

EVLA Memo No. 208

JVLA calibration stability at L-band over 5.5 years

Christopher A. Hales^{1,2} and Nathan Stephenson²

¹NRAO, Socorro, USA

²Newcastle University, Newcastle upon Tyne, UK

June 28, 2019

Abstract

We present the first systematic investigation of stability in JVLA calibration parameters over timescales of weeks, months, and years. We focus on a limited parameter space at L-band using an ideal science dataset, the recently completed CCP¹ survey, in which the same set of calibrators were observed many times over several years using the same array configuration and correlator setup. We examine electronic delay, bandpass, and polarization leakage calibration solutions, taking care to account for changes in reference antenna. Overall we find very stable solutions over timescales of months and, in the case of leakages, years. Exceptions are discussed, for example changes arising from electronic reset events or receiver changes. This memo is not intended to advocate for the use of calibration database solutions (e.g. in some subset of observing modes, perhaps default continuum modes), but merely to provide evidence to support discussions along these lines regarding both JVLA operations and design decisions for the proposed ngVLA.

¹<http://www.chilesconpol.com/>

Contents

1	Introduction	3
2	Data	3
3	Results and Discussion	7
3.1	Electronic Delay	7
3.1.1	Crosshand Delay	8
3.2	Bandpass	8
3.2.1	Crosshand Bandpass Phase	9
3.3	Leakage	10
3.3.1	3C286 Fractional Polarization	11

List of Figures

1	Antenna elevation when observing 3C286	13
2	Solar elevation when observing 3C286	13
3	Slant total electron content when observing 3C286	14
4	Station K indices for Boulder CO	14
5	Reference antenna per observation	15
6	Receiver changes	15
7	Electronics reset events	16
8	Parallel hand delay	17
9	Crosshand delay	18
10	Bandpass amplitude	19
11	Crosshand bandpass phase for ea24	20
12	Crosshand bandpass phase for ea26	21
13	Crosshand bandpass phase for ea21	22
14	Leakage amplitude	23
15	3C286 fractional polarization spectrum	24

1 Introduction

Scheduling blocks for the Karl G. Jansky Very Large Array (JVLA) are currently required to contain observations of all necessary calibrators together with the science target. An alternative approach that has been considered since the planning stages of the VLA upgrade is to store calibration solutions in a database and apply them when practical. In cases where the atmosphere does not contribute significantly to the calibration parameter of interest, or where the atmosphere can be accounted for separately, it is generally expected that the telescope hardware should be sufficiently stable to support the re-use of calibration solutions obtained days, months, or even years earlier (modulo changes following maintenance). This would minimize the number of scheduling blocks in which certain calibrator observations are required, and in turn improve overall science observing efficiency. However, to date, no systematic study has been performed to assess this capability (though see Murray et al. 2013, Perley & Sault 2014, and Schinzel 2018).

This memo presents the first such investigation, focusing on a limited parameter space using an ideal science dataset in which the same set of calibrators and target were observed many times over several years using the same array configuration and correlator setup. The data examined in this memo are obtained from the recently completed CHILES Continuum Polarization (CCP) survey (Hales et al., in prep.). The focus in this memo is on stability of calibration parameters associated with typical JVLA data reduction procedures, such as bandpasses and polarization leakages, over timescales of months and years.

Section 2 presents an overview of the CCP dataset and highlights several issues to be aware of when interpreting results. The results for several calibration parameters are presented in respective subsections in Section 3. Each subsection presents the metric used for assessing stability in the calibration parameter of interest, caveats for interpreting the resulting plots, and an examination of trends.

2 Data

CCP is a full-polarization continuum radio survey at 1.4 GHz that has captured an ultra-deep exposure over 0.2 square degrees within the COSMOS field, yielding the most sensitive survey of extragalactic radio sources ever performed with the VLA. Observations were performed concurrently with the COSMOS HI Large Extragalactic Survey (CHILES; HI spectral line, total intensity only) over a total observing time of 1 khr, the largest ever awarded to a PI project.

The key science aims of CCP are to improve our understanding of galaxy evolution by opening up radio detectability of typical star forming galaxies seen in the optical, X-ray, and infrared bands; probe large scale magnetic fields throughout the cosmic web; provide high quality data of legacy value to the astronomical community; and, through collaboration with a partner team CHILES Verdes, probe faint radio transient and variable sources.

