ACES Memo #7
Widefield polarimetric considerations for ASKAP

R.J. Sault

5 May 2015

Summary

This builds on the earlier memo on BETA polarimetric characterisation (hereafter memo#2). It considers issues with widefield polarimetric observations with the current BETA system, and makes suggestions for implementing widefield polarimetric processing on ASKAP. The conclusions are:

- The ASKAP antennas and BETA hardware have good and stable polarimetric response. The different BETA antennas show the same response, which will simplify the processing needed for widefield polarimetric imaging.

- The ability of BETA (and ultimately the full ASKAP) to produce quality widefield polarimetric images is critically dependent on the algorithm used to determine the beamformer weights. The algorithm currently used for BETA is inadequate. More broadly, the beamforming algorithm will affect the mosaicing capability of ASKAP. There is a need to develop better algorithms to determine the beamformer weights.

- Correction for both widefield instrumental polarization and ionospheric Faraday rotation can probably be achieved without the need for so-called *A-projection* algorithms.

An observation

To understand some of the characteristics of the widefield polarimetric response of BETA, a "standard" observation was made of 1934-638 on 26 March 2015 in the frequency range $700 - 1000$ MHz. Throughout the observation (for beamforming, bandpass calibration, and then main run on 1934-638) the roll axis of the BETA antennas were enabled to remove their parallactic angle rotation (i.e. to 'deparallactify' the the observation). With this deparallactification the responses of the antennas are akin to equatorial-mounted ones. In this mode, the alignment of the X dipole is always at $45°$ to the meridian (the quadrapod legs align with or are orthogonal to the meridian). As in memo#2, the roll axis on AK03 was intentionally misaligned by $5°$ to allow the system XY phase to be measured.

The footprint of nine formed beams of BETA (with the beams numbered from 0 to 8) were made in a square 3×3 pattern, with the separation between adjacent
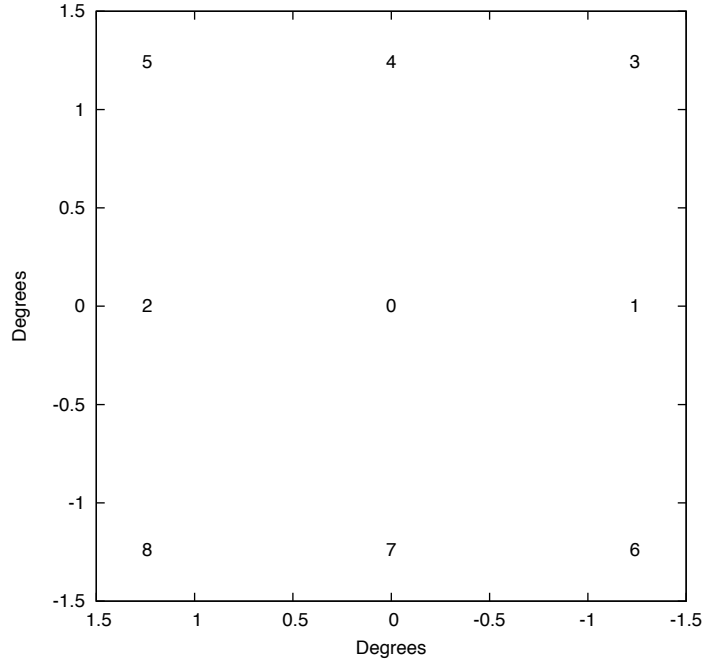
Figure 1: Beam footprint geometry

beams of $1.24°$ – see Fig 1. The so-called "maximum sensitivity" algorithm was used in determining the beamformer weights. This beamforming was done shortly before the observing session, using the Sun as the primary source and a "blank field" $15°$ south of the Sun in declination. During the beamforming, the elevation of the Sun and blank field were approximately $55°$ and $66°$ respectively, and the parallactic angle of the blank field was approximately $135°$.

A "bandpass calibration" sequence was performed on the nine formed beams by pointing each of the formed beams at 1934-638 for 15 min. These data were also used for the XY phase measurement and (on-axis) leakage calibration of each formed beam. XY phase and leakages were determined with a spectral resolution of 5 MHz. As with memo#2, the leakages were uniformly low and reasonably constant with frequency. The imaginary parts of the leakages were very low, whereas the real parts show that there is a small rotational misalignment of the PAFs on some antennas. These misalignments were consistent with previous results and the intentional misalignment introduced into antenna AK03.

The main observation was for 12 hours, with the central beam (beam 0) pointed at 1934-638. The other beams would see 1934-638 at pointing offsets of 1.24 or $1.75°$. The delay tracking centre was set to 1934-638 for all formed beams. This observation was taken 11 hours after the bandpass and beamforming sequences, and so some variation in the instrumental amplitudes and phases might be expected. The calibrations from the "bandpass calibration" sequence were first applied to main observation data. The residual amplitude and phase errors were largely independent of formed beam and frequency but dependent on antenna and time. To correct for this residual error, amplitude and phase calibrations derived from the central beam were applied to all the formed beams. After this correction, the final residual phase error was typically several degrees. It is

believed some of this is intrinsic to the off-axis response being investigated.

As a check, XY phase and leakage calibrations were also derived from the central beam of the main observation. It was found that there was no appreciable change in these between the main observation and the earlier "bandpass calibration" sequence. There is no reason to believe there was any appreciable change for other beams either.

The source 1934-638 can be taken as unpolarized for the purposes of this test. It has no detected linear polarization and its circular polarization is lower than the precision of the results presented here. Because it can be assumed unpolarized, ionospheric Faraday rotation can be neglected in the calibration. The apparent polarization in the observations is instrumental in origin, and is not affected by errors in the assumed ionospheric properties.

The analysis that follows is based on the instrumental response derived from data for different beams and correlations/Stokes parameters. They are averaged over the duration of the main observation and, generally, for all available baselines. For the offset beams in the main observation, in addition to the response of 1934-638 there are the responses (fringes) of confusing sources in the fields. These confusing sources have not been modelled. Combining the data over 12 hours averages out the response of these confusing sources. Similarly by combining all baselines together, confusing sources and any antenna-to-antenna variability are averaged out. Inspecting the baselines individually did not show any convincing differences between the antennas, although the confusing sources limits the authority of this statement.

In general, only the real part of the response is given as the imaginary parts were small and not believed to be significant.

Results: Parallel hands

Figure 2 shows the instrumental gain in beams 1 through 8. As the response is essentially real-valued, the plots are the real part only. Figure 2a gives the formed beams at the north, south, east and west positions, and Fig. 2b the north-east, north-west, south-east and south-west formed beams. These spectra show the primary beam response as a function of frequency for at offsets of 1.24 or $1.75°$ from the beam centre (Fig. 2a and 2b respectively).

