Chapter 2

Instrumentation and calibration

2.1   Introduction

The work presented in this thesis makes use of aperture synthesis radio telescopes, the basic principles of which are explained in Section 2.2 below. This thesis focuses mainly on the results of observations made with the Australia Telescope Compact Array (ATCA), near Narrabri, New South Wales, and the Australia Telescope Long Baseline Array (LBA), both of which are operated as part of the CSIRO's Australia Telescope National Facility (ATNF). For variability studies, accurate flux density measurements are important, hence the emphasis in this chapter on calibration. Important considerations for the use of the ATCA for variable source monitoring are discussed in Section 2.3, including the advantages and disadvantages of using an interferometer, as opposed to a single dish telescope, for flux density measurements. Polarimetric and gain calibration of the ATCA is also discussed. Section 2.4 describes the LBA and calibration of LBA data.

2.2   Fundamentals of radio interferometry and aperture synthesis

The published collection of lectures from the Sixth NRAO/NMIMT Synthesis Imaging Summer School (eds. Taylor, Carilli, & Perley 1999; hereafter TCP99) contains a great deal of detailed discussion and treatment of various topics in aperture synthesis. Lecture 1 (Clark) gives the fundamental relations and simplifying assumptions used in aperture synthesis. The main results are briefly summarised here, along with additional results which are important for the work presented in this thesis. Commonly used terminology and units of measurement are also introduced.

An interferometer measures the *spatial coherence function*, V_ν, of the electric field between pairs of elements (antennas), located at positions $\mathbf{r}_1$ and $\mathbf{r}_2$. This is related to the observed *intensity*, I_ν, of radiation from an astronomical source (assumed to be incoherent) in direction $\mathbf{s}$, by

$$V_\nu(\mathbf{r}_1, \mathbf{r}_2) = \int I_\nu(\mathbf{s}) e^{-2\pi i\nu \mathbf{s}(\mathbf{r}_1-\mathbf{r}_2)/c} d\Omega \tag{2.1}$$

where the integration is over the celestial sphere. The complex quantity V_ν is commonly called the *visibility*. The magnitude of V_ν has dimensions of spectral power flux density ($\mathrm{W\,m^{-2}\,Hz^{-1}}$), while the intensity I_ν is measured in $\mathrm{W\,m^{-2}\,Hz^{-1}\,sr^{-1}}$. In radio astronomy, the "Jansky" (Jy) is used as a measure of flux density, where $1\,\mathrm{Jy} = 10^{-26}\,\mathrm{W\,m^{-2}\,Hz^{-1}}$.

Equation 2.1 can be written in a coordinate system where $\mathbf{r}_1 - \mathbf{r}_2 = \lambda(u, v, w)$ (wavelength $\lambda = c/\nu$) and $\mathbf{s} = (l, m, n = \sqrt{1 - l^2 - m^2})$ as shown in Figure 2.1. For practical applications it is usually possible to set $w = 0$ and reduce Equation 2.1 to a two-dimensional Fourier transform:

$$V_\nu(u, v, w \equiv 0) = \int\int I_\nu(l, m) \frac{e^{-2\pi i(ul+vm)}}{\sqrt{1 - l^2 - m^2}} dl\, dm. \tag{2.2}$$

In many cases, Equation 2.2 can be inverted to obtain the source intensity distribution from V_ν. An additional factor in Equation 2.2 is the sensitivity of the interferometric elements as a function of direction, which is often called the *primary beam* of the antennas. This usually limits the field of view from a single pointing to a discrete region on the sky, but is generally well-determined and is a convenience rather than a problem, because setting $w = 0$ in Equation 2.2 often relies on the assumption that the radiation comes from a small region on the sky. Additionally, this implies l and m are small and $n = \sqrt{1 - l^2 - m^2}$ is close to unity. In the special case of coplanar baselines (see Section 2.2.2) it is always possible to choose coordinates such that $w = 0$.

2.2.1 The effect of discrete sampling

In practice, V is sampled only at particular places on the (u, v) *plane*. The sampling can be described by a *sampling function $S(u, v)$*, which is zero where no data have been taken. Then the function
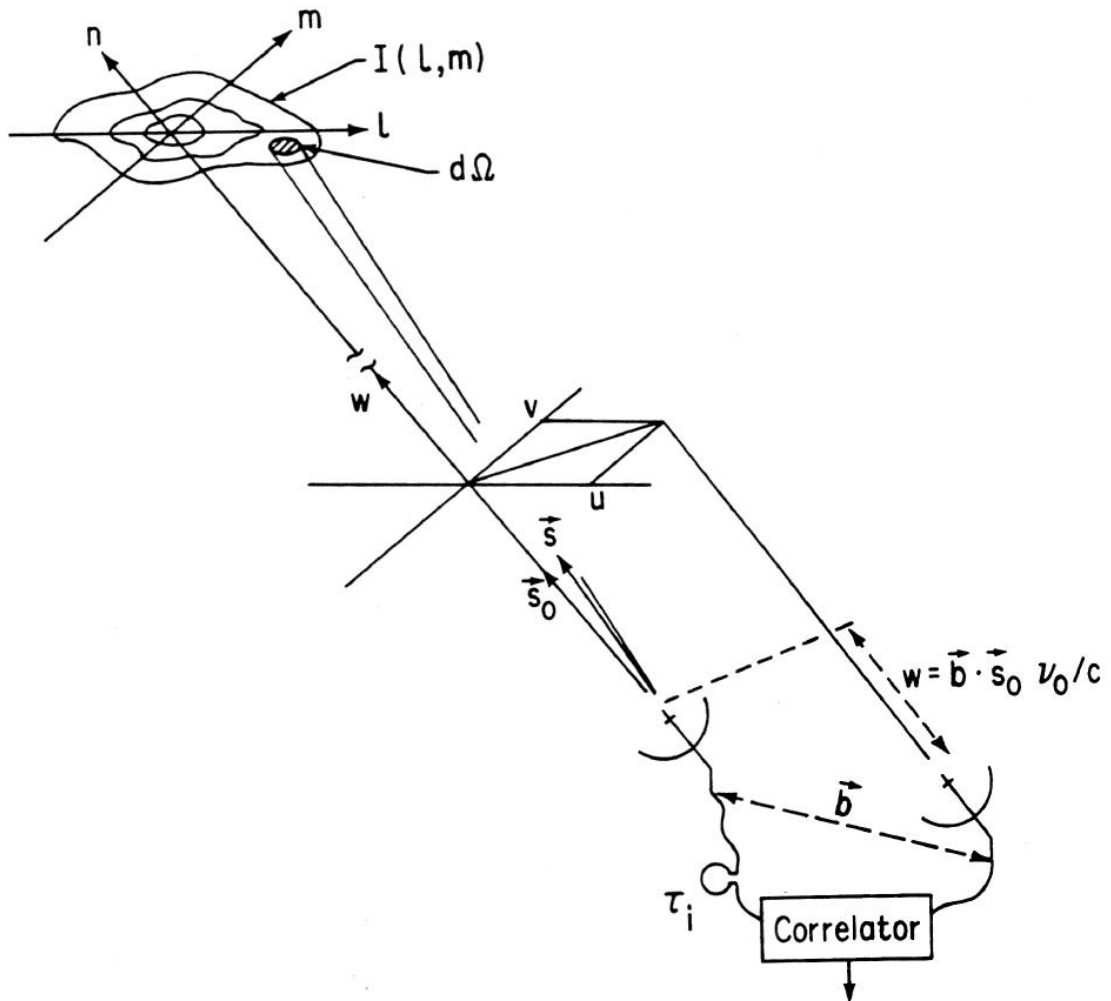
$$I_\nu^D(l, m) = \int\int V_\nu(u, v) S(u, v) e^{2\pi i(ul+vm)} du\, dv \tag{2.3}$$

is often referred to as the *dirty image*. $I_\nu^D(l, m)$ is the convolution of the actual intensity distribution $I_\nu(l, m)$, with the *synthesised beam* or *point spread function B*:

$$I_\nu^D = I_\nu * B, \tag{2.4}$$

where B is the Fourier transform of the sampling function:

$$B(l, m) = \int\int S(u, v) e^{2\pi i(ul-vm)} du dv. \tag{2.5}$$

Figure 2.1: Coordinate systems used for synthesis imaging. The (u, v, w) and (l, m, n) right-handed coordinate systems are used to express interferometer baselines and source brightness distribution, respectively. From Lecture 2 (Thompson), p. 20 of *Synthesis Imaging in Radio Astronomy II*, ed. Taylor et al. 1999.

The process of obtaining an estimate of the true source intensity distribution I_ν from the *dirty image* I_ν^D, is known as *deconvolution*, and a number of algorithms have been developed for this purpose. These are not described in detail here, but see for example Lecture 8 (Cornwell et al.) of TCP99. Deconvolution using CLEAN algorithms (e.g. Högbom, 1974) has been used for some of the data presented in this thesis, but was not required for the most part, as most observed sources are point-like and the data were analysed in the visibility domain.

