Frequency-switched bandpass calibration at the GMRT

Nirupam Roy[1] & Nissim Kanekar[2]
[1]*National Centre for Radio Astrophysics, Pune, India;*
[2]*National Radio Astronomy Observatory, Socorro, USA*

March 26, 2007

Abstract

We report GMRT H I 21cm absorption studies of three bright, compact background sources to test the spectral dynamic range that can be achieved with frequency-switched bandpass calibration at the GMRT. Two approaches were used, each on two sources, with switching at the first and fourth local oscillators (LOs) by 5 MHz and 1 MHz respectively. Switching at the fourth LO is found to yield a complex spectral baseline, with large-scale ripples and a relatively low dynamic range (< 300) for the detection of wide absorption features. Conversely, the spectral baseline obtained by switching at the first LO is quite flat, with RMS noise values and spectral dynamic ranges similar to those obtained with standard bandpass calibration. GMRT observations that require frequency-switched bandpass calibration should thus implement this switching at the first LO, to ensure spectral baselines of the highest quality.

1 Introduction

It is necessary to calibrate the frequency response of a radio interferometer, in amplitude and phase, to correct for changes in the complex gain as a function of frequency; this is referred to as bandpass calibration. The standard approach is to use observations of a strong source without any spectral features in the observing band (a "bandpass calibrator") to calibrate the frequency response. Ideally, the bandpass shape remains unchanged with time and a single observation of the bandpass calibrator is sufficient to measure this shape. In practice, however, the antenna bandpasses drift with time and calibration must be carried out more often, especially for projects that require a high spectral dynamic range ($\gtrsim 1000$). Of course, sufficient time must be spent on the bandpass calibrator to ensure that the noise on the final source spectrum is limited by the on-source time and not dominated by the noise on the bandpass. This essentially implies that the bandpass calibrator must be observed for a time interval of at least $t_{BP} = [S_{SRC}/S_{BP}]^2 \times t_{SRC}$, where S_{BP} and S_{SRC} are the flux densities of the bandpass calibrator and the target source respectively, and t_{SRC} is the total time spent on the target (e.g. [1]). Note that t_{BP} is only the time for which the bandpass calibrator

must be observed in order to make the signal-to-noise ratios on the calibrator and the source equal; it is preferable to observe the calibrator for more time than this to ensure that it does not make a significant contribution to the signal-to-noise ratio in the final cube.

The above "standard" route towards bandpass calibration fails in two cases, if (1) it is not possible to find a calibrator without spectral features in the observing band or (2) the target source is so strong that a large fraction of the total time is needed on a (weaker) bandpass calibrator to reach a signal-to-noise ratio comparable to that on the target. For example, the first case arises in studies of Galactic H I, where H I 21cm absorption is likely to be present towards any possible bandpass calibrator. In this situation, the final spectrum towards the target source (with standard bandpass calibration) would be corrupted by the absorption against the calibrator, making it impossible to detect features weaker in optical depth than those towards the calibrator.

The usual solution to this problem is to employ an alternative approach to bandpass calibration, involving frequency-switching. Here, one measures the shape of the bandpass at a frequency slightly offset from the frequency of interest, making the assumption that this shape is not a strong function of frequency. For studies of Galactic H I, the frequency offset is usually chosen to ensure that the observing band towards the calibrator is free of spectral features. For strong sources, the target source itself can be used as the bandpass calibrator, implying that only half the total time need be spent on calibration, instead of the far larger fractions for the case of a weaker calibrator. In fact, it is sometimes even possible to use a frequency offset which is a fraction of the bandwidth, to ensure that the frequency of interest remains within the observing band at all times ("in-band frequency-switching"; e.g. [3]). Here, of course, one is essentially determining the *local* shape of the observing band around the spectral feature of interest, with the advantage that half the observing time is not "wasted" on calibration.

The crucial questions for frequency-switched bandpass calibration are (1) the location at which the frequency-switching is implemented (i.e. which local oscillator is to be switched), (2) the timescale on which the switching must be carried out, and (3) the spectral dynamic range that can be achieved. We note that frequency-switching is routinely used at radio interferometers such as the Australia Telescope Compact Array (ATCA) and the Westerbork Synthesis Radio Telescope (WSRT) and that it has been possible to achieve spectral dynamic ranges of > 1000 at both these telescopes, with frequency-switching on a timescale of five minutes (e.g. [3, 5]). In the case of the GMRT, to the best of our knowledge, the only detailed study based on this technique is that of [6]. These authors achieved a spectral dynamic range of $\sim 50-140$ per ~ 3.3 km s^{-1} channel along 14 lines of sight, with frequency-switching carried out every two hours, at the fourth local oscillator (see below). Further, these observations were carried out in the early days of the GMRT, in 1999, with only $8-10$ working antennas. We have recently carried out a series of observations with the full GMRT array to test the quality of spectral baselines that can be achieved with frequency-switching at the GMRT and report here the results of these studies.

2 Observation and Data Analysis

The obvious approach to test the spectral dynamic range achievable with frequency-switching is to attempt to detect known, weak absorption lines towards a strong background source. The WSRT Galactic H I absorption profile towards the standard calibrator 3C286 shows three narrow spectral components with optical depths ranging from ~ 0.005 to ~ 0.007 [3]. This was chosen as the primary target for the project. In addition, we also attempted to detect the Galactic H I absorption towards the calibrators 3C295 [7] and 3C48. All observations used the GMRT L-band receivers, with a bandwidth of 1 MHz centred at 1420.4 MHz and sub-divided into 256 channels, with the USB-High-Res correlator mode. This gave a velocity resolution of ~ 0.8 km s^{-1} per channel, before any smoothing.

