URSI GASS 2021, Rome, Italy, 28 August - 4 September 2021

ASKAP: From Commissioning to Operations

V. A. Moss*[(1)(2)], M.T. Whiting[(1)], A.W. Hotan[(1)], D. C. J. Bock[(1)] and K. W. Bannister[(1)]
(1) ATNF, CSIRO Astronomy and Space Science, PO Box 76, Epping, New South Wales 1710, Australia
(2) Sydney Institute for Astronomy, School of Physics, University of Sydney, Sydney, New South Wales 2006, Australia

Abstract

The Australian Square Kilometre Array Pathfinder (ASKAP) is an SKA precursor located in the Murchison region of Western Australia. This 36-antenna array utilises innovative phased-array-feeds (PAFs) to provide a wide field-of-view, dramatically increasing the survey capabilities of the telescope. It is currently in the process of transitioning from commissioning into full operations, which is the focus of the overview presented here.

1 Introduction

ASKAP is a radio telescope array designed for maximum survey speed through a combination of sensitivity, spatial frequency coverage and wide field of view, thanks to the next-generation wide-field capabilities offered by PAFs. As of early 2019, the full system including 36 dishes equipped with PAFs was online and ready for the final stages of commissioning with both engineering and observational tests carried out using the full array. Since then, we have been in the process of working towards full operations while continuing to improve and finalise the existing system.

2 The ASKAP Surveys

In full operations, ASKAP will spend approximately 75% of its time on specific survey projects aligned with science goals across a variety of astronomical fields. These surveys were originally designed and approved in 2010, covering continuum, spectral and transient science. The current list of active ASKAP surveys is: EMU (continuum, [1]), WALLABY (spectral, [2]), DINGO (spectral, deep, [2]) POSSUM (polarisation, [3]), CRAFT (fast radio bursts, [4]), VAST (transients, [5]), GASKAP (spectral, Galactic, [6]) and FLASH (spectral, absorption, [7]). The exact survey allocations are to be reassessed based on the outcomes of the current transition period, which includes carrying out a series of 100 hr Pilot Surveys (described below).

In addition to these, ASKAP has also carried out two Observatory Projects. Firstly, the Rapid ASKAP Continuum Survey (RACS), which is an all-sky survey observed within a few weeks of on-sky time (eventually across all three ASKAP filterbands). Data for RACS-LOW were released in late 2020 [8]. We have also partnered with the eROSITA team within the scope of the AAL/eROSITA Memorandum of Understanding to carry out the Survey With ASKAP of GAMA-09 + X-ray (SWAG-X, complementary to eFEDS) at 888 MHz and 1296 MHz, with the data designed to be as commensal for existing ASKAP surveys as possible.

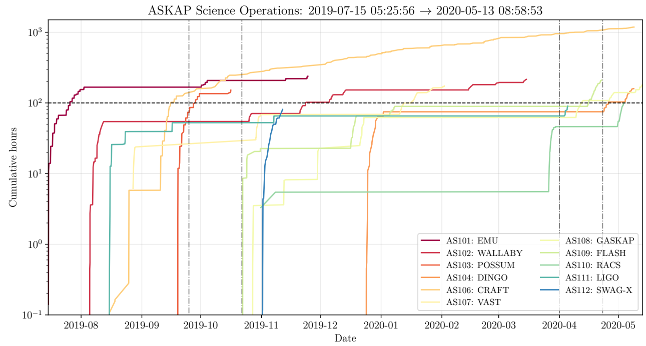

Figure 1. Cumulative histogram of hours observed for the different ASKAP SST projects over the course of Pilot Surveys Phase I, from July 2019 to May 2020.

3 Pilot Surveys Phase I

Each Survey Science Team (SST) was awarded 100 hr during Pilot Surveys Phase I and given the freedom to design their Pilot Survey to inform their future full survey plans. This gave the teams the chance to optimise parameters and ensure that the full survey will meet the science goals, as well as providing them with the data to justify their full survey requirements. This first round of Pilot Surveys (to be followed by Phase II starting early 2021) was thus not focused on addressing the questions of commensality or observational efficiency in any great detail. We expect the second round of Pilot Surveys to be more focused on examining the maximum amount of commensality we can obtain from these different science requirements.

We officially began the Pilot Survey period on 15th July 2019, starting with observations for the EMU Pilot. We then continued observing data for each of the different surveys where possible, completing Phase I on 14th May 2020. In total for Phase I, we spent 2216 hr successfully observing data for the Pilot Surveys (2696 hr including all observations, plus an additional 1184 hr running filler CRAFT), with a total observing efficiency fraction per day of 0.50, a total successful observing efficiency fraction per day of 0.37 and a relative success fraction of 0.74. The distribution of hours amongst the surveys is shown in Figure 1.

