ASTRONOMY AND SPACE SCIENCE
www.csiro.au

ASKAP antenna pointing calibration

David McConnell, Aidan Hotan, Michael Kesteven

ASKAP Commissioning and Early Science Memo 009
November 27, 2015

CSIRO Astronomy and Space Science
Cnr. Vimiera and Pembroke Roads
PO Box 76, Epping, NSW 1710, AUSTRALIA Telephone : +61 2 9372 4100
Fax : +61 2 9372 4310

Contents

Summary

This memo describes several methods that have been used on BETA to refine the pointing corrections that the Telescope Operating System (TOS) uses to ensure accurate tracking of astronomical sources. These methods were developed to take advantage of unique aspects of the ASKAP system, including the extended coverage of the focal plane provided by the Phased Array Feed (PAF) and the ability to rotate the reflector about its central axis (known as the roll axis, or the 3rd axis). Initially, we were concerned with only the 1st order pointing parameters, i.e. fixed offsets from the zero position in all three drive axes.

The first method involves observing the Sun and noting which PAF ports receive the most power. This effectively locates the reflector's Airy disk in the focal plane, allowing measurement of any offset from centre. We have shown that it is possible to achieve pointing accuracy of about 0.1°in azimuth and elevation using this method. Unfortunately, the roll axis is degenerate with the other two and any offset here will simply mix the azimuth offset into the elevation offset.

A second method resolves the roll-axis degeneracy. We observed the transit of the Sun across the surface of the PAF at its highest point in the sky. The path at this point should be symmetric in elevation, tracing a shallow curve. Any rotation of this curve can be used to estimate the roll axis offset. In addition, the location of the transit peak can be used to measure the accuracy of the timing system, once other corrections have been taken into account.

Finally, we describe a method using interferometry of calibrator sources with rotation of the roll axis to measure pointing relative to that axis. We report a test of this method and find sub-arcminute accuracy.

1 Introduction

Pointing calibration for radiotelescopes is usually done by observing the visibility amplitude (or total power for single dish telescopes) while pointing the radio beam close to a strong radio source of known celestial position. This can be done either as a continuous scan over the source or as a set of discrete pointings bracketing the source and spaced by about a beam-width. A function approximating the beam shape (usually Gaussian) is fit to the measured amplitudes and the true position of the beam derived.

In telescopes with a feed horn, the radio beam direction is fixed relative to the telescope's structure and defines its pointing direction. Imperfections in the mechanical pointing can be determined by measurements of the beam position on the sky. Any misalignment between the radio beam and the telescope's optical axis is absorbed in the ultimate pointing model.

The ASKAP antennas form their radio beams from a Phased Array Feed (PAF) and their precise direction relative to the antenna's optical axis, and their shape, can be controlled through the set of weighs used in the combination of PAF element signals. This, and the ability to form many simultaneous beams, is the advantage of the PAF for radio astronomy. However, for pointing calibration it introduces new sources of error not related to the mechanical pointing performance of the antenna.

ASKAP antennas have a third axis of motion—the "roll" axis—that allows the rotation of the primary reflector and the PAF about the optical axis. The roll axis is designed to hold formed beams that are inclined to the optical axis fixed on the rotating sky, but it also defines a reference direction, fixed to the mechanical pointing direction. In this note we describe a method for measuring apparent source positions relative to the direction of the roll axis.

In normal operation, radio beams are formed from the linear combination of signals from a number of PAF elements.[1] However, a radio beam can be formed from a single PAF element using a set of weights with only one non-zero value. A beam generated in this way (a so-called "single-port beam") has an offset position determined by the location of its element on the PAF, and may ultimately be usable to measure mechanical pointing errors in the conventional way (albeit with reduced sensitivity compared to a formed beam). In the present context, the single-port beams provide the off-axis response needed in this pointing measurement method.

This memo explores several methods of determining the relationship between the mechanical drive system and the optical pointing direction of the ASKAP antennas. As a starting point, we assume the same mathematical model used at the Australia Telescope Compact Array to translate between the two domains, with the addition of a parameter describing the zero-point offset of the roll axis.

2 The Antenna Pointing Model

The current pointing model is based on the one used at the Australia Telescope Compact Array (ATCA) and is described by nine parameters that include static offsets to position dependent corrections. The model parameters are defined as follows[2]:

[1] The Mark I PAFs used on BETA have 188 elements in two interleaved sets of 94, one set for each linear polarization. Typically X and Y beams are formed in pairs with the same pointing offset.

[2] It was discovered in late 2014 that the pointing model was being applied with an overall sign reversal (https://jira.csiro.au/browse/ASKAPTOS-3182). Since the model parameters were determined

```
pp[0] - cross-el error ("az squint").
pp[1] - az encoder zero
pp[2] - el axis skew (non-horizontality)
pp[3] - az axis tilt (I) = -phi*sin(K) tilt phi at az K
pp[4] - az axis tilt (II) = phi*cos(K)
pp[5] - el encoder zero
pp[6] - el gravitational deformation
pp[7] - residual refraction
pp[8] - pol encoder zero
```

As noted above, the presence of a PAF on the ASKAP antennas allows for some new ways of determining pointing corrections. The approach developed involves single-dish observations of the Sun to determine parameters 1, 5 and 8 to within roughly 0.1°, followed by interferometric observations of a number of pointing calibration sources distributed across the sky. These are fitted to the model to estimate the remaining parameters and refine the zero-point offsets.

Some of the antennas (notably AK15) exhibit large (of order 2°) zero point offsets in Azimuth as-built. This is large enough to significantly attenuate the power from an astronomical source if neglected, and caused significant confusion during early commissioning. For future work, it is recommended to perform the zero-point offset measurements as soon as possible (once an Array Covariance Matrix can be obtained). Measured offsets in Elevation have so far been much smaller. Larger offsets (of order 6°) are commonly present in the roll axis, but these are only important when considering polarisation.

3 Observing the Sun with a Single Antenna

3.1 Obtaining Zero–Point Offsets

This method was first used by Mike Kesteven to visually and roughly check antenna pointing accuracy after installation. When tracking the Sun, the shadow of the feed should fall directly in the centre of the reflecting surface. It is possible for the eye to perceive offsets as small as a few tenths of a degree, which can then be corrected with some trial and error.