A total of 210 observations were performed over 5 consecutive B configurations. The combined observing time was 1014 hours. A summary of CCP observing semesters and dates is presented in Table 1. The last 8 observations in semester 3 were performed during the initial

Semester	Start Date	End Date	# of Epochs	Total Observing Time (h)
1	2013-10-25	2014-01-21	45	182.3
2	2015-02-25	2015-05-04	39	215.5
3	2016-05-19	2016-09-27	49	188.0
4	2017-11-11	2018-01-29	48	236.7
5	2019-03-01	2019-04-11	29	191.6

Table 1 Summary of CCP observations.

Spw ID	Ch0 Freq. (MHz)	# of Chan's	Spw BW (MHz)	Baseband Pair
0	1000	64	128	B0D0
1	1384	64	128	B0D0
2	1640	64	128	B0D0
3	1768	64	128	B0D0

Table 2 Summary of CCP correlator setup.

stages of move time from B to A configuration. Observation lengths varied from 8 hours to as little as 1 hour; the typical length was 5 hours. All observations used the same correlator setup in which 4 spectral windows were observed in all four correlation products (RR , RL , LR , LL) from a common 8-bit baseband pair at L-band. It was not possible to access more spectral windows without exceeding the number of available WIDAR baseline board pairs shared between CHILES and CCP. The spectral setup is summarized in Table 2. Spectral window frequencies were selected to avoid significant RFI while also providing wide coverage in wavelength-squared space (necessary for optimized polarization analysis of the survey data) and, where possible, avoiding significant sensitivity loss at the band edges (Momjian 2017). CCP observations between 2016-08-12 and 2016-09-27 inclusive (19 observations) were affected by the atmospheric delay problem², though this does not appear to have had any significant impact on the quality of our data.

Every observation included the same calibrators: 3C286 for calibrating delay, bandpass, absolute flux density, crosshand delay, and crosshand bandpass phase (position angle of linear polarization); and the unpolarized ($< 0.05\%$) source J0943–0819 for calibrating complex gain and polarization leakage. The minimum angular separation between the sun and 3C286 during any observation was 37° ; for J0943–0819 it was 21° . The minimum angular separation between the moon and 3C286 was 33° ; for J0943–0819 it was 16° .

The data from each observation were calibrated in identical fashion using a custom-developed scripted pipeline operated in CASA release version 4.7.2. This was run in a semi-automated manner by a single person, subjecting the results from every calibration step to manual inspection and providing high quality assurance. High quality flagging was performed using `pieflag` (Hales & Middelberg 2014). CASA functionality in the pipeline was enhanced by the incorporation of the custom tasks `antintflag`³ (Hales 2016), `plot3d` (Hales 2016), and `interpgain` (Hales 2016). An additional assessment of calibration quality for each

²<https://science.nrao.edu/facilities/vla/data-processing/vla-atmospheric-delay-problem>

³This functionality is now available in the task `gaincal` in more recent releases of CASA.

observation was performed by applying the calibration solutions to 3C286 and measuring the spectrum of fractional linear polarization, and comparing this with the known spectrum from D configuration observations (Perley & Butler 2013). Measurements from the switched power system were not used in the calibration process. The pipeline will be explained in more detail elsewhere (Hales et al., in prep.).

The pipeline was run twice for each observation, the first time to obtain calibrated data and a quicklook image within a few hours of coming off the telescope (for transient searches and potential optical followup), and the second at least 2 weeks after the observation so as to incorporate GPS-derived ionospheric total electron content (TEC) data obtained from the International Global Navigation Satellite System Service (IGS), which has a latency period of up to 2 weeks. As a result, there are two complete CCP datasets: the first calibrated without accounting for ionospheric electron content data, and the second taking this into account (including ionospheric Faraday rotation). In general, only results from the second run (including TEC corrections) will be shown in this memo unless otherwise specified, because there are no differences other than occasionally some slightly improved flagging.

The following plots are provided to assist interpretation of the results shown in Section 3.

Fig. 1 displays antenna elevation when observing 3C286.

Fig. 2 displays solar elevation at the time when observing 3C286.

Fig. 3 displays slant TEC (i.e. vertical TEC from IGS corrected for zenith angle) at the time when observing 3C286. As expected, the TEC data do not show close correspondence with solar elevation (i.e. with time of day) because the VLA is at sufficient magnetic latitude to avoid significant diurnal fluctuations in TEC (e.g. Hales 2019). The TEC data do appear to correlate with Solar Cycle 24, in which sunspot activity peaked in early 2014 and dropped to quiet conditions by 2019.

