

# ARDG Test Report: Wideband Mosaic Imaging for VLASS

U. Rau & S. Bhatnagar

June – August, 2018

(Modified: Sept. 26, 2018)

## Executive Summary

This report describes the commissioning of wideband mosaic algorithms for VLASS Epoch I imaging. Two wide-field wide-band imaging options were evaluated against simulated as well as real data to demonstrate numerical correctness and practical usability. For successful demonstrations on real data, a series of issues with the VLASS data themselves had to also be understood and worked around. In summary, we have a working algorithm for VLA wideband mosaics, but development is still required at multiple stages (data acquisition, pipeline, ARDG and CASA) in order to handle previously overlooked consequences of the specific OTF observing setup chosen by VLASS.

**Algorithms and Software :** When the input mosaic data are self-consistent and properly calibrated and weighted, CASA/ARDG has a wideband mosaic imaging option that can operate on a VLASS-style mosaic to produce correct source positions, intensities and spectral indices at an adequate dynamic range and with expected noise levels in the image. W-term corrections are required for imaging at full angular resolution and image sizes larger than those employed for previous VLASS imaging tests are required to control PB aliasing. Compute costs of each stage of imaging are also consistent with expectations and documented. Recommendations for setting up imaging runs for VLASS are supplied, but accuracy and performance will have to be re-evaluated if VLASS chooses to scale down imaging parameters to reduce compute cost. A test build of the CASA package from the ARDG code branch has been made available for internal testing (both to the VLASS testers as well as to a few general testers).

**VLASS Data :** VLASS data has two types of inconsistencies between pointing direction and correlator phase center. One is per subscan (an erroneous offset for antennas with old ACUs due to a bug in the VLA real-time software) and one is per integration (a feature of the OTF mosaic mode, chosen on purpose to accommodate correlator rate limits, for all antennas). The resulting systematic pointing errors of about 10% of the primary beam (PB) width produce significant errors in the imaging results, especially in spectral index. A potential solution exists, but it will require development at multiple levels and may increase compute cost by a factor of up to 2. Other sources of artifacts in the mosaic images are traceable to data weighting schemes, insufficient phase calibration prior to imaging and under-sampling of the PSF.

**Size-of-computing Estimate:** Imaging with the full range of w-values will require about 128 w-projection bins. The extra computing load relative to 1 w-plane (which VLASS has used till date) is estimated to be up to 50 times higher. Additional computing to correct for the effects of antenna-dependent pointing offsets is estimated to be of the order of gridding load for the data from the affected baselines, adding up to another factor of 2. W-term corrections also add to the overall memory footprint. This needs to be determined empirically once the value for the w-projection planes parameter for VLASS imaging is settled on, but from the test VLASS data set used in this study (a typical VLASS data set), we estimate an additional 10 – 15 Gigabytes of extra memory requirement per process for 128 w-projection planes, compared to a few Megabytes for w=1.

**Deliverables and Recommendations :**

1. **Algorithms required** : A set of algorithm options to properly handle wideband PBs, w-term effects and error control due to aliasing have been documented for VLASS. For imaging of Stokes I<sup>1</sup> and spectral index at full angular resolution, the need for W-term corrections has been demonstrated and the 'awproject' gridded demonstrated to provide them along with accurate wideband PB correction. The associated compute costs of 'awproject' have been provided along with a list of remaining development items within CASA for both 'awproject' and 'mosaic' gridders in order to prepare both for VLASS imaging. The most significant new feature that both gridders lack is the recently discovered requirement for antenna-dependent pointing offset corrections. We show that the spectral index recovery is especially sensitive to the systematic pointing offsets designed into the VLASS OTF mosaic observing scheme, and that this feature cannot be ignored during imaging.
2. **Imaging Parameters** : Our tests were carried out assuming full angular resolution imaging. Based on recent discussions, if VLASS chooses to scale down some parameters in order to improve compute load, we recommend that the numerical consequences be carefully re-evaluated. In particular, decreasing the imaging angular resolution by tapering and de-selecting baselines with the largest w values is one option being considered that can potentially decrease w-term compute load while also reducing image sizes. Tests to evaluate source position offsets relative to the cell size must then be repeated. Once these parameters are finalized, compute costs for full-scale VLASS imaging must also be re-evaluated. Early versions are already in progress (by J. Robnett). In case the compute costs are still prohibitive, or the angular resolution loss due to the above types of approximations are unacceptable for the target science, we suggest that other alternate methods be explored (perhaps borrowing from some older NRAO low frequency surveys, or via alternate approaches to handle w-term effects).
3. **Data Pre-processing** : The data weighting and phase calibration schemes within the VLASS calibration pipeline must be modified in order to better prepare the data for imaging at the required dynamic range. In particular, the artificial downweighting of OTF mosaic data around all bright sources should be avoided by a more strategic use of 'statwt' for reweighting, and options to identify and correct phase errors possibly due to 'delay clunking' must be explored in order to achieve noise-limited imaging for VLASS.
4. **OTF Mosaic Pointing offsets** : For current and future VLASS data, the offsets between pointing and phase centers in VLASS data must be accounted for at all levels of the system. This will require development efforts to correct erroneous offsets within the real-time data acquisition system for antennas with old ACUs, to create accurate pointing tables for use with existing data, and to correctly use these pointing tables during imaging. The gridding compute load will also increase as the pointing offsets will have to be additionally computed per baseline and per integration.

To support the generation of pointing tables for data already taken, a self-calibration scheme has been demonstrated to work on subsets of VLASS data near known calibrators, for the purpose of determining which antennas show the above pointing offsets and which do not. This would also pick out antennas with significant enough pointing offsets (due to other reasons).

Given the nature of the pointing and phase center offsets (and their relation to delay errors and w-term effects for single bright sources), some tests to determine the final imaging parameters for VLASS will have to be repeated once a solution is developed. Supporting data and test scripts can be supplied to aid future testing.

---

<sup>1</sup>Note : This study did not investigate Full-Mueller imaging for VLASS.

# Contents

|          |  |           |
|----------|--|-----------|
| <b>1</b> | <b>Wideband Mosaic Algorithm Validation</b>                    | <b>4</b>  |
| 1.1      | Data Sets Used   | 5         |
| 1.1.1    | Simulation of VLA L-Band D-config Pointed Mosaic               | 5         |
| 1.1.2    | Simulation of VLASS-style Pointed Mosaic                       | 5         |
| 1.1.3    | Real VLASS OTF Mosaic Data                                     | 5         |
| 1.2      | Gridder Choice   | 5         |
| 1.2.1    | Current state of the 'awproject' gridder                       | 5         |
| 1.2.2    | Current state of the 'mosaic' gridder                          | 6         |
| 1.3      | Accuracy of Wideband PB Models                                 | 7         |
| 1.4      | Expected Noise Levels With conjbeams=True                      | 8         |
| 1.5      | Image/Cell Size And The Effect Of Aliasing                     | 8         |
| 1.6      | W-Term Effects   | 9         |
| 1.6.1    | How many w-planes?   | 10        |
| <b>2</b> | <b>Understanding VLASS Data</b>                                | <b>11</b> |
| 2.1      | Data Weights   | 12        |
| 2.2      | Self-Calibration   | 12        |
| 2.2.1    | Dynamic Range Improvements                                     | 12        |
| 2.2.2    | Interpreting the self-cal solutions                            | 12        |
| 2.3      | VLASS Pointing Offsets   | 15        |
| 2.3.1    | Type 1 : Offsets per subscan with old vs new ACUs :            | 15        |
| 2.3.2    | Type 2 : Offsets per integration within each subscan           | 15        |
| 2.3.3    | OTF scans and PB Smearing                                      | 17        |
| 2.3.4    | Effects on Wideband Imaging                                    | 17        |
| <b>3</b> | <b>Required Development Efforts</b>                            | <b>22</b> |
| 3.1      | CASA, Gridders, Imaging Infrastructure                         | 22        |
| 3.2      | Addressing phase/pointing offsets at all levels                | 23        |
| 3.3      | VLASS Pipeline   | 23        |
| <b>4</b> | <b>Computing Estimates</b>                                     | <b>24</b> |
| 4.1      | Gridding Load With AWProjection                                | 24        |
| 4.2      | Correction Of Antenna-dependent Pointing Errors                | 24        |
| 4.3      | Measurements Of Run-time Performance (Added on Sept. 26, 2018) | 26        |

## Chapter 1

# Wideband Mosaic Algorithm Validation

**Background** : Wideband mosaic imaging in CASA is done via the Multi-term Multi-Frequency-Synthesis (MTMFS) reconstruction algorithm (`deconvolver='mtmfs'`) along with either one of two wide-field gridders for mosaic imaging (`gridder = 'awproject' or 'mosaic'`). Over the past two years, the `'mosaic'` gridder has been used and improved upon for VLASS imaging (i.e. for commissioning and Quick-look imaging). At this time, it does not support w-term corrections and until now, VLASS imaging has been proceeding with the assumption that w-term corrections would not be needed. However, VLASS tests for Epoch-1 imaging produced inconclusive results on the accuracy of source position, intensity and spectral index ( $\alpha$ ). The question of which gridder to pick as well as whether w-term corrections were needed or not remained open and unresolved. Observed increases in noise levels with the `conjbeams=True` setting of the Wideband A-Projection algorithm to perform accurate wideband primary beam (PB) correction were also noticed and remained unresolved.

**Tests** : This study began by re-evaluating the correctness/readiness of both algorithms and their implementations, the need for w-term corrections specifically for VLASS imaging and the consequences of recently understood pointing offsets due to the chosen observing setup for VLASS. The metrics that were evaluated were:

1. The accuracy of source intensity and spectral index in single pointing imaging as well as mosaic imaging
2. The effect of the w-term and its correction on source position accuracy
3. The effect of `conjbeams=True` on noise levels across the image (for a single pointing and a mosaic)
4. General usability along with parallelization.

