National Radio Astronomy Observatory
Green Bank, WV

Cryostat Cavity Noise and the Impact on Spectral Baselines

Electronics Division Internal Report
No. 318

Roger D. Norrod
April 26, 2007*

Abstract

From the time of construction, the GBT Ka-band and Q-band receiver front-ends exhibited a spectral baseline instability of unknown cause. This report describes a methodical investigation into the problem. The cause is found to be due to thermal noise radiated from the cryostat interior cavity surfaces. Due to multimoding and multiple reflections in the cryostat cavity, the noise radiation exhibits complex frequency structure sensitive to environmental factors. This radiation can couple into the receiver input signal via leaky waveguide flanges, leaky component bodies, or reflections off of the imperfectly matched vacuum windows. The detrimental baseline effects can be mitigated by using microwave absorber to smooth the cavity noise ripple, and by shielding the signal path from the noise radiation.

* This version corrects two minor errors from the original: The absorber part number is HR10, and the Kerr calculation of mode spacing was for a 50cm cube.

Table of Contents

List of Figures

1 Introduction

One of the primary fields of radio astronomy is spectrometry, and the 100-meter Robert C. Byrd Green Bank Telescope is extremely active in this field. In spectrometry the astronomer is generally trying to detect quite weak spectral features (e.g. molecular lines), much smaller in amplitude than unavoidable frequency variations in system noise temperature and gain. Detecting these weak signals requires observing techniques that cancel out the larger systemic variations, and to be effective all these techniques require extraordinary stability of the system gain and noise temperature. For typical GBT observations the rms noise in an observed spectrum is roughly 2 parts in 10^4 for a one-minute integration, and the stability must be several times better than this value over several minutes to avoid masking the target signals.

Many causes can degrade the spectral stability, and much effort has gone into finding and correcting problems in the GBT receiving systems[1]. During commissioning of two GBT receivers, the 40-50 GHz Q-band and the 26-40 GHz Ka-band, a spectral instability was noted that was not seen with other receivers, had unique characteristics, and had an unknown cause. This document reports on the search for the cause of this unique effect, explains how it arises, and how it can be eliminated.

2 Background

A radio astronomy spectrometer is a specialized data acquisition system designed to produce power spectra averaged over a user-specified time period. The main spectrometer at the GBT is an auto-correlation device and has at the heart 256 auto-correlation ASICs, each with 1024 lags. Although many modes are available, one common GBT Spectrometer observing configuration accepts one or more IF signals and produces 800 MHz wide output spectra, with 2048 frequency resolution elements (390 kHz resolution). This is one of the more challenging modes, due to the broad bandwidth, and is used throughout this investigation. For the laboratory tests described in this report, our Lab Spectrometer, which uses the same ASIC and has design similar to the GBT Spectrometer, was used.

One of the simplest observing procedures is called *position switching*, in which the telescope tracks a target source for a period of a few minutes, and then the telescope position is changed to a nearby off-source reference position tracked for the same duration. During these tracking periods the system output spectra is separately measured, averaged, and stored. Figure 1 illustrates two spectra, $S_1(f)$ and $S_2(f)$, similar to what might be observed on the telescope at the on-source and off-source positions, respectively. Since a desired astronomical signal will typically be much smaller than the frequency variations seen in the detected spectra, it is necessary to cancel out the systemic variations by calculating the spectral baseline function $B_L(f)$:

$$B_L(f) = \frac{S_1(f) - S_2(f)}{S_2(f)} \tag{1}$$

The system output power is given by:

$$S_i(f) = kBT_i(f)G_(f) \tag{2}$$

where k is Boltzmann's constant, B is the detected bandwidth, and G is the system gain. T_i is the sum of the effective sky temperature, antenna temperature (spillover, etc.), and receiver noise temperature. If the receiver system is working well and the gain is constant during the measurement period, then substituting (2) into (1) yields:

$$B_L(f) = \frac{T_1(f) - T_2(f)}{T_2(f)} \tag{3}$$

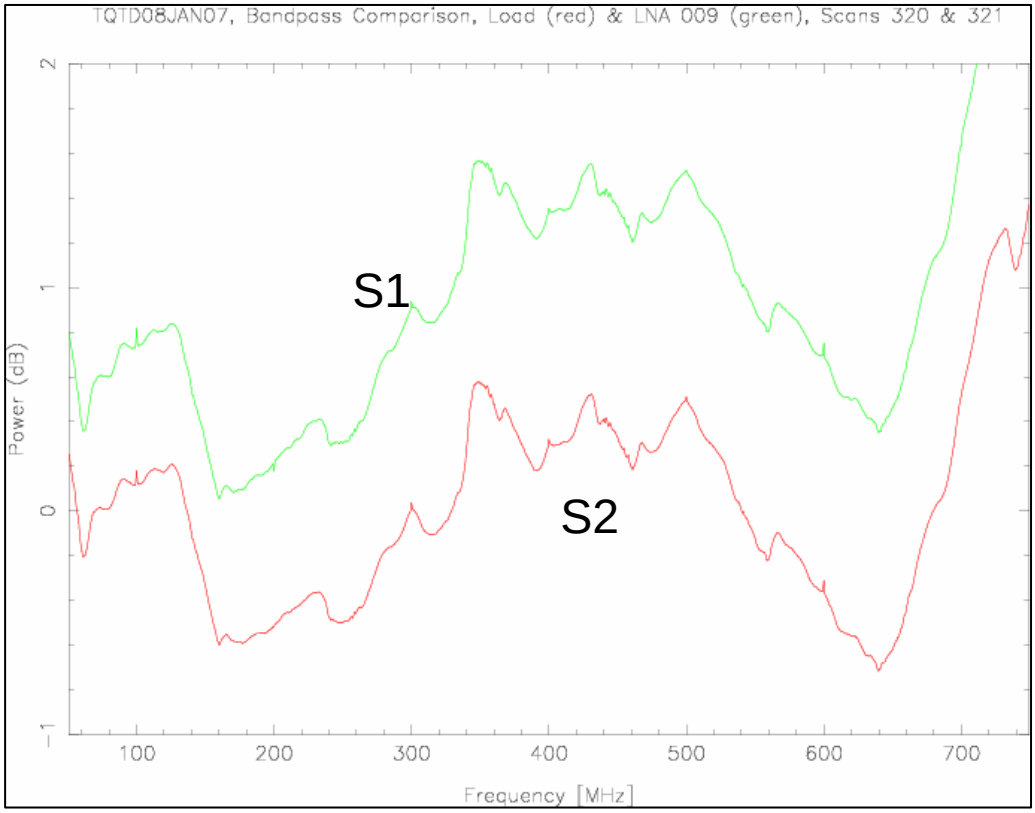

Figure 1: Typical power spectra measured by the GBT Spectrometer or Lab Spectrometer.

In astronomical observations, additional steps are needed to calibrate the results, but this simple function is convenient when evaluating instrumentation. The baseline function as given by (3) is a measure of $\Delta T/T$, and the expected rms variation, with time at each frequency, can be calculated using the radiometer equation.[2] The output system noise power at any frequency should be independent of that at any other frequency, so the radiometer equation should also govern the rms noise variation across the frequency range for any given time period. That is, the baseline function should be flat with frequency, except for random noise variations, and the amplitude of the noise variations

should average down with the square-root of time. Figure 2 shows a baseline function measured on the GBT Q-band receiver approaching this ideal.

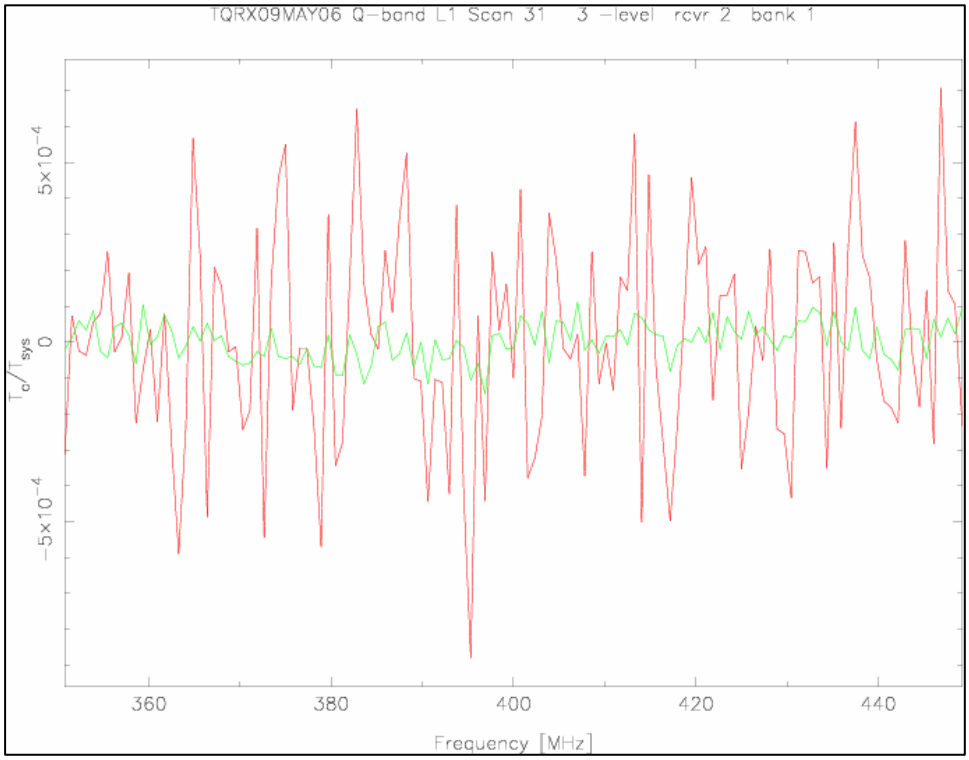

Figure 2: A baseline function; 1 minute (red) and 1 hour (green) averages.

In order to evaluate a receiver's performance, it is useful to simply take a long series of spectra with the receiver input terminated to a constant temperature, or to track a constant sky position on the telescope. For each pair in the series, the baseline function is calculated using (3) and plotted in a waterfall plot. Figure 3 shows such a plot measured on one channel of the Q-band receiver. Each trace in this plot is the baseline function for one pair of one minute averaged spectra. A mean value is subtracted from each baseline removing any constant offset, the first pair's baseline is plotted at the bottom, and each subsequent baseline is offset upward by a small amount. One hopes that each baseline in the waterfall will be flat with frequency, except for the random noise variation. The performance in Figure 3 is fairly good, except for the occasional trace that shows ripple a few times larger than the noise (e.g. the orange trace near 0.02 on the vertical scale). Still, the rms is within ten or fifteen percent of theoretical most of the time, and integrates down, as expected, for an hour or two. Unfortunately, even this level of quality was not typical for the Q-band receiver.

