



OPEN ACCESS

CASA on the Fringe—Development of VLBI Processing Capabilities for CASA

Ilse M. van Bemmél¹, Mark Kettenis¹, Des Small¹, Michael Janssen², George A. Moellenbrock³, Dirk Petry⁴, Ciriaco Goddi^{5,6}, Justin D. Linford³, Kazi L. J. Rygl⁷, Elisabetta Liuzzo⁷, Benito Marcote¹, Olga S. Bayandina^{1,8}, Neal Schweighart³, Marjolein Verkouter¹, Aard Keimpema¹, Arpad Szomoru¹, and Huib Jan van Langevelde^{1,9,10}

¹ Joint Institute for VLBI ERIC (JIVE), Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands

² Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, D-53121 Bonn, Germany

³ National Radio Astronomy Observatory, P.O. Box O, Socorro, NM 87801, USA

⁴ European Southern Observatory, Karl-Schwarzschild-Strasse 2, D-85748 Garching, Germany

⁵ Dipartimento di Fisica, Università degli Studi di Cagliari, SP Monserrato-Sestu km 0.7, I-09042 Monserrato, Italy

⁶ INAF—Osservatorio Astronomico di Cagliari, via della Scienza 5, I-09047 Selargius (CA), Italy

⁷ INAF-Istituto di Radioastronomia & Italian ALMA Regional Centre, Via P. Gobetti 101, I-40129 Bologna, Italy

⁸ INAF—Osservatorio Astrofisico di Arcetri, Largo E. Fermi 5, I-50125 Firenze, Italy

⁹ Leiden Observatory, Leiden University, Postbus 2300, 9513 RA Leiden, The Netherlands

¹⁰ University of New Mexico, Department of Physics and Astronomy, Albuquerque, NM 87131, USA

Received 2022 June 7; accepted 2022 July 18; published 2022 November 15

Abstract

New functionality to process Very Long Baseline Interferometry (VLBI) data has been implemented in the CASA package. This includes two new tasks to handle fringe fitting and VLBI-specific amplitude calibration steps. Existing tasks have been adjusted to handle VLBI visibility data and calibration meta-data properly. With these updates, it is now possible to process VLBI continuum and spectral line observations in CASA. This article describes the development and implementation, and presents an outline for the workflow when calibrating European VLBI Network or Very Long Baseline Array data in CASA. Though the CASA VLBI functionality has already been vetted extensively as part of the Event Horizon Telescope data processing, in this paper we compare results for the same data set processed in CASA and AIPS. We find identical results for the two packages and conclude that CASA in some cases performs better, though it cannot match AIPS for single-core processing time. The new functionality in CASA allows for easy development of pipelines or Jupyter notebooks, and thus contributes to raising VLBI data processing to present day standards for accessibility, reproducibility, and reusability.

Key words: Astronomy software – Very long baseline interferometry – Radio astronomy

1. Introduction

1.1. Background

For decades the Very Long Baseline Interferometry (VLBI) community primarily processed data using the National Radio Astronomy Observatory (NRAO) Astronomical Image Processing System (AIPS) software package (Greisen 2003). However, over the last decade the Common Astronomy Software Application (CASA; McMullin et al. 2007) has replaced AIPS for most *non*-VLBI applications in radio astronomy. Requiring users wishing to process VLBI observations to learn AIPS is a significant barrier to entry. In addition, with the changing hardware architecture and ever growing data volumes, AIPS is running into limitations that are increasingly harder to overcome in software. A particular challenge was the calibration of

high-frequency, global VLBI observations carried out with the Event Horizon Telescope (EHT; Event Horizon Telescope Collaboration 2019b). To overcome this challenge the ERC-funded *BlackHoleCam* project (Goddi et al. 2017), initiated the development of the CASA-VLBI functionality which led to the development of the first CASA-based calibration pipeline for VLBI data (Janssen et al. 2019; Event Horizon Telescope Collaboration 2019a).

This development was preceded by a detailed comparative study among the main radio-interferometric data processing software packages currently in use. The comparison scored each package on its suitability to build a pipeline for VLBI observations by comparing reliability, flexibility, sustainability, user access and support. From this exercise, CASA and AIPS were found to be the best options, with CASA being the prime choice due to continuous and future software development and extensive support for users of large observatories such as ALMA and VLA (see Appendix). The development of the CASA VLBI functionality has in the meantime matured into a



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

joint effort between the Joint Institute for VLBI ERIC (JIVE) in the Netherlands and NRAO in the US to improve the accessibility for (new) users of the European VLBI Network (EVN) and the Very Large Baseline Array (VLBA). In recent years also the option to run CASA in a Jupyterhub environment provides a many new benefits (Keimpema et al. 2020).

Recently the CASA 6 series was released, the Python3-based distribution of the CASA package, available as a modular and a monolithic distribution. A detailed description of the CASA 6 software is presented in CASA Team et al. (2022), hereafter called the CASA reference paper. In this paper we describe the additional functionality required for the processing of VLBI data as it is implemented in CASA 6.4.1. It is therefore entirely complementary to the CASA reference paper.

1.2. CASA Calibration Framework and VLBI

The ongoing introduction of VLBI capabilities in CASA takes full advantage of the modularity of the general CASA calibration model. As described in more detail in the CASA reference paper, this calibration model is based on the Jones matrix formalism of the Hamaker–Bregman–Sault Measurement Equation (Sault et al. 1996). Adding VLBI-specific terms (e.g., fringe fitting) mainly requires introducing calibration-type-specific specializations to the CASA calibration framework. These specializations transparently inherit general interfaces and features, and need only implement their specific properties, including the details of their time- and frequency-dependence, the algebra for calculating their Jones matrix elements, specialized solvers for deriving solutions, and any other specifics. Most VLBI-specific code is therefore well-isolated from existing calibration code, and adopts existing mechanisms for data storage (the MeasurementSet), data iteration for calibration solving and application, pre-calibration for solving, calibration solution storage, plotting, etc. VLBI data processing may also make use of the extensive suite of existing calibration (complex gain, bandpass, instrumental polarization, etc.) and imaging tasks. As a result of this development strategy, the VLBI-specific terms fit naturally into CASA’s generalized (self-)calibration mechanism. At the same time, VLBI-specific capabilities are fully-integrated by default and become seamlessly available to more general CASA processing contexts (e.g., ALMA, JVLA). Additionally, the VLBI development exercise has had a favorable influence on some of the more general aspects of the CASA calibration infrastructure, including a mechanism for creating empty solution tables for any calibration type, solution cadence improvements, solution interpolation, definitions for ancillary calibration (e.g., gain curves), etc.

1.3. This Paper

This paper is organized as follows. In Section 2 we describe the development of new tasks and updates of existing tasks. In

Section 3 we present the main use cases for CASA-VLBI, and a rudimentary comparison with AIPS is given in Section 4. In Section 5 we discuss future development plans, and we wrap up in Section 6 with a brief summary. Throughout the paper we will refer to CASA tasks in **bold**, task parameters in *italic* and AIPS tasks in ALL CAPS.

