ASKAP Science Observation Guide

David McConnell, Aidan Hotan, Aaron Chippendale, Emil Lenc, Elizabeth Mahony, Wasim Raja, Craig Anderson, Chenoa Tremblay and Matt Whiting for the ACES[a] and SDP[b] teams.

Version 1.0 – Pilot Survey Preparation Phase
March 2019

[a]Craig Anderson, Keith Bannister, Shivani Bhandari, Aaron Chippendale, Jimi Green, Aidan Hotan (Project Scientist), Karen Lee-Waddell, Emil Lenc, Elizabeth Mahony, Andrew O'Brien, Wasim Raja, Chenoa Tremblay, Matthew Whiting
[b]Eric Bastholm, Daniel Mitchell, Steve Ord, Wasim Raja, Max Voronkov, Matthew Whiting, Mark Wieringa

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

Contents

1 Introduction

Here we collect notes on ASKAP's capabilities and performance in advance of accepting specifications for pilot surveys from Survey Science Teams. The capabilities described represent the state of the telescope in the 1st quarter of 2019. Modes of operation (such as split frequency bands or tied array beams) not covered in this document will not be available for pilot surveys. Nevertheless, the ASKAP team welcomes suggestions for future enhancements and these will be prioritised in consultation with the science teams at a later date. This document will evolve over time into a comprehensive guide that can be used to plan science observations with ASKAP and as a reference for understanding survey data released through the online archive.

1.1 Pilot surveys

The idea of pilot surveys is to test whether an observing strategy based on current capabilities could feasibly be used to achieve the major science goals of a survey team. Given the multi-year duration of the survey projects it is likely that the telescope will improve in capability while surveys are underway. Survey teams should keep this in mind and build appropriate flexibility into their planning process.

This document describes the telescope, its operating model and performance and also outlines survey strategies appropriate for ASKAP's large field of view, including options for tiling the phased array feed "footprint" over large areas of sky. We also introduce a repository of software tools that can be used to experiment with survey parameter space.

As this document is being produced, the ASKAP operations team is observing test fields that have been requested by the science teams. These will be processed and uploaded to CASDA for investigation. One of the prerequisites for beginning pilot surveys is satisfactory production of an example data set for a test field in the mode of operation required for the survey. This process may take some iteration, but more general aspects of planning the large-scale survey strategy should be able to continue on the basis of advice herein.

We acknowledge that a highly detailed analysis of ASKAP's performance in the context of each science project is beyond the scope of this document. The most thorough investigations will no doubt be done on the science data itself, largely by the science teams. We hope that the results of these investigations will be shared with the telescope operations team and the rest of the community, either through publications or volunteer submissions to the ACES memo series.

If there is any background information required for planning a pilot survey that is not present in this document, please contact the authors.

2 ASKAP operation

2.1 Overview

ASKAP was designed to be a survey telescope operated by the Australia Telescope National Facility, providing calibrated images as its primary data product. This is very different to the user-driven operations model adopted by other ATNF telescopes which supply visibilities or some other form of raw data directly to an astronomer who operates the telescope themselves.

There is an expectation that a telescope designed for surveys should be more reliable and automated than a traditional instrument. While the ASKAP subsystems were designed with these factors in mind, the complexity of the phased array feeds and the large number of antennas in the array mean that reliability (and therefore automation) will naturally improve over a long timescale as updates are made. The pilot survey process reflects this with its small allocation of observing time. We wish to keep the emphasis on quality control rather than obtaining the first survey science results. However, we expect that pilot survey data should be publishable.

Treating calibration and imaging as part of observatory operations has proven to be one of the most challenging aspects of ASKAP. The original real-time processing model proved impossible to realise for the first generation of surveys due to the lack of a sky model for calibration and the need to thoroughly test and tune algorithms using real data.

For pilot surveys we have decided to operate in batch processing mode, which will exchange operational efficiency for the ability to re-process data until image quality criteria are met. This also allows the imaging algorithms used by ASKAPsoft to be tested while development continues and we wait for an upgrade to the capability of Pawsey's supercomputing facilities in 2020.

The default operational process will be as follows:

- A survey team will submit detailed descriptions of their survey strategy to the ASKAP operations team.
- The ASKAP operations team will create a set of scheduling blocks on the basis of the survey description, consulting with the survey team ACES representative as necessary.
- The ASKAP operations team will observe each block according to scheduling constraints. Progress will be visible to all teams via the web-based Observation Management Portal[1].
- The ASKAP operations team will process each observed scheduling block using ASKAPsoft in a mode specified by the science team, iterating several times and adjusting parameters until the data products are deemed to be of suitable quality. If problems or quality concerns arise during processing, the ASKAP operations team will consult with the science team ACES representative and revise the observing or processing strategy as necessary.
- Outputs from the ASKAPsoft pipeline will be uploaded to CASDA[2] as level 5 data.
- Science team ACES representatives will work with ASKAP operations to release each observation uploaded to CASDA as level 6 data products once they pass quality control.
- Science teams may create additional, value-added level 7 data products (e.g. additional catalogues, mosaics, etc.) and upload these to CASDA with an optional embargo period prior to publication.

The set of data products uploaded to CASDA will include calibrated continuum visibilities by

[1] https://apps.atnf.csiro.au/OMP/login.jsp
[2] https://data.csiro.au/collections/#sciDomainSearch

default. This will allow science teams to continue experimenting with imaging methods using their own computing resources, with ASKAPsoft or other tools. Most of the CPU time on Galaxy at Pawsey will be dedicated to operational processing and only a small queue of 50 nodes will be available for general use.

In special cases, it will be possible to arrange for transport of large visibility data sets to perform tasks that ASKAPsoft is not yet capable of (for example, joint de-convolution with single-dish data). This should be arranged with the ASKAP operations team in advance, since excessive file access can interfere with observations in progress.

Some of the ASKAP survey science projects do not use data from the correlator. For example, the CRAFT project obtains data from an alternative part of the signal processing system and operates its own processing pipeline. The stages listed above therefore do not apply, but we encourage all teams to consider how these alternative data products can be archived for public release in a timely manner.

3 ASKAP telescope specifications

ASKAP is an array of 36 antennas, each 12m in diameter. Detailed information about the configuration can be found on the ATNF web pages[3]

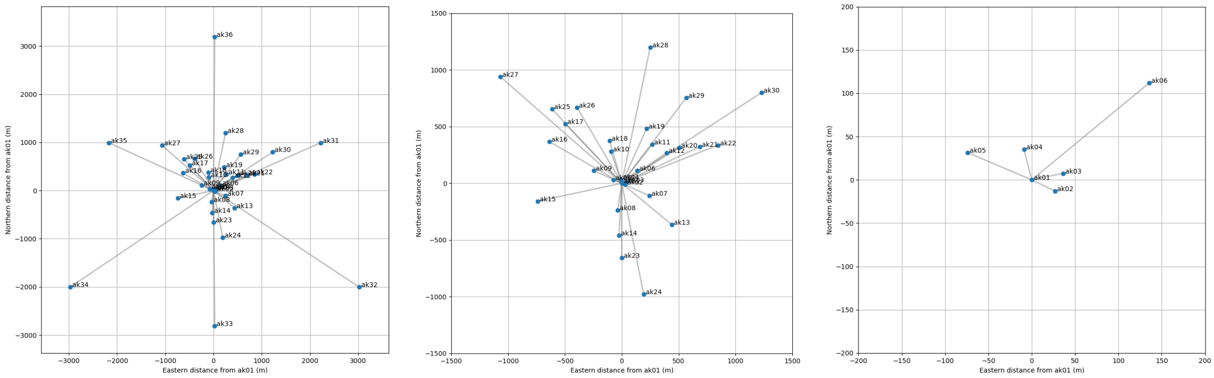

Figure 3.1: ASKAP array configuration at various scales

By default, data from all baselines is recorded. UV distance thresholds can be applied during processing to exclude short or long baselines as necessary. The shortest baseline is 22 m and the longest is 6.4 km. The array configuration is optimised for sensitivity to structures on 30” scales. With all antennas included, the synthesised beam size is approximately 10”.

ASKAP is sensitive to frequencies between 700 and 1800 MHz, with an instantaneous bandwidth of 288 MHz. Phased array feeds give each antenna a field of view of 30 square degrees (although after imaging and mosaicking, the equivalent area of the field of view is closer to 25 square degrees as described below). Within this field of view, we electronically form 36 simultaneous beams. At the low end of the frequency range, 36 beams are sufficient to cover the entire field of view. At higher frequencies this is not the case, leading to a decrease in survey speed unless other measures are taken (see below). We are investigating future capability enhancements to exchange bandwidth for additional beams, but this will not be possible for pilot surveys.

[3] https://www.atnf.csiro.au/projects/askap/config.html

We currently use maximum sensitivity beam-forming. This provides the best performance in terms of point source sensitivity, but allows each beam some degrees of freedom that are not present when using a mechanical feed. Through holography and source flux matching we have determined that the most significant difference between maximum sensitivity beams and a simple Gaussian illumination approximation is an overall beam width scale factor of roughly 10%. This is included in the primary beam correction done by ASKAPsoft.

At a lower level there will also be variation in beam shape across the field of view (with edge beams elongated radially due to coma distortion) and to an even lesser extent there may be some variation between beams from one antenna to the next. Aside from occasional anomalous points, holography measurements show that these variations are small. It is possible to enforce self-similarity of all beams, but at the cost of sensitivity. This trade-off will be investigated over time, but for pilot surveys we intend to use maximum sensitivity beam-forming and search for any systematic errors that may arise in the data as a result.

ASKAP records dual orthogonal polarisations and can form Stokes images as part of the standard processing pipeline. The X and Y polarisation beams are aligned in phase during formation and leakage terms (at least on-axis) are small. Detailed off-axis polarisation performance has not been studied and this is one of the objectives of the pilot survey process.

4 ASKAP performance

4.1 Sensitivity

The sensitivity of the telescope depends upon the system temperature and efficiency of each ASKAP antenna: T_{sys}/η. In practice we measure the system equivalent flux-density (SEFD) that is related to T_{sys} as:

$$\text{SEFD} = \frac{2kT_{sys}}{A_{\text{eff}}}$$

where A_{eff} is the effective antenna area. Figure 4.1 below gives measurements of SEFD across the ASKAP band for a central beam.