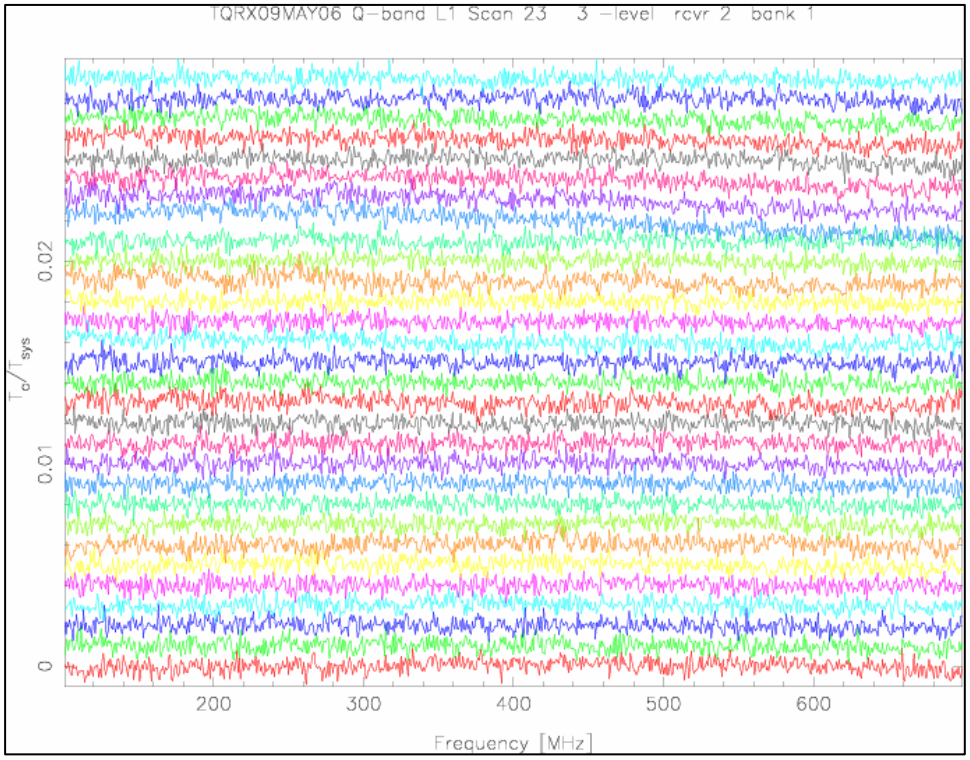

Figure 3: Baseline waterfall plot showing good performance.

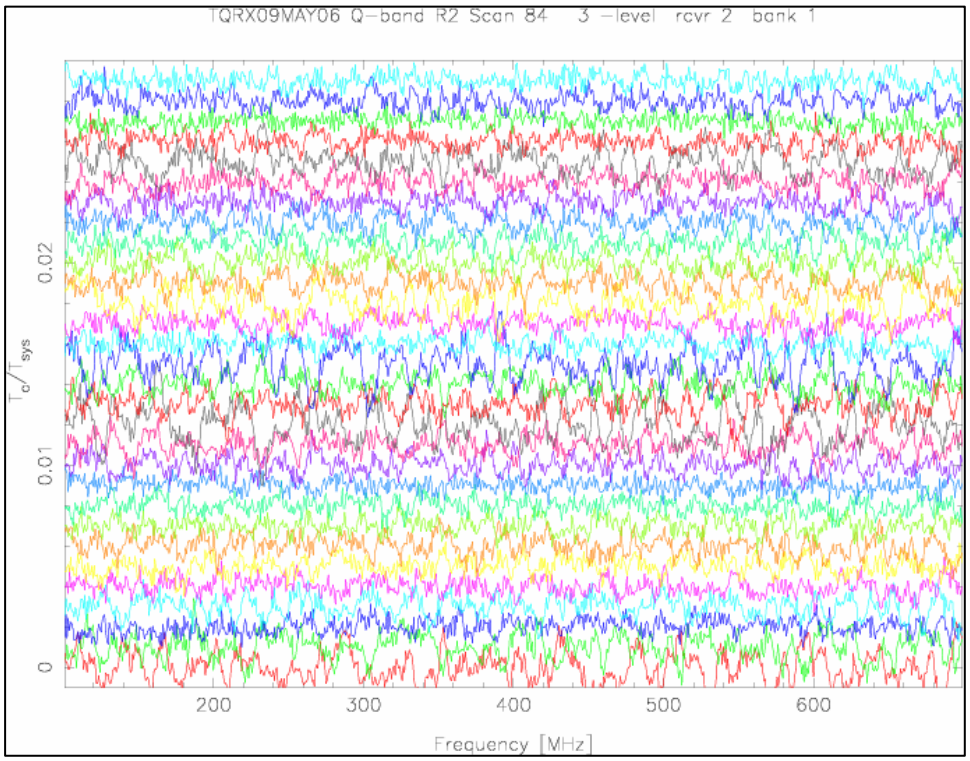

Figure 4: Baseline waterfall showing poor performance.

Figure 4 shows a waterfall for another Q-band channel. Note the irregular ripples that wax and wane from trace to trace. Extensive testing during Q-band commissioning and afterward showed that the strength of this behavior varied with time and with receiver channel. During commissioning observations, evidence showed that the Ka-band receiver showed similar behavior at even larger magnitude. Causes related to the GBT optics were quickly ruled out, because the problem was clearly visible with absorber over the receiver vacuum window. (Both the Q-band and Ka-band receivers have cooled feedhorns that view the subreflector through a vacuum window, and are unique among current GBT receivers in this respect.) A Fourier transform on the baselines showed the spectral power of the ripples to be concentrated in the 10-30 MHz region. The slow time variations, irregularity, and relatively short periods presented a unique signature and ruled out simple reflections in the receiver transmission lines - such short periods require path lengths much longer than are physically present.

During the summer of 2006, an intensive effort was undertaken to understand and correct the baseline problems seen in the Q-band and Ka-band receivers. Many possible causes were suggested (LO noise, amplifier bias oscillations, microphonics, spectrometer problems …) – all were investigated and ruled out. However, progress was made. The findings as of October 2006:

- The problem arose right at the input to the receiver, in or before the cryogenic low-noise amplifier (LNA).
- The magnitude varied from channel to channel, and the problem followed specific amplifiers, at least in two cases.
- The irregular ripples were only seen when the system was cold (below 150K).
- Rise and fall of the ripple amplitude tended to be correlated across multiple receiver channels.
- In $\Delta T_{sys}/T_{sys}$, the ripple was roughly of the same amplitude for several values of T_{sys}, where the variable portion of T_{sys} was the receiver input source temperature.
- By selecting amplifiers (Q-band), tightening loose waveguide flanges (Ka-band), and adding absorber in the cryostats, we were able to get both receivers to acceptable performance.

In October 2006 both receivers showed greatly improved performance and were installed on the GBT for the winter observing season. However, the situation was not satisfactory because the source of the problem was not at all understood. The fact that the magnitude of the problem seemed to follow certain LNA units led to much speculation about how it could arise. The LNA 1/f gain fluctuations came under strong suspicion, but mostly because it was thought to vary from unit to unit and because resources were not available to measure the 1/f performance independently. Another hypothesis was that noise waves flowing out of the LNA input port, reflected off the vacuum window and coupled to the cryostat cavity, could contribute to the instability. However, none of the speculated causes fit all the observed behaviors or seemed likely. So, it was decided to assemble a test cryostat setup and continue the investigations after the GBT receivers were reinstalled, using components on hand and the Lab Spectrometer.

3 Test Setup

It was thought that a comparison radiometer was the approach needed, as it allows performance comparisons between well understood source components (i.e. waveguide loads) and devices to be measured. Hence a mechanical WR22 waveguide transfer switch operable at 15K was located and built into a test setup as shown in Figure 5. While the mechanical switch only allows slow switching, it was felt that comparisons at periods of a few minutes were sufficient.

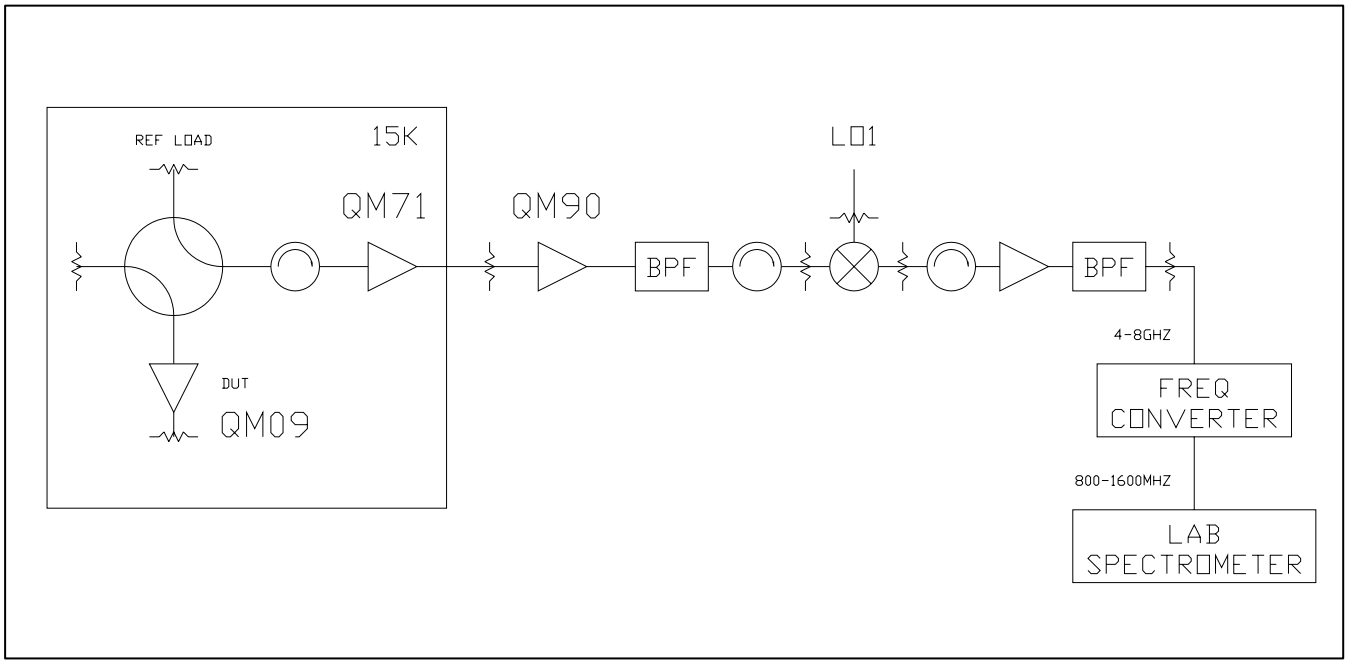

Figure 5: Comparison radiometer and test setup.

