Quantifying and mitigating correlated noise between formed beams on the ASKAP Phased Array Feeds

Ian Heywood*, Paolo Serra, Aidan Hotan, David McConnell
CSIRO Astronomy & Space Science

Abstract

Pairs of beams formed from the weighted sum of receptors in the ASKAP Phased Array Feed share signals from common elements in the chequerboard. The noise characteristics of the two formed beams are thus not completely independent and share some correlated component. Intuitively the amount of correlated noise should increase as the angular separation of the directions of maximal sensitivity of the two beams decreases. The effect should also be frequency dependent as the optics dictate that fewer PAF elements contribute to each beam as the observing frequency increases. A correlated noise component has a negative effect on the survey speed of the instrument. With the Boolardy Engineering Test Array (BETA) now performing observations routinely the significance of the correlated noise issue can be assessed. Two 12 hour observations in Band 1 (711 - 1015 MHz) and Band 3 (1224 - 1528 MHz) are recast as pairs of coincident and anti-coincident measurements in order to simulate regular observing and an approximately equivalent set of fully independent, uncorrelated measurements as a reference point. The image noise for these observations is measured for beam pairs, expressed as a fraction of the fully uncorrelated measurements, as a function of frequency and beam separation. ***For any arrangement of beams that would be considered suitable for a typical large area survey, the increase in image noise due to correlated signals between pairs of beams is ∼10% in Band 1. The effects are negligible in Band 3. For BETA, the few percent increase in noise in Band 1 can be mitigated by deploying widely spaced beams and utilising an interleaved pointing strategy.*** The viability of this technique is demonstrated by imaging 50 square degrees of sky using only two BETA observations, each cycling between two pointing positions.

*ian.heywood@csiro.au

1 Introduction

The strength of the Australian Square Kilometre Array Pathfinder (ASKAP) as a survey instrument comes primarily from the ability of its Phased Array Feed (PAF) to form several beams on the sky. Single pixel feed telescopes require the whole array to be mechanically repointed in order to conduct a wide area survey whereas ASKAP can observe N pointings simultaneously where N is the number of formed beams. However since single pixel feed telescopes obtain their areal coverage sequentially the individual pointings exhibit completely independent noise properties. A pair of beams formed from the weighted sums of elements on the PAF have individual receiver chains that are common to both beams at a level which intuitively depends on the beam separation and the observing frequency. Thus the signals between two formed beams may not be completely independent of one another. Specifically they may exhibit correlated noise on some unknown level. When the data from these beams are eventually combined in the image domain, the noise in the image may integrate down slower than $1/\sqrt{t}$ where t is the observing time, having an adverse affect on the survey speed of the instrument.

With the Boolardy Engineering Test Array (BETA; [**1**]) now performing observations routinely an empirical investigation of this effect can be conducted and is presented in the sections that follow. The method adopted is described in Section 2 and the results are presented and discussed in Section 3. The article concludes in Section 6. A demonstration of conducting surveys with interleaved pointings [**2**] is provided in Section 4 and the article concludes in Section 6.

2 Method

The basic premise of this test is to split an observation into two approximately equal halves, one of which retains the correlated noise properties of the beams while the other mimics an equivalent observation with completely independent noise properties exhibiting no correlation. The two data sets can then be compared to determine at what level the shared signals between formed beams affect the survey speed of a PAF-equipped radio interferometer.

2.1 Observations and external calibration

Two BETA commissioning observations are used for this experiment, one observed in Band 1 (711 – 1015 MHz) and one in Band 3 (1224 – 1528 MHz). These two long track data sets were selected because of the range of beam spacings that were used, including very close packed pairs. The Band 1 observation was made with a circularly symmetric ring of eight beams around a central on-axis ('boresight') beam, and a linear arrangement of nine beams was deployed for the Band 3 observation. These beam configurations result in 5 and 8 unique beam pair spacings respectively.