Fig. 4 displays Station K indices measured at Boulder CO at the times when 3C286 was observed, obtained from the Space Weather Prediction Center⁴. Station K indices are used to characterize the magnitude of geomagnetic storms, tied to specific geographic locations⁵, and in turn to identify periods of local ionospheric instability during which the TEC data obtained from the IGS may poorly capture the true state of the ionosphere. Planetary K indices can range from 0–9 with 1 being calm and 5 or more indicating a geomagnetic storm, though note that Station K indices represent raw data for a given geomagnetic observatory and are not normalized to the Planetary scale. The Station K indices in Fig. 4 do not indicate the presence of any significant geomagnetic storms during CCP observations. As expected, the indices are elevated during daytime observations and are suppressed at night. Note also that the K indices observed in semester 5 are typically smaller than those in semester 4, the former observed just after sunset when the ionosphere was recombining, while the latter were observed during the quieter hours before sunrise. Curiously, there does not appear to be any correlation with Solar Cycle 24, perhaps reflecting the low duty cycle of significant flaring activity⁶.

⁴<https://www.swpc.noaa.gov/products/station-k-and-indices>

⁵An ensemble of Station K indices from locations worldwide are used to calculate the Planetary K index, or K_p . There are no significant differences between the K_p and Boulder Station K indices during CCP observations.

⁶For example, examine X-ray data from the *GOES*-series satellites, or the frequency of transient alerts for solar flares and coronal mass ejections from the CACTus quicklook catalog derived from the LASCO instrument on board the *Solar and Heliospheric Observatory* satellite.

Fig. 5 displays the reference antenna selected for each observation, on which the phase was fixed to zero in both polarizations during the calibration process. Antennas ea21, ea24, and ea26 were frequently referenced. As identified in Section 3, where possible (and relevant), the calibration solutions under examination have been modified to fix antenna ea25 as the reference antenna for all 210 observations. Fortuitously, this antenna (and only this antenna) was available and remained unflagged in all 210 observations.

Fig. 6 displays receiver changes. These were detected by tracking changes in T_{cal} values, accessed from the measurement set data from each observation. These data were (crudely) verified against historical weekly reports of receiver serial numbers and their locations in the array. Note that a small number of T_{cal} revisions could reflect database updates (e.g. following correction of internally swapped polarization cables; this particular issue was largely resolved by mid 2016) rather than true receiver changes.

Fig. 7 displays times when internal delays were adjusted (e.g. following digital transmission system resynchronization events) for IFs B and D corresponding to polarizations R and L , respectively. These data were obtained from the EVLA Parameters Database. Delays calibrated from the observational data may sometimes, but not always, be different following these electronics reset events.

Finally, some general comments to assist interpretation of the results presented below.

- The time axis is shown compressed in all figures. Observations are displayed sequentially. Each observation is assigned the same plotting width, regardless of observation length. Time gaps between observations (typically one to a few days) are not shown.
- Dashed vertical lines indicate breaks of approximately 13 months between observing semesters. Start and end dates for each semester are displayed.
- Missing antenna data, whether arising from flagging (online, RFI, egregious calibration solutions) or simply because the antenna was unavailable (maintenance, out of service), are indicated by white.
- There is inherent low-level noise (fluctuations) in many of the results shown below, arising from the various normalizations used to derive statistics. For example, each bandpass amplitude spectrum per spectral window, polarization, and observation was initially normalized to unit median over the central 42 channels (avoiding 1/3 of channels at the spectral window edges) so that changes in bandpass shape relative to other observations could be easily detected. However, if different channels were flagged between observations, then their normalizations could differ, even for otherwise identical bandpass spectra. In turn, the resulting comparisons would then indicate a (spurious) change in bandpass shape. Therefore, in general, attention should only be paid to trends in the plots that are significant and coherent over multiple observations. Small variations between observations should be interpreted as upper limits to the true underlying variance.

3 Results and Discussion

Three types of calibration solutions are examined in the subsections below: (i) electronic delay, (ii) bandpass, and (iii) polarization leakage. Notation follows the CASA measurement equation formalism⁷, for example demonstrated by Hales (2017).

3.1 Electronic Delay

To enable results to be easily compared over all observations, the delay solutions were recalculated per observation setting ea25 as the new reference antenna. Modified parallel-hand delays K' were calculated from the original delay K per observation following

$$K'_{ijp} = K_{ijp} - K_{iAp} \quad (1)$$

for each i observation, j antenna, and p polarization (R and L for VLA L-band), where A is the index of the new reference antenna (ea25).

