

Research Paper

Australian square kilometre array pathfinder: I. system description

A. W. Hotan¹*, J. D. Bunton², A. P. Chippendale², M. Whiting², J. Tuthill², V. A. Moss², D. McConnell², S. W. Amy², M. T. Huynh¹, J. R. Allison^{2,11}, C. S. Anderson^{2,20}, K. W. Bannister², E. Bastholm¹, R. Beresford², D. C.-J. Bock², R. Bolton², J. M. Chapman², K. Chow², J. D. Collier^{5,6}, F. R. Cooray², T. J. Cornwell^{2,7}, P. J. Diamond^{2,8}, P. G. Edwards², I. J. Feain², T. M. O. Franzen¹⁹, D. George², N. Gupta^{2,9}, G. A. Hampson², L. Harvey-Smith^{2,6,10}, D. B. Hayman², I. Heywood^{2,11,12}, C. Jacka², C. A. Jackson^{2,19}, S. Jackson¹⁸, K. Jeganathan², S. Johnston², M. Kesteven², D. Kleiner¹³, B. S. Koribalski², K. Lee-Waddell², E. Lenc², E. S. Lensson¹⁴, S. Mackay², E. K. Mahony², N. M. McClure-Griffiths¹⁵, R. McConigley⁴, P. Mirtschin¹⁶, A. K. Ng², R. P. Norris^{2,6}, S. E. Pearce², C. Phillips², M. A. Pilawa², W. Raja², J. E. Reynolds², P. Roberts², D. N. Roxby², E. M. Sadler^{2,3}, M. Shields², A. E. T. Schinckel², P. Serra^{2,13}, R. D. Shaw², T. Sweetnam², E. R. Troup², A. Tzioumis², M. A. Voronkov² and T. Westmeier¹⁷

¹CSIRO Astronomy and Space Science, PO Box 1130, Bentley, WA 6102, Australia, ²CSIRO Astronomy and Space Science, PO Box 76, Epping, NSW 1710, Australia, ³Sydney Institute for Astronomy, School of Physics A28, University of Sydney, Sydney, NSW 2006, Australia, ⁴CSIRO Astronomy and Space Science, PO Box 2102, Geraldton, WA 6530, Australia, ⁵The Inter-University Institute for Data Intensive Astronomy (IDIA), Department of Astronomy, University of Cape Town, Private Bag X3, Rondebosch 7701, South Africa, ⁶School of Science, Western Sydney University, Locked Bag 1797, Penrith, NSW 2751, Australia, ⁷Tim Cornwell Consulting, 17 Egan Crescent, Sandbach, CW11 1LD, UK, ⁸SKA Organisation, Jodrell Bank, Lower Withington, Cheshire, SK11 9FT, UK, ⁹Inter-University Centre for Astronomy and Astrophysics, Post Bag 4, Ganeshkhind, Pune 411007, India, ¹⁰Deans Unit, Faculty of Science, Dalton Building F12 UNSW Sydney, NSW 2052, Australia, ¹¹Sub-Department of Astrophysics, Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Rd., Oxford, OX1 3RH, UK, ¹²Rhodes University, PO Box 94, Makhanda (Grahamstown) 6140, Eastern Cape, South Africa, ¹³INAF - Osservatorio Astronomico di Cagliari, via della Scienza 5, Selargius, CA 09047, Italy, ¹⁴CSIRO Astronomy and Space Science, PO Box 276, Parkes, NSW 2870, Australia, ¹⁵Research School of Astronomy & Astrophysics, The Australian National University, Canberra, ACT 2601, Australia, ¹⁶CSIRO Astronomy and Space Science, 1828 Yarrle Lake Rd, Narrabri, NSW 2390, Australia, ¹⁷ICRAR, M468, The University of Western Australia, 35 Stirling Highway, Crawley WA 6009, Australia, ¹⁸CSIRO Astronomy and Space Science, PO Box 2225, Ellenbrook, WA 6069, Australia, ¹⁹ASTRON, the Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, Dwingeloo, PD 7991, the Netherlands and ²⁰Jansky fellow of the National Radio Astronomy Observatory, NRAO, 1003 Lopezville Rd, Socorro, NM 87801, USA

Abstract

In this paper, we describe the system design and capabilities of the Australian Square Kilometre Array Pathfinder (ASKAP) radio telescope at the conclusion of its construction project and commencement of science operations. ASKAP is one of the first radio telescopes to deploy phased array feed (PAF) technology on a large scale, giving it an instantaneous field of view that covers 31 deg² at 800 MHz. As a two-dimensional array of 36 × 12 m antennas, with baselines ranging from 22 m to 6 km, ASKAP also has excellent snapshot imaging capability and 10 arcsec resolution. This, combined with 288 MHz of instantaneous bandwidth and a unique third axis of rotation on each antenna, gives ASKAP the capability to create high dynamic range images of large sky areas very quickly. It is an excellent telescope for surveys between 700 and 1800 MHz and is expected to facilitate great advances in our understanding of galaxy formation, cosmology, and radio transients while opening new parameter space for discovery of the unknown.

Keywords: radio interferometers – wide-field telescopes

(Received 23 July 2020; revised 8 December 2020; accepted 4 January 2021)

Author for correspondence: A. W. Hotan, E-mail: aidan.hotan@csiro.au

Cite this article: Hotan AW, Bunton JD, Chippendale AP, Whiting M, Tuthill J, Moss VA, McConnell D, Amy SW, Huynh MT, Allison JR, Anderson CS, Bannister KW, Bastholm E, Beresford R, Bock DCJ, Bolton R, Chapman JM, Chow K, Collier JD, Cooray FR, Cornwell TJ, Diamond PJ, Edwards PG, Feain IJ, Franzen TMO, George D, Gupta N, Hampson GA, Harvey-Smith L, Hayman DB, Heywood I, Jacka C, Jackson CA, Jackson S, Jeganathan K, Johnston S, Kesteven M, Kleiner D, Koribalski BS, Lee-Waddell K, Lenc E, Lensson ES, Mackay S, Mahony EK, McClure-Griffiths NM, McConigley R, Mirtschin P, Ng AK, Norris PR, Pearce SE, Phillips C, Pilawa MA, Raja W, Reynolds JE, Roberts P, Roxby DN, Sadler SE, Shields M, Schinckel AET, Serra P, Shaw RD, Sweetnam T, Troup ER, Tzioumis A, Voronkov MA and Westmeier T. (2021) Australian square kilometre array pathfinder: I. system description. *Publications of the Astronomical Society of Australia* 38, e009, 1–31. <https://doi.org/10.1017/pasa.2021.1>

1. Introduction

ASKAP is a Square Kilometre Array (SKA^a) precursor telescope located at Australia's SKA site, the Murchison Radio-Astronomy Observatory (MRO). This radio telescope was developed using new phased array feed (PAF) technology to achieve high survey speed by observing with a wide instantaneous field of view. The design concept is described in more detail by DeBoer *et al.* (2009) and the science goals in Johnston *et al.* (2007). Figure 1 shows a photograph of ASKAP's core taken in 2018.

After more than 10 years of design, prototyping, construction, and commissioning, ASKAP became fully operational in 2019. During construction, we developed a second generation of electronics (Hampson *et al.* 2012) based on lessons learnt from prototype hardware (Schinckel *et al.* 2011; Hotan *et al.* 2014; McConnell *et al.* 2016), along with many iterations of firmware and software. Thousands of components have been installed and connected together to form one of the most complicated and powerful radio astronomy signal processors ever developed.

The first interferometry between PAF-equipped antennas at the MRO was conducted using a prototype system known as the Boolardy engineering test array (BETA, Hotan *et al.* 2014). This consisted of six antennas fitted with first-generation receiver systems (Schinckel *et al.* 2011). Beyond teaching us how to improve ASKAP's system design, BETA contributed to astrophysical research (e.g., Serra *et al.* 2015b; Harvey-Smith *et al.* 2016; Hobbs *et al.* 2016; Heywood *et al.* 2016; Allison *et al.* 2017; Moss *et al.* 2017) including the discovery of neutral hydrogen in a young radio galaxy at redshift $z = 0.44$ through an absorption line search (Allison *et al.* 2015). BETA also demonstrated real-time spatial radio-frequency interference (RFI) mitigation (Hellbourg, Bannister, & Hotan 2016).

A diverse early science programme conducted on a subset of 12–18 antennas fitted with the second-generation PAF systems also produced a wide range of results, showcasing the utility of a wide-field radio telescope at gigahertz frequencies. See Section 15 for more details.

In 2019, we commenced a programme of pilot surveys with the full array and an all-sky survey known as the Rapid ASKAP Continuum Survey (RACS^b), demonstrating the full science capability of the telescope for the first time.

With up to 36 beams per antenna and 36 antennas, ASKAP produces a torrent of raw data (approximately 100Tbit s^{-1}). This is correlated and averaged at the observatory, producing an output visibility data stream of up to 2.4GB s^{-1} that is sent to the Pawsey Supercomputing Centre^c in Perth, some 600 km south of the telescope, for image processing. This gives the research community a hint of the challenges to come in the era of the SKA.

This paper describes the technical details of ASKAP and documents key performance metrics based on commissioning data. Future papers in the series will describe the image processing software and performance metrics in more detail. We describe ASKAP's design (Section 1.1) followed by information about the observatory site (Section 2), an overview of the system and its components (Section 3), and more detailed information about key subsystems (Sections 4–12). We report on measured performance metrics in Section 13, the site radio frequency environment



Figure 1. CSIRO's Australian Square Kilometre Array Pathfinder (ASKAP) telescope.

in Section 14, the telescope operations model in Section 15, and future upgrade options in Section 16.

1.1. Designing a radio telescope for surveys

The primary design goal of ASKAP was to maximise survey speed, the rate at which the telescope can observe a given area of sky to a certain sensitivity limit (Johnston *et al.* 2007). Compared to existing telescopes, it was clear that high survey speed could be achieved by building an array of many small antennas to keep the primary beam size large, provide sensitivity over multiple spatial scales, and achieve good surface brightness sensitivity. Such an array would have many baselines (antenna pairs) and would, therefore, produce a very large amount of data and cause the computational cost of imaging to dominate array design.

ASKAP uses PAF receivers to achieve a 31 deg^2 instantaneous field of view at 800 MHz. PAFs capture more information from each antenna and can be a more cost effective way to increase survey speed than using a much larger number of smaller antennas with cryogenic single-pixel feeds (Chippendale, Colegate, & O'Sullivan 2007).

The survey speed of a radio telescope array also depends on bandwidth (Bunton 2003a) and the precise nature of the survey parameters (Johnston & Grey 2006; Johnston *et al.* 2007), but can be broadly represented by the figure of merit (Bunton 2003a; DeBoer *et al.* 2009; Bunton & Hay 2010)

$$SS = \left(\frac{A_e}{T_{\text{sys}}} \right)^2 \Omega_{\text{FoV}}, \quad (1)$$

where A_e is the total effective area of all antennas, T_{sys} is the system equivalent noise temperature, and Ω_{FoV} is the instantaneous (processed) field of view. Effective area is related to the physical area A of the antennas via $A_e = \eta A$ and η is the antenna efficiency.

Table 1 summarises key indicative performance parameters of the ASKAP array using (1) to measure survey speed and A_e/T_{sys} to measure sensitivity. A nominal measured value of 75 K is used for the effective system temperature T_{sys}/η . With 36 dual-polarisation beams, ASKAP achieves a 31 deg^2 field of view at 800 MHz but this reduces to 15 deg^2 at 1700 MHz (McConnell 2017b). Field of view and survey speed may be increased at the high end of the band in the future by processing more than 36 beams. See Section 13 for more details on the definition and measurement of sensitivity, field of view, and the resulting survey speed.

^a<https://www.skatelescope.org>.

^b<https://www.atnf.csiro.au/content/racs>.

^c<https://pawsey.org.au>.

Table 1. Key parameters of the ASKAP telescope

Number of antennas	36
Antenna diameter	12 m
Focal ratio f/D	0.5
Total collecting area	4071.5 m ²
Maximum baseline	6 km
Angular resolution	10'' at 1 GHz
Observing frequency	0.7–1.8 GHz
Processed bandwidth	288 MHz
Frequency channels	15 552
Frequency resolution	18.5–0.58 kHz
Effective system temperature	75 k
Sensitivity	54 m ⁻¹ K ⁻¹
Dual-polarisation beams	36
Field of view ^a (800 MHz)	31 deg ²
Field of view ^a (1 700 MHz)	15 deg ²
Survey speed ^b (800 MHz)	91400 m ⁴ deg ² K ⁻²
Survey speed ^b (1 700 MHz)	44 200 m ⁴ deg ² K ⁻²

^aCalculated via (15).^bCalculated via (14).

ASKAP's large field of view and increased survey speed came at the expense of sensitivity since cooling each PAF to cryogenic temperatures was not economical on the scale required at the time of design. This resulted in a beamformed T_{sys}/η in the range 60–80 K over most of the frequency range, or a system equivalent flux density (SEFD) of approximately 1 800 Jy for a single antenna (see Section 13).

ASKAP's T_{sys} is dominated by low-noise amplifier (LNA) noise, in part because the LNA uses a transistor (ATF-35143, circa 2006) that is now quite old. Sensitivity could be improved significantly in the future by updating the room temperature LNA design with new transistors (Shaw & Hay 2015; Weinreb & Shi 2019) or by scaling up the manufacturability and affordability of cryogenic PAF technology like that under development for the Parkes 64 m telescope (Dunning et al. 2016, 2019).

Pilot survey observations (see Section 15) have tested many different observing modes, providing direct experience of the sensitivity that can be achieved within practical constraints. In many cases, we approach the thermal noise limit. However, for some observing modes (particularly with short integration times), other factors such as deconvolution errors contribute to an elevated noise floor. The values given in Table 2 should provide a realistic basis for planning observations, and may improve as updates are made to the telescope. The data used to compile Table 2 were obtained using different beam arrangements and processing strategies, so some variation is expected on top of the intrinsic spectral behaviour, as shown in Figure 22.

Importantly, the measured noise can be increased above the theoretical value by weighting the data differently to natural weighting. ASKAP uses a preconditioning approach with robust weighting. The robustness parameter has a similar effect to the description in Briggs (1995), ranging from -2.0 (uniform weighting) to 2.0 (natural weighting). The robustness values typically used for pilot survey processing were: 0.0 for continuum imaging at low frequencies, -0.5 for continuum imaging in the mid-band, and 0.5 for spectral imaging in the low and mid-bands. The

Table 2. RMS noise measured in pilot survey phase I data. RMS noise per beam is given as the minimum and average over all observations in CASDA, then the minimum scaled to a standard 1 h and 288 MHz (for continuum, above the line) or 18.5 kHz (for spectral line, below the line)

ν (MHz)	$\Delta\nu$ (MHz)	t_{int}	Min (μJy)	Avg (μJy)	Scaled (μJy)
944	288	10 h	24	37	74
864	288	12 m	202	420	90
1 368	144	8 h	40	70	80
1 665	9	8 h	186	582	95
856	18.5 kHz	2 h	4 860	5 700	6 870
1 368	18.5 kHz	8 h	2 000	2 500	5 660

noise increases by factors of approximately 2.5, 1.5, and 1.2 for robustness values of -0.5 , 0.0 , and 0.5 , respectively.

Most of ASKAP's observing time will be dedicated to large-scale survey projects and the observatory will provide science-ready data products via a public archive (see Section 12). ASKAP's deep survey projects are expected to discover millions of new radio sources in the southern sky (Johnston et al. 2007). The large instantaneous field of view also opens up new parameter space for the study of transient sources. ASKAP will excel at wide-area, high cadence surveys for slow transients; wide-area surveys for spectral line emission and absorption; rapid, wide-field searches for gravitational-wave counterparts; and the discovery and localisation of fast transients via 1 ms cadence autocorrelations and a deep voltage buffer (up to 14 s).

2. The murchison radio-astronomy observatory

ASKAP is located on a remote site in Western Australia, specifically established as the Australian radio quiet zone (WA) for the SKA and its precursors (Wilson, Storey, & Tzioumis 2013; Wilson et al. 2016). The total power in RFI signals over the ASKAP band at this site is typically more than an order of magnitude less than ASKAP's system noise power (Chippendale & Wormnes 2013). Legislation regulates the use of radio transmitters within 260 km of the site (ACMA 2014), helping to ensure that the environment will remain as favourable as possible for radio astronomy into the future. Further protection is afforded by carefully testing and, if necessary, shielding all electronic equipment installed at the site (Beresford & Li 2013).

The observatory site (see Figure 2) is roughly 315 km north-east of Geraldton, the largest town in the region. It is operated by the Australian Commonwealth Scientific and Industrial Research Organisation (CSIRO) and hosts the Murchison Widefield Array (MWA) (Tingay et al. 2013; Wayth et al. 2018) and EDGES (Bowman et al. 2018) telescopes in addition to ASKAP. In future, the low-frequency component of the SKA will be established nearby.

In order to prevent ASKAP's infrastructure from creating RFI, all processing hardware is contained within a dual-shielded central control building (Abeywickrema et al. 2013) that attenuates any radio emission by 160 dB. Signal processing and computational systems are housed inside the second layer of shielding, while office and workshop space are inside the first layer. The building provides working space for maintenance crews but no on-site accommodation. The number of people on site is kept

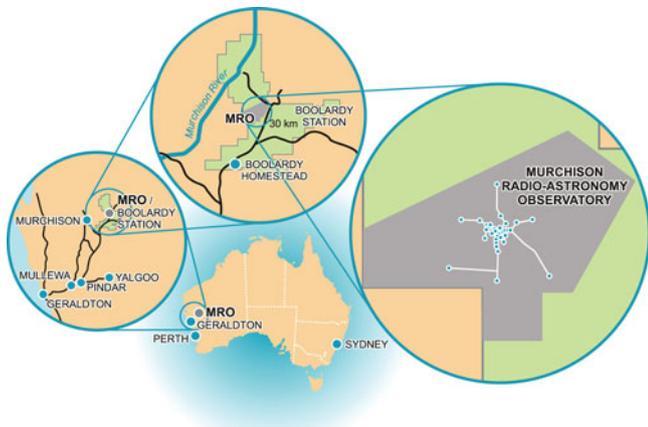


Figure 2. Diagram showing the relative location and size of the Murchison Radio-Astronomy Observatory, with the final inset showing ASKAP antennas as blue dots and service tracks as white lines.

to a minimum and crews are housed in the Boolardy station accommodation complex, roughly 40 km from the telescope.

3. System overview

Figure 3 shows the major components of the telescope and highlights how data flows from the antennas through to the image archive. Each subsystem is described in more detail in the sections below.

ASKAP consists of 36 paraboloidal reflector antennas, each 12 m in diameter with a chequerboard PAF at the primary focus. Within each PAF, the radio frequency signals received by 188 active feed elements (94 per polarisation) are converted into analogue optical signals. The optical signal from each element is transmitted over its own dedicated optical fibre back to the central control building. See Section 5 for more details.

Inside the control building, each PAF element signal is digitised and then passed through a coarse oversampled polyphase filter bank with 1 MHz channel spacing. This generates 640 or 768 channels depending on which of three frequency bands is being used. The digital receiver system selects 384 of these channels and sends them via digital optical communications links and a passive optical cross-connect to the beamformers.

The beamformers compute the weighted sum (see Section 9) across PAF elements for each frequency channel and antenna, forming 36 dual-polarisation beams over 336 MHz. The beamformed signals then pass into a fine filter bank that divides each 1 MHz channel into 54 channels of 18.5 kHz width, for a total of 18 144 fine channels. Bandwidth can be exchanged for higher frequency resolution by up to a factor of 32 using zoom modes. See Section 7 for more details. The beamformer also generates 1 ms power averages of the signal for each 1 MHz channel of each beam. This is processed on site to search for fast transients.

The fine channels are sent via digital optical links and an optical cross-connect to the correlator, which computes and accumulates the visibilities for each baseline. The correlator has the capacity to process 15 552 fine channels for a maximum bandwidth of 288 MHz. These visibilities are sent directly to the Pawsey Supercomputing Centre in the city of Perth via a long-haul underground optical fibre network, with no buffering at the observatory. Three petabytes of disk buffer is available at Pawsey, where raw visibilities are formatted and written to a number of Common

Astronomy Software Applications (CASA) MeasurementSet files (version 2, Kembell & Wieringa 2000) in parallel.

The visibility data are imaged offline in batch processing mode using a custom software package called ASKAPsoft (Guzman *et al.* 2019). The Pawsey Supercomputing Centre currently maintains a dedicated compute cluster called *Galaxy* for the purpose of imaging ASKAP science data. This system is due to be replaced with a more powerful supercomputer in 2021. See Section 11 for more details.

Data products are made available to astronomers via a public archive, described in Section 12. As shown in Figure 3, the archive can store calibrated visibility data (averaged in frequency to reduce disk space usage), mosaicked image cubes (but typically not individual beam images) and source catalogues.

3.1. Control and monitoring

ASKAP's control system is built on the Experimental Physics Instrument Control System (EPICS) library.^d Each major hardware subsystem has an associated EPICS input/output controller (IOC), which is a software application that issues commands and gathers monitoring information via network transactions. Several servers located at the MRO run these IOCs and there is a dedicated monitoring and control network in addition to the office network and astronomical data network. The design of the control system is described in Guzman & Humphreys (2010).

The digital signal processing hardware was designed with a common command and control interface that translates instructions from the network into firmware registers on field-programmable gate arrays (FPGAs) throughout the system. A common software library was written to support the hardware layer and this library is used by the EPICS IOCs. Various other commodity electronics devices (such as drive controllers and power supplies) are also used, requiring different network protocols at the IOC level.

3.2. Power infrastructure

Due to its remote location, the MRO is not connected to a mains power grid. The site is supplied by a hybrid solar–diesel power station purpose-built with sufficient RFI shielding to protect the observatory from stray radio emissions. The power station is located approximately 6 km from the core of the telescope to further aid in RFI mitigation. The solar component of the power station can supply 1.6 MW (enough to run the entire site during the day) and is linked to a battery bank with 2.5 MWh capacity. Several redundant diesel generators are the main power source at night.

Power is transmitted via underground high-voltage cable to the central control building. Distribution to the antennas is on six independently switched underground tracks that roughly follow the surface roads and power several antennas each. A transformer at each antenna pad steps the high-voltage input down to three-phase 415 V for the equipment in the antenna pedestal.

3.3. Cooling infrastructure

ASKAP's digital signal processing hardware in the central control building consumes 280 kW of power and requires a commensurate amount of cooling. This is achieved by circulating chilled water to each electronics cabinet in the central building, where internal

^d<https://epics.anl.gov>.