With a PAF on the antenna under test, a similar operation can be performed quantitatively, by measuring the total power received by each PAF port when the antenna is tracking the Sun. This only works because the Sun radiates significantly more power than the system temperature of the receiver itself, and can therefore be detected (and spatially resolved on the surface of the PAF) with a single antenna. Because there are only 94 active PAF elements per polarisation, spatial resolution in the focal plane is quite poor. This can be overcome by fitting a two-dimensional Gaussian profile to approximate the central part of the reflector's Airy disk (see Figure 3.1). Due to the fact that the BETA PAFs are physically aligned with their axes of polarisation at 45 degrees to the horizon, the Azimuth and Elevation axes cut across the diagonals of Figure 3.1. The centre of the fitted Gaussian can be projected onto these two axes to measure the offsets in azimuth and elevation.

by experiment, the overall inversion resulted in the negation of several model parameters when compared with expectations. While this was mostly transparent, it caused a problem with the application of a correction for atmospheric refraction. The model was corrected and pointing parameters updated in the Facility Configuration Manager (FCM) on the 17th of December, 2014, revision 257.

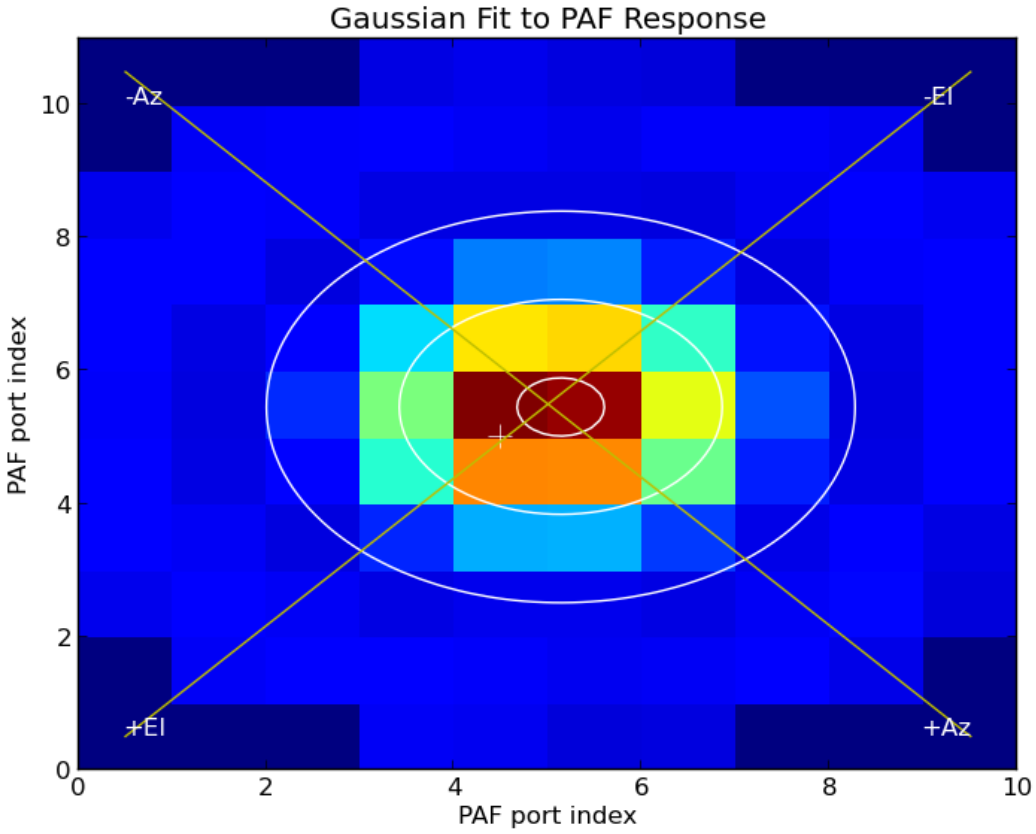

Figure 3.1: Example of the total power response of each PAF element to an observation of the Sun. The white contours correspond to a two-dimensional fitted Gaussian, showing 95%, 50% and 10% of the peak response.

A python script called acm_pointing_offsets.py was developed to make Figure 3.1. One important step in the development of the technique was the addition of a flat-field calibration stage. During the observations, it is necessary to include a blank field (far from the Sun) which is then divided into all Solar observations to flatten the response of the PAF. This improves the quality of the fitting process.

3.2 Observing Solar Transit

With a single observation of the Sun, it is not possible to separate a zero-point error in the roll axis from a combination of Az and El errors. The roll axis offset would be projected onto the other two axes, leading to incorrect estimates of their zero-point offsets. However, by making multiple observations of the Sun at different positions on the sky, it is possible to break this degeneracy. In particular, observations of the Sun over roughly 30 minutes on either side of its transit (crossing North when observed from the Southern Hemisphere) can reveal a roll axis offset and also a timing offset in the antenna control system (though this is in turn degenerate with an Azimuth zero point error).

The script used to calculate offsets from a single observation of the Sun was extended to loop over ACM cycles in a longer observation, producing independent offset measurements as a

function of time. This new script is called acm_transit_offsets.py. It produces a list of numbers similar to the following:

```
 Az       El      Time                          dE(P0) dA(P0) dE(P1) dA(P1)
 43.92,   25.03, 2014-06-12T01:21:39.0000 UT, -0.10,  0.41, -0.15,  0.43
 39.49,   28.38, 2014-06-12T01:44:10.0000 UT, -0.09,  0.42, -0.13,  0.45
 35.93,   30.67, 2014-06-12T02:00:57.0000 UT, -0.08,  0.42, -0.12,  0.46
 29.73,   33.92, 2014-06-12T02:27:45.0000 UT, -0.06,  0.42, -0.09,  0.46
 22.84,   36.60, 2014-06-12T02:54:53.0000 UT, -0.04,  0.43, -0.06,  0.48
 10.55,   39.42, 2014-06-12T03:38:44.0000 UT,  0.00,  0.43, -0.01,  0.51
 01.09,   40.14, 2014-06-12T04:10:31.0000 UT,  0.03,  0.43,  0.04,  0.53
-10.56,   39.42, 2014-06-12T04:49:34.0000 UT,  0.06,  0.43,  0.09,  0.55
-18.11,   37.96, 2014-06-12T05:15:58.0000 UT,  0.09,  0.42,  0.13,  0.55
-26.46,   35.29, 2014-06-12T05:47:20.0000 UT,  0.11,  0.41,  0.16,  0.55
-37.51,   29.68, 2014-06-12T06:34:40.0000 UT,  0.13,  0.41,  0.20,  0.56
-44.33,   24.67, 2014-06-12T07:08:52.0000 UT,  0.14,  0.40,  0.21,  0.56
```

The list above shows the measured offset as a function of antenna Az, El and time. From a list such as this, pointed observations of the Sun can be used to compute an average offset and uncertainty (see plot_offsets.py) or transit observations can be plotted as a function of time (see plot_pointing.py) for fixed antenna coordinates. An example of one such transit observation is shown in Figure 3.2.

