

Technical Implementation Plan for the Jansky Very Large Array Sky Survey (VLASS)

Steven T. Myers, Claire J. Chandler, & Casey Law

February 2, 2015

Contents

1	Overview	4
2	Sensitivity, Angular Resolution, & Overheads	4
2.1	Point-Source Sensitivity Assumptions	4
2.1.1	Useable bandwidth	4
2.1.2	Sensitivity Loss at Low Elevation	5
2.2	Angular Resolution	5
2.2.1	Case Study: Stripe-82 and COSMOS at $\delta = 0^\circ$	5
2.3	Overhead Assumptions	6
2.3.1	Case Study 1: Stripe-82	6
2.3.2	Case Study 2: A simplified model	7
2.3.3	Applying the Simplified Model to VLASS	8
2.4	Scanning Speed	8
2.5	Cadences for Transients & Variability	10
2.6	Source Spectrum, Angular Size & Sensitivity	10
2.6.1	Source Spectrum & Sensitivity	10
2.6.2	Source Angular Size and Surface Brightness Sensitivity	12
2.6.3	Measurement of the Source Spectrum	13
3	Survey Structure	16
3.1	Tier 1: ALL-SKY	16
3.2	Tier 2: DEEP	17
3.2.1	Tier 2 DEEP — COSMOS	17
3.2.2	Tier 2 DEEP — ECDFS	17
3.2.3	Tier 2 DEEP — Elais-N1:	18
4	Data Products	19
4.1	Basic Data Products	19
4.1.1	Raw Visibility Data	20
4.1.2	Calibrated Data	20
4.1.3	Quick-Look Images	20
4.1.4	Single-Epoch Images	21
4.1.5	Single-Epoch Basic Object Catalogs	22
4.1.6	Cumulative Images	23
4.1.7	Cumulative Basic Object Catalogs	24
4.2	Enhanced Data Products	24
4.3	Enhanced Data Services and the VLASS Archive	25
4.4	Data Formats	26
4.4.1	Visibility Data	26
4.4.2	Images and Image Cubes	26
4.4.3	Catalogs	26
4.4.4	Plots	26

5	Observing	26
5.1	Mosaicking	27
5.2	Scheduling Considerations	28
5.3	Schedule Construction	29
5.4	Overall Observing Schedule	29
5.4.1	Scenario: 7-Year Schedule	29
5.5	Schedule Pressure by LST	30
5.6	Impact of VLASS on PI science	31
6	Calibration	32
6.1	Observations of Calibrators	33
6.2	Calibration Issues for the VLASS	34
6.2.1	RFI at S-band	34
6.2.2	The Ionosphere at S-band	35
6.3	Algorithm and Software Development for the Calibration Pipeline	35
6.4	Calibration Processing Requirements	35
6.4.1	Case Study: Calibration of Stripe-82 datasets	36
6.5	Temporary Storage Requirements for Calibration	36
7	Imaging	36
7.1	Imaging Requirements	37
7.2	Algorithm and Software Development for Imaging	39
7.3	Processing Requirements for Imaging	39
7.3.1	Test Case: Imaging of Stripe-82	39
7.4	Temporary Storage Requirements for Imaging and Analysis	40
8	Image Analysis and Sky Catalogs	41
9	Archiving and Data Distribution	41
9.1	Estimated Archive Space	41
9.2	Raw Visibility Data	42
9.3	Calibration Data	43
9.4	Quick-Look Continuum Images	43
9.4.1	Tier 1 ALL-SKY	43
9.4.2	Tier 2 DEEP — COSMOS	43
9.4.3	Tier 2 DEEP — ECDFS	43
9.4.4	Tier 2 DEEP — Elais-N1	43
9.5	Single-epoch images and cubes	44
9.5.1	ALL-SKY Tier 1	44
9.5.2	DEEP Tier 2	44
9.6	Single-epoch basic object catalogs	45
9.7	Final cumulative “static sky” images and image cubes	45
9.7.1	ALL-SKY Tier 1	45
9.7.2	DEEP Tier 2	46
9.8	Cumulative “static sky” basic object catalogs	46
10	Test and Development Plan	46
10.1	General Flagging, Calibration, and Imaging Tests	46
10.1.1	RFI, Bandwidth, and Sensitivity	47
10.1.2	Wide-band Wide-field Continuum Imaging	47
10.1.3	Polarimetry	49
10.1.4	Transients and Variability	49

10.2 ALL-SKY Tier 1 OTFM Testing	50
10.3 Tier 2 Deep Imaging Tests	52
10.4 Other Logistical Tests	53
10.5 Pipeline development	53
11 Project Schedule and Resource Estimates	54
11.1 Schedule	54
11.2 Resource Requirements	54
12 Acknowledgements	56

1 Overview

This working document details the Technical Implementation Plan (TIP) for the VLA Sky Survey (VLASS).

The main Survey Science proposal “The Jansky Very Large Array Sky Survey (VLASS)” contains the science justification and overall description of the survey. This supplementary document describes in more detail the technical issues, project plan, and risks associated with the proposed VLASS.

Version Note: This draft contains the revised numbers for the new survey proposal after the NRAO review.

2 Sensitivity, Angular Resolution, & Overheads

The following assumptions about the sensitivity, angular resolution, and calibration and slewing overheads are used to calculate the parameters of the survey.

2.1 Point-Source Sensitivity Assumptions

The point-source sensitivity of the Jansky VLA for mosaicking is computed using the procedure given in the Guide to VLA Observing: Mosaicking¹. For continuum (Stokes I) at S-band (2–4 GHz) we assume a Survey Speed (SS) of

$$SS = 16.55 \left(\frac{\sigma_I}{100 \mu\text{Jy}/\text{beam}} \right)^2 \text{ deg}^2 \text{ hr}^{-1} \quad (1)$$

of on-sky integration time for an assumed image rms of σ_I ($\mu\text{Jy}/\text{beam}$). This assumes 25 antennas available at all times, 1500 MHz of useable bandwidth (after RFI excision), and an image averaged over the band using multi-frequency synthesis. The integration time needed to survey a given area to a depth is given by dividing that area by the survey speed.

2.1.1 Useable bandwidth

A key parameter that factors into the sensitivity is the amount of useable bandwidth from 2–4 GHz. Standard JVLA guidelines and the VLA Exposure Calculator Tool² recommend a value of 1500 MHz, and therefore this is what we used in Equation 1, and was employed to calculate the required observing time to reach a given depth. However, the exact amount of RFI affected frequency space is not yet carefully quantified, and also depends upon direction (antenna pointing direction in Azimuth and Elevation in many cases) and time of day.

As a case in point, the Stripe-82 surveys referenced throughout this TIP (12A-371 PI: Kulkarni, and 13B-370, PI: Hallinan) employed fairly brutal excision, throwing away whole spectral windows. The effective bandwidth used for the imaging of these datasets was around 1350 MHz. If the VLASS employed this level of bandwidth reduction, then we would require an 11% increase in required observing time to reach the specified depths of the survey, or would incur a 5.4% penalty in achieved rms imaging sensitivity. On the other hand, even one extra antenna (26 versus the assumed 25) would provide an effective time increase of 8.3% which gets back most of this loss. In addition, it is worth noting that for many S-band continuum experiments that are run through the production pipeline, only two 128 MHz spectral windows are lost. The effective bandwidth is a function of position on the sky and time of day of observation.

¹<https://science.nrao.edu/facilities/vla/docs/manuals/obsguide/modes/mosaicking>

²<https://obs.vla.nrao.edu/ect>

Risk: Medium. We consider this to be medium risk for the VLASS. This is one of the key issues to be resolved in the Test and Development Program (§ 10).

2.1.2 Sensitivity Loss at Low Elevation

The VLA exhibits sensitivity loss at low elevations, in the S-band primarily due to an increase in system temperature from spillover from the over-illuminated secondary. In the calculations used to estimate the required integration time, we have used measured values for this noise increase provided by Rick Perley (private communication).

A more detailed explanation of the sensitivity and observing time calculations are available in an online iPython Notebook.³

2.2 Angular Resolution

The Observational Status Summary (OSS)⁴ for the VLA lists the angular resolution (FWHM) for S-band (at 3GHz) as 2.1'' in B-configuration and 0.65'' in A-configuration. These apply to long tracks at high declination using *uniform* weighting (equivalent to a fully filled uv-plane out to the maximum baseline). The current advice given to users in the OSS is “The listed resolutions are appropriate for sources with declinations between -15 and 75 degrees...The approximate resolution for a naturally weighted map is about 1.5 times the numbers listed for θ_{HPBW} . The values for snapshots are about 1.3 times the listed values.” These values are derived from the 1995 PhD thesis of Dan Briggs⁵ based on classic VLA tests using the narrow continuum bandwidths then available. More careful examination of Briggs’ thesis suggests that the resolution for naturally-weighted snapshots is 1.4 times the numbers listed for uniformly-weighted, long track observations in the OSS.

The above standard numbers for uniform weighting, coupled with the Briggs factor of 1.4 for naturally weighted snapshots, imply that we should obtain a naturally weighted resolution of 2.9'' in B-configuration and 0.9'' in A-configuration. These are expected to be valid for regions with $-15^\circ < \delta < 75^\circ$, fields outside this range will have degraded resolution (unless observed in a hybrid configuration).

It is not clear how applicable the Briggs thesis numbers are to the wideband JVLA observations proposed here, and so this issue will be further investigated as part of the Test & Development Program (§ 10.1.2). In particular, the Briggs thesis found that “super-uniform” weighting of snapshots would yield a rms sensitivity 1.11 times the naturally weighted values, with an increase of the beam size of only 1.2 times that for uniform weighting. If this holds for VLASS continuum imaging, this would yield resolution of 2.5'' in B-configuration and 0.8'' in A-configuration at an rms image sensitivity level 1.11 times the natural weighting values quoted for the survey components in § 3.

It is the baseline plan to use Briggs weighting with a “robust” parameter $r = 1$ for the Basic Data Products, with differently-weighted images falling under Enhanced Data Products (see below).

2.2.1 Case Study: Stripe-82 and COSMOS at $\delta = 0^\circ$

In the 13B-370 Stripe-82 B-configuration data⁶, we found (at Declination $\delta = 0^\circ$) snapshot synthesized beam sizes (FWHM) of around 3.4''2.3'' for natural weighting, with a geometric mean beam

³<http://goo.gl/znDX0i>

⁴<https://science.nrao.edu/facilities/vla/docs/manuals/oss/performance/resolution>

⁵<http://www.aoc.nrao.edu/dissertations/dbriggs/>

⁶SB 13B-370.sb28581653.eb28626177.56669.781848645835

size of $2.8''$. This is encouragingly close to the expected values based on the Briggs study. In a small sub-mosaic in a relatively empty part of the sky imaging tests gave geometric mean PSF sizes of $2.6''$ (Briggs robust $r = 1.0$) down to $2.2''$ ($r = 0.5$) with image rms of $81 \mu\text{Jy}$ ($r = 1.0$) and $79 \mu\text{Jy}$ ($r = 0.5$) respectively. Note that $r = 0.5$ gives slightly better rms due to better sidelobe suppression from the sources that are in the field. Using the assumptions of sensitivity for VLASS (1500MHz bandwidth, survey speed as in Eq. 1) and the survey speed used for the Stripe-82 observations ($10.6 \text{ deg}^2/\text{hr}$), we would expect a “naturally weighted” image rms of $80 \mu\text{Jy}$. Thus we conclude that using the naturally weighted survey speed sensitivity given by Equation 1 and listed in the proposal and in the next section is valid, within uncertainties given by the issues raised in § 2.1.1 and the calibration of the Stripe-82 data.

2.3 Overhead Assumptions

In the estimates for total observing time, the assumption is made that there is an overhead for slewing, setup, and calibration applies to the components of the survey. For example, for general VLA observing, we recommend use of a 25% overhead, where integration times are to be multiplied by 1.25 in order to arrive at the “clock time” needed to execute observations. In practice, the overhead will depend on exactly how the survey components are scheduled and how much calibration is required, and what calibration can be shared between blocks.

Below we present two case studies based on an actual schedule and on a simple model respectively. Based on these studies, we fully expect to be able to meet (or come in under) the global 25% overhead for (i) observing blocks of 6 hours or longer, and (ii) short blocks that can share calibration with other blocks observed around the same time, (iii) blocks of 3 hours or longer that are fixed-time scheduled. More detailed verification of overheads, including full calibration, are scheduled as part of the Test Plan.

2.3.1 Case Study 1: Stripe-82

As an example, the S-band Stripe-82 observations of program 13B-370 (Hallinan et al.) were observed using dynamic scheduling and independent 3-hour blocks with 2.25 hours of on-target observing, and thus had a 33% overhead. This particular base schedule were self-contained, and included standard calibrators 3C48 (for flux density and polarization angle) and 3C84 (low-polarization leakage calibrator). The observations were made in On-The-Fly Mosaic (OTFM) mode, with the OTFM scans comprising 15 stripes to cover the Declination range, with each stripe covering 36min (9 degrees) in RA. At the scan rate chosen, each stripe took 9.25 min to get 9 min on-source integration (3% scan overhead). Groups of two stripes were interspersed with calibrator scans (with a singlet stripe left over in the block). There were extra observations of calibrators from adjacent blocks included to help link the blocks. This was a fairly conservative strategy, but control of calibration errors was important for the program.

Starting with the actual SBs used for the observations, a number of example SBs were constructed under a range of different assumptions. The overheads ranged from 24% to as low as 14% under different less conservative criteria. See Table 1 for results from the example schedule construction exercise.

All blocks contained full calibration. Note that whether an example is “dynamic” or not depends on whether this SB could be submitted over a significant range of LST and thus includes extra padding for possibly long slews or wraps. If it needed a narrow range of starting LST then it was deemed to be essentially a fixed time schedule.

Table 1: Overheads for Example Scheduling Blocks

Duration (hours)	Dwell (hours)	Overhead	Dynamic?	Notes
3.0	2.25	33%	yes	(1)
2.75	2.25	22%	no	(2)
5.6	4.5	24%	yes	(3)
5.25	4.5	17%	no	(2)
7.7	6.75	14%	no	(2)

Table Notes: (1) 13B-370 original schedule, multiple gain calibrators, calibration every 20min; (2) single gain calibrator, calibration every 30min; (3) multiple calibrators, calibration every 20min.

2.3.2 Case Study 2: A simplified model

We now present a simplified model of overheads in SBs, guided by the first case study above. In this model we assume:

- a fixed startup overhead (to slew to first source) of 10 min.
- a fixed time on per calibrator observation of 2 min.
- for each hour of on-source integration, a number N of 2 min. calibration scans (e.g. N=2 or N=3)
- a fixed flux density and polarization calibration (on a low polarization source) of 15 min.
- (if needed, an additional fixed polarization leakage calibration of 15 min with 3 scans at different PA with 5 min. per scan)

Thus in this model, there is a fixed overhead of 27 minutes, plus a possible additional overhead of 15 minutes for the polarization calibration, so the fixed per-block overhead is either 30 minutes or 45 minutes (rounded).

These assumptions are meant to be appropriate for large areas over a range of Declinations, where slew times to calibrators may vary. Small areas (e.g. Deep fields) may or may not conform to this model.

Assuming our calibration every 20 min (N=3), and allowing for a single leakage scan on an unpolarized source (or use of the phase calibrators for PA coverage of a polarized source), then we obtain the overheads in Table 2. The overhead for calibration every 30 min. (N=2) is about 3%

Table 2: Overheads for Simplified Model A — Calibration every 20min, single leakage scan

On-source time (hours)	Duration (hours)	Overhead
3.0	3.75	25%
4.0	4.75	19%
6.0	7.05	18%
10.0	11.45	15%

lower in all cases.

Table 3: **Overheads for Simplified Model B — Calibration every 20min, 3 leakage scans**

On-source time (hours)	Duration (hours)	Overhead
3.0	4.00	33%
4.0	5.00	25%
6.0	7.30	22%
10.0	11.70	17%

If we add the additional 15 min. fixed overhead for 3 leakage scans, then we obtain the overheads in Table 3. The overhead for calibration every 30 min. (N=2) is again about 3% lower for all cases.

2.3.3 Applying the Simplified Model to VLASS

We will use Simplified Model B as a conservative estimate of overheads for the VLASS. We assume the overheads listed in Table 4 for the various elements of the survey. These numbers are adopted

Table 4: **Overheads for VLASS under Simplified Model B**

Component	Overhead	Comments
ECDFS	25%	5-hr blocks, contained calibration
COSMOS,E-N1	22%	6–8 hour blocks, contained calibration
ALL-SKY	17%	long blocks, shared calibration

as guidance in the per-Tier times given below. Note that the overheads actually used for scheduled hours for the DEEP fields are close to these, but rounded to make the schedules of length rounded to nearest 15min. For example, the COSMOS field (§ 3.2.1) will be observed with 25% overhead rather than 22% due to this rounding.

It is likely possible to further reduce overheads if needed, particularly for the DEEP fields which are observed repeatedly over multiple days per cycle, by sharing primary calibration (flux density, polarization angle, possibly leakage). This will require testing to determine the optimum strategy. Note that this use of shared calibration is not included in the proposed VLASS plan, which employs standalone calibration for each block.

2.4 Scanning Speed

The Survey Speed (SS) defined above is the sky area covered per unit time to reach a given continuum image rms sensitivity (16.55 deg²/hour for 100μJy/beam rms). For an On-the-Fly Mosaic (OTFM) this translates to a Scanning Rate (SR) in arcmin/second at which the telescopes move while taking data. There are limitations on how fast one can scan, mostly dependent on the data rates resulting from the need to support short data dump times during fast scanning.

In OTFM mode, the survey area is covered by a series of stripes (“rows”) spaced by θ_{row} . For nearly uniform sensitivity coverage over the band, we adopt a value

$$\theta_{row,max} = \theta_{pb,min} / \sqrt{2} = 8.84' \quad (2)$$

where $\theta_{pb,min} = 12.5'$ is the primary beam FWHM at the highest useful frequency within S-band of 3.6 GHz (most of the spectrum above this frequency is lost to RFI). *Note: In practice we will adjust*

this spacing to a value to put the rows on a more sensible, grid, e.g., $\theta_{\text{row}} = 8.75'$. This striping parameter translates the survey speed to an antenna slewing speed relative to the sky:

$$SS_{\text{max}} = \dot{\theta}_{\text{max}} \theta_{\text{row,max}} \quad (3)$$

$$= 8.84' \dot{\theta}_{\text{max}} \quad (4)$$

where $\dot{\theta}_{\text{max}}$ is the maximum allowed scanning rate.

In an ideal case, OTF enables any survey to be conducted as N -epochs, each with depth \sqrt{N} of the final depth. The VLASS science case assumes these can be imaged individually and jointly to produce both deep maps and information about the transient radio sky. However, a critical limitation is that the correlator dump time must be fast enough to temporally resolve the amplitude envelopes of sources as they move through the primary beam pattern. The dump time must sample sources faster than roughly 10 times its motion through the primary beam:

$$t_{\text{dump}} * \dot{\theta} < 0.1 \theta_{\text{pb,min}}. \quad (5)$$

or

$$\dot{\theta}_{\text{max}} = 0.1 \theta_{\text{pb,min}} t_{\text{dump,min}}^{-1} \quad (6)$$

$$= 1.25' t_{\text{dump,min}}^{-1} \quad (7)$$

for our primary beam at 3.6GHz.