4 A current view of the ASKAP data flow

While we are in this transition period, we are using the experiences and lessons learned to inform our plans for full operations of ASKAP. The snapshot view provided here gives some insight into the current operations of ASKAP, but we expect this data flow model to adapt and develop over the course of the coming months as we learn more about how best to operate the full system while maximising efficiency and science quality. This developing operations model of ASKAP is new to CSIRO Astronomy and Space Science when compared to operation of previous instruments within the Australia Telescope National Facility. The amount of data generated by a multi-beam, multi-antenna interferometer is too great for easy transport or processing by individual astronomers and the complexity of scheduling multiple large-scale survey projects means that observing, calibration and imaging are all performed by the observatory (ideally, as automatically and autonomously as possible). This presents new challenges, especially in terms of data throughput and quality control. Internally, we have two development teams focused on ASKAP (ASKAP-X Earth and ASKAP-X Air), dedicated to implementing new features necessary for long-term ASKAP operations as well as increasing robustness, stability, reliability and automation.

4.1 Coordination

ASKAP operations depends on careful coordination between science operations, maintenance, development and processing. We currently carry out a significant amount of this coordination on a weekly basis with a 30 min scheduling meeting bringing together key representatives from the different groups requiring access to the system, to gain an overview of what is needed for the week. This has been a very effective way of balancing priorities and ensuring everyone is kept in the loop. We also generally work according to a structured weekly schedule, where the priorities for different activities (development, maintenance, science) vary throughout the week. Daily communication is primarily (and effectively) carried out via Mattermost online chat. Over time, we expect development requirements to decrease, though we will always need to balance science observations against both preventative and reactive maintenance plus development access for continuous improvement and future upgrades into the era of full operations.

4.2 Specification

Pilot Surveys Phase I served as a useful experiment in gathering the different requirements from SSTs for their science and ensuring we can perform the necessary observations on the system. Though we gained much experience during the commissioning phase, the transition to operational Pilot Surveys has required more regimented structure of how surveys are specified and where the responsibilities divide between SSTs and the observatory. Specification has been further developed to account for various observational modes: continuum, spectral-line, interleaving (nightly cadence), interleaving (short cadence), zoom modes, snapshot observations and LST-matched observations, to name a few. This has been a useful and instructive learning process that also helped define clear survey constraints going forward. The need to specify observations in a more automated way has involved developing additional scheduling tools which are capable of automatically filling a given time period with appropriate observations and sending those to the system with minimal human intervention. As of late 2020, these structures were extended into a dynamic autonomous scheduler (SAURON: Scheduling Autonomously Under Reactive Observational Needs, to be described in more detail elsewhere). SAURON is designed to act based on system status, environmental conditions, the existing observation pool and survey constraints, and will be further refined and improved over the course of Pilot Surveys Phase II. While it is early in our adoption of autonomous scheduling, SAURON has already shown a dramatic improvement with an increase in efficiency of $\sim 30\%$ and in success rate of $\sim 25\%$.

4.3 Initialisation

Previously, the system required a large degree of human intervention in the initialisation process to prepare it for a given set of science observations. Tasks associated with initialisation included manually putting the system in the right frequency configuration, ensuring everything is online and working normally, updating existing beamweights using an on-dish calibrator (ODC), calibrating delays on a bright compact source and performing a test observation to verify system functionality. Depending on the system stability on a given day, this process would take as little as 30 min but generally between 1-2 hr, or longer if something went significantly wrong. Since each system setup required human intervention, this limited our ability to carry out different modes of science observation in close prox-

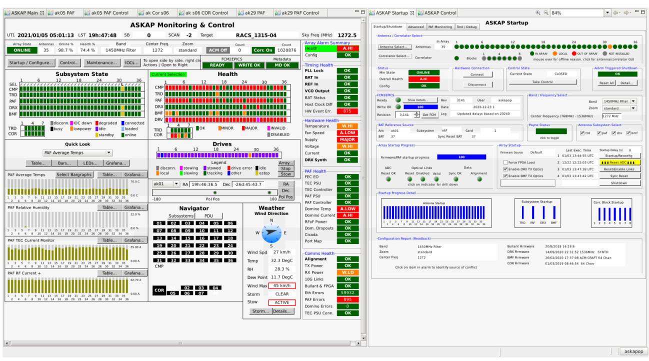

Figure 2. Screenshot of the main ASKAP CSS monitoring screen, which gives a quick overview of the entire system state while also allowing drill-down for debugging and manual control of the system if needed.

imity. As part of ASKAP-X Earth, significant effort has been invested to improve the system stability and increase the amount of automation in the setup process to facilitate much more automated system startups, with parameters specified from the parset to determine what kind of setup steps are needed. Many of these improvements have been successfully in place since late 2020, and are key to the increase in observational success rate and overall efficiency.