The fact that the two lines in Figure 3.2 do not cross at the origin indicates that offsets are present in both axes. A rotation of the roll axis will project differently into the Az and El axes, therefore a useful measure of the roll axis offset is to take the tangent of the fractional slope (El/Az) of the lines of best fit. Assuming no timing errors in the drive system, the horizontal intercept of the Azimuth track can be used as another measure of the zero- point offset in that axis. The vertical intercept of the Elevation track gives its zero-point offset.

4 Interferometric Method

In its simplest form, the measurement is made by pointing two antennas at a bright radio source. One, the reference, is kept fixed, tracking the sidereal motion of the source with its PAF feed angle (Reynolds 2014) fixed at $FA = 0$. The antenna to be measured—the target antenna—also tracks the source, but is rotated about its roll axis. If its pointing is true and the roll axis points accurately at the radio source, the visibility amplitude on this interferometer will be constant (but note the discussion below on the polarised response). If the target antenna has a pointing error, the position of the source in the rotated single-port beam will vary causing a corresponding variation in the visibility amplitude, as illustrated in Figure 4.1.

4.1 Method implementation

The method used in practice needs some extra explanation.

4.1.1 Selection of single port beams

The PAF elements lie in a “chequerboard” pattern, symmetrically arranged about the optical axis, but with no element lying at the centre (Figure 4.2). However, all the elements close to the centre have significant response on the antenna boresight, so any of the central four elements could be chosen for the reference beam. That reference can be paired with any element on the target antenna; each pairing should give different but consistent information about the pointing error.

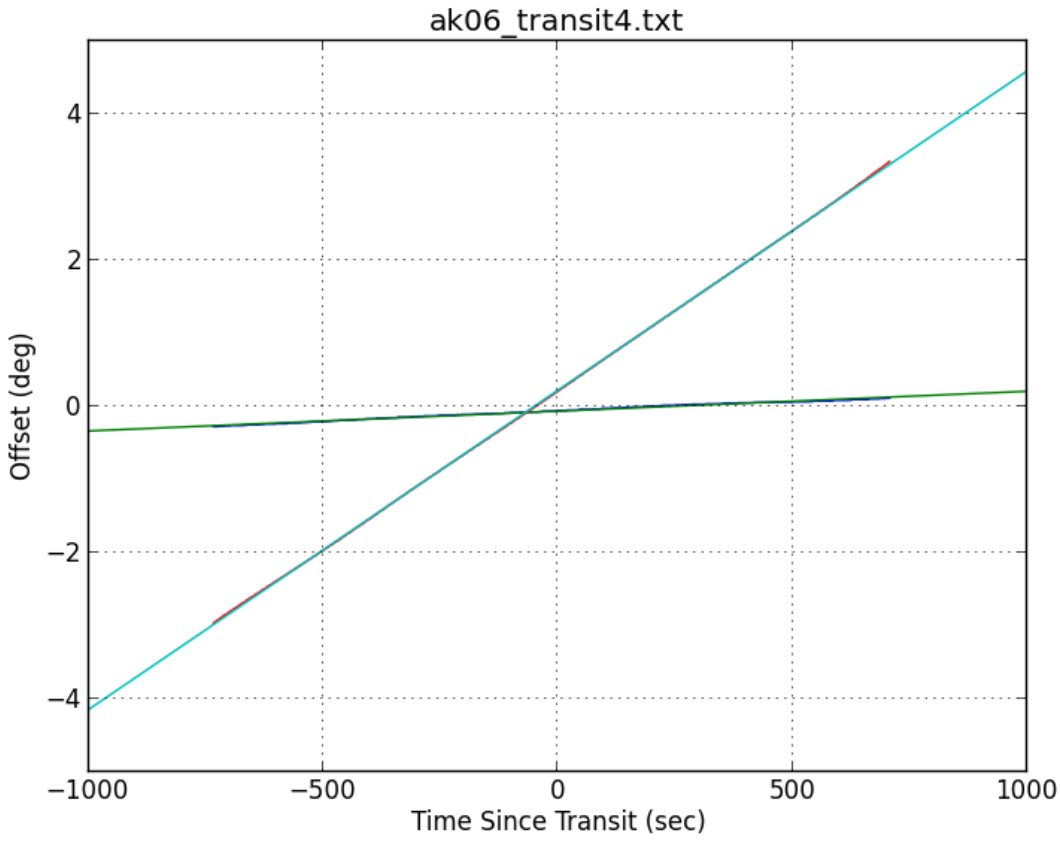

Figure 3.2: Example of a Solar transit observation, taken with antenna AK06 pointed at fixed Az, El coordintes of 0.0°, 39.6°; these being the expected coordinates of the Sun at its highest point in the sky on that day. The line closest to horizontal is the Elevation offset (averaged over both polarisations), with the data shown in darker blue, covering only the central portion of the fitted line. Axis intercepts give the Az and El zero-point offsets and the slopes can be used to determine the roll axis offset.

So, we form nine pairs, each with the same element on the reference antenna and nine different ports on the target antenna. We duplicate this scheme; one set of nine pairs for each of X and Y polarizations. Of these, we use only eight ports of each polarization to allow a balanced, symmetrical pattern around the boresight. Figure 4.3 shows the X and Y ports used in the target antennas.

4.1.2 Polarization

The method involves measuring the visibility amplitude as the target antenna is rotated about its roll axis. The two feeds are linearly polarized, and so the correlated response will vary as the cosine of the angle between them. This effect will combine multiplicatively with the variation from the source's motion in the response pattern of the target antenna, as illustrated in Figure 4.4.

If there is imperfect isolation between the X and Y feeds (i.e. polarization leakage), the correlated response will deviate from the ideal cosine function, and the polarization characteristics of the source will perturb the result.

