ngVLA Memo #91
Status of the VLA CWVRs

Bryan J. Butler
NRAO
Version 2; July 23, 2021

Abstract

Water vapor radiometry (WVR) has been chosen as the baseline phase (delay) calibration technique for ngVLA (Hales 2019). Five Compact WVR (CWVR) units were built as part of the ngVLA development program and installed on VLA antennas in 2019 (I will refer to them herein as nCWVR units). This memo will discuss the current status of those five units, and present some preliminary data from them. Their performance does not meet the sensitivity requirement, and they do not appear to be accurate (not on a proper temperature scale), so it is not possible to draw any conclusion about their applicability to ngVLA phase calibration at this point in time.

1. Introduction

The use of WVR for visibility phase correction has been described in a number of documents, so I will not cover the background here. See Welch (1999) for a good description of pre-1999 efforts, and for efforts at the VLA prior to those described here, see Gill+ (2017), Chandler+ (2004a,b) and Butler (1999). The nCWVR units are described at some level in Butler (2020) and Towne (2020a). In this memo I will describe the overall characteristics of the nCWVR units, and some data from them. I will note here that there are many details of the nCWVR unit design, implementation, and testing that are mostly lost since the retirement of Nathan Towne. He left memos, notebooks, software, and emails, and some information can be gleaned from them, but his detailed knowledge of the system is difficult to replace.

2. The ngVLA CWVR

The nCWVRs were designed to have similar functionality as the previous generation CWVRs, and built around the same core MMIC spectral filter (Gill+ 2017), but the mechanical and electrical design was modified significantly. In this section, I describe those designs (at the highest level), the monitor and control of the units, testing and acceptance, installation on VLA antennas, and operation.

2.1 Mechanical Design

The mechanical enclosure was designed in two hinged halves, one housing the IF section and related electronics (including the temperature controller and 'Cal' board, which controls power and noise diodes), the other housing the cold plate, with actual noise diodes, the MMIC splitter, amplifiers, and VFC converters. Figure 1 shows these two halves, and the entire hinged enclosure on a lab bench.

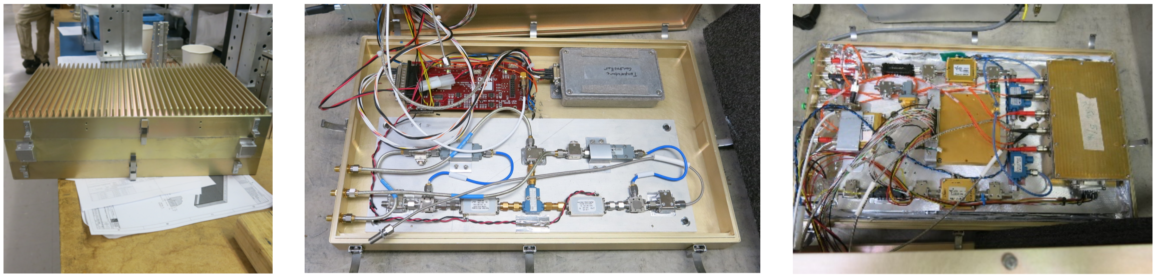

Figure 1. The mechanical CWVR enclosure. Left: the entire hinged enclosure, closed. Middle: The IF section. Right: The cold plate section.

The entire enclosure is bolted to the side of the K-band dewar in the vertex room of a VLA antenna. Figure 2 shows a unit bolted in that state in an antenna.

Figure 2. An ngVLA CWVR unit bolted to a K-band dewar in a VLA antenna.

2.2 Electrical Design

A detailed description of the electrical design of the nCWVRs is beyond the scope of this memo, but I will describe in general how it fits into the signal path, and some of the internal electronics here.

The nCWVR taps directly into the RF of the VLA K-band receiver, with a splitter, as shown in Figure 3.

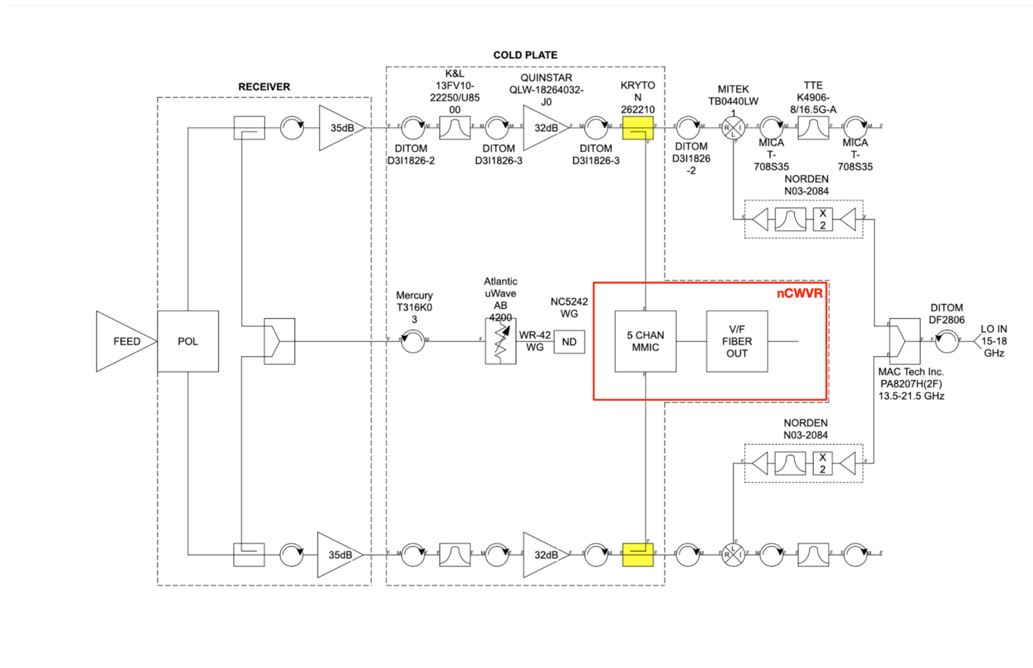

Figure 3. Block diagram of how the nCWVR fits in to the VLA K-band signal path.

The internal nCWVR electronics block diagram is shown in Figure 4, taken from Towne 2020b.

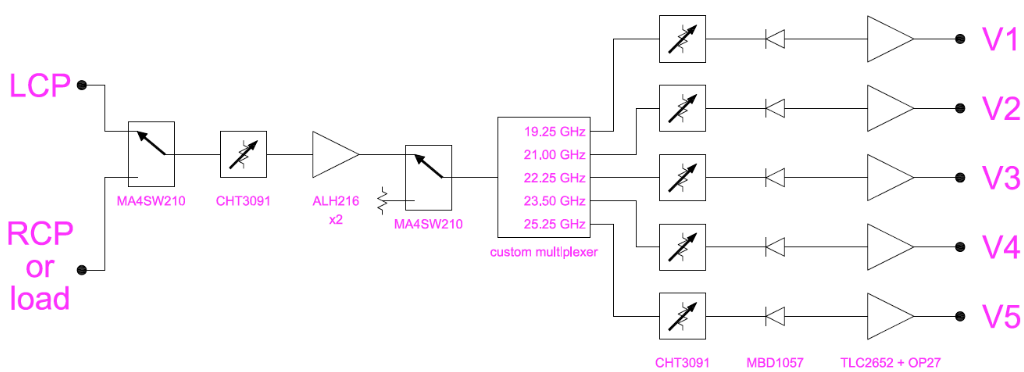

Figure 4. The block diagram of the electronic signal path for the nCWVR (from Towne 2020b).

The voltage counts for the five channels of the nCWVR (see Butler 2020 for definition of frequencies) are the V1-V5 outputs in Figure 4.

2.3 Monitor and Control

The monitor and control (M&C) for the nCWVR is via the VLA F317-1 and F318 MIBs. A detailed description of nCWVR M&C is in Towne (2020b), but I will note the most important parts here. The F317-1 provides control of the front end, including switched-power switching, and readout of the CWVR board temperature as well as the F317-1 module physical temperature. The F318 provides all of the rest of the M&C for the nCWVR. Notably, attenuation on or off and levels for each channel (in order to get count rates in the right range), polarization choice, on or off for the whole unit, and the V1-V5 counts are all accessed via the F318. Specific M&C point names are listed in Towne (2020b), but I note here the monitor point name of the measured counts.

