Calibration of the Australia Telescope

R.P.Norris
22nd February 1988

This document proposes a system of calibrating data from the AT. It is provisional, in that it is designed to invite comments rather than propose a final system. In particular, AIPS is undergoing major changes at the moment, as a result of which some of the ideas here may need re-examination. Furthermore, there will inevitably be calibration parameters that we haven't yet thought of.

1) INTRODUCTION

1.1) The scope of this document

This document details a scheme for routine calibration of data from the AT compact array. It is not concerned with the initial fringe-searches and debugging of the instrument, although parts of the system described here may well be used during that process, but instead looks at how we should aim to calibrate the data once the AT is running more or less routinely.

An earlier document (AT/25.1.1/047) reviewed the more general question of off-line software. By contrast, this document is concerned not only with software, but also hardware and operational procedures. It should also be noted that this document concentrates on the calibration of data from the Compact Array (CA). This is because calibration of LBA data will largely follow existing VLBI (and future VLBA) practices, and so detailed planning at this stage is probably not warranted.

This document does not consider polarisation calibration explicitly, since that has already been considered in great detail (AT/10.1/039 and AT/01.17/017) and incorporated into the on-line array processor software. The 'raw' data produced by the AT will therefore already have been converted to Stokes parameters, and the only further interest in polarisation calibration information will be for the purposes of re-calibration within AIPS.

1.2) Philosophy

It was originally intended that the u-v data from the AT should be corrected, using the mid-observation calibration data described below, before recording. However, since that decision was made AIPS (which will form the basis of our analysis and calibration package) has undergone considerable changes (and is still changing), as a result of which I now propose that the calibration information should not be applied directly to the data, but should instead be recorded along with the data in the form of (a) observations of calibrator sources, recorded (as FITS multi-source uv data) in exactly the same way as observations of target sources, and (b) FITS extension tables appended to the data. The first stage of reading the data into AIPS will then be to apply the calibration to the data.

This change has the advantage that the data will look very similar to data produced by the VLA, and hence can be processed by those standard calibration routines which are now being implemented within AIPS. Since these routines are standard, it also means that we retain the option that any astronomer can take AT data back to his home institution and process it there.

It also has the minor advantage that, because all the calibration data is recorded along with the u-v data, the Convex computer used for data reduction (previously called the Data Reduction Computer, or DRC) need access the AT database only when an astronomer wishes to use updated calibration information, rather than that available at the time of the observations.

It should be noted that a certain amount of on-line 'calibration' is intrinsic to the operation of a tracking interferometer, in that the algorithms that calculate the phase and delay rates to be applied inevitably contain some model of the interferometer and the source. It seems sensible to make this model as accurate as possible, so that, in effect, the 'raw' data from the AT will already have been corrected for such variables as earth rotation, weather, antenna positions, etc.

Note that the scheme proposed here retains the philosophy that any calibration applied can always be undone and re-done. Where necessary, this will be done by inserting COMMENT records into the FITS headers which will point at entries in the AT database.

1.3) Overview of the Calibration Process

There are three main stages in the determination of calibration parameters, which are illustrated in Figure 1:

1.3.1) Array Calibration
(previously called primary calibration)

First, after an array change, positions and other parameters of the telescopes need to be determined. This will be achieved by a series of observations of calibrator sources scattered around the sky, taking perhaps 6 - 12 hours of observing time. Antenna pointing is not included here as it is considered to be a separate process. As a result of these observations, the location and orientation of the antennas and stations will be determined (by the program CALFIT which has not yet been written), entered into the AT database, and subsequently used for on-line phase and delay calculations and written into the antenna table appended to the data.

1.3.2) Mid-Observation Calibration
(previously called secondary calibration)

During the observation of a target source, a number of calibrator sources close to (within a few degrees of) the target will be observed at intervals of perhaps ten minutes or so. The aim of this process is to measure the short-term changes in the complex gain of the antennas, to measure the complex bandpass response, and to determine random (i.e. those not accounted for by meteorological data) atmospheric, ionospheric, and residual effects. The data from these observations will be recorded along with the data in FITS format, and can subsequently be used to write calibration tables within AIPS. No processing of this data will occur on-line, except that required for on-line monitoring.

