

Zspectrometer commissioning report, Fall 2006

NRAO GBT Memo 245

18 April 2007

A. Harris^a, S. Zonak^a, K. Rauch^a, A. Baker^{a,b,c}, G. Watts^b, K. O'Neil^b,
R. Creager^b, B. Garwood^b, P. Marganian^b

^aUniversity of Maryland, ^bNational Radio Astronomy Observatory, ^cRutgers University

ABSTRACT

We report on the initial commissioning phases of the Zspectrometer, a spectrometer to cover the entire Ka band at modest resolution. Our report provides a summary of the first tests. All important functional tests were successful. Spectra show significant nonideal baseline structure, and we identify two problems with the Ka band correlation receiver that are potentially important. First, loss imbalance in the receiver's input components produces an offset at the correlator output. Second, fluctuating receiver gains multiply this offset to produce fluctuations at low lags in the cross-correlation functions. These two effects combine to produce lumpy spectral baselines. A third area of investigation is the generation of the wideband tone needed to calibrate the correlator's internal phase. Either the test tone's sweep bandwidth or the receiver components will set limits on the observing band.

1 OVERVIEW

This report covers the first commissioning phase of the Zspectrometer project in the Fall of 2006¹. All functional aspects of the commissioning were successful: the interfaces between the receiver, correlators, data acquisition system, and calibration systems are compatible and are nearly complete. We were able to make demonstration observations of strong astronomical sources. Laboratory testing revealed systematic instrumental effects that limit the astronomical performance of the overall system, however, and systematic spectral structure precluded deep spectroscopic observations.

Much of this report covers the characterization of the systematics and their origins. In summary, it appears that differential loss in the circuits at the very front of the receiver leaves an effective temperature imbalance that causes a nonzero signal at the correlator outputs. Ideally, the correlation receiver perfectly differences the signals at its two inputs, producing an output of zero when both inputs view the same temperature. This property reduces the sensitivity to amplifier gain fluctuations: zero scaled by a fluctuation is still zero. Non-zero offset signals are amplified by fluctuating gains, adding relatively large and unstable spectral baseline structure to the spectra. Improving the spectral baselines will involve reducing the differential loss and improving the gain stability. Sections 3 and 4 cover these topics in detail. Section 5 covers the spectral range limit imposed by the calibration signal generator.

2 THE ZSPECTROMETER AND SOME FIRST RESULTS

The overall Zspectrometer system, shown in Figure 1, consists of the NRAO Ka-band correlation receiver and a correlator package of eight WASP2 analog lag correlators [1]. The correlators

¹Laboratory work from September 16–19, September 21, October 3–5; mechanical installation on the telescope in mid-October with functional checkout on October 31; and on-sky tests on November 16, 17, and 29

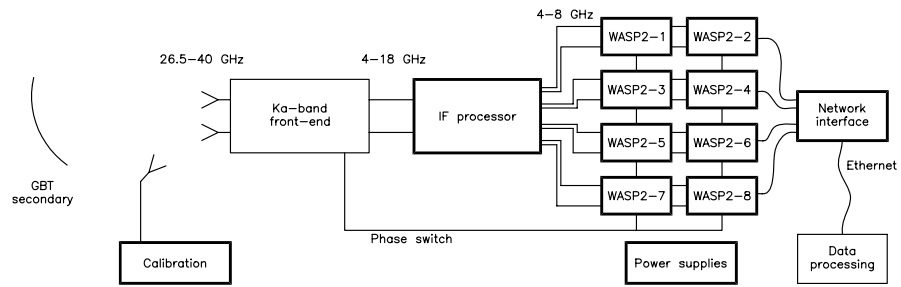


Figure 1: Instrumental block diagram. The Ka-band correlation receiver has two input horns and two IF outputs. An IF processor converts the 4-18 GHz band to four 4-8 GHz bands for the correlators. Four sets of WASP2 correlators, paired to produce 256 lags each, find the 1024 lag cross-correlation function.



Figure 2: The Zspectrometer correlator package in the lab.




Figure 3: A view of the Zspectrometer mounted on the Green Bank Telescope receiver turret.

are configured in a 4×2 arrangement with four sub-bands, each with 256 lags over 3.5 GHz, or about 16 MHz resolution over 14 GHz. Figure 2 shows the correlator package in the lab. From top to bottom, the chassis are the correlator power supply, two correlator pairs (in Slots 1 and 2), the IF distribution box, and two further correlator pairs (in Slots 3 and 4); network hubs and fiber conversion boxes are at the very bottom. All of the cables connect to the individual chassis' front panels to simplify the enclosure design, at the cost of adding some clutter when the enclosure is open. The enclosure, designed and fabricated by NRAO Green Bank, eliminates radio frequency interference (RFI) when its front cover is in place. Putting the wheels on outriggers adds stability and allows the enclosure to roll through standard height doors. Figure 3 shows the overall system on the GBT's receiver turret: the correlator package, with its front cover in place, to the left, and the bottom of the Ka-band receiver to the right.