Frequency-switching at the GMRT can be implemented at either the first or the fourth local oscillators (LOs). The first LO can be tuned at steps of 5 MHz at frequencies above 350 MHz and at steps of 1 MHz below 350 MHz, while the fourth LO is tunable at steps of 0.1 kHz at all frequencies covered by the GMRT. Note that the smaller step-size at the fourth LO allows the possibility of the in-band frequency-switching mentioned above; in-band switching is unlikely to be possible at the first LO except for fairly wide lines, due to the large step-size. In addition, changing the fourth LO takes only a few seconds while the standard method of changing the first LO takes of the order of 1.5 minutes, far too large if the switching is to be carried out on a timescale of five minutes. The best situation would thus be if a high spectral dynamic range might be achieved with frequency-switching at the fourth LO. Both approaches were tried for 3C286, in order to compare the resulting spectral dynamic range, while 3C48 was observed with switching only at the first LO and 3C295, with switching at the fourth LO.

The observations of 3C286 and 3C295, with switching at the fourth LO, were carried out on August 6, 2006. In each case, the switching was carried out by alternately shifting the band above and below the central observing frequency by 1.0 MHz every five minutes (within the command file). The typical timescale required to change frequencies was 2–3 seconds. The total on-source, on-frequency times were ~ 1 and ~ 2 hours for 3C286 and 3C295, respectively, with equal off-frequency times for each source. Half the off-frequency time was alternately spent at frequency settings above and below the target frequency.

The observations of 3C286 and 3C48, with switching at the first LO, were carried out on September 2 and 4, 2006, respectively. Here, the switching was carried out by alternately shifting the band above and below the central observing frequency by 5.0 MHz every five minutes (again within the command file). The typical timescale required to change frequencies was found to be ~ 40 seconds, after modifying the code used to set the first LO (kindly done by Jitendra Kodilkar). The total on-source, on-frequency time for each source was ~ 1 hour, again with equal time being spent on off-frequency settings.

In passing, we have also, more recently, tested the new "intelligent Monitor and Control Module" (iMCM) scheme, which, in principle, allows the first LO to be tuned on a shorter

timescale. Using this scheme from the standard *USER4* control window reduces the time taken to set the first LO to $\sim 30 - 35$ seconds (i.e. marginally better than the original approach). Using the same iMCM scheme from the *USER0* control window reduces the above switching time to ~ 1 second (if command monitoring is disabled), thereby significantly reducing the observational overheads. Unfortunately, it is currently not possible to control the GMRT Data Acquisition System from *USER0* via a command file. We hence conclude that the frequency-switching at GMRT is best carried out via the iMCM scheme from the *USER4* control window; this is the approach detailed in the Standard Operating Procedure set out in Appendix 1.

All data were analysed in NRAO AIPS, using standard procedures. About half the baselines in all data sets were found to be dead or misbehaving for a variety of reasons (e.g. correlator spikes in one side-band) and were edited out (Table 1 lists the number of baselines used in the final spectrum for each source). Gain calibration was straightforward for 3C286 and 3C48, which are phase calibrators for the GMRT. In the case of 3C295, which is a phase calibrator only on the shorter baselines, a few rounds of phase and then amplitude-and-phase self-calibration were carried out to obtain the antenna-based complex gains and only baselines within the range $0 - 15$ kλ were used. Bandpass calibration was applied using interpolation between bandpass solutions (i.e. DOBAND=3 in AIPS). Next, continuum subtraction was carried out with the task UVLIN, fitting a linear baseline to line-free channels for each visibility spectrum, and the task CVEL then used to shift the residual U-V data to the LSR frame. The residual data were imaged in all channels (using ROBUST=4 in AIPS), and the final spectrum then obtained by a cut through the dirty cube at the location of the target source. Spectra were also obtained by directly vector-averaging the complex visibilities, for comparison. In addition, in the case of 3C295, which is an extended source, continuum-subtraction was also carried out using the task UVSUB, followed by UVLIN, before imaging; this did not produce any noticeable difference. The spectrum from the original (non-UVSUB) analysis will be used below.

Finally, we have also recently carried out somewhat deeper (~ 4 hours on-source, on-line) frequency-switched observations of the calibrator 0318+164, in GMRT proposal 11NRb01, with switching at the first LO and a spectral resolution of ~ 0.4 km/s (256 channels across a bandwidth of 0.5 MHz). While the results of these observations will be discussed elsewhere, we merely note here that a slightly different scheme of determining the bandpass (using an antenna-based division by channel-0 during the determination of the bandpass shapes, setting BPASSPRM(5)=1, BPASSPRM(10)=3 in the AIPS task BPASS) gave a significantly improved spectral baseline. We also attempted this calibration strategy for the present datasets and found no significant change in the results (except for a 20% improvement in the RMS noise for the LO-1-switched dataset towards 3C286). The results discussed below use the original analysis, with baseline-based division by channel-0 during the determination of the bandpass shapes.

3 Results

The final spectra towards 3C286 and 3C295, with switching at the fourth LO, are shown in Figures (1) and (3), while those towards 3C286 and 3C48, with switching at the first LO, are shown in Figures (2) and (4). Observational details and results are summarized in Table 1 while the following two sub-sections discuss the results from the two different switching approaches at greater length.