**Conclusions** : The conclusion of this evaluation is that the `'awproject'` gridder is currently closer to being a viable solution for VLASS imaging. It is currently able to perform accurate wideband PB correction for continuum mosaic imaging while also performing accurate W-term corrections. It has performance control features that are particularly suited to survey imaging. Its wideband PB causes the reconstructed spectral index to be steeper by 0.02 (compared to polynomial-fit PB models for S-band that are assumed to be more accurate). Minor usability improvements and peripheral bug fixes remain.

To be usable for VLASS imaging, the `'mosaic'` gridder will need W-term support (either by duplicating what `'awproject'` does to include w-projection or via a faceting approach) as well as an accurate `conjbeams=True` implementation to get correct spectral indices for all sources.

In addition, in light of recently understood algorithm requirements (to account for pointing offsets introduced during OTF mosaic data acquisition) both gridders will need development and testing efforts to ensure accurate use of the pointing table to apply corrections for antenna-dependent pointing offsets.

The following sections describe details of the datasets used, the current state of both wide-field gridders, and the process of determining optimal imaging parameters for VLASS that balance numerical accuracy with practical compute costs.

## 1.1 Data Sets Used

Two sets of data were simulated and imaged using the MT-MFS deconvolver along with either one of the two gridded options ('mosaic' and 'awproject'). A row of pointings of real VLASS data were also used.

### 1.1.1 Simulation of VLA L-Band D-config Pointed Mosaic

VLA L-Band D-config 2-pointing pointed-mosaic observations were simulated for one bright source with an amplitude of 1 Jy and spectral index of -0.5 using both the 'awproject' and 'mosaic' gridders model prediction routines. Predicted model visibility values were examined to evaluate correctness as well as differences due to the use of different PB models. Imaging was performed using both gridders on both datasets in order to prevent erroneous conclusions arising from using the same code for simulation as well as imaging and to evaluate the effect of differences in PB models between the two. W-term effects were ignored in these initial tests.

### 1.1.2 Simulation of VLASS-style Pointed Mosaic

A second set of simulations used a Measurement Set from a VLASS BnA configuration OTF observation containing 1 row of pointings covering a known calibrator source as the observing setup. Visibilities corresponding to the flux and spectrum of the known calibrator (Stokes-I flux=1.38 Jy,  $\alpha=-0.46$ ) were predicted for each pointing using the 'awproject' prediction to include appropriate PB values for each spectral window and pointing in the simulation along with w-term effects. This was a pointed mosaic simulation using the meta-data from a real observation, and this was sufficient to get realistic point spread functions and W-term effects. In a separate simulation, noise was added to test the effect of conjbeams=True setting on the continuum noise level (see sec. 1.4).

### 1.1.3 Real VLASS OTF Mosaic Data

A BnA-config row of pointings around the J1448-1620 calibrator was used for real-world tests in order to evaluate the required degree of W-term quantization (via the 'wprojplanes' parameter in CASA) and to assess practical noise levels and the effect of the algorithm on the total image RMS noise. These data were also used to evaluate the consequences of two types of recently discovered (and understood) pointing offsets, as well as to experiment with potential solutions.

## 1.2 Gridder Choice

Below is a summary of the current state of both the wide-field gridders when used along with the multi-term wideband deconvolution algorithm. Some infrastructure-code modifications were made as part of this evaluation. In scaled-down tests designed to evaluate only numerical accuracy, the runtime per major cycle for both gridders is comparable, but full-scale VLASS compute load tests (possibly via J.Robnett) will provide more realistic results.

### 1.2.1 Current state of the 'awproject' gridded

1. Works correctly for continuum imaging with conjbeams=True for internal wideband PB correction.
2. A meta-data bug fix was required to enable accurate cube model prediction (for accurate simulation) and some programmatic issues remain for cube imaging.
3. W-term support exists. On the VLASS simulation (and also on real VLASS data) it was able to shift the positions of the known calibrator source per OTF pointing to its known correct location.
4. Memory use is as expected for different stages of the imaging run.
5. The use of a pointing table has a bug that causes a misalignment of pointing information at subscan edges. For our tests, all pointing table rows were deleted (this is what the VLASS pipeline currently does too) to avoid this problem.

6. The ability to use a disk-cache for convolution functions (CF) was crucial to performing repeated imaging tests and a few usability improvements have been made to enable its efficient setup and re-use, especially with parallelization. This on-disk CF cache can be very large when W-Projection is combined with WB A-Projection, but it needs to be computed only once and can be re-used for all of VLASS imaging. Only a fraction ( $1/N_{spw}$ ) of it is loaded in the memory at any time during processing. For our tests with 10kx10k images and 0.3 arcsec cell sizes, the on-disk CF Cache sizes were 55MB for  $w=1$ , 46GB for  $w=64$ , 91GB for  $w=128$ , ( 200GB for  $w=256$ ) and 359GB for  $w=512$ .
7. The PB model used is a ray-traced beam evaluated from VLA dish parameters at the central frequencies of each spectral window. This model supports parallactic angle rotation and beam squint. The frequency scaling of the beam width is linear and evaluated at spectral window (SPW) resolution. The model includes azimuthally asymmetric sidelobes that appear at  $\sim 5\%$  gain level.
8. Currently, this gridded can operate only with square images.
9. This gridded does not have special support for VLASS-specific features that *may* have recently been added to the 'mosaic' gridded (see last point below).
10. Further development is required to implement the application of antenna-based pointing offsets based on information in the pointing table.

### 1.2.2 Current state of the 'mosaic' gridded

1. `conjbeams=True` for internal wideband PB correction is currently incorrectly implemented for mosaics. A few programmatic issues were noticed and modified in the code but the results are not yet accurate. CAS-11902 contains a description of this error as well as test programs for use in verifying a fix. Note that although `gridded='mosaic'` starts out with uncorrected spectral indices, it converges to the correct value after a couple of major cycles. This means that sources that are deconvolved will eventually converge (given enough major cycles), but undeconvolved sources will have incorrect spectral indices.
2. No support for w-term yet. This gridded is therefore usable for VLASS only if it can be demonstrated that w-term position shifts are smaller than a pixel.
3. The pointing table is used correctly (with these simulated data) when pointing offsets are identical for all antennas. However, all subsequent tests were done after deleting pointing table rows (as is currently done in VLASS imaging pipeline).
4. The PB model is an azimuthally symmetric model evaluated from a polynomial fit to EVLA holography data. The model does not include sidelobes, although there are plans to re-fit polynomial models from the original holography data and use them instead. The PBs are scaled per channel, and the frequency dependence across the band includes measured deviations from a linear scaling that are noticeable at the higher end of S-Band (i.e. closer to reality than a linear scaling of beam width).
5. The computed PBs show large scale Gibbs oscillation (a.k.a "ringing"), which also leads to strong aliasing along the edges of the field-of-view. This can be minimized with the use of anti-aliasing operators, which is not yet included in this gridded. CAS-10701 and CAS-11183 discuss fixes.
6. This gridded can operate with rectangular as well as square images.
7. Further development is required to implement the application of antenna-based pointing offsets based on information in the pointing table. Currently, the pointing offset for the first antenna in each antenna-pair is applied to the entire baseline.
8. There are a few other VLASS-specific features that have recently been added and enabled only for the mosaic gridded. They include support for declination-specific PB projection/evaluation, choosing position-dependent PSFs for deconvolution and field-dependent data weighting.




Figure 1.1: Comparison of the PB spectra between 2.4 and 3.9 GHz, at the half-power point of the PB at 3 GHz. Blue curve: Ray-traced model with linear scaling in frequency. Green curve: Polynomial-fit azimuthally symmetric model that shows non-linear structure in the frequency dependence of the PB. Overall, the PB spectral index of the polynomial-fit model is about 0.02 steeper than the ray-traced model.

### 1.3 Accuracy of Wideband PB Models

There are two Primary Beam models in the CASA code. The 'mosaic' gridded uses model derived from a polynomial fit to the holography data. The 'awproject' gridded uses a ray-traced model the parameters of which are derived from fitting the ray-traced beams to the holography data. The frequency scaling in the ray-traced model is linear.

The accuracy of the PB models affect the derived flux and spectral indices. While both these models are within a percent of each other (and hence also to the holography data), they differ in the outer parts of the PB. Specifically, only the ray-traced model includes the PB sidelobes which are azimuthally asymmetric. The accuracy of the PB models used by both gridders was evaluated in two ways.

#### 1. Differences between the ray-traced and polynomial-fit PB models:

The PB models were computed using the two gridders and saved to the disk and compared. To check the difference in frequency scaling, the PB spectrum near the half-power beam width, where the frequency dependence is strongest, was compared. Plot in Fig. 1.1 shows the two spectra. The blue curve is the ray-traced model and the green is the polynomial-fit model. The overall difference in slope corresponds to a spectral index error of 0.02. The imaging tests above also show this difference. Visibilities simulated with the polynomial-fit model and imaged using the ray-traced model produced spectral indices that were steeper by 0.02 compared to a reconstruction using the polynomial-fit model.

#### 2. Once issues with the real VLASS data were understood and temporary work-around found via shifting the phase center to match the true pointing center for most antennas (see Sec. 2.3), one single integration from pointing number 18958 was imaged separately. With the source at the 50% gain level of the PB, the PB-corrected output intensity and spectral index of the calibrator came to 1.31 and -0.45 which was the closest to the true values obtained so far on the real data. This is tentative confirmation that spectral structure in the PB models are good enough for the level of accuracy offered by the current VLASS data.

