D_1 then these relations will break down and we will not be able to get an accurate estimate of the safe CLEAN limit. In this implementation of DD the *a priori* knowledge that allows a successful separation between the two distributions ultimately lies in the implicit assumptions made about the structure of the sky.

5.3.5 Summary

In this section we have considered the performance of the DD algorithm. We have seen that the solution that emerges from DD should always converge to fit the uv data provided that some simple properties hold. These properties are normally present.

We have also considered the conditions under which DD can successfully separate between I and $I\alpha'$ distributions. In the oversampled case we have proved mathematically that in the absence of noise the DD solution is unique and correct. In the undersampled case the situation is more complicated, the particular solution chosen will depend on the set of CLEAN limits chosen. In Appendix F we have considered a criterion which defines an optimum CLEAN limit for a particular half-cycle, unfortunately this optimum CLEAN limit will depend on quantities which cannot be known exactly. Schemes can be devised however to estimate the 'optimal' CLEAN limits roughly, based on the statistics of the residual map and implicit assumptions about the morphology of the true I and $I\alpha'$ distributions.

Limited iteration CLEAN is biased against including flux that has been scattered over the map, because such flux is incorporated into the models only slowly. The error effects are both scattered and uncorrelated with the dirty beam so their level does not decrease very rapidly with CLEANing. The limited area that can be CLEANed in each half-cycle in a reasonable time ensures that the algorithm can recover from CLEANing errors into the image. DD can be optimally biased against including unwanted errors, and hence convergence be made faster and more assured, if we can estimate the error brightness level and stop CLEANing in each half-cycle before this level is reached.

Future theoretical work should concentrate on developing the termination criteria for the half-cycles within DD. Some work has already been done on the residual limit method and this should be extended, most important however is to develop the theory of the area limit method of control first described in section(5.2.4). Another objective is to understand the treatment of noisy

data. Although the simulations in section(5.4) will show that the noise in the solutions are very close to the noise in the dirty maps it would be best to understand the effects of noise theoretically.

Despite the extensive discussion in this section about the methods of controlling DD, in practice it is likely that the actual levels of the residual limits set will not be too critical. We will find that we get good solutions though slower convergence from DD even if the limits are only roughly estimated. This is because of the large numbers of iterations that are needed to CLEAN in a significant fraction of spectral sidelobes. CLEAN limits set by the user examining the residual maps intelligently and estimating the error contribution have been found to produce good results (see section(5.4)). It is interesting to note that at present the self-calibration technique is usually operated in such an interactive fashion with great success.

5.4 Simulations of Double Deconvolution

5.4.1 Software Implementation

In this section we consider the result of applying the double deconvolution algorithm to simulated MFS data from our standard ‘target source’ model. This particular implementation of the MFS philosophy will use CLEAN as the deconvolution algorithm and will use D_0 and D_1 as the convolving beams for each half cycle, as described in section(5.2.1) and analysed in section(5.2.4).

The DD algorithm was implemented within the NRAO AIPS software package. This particular software system was chosen because ‘user-friendly’ routines already exist which correspond closely to most of the steps of the double deconvolution algorithm, such as CLEAN deconvolution with a particular dirty image and dirty beam (task APCLN) and the subtraction of the effects of a set of CLEAN components from uv data (task UVSUB). Only one specialist piece of software was needed before the DD algorithm could be implemented as a succession of routines. This software was required to implement an exact convolution between a dirty image and the D_1 beam his

was done via the uv plane by multiplying together, uv point by uv point, the data sets corresponding to both the dirty image and the D_1 beam and then Fourier inverting the resulting uv data. A general task UVMAN was written by the author which has two inputs each of which is a uv data file. These should contain uv data points which are coincident in the uv plane. The task takes matching pairs of uv points and then manipulate in a variety of ways the two complex values from the two inputs to produce an output uv data point. One option within the UVMAN task is to output a complex multiplication of the two inputs; the task UVMAN run in this mode will then provide a mechanism for implementing the exact convolution described above.

The strategy adopted of using standard software as much as possible was chosen so as to minimise the amount of software effort needed in development and testing etc, whilst the practicality of the algorithm was still to be demonstrated. A disadvantage of this approach however is that the implementations was very inefficient with respect to speed and computer resources. It is not intended that this implementation be used in practice, only in the initial proving of the algorithm. In section(5.5) the future practical implementation of the method is discussed together with some suggestions as to variations of the basic algorithm.

5.4.2 The Simulation Target

The model source data used for the simulations is based on the target source model first presented in Chapter 1. This complex source should provide a rigorous test of the MFS technique; with the expected thermal noise of extended MERLIN, the dynamic range between the bright hotspots and the noise amounts to more than 2×10^4 and the emission at a brightness of several times the thermal noise occupies of order 6×10^4 beam areas. In Chapter 1 it was shown that under conventional imaging (with no phase errors) the reconstruction of this model source using data from the 8 telescope extended MERLIN is poor especially in the extended diffuse areas (see fig(1.3a)). These

problems are due to inadequate uv coverage. The uv coverage can be dramatically improved using MFS. The improved uv coverage should allow images of the quality of that shown in fig(1.3b) to be achieved. In realistic cases there will of course be spectral effects, the objective of DD will therefore be to substantially remove these effects. The image in fig(1.3b) will therefore be a target against which the success of DD can be judged.

In Chapters 2 and 3 we discussed at some length the effects of spatially varying power law spectra on MFS dirty images. It was shown that varying spectral indices across a radio source cause the addition of ‘spectral sidelobes’ on top of the normal sidelobes. If we attempt to process the resulting dirty map by CLEANing with the conventional dirty beam (see section(3.3)) the resulting image is to a reasonable approximation equal to the image that would be recovered in the absence of spectral effects overlaid by the spectral sidelobes. In fig(5.5a) is shown such a CLEAN map, comparing this to fig(1.3b) we can see that difference between the two images is very similar to the expected first order spectral contribution as presented in fig(3.8), namely the addition of features which surround the bright regions of the image where the values of $I\alpha'$ are large.

It is interesting to note that this image is already in some ways superior to the image that emerges from conventional mapping (see fig(1.3a)), Some real details within the extended lobes can be seen in fig(5.5a) which cannot be seen in fig(1.3a), the single frequency reconstruction. The effect of going from conventional imaging to MFS imaging is to exchange one type of image defect for another. It so happens that in this case the image defects due to spectral effects are less damaging than those due to poor uv coverage. The characteristic effects of inadequate uv coverage can be readily seen in fig(1.3a) including the breakup of structure within the diffuse lobes and the presence of extensive intersecting ‘stripes’ in the image. In contrast in fig(5.5a) the defects due to spectral effects consist of circular arcs rather than linear stripes.

In the simulation to be described the input MFS uv data to the double deconvolution algorithm consisted of the equivalent of five separate day long

observations at each of 5 frequencies (within a total bandwidth of 25%) with the extended eight telescope MERLIN. The frequencies chosen were the ‘standard’ frequencies used in the simulations of Chapters 2 and 3, i.e., 1350, 1435, 1540, 1615, 1710 MHz. The declination was chosen to be the standard 50° , all the observing parameters were therefore the same as in many of the simulations of Chapter 2 and 3. The uv data contained thermal noise at a level which would produce an expected map noise of $22.4 \mu\text{Jy beam}^{-1}$ for a 5 day observation. The data did not however contain realistic telescope based phase or amplitude errors. Note that recent estimates of the noise performance of extended MERLIN would make the array up to a factor of 2 more sensitive than was assumed in the simulations (R.J.Davis, private communication). It is assumed that before applying the double deconvolution algorithm the effect of the core had been successfully removed from the uv data at each observing frequency. It is also assumed that the data is corrected for an ‘overall’ spectral index α_o of 0.5 as described in section(2.3.2). This value of α_o was chosen to approximate the spectral index of the bright hotspot regions whose contributions dominate the spectral errors in the dirty map.

5.4.3 Control of the DD Simulation

As described in section(5.3.4) in the case that the the uv plane is not over-sampled by the MFS uv coverage it should be expected that the actual I and $I\alpha'$ distributions which will emerge from the double deconvolution algorithm will depend on the choice of the termination conditions of the half-cycles. It is within these termination criteria that the *a priori* knowledge about the source must be embodied. It is this *a priori* knowledge that must be used to separate the effects of the I and $I\alpha'$ distributions. We have seen that there are different ways of setting the termination criteria. The termination can either be set to be when the residuals in a CLEAN image reach a certain level or alternatively when the total area over which CLEAN components are found reach a preset level. Because the first of these methods is trivial to implement it was this method of control that was used in the simulations. In section(5.3.4) it was

shown that the size of the residual limit should be chosen to be approximately the level of the errors due to the unwanted distribution, the problem then becomes one of estimating this contribution. For the initial half-cycle the estimate of spectral errors based on the simple rules of Chapters 2 and 3 could be used. These rules predict a peak first order error of $450\mu\text{Jybeam}^{-1}$, therefore for the first (I CLEANing) half-cycle a clean limit of 0.5 mJy beam^{-1} was set.

The limits in subsequent I and I α half-cycles were set by examining the residual images and estimating the spectral error contribution away from the source by eye. The CLEAN limit was then set to be twice this value. This procedure is an informal version of the termination control method suggested in section(5.3.4). The values of the clean limits chosen in the various half-cycles are shown in table(2.1). This highly interactive method of controlling DD is sufficient to demonstrate the feasibility of the DD method, in a working implementation it may well be both desirable and possible to automate the selection of residual limits.

5.4.4 Adding VLA Data

As pointed out in section(1.6.2) the MFS technique is very good at filling in the outer parts of the uv plane, in the very centre of the uv coverage however there will be left a large hole. This hole will become the most significant defect of the uv coverage and can therefore limit the fidelity of the images possible. This defect will make itself felt in the absense of smooth large scale features with low spatial frequency. For the best possible images it is desirable to fill this hole by observing with the VLA in A-array and adding this data to the MFS data. Since this may be a common mode of operation for high fidelity imaging it was decided to include the effects of adding VLA data in the simulations.

The model VLA data consisted of three short observations (or ‘snapshots’) made at the reference frequency and spread over three hours.

After applying two cycles of DD to remove spectral effects with termination parameters as described in Table(2.1) the above VLA data were added and

Cycle No	Half Cycle Type	Critical Residual (mJy.Beam ⁻¹)	No of Clean Components	Total Flux (Jy)
1	I	0.5	11,520	8.139
„	I α'	7.0	3,260	3.33
2	I	0.25	30,102	8.095
„	I α'	4.4	3,153	2.411
3	I	0.25	29,342	8.378

Table 5.1: Half-cycle limits for MERLIN only DD test

Cycle No	Half Cycle Type	Critical Residual (mJy.Beam ⁻¹)	No of Clean Components	Total Flux (Jy)
1	I	0.5	12,032	8.112
„	I α'	3.0	2,075	2.216
2	I	0.30	30,056	8.092
„	I α'	1.4	6,792	2.602
3	I	0.27	36,392	8.405

Table 5.2: Half-cycle limits for VLA assisted DD

then the whole data set CLEANed to give the final MFS image. The resulting image is shown in fig(5.5b). The MFS map is superior to that produced by conventional single frequency methods (see fig(1.3a)). The processed MFS images show interesting low contrast filament features, whilst in the single frequency image these features are obscured by the presence of reconstruction errors.

5.4.5 The Performance of the DD Simulation

To decide whether DD has been successful we should compare the final image with that available before applying DD (fig(5.5a)), with the image possible when there are no spectral effects fig(1.3b) and the best single frequency image fig(1.3a) . From these comparisons we can see that the DD algorithm has been very successful in removing spectral effects and that the final image is very close to the ultimate that can be achieved in the absence of any spectral errors. The total improvement over the conventional single-frequency image in fig(1.3a) is quite dramatic.

To examine more closely the effects of the DD algorithm in reducing spectral effects we can compare fig(5.6a) which shows the original first order spectral contribution plus thermal noise to fig(5.6b) which shows the residual first order spectral error after the two cycles of double deconvolution. It can be seen that the algorithm has reduced the spectral contributions sufficiently that the residual effects are comparable to the thermal noise over most of the image, there exists only a faint hint of spectral sidelobes close to the hotspots. Overall there is a reduction of about 4 in the rms level of spectral errors. This reduction is sufficient to reduce spectral effects to be close to the thermal noise over the whole image. Since this test is close to being the most difficult problem with which we are likely to be faced with, this test is sufficient to demonstrate that the MERLIN MFS problem is essentially solved.

The quantitative performance of DD can be estimated by considering the difference between the accumulated $I\alpha'$ model and the true $I\alpha'$ image. The level of the distinct components in the $I\alpha$ map described first in section(3.2.2)






Figure 5.4: The model source. a) The intensity distribution at the reference frequency (1530 MHz). b) The spectral index distribution. Beam has a FWHM of 0.2".

a)



b)




Figure 5.5a: Results of CLEANing image containing spectral errors. Linear plot from -1 mJy.beam^{-1} to 1 mJy.beam^{-1} . Resolution $0.15''$.

b: Result after two cycles of double deconvolution linear plot from scale as above.

(Nb A small amount of VLA data has been added into the final deconvolution as explained in sect(5.5.4))

a)



b)




Figure 5.6: Comparison of the contribution from first order spectral errors.

a) Before applying DD.

b) After two cycles of DD.

Both images are linear from $-0.5 \text{ mJy.beam}^{-1}$ to $0.5 \text{ mJy.beam}^{-1}$.

can be compared before applying DD to their values after applying DD and so the reduction achieved can be quantified. The components with angular diameter 2''-3'' (i.e. components D,E,F,K,L,M) are found to be reduced by factors of 6 or 7 after two cycles but the level at a very compact hotspot such as component G is only reduced by a factor of just over 4. The reduction for the medium size components is close to the maximum possible given that the final CLEAN limit in the $I\alpha'$ domain is set at the comparatively high level of $4.38 \text{ mJybeam}^{-1}$. The reduction at the hotspot however falls short of the expected value based on taking the ratio of the peak $I\alpha'$ to this clean limit.

The difference between the performance for these two different types of component can be explained by the the fact that there may be some 'overcleaning' occurring when cleaning in the I domain. Because the initial residual limit was so low a very large area of the source was CLEANed and so it is possible that a significant fraction of the spectral sidelobes surrounding the hotspot has been cleaned into the I model. This is an example of the effect described in section(5.2.4) where, despite the fact that the residual is chosen to be equal to the expected spectral error, spectral sidelobes are 'CLEANed' in because they overlay bright regions. For the compact components the area CLEANed will be much larger in terms of the area of the component than is the case for diffuse components. Overcleaning will therefore have a much bigger impact on these compact components. Overcleaning in the I domain will reduce the size of the peak that occurs in the $I\alpha'$ domain dirty map, this effect will be worse for the more compact components. Therefore less $I\alpha'$ flux will be found in the subsequent deconvolution in the $I\alpha'$ domain and so the reduction in the fractional error in $I\alpha'$ will be small. One way to reduce the effects of overcleaning may be to start off with more gradual cleaning. However controlling DD by area limits rather than residual limits is the more logical way to tackle the root problem.

A second effect limiting the performance of DD is the high level of residual limit set in the $I\alpha'$ domain. Whilst the problem in the I domain is overcleaning, the problem in the $I\alpha$ domain is undercleaning. The large final residual limit

of 4.38mJy beam^{-1} prevents cleaning and sets an absolute minimum to the size of the $I\alpha'$ error. The reason why the residual limit is so large is that it is set to be equal to the size of the error due to the uncleaned flux in the other domain. This error contribution is dominated by low spatial frequency errors which are caused by insufficiently cleaning the extended structure in the I domain.

It is well known that CLEAN takes many iterations to fit the uv data in the centre of the uv plane, especially when there is a large hole in the central uv coverage. The data fit in the centre of the uv plane will be especially poor because only a limited number of iterations can be carried out in the I domain. This limited iteration CLEANing is of course essential to the convergence of DD because we must not transfer too many spectral errors into the I model. The resulting undercleaning of extended structure will however leave large errors in the $I\alpha'$ domain. These large errors will in turn limit the depth of CLEANing possible in the $I\alpha'$ domain.

Unaccounted for diffuse I components will leave especially bad errors in the $I\alpha'$ domain when the central multi-frequency uv coverage is poorly 'mixed'. Badly mixed uv coverage occurs when there is not much overlap in the uv plane between observations made at different frequencies on different physical baselines. When this occurs we get regions of the uv plane only sampled by the same physical baseline at different frequencies. The lowest frequency observations will be associated with the uv tracks of smaller radius and the higher frequency observations will give those of larger uv radius. An unaccounted for extended I component will give a feature in the uv plane which this type of uv coverage cannot distinguish from an extended $I\alpha'$ component. When we then convolve with the D_1 beam to recognise a $I\alpha'$ component the result will be a large contribution to the $I\alpha'$ domain.

5.4.6 Reducing the Effects of Undercleaning

How should we deal with the undercleaning problem? One method is simply to ignore the residual limit criterion and simply to CLEAN deeper, despite

the fact that this means CLEANing below the level of the extended structure due to uncleaned I. This will introduce errors into the $I\alpha'$ model but it can be shown that their subsequent effects on the I image will be quite small. The introduction of these errors can be justified if the deeper CLEANing allows other real $I\alpha'$ structures to be more readily found and so the eventual I image improved. Again it appears that the residual level method of control does not appear optimum in practice and other methods of control may be superior.

A way of reducing the contribution due to uncleaned extended I flux is to utilise short spacing single frequency VLA data. For the best quality images it will probably be advisable to supplement the MFS data with such short spacing data. In the simulations described in section(5.4.3) the data were simply added after the MFS data had been processed to remove spectral effects. This data can however also be used within the DD algorithm itself. The VLA data alone can be CLEANed and some corresponding model uv data formed. At each cycle of DD a MERLIN-only I image will be available and this can be likewise Fourier transformed to provide model uv data. An improved uv plane I model can be created at each cycle by forming a linear sum of the two sets of model I uv data with the VLA model dominating at short spacings and the MERLIN model dominating at long spacings. This composite uv model will not have 'missing' flux in the centre of the uv plane. On removing this composite model from the uv data at the beginning of a $I\alpha'$ cycle there will therefore be no problem due to extended I components. This will in turn mean that deeper CLEANing in the $I\alpha'$ domain will be possible.

5.4.7 Testing VLA Assisted Double Deconvolution

The new simulations used exactly the same uv data as those described in section(5.4.3). Again the data was assumed to be corrected for an 'overall' spectral index of $\alpha_o = 0.5$ and again it was assumed that the effects of the core component had been successfully removed from the data prior to processing. The same VLA data were also used. This data set was mapped and the resulting CLEAN components Fourier inverted to form an associated I

uv model, \bar{I}_{VLA} . The processing of the MFS data followed exactly the same pattern as before except that the uv model subtracted from the data prior to forming the $I\alpha$ domain dirty map was


$$\bar{I} = g\bar{I}_{MFS} + (1 - g)\bar{I}_{VLA} \quad (5.15)$$

where the subscripted quantities are the FT's of the images made from deconvolving the corrected MERLIN MFS data and the VLA only data respectively. The function g was a weighting that depended on uv radius. In these simulations g was taken to be a gaussian with a FWHM of 100,000 wavelengths.

The half-cycles were again terminated by setting residual flux limits, these limits being chosen by estimating the size of the contribution from the unwanted distribution as before. The details of the processing are given in Table(5.2). Note that the residual limit in the final $I\alpha'$ deconvolution is reduced from $4.38 \text{ mJy.beam}^{-1}$ in the previous simulation to $1.404 \text{ mJy.beam}^{-1}$. This lower level is allowed because of the reduced level of effects due to uncleaned extended I flux. The reduction achieved by using the modified algorithm can be seen by comparing fig(5.7a) and fig(5.7b) which show the residual $I\alpha'$ domain image at the end of the first cycle for both the original and the present simulations. In the original simulation the residual maps are dominated by low spatial frequency structure due to uncleaned I components whilst for the modified algorithm these are almost eliminated.

In terms of the visual quality of the resulting final image (not shown) it is hard to detect much difference between the results of the two simulations. This is because the major spectral effects were almost completely reduced to below the thermal noise even in the original test. Improvements are apparent however if the estimates of $I\alpha'$ are compared to the true values. The unaccounted for $I\alpha'$ flux at the hotspot G for instance is reduced from 34.37 mJy to 4.2 mJy , a reduction of a factor of 8.18 after 2 cycles. The general level of the first order spectral errors are likewise reduced by a factor of about 8, a factor of 2 improvement upon the performance of the MERLIN only DD test.

a)



b)




Figure 5.7: Effects of residual I in the $I\alpha'$ half-cycle residual map after one full cycle of DD.

a) Normal DD

b) VLA-assisted DD

Linear plots from -5 mJy.beam^{-1} to 5 mJy.beam^{-1} .

5.5 Future Developments

5.5.1 Improved Software Implementations of DD

The present software runs quite slowly and this has limited the amount of testing of the algorithm that was done on the then available hardware of a VAX 11/780. This slowness of the algorithm is in no way fundamental. Because in DD we always map only two distributions, viz I and $I\alpha'$ no matter how many frequency channels are used, we can expect that DD in principle should only take twice as long as would CLEANing an equivalent single frequency data set. It is certainly likely to be very much quicker than the sky plane based methods described in Chapter 4 in which each frequency channel must be deconvolved separately.

One reason for the slowness of the present software is that the method is implemented as a series of separate tasks within the AIPS environment. Once the method is coded into a single AIPS ‘task’ as a stand-alone program a large improvement can be expected because of the removal of overheads.

Another reason for the slow performance is that when, in every half-cycle, we deconvolve we must build up the new I or $I\alpha'$ image from scratch. Certain bright regions will repeatedly be CLEANed in every cycle, with very little difference from cycle to cycle. This duplication will take up a large amount of the running time. A similar effect occurs when combining CLEANing and self-calibration. Duplication of effort in this case can be avoided by using a ‘difference mapping’ strategy as in the OLAF program MAP. After carrying out a self-calibration step the effects of the best estimate model are first removed from the data before CLEANing. The CLEAN components found in the residual map are then used to update the model. The name ‘difference mapping’ is appropriate because we effectively map the difference between the corrected data and the best image estimate rather than the corrected data itself. An exactly analogous method can be implemented in DD. Before searching for flux in either the I or $I\alpha'$ domains the effects of *both* the current I and $I\alpha$ models can be removed from the data. This differs from the original method

in which only the effects of the undesired distribution are removed from the data. As in the original method the data will then be appropriately weighted so that the desired distribution will be emphasised and then CLEANed. The CLEAN components found will then be used to update the model of the desired distribution. The new modified 'Difference' DD algorithm is illustrated in fig(5.8).

A second class of important improvements to DD concerns the calculation and enforcement of the termination conditions of the half-cycles. The setting of these termination conditions should certainly be automated. At present the interactive setting of CLEAN limits is inconvenient especially given the slow speed of the software. Methods of termination other than the residual limit method should also be tested. Methods based on area limits seem natural given the heuristic idea of the mechanism of convergence presented in section(5.2.4).

Both the ideas of difference mapping and of automatic termination are incorporated into the schematic flow diagram for a practical DD presented in in fig(5.8). This implementation is based on the form of CLEAN first suggested by Clark(1980) and used in the AIPS program MX. In this algorithm the CLEANing is accomplished in blocks known as 'major cycles'. Within each major cycle CLEAN components are found by searching the residual map and approximately removing their sidelobes. When the size of the peak residual reaches a pre-determined limit the search for CLEAN components is stopped. The major cycle finishes by Fourier transforming the CLEAN components found and removing their effects from the uv data, i.e. effectively doing an exact removal of the sidelobes. A new major cycle then starts by forming a new residual map by Fourier inverting the residual data. A limit for termination is set and a new search phase begins. Note that internally MX is already in practice a 'difference' CLEAN.

The above algorithm can be modified so that as well as the normal limit a second peak residual or area limit can be set at which CLEANing in that particular domain should cease. When either of the two limits is reached

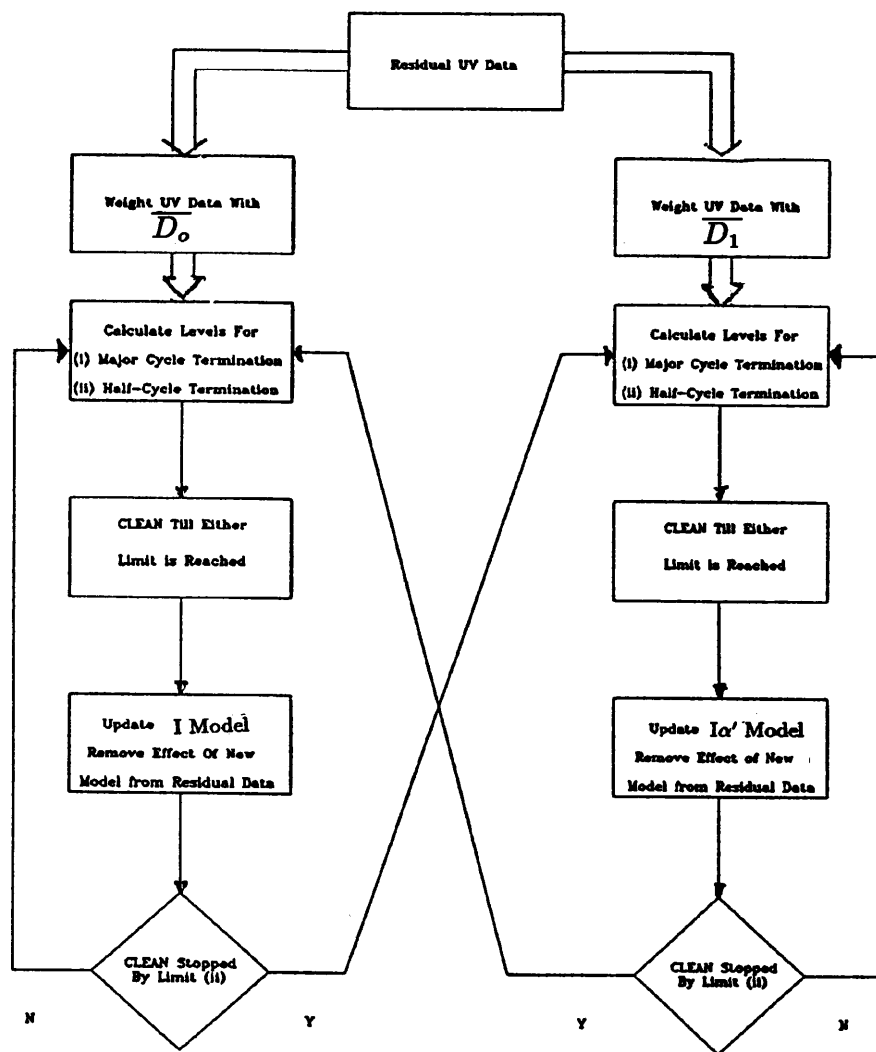


Figure 5.8: Flow diagram for a 'Difference Mapping' implementation of Double Deconvolution

the major cycle will finish. If the first limit was reached first then the next major cycle is started in the same domain but if the second limit was reached first the domain in which CLEANing occurs is switched.

5.5.2 Other Implementations of the DD Philosophy

Implementations of the DD philosophy other than that based on CLEAN and the D_1 convolving function should be investigated.

So far we have only considered recognizing the $I\alpha' * D_1$ term by convolving with the D_1 beam. This has been justified by considering CLEAN as a pattern recognition method, since the term we wish to find is convolved with this beam it makes sense to recognise it by finding the highest correlation between the dirty image and the beam. Other convolving beams (D'_1) may have advantages however and should be considered. One problem with using such beams with a CLEAN based DD algorithm will be that for guaranteed convergence CLEAN requires that the FT of the resulting effective beam which is convolved with the $I\alpha'$ distribution should be positive at all sampled uv points (Schwarz 1978). This may be difficult to arrange in practice.

The MEM deconvolution method (e.g. Narayan and Nityananda 1986) is one clear alternative to the use of CLEAN as the deconvolution component of the DD algorithm. So far no tests have been carried out using it and it would be interesting to see how it performs. Although it has advantages in that a wider range of convolving beams are allowed, there will be other restrictions. For instance, the overall spectral index α_o that is removed from the data (see sect(2.3.3)) must be such that the $I\alpha'$ image produced is strictly positive, because MEM can only deal with positive images. For some very complex images on which the MFS technique is likely to be employed (i.e. images of size 1024 pixels square and greater), the MEM method will be faster than the CLEAN method and this may be a factor in deciding to use this algorithm rather than CLEAN. Note that there is always the option of using MEM as the final stage of deconvolution on data from which the effects of an $I\alpha'$ model have already been removed by applying a CLEAN based DD.

Note finally the connection between MEM based ‘double deconvolution’ as described above and the ‘MEM methods’ briefly described in section(4.5). In this method a reference frequency image and a spectral index image are optimised so that the images produced at all the observing frequencies have the highest possible total ‘entropy’ consistent with the data at each frequency.

This ‘MEM method’ and an MEM based DD are similar in that both are iterative algorithms in which first the model reference frequency intensity image is improved and then an image defining the spectral behaviour is improved. In the first algorithm the $\alpha(x, y)$ image defines the spectral behaviour whilst in the second case it is a $I\alpha'$ image.

Under both algorithms the models are fitted to the data whilst seeking to fit *a priori* constraints, in both cases these *a priori* constraints are taken to be measures of entropy. The entropy measures for the spectral half-cycle are however taken over different images. In the DD MEM method the entropy of the $I\alpha'$ image is maximised subject to the image fitting the residual data. In contrast in one implementation of ‘MEM methods’ the α image is changed so that it is the sum of the entropies of the resulting frequency channel images that is maximised.

Chapter 6

Practical Aspects of MFS

6.1 Introduction

In this chapter we consider some of the practical aspects that must constrain any future MFS system on MERLIN.

The frequency bands in which MFS can be used will be strongly affected by the presence of RF interference. In section (6.2) the prospects for finding sufficient numbers of channels free from interference at each of the major MERLIN observing bands are considered. The critical levels of harmful interference are calculated and are compared against the results of interference surveys at L and P bands.

In section(6.3) some of the engineering problems of observing at many frequencies are discussed. The RF technology of the present MERLIN RF system is considered and the limitations that determine the present maximum fractional bandspread of 7% outlined. The changes required to produce an RF system capable of observing over the whole of L-band are described.

In section (6.4) the problem of phase self-calibration as it applies to processing of MFS data is considered. The simulations of MFS done so far in Chapter 5 have dealt only with data which possess perfect uncorrupted phase information whereas real observations are affected by tropospheric and ionospheric phase errors. Ways of incorporating self-calibration into the processing of MFS data are briefly considered.

Also considered is the possibility of linking the solution of the phase errors at the different observing frequencies (see sect(6.4.4)). In the mode of MFS observing in which we rapidly switch from frequency to frequency within an atmospheric coherence time the phase error at the different frequencies will be correlated. Using this correlation we can hopefully improve the accuracy

of the phase correction. Preliminary observations are presented in which such frequency switching is carried out and the expected degree of correlation is observed.

After phase calibration we next consider amplitude calibration (see sect(6.5)). It appears that less reliance on methods of amplitude self-calibration and more use of external amplitude calibration will probably be required if we wish to successfully use MFS on data which contain significant spectral effects. Although improvements in amplitude calibration are necessary there does not appear to be any fundamental reason why the sufficient accuracy should not be achievable if enough care is taken.

Lastly the effects on the MFS method of regions of the source with non power-law spectra are considered (see sect(6.6)), in particular we will consider the effects on MFS images of bright self-absorbed cores and bright hotspots which show the effects of ‘spectral ageing’.

6.2 Interference and MFS

6.2.1 Introduction

The radio frequency spectrum is a finite resource with many competing pressures on its use for communications, radar, broadcasting etc. The availability of bands with sufficiently clear frequency channels will be an important requirement for the success of an MFS system. In a country such as the U.K with a large pressure on the frequency spectrum such clear bands may be hard to find.

