RFI Flagging and Replacement for Calibration

Daniel Bautista

Date: September 4, 2025
Keys: rfi, satellite constellation

Abstract

This memo describes a method for calibrating the intensity of transient broadband RFI in GBT data. This method relies on the RFI being statistically detectable from rest of the spectrum and replaces flagged calOFF integrations containing RFI with nearby calOFF integrations that are free of RFI. The technique works best when there are intermittent RFI emissions throughout the scan and does not perform better than standard calibration methods when the RFI persists for the duration of a scan. `dysh`[1], a Python-based spectral line data analysis program, was used to develop this method.

Contents

[1] https://dysh.readthedocs.io/en/latest/

Changelog

320.0 Bautista (2025-09-04) — Initial published version.

1 Introduction

It should come as no surprise to anyone who has taken radio observations that radio frequency interference (RFI) is an ever growing problem that needs to be addressed during data reduction. For a spectral line observer, the affected frequencies or integrations can be flagged and excluded from the calibrating and stacking of spectra. Pulsar observers can flag frequency channels to exclude during the pulse stacking of data. But in the case of studying the efficacy of RFI mitigation efforts, the strength of the interfering signals is the interesting part of the spectrum that should be preserved. The rapid time variation of signals from commercial satellites in low earth orbit cannot be accurately calibrated using the standard noise diode calibration methods (described below), so the data need to be cleaned to obtain an interference-free reference before calibrating.

2 Standard Calibration Method

The standard method for calibrating GBT spectral line data (Braatz, 2009) is to use noise diodes with a known contribution (T_{cal}) to the system temperature (T_{sys}) to determine the flux scale. This calibration is done on an individual integration basis and then averaged over the entire scan. Figure 1 shows a neutral hydrogen line captured during an RFI scan, and calibrated using this method. Thanks to the allocation of the radio astronomy protected band (1400 - 1427 MHz) this frequency range is clean of RFI. Unfortunately, outside of these protected bands, many of the low frequency receivers experience RFI across the band.

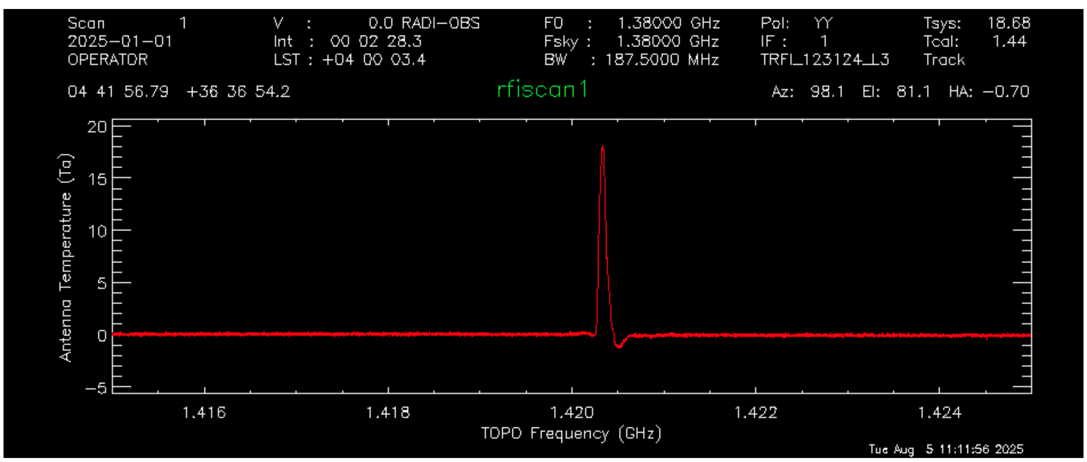

Figure 1: Calibration of an L-Band RFI scan using scan 1 as the signal scan and scan 2 as the reference scan, and zooming in to show a frequency range of 1415 - 1425 MHz.

This method of calibration does not perform well in the presence of RFI, especially emissions that vary significantly on short timescales. A strong signal in the "off" state and a weaker signal in the "on" state will result in a negative flux being calculated. The $\frac{ON-OFF}{OFF} \times T_{sys}$ method is prone to producing artifacts when calibrating channels with a high RFI density. Additional artifacts are introduced when dividing by the off scan, which may not fully remove any time-varying RFI.

RFI scans differ from the usual scientific observation setup due to the nature of the signals we aim to detect. A scientific observation will typically have four states that the data comes in. The calON and calOFF data specify whether the noise diode is in the "on" or "off" state. The other two states are "signal" and "reference," which specify whether the telescope is observing the target of interest (signal) or the something else (reference). For spectral line observations, the "reference" is typically either blank sky, for a position switched observation, or the "off" frequency for frequency switched observations. RFI is transient and unpredictable, which makes it difficult to obtain a reference scan that is free of interfering signals. As a result, standard RFI runs do not bother performing reference scans, and rely on the noise diodes to set the flux scale.

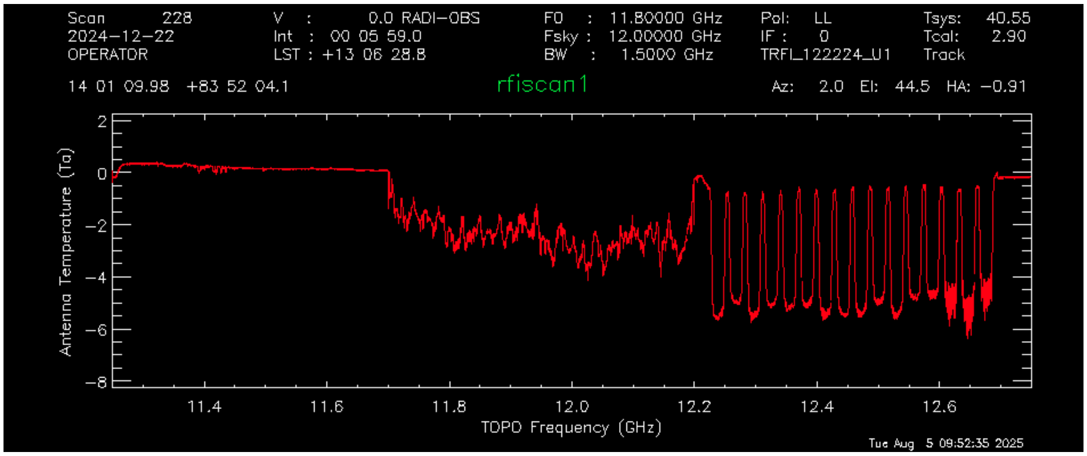

Figure 2: Calibration of a Ku-Band RFI scan using scan 228 as the signal scan, and scan 226 as the reference scan. Scan 228 had a strong RFI source localized in the southern part of the sky, and slewed across the location of the emitter as part of the standard RFI scan process. The negative system temperatures are a result of the reference data having a greater intensity than the signal data.