Figure 6 is a photograph of the cryostat interior as initially setup. The large black component is the transfer switch. A fiberglass tube extends from the bottom of the switch to a rotary feed through the cryostat bottom plate, where a knob allows manual operation of the switch. Clockwise from the left switch port in this view is the device under test, the receiver LNA (QM90), the reference waveguide termination, and the off-circuit termination. Installation of a radiation shield and the cryostat outer shell completes the setup and allows cooling of the system.

Using a waveguide load at an elevated temperature (~45K) at the DUT (Device Under Test) port, it was possible to do a hot/cold noise temperature measurement of the radiometer. With this information it is possible to determine the effective source temperature of any DUT.

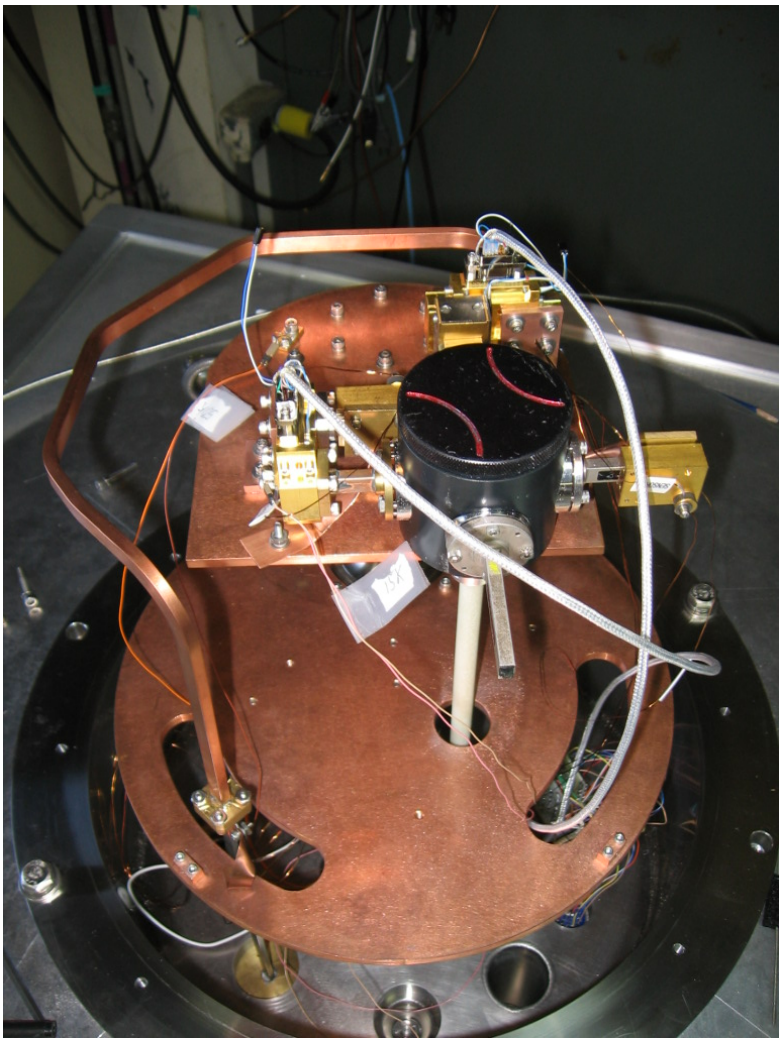

Figure 6: Photo of the initial LNA test setup cold electronics.

4 Initial Results

4.1 LNA Input Port Noise

To characterize the noise power flowing out of a LNA input port, one of the Q-band LNA units was connected to the transfer switch DUT port. (This is the configuration shown in Figure 5 and Figure 6). The setup allowed characterization of the noise power flowing out of the LNA input port: mean effective source temperature, frequency structure, and stability.

Figure 7 and Figure 8 show representative waterfall plots for the reference load and the LNA input port tests respectively, measured at 42 GHz. Two and three hours of data (series of one minute scans) were measured on each of the two sources. For a qualitative comparison, it is useful to calculate the standard deviation of each baseline function across the total bandwidth, divide by the theoretical rms deviation, and plot versus baseline (scan pair) number. Figure 9 shows the result for all the data taken on the

reference load (red) and the LNA input port (green). The LNA input port mean effective source temperature at 42 GHz was found to be about 30K, which seems reasonable. But, there is no significant difference in the baseline stability between the two sources, and so we conclude the LNA input port noise was not the source of the baseline instabilities.

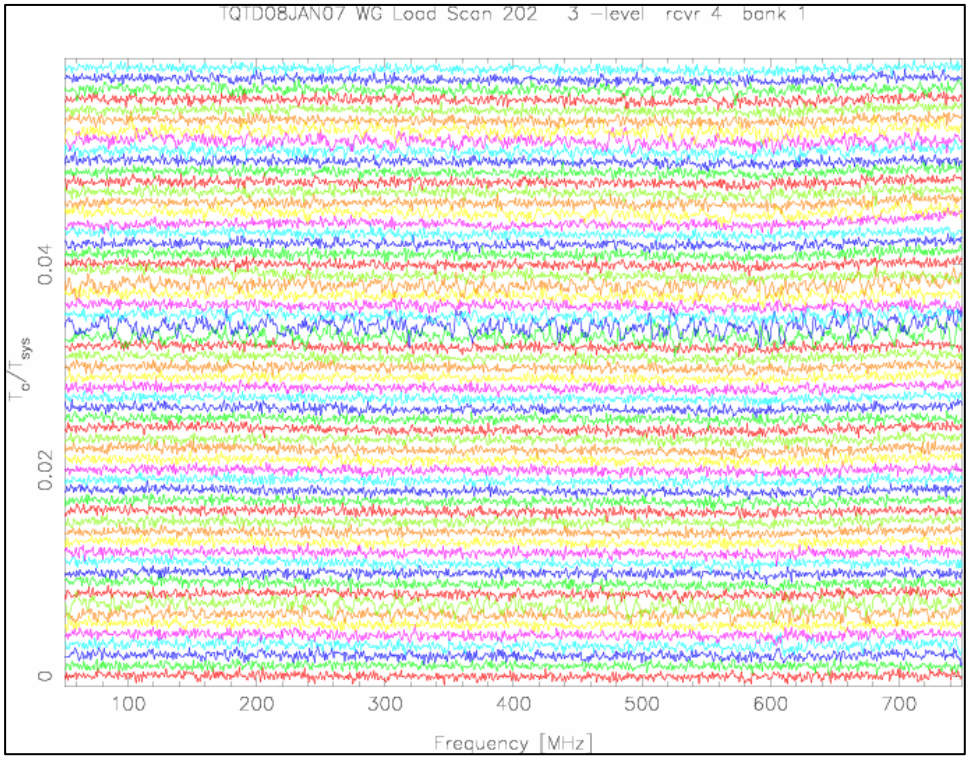

Figure 7: Reference load connected to radiometer.

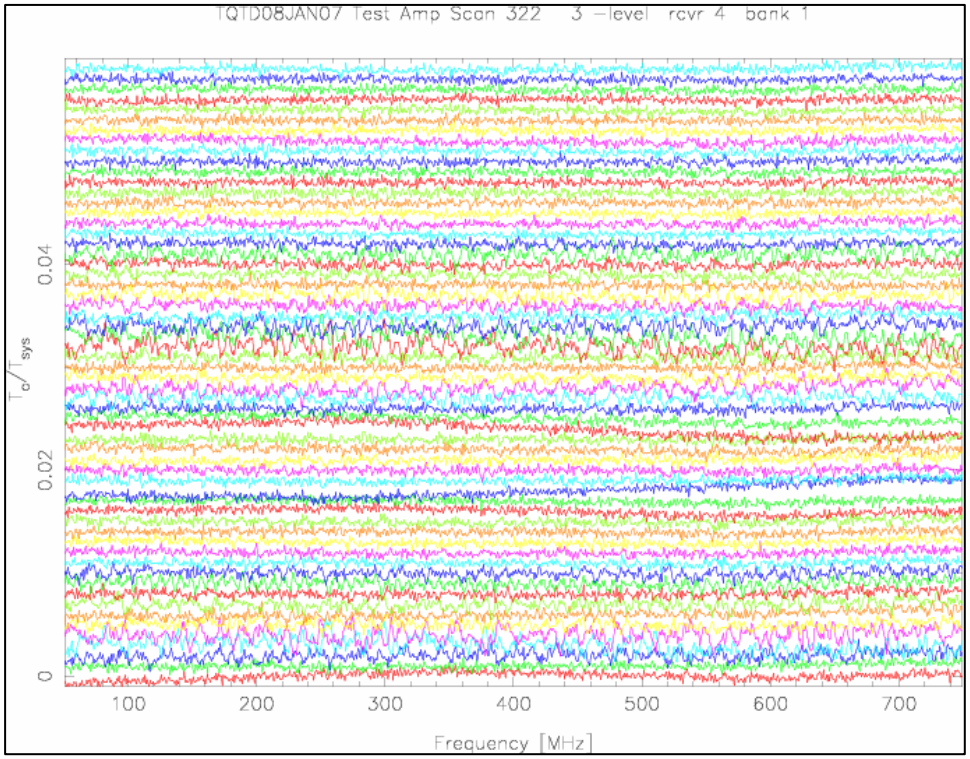

Figure 8: LNA input port connected to radiometer.

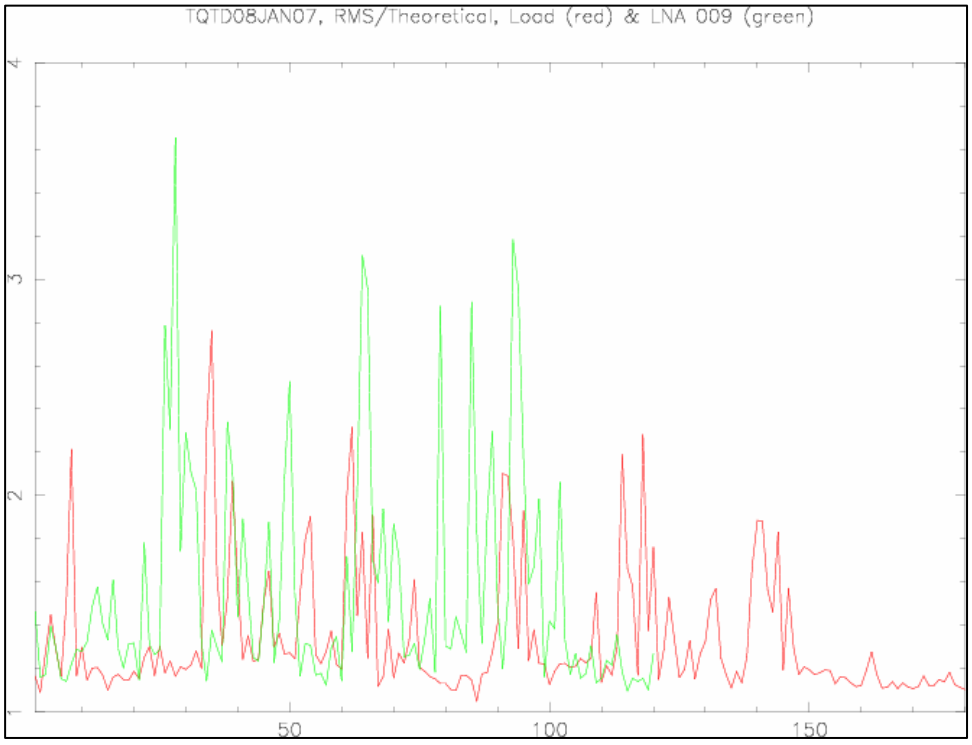

Figure 9: RMS comparison, Load and LNA.

