NGVLA 110

IMPLICATIONS FOR THE NGVLA FROM VLBA PHASE STABILITY AT 43 GHz

R. Craig Walker (NRAO)

Apr 25, 2023

ABSTRACT

Two full-track observations of M87, made using the VLBA at 43 GHz during winter, have been analyzed for RMS phase fluctuations on timescales from 2 to 60 seconds. The intent is to try to understand with real, nearly-noiseless data on strong sources, how well phase referencing with fast switching and paired antenna calibration will work to maintain coherence on the ngVLA at a more diverse collection of sites than just the VLA. These time series data do not actually sample the phase offsets either between close sources or close antennas at the VLBA stations. But reasonable inferences can be made by assuming that the phase variations are caused by a fixed screen blowing over the site at a nominal wind speed aloft and that the fluctuations in that screen are statistically similar between possible antenna locations and along the path that passes over a single antenna. The most important conclusions are (a) that phase referencing, even at 86 GHz, should work well at most sites and work some of the time at all sites, (b) the antennas at poor sites should be placed as close together as possible for the best paired antenna calibration, even at the cost of some shadowing, (c) for coherence, 4 antennas is best but 3 antennas per station may be adequate to maintain coherence, although the possibility of rapid gradient changes still needs to be addressed, and (d) for referencing between frequency bands, fast switching may work as well or better than paired antennas if the cycle can be kept very short. This memo, and the data upon which it is based, do not address phase offsets between sources that are not reflected in the variations on one source. Such gradients will need to be calibrated for the best coherence and certainly for astrometry. The time scales required for such calibration will impact the choice of number of antennas per station. They will be explored in a future memo.

Introduction

Phase calibration using calibration sources, commonly called phase referencing in the VLBI context where it is used to allow detection of weak sources and to enable accurate relative astrometry, will be arguably the most important observing technique used on the ngVLA. The effectiveness at the VLA site is well known from long history with the VLA for 43 GHz and lower frequencies. The

effectiveness at 86 GHz has been addressed in earlier NGVLA Memos including the first by Carilli (2015; hereafter C15). The stations of the ngVLA LONG, the VLBI segment of the ngVLA, present two factors that are not well covered in these previous studies. Some of the sites are not as good as the VLA. In fact some chosen for their UV coverage may be downright poor at high frequencies. Also the ngVLA LONG sites will have multiple antennas which enable the use of calibration techniques not possible for the rest of the array.

The default plan for ngVLA LONG is to have 3 antennas per station, driven originally by sensitivity. But a calibration method known as multiview (Rioja and Dodson, 2020), based on measuring atmospheric gradients, should allow the best referencing and may work best with 4 antennas per station. In either case of numbers of antennas, it is expected that some variant on paired antenna calibration, where calibrators and targets are observed simultaneously on different antennas, will work better than the fast switching that must be used on the single antenna stations of other portions of the array. With 4 antennas, three antennas on calibrators can measure the gradient while the fourth tracks the target. With 3, there will need to be some form of switching or other scheme to determine the full parameters of the gradient. Multiple possible observing modes are possible with 3 antennas.

It would be useful to test these modes with real data. But a full test requires finding at least two existing instruments with multiple, close-packed antennas and the ability to record the VLBI data streams from multiple antennas. It will not be trivial to arrange such an observation.

An alternate way to address the issue is to look at data from single-antenna stations as a function of time. If we assume that the phase fluctuations are the result of fixed atmospheric structures blowing over the antenna (Taylor, 1938 via Thompson, Moran, and Swenson, 2001; TMS) and assume a nominal wind speed aloft to turn the time variations into spatial variations, we can get a handle on the relative phase fluctuations that might be expected between closely spaced antennas. Making an additional assumption about an atmospheric scale height will in turn allow an estimate of the phase fluctuations between close sources. The analysis of VLA data discussed elsewhere such as C15 is similar, or even better because the the phases between antennas could actually be measured by the interferometer. But those studies only cover a single, good site.

The Program and Data

For this project, I developed a quick-and-dirty program (RMSLISTR) that parses a phase listing from the AIPS task LISTR. For each data point, it looks at the phase difference between that point and the data point a number of seconds later. It does that for all points for which the later point exists (which it doesn't near the end of a scan) and accumulates the numbers needed to calculate an RMS. That process is repeated for time differences from 2 seconds (one record) to 60 seconds. LISTR is run in a mode where the phases are given in tenths of a degree. In addition to the RMS described, for each pair of points, the program determines the RMS deviation of the phase of the data point half way between them after subtracting the average phase of the original pair. That mimics what is normally done in phase referencing where the original pair represent adjacent calibrator scans while the half way point represents the target scan in between. If the gradients were strictly linear, that interpolated RMS would be near zero. As it is, it is a measure of the non-linearity which turns out to be significant.

The data files used for this were processed through the usual VLBI data reduction sequence in AIPS to just before any short term self calibrations, including fringe fitting, were done. All of the calibration steps that were done either did not affect the phases , applied smoothly varying phases derived from

external information, or were based on a single self calibration with constant results applied to the whole data set. The steps were the ionospheric calibration based on GPS data (TECOR), correction for updated Earth Orientation Parameters (EOP), parallactic angle calibration, manual pulse cal (removal of delay and phase offsets for each IF for one calibrator scan using a single interval fringe fit), atmospheric corrections based on geodetic segments (DELZN), and amplitude calibration based on gains and tsys (APCAL). A bandpass calibration was also done based on a single calibrator scan. As a result of these calibrations, the channels, IFs and polarizations were aligned in phase and could be averaged. The residual phases at this point were not showing rapid phase winding. The multi-source, multi-frequency data were then averaged across channels and IFs to form new data sets (SPLIT), one per source, with a single data point per 2 second record. That file was read by LISTR which listed the phases for stokes 'I'. Because of the 132 column limit for LISTR output, not all baselines could be listed. I determined two of the most stable stations based on looking at all of the data, and listed the phases on all 17 baselines involving either of those two. With some more hoop jumping – three LISTR passes and output files – all baselines could be listed but that did not seem necessary.

The program RMSLISTR generates the data for a number of tables. The ones that will be shown in this memo are the RMS phase vs time offset for offsets from 2 to 60 seconds and the RMS phase of the interpolated midpoints as a function of the half cycle time (one end to the mid point). The RMS values are relative to an average but deviate by insignificant amounts from the RMS relative to zero. These RMS values are of the baseline phases. The individual antenna calibration phase fluctuations will be lower as the two antennas contribute in quadrature. For the best baselines, the antenna phases will be lower by about $\sqrt{2}$ while the noisy baselines, like those involving SC, will be dominated almost entirely by the high-noise antenna. Additional tables available from RMSLISTR include the coherence loss that would be implied by the RMS phase and the fractional improvement in RMS provided by the interpolation (typically 0.7 and very constant with time on each baseline). There are also a couple of tables examining the changes in phase slope with time, but those turned out to be basically redundant with the other tables. A subtlety used in the program is that, to make resolving ambiguities more reliable, the phase difference between two points is calculated by stepping through all the intervening points and adding up the differences. Whenever adjacent points were more than 180° apart, it was assumed that 360° should be added or subtracted to bring the step back under 180. In practice, there are essentially no real steps more than 180° in 2 seconds in these data so this scheme works well.

The two data sets were ones for which I did a full imaging analysis in the past and have the calibrated data handy on disk. They both used the new-at-the-time wideband system on the VLBA. The observing frequency is 43 GHz and all 10 VLBA stations produced data in both cases. After editing of edge channels, the total bandwidth was 192 MHz per polarization and, as noted, LISTR displayed stokes 'I' so the effective bandwidth was twice that. The data sets were taken on 12 Jan 2013 (BW098A) and 14 Mar 2016 (BW115D) so this is winter data. Most of the data were taken at night, but some extended into the morning. These data have been published in Walker et al. 2018. For this study, two sources have been used. One is 3C279 which was used as a calibrator and observed for a few scans of about 2 minutes each for a total of about 13 minutes in each project. It has correlated flux densities between 4 and 20 Jy, so so the thermal noise contribution to the phase is very small. The phase scatter is dominated by the atmosphere. The other source is 3C274 (M87) which was the main target and which was observed for a total of about 5.4 hours in each project. It has a correlated flux density between about 0.3 and 1.5 Jy. With the full bandwidth, the phase fluctuations for this source are also dominated by the atmosphere. The ngVLA will have about 100 times the bandwidth used here, so these results are comparable to what could be obtained on sources 10 times weaker with the new instrument.

Table 1: The RMS phase difference (deg) between points separated by "Time sep" seconds vs baseline for 3C274 in BW098A. These are the RMS baseline phase relative to an average so are the quadrature sum of the individual antenna phase RMS values. They differ from the RMS relative to zero by less than 0.3% for the vast majority of times and baselines, with no differences over 0.5% at times under 30 seconds and none over 1.4% at 60 seconds.