The minimum dump time (and thus number of possible epochs) is limited by the output data rate, which should not exceed 25 MB/s in order to limit the archiving requirements. The VLA produces data at a rate R_{data} of:

$$R_{\text{data}} = 45 * (n_{\text{spw}} * n_{\text{ch}} * n_{\text{pol}} / 16384) / t_{\text{dump}}. \quad (8)$$

Conservatively assuming that all 16 subbands are useful (2 GHz bandwidth), each with 64 2MHz channels and 4 polarization products, the VLASS will produce 11.25 MB per dump. For a data rate limit of 25 MB/s, the shortest full-VLASS dump time is 0.45s for the nominal VLASS configuration bringing back the full 2048 MHz bandwidth ($n_{\text{spw}} = 16$). Thus,

$$t_{\text{dump,min}} = 0.45 \text{ sec } (n_{\text{spw}} / 16) \quad (9)$$

and thus

$$\dot{\theta}_{\text{max}} = 1.25' t_{\text{dump,min}}^{-1} \quad (10)$$

$$= 2.778' / \text{sec } (16 / n_{\text{spw}}). \quad (11)$$

Putting these together, we get:

$$SS_{\text{max}} = 8.84' \dot{\theta}_{\text{max}} \quad (12)$$

$$= 24.56 \text{ arcmin}^2 / \text{sec } (16 / n_{\text{spw}}) \quad (13)$$

$$= 24.56 \text{ deg}^2 / \text{hr } (16 / n_{\text{spw}}). \quad (14)$$

We established earlier that for a continuum image rms of $\sigma_1 = 100 \mu\text{Jy}/\text{beam}$ sensitivity, the implied survey speed was $16.55 \text{ deg}^2 / \text{hour}$. Therefore, the scan-rate limited image rms for the JVLA is

$$\sigma_{1,\text{max}} = 100 \mu\text{Jy}/\text{beam } (SS_{\text{max}} / 16.55 \text{ deg}^2 / \text{hr})^{1/2} \quad (15)$$

$$= 122 \mu\text{Jy}/\text{beam } (16 / n_{\text{spw}})^{1/2} \quad (16)$$

and therefore we have set the imaging sensitivity limit for the ALL-SKY tier to be $120 \mu\text{Jy}/\text{beam}$. Using this, the difference from the value in Equation 16 will allow us to use a more practical value $\theta_{\text{row}} = 8.75'$ as described above.

Risk: Medium We have not done extensive testing of OTFM near this scanning limit. We do not expect to encounter issues that cannot be fixed in our observing system, but this is a key element of the testing plan below.

2.5 Cadences for Transients & Variability

One of the headline science goals of the VLASS is the exploration of the dynamic radio sky on timescales of days to years. There is some latitude with which the scheduling of the VLASS can influence the sensitivity to transient and variable phenomena through the selection of cadences for fields.

The ALL-SKY component will be carried out in 3 total epochs (each area is observed three times) over the six configuration cycles spanning 7 years. The most straightforward approach is to divide the sky into two (equal observing time) sets of regions, say N and S (possibly, but not necessarily, referring to North and South), and alternate between them each cycle, e.g. NSNSNS. In this scheme there is a 32-month cadence between epochs. For known transient classes at 2–4 GHz this means that each epoch is essentially independent and will observe fresh cases of these transients on the sky. In the Test & Development plan (in § 10.1.4) we propose to investigate optimal schemes. For example, an alternate approach might be to introduce two cadences, e.g. NNSNSS where both 16-month and 32-month separations are measured.

The DEEP fields will be observed on multiple passes each cycle, and thus timescales from a single day to the full 7 years will be explored. Note that different fields will get different cadences, as they have differing number of passes given different block lengths as well as extra overhead for ECDFS.

In addition to finding and characterizing transient and variable sources, the VLASS aims to make deep images of the static sky. For this goal, the effects of time-variable sources (making up $\sim 1\%$ of the population at 1 mJy) must be accounted for in imaging. This is most likely best done by observing the ALL-SKY areas at fixed LST and thus with identical uv coverage per pass. This may also be needed for the DEEP fields (though by necessity ECDFS will be almost fixed scheduled anyway). Note that identical uv coverage will also be useful for controlling systematics in the transient and variable studies, particularly for polarization.

2.6 Source Spectrum, Angular Size & Sensitivity

The basic calculations presented in § 2.1 are valid for sources unresolved by the VLA in a given configuration, and with a spectrum that is either constant or linearly varying with frequency. Sources that break one or both of these assumptions will have an apparent brightness in the VLASS that is different from their flux density at band center. We deal with these two effects below.

2.6.1 Source Spectrum & Sensitivity

If the spectral energy distribution (SED) of a radio source were constant or a linear function of frequency, and the VLA sensitivity were constant over the band, then the measured signal-to-noise ratio for a source would be given by the ratio of the flux density at band center to the rms noise over the band. However, this is not really the case for actual radio sources, or for the sensitivity of the JVLA S-band system.

The SED of a radio source is generally taken to be a power-law as a function of frequency

$$S(\nu) = S_0 \left(\frac{\nu}{\nu_0} \right)^\alpha \quad (17)$$

where α is the spectral index and S_0 is the flux density at the reference frequency ν_0 . For example, the flux density of an average radio source with $\alpha = -0.7$ measured at the VLASS band center of 3 GHz is 0.587 times the flux density measured at the FIRST/NVSS frequency of 1.4 GHz. When we make a continuum image (of an unresolved source) we form the band-weighted average

$$\bar{S} = \frac{1}{Z} \int_{\nu_{min}}^{\nu_{max}} d\nu w(\nu) S(\nu) \quad Z = \int_{\nu_{min}}^{\nu_{max}} d\nu w(\nu) \quad (18)$$

where Z is the normalization factor (integral of the weights w). Note that the rms noise is also a band averaged quantity

$$\bar{\sigma}^2 = \frac{1}{Z} \int_{\nu_{min}}^{\nu_{max}} d\nu w(\nu) \sigma^2(\nu) \quad (19)$$

where $\sigma(\nu)$ is the rms image noise (in flux density rms units) as a function of frequency. The calculations that lead to Equation 1 assumed a particular form of the VLA system equivalent flux density (SEFD) in the online calculator, and a (constant) assumed integration time on source. The true performance of the JVLA over S-band (versus the assumed sensitivity) is a Test & Development target. For now we will assume that this is a known function.

We first consider the uniform case where $w = 1$ with uniform frequency coverage between ν_{min} and ν_{max} . In this case $Z = \Delta\nu = \nu_{max} - \nu_{min}$. Then it is straightforward to show for the power-law model (Eq. 17) that

$$\bar{S}_{unif} = \frac{1}{Z} \int_{\nu_{min}}^{\nu_{max}} d\nu S_0 \left(\frac{\nu}{\nu_0} \right)^\alpha \quad (20)$$

$$= S_0 \frac{\nu_0^{-\alpha}}{\alpha + 1} \frac{\nu_{max}^{\alpha+1} - \nu_{min}^{\alpha+1}}{\Delta\nu} \quad (21)$$

$$(22)$$

for all $\alpha \neq -1$ (for $\alpha = -1$ the integral becomes a logarithmic function). Note that for a flat spectrum source $\alpha = 0$ we recover the constant flux density S_0 . For other values of the spectral index we can define an equivalent frequency ν_{eq} for which

$$S(\nu_{eq}) = S_0 \left(\frac{\nu_{eq}}{\nu_0} \right)^\alpha = \bar{S}_{unif} \quad (23)$$

which gives

$$\nu_{eq} = \left[\frac{1}{\alpha + 1} \frac{\nu_{max}^{\alpha+1} - \nu_{min}^{\alpha+1}}{\Delta\nu} \right]^{\frac{1}{\alpha}} \quad (24)$$

for $\alpha \neq 0$ (as well as $\alpha \neq 1$). For $\alpha = 1$ this gives $\nu_{eq} = 0.5(\nu_{min} + \nu_{max})$ the average frequency in the band. For other reasonable values we obtain an equivalent frequency close to but below the mean frequency. For example, for a typical radio source $\langle \alpha \rangle = -0.7$ and the full S-band frequency range one finds $\nu_{eq} \approx 2.9$ GHz. Thus, average sources are slightly brighter than one might expect by just using the mid-band frequency of 3 GHz (0.601 times the 1.4 GHz flux density, versus the 0.587 factor we found above).

A more careful calculation involves using optimal weights as well as a more realistic model for the noise part of the signal-to-noise ratio. There are two effects: the system noise which for JVLA

gets worse at the lower frequencies, and the effective mosaic time-on-source which gets better at the lower frequencies where the primary beam is larger. At this time no detailed calculation can be done due to lack of system noise curves, and thus for the purposes of this proposal we assume those effects cancel out, leaving a roughly uniform $\sigma(\nu)$.

2.6.2 Source Angular Size and Surface Brightness Sensitivity

To predict the sensitivity of the VLASS to extended sources such as star forming galaxies and extended AGN-related emission regions, we have to take into account the match between the resolution of the VLASS and the size of the source.

Surface brightness in the radio regime is typically given in Rayleigh-Jeans equivalent brightness temperature T_b (in K), or in flux density per unit solid angle S/Ω (in MJy/sr or mJy/sq.arcsec). These are related by

$$T_b = \frac{c^2}{2k_b \nu^2} \frac{S}{\Omega} = 32.6 \text{ K} \left(\frac{S/\Omega}{\text{MJy/sr}} \right) \left(\frac{\nu}{1 \text{ GHz}} \right)^{-2} \quad (25)$$

where $k_b = 1380 \text{ Jy K}^{-1} \text{ m}^2$ is Boltzmann's constant in radio astronomy units (where brightnesses are typically reported as Jy/beam). For a Gaussian source of peak flux density S_p and angular (beam or source) size full-width half-maximum θ_F ,

$$\Omega = \frac{\pi}{4 \ln 2} \theta_F^2 = 1.1331 \theta_F^2 \quad (26)$$

we get

$$T_b = 1223 \text{ K} \left(\frac{S_p}{1 \text{ mJy}} \right) \left(\frac{\theta_F}{1''} \right)^{-2} \left(\frac{\nu}{1 \text{ GHz}} \right)^{-2}. \quad (27)$$

Thus, the 1.4 GHz FIRST survey flux density limit of $150 \mu\text{Jy}/\text{beam}$ ($\theta_F = 5.4''$) corresponds to a $T_b = 3.21 \text{ K}$ limit, while the 3 GHz VLASS ALL-SKY limit of $69 \mu\text{Jy}/\text{beam}$ ($\theta_F = 2.5''$) corresponds to $T_b = 1.50 \text{ K}$. Where brightness temperatures and limits are given in the proposal, they were calculated using the above formulation.

However, the emission targeted by the VLASS is either not thermal in origin, or does not follow a Rayleigh-Jeans relation. The predominant extragalactic radio sources of interest have a non-thermal synchrotron emission spectrum, while thermal bremsstrahlung emission from the Galaxy typically shows a flat spectrum in the radio bands. The power-law spectrum for these cases is discussed in the previous sub-section and given by Equation 17. A Rayleigh-Jeans thermal spectrum is equivalent to a power-law flux density spectrum with $\alpha = 2$. Likewise, a non-thermal power-law spectrum is equivalent to a power-law thermal spectrum

$$T_b(\nu) = T_b(\nu_0) \left(\frac{\nu}{\nu_0} \right)^\beta \quad \beta = \alpha - 2 \quad (28)$$

with $\alpha = -0.7$ corresponding to $\beta = -2.7$.

For the VLASS, it is most convenient to describe radio sources with these spectra with some total flux density (e.g. mJy) and characteristic angular size, rather than some peak brightness (e.g. mJy/beam). To quantify the yield of the VLASS for such science, the brightness limits of the survey to these sources must be calculated. To do this, the method used for finding sources must also be factored in.

The most straightforward approach is to look for peaks above some threshold, typically above some peak brightness limit, e.g. 5σ , set by the rms image noise level. The PSF width θ_F (a Gaussian "synthesized" or "restoring" beam FWHM used in the imaging) is convolved with the sky

brightness of the source, which then appears more extended. The peak brightness is thus reduced by the apparent area. A Gaussian source of total flux density S and intrinsic FWHM ϕ_F will show up in the image with a peak brightness

$$S_p = S \left(\frac{\theta_F^2}{\phi_F^2 + \theta_F^2} \right) \quad (29)$$

in usual units (Jy/beam etc.). Thus, a face-on (circular) star-forming galaxy disk at $z \sim 2$ might have an apparent angular size of $\phi_F = 3.1''$ (5 kpc), which will appear in the DEEP fields ($\theta_F = 0.8''$) with peak brightness $S_p/S = 0.0624$. This reduction, a factor of 16 in this case, biases the high-resolution images against finding more nearby or extended galaxies. Note that for ALL-SKY, with $\theta_F = 2.5''$, the penalty in this case is only $S_p/S = 0.394$ (disregarding whether such a galaxy is actually bright enough to detect in ALL-SKY). These are the factors used in Figure 1 and Table 2 of the main proposal to calculate the source yields. Note that more edge-on disks will not be as adversely affected, with factors approaching the square-root of those calculated above.

This angular size penalty is harsh, going as the square of the sizes of the PSF to the convolved source size. This, however, can be partly mitigated by employing matched filtering in the source finding. In the example above, the input image could be convolved with a Gaussian of size θ_c , where $\theta_c^2 \approx \phi_F^2$ producing a finding image with resolution matched to the angular size of the source and detecting the full flux density S . However, there is a penalty in the noise of the convolved map (after renormalization to conserve flux) with the increase in the convolved image variance given by the inverse of the peak dilution factor above, or the square root of this for the rms noise

$$\sigma_c = \sigma \left(\frac{\phi_F^2 + \theta_F^2}{\theta_F^2} \right)^{1/2} \quad (30)$$

where σ is the rms noise in the input image. This filtering gives a net gain over the unfiltered peak finding algorithm, falling only roughly linearly with the ratio of source size to PSF rather than the square. Note that this image convolution filtering is equivalent to imaging with a uv-taper applied. In our above example, the reduction of S/N due to the filtered noise increase to measure a $3.1''$ galaxy would be $\sigma/\sigma_c = 0.250$ for DEEP and $\sigma/\sigma_c = 0.628$ for ALL-SKY. In principle, this could significantly increase the source yields over those given in the proposal. There is certainly some limit to how large a source can be recovered this way, based on the (wideband) uv-coverage of the array. This simple calculation is probably applicable to source sizes out to ten times the PSF size.

The price of this matched filtering is the added complexity to devise and use a multi-scale source finder. Development and testing of this is probably beyond the scope of our Test & Development plan unless contributed effort is provided by the community. If such a source finder were provided and worked efficiently, it could be used for the final catalogs, most likely as a Enhanced Data Product (§ 4.2).

Note that the most effective approach to increasing the sensitivity to extended (few arcsecond) sources, in the DEEP observations in particular, would be to observe roughly 25% of the time in a more extended configuration such as the C-configuration ($8.6''$ resolution at 3 GHz) — this would essentially reduce the penalty in signal-to-noise ratio to only 0.5 in peak brightness for sources of angular sizes out to around 3 arcseconds and 0.25 at 9 arcseconds, even without matched filter source-finding. This should be considered as an option for the VLASS DEEP if extra sensitivity to faint extended sources becomes a strong science driver.

2.6.3 Measurement of the Source Spectrum

One of the science goals of the VLASS is to measure the frequency spectrum of source emission across the 2–4 GHz band. This can be done using various techniques: by imaging the data in

spectral channels across the band and fitting source spectra to the cubes, by uv-plane model fitting to the channelized visibilities, or by constructing multiple Taylor-term images using “Multi-frequency Synthesis” (MFS) as a part of the continuum imaging. We plan to use the last of these as a main data product of VLASS, and we now derive the sensitivity limits to this measurement. Note that the other approaches should have comparable performance.

In MFS, one uses a frequency dependent gridding kernel to appropriately combine the visibilities to form images corresponding to the Taylor expansion of the frequency spectrum. We now analyze a toy model for this process that represents the salient characteristics of the algorithm. We assume the source follows a power-law spectral energy distribution (Equation 17). The Taylor expansion of the spectrum with α as a function of frequency is

$$S(\nu, \alpha) = S_0 + \alpha \left. \frac{dS}{d\alpha} \right|_{\alpha=0} + \dots \quad (31)$$

where for our power-law

$$\frac{dS}{d\alpha} = \frac{d}{d\alpha} S_0 e^{\alpha \ln(\nu/\nu_0)} = S(\nu) \ln(\nu/\nu_0). \quad (32)$$

The series expansion of the logarithm is

$$\ln(\nu/\nu_0) = \left(\frac{\nu}{\nu_0} - 1 \right) - \frac{1}{2} \left(\frac{\nu}{\nu_0} - 1 \right)^2 + \dots = \left(\frac{\nu - \nu_0}{\nu_0} \right) - \dots \quad (33)$$

Thus, to first order, expanding about $\alpha = 0$,

$$S(\nu, \alpha) = S_0 \left[1 + \alpha \left(\frac{\nu - \nu_0}{\nu_0} \right) + \dots \right] \quad (34)$$

and it is apparent if we integrated these first two terms over the frequency band setting for convenience $\nu_0 = 0.5(\nu_{min} + \nu_{max})$ and $\Delta\nu = (\nu_{max} - \nu_{min})$, we find

$$M_0 = \frac{1}{\Delta\nu} \int_{\nu_{min}}^{\nu_{max}} d\nu S(\nu, \alpha) \quad (35)$$

$$\simeq \frac{1}{\Delta\nu} \int_{\nu_{min}}^{\nu_{max}} d\nu S_0 \left[1 + \alpha \left(\frac{\nu - \nu_0}{\nu_0} \right) \right] \quad (36)$$

$$= S_0 \quad (37)$$

as the integral over the kernel in the second term to zero. Likewise, we can form the next order integral

$$M_1 = \frac{1}{\Delta\nu} \int_{\nu_{min}}^{\nu_{max}} d\nu \left(\frac{\nu - \nu_0}{\nu_0} \right) S(\nu, \alpha) \quad (38)$$

$$\simeq \frac{1}{\Delta\nu} \int_{\nu_{min}}^{\nu_{max}} d\nu S_0 \left[\left(\frac{\nu - \nu_0}{\nu_0} \right) + \alpha \left(\frac{\nu - \nu_0}{\nu_0} \right)^2 \right] \quad (39)$$

$$= \frac{\alpha S_0}{\Delta\nu} \int_{\nu_{min}}^{\nu_{max}} d\nu \left(\frac{\nu - \nu_0}{\nu_0} \right)^2 \quad (40)$$

$$= \frac{\alpha S_0 \nu_0^2}{\Delta\nu^2} \int_{-1/2}^{1/2} dx x^2 \quad x = \frac{\nu - \nu_0}{\Delta\nu} \quad (41)$$

$$= \frac{\alpha S_0 \nu_0^2}{12 \Delta\nu^2} \quad (42)$$

where in this case the first term vanishes in the integral. Thus, the unity weighted integral M_0 (Taylor term “tt0”) measures the intensity S_0 , while the frequency weighted integral M_1 (Taylor term “tt1”) measures αS_0 , in particular

$$\alpha \approx \frac{12 \Delta v^2}{v_0^2} \frac{M_1}{M_0}. \quad (43)$$

Note that there are higher order corrections to this, so one needs more careful simulation testing on MFS recovery of spectral index, but this formula captures the essence of the method.