2.1 Interfaces

Planning for the interfaces to the GBT system started early in the project, and the implementation went smoothly. The Zspectrometer provides two signals that drive the Ka-band receiver's phase switches in quadrature for an effective switch frequency of 10.4 kHz². Two 0.141" diameter conformable cables, matched in length and about 1 m long, carry the 4–18 GHz IF signals from the receiver to the correlator. Over this band, the receiver output power is about –20 dBm. Power across the band is flat to about 10 dB, and we trim the average for each correlator sub-band with attenuators. We measured the slope across the individual bands, and have ordered gain equalizers to compensate.

By the end of the test period we had completed all software integration with the GBT system. Manager and Astrid scripts set up and control configurations and integrations. A set of GBTIDL scripts form the beginning stages of a full data reduction package. Although real-time monitoring of data and spectra is not yet possible, it should be implemented in summer 2007.

Apart from the usual startup and teething problems, we had two noteworthy computer-related problems. First, one of the Zspectrometer's internal microcontrollers hung once during testing on the telescope, and had to be manually reset. This is very unusual, occurring once in the month or so that the system was on the telescope, and never in the lab. We have changed the microcontroller code to deal with combined packets; if the telescope network is very busy, it is possible that two commands are sent to the microcontrollers in one packet, but the microcontrollers expected a single command. The other problem was that the GBT system's requests failed to trigger data collection from all of the microcontrollers for a period of about an hour on the last night of observation. This was likely a Manager configuration problem, since it happened after reconfiguration for pointing. It appeared that the Manager did not request the microcontrollers to set up hardware or to start integrating. Simple communication problems can be ruled out as the cause of the problem since the Zspectrometer control software did not generate errors and the microcontrollers all returned monitor data to the monitor file and to a terminal through this period.

²Tests with lower frequency switching did not change the CCF's shape or offset within fluctuations, but the rms of the difference from the standard increases slightly with decreasing switch frequency. For correlator #2, the variance of the residual was 1.2 counts at 10.4 kHz (Scans 12097–12100, 10/4/06). The variance of the residual was 2.9 counts for a 6.25 kHz frequency and 7.0 for 4.17 kHz (Scans 12101–12104 and 12105–12108 compared with 12109–12112, 10/4/06). Based on these few data points, it seems that switching frequencies above 6 kHz are necessary for optimum sensitivity.




Figure 4: Test spectrum of Jupiter calibrated against 3C48. Different colors denote the three separate sub-bands.




Figure 5: Raw, unnormalized spectrum of 3C48. Different colors denote the three separate sub-bands.

2.2 First results

Figure 4 is our first spectrum of Jupiter, taken during the second on-sky commissioning period. It includes data from three of the four correlator sub-bands; the power level from the phase calibration source drops rapidly beyond 37 GHz (see Sec. 5). This power drop, possibly combined with deteriorating receiver sensitivity, makes the top 3 GHz of the band difficult to observe at present. We divided the spectrum of Jupiter by the spectrum of the quasar 3C48 (~ 1 Jy, $z = 0.367$) to remove the passband gain, but made no other corrections. The sub-band edges match well, indicating that the receiver and spectrometer systems are very linear. The spectral index across the 26–36 GHz spectrum is $+3$, which plausibly matches an index of $+2$ (blackbody) for Jupiter and -1 for 3C48. Figure 5 is the raw, unnormalized spectrum of 3C48. After full reduction, a spectrum of the Orion A HII region had considerable baseline structure but did not show H α recombination lines. A possible reason for the nondetection is that the emission region is extended on the scale of the $1.8'$ beam separation.

On the last commissioning night we made a 2.2 hour integration toward SMMJ02399-0136, a source in one of our previously TAC-approved programs. These spectra were dominated by systematic structure in the baselines, precluding the detection of weak lines.