4.4 Monitoring

Monitoring is primarily done through four avenues: 1) a graphical user interface designed for engineering development, built on EPICS Control System Studio (CSS), 2) Grafana + InfluxDB web-based time-series monitoring, 3) real-time visibility visualisations (VIS, SPD), and 4) a web-based Observation Management Portal (OMP). CSS gives us the most direct insight into the hardware state (see Figure 2), allowing easy tracking of the PAFs, digitisers, drives, correlator, beamformers and more. From this console it is also possible to restart subsystems or deal with certain kinds of errors, though if an error is particularly severe then it may be impossible to resolve without on-site intervention. The Grafana monitoring accesses a database containing the history of many monitoring points - for operations, we primarily use it to track scan status, weather behaviour, astronomical data ingest and (as needed) hardware diagnostics. The visibility visualisations give us live insight into the data being recorded, based on existing visualisations originally developed for the Australia Telescope Compact Array, and as part of ASKAP-X we hope to see a new, more robust system for this purpose to give us more flexibility in visualising (big) data. Finally, we make use of the OMP (which is a visual interface to the scheduling block database) to track the history of observations as well as note the current status of observations, and to dig into the logs if an observation fails for some reason. The OMP is web-accessible with an account, while the other three tools are accessed via a VNC session, and between the four tools we have good visibility into the operational state of the system. The long-term plans for monitoring are in flux, but will be driven by the same goals of having the system be as automated and autonomous as possible with human effort only as necessary.

4.5 Diagnosis

ASKAP, like many modern telescopes, is an extremely complex system. Additional challenges we face are the relatively recent development of PAF technology for radio astronomy purposes, which we are still in the process of understanding and optimising, and the immense size of data that we have to work with. Diagnostics for assessing ASKAP must therefore come in at different places in the system, starting with the hardware (e.g. PAF domino health), then moving towards raw data prior to any processing (e.g. SEFD measurements, raw data diagnostics) and then ending with diagnostics related to processed data (e.g. bandpass calibration plots). We are in the process of exploring various diagnostics and developing new ones as necessary, as well as ensuring that we are seeing consistent results across the system. For example, a newly developed tool to track health on the PAF means it is possible to check bad elements on particular antennas against the raw data for beams formed using that element, and from this we have seen that issues with RX power have a worse impact on raw data than LNA variability. Due to the scale and complexity of the data, it has generally been necessary to develop diagnostics targeting specific questions and slice through the data to directly answer those questions, as it is otherwise difficult to see the big picture. That noted, in the quest to obtain a truly broad overview of the raw data, the "everything" plot was developed. This plot uses the autocorrelation waterfall data for each antenna and beam to really highlight issues which are across an entire antenna vs. affecting certain beams on a particular antenna (e.g. bad PAF elements). An example of this overall system state in YY polarisation is given in Figure 3. Broad trends for antennas are clearly visible, as well as localised issues. Based on existing diagnostics, we are investigating better ways to automatically capture and convey system state.

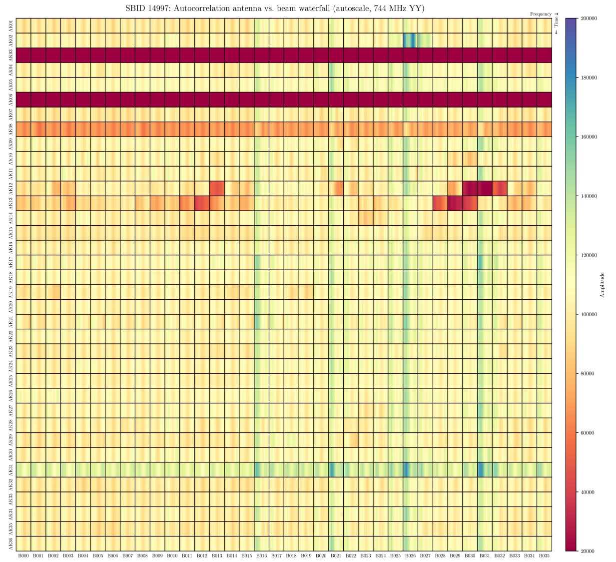

Figure 3. Example of YY-polarisation "everything" plot for ASKAP raw data, a useful diagnostic tool developed to visually highlight issues across entire antennas (e.g. ak03, ak06, ak08, ak31) or affecting particular antenna-beam combinations (e.g. ak12 B013/B030/B031, ak13 B012/B029). Each box is an auto-correlation waterfall plot (time y-axis, frequency x-axis) with beams increasing horizontally across and antennas vertically down.