The sensitivity of ASKAP beams is expected to fall with increasing angular distance from the boresight. A large part of this decrease is due to increasingly large parts of the antenna's diffraction pattern on the focal plane moving off the sensitive portion of the PAF as the beams' angular displacement increases. Thus the sensitivity over ASKAP's field-of-view will decline from a central maximum to the edge of the footprint. In ACES Memo-15[4] we reported estimates of this envelope of sensitivity measured using the six-antenna Boolardy Engineering Test Array (BETA). We now give recent estimates based on both SEFD measurements of offset beams and image noise variation across 36-beam images. We find that the two methods disagree by about ten per cent in the equivalent area of the PAF field-of-view: 24.7 square degrees as determined from image noise against 27.5 square degrees from SEFD measurements. This discrepancy is not yet understood.

Section 6.1 below describes how the system noise relates to the final image sensitivity.

[4] http://www.atnf.csiro.au/projects/askap/memo015_a.pdf

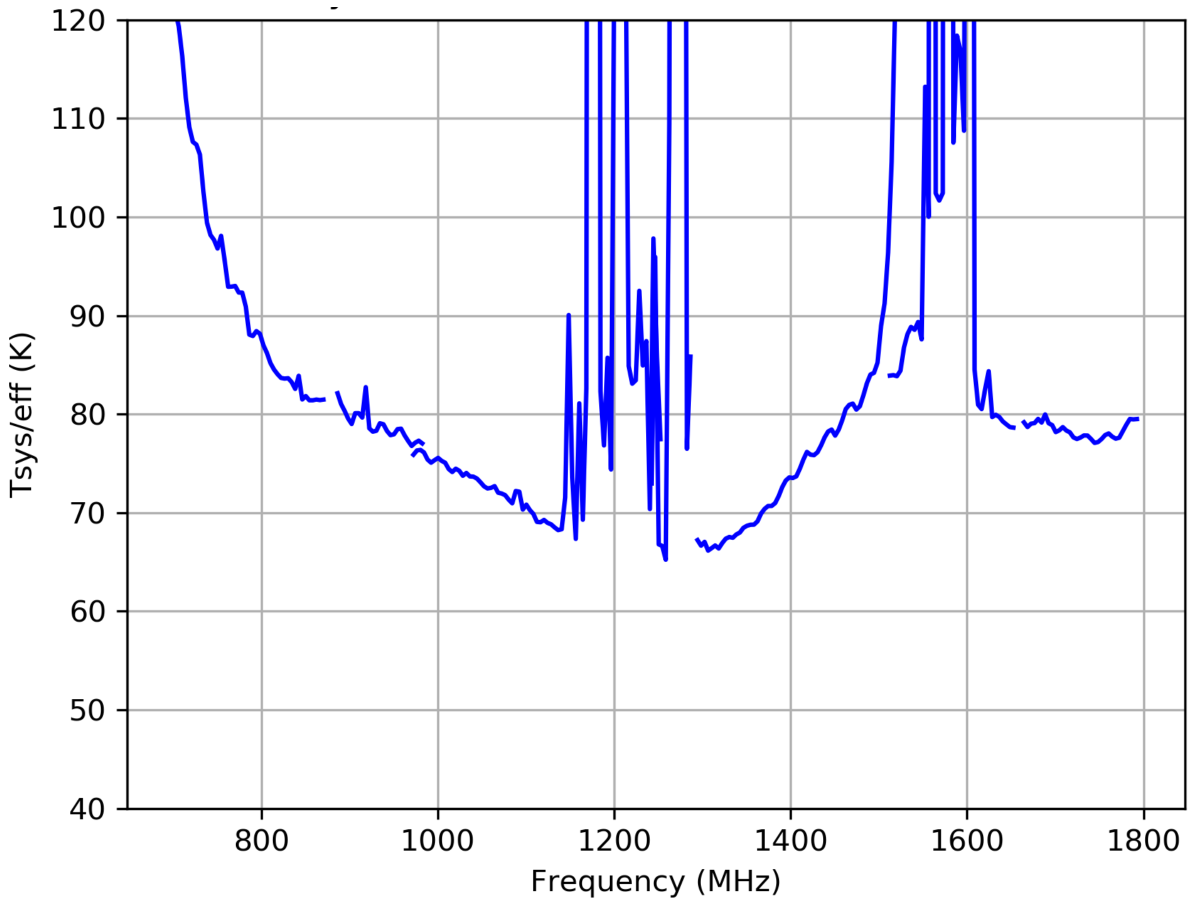

Figure 4.1: System temperature over efficiency across the ASKAP band. The median value over all antennas is plotted for a beam close to the antennas' boresight. Anomalously high values are typically due to RFI.

5 ASKAP data processing

In its wide-field spectral line mode, ASKAP produces 1.8 GB/s of visibility data. These are sent via a dedicated network link to the Pawsey centre in real-time, where they go through an ingest pipeline. Visibility data are subject to basic flagging before being merged with metadata and written to a number of measurement sets. The ingest pipeline has several tasks that can be switched on or off to define specific modes of operation.

5.1 Ingest modes

In full resolution wide-field mode, we write 6 measurement sets for each beam (to sustain the necessary disk write speed), creating a total of 216 measurement sets. The science data processing pipeline automatically merges these together during calibration and imaging. The ingest pipeline can also be run in continuum mode, in which it averages together 54 channels and writes only a coarse spectrum to disk. In this mode, we write only one measurement set per beam.

It is commonly requested that ASKAP offer additional flexibility at ingest, including the ability to average a configurable number of channels and discard parts of the band impacted heavily by RFI. ASKAP's data ingest system was designed to operate at high capacity at the expense of offering this kind of flexibility. Re-configuring the way data are written to disk requires altering

Task(s)	Approximate time [min]	Notes
Data preparation	150	Splitting, applying bandpass calibration, flagging, averaging and merging.
Continuum imaging, self-calibration	350	1 self-cal loop
Continuum mosaicking	240	Several jobs running simultaneously
Continuum source-finding	20	
Continuum-cube imaging	90	One cube per Stokes parameter
Preparation of spectral data	270	Applying gains calibrations from self-calibration, subtraction of continuum, and merging
Spectral Imaging	330	Image sizes = 1024x1024
Image-based continuum-subtraction	190	Serial job
Spectral mosaicking	360	Several separate jobs

Table 5.1: Approximate times for different aspects of the pipeline, making use of the new capability of splitting the datasets into small time chunks for the pre-processing. These numbers are rough averages taken from full-scale processing of a 28-antenna dataset.

the distribution of jobs across a cluster of machines and each configuration must be tested and re-tested every time any aspect of the system changes. To limit maintenance overheads we do not intend to offer more than the two simple modes described above. However, it is possible to split out smaller amounts of data after it has been written to disk and discard the original files.

5.2 Science data processing pipeline

The ASKAPsoft package was designed to efficiently image ASKAP data on a parallel processing platform. It is based on casacore and implements its own version of multi-scale, multi-frequency clean, with major and minor cycles that can be tuned using an extensive set of parameters. Taylor term imaging is used to deal with sources that exhibit spectral curvature across the band.

The optimum parameters for processing of 36-antenna data are still being determined, with a view to giving the best image quality subject to processing-time constraints. There have been a number of enhancements to the processing pipeline to improve the wall-clock execution times. These allow us to make better use of the supercomputing resources and more rapidly get data processed.

We show in Table 5.1 some indicative processing times for key blocks of the pipeline. These figures are derived from tests that were sized to fit all beams on the supercomputer at once, where each beam was limited to 13 nodes for the imaging (thereby utilising up to 468 of the 472 compute nodes). This is representative of how we intend to structure the processing for pilot surveys. The specifics of the parameterisation (such as number of self-calibration loops or major cycles in the imaging/deconvolution) will lead to variations on these times, but the scenario described in the table was found to give reasonably good images for well-behaved fields.

Flagging, calibration, imaging and mosaicking are all implemented as individual tools within ASKAPsoft. These tools are combined to form a pipeline by a set of scripts written in bash, which is well-matched to the slurm scheduling system used by Pawsey's supercomputers. The pipeline itself is therefore not very portable, but the individual applications are written in C++ and should run on a wide variety of platforms. Support for installation of ASKAPsoft outside of Pawsey is expected to be provided on some level, with the first tests being planned for OzSTAR at Swinburne.

5.3 CSIRO ASKAP Science Data Archive (CASDA)

CASDA is intended to be the astronomer's primary interface to the telescope. Each scheduling block that is observed will be run through the ASKAPsoft pipeline with a set of parameters defined by the observing mode and science team requirements. Outputs that pass quality control will be released on CASDA and made available to the public. The exact process of quality control is still under construction - the initial design of the archive assumed that science teams would be involved in releasing each individual scheduling block, but current thinking is that this will be done primarily by the ASKAP operations team.

Science teams are anticipated to do most of their analysis on products that combine several scheduling blocks, using their own computing resources to access data from CASDA and combine as necessary. These "value added" products can be uploaded to CASDA for archiving. The pilot surveys will test whether this model works well in practice or whether changes need to be made before commencing large-scale surveys.

6 ASKAP imaging performance

Most ASKAP time will be spent doing continuum or spectral line imaging. We have developed some additional modes for performance analysis and operational maintenance (such as antenna pointing measurement, holography for beam shape imaging, Solar observations for beamforming and use of the on-dish calibration system for beam weight updates). These will be used by the operations team to keep the telescope well-calibrated and continue research into its performance and characteristics. The rest of this section will consider imaging-mode performance in various ways.

6.1 Image noise compared to thermal predictions

ASKAP thermal noise can be estimated for a given observation based on SEFD measurements. For a 10 h, 288 MHz observation with 36 antennas we expect our thermal sensitivity limit (natural weighting) to be 12 μJy beam^{-1} assuming $T_{\mathrm{sys}}/\eta = 80$ K. Efforts during commissioning have been made to estimate sensitivity with the 28 antenna array by observing long-tracks of PKS B1934$-$638. While this is challenging from a dynamic range point-of-view, since the source has a flux density of $\sim$14 Jy, it provides a convenient in-beam calibration source to ensure that our flux scale is true. These test observations reveal that we can achieve a near-theoretical sensitivity of 13 μJy beam^{-1} with natural weighting (robust $= +1$), 16 μJy beam^{-1} with near-natural weighting (robust $= +0.5$), 20 μJy beam^{-1} with intermediate weighting (robust $= +0.5$), and 30 μJy beam^{-1} with near-uniform weighting (robust $= -0.5$) over a 10 h integration. Note that while sensitivity improves with larger values of robustness, the synthesised beam broadens and has increased side-lobe structure that complicates calibration and de-convolution and ultimately limits sensitivity and dynamic range.

Tests with snapshot images on shorter scales (10 seconds through to 10 minutes) and integrating multiple long tracks of PKS B1934-638 confirms that sensitivity scales appropriately as a function of $\sqrt{T_{\text{int}}}$. The exception being that for long integrations, when employing a natural-like weighting scheme, ASKAP approaches the confusion limit towards the beam centre in full-track integrations (> 10 h).