3.1 Frequency switching at the fourth LO

The top two panels of Figure 1 show the vector-averaged RR and LL visibility spectra towards 3C286, while the bottom panels of the figure show these spectra towards 3C295, with frequency-switching carried out at the fourth LO; in each case, continuum-subtracted optical depth is plotted against channel number. Figure 3 presents the final Stokes I spectra towards the two sources, obtained from a cut through the spectral cube at the source location, with optical depth plotted against LSR velocity, in km/s. While the known narrow Galactic absorption features [3, 7] are indeed detected towards both sources, each pair of spectra in Figure 1 can be seen to have a complex baseline, with large-scale ripples and a sharp drop in the band beyond $\sim$ channel 170. Further, the baseline shapes are very different in the two polarizations and it is not possible to model these shapes with a low-order (order ≤ 2) polynomial. In passing, we note that a fifth-order polynomial was found necessary to obtain a flat baseline for the spectrum towards 3C286; even after this, the relative strengths of the absorption features in the residual spectrum was found to not match those of the WSRT spectrum of [3], implying that the baseline had still not been modeled correctly.

The detectability of weak features in a spectrum is quantified by the spectral dynamic range, SDR, usually defined as the ratio of the continuum flux density to the RMS noise. However, this definition does not provide a good indicator of the ability to detect weak features in a spectrum contaminated by the presence of systematic effects such as ripples. For example, the "local" RMS noise values in the Stokes I spectra of Figure 3, obtained from a small, relatively flat portion of the band (see Table 1), yield a fairly high spectral dynamic range for the detection of narrow, isolated features. However, these RMS values are in no way representative of the quality of the spectral baselines or the ability to detect wide features. Specifically, the large-scale ripples in the band imply that the RMS would not reduce $\propto N^{-1/2}$ (and, in fact, might not reduce at all), on averaging N channels together.

A rough estimate of the spectral dynamic range for wide features, SDR_{wide} can be obtained by the ratio of the continuum flux density to $\sim (1/2.8)$ times the peak-to-peak spread of the ripple in the baseline. This uses the fact that the peak-to-peak amplitude of the simplest ripple (a sine-wave) is ~ 2.8 times the RMS of the wave. Thus, for a baseline with a single sine-wave, SDR_{wide} is just the ratio of the continuum flux density to the RMS noise of this sine-wave, in analogy with the standard definition. Of course, the ripples in real spectra,

such as those in Figure 3, are far more complex than a single sine-wave, but this definition of SDR_{wide} at least provides a handle on the detectability of wide features as limited by the largest-amplitude sine-wave present in the spectral baseline. Note that smoothing to coarser resolutions would not usually significantly reduce the peak-to-peak value of a systematic ripple; SDR_{wide} hence does *not* increase $\propto N^{1/2}$ for N-channel-wide features.

For the spectra in Figure 3, $SDR_{wide}(N)$ was estimated after subtracting out the best 2nd-order polynomial fit to line-free channels, excluding ~ 35 channels at the edges of the band. The peak-to-peak for the residual Stokes I spectrum towards 3C286 is ~ 160 mJy, giving $SDR_{wide}(N) = (2.8 * 15.5/0.16) \sim 270$, while that towards 3C295 is ~ 260 mJy, giving $SDR_{wide}(N) = (2.8 * 23.4/0.27) \sim 240$. Both values are far lower than the spectral dynamic range for narrow features, SDR; clearly, these spectra are not very sensitive to wide absorption. We conclude that the complexity of the spectral baseline obtained with fourth-LO switching would render impossible the detection of weak, wide absorption features.

3.2 Frequency switching at the first LO

The vector-averaged RR and LL visibility spectra towards 3C286 and 3C48, with frequency-switching at the first LO, are shown in the top and bottom pairs of panels in Figure 2, with continuum-subtracted optical depth plotted against channel number. The final Stokes I spectra towards the two sources are in Figure 4, with the known Galactic absorption features again clearly detected. Note that no baseline has been fit to these spectra, beyond the linear baseline subtracted from each visibility spectrum in UVLIN. The single-channel RMS noise is similar to the "local" RMS on the spectra of the previous section, obtained from switching at the fourth LO. However, the major difference here is that the spectral baselines are quite flat (excluding edge channels), without any obvious systematic effects. Further, the RMS noise values for these spectra (see Table 1) were computed over a large fraction of the observing band (channels $\sim 10-230$, excluding absorption channels and a few channels affected by RFI) and are hence representative of the quality of the entire spectrum. The peak-to-peak spreads in both spectra are also consistent with random noise characterized by the RMS values listed in Table 1. The spectral dynamic range for an N-channel-wide feature (where N is much smaller than the total number of channels) should thus be $SDR_{wide}(N) \sim SDR/N^{1/2}$. This is comparable to the best values obtained with normal bandpass calibration (using a separate bandpass calibrator at the same observing frequency) at the GMRT (e.g. [4]). We conclude that it is possible to obtain spectral dynamic ranges $\gtrsim 1000$ for narrow bandwidths, with frequency-switching at the first LO.

3.3 Comparison with theoretical expectations

The eighth line of Table 1 lists the expected theoretical RMS values per 0.8 km/s channel, assuming a gain of $G = 0.22$ K/Jy and a system temperature of $T_{\rm sys} \sim 76$ K (plus the H I

emission brightness temperature in each direction, also listed in the table, and a $\sim 4-5$ K contribution from the flux density of each source). The measured RMS noise values are in all cases significantly larger than the theoretically-expected ones, typically by a factor of ~ 2.5. The source of this discrepancy is unclear. One obvious issue is that the above $T_{\rm sys}$ and gain values are for the centre of the GMRT L-band (~ 1280 MHz) where the sensitivity is higher than at 1420 MHz. However, the difference in sensitivities is nowhere near as much as a factor of 2.5. We note that other short observations with "normal" bandpass calibration at ~ 1420 MHz have also obtained RMS values significantly larger than theoretical ones ([2]; Ekta, private communication) and there are suggestions (e.g. the GMRT specifications page) that a "fudge-factor" of ~ 2 is typical for short observations. While this is an issue that needs further investigation, we note that our more-recent deeper (~ 8 hour) frequency-switched observations of other calibrators (e.g. 0318+164; see below) have indeed obtained RMS noise values within a factor of $\sim 20\%$ of the theoretically-expected ones.