3 ORIGIN OF THE OFFSETS IN THE CROSS-CORRELATION FUNCTIONS

An ideal correlation receiver differences its inputs for a zero output when the input temperature is the same for both inputs X and Y . Compared with single-pixel total power receivers, this property reduces the sensitivity to gain fluctuations because the signals are very small. The output u of an ideal correlation receiver is

$$u \propto \alpha\beta(T_X - T_Y)g_A g_B^* g_M \quad (1)$$

in the notation of [2], where α and β are the voltage transmission coefficients through the first hybrid, $T_{X,Y}$ are the temperatures at the two inputs, $g_{A,B}$ are the complex voltage gains of the amplifiers after the first hybrid (the star denotes a complex conjugate), and g_M is the multiplier gain. Equation (1) is easily generalized for cross-correlation functions (CCFs), u as a function of delay τ .




Figure 6: Plots of typical cross-correlation functions (CCFs) for the four sub-bands with both receivers terminated at ambient temperature. Lag decreases with lag number; a dotted line marks the zero (total power) lags. The CCF amplitudes are dominated by systematic structure. Slots 1 through 4 cover sky frequencies of 36.1–39.6 GHz, 32.6–36.1 GHz, 29.1–32.6 GHz, and 25.6–29.1 GHz. The nonideal structure amplitude decreases with sky frequency.

Zspectrometer measurements in the lab and on the telescope show nonzero multiplier outputs with characteristic shapes in the cross-correlation function $u(\tau)$; Figure 6 summarizes the CCF’s typical shapes. Such large offsets seriously compromise the receiver stability. At present we believe that electronic pickup in the amplifier bias causes moderate offsets over all lags, and that loss imbalance in the circuits before the first hybrid causes the higher amplitude structure at low lags. We discuss each in turn.

3.1 Approximately constant offsets across all lags

In addition to few analog to digital converter (ADC) counts of stable intrinsic correlator offset, all of the cross-correlation functions have a more or less constant offset that depends on IF power and extends over all lags. The offset disappears when the first-stage low-noise amplifiers (LNAs) are off, when the correlator input attenuation is high, or when one or both IF cables between the front-end and Zspectrometer are disconnected. This dependence on IF power level indicates that the offsets are from the front end. Offsets for correlators³ that had high input attenuation were unaffected in all of the tests, confirming that the offsets are not electronic pickup within the correlators.

We initially drove only one of the receiver’s phase modulators (this scheme has worked well

³Scans 13-16, 85-88, 9/21/06




Figure 7: CCF for Slot 2 with fundamental switching in each phase modulator (lines away from zero counts, blue and black) and quadrature switching (line close to zero, red).




Figure 8: CCF as in Figure 7, but with one trace inverted and offset. While the structure at high lags (low lag numbers) changes with quadrature modulation, the low-lag structure remains unchanged.

with WASP2 correlators in the past), but saw large offsets over all lags. Curves in Figure 7 compare simple switching of individual phase modulators and quadrature switching of both. In the single switching mode the correlator generates the same phase reference signal whichever phase switch is active; electronics within the front end distribute the reference to the appropriate phase switch and generates the drive current. Changes in the CCF that depend on which phase switch is running must be caused by changes in the front end rather than by changes in internal correlator pickup.

Compared with modulation using only one phase switch, these offsets are much smaller when both of the front end's phase switches run in a quadrature-switching scheme⁴ (Figure 7). The offset amplitude drops to within a few counts of the level without an IF signal⁵, which is typically -4 . For correlator #2, a chi-square minimization of the amplitude difference between the low lags of the single- and quadrature-switched CCFs neatly removes the low lag structure⁶, clearly separating the structure at low lags from the sawtooth pattern that appears with a single phase switch (Figure 8). The other correlators behave in the same way with different scaling and offset values⁷.

For the quadrature switching tests above we used a bench power supply for the phase switch driver, which normally shares the internal power supply with the front end LNA bias cards and other front end subsystems. Reconnecting the phase switch drivers to the internal supply did not significantly change the CCF shapes⁸. A small and consistent gain change for all correlators between the measurements with an external and internal power supplies most likely indicates a change in load temperatures or gain. Injecting additional phase-switch drive current into the front end's chassis near the phase switches has a small but visible effect, but one that is far from accounting for the entire offset.