No differences were found between the TEC-corrected and uncorrected data, as expected. The results from the TEC-corrected dataset are displayed in Fig. 8.

The data indicate that delays are typically stable to within less than 1 nanosecond over monthly timescales. Abrupt changes are seen between observing semesters, and sometimes during observing semesters (e.g. ea18 in semester 2). There are several potential origins for these changes. Those between semesters arise from delay resets that take place whenever the array configuration changes (Fig. 7). Another factor is receiver changes (Fig. 6); for example, delays are systematically affected over the array in semester 5, following the only receiver change on the reference antenna ea25 over the 5.5 year timespan. Receiver changes are always accompanied by an internal delay reset. Internal delay adjustments (Fig. 7) are not always correlated with the delays displayed in Fig. 8, as expected, because the aim of an internal adjustment is to keep the system as stable as possible, compensating for internal step changes as they occur and ensuring that celestially-calibrated delays will fall within a 16 nanosecond window. However, many of the step changes in Fig. 8 (or curious gradients, e.g. ea17 in semester 1) are not clearly correlated with receiver changes or internal delay adjustments. These likely arise from subtleties in the electronic architecture, requiring a deeper investigation that is beyond the needs of this memo. Section 3.1.1 will examine a subtlety regarding the internal reference antenna, with implications for some (but again, not all) of the features seen in Fig. 8.

In principle, if perturbing events can be tracked, and if the overall system is otherwise sufficiently stable, then it should be possible to maintain stability in delays, store delays in a database and update them as required, and supply the database delays to science scheduling blocks. However, this may be more work than it is worth; delays can be easily and quickly (and perhaps preferably) measured from celestial (secondary) calibrators that will always be present in any scheduling block.

⁷<https://casa.nrao.edu/casadocs/latest/casa-fundamentals/the-measurement-equation-calibration>

3.1.1 Crosshand Delay

Unlike for parallel hand delays, CASA cannot (yet) solve for multi-band crosshand delay solutions (i.e. combining all spectral windows). As a result, the crosshand delay on the reference antenna was calibrated per spectral window, per observation. However, the results are practically identical between spectral windows, so only the results from spectral window 1 will be examined here (see Table 2). To enable results to be easily compared over all observations, the crosshand delay solution was recalculated per observation setting ea25 as the new reference antenna. Modified crosshand delays K'_x were calculated from the original crosshand delay K_x per observation following

$$K'_{x,is} = K_{x,is} + (K_{iAR} - K_{iAL}) \quad (2)$$

for each i observation in spectral window $s = 1$, where A is the index of the new reference antenna (ea25) and K_{ijp} follows the notation for parallel hand delays from the previous section.

No differences were found between the TEC-corrected and uncorrected data, as expected. The results from the TEC-corrected dataset are displayed in Fig. 9.

The data indicate that the crosshand delay on ea25 is typically very stable to well within 1 nanosecond over monthly timescales, modulo some abrupt step changes. These changes are not correlated with receiver swaps (Fig. 6) or with delay reset events on ea25 (Fig. 7). However, as mentioned in the previous section, step changes in delay can arise from more subtle aspects of the electronic architecture. One such subtlety, which explains many of the observed step changes, arises because internal crosshand delays are tied to the X-band receiver located on the internal reference antenna. For any given observation, this internal reference antenna is identified as one of the three possible antennas located on pads N8, E8, or W8. Fig. 9 highlights observations performed after a delay reset event was recorded on any of the antennas located on these pads. Many of the step changes are correlated. A deeper investigation to identify the origins of all step changes is beyond the scope of this memo.

As with parallel-hand delays, it should be possible to maintain crosshand delay solutions in a database. However, a quick observation of an unresolved linearly polarized celestial calibrator (not necessarily with known position angle) in each scheduling block can replace the overheads associated with implementing such a scheme.

3.2 Bandpass

Only bandpass amplitude stability is investigated here, and specifically bandpass amplitude shape. Bandpass phases are less revealing because they are typically close to zero following parallel-hand delay calibration, with deviations that are correlated with parallel-hand delays.

For each individual spectral window per polarization, antenna, and observation, the bandpass amplitude spectrum was normalized to unit median over the central 42 channels (avoiding 1/3 of channels at the spectral window edges). Then, taking all amplitudes for a given channel, polarization, spectral window, and antenna, the median amplitude over all observations was subtracted and the result was normalized by the median. At this stage of processing, the data for each antenna reflects the spectrum of fractional amplitude deviations per observation and polarization. Finally, to compress this data for each antenna into a single datum per

observation and polarization, the mean deviation over all channels in all spectral windows was calculated.