4.2 Open-ended Waveguide

The results above led to the hypothesis that white noise flowing into the cryostat cavity from a LNA input port or an isolator load, as the case may be, could somehow acquire short-period frequency structure and get reflected back into the receiver in an unstable way, perhaps influenced by environmental factors. To test this idea, the DUT LNA was removed and replaced with a waveguide attenuator (~3dB value), with one end open to the cryostat. Test results with this configuration were remarkable in two ways. First, the Y-factor measured between the reference load and the open-ended waveguide attenuator was 4.5 dB. Second, the baseline ripple instability was several times larger than previously observed, yet with ripple periods consistent with that always seen. Noise power from the attenuator was 4.5 dB higher than power from the reference attenuator at 15K, implying an effective source temperature of about 70K! Noise power flowing out of the isolator at the radiometer input and from the cold attenuator should be about 20K, so the much larger measured level implied that the warm cryostat wall surfaces were radiating a large noise power, and this was efficiently coupling into the receiver input impedance - not an expected result. Of course, switching between the well-matched reference load and the probably not well-matched open-ended waveguide could lead to LNA gain changes and measurement errors, but the input isolator should minimize this effect and subsequent results under many conditions convinced us gain changes were not a significant source of measurement errors.

Because the waveguide attenuator value and physical temperature were not well known and had strong influences on the test results, a decision was made to simplify the setup. The open-ended waveguide attenuator was replaced with a simple pyramidal horn.

4.3 Horn

Figure 10 is a photo of the setup with a pyramidal horn on the switch DUT port. The setup at this point included a 70 K radiation shield open at the top and the horn illuminating mostly the 300 K cryostat top plate. Seen in Figure 11 is the cryostat with a heater pad strapped to the outer cylinder, used to induce a temperature rise in the shell, and testing for presence of an environmental link to the baseline instabilities. Figure 12 shows a resulting baseline waterfall. The ripple amplitude was clearly larger than previously seen. Figure 13 plots the baseline Rms/Theoretical for three data sets. The red trace is for the reference load (compare the vertical scale of this figure with that in Figure 9). The green trace is for the horn, with no intentional perturbations into the system, and the blue trace shows what happens when heat is applied to the cryostat outer cylinder. Power was applied to the heater pad during scan pairs 5 through 10, and again during pairs 20 through 23. A temperature sensor on the cylinder outside the heater edge registered temperature changes associated with these heat applications of 17 C and 8 C, respectively.

Figure 14 shows the spectra measured from the reference load and from the horn, centered at 42 GHz. It was necessary to increase an IF attenuator by 7 dB when switching to the horn in order to keep a satisfactory level into the spectrometer sampler,

and the plots in this figure have also had a mean value subtracted so that the traces overlay. The mean total power from the horn was 7.2 dB higher than from the reference load (at 19 K), giving the effective horn temperature as 181 K.

It is clear that the majority of the noise power seen by the horn arises in the cryostat, not flowing out of the isolator load, and the ripple in the horn spectra changes with the temperature of the cryostat walls. The cause of what is described here will be discussed in Section 5, but first the result of another test is presented.

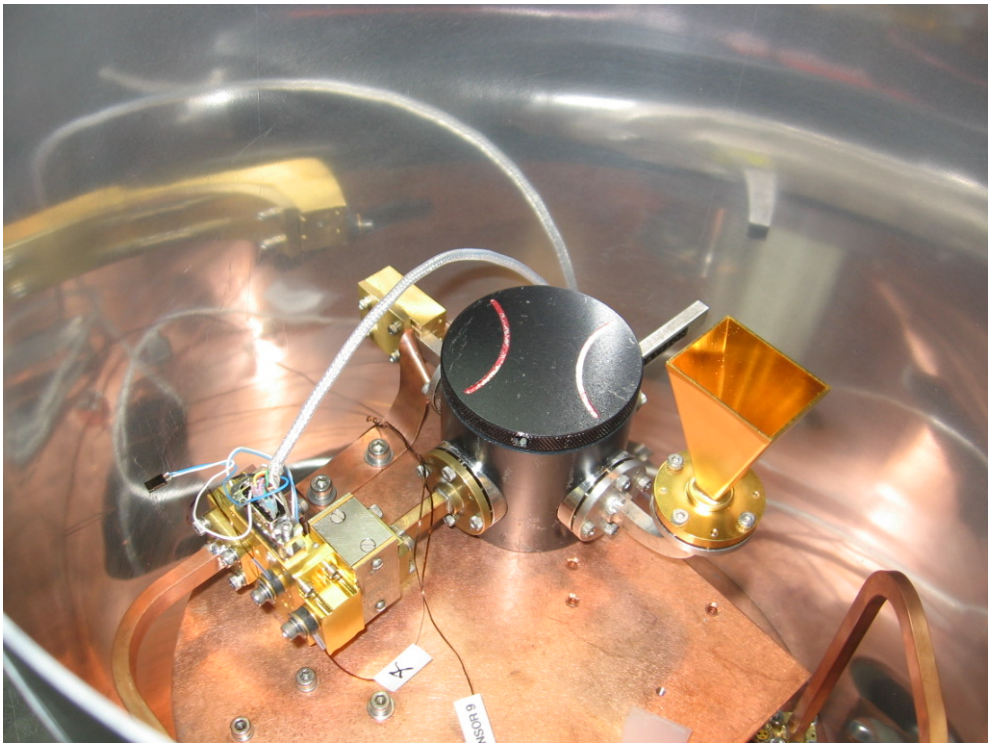

Figure 10: Photo showing pyramidal horn in setup.

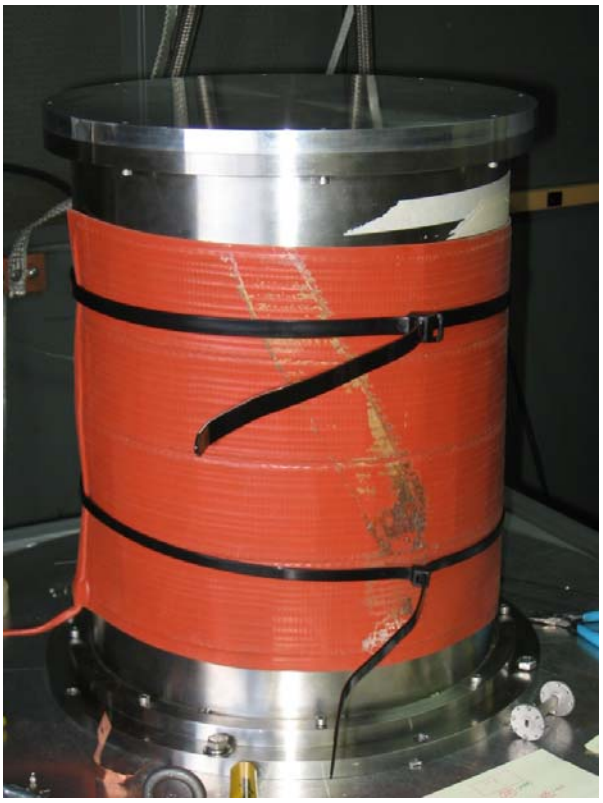

Figure 11: Test cryostat with heater pad attached.

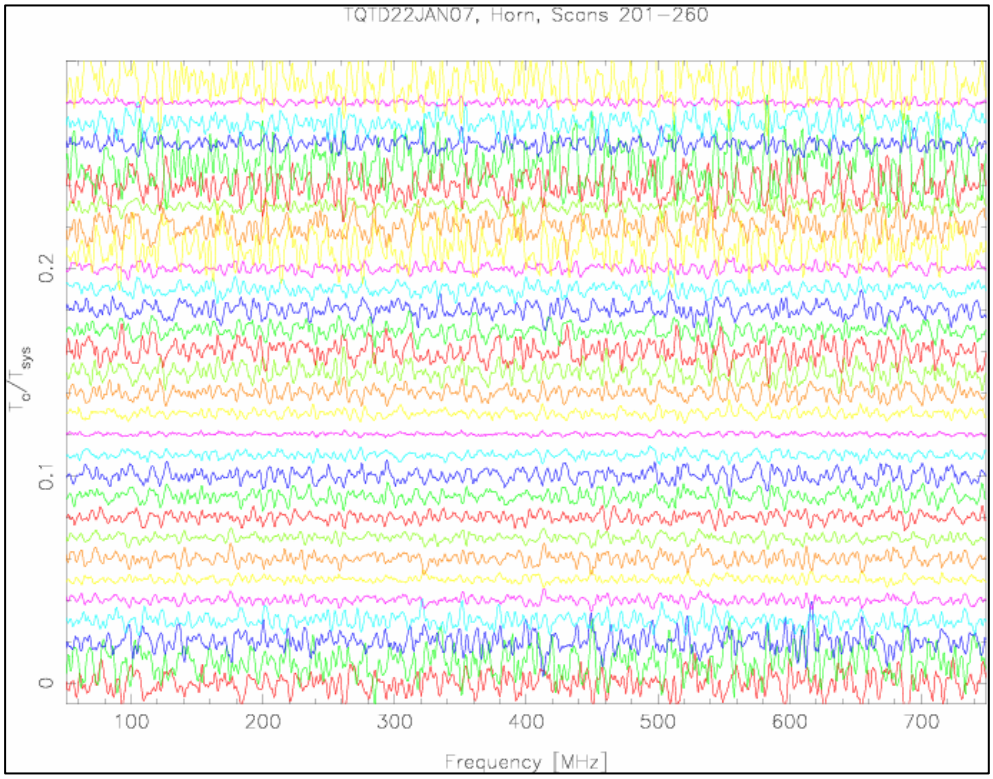

Figure 12: Horn baseline waterfall. Note scale is 5X previous waterfalls.