Comparison of simple fields (cosmology fields and UV Ceti observations) and fields containing relatively complex structure (SN1006) show that some degradation in sensitivity occurs when imaging a complicated field ($14 - 40$%). It is anticipated that this degradation can be reduced with improved de-convolution and in-field calibration.

6.2 Dynamic range in the presence of bright sources

Tests of dynamic range have also been performed using the field surrounding PKS B1934-638. Fundamentally, the dynamic range is limited by the quality of the calibration and the sky model. In the case of PKS B1934-638 we can routinely achieve a dynamic range of $> 500,000 : 1$ and have exceeded $1,400,000 : 1$ in a combination of three 14.5 h observations of the source. In the deepest observation, the only notable artefacts were around sources towards the edge of the field and were likely due to minor pointing errors. The dynamic range of more complex fields are currently limited by calibration and de-convolution errors - efforts to improve the dynamic range of such fields are on-going.

6.3 Performance near the celestial equator

To investigate how the synthesised beam, or point-spread-function (psf), changes across the equator, the GAMA 09hr field was observed with 36 beams centred at 09:00:00, +00:30:00 using the standard square_6x6 footprint. The observations were carried out on 2019-01-22 (SBID 7705) using a central frequency of 864.5 MHz and 1 MHz channels for an integration time of 5.5 hrs. Seventeen antennas were included in the array with baselines ranging from 25 m - 6 km. The data was processed with ASKAPsoft using a pixel size of 2 arcsec, image size of 6000 pixels, robust weighting of -0.5 and 2 Taylor terms (no multi-scale clean was used). Each beam was processed independently and the bmaj and bmin parameters extracted from the headers to compare how the psf changes with declination from -2 to $+2$ degrees. The final mosaiced image, overlaid with the psf measured for each beam is shown on the left in Figure 6.1. The rms at the centre of the field is $\sim 120\mu$Jy/bm.

A second observation with all 36 antennas was taken on 2019-03-05 (SBID 8116), again using a central frequency of 864.5 MHz, 1 MHz frequency resolution but observed for only 2 hrs. The data was processed with ASKAPsoft using the same parameters as the previous observation and reaches an rms at the centre of the field of $\sim 80\mu$Jy/bm. This observation, shown on the right in Figure 6.1, confirms that the psf stays roughly constant across the entire ASKAP field of view at the equator.

The presence of artefacts around bright sources in the field become more apparent at the equator due to the limited u,v coverage, but these would likely be reduced with more careful processing (e.g. using multi-scale clean, peeling of very bright sources). These artefacts are more apparent in the 36-antenna image due to the shorter integration time. The default parameters used by ASKAPsoft continue to be refined as more data using the full array becomes available.

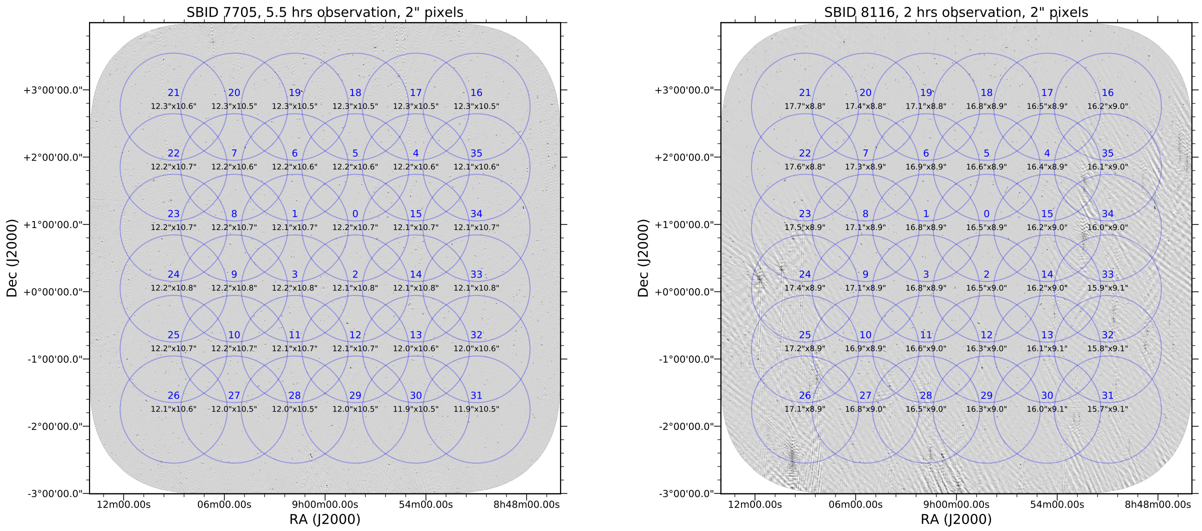

Figure 6.1: Psf for each of the 36 beams derived from a 17-antenna observation of the equatorial GAMA-09hr field (left) and a 36-antenna observation (right). The greyscale for both images is shown on a linear scale from -1 to 5 mJy. Each beam was processed independently, but all show similar synthesised beam sizes.

6.4 Effect of beam-forming frequency interval

Beamformer weights are calculated using channel widths of 1 MHz. At 850 MHz this corresponds to a velocity width of 350 km/s which, in some cases, can match the width of an astrophysical line. This could lead to the removal of spectral lines with widths larger than or equal to this interval. An option has been implemented to repeat beamformer weights over multiple 1 MHz channels to produce beamforming intervals that are much wider than expected astrophysical lines. This beamforming interval must be an odd number of 1 MHz channels. The weights for the central 1 MHz of each interval are used for all 1 MHz channels in each interval. This means that the weights are exact at the center of each interval and errors due to fixing the weights increase gracefully towards the edges of each interval. Test observations were successfully carried out using 5 MHz (SBID 8100) and 9 MHz beamforming intervals (SBID7210). The beamforming interval is specified by the `bf_resolution` parameter of the weights calculation script `osl_a_weightscalc.py`.

6.5 Spectral ripple - description and current state of understanding

Previous ASKAP observations of bright continuum sources resulted in higher spectral noise than expected. This excess noise has a very similar character to the pass-band response of the polyphase filterbank used to channelise the observing band into coarse 1 MHz channels – the Coarse Filter-Bank (CFB).

The transfer function design of the CFB is a trade-off between a flat bandpass and high stop-band attenuation with sharp transition bands to minimise leakage into neighbouring channels. This compromise leads to a small (0.2 dB) pass-band spectral ripple across each 1 MHz channel. The ripple is well-characterised and can, in principle, be calibrated out using an “anti-ripple” correction. The correction is applied to the beamformed data immediately after the fine filter-bank but before any fringe rotation is done.

However, observations towards bright continuum sources ($>$10 Jy) have shown that while this anti-ripple correction removes the majority of the PFB response a small residual ripple remains at the level of 0.2 – 0.5% of the continuum flux. Figure 6.3 shows this residual spectral ripple across 1 MHz (blue solid line) compared to the predicted response of the CFB (orange dashed line).

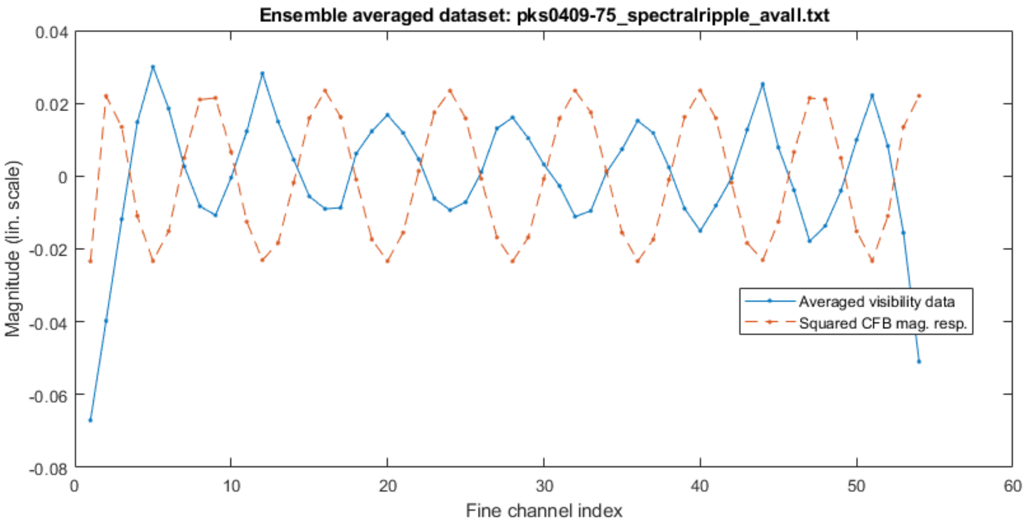

Figure 6.2: Spectral ripple seen across 1 MHz (54 fine channels) towards a 20 Jy continuum source (blue solid line) compared to the predicted response of the PFB (orange dashed line). The y-axis is squared magnitude, however, the mean value of each data-series has been removed so that the characteristics can be seen together on the same plot. Image credit: J. Tuthill, CASS.

The latest investigations by ASKAP's digital engineering team show that the correction coefficients were not convolved with the fine filterbank response in the same way that the visibility data are, leading to a residual ripple as shown above. An adjustment to the correction coefficients should be able to decrease the amplitude of this residual significantly, as shown below. This has yet to be tested on the telescope itself.

6.6 Astrometry and frequency accuracy

To investigate the accuracy of the astrometry and frequency/velocity of the spectral lines, a 10.5 hr observation was completed on 05-02-2019 along the Galactic plane. The observations centred on G335.2 +0.1, had a frequency resolution of 1.157 kHz ($\approx 0.25\, km\, s^{-1}$) and contained 20 known OH masers with emission at 1665.40 MHz. Of the 20 masers in the field, 18 were detected and all the detected OH masers were within 5 arc seconds of their previously published coordinates.

An additional check on the astrometry of the field was done by creating a shallow continuum image of the first 21 beams and cross matching the source positions with the Sydney University Molonglo Sky Survey (SUMSS) source catalogue. All of the detected continuum sources were offset by only a few arc seconds, smaller than the size of the synthesised beam.

