VLASS Project Memo: 24

The VLASS-II project: processing the VLASS Sky

M. Lacy, J. Ott, A. Kimball, S. Myers (NRAO),

August 21, 2025

Abstract

The high resolution and wide bandwidth of the VLA Sky Survey has resulted in a uniquely powerful dataset for studying the Universe at radio wavelengths, but also has led to a data processing challenge. The VLASS-II project was started in order to reduce the projected processing timescale of VLASS from more than a decade to ≈ 3 years (given access to sufficient computing resources). This Memo describes how this could be achieved while maintaining adherence to the survey science requirements. We will employ both new algorithms and new implementations of existing algorithms to speed up the processing and also optimize the processing parameters so as to increase processing throughput. We also describe the other aspect of the VLASS-II project, namely software improvements to enable bulk processing of VLASS images on remote compute clusters, user-defined reprocessing of individual regions of the survey and to establish a legacy archive for the survey.

1 Background

Processing the data from the VLA Sky Survey has proven to be a significant challenge. Although calibration and Quick Look Stokes I images can be produced within 2-4 weeks of observations, to obtain images that meet the survey requirements (the Single Epoch, SE, imaging) requires a large amount of computing resources. In this Memo, we describe the steps taken to enable VLASS processing within a goal timeframe of ≈ 3 years (if given access to sufficient computing resources from both in-house and external partner computing centers, as detailed in Appendix B). These steps include decisions that were made to increase imaging throughput and reduce the overall computing need to an attainable level, while maintaining adherence to the survey requirements.

2 Single Epoch imaging

2.1 SE Continuum imaging

The Single Epoch Continuum image (SECI) quality suffers if the w-terms, arising from the curvature of the sky, are uncorrected. Stokes I continuum images can be corrected to sufficient accuracy (≈ 0.2 arcsec positional error) using a simple position shift applied to the center of each 1 deg^2 image. Image-plane corrections, however, cannot fix errors in the in-band spectral index. As described in VLASS Memo 14,[1] the positional error from neglecting the w-terms is proportional to the square of the frequency, thus, the low and high frequency beams in the 2-4GHz band are spatially displaced from each other, introducing a spurious spectral index gradient across sources.

The w-terms are proportional to the tangent of the zenith distance, ζ. For $\zeta \lesssim 45$ deg., the effects of neglecting the w-terms are small and within the survey requirement of a systematic uncertainty in the in-band spectral index, α of $\Delta\alpha < 0.2$. At higher ζ though, the $\tan(\zeta)$ term grows rapidly and spectral indices are strongly affected. We thus made the decision to ignore w-terms for the $\approx 50\%$ of the survey whose tiles

[1] https://library.nrao.edu/public/memos/vla/vlass/VLASS_014.pdf

had a mean $\zeta < 45$ deg. and a maximum of 50 deg. in any one image from the tile. The rest, however, will be imaged with w-term corrections to meet the survey requirements.

Imaging with w-term corrections is possible in CASA using the awp gridder (which also includes the a-terms for improved correction of the primary beam, though these have a negligible effect in VLASS except around the very brightest ($\gtrsim 1$ Jy) sources [2]). The original, CPU-based, implementation of awp took several weeks to run for a VLASS image, however, a GPU-based implementation (awphpg) was developed that is many times faster. This is combined with a new multifrequency synthesis via cube (mvc) algorithm (which is also useable with the CPU-based awp2 gridder, an improved version of the original awp). The use of awphpg speeds the imaging considerably, (≈ 50 min/major cycle or 3-5 days for a typical image, largely independent of the number of w-planes because the runtime is dominated by the need to move data between the CPU and GPU). Each job takes one GPU and one CPU to run. 30 Nvidia L4 GPUs have been installed on the VLASS clusters and we are also testing on H100 GPUs in the Texas Advanced Computing Center as part of a joint initiative with the ngVLA. Power supply restrictions mean it is not cost effective to install more GPUs on the VLASS clusters, which restricts us to a maximum throughput of ≈ 300 images/month from the NRAO resources (taking ≈ 20 years for the $\approx 60,000$ images needed). More GPU resources will therefore be required, for example, from the Texas Advanced Computing Center (TACC) or the RADIAL node planned for the University of Puerto Rico Mayaguez (which would add another 15 GPUs). Another possibility, if GPU resources are at a premium, is to use CPUs, especially for areas of the sky where smaller numbers of planes than the current default of 32 would be adequate. For example, 16 w-planes with 12 cores could be run and only take ~ 3 times longer than the same image using awphpg. In the long term, optimization of the imaging software to exploit the shared memory of, for example, the Nvidia Grace Hopper chips, should result in much faster GPU-based gridding, but this is likely to only be implemented in the RADPS software designed for use with ngVLA data and is not expected to be available for about a decade.

