EVLA Memo 240

Relation Between Physical Temperature and System Temperature and the Impact on EVLA Performance

Pedro P.B. Beaklini
March 10, 2026

Abstract

This memo aims to quantify the extent to which an increase in the physical temperature at which the EVLA receivers operate affects individual antenna performance and the overall thermal noise of the array. I cross-matched the monitor catalog of the temperature data points with the uncalgains and the system temperatures obtained in several stress tests. The results show an expected dependence of antenna performance and the system temperature ($T_{\rm sys}$) on the physical temperature ($T_{\rm cryo}$). I derived a model that describes $T_{\rm sys}$ as a function of $T_{\rm cryo}$ and estimated the overall array thermal noise for each temperature measurement. The primary conclusion is that an individual antenna's performance does not severely degrade with a 20 K increase in temperature, but the entire array's sensitivity can be severely impacted if several antennas are only a few degrees out of range. I expect that the present memo will help determine priority when several compressors need attention at the same time.

1 Introduction

It is well known that to achieve the sensitivity required to observe astronomical radio signals, we need to reduce thermal noise as much as possible. This is valid from centimeters to submillimeter wavelengths, although at centimeter wavelengths, operating at room temperature is not a blocker; it compromises the instrument's sensitivity considerably.

At the EVLA, the cryogenic coldhead connected to the receivers has two monitored data points, one at each stage, as shown in Figure 1. The labels indicate the nominal operating temperatures of each stage: 15 K and 50 K. Maintaining such low temperatures requires regular charging of the compressors, and without careful supervision the temperature can easily increase beyond the desired range. Although operating at 15 K is the goal, it is rare for all antennas to operate at those temperatures across all bands simultaneously. When the refrigeration breaks down completely, the receiver reaches room temperature and the operators place it in standby. For a more detailed description of the cryogenic systems at the EVLA, I refer the reader to EVLA Memo 196 (Urbain et al. 2016).

In the VLA, excluding the antenna in the barn undergoing maintenance, we have 27 antennas with 8 receivers each, which should operate at 15 K, totaling 256 $T_{\rm cryo}$ monitoring points only for the T15 stage and another 256 for the T50 stage. And surely, it is common that more than one receiver starts to get warmer simulta-

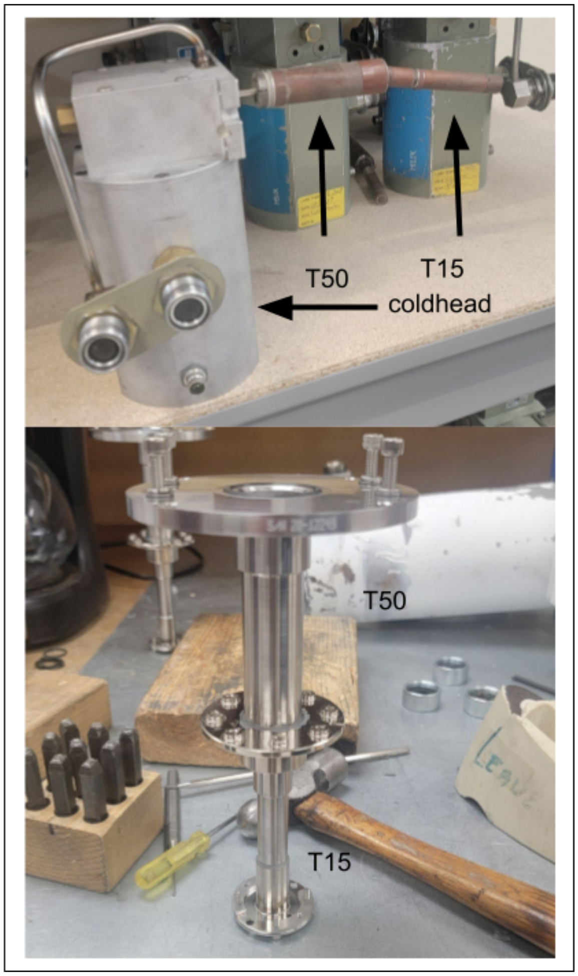

Figure 1: Coldhead and its connection to the receiver chamber. Top: photograph indicating the location of the T50 and T15 temperature stages. Bottom: photograph showing how the coldhead are connected to the receiver.

neously and creates a priority criterion that is crucial. Today, the alert thresholds are 40K to the T15 stage and 100K for the T50 stage, meaning that when those monitored points reach the respective temperatures, an alert is sent to the operators and the cryogenic team (cryo team hereafter). That means that a receiver operating at 30 K is considered normal operation, although the cryo team works hard in an attempt to keep everything within the requirements.