⁴Scans 1129–1193 with quadrature, 11949–11952 for simple switching, 10/4/06

⁵Scans 11933–11936, 10/4/06

⁶Subtracting the quad-switched 11930.tpr from a scaled simple-switched 11950.tpr scaled by -0.963 and minus 238.3 counts

⁷Correlator #1: scale = -0.975 , offset = 302.8; #2: -0.963 , 238.3; #3: -0.949 , 53.5; #4: -0.86 , 197.4

⁸Scans 11985–11988, 10/4/06

Taking these measurements together, we rule out pickup that couples through the power supply as the origin of the moderate offset. The small (although visible) change with switching current injected into the receiver frame shows that ground currents are not the main problem. It is unlikely that the offset comes from spectral changes from the switching: a nearly constant signal versus lag would indicate a spike at zero frequency, where the correlators have no response. It is most likely that the offset is due to phase-switch pickup by the amplifier bias lines. The amount of pickup depends on which phase switch is running. Changing to quadrature switching reduces the amount of pickup because there is no drive signal at the demodulation frequency, but instead two orthogonal signals, each at half the demodulation frequency. Generating the modulation frequency requires a multiplication rather than an addition, so as long as the pickup is small and linear the demodulation frequency is not present on the front end chassis. Driving the phase switches with orthogonal signals seems sufficient at this stage.

If the offset is due to inadequate shielding on the LNA bias lines, pickup from other sources is a concern because it can contribute to the overall system stability.

3.2 Large offsets at low lags

With both receiver inputs terminated at the same temperature, all of the correlators' cross correlation functions have significant structure that is indistinguishable from a spectral difference between the two inputs. The two high-frequency bands (low end of the IF band, Fig. 6) have three or four large peaks that are concentrated at low lags and correspond to an input temperature imbalance of about 20 K⁹. The two low-frequency bands contain rapid variation with an overall modulation extending to high lags, with lower peak amplitudes by a factor of about two¹⁰ (Fig. 6). In all cases the zero lag, which gives the total power, has low amplitude compared with its neighboring lags. Since delays in the front end have been trimmed to minimize the total power signal for the CCB this is not surprising: the CCB's detectors are equivalent to the Zpectrometer's zero lags. For the Zpectrometer, delay match over each spectral channel is necessary, not just an average match across the input band.

As is the case for the offsets that are approximately constant with lag (Sec. 3.1) these offsets also scale with IF power and disappear when the first stage of amplification (LNA) is off¹¹. Unlike the other offsets, these large offsets are unaffected by different phase switching schemes (see Figures 7 and 8). The amplitudes do scale with the temperature of the input common to both inputs; this offset has a well-defined Y -factor.

If the structure were constant in time it would nod out and be of no concern. As discussed in Section 4, however, the offsets are not entirely stable in time. This limits the spectrometer performance, and it is important to identify the source of the offset and remove the effect.

It is most likely that these offsets arise from a frequency-dependent loss in the front end's input circuit. Allowing for power gains $G_{X,Y}$ and additional offsets $T_{OX,OY}$ before the hybrid modifies equation (1) for the correlator output to

$$u = \alpha\beta [G_X (T_X + T_{OX}) - G_Y (T_Y + T_{OY})] g_A g_B^* g_M, \quad (2)$$

where $T_{X,Y}$ are the radiation temperatures at inputs. For equal input temperatures $T_X = T_Y = T_{in}$ (2) can be rewritten as

$$u = \alpha\beta G_X \left[\left(1 - \frac{G_Y}{G_X} \right) T_{in} + T_O \right] g_A g_B^* g_M, \quad (3)$$

⁹Scans 12061–12062, ambient load, 10/5/06

¹⁰Scans 12063–12064, ambient load 10/5/06

¹¹Scans 00017–00020, 9/21/06

with offset temperature

$$T_O \equiv T_{OX} - \frac{G_Y}{G_X} T_{OY} . \quad (4)$$

In the ideal case $G_X = G_Y$ and $T_O = 0$. Although the equivalent offset temperatures may not be known, a scaled version can be measured with the Y -factor (ratio) method with hot and cold loads, yielding

$$T_O = \left(1 - \frac{G_Y}{G_X}\right) \frac{T_H - Y T_C}{Y - 1} . \quad (5)$$

If the offset temperatures come from resistive loss at temperature T_{res} , then

$$T_O = \left(\frac{1}{G} - 1\right) T_{res} . \quad (6)$$

Some algebra gives the input gain ratio G_Y/G_X in terms of the measured correlator outputs with combinations of hot-hot, cold-cold, and hot-cold loads at the inputs (u_{HH} , u_{CC} , u_{HC} ; $Y = u_{HH}/u_{CC}$):

$$\frac{G_Y}{G_X} = \frac{\frac{u_{HC}}{u_{HH}} Y (T_H - T_C) - (T_H - Y T_C) - (Y - 1) T_H}{\frac{u_{HC}}{u_{HH}} Y (T_H - T_C) - (T_H - Y T_C) - (Y - 1) T_C} . \quad (7)$$

Table 1 contains the Y -factors of the structure for hot-hot and cold-cold measurements, the zero-lag value for hot-cold, and the input power gain ratio derived from (7).