A very small part of the radio spectrum is specifically assigned to radio-astronomy use by national and international agreement. Most of the allocated frequencies have only a low level of protection, the only bands with full protection lie in the narrow range 1400-1427 MHz covering the important HI line and from 150.02 to 152 MHz. The operation of a 25% MFS system will obviously require observations outside of these narrow protected bands.

Of the present standard MERLIN frequencies of 151, 408, 1662 MHz and

5 GHz, MFS operating in the 408 MHz band (P-band) and 1662 (L-band) are of most interest. Although there have been problems at 151 MHz in reconstructing the steep spectrum emission from extended radio-galaxies and quasars (Stephens 1987) and so there is obviously a strong need for improved uv coverage at this frequency, practical experience of radio interference at this frequency leads to the conclusion that MFS observations would be impracticable (J.A.Battilana, private communication). Even observing within the nominally assigned band from 150-152 MHz has proved to be very difficult.

At 5 GHz the interference situation may be tolerable, however there is less pressing need for a MFS system than at L or P bands. Because extended sources are steep spectrum this tends to mean that the number of filled beam areas is larger at lower frequencies and so the need for MFS is greater. Most of the scientific interest shown so far for enhanced MERLIN at 5 GHz lies in imaging weak compact galactic objects which will not benefit much from MFS, (MERLIN phase 2 proposal). This does not mean that MFS will not be useful for extragalactic objects such as radio jets, only that MFS at 5 GHz is a lower priority than at lower frequencies.

The present plans for the RF system of improved MERLIN mean that any MFS system at 5 GHz will be limited to a 15 % bandwidth (set by the performance of the amplifier systems) (R.J. Davis, private communication). This will give a significant improvement in uv coverage at 5 GHz band (C band), although not as good as that achievable by a 25% bandwidth (see fig(1.7)).

The primary bands of interest for implementing MFS, at least initially, appear to be P and L. The planned RF capabilities of enhanced MERLIN should allow for a future MFS system which allows observations to be made over a 20%-25% frequency range in either of these bands. A bandwidth of this size is sufficient to 'fill in' the MERLIN uv coverage as described in sect(1.6.2).

In section(6.2.4) and section(6.2.6) respectively we consider the results of interference surveys at P and L bands and the impact of the interference environment on the prospects for MFS. Before examining these results we

first consider the different effects that interference can have on interferometer observations and the critical levels of each type of interference. These critical levels will be useful in evaluating the results of the interference surveys.

6.2.2 Effects of Uncorrelated Interference

This type of interference occurs when only one antenna on a given baseline receives the interfering signal. This form of interference is so-called because the interfering signal will not be correlated with the noise signals nor the astronomical signals received by the unaffected antenna.

The effect of uncorrelated interference is simply to increase the total power received at the affected antennae. Arrays such as MERLIN find the amplitude of their visibilities from the measured correlation coefficients by assuming a known antenna noise. These antenna noises are most easily found by making calibration observations of point sources with known flux. The effect of additional noise from interference is to apply an antenna based *multiplicative* error to the amplitudes of the visibilities on baselines in which the affected antenna participates.

We assume that the interfering signal enters one of the the far sidelobes of an antenna so that the gain of the antenna is equal to the effective isotropic gain (Thompson, Moran and Swenson 1986). The power dissipated at the frequency of observation will then be equal to that which would be dissipated in an approximately isotropic survey antenna. If for the sake of argument we set the criterion for significant interference to be an interference power of 1% of the thermal noise power then this gives us a limit for allowed interference F as measured by the survey antenna of

$$F(Watts) = 0.01kT_{sys}\Delta\nu \quad (6.1)$$

where T_{sys} is the system temperature and $\Delta\nu$ is the observing bandwidth. If at L-band $\Delta\nu = 16$ MHz and $T_{sys} = 35$ K this gives us a limit for F of 7.7×10^{-15} W or -131 dBm (dBm is equal to decibels of power in units of milliwatts). At P-band if we have $\Delta\nu = 5$ MHz and $T_{sys} = 75$ K then $F = 1.66 \times 10^{-15}$ or

-127 dBm. Note that if the above levels of uncorrelated interference occur on one antenna of a baseline then this results in a 0.5% change in amplitude on that baseline.

If the increased power from an interfering signal were constant then it could, in principle, be calibrated out. In general however the size of the interference effect will vary as the telescope tracks a radio source and the source of interference moves through an antennas sidelobes. Forms of continuous amplitude calibration such as noise diode calibration (see sect 6.5) could in principle calibrate modest time variable amplitude errors such as these. At present methods of amplitude self-calibration are often successfully used to remove the effects of telescope based time variable amplitude errors of up to 5% (though a few percent is more normal). For high fidelity imaging extensive amplitude self-calibration should however be avoided. When processing MFS data it appears that amplitude self-calibration should be avoided even more because of possible interactions between self-calibration and MFS algorithms (see sect 6.5).

6.2.3 Effects of Correlated Interference

Correlated interference will only occur when the same signal is received by both of the antennae that participate in a particular baseline. The effect of this interference is to introduce an *additive* error into the data.

Correlated interference error is more serious than the uncorrelated variety. Whilst moderate amplitude errors can be subsequently corrected, data affected by significant additive errors comparable to the thermal noise must normally be deleted. Further, because of its correlated nature such interference will be averaged coherently as compared to uncorrelated interference which will be averaged incoherently, this means that much lower levels of correlated interference are sufficient to ruin observations. The sensitivity to interference is not as great as with single dish antenna however. First, the effect of the interference is reduced because the source does not move at the sidereal rate and so is averaged with a significant phase rate. Secondly interfering signals

of significant bandwidth have a reduced effect because the delays introduced into the arms of the interferometer to equalize the signal paths from the source will cause the signal path for signals from other directions to be unequal (with effective time delay $\Delta\tau$ so causing a decorrelation if $\Delta\tau > \Delta\nu^{-1}$

We will assume, as before, that the relevant gain is the isotropic gain of the antenna. The criterion for significant interference will be taken to be that the effect in the dirty map of 24 hours of observations affected by interference on a particular baseline introduces an error of less than 10% of the thermal noise, so that the resulting images remain noise dominated. Following Thompson, Moran and Swenson(1986) we obtain the critical interference power level (assumed equal at the both antenna of the affected baseline) of

$$F_{int} = 0.1kT_{sys}(2\Delta\nu\nu_e q)^{1/2} \quad (6.2)$$

where q is the mean length of the baseline in wavelengths and ν_e is the angular frequency of rotation of the Earth. The above equation successfully takes into account the effect of the non-sidereal phase rate for declinations $\delta \leq 70^\circ$. For higher declinations the phase rates of the source and the terrestrial interference will become more and more similar because a stationary interference source appears to come from the NCP. When observing at such high declinations the interferometer will therefore be more sensitive to correlated interference.

If we assume the worst case of narrow band interference (less than a few kilohertz) then there will be no additional reduction for temporal decorrelation on top of the above phase rate decorrelation. From equation(6.2) we find that for the Lovell-Defford baseline of MERLIN at L-band, the critical limit for correlated interference is 16 dB more stringent than that for uncorrelated interference at -147 dBm; at P-band the limit is 16.4 dB more stringent at -143.4 dBm.

Unlike the case of uncorrelated interference there is no simple procedure for reducing the effects of correlated interference in data, the only effective way of reducing the effects of such interference is to remove the portions of the data most affected. It is generally found that correlated interference will have its largest amplitude near the v axis where the sidereal phase rate is low and so

it is possible to significantly reduce the effects of this type of interference by deleting a band of visibilities along this axis. This however is not a desirable procedure to apply to MFS data in which the improvement of uv coverage is the aim.

6.2.4 Results of Interference Surveys in L-band

In order to see if the interference environment in L band would be tolerable for MFS, interference surveys have been conducted by Holmes and Woodall (1987). Surveys were carried out with a simple antenna at three of the MERLIN sites, i.e. at Jodrell (from the top of a tower of the Lovell telescope), and at ground level at the Knockin and Defford sites. The measurements were carried out in November 1987. The results of these surveys plotted as interference strength against frequency are shown in fig(6.1) and fig(6.2).

The survey sensitivity limit was -125 dBm. According to earlier calculations this limit corresponds to an amplitude error due to uncorrelated interference of 2%. The above survey is therefore almost sensitive enough to rule out significant uncorrelated interference. The sensitivity limit is however inadequate for detecting significant correlated interference at levels above the critical limit of -143 dBm.

In addition to the above survey which covers the whole of L-band at low sensitivity, high sensitivity data from AOS (Acousto-Optic Spectrometer) observations are also available, however these data only cover a narrow range of L-band. The results of AOS observations made with the Lovell telescope at frequencies around 1420 MHz (kindly provided by G.C. Simpson) are presented in fig(6.3). The spectrum in fig(6.3) is formed by taking the difference between beam switched observations. Easily seen in fig(6.3) is local HI (at 0 kms^{-1}) and the emission from the target source (at 2000 kms^{-1}) but there also several interference signals. The recorded level of interference is of course the difference in signal between the levels of interference at the two pointing positions. If it is assumed that the interference enters through a far sidelobe




Figure 6.1: Plots at L-band interference strength against frequency at three MERLIN sites. After Holmes and Woodall (1987).

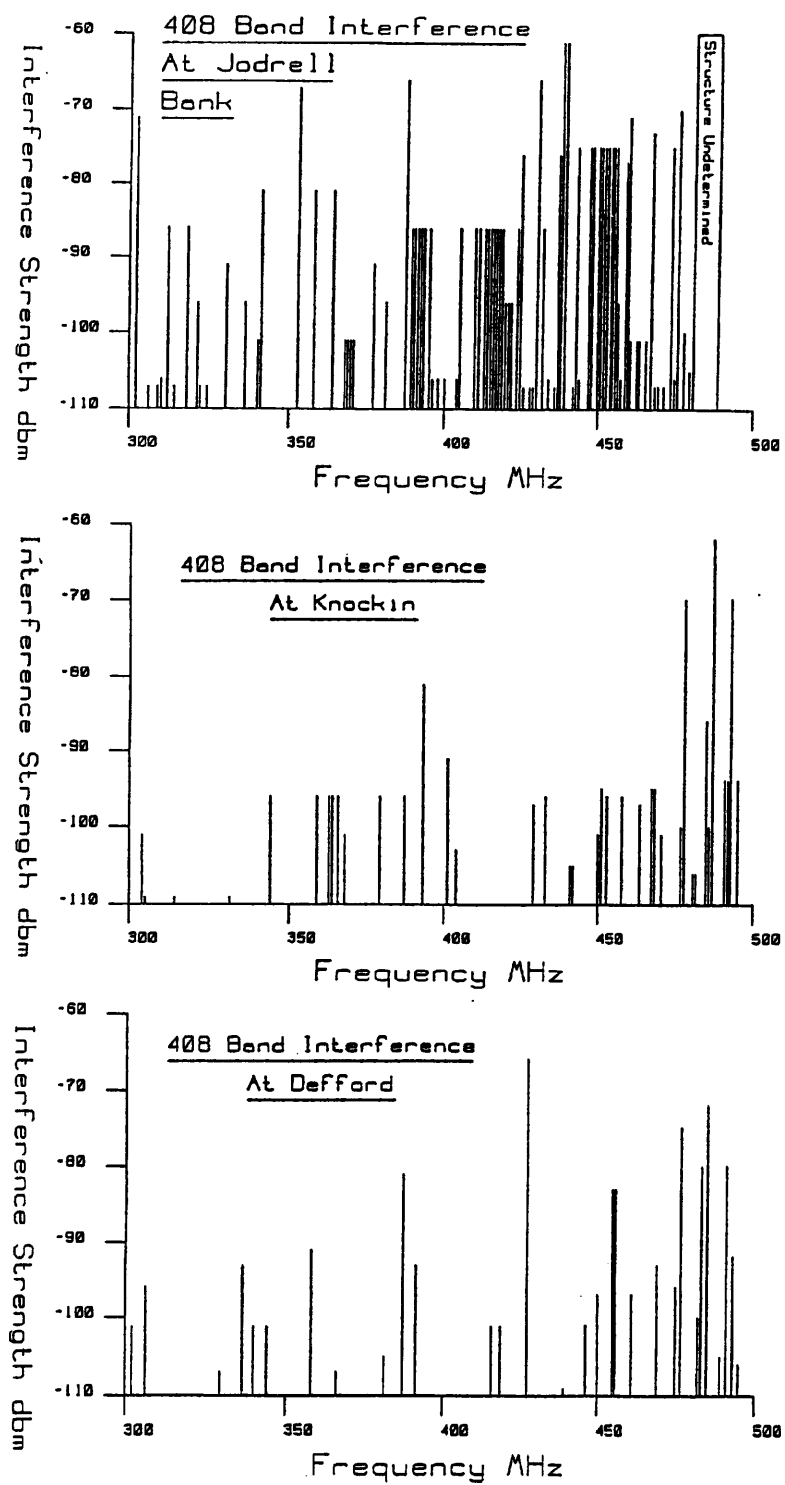


Figure 6.2: Plots at P-band interference strength against frequency at three MERLIN sites. After Holmes and Woodall (1987).

and the position shift is much larger than a sidelobe spacing then this difference should be comparable to the actual level of interference. Although the frequency range of these spectroscopic observations are only of order 80 MHz they are very sensitive, being able to observe signals down to levels below -175 dBm, so allowing us easily to detect signals that would cause correlated interference during interferometry observations.

The results of the interference survey of Holmes and Woodall (fig(6.1) and fig(6.2)) show that the interference environment is apparently worst at the Jodrell site of the three sites surveyed. This is probably because Jodrell Bank is nearest to large centres of population. In addition the Jodrell survey, unlike the others, was taken from the top of the telescope, and so it has a much more distant horizon for interference. It is also well known that radio observatories are often their own worst sources of interference.

Most of the interference found at the three sites lies in three bands 1300-1365 MHz, 1437-1530 MHz and 1700-1800 MHz. According to the assigned frequency allocations (fig(6.4)) these three bands contain respectively radars, private FM links and communication channels assigned to the emergency services. Of the private FM links note in particular two frequencies. First at 1497 MHz there are the L-band links operated by Jodrell to communicate LO phase information to the antennae (Davies, Anderson and Morison 1980), these signals are naturally present at all sites. Secondly at the Jodrell site alone there is a very strong signal at 1458 MHz (-81 dBm) from an L-band link operated by the UKAEA whose line of sight passes very close to Jodrell Bank. Interfering signals of such large power will cause special problems if they are able to reach the receiver electronics and special filtering will be needed to deal with such signals (see 6.3.2). Large signals comparable to the above also occur in the emergency services band at 1766 MHz at Knockin (-80 dBm) and in the radar band at Jodrell at 1314 MHz (-89 dBm). From the above survey it appears that interference found at one site is usually absent at other sites and therefore its effects are those of uncorrelated interference. Even when the interference is shared at different sites (i.e. at 1305,1340,1470 and 1671 MHz) it is impossible




Figure 6.3: Beam switched Acousto-Optic Spectrometer observations made on the Lovell telescope.

civil aviation radars
1215-1365MHz

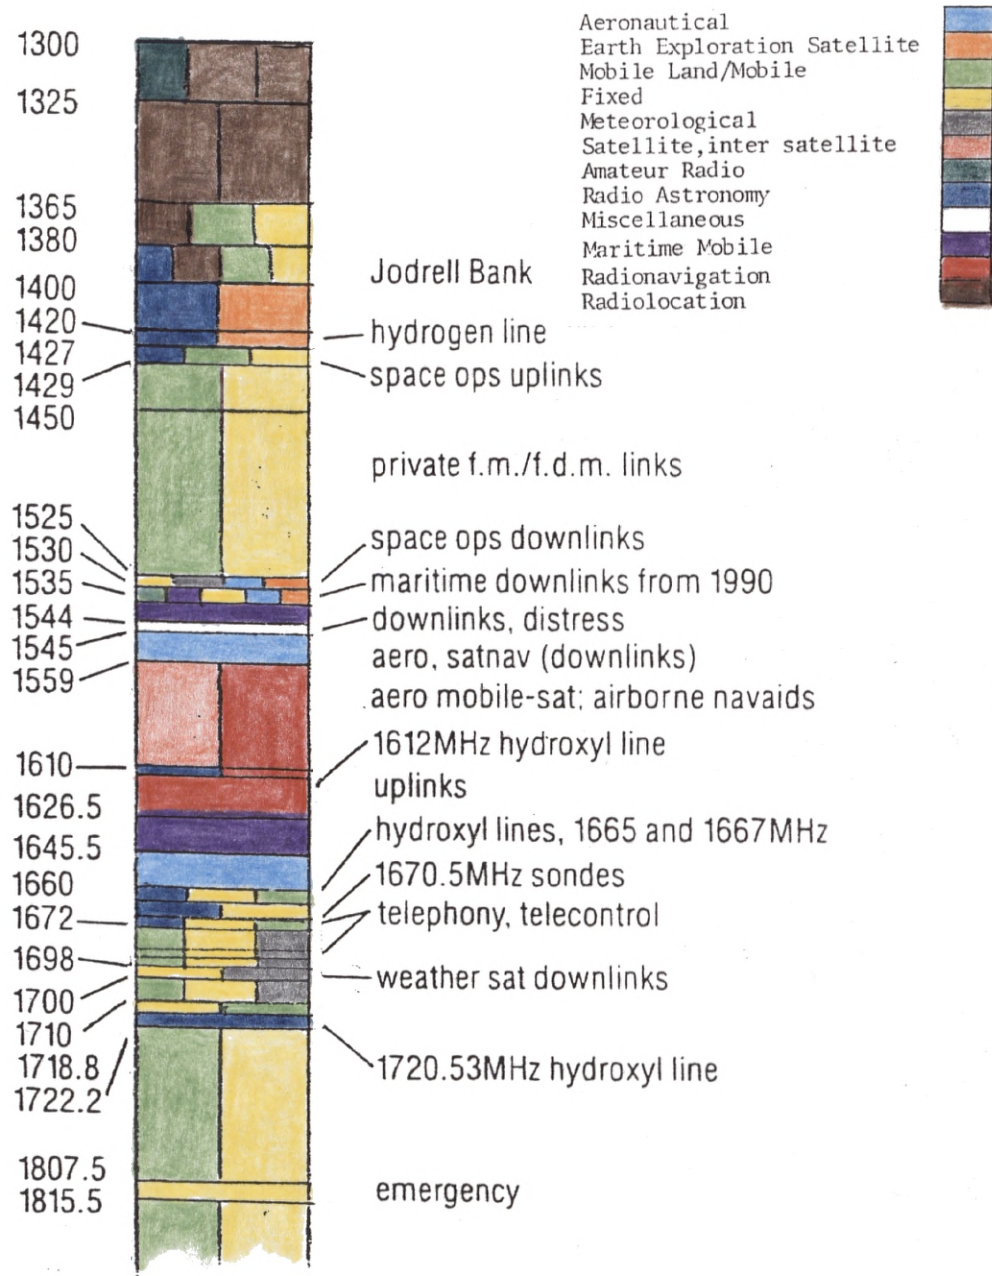


Figure 6.4: Frequency allocations in the U.K. around L-band.

to tell from the survey whether the interference is correlated or simply due to two separate sources of interference radiating at the same frequency. It is quite plausible that the widely separated sites which participate in the long baselines such as Jodrell, Knockin and Defford will see only uncorrelated interference because these telescopes do not share common horizons (exceptions being interference from satellites which will be visible from all antennae and also possibly interference from aircraft). For baselines connecting the above telescopes it is therefore likely that the low sensitivity survey provides an almost comprehensive description of the significant interference to be avoided when planning an MFS system.

For the shorter baselines involving the nearby antennae at Jodrell, Wardle, Tabley and Darnhall it might be that at least some of the interference will be similar to that found in the Jodrell Bank survey and there may be correlated interference effects. Despite the fact the the full L-band survey was not sensitive enough to rule out significant correlated interference there are grounds for optimism. Observations (of 10MHz bandwidth) made at 1604, 1632, 1662, 1690 and 1720 MHz with the present MERLIN (see chapter 7) show no sign of any correlated interference, this is in a frequency range in which the survey shows a significant number of high powered interference signals (fig(6.1)). The results of AOS spectra (fig(6.3)) taken with the Lovell telescope likewise show no significant signals in the range 1370-1440 MHz. If these two regions are typical of the 'empty' regions found in the low sensitivity survey then there should be no problems with MFS observations in these parts of the L-band.

6.2.5 Implications of L-Band Interference for MFS

Figure(6.5a) shows the sum of the interference plots at the three sites surveyed. Using this plot frequencies within L-band at which clear 16 MHz regions exist (the proposed bandwidth for the improved MERLIN) can be identified, these will be possible frequencies at which we could site our operational MFS channels. Below 1350 MHz radars prohibit observations whilst L-band links rule out MFS observations from 1430 MHz to 1550 MHz.

The prohibited frequency range from 1430-1550 MHz will determine a maximum useful number of frequencies for MFS. Because the quality of uv coverage is determined by the largest gaps in the uv coverage it will not be worthwhile to sample the frequency spectrum too densely away from the prohibited region. The prohibited region has a width of 120 MHz, this leaves 280 MHz free in the range 1350-1750 MHz, If the channels are spaced 50 MHz apart this gives a maximum of 6 or 7 channels. Possible sites for a 5 frequency MFS system would be at 1350,1430,1560,1640 and 1740MHz, this set includes two frequencies at or near to frequencies allocated to radio astronomy use (1420 and 1662 MHz). Note that the final decisions about which frequencies to use with MFS on the extended MERLIN, which is planned to include antennae at Cambridge and Chilbolton (in the South of England) will also depend on the interference environments at these sites.

The future availability of frequencies which are presently free is uncertain, the only certainty is that pressure on the radio spectrum will increase. The bands 1350-1430 MHz and 1620-1750 MHz are allocated to fixed and mobile land communications and there are plans for considerable expansion in this application (DTI annual report 1987).

The band 1559-1626 MHz is partly assigned to satellite communications and satellite radionavigation. The Soviet GLONASS position finding system is already operational at frequencies centred near 1610 MHz. Although it greatly disrupts single dish observations (R.J. Cohen,private communication) it has little impact on interferometry because its signals are of the 'spread spectrum' type (Thompson 1982) with wide bandwidths. Spikes in the the spectrum of order 2MHz wide are the only possible source of interference to interferometry. Even these are very substantially decorrelated because of delay decorrelation and the effects of fringe rate. No interference effects on MERLIN observations from GLONASS have been found when observing in the frequency range 1604-1720 MHz (see chapter 7). A very similar system operated by the U.S called NAVSTAR/GPS is centred at a frequency of 1575 MHz.

Navigation satellites do not seem to be a very important threat to MFS




Figure 6.5: Plots of total a) L-band and b) P-band interference summed over the three MERLIN sites surveyed in fig(6.1) and fig(6.2). After Holmes and Woodall (1987)

because of their wide bandwidths. Other satellites operating with narrow bandwidths in the band between 1559 MHz and 1626 MHz might however cause significant correlated interference as might weather satellites in their allocated band. The operational METEOSAT system for example will use 1686 MHz. Such systems may be an increasing threat to MFS.

The survey of Holmes and Woodall(1987) has shown that the extent of L-band interference is not as bad as was feared by some. Combined with the experience of successfully observing in the range 1604-1720MHz and of spectroscopy around 21cm this survey indicates that the prospects for MFS in L-band are good at present. To confirm this assessment rigorously will require surveys at least 20dB more sensitive to detect possible harmful correlated interference. To achieve such sensitivities long integrations at each of the sites will be required. Another priority for the future is to carry out interference surveys at more MERLIN sites. Surveys at Cambridge and Chilbolton are particularly important because the interference environments might be quite different at these sites.

6.2.6 Results of Interference Surveys Around 408 MHz

Given the low sensitivity of the interference surveys in the 408 MHz band (fig(6.2) and fig(6.5a.)) it is hard to make definite statements about whether enough frequencies will be sufficiently free to allow MFS. As with the L-band survey the Jodrell site again shows the most interference and will largely determine whether MFS is feasible or not.

Except for the partially protected band at 408 MHz the survey indicates that observations above 380 MHz would be very difficult. From 400-470 MHz the band is allocated to a wide variety of different uses including radio-positioning, personal mobile radio, radio paging and telemetry and control. At frequencies higher than 470 MHz the band is allocated to TV transmissions. Frequencies below 400 MHz are used for satellite, mobile and fixed communications including mobile to satellite communications and aircraft landing systems.

A set of 5 frequencies from 327 MHz at one extreme to 408 MHz at the other could be fitted into the gaps in fig(6.5b) if bandwidths of 5 MHz are assumed. Such a set of frequencies would form a useful MFS system with a 25 % bandspread. Unfortunately the present and proposed future feed systems have their performance optimised for 408 MHz and are unlikely to have sufficiently good performance at the lower end of the band. Observations in the higher half of P-band might be possible if narrow bandwidths (i.e. 1 MHz) were used although it is technically very difficult to provide the necessary RF filtering to exclude the very strong signals found there.

6.3 RF Technology and the Implementation of MFS

6.3.1 The Present L-Band MERLIN RF System

The present L-band RF system is presented in diagrammatic form in fig(6.6). The signal from the radio source enters the feed of the antenna, encounters a quarter wave plate (QWP) to separate the polarizations which are then transferred from the feed system to wires by means of an ortho-mode transition (OMT). The signals from the two polarizations then follow separate paths each passing in turn through a pre-amplifier, an isolator, a RF filter and another isolator before being mixed with a suitable local oscillator frequency. Each signal is mixed with an appropriate LO signal. The L-band LO signal is generated by taking an input originating from the MERLIN frequency synthesiser at around 454 MHz, which is phase-locked to the station frequency standard, and then multiplying this by a factor of 4. After mixing the signals undergo another stage of amplification before being low-pass filtered. The signals are now in a form ready to be sent over the microwave link to the MERLIN correlator.

The present system has been used for narrow bandspread MFS observations as will be described in Chapter 7. The range of frequencies accessible

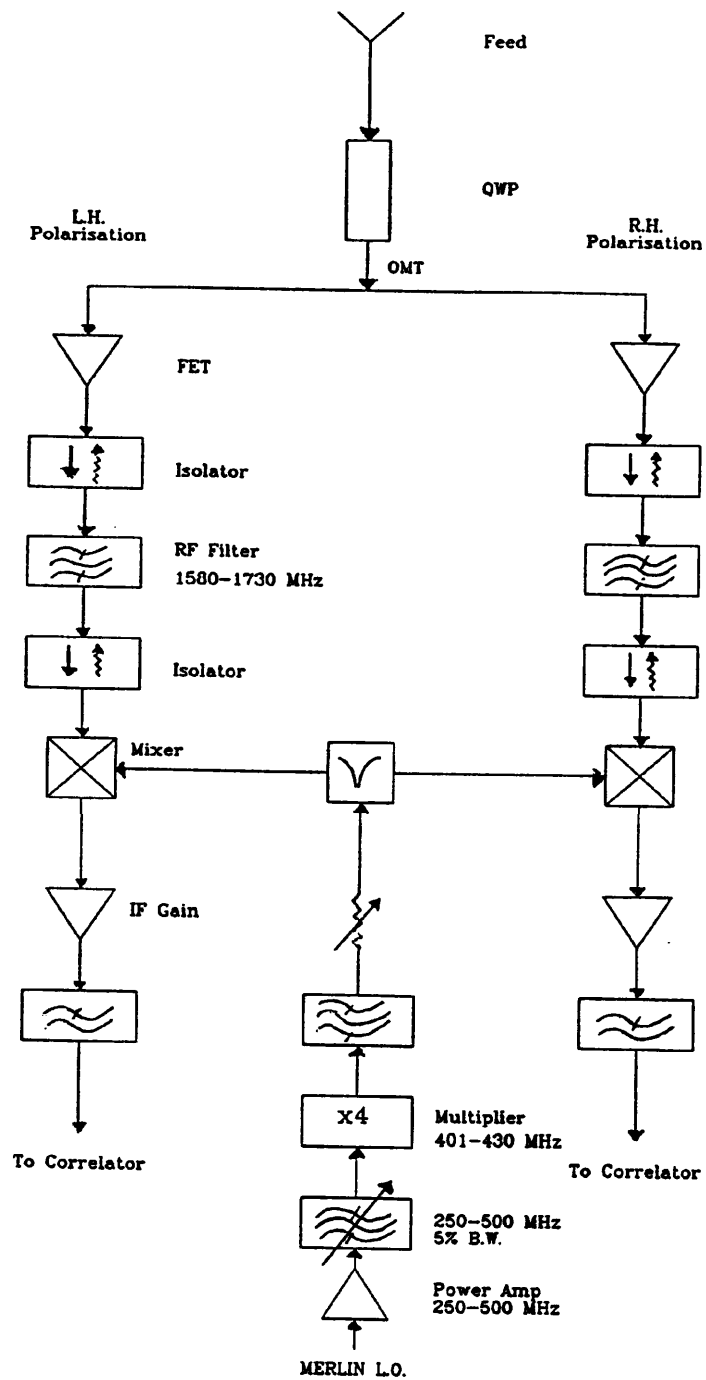


Figure 6.6: Diagram of the present MERLIN L-band RF system

with the above system is limited to the range 1604-1720 MHz, a fractional change of less than 7 % which only allows a modest improvement in uv coverage. In Chapter 1 it has been demonstrated that only by having bandwidths greater than about 15% can dramatic improvements in uv coverage be achieved. The fractional bandwidth is at present limited by the frequency performance of the front end including feed, QWP, OMT and pre-amplifier, by the bandwidth of some filters but most of all by the limitations in the present LO system. The major problem is the use of the frequency multiplier, this device operates only over a narrow range of frequencies. It is this component that presently limits the usable range of frequencies to the range 1604 MHz to 1720 MHz.

At present the MERLIN RF system is undergoing substantial development and improvement as part of the general enhancement of the MERLIN system. The first part of this improvement is presently being implemented. In phase 1 the improvement that most directly affects MFS is the replacement of the present front end including feed, QWP, OMT and pre-amplifier with a system that can operate efficiently over the range 1350-1750 MHz. Observations over this range will not be possible however till the present LO system and RF filters are replaced. These changes are part of the second phase of the improvements. If all goes well however a wide bandwidth MFS system could be operational by mid-1990 (P.N. Wilkinson, private communication).

6.3.2 LO Generation and RF Filters for MFS

The new LO system in L-band will use an electrically tunable YtYAG oscillator (M. Bentley private communication) that can produce signals directly in L-band without the need for multiplier stages. Using this device the LO frequency will be tunable across L-band simply by changing the magnetic field across the oscillator. To achieve the necessary phase stability the oscillator will be linked to the MERLIN synthesiser output by the use of a phase locked loop. The changes in the LO system are incorporated into the diagram of the new RF system presented in fig(6.7).