Figure 2 shows the result of calibrating data with strong RFI in the reference data. The strong RFI in the reference data result in a negative system temperature in the calibrated spectrum, which is quite obvious in this case, because the signal data had relatively little RFI. However, when both the signal and reference data have RFI distributed throughout the integrations, it becomes much harder to determine the true contributions of the RFI in the calibrated and averaged spectrum.

3 RFI Flagging and Replacement

The simplest method of calibrating RFI-dense data is to identify a "good" reference scan and apply it to the entire scan. With a well-chosen (RFI free) reference scan, this performs better than using the RFI-dense calON and calOFF pairs for each integration, but it can result in broad gradients across the spectrum as a function of time. This is not apparent in the average spectrum, but when looking at the data as a waterfall plot, it becomes obvious that there is a slow change in the sky across the scan(s). Artifacts introduced under the standard calibration technique demonstrate the need for cleaning the reference data before calibrating. The proposed alternate method is to flag the integrations that are affected by RFI and replace them with similar data that are free of RFI. These RFI-free reference scans can then be used to calibrate the data.

3.1 Assumptions

We assume that the noise diodes will remain stable for the course of the RFI scan and that the continuum contribution of the sky varies slowly over the course of the scan. The noise diode contribution and RFI-free "off" scan can then be used to perform the $\frac{ON-OFF}{OFF} \times T_{sys}$ calibration. A potential pitfall of this method could arise if there are so few integrations that are free of RFI that a clean "off" scan cannot be identified. In this case, there would be no RFI replacement, and the calibration would be the standard noise diode calibration.

3.2 Intended Use Cases

It is important to note that this calibration method is only intended to use to calibrate the intensity of transient RFI signals. If used on persistent RFI that has approximately constant emission strengths, the long-term lack of variability in the emissions will cause it to be largely missed[2] in the flagging step. This will allow for the same calibration artifacts that are seen in the standard calibration methods. It should also not be used for calibrating spectral line data from astronomical sources. The resampling of the calOFF data will introduce correlations in time and the reused data will affect the average spectrum.

3.3 Flagging

When observing with the noise diodes firing, the data comes in two "phases," cal state ON and cal state OFF. To identify the presence of strong RFI on short timescales (of order a few integrations), the calON and calOFF data need to be combined so the ordering of the data matches the order the data were taken. The calON data also has the contribution from the noise diode, which needs to be removed before searching for RFI. The noise diodes are very stable over the course of a scan, so the average noise diode spectrum is obtained by taking the channel-wise average over the duration of the scan. Assuming a stable noise diode, gradual changes in sky brightness with intermittent RFI, the average noise diode contribution is calculated by taking the median of the difference between the calON and calOFF data. Taking the median of the data prevents interfering signals from skewing the noise diode spectrum. Once the noise diode contribution has been subtracted from the data, it is ready to be checked for RFI. The time-wise difference of the data is taken. This will reveal when the received power changes significantly from one integration to the next. This change in power can be an increase (e.g. transmitting satellite entering the telescope beam) or a decrease (e.g. transmitting satellite leaving the telescope beam). This method is less effective with persistent RFI signals, such as emissions from a geosynchronous satellite, which may not move much relative to the telescope beam for a particular pointing.

For emissions with a well-known frequency allocation, the frequency range in question can be provided and searched for power levels exceeding the specified threshold which is set to one standard deviation above or below the median by default. The specified frequency range is averaged to a single value for each integration, and integrations with average powers N standard deviations away from the median are flagged as containing RFI. The value of N is provided by the user to set the threshold for RFI signals. The process of selecting a frequency range from the observation bandpass and identifying when RFI is present can be repeated for as many channel allocations as the user wants. Once all channels have been checked for RFI, the indices are collected and reduced down to the set of unique calOFF indices. At this point, the RFI flagging is complete and the flagged calOFF integrations can be replaced with cleaner data. It is worth noting that only the calOFF indices are selected for RFI replacement.

Figure 3 shows examples of when the flagging procedure would perform well and when it would fail to flag all the RFI. For clarity, all integrations are normalized to unity at a frequency that falls outside the searched frequency range. This allows the normalization to show the standard receiver bandpass and also to make apparent the integrations that contain RFI. Figure 3a shows all integrations from a single Ku-Band scan. In each of the labeled frequency ranges, there are some integrations that are elevated above the majority of the integrations. These are the integrations that contain RFI, and are the ones that are flagged and replaced. Figure 3b shows all integrations from a single Ultra Wideband Receiver (UWBR) scan. In the frequency range labeled "DtC" all integrations have power levels elevated above where the baseline should be. In this case, there are no integrations that are free of RFI, and the flagging process will fail to accurately identify and remove RFI and the new calibration procedure performance cannot be expected to do any better than the standard calibration.

3.4 Replacement

Once the offending RFI integrations have been flagged, replacement data needs to be identified. This method of replacement exploits the high time variability of these types of RFI, so that a suitable replacement integration can be sampled from integrations shortly before and after the offending integration. I identify the closest non-flagged integrations preceding and following the flagged integration. From these two "good" integrations, the minimum spectrum is selected on a channel-wise basis. This step of minimizing the amplitude serves to further protect

[2] It is possible that the first and last integrations containing the RFI will be flagged, if a source of emissions enters or leaves the telescope beam at the start or end of a scan

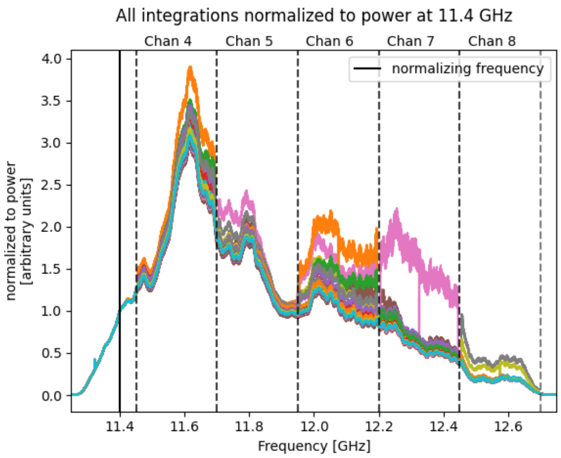

(a) Normalized Ku-Band spectral window.