The F318 records the V1-V5 counts (see Figure 4) both as "instantaneous" (measured over 0.0425 seconds) and "accumulated" (averaged over 1 second) values. It records these for both the 'lo' (T_{cal} off) and 'hi' (T_{cal} on) states. These are multicast out over the VLA astream, and hence stored in the M&C archive database. For SNR reasons, it is always better to use the accumulated values. These values are multicast and stored with the host/dev/mon name: eaXX-f318.WVR.ChN_Accum_YY, where XX is the antenna number, N is 1-5, and YY is either 'hi' or 'lo'. So, for example, to retrieve a value for the T_{cal} on state, for channel 3, for antenna 7, you would use ea07-f318.WVR.Ch3_Accum_hi.

2.4 Testing and Acceptance

Testing and acceptance for each of the five nCWVR units consisted of four main activities: 1 – basic signal path integrity; 2 – check of MMIC frequency response (all five channels); 3 – determination of power count rates-to-power linearization coefficients; 4 – determination of T_{cal} values. This was all done by Nathan in the lab. The results for item 2 were emailed out as internal reports; those from items 3 and 4 were combined into Towne (2020b).

2.5 Installation on VLA Antennas

The first four completed units were installed on VLA antennas as listed in Table 1. The fifth unit was completed and tested (and calibrated), but has not been put in an antenna.

Table 1. nCWVR installation on VLA antennas.

Serial #	Antenna	Date
001	ea07	2019-Jul-24
002	ea10	2019-Jun-19
003	ea12	2019-Sep-19
004	ea25	2019-Nov-13
005	lab	n/a

2.6 nCWVR Operation

In order to operate the nCWVRs, one must:

1. Make sure that the unit is powered on.
2. Make sure that switching is turned on.
3. Make sure the desired polarization is selected (LCP was found to be better behaved).
4. Make sure the attenuators are set to the right level, to get count rates in the valid range (as noted in Towne 2020b).

Nathan provided a suite of software (all in Python) to do these various things, as described in detail in Towne (2020b) and I have copied and modified some of that software. Note that Nathan preferred to run the rates at particular values, different for each unit on each antenna (item 4 above), but my own experiments showed that as long as the extrema of the valid ranges were avoided the performance of the units was not affected. Figure 5 shows the T_{sys} and T_{obs} (see below for definition of these quantities) RMS for all of the nCWVR units during a test (on April 23, 2020) where I ran the count rates across their valid ranges (and a bit beyond). There is no discernable trend of RMS with count rate, meaning the count rates can be set anywhere in their valid ranges (ea07 – 0.2-2.5 MHz; ea10 – 0.12-1.2 MHz; ea12 – 0.2-2.5 MHz; ea25 – 0.4-4.0 MHz).

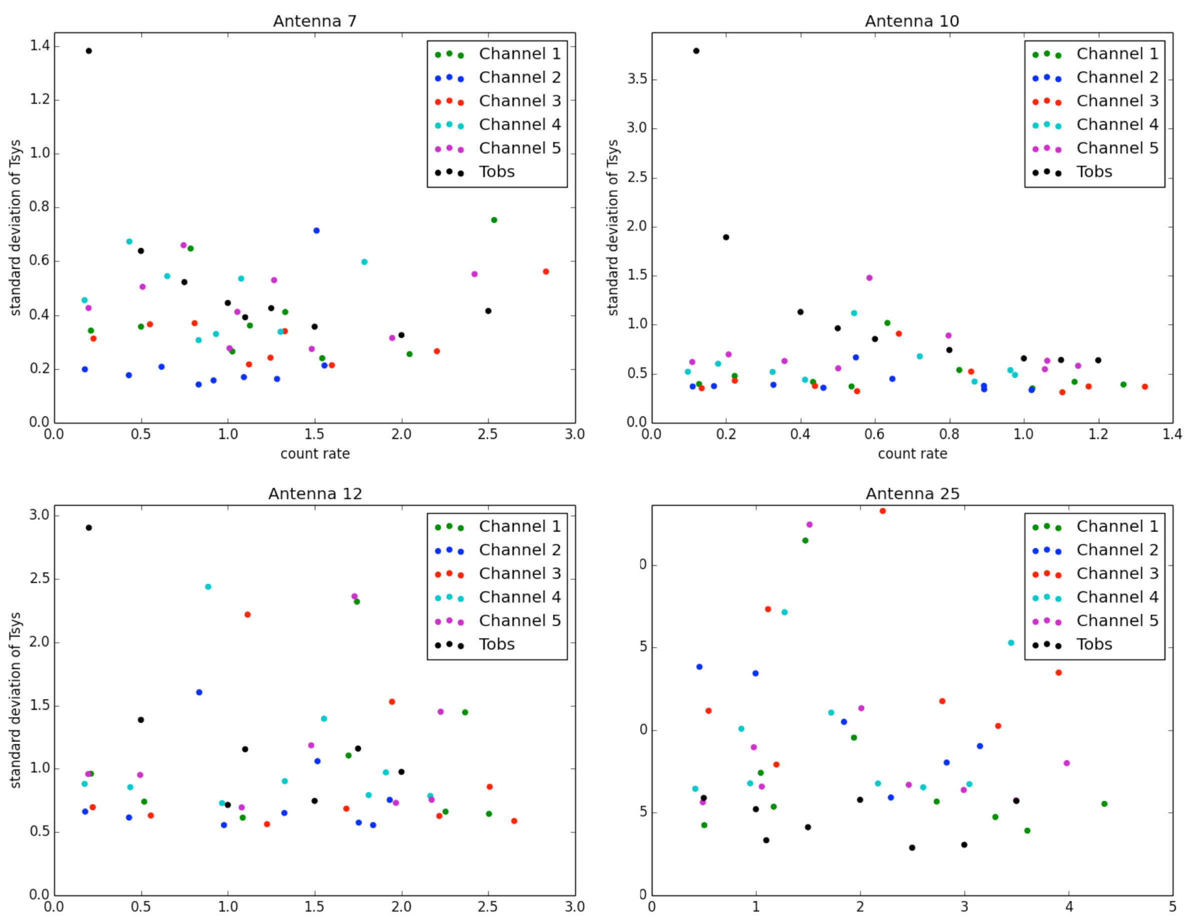

Figure 5. T_{sys} and T_{obs} RMS as a function of count rate for the four nCWVR units, from a test on April 23, 2020.

3. Turning nCWVR Measurements into T_{sys} Spectra

Given a measure of both the T_{cal} on (hi) and T_{cal} off (lo) counts of an nCWVR channel, there are then three steps to turn those into a measurement of T_{sys} (in K):

1. Divide by the sample interval (0.0425 sec) to turn the count into a count rate, then divide by 10^6 to get that rate in MHz.
2. Convert the count rate to power (in nW) using the linearization conversion coefficients measured for each unit (for the specific polarization).
3. Turn that into a T_{sys} measurement via:

$$T_{sys} = T_{cal} \frac{P_{lo}}{P_{hi} - P_{lo}}$$

Where the T_{cal} is that measured in the lab for the appropriate unit, channel, and polarization during the acceptance testing (Towne 2020b).

Note that this is only a first order conversion. Towne (2020b) notes that the 0.0425 value is slightly different for each unit, and for the lo and hi states (see Table 3 from that document). The differences are very small, however. A bigger issue, however, is that the lo and hi states are coupled due to response times in the analog electronics of the nCWVRs (Towne 2019a,b). The main effect of this is an increase in the effective noise of the values for T_{sys}.

4. Turning T_{sys} Spectra into Visibility Phase Corrections

Once T_{sys} are obtained on two antennas, they should be able to be used to make a correction to the visibility phase on the baseline between those two antennas as a function of time. There are several ways to do this (see the discussion in Butler 2020); here I will only use the traditional way – creation of an "observable" (T_{obs}) on each antenna as a weighted sum of the T_{sys} on the five channels, and then using the difference between those values on the two antennas as a measure of phase variation (hopefully correction).

In the end, we expect that we will have to determine a "correction factor" during calibrator observations, for visibility phase correction (see the discussion in Chandler+ 2004); determination of that factor is beyond the scope of this memo.

