

**A PRELIMINARY REPORT: March 16, 2003
CALIBRATING THE GBT AT L, C, AND X BANDS**

SOC Library

Carl Heiles, Tim Robishaw, Tom Troland, D. Anish Roshi

ABSTRACT

During January 2003 we made extensive observations to obtain complete characterizations of the main beam parameters and polarization properties of L-, C-, and X-band systems on the GBT. This effort was successful, and we present this preliminary report whose purpose is to summarize what we consider to be the most interesting essentials without much explanation about details. A complete and detailed report is in preparation, which will present the complete smorgasbord of results together with commentary.

For Stokes I we display and discuss system temperatures, pointing errors, point-source gain, main beam efficiency, and beam ellipticity. For the polarized Stokes parameters (Q, U, V) we display and discuss beam squint and squash.

For polarization calibration, we present tabular results on the Mueller matrix coefficients for polarization calibration. We also discuss the frequency derivative of the relative cal phase with frequency, which gets as high as 0.9 rad/MHz.

1. EXECUTIVE SUMMARY

We quickly summarize only the points that we suspect are of interest to most people.

- In section 3.2, we estimate the accuracy of the absolute cal values by assuming that the point source gain is independent of frequency within a receiver band. If this assumption is correct, then the range of cal inaccuracy is about 10% at L band, 17% at C band, and 26% at X band.
- There are significant pointing errors, particularly for X band at high zenith angles (§3.3).
- The beam is significantly elliptical at X band. The ellipticity depends on zenith angle; near the zenith, the beam gets more elliptical (§3.4).
- The beam exhibits significant coma at X band. The coma varies with zenith angle, but there are too few data to determine the systematic behavior (§3.5).
- Stokes V beam squint is a function of frequency within each receiver band, which is a surprise (§4.1). Stokes (Q, U) beam squints are larger for L and C bands and exhibit similar behavior. (§5.1).

- Stokes (Q, U) beam squashes are much worse for L and C bands than for X band. This suggests that the L- and C-band feeds produce asymmetric primary surface illumination in orthogonal linear polarizations. The beam squash tends to flop randomly by 90° , which is equivalent to a change in sign; this renders its effect on observations unpredictable (§5.2).
- The X-band receiver, whose nominal polarization is circular, is nearly perfectly so at 8 GHz but significantly not so at 10 GHz (§6 and 6.3).

2. INTRODUCTION

During January 2003 we made extensive observations to obtain complete characterizations of the main beam parameters and polarization properties of L-, C-, and X-band systems on the GBT. We used the spectral processor in Stokes mode with four simultaneous inputs, each with bandwidth 5 MHz, whose frequencies spanned the range of each receiver band. The central frequencies are: L band, 1160, 1420, 1666, 1790 MHz; C band, 4034, 4700, 5400, 6032 MHz; X band, 8000, 9000, 9495, 10000 MHz. Typically we had several sources for each receiver.

We observed several sources over a range of zenith angle¹. We observed by taking two sets of crosses across the source, one cross aligned with (AZ, ZA) and one at 45° . This allows us to derive full information about beam properties including ellipticity and coma, the first sidelobe, and polarization calibration parameters.

3. STOKES I BEAM PARAMETERS

3.1. System temperatures

Figure 1 shows system temperatures, which are the average of the two polarization channels. Each symbol represents a different calibration source. For each band, the system temperatures cluster into 3 or 4 values. These clusters are for different frequencies within the band, of which there are four. Almost certainly, it is not the case that the system temperature changes so much with frequency. Rather, the derived system temperature depends on the cal value. These values are not perfectly accurate. The error in the adopted cal value changes with frequency, resulting in different system temperatures.

If the actual system temperatures are indeed independent of frequency, so that our derived ones reflect inaccuracies in the adopted cal values, then the range of errors in the cal values is equal to the range of derived system temperatures. These ranges are 22% for L band, 33% for C band, and 100% (!) for X band. For X band, the odd point at 43 K is for 8000 MHz.

¹In this report we use zenith angle ZA instead of elevation.




Fig. 1.— System temperatures (average of the two channels). Each source is represented by a different plot symbol.

Of course, there may also be some real variation in system temperature. In fact, we conclude this to be the case in §3.2 below.