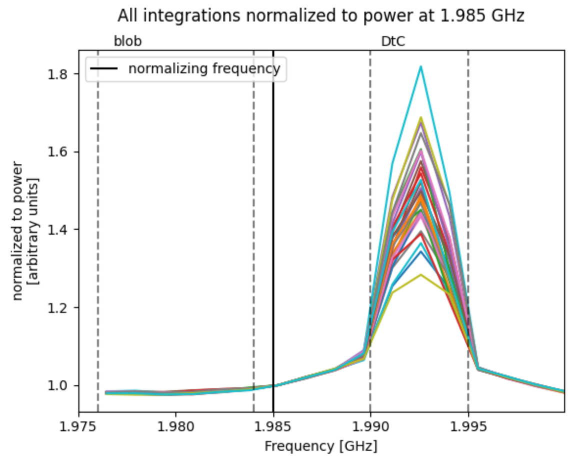

(b) Normalized excerpt of UWBR spectral window.

Figure 3: Left: all integrations from a single scan taken during an RFI scan using the Ku-Band receiver and normalized to unity at 11.4 GHz. Right: all integrations from a single scan of UWBR data taken during test observations and normalized to unity at 1.985 GHz.

against interference that was missed by the initial flagging process. Figure 4 shows the intermediate steps of flagging an RFI afflicted spectrum and identifying clean spectra shortly before and after the flagged integration to use as a replacement spectrum.

Occasionally, strong RFI is present in the first or last integration of a scan. In these cases, it is impossible to identify a clean integration before or after the flag. In these cases, a second clean integration is searched for by going an additional index after or before the first selected clean integration.

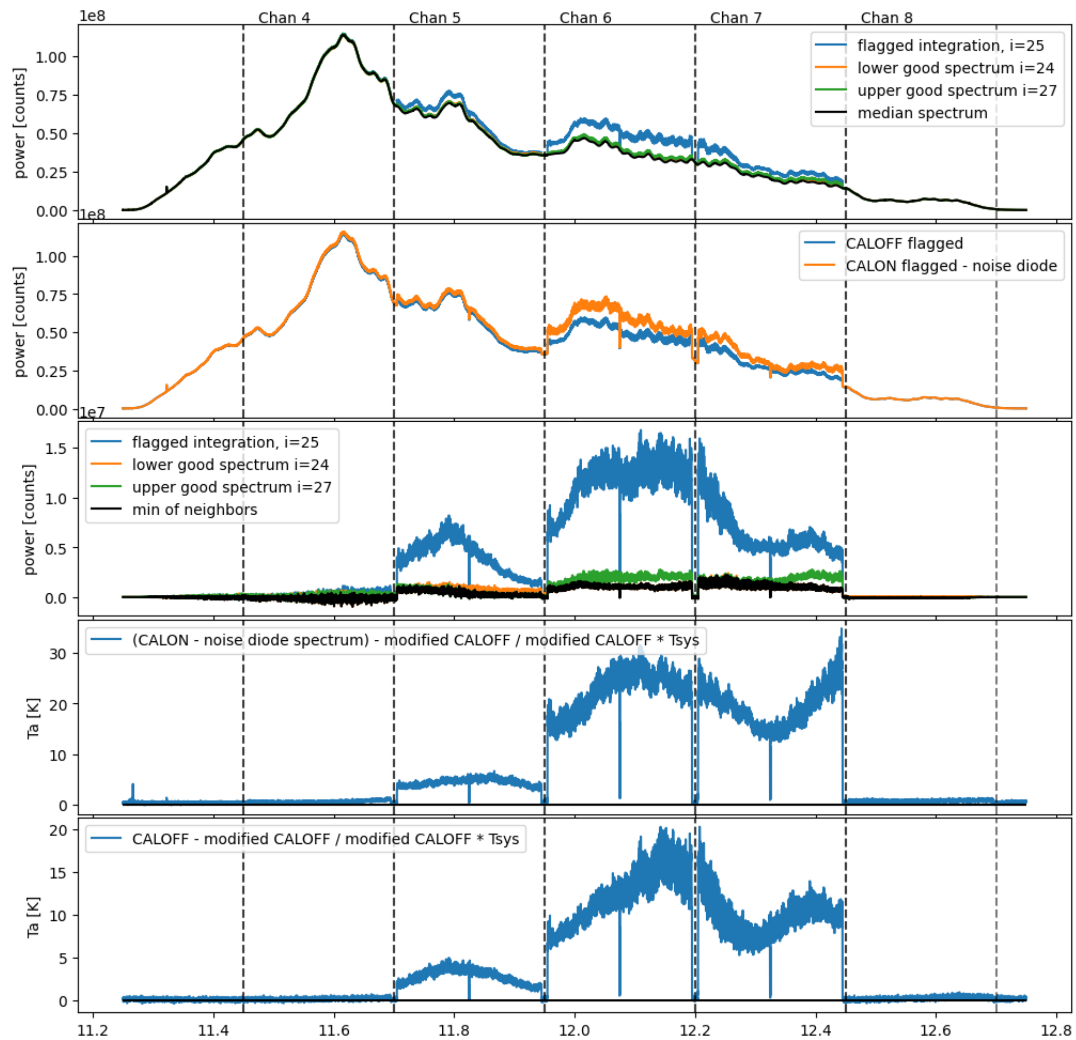

Figure 4: Plots from top to bottom: 1. Total power (counts) plot of the flagged spectrum, the two candidate replacement spectra, and the median spectrum over the duration of the scan. All spectra in this panel are from the calOFF state. 2. The flagged calOFF spectrum, and its corresponding calON spectrum, minus the noise diode contribution. 3. The median spectrum over the duration of the scan has been subtracted from all spectra in this panel. The displayed spectra are the flagged spectrum, the two neighboring, free of RFI spectra, and the minimum of these two good neighbors. If good neighbors have been chosen, their minimum, minus the median spectrum, should be flat and near zero. 4. The calON data calibrated using the standard calibration process. 5. The original calOFF data calibrated using the standard process, with the replacement spectrum, obtained from good neighboring data, used as the reference data.

4 Comparison to Standard Calibration

Appendix Figures 5 - 11 show the progression of calibration techniques starting with raw data and stepping through more advanced methods, finishing with the flag and replace method described in this memo. Each figure is displayed as a waterfall plot, with frequency on the horizontal axis and time increasing up the vertical axis. The spectrum displayed above the waterfall plot is the averaged over the duration of the scan. The plot on the right shows the average power for each integration during the scan. Each set of figures was generated using the same scan data, but treated with different calibration methods. Figures 5 - 8 show a successful flag and replace result, while Figures 9 - 11 show the result of the flagging process on a scan that contains an RFI band that receives strong emissions through the entire duration of the scan.