5. Some Issues with Existing Data

5.1 Synchronization of Measured Values

It is important that the measured accumulated values (which are one per second) are measured at the same time on all the antennas. This was established during discussions in early May 2020 and implemented in the F318 software on May 26, 2020 – prior to that the measurements were not synchronized. However, from investigating the values in the database since then, it is clear that at times this synchronization goes awry; one or more antennas are not synchronized as they should be (sometimes none of them are synchronized with each other). The periods when they were all synchronized are: 2020-May-26 to 2020-Jun-22, and 2021-Feb-03 to 2021-May-27. It appears that the disruption of the synchronization might occur when there is a loss of the interrupt signal in the F318. Not necessarily a reboot of one or more of the F318s (because on reboot they are supposed to get re-synchronized), but loss of the interrupt signal without a reboot, perhaps caused by a power glitch. But the true cause of this is not known yet, and is under investigation.

This does not prevent analysis, and in any case the values can be smoothed in time so that the offsets (which are a few tenths of seconds) are a small fraction of the integration time, but in the end the corrections do need to be done on ~1 second timescales to get the accuracy needed (Hales 2019), so for any future WVR system this needs to be a carefully considered.

5.2 Database Gaps

The values noted in section 2.3 above are stored in the M&C archive database, which can be accessed either directly via code, or through the M&C query interface. At this time, the production and standby databases go back to March 1, 2020. However, there are gaps in the data – notably from April 1, 2020 to June 1, 2020 – in both of those databases. The data exists in the Postgres backups on disk, and I have made code to access the data from this long gap specifically. The cause of these gaps is under investigation.

5.3 Aberrant Values

There are times when the nCWVR data makes no sense. At times, the count rate is outside the valid range; at times even with valid count rates, the data just makes no sense. Values of T_{sys} can be negative (meaning the lo counts are higher than the hi counts), or unphysically high (hundreds of K – note that even if the atmosphere were entirely opaque at K-band, which it never is at the VLA site, the T_{sys} value should not be higher than a bit more than 300 K or so). Or one channel has values that are far different than the others.

6. Early nCWVR Results

Some very early nCWVR results are shown in Towne (2020b). Here I investigate further some characteristics of measured nCWVR data over the past year. First, though, where were the antennas with nCWVR units for that period?

6.1 Locations of Antennas with nCWVR Units

Table 2 shows the pad locations for the four antennas with nCWVR units from 2020-Jan-29 to present. The values for L_x, L_y, L_z, and R are taken from the Green Book (Hjellming 1992). Antenna ea25 went into the AAB on 2021-Jun-07 and is still there; when it comes out it is planned to go to pad N4. Figure 6 shows these pads graphically, demonstrating that we have had short and long baselines in both the E-W and N-S directions. Table 3 and Figure 7 show the position of all four nCWVR antennas as a function of time for the past 18 months.

Table 2. Pad locations for the four antennas with nCWVR units since 2020-Jan-29.

Antenna	Begin	End	Pad	L_x (ns)	L_y (ns)	L_z (ns)	R (m)
ea07	2020-Jan-29	2020-Dec-08	E12	765.39	2933.01	-1133.62	970.21
	2020-Dec-08	2021-Mar-08	E48	8324.92	31661.66	-12190.73	10472.93
	2021-Mar-08	2021-Jun-01	E9	465.79	1790.89	-692.95	592.38
	2021-Jun-01	Present	E18	1548.02	5883.17	-2264.55	1946.03
ea10	2020-Jan-29	Present	E8	381.68	1463.33	-565.35	484.02
ea12	2020-Jan-29	2020-Jun-11	N1	2.24	0.05	1.71	0.84
	2020-Jun-11	2021-Mar-16	N24	-5538.93	-865.16	8187.02	2974.68
	2021-Mar-16	2021-Jun-03	N3	-174.91	-27.56	262.39	94.90
	2021-Jun-03	Present	N16	-2673.19	-416.88	3943.10	1433.62
ea25	2020-Jan-29	2021-Jun-07	N8	-812.58	-126.88	1200.98	436.37
	2021-Jun-07	Present	AAB	N/A	N/A	N/A	N/A

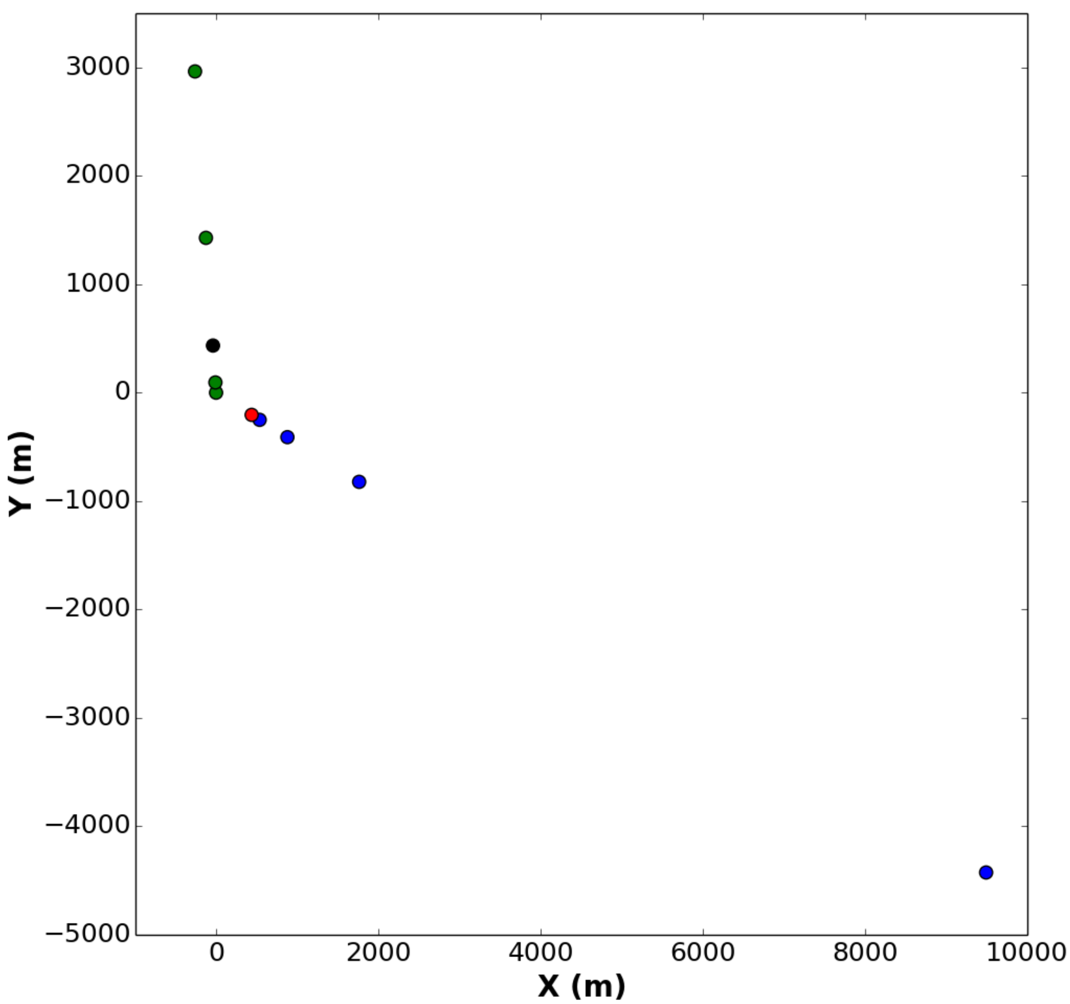

Figure 6. Locations of antennas with nCWVR units over the period 2020-Jan-29 to present. Blue dots are ea07, red is ea10, green are ea12, and black is ea25 (in the AAB as of 2021-Jun-07).

Table 3. Positions of the antennas with nCWVR units as a function of time for the past 18 months.