The present memo aims to quantify the effect of an increase in the physical temperature $T_{\rm cryo}$ at which the receiver operates on the system temperature $T_{\rm sys}$, and consequently on the thermal noise and the sensitivity of the EVLA. The analysis was motivated by a request from the cryo team to evaluate how changing the current threshold alerts would impact operations. Although it is obvious that there is a correlation between the physical temperature and the system temperature, an accurate relation obtained during regular EVLA operations has not been determined to date.

In Section 2, I will show the methodology followed to evaluate the relation between physical temperature and EVLA performance, while in Section 3 I will focus on empirically determining the relation with $T_{\rm sys}$, and in Section 4 I will use it to obtain the thermal noise. Final considerations are stated in Section 5.

2 Methodology

The temperature data were extracted from the monitor database, where data are recorded approximately every 10 seconds. I performed a cross-match between the temperature data and the stress test records (EVLA Memo 186; Morris & Momjian 2014) using the Julian day of the observation, without splitting the data per scan. The time information was taken from the stress test records, and I assumed the average temperature over a one-hour interval across all bands. There is a caveat to this approach: the time reported in the stress test corresponds to when the data were processed, not to the actual observation time.

However, stress tests are typically processed shortly after completion by the operators, and I assume that the temperature during the hour preceding processing is sufficient to cover the full stress test interval (36 minutes). I do not expect significant temperature variations on minute timescales; temperature changes are generally slow, and this approximation is representative of the conditions during SB execution and adequate for our purposes.

Occasionally, longer delays between data acquisition and processing, combined with abrupt temperature changes associated with severe failures, could contribute to additional scatter in the data. The standard deviation within the one-hour window was therefore used to estimate the uncertainty in the temperature and to identify such cases. Nevertheless, as shown in the next section, no complex data filtering was required to reveal the correlation.

I obtained the temperatures from the T15 and T50

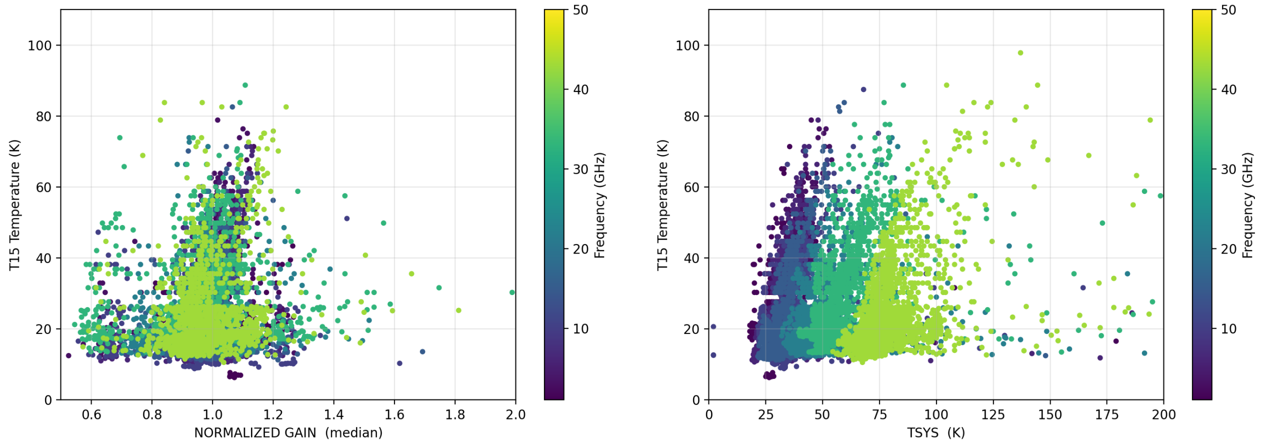

Figure 2: Physical temperature (K) at the T15 stage as a function of antenna performance (normalized stress-test uncalgains; left) and system temperature $T_{\rm sys}$ (right). The plots show a weak dependence of antenna performance on temperature but a clear increase of $T_{\rm sys}$ (band dependent) with increasing cryogenic temperature.

stages in the monitor database, filtered them to the time corresponding to the stress test, and selected the data for autumn and winter. The strategy is to ensure a larger amount of good weather data. I used all stress data from January to April 2025 and from October 2025 to January 2026, without any arbitrary selection. The only obvious exclusion was the January 8, 2026, stress test, when something clearly went wrong due to abnormally high $T_{\rm sys}$ values across several antennas and bands. I correlated $T_{\rm cryo}$ (both T15 and T50 stages) with uncalgains and $T_{\rm sys}$. I did not take data prior to 2025 because data older than 13 months are moved to a long-term storage archive.

