



ASKAP update for December 2022

In this issue we discuss the outcome of ASKAP's first full survey trial and plans for continuing to expand survey operations

Full survey trial highlights

The ASKAP full survey trial commenced on November 16th, with the first full science field being EMU_1254-55, observed in scheduling block ID 45638. This observation has been processed and the data are currently being validated in CASDA. Thanks to the efforts of observatory site staff, the telescope was in excellent condition with all antennas available for the beginning of the trial. The observing pool consisted of the full survey specification for EMU, WALLABY, and FLASH, with one of the VAST dedicated survey components included to manage epoch requirements. CRAFT continued to run using incoherent commensal mode as before. The trial ran for 11 days (271 hr) before using most of the available disk space.

While attempting to run data processing on Setonix we encountered several platform issues, making it necessary to revert to Galaxy. This caused an initial delay, which was compounded by the need to address pipeline issues introduced with changes to support feature requests.

New features for ASKAP

Several new features were deployed for the first time during this trial. These observations marked the first usage of RASSP outcomes to determine SAURON scheduling priorities, previously tested only via simulation. SAURON now opts for preferential observing of paired (in the case of interleaving) or neighbouring fields, to best support prompt availability of value-added survey outputs. New software catalogues holography, bandpass and science field information, then crossmatches this information to select the best calibration data before launching processing jobs for each science field. Alongside this, relevant parameters were added to the scheduling

blocks to make automated processing more systematic and efficient. Several software components switched to Gitlab prior to the trial start, with no adverse effects on the operational system. SAURON now has the capacity to interrupt observations based on ducting or time critical targets of opportunity, with additional automated notifications implemented in the workflow.

Observing efficiency

During the active trial period, observations were managed and scheduled by SAURON. Of the 271 hr, 238 hr were successfully observed (88% efficiency, 96% success rate), of which 189 hr were allocated to surveys. The observed pool included EMU (60 hr), WALLABY (80 hr), VAST (8 hr) and FLASH (18 hr), along with some dedicated CRAFT (23 hr) filler time. Additional observing time was allocated to weights/holography (22 hr) and calibration (27 hr). Our observing KPI for 2022/2023 is 60% successful observing, so these initial results are encouraging as we look towards full surveys. However, long-term ASKAP efficiency depends on both observing and processing.




Figure 1: ASKAP observing efficiency over the period of continuous full survey trial observing, with areas below our target shaded in grey. Figure made by Vanessa Moss.

Full survey data processing

The increased computing power of Pawsey's new supercomputer Setonix will be required to keep up with ASKAP's full survey data rate. The recent trial has provided better estimates of the likely performance on Setonix, although we have yet to transition data operations to the new platform due to reliability issues.

After regaining access to an upgraded Setonix Phase I system a few days before the full survey trial commenced, we encountered several stability issues manifesting in unpredictable job failures. The ASKAP processing pipeline launches hundreds of linked tasks using the slurm scheduling system, and a significant number of these are currently failing for unexpected reasons on every attempt.

Other Setonix users have reported similar issues and Pawsey staff have been investigating the configuration of Setonix with the platform vendor. One major hardware issue has been resolved, but this did not fix all the failure modes. Until the underlying issues can be resolved, the job failure rate on Setonix is too high to be operationally feasible. We have therefore reverted to processing on Galaxy, which allows us to verify the full survey processing strategies and begin quality control.

Setonix Phase I will be progressively upgraded and Phase II deployed over the coming months. It is hoped that this will resolve the remaining issues. Once the existing workflow operates reliably on the new platform, we may look at optimisations such as using faster scratch storage space to store intermediate processing products.

Since EMU and POSSUM are using a combined commensal strategy for their full survey, the associated processing load is greater (per scheduling block) than for Pilot Surveys. This means that processing on Galaxy takes about twice the observing time even for EMU+POSSUM, and longer again for spectral line surveys such as WALLABY. We will be unable to keep up with full survey observations using Galaxy and will incur a significant overall efficiency penalty as a result. It is therefore important to transition to Setonix as soon as possible and this will be one of our top priorities for early 2023.

Plans for ongoing survey operations

ASKAP full survey operations will rely on end-to-end automation to achieve maximum efficiency. During the trial we managed to demonstrate the full workflow using the new processing manager software but had to switch off automatic processing when pipeline issues emerged after the first few attempts. Now that EMU+POSSUM and VAST mode have been verified through to the CASDA deposit stage, we are confident that fully automated operations on Galaxy should be possible, albeit with limited efficiency. Therefore, over the holiday period we will conduct a further trial of fully automated operations with EMU+POSSUM and VAST only. Concentrating on continuum projects reduces the rate at which data will accumulate and decouples observing efficiency from processing efficiency, since even with no deletion of processed data we should not fill the remaining disk space over the holiday period. We are currently investigating ways to ensure the slurm queue does not get overloaded with an accumulating number of tasks since we will be unable to keep up with incoming data on Galaxy.

The holiday observing campaign will provide statistics on processing success rates to complement observing efficiency. This will help inform the development priorities for early in the new year. One of the first priorities will be to support automatic removal of raw and intermediate files upon completion of a CASDA deposit, which should keep data flowing provided jobs succeed.

Towards the new year

Over the last few years, Pilot Surveys have brought us to the point where several of the SSTs are ready to embark on their full survey program. It has been extremely encouraging to see examples of high-quality value-added data products demonstrating the ongoing impact that ASKAP will have as full survey data flows into CASDA.

As we move into 2023, we will continue to scale up survey operations and work to support all the remaining SSTs. We would like to thank the ASKAP community for your support during 2022 and look forward to the exciting discoveries sure to emerge in 2023 and beyond.

As Australia's national science agency and innovation catalyst, CSIRO is solving the greatest challenges through innovative science and technology.

CSIRO. Unlocking a better future for everyone.

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>