Begin Date	End Date	ea07	ea10	ea12	ea25
2020-Jan-29	2020-Jun-11	E12	E8	N1	N8
2020-Jun-11	2020-Dec-08	E12	E8	N24	N8
2020-Dec-08	2021-Mar-08	E48	E8	N24	N8
2021-Mar-16	2021-Jun-01	E9	E8	N3	N8
2021-Jun-06	Present	E18	E8	N16	AAB

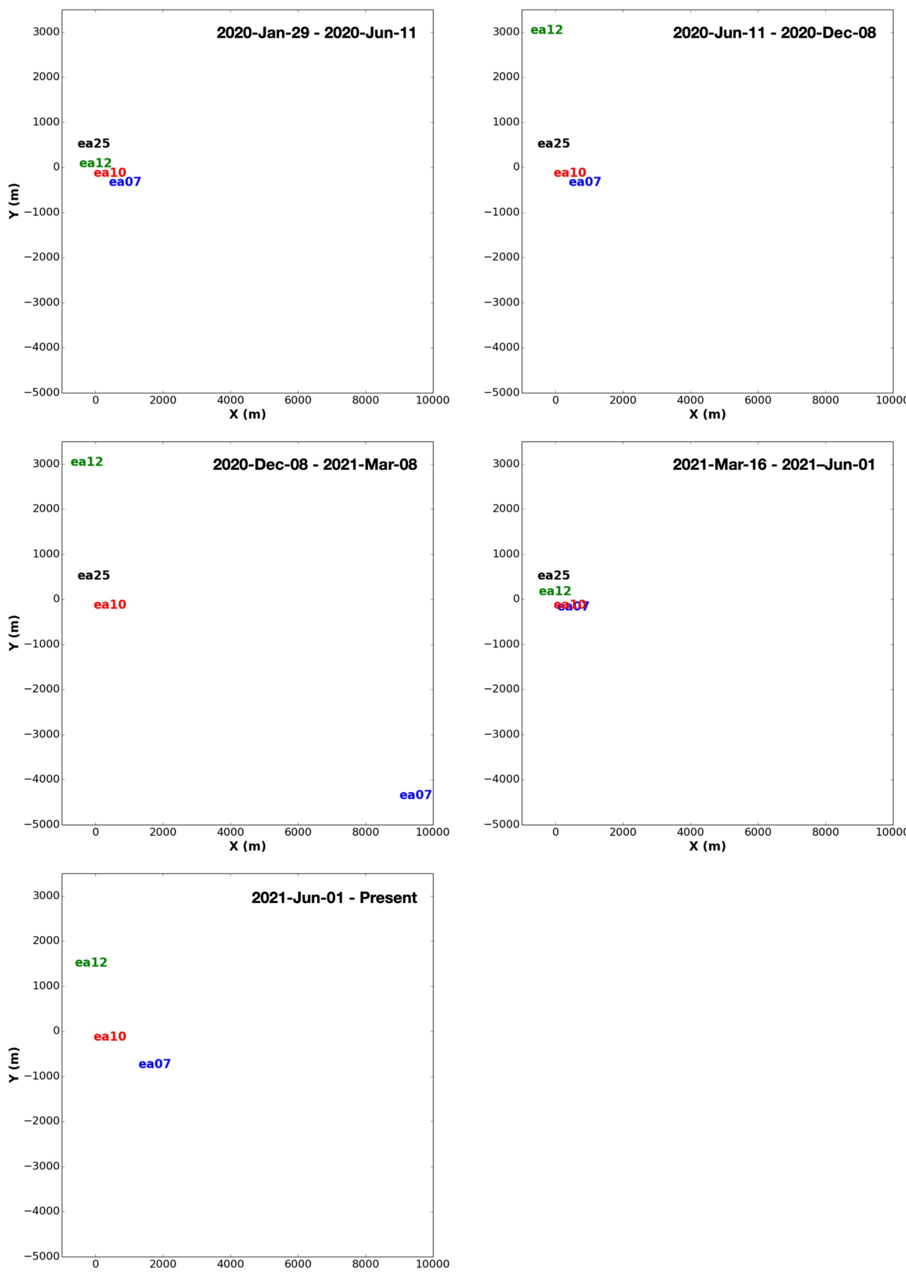

Figure 7. Positions of the antennas with nCWVR units as a function of time for the past 18 months.

6.2 Spectra and Possible Errors in Measured T_{cal}

For WVR phase correction to work, the T_{sys} measurements across antennas and channels should be on an absolute temperature scale. This is accomplished by accurately measuring the values of the T_{cal} for each unit and channel, and measuring the on and off (hi and lo) count values (see the equation in section 3 above). It also depends on the counts being accurately measured, of course. In some sense, if we allow for arbitrary scaling factors when forming the difference quantities between antennas (which we may in the end), an overall scale uncertainty may not be critical, though separation of atmospheric from other components of T_{sys} (like spillover and receiver temperature) will not be possible in that situation. However, for a particular antenna (nCWVR unit), the T_{sys} values across the five channels should be on the same absolute scale. Without that property, it is not possible to do precise WVR phase correction. So, at any point in time, the spectrum of the T_{sys} values on a particular antenna (nCWVR unit) should be a true spectrum, representing the shape of the water line (along with other contributions to T_{sys}).

I have found, unfortunately, that this is not the case. In fact, none of the units produce spectra that seem to be particularly well calibrated, channel-to-channel, and they have not since they were first installed. Figure 8 shows a plot of the spectrum for each of the four nCWVR units, every three months, on the first day of the month, since March 1, 2020. These spectra were formed by taking the median T_{sys} value for each channel for the entire day. Shorter time periods within each day were examined, and though the values fluctuate slightly, the overall spectral shape remains roughly the same. Other days were examined, and similar results were found – taking the first day of every third month is just meant to give a sampling of the spectral shapes.

There are some consistent things – like the units on ea10 and ea25 seem to produce spectra closest to reality; ea12 channel 1 is often clearly higher than it should be (but changes over time); etc. It is difficult to know whether this is a problem with the measured T_{cal} values for the units and channels, or to know whether it is a more fundamental problem with the measurement itself. If it were just T_{cal}, it would seem that the spectral shapes would be more consistent over time – the fact that this is not quite true probably means that it is some combination of the two factors. More investigation is needed.

It may actually be possible to calibrate the T_{cal} values separately. Note that we have found in the past for the normal T_{cal} values that what is measured in the lab is often not what is seen once a receiver is installed on an antenna. This is true for both pre- and post-upgrade. I believe that a careful comparison of measured nCWVR spectra (as shown in Figure 8) with one or all three of the following quantities could be used to perform this calibration: 1 – a theoretical atmospheric spectrum, given surface conditions, augmented with information on other contributions to T_{sys}; 2 – an estimate of the T_{sys} spectrum from the switched power measurement of the regular VLA system, using an observation of K-band continuum (perhaps a median across all other antennas, to average across the other T_{sys} contribution factors); a measure of the T_{sys} spectrum from the radiometer, augmented with information on other contributions to T_{sys}.