Each observation includes scans of the standard calibrator source B1934-638 for flux scale and bandpass calibration purposes. The calibrator source is observed at the centre of each formed beam for approximately ten minutes. Following data flagging and spectral averaging to 304 × 1 MHz channels, these scans are used to derive bandpass and amplitude scale corrections for each beam by averaging in time and solving against a model of the spectrum of B1934-638 [**3**]. These corrections are then applied to the relevant beams in the target fields.

2.2 Preparing the images

The Measurement Set containing the visibility data from each of the nine beams is processed as follows, and as illustrated in Figure 2. Two copies are made (A and B) of the data from each beam, and both are sliced into time intervals of 500 seconds. Every other interval is discarded, however for the second copy of the data the retention pattern is inverted. The two data sets for each beam now do not share any common time slots, however they retain an equal number of scans (and therefore sensitivity), and approximately equal uv coverage (and therefore imaging quality).

To capture the frequency dependence of the correlated noise the data are split further into eight spectral windows (SPWs) of 38 MHz each. Large sky area, deconvolved images are then formed from each of these data sets: 13.65 × 13.65 degrees for the Band 1 data, and 6.83 × 6.83 degrees for the Band 3 data. The rationale for the large area is given in Section 2.3. The principal set of data products is thus 144 images per band: 9 beams × 8 spectral windows × 2 interleaved epochs.

2.3 Determining the F ratio

The key metric used in this analysis is the F ratio, defined as the fractional increase in the image noise due to the correlated signals between beams over the noise in an image formed from completely independent measurements.

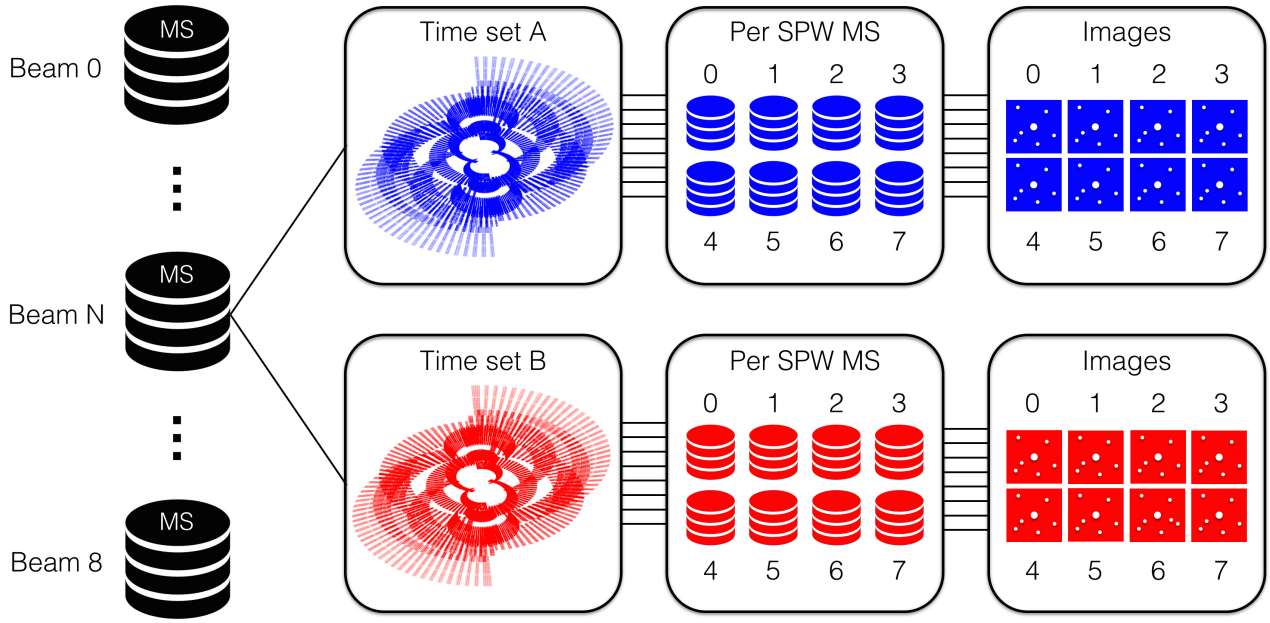

Figure 1: Illustration of the processing steps for each beam. The Measurement Set is split into two interleaved sets of scans. Actual uv coverages for time sets A and B are shown in the figure, demonstrating the very similar Fourier plane coverage despite the measurements in each set being completely anticoincident. The data are then split into eight separate spectral windows in order to capture the frequency dependence of any correlated noise effect. The individual spectral windows are then imaged, resulting in 16 unique maps per beam.

