Cross-calibration of VAST pilot survey epochs 8 & 9: tests on a small sample of data

June 29, 2020

Mark A. Walker (Manly Astrophysics)

1 Background

LST-matched observing mitigates the influence of steady, but unmodelled flux in the primary beam and may therefore be desirable in precision studies of variable radio sources. Combining calibrated, LST-matched visibilities is a means of increasing the signal-to-noise of data without increasing storage requirements and may therefore be attractive for spectral line imaging programs. Consequently the use of LST-matched observing strategies may help to achieve commensality in ASKAP surveys.

This study explored novel methods of calibration for LST-matched data, attempting to cross-calibrate one epoch from the other — i.e. without the overhead of imaging-plus-deconvolution that is needed for a self-calibration. All cross-calibration attempts reported here used purely antenna-based complex gains. Time- and frequency-dependent models, and direction-dependent models were considered.

Data for one beam of one field were supplied by Emil, exported to uvfits by Keith, and further exported to ascii files by Tyoma - thanks Emil, Keith and Tyoma. Subsequent processing was by MW in Mathematica using code written specifically for the job.

Notation: $V_{ij}(8,9)$ are the visibilities between antennas i and j for Epoch 8,9, and g_i is the complex gain of antenna i. The gains may simply be constants, with one complex number per antenna, or else linear functions of frequency, ν, and time, t (three complex numbers per antenna), or else linear functions of right-ascension, α, and declination, δ (three complex numbers per antenna).

Phase centre: $(\alpha, \delta)_{2000} = (39.8182°, -5.7745°)$, Equatorial; $(l, b) \simeq (178°, -56°)$, Galactic.
Zenith Distance: $23°$
Hour Angle: $2^h 35^m$
Integration time: $12\ \mathrm{min}$
Observing frequency: $888\ \mathrm{MHz}$ (band centre)
Observing epochs: JD2458866 and JD2458867 (offset of 1 sidereal day)

2 Pre-processing

Emil applied flagging and bandpass calibration prior to export.

I flagged the data at 5-sigma, based on the statistics of the XY polarisation products. A robust estimate of sigma was used: the root-median-square, divided by $\log_e 2$ so that its expectation corresponds to the root-mean-square. Flagging was matched across both epochs, meaning that if a sample was flagged at one epoch then the corresponding sample was flagged at the other epoch. A consequence of matched flagging is that the impact of bad data is approximately twice as large in the matched dataset: only 60% of data survived (different antennas were "bad" on different days).

Although the overall observing windows were LST-matched, the correlator sampling times were not aligned, with the offset being approximately 0.63 samples for these data. The visibility estimates for Epoch8 were therefore interpolated to match the sampling times of Epoch9. This procedure leads to some loss of signal amplitude on long baselines, and to correlated noise in adjacent samples; it would be better to start with samples that are precisely synchronised in LST.

3 General approach

All gain models included both amplitude and phase (phase-only models lead to non-convex optimisations), and solutions were obtained by minimising a figure of demerit that measures the sum of squared-modulus-of-differences between one epoch and the other. Minimisation was done with the "FISTA" algorithm of Beck and Teboulle (2009)[1] with the regularisation parameter set to zero (i.e. no l_1 contribution to the figure of demerit). For each optimisation, a good estimate of the value of the Lipschitz constant could be made beforehand, so it was not necessary to employ backtracking. Optimisation problems that are convex are guaranteed to converge rapidly with this algorithm. Wirtinger derivatives were used to construct the gradient, because we're minimising a real function of complex variables, $D(g_i)$, which is necessarily a non-analytic function. Separate gain solutions were obtained from the XX and YY polarisation products; XY and YX products were not used. A restarting procedure was employed in all cases, to avoid the oscillations in D (vs. iteration number) which are inherent in the convergence dynamics of FISTA. Restarts can be identified as abrupt changes in the gradient of D as a function of iteration number — e.g. near iterations 20 and 90 in figure 4. In all cases the antenna phases are measured relative to antenna $\#1$, which is near the centre of the array.

4 Fitting in the uv-plane

The simplest possible approach was tried first: define a residual

$$R_{ij} = g_i g_j^* V_{ij}(9) - V_{ij}(8), \tag{1}$$

and a figure of demerit

$$D = \sum_{i,j} R_{ij}^* R_{ij}, \tag{2}$$

then find the set of constant gains $\{g_i\}$ that minimise D. I note that D is convex. Unfortunately this simple approach introduces a large noise bias in the resulting solutions. The reason is that there are two contributions to the residuals: one is from the mismatched signal – we aim to eliminate this contribution by the correct choice of antenna gains – and the other is from the measurement noise. The Epoch 9 component of the latter contributes to D in proportion to the squared-moduli of the

[1] SIAM J. Imaging Sci., 2 (1), 183

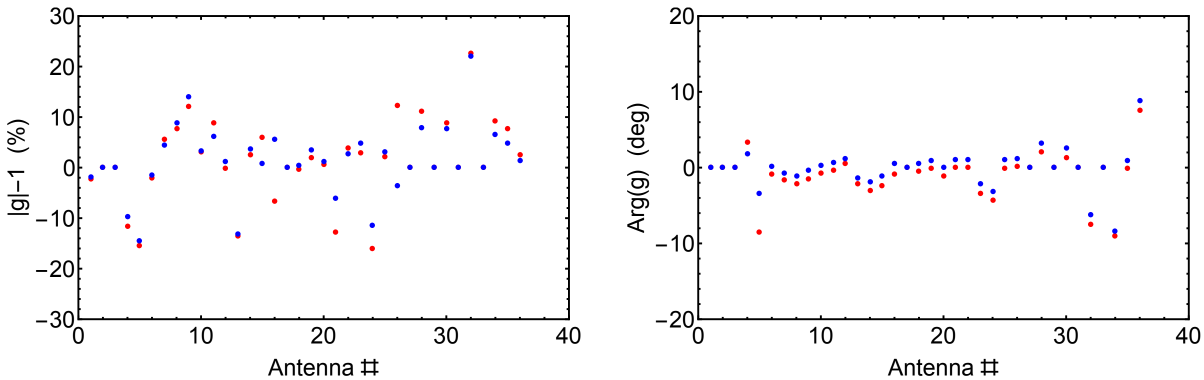

Figure 1: Amplitude (left) and phase (right) of the constant gain solutions obtained by minimising equation (3). Solutions for the X (Y) signal chain are shown in red (blue). Seven antennas ($\#2, 3, 17, 27, 29, 31, 33$) had no valid data and their gains are set to unity.

model antenna gains, thus causing the minimum of D to be displaced towards smaller $|g|$ relative to the solution that is sought. Because the signal-to-noise ratio on the individual visibilities themselves is only of order unity, this bias is a large effect.