4.1 Successful Flagging

Appendix Figures 5 - 8 were generated from the data shown in the normalized spectra of Figure 3a. The waterfall plot of the raw data in Figure 5 is dominated by the instrument bandpass, and only strongest emissions can be seen in the waterfall plot. Figure 6 shows the same data, but with the median spectrum subtracted from each integration. This makes the emissions stand out from background, but the data is still in "counts" and the physical units quantifying the emission strengths are not not known. Figure 7 shows the data calibrated using the standard method. It has strong artifacts in both the waterfall and average spectrum, demonstrating the need for a way to calibrate the RFI without its presence detrimentally affecting the calibration. Figure 8 shows the result of flagging and replacing RFI in the reference data before calibrating. This result preserves the structure of the interfering signals in the average spectrum and lacks the "negative" system temperatures that are seen in the previous figure.

4.2 Unsuccessful Flagging

Appendix Figures 9 - 11 were generated from the data shown in the normalized spectra of Figure 3b. The median subtracted data from this scan was left out because it did not provide any new information to discuss. The raw waterfall plot of Figure 9 shows strong emission spanning 1990 - 1995 MHz for the duration of the scan. It is this feature that makes it impossible to identify an RFI free integration to use as the reference data. Any subsequent calibration will contain RFI emissions in both the signal and reference data, leading to artifacts in the calibrated data. Figure 10 shows the data calibrated to antenna temperature using the standard calibration method. The antenna temperatures between 1990 - 1995 MHz frequently dip into "negative" temperatures. This also leads to a dip in the average spectrum. Figure 11 shows the result of calibrating the data after flagging and replacing the RFI. This waterfall also has the "negative" temperatures, but the average spectrum remains positive. Despite the lack of "negative" antenna temperatures, the calibrated temperatures of the 1990 - 1995 MHz RFI should not be trusted. The flag and replace method was able to identify the worst offenders, but for this dataset, the RFI was so dense that the replacement data also contained strong RFI in that range.

5 Conclusions

The flag and replace method for calibrating RFI emissions performs well when the interfering emissions are intermittent in time, but the method breaks down when the RFI is persistent and fills the frequency range for the duration of the scan. Therefore it will be important to compare the cleaned and calibrated data to the raw data. The code is available on github at: https://github.com/dbautista98/water-dysh. It comes with the ability to generate waterfall plots of raw (counts) data, median subtracted (counts) data, and calibrations with or without the RFI flagging and replacement.

6 Acknowledgments

The National Radio Astronomy Observatory is a facility of the U.S. National Science Foundation operated under cooperative agreement by Associated Universities, Inc. This work is supported by the NSF's SII-NRDZ (AST-2232159) grant. The author would like to thank many individuals at NRAO and GBO for fruitful conversations: Karen O'Neil, Pedro Salas, Evan Smith, Jay Frothingham, Brenne Gregory, Chris De Pree, Bang Nhan.

References

Braatz, Jim (Oct. 2009). *Calibration of GBT Spectral Line Data in GBTIDL v2.1*. URL: https://www.gb.nrao.edu/GBT/DA/gbtidl/gbtidl_calibration.pdf.

A Successful Flagging

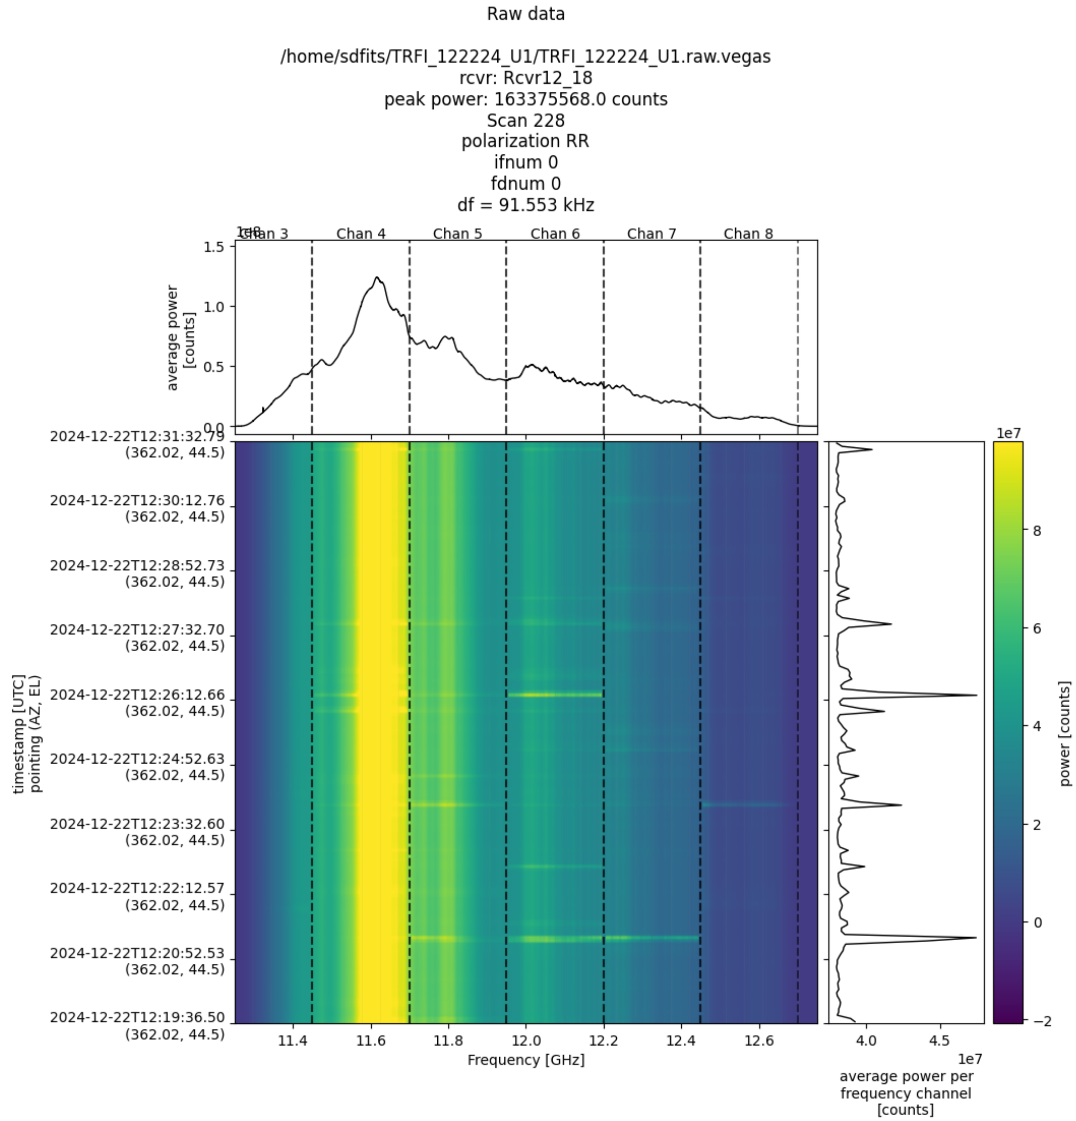

Figure 5: Raw calON data displayed as a waterfall plot. This spectrum contains the contribution of the receiver bandpass, the noise diode, and the signal received from the sky (including RFI sources).