The value of F is computed as a function of beam separation and frequency by combining appropriate pairs of the 144 images generated as described in Section 2.2. Figure 2 illustrates this process.

For every possible pair of beams in the observation, and for each of the frequency bins, two maps are formed by stacking (i.e. taking the mean of) two images that have coincident times and two images that have anticoincident times. Since the latter has completely independent measurements between the two beams the noise should replicate that of an observation with no correlated signals, and the root mean square (rms) measurement of the image noise forms the denominator of the F ratio. The noise of the coincident combination of images will be higher if there is a correlated component between the two beams that were used to form the images, and the rms of the background in this combined image forms the numerator of the F ratio.

The greyscale zooms in Figure 2 show the reason for imaging over a large sky area. The central panel shows the rms noise across the full field. Measurements of the background noise are made over the area shown by the yellow rectangle. The noise measurements in the central field are artificially raised above the true noise level due to the presence of sources and confusion (likely dominated by sidelobe confusion rather than classical for images of this depth made with a six element array), as shown by the right hand zoom.

3 Results

There are 72 ways to combine the images from the 9 beams, however since the beams are regularly spaced there is a smaller number of discrete separations. The discrete separations have numerous redundant measurements which are used to provide constraints on the scatter in the F ratio. The Band 1 observations provide five unique spacings and the Band 3 observations provide eight. Note that comparing the M-N pair of beams is not the same as comparing the N-M pair.

Evaluating the F ratio for all possible beam pairs and each of the spectral windows for both sets of observations results in the plots shown in Figure 3. The spectrum of the F ratio is shown for Bands 1 and 3 (left and right panels respectively), with a single line for each of the unique beam separations that the observations provide. Further details are provided in the caption.

An F value of unity is the ideal case whereby the combination of the data from two beams is consistent with the noise being completely uncorrelated. The figure shows as one would expect that the value of F is strongly dependent on both frequency and beam separation. For a critically sampled mosaic in Band 1 the correlated noise increases the map noise by $\sim$10%. For Band 1 the effect is negligible for beam separations above approximately 1 degree.