The next ingredient is the error estimates on these integrals. We assume that the S in the integral is now a measurement with a uniform (zero-mean Gaussian) noise given by the spectral noise power density σ_v . Thus,

$$\sigma_{M0}^2 = \frac{1}{\Delta v^2} \int_{v_{min}}^{v_{max}} dv \sigma_v^2 \quad (44)$$

$$= \frac{\sigma_v^2}{\Delta v} \quad (45)$$

$$\equiv \sigma_I^2 \quad (46)$$

where we define σ_I as the (measured) rms intensity image noise, and

$$\sigma_{M1}^2 = \frac{1}{\Delta v^2} \int_{v_{min}}^{v_{max}} dv \left(\frac{v - v_0}{v_0} \right)^2 \sigma_v^2 \quad (47)$$

$$= \frac{\Delta v_0^2}{12 v_0^2} \frac{\sigma_v^2}{\Delta v} \quad (48)$$

$$= \frac{\Delta v_0^2}{12 v_0^2} \sigma_I^2 \quad (49)$$

Using standard propagation of errors on the ratio in Equation 43, we get

$$\frac{\sigma_\alpha^2}{\alpha^2} \approx \frac{\sigma_{M0}^2}{M_0^2} + \frac{\sigma_{M1}^2}{M_1^2} \quad (50)$$

$$= \frac{\sigma_I^2}{S_0^2} \left[1 + \frac{\Delta v^2}{12 v_0^2} \left(\frac{12 v_0^2}{\alpha \Delta v^2} \right)^2 \right] \quad (51)$$

$$= \frac{\sigma_I^2}{S_0^2} \left[1 + \frac{12 v_0^2}{\alpha^2 \Delta v^2} \right] \quad (52)$$

$$(53)$$

or

$$\sigma_\alpha^2 \approx \frac{\sigma_I^2}{S_0^2} \left[\alpha^2 + \frac{12 v_0^2}{\Delta v^2} \right]. \quad (54)$$

Since $\alpha^2 < 1$ for the usual range $-1 < \alpha < 1$, the second term dominates. Thus,

$$\sigma_\alpha \approx \frac{\sigma_I}{S_0} \sqrt{12} \frac{v_0}{\Delta v}. \quad (55)$$

For the ideal full S-band VLASS $\Delta v = 2$ GHz, $v_0 = 3$ GHz, and thus

$$\sigma_\alpha \approx 5.2 \frac{\sigma}{S_0} \quad (56)$$

and thus in order to obtain a spectral index uncertainty of $\sigma_\alpha < 0.1$ a source bright enough in intensity such that $S_0/\sigma_I > 52$ is needed, or approximately 10 the 5σ noise cutoff on object finding.

3 Survey Structure

The VLASS science case is comprehensive and motivates a dual-tiered survey: ALL-SKY (Tier 1) at $69 \mu\text{Jy}$ rms ($120 \mu\text{Jy}$ per epoch in each of 3 epochs), and DEEP (Tier 2) at $1.5 \mu\text{Jy}$ rms. The observing time required, including overhead, is 5436 hours for ALL-SKY, and 3391 hours for DEEP. **The total scheduled observing time in VLASS, including overhead, is 8827 hours for our proposed 7-year 6-cycle schedule.**

3.1 Tier 1: ALL-SKY

The ALL-SKY tier covers the entire sky visible to the VLA north of Declination -40° , 33885 square degrees (for 82.14% of the sky). The observations for ALL-SKY will be carried out using On-The-Fly Mosaicking (OTFM). The scan rate will be adjusted depending on declination (elevation) to provide uniform sensitivity, requiring an extra 9% net integration time over the entire area. A given area of sky will be passed over three times during the survey duration, making this a 3-epoch transient survey at $120 \mu\text{Jy}/\text{beam}$ rms per epoch. This uses the fastest scanning rate possible at full data rate (see § 2.4). The ALL-SKY area includes the areas of the Galactic plane and bulge visible to the VLA.

The summary statistics of the ALL-SKY tier are:

- B-configuration ($2.5''$ resolution)
- total visible sky area of 33885 square degrees, $\delta > -40^\circ$
- per epoch continuum image rms (Stokes I) $\sigma_I \geq 120 \mu\text{Jy}/\text{beam}$
- net survey speed is $23.832 \text{ deg}^2/\text{hr}$ at this depth
- effective integration time required is 1421.8 hours per epoch (4265.4 hours total for 3 epochs)
- multiply by factor 1.09 for sensitivity loss at low elevation for low declinations (§ 2.1.2), requiring 4649.3 hours total true integration to reach uniform sensitivity over entire area
- total of 5440 hours observing required with 17% overhead
- observations are spread through the duration of the survey comprising 6 configuration cycles, rounded to 5436 hours total (906 hours per cycle)
- a given area of sky is observed in 3 cycles (each cycle half the ALL-SKY area is observed), mean cadence is 32 months between epochs

Special Considerations: The data will be taken in a single-pass each, near the maximum slew rate practical for OTFM at the maximum allowed data rate. Thus, there is no opportunity for observing at two hour angles for improved snapshot uv coverage. For the continuum observations, the MFS uv-coverage improvement should partly compensate for this in the imaging quality. This will be verified and quantified as part of the test plan.

Use of BnA Hybrid: If practical, we would like to observe the southern area ($\delta < -10^\circ$) using the BnA hybrid to improve the point spread function. This would require observing 3565 total hours in B-configuration, and 1871 total hours in BnA. This is around 624 hours per epoch, or if split evenly among all the six configuration cycles (§ 5.4), 312 hours per cycle. Note that the time is more heavily weighted towards BnA than the relative fractional area due to the compensation for sensitivity loss at low elevations which impacts VLASS observations only in the low-declination region.

3.2 Tier 2: DEEP

The DEEP tier comprises observations of 3 deep fields with multi-wavelength coverage in a total area of 10 square degrees. The fields are: COSMOS (2 square degrees, 10h00m28s, 02d12m21s), ECDFS (4.5 square degrees, 03h32m28s, -27d48m30s), and Elais-N1 (3.5 square degrees, 16h08m44s, 56d26m30s). The cumulative depth for each field will be an image rms (Stokes I) of $1.5 \mu\text{Jy}$ in A-configuration. There is a multiplicative factor of 1.3 applied to the integration time needed for the ECDFS, to compensate for its low elevation assuming observations within 2 hours of transit (§ 2.1.2). For COSMOS, there is an existing $2 \mu\text{Jy}$ S-band C plus A-configuration dataset (12B-158, Smolcic et al.). Note that although these fields are all included in ALL-SKY, the amount of time spent in that tier on the deep fields (< 1 hour total per square degree) is negligible compared to the total required for DEEP, so we treat DEEP as independent. These smaller fields can be efficiently mapped using pointed mosaics, so OTFM need not be used for the deep fields unless it is desirable to cover these multiple times in a single day for more uniform uv-coverage and sensitivity to fast transients and variability.

The total scheduled observing time in DEEP is 300 hours for COSMOS, 1960 hours for ECDFS, 1131 hours for Elais-N1, for a total of **3391 hours**.

3.2.1 Tier 2 DEEP — COSMOS

The summary statistics for Tier 2 COSMOS are:

- total area of 2 square degrees, centered at 10h00m28s, +02d12m21s
- previous observations in 2013-14 (12B-158) to depth of $2 \mu\text{Jy}$ rms equivalent to 302 hours of integration
- combined continuum image rms (Stokes I) $\sigma_I = 1.5 \mu\text{Jy}$ (added to previous data from 2012-14)
- total on-sky time of 234 hours (clock time 292 hours with 25% overhead)
- combined continuum image rms (Stokes I) $\sigma_I = 2.27 \mu\text{Jy}$ (this data only)

We can observe for 6 hours centered near transit without suffering elevation induced sensitivity losses, so a plausible scenario is to break the observations into 7.5 hour blocks (25% overhead). For 300 hours total, we need 40 blocks. To spread the observing evenly for transient science, we observe this in five cycles:

- each pass is 7.5 hours observing each (6 hours each integration time per pass)
- observe in five A configuration epochs, 8 passes per cycle
- single-epoch continuum image rms (Stokes I) $\sigma_I = 5.0 \mu\text{Jy}$
- single-pass continuum image rms (Stokes I) $\sigma_I = 14 \mu\text{Jy}$
- total scheduled time 300 hours (60 per epoch)

3.2.2 Tier 2 DEEP — ECDFS

The summary statistics for Tier 2 ECDFS are:

- total area of 4.5 square degrees, centered at 03h32m28s, -27d48m30s
- combined continuum image rms (Stokes I) $\sigma_I = 1.5 \mu\text{Jy}$

- this would normally take 1208 hours of integration at optimal elevation
- if observed within 2 hours of transit, extra overhead of 1.3 to reach stated sensitivity for elevation effect (§ 2.1.2)
- total on-sky integration time of 1570 hours (1962 hours observing calculated using 25% overhead, rounded to 1960 hours for scheduling convenience)

Since observations should occur within 2 hours of transit, we break the observations of this field into 5 hour blocks (each should get 4 hours integration, 25% overhead, see § 2.3). We will need 392 such blocks total. A plausible scenario is to spread these over six cycles of A configuration:

- six A configuration epochs
- 65 passes of 5 hours each (325 hours) per cycle for cycles 1–5, final cycle get 67 passes (335 hours)
- single-epoch continuum image rms (Stokes I) $\sigma_I = 3 \mu\text{Jy}$
- single-pass continuum image rms (Stokes I) $\sigma_I = 29.7 \mu\text{Jy}$
- total scheduled time is 1960 hours

Note that by observing closer to transit (e.g. within 1.5 hours) one could decrease the elevation losses incurring less overhead (1.28 increase). However, 3.75 hour blocks would require 512 passes total, which would require 85 passes per cycle. This is likely to have serious logistical problems in scheduling, and thus we choose to take an extra 2% loss in efficiency to observe longer blocks and get back around 8% in overhead which is a net win.

We note that the efficiency loss due to the low elevation has been estimated from a plot of system temperature versus antenna elevation angle at various frequencies in S-band (courtesy R. Perley). Our estimate is uncertain at the few percent level, and could be either too lenient or too conservative. Further testing on the VLA through a trial observation of the ECDFS (e.g. a single pass) should be scheduled as soon as is practical.

3.2.3 Tier 2 DEEP — Elais-N1:

The summary statistics for Tier 2 Elais-N1 are:

- total area of 3.5 square degrees, centered at 16h08m44s, 56d26m30s
- combined continuum image rms (Stokes I) $\sigma_I = 1.5 \mu\text{Jy}$
- total on-sky time of 937 hours (1143 hours observing calculated for 22% overhead, rounded to 1131 hours for scheduling of 7.25 hour blocks, see below)

This field can be observed for long blocks due to its high declination. However, we restrict to around 6 hour integration per block so that they better fit in with dynamic scheduling (but can be observed over a wide range of start times). These blocks will be 7.25 hours long with overhead (in this case 21%), and thus we need 156 blocks total (giving 936 total hours integration, just under our target) for 1131 total hours. Again we split these into four cycles:

- six A configuration cycles
- each cycle 26 passes of 7.25 hours observing each (188.5 hours per epoch)
- single-epoch continuum image rms (Stokes I) $\sigma_I = 3.7 \mu\text{Jy}$
- single-pass continuum image rms (Stokes I) $\sigma_I = 18.7 \mu\text{Jy}$
- total scheduled time is 1131 hours

4 Data Products

The VLASS data products are described below, along with risks and considerations in the production and assessment of these products.

Products are broken into classes of “Basic” and “Enhanced”, with the former being simple enough to produce by NRAO’s standard (or soon to be standard) data processing system. Enhanced data products will require domain expertise and/or extra resources, so they will be left for community members to define and produce. Both kinds of products will be curated and served to the public by the NRAO, as described below.

4.1 Basic Data Products

The Basic Data Products (BDP) of the VLASS consist of:

1. raw visibility data
2. calibration data and process to generate calibration products (current best version as well as past released versions maintained in archive)
3. quick-look continuum images
4. single-epoch images and image cubes
5. single-epoch basic object catalogs
6. cumulative “static sky” images and image cubes (generated after epoch beyond the first)
7. cumulative “static sky” basic object catalogs (generated after epoch beyond the first)

The resources for processing, curating, and serving the data products will be provided by NRAO, as described in §4.3. Teams led by NRAO, but including external community members where possible, will carry out the activities required for the processing and Quality Assurance (QA) of the products. Table 5 summarizes the BDP.

Table 5: Summary of Basic Data Products

Product	Timescale	Notes
Raw visibility data	immediate	in standard archive
Calibrated data	1 week	from standard archive
Quick-Look Images	48 hrs	I continuum only
Single-Epoch Images	6 mos. (I)	12mos. for pol. and spectra
Single-Epoch Catalogs	w/Single-Epoch Images	
Cumulative Images	12 mos. (I)	16mos. for pol. and spectra
Cumulative Catalogs	w/Cumulative Images	

Note that the polarization continuum images and the spectral cubes for the single-epoch and cumulative data products will be delivered on longer timescales than the Stokes I continuum intensity images.

Details of individual BDP are now described.

4.1.1 Raw Visibility Data

The raw visibility data for the VLASS will be stored in the standard VLA archive. These data will be available immediately after observation with no proprietary period. As the VLASS is being observed using standard data rates (25 MB/s maximum) there are no special resources required for the storage and distribution of these data. Data will be downloadable by users from the archive web pages as normal.

Risk: None. These is standard data comprising a volume that is normal for the amount of time spanned by the VLASS.

4.1.2 Calibrated Data

VLASS data will be processed using a modified version of the normal CASA-based VLA calibration pipeline. By the time of the VLASS observations, the VLA pipeline will have the requisite functionality to process VLASS data (full polarization, many individual target fields generated in OTF mode). The VLA pipeline is currently run on all VLA observations and is well tested. The VLASS will use a version specifically tested on VLASS pilot observations. See the description of the Calibration Pipeline below for more details.

In summary, the output of the pipeline will consist of:

- a set of final calibration tables
- (possibly) a set of sky models used in calibration
- one or more sets of flagging commands
- the final flag column
- QA reports and plots
- the set of pipeline control instructions or script necessary to calibrate the raw data

Once the pipeline has run and after QA assessment, users will gain access to the calibrated data from the archive. Physically archiving the calibrated visibility data for long-term storage would double the amount of archive space required, which is too costly. Instead, the VLASS will archive calibration products and maintain scripts to generate the calibrated dataset. Scripts to apply calibration products will be based on CASA. Upon first processing of a Scheduling Block the calibrated data will be available for a short time (one to two weeks) for transfer to subsequent processing teams.

Risk: Low. The VLA Calibration Pipeline is in regular current operation, and its code is downloadable by users. We are currently upgrading this pipeline to use the infrastructure developed for the ALMA pipeline. Thus, by the time of the VLASS, this pipeline will have been in regular use and extensively tested. The addition of polarization calibration capability to the pipeline is the only significant addition, and this is expected to be available in the next year also.

4.1.3 Quick-Look Images

The identification of transient and variable objects is a key science goal of the VLASS. This requires the capability to process and image the data within a short period after the data is taken. We have set the requirement on this to be 48 hours maximum, with a goal of 24 hours. If possible we will use the standard VLA Calibration Pipeline described above in a streamlined mode (highly parallelized where appropriate). As a fallback we can use a separate pipeline based on different software, such as the AIPS-Lite based pipeline developed and used by the Caltech group (PI: Hallinan) used for the S-band Stripe-82 observations taken and processed by that group (VLA Project 13B-370).

The output of the Quick-Look Calibration and Imaging Pipeline (QLP) will be continuum images (Stokes I) that can be used by the Transient Object identification teams (see below under Enhanced Data Products) to produce alerts for transient objects and arrange for follow-up studies. As these images will be constructed through mosaicking, there will also be a corresponding rms sensitivity (noise) image. Polarization images (Stokes QUV) are not part of the Quick-Look Basic Data Products.

There are thus a total of 2 images to potentially be produced by the QLP per trigger: an I image, and the corresponding noise map. As detailed in § 9.4 these images take 35TB and can easily be stored and served from the archive.

To manage the computational load, we do not plan to produce the higher-order MFS term images for spectral index ($\tau\tau1$) or curvature ($\tau\tau2$) per epoch. These could be made available as EDP. We do not plan to produce any image cubes in the QL pipeline as a BDP due to the high data volumes possible. This could also be part of a EDP proposal, most likely in a “processing on demand” (POD) service offsite.

The BDP Quick-Look images will be publicly available from the VLASS archive promptly upon production.

Risk: Medium Production of Quick-Look images will require prompt execution of a Calibration Pipeline, and if the main pipeline (see above) is not fast enough a separate one will need to be developed. Personnel will be needed to oversee running of the QLP and monitor this output, including astronomers and data analysts. Data storage needs are modest.

4.1.4 Single-Epoch Images

Within 6 months of completion of the observations of a given “epoch” (observations in a given configuration), fully calibrated and quality-assured continuum intensity images as detailed below will be produced and available in the archive. These will be generated using a specialized CASA-based Imaging Pipeline. This pipeline will be developed by the NRAO staff, with the involvement and guidance of the VLASS teams. The polarization images and the intensity and polarization cubes will be produced within 12 months of the end of an epoch. This delay is due to the added complexity and processing required.

The VLASS single-epoch wide-band continuum images will include:

1. Flux density calibrated beam-corrected Stokes I continuum intensity (band averaged) images covering the full mosaic area
2. Sensitivity (rms noise) images for the intensity I continuum image mosaic (the noise will vary over the mosaic)
3. Spectral Index and uncertainty images for Stokes I (generated using Multi-Frequency Synthesis)
4. Flux density calibrated beam-corrected Stokes QU continuum polarization (band averaged) images covering the full mosaic area
5. Sensitivity (rms noise) images for the QU polarization continuum images (the polarization is represented by a complex $Q+iU$ image, and thus Q and U share the same rms noise image)

There are thus 7 single-epoch continuum image products (I, I_{rms} , I_{α} , $I_{\alpha_{rms}}$, Q, U, pol_{rms}). We do not plan to produce V (circular polarization) images for archiving, as these require different (and more expensive) processing options deal with instrumental effects.

Although the wide-band MFS imaging can produce spectral index (I_{α}) images, we suggest that instead of alpha, that the MFS “ $\tau\tau1$ ” (Taylor-term order 1) image be natively stored, as this is

effectively I^* and will not have the zeroes that come from dividing by I in noisy regions. The archive server can produce a masked α image from the native I and $tt1$ images.

In addition to continuum images, there are spectral image cubes of up to 1024 channels (at full spectral resolution) for each of:

1. Flux density calibrated beam-corrected Stokes IQU continuum per-channel image cubes covering the full mosaic area
2. Sensitivity (rms noise) cube images for I and polarization ($Q+iU$)

There are thus 5 image cube products.