Slot	Sky Freq. [GHz]	Y -factor	T_O [K]	Peak (H-C) [cts]	G_Y/G_X [dB]
1	36.1–39.6	2.81	36.4	1.167×10^4	0.4
2	32.6–36.1	1.55	300.1	-1.456×10^4	0.2
3	29.1–32.6	3.01	24.4	6.104×10^3	0.2
4	25.6–29.1	3.18	16.5	5.195×10^3	0.2

Table 1: Values for evaluating the front end loss. The columns are: the slot number, corresponding sky frequency, the Y -factor measured with ambient and liquid nitrogen loads over both inputs, the equivalent temperature imbalance assuming 80 and 290 K for the load temperatures, the zero-lag counts for a ambient-cold load difference, and the gain imbalance in dB.

Some differential loss between the two inputs is guaranteed by asymmetrical front-end circuitry; this is not optimal for a balanced receiver architecture. Table 2 summarizes components in the path between the horns and first hybrid. Data in Figure 6 indicate that the nonideal effects are larger at higher frequencies, where loss mismatches can be expected to be larger. To minimize optical cross-coupling, correlation is between orthogonal polarizations: one signal comes from the “through” port of the orthomode transducer (OMT) on one horn, and the other from the “side” port of the other OMT. An obvious worry is the differential band shape between different ports of the OMTs. Measurements of other OMTs show transmission matches within a few hundredths of a dB [3], however, apart from narrow resonances that would introduce a number of weak sinusoidal components in the spectral baseline. The waveguide circuit is asymmetric, with two waveguide bends and a twist in one input path, but not the other. Ohmic loss sets a minimum of 0.11 dB loss at 26.5 GHz to 0.08 dB loss at 37 GHz for copper WR-28

Input 1	Input 2
Horn 1	Horn 2
Circular-square transition 1	Circular-square transition 2
Polarization phasing section 1	Polarization phasing section 2
OMT 1 (side output)	OMT 2 (thru output)
45° swept E-plane bend	Short straight section
Cross-guide coupler	Cross-guide coupler
90° twist	Short straight section
180° swept E-plane bend	

Table 2: Front-end components before the first hybrid for the two inputs.

waveguide of this length [4, Ch. 8.02]. Since the CCF nonideal structure is larger at high frequencies and the loss is not ohmic (compare Table 1 and Eq. (6)), reflection loss may dominate. Ohmic and reflection loss mismatches between the horns, transitions, phasing sections (which have significant internal structure), and cross-guide couplers also be important. An fluctuating offset from differential coupling to radiation modes within the cryostat cavity is an additional source of potential offsets [5], especially if the coupling is via the LNAs.

4 CONSEQUENCES OF THE STRUCTURE IN THE CROSS-CORRELATION FUNCTION

Even large offsets would not be a problem if they were stable on the timescale of a telescope nod, and subtracted out. Figures 9 and 10 show that the offsets are not stable, however. Each figure shows a time series of data from a lag with a large offset in each of the four correlators. The outputs have been scaled to simplify comparison. While we typically saw waveforms similar to Fig. 9 in the laboratory, Fig. 10 also captures a period with low fluctuations.

The figures show that outputs from the four correlators share a common waveform, implying a common origin to the structure. Since the individual correlators and much of the electronics in the channelizing downconverter are independent, the common structure must come from the receiver’s output power or from the correlator or IF processor power supplies. We can rule out the correlator power supply since the correlator electronics use the power supply as a reference to convert the bipolar correlation function into a unipolar signal for the ADCs, and in the absence of microwave input power the offsets are small and nearly constant. It is possible but unlikely that one of the amplifiers in the downconverter is damaged and unstable. Stability is a problem for Ka band observations with the Zpectrometer, the Spectrometer, and the CCB, however; all of these point to an unstable front end.

Instabilities persist but retain a basic shape over longer time periods. Figure 11 is a waterfall image of 263 raw (undifferenced) cross-correlation functions from Correlator #2 with the telescope tracking SMMJ02399-0136 for 3.33 hr of elapsed time¹². We chose a short integration time of 30 s for individual integrations as a balance between rapid switching and integration efficiency. With 15 s of overhead for each scan, the observing efficiency was 66%. The vertical stripes in the image show that the data are dominated by a nearly fixed pattern structure. Differencing successive scans reduces the structure’s amplitude by a factor of about 500, but still leaves substantial structure in the correlation function (heavy line in Fig. 12). Much of the high amplitude structure is close to the zero lag (lag numbers above 200 in the figure), and

¹²from 20:57 on 29 Nov 2006 to 00:36 on 30 Nov 2006




Figure 9: Total power from all correlators versus time for the receiver on the sky and the telescope in its service position. Correlator outputs have been scaled and shifted for comparison.