No differences were found between the TEC-corrected and uncorrected data, as expected. The results from the TEC-corrected dataset are displayed in Fig. 10.

The data indicate that bandpass amplitudes are stable to within 1% over weeks and months. This is likely an upper limit to any true variations, due to the artificial noise contribution from processed highlighted at the end of Section 2. Similarly, the correlated deviations over all antennas seen in some observations are likely caused by RFI, where a substantial number of channels were flagged resulting in a systematic normalization offset. No correlation is found between receiver swaps and changes in bandpass amplitude shapes, consistent with the findings from Morris & Momjian (2012) where characteristic bandpass shape was traced to the downconverter (T304).

These data support the concept of storing bandpass solutions in a calibration database at low frequencies where the atmosphere does not contribute significantly. At higher frequencies, if the atmospheric contribution to the bandpass can be estimated, then this could be used to adjust the database bandpass solutions as needed. To implement such a scheme, it will be necessary to carefully account for various filter responses (baseband, subband) so that the standardized database solutions can be propagated to the nominated observing mode; good bookkeeping will be essential.

3.2.1 Crosshand Bandpass Phase

Crosshand bandpass phases are specific to the reference antenna. It is not possible to recalculate these for a different reference antenna without recalibrating the data. While possible, this has not been performed here. Rather, stability in crosshand bandpass phases are investigated for the 3 antennas that were most often set as the reference antenna: ea24 in semesters 1, 3, and 4; ea26 in semester 2; and ea21 in semester 5.

For each spectral window per observation, the mean phase (accounting for any wraps) over the central 42 channels (avoiding 1/3 of channels at the spectral window edges) was calculated. In the circular feed basis (as is the case here for VLA L-band), crosshand bandpass phase calibration is synonymous with absolute position angle calibration, and is therefore affected by ionospheric Faraday rotation (e.g. Hales 2017). The mean phase per spectral window is therefore expected to be rotated accordingly when comparing between results calculated from the TEC-corrected and TEC-uncorrected data.

Ionospheric Faraday rotation will cause the position angle of linear polarization to be rotated by

$$\Delta\theta_f \approx 1.35 \left(\frac{\nu}{1 \text{ GHz}} \right)^{-2} \left(\frac{B_{los}}{1 \text{ G}} \right) \left(\frac{N_e}{1 \text{ TECU}} \right) \text{ deg} , \quad (3)$$

where ν is the observing frequency in units of GHz, N_e is the slant TEC in units of TECU, and B_{los} is the line of sight magnetic field in units of gauss which is positive when pointed toward the observer. B_{los} can be suitably estimated for a characteristic ionosphere approximated by a thin shell at a weighted altitude of ~ 400 km by taking the Earth's $B_{los}(l)$ at the height of the thin shell (Hales 2019). The Earth's line of sight magnetic field at the VLA site ranges

between 0 G to +0.4 G with typical value +0.2 G over the sky at the IGS-assumed 450km height of the thin-shell ionosphere (Hales 2019). Following Equation 3, and assuming a typical slant TEC ≈ 20 TECU from Fig. 3, the expected rotation at 1 GHz is approximately 5° , while at 2 GHz it is approximately 1° .

Fig. 11 displays both the TEC-corrected results and the difference between corrected and uncorrected data for reference antenna ea24. Similarly, Fig. 12 displays results for ea26, and Fig. 13 displays results for ea21. Some of the outliers in the lower panels arise from flagging differences between the initial and final pipeline calibrations, where the latter contains improved solutions.

The results displayed in the upper panels reveal that crosshand bandpass phases are stable to within a few degrees over monthly timescales. This is likely an upper limit to any true variations, due to an artificial noise contribution from data processing similar to that described earlier for bandpass shapes. The data reveal abrupt step changes that are correlated with those seen in Fig. 9 for crosshand delays. The lower panels reveal frequency-dependent rotations that are consistent with the expectations from ionospheric Faraday rotation discussed above. These data also verify that the various calculations necessary to account for ionospheric Faraday rotation in CASA are performing as intended.

These data support the concept of storing crosshand bandpass phase solutions in a calibration database, as long as Faraday rotation corrections are applied to the database values as required. A complication will be the need to account for step changes as described in Section 3.1.1 for crosshand delays.