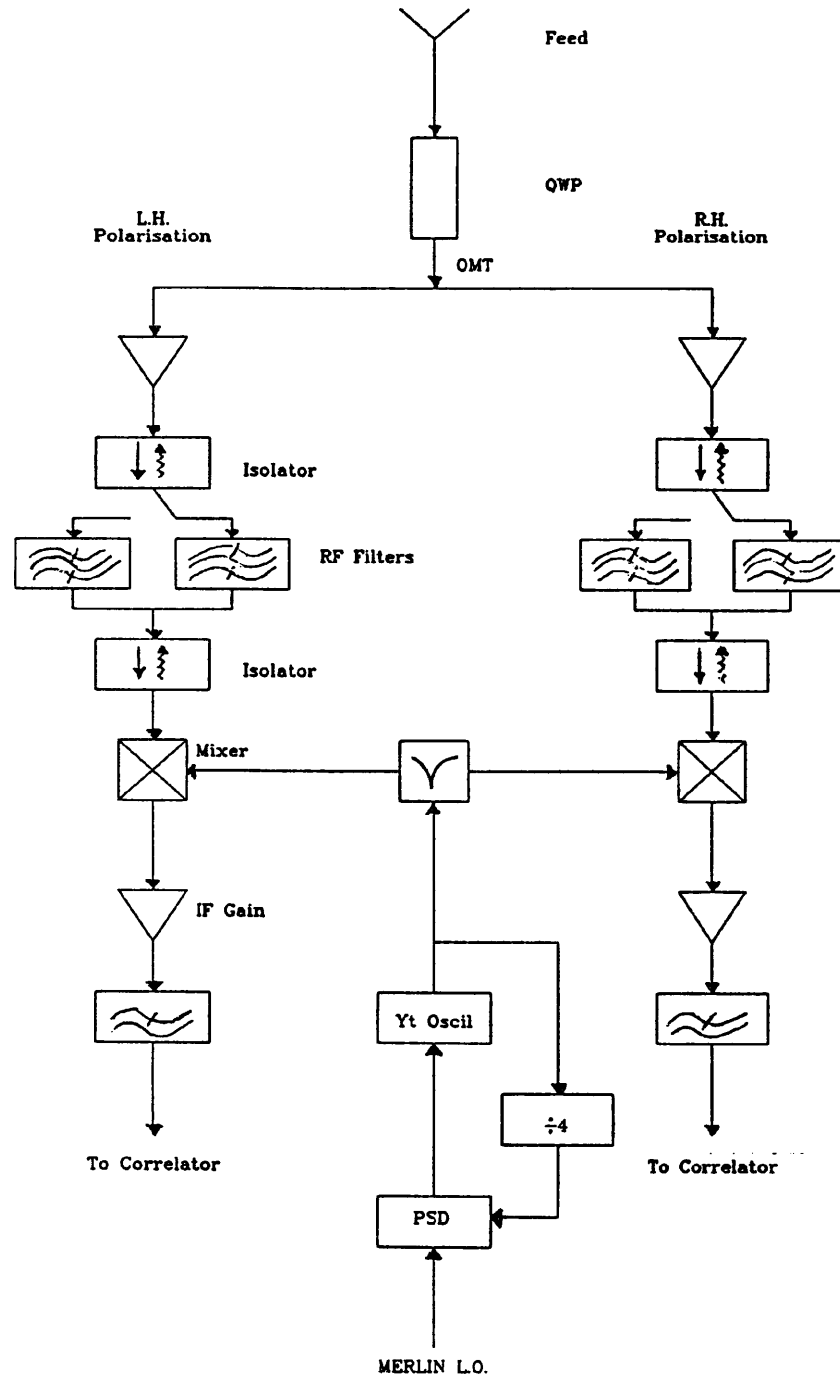


Figure 6.7: A proposed MFS L-band RF system for the enhanced MERLIN.

The other vital components required are suitable RF filters at each antenna. It has been found that if very strong interference signals outside of the observing bandpass enter and get as far as the mixer and the second stage of amplification then they can cause serious problems. The strong signals can saturate the amplifiers and generate harmonics which will effect observations within the chosen bandpass. Such strong interference presently occurs at Jodrell at 1458 MHz(-81 dBm). Largely in order to exclude this particular frequency the present system has an RF filter with a passband from 1580-1730 MHz early in the signal path. For successful observations over the full L-band some scheme of RF filtering is required which allows access to all the MFS frequencies yet ensures the very strong interference is filtered out. According to the L-band interference survey there are at present 3 particular frequencies to be avoided i.e. 1314, 1458 and 1766 MHz, each corresponding to the worst example of interference in each of the three main interference bands. In addition the whole frequency range from 1430-1550 MHz which contains many strong signals is probably best filtered out. At the very least the RF filtering should allow us to avoid the particular areas of the spectrum described above, at best it should be flexible enough so that it would be possible to deal with any future strong interference that might arise. There are three possible ways to achieve the necessary RF filtering (J.A. Battilana, private communication)

(i) Use electronically tunable narrow band RF filters that only allow through the 16 MHz corresponding to the observing bandwidth. (ii) switch between a battery of fixed narrow band filters (with 16 MHz bandpass) at each antenna, each filter being centred on one of the chosen set of MFS frequencies. (iii) switch between several very wide (100-200 MHz) filters which avoid the worst signals.

Methods (i) and (ii) provide the best protection since only the necessary band is allowed to propagate through the receiver electronics. Method (i) is the most flexible since it is not tied to a preset set of observing frequencies. Unfortunately it appears technically impossible to produce electrically tunable filters that will return repeatably to precisely the same centre frequencies at all

the different antennae. Such bandpass mis-matching would cause serious non-closing phase and amplitude errors (Thompson and D'Addario 1982) which are not correctable with self-calibration methods. Method (ii) avoids the above problem but it lacks flexibility. The main drawback of method(ii) however is the cost, such a system requires one filter per frequency per antenna.

Method (iii) is the simplest and the cheapest. With the present distribution of strong interference, problems can be avoided if we have two wide passbands covering the ranges 1320-1440 MHz and 1530-1750 MHz respectively. This solution is shown incorporated into the RF system in fig(6.7) with an electronically controlled switch which will allow observations to be made either in the lower or upper halves of L-band. Whilst the above system will suffice at present, if very strong interference were to occur at other frequencies then it might be necessary to have more, narrower, RF filters till a system like (ii) emerges.

What about implementing MFS in P-band? Because the input phase signal from the MERLIN synthesiser is itself in P-band there is no need for multiplier systems to generate the LO. Only simple modifications will be required to produce an LO system tunable over 25 % in frequency. Whilst the LO problem is simpler than at L-band the RF filtering problem is however much worse because of the much higher levels of interference. The only practicable filtering scheme appears to be type (ii) with many narrow channels.

6.4 Phase Self-Calibration and MFS

6.4.1 Introduction

So far in this thesis we have considered MFS data which is uncorrupted by the effects of phase errors. Such phase errors are caused by propagation effects as the radio signals pass through the ionosphere and troposphere. In order to produce high quality images data affected by such errors must be processed by so called 'self-calibration' methods which substantially remove the effects of these phase errors (for a review see Pearson and Readhead 1984). These

methods utilise the fact that the phase errors on any particular baseline can be considered to be equal to the difference in phase above each antenna i.e. they assume that phase errors are ‘antenna factorizable’. Self-calibration is a mandatory part of the processing of uv data to produce high quality images and so it is important to consider how it fits into the processing of MFS data.

6.4.2 Self-Calibration and Reconstruction Errors

The simplest possible implementation of self-calibration with MFS data is illustrated in fig(6.8a). In this ‘separate’ self-calibration each individual frequency data set is separately self-calibrated in a conventional manner. The data from the different frequencies is then combined and a MFS algorithm applied to eliminate spectral effects and produce a high fidelity intensity image at the reference frequency. The reasoning behind this implementation is that the self-calibration and spectral error removing steps are best separated to stop any possible interactions between the algorithms.

Whether the above method is acceptable or not will depend on the size and nature of the errors that are introduced by the single frequency reconstructions that accompany the separate self-calibrations at the different frequencies. In considering such errors it is useful to consider the Fourier transform domain and to compare the FT of the final image and the FT of the true source distribution. In anything other than imaging simulations exact knowledge of the true visibility will of course be unknown, yet in classifying the errors it will be useful to consider the difference between these two distributions.

Single frequency image reconstruction errors can be divided into two types. In the first type the model uv data differs from the true uv data only in regions of the uv plane not sampled by the uv coverage , i.e. we get ‘interpolation’ errors in the uv plane. The interpolation errors which occur when CLEAN is used as the deconvolution method tend to have the form of ‘ripples’. The above errors occur even in situations in which no self-calibration is applied, they occurred for instance in the CLEAN deconvolutions of perfect phase data presented in section(4.2.2).

The second type of reconstruction error differs in that the model uv data and the true uv data do not agree even on the sampled uv points. These errors occur because the phase corrections found and applied in the self-calibration loop are not exactly correct. This second type of reconstruction error will be referred to as 'self-calibration based' although in fact they are caused by complex interactions between the deconvolution and self-calibration algorithms. The complicated nature of this second type of error means that it is almost impossible to characterise analytically. Reconstruction errors of this type can however be easily distinguished from 'interpolation' errors because they cause errors that are 'lobe-like' rather than 'ripple-like'. These errors are more serious than interpolation errors because, whilst the former errors can easily be recognised as artifacts, 'lobe-like' errors can be mistaken for real structure.

In situations in which 'interpolation' reconstruction errors dominate single-frequency images the simple self-calibration scheme outlined above will be sufficient. Adequate self-calibration of the data is achieved by separately processing the data at the different frequencies. The interpolation errors do not propagate to the MFS stage and when the data are deconvolved with the MFS uv coverage, interpolation errors will be dramatically reduced.

When the single-frequency maps are dominated by 'self-calibration based' errors then the above method will probably not be adequate however. The errors accumulated in the single frequency reconstruction/self-calibration process will be carried through into the MFS data set. The fidelity of the MFS reconstruction will therefore be tied to that of the single-frequency reconstructions. The best that can be hoped for from MFS is a statistical reduction in the final level of reconstruction error on combining the data from the different frequencies.

It is notable that whereas for the VLA images 'interpolation based' ripple errors are common on MERLIN images they are not as prominent, instead 'lobe like' errors dominate (see fig(7.6)). The difference can probably be explained by the much larger number of antennae in the VLA making self-calibration more reliable. At present MERLIN images of complex sources

usually seem to be dominated by ‘self-calibration’ based errors so that ‘separate’ self-calibration doesn’t seem a very good approach.

6.4.3 Parallel Self-Calibration

If MERLIN single-frequency reconstructions are indeed dominated by ‘self-calibration’ errors then the dramatic improvements in fidelity seen in going from SF (single frequency) to MFS imaging in chapter 5 may not be achievable in practice by using simple ‘separate’ self-calibration. An understanding of the effect of phase errors on MFS, and of better methods of self-calibration are therefore of the greatest importance for implementing practical MFS.

In order to avoid the overall fidelity being bound to that of the single frequency reconstructions some way of using the MFS data *in toto* to constrain the phase solution at a particular frequency is needed. The situation is similar to that of the removal of spectral effects. The methods based on single frequency reconstructions presented in Chapter 4 tied the MFS image fidelity to that of the single frequency images, only by deconvolving the total MFS data was a dramatic improvement in fidelity achieved. Simply combining the frequencies and then self-calibrating the whole data set will not work when there are significant differences between images at the different frequencies. This approach may, however, be all right if the bandspread is small (see Chapter 7).

One possibility for linking phase corrections is to use the estimates of I and $I\alpha'$ that emerge from a MFS processing algorithm like Double Deconvolution (which uses all the data) to generate estimates of the source structure at all the observing frequencies. It is then possible to self-calibrate the data at each of the different frequencies against these separate images. Such an algorithm is outlined in fig(6.8b). By using the I and $I\alpha'$ images created from the full MFS data it is hoped that the effects of constraints at one frequency can be communicated to the phase solutions at other frequencies. It has to be admitted that the complexity of the above method makes it difficult to follow how information is transferred between frequencies. It can be hoped however

that the method will be significantly better than separate self-calibration at reducing the effects of self-calibration based errors. It will however probably fall short of the performance of a conventional array with similar uv coverage to that of MFS MERLIN.

Although the above method of ‘parallel self-calibration’ seems hopeful as an improvement upon ‘separate self-calibration’ we cannot be certain that DD can be used successfully in combination with self-calibration until simulations are carried out.

It is quite probable that the most efficient implementation of self-calibration would be a hybrid method. Initially separate self-calibrations would be done to get the phase errors approximately correct and these would be followed by parallel self-calibration with DD in the processing loop. The self-calibration process can also be speeded up by first doing conventional self-calibration processing at the central frequency, the single-frequency image that emerges can then be used as the starting point for the separate self-calibration of each of the other frequencies. Errors initially introduced by self-calibrating a uv data set against an image made at a different frequency would probably be removed at a later stage.

6.4.4 Linking Phase Errors

As mentioned first in Chapter 1 there are two proposed ways of operating MFS on MERLIN. In the first mode a whole day is spent at each frequency so that, if we wish to form the uv coverage for f frequencies, f days of observation will be required. Whilst this has the advantage of producing a lower image noise than observations that are confined to a single day it greatly restricts the operational flexibility. Because of the large amount of observing time required to obtain good uv coverages only a minority of sources will warrant such many-day MFS observations.

In the second implementation the frequency of observation is rapidly varied (on timescales of a few minutes) and so a full MFS observation can be completed within a day. Note that in this implementation no more data is




Figure 6.8: Possible methods of phase self-calibrating MFS data illustrated for the case of three frequencies. a) Separate self-calibration b) Parallel self-calibration.

collected than with conventional methods. The same number of uv points are however more uniformly spread over the uv plane by the frequency switching, so improving the uv coverage.

A by-product of the rapid switching implementation will be that the phase errors at the different frequencies will be correlated if the switching is fast enough. If we cycle through all the frequencies before conditions above the antennae have time to change significantly then the known behaviour of phase errors with frequency can be used to connect the phase errors at the different frequencies. Such extra constraints may be useful given the possible problems of self-calibration with MFS.

The phase errors that occur above an antenna can be separated into two parts, a tropospheric part and an ionospheric part. The origin and frequency dependence of both components of phase error are described by Thompson, Moran and Swenson(1986). The ionospheric phase error is caused by differing electron densities in the ionosphere which cause changes in the refractive index. The resulting phase difference (relative to free space propagation) at antenna i is from Thompson et al (1986) equal to

$$\phi_i = \frac{8.49 \times 10^{-7} N_i}{\nu} \quad (6.3)$$

where ϕ_i is the phase error, ν the observing frequency in Hz and N_i the column density of electrons above antenna i (m^{-2}). The phase error on a baseline ij will depend on the difference between ϕ_i and ϕ_j and therefore on the difference in the column density of electrons above the two sites $N_{i,j}$. For small differences in frequency the baseline phase error will obey

$$\phi_{i,j}(\nu_h) - \phi_{i,j}(\nu_l) = -8.49 \times 10^{-7} \Delta N_{i,j} (\Delta \nu / \nu_o^2) \quad (6.4)$$

$$\Delta \nu = \nu_h - \nu_l \quad (6.5)$$

where ν_h and ν_l are the upper and lower frequencies respectively and ν_o their mean.

The differences in electron column densities above the antennae can be caused by several effects in the ionosphere including the propagation of gravity-acoustic waves (Travelling Ionospheric Disturbances or TID's). These waves

have periods of 20-50 minutes, typical wavelengths of hundreds of kilometres (Kelder and Spoelestra 1987) and can have amplitudes of up to 5% of the total electron column density.

The second major source of phase error is due to propagation through the troposphere. The refractive index of water vapour at centimeter wavelengths is significant and so the effective path length above each antenna will be affected by the water vapour density and temperature above each antenna. There is also a smaller contribution due to the effects of the gaseous components of the atmosphere which depends only on the total air density above each antenna. The extra path length introduced will be almost independent of frequency and therefore the phase error on baseline ij will be proportional to frequency, so that

$$\phi_{ij} = (2\pi/c)(L_i - L_j)\nu \quad (6.6)$$

where L_i and L_j are the extra path lengths above each antenna. Such variations in path length occur due to different phenomena (Hinder and Ryle 1971) on different sized baselines. On scale sizes smaller than 700 m or so they are due to the poor mixing of water vapour. On longer baselines comparable with those of MERLIN, there are differences caused by weather systems and the curvature of the troposphere. Effects due to weather systems can be partially corrected by modelling the troposphere based on ground measurements of temperature, pressure and humidity, likewise the effects of different zenith angles can be removed by assuming a model troposphere. Phase errors which are proportional to frequency of the type caused by the troposphere are also introduced by propagation along cables or by other time delays applied to the signals after reception but before correlation.

It has been seen that there are two components to the phase error, one of which is proportional to $(1/\nu)$ the other proportional to ν . The phase error on a particular antenna can therefore in principal be modelled completely by only two parameters, one proportional to N_i the electron column density and the other proportional to L_i the tropospheric excess path length. This means that instead of self-calibration being required to solve for Nf antenna-based

unknown phases, where N is the number of antennae and f the number of frequencies used, if we can assume linked phases only $2N$ parameters need be estimated. This decrease in the number of unknowns to be solved for can potentially increase the fidelity of MFS self-calibration.

6.4.5 Results of Frequency-Switching Observations

In order to see whether the predicted correlation between phase errors described above actually occurs in practice some preliminary observations have been made.

Frequency switching observations were made in October 1986 whilst the MERLIN system was operating in L-band. The frequency was switched every 2 minutes between 1612, 1632 and 1662 MHz. This 3.1% range of frequencies is the largest change of frequency that can be produced with the present RF system without having to alter filters at the antennae (see section(7.1.1)). The chosen target, the bright point source calibrator 3C48, was observed for over 4 hours and samples of amplitude and phase were produced every 30 seconds. Because the present LO system contains multiplier stages which require at least 30 seconds to settle down, the first 30 seconds of each 2 minute period on each frequency had to be discarded. The switching time of two minutes was chosen as a compromise between being as fast as possible to follow phase variations and not being so short as to lose too large a fraction of data due to the above effect.

Plots of phase against time are given in fig(6.9) for the 134km Defford-Tabley baseline at each of the 3 observing frequencies. Figure(6.10) shows a composite plot of the same data in which the data at different frequencies have been superimposed after applying a constant phase shift to align the data sets up. The time varying part of the phase shows a strong correlation from one frequency to another which suggests that the idea of linking phase errors is practical. The 3% range of frequencies is however far too small to allow us to distinguish between an ionospheric or tropospheric origin for these errors on the basis of their frequency dependence.




Figure 6.9: Results of frequency switching observations on the Defford- Tabley baseline at each of the three frequencies used.




Figure 6.10: The result of applying constant phase shifts to the data at each frequency in fig(6.9) so as to align the data sets up. The time variable part of the phases can be seen to be highly correlated from frequency to frequency.

UT

The time constant phase differences between the frequency channels seem to be roughly proportional to frequency and to be antenna- based. The antenna-based nature of these phases follows from the fact that they ‘close’ (i.e. sum to zero) around closed loops of baselines. The size of the observed phase gradients with frequency however seems difficult to reconcile with an ionospheric or tropospheric origin to these time constant phase errors. In addition, phase gradients of the magnitude observed would also cause serious decorrelation across the 10 MHz passband at each frequency, such effects would be obvious in the recorded amplitudes but are not observed. The above fact points to an origin for the phase errors in the frequency dependence of the LO phase rather than the received signals. Because the LO is a very narrow band signal no decorrelation will result from delays in the LO signal path. The maximum observed phase gradient of 1 turn in 20 MHz corresponds to a time delay of 35ns or a path length of 10 m or so such path lengths tie in well with these phase errors being due to the cables used to carry the LO signal from the ground to the telescope focus boxes.

6.4.6 Possible Methods of Phase Linking

Consider the total phase error (ionospheric plus tropospheric) on a particular baseline at frequency ν

$$\phi(\nu) = \phi_{ion}(1 + \Delta\nu/\nu_o)^{-1} + \phi_{trop}(1 + \Delta\nu/\nu_o) \quad (6.7)$$

where ϕ_{ion} and ϕ_{trop} are respectively the ionospheric and tropospheric phase error at the reference frequency on the baseline. Note that each of the above phase errors are telescope factorizable, i.e. they can be written as the difference between phase errors associated with each antenna. Expanding the above equation over a narrow range of frequency we obtain.

$$\phi(\nu) = (\phi_{ion} + \phi_{trop}) + (\phi_{trop} - \phi_{ion})(\Delta\nu/\nu_o) + \frac{1}{2}\phi_{ion}(\Delta\nu/\nu_o)^2 \dots \quad (6.8)$$

The first term contains the major part of the phase error contribution whilst the second and third terms contain the significant frequency dependent

contributions. We can readily see that in order to estimate the first order contribution from the frequency dependence of phase alone we will be required to detect both the second and third order terms. For the narrow frequency bandspreads envisaged for MFS (of say $\pm 10\%$) finding an error in the overall phase error solution of 5° would require the detection of differences in phase with frequency of better than 0.025° . If the estimates of phase error at each frequency are affected by structure or by significant thermal noise then it can be seen that the solution for the phase errors will be very poor. The above method of phase recovery therefore appears to be very poorly conditioned.

Another approach to phase linking is to note that, provided the residual ionospheric phase contribution after removing the results of any ionospheric modelling are small enough, the phase at each antenna can be modelled as a constant plus a linear term in frequency. The standard self-calibration algorithms can then be modified to solve for an antenna based constant phase plus an antenna based phase gradient with frequency. This method combines the principles of antenna based self-calibration methods and of frequency dependent phase, rather than relying only on the frequency dependence. We can hope that this approach is much more robust than the former method.

If the above modification of self-calibration appears practical then the number of unknowns (telescope phase errors) will be reduced from Nf when the phase errors are uncorrelated to only $2N$ when they are linked. The total number of good closure constraints will then be

$$\frac{Nf(N-1)}{2} - (2N-1), \quad (6.9)$$

The ratio of these 'good' constraints to the number of baseline phase measurements $fN(N-1)/2$ will be

$$1 - \frac{(4N-2)}{fN(N-1)} \quad (6.10)$$

as compared to the ratio when there is no phase linking of

$$1 - \frac{2}{N} \quad (6.11)$$

For a single frequency array with 8 elements (i.e. 28 baselines) the 'self-cal' ratio is 0.75 whereas if there are 5 frequencies ($f=5$) (with 140 baselines) which

are linked, then the 'self-cal' ratio for the same array will be 0.89. This value is only achieved on a conventional array of 20 antennae. Note that in terms of uv coverage the MFS observations will 'only' be equivalent to a conventional array of 17 telescopes. Phase linking therefore appears to have the potential in principle to transform a situation in which an MFS observation is more susceptible than an equivalent single frequency observation to self-calibration based errors into a situation in which it is more resilient to such errors.

The above argument ignores the conditioning of phase linking methods with respect to noise. In realistic situations in which there is noise present the improvements in fidelity may not be as dramatic as described above. Note that as well as thermal noise there will also be effects due to the finite rate of frequency switching. Phase linking assumes that the conditions above the antennae remain constant over the frequency cycling time or longer, at some level this assumption will break down. There will for instance be short timescale (shorter than 5 minutes) phase variations dominated by geomagnetic effects (Kelder and Spoelstra 1987) and the effects of local water vapour fluctuations. This second tropospheric effect has been shown to have a power law time-frequency spectrum with index -1.2 (Armstrong and Sramek 1982) so that the size of phase variations only decreases quite slowly at short timescales.

Fast time variations in phase which are much faster than the switching rate will have effects similar to noise, variations with timescales of order the cycle time will in contrast introduce systematic errors into the phase solutions. Increasing the switching rate will reduce the above problems but this remedy is ultimately limited by the capabilities of the hardware and the necessity of having a reasonable signal to noise per integration if self-calibration is to be used.

Further study of short timescale variations are needed, for instance by looking at the spectrum of time variations in MERLIN observations of point dominated sources. In addition the typical size of phase errors causing 'self-cal based' reconstruction errors in single frequency images must be investigated. Clearly if the phase solution errors introduced in linked self-calibration due

to thermal noise or switching effects are larger than these errors then phase linking is unlikely to give any improvement over separate self-calibration.

6.4.7 Phase Referencing

Given the possible problems of using phase self-calibration with MFS data it may be worthwhile to consider using phase referencing even on strong sources. Already this mode of observing is becoming more common with MERLIN. With the improved sensitivity of extended MERLIN a much larger number of calibrators will be observable sufficiently close to the target source that they share the same phase errors as the target source, i.e. so that they lie in the same 'isoplanactic' patch. There should be at least one source of at least 100mJy within 2° of most target sources at L-band (P.N. Wilkinson, private communication). Because the primary beams of the MERLIN 25m antennae are of size 0.5° at L-band this means that phase referencing observations of calibrators of such flux will need to be made in a nodding mode, spending say 2 minutes out of every 10 on the calibrator. Thermal noise limits on the accuracy of the resulting phase calibration will limit images of the target source to a dynamic range of 2000:1. Errors will also occur due to variations of phase with time and variations between the phase error at the target source and calibrator, i.e. due to the isoplanatic assumption being imperfectly correct. From fig(6.10) variations of 10° in phase in 10 minutes seem reasonable, the effects of non-isoplanatism on the angular separation of 2° are unknown but are probably less than this. Errors of this size on this timescale should give rise to residual errors in the map at a dynamic range of 500:1 if the errors are independent from each 10 minute period to the next.

Although the above argument is rather crude it does suggest that phase referencing could be a viable alternative to phase self-calibration for observations where only moderate dynamic range is required. Again more work is needed, especially valuable would be more practical experience of phase referencing with the present MERLIN. Even if phase referencing is not a complete answer

it may well be sensible to apply it to data to get the phase approximately correct before later applying self-calibration.

6.4.8 Summary

We have seen that there are special problems when using phase self-calibration with MFS data. Doing simple separate self-calibration at each frequency may well introduce errors which are not easily recovered from and which may potentially limit the fidelity of any MFS image. So-called parallel self-calibration may limit this problem, but this method must first be fully tested to ensure there are no problems with combining DD and self-calibration.

When MFS observations are made by the frequency switching method there is an additional possibility, that of using the frequency dependence of phase errors with frequency to constrain the solutions. Again tests are required to check the feasibility of this approach.

Phase referencing may be a solution for sources of low to moderate dynamic range. Both simulations and test observations are required to investigate this technique further.

An additional point not mentioned above is that the inclusion of VLA data, necessary to increase the fidelity of reconstructed extended emission may also ease the phase solution problem. As described in section(7.2.4), VLA data will allow a phase solution to be made for the four MERLIN antennae which participate only in the short baselines. This will greatly ease the scale of the problem which self-calibration needs to solve and so reduce the size of self-calibration based reconstruction errors.

The whole question of phase self-calibration is probably the most important aspect of MFS imaging still to be fully understood and should be a priority for future work.

6.5 Amplitude Calibration

6.5.1 Introduction

As well as applying phase calibration to the MFS data, amplitude calibration will also be required. Amplitude calibration will be applied in two ways. First the observations of the target source will be compared against observations of calibrator sources of known flux and secondly any residual amplitude errors will be removed by applying amplitude self-calibration (see e.g. Pearson and Readhead 1984).

6.5.2 External Calibration

Errors in the external calibration of the data at the different frequencies are potentially serious for the successful processing of MFS data. Such amplitude errors will be of long duration, possibly the whole of the observation, and will have signatures in the uv plane very similar to those expected from spectral effects. These ‘time constant’ amplitude errors associated with the different frequencies will cause confusion if it is necessary to apply processing algorithms (such as DD) to remove the results of spectral effects. If the amplitude at a frequency ν_i is in error by the factor $1 + e_i$, then the resulting dirty map will be

$$d = I * (D_o + E) + I\alpha' * (D_1 + (E * D_1)) + \quad (6.12)$$

where the error beam E is such that \bar{E} is non-zero only on sampled points of the uv plane where it has a value of e_i where ν_i is the frequency at which the observation was made. The error beam E should be very similar to D_1 in appearance, the ratio of the magnitudes of the two beams can be expected to be roughly in the ratio of the magnitudes of the weights of the two uv coverages, i.e. roughly

$$\frac{\sigma_i}{(\Omega/2\nu_o)} \quad (6.13)$$

where σ_o is the mean of the amplitudes of the factors e_i and Ω is the total bandsread. If we require that the effects of amplitude errors are say 1/10 th

of the spectral error then we require that

$$I * E \leq ((1/10)I\alpha' * D_1) \quad (6.14)$$

which gives $\sigma_o \leq (1/10)\alpha'_{typ}(\Omega/2\nu_o)$ where α'_{typ} is the typical value of α' on the dominant bright parts of the radio source (say 0.2). If the bandspread is taken as 20% then we obtain that σ_o is smaller than 0.002, i.e. we require the relative overall amplitude calibration of the different frequencies to be correct to better than 0.2%.

Amplitude calibration at this level of accuracy should be possible though it will require some care. Note that the accuracy of calibration will probably be better in the time switching implementation than in the sequential implementation of MFS (see sect(1.5)). Calibration can be done either against external point source calibrators or using noise diode calibration. In this latter method of amplitude calibration known amounts of uncorrelated noise are injected into the receivers periodically and the resulting changes in the measured correlation coefficients used to calibrate the baselines.

Point source calibrator sources should be chosen with care, slightly resolved calibrators will give systematic effects because observations at higher frequency will resolve the calibrators more than observations at lower frequencies. The spectra of the calibrators are also of obvious importance, calibrators which are assumed to have good power law spectra but contain slight non-power-law or 'curved' components in their spectra will introduce errors into the calibrated data and should be avoided.

6.5.3 Self-Calibration

At present MERLIN L-band images with dynamic ranges of greater than 200:1 require the use of amplitude self-calibration (SC) to remove the effects of time varying amplitude errors, or the residual effects of inaccurate external calibration. Although it should be possible to reduce the size of the above amplitude

errors somewhat by careful hardware design and more use of external calibration it is likely that to produce MFS images with dynamic ranges greater than say 1000:1 will require the continued use of amplitude self-calibration.

There are two ways to implement such self-calibration with MFS data. In exactly the same way as for phase one can use either separate or parallel self-calibration. Both of these methods have their own advantages and disadvantages.

One effect that may cause problems with the ‘separate’ amplitude self-calibration is the wander in the overall amplitude scale of the data at the different frequencies. This is the amplitude analogue of the position wander that we get with phase self-calibration. As we have seen the overall amplitude calibration needs to be maintained to high accuracy if spectral effects are to be successfully separated out and so this will be a potential problem with this implementation. In order to stop the above amplitude wander it is vital that the separate SC’s at the different frequencies utilise gain-normalisation so that the mean amplitude change applied to any baseline in any data set is unity.

The ‘parallel’ form of SC processing has the worry that there may be substantial interaction between the amplitude SC and the DD algorithm. The problem can be expected to be worse than in the phase case. Whilst sky plane errors due to errors in phase consist of the true image convolved with an anti-symmetric beam, amplitude errors will instead be equal to the true source convolved with a symmetric beam (see e.g. Ekers 1985). When in DD we attempt to estimate the $I\alpha'$ distribution by first convolving the dirty image with the symmetric D_1 beam, amplitude-error contamination will tend to have more effect than phase-error contamination because of the symmetry of the error contributions.

It appears that the interaction of DD with amplitude errors is likely to be a serious problem for the ‘parallel’ mode of applying SC. It appears therefore that (provided gain normalisation is applied) separate self-calibration is the better option for amplitude SC.

6.6 Non-Power Law Spectra

6.6.1 Introduction

So far it has been assumed that the radio spectrum at every point on a radio source is exactly power law. This power law spectrum has been shown to hold at moderate frequency (see sect(2.2.2)) provided the underlying electron energy spectra is also power law. However more detailed analysis of the theoretically expected synchrotron spectra (Pacholczyk 1970) shows that, particularly at high and low frequencies, there will be significant departures from a power law spectrum even if the underlying electron energy distribution is power law after initial electron acceleration. The extent to which these departures from the power law affect MFS are considered in this section.

In chapter 2 it was shown how an arbitrary spectrum could be analysed as a power series in either $(\Delta\nu_i/\nu_o)$ (the linear expansion) or in terms of $\ln(\nu_o/\nu_i)$ (the logarithmic expansion). These expansions were then used to show the form of the dirty map that arises from inverting MFS data. Most MFS algorithms which remove spectral effects assume that the two order approximation to this expansion is valid. Although the logarithmic expansion is to some extent optimised for a power law spectrum (see Appendix B), so that a power law spectrum gives terms which rapidly decrease with order, departures from a power law spectrum are not critical provided the spectra remain smooth. In the coming sections we consider the various sources of non-power-law spectra and the effect of each on MFS. The impact of each is described by the size of the second order terms of the spectral expansion that these departures introduce. The effect of non-power law components become important at dynamic ranges such that the resulting second order contributions become larger than the thermal noise. Above these critical dynamic ranges the two order approximation will no longer be sufficient and the effects of a non-power-spectra will then become a problem.