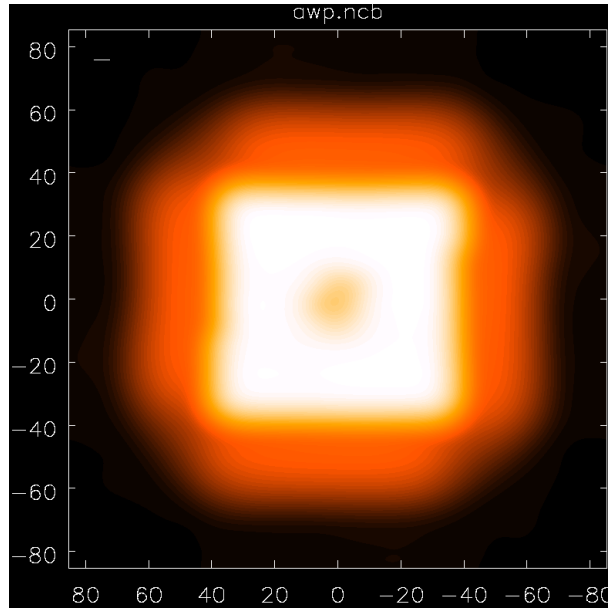


Figure 1.2: Figure shows an example of wide-band sensitivity pattern for a typical VLASS mosaic image. The labels along the two axis are in arcmin. The band around the main sensitivity pattern is the cumulative sensitivity due to the wide-band sidelobes of the PBs from various pointings in the mosaic.

## 1.4 Expected Noise Levels With `conjbeams=True`

In a VLASS simulation, noise consistent with VLASS data was added to the ideal visibilities in all frequency channels. For a single pointing, the use of `conjbeams=True` resulted in an increase in noise level of 10% on the edge of the beam compared to the center of the pointing.

This is consistent with theoretical expectations of thermal noise in a wide-band continuum image. The use of `conjbeams=True` is equivalent to performing a PB correction (moving from flat-noise to flat-sky normalization) along the spectral axis. At any location away from the center of the PB, there will always be a lower gain at higher frequencies, and therefore higher noise in the PB-corrected images. This effect will be present all across a mosaic image because at any given location there is always more overlap between PBs at the lower frequencies and less total gain at the higher end of the band.

Calculations based on a gaussian PB model and a linear frequency scaling suggest a noise increase of about 10% at the 10% PB point at the reference frequency and this is consistent with the results from our simulations.

VLASS data contain higher noise at the upper end of S-band compared to the lower end. In our tests, data weights were set to 1.0. Two pointings (field IDs 18957 and 18958) were imaged with `conjbeams=True` (Fig. 2.2). The noise level obtained at the mosaic edges were about 50% higher than at the center, a value consistent with what was observed by the VLASS team as part of their initial tests.

Given the interpretation of `conjbeams=True` as a PB correction along the spectral axis, and our independent validation of the software producing the expected 10% noise increase when supplied with a similar amount of noise across the band, an increase of 50% for VLASS data (which has more noise at higher frequencies) may be consistent with fundamental expectations. I.e., this is the expected noise in wide-band PB-corrected continuum image, independent of the algorithm used for imaging. Note that in this test, only two OTF pointings (separated by 3 arcmin) were used.

## 1.5 Image/Cell Size And The Effect Of Aliasing

For this study, we started with the imaging parameters as used in the VLASS imaging pipeline (image size=6K pixels, pixel size=0.6 arcsec and about 20 pointings across two scan rows). Those were found to result in aliasing strong enough to affect reconstructed flux values on source. We think that similar errors due to aliasing (which can also lead ultimately to divergence of deconvolution iterations) may be present in images made using these parameters.



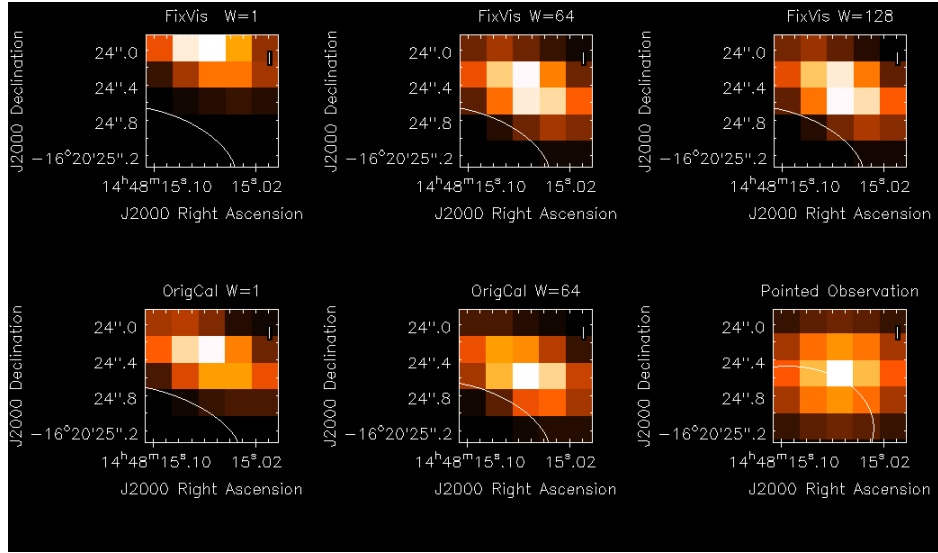


Figure 1.3: Position offsets with and without W-term correction, compared to the true source position.

Figure 1.2 shows the wide-band continuum sensitivity pattern for VLASS mosaic image. This shows that the image size should be such that it includes the PB sidelobes and a buffer around it to ensure that the aliasing effects are minimized in the regions of the image with significant sensitivity (this includes the PB sidelobes). Specifically, we estimate that we need a buffer of  $\approx 30$  arcmin between the centers of the edge-pointings and the edge of the mosaicked field for joint imaging. With a frequency bandwidth ratio of 2:1, a reasonable way to determine the image size for wide-band continuum imaging is to use single-pointing beam-width as  $2\lambda/\text{AntennaDiameter}$  (instead of  $\lambda/\text{AntennaDiameter}$ ).

With these considerations, and to minimize the effects of aliasing, we used an images size of 10Kx10K pixels in imaging tests with few pointings. For the VLASS data taken in BnA configuration, we needed a cell size of 0.3 arcsec for critical sampling of the PSF along the shorter axis. Further tapering or uvrangle limits could allow the use of a larger cell size but this has not been evaluated explicitly in this study.

## 1.6 W-Term Effects

Until the time of this investigation, W-term effects were not being considered for VLASS imaging (based on independent simulations done a year ago). However, position offsets were recently reported via independent tests by the VLASS team, and our simulations (described in Sec. 1.1.1 and 1.1.2) that included the effects of the w-term also showed position offsets that required W-projection to correct them.

Figure 1.3 shows position offsets on real VLASS data using W-planes = 1, 64 and 128. The cell size used was 0.3 arcsec and only one pointing was imaged separately to measure the offset. The source intensity in all the images shown in the figures is scaled to emphasize the location of the peak. The bottom right image is from the pointed observation of the calibrator source and this position is taken as the truth. The left and middle images in the bottom row show the positions with W-planes=1 and 64 using the original VLASS data where the source was at the 70% point of the PB gain. In this case, 64 planes sufficed to move the source to the correct location. A second test was performed with the same pointing after application of 'fixvis' to partially correct an offset between pointing and phase reference center (see Sec. 3.2). This moved the source to the 52% point of the PB and at this distance from the phase reference (and pointing) center, it required 128 W-planes to move the source back to the correct location. The top row of images illustrates the shift versus number of w-planes for this modified dataset. For the purpose of algorithm validation, this sufficed to demonstrate that W-projection is working correctly with the 'awproject' gridded. Previous independent tests using other VLA low-frequency data also confirm this.

### 1.6.1 How many w-planes?

Our estimate of the number of w-planes required for the VLASS data taken in B and BnA configurations is 128 and 512 respectively. This must be re-evaluated after resolving pointing and phase center issues (Sec. 3.2) and after decisions about image cell size and PSF shaping via tapering and uv-range selection are finalized by the VLASS team. In particular, the number of w-planes required may be reduced by increasing the cell size to 0.6 arcsec as desired by the VLASS team and to also select out (or taper down) long baseline data to produce a PSF that is adequately sampled with a 0.6 arcsec cell size for both B-config and BnA config.

Note that selection on w-range (as against uv-range) to decide the number of W-planes is a more effective way of reducing the compute load and inaccuracy of using too few w-planes. The W-Projection implementation in AW-Projection bins w-values using the full range of w-values present in the selected data. With w-selection, computing for larger w-values (the dominant source of the extra computing) will be eliminated (see Sec. 4.1). Tapering for PSF shaping on the other hand would still require the full number of w-planes to be calculated (so that the quantization for the intermediate w-values is accurate enough) and data gridded, only to be finally weighted-down. A selection based approach would at least prevent the most expensive visibilities from being gridded and then down-weighted.

## Chapter 2

# Understanding VLASS Data

**Test Data Subset :** VLASS data for a pointed observation of the J1448-1620 calibrator as well as one OTF scan row covering the same region were used for these tests (see Fig.2.1). The series of pointings considered were field IDs in the range 18952 - 18960, with the source passing closest to the center of the beam between pointings 18956 and 18957 at about the 85% PB point and extending out to the 20% PB gain at both edges of this row of OTF pointings. Wideband imaging of the pointed calibrator observation yielded an intensity of 1.38 Jy and spectral index of -0.46. This was independently confirmed by comparing with the VLASS Quick-look pipeline weblogs as well as calibrator monitoring data from L. Sjouwerman that included this VLASS dataset. These values and source position were then treated as the truth against which all other imaging was compared.

**The Problem :** Imaging with `deconvolver='mtmfs'` and `gridding='awproject'`, using code that was tested to be working on simulated VLASS data did not yield the correct source intensity or spectral index when applied to the real OTF VLASS data as obtained from the VLASS calibration pipeline. The intensity was lower than it should be and spectral index was steeper by different amounts depending on which pointings were included in the imaging. Additionally numerous artifacts were present around the source, preventing an evaluation of noise levels. A series of tests were done to investigate this (sections 2.1,2.2, 2.3).

**The Primary Cause :** The source of the dominant problem was eventually traced (along with parallel efforts from the VLASS and VLA data acquisition teams) to what is an approximately one integration offset between the antenna pointing directions and the correlator phase reference direction for antennas with old ACUs as well as an additional half-integration pointing offset per integration as an OTF mosaic design choice (see Sec.2.3). The generation and use of an accurate high time-resolution antenna pointing table is a straightforward solution to both problems for existing as well as future OTF data, but this will require extra development effort at multiple stages and will also results in an extra computing load during gridding. Imaging tests must be repeated once the above is resolved.