3.2. Gain, beam efficiency

Figure 2 shows gains and beam efficiencies for the three bands. The left column of figures shows $KperJy$. Each figure tends to have three horizontal lines of points. These correspond to different frequencies. We sampled four frequencies for each receiver. The cal values for each frequency are in error by different amounts. Because $KperJy$ is dependent on the adopted cal value, errors in the cal values produce different values for $KperJy$.

If the actual $KperJy$'s are indeed independent of frequency, so that our derived ones reflect inaccuracies in the adopted cal values, then the range of errors in the cal values is equal to the range of derived $KperJy$. These ranges are 10% for L band, 17% for C band, and 26% for X band. These are smaller ranges than those obtained from system temperatures in §3.1 above, which argues that there is some real frequency variation in system temperature.

Mean values of $KperJy$ are about 1.9, 1.5, and 1.8 for L, C, and X bands respectively. This range is comparable to the range with frequency within each band. We conclude that there is no reason to suspect that the C-band point-source response is low compared to the other bands.

The right-hand column of figures shows the ratio of main beam efficiency η_B to $KperJy$, i.e. $\frac{\eta_B}{KperJy}$. This ratio is a convenient measure of the integrated solid angle of the main beam, and is independent of the adopted cal value. L-band values cluster around 0.47; if the point source response were really as high as $KperJy = 2.0$, as is indicated by one row of points, then this would mean that the beam efficiency $\eta_b = 0.94$, which would be phenomenal. C-band values look a bit smaller and exhibit a systematic increase with ZA , spanning the range of about 0.43 near the zenith to about 0.5 near the horizon.

The X-band ratios are smaller and cluster around 0.4 for intermediate ZA and decrease to about 0.35 near the zenith. This might be related to the beam becoming elliptical near the zenith (see §3.4).

3.3. Pointing errors

Figure 3 shows pointing errors ($delAZ$, $delZA$) in arcmin for the three bands. The left-hand column is $delAZ * \sin(ZA)$; the multiplication is necessary to convert to great-circle pointing error. The right-hand column is $delZA$. Errors are generally comparable to or smaller than the size of the plotted points.

AZ pointing errors are typically small enough compared to the beam. However, this is not




Fig. 2.— Left column is $KperJy$ and right column is $\frac{beam\ efficiency}{KperJy}$. For C band, points for $KperJy$ are absent where the source fluxes are not well enough known. Each source is represented by a different plot symbol.




Fig. 3.— Pointing errors in arcmin. Left column: AZ pointing errors [times $\sin(ZA)$] for the three bands. Right column: ZA pointing errors. Each source is represented by a different plot symbol.

always true for ZA . At L band the ZA pointing errors average around 0.6 arcmin near $ZA = 30^\circ$. We have $HPBW \sim 9$ arcmin, so the pointing error is about $\frac{HPBW}{15}$, which is “good enough”. C-band pointing errors can get as high as ~ 0.3 arcmin and the beamwidth is about 2.4 arcmin, so pointing errors are about $\frac{HPBW}{8}$; this is getting marginal. Finally, X-band pointing errors for small ZA can be as high as 0.5 arcmin, while near $ZA = 30^\circ$ they can be 0.2 arcmin. With $HPBW \sim 1.2$ arcmin, the pointing errors are typically $\frac{HPBW}{6}$. This is quite serious—for a point source, the flux is made noticeably smaller. We visually noticed this degradation for our one-dimensional crosses across the nominal source center.

3.4. Beam ellipticity

Figure 4 shows the beam ellipticity. The left-hand column of figures shows the fractional beam ellipticity, defined as

$$\text{fractional beam ellipticity} = \frac{HPBW_{max} - HPBW_{min}}{HPBW_{max} + HPBW_{min}}. \quad (1)$$

The right-hand column shows the angle of orientation ϕ of the major axis of the ellipse relative to $AZ = 0$. Errors on these points are most easily judged by the scatter—we didn’t want to clutter the plot with errorbars.

L-band results are in the top row. The ellipticity is small and, judging from the essentially uniform distribution in the position angle ϕ , too small to measure. C-band results are in the middle row; note the increase in scale from 0.1 to 0.2 for the ellipticity. The angle ϕ exhibits a bit less scatter, particularly at large ZA , meaning that our measurements of ellipticity are becoming meaningful (marginally so).