As a final consideration, since the uncalibrated gains are strongly affected by weather conditions and adverse weather can still occur during the winter, I normalized the data on a daily basis for each band. Specifically, for each band and each day, I computed the average uncalibrated gain and expressed the values relative to the daily median. This strategy has the advantage of allowing the temperature to be compared with the average system performance while largely mitigating weather-related effects. However, it also has the drawback of potentially masking gradual performance degradation, since the average performance may differ from one day to another. Nevertheless, we do not expect year-long timescales to be sufficient to reveal significant overall performance changes, and, for the analyses presented here, this approach is appropriate.

3 Results: $Tsys \times Tcryo$

In Figure 2, I show the dependence of EVLA antenna performance and $T_{\rm sys}$ on $T_{\rm cryo}$ measured at the T15 stage. The antenna performance is described in terms of normalized uncalgains, which I labeled as normalized gains for simplicity. Both plots are colorized by frequency (band). It seems straightforward to conclude that performance is not severely band-dependent, whereas the $T_{\rm sys}$ correlation depends strongly on the observed frequency. I have excluded those data with uncalgains and system temperatures higher than 200 and physical temperatures lower than 20 K because they obviously have issues not related to a warm receiver. Also, keep in mind that $T_{\rm sys}$ depends on elevation. However, since we are analyzing stress test data obtained from circumpolar targets, the elevation does not vary significantly across the sample. Elevation effects may contribute to the observed scatter, but they do not alter the observed correlation.

Starting with the plot showing EVLA antenna performance, it represents a clear prediction of radio astronomy instrumental theory: sensitivity decreases as receivers operate at higher temperatures. But more than confirming what we already know, the plot shows a trend, how much the EVLA performance scales with the physical temperature. Considering that the current temperature alert is 40 K for the T15, a change to 60 K alone does not seem to have a very significant impact on our performance compared to what is currently

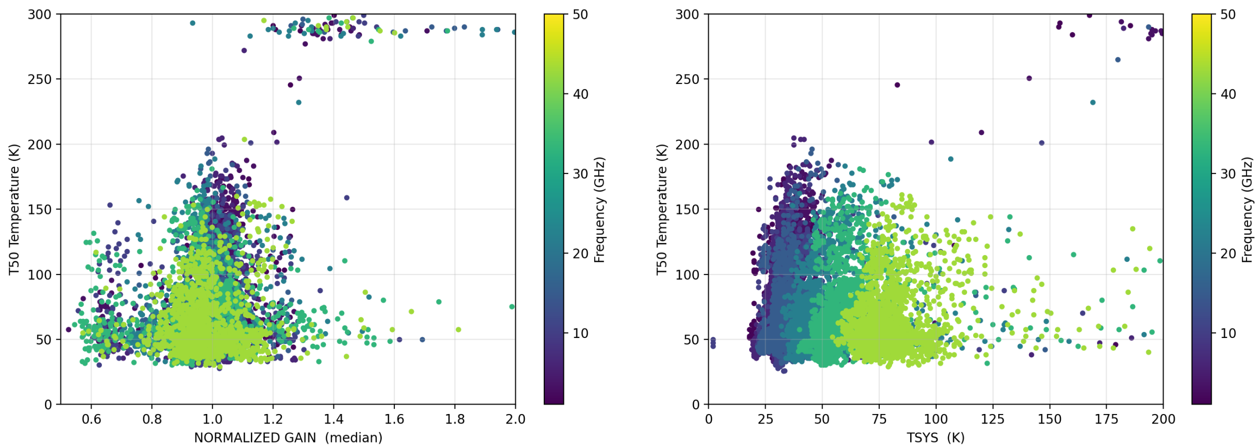

Figure 3: Physical temperature (K) at the T50 stage as a function of antenna performance (normalized stress-test uncalgains; left) and system temperature $T_{\rm sys}$ (right).

done. However, it is a significant loss of efficiency when compared with the 15 K requirement.