The problem just described can be mitigated by using a different figure of demerit that removes the expected noise bias:

$$D = \sum_{i,j} R_{ij}^* R_{ij} - |g_i|^2 |g_j|^2 \sigma_{ij}^2(9), \tag{3}$$

where $\sigma_{ij}^2(9)$ is the sample variance for the baseline (and time, frequency, and polarisation) under consideration. It is important to have a precise estimate of the noise, because the noise bias has an influence of order unity on the derived gain amplitudes, so the antenna autocorrelation signals were used to calculate the variance for each sample entering the sum. Although the demerit specified in equation (3) deals with the noise bias issue, the additional term is non-convex and therefore has the potential to cause problems in finding the global minimum of the demerit. The resulting gain solutions showed surprisingly large amplitude deviations from unity, exceeding 20% in one case — as shown in figure 1.

The solutions exhibit a clear correlation between the X and Y gains, both in amplitude and phase. The influence of the troposphere and ionosphere should be highly correlated between polarisation modes, so a correlation between phases is unsurprising. I note that figure 1 also shows an offset of approximately 1° between the phases of the X and Y gain solutions; that is readily understood, as the zero-point of the phases is not absolute — it is defined by the phase of antenna $\#1$, and can thus differ between the two polarisation modes.

The origin of the correlated amplitude structure is less obvious. My starting assumption was that the amplitude of the gains would likely be dictated by the amplifiers, but I'm unclear on why they're so well correlated. In any case the high level of amplitude structure displayed by these solutions suggests to me that these solutions are not credible, so I have not attempted imaging/photometry using this calibration.

A second, more general attempt was made at calibrating the visibilities without reference to the structure of the field being observed. In this second attempt I allowed the gain for each antenna to be a linear function of frequency and time, so that

$$g_i(\nu, t) = g_i(\nu_0, t_0) + (\nu - \nu_0)\, g_{i,\nu}(\nu_0, t_0) + (t - t_0)\, g_{i,t}(\nu_0, t_0), \tag{4}$$

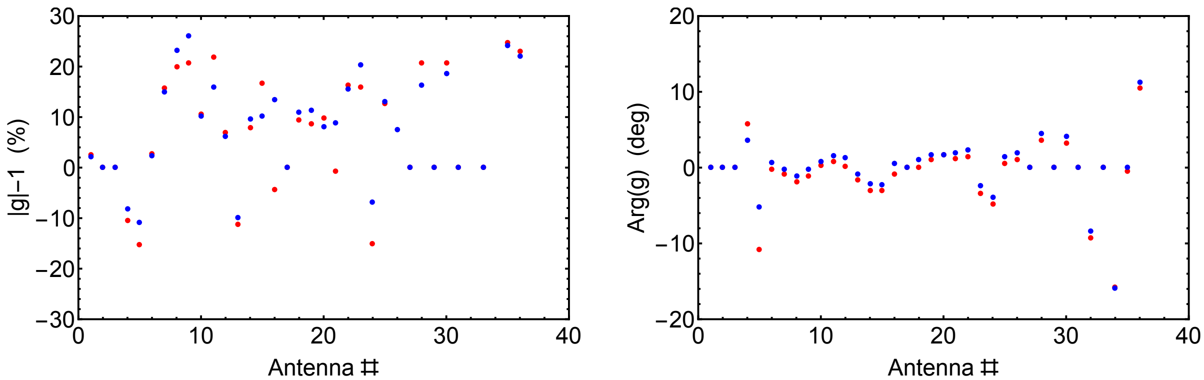

Figure 2: Amplitude (left) and phase (right) of the constant terms, $g_i(\nu_0, t_0)$, in the gain model given in equation (4), obtained by minimising equation (3). Solutions for the X (Y) signal chain are shown in red (blue). Seven antennas ($\#2, 3, 17, 27, 29, 31, 33$) had no valid data and their gains are set to unity. Note that the amplitudes of antennas $\#26(X), 32$ and 34 are off-scale in this plot.

where $g_{i,\nu}$ and $g_{i,t}$ are the partial derivatives of the gain of antenna i with respect to frequency and time. I then minimised the demerit (3) to obtain the 3×36 complex numbers that describe the antenna gains. (The two additional parameters that were optimised were scaled versions of the derivatives, with the scaling chosen so as to make the curvature of the demerit surface approximately isotropic. This step, known as "preconditioning" the optimisation, helps to achieve a rapid convergence.)

The constant part of the solution, i.e. $g_i(\nu_0, t_0)$, obtained in this second optimisation is displayed in figure 2. There we see that the overall pattern of antenna phases is similar to that of figure 1, but with noticeable differences between the two solutions. The amplitudes are very different from those of figure 1, and display even larger deviations from unity. Consequently these solutions also lack credibility and have not been taken to the stage of imaging/photometry. As with the first set of solutions, it is unclear why this approach failed. One possibility is that there are significant direction-dependencies in the antenna gains, which cannot be compensated in the scheme I've used; however, the results of the next section suggest that the direction-dependent terms in the antenna gains are not playing a major role. Another possibility is that the non-convex term in the demerit function has caused problems in finding the global optimum; that hypothesis is consistent with the fact that the results of this optimisation appear to be worse than those of our first attempt (where the parameter space is only $1/3$ the size). Finally I note that the noise bias correction that I've employed is based on an expectation value, whereas the true influence of the noise bias will actually differ from that expectation. Given that the noise bias is large when we are fitting directly to the uv data, that difference could introduce substantial systematic errors in the gain amplitudes.

5 Calibration via recentred summation

The principal difficulties encountered in the previous section are associated with noise bias and the steps taken to deal with that bias. Absent a sky model we are stuck with noise bias at some level, but we can decrease its influence if we restrict attention to those regions of the sky where there are strong sources. If we were starting from scratch, with no prior information on the field, this path would take us back to the conventional steps of imaging, deconvolution and self-calibration. However, given a catalogue of bright sources we can instead proceed as follows.

