



ASKAP Science Observation Guide

Prepared by the AKVET^a, SDP^b and ASKAP Operations^c teams.^d

Version 1.3 – Towards Full Survey Operations

April 2023

^aCraig Anderson, Keith Bannister, Shivani Bhandari, Aaron Chippendale, Jimi Green, Karen Lee-Waddell, Emil Lenc, David McConnell, Elizabeth Mahony, Andrew O'Brien, Chenoa Tremblay

^bEric Bastholm, Juan-Carlos Guzman, Daniel Mitchell, Steve Ord, Wasim Raja, Max Voronkov, Mark Wieringa

^cVanessa Moss, Matthew Whiting, Minh Huynh, Aidan Hotan

^dContact authors: Aidan.Hotan@csiro.au, Vanessa.Moss@csiro.au and Matthew.Whiting@csiro.au

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

Copyright and disclaimer

© 2022 CSIRO To the extent permitted by law, all rights are reserved and no part of this publication covered by copyright may be reproduced or copied in any form or by any means except with the written permission of CSIRO.

Important disclaimer

CSIRO advises that the information contained in this publication comprises general statements based on scientific research. The reader is advised and needs to be aware that such information may be incomplete or unable to be used in any specific situation. No reliance or actions must therefore be made on that information without seeking prior expert professional, scientific and technical advice. To the extent permitted by law, CSIRO (including its employees and consultants) excludes all liability to any person for any consequences, including but not limited to all losses, damages, costs, expenses and any other compensation, arising directly or indirectly from using this publication (in part or in whole) and any information or material contained in it.

Contents

1 Introduction	1
2 ASKAP operations workflow	1
2.1 Overview	1
2.2 Scheduling ASKAP observations	3
2.3 Specifying survey projects	3
2.4 Quality gates	3
3 ASKAP telescope specifications	4
3.1 The array	4
3.2 Operating frequencies	4
3.3 Polarisation	5
3.4 Primary beamforming	5
3.5 Effect of beamforming frequency interval	6
3.6 Beam footprints	7
4 ASKAP polarisation characteristics	7
4.1 On-axis polarisation leakage	7
4.2 Off-axis polarisation leakage	8
4.3 Angle calibration	9
4.4 Faraday rotation measure precision	11
4.5 Ionospheric Faraday rotation	12
5 ASKAP data processing	13
5.1 Ingest modes	13
5.2 Science data processing pipeline	13
6 CSIRO ASKAP Science Data Archive (CASDA)	15
7 ASKAP performance	15
7.1 Sensitivity	15
7.2 Field of view	15
7.3 ASKAP imaging performance	16
7.3.1 Image noise compared to thermal predictions	16
7.3.2 Dynamic range in the presence of bright sources	18
7.3.2.1 Artifacts at the level of 1% in wide-field images	18
7.3.2.2 Incomplete calibration of ASKAP data	18
7.3.3 Performance near the celestial equator	19
7.4 Spectral ripple	19
7.5 Astrometry and frequency accuracy	21
8 Radio Frequency Interference (RFI)	21

9 Sampling the sky	23
9.1 Tiling the whole sphere	24
9.2 Tiling polygons	25
9.3 The optimum tile size	25
9.4 Example results from test RACS observations	26
A Software tools for observation planning	31
B Project codes	31
C Antenna positions	32
D ASKAP polarisation	34
E Effects of increasing the number of beams	36

1 Introduction

This document provides a comprehensive guide to planning and interpreting ASKAP observations. It contains information about ASKAP's configuration, capabilities, operating modes and parameters. We also document measured performance and describe the telescope's autonomous operations model. After reviewing the various configuration options, we present example survey strategies appropriate for ASKAP's large field of view, including options for tiling the phased array feed beam footprint over large areas of sky.

A more detailed description of the telescope's hardware and the observatory's infrastructure was published in Hotan et al. [2021], which may be a useful companion to this guide.

2 ASKAP operations workflow

ASKAP is operated by the Australia Telescope National Facility (ATNF) and was designed to be a survey telescope, dedicated primarily to large-scale Survey Science Projects (SSPs). Due to the large data rates associated with having a wide field of view, ASKAP's workflow includes an integrated super-computing pipeline that produces science ready data products which are accessible via a public archive. 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 and is responsible for all data processing.

Incorporating calibration and imaging into ASKAP operations has proven to be challenging. The original real-time processing model proved impossible to realise for early science and the first generation of pilot surveys due to lack of a sky model for continuous calibration and the need to thoroughly test and tune processing algorithms using multiple passes over the data.

Pilot survey data were processed in batch mode, exchanging operational efficiency for the ability to re-process data until image quality criteria were met. Sustaining full survey operations requires processing to keep up with the incoming data rate, since the disk buffers only support a few days of observing (depending on the mode).

Some fraction of ASKAP's time may be available for guest science projects of smaller scope than the SSPs. Applications for guest science time will be made via the central ATNF proposal process and assessed by the ATNF Time Assignment Committee. ATNF will announce the availability of guest science time on ASKAP via its usual mailing lists and communication channels and applications will be managed via the OPAL¹ web interface.

2.1 Overview

The atomic unit of ASKAP observations is known as a scheduling block. Each scheduling block is a single continuous integration of one position on the sky, with up to 36 beams recorded simultaneously. The only exception is rapid interleaving mode, in which case two or three nearby locations are observed in an alternating pattern that repeats several times over the full duration. Other, more complicated observing modes are available internally (such as beamforming and holography for maintenance of the phased array feed beams) and advanced procedures could be developed for exceptional cases in consultation with the operations team if there is a compelling reason. In general, the single target per scheduling block approach optimises the combined scheduling, processing and archiving workflow.

Each scheduling block is created with a set of parameters that determine the telescope's config-

¹<https://opal.atnf.csiro.au/>

uration for that observation. In addition, scheduling blocks are associated with a survey project ID code (e.g. AS201) that defines a set of constraint parameters, which represent the state of the system under which the observations could be scheduled. The core parameters that determine a survey specification are described here and also in another document, the ASKAP Survey Plan². The science data processing pipeline also accepts a vast array of parameters that can be tuned to optimise imaging and tailor the output data products. These parameters and the pipeline workflow are documented online³. Each scheduling block includes a parameter that links to a processing strategy, which is itself a list of pipeline parameters.

The default operational workflow is as follows:

- A survey team or guest science user will submit detailed descriptions of their observing strategy to the ASKAP operations team in a machine-readable format.
- ASKAP's autonomous scheduler, SAURON, will queue each nominated field for observation according to the specified parameters. Selections are made based on compatibility with the telescope's current configuration, the sidereal time, etc. Telescope activity and history will be visible to all users via the web-based Observation Management Portal⁴.
- Science data products will be created at the Pawsey supercomputing centre by the ASKAPsoft pipeline, using processing parameters specified by the associated science team and archived in a central repository. Processing is event-driven and controlled by the Central Processor Manager software. Most scheduling blocks will be processed automatically, as soon as they have been observed and data has been automatically transferred from the ingest cluster to the processing buffer.
- Outputs from the ASKAPsoft pipeline will be automatically uploaded to CASDA⁵. Virtual Observatory compatible events⁶ are published when new data become available for validation.
- Science teams will conduct quality checks and validate each CASDA deposit using a web interface or script, ideally within 2 weeks of the deposit.
- The ASKAP operations team will release or reject validated scheduling block data, typically within 48 hours of receiving the validation report, based on the quality assessment.
- Science teams may create additional, value-added data products (e.g. additional catalogues, mosaics, etc.) and upload these to CASDA with an optional embargo period.

The set of data products uploaded to CASDA will always include calibrated visibilities with 1 MHz frequency resolution. This will allow science teams to continue experimenting with imaging methods using their own computing resources, with ASKAPsoft or other tools. ASKAP has dedicated access to a portion of the Pawsey Setonix platform for operational processing, but there will also be resources available for general use and science post-processing via the JA3 project.

Survey science teams may need to apply for additional post-processing resources. Collaboration with other facilities and the Australian SKA Regional Centre is encouraged, however the use of super-computing resources outside the Pawsey centre may require significant data transfer,

²<https://confluence.csiro.au/display/askapsst/Full+Survey+Plan>

³<https://yandasoft.readthedocs.io/en/latest/introduction/index.html>

⁴<https://apps.atnf.csiro.au/OMP/login.jsp>

⁵<https://data.csiro.au/collections/#sciDomainSearch>

⁶https://casda.csiro.au/casda_data_access/observations/events

which should be considered when planning resource requirements.

Some observing modes do not require visibilities. For example, the CRAFT project obtains high time resolution data from an alternative part of the system and operates its own processing pipeline. These alternative modes are run in collaboration with the ASKAP operations team, but support is shared with the associated team and the outputs may not be suitable for CASDA. We still encourage all teams to consider how such alternative data products can be archived for public release in a timely manner.

2.2 Scheduling ASKAP observations

Survey observations will be scheduled autonomously according to constraints set down by the science teams, technical constraints arising from system configuration, and science priorities determined by an external review panel. Science-oriented constraints should be described in the submission of survey specifications. These include the minimum number of antennas acceptable in terms of sensitivity or (u,v) 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.