Figure 10: Continuation of time series in Figure 9 after a brief break.

looks like a scaled version of the raw data. This indicates imperfect subtraction due to relatively small variations in the structure’s amplitude with time. Subtracting a least-squares fit of the summed data to the differenced data (light dotted line in (Fig. 12)) reduces the structure in the high-amplitude lags by a further factor of a few, but still leaves substantial structure. While the fit matches much of the structure, it does not fit in detail, indicating that band shape changes slightly in time.

Figure 13 is a waterfall plot of 131 differenced pairs of data from Fig. 11, each with a least-squares fit of the summed data from each pair removed. This image shows the imperfect subtraction of the systematic structure as a function of time: there are times when the subtraction is successful, leaving little structure in the low lags toward the right of the image, and other times when the subtraction leaves residual structure. We saw no relationship with antenna elevation or other activity with the degree of instability. Figure 14 is a companion image to Fig. 13, showing the differenced data’s covariance matrix. In the case of CCFs composed of uncorrelated random noise, this matrix would contain a series of uniformly bright spots along the diagonal, indicating each function’s variance, and off-diagonal elements that would be zero within statistical uncertainty. While the spots along the diagonal are somewhat visible in the upper left part of the figure, considerable horizontal and vertical striping dominates the image. This shows the degree of covariance (mutual variance, unnormalized correlation) between different differenced correlation functions. White denotes positive correlation, and black negative correlation. The covariance between different differenced functions is strong when the structure within a correlation function is large: the horizontal and vertical stripes are strongest when the diagonal spot is brightest. An absence of localized sets of striping shows that the structure within the correlation functions varies in amplitude, but not in basic shape for all of the data. A principal components analysis of the differenced and fit data confirms that most of the structure is common-mode, as expected from the structure of Figure 14. Almost all (85%) of the variance is in the first principal component, and essentially all (90%) of the variance is within the first four components. The shape of the first principal component is similar to the undifferenced data at lags away from the zero lag. A spectrum of the first principal component




Figure 11: Waterfall plot of raw data from the telescope for Correlator #2. Each row is a cross-correlation, with the zero lag toward the right edge of the plot. Time increases toward the bottom of the image; each row represents a 30 sec. integration.




Figure 12: Typical difference for Correlator #2 (heavy lines) and least-square fit of sum data (light dotted lines). This plot is of the first readout (top row of Fig. 11), a relatively low noise correlation function.




Figure 13: Waterfall plot of differenced data with least-squares fit of sum term removed.




Figure 14: Covariance matrix for differenced and fit data of Figure 13. The amplitudes have been clipped to show low-level structure.

shows lumpy baseline structure with typical widths of 300–400 MHz across the band.

4.1 Allan variance noise spectrum characterization

We used Allan variance calculations to quantify our measurement of the system stability. The Allan variance was invented to characterize the stability of frequency standards [6, 7, 8]. Schieder and collaborators [9, 10] popularized and extended its use for in radio astronomy. The Allan variance is the variance computed from differences of successively larger subsets of time-series data; this is exactly the fluctuation one is most interested in for finding efficient chopping or nodding rates.

A plot of the Allan variance versus integration time T for the subsamples often follows a form

$$\sigma_A^2(T) = aT^\alpha + bT^\beta, \quad (8)$$

where $\alpha = -1$ represents the ideal case where the deviation integrates down as the square root of integration time and a and b are the values for each power law term at a time of one second. For some devices $\beta = 0$, representing $1/f$ noise, but most receiver systems show “drift noise”, correlated noise and drift with β usually between 1 and 2. Curves following expression (8) have a minimum at the time when the decreasing radiometric and increasing nonideal noise are equal.

We wish to observe with switching times that are firmly in the radiometric noise regime, or times of about one third of this minimum in the Allan variance versus integration time. With telescope and software overheads of 15 to 20 s per integration, reasonable efficiency dictates a nod-side time of 60 s. Ideally, then, we would like to push the Allan variance time to about 180 s.

Figure 18 contains data from lab measurements with an ambient temperature absorber across both horns. The upper panel is amplitude versus time, showing white noise fluctuations superposed on an underlying slowly varying signal. The lower panel shows the Allan variance versus time, with the dark line a least-squares fit to (8) with α constrained to -1 . The noise is radiometric to a few seconds, with a minimum in the Allan variance at an integration time of 4.1 seconds and a drift slope power law of $\beta = 1.9$.