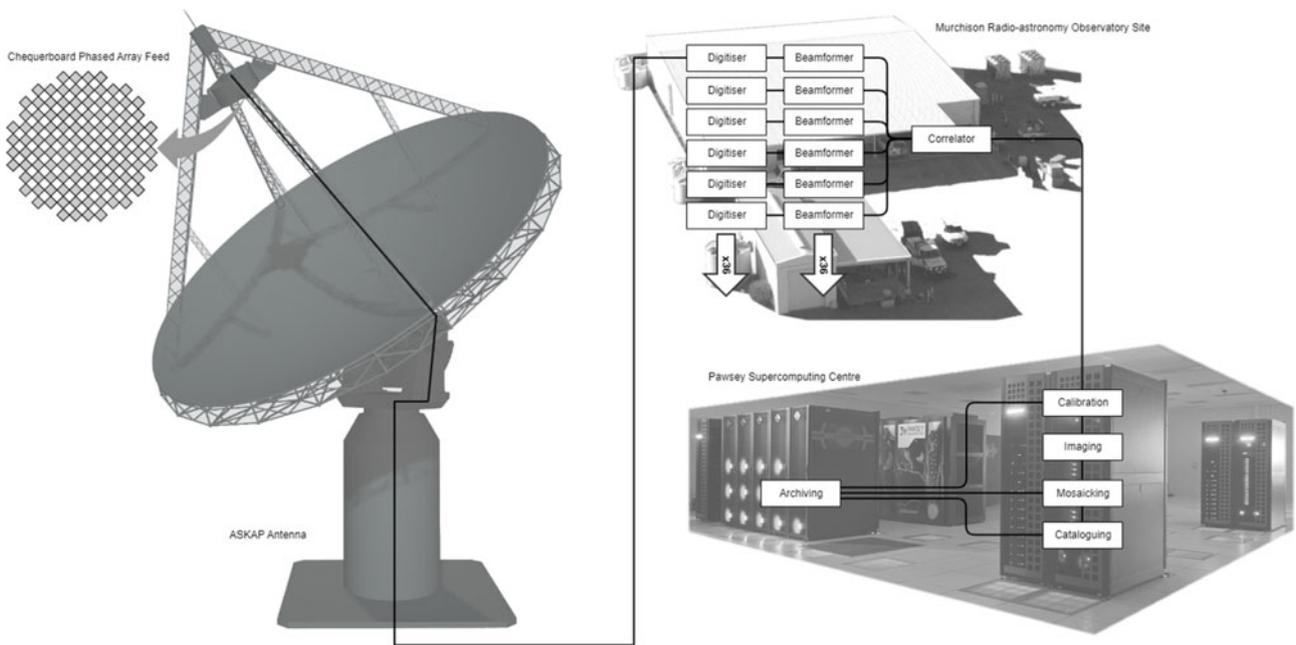


Figure 3. Overview of key ASKAP subsystems and data flow.

heat exchange units extract heat from the air. The cooling system vents this waste heat into a geothermal exchange system. When the air temperature is low, a water-to-air heat exchanger can be put in series, extending the life of the geothermal system. The entire cooling system was designed to minimise radio frequency emissions, as described in Abeywickrema et al. (2013).

3.4. Network infrastructure

The network between Perth and the MRO consists of approximately 900 km of optical fibre using carrier-grade transmission equipment. The network was designed and is managed jointly by CSIRO and AARNet (the Australian Academic and Research Network). Using dense wavelength division multiplexing, the system can support up to 80 channels on a single fibre pair. Over this distance, optical amplification is required every 80–100 km. Only two of these channels are used for production services, each lit by a 100 Gbit s⁻¹ transponder split out into 10 × 10 Gbit s⁻¹ services. The first runs direct to the Pawsey Supercomputing Centre in Perth with 4 × 10 Gbit s⁻¹ used for ASKAP data and one 10 Gbit s⁻¹ service to carry MWA data. The second 100 Gbit s⁻¹ link provides services between Perth, Geraldton, and the MRO such as CSIRO internal network connectivity, IP telephony, and monitor and control traffic.

There are additional transponders to support SKA development activities at Curtin University and for long-haul network performance testing undertaken by AARNet and CSIRO. Currently, there is 500 Gbit s⁻¹ of lit capacity between Perth and the MRO using 4 of the 80 available channels. There is also a parallel network consisting of two 1 Gbit s⁻¹ links directly patched to fibre pairs between Geraldton and the MRO. This provides network services at the various locations en route and acts as a backup network should the transmission system fail, however, it does not protect against a fibre cable cut.

The network at the MRO itself is a dual-core design. One core handles all the telescope data and connects the digitisers and

beamformers. The second core handles the monitoring and control of the telescope as well as the general purpose network at the MRO. The dual-core allows for separation of the different systems providing better control and resilience. The data network is large, the core router has over 2 400 network ports.

The second core makes extensive use of virtual local area networks (VLANs) to partition traffic. This avoids telescope monitor and control commands being mixed with unrelated network traffic. VLANs provide different network connections at relevant locations across the MRO.

The conditions at the MRO present some challenges to network equipment; the antenna pedestals are not cooled, so industrial switches are used, providing greater reliability and enabling connectivity for the varied range of equipment that has been installed at the observatory. The flexible network design allows for services to be provided to several clients located at the observatory, helping to broaden the research taking place at the site.

4. Antenna design

ASKAP’s antennas were designed, constructed, and installed by the 54th Research Institute of China Electronics Technology Group Corporation (CETC54) to CSIRO’s specification (CSIRO 2008). The antennas are 12 m diameter unshaped prime focus paraboloidal reflectors with a focal ratio f/D of 0.5. They have a solid surface, specified up to 10 GHz, and are mounted on a novel ‘sky-mount’ that can roll the entire reflector, quadrupod and feed structure about the optical axis (DeBoer et al. 2009). Slew rates of 1deg s⁻¹ in elevation and 3deg s⁻¹ in azimuth and roll mean that ASKAP can slew to and track an arbitrary position on the sky within one minute of request.

Continuous rise-to-set tracking of most sources is possible above the antenna elevation limit of 15°. From the observatory’s position at 26.7° S latitude, the telescope can observe sources between -90° and +48° declination, although with decreasing

time above the horizon for sources closer to the northern limit. The PAF's wide field of view extends this reach by another 2.5° .

4.1. Mechanical design and axes of motion

The antenna structure consists of a steel pedestal and support frame, topped with solid panels made of non-metallic honeycomb sandwiched between aluminium sheets. The feed is located at the prime focus.

Due to the small size of the reflectors and the comparatively large size of the PAF, special consideration was given to the strength and rigidity of the prime focus support structure. The robust feed legs, specified to support a 200 kg PAF receiver, block 4% of the total collecting area with the symmetric antenna design used. However, the symmetric design produces more uniform offset beams, in both total intensity and polarisation.

The use of a PAF with offset beams also makes it necessary to track parallactic angle on the sky during an observation. This is typically done with an equatorial mount, which creates extra mechanical complexity. For ASKAP, a unique hybrid approach was developed. CSIRO designed a simple azimuth–elevation mount that included a third axis of rotation for the reflector itself.^e

While it would be possible to track parallactic angle by continuously updating the electronic beamformer weights or rotating just the feed, the third axis exchanges a small amount of mechanical complexity for greatly reduced computational complexity at the time of calibration and imaging, since it maintains the angular relationship of the feed elements with respect to the support structure.

Unlike the azimuth axis, the antenna polarisation axis (also known as the roll axis) does not have overlapping range limits and can only be driven slightly less than $\pm 180^\circ$ from its neutral position. This means it is not possible to continuously track northern sources without unwrapping at transit or beginning the track at an offset roll angle.

The antennas were specified to operate up to 10 GHz, providing the flexibility to consider higher frequency receiver upgrades in future. A single ASKAP antenna was successfully used for several years with a single-pixel feed for VLBI observations at 8.4 GHz (Kadler *et al.* 2016). Preparation for a large-scale receiver system upgrade to higher frequencies should include direct measurements of atmospheric water vapour to assess the suitability of the site itself. Photogrammetry was used to measure both the location of the PAF at the focus and the reflector accuracy after installation. The rms deviation from the expected parabolic shape is typically less than 1 mm. This is consistent with measurements of the first antenna during test assembly at the factory (Feng, Li, & Li 2010).

4.2. Wind limits and storm monitoring

The ASKAP antennas were specified to operate under conditions similar to other radio telescopes. The antennas will stow if the peak wind speed exceeds 45 km h^{-1} during the last 30 min. We find that very little observing time is lost as a result.

Strong winds at the MRO are usually associated with oncoming storm fronts. These can cause the wind speed to rise extremely rapidly, sometimes faster than the time it takes the antennas to

^eThe third antenna axis was first suggested by Dr. Peter Dewdney and colleagues from the Dominion Radio Astrophysical Observatory in Canada.

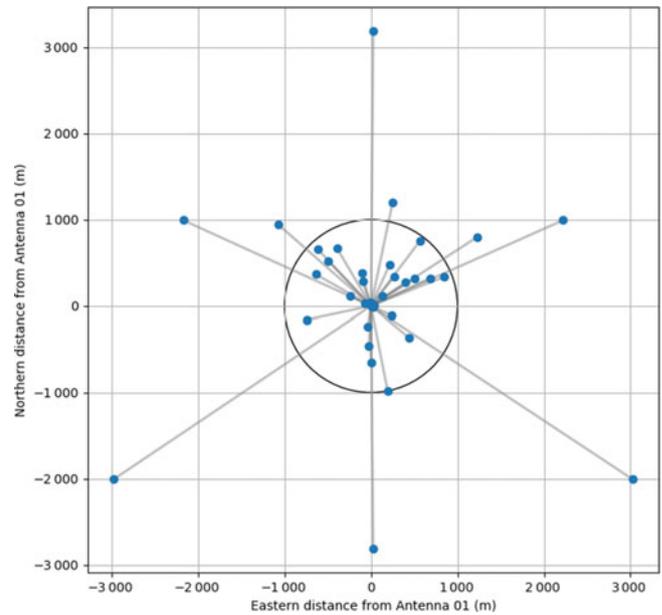


Figure 4. Location of each ASKAP antenna plotted relative to Antenna 1. A circle of radius 1 km is drawn for scale. The outer six antennas form a Reuleaux triangle and provide approximately $10''$ resolution. The dense cluster of antennas in the core provides excellent surface brightness sensitivity.

stow. On a few occasions, this has led to drive errors during the stow procedure as wind loading exceeded the safe operating range of the motors. To prevent circumstances such as this, we have implemented a storm stow system that uses satellite meteorological data and lightning detectors to sense approaching storm fronts and stow the antennas in advance of their arrival (Indermuehle *et al.* 2018a).

Not all nearby storms lead to high winds at the telescope, but associated lightning activity usually produces extensive radio-frequency interference. False positive stow triggers are, therefore, not a major concern as astronomy data would be impacted by a storm in any case.

4.3. Array configuration

Of the 36 ASKAP antennas, 27 were placed to provide a Gaussian distribution of spatial scales with a point spread function of $30''$ (Gupta *et al.* 2008). Three additional antennas were added to the core of the array to increase surface brightness sensitivity and another six antennas were added on longer baselines (up to 6 km) in a Reuleaux triangle (see Figure 4) to provide improved resolution (approximately $10''$) for compact sources. Antennas are plotted relative to ASKAP antenna 1, which is located at $26.6970007225^\circ\text{S}$, $116.6314242861^\circ\text{E}$ and an elevation of 361 m.

ASKAP has excellent instantaneous (u,v) coverage (Gupta *et al.* 2008), making snapshot surveys and equatorial imaging possible.

5. ASKAP phased array feed

5.1. Design background

The first production prototype PAFs for ASKAP used copper coaxial cables for signal transport, with digitisation inside the antenna pedestal (Hotan *et al.* 2014). Six of these units were constructed for assessment in the field. Aperture array measurements (Schinckel *et al.* 2011; Chippendale *et al.* 2014), tests at the focus of a 12 m

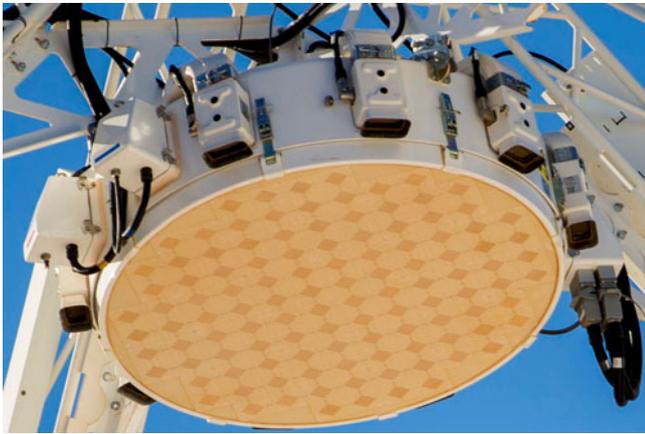


Figure 5. Photograph of a Mk II PAF installed on one of the ASKAP antennas at the MRO. The chequerboard surface is visible, along with the composite case and air vents for the cooling system. Power and optical cables attach via several bulkheads on the side of the case.

antenna located at the Parkes observatory (DeBoer et al. 2009; Chippendale et al. 2010; Sarkissian et al. 2017) and the BETA array at the MRO (Hotan et al. 2014) revealed several issues with the Mk I PAF. These included effective system temperatures in excess of 150 k (more than twice the design requirement) in the upper half of the frequency band (McConnell et al. 2016) and the need to fully remove the PAF from the focus of the antenna to perform internal maintenance.

In order to better assess the performance of prototype PAF designs, Chippendale et al. (2014, 2015, 2016a) developed an aperture array method of determining the noise temperature of a PAF beam using measurements of a microwave absorber, the radio sky, and broadband noise transmitted from a reference antenna. This was a valuable step in the testing process as it helped validate electromagnetic modelling of the PAF that in turn enabled improvements in sensitivity (Shaw, Hay, & Ranga 2012). Aperture array testing was also practical as it allowed rapid testing of PAF prototypes on the ground instead of at the focus of an antenna.

During these early tests, research into RF over fibre showed that the PAF signal could be transported to the central site for processing rather than digitising the signals at the antenna. A major redesign was initiated to incorporate experience from the Mk I tests and take advantage of new technological developments. The ASKAP Design Enhancements (ADE) project developed a Mk II system with lower overall cost, improved maintainability, and effective system temperature less than 80 k across most of the band. LNAs could be changed without removing the PAF from the dish and sensitive electronics were relocated from the antenna pedestal to the climate-controlled central building. All 36 ASKAP antennas are now fitted with Mk II PAFs and digital systems.

5.2. Planar chequerboard connected array

The ASKAP PAFs fill a region of the antenna's focal plane with small receptors that consist of flat, square conductive patches (see Figure 5) printed on a circuit board substrate to form a planar connected-array antenna (Hay et al. 2007; Hay & O'Sullivan 2008). The distance between each patch is 90 mm. A row of patches can be considered as a row of bow-tie antennas that are connected edge to edge, forming a linearly connected array. The feed point at the centre of each 'bow-tie' is differential and both orthogonal

linear polarisations are available due to the two-dimensional nature of the grid. Each element has a radiation pattern that would over-illuminate the dish surface, but the beamforming process combines these elements to create an efficient illumination pattern for a given direction on the sky.

The precise geometry of the chequerboard surface was determined through electromagnetic simulation (Hay, O'Sullivan, & Mitra 2011). Matching the simulation results to experimental data from the Mk I system was a key step to ensure that design modifications could be assessed prior to manufacturing. Co-design of the chequerboard array and its LNAs was critical for achieving high sensitivity. An initial impedance target for the LNA and matching network design was derived by optimising for maximum beamformed sensitivity over the field of view (Hay 2010). Slight modifications were made to the shape of the surface elements on the Mk II design to improve the impedance match with the LNAs in the upper observing band (1 400–1 800 MHz), and therefore reduce the system temperature.

The surface panels are bonded to a ground plane via several centimetres of non-conductive honeycomb backing, which provides structural rigidity as well as insulation. Twin-wire feed lines are used to connect the corners of each 'bow-tie' element to LNAs housed underneath the ground plane. A conformal coating is applied to the outward-facing surface for protection from dust and weather.

5.3. Composite enclosure

The density of electronics inside the PAF and the sometimes harsh environmental conditions at the focus of the antenna makes the design of the enclosure extremely challenging. As well as keeping the internal systems dry, free of dust, and thermally isolated from the external environment, the enclosure must provide more than 30 dB of extra RF shielding (Beresford & Bunton 2013). Experience with the BETA array showed that poorly shielded enclosures led to significant leakage of noise that could correlate, including between PAF elements and nearby antennas.

The Mk II enclosure is a moulded composite structure that is bonded to the metal ground plane and shielded with carbon fibre. Four access hatches at the top are fitted with RF gaskets to reduce radio frequency leakage. The access hatches allow replacement of faulty electronics modules, even when a PAF is installed at the focus of an antenna.

5.4. Dominoes and monitoring

Each pair of LNAs is housed in a self-contained RF shielded enclosure known as a domino. The domino design was developed for the Mk II PAF to improve shielding of the components with the highest gain, and also to improve modularity and maintainability. The dominoes are bolted directly to the ground plane to provide conductive heat transfer and can be removed individually for maintenance.

Each domino has three layers that contain differential LNAs, filters, and electro-optical converters stacked in order of distance from the ground plane. A photograph of the interior of each layer is shown in Figures 6, 7, and 8, and a transparent view of the assembled stack is shown in Figure 9. The signal path through the domino layers is shown in the upper left of Figure 10 in the boxes marked LNA, FILTER, and TRANSMITTER.

The top of each domino provides connections for two fibre-optic cables that carry the analogue RF output, and for a copper

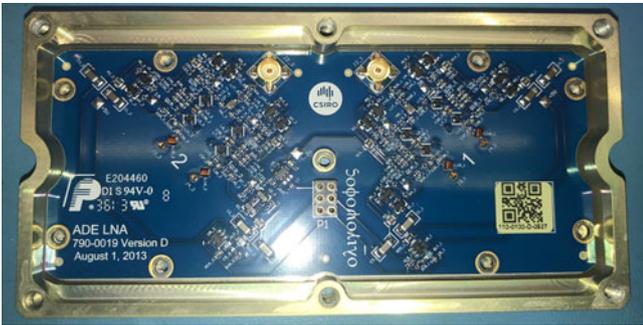


Figure 6. Photograph of the low-noise amplifier layer in the domino, referred to elsewhere as LNA.



Figure 7. Photograph of the filter layer in the domino, referred to elsewhere as FILTER.

ribbon cable that supplies DC power and supports monitoring and control within the PAF. Due to the proximity of the LNAs, front-end monitoring has a very high risk of generating RFI. ASKAP's PAF monitoring system was designed to be low impact, with no continuous clock signal and RFI filters on all connections to wires within the enclosure. The monitoring system is based on the Serial Peripheral Interface (SPI) bus and is shown in the lower right part of [Figure 10](#). Transmission of data packets may generate RFI, but most monitoring is disabled during an observation, with short updates between scheduling blocks to maintain a record of the system health. Control of each PAF is managed by a front-end controller (FEC) unit that converts instructions from an EPICS IOC into optical signals for the PAF and returns monitoring data to the software layer.

5.5. Ground plane and cooling

The Mk I PAFs employed water cooling inside their enclosures, using a heat exchange unit in the PAF and fans to circulate cool air throughout the electronics. Although this provided good cooling capacity, it required housing a water chiller at each antenna pedestal (with associated RFI issues) and pumping water up to the focus of the antenna. In addition, the possibility of condensation around the chilled water inlet or water leakage inside the PAF enclosure was a significant concern.

The Mk II PAFs use heat pipes embedded in the ground plane to transport heat to the edge of the enclosure. Further heat pipes connect to thermoelectric coolers (TECs, as shown in [Figure 10](#) and seen around the edge of the PAF in [Figure 5](#)) that transport the heat to the outside of the case where a ducted fan blows ambient air over heat sink fins. This system has proven effective at keeping the internal temperature as much as 15°C below ambient. It can be



Figure 8. Photograph of the optical transmitter layer in the domino, referred to elsewhere as TRANSMITTER.

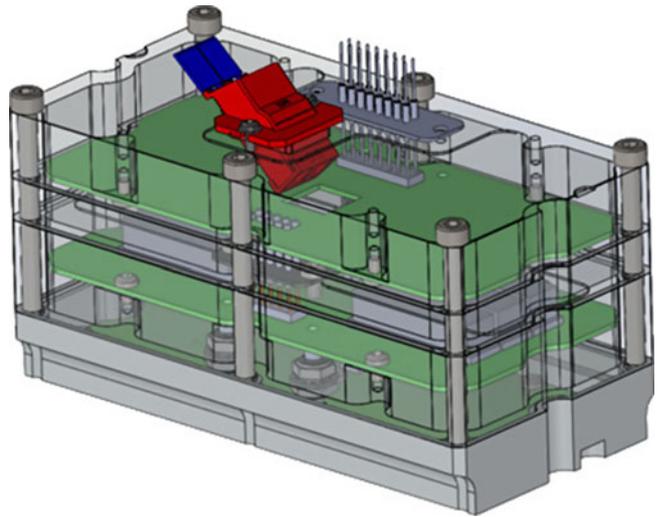


Figure 9. Diagram showing the assembled domino with transparent walls. The LNA card is closest to the bottom of the stack, followed by the filter card in the middle and the optical transmitter at the top.

regulated via the voltage applied to each TEC. A servo loop keeps the internal temperature at a steady 25°C set point (which can be changed in software if necessary).

The PAFs were designed to operate in ambient temperatures up to 45°C. This threshold is rarely exceeded except during the most extreme summer conditions. The cooling system is capable of maintaining safe internal operating temperatures in all conditions experienced to date. However, it is not always possible to keep the internal temperature fixed. The optimal internal temperature of 25°C is exceeded when ambient temperatures rise above 35°C, which is quite common during summer. The PAF is fitted with an automatic safety shutdown that is triggered when the internal air temperature exceeds 50°C.

The exact cooling capacity is slightly different for each unit. When the cooling capacity is exceeded, the internal temperature of the PAF begins to track the external ambient temperature, with a constant offset of approximately 10°C. This can cause amplitude calibration errors during astronomical data processing as the gain of the analogue electronics depends on their operating temperature. The impact of this on image quality has yet to be determined and it may be possible to correct using time-dependent self-calibration, a temperature-based lookup table, or the on-dish calibration (ODC) system described in the next section.