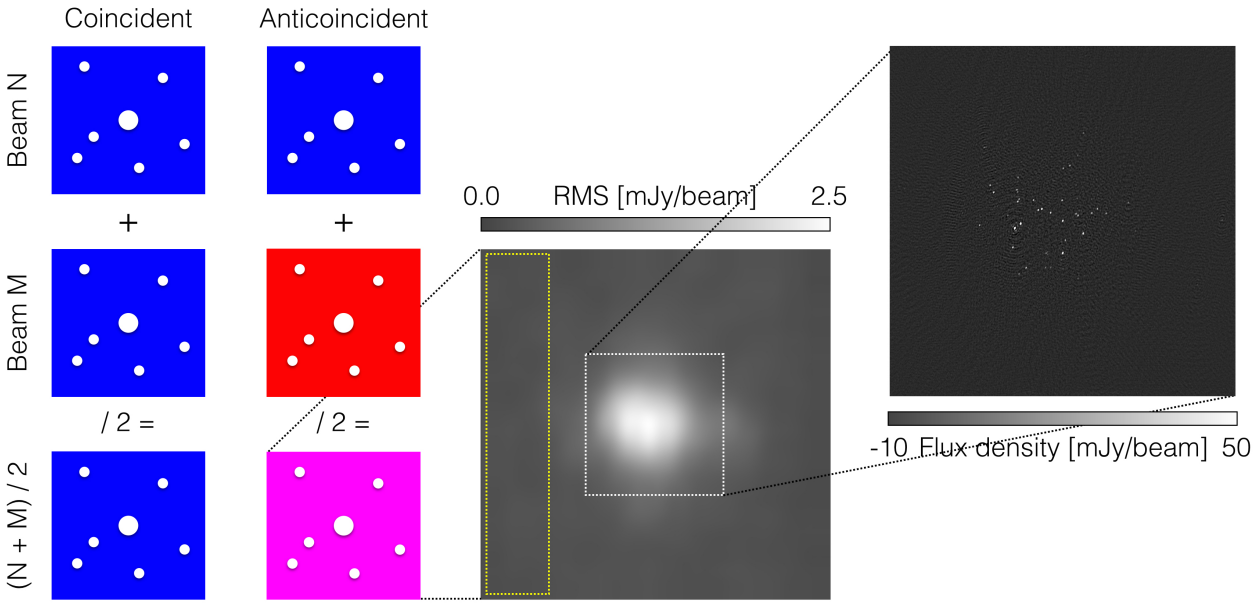

Figure 2: A measurement of F, the factor by which the correlated signals between beam pairs raises the noise above the uncorrelated case, is made by combining pairs of images generated from the beams in question across the frequency band. For example, to determine F on the beam separation exhibited by beams M and N, the images pairs formed from data that are coincident and anticoincident are averaged as illustrated. The rms of the background noise is measured far away from the main lobe of the beam in order to avoid contamination by astronomical sources and their sidelobes. This would skew such a measurement as shown by the rms map in the centre of the above figure, and the sky image inset on the right. These panels show SPW 0 of beam 0 for the Band 1 data. The ratio of the noise in the coincident map to that measured in the anticoincident map is the F ratio.