Solving (8) for the time at the minimum variance, T_{min} , gives

$$T_{min} = \left(-\frac{\beta b}{\alpha a} \right)^{1/(\alpha-\beta)}. \quad (9)$$

Assuming that the input power is as high as possible while preserving linearity, a will be as large as possible, so increasing the time to the minimum requires reducing the drift noise amplitude, b , by minimizing system gain variations. For $\alpha = -1$ and a typical value of $\beta = 1.5$, the time to minimum scales as $T_{min} \propto b^{-0.4}$.

For comparison, a fit to the Allan variance for the -4 counts average correlator offsets with no microwave input signal gives $\alpha \sim -0.95$ to more than 70 s. In operation, the additional receiver noise brings $\sigma_A^2(1s) = 150 \text{ counts}^2$ even at high lags, so input signals will always dominate the stability.

We see the same characteristic time of a few seconds in the lab and on the sky. Figures 15, 16, and 17 are the peak amplitude lag, mid-lag, and high-lag Allan variance plots of the time-series data of Fig. 9 for Slot 3. The time series for all lags show some identical underlying structure, but the higher lags contain noticeably more uncorrelated (white) noise. While the zero lag shows pure drift with $\beta = 0.85$, the higher lags have minimum variance times are 6 to 8 s with $\beta = 1-1.6$.




Figure 15: Slot 3, ADC 253, peak amplitude lag, on the telescope. A fit indicates pure $\alpha = 0.85$ drift (straight line).




Figure 16: Time series and Allan variance plots for Slot 3, ADC 129, on the telescope. ADC 129 is a low-amplitude lag showing structure from band edges. The smooth curve shows a fit with $\alpha \equiv -1$, $\beta = 0.96$, and $T_{min} = 6.5$ s.




Figure 17: Slot 3, ADC 129, low amplitude lag, on the telescope. A fit with $\alpha \equiv -1$ yields $\beta = 1.6$ and $T_{min} = 8.0$ s




Figure 18: Time series and Allan variance plots for Slot 2's zero lag in the lab. A fit (curve) with $\alpha \equiv -1$ gives $\beta = 1.9$ and $T_{min} = 4.1$ s.

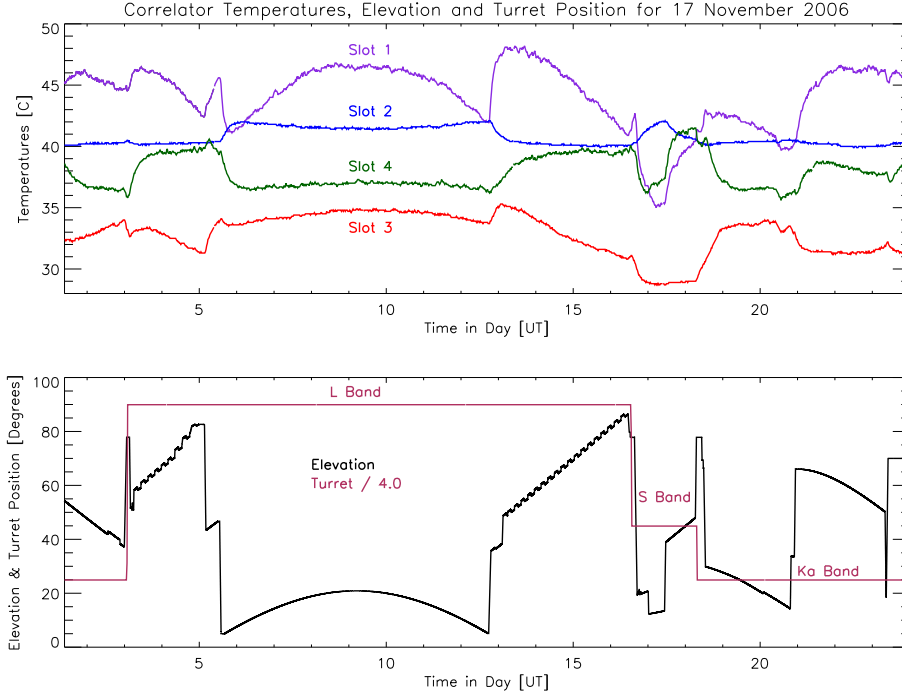


Figure 19: Correlator card internal temperatures (top panel) and package orientation information (bottom panel) versus time.