4 Future work

The present observations only tested the spectral dynamic range achievable with frequency-switching using a bandwidth of 1 MHz; we have also more recently found that similarly high values can be obtained with a bandwidth of 0.5 MHz. It would be interesting to test whether the high spectral dynamic range persists for other (especially larger) bandwidths. Next, the observational overheads presently contain a large contribution from the actual time required to switch frequencies ($\sim 30-35$ seconds); this could be reduced to negligible levels if the command file could be run from the *USER0* control window, which is presently not possible. All the observations discussed here were limited by their relatively short spans (~ 1 hour of on-frequency time in most cases) as well as problems with antennas and the correlator, resulting in typically half the baselines being unusable (but see below). This limited the dynamic range to $\sim 1000-1400$, at a resolution of ~ 0.8 km/s. Further, as noted in the preceding section, the observed RMS noise was found to be a factor of ~ 2.5 higher than the theoretically-expected value. It would be interesting to carry out a full-synthesis observation of bright sources like 3C48 or 3C286, with switching at the first LO, to test whether the dynamic range continues to improve as expected or whether any systematic effects are observed that might be limiting factors in frequency-switched spectral line observations with the GMRT.

We note, finally, that deeper observations of the calibrator 0318+164, in GMRT project 11NRb01, obtained a spectral dynamic range of ~ 1400 per ~ 0.4 km/s channel in an ~ 8 hour integration, with switching at the first LO; this is quite close to the theoretical value. While this suggests both that the spectral dynamic range continues to improve as expected for long integrations and that the fudge-factor is close to unity for such integrations, further deep observations should be carried out to confirm this result.

5 Summary and conclusions

We have carried out Galactic H I 21cm absorption studies of three bright, compact background sources to test the spectral dynamic range that can be achieved with frequency-switched bandpass calibration at the GMRT. Two approaches were used, each on two sources, switching at the first and fourth LOs by 5 MHz and 1 MHz respectively. The switching was carried out every five minutes, alternately shifting the band above and below the central observing frequency by the above offsets. We find that switching at the fourth LO yields a complex spectral baseline, with large-scale ripples. While the "local" RMS noise on these spectra gives a high spectral dynamic range ($\gtrsim 1000$) for isolated narrow absorption features, the ripples in the baseline imply a relatively low dynamic range ($\lesssim 300$) for the detection of wide absorption. On the other hand, the spectral baseline obtained by switching at the first LO is quite flat, with RMS values and peak-to-peak spreads consistent with a normal distribution. The resulting spectral dynamic range is $\gtrsim 1000$ for narrow features and $\sim 1000 \times N^{1/2}$ for features that are N channels wide (where N is significantly smaller than the total number of channels). We conclude that GMRT observations that require frequency-switched bandpass calibration should implement this switching at the first LO, to ensure spectral baselines of the highest quality.

Acknowledgments

We thank Jitendra Kodilkar for modifying the code used for setting the first LO, Jitendra Kodilkar and Sanjay Kudale for help with the observations, Jayaram N. Chengalur, T. L. Venkataramani and Yashwant Gupta for discussions about the observational procedure, and Jayaram N. Chengalur, Ekta and Viswesh Marthi for comments on an earlier version of this manuscript.

References

[1] Rupen M. P. 1999, in *Synthesis Imaging in Radio Astronomy II*, G. B. Taylor, C. L. Carilli & R. A. Perley (eds.), ASP Conference Series Vol. 180, 229

[2] Begum A. & Chengalur J. N. 2005, MNRAS, 362, 609

[3] Braun R. & Kanekar N. 2005, A&A, 436, L53

[4] Kanekar N. & Chengalur J. N. 2002, A&A,

[5] Kanekar N., Subrahmanyan R., Chengalur J. N. & Safouris V. 2003, MNRAS, 346, L57

[6] Mohan R., Dwarakanath K. S., Srinivasan G. & Chengalur J. N. 2001, JApA, 22, 35

[7] Dwarakanath K. S., Carilli C. L. & Goss W. M., 2002, ApJ, 567, 940

A Standard Operating Procedure for Frequency-Switching at the GMRT

It is quite straightforward to carry out frequency-switching at the GMRT, with minimal changes in the standard observational procedure. The only real difference is that the first LO and the TPA must be set within the command file. In addition, one should, as far as possible, reduce the use of non-critical commands such as "*stabct*", "*sndsacsrc*", etc, as each of these take up $\sim$ 20 seconds of time, thus significantly adding to the overheads (as the actual scans are only five minutes in duration). We have hence used the *stabct* and *sndsacsrc* commands only once apiece for each set of four scans.

For convenience during later analysis, it is simplest to use different source names for the on-line (i.e. target frequency) and off-line (i.e. bandpass) scans of the same source. These names, along with source co-ordinates, must be added to the external file containing the list of sources. We also recommend that switching be carried out alternately at frequencies above and below the target frequency, by the lowest acceptable first LO increment (i.e. 5 MHz at frequencies above 350 MHz and 1 MHz at lower frequencies, if possible).