The reader must also keep in mind that the plot shows the performance of a single antenna and does not represent the entire array. It shows that, under typical EVLA operating conditions, efficiency losses reach a few percent due to higher temperatures. If we consider other losses besides thermal losses, such as optical issues or motor failures, a variation of 10% from the average of a given antenna is typically considered acceptable, as we can see in the pointing spread showing performance values up to 1.2 at low $T_{\rm cryo}$. We typically raise a red flag when an antenna is much worse than others by more than 20%. Of course, all the efficiency losses are added to the final performance of an antenna, but an increase of a few K alone will not be the reason for a severe loss of performance in a given antenna. The conclusion may differ if we consider an increase across the entire array, as I will discuss in the next section.

The right plot in Figure 2 correlates the physical temperature with a direct quantity that requires no further explanation: the system temperature ($T_{\rm sys}$). We expected that the increase in temperature would raise $T_{\rm sys}$, but how much it would increase was still unclear under regular EVLA observation conditions (i.e., outside a lab). The plot shows that the dependence differs across bands, with a steeper slope at high frequency. It also shows that a linear model fits well if we separate the band dependence.

Before discussing the band-dependent behavior of $T_{\rm sys}$, it is worth commenting on the behavior of the T50 stage. In cases of major cryogenic issues, such as when a band requires coldhead replacement, both temperature data points tend toward ambient temperature, resulting in a flat relation. However, in borderline cases, the antenna performance dependence on the T50 reproduces the correlation observed with T15, although with a larger scatter (Figure 3).

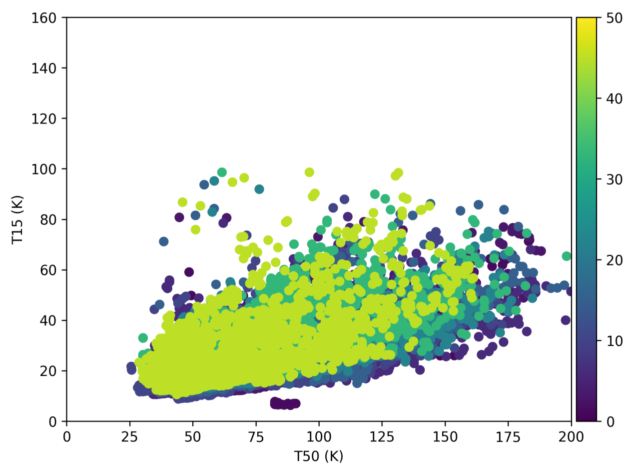

Figure 4: Relation between the two cryogenic temperature monitoring points, T15 and T50.

Figure 4 shows the relation between the T15 and T50 monitoring data points. Since both stages are strongly correlated, I proceed with the analysis using T15 as $T_{\rm cryo}$. The T50 is more of an indication that something is about to happen to the receiver than a relevant issue

