



ASKAP update for January 2023

This month we report our experience with automated operations during the holiday period and plans for ongoing survey operations in 2023.

ASKAP full survey operations status

As reported in December 2022, the first month-long full survey trial integrated several new components into ASKAP's workflow and tested the observatory's capacity to conduct survey operations with a large pool of fields. Automated scheduling and processing systems are already working well and quickly evolving with operational experience. However, issues with Setonix Phase I prevented migration of data processing to the new supercomputer. We are still using Galaxy as our primary data processing platform, which means we cannot keep up with incoming data in most survey modes.

As shown in Table 1, we have processed at least a few scheduling blocks for all the Survey Science Teams currently active in the observing pool.

SST	Observed	Processing	Awaiting Validation	Released	Rejected
EMU	41	28	13	0	0
WALLABY	10	8	2	0	0
POSSUM	51	36	15	0	0
VAST	223	0	189	0	34
FLASH	9	7	2	0	0

Table 1: Status of scheduling blocks observed since the start of the full survey trial. The first column shows the total number of fields observed per survey, the second column shows the number currently progressing through various stages of processing, the third column shows the number that have been deposited into CASDA, and the final two columns show the number released to the public or rejected. POSSUM products are derived from both EMU and WALLABY observations.

Due to the large number of observations still being processed and the expiry of validation deadlines, we have suspended further observing until the validation queue clears, and Setonix is ready to process newly recorded data. This is to limit the total number of simultaneous processing jobs on Galaxy and to avoid gathering too much data before receiving feedback. We have managed

to avoid using too much of the ingest disk buffer, but all processing partitions are currently at maximum capacity.

Operations during the holiday shutdown

Upon concluding the initial full survey trial, we aimed to continue survey operations in a limited capacity across the holiday shutdown period. With reduced staff availability across the entire observatory, our goal was to rely on ASKAP's automated systems with minimal supervision and resort to human intervention only for hardware protection in the event of a major issue.

To minimise pressure on disk space and avoid the need for automated data deletion, we restricted the observing pool to continuum projects (EMU+POSSUM and VAST). Observations in this mode proceeded relatively well, demonstrating the telescope's capacity to observe, process and archive data with minimal supervision once a processing strategy has been established. We did encounter a few software issues that combined to cause one VAST epoch to be processed with an outdated bandpass calibration solution, resulting in the rejected observations in Table 1. These issues have been corrected and an additional safeguard put in place to ensure that calibration data must normally be separated by no more than 24 hours from corresponding science data.

Observing efficiency

The shutdown period was in effect from 20th December to 9th January, for a total of 475 hours. During this time, we were able to achieve an efficiency of 94.4% for successful observing, with 88.2% efficiency for surveys. Our success rate for all attempted observing in this period was 96.9%, however this only captures the observing success and does not consider any data quality or processing issues. Most failed observing was due to stormy weather early in the shutdown period. This performance bodes well for future survey operations, though improvements will likely be required to manage the full SST pool.

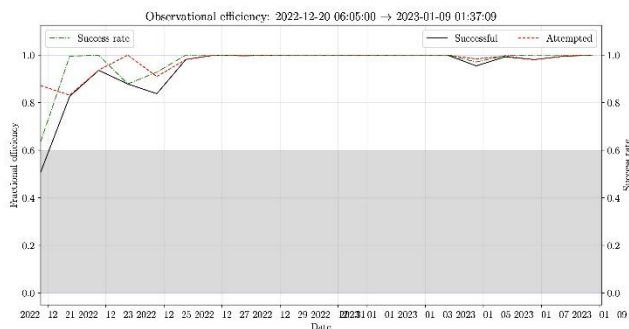


Figure 1: Plot of ASKAP observing efficiency during the 475-hr shutdown period, showing attempted (red dashed), successful (black solid) and success rate (green dot-dashed). The top of the grey region indicates ASKAP's target KPI for 2022/2023. Figure: Vanessa Moss.

Processing efficiency

Experience shows that we have been able to keep up with incoming data from VAST epochs, due in part to previous processing optimisations made during Pilot Surveys and for RACS. VAST's duty cycle is also quite low, given the epoch-based approach of the dedicated component. All VAST scheduling blocks taken under the full survey project code AS207 have been uploaded to CASDA and are awaiting validation.

On Galaxy, the full suite of processing for EMU and POSSUM takes significantly longer than the time to observe a single field. We have been able to upload 12 of 41 observations to CASDA so far, though many more are in various stages of completion. WALLABY processing takes longer still, but disk space limitations meant that fewer WALLABY fields were observed in the first place. Automated deletion of data upon successful deposit into CASDA will be required to achieve a sustainable spectral line workflow and this will be implemented soon.

The ASKAPsoft pipeline has been updated to retry job submission several times after a delay if an initial attempt is rejected by the slurm scheduling system. This has proven an effective way to work around errors that occur when many requests are received in a short space of time.

Further refinements may be made to the job scheduling system in future. Currently, processing jobs using many nodes (such as imaging) are automatically given a higher priority than smaller jobs that gather information and finalise CASDA deposits. When the queue is heavily utilised, this can delay CASDA upload for some time after the bulk of processing is complete. We would prefer something closer to a first-in-first-out approach where

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 | csiro.au/contact | csiro.au

several scheduling blocks may be processing at the same time but archiving proceeds as soon as all required data products are available for a given block.

Preparing and testing Setonix

The Pawsey Supercomputing Research Centre is currently installing additional hardware and upgrading existing components of the new Setonix supercomputer. We will need to create new software containers to run ASKAP tools on the upgraded hardware. Our tests on Setonix Phase I showed that container management is challenging due to the need for vendor-specific communication libraries that support the platform's high-speed interconnect. Pawsey staff supply a base container that we can build on, but care is required to ensure none of the key libraries are pulled from elsewhere. Once this work is complete and Setonix is available again, we will resume testing ASKAP's processing workflow on the new platform. We expect that the upgrades currently underway will address issues encountered last year and prevent spurious job failures.

Next steps

Once tests show that jobs complete reliably on Setonix and output meets expectations, we will resume survey observations using Setonix as the primary processing platform. If the queue of jobs on Galaxy empties before Setonix is ready, we will seek to resume survey operations using Galaxy. The observing pool will remain restricted to EMU, POSSUM, VAST, WALLABY and FLASH (with CRAFT incoherent commensal mode) until we have tested and verified improvements to velocity correction for GASKAP-OH, demonstrated a viable (u,v) grid compression scheme for DINGO, established a way to conduct joint deconvolution across the full ASKAP field for GASKAP-HI and verified the correlator firmware changes required to support CRAFT coherent search mode.

Data validation

All Survey Science Teams should note that we require validation of CASDA deposits no later than 28 days after upload to CASDA. Further observations will be suspended for individual Survey Science Teams if there are outstanding deposits, to ensure that the telescope's output continues to meet the required level of quality and fulfil the observatory's commitment to make data public.

For further information

Insert Business Unit name
 Insert contact name
 +61 0 0000 0000
 first.last@csiro.au
 csiro.au/businessunit