Scheduling is done autonomously by a software package called SAURON. Autonomous scheduling selects fields from a large pool of possibilities based on quantitative metrics and is robust and efficient, eliminating the possibility of human error when managing thousands of fields. Autonomous reactive scheduling means that we do not know in advance when an individual field will be observed.

Ultimately, the CASDA validation process provides an opportunity to reject an observation that does not meet requirements. Rejected fields will be returned to the survey pool and re-observed.

2.3 Specifying survey projects

For the purpose of input to SAURON, surveys are determined by the following:

- A detailed description of the observing mode (see all of Section 3)
- A footprint configuration (see Section 3.6 in particular)
- A list of all field names, locations and associated processing strategies
- A list of any constraints that should be respected by SAURON

Each survey should contain a list of field IDs (as a string) that represent each location at which the telescope will be pointed. This allows easy reference to any individual part of the survey. This may be an excerpt from a larger tiling (of the whole sky for example). Each field can have a unique processing strategy, but we recommend keeping the total number of processing strategies to a minimum where possible. These processing strategies are stored in an online repository ⁷ to aid with configuration management.

2.4 Quality gates

In general, it will only be possible to process a scheduling block once during normal operations. Pilot Surveys Phase II introduced the concept of quality gates, which are observations that block further observing in the same mode and can be kept on disk and iterated over until the data products are deemed to be of suitable quality for the science goal.

⁷<https://bitbucket.csiro.au/projects/ASKAPSDP/repos/askap-sst/browse>

Quality gates are not part of the full survey workflow, but may be triggered if circumstances demand (e.g. we experience a high rate of CASDA rejections, or commission a new mode).

3 ASKAP telescope specifications

3.1 The array

ASKAP is an array of 36 antennas, each 12 m in diameter. Detailed information about the configuration can be found on the ATNF web pages⁸




Figure 3.1: ASKAP array configuration for three different distances from the core, with the outermost antennas shown in the left panel and the innermost on the right.

By default, data from all baselines are recorded. Distance thresholds in (u,v) coordinates can be applied during processing to exclude short or long baselines. 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''.

3.2 Operating frequencies

ASKAP is sensitive to frequencies between 700 and 1800 MHz, with an instantaneous bandwidth of up to 288 MHz. The observed band must lie in one of three ranges determined by the ASKAP hardware, as described in Table 3.1. See also Figure 7.1. Frequency tuning is done by selecting a subset of channels from the output of the first stage filterbank, which has 1 MHz channel resolution. Therefore, frequency tuning steps are also limited to 1 MHz, even for zoom modes. Bands 1 and 2 are inverted upon sampling, but get re-ordered before data are written out. All ASKAP visibilities are recorded in ascending frequency order.

Table 3.1: ASKAP's frequency band and filter specifications.

Band	Name	DRX sample rate (MHz)	Nyquist band (MHz)	Sky frequencies (MHz)	Tuning range (MHz)
1	Low	1280	640-1280 (2nd)	700-1200	500
2	Mid	1536	768-1536 (2nd)	840-1440	600
3	High	1280	1280-1920 (3rd)	1400-1800	400

Table 3.2 lists the available frequency resolutions. The ASKAP correlator produces 15554 channels with 18.5 kHz width, but it is also possible to do on-the-fly averaging as the data are

⁸<https://www.atnf.csiro.au/projects/askap/config.html>

recorded. This is offered as an optional “continuum” observing mode, where the standard frequency resolution is averaged to 1 MHz channels over 288 MHz of bandwidth, reducing the data rate by a factor of 54. Continuum imaging is typically done using two Taylor terms, though it is possible to create a continuum cube with 1 MHz channels. Regardless of the observing mode, we archive the calibrated visibilities with 1 MHz frequency resolution.

The observing frequency is specified by the quantity `sky_frequency` (f_p) in the scheduling block parameter set. At the time of writing, this is offset from the true sky centre frequency (f_s) as follows:

$$f_s = f_p + \frac{1}{2} \frac{48 \text{ MHz}}{\text{Zoom factor}} - 0.5 + f_z$$

Here, f_z is a correction factor that depends on the zoom mode, as shown in the last column of Table 3.2:

Table 3.2: ASKAP zoom modes

Zoom mode	Zoom factor	Total BW (MHz)	Channel BW (kHz)	Offset f_z (MHz)
Standard (Zoom 1x)	1	288	18.516	-0.5
Zoom 2x	2	144	9.259	0.5
Zoom 4x	4	72	4.630	0.0
Zoom 8x	8	36	2.315	0.0
Zoom 16x	16	18	1.157	0.0
Zoom 32x	32	9	0.579	0.0

The effective centre frequency of ASKAP data may be further influenced by flagging strategies employed in the processing pipeline.

3.3 Polarisation

ASKAP is sensitive to orthogonal linear polarisations, referred to as X and Y ⁹, and records all correlation products XX , XY , YX , YY . These are transformed to the four Stokes parameters I , Q , U and V by ASKAPsoft. The transformation to Stokes quantities in the antennas’ reference frame is determined by the feed geometry and is given in Appendix D. Transformation from this to the celestial reference frame requires correcting for the roll axis angle of the antennas.

The X and Y polarisation beams are aligned in phase during weight calculation [Chippendale and Anderson, 2019] and residual on-axis leakage can be corrected as part of standard calibration procedures. Detailed off-axis polarisation performance is measured using holography and corrections for variations over the field of view are made as part of the linear mosaicking process when matching holography is available.

3.4 Primary beamforming

ASKAP’s phased array feeds give each antenna a field of view of roughly 30 square degrees. 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 with nearly uniform

⁹ X and Y are embedded in all ASKAP system descriptions and software in spite of them rarely corresponding to horizontally and vertically polarised waves

sensitivity. At higher frequencies this is not the case and there is a corresponding decrease in survey speed. Science goals that require Nyquist spatial sampling may additionally need to interleave a single footprint using several closely-space field centres.

The maximum sensitivity algorithm [Hotan et al., 2014] is used when forming beams. 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. Using holography and source flux comparison we have determined that the most significant difference between maximum sensitivity beams and a simple Gaussian approximation is an overall beam width scale factor of roughly 10%. This is included in the default circular Gaussian primary beam correction available in ASKAPsoft. There are also significant departures from circular symmetry, most notably elongation due to coma distortion for beams near the edge of the field.

ASKAPsoft can apply an array-average of the holography-determined primary beam shapes during mosaicking, which improves the flux scale over the entire field of view. This is the mode used for typical science observations. Each newly-formed beam footprint is measured using holography so that accurate primary beam corrections are available. The association between a science field and its corresponding beam shape measurement is tracked by the processing manager. We use an on-dish calibration system to maintain the characteristics of a given set of beams over time, so holography is only needed when a new set of beams are made (roughly every few months). Some hardware level changes (such as LNA failure) may still cause differences between the true beam shape and the matching holography measurements.

There may also be some variation between beams on different antennas, or between beam weight updates, which is not currently accounted for. It is theoretically possible to enforce beam self-similarity by applying additional constraints to the algorithm, but at some cost to sensitivity. This may be investigated in future if beam variations become a limiting factor. In general, when mosaicking with matching holography measurements, we expect any flux scale errors due to primary beam correction to be less than a few percent even at the edges of the field, and much better elsewhere.

3.5 Effect of beamforming frequency interval

Beamformer weights are calculated for frequency intervals 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 accidental removal of spectral lines with widths larger than or equal to this interval due to the way bandpass calibration is done. 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 the 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, though with a visible discontinuity in phase, amplitude or both at the interval boundaries. 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 during weights calculation.

Since each PAF port has an intrinsic phase slope with respect to its neighbours, using larger beamforming intervals does run the risk of degrading beam performance across the band when large relative delays are present. Inter-port delays are a function of both the PAF geometry and

the digital system configuration, so they are not fully predictable. Improvements to the digitiser reset mechanism have reduced relative port delays to a minimum, but we recommend using 5 MHz intervals as a good compromise that provides sensitivity to broad lines with less risk of degrading the bandpass.

3.6 Beam footprints

ASKAP beams are arranged into patterns within the field-of-view which we refer to as “footprints”. Two 36-beam footprints are in common use, and most imaging experience has been gained using these. The diagrams in Figure 3.2 give the numbered beam positions within the footprints.