Figure 2 shows a periodicity of about 25 MHz at the low frequency end of the band. It is believed this is caused by a "standing wave" effect between the dish surface and PAF (note the distance between PAF and dish surface is 6 m, implying a "standing wave" frequency of 25 MHz). Figure 2 also shows regular sharp discontinuities. The weights used the in beamformer are fixed over a frequency range of 4 or 5 MHz (where it is 4 MHz and where it is 5 MHz is a detail which is not important here): the discontinuities happen exactly at the boundary where the beamformer weights change.

In Fig. 2a, the XX and YY responses are moderately similar, whereas in Fig 2b they are appreciably different. This is caused by the orientation of the elemental "dipoles" that are used to form the beams. For those in Fig 2a, the beams are offset from the antenna boresite in directions that are at $45°$ to both the X and Y dipole axes. However in Fig. 2b, the offset is aligned with one dipole axis and normal to the other dipole axis. Clearly Fig 2b shows that the formed beams are quite elliptical, with the major axis of the XX and YY formed beams being

normal to each other. Note that which dipole is aligned and which dipole is normal changes between the four plots of Fig. 2b (for beams 3 and 8, the YY response is larger than XX, whereas the reverse is true for beams 5 and 6).

Results: Cross-hands

Figure 3 shows the XY and YX correlations of the offset beams as a function of frequency. The response is largely real-valued – the figure gives the real part only. Note the difference in scales between Fig 3a and 3b. Note that the response of beams 1 and 2 are very similar, and that response of beams 4 and 7 are approximately the negative of beams 1 and 2. For a given beam, the XY and YX correlations are always similar.

Some simple modelling

The observed responses have been compared with the expected polarimetric response of a simple model antenna system: an ideal parabolic antenna with a circularly-symmetric illumination pattern. The illumination pattern is assumed to be a gaussian with a central blockage. The amount of tapering of the pattern was tuned to give reasonable agreement between the model system and the actual response of beams 1, 2, 4 and 7. Otherwise the parameters used to determine the response are taken from the ASKAP antennas (antenna size, focal ratio and central blockage caused by the PAF). The response in this case has been derived from the analysis of Ghobrial (1976, IEEE Trans Ant. Prop., vol. 24). Figure 4 gives the model illumination pattern.

This model of the response is a simple one. It assumes circular symmetry, and assumes all the polarimetric properties of the response are determined by the reflection off the dish surface. It ignores blockage by the quadrapod legs (which is significant). It also assumes illumination which is circularly symmetric, which the "dipole" elements that make up the PAF are unlikely to produce. The induced instrumental polarization from the simple model system will be purely radially directed linear polarization. When this circularly symmetric response is expressed in a Cartesian system (as is natural when the probes are nominally linear and orthogonal), the results are "clover leaf" patterns. For this set of observations, the offset beams either fall at the peaks (positive or negative) or zero crossings of the clover leaf pattern.

Figures 5, 6, 7 and 8 compare the measured with the model response for the offset beams. These figures show:

- the total intensity primary beam response,
- $XX - YY$ ("Q response in the frame of the antennas"),
- XY and YX correlations ("U response in the frame of the antennas") and
- Stokes V responses

respectively.