2. Development

This was not the first attempt to implement VLBI tools in CASA. Past efforts have ended prematurely mainly because they attempted to improve on the existing algorithm. The scope of this work therefore was to replicate the AIPS functionality as closely as possible, and optimize later. This makes the CASA implementation of certain VLBI-tasks slower compared to the AIPS implementation (see Section 4). However, parallelization through the CASA message-passing interface (MPI) implementation can overcome this problem.

We assessed the CASA calibration framework in order to identify the work needed to enable data processing as it is done in the EVN pipeline (see Section 3.1). We started with the data products found in the EVN archive, and assessed all calibration steps. Imaging and further analysis was excluded from this work as these processes are highly dependent on the science goal. Polarisation calibration has also not been handled in detail, though CASA is based on the Measurement Equation and preserves all information needed for this step (for details see Martí-Vidal et al. 2021).

For EVN data all pipeline steps were found to have equivalent tasks in CASA except for the fringe fitting. For VLBA data it was found that CASA had no equivalent task for ACCOR,¹¹ which is required for data processed in the Distributed FX (DiFX) correlator (Deller et al. 2007). For each of these steps, a new CASA task was developed: **fringefit** and **accor**.

Additional work was needed to ensure the proper handling of VLBI meta-data and smooth operations of the new tasks within the CASA calibration framework. This was done in close collaboration with the CASA development team to ensure that CASA would maintain pre-existing functionality.

2.1. New Task: *fringefit*

2.1.1. Solving

The general strategy we employ is closely modeled on that of Schwab & Cotton (1983), which is a description of the mathematical framework behind the AIPS FRING task. In a preliminary step, for each baseline to a given reference antenna, the visibilities are Fourier-transformed in frequency and time and the peak of this transform is used to identify candidate parameter values in the two-dimensional space of delay and

¹¹ ACCOR corrects amplitudes in cross-correlation spectra due to errors in sampler thresholds using measurements of auto-correlation spectra.

delay rate. The height of the peak is used to calculate a signal-to-noise ratio, and antennas whose baseline to the reference antenna do not exceed a user-specified threshold are excluded from further steps and will be flagged when the resulting calibration table is applied to the data.

In a second (optional) step, these parameter estimates are refined using a non-linear least-squares solver that uses all baselines to all stations whose signal-to-noise ratio exceeded the threshold used in the FFT stage. While Schwab and Cotton implemented a specialized least-squares solver for AIPS (Schwab & Cotton 1983), the CASA implementation uses an implementation from the GNU Scientific Library.¹²

In keeping with CASA’s Measurement Equation based approach to calibration the parameters solved are stored in a calibration table from which Jones matrices can be calculated on demand. The fringe fit results are stored in a new type of *G*-Jones table, customised for the specific parameters that need to be stored; which are the phase (a.k.a. secular phase), the delay (phase slope as function of frequency), and the delay-rate (phase slope as function of time) and a dispersive delay term to characterize ionospheric delays proportional to the square of the wavelength (Small & Moellenbroeck 2022).

It is common practice in radio astronomy to perform a fringe fit on a bright source for a short interval with each spectral window separately; this can characterize and allow correction of any instrumental delays between the bands. Since the corrections calculated for this effect are typically applied to the whole data set, it is convenient to be able to zero the delay-rate term in the parameters, which would otherwise be extrapolated in time and dominate the correction. Yet, the delay-rate should always be included in the solve step to maximize the signal-to-noise value of the data. The CASA task **fringefit** has been outfitted with a dedicated parameter *zerorates* to do exactly this.

Additionally it is possible to control which of the delay, delay-rate and dispersive delay parameters are included in the fringe fit solution. The solution for phase cannot be switched off. By default, delay and delay-rate are solved for and dispersive delay is not. The motivating use-case for this option was to allow delay to be omitted when fitting spectral line data, but the functionality is quite general.

As is the case throughout CASA’s calibration framework, it is straightforward to use a source model (via the `MODEL_DATA` column of the `MeasurementSet`), and source models can be imported from images produced via self-calibration in various formats.

Not every feature available in AIPS has been implemented at the time of writing; notable omissions are that it is not yet possible to “stack” baselines, or combine correlations (LL and RR, for example) in a single fringe fit stage for added

sensitivity, but the **fringefit** task is under active development, and new features continue to be added.

2.1.2. Applying and Interpolating

CASA’s calibration framework includes generic methods to interpolate and apply solutions determined by the various calibration tasks to the data to be calibrated. The solutions determined by the **fringefit** task differ from the other calibration tasks in the sense that the solutions include the time-derivative of the phase (phase rate) as well as the phase itself. This time-derivative needs to be taken into account when interpolating (and extrapolating) fringe fit solutions in time in order to resolve phase ambiguities. This means the standard time interpolation implementation that just does a linear interpolation of the individual parameters of a solution is not sufficient. Therefore CASA’s interpolation framework was extended such that individual calibration classes that form the implementation of the various calibration tasks can override the time interpolation mechanism. A “rate aware” interpolation mechanism was added to the **fringefit** implementation. The interpolated phase is formed by extrapolating the nearest solutions in time in either direction using the time derivative and taking the average of these extrapolated solutions weighted according to their distance in time. This yields a smooth solution that correctly tracks the phase evolution in time as long as the phase rates of the different solutions are comparable. The delay is interpolated and applied independently.

2.2. New Task: *accor*

Not all VLBI correlators normalize the visibility amplitudes, which is a pre-requisite for proper amplitude calibration. A new CASA **accor** task has been written to do this normalization which divides the visibilities by the average of the auto-correlations on a timescale that can be specified by the user. But since the task uses the common CASA interfaces for data selection it not only provides the functionality of the AIPS ACCOR task, but also the functionality of the AIPS ACSCCL task¹³ which is needed to correct wide-band VLBA data.

2.3. Upgrading Other Tasks

The CASA **genca1** task already included code to use system temperature (T_{sys}) measurements to calibrate visibility amplitudes. This code made the assumption that such measurements are made at the same cadence for all antennas in the array. This is almost never the case for VLBI arrays so the code was changed to allow for measurements with a more irregular pattern by setting the parameter *uniform* set to True (the default value). In addition support was added for storing gain curves (which

¹² <https://www.gnu.org/software/gsl/>

¹³ ACSCCL is similar to ACCOR but uses only the inner part of the spectral bandpass to avoid a bias from bandpass falloffs, and should be performed after applying any bandpass corrections.

describe the elevation-dependency of the antenna gain) in the MeasurementSet and using these in the **genical** task to calibrate visibility amplitudes.

