ASKAP update for August 2022

In this issue we discuss ongoing ASKAPsoft commissioning on Setonix, readiness for full surveys and the latest thoughts on a possible timeline.

Commissioning ASKAPsoft on Setonix

An image of supernova remnant G261.9+5.5 within the first ASKAP continuum field processed on Setonix has gained significant media attention over the last week. Although it only occupies a small part of ASKAP's field of view, the detailed structure of the supernova remnant demonstrates ASKAP's ability to detect extended emission and the large number of background radio sources demonstrates its sensitivity and spatial resolution.

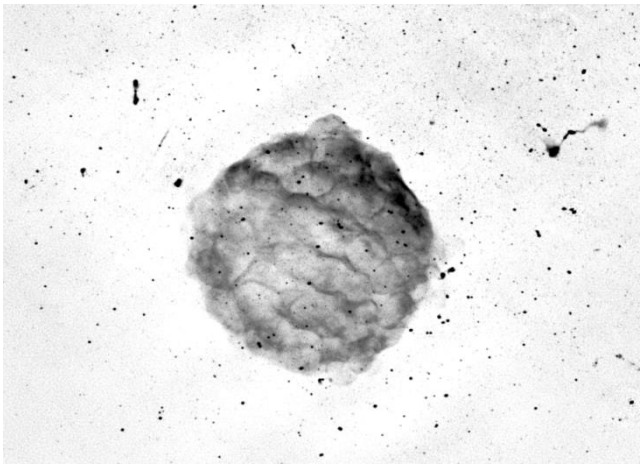

Figure 1: ASKAP Stokes I continuum image of supernova remnant G261.9+5.5 and its surroundings, extracted from a larger test field processed on Setonix. The observation was made at a centre frequency of 944 MHz with a duration of 10 hours. Image provided by Wasim Raja.

While impressive, the image represents only one of several modes required to deliver science ready data products. The pipeline infrastructure for all standard modes has been migrated to Setonix and updated to best utilise the new platform architecture. Tests of these modes are underway. Continuum imaging is significantly faster than on the previous supercomputer, Galaxy, and should easily meet the requirement of being able to process each observation in less time than it took to record the data. However, we have encountered some challenges running full-scale spectral line jobs.

The first was a problem with parallelisation via MPI which caused a high rate of failure in jobs using many MPI channels. This issue was resolved with a change to the MPI environment variables. Unfortunately, we are now seeing a different kind of failure mode that impacts jobs running on specific nodes. The cause of these failures was difficult to diagnose, but investigation with Pawsey staff revealed that some nodes are failing to clear shared memory on job completion, leading to "out of memory" errors. Other Setonix users have reported the same issue, so it is not specific to ASKAPsoft. Possible solutions range from patching the associated system libraries to upgrading the hardware interconnect, which would involve about a month of downtime. Pawsey staff are in discussions with the vendor to find the best solution. Meanwhile, the ASKAPsoft team is investigating changes to our code that may avoid the circumstances triggering the memory issue.

It currently takes at least 7 executions of the spectral line processing pipeline to obtain all required data products, with each subsequent run filling in some of the gaps left by previous failures. The continuum pipeline uses fewer nodes and completes successfully in one or two passes.

Although we are not yet ready to run spectral line jobs efficiently, we do have enough information to understand the performance of this mode. Using the fast solid-state scratch disks on Setonix, it is possible to run a WALLABY-style job from end-to-end in about 12 hours, processing all

15552 spectral channels. When the /askapbuffer disk is made available on Setonix, we can test the Pilot Surveys Phase II workflow (which did not make use of intermediate scratch space) and determine whether the pipeline needs to be modified to use Setonix's new solid-state scratch space routinely.

Although 12 hours is longer than the WALLABY observation time of 8 hours, we only need to process 7000 channels due to satellite interference in the lower half of the frequency band. This should reduce the execution time, getting closer to the 8-hour requirement.

During full survey operations, we also need to produce commensal data products. While the initial calibration and continuum imaging stage is common to all modes, spectral line imaging and production of polarisation cubes can be run in parallel on our Setonix allocation. It takes about 90 minutes to produces Stokes I, Q, U and V cubes with 1 MHz channels, then another 5 hours to extract spectra and perform RM synthesis. It may therefore be possible to produce POSSUM commensal products without increasing the overall execution time. If necessary, the extraction of source spectra could be made much more efficient with better parallelisation, but this requires development.

Survey readiness

Recent discussions around software features required for the full surveys highlighted some gaps between Pilot Surveys Phase II, the consolidation plan, and the expected full survey outputs. These primarily involve storing gridded (u,v) data to improve the quality of deep imaging surveys and storing spectral line visibilities from beams with bright sources to allow future post-processing. Although several Survey Science Teams requested these products in their RASSP documents, they are not part of the standard modes and have not been prioritised in the consolidation plan.

There is a desire among the SSTs to implement methods that may improve ASKAP's image quality, especially by de-convolving bright sources outside the boundaries of the individual beam images. Work is also required to improve the storage and retrieval of holography beam maps, since these have become a key part of the processing workflow.

These features are in addition to pipeline automation and checkpointing development that was already planned as an essential aspect of maintaining survey data throughput. It will be challenging to deliver all these features before the end of the year. We are working with the SSTs to ensure that priorities are understood.

Pilot Surveys Phase II Progress

Re-observations for POSSUM are underway, with the goal of determining whether we can meet off-axis polarisation leakage calibration requirements. Changes are also being made to trial source finding and spectrum extraction on each individual beam. These spectra would be used in a validation process where the properties of field sources prior to primary beam correction are compared with holography beam maps applied during mosaicking. This would avoid the need to store individual beam images as well as the mosaic.

Processing of FLASH re-observations designed to replace fields impacted by spectral ripple (introduced in bandpass smoothing) is largely complete, although some of the new observations were impacted by ducted RFI. We have implemented an interrupt mechanism to abort future observations if ducting begins during the requested integration time (if ducting avoidance is active). GASKAP-OH have now received a version of their first quality gate mosaicked using Galactic coordinates. Observations of other OH lines are queued to proceed when the fields exit the solar exclusion zone.

Full survey timeline

Our initial timeline required 6 months from availability of Setonix Phase I to operational readiness. The late delivery of Setonix pushes this window close to the end of the year. Although our goal is to automate ASKAP's workflow to the greatest possible extent, the first few months of survey operations will undoubtedly reveal new issues requiring human intervention. It would therefore be difficult to start survey operations immediately prior to the end of year holidays.

Our current goal is to begin a trial month of full surveys in mid-November, including all SSTs willing to start given the list of available data products and modes at the time. If things are running smoothly at the end of this period, we may continue survey operations in an automated way over the holidays. The availability of spectral line mode will depend on resolving issues with Setonix as above.

The operations team will keep the community informed as the situation evolves.

Contact us

1300 363 400
+61 3 9545 2176
csiroenquiries@csiro.au
csiro.au

For further information

Space & Astronomy
Aidan Hotan
+61 8 8643 8543
aidan.hotan@csiro.au
https://www.csiro.au/en/research/technology-space/astronomy-space