Figure 5 shows reasonable agreement. Recall the illumination taper has been tuned to give reasonable agreement in the upper panel, although agreement in
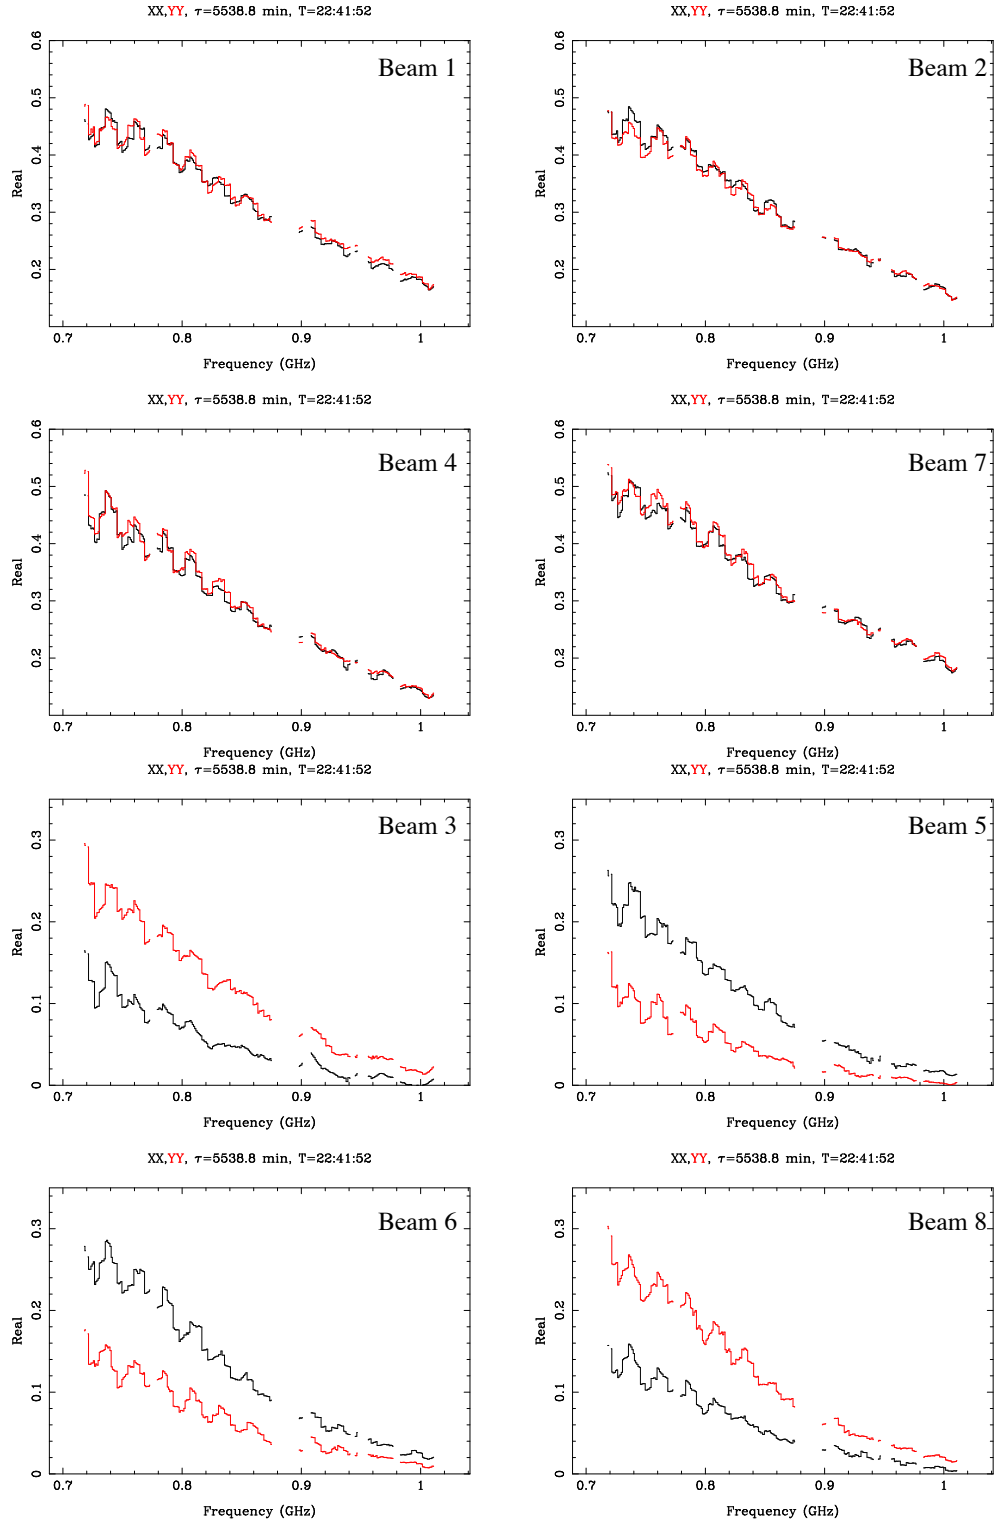
Figure 2: Response in XX and YY for the eight offset beams. (a) The upper four panels: Beams 1, 2, 4 and 7. (b) The lower four panels: Beams 3, 5, 6 and 8. Note the difference in scales between (a) and (b).
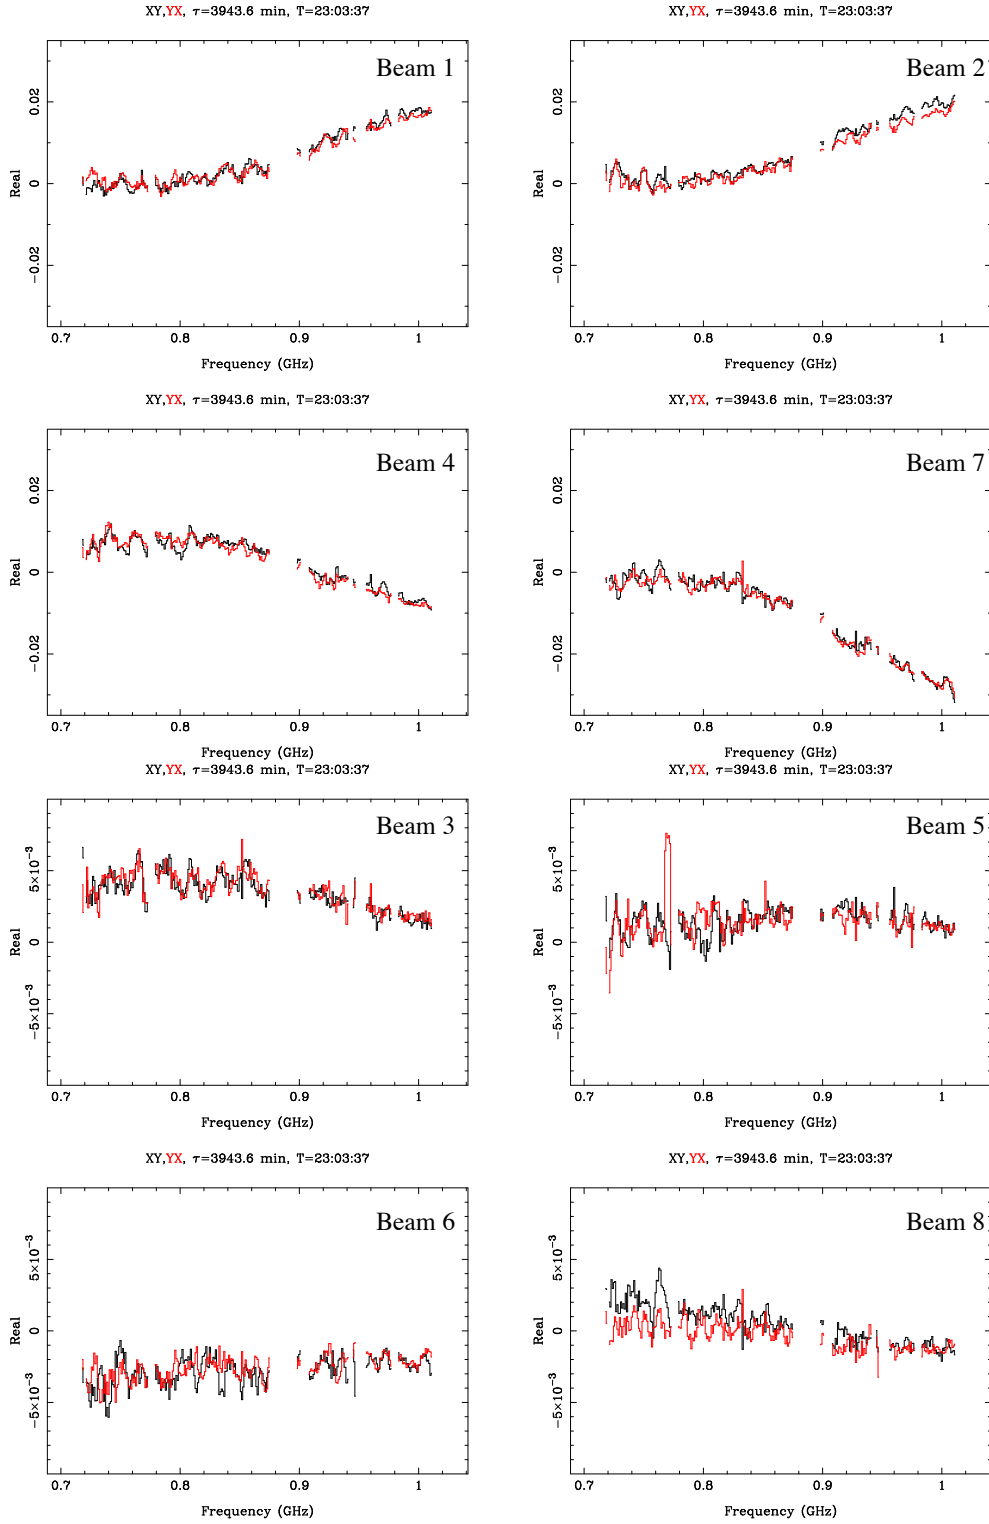
Figure 3: Response of XY and YX correlations for the eight offset beams. (a) The upper four panels: Beams 1, 2, 4 and 7. (b) The lower four panels: Beams 3, 5, 6 and 8. Note the difference in scales between (a) and (b).
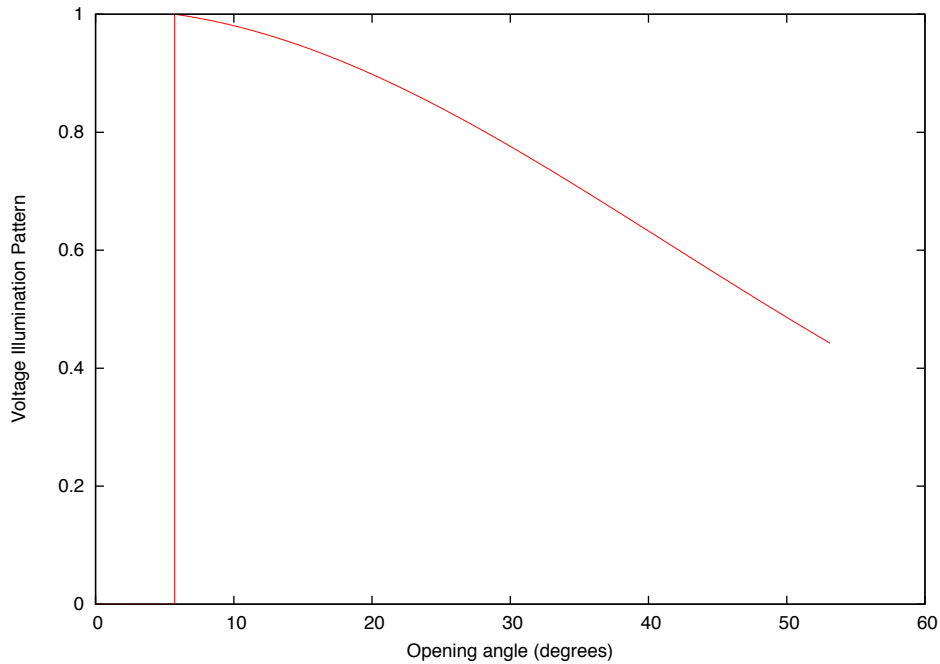
Figure 4: Model illumination pattern.