1.3.3) Non-astronomical Calibration

In addition to the two types of calibration discussed above, there will be a quantity of data available from literature and instrumental measurements. These will all be entered into the AT database, used for on-line calculation of phase and delay, and subsequently written into tables appended to the data. These types of calibration data include:

(a) Instrumental data - This includes the complex gains, cross-polarisation, T_{sys}, etc. of the various receivers and other bits of hardware, as determined by the one-bit correlators and phase-switched detectors included specifically for this purpose.These will normally be measured continuously. The interval between samples to be recorded will be determined by experience - a working hypothesis would be to assume they will be sampled every minute or so.

(b) Ephemeris and geophysical data - This includes all the data obtained from the ephemeris, USNO circulars, etc., and is updated periodically by reference to the literature.

(c) Meteorological data - This includes data from the local meteorological instruments, and possibly ionosonde or satellite Faraday-rotation data, and is updated on a time scale of minutes-hours. It is not yet known whether Rdb will cope with the volume of meteorological data. At present it is assumed that it will, but that these data may need to be purged at (perhaps daily) intervals. Alternatively, the meteorological data could be stored in a simple sequential file.

Once these parameters have been determined, they are used to correct the data in a number of separate places:

(a) Implicitly, during the phase tracking process, the data are corrected for array and instrumental calibration parameters.

(b) During observing, the data will be corrected on-line for the polarisation parameters, to produce the data in terms of Stokes parameters.

(c) Subsequent correction for information from the Mid-observation Calibration process will take place within AIPS.

(d) Re-calibration for different values of array or instrumental parameters may be performed by a new (and unwritten!) task in AIPS called RECAL. This will use the tables, together with comment cards in the data header, to establish which parameters were used on-line or in subsequent calibration, and will access the AT database to correct the data for the difference between the new and the old parameters.

(e) Finally, self-calibration may be used in most cases within AIPS.

1.4) The AT Database

Much of the data determined from measurements and the literature will be contained in a relational database (the DEC product Rdb) kept on one of the VAX 8250's at Culgoora. A 'slave' version of this database will also be maintained at Epping. A first version of this database has already been designed and set up by Mark Calabretta, and it is on his version of the database that most of the following references to 'the AT database' are founded.

1.5 The Signal Down-Conversion Chain

The route by which the local oscillators are mixed with the signal to convert it down to video is a tortuous one, involving many filters and phase shifts. There is a potential problem in that the route (and choice of sideband within the down-conversion) depends critically on the observing frequency, so that, in the worst case, the usual changing doppler correction (for l.s.r. etc) can shift the observing frequency across a boundary, resulting in a sudden reconfiguration of the down-conversion chain. The result will be an arbitrary phase and amplitude change, possibly together with a change in the bandpass response. Clearly, this must be avoided as far as possible within an observation.

I raise the problem here to draw attention to it, so that the LO system may have enough overlap built in to it to prevent this problem, and the observing software may be written to take advantage of this overlap to maintain the same LO configuration for the length of an observation. In any cases where this is not possible, then the scheduling program should issue a warning , so that additional calibration may be incorporated into the schedule.

2) Array Calibration

The aim of array calibration is to determine the precise location and orientation of each of the antennas after they have been moved to a different station on the track, and to measure the transmission delay back to the central site. This is achieved most efficiently by a series of observations over 6 - 12 hours in which a large number of calibrator sources distributed over the sky are observed for a few minutes each. The resulting data is then analysed by an (as yet unwritten) program called CALFIT, which attempts to fit the observed phase to a model which is described by a large number of free parameters. In practice, some of these parameters will already be well-known, and so CALFIT must allow the user to label each parameter as free, fixed, or constrained. Typically, hundreds of phases (and their time derivatives) will be measured on each baseline, and about ten parameters per antenna will be determined. To do so, however, it is of course necessary to have beforehand a framework of 20 or 30 calibrators with accurately known positions.