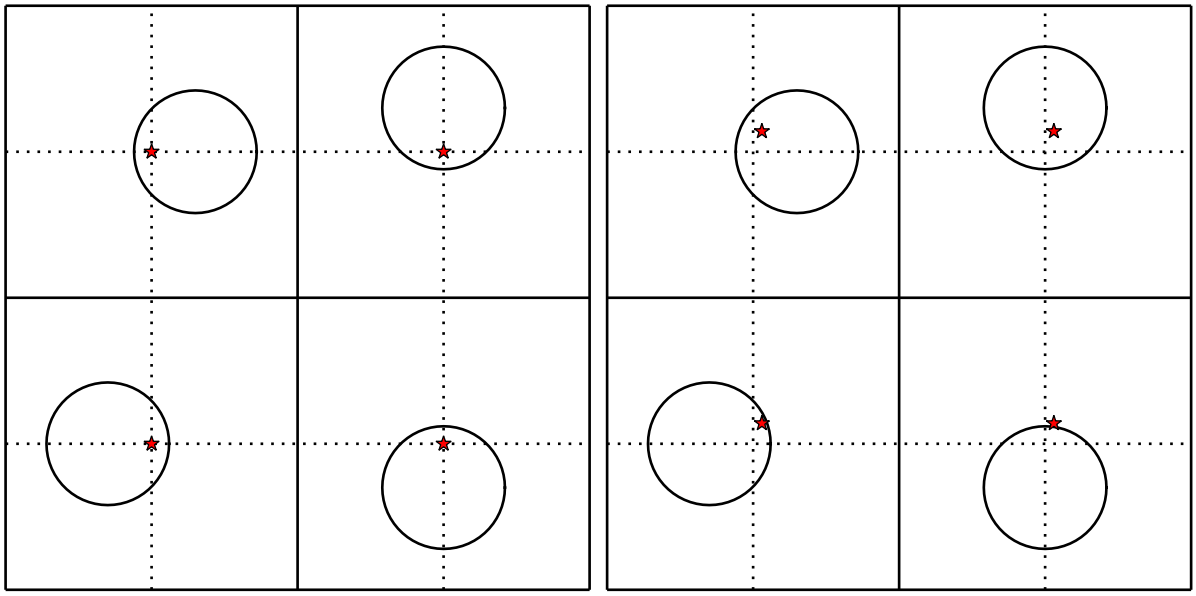

Figure 4.1: In the left panel a calibration source, marked as a red star, lies on the roll-axis; the antenna pointing has no error. The extent of a single-port beam is shown as a circle, offset from the roll (and optical) axis. As the antenna rotates about the roll axis and the beam moves about the optical boresight, the response to the calibration source should be constant; its offset from beam centre is fixed. In the right-hand panel a pointing error is indicated by the source being marked off-centre. As the antenna rotates, the source moves in and out of the beam. The corresponding variation in response to the source is used to determine the pointing error.

Let the X signal from the reference antenna be $v_{1_X} = X_1 + d_{1_X} Y_1$ where d_{1_X} is the leakage of $Y-$ polarized signal into the X feed. Similarly, $v_{2_X} = X_2 + d_{2_X} Y_2$ and $v_{2_Y} = d_{2_Y} X_2 + Y_2$ are the X and Y signals from the target antenna. From the correlator we read the visibility r_x as the time-averaged quantity

$$r_x = < v_{1_x} \times v_{2_x} \cos\theta > + < v_{1_x} \times v_{2_y} \sin\theta >$$

where θ is the feed angle ($\theta = FA$).

Expanding this, and dropping the time-averaging symbols $<>$, we get

$$\begin{aligned} r_x \;=\; & X_1 X_2 \cos\theta - X_1 Y_2 \sin\theta \\ & + d_{2_x} X_1 Y_2 \cos\theta - d_{2_y} X_1 X_2 \sin\theta \\ & + d_{1_x} Y_1 X_2 \cos\theta - d_{1_x} Y_1 Y_2 \sin\theta \\ & + d_{1_x} d_{2_x} Y_1 Y_2 \cos\theta - d_{1_x} d_{2_y} Y_1 X_2 \sin\theta \end{aligned}$$

If the source is unpolarized, $< XY >= 0$, $< X_1 X_2 >=< Y_1 Y_2 >= I/2$, where I is the total intensity (Stokes I) and

$$r_x = \frac{I}{2} \left[\cos\theta - (d_{1_x} + d_{2_y}) \sin\theta + d_{1_x} d_{2_x} \cos\theta \right]$$

showing that the ideal cosine variation is perturbed by the sum of reference and target antenna leakages.

Leakages have been measured on the BETA antennas (Sault (2014)) and have amplitudes less than 0.01, except on antennas with known roll-axis zero-point errors. For unpolarised sources, the error induced in the pointing measurement by instrumental leakages, given the gradient near half-power of the $\sim$1 GHz beam, will be $\lesssim$ 0.3 arcminutes. Note that the effect of leakage in the reference antenna (d_{1_x} in the above) is diminished by the scheme of measuring each antenna against three different references.

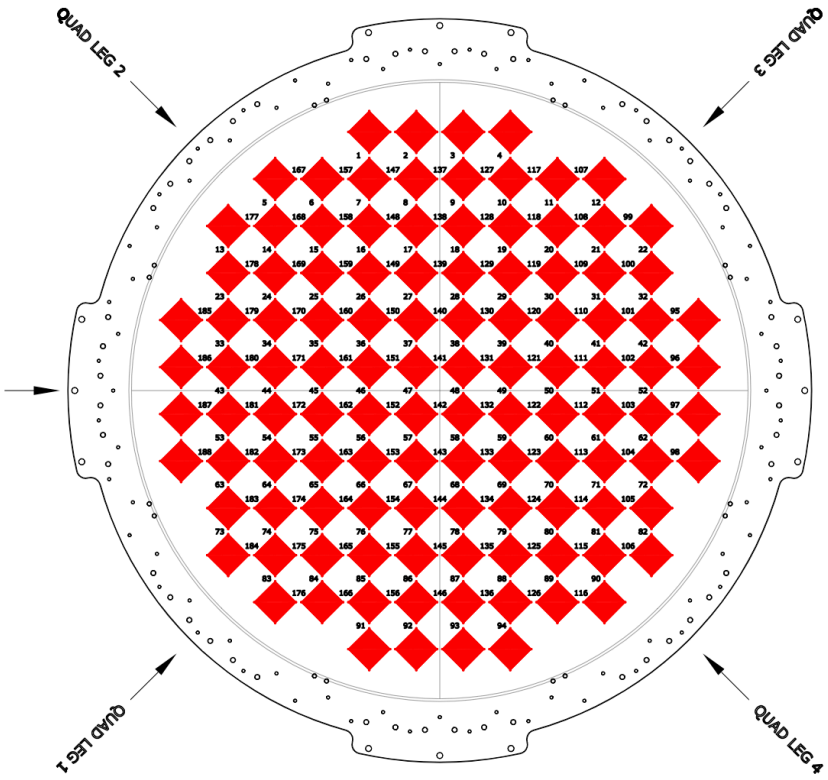

Figure 4.2: Layout of the Mark I Phased Array Feed used on the BETA antennas. This is the view from the dish vertex, and Quad leg 2 is closest to the zenith when the Feed Angle is set to zero $FA = 0$.