the lower panel is also reasonable. The 25 MHz standing wave response and the beam former discontinuities are readily apparent, with a peak-to-peak response of approximately 10% at the low frequency end.

Figure 6 gives the difference between the XX and YY responses, which is the Stokes Q response in the frame of the antennas (i.e. the frame relative to the X feed orientation)[1]. The expected clover leaf pattern suggests that beams 1, 2, 4 and 7 are at the zero crossing in the response, and that the response of beams 3 and 8 will have the opposite sign to beams 5 and 6. So the plot presents the negative of beams 3 and 8. Again the standing wave and beamformer discontinuities are apparent, with peak-to-peak variations of about 2% at the low frequency end. This is appreciably less than the 10% seen in the sum of XX and YY, i.e. the standing wave response is largely common to both XX and YY. It is reasonable to hypothesise that the difference between antennas would comparable to the difference between XX and YY.

Figure 7 compares the XY and YX correlations with the model prediction. These are the Stokes U response in the frame of the antennas. Again, the expected clover leaf pattern suggests beam 4 and 7 will have the opposite sign to beams 1 and 2, and that the other beams are at the zero crossing. So the negative of beams 4 and 7 have been plotted. For a given beam, the XY and YX correlations are very similar. Consequently the plots do not distinguish between XY and YX. Similarly beams 1 and 2 are very similar, and the plot does not distinguish between them.

In the upper panel of Fig 7, agreement between the model and actual response are quite poor – the model system response is too simple. Although the agreement between actual and model is poor, Figures 6 and 7 taken together show