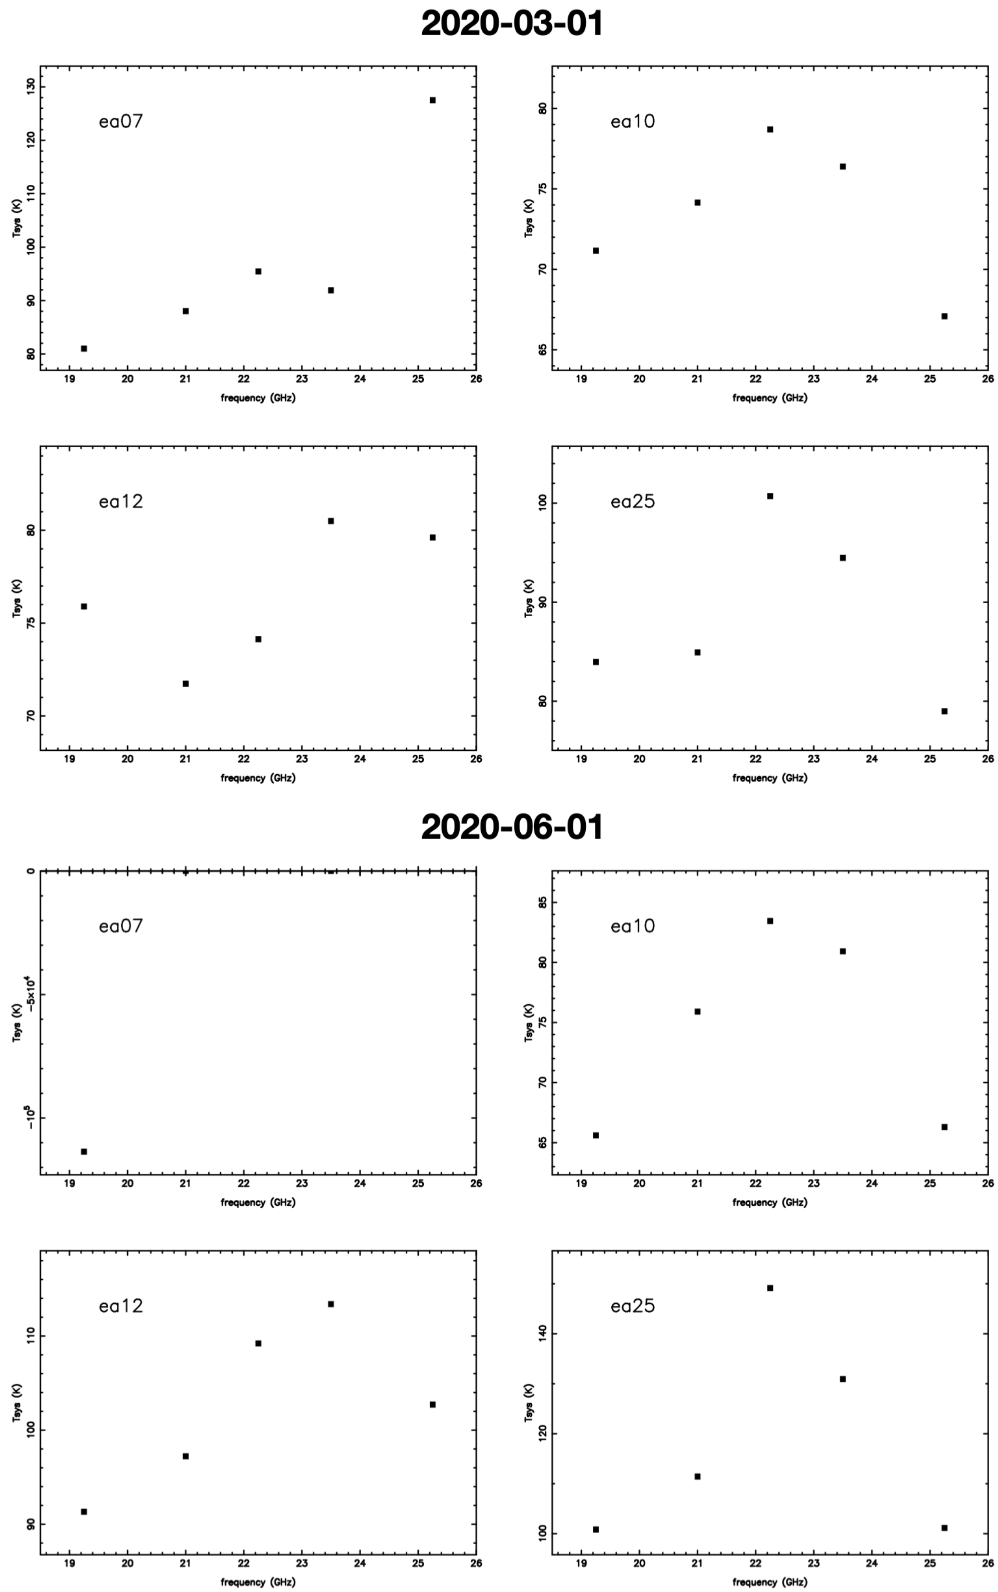

Figure 8. T_{sys} spectra for the four nCWVR units on the first day of the month, every third month, since March 2020.

2020-09-01

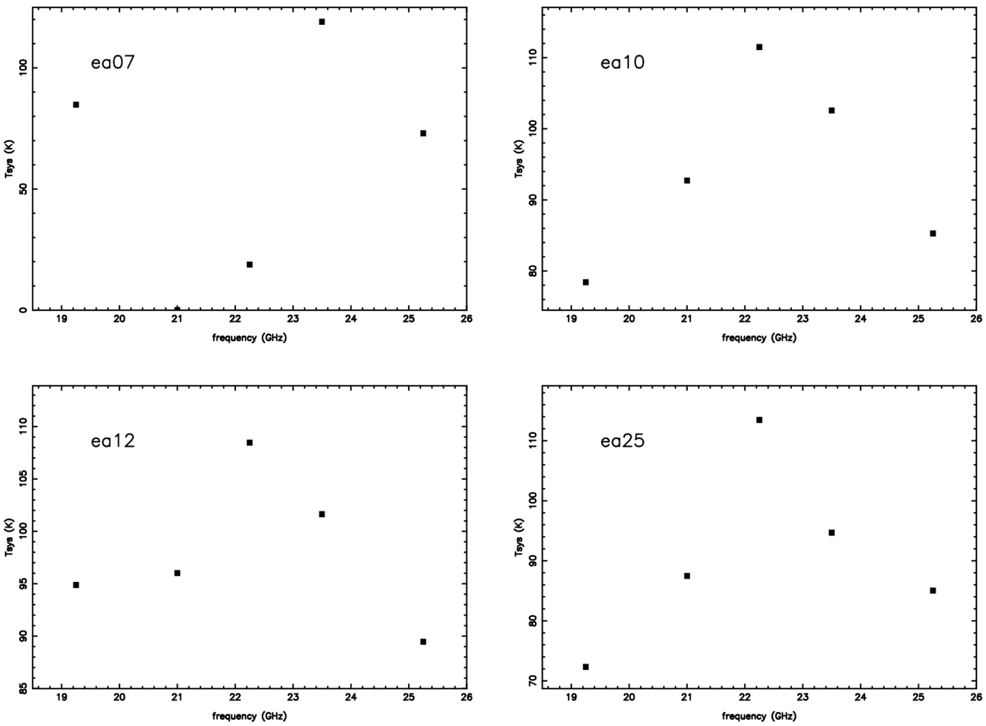

2020-12-01

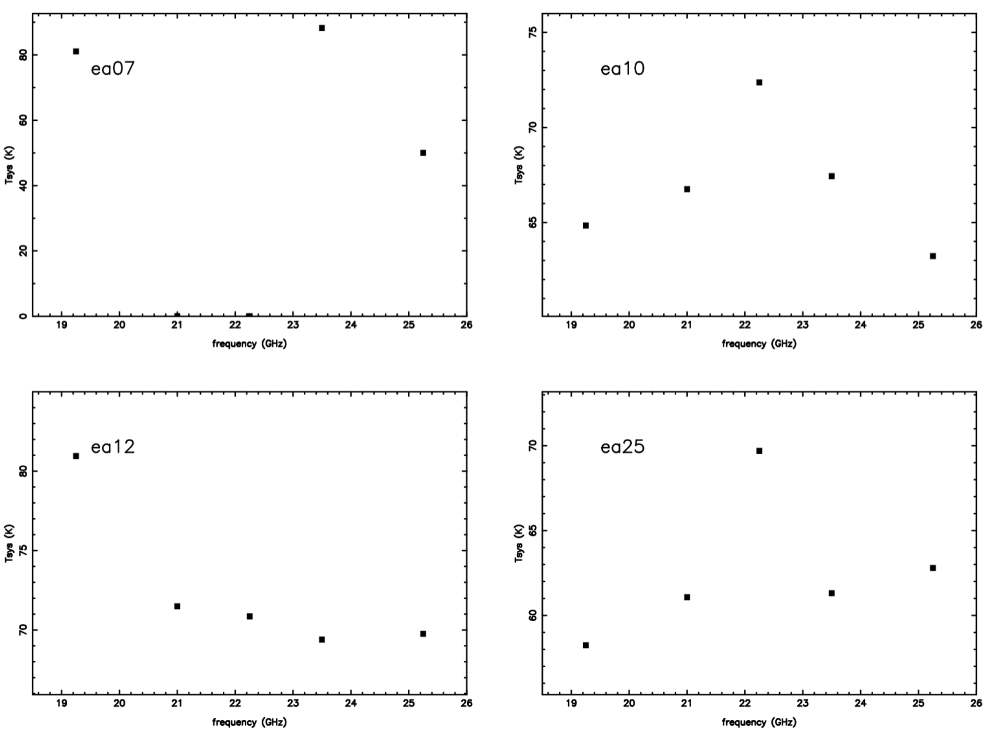

Figure 8. Continued...