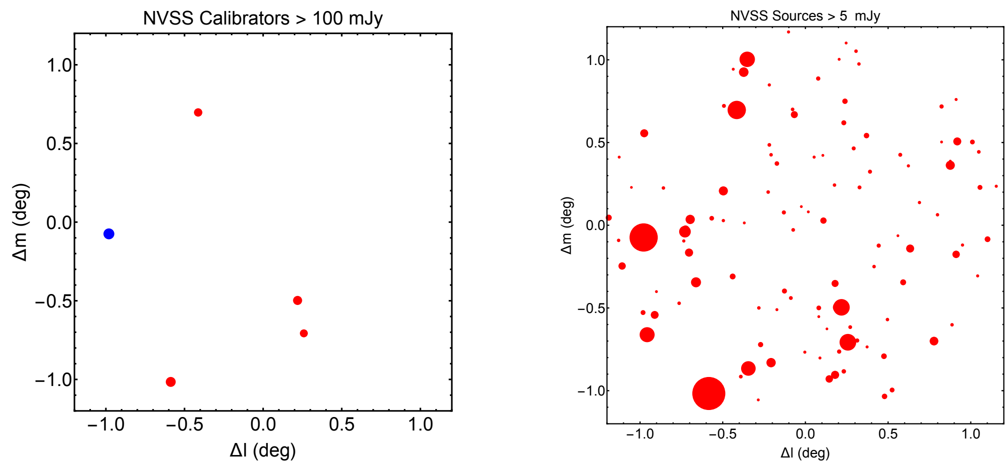

Figure 3: Left panel: location of the five NVSS sources expected to make the largest visibility contributions (all expected to be above $100\ \mathrm{mJy}$) to our test data. The blue source is PKS J0243-0550, which was discovered to be vary between Epochs 8 and 9, and was therefore dropped from the list of calibrators. Right panel: Locations of all 114 NVSS sources with catalogue fluxes above $5\ \mathrm{mJy}$, and within $1.2°$ of the pointing centre; in this panel point size indicates catalogue flux. For both panels North is up and East to the left.

A compact source at the phase centre should yield a very simple set of visibilities: $V_{ij}(\nu, t) = F$. (For brevity I'm assuming a flat spectrum source in this argument, but the validity of the method does not depend on that assumption.) In practice, of course, the antenna gains are not unity so the measured visibilities are actually $V_{ij}(\nu, t) = F\, g_i g_j^*$, and a simple way of calibrating is to form $\sum_{\nu,t} V_{ij}(\nu, t)$ and look for the best-fit to an outer-product structure (i.e. the antenna gain vector multiplied by its complex conjugate), with a scaling factor that is the source flux. In our case the bright sources are not located at the phase centre, but we can deal with that by rephasing the visibilities so that the phase centre does correspond to a bright source, and then forming the sum over frequency and time. This process can be repeated for each of the N_{cal} sources that we want to use for calibration, leading to a set of N_{cal} matrices, each of size 36×36,

$$B_{ij}^k \equiv \sum_{\nu,t} V_{ij}^k(\nu, t),$$

where k specifies the source being used for calibration, and V^k denotes a visibility that has been rephased so that the phase centre corresponds to the calibrator position. Summing over time and frequency tends to suppress the contribution of every source that is not at the phase centre, because the visibilities of those sources are oscillatory in time and frequency.

For my investigation I chose the five NVSS sources which, after applying an approximate primary beam attenuation, were expected to be the brightest in the field; all are expected to be brighter than $100\ \mathrm{mJy}$ in the VAST data. Their distribution over the field is shown in figure 3. Most of these sources are not sufficiently compact that they could be usefully modelled as point-like to ASKAP (see later for a counterexample). However, they can all be used as templates for evaluating the relative gains of the two epochs, just as in the previous section but using the summed, rephased visibilities instead of the raw visibilities. Thus instead of the residual given in equation (1) we now have

$$R_{ij}(k) = g_i(k) g_j^*(k) B_{ij}^k(9) - B_{ij}^k(8), \tag{5}$$

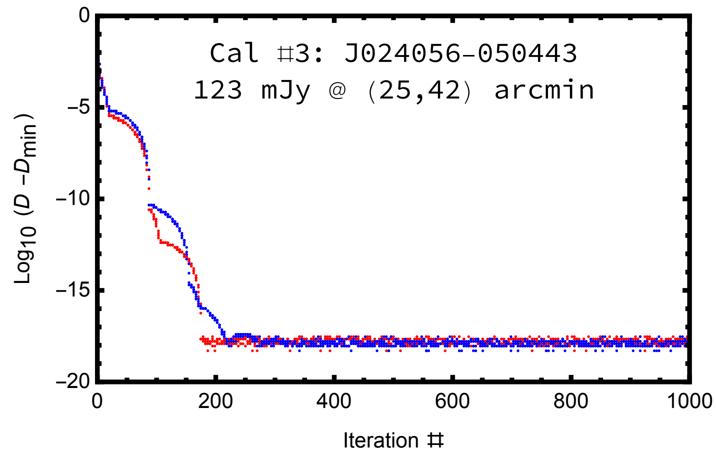

Figure 4: Progress of the gain optimisation for the most point-like source in the field of view: the (log of the) demerit relative to the final demerit is shown in red (blue) for the X (Y) polarisation. The limits of machine precision are quickly reached.

where the g_i are now additionally labelled by the calibrator under consideration, to allow us to investigate gain variations across the field. By using the B_{ij}^k, instead of the raw visibilities, the influence of the noise bias is reduced by a large factor (approximately 140) and I therefore chose to minimise a (convex) demerit analogous to equation (2),

$$D(k) = \sum_{i,j} R_{ij}^*(k) R_{ij}(k), \tag{6}$$

when solving for the relative antenna gains toward each of the calibrator sources.

Because my aim was to solve for the gains in specific directions, the summation in equation (6) was restricted to a subset of baselines whose lengths in the (u, v) plane are in the range $4{,}000$ to $15{,}000$ wavelengths. The lower limit was imposed to limit the influence of source confusion, which is large for short baselines, and which interferes with the desired directional specificity. (Each baseline can be thought of as a one-dimensional array in its own right, for the purposes of continuum imaging – because of the wide range of radio-frequencies that are sampled – with an instantaneous beam that is narrow in the direction parallel to the baseline, but almost as wide as the primary beam in the perpendicular direction.) The longest baselines were excluded in order to minimise the effects of the mismatch in sampling times between the two epochs. All antennas have some baselines within the required range so relative gains for all antennas can be determined. Figure 4 shows the progress of the optimisation for the most point-like of the chosen calibrators, and figure 5 shows the gain solutions for that calibrator obtained for each antenna.