The current ASKAPsoft pipeline offers a barycentric standard-of-rest correction to the sky frequency but most published results are corrected to the radio local standard-of-rest (LSR) frame. Therefore, the data were processed in ASKAPsoft with no correction to the sky frequency and it was done upon creation of the spectra. To test the accuracy of the frequency and therefore the velocity of the OH masers detected, the observations from ASKAP were compared to obser-

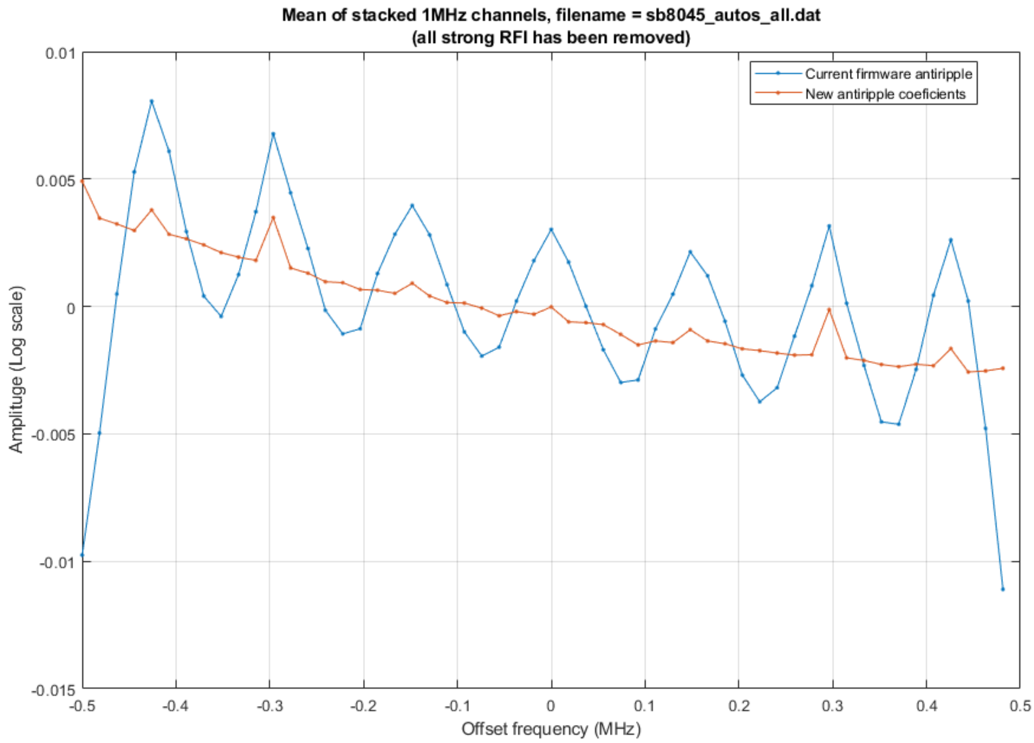

Figure 6.3: Comparison of spectral ripple remaining after using the existing (blue) and proposed (red) correction coefficients. Image credit: J. Tuthill, CASS.

vations taken with the Australia Telescope Compact Array on 29-02-2019. This is important as the maser profiles can change depending on the activity of the object. As shown in Figure 6.4, the profiles and peak velocities match within the differences of each telescope's spectral and spatial resolution. This comparison work is ongoing and will be updated in a future version of this document.

7 Radio Frequency Interference

Although the MRO site is largely free from terrestrial RFI, it is still susceptible to airborne and satellite-based sources of interference. Global navigation system satellites are a major source of interference in bands around 1200 and 1600 MHz, along with aircraft transponders at 1090 MHz. The worst-impacted part of ASKAP's frequency range is 1150 – 1300 MHz, with fractional occupancy averaged over time and baselines exceeding 10% nearly everywhere and rising to 80% in some channels. While it may be possible to mitigate these satellite signals with adaptive beam-forming in future, this technique is not ready for operational deployment and will likely have the side-effect of making the primary beam vary with time (making calibration more difficult). In the meantime, traditional flagging will be necessary.

As discussed previously, split bands are not available yet. However, if astronomical sources of interest occupy less than 288 MHz of bandwidth, zoom modes (even combined with online frequency averaging) could be used as a form of RFI mitigation by narrowing the total frequency range recorded.

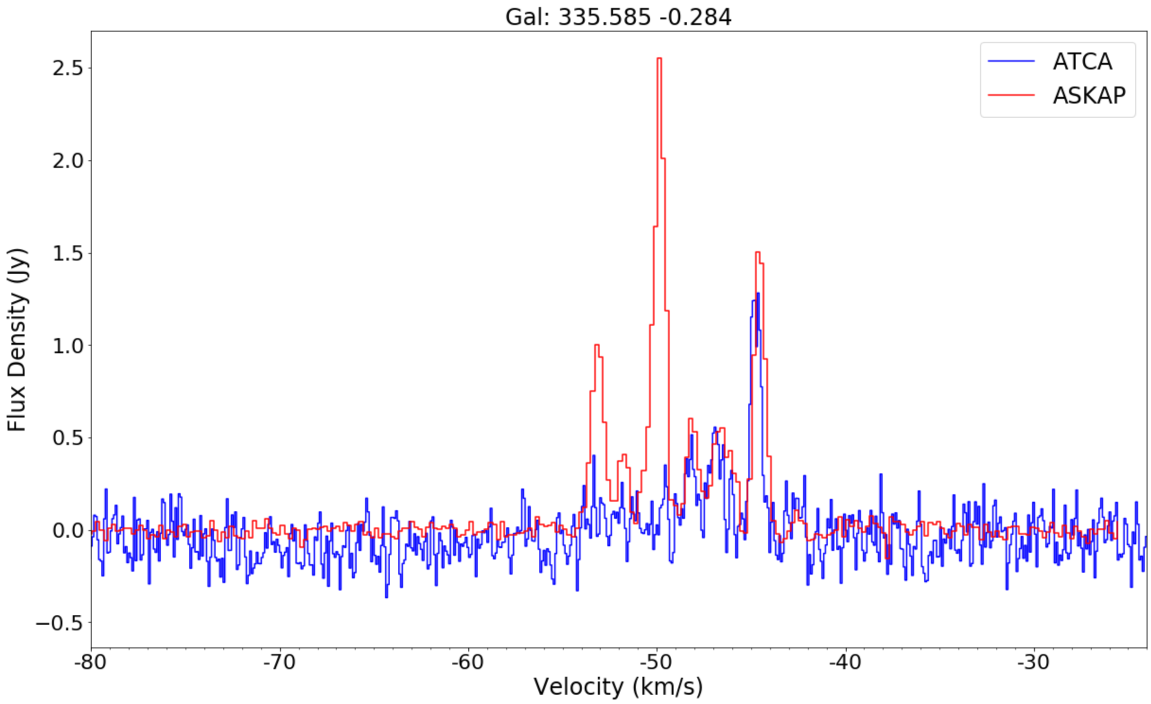

Figure 6.4: A comparison of the OH maser profile of OH335.585–0.284 between ATCA and ASKAP (both obtained February 2019). The y-axis is the flux density but the values are not scientifically accurate. The image shows that the velocity and peak profiles match within the differences of each telescopes resolutions (the ASKAP spectrum includes the neighbouring source OH335.585–0.289, peaking at $\sim$ –50km/s). Image credit: C. Tremblay, CASS.

Atmospheric ducting can also create a path by which distant transmitters interfere with telescopes at the MRO. This typically manifests as the occasional presence of mobile phone transmission tower signals around 850 MHz in the otherwise extremely clean 700 – 1000 MHz frequency range. This can be predicted using meteorological data and we will attempt to schedule around ducting events where possible.

Early attempts have been made to measure RFI occupancy by collecting flagging statistics over all of the observing bands during observations of PKS B1934-638. Figure 7.1 shows the flagging occupancy over the entire ASKAP frequency range measured for each 18.5 kHz channel. At the moment no attempt has been made to separate data flagged as a result of instrumental errors from those flagged as a result of RFI so several of the dominant lines at 100% occupancy are in fact not caused by RFI. Effort is on-going to improve the method used to measure the RFI occupancy as a function of baseline and as a function of time.

ASKAPsoft includes a basic flagging tool which is run as part of the pipeline. It is possible to use AOflagger instead (also as part of the standard pipeline) and several science teams are working on ASKAP-specific flagging strategies for this tool.

8 Sampling the sky

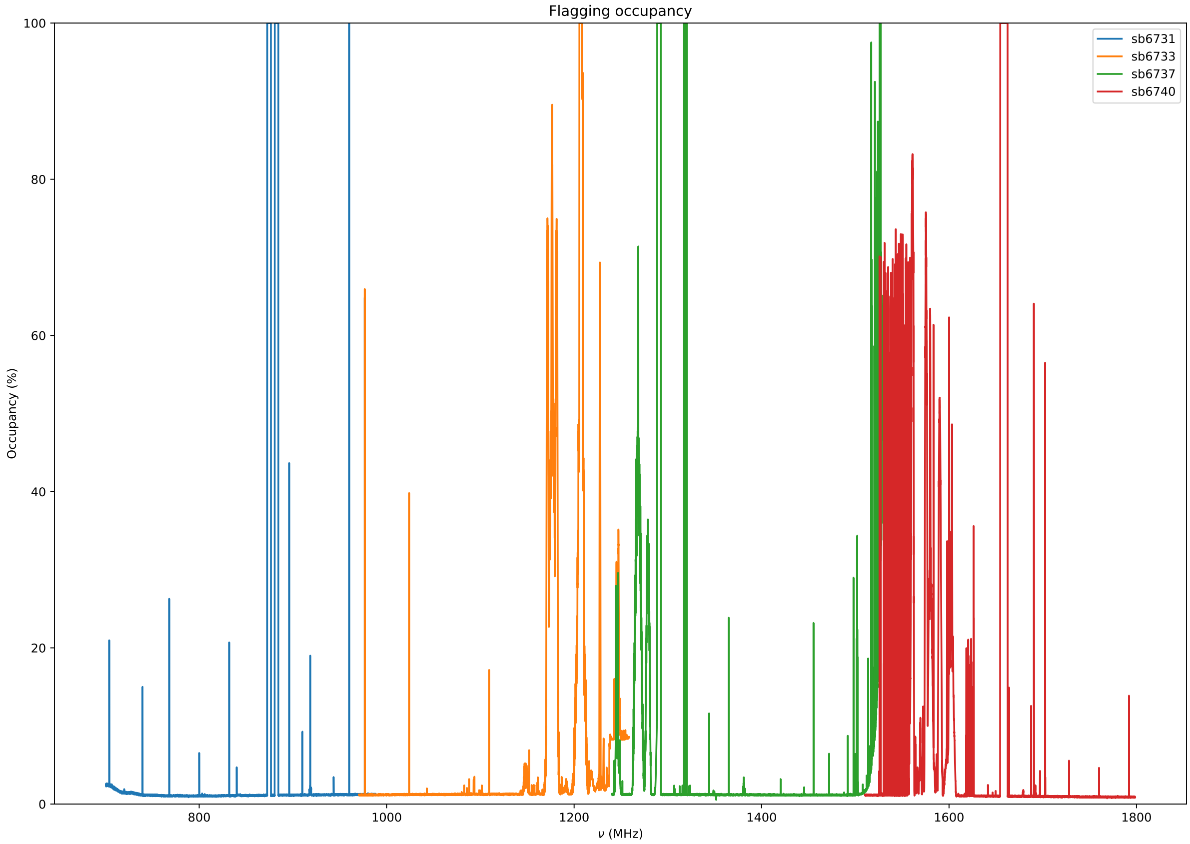

Figure 7.1: Flagging percentage (as a proxy for RFI contamination) averaged over time and baselines for observations taken at different centre frequencies. Some of the peaks that persist at 100% for several channels (most notably around 900 MHz) are in fact due to correlator mis-alignment that occurred during the scheduling block used to obtain the flagging statistics. The band between 700 and 1080 MHz is almost entirely free of RFI most of the time.