Time sep	BR-OV	FD-OV	HN-OV	KP-OV	LA-OV	MK-OV	NL-OV	OV-PT	OV-SC	BR-PT	FD-PT	HN-PT	KP-PT	LA-PT	MK-PT	NL-PT	PT-SC
2	4.9	9.0	15.4	4.3	4.6	5.8	11.3	4.5	19.4	5.3	8.4	10.3	3.4	3.7	5.7	10.8	17.9
4	6.6	14.5	17.4	6.0	6.5	7.0	18.1	6.4	29.8	6.5	14.2	12.6	5.1	5.5	7.0	17.6	28.1
6	8.1	19.1	19.5	7.6	8.1	8.5	23.2	8.0	39.7	7.7	18.9	15.2	6.5	7.0	8.4	22.8	37.7
8	9.4	23.2	21.6	9.0	9.4	9.9	27.5	9.4	48.9	8.9	23.1	17.7	7.8	8.3	9.9	27.1	46.4
10	10.6	26.8	23.7	10.2	10.7	11.4	31.2	10.7	57.4	9.9	26.8	20.0	9.0	9.5	11.3	30.8	54.4
12	11.8	30.1	26.1	11.3	11.8	12.7	34.6	11.8	65.4	10.9	30.1	22.3	10.1	10.7	12.7	34.2	61.6
14	12.8	33.3	28.0	12.4	12.9	14.2	37.5	12.9	72.8	11.8	33.2	24.4	11.2	11.7	14.2	37.2	68.4
16	13.8	36.2	30.2	13.4	13.9	15.5	40.0	13.8	79.8	12.7	36.2	26.6	12.2	12.8	15.6	39.7	74.8
18	14.7	39.0	32.1	14.4	14.7	16.9	42.2	14.7	86.5	13.5	38.9	28.5	13.2	13.7	17.0	41.9	80.9
20	15.6	41.7	33.9	15.3	15.6	18.2	44.3	15.6	92.6	14.2	41.6	30.4	14.1	14.6	18.3	44.0	86.9
22	16.4	44.3	35.8	16.2	16.4	19.6	46.3	16.6	98.5	15.0	44.2	32.3	15.0	15.5	19.7	45.9	92.6
24	17.3	46.8	37.6	17.1	17.2	21.0	48.2	17.5	104.1	15.7	46.7	34.1	15.9	16.3	21.1	47.8	98.2
26	18.1	49.3	39.5	18.0	18.0	22.4	50.0	18.4	109.3	16.4	49.2	35.9	16.8	17.0	22.4	49.6	103.5
28	18.9	51.7	41.2	18.9	18.8	23.7	51.8	19.2	114.3	17.1	51.7	37.6	17.7	17.8	23.8	51.3	108.5
30	19.6	54.0	43.0	19.8	19.5	25.1	53.5	20.1	118.9	17.8	54.0	39.4	18.5	18.6	25.1	53.0	113.2
32	20.4	56.3	44.9	20.7	20.3	26.5	55.1	21.0	123.3	18.5	56.3	41.1	19.3	19.3	26.5	54.6	117.8
34	21.1	58.5	46.6	21.5	21.0	27.8	56.7	21.8	127.5	19.1	58.5	42.7	20.1	20.1	27.8	56.2	122.1
36	21.7	60.6	48.4	22.3	21.7	29.2	58.3	22.6	131.5	19.8	60.6	44.4	20.9	20.8	29.2	57.7	126.4
38	22.4	62.8	50.1	23.1	22.3	30.5	59.9	23.4	135.4	20.4	62.7	46.0	21.7	21.5	30.5	59.3	130.4
40	23.0	64.8	51.8	23.9	22.9	31.9	61.4	24.1	139.1	21.0	64.8	47.6	22.4	22.2	31.8	60.9	134.3
42	23.6	66.9	53.6	24.6	23.6	33.2	63.0	24.9	142.6	21.7	66.8	49.2	23.2	22.9	33.2	62.4	138.1
44	24.2	68.9	55.2	25.3	24.2	34.5	64.5	25.6	145.9	22.3	68.8	50.8	23.9	23.5	34.5	63.9	141.7
46	24.7	70.7	56.8	26.1	24.8	35.9	66.1	26.3	149.1	22.8	70.6	52.3	24.6	24.2	35.8	65.4	145.2
48	25.3	72.6	58.4	26.8	25.4	37.2	67.6	27.0	152.1	23.4	72.4	53.8	25.3	24.8	37.1	66.8	148.5
50	25.8	74.3	60.0	27.5	26.0	38.6	69.0	27.7	155.0	24.0	74.2	55.3	26.0	25.5	38.4	68.3	151.7
52	26.3	76.0	61.5	28.2	26.6	39.9	70.5	28.4	157.7	24.5	75.9	56.7	26.6	26.1	39.7	69.6	154.7
54	26.7	77.7	63.0	28.9	27.2	41.2	71.9	29.1	160.3	25.0	77.5	58.1	27.3	26.7	41.0	71.0	157.6
56	27.2	79.2	64.5	29.6	27.8	42.6	73.3	29.7	162.8	25.5	79.0	59.5	27.9	27.2	42.3	72.4	160.4
58	27.7	80.7	66.0	30.2	28.3	43.9	74.7	30.3	165.4	26.0	80.5	60.9	28.6	27.8	43.6	73.7	163.2
60	28.1	82.2	67.5	30.9	28.8	45.2	75.9	30.9	167.8	26.5	81.9	62.3	29.2	28.4	44.9	75.0	165.9

Table 1 gives the RMS phase difference measurements for BW098A on 3C274 for point separations from 2 to 60 seconds. It is apparent that there is significant variation between stations. The best, BR, KP, LA, OV, and PT, are all similar. SC is about 6 times worse while the other stations are in between, mostly closer to the better stations. Tables 2, 3, and 4 give similar data for the first 30 seconds for 3C279 in BW098A, and 3C274 and 3C279 in BW115D. The tables stop at 30 seconds to save space and because the longer intervals are only of interest for this memo for the VLBA fast switching case. For BW115D, the better stations are FD, KP, LA, NL, and PT. OV is extremely poor despite no obvious indicators of trouble in the weather data, Tsys data, or logs (wind speeds low, no precipitation, Tsys stable). The phase vs time plot for one of the 3C274 scans of BW115D is shown in Figure 1. The rather wild phases of OV for this project are clear. In BW098A, it was one of the best stations. I'll hazard a guess that there was rotor from the Sierra wave affecting the data on the day BW115D was observed. The two projects were dynamically scheduled with high priority and had roughly half the antennas in really good conditions. These are small number statistics but perhaps they are an indications of what to expect on a widely distributed array.

For both projects, the 3C279 data are somewhat lower noise then 3C274. This is likely a result of some combination of the higher SNR of 3C279, which is close to an order of magnitude stronger than 3C274,

Table 2: The RMS phase difference (deg) vs baseline and time separation (seconds) for 3C279 in BW098A.

Time sep	BR-OV	FD-OV	HN-OV	KP-OV	LA-OV	MK-OV	NL-OV	OV-PT	OV-SC	BR-PT	FD-PT	HN-PT	KP-PT	LA-PT	MK-PT	NL-PT	PT-SC
2	3.1	8.6	10.3	2.6	2.7	2.3	9.8	2.9	14.2	3.4	8.8	10.1	2.8	2.9	2.6	9.9	14.2
4	5.2	15.6	18.4	4.3	4.2	3.8	17.5	4.8	27.3	5.6	15.9	18.3	4.7	4.6	4.3	17.5	27.3
6	6.8	21.6	25.6	5.7	5.2	5.0	23.9	6.2	39.5	7.4	22.0	25.6	6.3	5.9	5.8	23.8	39.3
8	8.2	27.1	32.4	6.9	6.0	6.1	29.6	7.2	50.7	8.9	27.6	32.5	7.7	7.1	7.0	29.4	50.5
10	9.5	32.2	39.0	8.0	6.7	7.0	34.8	8.1	61.4	10.3	32.9	39.2	8.9	8.1	8.1	34.4	61.2
12	10.7	36.8	45.5	9.0	7.5	7.9	39.3	8.9	71.6	11.6	37.8	45.5	10.0	9.0	9.2	38.8	71.4
14	11.8	41.3	51.6	10.0	8.2	8.7	43.3	9.7	81.3	12.7	42.5	51.2	11.0	9.9	10.3	42.7	81.1
16	12.7	45.5	57.2	11.0	8.8	9.4	46.9	10.2	90.7	13.7	47.0	56.6	11.9	10.7	11.2	46.1	90.3
18	13.5	49.7	62.5	11.8	9.4	10.1	50.0	10.7	99.3	14.7	51.5	61.6	12.7	11.3	12.0	49.1	99.1
20	14.4	53.6	67.0	12.7	10.0	10.7	52.6	11.1	107.1	15.7	55.8	66.1	13.5	12.0	12.9	51.7	107.1
22	15.3	57.2	71.3	13.5	10.5	11.4	54.8	11.6	114.1	16.8	59.7	70.4	14.2	12.6	13.7	53.9	114.4
24	16.2	60.5	75.6	14.3	11.0	11.9	56.5	12.1	120.4	17.9	63.4	74.6	14.8	13.3	14.6	55.5	121.0
26	17.2	63.9	79.7	15.0	11.6	12.5	57.9	12.7	126.1	19.1	67.0	78.7	15.4	14.0	15.5	56.7	127.0
28	18.2	66.9	83.5	15.7	12.2	13.1	58.9	13.2	131.3	20.3	70.4	82.5	15.9	14.6	16.4	57.5	132.3
30	19.2	69.8	87.0	16.5	12.6	13.6	59.6	13.6	136.0	21.5	73.5	86.2	16.5	15.2	17.4	58.0	137.2

Table 3: The RMS phase difference (deg) vs baseline and time separation (seconds) for 3C274 in BW115D.

Time sep	BR-LA	FD-LA	HN-LA	KP-LA	LA-MK	LA-NL	LA-OV	LA-PT	LA-SC	BR-PT	FD-PT	HN-PT	KP-PT	MK-PT	NL-PT	OV-PT	PT-SC
2	6.4	5.2	9.0	4.4	7.8	5.9	18.1	4.3	11.5	6.2	4.6	9.2	3.7	7.3	5.7	18.2	11.6
4	9.9	7.9	13.3	6.7	10.1	8.0	32.2	6.6	17.9	9.5	6.9	13.3	5.6	9.2	7.4	32.5	18.0
6	13.1	10.4	17.3	8.7	12.2	10.0	44.2	8.7	23.7	12.6	8.9	17.3	7.2	11.0	9.2	44.4	23.7
8	16.0	12.5	20.9	10.5	14.1	11.8	54.8	10.6	28.7	15.4	10.6	21.0	8.7	12.8	10.8	54.9	28.8
10	18.7	14.5	24.2	12.1	16.0	13.5	64.3	12.4	33.3	18.0	12.3	24.4	10.0	14.5	12.3	64.3	33.3
12	21.2	16.3	27.4	13.7	17.7	15.0	73.0	14.0	37.4	20.5	13.8	27.6	11.2	16.0	13.7	73.0	37.4
14	23.7	17.9	30.3	15.2	19.4	16.5	81.0	15.6	41.1	22.8	15.3	30.7	12.4	17.4	15.1	80.9	41.1
16	25.9	19.5	33.1	16.6	21.0	17.9	88.5	17.0	44.7	25.1	16.6	33.5	13.5	18.8	16.3	88.3	44.6
18	28.1	21.0	35.7	17.9	22.4	19.2	95.5	18.3	47.9	27.3	17.8	36.2	14.5	20.1	17.6	95.3	47.8
20	30.2	22.3	38.2	19.2	23.9	20.4	102.1	19.6	51.0	29.4	19.0	38.6	15.5	21.3	18.8	101.9	50.8
22	32.2	23.6	40.6	20.4	25.2	21.6	108.4	20.8	54.0	31.4	20.1	41.0	16.5	22.5	19.9	108.2	53.7
24	34.2	24.8	42.8	21.6	26.4	22.7	114.2	21.9	56.8	33.3	21.1	43.2	17.4	23.6	21.0	114.1	56.5
26	36.0	25.9	44.9	22.8	27.6	23.8	119.8	23.0	59.5	35.1	22.1	45.3	18.3	24.6	22.1	119.6	59.2
28	37.8	27.0	46.9	23.9	28.7	24.9	125.2	24.1	62.0	36.8	23.0	47.3	19.1	25.6	23.1	124.7	61.8
30	39.6	28.0	48.9	25.0	29.8	25.9	130.4	25.1	64.5	38.5	23.8	49.3	19.9	26.6	24.0	129.6	64.2

Table 4: The RMS phase difference (deg) vs baseline and time separation (seconds) for 3C279 in BW115D.