The corresponding details for the other four calibrators are not shown here because they are qualitatively similar to the behaviour seen in figures 4 and 5. The most striking feature of the solutions is that the phases are very well correlated between the two polarisation states, as expected if they're determined by propagation through the ionosphere/troposphere. The range of phase variation is, however, modest. There is also some degree of correlation between the amplitudes of the X and Y gains. Importantly, all the gain amplitudes show only modest deviations from unity, and are thus much more plausible than those described in §4.

It is important to be clear on the meaning of the solutions obtained in this way: they are the gains that best match the change in the visibilities between the two epochs, and consequently they should provide a good photometric match. However, as neither epoch has been calibrated for imaging

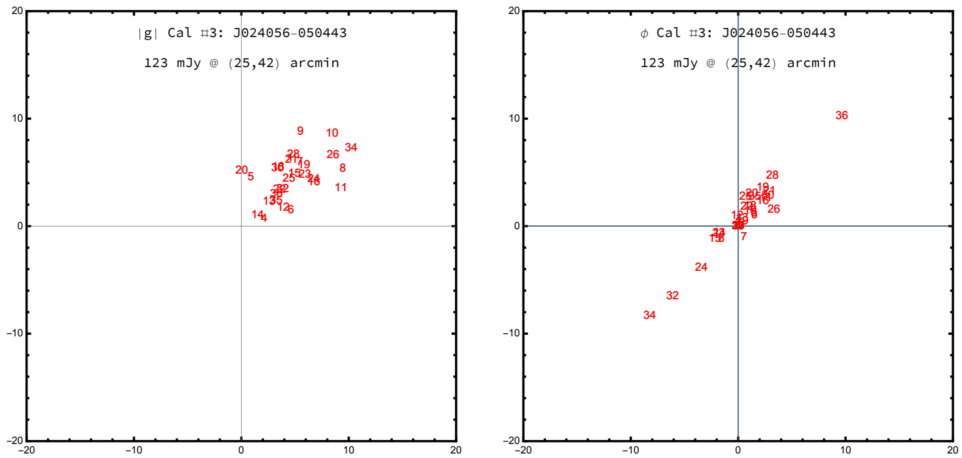

Figure 5: Optimised antenna gains for the most point-like source in the field. Left panel: the amplitude of the relative gain for each antenna (as numbered), showing the deviation from unity, in percent, with the X (Y) polarisation on the x- (y-) axis. Right panel: same, but for the phase (in degrees), relative to antenna $\#1$.

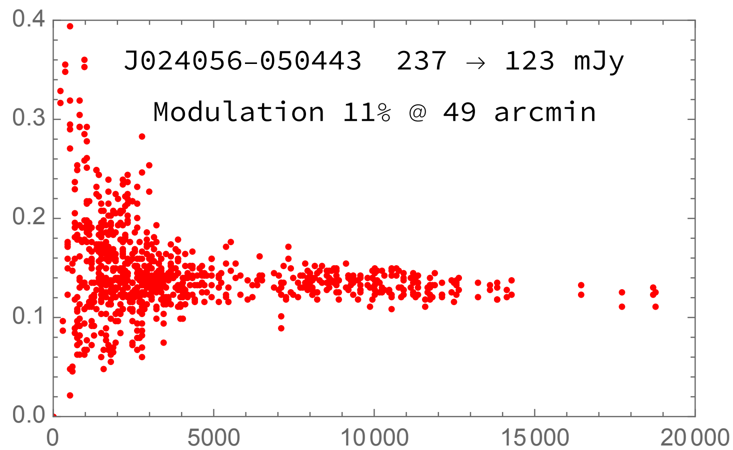

Figure 6: Recentred-and-summed (actually, averaged) visibility amplitudes versus baseline length (in wavelength units) for the calibrator that appears most point-like. This source has a catalogue flux of $237\ \mathrm{mJy}$ in NVSS, and is expected to contribute approximately $123\ \mathrm{mJy}$ to our visibilities. Confusing sources dominate the visibilities on short baselines.

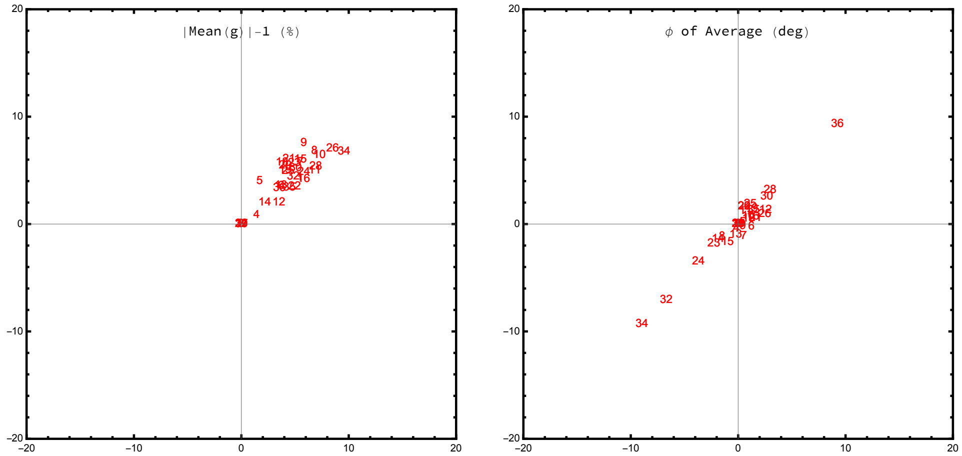

Figure 7: Antenna gains averaged over four calibrator sources; layout as per figure 5

purposes, an image formed from the differenced visibilities will exhibit aberrations in the mapping of any residual flux. These aberrations could be large and therefore need to be determined by a separate calibration step. Fortunately one of the five calibrators in the field is fairly compact, as evidenced by a small dispersion (11%), and small gradient, in the recentered-and-summed visibility amplitudes (B^k_{ij}) as a function of baseline length — see figure 6. A calibration for imaging was determined by fitting those visibilities to a point-source model, over the same range of baseline lengths specified above; this calibration was applied to the visibility differences prior to imaging the residuals. In subsequent discussion we will not make further reference to this calibration-for-imaging step, as it is conventional. All images were made with natural weighting, to maximise point-source sensitivity.