2021-03-01

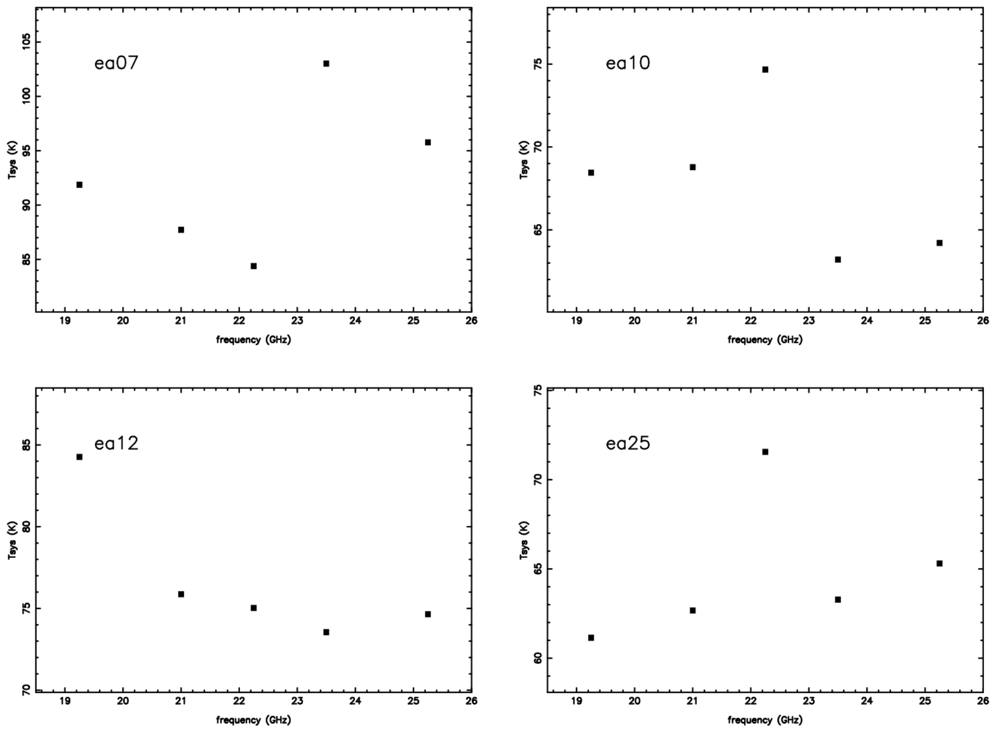

2021-06-01

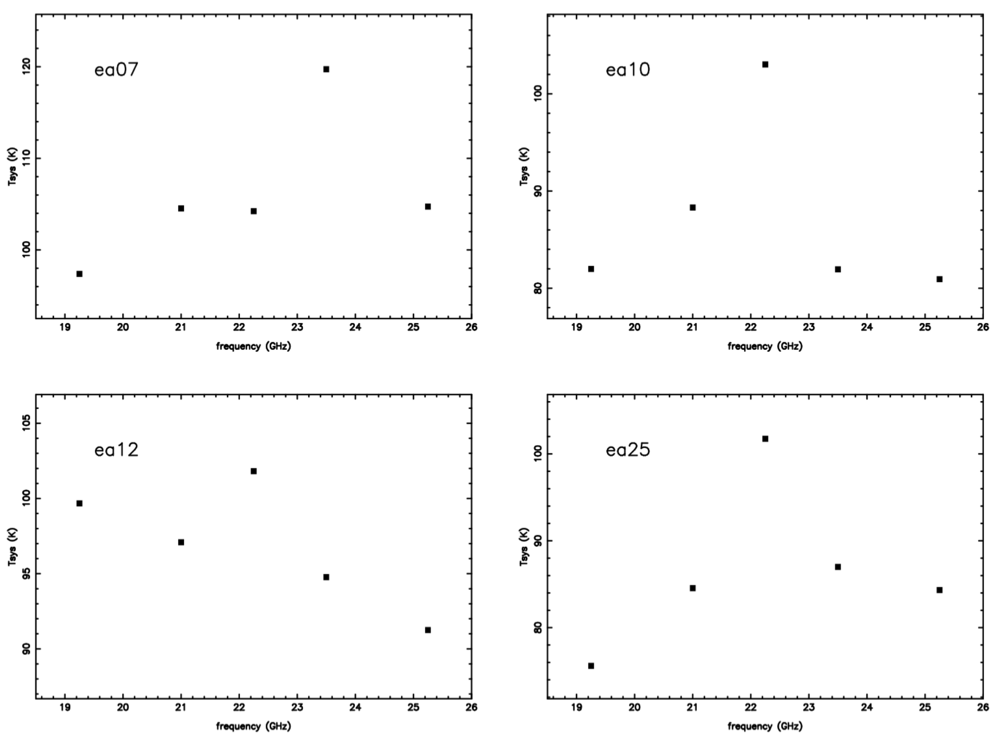

Figure 8. Continued...

6.3 T_{obs} Data and Sensitivity

Despite the fact that the measurements might not be on a good absolute scale (see section 6.2 above), it still may be possible to create some difference quantity between pairs of antennas which would track phase, similar to what was found in Towne (2020a). I performed three tests on May 28 and June 2, 2020 to attempt this; one at LST 04h, one at 14h, and one at 20h, to potentially get different observing conditions. Figure 9 shows the median channel T_{sys} values (similar to figure 8) during one of these tests (the first on May 28) – you can see that the nCWVR units on ea10 and ea25 are at least close to the expected shape for the water vapor line.

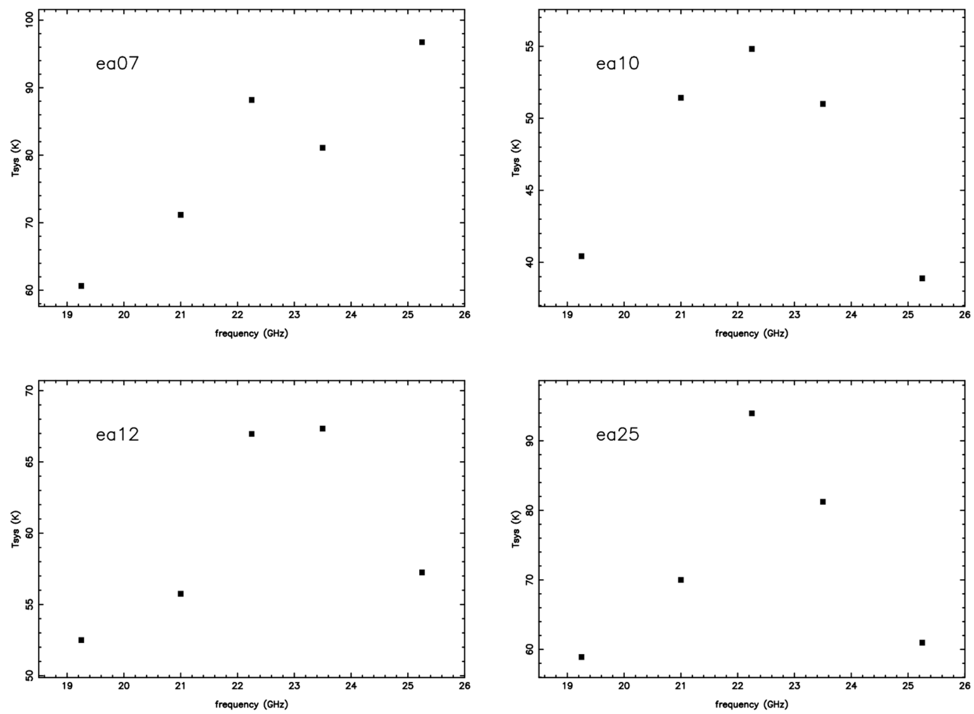

Figure 9. Median T_{sys} spectra for the four nCWVR units during a test on May 28, 2020.