Confusion

Confusing sources are those which lie within the primary beam response area on the sky, that are not related to the target. If a small field of view is imaged, then confusing sources outside this area may be *aliased* into the image, or *sidelobes* of the synthesised beam response to the confusing source may produce artefacts in the image. If confusing sources can be well modelled, for example by making a low resolution image of the full primary beam area, then it may be possible to subtract them from the visibilities. For sparsely sampled data it is not always possible to accurately model confusing emission, and confusion may limit the sensitivity of an image or the accuracy of a flux density measurement. The effect of confusion on flux density measurements is further discussed in Section 2.3.4.

Techniques for sparsely sampled data

Image deconvolution effectively involves extrapolating V over unsampled regions of the (u, v) plane. If the (u, v) coverage is sparse, the reliability and quality of an image are limited; artefacts may be introduced by poor "guesses" of V where no data exist. In this case, provided there is sufficient signal-to-noise and provided the source can be described by a simple model, it is often better to fit a model directly to the measured visibilities using as few parameters as possible to obtain a good fit. This process is known as *model-fitting*. An initial estimate of the model can often be obtained by inspection of the visibility data, using simple Fourier transform relations (e.g. Fomalont & Wright, 1974). While a model obtained from sparse data may be a simplified representation of the true source brightness distribution, some source parameters can always be constrained in the visibility domain. Model-fitting is sometimes also used in conjunction with deconvolution, if the goal is to produce an image.

2.2.2   Earth-rotation synthesis

Earth-rotation synthesis arrays make use of the Earth's rotation to vary the sampling of the (u, v) plane. As the Earth rotates over the course of a day, the projected baseline between each pair of interferometer elements changes, so that the (u, v) locations being sampled are constantly changing. This makes it possible to increase (u, v) coverage with only a small number of elements, and without having to physically move antennas around on the ground. In the case of an east-west array such as the Australia Telescope

Compact Array, the components of the baseline vectors parallel to the Earth's rotation axis are zero; the baseline vectors remain coplanar as the Earth rotates. This means that it is always possible to set $w = 0$ by choosing the w-axis to lie along the celestial pole.

2.2.3 Calibration

The quantities $\tilde{V}$ measured with a real interferometer are somewhat corrupted versions of the true complex visibilities V. This statement can be expressed as:

$$\tilde{V}_{ij}(t) = \mathcal{G}_{ij}(t)V_{ij} + \text{noise}, \tag{2.6}$$

where $\mathcal{G}_{ij}(t)$ is a baseline-based complex gain factor. $\mathcal{G}_{ij}(t)$ can be approximated by the product of the two associated antenna-based complex gains $g_i(t)$ and $g_j(t)$:

$$\mathcal{G}_{ij}(t) = g_i(t)g_j^*(t) = a_i(t)a_j(t)e^{i(\theta_i(t)-\theta_j(t))}, \tag{2.7}$$

where $a_i(t)$ is the antenna-based amplitude correction and $\theta_i(t)$ is the antenna-based phase correction. In a well-designed system any residual baseline-based complex gain, which cannot be factored into a product of antenna gains, has amplitude close to unity and phase close to zero, and can usually be ignored. The ATCA and LBA correlators use well-matched filters, digital filters in the case of the LBA, to minimise baseline-based errors, and the spectral channels at each end of the pass-band, where errors may be significant, are discarded in the data reduction process.

The relative antenna gains, and some source parameters, can be determined from observations using *self-calibration* techniques, usually by observing a source for which some parameters are known; ideally one which is unresolved, not affected by confusion and which does not vary over the course of the observation. Self-calibration is a process whereby antenna gain corrections are derived in order to fit the data to a model source brightness distribution. Generally, the amplitude gain corrections are scaled using a standard calibration source, for which the flux density and spectrum are known accurately and measured to be stable on long time-scales.

The antenna-based phase corrections, $\theta_i(t)$, are generally much more variable than the amplitude, as the phase is affected by atmospheric fluctuations. For high signal-to-noise observations, which is the case for most of the data presented here, phase-only self-calibration can often be used to determine $\theta_i(t)$ for an unknown source brightness distribution. An iterative procedure of deconvolution or model-fitting, and phase self-calibration, is used until convergence is reached between the data and the source model. An important constraint in this process is that the model must be consistent with the *closure phases* (see below).

2.2.4 Closure quantities

The complex visibility between two antennas i and j may be written as $V_{ij} = |V_{ij}|e^{\phi_{ij}}$. Then Equations 2.6 and 2.7 imply that the observed phase $\tilde{\phi}_{ij}$ is related to the true

visibility phase ϕ_{ij} by

$$\tilde{\phi}_{ij}(t) = \phi_{ij}(t) + \theta_i(t) + \theta_j(t) + \text{noise term.} \tag{2.8}$$

The *closure phase* is the sum of visibility phases around a closed loop of baselines. For a loop of three baselines formed from antennas i, j, and k, the observed closure phase $\tilde{C}_{ijk}(t)$ is given by

$$\begin{aligned}
\tilde{C}_{ijk}(t) &= \tilde{\phi}_{ij}(t) + \tilde{\phi}_{jk}(t) + \tilde{\phi}_{ki}(t) \\
&= \phi_{ij}(t) + \phi_{jk}(t) + \phi_{ki}(t) + \text{noise term} \\
&= C_{ijk}(t) + \text{noise term.}
\end{aligned} \tag{2.9}$$

Thus the observed closure phase, as first noted by Jennison (1958), is independent of antenna-based errors such as phase variations due to atmospheric turbulence. Closure phases are particularly useful in Very Long Baseline Interferometry (VLBI), where each station has an independent local oscillator, and hence absolute phase information is lost.

A *closure amplitude* Γ_{ijkl} can be defined for any loop of 4 antennas:

$$\Gamma_{ijkl}(t) = \frac{|\tilde{V}_{ij}(t)||\tilde{V}_{kl}(t)|}{|\tilde{V}_{ik}(t)||\tilde{V}_{jl}(t)|} \tag{2.10}$$

The amplitudes of the complex gains cancel out of these ratios, and thus the observed and true closure amplitudes should be identical apart from noise. For a point source, the closure amplitude is equal to unity.

Closure quantities are independent of antenna-based errors, which makes them "good observables" that can be used to constrain source parameters. In Section 6.5.5, closure phases are used explicitly to constrain the structure of a source exhibiting rapid changes in flux density, where standard imaging was not possible. Self-calibration procedures used to solve for the complex antenna gains also implicitly make use of closure relations (Cornwell, 1982).

Estimation of extended structure from closure phases

The majority of sources observed in this thesis are mostly unresolved on arcsecond angular scales. Much of the analysis is performed directly on the visibility data. In this case it is extremely useful to have an estimate of how well the data can be fitted by a point source model, as this greatly simplifies self-calibration and analysis. For examining flux density variability, particularly intraday variability, it is often necessary to remove the effects of extended structure and confusing sources from the visibilities.

Provided signal-to-noise is high ($S/N \gg 1$), the flux density in extended components, S_{ext}, can be estimated from the closure phases using

$$S_{\text{ext}} \propto S(\sigma_{\text{obs}}^2 - \sigma_{\text{theo}}^2)^{1/2} \tag{2.11}$$

where:

S is the (average) total flux density;

σ_{obs} is the observed root-mean-square (rms) variation of closure phases measured for all independent baselines;

σ_{theo} is the theoretically predicted rms due to thermal noise of the closure phases for a point source (for the ATCA, ~ 0.02 radians for a 1 Jy source);

and the proportionality constant is of order unity (Kedziora-Chudczer et al., 2001b). This method allows sources which are essentially unresolved to be recognized quickly. It also provides an estimate of the amount of extended structure or confusion, which can then be followed up with imaging and subtracted from the visibilities if required. Alternatively, extended structure may sometimes be discarded by selection based upon (u, v) range.

2.3 Use of the ATCA for flux density measurements

The Australia Telescope Compact Array (ATCA) is an earth-rotation aperture synthesis interferometer, consisting of six 22-m antennas on an east-west track, with a maximum baseline of 6 km (Frater et al., 1992). The ATCA has a number of advantages over a single dish telescope for flux density monitoring of compact sources. These include:

- continuous, high-accuracy measurements on-source, without the need for on/off source switching

- accurate flux density measurements: gain amplitude errors are minimised by on-line system temperature corrections

- sensitive measurements of the full set of Stokes parameters.

- simultaneous dual-frequency measurements, and efficient turret rotation $\longrightarrow$ broadband instantaneous frequency coverage

The main disadvantage of using the ATCA for flux density measurements is the varying response of the synthesised beam to extended source structure and the presence of possible confusing sources. These advantages and disadvantages are explained in more detail below, with discussion of calibration procedures for ATCA data.

2.3.1 Sensitivity and efficiency

Being an interferometer, the ATCA enables very high signal-to-noise flux density measurements without the need for the background or off-source subtraction necessary for a single dish. The rms flux sensitivity of the ATCA, ΔS_{rms}, is given by (ATCA Users Guide 2002, equation 6):

$$\Delta S_{\mathrm{rms}} = 0.0195\, S_{\mathrm{sys}}/\sqrt{tB} \quad \mathrm{mJy}, \qquad (2.12)$$

where:

S_{sys} is the system sensitivity in Jy, defined as the signal which doubles the system temperature at elevation 40°. This is dependent on observing frequency but is typically ~ 400 Jy for the frequencies used here (1.4, 2.5, 4.8 and 8.6 GHz);

t is the integration time in minutes;

B is the bandwidth in MHz.