---

[1] For the standard observing mode of ASKAP, the roll axis maintains X feed orientation to be at 45° to the meridian of an observation.
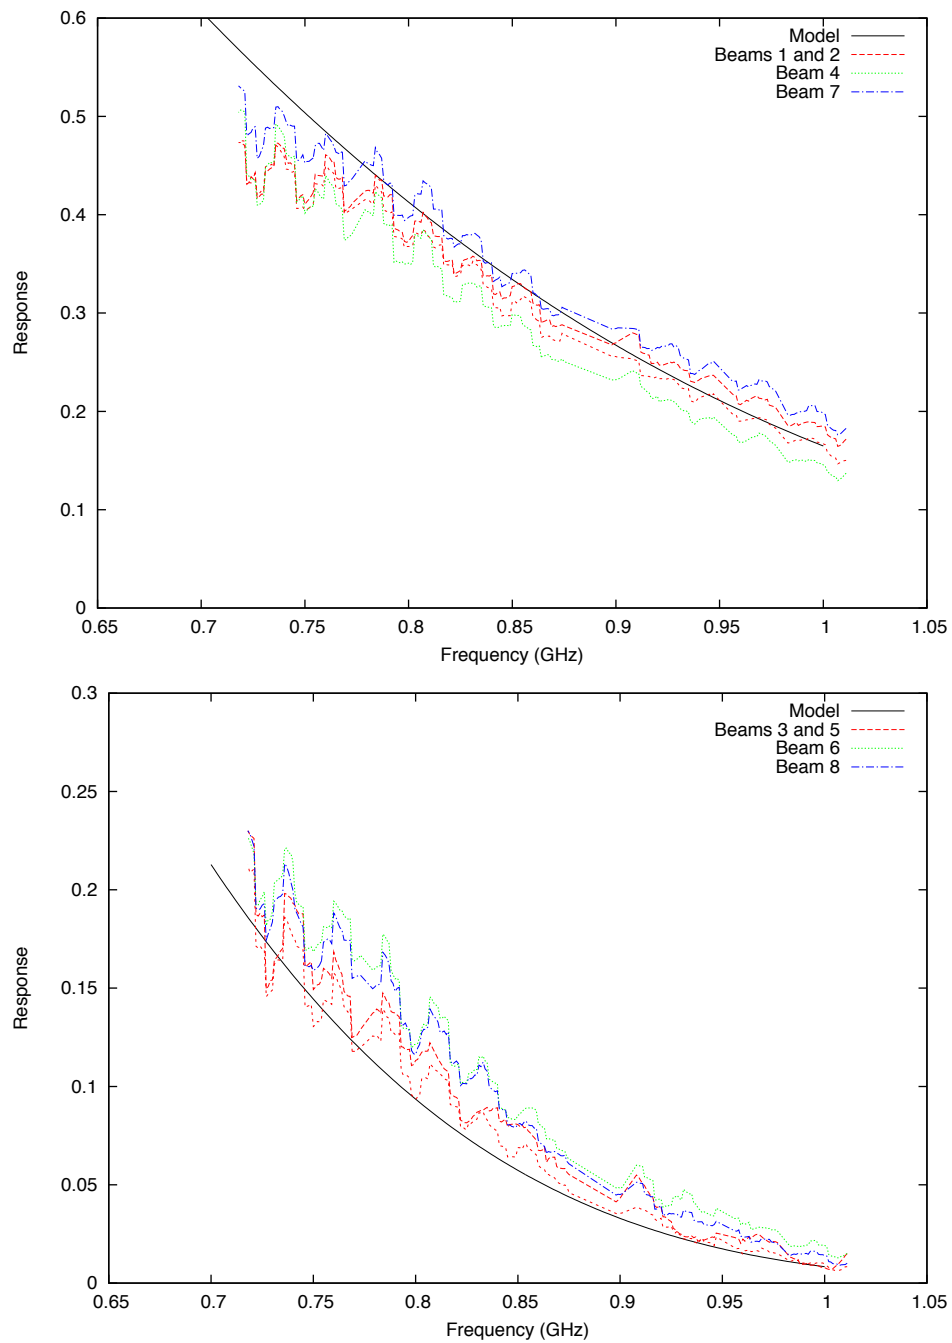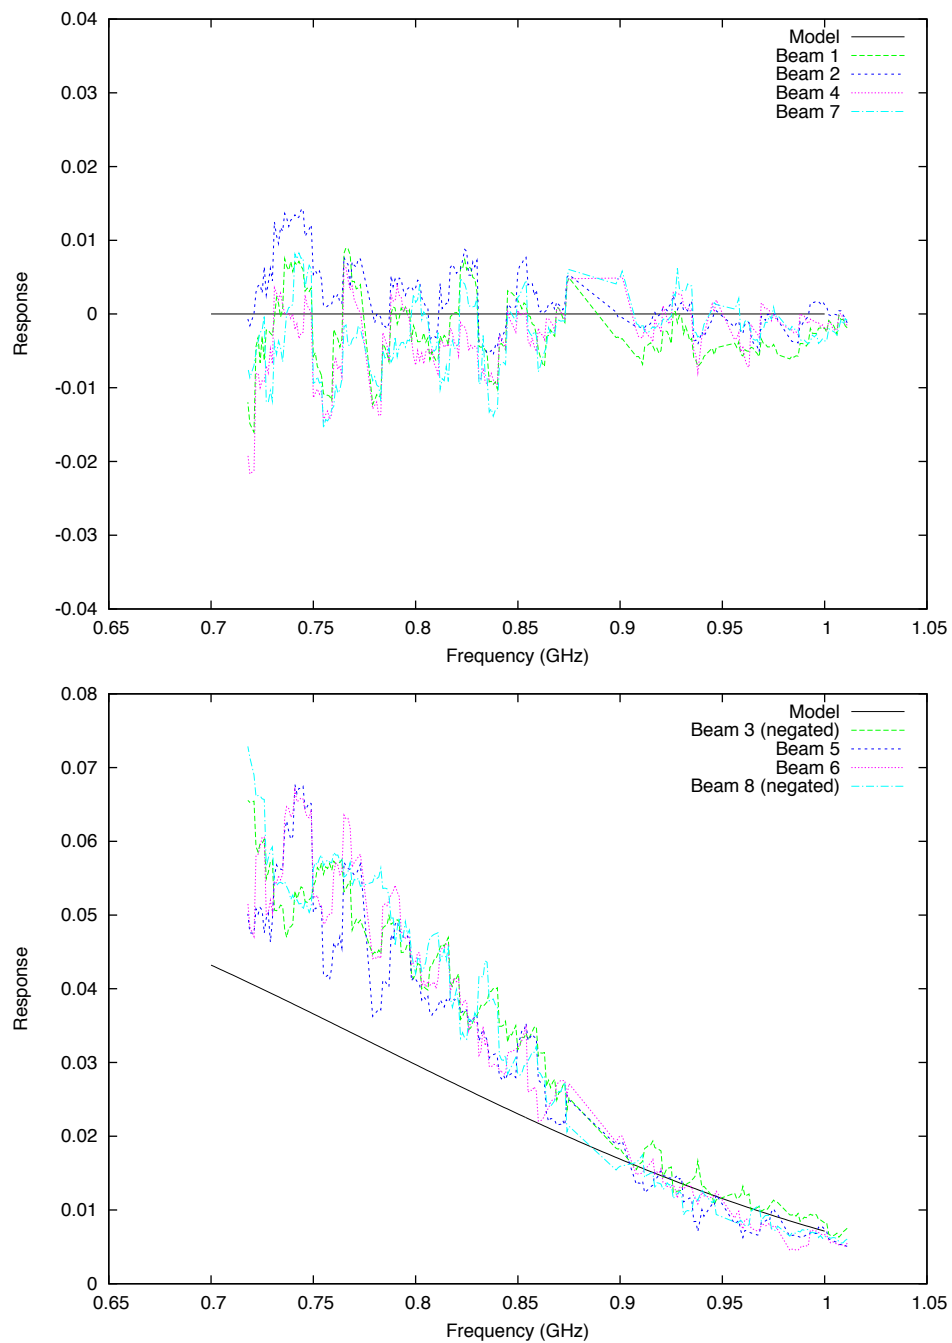
Figure 5: Stokes I ("primary beam") response.
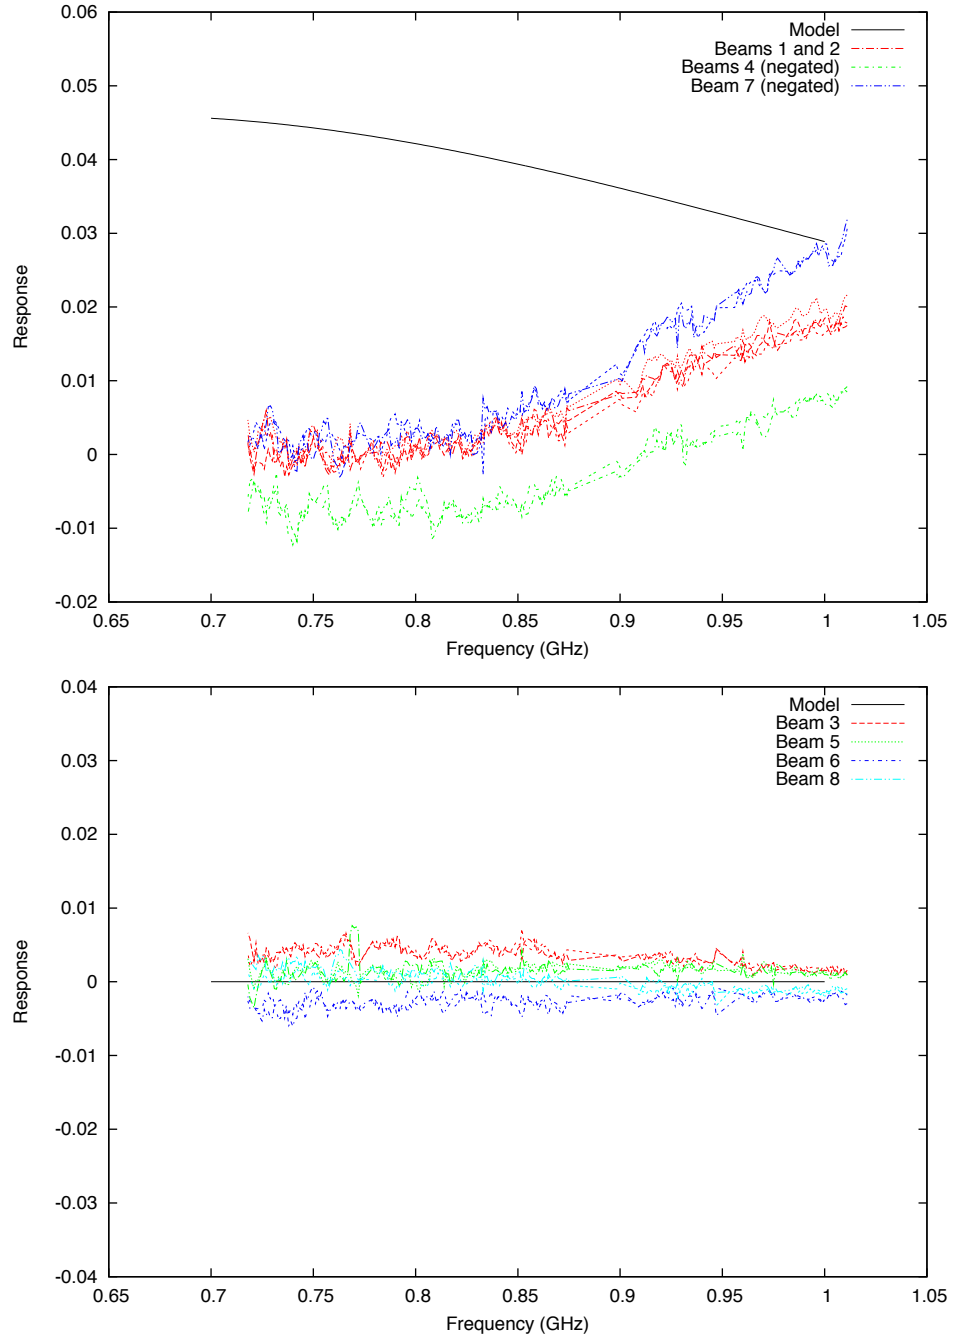
Figure 6: $XX - YY$ response.