Figure 9 shows the difference of the "observable" – a weighted sum of the channel values of T_{sys} with weights as determined in Butler (2020) – for two antennas (10 and 25) as a function of time during that test. While those weights will not be right if the data are not on a proper temperature scale, since the spectra are at least close to what is expected they might give a proper difference quantity. Figure 10 shows this quantity with its intrinsic 1 second sampling, then smoothed to 10 seconds, then smoothed to 20 seconds. It is clear from the raw data (1 second sampling) that the noise is far too high to do reasonable phase correction with these values. The noise (rms of the observable) is ~2 K; this is almost a factor of 100 higher than is required (Hales 2019). When smoothed to 20 seconds, what might be real fluctuations can be seen. This same issue was noted by Nathan (Towne 2019b), where he suggested smoothing only the $P_{hi} - P_{lo}$ (what he called P_{Δ}) portion of the T_{sys} calculation. I am currently investigating this. Note that we found the same thing with earlier VLA WVR data, where we smoothed both the equivalent of $P_{hi} - P_{lo}$ (to 10 minutes) and T_{obs} (to 20 seconds; Chandler+ 2004a).

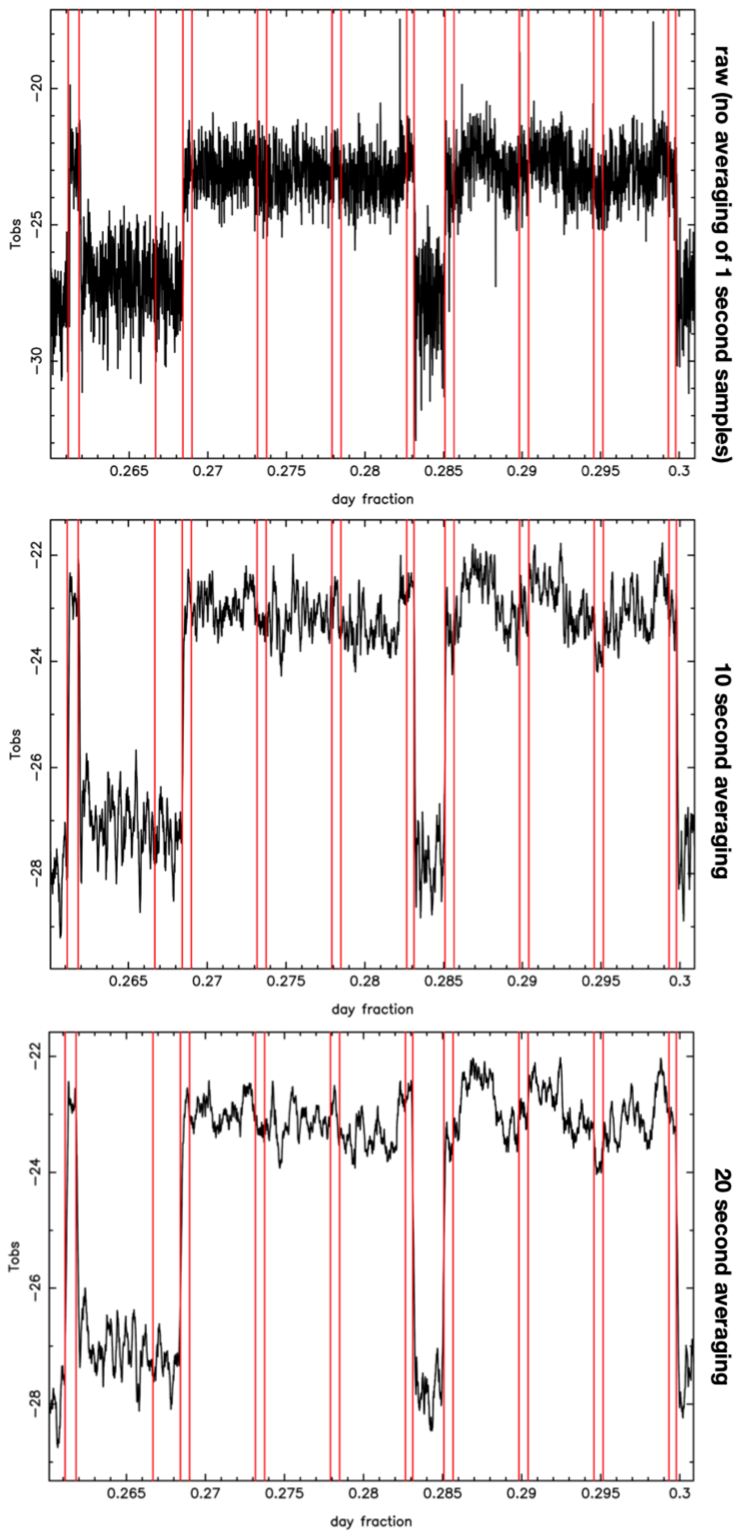

Figure 10. The difference of the "observable" (Tobs) for antennas 10 and 25, during an observation on May 28, 2020. The duration of the test was ~1 hour. Top panel is with no smoothing, middle panel is 10 second smoothing, lower panel is 20 second smoothing. The red lines are scan boundaries (which are almost all source changes, though the positions are similar for all sources).

6.4 Comparison to Visibility Phase

If we knew the atmospheric physics perfectly, and were doing proper radiometry, then we could convert the difference values of the WVR T_{obs} values on a baseline directly to an estimate of visibility phase with a correction factor for Kelvin to degrees of phase. However, neither of those is true – see Figure 7 and related discussion in Chandler+ 2004a (and references therein) regarding the first point, and the discussion above regarding the second point. However, we may still attempt to do a simple brute-force fit of the T_{obs} values to the visibility phases, allowing for an offset and a scaling factor, to investigate whether atmospheric phase is being tracked at all.

Denote the visibility phases on the i-j baseline as ϕ_{ij}. Then if we want to allow for a best-fit of the T_{obs} difference values to the visibility phases over time of the functional form:

$$\phi_{ij}^{WVR} = a + bT_{ij}$$

(allowing for an offset and scaling factor) where ϕ_{ij}^{WVR} can be thought of as the WVR estimate of visibility phase, and T_{ij} is the T_{obs} difference on that baseline, we can derive the solution for the coefficients a and b in a least-squares sense by forming the sum of squared differences, differentiating with respect to a and b, setting those equal to 0, and jointly solving for the coefficients. That solution is:

$$b = \frac{N\sum \phi_{ij}T_{ij} - \sum \phi_{ij}\sum T_{ij}}{N\sum T_{ij}^2 - \left(\sum T_{ij}\right)^2} \quad ; \quad a = \frac{\sum \phi_{ij} - b\sum T_{ij}}{N}$$

where the sums are over the N values of the visibility phase and WVR difference measurement over time. One could go even further and derive the channel sum coefficients used in the creation of T_{obs} for each antenna independently, given the visibility phases (as was done in Towne 2020a), rather than taking them as given from an *a priori* analysis (from Butler 2020, for example, as I've done above). Given the other problems with the system, I see no compelling reason to go that far at this point.

For the observation discussed above, and plotted in Figure 10, I have done this. Figure 11 shows the results, where visibility phase from one spectral window (near the center of K-band) is compared to the fitted T_{obs} difference for baseline 10-25, for the two "targets" (J1310+3233 and J1326+3154), each observed for three 6 minute scans (initial calibration done by observing the nearby calibrator J1310+3220). I see very little correlation of the T_{obs} difference to visibility phase, despite there being large excursions in phase (especially on J1326+3154). I have investigated other baselines, and data from the two other experiments of this type done around that time, and see similar results – the T_{obs} difference values simply do not seem to track visibility phase, at least at that point in time. I cannot reproduce the successful results shown in Towne (2020a) with these data.