An example of a command file to observe the Galactic H I 21cm line towards 3C48, with frequency-switching above and below the line frequency (1420.4058 MHz) every five minutes, is given below. Note that the fourth LO is assumed to have been set (to 65.4058 MHz) *before* running the command file. The source has been labelled as "3C48" for on-line scans (LO-1 = 1355 MHz), "3C48A" for below-line scans (LO-1 = 1350 MHz) and "3C48B" for above-line scans (LO-1 = 1360 MHz) and these source names have been added to the external file containing the list of sources (here, `/temp2/data/source/nirupam.list`).

Finally, if the data are to be analysed in the standard package NRAO AIPS, the scan keywords `FIRST_LO` and `FREQ` of the off-frequency scans must be changed to those of the target frequency using the GMRT off-line program `ltaedit` (written by Jayaram N. Chengalur) before converting the data into FITS format. One may then follow the standard AIPS reduction procedure, using the off-frequency scans for bandpass calibration.

```
*BEGIN COMMAND FILE
cmode 1
subar 4
addlist '/temp2/data/source/nirupam.list'
lnkndasq
mpa 2 2 3
comm 29;dest 17;t3v
allant
cp 0;defs 4;subar 4
subar 4
goout
```

```
gosacout
$1

ana 1350 1350; comm 160;dest 2;t3v; dest 3;t3v
tpa 1415.4058 1415.4058 1350 1350 65.4058 65.4058
prjfreq
gts '3C48A'
sndsacsrc(1,12h)
stabct
strtndasc
time  5m
stpndasc

ana 1355 1355; comm 160;dest 2;t3v; dest 3;t3v
tpa 1420.4058 1420.4058 1355 1355 65.4058 65.4058
prjfreq
gts '3C48'
strtndasc
time  5m
stpndasc

ana 1360 1360; comm 160;dest 2;t3v; dest 3;t3v
tpa 1425.4058 1425.4058 1360 1360 65.4058 65.4058
prjfreq
gts '3C48B'
strtndasc
time  5m
stpndasc

ana 1355 1355; comm 160;dest 2;t3v; dest 3;t3v
tpa 1420.4058 1420.4058 1355 1355 65.4058 65.4058
prjfreq
gts '3C48'
strtndasc
time  5m
stpndasc

goto 1
end
*END COMMAND FILE
```

Table 1: Summary of the observations and results

Frequency-switching at	First LO		Fourth LO	
Sources	3C286	3C48	3C286	3C295[c]
Frequency increment[a] (MHz)	± 5	± 5	± 1	± 1
Bandwidth (MHz)	1	1	1	1
Spectral resolution (km/s)	0.8	0.8	0.8	0.8
On-frequency time (Hours)	$\sim$1	$\sim$1	$\sim$1	$\sim$2
Number of baselines[b]	467	420	437	120
Continuum flux density (Jy)	15.1	16.0	15.5	23.4
Brightness temperature T_b	3 K	10 K	3 K	2 K
Theoretical RMS (mJy)	~ 4.6	~ 5.3	~ 4.8	~ 6.5
Single-channel RMS[c] (mJy)	14.8	13.8	13.1	16.7
Dynamic range[d], SDR	~ 1020	~ 1160	~ 1190	~ 1400
Dynamic range[e], $SDR_{wide}(N)$	$\sim 1020 * N^{1/2}$	$1160 * N^{1/2}$	~ 270	~ 240

[a] Alternately above and below the central frequency, every five minutes

[b] Total number of baselines used in the final spectrum, summing the two polarizations. In the case of 3C295, only baselines out to 15 kλ were used.

[c] For switching at the first LO (columns $1-2$), this is the RMS noise across the entire line-free band, excluding edge channels. For switching at the fourth LO (columns $3-4$), this is the "local" RMS noise, from a small, "flat" part of the band.

[d] The spectral dynamic range for the detectability of narrow features, the ratio of continuum flux density to single-channel RMS.

[e] The spectral dynamic range for the detectability of N-channel-wide features. See text for the definition of this quantity and discussion.