To a first order, the position of each antenna is known in terms of a sum of two components:

(a) the nominal position of each station in array-centred Cartesian coordinates. This is contained in the relation CA_STATIONS (containing 6 parameters) on the AT database.

(b) the nominal position of the antenna phase centre relative to its fiducial reference point. This is dependent on the position of the feed-horn and subreflector, and is contained in the relations AT_ANTENNAS (11 parameters), AT_FEED_HORNS (3 parameters), and SUBREFLECTOR (6 parameters).

Similarly, the transmission delay is approximated by the sum of two components: the delay in the fibre-optic cable from the sampler to the Stratos connectors (stored as 4 parameters in AT_ANTENNAS) and the delay from the station post to the central site (stored as 4 parameters in CA_STATIONS).

The data contained in the four relations named above would initially be determined from engineering data, and subsequently refined in the early days of AT operation by a series of calibration measurements which allowed them to vary as free parameters. Subsequently, they would not normally be updated very often, unless an analysis of error terms over an extended period revealed a systematic offset which could be traced to one of these data.

After moving an antenna, its position on the station is measured optically, and placed in the relation CA_CONNECTION. There then remain only the following:

(a) the residual effects which depend on the way it is placed on the station. These are contained in six parameters per antenna to be determined by CALFIT, and which are contained in the relation AT_CALIBRATION.

(b) the transmission delay from the central site to the antenna. This is contained in one parameter per IF in the relation PHASETC, which is determined by examining the phase slope across the band when observing a calibrator. Note that this is a coarse bulk instrumental delay: fine variations are also monitored via the mid-observation calibration. Note also that this term includes both the physical delay and also any offset in the sampler clock; the two effects are observationally indistinguishable.

The parameters derived from array calibration are subsequently used in the calculation of phase and delay rates to be applied to the data during observing, and it is therefore at that stage that the data are effectively corrected for these effects. In addition, the antenna coordinates are also written into the antenna table for subsequent use.

3) Mid-observation Calibration

The scheme of mid-observation calibration proposed here is based on that used by the VLA. During (say) a 12-hour observation of a target source, a nearby calibrator source will be observed for (say) a minute every ten minutes. The actual duration and frequency of these observations will depend on observing frequency, weather conditions, and the source characteristics. During the observation, no special attention need be paid to these data as the calibration from them will be done by tasks such as CALIB within AIPS.

This calibration is used to measure the antenna complex gain, clock error, bandpass calibration, and baseline-dependent errors.

3.1) Antenna complex gain

The phase and amplitude response of each antenna will depend upon:

a) Atmospheric and ionospheric delays

b) Receiver phase drifts

c) The arbitrary phase offset of the local oscillator (this is frequency dependent, and must be measured separately for each LO configuration).

All three of these effects can be lumped into the antenna complex gain.

3.2) Antenna clock error

This error includes (a) an error in the assumed transmission delay from the central site to the antenna, and (b) any offset in the sampler clock. It is hoped that experience will find this error to be negligible, but it must be monitored until we have that experience. The effect of this error will be seen in two equivalent ways: (a) a phase gradient across the bandpass, and (b) a shift in Right Ascension of calibrator sources. Gross errors will be detected most effectively by (a), but (b) means that normal mid-observation calibration will remove the effects of small errors, and so this effect (at least for small errors) can also be lumped into the antenna complex gain.

3.3) Bandpass Calibration

A spectral line observation clearly needs correcting with a bandpass. Less obviously, continuum observations also need this correction to avoid 'closure offsets', which reduce the effectiveness of self-calibration. This complex bandpass response can be determined from short observations of strong calibrator sources. However, at present it is not clear how often such a calibration needs to be undertaken. At best, filters may be so stable that they need to be calibrated only at very infrequent intervals, and then standard tables used to correct all observations. At worst, filter responses may change on hourly timescales, so that bandpasses have to be calibrated from the mid-observation calibrations (described below). Even worse, the responses may depend not only on the bandpass-defining filters but also on other filters in the down-conversion chain, so that, in the very worst case, the response may depend on the l.s.r. velocity correction (see section 1.5 above).