The fact that each of the five calibrators gave solutions similar to those in figure 5 indicates that the relative antenna gains – in effect, the changes between the two epochs – are not strongly direction-dependent. This is encouraging because it suggests that the instrument itself is fairly stable, despite the complexity of the PAF. It also suggests that whatever direction-dependence there is could perhaps be approximated as a small correction on top of a constant gain. In what follows I've considered two distinct cases: first is the approximation of a direction-independent relative gain for each antenna; second is the more general case of a gain that is a constant plus a uniform gradient across the field, i.e.

$$g_i(\alpha,\delta) = g_i(\alpha_0,\delta_0) + (\alpha-\alpha_0)\; g_{i,\alpha}(\alpha_0,\delta_0) + (\delta-\delta_0)\; g_{i,\delta}(\alpha_0,\delta_0). \tag{7}$$

These two models were fitted separately to the set of relative gains for each antenna (and each polarisation) in the directions towards the calibrators shown in the left panel of figure 3.

5.1 Direction-independent gain model

By averaging the gain solutions for the five calibrator sources I obtained a direction-independent gain model, applied that model as a relative calibration and imaged the result. This immediately showed that the one of the calibrators varied significantly in the 24 hours between observations, at a level (4%) that would significantly bias the average gains. Consequently I removed that source from the set of calibrators and subsequently worked exclusively with the four remaining calibrators, for both direction-independent and direction-dependent models.

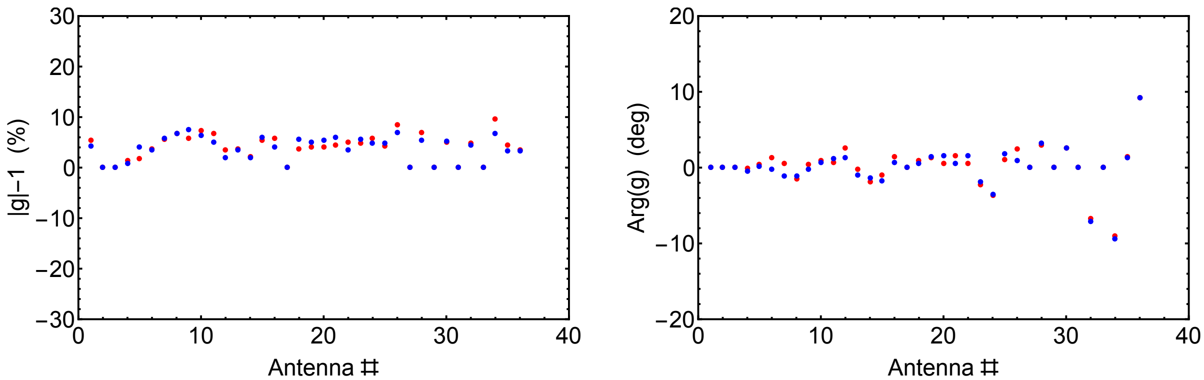

Figure 8: As figure 1, but for the gains shown in figure 7.

Averaging the gains of those four calibrator sources yields the direction-independent gain model shown in figures 7 and 8 (different representations of the same information). Applying this gain model to the Epoch 9 visibilities and imaging the result allows us to test the success of the relative calibration, through photometry of the sources in the field and comparison of these fluxes with those derived from Epoch 8. Suitable sources were taken from the NVSS catalogue, applying a catalogue flux cutoff of $5\ \mathrm{mJy}$ (lower limit) and restricting attention to sources within 1.2° of the pointing centre; there are 114 such sources, as shown in the right-hand panel of figure 3. ASKAP beam centres are separated by approximately 0.9°, so our check sources cover an area that is larger than strictly necessary.

I used the following photometric method for all sources. The nominal source positions in NVSS were used to extract a postage stamp of 32×32 pixels (approx 2 arcmin on a side) around each check source. The location and value of the peak intensity was then determined from the image at one epoch, and the intensity in the same pixel was also determined for the other epoch. The differences between these measures reflect a combination of: real changes in source flux, in the case of some compact sources; errors in the calibration model relating the two epochs; and, thermal noise. The latter was determined from the statistics of the difference image, which should be almost entirely noise if we've done a good job with the calibration. Figure 9 shows the postage stamps for the brightest 64 sources (down to about $10\ \mathrm{mJy}$ in catalogue flux); figure 10 shows average and difference images of the field out to $\pm 1.2^\circ$ from the pointing centre.

The results of the photometry are shown in figure 11. Taken together, the standard deviation of the flux differences for the 114 NVSS check sources is 2.64 — only slightly larger than the standard deviation of 2.49 that applies to the difference image as a whole. The fact that on-source pixels have only a small additional variance points to a low level of residual calibration error. To estimate the size of that error we computed the median value of $\alpha^2 = ((\Delta F)^2 - m^2)/F^2$, where ΔF is the measured flux difference for each source, of mean flux F, and m^2 is the median value of the square of all pixels in the difference image. The result is $\sqrt{\langle \alpha^2 \rangle} = 0.0018$, and the corresponding curve as a function of source flux is shown in the right panel of figure 11.

One source stands out in figure 11 (and in figures 9 and 10) as a likely variable source: NVSS J024312-055055 (PKS J0243-0550, a flat-spectrum quasar at $z \simeq 1.8$), which has a catalogue flux of $562\ \mathrm{mJy}$, has a measured flux difference in excess of 30 times the standard deviation of the difference image and more than 20 times the estimated residual calibration error. It is the second most compact of the five sources originally selected for calibration purposes,[2] as gauged by the root-

[2]This source was dropped from the list of calibrators, once it was discovered to be variable.

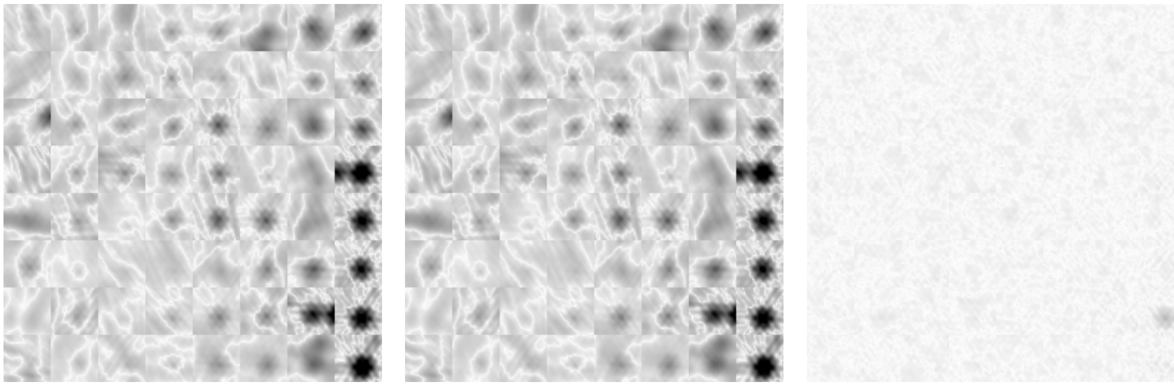

Figure 9: Compilation of postage stamps centred on the brightest 64 NVSS sources ($\gtrsim 10$ mJy). In many cases the fainter NVSS sources are not obvious amidst the sidelobe structure from the brighter ones. Left panel: Epoch 8. Middle panel: Epoch 9. Right panel: difference image. The transfer function is the same in each panel.