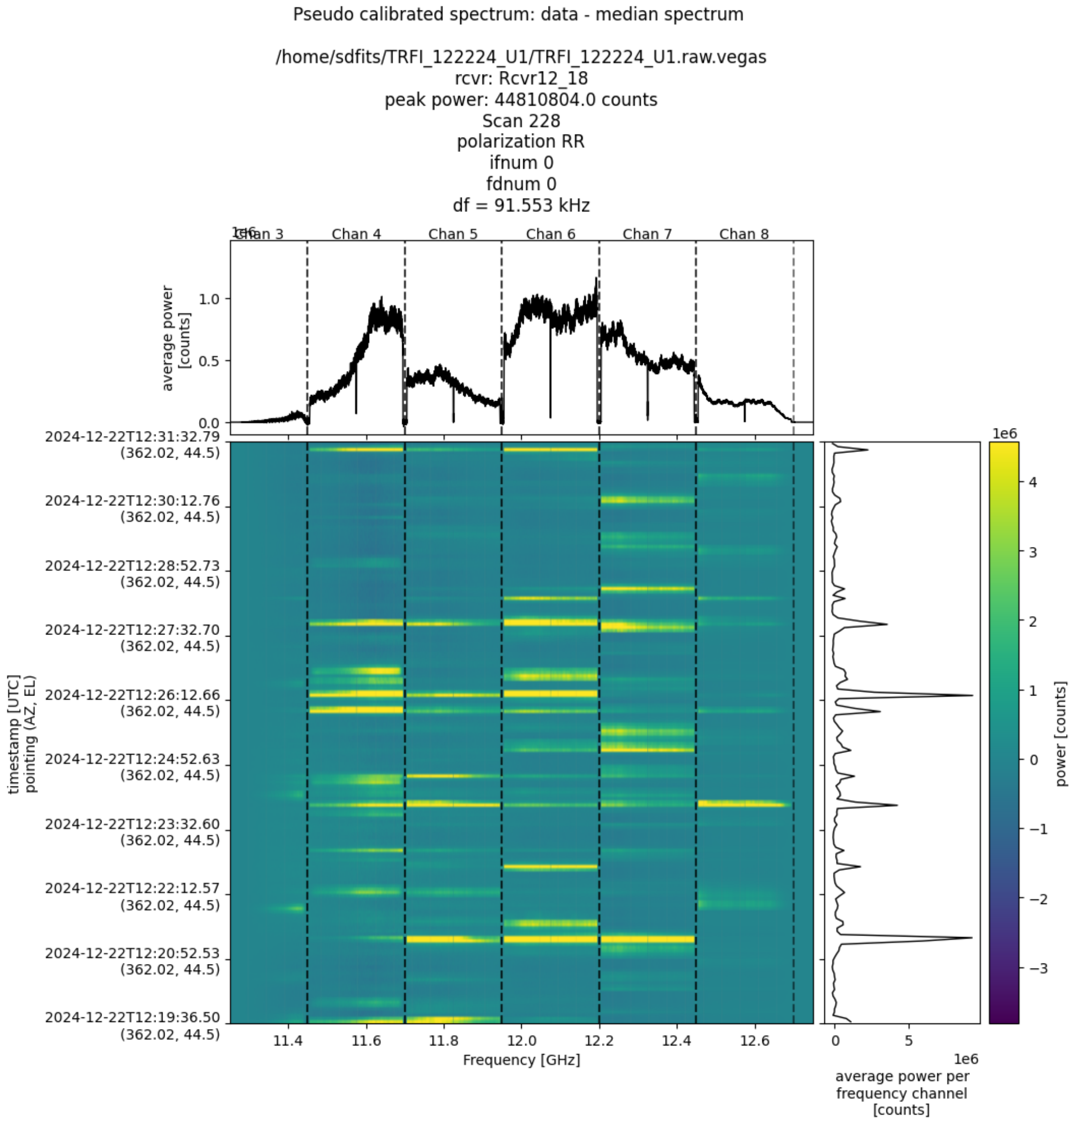

Figure 6: Raw data with the median spectrum subtracted from each integration. This was an early method of analyzing time-varying RFI, as it makes the interfering signals stand out more than they would in the raw data. However this data has not been calibrated to physical units, so the displayed units are still “counts.”