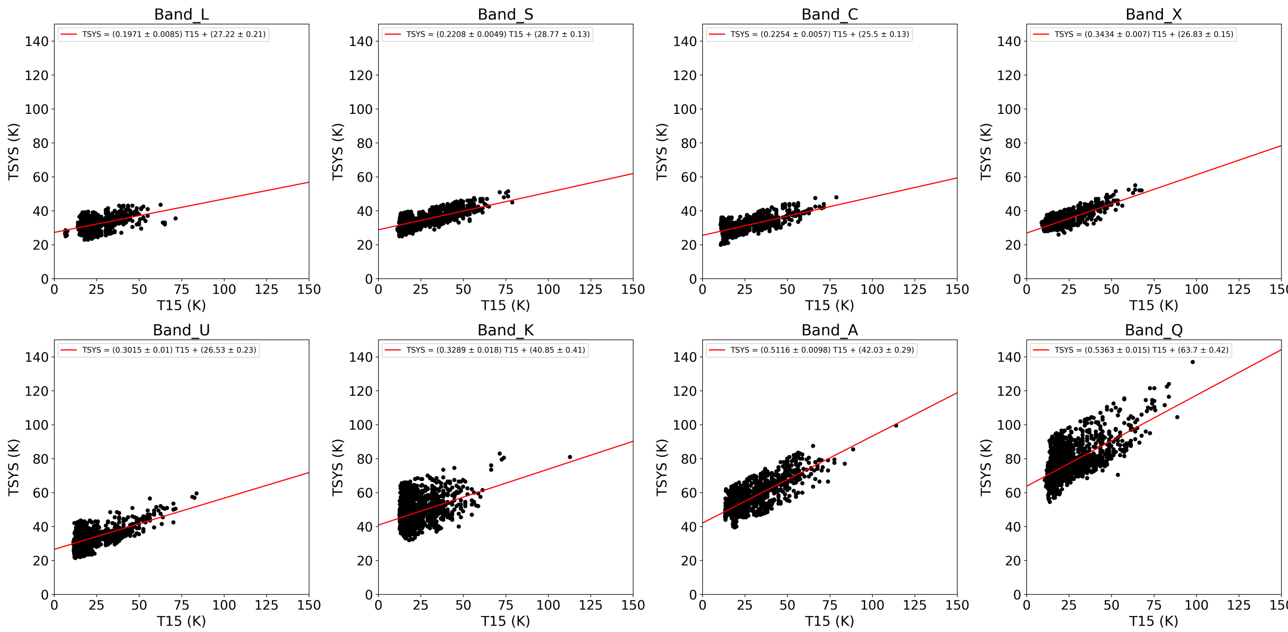

Figure 5: Relation between $T_{\rm sys}$ and $T_{\rm cryo}$ per band. The best-fit parameters are indicated in the legend. Note that the axes are inverted relative to Figures 1 and 2 so that $T_{\rm sys}$ is expressed as a function of the T15 monitoring point.

affecting the data when only T50 is high. This becomes evident when we return to Figure 1 and see how the monitored data points are physically close to each other. Modeling $T_{\rm sys}$ as a function of both T15 and T50 seems pointless and I will consider T15 as a representative of the physical temperature at the receiver location, which is evidently more appropriate for predicting $T_{\rm sys}$.

In Figure 5, I show the fits for all bands considering all antennas. There is a spread in the high-frequency bands that points toward antenna dependence, while the slope remains the same. In the low-frequency bands, the spread is very narrow. I performed a linear fit for all cases, and although the solutions are shown in the plots, I also added them to Table 1 for better visualization. I chose to model $T_{\rm sys}$ as a function of $T_{\rm cryo}$ (rather than the other way around) because this formulation is more intuitive and directly reflects how the physical temperature affects the system.

Table 1: Best fit for each band, where $T_{\rm sys}(K) = a\,T_{\rm cryo}(K) + b$ and $T_{\rm cryo}$ represents the T15 stage.

Band	a	b
L	0.20 ± 0.01	27.22 ± 0.21
S	0.22 ± 0.01	28.77 ± 0.13
C	0.23 ± 0.01	25.50 ± 0.13
X	0.34 ± 0.01	26.83 ± 0.15
Ku	0.30 ± 0.01	26.53 ± 0.23
K	0.33 ± 0.02	40.85 ± 0.41
Ka	0.51 ± 0.01	42.03 ± 0.29
Q	0.54 ± 0.02	63.70 ± 0.42

Table 2: The aperture efficiency values used to match the thermal noise predicted by the exposure calculator.

Band	efficiency	Band	efficiency
L	0.256	Ku	0.504
S	0.461	K	0.292
C	0.616	Ka	0.401
X	0.664	Q	0.329

4 Discussion: Thermal Noise

Once we have modeled $T_{\rm sys}$ as a function of $T_{\rm cryo}$, we can obtain the thermal noise at the EVLA for each $T_{\rm cryo}$ data point (i.e., every $\sim$10 seconds). For that, we

Table 3: Time required (s) to reach the reference thermal noise for a 128 MHz bandwidth under different cryogenic conditions.