The spectral cubes will contain a wealth of information on the SED and polarization necessary for deriving RM and other computed products. In addition, these will be important for QA purposes and diagnosing RFI problems. Thus they will need to be generated with the processing of a given epoch. However, we note that it will be expensive to store these long-term in the archive for each epoch. In § 9.5 below, we describe the storage requirements for the single-epoch image cubes and propose to only store and serve cubes averaged over frequency, e.g., in 14 128MHz spectral windows for DEEP, further compressed (10:1, likely through image cutouts around bright sources only) cubes for ALL-SKY. These cubes will still allow continuum polarimetry and rough rotation measure fitting to bright sources. However, full rotation measure synthesis will not be possible. The cumulative images cubes will be used for this science goal.

Given these limitations, it is attractive to consider a POD on-the-fly processing service for these cubes for ALL-SKY, as described above for the QL continuum images (§ 4.1.3), hosted on-site or externally (e.g. through XSEDE). In addition to efficiency, POD imaging will provide flexibility for choosing coarser frequency resolution in cubes, and the ability to control whether beam-corrections are done or not. The single-epoch DEEP images are sufficiently deep and costly that we will assume that cubes are created once and stored in the archive. In practice we have found that the channels in two or more entire spectral windows with 64 channels each will be entirely flagged, and so we assume we need cubes for only 896 channels maximum.

We assume that only temporary storage is required for the large full-resolution cubes for QA purposes and compression (see § 7.4). For DEEP we will archive all of the per-epoch cubes.

Risk: Low to Medium. Currently, OTF imaging is done using special CASA scripts, with ongoing testing in development code. By the time of the VLASS, it is expected that all needed capability to produce the per-epoch images will be available and tested in a standard CASA release. The fallback is to continue to use development code. Processing for shallow to medium deep observations in ALL-SKY is of low risk. Single epoch images for the DEEP fields are of medium risk. Development of a robust POD service for Tier 1 cubes is of medium risk, with a fallback to long-term archiving of the cubes. Storage of these images is of medium risk, with a number of fallbacks available using compression or reduced resolution. Investigation of acceptable spectral and spatial compression schemes are part of the Test & Development Plan.

4.1.5 Single-Epoch Basic Object Catalogs

There are many object finders used in the community for identification and classification of objects from images. It is expected from surveys such as the VLASS that basic object catalogs be produced and released along with the images. Over the next year, we will carry out a study of the available object finders suitable for use on the VLASS data and select and test one (or more if necessary) for production use by the team.

We consider the Basic Object Catalog entry for an “object” to contain:

1. Position, and uncertainty (likely centroid of I emission)
2. Peak Flux Density (continuum) in IQU, and uncertainty

3. Spectral Index at Peak (Stokes I) and uncertainty
4. Integrated Flux Density (continuum) in IQU, and uncertainty
5. Integrated Spectral Index (Stokes I) and uncertainty
6. Basic Shape information IQUV (TBD)

For the linear polarization Stokes Q and U values, we can instead report these as the more traditional polarized flux density (amplitude of $Q+iU$) and polarization angle ($0.5 \tan^{-1}(U/Q)$).

It is also likely that for bright sources we will be able to compute IQU spectra (averaged over the source region). These would be extracted from the temporary full-resolution cubes.

Risk: Low. This has minimal requirements, and it is likely that object finder tools used for previous surveys such as NVSS, FIRST, and the Stripe-82 surveys will be sufficient, if not ideal. It is likely that the studies will identify a superior tool.

4.1.6 Cumulative Images

Within 12 months of the completion of observations of each epoch after the first (16 months for polarization and cubes), continuum images for the cumulative data will be made. These will be largely the same set of images as described above for the single epochs. Accommodation for variable objects in the imaging will be required across the multiple epochs.

The VLASS cumulative wide-band continuum images will include:

1. Flux density calibrated beam-corrected Stokes IQU continuum (band averaged) images covering the full mosaic area
2. Sensitivity (rms noise) images for the I and linear polarization ($Q+iU$) continuum images
3. Spectral Index and uncertainty images for Stokes I (generated using Multi-Frequency Synthesis)
4. Spectral Curvature and uncertainty images for Stokes I (generated using Multi-Frequency Synthesis)

There are thus 9 cumulative continuum image products. These are the images produced for the individual epochs (§ 4.1.4), with the addition of the next Taylor Term (tt2) in the MFS expansion for spectral curvature, plus its uncertainty.

In addition to continuum images, there are cubes of up to 1024 channels for each of:

1. Flux density calibrated beam-corrected Stokes IQU continuum per-channel image cubes covering the full mosaic area
2. Sensitivity (rms noise) images for each of the cube images in I and $Q+iU$

There are thus 5 cumulative image cube products.

The same considerations apply to the cumulative image cubes as for the per-epoch cubes (see § 4.1.4). Storage requirements are described below in § 9.7. The archive will store and serve only the most recent cumulative image cube for each Tier. If the POD capability is available, users can request images from previous cumulative epochs (or possibly other selections). Note that the storage needs for these full-resolution cubes is considerable, and fallback options to reduced resolution or compressed cubes will be necessary to implement in the absence of enhanced archive resources. See the discussion of the required compression in § 9.7.

Risk: Low to High. The cumulative ALL-SKY image rms will be only 58% of the single epoch image rms and thus low risk outside the Galactic plane, and has medium risk in the confused

Galactic plane, bulge, and center region. DEEP image production will be challenging (high risk), and some algorithm development will need to take place in order to make automated deep imaging (such as that currently manually done on deep fields) work with larger mosaics and higher dynamic ranges required for the Deep tier. This could cause delays in production of final images. Development of a robust POD service for the cumulative data is highly desirable, but not critical as only the most recent cumulative images are of high scientific priority. Data storage for full-resolution cubes is of high (resource) risk. Fallback to reduced resolution or spatial coverage will have to be implemented in case additional archive resources (e.g. through EDP) are not identified.

4.1.7 Cumulative Basic Object Catalogs

As in the case of the single-epoch catalogs, these basic catalogs would be made available along with the cumulative images. These are largely based upon the continuum images, we do not foresee using the cubes for these.

We consider the Basic Object Catalog entry for an “object” in the cumulative catalogs to contain:

1. Position, and uncertainty (likely centroid of I emission)
2. Peak Flux Density (continuum) in IQU, and uncertainty in I, Q+iU
3. Spectral Index at Peak (Stokes I) and uncertainty
4. Spectral Curvature at Peak (Stokes I) and uncertainty
5. Integrated Flux Density (continuum) in IQU, and uncertainty in I, Q+iU
6. Integrated Spectral Index (Stokes I) and uncertainty
7. Basic Shape information IQU (TBD)
8. Spectrum in IQU integrated over the source region (for bright sources)

Risk: Low to Medium. The main risk is in the imaging. All-sky and Wide-Extragalactic are low risk. Will need to handle increased object density and complexity for the deeper tiers and Galactic regions (medium risk).

4.2 Enhanced Data Products

The Enhanced Data Products (EDP) are those that require more domain expertise, and so will be defined and produced by the VLASS community outside the NRAO. These data products will require external support to define, produce, and validate. However, these products are seen as essential to the VLASS science case, so both BDP and EDP will be curated and served by the NRAO.

The initial list of possible EDP are:

1. Transient Object Catalogs and Alerts
2. Full Spectral Resolution Image Cubes
3. Rotation Measure Images and Catalogs
4. Improved Object Catalogs
5. Matched filter source finding for extended objects

6. Light curves (intensity and polarization) for objects and/or image cutouts
7. Catalogs of multi-wavelength associations to VLASS sources

Beyond requiring extensive domain expertise, these areas are also ideal for nucleating multi-wavelength community groups and resources to work with and enhance the VLASS.

A special case of an EDP is the support for commensal observing at P-band (230–470MHz) using the VLITE system. There is no NRAO processing or archive support currently budgeted for VLITE data products, and thus use of VLITE with VLASS should be considered as an EDP (and Enhanced Data Services, EDS) provided in partnership with NRL.

Other areas for EDP will undoubtedly become apparent. Once the survey is approved, we will take proposals for new EDP to be included in the list above. Criteria for including EDP in the VLASS archive will be relevance to the VLASS science case and cost of curating and serving the products. As an incentive to include EDP in the VLASS archive, we ask that the NRAO encourage authors that use EDP to acknowledge groups that produced them.

4.3 Enhanced Data Services and the VLASS Archive

A comprehensive survey like the VLASS will produce a diverse set of data products and will require a full-featured archive to serve it to the public. A baseline plan is to serve products from a website hosted by the NRAO. This site will feature basic search capabilities of catalogs and products, as has been done for FIRST⁷ and NVSS⁸. Previous VLA surveys only provided catalogs and images, so at a minimum the VLASS will extend that search capability to visibility data, calibration products, and deep/multi-epoch images. The NRAO will also provide data analysis scripts to apply calibration to raw data.

However, astronomy is increasingly a multi-wavelength discipline with a diverse set of tools for comparing observations from different observatories. If the VLASS archive exists only as a stand-alone NRAO-hosted service, it would not be as useful as one integrated with the tools available at places like IPAC⁹ or the Virtual Observatory¹⁰.

We are investigating options for having catalogs and/or images served by organizations outside the NRAO. This would extend the reach of the VLASS outside the radio community and open access to powerful tools for multi-wavelength analysis. The VLASS community, including the co-authors of the VLASS proposal, will be writing an NSF proposal to support VLASS data analysis and a more effective archive. These Enhanced Data Services (EDS) will greatly augment the utility of the VLASS and its basic and enhanced data products to the wider astronomical community.

Another area that would be greatly improved through EDS is the capability for “processing on demand” (POD) of images or image cubes. This would alleviate storage volume concerns, and enable more flexible angular and spectral resolution of the resulting products. We expect to utilize the NSF XSEDE network for modest use of POD-like processing for the pipeline. Fully enabled POD for VLASS could be carried out through partnerships with NSF supercomputing centers or with DOE science labs. Exploration of these options will commence upon approval by NRAO for VLASS.

As noted above, data archive and distribution support for commensal observing at P-band (230–470MHz) using the VLITE system is not currently budgeted for support by NRAO. Archive serving of commensal VLITE data taken at the same time as the VLASS should be considered an EDS provided in partnership with NRL.

⁷<http://sundog.stsci.edu/cgi-bin/searchfirst>

⁸<http://www.cv.nrao.edu/nvss/postage.shtml>

⁹<http://www.ipac.caltech.edu>

¹⁰<http://www.us-vo.org>

4.4 Data Formats

We now describe the formats that we expect to use and serve as data products from the VLASS.

4.4.1 Visibility Data

The VLA archive holds the raw data in the native SDM format. Users can request CASA MS format data that are produced from the SDM by the archive server.

Users of VLASS can also request the calibration products (tables, flags, pipeline instructions) that can be used to apply to a raw SDM or MS downloaded from the archive.

Calibrated visibility data will be served in CASA MS format upon request through application of the tables, flags, and pipeline instructions.

4.4.2 Images and Image Cubes

The most common image format in astronomy is FITS. CASA can recognize and use (for most operations) FITS images, although it has its own image format based on CASA tables.

The VLASS data services will need to provide the capability to serve FITS and CASA format images upon request for sub-areas (postage stamps, simple rectangular regions) from the available images.

The question of a “native” format in which to store images is more complex. The most straightforward approach would be to store a set of FITS (or CASA) images that tile the survey areas. Users could request sets of these directly, or ask for a region that would be assembled by the data server to one or more requested FITS or CASA images.

It might be advantageous to take more forward-looking approach and store the image data in a hierarchical data format such as HDF5 (e.g. in the context of LOFAR, Anderson et al. 2010, arXiv:1012.2266). For the all-sky Tier 1, representation in HEALPIX might be useful also. Investigation of options such as these should occur in the year leading up to the commencement of VLASS.

4.4.3 Catalogs

Catalogs could be stored in some internal manner using flat ASCII files, in a relational database (RDB), or in a hierarchical data format (such as HDF5).

Users must be able to obtain catalogs in simple formats such as flat ASCII files, XML files, or other basic formats.

4.4.4 Plots

Where the data services have the capability of providing plots, these should be in standard formats such as PNG, JPEG, GIF, or even FITS in some cases.

5 Observing

In order to carry out the VLASS, we will need to observe the sky using a large number of mosaicked pointings of the VLA. At 2–4 GHz, the VLA has a field-of-view given by the primary beam response of the 25-meter diameter antennas. This approximately follows a Gaussian response, with a full-width at half-maximum (FWHM) given by

$$\theta_{FWHM} \approx 45' \left(\frac{1 \text{ GHz}}{\nu} \right) \quad (57)$$

at observing frequency ν , and thus over the S-band the FWHM varies from 22.5' at 2 GHz to 11.25' at 4 GHz, with FWHM of 15' at 3 GHz mid-band. In order to optimally cover a given sky area in an efficient manner, the array must either conduct a raster scan using “on-the-fly” mosaicking (OTFM), or tile the area with a number of discrete pointings in a hexagonal packed configuration (“Hex-pattern Mosaicking”). The choice between these techniques is determined by the extent to which the extra overhead (from 3 to 7 seconds) needed to move the array and settle at each pointing in the discrete hex-pattern mosaic becomes a burden on the observations, and thus OTFM is favored.

5.1 Mosaicking

The techniques of OTFM and Hex-pattern Mosaicking and the calculations and procedures needed to set these up are described in the Guide to VLA Observing: Mosaicking¹ section. The salient features are:

OTFM: There is very little move-and-settle overhead as the array is in continuous motion over a row of a raster with only a small start-up (~ 10 – 15 sec) at the start of each row. In OTFM the phase center of the array is discretely stepped on timescales of a few seconds or longer, so no phase smearing of the images results. However, because the primary beam response pattern is moving with respect to the sky, there are errors introduced in the amplitudes by the moving beam in a single visibility integration time. Thus, the main cost of OTFM therefore is that for fast scanning rates the fundamental integration (“dump”) times in the dataset must be short (10% or less) compared to time it takes to cross the FWHM of the primary beam. This in turn increases data rates from those that would otherwise be required (e.g. to avoid time-smearing of the loci in the uv-plane). The secondary cost of OTFM occurs in the imaging process, where the effects of the moving primary beam over the time at which the phase center is fixed must be compensated for by the imaging algorithm at significantly increased computational cost over that required for a similar observation taken with a fixed pointing center. This is currently done by the CASA software package in its `clean` deconvolution task. Testing of the efficacy and efficiency of the use of the CASA imaging for OTFM is underway, and is part of the testing plan given below. For the purposes of this plan, we will assume that OTFM datasets can be imaged with sufficient accuracy to be practical for at least Tier 1 and 2 observations.

Hex-pattern Mosaicking: It takes the VLA 3–7 seconds (depending on the direction of motion in azimuth-elevation coordinates, usually around 6–7 seconds if not optimized) to move and settle between nearby pointings. Thus, if less than 28 seconds is spent integrating on each pointing, the overhead from this motion itself is 25% or higher in the worst case. For a hex-pattern, each field gets 67% of the total integration time desired on-sky, so observations where the VLA Exposure Calculator¹¹ indicates an on-source time of 42 seconds or less will incur significant overhead if not done with OTFM. For VLA S-band, the calculated exposure time is 7.7 seconds at a rms image sensitivity of $100 \mu\text{Jy}$, and thus observations desiring a depth shallower than around $43 \mu\text{Jy}$ in a single pass will prefer OTFM. Thus, our Tier 1 observations with single-pass depths $> 100 \mu\text{Jy}$ will require OTFM, while the deep Tier 2 observations can be carried out with straightforward Hex-pattern Mosaicking. Note that it is in principle possible to arrange the order in which the fields are observed to make the motions be predominantly in the elevation axis which damps the telescope settling, and has been shown to lose only 3–4 seconds. However, this optimization is dependent on the local time at which the observations are made requiring restricted LST scheduling blocks, and is difficult to arrange over large areas of the sky. We will try where possible to use this optimization to further reduce overheads in the Tier 2 mosaics.

Risks: Low to Medium Other than increased overheads, the Hex-pattern Mosaicking for Tier 2 does not incur significant risk. OTFM has been used successfully for shallow (single-pass image

¹¹<https://science.nrao.edu/facilities/vla/docs/manuals/propvla/determining/source>

rms $\sim 100 \mu\text{Jy}$) in S-band in the Stripe-82 observations by Hallinan et al. (program 13B-370). Images with rms $\sim 60 \mu\text{Jy}$ from the three passes spanning 2 months are being made as part of the testing program, and current indications are that they are of acceptable quality. Thus, for Tier 1, we deem the risk associated with using OTFM to be low overall on the basis of quality. The further consideration is the extent to which more expensive imaging costs are required in areas in which there are bright sources, particularly for polarimetry, and for accurate determination of fainter source spectral indices. This is also being assessed in the testing. As the new imaging algorithms are only recently available, we assign this as medium risk to resourcing for the VLASS, as extra computational power or longer processing times are likely to be required in some regions.

5.2 Scheduling Considerations

The Jansky VLA is normally operated using a “Dynamic Scheduling Queue” where the individual Scheduling Blocks (SBs) are created in the VLA Observing Preparation Tool (OPT) to be able to be executed in a prescribed range of LST, and submitted to the VLA Scheduler software (OST) to be queued up for observation by the array in a manner dictated by weather and priority. It is our intent that Tier 2 be observable using standard Dynamic Scheduling as much as possible. This requires that these blocks contain sufficient calibration to stand alone or be boot-strapped from other VLASS SBs executed nearby in time.

For Tier 1, it would be advantageous to construct the schedules in large blocks to be observed at specific LST start times. This would allow maximum efficiency in calibration and control of slewing (e.g., telescope wraps). In practice, due to considerations such as the ability to allow interrupts for target-of-opportunity observations, and fault tolerance (e.g., for power outages, weather, etc.), the schedule will need to be broken in to modestly sized blocks. Ideally, the VLA Scheduler software would be able to handle sets of SBs that are linked (e.g., in a particular order or in alternate sets). However, this capability does not currently exist, and there may not be resources in the software group to allow this. Instead, the most straightforward plan is to make sets of SBs for submission (see below) and submit only a day ahead. This will require an “Astronomer on Duty” (AoD) for VLASS who will keep track of what SBs are ready to observe, make sure they are submitted, make sure that they run, and make any modification necessary (e.g., due to TOO or weather interrupts).

Target-of-Opportunity Interrupts: For the long-block observations in Tier 1, provision will be made for the possibility that observations will be interrupted for time critical TOO programs (e.g., for triggered transient observations). There is currently no mechanical provision in the way schedules are constructed or executed for the suspension and restarting of schedule blocks. Therefore, the most straightforward implementation is to break all schedules into blocks of 2–3 hours in length, and to allow TOO interrupts to simply stop the execution of the current schedule and possibly pre-empt the execution of the following one or more SBs. After TOO observations are complete, the VLASS schedule would resume with the next appropriate block. The AoD would be informed of this interruption, and would examine the archive record to determine the missing observations and construct a “make-up” SB to be run at the first appropriate opportunity. The plan for the construction of schedules (see below) will take this need into account.

Risk: Medium The VLASS scheduling and observation monitoring will be a vast book-keeping exercise. As described above we believe this is controllable, either through some modest improvements in operations software, or at worst with some workarounds in the way we carry out the scheduling. There will need to be a pool of VLASS astronomers and data analysts to fulfill AoD staffing throughout the survey. This ideally might include students participating in the VLASS from the community. Some NRAO resourcing will be required to support these activities both directly and for supervision.