Problems with the antenna pointing in VLASS1.1 resulted in poor calibration. This can be corrected, but only in the awphpg/awp2 gridders and via a very compute-intensive process. Consequently, SE products (and the combined products derived from them) will only use data from VLASS1.2,2.1,2.2,3.1,3.2 and 4.1.

2.2 SE Coarse Cubes

The VLASS coarse cubes are Stokes I, Q, U cubes made for each of the 16 spectral windows of the VLASS data (or fewer in cases where one or more spectral windows are fully flagged due to RFI). Each spectral window corresponds to a 128 MHz channel of the coarse cubes. Each of these can be imaged without the need for w-term correction and still meet the survey requirements after a frequency-dependent position fix is made to the image for each channel (see VLASS Memo 14). Nevertheless, with up to 16 cube planes and four polarizations per image, producing the coarse cubes still represents a considerable computing challenge.

The original plan for Single Epoch Coarse Cube (SECC) imaging was to make them at the same pixel scale as the SECI images, 0.6 arcseconds/pixel. However, J. Tobin recognized that the memory requirement for imaging a VLASS per-image measurement set (PIMS) is inversely proportional to the image size in pixels. We therefore decided to increase the pixel scale for cubes from 0.6 arcsec per pixel to 1 arcsec per pixel. This reduces the memory footprint per job by a factor of 0.36, enabling a much higher throughput of jobs in the cluster (which is limited by the available memory on each node). The science impact of this will be minimal as current analysis algorithms (e.g. RMTOOLS [3]) rely on the images being convolved to a common beam, usually the lowest frequency beam as that is the one with the lowest resolution (in the absence of heavy flagging in higher frequency channels). We will apply a 2 arcsec taper to the cube images to avoid undersampling the highest frequency channels.

The default SECC jobs are run in MPI CASA with five CPUs (four processing plus one management). Prior to the pixel scale change, we could only fit two jobs per node and use ten of the 24 CPUs. After the change, we can run four jobs per node and use 20 CPUs. These smaller memory requirements will also mean the jobs will fit better onto other, non-NRAO, HPC facilities that typically have less memory per core than the VLASS processing clusters. Even without additional external resources, assuming we can run cube jobs at the same time as SECI GPU jobs on the 60 available nodes the cube imaging can be completed within 3-4 years.

[2] https://library.nrao.edu/public/memos/vla/vlass/VLASS_017.pdf
[3] https://github.com/CIRADA-Tools/RM-Tools

3 Combined imaging

After the SE image sets have been made for each epoch of VLASS, they will be combined into Combined Epoch images. This will be done as a simple image-plane combination. Tests reported in VLASS Memo 21 [4] show that the results from image plane combination are comparable to those from imaging using the combined visibilites. In particular, the RMS noise and the beam sizes are almost the same, with the RMS scaling inversely as the square root of the number of epochs, as expected (Table 1 and Figure 1 of VLASS Memo 21). Combined visibility imaging results in a slightly better dynamic range as the clean threshold can be lowered relative to the individual epochs. There are, however, two reasons why image combination is preferred. The first is computational cost - because each epoch was taken at a different elevation, the w-term corrections need to be applied (using awp) over the whole survey area. The large data volumes mean that this imaging may take up to three times longer than imaging an entire individual epoch ($>$ 60 years if only in-house resources are available). The second issue relates to variable sources - if a source varies between epochs the difference in flux cannot be accurately cleaned, reducing the dynamic range. Although algorithms could be developed to deal with this, several person-years of development and testing would be needed to implement them.

4 Overall Computing needs

The VLASS compute cost to complete processing is estimated by pipeline in Table 1. In total, the needs will be approximately 46 million core hours and 5 million GPU hours. With about 10 million CPU hours/year available in house and the more efficient utilization discussed in this document, this will be sufficient to complete CPU-based VLASS processing in 4-5 years even if no external resources are used, or within 3 years if we can outsource 30-40% of the processing. However, the GPU requirements (5 million GPU hours) are much larger than available in-house (0.25 million GPU hrs/year) and we will need external computing to complete the GPU-based processing in a timely manner. Our collaborations with the TACC and University of Wisconsin, plus the possibility of a RADIAL node at the University of Puerto Rico Mayagüez should mitigate this issue, however, it will also be possible to run the SECI imaging with awp2 on CPUs if sufficient GPU resources are not forthcoming, or if it is much easier/cheaper to use CPUs on external resources than GPUs.