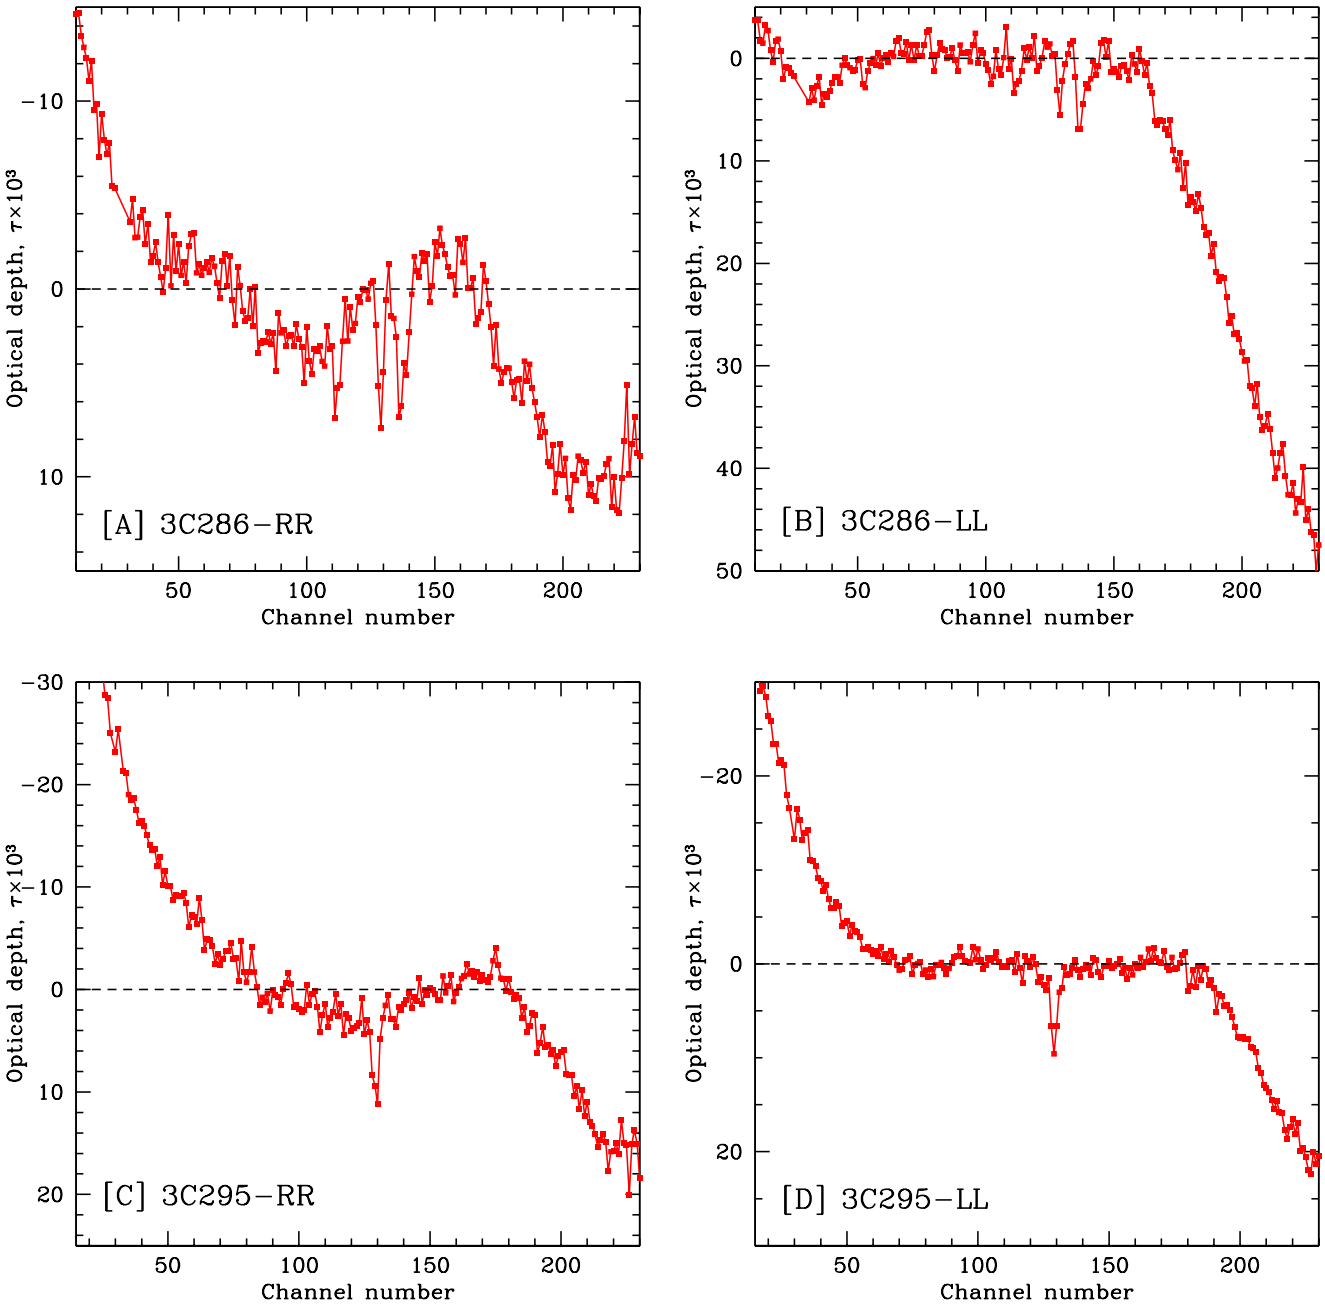

Figure 1: GMRT H I 21cm spectra in RR and LL polarizations towards (top panels) 3C286 and (bottom panels) 3C295, with switching at the fourth LO. All spectra contain optical depth, $10^3 \times \tau$, plotted against channel number. The large-scale ripples in the spectra as well as the sharp drop beyond $\sim$ channel 170 are clearly visible. See text for discussion.