Many VLBI arrays use the standardized FITS Interferometry Data Interchange Convention (FITS-IDI; Flatters 1998; Greisen 2009). In preparation for future VLBI work, a CASA task **importfitsidi** had already been developed by the CASA team for CASA 3/4 in 2010–12 as part of the European RadioNet-funded ALBiUS project. This task converts FITS-IDI data into a MeasurementSet. The initial implementation did not include a number of optional (according to the format definition) metadata tables. Further upgrades were made to the task in later CASA versions, and the current version will correctly import system temperature measurements, gain curves and weather data from FITS-IDI (see also Section 3). Functionality to import the correlator model and pulse-cal measurements is still missing, but these measurements are not essential for the majority of VLBI observations. Another important addition is that the **importfitsidi** task will now apply the so-called digital corrections required for DiFX correlated data in the same way as the AIPS **FITLTD** task. These corrections are necessary to correct the effects of the coarse (2- or 4-level) signal quantization used by most VLBI data acquisition systems (generalization of the Van-Vleck correction as described in Appendix 8.3 of Thompson et al. 2016). These corrections are non-linear in the correlation amplitude and require information not available to CASAs generalized calibration framework, which is why the status quo of applying these in the “filler has been retained. The current implementation applies this correction for 2- and 4-level quantization, and combinations thereof.

2.4. Supporting Scripts

The traditional way to calibrate visibility amplitudes for connected-element radio interferometers is to observe a point-like source with known brightness and adjust the antenna gains to match the corresponding expected model visibilities. In CASA this is done using the **setjy** and **gaincal** tasks. This strategy does not work very well for VLBI since most bright calibrators are variable or have extended structure at the scales probed by VLBI. Instead visibility amplitudes are calibrated based on system temperature (T_{sys}) measurements at the individual antennas and the antenna gain (Moran & Dhawan 1995). Antenna gains are receiver-dependent and usually dependent on elevation as well. Ideally both T_{sys} and gain curves are distributed within the FITS-IDI data files. Unfortunately many VLBI arrays distribute this meta-data separately, usually in AIPS ANTAB format. To make this meta-data available to CASA a set of scripts were developed that add T_{sys} measurements and gain curve information to FITS-IDI files that lack them. These scripts are distributed separately from CASA.¹⁴

¹⁴ <https://github.com/jive-vlbi/casa-vlbi>

3. VLBI Applications

The CASA-VLBI functionality allows for calibration of the majority of VLBI observations. In this section we describe a few science cases for the EVN and VLBA. These instruments were chosen because of the differences in the processing steps caused by the differences in the correlators they use. The processing for continuum and spectral lines is discussed. Time domain and geodesy require a fundamentally different functionality (see also Section 5).

3.1. Instruments

The majority of observations with the EVN is correlated at the Joint Institute for VLBI ERIC (JIVE; the Netherlands) and the pre-processed data are made available through the EVN Data Archive in FITS-IDI format.¹⁵ The a-priori calibration information (gain and system temperature) as well as observational flags for all antennas can be applied by using the supporting scripts detailed in Section 2.4. The visibility data can be imported in CASA with the task **importfitsidi**. The full calibration and imaging can then be performed within CASA with consistent results with respect to the calibration done in AIPS. JIVE provides a Jupyter Notebook for continuum calibration of any EVN observation, and a full guide to calibrating EVN data within CASA is also available at the EVN website.¹⁶

Observations with the Very Long Baseline Array (VLBA) can also be calibrated with CASA. VLBA observations are correlated with the DiFX correlator (Deller et al. 2007) at the NRAO Pete V. Domenici Science Operations Center (Socorro, NM, USA) and the correlated data are available through the NRAO Archive Access Tool.¹⁷ As of CASA 5.8/6.2, the CASA task **importfitsidi** is able to import the gain curve and system temperature tables from the FITS-IDI file. The CASA **accor** task is capable of reproducing the behavior of the AIPS task **ACSCCL**, allowing for better amplitude calibration of wide-band VLBA data. A guide to calibrating VLBA data with CASA is available (Linford 2021), and a CASA Guide tutorial is in development. Astronomers at the US Naval Observatory (USNO) have developed a pipeline for calibrating and imaging VLBA observations with CASA and have demonstrated that the results compare favorably with calibration done in AIPS and imaging done in Difmap (Hunt et al. 2021).

The flexibility of the CASA framework and its convenient interface to the MeasurementSet data makes it easy for users to build high-performance data processing pipelines with a Python front end. In Janssen et al. (2019) *rPICARD* is presented, the first generic VLBI calibration and imaging

¹⁵ <http://archive.jive.nl/scripts/portal.php>.

¹⁶ <https://www.evlbi.org/evn-data-reduction-guide>.

¹⁷ <https://data.nrao.edu/>

pipeline that has been built fully on top of CASA and makes use of the new VLBI functionalities presented in this paper. With *rPICARD*, the first VLBI data set fully calibrated and imaged with CASA has been published (Janssen et al. 2019). The M 87 jet has been observed with the VLBA at 43 GHz and imaged with the **tclean** task and CASA self-calibration methods implemented in *rPICARD*. The pipeline is also used as a calibration pathway for the Event Horizon Telescope (e.g., Event Horizon Telescope Collaboration 2019b; Janssen et al. 2021; Event Horizon Telescope Collaboration 2022). With a significantly shorter track record, and continuous development, CASA performed as well as HOPS, the standard package for high-frequency VLBI calibration for the past decade (Whitney et al. 2004; Blackburn et al. 2019). More broadly, the calibration pipeline has been successfully applied to the EHT, GMVA, VLBA, EVN, and synthetic VLBI data from *MeqSilhouette* (Blecher et al. 2017; Roelofs et al. 2020; Natarajan et al. 2022). Input files for these use-cases and a detailed documentation are available in the online repository of the *rPICARD* pipeline.¹⁸

3.2. Main use Cases

The continuum data calibration in CASA follows the standard steps, which are described in the EVN User Manual and the VLBA User Manual (see also Section 3.1). A notable difference with AIPS is that CASA writes the calibration solutions to an external calibration table, which is a separate directory on disk, and not a table associated with the *UV data*. It is up to the user to specify which calibration tables to use when on-the-fly calibration is performed in tasks like **fringefit**. There is no cumulative calibration table in CASA, after all calibration steps are complete, the task **applycal** is used to apply all the calibration tables to the data.

Spectral line data in principle require the same calibration steps as continuum data, although they introduce a few additional complications. Spectral line observations require high spectral resolution, and the EVN correlator provides either one or two correlation passes. In the latter case, a separate low spectral-resolution pass is provided for calibrator sources. In the case of one correlator pass, the calibrator data needs to be manually averaged down to optimize signal-to-noise and processing time. Spectral line calibration requires an excellent bandpass correction, especially in the case of HI-absorption line studies, which is possible with the very flexible CASA **bandpass** task. The **mstransform** can be used to re-reference the data in frequency domain. Self-calibration of a target maser source as well as inverse phase referencing is usually performed using only a few narrow spectral channels. This functionality is also available in CASA **fringefit**.