3.3 Leakage

Notation for complex polarization leakages in this work follows Hales (2017) where d_{jR} indicates the fraction of orthogonal polarization L sensed by feed R on antenna j (and vice versa for d_{jL}). The CCP data were calibrated assuming an unpolarized calibrator, yielding relative leakages in which the real and imaginary leakage components on the R feed of the reference antenna were set to zero. To enable results to be easily compared over all observations, the leakage solutions were recalculated per observation setting ea25 as the new reference antenna. Modified leakages d' were calculated from the original leakages d per observation following

$$d'_{ijRsc} = d_{ijRsc} + \beta \quad (4)$$

$$d'_{ijLsc} = d_{ijLsc} - \beta^* \quad , \quad \text{where} \quad (5)$$

$$\beta = -d_{irRsc} \quad (6)$$

for each i observation, j antenna, s spectral window, and c channel, where r is the index of the new reference antenna (ea25).

Only leakage amplitude stability is investigated here. While it is possible to perform a similar analysis for leakage phases, this is more challenging because of correlations with the parallel-hand and crosshand delays examined earlier, and complications associated with fluctuations in the presence of phase wraps (not insurmountable, but not compelling if amplitudes are available to indicate instrumental instability).

For each individual channel per spectral window, polarization, and antenna, the median leakage amplitude over all observations was subtracted. At this stage of processing, the data for each antenna reflects the spectrum of leakage amplitude deviations per observation and polarization relative to a median spectrum for that polarization. To compress this data for each antenna into a single datum per observation and polarization, the mean deviation over all channels in all spectral windows was calculated.

No differences were found between the TEC-corrected and uncorrected data, as expected. The results from the TEC-corrected dataset are displayed in Fig. 14.

The data indicate that leakage amplitudes are very stable to within fractions of a percent over weeks, months, and even years (over multiple semesters). Some step changes are correlated with receiver swaps (Fig. 6), for example ea17 in semester 1. Examination of receiver serial numbers indicates that T_{cal} values for ea01 in semester 3 were updated several observations after the receiver was replaced; the true receiver swap date is correlated with the step change seen in leakage amplitudes. Some step changes are more difficult to diagnose, for example the $\sim 1\%$ change in semester 3 for d_L (but not d_R) on ea13.

The data support the concept of storing polarization leakage solutions in a calibration database, requiring only periodic re-evaluation on minimum \sim monthly timescales modulo qualifying events such as receiver swaps or (less likely) changes in optical alignment.

3.3.1 3C286 Fractional Polarization

To provide rapid assessment of the quality of full polarization calibration per observation, the CCP pipeline applies full calibration to 3C286 and measures the spectrum of fractional linear polarization. Recall that leakages were calibrated using J0943–0819.

For each individual channel per spectral window, the median fractional polarization over all observations was subtracted. At this stage of processing, the data reflects the spectrum of fractional polarization deviations per observation relative to the median spectrum. To compress this data into a single datum per observation, the mean deviation over all channels in all spectral windows was calculated.

No differences were found between the TEC-corrected and uncorrected data, as expected. The results from the TEC-corrected dataset are displayed in Fig. 15. For comparison, the fractional polarization spectrum from D configuration observations reported by Perley & Butler (2013) is also displayed (consistently) as a deviation from the median CCP spectrum. Note that the data plotted from Perley & Butler (2013) does not include time evolution (this is examined below).

The data indicate that the spectrum of fractional polarization is typically very stable, with scatter at any point in time typically within $\pm 0.02\%$. The observed levels of fractional polarization are all systematically higher than the value expected from Perley & Butler (2013), likely arising from higher angular resolution and thus decreased depolarization when observing 3C286 in B configuration. Similarly, the data during move time reveal systematically higher levels of fractional polarization. The scatter in fractional polarization levels seen during this period likely reflects the different hour angles (uv coverage) for these observations (e.g. see Fig. 1). The data reveal a slow systematic rise over time. This does not appear to arise

from changes in J0943–0819, though this is approaching the accuracy with which this can be examined (comparing leakages calibrated in semesters 1 and 5 using parallactic angle coverage). Instead, the fractional polarization for 3C286 appears to be increasing with time at a rate of approximately 0.1% over 5.5 years, or 1.8% per century. This is roughly consistent with the rise of 0.7% per century inferred from observations over a 30 year baseline reported by Perley & Butler (2013). The data presented here appear to exhibit temporal structure, rising between semesters 1 and 3, falling slightly in semester 4, and continuing to rise again in semester 5. It is plausible that the gradual rise of 0.7% over decades is punctuated by slightly variable behaviour on intermediate yearly timescales, particularly for higher angular resolution observations.