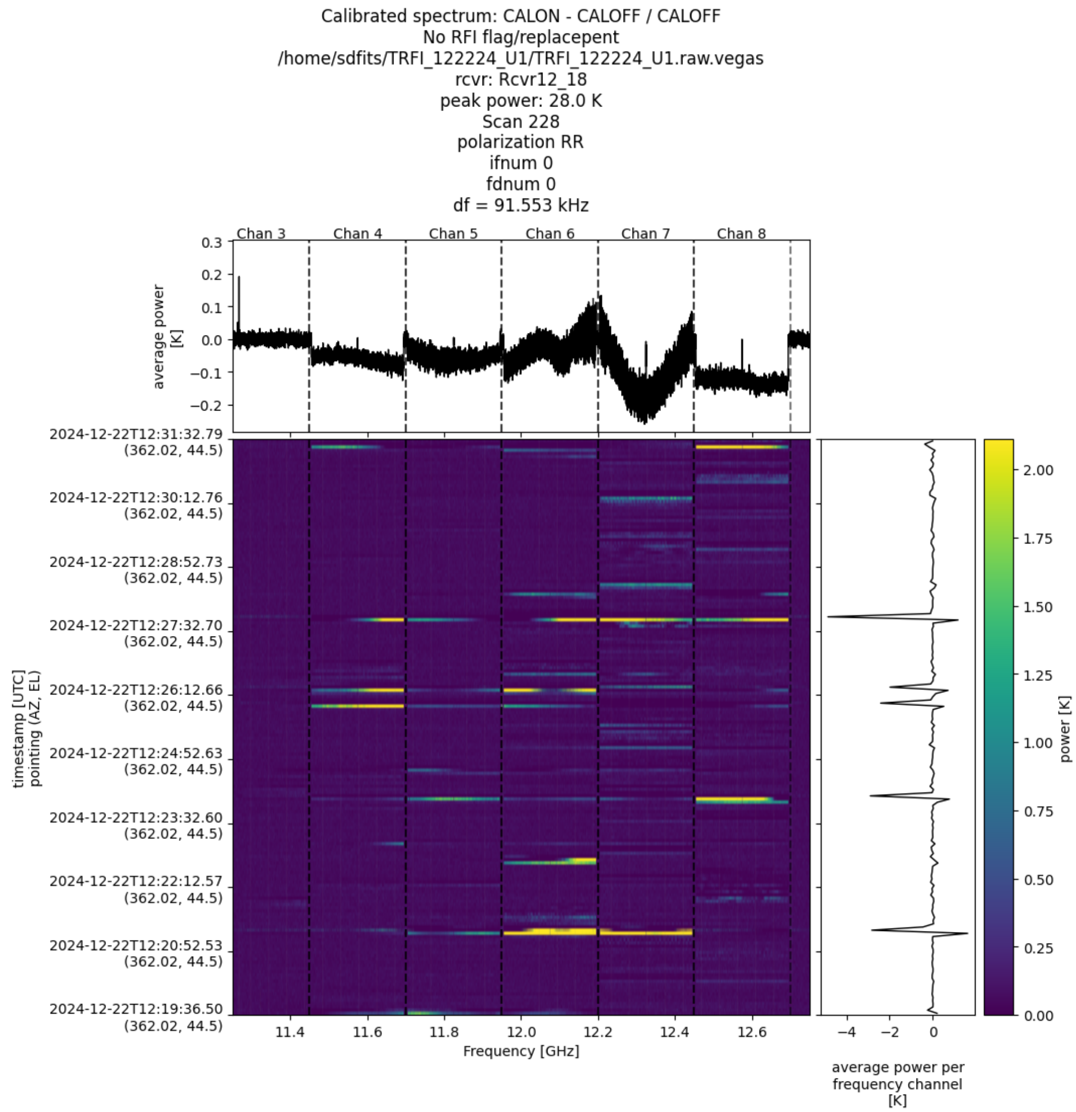

Figure 7: Data calibrated with the standard $\frac{ON-OFF}{OFF} \times T_{sys}$ method. Individual integrations containing RFI can be seen in the waterfall plot as yellow or purple stripes, and in the time-series power plots as spikes above or below the average line. The instability across the average spectrum is indicative of interference during the scan, but the structure and strength of the emissions cannot be seen due to the averaging of positive and negative system temperatures.

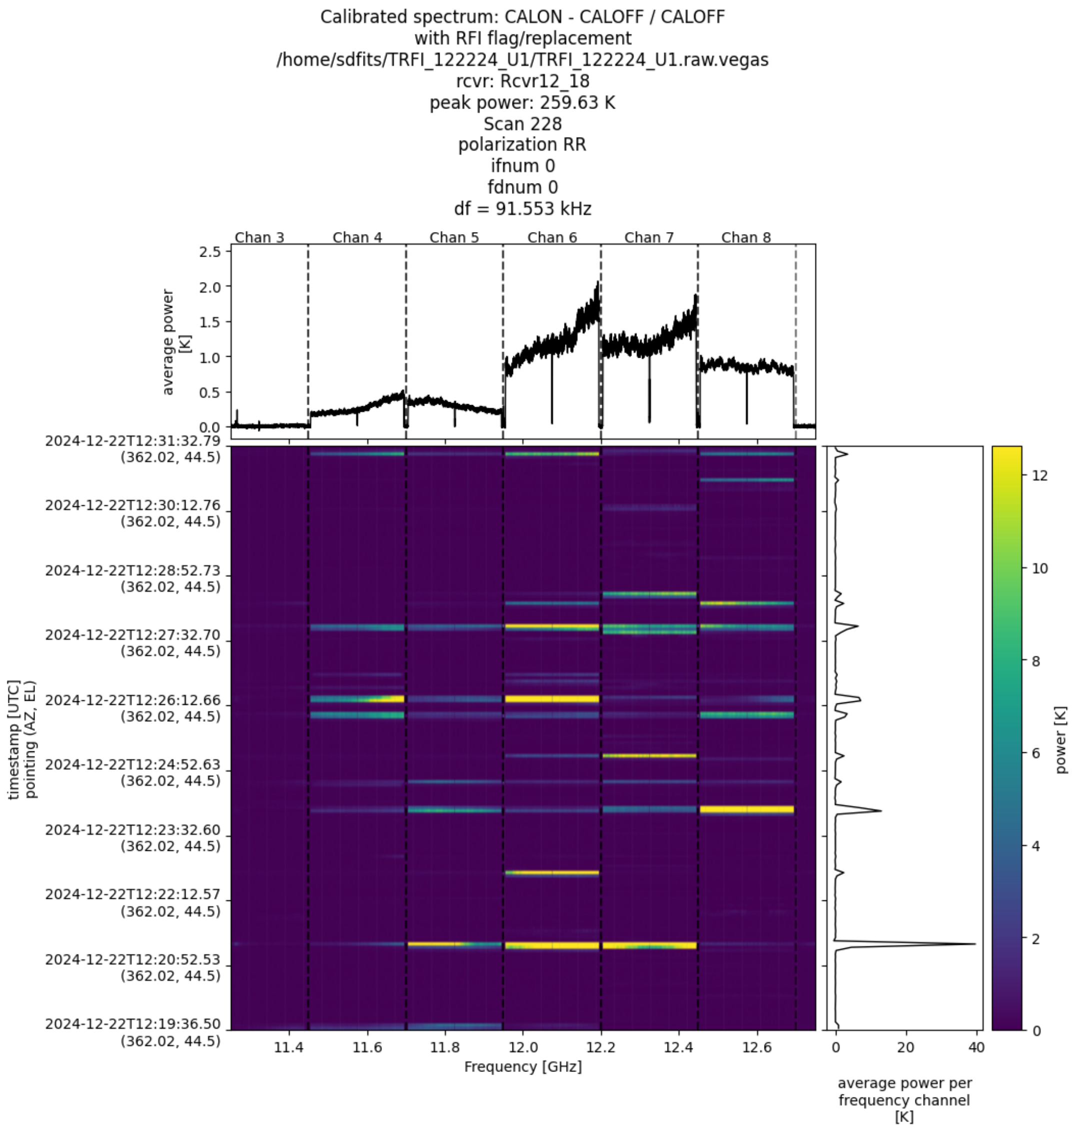

Figure 8: Data that has been treated with the RFI flagging and replacement prior to calibration. The structure of the interfering signals has been preserved and is visible in the average spectrum. The antenna temperature does not dip into “negative” values.