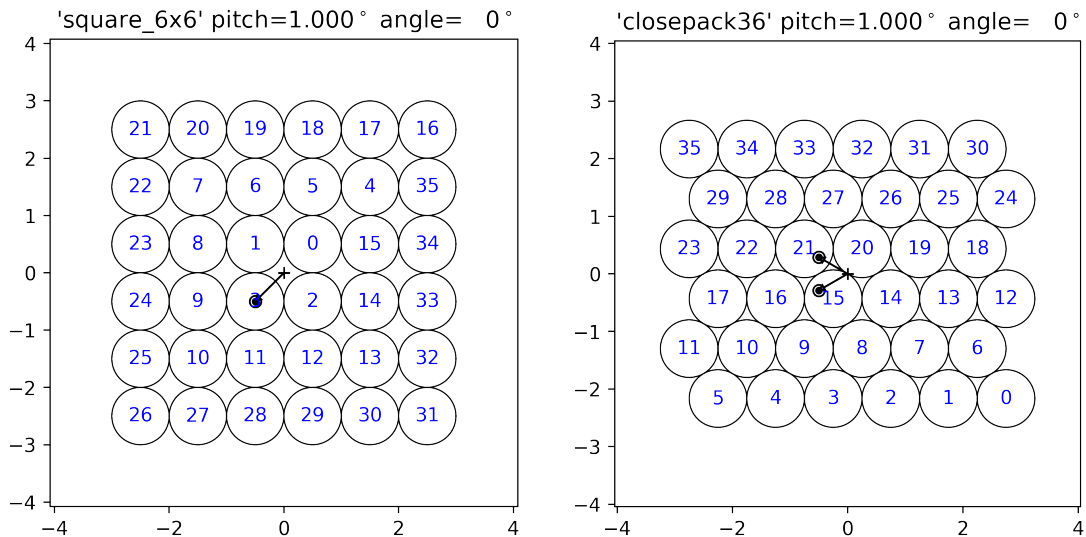


Figure 3.2: The two commonly-used footprints: `square_6x6` (left) and `closepack36` (right). In this case both have beam spacings (pitch) of 1.00° and a position angle of zero. The scales are in degrees, and celestial north (west) is to the top (right) of both diagrams. Lines radiating from the centre indicate the shift(s) required when interleaving the footprint to achieve more uniform sensitivity.

4 ASKAP polarisation characteristics

A description of ASKAP’s instrumental polarisation and how it relates to image products produced by ASKAPsoft appears in Appendix D. The following sections outline our current knowledge of ASKAP’s polarimetric performance, but this is an active topic of research and can therefore change rapidly.

4.1 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 4.1). 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 evi-

dent, 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 band-pass calibration observations of PKS B1934-638, by assuming that it is unpolarised (accurate to 0.1%).




Figure 4.1: 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.

4.2 Off-axis polarisation leakage

The max-SNR beam-forming algorithm (Section 3) does not optimise for beam shape or off-beam-axis performance. However, the max-SNR formed beams are expected to be stable in time and similar across antennas¹⁰, so a correction of the off-axis polarisation response is possible in the image plane. It is important to note, however, that when beamformer weights are refreshed we have found that the beams can change slightly in position and shape. It is therefore critical to match the leakage model to the associated beams used for observing when applying an image-plane correction.

To date, we have only investigated leakage from Stokes I into Q , U , and V ; which is expected to be the dominant leakage systematic. The Observatory performs holographic measurements of the primary beams (on unpolarised reference sources) after every beam refresh. In this way, holographic measurements are available for every observation. The ASKAPsoft tool LINMOS has been updated to accept primary beam models of the form produced by the ACES python tools¹¹ and apply a leakage correction to a given set of images. Additionally, both the RACS and POSSUM teams have also demonstrated that, given a sufficient number of observations, primary beam models can be derived from field sources themselves.

We have used models of field sources as an effective statistical probe of ASKAP's off-axis leakage response (see Figures 4.2 and 4.4) and as well as holographic measurements (Figure 4.3). Analysis of these data indicates that:

- The severity of off-axis effects apparent in any given beam depends on its distance from

¹⁰A quantification of this will be presented in a future version.

¹¹<https://bitbucket.csiro.au/projects/ACES/repos/aces-apps>

the optical axis of the telescope as well as the given frequency and Stokes parameter.

- In Stokes V the leakage from I is less than in Q or U . Further, again in Stokes V , the inner beams are least affected. Analysis suggests that leakages of $\sim 0.1\%$ of Stokes I can be attained with typical processing.
- For beams further from the optical axis, the apparent leakage varies from near-zero at the beam centre to 1% of Stokes I at beam-centric radii of ~ 0.5 degrees (where the individual contribution of a given beam to the final mosaic is dominant), and up to a few percent at or beyond the half power point. At frequencies above around 800 MHz, the spatial structure of the leakage surface appears to shrink in line with the shrinking total power beam. At lower frequencies, however, we have observed rapid sign changes of the leakages patterns.
- The leakages observed in beam mosaics are substantially smaller than those measured from the individual beam maps. This is due to two effects: Image data from beam-centric radii greater than ~ 0.5 degrees are strongly down-weighted in the mosaic, and off-axis effects in these regions differ between beams and therefore tend towards further cancellation.
- Given matched beam-forming weights between leakage models and images, residual leakage on the order of 0.1% can be achieved over a given mosaiced field with a suitable choice of footprint.

We note that, as of writing, we have identified systematic issues in Stokes V holography that has rendered those data unsuitable for leakage corrections. Until the source of this issue is understood and rectified, we recommend using a field source model of the Stokes $I \rightarrow V$ leakage surfaces.

Observations aimed at high-fidelity reconstruction of linear and circular polarisation should pursue a conservative strategy, by observing sources relatively close to beam centres, and relatively close to the optical axis of the dishes. The AKVET team recommends either:

1. a `closepack36` footprint with a pitch of 0.9° , ensuring that all sources lie within 0.52° from the nearest beam centre; or
2. a `square_6x6` footprint with a pitch of 0.9° and AB interleaving, ensuring that all sources lie within 0.45° from the nearest beam centre.

The full width at half maximum (FWHM) is 0.51° at 1.4 GHz, and 0.84° at 850 MHz. The POSSUM SST's own simulations¹² 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 $\sim 1\%$ of Stokes I .

4.3 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 beams, and with the mosaic of beams (for a `closepack36` footprint with 0.9° pitch at 850 MHz), whenever sources are located less than $\sim 0.8^\circ$ 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 (`closepack36`, or `square_6x6` with AB interleaving).

¹²ASKAP POSSUM Report 19: Simulating the Primary Beam Response of the ASKAP Telescope over 500MHz Bandwidth, by Tony Willis




Figure 4.2: Maps of the fractional Stokes $I \rightarrow V$ leakage for the 36 ASKAP beams in a `closepack36` footprint, averaged over 288 MHz centred on ~ 1.3 GHz. These maps are derived from point sources observed in a single beam-forming weight in RACS-mid (Duchesne et al. submitted). The surfaces shown are fitted Zernike polynomials to stacked measurements of V/I for sources detected in multiple observations.

Absolute calibration is now handled in the ASKAP processing pipeline. For scheduling blocks with $\text{SBID} \geq 16000$, the correction from instrument to sky frame is handled automatically by the ingest pipeline and is applied to the visibilities. For older observations ($\text{SBID} < 16000$), a correction for the feed rotation must be applied to obtain the true Stokes parameters in the sky frame, which is handled automatically in the ASKAP pipeline. For manual re-imaging with ASKAPsoft, the relevant `cimager`¹³ parameters are `Cimager.gridder.parotation`, `Cimager.gridder.parotation.angle`, and `Cimager.gridder.swappols`.

¹³<https://www.atnf.csiro.au/computing/software/askapsoft/sdp/docs/current/calim/cimager.html>




Figure 4.3: Maps of the fractional Stokes $I \rightarrow U$ leakage for the 36 ASKAP beams in a `closepack36` footprint, averaged over 1 MHz centred on ~ 1.3 GHz. These maps are derived directly from holography observations.

4.4 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 4.5 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 published¹⁴ value of 0 rad m^{-2} , 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.

¹⁴Perley and Butler (2013) “Integrated Polarization Properties of 3C48, 3C138, 3C147, and 3C286”




Figure 4.4: Reproduced from Duchesene et al. submitted: Residual Stokes $I \rightarrow V$ leakage in RACS-mid after a source-based correction. This map shows an equatorial representation of the residual leakage, $|V/I|$, in $\sim 13.4 \text{ deg}^2$ HEALPix bins. Bin values are the mean $|V/I|$ per bin for $S_I > 500\sigma_{\text{rms},I}$ sources. Galactic latitudes of $|b| < 5^\circ$ are enclosed by the dashed, black lines.

4.5 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. 4.5). We may revisit this in future 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. 4.5), (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, IFR correction will not be performed in standard pipeline processing.

The most accurate correction method would rely on measured ionospheric parameters made available some time after the observations took place. Since all ASKAP data must be promptly processed, it is unlikely that we will be able to implement such a correction prior to CASDA upload.

Given the required delay between ASKAP observations and accurate ionospheric measurements POSSUM has developed FRion¹⁵: a set routines for correcting ionospheric Faraday rotation

¹⁵<https://github.com/CIRADA-Tools/FRion>

in the image plane. This package uses RMextract to derive the time-dependent effects and computes an integrated correction. As described by Van Eck [2021], this approach has the advantage of only requiring image cubes with a sufficiently complete header, but comes at the cost of some sensitivity in linear polarisation.