5.3 Schedule Construction

It will be impractical to use the VLA OPT in standard interactive mode to construct the thousands of hours of schedules for thousands of pointings that must be in the Source Catalog Tool. Instead, we will create some lightweight software in Python to construct the ascii lists that the OPT and SCT can read. This feature has been used by the 13B-370 Stripe-82 observers (mainly Caltech graduate student Kunal Mooley) with success to schedule those observations. Our plan is to use and modify as needed the Python code developed by Kunal for that project. These scripts and code would be made available to the community for their own use for similar surveys.

Risk: Low We can build upon the process used by 13B-370 for the Stripe-82 survey for schedule and catalog construction. Some rules for keeping track of schedules and archiving will need to be devised and followed.

5.4 Overall Observing Schedule

The VLASS as proposed will be carried out over the course of at least 6 cycles of JVLA in its A and B configurations, spanning a total of 7 years (88 months or more). We present here one possible schedule for observing the VLASS. Other scenarios are possible, should new constraints be imposed (e.g. based on design reviews).

5.4.1 Scenario: 7-Year Schedule

In the following scenario, we spread the VLASS over 7 years and parts of 6 cycles (88 months: Table 6). In this scenario, the survey starts with B configuration in 2016A in May 2016, and completes

Table 6: Possible distribution of VLASS observing hours by configuration and configuration cycle

Cycle	Config	ALL-SKY	COSMOS	ECDFS	Elais-N1	Total
1	B	906	0	0	0	906.0
1	A	0	60	325	188.5	573.5
2	B	906	0	0	0	906.0
2	A	0	60	325	188.5	573.5
3	B	906	0	0	0	906.0
3	A	0	60	325	188.5	573.5
4	B	906	0	0	0	906.0
4	A	0	60	325	188.5	573.5
5	B	906	0	0	0	906.0
5	A	0	60	325	188.5	573.5
6	B	906	0	0	0	906.0
6	A	0	0	335	188.5	523.5
Total		5436	300	1960	1131	8827.0

in August 2023. Each configuration cycle gets no more than 1480 hours of scheduled VLASS. As noted in § 3.1 earlier, however, the low declination portion of ALL-SKY could benefit from being undertaken in the BnA configuration (rather than B configuration), in which case the total number of hours given for B configuration in Table 6 would be split into 312 hours for BnA, and 594 hours for B.

5.5 Schedule Pressure by LST

Using the proposed schedule (§ 5.4.1) for 8827 hours spanning 6 configuration cycles over 7 years, we have estimated the total LST pressures in the B and A configurations for the various components. These pressures represent the average number of “passes” at a given LST needed to carry out the VLASS.

Formally, we approximate each component of the VLASS as a total time T broken into observing blocks that can be observed in a “window” of length W starting at some LST H . All times from H to $H + W$ for that component are assigned a pressure $P = T/W$. We estimate the parameters in Table 7 for the VLASS:

Table 7: LST pressure for components of the VLASS

Component	Config	T(hrs)	W(hrs)	H	P
ALL-SKY	B	5436.0	24.0	0 ^h	226.5
DEEP COSMOS	A	300.0	9.0	5.5 ^h	33.3
DEEP ECDFS	A	1960.0	5.0	1 ^h	392.0
DEEP E-N1	A	1131.0	12.0	10 ^h	94.2

Plots of the total VLASS LST pressures are presented for B-configuration (Fig 1) and A-configuration (Fig 2).

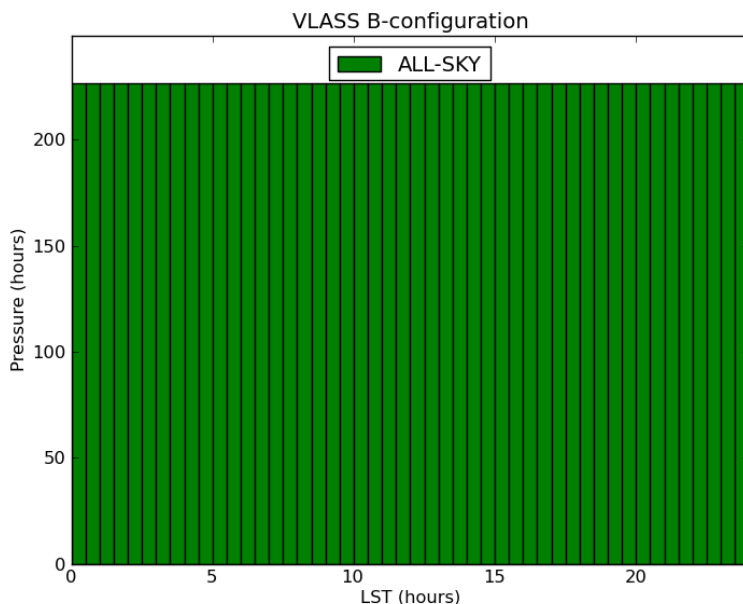


Figure 1: Total pressure by LST for the ALL-SKY component of the VLASS observed in B (and possibly BnA) configuration. Hours are calculated for the scenario with observations spread through 6 configuration cycles spanning 7 years, using the methodology described in § 5.5. Note that use of BnA in ALL-SKY (1872 hours) would be an equivalent pressure of 77.8 hours on this plot uniformly spread over all LST.

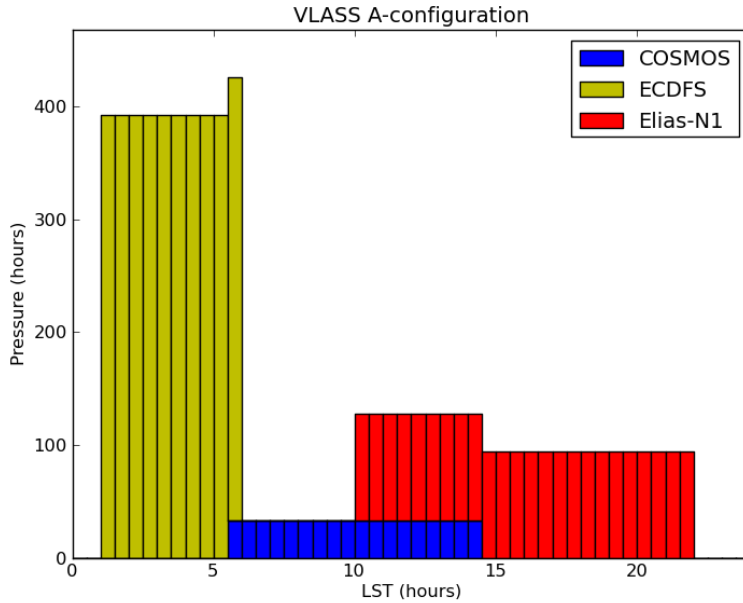


Figure 2: Total pressure by LST for the DEEP components of the VLASS observed in the A-configuration. The hours calculated are for the proposed scenario spread over 6 configuration cycles spanning 7 years. The large peak in the DEEP pressure corresponds to the ECDFS window.

5.6 Impact of VLASS on PI science

When spread over six configuration cycles the total number of observing hours that will be needed for the VLASS amounts to 18% of the total number of hours typically used for science observing, given the *current* observing efficiency of the VLA, which includes 50 hours per week for engineering, software maintenance, and other commissioning tests. While it is probably possible to reduce the time spent on software and commissioning tests while undertaking the survey by making the decision to offer only stable, fully-commissioned capabilities to the community during this period, thereby minimizing the need for commissioning test time, this is a decision that will have to be made by NRAO management rather than the proposers of the VLASS. We therefore assume that operations continue as currently in the following analysis.

Table 8 shows the current distribution of observing hours per configuration within a 16-month configuration cycle (move time between configurations is not included), along with the approximate number of hours for PI science. It is important that the 16-month cycle be maintained, so that targets at a particular LST can be observed at night time at some point during three configuration cycles, for both the VLASS and regular PI science. However, it is clear from comparing column 3 in Table 8 with Table 6 that, if the timing of the VLA configurations within a configuration cycle were to stay the same during the VLASS, then the B/BnA/A configurations will be adversely impacted in terms of the number of hours available for PI science compared with the other configurations. In addition to these considerations, we note that NRAO is currently investigating whether it is possible to eliminate the hybrid configurations from the regular configuration cycle, because the usual proposal pressure as a function of LST for the hybrids is currently very skewed and poorly matched to the available hours at certain times of year.

We therefore propose that, while undertaking the VLASS, the configuration lengths within a

configuration cycle be modified as shown in Table 9. Two options are given, one assumes that ALL-SKY uses only B-configuration, the other assumes that low declinations are observed using the BnA hybrid. With these models, the impact of the VLASS is spread over all configurations, rather than confining them to those used by the VLASS. Note that in the A-configuration the availability of time for PI science is not a uniform function of LST (see Fig. 2).

Table 8: **Current configuration lengths and hours for PI science**

Configuration	Length (days)	Hours for PI science
D	84	1411
DnC	18	302
C	80	1344
CnB	18	302
B	95	1596
BnA	18	302
A	100	1680

Table 9: **Proposed configuration lengths and hours for PI science during the VLASS**

Config.	Length (days)	PI sci. hours	Length (days)	PI sci. hours
	(ALL-SKY B-only)	(ALL-SKY B-only)	(ALL-SKY BnA+B)	(ALL-SKY BnA+B)
D	79	1327	79	1327
DnC	0	0	0	0
C	79	1327	79	1327
CnB	0	0	0	0
B	135	1362	117	1372
BnA	0	0	18	0
A	121	1459	121	1459

6 Calibration

The goal of the VLASS Calibration process is to determine, on the basis of *a priori* factors and from observations of standard calibration sources, the corrections to the raw data amplitude, phase, and visibility weights to be applied to the data. This process also determines the flags that are needed to remove bad data due to instrumental faults, RFI, and other causes of error. When applied to the VLASS data, this calibration will allow the production of images in the next processing stage. This process only includes the derivation of the complex gain and bandpass calibration factors known through previous measurements or determined by the observations of calibrators and transferred to the VLASS target observations. The self-calibration of VLASS data is included in the imaging stage of processing.

VLASS data will be processed using a modified version of the normal CASA-based VLA calibration pipeline. By the time of the VLASS observations, the VLA pipeline will have the requisite functionality to process VLASS data (full polarization, many individual target fields generated in OTF mode). The VLA pipeline is currently run on all VLA observations and is well tested. The VLASS will use a version specifically tested on VLASS pilot observations.

The VLASS Calibration Pipeline will carry out

1. Application of initial online flags (off-source, focus error, subreflector error)
2. Determination and application of derived flags (RFI, bad antennas, shadowing, other)
3. Switched power amplitude calibration and antenna gain curves
4. Flux scale calibration (using standard sources)
5. Complex Delay and Bandpass Calibration
6. Complex Gain Calibration
7. Flux density bootstrapping (from primary to secondary calibrators)
8. Polarization On-Axis Leakage Calibration
9. Linear Polarization Angle Calibration
10. Interpolation and Application of Cumulative Calibration
11. Final Flagging of Data (insufficient or failed calibration, RFI)
12. Output of Quality Assurance (QA) information, plots, images

As mentioned earlier when discussing Data Products, it is more efficient to store the calibration tables, flags, and pipeline commands and then create the calibrated dataset upon request from the archive, rather than to store both raw and calibrated datasets. Should it be deemed appropriate and possible, we might consider also archiving off-site the full calibrated dataset (e.g. through Enhanced Data Services by a community partner).

Risk: Low to Medium The VLASS is well-suited to pipeline processing and will be able to make use of the VLA Calibration Pipeline as it is currently being developed. This pipeline is now being re-implemented in a new architecture based on the ALMA pipeline development, and by the time of VLASS will be well tested. Incorporation of polarization in this pipeline will be a primary goal over the coming year, and is necessary not just for VLASS. There are minor medium risk issues in the calibration for the Quick Look images.

6.1 Observations of Calibrators

The first concern for calibration is setting up the observing such that sufficient calibration can be performed on the data. The calibration of the Tier 2 DEEP fields is straightforward, as they require only one or two good calibrators per field.

The calibration of the entirety of Tier 1 will be more difficult. We will require the availability of a “network” of suitable calibration sources spread across the entire visible sky with a spacing of one every few degrees — a density of around $0.05\text{--}0.25\text{ deg}^{-2}$ giving a list 1700–8500 in total. These will need to be compact (for the VLASS A and B configuration observations). Initially, the obvious starting place is the current VLA Calibrator Database, currently containing 1865 sources. Of particular interest is the VLBA calibrator database, with around 4700 sources in the VLASS area, as these sources are known to be compact and have astrometric precision positions. They will need to be verified as point-like to the VLA, with no source structure outside the VLBA delay beam.

Another good starting set is the CLASS catalog of compact sources measured at 8.4GHz in A-configuration Myers et al. (2003). The CLASS database contains over 13000 distinct detected sources in the northern celestial hemisphere, a good fraction of which would be suitable as VLASS

calibrators. There are extensions to the VLA-visible southern celestial hemisphere, including the compact calibrator catalog of Winn et al. (2003). These have some overlap with the VLA and VLBA calibrator databases but should yield a substantial number of excellent calibrators.

Prior to starting the VLASS Tier 1, it would be prudent to observe candidate calibrators from these lists to identify those suitable for our survey. This could be carried out for 6000 targets with around 60 hours of observations. Ideally, this would be done at S-band in B or A configuration, but it may also be practical to carry out at C- or X-band in the C-configuration prior to the start of the VLASS.

Identification of calibrators for Tier 1 in the Galactic Plane will be more difficult. Most of our known calibrators come from parent surveys that avoided the Galactic Plane and bulge. Some additional work will likely be required to find a good dense set of calibrators in this region. For example, one could contemplate a shallow fast-scanning survey covering the $\sim 3200 \text{ deg}^2$ area in around 40 hours, sacrificing imaging quality and spectral bandwidth and sensitivity for scan speed within allowed data rates. This would also be of general benefit to VLA users.

Note that once the first epoch of Tier 1 is complete, the survey itself will yield a full (unbiased) list of calibrators for use in subsequent epochs.

Identification and characterization of the VLASS calibrator list is a key item in the Test & Development Plan below (§ 10).

Risk: Low to Medium. There are viable paths to obtaining an initial calibrator database, although ideally verification of this database would require up to 100 hours of test observing prior to the VLASS.

6.2 Calibration Issues for the VLASS

Because the VLASS covers the entire sky usefully visible to the VLA, there are a number of issues related to calibration that require special attention and investigation. These are mostly related to the RFI known to exist at 2–4 GHz, but also include possible ionospheric issues.

6.2.1 RFI at S-band

As described above in § 2.1.1 substantial parts of the 2–4 GHz band are plagued by RFI. Some (though not all) of the most pernicious interference comes from geostationary satellites in the “Clarke” belt, which appears at Dec $\delta \approx -5.5^\circ$ from the VLA. This will affect VLASS observations near (within 10 degrees from our current estimation) this Declination. In addition to reducing the usable bandwidth, RFI can affect the calibration of the data. Investigation of this effect and the devising of observing and calibration strategies to mitigate RFI induced problems are part of the Test and Development Plan (§ 10.1.1).

Case Study: Stripe-82. The B-configuration S-band observations (12A-371, 13B-370) in Stripe-82 encountered calibration issues attributed to RFI. In the first set of observations in 2012 (12A-371) covering 50 square degrees, the sub-stripe was broken into two sections that were calibrated using different calibrators, one slightly north of the stripe (J2212+0152) and the other to the south (J2323–0317). During processing, it was noticed that the observed amplitudes on the southern calibrator J2323–0317 varied greatly as a function of time (azimuth), and that application of the gains derived from these scans to the target data introduced spurious variations in the flux densities and rms variations in off-source areas of the images. The observations of the northern calibrator J2212+0152 showed little or no variation. There did appear to be some variation in the Stripe-82 target observations, and thus the effect was still present on the equator. Some of the hour angles for J2323–0317 most strongly affected appeared to correspond to azimuths of known geostationary satellites, and thus compression in the receiver system due to RFI from these satellites was strongly suspected. When the observations for the full 270 square degrees of Stripe 82 for 13B-370 were carried out, schedules were constructed avoiding the hour angles most strongly impacted

in the 12A-371 data. In addition, calibrators to the south of Stripe-82 were not used. The discovery and mitigation of this effect in the 50 square degree sub-stripe are described in more detail in the paper by Mooley et al. 2015 (in preparation). More extensive tests on the full Stripe-82 from 13B-370 are still underway, and will also be part of the VLASS Test and Development Program.

Risk: Medium RFI is clearly going to be a problem near the Clarke belt, and possibly elsewhere. Observations in these affected regions will likely need to be treated specially during processing, with extra resources and personnel requirements.

6.2.2 The Ionosphere at S-band

There will probably be some times when ionospheric disturbances will impact the calibration and imaging of the data. Some of these can be avoided by use of dynamic scheduling, or re-observing of blocks damaged beyond repair. Use of pre-calibration (e.g. using the TECOR algorithm) and imaging/self-calibration iterative schemes coupled with global sky models (which will evolve as the survey progresses) will be employed as needed for the rest of affected data.

Risk: Low to Medium. The extent of this is unknown. Recent observations and reduction of data from Stripe-82 and COSMOS at S-band do not appear to be strongly impacted by the ionosphere, although occasional times requiring careful self-calibration were encountered, and were possibly due to ionospheric issues. The main risk will be in the time and effort required to identify affected data and to re-process using mitigation techniques.

6.3 Algorithm and Software Development for the Calibration Pipeline

Only a modest amount of additional software development is required for the VLA Calibration Pipeline to handle the VLASS data. In particular, accommodation of OTF scans in the pipeline (e.g. suppression of flagging at beginning and end of scans for the short OTF scans) is the primary issue. Otherwise, the VLASS data can be treated as any other VLA dataset.

The other issue for use of this standard VLA pipeline for VLASS is its overall speed and efficiency. In particular, use of the pipeline within the Quick Look image production path will require it to process data promptly upon observation. The current VLA Calibration Pipeline is too slow, although bottlenecks have been identified (e.g. the speed of `applycal`, production of QA plots). It is likely that development of a version of the pipeline that streamlines processing for VLASS will be necessary. Exploration of options for speed-up for QL and for general use are part of the Test and Development Program.

Risk: Low to Medium The modest changes in the standard pipeline here should not require significant resources. Use of the pipeline for normal VLASS use is straightforward and low risk. Use for Quick Look processing might be medium risk if our efforts to speed it up are insufficient. Fall-backs include using more cluster nodes for parallel processing of single SBs, or use of the Caltech AIPS-Lite calibration pipeline.

6.4 Calibration Processing Requirements

Estimation of the processing resource requirements required for carrying out the VLASS operations outlined above is fairly straightforward. We now have experience running the current (script-based) JVLA Calibration Pipeline for all observed projects. In addition, the Stripe-82 and COSMOS programs have used custom calibration scripts as well as the standard calibration pipeline. All of these use CASA in the currently available (version 4.2.2 and earlier) standard “serial” processing mode. This forms the baseline for our processing time estimates.

6.4.1 Case Study: Calibration of Stripe-82 datasets

For the purposes of VLASS testing, we have created a custom crafted CASA calibration script for the 3-hour blocks of the 13B-370 program. This goes end-to-end from the importing of a raw SDM file to a CASA Measurement Set, through to the creation of a calibrated MS containing only the target data. As of this time, it does not include polarization calibration (though this should only increase the processing time by a small factor, perhaps 10%). It also does not create plots equivalent to the standard pipeline “weblog”.