5 Archive and workflow enhancements

5.1 Support for multiple SECI job types and combined images in the VLASS manager

The VLASS manager product types currently include only one type of SECI (mosaic gridder) and no support for cumunlative imaging. Changes to allow multiple SECI types (mosaic, awphpg, awp2) will need to be introduced, along with support for the combined products.

5.2 Remote processing

Given that much processing will need to be performed remotely, the current VLASS manager software will need to enable runs on resources outside of Socorro in an automated manner. The Open Science Grid [5] provides a software stack to enable copying of data to and from remote nodes and is compatible with our likely partners (the Center for High Throughput Computing is heavily involved with the OSG and the TACC say they could install their software). However, development work at NRAO will be needed to implement this stack locally and connect the OSG software to the VLASS manager. We will also need software effort from the CASA and pipeline teams to adapt the VLASS CASA pipeline to the different architectures, GPUs and operating systems used by partner facilities (which can also change with time as new clusters become available and old ones are retired).

[4] https://library.nrao.edu/public/memos/vla/vlass/VLASS_021.pdf

[5] https://osg-htc.org

5.3 On-demand processing

We will make available an on-demand imaging service for specialist use cases that go beyond the standard VLASS data products for small areas of sky, for example, cubes of bright objects at finer spectral resolution (down to the 2 MHz channelization of the data), or combined visibility imaging if needed for high dynamic range use cases. This service will be based on the VLA on-demand imaging service (VUDI) that will be developed by NRAO in the next 1-2 years. Ideally, this would be supported by an ≈ 2 PB storage system that would allow all VLASS calibrated measurement sets to be available with a short (few hour) latency.

5.4 Cutout and catalog servers

A cutout service will be developed to extract spatial regions from the SE and combined image data products. This will have both an interactive (webform) interface and an API for scripted queries. The API will enable bulk downloads of cutouts, for example, for AI/ML related work on VLASS images. The API should be able to handle large numbers of queries rapidly (goal of 1000 cutouts per hour). A catalog server will also be provided to allow downloads and searches of the VLASS catalogs by position (cone and polygon searches) and parameters (e.g. flux density and spectral index). These services should follow Virtual Observatory protocols so far as is practical for easy interoperability with external services.

5.5 Polarization tools and services

The VLASS polarization cubes are especially challenging to analyze. We will provide tools to regrid and convolve polarization cubes and perform Faraday synthesis. Ideally, these would be provided by an external collaboration with specific expertize in polarimetry, for example, using the pipeline developed for POSSUM and VLASS by the Canadian CIRADA project. It could be a Canadian contribution, or other partners e.g. NAOJ or IDIA in South Africa might be interested. If not, we do have the expertize in house to develop scripts and notebooks based on CASA tools/tasks and the polarization analysis tools such as Faraday Synthesis available in the RMTOOLS package. These could be adapted into web or API-driven tools at a later date.

6 Initial data release

Based on a recommendation from the VLASS 4th Epoch review panel, we will initially focus on creating a set of full products for a specific region of sky, namely the equatorial stripe (R.A. 0h-24h; Dec $-4^{\circ} - +4^{\circ}$). Nearly all of the tiles in this region can be imaged with the mosaic gridder for SECI and the remainder will be prioritized for awphpg processing. This region is well covered by surveys at other wavebands (SDSS Stripe 82, HyperSuprimeCam etc.) and will give users nearly 10% of the VLASS sky in all three epochs of SE imaging and combined imaging.

7 Uncertainties and risks

The 3 year timescale is very ambitious and depends on a number of factors:

- Access to sufficient compute. In-house CPU resources are sufficient to complete SECI mosaic gridder and SECC processing in 4-5 years and about 20 million core hours from external resources would be needed to complete this processing within 3 years. The 3 year timeline is quite feasible given the low cost of CPU-based processing on external clusters ($\sim$ \$0.003/CPUhr), though it depends how our providers would charge if we were only able to utilize a small fraction of the cores on a node due to memory limitations. The GPU processing presents a much bigger challenge. With the 30 GPUs available in-house, giving 0.25 million core hours per year, processing would take ≈ 20 years. Access to external resources are therefore essential for this aspect of the processing, but 5 million GPU hours at current rates have a value of about \$10 million. As discussed above, it may also be possible to run awp2 with CPUs, this would have a large memory requirement, but jobs would still fit on the NRAO nodes and could be run there if more of the SECC processing could be moved to external facilities.