Equation 2.12 implies that the theoretical rms noise for a 1-minute integration with the usual 128 MHz continuum bandwidth is < 1 mJy. In practice, the presence of confusing sources in the primary beam often makes the error in a flux density measurement much larger than ΔS_{rms}, especially at observed frequencies of $\sim 2.5\,\mathrm{GHz}$ and below (see Section 2.3.4).

As well as signal-to-noise, calibration is important for accurate flux density measurements. System temperature at the ATCA is measured on-line with the use of a thermally stabilised noise diode (Nelson, 1992), which makes for well-calibrated flux density measurements which are relatively stable with time. Accuracies of 1% or better have proven readily achievable (e.g. Kedziora-Chudczer et al., 2001b, and see Section 2.3.5).

2.3.2 Frequency coverage

The ATCA currently operates in four cm wavelength bands; 20, 13, 6 and 3 cm. In each antenna the 20 and 13 cm receivers use one wideband feedhorn, and the 6 and 3 cm receivers use another. It is therefore possible to observe simultaneously either at 20 and 13 cm wavelengths, or at 6 and 3 cm wavelengths. The feedhorns are mounted on a rotating turret, and it is possible to automatically switch between them in 30 seconds, although turret rotation is allowed no more than once per 15 minutes on average, due to the stress it places on the bearings. The set-up allows almost instantaneous, four-frequency spectra to be measured. In continuum mode, the ATCA correlator provides 128 MHz bandwidth at each frequency, made up of 32 overlapping, 8 MHz channels.

2.3.3 Polarization

The ATCA uses on-axis, orthogonally polarized linear feeds. This allows a complete description of the polarized state of the incident radiation. It is common to describe the polarization state of a radiation field in terms of the four Stokes parameters, I, Q, U, and V (first introduced by G. G. Stokes in 1852). Stokes I corresponds to total intensity, the sum of all polarization states. Q and U describe the linear polarization, and V represents the circularly polarized component, where positive V is by convention defined as right-hand circular polarization (see Born & Wolf 1975 for a complete description of the Stokes parameters). Linear polarization can be equivalently represented in terms of the linearly polarized flux density, p, and position angle ψ of the electric vector of linear polarization, where

$$p = \sqrt{Q^2 + U^2}, \tag{2.13}$$

and

$$\psi = \frac{1}{2} \arctan\left(\frac{U}{Q}\right). \tag{2.14}$$

It is also common to quote fractional linear polarization, $m_p = p/I$, or fractional circular polarization, $m_V = |V|/I$.

For ideal linearly polarized feeds (X, Y), and antennas with an altitude-azimuth ('alt-az') mount such as those of the ATCA, the four polarization products, corresponding to the cross-correlations between the feeds on two antennae, can be converted to Stokes parameters using the following equations (see, e.g., Sault et al., 1991):

$$\begin{aligned}
I &= \frac{1}{2}(XX + YY) \\
Q &= \frac{1}{2}[(XX - YY)\cos(2\chi) - (XY + YX)\sin(2\chi)] \\
U &= \frac{1}{2}[(XX - YY)\sin(2\chi) - (XY + YX)\cos(2\chi)] \\
V &= \frac{1}{2}i(YX - XY)
\end{aligned} \tag{2.15}$$

where χ is the parallactic angle, χ_{par}, plus a constant. Parallactic angle is defined as:

$$\chi_{\mathrm{par}} = \arctan\left(\frac{\cos(\mathcal{L})\sin(h)}{\sin(\mathcal{L})\cos(\delta) - \cos(\mathcal{L})\sin(\delta)\cos(h)}\right) \tag{2.16}$$

where δ is the source declination, $\mathcal{L}$ is the geodetic latitude of the antenna, and h is the source hour angle. The parallactic angle describes the rotation of an alt-az mounted antenna, and hence its feeds, with respect to the source, over the course of an observation. For the ATCA, the X feed is offset by $45°$ from Celestial North at transit, so $\chi = \chi_{\mathrm{par}} + 45°$.

In reality, feeds are never ideal and it is necessary to solve for "leakage terms", which can be thought of as the leakage of a small amount of one component of the electric field into the orthogonal component's feed. For the ATCA, leakage terms are found to be relatively stable over periods of several months, and hence are assumed to be time-independent for calibration of data from a given observing session. The leakages can be derived with high precision using observations of calibration sources - strong, unresolved sources which are assumed not to vary within an observing session.

The advantage of the relative rotation of the feeds with parallactic angle for an alt-az telescope, as a source is tracked across the sky, is that by observing just one calibration source at several different parallactic angles, it is possible to decouple the leakages and the calibrator linear polarization. Thus the polarization characteristics of the calibrator need not be known *a priori*, provided it is observed over a range of parallactic angle.

Polarimetric calibration also requires solving for the time-dependent complex antenna gains, which can generally be done on the same calibration source. It is also necessary to know the phase difference between the X and Y polarization channels, which is measured on-line at the ATCA using a noise-diode system (Nelson, 1992).

As some sources exhibit short time-scale variability, calibration sources should ideally be selected from well-studied sources which have been observed to be non-variable over long periods of time (at least weeks). Many sources known as Compact Steep-Spectrum (CSS) and Gigahertz Peaked Spectrum (GPS) sources show little evidence of variability on time-scales of several years, and thus are often reliable calibrators (King, 1994).

An alternative method of solving for the leakage terms is to observe a source with known polarization characteristics. The ATCA primary calibrator, PKS 1934−638, has been measured to have $\lesssim 0.1\%$ linear polarization (Komesaroff et al., 1984). Leakage solutions adequate for measurements of linear polarization of the target sources can therefore be found from a single scan on PKS 1934−638. This was the approach used to calibrate most data presented in this thesis.

Stokes V is often an order of magnitude weaker than the linear polarization, and very accurate calibration is required to detect and measure it. Calibration for Stokes V involves solving the 'strongly polarized' equations (Sault et al., 1991), which contain some additional terms that are not usually important when solving only for linear polarization. Solving the 'strongly polarized' equations requires observation of a calibration source which has a few percent linear polarization. Under good conditions, fractional circular polarization, V/I, can be measured to an accuracy of $\sigma_V/I \sim 0.01\%$ (Rayner et al., 2000). Circular polarization measurements are presented in this thesis only for a subset of the data, in Chapter 5.

The necessary tasks to find gain and leakage solutions, and apply the corrections to the data, are incorporated into the MIRIAD software package (Sault et al., 1995) and are part of standard procedures for reduction of ATCA data. The specific calibration "recipes" used for data presented in this thesis are described in Sections 3.2.3 and 5.4.

2.3.4   Source structure, confusion, and observing with "cuts"

The majority of sources observed for this thesis contain nearly all of their flux density in a compact component, unresolved with the ATCA on arcsecond scales. For monitoring variability of such sources, accurate measurements of point source flux density are of more concern than the full synthesis imaging capabilities of the ATCA. For an isolated point source, the visibility amplitude is uniform everywhere and thus the measured flux density is independent of (u, v) coverage; in this idealised case the flux density can be measured from just a single short scan. However, there are often "confusing" sources close on the sky to the target, within the primary beam of the 22-m antennas. Also, many compact sources have weakly extended emission on arcsecond scales.

For a linear east-west array such as the ATCA, the instantaneous synthesised beam corresponding to the sampling of the visibility function (see Section 2.2.1) is a fan-beam, whose orientation varies with time as the Earth rotates. Thus, the contribution of extended structure varies over the course of a day. A 12-hour integration is required to obtain full two-dimensional information on the sky brightness distribution. However, when the source structure is simple and compact, a limited number of short scans or "cuts" observed over a 12-hour period is enough to obtain a good flux density

measurement and estimate of structure. Thus it is possible to observe a large number of compact sources in a 24 hour period, as was done, for example, in the observations presented in Chapter 3.

For point sources, the flux density can be measured directly from the visibilities. To estimate the accuracy of a flux density measurement it is necessary to know the contribution of extended structure and confusion, which can be estimated from the closure phases as described in Section 2.2.4. An estimate of source structure or possible confusion can also be obtained by inspecting the visibilities, e.g. by plotting amplitude as a function of (u, v) distance.

The severity of confusion increases sharply toward lower frequencies because of the larger primary beam, and also because most confusing emission is steep-spectrum and hence stronger at lower frequencies. Based on radio source counts, the predicted strongest confusing source in the full-width at half-maximum (FWHM) of the ATCA primary beam at 20 cm is ~ 140 mJy, while at 3 cm, it is only 0.4 mJy (ATCA Users Guide 2002), although there are large fluctuations in practice. Extended source structure is also steep-spectrum in general. The large increase in confusing emission towards low frequencies often makes flux density measurements with the ATCA difficult at 13 and 20 cm. Confusion is by far the dominant source of error in flux density measurements at these wavelengths.

2.3.5   ATCA antenna gain-elevation dependence at 3 cm

A large part of this thesis relates to interstellar scintillation and intraday variability (IDV). The level at which IDV is detectable in ATCA data was found to be often limited by systematic effects. The on-line system temperature measurements at the ATCA correct for changes in system temperature which may occur with time and with changing elevation. The behaviour of the ATCA antenna gains as a function of elevation had not previously been characterised, although the effect on the measured visibity amplitudes was known to be less than a few percent. This lack of a standard elevation-dependent gain correction at cm wavelengths is a reflection of the fact that the ATCA is used primarily for synthesis imaging, rather than flux density monitoring. It is more often phase errors, rather than amplitude errors, which limit image fidelity. However, flux density measurements, and in particular detection of IDV, are often limited by gain amplitude errors.