The test data set was one of the 3-hour blocks from the 13B-370 Stripe-82 program¹². Running the custom calibration script took 10.8 hours of elapsed time, or a factor of 3.6 times the observing block duration of 3 hours. This implies that it is feasible to run a serial calibration process in less than 4 times the observing time, and thus it is practical to keep up with observing by running 4 independent calibration processes on a single cluster node (the memory footprint for calibration appears to be sufficiently small). For reference, the standard VLA Calibration Pipeline¹³ took 66.25 hours to process this same 3 hour block (22 times the block duration), so some substantial efficiency improvements (including disabling of much of the plotting) would be required to employ this for the VLASS.

The calibration for Quick-Look imaging is the most constrained, as our specification is to fully process the data and make images available within 48 hours of the completion of an observation block, with a goal of 24 hours or less. If the serial processing of a block takes 4 times the block duration, then observing blocks would need to be 6 hours or less in duration to process within 24 hours (leaving 24 hours for imaging to meet the specification), or 3 hours if you want to process within 12 hours (half of the goal of 24 hours total). Implementation of an truly parallel calibration pipeline would allow use of longer observing blocks.

Risk: Low The modest changes in the standard pipeline here should not require significant resources. Use of the pipeline for normal VLASS use is straightforward and low risk. Use for Quick Look processing might be medium risk if our efforts to speed it up are insufficient, although custom scripts are clearly feasible. Implementation of intrinsic parallelization would be ideal and greatly beneficial, but is not required to meet our specifications.

6.5 Temporary Storage Requirements for Calibration

Carrying out the calibration and quality assurance of the VLASS will require temporary storage of the calibrated data. This is through use of our Lustre filesystem, a highly parallel high throughput RAID array system. We expect that our Lustre system will have a total capacity of 1200TB in 2017, and 2500TB in 2020. We would like the total temporary VLASS Lustre usage during the survey to not exceed 10% of the capacity, or 120TB to 250TB total. As described below in § 9.2, the raw visibility storage volume of VLASS will be 132TB per configuration cycle (with only around 111TB on target after calibration backing out the overhead). Thus, we will not be able to keep a full cycle worth of calibrated data on live disk during long-term processing (e.g. excluding the Quick Look imaging and analysis). If we stage the processing carefully, we should be able to handle keeping less than 25% live, or around 25TB for the target data.

7 Imaging

The VLASS is at its heart a wide-band continuum imaging survey. The science goals of the survey are predicated on the ability of the instrument and data processing to deliver images of sufficient

¹²SB 13B-370.sb28581653.eb28626177.56669.781848645835

¹³<https://science.nrao.edu/facilities/vla/data-processing/pipeline>

quality to be able to identify objects and measure the salient properties (e.g. flux density, position, spectral index, polarization, light curve).

Starting with calibrated visibilities output from the Calibration Pipeline, the Imaging Pipeline must:

1. Select a sub-region of the sky to image
2. Gather the visibility data that are relevant to that sub-region of sky
3. Carry out wide-band continuum imaging for that sub-region
4. Assess the quality of the imaging, determine whether further imaging, self-calibration and/or data flagging iterations are required
5. If required, perform iterative imaging steps
6. Output final images, sky models, and QA information and plots

We now describe the requirements on the VLASS Imaging Pipeline, discuss software development needed to carry this out, and estimate the processing resources that will be used to image the VLASS.

7.1 Imaging Requirements

In order to keep up with the observing, the Imaging Pipeline must be able to process the data at a rate commensurate with the observing rate. This will be effected through the parallel image processing of sub-mosaics on NRAO-based clusters or externally provided systems (e.g. through XSEDE).

There are three imaging processes that need to be handled by the pipeline:

- Quick Look (QL) imaging triggered after every scheduling block is observed (e.g. for transient identification)
- per-epoch imaging triggered after the last observation each configuration
- cumulative imaging triggered after each epoch beyond the first, incorporating all previous data

For each of these, there are three kinds of images that may be produced:

- Wide-band (2–4GHz) continuum images
- Full-resolution (2MHz channels) image cubes
- Coarse-resolution (128MHz spectral windows or similar) image cubes

Imaging is done in all Stokes parameters (IQUV) for polarimetry capability.

Continuum imaging can include higher-order Taylor terms in the spectral dimension (e.g. spectral index, spectral curvature) depending on image depth (e.g. for processing beyond the Quick Look). CASA has algorithms for this that have been used in past programs, and further development of these capabilities is underway. Full-polarimetric imaging is a key part of VLASS, and the use of accurate polarized “primary beam” maps of the VLA field-of-view during imaging and analysis are critical to the production of science ready images. Self-calibration (through the use of previous sky models as well as true self-calibration from iterative imaging) is also an integral part of the image processing.

We assume for all image size calculations that in the ideal case the images will be pixellated at a sampling level 0.4 of the (robust weighted) resolution at the highest frequency of the band (4 GHz), rounded to a convenient value. For A-configuration, this is 0.2'' (resolution 0.49'' at 4GHz), giving 324Mpix per square degree. For B-configuration, this is 0.6'' (resolution 1.58'' at 4GHz), or 36Mpix per square degree. In practice we will refine this based on imaging performance and we may be able to get by with less oversampling. In addition, individual sub-images will likely need some amount of extra padding to accommodate odd shapes. Overall we should treat these estimates as reasonably conservative, uncertain at the around the 25% level. However, this optimal level of resolution leads to large image archive sizes (see § 9) and thus a key issue for testing is the determination of the lowest acceptable resolution in the images and image cubes that will still enable the key science with the Basic Data Products. It may be possible to reduce the image data volumes by significant factors (2–8) in this manner. In addition, image compression algorithms (lossless and lossy) can also improve storage efficiency and will be investigated in the Test and Development Program.

The image cubes are needed as input to more advanced processing for Rotation Measure determination, spectral line surveys, and more detailed SED modeling of sources. Most of these would be provided as Enhanced Data Products and Services. Note that the storage and distribution of the large full-resolution cubes will be a challenge for the archive (see § 9), and options for external hosting and “on-demand” image processing should be explored (e.g. as an Enhanced Data Service). As a fallback we would carry out compression through a combination of reduced angular resolution (e.g. 0.3'', Nyquist at 3GHz) and spectral resolution (average 8 channels, 32MHz).

The imaging might be expected to be particularly challenging in the Galactic Plane, due to poorly-sampled extended emission. Here, we use the CORNISH survey as an example of what to expect Purcell et al. (2013). CORNISH was a 110 deg² continuum survey of the Galactic Plane from 10° < *l* < 65° carried out in the B and BnA configurations at 5 GHz in 2006–2008, before the WIDAR correlator was available. Purcell et al. reported that areas exhibiting imaging problems due to extended emission amounted to less than 2% of the surveyed area, and that these problems could be dealt with through special weighting schemes and manual inspection of the uv data. The CORNISH survey did not include the complex Galactic Center region, but the service monitoring of this region for the Sgr A*/G2 encounter in 2014¹⁴ showed that snapshot imaging in the B-configuration is relatively straightforward, helped by the increased fractional bandwidth of the JVLA and MFS imaging algorithms. We can therefore expect that some of the Galactic Plane imaging in Tier 1 will require manual intervention, but that the fraction of the data requiring such attention will be small. In these regions imaging may be dynamic range limited, rather than sensitivity limited.

Risk: Medium to High Previous Jansky VLA programs (e.g. Stripe-82 and COSMOS observations) have been imaged using scripts and software tailored to the individual projects. As of this time, there is no standard CASA-based JVLA imaging pipeline available. The development of this capability is a high priority for NRAO in the next two years, and thus the VLASS can build upon and help pioneer this effort. We expect this to take substantial resources (design, development, implementation, execution, and QA), some of which will need to take place also as part of the normal JVLA operational plan. Due to their similarity to Stripe-82 processing, we expect the pipelining of Tier 1 extragalactic imaging to be only medium risk, with the possibility of falling back to the processed used as part of 13B-370. We assess the risk for Tier 1 imaging in the Galactic Plane and Tier 2 Deep Fields to be higher, due to the considerably more stringent dynamic range constraints, the complexity of the emission in the Galactic fields, and the relative lack of previous experience in automated image processing for large deeper mosaics. These risks will be mitigated through the VLASS Test and Development Program (§ 10) using existing data taken in similar modes of observing.

¹⁴<https://science.nrao.edu/science/service-observing>

7.2 Algorithm and Software Development for Imaging

The development and implementation of improved software in CASA for imaging is critical for the success of the VLASS. For example, the Quick-Look imaging is the most time-critical, while the DEEP and Galactic Plane imaging will be the most difficult from dynamic-range and image complexity perspectives.

The issues to be addressed include:

- the robust parallelization of imaging with minimal memory footprint (so multiple processes can be run on a single node) and optimal I/O (to keep disk access traffic to only what is required);
- the efficient joint mosaicking over a modest area (44 primary beams) with multiple Taylor terms, including W -projection (and A -projection), multi-scale, for full polarimetry;
- handling of primary beam effects during imaging (e.g. A -projection) and for correction of the final images (“wfpbcor”) at the appropriate level and accuracy for the VLASS components;
- the use of pre-boxing (based upon other catalogs or evolving VLASS global sky models) and auto-boxing;
- ability to carry this out in a semi-automated pipeline that includes fault detection and handling, robust and frequent check-pointing of the run to allow restarting of the imaging upon failures or crashes.

Note that these are also general development targets needed to fully realize the potential of the JVLA and ALMA for high-quality imaging. The needs of the VLASS are not special (the DEEP component is not as deep as existing and planned observations) but the large volume of data and the requirement to produce science capable images on our schedule do place constraints on the timescales for this work, and thus require substantial resources for the software development and testing. Naturally, aspects of this work are part of the Test & Development Plan (§ 10.2 and 10.3).

7.3 Processing Requirements for Imaging

The imaging of the VLASS data will clearly be the computational bottleneck for processing. However, it is also the most straightforward part of the processing to run parallel serial scripts on, as individual fields or sub-mosaics can be processed independently and combined together in a final linear mosaic (or left as “postage stamp” sub-images). Careful benchmarking using representative test data for the various VLASS Tiers and fields will be required to obtain accurate estimates of required resources. This will be a critical task under the Test & Development Program, and will require significant effort to carry out.

As an illustrative example, we consider a simple test case based upon existing data.

7.3.1 Test Case: Imaging of Stripe-82

A custom simple imaging script for the same observing block used for the calibration benchmarking (§ 6.4.1) was run alone on a single cluster node. The observations contained in this 3 hour observing block covered approximately 22.5 square degrees of the sky reaching a depth of $\sim 90 \mu\text{Jy}/\text{beam}$ rms for the full band. On average, it took 145 seconds per field, or 81.56 hours for all 2025 OTF phase center fields. Thus, the basic imaging for our 3 hour observing block took approximately 27 times the block duration. Thus, imaging could keep up with observing by running the imaging of individual fields in parallel using 7 nodes with 4 processes per node.

Note that the script used for this test only did very basic imaging of each OTF phase center as if they were single fields. No joint mosaicking was employed, nor was any self-calibration carried out. This did a single MFS Taylor term and single scale deconvolution, did not include any auto-boxing, but did use W-projection to deal with the wide-field aberrations. The cell size employed was 1'' matched to the band center, rather than the more optimal 0.6'' proposed for VLASS (§ 9.4), so the VLASS images would have 2.8 times the number of pixels compared to images generated in this test. Our test thus provides a lower limit on the processing required for the Tier 1 ALL-SKY imaging, and an approximation of the requirements for Quick-Look imaging in both tiers. Given the above considerations, one might conservatively extrapolate that VLASS imaging will require around 4 times the resources indicated by this test, requiring around 30 nodes (with 4 processes per node) to image in the same time as observing (e.g. 6 hours to image a 6 hour block).

These preliminary estimates carry significant uncertainty, and more careful studies are required to provide accurate projections of resources needed for VLASS imaging. In particular, detailed tests using the deeper COSMOS data and test observations in the Galactic plane, are needed to give reliable estimates for the imaging of the galactic plane in ALL-SKY and the field of the DEEP component of the survey.

7.4 Temporary Storage Requirements for Imaging and Analysis

In § 6.5 we estimated that we could handle keeping around 25TB of the calibrated target data live on Lustre storage at any given time for further processing. We also argued that we wanted VLASS to use no more than around 120–250TB (from 2017–2020 onward) at any given time. The image data products will likewise take up temporary storage, in particular for further analysis and production of catalogs, as well as for quality assurance before ingestion into the NGAS archive.

The following numbers are derived from the archive space requirements given in § 9 below:

Quick-Look Images: The Quick-Look images will take on 35TB total, or only 6TB per cycle, and thus are easily accommodated. We plan to keep all 6TB for a given cycle live.

Single-Epoch Images: There are 207TB of single epoch images to be archived in the VLASS (§ 9.5). Note that the DEEP Tier 2 uses a total of only 6TB and can be kept on disk. The ALL-SKY Tier 1 uses the remaining 201TB or 67TB per epoch (3 epochs), produced at a rate of 33TB per configuration cycle. If we assume that two cycles (previous and current) are on disk at a given time for processing and comparison, that requires 67TB live. We expect to need a total of 73TB for single-epoch images.

Cumulative Images: There are 253TB of cumulative images to be archived (§ 9.7). As with single epoch images, the DEEP fields take only a modest amount of storage, in this case 59TB (almost all for the spectral cubes). A total of 184TB is required for ALL-SKY cumulative image storage, with 43TB for continuum images and 141TB for spectral cubes. We assume all of DEEP and ALL-SKY continuum, and 25% of ALL-SKY need be on disk at a given time (35TB), and thus require a total of 137TB. Note that the cumulative imaging does not kick in until after the second cycle for DEEP (2018) and the second epoch for ALL-SKY (2019).

The total temporary storage for imaging and analysis estimated above is 216TB.

General Scratch Space: In processing all of the above, some amount of scratch space will be required to handle trials of processing options and comparisons, as well as higher resolution spectral cubes before compression. Note that processing for the B-configuration ALL-SKY and A-configuration DEEP described above will be naturally staggered in time (though with some overlap), and thus we will not need the full 216TB live all the time. Thus the 216TB can accommodate a 67% of that total steady state plus a 50% extra for scratch. Thus we do not add extra for scratch.

The 216TB for imaging is to be added to the 25TB estimated in § 6.5 for calibration. Thus, we estimate that **a total of 241TB of temporary Lustre storage** will be sufficient to carry out the

VCLASS. This fits within the long-term envelope of 10% of predicted Lustre storage reached in 2020 (250TB).

8 Image Analysis and Sky Catalogs

The main image analysis task for the VCLASS is the production of the basic object catalogs for the Quick-Look and standard images.

A good study of the performance of radio continuum image source finders is Hancock et al. (2012), which considers the available options in the context of ASKAP. The upshot is that there are options available that should have acceptable performance for the basic catalogs for VCLASS. Note that inclusion of the spectral index images and polarimetric images from VCLASS will likely require some extensions to these source finders, which in turn will require some developer or astronomer time.

Also available as a proof of concept is the source finding carried out for the JVLA Stripe-82 surveys by Mooley et al. (in preparation). There is also a comprehensive discussion in Mooley et al. (2013) in the analysis of archival VLA ECDFS multi-epoch data. It is our current assessment that one or more of these methods will be suitable for the basic catalogs from the VCLASS.

More advanced catalogs and source finding algorithms could be developed and produced as an Enhanced Data Product.

Risk: Low to Medium There are a number of options available which will do some of what we want, if not all. Dealing with the wide-band spectral index and polarimetric data products of the JVLA will likely require some amount of extra work.

9 Archiving and Data Distribution

The primary interface that the user community will have to the VCLASS is through the archive and data distribution system. Raw data will be served via the normal JVLA archive, available with no proprietary period as soon as it has been ingested into the archive system.

The archive, or at least some archive, will have to also serve the VCLASS data products as described above. It is the responsibility of NRAO and the VCLASS to make the Basic Data Products available through this archive mechanism. Enhanced Data Products may be made available through the NRAO-hosted VCLASS archive, this will need to be negotiated and is largely dependent upon resources required. The VCLASS products, either in basic form or further processes, may also be made available via alternative Enhanced Data Services, as described above.

9.1 Estimated Archive Space

The estimated ideal data volumes required for storage of the VCLASS data and data products are:

1. raw visibility data — 800TB
2. calibration data — not significant
3. quick-look continuum images — 35TB
4. single-epoch continuum images — 101TB
5. single-epoch image cubes — 106TB
6. single-epoch basic object catalogs — not significant

7. cumulative “static sky” continuum images — 43TB
8. cumulative “static sky” image cubes — 200TB
9. cumulative “static sky” basic object catalogs — not significant
10. **Total: 800TB of data, 485TB of images** (continuum images — 179TB; image cubes — 306TB)

These volumes have been calculated assuming the angular and spectral resolutions defined in § 7 with compression in some cases as noted below, for the data products listed in § 4.1. We consider for the above calculation that the storage needed for the calibration data and catalogs to not be significant compared to the imaging and visibility data, and should be less than 10TB total.

The 800TB of data storage is slightly above what we expect for normal science data use (it uses the maximum allowed data rate of 25MB/s for the ALL-SKY observations), and thus will be accommodated within the normal JVLA archive data storage plans. The required storage volume of nearly 500TB for the images, however, is significant, but can be accommodated in storage expansion plans. Based on information provided by James Robnett, the JVLA NGAS archive is planned to have the capacity for around 3300TB in 2017 (1 year after start of VLASS) and 5700TB in 2020. Thus, our image data will use less than 10% of the total NGAS storage in 2020. The NRAO support for storage beyond this level would require extremely strong recommendation by the VLASS science review based on strong science drivers. Beyond this, partnerships with the community and other agencies will be necessary (see § 4.2).

Risk: Medium to High Details of the overall NRAO archiving plan as they extend into the future are yet to be developed fully and resourced. The VLASS archive should fit in with this overall plan, and perhaps be used as a testbed to explore options. The risk is mainly due to this uncertainty in overall NRAO data management strategy at this time. As noted above, the clear high (resource) risk issue is the large volume taken by the image cubes, and the current plan relies upon significant (spatial and frequency) compression. We believe the proposed compression schemes can be carried out without unduly impacting the science goals (spectral intensity, polarization and rotation measure studies).

We now provide details on our estimates of the data volumes required for storage of each of the individual VLASS data products.

9.2 Raw Visibility Data

The VLASS will take a total of 8827 hours of observing time. If we assume that the entire survey uses the maximum allowed data rate of 25MB/s, then **the total visibility data volume for the VLASS is approximately 800 TB**. We observe around 1480 hours per configuration cycle, and thus a full A and B cycle of data takes 132TB of storage.

In practice, different parts of the survey may end up using different correlator configurations. For example, the ALL-SKY passes at shallow depth (120 μ Jy per pass) use shorter integration times and the maximum data rate of 25MB/s. However, the DEEP observations if not using OTFM could use longer integration times and thus lower data rates. Thus our 800 TB is an upper limit to the VLASS data volume.

Note that if it is necessary for us to reduce data rates and volumes further, we could employ frequency averaging in the correlator backend or archive. This has potential undesirable effects due to delay losses and loss of instantaneous field of view which could complicate processing, as well as limiting the usefulness of the data products for Rotation Measure analyses, and we do not recommend taking this approach unless truly necessary.