Bands	L	S	C	X	Ku	K	Ka	Q
Reference noise (μJy)	215	125	88	89	115	284	228	403
27 at 15 K	60	60	60	60	60	60	60	60
24 at 15 K + 3 at 50 K	62.5	62.7	63.0	63.8	63.5	62.8	63.7	62.9
24 at 15 K + 3 at 70 K	63.6	63.9	64.2	65.3	64.9	64.0	65.2	64.1
27 at 22 K	65.5	66.0	66.7	69.2	68.3	66.3	68.9	66.7
27 at 35 K	72.0	73.3	74.9	80.6	78.4	73.9	79.9	74.2

can use Equation (11) from EVLA Memo 127 (Perley 2008), which estimates the thermal noise of a baseline (or a two-element array). We can assume no source contribution to set the $r_i + r_j + 2r_ir_j$ term to zero and obtain a reference value, but the thermal noise can also be computed including the contribution of a particular source if needed. The most important assumption is that the slope between $T_{\rm sys}$ and $T_{\rm cryo}$ is independent of elevation, which is likely the case. Elevation should affect the intercept rather than the slope. We can obtain the total VLA noise by considering the contribution of each individual baseline.

The second most important assumption concerns aperture efficiency, since the system equivalent flux density (hereafter, SEFD) depends on it and is strongly dependent on wavelength (even within a particular band). However, I do not really need it, since I do not aim to measure the absolute thermal noise but rather to understand how much an increase in thermal noise degrades performance. I assumed that if all antennas operate at 15 K, we would obtain the same sensitivity as provided by the EVLA exposure calculator, which is what we offer to our users[1]. The aperture efficiencies are therefore tuned only to reproduce the reference sensitivity provided by the EVLA exposure calculator. Although these coefficients are not required for the analysis itself, the fact that they remain close to the expected EVLA efficiencies (see Table 2) provides a useful self-consistency check[2].

Finally, I obtain the thermal noise from the measured $T_{\rm cryo}$ and compare it with the ideal exposure calculator scenario to determine how much operating above 15 K degrades performance. We can discuss thermal noise directly in terms of noise increase, bandwidth, or as a function of observation time[3]. I therefore evaluate how much observing time is required to reach the noise level predicted by the exposure calculator (for a 128 MHz bandwidth and a one-minute observation using 27 antennas) given the measured $T_{\rm cryo}$. I present the numbers in Table 3 under different conditions. The first column indicates the number of antennas at different temperatures.

Table 3 also reveals an important fact. We currently set alert thresholds per antenna. However, we can see that having three antennas with higher thermal noise is not as problematic as having all antennas operating at 25 K, and we currently do not have any alerts for a general increase in thermal noise across the array. I suggest changing our approach to prioritize the cryo team's efforts and, instead of focusing on individual antenna alerts, consider the overall thermal health of the array in a given band.

It is true that when one antenna is at a higher temperature, even though its contribution to the overall array thermal noise is not significant, it implies a higher noise level for the $N-1$ baselines that the antenna is part of. That can have a major impact on the observation if the affected baseline is the longest one. Nevertheless, if we return to Figure 2, it shows that an increase in $T_{\rm cryo}$ from 30 K to 60 K does not result in a performance downgrade greater than 20%, which is comparable to our regular performance spread when considering all other losses. We could discuss prioritizing the antennas at the ends of the arms, but we do not do that in other maintenance prioritization efforts across the groups. During Y1 observations, operations and science

[1]The assumption is not strictly true, since the exposure calculator uses SEFD values derived empirically from real observations, which depend on the values of $T_{\rm cryo}$ on a given day. However, this approximation is not far from reality.

[2]The only exception is the K band, where $T_{\rm sys}$ is strongly affected by the atmospheric water line, and the tuned efficiency was lower than the actual value.

[3]thermal noise scales as $\sigma \propto T_{\rm sys}/\sqrt{\Delta\nu\, t}$, where $\Delta\nu$ is the bandwidth and t is the observation time.