Flagging can be done using the automated *rflag* option in the CASA **flagdata** task, by hand, or with any additional package. When flagging by hand, it is recommended to maintain a CASA flagmanager file to reproduce the flags at a later stage. The calibration tables can also be flagged, though it is better practice to flag the data that gives rise to poor calibration solutions, as this is generally bad data anyway. Particular attention has to be paid to the target flagging steps of spectral line observations, as narrow spectral lines (especially in masers) can be confused with man-made radio frequency interference (RFI) by automatic flagging algorithms. In this case manual flagging is recommended. After applying all the calibration tables to the data, the user can continue with imaging using CASA **tclean** or another package of their choice. We recommend to take care when changing the data format, and verify that all calibration and flags are properly applied in the new format before proceeding.

During the entire process of calibration, frequent quality control is necessary. The CASA task **plotms** can plot both visibility data and calibration solutions, also for the VLBI-specific tables. As the CASA task **plotcal** is no longer included in CASA 6, the EVN Jupyter notebook environment uses a custom-made replacement called **plotcalng**¹⁹ as the Matplotlib-based plots produced by **plotcalng** are better suited for inclusion in a Jupyter notebook. Note that **plotcalng** is not part of the CASA package.

4. Comparison

The EHT data calibration pipelines (Event Horizon Telescope Collaboration 2019b) have demonstrated extensively that CASA performs as well as HOPS (Whitney et al. 2004) and AIPS (Greisen 2003). The EHT data are difficult to calibrate compared to data from better understood instruments such as EVN and VLBA. A comparison with VLBA observations is done in Hunt et al. (2021), finding no differences in the results. We include here a basic comparison for EVN observations, which observes at lower frequencies than EHT, and is a more in-homogeneous array than the VLBA.

A comparison between AIPS and CASA calibration was performed using EVN observations from the Network Monitoring Experiment with project code N14C3. This observation has been extensively used in testing and verification, and all outcomes and issues are known and understood. The observation includes a phase-referencing setup which is a standard continuum observation technique for VLBI. The data include a fringe finder source, and two pairs of a target and phase reference source. All sources are bright, which makes this data set ideal for testing the frequently used phase-referencing mode of the EVN.

¹⁸ https://bitbucket.org/M_Janssen/picard

¹⁹ <https://github.com/aardk/evn-tools>

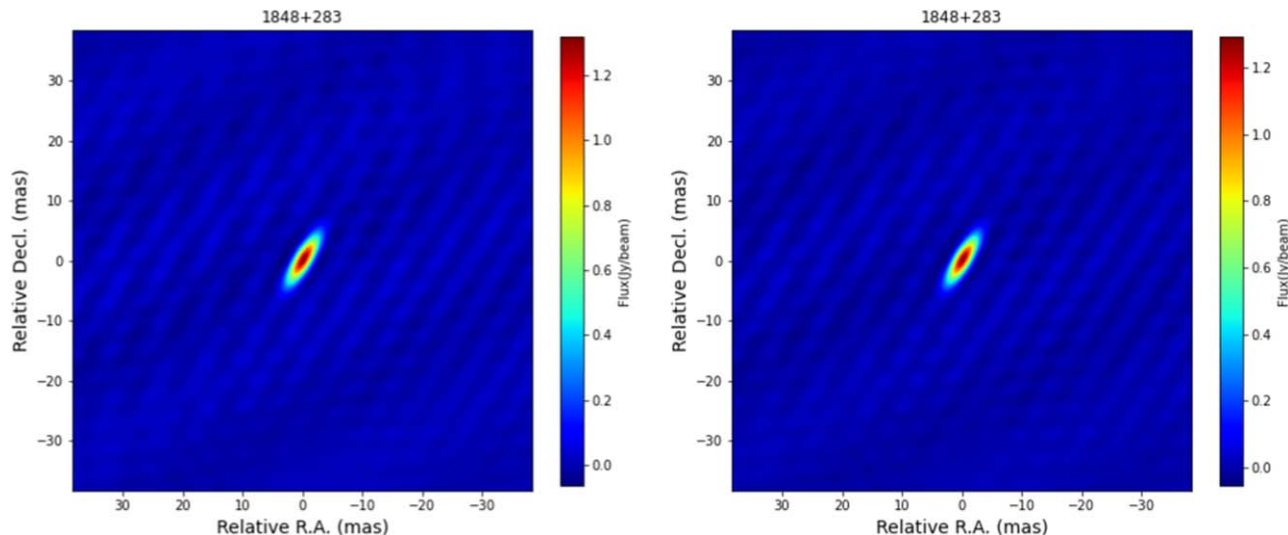


Figure 1. Total intensity images of calibrator 1848+283. Data calibrated in AIPS (left) and CASA (right). Both calibrated data sets were imaged with `tclean` using the same parameters.

The data were taken in 2014, in C band. The participating telescopes were Effelsberg, a single dish from the Westerbork array, Jodrell Bank Lovell telescope, Onsala, Toruń Svetloe, Zelenchukskaya, Badary, Sheshan, Hartebeesthoek and Yebes. The data were correlated in eight spectral windows of 32 channels each, and include all four polarization products. We used AIPS 31DEC22 (Greisen 2003), ParselTongue3 (Kettenis et al. 2006) and CASA 6.4.1 (McMullin et al. 2007; CASA Team et al. 2022). Two scripts were developed, one for AIPS and one for CASA and are publicly available at the JIVE code repository.²⁰ The scripts use Python 3.8. The CASA script is based on the EVN Jupyter notebook for continuum data processing,²¹ and the EVN Users Manual (see Section 3.1). The AIPS script also follows the online EVN Users Manual, but in places where the order of steps or the parameters deviate from CASA, we give priority to the CASA version and change the AIPS settings to reflect the CASA procedure as closely as possible.

Both scripts handle the visibility calibration, using the standard steps (in this order): system temperature and gain corrections, flagging of data based on the telescope logs and known bad stations, fringe fitting for the delay (also referred to as instrumental delay), fringe fitting for the delay-rate (also referred to as multi-band delay), and finally a complex bandpass correction.

Though we note that AIPS is generally faster than CASA, the processing speed of the two packages is difficult to compare due to differences in architecture, calibration model, and data format. These make CASA more user friendly, and more

flexible for larger and complex data sets, at the expense of processing speed and memory use. For example, converting the data from FITS-IDI to MeasurementSet format is very time consuming, but typically this is only done once in the calibration procedure. To give a rough idea of the difference, we compared the processing speed of the fringe fitting step, where the implementation in the two packages is fairly similar, and for the entire calibration process. On a personal laptop CASA on a single core is a factor ~ 15 – 20 slower in fringe fitting. For the full calibration procedure a similar difference is measured. When using the MPI option, this difference reduced to a factor ~ 2 , see Appendix F in Janssen et al. (2019).