Phase self-calibration can generally be used to correct phases in the case of simple source structure and high signal-to-noise. The standard procedure of observing a non-IDV calibration source, or *secondary calibrator*, close on the sky to the target source, then interpolating gain solutions in time from the secondary on to the target source data, removes most systematic errors. However, for the purpose of measuring flux densities with high accuracy for a large number of sources, this is not an efficient procedure. Furthermore, it is not always possible to find a reliable calibration source within a few degrees of every target.
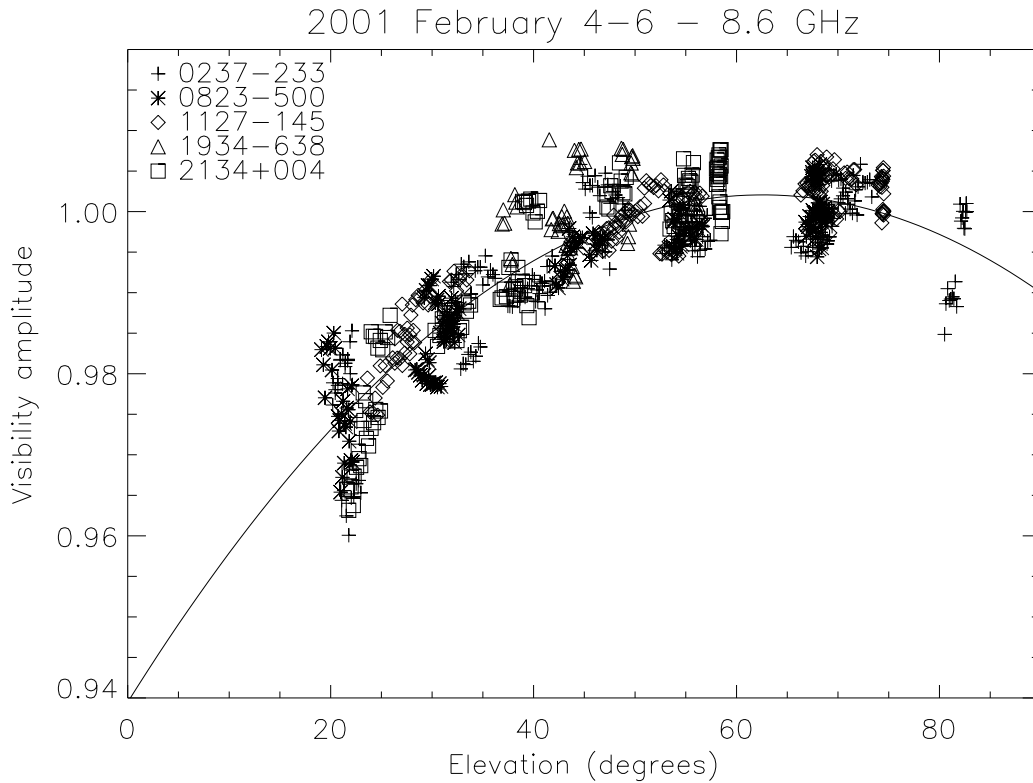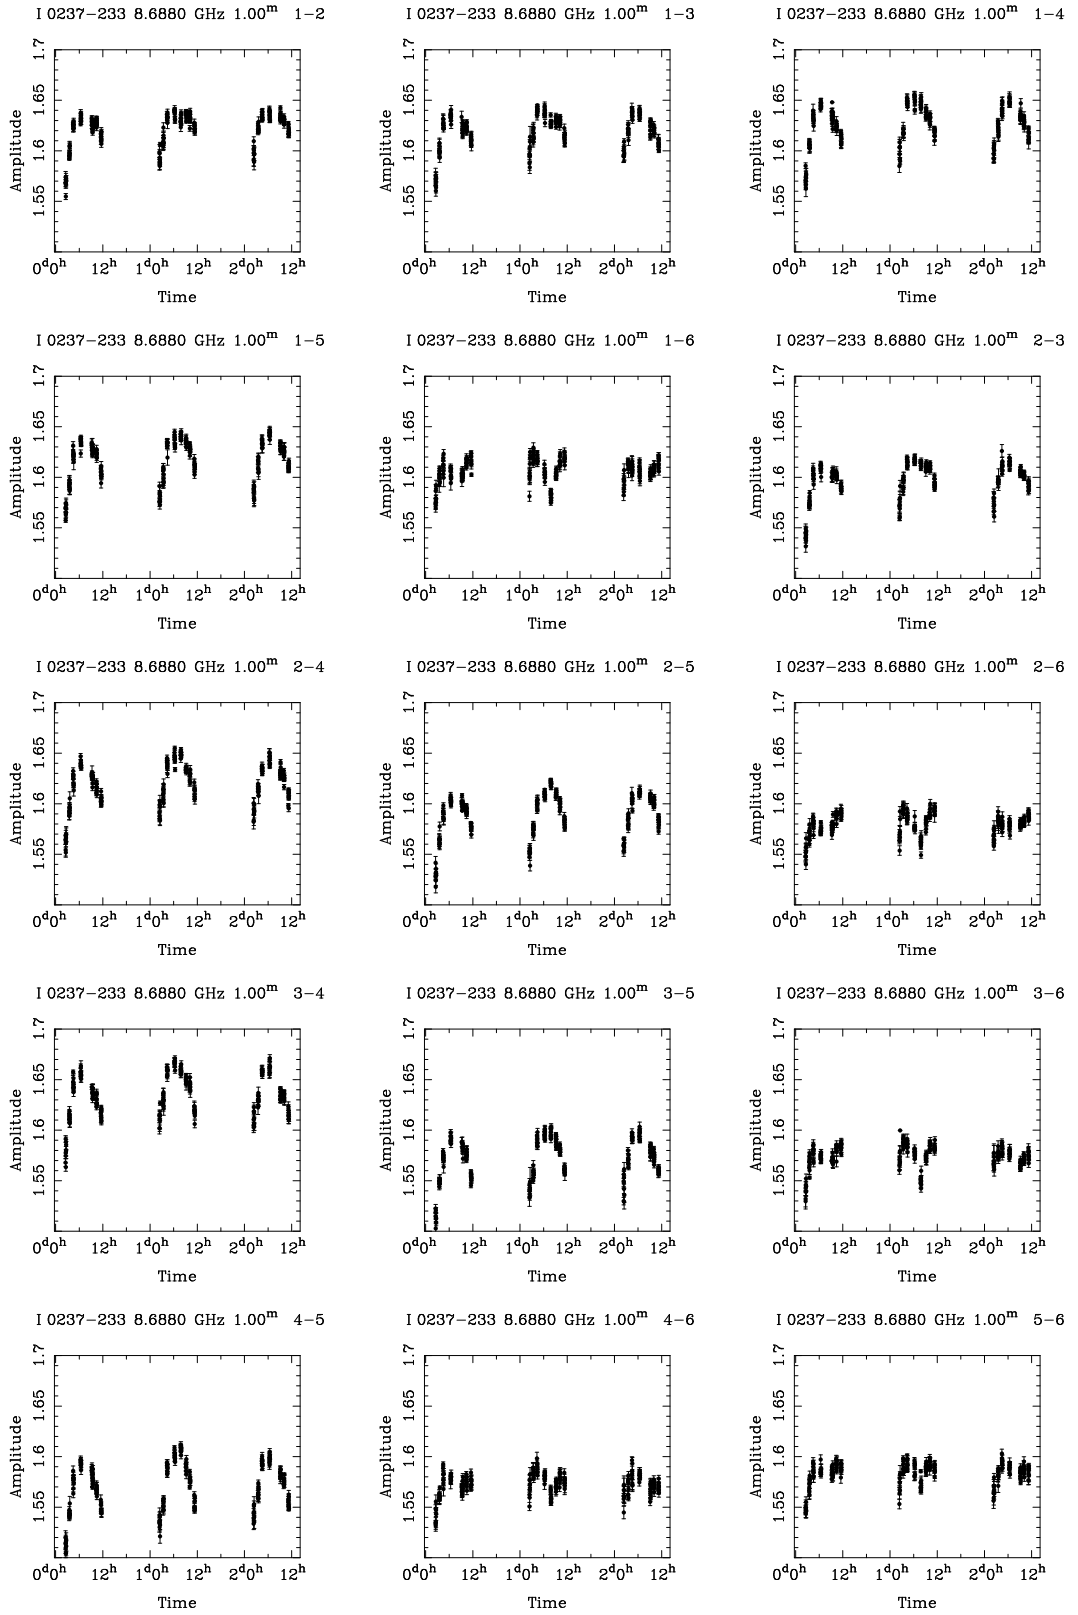
Figure 2.2: Variation of visibility amplitudes with elevation over 3 days for data at 8.6 GHz, averaged over all ATCA baselines and normalised to unity at high elevation. The solid line shows the best-fit quadratic dependence.