Figure 2.1: J1448-1620 source position w.r.to a row of OTF pointings used for these tests. At closest approach, the source was at the 85% PB point, going out to the 20% PB gain at the farthest pointings included in these tests. The blue circles indicate the pointing directions at the middle of each OTF subscan and the red circle is the calibrator pointing.

## 2.1 Data Weights

The first issue encountered was the uneven distribution of data weights resulting in a dip in mosaic PB gain down to the 40% PB gain for the pointings nearest the 1.38 Jy calibrator source. This is due to the use of the 'statwt' task in CASA, which is intended to modify the data weights to be proportional to System Temperature. However, as implemented, it leads to demonstrably incorrect reduction of the weights in an OTF mosaic observation. In particular, it amounts to introducing closure errors which leads to serious and significant problems in imaging (incorrect or insufficient PSF removal, and divergence in extreme cases). Also, an artificial decrease in sensitivity (and increased error bars) for the brightest sources is not generally desirable, especially along with the resulting position-dependent PSF shapes. ( Varying degrees of RFI across the OTF scans is another reason for uneven data weights (and potentially position-dependent PSFs) but this is estimated to be a smaller effect than the above in terms of the fraction of data affected. )

Therefore for our tests, we set all data weights to 1.0.

The following are suggestions to the VLASS team about ways to avoid this.

1. Run the 'statwt' task on residual data instead of the raw (or calibrated) data. The reason is that the 'statwt' algorithm is an outlier detector that assumes that the underlying signal (the mean value) is invariant. When presented with multiple pointings of OTF mosaic data, it sees a source moving through the primary beam as a transient increase in noise and therefore chooses to weight-down the data for those pointings. Running 'statwt' on residuals is likely to be more robust to the presence of bright sources. (As of this writing, the VLASS team has done preliminary tests of this approach, with some success).
2. Request the CASA team that the 'statwt' task be augmented to use a running mean (or similar) approach in order to be immune to smooth variations of the underlying true signals.

Once the data weights in the data are improved, it also may be useful to re-evaluate the need for features recently added to the 'mosaic' grider to handle the consequences of uneven weighting (position-dependent PSFs and field-dependent weighting schemes).

## 2.2 Self-Calibration

Two adjacent pointings (18957,18958) of the original VLASS data were imaged using multi-term wideband imaging along with grider 'awproject' and 64 W-planes (as determined from VLASS simulations). The resulting images are shown on the right panels of Fig.2.2. Intensities were lower than expected, the spectral index was too steep, and artifacts are clearly present high above the noise floor.

### 2.2.1 Dynamic Range Improvements

Bandpass self-calibration per pointing (or subscan) was attempted in order to assess the root cause of these artifacts. The known true intensity and spectral index of the calibrator was used to predict model visibilities for each of the OTF scans using the 'awproject' grider with wprojplanes=64, against which bandpass amplitude and phase solutions were then computed.

As seen from the left panels of Fig.2.2 the application of the solutions did reduce artifacts enough to see and therefore evaluate noise levels. However, self-calibration is not a viable option for all VLASS OTF data and so the root causes must be addressed.

### 2.2.2 Interpreting the self-cal solutions

**Amplitudes :** Bandpass solutions showed spectral slopes in amplitudes for all antennas with old ACUs (see Fig.2.3). This was eventually traced to the issue described in Sec.2.3 where these antennas had systematic pointing offsets. In particular, bandpass amplitude slopes were absent only in antennas ea01, ea02, ea08, ea14, ea17, ea21, which happen to be a subset of the list of antennas with new ACUs as of Sept 2018 (antennas 1, 2, 7, 8, 14, 16, 17, 21, 27, 28). A partial workaround (for debugging purposes) was evaluated independently by F.Schinkel (VLASS) and replicated for these

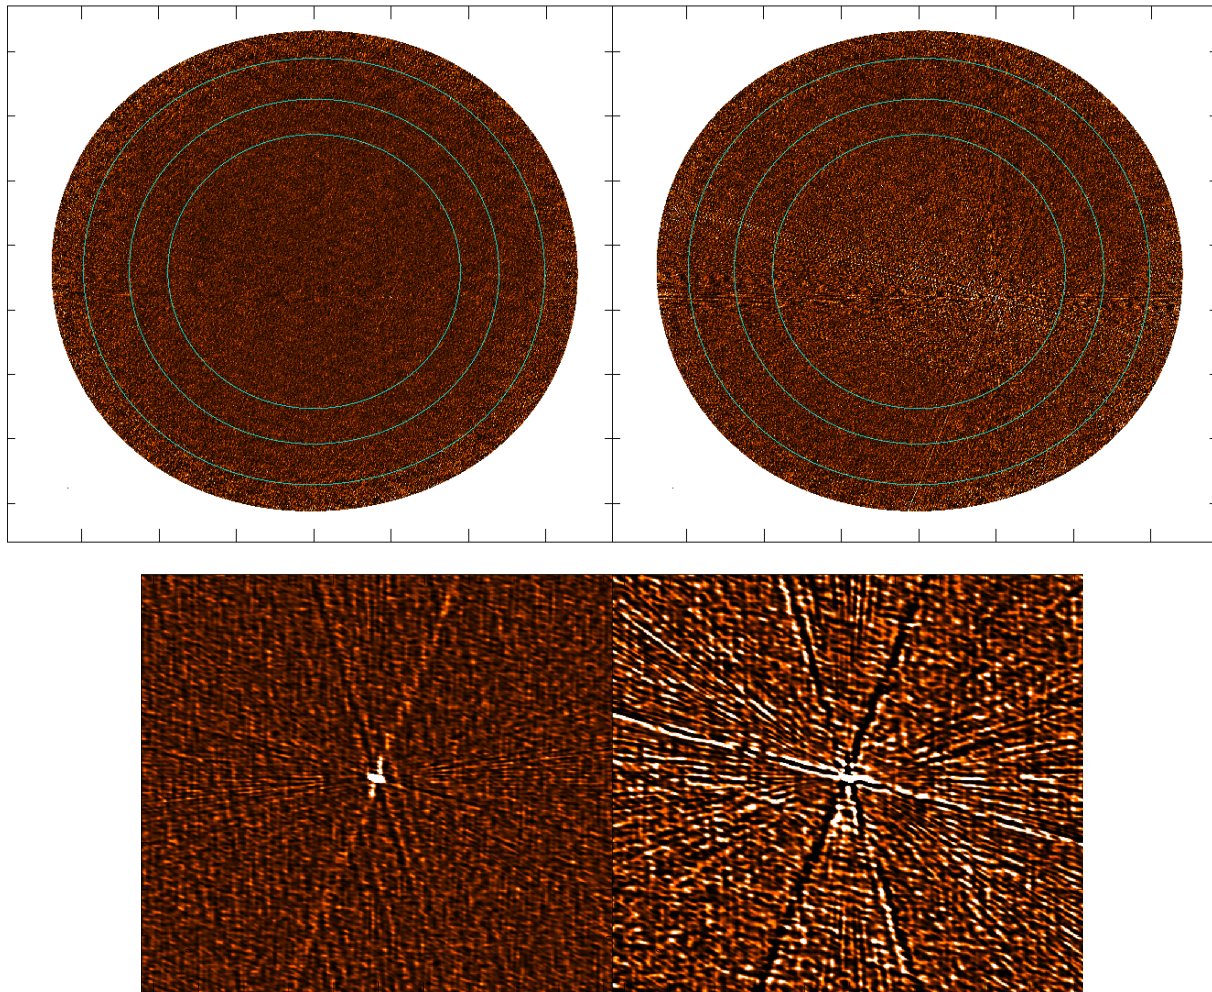


Figure 2.2: Two-pointing mosaic after (LEFT) and before (RIGHT) self-calibration. Artifacts present above the noise floor in the originally calibrated data reduced sufficiently to be able to see the noise floor. The top pair of images show the full field and the bottom pair are zoomed in to the calibrator source. The self-calibration solutions from this test were used to diagnose the root causes of the artifacts. Bandpass phase solutions showed systematic gradients per spectral window and their correction made the most difference in converting dynamic-range limited images to noise-limited image. Amplitude solutions had features consistent with pointing offsets, and once they were worked-around (for this data set with only one bright source), the solutions showed no further systematic effects.



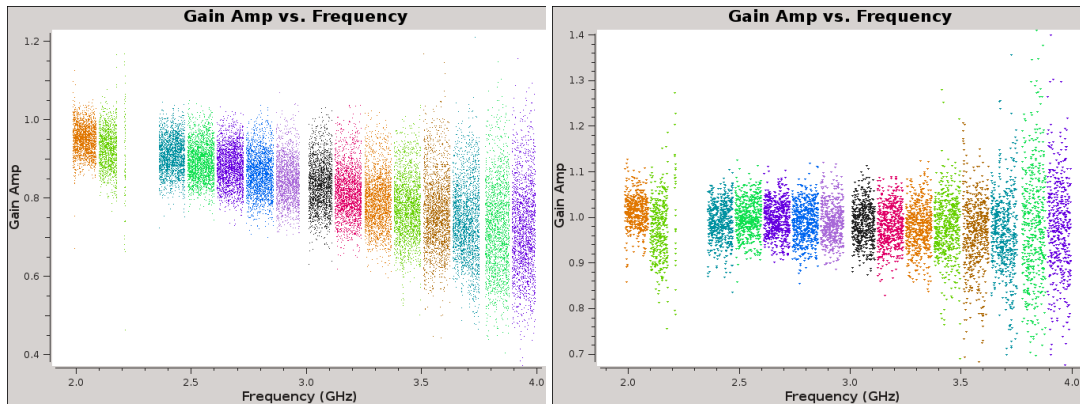


Figure 2.3: Original VLASS data : Self-Calibration solution amplitudes : Old ACUs (LEFT) and New ACUs (RIGHT). Antennas with old ACUs show evidence of a systematic pointing offset. Other antennas do not.