In either case though, significant external resources would need to be obtained to achieve the 3 year processing target.

- Another area of risk is the changes to the VLASS manager software that will be needed to support the processing of multiple types of job on external resources. Current planning for the Science Support and Archive (SSA) team that would perform the bulk of this work extends > 1 year into mid-2027. Small amounts of work (e.g. enabling new SECI modes) could be scheduled before then, but significant changes such as enabling fully remote processing may take longer and will need collaboration with Science Information Support (SIS). The SSA group will also be responsible for the archive enhancements and on-demand processing that are also part of the VLASS-II plan, though these needs are less urgent.
- CASA and the VLASS pipeline that runs in it are only compiled for a small selection of operating systems, computing architectures and GPUs. Many of our potential external computing partners are using different systems not covered by the current CASA builds. Effort will be needed from the already oversubscribed CASA/pipeline team to enable builds on these systems and from the VLASS operations team to test them for scientific validity.

8 Summary

The VLASS-II project has the dual aims of accelerating the processing of VLASS and providing archive services for VLASS products that will ensure long-term accessibility to the survey data. For processing, the aim is to set up piplines and infrastructure that will allow processing on a short (3 year) timescale given the available internal computing and a realistic amount of external computing using facilities from existing and likely future partners (subject to budgetary constraints). For archive services, a key contribution will be an on-demand imaging service that will allow users to image small ($<< 1$ deg^2) areas of sky around individual sources of interest at finer spectral resolution and/or combine the three epochs of visibility data and image. An image cutout service and catalog server are also needed to improve accessibility as well as polarization tools and services.

A List of acronyms

API - Application Program Interface
CIRADA - Canadian Initiative for Radio Astronomy Data Analysis
IDIA - Inter-university Institute for Data Intensive Astronomy (South Africa)
LCCF - Leadership Class Computing Facility (successor to the TACC).
NAOJ - National Astronomical Observatory of Japan
OSG - Open Science Grid
PIMS - per-image measurement set. These contain the calibrated visibility data split out for each VLASS image.
RADPS - radio astronomy data processing system (data processing software being developed for ALMA WSU and ngVLA)
SECI - Single Epoch Continuum Image
SECC - Single Epoch Coarse Cube
SIS - NRAO's science information support team
SSA - NRAO's science support and archive group.
SSG - the VLASS Survey Science Group.
TACC - Texas Advanced Computing Center
VLASS - Very Large Array Sky Survey

B Processing requirements for VLASS-II

The pipelines used for VLASS consist of an SE continuum and an SE cube pipeline, both using CPUs only. We will soon be adding an additional SE continuum pipeline that utilizes a GPU gridder and can enhance the

Parameter	SE-continuum* (no w-projection)	SE-continuum (w-projection)	SE-Cube	Total
Number of jobs	40,000	60,000	120,000	220,000
Total CPU hrs	5 million	5 million	36 million	46 million
Total GPU hrs	-	5 million	-	5 million
Total data in	0.35 PB	0.5 PB	1PB	2 PB
Total products	35 TB	50 TB	2 PB	2.1 PB
RAM/CPU core	64 GB	128 GB	30 GB	-
CPU Cores/job†	1	1	4	-
GPUs /job (L4)	-	1	-	-
Job duration‡	6 days	3.5 days	3 days	-
Storage/job	150 GB	150 GB	170 GB	-

Table 1: Remaining processing needs for VLASS-II, including a 13% contingency for reruns.
* Excludes the $\approx 18,000$ images already made.
† Core numbers can be varied for CPU-based jobs
‡ Job durations can run 2-3 times longer on complicated fields with bright emission.

accuracy of the products by using w-plane projection. These enhanced accuracy products are needed over about half the sky (corresponding to observations made at >45 deg. zenith distance), where the w-terms are relatively large and ignoring them results in the products failing to meet the survey science requirements. Details of the processing requirements for these are given in Table 1. Note that these numbers assume a 13% reprocessing fraction and exclude the SE continuum images already made.