6.6.2 Spectral Ageing

Synchrotron emitting electrons will radiate and lose energy. Because this energy loss is proportional to the energy of the electron squared this mechanism will tend preferentially to reduce the population of high energy electrons. These high energy electrons dominate the high frequency radio emission and so the net effect of 'electron ageing' is to reduce emission at high frequencies below the level expected from a power law spectrum.

The full theory describing the emission spectrum that emerges from an ensemble of electrons evolving with time has been presented in several reference works (e.g. Kardeshev 1962, Pacholczyk 1970). In this section we will consider only the basic parameters of 'aged' spectra.

The form of spectrum that emerges from an aged electron population is presented in fig(6.11). The three curves shown depend on which assumptions are made about the distribution of electron pitch angles as they spiral around lines of magnetic field and about whether there is a continuous injection of fresh electrons with a power law energy distribution. In the latter case we obtain a 'continuous injection' spectrum (Pacholczyk 1970) marked CI in fig(6.11). Because of the continuous injection of fresh electrons the fall off in emission at high frequencies is much less than in the other two models. Instead of a rapid decrease, the spectrum after the 'break' or knee becomes power law again with a spectral index of $\alpha + (1/2)$ where α is the spectral index at low frequencies.

If we assume there is no injection of electrons then there are two possibilities, the curve marked KP refers to the so-called 'Kardeshev-Pacholczyk' model in which it is assumed that the pitch angle of an electron stays constant throughout its radiative lifetime. The alternative model marked JP is known as the 'Jaffe and Perola' model (Jaffe and Perola 1974) in which it is assumed that the pitch angle distribution is continually isotropized by the scattering of electrons off MHD waves or other irregularities in the magnetic field. The distribution of electron pitch angles is important because the rate of energy loss depends on $(\sin\theta)^2$, where θ is the pitch angle. In the KP case



Theoretical Synchrotron Spectra.

Injection index = 2.

CI Continuous injection of electrons.
 KP Kardashev/Pacholczyk model (pitch angle anisotropy).
 JP Jaffe/Perola model (pitch angle isotropy).

Figure 6.11: Model spectra affected by spectral ageing.

After Stephens (1987)

it is possible for high energy electrons with small pitch angles to survive a long time whereas in the JP case the continual isotropization rules this out. The ability for high energy electrons to survive much longer in the KP model means that the reduction in flux is reduced compared to the JP model. In the latter case the high frequency tail shows an exponential decrease whilst in the KP model the spectrum eventually becomes straight with a spectral index of $\alpha = (2\gamma + 1)/3$, γ being the electron energy injection index. Despite the large differences between the models at high frequency, at moderate frequencies (say less than 15 GHz, Stephens 1987) the two models have very similar behaviours. This is true to the extent that it has not been ^{possible} to decide which of the models is more realistic despite making extensive observations (Stephens 1987).

Given the similarity between the spectra of the three models at moderate frequency, it should be possible to get some idea of the effects on MFS of spectral ageing by analysing a single model. In the following analysis the KP model is assumed and the approach of Myers and Spangler (1985) adopted. According to these authors the spectrum that results after a time t in the presence of a magnetic field B if the initial electron energy index is γ will be

$$I = I_o(\nu/\nu_o)^{-(\gamma-1)/2} Q(\gamma, \chi(\nu, B, t)) \quad (6.15)$$

where the dimensionless quantity χ is defined as

$$\chi = 8.9 \times 10^{-4} (1 + z) \nu_{GHz} B_{-5}^3 t_{Myr}^2 = (\nu/\nu_T) \quad (6.16)$$

where ν_{GHz} is the observing frequency in GHz, z the redshift, B_{-5} is the magnetic field in units of 1×10^{-5} G and t_{Myr} the age in millions of years; ν_T is defined to be the ‘break frequency’ of the spectrum. The function $Q(\gamma, \chi)$ must be evaluated numerically, a graph showing Q against $\chi^{1/2}$ for several different injection indices is presented in fig(6.12). For the injection index $\gamma = 2.5$ the graph is approximately linear in the range $0 \leq \chi \leq 0.7$ allowing the use of the approximation

$$Q = 1 - \chi^{1/2} = 1 - (\nu/\nu_T)^{1/2} \quad (6.17)$$

If we attempt MFS observations of an image dominated by a point source with the above spectrum which has been corrected for an ‘overall’ spectrum of $(\gamma - 1)/2$ then for the logarithmic expansion (see Appendix B) we obtain a second order contribution of

$$(1/2) \frac{d^2 I'}{dy^2} * D_2 \quad (6.18)$$

where

$$y = \ln(\nu_o/\nu), \quad (6.19)$$

Since we have after correction for the low frequency spectral index that $I' = QI_o$ we obtain using the above expression for Q that the second order contribution will be

$$I_o(1/4)(\nu_o/\nu_T)^{1/2} * D_2 \quad (6.20)$$

centred at the position of the point component. From Chapters 2 and 3 we have seen that for a fractional bandspread of 20%-25% the D_2 beam is similar in structure to that of a normal dirty beam and has a peak level of around 10^{-2} with a sidelobe level a factor of 10 smaller still. The peak second order effect and the peak sidelobe effects will therefore be respectively 10^{-2} and 10^{-3} times the above expression.

The thesis work of Stephens(1987) contains an analysis of the spectra of hotspots over two decades of frequency, the estimated values of V_T found can be used to estimate the size of spectral errors introduced by aged spectra. The lowest break frequency found in a sample of 25 hotspots (assuming the KP model) was 7.9GHz for a hotspot in the source 3C47, the next lowest was 18 GHz in 3C265, the rest of the hotspots had ν_T in the range 20GHz to greater than 1417 GHz. In fig(6.13) are shown the fitted spectra for a selection of the hotspots observed in the above study.

For MFS observations centred at 1660 MHz and $\nu_T=7.9\text{GHz}$ ($\chi = 0.2$) we obtain $(1/4)((\nu_o/\nu_T)^{1/2})=0.057$ with a peak level for D_2 at 10^{-2} this gives a peak second order contribution at the point component of $2.85 \times 10^{-4} I_o$ where I_o is the flux of the point at the reference frequency. As noted in chapter 2 this ‘strong limit’ is probably overly pessimistic and images with ‘dynamic ranges’




Figure 6.12: Graphs of the quantity Q against $\chi^{1/2}$ for different values of electron injection index γ assuming the KP model. From Myers and Span-
gler(1985).

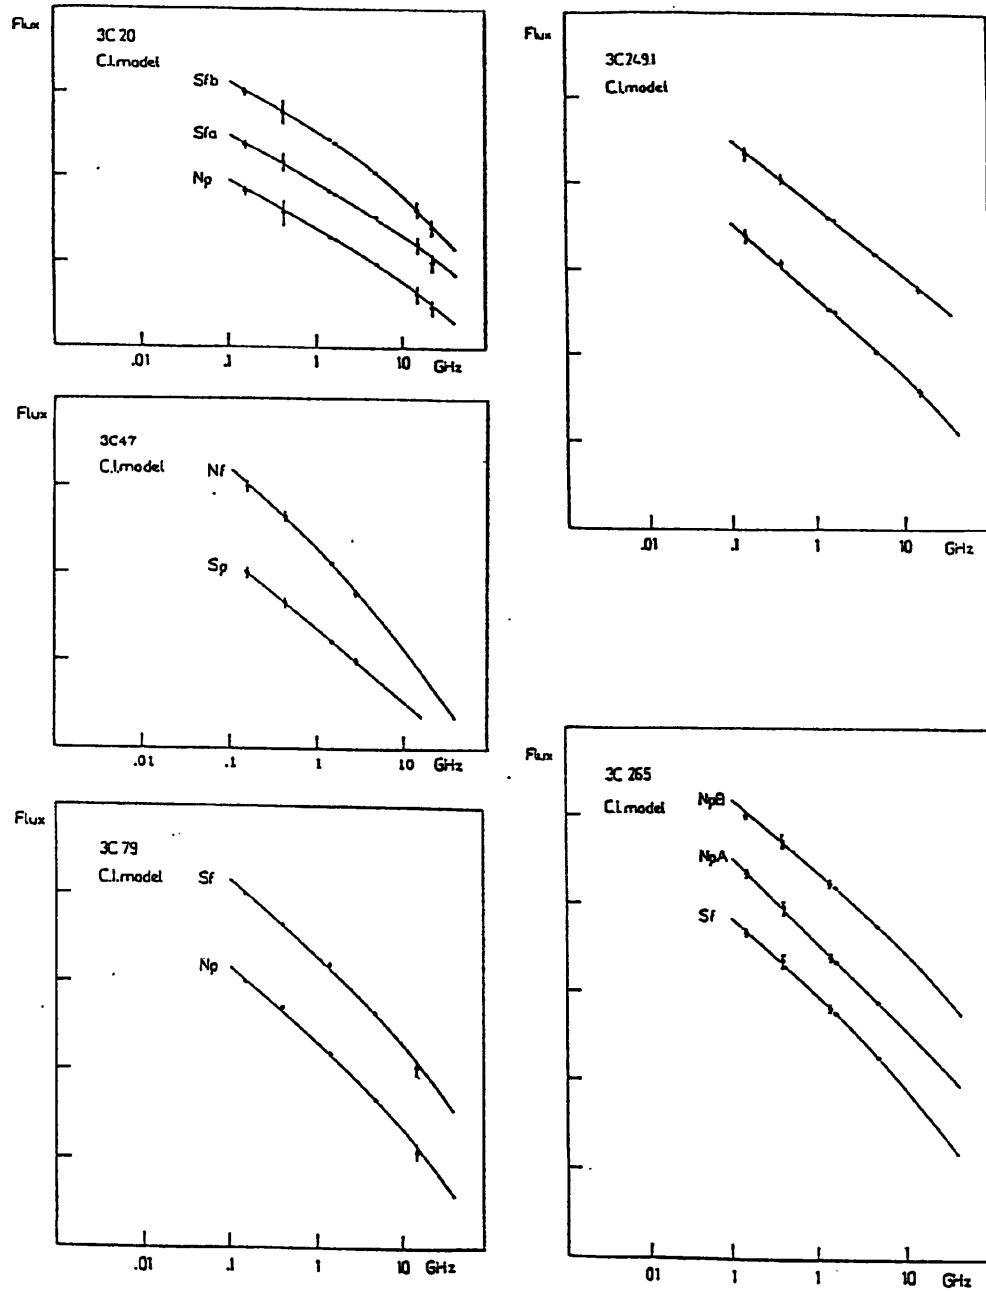
of a factor of 10 greater still can be successfully processed by MFS algorithms giving a limit of 40,000:1.

Away from the hotspots the values of ν_T can decrease to lower levels than assumed above. Much of the extended emission will be due to electrons that have long ago passed through the hotspots and so have had much longer to age, this will naturally tend to decrease ν_T . Counteracting this effect however will be the fact that the magnetic fields will be much smaller in the lobes and this will partly compensate for the above effect (equ(6.16)). From Stephens(1987) it appears that ν_T can decrease by a factor of 5 in going from the hotspots to the lobes. Note that because the intensity will fall by factors of hundreds or thousands in going from the hotspots to the lobes the dynamic range constraints due to non-power-law spectra will be set by the hotspots, despite the lower spectral break frequencies in the lobes.

6.6.3 Synchrotron Self-Absorption

The other major effect causing non-power-law spectra¹⁵ is synchrotron self-absorption. It can be shown (see e.g. Lang 1974) that a plasma that emits synchrotron radiation will also reabsorb a fraction of the radiation emitted, reabsorption being most efficient at low frequencies. The optical depth for absorption will depend on the length of the line of sight and the properties of the plasma along the line of sight. In very compact regions such as radio cores the optical depth can exceed unity and the spectrum can 'turnover' (see fig(6.14)) so that the emission decreases with increasing wavelength.

The main features of a Synchrotron Self-Absorbed (SSA) spectra are shown in fig(6.14). Because the absorption coefficient for SSA scales as $\nu^{-(\gamma+4)/2}$ (Pacholczyk 1970) the effects of SSA are seen mainly at low frequency. At very low frequencies the theoretical spectrum will obey $I(\nu) \propto \nu^{5/2}$ no matter what the value of the electron injection index (γ) whilst at high frequencies it becomes the conventional power law spectra with a spectral index which depends on γ .



Hot Spot Spectra.

Figures show the observed spectra with fitted continuous injection curves. Flux densities are on an arbitrary scale.

Figure 6.13: Results of fitting model aged spectra to multi-frequency observations of hot spot flux (from Stephens(1987)).




Figure 6.14: Theoretical Synchrotron Self-Absorbed (SSA) spectra for different values of the electron injection index γ . From Pacholczyk(1970).

At centimeter wavelengths significant effects from SSA occur only in compact radio cores such as unresolved radio cores and possibly also in compact hotspots. The spectra of cores rarely if ever follow the theoretical spectra shown in figure(6.14) however (Kellermann and Pauliny-Toth 1981). The complex spectra actually seen can however be modelled as a superposition of several self-absorbed components with their frequencies of peak emission centred at different frequencies. Typical complex core spectra are presented in fig(6.15). Because cores are often very bright and have complex spectra it will often be necessary to correct for the effects of these cores at each frequency. This can be done most simply by imaging the core at each frequency and then removing the effects of these model cores from the corresponding uv data.

A core with a spectrum similar to 3C84 as shown in fig(6.15) is probably the worst case for MFS. Observations at 1.6 GHz would occur near a minimum in the spectrum, i.e. near a maximum value of $(1/2)\frac{d^2I}{dy^2}$ where $y = \ln(\nu_o/\nu)$. For 3C84 the value of this quantity can be estimated to be 20 Jy (when the value of $I(\nu_o)=10\text{Jy}$), if the sidelobe level of the D_2 beam is assumed to be 10^{-3} then the core second order spectral sidelobes will occur at a dynamic range smaller than 500:1. If dynamic ranges larger than this value are required then special core correction methods must be employed. Although core correction is conceptually easy, in practice it may require some care to achieve the necessary accuracy. Core correction might be especially difficult if the core is embedded in extended emission. Core correction is yet another aspect of MFS which will benefit from simulations in the future.

Apart from cores the other class of object in which SSA is possibly significant is in compact hotspots. For realistic values of physical parameters in a hotspot such as a magnetic field B of $4 \times 10^{-4}G$ (Stephens 1987), diameter $\theta = 0.1''$ (Dreher 1981) and a flux of 100mJy at 1660 MHz (ν_o) then the frequency at which the optical depth is unity ν_s will (following Lang 1974) be equal to

$$\nu_s^{6/5} = 34\left(\frac{S(\nu_o)}{\theta^2}\right)^{2/5} B^{1/5} (1+z)^{1/5} \nu_o^{1/5} \quad (6.21)$$




Figure 6.15: Radio spectra of core dominated sources.
(Kellermann and Pauliny-Toth 1981).

which for a redshift $z=0$ is 38MHz. Given this low value of ν_s any effects on MFS at P or L bands due to spectral curvature will be small and we can use the approximation that $(\nu/\nu_s) \gg 1$. From Pacholczyk(1970) the self-absorbed spectrum obeys

$$I(\nu) = I(\nu_s)J(\nu/\nu_s, \gamma) \quad (6.22)$$

where $I(\nu_s)$ is the intensity at ν_s if the low frequency power law spectrum extended to the frequency ν_s . Taking $r = (\nu/\nu_s)$ the function J is such that

$$J(r, \gamma) = S(\nu_s).r^{5/2}.(1 - \exp(-r^{-(\gamma+4)/2})) \quad (6.23)$$

If we assume that $r \gg 1$ then the coefficient of the exponential is small and the expression may be expanded so that

$$I(\nu) = I(\nu_s)r^{-(\gamma-1)/2}(1 - (1/2)r^{-(\gamma/2+2)} + (1/6)r^{-(\gamma+4)}....) \quad (6.24)$$

for large r all terms in the above expansion except the zeroth and first will be negligible. If we observe a point source with the above spectrum and then 'correct' for an overall power-law spectrum of spectral index equal to the high frequency spectral index $(\gamma - 1)/2$ then the resulting uv data if $\gamma = 2.0$ will correspond to the Fourier transform of a point source with the effective spectrum

$$I'(\nu) = I_o(1 - (1/2)(\nu/\nu_1)^{-3}...) \quad (6.25)$$

Setting $y = \ln(\nu/\nu)$ and finding the second order derivative of I' with respect to y we obtain for the coefficient for the second order spectral effect

$$(1/2)\frac{d^2 I'}{dy^2} = -(9/4)(\nu/\nu_s)^{-3} \quad (6.26)$$

Taking the example of a hotspot with $\nu_s = 38MHz$ then at 1660 MHz the above coefficient equals 2.6×10^{-5} whilst at 408MHz it is 1.8×10^{-3} . With the D_2 sidelobes being at a level of 10^{-3} then even at 408 MHz the MFS dirty map should be free of errors due to spectral curvature down to a dynamic range limit of 5×10^5 . It can be concluded that SSA does not provide any significant problems for observations of hotspots at P or L bands.

6.6.4 Free-Free Absorbtion

The above mechanisms of spectral ageing and SSA are the most important sources of non-power law spectra, for completeness however we examine the effect of free-free absorbtion on MFS. Free-free absorbtion (FFA) occurs when radio waves pass through a region of relatively cool ionized gas, photons being absorbed by thermal electrons as they interact with heavy nuclei in the gas. From Lang(1974) the optical depth will be equal to

$$\tau \approx 8.235 \times 10^{-2} T_e^{-1.35} \nu^{-2.1} \int N_e^2 dl \quad (6.27)$$

provided that $\nu < 10^{10}$ Hz and $T_e < 9 \times 10^5$ K where T_e is the thermal temperature of the intervening gas, ν is the observing frequency in GHz and the integral of the electron density squared (N_e is in units of cm^{-3}) is evaluated over the line of sight in parsecs. For low τ (high ν) the effect of the above absorbtion on an initially power law spectrum will be to produce

$$I(\nu) = I_o(\nu/\nu_o)^{-\alpha_o}(1 - (\nu/\nu_{FF})^{-2.1}..) \quad (6.28)$$

where ν_{FF} is the frequency at which the free-free optical depth is unity

$$\nu_{FF} = 0.3(T_e^{-1.35} N_e^2 l)^{1/2} \quad (6.29)$$

assuming a constant electron density N_e and total path length l parsecs. Proceeding as usual to remove an overall spectral index α_o the coefficient multiplying the D_2 beam will be

$$2.2 \times (\nu/\nu_{FF})^{-2.1} I_o \quad (6.30)$$

Stephens (1987) found that the source 3C123 has a low frequency absorbtion feature probably due to the free-free mechanism with a $\nu_{FF} \approx 500 \text{ MHz}$, (possibly caused by gas with $T_e=100\text{K}$, $l=17 \text{ kpc}$ and $N_e=1 \text{ cm}^{-3}$). At 1660 MHz assuming D_2 sidelobes of 10^{-3} this gives a curvature limited dynamic range limit of 6000:1. At P-band the effects of FFA on MFS will be very considerable, calculating the effect using the above formulae will not be valid since the assumptions made will not fully hold, however it can be expected

that the dynamic range limit will be a least a factor of ten more severe than at L band. Note that although in the above case there might be significant effects due to FFA, none of the rest of the hotspots in the sample of Stephens(1987) showed any sign of FFA absorbtion down to 151 MHz and so it appears that free-free absorbtion is unlikely to be an important factor affecting MFS.

6.6.5 Superposition of Power Law Spectra

The superposition of power-law spectra from different components either along the same line of sight or from different components separated by less than the resolution will produce a total spectrum which is not exactly power law. The coefficient of the second order contribution that results is trivially equal to the sum of the second order coefficient of each separate component. If the total spectrum is

$$I(\nu) = I_1(\nu/\nu_o)^{-\alpha'_1} + I_2(\nu/\nu_o)^{-\alpha'_2} \quad (6.31)$$

then the corresponding coefficient multiplying the D_2 beam will be equal to

$$\frac{I_1(\alpha'_1)^2}{2} + \frac{I_2(\alpha'_2)^2}{2} . \quad (6.32)$$

This will be less than the coefficient due to a single component of flux $I_1 + I_2$ at the reference frequency and effective spectral index equal to the larger of α'_1 or α'_2 . The second order effects of such a composite spectrum are therefore not significantly larger than those from a perfect power law spectrum and should cause no special problems for MFS.

6.6.6 Summary

In this section possible problems due to non-power-law spectra have been identified and quantified. In practice it has been found that at centimeter wavelengths these effects are typically less than the thermal noise and can therefore be ignored. Despite this general rule there may be rare cases of objects in which the effects of spectral curvature due to spectral ageing, SSA or FFA are important in images of high dynamic range (i.e. greater than

10,000:1). The above estimates of spectral effects should allow us to identify these situations.

The one class of component in sources which will often cause problems are compact cores. In the worst case these cores will introduce sidelobe-like features at a dynamic range level of 500:1 and above. To remove the significant spectral effects around such components it is fortunately only necessary to determine the total flux of the unresolved core at each observing frequency and then remove the effects of the core at each frequency from the MFS data. By imaging the separate channels it should in principle be easy to determine the core flux at each frequency and so carry out a successful core correction. Simulations of this process should be carried out in the future.

Chapter 7

MERLIN MFS Observations of 3C179

7.1 Introduction

7.1.1 Previous MFS Observations

In the last chapter we discussed the practical aspects of MFS, however to fully understand the practical problems and benefits of MFS there can be no substitute for making actual MFS observations. In this chapter we will consider the processing of such real MERLIN MFS observations of the quasar 3C179.

Some preliminary MERLIN MFS observations have already been made; the source 3C249.1 was observed with 5 frequencies in L-band with MERLIN in 1984. The results of these observations (Conway 1985) were encouraging but inconclusive. The main problem was that at most frequencies only short observations of 8 hours were available so that the slight improvement in radial filling of the uv plane gained by using MFS was completely dominated by the large gaps in hour angle coverage.

The new MFS observations described in this chapter have almost complete hour angle coverage at each frequency and so the major drawback of the 1984 observations should be avoided. Unfortunately, as in 1984, the present MERLIN RF technology limits MFS L-band observations to the frequency range 1604 MHz to 1720 MHz, a fractional bandspread of less than 7% (see Sect(6.3.1). Although this bandspread allows a modest improvement in total uv coverage it falls a long way short of the dramatic improvements obtained from a 25% bandspread, which is the value assumed in the simulations presented in chapter 5. Despite this it is still a very useful exercise to attempt MFS with the present system if only to gain some experience with dealing with MFS

data. Only by experiencing 'real life' can all aspects of the MFS problem be appreciated.

7.1.2 MFS and the Present MERLIN RF System

A schematic diagram of the present RF system was presented in fig(6.6). Three aspects of this system limit the range of frequencies over which MFS observations can be made. The first limit is set by the frequency range of the RF filters immediately after the feed system, which have a passband set to the range 1580-1730 MHz. The purpose of these filters is to reject strong interfering signals which, once they enter the receiver system, might cause saturation of amplifiers and the consequent production of strong harmonics.

The second and most important restriction affecting the MFS performance is the limited range of local oscillator frequencies that can be generated. Although the MERLIN frequency synthesiser can generate frequencies over an octave in frequency (250 - 500 MHz) these frequencies must be multiplied by 4 in order to provide an LO at L-band. This multiplication is achieved by a device which is limited to processing LO signals for the RF frequencies which lie in the approximate range 1604 MHz to 1720 MHz (a range of less than 7% in frequency). This device is the limiting factor in the range of frequencies that can be accessed.

The performance of the multipliers on the different antennae are variable. In the observations to be described there was some difficulty reaching the bottom end of the above band on some antennae. The performance of all the multipliers is thought to be deteriorating with time as the devices age.

Finally the observations were restricted by the presence of a filter in the LO system. This filters the signal coming from the MERLIN synthesiser before it enters the multiplier stage. Due to the limitations in the performance of the multipliers this filter must have a bandwidth of less than 4%. If we wish to observe over the full range of 7% in frequency this means that at some point the centre frequency of the filter must be changed. This can only be done by manually adjusting the filter at each antenna. Given this restriction the only

practical implementation of MFS consisted of making separate observations at each observing frequency. This meant that to obtain observations over 5 frequencies required 5 days of observation.

7.1.3 Selecting a Target Source

The uv coverages that can be achieved with the 6 antenna MERLIN at various declinations of source with 5 frequencies distributed over the range 1604-1720 MHz are shown in fig(7.1). The frequencies assumed are 1604, 1632, 1662, 1690 and 1720 MHz. These frequencies are chosen so as to be approximately evenly spaced and also to avoid known interference. The best filling of the uv plane in this configuration is achieved for high declinations between say 65° and 80° . For testing MFS a source in this declination range should therefore be chosen.

In addition to having a high declination the target source should also be large enough and complex enough that single-frequency imaging with MERLIN is difficult because of reconstruction errors. The single frequency reconstruction must be 'complexity limited' so that improvements in uv coverage will improve imaging quality.

The target source should also have enough unresolved flux on the longest baselines so that self-calibration can determine good phase and amplitude solutions for all antennae. Low signal-to-noise ratio was a major problem with the earlier observations of 3C249.1 (Conway 1985), especially on the longest baselines where the filling due to MFS was most effective.

The source that best fits all the above requirements is the quasar 3C179. This source has a declination of 67° and the core flux of around 300mJy which will ensure that there is a good signal-to-noise ratio on all baselines. The source also shows emission over a large area as can be seen in the 18cm combined VLA/MERLIN image of Shone et al (1985) (fig(2b)). The total area of the source above the lowest contour of fig(7.2b) is greater than 1600 beam areas (assuming a restoring beam FWHM of $0.2''$). This is much larger than the 30×30 beam areas suggested by Wilkinson(1983) as a 'safe' limit on the size of image that can be made reliably with single-frequency MERLIN.

Although this limit is probably cautious, given prior knowledge of the emission which allows us to use CLEAN ‘windows’, it is clear that 3C179 is a difficult object to image conventionally with MERLIN.

7.1.4 3C179 as an Astronomical Target

As well as being a good source for testing the present MFS system the source is very interesting astrophysically.

The source 3C179 is a well known superluminal on VLBI scales (Porcas 1981 and Porcas 1984). Of the population of superluminal sources with ‘strong’ cores 3C179 stands out (together with 3C345, Browne 1987) as having a large ratio of extended flux to core flux.

Given that the source shows one-sided superluminal motion at VLBI scales it is plausible to suggest that the apparently one-sided kiloparsec scale jet seen in fig(7.2) is due to doppler beaming. An alternative explanation of the one-sided jet is the so-called ‘flip-flop’ mechanism in which the jet is always intrinsically one-sided but periodically changes direction (Rudnik and Edgar 1984). If this model is correct then we would expect fresh electrons to be being injected into the W lobe but not into the E lobe and the effects of this should be detectable in the spectra of the two lobes.

A very striking feature of the source is the dramatic change of direction of the jet at a point in the W lobe; this may be due to the jet interacting with some obstacle. An alternative interpretation is that the jet is splitting into two or more parts at this point in the W lobe possibly due to Kelvin-Helmholtz instability (e.g. Ray 1981). Deciding between the two mechanisms requires better intensity and polarisation images of the relevant region.

7.1.5 Expected Spectral Effects

What type of processing must be applied to the 3C179 MFS data to produce a noise limited image? Will it be necessary to remove spectral effects by using the processing algorithms described in Chapter 4 and 5 or can the data simply






Figure 7.1: Present 6-station MERLIN uv coverages at different declinations for an MFS system of 5 frequencies spread over 7% in frequency.



3C179 (0723+679) VLA 1.6 cm (after Owen & Puschell, 1984)
Contours: 0.6 mJy/beam \times (1,2,3,4,6,9,13,24,39,56)



After Shone(1985) Restoring Beam 0.28" circular
0723+679 (3C179) Combined MERLIN/VLA 1.8 cm
Peak = 364 mJy/beam
Contours: 1.1 mJy/beam \times (1,2,3,4,6,8,12,16,32,64,128,256)

Figure 7.2: Previous maps of the kiloparsec structure of 3C179. a) A VLA A-array map made by Owen and Puschell(1984). b) A combined VLA/MERLIN 18cm made by Shone et al 1985.

be combined and then processed conventionally? To answer this question we must calculate the size of the expected spectral error contribution to see if it is significant with respect to the noise. This contribution will depend both on the size of the spectral dirty beam and the expected $I\alpha'$ distribution. The maximum expected sidelobe level of the D_1 beam for the chosen set of frequencies at the declination of 3C179 is of order 10^{-3} (see fig(7.3)). This level of spectral effect fits in well with the results of chapters 2 and 3.

We consider next the expected $I\alpha'$ distribution across the source. Using the 18cm map of Shone et al (1985) convolved down to $0.35''$ resolution and a 6 cm VLA image (see fig(7.4)) convolved down to the same resolution, estimates of the I , α fluxes and the characteristic sizes of dominant components can be made. These quantities are presented in Table (7.1). In a similar manner to section(3.2.1) we can estimate the spectral sidelobes due to each component and hence the total level of spectral error to be expected.

According to the arguments in Chapter 2 in order to reduce spectral effects we should choose to remove an 'overall' spectral index close to that of the brightest component. In this case this is the core, which has a flat spectrum. We should, therefore, choose to remove an 'overall' spectral index of $\alpha_o = 0$, i.e. in practice we need make no correction for an 'overall' spectral index at all. The component whose spectral effects will then dominate, assuming the structure factors $f(\theta)$ (see section(3.2.1)) are modest (i.e. ≤ 5) will be the western hotspot. If this has $I=76.22$ mJy and $\alpha' = 0.53$ this gives us an $I\alpha'=40.4$ mJy and so a peak sky plane effect of order $40.4\mu\text{Jy}^{-1}$, assuming that the 'scale-size' factor is taken as unity for this compact component. The expected thermal noise for our MFS image from 104 hours of MFS observations with system temperatures of 40K on the MK1A and 75K on the other antennae of the array will be $120\mu\text{Jy.beam}^{-1}$. The spectral error contribution is therefore smaller than the thermal noise and so we can expect that the spectral errors will be undetectable. Algorithms which attempt to remove spectral effects will therefore not be needed. It should be sufficient simply to




Figure 7.3: North-South slice through the D₁ beam for the 3C179 observations.

Component	Peak I Flux (mJy)	α	Peak Ia (mJy)	Characteristic size (arcseconds)	Estimated Peak Spectral Effect ($\mu Jy Beam^{-1}$)
West Lobe A	38	0.85	32.3	< 0.6	32.3
„ B	20	0.84	16.8	< 0.6	16.8
„ C	16	0.85	23.6	< 0.6	13.6
Core	320	+0.11	35.2	unresolved	35.2
East Lobe A	76	0.53	40.4	unresolved	40.4
„ B	8.3	0.72	13.2	< 1.5	13.2

Table 7.1: Parameters of the components identified in fig(7.2b) and their contribution to the first order spectral errors

combine the MFS data together in the uv plane and then deconvolve as with conventional data.

7.2 Observations and Initial Processing

7.2.1 Observations

MERLIN MFS observations of 3C179 were made in May and June of 1986. The full six antennae MERLIN was used with the MK1A (now known as the Lovell Telescope), as the Jodrell antenna. Details of the configuration and operation of the MERLIN array can be found in Thomasson (1986).

Observations were made at the five frequencies, 1604,1632,1662,1690 and 1720 MHz, details of these observations are given in Table(7.2). The source 3C179 is a circumpolar source and observations covering all 24 hours of hour angle (minus small gaps for calibration) were made at most frequencies. The total hour angle coverage at each frequency is presented in the diagram in fig(7.4). Note that the total amount of data on 3C179 comes to the very large figure of 104 hours, this should give an improvement in map signal/noise compared to a conventional 24 hour observation of a factor of approximately 2.

The three lower frequencies were observed first with the filter in the LO system (described in section(7.1.1)) set to cover the lower end of the 7% band. After observing at these three lower frequencies this filter was adjusted manually at each antenna so that the higher half of the band was accessible, observations were then carried out at the two higher frequencies.