Figure 2.4: Fixvis on VLASS data : Self-Calibration Amplitudes : Old ACUs (LEFT) and New ACUs (RIGHT). Antennas with old ACUs are now correct, but the others have a new artificial pointing offset in the opposite direction.

data. It involved running 'fixvis' to rephase the data from each pointing (or subscan) to a position halfway between the current and next pointing (or subscan), i.e. a one-integration offset. This worked for all antennas with old ACUs as their solutions were now flat (see Fig.2.4). This however introduced the opposite problem for the antennas that were previously correct (new ACUs), but since this involved only 6 out of 26 antennas it was a minor effect in comparison to the original situation. This is of course not meant to be an actual solution, but it was sufficient to demonstrate and understand the root problem (described in Sec.2.3).

Note that this self-calibration scheme can be used for other existing VLASS data of OTF scans near pointed calibrator observations, specifically for the purpose of determining which antennas show the above effects and which don't. This might be useful in order to assist in the creation of accurate pointing tables for use for all existing VLASS data. (This would also pick out antennas with significant enough pointing offsets (due to other reasons)).

**Phases :** Several antennas showed strong phase gradients per spectral window, anchored to zero at the centers of each spectral window. Fig.2.5 shows plots of the phase solutions vs frequency for all antennas (LEFT) and for one antenna (RIGHT) in which the effect was most prominent. These plots were relatively unchanged before/after 'fixvis' and the phase solution features were not correlated with whether the antennas had new or old ACUs.

**Effect of pointing offsets vs phase gradients** Imaging tests were done on the VLASS data before and after 'fixvis' and also before and after applying self-calibration solutions. The results suggest that the pointing offsets were the main reason for on-source errors in reconstructed intensity and spectral index but that the residual phase errors were the dominant source of imaging artifacts around the bright source that limited the dynamic range.

Therefore, in addition to properly handling pointing offsets, better a-priori phase calibration is essential for VLASS

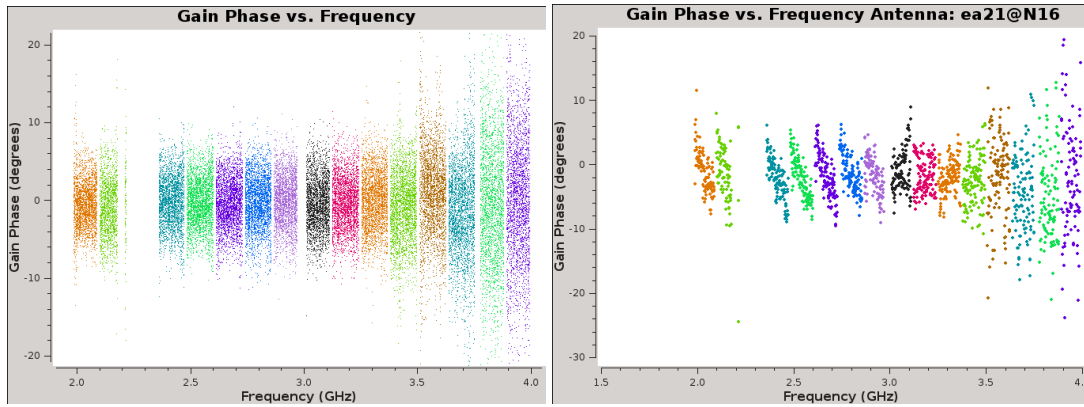


Figure 2.5: Self-Calibration Phase solutions : All antennas (LEFT) and details on one antenna showing possible delay-clunking errors (RIGHT)

Epoch 1 imaging to reach the desired noise levels and dynamic ranges. The specific features in the right panel of Fig.2.5 suggest the presence of the already-known VLA 'delay-clunking' phase error that has not been calibrated out by the VLASS pipeline. This detail must be revisited after the issues in Sec.2.3 are resolved. (Obit Memo #55 by B.Cotton also details various similar calibration effects noticed in VLA S-Band data and may be a useful reference.)

## 2.3 VLASS Pointing Offsets

VLASS data has two types of inconsistencies between the pointing direction of antennas and the correlator phase center. One is an error during data acquisition and one is a design choice. Both have significant effects on the final imaging if left unaccounted for, and neither have ready-to-use solutions in the software. This study has evaluated and quantified the effects of these pointing offsets along with predictions that match what is seen in the real VLASS data.

VLASS observations are set up with an OTF antenna scanning rate of 3.31 arcmin/second, 0.45 second integrations and 2 integrations per subscan. Each subscan is set up to have a unique correlator phase center matched to the middle of the range of positions covered by the continuously moving antenna primary beams in that subscan. Therefore, each integration effectively has a pointing offset of upto 0.74 arcmin.

Fig.2.6 illustrates our current understanding of the situation. Time is shown on the X-axis and position on the sky is shown on the Y-axis. The green dotted line is where the antennas are pointing in their continuous scanning mode. The blue circles represent individual integrations, with each pair belonging to one subscan. Their abscissa (Y-coordinate) is the correlator phase center direction. The left panel shows the intended observing setup with the correlator phase center for each pair of integrations straddling the true pointing direction.

### 2.3.1 Type 1 : Offsets per subscan with old vs new ACUs :

The bottom panel of Fig.2.6 shows what the VLASS data contain for all antennas with old ACUs. Antennas with new ACUs exhibit the desired behaviour (shown on the left panel). For antennas with old ACUs, the correlator phase center is not matched with the center of the antenna scanning position. The root cause is currently being investigated (and addressed) by the VLA data acquisition team, and possible solutions are being investigated to (a) fix the problem before further observations and (b) generate accurate pointing table information for use when imaging existing data.

### 2.3.2 Type 2 : Offsets per integration within each subscan

Even with the desired behaviour (top panel of Fig.2.6), the correlator phase center for each integration is still offset from the true pointing direction. In the top panel, the blue arrows represent the distance on the sky over which the PB has moved for each integration that has been phased up at the location of the blue circles. This is a one-sided PB smearing that produces an effective pointing offset with every alternate integration being offset in opposite directions. Section 2.3.3

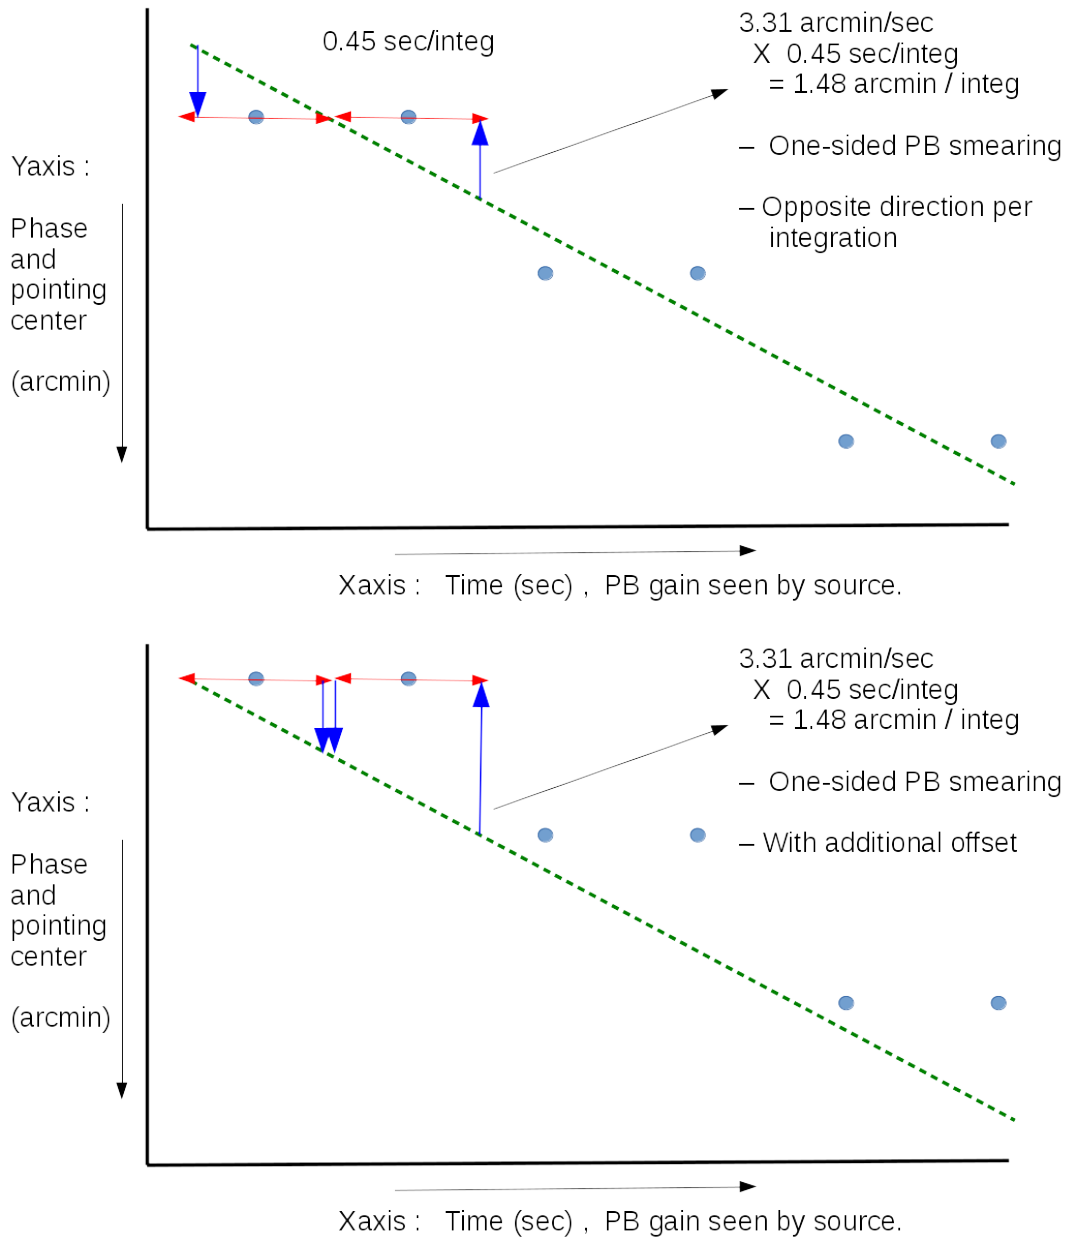


Figure 2.6: TOP: Intended VLASS observation setup. The correlator phase center is fixed for every subscan (one pair of integrations) and synchronized with the true pointing direction at the middle of the subscan. This is the situation for antennas with new ACUs. BOTTOM : Actual VLASS observation for antennas with old ACUs. The correlator phase center is offset by (roughly) one integration, causing an additional effective pointing offset.