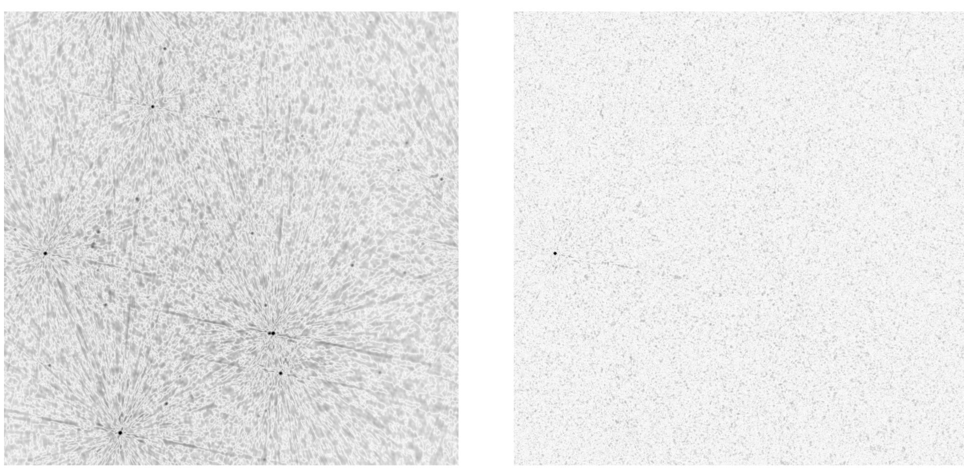

Figure 10: Average (left panel) and difference images of the two epochs, obtained with the direction-independent gain model of §5.1. Each image is $2.4°$ on a side; the grey-scale transfer function is the same for both. Only one source stands out in the difference image: PKS J0243-0550.

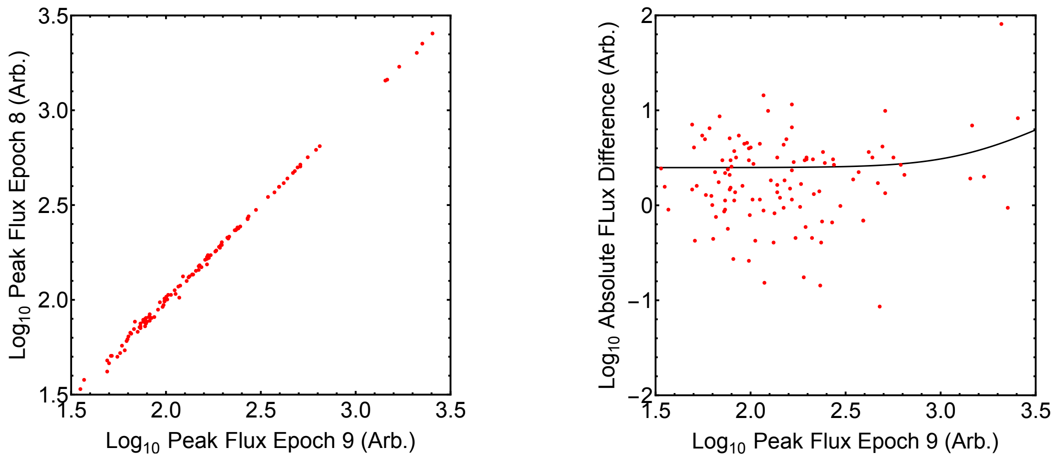

Figure 11: Fluxes of 114 NVSS sources (> 5 mJy) derived from the optimum direction-independent calibration of §5.1. Left panel: Epoch 8 versus Epoch 9. Right panel: absolute differences versus Epoch 9. Also plotted in the right panel is the standard deviation of the difference image, with a residual calibration error of 0.18% added in quadrature.

mean-square amplitude of the summed, rephased visibilities for baselines whose length is within our selection window. However, the measured difference is only 4% of the mean flux, making it an unremarkable variable.

The second-most significant change ($5.8\,\sigma$) is seen in NVSS J023833-053151 (catalogue flux: 8.5 mJy). At $18'$ this source is closer to the field centre than any of the five calibrators, so it should not have an unusually large calibration error, and it is more than a degree from PKS J0243-0550 so there should not be much contribution from sidelobes of the flux variation in that source. Although the change in J023833-053151 is less statistically significant than that in PKS J0243-0550, it is potentially more interesting as it amounts to 12% of the source flux (perhaps more, if temporally undersampled) in 24 hours.

There are several other sources whose flux changes are only marginally statistically significant ($< 5\,\sigma$); I have not examined those cases.

5.1.1 Interpretation of the measured gain changes

The good photometry obtained for the optimum direction-independent gain model suggests that it has captured the majority of the changes between the two epochs, so it is appropriate to consider the physical origin of those changes.

Antenna phase changes are usually attributed to the atmosphere, and that seems likely here because the phases seen in figure 7 are very tightly correlated between X- and Y-polarisations. Anyone familiar with the layout of the array will also notice that the Northernmost antenna (#36) has the most positive phase (measured relative to antenna #1), whereas the most negative phases belong to antennas #32 and #34, which are Southern outriggers. This pattern suggests a difference in the refractive shift between the two epochs, so I compared the observed phases to that pattern and found a best fit for a refraction of $0.67''$ at an azimuth of 9.4° (N→E). Residual phase errors, relative to that model, are as shown in the right-hand panel of figure 12; the root-mean-square of those residuals is 0.9°.