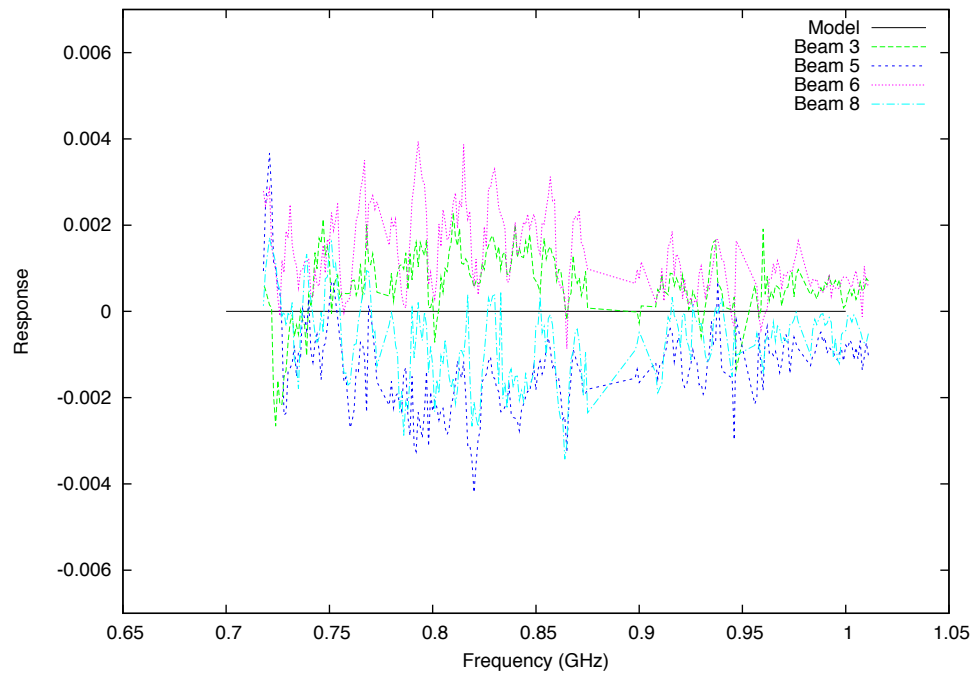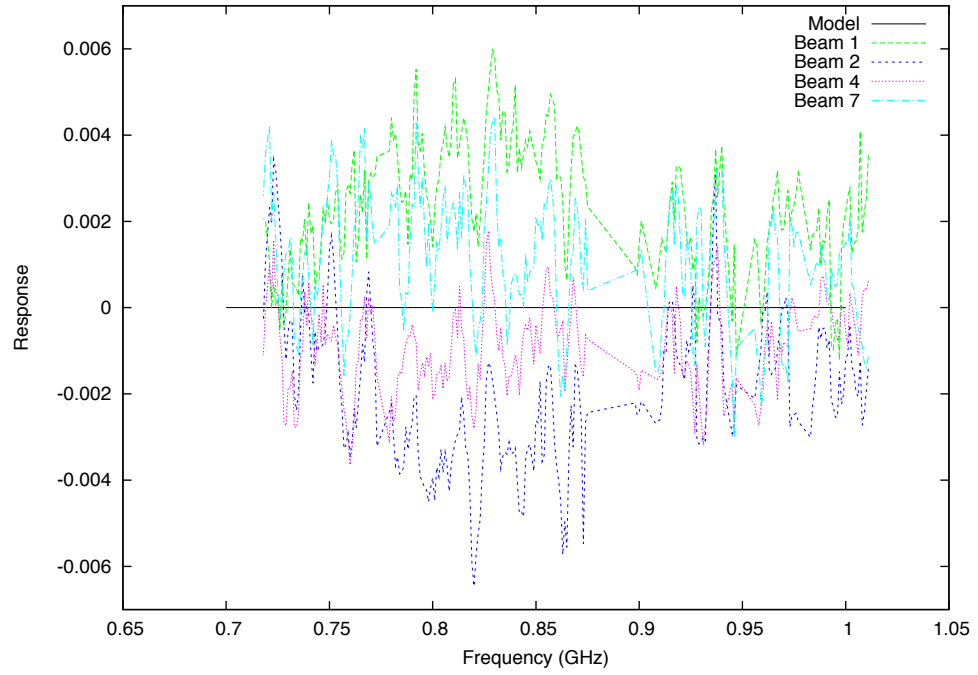Figure 7: XY and YX correlation responses.