X-band results are at the bottom. Note the change in vertical scale! The X-band beam is highly elliptical, and the ellipticity increases to about 0.4 when looking straight up. The orientations of the ellipse tend to cluster at $\sim 65^\circ$ and 110° , which are fairly close to 90° , where it would be perpendicular to a line connecting the point of observation to the zenith.

We find this ZA variation of the ellipticity distinctly unusual. It perhaps indicates that the active surface adjustment, which was turned on for these measurements, is not perfect.

3.5. Beam coma

Figure 5 shows the beam coma. The left-hand column of figures shows the fractional beam coma in units of $HPBW$; for the precise definition, see Heiles et al. (2000) or Heiles et al. (2001). The right-hand column shows the direction of the coma relative to $AZ = 0$.




Fig. 4.— Beam ellipticity. Left column: fractional beam ellipticity (equation 1) for the three bands. Right column: position angle of major axis of ellipse relative to $AZ = 0$. Each source is represented by a different plot symbol.




Fig. 5.— Beam coma. Left column: fractional beam coma for the three bands in units of HPBW. Right column: position angle of the coma relative to $AZ = 0$. Each source is represented by a different plot symbol.

L-band results are in the top row. The coma is small and, judging from the essentially uniform distribution in the position angle ϕ , too small to measure. C-band results are in the middle row, and again the coma is negligible.

X-band results are at the bottom. The X-band beam has significant coma, as shown by the clustering in the coma angle ϕ . The coma changes with ZA . Near the zenith the coma clusters around 0.1 and its direction ϕ clusters around 230° . At higher ZA the coma exhibits much more scatter but a majority of angles cluster around $\phi = 160^\circ$, indicating that at least some of the coma directions are real. Understanding the coma behavior would require more data.

4. STOKES V BEAM PARAMETERS

4.1. Stokes V beam squint

Figure 6 shows the Stokes V beam squint. Beam squint is the angular difference between the directions that the two circular polarizations point. The left-hand column shows the magnitude of the beam squint in arcsec and the right-hand column shows ϕ_{squint} , the position angle of the beam squint relative to $AZ = 0$. Note that beam squint has a *direction*, not an orientation, so that the angles are modulo 360° .

L-band results are in the top row. A noticeable peculiarity is the frequency dependence of beam squint within the band. The lowest frequency, 1150 MHz, has typical beam squint ~ 1.3 arcsec with $\phi_{squint} \sim 240^\circ$. The HI line frequency has typically ~ 1.2 arcsec, with $\phi_{squint} \sim 180^\circ$. The OH line frequency has smaller beam squint, ~ 0.7 arcsec, with $\phi_{squint} \sim 0^\circ$. The fact that ϕ_{squint} rotates by $\sim 180^\circ$ —i.e. *reverses direction*—between the HI and OH line frequencies implies that the *magnitude* of the beam squint becomes zero at some frequency between those two. Finally, at 1789 MHz the beam squint goes to pot.

C band has lower beam squint than L band, as reflected in the half-size plotting scale. Like L band, C band also exhibits frequency-dependent beam squint. The squint is highest at the band edges. In fact, the $\sim 180^\circ$ rotation of ϕ_{squint} suggests that there is some frequency between 5400 (circles) and 6032 MHz (diamonds) where the squint magnitude becomes zero.

The maximum squint at X band is smaller than for the other bands, which is again reflected in the plotting scale. Again, the angle ϕ tends to rotate by $\sim 180^\circ$ between 9495 and 10000 MHz, suggesting the squint magnitude becomes zero.

We find the frequency dependence within each band highly surprising. The classical reason for beam squint in the prime focus situation is the feed not pointing directly at the vertex of the paraboloid. Thus, beam squint is a *geometrical* phenomenon. The secondary optics of the GBT were designed to minimize beam squint—again, a *geometrical* design choice. We emphasize that all results discussed in this section are derived for the properly calibrated Stokes parameters, so




Fig. 6.— Beam squint for Stokes V . Left column: beam squint magnitude in arcsec; *note* that the vertical scales differ for the three bands. Right column: position angle of beam squint relative to $AZ = 0$. Each frequency is represented by a different plot symbol.

there can be no question of contamination of Stokes V by (Q, U) .