8.1 Beams and footprints

ASKAP beams are arranged into patterns within the field-of-view referred to as “footprints”. Two 36-beam footprints are in common use, and most imaging experience has been gained using one or other of them. The diagrams in Figure 8.1 give the numbered beam positions within the footprints.

8.2 Tiling the sky

ASKAP surveys will comprise a series of observations, each sensitive to the sky over the PAF field-of-view as determined by the footprint. We refer to the placement of these observed fields on the sky as “tiling” the surveyed area. The details of the PAF field-of-view and the relative merits of various footprints are described in a later section. Here we present options for tiling the celestial sphere.

For AKSAP-BETA operation and for the early period of the Early Science observations a simple quasi-rectangular tiling scheme was used and implemented by the TILESKY script. This did not allow for the efficient tiling of the whole celestial sphere needed for more extensive surveys. Following suggestions made in the “WALLABY Memo 22 v1.0: Sky Tiling” by Aaron Robotham, we have developed a new scheme and software for its implementation.

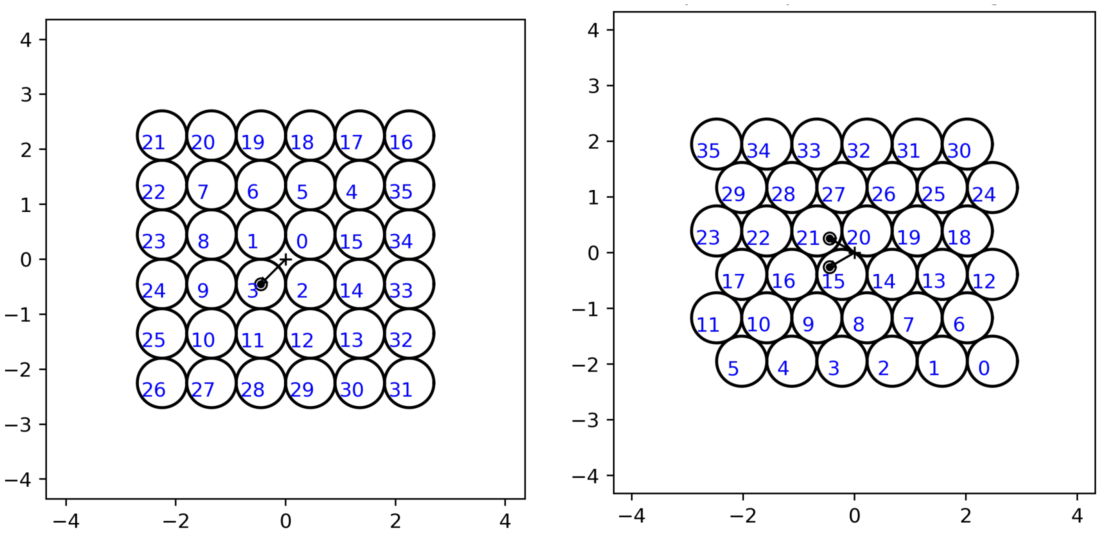

Figure 8.1: The two commonly-used footprints: `square_6x6` (left) and `closepack36` (right). In this case both have beam spacings (pitch) of 0.9° and a position angle of zero. The scales are in degrees, and celestial north (west) is to the top (right) of both diagrams.

An approximate description of the new scheme follows. The celestial sphere is divided into three zones: a central zone and two polar zones. Positions on the sphere are denoted by their longitude and latitude (Θ, Λ). Within the central zone, the sky is tiled with a series of bands (along small circles) parallel to the equator and separated by one tile-width. Tiles are arranged evenly around each band to touch or overlap at the lower-latitude edge. If the longitudinal tile width is w, and the low-latitude edge of the band is at latitude Λ, there will be N_Λ tiles such that N_Λ is the smallest integer $N_\Lambda \geq \frac{2\pi \cos \Lambda}{w}$ and so the overlap between adjacent tiles will be $w - \frac{2\pi \cos \Lambda}{N_\Lambda}$. The bands of tiles are arranged without overlap in the range $[\Lambda_-, \Lambda_+]$. The values of $\Lambda_\pm$ are adjusted to ensure an integral number of bands: if the tile north-south dimension is height h, then $\frac{\Lambda_+ - \Lambda_-}{h}$ is integral.

Tiles are arranged in the polar zones in a similar way along segments of small circles centred at $(\Theta, \Lambda) = (0, 0)$. Figure 8.2 shows the tiling scheme.

The scheme implemented differs slightly from that described above.

1. In the longitude direction, a tile is placed so that its eastern edge intersects the western edge of its neighbour half a beam-pitch polewards from the low-latitude edge. This allows a slightly greater tile spacing without letting beam separation across tile boundaries exceed the footprint beam pitch.

2. In the latitude direction, ranks of tiles are placed slightly closer together in the high latitudes to avoid regions of lower sensitivity along tile north and south boundaries. Figure 8.3 illustrates the geometry involved.

There is no tiling scheme that avoids some inefficiency in sky coverage: most tiles overlap with their neighbours. The overlap increases with increasing distance from the equator, and is large along the boundary of the polar zones. The fractional inefficiency is a function of tile size and the values of $\Lambda_\pm$. Figure 8.4 shows the inefficiency fraction as a function of Λ_+ (assuming $\Lambda_- = \Lambda_+$) for several typical values of beam-pitch and therefore the tile size.

These graphs show that there is no strongly preferred value for $\Lambda_\pm$; the main significant trend is

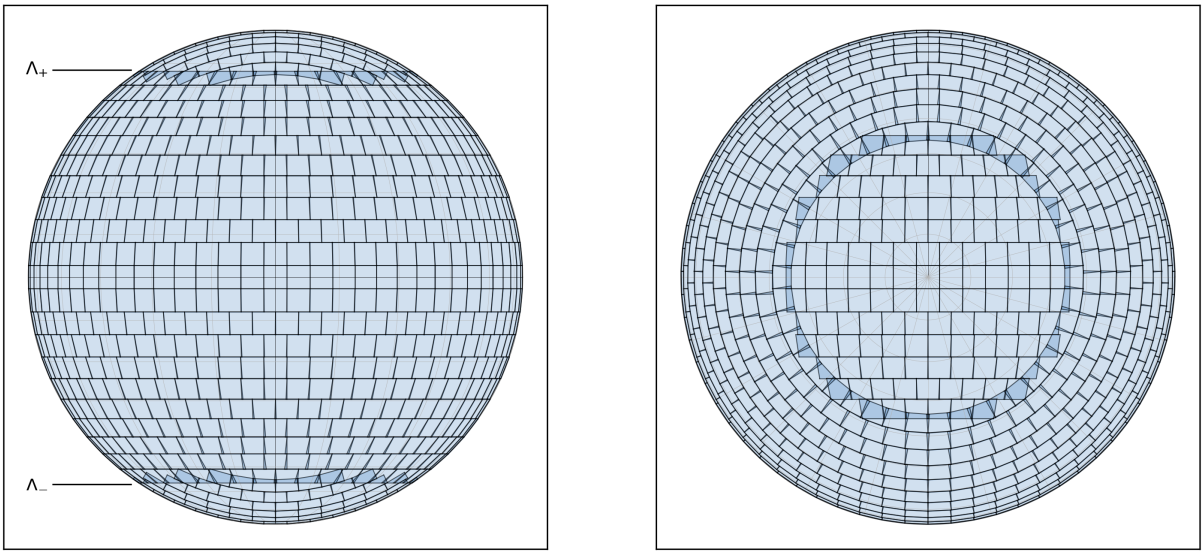

Figure 8.2: Two views of the tiles sphere: equatorial (left) and polar (right). The boundaries of the two polar zones at Λ_- and Λ_+ are indicated in the left-hand panel.