B Unsuccessful Flagging

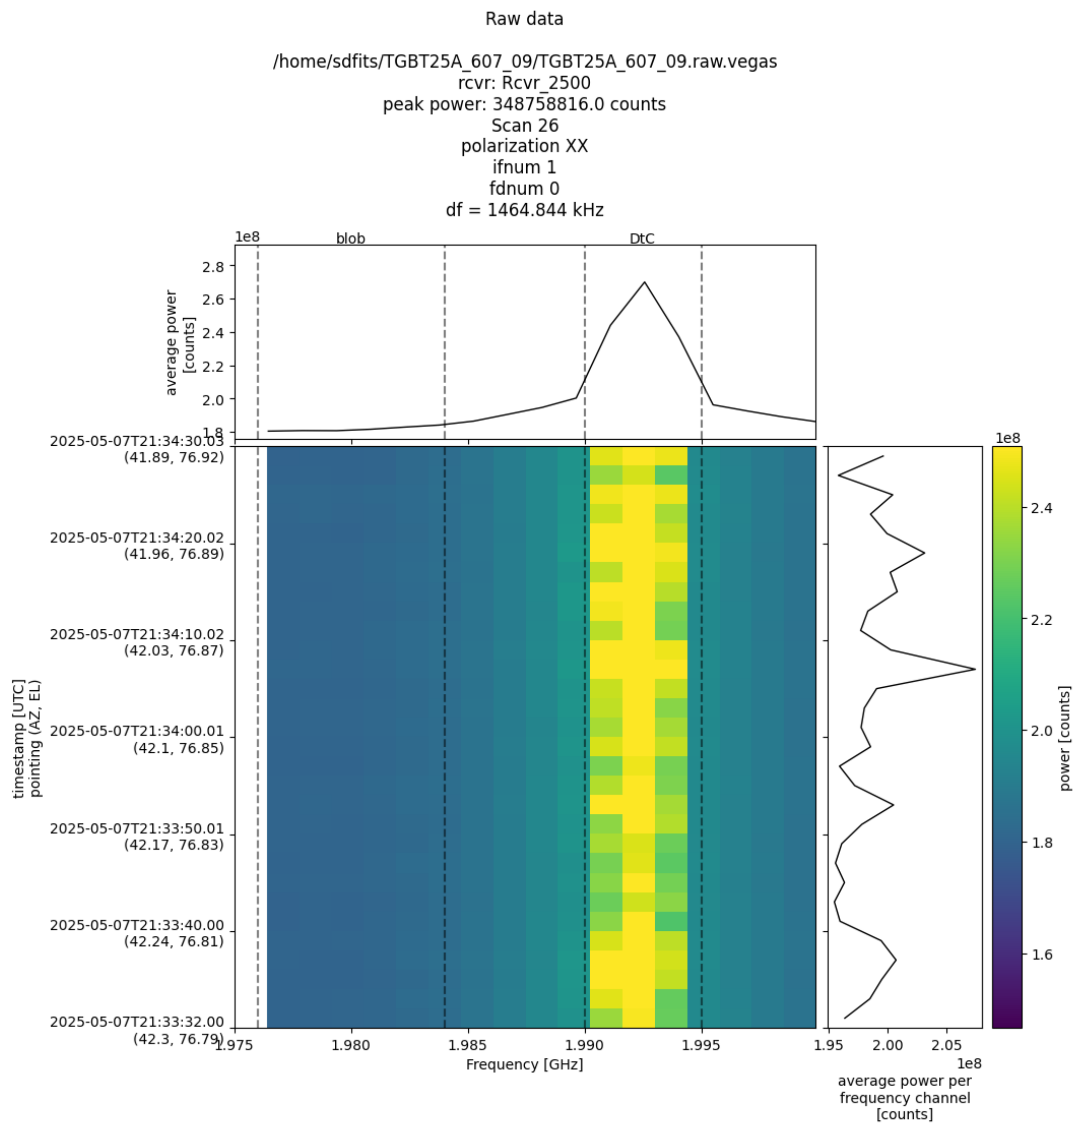

Figure 9: Raw calON data displayed as a waterfall plot. This spectrum contains the contribution of the receiver bandpass, the noise diode, and the signal received from the sky (including RFI sources). Strong emissions can be seen between 1990 - 1995 MHz throughout the duration of the scan.