Figure 8: Stokes V response.

that the actual response is a clover leaf pattern, although the peaks (positive and negative) of the clover leaf can differ significantly from the model. Also this clover leaf pattern does not represent a circularly symmetric polarimetric response[2].

Note, also in the upper panel of Fig 7, that while beams 1 and 2 are very similar, beams 4 and 7 differ from them and from each other. This implies a lack of circular symmetry in the response. This lack of symmetry is built into the response when the beamformer weights are determined, and is probably caused an asymmetric broadband noise field in the beamforming process. Some possible sources for an asymmetric noise field are the Sun, ground spillover or broadband PAF/telescope electronics RFI. Broadband RFI over 300 MHz seems unlikely. Recall that in the beamforming, the blank field was 15° south in declination from the Sun, and that the parallactic angle at the time was approximately 135°. Given this, the nature of the asymmetry suggests an effect in the direction of the Sun during the beamforming process, whereas a possible effect of the ground would be 45° of rotation to what is seen. This suggests the asymmetry may be related to Sun's presence in the sidelobes of the antennas during the observation of the blank field.

Figure 8 plots the real part of the Stokes V response. Stokes V equals $i(XY - YX)/2$, i.e the imaginary part of the difference between XY and YX correlations. Given the symmetry of the X and Y probes, a circularly polarized instrumental response is not expected. Put another way, there is no break in symmetry that would lead one *a priori* to expect instrumental circular polarization to be more likely to be left-handed over right-handed. The measured response is low ($\sim$0.5%): it is not clear that this is not an artefact of some calibration error in the data processing.

Widefield polarimetric calibration of ASKAP

If all antennas can be treated as having the identically same polarimetric response that are independent of time, the polarimetric correction of an ASKAP-like array (or an array with equatorial antennas) can be relatively simple[3]. This is because the deparallactification that the ASKAP roll axis provides means that the polarimetric response does not vary throughout the observation as parallactic angle changes. The instrumental polarimetric response of a source is thus independent of time and baseline, and can be corrected in the image domain after deconvolution. The correction would be implemented by a Mueller matrix, the coefficients of which would be dependent on a pixels location relative to the pointing centre. This is akin to primary beam correction being performed in the image domain[4].

---

[2] Both the quadrapod legs and the illumination pattern caused by dipole feeds are expected to produce clover leaf instrumental polarization patterns, although these will not be circularly symmetric.

[3] If the polarimetric response varies with time or between antennas, computationally more expensive techniques, such as the so-called *A-projection* algorithm may be needed.

[4] If the instrumental response is circularly symmetric, image-based correction can also be used for arrays with alt-az antennas. This would be the case for the polarimetric response of the simple model antenna used in this memo. Historically the primary beam of alt-az arrays is approximated as being circularly symmetric to allow image-based correction. Interestingly the circularly symmetry approximation of the polarimetric response of alt-az antennas does not appear to have been used in any widefield observations. Circular symmetry in the widefield polarimetric response of the VLA is a poor assumption, and the WSRT is an equatorial instrument, and so the approximate is not relevant for them. However in some of the observing

To use this approach, the polarimetric response must be known and time independent (during at least the course of an observation), antenna-independent and preferably reasonably "clean". This is not so with the current BETA system. Whereas Fig 6 shows that the majority of the standing wave response is common to XX and YY, they do differ at the $\sim$2% level. The nature of the standing wave introduces a periodicity into the polarimetric response that will likely contaminate some rotation measure experiments. In general the illumination pattern and primary beam response of the BETA system is poorly known, and the response is far from clean. A significant reason for these characteristics is that the beamforming used is the maximum sensitivity algorithm. This algorithm optimises the performance at a point – not across field of the formed beam. The off-axis response of the formed beam is no part of the optimisation. The off-axis response can therefore be poor. The off-axis effect of the standing wave, the beamformer discontinuities or the Sun in the sidelobes is not part of the optimisation. Variations in LNA performance across the PAF or between antennas will also affect beam characteristics.

This is not an issue for polarimetry only. The dynamic range achievable with mosaicing will also be affected when the primary beam pattern is poorly known, rapidly varying response or has appreciable sidelobes.

An approach to ameliorate this may be to use a beamforming algorithm which optimises the response of the formed beams to a particular shape. Notionally this would improve beam shape and made clear the form to be used in the downstream data processing. It remains an open issue as to the level to which the different beams can be made to represent a given form.

Ionospheric Faraday rotation

For ASKAP it is assumed that the ionospheric Faraday rotation is constant across the field of view of a formed beam and that it is known independently from the astronomical observation (e.g. by modelling of GPS data). Ionospheric Faraday rotation correction is a complication of the scheme to apply widefield polarimetric corrections in the image plane. This is because usually the ionosphere Faraday rotation will vary significantly over the course of an observation, at least at the low frequency end of ASKAP's band.