9.3 Calibration Data

The total primary and secondary calibration data for VLASS must come from the assumed overhead time not on target, and in practice will be only around 10% of the total data volume. Furthermore, the calibration and flagging tables derived from this data will be significantly smaller than this (being per antenna rather than per correlation), comprising much less than 1% of the total data volume. Thus we conclude that calibration data is an insignificant contributor to archive volume.

As noted earlier, the archive will need the facility to store, retrieve, and apply this calibration data to raw visibilities when requested by users and further processing pipelines.

9.4 Quick-Look Continuum Images

These images are produced for every scheduling block for purposes of QA and transient object detection. We assume for all image size calculations that the images will be pixellated at a sampling level 0.4 of the (robust weighted) resolution at the highest frequency of the band (4 GHz), rounded to a convenient value. For A-configuration, this is $0.2''$ (resolution $0.49''$), giving 324Mpix per square degree. For B-configuration, this is $0.6''$ (resolution $1.58''$), or 36Mpix per square degree. In practice we will refine this based on imaging performance and we may be able to get by with less oversampling. In addition, individual sub-images will likely need some amount of extra padding to accommodate odd shapes. Overall we should treat these estimates as reasonably conservative, uncertain at the around the 25% level.

For QL production images, we assume 2 BDP images for Stokes I plus an uncertainty map. See § refsec:qlimages for a description of these products.

The total storage required for the QL images is 35TB.

The rates and volumes of data production depend upon the specific Tier:

9.4.1 Tier 1 ALL-SKY

The 33885 square degrees will be mapped in B-configuration in three epochs. An all-sky image comprises 1.2Tpix, with size 4.8TB assuming single precision storage. For a single epoch of the ALL-SKY, the 2 QL images will take 9.6TB per epoch, with 28.8TB required for all three epochs.

9.4.2 Tier 2 DEEP — COSMOS

The 2 square degrees in A-configuration has 648Mpix (2.6GB) per image. The 2 images in a given QL pass will take 5.2GB for each block. There are a total of 40 observing blocks, so the total storage required for COSMOS is a modest 208GB.

9.4.3 Tier 2 DEEP — ECDFS

The 4.5 square degrees in A-configuration takes 1.46Gpix (5.83GB). The 2 QL images per pass take 11.7GB. There are 392 passes total requiring 4.6TB.

9.4.4 Tier 2 DEEP — Elais-N1

The 3.5 square degrees in A-configuration takes 1.13Gpix (4.54GB) per image. The 2 QL images per pass take 9.1GB. There are 156 passes total requiring 1.4TB.

9.5 Single-epoch images and cubes

After each epoch, we can potentially produce refined images and cubes for the following:

Continuum Intensity Images: For these we assume there will be 2 Taylor terms ($\tau\tau_0$ and $\tau\tau_1$) and uncertainty maps for each, giving a total of 4 wide-band continuum images (I, I α , plus uncertainty maps).

Continuum Polarization Images: We assume that only 3 images are required, Q, U, and an uncertainty (rms) map (valid for both Q and U). We do not plan to create V circular polarization images for VLASS as a BDP, due to the required care of careful gain calibration and added expense of beam squint corrections during imaging.

There are thus a total of 7 continuum image products.

Coarse Spectral Cubes: As a compromise, we can store coarse cubes (in the 14 viable spectral windows, each 128MHz wide). These would be for the 5 standard images (I, Q, U, plus uncertainty maps for I and polarization) resulting in a total of 70 image cube planes. This takes 1/64 the storage of the full cubes. Note that for the ALL-SKY Tier 1, we cannot as BDP store the “full” coarse cubes, and offer only a compressed (10:1) version. This is most straightforwardly done as “postage stamp” cutouts around identified sources.

Full Spectral Cubes: As described in § 4.1.4, the image cubes can potentially have 1024 channels at full resolution (2MHz per channel), of which expect to only have unflagged data in a maximum of 896 channels (2 spectral windows totally lost and dropped). We would expect to create 5 cubes (I, Q, U and uncertainty maps in I and polarization) for the 896 channels, producing 4480 total image cube planes. Storage of cubes this size will be prohibitively large for ALL-SKY, and still too large for DEEP (336TB) to store and serve as a BDP.

Total storage required for single-epoch images and coarse spectral cubes is 207TB. The continuum images require 101TB and (coarse and compressed) cubes use 106TB.

9.5.1 ALL-SKY Tier 1

An all-sky image in B-configuration comprises 1.2Tpix, with size 4.8TB assuming single precision storage. This requires a total of 14.4TB for the 3 epochs.

Continuum Images: The 7 continuum images (4 intensity, 3 polarization) will take 100.8TB total storage.

Full Spectral Cubes: For the 4480 total cube planes (5 cubes, 896 channels) this would require 64.5PB total (10.75TB per cycle). *We do not plan to store or serve these ALL-SKY full spectral cubes as part of the BDP.*

Coarse Spectral Cubes: The coarse spectral image cubes with 70 planes total (5 cubes, 14 windows) will take a total of 1008TB for all 3 epochs (336TB per epoch, or 168TB per cycle). *We do not plan to store or serve these ALL-SKY coarse spectral cubes as part of the BDP.* Because polarization science requires some form of spectral cubes, we will provide an even more compressed (10:1) coarse cube, most likely by providing “postage stamps” around identified sources. This would require 100.8TB.

The total storage for all three of the single epochs of ALL-SKY under BDP is 201.6TB, or 67.2TB per epoch.

9.5.2 DEEP Tier 2

There is a total of 3.24Gpix in the 10 square degrees of DEEP, or a total of 18.8Gpix epochs (COSMOS gets 5 epochs, the others 6) requiring a total over the survey of 75.2GB for each image or cube plane.

Continuum Images: The 7 continuum images require a total of 526.4GB.

Coarse Spectral Cubes: The 70 total planes (5 cubes, 14 windows) require a total of 5.3TB.

Full Spectral Cubes: The 4480 total planes (5 cubes, 896 channels) require a total of 336.9TB (56.15TB per cycle). *We do not plan to store these DEEP full spectral cubes as part of the BDP.*
Total storage required for all single epochs of DEEP is 5.8TB.

9.6 Single-epoch basic object catalogs

The data volume for catalogs will not be significant. Flat catalog files in simple formats can easily be served from the archive. If served by a Relational Database (RDB), the number of entries may be large and may need consideration in implementation.

9.7 Final cumulative “static sky” images and image cubes

After each epoch after the first, we will produce refined images and cubes summed over the preceding decades, eventually arriving at the “final static-sky image”. For these we assume there will be potentially:

Continuum Images: For these refined deeper images we can make 3 Taylor-Term images, for 9 total (I, Ialpha, Icurv, Q, U and uncertainty for I, Ialpha, Icurv, and polarization).

Coarse Spectral Cubes: For these we make 5 cubes (I, Q, U, plus uncertainty maps for I and polarization) in each of 14 spectral windows, or a total of 70 image planes. For ALL-SKY, as with the single epoch images, we plan to provide compressed (10:1) coarse cubes, likely as cutouts in regions containing sufficiently bright emission. The delivery of the uncompressed coarse resolution cubes for ALL-SKY is a possible target for an Enhanced Data Product or Service (§ 4.2).

Full Spectral Cubes: For these we make 5 cubes (I, Q, U, plus uncertainty maps for I and polarization) in each plane. There are potentially 896 channels in a 14 spectral window dataset, giving a total of 4480 image planes. As noted below this requires more storage than we expect to be available, so some channel compression will be necessary for the BDP cubes. The delivery of full spectral resolution cubes (even in “postage stamp” cutouts around bright sources) is a target for an Enhanced Data Product or Service (§ 4.2).

Highly Compressed Spectral Cubes: Given that full resolution spectral cubes are too large to store with projected resources, we entertain the possibility of providing compressed cubes at sufficiently high spectral resolution to carry out the primary spectral science goal of rotation measure studies. The most likely approach is a combination of spatial compression (cutouts around bright emission regions) and frequency compression. The frequency channel compression will need to be done carefully, likely with the higher frequency channels compressed more. There may also be lossless compression schemes that can reduce the lossy compression required. We assume that a total compression factor of 200:1 is possible. The possible compression factors and optimal scheme is a Test & Development Plan item.

We need to archive these only for the latest cumulative release.

The cumulative image storage total for VLASS is 243TB. The continuum images take up around 43TB, and the cubes will require 200TB.

9.7.1 ALL-SKY Tier 1

There is a total of 1.2Tpix in ALL-SKY or 4.8TB per plane.

Continuum Images: A total of 43.2TB required for the 9 images.

Coarse Spectral Cubes: A total of 336TB would be required for 70 planes. Thus, we propose to provide compressed (10:1) coarse cubes as in the single-epoch image case. This requires 33.6TB.

Full Spectral Cubes: For full resolution cubes with 4480 total planes (5 images, 896 channels) a total of 21.5PB would be required. This is well above our projected storage available (§ 9.1). *Full spectral resolution cubes for ALL-SKY will not be provided as a BDP.* We propose to provide highly

compressed (200:1) high frequency resolution spectra around objects as described above. This will require 107.5TB of storage.

A total of 184.3TB is required for ALL-SKY cumulative image storage, with 43.2TB for continuum images and 141.1TB for spectral cubes.

9.7.2 DEEP Tier 2

The 10 square degrees of DEEP requires 3.24Gpix, or 13.0GB per plane.

Continuum Images: A total of 116GB required for the 9 images.

Coarse Spectral Cubes: A total of 910GB required.

Full Spectral Cubes: A total of 58.2TB required.

A total of 59.2TB is required for DEEP cumulative image storage, with 0.1TB for continuum images and 59.1TB for spectral cubes.

9.8 Cumulative “static sky” basic object catalogs

The data volume for catalogs will not be significant. Flat catalog files in simple formats can easily be served from the archive. If served by a RDB, the number of entries may be large and may need consideration in implementation.

10 Test and Development Plan

There are a number of issues related to the VLASS that must be addressed before the survey can be observed on the telescope. We feel that none of these are “show-stoppers” that are likely to prevent the survey from being carried out at all, and most have obvious work-arounds. Fundamentally, these are schedule and resource risks rather than functionality risks, in that it will take longer and will be more costly in computing and human resources to process the survey, impacting the data product delivery schedule. However, they do need to be addressed, and we propose a VLASS Test and Development Program leading up to and through the survey start. For example, short test observations or larger pilot observations are indicated in some areas, while analysis of archival data from previously observed projects such as Stripe-82 13B-370 will serve in others. These will require significant astronomer resources to carry out, and thus we are unable to fully execute this program before submission of the VLASS proposal — approval for the observation of VLASS would be necessary before allocating the resources to carry out this test program. This is particularly true for the issues in Tier 2 imaging. There will be a final critical design review before survey observations commence, and by that point we will have dealt with the high and medium risk issues sufficiently to proceed.

The areas we have initially identified for testing, exploration, and development include: general flagging, calibration, and imaging issues; OTFM tests; deep mosaic imaging; and pipeline development. Many of these are currently being explored by VLA staff and resident observers as part of normal development and science support, and by the user community in their research activities. However, new resources (additional staff, post-doc, or student time, computing, and telescope test time) will be required to fully implement this plan before the start of the VLASS. We do also note that some of the tests and developments do not have to be completed before the start of survey. We now describe more details pertinent to selected key areas.

10.1 General Flagging, Calibration, and Imaging Tests

The main tests to be performed are related to RFI, calibration, and wide-field polarimetric imaging performance of the JVLA. These are issues that are important for VLA observers in general, not

just the VLASS. The VLASS provides us an opportunity for directed testing of these issues, with the results to be made available to the user community (e.g. through VLA memos and improved software tools).

Also to be studied in this area are the trade-offs that can be made in the image and image cube sizes (angular and spectral resolution) in order to reduce archive and distribution costs but still maximize the science capabilities of the survey. Also under investigation are the utility and limitations of more general lossless and lossy compression algorithms on the images.

We highlight a few of the key challenges below.

10.1.1 RFI, Bandwidth, and Sensitivity

As described in § 2.1, the main contribution to the sensitivity of the JVLA for continuum observations is the effective useable bandwidth that one must enter into the radiometer equation or the Exposure Calculator Tool. Past experience has shown that good (automated) imaging can be carried out with fairly brutal flagging of affected sub-bands, leaving around 1350 MHz of effective bandwidth. On the other hand, past careful reductions have shown that 1500 MHz should be achievable, and this is what we recommend to users. As this difference is equivalent to 11% in observing time or 5.4% in sensitivity, this issue constitutes a medium risk for the VLASS and will require significant testing as well as the development of automatic RFI heuristics (and possibly new algorithms).

An example RFI sweep in S-band carried out as part of VLASS testing is shown in Figure 3. This was carried out at Declination $\delta = -5^\circ$ which is the location of the XMM-LSS deep field (an early candidate for a DEEP field), and is also located in the heart of the “Clarke” geostationary satellite belt as viewed from the VLA. The XMM-LSS deep field was a high priority candidate for the VLASS DEEP component. It was rejected based on the RFI at that Declination as seen in that sweep at S-band (L and C band have similar concerns). It might be prudent to perform a test observation of this field to verify that the impact of RFI precludes it from being part of DEEP; if it turns out that this field is viable (e.g. by adjusting the baseband edges to exclude the worst RFI and recording only the good spectral windows) then it might be beneficial to observe this instead of the more Southern ECDFS field, which could save time to reach the desired sensitivity. More such scans are needed covering a range of Declinations.

Risk: Medium. Data can easily be obtained, but requires careful testing to establish the performance at all Declinations. Mitigation could involve loss of data and sensitivity, or development of more complex heuristics or algorithms.

10.1.2 Wide-band Wide-field Continuum Imaging

The substantial increase in the continuum survey speed for the JVLA comes from the greatly increased instantaneous bandwidth provided by the upgrade. However, in order to produce a continuum image from this bandwidth, the system response and the source spectrum must be folded in to the calculation. This includes effects such as:

- the system response over the full bandwidth (e.g. the complex bandpass and system noise spectrum);
- the antenna primary beam responses (in the R and L polarizations) over the field-of-view (initially in the main beam only, but ultimately out into the near sidelobes);
- the array synthesized beam (PSF) as a function of frequency, for extended sources;
- the intrinsic source spectrum over this band (intensity, spectral index, spectral curvature, etc.);

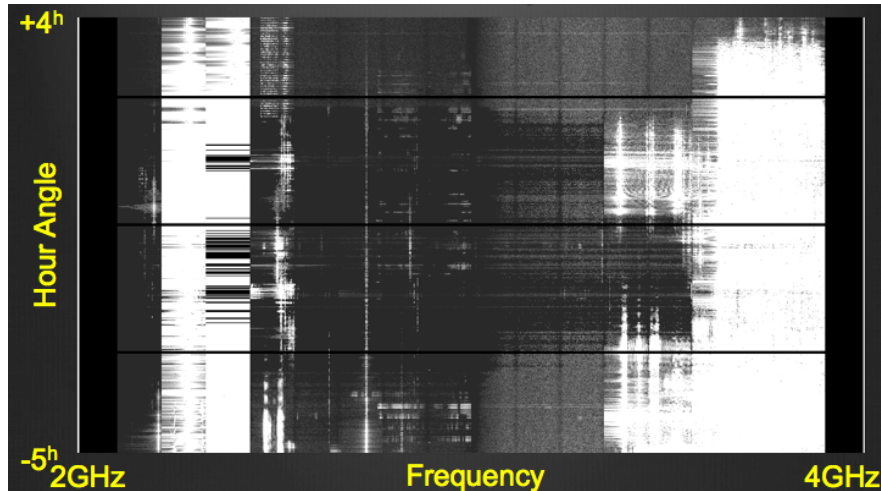


Figure 3: A sweep using fast OTF scanning across the sky at Declination $\delta = -5^\circ$. This is the location of the XMM-LSS deep field which was a candidate for a DEEP target. The horizontal axis is frequency channel from 2–4 GHz, and the vertical axis is time, covering hour angles from -5 hours to $+4$ hours. At the lower and upper ends of the spectrum we clearly see strong satellite interference at all azimuths, and over wider areas of the spectrum at certain hour angles corresponding to specific satellites. On the basis of this test, the XMM-LSS field was disfavored. Sweeps like this will need to be done at various declinations of interest.

- the amount of flagged channels (e.g. for RFI) as a function of frequency over the band;
- the distribution of visibilities (and weights) in the uv-plane after gridding.

These effects will all contribute to the final continuum image rms, resolution (“synthesized beam width”, sidelobe level and structure, response to extended emission, and reconstruction of the source spectrum (through the multi-Taylor-Term imaging).

The default way of dealing with the source spectrum over the bandwidth is to break the band into smaller coarse channels, such as imaging each 128 MHz spectral window separately as a sub-band. In this approach, the spectral index can be fitted across these sub-bands, and incorporated into the model. This can be done iteratively during imaging, until no more emission is detectable above the noise in each sub-band. The drawback is that deconvolution is limited to what is detectable in a sub-band rather than the combined full-band image.

In principle, the technique of Multi-Frequency Synthesis (MFS) can be used to image the Taylor expansion terms (TT) with frequency over the full band. This is implemented in CASA as part of the `clean` task, with the performance as described in Rau et al. (2014). However, very recent tests using simulated point sources and the L-band CHILES data (Gim et al. private communication) may indicate at low signal-to-noise ratios (SNR) there may be a bias in the recovered spectral indices at $\text{SNR} < 20$. In this event, we would have to fall back to using the sub-band imaging approach described above.

Another aspect of wideband continuum imaging performance to be evaluated is the mosaic sensitivity to sources of different sizes and spectral indices. In a mosaic, the effective survey speed as a function of frequency goes as the area of the primary beam, which varies by a factor of 4 over 2–4 GHz. Thus, the lower end the band gets more integration time when summed up over the mosaic than the upper end. When combined with weighting of the spectrum taking

this factor into account, this moves the effective mean frequency of the wideband continuum flux density downward from the nominal 3 GHz band center, giving more sensitivity to steep spectrum sources than one would otherwise expect. However, the exact weighting function and thus the magnitude of this effect also depends upon the relative instrumental noise (SEFD) and amount of channel flagging done over the band. In addition to this purely spectral effect is the fact that the synthesized beam also varies as a function of frequency in the same way, and thus more extended sources also get more sensitivity at the lower end of the band. Thus, careful test measurements will be needed before we can establish the true sensitivity to sources as a function of size and spectral index.

Risk: Medium The tests should be straightforward, and if software or algorithm flaws are identified CASA will need to apply bug fixes or re-engineering in any event. The fall-back to sub-band imaging may incur penalties in the achievable imaging rms.

10.1.3 Polarimetry

A key aspect of the VLASS as compared to previous surveys is the capability for wide-band polarimetry at sub-mJy flux density levels. This science case was highlighted in the VLASS White Paper “A Wideband Polarization Survey of the Sky at 2–4 GHz” Mao et al. (2014).

The use of the Jansky VLA for deep wide-band polarimetric is being pioneered through several RSRO programs. For example, graduate student Preshanth Jaganathan is carrying out C-band observations in Elais-N1 that will inform our development of VLASS processing. The CHILES-ConPol L-band processing is also ongoing (C. Hales, NRAO). Of more direct interest are the COSMOS S-band observations (PI: Smolcic). These programs have identified issues and paths forward, in particular the need for high-quality full-Stokes primary beam measurements and models. Observations obtaining these and the processing are underway.