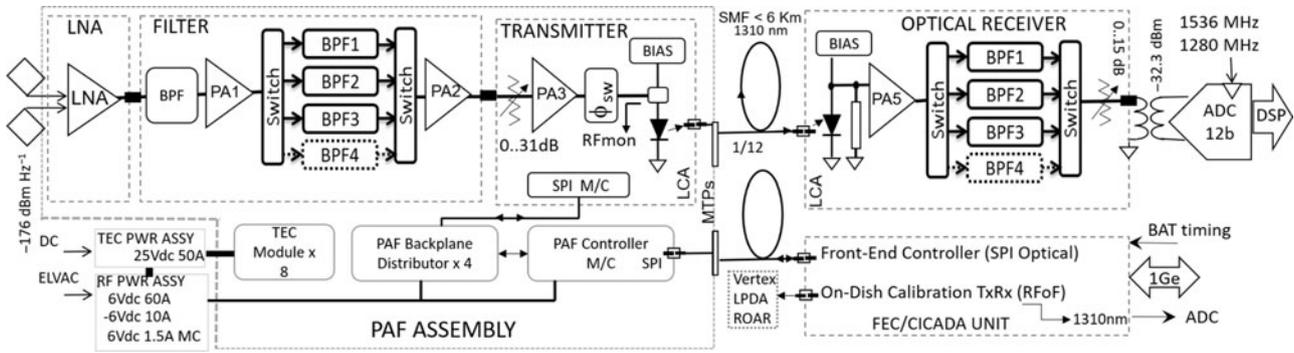


Figure 10. The analogue signal path for a single PAF element from the feed to the ADC in the digital receiver. After the chequerboard feed elements, the signal encounters a low-noise amplifier (LNA), followed by switchable bandpass filters (BPF) and power amplifiers (PA). The transmitter module provides an option to introduce phase switching (ϕ_{sw}), after which the signal is used to modulate a laser diode for transmission over single mode optical fibre (SMF) from the antennas to the control building. After analogue-to-digital conversion (ADC), the digital signal processing (DSP) stage begins. PAF monitoring and control (M/C) is distributed to all elements via a system based on the serial peripheral interface (SPI) standard. Cooling inside the PAF is provided by a series of 8 thermoelectric cooler (TEC) modules. Direct current (DC) and extremely low voltage alternating current (ELVAC) powers the various subsystems inside the PAF. The on-dish calibration system consists of a log-periodic dipole antenna (LPDA) mounted at the dish vertex and connected to an optical receive module (ROAR) fed by a noise source (CICADA) located in the front end control (FEC) module. This can be driven synchronously using binary atomic time (BAT) events loaded via a gigabit ethernet (1Ge) interface.

5.6. On-dish calibration

For a dish with a single feed, amplitude calibration is often obtained by placing a switched noise source in the feed that adds a known amount of power to both polarisations when activated. To determine the system amplitude gain, the power induced by the noise source is synchronously demodulated and measured with respect to the background.

Individually injecting a calibration signal into each PAF element is non-trivial due to the number of active elements. Instead, a common noise signal is radiated from the dish vertex to the PAF, illuminating all elements as uniformly as possible. The calibration signal is typically set to inject around 15 k equivalent noise temperature into the PAF, but this can be adjusted within the range 100 mK–200 K as required. We expect to reduce the calibration signal level as we gain experience in its use. The calibration noise is currently on continuously during all observations, but it is largely cancelled by beamforming with maximum signal-to-noise ratio beamformer weights resulting in less than 1% increase to beamformed system temperature (Chippendale, Button, & Lourenço 2018).

Using a radiated calibration signal has the disadvantage that it can leak into neighbouring antennas, but the signal power is such that the correlated contribution in a nearby antenna is less than that due to leakage of amplifier noise. Such common noise contributions could be further mitigated via Walsh modulation of the phase switches incorporated into each domino module, however this has not been found necessary.

Processing one or two calibration signals is a small engineering overhead compared to the 188 signals generated by each PAF. This allows the injected calibration signal to be digitised and correlated with the signals from all PAF elements, yielding an amplitude and phase calibration for each port. Corrections are then made to the beamformer weights to maintain stable beam patterns and sensitivity as described in Section 9.6.

The on-dish calibration system consists of several modules, as shown in Figure 11. There is one Cicada module per antenna located in the central building. During normal operation, it generates an independent noise signal that is directly modulated onto an

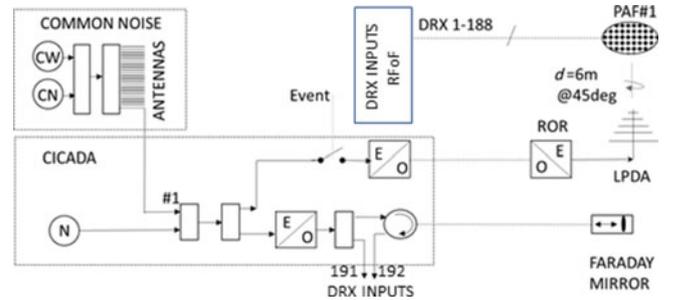


Figure 11. Diagram showing key components and signal paths for ASKAP’s on-dish calibration (ODC) system.

optical carrier and sent to the antenna on single-mode fibre. At the antenna, the optical signal is detected and the resulting electrical signal is sent to a small log-periodic dipole array antenna (LPDA, $G = 4$ dBi) located at the vertex of the 12 m dish, which points at the PAF and is oriented so that its linear polarisation is at 45 to that of the PAF elements. The illumination of the PAF is close to uniform but there is a path length difference of up to 21 mm from the central elements to the edge. To a first approximation the signal injected into all PAF elements is the same.

The Cicada also generates an optical copy of the noise signal. Half the power is sent directly to the corresponding antenna’s digital receiver and half is sent to the antenna on the same multicore fibre cable as the PAF signals. This copy is reflected by a Faraday mirror in the antenna pedestal and returns to the Cicada, then continues into the digital receiver. In the digital receiver, both optical signals are processed in the same way as the PAF signals, becoming inputs to the array covariance matrix and a higher sensitivity calibration correlator that forms all correlation products between a nominated reference signal and the PAF signals.

The direct signal from the Cicada gives the best estimates of gain and phase variation. The signal returned from the Faraday mirror allows the optical path to the LPDA to be calibrated via a round-trip phase measurement. This allows the phase variation of the full analogue optical path to be characterised. The system also

provides a way to inject a common calibration signal into the PAFs on all dishes, though this is not routinely used.

Although the main purpose of the ODC system is to measure complex gain and provide updates to the beamformer weights, it can also be used to improve system diagnostics in areas of RF continuity, RF distortion assessment and filter-bank frequency integrity. The main contribution of the ODC at present is to stop beams from being completely destroyed by phase slopes corresponding to unpredictable delay changes that occur with each full reset of the analogue-to-digital converters (ADCs). With further research it is hoped to use the ODC to temporally stabilise PAF beam (voltage) patterns to better than 1% at the half power points (Hayman *et al.* 2010).

6. Analogue signal processing

6.1. Design background

In the BETA version of the ASKAP PAF (Schinckel *et al.* 2011; Hotan *et al.* 2014), all analogue signal processing and digitisation occurred in the antenna pedestal. At the time of the BETA design, photonics for radio frequency over fibre (RFOF) signal transport appeared expensive and sufficiently unproven to deploy on the scale of a 6768 element array for radio astronomy.

The BETA system comprised two shielded cabinets packed full of equipment and occupying most of the space in each antenna pedestal, along with a chiller plant located on the concrete antenna foundation pad to keep both the cabinets and PAF cooled. The chillers required continuous maintenance and the pedestal electronics were not easily accessible due to the confined space. Containment of radiated emissions was less than satisfactory due to the single layer of shielding and the large number of filtered conductive connections to the outside world. Even so, BETA served well to verify basic PAF performance and signal processing algorithms and was a pivotal step to the more advanced Mk II design.

For the Mk II version of ASKAP described here, the price of high linearity uncooled distributed feedback (DFB) laser diodes had fallen. Investigations of performance of low-cost RFOF analogue signal transport using these DFB components were undertaken with fibre spans up to 6km. It was found that, with suitable precautions, a directly modulated DFB laser solution could meet astronomy requirements (Beresford, Cheng, & Roberts 2017a). The cost of the optical transmitter was reduced to one hundred dollars rather than several hundred dollars for cooled laser devices or thousands of dollars for high quality externally modulated lasers traditionally used in broadband applications.

This allowed all analogue electronics, except the PAF itself, to be moved to the central building, which eliminated self RFI problems, lowered cost, and improved accessibility and maintainability of the analogue signal processing hardware.

6.2. LNA

Each domino houses the low-noise amplifiers (LNAs) for a pair of PAF elements (Figure 6) and there are 94 dominoes in each PAF. ASKAP's LNA design forms an active balun by connecting the twin-wire feed line of one connected-array element on the PAF surface directly into two single-ended LNAs with a passive output-side balun (Shaw *et al.* 2012). A passive input-side balun would have introduced too much loss. The differential input complicates the measurement of LNA noise temperature, so an innovative method was developed by Shaw *et al.* (2012). The measured LNA

Table 3. ASKAP receiver bands

Band	Sampled RF band (MHz)	Sample clock frequency (MHz)	1 MHz channels
1	700-1200	1 280	640
2	840-1440	1 536	768
3	1400-1800	1 280	640
4	600-700	1 536	768

minimum noise temperature T_{\min} of the current design is 30 k in the middle of the frequency range (Shaw & Hay 2015).

6.3. Filter

After the LNA, the signal encounters selectable anti-aliasing filters (Beresford *et al.* 2017b) as shown in Figure 7. These filters define four observing bands summarised in Table 3 and Figure 14. The fourth band was included for assessment of the low-frequency performance of the PAF and has not been used for astronomy.

6.4. Transmitter

After filtering, the signal is converted from electrical to optical for transmission back to the central building. The optical RFOF transmitter (Beresford *et al.* 2017c,b; Hampson *et al.* 2012) is shown in Figure 8. The RF signal is directly added to the laser diode DC bias current and modulates the optical power. An average power control loop ensures constant optical level, compensating for component ageing. The low loss in the single-mode fibre allows the signal to be transported several kilometres. Each antenna connects to the central processing building via underground fibre-optic cables containing 216 cores per antenna, 188 of which carry RFOF PAF signals. Three additional fibres are used for the RFOF ODC system described in Section 5.6.

6.5. Optical link design considerations

All observing bands are suboctave as they are in the second or third Nyquist band of the analogue-to-digital converter (ADC). This mitigates the second-order harmonic distortion of the RFOF link. Radio astronomy signal paths are designed to respond linearly to their input and this was an important consideration for the optical link as well. For systems operating close to linearity, small departures are often characterised using a measure known as the input or output intercept point (IIP or OIP) for various polynomial orders. The ASKAP RFOF link has an OIP1 of -10 dBm, an OIP2 of 8.3 dBm, and an OIP3 of 6 dBm (Beresford *et al.* 2017a) as shown in Figure 12. This marginally degrades as the fibre span increases (Beresford *et al.* 2017c). The nominal RFI-free RFOF link output signal level is -47 dBm (0 dBc), relative intensity noise (RIN) and shot noise define the output noise floor at -72 dBm (-25 dBc).

For narrow-band RFI equal in power to the system noise at the RFOF output (0 dBc), the second-order distortion is -53 dBc and the third is -102 dBc. The strongest narrow-band RFI is due to aviation Automatic Dependent Surveillance-Broadcast (ADS-B) transmitters, located on aircraft flying over the observatory at high altitude. This can be as high as 29 dBc for short periods of time (minutes) if the aircraft passes overhead.

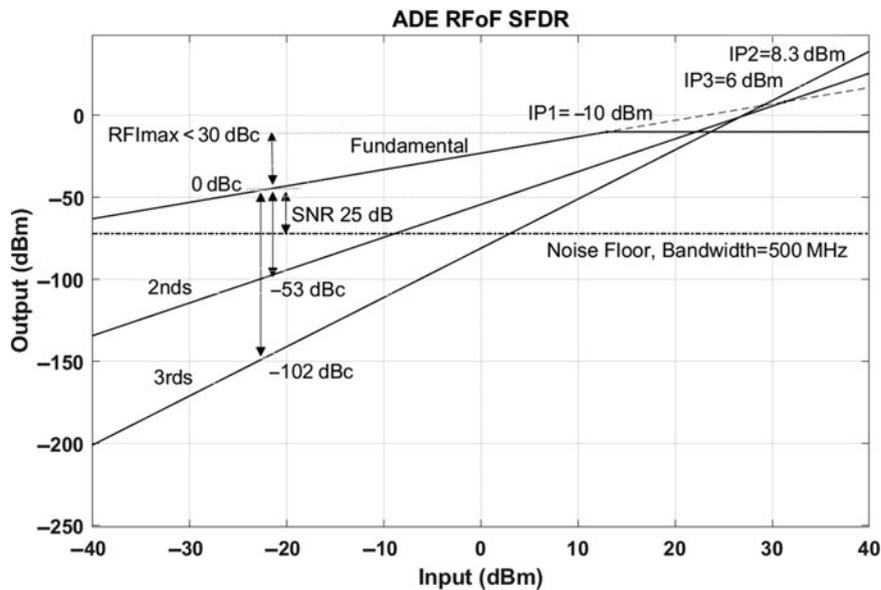


Figure 12. Distortion performance, measured as the spurious-free dynamic range (SFDR), of the RF over fibre link.

Placement of half the anti-aliasing filter after the RFoF link eliminates the severe in-band second-order response from the RFoF link as well as out of band relative intensity noise and shot noise (Beresford et al. 2017c). This allows the system to handle an extra 17 dB of RFI (third-order intermodulation at -53 dBc, see Figure 12). The other half of the anti-aliasing filter preceding the RFoF link removes out of band RFI that could corrupt the link. Adjustment of the attenuator in the domino transmitter (see Figure 10) can reduce the signal level on the RFoF link to increase the compression headroom, but this comes at the cost of increased additive noise. With nominal settings, the RFoF link adds 1.5 k to the receiver noise temperature. Decreasing the power on the RFoF link by 3 dB gives 3.7 dB extra headroom for RFI but adds 1.9 k to the receiver noise temperature for the most distant antennas.

6.6. Optical receiver

The conversion from optical back to electrical signals occurs in the digital receiver (see Section 7.1). Further amplification increases the electrical signals to the level needed to drive the analogue-to-digital converter (ADC) in the digital receiver (-35 dBm). The ADC has a specified spurious-free dynamic range (SFDR) of approximately 61 dB. The RF signal chain has an SFDR of approximately 48 dB. Marginal degradation of SFDR occurs for longer fibre spans due to the optical loss of 0.35 dB km^{-1} . The longest fibre span in the array is 6km.

7. Digital signal processing

The next stage in the ASKAP signal chain consists of a distributed, custom digital signal processing (DSP) system built on field-programmable gate array (FPGA) technology (Hampson et al. 2012). This system is responsible for the following key processing steps:

1. sampling the PAF analogue RF signals;
2. channelising the band with 1 MHz coarse resolution;
3. selecting channels for further processing;
4. beamforming on the selected coarse channels;

5. applying time-varying coarse delays to align wavefronts with an error of at most half a sample at 1 MHz;
6. further channelising the beamformed output to the final frequency resolution of the correlator;
7. applying a time-varying phase slope across 1 MHz channels to provide fine delay control (fraction of coarse delay step); and
8. for each beam, cross-correlating the beam voltages.

Processing is implemented in three stages, first, the digital receiver, then the beamformer and finally the correlator. The result is the cross-correlation of beam voltages for each fine channel across all antenna pairs and polarisation products (the raw visibilities) for the specified phase centre. Each beam is correlated with itself (autocorrelation) and with each corresponding beam that points in the same direction from every other ASKAP antenna. Correlations between beams pointing in different directions are not calculated.

7.1. Digital receiver

ASKAP’s digital receiver (Brown et al. 2014) is implemented with 12 Dragonfly modules (see Figure 13) for each antenna, each processing 16 optical signals. The 192 signals include 188 from each PAF, 2 calibration signals, and 2 spares for future applications such as RFI mitigation.

The ADC devices used for ASKAP are National Semiconductor ADC12D1600 parts that have 12-bit resolution with an effective number of bits (ENOB) of 9 and a full-power analogue bandwidth of 2.8 GHz. The sample clock is generated in the Dragonfly synthesiser module as described in Section 7.6.

The ASKAP digital receiver has a bandpass sampling architecture. It directly samples the PAF signals at RF (Figure 10) without the need for analogue frequency conversion. The full 700–1 800 MHz frequency range of the PAF is covered by three selectable and overlapping sampling bands (see Table 3 in Section 6).

Each sampling band is designed to match a Nyquist zone of the ADC. Two sampling frequencies, 1 280 and 1 536 MHz, are used to sample at RF with an instantaneous bandwidth of 640 MHz or 768 MHz. Figure 14 illustrates how the sampling

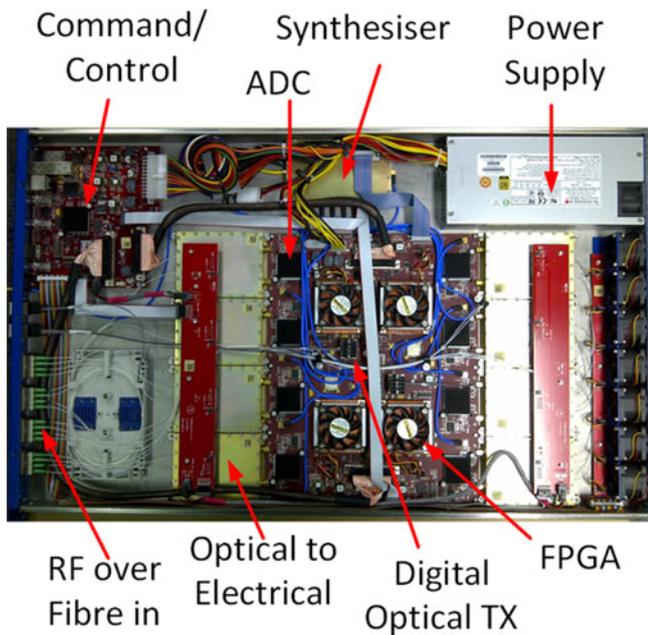


Figure 13. The Dragonfly digital receiver module, designed by CSIRO specifically for ASKAP. There are 12 per antenna.

bands correspond to sampling frequency, how they overlap, and how much usable RF band is left after filtering. The lower and upper observing bands are located in the second and third Nyquist zones, respectively, of the ADC running at 1 280 MHz, while band 2 is located in the second Nyquist zone of the ADC running at 1 536 MHz, which overlaps the adjacent bands at the other sampling frequency. The sampling frequency and corresponding FPGA firmware are configured at the start of an observation.

An overview of the signal path following the ADC, for a single PAF element, is shown in Figure 15. Also shown, in unshaded boxes, are a number of other ancillary signal statistics and monitoring modules. In operation, the ADC histogram is particularly useful as it shows whether the signal is Gaussian and helps to set the signal level into the ADC.

The sampled PAF element voltages feed directly into a polyphase filter bank (PFB) with an oversampling ratio of 32/27 (Tuthill *et al.* 2012). The PFB forms channels with a sample rate of 1.185 and 1 MHz spacing. As the channels have a noise bandwidth of 1 MHz and a channel spacing of 1 MHz they will be referred to as 1 MHz channels. The oversampling provides substantial gains in sub-band fidelity and downstream processing capabilities with only a modest increase in system complexity (Bunton 2003b; Tuthill *et al.* 2012). The channel-dependent frequency rotation within each output channel that is introduced by oversampling is compensated for within the PFB implementation so that the output 1 MHz channels are all correctly DC-centred before being forwarded on to the beamformer subsystem (Tuthill *et al.* 2015).

The PFB produces either 768 or 640 1 MHz channels across the selected observing band (Figure 14), depending on the sampling frequency. In both cases, the channel sample rate is the same and the channel centre frequency is an integer multiple of 1 MHz. The output of the PFB for each channel is the complex-valued analytic envelope of the full-band signal component within that sub-band, i.e. the single-sided spectrum in that sub-band, translated to DC (Rice 1982). The equivalent operations to generate

the output of the PFB for each channel are a multiplication of the input by $e^{2\pi if t}$, where f is the channel centre frequency, followed by low-pass filtering and decimation to the output sample rate.

Of the total number of output channels from the PFB, only 384 channels are selected for transport to the beamformer subsystem for further processing. These are transported as 8 data streams each with 48 channels. However, cost constraints and computing limitations mean that only seven of these data streams are beamformed and only six are correlated. This limits the total bandwidth available for imaging (see Section 11) to 288 MHz.

The channel selection process is very flexible and permits contiguous or non-contiguous groups of channels to be selected and also multiple copies of the same channels to be selected and transported to different downstream end points. The complex task of managing channel selection is done by the Telescope Operating System (TOS) software. At present, only contiguous bands are implemented to simplify visibility data storage and imaging. In future, we plan to support two or more split bands for purposes such as avoiding satellite interference and observing widely separated spectral lines.

Bit growth is permitted through the PFB to ensure that the DSP system meets dynamic range and noise floor requirements, so the output number representation for the coarse 1 MHz channels is 16-bit real and 16-bit imaginary two's complement (signed) integers. The resulting aggregate data rate at the output of the ASKAP digital receiver and PFB is approximately 100Tbit s⁻¹. This is efficiently transported to the beamformer using a custom point-to-point packetised streaming protocol that has very low overheads. The physical transport medium consists of 10 368 multimode optical fibres in 864 12-fibre ribbons. All fibres in a ribbon route to different locations and the fibre reordering is achieved by using custom passive optical cross-connects.

7.2. Beamformer

Output from the digital receiver for a single antenna is processed by seven Redback modules (Hampson *et al.* 2014), as shown in Figure 16. Each receives 48 MHz of bandwidth from all PAF elements for a single antenna and can generate 36 independent dual-polarisation beams. With seven modules, 336 MHz out of 384 MHz available from the digital receiver is processed. Wiring for the eighth module is in place, but cost constraints meant it was not installed. Each Redback has six processing FPGAs, which are fully interconnected electrically. The interconnect is used to distribute data so that each FPGA has 8 MHz of data from all 192 digitised ports (188 PAF elements + 2 calibration signals + 2 spare ports).

The signal path through the beamformer is shown in Figure 17. The beamformer engine produces up to 72 independent beams on the sky, each a linear combination of M digitised port voltages $\mathbf{x} = [x_1, x_2, \dots, x_M]^T$. The beamforming operation is defined by

$$y(n) = \mathbf{w}^H \mathbf{x}(n), \quad (2)$$

where \mathbf{w}^H is the conjugate transpose of the beamformer weights and $\mathbf{x}(n)$ is the vector of port voltages for a single frequency channel at time sample n . The beamformer weights are complex valued and may be varied independently for each 1 MHz channel of each beam. The weights are represented as 14-bit real and 14-bit imaginary two's complement (signed) integers.