This refractive shift is about twice as large in magnitude as might naïvely be expected from the changing refractive index of the bulk of the neutral atmosphere, based on a difference of only 4° C

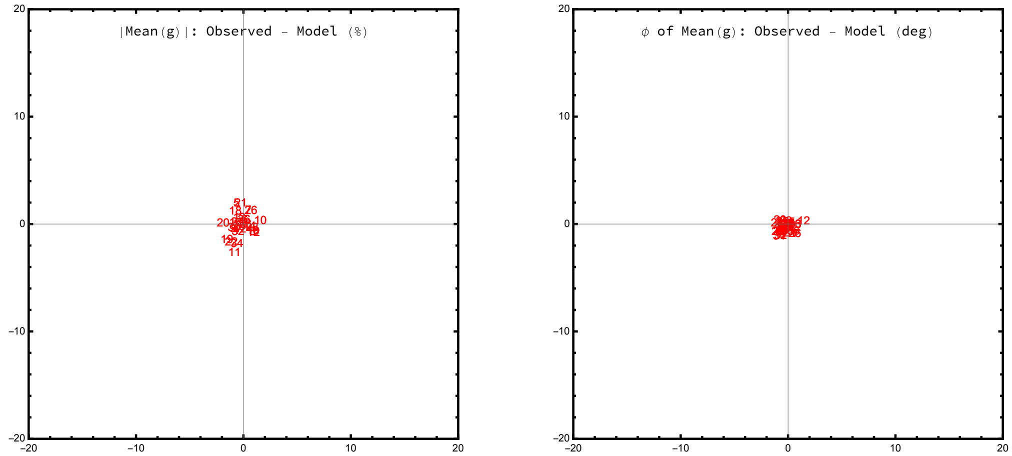

Figure 12: The amplitude and phase of the relative gains shown in figure 7, after subtraction of the physical models described in §5.1.1; it appears that those models can explain the majority of the gain changes between the two epochs.

in the overnight minimum temperature between the two days (as recorded at Murchison weather station). Furthermore the inferred azimuth of the refraction is 15° from the elevation vector for these observations, so a change in refractive index alone appears inadequate to explain the measured phase structure. This issue actually merits closer examination, because the real situation has more going on than a simple static atmosphere with a time-varying refractive index. Rather there is a diurnal variation that can be considered as an atmospheric wave that is travelling past the observatory, from East to West, and the diurnal variation has a large amplitude (more than 12° C, peak-to-trough), so with a modest LST offset in the harmonic structure of that atmospheric wave one can imagine that the observed magnitude and direction of the change in refraction might perhaps be accounted for. I expect that the level of detail needed for an accurate calculation of the bulk atmospheric refraction is captured by existing meteorological models and it might be useful to explore their use in this context.

The *differential* refraction, due to the neutral atmosphere, is expected to amount to $1.2\ \mathrm{arcsec}$ per degree of elevation in this test field. That effect was apparent when comparing the antenna gains towards the most point-like calibrator (shown in figure 6) with those of the variable source PKS J0243-0550, on a single epoch. However, the relative gains between epochs are insensitive to the average direction-dependent refraction: they respond only to the change in the differential refraction. If the measured change in refraction angle is due to the bulk properties of the neutral atmosphere, as suggested above, then the change in differential refraction is expected to be less than $0.03\ \mathrm{arcsec}$ per degree of elevation. I note that a physical model of the antenna phases introduced by the neutral atmosphere would also predict the direction dependence of those phases, and that information could be used to improve the imaging and photometry even if we don't actually solve for directional dependence in the gains.

In addition to the phase gradients described above, atmospheric refraction can influence the measured fluxes of compact sources — because the atmosphere is acting as a lens; it introduces magnification. The idea of "magnification" is in fact another way of thinking about the differential refraction across the field, so for sources that are resolved by the array the lensing effect can be completely removed by correcting for the differential refraction. However, for sources that are unresolved the image distortion is also unresolved and therefore remains uncorrected: one is left with a small change in the measured flux. The effect amounts to only about 0.4% even at the

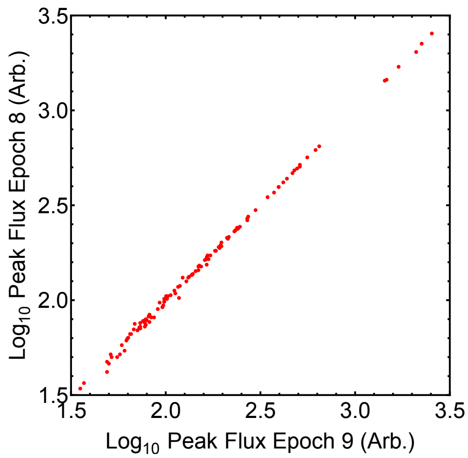

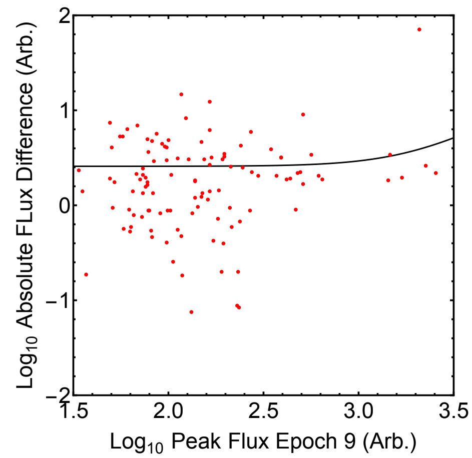

Figure 13: Fluxes of 114 NVSS sources (> 5 mJy) derived from the optimum direction-dependent calibration of §5.2, using the gain model of equation 7. Left panel: Epoch 8 versus Epoch 9. Right panel: absolute differences versus Epoch 9. Also plotted in the right panel is the expected median difference for a residual calibration error of 0.14%.

elevation limit of ASKAP ($15°$), and is an order of magnitude smaller than that near the zenith; it is an important correction only for high dynamic range imaging or high precision photometry.

The modulus of the relative gains, shown in figure 7, is well correlated between polarisations. As noted in connection with a similar finding in §4, I'm a bit surprised by this as I was expecting the amplifiers to be principally responsible for those changes and anticipating that the different signal chains would exhibit different fluctuations. It is also possible for the atmospheric opacity to affect the measured gain amplitude, although at 1 GHz the expected variations are below the 1% level whereas variations up to about 15% (and varying across the array) would be required to explain the left-hand panel of figure 7. As a sanity check I examined the auto-correlation signals for each antenna, as these should differentiate between atmospheric opacity and amplifier gains as the source of the measured gain variations: increasing opacity leads to a decrease in the cross-correlation signal, but adds atmospheric emission to the noise power that is recorded in the autocorrelations; whereas increasing the amplifier gain should lead to an increase in the cross-correlation signal and an increase in the autocorrelation signal. This check revealed a rough proportionality between the measured change in cross-correlations and the measured change in the geometric mean of the autocorrelations for the relevant antennas, so I conclude that it is indeed fluctuations in the amplifier gains that are responsible for the epoch-to-epoch differences in antenna gain moduli. The left-hand panel of figure 12 shows the residuals of the gain amplitudes after subtracting the gain estimate derived from the antenna autocorrelations; these residuals are small compared to the variations seen in figure 12. Presumably the amplifier gain can actually be tracked to high accuracy through monitoring of the on-dish noise calibration source; that information was not utilised in this study.