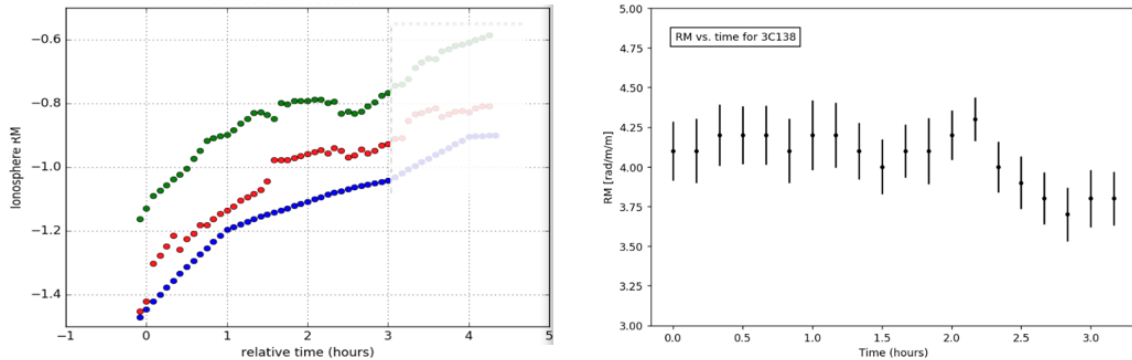


Figure 4.5: Predicted ionospheric Faraday rotation measure (RM) contribution for 3C138 during a December 2018 observing run (left panel) vs. the observed RM for 3C138 during the same time window (right). Note the different axes scaling. The time window corresponds to a change in source elevation from +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.

5 ASKAP data processing

In its wide-field spectral line mode, ASKAP currently 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, six measurement sets are written 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 as a single 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 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

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.

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 ASKAP data were determined during two iterations of Pilot Surveys, though it has proven difficult to find a single set of parameters that suit all regions of sky. Most surveys use one set of parameters for regions with few extended sources and another for regions (especially near the galactic plane) with more complex extended emission. Full documentation of the ASKAPsoft package and processing pipeline is available online¹⁶.

We show in Table 5.1 some indicative processing times for key blocks of the pipeline, as measured on the Galaxy supercomputer. These figures are derived from tests that were sized to fit all beams on Galaxy 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 processed data for pilot surveys. The specifics of the parameterisation (such as number of self-calibration loops or major cycles in the imaging/de-convolution) 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.

We are now commissioning the pipeline on Pawsey’s latest supercomputer, Setonix, with the expectation that it will be a factor of two faster. ASKAP’s dedicated allocation on Setonix has fewer nodes, but each one has significantly more cores than on Galaxy. This requires some re-organisation of the processing tasks and optimisation is ongoing.

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. Some support for installation of ASKAPsoft outside of

¹⁶<http://www.atnf.csiro.au/computing/software/askapsoft/sdp/docs/current/>

Pawsey may be available, please contact the data operations team for further information.

6 CSIRO ASKAP Science Data Archive (CASDA)

CASDA¹⁷ is intended to be the astronomer's primary interface to the telescope. Each processing parameter set defines which output files are uploaded to CASDA for permanent storage. The ASKAPsoft pipeline can also run a variety of validation scripts that create summary plots and other diagnostic material, which will be presented as html pages linked to the deposit itself.

Once a deposit completes, it must be validated by someone associated with the relevant project code. After inspecting the data and validation reports, the responsible astronomer completes a web-based validation report and assigns a status to the data files, either GOOD, BAD or UNCERTAIN. It is common to assign the same status to all files in a deposit, but CASDA does provide the capacity to assign an individual status to each file if necessary. Outputs that pass quality control (with a status of GOOD or UNCERTAIN) will be released on CASDA and made available to the public.

CASDA also provides a mechanism to upload value-added data products derived from further analysis of ASKAP deposits.

7 ASKAP performance

7.1 Sensitivity

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

$$\text{SEFD} = \frac{2kT_{sys}}{A_{\text{eff}}} \quad T_{sys}/\eta = A \frac{\text{SEFD}}{2k} \quad A_{\text{eff}} = \eta A$$

where A is the antenna area. Figure 7.1 below gives measurements of T_{sys}/η across the ASKAP band for a central beam.

Section 7.3.1 below describes how the system noise relates to the final image sensitivity.

Note that it is normal to operate with less than the full number of antennas - the allowed number can be specified as an observing constraint, either in general or with respect to antennas providing the longest baselines. The ASKAP control system should automatically flag antennas that experience faults during an observation. However, there may be failure modes that the control system does not catch. In general we will not be monitoring live data from the telescope at all hours and some problems may go undetected. The ASKAP operations team will occasionally inspect raw data diagnostics and calibration solutions and will refine automated checks of these quantities over time. We also plan to improve connections between the telescope control system and pipeline flagging directives.

7.2 Field of view

The size of ASKAP's field of view is related directly to the size of the array of receptors in the focal plane. A measurement of the field of view was made using the ASKAP-12 array and reported in ACES memo 15 [McConnell, 2017]. More recently, that measurement was compared with the

¹⁷<https://research.csiro.au/casda/>

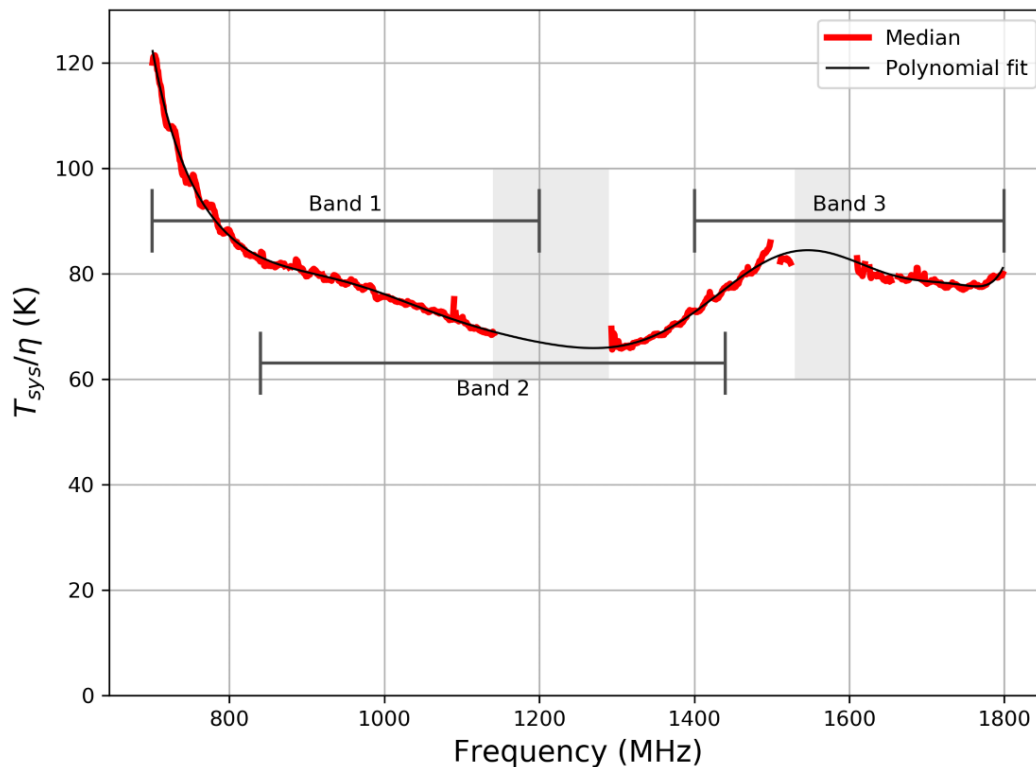


Figure 7.1: Effective system temperature T_{sys}/η across the ASKAP band. The median value over all antennas is plotted (red) for a beam close to the antennas’ boresight. The black curve is a polynomial fitted to the spectrum. Sensitivity estimates in the shaded frequency ranges are difficult because of persistent RFI. The ASKAP operating bands are shown as horizontal lines and are also listed in Table 3.1

observed variation of image noise over 36-beam tiles. Figure 7.2 shows this comparison in one dimension as an east-west profile of image sensitivity across the median of a set of 44 36-beam RACS¹⁸ tiles. In this case the RACS observations were made with the `square_6x6` footprint using a beam spacing of 1.05° . Relative sensitivity is calculated as the reciprocal of the image *rms*, divided by its maximum.

7.3 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. These will be used by the operations team to keep the telescope well-calibrated and to continue research into its performance and characteristics. The rest of this section will consider various aspects of imaging performance.

7.3.1 Image noise compared to thermal predictions

ASKAP thermal noise can be estimated for a given observation based on the SEFD measurements. For a 10 h, 288 MHz observation with 36 antennas, the expected thermal sensitivity limit (natural weighting) is $12 \mu\text{Jy beam}^{-1}$ assuming $T_{sys}/\eta = 80 \text{ K}$. Shifting the weights towards uni-

¹⁸RACS — Rapid ASKAP Continuum Survey.




Figure 7.2: The sensitivity profile over the field of view. The black line traces the observed sensitivity (see text), and the red line shows an analytic approximation of the sensitivity estimated from SEFD observations.

form, this limit is expected to increase as the effective area of the array decreases. Results given in Table 7.1 quantify the trend for ASKAP with the measured *rms* of a series of Stokes-V images made of the same data using weights ranging from natural (robustness $r = +2$) to uniform ($r = -2.0$). The Table also gives the size of the point-spread-function (PSF¹⁹) for each value of r .