The number of ports M that are weighted into each beam can be as many as all 192 ports from the digital receiver for 1

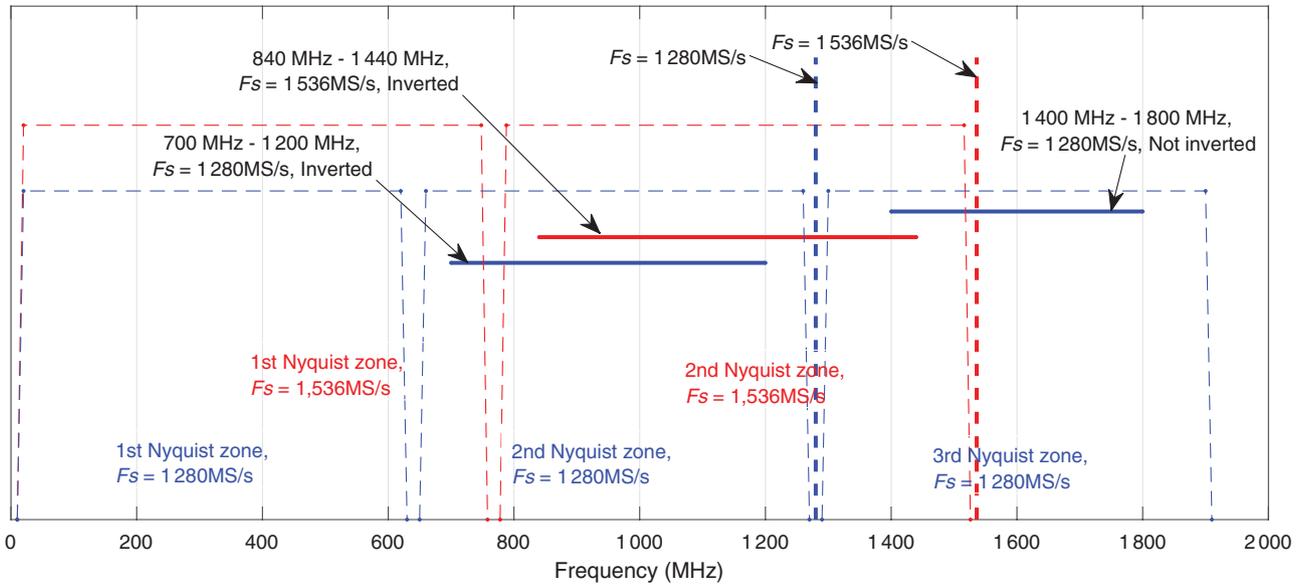


Figure 14. ASKAP digital system sampling bands. For convenience, we refer to these as bands 1, 2, and 3 in left-to-right order elsewhere in this document. The labels indicate whether the frequency channel order is inverted or not, with respect to the natural ascending order.

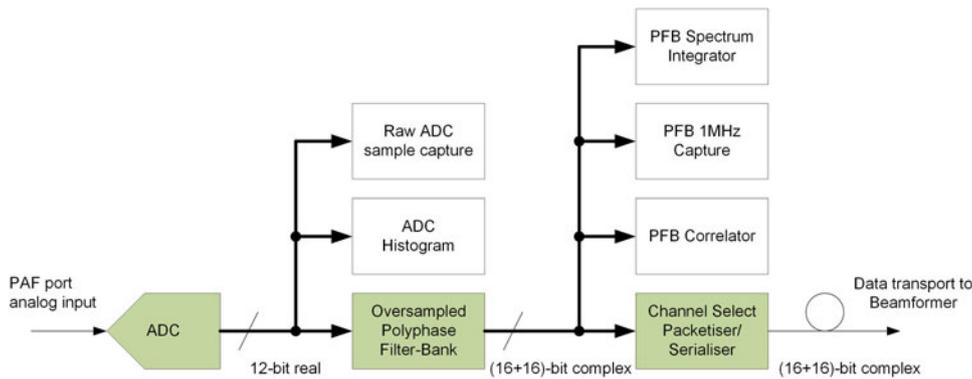


Figure 15. The signal path for a single PAF element through the ASKAP digital receiver ADC and PFB.

Table 4. Number of ports per polarisation per beam for various numbers of dual-polarisation beams

Dual-polarisation beams	Ports per beam M
1–9	192
10–18	126
19–27	82
28–36	60

ASKAP antenna, but it reduces to 60 ports as the number of beams increases to the full complement of 72 single polarisation beams. The reduction of input ports with beams is summarised in Table 4 and is required due to hardware resource limitations in the beamformer. The selection of M ports from 192 can be specified arbitrarily for each 1 MHz channel of each beam.

The system is typically configured to produce 36 dual-polarised beams. The beamformed voltages for each 1 MHz channel pass into a second polyphase filter bank, which performs a final fine frequency channelisation. For normal operation the frequency resolution is 18.5 kHz.

A key advantage of beamforming on 1 MHz channels is that the beamformer can be most efficiently implemented mid-stage as 336 separate narrow-band beamformers (Van Veen & Buckley 1988). This allows beamforming with complex weights as opposed to frequency dependent weighting with FIR filters, which would be needed if the beamforming occurred on the ADC voltages. Furthermore, since the fine filter bank (FFB) is operating on beam voltages rather than raw PAF element voltages, there is a reduction in the FFB data throughput by a factor of approximately three (188 PAF elements/72 beams). This reduction also means that although three 12-fibre ribbons are needed to input data to a Redback module only one is needed at the output. That is, two fibres per FPGA each carrying 4 MHz of beamformed data.

A key data product for many beamforming algorithms is the array covariance matrix (ACM), which is defined as

$$\mathbf{R} = \langle (\mathbf{x} - \mu_{\mathbf{x}}) (\mathbf{x} - \mu_{\mathbf{x}})^H \rangle, \tag{3}$$

where $\mathbf{x} = [x_1, x_2, \dots, x_K]^T$ is a vector of K array element voltages assumed to be stationary discrete-time stochastic processes, $\mu_{\mathbf{x}}$ is the element-wise mean value of \mathbf{x} , and $\langle \cdot \rangle$ denotes expectation. For

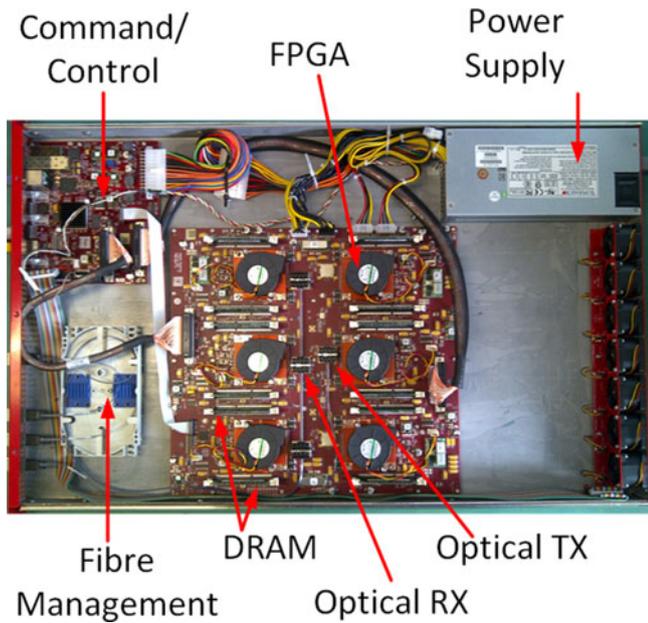


Figure 16. Redback module used for beamformer and correlator. Field-programmable gate arrays (FPGAs) used for digital signal processing are hidden beneath cooling fans. Banks of dynamic random access memory (DRAM) are visible beside each FPGA.

the ASKAP signal model the x_i are the frequency-channelised PAF element voltages, which are complex-valued random time series whose statistics are assumed to be stationary with zero mean over the beam calibration interval. Under these conditions, the elements of \mathbf{R} in (3) are adequately approximated for each 1 MHz channel by the time average

$$\hat{r}_{ij} = \frac{1}{N} \sum_{n=1}^N x_i(n) x_j^*(n), \quad (4)$$

where $x_i(n)$ is the n th voltage sample of the i th input port of $K = 192$ and x^* is the complex conjugate of x . With this definition, the ACM estimate $\hat{\mathbf{R}}$ is a 192×192 matrix where each element \hat{r}_{ij} is the correlation between x_i and x_j . This is computed for all 1 MHz frequency channels.

The full computation of the ACM has a higher compute load than the final correlator for each antenna. This load is reduced by using every fourth time sample and computing ACMs for 42×1 MHz frequency channels concurrently for each PAF. The ACM estimation process must be run eight times to cover the entire 8×42 MHz = 336 MHz of processed bandwidth. Raw ACM data products are transported to conventional compute hardware where they are used to calculate the beamformer weights offline as described in Section 9. The beamformer weights are then uploaded back to the beamformer block, which is shown below the ACM block in Figure 17.

Depicted above the ACM in Figure 17 is the calibration correlator. This calculates the correlation between a reference calibration signal and all PAF element outputs. This is mathematically equivalent to computing one row of the ACM, but in the calibration correlator, it is done for a specific reference input without time or frequency decimation. This achieves the best possible SNR for tracking PAF element gains. The calibration system is described in more detail in Section 5.6.

Table 5. ASKAP frequency zoom modes. In all cases, the number of frequency channels correlated is 15 552 for the current hardware deployment

Mode number	FFB output channels	Channel width (kHz)	Correlated Bandwidth (MHz)
1	64	18.5185	288
2	128	9.2593	144
3	256	4.6296	72
4	512	2.3148	36
5	1024	1.1574	18
6	2048	0.5787	9

The beam voltages output from the beamformer block are fed into two separate processing modules:

1. a deep ‘corner-turn’ buffer; and
2. a transient ring buffer.

Both of these modules share a large double data rate (DDR) memory resource for temporarily buffering large amounts of beam voltage data. The transient module implements a ring buffer that can be frozen by an external trigger and is described in Section 7.3.

The primary signal flow for ASKAP, however, is the corner turn buffer that implements a streaming data transpose, supplying large blocks of sample data for the same beam and frequency channel to the FFB, which is operated in block-mode for efficiency. The FFB is a second-stage frequency channeliser implemented as a critically sampled polyphase filter bank on the 1 MHz channels which have a sample rate of 1.185 MHz. This module can be configured at the start of an observation to provide a standard observing frequency resolution, or one of five ‘zoom’ resolutions, according to Table 5.

Note that since each of the 42 FPGAs in a single PAF beamformer subsystem processes the 8×1 MHz channels that are routed to it from the digital receiver, different zoom modes can be applied to different or even the same parts of the total processed bandwidth by loading the appropriate zoom mode firmware on each FPGA and routing the corresponding 1 MHz channels to it. Control software to implement this is not available currently. Each FPGA always generates 423 frequency channels independent of zoom mode. Thus, the total bandwidth decreases as frequency resolution improves as shown in Table 5. Processing of a full 1 MHz channel, for zoom modes, requires multiple correlator FPGAs to process the same 1 MHz channel.

7.3. Fast transient detection

As shown in Figure 17, a secondary beamformer output is sent to the fast-transient detector. This operates with the 1 MHz channels and has an adjustable time resolution that is typically set to 864 s or 1.7 ms. Power averages are sent over the MRO network to a dedicated computer that searches for dispersed pulses in real time (Bannister *et al.* 2019a). In the search algorithm, these data are currently added incoherently across all antennas.

If a pulse is detected, a trigger is sent to the transient ring buffer in the beamformer, telling it to freeze the contents of the ring buffer (which stores an adjustable amount of raw voltage data from all beams, corresponding to the last 0.9–14.2 s depending on data precision selected) and download to disk. The voltage data are then correlated offline using the DiFX software correlator (Deller

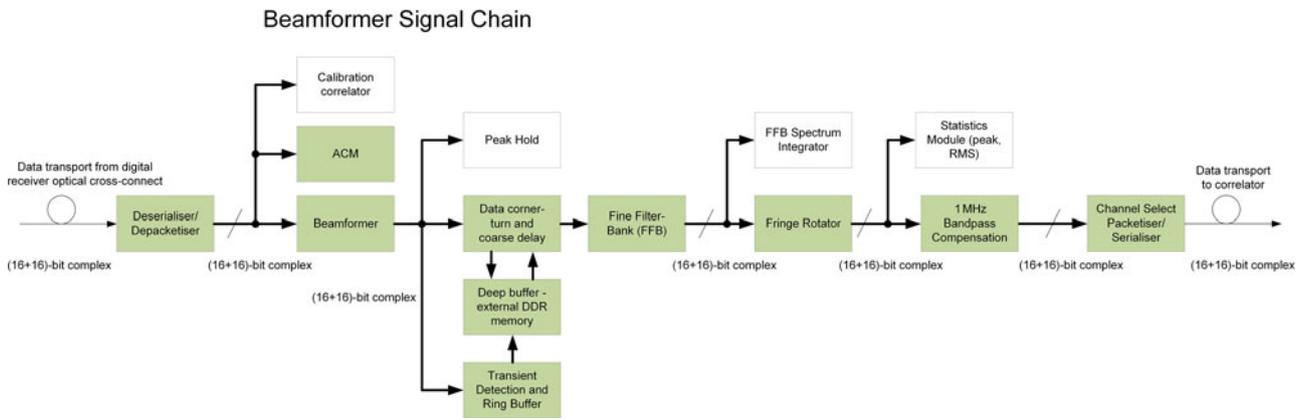


Figure 17. Major components of the signal path through the ASKAP beamformer.

et al. 2011) and imaged in order to localise the position of the event with sub-arcsecond resolution and provide high sensitivity data on the pulse (Bannister et al. 2019b).

7.4. Delay tracking and fringe stopping

Distant astronomical sources remain effectively stationary in the sky, while the Earth (and therefore the telescope) rotates once per day. This rotation means that the projected baseline vector between any pair of antennas in the array is constantly changing. Interference fringes, therefore, move during an observation, which would decorrelate the signal from the interferometer if not compensated for. The process of adjusting the delay and phase of each antenna to keep interference fringes stationary is known as delay tracking and fringe stopping.

Since ASKAP is a multi-beam instrument, we independently stop fringes in the nominal pointing direction of each beam. This is straightforward when using the antenna polarisation axis to keep beams pointing at a fixed position on the sky. We also offer the option to stop fringes in the boresight direction of the antenna for all beams, which can be useful for experiments such as holography where polarisation axis tracking is not used.

The ASKAP digital hardware implements:

1. a pure time delay to a precision of 0.84 s applied across all beams and 1 MHz coarse frequency channels; and
2. a time-dependent and frequency-dependent phase correction applied to each fine frequency channel.

A pure ‘coarse’ time delay (with precision of 0.84 s) is implemented between the beamformer and the FFB by offsetting the read pointer into the corner turn sample buffer for the FFB frame input. The coarse delay is the same for all beams and 1 MHz coarse frequency channels for a single PAF system, but is typically different between antennas in the array to allow for a coarse compensation of the difference in arrival time (up to 20 s) of a wavefront at different antennas. During an observation, the required coarse delay is computed repeatedly and set in the hardware at times synchronised with signals from the timing and synchronisation subsystem (see Section 7.6). Changes occur at the start of each data frame and do not introduce any discontinuities.

The second ‘fine’ stage of fringe stopping occurs in the *fringe rotator* (Figure 17), which uses complex multipliers to rotate the

phase of each fine frequency channel time series as they come out of the FFB. The rotation phase is calculated from coefficients loaded via the control system software and permits a unique phase correction to be applied at each of the four levels:

1. coarse 1 MHz channels;
2. dual-polarised beams;
3. fine frequency channels; and
4. particular time samples.

The fringe rotator module exploits the fact that the rotation phase, for any beam, is linearly dependent on frequency. Furthermore, the module parameters are updated often enough that the time variation of phase is adequately approximated with a piecewise linear function. Under these conditions, the rotator can calculate the required phase, for a beam, from a unique starting phase for each coarse frequency channel and two increment parameters. These parameters are the change in phase with time increment and change in phase with frequency increment.

For each coarse channel and beam combination, the phase correction applied for an individual time sample Δt units after the update epoch and Δf fine frequency channels from a reference channel, is given by

$$\phi(\Delta t, \Delta f) = \phi_0 + \phi_T \Delta t + \phi_F \Delta f, \quad (5)$$

where ϕ_0 is an initial phase offset value and ϕ_T and ϕ_F are the time slope and frequency slope values, respectively.

The array-wide synchronous event that updates the coarse time delay has an appropriate processing delay applied before being used to update new fringe rotator parameters ϕ_0 , ϕ_T and ϕ_F so that the coarse delay and fringe parameters are consistent for each update interval.

The smallest incremental change that can be made to the fringe stopping parameters at the fine control stage (corresponding to 18.5 kHz channel bandwidth) is 0.206 ns in delay, 0.044deg in phase, and 0.0248deg s⁻¹ in phase rate. The fine-level delay slope across the band is implemented as a combination of delay steps and matching phase offsets with the increments given above.

At the output of the FFB, the beamformer subsystem produces $2^{M-1} \times 64$ fine frequency channels for each 1 MHz channel of each PAF beam, where M is the zoom mode number in Table 5 and 336×1 MHz channels are processed in the beamformer. As a result of the oversampling ratio of 32/27 in the coarse filter bank,

each 1 MHz of processed bandwidth is contained within the centre $2^{M-1} \times 64 \times 27/32 = 2^{M-1} \times 54$ fine channels. The redundant $2^{M-1} \times (64 - 54) = 10 \times 2^{M-1}$ fine channels at the band edges of each 1 MHz sub-band are discarded before further processing. The number of channels processed is independent of zoom mode.

7.5. Correlator

The correlator subsystem is the final stage in the FPGA-based signal processing chain at the MRO central site. It takes in fine frequency channelised beam voltage time-domain data, streaming from the 36 antenna beamformer subsystems, and generates the raw visibilities (cross-correlations) for the final science data processing stage. The data are carried from the beamformer to the correlator via multimode optical fibre and are transported using a custom, unidirectional streaming protocol for maximum link efficiency.

The outputs from the beamformer Redback modules connect to the correlator subsystem via several 12-fibre ribbon cables. A single correlator subsystem (also known as a block) processes 48 MHz of bandwidth and currently six of a possible eight blocks are installed. Each correlator block consists of 12 Redback modules (programmed for correlator functionality) and passive optical cross-distribution networks. The numbers in this section represent the installed capacity.

The accumulation operation in computing the visibilities results in a large reduction in the data rate such that the correlator output can be sent as Ethernet data over conventional network infrastructure to the remote high-performance computing facility in Perth, which is approximately 600 km from the telescope site.

The aggregate data arriving at the correlator input consists of 15 552 fine channels for 72 beams from each of the 36 antennas. To provide the capacity for polarimetry, these are configured as 36 dual-polarisation beams (with a linear basis). For each fine channel of each dual-polarisation beam, the correlator computes the correlations between all 36 antennas for both polarisations, resulting in $(2 \times 36 + 1) \times (2 \times 36)/2 = 2,628$ products. This is full Stokes cross-correlation for each of 630 baselines and full Stokes auto-correlation for each of 36 antennas. The total correlator output in one integration period, for all 36 dual-polarisation beams and all fine channels, is $2\,628 \times 36 \times 15\,552 = 147\,134\,3616$ complex-valued visibilities in the current 6-block configuration with 15 552 fine channels.

The complex-valued visibilities are streamed from the correlator as pairs (real/imaginary) of 32-bit floating-point numbers. The data streaming cycle from the correlator is locked to the integrate-and-dump cycle which, in turn, is controlled through the timing event subsystem (see Section 7.6). For the minimum allowed integration interval of 5 s, the total output data rate from the correlator is $147\,134\,3616 \times 32 \times 2 / (8\text{bit}/\text{Byte} \times 5\text{s}) \approx 2.4\text{GB s}^{-1}$. See Section 8 for more information on system-level constraints.

Data from the correlator are transmitted as user datagram protocol (UDP) packets, each containing up to 657 complex floating-point numbers corresponding to a single spectral channel, beam and polarisation. Partitioning the data in this way makes it easy to scale up the size of the telescope, but means the data must be reordered before further processing. Each UDP packet contains a total of 5 256 bytes of raw data, to which an additional 48 bytes of metadata are added, including the reference time of the accumulation cycle. The metadata expands the total amount of data leaving the correlator hardware by approximately 1%.

7.6. Timing and synchronisation

For a synthesis array radio telescope, the accurate distribution of a stable clock signal to multiple endpoints is a critical part of the instrument design. ASKAP is an array of arrays and this aspect of the system architecture is particularly important and challenging, where signal processing is distributed across some 4 000 separate FPGA devices and 8 000 separate ADCs. The station reference and timing system (Figure 18) must ensure synchronous sampling and processing of the data across the entire array and throughout the whole DSP chain.

ASKAP's timing system is based on a general architecture used at CSIRO's observatories (Hoyle & Mirschin 2015). It distributes a 16 MHz reference clock to phase reference the signals from the antennas and a serial-encoded timestamp to enable absolute timing of synchronous events across the array.

The 16 MHz reference clock is provided by a GPS (global positioning system)-disciplined rubidium time standard. The rubidium clock provides the necessary short-term stability while the GPS corrections ensure absolute long-term stability of the reference. Also present on site is a hydrogen maser that can serve as an external reference input to the rubidium sources. This is not used for typical ASKAP operation, but is used for VLBI (very long baseline interferometry) experiments (e.g., Kadler *et al.* 2016).

The rubidium time standard provides a one pulse per second (1PPS) signal as well as the 16 MHz reference clock to the ASKAP timing reference control computer (ATRCC), which generates a serial-encoded timestamp known as binary atomic time (BAT). The BAT serial data stream is phase-locked to the 16 MHz reference clock and encodes a 64-bit integer representing the count of microseconds since the epoch MJD 0.0, which is 00:00 UT on 1858 November 17.

The BAT serial stream and the 16 MHz reference clock are distributed together in a star topology from the ATRCC by the optical time and reference distributor, Figure 18, to the rest of the ASKAP DSP hardware subsystems (digital receivers, beamformers and correlators). The transmission is by single-mode, length-matched optical fibres using optical splitters. The BAT signal is transported using Manchester encoding to ensure no DC bias for compatibility with optical transmission.

As shown in Figure 18, the system has dual-parallel redundancy, and is remotely configurable via the Optical Time and Reference Distributor. Either of the redundant time standards can be selected to provide the primary 16 MHz and BAT references for distribution to ASKAP DSP endpoints. The final stage in the time and reference distribution occurs at the digital receiver, beamformer or correlator subsystems where the optically distributed 16 MHz and BAT signals are used to generate all the timing signals needed for the subsystem.

In each beamformer and correlator subsystem, the timing signals are received by the timing reference distributor (TRD). The digital receiver subsystems use the low-noise reference distributor (LRD). The architecture of the TRD and LRD are shown in Figure 19. In the LRD and TRD, the 16 MHz clock is multiplied up to 64 MHz and the 64 MHz and BAT signal are distributed by CAT-6 cables to all modules in the subsystem. These signals are used to generate synchronous processing events in all Redback and Dragonfly modules. Each TRD can supply signals to eight Redback modules.