Fig. 15 provides confirmation that the CCP calibration solutions are of high quality.

Acknowledgements: We thank Vivek Dhawan, Wes Grammer, Emmanuel Momjian, and Ken Sowinski for insightful discussions and feedback. CAH acknowledges support from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 705332. NS acknowledges support provided by an undergraduate summer studentship from the School of Mathematics, Statistics & Physics at Newcastle University.

References

- Hales C. A., Middelberg E., 2014, Astrophysics Source Code Library, ascl:1408.014
- Hales C. A., 2016, antintflag, v1.0, doi:10.5281/zenodo.163546, as developed on GitHub
- Hales C. A., 2016, interpgain, v1.1, doi:10.5281/zenodo.163002, as developed on GitHub
- Hales C. A., 2016, plot3d, v1.5, doi:10.5281/zenodo.162986, as developed on GitHub
- Hales C. A., 2017, AJ, 154, 54
- Hales C. A., 2019, ngVLA Memo 62 (in prep.)
- Momjian E., 2017, EVLA Memo 204
- K. Morris, E. Momjian, 2012, EVLA Memo 154
- Murray C., Goss W. M., Stanimirovic S., 2013, EVLA Memo 171
- Perley R. A., Butler B. J., 2013, ApJS, 206, 16
- Perley R., Sault B., 2014, EVLA Memo 178
- Schinzl F. K., 2018, EVLA Memo 205




Figure 1 Antenna elevation when observing 3C286.




Figure 2 Solar elevation when observing 3C286.




Figure 3 Slant total electron content when observing 3C286. 1 TECU = $10^{16} \text{ e}^-/\text{m}^2$.




Figure 4 Station K indices for Boulder CO.




Figure 5 Reference antenna used for calibration.




Figure 6 Observations where antenna T_{cal} values differ from the previous observation. In most cases this indicates that an L-band receiver was replaced. It is possible that a small number of T_{cal} revisions reflect database corrections rather than receiver changes.




Figure 7 Observations performed after internal delay reset events for IFs B (top) and D (bottom), corresponding to polarizations R and L respectively.





Figure 8 Stability of parallel hand delays for polarization R (top) and L (bottom). The reference antenna for all observations is ea25. See Section 3.1 for details.




Figure 9 Stability of crosshand delays for spectral window 1 (others are practically identical) on the reference antenna ea25. The vertical orange lines indicate observations performed after internal delays were adjusted for any of the antennas located on pads N8, E8, or W8. See Section 3.1.1 for details.





Figure 10 Stability of bandpass amplitude shapes for polarization R (top) and L (bottom). See Section 3.2 for details.