Imperfections in antenna gain calibration, including the effects of time-dependent or direction-dependent gains lead to difficulties in properly de-convolving sources and an increased image noise, although in practice it is possible to get close to the expected values under most circumstances. There are some remaining concerns about dynamic range close to bright sources.

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 in future.

¹⁹Or synthesised beam in radioastronomy terminology

Table 7.1: Effect of weighting schemes. Images were made from a 10.5-hour observation on 2019-03-07 of UV ceti (declination $-17:33$) using 33 antennas and the full 288 MHz bandwidth. The noise in the Stokes-V image is tabulated as σ_{33} . Column σ_{36} gives the values expected for the full array.

	r	PSF		σ_{33}	σ_{36}
		arcseconds		μJy	μJy
Natural	+2.0	25.1	× 20.7	13	11
	+1.0	22.9	× 18.6	14	12
	+0.5	17.4	× 13.6	16	13
	0.0	12.8	× 10.7	20	16
	-0.5	9.2	× 8.6	35	29
	-1.0	8.4	× 7.1	59	48
Uniform	-2.0	8.0	× 6.7	74	61

7.3.2 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. In the deepest image, the only notable artefacts were around sources towards the edge of the field and were likely due to minor pointing errors.

7.3.2.1 Artefacts at the level of 1% in wide-field images

In general, we do not achieve the same level of dynamic range around bright sources in mosaics covering ASKAP's full field of view. Typically, we see residual artefacts (often having the form of radial streaks) around bright sources at the level of 1% of the source flux. This means the effective noise level can vary dramatically across the full field of view. In regions relatively free of bright sources (such as the GAMA23 field) the effect is barely noticeable as the 1% artefacts are close to the thermal noise limit, but in fields containing sources with flux density around 1 Jy/beam, the artefacts are very obvious. Since these artefacts appear as residuals, they have structure related to the synthesised beam shape and are therefore much more obvious close to the equator. This gives the impression of degraded imaging performance for equatorial fields, but in truth this is likely to be another symptom of a single underlying problem.

7.3.2.2 Incomplete calibration of ASKAP data

The underlying problem leading to the 1% residuals appears to be an instrumental component of the visibilities that cannot be solved using standard self-calibration techniques. Hints of such an instrumental error first arose during our attempts to fit for the positions of the antennas using long tracks of calibrator sources. Although this process converges to a solution, the residuals have not been noise-like and the quality of the fit has not met expectations. In turn, we now see that the self-calibration solutions on a typical field sometimes require significant time-dependent corrections that can be of order 100 degrees in phase over 12 hours.

Investigations into the source of the problem will continue until it is understood. It has been suggested that variations in beam shape may be responsible for calibration errors, but we believe this is not the primary cause of the 1% problem. The artefacts have similar structure across the entire field of view and appear around bright sources regardless of their proximity to the centre of any given beam. This suggests the problem is antenna-based rather than beam-based.

We also see semi-random position offsets in astronomical sources with respect to existing cata-

logues. These are typically of order the resolution element (about 2") and are most likely related to self-calibration. Due to the overhead of having multiple beams, ASKAP does not currently use a secondary phase calibrator. We rely on the solution derived from an initial bandpass calibration to remain valid for up to 24 hours and use self-calibration to correct for any deviations from this assumption. However, in the absence of a sky model, self calibration acts to place the brightest sources in the centre of a pixel, which can arbitrarily shift the image of each beam by about the pixel width. Under these circumstances, source position accuracy is not a function of signal-to-noise.

7.3.3 Performance near the celestial equator

To investigate how the PSF changes across the equator, the GAMA 09 field was observed with 36 beams centred at 09:00:00, +00:30:00 using a `square_6x6` footprint. The observations were carried out on 2019-01-22 (SBID 7705) using a frequency setting²⁰ 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", 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 7.3. The *rms* at the centre of the field is $\sim 120\mu\text{Jy/bm}$.

A second observation with all 36 antennas was taken on 2019-03-05 (SBID 8116), again using a frequency setting of 864.5 MHz, 1 MHz frequency resolution but observed for only 2 hrs. An image was produced with ASKAPsoft using the same parameters as the previous observation and reached an *rms* at field centre of $\sim 80\mu\text{Jy/bm}$. This image, shown on the right in Figure 7.3, 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 becomes more apparent at the equator due to the limited (u,v) coverage. It was hoped that these would be reduced with more careful processing (e.g. using multi-scale clean, peeling of very bright sources) but they are most likely a manifestation of the 1% problem described in the previous section, exacerbated by the complicated PSF. We expect that solving the underlying problem will therefore improve the quality of equatorial images.

7.4 Spectral ripple

During early science, spectral line survey teams identified a significant short-period ripple in ASKAP's amplitude spectra. The transfer function design of the coarse filter bank (CFB) used by ASKAP 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 \equiv 5%) 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 fringe rotation. Observations towards bright continuum sources demonstrated imperfect removal of the ripple, the residual having an amplitude of 0.2 – 0.5% of the continuum flux.

ASKAP's digital engineering team updated the correction function to solve this problem. Figure 7.4 shows a stacked spectrum from an observation of PKS B1934–638 after the fix. This places

²⁰See Table 3.1 for the relationship between the requested frequency and the band of sky frequencies observed.

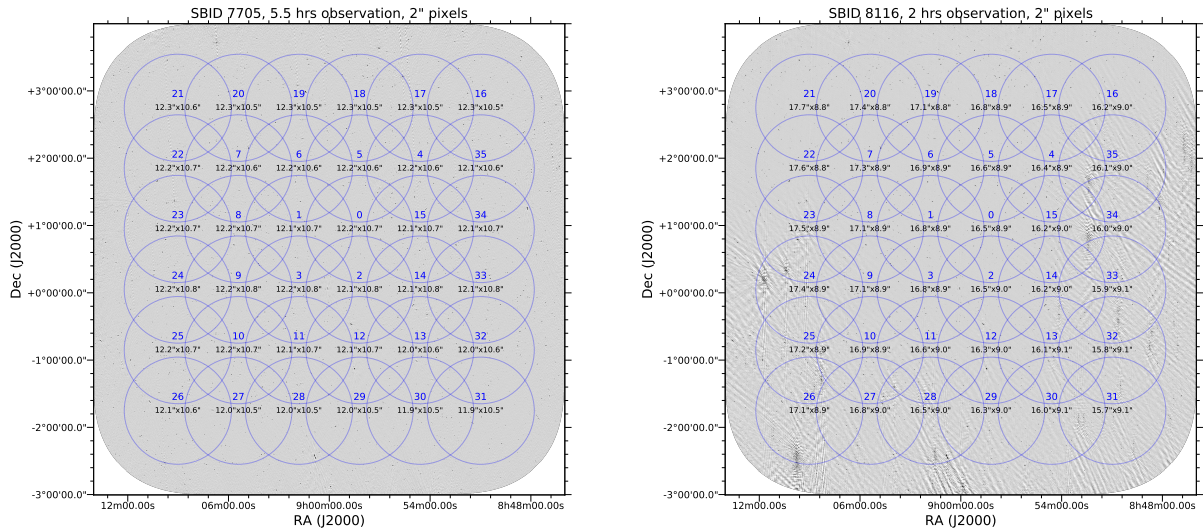


Figure 7.3: PSF for each of the 36 beams derived from a 17-antenna observation of the equatorial GAMA-09hr field (left) and a 36-antenna observation (right). The greyscale for both images is shown on a linear scale from -1 to 5 mJy. Each beam was processed independently, but all show similar PSF sizes.

an upper limit on any remaining ripple of about 0.04%

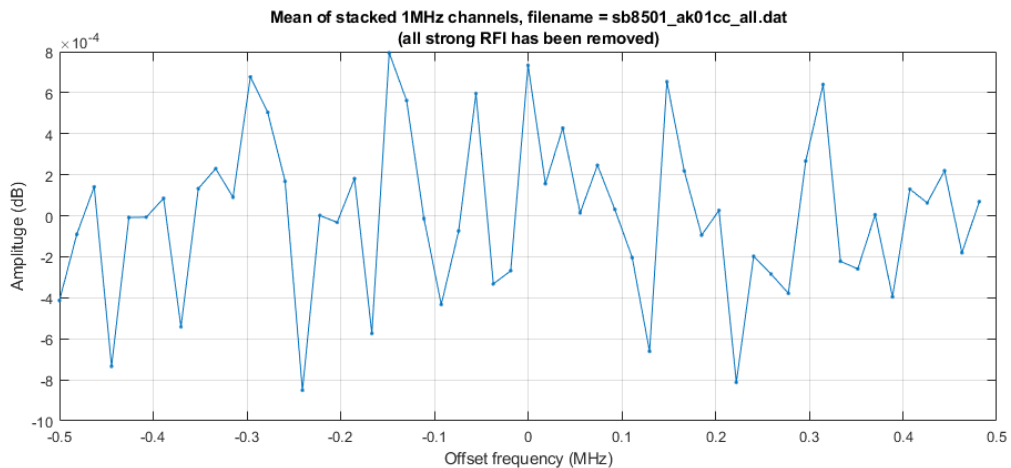


Figure 7.4: Stacked spectrum across 1 MHz intervals (54 fine channels) towards a 15 Jy source. The y-axis is relative power in dB; the peak-to-peak variation of 1.6×10^{-3} corresponds to about 0.04% of the source signal. Image credit: J. Tuthill, CASS.