Special care will be taken in investigating strategies for dealing with variable polarized sources when combining passes and epochs in ALL-SKY and DEEP imaging. It is likely that this may push us towards observing fields at the same LST and thus with the same uv coverage on each occasion.

Risk: Medium to High. Given the unique science that polarimetry brings to the VLASS makes this a high science priority issue. That this JVLA capability is only recently being fully explored, makes this medium to high risk at the current time. Mitigation of this risk over the next year is a top priority test and development target.

10.1.4 Transients and Variability

Exploring the time domain is a key science frontier for the VLASS. The survey has been designed to survey the sky in several epochs with sufficient cadence to identify candidate varying sources. The requirements on, and pitfalls of, a transient survey are described by Frail et al. (2012). Previous and current JVLA programs (12A-371, 12B-158, 13B-370) are employing this capability and provide us a set of on-sky test cases. We will work with these groups to carry out the required tests for the VLASS.

As part of the Test & Development Program, the choice of cadences of the ALL-SKY and DEEP observations for classes of transients and variables will be investigated. For example, if ALL-SKY is broken into two regions of the sky (say N and S) then the three epochs spread among 6 cycles would uniformly sequence as NSNSNS giving 32-month cadence between epochs of a given region. It may be advantageous to sequence this to NNSNSS so there are 16-month and 32-month separations sampled between epochs.

Risk: Low Existing pilot programs are underway. No obvious show-stoppers. The main need is modest personnel time to investigate and document.

10.2 ALL-SKY Tier 1 OTFM Testing

The VLASS science case in the ALL-SKY tier is built around the efficiency of OTF mapping. The overhead time during OTF mapping is driven by time to slew to calibrators, which scales with time on sky. Overhead for pointed mapping scales with the number of points visited on the sky (beams times visits), so shallow and/or repeated visits to the same part of the sky would be very costly. However, OTF imaging has until recently been considered an experimental observing mode, so more work is needed to confirm in detail that the VLASS science goals are achievable in this mode. Here we describe the open questions to achieving the full VLASS science case (particularly transients) via OTF mapping and a plan to resolve these questions.

As detailed in the Guide to Observing with the VLA, survey speed is equal to the beam area divided by the integration time needed to reach a given sensitivity:

$$SS(\text{deg}^2\text{hr}^{-1}) = 0.5665\theta_{pb,ref}^2/t_{int} \quad (58)$$

where $\theta_{pb,ref} = 15'$ is the primary beam FWHM at the band center reference of 3 GHz, and t_{int} is the integration time reported by the VLA Exposure Calculator Tool for the required image rms. In OTF mode, the survey area is covered by a series of stripes (“rows”) spaced by θ_{row} . For nearly uniform sensitivity coverage over the band, we adopt a value

$$\theta_{row} = \theta_{pb,min}/\sqrt{2} \quad (59)$$

where for an upper frequency (RFI free) of 3.6 GHz we find $\theta_{row} = 8.84'$ for $\theta_{pb,min} = 12.5'$. To cover a given area in a given time using OTFM, the survey speed is given by

$$SS = \theta_{row} \dot{\theta} \quad (60)$$

where $\dot{\theta}$ is the scan rate of the telescopes across the sky (in arcmin/sec).

In Equation 1 we derived a nominal survey speed of $SS = 16.55 \text{ arcmin}^2/\text{sec}$ (equivalently deg^2/hr) for image rms of $\sigma_{I0} = 100\mu\text{Jy}/\text{beam}$. Using our chosen row separation and Equation 60, we find for this fiducial depth

$$\dot{\theta}_0 = 1.87 \text{ arcmin s}^{-1} \quad (61)$$

For a given depth (rms image noise level) of σ_I

$$\dot{\theta} = SS/\theta_{row} \quad (62)$$

$$= \dot{\theta}_0 (\sigma_I/\sigma_{I0})^2 \quad (63)$$

$$= 1.87 \text{ arcmin s}^{-1} (\sigma_I/\sigma_{I0})^2. \quad (64)$$

$$(65)$$

A critical limitation to OTFM is that the correlator dump time must be fast enough to temporally resolve the amplitude envelopes of sources as they move through the primary beam pattern. We have empirically found that the dump time must sample sources roughly faster than 10 times its motion through the primary beam:

$$t_{\text{dump}} * \dot{\theta} < 0.1 \theta_{pb,min}. \quad (66)$$

The minimum dump time (and thus number of possible epochs) is limited by the output data rate, which should not exceed 25 MB/s. See § 2.4 for a discussion of the limitations on scanning speed and thus survey depth per epoch imposed by this 25 MB/s data rate. There, we derive the limit of $120\mu\text{Jy}/\text{beam}$ in a single pass. As argued above, this is the shallowest (and fastest) scanning that can be achieved without dropping spectral windows. This will require some testing to make sure

that the correlator back-end can deliver data for long periods at the short dump times and that the lost data does not reduce the image sensitivity below the desired values.

With this sketch of how a multi-epoch OTF survey could work, we outline the tests needed to address remaining questions:

1. Quantify the imaging errors introduced by fast OTF slewing
 - We have test data taken at a range of slew rates that will define the magnitude and scaling of imaging artifacts for OTF mode.
 - This test will give us a maximum safe slew rate for a given required image sensitivity.
2. Evaluate whether the sensitivity of three co-added $120 \mu\text{Jy}$ epochs scales down to $69 \mu\text{Jy}$
 - Science (13B-370 PI: G. Hallinan) and test OTF data in S-band are in hand to study whether shallow co-added epochs can reach the theoretical sensitivity at a level of $69 \mu\text{Jy}$ or better.
 - Included in this test will be an investigation of whether observations at similar or different hour angles as a function of epoch provide the best results
 - This test will benefit from ongoing development of joint deconvolution image cleaning in CASA.
3. Confirm whether the correlator can observe with 16 spectral windows and a dump time of 0.45s
 - A correlator test observation would quickly answer whether this is possible.
4. Investigate whether certain spectral windows are consistently hit by RFI for ALL-SKY
 - Search of archival data for S-band RFI.
 - If a spectral window is unusable for a large fraction of the sky, we will exclude it and recalculate the minimal single-epoch sensitivity.
5. Verify increase in integration times needed to reach uniform depth for low-declination fields
 - These would be carried out through test observations in the regions under study to see the loss of sensitivity in imaging compared to equivalent observations at higher elevation. In addition to these direct imaging rms tests, we will use the “Switched Power” data in the VLA dataset to track system temperature variations.
 - For OTF mosaics, longer integration times translate to a slower scanning speed compared with the rest of ALL-SKY, which will need to be quantified and tested
6. Quantify polarization bias/uncertainty is introduced by OTF mapping, focusing on answering the following questions:
 - Are observations of polarization calibrators in OTF and pointing mosaic mapping modes required?
 - What is the polarization bias as a function of frequency and pointing location? Polarimetric mapping in both OTF and pointed mosaic modes may require further algorithm development to accommodate large fractional bandwidths.
 - Are more observations of the in-beam polarization required to support polarimetric calibration in general?

7. Investigate strategies for handling time-variable sources for intensity and polarization imaging when combining the 3 epochs of ALL-SKY
 - The (13B-370 PI: G. Hallinan) data can be used for this; it is likely that the best strategy will involve observing the ALL-SKY area at fixed LST scheduling, so that each epoch has exactly the same uv-coverage and thus will average during gridding.
8. Evaluate more accurately the calibration and slewing overheads for ALL-SKY, in order to optimize the scheduling
 - Construction, observation, processing and analysis of realistic VLASS test blocks is needed.
9. Finalize estimate of computing resources are needed for ALL-SKY image processing
 - The (13B-370 PI: G. Hallinan) data and CASA scripts used to estimate the processing time given in § 7.3.1 above can be used for this.
 - This test will guide, test, and track speed and performance (e.g. memory footprint) improvements in CASA imaging code as they are implemented.
10. Quantify sensitivity degradation of imaging complex structures in the Galactic Plane and Bulge region
 - Archival data (e.g., the service monitoring of the Galactic Center) can be used for this test, but test observations in OTF will also be required

The timescale on which the answers to some of these questions is needed is quite extended, as the second epoch of ALL-SKY does not start until the third configuration cycle of the VLASS.

10.3 Tier 2 Deep Imaging Tests

The Jansky VLA has successfully been used to take data, which has been processed using AIPS and CASA (sometimes in combination) to make deep ($\sim 1 \mu\text{Jy}/\text{beam rms}$) images from wideband continuum data at S-band. However, deep imaging at this level is not routine and is heavily manual in nature. Most of the past experience has been in imaging single pointings (Owen et al. A2256 and GOODS-N, Van Gorkom et al. CHILES COSMOS pointing at L-band). The most relevant current mosaic is the 12B-158 Smolcic et al. COSMOS mosaic that is the basis for the COSMOS DEEP field. This is currently being processed by the Smolcic group, and lessons learned in this processing will be critical to developing the plan and pipeline for Tier 2 processing.

The VLASS Deep Fields will present a challenge to our ability to reliably and efficiently pipeline image deep wideband polarimetric mosaics. The key high-level issues that we foresee include:

1. What is the true system performance (RFI-free bandwidth, system temperature) in the low-declination ECDFS field? What is the actual added integration time over the nominal needed to reach the $1.5 \mu\text{Jy}$ uniform depth (1.3 was assumed)?
2. Can the full sensitivity required for DEEP be realized in CASA for both Stokes I and polarization products?
3. Is the VLA instrument (e.g. primary beam) sufficiently characterized to allow imaging and calibration of the wide-band wide-field polarimetric mosaics at the depth and dynamic range required for Tier 2?

4. What is the best strategy to deal with time-variable sources in the DEEP fields, to both characterize them for variability studies and to mitigate the effects on the combined deep images?
5. What computing resources are necessary to image Tier 2? New methods of averaging the uv-data are being implemented to support the imaging of large datasets, and will need to be tested.
6. What is the optimal scheduling strategy to support deep imaging (e.g., maximize uv-coverage, vs. minimizing the effects of variable sources on the imaging)? Archival observations of the COSMOS field will be used to support this investigation.

10.4 Other Logistical Tests

In addition to the tests already described, the following will need to be investigated:

1. Calibrators for VLASS
 - In order to optimize the observing schedule nearby calibrators are needed; the VLA calibrator catalog has several areas of the sky where the density of calibrators could be significantly improved, and some effort will be needed to compile a set of calibrators for the VLASS from other catalogs
2. Data storage and image cube compression
 - Determination of optimize vs. practical options for image cube pixel and spectral resolution compression (needed for archival storage planning)
3. Source finding algorithms
 - Assessment of source finding algorithm performance for both Tier 1 and Tier 2 images is needed, and any needed improvements identified

10.5 Pipeline development

The following modifications are expected to be needed for the VLA calibration pipeline to support the VLASS:

1. Polarization calibration
 - The addition of polarization calibration, based on user-defined scan intents, will be needed for both Tier 1 and Tier 2. The calibration heuristics for Tier 2 are expected to be the same as is needed for standard VLA observations. The calibration heuristics for Tier 1 may need modification pending the outcome of the polarization tests described in § 10.2.
2. Imaging
 - A modified version of the VLA calibration pipeline is already in use for 13B-370 that includes Stokes I imaging, so automated production of polarization products will need to be developed and tested
3. Quick-look calibration and imaging pipeline

- A stripped-down version of the full CASA-based calibration and imaging pipeline will be developed and tested as the Quick-Look pipeline. If this proves to be too slow for transient detection we may have to use the AIPS-Lite pipeline already developed for 13B-370. This item will primarily comprise timing tests.

4. Source catalog production

- The output of source finding algorithms will need to be cataloged and integrated into the NRAO archive system

11 Project Schedule and Resource Estimates

11.1 Schedule

The key high-level **milestones** and notional dates, as well as *opportunities for testing*, following the March 2015 VLASS Review and a subsequent decision by the NRAO Director to carry out the VLASS, and leading up to the start of VLASS observing, are:

- **2015 May** Pending outcome of Community Review, VLASS Project starts. Establishment of Project Office and initial staffing.
- *2015 Feb 6 – 2015 May 11* 2015A B-configuration. Opportunity for on-sky tests of VLASS B-configuration observing, as well as calibrator searches.
- *2015 May 15 – 2015 Jun 1* 2015A BnA-configuration. Opportunity for on-sky tests of VLASS BnA hybrid configuration observing.
- *2015 Jun 12 – 2015 Sep 21* 2015A A-configuration. Opportunity for on-sky tests of VLASS A-configuration observing, as well as calibrator searches.
- **2015 August 15** VLASS Preliminary Design Review. Main goal is to identify key software developments (e.g. in CASA) needed over the next 6 month development cycle in order to be ready by the start of VLASS observing.
- *2016 Feb 5 – 2016 Apr 25* 2016A C-configuration. Opportunity for final on-sky tests and preparatory calibrator searches.
- **2016 April 15** Final VLASS Critical Design Review. A positive outcome of the CDR will be approval for commencing observing, possibly with modifications of plans and delay of start.
- *2016 May 27* Start of 2016A observing in B configuration, and earliest possible start for VLASS observing.

11.2 Resource Requirements

The telescope time, computing resources required for data processing, and archive storage, have been described in detail above. Here we address the effort required to execute the Test and Development Plan, and to deliver the VLASS and its basic data products. These are FTE estimates based on local experience and expertise.

Test and Development Plan, as described in § 10:

- General flagging, calibration, and imaging issues

- 0.7 FTE-yrs (risk: low)
- OTFM tests and development
 - 0.8 FTE-yrs (risk: medium)
- Deep Mosaic Imaging tests
 - 0.5 FTE-yrs (risk: medium)
- Other Logistical tests
 - 0.3 FTE-yrs (risk: low)
- Calibration and imaging pipeline development and testing
 - 0.8 FTE-yrs (risk: low)

The total effort required for the Test and Development phase of the VLASS is 3.1 FTE-yr, but as noted elsewhere, some of the test items extend beyond the start of VLASS observing.

VLASS operations and production of BDPs (May 2016 – Aug 2023):

- Management of project office
 - 0.3 FTE/yr (risk: low)
- Scientific staff to oversee observing execution and pipeline products (“Astronomer on Duty”)
 - 0.5 FTE/yr (risk: low)
- Data analysts to oversee pipeline product quality assurance
 - 3 FTE/yr (risk: low)
- Scientific staff to develop and oversee catalog production
 - 0.5 FTE/yr (risk: high)

The total effort required during VLASS operations is 4.3 FTEs/year.

Development of “processing on demand” capabilities:

- Requirements definition (scientific staff)
 - 0.3 FTE-yr (risk: medium)
- Software development and implementation
 - 1.5 FTE-yr (risk: medium)
- Testing (scientific staff)
 - 0.3 FTE-yr (risk: medium)

The total effort needed in order to implement POD is estimated to be 2.1 FTE-yr.

Beginning with the first cycle of observations, in 2016, effort will be need to support Communication, Education, and Outreach activities:

- VLASS “picture of the week,” interface to NRAO EPO office

– 0.5 FTE/yr (risk: low)

These numbers reflect the NRAO staff time required to carry out or supervise these activities. For NRAO scientific staff, the FTEs listed are for VLASS-related functional work only, and do not include the additional time required for self-directed research.

There are a number of areas that can be addressed and augmented through the participation of members of the astronomical community (SSG members, post-docs, graduate students, NRAO REU students in a few cases). Note that use of community shared effort in these technical areas is unlikely to significantly reduce the required NRAO staff levels given above, as careful coordination and support will be required.

12 Acknowledgements

The authors of this document would like to acknowledge the hard and effective work done by the VLASS Survey Science Group (SSG) and the numerous White Paper authors and working group participants who contributed to the proposal for which this is but a supporting document. We also acknowledge the input from SSG member Jim Condon, who pointed out the issues of setting survey limits and sensitivity to extended sources and sources with non-flat spectra, and provided many of the calculations presented in § 2.6. Finally, we would also like to acknowledge the constructive input given by the NRAO reviewers of this plan and the VLASS proposal.

References

- Frail, D. A., S. R. Kulkarni, E. O. Ofek, G. C. Bower, and E. Nakar, 2012: A Revised View of the Transient Radio Sky. *Astrophysical Journal*, **747**, 70.
- Hancock, P. J., T. Murphy, B. M. Gaensler, A. Hopkins, and J. R. Curran, 2012: Compact continuum source finding for next generation radio surveys. *Monthly Notices of the Royal Astronomical Society*, **422**, 1812–1824.
- Mao, S. A., J. Banfield, B. Gaensler, L. Rudnick, J. Stil, C. Purcell, R. Beck, J. Farnes, S. O’Sullivan, D. Schnitzeler, T. Willis, X. Sun, E. Carretti, K. Dolag, D. Sokoloff, R. Kothes, M. Wolleben, G. Heald, J. Geisbuesch, T. Robishaw, J. Afonso, A. M. Magalhães, B. Lundgren, M. Haverkorn, N. Oppermann, and R. Taylor, 2014: A Wideband Polarization Survey of the Extragalactic Sky at 2-4 GHz: A Science White Paper for the VLA Sky Survey. *ArXiv e-prints*.
- Mooley, K. P., D. A. Frail, E. O. Ofek, N. A. Miller, S. R. Kulkarni, and A. Horesh, 2013: Sensitive Search for Radio Variables and Transients in the Extended Chandra Deep Field South. *Astrophysical Journal*, **768**, 165.
- Myers, S. T., N. J. Jackson, I. W. A. Browne, A. G. de Bruyn, T. J. Pearson, A. C. S. Readhead, P. N. Wilkinson, A. D. Biggs, R. D. Blandford, C. D. Fassnacht, L. V. E. Koopmans, D. R. Marlow, J. P. McKean, M. A. Norbury, P. M. Phillips, D. Rusin, M. C. Shepherd, and C. M. Sykes, 2003: The Cosmic Lens All-Sky Survey - I. Source selection and observations. *Monthly Notices of the Royal Astronomical Society*, **341**, 1–12.
- Purcell, C. R., M. G. Hoare, W. D. Cotton, S. L. Lumsden, J. S. Urquhart, C. Chandler, E. B. Churchwell, P. Diamond, S. M. Dougherty, R. P. Fender, G. Fuller, S. T. Garrington, T. M. Gledhill, P. F. Goldsmith, L. Hindson, J. M. Jackson, S. E. Kurtz, J. Martí, T. J. T. Moore, L. G. Mundy, T. W. B. Muxlow, R. D. Oudmaijer, J. D. Pandian, J. M. Paredes, D. S. Shepherd, S. Smethurst, R. E. Spencer, M. A. Thompson, G. Umana, and A. A. Zijlstra, 2013: The Coordinated Radio and

Infrared Survey for High-mass Star Formation. II. Source Catalog. *Astrophysical Journal Supplement*, **205**, 1.

Rau, U., S. Bhatnagar, and F. N. Owen, 2014: Wideband Mosaic Imaging with the VLA - quantifying faint source imaging accuracy. *ArXiv e-prints*.

Winn, J. N., A. R. Patnaik, and J. M. Wrobel, 2003: Compact radio sources at Dec -30 to 0 (Winn+, 2003). *VizieR Online Data Catalog*, **214**, 50083.