The level of systematic variations in ATCA data

Inspection of visibilities for compact sources which were observed over two or more consecutive days showed amplitude variations which repeated on different days and for different sources. Particularly at 8.6 GHz there was an evident elevation dependence of the measured flux density, with the observed systematic variation at a level of ~ 2–3%. To illustrate the effect, Figure 2.2 shows uncalibrated 8.6 GHz data from 5 calibration sources observed over 3 days in February 2001. The visibilities plotted are averaged over all baselines, over the frequency band, and over 1 minute in time, and normalised to unity at high elevation. While Figure 2.2 shows the behaviour of the array as a single telescope, not all antennas show the same behaviour, as seen in Figure 2.3, which shows each baseline plotted separately for the source PKS 0237−233. In particular, the behaviour of Antenna 6 is quite different from the others.

Possible causes of systematic variation

No opacity correction has so far been made to the data. Correction for atmospheric opacity is not a standard procedure for cm-wavelength ATCA data, although a MIRIAD

Figure 2.3: Variation of visibility amplitudes at 8.6 GHz on separate ATCA baselines. The repeated pattern over the 3 days shows the elevation dependence. Baselines to Antenna 6 are clearly different from the others.

Table 2.1: Sources used to derive gain-elevation dependence for ATCA antennas

Name (B1950)	RA (J2000)	Dec (J2000)	$S_{4.8\,\text{GHz}}$ (Jy)	$S_{8.6\,\text{GHz}}$ (Jy)
PKS 0237−233	02:40:08.17	−23:09:15.7	2.8	1.6
PKS 0823−500	08:25:26.87	−50:10:38.5	3.0	1.5
PKS 1127−145	11:30:07.05	−14:49:27.4	4.2	3.5
PKS 1921−293	19:24:51.06	−29:14:30.1	11	13
PKS 1934−638	19:39:25.03	−63:42:45.6	5.8	2.8
PKS 2126−158	21:29:12.18	−15:38:41.0	1.2	1.2
PKS 2134+004	21:36:38.59	+00:41:54.2	10	9.0

task exists to handle such a correction. Opacity is expected to slightly attenuate the signal at low elevations, and indeed a drop-off in signal toward low elevations is evident in Figure 2.2. However, opacity cannot be the sole cause of the observed effect, since opacity would affect all antennas in the same way, and Figure 2.3 shows that different antennas behave differently. In particular, Antenna 6 baselines show variations in the opposite sense to that expected from opacity. Also, Figure 2.2 suggests a possible drop-off in signal at high elevation, which cannot be due to opacity.

Errors in the pointing model may also result in elevation-dependent systematic effects. However, tests using reference pointing on the ATCA primary calibrator, PKS 1934−638, revealed the same elevation dependence using reference pointing as for non-reference pointed observations (D.P. Rayner 2001, private communication). Pointing therefore cannot be the major cause of the observed systematic variation.

The ATCA antenna design employs Cassegrain optics using shaped reflectors. At some level, all antenna gains change with elevation, due to surface deformation and focal length changes. This is probably the dominant source of amplitude gain variations in 3 cm data.

### Derivation of 3 cm gain curves from data

The gain-elevation dependence of the ATCA antennas in the 3 cm band was derived using calibration source data from three observing sessions in early 2001. No separate opacity correction was done, therefore opacity effects are incorporated into the derived gain curves. In the 6 cm band, the gain variations are significantly smaller, and no elevation correction has so far been derived for 6 cm data.

A gain curve was derived for each individual antenna rather than the array as a whole. Sometimes not all six antennas are on source, or data on particular baselines are discarded before visibilities are averaged, in which case an overall gain-elevation correction, such as the one shown in Figure 2.2, would not be appropriate. The observations used to derive the gain curves are described in Chapter 5. Each of the three observing sessions lasted 48 hours or more. The seven sources selected to characterise gain-elevation dependence, listed in Table 2.1, show no evidence of IDV, are unresolved, and are at various Right Ascensions and Declinations.

The method used to derive the gain curves was as follows.

- Gain solutions were found for each source and in each epoch using task GPCAL in MIRIAD, after first applying the time-independent polarization leakage corrections from data on PKS 1934−638. Note that gain solutions here refer to the voltage gain correction which needs to be applied to each antenna, $1/g_i$ in the notation used in Equation 2.7. Although the weather conditions were relatively good for the relevant observing sessions, time-dependent effects may lead to increased scatter in the data. However, there is not a straight-forward way to disentangle time-dependent and elevation-dependent effects, and as Figure 2.2 shows, the variation at 8.6 GHz is clearly dominated by elevation-dependent effects. Gain solutions were found for each 10-second integration cycle. Using the shortest possible interval for solutions avoids any decorrelation of the signal due to atmospheric phase instabilities, which may be a problem if vector averaging is performed over a longer time interval. As all the sources are > 1 Jy, there is sufficient signal in 10 seconds for a solution to be found. The gain solutions were found for the X and Y polarization channels separately.

- Scalar box-car averaging over a 1-minute interval was subsequently applied to the gain amplitudes, which were then output to text files using MIRIAD task GPPLT.

- Procedures written in IDL (*Interactive Data Language* by Research Systems, Inc.) were used to fit the elevation dependence. This dependence was observed to be the same in both polarization channels, therefore the X and Y amplitude gains were averaged together for the purposes of deriving gain curves. Curves were fitted to the amplitude gain solutions as a function of zenith angle ($90°−$ elevation) for each antenna, and each source, in each observing session. Quadratic functions were found to fit the data well. The POLY_FIT function in IDL was used to perform a least-square quadratic fit.

- The fitted quadratic functions were each normalised by a constant factor derived separately for each antenna, source and epoch, so that the gain correction was equal to 1 at the zenith. This enabled gain solutions from all sources and epochs to be combined, so that all the data could be fitted simultaneously. In a few cases where the elevation coverage for a particular source was poor, and led to a fit which looked significantly different from the average, that particular dataset was discarded.

- A least-square quadratic fit was then performed on the combined set of gain solutions for each antenna.

Figure 2.4 shows the combined gain solutions from the data and the fitted quadratic function for each antenna. The derived dependence of the antenna gains on elevation appears to be reasonably consistent in different epochs, and for the different sources covering a range of RA and Dec. Table 2.2 shows the quadratic coefficients derived for each antenna. The results suggest that Antenna 6 is poorly focussed compared with the other antennas, probably due to the subreflector not being optimally positioned.
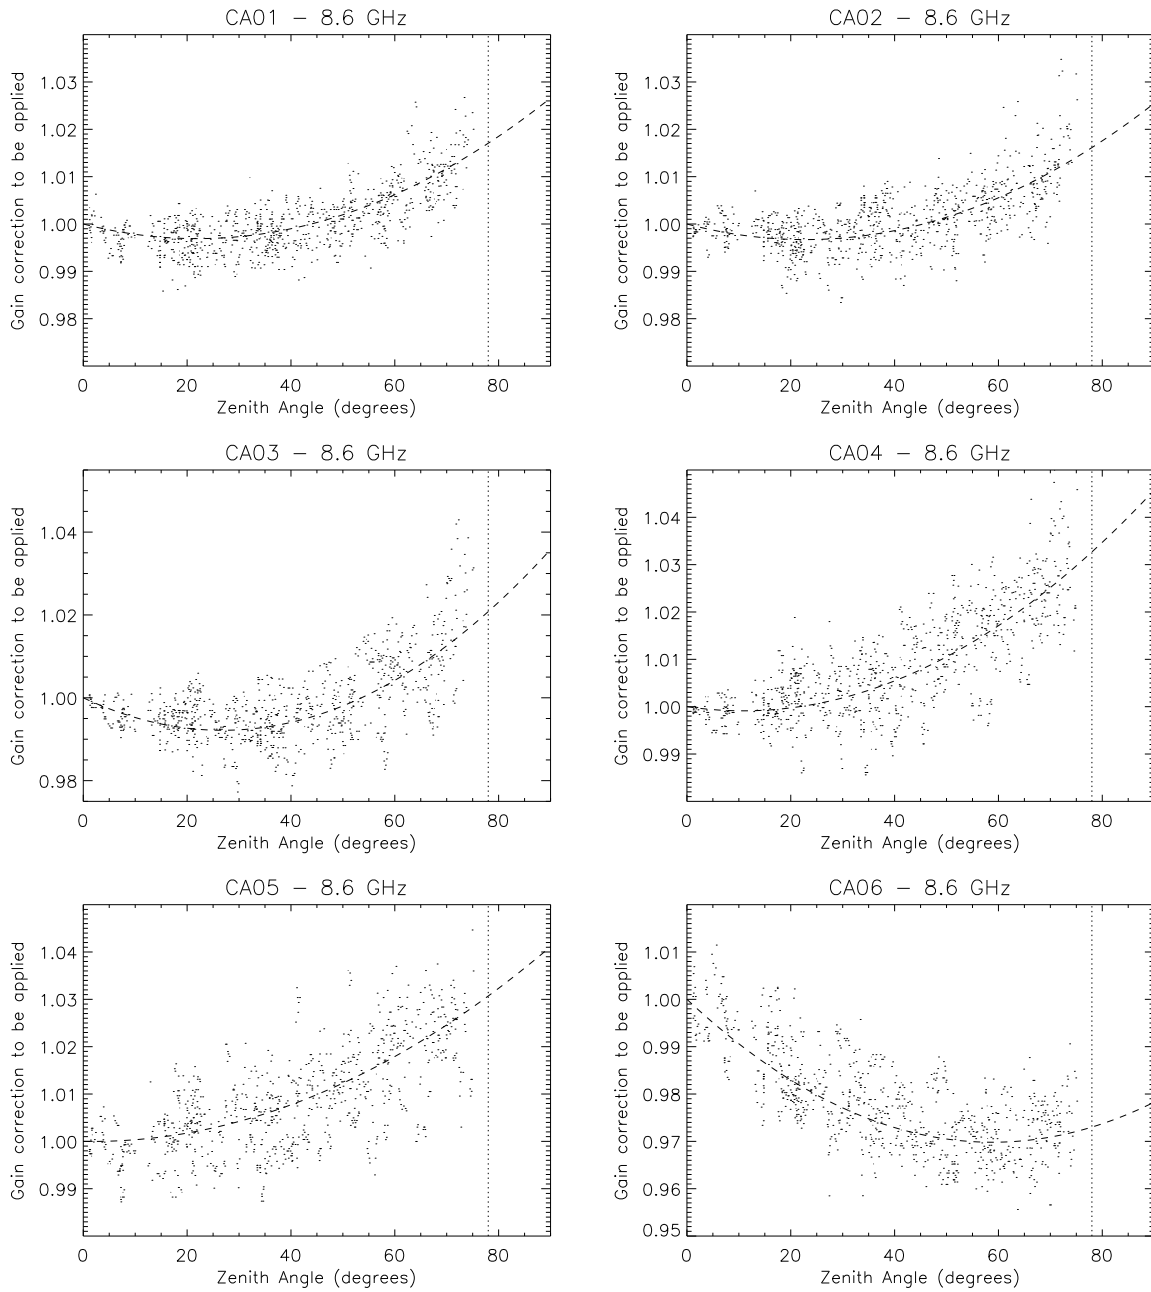
Figure 2.4: Dashed curves show polynomial fits to ATCA antenna gains at 8.6 GHz as a function of zenith angle, using data from a number of different epochs and sources, with measured gain values shown as points. Vertical dotted lines show the elevation limit of the antennas, corresponding to a zenith angle of 78°. See Section 2.3.5 for details.