Figure 2.7: Effect of 1.48 arcmin of PB smearing. LEFT : Asymmetric smearing. RIGHT : Symmetric smearing. In each set of plots, the left panels show the PB gain and spectral index in solid lines and their smeared counterparts as dotted lines. The right panels show differences between the original and smeared beams for intensity(top) and spectral index(bottom). The black horizontal and vertical lines are meant to guide the eye by marking the magnitudes of various terms at the half power position of the PB gain.

quantifies this effect. This was a design feature chosen on purpose to accommodate data rate limits within the correlator. It's significance on the final imaging was not understood prior to this study and is discussed in Sec.2.3.4.

### 2.3.3 OTF scans and PB Smearing

1-D cuts through the VLA S-Band average primary beam were taken, representing absolute gain as well as PB spectral index. Fig.2.7 shows these 1-D plots along with the results of one-sided smearing by 1.48 arcmin (simulated by averaging the appropriate set of samples on one side of each sample) on the LEFT and symmetric smearing on the RIGHT. The difference between the original and smeared beam shapes are shown in the cyan and yellow plots. Following the solid black lines, one can see that a one-sided smearing at the half power level at the reference frequency will cause a spectral index error of upto 0.35. Symmetric smearing however has negligible effect.

The imaging results in the next section show offsets in intensity and spectral index that match these calculations for one-sided PB smearing, thus confirming our understanding of the consequences of the intended VLASS pointing offsets on imaging. The results for symmetric smearing also suggest that the solution only needs to account for the pointing offset per integration, moving the primary beam to the middle of the OTF scan range per integration, and that it need not include the actual smearing (this last detail is consistent with initial analysis by the VLASS team that concluded that OTF PB smearing may be ignored in imaging).

### 2.3.4 Effects on Wideband Imaging

During imaging, the above pointing offsets currently result in the imaging code placing primary beams at incorrect locations compared to where the antennas were actually pointing. These offsets result in PB spectral index differences of up to 0.6 (>100%) for pointings where the source was near the 30% PB gain, producing differences in imaging results that depended on which pointings were included in the mosaic.

For a joint mosaic, the expectation was that all these errors will average out. However, it does not. This is because the location of the source in the primary beam is also constantly changing with time and the primary beam spectral index is a very strong function of position in the primary beam. Therefore, a one-sided smearing at (say) the 70% PB gain does not produce the same offset as a smearing at the 50% PB point which is where the source would be after 0.45 seconds (i.e. one integration). In order to ensure that averaging does indeed solve the problem, every source must be imaged with the exact same number of pointings (and weights) on either side of it in each OTF scan row. Even if the errors do average out perfectly, the error bars on spectral index will still be several 100% even considering only those pointings where the source is in the main lobe. Note that there are also natural random antenna pointing errors of a magnitude similar to this (upto 2 to




Figure 2.8: Antenna Azimuth vs Phase center for the series of integrations in the VLASS data used for these tests. RED dots indicate the phase-centers of the original data.

3 arcmin on some antennas). A similar argument will apply in the context of deriving error bars on recovered spectral index.

**Single Pointing Wideband Imaging :** The test VLASS data were imaged one pointing at a time in order to understand and demonstrate these effects clearly. Three versions of the data were used depending on different ways of using 'fixvis', and one set of data was a simulation using the known true source parameters. Fig.2.8 shows the phase centers per integration for the three versions of the real VLASS data, similar to the diagrams shown in Fig.2.6. RED dots indicate the phase-centers of the original data. BLUE dots represent the phase centers after applying a 'fixvis' operation that re-phases the array to the intended direction on the sky per subscan (a shift of 1 integration for all antennas, so it is still only approximate). GREEN dots represent the result of applying 'fixvis' per integration by 1/2 and 3/2 integration for every alternate sample (also for all antennas), and is likely to be closest to the true pointing directions.

Each integration was then imaged separately using wideband AWPProjection to recover the sky intensity and spectral index. Intensities, spectral indices and PB gain values are shown in Fig.2.9.

1. TOP LEFT : The original data show a clear systematic global offset that is minimized when the source is at the minimum distance from the OTF scan line. This is the erroneous offset present on all antennas with old ACUs. On top of this feature is an alternating positive and negative offset, arising due to each integration being correlated at phase centers either ahead of or behind the true pointing direction.
2. TOP RIGHT : The application of 'fixvis' to all antennas corrects for the large systematic offset, although since this correction was applied to all antennas, it produces a minor over-correction (there were 6 antennas with new ACUs and the rest had old ACUs). The offsets within each subscan remain. This suggests that even when the error in the online system is fixed, the per-integration offsets are still significant.
3. BOTTOM LEFT : In the bottom-left figure, these offsets were smoothed out by re-phasing each integration within each subscan to the correct PB location. Here too, the corrections were applied to all antennas, leading to a systematic over-correction. But, this test shows that if the correct pointing directions are used during imaging (per antenna), the wideband imaging is likely to be accurate.
4. BOTTOM RIGHT : A simulation onto these same set of pointings yielded accurate recovered intensity and spectral indices. This test was a control experiment to ensure self-consistency of the software on the same VLASS observing setup.

**Joint Mosaic Wideband Imaging :** Joint mosaic imaging of the entire scan row got the intensity and spectral index values close to the correct value only when care was taken to include a symmetric set of pointings on either side of the source. Intensity = 1.31 and alpha = -0.56 were obtained and this matches the calculations shown by dotted lines in Fig.2.9. Fig.2.10 are example images from the original data as well as a control simulation. The artifacts present in the

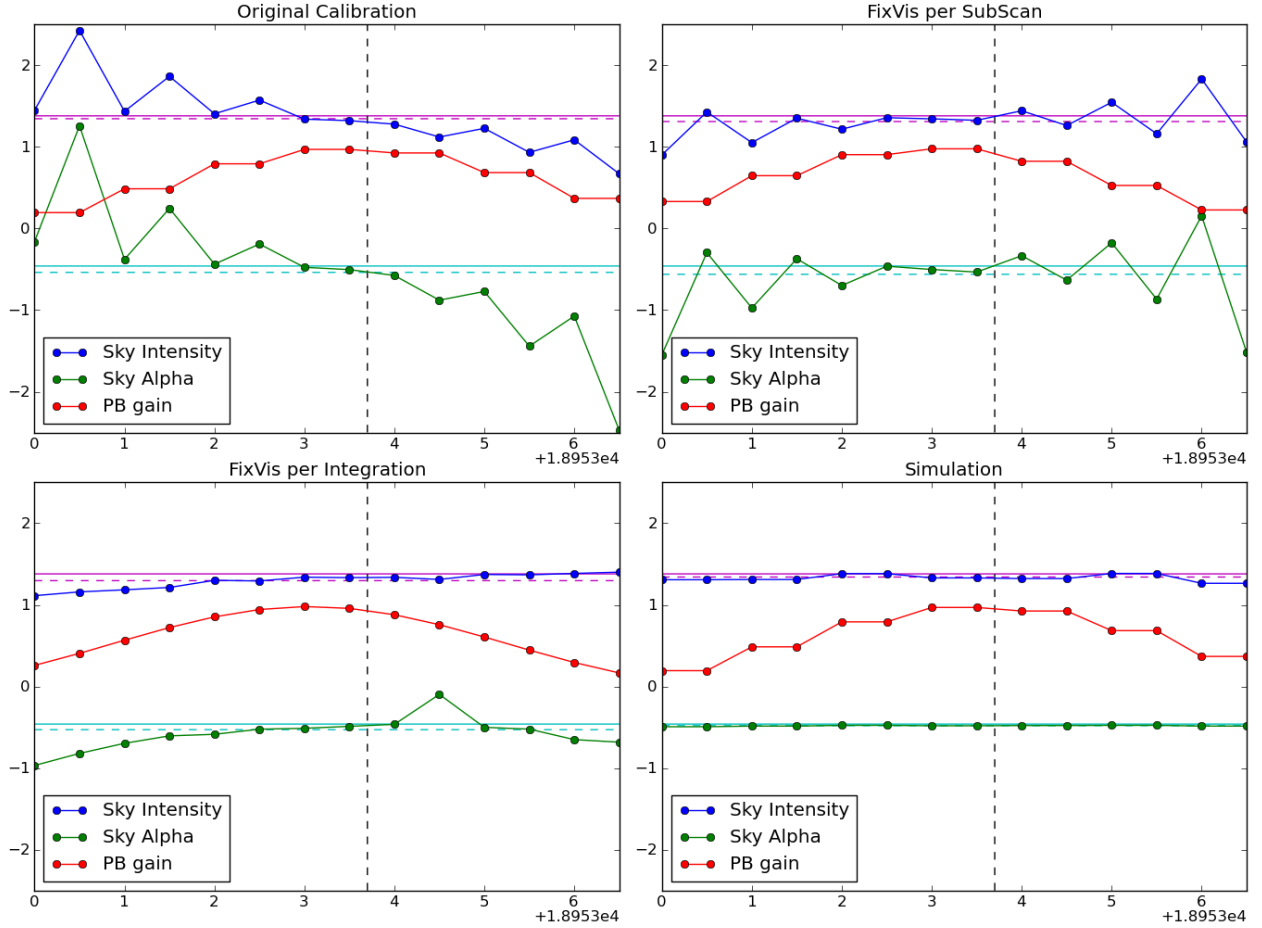


Figure 2.9: Imaging each pointing separately. TOP LEFT : Original Data. TOP RIGHT : Fixvis shift by 1 integration on all antennas. BOTTOM LEFT : Fixvis shift by 1/2 and 3/2 per integration pair, for all antennas. BOTTOM RIGHT : Simulation onto original VLASS data. The X-axis shows the field IDs of the data used (one scan row). The Blue dots show sky intensities with the magenta solid line showing the known true value. The magenta dotted line is an average across all integrations in the test, weighted by the PB gain. Green dots represent spectral index for each integration, with the cyan solid and dotted lines showing the ideal and average obtained values. Red dots show the PB gain level. The pairing in the values indicates the pair of integrations per subscan for which the imaging software places the PB at the same location. The vertical dotted line indicates the position at which the source is at its closest to the pointing center. The true intensity is 1.38 Jy and true spectral index is -0.46.

simulated image are consistent with the PSF shape, lower than the VLASS noise threshold of  $100\mu Jy$ , and shown via independent tests to decrease further with more iterations. However to maintain consistency between our tests, a fixed set of iterations were used.