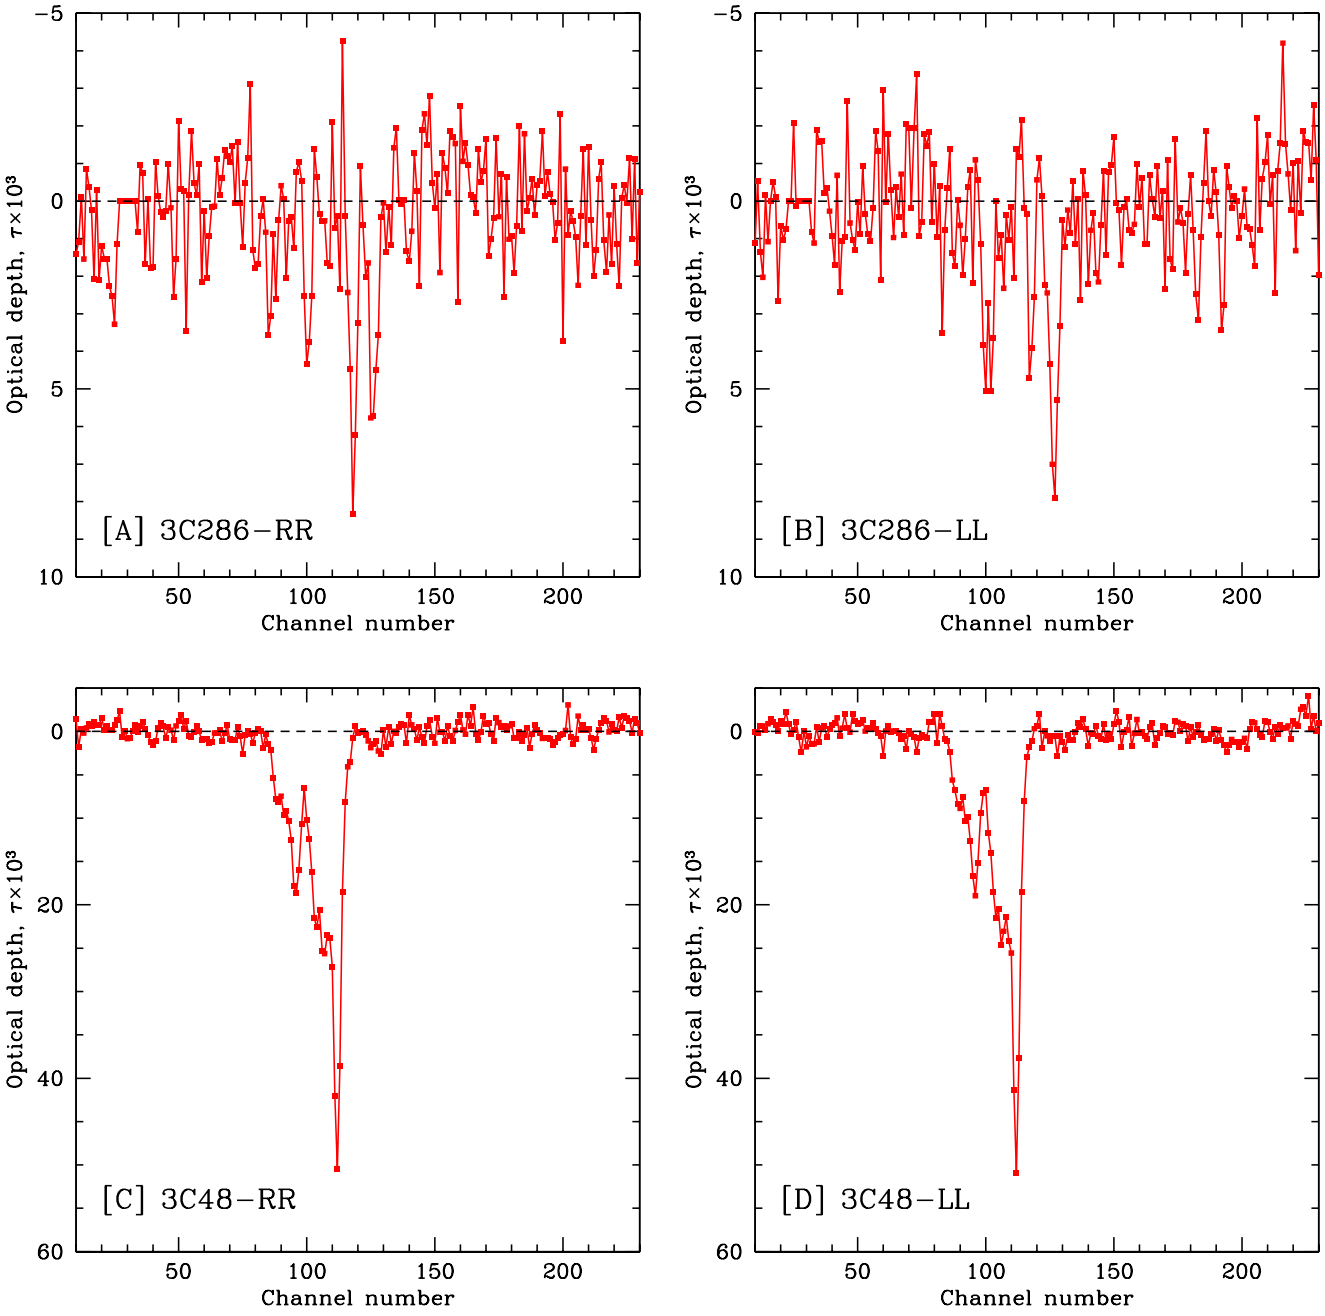

Figure 2: GMRT H I 21cm spectra in RR and LL polarizations towards (top panels) 3C286 and (bottom panels) 3C48, with switching at the first LO. All spectra contain optical depth, $10^3 \times \tau$, plotted against channel number. The spectral baselines are seen to be much flatter than those in the latter figure. See text for discussion.