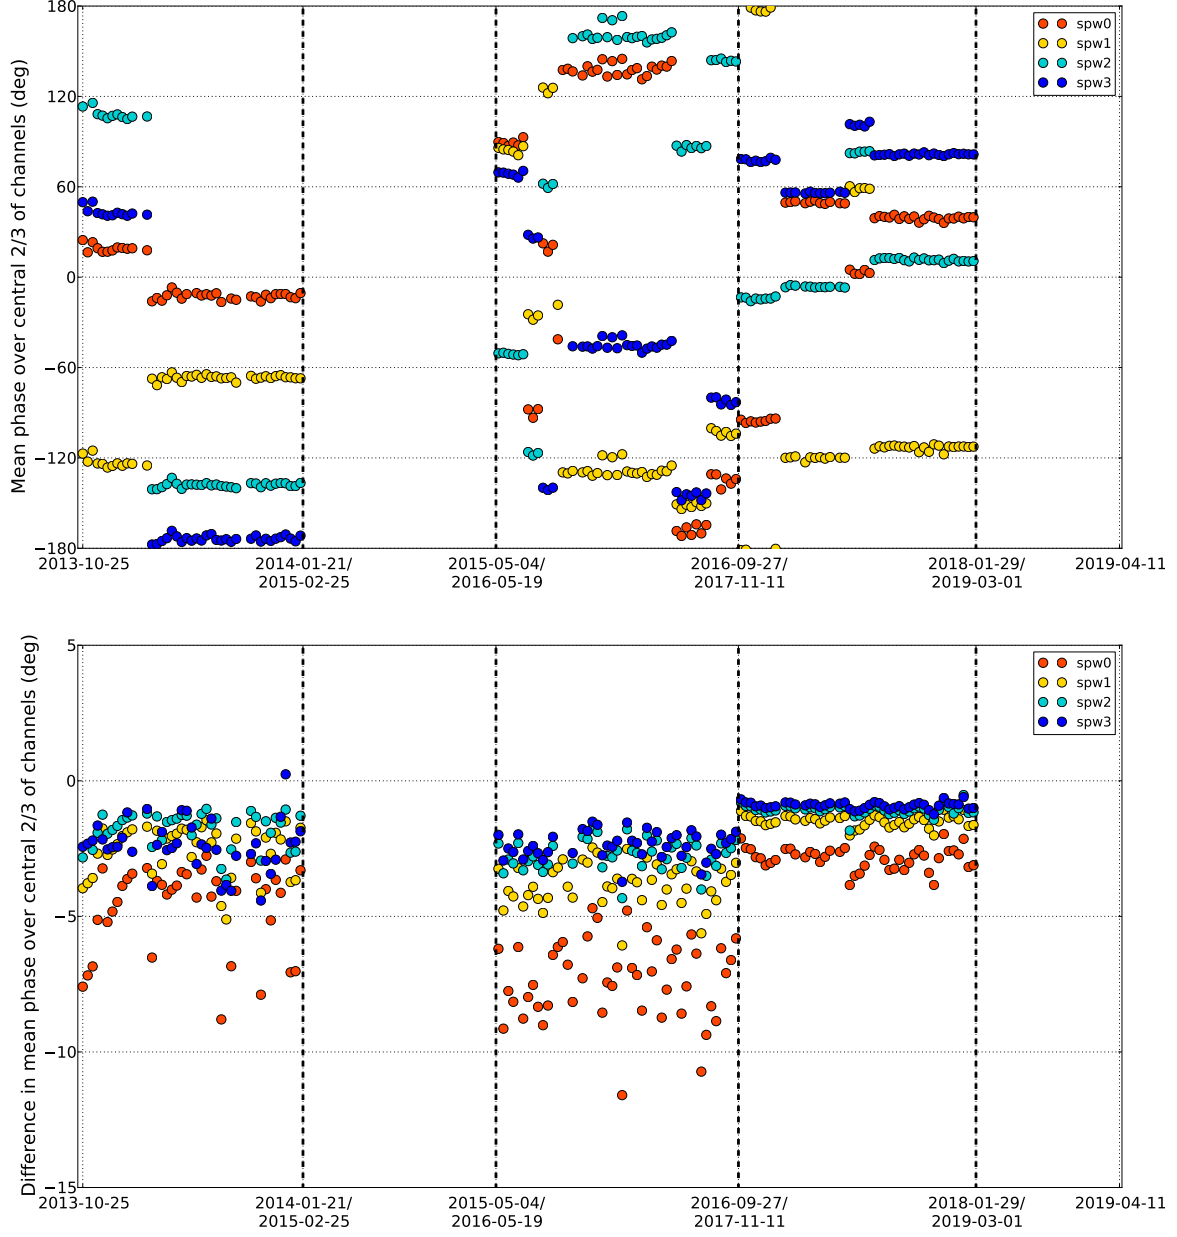


Figure 11 Top: Stability of crosshand bandpass phases calibrated with TEC corrections on reference antenna ea24. Bottom: Difference between crosshand bandpass phases with–without TEC corrections on reference antenna ea24. See Section 3.2.1 for details.




Figure 12 Top: Stability of crosshand bandpass phases calibrated with TEC corrections on reference antenna ea26. Bottom: Difference between crosshand bandpass phases with–without TEC corrections on reference antenna ea26. See Section 3.2.1 for details.




Figure 13 Top: Stability of crosshand bandpass phases calibrated with TEC corrections on reference antenna ea21. Bottom: Difference between crosshand bandpass phases with–without TEC corrections on reference antenna ea21. See Section 3.2.1 for details.





Figure 14 Stability of leakage amplitudes for d_R (top) and d_L (bottom). The reference antenna for all observations is ea25. See Section 3.3 for details.




Figure 15 Stability of 3C286 fractional polarization spectrum. The scale is in units of percentage fractional linear polarization. The orange dashed line indicates the D configuration fractional polarization spectrum from (Perley & Butler 2013); this line does not account for time evolution. Note that the last 8 observations in Semester 3 were during initial stages of move time from B to A configuration. See Section 3.3.1 for details.