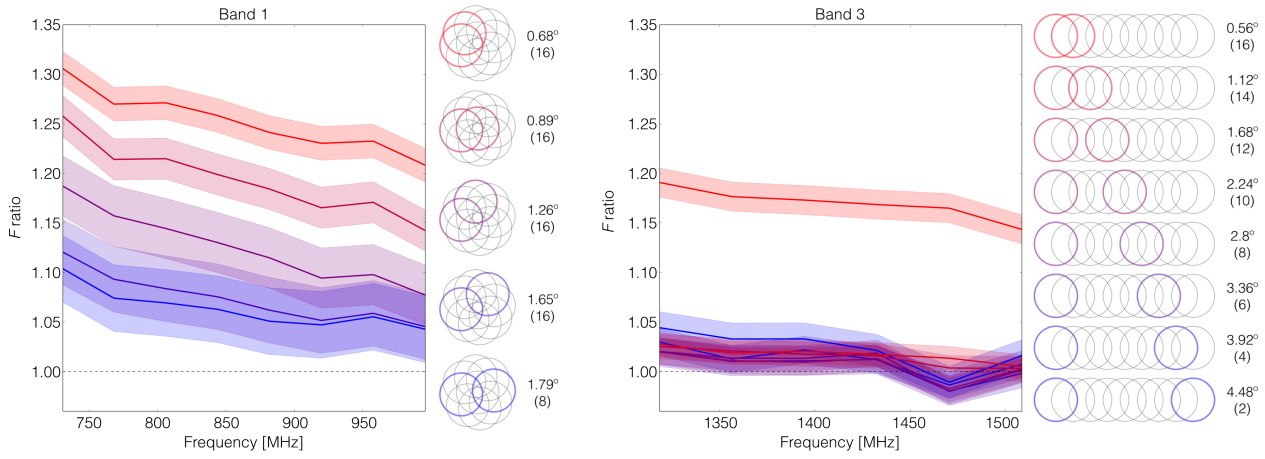

Figure 3: Plots of the F ratios as a function of frequency for the 5 and 8 unique beam separations from the Band 1 (left panel) and Band 3 (right panel) observations respectively. The diagrams to the right of each panel show the arrangement of beams with an example pairing that provides a given separation. The circles show the approximate half power point of the beam in the centre of the band, demonstrating the significant oversampling of these particular arrangements. The colour of the highlighted beams corresponds to that particular separation on the plot, moving from red to blue as the separation of the beam pair increases. The number of redundant measurements for each spacing is given in the parentheses. The standard deviation of the spectra of the F values is computed from the redundant measurements, and the shaded area associated with each line has a *total* height of one standard deviation. The horizontal line follows an F value of unity, i.e. the ideal ratio for the combination of completely uncorrelated noise fields.

4 Observing with interleaved pointings

This section briefly presents a demonstration of the use of appropriate beam spacings and interleaved telescope pointings to minimise F and survey a large area in the constellation Tucana. Nine beams are deployed in Band 1 in a 3×3 square formation with 711 to 1015 MHz of frequency coverage. The individual beam positions are shown on Figure 4, with the cyan and yellow points showing the two partially overlapping fields with 12 and 10+7 hours of integration time. The pointing direction of the telescope is coincident with the central beam and the pair of positions used per epoch are highlighted in red. The telescope switched from one pointing to another every 15 minutes. The separation of the beams is 1.46 degrees, which according to Figure 3 corresponds to $F \sim 1.1$.

Following editing and spectral averaging to 304 × 1 MHz channels the standard calibrator B1934-638 [**3**] is used to calibrate the flux scale and the bandpass of each beam independently. The data are then split into eight frequency bins and phase self calibrated using MeqTrees [**4,5**]. Three epochs, two pointings per epoch, nine beams and eight spectral windows thus results in 432 individual Measurement Sets. These are imaged and deconvolved using the CASA imager [**6**], and primary beam corrected (naively assuming a Gaussian) and the result of linearly mosaicking [**7**] these 432 images is shown in Figure 4.

The PyBDSM [**8**] source finder returns approximately 2000 sources above the 5σ level, where σ is the root mean square (rms) noise level in the image, typically about 0.6 mJy beam^{-1}. The noise as a function of area is shown in Figure 5. Eleven slices across the noise image spaces 100 pixels apart are shown in the right hand panel of Figure 5 showing the uniformity across the survey area. The mean noise profile is shown by the black line. Elevated noise levels are coincident with the brightest sources in the image where residual point spread function sidelobes remain due to the blind deconvolution. Further amplitude self calibration would lessen these.