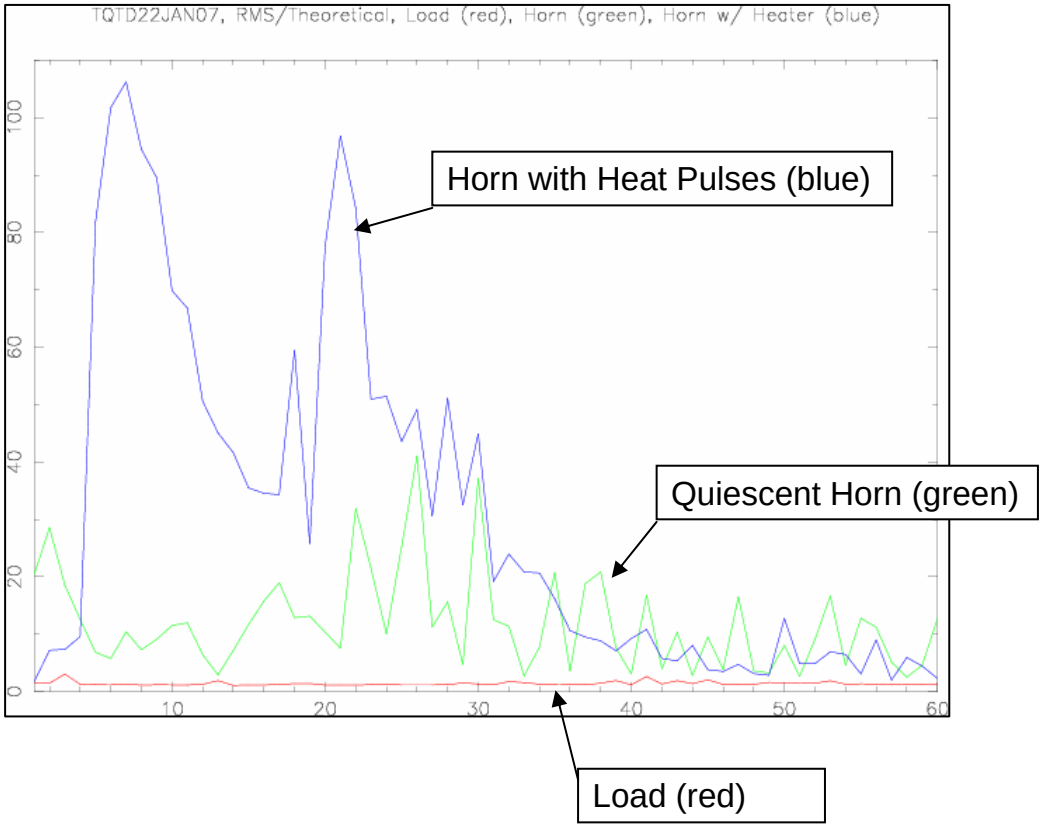

Figure 13: Feedhorn and reference load RMS comparison.

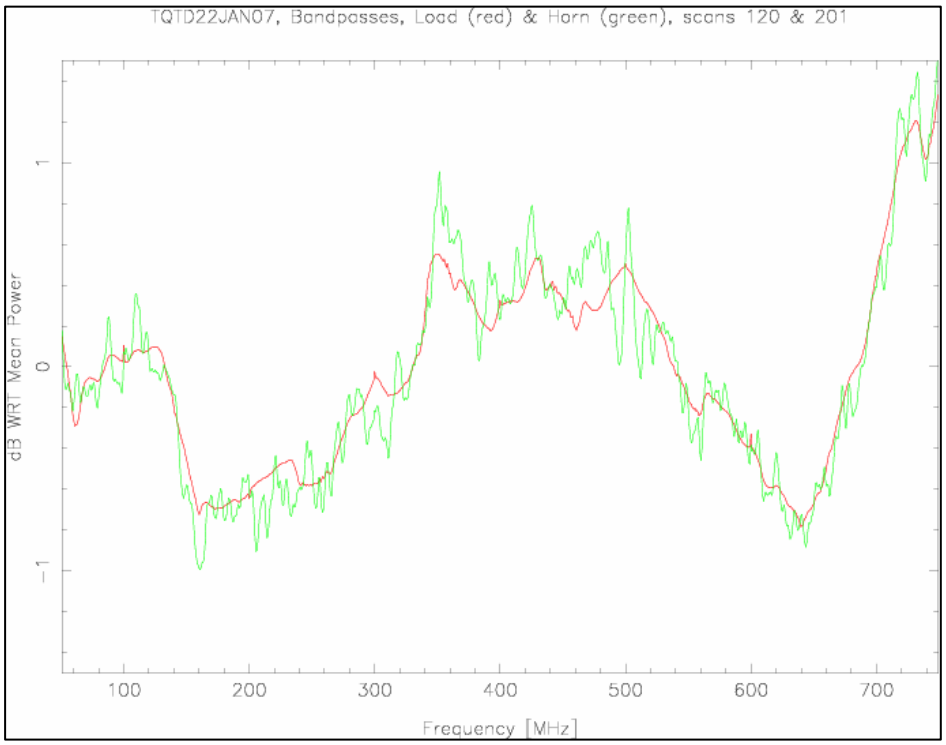

Figure 14: Raw spectra from the horn (green) and from the reference load (red).

4.3.1 Absorber

It was mentioned in Section 2 that placing absorber in the Q-band and Ka-band cryostats near the feedhorns appeared to reduce the baseline instabilities. The use of absorber was now explored in the test cryostat. Emerson & Cuming Eccosorb HR10 was selected because it is relatively inexpensive, easy to work with, and exhibits return loss in Q-band of about 20 dB when backed with a metal plate. The radiation shield was lined with this material (Figure 15), a metal radiation top plate was attached, and most of the seams in the radiation shield were sealed with copper tape. Figure 16 through Figure 18 shows the resulting performance. The baseline quality is the best ever seen, there is little if any difference between the horn and reference load stability, and the short-period ripple is no longer seen in the horn spectra. The total power from the horn is 5.5 dB higher than from the reference load (at 13 K), giving the effective horn temperature as 109 K.

More test results will be presented in Section 6, but first an explanation is proposed for what has been observed.

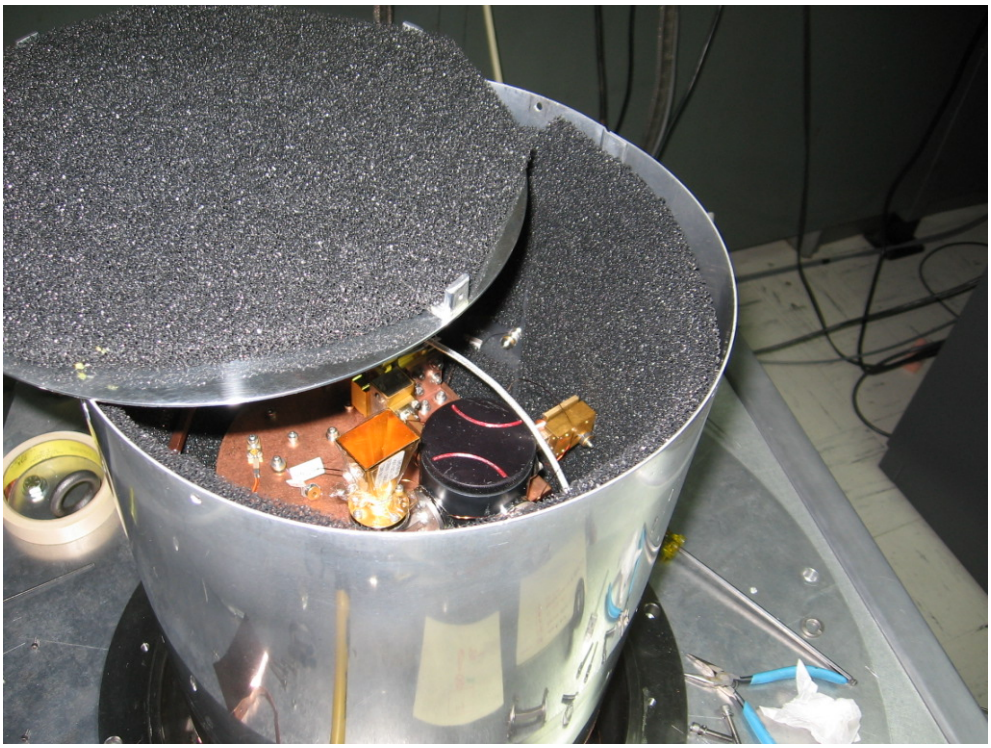

Figure 15: Absorber lining the radiation shield.

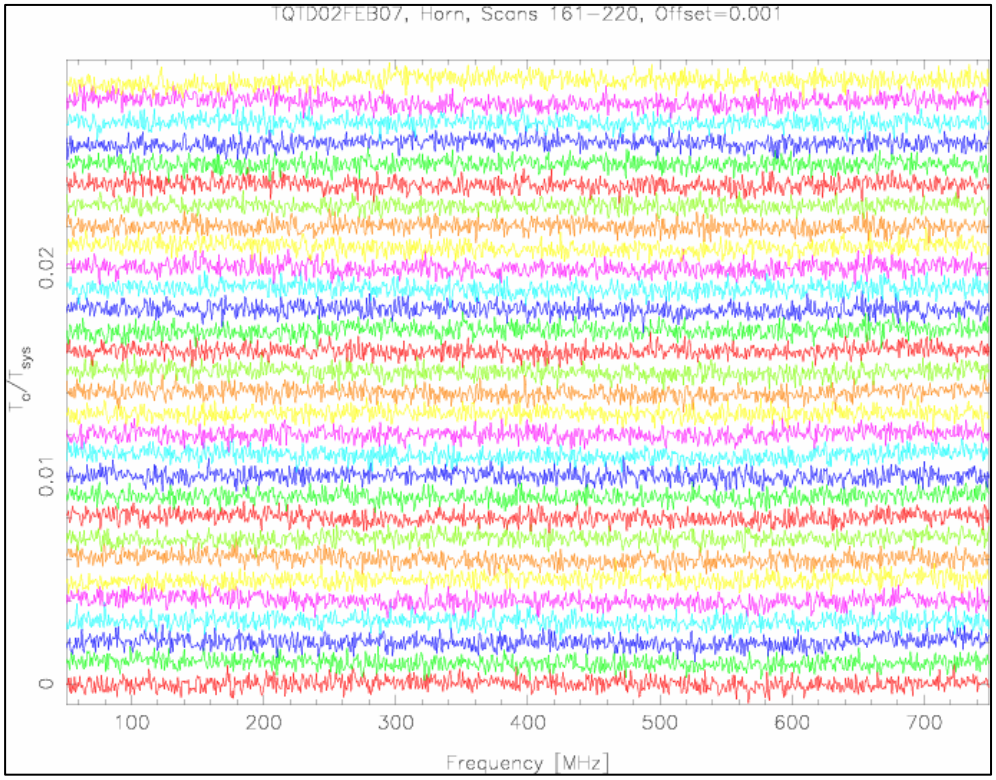

Figure 16: Baseline waterfall for absorber lined radiation shield.