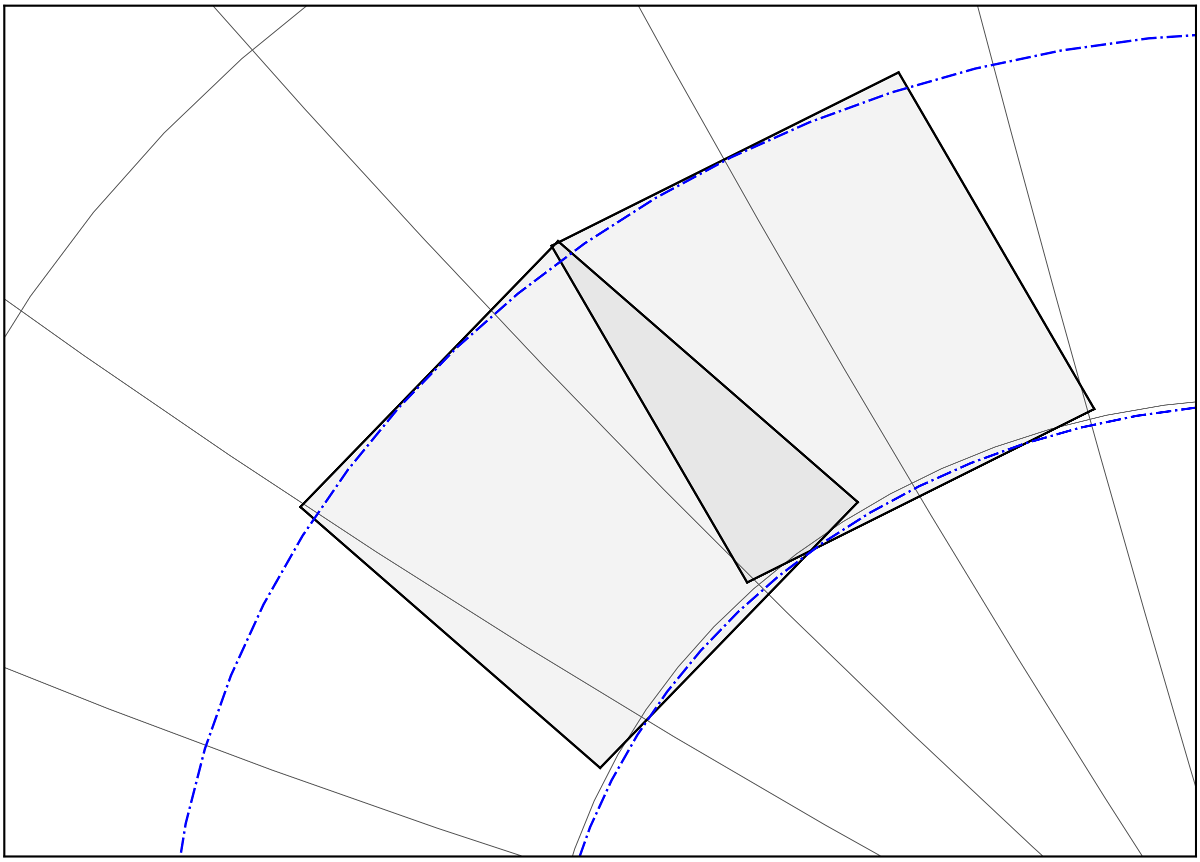

Figure 8.3: Two tiles at a high latitude. The north and south boundaries of the tiles are shown as dashed lines, and are separated by less than the tiles' latitudinal extent. This size decrement increases with latitude.

the expected, but minor increase of inefficiency with tile size. Also, tilings with very large polar zones are less efficient. The simplest arrangement, one with the smallest possible polar zones, is comparable to others and so perhaps is best.

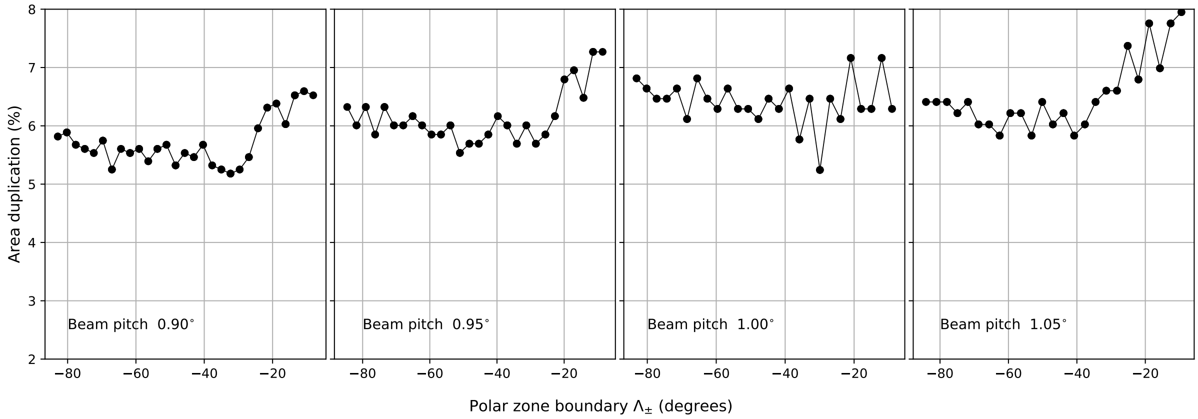

Figure 8.4: Graphs of the survey area observed by more than one tile, expressed as a fraction of the total survey area, and computed for a range of Λ_+ values. Here a SQUARE_6X6 footprint is assumed so the tile dimensions are $w = h = 6 \times p$ where p is the beam spacing.

The tiles can be laid down in Equatorial, Galactic or Magellanic (ref) coordinate systems. The tiling script converts tile positions to equatorial (J2000) coordinates and computes the tile position angle relative to the NCP for entry into the ASKAP control system. The origin of the tile pattern, normally at $(\Theta_o, \Lambda_o) = (0, 0)$, can be shifted, allowing the placement of the quasi-rectangular part of the pattern at some chosen point while maintaining the tile grid parallel to the underlying coordinates. Figure 8.5 gives all-sky views for the three supported coordinate systems and a set of equatorial tiles with a shifted origin.

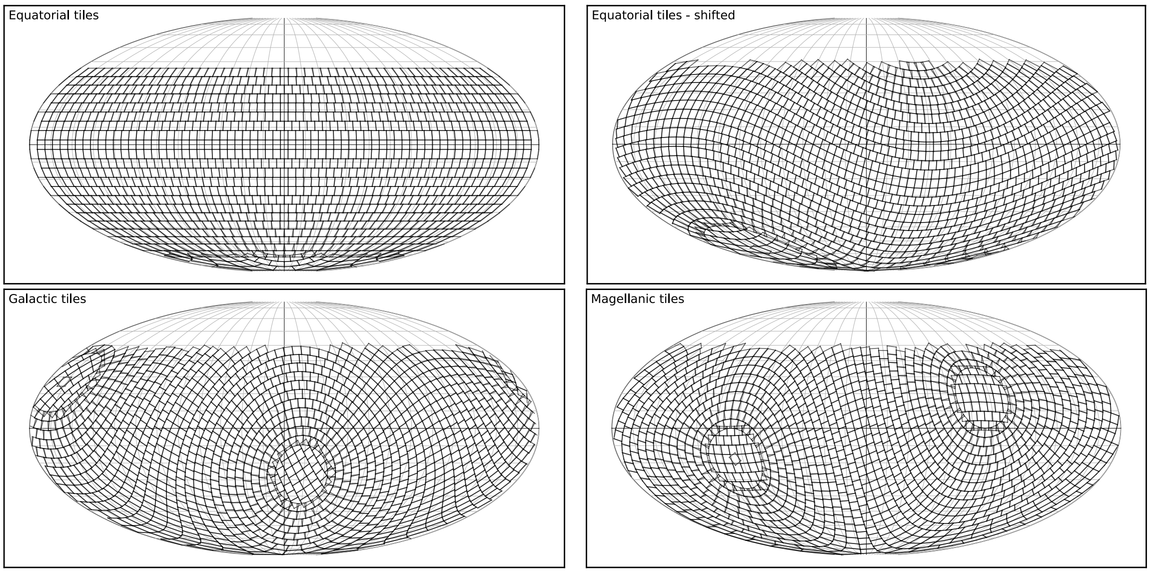

Figure 8.5: All-sky views of four different tilings, all constructed for $6^\circ \times 6^\circ$ footprints with $\Lambda_\pm = 72^\circ$. Top left: tiles defined in equatorial (J2000) coordinates; top right: the same set of tiles with origin shifted to $(\Theta_o, \Lambda_o) = (30^\circ, -23^\circ)$; lower left, right: tiles defined in Galactic, Magellanic coordinates. The tiles inaccessible to ASKAP are not shown. The heavier grid lines indicate the celestial equator and zero of the Right Ascension scale.

8.3 The optimum tile size

Following the method outlined in (ACES memo15), the estimate of the solid area of ASKAP's field-of-view (see section 4.1), we have estimated the optimum beam spacing as a function of observing frequency. Figure 8.6 also shows the optimum survey speed expected and the degree of sensitivity variation (ripple) over the 36-beam mosaics.

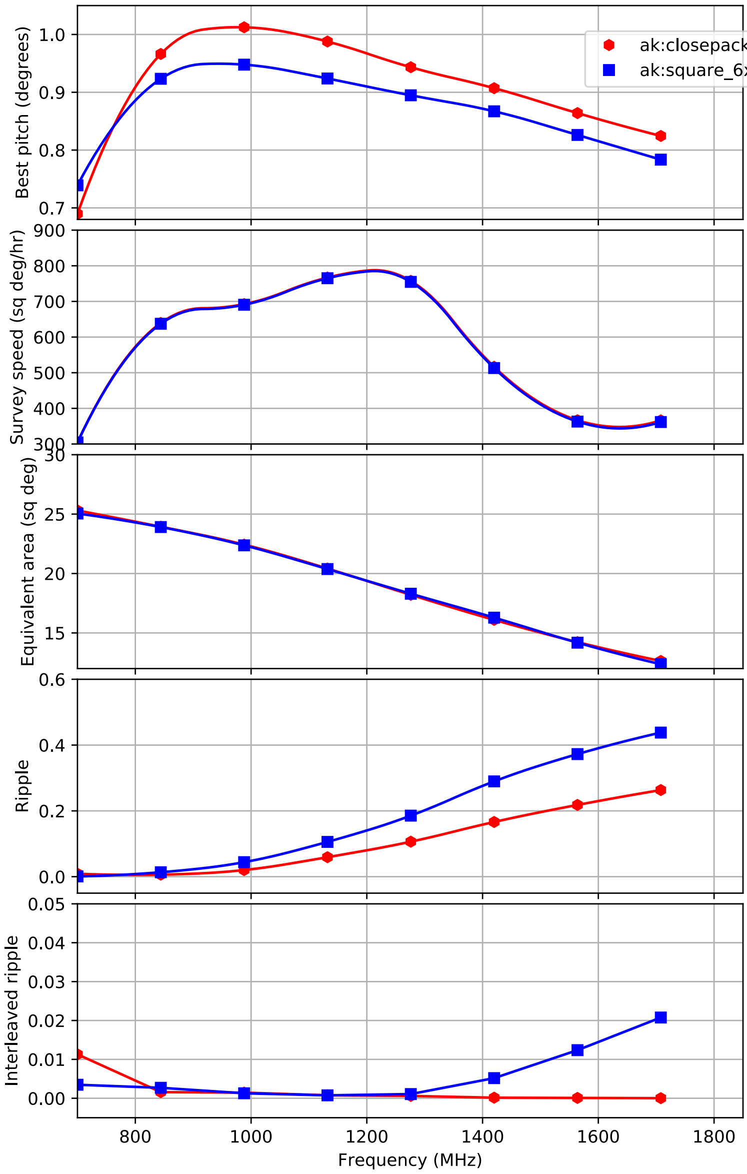

Figure 8.6: From top to bottom the plots show as a function of frequency: the beam pitch that optimises survey speed; the optimised survey speed; the equivalent area; the ripple of the optimised footprint before and after interleaving. The model was evaluated at the indicated frequencies; a smooth curve joins the calculated values. The ripple is computed as (max-min)/mean over a portion of the centre of the mosaiced field. The equivalent area of the single observation (not interleaved) is shown and is calculated as $\int w(l,m)dldm$ where $w = \frac{1}{\mathrm{SEFD}^2}$, varying across the field of view. The survey speed values were computed using values of T_{sys}/η evaluated from a polynomial fit to the data in Figure 4.1, for B = 288MHz, $n_p = 2$, $N_a = 36$ and $\sigma_s = 200\mu$Jy.