More recently, another source of bandpass ripple was discovered. The on-dish calibration signal used to update beamformer weights was causing a standing wave between the PAF and dish vertex that resulted in a much longer-period bandpass ripple. For a vertex-to-focus distance of 6 m, the fundamental frequency of a standing wave in a cavity with fixed ends is 25 MHz, which clearly identifies this particular type of ripple. We modified the control system to turn off the on-dish calibration signal during science observations and only activate it during beamforming observations. Data from Pilot Survey Phase II onward should be free of the ODC-induced ripple, but we have seen evidence of a similar ripple occurring at a much lower amplitude. This may be due to the same standing wave mechanism, excited by different sources.

It is also possible for residual continuum emission or artifacts from incomplete de-convolution of

sources outside the image boundaries to cause ripple-like effects in spectral baselines. These have less predictable characteristics and must currently be mitigated after data release. Various ways to improve ASKAP's image quality are under investigation, including use of RACS as a sky model.

7.5 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 for a pointing within the Galactic plane. The observations centred on Galactic longitude 335.2° , latitude $+0.1^\circ$, had a frequency resolution of 1.157 kHz ($\approx 0.25 \text{ km s}^{-1}$) and contained 20 known OH masers with emission at the ground-state mainline transition at 1665.4018 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 (from observations made with the Australia Telescope Compact Array, ATCA).

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 arcseconds, smaller than the size of the PSF.

The ASKAPsoft pipeline offers barycentric standard-of-rest and radio local standard-of-rest (LSR) velocity corrections. To test the accuracy of the frequency, and therefore the velocity, of the OH masers detected, the observations from ASKAP were compared to observations taken with ATCA on 2019-02-29. The ASKAP data were processed through the ASKAPsoft pipeline and the ATCA data were processed through a standard MIRIAD pipeline.

A comparison of recent observations instead of detections from literature is important as the maser profiles can change depending on the activity of the object. As shown in Figure 7.5, the profiles and peak velocities match. The ATCA data have been smoothed to the velocity resolution of the ASKAP observations.

More recently, the Rapid ASKAP Continuum Survey has provided a wealth of astrometric data for comparison with existing catalogues. See McConnell et al. [2020] and Hale et al. [2021] for more information. In general, we find that ASKAP's astrometry is usually good to within about half a pixel width in the image domain, with offsets below this level introduced on a per-beam basis due to the nature of the self-calibration process and the fact that we do not observe a phase reference source.

The GASKAP-OH team also noticed that velocity corrections within ASKAPsoft were done on a per-field basis using a single time reference, which leads to slightly different frequency scales on different interleaved observations and smearing over the duration of a long track. These issues were corrected by software updates made after the conclusion of Pilot Surveys Phase II.

8 Radio Frequency Interference (RFI)

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

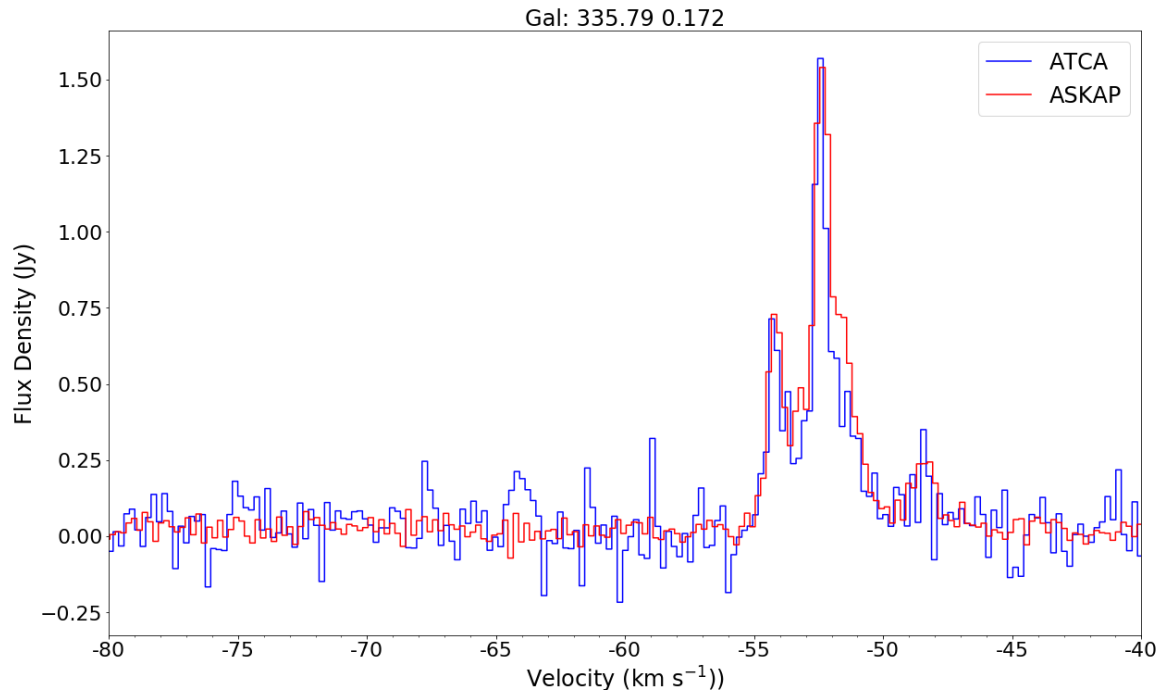


Figure 7.5: A comparison of the OH maser profile of 335.79+0.172 between ATCA (blue) and ASKAP (red), both obtained February 2019. The y-axis is the flux density (in Jy) and the x-axis is the Local Standard of Rest corrected velocity. The image shows that the velocity and peak profiles match when the ATCA data are smoothed to the same velocity as the ASKAP observations. Image credit: C. Tremblay, CASS.

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.

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 is somewhat predictable using meteorological data, but ducting is also easy to detect using the on-site RFI monitoring system. Real-time alerts are raised when ducted RFI power levels exceed certain thresholds and the autonomous scheduling system will avoid using the low-frequency band when ducting is active.

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 8.1 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.




Figure 8.1: 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.

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.

9 Sampling 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.

The script `TILE_SKY.PY` provides a means to tile an area to be surveyed using several approaches, depending on the size and location of the survey area. For efficient tiling of the whole celestial sphere, the scheme suggested by Aaron Robotham in “WALLABY Memo 22 v1.0: Sky Tiling” is used. The script also supports smaller surveys by tiling a polygon defined by the user on the sky. For intermediate sized surveys there is also the option of defining a mask to exclude parts of the whole-sphere tiling. Detailed instructions for the script are given in Appendix A.

9.1 Tiling the whole sphere

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 Λ_\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 9.1 shows the tiling scheme.




Figure 9.1: Two views of the tiled sphere: equatorial (left) and polar (right). The boundaries of the two polar zones at Λ_- and Λ_+ are indicated in the left-hand panel.

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.

There is no tiling scheme that avoids 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 Λ_\pm . However, the dependency on Λ_\pm is weak, so a value is chosen that is a compromise between size of polar cap and the overlap along latitude ranks of tiles.

The tiles can be laid down in Equatorial, Galactic or Magellanic [Nidever et al., 2008] 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 9.2 gives all-sky views for the three supported coordinate systems and a set of equatorial tiles with a shifted origin.

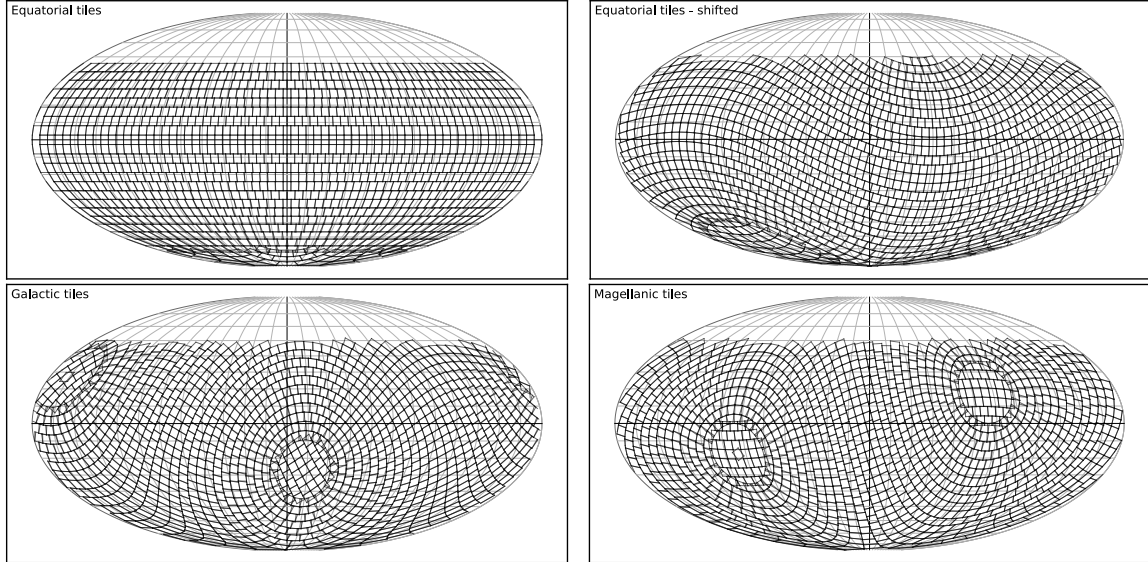


Figure 9.2: All-sky views of four different tilings, all constructed for $6^\circ \times 6^\circ$ footprints with $\Lambda_{\pm} = 72^\circ$. Top left: tiles defined in equatorial (J2000) coordinates; top right: the same set of tiles with origin shifted to $(\Theta_o, \Lambda_o) = (30^\circ, -23^\circ)$; 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.