Time sep	BR-LA	FD-LA	HN-LA	KP-LA	LA-MK	LA-NL	LA-OV	LA-PT	LA-SC	BR-PT	FD-PT	HN-PT	KP-PT	MK-PT	NL-PT	OV-PT	PT-SC
2	4.9	4.7	6.9	3.9	5.0	3.9	22.3	3.6	11.7	4.6	4.3	6.3	3.1	4.9	3.4	22.7	11.3
4	8.2	8.1	12.0	6.6	7.8	6.5	40.6	6.0	21.7	7.8	7.5	11.2	5.2	7.9	5.7	41.4	20.7
6	10.9	11.1	16.2	8.7	10.2	8.8	56.5	7.9	30.0	10.5	10.3	15.1	6.9	10.6	7.7	57.4	28.2
8	13.5	13.9	19.9	10.4	12.2	10.7	70.9	9.6	37.0	13.0	13.0	18.8	8.4	13.0	9.4	71.8	34.2
10	16.0	16.5	23.4	11.9	14.1	12.5	84.6	10.9	43.1	15.3	15.4	22.2	9.7	15.3	10.9	85.5	39.1
12	18.2	19.0	27.0	13.1	15.8	14.1	98.0	12.0	48.7	17.5	17.6	25.7	10.9	17.3	12.2	99.1	43.1
14	20.3	21.3	30.6	14.1	17.3	15.6	111.3	13.0	54.0	19.5	19.7	29.4	11.9	19.0	13.2	112.0	46.5
16	22.0	23.5	34.1	14.8	18.5	17.0	124.3	13.9	59.1	21.4	21.6	32.7	12.7	20.3	14.0	124.3	49.1
18	23.6	25.7	37.2	15.6	19.6	18.4	136.4	14.8	64.2	23.1	23.4	35.7	13.5	21.4	14.3	135.8	51.1
20	25.0	27.7	40.0	16.4	20.6	19.7	147.8	15.7	69.5	24.4	25.1	38.5	14.2	22.5	14.4	146.4	52.7
22	26.2	29.6	42.4	17.2	21.6	21.0	158.6	16.4	75.0	25.5	26.8	40.7	14.9	23.5	14.3	156.6	54.4
24	27.4	31.4	44.7	18.1	22.8	22.4	169.0	17.1	80.3	26.6	28.4	42.5	15.6	24.6	14.3	166.6	55.7
26	28.5	33.0	46.8	19.0	23.9	23.8	179.1	17.6	85.3	27.8	29.7	44.1	16.4	25.6	14.6	176.4	55.5
28	29.4	34.4	48.6	19.8	24.9	25.2	188.5	18.1	90.1	28.7	30.7	45.1	17.1	26.5	14.9	185.6	57.2
30	30.2	35.6	50.2	20.7	25.9	26.5	197.5	18.6	94.8	29.4	31.6	45.9	17.9	27.3	15.2	194.7	58.7

Figure 1: Phase vs time for 3C274 in BW115D over 14 minutes. The abnormally high fluctuations at OV are apparent. This is a relatively good period at SC.

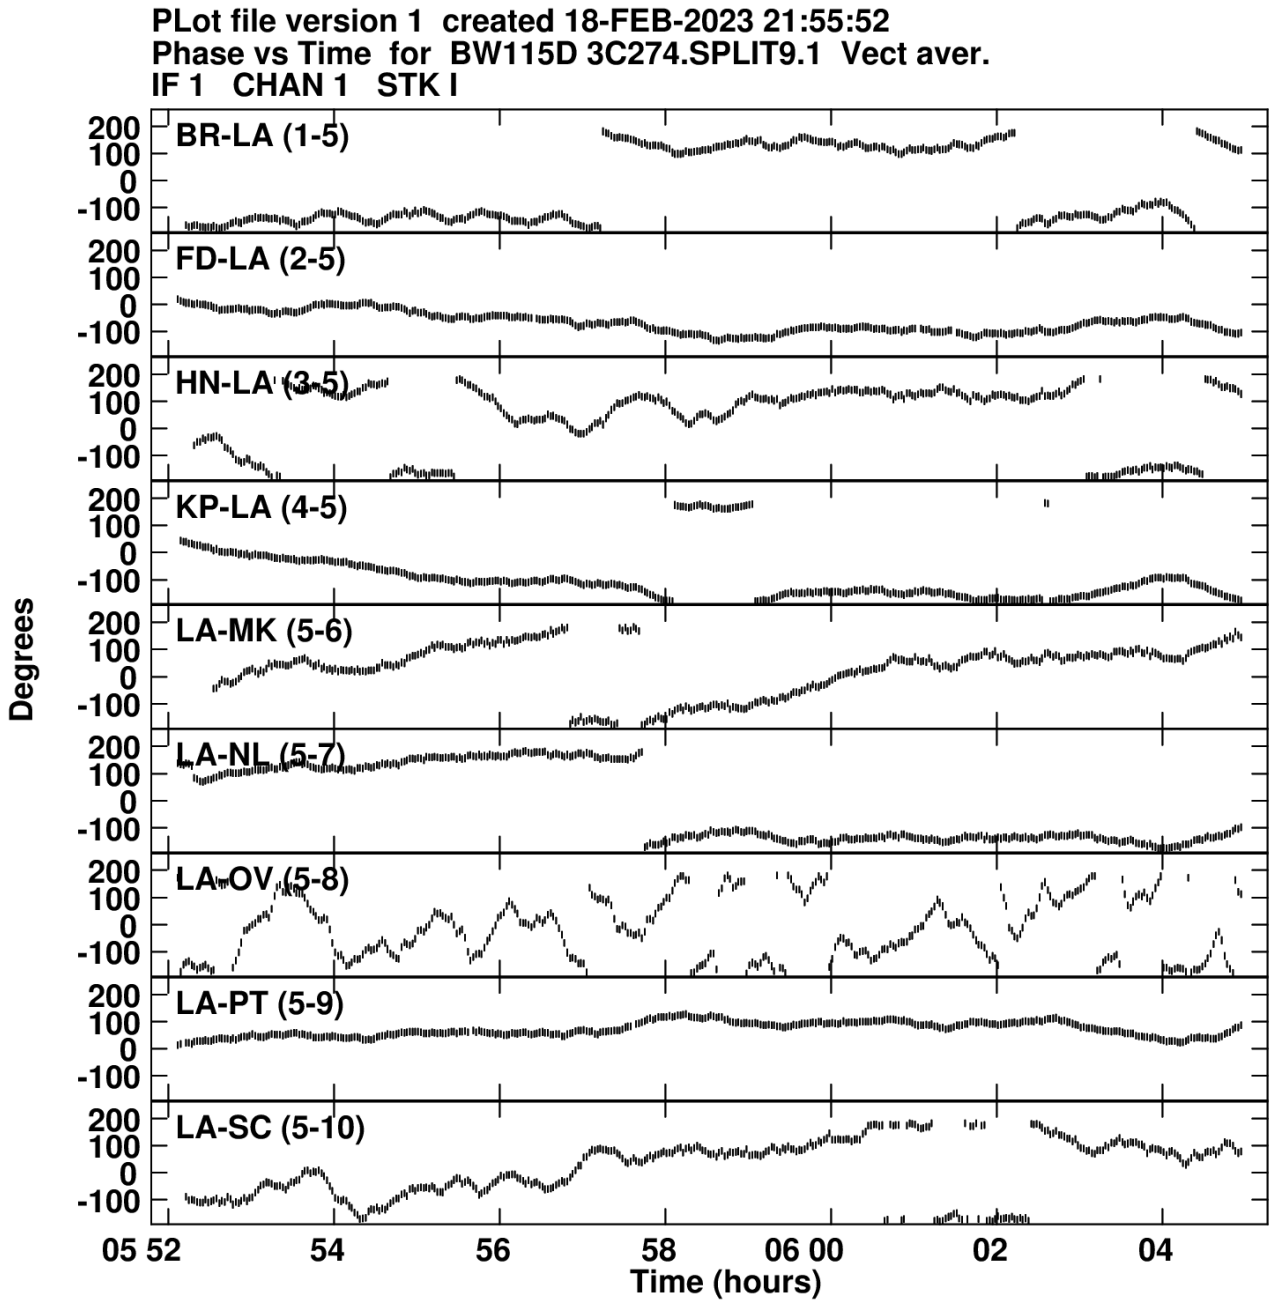

and of the significantly larger amount of data on 3C274, including some periods of low elevations or shorter term atmospheric issues.

A table similar to Table 1 was made, but not shown here, for 3C279 in BW098A using a data set that still had 8 separate IFs and had been self-calibrated record-by-record (2 second solution interval). The RMS phase differences ranged from 0.4 to 3.5° across baselines with only the HN baselines exceeding

1.4. Within each baseline, the RMS was constant with time interval to within 0.1° for intervals of 2 to 60 seconds. These self-calibrated data should retain the thermal noise but not the tropospheric fluctuations. Therefore, especially considering they are made from data with 1/8 the bandwidth (single IF data) of what was used for the other tables, they are a good indication that thermal noise does not play a significant role in the tabulated RMSs for these strong sources.

Plots of log(RMS) vs log(time interval) based on the data of Tables 1 and Table 3 extended to 60 s are shown in Figure 2. This shows the power law nature of the RMS variations for most baselines. A few baselines, most notably those in BW098A involving HN, flatten somewhat toward short times suggesting a noise floor. But the better baselines give rather straight power laws. The power law indices were calculated in two ways. An analytic slope was obtained as ($\alpha = (\log(RMS_{30} - RMS_6) / (\log(30) - \log(6))$) where RMS_{30} and RMS_6 are the values from Table 1 or Table 3 for 30 and 6 seconds. The average of those slopes are 0.60 ± 0.7 for BW098A and 0.63 ± 0.03 for BW115D using all baselines. Restricting to the baselines involving the best stations gives 0.58 ± 0.01 for BW098A and 0.62 ± 0.004 for BW115D. The second method was to do a non-linear least squares fit of the RMS phase differences to the parameters of the equation ($RMS = a + b\ t^{c}$). Here a is a noise floor, b is a scale factor, and c is the power law index. This was done after a warning that including the noise floor could make a significant difference (Dodson, private communication). This was only done for the baselines between the 5 best stations in each epoch. They used all points between 2 and 30 seconds. Beyond 30 seconds, the plots suggest the beginning of a flattening which is not described by the equation. The results are in the table in the bottom of Figure 2. The errors on each parameter are the formal errors from the fits and seem rather optimistic for overall errors given the scatter between baselines and days. That scatter, along with the presence of unphysical negative noise floors on a few baselines, suggests that, while the a power law is a good description of the phase noise behavior, it is not perfect and is likely somewhat time variable. The average index for the baselines shown is 0.64 ± 0.04 for BW098A and 0.63 ± 0.05 for BW115D. These are not significantly different from the analytic results.

The data shown here can be compared to the root phase structure function of Figure 1 of C15 which is Figure 6 of Carilli and Holdaway (1999; hereafter CH). To make that comparison, the time offsets of the VLBA data have to be translated to baseline lengths. That is done by assuming a nominal wind speed aloft of 10 m/s, the same value used by CH. The scale on the top of Figure 1 is such a distance scale. The CH data are mostly for larger separations than our time offsets times the wind speed. The data also need to be scaled by a factor of 2 because we observed at 43 GHz while the CH data are for 22 GHz. Finally the VLBA data sample the fluctuations over two independent locations while the VLA measures the gradients at a single site. The VLBA baseline RMS^2 will thus be assumed to be the quadrature sum of the RMS^2 over each site. For the best baselines, the two sites are likely to be similar so a correction of $1/\sqrt{2}$ or about 1/1.4 is applied for the comparison. As an example, if we take our 30 second offsets at 10 m/s, we can compare with the 300 m baselines of CH. Reading off the CH figure, for a 300 m baseline, the RMS phase is about 12° uncorrected for instrumental noise or 6.3° if corrected. For the VLBA, the average RMS phase difference for the best stations for a 30 second interval is about 19° for BW098A and 24° for BW115D. Scaling by 1/(2*1.4) gives 6.8° and 8.6° for the two days. That is in good agreement with the corrected VLA data given the small number of dates involved and the uncertainties, especially in the conversion between distance and time. In this comparison, no correction is made for noise floor. The fit results in Figure 2 indicate a very small noise floor and I am not aware of a known noise floor on the VLBA that would be apparent here (maser noise would be lower).