4.1. Images

The final calibrated data are split into separate, calibrated data sets for each calibrator and target. From AIPS the calibrated data are exported to UV-FITS and then imported into CASA using `importuvfits`. Since we do not wish to compare the imaging tools, we choose to image all the data sets with CASA `tclean`. The imaging is done with a mild clean of 100 iterations, a cell size of 0.3 mas, image size of 512×512 pixels and a clean box of 20 pixels in RA and decl. centered on the central pixel. The other parameters are set to the task defaults. This ensures that any differences in the resulting images are solely due to the calibration process. Since producing science quality images of VLBI observations is done in many different ways depending on the science goal, experience, and taste of the user, we perform no self-calibration or further imaging. All sources were imaged. We found the signal-to-noise ratio was best for the calibrator 1848+283, and therefore use this source for further comparison. The reported

²⁰ https://code.jive.eu/bemmel/Comparison_AIPS_CASA

²¹ https://code.jive.eu/bemmel/EVN_CASA_pipeline

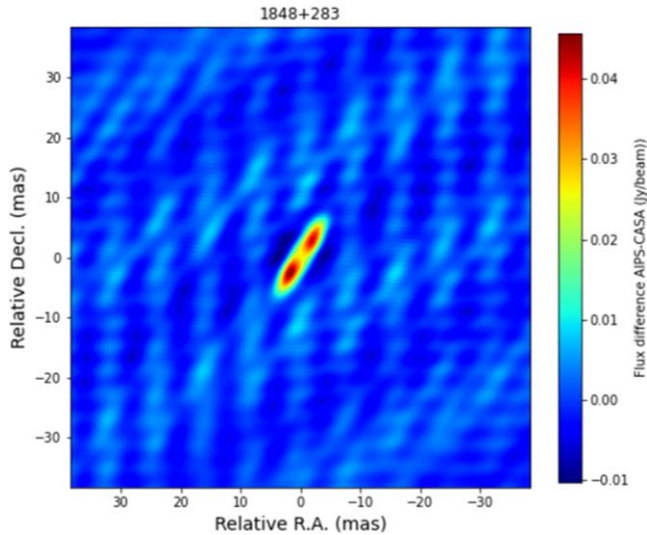


Figure 2. Difference image of the total intensity for 1848+382 (AIPS – CASA). Note the narrow range of the color scale.

trends are seen in all sources, but most obvious in the bright and compact calibrators.

In Figure 1 the resulting images are shown of the calibrator source 1848+283. For the AIPS calibrated data the peak brightness is $1.32 \pm 0.05 \text{ Jy beam}^{-1}$, for the CASA calibrated data the peak flux is $1.30 \pm 0.05 \text{ Jy beam}^{-1}$. The rms noise in the AIPS image is 0.054 Jy , versus 0.051 Jy in the CASA image.

Though the images look identical at first glance, a subtraction of the two reveals an extended residual (see Figure 2). The same pattern is visible in the background at a much lower level. The residual flux is well below the rms noise level of the original images. Though the majority of pixels shows a positive flux in this difference, implying that the AIPS flux calibration overall results in a higher flux, the distribution of the pixel values in the difference image is Gaussian and does not show a distinctive skew or offset (see Figure 3). This implies that the residuals originate from visibilities with similar noise realisation, but with a small (few percent) difference in the scaling.

4.2. Amplitudes and Phases

To compare the calibration quality without imaging the data, plots are generated of the calibrated amplitudes and phases as a function of UV -distance. Figure 4 shows the AIPS and CASA calibration results. Individual data points are plotted for each scan on 1848+283, both LL and RR polarizations, and for spectral windows number five and six (in CASA zero-based counting). For clarity not all spectral windows are included.

As for the total intensity images the differences between the AIPS and CASA calibration are not visible by eye, and a

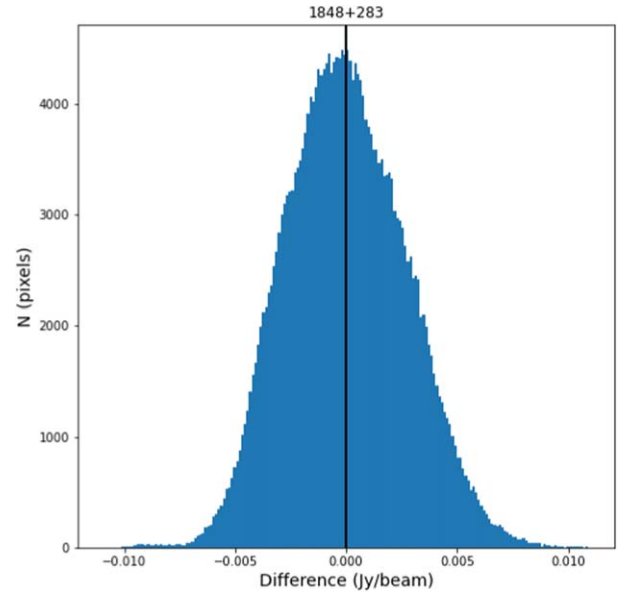


Figure 3. Histogram of the pixel flux in the AIPS–CASA difference image.

difference image was produced (see Figure 5). This demonstrates the same flux offset as seen in the total intensity images: the AIPS amplitudes are consistently higher than the CASA amplitudes by a few percent. This is somewhat larger than the differences found by Hunt et al. (2021), and we ascribe this to manual editing of the gain information in the archival data. Both packages should take their gain and system temperature information from the archival ANTAB file, however, for AIPS this resulted in errors in the gain calibration. This was overcome by using the calibration information in the AIPS CL2 table provided in the EVN archive. The generation of that particular CL2 calibration table cannot be reproduced with the current version of AIPS. The phases differences are all clustered around 0° phase, with at most a few degrees offset. The absolute amplitude and phase values are well within the empirically established accuracy limits for self calibration of 10% amplitude errors and 10° phase errors. The baselines with Westerbork show a split in the phase differences, which has been traced back to the instrumental delay calibration in AIPS. The odd and even spectral windows have respectively a positive and negative phase offset after the instrumental delay has been applied. These correspond to upper- and lower-sideband data, and indicate a small phase difference between those two, which is not properly corrected in AIPS. The offset is no more than 2° in either direction, a value much smaller than typical residual source errors. The fact that this offset is spotted in the phase differences indicates that the quality of the calibration in both packages is comparable and accurate to within less than a degree.

