CSIRO ASTRONOMY AND SPACE SCIENCE
www.csiro.au

ASKAP Commissioning Update, August 2018

In this issue, we report the successful integration of 12 additional antennas into ASKAP, for a total of 28. Since this exceeds the requirements for array release 3, we also discuss how a larger array will change the format of early science phase two.

Towards completion of ASKAP

ASKAP's commissioning plan involved sequential array releases of increasing scale. Each release was associated with minimum requirements, including a number of antennas and certain key capabilities or features.

We are currently working to deliver array release three, which includes a new delay tracking system and enhanced frequency resolution modes, along with a modest increase in the number of antennas to a total of 18.

Strong progress on hardware installation and trouble-free integration work over the last two weeks has allowed us to expand the size of the array ahead of schedule.

Antenna integration activities

After fit-out, each ASKAP antenna goes through an acceptance test process involving checks of signal integrity and stability. During early science observations, we have been conducting these acceptance tests on new antennas using an isolated control system environment known as the commissioning array.

Transferral of antennas from the commissioning array to the active array involves integration into the delay tracking model, which needs precise knowledge of each antenna's position. Although geographic survey data provides an initial estimate of this, it is not sufficiently accurate for astronomical delay tracking. Millimetre precision is achieved by fitting phase trends obtained during lengthy observations of calibrator sources. This is an iterative process that can take roughly a week.

To make the most efficient use of telescope time, it is best to integrate several antennas in one batch, as the required observations are common to all (and can also be used to refine the positions of existing antennas).

Progress update and scaling considerations

During the past two weeks, we attempted to integrate all antennas that had passed their acceptance tests into the array and were pleased to encounter no major problems.

Of particular note is that the new set of antennas includes 5 of the 6 outermost ASKAP stations, providing 6 km baselines. Deployment of the new delay tracking system was a prerequisite for inclusion of the outermost antennas and was completed in time for this expansion.

One factor motivating early expansion was to provoke potential future issues of scale well in advance of array release four (with all 36 antennas).

Operating with 28 antennas did cause some packet loss in the control system network (with a larger number of hardware drivers running on each server) but these were quickly mitigated by redistribution of jobs and staggering of network access requests.

We have demonstrated stable correlator data capture at a significantly increased rate (in excess of 1 GB per second for the first time) from the expanded array. This required a planned split of the frequency axis into four different files for each beam, distributing load over additional processing nodes at Pawsey.

As a result, we have decided to continue with 28 antennas for array release three instead of scaling back to 18. We have an opportunity to conduct early science observations at nearly full scale and this will be extremely valuable when planning pilot surveys for 2019. However, the larger array will place increased strain on limited disk resources and will therefore require a more structured approach to early science phase two.

Plans for early science phase two

With 28 antennas, 36 beams, 4 polarisation products, 15552 frequency channels and 10-s cycles, ASKAP now produces just over 1 GB of visibility data every second. Typical 12-hour observations of a single field occupy 40 TB of disk space. At this rate, we would fill the entire 1 PB online storage buffer in two weeks of operation.

Early science phase two will be a small and focussed campaign designed to ensure that all science teams can make informed plans for pilot surveys in 2019. It is not intended to provide large amounts of science data on its own. We expect the main observing campaign to begin in November 2018 and this timescale implies that most observing should be aimed at preliminary verification and establishment of methods rather than publication.

Given that we will be processing early science data and pilot surveys in batch mode (not in real time), we must ensure rapid turn-around to maintain a reasonable on-sky duty cycle. This will be done through a combination of careful resource allocation and the provision of new operating modes, in addition to increased assistance from the ASKAP operations team.

It is important to point out that from November, we will cease to archive the full-resolution visibility data to tape. It will not be possible to restore data cleared from the online visibility buffer thereafter.

Allocation of disk space and CPU time

Future early science proposals will need to factor in disk usage requirements. We will ensure that each science team has access to some amount of disk space for raw visibility storage, but the expectation will be that once this is used, no new data will be recorded for that team until older observations have been removed.

It will therefore benefit science teams to use more efficient observing modes (described below) where possible, or to collaborate on observing specific fields that are of interest to several different teams.

Bulk processing will increasingly be conducted by the ASKAP operations team, with science team processing limited to small, exploratory jobs designed to determine optimal parameters. This will be reflected in a layered group-based priority structure. Over time we expect that CASDA will become the primary interface to ASKAP data.

Space-efficient operating modes

One way to reduce the strain on disk space is to offer additional observing modes that decimate data at the point of capture. These provide a way to obtain more hours of observing time, with a potential cost to data quality or elimination of alternative science goals.

One such mode is online continuum averaging. This has now been tested and will be offered as a standard mode for array release three. The averaging width can be changed, but will provide 1 MHz channels by default. Some basic flagging of pre-averaged data will be done by the ingest pipeline, but any low-level RFI will contaminate the entire 1 MHz channel it falls into.

It may also be possible to decrease the number of channels captured without reducing their frequency resolution – this will be significantly more difficult to configure (as the choice of channels is important and they are distributed across multiple processors). More information will be made available in future.

Removal of existing raw data

Due to ongoing processing, ASKAP's visibility storage buffer still contains data from the first round of early science observations nearly two years ago. The ASKAP operations team is working with the science teams to image these data and archive the results on CASDA, but the process has been slow due to processing difficulties and the need to iterate in order to determine optimal algorithms and parameters.

With array release three, we will need a clean slate. We will therefore be removing all visibility data from early science phase one on the 6th of November 2018 (to coincide with planned Pawsey maintenance activities). Science teams will have until then to complete processing (with the assistance of operations), average and store the data on scratch space at Pawsey, or apply for additional storage at Pawsey or other institutions.

We expect that data from array release three will greatly surpass pre-existing observations (see figure below). In most cases it will be desirable to re-observe rather than continue working on older data.

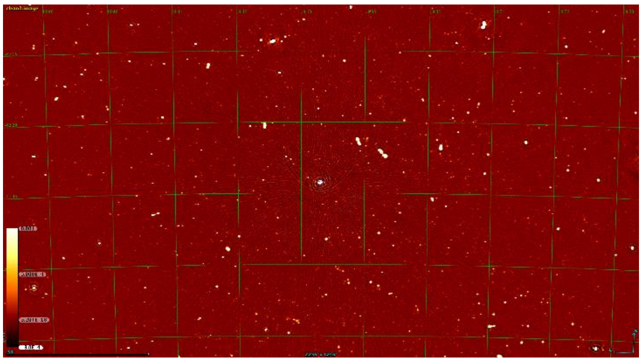

Image of the field surrounding calibrator source PKS B1934-638, made from 15 hours of data with 28 antennas and 288 MHz of bandwidth. Even with preliminary calibration, the RMS close to 1934 is 60 µJy/beam, giving a dynamic range of roughly 230,000:1.

Prior to this deadline, we will release resource allocation information for early science phase two and more details regarding how it will be accounted and managed.

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

FOR FURTHER INFORMATION
CSIRO Astronomy and Space Science
Aidan Hotan
t +61 8 6436 8543
e aidan.hotan@csiro.au
w www.csiro.au/askap