Figure 2: Plots of the data from Table 1 and 3 showing the RMS phase difference vs time interval for BW098A (left) and BW115D (right) – in log-log form. The top scale is the time offsets converted to a distance using a nominal wind speed aloft of 10 m/s. The table below gives the results of a least squares fit to the points between 2 and 30 seconds.

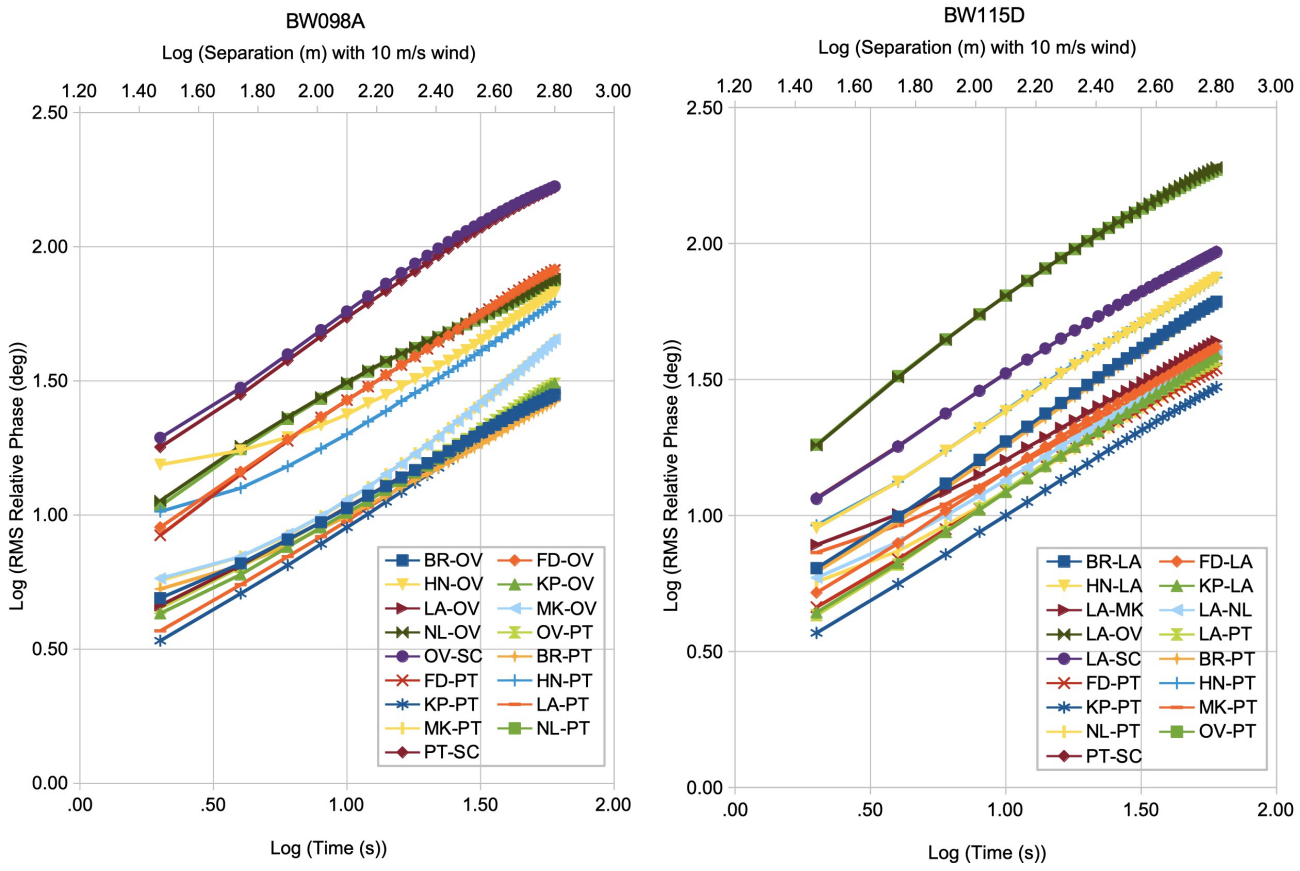

Results of non-linear least squares fits for the parameters of the equation (RMS = a + b t^c) on the best baselines:

BW098A				BW115D			
Baseline	a: Noise floor	b: Scale	c: Index	Baseline	a: Noise floor	b: Scale	c: Index
BR-OV	1.5 ± 0.12	0.51 ± 0.025	0.624 ± 0.008	FD-LA	-1.6 ± 0.40	1.25 ± 0.131	0.557 ± 0.016
KP-OV	1.2 ± 0.14	0.42 ± 0.026	0.662 ± 0.009	KP-LA	0.3 ± 0.08	0.55 ± 0.015	0.666 ± 0.004
LA-OV	0.4 ± 0.15	0.77 ± 0.040	0.564 ± 0.008	LA-NL	1.5 ± 0.25	0.61 ± 0.048	0.648 ± 0.012
OV-PT	1.0 ± 0.23	0.54 ± 0.048	0.622 ± 0.014	LA-PT	-0.8 ± 0.34	0.79 ± 0.074	0.614 ± 0.015
BR-PT	2.7 ± 0.18	0.32 ± 0.031	0.675 ± 0.015	FD-PT	-0.9 ± 0.43	0.92 ± 0.108	0.577 ± 0.019
KP-PT	0.7 ± 0.06	0.34 ± 0.009	0.695 ± 0.004	KP-PT	-0.0 ± 0.12	0.56 ± 0.026	0.628 ± 0.007
LA-PT	0.1 ± 0.16	0.57 ± 0.036	0.611 ± 0.010	NL-PT	2.3 ± 0.22	0.38 ± 0.033	0.711 ± 0.014

A difference between the data sets is the power law index at short spacings – all we are looking at on the VLBA. CH see an index of 0.85, close to the 0.83 predicted for short baseline phases by Kolmogorov turbulence in the 3D case (short baselines). The VLBA data imply an index near 0.60, which is lower . The index predicted for Kolmogorov turbulence in the 2D case is 0.33 which is closer to the 0.43 that CH see on scales of 1.2 to 6 km. It is not clear to me why the VLBA data consistently fall between these cases. As can be seen in Figure 2, the index does not vary much across baselines or between days, so the effect is clearly systematic. The index is not affected by simple scale changes like the assumed wind speed. I suspect that the explanation for the index lower than Kolmogorov will involve the fact that the VLBA baseline involves two independent samples of the atmosphere. If anyone can provide an explanation for our index, that would be welcome.

In an attempt to explore the efficacy of using interpolation to correct phases, Table 5 shows the RMS deviation of the phase of samples after subtracting the mean of two the points t_c seconds ("Time sep" in the tables) before and after the sample. This mimics the fast switching calibration method with t_c being the time between calibrator and target . The data sets are those of Tables 1 and 3, namely 3C274 in BW098A and BW115D. If the RMS phases of the earlier tables were primarily the result of persistent linear phase slopes, the phases corrected by interpolation would have RMS fluctuations significantly smaller than those in Table 1. It is also likely that the power law index would be different for short times. RMSLISTR was tested with fake data to make sure predominantly linear phase slopes resulted in low RMS values for the interpolated data. The result was as expected. The middle subtable of Table 5 shows the ratio of the RMS of the interpolated data to the RMS of the uncorrected phases from Table 1 for the same time interval, t_c, for BW098A. The ratios are typically near 0.65 and are rather constant across baselines and especially across time on each baseline.

It turns out that the ratio of 0.65 of the RMS fluctuations of the interpolated points to the RMS fluctuations of the unmodified points can be derived from the 0.6 power law structure function seen in these data. To show that, consider that the total phase at time separation t_c is the sum of the interpolation phase (average of points t_c before and after) plus the interpolated phase (residual after subtraction of the interpolation phase). The RMS of that total will be the quadrature sum of one half the RMS for separations of 2 t_c (the phases are divided by 2 in the interpolation) and the RMS of the interpolated phase. Rearranging to solve for the RMS of the interpolated phase and using RMS = a t^{α}, where α is the power law index, the RMS of the interpolated phase will be

$$\text{RMSi} = \sqrt{(a\, t_c^{\alpha})^2 - (\tfrac{1}{2}\, a\, (2\, t_c)^{\alpha})^2} = a\, t_c^{\alpha} \sqrt{1 - (\tfrac{1}{2}\, (2^{\alpha})^2}$$

For $\alpha = 0.60$, this is the total RMS at t_c times 0.65 as observed.

The ratio calculated above depends on the power law exponent of the structure function and would be different for other power laws. For example, it would be 0.46 for the Kolmogorov 3D index of 0.83. For a fake data set with the fluctuation being set entirely by a gaussian distributed phase change between any two adjacent points (no longer term correlations), the index is 0.5 and the extra component to the RMS fluctuations after interpolation is 0.707 of the uninterpolated value. This is the same as the $1/\sqrt{2}$ expected for the quadrature sum of the RMS fluctuations of two calibrator values weighted by ½ each. While calculating the performance of some calibration methods below, I will take advantage of the similarity of the interpolated result to just treat the two "calibrator" points as independent estimates of the calibration phase even for the 0.6 index. Given the complexities that will be discussed in determining the line-of-sight separations, this is not likely to be the limiting factor in the quality of the RMS fluctuation estimates.

Table 5. RMS phase for 3C274 points corrected by the average of points "Time sep" seconds before and after. the middle table is the ratio of the points with the interpolated correction to the uninterpolated points from Table 1. These top two tables are for BW098A. At the bottom is the same table as the top, but for BW115D which had better conditions at SC.

Time sep	BR-OV	FD-OV	HN-OV	KP-OV	LA-OV	MK-OV	NL-OV	OV-PT	OV-SC	BR-PT	FD-PT	HN-PT	KP-PT	LA-PT	MK-PT	NL-PT	PT-SC
2	3.7	5.3	12.7	3.1	3.2	4.6	6.8	3.1	12.5	4.2	4.6	8.2	2.3	2.4	4.5	6.2	11.1
4	4.6	8.6	13.7	4.0	4.4	5.0	11.6	4.3	16.9	4.8	8.2	9.0	3.2	3.6	5.0	11.2	15.8
6	5.5	11.7	14.5	5.0	5.5	5.6	15.4	5.4	22.3	5.5	11.4	10.3	4.1	4.6	5.6	14.9	21.6
8	6.4	14.4	15.5	6.0	6.4	6.2	18.7	6.4	27.8	6.2	14.3	11.7	4.9	5.4	6.1	18.2	27.1
10	7.3	16.8	16.6	6.8	7.3	6.9	21.8	7.3	33.3	6.9	16.7	13.0	5.6	6.2	6.7	21.4	32.4
12	8.0	18.8	18.1	7.5	8.1	7.4	24.7	8.0	38.8	7.6	18.8	14.4	6.3	6.9	7.3	24.4	37.0
14	8.7	20.8	19.0	8.1	8.8	8.0	27.0	8.6	44.1	8.2	20.7	15.6	7.0	7.7	7.9	26.8	41.5
16	9.3	22.9	20.4	8.7	9.4	8.5	29.0	9.1	49.3	8.7	22.6	17.0	7.5	8.4	8.6	28.7	45.9
18	10.0	24.9	21.4	9.2	10.0	9.0	30.7	9.6	54.3	9.2	24.5	18.3	8.1	9.0	9.1	30.4	50.2
20	10.6	26.8	22.4	9.7	10.5	9.4	32.2	10.1	59.0	9.7	26.5	19.4	8.7	9.6	9.6	31.9	54.5
22	11.2	28.7	23.5	10.3	11.0	9.9	33.5	10.7	63.6	10.1	28.4	20.7	9.2	10.1	10.2	33.3	58.7
24	11.8	30.7	24.6	10.8	11.5	10.4	34.9	11.3	68.1	10.5	30.6	21.9	9.8	10.6	10.7	34.6	62.9
26	12.3	32.8	25.9	11.3	11.9	10.9	36.1	11.8	72.6	11.0	32.8	23.1	10.3	11.1	11.2	35.8	67.3
28	13.0	35.0	26.8	11.8	12.3	11.4	37.2	12.4	77.2	11.5	35.1	24.3	10.8	11.6	11.8	36.9	71.7
30	13.6	37.1	28.0	12.4	12.8	11.9	38.3	13.1	81.5	11.9	37.4	25.4	11.3	12.0	12.4	38.0	75.8

Ratio of above to points in Table 1.