4.1.3 Observations

A single pointing observation makes measurements for pointing error determination on half[3] of the antennas; the other half are used as the reference antennas. All antennas are pointed at the chosen pointing calibration source and visibilities are recorded in the standard manner, with the phase-tracking centre set to the position of the pointing source. The observation is made in four phases, each with a different FA set on the target antennas; the reference antennas have $FA = 0$. Two sets of beam weights are defined, all with "trivial" weights, that is $w_j = 1$ and $w_{i\neq j} = 0$, $i \in 1, ..., 94$ to select PAF element j. The weights used for the reference antennas selects elements 47 (X), and 142 (Y) for all beams. The weights used for the target antennas selects

[3]The optimum division of antennas between target and reference is equal: it minimises the number of observations needed to measure pointing errors on all antennas, and keeps the sensitivity of all measurements equal.

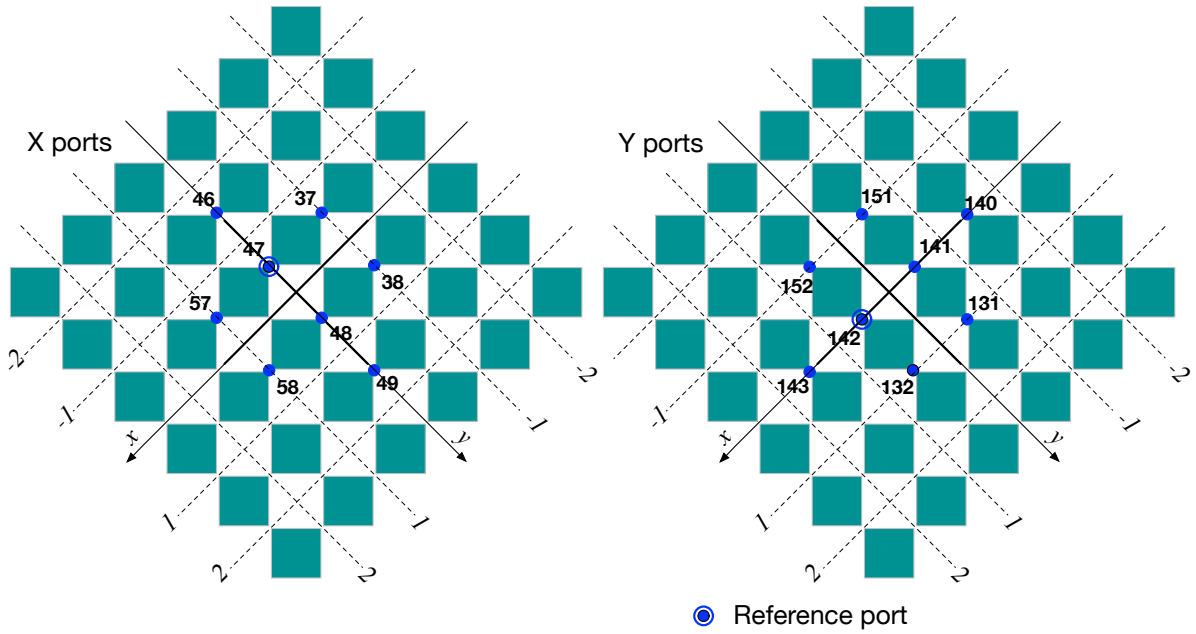

Figure 4.3: The central part of the PAF, showing the X (left) and Y ports used in the target antennas. In this view the Zenith is towards the top of the diagram when the Feed Angle is set to zero $FA = 0$. The diagrams are overlaid with the coordinate grid used in the fitting, following Reynolds (2014), shown here in units of the 90mm pitch of the PAF pattern. In these coordinates, the position of port 57 is $(x, y) = (1.0, -0.5)$.

elements {38,48,58,49,47,57,37,46} and {140,141,142,131,151,152,132,143} for the X and Y polarizations respectively. Usually, a second observation is made of the same calibration source, with the roles of the antennas swapped between target and reference.

4.1.4 Data preparation

The observed data are written to a CASA-style MeasurementSet, and are pre-processed to give a new condensed MeasurementSet for each beam. A python script is run in the CASA environment to reduce these pre-processed data to neasurements of pointing errors. The python script `$ACES/pythonlib/aces/pntcasa.py` hold all the procedures mentioned in this section. The data processing has two phases: data preparation and fitting (§4.1.5).

Visisbility measurements are read from the MeasurementSets by routine `getData`, which is called once for each pointing observation (see §4.1.3). It returns data structures for each of X and Y polarization: `datax` and `datay`. Each of these hold a set of masked arrays (`numpy.ma.masked_array`) of shape `(nch,nt)` where `nch` and `nt` are the number of frequency channels and the number of time intervals. The structure of `datax` and `datay` is: $\texttt{datax}[\texttt{k}_\texttt{a}][\texttt{i}_\texttt{ref}][\texttt{k}_\texttt{p}][\texttt{i}_\texttt{pa}]$ where

$\texttt{k}_\texttt{a}$	dictionary key	target antenna number
$\texttt{i}_\texttt{ref} \in \{0, 1, ..., N_{refant} - 1\}$	list index	reference antenna sequence number
$\texttt{k}_\texttt{p}$	dictionary key	PAF port number
$\texttt{i}_\texttt{pa} \in \{0, 1, 2, 3\}$	list index	Position angle sequence number

The port numbers for the X and Y polarizations are given in §4.1.3.