The overall squint tends to decrease from L to X band. This trend is in line with the classical prime-focus model of beam squint, which predicts that beam squint should be proportional to HPBW. Nevertheless, this trend is overshadowed by the variations within each band. Therefore, it seems that the classical model has little relevance for the GBT.

4.2. Stokes V beam squash

Figure 7 shows data on beam squash. Beam squash is the angular difference between the HPBWs of the two circular polarizations. The left-hand column shows the beam squash in arcsec and the right-hand column shows the PA of the beam squash relative to $AZ = 0$. In contrast to beam squint, beam squash has fourfold symmetry with PA so it has an *orientation*, not a direction, so that the angles are modulo 180° —the same as polarization position angle.

As with beam squint, the frequency dependence of beam squash within each band is very noticeable. There is no classical model for beam squash in Stokes V . Thus, in contrast to beam squint, we have no prior expectation about frequency dependence of beam squash.

5. STOKES (Q, U) BEAM PARAMETERS

5.1. Stokes (Q, U) beam squint

Stokes (Q, U) exhibit beam squint. For L and X bands, the (Q, U) amplitudes are about three times larger than V . For X band, the (Q, U) amplitudes are comparable to V . In all bands there is frequency structure as there is in Stokes V . We don't present the figures here to save space. In the classical prime focus model, there should be no beam squint in linear polarization.

5.2. Stokes (Q, U) beam squash

Figures 8 and 9 show beam squash for (Q, U) . The left-hand column shows the beam squash in arcsec and the right-hand column shows the PA of the beam squash relative to $AZ = 0$. In contrast to beam squint, beam squash has fourfold symmetry with PA so it has an *orientation*, not a direction, so that the angles are modulo 180° —the same as polarization position angle.

Beam squash is the angular difference between the HPBWs of the two linear polarizations. This is caused by two effects. One is the interaction of the linearly polarized vectors with the curvature of the reflecting surface. The other is a difference between the two polarizations in illumination of the primary surface, which occurs from pseudo-waveguide modes in a feed.




Fig. 7.— Beam squash for Stokes V. Left column: beam squash magnitude in arcsec; *note* that the vertical scales differ for the three bands. Right column: position angle of beam squash relative to $AZ=0$. Each frequency is represented by a different plot symbol.




Fig. 8.— Beam squash for Stokes Q . Left column: beam squash magnitude in arcsec; *note* that the vertical scales differ for the three bands. Right column: position angle of beam squash relative to $AZ=0$. Each frequency is represented by a different plot symbol.




Fig. 9.— Beam squash for Stokes U . Left column: beam squash magnitude in arcsec; *note that the vertical scales differ for the three bands*. Right column: position angle of beam squash relative to $AZ=0$. Each frequency is represented by a different plot symbol.

There are several striking features of beam squash. By far the most striking feature is the drastically larger beam squash for L and C bands compared to X band: a factor ~ 20 ! This suggests that the L- and C-band feeds produce asymmetric primary surface illumination in orthogonal linear polarizations.

The second striking feature is the tendency of the squash angles ϕ to cluster into two groups separated by 90° . This is most striking for C band, but it also occurs for L and X band. A difference of 90° corresponds to the squash *changing sign*. This feature is most striking in the C-band plots where this change in sign tends to occur at will. This makes the squash angle ϕ unpredictable. This is not good for those who want to measure polarization of extended sources! Perhaps a more detailed study would reveal predictable trends, but this seems unlikely from the present data.

6. MUELLER MATRIX PARAMETERS

We do not derive the Mueller matrix elements directly. Rather, we derive five equipment-related parameters from which we calculate the matrix elements. As explained in §2, for each receiver we observed simultaneously at four frequencies. Typically we had several sources for each receiver. If a parameter varied with frequency for a receiver, then we fit a least squares polynomial.

We present the parameters in Table 1. For polarization measurements, the elements of primary interest are the leakage between linear Stokes parameters (Q, U) and the circular one V . These are expressed in terms of relevant matrix elements; for example, for a source with linear polarization expressed by Stokes parameter Q , m_{QV} expresses the amount of power from Q that falsely contributes to Stokes V .

6.1. Equipment-related parameter definitions

First we reiterate the definitions of the parameters; for more detailed information about these, see Arecibo's Technical and Operations Memo ATOM 2000-05 (Heiles et al. 2000) or Heiles et al. (2001). One uses these parameters to derive the Mueller matrix and then corrects the measured values with the inverse Mueller matrix to obtain the true ones. We emphasize that nonideal values for measured parameters do not constitute a problem for polarization measurement. The only difficulty arises if the parameters change with time or telescope position.