Time sep	BR-OV	FD-OV	HN-OV	KP-OV	LA-OV	MK-OV	NL-OV	OV-PT	OV-SC	BR-PT	FD-PT	HN-PT	KP-PT	LA-PT	MK-PT	NL-PT	PT-SC
2	0.74	0.59	0.82	0.71	0.70	0.79	0.60	0.70	0.64	0.79	0.54	0.79	0.67	0.66	0.79	0.57	0.62
4	0.70	0.59	0.78	0.66	0.69	0.71	0.64	0.68	0.57	0.73	0.58	0.72	0.64	0.65	0.72	0.64	0.56
6	0.69	0.61	0.74	0.66	0.68	0.66	0.66	0.67	0.56	0.71	0.60	0.68	0.63	0.65	0.66	0.65	0.57
8	0.68	0.62	0.72	0.66	0.68	0.63	0.68	0.68	0.57	0.70	0.62	0.66	0.63	0.65	0.62	0.67	0.59
10	0.68	0.63	0.70	0.66	0.68	0.61	0.70	0.68	0.58	0.70	0.63	0.65	0.63	0.65	0.59	0.69	0.60
12	0.68	0.63	0.69	0.66	0.68	0.58	0.71	0.68	0.59	0.70	0.62	0.64	0.63	0.65	0.57	0.71	0.60
14	0.68	0.63	0.68	0.65	0.68	0.57	0.72	0.67	0.61	0.69	0.62	0.64	0.62	0.65	0.56	0.72	0.61
16	0.68	0.63	0.68	0.65	0.68	0.55	0.72	0.66	0.62	0.69	0.63	0.64	0.62	0.66	0.55	0.72	0.61
18	0.68	0.64	0.67	0.64	0.68	0.53	0.73	0.65	0.63	0.68	0.63	0.64	0.62	0.66	0.54	0.73	0.62
20	0.68	0.64	0.66	0.64	0.67	0.51	0.73	0.65	0.64	0.68	0.64	0.64	0.62	0.66	0.52	0.73	0.63
22	0.68	0.65	0.66	0.63	0.67	0.50	0.72	0.65	0.65	0.67	0.64	0.64	0.62	0.65	0.52	0.72	0.63
24	0.68	0.66	0.65	0.63	0.66	0.49	0.72	0.65	0.65	0.67	0.65	0.64	0.61	0.65	0.51	0.72	0.64
26	0.68	0.67	0.66	0.63	0.66	0.49	0.72	0.65	0.66	0.67	0.67	0.64	0.61	0.65	0.50	0.72	0.65
28	0.69	0.68	0.65	0.63	0.66	0.48	0.72	0.65	0.68	0.67	0.68	0.65	0.61	0.65	0.50	0.72	0.66
30	0.69	0.69	0.65	0.63	0.66	0.47	0.72	0.65	0.69	0.67	0.69	0.65	0.61	0.65	0.49	0.72	0.67

For BW115D:

Time sep	BR-LA	FD-LA	HN-LA	KP-LA	LA-MK	LA-NL	LA-OV	LA-PT	LA-SC	BR-PT	FD-PT	HN-PT	KP-PT	MK-PT	NL-PT	OV-PT	PT-SC
2	4.1	3.3	6.1	2.9	6.0	4.4	8.3	2.7	7.2	4.0	3.1	6.3	2.5	5.7	4.3	8.3	7.4
4	5.9	4.9	8.2	4.1	7.1	5.3	17.1	4.0	10.7	5.7	4.4	8.3	3.5	6.6	5.0	17.4	10.8
6	7.7	6.5	10.6	5.3	8.3	6.4	24.9	5.2	14.6	7.4	5.6	10.6	4.5	7.6	5.9	25.4	14.7
8	9.4	7.9	12.8	6.5	9.5	7.4	32.1	6.4	18.2	9.0	6.7	12.8	5.5	8.7	6.8	32.7	18.2
10	11.1	9.3	15.0	7.5	10.6	8.5	38.7	7.7	21.3	10.6	8.0	15.2	6.4	9.8	7.7	39.5	21.4
12	12.7	10.7	17.2	8.5	11.7	9.4	44.9	8.9	24.1	12.2	9.2	17.5	7.2	11.0	8.6	45.9	24.3
14	14.4	12.1	19.3	9.4	12.9	10.4	50.7	10.1	26.8	13.9	10.3	19.9	8.1	12.1	9.5	51.9	26.8
16	15.9	13.3	21.4	10.3	14.0	11.2	56.1	11.2	29.4	15.6	11.4	22.1	8.8	13.3	10.3	57.5	29.2
18	17.5	14.5	23.3	11.2	15.1	12.1	61.2	12.2	31.9	17.3	12.4	24.2	9.6	14.3	11.2	62.7	31.6
20	19.0	15.6	25.1	12.0	16.2	12.8	66.0	13.1	34.4	19.0	13.3	26.0	10.3	15.4	12.0	67.6	34.1
22	20.5	16.6	26.7	12.7	17.2	13.6	70.5	14.0	36.9	20.6	14.2	27.8	11.0	16.4	12.7	72.5	36.4
24	21.9	17.5	28.2	13.4	18.1	14.3	74.8	14.9	39.3	22.1	15.1	29.4	11.7	17.3	13.5	77.2	38.8
26	23.2	18.2	29.5	14.1	18.9	14.9	79.0	15.7	41.5	23.5	15.9	30.8	12.3	18.2	14.2	81.7	40.9
28	24.5	18.9	30.7	14.7	19.7	15.5	83.3	16.5	43.6	24.9	16.7	32.1	12.9	19.0	14.9	86.3	42.8
30	25.9	19.6	31.9	15.4	20.5	16.0	87.6	17.2	45.3	26.2	17.4	33.3	13.5	19.7	15.6	90.9	44.5

The discussion above indicates that a calibration using interpolation will only reduce the RMS fluctuations by about the factor of 0.65 over a simple calibration using one calibrator. This suggests that, for coherence improvement, methods involving short term gradient measurements do improve the results, but the factor is only by around 35%. Use of such methods will mainly be of benefit at sites where that improvement may make the difference between successful and unsuccessful phase referencing. As will be seen, the SC site of the VLBA is likely to be in that regime much of the time. These statements are about short term measurements to improve coherence in cases where the troposphere is the main contributor of phase fluctuation and there are limited or corrected systematic gradients. It does not address ionospheric fluctuations, which may have longer time scales. And especially, it does not address the need to measure even subtle gradients, probably on longer time scales, for the purpose of improving astrometric accuracy.

Phase Referencing

The data shown in this memo are based on the fluctuations in time of real VLBA data phases at 43 GHz measured on 2 different dates. For relative astrometry, and even just phase referencing for imaging, the phase measured on one or more calibrators is used to correct the phase on a target source. The method used to derive such corrections must deal with gradients between the target and calibrator positions, with gradients in time if the calibrator and target are not being observed simultaneously, and gradients in space if multiple antennas are being used. Below various calibration schemes are analyzed for how well they are expected to work in the presence of the fluctuations seen in our data and described above. The analysis is based on an assumption that the phase fluctuations are statistically uniform in all directions and can be described as a fixed screen blowing over the antennas at an assumed speed. The analysis, given the nature of the data, cannot capture the effects of offsets that do not show up in the time series. For example slowly variable offsets that are a function of zenith angle would provide additional phase errors above what is deduced in the discussions below. Such effects will be considered more fully in future memo(s) based on theory and/or data sets that that have multiple sources and fast switching. Note that any such effects will almost certainly increase the phase offsets over what is derived here, although the time scales should be longer making the effects easier to calibrate.

For a station with a single antenna, such as those of the VLBA or ngVLA MAIN (all but LONG), phase referencing is done with fast switching (I do not address Water Vapor Radiometry in this memo). For fast switching, the main concerns are the interpolation of the calibrator phase to the time of the target observation and the phase offsets between the positions on the sky of the calibrator and target. For stations with multiple antennas such as are planned for ngVLA LONG, fast switching could be used, but another, presumably better, option is available. That option is paired antenna calibration, where targets and calibrators are observed simultaneously. Interpolation in time is not required. But now the lines-of-sight do not originate in the same place, so phase offsets between the locations of the antennas on the ground are of concern.

Paired antenna calibration provides some significant benefits over fast switching beyond just not having to interpolate in time. The schedules are more relaxed. The scans are longer and hence easier to assess for quality and less subject to data loss when the source switching process does not go smoothly. The significant dead time due to slewing and settling between sources is not lost. There is more flexibility to use a variety of calibration styles with multiple calibrators. However note that, if two antennas are present and both are switching in the fast switching case, the sensitivity is similar to

Table 6: Summary of expected difference phase RMS using various phase referencing techniques. The numbers shown are for a source separation of about 2°, a troposphere scale height of 1.4 km, a 10 m/s wind, and antenna separations of 100 m or 50 m. Fast switching assumes a 30 second full cycle on the ngVLA. Estimates are shown for SC for BW098A and for BW115D, which had rather different conditions. The sites with intermediate conditions in both observations were HN and MK(!). OV was very good in BW098A and twice as bad as SC in BW115D. The good site numbers are eyeball averages.