The observations made at 1632,1690 and 1720 MHz were straightforward and there were no major problems with these data. Observations at 1662 MHz suffered from several technical problems with the microwave links and the central control computer, observations also had to be curtailed at times due to high winds. The observations at 1604 MHz presented the most difficulty however. Both Defford and Knockin suffered with problems with their local oscillator synthesisers, these problems were probably due to the poor

Run No	Frequency (MHz)	Start Date	Start Time (UT)	Stop Date	Stop Time	Observing Time (hrs)	Comments
1	1662	21/05/86	05:00	22/05	05:00	10	No Defford Data
2	1662	22/05	07:40	22/05	20:19	12	
3	1632	24/05	01:35	25/05	00:28	22	
4	1604	25/05	01:35	25/05	09:50	8.5	
5	1604	25/05	10:35	27/05	07:22	30.5	
6	1720	6/06	13:30	7/06	11:30	21	
7	1690	7/06	13:30	8/06	11:30	22	

Table 7.2: Details of MFS observations of 3C179.




Figure 7.4: Hour angle coverage of observations at each of the 5 frequencies.
Note day/month of observations.

performance of the frequency multipliers (see sect(7.1.1)). By gradually reducing the frequency step by step from 1632 MHz to 1604 MHz whenever a ‘glitch’ in the LO system occurred it was found that it was possible to get good observations at 1604 MHz on all baselines. Overall it was found that the data was of good quality and *despite fears to the contrary no obvious interference effects were found at any of the frequencies.*

7.2.2 Calibration

As noted in Chapter 6 it appears that the relative amplitude calibration between the different frequency channels of an MFS data set will be an important requirement for the success of MFS. This will be especially true if we attempt to estimate spectral effects from the MFS data. It appears (see sect(7.1.5)) that such methods will not be needed with the present observations and so the observations will not be too susceptible to amplitude errors. Despite this the best possible calibration was attempted so as to experience the problems and limitations of amplitude calibration at first hand.

Three calibration sources were observed at least once per day at each frequency, these sources were 3C48, 3C84 and DA193 (0552+405). The source 3C48 is the standard absolute flux calibrator (Baars et al 1977). It was taken to have a flux of 13.8 Jy at 1662 MHz. The remaining two sources were observed as point source calibrators, their fluxes to be estimated using the 3C48 data.

Before doing any calibration all the data sets were first corrected for gain-elevation effects. The fluxes of 3C84 and DA193 at 1662MHz were then estimated to be 22.56 Jy and 1.984 Jy by calibrating these sources against the 3C48 data with its assumed flux. The next stage was to carry out calibration of the 3C179 data against one of the unresolved calibrators.

Despite the higher flux and the longer observation time of 3C84 it was decided to use DA193 as the point source calibrator. There were a few reasons for this decision. First, DA193 is a VLBI calibrator and a better point source at MERLIN resolution than 3C84 which is slightly resolved on the longest

MERLIN baselines. Second it is known that 3C84 has a non-power law spectrum at L-band (Kellermann and Pauliny-Toth 1981) and from section(6.5.2), we know that for the best calibration such sources should be avoided. Lastly and most importantly it was found that if 3C84 was used as a calibrator it gave systematically lower values of flux on MK1A baselines by up to 13%. This effect was subsequently traced back to an inaccuracies in the calibration program. The MERLIN system uses one bit quantisation of the signals coming from the antenna, any direct amplitude information is therefore lost. The amplitude of the signals are found as in VLBI by assuming a constant noise temperature which can be calibrated (Cohen 1973). Use of strong calibrators like 3C84 complicates matters because the calibrator is itself a source of noise, it is this feature which the calibration routine failed to properly take account of.

Using the 3C48 data a flux of 1.947 Jy was estimated for DA193 at 1662 MHz, this agrees well with the estimate of 1.950 Jy \pm 50mJy made during a VLBI observation in 1982 (P.N. Wilkinson, private communication). The closeness of these flux estimates strongly suggests that the source is unlikely to have varied much since 1982 and so it is also unlikely that the spectrum has changed significantly. A spectrum of DA193 (P.N.Wilkinson, private communication) based on observations made around this epoch is presented in fig(7.5). The spectrum is inverted with a spectral index around 1662 MHz, estimated to be $\alpha = -1.2 \pm 0.1$. This value of α together with the estimated flux at 1662 MHz was used to estimate the flux of DA193 at each of the 5 observing frequencies.

The DA193 data at each frequency was used with the appropriate estimated flux to calibrate each of the different frequency data sets of 3C179. This calibration consisted of two main parts. First baseline-related phase errors (non-closing phase errors) were determined and applied to the different baselines, the result being that all closure phases should be zero for observations of the point source calibrators. The second step consisted of using




Figure 7.5: A spectrum of the calibrator source DA193.

the estimated fluxes of DA193 to convert between measured correlation coefficients and actual fluxes. The noise limited accuracy of this amplitude calibration should be about 0.45% at each frequency.

7.2.3 Processing the MFS Data

Given that spectral effects are not important the question of processing strategy hinges on the procedure adopted for self-calibrating the data at each frequency.

One option discussed in sect(6.4.2) is ‘separate’ self-calibration in which the individual frequency channels are separately imaged and self-calibrated before combination. This method is very time consuming but has the useful side-effect that, by comparing the different images, we can get some idea of the degree of fidelity of the single frequency reconstructions.

An alternative approach, possible only with data taken over a narrow bandspread as with the present data, is to combine all the data together before the self-calibration step. This is a much quicker operation than the former method and for data observed over a narrow bandspread it should be sufficiently accurate. For the 3C179 observations both approaches were tried and both gave equally good results.

Because a good image of 3C179 at around 1662MHz already existed fig(7.2b) (Shone et al 1985), this was used as a starting model for the phase self-calibration of the individual frequency channels. The data at the each of the frequencies and the above image were submitted in turn as the inputs to the OLAF MAP program which carries out an iterative combined deconvolution and self-calibration of the data. The output CLEAN components from this process, at each frequency, were then used as starting models for the deconvolution/self-calibration of each of the frequency channels in turn.

At each frequency the MAP algorithm was applied several times in succession, each time using the corrected data and final image of the previous application as the starting model for the next. At the beginning of each application of MAP ‘windows’ (Högbom 1974) were placed based on the presence

of significant residuals at the end of the previous application. Only CLEAN components contained within these windows were used in the subsequent application of MAP.

7.2.4 The Single-Frequency Images

The single frequency images made at each of the five observing frequencies are shown in fig(7.6), also shown in the same figure is the best MERLIN only map achieved by Shone(1985). It is interesting to note that these new images are very much better than the MERLIN-only image of Shone(1985) also shown in fig(7.6). The uv coverages and signal-to-noise ratio of the data which give these new images are however not significantly better than those that gave rise to this earlier map. It is probable that the main source of the improvement comes from self-calibrating the new data against a good starting model which was made from data containing VLA observations. The FT of this starting image will have phase which is very close to the correct phase at uv radii sampled by VLA baselines. Four MERLIN antenna participate in baselines of this sort of length ($\leq 35\text{km}$) and so they will have a good model for their self-calibration. Because it appears that the dominant errors in the single frequency maps appear to be 'lobe like' and hence 'self-calibration based' rather than 'interpolation based' (see sect(6.4.2)) the image fidelity can be greatly improved by using VLA data to improve the phase solutions.

The set of images made at each of the five observing frequencies (fig 7.6) provides an interesting clue as to the reliability of conventional MERLIN images of complex sources. Previous estimates of the MERLIN image fidelity achieved in practice have relied upon imaging the same uv data independently and comparing the results (Thomasson 1986). This procedure ignores the fact that there may be systematic errors caused by employing a particular fixed uv coverage, i.e. there may be reconstruction errors that are reproducible and therefore not recognisable by comparing independent reconstructions made from the same data. In the present case the same processing was applied at

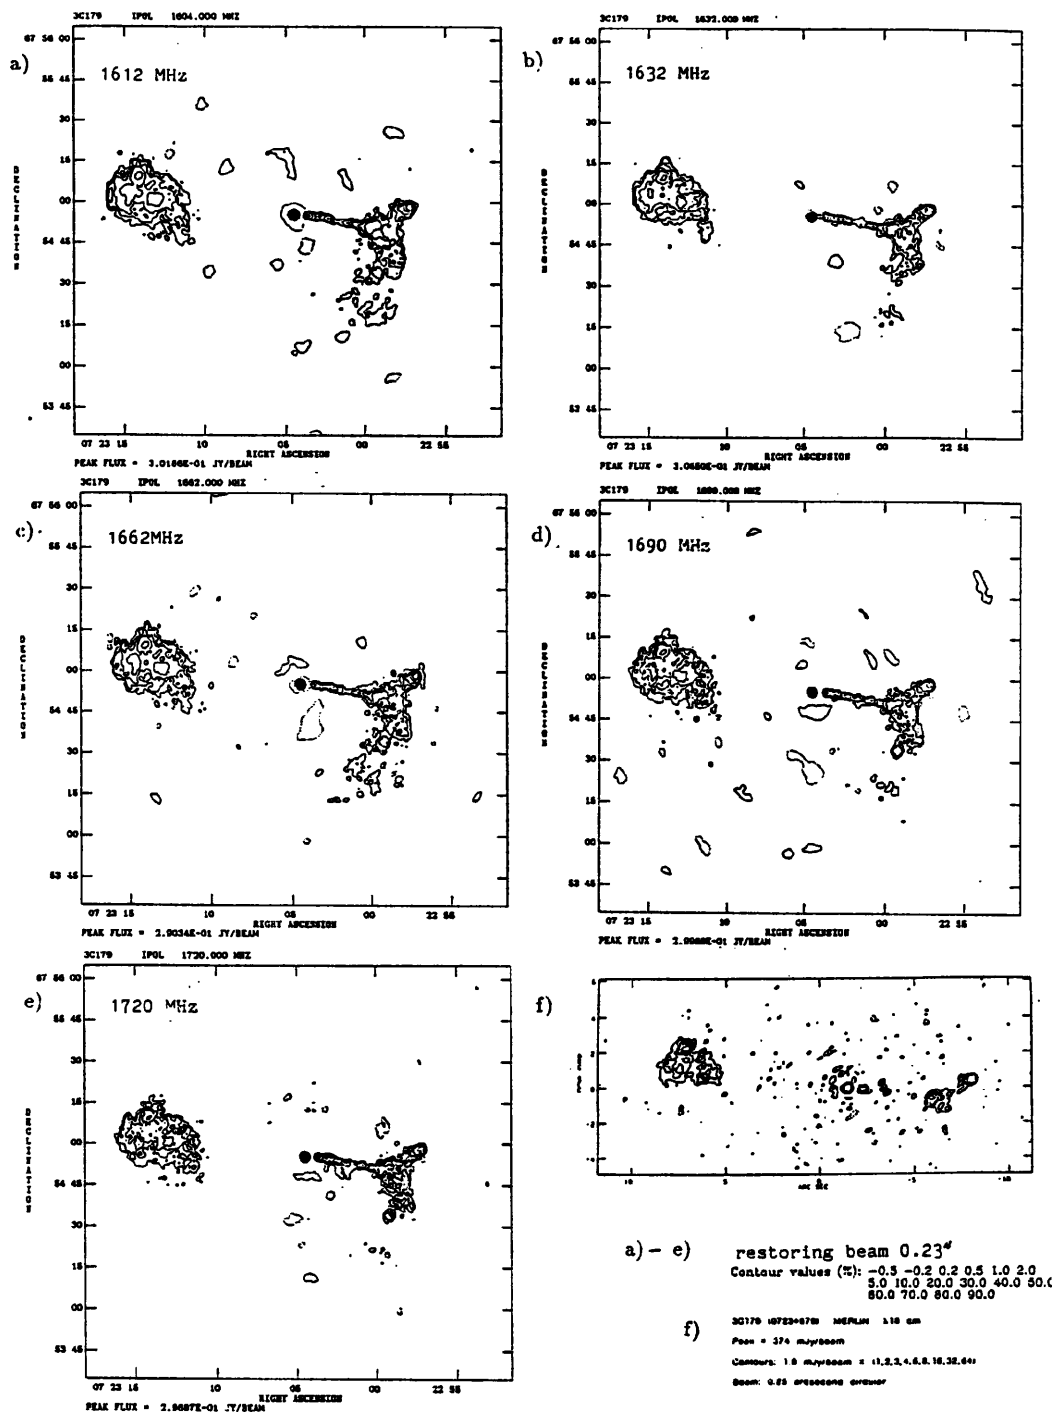


Figure 7.6: Single frequency MERLIN images at each MFS frequency. Also shown in (f) is the best MERLIN only image produced by Shone(1985).

each frequency but the uv coverages were slightly different. Noting the variations between the different reconstructions may show up influences due to these subtle changes in uv coverage.

The differences between the images are particularly obvious along the jet, the presence and prominence of the different knots within the jet varies quite considerably from one frequency to another. Given that the most common form of reconstruction error consists of a ripple-like error it is perhaps not surprising that the effects of reconstruction errors are most prominent in structures, such as the jet, which consist of approximately periodically spaced knots. It is even plausible that such a periodic structure might excite such ripple-like errors.

As well as apparently ‘CLEAN based’ errors there are also the effects of residual phase and amplitude errors which produce sidelobe-like errors; such errors are most prominent around the bright core. The structure of the knot adjacent to the core may be significantly affected this type of artifact.

7.3 MFS Imaging

7.3.1 Combining the Frequencies

Having done the preliminary phase self-calibration of the individual frequency channels we can now combine the frequencies together. Because the spectral index of the brightest point in our image (the core) was close to zero it was decided that the optimum overall spectral index correction to apply was zero, i.e. no overall correction was needed.

After combining together the uv data sets the total combined data set was transferred into the AIPS package where it was deconvolved using the MX CLEAN routine. It was found that the resulting image contained artifacts which were apparently due to residual phase errors in the uv data. To remove these phase errors one cycle of phase self-calibration was performed in AIPS using the task ASCAL. One further cycle was then executed, self-calibrating in amplitude only, and correcting amplitudes with a time-constant set to several

hours. The idea behind this final self-calibration is to remove any effects due to differences in the overall calibration levels of the individual frequencies.

The MFS image that results from the above processing is shown in fig(7.7), the restoring beam chosen was $0.25''$. This image should be compared with the single-frequency images in fig(7.6). This comparison will reveal that the MFS image is smoother, along the jet especially the reconstruction appears better. Detailed comparison between the conventional imaging and MFS imaging will however be left till section(7.3.4) where we consider SF and MFS images after adding VLA data to help in the reconstruction of extended structure.

The above image (fig(7.6)) is the end point of the processing in which all the separate frequencies are separately self-calibrated. The main advantage of the above approach is that cross-self-calibration between data at one frequency and an image at another is avoided, which is good because the effects of such processing are unknown.

The alternative to separate self-calibration is to combine the frequencies before self-calibration. This method of processing the uv data was tested by putting the combined observed uv data into MAP and using the Shone et al (1985) map as the starting model. Several applications of MAP were made using the same mapping strategy as before (see sect(7.2.3)). Finally the processed data was transferred to AIPS and a couple of cycles of MX and ASCAL executed to produce the final image. The result was very similar to that obtained by the previous processing method.

It appears that for a narrow fractional bandspread of 7%, self-calibration between data and images at different frequencies does not result in any obvious artifacts.

7.3.2 Complementary VLA Observations of 3C179

To provide complementary observations, 3C179 was observed at two wavelengths (18cm and 6cm) whilst the VLA was in 'A' array. There were two main aims behind these observations.

The observations at 18cm could be directly combined with the MERLIN




Figure 7.7: The result of combining the MERLIN MFS data and CLEAN-
ing. The restoring beam has a FWHM of 0.25".

MFS data and so fill in short uv spacings to increase the sensitivity to extended emission. In order to get the best possible image it will therefore be advantageous to combine the MERLIN MFS data with 'A' array VLA data.

From the observations at 6cm a high fidelity image of 3C179 can be obtained at a resolution between $0.3''$ and $0.35''$ i.e. at a comparable resolution to that expected from the MERLIN MFS/VLA 18cm combined image. The 6cm VLA image could then be used as an independent estimate of the source against which the relative quality of the MERLIN SF and MFS reconstructions could be judged. The multiple snapshot VLA 6cm observations have both a much better uv coverage and much better conditioning with respect to self-calibration than the SF or MFS MERLIN data. This means that the reconstruction errors in the VLA 6cm map will be negligible compared to those in the MERLIN/VLA maps and so the VLA image will be a good 'reference' image.

The details of the 18cm and 6cm VLA 'A' array observations are given in Table(7.3) in which the durations and hour angles of snapshots at 6cm and 18cm are listed. At 6cm the observations consisted of two observing bands of bandwidth 50MHz separated by 100MHz. Because the 18cm observations were to be combined with the MERLIN data the observations at this waveband consisted of bandpasses centred close to the MERLIN MFS central frequency of 1662 MHz, i.e. at 1660 MHz and 1640 MHz, both with bandwidths of 12.5 MHz. The relatively narrow bandpass widths were chosen so as to avoid possible effects from interference at these non-standard VLA observing frequencies.

Between the snapshots described in Table(7.3) calibration observations were made of the point source calibrator 0917+624, a source comparatively close to the target source 3C179 (0723+679). This source was used for both amplitude and phase calibration and was also the instrumental polarisation calibrator. Because of this last role the scheduling of observations of 0917+624 was largely dictated by the need to achieve a good parallactic angle coverage on this source (Bignell and Perley 1985). Finally a couple of observations

Wavelength (cm)	Time Observed (minutes)	HA Range	Bandwidth (MHz)
6	10	0.1-0.3	2x 50
18	8	0.3-0.4	2x 12.5
6	10	0.9-1.1	2x 50
18	8	1.1-1.3	2x 12.5
6	10	2.3-2.5	2x 50
18	8	2.6-2.8	2x 12.5

Table 7.3: Details of VLA A-array observations of 3C179

of 3C286 (1328+307) were made, this source acted as both the primary flux calibrator and as the calibrator source for polarisation position angle.

Calibration followed a standard pattern. After initial editing to remove bad uv data, the flux density of 0917+624 at each observing frequency was estimated by calibrating the calibrator source against observations of the primary flux calibrator, 3C286. Having determined a flux for 0917+624 at each of the observing frequencies the phase and amplitude of the target source data was calibrated using the 0917+624 observations with the appropriate fluxes set. To achieve polarisation calibration the cross-hand visibilities recorded when observing 0917 were processed so that the instrumental polarisation ‘cross-terms’ (Bignell and Perley 1985) between the left hand and right hand polarisation observations could be determined and then removed from the uv data. Finally the unknown rotation in polarisation position angle caused by the Faraday rotation that occurs as the signals from the source propagate through the Earths ionosphere were calibrated out. The required correction was found by comparing the observed position angle of polarisation on 3C286 with its standard value, this correction was then applied to the 3C179 data.

Mapping of the calibrated VLA data proved relatively straightforward. The data were partially phase-stable after the phase calibration so that initial inversion and CLEANing with the AIPS task MX produced reasonable images. These model images were then used to self-calibrate the uv data, two or three iterations of the above self-calibration/CLEAN cycle were sufficient to achieve good images with the measured off-source noises relatively close to that theoretically expected. At 6cm the off-source noise was measured to be $80\mu\text{Jybeam}^{-1}$ compared to the expected value of $50\mu\text{Jybeam}^{-1}$ whilst at 18cm the measured value was $110\mu\text{Jybeam}^{-1}$. At 6cm there appeared to be some low-level stripes in the final image which were apparent equally on-source and off-source. Extensive examination of the uv data failed to reveal the presence of obviously bad data or of the effects of a confusing source. This effect was eventually removed by editing out the part of the uv plane corresponding to the stripe errors.

The final 18cm map is shown in fig(7.8a) whilst the 6cm map is shown in fig(7.8b). The 18cm image shows much more extended emission than in previous images at this resolution (Shone et al 1985). The 6cm image is likewise a great improvement over the previous best published image (Owen and Purschell 1984), greyscale representations however show the image embedded in a low level negative 'bowl', a classic symptom of using data with insufficient short-spacings. This effect could be largely removed by adding 'B-array' data.

7.3.3 Combining MERLIN and VLA Data

The self-calibrated VLA 18cm data (section(7.3.2)) was next combined with the MFS data that emerged from separately self-calibrating the individual frequency channels (see sect(7.2.3)). The VLA data set was also combined with the best single frequency MERLIN data to produce a comparison image. In both cases the first step was to ensure that the MERLIN and VLA data sets were correctly aligned, this is necessary because the process of self-calibration causes the absolute position of the data sets to vary so that the positions of the MERLIN and VLA data sets on the sky will be different. To align the VLA data to be consistent with the MERLIN data the VLA data were self-calibrated against the map made from the MERLIN-only data.

The two data sets from the two instruments were now combined, whilst ensuring that the weights attached to uv points from the two data sets were the same. The combined data sets were then CLEANed using the AIPS task MX using uniform weighting to compensate for the increased density of uv points on the short VLA dominated uv spacings. The resulting image was then used as the model for the self-calibration of the input uv data, the deconvolution/phase self-calibration procedure was repeated several times until no further improvement occurred. Finally one cycle of amplitude self-calibration was applied before the final image was produced. The resulting MFS MERLIN/VLA image is shown in fig(7.9a) whilst the comparison SF MERLIN/VLA image is presented in fig(7.9b).

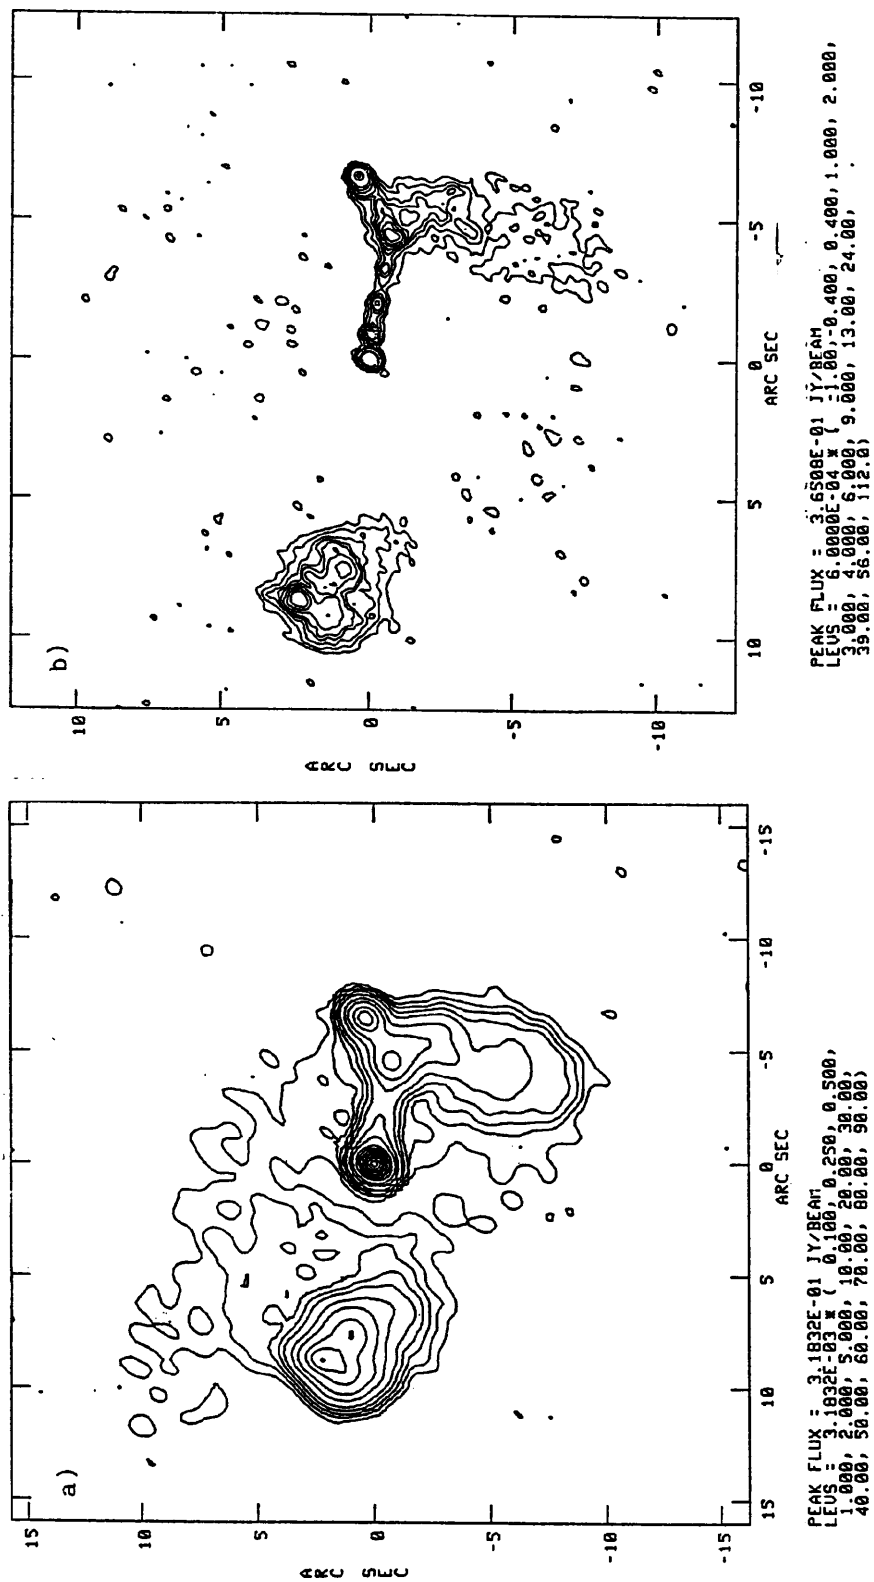



Figure 7.8: a) The VLA A-array map at 18cm. Restoring beam is 1.01'' by 0.89'' with a position angle of -31.9°. b) The VLA A-array map at 6cm. The restoring beam is circular with diameter 0.35''.

a)



b)




Figure 7.9 a) Greyscale of the final MERLIN MFS/VLA image.

b) Comparison greyscale of the best single frequency MERLIN/VLA image.

Both images are linear plots from $-10 \text{ mJy.beam}^{-1}$ to 10 mJy.beam^{-1} .

7.3.4 Comparison of MFS and Conventional Images of 3C179

Comparison between the greyscale plots in fig(7.9a) and fig(7.9b), which show the final MFS image and the final SF image respectively, shows the slight but definite superiority of the former image. Remember that a dramatic improvement is not expected given the differences between the uv coverages corresponding to each image. The uv plots in fig(7.10) show only a slight decrease in the size of the large gaps in the uv plane when we go from conventional observations to 7% MFS observations.

Direct comparison between the two images is made difficult because the MFS image is made from a great deal more data than the SF image and so the thermal noises in the two images are different. Away from the source the thermal noise in the MFS map is measured to be $175 \mu\text{Jy beam}^{-1}$ whilst for the single frequency image the noise is $272 \mu\text{Jy beam}^{-1}$. Despite this difference it is possible to see that, especially in the western lobe, the SF image contains appreciably more ripple-like reconstruction errors. In the eastern lobe, whereas the MFS image appears dominated by a single ripple error running approximately north-south, the SF image shows much more complicated behaviour possibly explained by the superposition of many ripples of different orientation (see section(4.2.2)).

Reconstruction errors of the ripple-like form described above are much more obvious in fig(7.11) which shows the effect of subtracting the MFS map from the SF map. Given that each of the above images can be considered to be equal to the true image plus reconstruction errors the above difference image represents the difference between the reconstruction error in the MFS and SF images. From this difference image it is strictly impossible to tell within which of the two images the major reconstruction errors occur. We can however expect the reconstruction errors will be larger in the SF image. It also seems plausible that ripples that occur in the MFS image will also occur in the SF image but not vice-versa. These facts mean that the difference image




Figure 7.10: The uv coverage achieved on 3C179 with (a) Combined single-frequency MERLIN/VLA observations (b) Combined MFS MERLIN/VLA observations.

should be a rough measure of the improvement achieved in fidelity in going from a SF to a MFS reconstruction.

The largest values in the difference image are found along the jet (see fig(7.11)). The difference slice along the jet shown in (fig(7.12)) shows errors with a definite wavelength with an amplitude which is modulated by the underlying amplitude of the jet emission. The wavelength of the ripple error corresponds to a gap in the SF uv coverage which is partially filled in the MFS uv coverage. Although there is general agreement between the positions of the knots in the two images the quantitative errors in the fluxes of the knots are substantial, with differences of up to 25% in the reconstructed fluxes of the knots. This means that if a high resolution spectral index map, between say a 6cm image and a 18cm image at the same resolution, were made the differences in the spectral indices determined with a MFS and SF MERLIN 18cm image would be up to 0.25. Such errors in spectral index would make it difficult to use observations made in different bands to constrain the physics of the knots. For example it would be difficult to detect the effects of self-absorption or 'spectral ageing' caused by the adiabatic expansion of the components which might constitute the knots. Using MFS might well allow these quantitative projects to be possible whereas they are impossible with conventional SF imaging.

Attempts to estimate the comparative fidelity of the two distributions by comparing both images with the 6cm VLA data have not proved as useful as was hoped when the observations were initially planned. Spectral index maps made between the MFS 18cm map and the 6cm VLA image are only very slightly smoother than the corresponding spectral index distributions made between the single-frequency 18cm map and the 6cm VLA image. The problem is that the resolution of the VLA image at $0.35''$ is quite a lot less than the resolution of the MERLIN MFS/VLA 18 cm image which has a resolution of $0.2''$. In order to make spectral index maps the MERLIN images must be convolved down and in doing this the differences between the SF and MFS images are greatly reduced.




Figure 7.11 Greyscale of the difference between Fig.(7.9a) and Fig.(7.9b). The greyscale is linear from $-2.5 \text{ mJy.beam}^{-1}$ to $3.0 \text{ mJy.beam}^{-1}$.




Figure 7.12: Slices along the jet from the first knot westward. a) A slice through the single frequency image of fig(7.9a). b) A slice through the the MFS image of fig(7.9b). c) Slice through the difference image of fig(7.11).

7.4 Conclusions

We have in this chapter considered successful MFS observations of the source 3C179. Despite the relatively narrow range of frequencies over which the present MERLIN can observe there was a slight but definite improvement in image quality over that achievable from single frequency observations. The main conclusion is that there were no unforeseen problems in the processing of the MFS data. Even in the area of phase and amplitude self-calibration, an area which was identified as causing possible problems in chapter 6 no significant problems were actually found.

Despite the success of these test observations we must however be cautious in extrapolating to a 25% MFS system where self-calibration and other problems may be much more significant. We must also recall that these observations did not require the use of an MFS algorithm to remove spectral effects. For a wide bandwidth system such processing would be required on the majority of complex targets whilst with a 7% system it appears that it would not usually be necessary. It seems that the important job of testing these MFS algorithms under realistic conditions can only be done in simulations until MERLIN is upgraded.

Despite the above qualifications the 3C179 observations have shown that MFS works in practice, albeit in the restricted but important case when the spectral effects are below the thermal noise.

Chapter 8

Conclusions and Suggestions for Future Work

8.1 Summary

In Chapter 1 a brief history of the development of aperture synthesis imaging was presented, in which the evolution from filled aperture techniques to the use of arrays with a small number of immobile randomly distributed antennae (as found in MERLIN and VLBI) was described. In order to produce high quality images from these sparse arrays non-linear processing algorithms which use *a priori* information have become vital. Although careful application of such non-linear algorithms is beneficial, if we are to image ever more complex sources with higher fidelity there must be a limit to their performance with any given array. If the imaging of ‘over-complex’ sources is attempted then significant reconstruction errors will be generated by the processing algorithms.

If we wish to achieve higher fidelity there is eventually no alternative but to improve the uv coverage of the array. The best way of achieving this is to build more antennae, however such antennae are very expensive and so very significant improvements in uv coverage by this route are ruled out by economic factors. An alternative way to achieve improved uv coverage is to note that because the dimensions of the uv plane are measured in units of wavelengths, varying the frequency of observation will allow filling of the uv plane. This technique we refer to as ‘Multi-Frequency Synthesis’.