8.4 Example results from test RACS observations

The observatory-led project RACS (Rapid ASKAP Continuum Survey) is being planned to make a shallow survey of the whole accessible sky, motivated by the need to initialise ASKAP's global sky model, and the complementary science that would result. The tiling scheme described above has been used to define test RACS observations. Figure 8.7 illustrates the results.

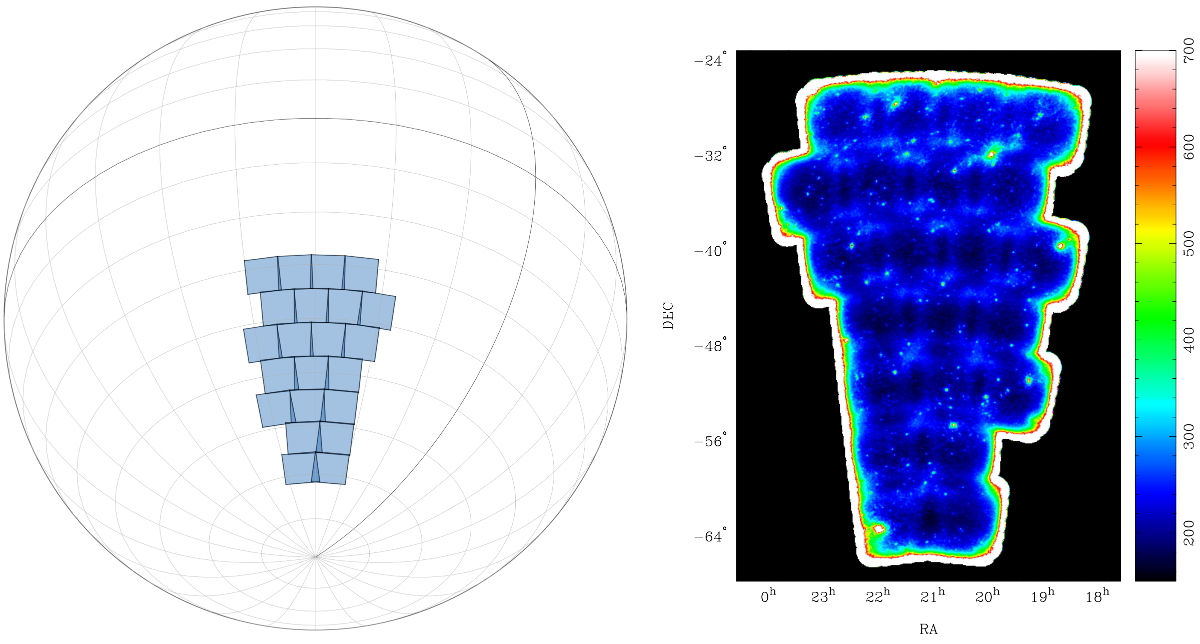

Figure 8.7: A display of multi-tile observation. The left panel shows the tile positions defined by TILESKY for a SQUARE_6X6 footprint with 1.05° beam spacing. The right panel shows a map of the image noise across the resulting mosaic. The total observing time was 7.5 hours and the colour bar gives the scale in μJy/beam. This observation was made with a relatively large beam spacing and the lower sensitivity towards the edge of each tile is evident, as is the improved sensitivity in the narrow overlap zones between tiles.

9 Polarisation

A full description of the instrumental polarisation and how it relates to image products produced by ASKAPsoft is being developed and will be available in a future version of this document. The following sections outline our current knowledge of ASKAP's polarimetric performance.

9.1 Off-axis polarisation leakage

The max-SNR beam-forming algorithm (Section 3) does not optimise for off-beam-axis beam shape or performance. Alternative algorithms have not yet been tested on the array, and will not be available for the pilot surveys. Nevertheless, the max-SNR formed beams are stable in time and similar across antennas[5], meaning it should be possible to correct for the off-axis cross-polarisation response in the image plane. This correction will not be implemented in the standard processing pipeline at the time of the pilot surveys — it will need to be experimented with in post-processing. Detailed measurements of the off-axis leakage will be conducted prior to the pilot surveys, by observing the polarisation calibrator 3C138 in a dense grid of pointings covering the square 6x6 beam footprint. The required corrections will also be derived from these data.

In general, off-axis leakage is minimised by ensuring that sources are observed relatively close to beam centres. For the pilot surveys, ACES recommends that POSSUM consider using either (a) a closepack 36 beam footprint with a pitch of 0.9°, which ensures that sources can never lie further than 0.52° from the nearest beam centre, or (b) a square 6x6 beam footprint with a pitch of 0.9° and AB interleaving, which ensures that sources can never lie further than 0.45° from

[5]A quantification of this will be presented in a future version.

the nearest beam centre (cf. the beam half power point of 0.51° at 1.4 GHz, and 0.84 ° at 850 MHz). The POSSUM SST's own simulations[6] and experiments suggest that for these limits, the inter-beam polarisation angles will be consistent, and the off-axis leakage response is generally limited to ~1% of Stokes I. More detailed characterisation of off-axis leakage will be supplied to POSSUM at the conclusion of the 3C138 grid experiment cited above.

9.2 On-axis polarisation leakage

The max-SNR beam-forming algorithm has been modified to form X and Y beams with zero relative phase by construction using an on-dish calibration (ODC) system. Since the completion of antenna integration, a functional ODC system has been an operational requirement for including an antenna in the array, so all observations should now have zero relative phase between the two polarisations.

The *uncorrected* on-axis polarisation leakages for ASKAP are generally less than 1% of Stokes I (see Figure 9.1). They are dominated by their real part, which is in turn dominated by a (nearly) DC offset component. However, a high-frequency ripple (with periods 10–20 MHz) is also evident, with measured median absolute deviations of 0.5% and 0.1% for the real and imaginary parts of the leakage, respectively. The coherent high frequency sub-structure is currently being investigated, but is possibly due to beating of the ODC signal from dish-PAF reflections. Regardless, the on-axis leakages can be corrected on a per-beam basis using the standard bandpass calibration observations of PKS B1934-638, by assuming that it is unpolarised (accurate to 0.1%). This correction will be implemented in the standard processing pipeline for pilot survey data processing.

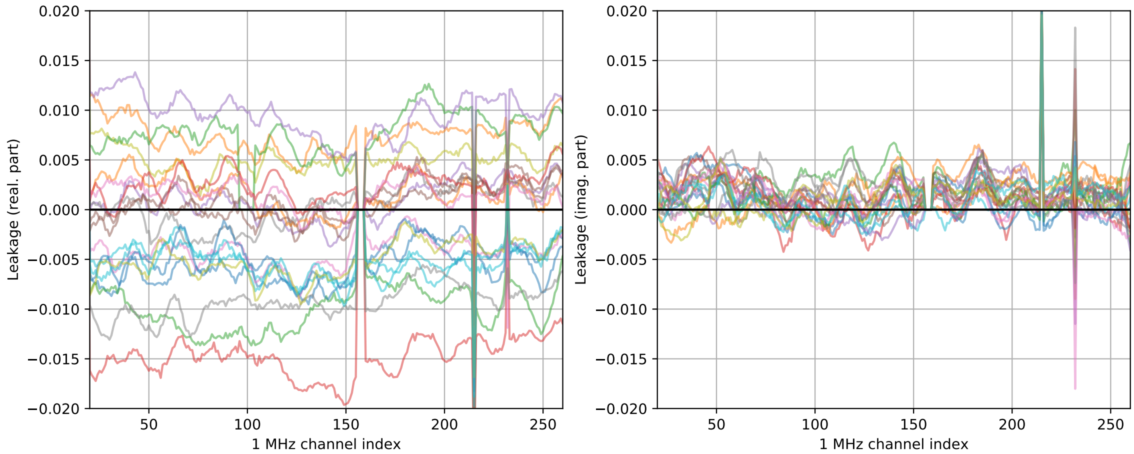

Figure 9.1: The real (left) and imaginary (right) parts of the frequency-dependent on-axis polarisation leakages, as a fraction of Stokes I. The leakages were calculated in each 1 MHz channel for 20 active antennas and a single beam from observations the calibrator PKS B1934-638 by assuming that the source is unpolarised.

9.3 Angle calibration

With the current beam-forming approach, POSSUM SST experiments have demonstrated that the relative, frequency-dependent polarisation angles of field sources are consistent between

[6] ASKAP POSSUM Report 19: Simulating the Primary Beam Response of the ASKAP Telescope over 500MHz Bandwidth, by Tony Willis

beams, and with the mosaic of beams (for a closepack 36 footprint with 0.9° pitch at 850 MHz), whenever sources are located less than $\sim$0.8° from a beam centre (roughly the half-power point). This should continue to hold true throughout the beam mosaics for the footprint configurations suggested above (closepack 36, or square 6x6 with AB interleaving).

The absolute polarisation angle calibration is not yet implemented in the processing pipeline (we note that no requirement for the post-calibration accuracy of this number was supplied by POSSUM, though 1° was listed as a desirable value[7]). Absolute angle calibration should be possible with the on-dish calibration system but has not yet been tested; we will advise on further progress here.

9.4 Faraday rotation measure precision

Experiments to determine the accuracy of ASKAP-derived Faraday rotation measures (RMs) are presently underway. The right-most panel of Figure 9.2 shows the results on one such experiment: The RM of 3C138 observed in a single beam over a three hour track and 25 degrees in elevation (dating from December 2018). The observed RM is $4.05(5)$ rad m^{-2} (standard error of the mean), compared to its published value of 0 rad m^{-2}[8], and a value of $-1.5(7)$ rad m^{-2} derived from L-band ATCA observations (dating from 2017). Observations are being planned to determine whether the offset is due to a change in the source (which has flared in recent years) or an issue with the data or observations.

9.5 Ionospheric Faraday rotation

An experiment to search for ionospheric Faraday rotation (IFR) in ASKAP data was unable to recognise the effect, since the expected magnitude was similar to the measurement errors (see Fig. 9.2). Similar joint MWA-ASKAP detection experiments are being planned, but may also result in non-detections, owing to the presently occurring minimum between Solar cycles 24 and 25. This minimum is expected to last until late 2019–early 2020 (i.e. beyond the pilot survey observations), with Solar cycle 25 rising to a peak of activity in 2024.