Case	Phase Referencing Method	Approximate calibration phase RMS error at 43 GHz (deg)		
		Good sites BW098A	SC BW098A	SC BW115D
0	Contribution from 2° calibrator-target offset	10	46	29
1	Fast switching, VLBA (50 sec). Includes calibrator interpolation error and 2° offset source error.	15	80	49
2	Fast switching, ngVLA single antenna (30 sec). Includes calibrator interpolation error and offset source error.	13	64	40
3	Fast switching, ngVLA single antenna (30 sec), single source interpolation only. Used for other calculations.	9	44	28
4	Paired antennas 100 m between antennas, no gradient used. Use 10 sec points in Tables 1 and 3.	9	44	33
5	Paired antennas 50 m, no gradient used. Use 5 sec points.	7	33	21
6	Four antennas. Weighted quadrature sum of 100 m paired antenna RMSs (Case 4) with 3 calibrator antennas. Weights 0.333 each. Results applied simultaneously to target observed with the fourth antenna.	5	27	19
7	Four antennas. Same as Case 6 but with 50 m spacing.	4	19	12
8	Four antennas. 100 m. Gradient interpolation like Table 5, 10 sec. Alternate calculation for the same data as Case 6.	6	32	21
9	Four antennas. Like 8 but for 50 m. Table 5, 5 sec.	5	19	13
10	Three antennas. Weighted quadrature sum of 100 m paired antenna RMSs at 2 calibrator antennas. Weights 0.5, 0.5.	6	32	23
11	Three antennas. Same as 8 but 50m spacing of antennas.	5	23	15
12	Three antennas, 100 m. one each continuously observing the primary calibrator and target, and one fast switching between the two other calibrators. Determine gradient after interpolation. One antenna Case 3, plus Case 6.	6	30	21
13	Three antennas, Case 12 but with 50 m spacing.	5	24	15
14	Three antennas, Switch between gradient and target scans. Fluctuations similar to Cases 1, 2, 4, or 5. Mainly for astrometry, not coherence.			

the paired antenna case. For the fast switching case, the on-source time for the target is about a quarter of the on source time for paired antennas assuming a typical schedule where the scan lengths are similar to the slew/settle times. But there are two antennas on source doubling the collecting area which makes up, in sensitivity, for the lost time. That said, it is unlikely that fast switching two antennas will be a mode that is used.

In the next two sections, various calibration styles will be analyzed for their RMS fluctuations based on the time series data presented above. Most of those styles are summarized as cases in Table 6 and are discussed in the subsections below.

Case 0 and 1: Fast Switching, VLBA, Including Effect of Calibrator-Target Offset.

The VLBA has one antenna per site so must use fast switching. The minimum time between calibrator scans is limited by antenna slew rates and settling times and by the tolerance for loss of integration time to the switching overhead. My experience is that, with the VLBA, for cycle times (time between centers of two calibrator scans) of much less than about 50 seconds, significant numbers of scans are lost and the on-source integration is small and variable on what is left. The time lost to the switch, for 1.5° (an example for which I have data), is between about 11 and 18 seconds, with the longer switches happening when the sources are near the zenith at some station giving a long azimuth slew. It is also hard to assess the quality of the data with very short scans. Other observers may switch faster, but I'll use 50 seconds for this analysis. Interpolation between calibrator scans is typically used so, to estimate the performance on the VLBA, the values from Table 5 are appropriate and the rows for around 25 seconds should be used. That gives the RMS phase at the position of the calibrator at the time of the target source observation. For the best antennas that is about 11°. For SC it is about 65° for BW098A and 40° for BW115D, which had better conditions at that site. The SC values are already high and they still have to be added in quadrature to the RMS phase due to the offset position of the target.

As pointed out by C15, this offset position gives rise to an “effective baseline” that is about the source-calibrator separation angle times the scale height of the troposphere. That product is the separation of the lines-of-sight at the scale height. For a two degree source separation and a 1.4 km scale height (a value used by Butler 2002), this baseline is about 50 m at the zenith. At higher zenith angles (Z) the distance to where the diverging lines-of-sight cross the scale height, and hence their separation, rises with sec(Z). In addition, the separation in the horizontal plane, which is at an angle of sec(Z) to the lines-of-sight, increases with an additional factor of between 1 and sec(Z) depending on how the sources are separated in azimuth vs elevation. That may matter if the wind is primarily horizontal so the fluctuations scale with the horizontal separation. The path separation will result in phase differences between the target and calibrator similar in magnitude to what is seen on the calibrator in the time it takes the wind to carry the screen a similar distance, assuming the gradients are statistically similar in all directions. Just to add another complication, the actual path length along the line of sight, and hence the total tropospheric delay also scales with sec(Z). The effect on fluctuations is not obvious, but that may not matter because the data sets used here are subject to the same effect so, to first order, it is already accounted for. If we use a nominal 45° zenith angle, sec(Z) is 1.41 and the line of sight separation in the horizontal plane is between that (for sources separated in azimuth) and $\sec^2(Z)$ is 2.00 (for sources separated mainly in elevation).

Because of these complexities, it is not clear what separation should be used, but for a representative example, choose 80 m for the separation of the lines-of-sight for the calculations for the nominal 2° source separation and 1.4 km scale height. With this assumption, the contribution from the source separation is similar to the data in the 8 second row in Table 1 and is shown in Case 0 of Table 6. The

"Good sites" value is averaged over the two observations. Adding those values in quadrature to the interpolation values gives the values shown in Case 1 of Table 6. So the referencing should work for most VLBA antennas although rather marginally and only on good days for SC, confirming experience. Scaling to 86 GHz should also work for the good antennas, but not really for SC. The tables can be examined to see what intermediate cases were seen.

Cases 2 and 3: Fast Switching on the ngVLA.

For the ngVLA, it should be possible to switch faster. The system is being designed to slew 3° and settle in a total of 7 seconds (Selina 2022). A minimum total cycle time would likely be about 30 seconds giving, for example if split evenly between calibrator and target, 8 seconds on each source plus 14 seconds total for the switches (a bit less for a 2° separation). Care should be taken in the hardware and software design so that scans of a small number of seconds can be observed reliably. If the calibrator is strong, the user may well prefer even shorter calibrator scans. Variable slewing and settling delays are part of what make short scans on the VLBA problematic. To estimate the expected ngVLA RMS calibration phase fluctuations at a given position, lines in Table 5 for near 15 seconds can be used. This assumes the use of interpolation between adjacent calibrator scans. Reading from Table 5 gives the values in Case 3 of Table 6. Again, the number for good sites is an average between both observations. Case 3 is kept as a separate case because it will prove useful in some of the gradient and frequency switching discussions below. For calibrator-target situations, the values that need to be added for the source position differences when calibrating a target are the same as for the VLBA (Case 0), so the quadrature sums are those given in Case 2 of Table 6. The data show that one ngVLA antenna can get, on good days only, below a radian for the RMS phase at 43 GHz at SC, but barely. The fluctuations at 86 GHz with fast switching will still be too large at SC and cannot be fixed with gradient measurements because they cannot be made fast enough with one antenna. The best stations will work, consistent with C15 who studied a good site, the VLA. Some of the intermediate stations will be somewhat marginal. Note that finding closer calibrators improves the calibrator-target fluctuations (Case 0) but not the switching (time offset) fluctuations (Case 3). At a 30 second cycle time and 2° calibrator-target separation, the two contributions are about equal. As the calibrator-target separation is decreased, the Case 0 contribution goes down, but not the Case 3 contribution.

Cases 4 and 5: Paired Antenna Calibration Without Gradients.

The LONG segment of the ngVLA will have another option – paired antenna calibration. There will be more than one antenna at each station located close to each other. One of the antennas can observe a calibrator all the time and its phases can be used to correct the other antennas, which are looking at the target or other calibrators. Or data from multiple antennas looking at different calibrators can be combined to beat down the fluctuations on the target. Occasionally, all of the antennas will need to observe the same calibrator to keep the instrumental phases aligned. The quality of paired antenna calibration will depend on several factors including the separation between the lines-of-sight to calibrators and target, atmospheric gradients between those lines-of-sight, and the wind speed aloft as turbulence is carried between those lines-of-sight.

For the analysis of the multiple antenna cases, an important factor is the separation of lines-of-sight between one antenna looking at the calibrator and another looking at the target. That separation is variable along the line of sight and is sensitive to several effects. The component of the separation due to antenna separation will be reduced by foreshortening of the baseline between the antennas with increasing zenith angle by an amount that depends on the angle between the baseline and the source azimuth. The component of the separation at the scale height that is due to source separation will

increase with zenith angle due to the angle between the sources and the increasing distance. The relative angle of the baseline between the antennas and the offset between sources will influence how the separation behaves with distance and that can be influenced very significantly by which antenna observes which source. That choice, for example, can determine if lines-of-sight converge or diverge with distance. It is likely that the separation that matters is in the horizontal plane in which the wind is assumed to blow. The horizontal separation due to the source separation will increase with zenith angle by an amount that depends on the relative angle of the source separation and the azimuth.

I have explored the range of these effects with a spreadsheet. The separation at the scale height can range from near zero if the lines-of-sight cross to about double the antenna separation for the 50 m antenna separation, a 45° representative zenith angle, a 2° source separation and the optimal choice of which antenna observes which source. For the 100 m antenna separation with the same zenith angle and source separation, the lines-of-sight separation can be up to around 30% more than the antenna separation with an optimal source-antenna pairing. The difference between the cases is the result of the relative dominance of the source separation and the antenna separation. The separations can be much larger if the wrong allocation of sources to antennas is made. It is clear that any scheduling software needs to be taught how to pick which antenna observes which source because a mistake there can make differences on the order of a factor of 2 or more. A more thorough study, perhaps with Monte Carlo simulations, should probably be done at some point.

This is clearly a complex situation, but for purposes of the calculations in this memo, I will assume that the site lines are separated by about the antenna separation in the multi-antenna case. I will also continue to assume a nominal 10 m/s wind speed aloft. Two cases, those for 50 m and 100 m separations will be explored. With these assumptions, the lines in Table 1 for 5 (average 4 and 6) and 10 seconds give estimates of the residual RMS fluctuations when using paired antenna calibration in the absence of gradient measurements. Cases 4 and 5 in Table 6 show those values. For the 100 m separation, the gain over fast switching without changing source (Case 3), such as one might do if band switching, is not great. This has interesting implications for multi-band observations to be discussed later. For this situation, it appears that the losses to the switch cycle in fast switching, which are somewhat lessened by the interpolation that is possible, are comparable to the loss due to the separated antennas for paired antenna calibration. The situation changes if the fast switching is to another source, or if the 50 m antenna separation is used in either case. There is a clear advantage for the paired antenna method. These results indicate that there is a significant benefit to making the antenna separation as small as possible. The paired antenna calibrations also have significant advantages in terms of reducing calibration overhead and providing continuous data sets. However it is clear that trading time offsets (fast switching) for antenna position offsets (paired antennas), while beneficial, does not completely eliminate the tropospheric phase fluctuations.

Phase Referencing With Gradients

The rest of the calibration methods to be considered here involve attempts to use gradients, or at least multiple calibrators, to improve the calibration.

There are two regimes of gradients of concern for calibrating phase reference data. One is the systematic gradients, which vary slowly and are caused by errors in the model such as source position errors, station position errors, Earth orientation errors, tropospheric and ionospheric model errors etc. These don't have significant impact on coherence in the sense that they don't seriously reduce the allowed integration times for data averaging or self calibration (or fringe fitting) intervals. But they

can be large enough to significantly degrade imaging or flux recovery, especially for larger source-calibrator separations. Also they must be dealt with for the highest accuracy astrometry. See Reid (2022) and Rioja and Dodson (2020) for discussions of how to deal with these issues, mainly in the context of single antenna stations. There is a current effort to address these issues when multiple antennas per station are available and there should be an NGVLA memo on the topic soon. My 43 GHz, single-source, phase stability data provide no insight into dealing with long term gradients so they will not be discussed further here.