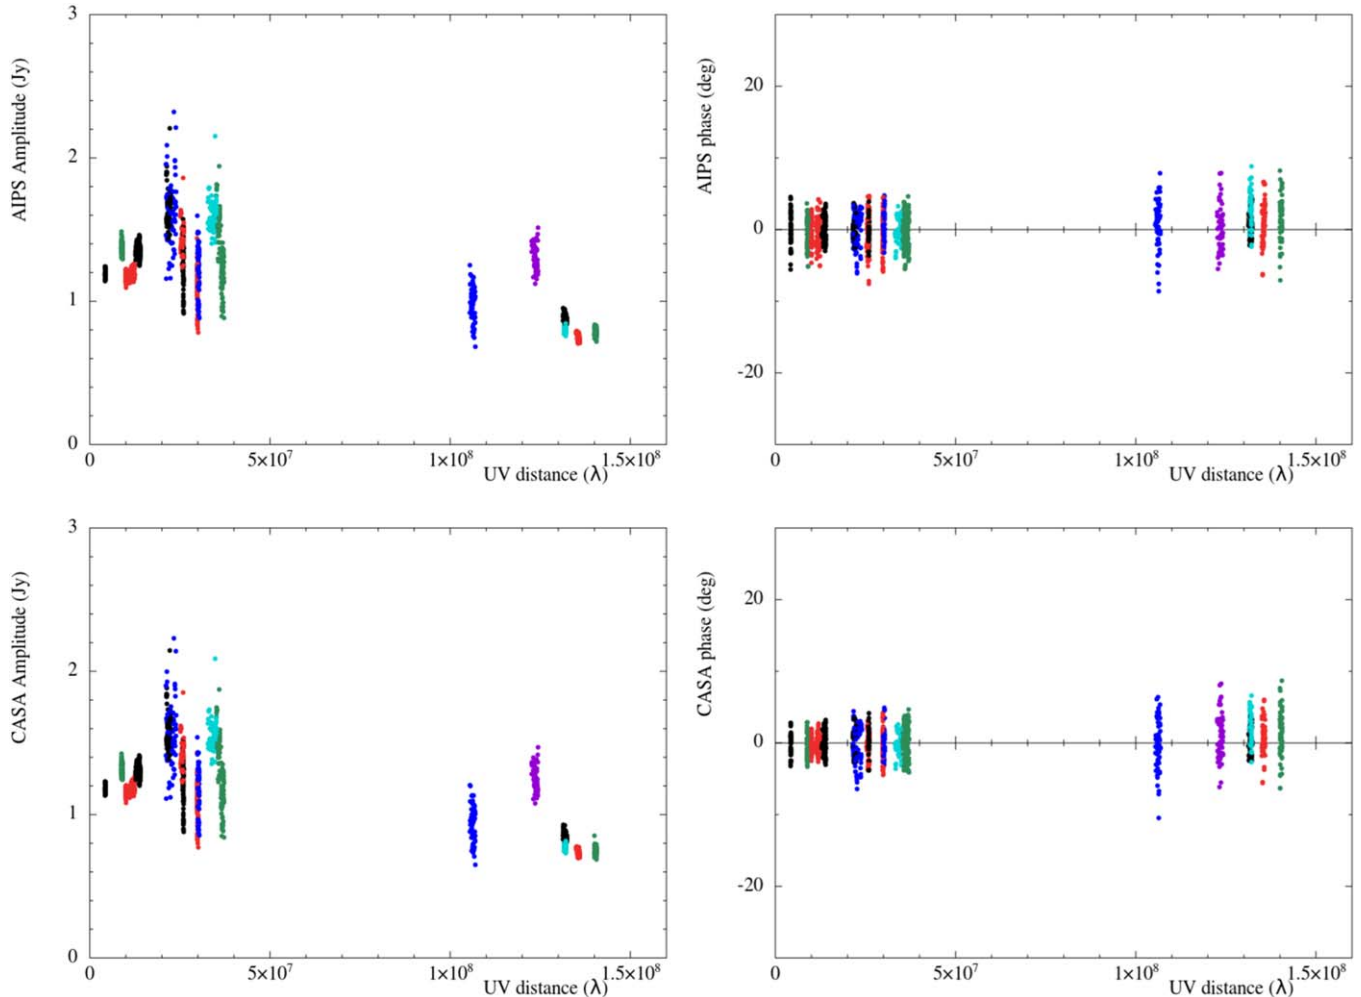


Figure 4. Plots of calibrated amplitudes (left) and phases (right) as a function of UV distance for the calibrator 1848+283. The data include LL and RR polarization, and spectral windows 5 and 6. Top row is AIPS calibration, bottom row is CASA calibration. The colors indicate the first station in the baseline. Black: Effelsberg; red: Westerbork; dark green: Onsala; blue: Noto; cyan: Torun; purple: Hartebeesthoek.

5. Future Plans

While the functionality implemented in CASA 6.4 should be sufficient to achieve the scientific goals of many types of VLBI observation, some more advanced functionality is still missing. Work continues to implement this. Here we review several upcoming changes and improvements. Note that this list is not complete, and as instruments advance, the software development will follow.

1. *Pulse cal tones.* This may help taking out instrumental delays by making use of calibration signals inserted at the receivers and measured at the VLBI correlator.
2. *Fringe fit improvements.* Several improvements are planned such as baseline stacking, stacking of polarization products, and multi-band fringe fitting of irregularly spaced spectral windows.
3. *Polarization calibration.* Full-stokes imaging requires calibration of the polarization leakage (the so-called D-terms). For polarization observations with VLBI instruments the calibration algorithm needs to be able to handle resolved polarization calibrator sources. This is something that CASA currently does not provide. A possible candidate software package is PolSolve (Martí-Vidal et al. 2021) since it fully exploits the Measurement Equation.
4. *Ionospheric corrections.* While CASA implements ionospheric corrections in the **gencal** task, preliminary evaluation of this task suggests that improvements can be made for VLBI.
5. *Delay model accountability and EOP corrections.* For astrometric VLBI observations it is desirable to be able to make corrections to the delay model used by the