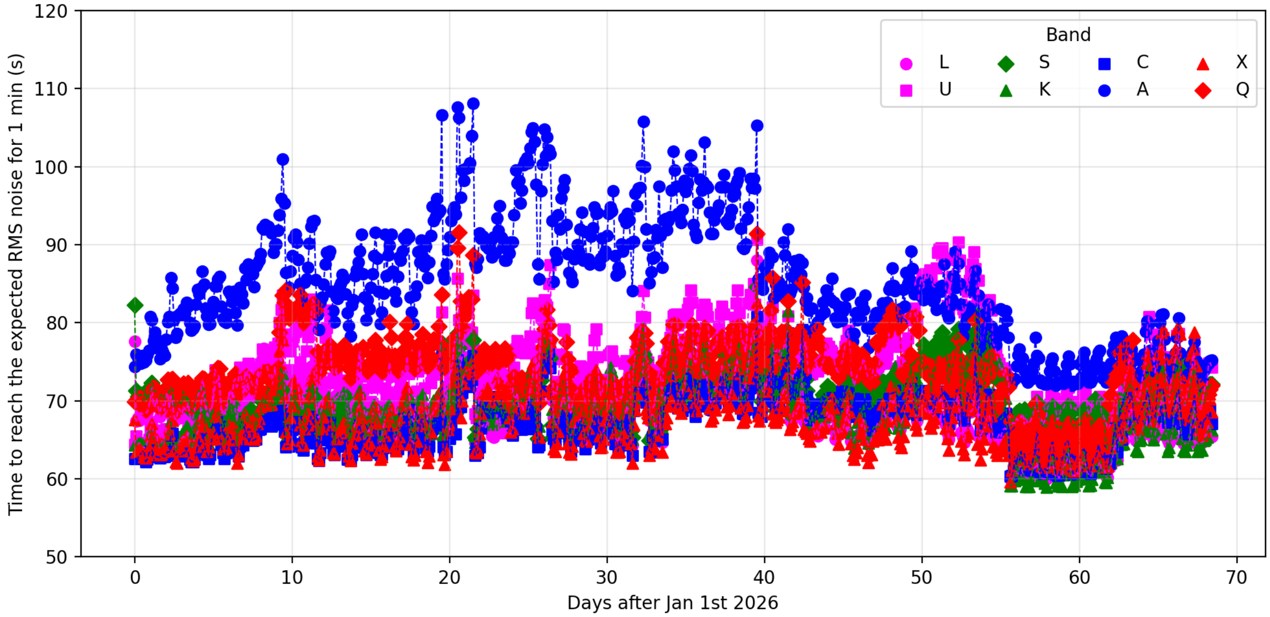

Figure 6: VLA thermal noise expressed as the time required to reach the sensitivity expected by the EVLA exposure calculator for an observing time of 60 s, a bandwidth of 128 MHz, and using 27 antennas. When 27 antennas operate at $T_{\rm cryo} = 15\,$K, the required time matches the value obtained from the EVLA exposure calculator. During the period between February 26 and March 3, the VLA operated with 28 antennas, with antenna ea04 on the master pad, which appears in the plot as values below 60 s.

staff typically check the health of the antenna systems, which includes $T_{\rm cryo}$, and this practice will not change. These are not blockers to revisiting our thresholds in terms of overall array health.

In Figure 6, I show the observing time loss we are experiencing in 2026. The reader should have in mind that I did not compare with the previous epoch, and that it represents only a snapshot of the current situation. More work is needed to determine how the current thermal noise relates to other epochs. The goal of Figure 6 is to show that having one number per band describes the array thermal noise better than chasing particular antennas.

The Ka band needs a dedicated discussion. The data shows a higher thermal noise than in any other band. Between the B and A array reconfiguration in the 2026 winter, we spent two weeks on the BnA configuration for the VLA Sky Survey, which was conducted at S band. Therefore, the S-band performance evaluation and maintenance work had higher priority. During the same period, the cryo team had a shortage of coldheads, and several antennas were not being cooled at Ka band. One weekend, we were out of six Ka-band antennas. Therefore, the increases in the Ka band thermal noise between January 1 and January 10 and between January 15 and January 20 are real. The thermal noise at Ka band started to decrease considerably by March.

However, even in March, and when examining the first two weeks of 2026, the Ka band still shows higher thermal noise than the other bands, although less prominently. This could mean that we are underestimating the RMS noise level in the EVLA exposure calculator for that band. It may be worth carrying out a more careful analysis to correlate the thermal noise estimated from $T_{\rm cryo}$ with the empirical SEFD at Ka band.[4].

It is also notable in the figure that the thermal noise decreases when one extra antenna is available. The reference values for thermal noise, either in the plot or in Table 3, assume 27 antennas, since that is the number of antennas we normally operate. Between February 26 and March 3, the array temporarily operated with 28 antennas, with antenna ea04 on the master pad. This

[4] Emmanuel Momjian and I are planning a test to measure the Ka-band sensitivity of the array in a field without strong sources to determine whether it correlates with $T_{\rm cryo}$ as expected.

reduces the thermal noise and produces values below 60 seconds in the figure.