The other type of gradient is that associated with the short term fluctuations that affect coherence. My data can supply insight about these as they manifest as phase slopes in time as the pattern blows over the antennas. They are the topic of the rest of this section. One aspect of such gradients, that I cannot directly address, is whether the assumption of a mostly fixed pattern blowing over the antennas is actually true. That assumption underlies the assumed relationship between the temporal gradients seen in the data and the derived spatial gradients. That assumption is used by other studies as noted in the Introduction. The discussion above about the interpolation results in Table 5 shows that the linear gradient approximation is only approximate.

From the tables and the discussions of calibration methods above, it is apparent that most baselines should be sufficiently stable that decent coherence can be maintained, even at ngVLA Band 6 (70-116 GHz), with phase referencing, either using paired antennas or using fast switching with interpolation. Explicit gradient calibration would not be needed unless there are gradients between sources that are not captured in our time series data or unless trying for the best astrometric accuracy. But some stations and some conditions require more than the basic phase referencing calibration. The SC data shown, for example, would allow phase referencing at 43 GHz (Band 5), but not 86 GHz (Band 6). To improve the situation for a single antenna, significantly faster switching would be required, which is not likely to be possible. For paired antenna calibration, minimizing the separations between the antennas at at station could help significantly. Of course, any such efforts approach a limit as the fluctuations caused by the calibrator to target offset come to dominate over the interpolations. In the marginal cases, calibration of gradients using multiple antennas is a promising method for reducing the fast residual phase fluctuations. This requires stations that have multiple antennas, which is the plan for ngVLA LONG. These methods will ultimately be limited by the non-linearity of the gradients and by uncertainties in how to weight the different calibrator data thanks to the variable separation of the lines-of-sight with distance along the line-of-sight and the likely variable altitude dependence of the tropospheric effects. At the very least, the different antenna-source combinations provide multiple measures of the calibration phase and so should beat down the fluctuations.

Cases 6 and 7: Four Antenna Gradient Correction Using 3 Calibrators and one Target.

In principle, a gradient can be determined by the measurement of phase in three directions through observations of three calibrators. Then the phase to use on a target is a weighted sum of those phases where the weights are dependent on the geometry of the observation. The RMS fluctuations of the resulting sum can be estimated by a weighted quadrature sum of the RMS fluctuations expected at the target position from calibration by each of the calibrators.

Cases 6 and 7 show the results of such a calculation assuming 3 calibrators are being observed. The two cases are distinguished by the antenna separations – 100 m for Case 6 and 50 m for Case 7. The weights, just for a useful example, are 0.333 each. The numbers shown are those simply for three separate calibration determinations, one from each calibrator, which the results of Table 5 showed are about what to expect. The RMS fluctuations for each calibrator are those of the paired antenna

calibration of Cases 4 and 5. Given the nature of the calculation, it is not a surprise that the results are about $1/\sqrt{3}$ of the results of Case 4 or 5. To use such data without incurring further increases of fluctuations from time interpolation, you would need 4 antennas so that the target can be observed simultaneously with the calibrators.

Cases 8 and 9: Four Antenna Treated as Interpolation.

As a sanity check, an alternate way to calculate the expected fluctuations would be to consider the 3 calibrator measurements to be equivalent to the adjacent calibrator scans in fast switching with a switching time roughly equal to the wind crossing time for 100 m line-or-sight separations and 2-3º separations on the sky. For this, use the RMS fluctuation values from Table 5 for 10 seconds. Those are Case 8. For 50 m separations, use the 5 second values from Table 5 to get Case 9. The results are very close to those of Cases 6 and 7. That is not surprising given that the Table 5 results for the interpolated data are very close to the quadrature sums of the paired antenna results. This scheme likely depends on the calibrators not being close together compared to the offsets to the target.

Cases 10 and 11: Three Antennas, Weighted Quadrature Sum of Two Antennas:

With 3 antennas, one could simply use a variant on the above technique which is to use a weighed sum of the phases from two calibrators to calibrate the target. This should reduce the residual errors by about $\sqrt{2}$ compared to the simple paired antenna calibration of Cases 4 and 5. The values obtained with 0.5 weight on each calibrator are shown in Cases 10 and 11. This method is only a partial (one dimension) correction for gradients, but is better than no correction. The quality of the calibration will depend on the geometry. A second calibrator very close to the first and both offset by a larger amount from the target is unlikely to help much. On the other hand, when the second calibrator is opposite the target from the first, the full benefit should be obtained. Intermediate cases need study, but a guess would be that there will be a significant benefit as long as the calibrators are separated by more than each of them is separated from the target.

It might be possible to enhance the method significantly using the weighted sum of two calibrators to interpolate to the location where the point in the phase screen that will cross the target crosses the line between the calibrators. Then use a time lag determined somehow, perhaps by occasional 3 calibrator observations, to effectively interpolate to the target. The system calibration observations where all antennas observe the same source might be used to determine the relevant direction and speed of the winds aloft. With very steady wind, this calibration method might work as well as 3 calibrator observations but with the RMS offset for the virtual calibrator scaled up to account for the errors in the interpolation between the other two. I would expect the results to be slightly worse than for 3 real calibrators.

Cases 12 and 13: Three antennas, fast switch one between two gradient calibrators.

Another 3 antenna method is to observe the target and primary calibrator continuously and fast switch between the two gradient calibrators. This method is similar to Cases 6 and 7, except one of the antennas has an additional RMS term, added in quadrature with a weight, for the fast switching. At any given time, one antenna is in that state, so the fact that it is not always the same antenna does not matter for the representative calculation of the predicted RMS fluctuations. Those are shown as cases 12 and 13. Effectively the gradient is calculated for each time based on observations of two real calibrators and one interpolated calibrator. It might be possible to relax the switching cycle by using the paired antenna calibration to improve the interpolation, but there are likely to be problems with correlated variations. Certainly if you use the paired antenna calibration to calculate virtual phases for all three calibrators, you end up reverting to the normal paired antenna calibration of the target as given in Cases

4 and 5, not to an equivalent to the 4 antenna case.

Case 14: Three antennas, switch between three-antenna gradient scans and target scans.

Finally, a method that could be used with 3 antennas would be to switch between two scans, a gradient scan and a target scan. For the gradient scan, all three antennas would observe calibrators and a phase offset between the target position and the primary calibrator position could be determined with RMS fluctuations something like Cases 6 and 7. That offset would then be used to bias paired antenna calibration during the target scan. During the target scan, paired antenna calibration would be used to smooth out the fluctuations as in Cases 4 and 5. The RMS fluctuations of interest for coherence are those of the paired antenna calibration. The gradient calculation is only appearing as one point per gradient scan with some linear interpolation between gradient scans, so it doesn't influence the short term fluctuations and the method, for coherence improvement, is no better than paired antenna calibration. Where it is likely to be useful is for when other types of gradient calibration are not being done and there is a desire to remove the long term gradients that must be removed for the best astrometry or even best imaging. For that, it will likely be best to have scans of order minutes rather than seconds long.

Gradient Summary

As discussed earlier, any 3 or 4 station method that involves determining weights to use to determine the gradient-corrected offset at the target will run into interesting issues related to the differing behavior with altitude of the portion of the line-of-sight separations due to source separation and the portion due to antenna separation. They will also have to deal with the effect that the instrumental calibration scans, where all antennas look at the same source, may take out the systematic tropospheric or model gradients due to the antenna separations in a manner not easily isolated from the pure instrumental offsets. Disentangling this is of serious interest for astrometry, but not so much for coherence improvement. This is an interesting issue that I won't address here other than to point it out.

Looking over the results in Table 6, it is clear that the best case involves four antennas and very close spacing of antennas. The best cases with 3 antennas are to either just to observe 2 calibrators and take out one dimension of the gradient or to fast switch two of the gradient calibrators. I suspect that using 2 calibrators and attempting to deal with the other dimension of the gradient with a time lag would give somewhat similar or maybe better results depending on how steady the wind is. Comparing the 4 antenna case with the 3 antenna case with fast switching 2 calibrators, the extra antenna provides a 10 to 20% improvement in the RMS fluctuations at SC. At the good sites, this is only about 1° out of an already low value of 4 to 6. All of this depends on there not being significant, fast-changing gradients that do not show up in the single source/single antenna time series, a possibility that will be considered, as previously mentioned, in a future memo(s). Meanwhile separating the antennas by only 50 m instead of 100 m provides a 20-40% improvement. For the BW098A results at SC, pushing to 50 m separations makes the difference between being able to phase reference at 86 GHz and not being able to do so with any method involving only one frequency band.

Source Frequency Phase Referencing

There is another method that can be used to extend coherence at high frequencies. That is source-frequency phase referencing (SFPR: Rioja et al. 2020). Ideally the high band and one or more lower bands would be observed simultaneously and the results at the lower band would be multiplied up to correct the higher band. To avoid problems with 2π ambiguities, the ratio of frequencies of the bands

should be an integer, although on the ngVLA, with its very wide bandwidths, it may be possible to resolve the ambiguities using delay. For the ngVLA, SFPR is most likely to be of interest to assist calibration at Band 6 (70 to 116 GHz). The tropospheric phase fluctuations calculated in this memo for various calibration methods can simply be scaled to the target source frequency to estimate the RMS phase fluctuations at, for example, 86 GHz. The estimates should be independent of the frequency at which the calibrators are observed other than the scaling by frequency – ignoring the contribution of noise and as long as the calibration frequency is high enough so that the ionosphere doesn't contribute significantly. It should be possible to observe the calibrators at the lower frequency and the target at only the higher frequency. The only time that multiple frequencies would need to actually be observed on the target is when core shifts or some other frequency effect is of scientific interest or when calibrating the instrumental offsets. The SFPR technique would mainly be of interest in situations when, for any of a number of reasons, the phase referencing can be done better, or even at all, at a lower band, like 43 GHz, rather than at the higher band, or when there are scientific reasons to want both bands. One possibility is that observations at both bands would be desired for the best astrometry if trying to remove ionospheric effects but that should not be an issue for coherence.

In the cases where the relative phase of the target at different bands is of interest, there is a significant contribution to the RMS of the relative fluctuations from the observing styles required on the ngVLA. The best scheme is to have coaxial feeds, such as are on the KVN, so that the same source can be observed simultaneously at different bands. The ngVLA optics do not allow that. Therefore either fast switching or paired antennas must be used and the offsets in time for fast switching or the offsets in space with paired antennas introduce fluctuations essentially the same as those discussed above for those techniques in phase referencing.

First consider the situation when band switching is used, as it must be on all but the LONG stations of the ngVLA. The ngVLA band switch time for the high frequencies is likely to be near 5 seconds so it is not much faster than the position switch time of 7 seconds and will introduce similar RMS relative phase fluctuations. But, when observing the same source there is no offset source position so the numbers for Case 3 of Table 6 can be used, for a 30 second total cycle time. For the best antennas, it will work reasonably well. At SC, the RMS with a 30 second cycle time is shown as 28° which scales to 56° at 86 GHz.. That is right at the limit of one radian for "successful" phase referencing, leaving little room for other sources of fluctuations. The somewhat faster switching (5 seconds vs 7 for source switching) could allow a shorter cycle. Plus it might be possible for the hardware to provide faster band switching. For example, if the switching time could be reduced to 3 seconds and the scan times reduced to 3 seconds per band for a 50% duty cycle, the full cycle would be 12 seconds and, for the SC BW115D case (6 second case in Table 5), the RMS fluctuations would come down to 15° at 43 GHz which translates to 30° at 86 GHz. That is a 13% loss of coherence which, while not great, is usable. That is something of a best case, so this may work for 86 GHz some of the time, but it will be marginal.