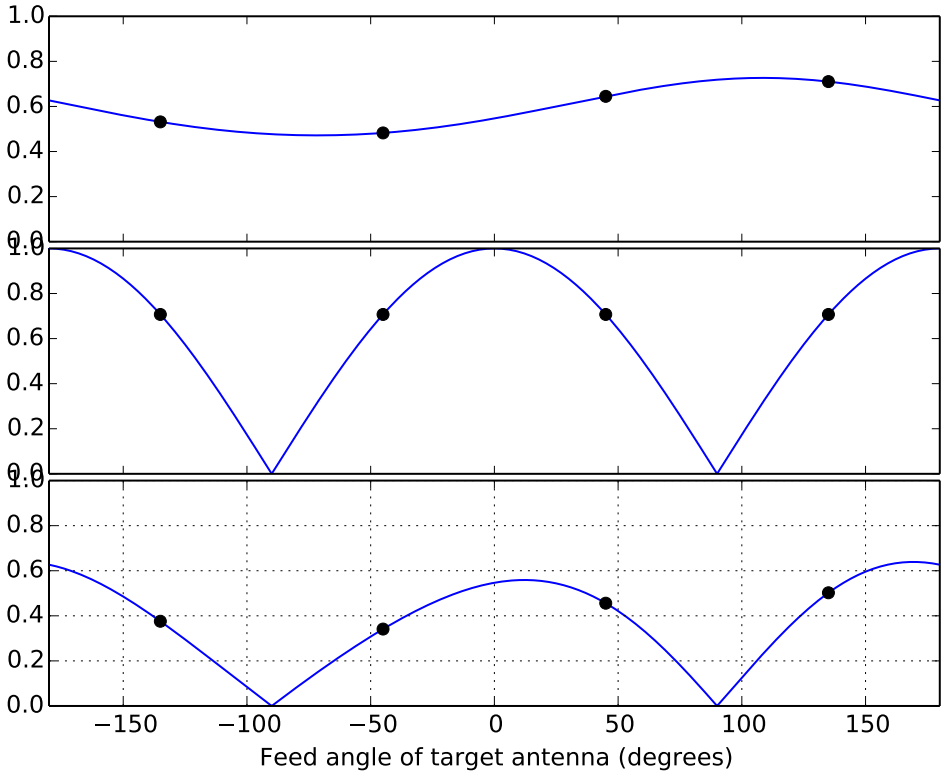

Figure 4.4: An illustration of how the visibility amplitude might vary with FA when observing a source with a target antenna suffering a pointing error. The top panel shows the power received through a sample single port beam on the target antenna, normalised by the power received at the beam peak. The middle panel shows the expected visibility amplitude with no pointing error, the variation arising from the changing alignment between the two linearly polarized feeds. The bottom panel is the visibility amplitude with a non-zero pointing error; it is the product of the two functions above. As currently implemented, the technique uses measurements made at the points marked: $FA = -135°, -45°, +45°, +135°$, where all the variation should arise from a pointing error.

Each of the complex visibility arrays are reduced to a single representative signal amplitude in routine `avgData`. The result is a vector of four amplitudes (one for each position angle) for each antenna–PAF port combination, stored in a structure $\texttt{datax}[\texttt{k}_\texttt{a}][\texttt{k}_\texttt{p}][\texttt{i}_\texttt{pa}]$ where

$\texttt{k}_\texttt{a}$	dictionary key	target antenna number
$\texttt{k}_\texttt{p}$	dictionary key	PAF port number
$\texttt{i}_\texttt{pa} \in \{0, 1, 2, 3\}$	list index	Position angle sequence number

The reduction proceeds in these steps:

1. identify possible RFI in the spectrum and flag;
2. fit for non-zero delay (a phase slope across the spectrum) and correct;
3. compute the amplitude as the magnitude of the vector average;
4. compute the average amplitude across all reference antennas.

The procedure assumes that the array is not calibrated beyond the antenna position and delay model. The signals received through individual PAF ports may have small relative delays (typically $\sim$1 ns), manifested as a phase gradient across the band. Also, the visibility phase should be constant because data is taken with phase tracking enabled and the phase centre set to the position of the calibration source; however, the phase cannot be relied upon to be zero, and in general has different values on different baselines. This means that the combination of measured responses between a target and its several reference antennas must be a scalar average. One of the advantages of measuring the target antenna against several references is that only the responses to the source at the phase centre add coherently; confusing sources in the field will result in complex additions to the visibility with baseline-specific phase. Ideally, the phase of the calibration source would be the same (and zero) on all baselines, so that a vector average across reference antennas would preserve its response, but reduce the out-of-phase responses to confusing sources. However, pointing calibration sources are chosen to be bright and typically have ten times the flux-density of the next strongest source in the field. Under these circumstances, there is little difference between the scalar and vector averages.

4.1.5 Fitting

For each antenna, data to be modelled are the set of $4N_p$ amplitudes corresponding to the measured responses from the N_p PAF elements at the four position angles. The response pattern for each element is assumed to be a Gaussian function whose centre is determined by the element's position on the PAF, and whose width is a function of the antenna diameter and wavelength of the observed radiation. The gains of the PAF elements are known to vary from element to element, and so the height of each response pattern is allowed to vary in the fit. The fitting model has $N_p + 2$ parameters: the two components (e_x, e_y) of the source position relative to the grid defined by the element responses, and the N_p response heights.

The coordinates of the problem are derived from the description of the PAF geometry given in Reynolds (2014) and are illustrated in Figure 4.3. The conversion from the PAF grid to angular units is a multiplication by the factor $BDF \times p/f$ where BDF is the "beam deviation factor", p is the size (pitch) of the square tesselations on the PAF, and f is the focal length of the primary reflector. ASKAP has values $f = 6000$mm and $p = 90$mm. The beam deviation factor is a function of the paraboloid focal ratio; for ASKAP antennas $BDF \simeq 0.86$ (Reynolds 2014). The assumed beam centre for each PAF element is calculated as its grid position multiplied by this factor to give its position on a rectangular frame in angular units on the sky which is the analogue of the PAF grid illustrated in Figure 4.3.

The model fitting is organised in routine `rollFit`, which uses the python routine `curve_fit` in `scipy.optimize` to find the optimal parameter values. `Scipy.optimize.curve_fit` performs a non-linear least-squares fit using the Levenberg-Marquardt algorithm, and is called as

```
func_yy = [avg[ant][jp] for jp in ports]
func_yy = np.array(func_yy).flatten()
hs = np.array(len(ports)*[max(func_yy)])
p0  = np.concatenate((exey,hs))
w = op.curve_fit(f2, func_xx, func_yy, p0)
```