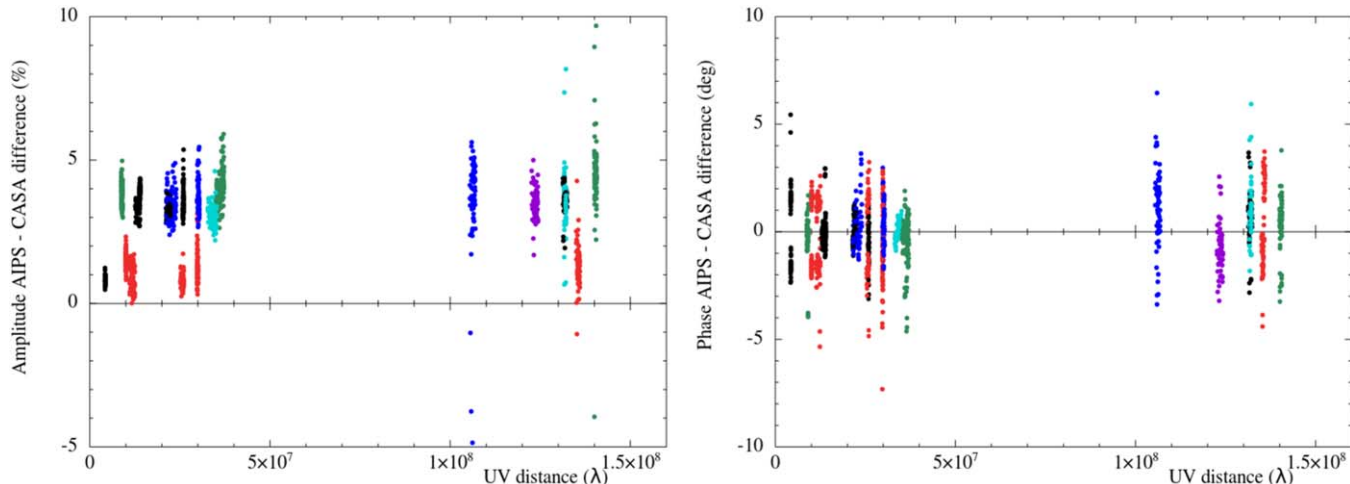


Figure 5. Difference of the calibrated amplitudes (left) and phases (right) for the calibrator 1848+283, using CASA – AIPS as a function of UV distance. The data include LL and RR polarization, and spectral windows 5 and 6. The colors indicate the first station in the baseline. Black: Effelsberg; red: Westerbork; dark green: Onsala; blue: Noto; cyan: Torun; purple: Hartebeesthoek.

correlator. The most important of these are corrections to the Earth Orientation Parameters (EOPs) since the most accurate measurements of these parameters are only available some time after observations.

Our focus has been primarily on processing VLBI radio astronomy observations. But there are opportunities for using CASA for calibration of geodetic observations as well. The **fringefit** task is well-positioned to handle the increased bandwidth available in the VGOS system (Petrauchenko et al. 2012). And plans to include resolved sources in geodetic observations will require the capability to fringe fit based on a source model, as well as the imaging capabilities to determine a source model in the first place. CASA already provides that functionality, unlike the HOPS package that is currently used to process most geodetic observations. How to reconcile the antenna-based solutions from a global fringe fitter with the baseline-based solutions required by geodetic analysis software is an open question though. Note that the CASA fringe fitter has a mode to skip the least-squares globalization step and provide just the baseline-based solutions in a calibration table, but it does not currently convert this into the information needed for geodetic experiments.

While CASA is undeniably based on more modern technologies than AIPS, its data access layers provided by *casacore* are limiting its scalability. This already is an issue for larger observations with current instruments but will be a problem for future instruments such as ngVLA. At NRAO, the CASA Next Generation Infrastructure (CNGI) project studied new technologies for replacing *casacore* and the existing MPI-based parallelization framework. The

conclusion from this study is that Python-based Xarray²² and Dask²³ technologies are good choices to develop a Next Generation CASA (ngCASA). The current plan is to introduce these technologies step by step into CASA. While scalability is less of an issue for typical VLBI observations, more demanding VLBI observations (such as wide-field observations) can benefit from these new technologies as well. Therefore we expect to adopt these technologies in future developments of VLBI-specific CASA tasks and ultimately reimplement already existing tasks (such as **fringefit**) on top of these new technologies (see CASA Team et al. 2022).

6. Summary

The implementation of VLBI functionality in CASA is an important step toward sustainable software for VLBI data processing. It also opens the route toward a Findable, Accessible, Interoperable and Reusable (FAIR) software and data policy through Jupyter notebooks²⁴ that can be linked to the science data from a specific experiment, improving scientific reproducibility as well. This paper presents the current status of the VLBI functionality of the CASA package, based on CASA 6.4.1. New tasks have been added to handle the fringe fitting and specific amplitude calibration steps for VLBI use cases. Existing tasks have been updated to handle the calibration meta-data properly, and underlying processing has been adjusted to work in the extreme limits of VLBI observations.

²² <https://docs.xarray.dev/en/stable/>

²³ <https://docs.dask.org/en/stable/>

²⁴ https://code.jive.eu/bemmel/EVN_CASA_pipeline

The CASA package is now capable of handling the majority of VLBI science cases, and has already undergone intense verification as part of the EHT data processing. The Python3 base allows for easy development of automated pipelines, and there is a Jupyter notebook kernel including CASA, which is extremely suitable for training purposes and less experienced users. We have presented an outline of the workflow for continuum and spectral line observations in CASA. Detailed recipes are available for the EVN and VLBA on their websites, and other telescopes are starting to provide similar resources.

A detailed comparison with the AIPS package has demonstrated that CASA performs equally well, though slower on a single CPU core. This can be overcome by using the MPI infrastructure, which enables multi-threading and can speed up the processing by orders of magnitude, depending on the underlying hardware.

As the technology of VLBI continues to evolve, the software is also under constant development. We have listed several improvements that are planned for the coming year. Longer term plans are developed jointly between JIVE and NRAO, and involve feedback from users and other stakeholders.

The CASA VLBI functionality is not only suitable for radio astronomy applications, but can potentially also serve geodetic experiments. In the more distant future, the need for handling large data will require a significant overhaul of the underlying CASA infrastructure, and plans for this are already under development. We expect that the functionality presented in this paper will remain available.

We would like to thank Sophia Vaughan for her contribution in testing specific functionality of the fringe-fit task for spectral line applications during her Summer Studentship at JIVE in 2017. We thank the anonymous referee for a prompt and helpful assessment of the paper. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc. The European VLBI Network is a joint facility of independent European, African, Asian, and North American radio astronomy institutes. Scientific results from data presented in this publication are derived from the following EVN project code(s): N14C3. This work is supported by the ERC Synergy Grant “BlackHoleCam: Imaging the Event Horizon of Black Holes (grant 610058). This work is also supported by ESCAPE—The European Science Cluster of Astronomy & Particle Physics ESFRI Research Infrastructures, which has received funding from the European Unions Horizon 2020 research and innovation program under the grant Agreement 824064. B.M. acknowledges support from the Spanish Ministerio de Economía y Competitividad (MINECO) under grants AYA2016-76012-C3-1-P and MDM-2014-0369 of ICCUB (Unidad de Excelencia “María de Maeztu”).