Although the stability measurements in the lag domain cover the full band, we are more interested in the stability on bandwidths comparable to a few times a typical galaxy widths. The minimum time increases with decreasing bandwidth as [10]

$$T_{min2} = T_{min1} \left(\frac{B_1}{B_2} \right)^{1/(\beta+1)}. \quad (10)$$

Scaling the stability results from the ~ 200 spectral channels implied by the full-bandwidth measurements in the lag domain to the 10 or so channels that include a galaxy's emission and some baseline gives an increase in stability time by a factor of about $20^{0.4} \approx 3$, or about 15 s. This still falls well short of stability times that allow efficient integrations.

One open question is still why, given the pure drift of the zero lag, the average difference baselines are so close to zero (e.g. Figure 4). The answer may well be that the fractional drift is large but the absolute drift amplitude is small.

5 FREQUENCY CALIBRATION AND BANDWIDTH

The Zspectrometer has small internal phase shifts, so a simple Fourier transform would produce a distorted spectrum. We obtain the coefficients for a discrete Fourier transform by injecting monochromatic signals at known frequencies across the band [1]. This phase calibration is needed only occasionally, when the system changes. Changes in temperature affect the cable and microstrip delays, but measurements show that temperature changes with elevation (Figure 19) are small enough that a single calibration at a representative temperature is sufficient.

A falloff in power from the doubler that generates the continuous wave (CW) calibration tone sets the limits on the spectral bandwidth and spectral coverage. At present approximately

three quarters of the band is usable, from about 27 to 37 GHz ($z_{CO1-0} = 2.1$ to 3.3). Some regions within this bands also have low calibration power levels, which makes them difficult to use for high signal-to-noise calibrations. We found that some of the low power regions come from loss in the cross-guide couplers. Work in this area is still in the early stages and we expect improvement.

ACKNOWLEDGMENTS

In addition to acknowledging the contributions of many GBT staff members to this work, we especially note the contributions that P. Jewell, R. Prestage, R. Norrod, J. Ford, and A. Shelton made in their efforts in supporting, coordinating, and managing the GBT efforts for this project. D. Egan and J. Cromer led the design and fabrication of the Zpectrometer's housing and its mounting on the GBT receiver turret. C. Beaudet provided advice on the instrument's electromagnetic compatibility and performed the EMC testing in the Green Bank anechoic chamber. B. Mason provided useful insights in system properties gleaned from measurements with the Caltech Continuum Backend instrument. We appreciate the scheduling flexibility that C. Bignell, the engineering staff, and the commissioning scientists on other projects provided.

This work was supported by the National Science Foundation's Advanced Technologies and Instrumentation program through grant AST-0503946 to the University of Maryland.

REFERENCES

- [1] A. Harris and J. Zmuidzinas, “A wideband lag correlator for heterodyne spectroscopy of broad astronomical and atmospheric spectral lines,” *Rev. Sci. Inst.*, vol. 72, pp. 1531–1538, Feb. 2001.
- [2] A. Harris, “Spectroscopy with multichannel correlation radiometers,” *Rev. Sci. Inst.*, vol. 76, pp. 4503–+, May 2005.
- [3] E. Wollack, W. Grammer, and J. Kingsley, “The Bøifot orthomode junction.” ALMA Memo 425, 2002.
- [4] S. Ramo, J. R. Whinnery, and T. van Duzer, *Fields and waves in communication electronics*. New York: John Wiley & Sons, Inc., 1965.
- [5] R. Norrod, “Cryostat cavity noise and the impact on spectral baselines.” NRAO Electronics Division Internal Report No. 318, 2007.
- [6] D. W. Allan, “Statistics of atomic frequency standards,” *Proc. IEEE*, vol. 54, pp. 221–230, 1966.
- [7] J. A. Barnes, “Atomic timekeeping and the statistics of precision signal generators,” *Proc. IEEE*, vol. 54, pp. 207–220, 1966.
- [8] J. Rutman and F. Walls, “Characterization of frequency stability in precision frequency sources,” *Proc. IEEE*, vol. 79, pp. 952–960, 1991.
- [9] R. Schieder, G. Rau, and B. Vowinkel, “Characterization and measurement of system stability,” in *Instrumentation for submillimeter spectroscopy* (E. Kollberg, ed.), vol. 598 of *Proc. SPIE*, pp. 189–192, Society of Photo-Optical Instrumentation Engineers, Bellingham, WA, 1986.
- [10] R. Schieder and C. Kramer, “Optimization of heterodyne observations using Allan variance measurements,” *Astron. Astrophys.*, vol. 373, pp. 746–756, 2001.