The LRD has an additional low-noise distributor for the 64 MHz clock, which uses length-matched coaxial cable to distribute

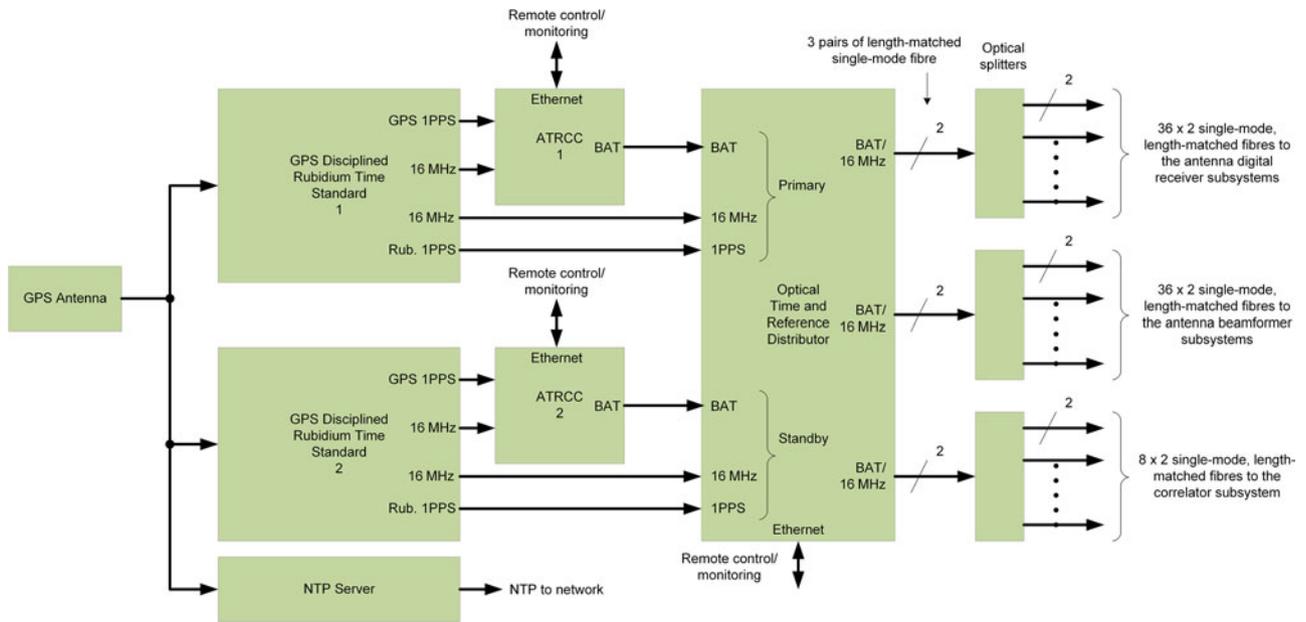


Figure 18. Station reference: ASKAP’s master clock generation and event timing system. The ASKAP Timing Reference Control Computer (ATRCC) block is an industrial PC with a custom PCIe card that generates a serial encoded timestamp referred to as Binary Atomic Time (BAT) that provides a precise, absolute time reference with a resolution of 1 μ s.

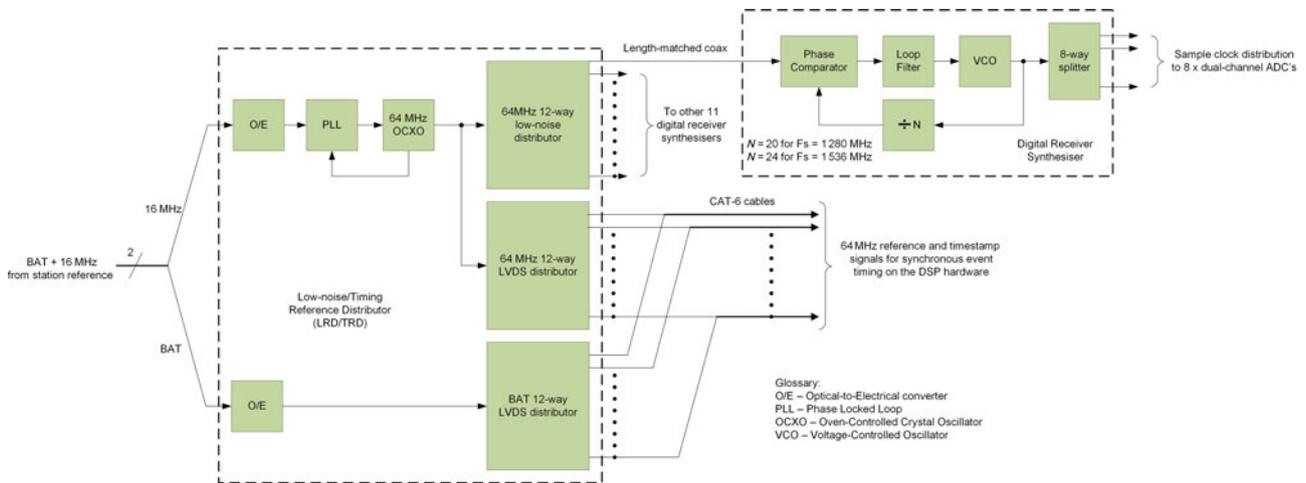


Figure 19. End point of the reference/timing distribution system. The low-noise reference distributor (LRD) generates and distributes a 64 MHz reference clock for the digital receiver sample clocks plus $12 \times (64 \text{ MHz} + \text{BAT})$ signals for synchronous event timing. The timing reference distributor (TRD) is the same hardware as the LRD without the separate 12-way 64 MHz reference distributor. The TRD is used for synchronous timing of events in the beamformer and correlator subsystems.

a low phase-noise copy of the 64 MHz signal to the 12 digital receiver Dragonfly modules. This generates a high quality ADC sample clock with low jitter.

8. Operational modes

ASKAP’s digital system is very modular and was designed to exceed the nominal bandwidth of 300 MHz for a full hardware deployment. For reasons outlined here, the existing system does not have a full complement of hardware and does not exercise all possible configuration flexibility, though more may be unlocked in future upgrades.

8.1. Digital system hardware considerations

The digital receiver outputs 384 channels of 1 MHz bandwidth each. Processing all of these would require eight beamformers per antenna, but only the first seven were installed. Wiring for the final beamformer is in place, but the construction budget did not allow the additional boards to be manufactured and they were not needed to meet the nominal bandwidth specification. Similarly, only sic of the full eight sets of correlator redbacks were installed. Since the correlator works in blocks of 48 MHz, meeting the nominal 300 MHz specification would have required installation and configuration of a partial block. While possible, this would have increased the complexity of control system logic and configuration

management, for only a few percent improvement in continuum sensitivity. Using all available test hardware and reducing the stock of spare parts, it may be possible to install a seventh correlator block. However, the data ingest and image processing system was only designed to work with up to 300 MHz of bandwidth, so expanding to 336 MHz will be left for a future upgrade once operational experience shows how much spare capacity exists in the post-correlator stages.

8.2. Correlator output modes

The flexibility of ASKAP's digital processing system allows each beamformer FPGA to operate in a different frequency resolution mode with arbitrary channels from the digital receiver. However, the complex channel routing logic required to implement a generalised configuration manager is prohibitive, and the control system is currently restricted to a single frequency resolution mode with contiguous bandwidth across all hardware. A summary of the available modes and the resulting bandwidth available for imaging is shown in Table 5. Regardless of frequency resolution, the 6-block correlator always outputs 15 552 spectral channels.

To reduce the amount of storage space required for initial capture of data intended for continuum science, the ingest pipeline software can perform online averaging down to 1 MHz frequency resolution. This reduces the storage footprint by a factor of 54 and can be used to balance resource requirements when scheduling observations. This averaging mode is only intended for use with the full bandwidth, not zoom modes.

At the time of writing a 10 s integration interval is used in the correlator accumulation stage, reducing the data volume by a factor of two, for a total of 1.2GB s⁻¹ leaving the observatory. If pilot surveys show that this accumulation interval does not introduce any significant smearing, it will be used for large-scale surveys as well. We have tested cycle times down to roughly 7 s successfully with the existing system, so there is some room for improvement. Attempts to record with the nominal 5-second cadence lead to data loss that will require more investigation to understand. An upgrade to the ingest computing hardware at Pawsey scheduled for November 2020 should also improve performance.

8.3. Data processing assumptions

At the time of design, ASKAP presented an extreme computational challenge. Short-term storage of the correlator output was deemed impossible (and still presents challenges, though advances in technology have made temporary storage possible) and efforts were made to minimise resource-intensive aspects of downstream processing. One such limitation is that the outer six antennas were not expected to be used in spectral line mode, since the size of the resulting cube would be too large. Advances in visualisation and analysis software make it desirable to consider creating high-resolution sub-cubes, at least for specific regions of interest.

9. Digital beamforming

The use of digital beamforming allows ASKAP to alter the illumination of the reflector antennas (and therefore the primary beam shape and direction) by changing numerical coefficients. This level of flexibility is of great benefit, but also significantly increases the complexity of the system. Unlike a mechanical feed, ASKAP's primary beam shape and receiver noise can change due to drifts in the complex gains of the elements that combine to form a beam.

Radio astronomy imaging algorithms typically assume that the primary beam is fixed, symmetric, and identical from one antenna to the next. For ASKAP, extra calibration steps are needed to ensure that these assumptions are met. The on-dish calibration system (ODC), described in Section 5.6, is used to monitor the complex gain of each PAF element. Although this does not provide a direct measure of the beam shape, it can be used to compensate for individual-element gain changes so that beams are stable over time.

ASKAP uses measurements of an astronomical reference source to determine the numerical weights used by the digital beamformers. A strong source is placed in the desired beam centre for each of 36 offset positions and the array covariance matrix $\hat{\mathbf{R}}_{n+s}$ is recorded. These observations are compared with a noise field $\hat{\mathbf{R}}_n$ that does not contain any strong sources.

9.1. Maximum sensitivity weights

We compute the maximum sensitivity (maxSNR) beamformer weights following the procedure described by Hotan *et al.* (2014) and McConnell *et al.* (2016), using the Sun as our reference source. The estimated array response $\hat{\mathbf{a}}$ of the PAF to a far-field signal incident from the direction of the desired beam is given by $\hat{\mathbf{a}} = \mathbf{u}_1$ where \mathbf{u}_1 is the dominant eigenvector of $(\hat{\mathbf{R}}_{n+s} - \hat{\mathbf{R}}_n)$ (Jeffs *et al.* 2008). The maxSNR weights (Lo, Lee, & Lee 1966; Applebaum 1976) are then computed as

$$\mathbf{w} = \hat{\mathbf{R}}_n^{-1} \hat{\mathbf{a}}. \quad (6)$$

The *noise-plus-signal* covariance estimate $\hat{\mathbf{R}}_{n+s}$ is measured via (4) with the antenna pointed so that the calibration source (Sun) is at the same position within the field of view as the desired beam pointing. The *noise* covariance estimate $\hat{\mathbf{R}}_n$ is measured with the antenna pointing at an 'empty' piece of sky, typically 15° south of the calibration source.

To ensure there is only one dominant eigenvector, we use only X-polarisation ports to make X-polarisation beams and only Y-polarisation ports to make Y-polarisation beams. One eigenvalue equation is solved for the sub-array of X-polarisation ports and then a second eigenvalue equation is solved for the near-independent sub-array of Y-polarisation ports. This means that the polarisation states and angles of the X and Y beams will be tied to the inherently linear and orthogonal polarisations of the X and Y elements of the chequerboard array. Commissioning measurements have shown that this technique results in beams with low and stable polarisation leakages (Sault 2014, 2015; Deng *et al.* 2017) and that ASKAP already has excellent polarisation imaging capability (Anderson *et al.* 2018).

Weights are calculated independently for each 1 MHz channel of each beam, which can lead to a randomised phase response. As described in McConnell *et al.* (2016), we adjust the phase of the weights at each 1 MHz channel to ensure a smooth phase response with frequency for each beam.

Digital beamforming with 1 MHz resolution introduces quasi-periodic bandpass errors with a 1 MHz period. To work around this, for spectral-line observations, we sometimes reduce the beamformer resolution by fixing the beamformer weights to be identical over sub-bands of $N \times 1$ MHz where N is an odd integer. This increases the magnitude of the bandpass errors, but ensures their smoothness over the larger $N \times 1$ MHz sub-bands. This is advantageous when looking for astronomical features at 1 MHz

resolution. It also reduces the probability of an astronomical feature falling on a sub-band boundary.

Since McConnell et al. (2016), we have extended our maxSNR calibration algorithm to:

- calibrate the XY phase of beams to simplify polarisation observations;
- accept the limit on the number of ports that our hardware can weight into each beam;
- avoid bias due to ports with abnormally high or low signal levels;
- correct weights for gain, phase and delay changes in the receiver electronics; and
- avoid bias due to RFI by interpolating weights with frequency.

9.2. Calibrating XY phase

Some of these changes were required to improve the robustness and reliability of the beamforming process for operational deployment. However, one great advantage of digital beamforming is that it allows control over aspects of the primary beam that would normally be fixed by construction. One such example is the relative phase of the two orthogonal polarisations.

We use the ODC system and the technique of Chippendale & Anderson (2019) to adjust the beamformer weights so that the XY phase of each dual-polarisation beam pair is near zero by design. Calibrating XY phase up-front in the beamforming allows polarisation calibration and imaging to take place in the standard ASKAP software pipeline, so that polarisation studies can be made commensally with all ASKAP observations.

9.3. Beamforming with limited port selection

ASKAP’s beamformer (see Section 7.2) can simultaneously form 36 dual-polarisation beams compared to BETA’s nine. However, as the number of processed beams increases towards the full 36, the number of ports that can be weighted into each beam reduces towards 60 according to Table 4. This is due to hardware resource limitations in the beamformer.

If all ports were functional and had the same gain and phase response, we could just select the weights with the largest amplitude when selecting the finite number of weights M that can be processed by the beamformer. However, the inversion of the noise covariance in (6) can lead to very high numeric weights for ports that have low gain. A first step towards avoiding these problems is to use RF attenuators along the signal-chain of each port to normalise the signal level that is presented to the digital backend from each PAF element.

9.4. Robust port selection

In the case of a malfunctioning port, with near-zero gain even after attenuator adjustment, the covariance matrices are poorly conditioned and can further complicate the solution for maxSNR beamformer weights. To avoid this problem, we use a diagonal loading process to avoid singularities in the inversion of the full covariance matrix. This adds a small amount of white noise to the noise covariance matrix before its inversion and application in (6) to calculate the weights

$$\mathbf{w} = (\hat{\mathbf{R}}_n + \alpha \mathbf{I})^{-1} \hat{\mathbf{a}}. \tag{7}$$

The diagonal loading places an upper limit on the weight that will be ascribed to any malfunctioning port with near-zero gain.

We use diagonal loading in a two-step process to robustly exclude malfunctioning ports from the beamformer weights. First, we select M ports with the largest signal contribution to the beam defined by (7). The signal contribution of each port is estimated with the aid of the ODC (see Section 9.5). Second, we recalculate the basic maxSNR weights using (6) on a reduced $M \times M$ covariance matrix of the selected ports.

The first step of the process tries to select the ports that are important to the beam, using diagonal loading to exclude malfunctioning ports. The second step applies the basic maxSNR algorithm, with no diagonal loading and therefore no bias, to the selected ports.

There is a sizeable literature on the appropriate selection of α in the diagonal loading process. We use an empirical technique guaranteed only to make improvements upon the maxSNR solution. We set

$$\alpha = \beta \text{median}(\text{diag}(\hat{\mathbf{R}}_n)) \tag{8}$$

and iterate the above solution over diagonal loading values of $\beta = [0, 0.1, 0.2, 0.3, 0.4]$, calculating the Y-factor (signal plus noise to noise ratio) at each step. Only solutions that improve Y-factor upon previous trials are retained. The first trial, with $\beta = 0$, is equivalent to the maxSNR solution with no diagonal loading, so the iterative process is guaranteed to achieve or improve upon the sensitivity of the basic maxSNR solution.

Comparison of SEFDs achieved by this two-step process to those achieved by the raw maxSNR algorithm show that it improves sensitivity over the raw maxSNR algorithm in almost all scenarios. The improvement is most noticeable when there are low-gain ports with significant energy falling on them due to the focal plane field of the desired PAF beam.

9.5. Focal plane field estimation

We use the technique of Chippendale et al. (2016a) to decouple the PAF element gains from the beamformer weights and estimate the focal plane field \mathbf{s} matched by beamformer weights \mathbf{w}

$$\mathbf{s} = \mathbf{d}_{\text{cal}}^* \circ \mathbf{w}. \tag{9}$$

Here, \circ is the Hadamard (element-wise) product and $\mathbf{d}_{\text{cal}}^*$ is the complex conjugate of the PAF’s response to a plane wave. The plane-wave response is estimated by correlating each PAF element with a copy of the broadband calibration noise that is radiated into the PAF

$$\mathbf{d}_{\text{cal}} = \frac{\langle \mathbf{x}_{\text{paf}} \mathbf{x}_{\text{rtn}}^* \rangle}{\langle |\mathbf{x}_{\text{rtn}}|^2 \rangle}. \tag{10}$$

Here, \mathbf{x}_{paf} is a vector of the PAF element voltages sampled at the digital receiver and \mathbf{x}_{rtn} is the sampled copy of the calibration noise that has been sent out to the antenna and reflected back to the digital receiver via a Faraday mirror as shown in Figure 11. We commonly refer to \mathbf{d}_{cal} as the ODC response. In fact \mathbf{d}_{cal} are the conjugate field match beamformer weights that would form an aperture array beam matched to the ODC source.

Ideally, the calibration source would present a plane wave to the PAF, illuminating all ports equally and in phase. In this case, \mathbf{s} is a good estimate of the focal plane field excited when the telescope observes a point source in the centre of the beam created

with weights \mathbf{w} . The ports with the largest component amplitudes in \mathbf{s} are the ports with most significant focal plane field excitation by the desired signal. They are the ports with the largest signal contribution to the beam.

9.6. Correcting PAF element gain drifts

We infrequently observe the Sun to make new beamformer weight solutions; typically after several months have passed or when significant changes have been made to the hardware. Between these solution intervals we use the ODC system to correct existing beamformer weights for changes in the complex gains of the PAF elements. This is effected by the Hadamard product of the original weights by the Hadamard (element-wise) division \oslash of ODC responses at the beamforming (subscript 0) and update (subscript 1) epochs

$$\mathbf{w}_1 = \mathbf{d}_{\text{cal},0}^* \oslash \mathbf{d}_{\text{cal},1}^* \circ \mathbf{w}_0. \quad (11)$$

The phase component of this correction is applied independently for each 1 MHz channel, but the median value over the 336 MHz beamformer bandwidth is used in amplitude. This forces the correction to be frequency independent in amplitude despite the ODC source having some undesired variation in spectral shape with dish antenna motion. We will revise the need for this smoothing in amplitude when the stability of the noise source is improved. The problem appears to be related to mechanical stability of the ODC vertex antenna and its RF connection to the noise source.

Since the ODC system is constantly injecting a low-level broadband reference signal into the feed (Chippendale *et al.* 2018), the array covariance measurements used for beamforming also capture a measure of the complex response of each element to the ODC signal. When we wish to update the beamformer weights, we record a smaller data product from part of the beamformer known as the calibration correlator (see Sections 5.6 and 7.2). Updated weights are calculated via (11) and uploaded to the beamformer.

This process requires only a small amount of data and since the beamformer weights are loaded via a double-buffered mechanism in the beamformer firmware, it is possible to update the weights without interrupting an observation in progress. The minimum update timescale is set by the time taken to compute, upload and switch to a new set of weights (a few seconds). However, the weights are currently left fixed during scheduling blocks. We perform an update before each observing session and leave the weights fixed until it is necessary to restart the digital receivers, typically when changing sampling bands.

Currently, the ODC correction in (11) stops beams from being completely destroyed by unpredictable delay changes that occur with each full reset of the ADCs. In the future we hope to use the ODC to stabilise PAF beam (voltage) patterns to better than 1% at the half power points (Hayman *et al.* 2010).

9.7. Avoiding RFI bias of beam direction

Basic maxSNR weights can be biased at frequencies with strong RFI. If there is RFI stronger than the beamformer calibration source in the measurement of $\hat{\mathbf{R}}_{n+s}$, the resulting beam will try to point in the direction of the RFI instead of the calibration source. These weights will have a low gain in the desired beam direction at the frequencies with strong RFI. The astronomical signal in the intended beam direction cannot be recovered when using these weights, even with careful flagging of the RFI-impacted data in the time domain.

We avoid this problem by interpolating valid beamformer weights for ASKAP at interference-affected channels following the technique in Chippendale & Hellbourg (2017). This uses iterative fitting of low-order polynomials to the weight amplitude spectrum of each individual port, taking advantage of the inherent smoothness of maxSNR weight amplitudes with frequency.

9.8. Using maxSNR beams for astronomy

Optimisation for sensitivity tends to create stronger primary beam side lobes than a mechanical feed and allows the beam shape to take on coma distortion off axis. Both of these can introduce undesirable effects in the imaging process.

Holography observations show that the total intensity beam shape is consistent with a circular Gaussian model to within roughly 10%, according to measured eccentricities (McConnell 2016) for a beam spacing of 0.9° . However, the level of coma distortion is a strong function of distance from the field centre, which means that the consequent variation in beam shape needs to be well characterised to form an accurate linear mosaic.

Over time we intend to develop other optimisations that may improve the symmetry, polarisation and other properties of the beams. A starting point will be linearly constrained minimum variance (LCMV) beamforming (Frost 1972), which can be used to optimise sensitivity with additional constraints on main beam shape. See Elmer, Jeffs, & Warnick (2014) for an example of constrained beamforming with a PAF.

9.9. Measurement of beams using holography

To assist with beam optimisation research, we have developed a raster grid holography method that provides a measure of the beam shapes over a large fraction of the field of view (Hotan 2016). Conducting a full measurement requires roughly six hours of observing time. This can be used to measure the stability and reliability of the beamforming process and assess the impact of changes to the constraints used when computing beamformer weights.

One of the next planned improvements to the imaging software will allow primary beam correction in the mosaicking stage using beam shapes measured via holography, instead of the circular Gaussian model currently applied.

9.10. Beam footprints for science observations

Digital beamforming provides the freedom to place primary beams anywhere within the field of view of the antenna. However, in practice, it is important to maximise the packing efficiency of the beams and develop standard arrangements that minimise sensitivity fluctuation across the field of view and also tessellate well when covering larger areas with multiple pointings.

We call the arrangement of beams within the field of view the footprint. Two footprints are in common use on ASKAP: `square_6x6` that follows a square grid and `closepack36` that offsets alternating rows to provide more uniform sensitivity as shown in Figure 20.

Note that neither of these footprints has a beam located on the optical axis of the antenna, which would have been desirable for calibration purposes. Having an extra beam would have allowed the use of a hexagonal arrangement containing 37 beams, including one on the optical axis. One of the future enhancements we are considering would allow the exchange of bandwidth for additional beams. As part of this work, we may investigate adding one

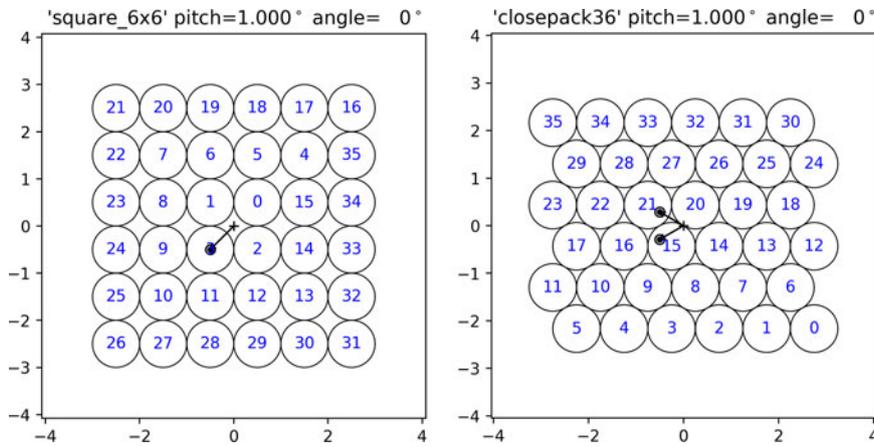


Figure 20. The two commonly used footprints: square_6x6 (left) and closepack36 (right). In this case, both have beam spacings (pitch) of 1.0° and a position angle of zero. The scales are in degrees, and celestial north (west) is to the top (right) of both diagrams. The + symbol represents the optical axis of the reflector and the lines extending from this position ending in filled circles represent the pointing shift required to optimally interleave the footprint by using other observations to fill in the least sensitive parts. The square arrangement has larger intrinsic sensitivity variation but requires only one interleaving position, while the closepack arrangement is more uniform initially but requires two interleaving positions for improvement.