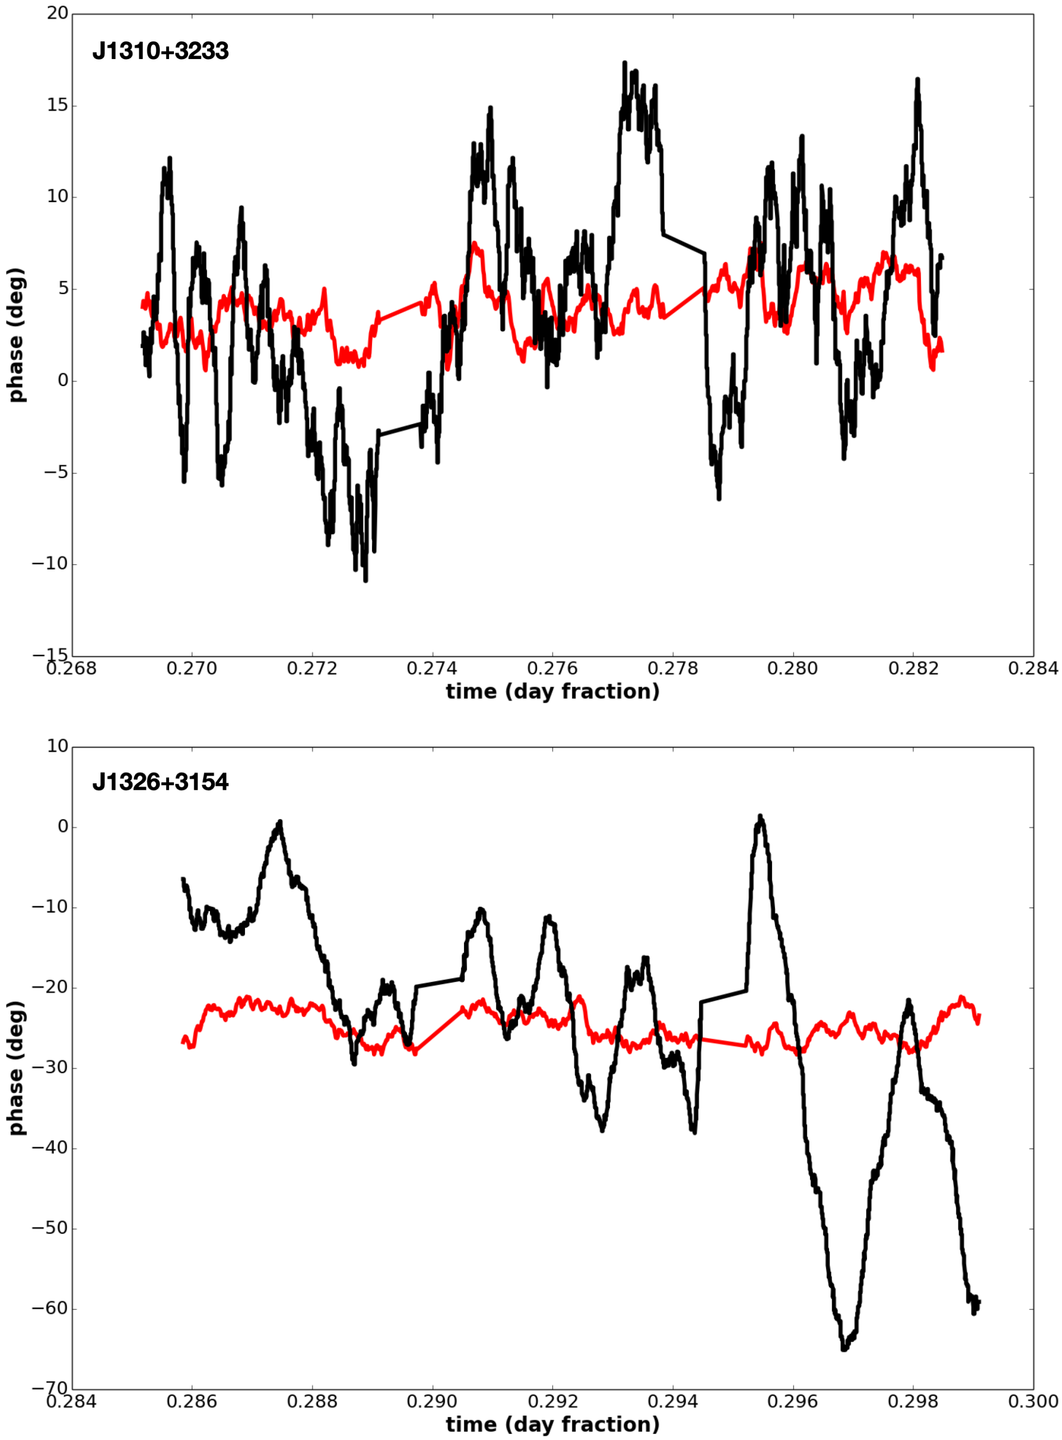

Figure 11. Visibility phase compared to the scaled T_{obs} difference on baseline 10-25 during an observation on May 28, 2020. Visibility phase is in black, Tobs difference in red. Both "target" sources (which are really calibrators themselves) were initially calibrated with an observation of J1310+3220.

7. Next Steps

The general state of the nCWVR units sometimes suffers, so it would be good to have somebody who understands the units well normally checking on that, and looking into problems as they come up. The values noted in section 5.3, for instance, should be looked into as they arise. They are not uncommon, and make application of the nCWVR data to visibility phase correction impossible.

It seems imperative that the cause of the improper scaling of the temperature values for the channels be figured out. Otherwise, we are not doing radiometry, and so are not really doing water vapor radiometry. If it is just improper T_{cal}, then perhaps those values can be calibrated, as discussed in section 6.2. If there are other issues, as I suspect, then further investigation is warranted.

Unfortunately, it does not appear that there is anything to be done about the sensitivity of the nCWVR measurements, without a significant re-design and/or re-work. It remains to be seen whether smoothing will work to give the needed sensitivity. I need to implement smoothing of $P_{hi} - P_{lo}$, and investigate whether that will get us the sensitivity needed at the time sampling interval we need.

We also need a simpler way to extract nCWVR data, and compare to visibility phase. This is particularly onerous currently, and while I have some software to do parts of it, much of it is still hands-on. One thing that we have been working on for some time is the inclusion of WVR data in the SDM. A table has been defined, and jointly agreed to by ALMA, and is in the process of being implemented. Once that happens, all of the data will be in one place (in the combined SDM+BDF) and it will be much easier to write software to access it. That does not mean that writing that software is fundamentally easy, it will just be easier.

Once the above is done, then a more complete analysis of how well the nCWVRs work under different observing conditions, different baselines, different primary observing frequency, etc., can be undertaken. Something similar to what was done in Chandler+ (2004a). This will be critical input to our understanding of how well WVR phase correction might work for ngVLA.

8. Conclusions

While the nCWVR units have perhaps been useful from an engineering perspective (gaining experience with building and installing this kind of electronics), the lack of a dedicated engineer assigned to the units for the past year has meant that their performance has suffered. Not only that, but some of the experience gained has been lost. The units are not accurate (not on an absolute temperature scale) so are not radiometers in the proper sense. They do not meet the required sensitivity requirement, and comparison with visibility phase is difficult. It is therefore not possible to draw conclusions about the use of WVR for ngVLA phase (delay) calibration based on these units at this point in time.

References

Butler, B., Some Issues for Water Vapor Radiometry at the VLA, VLA Scientific Memo 177, 1999

Butler, B.J., Channel Weights for the VLA CWVRs, ngVLA memo 74, 2020

Chandler, C.J., and 4 others, Results of Water Vapour Radiometry Tests at the VLA, EVLA Memo 73, 2004a

Chandler, C.J., and 5 others, A Proposal To Design and Implement a Compact Water Vapour Radiometer for the EVLA, EVLA Memo 74, 2004b

Gill, A., and 9 others, A Study of the Compact Water Vapor Radiometer for Phase Calibration of the Karl G. Jansky Very Large Array, EVLA Memo 203, 2017

Hales, C., Calibration Strategy and Requirements, ngVLA Document 020.22.00.00.00-0001-REQ-A-ARRAY_CALIBR_STRATEGY_REQS, 2019

Hjellming, R.J., An Introduction to the NRAO Very Large Array, 1992

Towne, N., Switched Power Correction for the Long Analog Response Times in the CWVR, Internal Memo, 2019a

Towne, N., Early Data on Achievable T_{sys} Precision in the CWVR, Internal Memo, 2019b

Towne, N., Synthesis of Optimal Estimators for Water Vapor Radiometry, ngVLA Memo 79, 2020a

Towne, N., Pathfinder WVR Monitor and Control Points, Internal Memo, 2020b

Welch, W.J., Correcting Atmospheric Phase Fluctuations by Means of Water-Vapor Radiometry, Review of Radio Science 1996-1999, ed. W.R. Stone, Oxford, 1999