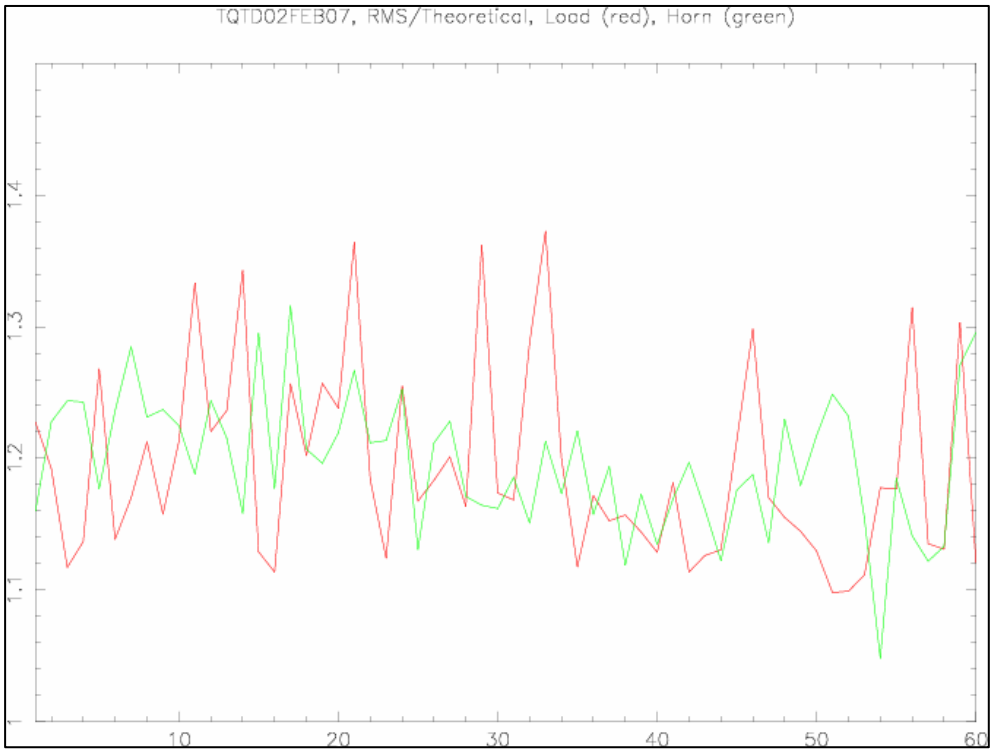

Figure 17: RMS performance for absorber lined radiation shield.

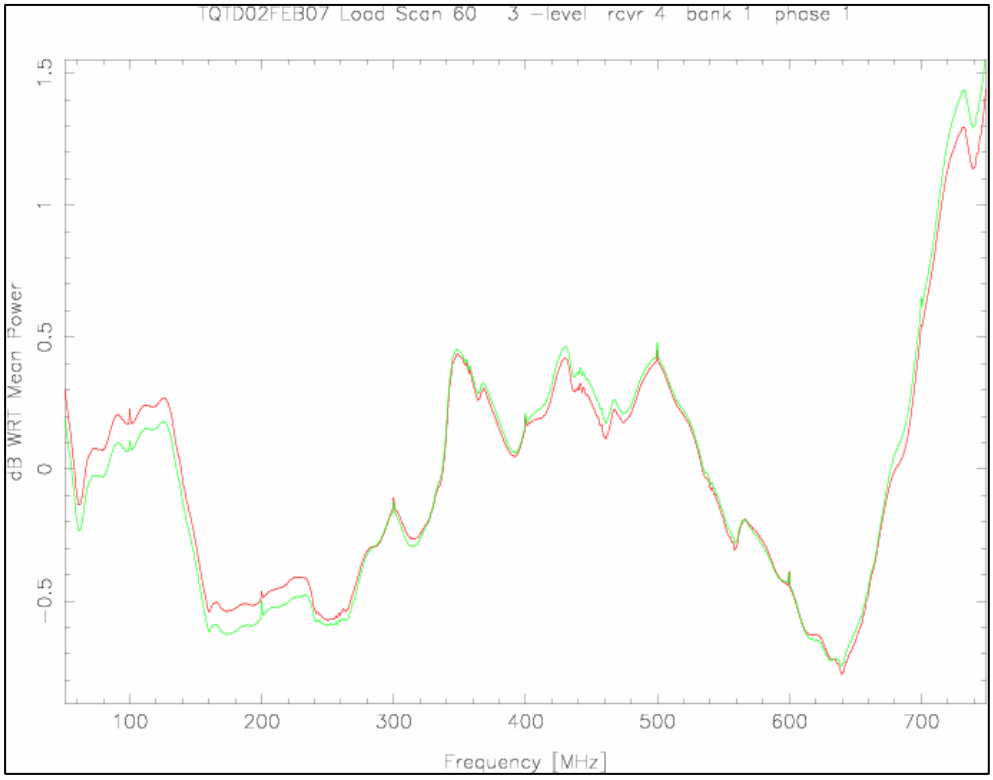

Figure 18: Raw spectra for absorber lined radiation shield. Red is from the reference load; green is from the horn.

5 Explanation & Theory

Two aspects of the results presented in the previous section were not anticipated when this investigation began. First, when the findings led to measurements using a horn to sample the noise environment in the cryostat, the expected mean noise level was on the order of 20 K. Yet, levels approaching 200 K were seen. Second, while there were suspicions that environmental effects were contributing to the baseline ripples, the ripple magnitude and the degree of sensitivity to the cryostat wall temperature seen by the horn measurements were surprising. The only explanation that seems credible is that the cryostat internal surfaces radiate thermal microwave radiation that is fairly efficiently coupled into the horn but with a nasty frequency structure. This supposition led to an exploration of the literature related to cavity noise.

There is extensive literature on cavity noise as related to oscillators and filters, but none relevant to this problem was found. However, in a paper related to passive primary noise standards using microwave cavities, Daywitt[3] explains that each differential cavity wall area $d\Omega$ radiates energy proportional to its thermodynamic temperature T. Assuming the cavity has an output port, (e.g. a waveguide transmission line), a small part of the energy radiates directly to the port, but most is indirect, making multiple bounces off walls before reaching the port. The port total noise temperature, T_n, is given by:

$$MT_n = \int TSd\Omega \tag{4}$$

where S contains the effects of both the direct and indirect energy components. M is the port mismatch factor, giving how well the cavity impedance is matched - equal to one if perfectly matched, zero if totally mismatched. The integral is taken over the total cavity surface area.

The test cryostat is a cylinder approximately 55 cm long by 32 cm inside diameter, with a complicated internal geometry of radiation shields, thermal straps, electronic components, and a refrigerator cylinder. The dimensions and geometry of the GBT receiver cryostats differ, but not greatly. Materials present include aluminum, copper, stainless steel, and small amounts of dielectrics. It would be extremely difficult to calculate Equation 4, but a swept measurement of the embedded horn's return loss was quite simple and informative. It was necessary to move the horn a few inches from the position used in the baseline tests and connect it directly to a waveguide exiting the cryostat. One might suspect the impedance seen by the horn to be close to a short circuit, but that proves not to be the case, as seen in Figure 19. The return loss over the frequency range shown averages close to -20 dB, and a variety of short ripple periods are seen - the one marked is 12.5 MHz. Small changes in the cryostat wall temperature and movements of small metal interior items made clearly visible changes in the ripple pattern. It is clear that little of the power in the spherical wave transmitted from the horn reflects directly back into the horn (the horn aperture was about 24 cm below and facing the cryostat top plate), but undergoes multiple reflections (perhaps hundreds) off the internal surfaces, losing a little power at each reflection from the imperfectly conducting walls.

The cryostat dimensions will not support a simple fundamental resonant mode anywhere near frequencies as low as 20 MHz, but there is extensive literature related to antenna measurement and other fields using highly overmoded chambers, showing that higher-order modes can be quite closely spaced and cover a wide range of periods. Kerr[4] has made a specific calculation for a simple 50 cm cubical cavity giving an average mode spacing of 9.5 kHz at 30 GHz. It is certainly then not unreasonable to see ripple periods in the 10-30 MHz region in the cryostats although it is not clear why these periods stand out.

Therefore, it seems a viable hypothesis that the receiver baseline problems discussed here are due to thermal (Johnson) microwave noise radiation from the cryostat interior cavity surfaces, and that the noise radiation exhibits complex frequency structure sensitive to environmental factors such as temperature and geometry. This radiation can couple into the receiver input signal via leaky waveguide flanges, leaky component bodies, or reflections off of imperfectly matched vacuum windows into the feedhorn aperture. In practical situations, the coupled radiation is too weak to significantly impact the receiver total noise temperature, but the fact it has short-period, unstable frequency structure can seriously impact spectral line observations. This hypothesis explains the things observed when investigating the GBT receivers, and several measurable effects can be predicted. Additional test results are presented in the following section, helping to confirm the hypothesis.