Table 2.2: ATCA antenna gain polynomial coefficients, fitted as a function of zenith angle (ξ) in degrees, for 3 cm data. For the ith antenna, $1/g_i = a_0 + a_1\xi + a_2\xi^2$.

Antenna	a_0	$a_1(10^{-4})$	$a_2(10^{-6})$
1	1.00	-2.8	6.4
2	1.00	-2.9	6.4
3	1.00	-5.9	10.9
4	1.00	-1.6	7.4
5	1.00	-0.18	5.3
6	1.00	-10.3	8.7

In order to test the effect of applying the derived antenna gain curves on the data, the gain polynomials shown in Table 2.2 were applied to the "raw" calibration source data, initially using IDL routines. In addition to the gain polynomials, which were normalised to the zenith, a single correction factor (for each epoch) was derived for each antenna and applied to the data. Results showed typically a factor of ~ 2 reduction in rms gain variation, after applying the corrections, for several different sources observed over a large range of elevations. For the purposes of detecting low-level IDV, this is a significant improvement.

Subsequently, a MIRIAD task called ELEVCOR was written by D. Rayner (2001) to correct visibilities using the polynomials derived for each antenna. This step is performed before any other calibration. Following the elevation-dependent correction, scaling factors can then be derived for each antenna from observations of an unresolved source, e.g. the primary calibrator, to equalise the visibilities on all baselines. Figure 2.5 shows the data from Figure 2.3, after applying the polynomial corrections. Figure 2.6 shows the effect of the gain-elevation correction on data for PKS 2126−158 observed in 2001 November. These data have been averaged over all baselines, and a constant antenna gain correction factor derived from PKS 1934−638 has been applied in both cases. The only difference between the two datasets is that for the lower plot, the correction from ELEVCOR has been applied to the data. Clearly the systematic variation has been substantially reduced, although not completely removed, showing that the correction is appropriate for data collected more than 6 months after those from which the gain-elevation correction was derived.