Note that dynamic range limits still exist even after various 'fixvis' style corrections and it is expected that phase-self-calibration solutions (similar to Fig.2.2) will be required before the true noise level in the mosaic will be visible.

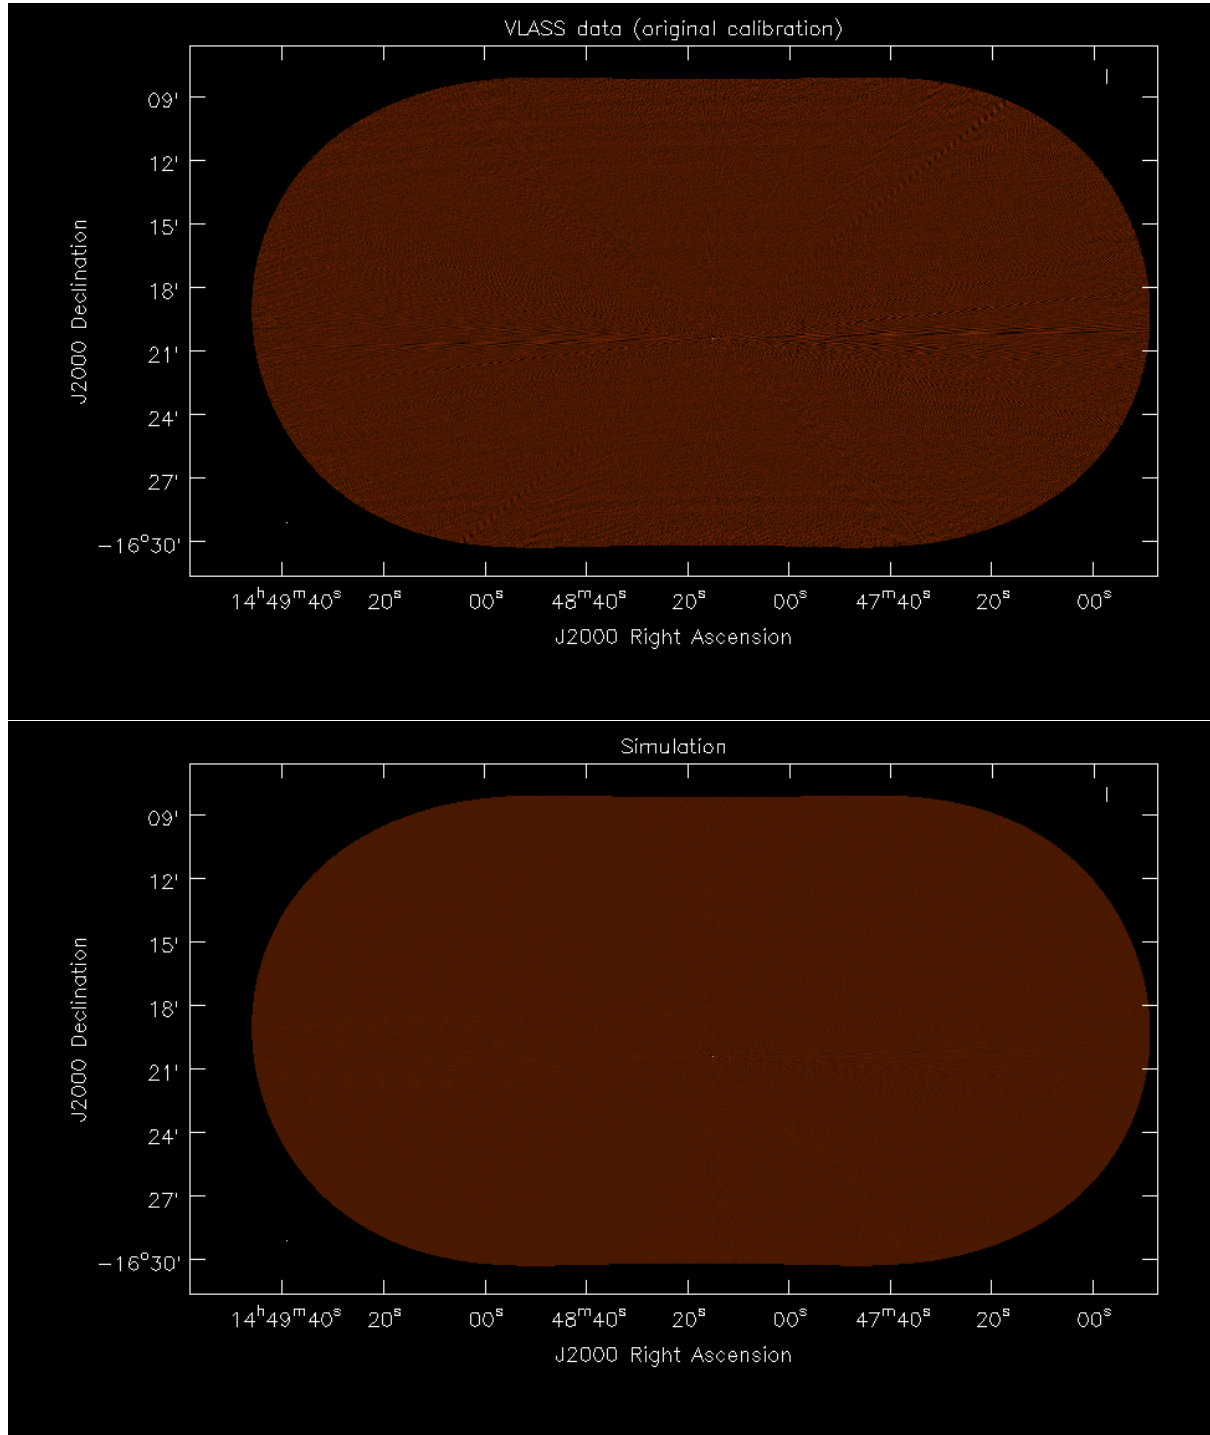


Figure 2.10: Joint mosaic of a row of pointings (with niter=500). TOP : Original VLASS data. BOTTOM : Simulated VLASS data. The numerical noise/artifacts in the simulated image are the  $100 - 50 \mu Jy$  level. This is below the desired/expected noise level of  $100 \mu Jy$  for VLASS, especially if the data are calibrated well enough to allow deeper deconvolution. The real VLASS data currently produces a noise or artifact level that is 5 to 10 times higher than this, at approximately  $600 \mu Jy$ . With the original data, intensity = 1.31 Jy and  $\alpha = -0.57$  when pointings are carefully selected to be symmetric about the source in order to average out the errors due to phase and pointing offsets. (Ideal values are intensity=1.38 Jy and  $\alpha=-0.46$ )

## Chapter 3

# Required Development Efforts

### 3.1 CASA, Gridders, Imaging Infrastructure

The following is a list of issues that need work. As per our tests, items marked with ★ are critical for VLASS.

1. To-Do list for AW-Projection code (gridder='awproject'):
  - BUG: Fix cube imaging bug.
  - ★BUG: Fix pointing table reading bug
  - Improvement: Write-out the PS-term as a separate image for use in making final flat-sky images
  - Improvement: Acquire PB frequency scaling per channel and using measured holography models (fix steepening of 0.02)
  - Improvement: Make CF cache reusable for different numbers of w-planes
  - Improvement: Package script to make CF Cache independent of imaging major cycle. ( via a 'makeimage' like method )
  - Improvement: Enable the use of non-square images
  - ★New Feature : Ability to apply antenna-dependent pointing offset corrections
2. To-Do list for MosaicFT code (gridder='mosaic') :
  - ★BUG: Fix use of the conjbeams parameter (CAS-11902)
  - ★BUG: Fix Gibbs ringing in the images<sup>1</sup> (Test fixes from CAS-10701 and CAS-11183 for VLASS).
  - Improvement: Allow use of the PS-term
  - Improvement: Use PB model that includes PB sidelobes correctly
  - ★New Feature : W-term support (if this gridder is to be used for VLASS at full resolution)
  - ★New Feature : Ability to apply antenna-dependent pointing offset corrections
3. To-Do list for imaging infrastructure code
  - Factor-out the code for using the POINTING table so that same code is used everywhere it is required
  - Evaluate possible issues with PB sky projection and solution in the infrastructure code
  - Unify features such as field-dependent data weighting and position dependent PSFs between the gridders.
4. Fix the 'statwt' task
  - If the use of 'statwt' task is deemed essential, use robust statistics and move-averages to compute the modification to the data weights
5. ★Make VLASS pointing tables ready for use by imaging (CAS-11191).

---

<sup>1</sup>See [http://www.aoc.nrao.edu/~sbhatnag/misc/AWP\\_Mos.pdf](http://www.aoc.nrao.edu/~sbhatnag/misc/AWP_Mos.pdf)

### 3.2 Addressing phase/pointing offsets at all levels

As of this writing, solutions to the VLASS pointing offsets described in Sec.2.3 are being evaluated both at the data acquisition stage as well as in the imaging stage. Note that even if future observations do not have the erroneous  $\sim 1$  integration offset for antennas with old ACUs, accurate pointing tables and related algorithmic features will still be required for imaging in order to handle the remaining intentional  $\sim 1/2$  integration offsets. This will require an independent commissioning effort involving multiple development teams.