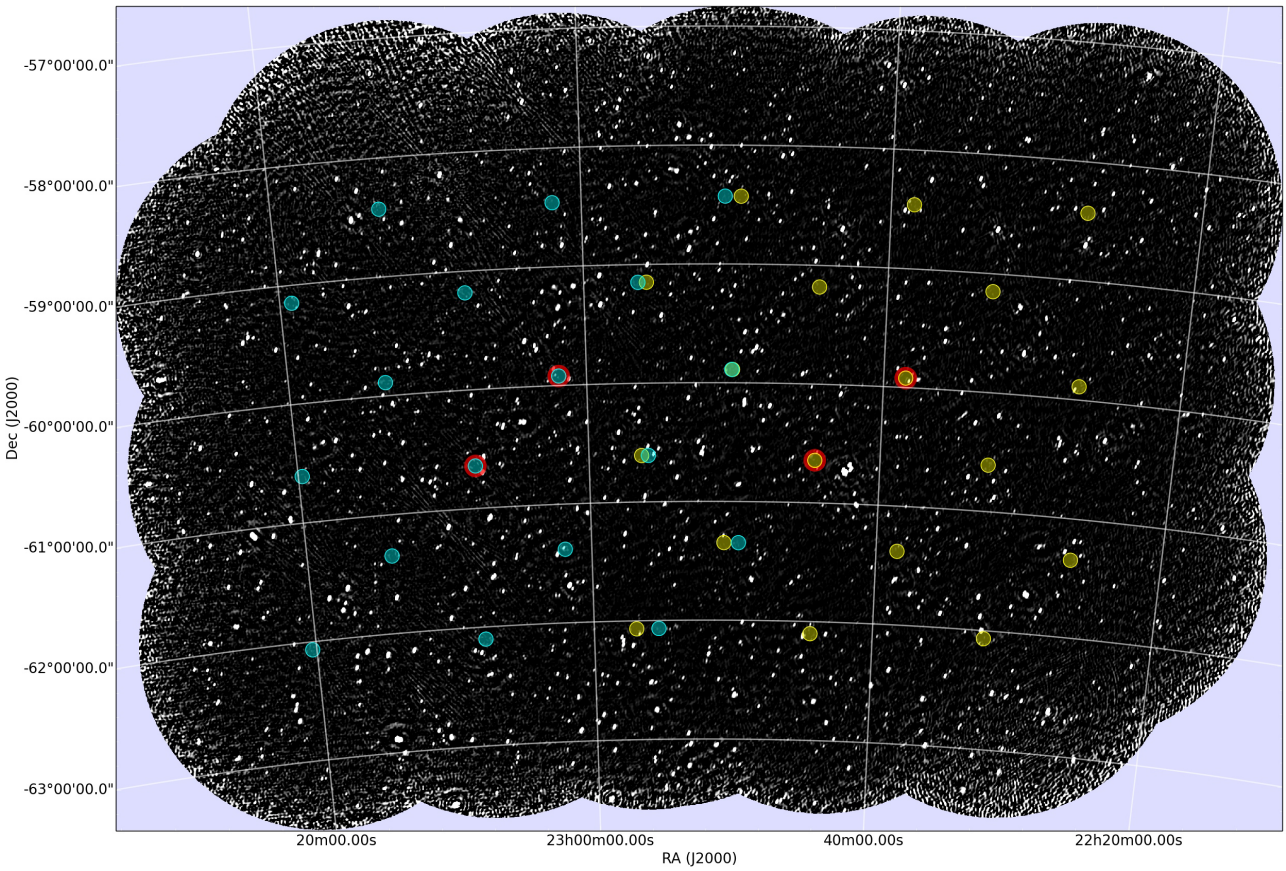

Figure 4: Example of using interleaved pointings and broad beam spacings over the course of an observation to cover a large survey area. Two fields are covered (the cyan and the yellow points) with an 3 × 3 arrangement of beams, and the array moving between two pointing directions on a fifteen minute cycle in order to achieve uniform sensitivity while minimising the correlated noise effect in Band 1. The image contains approximately 2000 sources above 5σ where σ is the rms noise in the image, typically 0.6 mJy beam^{-1}, see Figure 5.