The POSSUM SST has supplied CSIRO with the ALBUS software package to predict the time- and position-dependant IFR. Several other such packages are also publicly available (e.g. RMextract, IonFR). Any of these can be run by CSIRO to make IFR predictions to feed into a pipeline-based correction. However: (i) the different packages differ somewhat in their predictions (Fig. 9.2), (ii) there are different ways to make the correction, with different levels of associated computing demands (e.g. operations in the visibility domain with various time resolutions, or image plane corrections), and (iii) we have failed to detect IFR in ASKAP data at levels above measurement errors to test these corrections. Thus, there are decisions to be made with respect to the implementation of IFR correction, and it will not be available in standard pipeline processing for the pilot surveys.

For the pilot surveys, ACES recommends the following: To minimise the effect of IFR, POSSUM observations should be scheduled during night time.

[7]ASKAP POSSUM Report 66: POSSUM Polarisation Characterisation Tests for BETA and ASKAP

[8]Perley and Butler (2013) “Integrated Polarization Properties of 3C48, 3C138, 3C147, and 3C286”

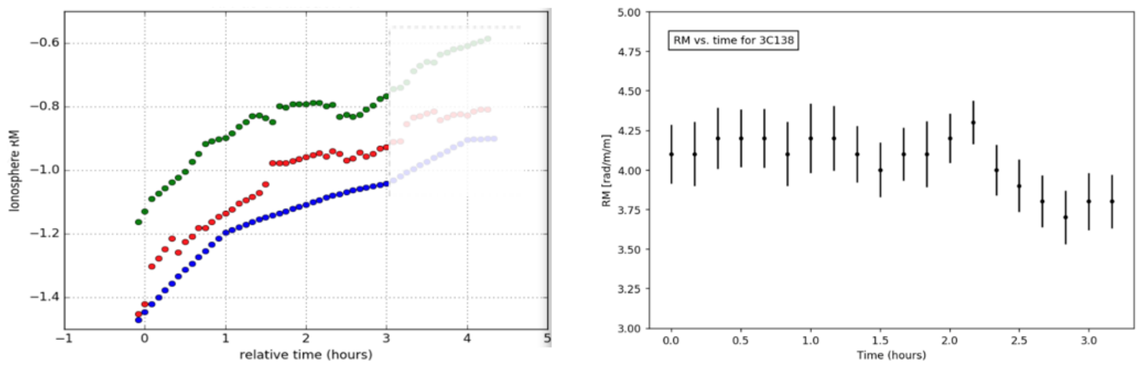

Figure 9.2: Predicted ionospheric Faraday rotation measure (RM) contribution for 3C138 during a December 2018 observing run (left panel) vs. the observed RM for 3C138 during the same time window (right). Note the different axes scaling. The time window corresponds to a change in source elevation from +20 to +45 degrees. The three different predictions originate from different software packages and ionospheric models. All models predict a $\sim$+0.4 rad m^{-2} change in the source RM, whereas the observed value which remained flat, or possibly decreased slightly.

9.6 Spectral line (zoom mode) polarisation

The polarisation performance of the high spectral resolution (zoom mode) observations has yet to be investigated.

10 Scheduling the pilot surveys

Pilot survey observations will be scheduled and conducted by the ASKAP operations team according to constraints set down by the science teams. By default we will avoid known zones of Solar interference and will attempt to avoid (or re-observe) observations in the lowest frequency range that experience RFI due to atmospheric ducting. Any additional constraints should be described in the submission of the pilot survey specifications. This should include the minimum number of antennas acceptable in terms of sensitivity or UV coverage requirements. Routine maintenance or unexpected outages may mean that the majority of observations are done with fewer than the full complement of 36 antennas.

The ASKAP control system should automatically flag antennas that experience faults. However, at the time of writing there are still several failure modes that the control system does not catch or does not deal with appropriately. ASKAP operators will endeavour to add comments to each scheduling block if they notice these conditions, but in general we will not be monitoring live data from the telescope at all hours and some problems may go initially undetected. We will inspect calibration solutions and run simple checks for bad visibility data during pipeline processing in an effort to detect any failures that reach the measurement sets. The control system will improve over time, providing more complete automatic flagging.

Pilot surveys are expected to occupy roughly 50% of the total available telescope time, alongside observatory projects, system tests (run by the ACES team) and commissioning of new modes or control system features. Some of the commissioning activities will require intensive campaigns during which no other observing will be possible (e.g. continued investigation of array phase stability using long tracks of calibrator sources), so the availability of new pilot survey data will be dependent on this.

Pilot survey throughput may be limited by processing resources and disk space. As much as possible, we will manage visibility storage space so that all pilot surveys may run in parallel. This will require that existing scheduling blocks be deleted after archiving to CASDA to make way for new data. We do not expect to be able to keep 100 hrs of visibility data from all 8 teams on disk at once. This means that teams requiring spectral line data may find themselves limited more by disk space than teams using continuum averaging mode. We expect that processing will dominate the total time required to go from observation to data release. The processing will be done by the ASKAP operations team in close collaboration with the relevant ACES team representatives, with scope for several passes if the initial results show room for improvement. However, in the case of spectral line surveys, iterative re-processing will likely cause a halt to further observing due to the amount of disk space involved.

11 Specifying the pilot surveys

It is anticipated that pilot surveys will not start until data quality in the required observing mode has been assessed and verified using one or more test observations. These test observations are requested through the science team confluence page and many have already been observed. The ACES and ASKAP operations teams will process each of these test observations and the output from the ASKAPsoft pipeline will be uploaded to CASDA so that the associated science team can assess its quality and whether all necessary data products are present.

Once test observations have verified the observing mode, the operations team will need to receive a specification for the pilot survey itself. This includes:

- A detailed description of the observing mode
- A list of all field centres, or alternatively a reference position and tiling scheme for use with survey-plan.py
- A list of any constraints on observing parameters not described above

We will develop a standard confluence page format for this specification and provide an editable page to each science team to provide broad visibility of the strategy. Each survey should contain a list of unique field IDs (as a string or number) that represent each location at which the telescope will be pointed. This allows easy reference to any individual part of the survey for re-observation or processing. This may be an exerpt from a larger tiling (of the whole sky for example) but the pilot survey specification should clearly identify 100 hrs worth of fields that will be observed first.

12 ASKAP BitBucket repository

A set of python-based tools for planning ASKAP observations has been developed and placed in a CSIRO BitBucket repository. For access and instructions on how to use the package, please contact Aaron.Chippendale@csiro.au.

A Antenna positions

The following is a list of longitude, latitude (degrees) and elevation (metres) for each antenna, taken from the facility configuration database. The values are specified with respect to the WGS-

84 coordinate reference frame. Note that these are not the geocentric X,Y,Z coordinates used by the phase tracking model, but should be sufficiently accurate for estimating U,V density and other parameters for the purposes of planning an observation.

```
ant1.location.wgs84 = [116.631424, -26.697000, 360.99]
ant2.location.wgs84 = [116.631695, -26.697119, 378.00]
ant3.location.wgs84 = [116.631785, -26.696934, 360.43]
ant4.location.wgs84 = [116.631335, -26.696684, 379.00]
ant5.location.wgs84 = [116.630681, -26.696714, 381.00]
ant6.location.wgs84 = [116.632791, -26.695960, 367.89]
ant7.location.wgs84 = [116.633861, -26.697967, 367.89]
ant8.location.wgs84 = [116.631037, -26.699120, 370.81]
ant9.location.wgs84 = [116.628965, -26.695968, 370.11]
ant10.location.wgs84 = [116.630464, -26.694442, 370.00]
ant11.location.wgs84 = [116.634127, -26.693933, 370.00]
ant12.location.wgs84 = [116.635412, -26.694578, 375.00]
ant13.location.wgs84 = [116.635835, -26.700266, 370.00]
ant14.location.wgs84 = [116.631162, -26.701120, 370.00]
ant15.location.wgs84 = [116.623984, -26.698407, 374.00]
ant16.location.wgs84 = [116.625041, -26.693651, 370.00]
ant17.location.wgs84 = [116.626445, -26.692237, 370.00]
ant18.location.wgs84 = [116.630362, -26.693575, 360.00]
ant19.location.wgs84 = [116.633628, -26.692626, 370.00]
ant20.location.wgs84 = [116.636523, -26.694137, 367.89]
ant21.location.wgs84 = [116.638331, -26.694078, 367.89]
ant22.location.wgs84 = [116.639936, -26.693956, 367.89]
ant23.location.wgs84 = [116.631398, -26.702900, 367.89]
ant24.location.wgs84 = [116.633326, -26.705803, 370.00]
ant25.location.wgs84 = [116.625276, -26.691051, 370.81]
ant26.location.wgs84 = [116.627491, -26.690931, 370.00]
ant27.location.wgs84 = [116.620692, -26.688443, 370.00]
ant28.location.wgs84 = [116.633970, -26.686153, 370.00]
ant29.location.wgs84 = [116.637135, -26.690184, 367.89]
ant30.location.wgs84 = [116.643802, -26.689790, 370.00]
ant31.location.wgs84 = [116.653776, -26.688091, 370.00]
ant32.location.wgs84 = [116.661764, -26.719712, 370.00]
ant33.location.wgs84 = [116.631594, -26.722380, 370.00]
ant34.location.wgs84 = [116.601456, -26.715008, 360.00]
ant35.location.wgs84 = [116.637135, -26.690184, 367.89]
ant36.location.wgs84 = [116.631750, -26.668198, 370.00]
```

B ASKAP frequency spectrum

ASKAP has three useable frequency bands as described in Table B.1. Table B.2 lists the available frequency resolutions.

The observing frequency is specified by the quantity `sky_frequency` (f_s) in the observing parset. The band centre frequency f_c relates to f_s as

$$f_s = f_c - \frac{1}{2}\frac{48\,\text{MHz}}{\text{Zoom factor}}$$

.

Table B.1: Operating bands and DRX and PAF filter settings. The inversion is defined at the output of the sampler.

Band (MHz)	DRX sample rate (MHz)	Nyquist band (MHz)	Inverted?	Sky frequencies (MHz)	Available bandwidth (MHz)
1	1280	640-1280 (2nd)	Y	700-1200	500
2	1536	768-1536 (2nd)	Y	840-1440	600
3	1280	1280-1920 (3rd)	N	1400-1800	400

Table B.2: ASKAP zoom modes

Zoom mode	Zoom factor	Total BW (MHz)	Channel BW (kHz)	
1	1	288	1000.	On-line averaged
1	1	288	18.516	
2	2	144	9.259	
3	4	72	4.630	
4	8	36	2.315	
5	16	18	1.157	
6	32	9	0.579	

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