2.3.6 Summary

A discussion of the use of the ATCA for compact source flux density measurements has been presented. Although the ATCA is primarily a synthesis imaging telescope, it can be used to measure compact source flux densities and polarization parameters with high efficiency and accuracy. Polarimetric calibration, as described in Section 2.3.3, is standard for ATCA data and is implemented in the MIRIAD software package. The elevation dependence of antenna amplitude gains had not previously been characterised at the ATCA for cm-wavelength observations, but was found to be a significant effect in the 3 cm band. The effect is mainly important for observations of intraday variabil-
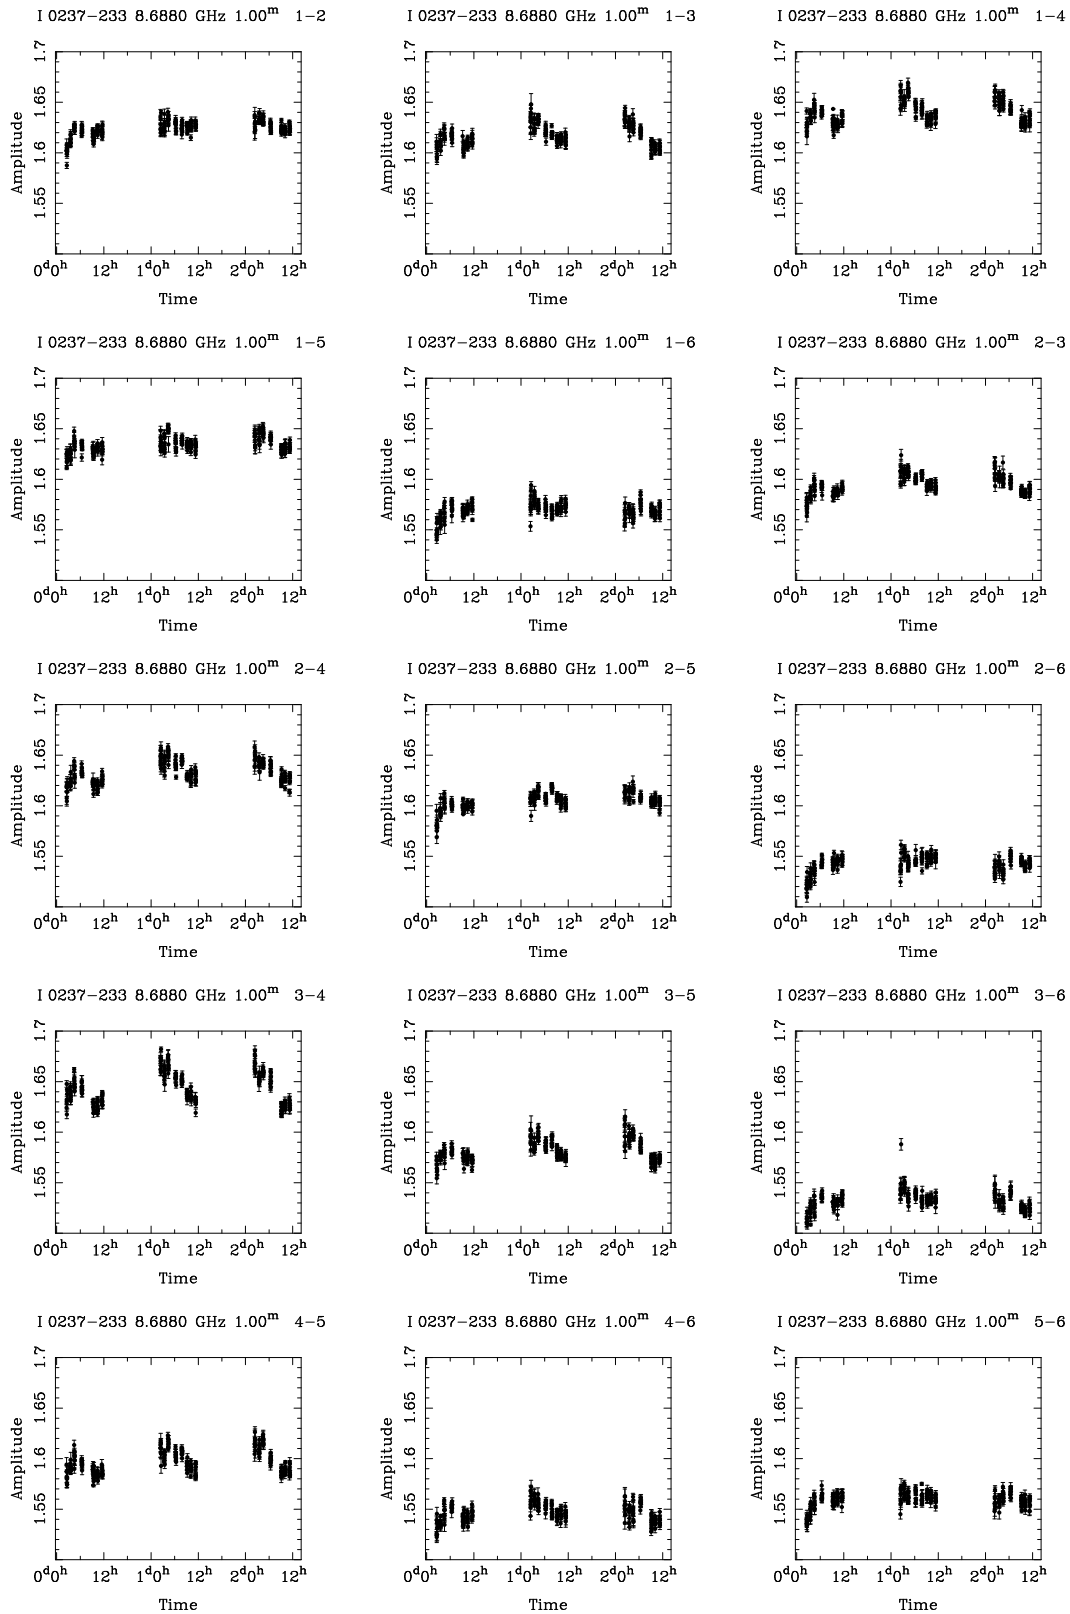
Figure 2.5: The data shown in Figure 2.3, with elevation corrections applied. Much of the systematic elevation-dependent variation has been removed. Note that a constant correction factor for each antenna has not yet been applied to these data.
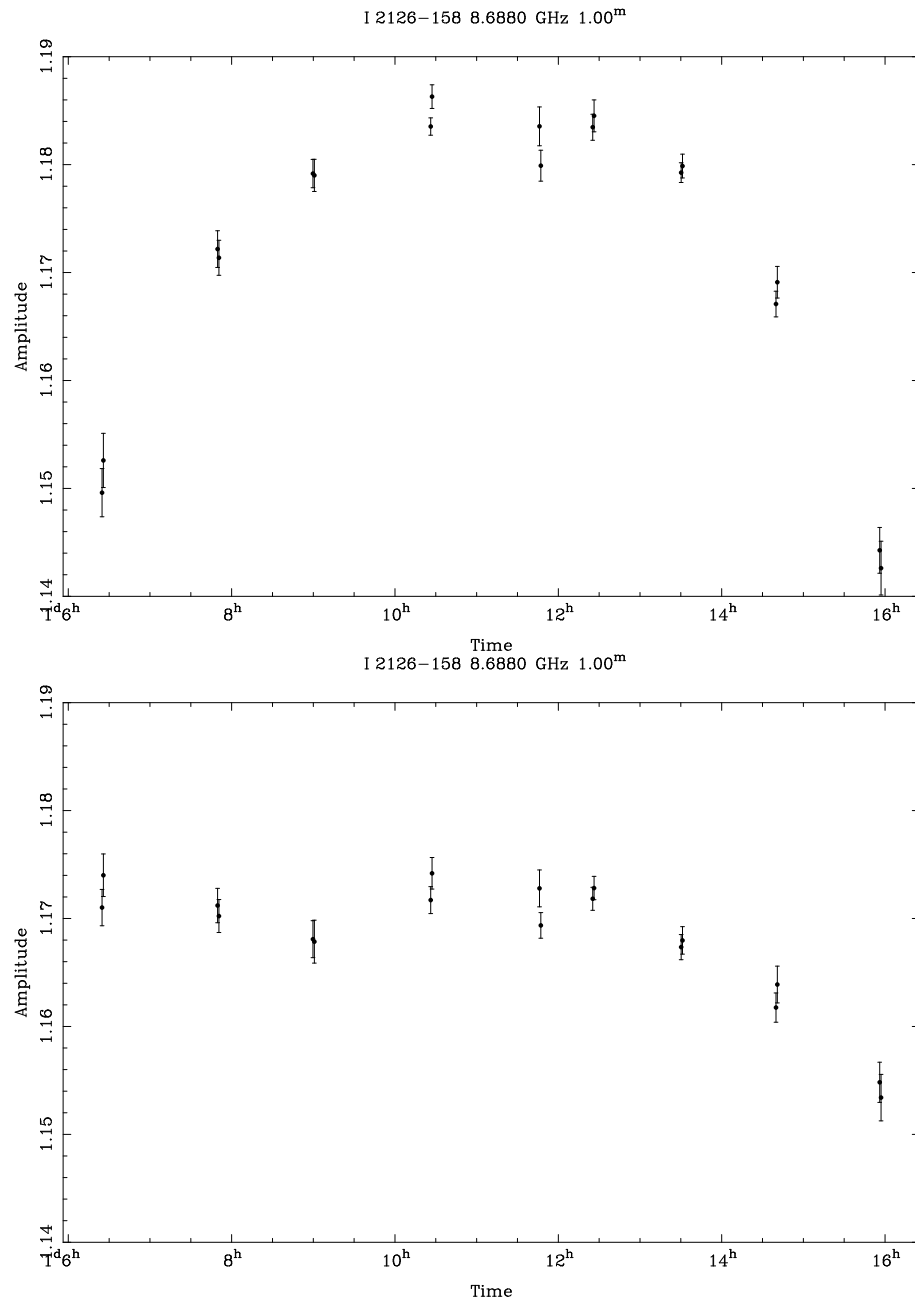
Figure 2.6: Data on a calibration source observed in 2001 November, averaged over all baselines, before (above) and after (below) applying elevation-dependent correction derived from earlier data. Error bars show the rms scatter in the mean; data are averaged over 1 minute integrations. In each case, a time-independent correction factor, derived from PKS 1934−638, has also been applied to the gains of each antenna.

ity, presented in Chapter 5 of this thesis. Correcting flux density measurements for elevation-dependent effects has therefore been investigated here, and incorporated into the calibration procedure for data presented in Chapter 5.

2.4 The Australia Telescope Long Baseline Array

Very Long Baseline Interferometry (VLBI) is an aperture synthesis technique used to achieve high angular resolution. In a VLBI array, antennas are not electrically connected, hence they can be separated by large distances; each station has an independent local oscillator. The Australia Telescope Long Baseline Array (AT-LBA or just LBA) is a set of antennas in the Southern Hemisphere, used for VLBI. It is an *ad hoc* array rather than a dedicated instrument. LBA observations are typically scheduled for about one week during each four-month observing term. Locations and parameters of the antennas used for VLBI observations presented in this thesis are shown in Table 2.3 (this and other information on the LBA is available at the URL `http://www.atnf.csiro.au/vlbi/`). The baseline lengths, and corresponding FWHM of the synthesised beam obtained at the frequencies observed for this thesis, 2.3 and 8.4 GHz, are shown in Table 2.4.

The facilities at Parkes, Mopra and the ATCA are operated by the ATNF. The Hobart and Ceduna antennas are operated by the University of Tasmania. DSS43 and DSS45 are located at the Canberra Deep Space Communication Complex at Tidbinbilla, which is part of the National Aeronautical and Space Agency's (NASA) Deep Space Network. As part of the Host Country agreement with NASA, a fraction of time on these antennas is allocated for independent scientific activities sponsored by the Australian Government, so one of the Tidbinbilla (Tid.) antennas is sometimes available for VLBI observations. The 26 m antenna at Hartebeesthoek (Hart.) Radio Astronomy Observatory, South Africa, is also sometimes used in conjunction with the LBA, for higher resolution observations.

Since the array elements are not connected, data are recorded on video tapes along with high-precision time stamps, and correlated at a later time. To keep time very precisely, and to provide very high phase stability, each station uses a Hydrogen maser frequency standard. The procedures involved in obtaining calibrated visibilities from VLBI observations are described below in Sections 2.4.1, 2.4.2 and 2.4.3.

The ATCA tied array

For VLBI, the ATCA operates in "tied array" mode, where the signals from each antenna are added together before being recorded on tape, rather than multiplied, as in normal interferometric mode. Using the tied array rather than a single, 22 m antenna provides greater sensitivity as the collecting area is increased, e.g. including 5 of the ATCA antennas in the tied array provides a collecting area approximately equivalent to a single dish 50 m in diameter.

"Tying" antennas together requires the individual antenna phases to track each

Table 2.3: Location and parameters of Southern Hemisphere VLBI antennas used in blazar monitoring. Location is described by geodetic latitude and longitude. Altitude (Alt.) corresponds to approximate height above sea level. Parameters shown are antenna diameter, type of mount, elevation limit, and approximate *System Equivalent Flux Density* (SEFD) at the frequencies observed for this thesis.

| Station | Location | | | Parameters | | | | |
| | Long. E (°) | Lat. (°) | Alt. (m) | Diam. (m) | Mount | El lim (°) | SEFD (Jy) | |
							2.3 GHz	8.4 GHz
Narrabri	149.57	−30.31	217	$n \times 22^a$	AZEL	12	$400/n$	$400/n$
Mopra	149.07	−31.30	840	22	AZEL	12	350	400
Parkes	148.26	−33.00	392	64	AZEL	30	90	90
DSS43	148.98	−35.40	670	70	AZEL	6	15	20
DSS45^b				34	AZEL	8	165	130
Hobart	147.44	−42.80	100	26	XYEW	$\sim 16^c$	650	750
Ceduna	133.81	−31.87	160	30	AZEL	~ 10	600	600
Hart	27.68	−25.89	1416	26	HADEC	$\sim 10^c$	390	1100

$^a n \equiv$ number of antennas in tied array
^{b} DSS45 is 548 m from DSS43.
^{c} Elevation limit is position-dependent.

other, as the signals are added vectorially. If the phases differ by more than $\sim 30°$, significant decorrelation occurs, resulting in reduced amplitudes at the output. Phases commonly vary with time, particularly at high frequencies and in poor weather conditions, which can lead to time-variable amplitude errors in the VLBI visibilities. Antenna 6 was not included in the tied array as the phases on the longer baselines are much more variable. It is possible to continuously calibrate phases, using the ATCA on-line CACAL procedure, on strong, unresolved sources. Using one antenna as a reference, phase corrections are applied approximately every minute when a designated "calibrator" source is being observed, to keep the visibility phases close to zero. CACAL is also used to insert a 90° phase offset between the X and Y polarization outputs of each antenna, which forms circular polarization at the tied array output. Since the sources observed for this thesis are generally strong and unresolved on ATCA baselines, the option of continuous calibration was used in some observing sessions to keep the tied antenna phases within a few degrees of each other, and this generally works well provided the phases do not vary too rapidly. In conditions of very unstable phase, antennas can be selectively removed from the tied array, which reduces the collecting area but removes the effect of decorrelation of the signal.