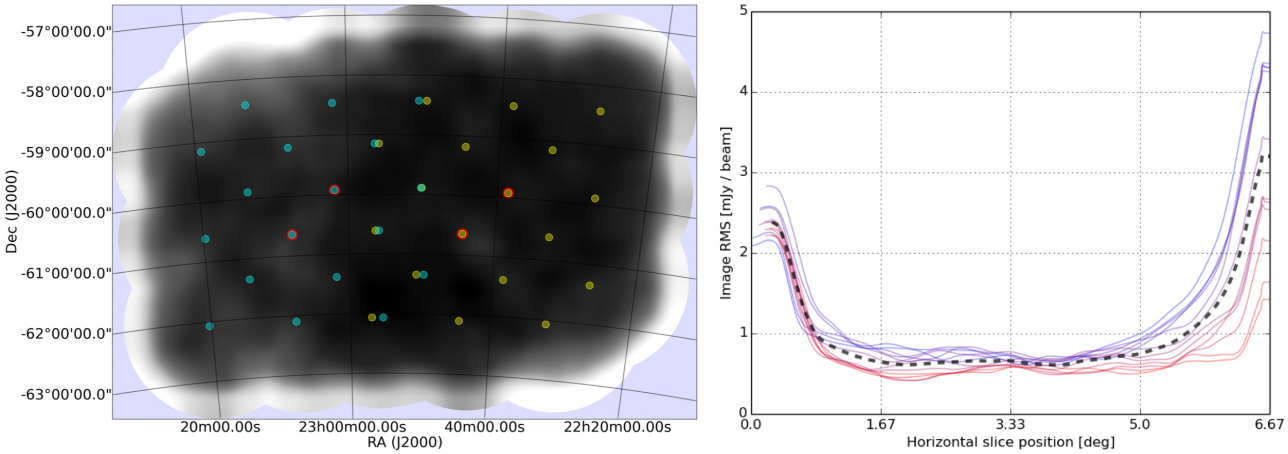

Figure 5: The left hand panel shows the rms noise level of the image shown in Figure 4. The right hand panel shows horizontal cuts made across the centre of the left hand panel every 100 pixels. The sensitivity of the image is uniform across the survey region demonstrating the viability of the interleaving strategy.

5 A correlation-weighted imaging mosaic scheme

Serra et al. **[9; Sec. 3.3]** present a scheme for linear mosaicking where the weights are derived from measurements of the noise covariance between the beams. A 5% improvement in the noise in their HI cube is achieved by using this scheme instead of the traditional variance-weighted mosaicking. Correlation factors (F) of 0.13–0.2 are measured for adjacent beam pairs (0.78 deg. separation), the spread of these numbers being ascribed to departures from circular symmetry depending on the beam pairs. F values for beam pairs beyond this separation were found to be negligible. We note the value of using such a weighting scheme (assuming pure noise regions of the map can be found in which to compute the noise covariance), and that the results are consistent with our findings.

6 Conclusions

The effects of the correlated noise component between beams formed using the ASKAP BETA PAF has been quantified as a function of beam separation and observing frequency. For a beam spacing typical of a critically sampled sky survey at the centre of the band, Band 1 (711 – 1015 MHz) exhibits a noise increase of about 10%. The effect is negligible for such a survey in Band 3, however the Band 3 plot suggests that the value of F may lie above unity, even for beams which are formed from receptors at opposite ends of the PAF. This warrants further investigation, although it could plausibly be a correlated noise-like signal that originates outside of the PAF.

Deploying a set of beams with greater separation and executing the survey with a series of interleaved sky pointings can mitigate the effect at low frequencies. A test of this observing mode has been completed. For full ASKAP using the maximum number of beams in Band 1 the beams will completely fill the field of view. An alternative approach may be to reconfigure the digital systems to reduce the number of beams on the sky but increase the number of elements that are used to form a beam, resulting in fewer but more sensitive or better shaped beams.

Certain imaging applications may require a close packed arrangement of beams, e.g. imaging very extended and morphologically complex structures, however such targets are not likely to benefit from maximisation of the survey speed in the way that for example a deep extragalactic survey does as they will not be noise limited.

A sensible way to extend the study presented here, and to see if the value of F approaches unity, would be to repeat the process on many more BETA data sets covering the full frequency range, broader range of beam spacings and improving the statistics. This would result in a useful tool for planning large surveys with ASKAP, delineating the (likely very few) observing regimes where correlated noise may be of concern.

References

[**1**] Hotan A., et al., 2014, PASA (preprint: arXiv 1409.1325H)
[**2**] Bunton, J.D. & Hay, S.G., 2010, International Conference on Electromagnetics in Advanced Applications (ICEAA), 728
[**3**] Reynolds, J., 1994, ATNF Technical Memos, 39.3/040
[**4**] Noordam, J. E., & Smirnov, O. M. 2010, A&A, 524, A61
[**5**] `http://www.meqtrees.net`
[**6**] `http://casa.nrao.edu`
[**7**] `http://montage.ipac.caltech.edu/index.html`
[**8**] `http://www.lofar.org/wiki/doku.php?id=public:user_software:pybdsm`
[**9**] Serra, P., Koribalski, B., Kilborn, V., et al., 2015, MNRAS, 452, 2680