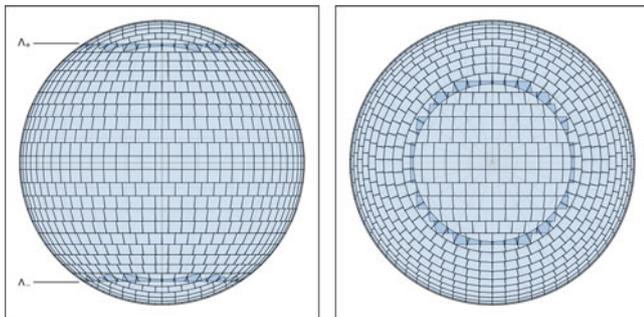


Figure 21. Two views of the survey tiling scheme: equatorial (left) and polar (right). The boundaries of the two polar zones at Λ_- and Λ_+ are indicated in the left-hand panel. The whole sphere is shown here; the tiling can be defined for any coordinate system (Equatorial, Galactic, etc.) and then transformed into the natural operating frame.

more beamformer engine to the standard full-bandwidth mode if resources on the FPGAs will allow it.

9.11. Tiling footprints for surveys

Since ASKAP will spend most of its time conducting all-sky surveys, it is useful to define a standard tiling scheme to cover the celestial sphere with the footprints described above. In consultation with ASKAP’s survey science teams^f, we have developed a tiling scheme in which most of the sky is covered by bands of footprints stacked end to end at constant declination (Figure 21). This becomes less efficient towards the celestial poles, so at 72 degrees above or below the equator we switch to polar caps consisting of a square grid that is truncated at the overlap with the outermost declination band. This approach allows a fixed beam footprint to be used for the entire survey. Methods such as HEALPix (Górski et al. 2005) were considered, but deemed unsuitable. HEALPix was designed to efficiently partition a spherical surface by allowing the shape of each partition to vary. Matching the different partition shapes would greatly increase the operational overheads associated with beamforming.

^f<https://www.atnf.csiro.au/projects/askap/ssps.html>.

The number of tiles required to cover the sky depends on the beam spacing within each tile, which can be adjusted to trade between survey speed and uniformity or mitigate widefield polarisation leakage. This optimisation is also a function of the observing frequency.

The first ASKAP all-sky survey, known as the Rapid ASKAP Continuum Survey (RACS), used a square_6x6 footprint with beam spacing of 1.05° at a centre frequency of 888 MHz. RACS consisted of 904 tiles observed for 15 min each, covering the entire sky south of +40° declination. A future publication will describe the forthcoming data release and source catalogue.

9.12. Beamformer weight storage and management

Beamformer weights are stored in a custom HDF5 (The HDF Group 2020) file format. Currently, each scheduling block specifies which weights file to use. We have developed, and are now commissioning, a database service that tracks weights files and can automatically retrieve the best available weights for a requested footprint and frequency setup. It is important to ensure that details of the beamformer weights are included in the metadata associated with an observation, to provide a complete record of the system configuration.

10. Raw data ingest

Visibility data leaves the ASKAP correlator hardware as UDP packets which are transmitted over a standard Ethernet network. An EPICS IOC running at the observatory gathers data from each correlator block and transmits it to the Pawsey supercomputing centre. Pawsey hosts a dedicated ingest cluster for ASKAP that is isolated from general purpose activity to ensure predictable load.

Each UDP packet sent from the observatory contains one fine channel (packed in a non-contiguous way), one beam and a range of baselines (the whole baseline space is split into four groups). A specially designed ingest pipeline software application receives the data using multiple threads, tags it with metadata from the telescope control system and reorders it into the form required for further processing. Data sent from each correlator hardware module are staggered in time to spread the network transmission load

out over the full accumulation cycle. The ingest pipeline can also perform basic tasks such as channel averaging if required.

Data are written in parallel to a series of CASA MeasurementSet (version 2, [Kemball & Wieringa 2000](#)) files on a Lustre^g file system that is common to all ingest nodes. We need to split each of the PAF beams into its own measurement set to avoid writing to a single file at an excessive rate. For full spectral resolution mode, we must further divide each beam into six frequency channel groups. ASKAP's image processing pipeline builds a complete image by combining data from all of these independent measurement sets.

Although the performance specifications of the file system are adequate for capturing ASKAP data at full rate, we have run into practical problems where other processes operating on the same Lustre hardware can impact performance.

11. Science data processing

The plan ([Cornwell et al. 2016](#)) for ASKAP science data processing (SDP) was to have a near-real-time autonomous pipeline. This would handle the calibration, imaging, cataloguing and archiving of the data, in a way that would keep up with the observing rate. The plan hinged on the ability to forward-predict calibration parameters for the array using a sky model and apply the solution on-the-fly.

Through commissioning and early science, we were unable to bootstrap this method of calibration and instead reverted to the use of dedicated observations of a standard calibrator source in each beam. Calibration can therefore be applied offline and improved iteratively with additional flagging, which has been essential to meeting data quality requirements. Online calibration is now considered a possible future improvement, once a suitable sky model has been obtained from the first all-sky surveys calibrated using traditional methods.

Using dedicated calibration observations increases operational overheads. Due to the large number of beams, the current calibration procedure uses about 2.5 h of observing time and must be done whenever the beamformer weights are changed or updated, which is typically once every day or two. This is feasible for survey observations, but adds scheduling complexity and reduces overall efficiency by about 10%. Although the time commitment for per-beam calibration using a reference source is significant, the storage requirements can be mitigated by keeping only the bandpass solution or discarding visibility data from beams when they are not pointing at the calibrator.

The batch-processing pipeline that has been developed provides the functionality required to process pilot survey data to a science-ready state. This section describes the high-level functionality of the pipelines, with more detail provided elsewhere ([Whiting 2020](#)).

The pipeline and the software it runs are collectively known as ASKAPsoft. The calibration and imaging software of ASKAPsoft^h has been custom-written to address the particular imaging requirements of ASKAP, especially imaging over a wide field-of-view and in near-real-time.

^g<http://lustre.org>.

^hThe core ASKAPsoft imaging software is available as the *Yandasoft* package <https://bitbucket.csiro.au/projects/ASKAPSDP/repos/yandasoft>. *Yanda* is the Wajarri word for *image*.

11.1. Pipeline workflow

The processing pipeline currently runs on *Galaxy*, the Cray XC30 at the Pawsey Supercomputing Centre. It creates a workflow from a series of connected compute jobs, running under the Slurm workload manager, with appropriate dependencies to ensure completion in the correct order.

The processing is mostly performed on individual beams independently, before combining images in a linear mosaic. The top-level workflow for a single beam can be expressed as follows:

1. Bandpass calibration is done using a separate observation of PKS B1934–638, establishing the overall flux scale. Frequency-dependent complex gains are determined for each beam and later applied to the science observation. See [Section 11.2](#).
2. Visibility data are prepared by splitting raw data measurement sets into single beam datasets (where necessary), then applying the bandpass calibration. The splitting may also create datasets for short time windows (typically one hour in length), which may be processed independently.
3. RFI is flagged, before the data are averaged to 1 MHz resolution (this will be used for the continuum imaging). Coarse-resolution data may be flagged again. If time splitting was used, the time windows are recombined.
4. The coarse-resolution measurement set is imaged and a sequence of self-calibration is performed. The self-calibration results in a fully calibrated image and table of time-dependent (but frequency-independent) complex gains. See [Section 11.2](#).
5. Both the coarse and full-resolution measurement sets are calibrated with the self-calibration gains.
6. If necessary, continuum emission is subtracted from the full-resolution visibilities. This emission may be characterised by clean components, fitted components identified through source finding, or through fitting to the visibility spectra. The optimal strategy for a given science case is being determined in conjunction with the Survey Science Teams, using the pilot survey processing.
7. The calibrated visibilities are imaged, producing final continuum images, coarse-resolution spectral cubes or full-resolution spectral cubes as required.
8. If necessary, any residual continuum emission remaining in spectral cubes (as a result of incomplete modelling of the continuum emission) is removed through an image-based continuum subtraction step.
9. Diagnostic and quality control information is assembled into a set of reports that are archived along with the image data.

Once all the beams have been imaged, full-field images and cubes are created through linear mosaicking. This corrects for the primary beam attenuation by applying a two-dimensional Gaussian primary beam model (see [Section 9.9](#) for efforts to develop more realistic models) centred on the position of each beam in the footprint. Once the mosaics have been created, source finding is run to create catalogues. The data products can then be uploaded to the archive (see [Section 12](#)).

11.2. Calibration

Most ASKAP observations are calibrated in two steps. First, the flux scale and frequency-dependent complex gains (the

'bandpass') are determined from a dedicated observation of PKS B1934–638. This is the primary southern-hemisphere flux calibrator, and has a well-determined spectral energy distribution (Reynolds 1994). The bandpass observation points each PAF beam in turn at PKS B1934–638 and observes for several minutes per beam. Although we currently assume a simple point source model for the calibrator source, Heywood et al. (2020) have shown that slight improvements could be made using a model of the full field that includes confusing sources within the primary beam. The calibration parameters derived from this observation are applied to the raw science data prior to any other processing.

Since the bandpass gains are determined from a different observation, some refinement to them is often necessary due to time-variable factors such as atmospheric and ionospheric conditions that can impact the calibration solution at a low level. This is currently done through a self-calibration process, using the continuum-averaged dataset. A continuum image is made, then calibration is performed using a shallow sky model derived from that image. The complex gains are solved for within short (~ 5 min) intervals across the observation. The data are then re-imaged using this updated gain calibration to provide an improved result. This process can be iterated several times if required. We typically solve for the phases of the complex gains and leave the amplitudes fixed, in order to preserve the overall flux scale. Once complete, the gain solution can be applied to the full-spectral-resolution visibilities to facilitate spectral imaging.

11.3. Wide-field imaging

Wide-field imaging is enabled via the W -projection algorithm (Cornwell, Golap, & Bhatnagar 2005), allowing handling of image-plane effects due to non-coplanar baselines. This gives accurate imaging over the field of view of a single beam. Multiple beams could be combined via AW -projection (Bhatnagar et al. 2008; Bhatnagar, Rau, & Golap 2013), taking into account the aperture illumination, although this approach is currently not used due to its large memory requirements. As described above, the aperture correction is instead applied when mosaicking the individual beams together.

The near-real-time imaging requirement has driven a number of design decisions within the code. The weighting of visibility data prior to imaging is traditionally done with an initial pass over the data to determine the distribution of visibilities as a function of baseline length and hence the weighting function, and then a subsequent pass for the imaging. To prevent repeated iteration over the large ASKAP data products, the ASKAPsoft code incorporates the *preconditioning* approach (Rau 2010), where the re-weighting is performed after the construction of the dirty image and point spread function (PSF). ASKAPsoft imaging makes use of robust weighting, similar to that described in Briggs (1995), allowing a tradeoff between the suppression of sidelobes and the overall sensitivity.

Continuum imaging is performed using Taylor-term multi-frequency synthesis (Rau 2010; Rau & Cornwell 2011). This creates an image of the total intensity at a reference frequency, along with images of the spectral variation across the band, from which can be determined the spectral index (and curvature, if enough terms are used) of sources. For spectral imaging, each channel is imaged independently, using either the averaged (continuum) data to create 'continuum cubes', or the full-resolution data to create spectral cubes. Continuum cubes can be made in all Stokes parameters, allowing polarisation imaging.

The imaging software has distributed processing designed in to the code at a low level. This allows the processing to be distributed over many compute nodes. The continuum imaging has the gridding distributed per coarse channel and Taylor term, with the deconvolution currently performed on the head node after combination of all channels. The spectral imaging is performed independently on each channel, and so is extremely parallel, allowing distribution over as many nodes as can be accommodated.

It should be noted that this extremely parallel case assumes a coarse form of Doppler correction done per beam and with only whole-channel offsets. This may be a limiting factor for some science applications.

Images are currently stored in FITS format (Pence et al. 2010) for broad compatibility and because the science data archive was designed around this format. ASKAP data cubes can be very large (tens of terabytes) and it is possible that future research and development may provide a more efficient data storage format for use in parallel processing environments with remote visualisation services. If a new standard is established, we will investigate upgrading ASKAP's software platforms to incorporate it.

11.4. Source finding and cataloguing

As part of the pipeline processing, we make catalogues of the sources in the images, using the Selavy source finder (Whiting & Humphreys 2012). Selavy uses the source-finding algorithms of Duchamp (Whiting 2012), a source finder developed for three-dimensional datasets. It can produce catalogues of spectral-line sources from the spectral cubes, largely following the Duchamp methodology, as well as catalogues of continuum sources.

The continuum source finding produces two different catalogues. The first is the catalogue of *islands* in the image. These are groupings of image pixels that are above the detection threshold. Each island is parameterised by its location and shape, determined solely from the detected pixels. To each island, we fit some number of two-dimensional Gaussians—these are the *components*, and the second catalogue reports their parameters. Each component will have a location, flux and shape, as well as information on the spectral index. This is obtained by fitting the Gaussian in question to the Taylor-1 image, which gives the product of the Taylor-0 flux and the spectral index.

Source finding is an active topic of research within the broader community and some of the ASKAP survey science teams will produce their own catalogues using different methods and tools. These are referred to as 'value-added' data products. For example, the SoFiA spectral-line source-finding software developed by Serra et al. (2015a) is being used on neutral hydrogen (HI) emission-line pilot survey data, and the FLASHfinder line-finding software developed by Allison, Sadler, & Whiting (2012) is being used to search for HI and OH absorption lines on sightlines to bright radio continuum sources in ASKAP spectral-line cubes.

12. Data archiving and access

Science data products produced by the ASKAPsoft pipeline are archived and made available to the astronomy community by the CSIRO ASKAP Science Data Archive (CASDA¹). Data deposited into CASDA are made available to the survey science team (SST) members initially, and then made publicly available after quality control and validation.

¹<https://research.csiro.au/casda>.

CASDA is described in detail in Chapman *et al.* (2017) and Huynh *et al.* (2020). In brief, CASDA is implemented across two data centres, Pawsey and the CSIRO data centre in Canberra. The CSIRO data centre runs the CSIRO data access portal, an enterprise-wide system that archives and provides access to data across many areas of CSIRO research, including ASKAP/CASDA. The ASKAP data are deposited on to tape and disk storage at Pawsey, with storage managed by the next-generation archive system (NGAS) (Wu *et al.* 2013). Functions such as metadata search and authentication are implemented at the Canberra data centre, while functions that need to be close to the data, such as data deposit and data access, are implemented at Pawsey.

At full operations the ingest data rate of ASKAP is expected to be approximately 200TB day^{-1} , or more than 70PB year^{-1} , for 100% duty cycle. Given this extremely high data rate, the full resolution uncalibrated visibilities are not stored. The data products archived and served by CASDA are as follows:

- calibrated visibilities (continuum resolution);
- images and cubes, including intermediate imaging products such as clean models;
- catalogues; and
- metadata from ASKAPsoft and quality assessment.

In addition, derived data products such as moment maps and spectra of detected sources will also be stored in CASDA, once they are produced by ASKAPsoft (planned in future development). Given $15\text{--}20\text{ TB day}^{-1}$ of expected science data products from ASKAPsoft, we expect to have more than 5PB year^{-1} of data flowing into CASDA during full operations. CASDA has a current allocation of 10 PB of long-term tape storage (+10PB for redundancy) at Pawsey.

CASDA data search and access is possible through both the data access portal web user interface¹ and virtual observatory (VO) services. CASDA implements a wide suite of VO services to maximise the usability and interoperability of ASKAP data products. For example, the table access protocol (TAP) can be used to search for specific observations via an application such as TOPCAT (Taylor 2005). Uploaded catalogues can be filtered for desired radio sources using the astronomical data query language (ADQL). Python scripts can be used for automated data discovery and retrieval, including the production and downloading of cube cutouts.^k A CASDA module has also recently been added to the `astropy astroquery`^l package.

13. Performance measurements

In this section, we provide a broad overview of measured key performance parameters. Detailed analysis of ASKAP pilot surveys will yield further performance information that will be summarised in future papers.

13.1. Sensitivity

The sensitivity of each ASKAP antenna depends on the physical area A of the reflector aperture and the effective system temperature T_{sys}/η achieved by the receiver as installed on the reflector.

¹<https://data.csiro.au>.

^k<https://github.com/csiro-rds/casda-samples>.

^l<https://astroquery.readthedocs.io>.

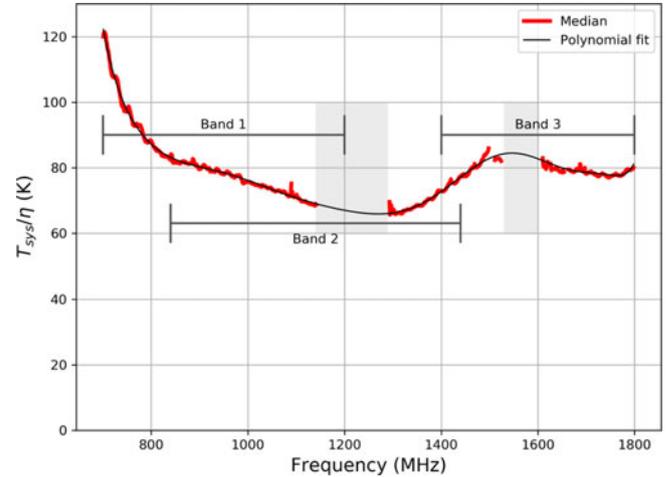


Figure 22. Effective system temperature T_{sys}/η across the ASKAP band. The median value over all antennas is plotted (red) for a beam close to the antennas' boresight. The black curve is a polynomial fitted to the spectrum. Sensitivity estimates in the shaded frequency ranges are difficult because of persistent RFI. The ASKAP operating bands are also shown.

Here, T_{sys} is the system equivalent noise temperature and η is the antenna efficiency.

Both T_{sys} and η vary with beamformer weights for a PAF receiver, so we report results for beamformer weights that maximise sensitivity (as used in operation; see Section 9). Figure 22 gives T_{sys}/η across the ASKAP band measured with observations of PKS B1934–638 using the same technique as McConnell *et al.* (2016). We use 75 k as a nominal value of T_{sys}/η when summarising ASKAP performance.

The mid-band T_{sys} of an ASKAP PAF beam, excluding spillover contribution from the ground, is 50 k (Chippendale *et al.* 2016a). This measurement was made with the PAF on the ground pointing at the zenith, but using the same MaxSNR beamformer weights as are applied when the PAF is installed on the reflector.

Astronomers characterise sensitivity of each antenna plus receiver system in terms of the system equivalent flux density (SEFD) defined by

$$\text{SEFD} = \frac{2kT_{\text{sys}}}{\eta A}. \quad (12)$$

This is the flux an unpolarised point source must have to yield a signal-to-noise ratio of one at the output of a given beam. The factor of two appears in (12) because a single-polarisation beam can only collect half of the total incident flux from an unpolarised source (Wrobel & Walker 1999). ASKAP's nominal T_{sys}/η of 75 k corresponds to an SEFD of approximately 1 800 Jy for a single ASKAP antenna.

The sensitivity of the telescope can also be defined by the effective area divided by the system temperature

$$S(\Omega) = \frac{A_e}{T_{\text{sys}}}. \quad (13)$$

This is generally a function of direction Ω within the field of view and antenna efficiency η achieved together with the PAF receiver. The sensitivity of the full ASKAP array evaluates to $54\text{ m}^{-1}\text{K}^{-1}$ for a beam close to the boresight.

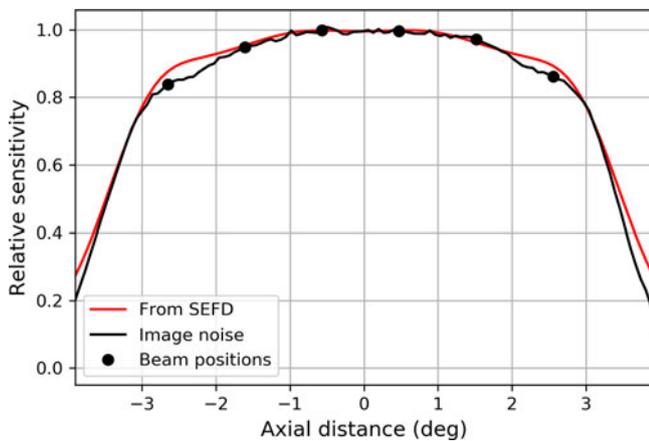


Figure 23. The sensitivity profile over the field of view. The black line traces the observed sensitivity (see text), and the red line shows an analytic approximation of the sensitivity estimated from SEFD observations. These data were obtained from observations using the square_6x6 footprint with a beam pitch of 1.05° at a centre frequency of 888 MHz.

13.2. Field of view

ASKAP’s field of view is set primarily by the size of the array of receptors at the focal plane and the size and focal ratio of the reflector. It is also modified by the aperture illumination, which is in turn controlled by the beamformer weights.

A measurement of the field of view with maxSNR weights was made using the ASKAP-12 array and reported in McConnell (2017b). More recently, that measurement was compared with the observed variation of image noise over the field of view. Figure 23 shows this comparison in one dimension as an east-west profile of image sensitivity across the median of a sample containing 44 36-beam fields. The observations were made with the square_6x6 footprint using a beam spacing of 1.05° at a centre frequency of 888 MHz. Relative sensitivity is calculated as the reciprocal of the image rms, divided by its maximum. Sensitivity is maximum near the centre of the field of view and is reduced by less than 20% at the edge.

These measurements show that ASKAP achieves its design specification of a 30 deg² field of view at 850 MHz. The field of view reduces to 15 deg² at 1 700 MHz due to the current limit of 36 dual-polarisation beams. The field of view at the high end of the band may be increased to 30 deg² in the future by adding or reconfiguring digital hardware to process more than 36 beams.

13.3. Survey speed

As seen in Figure 23, the shape of ASKAP’s field of view is very different from that of a dish with a single feed, which is approximately Gaussian. The traditional survey speed figure of merit (1) should not be calculated using the half power field of view. When the sensitivity varies significantly over the field of view, (1) should be viewed as a point estimate that is integrated over the field of view to yield the survey speed (Bunton & Hay 2010)

$$SS = \int S^2(\Omega)d\Omega. \tag{14}$$

Equivalently, (1) may be used with the maximum sensitivity S_{\max} and a survey speed weighted field of view

$$\Omega_{\text{FoV}} = \frac{1}{S_{\max}^2} \int S^2(\Omega)d\Omega. \tag{15}$$

All ASKAP survey speed and field of view numbers are calculated by (14) and (15) in this paper.