On ngVLA LONG, paired antenna calibration, where two antennas look at different bands, is available. Here the separate lines-of-sight due to antenna separation is the source of relative fluctuations. Plus there is no way to make gains similar those obtained from a faster cycle for fast switching. Also having the same target at the different bands doesn't help like it does for fast switching – the different lines-of-sight are based on the different antennas, not different sources. Therefore the numbers of Cases 4 and 5 are appropriate pretty much as is other than the frequency scaling. In the best 2 antenna situation at SC, which is BW115D with 50 m spacings, the fluctuations are 21° at 43 GHz and 43° at 86 GHz, the latter causing a 24% loss of coherence. If the third antenna (or third and fourth) is not needed for phase

referencing to other sources, they could be used in bring down the RMS fluctuations by doubling up on one or both bands in analogy with the gradient measuring schemes for calibrator(s)-target phase referencing. From Table 5 the hypothetical switching time for fast switching that gives the same RMS as the paired antenna case for BW115D, 50 m (21°) is 10 seconds for the half cycle or 20 seconds for the full cycle. That could be 5 second switches and 5 second scans which matches the current best guess switch time. So the interesting implication is, if measuring the relative phase on a single source, the results will be the same or better using fast switching as simple paired antennas. Since fast band switching doesn't require extra antennas, it could be done in combination with the multi-antenna gradient based phase referencing schemes. These results also suggest we should push for especially fast band switching, at least between the top bands.

LONG Station Antenna Separation

The measurements here make clear that the quality of the paired antenna calibration will be a strong function of the antenna separation at the stations. Looking at SC in Table 1, if the antennas are 100m apart and that ends up being about the separation of the lines-of-sight at the scale height of the troposphere (lots of room for geometry variations here), then the RMS phase difference between sources with 10 m/s wind will be on the order of 54° for a station like SC in BW098A. If instead, the antennas were within 50 m of each other, the RMS phase difference would be about 33°. Those numbers give a coherence loss of 36% for 100 m and 15% for 50 m. That would make a big difference. For the better date of BW115D, the numbers are 33° for 100 m and 21° for 50 m. That is the difference between 15% and 7% loss of coherence at 43 GHz. Scaling to 86 GHz, for 100 m separation the RMS phase will be over a radian which is generally considered to not be working. For 50 m, it would be 42° for a coherence loss of about 25% - not nice but usable. Thus there are days when paired antennas calibration could work at 86 GHz at SC with 50 m separation, but likely not with 100 m separation.

The flip side of having the antennas that close is that there will be shadowing, although only for one antenna at a time if station internal baselines are not parallel. For 50 m separation, shadowing will occur at up to 21° elevation maximum. At a poor site, observing at low elevations is somewhat problematic in any case, plus only certain directions are affected. So I maintain that, for a poor site, the NGVLA should place the antennas very close to each other and recommend 50 m or less. It might even be worth going as short as 38 m, the minimum allowed by the antennas, which can risk shadowing up to 28° elevation for one antenna at a time. Better sites could be more relaxed although there is no point in going over 86 m because the antennas cannot point low enough to be shadowed beyond that. Something like 70 m with the highest elevation shadowing at 15° might be a good compromise for the good stations.

CONCLUSIONS

The RMS relative phase fluctuations of 43 GHz VLBA data agree well, for sites like the VLA, with the predictions based on structure functions of the atmosphere over the VLA presented in C15 and CH if a nominal winds aloft speed of 10 m/s is assumed to translate between time and spatial separation. Therefore the conclusions of C15 that phase referencing with fast switching will work applies to the ngVLA LONG, at least for most stations in decent weather.

For the two projects studied, which were dynamically scheduled with high priority in winter, about half or a bit more, of the antennas had really good conditions. One antenna, SC, had poor conditions both times although they differed significantly. This is no surprise for this sea level, tropical site. Another site, OV, was one of the best stations in BW098A but had conditions significantly worse than SC during BW115D despite no indications of adverse conditions in the logs or monitor data. With an array of diverse locations, the observations will need to be designed to be robust against poor performance of some sites.

A strong conclusion is that the antennas at stations likely to have poor phase stability most of the time should have the antennas placed as close together as possible even if that results in some losses to shadowing. Even at good sites, close spacings can be helpful. The stability gains are large – roughly 20-40% for paired antenna calibration for 50 vs 100 m separation. The shadowing will only affect one antenna at a time as long as none of the lines separating antennas are parallel. Placing antennas close together does not have cost impact or impact on hardware designs already developed. It is not completely clear how much better a site like Arecibo at 1000 ft in Puerto Rico or a site many miles from the ocean in Florida will be than SC, but those sites should either have very close antennas, or have site testing done that is adequate to show that such spacings are not needed. Other low elevation or high water vapor sites should also have closely spaced antennas.

The RMS residual phase fluctuations after phase referencing with fast switching between sources will be larger than for phase referencing using paired antenna calibration with one calibrator and antenna separations of 100 m. The paired antenna scheme gets considerably better for 50 m antenna separations. The fact that there is not an even larger difference between the methods is related to the ability to use interpolation between calibrator scans for fast switching and to the fluctuations introduced by the antenna separation for paired antenna calibration.

An factor that needs to be considered in the detailed scheduling of projects that use paired antenna calibration is the choice of which antenna at each station observes which source. That choice can have a strong impact on the separation of the lines-of-sight and, hence, on the quality of the calibration. This will affect the scheduling software design as the scheduler must know how to allocate the sources to antennas.

Phase referencing schemes that address gradients can provide improved results. This is somewhat analogous to the interpolation used in fast switching. The best results involve 4 antennas at a site so that 3 can be used to continuously measure the gradients with nearly 100% on-source times. The gradient measurements cannot be perfect because such gradients are not fully linear as shown by the RMS scatter of the interpolated points in Table 5. But measuring gradients provides a significant improvement over not doing so, at least for a poor station. At good stations, the fluctuations may be sufficiently small that gradient calibration is not productive for maintaining short term coherence. Calibration of persistent gradients is likely to be critical for the best astrometry and even for imaging if they are large, but the short term phase fluctuation data analyzed here are not sensitive to such gradients. They will be addressed in a future memo(s).

With 3 antennas per site, there are several calibration strategies to deal with the fact that the full 2D gradient cannot be measured continuously while the target is also being observed. Similar results are obtained by using only 2 calibrators and by using one antenna to switch between the two gradient calibrators while the other antennas continuously observe the primary calibrator and the target. It might be possible to improve the 2 calibrator case by utilizing a known (likely measured with calibrator observations) wind speed aloft and direction along with a time lag to get the gradient along the wind

direction and across the line between the calibrator lines-of-sight. In any case, these related techniques allow results 10-20% worse than those provided by 4 antennas.

Another 3 station calibration scheme would be to switch between gradient scans and target scans. For the gradient scans, all 3 antennas would observe different calibrators to obtain the gradients. For the target scans, the gradient would simply be interpolated and used as a bias for the paired antenna calibration. This scheme, as far as short term fluctuations is concerned, is basically the same as the paired antenna case without gradient measurement. But this scheme, perhaps with a more relaxed switching cycle, may well be the method of choice for removing the long term gradients that are important for astrometry and otherwise to deal with model errors.

Given that the gains provided by 4 antennas are only in the 10-20 % range, and the fact that even the 3 antenna schemes allow a station like SC to phase reference at 86 GHz on at least some days, the ability to improve calibration of rapid phase fluctuations may not justify the disruption and cost of changing the base plan of the ngVLA LONG from 3 to 4 antennas per station. However, this conclusion is based on fluctuations seen on a single source. Calibrating the gradients between sources to which the data analyzed are not sensitive, but that can degrade phase referencing, may provide stronger justification for the extra antenna. It will depend on the time scales of such gradients. Such time scales are likely to be relatively long because, since a gradient is seen as a difference in phases at two positions, rapid gradient variations should show up as rapid phase fluctuations, which is what is studied here. The conclusion about 4 vs 3 antennas is also not about the needs for microarcsecond astrometry, which will be the subject of another memo in preparation. It is also not about possible geodetic observing modes that measure parameters of interest (EOP, UT1-UTC, clock) continuously.

For Source Frequency Phase Referencing, it appears, for observations at different bands on the same source, that fast switching works as well or better than paired antenna calibration to relate the phases of different bands on a source. That depends on getting the total cycle time (center of scan at one frequency to center of the next scan at that frequency) under around 20 seconds at a site like SC, which should be possible with the predicted switch time of 5 seconds between the highest frequency bands. Shortening those times even more would be of significant benefit and should be a technical goal. This opens the possibility that all the antennas involved in paired antenna phase referencing could also be band switching. This is especially true since the best multi-antenna phase referencing schemes don't depend strongly on rapid source switching.

For the single antenna stations in the central parts of the ngVLA, any phase referencing will have to involve fast switching, probably with cycle times of about 30 seconds. That is Case 2 of Table 6. The RMS phase fluctuations using that method will be about twice as large as those at the multi-antenna stations. This should still be ok at good sites such as nearly everything in the general vicinity of the array center. But this should be a warning that we need to avoid including likely poor sites in MID. At sites like SC, and to a lesser extent some of the others, the fast switching would not be adequate to allow phase referencing at Band 6 while the multi-antenna schemes will allow such observing.

So the final summary is that phase referencing at 86 GHz should work at most sites and work some of the time at all sites. The antennas at poor sites should be placed as close together as possible even at the cost of some shadowing. Even at good sites, close spacings are beneficial. And, for dealing with coherence in the face of the kinds of rapid phase fluctuations on a single source studied here, the gains provided by 4 antennas per station, while significant, are may not be adequate to justify the cost and disruption that changing from the established 3 antenna scheme. However the data this conclusion is based on are not sensitive to systematic gradients between sources. Dealing with such gradients may

yet justify 4 antennas, but that is a topic for another memo based on different data.

ACKNOWLEDGEMENTS

I would like to thank Joan Wrobel, Mark Reid, Richard Dodson and Maria Rioja for very useful comments and discussion. The non-linear least squares regressions used to obtain the structure function parameters is provided by John C. Pezzullo on the interactive web page https://statpages.info/nonlin.html.

REFERENCES

Butler 2002, VLA Test Memo. 232

Carilli, C. and Holdaway, M, 1999, Radio Science, 34, 817 (CH).

Carilli, C., 2015, ngVLA Memo 1 (C15).

Reid, M. 2022, PASP, 134, 123001.

Rioja, M. and Dodson, R. 2020: The Astronomy and Astrophysics Review. Also arXiv:2010.02156v1

Selina, R., 2022, NRAO Doc. #: 020.10.25.00.00-0005-REP "System Concept Options and Trade Offs"

Taylor, G. I. 1938, Proc. Roy. Soc. 164A, 476.

Thompson, A. R., Moran, J. M., and Swenson, G. W., 2001, Interferometry and Synthesis in Radio Astronomy, v2, John Wiley & Sons, Inc. (TMS; More recent versions exist).

Walker, R. C., Hardee, P. E., Davies, F. B., Chun, L.., and Junor, W. 2018, ApJ, 855, 128.