9.2 Tiling polygons

For tiling less than the full sphere `TILE_SKY.PY` can either return a masked full-sky tiling, or construct a tiling within a user-defined area. In either case the survey area is defined by a spherical polygon; that is an n -sided region on the sphere bounded by a set of great-circle arcs. The n vertices of the polygon are provided by the user²¹.

The two options that use a polygon are best described with an illustration. The two parts of Figure 9.3 both use a polygon with five vertices, given to `TILE_SKY.PY` as five coordinate pairs (lon, lat) in degrees:

```
15.0 -22.0, 47.0 5.0, 40.0 22.0, 20. 7.0, 4.0 0.0
```

9.3 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 7.2), we have estimated the optimum beam spacing as a function of observing frequency. Figure 9.4 also shows the optimum survey speed expected and the degree of sensitivity variation (ripple) over the 36-beam mosaics. The survey speed values were computed using values of T_{sys}/η evaluated from a polynomial fit to the data in Figure 7.1, for $B = 288\text{MHz}$, $n_p = 2$, $N_a = 36$ and $\sigma_s = 100\mu\text{Jy}$, and assume that the calibration and imaging process is able to reach the thermal limits in image noise. The primary beam size used was $1.09\lambda/D$, the standard used in all ASKAP data imaging.

²¹Note that the smaller region defined by those points is tiled. On the sphere, any polygon defines two regions whose areas add to 4π .

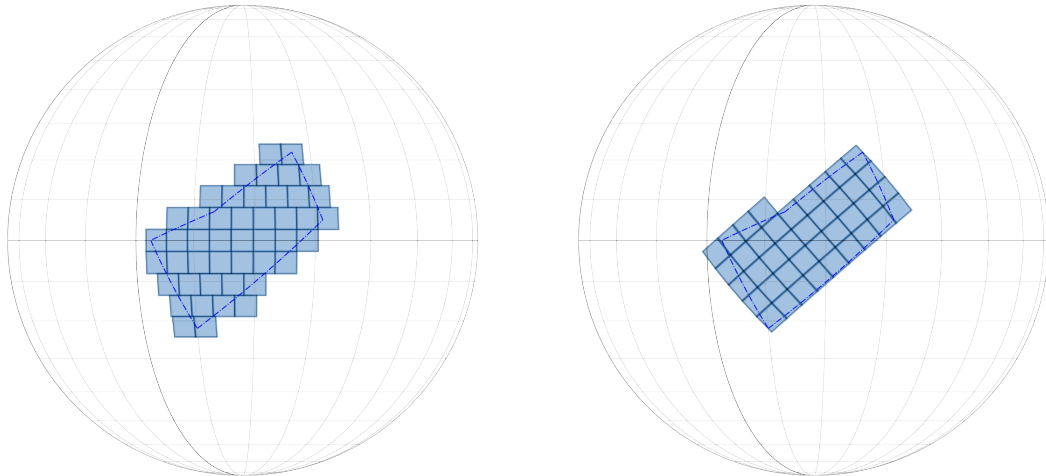


Figure 9.3: An illustration of survey specification using a polygon. On the left, the polygon is used to select only the tiles from an all-sky tiling. On the right, `TILE_SKY.PY` has determined an optimal placement of tiles within the polygon.

9.4 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 9.5 illustrates the results.




Figure 9.4: 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 $(\max - \min) / \text{mean}$ 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 $w = \frac{1}{\text{SEFD}^2}$, 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 7.1, for $B = 288\text{MHz}$, $n_p = 2$, $N_a = 36$ and $\sigma_s = 100\mu\text{Jy}$.

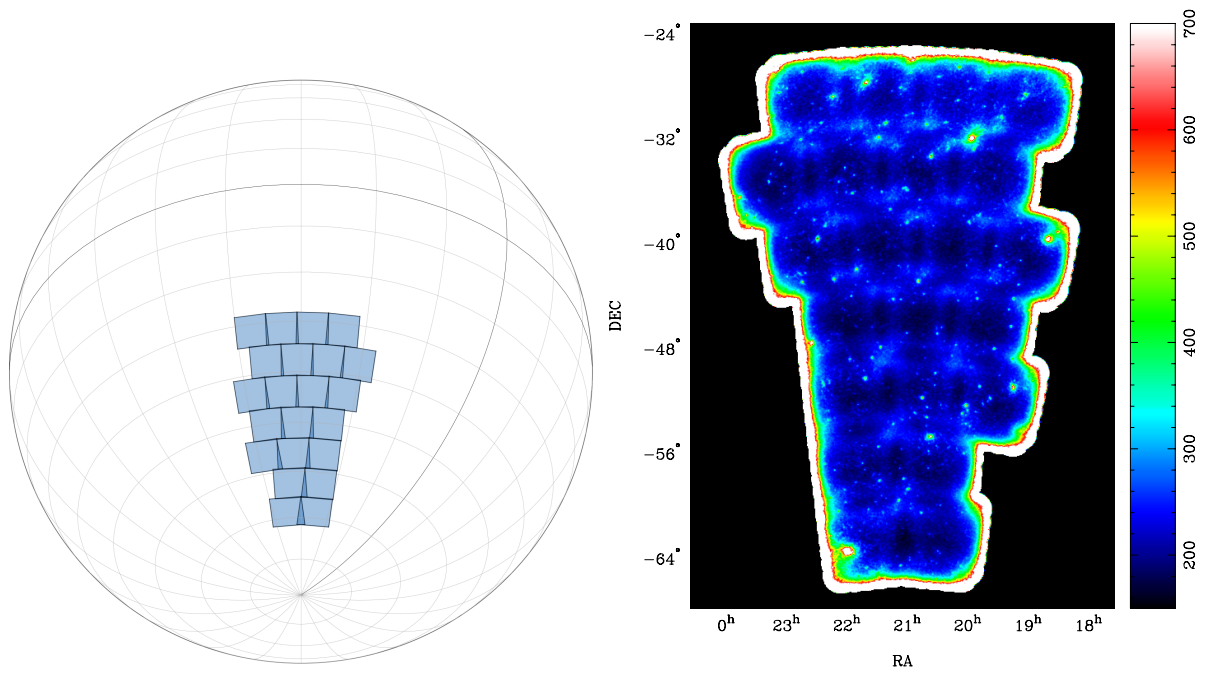


Figure 9.5: A display of multi-tile observation. The left panel shows the tile positions defined by `tile_sky` for a `square_6x6` footprint with 1.05° beam spacing. The right panel shows a map of the image noise across the resulting mosaic. The total observing time was 7.5 hours and the colour bar gives the scale in $\mu\text{Jy}/\text{beam}$. 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.