Using this definition, the survey speed of ASKAP is currently 91400 m⁴ deg²K⁻² at 800 MHz with a smooth reduction to 44 200 m⁴ deg²K⁻² at 1 700 MHz. It may be possible to double survey speed at the high end of the band in future by processing more than 36 dual-polarisation beams. However, with current hardware, the extra beams would come at the expense of bandwidth.

As a point of comparison, evaluating (14) for a Gaussian beam yields a survey speed of 0.721Ω_GS_{max}², where Ω_G is the half power field of view of the Gaussian beam. Using this measure, the survey speed of ASKAP at 1.4 GHz is 245 times faster than the Parkes 64 m dish with a cryogenic single-pixel receiver having T_{sys}/η of 40 k (Hobbs et al. 2020).

14. Radio-frequency interference

ASKAP’s location on a protected radio quiet site makes it possible to observe in parts of the spectrum that are no longer usable at other radio telescopes. In particular, the band from 700 to 1 080 MHz is almost entirely free of persistent RFI. Occasionally, atmospheric conditions can lead to ducting of distant mobile communications signals, but this can be predicted to some extent based on meteorological data (Indermuehle et al. 2018b). and impacts relatively few observations.

Satellites and aircraft have been the most common sources of RFI in observations to date (Indermuehle et al. 2016). Aircraft automatic dependent surveillance broadcast (ADS-B) transponders transmit at 1 090 MHz and are visible when flights pass over the site. This signal is strong, but has a relatively low duty cycle so future real-time interference mitigation methods may be able to remove it.

Satellite navigation systems are visible between approximately 1 150–1 300 MHz. Occupancy of interference approaches 100% at several places within this range and typical flagging algorithms leave very little usable data. The band from 1 500 to 1 620 MHz also contains significant satellite interference, with up to 70% occupancy.

In future, adaptive beamforming could be used as a form of RFI mitigation on ASKAP, by placing nulls at the locations of transmitting satellites (Black et al. 2015; Hellbourg et al. 2014, 2016). This would involve changing the beams during an observation, which would need main-beam constraints and corrections for pattern rumble bias (Jeffs et al. 2008) to avoid invalidating array calibration solutions.

15. Telescope operations

In a departure from other Australia Telescope National Facility (ATNF) telescopes, 75% of observing time on ASKAP has been pre-allocated to several international survey science teams with 5-year observing plans. The remaining 25% of available time will be allocated to guest science projects, subject to peer review through the ATNF time allocation committee (TAC). ASKAP will be operated by CSIRO staff and astronomers will not interact directly with the telescope. Survey plans and guest science proposals will be converted into scheduling blocks that are observed, calibrated and imaged as part of operations. Science-ready images will be made available on CASDA. This means that calibration, imaging, archiving and initial quality control all fall within the domain of the observatory.

Multi-year surveys are commonly considered legacy projects at other observatories and are usually done after the telescope

Table 6. ASKAP science operations timeline

Dec 2016	Early science begins
July 2018	Early science ends
Feb 2019	Full array operational
July 2019	Pilot Surveys Phase I observing begins
May 2020	Pilot Surveys Phase I observing ends
Q1 2021	Pilot Surveys Phase II planned start
Q3 2021	Pilot Surveys Phase II planned end
Q4 2021	Full Surveys planned start

has been operational for many years. Since ASKAP is conducting large-scale survey projects as one of its first activities, we devised a sequence of steps to test the telescope's readiness. See Table 6 for a timeline showing significant milestones on the road to full survey operations. Descriptions of the various phases can be found below.

One of the most important lessons from the early science programme was the need to maintain close engagement with the science community. Taking on the responsibility of providing science-ready data products means that the observatory must know exactly what is required by its users. We developed a commissioning team that included representatives from the observatory and the survey science teams working together.

Due to the high data rate, we do not plan to keep all visibilities long term, although we will archive visibilities that have been calibrated and averaged to 1 MHz frequency resolution. We have also found it useful to keep raw calibration data, which accumulates at a more modest rate than the science observations. Outputs from the imaging pipeline are tested for quality and reprocessed if necessary, with raw data being deleted to make way for the next project once the image products have been archived.

15.1. Early science

The ASKAP early science programme started with 12 antennas equipped with Mk II PAFs in 2016. We planned to run two survey projects on behalf of all the survey science teams, but found that this was not feasible so early in the life of the instrument. Instead we devoted time to smaller test observations for each science team, using the data to test ASKAP's image processing software and identify problems with the visibility data. These test observations were separated by commissioning periods of a few weeks duration, where antennas were added to the array and features necessary for full operations were implemented.

During this time, various existing software tools were used to cross-check and supplement ASKAP's custom processing software as it was developed. Eventually we were able to complete a modest continuum survey known as the cosmology survey, which is available on CASDA.

Many publications arose from the early science programme, as a result of identifying projects that the telescope could excel at with only a third of its collecting area. These include a continuum survey and detection of cold gas outflows from the small magellanic cloud (McClure-Griffiths *et al.* 2018; Joseph *et al.* 2019), studies of neutral hydrogen in nearby galaxy groups (Reynolds *et al.* 2019; Lee-Waddell *et al.* 2019; Elagali *et al.* 2019; Kleiner *et al.* 2019; For *et al.* 2019), studies of absorption lines (Glowacki *et al.* 2019; Allison *et al.* 2020), continuum observations of the galaxy and mass assembly (GAMA) G23 field (Leahy *et al.* 2019), searches

for transient and variable sources (Bhandari *et al.* 2018), and a few targeted studies of individual radio galaxies (e.g. Seymour *et al.* 2020), including the polarisation characteristics of Centaurus A (Anderson *et al.* 2018). ASKAP was also the first telescope to localise a non-repeating fast radio burst to its host galaxy (Bannister *et al.* 2019b).

15.2. Pilot surveys

Integration of all 36 antennas with Mk II PAFs and digital back-ends was completed in February 2019. Soon after, we commenced a series of test observations designed to verify all the modes required for survey operations with the full array. This involved relatively simple imaging of representative fields. On 2019 July 15, we began a pilot survey programme to demonstrate readiness for extended science operations, and to highlight places of improvement for sustainable full operations. Pilot surveys involved observing small parts (100 h per project) of the larger survey plans developed by the survey science teams. The pilot survey concept arose as a natural extension of the early science programme. Reaching peak operational efficiency is a process that takes time and these pilot surveys provide useful science data while improvements and updates to the telescope's core systems are ongoing.

There are several goals associated with pilot surveys:

- Provide representative data to each SST.
- Test SST survey strategies.
- Push operational limits to find pressure points.
- Develop operational plans for full surveys.
- Assess and improve processing pipelines.

Observations for phase I of these pilot surveys were completed on 2020 May 14, although processing is expected to take several additional months. Upon completion of pilot surveys phase I observations we entered a 6-month consolidation period, in which the priority for telescope access shifted back to development, testing and maintenance. This provided time to consider lessons learned and deploy various software and firmware updates, while processing the data backlog. The primary goal of pilot surveys phase I was to demonstrate production of science-ready data products. Phase II of pilot surveys will focus on optimising commensality and overall data processing efficiency. We are not currently planning a third phase of pilot surveys, but this may be considered if significant issues arise from phase II.

Full survey operations are expected to begin by the end of 2021 and ramp up in efficiency by 10% each year, reaching a limit of 70% efficiency in 2024.

15.3. Specifying observations

The ASKAP control system implements a flexible observation management system built around the concept of observing procedures and scheduling blocks. Procedures are implemented as Python scripts using a library of functions that provide access to the hardware. The simplest (and most commonly used) procedure simply tracks a source position specified in equatorial coordinates. More complicated procedures have been developed for band-pass calibration and holography observations where on-the-fly calculation of offset tracking positions is required.

A scheduling block is specified as a list of parameters and an associated procedure. The list of parameters currently includes source coordinates and system configuration options. Eventually,

we plan to expand this list to include data processing parameters, which will allow automated execution of the image processing pipeline upon completion of a scheduling block. All-sky surveys require roughly 1 000 individual fields, but we have found in practice that scheduling each one individually (rather than specifying multiple fields in a single block) is the most effective approach, since it provides a level of fault tolerance and allows for more predictable timing.

15.4. Observation scheduling

Over the course of the pilot survey period, various scheduling requirements have been incorporated to meet the needs of different survey science cases. Developing a working model for efficient and effective scheduling is an ongoing process that depends on the level of system stability, robustness and automation. For now, scheduling is carried out in one of two modes: (i) transit-centric or (ii) target of opportunity. The first approach applies largely to long tracks, where the requirement is to schedule a long observation while the target is above the horizon and pair it automatically with a bandpass calibration observation. The second approach is more suitable for short tracks such as the 15 min observations used for the Rapid ASKAP Continuum Survey. It selects an appropriate target from a specified list based on a series of constraints (e.g. above the horizon, solar distance, lunar distance, hour angle, etc.). The scheduling algorithms are being refined and evolved as we gain experience with the requirements of survey projects and ensure the system is able to handle these in a way that is as automated and efficient as possible.

Our current intention is to aim for autonomous scheduling, in that the scheduling tool will detect system status and environment and have access to the pool of survey observations as well as survey constraints and historical observational data. It would then use this information to decide on the next observation to schedule and ensure the telescope is in the right state to complete this observation. The scheduling tool would also provide a forecast of the coming observations if the system status remained the same, but it would maintain the flexibility to adapt as necessary based on the changing system and environment. We intend to use the next period of pilot survey observations to prototype this approach and determine its feasibility for full surveys, ultimately maximising automation as much as possible which will help us to maximise efficiency gains. Recent developments by the ASKAP team to enable automated array start-ups based on the input parset have been a key milestone towards this vision for scheduling, and encourage us that our goal of autonomous scheduling may be realised for ASKAP.

15.5. Data quality diagnostics

ASKAP presents an incredible big data challenge, marking a noticeable shift in the ways astronomical data can be managed and interacted with. It is quite difficult to gauge data quality in meaningful ways when data sizes are tens of terabytes, including complex multidimensional axes across time, frequency, antennas, beams and polarisations. While working within the constraints of no longer being able to visualise all data, effort has been invested to ensure that certain aspects of the data can be visualised and used for diagnostic purposes, particularly to ensure that the system is working as expected in a timely way.

For the raw data, a series of plots have been designed and implemented to allow insight into the recorded data regardless of their

size. Though currently the amount of time for processing does scale with data size (improving this is the subject of ongoing investigation), we are able to obtain these raw data diagnostic plots automatically in close to real time for short observations, which has been important for determining system status. These plots are extensions of standard astronomical visualisations of the visibilities, but necessarily more complex due to the amount of data needing to be represented. The exact form of ASKAP diagnostics is under development as part of the continuing transition between commissioning and operations. The ultimate goal is to assess data quality in an automated way, making use of anomaly detection algorithms and other machine learning techniques where feasible, and this is the subject of further investigation. Currently, diagnostic plots are visually inspected where necessary and allow human pattern recognition to spot anomalous conditions.

Some examples of diagnostic plots for a continuum dataset (SBID 13830 and 14997) are shown in Figure 24, demonstrating visualisation of (correlation) amplitude versus time, amplitude versus frequency channel, and time-frequency autocorrelation waterfall visualisations that combine the two. In particular, the last two plots are an attempt to visualise the contents of the entire dataset at once, showing a time-frequency waterfall plot of autocorrelation amplitude for each antenna and beam within the array. This is extremely effective at highlighting particular hardware issues, but less so for issues that depend on cross-correlation baselines. Although in this visualisation each subplot is too small to be used for quantitative analysis, broad trends or hardware faults stand out clearly against a background of relatively uniform behaviour. Conversely, the data used to generate these visualisations is extremely rich in diagnostic information, and will likely form a useful input dataset for either classification or machine-learning algorithms.

Extensive quality control is also done post-imaging. Diagnostics are being developed by the science teams and will be incorporated into the standard observatory-driven processing pipelines over time. CASDA includes a validation stage that must occur prior to data release. Currently this requires human inspection of diagnostic information and selection of a three-tier quality grade. As diagnostic algorithms improve, we will seek to automate this process as much as possible.

16. Future enhancements

One of the first upgrades to ASKAP's systems could be a coherent transient detection module that improves the sensitivity of fast radio burst searches. This would provide access to an alternative search data stream from the correlator rather than the beamformers, increasing the sensitivity to fast transients by a factor of six and providing near-instantaneous arcsecond localisation.

Another possible improvement is the addition of a tied-array module to create phased array beams from all 36 antennas. This would allow the full array to be used for very long baseline interferometry, pulsar timing experiments, and the search for extraterrestrial intelligence.

Satellite interference breaks up ASKAP bands 2 and 3 into clean sections smaller than the telescope's instantaneous bandwidth, so it could be beneficial to split the observing band itself into two or more disconnected spectral windows. Although this is not currently possible, we are planning to develop such a split-band mode through software and control system upgrades. This would only be possible within one of the band-limiting filters

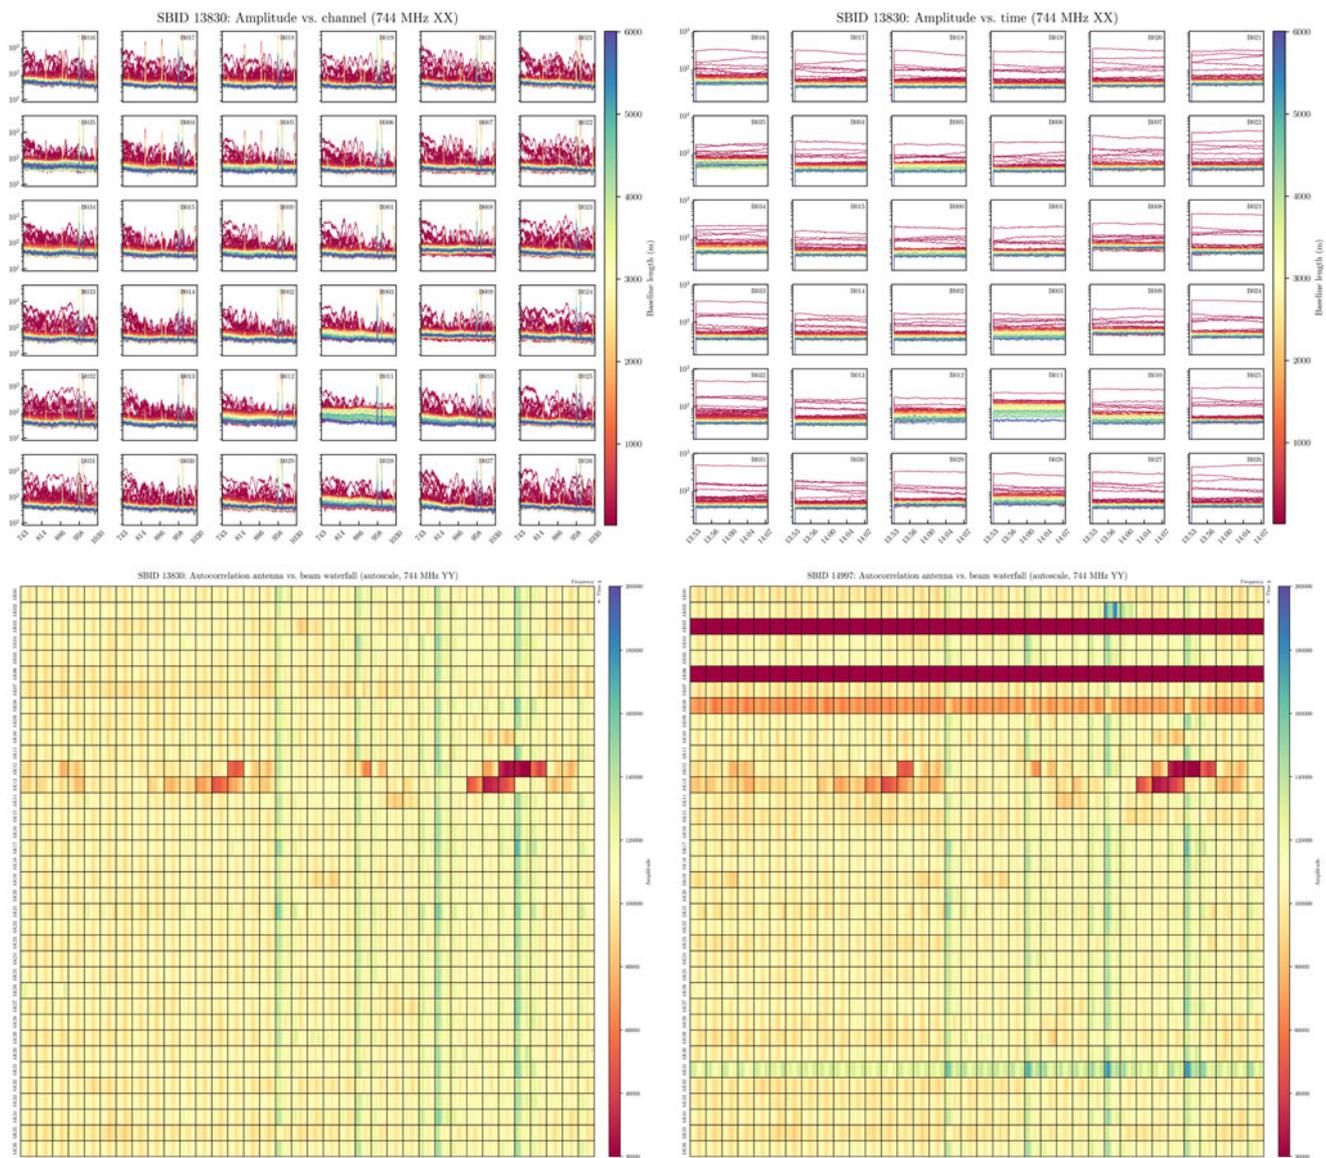


Figure 24. Example ASKAP raw data diagnostic plots for Scheduling Block IDs 13830 and 14997, demonstrating visualisations of correlation amplitude versus frequency (*top-left*), correlation amplitude versus time (*top-right*), and time-frequency waterfall of autocorrelation for all beams and antennas (*bottom-left*) for SBID 13830. For comparison, we show the same diagnostic plot for SBID 14997. These diagnostics give us a direct insight into the recorded data prior to processing, and are used to inform overall data quality assessment. In the top two plots, each line represents a baseline colour-coded by baseline length (where red is shortest and purple is longest). In the autocorrelation waterfall plots, orange and red sections indicate beams on particular antennas that may be affected by faulty PAF dominos (in this case, the same bad YY elements are seen on ak12/ak13 for both observations), while the four vertical green stripes across entire beams trace the less-sensitive beams at the corners of the field of view. Dark red stripes indicate a missing antenna, and any antennas with visually outlying amplitudes are likely to have scaling issues. The vertical spikes seen in amplitude versus channel and the baseline-average waterfall plot are RFI.

(see Table 3). Split-band mode would also allow simultaneous observations of widely separated spectral lines and may also be beneficial for determining spectral indices and rotation measures.

Another possible improvement would be allowing the exchange of bandwidth for additional PAF beams. This would be particularly beneficial at the high-frequency end of ASKAP's frequency range since, above 850 MHz, 36 dual-polarisation beams cannot fully sample the 30 deg² field of view of the PAF (Bunton & Hay 2010; McConnell 2017b). At 1 700 MHz, the field of view is currently limited to 15 deg². Processing more beams may therefore double the current field of view at 1 700 MHz. Additional beams would also reduce the need for interleaved observations,

but further analysis of the effects of correlated noise between closely spaced beams may be needed to understand the exact benefit (Serra *et al.* 2015b; McConnell 2017a,b).

An additional 48 MHz correlator block could be installed to bring the instantaneous bandwidth to 336 MHz. To achieve the full 384 MHz bandwidth a full complement of Redback modules for the beamformer and correlator are needed. The additional data rate would be difficult to support on the current supercomputing platform and would diminish the amount of spare hardware modules, so any such upgrade needs careful consideration. Alternatively, upgrading the correlator with new hardware would release sufficient hardware to expand the beamformer (since the

same physical boards are currently used in both subsystems). An option for a new correlator is the Gemini (Kooistra et al. 2017) board (currently in development), which is a good match to the current Redback modules. An upgraded correlator would allow increased frequency resolution across the band to 9.2 kHz or 4.6 kHz. Implementing tied-array beams would also be easier with new correlator hardware, and the extra resources would likely make an increased number of tied array beams possible. The upgrade would also allow full bandwidth ‘standard’ resolution observing with simultaneous zoom bands, increased DM range for transient searches and integration of transient searches within the correlator hardware.

Increased sensitivity could be obtained with an increase in antenna diameter. The diameter of the existing antennas could be increased to 15 m or more while still maintaining ground clearance at minimum elevation. Initial work shows the existing counterweights would be insufficient to support a heavier reflector and a lightweight design would be needed. The implications of additional wind loading on the drives would have to be carefully considered as well. At 15 m, the ASKAP dish has an f/D similar to Parkes, which has been tested successfully with an ASKAP PAF (Chippendale et al. 2016b; Chippendale & Hellbourg 2017; Deng et al. 2017; Reynolds et al. 2017). With the additional collecting area, the instantaneous sensitivity of ASKAP increases by 56%. Survey speed is also increased, but more beams may be required to compensate for the reduced primary beam size.

Sensitivity could also be improved significantly by updating the room temperature LNA design with new transistors (Shaw & Hay 2015; Weinreb & Shi 2019) or by scaling up the manufacturability and affordability of cryogenic PAF technology like that under development for the Parkes 64 m telescope (Dunning et al. 2016, 2019), as mentioned earlier in Section 1.1.

17. Conclusion

ASKAP is one of the first radio telescopes to employ PAF technology, giving the telescope a wide field of view, rapid survey speed, and excellent polarisation characteristics. This has been demonstrated in several pilot surveys and observatory projects that will be published elsewhere. We have shown that PAF technology is a powerful solution for wide-field radio astronomy. So far we have used ASKAP’s PAFs in familiar ways, creating beams that are treated independently and kept fixed for the duration of an observation. PAFs offer much more power and flexibility than we hope to explore in future, including adaptive interference mitigation.

ASKAP’s survey science projects will occupy most of the next few years of observing time and all data will be released to the public after quality control. At the present time, ASKAP pushes the boundaries of signal processing technology and requires a combination of techniques to convert the torrent of raw data into calibrated, science-ready images that can be archived permanently. As computing technology improves, we expect that it will be possible to extract even more information from the data. One important design consideration for other radio telescopes intending to deliver science-ready output is to budget for reprocessing. Radio astronomy imaging often requires iteration to perfect, which can be very resource-intensive.

As a precursor telescope for the Square Kilometre Array, ASKAP has demonstrated several key lessons. It is important to consider hardware, firmware and software development as connected activities and integrate subsystems as early as possible.

Design aspects may need to be changed on the basis of experience in the field, so it is important to begin science-oriented test observations as soon as possible. Commissioning a new radio telescope begins during subsystem tests and continues well into operations. Automation has played a key role in smoothing the transition from commissioning to operations, and will continue to be a core value applied to ASKAP operations. Based on our experience, automation acts as a catalyst for improving all aspects of system behaviour, identifying issues and fixing where possible or building robust work-around procedures otherwise. This is, especially, important for a telescope designed to spend a majority of its time conducting systematic surveys.