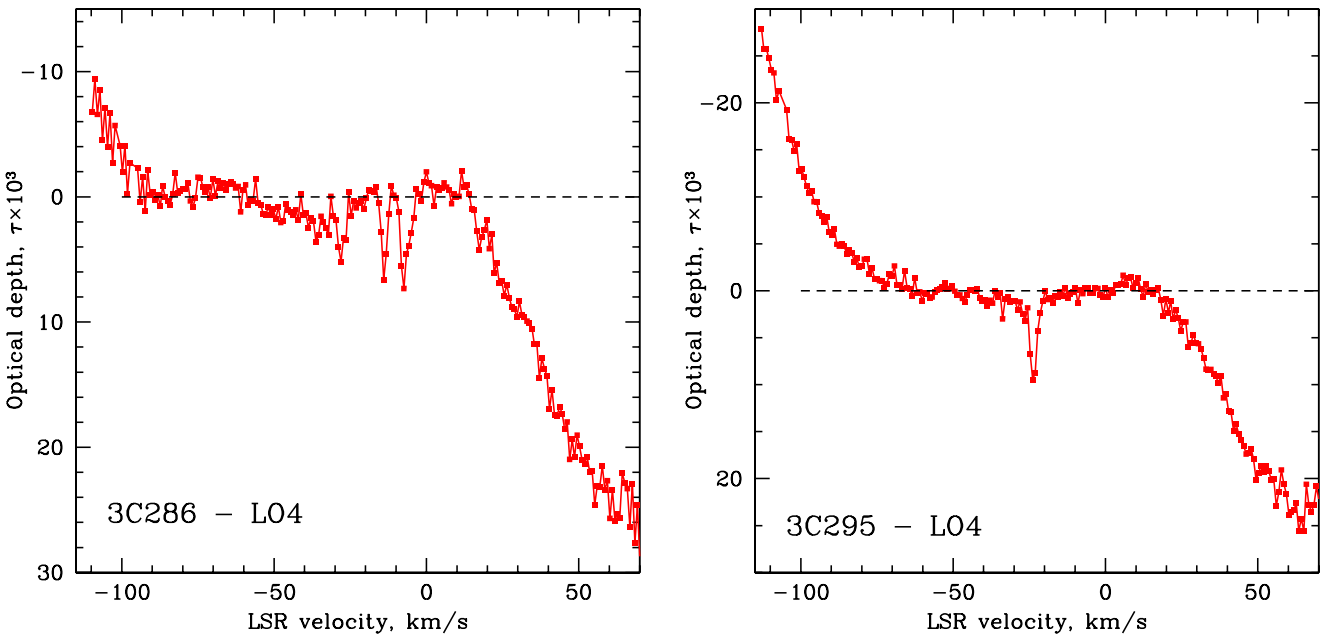

Figure 3: Final GMRT H I 21cm Stokes I spectra towards [A] 3C286 (left panel) and [B] 3C295 (right panel), with switching at the fourth LO; in both cases, optical depth, $10^3 \times \tau$, is plotted against LSR velocity, in km/s. See text for discussion.

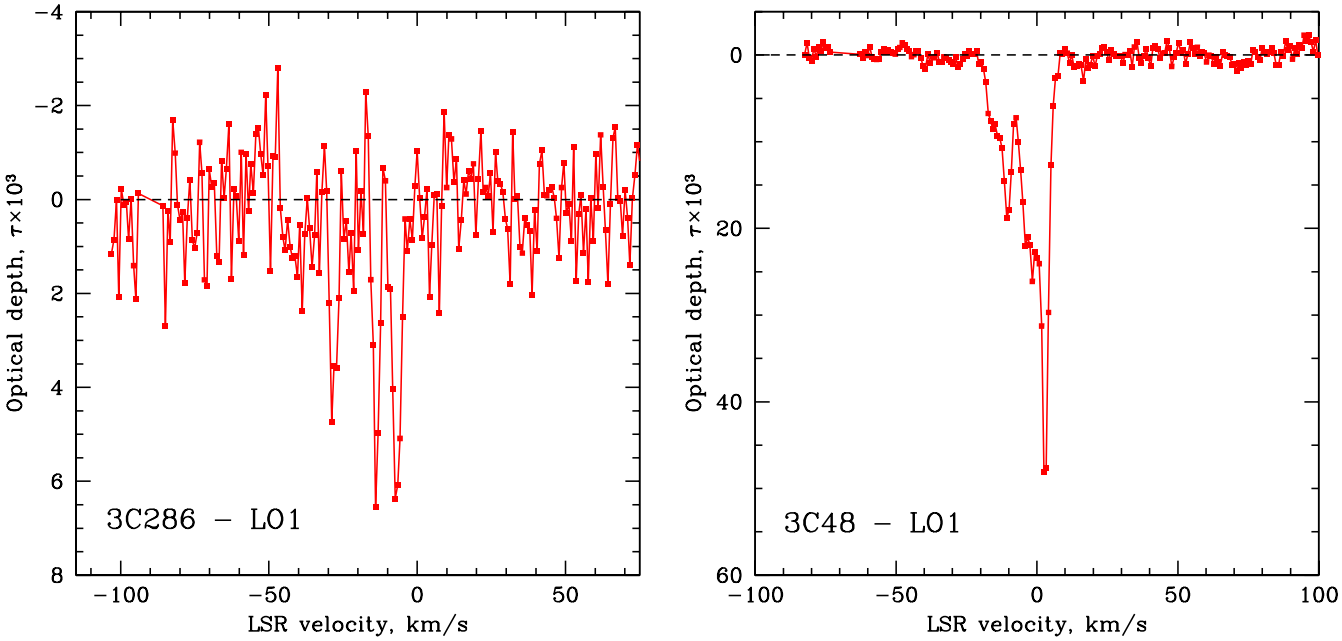

Figure 4: Final GMRT H I 21cm Stokes I spectra towards [A] 3C286 (left panel) and [B] 3C48(right panel), with switching at the first LO; in both cases, optical depth, $10^3 \times \tau$, is plotted against LSR velocity, in km/s. See text for discussion.