References

- A. P. Chippendale and C. Anderson. On-dish calibration of xy phase for askap's phased array feeds. ACES Memo 19, CSIRO, April 2019. URL <http://www.atnf.csiro.au/projects/askap/ACES-memos>.
- Catherine L. Hale, D. McConnell, A. J. M. Thomson, E. Lenc, G. H. Heald, A. W. Hotan, J. K. Leung, V. A. Moss, T. Murphy, J. Pritchard, E. M. Sadler, A. J. Stewart, and M. T. Whiting. The Rapid ASKAP Continuum Survey Paper II: First Stokes I Source Catalogue Data Release. PASA, 38:e058, November 2021. doi: 10.1017/pasa.2021.47.
- A. W. Hotan, J. D. Bunton, L. Harvey-Smith, B. Humphreys, B. D. Jeffs, T. Shimwell, J. Tuthill, M. Voronkov, G. Allen, and S. Amy. The Australian Square Kilometre Array Pathfinder: System Architecture and Specifications of the Boolardy Engineering Test Array. PASA, 31:e041, Nov 2014. doi: 10.1017/pasa.2014.36.
- A. W. Hotan, J. D. Bunton, A. P. Chippendale, M. Whiting, J. Tuthill, V. A. Moss, D. McConnell, S. W. Amy, M. T. Huynh, J. R. Allison, C. S. Anderson, K. W. Bannister, E. Bastholm, R. Beresford, D. C. J. Bock, R. Bolton, J. M. Chapman, K. Chow, J. D. Collier, F. R. Cooray, T. J. Cornwell, P. J. Diamond, P. G. Edwards, I. J. Feain, T. M. O. Franzen, D. George, N. Gupta, G. A. Hampson, L. Harvey-Smith, D. B. Hayman, I. Heywood, C. Jacka, C. A. Jackson, S. Jackson, K. Jeganathan, S. Johnston, M. Kesteven, D. Kleiner, B. S. Koribalski, K. Lee-Waddell, E. Lenc, E. S. Lensson, S. Mackay, E. K. Mahony, N. M. McClure-Griffiths, R. McConigley, P. Mirschin, A. K. Ng, R. P. Norris, S. E. Pearce, C. Phillips, M. A. Pilawa, W. Raja, J. E. Reynolds, P. Roberts, D. N. Roxby, E. M. Sadler, M. Shields, A. E. T. Schinckel, P. Serra, R. D. Shaw, T. Sweetnam, E. R. Troup, A. Tzioumis, M. A. Voronkov, and T. Westmeier. Australian square kilometre array pathfinder: I. system description. PASA, 38:e009, March 2021. doi: 10.1017/pasa.2021.1.
- D. McConnell. Observing with askap: Optimisation for survey speed. ACES Memo 15, CSIRO, October 2017. URL <http://www.atnf.csiro.au/projects/askap/ACES-memos>.
- D. McConnell, C. L. Hale, E. Lenc, J. K. Banfield, George Heald, A. W. Hotan, James K. Leung, Vanessa A. Moss, Tara Murphy, Andrew O'Brien, Joshua Pritchard, Wasim Raja, Elaine M. Sadler, Adam Stewart, Alec J. M. Thomson, M. Whiting, James R. Allison, S. W. Amy, C. Anderson, Lewis Ball, Keith W. Bannister, Martin Bell, Douglas C. J. Bock, Russ Bolton, J. D. Bunton, A. P. Chippendale, J. D. Collier, F. R. Cooray, T. J. Cornwell, P. J. Diamond, P. G. Edwards, N. Gupta, Douglas B. Hayman, Ian Heywood, C. A. Jackson, Bärbel S. Koribalski, Karen Lee-Waddell, N. M. McClure-Griffiths, Alan Ng, Ray P. Norris, Chris Phillips, John E. Reynolds, Daniel N. Roxby, Antony E. T. Schinckel, Matt Shields, Chenoa Tremblay, A. Tzioumis, M. A. Voronkov, and Tobias Westmeier. The Rapid ASKAP Continuum Survey I: Design and first results. PASA, 37:e048, November 2020. doi: 10.1017/pasa.2020.41.
- David L. Nidever, Steven R. Majewski, and W. Butler Burton. The Origin of the Magellanic Stream and Its Leading Arm. *ApJ*, 679(1):432–459, May 2008. doi: 10.1086/587042.
- A. Richard Thompson, James M. Moran, and Jr. Swenson, George W. *Interferometry and Synthesis in Radio Astronomy, 3rd Edition*. 2017. doi: 10.1007/978-3-319-44431-4.
- C. Van Eck. Time-integrated ionospheric correction derivation. Technical report, Dunlap Insti-

tute for Astronomy and Astrophysics, University of Toronto, 2021. URL <https://frion.readthedocs.io>.

A Software tools for observation planning

A set of python-based tooling for planning and processing ASKAP observations has recently be made publicly available for python version 3.9 and above. A pip installable wheel is made available through PyPi²², and source code is hosted on the CSIRO managed bitbucket²³. Full documentation of the code is available on Read the Docs²⁴. Some functionality depends on internal ASKAP databases which do not have an externally accessible interface. For this reason some components of the python ACES module are currently not supported if being used outside of the CSIRO network.

B Project codes

Pilot surveys phase 1 project codes:

- AS101: EMU
- AS102: WALLABY
- AS103: POSSUM
- AS104: DINGO
- AS106: CRAFT
- AS107: VAST
- AS109: FLASH
- AS111: LIGO

Pilot surveys phase 2 project code changes:

- AS105: GASKAP-OH
- AS108: GASKAP-HI

Observatory project codes:

- AS110: The Rapid ASKAP Continuum Survey (RACS)
- AS112: Survey With ASKAP of GAMA-09 and X-ray
- AS113: Other ASKAP science including tests, target of opportunity or guest science

Full survey project codes:

- AS201: EMU
- AS202: WALLABY
- AS203: POSSUM
- AS204: DINGO
- AS205: GASKAP-OH
- AS206: CRAFT
- AS207: VAST

²²<https://pypi.org/project/aces-apps/>

²³<https://bitbucket.csiro.au/projects/ACES/repos/aces-apps/>

²⁴<https://aces-apps.readthedocs.io>

- AS208: GASKAP-HI
- AS209: FLASH

ASKAP is planning to issues its first official call for guest science proposals in May 2023. These proposals will be given individual project codes.

C 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]
```

D ASKAP polarisation

Figure D.1 shows the alignment of the polarised feeds with the antenna and sky coordinates. The angles ‘pol’ and ‘PA’ are measured counter-clockwise from the directions to the Zenith and North Celestial Pole respectively. In an ASKAP antenna’s frame of reference, the X and Y feeds are inclined at -45 and $+45$ degrees to the antenna’s reference direction $\text{pol} = 0$. During normal observations, the antenna is rotated continuously about its third (“roll”) axis to keep the polarisation vectors fixed relative to celestial coordinates; the angle PA in the diagram is kept fixed to a value specified in the observing parset. Then the X and Y feeds are inclined at angles $\text{PA} - 45$ and $\text{PA} + 45$ to celestial north.




Figure D.1: Alignment of the polarised feeds with the antenna and sky coordinates. The black vectors show the antenna’s reference direction and the two polarisation vectors when the roll axis angle is $\text{pol} = 0$, and the grey vectors shows these for $\text{pol} \neq 0$.

Table 4.1 in Thompson, Moran, and Swenson [2017] gives the correspondence between polarisation products and Stokes parameters expressed in the antenna’s frame or on the sky with $\text{PA} = 0$. Referring to that table, and noting that the X dipole at -45 degrees is equivalent to one at $+135$ degrees with a change of sign, we find²⁵:

$$\begin{aligned}
 2XX &= I - U & I &= (XX + YY) \\
 2XY &= Q + jV & U &= (YY - XX) \\
 2YX &= Q - jV & Q &= (XY + YX) \\
 2YY &= I + U & V &= j(YX - XY)
 \end{aligned}$$

²⁵Here and in the following expressions we follow the convention used by ASKAPSOFT in which $I = XX + YY$, rather than $I = (XX + YY)/2$.

The set of equations on the right above can be generalised to any value of PA = θ :

$$\begin{pmatrix} I \\ Q \\ U \\ V \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 1 \\ \sin 2\theta & \cos 2\theta & \cos 2\theta & -\sin 2\theta \\ -\cos 2\theta & \sin 2\theta & \sin 2\theta & \cos 2\theta \\ 0 & -j & j & 0 \end{pmatrix} \begin{pmatrix} XX \\ XY \\ YX \\ YY \end{pmatrix} \quad (\text{D.1})$$

The ASKAPsoft pipeline implements the corrections required to deliver polarisation data consistent with IAU conventions. The same is not guaranteed when using other software.

E Effects of increasing the number of beams

It is believed to be possible to reconfigure ASKAP to allow the exchange of bandwidth for beams; that is to enable more than 36 beams at the expense of reducing the total bandwidth below 288 MHz. Before committing resources to exploring this option, the likely applications and benefits must be considered. In Figure E.1 are shown two hypothetical footprints, with 64 and 90 beams apiece. Figure E.2 compares the estimated performance of these with the standard 36-beam footprints, the estimates being made using the same assumptions of field-of-view and inter-beam correlation used in preparing Figure 9.4.

The critical quantity to compare is the equivalent field of view. It can be seen in the Figure that at the upper end of the operating range, the equivalent area increases approximately in proportion to the number of beams used.

It should also be noted that such large numbers of beams would present significant challenges to efficient operation. At present beam-forming and calibration for flux and bandpass are performed serially, so that the time taken is proportional to the number of beams to form or calibrate. For example, the calibration observation time would expand from the current 2h10m to 3h50m and 5h25m.




Figure E.1: Two proposed footprints with more than 36 beams: `square_8x8` with 64 beams and `closepack90` with 90 beams arranged in 10 rows of 9. Axis labels are in degrees.




Figure E.2: 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 $(\max - \min) / \text{mean}$ 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 $w = \frac{1}{\text{SEFD}^2}$, 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 7.1, for $B = 72\text{MHz}$, $n_p = 2$, $N_a = 36$ and $\sigma_s = 200\mu\text{Jy}$.

CONTACT US

t 1300 363 400
+61 3 9545 2176

e csiroenquiries@csiro.au

w www.csiro.au

WE DO THE EXTRAORDINARY EVERY DAY

We innovate for tomorrow and help improve today for our customers, all Australians and the world.

Our innovations contribute billions of dollars to the Australian economy every year. As the largest patent holder in the nation, our vast wealth of intellectual property has led to more than 150 spin-off companies.

With more than 5,000 experts and a burning desire to get things done, we are Australia's catalyst for innovation. WE IMAGINE. WE COLLABORATE. WE INNOVATE.