It is also important to keep in mind that the plot shows the thermal noise of the entire array; it does not imply that all antennas were actually in use at a given time. Only if a receiver is set to standby would the loss of one antenna be reflected in Figure 6.

5 Conclusion

The present memo quantifies the impact of an increase in the physical temperature at which the receivers operate on T_{sys}. As a consequence, it shows how a warming receiver impacts the EVLA sensitivity. The present analysis is important in supporting the cyo team in defining prioritization and maintenance strategies.

The cryo team constantly monitors the temperature of all receivers, but they receive specific alerts when the T15 stage temperature reaches 40 K and the T50 stage temperature reaches 100 K. The memo was intended to directly address the question of how changing those alerts to 60 K and 150 K, respectively, would impact science.

The short answer to the driving question is that the impact is minimal. It is true that the impact is not zero, since performance degrades by a few percent (Figure 2). However, as shown in Table 3, having three antennas above 60 K does not result in a significantly larger time loss than having three antennas above 40 K.

A more detailed answer is that not only is the impact minimal, but having all antennas operating at 25 K is actually worse than having three antennas operating at 60 K. Yet we currently have no alerts for that scenario.

Below is a summary of my conclusions, which can be thought of as a prioritization strategy.

– T_{sys} is highly correlated with the T15 stage temperature, but not independently correlated with the T50 stage temperature. The T50 stage is more an indication that the T15 stage will start to increase than an immediate issue.

– The impact of one antenna operating at 60 K is not worse for the thermal noise level than having several antennas at 30 K.

– Having an alert for overall thermal noise for a given band is more crucial than chasing particular antennas with slightly higher temperatures.

– There are no significant differences between the bands. A one-degree increase has a similar impact on the thermal noise level across frequencies (Table 3), although its effect on T_{sys} differs (Table 1).

– The overall Ka-band thermal noise level is constanttly higher than in other bands, as indicated by the EVLA exposure calculator. That requires more analysis.

Of course, it is up to the cryo group to define priorities and workflows to keep the system operating. My intention here is to show how these factors impact the data and which variables are most crucial for maintaining VLA operating conditions as we desire (or at least as we promise to our users).

To ensure some flexibility without compromising EVLA data quality, the present memo supports a new alert scheme, flagging antennas above 40 K as yellow and those above 60 K as red, rather than using a single alert threshold. The Cryo and VLA Operations groups agreed not to have more than three antennas above 40 K, and we expect that this will make it easier to allocate manpower and equipment to the most severe cryogenic issues. We are also discussing a new alert based on overall thermal noise, which would help identify specific bands that require more attention.

An underlying message of this memo is how the Engineering, Electronics, and Science Divisions can work together to identify improvements in EVLA operations. Any work that impacts antenna performance can be translated into overall operational costs. I decided to show the Table 3 as a function of time because it is easier to visualize the degradation in performance, and because time is directly correlated with money. If we need 1h01m to reach the 1h sensitivity in ideal conditions, that means VLA time being lost at a rate of 1m per hour.

Of course, we know that the user-allocated time is based on the calculator, and what is really happening is that they are getting noisier data. This highlights how crucial coordinated efforts across the different divisions are to keep the system close to its operational requirements, and how shortages of components can directly affect VLA operations.

Acknowledgements

I would like to thank Jesse Madigan, leader of the Cryo group, for bringing this question to be evaluated within the new VLA Operations Group. I also appreciate his careful explanation of the T15 and T50 stages, which allowed me to produce Figure 1.

I would like to thank Daniel Faes from the NM Software Group for creating a script to easily access the T15 and T50 databases without using any interface. This analysis would not have been possible without direct access to the monitor database.

I would like to thank my colleagues from the Science Division, Paul Demorest and Emmanuel Momjian, for our fruitful discussions regarding the relation between performance, $T_{\rm sys}$, and thermal noise, and for suggesting limits beyond the initial thresholds.

I also appreciate the support of Rich Moeser from the NM Software Group in implementing the two-color alert scheme in the new version of the ops screen package.

References

[1] Urbain, D., Grammer, W., Peck, G., Jackson, J., Durand, S., 2016, EVLA Memo 196
[2] Morris, K., & Momjian, E. 2014, EVLA Memo 186
[3] Perley, R. 2007, EVLA Memo 127