ΔG is the error in relative intensity calibration of the two polarization channels. It results from an error in the relative cal values (T_{calA}, T_{calB}). Our expansion currently takes terms in ΔG to first order only. Therefore, in cases like Rcvr8-10 at 8 GHz where $\Delta G = 0.21$, the other terms will be affected to some degree.

ψ is the phase difference between the cal and the incoming radiation from the sky. It redistributes power between (U, V) for a dual linear feed and between (Q, U) for a dual circular feed.

α is a measure of the voltage ratio of the polarization ellipse. Generally, the electric vector traces an ellipse with time; $\tan \alpha$ is the ratio of major and minor axes of the voltage ellipse. If a source having fractional linear polarization $P_{src} = \sqrt{Q_{src}^2 + U_{src}^2}$ is observed with a native circular feed that has $\alpha \neq 45^\circ$, then the measured Stokes V will change with $2PA_{\alpha z}$ and have peak-to-peak amplitude 4α .

ϵ is a measure of imperfection of the feed in producing nonorthogonal polarizations (false correlations) in the two correlated outputs. Our expansion takes ϵ to first order only. The only astronomical effect of nonzero ϵ is to contaminate the polarized Stokes parameters (Q, U, V) by coupling Stokes I into them at level $\sim 2\epsilon$; the exact coupling depends on the other parameters. For weakly polarized sources, this produces false polarization; for strongly polarized sources such as pulsars, it also produces incorrect Stokes I .

ϕ is the phase angle at which the voltage coupling ϵ occurs. It works with ϵ to couple I with (Q, U, V) .

6.2. Results

Table 1 shows the results. Data for C and X bands are consistent among sources and the fit quality was very good for all parameters.

However, for L band we have good data for only two sources, 3C138 and 3C286, and the data are not consistent for some parameters. The most serious inconsistency is with the parameter ΔG . At 1160 MHz, the worst frequency for consistency, the two sources had $\Delta G = (+0.070 \pm 0.006, -0.043 \pm 0.030)$ respectively. Only at 1666 MHz was the consistency good: the two sources had $\Delta G = (+0.064 \pm 0.005, +0.044 \pm 0.013)$ respectively. The discrepant results suggest a change in the cal or its coupling from one observation to the other.

6.3. The circularity of the X-band feed

The parameter α in Table 1 is the angle whose tangent specifies the ellipticity of the voltage ellipse of the output when observing a linearly polarized source. Thus, a perfect linearly polarized feed has $(\tan \alpha = 0)$ and a perfect circularly polarized one has $(\tan \alpha = 1)$.

Rcvr8-10 is supposed to be circularly polarized with $(\tan \alpha = 1)$ ($\alpha = 0.785$ radians). This expectation is fulfilled nicely at 8.0 GHz, but decays linearly with frequency above 8.0 GHz. These α values produce Mueller matrix element $m_{QV} = (+0.016, -0.061, -0.138)$ at frequency

Table 1: ADOPTED MUELLER MATRIX PARAMETERS, JAN 2003

RCVR	GHz	ΔG	ψ^a	α^a	ϵ	ϕ^a
R1-2	1.0	0.016 ^d	-0.14 ^b	3.0e-3	5.3e-3	-0.59
R1-2	1.4	0.016 ^d	-0.25 ^b	3.0e-3	5.3e-3	-0.59
R1-2	1.9	0.016 ^d	-0.38 ^b	3.0e-3	5.3e-3	-0.59
R4-6	4.0	+0.038 ^c	-0.13	2.2e-3	3.1e-3	-5.84 ^b
R4-6	5.0	-0.056 ^c	-0.13	2.2e-3	3.1e-3	-2.64 ^b
R4-6	6.0	-0.036 ^c	-0.13	2.2e-3	3.1e-3	+0.56 ^b
R8-10	8.0	-0.206 ^c	0	-0.78 ^b	7.3e-3 ^b	+3.18 ^b
R8-10	9.0	-0.032 ^c	0	-0.82 ^b	4.7e-3 ^b	-0.62 ^b
R8-10	10.0	-0.083 ^c	0	-0.86 ^b	2.1e-3 ^b	-4.42 ^b

a : Angles are in radians.

b : Linear fit.

c : Quadratic fit.

d : Internally inconsistent results and a poor fit; values are a straight average and not accurately representative.