Appendix Software Comparison

A detailed comparison of several software packages was done in late 2014 as part of the requirements analysis for a pipeline that could process high-frequency VLBI observations from the GMVA and EHT. The software packages assessed were AIPS, Miriad²⁵ (Sault et al. 1995), HOPS,²⁶ (Whitney et al. 2004; Blackburn et al. 2019), CASA (McMullin et al. 2007; CASA Team et al. 2022), the LOFAR software suite²⁷ and PIMA²⁸ (Petrov et al. 2011). The packages were compared on their abilities at that time to serve as a fully capable VLBI data processing package, with focus on the necessary steps to process correlated VLBI data into scientific results. To compare the packages a list of five aspects was made: reliability, flexibility, future prospects, user access and pipeline readiness. Each aspect included several requirements, which were weighted highest for critical requirements, to lowest for optional requirements. For each requirement the packages received scores between 1 (lowest) and 5 (highest). With this information a weighted average was calculated for each aspect. This resulted in an overall score with standard deviation for each package, see Table A1. The standard deviation is indicative of how much spread there is in the grades, and therefore the need for significant adjustments to meet some of the requirements. A lower standard deviation implies that overall there is less work to be done to meet individual requirements, while a high standard deviation means that though some requirements are met, significant discrepancies exist for others, which implies more work. The highest score was for CASA (4.5 ± 0.8), with AIPS a good runner-up (3.9 ± 1.2). The other packages scored well below 3, only the LOFAR toolkit scored 3.2, but required significant adjustments

Table A1

Weighted Mean Scores and Standard Deviation for all Packages Assessed for the Software Comparison. The Scoring was Done on a Range of Requirements, With 0 Being the Lowest and 5 the Highest Score. Weighting was Applied Based on the Necessity of Each Requirement

| Package | Mean | Standard Dev. |
|---------|------|---------------|
| AIPS | 3.9 | 1.2 |
| Miriad | 2.3 | 1.0 |
| HOPS | 2.3 | 1.1 |
| CASA | 4.5 | 0.8 |
| LOFAR | 3.2 | 1.3 |
| PIMA | 2.3 | 1.6 |

²⁵ <https://www.atnf.csiro.au/computing/software/miriad/>

²⁶ <https://www.haystack.mit.edu/haystack-observatory-postprocessing-system-hops/>

²⁷ https://www.astron.nl/lofarwiki/doku.php?id=public:user_software:start

²⁸ http://astrogeo.org/pima/pima_user_guide.html

given the very large difference in frequency, and it was not a singular software package but rather a collection of tools and packages. In further assessment with the development teams at NRAO it became clear that CASA would take priority for their future work. Combined with the better score, this led to the choice to use CASA as the basis for further development.

ORCID iDs

Ilse M. van Bemmell  <https://orcid.org/0000-0001-5473-2950>
 Mark Kettenis  <https://orcid.org/0000-0002-6156-5617>
 Des Small  <https://orcid.org/0000-0003-3723-5404>
 Michael Janssen  <https://orcid.org/0000-0001-8685-6544>
 George A. Moellenbrock  <https://orcid.org/0000-0002-3296-8134>
 Dirk Petry  <https://orcid.org/0000-0002-8704-7690>
 Ciriaco Goddi  <https://orcid.org/0000-0002-2542-7743>
 Justin D. Linford  <https://orcid.org/0000-0002-3873-5497>
 Kazi L. J. Rygl  <https://orcid.org/0000-0003-4146-9043>
 Elisabetta Liuzzo  <https://orcid.org/0000-0003-0995-5201>
 Benito Marcote  <https://orcid.org/0000-0001-9814-2354>
 Olga S. Bayandina  <https://orcid.org/0000-0003-4116-4426>
 Neal Schweighart  <https://orcid.org/0000-0002-3287-1535>
 Marjolein Verkouter  <https://orcid.org/0000-0003-2884-9834>
 Aard Keimpema  <https://orcid.org/0000-0002-5575-2774>
 Arpad Szomoru  <https://orcid.org/0000-0001-8525-4605>
 Huib Jan van Langevelde  <https://orcid.org/0000-0002-0230-5946>

References

- Blackburn, L., Crew, G. B., Fish, V. L., et al. 2019, *ApJ*, **882**, 23
 Blecher, T., Deane, R., Bernardi, G., & Smirnov, O. 2017, *MNRAS*, **464**, 143
 Deller, A. T., Tingay, S. J., Bailes, M., & West, C. 2007, *PASP*, **119**, 318
 Event Horizon Telescope Collaboration 2019a, *ApJL*, **875**, L2
 Event Horizon Telescope Collaboration 2019b, *ApJL*, **875**, L3
 Event Horizon Telescope Collaboration 2022, *ApJL*, **930**, L13
 Flatters, C. 1998, AIPS Memo, 102
 Goddi, C., Falcke, H., Kramer, M., et al. 2017, *IJMPD*, **26**, 1730001
 Greisen, E. W. 2003, Information Handling in Astronomy—Historical Vistas, Vol. 285 (Berlin: Springer), 109
 Greisen, E. W. 2009, AIPS Memo, 114
 Hunt, L. R., Johnson, M. C., Cigan, P. J., Gordon, D., & Spitzak, J. 2021, *AJ*, **162**, 121
 Keimpema, A., Kettenis, M., Small, D., Dijkema, T. J., & Szomoru, A. 2020, in ASP Conf. Ser., 527 (San Francisco, CA: ASP), 579
 Kettenis, M., van Langevelde, H. J., Reynolds, C., & Cotton, B. 2006, in ASP Conf. Ser., 351 (San Francisco, CA: ASP), 497
 Janssen, M., Goddi, C., van Bemmell, I. M., et al. 2019, *A&A*, **626**, A75
 Janssen, M., Falcke, H., Kadler, M., et al. 2021, *NatAs*, **5**, 1017
 Linford, J. 2021, VLBA Scientific Memo #38
 Martí-Vidal, I., Mus, A., Janssen, M., de Vicente, P., & González, J. 2021, *A&A*, **646**, A52
 McMullin, J. P., Waters, B., Schiebel, D., et al. 2007, in ASP Conf. Ser., 376 (San Francisco, CA: ASP), 127
 Moran, J. M., & Dhawan, V. 1995, in ASP Conf. Ser., 82 (San Francisco, CA: ASP), 161
 Natarajan, I., Deane, R., Martí-Vidal, I., et al. 2022, *MNRAS*, **512**, 490
 Petrachenko, W. T., Niell, A. E., Corey, B. E., et al. 2012, Int. Assoc. Geodesy Symp., 136, 999
 Petrov, L., Kovalev, Y. Y., Fomalont, E. B., & Gordon, D. 2011, *AJ*, **142**, 35
 Roelofs, F., Janssen, M., Natarajan, I., et al. 2020, *A&A*, **636**, A5
 Small, D., & Moellenbrock, G. 2022, CASA Memo, 12
 Sault, R. J., Hamaker, J. P., & Bregman, J. D. 1996, *A&AS*, **117**, 149
 Sault, R. J., Teuben, P. J., & Teuben, M. C. H. 1995, in ASP Conf. Ser., 77 (San Francisco, CA: ASP), 433
 Schwab, F. R., & Cotton, W. D. 1983, *AJ*, **88**, 688
 The CASA Team et al. 2022, *PASP*, **134**, 114501
 Thompson, R., Moran, J., & Swenson, G. 2016, Interferometry and Synthesis in Radio Astronomy (3rd edn.; Berlin: Springer)
 Whitney, A. R., Cappallo, R., Aldrich, W., et al. 2004, *RaSc*, **39**, RS1007