The parameters of a useful MFS system on MERLIN are determined by the need to achieve a substantial improvement in uv coverage and by practical aspects such as interference and the limitations of RF technology. In

section(1.6.2) it was shown that a bandspread of at least 25% is required to fill all the holes in the uv plane completely. Because the proposed feed systems for improved MERLIN at L-band limit the bandspread to be less than 30%, a practical system should therefore have a fractional bandspread in the range 25%-30%.

The number of frequencies within this frequency range is limited by the interference environment, this important practical aspect of MFS was discussed at length in Chapter 6. Although more work is required it appears that at present sufficient free spectrum at L-band exists to operate at least 5 frequencies within a 25% bandspread, certainly the maximum number of useful frequencies is less than 10. Note however that there is likely to be increasing use of L-band for land-mobile and satellite communications, so that the interference environment is likely to become worse and should be monitored.

At 151 MHz the interference situation seems so bad that MFS does not appear to be practical, At 408 MHz whilst the interference situation is bad it might be possible to operate some type of MFS system with narrow frequency channels. In C-band (around 5 GHz) although the interference may be tolerable the frequency performance of the OMT ($<15\%$) restricts the performance of MFS. Although MFS at C-band will be very useful it should be a secondary priority to a system at L-band where the possible performance is greater and the need is also greater.

The one obvious problem with the MFS idea is that the source structure is not usually identical at all of the observing frequencies. MFS must rely on a substantial correlation between the structures at all of the different frequencies in order to be useful. For sources of synchrotron emission at moderate frequencies (such as L-band) it appears that there is a simple frequency dependence which can be utilised when processing MFS data. Both observation and theory show that the emission at any point in a source obeys a power law spectrum with an exponent, the spectral index $\alpha(x,y)$ which in the general case varies across the source. In the special case that the index is constant it

was shown in section(2.3.2) that by rescaling the amplitudes of the observations the effects of this frequency dependence could trivially be removed from the MFS data.

In the case that there is a spatially varying index the problem is more complicated, the spectral effects should however be removable by the using some algorithm which estimates both of the unknowns at each pixel, viz $I_o(x,y)$ and $\alpha'(x,y)$, the intensity at a reference frequency and the residual spectral index (after correcting for an 'overall spectral index' α_o .) The resulting estimate of intensity can still be of higher fidelity than an estimate from conventional single frequency imaging because even though the number of unknowns has doubled the quality of the uv coverage scales with f , the number of observing frequencies, so that for $f > 2$ the MFS solution should be better than the conventional single frequency image.

In section(2.3.3) the effects of a space-varying spectral index on the dirty map that results from inverting MFS data was calculated. The resulting dirty map can be expressed as an expansion of terms, each term being equal to a convolution of a different image (formed from $I(x,y)$ and $\alpha'(x,y)$) with a different dirty beam, the spectral dirty beam of order n , D_n . This D_n is equal to the FT of the uv coverage with each uv measurement weighted by a quantity which depends on the frequency of observation and the order of the beam. Two different expansions were described, the so called 'linear expansion' and the superior 'logarithmic' expansion in which the contributions from the terms decrease faster with order.

For the proposed MERLIN L-band MFS system the expected contributions from spectral effects could be calculated based on the above expansions. As a rough estimate if the dynamic range required is less than $200/\alpha'_{max}$ where α'_{max} is the value of α' on the brightest parts of the image (assuming the effects of any bright core has been removed) then it can usually be assumed that the spectral effects are below the noise. Normal deconvolution applied to such dirty images will produce satisfactory images. The MERLIN MFS observations of the source 3C179 presented in Chapter 7 fell into this class and

simple combination and deconvolution seem to have been sufficient to produce an MFS image free of any spectral effects. This image showed a slight but definite superiority to the best single-frequency image.

For sources in which higher dynamic ranges are required the spectral effects will become significant. It can be shown that provided that the dynamic range is below $2000/\alpha^2$ then the dirty map is well described by a two order approximation to the logarithmic expansion, the first term being the normal expression for the dirty map ($I \cdot D_0$) and the second containing all the significant effects due to a space varying spectral index ($I\alpha' \cdot D_1$). Although errors significant with respect to noise may occur at dynamic ranges below this level these errors will only occur in bright regions in which the fractional error is very small and hence tolerable. Virtually all complex sources which can be imaged with MERLIN will fall within the above dynamic range limit.

The use of the above expansions allows quantification of the effects of departures from exact power law spectra. In chapter 6 it was shown that at L-band, except for unresolved cores, departures from power law spectra due to the effects of spectral ageing, synchrotron self-absorption and free-free absorption are not significant. The second order effects that are introduced by these sources of 'spectral curvature' are generally no larger than the second order effects due to perfect power law spectra. It does not appear that the effects of 'spectral curvature' are generally significant.

Radio cores have complex spectra over decades of frequency but even these have comparatively smooth spectra over the 25% range of frequency used for MFS. An estimate of the worst case suggests that at maximum such cores will cause harmful second order spectral errors at levels of 500:1 or higher. Thus if radio cores exist with complex spectra which are brighter than 500 times the noise then it will be necessary to remove the effects of these cores from the MFS data. This can be done by imaging the data at each frequency, estimating the core flux and then removing the effects of this unresolved component from the data at each frequency.

Assuming our MFS imaging problem is one in which the required dynamic

range (of the brightest region *after removing the core*, to the noise) is greater than $200/\alpha'_{max}$ but less than $2000/\alpha'^2_{max}$ then spectral effects will be important but the two order expansion will still provide an accurate description of the dirty map. For images in this range it is important that some algorithm be devised to remove the first order spectral effects.

In chapters 4 and 5 various approaches to the above problem were explored. One approach, first suggested by Cornwell(1984), is to weight the measured MFS uv data in such a way as to eliminate the spectral effects entirely, various variations on this theme were considered but all had the same fault; in eliminating spectral effects a serious degradation of the effective total uv coverage occurs. Since the improvement of this uv coverage is the sole aim of MFS this is a very serious fault. Whilst 'weighting methods' may be useful in cases of very good uv coverage or special geometries which internally scale with frequency these methods do not seem to be applicable to arrays like MERLIN, indeed tests show that the degraded effective uv coverage can be worse than the single frequency uv coverage making MFS pointless.

Another approach consists of imaging the individual channels separately, these images can then be combined in a variety of ways to form an image better than any of the single-frequency images. The most direct way to use the 'channel maps' is to combine them in the sky plane in such a way as to eliminate spectral effects, over a narrow bandspread this is equivalent to forming the mean of all the images. This approach is inadequate in that it does not exploit the non-linearity of deconvolution, the composite image will at best be only a statistical improvement in fidelity upon the single frequency reconstructions. A more subtle way of using the channel maps is to use them to estimate the spectral index distribution of the source, once this distribution is estimated it can be used to correct the spectral effects within the combined MFS uv data before the data is Fourier inverted and then deconvolved. The reasoning behind this approach is that the spectral index is best estimated in the sky plane using the images at each frequency whilst the deconvolution is best done upon the data *in toto*. The problem with this approach is that the

effects of reconstruction errors in the single frequency images are fed back into the MFS data, the performance of the algorithm is again therefore ultimately tied to the quality of the single frequency reconstructions.

In chapter 5 the 'Double Deconvolution' algorithm for removing spectral effects was presented. By appropriate weighting of the uv data either of the two distributions which contribute to the two order expansion can be emphasised. The contribution due to the desired distribution can be made to be compact whilst the effects of the undesired distribution are scattered over the image. By employing non-linear deconvolution algorithms which have an *a priori* bias toward compact structure the wanted distribution can be separated from the unwanted distribution. By emphasising in turn the I and $I\alpha'$ contributions a successful iterative algorithm for separating out the effects of the two images can be devised. For a particular form of the weighting functions the double deconvolution algorithm can be considered as a generalised form of CLEAN, i.e. as a method for recognizing a particular artifact in the dirty image and removing its effects. The algorithm uses the idea of searching for the highest correlation between a dirty image and a dirty beam as first proposed by Högbom(1984) and generalizes this principle to the situation in which there is more than one type of dirty beam in the image.

Some theoretical analysis of the DD method was attempted in section(5.3), amongst the conclusions are that

- (1) The DD algorithm is guaranteed to converge to find a solution that fits the data provided only that the normal requirements for a conventional CLEAN deconvolution as specified by Schwarz(1978) hold.
- (2) Under certain conditions of 'oversampling' the solutions for I and $I\alpha'$ will be unique and correct to within the noise i.e the DD algorithm will successfully separate the effects due to I and $I\alpha'$.
- (3) If the above condition does not hold then the division between I and $I\alpha'$ will depend on the decisions as to when to switch between estimating I and $I\alpha'$ distributions. Two different types of termination criteria can be considered, those which depend on stopping CLEANing at a particular residual level and

those based on stopping CLEANing when CLEAN components have been found over a set area. Although the second of these termination requirements is probably superior the first is easier to implement and there is also the advantage that it has proved possible to derive mathematically an optimum stopping criterion. In practice however it may prove difficult to implement this criterion.

Simulations with the DD method have showed that the algorithm can reduce the effects of spectral errors sufficiently so that even in the most demanding of circumstances spectral errors could be made smaller than the thermal noise. Images similar in fidelity to those possible if no spectral effects were present in the MFS data were shown to be recoverable. Although the reductions in spectral effects achieved by DD were not dramatic (less than a factor of 10) the spectral effects were reduced to less than the noise. More difficult models must be devised in order to determine the ultimate limits in performance of DD, the present simulations do however show that the application of DD to MERLIN MFS data is sufficient to eliminate spectral effects in most complex MERLIN targets. The combination of the intrinsically low level of spectral effects and DD ensures that the basic MERLIN MFS problem is essentially solved.

8.2 Priorities for Future Work

Although significant progress has been made both in understanding MFS and in demonstrating a practical MFS system, significant problems still exist and these must be explored. Below is a summary, in approximate order of priority of the major areas in which work needs to be done, both to investigate problems and to extend MFS.

8.2.1 Investigation of the Effects of Phase Errors on MFS

Although the simulations of MFS carried out so far are encouraging they do

as yet leave out the effects due to phase errors caused by propagation effects in the troposphere and ionosphere. Normally these phase errors are substantially removed by applying so called 'self-calibration' techniques, so far no attempt has been made to apply both DD and self-calibration together. The separate self-calibration and imaging of the data at the different frequencies prior to applying the DD algorithm is one simple way of combining the two algorithms. The problem with this approach is that the separate self-calibrations will allow the accumulation of 'self-calibration based' errors whose levels are controlled by the single-frequency uv coverage and not by the MFS uv coverage. These errors may cause any subsequent MFS image to have a fidelity which is only a statistical improvement upon the errors within each single-frequency reconstruction. An alternative to the above 'separate' self-calibration is 'parallel' self-calibration in which self-calibration is implemented in a loop containing a double deconvolution step, this will however require a more intimate relationship between self-calibration and DD with the possibility of an adverse interaction between the two algorithms. Even if the above implementation works properly the susceptibility of MFS to self-calibration errors will probably still be greater than for a single frequency array of similar uv coverage.

Simulations of conventional imaging should be done to find out the extent of 'self-calibration' based errors in reconstructions of complex sources with the improved MERLIN. Simulations of both 'separate' and 'parallel' self-calibration implementations should be done to see the effects of phase errors in practice. The time consuming process of doing simulations of complex sources with phase errors has up to recently been prohibitive, with the advent of larger computing resources and a more efficient, user friendly version of DD (see below), these simulations should however be accomplished more easily in the future.

An analysis of the time spectrum of phase errors in real data should be made. This, together with a clearer idea of the size and typical timescales of

the residual phase errors arising from self-calibration, should enable a decision about the usefulness of 'linked' self-calibration to be made (see section(6.4.4)).

8.2.2 Amplitude Errors and MFS

Errors in amplitude are potentially a serious problem for MFS. It is thought that time variable amplitude errors will interact more strongly with DD than do phase errors and so it may be necessary to do amplitude self-calibration separately on each frequency to prevent interaction. The relative overall amplitude calibration at each frequency will be yet another problem, possibly made worse by the need for separate self-calibration to remove the time variable errors. It is thought that the relative amplitude scales should be correct to at least 1%, and ideally 0.2%, if it is wished to successfully separate intensity and spectral effects. Much greater attention to amplitude calibration will be necessary if MFS is to be completely successful, however the necessary accuracies, although stringent, should be achievable with care and attention.

Simulations of the effect of amplitude errors are required. In addition a general effort should be made to improve the frequency and accuracy of external calibration and implement noise-diode calibration. The effects that cause time dependent amplitude errors should be traced and reduced so as to reduce the need for amplitude self-calibration. All of the above improvements will help conventional imaging as well as MFS and should be part of a general improvement in data calibration.

8.2.3 A User-Friendly Implementation of the DD Algorithm

At present the DD algorithm exists as a loose assortment of tasks. A coherent software package should be written in one of the standard analysis environments so as to increase speed and produce a user friendly and exportable facility. The implementation should be in the 'difference' mode as described in section(5.5) which will further decrease running time.

8.2.4 Understanding and Developing DD

Despite the success of DD in removing spectral sidelobes in the case of the target source it would be nice to know the ultimate performance limits of DD in practice and to have some clearer idea of the ‘gain’ of a cycle of DD. There is also much work to be done to optimise automatic stopping criteria and weighting functions.

Theoretical analysis should concentrate on optimising the criteria for cycle termination, especially in the area limit case. The various ways of implementing residual level and beam area termination should be implemented and tested, possibly with models which allow the ultimate limits of DD to be probed. Different weighting schemes can also be tested and a DD algorithm which uses MEM as the deconvolution step tried.

8.2.5 Developing ‘Phase Linking’

One possibility that in principle can be used if MFS is implemented in the rapid switching mode is to use the fact that the phase errors at the different frequencies will not be independent. There will be a frequency dependence of phase error which will depend on the relative contributions from the ionosphere and troposphere. A crude test with MERLIN, switching over only a small frequency range has shown that the phase errors are highly correlated as expected. Whether algorithms which use ‘phase-linking’ will be useful will depend on the rate at which phase errors change and on the size and timescale of the self-calibration errors which they seek to constrain. Depending on the results of the proposed analysis of the phase error time spectrum and the size of residual phase errors in self-calibrated data this technique might turn out to be useful. If it shows promise then algorithms along the lines suggested in Chapter 6 might be developed and tested.

8.2.6 Further Interference Surveys

Interference surveys are required at Cambridge and Chilbolton where the interference environment is possibly quite different from that at the present MERLIN antennae. In addition more sensitive surveys would be useful at the sites already surveyed. With the above information the RF system can be better planned and the expected increase in interference monitored.

8.2.7 Core Subtraction Simulations

When there is a strong core with a complex spectra it will be necessary to take special measures to remove the effects of this core. This can be done by mapping each frequency separately and then subtracting a core model, although in principle this should be easily done there may be special problems especially if the core is embedded in extended emission, simulations would again be useful.

8.2.8 Extensions of MFS

Present 408 MHz observations on MERLIN can suffer from a lack of short spacings, the ideal remedy would be to combine VLA data with the MERLIN data. Unfortunately because of the interference situation at the two sites the VLA cannot observe at 408 MHz and MERLIN cannot observe at 327 MHz, the VLA's standard P-band frequency. Combining MERLIN 408 MHz and VLA 327 MHz data whilst employing some form of MFS processing might however allow enhanced P-band images to be made. Some of the ideas presented in this thesis may be usefully employed in finding such a processing algorithm. The problem is more difficult than the MFS problem considered in this thesis because there will be much less overlap in the uv plane between the uv tracks at the different frequencies.

The problem of imaging a source whose structure changes significantly during the period of observation (i.e. SS433 with MERLIN/EVN or Cyg-X3 with MERLIN) can be analysed by an expansion method similar to the

spectral problem (R.C. Vermeulen, private communication). The equivalent two order approximation will contain a time constant image and an image equal to the first derivative of intensity with time, each coefficient image will have its own dirty beam, in principle an algorithm similar to DD could be employed to separate the two distributions and hence remove the artifacts introduced by time variability.

8.3 The Impact of MFS on MERLIN and VLBI

8.3.1 Implementing MFS on Long Baseline Arrays

MFS has its greatest impact on sparse arrays of telescopes which are distributed in non-regular ways and which cannot be moved. Although MFS can be usefully applied to other types of array (see section(8.4)) it is on arrays of the above type that the imaging potential is most enhanced.

This thesis has concentrated on MFS as applied to MERLIN, an array with a geometry similar to that of the sparse arrays found in VLBI, many of the conclusions drawn from MFS on MERLIN can therefore be directly transferred to the VLBI case. There are however differences particularly in the way MFS is put into practice.

Like MERLIN the ad-hoc VLBI arrays which have grown up (such as the EVN, US network etc.) consist of relatively small numbers of antennae whose positions are not optimised for uv coverage. Even the VLBA (Very Long Baseline Array) presently under construction in the US (Romney 1987), although designed deliberately as a dedicated VLBI instrument, will only have a total of 10 antennae. Although some effort has been made to optimise uv coverage the constraints of geography and economics ensure that the uv coverage falls far short of that achievable with connected element arrays such as the VLA. MFS applied both to the present arrays and to the VLBA can produce large improvements in uv coverage. Such improvements will be needed as the complexity of VLBI images increase. Until recently VLBI images have been 'dynamic range limited' and the theoretical thermal noise sensitivities have

rarely been reached. The effects of 'non-closing' errors which limit dynamic range are slowly being reduced, if we add the effects of expected sensitivity improvements then it seems certain that VLBI images will become ever more complex (e.g. Wilkinson 1987).

The design of the VLBA will allow many frequencies to be recorded simultaneously spread over a wide range of frequencies, this design has been adopted partly to allow for geodetic bandwidth synthesis (see section(1.4.1)). Even in the standard radio astronomy mode a multichannel bandwidth of 64 MHz will be standard and 128 MHz will be possible, at frequencies lower than L-band these large bandwidths will give a significant fractional bandspread and will be useful for MFS, although the problem of interference will be severe. The design of the feed systems will allow a 30% bandspread at L-band and a 10% bandspread at C band. The expected MFS uv coverage expected from the VLBA assuming a 30% bandspread and 5 frequencies is shown in fig(8.1), the declination is 40° . The improvement in uv coverage due to MFS is even more pronounced at lower declinations (i.e. $< 30^\circ$).

For the 'ad hoc' arrays such as the EVN the implementation of MFS is likely to be harder, each separate antenna will have its own feeds, filters and LO generation system so the available frequencies on the different antenna will all be quite different. The present MkIIIA recording system allows observations over a total bandwidth of 56MHz, with the bandwidth split into a number of contiguous channels. For low frequency observations at 327 MHz and 610 MHz such a 56 MHz bandspread corresponds to a fractional range of 17% and 9% respectively, implementation of MFS at these low frequency bands in this mode will almost certainly be impossible however due to the effects of interference.

At least initially the most practical way to operate MFS on the EVN would be to carry out separate observations in separate sessions at 18cm and 21cm respectively (a fractional bandspread of 16%). If the 56MHz MkIIIA system is used centered on each of the two frequencies then we obtain a many frequency MFS system with a total bandspread of 19%. Interesting observations of radio




Figure 8.1: The UV Coverage of the 10 station VLBA at a declination of 40°. a) Single frequency uv coverage. b) MFS UV coverage with 5 frequencies spread over 30%.

jets such as that in M87 would be possible using observations at 1420 MHz and 1660 MHz. Given the small variation of spectral index along jets (Bridle 1982) and the relatively small dynamic range of VLBI observations spectral effects would often be insignificant and no special MFS processing would therefore be needed. For the specific case of M87, where the bright knots have variations in α of less than 0.1 (Junor 1987), such a two frequency MFS observation would be free of spectral effects down to a dynamic range (after removing the core) of about 3000:1.

8.3.2 Astronomical Targets for MFS with MERLIN and VLBI

The prime targets for high fidelity imaging with MERLIN and VLBI are extragalactic radio galaxies and radio-loud quasars. Also of interest are some bright galactic objects i.e. SS433 and Cyg X-3 and also small separation gravitational lenses.

The improvements possible in images of typical extended radio sources by using MERLIN MFS have already been demonstrated in Chapter 5. By using MFS on the extended MERLIN such sources can be reliably imaged with a useful resolution of 0.12" at L-band, with single-frequency imaging however many of the details are lost. There should be the possibility of reliably detecting new radio features in these sources, such as the filaments seen in the test source. For other tasks such as determining the spectral properties of the radio emission and carrying out comparisons with images at other wavebands there will be particularly strong requirements on fidelity. To do more than simply detect distinct components will require much better aperture coverages than are presently available.

The imaging of radio jets with the 'stand-alone MERLIN' is one area of high fidelity imaging where MFS will be particularly useful. Such jets can also be imaged with the combined MERLIN/EVN, the VLBA and trans-Atlantic VLBI, providing the possibility of imaging the jets on all scale sizes. These radio jets appear continuous on all scales, e.g. as in 3C120 (Walker, Benson and

Unwin 1987). To detect the wide range of complex phenomena theoretically expected within these jets will require high fidelity images. Particularly at the highest resolutions proper motion studies will also become important, these will require particularly high fidelity and will especially benefit from MFS.

Because the objects of interest for MERLIN and VLBI are similar and because the expanded MERLIN will increasingly overlap and be combined with EVN observations the areas of astrophysics to which MFS will contribute on MERLIN and VLBI will be described together.

8.3.3 Observations of Jets

MFS imaging of jets will be relatively easy because spectral index variations along jets are usually quite small. At kiloparsec scales, Bridle (1982) finds that $\alpha=0.6 \pm 0.15$. At MERLIN/EVN L-band resolution the relatively nearby source M87 shows spectral index variations of the knots within this jet between $\alpha=0.37$ to 0.50 (Junor 1987). It also seems probable that the spectral variations along the jets at VLBI resolutions are small (Marscher 1986), at least for Earth-based centimeter wavelength VLBI, although there are as yet no convincing spectral index maps at milliarcsecond resolution. Only at the sub-milliarcsecond scales accessible to millimeter VLBI or Quasat is it possible that there might be very complex spectral effects (Aller, Hughes and Aller 1987) which could cause problems for MFS.

Amongst the overall properties of jets that require high fidelity imaging are the width and opening angle of the jet as a function of distance, possible structure across the jet (Junor 1987) and the overall brightness as a function of distance along the jet. These parameters can be used to constrain the regions where the synchrotron emission is generated, the ways in which jets are confined and the overall flow speed of the jet. Bends caused by source motion within a medium and wiggles within jets possibly caused by Kelvin-Helmholtz instability can also be used to constrain jet properties both at kiloparsec scales (Davis, Muxlow and Conway 1985) and at VLBI scales.

The possible effects of environment on radio jets can be studied with both

MERLIN/EVN and VLBI. The former instrument can probe the behaviour of jets whilst they transverse the ISM of the host galaxy and possibly interact with clouds, causing shocks observed as knots (e.g. Blandford and Königl 1979) or sudden deflections. Jet interactions with the ISM may be especially significant in the ‘compact steep spectrum’ (Van Breugel, Miley and Heckman 1984) sources, which one theory holds are caused by the radio jets being confined within a galaxy. Radio/high resolution optical comparisons may also be revealing in these sources (see sect(8.3.5)). Long baseline VLBI can investigate the effects of the BLR environment on jet behaviour on smaller angular scales.

All of the effects described above can be modelled on computers in large scale hydrodynamic simulations of radio jets (Norman, Smarr and Winkler 1982, Williams and Gull 1984, Wilson and Falle 1985, Lind 1986). It may well be that simple analytical analysis cannot be easily applied in the case of real jets. If this is the case then the comparison of high fidelity images and the results of simulations may provide the best constraints on jet physics.

8.3.4 Radio Spectral Studies

The simulations of chapter 5 showed that using MFS on a complex source allows images to be made which have much smaller fractional errors than those from single frequency observations. These improved images could be usefully used in spectral index studies of sources and in studies of spectral ageing at high resolution in which very accurate images are required. Previous studies which have used MERLIN and VLA images in spectral ageing studies (Stephens 1987) have found that the quality of the results have been limited by the fidelity of the MERLIN images. Full resolution MERLIN MFS at 408 MHz and tapered MERLIN MFS at L-band would be of comparable fidelity to VLA images at 0.3'' resolution and so give more accurate spectral age images.

Low frequency spectral effects such as ‘free-free’ absorption (Lang 1974) and synchrotron self-absorption may be detectable in compact regions by comparing images from the VLA at 2cm A-Array, MERLIN at L-band and the

EVN at 327 MHz. Particularly interesting would be the detection of low frequency turnovers in hotspots or knots. If due to SSA the frequency at which turnover occurs will constrain the source size whereas if due to free-free absorption it will constrain the presence of intervening thermal electrons.

8.3.5 Radio/Optical Comparisons

One of the most important areas in which high fidelity will be at a premium is that of radio/optical comparisons. With the advent of the Hubble Space Telescope (HST), diffraction limited optical continuum and line images will be available with resolutions of 50 mas at 5000 Å and up to 20 mas at 2000 Å in the near ultra-violet. These optical images will have comparable resolution to radio images from MERLIN/EVN at L-band and from MERLIN alone at C band. Images at the resolution of the HST will probably show interesting optical structures in quasars which are presently unresolved and these structures may be correlated with radio structures.

Many of the mechanisms that cause correlations between optical and radio emission will cause variations in the radio flux which are comparatively small compared to the background radio emission and so high fidelity radio images will be required to detect the radio/optical correlations.

Thermal gas emitting optical lines in front of or mixed in with radio emitting plasma will produce a variety of subtle effects. Kundt and Saripelli (1987) have considered free-free absorption by thermal gas as a mechanism for causing regions of reduced radio emission in Cygnus A. Kundt and Saripelli also consider Thompson scattering of radio emission by electrons in foreground thermal gas to explain these features.

The presence of foreground line emitting gas is often correlated with lower levels of fractional polarisation either due to internal or external Faraday depolarization (Heckman et al 1982).

Other direct correlations between optical and radio can occur due to optical synchrotron emission (Bridle and Perley 1984) and possibly optical emission

caused by inverse-Compton scattering of radio photons to optical frequencies (Okoye and Obinabo 1982).

Indirect correlations between optical and radio emission can occur due to the interaction of jets with interstellar clouds. Radio jets can have their magnetic fields compressed on hitting interstellar clouds so causing an increase in radio flux, likewise interactions can drive shock waves into the clouds causing shock ionization and hence optical line emission (Van Breugel, Heckman and Butcher 1984).

8.3.6 Proper Motions

High fidelity imaging will be required for proper motion studies and for monitoring the time evolution of radio components.

The milliarcsecond jet of M87 has recently been found to have a proper motion of 1.5 mas yr^{-1} (Biretta et al 1988) corresponding to a velocity of $0.3c$ ($H_0 = 75 \text{ km s}^{-1} \text{ Mpc}^{-1}$). If the kiloparsec jet continues to move at this velocity then the combination of the EVN and the extended MERLIN (resolution 20mas) could detect proper motions of knots in the jet of M87 in 5 to 10 years (MERLIN proposal 1988). Superb uv coverage would be required on both MERLIN and EVN length baselines in order to detect the motion of complicated knots which are much larger than their proper motion. Masson(1986) has shown that by using multi-epoch observations it is possible to detect proper motions of at least a factor of 10 less than the resolution, even in complex images, provided the uv coverage and hence the fidelity is excellent. MFS could be usefully employed on both MERLIN and the EVN to achieve the necessary uv coverage for proper motion studies in complex sources.

At VLBI scales components can have significant proper motions and show a general evolution with time, the monitoring of the time variability of structure is one area in which very high fidelity imaging is required, especially if it is desired to detect velocities well downstream where components rapidly fade and disperse. High fidelity multi-epoch imaging can also provide invaluable

astrophysical information if the deacceleration of components can be detected or if we can decide between components following ballistic or 'channel' paths.

8.4 The Impact of MFS on Other Arrays

8.4.1 MFS on the AT

The Australia Telescope (AT) presently being built in eastern Australia could usefully employ the MFS technique. The AT itself consists of two quite separate parts, the Compact Array (CA) at Culgoora will be an East-West array of conventional design with a maximum baseline of 6km whilst the long baseline array (LBA) is a 5 element array of antennae with a maximum baseline of 1400km (Norris 1987).

The number and distribution of antennae in the LBA is similar to the present MERLIN system and hence the quality of the uv coverage also similar except that the very pronounced North-South bias in the distribution of the antennae means that the uv coverage for sources at even moderate declinations shows a pronounced 'bow-tie' effect, so that the coverage at azimuths close to the v axis will be poor. Many of the improvements in uv coverage achievable with MFS on MERLIN can be duplicated on the AT LBA although the bow-tie effect will of course not be reduced by MFS. The addition of significant East-West baselines would substantially improve the impact of MFS on the LBA.

The compact array will consist of two sections of rail track orientated East-West, the first section 3km long will carry 5 antenna whilst the second track (approximately 100m long) will lie 3km from the end of the first track and will carry a single antennae. The 6km maximum baseline will give an array with a resolution of $6.2''$ at 21cm wavelength. The antennae will be movable but they will only be placed at discrete points or 'stations' which number in total 37. An almost complete synthesis will be possible by using observations taken over many days with antennae moved every day or every twelve hours or so. In order to achieve good uv coverage observations must be spread over large

amounts of time such as 20 days. By using MFS to alter the baseline lengths rather than moving antennae equivalent uv coverages could be achieved in a very much shorter time. With a many frequency MFS system excellent uv coverage could be achieved in a single 12 hour observation.

The advantages of using MFS on the AT compact array have already been considered by Manchester(1985) where the effects of MFS on reducing sidelobe levels of the synthesised beam were calculated and shown to be significant. Because of the inconvenience of having to move antenna to obtain good uv coverage MFS may be a very useful technique for the AT compact array. The possibility of using MFS has already been recognised and a wide bandspread capability has already been built into the hardware so that a 30% system at L-band is easily achieved. Given the priority given to spectroscopic observations with the AT the site has been chosen so that interference is small over a wide range of frequencies and this is obviously encouraging for the possibility of using MFS.

The astrophysical aims of AT compact array continuum imaging will include the imaging of the larger southern radio galaxies and in this MFS will play a useful role. One other class of objects which will make especially good targets for MFS will be supernova remnants (SNR's) in our galaxy and in the Magellanic clouds. SNR's will be amongst the most complex objects that will be observed with the AT and they will therefore require excellent uv coverage. Because the spectral index distribution across the face of SNR's appears to be remarkably uniform (Sieber et al 1981) they will be particularly easy targets for MFS observations; processing will simply consist of rescaling the amplitudes of the data at each frequency as described in section(2.3.2). Of particular interest in the southern hemisphere are the SNR's in the Magellanic clouds (Mathewson et al 1983) which provide a complete sample of SNR's much less effected by selection effects than SNR's within our own galaxy. There should be dozens of suitable SNR targets in the Magellanic clouds which will benefit from MFS observations.

8.4.2 MFS on the VLA

MFS observations have already been carried out with the VLA (Braun, Gull and Perley 1987) of very complex supernova remnants (such as Cas A), these sources have already been noted as good targets for MFS because of their simple frequency behaviour. Because the long track conventional uv coverage of the VLA is so good it is only on the the most complex sources such as the most complex SNR's that MFS will make a really significant impact. SNR's are a particularly important target for the highest fidelity imaging because the dynamics of shell expansion can be found only with the highest fidelity images at several epochs. Other targets in the same class of complexity are the galactic centre and the radio galaxy Cygnus A.

Although useful when mapping simpler sources with shorter tracks MFS will not provide the really dramatic improvements in uv coverage possible with sparse arrays, MFS tends to help fill radial gaps in uv coverage, for the VLA the radial filling is already quite good. Short 'snapshot' observations will suffer from azimuthal gaps and increasing the radial filling by MFS will therefore only have a limited impact on these observations. MFS might be useful in improving the uv coverage of a 4 hour observations to the quality of a 24 hour synthesis and so MFS may be of use in increasing the operational flexibility of the instrument but except for the most complex sources it will not itself improve upon what is possible by conventional imaging.