We have found with ASKAP that effective ongoing communications between the design, construction, operations, science, and management teams is critical to achieving the goal of delivering a productive radio telescope, and expect this to be one of the biggest challenges to solve for the next generation of large-scale science facilities to come.

Acknowledgements. The Australian SKA Pathfinder is part of the Australia Telescope National Facility which is managed by CSIRO. Operation of ASKAP is funded by the Australian Government with support from the National Collaborative Research Infrastructure Strategy. ASKAP uses the resources of the Pawsey Supercomputing Centre. Establishment of ASKAP, the Murchison Radio-Astronomy Observatory, and the Pawsey Supercomputing Centre are initiatives of the Australian Government, with support from the Government of Western Australia and the Science and Industry Endowment Fund. We acknowledge the Wajarri Yamatji as the traditional owners of the Observatory site.

Author Contributions. CSIRO’s ASKAP telescope was built and brought into operation by a large team working over a number of years. The authorship of this paper and its ordering reflects the effort invested by team members in preparing the paper, in addition to their contributions towards the construction and commissioning of ASKAP itself.

References

- ACMA 2014, Radiocommunications assignment and licensing instruction MS 32, Coordination of Apparatus Licensed Services Within the Mid West Radio Quiet Zone Western Australia, <https://www.acma.gov.au/publications/rali-ms32-mid-west-radio-quiet-zone>
- Abeywickrema, S., Allen, G., Ardern, K., Schinckel, A., Leitch, A., Wilson, C., & Beresford, R. 2013, in 2013 Asia-Pacific Symposium on Electromagnetic Compatibility (APEMC), 1–4, doi: [10.1109/APEMC.2013.7360649](https://doi.org/10.1109/APEMC.2013.7360649)
- Allison, J. R., Sadler, E. M., & Whiting, M. T. 2012, *PASA*, 29, 221
- Allison, J. R., et al. 2015, *MNRAS*, 453, 1249
- Allison, J. R., et al. 2017, *MNRAS*, 465, 4450
- Allison, J. R., et al. 2020, *MNRAS*, 494, 3627
- Anderson, C., et al. 2018, *Galaxies*, 6, 127
- Applebaum, S. P. 1976, *IEEE TAP*, 24, 585
- Bannister, K., Zackay, B., Qiu, H., James, C., & Shannon, R. 2019a, FREDDA: A fast, real-time engine for de-dispersing amplitudes, astrophysics source code library, record [ascl:1906.003](https://arxiv.org/abs/1906.003)
- Bannister, K. W., et al. 2019b, *Sci*, 365, 565
- Beresford, R., & Bunton, J. 2013, in 2013 Asia-Pacific Symposium on Electromagnetic Compatibility (APEMC), 1–4, doi: [10.1109/APEMC.2013.7360638](https://doi.org/10.1109/APEMC.2013.7360638)
- Beresford, R., Cheng, W., & Roberts, P. 2017a, in 2017 XXXIInd General Assembly and Scientific Symposium of the International Union of Radio Science (URSI GASS), 1–4, doi: [10.23919/URSIGASS.2017.8105422](https://doi.org/10.23919/URSIGASS.2017.8105422)

- Beresford, R., Ferris, D., Cheng, W., Hampson, G., Bunton, J., Chippendale, A., & Kanapathippillai, J. 2017c, in 2017 International Topical Meeting on Microwave Photonics (MWP), 1–4, doi: [10.1109/MWP.2017.8168641](https://doi.org/10.1109/MWP.2017.8168641)
- Beresford, R., & Li, L. 2013, in 2013 Asia-Pacific Symposium on Electromagnetic Compatibility (APEMC), 1–4, doi: [10.1109/APEMC.2013.7360637](https://doi.org/10.1109/APEMC.2013.7360637)
- Beresford, R., *et al.* 2017b, in 2017 XXXIIInd General Assembly and Scientific Symposium of the International Union of Radio Science (URSI GASS), 1–4, doi: [10.23919/URSIGASS.2017.8105423](https://doi.org/10.23919/URSIGASS.2017.8105423)
- Bhandari, S., *et al.* 2018, *MNRAS*, 478, 1784
- Bhatnagar, S., Cornwell, T. J., Golap, K., & Uson, J. M. 2008, *A&A*, 487, 419
- Bhatnagar, S., Rau, U., & Golap, K. 2013, *ApJ*, 770, 91
- Black, R. A., Jeffs, B. D., Warnick, K. F., Hellbourg, G., & Chippendale, A. 2015, in 2015 IEEE Signal Processing and Signal Processing Education Workshop (SP/SPE), 261, doi: [10.1109/DSP-SPE.2015.7369563](https://doi.org/10.1109/DSP-SPE.2015.7369563)
- Bowman, J. D., Rogers, A. E. E., Monsalve, R. A., Mozdzen, T. J., & Mahesh, N. 2018, *Nature*, 555, 67
- Briggs, D. S. 1995, PhD thesis, New Mexico Institute of Mining and Technology, <http://www.aoc.nrao.edu/dissertations/dbriggs/>
- Brown, A. J., *et al.* 2014, in 2014 International Conference on Electromagnetics in Advanced Applications (ICEAA), 268–271, doi: [10.1109/ICEAA.2014.6903860](https://doi.org/10.1109/ICEAA.2014.6903860)
- Bunton, J. D. 2003a, SKA Memo 40, Figure of merit for SKA survey speed. https://www.skatelescope.org/uploaded/51368_40_memo_Bunton.pdf
- Bunton, J. D. 2003b, ALMA Memo 447, Multi-resolution FX correlator. <http://library.nrao.edu/public/memos/alma/main/memo447.pdf>
- Bunton, J. D., & Hay, S. G. 2010, in 2010 International Conference on Electromagnetics in Advanced Applications, 728, doi: [10.1109/ICEAA.2010.5651120](https://doi.org/10.1109/ICEAA.2010.5651120)
- Chapman, J. M., Dempsey, J., Miller, D., Heywood, I., Pritchard, J., Sangster, E., Whiting, M., & Dart, M. 2017, in Astronomical Society of the Pacific Conference Series, Vol. 512, Astronomical Data Analysis Software and Systems XXV, ed. N. P. F. Lorente, K. Shortridge, & R. Wayth, 73, <http://aspbooks.org/custom/publications/paper/512-0073.html>
- Chippendale, A., O'Sullivan, J., Reynolds, J., Gough, R., Hayman, D., & Hay, S. 2010, in 2010 IEEE International Symposium on Phased Array Systems and Technology, 648, doi: [10.1109/ARRAY.2010.5613298](https://doi.org/10.1109/ARRAY.2010.5613298)
- Chippendale, A., & Wormnes, K. 2013, in 2013 Asia-Pacific Symposium on Electromagnetic Compatibility (APEMC), 1–4, doi: [10.1109/APEMC.2013.7360640](https://doi.org/10.1109/APEMC.2013.7360640)
- Chippendale, A. P., & Anderson, C. 2019, ACES Memo 19, On-dish calibration of XY phase for ASKAP phased array feeds. CSIRO, <http://www.atnf.csiro.au/projects/askap/ACES-memos>
- Chippendale, A. P., Beresford, R. J., Deng, X., Leach, M., Reynolds, J. E., Kramer, M., & Tzioumis, T. 2016b, in 2016 International Conference on Electromagnetics in Advanced Applications (ICEAA), 909, doi: [10.1109/ICEAA.2016.7731550](https://doi.org/10.1109/ICEAA.2016.7731550)
- Chippendale, A. P., Button, C., & Lourenço, L. 2018, ACES Memo 18, Measuring ASKAP's on-dish calibration signal level and its impact on beam sensitivity. CSIRO, <http://www.atnf.csiro.au/projects/askap/ACES-memos>
- Chippendale, A. P., Colegate, T. M., & O'Sullivan, J. D. 2007, SKA Memo 92, SKAcost: A Tool for SKA Cost and Performance Estimation, https://www.skatelescope.org/uploaded/33760_memo_92.pdf
- Chippendale, A. P., Hayman, D. B., & Hay, S. G. 2014, *PASA*, 31, e019
- Chippendale, A. P., & Hellbourg, G. 2017, in 2017 International Conference on Electromagnetics in Advanced Applications (ICEAA), 948, doi: [10.1109/ICEAA.2017.8065413](https://doi.org/10.1109/ICEAA.2017.8065413)
- Chippendale, A. P., *et al.* 2015, in 2015 International Symposium on Antennas and Propagation (ISAP), 1–4 ([arXiv:1509.05489](https://arxiv.org/abs/1509.05489))
- Chippendale, A. P., *et al.* 2016a, in 2016 10th European Conference on Antennas and Propagation (EuCAP), 1–5, doi: [10.1109/EuCAP.2016.7481741](https://doi.org/10.1109/EuCAP.2016.7481741)
- Cornwell, T. J., Golap, K., & Bhatnagar, S. 2005, in Astronomical Society of the Pacific Conference Series, Vol. 347, Astronomical Data Analysis Software and Systems XIV, ed. P. Shopbell, M. Britton, & R. Ebert, 86, <http://aspbooks.org/custom/publications/paper/347-0086.html>
- Cornwell, T., Humphreys, B., Lenc, E., Voronkov, M., Whiting, M., Mitchell, D., Ord, S., & Collins, D. 2016, Technical report, ASKAP science processing. CSIRO, <https://www.atnf.csiro.au/projects/askap/ASKAP-SW-0020.pdf>
- CSIRO 2008, Public specification, ASKAP antenna specification and operating parameters. CSIRO, https://www.atnf.csiro.au/projects/askap/ASKAP_Antenna_public_specification_Nov08_v0.0.pdf
- DeBoer, D. R., *et al.* 2009, *Proc. IEEE*, 97, 1507
- Deller, A. T., *et al.* 2011, *PASP*, 123, 275
- Deng, X., *et al.* 2017, *PASA*, 34, e026
- Dunning, A., Bowen, M. A., Hayman, D. B., Kanapathippillai, J., Kanoniuk, H., Shaw, R. D., & Severs, S. 2016, in 2016 46th European Microwave Conference (EuMC), 1568, doi: [10.1109/EuMC.2016.7824657](https://doi.org/10.1109/EuMC.2016.7824657)
- Dunning, A., *et al.* 2019, in PAF Workshop 2019
- Elagali, A., *et al.* 2019, *MNRAS*, 487, 2797
- Elmer, M., Jeffs, B. D., & Warnick, K. F. 2014, *IEEE TAP*, 62, 6067
- Feng, Q.-Q., Li, Z.-C., & Li, G.-Y. 2010, in *Proc. SPIE*, 782022, doi: [10.1117/12.866765](https://doi.org/10.1117/12.866765)
- For, B. Q., *et al.* 2019, *MNRAS*, 489, 5723
- Frost, O. L. 1972, *Proc. IEEE*, 60, 926
- Glowacki, M., *et al.* 2019, *MNRAS*, 489, 4926
- Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, *ApJ*, 622, 759
- Gupta, N., Johnston, S., Feain, I., & Cornwell, T. 2008, Memo 1, The initial array configuration for ASKAP. CSIRO, <https://www.atnf.csiro.au/projects/askap/newdocs/configs-3.pdf>
- Guzman, J. C., & Humphreys, B. 2010, in *Proc. SPIE*, 77401J, doi: [10.1117/12.856962](https://doi.org/10.1117/12.856962)
- Guzman, J., *et al.* 2019, ASKAPsoft: ASKAP science data processor software, astrophysics source code library, record [ascl:1912.003](https://arxiv.org/abs/1912.003)
- Hampson, G., *et al.* 2012, 807, doi: [10.1109/ICEAA.2012.6328742](https://doi.org/10.1109/ICEAA.2012.6328742)
- Hampson, G. A., Brown, A., Bunton, J. D., Neuhold, S., Chekkala, R., Bateman, T., & Tuthill, J. 2014, in Proc. XXXIth URSI General Assembly and Scientific Symposium, 1–4, doi: [10.1109/URSIGASS.2014.6930062](https://doi.org/10.1109/URSIGASS.2014.6930062)
- Harvey-Smith, L., *et al.* 2016, *MNRAS*, 460, 2180
- Hay, S. 2010, *IJMOT*, 5, 375
- Hay, S. G., & O'Sullivan, J. D. 2008, *RS*, 43
- Hay, S. G., O'Sullivan, J. D., Kot, J. S., Granet, C., Grancea, A., Forsyth, A. R., & Hayman, D. H. 2007, in The Second European Conference on Antennas and Propagation, EuCAP 2007, 1–5, doi: [10.1049/ic.2007.0899](https://doi.org/10.1049/ic.2007.0899)
- Hay, S. G., O'Sullivan, J. D., & Mitra, R. 2011, *IEEE TAP*, 59, 1828
- Hayman, D. B., Chippendale, A., Qiao, R., Bunton, J. D., Beresford, R. J., Roberts, P., & Axtens, P. 2010, in 2010 International Conference on Electromagnetics in Advanced Applications, 418, doi: [10.1109/ICEAA.2010.5653177](https://doi.org/10.1109/ICEAA.2010.5653177)
- Hellbourg, G., Bannister, K., & Hotan, A. 2016, in 2016 Radio Frequency Interference (RFI). IEEE Conference Series, 37, doi: [10.1109/RFINT.2016.7833528](https://doi.org/10.1109/RFINT.2016.7833528)
- Hellbourg, G., Chippendale, A. P., Kesteven, M. J., & Jeffs, B. D. 2014, in 2014 IEEE Global Conference on Signal and Information Processing (GlobalSIP), 1286, doi: [10.1109/GlobalSIP.2014.7032330](https://doi.org/10.1109/GlobalSIP.2014.7032330)
- Heywood, I., *et al.* 2016, *MNRAS*, 457, 4160
- Heywood, I., *et al.* 2020, *MNRAS*, 494, 5018
- Hobbs, G., *et al.* 2016, *MNRAS*, 456, 3948
- Hobbs, G., *et al.* 2020, *PASA*, 37, e012
- Hotan, A. W. 2016, ACES Memo 11, Holographic Measurement of ASKAP Primary Beams. CSIRO, <http://www.atnf.csiro.au/projects/askap/ACES-memos>
- Hotan, A. W., *et al.* 2014, *PASA*, 3, e041
- Hoyle, S. A., & Mirtschin, P. L. 2015, in *Proc. 15th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALPCS'15)*, 660, doi: [10.18429/JACoW-ICALPCS2015-WEM301](https://doi.org/10.18429/JACoW-ICALPCS2015-WEM301)
- Huynh, M., Dempsey, J., Whiting, M. T., & Ophel, M. 2020, in Astronomical Society of the Pacific Conference Series, Vol. 522, Astronomical Data Analysis Software and Systems XXVII, ed. P. Ballester, J.

- Ibsen, M. Solar, & K. Shortridge, 263, <http://aspbooks.org/custom/publications/paper/522-0263.html>
- Indermuehle, B., Harvey-Smith, L., Wilson, C., & Chow, K. 2016, in 2016 Radio Frequency Interference (RFI) workshop. IEEE Conference Series, 43, doi: [10.1109/RFINT.2016.7833529](https://doi.org/10.1109/RFINT.2016.7833529)
- Indermuehle, B. T., Harvey-Smith, L., Marquarding, M., & Reynolds, J. 2018a, in Proc. SPIE, 107041W, doi: [10.1117/12.2311926](https://doi.org/10.1117/12.2311926)
- Indermuehle B. T., Harvey-Smith L., Marquarding M., & Reynolds J., 2018b, in Proc. SPIE, 107042S, doi: [10.1117/12.2311917](https://doi.org/10.1117/12.2311917)
- Jeffs, B. D., Warnick, K. F., Landon, J., Waldron, J., Jones, D., Fisher, J. R., & Norrod, R. D. 2008, IEEE JSTSP, 2, 635
- Johnston, S., & Grey, A. 2006, SKA Memo 72, Surveys with the xNTD and CLAR, https://www.skatelescope.org/uploaded/24448_72_Johnston.pdf
- Johnston, S., et al. 2007, PASA, 24, 174
- Joseph, T. D., et al. 2019, MNRAS, 490, 1202
- Kadler, M., et al. 2016, Natur Phys., 12, 807
- Kemball, A., & Wieringa, M. 2000, AIPS++ Memo 229, MeasurementSet definition version 2.0. NRAO, <https://casa.nrao.edu/Memos/229.html>
- Kleiner, D., et al. 2019, MNRAS, 488, 5352
- Kooistra, E., Hampson, G. A., Gunst, A. W., Bunton, J. D., Schoonderbeek, G. W., & Brown, A. 2017, in 2017 XXXIInd General Assembly and Scientific Symposium of the International Union of Radio Science (URSI GASS), 1–4, doi: [10.23919/URSIGASS.2017.8104976](https://doi.org/10.23919/URSIGASS.2017.8104976)
- Leahy, D. A., et al. 2019, PASA, 36, e024
- Lee-Waddell, K., et al. 2019, MNRAS, 487, 5248
- Lo, Y. T., Lee, S. W., & Lee, Q. H. 1966, Proc. IEEE, 54, 1033
- McClure-Griffiths, N. M., et al. 2018, NatAS, 2, 901
- McConnell, D. 2016, ACES Memo 10, Field structure at the focus of paraboloidal reflectors and comparison with ASKAP beams. CSIRO, <http://www.atnf.csiro.au/projects/askap/ACES-memos>
- McConnell, D. 2017a, ACES Memo 14, Estimating ASAKP beam to beam correlation. CSIRO, <http://www.atnf.csiro.au/projects/askap/ACES-memos>
- McConnell, D. 2017b, ACES Memo 15, Observing with ASKAP: Optimisation for survey speed. CSIRO, <http://www.atnf.csiro.au/projects/askap/ACES-memos>
- McConnell, D., et al. 2016, PASA, 33, e042
- Moss, V. A., et al. 2017, MNRAS, 471, 2952
- Pence, W. D., Chiappetti, L., Page, C. G., Shaw, R. A., & Stobie, E. 2010, A&A, 524, A42
- Rau, U. 2010, PhD thesis, New Mexico Institute of Mining and Technology, http://www.aoc.nrao.edu/rurvashi/DataFiles/UrvashiRV_PhDThesis.pdf
- Rau, U., & Cornwell, T. J. 2011, A&A, 532, A71
- Reynolds, J. E. 1994, Technical Report 39.3/040, A revised flux scale for the AT compact array. Australia Telescope National Facility, <https://www.atnf.csiro.au/observers/memos/d96783-1.pdf>
- Reynolds, T. N., Staveley-Smith, L., Rhee, J., Westmeier, T., Chippendale, A. P., Deng, X., Ekers, R. D., & Kramer, M., 2017, PASA, 34, e051
- Reynolds, T. N., et al. 2019, MNRAS, 482, 3591
- Rice, S. O. 1982, Proc. IEEE, 70, 692
- Sarkissian, J. M., Reynolds, J. E., Hobbs, G., & Harvey-Smith, L. 2017, PASA, 34, e027
- Sault, B. 2014, ACES Memo 2, Initial characterization of BETA polarimetric response. CSIRO, <http://www.atnf.csiro.au/projects/askap/ACES-memos>
- Sault, B. 2015, ACES Memo 7, Widefield polarimetric considerations for ASKAP. CSIRO, <http://www.atnf.csiro.au/projects/askap/ACES-memos>
- Schinkel, A. E., et al. 2011, in Asia-Pacific Microwave Conference 2011, 1178
- Serra, P., et al. 2015a, MNRAS, 448, 1922
- Serra, P., et al. 2015b, MNRAS, 452, 2680
- Seymour, N., et al. 2020, PASA, 37, e013
- Shaw, R. D., & Hay, S. G. 2015, in 2015 9th European Conference on Antennas and Propagation (EuCAP), 1–4, <https://ieeexplore.ieee.org/document/7228365>
- Shaw, R. D., Hay, S. G., & Ranga, Y. 2012, in 2012 International Conference on Electromagnetics in Advanced Applications, 438, doi: [10.1109/ICEAA.2012.6328666](https://doi.org/10.1109/ICEAA.2012.6328666)
- Taylor, M. B. 2005, in Astronomical Society of the Pacific Conference Series, Vol. 347, Astronomical Data Analysis Software and Systems XIV, ed. P. Shopbell, M. Britton, & R. Ebert, 29
- The HDF Group 1997–2020, Hierarchical data format, version 5, <https://www.hdfgroup.org/HDF5/>
- Tingay, S. J., et al. 2013, PASA, 30, e007
- Tuthill, J., Hampson, G., Bunton, J., Brown, A., Neuhold, S., Bateman, T., de Souza, L., & Joseph, J. 2012, in 2012 International Conference on Electromagnetics in Advanced Applications, 1067, doi: [10.1109/ICEAA.2012.6328788](https://doi.org/10.1109/ICEAA.2012.6328788)
- Tuthill, J., Hampson, G., Bunton, J. D., Harris, F. J., Brown, A., Ferris, R., & Bateman, T. 2015, in 2015 IEEE Signal Processing and Signal Processing Education Workshop (SP/SPE), 255, doi: [10.1109/DSPSPE.2015.7369562](https://doi.org/10.1109/DSPSPE.2015.7369562)
- Van Veen, B. D., & Buckley, K. M. 1988, IEEE ASSPM, 5, 4
- Wayth, R. B., et al. 2018, PASA, 35
- Weinreb, S., & Shi, J. 2019, in PAF Workshop 2019, <https://events.mpifr-bonn.mpg.de/indico/event/108/session/17/contribution/27>
- Whiting, M. T. 2012, MNRAS, 421, 3242
- Whiting, M. T. 2020, in Astronomical Society of the Pacific Conference Series, Vol. 522, Astronomical Data Analysis Software and Systems XXVII, ed. P. Ballester, J. Ibsen, M. Solar, & K. Shortridge, 469, <http://aspbooks.org/custom/publications/paper/522-0469.html>
- Whiting, M., & Humphreys, B. 2012, PASA, 29, 371
- Wilson, C., Chow, K., Harvey-Smith, L., Indermuehle, B., Sokolowski, M., & Wayth, R. 2016, Proc. 2016 International Conference on Electromagnetics in Advanced Applications (ICEAA), 922, doi: [10.1109/ICEAA.2016.7731554](https://doi.org/10.1109/ICEAA.2016.7731554)
- Wilson, C., Storey, M., & Tzioumis, T. 2013, in 2013 Asia-Pacific Symposium on Electromagnetic Compatibility (APEMC), 1–4, doi: [10.1109/APEMC.2013.7360651](https://doi.org/10.1109/APEMC.2013.7360651)
- Wrobel, J. M., & Walker, R. C. 1999, in Astronomical Society of the Pacific Conference Series, Vol. 180, Synthesis Imaging in Radio Astronomy II, ed. G. B. Taylor, C. L. Carilli, & R. A. Perley, 171, <https://ui.adsabs.harvard.edu/abs/1999ASPC..180..171W>
- Wu, C., Wicenc, A., Pallot, D., & Checcucci, A. 2013, ExA, 36, 679