5.2 Direction-dependent gain model

To construct a set of direction-dependent gains I optimised a fit to the linear model in equation 7, and then evaluated the photometric performance of that solution in the same way as described for the direction-independent gain model. In this case the standard deviation of the flux differences for the 114 NVSS check sources is 2.78, which is only slightly larger than the standard deviation of 2.57 that applies to the difference image as a whole. The ratio of these two figures is 1.08, which is slightly larger than the corresponding figure (1.06) for the direction-independent model. The individual photometric points for this calibration are shown in figure 13. Adopting the same method as in §5.1 I estimate a value of $\sqrt{\langle\alpha^2\rangle} = 0.0014$ for the magnitude of the residual gain calibration error. That figure appears slightly better than the result obtained for the direction-independent model, but I don't set much store by that because it's a figure-of-merit that's really only sensitive to the brightest sources, most of which are actually being used to determine the calibration model. Consequently the four additional parameters of the direction-dependent model will make a difference to the estimate of $\sqrt{\langle\alpha^2\rangle}$ even if there is no direction-dependence in the relative gains.

It might be possible to get a better gauge of success of the direction-dependent model by extending the photometry out to larger distances from the phase centre, where there is greater sensitivity to pointing errors, for example. I have not tried that, for three reasons. First, the crop of sources being used here already extends beyond the point where these data would be considered as the primary source of photometric information. Secondly, although the fractional gain changes ought to be larger as one moves out towards the first null of the primary beam, the absolute response is low and it becomes hard to find sources that are bright enough to yield meaningful constraints at the level of $\alpha \lesssim 10^{-3}$. Thirdly, the simple, linear form of equation 7 is increasingly suspect as one moves further away from the pointing centre.

6 Summary and conclusions

Cross-calibration of an LST-matched dataset has been attempted and appears to be feasible. Direct cross-calibration of the uv-data themselves, without reference to the field structure, failed; reasons for the failure are not clear but are likely related to the problem of noise-bias in the figure of demerit.

The issue of noise bias was circumvented by working with summed, recentred visibilities — i.e. shifting the phase centre to match the position of a known, bright source and summing over time and frequency for each baseline separately. This procedure requires information on the positions of bright sources, but does not require any knowledge of their structure. Five bright sources in the field of the test data were identified from the NVSS catalogue to use in this type of calibration. On the first pass, one of these sources (PKS J0243-0550) proved to be significantly variable and was subsequently dropped from the list of calibrators.

Changes between the two observing epochs were modelled with a complex, relative gain for each antenna, both with and without direction-dependence. Optimum gains for each model were obtained by a least-squares minimisation procedure using the unregularised FISTA algorithm, with Wirtinger derivatives to form the gradient of the demerit with respect to the complex parameters being optimised. The optimisation is convex and FISTA converged rapidly in each case. The direction-independent model provided a good calibration, as determined by the measured flux differences of 114 bright ($> 5\ \mathrm{mJy}$) NVSS sources within 1.2° of the pointing centre. Residual calibration errors in the relative gain are at the level of $\lesssim 0.2\%$ for this model. This is sufficiently low that it is difficult to detect the inadequacies of the direction-independent model, because the brightest sources in the field have signal-to-noise $\sim 10^3$.

The best-fit direction-dependent model yielded comparable photometric performance to the direction-independent model, confirming that the direction-dependence of the antenna gain changes is small. This dataset provides little motivation for the extra generality of a direction-dependent model, but that is not to say that it would not be needed for other datasets in which the ionosphere or troposphere might be less benign. Indeed some fields in the VAST Epoch8/9 survey region exhibit particular calibration problems, as revealed by Emil's analysis; it would be valuable to investigate those fields, to determine what calibration freedoms must be permitted in order to achieve good photometry (and imaging) during these observing conditions. The calibration method used in the present study could perhaps be adapted for such an investigation as it's likely that the problematic conditions were manifest in only one of the two epochs, so the other should provide a good reference.

For the data examined in this study the bulk of the change in antenna phases between epochs is due to a change in the refraction angle of $0.67\ \mathrm{arcsec}$ — a small fraction of the $26\ \mathrm{arcsec}$ elevation shift that is expected at this zenith distance from the neutral atmosphere. The origin of this additional refraction could not be conclusively identified: it might be from the ionosphere or from the troposphere. Models already exist for the dynamic behaviour of both components and it would be useful to compare the predictions of those models with the phases that are measured from radio interferometric data. Although these atmospheric models are complex entities in their own right, they have the potential to provide good starting points for calibrating the antenna phases. Bearing in mind that the physical models predict time-, frequency- and direction-dependence, there is the potential for lowering the dimensionality of the space in which self-calibration solutions must be constructed — at least in the case of benign observing conditions.

It is encouraging that the telescope itself is very stable over an interval of 24 hours. It would be useful to investigate whether comparable photometric stability can typically be achieved over longer temporal baselines. If so then co-adding calibrated, LST-matched visibilities is an attractive option for spectral line surveys, providing greater sensitivity without increasing data storage requirements; this avenue is particularly attractive if the visibilities cannot be imaged as fast as they are acquired (as is currently the case). Even in the case where there is measurable evolution in the direction-dependent contributions to the antenna gains it's possible that co-adding LST-matched data might still be a good strategy. For example, in this study I modelled the direction dependence as a uniform gradient – a description that should capture differential refraction (relative phase gradient across the field) and small antenna pointing errors (relative amplitude gradient across field) – so one can imagine that averaging might well result in a lower-level of direction dependence for the combined dataset; and furthermore the average relative gradient ought to be predictable from the individual relative gradients that are measured at each epoch.

If LST-matching is a desirable way to proceed then ASKAP needs a facility for initiating correlator sampling at a pre-specified time, to a precision that is much better than $1\ \mathrm{ms}$.

Although feasible as a way of calibrating LST-matched data, the method described in this document ought to be inferior to calibrating both data-sets off a good sky model. The main advantage of the latter approach is that source confusion is not a problem because all sources are included in the one model. In turn that means that many more sources are contributing information to the calibration process, which should provide better constraints on any direction dependence in the gains. The FISTA algorithm, with Wirtinger derivatives for the complex parameters describing the antenna gains, worked well in the context of the summed, recentred calibration strategy employed in this study, but could equally well be used for calibrating against a sky model.