It expects the name of a function to call and optimize the parameters `p0` to minimize (least-squares) difference between the function of the independent variables `func_xx` and the measured dependent variables `func_yy`. The function `f2` must be callable with dependent variables as a linear array, and the parameters as individual scalars. It is defined as follows, and

returns a vector of model values of length equal to the number of ports used.

```
def f2(x, ex, ey, h1=1.0, h2=1.0, h3=1.0, h4=1.0, h5=1.0, h6=1.0,h7=1.0,h8=1.0):
    # This is the function whose parameters are optimized. The parameters are ex,ey
    # and a set of heights, one for each port. These heights are allowed to vary,
    # reflecting the expected port-to-port gain variations.
    # op.curve_fit, used in rollFit, requires the function being optimized to have a
    # 1D array holding the independent variable.  We have a function of two variables:
    # port and position angle, so a scheme using decodexx and encodexx is used to get
    # our 2D independent variables through a 1D hole.
    p = [ex,ey,h1,h2,h3,h4,h5,h6,h7,h8]
    xdat = [decodexx(a) for a in x]
    v = np.array([func(p,a[0],a[1],a[2]) for a in xdat])
    return v
```

For each port measured, `f2` calls `func` (which in turn calls `fn`)to compute the modelled response of port `pj`, with roll axis set to `pa` to a source shifted (e_x, e_y) = `p[0:2]` through a beam with peak response h = `p[j]`.

```
def func(p,pa,pj,j):
    # p: parameters [ex,ey,[h]]
    # pa: the position angle
    # pj: the port number
    # j: the index of pj, used to reference the correct height parameter

    ex = p[0]
    ey = p[1]
    h = p[2:]
    W = Wglob

    s1 = fn(ex,ey,pa,pj,W,h[j])
    return s1


def fn(ex,ey,pa,pj,W,h):
    # This function computes the voltage response from port pj for a source
    # at position (ex,ey) relative to bore-sight, with the PAF as position angle PA.
    # The position (error) vector (ex,ey) is in units of grid spacing and aligned
    # with the grid; this vector is rotated 45deg relative to a pointing error vector
    # expressed as (dAz*cos(el),dEl).
    # The fitted quantity is rotated to the horizontal/vertical coordinates in rollFit.
    # W is the single-port beam width (FWHM) in radians
    fac = plateScale*pitch
    dx = ex*fac
    dy = ey*fac
    if pj < 95:
        Xp = Xxp[pj-1]
        Yp = Yxp[pj-1]
    else:
        Xp = Xyp[pj-1-94]
        Yp = Yyp[pj-1-94]
    # Note that we rotate the nominated x,y coordinate by -PA; PAF has rotated +PA.
    cp = cos(-pa)
    sp = sin(-pa)
    xp = cp*dx-sp*dy
    yp = sp*dx+cp*dy
    rp = np.sqrt((xp-Xp)**2 + (yp-Yp)**2)
    # In the expression for s below, the square root enters because we are exploring
    # the voltage pattern of the single-port beam, not power.
    # b = 2*np.sqrt(np.log(2))/W
    # s = h * np.sqrt(np.exp(-(b*rp)**2))
    # Rewrite these to avoid the square root, by dividing b by sqrt(2).
    # b = np.sqrt(2)*np.sqrt(np.log(2))/W
    # and rewrite again to avoid repeated log and sqrt calls.
    b = 1.1774100225154747/W
    s = h * np.exp(-(b*rp)**2)
    return s
```

4.2 Testing the method

A test was conducted[4] to assess the success of the method. The position of the source B1934–638 was measured repeatedly with intentional pointing offsets added cyclically to the observations. Pointing offsets of ± 3 arcminutes were applied as described in the table below.

Scan	Offset (arcmin)	Targets	References
1	0	8, 15	1, 3
2	0	1, 3	8, 15
3	$+3$	8, 15	1, 3
4	$+3$	1, 3	8, 15
5	-3	8, 15	1, 3
6	-3	1, 3	8, 15
	repeat 3 times		
25	0	8, 15	1, 3
26	0	1, 3	8, 15

Table 4.1: Observing scans conducted for the test.

During this time the source moved over the hour angle range -6^h48^m to $+0^h55^m$ (elevation range 17°– 53°). The data were analysed using the method described above. The estimated pointing offsets were transformed to (HA,Dec) coordinates for easy comparison with the intentionally introduced offsets. Figure 4.5 shows the results for the four antennas measured. Three of the four antennas showed a marked trend in their pointing errors measured in the zero-offset scans $(1, 2, 7, 8, \ldots, 25, 26)$. We assume these trends are the result of an imperfect set of parameters loaded for the global pointing model, and we fit those trends with a quadratic function and subtract the fitted values from all scan results.

The effectiveness of the method described in this memorandum can be gauged by estimating the measurement error. We do this by taking the measured pointing errors in hour angle and declination and subtracting both the trends determined from the zero-offset scans, and the ± 3 arcminute declination offsets. The resulting residuals will include contributions from the following:

- random measurement error arising from the finite signal-to-noise in the data;
- systematic errors arising from false assumptions about the single-port beams, such as the positions of their centroids and their circularity;
- real non-repeatibility in the antenna pointing.

Figure 4.6 shows the residual errors; the means and standard deviations in $\Delta\mathrm{HA}/\cos\delta$ and $\Delta\delta$ are -0.04 ± 0.39 and 0.06 ± 0.31 arcminutes, respectively. The residual errors do not appear to belong to a Gaussian distribution.

References

Reynolds, J. E. 2014, ACES memorandum, 1
Sault, R. J. 2014, ACES memorandum, 2, 1

[4] Scheduling block 1970 executed on 2015-June-17.