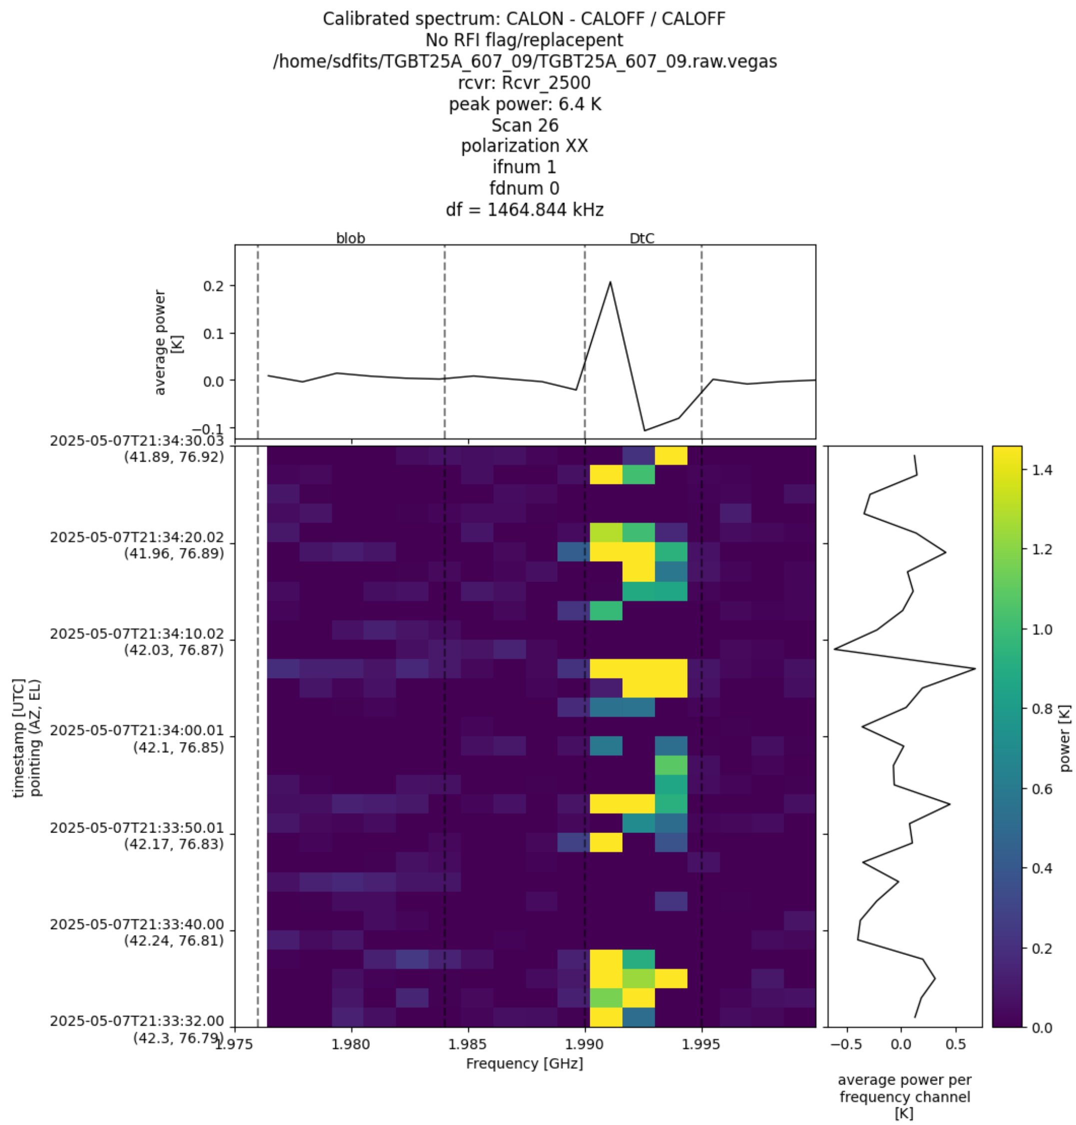

Figure 10: Data calibrated with the standard $\frac{ON-OFF}{OFF} \times T_{sys}$ method. In the waterfall plot, RFI between 1990 - 1995 MHz can be seen fluctuating between positive and "negative" antenna temperatures. The average spectrum in that range also dips down into the "negative" values.

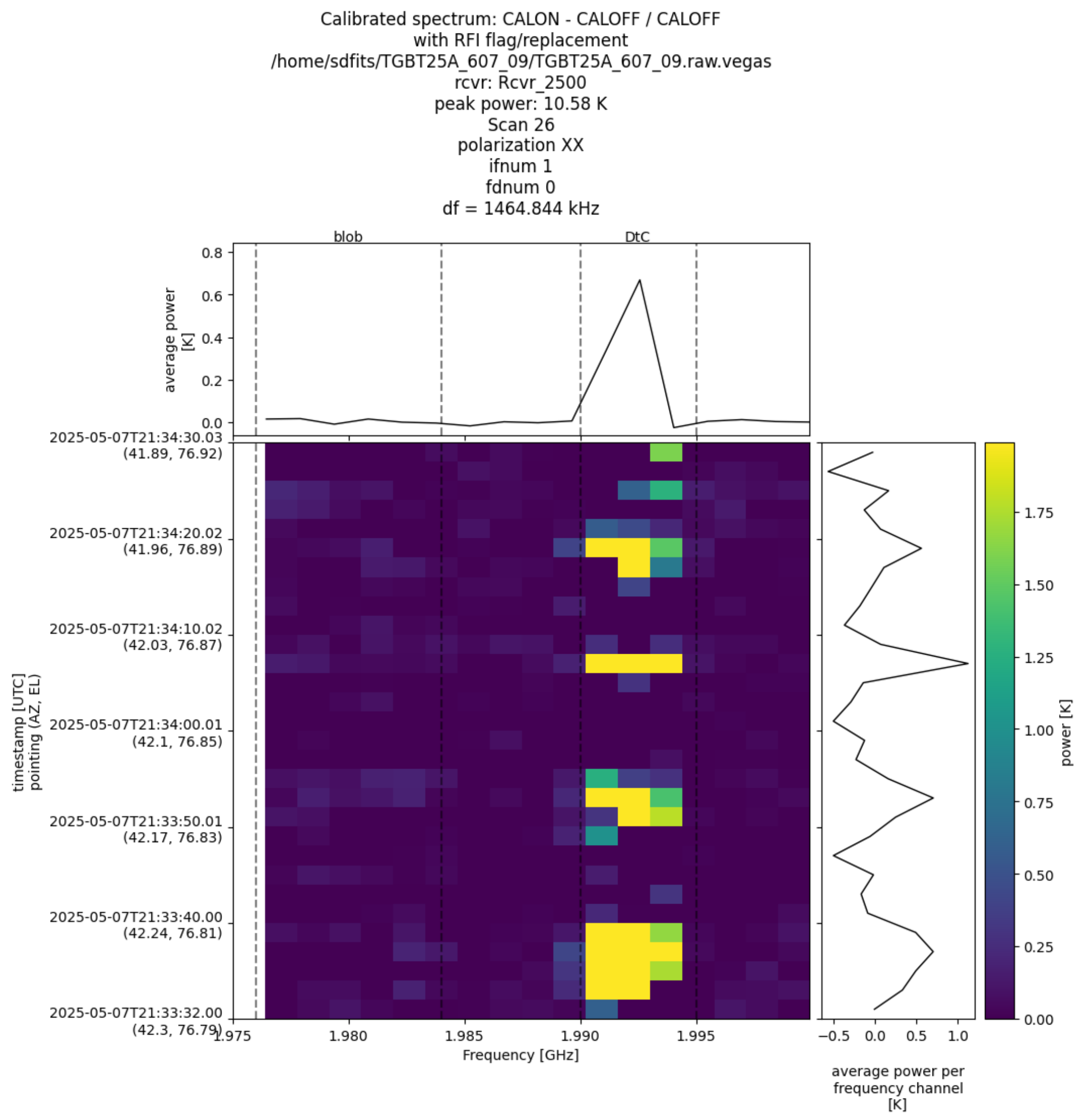

Figure 11: Data that has been treated with the RFI flagging and replacement prior to calibration. The RFI between 1990 - 1995 MHz still fluctuates between positive and "negative" antenna temperatures. The average spectrum in this range does not dip down to "negative" values, which is a slight improvement over the standard method.