(8.0,9.0,10.0) GHz. In other words, at 8 GHz the feed is almost perfect while at 10 GHz about 14% of the power of a linearly polarized source ends up in Stokes *V*.

6.4. "Cross-Coupling" between polarizations

Rcvr1-2 is linearly polarized, so Stokes (*U*, *V*) comes from the correlated outputs. If cables are perfectly cut so that the relative phases of the cal are identical to those of the incoming radiation through the feed, then $\psi = 0^\circ$ and, correspondingly, $m_{UV} = 0$. Such an ideal state is unlikely, and is one of the important reasons for deriving the Mueller matrix elements. For Rcvr1-2, we have $m_{UV} = (-0.138, -0.243, -0.370)$ at frequency (1.0, 1.4, 1.9) GHz.

We emphasize that these nonideal values should not be considered a problem or engineering fault; they are inevitable at some level for any system. The polarization calibration measures the cross coupling and makes accurate polarization measurements possible.

6.5. Relative cal accuracy

One uses the cal to determine relative gains of the two channels. These gains are complex, with an amplitude and phase.

6.5.1. *Relative cal amplitude*

If the relative cal amplitude is not perfectly calibrated, this induces a coupling between Stokes (I, Q) in a linearly polarized feed and between (I, V) in a circularly polarized one (both are removed by our calibration). This is reflected in ΔG in the table; we have, for example, $m_{I,Q} \approx \frac{\Delta G}{2}$ for a linearly polarized receiver.

The nonzero ΔG values are purely a result of imperfect *relative* cal calibration values. In fact, our measured values of ΔG can be used to correct the tabulated values. By comparing these values with those in our discussion of §3.2, we conclude that the absolute cal calibration is not as good as the relative calibration. This is hardly unexpected!

6.5.2. *The frequency derivative of the relative cal phase*

If the relative cal phase is not perfectly calibrated, this is equivalent to an error in ψ , which couples (U, V) for a dual linear feed and (Q, U) for a dual circular feed.

Figure 10 exhibits the frequency derivative of the relative cal phase between the two polarization channels for each band. This derivative is produced by the difference between the two path lengths between the correlated cal and the point where the cross-multiplication is performed. Contributions to this path difference are produced at r.f., i.f., and optical fiber. The frequency derivative is proportional to this path difference measured in units of wavelength.

The fact that no receiver shows a systematic increase of cal slope with frequency means that the contribution at r.f. does not dominate. The fact that the slopes differ markedly for the four spectral processor inputs means that most of the contribution comes from the i.f. *after* the optical fiber, because the optical fiber would contribute a common difference to all four inputs.

The presence of this slope is not serious because our polarization calibration procedure removes it (with exquisite accuracy). However, people should be aware that different spectral processor inputs have differing path lengths; one cannot use the calibration results from one input for another one.

The GBT staff went out of its way to assist us and bend the system to perform to our demanding specifications. We very much appreciate the efforts of Frank Ghigo, who was always available and provided essential help with general software issues, GLISH scripts, and the X-band system. Roger Norrod and Steve White sacrificed both days and evenings to provide invaluable help with electronics issues. Mike Stennis braved the cold to perform the helix-based calibration of circular polarization. Astronomers Dana Balsler, Rick Fisher, and Toney Minter helped extensively with observing scripts and the telescope system. Mark Clark and Ron Maddalena provided needed information on several software issues. Carl Bignell bent the schedule to give us replacement time when weather and equipment problems prevented us from observing. Phil Jewell provided the




Fig. 10.— Frequency derivative of the relative cal phase between the two polarization channels. Errors of each point are well below the size of that point.

encouragement and resources necessary to perform this project.

This work was supported in part by NSF grants AST-0097417 to CH and AST-9988341 to TT, and NSF award GSSP02-0011 from the NRAO to TR.

REFERENCES

Heiles, C. et al., 2000, ATOM2000-05, available at www.naic.edu/~astro/aotms/performance.shtml

Heiles, C. et al., 2001, PASP, 113, 1247 and 1274.