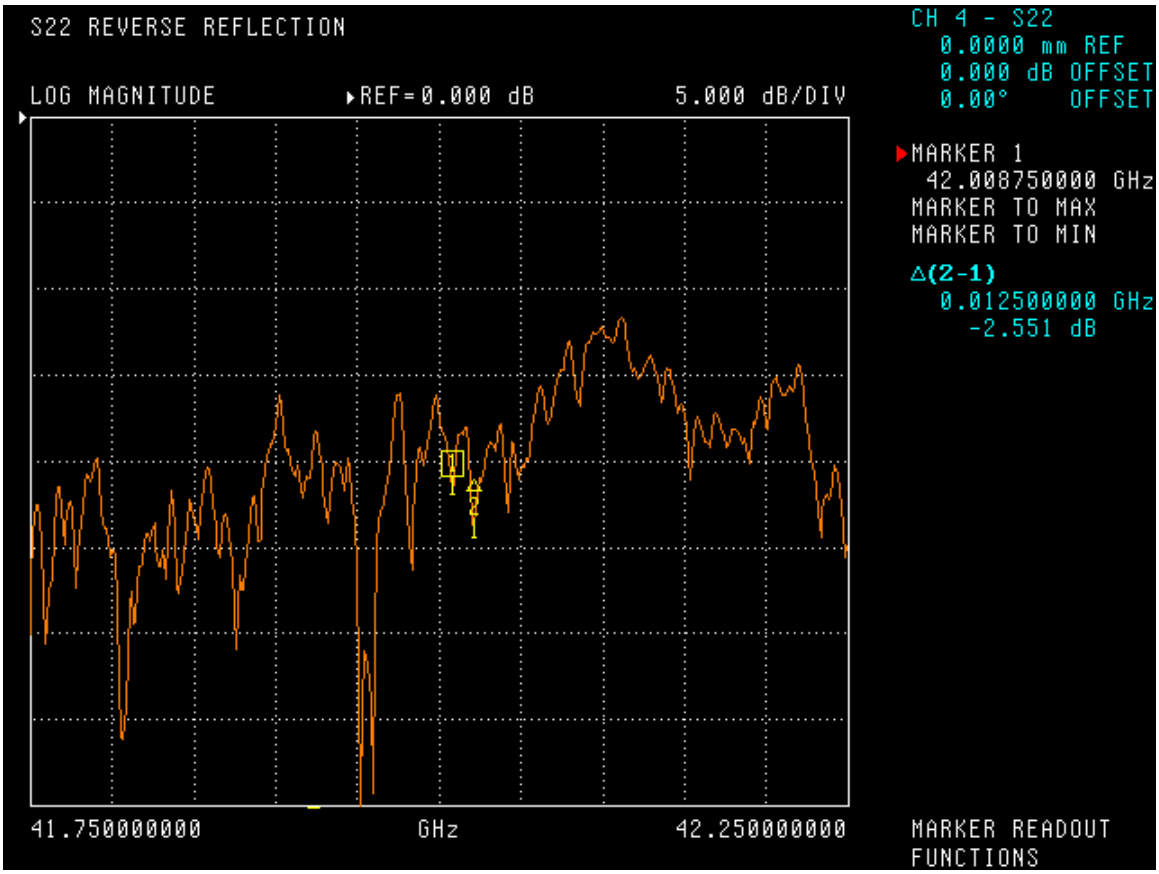

Figure 19: Return loss measured at room temperature with outer can and radiation shield in place. There was no top on the radiation shield, and no absorber in the cryostat.

6 Confirming Results

6.1 Effective Horn Temperature

If the deductions in Section 5 are correct, the horn mean effective temperature should be proportional to the surface temperatures radiating noise into its receiving pattern. Tests under several conditions were run and the results are presented in Table 1 and Figure 20. The effective horn temperature was measured with an averaging power meter over a 800 MHz bandwidth at the radiometer IF, using the transfer switch to alternate between the horn and the reference load at a known temperature. In the "Floating absorber" test, the radiation shield was lined with absorber, including a disk at the top, but there was no metal top on the shield so, due to radiation transfer, the top disk was at some temperature between 70 K and 300 K. In the "Shield w/ Top No Absorber" test, the radiation shield was held together by only four screws at the top and bottom. Such loosely fitting seams are known to provide poor microwave shielding effectiveness, and it appears significant

radiation reached the horn from the outer walls at 300 K. In the "Taped shield" tests, copper tape was used to seal all radiation shield seams, and most of the holes through the 70 K plate (the bottom of the radiation shield cylinder), and the horn temperature approaches the radiation shield physical temperature (~65 K). The fact the results are quite flat with frequency indicates that gain changes due to mismatch differences between the horn and load are not a major error factor. These results are consistent with the hypothesis.

Table 1
Effective Horn Temperature

Condition	Phorn/Pload	Thorn
Top off shield.	7.2dB	185K
Floating absorber	7.1dB	186K
Shield w/ Top No absorber	7.2dB	154K
Taped Shield No absorber	5.5dB	92K
Absorber in taped shield	5.5dB	91K

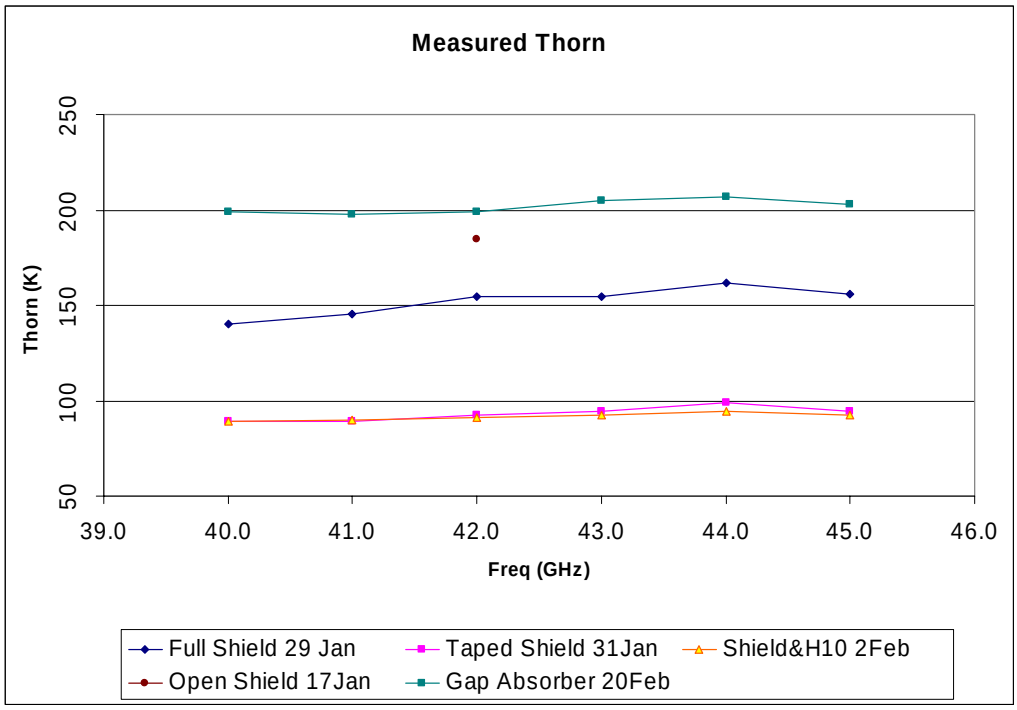

Figure 20: Effective horn temperature versus frequency for several conditions.

6.2 Waveguide Flange Gap

If the cryostat is filled with noise radiation, it should be possible for the noise to leak into the signal path through poorly mated waveguide flanges at the radiometer input. To test for this, a WR22 round UG383 flange between the reference waveguide load and the radiometer was purposely cocked, leaving a significant gap in the connection. Figure 21 shows the results. Strong responses were seen to the brief application of power to the heat pad on the cryostat outer wall, and even to application of heat on the wall with a blowing heat gun. The causes of the two peaks to the left of the heat gun peak are unknown. For comparison, typical rms data is also plotted as measured on the reference load (with the flanges as tight as possible) in an open radiation shield (green), and in an absorber lined shield (blue).

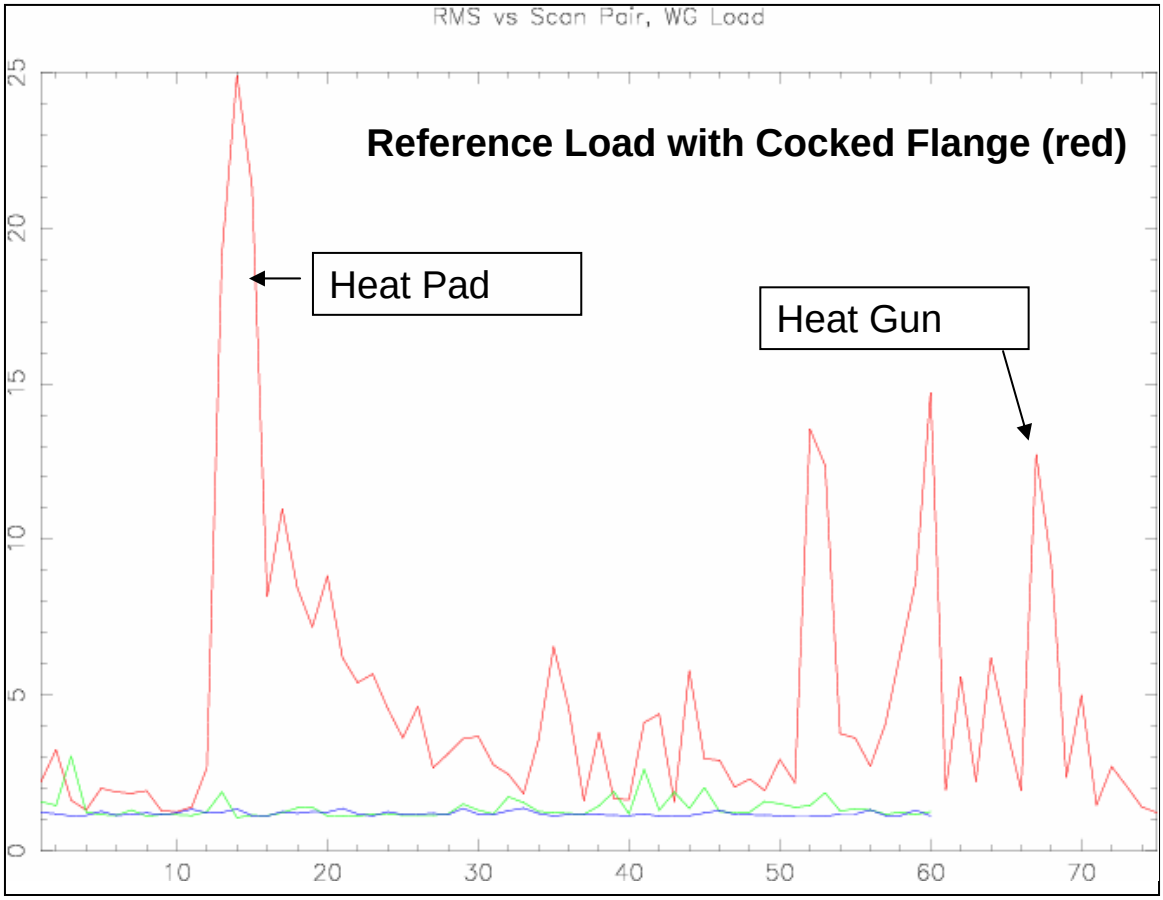

Figure 21: Baseline RMS data for load with a cocked waveguide flange.

6.3 Metal Movement

If the unstable baselines are caused by multimoding and/or multipath reflections of the noise radiation, then the baselines should be sensitive to changes in geometry inside the cavity. In order to test this idea, the fiberglass shaft was disconnected from the transfer switch, and a small metal "flag" fashioned from copper tape was attached to the shaft, Figure 22. The arrangement allows movement of the flag at specific times while the cryostat is closed and cold. Since the transfer switch shaft was disconnected, two cooldowns were required to measure both the horn and the reference load. Figure 23 shows the overlaid measurement results. Between the first and second scan of four scan pairs (on each source), the flag shaft was rotated. The largest peaks coincide exactly with these movements. Note that the vertical scale in this plot is in decibels - the peaks in the horn baselines are about 100 times the theoretical rms noise. Figure 24 shows how the ripple in the measured raw spectra changes; the flag was moved a few degrees between the red and green spectra measurements. These results are also in agreement with the noise radiation hypothesis.