For the present, we will assume that bandpasses need to be calibrated for each observation, but that they remain stable during an observation. Thus it is necessary merely to observe unresolved calibrator sources in the normal way (but remembering to observe in each LO configuration that will be used during the observation), and then derive bandpass (BP) tables from the calibration observations within AIPS.

3.4) Baseline-dependent errors

Baseline-dependent errors are generally errors produced by the correlator, but can also be produced by other effects such as bandpass mis-calibration. It is expected that AIPS tasks will be available to determine baseline-dependent errors (stored as a BL table) from the mid-observation calibration.

4) Non-Astronomical Calibration

4.1) Instrumental data

A number of instrumental parameters are monitored continuously by the hardware of the AT, and their results placed into the AT database, within the relation INSTRUMENT . These are described in AT/20.1.1/020 and AT/20.1.1/022 and include:

(a) PSD Gain and level monitors (2 per IF per antenna)

These devices measure the receiver gain and system temperature. They are used on-line by the array processor software, but should also be recorded for subsequent calibration. The output (derived from a dataset) should be averaged by the on-line software over a period of (say) 1 minute and then written into the database. They will change slowly over a time-scale of minutes and also at each frequency or receiver change. If these are to change rapidly (e.g. frequency switching for bandwidth synthesis) then it would be advantageous to have separate fields within the relation for each frequency used. Even so, the quantity of data might overwhelm the database, in which case this data might be stored better in a VMS file.

(b) Polarisation correlators ('the Komesaroff Machine': 2 per antenna)

These devices correlate each polarised output with the output in the orthogonal polarisation, to produce one number per feed at any frequency. This number is expected to be stable but will change with frequency, so that at any one time (when two frequencies are being observed simultaneously) we may expect two numbers for the four IF's. This number will again need to be stored as a number of separate fields in the case of bandwidth synthesis. All the polarisation information is used on-line in the array processor to produce Stokes parameters, but is included here for completeness.

(c) Culgoora time offset

This measures the difference between the time of the clocks at Culgoora (CAT) and International Atomic Time (TAI). This will be monitored daily (?) by GPS or other receivers, entered in the database, and used for on-line calculations. The required accuracy of this process has not yet been determined.

4.2) Ephemeris and Geophysical Data

These include the relation EOP_PREDICTION, which contains the prediction of polar motion and UT1-UTC . These will be obtained from the weekly USNO circulars by staff at Culgoora and entered into the database by hand. They also include the relation GEOPHYSICS, which contains the parameters for earth tides, ocean loading, and (for completeness only!) continental drift. These numbers will change only when a better model becomes available.

These data will be stored in the database and used for on-line calculations.

4.3) Meteorological Data

These include the relation IONOSPHERE, which will contain electron column density and mean magnetic induction field. The source of these parameters has not yet been determined. They will also include the relation WEATHER, which will contain the temperature, pressure, and humidity measured on the site. It may also include in the future a water vapour radiometer delay.

These data will be stored in the database and used for on-line calculations.

5) Calibration data to be output with uv data

At the end of each observation, an on-line task will write a series of FITS tables (such as the AN table) which will be a condensation of all the necessary information taken from the database. In addition, COMMENT records in the header will contain details of which entries from the various database relations were used for the on-line calculations, and any subsequent data correction. This will enable later re-calibration of the data.

Four tables are relevant to calibration: the antenna (AN) table, the baseline (BL) table, the bandpass (BP) table, and the calibration (CL) table. The CL, BL, and BP tables will be constructed within AIPS from the mid-observation calibration observations. The antenna table, on the other hand, must be appended to the data and contains:

(a) The position of the phase centre of each antenna, obtained from the sum of terms contained in relations CA_STATIONS, AT_ANTENNAS, AT_FEED_HORNS, and CA_CONNECTION.