If the antenna-based widefield polarimetric response can be ignored, handling ionospheric effects is straightforward: propagation through the ionosphere rotates the angle of linear polarization, which can be corrected by modifying the nominal angle of the feed system with respect to the observation's meridian. This angle is used when converting the correlations to Stokes parameters. That is, rather than using the physical feed angle of 45° to the meridian (which is the normal orientation of the feeds for BETA observations), this is modified to

$$\chi = \pi/4 + RM\lambda^2.$$

Correcting for ionospheric Faraday rotation should be done before the imaging of the Stokes parameters. Otherwise a time-varying rotation between Stokes Q and U will be introduced into the images. Thus imaging needs to be done in corrected Stokes parameters.

bands of the ATCA, the dominant off-axis polarimetric response is circularly symmetric, and so there is a computationally cheap correction for off-axis polarimetry is possible. However the generally good polarimetric purity of the ATCA off-axis response makes this a less pressing issue.

When the widefield response cannot be ignored, the issue is more difficult. Because antenna-based widefield effects happen after ionospheric Faraday rotation, they should be corrected before. Strictly, the order of correction is not commutative. That is, the corruptions are first the ionosphere and then by the antenna, and so the corrections should be first for the antenna and then the ionosphere. However image-based widefield polarimetric correction is clearly after imaging, and ionospheric correction is simplest before imaging.

The issue is perhaps less severe than it might at first seem. Although the corrections are strictly not commutative, two effects make them approximately so:

- The dominant correction of the widefield Mueller matrix is the leakage of Stokes I into Q, U and V. This leakage is unaffected by the ionospheric Faraday rotation. If done in the incorrect order, ionospheric Faraday rotation will affect the correction of widefield leakage of Q and U into each other and the other quantities. This might be modest enough to be ignored.

- If the widefield antenna response were circularly symmetric (as was the simple model used for comparison in the analysis), then the corrections can in fact be commuted. Although the widefield polarimetric response for these observations were not circularly symmetric, it is possible that a different beamforming approach can make them more so.

If this approximate commutativity is insufficiently accurate, approaches involving breaking the observation into several shorter time intervals ("snapshots") can be used. Over each time interval the ionospheric Faraday rotation can in some way be approximated as being constant. These snapshots are then imaged and polarimetrically corrected separately before being combined into an overall result. Two possible "snapshot" approaches to handling the correction present themselves:

- The visibility data are corrected for the ionospheric Faraday rotation and then the Stokes parameters imaged. In making the image-based widefield correction, the mean ionospheric Faraday rotation during the snapshot is accounted for in the widefield Mueller correction matrix. The Mueller matrix would include

    - a de-rotation of Q and U by the mean ionospheric Faraday rotation back to the frame of the antenna,

    - followed by the widefield polarimetric correction,

    - followed finally by a rotation of Q and U back to the astronomical frame.

- The visibilities are not corrected for ionospheric Faraday rotation before imaging. The widefield polarimetric correction is applied after imaging, and then Q and U are rotated by the mean ionospheric Faraday rotation.

The former of these two better treats the ionospheric Faraday rotation effects, whereas the latter better treats widefield antenna effects. When the approximate commutativity of ionospheric Faraday and widefield antenna effects is considered, possibly the first of the two snapshot approaches will be better from both image quality and computation cost prospectives.

Higher precision linear polarization with ASKAP

Figures 6 and 7 give the Stokes Q and U instrumental responses in the frame of the antennas. Normally the error in linear polarization measurements of a celestial object is dictated by the combination of these two. Note that the linear scales on Figs 6 and 7 are the same (although the range varies). The variation around the trend in Fig 6 is several times larger than that around Fig 7. That is, the systematic errors in Q in the antenna frame is appreciably larger than in U. This is usual in systems with linearly-polarized feeds, and indeed is one of the reasons that are put forward for the superiority of systems with circularly-polarized feeds. The larger systematic errors in the antenna-frame Q may limit the polarimetric purity.

It is interesting to note that ASKAP can avoid using the problematic $XX - YY$ combination. Ignoring ionospheric Faraday rotation for the moment, consider two ASKAP observations. The first could use the standard BETA deparallactification, where the X feed is set to be at $45°$ to the meridian, and then the second could deparallactify with the X feed at $0°$ to the meridian. In the first observation, the XY and YX correlations would measure astronomical Stokes U, and in the second observation they would measure astronomical Stokes Q. Although this avoids using $XX - YY$, the cost is a lost in raw sensitivity is a factor of $\sqrt{2}$ in Q and U compared to if $XX - YY$ had been used. This may not be an issue if the systematics in $XX - YY$ would have dominated over raw sensitivity. This is an observing strategy that should be considered where high purity polarimetry is needed, particularly if the source is not expected to be variable and if multiple observing epochs are required for other reasons. Depending on the data processing pipeline, it may be possible to decide whether or not to use $XX - YY$ after the observation.

Ionospheric Faraday rotation effects are again an issue with this scheme. When a time-varying ionosphere is also considered then the ionosphere, as well as the orientation of the PAF, dictates what combination of Stokes Q and U are measured by the XY and YX correlations. This can be handled as follows:

If χ is the angle between the X feed and the meridian (χ_0), plus any ionospheric Faraday contribution

$$\chi = \chi_0 + RM\lambda^2$$

then

$$XY + YX = U\cos 2\chi - Q\sin 2\chi.$$

Provided χ can be varied over an appreciable range (either resulting from ionospheric variability, or from actively changing the PAF alignment, χ_0), then it will be possible to decouple Q and U from $XY + YX$. To do this, consider forming the two visibility sets:

$$
\begin{aligned}
V_c &= (XY + YX)\cos 2\chi &= U\cos^2 2\chi - Q\cos 2\chi \sin 2\chi, \\
V_s &= (XY + YX)\sin 2\chi &= U\cos 2\chi \sin 2\chi - Q\sin^2 2\chi.
\end{aligned}
$$

Image these two sets to form two images

$$
\begin{aligned}
I_c &= U * B_c - Q * B_{cs}, \\
I_s &= U * B_{cs} - Q * B_s,
\end{aligned}
$$

where B_c, B_s and B_{cs} are point-spread functions ('dirty beams') formed with weights of $\cos^2 2\chi$, $\sin^2 2\chi$ and $\cos 2\chi \sin 2\chi$ respectively. These two images can then be deconvolved with these three point-spread functions. This can be

done by a process that is similar to decoupling spectral index in multi-frequency synthesis deconvolution. It is apparent that if χ were 0 and 45° (e.g. setting two PAF orientations and no ionospheric Faraday rotation), B_{cs} becomes zero, B_c and B_s become standard weightings of the two sets of data the two deconvolution equations decouple.

Clearly if ionospheric Faraday rotation needs to be considered, avoiding using $XX - YY$ becomes a good deal more involved and so may not be justified.