The ATCA correlator is generally used simultaneously with the tied array, so the usual interferometric ATCA data are obtained, which are extremely useful for calibrating the LBA data (see Section 2.4.3).

Table 2.4: Length of baselines, in km, between stations in Table 2.3 are shown in the top right-hand section of the table. The lower left-hand half of the table shows corresponding half-power beam-width, in milliarcseconds, at 2.3 and 8.4 GHz.

	Narrabri		Mopra		Parkes		Tid.		Hobart		Ceduna		Hart.
Nar	**2.3**	**8.4**	113		321		566		1396		1504		9853
Mop	240	65	**2.3**	**8.4**	207		459		1287		1449		9791
Pks	84	23	130	35	**2.3**	**8.4**	275		1089		1361		9673
Tid	47	13	58	16	98	27	**2.3**	**8.4**	832		1456		9595
Hob	19	5	21	6	25	7	32	9	**2.3**	**8.4**	1702		9175
Ced	18	5	19	5	20	5	18	5	16	4	**2.3**	**8.4**	8955
Har	3	0.7	3	0.7	3	0.8	3	0.8	3	0.8	3	0.8	-

2.4.1 Recording and correlation of LBA data

At each station, data are recorded using the Canadian-designed S2 recording system (Wietfeldt et al., 1991), at a maximum of $128 \, \mathrm{Mb \, s^{-1}}$. On the completion of each observing session, tapes are transported to the LBA correlator at ATNF headquarters in Marsfield, NSW. As with a connected element interferometer such as the ATCA, the VLBI correlator cross-multiplies signals from each pair of antennas, that correspond to the same arriving wavefront. It is therefore necessary to know the delay between the wavefront arrival times at each antenna, and the speed at which each antenna is moving along the direction to the source. This motion causes the signal to be Doppler shifted by some amount, known as the "fringe rate". These parameters are generally more difficult to determine in VLBI compared with connected element interferometry, because (i) the lines of sight from each antenna pass through totally uncorrelated atmospheres, (ii) many geophysical effects which are insignificant in connected-element interferometry become important for VLBI, and (iii) the clocks at each station operate independently of one another. Most important geophysical effects are now well-modelled and are taken into account at the correlator. Hydrogen maser frequency standards, which are now used at most VLBI stations, are highly stable. The atmosphere is in general more variable and more difficult to model.

At the LBA correlator, the first step in correlation is to "fringe-search" in order to make the clock models for each station as accurate as possible. This step involves searching for the delay corresponding to the peak of the measured cross correlation, using data on a strong, compact source. The delay is generally found at two different times in each 24 hour experiment, so that any drift of the clocks can be determined. For the experiments presented here, clock drift rates, if detectable, were found to be on the order of 10^{-6} microseconds per second, and delays could generally be determined at the correlator to an accuracy of a few nanoseconds.

Once the clock models are determined, the correlation is performed, choosing a suitable configuration for the number of baselines and recorded bandwidth. The data are output in the form of RPFITS files (Norris et al., 1992), the same format produced by the ATCA correlator. For the data presented here, either two- or five-second

integration cycles were used for the output data. These data are then loaded into the National Radio Astronomy Observatory's (NRAO) *Astronomical Image Processing System* (AIPS) software package for further processing.

2.4.2 Fringe fitting

Fringe fitting algorithms (e.g. Walker, Lecture 22 of TCP99) are designed to remove any residual delay and fringe rate from the data, so that the data can be averaged in time without resulting in reduced amplitudes due to phase variations. Since the delay on the i - j baseline, τ_{ij}, is related to the baseline phase, ϕ_{ij}, by:

$$\phi_{ij} = 2\pi\nu\tau_{ij}, \tag{2.17}$$

where ν is the observing frequency, delay errors can lead to large phase offsets. For the LBA data, a global fringe fit was performed as the first step after loading the data into AIPS, using task FRING. FRING was run for data on all sources, assuming a point source model, which generally worked well since all sources observed for this project are dominated by a strong, compact core. In all cases, a solution interval of approximately 3 minutes was found to be adequate. Under good conditions, the fitted delays were typically a few nanoseconds, with fringe rates of no more than a few mHz. Where the signal-to-noise was sufficient to find a good solution, the rms errors for each antenna over a solution interval were typically 0.1 ns in delay, 0.01 mHz in rate, and $< 1°$ in phase. Information on the absolute phase of the visibility function is lost in the fringe fitting process, since phase on a designated reference antenna is set to zero. Since finding absolute source positions was not an aim of this project, the absolute phase was of no concern. Solutions from FRING were then written to AIPS gain tables using task CLCAL.

2.4.3 Amplitude calibration of LBA data

The correlated flux density on the i - j baseline, S_{Cij}, is related to the correlation coefficient ρ_{ij} by

$$S_{Cij} = b\rho_{ij}\sqrt{\frac{T_{Si}}{K_i}\frac{T_{Sj}}{K_j}}, \tag{2.18}$$

where the T_S/K terms represent *system equivalent flux density*. T_S is the antenna temperature, and K is defined as the gain of the antenna in degrees Kelvin of antenna temperature per Jansky of flux density. b is a correlator dependent constant which for the LBA correlator is close to unity.

There were effectively two steps involved in amplitude calibration of the LBA data. The first step is a relative calibration, to find internally consistent gain solutions. The second step is to tie the calibration to a flux density scale; in this case the ATCA flux density scale, which is set by observations of the primary calibrator PKS 1934−638 (Reynolds, 1994) and tied to the widely-used absolute flux density scale of Baars et al. (1977).

Log files containing measured values for system temperature, T_s, are recorded at each station. On the completion of an observing session, these are collated and converted into a format suitable to be read by the AIPS task ANTAB. Also included in the ANTAB files are antenna gain curves. These were mostly not well determined for LBA antennas before the present project began, but were derived during the course of this work as part of the development of standard calibration procedures for the LBA. After running ANTAB, the task APCAL is used to produce a "solution" (SN) table of amplitude gains to apply to the data in AIPS. After applying this *a priori* calibration, data are inspected and a correction factor is derived for each LBA antenna from data on one or more calibration sources, assumed to be unresolved, so that the amplitudes are equal on all baselines. This is equivalent to using the *amplitude closure relations* (see Section 2.2.4). Following this step, the calibration is internally consistent.

For all LBA observations presented here, total source flux densities were obtained simultaneously with the ATCA. By observing a calibration source which is unresolved both on ATCA and LBA baselines, it is possible to calibrate the LBA visibility amplitudes against the ATCA flux density scale, to an accuracy of typically $5-10\%$, depending on the quality of the LBA data. Thus, the visibilities measured with the LBA for all sources can be compared with the total flux density as measured with the ATCA. This provides information on the source brightness distribution between arcsecond and milliarcsecond scales. Calibrating amplitude gains for the LBA data as accurately as possible was therefore an aim of this project.

The *a priori* calibration was found to be usually correct within $\sim 20\%$. On the intercontinental baselines to South Africa, no known calibration sources are completely unresolved, so the *a priori* calibration was assumed to be acceptable. Parameters for the Hartebeesthoek antenna are relatively well-determined.

2.4.4　Further processing of LBA data

After correlating, fringe-fitting and amplitude calibration, VLBI data are handled in just the same way as other interferometric data, such as those from the ATCA. In this case the data were split into individual source files and exported as FITS format files for editing and further analysis, including imaging and/or model-fitting, using standard procedures. The Difmap software package (Shepherd, 1997) was used for all further processing of the VLBI data. At this stage the phase part of the antenna gains has not yet been properly determined; phase self-calibration in Difmap is used to make these corrections.

2.4.5　Summary

This section presented an overview of the LBA and the steps involved in recording, correlation and calibration of LBA data. The techniques involved are standard for any VLBI, with some peculiarities specific to the LBA. The array itself is quite inhomogeneous, unlike for example the VLBA, and the differences between individual antennas must be taken into account for calibration purposes. Some development of

amplitude calibration procedures was undertaken during the course of this thesis. Post-correlation, the data can be processed using standard software packages such as AIPS, Difmap or MIRIAD.