(b) geophysical and timing parameters, from relations EOP_PREDICTION and GEOPHYSICS

(c) polarisation parameters, from AT_FEED_HORNS and the Komesaroff data.

6) Calibrator Sources

We require for both array and mid-observation calibration a network of unresolved sources spread uniformly throughout the sky. For array calibration, these sources must have positions that are known to a small fraction of a beamwidth. This latter requirement can be relaxed for mid-observation calibration provided that the same calibrator is used throughout the observations, and that we do not need to know the absolute position of the target. If we wish to use more than one calibrator, or if we wish to know the absolute position of the source, then the calibrator must have an accurately known position.

The AT will operate initially at wavelengths no shorter than 3 cm, giving a maximum resolution of 1 arcsec, but will operate in the future at wavelengths as short as 2.6 mm.We therefore need a framework of unresolved sources with 0.1 arcsec positions throughout the sky initially, and ultimately 0.01 arcsec positions.

Declinations north of -45° are adequately served by VLA and northern VLBI calibrator catalogues. South of -45°, there are at present about 20 sources known to a positional accuracy of about 0.008 arcsec (D.L.Jauncey, B.Harvey, et al.) with a further 30 with 0.3 arcsec positions from Molonglo (R.W.Hunstead et al.) and a further 100 or so with 0.3-1.0 arcsec positions from VLBI (Morabito et al.). No other southern sources have better than arcsec positions.

Those 20 sources with milliarcsec accuracy may be sufficient for array calibration at all AT wavelengths, although additional ones would be useful. By contrast, about 100 calibrators would be necessary to ensure that most target sources are within ten degrees of a calibrator. Both PTI (using phase referencing) and VLBI (using fringe rate/delay) are capable of measuring positions to the required accuracy, although recent MkIII VLBI results by B. Harvey et al. indicate that VLBI is probably the most efficient (in terms of manpower) of achieving the goal of 100 calibrators. On the other hand, the PTI is doubtlessly the most efficient way of identifying suitable sources. Such a program is already under way (R.A.Duncan, G.L.White, et al.), and can be expected to yield a suitable list by the end of 1988.

The subsequent VLBI measurements will take approximately 8 * 24-hour VLBI sessions, and about one man-year to analyse the data. The goal of 100 calibrators is therefore probably attainable by early 1990, given suitable manpower.