1. ★Fix the offsets during real-time data acquisition, for future OTF observations. This is already under way by the online team.
2. ★Generate accurate pointing tables for old as well as future VLASS datasets. For existing data, they must include the new/old ACU offsets as well as the per-integration intentional offsets. For future data they must include only the true pointing direction.
3. ★Accurate translation of these tables into CASA Measurement sets. Until now, a feature relating to time-stamping of pointing table data has prevented its use within CASA. CAS-11191 describes ongoing efforts.
4. ★Use the pointing information during imaging via the application of baseline-based phase gradients during gridding (see sec.4.2). This feature has been prototyped along with WB-AWProjection via pointing self-calibration algorithm R&D work, but is currently not part of either the ARDG code for 'awproject' being tested for VLASS imaging or the existing 'mosaic' grider.
5. ★Repeat the 'row of pointing' tests done as part of this study.

### 3.3 VLASS Pipeline

The following is likely to require VLASS commissioning effort and possibly development in VLASS pipeline.

1. ★**Improving 'statwt' usage** : The current pre-imaging data processing pipeline uses the 'statwt' task to modify the data weights in an attempt to make them proportional to the system temperature. This task was primarily designed for single-pointing data and leads to erroneous weights for OTF mosaic data and effectively introduces closure errors in the data from pointings with even moderately strong emission (see sec.2.1 for details).
2. **Improving Phase Calibration** : On the test data used for this study, a significant improvement in the imaging performance and dynamic range was observed after wide-band phase self-calibration was done (see Sec. 2.2). In this example, it converted artifact-limited imaging into noise-limited imaging. The root cause of why this was required needs to be determined and addressed, preferably prior to imaging.
3. ★**Re-determining Imaging Parameters** : As per recent conversations with the VLASS team, there are plans to re-evaluate imaging parameters in ways to reduce compute load. These include the downweighting (or deselection of) long baselines in order to produce a lower angular resolution PSF. This would then allow the use of larger image cell sizes and therefore a smaller total number of pixels. If chosen carefully, this may also have the desirable effect of reducing the effect of W-term related position offsets, requiring fewer W-planes during imaging and therefore less computing. We suggest that attention be paid to the following details when re-evaluating parameter choices.
  - Effect of aliasing for the chosen image cell size, image size, and number of OTF pointings included in each imaging run.
  - Number of W-planes required for a source at (say) the 20% PB gain level to get accurate positions. A large cell size is likely to hide small position shifts.
  - The effect of PSF shaping via uv-tapering vs uvrange selection on gridding compute load.
  - Effect of undersampling the PSF (similar to what was chosen for VLASS Quick look imaging).



## Chapter 4

# Computing Estimates

### 4.1 Gridding Load With AWProjection

The computing cost of Stokes-I and spectral index mapping of the OTF mosaic imaging is dominated by the cost of gridding the data on to a regular grid for wide-band imaging. The computing cost of gridding scales as the sum of the product of the number of visibility points and the square of the convolution function (CF) support used for gridding those visibilities. The number of data points and the CF support size varies with the value of the  $w$ -co-ordinate and frequency (SPW). In the discussion below, we refer to this cost as FLOP (FLoating-point OPerations).

$$FLOP \propto \sum_{w, spw} N_{vis}(w, spw) \times N_{Support}^2(w, spw) \quad (4.1)$$

where  $N_{vis}$  is the number of visibility points at a given  $w$  within a given  $spw$  and  $N_{Support}$  is the support size of the CF used to grid these data. The ratio of the FLOP for 'mosaicft' and 'awproject' gridders as a function of  $w$  for each spectral window is given by

$$R_{AWP/Mos}(w) = \frac{\sum_{spw} N_{vis}(w, spw) \times N_{Support,AWP}^2(w, spw)}{\sum_{spw} N_{vis}(spw) \times N_{Support,Mos}^2(spw)} \quad (4.2)$$

The aggregate extra computing load due to A- and W-projection together can be estimated by taking the ratio after summing over  $w$  the numerator and denominator of Eq. 4.2 separately. This ratio comes out to be  $\sim 50$  for the test data set used in this study. The computing load of W-Projection (in addition to A-Projection) is therefore expected to be a factor of 50 larger than with 'mosaic' grider.

Figure 4.1 shows  $N_{vis}$  and  $R_{AWP/Mos}$ , both as a function of  $w$  for a typical VLASS data set. The number of visibility points for larger  $w$ -values is relatively smaller.  $R$  increases with  $w$  due to increase in the area under the CFs with  $w$  (Fig. 4.2). This suggests that it may be possible to reduce to extra computing load by carefully selecting-out some baselines on the outer edges of the  $w$ -axis. Its impact on sensitivity and resolution needs further exploration to decide if this is useful.

### 4.2 Correction Of Antenna-dependent Pointing Errors

The effect of erroneous pointing offsets in VLASS data is discussed in Sec. 2.3. The effect of antenna pointing offsets is a Direction Dependent (DD) effect and fundamentally cannot be corrected as Direction Independent (DI) calibration term. I.e., its correction must be included during imaging. A possible solution (see Sec. 3.2 for details) is to generate a pointing calibration table and use it during imaging to apply the corrections for pointing offsets. The pointing offset due to OTF setup is global (array-wide) offset while the offset due to a bug in the online software (in the new ACUs) is antenna dependent. While the latter will be fixed before acquiring new data, the effective pointing error in existing VLASS data is antenna-dependent.

Correction for array-wide offset is relatively simple and can be applied during imaging with minimal extra computing load. Correction for antenna-dependent pointing offsets requires modifying the A-term per affected baseline as (written






Figure 4.1: The blue curve shows  $N_{vis}$  as a function of  $w$  for a typical VLASS dataset. The orange curve shows the ratio of FLOP for 'awproject' and 'mosaic' gridders ( $R_{AWS/Mos}$ ) as a function of  $w$ . It measures the extra compute-load of AW-Projection.




Figure 4.2: The plot shows the area under the convolution functions (CF) including the  $w$ -term as a function of  $w$ . The computing load is a strong function of the area under the CF and total cumulative load is dominated by CFs for large  $w$ .

for one dimension for clarity)<sup>1</sup>:

$$A_{ij} = A_{ij}^{\circ} \left[ e^{-\frac{(l_i - l_j)^2 \alpha^2}{2}} \right] e^{i\pi u(l_i + l_j)} \quad (4.3)$$

where  $A_{ij}$  is the modified A-term,  $A_{ij}^{\circ}$  is the original A-term computed without antenna-dependent pointing offsets,  $l_i$  and  $l_j$  are the pointing offsets for the two antennas and  $\alpha$  is  $2^{-1/2}$  times the inverse of the beam-width. Given  $A_{ij}^{\circ}$ , this modified A-term can be computed on-the-fly. However that adds extra computing load which can be comparable to the gridding load for the affected baselines. The alternative of pre-computing  $A_{ij}$  would lead to prohibitively large memory footprint or i/o load.

The complex exponential in Eq. 4.3 is the phase ramp across the A-term. This is similar to the mosaic-term currently being computed in both gridders, except that the mosaic-term is array-wide which is computed and stored (in memory) as a single copy for all baselines and computed only for finite change in the phase center. The phase term in Eq. 4.3 on the other hand needs to be computed per affected baseline and for every integration dump from those baselines.

The term in square brackets is the reduction in the amplitude of the correlation coefficient due to antenna-pointing related decorrelation and changes across the band (i.e., it is a frequency dependent correction). This is close to unity and can be ignored for small offsets (few percent of the beam-width) – it is unity for baselines where the two pointing offsets are the same. The offsets in VLASS data are in the range of 5 – 10% of the beam width and the effect may not be insignificant. The effect of this term therefore needs to be evaluated carefully before ignoring this term for VLASS also. Including this term will further add to the computing load.

### 4.3 Measurements Of Run-time Performance (Added on Sept. 26, 2018)

Robenett et al. did a test run with `grider='awproject'` and `wprojplanes=128` settings in the `tclean` task using the “weak” and “medium” VLASS test fields to make a measurement of the execution time<sup>2</sup>. Exact same test (using the same ARD-CASA build) were done with the ‘mosaic’ grider to determine the baseline run-time. This run is equivalent to `wprojplanes=1` and identical to what has been used by VLASS till now.

The tests with ‘awproject’ grider were run using the `parallel=True` feature of the `tclean` task, which parallelizes the major cycles. A 8-way parallelization was done. The tests with ‘mosaic’ grider were run in serial mode (`parallel=False`).

The total measured run time for the ‘awproject’ grider was about 60x longer compared to ‘mosaic’ grider. The ‘awproject’ grider for the “weak” case completed 29 major cycles and took 8 days with 8-way parallelization, so around 60 days if run serially. The ‘mosaic’ grider with 22 major cycles took ~1 day (22 hours).

The average major cycle run-time for the ‘mosaic’ grider for the “weak” case was ~50 minutes. For the ‘awproject’ case with 8-way parallelization it was ~310 minutes or 2480 minutes serially. The measured major cycle run-time is therefore about 50x longer for AW-Projection, in good agreement with the prediction in Sec. 4.1. Since the 50x scaling measured here is based off the run time for ‘mosaic’ grider, while the predicted 50x scaling in Sec. 4.1 is based off the run time for ‘awproject’ grider with `wprojplanes=1`, these measurements also imply that the underlying efficiency of both the gridders must also be very similar.

It was noted that for the “medium” field case, the ‘awproject’ converged in about half the major cycles (21 vs 43). This is consistent with theoretical expectations.

<sup>1</sup>Eq. 14 of Sec. 3 in Bhatnagar & Cornwell, 2017, AJ, 154, 197

<sup>2</sup>The runs for the “strong” test field did not converge with either of the gridders.