4.6 Processing

ASKAP raw visibility data is transferred in real time to the Pawsey Supercomputing Centre located in Perth, Western Australia. As of late 2020, it is written to a 500 TB ingest cluster (RUBY), and then, at the end of the observation, is copied automatically to a 4 PB disk consisting

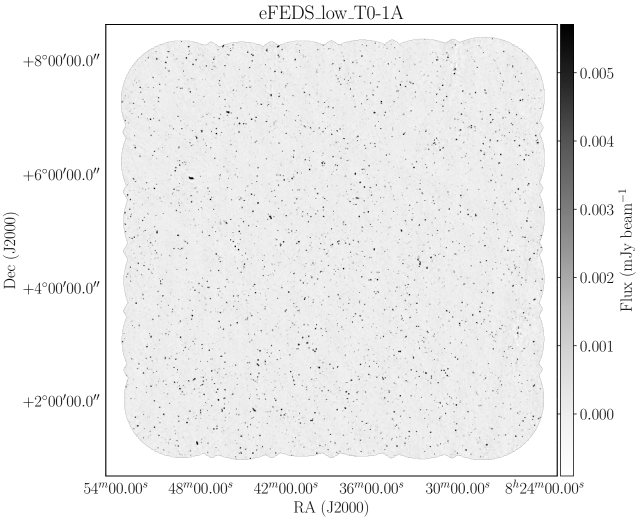

Figure 4. A continuum image from a single ASKAP observation at 888 MHz, produced by the ASKAPsoft pipeline. Note the field of view formed by the multiple PAF beams.

of two partitions (PAYNE and SCOTT). The recent switch to RUBY serves to isolate the data-intensive processing from ingest as much as possible, to avoid interruptions to the real-time ingest. Raw data is turned into science-ready data products (calibrated data, images, and catalogues) by the ASKAPsoft processing pipeline [9], making use of Pawsey's GALAXY Cray XC30. ASKAPsoft provides a suite of tasks to perform calibration, imaging, source-finding and analysis, plus specialised tasks relevant for ASKAP. It has been designed with high-performance processing built in, and can provide distributed processing to handle the very large datasets (the input data for a spectral-line observation can be > 30 TB).

The pipeline provides a streamlined workflow that runs the large number of jobs required, interacting with the Slurm queue manager to ensure appropriate dependencies across the workflow. The typical procedure involves: bandpass calibration, flagging, averaging to coarse resolution, continuum imaging and self-calibration, continuum subtraction from the full-spectral-resolution data, spectral imaging, mosaicking to form full-field images, and source extraction. Upon completion, the data (images, catalogues and calibrated continuum visibilities) are archived in the CSIRO ASKAP Science Data Archive (CASDA), where they are publicly available upon release.

Another key goal of Pilot Surveys is to better link the processes of observation and processing, which often take place largely in isolation. By doing so, we move towards a full operational workflow from specification to data release.

5 Summary

We have presented here an overview of current ASKAP operations, which have successfully led to high-quality science observations as part of Pilot Surveys Phase I. Over the coming months, we will continue to develop and improve all aspects of this model as we work towards full ASKAP operations. In particular, we have seen great developmental progress recently leading to useful new functionality, improvements in system robustness and automation, and effective autonomous system management, all of which are critical for the smooth, efficient and successful running of the full ASKAP surveys in the coming years.

6 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 people as the traditional owners of the Observatory site.

References

[1] R. P. Norris et al., "EMU: Evolutionary Map of the Universe" *PASA*, **28**, 3, January 2013, pp. 215–248

[2] A. R. Duffy et al., "Predictions for ASKAP neutral hydrogen surveys" *MNRAS*, **426**, 4, November 2012, pp. 3385–3402

[3] B. M. Gaensler et al., "Survey Science with ASKAP: Polarization Sky Survey of the Universe's Magnetism (POSSUM)", *Bulletin of the AAS*, **AAS Meeting 215, Vol 42**, January 2010, pp. 515

[4] J-P. Macquart et al., "The Commensal Real-Time ASKAP Fast-Transients (CRAFT) Survey", *PASA*, **27**, January 2010, pp. 272—282

[5] T. Murphy et al., "VAST: An ASKAP Survey for Variables and Slow Transients", *PASA*, **30**, February 2013

[6] J. M. Dickey et al., "GASKAP-The Galactic ASKAP Survey", *PASA*, **30**, January 2013

[7] J. R. Allison et al., "Tracing the neutral gas environments of young radio AGN with ASKAP", *Astronomische Nachrichten*, **337**, February 2016

[8] McConnell, D., Hale, C. L., Lenc, E., et al. "The Rapid ASKAP Continuum Survey I: Design and first results", *PASA*, **37**, December 2020, e048

[9] M. Whiting, M. Voronkov, D. Mitchell et al., "Early Science Pipelines for ASKAP", *ADASS XXV*, Ed. Lorente, N. P. F. and Shortridge, K. and Wayth, R., **512**, 431, 2017