8.4.3 MFS on the GMRT

The GMRT (Giant Metre-Wave Radio Telescope) presently under construction in India (Swarup 1987) will consist of 34 large antennae (45m diameter) primarily designed for making observations at 610 MHz and below, although observations above 1 GHz may be possible. Of the 34 antennae, 16 will be placed within a dense pack of diameter 1 km and will be primarily designed for pulsar searches whilst the remaining 18 will be distributed in a Y configuration similar to the VLA. Each arm of the Y will have 6 antennae and the

arms of the Y will be 14 km long, so that the array will have a resolution of 5'' at 610 MHz.

Although many of the most interesting scientific aims of the telescope consist of observing line emission (i.e. highly redshifted HI and deuterium) to which MFS cannot contribute, MFS can make a contribution to continuum studies of large diffuse extragalactic sources and to some galactic sources including SNR's. Large extragalactic sources particularly tend to be more complex (in terms of beam areas) at low frequency and with the large collecting area and sensitivity of the GMRT the uv coverage may well be usefully supplemented by MFS methods. The site of the GMRT in northern India will have an excellent interference environment which will be ideal for low frequency MFS. One effect that may complicate MFS at low frequency will be the impact of synchrotron self-absorption and free-free absorption described in chapter 6 which will cause departures from perfect power-law behaviour at low frequencies.

8.4.4 MFS on the 5km Telescope

The Cambridge 5km telescope (Ryle 1975) is presently being upgraded by increasing the number of baselines that are correlated and by dramatically increasing the observing bandwidth by a factor of 35 to 350 MHz. The larger bandwidth will greatly improve the sensitivity of the array in particular so that the cosmic microwave background can be observed and fluctuations searched for. A secondary objective for the improved array is to make observations of large radio galaxies.

The 350 MHz bandwidth will be split into many channels smaller than 20 MHz in width to reduce the effects of bandwidth smearing (Cotton 1985) and allow for the rejection of interference. The multichannel system that results will effectively be an MFS system. At the lower observing frequencies of 2.7 GHz and 5 GHz the fractional bandwidth will be 13% and 7% respectively (though interference at 2.7 GHz may partially restrict the use of all channels). Whilst improved uv coverage is not the prime motivation behind the increase

in bandwidth the filling of the uv plane that results is appreciable (see fig(8.2)). For the imaging of complex extragalactic sources the improvements in uv coverage that results will be useful given the small number of baselines that can be formed without moving antennae.

For the prime goal of measuring fluctuations in the microwave background the appreciable size of the fractional bandwidths may be a problem. The spectrum of the background is known so the effects of this can be easily removed. Very sensitive measurements of the microwave background however may be affected by foreground sources (especially in the region of clusters in which observations of the Sunyaev-Zeldovich effect are being attempted) which will contribute both conventional sidelobes and spectral sidelobes. It might be necessary to use the methods outlined in this thesis to remove the spectral effects of these confusing sources.




Figure 8.2: The uv coverage of the upgraded 5km telescope at 2.7 GHz and declination 50°. a) Single Frequency uv coverage. b) Effect of splitting the 350 MHz bandwidth into five frequencies.

Appendix A

The Derivative Relationship for Linear SDB's

In section(2.3.3) we saw that the n th order linear spectral dirty beam D_n is such that

$$D_n(\theta, \phi) = \frac{1}{N} \sum_{i=1}^N (\Delta\nu_i/\nu_o)^n D(\theta, \phi, \nu_i) \quad (\text{A.1})$$

where $D(\theta, \phi, \nu_i)$ is the single frequency dirty beam due to the uv data observed at frequency ν_i , expressed as a function of the polar coordinates, radial distance θ and azimuth ϕ , and N is the total number of frequencies observed. If it is assumed that there is exactly the same hour angle coverage at each of the observing frequencies then the single frequency beam at any frequency ν_i will be the same as a single frequency beam at the reference frequency ν_o except for a change in radial scale. We may write formally that

$$D(\theta, \phi, \nu_i) = D\left(\frac{\nu_i}{\nu_o}\theta, \phi, \nu_o\right) = D(\theta'_i, \phi, \nu_o) \quad (\text{A.2})$$

where $\theta'_i = (\nu_i/\nu_o)\theta$. let us define $\Delta\theta'_i = \theta'_i - \theta = [(\nu_i/\nu_o) - 1]\theta$ and consider a Taylor expansion of $D(\theta, \phi, \nu_i) = D(\theta'_i, \phi, \nu_o)$ about θ .

$$D(\theta, \phi, \nu_i) = D(\theta + \Delta\theta'_i, \phi, \nu_o) = D(\theta, \phi, \nu_o) + \frac{dD}{d\theta}\Delta\theta'_i + \frac{1}{2}\frac{d^2D}{d\theta^2}\Delta\theta'^2_i \dots\dots (\text{A.3})$$

substituting the above into equ(A1) we obtain

$$D_n(\theta, \phi) = \frac{1}{N} \sum_{i=1}^N (\Delta\nu_i/\nu_o)^n \left[D(\nu_o) + \frac{dD(\nu_o)}{d\theta}\theta(\Delta\theta'_i/\theta) + \dots \right] \quad (\text{A.4})$$

if we substitute

$$(\Delta\theta'_i/\theta) = [(\nu_i/\nu_o) - 1] \quad (\text{A.5})$$

and then carry out the summation over i we obtain

$$D_n(\theta, \phi) = \frac{1}{N} \sum_{m=0}^{\infty} C_{n,m} \frac{d^m D(\theta, \phi, \nu_o)}{d\theta^m} \theta^m / m! \quad (\text{A.6})$$

and $C_{n,m}$ is defined so that

$$C_{n,m} = \sum_{i=1}^N (\Delta\nu_i/\nu_o)^n (\nu_i/\nu_o - 1)^m \quad (\text{A.7})$$

$$(\text{A.8})$$

$$C_{n,m} = \sum_{i=1}^N (\Delta\nu_i/\nu_o)^{n+m} (-1)^m$$

We can see from the above equations that the n th order SDB will depend on the derivatives of the conventional reference frequency dirty beam with respect to radial distance. The relative contributions from the different derivatives at a given radius will depend on the size of the bandspread and the reference frequency dirty beam, and the higher terms will become more important at larger radii.

For the specific case of the D_1 beam, if ν_o , the reference frequency is taken as the arithmetic mean of the observing frequencies then $C_{1,0}=0$ and so the zeroth derivative of $D(\nu_o)$ does not contribute to the spectral dirty beam. This will mean that to a first approximation the D_1 beam will depend on the first derivative of $D(\nu_o)$ with respect to radial distance. This relationship is used in section(2.4.4) to estimate the peak size of the D_1 beam. Note however that at higher values of the radial distance θ higher derivatives will also start to give significant contributions.

Appendix B

The Logarithmic Expansion

In section(2.3.3) we saw that it is possible to describe the sky-plane effects of MFS data in terms of a linear expansion. This linear expansion can be viewed as arising from expanding the frequency dependent part of the intensity distribution at frequency ν_i in powers of $(\Delta\nu/\nu_o)$, where $\Delta\nu = \nu_o - \nu_i$. We start with $I(x,y,\nu_i)$ and then expand this as a Taylor series about ν_o

$$I(x, y, \nu_i) = I(x, y, \nu_o) + \frac{dI}{d(\Delta\nu/\nu_o)}(\Delta\nu/\nu_o) + \frac{1}{2} \frac{d^2 I}{d(\Delta\nu/\nu_o)^2}(\Delta\nu/\nu_o)^2 \dots \quad (\text{B.1})$$

where the derivatives are evaluated at $(\Delta\nu/\nu_o)=0$. Using

$$\frac{d^n}{d(\Delta\nu/\nu_o)} = (-1)^n \nu_o^n \frac{d^n}{d\nu^n} \quad (\text{B.2})$$

we obtain

$$I(x, y, \nu_i) = I_o + (-1) \frac{dI}{d\nu} \nu_o (\Delta\nu/\nu_o) + \frac{1}{2} \frac{d^2 I}{d\nu^2} \nu_o^2 (\Delta\nu/\nu_o)^2 + .. \quad (\text{B.3})$$

If we now assume that the source has an effective SVPL spectra so that $I = I(x, y, \nu_o)(\nu/\nu_o)^{-\alpha'(x,y)}$ then we recover the linear expansion of section(2.3.3) because for this functional form of spectral behaviour

$$\nu_o \frac{dI}{d\nu} = -I\alpha' \quad (\text{B.4})$$

and

$$\nu_o^2 \frac{d^2 I}{d\nu^2} = I\alpha'(\alpha' + 1) \quad (\text{B.5})$$

or generally

$$\nu_o^n \frac{d^n I}{d\nu^n} = (-1)^n I\alpha'(\alpha' + 1) .. (\alpha' + n - 1) \quad (\text{B.6})$$

Substituting the above expression into equ(B.3) we obtain the linear expansion of I, which was used in section(2.3.3). At each frequency the Fourier transform of this expansion sampled by the appropriate uv coverage at each

frequency will constitute the MFS data at frequency ν_i , On combining together the sampled data at each frequency and then back inverting we obtain an expression for the MFS dirty map as a linear expansion (see equ(2.3.3)). As noted in section(2.5.1) the terms of this linear expansion decrease only slowly with order.

Because it is possible to form a Taylor series in any function of frequency $f(\nu)$, it might be wondered whether there exist any alternative expansion that converges faster, especially in the case of sources with SVPL spectral distributions. Although the linear expansion is well suited to the case of a linear spectrum with which second and higher terms would be zero, the linear expansion appears to have no special relationship with SVPL spectra. It should be possible to create an alternative expansion which utilises the *a priori* knowledge that the source spectrum is likely to be power law and that the index of this power law α' is such that $|\alpha'| \ll 1$, to ensure rapid convergence.

In the linear expansion the α' dependance within the n th order coefficient is $\alpha'(\alpha'+1)...\alpha'+n-1/n!$. It is the presence of the integers in these expressions which mean that the coefficients are large even when α' is small. The situation would be improved if the expansion for $I(x,y,\nu)$ contained a power series in $\alpha'(x,y)$. With $|\alpha'| \ll 1$ the size of the terms would then decrease rapidly with order. Such a expression can be obtained if we expand in α' instead of $(\Delta\nu/\nu_o)$. We shall see later that the same expansion can also be derived by expanding $I(x,y,\nu)$ in terms of a particular function of frequency $f(\nu)$.

Consider the substitution

$$I(x, y, \nu) = I_o(x, y)(\nu/\nu_o)^{-\alpha'(x,y)} = I_o \exp[\ln(\nu/\nu_o)^{-\alpha'(x,y)}] = I_o \exp[-\alpha'(x, y) \ln(\nu/\nu_o)] \quad (B.7)$$

At any particular x,y we can expand the above expression in terms of α' about $\alpha' = 0$. i.e.

$$I(\nu) = I + \left. \frac{dI}{d\alpha'} \right|_{\alpha'=0} \alpha' + \frac{1}{2} \left. \frac{d^2 I}{d\alpha'^2} \right|_{\alpha'=0} \alpha'^2 + \dots \quad (B.8)$$

where

$$\left. \frac{dI}{d\alpha'} \right|_{\alpha'=0} = -I(x, y, \nu_o) \ln(\nu/\nu_o) \cdot \exp[-\alpha' \cdot \ln(\nu/\nu_o)] \Big|_{\alpha'=0} \quad (B.9)$$

and generally

$$d^n I / d\alpha'^n = I(x, y, \nu_o) [-\ln(\nu/\nu_o)]^n \exp[-\alpha' \ln(\nu/\nu_o)] |_{\alpha'=0} \quad (B.10)$$

so that

$$\frac{d^n I}{d\alpha'^n} = [\ln(\nu/\nu_o)]^n I_o \quad (B.11)$$

therefore

$$I(x, y, \nu) = I_o(x, y, \nu_o) + I_o \alpha' \ln(\nu_o/\nu) + I_o \alpha'^2 / 2 [\ln(\nu_o/\nu)]^2 \dots \quad (B.12)$$

the n th order term in this series being

$$\frac{I_o \alpha'^n}{n!} [\ln(\nu/\nu_o)]^n \quad (B.13)$$

with $I_o = I(x, y, \nu_o)$

Note that this expression can also be considered as a power law expansion in terms of the quantity $\ln(\nu_o/\nu_i)$, as well as being a power law expansion in α' . The above expansion can therefore be viewed as a generalisation of the linear expansion where instead of expanding in terms of the function $(\Delta\nu/\nu_o)$ we expand in the function $\ln(\nu_o/\nu_i)$.

Continuing with our derivation we note that the visibility function at frequency ν_i , $V(u, v, \nu_i)$ will from the Zenicke Van Cittert theorem be equal to

$$V(u, v, \nu_i) = \overline{I(x, y, \nu_i)} \quad (B.14)$$

$$= \overline{I_o} + \overline{I_o \alpha'} \ln(\nu_o/\nu_i) + \overline{I_o \alpha'^2} / 2 [\ln(\nu_o/\nu_i)]^2 \dots \quad (B.15)$$

The sampled uv data at a frequency ν_i will equal the above function sampled by the uv coverage function $S(u, v, \nu_i)$

$$V_s(u, v, \nu_i) = S(u, v, \nu_i) V(u, v, \nu_i) \quad (B.16)$$

If we combine all the observed uv data sets at all the frequencies into one data set we obtain V_T

$$V_T = \sum_{i=1}^N V_s(u, v, \nu_i) \quad (B.17)$$

$$V_T = \overline{I_o} \sum_{i=1}^N S(u, v, \nu_i) + \overline{I_o \alpha'} \sum_{i=1}^N S(u, v, \nu_i) \ln(\nu_o/\nu_i) \quad (B.18)$$

$$+ \overline{I\alpha'^2/2} \cdot \sum_{i=1}^N S(u, v, \nu_i) [\ln(\nu_o/\nu_i)]^2 \dots$$

On Fourier inverting the uv data we obtain a dirty map that can be expressed as

$$\overline{V_T} = (I_o * D'_o) + (I_o \alpha' * D'_1) + (I_o \alpha'^2/2 * D'_2) \dots \quad (\text{B.19})$$

or more generally

$$= \sum_{n=1}^{\infty} (I_o \alpha'^n * D'_n) / n! \quad (\text{B.20})$$

where the D'_n , the 'logarithmic dirty beams' of order n, are defined by

$$D'_n = \overline{\sum S(u, v, \nu_i) [\ln(\nu_o/\nu_i)]^n} \quad (\text{B.21})$$

The form of the logarithmic expansion is very similar to that of the linear expansion, each term of the expansion consists of a coefficient image convolved with its corresponding spectral dirty beam. The coefficient image for each order will however be a new function of I and α' .

The logarithmic spectral dirty beams will also be very similar to their linear counterparts, the only difference being that instead of uv points at frequency ν_i being weighted by $(\Delta\nu_i/\nu_o)$ they are instead weighted by $\ln(\nu_o/\nu_i)$. An important property of these new SDB's is that as we decrease the size of the total bandspread we can expect that they will become equivalent to the linear SDB's. To see this note that if $x = (\Delta\nu_i/\nu_o) = (\nu_o - \nu_i/\nu_o)$

$$\ln(\nu_o/\nu_i) = -\ln(\nu_i/\nu_o) = -\ln(1 - x) \quad (\text{B.22})$$

now

$$-\ln(1 - x) = x + x^2/2 + x^3/3 + \dots x^n/n \quad (\text{B.23})$$

and so

$$\ln(\nu_o/\nu_i) = (\Delta\nu_i/\nu_o) + (1/2) (\Delta\nu_i/\nu_o)^2 + \dots \quad (\text{B.24})$$

which means that as $\Delta\nu_i \ll \nu_o$ the higher terms become negligible relative to the first and the logarithmic weighting becomes approximately equivalent to the linear. If the fractional bandspread is small then the above condition

is applicable to all frequencies within the bandspread and so the logarithmic beam of order n becomes equivalent to the linear beam of the same order.

We can further see how the linear and logarithmic expansions are related by multiplying out the coefficients of the n th order term of the linear expansion given by

$$[I\alpha'(\alpha' + 1)(\alpha' + 2)\dots(\alpha' + n - 1)/n!] * D_n \quad (\text{B.25})$$

and then collecting together the terms from the different orders which contain α' raised to the same power, as in the different orders of the logarithmic expansion. If we consider the sum of terms containing $I\alpha'$ we obtain

$$I\alpha' * [D_1 + D_2/2 + D_3/3\dots D_n/n] \quad (\text{B.26})$$

This can be compared with the expression containing $I\alpha'$ in the logarithmic expansion which is, $I\alpha' * D'_1$. It follows that

$$D'_1 = (D_1 + D_2/2 + D_3/3\dots\dots) \quad (\text{B.27})$$

or generally

$$D'_1 = \sum_{n=1}^{\infty} D_n/n \quad (\text{B.28})$$

Appendix C

Proof that the MFS Dirty Image is In Range of D_o

An important consideration in deconvolving a dirty image which contains spectral errors is whether this dirty image is ‘in range’ of the ‘dirty beam’, D_o . Schwarz(1978) and Tan(1986) have shown that unless this condition holds the CLEAN algorithm will fail. The ‘in range’ condition is most easily understood by considering the deconvolution problem in terms of matrices and vectors (Schwarz 1978).

A P by Q image can be completely described as a vector of N elements (where $N=PQ$) with elements in the image being mapped into the vector by some scheme of ‘lexicographic’ ordering (Andrews and Hunt 1977). The vector representing the true image within the N element area or ‘window’ is symbolized by t_o and this will be connected with the N element dirty image vector d in conventional imaging by a matrix equation.

$$d = B_o t_o \quad (C.1)$$

Where each row of the N x N matrix B_o (the dirty beam matrix) consists of a vector which equals the dirty beam with a position shift which depends on the row number. Such a matrix is termed ‘Toeplitz’ (Andrews and Hunt 1977) in form.

The dirty image d is ‘in range’ of D_o if there exists an image vector such that when multiplied by the dirty beam matrix the result given is d . In conventional imaging the dirty map should always be ‘in range’ because given that the dirty map is defined by equ(C.1) there is at least one image vector that gives rise to the dirty map, viz the true image vector t_o .

The CLEAN algorithm assumes that the dirty image is created by the simple convolution as described in equ(C.1) and attempts to estimate the

underlying true image vector. If we attempt to CLEAN a dirty image created from inverting MFS data the dirty image is no longer given by an expression of the form of equ(C.1) and so there is no automatic guarantee that the dirty map is in range. The dirty image is in fact described by a two order approximation which can be written in the matrix form

$$d = B_o t_o + B_1 s_o \quad (C.2)$$

where B_o and t_o are as defined above, B_1 is the first order spectral dirty beam matrix with each row consisting of the D_1 beam and s_o is the $I\alpha'$ image vector. The first term of the above equation presents no problem, it is always in range, it is the second term which we must consider to see if it gives a vector which is out of range of B_o .

As a first stage in seeing whether d is in range of B_o we must consider decomposing d into the eigenvectors of the matrices B_o and B_1 . The matrices B_o and B_1 are real and symmetric, they are therefore 'orthonormal'. This means that the N eigenvectors of the matrices form a complete orthogonal set so that any vector can be represented as a linear combination of the eigenvectors and these eigenvectors are all mutually orthogonal (i.e. Bellman 1960). Any dirty image can therefore be split up into a linear combination of the eigenvectors of the matrices B_o or B_1 as desired. If the MFS dirty image d , contains a zero eigenvector (i.e. a eigenvector with zero eigenvalue) of B_o then it follows that the dirty image is out of range of the D_o beam, *because no vector exists such that when multiplied by B_o the result contains a zero eigenvector of B_o .*

To show that the MFS dirty beam is in range of B_o it is sufficient to show that all the zero eigenvectors of B_o are also zero eigenvectors of B_1 . If this is the case there will be no vector s which can generate zero eigenvectors of B_o in the contribution $B_1 s$ to d . Without these zero eigenvectors of B_o the MFS dirty image will be in range of the dirty beam D_o .

The case in which the size N of the image goes to infinity can be trivially dealt with. The convolution problem can then be posed in terms of infinite

vectors and matrices.

$$d' = B'_0 t'_0 + B'_1 s'_0 \quad (C.3)$$

where the primed quantities are infinite vectors and matrices, with each row of the beam matrices being equal to the sampled D_0 or D_1 beams. Note that the finite $N \times N$ beam matrices in equ(C2) are both submatrices of the above infinite matrices and the vectors of size N are truncated versions of the above infinite vectors.

The eigenvectors of the infinite matrices will correspond to sine and cosine ripples in the sky plane. Only ripples whose spatial frequencies are sampled by the uv coverage will be non-zero eigenvectors. Since both the uv coverages which give rise to D_0 and D_1 beams have the same non-zero uv points, it follows that any non-zero eigenvector of one beam matrix is a non-zero eigenvector of the other, likewise any zero eigenvector of one is a zero eigenvector of the other. It follows that in the infinite case the MFS dirty map is always in range.

We now consider the case of a non-infinite image. We can divide the problem into two depending on the 'rank' of the finite matrices (Bellman 1960). First we consider the rank of the parent infinite matrices. We note that for each sampled uv point there will be one non-zero eigenvector of B'_0 and B'_1 . If there are R' sampled uv points then there will be R' non-zero eigenvectors of each matrix, this number is equal to the 'rank' of the matrix. For the $N \times N$ submatrices B_0 and B_1 it is a theorem that the rank will be equal to the rank of their parent matrices or their order (N) whichever is smallest (Bellman 1960). There are now two cases, the rank R can be equal to N or it can be smaller than N . In the former case there will be no zero eigenvectors in either of the two beam matrices because the number of zero eigenvectors is equal to the order minus the rank. It follows that the sets of zero eigenvectors for the two beams are identical and so the dirty map is in range.

The other case, of undersampling, requires further analysis. In this case there will exist zero eigenvectors of both B_0 and B_1 . Consider a particular zero eigenvector (ZEV), say of B_0 denoted by z_0 , then by definition $B_0 z_0 = 0$. Consider 'padding' the vector z_0 of size N with zeros out to infinity to form

the infinite vector z'_o . It is possible to show that because z_o is a ZEV of B_o then z'_o is a ZEV of B'_o .

Consider multiplying out $B'_o z'_o$. Because $B_o z_o = 0$ the resulting vector must be zero over a field of size N , what however about outside of this range? It is a theorem that any row of a matrix such as B'_o can be considered as a linear combination of any R' other rows, where R' is the rank of B'_o (Bellman 1960). Because $N > R$ it follows that we can consider any row as a linear combination of any R of the N rows involved in the submatrix B_o . Because $B'_o z'_o = 0$ in the N element window, z'_o will be orthogonal to each of these N rows. It will therefore be orthogonal to any linear combination of the rows and hence to any row of B'_o . It follows that the result of the matrix multiplication $B'_o z'_o$ is zero and therefore that z'_o is a ZEV of B'_o .

To be a ZEV of B'_o means that the FT of z'_o is zero at all sampled uv points, it therefore follows that it is also a ZEV of B'_1 . Because $B'_1 z' = 0$ it follows that over the N element field corresponding to the submatrix B_1 we have $B_1 z_o = 0$ and so z_o is a ZEV of B_1 as well as B_o .

It is possible to show by an analogous argument to the above that a ZEV of B_1 is a ZEV of B_o . It follows that any ZEV of one matrix is also a ZEV of the other matrix and so the two matrices share the same 'null space'. The required condition to ensure that the MFS dirty image is in range of the dirty beam therefore holds in all cases.

Note that although we have concentrated on showing that the first order spectral errors do not cause the MFS dirty image to be out of range a similar argument holds for higher order error contributions as well. This is because the argument used depends only on the the FT of the different order dirty beams being non-zero on the same uv points, a situation which holds for all orders.

Appendix D

The Gridded Cornwell Algorithm

In the gridded version of the Cornwell algorithm we use the gridding mechanism as a means to introduce *a priori* knowledge about the source size, and so relax the ‘double sampling’ requirement of the exact algorithm. In this appendix we seek to describe in more detail the operation of the gridded algorithm.

As a preliminary we consider the gridding process generally as it applies to normal mapping. The uv sampling function we denote by \overline{D}_o and the visibility by \overline{I} , the ungridded measured data is thus $\overline{D}_o \times \overline{I}$. We wish to interpolate this onto a uniform grid denoted by the Shah function $III(u,v)$ (Bracewell 1965); we may do this by first convolving the uv data with some gridding function $g(u,v)$ and then sample the result with III , i.e. we form

$$III \times [g * \overline{D}_o \overline{I}_o] \quad (D.1)$$

This expression is a convenient form from which we can explore the map plane effects of gridding (Sramek and Schwab 1985). Computationally full convolution in the uv plane is unnecessary, what is normally done is to take each grid point in turn and then for all uv points that are within a position vector \underline{r} from the grid point such that $g(\underline{r}) \neq 0$, we accumulate a contribution $g(\underline{r})\overline{I}(\underline{r})$ to the value at the grid point. If $g(\underline{r})$ is non-zero only within a region S and there exists two uv data points within this region then the resulting value at the grid point will be

$$V_g = g(\underline{r}_1)\overline{I}(\underline{r}_1) + g(\underline{r}_2)\overline{I}(\underline{r}_2) \quad (D.2)$$

which we will abbreviate as

$$V_g = g_1\overline{I}_1 + g_2\overline{I}_2 \quad (D.3)$$

a result which is easily generalized if we have more than two points within S.

We now consider the modifications that gridding introduces into the Cornwell algorithm. As described in the main text the gridded algorithm's effective dirty map d is defined by analogy to the exact algorithm as

$$d = (D_2 \mathfrak{F} d_o) - (D_1 \mathfrak{F} d_1) \quad (\text{D.4})$$

and the associated 'effective dirty beam' b_{eff} by

$$b_{eff} = (D_2 \mathfrak{F} D_o) - (D_1 \mathfrak{F} D_1) \quad (\text{D.5})$$

where D_1 and D_2 have their normal definitions (see Chapter 2), and where d_o and d_1 are defined in section(4.4.2). The operation denoted by \mathfrak{F} is 'gridded convolution'. This operation, $A \mathfrak{F} B$, consists of gridding the uv data corresponding to the images A and B onto a regular grid in the uv plane, these two gridded functions are multiplied together grid point by grid point and then a back FT is executed to produce the final result. By using the gridding mechanism *a priori* information about the source support size can be introduced into the Cornwell algorithm. To contribute toward the final image two uv points now need only be within the same grid cell, they no longer need to be exactly coincident.

We proceed by considering the FT's of d and b_{eff} . Using the definition of gridded convolution and using the symbol $\overline{G[C]}$ to denote the 'gridded' uv data corresponding to the image C, we have

$$\overline{d} = (\overline{G[D_2]} \cdot \overline{G[d_o]}) - (\overline{G[D_1]} \cdot \overline{G[d_1]}) \quad (\text{D.6})$$

and

$$\overline{b_{eff}} = (\overline{G[D_2]} \cdot \overline{G[D_o]}) - (\overline{G[D_1]} \cdot \overline{G[D_2]}) \quad (\text{D.7})$$

Consider now the values of \overline{d} and $\overline{b_{eff}}$ at a given grid point when there are 0, 1 or 2 uv points respectively within the region S about that grid point

If there are no points in S then trivially the value of \overline{d} at the grid point is zero, likewise $\overline{b_{eff}} = 0$. If instead there is a single uv point observed at a frequency ν_1 then from the logarithmic expansion the observed visibility will

then be approximately equal to $\bar{I} + f\bar{I}\alpha'$ where $f = f(\nu_1) = \ln(\nu_o/\nu_1)$, \bar{I} is the Fourier transform of the true intensity distribution at the position of the sampled uv point and $\bar{I}\alpha'$ the Fourier transform of $I\alpha'$. The resulting values of the quantities which comprise equ(D.6) and (D.7) are

$$\overline{G[D_o]} = g \quad (D.8)$$

$$\overline{G[D_1]} = gf$$

$$\overline{G[d_o]} = g\bar{I} + gf\bar{I}\alpha'$$

$$\overline{G[d_1]} = gf\bar{I} + gf^2\bar{I}\alpha'$$

because for a single sample within the grid cell $\overline{G[C]} = g\bar{C}$. From equ(D.6) the value of d at the grid point will be

$$\bar{d} = gf^2(g\bar{I} + gf\bar{I}\alpha') - gf(gf\bar{I} + gf^2\bar{I}\alpha') = 0 \quad (D.9)$$

likewise we can show that $\overline{b_{eff}} = 0$. Hence grid points with either one point or no point in their region S contribute nothing to the effective dirty map or dirty beam.

Consider now the case of two points within S . The frequencies of observation of these two points will be ν_1 and ν_2 and the values of the visibility at these two points will be $\bar{I}_1 + f_1\bar{I}\alpha'_1$ and $\bar{I}_2 + f_2\bar{I}\alpha'_2$ respectively. The corresponding gridded values will be

$$\overline{G[d_o]} = g_1\bar{I}_1 + g_1f_1\bar{I}\alpha'_1 + g_2\bar{I}_2 + g_2f_2\bar{I}\alpha'_2 \quad (D.10)$$

$$\overline{G[d_1]} = g_1f_1\bar{I}_1 + g_1f_1^2\bar{I}\alpha'_1 + g_2f_2\bar{I}_2 + g_2f_2^2\bar{I}\alpha'_2 \quad (D.11)$$

$$\overline{G[D_1]} = g_1f_1 + g_2f_2 \quad (D.12)$$

$$\overline{G[(D_2)]} = g_1f_1^2 + g_2f_2^2 \quad (D.13)$$

From equ(D.6) \bar{d} at the grid point can be shown after some manipulation to be

$$\begin{aligned} \bar{d} &= g_1g_2(f_1 - f_2)^2\bar{I} \\ &+ g_1g_2(f_1 - f_2)(f_1\Delta\bar{I}_2 - f_2\Delta\bar{I}_1) \end{aligned} \quad (D.14)$$

$$+g_1g_2(f_2 - f_1)f_1f_2(\Delta\overline{I\alpha'_1} - \Delta\overline{I\alpha'_2})$$

where $\overline{I_1} = \overline{I} + \Delta\overline{I_1}$, $\overline{I_2} = \overline{I} + \Delta\overline{I_2}$, where \overline{I} is the value of the true visibility at the grid point. Likewise $\overline{I\alpha'_1} = \overline{I\alpha'} + \Delta\overline{I\alpha'_1}$, $\overline{I\alpha'_2} = \overline{I\alpha'} + \Delta\overline{I\alpha'_2}$ with $\overline{I\alpha'}$ being the FT of $I\alpha'$ at the grid point. We obtain an analogous result for $\overline{b_{eff}}$, the effective dirty beam or point spread function by choosing $\overline{I_1}, \overline{I_2} = 1$ and likewise $\overline{I\alpha'_1}, \overline{I\alpha'_2} = 0$,

$$\overline{b_{eff}} = g_1g_2(f_1 - f_2)^2 \quad (D.15)$$

For the case of two points within S, \overline{d} and $\overline{b_{eff}}$ will be non-zero at the grid point *provided that the observing frequencies of the two points are different* so that $f_1 \neq f_2$.

If the second and third terms of equ(D.14) are small compared to the first then

$$\overline{d} = \overline{b_{eff}}.\overline{I} \quad (D.16)$$

and in the sky plane the problem is one of the true I image convolved with an effective dirty beam, with no spectral effects.

Note that if the two uv points within S approach each other until they are coincident then $\overline{I_1} = \overline{I_2}$ and $\overline{I\alpha'_1} = \overline{I\alpha'_2}$ and all terms in equ(D.14) except the first disappear. In the above limit we eventually obtain the result of equ(4.4) in the main text which applies to exactly overlapping uv points.