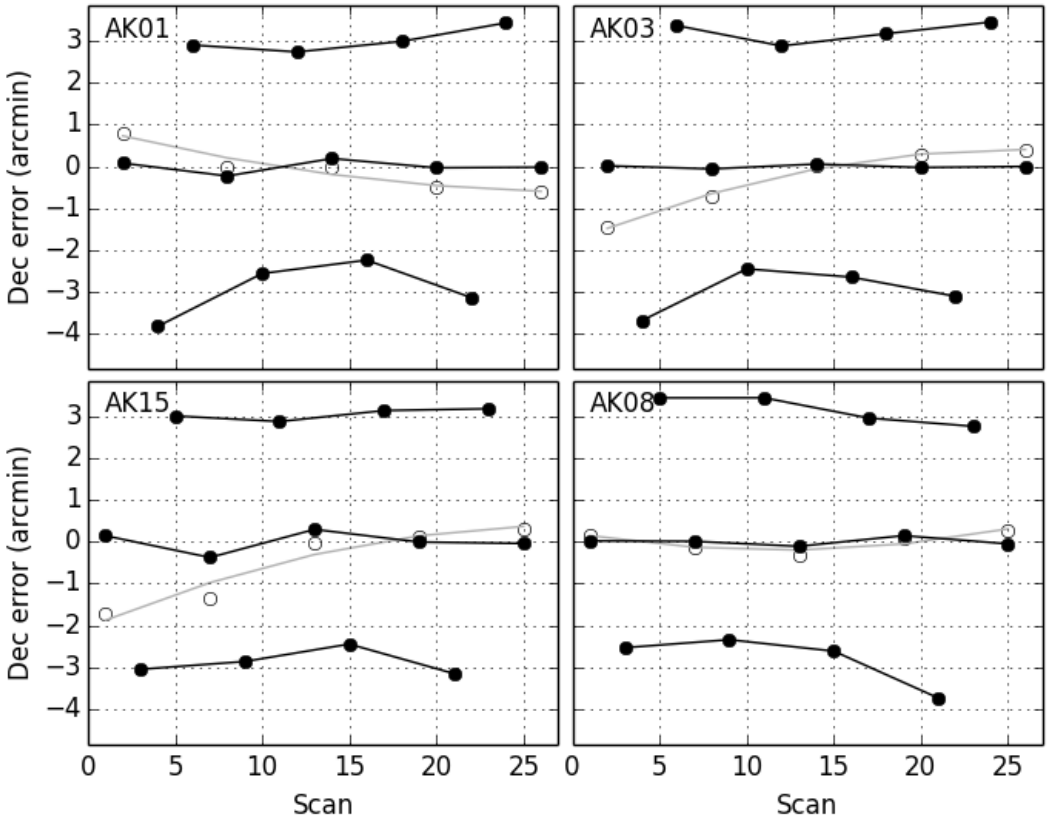

Figure 4.5: Results from the declination offset test described in §4.2. The open grey points show the declination error measured with zero-offset pointings. The grey line is the quadratic fit to these; the trend is assumed due to an imperfect set of pointing model parameters. The filled black symbols show the results with these trends removed, connected in groups for the 0, -3 and $+3$ arcminute offsets. The abscissae are labelled by the scan numbers of Table 4.2.

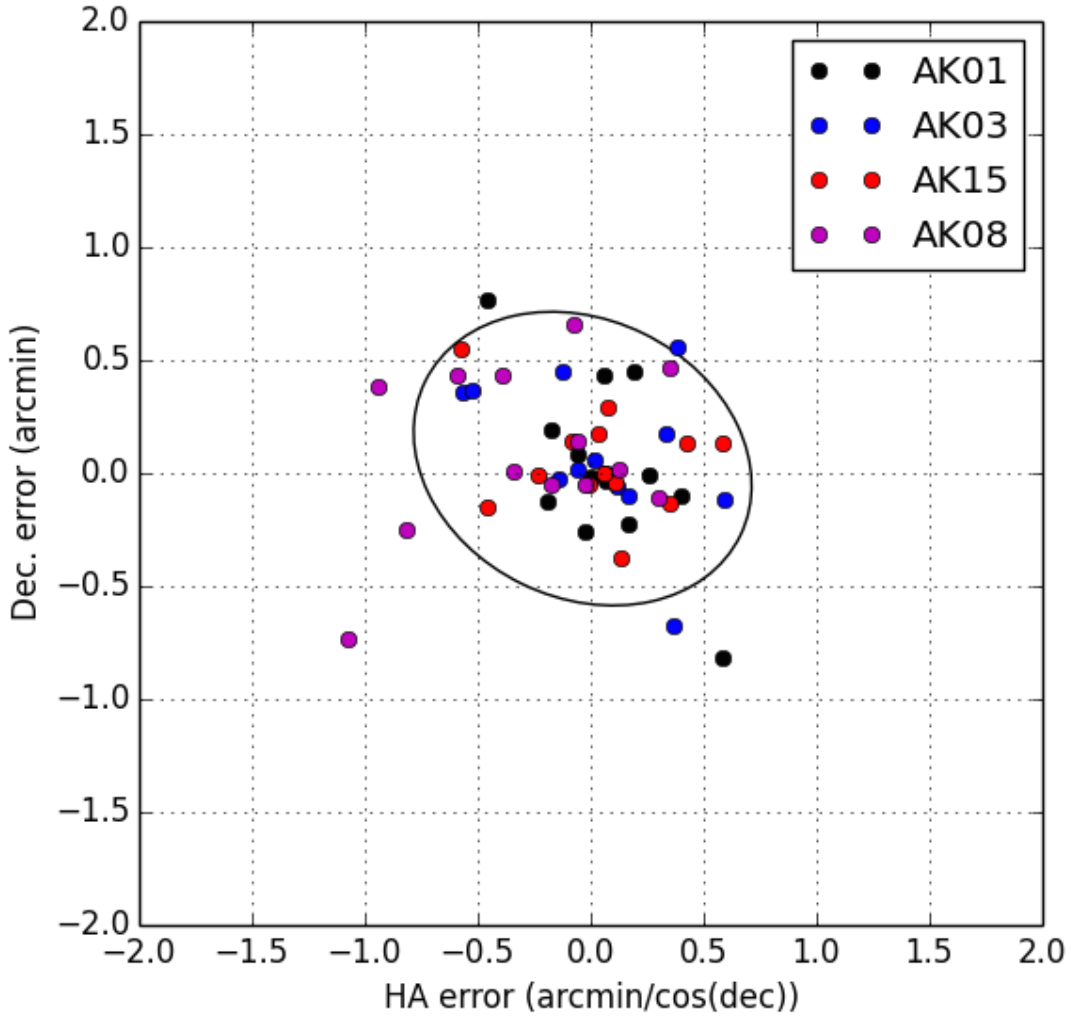

Figure 4.6: Residual measurement errors $(\frac{\Delta_{\mathrm{HA}}}{\cos\delta}, \Delta\delta)$ give an indication of the efficacy of the method (see §4.2). The ellipse is computed from the covariance matrix and is drawn at the 2σ level.

CONTACT US

t 1300 363 400
+61 3 9545 2176
e enquiries@csiro.au
w www.csiro.au

YOUR CSIRO

Australia is founding its future on science and innovation. Its national science agency, CSIRO, is a powerhouse of ideas, technologies and skills for building prosperity, growth, health and sustainability. It serves governments, industries, business and communities across the nation.

FOR FURTHER INFORMATION

CSIRO Astronomy and Space Science
David McConnell
t +61 2 9372 4132
e David.McConnell@csiro.au
w Astronomy and Space Science

CSIRO Astronomy and Space Science
Aidan Hotan
t +61 2 9372 4425
e Aidan.Hotan@csiro.au
w Astronomy and Space Science