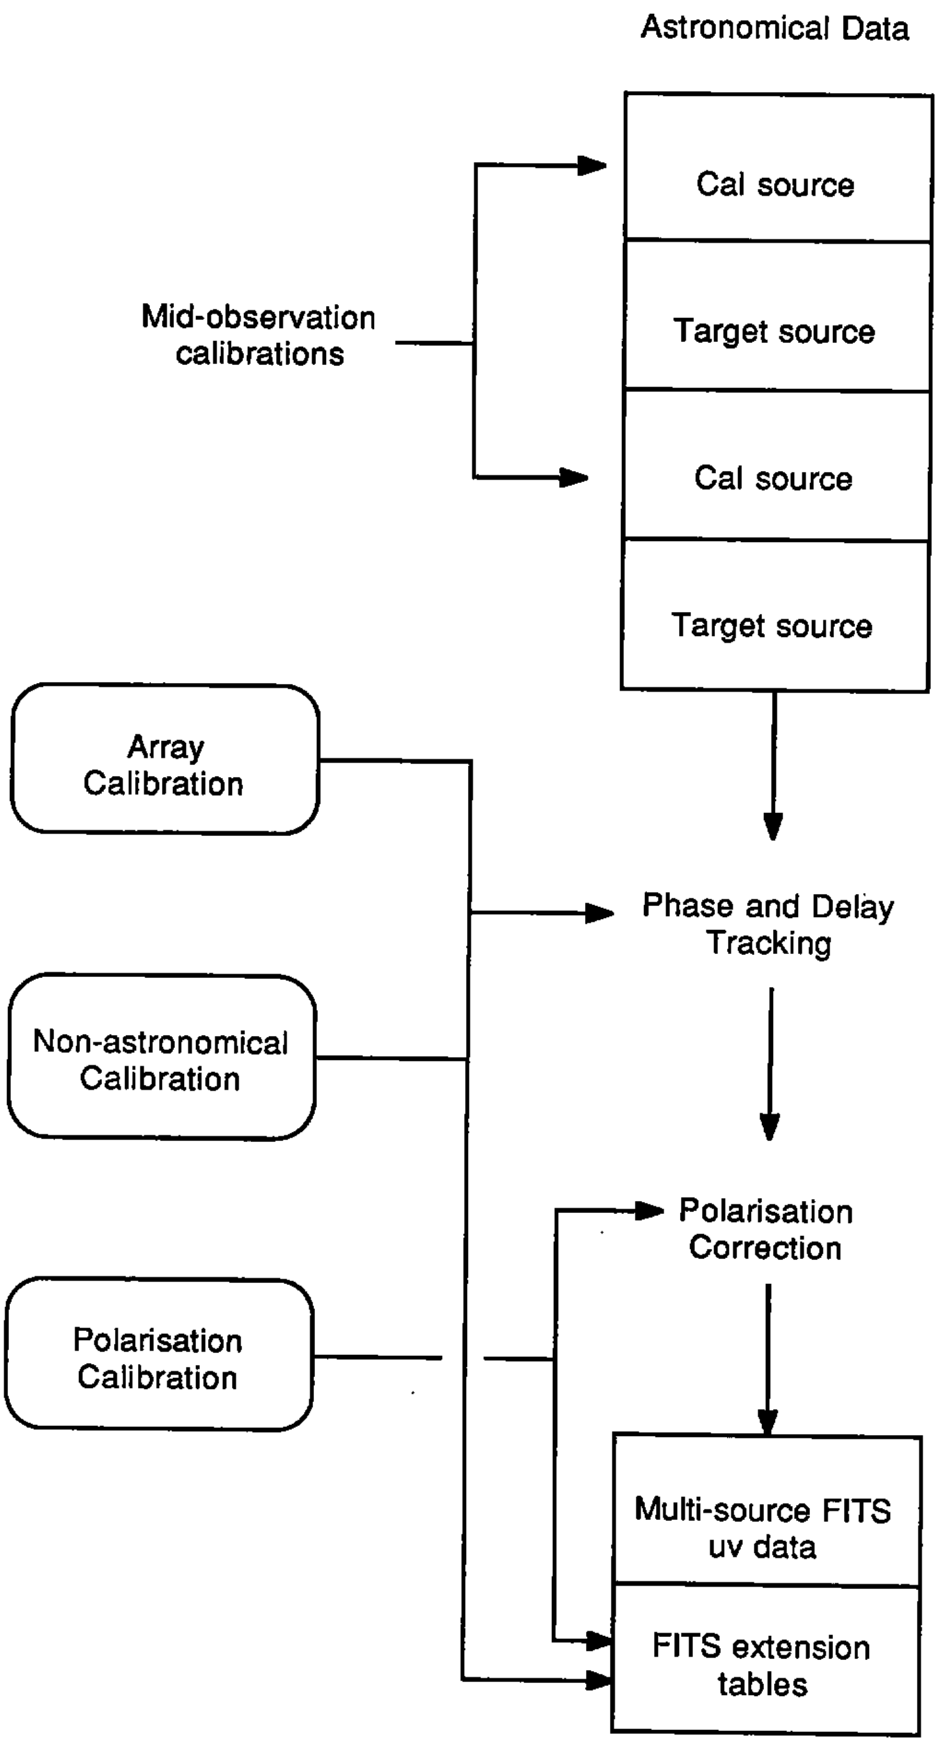

Figure 1: On-line calibration processes