When the two uv points are not coincident the second and third terms of the above expression are not zero. In such circumstances the effective dirty map \overline{d} will not strictly equal an image convolved with the effective dirty beam, $\overline{b_{eff}}$. That is to say the dirty map will not strictly be ‘in range’ of the dirty beam. As pointed out in Appendix C in this situation the application of algorithms which assume that the dirty map is in range may generate serious reconstruction errors. Only if the second and third order terms are below the noise can this problem be avoided.

The third term will contribute residual spectral effects due to the quantities $\Delta\overline{I\alpha'_1}$ and $\Delta\overline{I\alpha'_2}$. In contrast to the exact algorithm this gridded algorithm does not therefore eliminate spectral effects entirely, it does however reduce the relative contribution from the $I\alpha'$ terms greatly. The third term will

depend on the difference between $I\alpha'$ at the two observing points, which when S has a smaller diameter than $1/\theta$ (where θ is the typical source size) will be necessarily small.

The second term in the expression also depends on the difference between two similar quantities i.e. \bar{I}_1 and \bar{I}_2 within S . This second term can be thought of as an additional form of gridding error introduced in addition to the effect of multiplication by the quantities g_1 and g_2 at the grid points.

The second term contribution will be larger than the third term, both because $\bar{\Delta I}$ will tend to be larger than $\bar{\Delta I\alpha'}$ and because the third term contains an extra weighting factor f . The errors introduced by this second term are hard to analyse in the general case but by making some assumptions a very rough estimate can be given. First we will assume that the gridding functions g_1 and g_2 are 'box-car' so that the region S is a square centred on the grid points with the dimensions of a grid cell. Within S the gridding function is taken to have value unity but outside S the value is zero. On examining the form of the second order term we will see that the size and sign of the contribution will depend on f_1 and f_2 and hence on the frequencies of observation. We will assume that for any double-sampled cell the two frequencies involved are randomly selected from those involved in the MFS observation. Given this assumption the sign of the contribution from each cell will not be correlated with its neighbours and so the resulting sky plane effect will be a statistical sum of the contributions from the cells. In the centre of the image the contribution will equal

$$\frac{[\sum_i^N ((f_{1,i} - f_{2,i})(f_{1,i}\Delta\bar{I}_2 - f_{2,i}\Delta\bar{I}_1))^2]^{1/2}}{\sum_i^N (f_{1,i} - f_{2,i})^2} \quad (D.17)$$

where the subscript $(1,i)$ refers to the first uv point within the i th cell and $(2,i)$ the second uv point within the i th double sampled cell. Given the normalisation factor on the bottom of the above equation the effect of the f factors are approximately cancelled out. The level of the contribution will then be roughly

$$\frac{\Delta\bar{I}_{rms}}{N^{1/2}} \quad (D.18)$$

where the numerator is the rms value of the variation of \bar{I} across a grid cell and N is the number of double sampled cells.

Appendix E

The Convergence of Double Deconvolution

In this appendix we will consider the progress of the Double Deconvolution (DD) algorithm in fitting MFS data and show that given certain conditions we can expect that the resulting solution that emerges from DD will fit the observed data. We will follow the approach and nomenclature of Schwarz(1978) and define a norm Q_T such that

$$Q_T = [B_o(t - t_o) + B_1(s - s_o)]^T B_o [B_o(t - t_o) + B_1(s - s_o)] \quad (\text{E.1})$$

Where t and s are the vectors representing the current estimates of the I and $I\alpha'$ images, s_o and t_o are the true I and $I\alpha'$ images and B_o and B_1 are Toeplitz matrices representing the D_o and D_1 spectral dirty beams, note the superscript T denotes the transpose operation. The square brackets contain the residual image after removing the effects of the present estimate of I and $I\alpha'$ from the dirty map. The norm measures the mean square of this residual map. It can be shown that the norm Q_T also equals the square modulus error between the model and observed data at the sampled uv points.

$$Q_T = \sum \overline{D_o} |(\overline{t} - \overline{t_o} + (\overline{s} - \overline{s_o})\overline{D_1})|^2 \quad (\text{E.2})$$

where the overlined quantities refer to the FT's of the sky plane quantities and where the summation is over sampled points in the uv plane.

The change of Q_T as both halves of the DD algorithm are applied will now be considered. First we consider the behaviour of the norm in the I deconvolving half of DD, i.e. in the half where t is updated. It follows from the above definition of Q_T that

$$Q_T = (B_o\Delta t + B_1\Delta s)^T B_o (B_o\Delta t + B_1\Delta s) \quad (\text{E.3})$$

where

$$\Delta t = t - t_o \quad (E.4)$$

$$\Delta s = s - s_o$$

Consider multiplying out equ(E.3) and simplifying it using the facts that B_o and B_1 are symmetric (so that $(B_1^T = B_1$ and $B_0^T = B_0))$ and that $B_n B_m = B_{n+m}$ (Given that matrix multiplication corresponds to a convolution of the underlying dirty beams this identity follows from the identity $D_n * D_m = D_{n+m}$). We obtain

$$Q_T = \Delta t^T B_o \Delta t + 2\Delta t^T B_1 \Delta s + \Delta s^T B_2 \Delta s \quad (E.5)$$

Following Schwarz we consider how the norm changes if the estimate of t is updated, from t' to $t' + \delta t$, so that Δt changes from $\Delta t'$ to $\Delta t' + \delta t$. The resulting change in Q_T , δQ_T is

$$\delta Q_T = \delta t^T B_o \delta t + 2\delta t^T (B_o \Delta t + B_1 \Delta s) \quad (E.6)$$

The residual image $r_o = -B_o \Delta t - B_1 \Delta s$ and so

$$\delta Q_T = \delta t^T B_o \delta t - 2\delta t^T r_o \quad (E.7)$$

If we consider selecting some point in the residual map by some criteria (such as being the largest residual) and updating our estimate t by some loop gain g times the value at the selected point r_m we obtain

$$\delta t = (0, 0, 0, \dots, g r_m, \dots, 0, 0) \quad (E.8)$$

so that

$$\delta Q_T = g^2 r_m^2 B_o(i, i) - 2g r_m^2 \quad (E.9)$$

where $B(i, i)$ is the value of B_o on a diagonal element, i.e. unity. For convergence we require that $\delta Q < 0$ which will occur whatever r_m provided that $(g^2 - 2g) < 0$ i.e. when $0 < g < 2$. Since the uv weights corresponding to the beam D_o are positive this will ensure that not only will Q_T continuously decrease but that the norm will have a lower bound of zero (see equ(E.2)).

We now consider the $I\alpha'$ deconvolving half of the DD algorithm. On starting the $I\alpha'$ cycle the residuals are convolved with D_1 and the new residuals become.

$$r_1 = B_1 r_o = -B_1 \Delta t - B_2 \Delta s \quad (\text{E.10})$$

so that

$$\delta Q_T = \delta s^T B_2 \delta s - 2 \delta s^T r_1 \quad (\text{E.11})$$

The change in s , δs will be g times the largest point in the normalised $I\alpha'$ half-cycle residuals ($r_1/D_2(0,0)$)

$$\delta s = (0, 0, \dots, g r_{1,m}/D_2(0,0), \dots, 0, 0) \quad (\text{E.12})$$

so

$$\delta Q_T = g^2 r_{1,m}^2 B_2(i,i)/D_2^2(0,0) - 2 g r_{1,m}^2 / D_2(0,0) \quad (\text{E.13})$$

Since the diagonal elements of B_2 equal $D_2(0,0)$ it follows that the condition for $\delta Q < 0$ are the same as for the I half cycle i.e. $0 < g < 2$.

If the above condition holds Q_T must decrease in both halves of a DD cycle. Given that uv weights that correspond to the beam D_o are positive (i.e. the matrix B_o is at least positive semi-definite) this means that the norm Q_T is bounded below by zero (see equ(E.2)). The above two conditions therefore mean that the fit to the data must continually improve in both half-cycles.

Consider if the I and $I\alpha'$ distributions are contained within a 'window' of N elements and assume that in the two half cycles of DD CLEANing is restricted to within such an N element window. Even though the windows are not infinite it was shown in Appendix C that the residual image is always in range in both halves of the cycle and therefore there will be no runaway increase in the zero eigenvectors on applying DD. Provided B_o is at least positive semi-definite and $0 < g < 2$ then the above proof will apply and Q_T will continuously decrease in both half cycles. Since the windows are larger than the source there will be at least one solution such that $Q_T = 0$, viz $t=t_o$ and $s=s_o$ so therefore Q_T can become zero and the data can be fitted exactly. In equ(E.2) it is apparent that if Q_T goes to zero then at the sampled uv points (i.e. where $\overline{D_o}$ is non-zero and positive) the contents of the square

bracket must be zero. This means that the model and real uv data sets are the same at the sampled uv points i.e. the observed data is indeed fitted by the model. Given the above fitting of the uv data the residual images r_0 and r_1 must therefore be zero when $Q_T=0$.

It appears that provided (a) the residual map in each cycle is in range of the dirty beam being CLEANed, (b) the matrix B_0 is at least positive semi-definite and (c) the loop gain g is such that $0 < g < 2$ in both half-cycles, then DD can converge to fit the uv data exactly.

Appendix F

Separating the Effects of the I and $I\alpha'$ Distributions with Double Deconvolution

In Appendix E it was shown that DD will always converge to fit the data provided certain common conditions hold. Given this convergence occurs the remaining question is whether the algorithm successfully separates between the effects of the I and $I\alpha'$ distributions. The previous appendix showed that the total effect at any sampled uv point due to the model I and $I\alpha'$ distributions equals the observed data, the remaining question is how successful is it in dividing the data between the two models at the sampled uv points?

If Q_T tends to zero as shown in Appendix E then the observed uv data is fitted by the model and the residual map described by

$$r_o = -B_o\Delta t - B_1\Delta s \quad (F.1)$$

will tend to zero. The r_1 residual map which depends on a reweighting of the residual uv data will also tend to zero.

$$r_1 = -B_1\Delta t - B_2\Delta s = 0 \quad (F.2)$$

The above equations can be presented as a single matrix equation

$$\begin{pmatrix} B_o & B_1 \\ B_1 & B_2 \end{pmatrix} \begin{pmatrix} \Delta t \\ \Delta s \end{pmatrix} = 0 \quad (F.3)$$

If the matrix on the left hand side, C, is non-singular then there will be a unique solution to the above equation viz $\Delta t=0$ and $\Delta s = 0$. This means that $t = t_o$ and $s = s_o$, so that DD produces a solution which is both unique and correct.

A necessary but not sufficient condition that C is non-singular will be that the number of different uv points that give rise to the B matrices should be at least $2N$, where N is the number of pixels in the image. If the $2N$ points are distributed almost randomly about the uv plane, as in realistic observations, then in nearly all cases the resulting C matrix will be non-singular. The above limit of creating a non-singular C corresponds to the ‘over-sampled’ limit described in section (4.1.4) in which there are sufficient observations to uniquely define a solution for I and $I\alpha'$, i.e. to successfully separate out the effects of the I and $I\alpha'$ distributions.

We have seen above that in the ‘oversampled’ limit the successful separation between I and $I\alpha'$ effects is guaranteed. What if we are not in this regime, under what conditions if any can we be sure that the DD algorithm finds the correct solution?

Using the nomenclature of Appendix E we will consider the norms

$$Q_o = (t - t_o)^T B_o (t - t_o) = \Delta t^T B_o \Delta t \quad (F.4)$$

$$Q_1 = (s - s_o)^T B_2 (s - s_o) = \Delta s^T B_2 \Delta s \quad (F.5)$$

These two norms measure respectively the fit of the estimate of I to the true I image and the fit of the model $I\alpha'$ to the true $I\alpha'$ distribution. Given that the aim of DD is to find the unknowns t_o and s_o the values of these norms will be inaccessible to us when we apply DD to MFS data. This is in contrast to Q_T which is quantifiable at all stages. The norms Q_o and Q_1 will however still be useful because by considering under which conditions they continuously decrease we can define the situations in which the separation between I and $I\alpha'$ is assured.

Consider the behaviour of the norms in the I deconvolution half cycle, since s does not change Q_1 will not change, we therefore consider only Q_o . The change on updating the model t by δt , δQ_o will be

$$\delta Q_o = \delta t^T B_o \delta t + 2\delta t^T B_o \Delta t \quad (F.6)$$

The residual in the I half of the cycle equals

$$r_o = -B_o \Delta t - B_1 \Delta s \quad (F.7)$$

Substituting for $B_o\Delta t$ in the above equation we obtain

$$\delta Q_o = 2\delta t^T(-r_o - B_1\Delta s) + \delta t^T B_o\delta t \quad (\text{F.8})$$

The vector δt will be $(0, 0, \dots, gr_m, \dots, 0)$ where r_m is the peak residual, so

$$\delta Q_o = g^2 r_m^2 - 2gr_m^2 - 2gr_m(B_1\Delta s)_m \quad (\text{F.9})$$

where $(B_1\Delta s)_m$ is the value of $B_1\Delta s$ at the position at which the peak residual was found. If we choose $0 < g < 2$, then δQ_o is negative provided that either

$$r_m(B_1\Delta s)_m > 0 \quad (\text{F.10})$$

or if the above condition does not hold then if

$$|(g^2 - 2g)r_m^2| > |2gr_m(B_1\Delta s)_m| \quad (\text{F.11})$$

so if

$$(2g - g^2)|r_m|^2 > 2g|r_m|| (B_1\Delta s)_m| \quad (\text{F.12})$$

then $\delta Q_o < 0$ no matter whether $r_m(B_1\Delta s)_m$ is greater or less than zero. The convergence condition then becomes that the magnitude of the peak residual should be such that

$$|r_m| < \frac{2g}{2g - g^2} |B_1\Delta s|_m \quad (\text{F.13})$$

The above condition will automatically hold if the residual is such that

$$|r_m| < \frac{2g}{2g - g^2} |B_1\Delta s|_{max} \quad (\text{F.14})$$

where $|B_1\Delta s|_{max}$ is the peak contribution due to the uncleaned $I\alpha'$ distribution. If g is small then $(2g/(2g - g^2)) \approx 1$ and so we obtain the plausible condition that in order for a successful separation to occur CLEANing should stop when the peak residual is larger than the peak effect due to the residual unwanted distribution.

Carrying out an analogous analysis for the $I\alpha'$ CLEANing half cycle of DD then we obtain a similar condition for $\delta Q_2 < 0$ i.e. that

$$|r_{1,m}| < \frac{2g}{2g - g^2} |B_1\Delta t|_{max} \quad (\text{F.15})$$

where $|B_1 \Delta t|_{max}$ is the peak $I\alpha'$ domain effect due to uncleaned I. The condition is again to stop before the effects due to the unwanted distribution become dominant in the residuals.

The above conditions provide the basis for implementations of DD which are 'residual limit' controlled, i.e. in which we switch from CLEANing in one domain to CLEANing in the other when the residuals reach a given value. Provided the critical limit is not breached in each half cycle both Q_0 and Q_2 will continuously decrease and can eventually be made zero. The contributions from the I and $I\alpha'$ models at each uv point would then be correct, and the effects of the two distributions would have been correctly separated.

The problem with the above method of control however is that the important criteria; the level of the errors due to the unwanted distributions, are not directly accessible. There are ways however of estimating upper limits to these contributions, for instance by considering the far sidelobes in the residual map (see section(5.3.4)). One other problem is that there is no guarantee that when a critical value of peak residual is reached and we switch to CLEANing the other distribution that the peak residual in this new map is above the new critical limit. If this situation arises there will be no guarantee that we can successfully separate out the two distributions.

Bibliography

- Achterberg, A., 1986, *Proc 'Astrophysical Jets and Their Engines.'*
NATO ASI conf. Ed Kundt, W. Publ Reidel, D. p. 223.
- Alfvén, H. and Herlofson, N., 1950, *Phys. Rev.*, **78**, 616.
- Aller, H.D., Hughes, P.A. and Aller, M.F., 1987, *Proc 'Superluminal Radio Sources', Big Bear, California.* Eds. Zensus, J.A. and Pearson, T.J. p 211.
- Andrejanov, V.V., Gurvits, L.I., Kardeshshev, N.S., Pogrebenko, S.V.,
Rudakov, V.A., Sagdeev, R.Z. and Tsarevsky, G.S., 1984,
'*Quasat-A VLBI Observatory In Space*', *Proc. Workshop. Gross Enzersdorf. Austria.*, p 161.
- Andrews, H.C. and Hunt, B.R., 1977, '*Digital Image Restoration*',
Publ, Prentice-Hall Inc.
- Armstrong, J.W. and Sramek, R.A., 1982, *Radio Sci*, **17**, 1579.
- Baars, J.W.M., Genzel, R., Pauliny-Toth, I.I.K. and Witzel, A., 1977,
Astron. Astrophys., **61**, 99.
- Baars, J.W.M., Van der Brugge, J.L., Casse, J.L., Hamaker, J.P.,
Sondaar, L.H., Visser, J.J., Wellington, K.J., 1973, *Proc. I.E.E.E*,
61, 1258.
- Baldwin, J.E., 1971, *IAU Symp.*, **46**, 22.
- Bell, A.R., 1978, *Mon. Not. R. astr. Soc.*, **187**, 147.
- Bellman, R., 1960, '*Introduction to Matrix Analysis*', McGraw-Hill Book Co.,
New York, New York.
- Bignell, R.C. and Perley, R.A., 1985, *Chapter 4, 'Synthesis Imaging'*,
Publ N.R.A.O. Eds. Perley, R.A., Schwab, F.R. and Bridle, A.H., 49.
- Biretta, J.A, Reid, M.J., Junor, W., Spencer, R.E. and Muxlow, T.W.B.,
1988, '*The Impact Of VLBI On Astrophysics And Geophysics*'
I.A.U. Symp. **129**.
- Blandford, R.D. and Königl, A., 1979, *Astrophysical Lett*, **20**, 15.
- Blandford, R.D. and Rees, M.J., 1974, *Mon. Not. R. astr. Soc*, **169**, 395.

- Bracewell, R.N., 1958, *Proc. I.R.E.*, **46**, 97.
- Bracewell, R.N., 1965, *'The Fourier Transform and its Applications'*.
Publ McGraw Hill. New York.
- Bracewell, R.N., 1984, *'The Early Years Of Radio Astronomy'*. Ed. W.T.
Woodruff. Pub C.U.P., p. 167.
- Braun, R., Gull, S.F. and Perley, R.A., 1987, *Nature*, **327**, 395.
- Bridle, A.H., 1982, *I.A.U. Symp.*, **121**.
- Bridle, A.H. and Perley, R.A., 1984, *Ann. Rev. Astrophys. Astron.*,
22, 319.
- Browne, I.W.A., 1987, *Proc 'Superluminal Radio Sources', Big Bear,*
California. Eds. Zensus, J.A. and Pearson, T.J.
- Clark, B.G., 1980, *Astron. Astrophys.*, **89**, 377.
- Cohen, M.H., 1973, *Proc I.E.E.E.* Vol **61**, No **9**, 1192.
- Cole, T.W., 1979, *J.Opt.Soc.Am*, **69**, 554.
- Coleman, C.S. and Bicknell, G.V., 1988, *Mon. Not. R. astr. Soc*,
230, 497.
- Conway, J.E., 1985, *Dipl. Ad. Stud. Sci. University Of Manchester*
- Conway, R.G., Kellerman, K.I. and Long, R.J., 1963, *Mon. Not. R.*
Soc., **125**, 261.
- Conway, R.G. and Stannard, D., 1975, *Nature*, **255**, 310.
- Cornwell, T.J., 1982, *PhD Thesis, University of Manchester.*
- Cornwell, T.J., 1983, *Astr. Astrophys.*, **121**, 281.
- Cornwell, T.J., 1984, *VLBA Memo* **324**
- Cornwell, T.J., 1985, *Chapter 9, 'Synthesis Imaging'*,
Publ N.R.A.O. Eds. Perley, R.A., Schwab, F.R. and Bridle, A.H., p 137.
- Cornwell, T.J., Evans, K.F., 1983, *Astr. Astrophys.*, **143**, 77.
- Cornwell, T.J. and Wilkinson, P.N., 1981, *Mon. Not. R. astr. Soc.*
196, 10.
- Cotton, W.D., 1985, *Chapter 8, 'Synthesis Imaging'*,
Publ N.R.A.O. Eds. Perley, R.A., Schwab, F.R. and Bridle, A.H., p 123.
- D'Addario, L.R., 1985, *Chapter 3, 'Synthesis Imaging'*,

- Publ N.R.A.O. Eds. Perley, R.A., Schwab, F.R. and Bridle, A.H.*, p 31.
- Davies, J.G., Anderson, B. and Morison, I, 1980, *Nature*, **288**, 64.
- Davis, R.J., Muxlow, T.W.B. and Conway, R.G., 1985, *Nature* **318**, 343.
- Douglas, J.N., Bash, F.N., Ghigo, F.D., Mosely, G.F. and Torrence, G.W., 1973, *Astrophys. J.*, **78**, 1.
- Dreher, J.W., 1981, *Astron. J.*, **86**, 833.
- Drury, L.O'C., 1983, *Rep. Prog. in Physics*, **46**, 975.
- Edge, D.O. and Mulkay, M.J., 1976, 'Astronomy Transformed'. *Publ John Wiley*
- Ekers, R.D., 1985, *Chapter 10, 'Synthesis Imaging'*,
Publ N.R.A.O. Eds. Perley, R.A., Schwab, F.R. and Bridle, A.H., p 149.
- Fanselow J.L., Sovers, O.J., Thomas, G.H., Purcell, G.H., Cohen, E.J., Rogstad, D.H., Skjerve, L.J., Spitzmesser, D.J., 1984, *Astron. J.*, **89**, 987.
- Fomalont, E.B. and Wright, M.C.H., 1974, 'Galactic And Extragalactic Astronomy'. *Eds. Verschuur, G.I. and Kellermann, K.I.*, *Publ Springer-Verlag, Berlin, West Germany.*, p. 256.
- Ginzburg, V.L. and Syrovatski, S.I., 1969, *Ann. Rev. Astron. Astrophys.*, **7**, 375.
- Gull, S.J. and Skilling, J., 1984, *Proc. Int Symp 'Indirect Imaging', Sydney Australia 1983. Ed. Roberts, J.A.*, *Publ C.U.P.*
- Hanbury Brown, R., Jennison, R.C. and Das Gupta, M.K., 1952, *Nature*, **170**, 1061.
- Heavens, A.F. and Meisenheimer, K., 1987, *Mon. Not. R. astr. Soc.*, **225**, 235.
- Heckman, T.M., Miley, G.K., Balik, B., Van Breugel, W.J.M. and Butcher, H.R. *Astrophys. J.*, 1982, **262**, 529.
- Hinder, R. and Ryle, M., 1971, *Mon. Not. R. astr. Soc.*, **154**, 229.
- Högbom, J.A., 1974, *Astr. Astrophys. Suppl.*, **15**, 417.
- Holmes, G.F. and Woodall, P., 1987, MSc Technical Project,

Univ of Manchester.

- Jaffe, W.J. and Perola, G.C., 1974, *Astron. Astrophys*, **26**, 423.
- Junor, W., 1987, *PhD Thesis, University Of Manchester*.
- Kardeshev, N.S., 1962, *Soviet Astronomy*, Vol **6**, No **3**, 317.
- Kelder, H. and Spolestra, T.A.T., 1987, *Journal of Atmospheric and Terrestrial Physics*, **49**, 7.
- Kellermann, K.I. and Pauliny-Toth, I.I.K., 1981, *Ann. Rev. Astron. Astrophys*, **19**, 373.
- Kock, W.E. and Stone, J.L., 1958, *Proc I.R.E.*, **46**, 499.
- Kundt, W. and Saripelli, L., 1987, *Jl. Astrophys. Astron.*, **8**, 211.
- Laing, R.A., 1984, 'Physics of Energy Transport in Extragalactic Radio Sources' Proc. N.R.A.O. Workshop No 9. Ed Bridle, A.H. and Eilek, J.A.
- Lang, K.R., 1974, 'Astrophysical Formulae', Publ. Springer-Verlag, Berlin, W. Germany.
- Lind, K.R., 1986, *PhD Thesis, California Institute of Technology*.
- Longair, M.S., 1981, 'High Energy Astrophysics.' Publ C.U.P.
- MacDonald, G.H., Kenderdine, S. and Neville, A.C., 1968, *Mon. Not. R. astr. Soc*, **138**, 259.
- Manchester, R.N., 1985, *Australia Telescope Memo*, **AT/20.1.1/014**.
- Marscher, A.P, 1987, *Proc 'Superluminal Radio Sources', Big Bear, California*. Eds. Zensus, J.A. and Pearson, T.J. p. 280.
- Masson, C.R., 1986, *Astrophys. J. Lett.*, **302**, L27.
- Mathewson, D.S., Ford, V.L., Dopita, M.A., Tuohy, I.R., Long, K.S. and Helfand, D.J., 1983, *Astrophys. J. Suppl*, **51**, 345.
- Mayer, C.H., McCullough, T.P. and Sloanaker, R.M., 1957, *Astrophys. J.* **126**, 468.
- MERLIN Phase 2 Proposal, 1987, *A Proposal Submitted to the S.E.R.C. The Jodrell Bank Staff*.
- MERLIN Phase 2 Proposal, 1988, *A Proposal Submitted to the S.E.R.C. The Jodrell Bank Staff*.

- McPhie, R.H., 1967. *VLA Scientific Memo*, **3**
- McPhie, R.H., 1979. *Radio. Sci.*, **14**, 1185.
- Murphy, D.W., 1988, *PhD Thesis. University Of Manchester.*
- Myers, S.T. and Spangler, S.R., 1985, *Astrophys. J.*, **291**, 52.
- Napier, P.J., 1983, Thompson, A.R. and Ekers, R.D., *Proc. I.E.E.E.*,
71, 1295.
- Narayan, R. and Nityananda, R., 1986, *Ann. Rev. Astron. Astrophys.*,
24, 127.
- Nityananda, R. and Narayan, R., 1982, *Jl. Astrophys. Astron. (India)*,
3, 419.
- Norman, M.L., Smarr, L. and Winkler, K.H.A, 1982, *Astron. Astrophys.*,
113, 285.
- Norris, R.P., 1987, '*The Impact Of VLBI On Astrophysics And Geophysics*'
I.A.U. Symp. **129**.
- O'Brien, P.A., 1953, *Mon. Not. R. astr. Soc.*, **113**, 597.
- Okoye, S.E. and Obinabo, O., 1982, *I.A.U. Symp.* **97**, 71.
- Owen, F.N. and Purschell, J.J., 1984, *Astron. J.*, **89**, 932.
- Pacholczyk, A.G., 1970, '*Radio Astrophysics*'. *Publ. Freeman.*
- Pearson, T.J. and Readhead, A.C.S., 1984, *Ann. Rev. Astron.*
Astrophys., **22**, 97.
- Perley, R.A., Dreher, J.W. and Cowan, J.J., 1984, *Astrophys. J.*,
285, L 35.
- Porcas, R.W., 1981, *Nature*, **294**, 47.
- Porcas, R.W., 1984, '*VLBI and Compact Radio Sources*', *I.A.U. Symp.* **110**
Ed. Fanti. R., Kellermann, K.I. and G.Setti. *Publ. G.Reidel, Holland.*
p 157.
- Ray, T.P., 1981, *Mon. Not. R. astr. Soc.*, **196**, 195. Readhead, A.C.S. and
Wilkinson, P.N., 1978, *Astrophys. J.*,
223, 25.
- Rodgers, E.E., 1970, *Radio. Sci.* **5**, 1239.
- Romney, J.D., 1987, '*The Impact Of VLBI On Astrophysics And Geophysics*'

I.A.U. Symp. 129.

- Rowson, B., 1963, *Mon. Not. R. astr. Soc*, **125**, 177.
- Rudnik, L. and Edgar, B.K., 1984, *Astrophys. J.*, **279**, 74.
- Ryle, M. 1962, *Nature*, **194**, 517.
- Ryle, M. 1972, *Nature*, **239**, 435.
- Ryle, M. and Hewish, A., 1960, *Mon. Not. R. astr. Soc*, **120**, 220.
- Scheuer, P.A.G., 1984, 'The Early Years Of Radio Astronomy.'
Ed. W.T. Woodruff. Pub C.U.P., p.249.
- Schwarz, U.J., 1978, *Astr. Astrophys*, **65**, 345.
- Schwarz, U.J., 1984, 'Indirect Imaging', p 333. Ed. Roberts, J.A.
Publ. Cambridge University Press.
- Shklovaski I. 1953, *Dokl. Akad. Nauk. SSSR.*, **90**, 983.
- Shone, D.L., 1985, *PhD Thesis, University of Manchester*
- Shone, D.L., Porcas, R.W. and Zensus, J.A., 1985, *Nature*, **314**, 603.
- Siebler, W., Salter, C.J. and Mayer, C.J. 1981, *Astr. Astrophys.*,
103, 393.
- Somerén Greve, H.W. Van., 1973, *Astr. Astrophys. Suppl.*, **3**, 419.
- Sramek, R.A. and Schwab, F.R., 1985, *Chapter 5, 'Synthesis Imaging'*,
Publ N.R.A.O. Eds. Perley, R.A., Schwab, F.R. and Bridle, A.H.
- Stephens, P.J., 1987, *Ph.D. Thesis. University Of Manchester.*
- Swarup, G, 1987, *Giant Metre-Wave Radio Telescope, Status Report.*
- Swenson, G.W. and Mathur, N.C., 1968, *Proc I.E.E.E.*, **56**, 2114.
- Swenson, G.W. and Mathur, N.C., 1969, *Radio. Sci.* **4**, 69.
- Tan, S.M., 1986, *Mon. Not. R. astr. Soc.*, **220**, 971.
- Thomasson, P., 1986, *Q. Jl. R. astr. Soc.*, **27**, 413.
- Thompson, A.R., 1982, Proc Conf 'Interference Identification and
Excision', Publ. N.R.A.O. p.46.
- Thompson, A.R. and D'Addario, L.R., 1982, *Radio. Sci.*, **17**, 357.
- Thompson, A.R., Moran, J.W., Swenson, G.W., 1986, 'Interferometry
And Synthesis In Radio Astronomy.' Publ J.Wiley
- Titterton, G.M., 1985, *Astron. Astrophys.*, **144**, 381.

- Van Breugel, W., Heckman, T. and Butcher, H., 1984, *Ap.J.*, **277**, 82.
- Van Breugel, W., Miley, G. and Heckman, T., 1984, *Astron. J.*, **89**, 5.
- Walker, R.C., Benson, J.M. and Unwin, S.C., 1987, *ApJ*, **316**, 546.
- Webb, G.M., Bogdan, T.J., Lee, M.A. and Lerche, I., 1985, *Mon. Not. R. Soc.*, **215**, 341.
- Whitfield, G.R., 1957, *Mon. Not. R. astr. Soc*, **117**, 680.
- Wilkinson, P.N., 1983, 'Very Long Baseline Interferometry Techniques.' *Publ CEPAD. Toulouse.* p. 375.
- Wilkinson, P.N., 1987, *Proc 'Superluminal Radio Sources', Big Bear, California.* Eds. Zensus, J.A. and Pearson, T.J. p 211.
- Williams, A.G. and Gull, S.F., 1984, *Nature*, **310**, 33.
- Wilson, M.J. and Falle, S.A.E.G., 1985, *Mon. Not. R. astr. Soc*, **216**, 971.

ProQuest Number: 30202995

INFORMATION TO ALL USERS

The quality and completeness of this reproduction is dependent on the quality and completeness of the copy made available to ProQuest.



Distributed by ProQuest LLC (2022).

Copyright of the Dissertation is held by the Author unless otherwise noted.

This work may be used in accordance with the terms of the Creative Commons license or other rights statement, as indicated in the copyright statement or in the metadata associated with this work. Unless otherwise specified in the copyright statement or the metadata, all rights are reserved by the copyright holder.

This work is protected against unauthorized copying under Title 17,
United States Code and other applicable copyright laws.

Microform Edition where available © ProQuest LLC. No reproduction or digitization of the Microform Edition is authorized without permission of ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 - 1346 USA