The flag movement is a repeatable and strong indicator, and it was used to track down how the baseline ripple gets into the reference load data. By eliminating components one by one, it was found that about half the reference load baseline ripple amplitude was associated with the waveguide transfer switch. Yet, with a waveguide load directly connected to the LNA input port, baseline ripple at about 2.5 times theoretical remained. Suspecting an imperfect waveguide connection at the load to LNA interface, an indium gasket was fashioned and installed in the flange, but still the baseline ripple remained. Finally, by packing a layer of the HR10 absorber just around the LNA body, all baseline response to the moving flag was eliminated. (In this arrangement, there was not sufficient absorber to eliminate ripple on the noise radiation. That was confirmed by replacing the load with the horn while leaving the absorber still in place.) The exact means by which the cryostat noise radiation is coupled into the LNA body has not been determined.

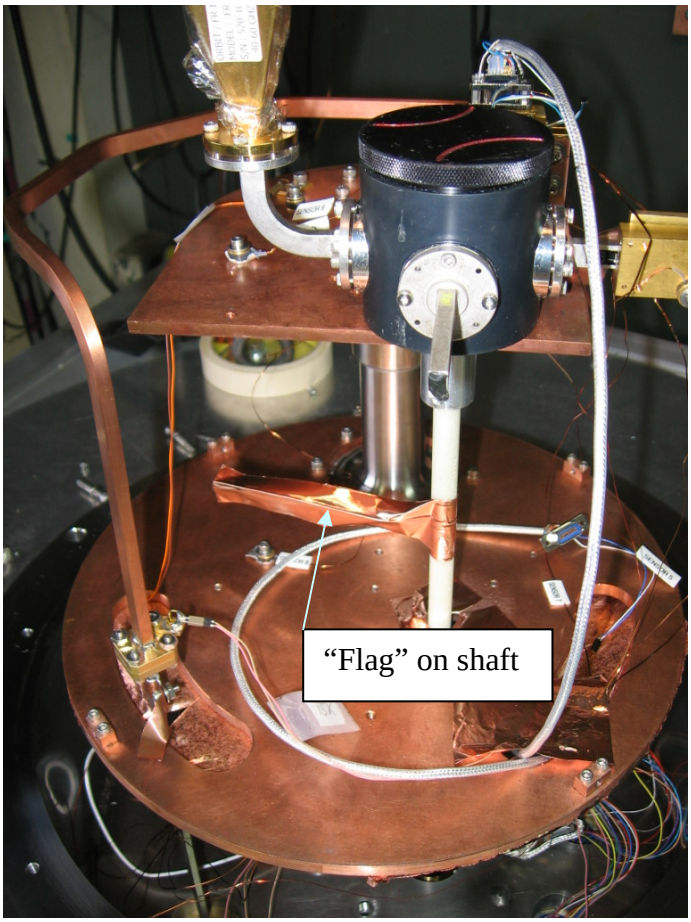

Figure 22: A metal flag that can be moved from outside the cryostat.

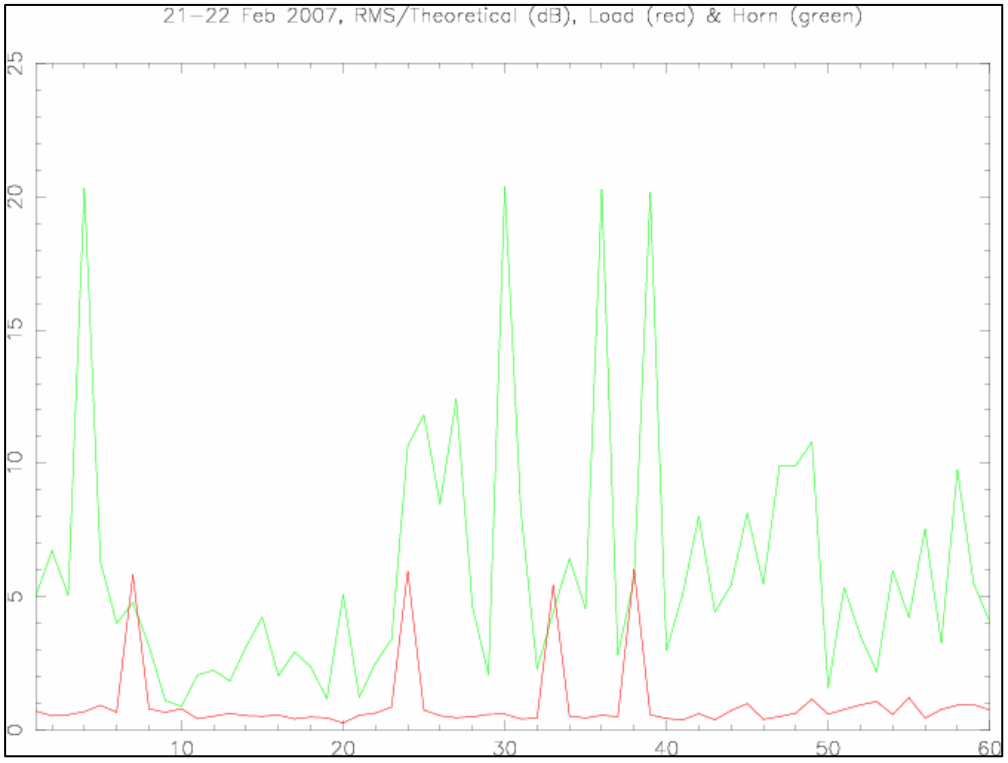

Figure 23: RMS peaks associated with moving the metal flag.

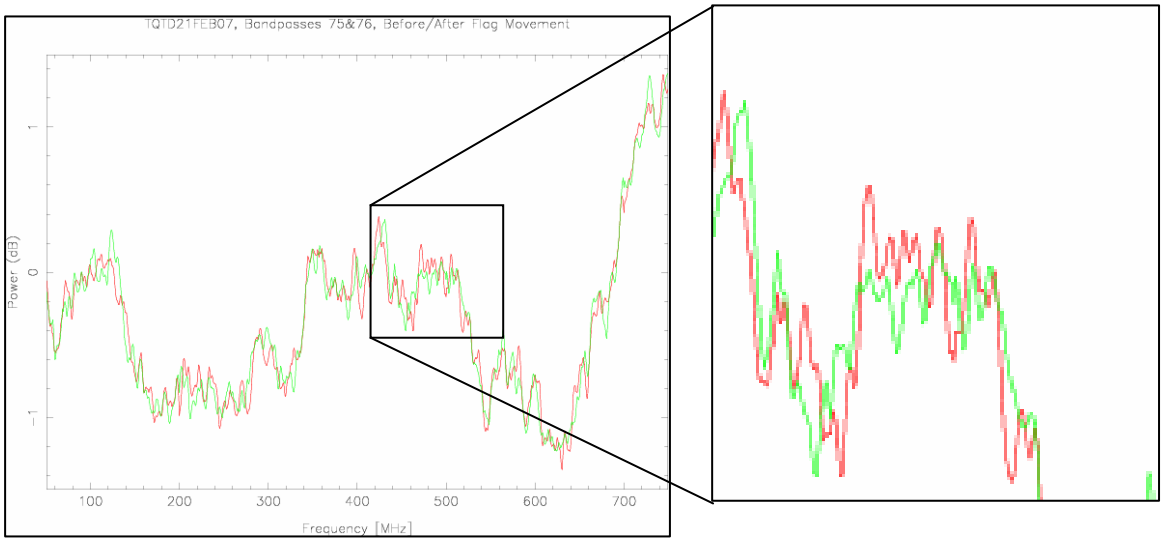

Figure 24: Raw spectra showing the effect of moving the metal flag.

7 Conclusions

It seems conclusive that the baseline problems which plagued the GBT Ka-band and Q-band receivers were due to thermal noise radiating from the cryostat interior cavity surfaces. Due to multimoding and multiple reflections in the cryostat cavity, the noise radiation exhibited complex frequency structure sensitive to environmental factors such

as changes in temperature and geometry. Small amounts of the radiation coupled into the receiver input signal path via leaky waveguide flanges, leaky component bodies, or reflections off of the imperfectly matched vacuum windows and into feedhorn apertures. The detrimental baseline effects can be mitigated by using microwave absorber to smooth the cavity noise ripple, and by other means to shield the signal path from the noise radiation.

It is difficult to estimate what level of shielding is necessary, but since it was relatively easy to induce baseline ripple at more than 100 times the theoretical rms noise level for one minute integrations, shielding effectiveness of more than 30 dB seems necessary. That degree of shielding is not too difficult to achieve, and may be one reason the same problem has not been seen in other GBT receivers. No others have the feedhorns inside the cryostat (probably scattering into the feed is a dominant coupling means and the most difficult to eliminate), and most have well-shielded coaxial transmission lines before the low-noise amplifiers.

To my knowledge, this cause of spectral baseline problems has not been previously reported, and it is not clear why. Certainly there have been many radio astronomy receivers built with cooled feeds, but most if not all at 80 GHz and higher. Up there, the cavity resonant modes will be spaced much closer, perhaps making the effect less of a problem and more difficult to recognize. Depending on the cryostat size, a spectrometer with at least 100-200 MHz of bandwidth and with resolution less than 1 MHz is needed to see the ripples in detail. Until recently, such spectrometers were rare. And, the GBT secondary focus optics is fairly low in magnification as antennas go, resulting in a broader feedhorn beam pattern and perhaps more susceptibility to scattering into the horn. Finally, it may be that others have indeed seen and recognized the source of such instabilities.

In any case, it is hoped that this report proves useful to other receiver designers, and helps them avoid the same pitfalls.

8 Acknowledgements

Without the contributions of many NRAO colleagues, this work would not have been successful. The author gained much experience and knowledge working with Rick Fisher and Dana Balser on general baseline issues, and from discussions with Rick, Marian Pospieszalski, Tony Kerr, Geoff Ediss, and others on this particular problem. Marian suggested the useful cavity return-loss measurement. Galen Watts and Gary Anderson led the work on the Ka-band and Q-band receivers, and Jonah Bauserman constructed and modified (many times) the test cryostat. It would not have been possible to unravel the problem without access to the Lab Spectrometer – it was designed by Ray Escoffier, and brought into operation at Green Bank by John Ford, Rich Lacasse, and others.

[1] *Investigation of Spectral Baseline Properties of the Green Bank Telescope*; NRAO Electronics Division Internal Report No. 312; R. Fisher et. al.; September 2003.

[2] *Radio Astronomy, 2nd Ed.*, J. Kraus, Section 3-19, 1986.

[3] *The Noise Temperature of an Arbitrarily Shaped Microwave Cavity with Application to a Set of Millimetre Wave Primary Standards*, W. C. Daywitt, Metrologia, 1993/94, **30**, pp471-478.

[4] *Spectral Density of Modes in a Rectangular Metal Box,* A. R. Kerr, NRAO Electronics Division Technical Note No. 207, March 2007.