VLA Test Memorandum #178

Report on VLA Technical Issues Seminar #3 held on September 22, 1993

Juan M. Uson

Spectral Line Observations with the VLA

The discussion was led by J. van Gorkom and J. Uson and what follows is my summary of it.

First, J. van Gorkom reviewed the problems involved in making high spectral dynamic range observations with the VLA. The discussion dealt with the specific issues of bandpass calibration and continuum subtraction. Recommended references are Cornwell, Uson and Haddad (Astron. Astrophys. 258, 583 [1982]) and Carilli (VLA test memo #158).

We need "perfect" data for bandpass calibration which implies a need for very careful excision of any visibilities contaminated by non statistical errors such as those produced by interference (RFI). Once the data have been cleaned, the usual procedures determine an antenna-based complex gain as a function of frequency (the bandpass table).

J. Uson: It should be kept in mind that this procedure does not preserve the correct normalization if the visibility functions of the problem field and bandpass calibrator are different (unless the difference is only a scaling factor) although this might in general be a small effect.

J. van Gorkom: What we usually do is get very high sensitivity data on a strong source which provides the bandpass calibration. If the data are not perfect, the bandpass calibration might ruin the entire observation. A new method for effectively cleaning the data consists of using the (experimental) task UVMLN which implements the visibility-based subtraction of the continuum using the algorithm described in the reference by Cornwell *et al.* listed above and produces entries in a flagging table attached to the multi-source data base by computing the rms noise of the residuals (after continuum subtraction). The user decides which threshold to use. The rejection of the contaminated data can greatly improve the bandpass calibration.

M. Goss: The method has worked very well on a contaminated data base with observations of high-velocity HI against CAS-A.

J. van Gorkom: If the observations are made using narrow bands it is hard to obtain enough sensitivity on the calibrator so it might be desirable to smooth the bandpass table. The problem here is that there are features in the spectral response of the system at a level of 1% which are narrow, perhaps 10-30 kHz wide.

J. Uson: As long as the errors are statistical, there should not be a reason to smooth the bandpass table. If the SNR of the resulting spectra are low, smoothing can be done afterwards.

J. van Gorkom: In any case, these narrow features prohibit frequency-switching when high spectral dynamic range is wanted (and it is not necessary for low dynamic range problems). Is the reason for this understood?

B. Clark: Does this also happen with observations at high frequencies? Perhaps this should be tested.

P. Napier: Could this be due to the L-band polarizer? (we know the polarizer has glitches).

P. Palmer: These narrow features in the bandpass are seen in many receivers at many telescopes.

J. Uson: There are a number of circumstances that can produce resonances, so that is not surprising.

J. van Gorkom: Observations that require high spectral dynamic range are limited by a variable standing wave which produces the well-known "3 MHz" ripple (see for example Carilli's test memo). This standing wave changes with time and (possibly) with position and perhaps even with source strength. Recent data taken by Schiminowitz and van Gorkom show a smooth change of the 3 MHz ripple with time. The amplitude of the ripple (21-cm data from June 6 and 7, 1992) can vary between 1% and 4% and that particular set of data showed two antennas with particularly bad instability.

F. Owen: Is it necessary that the spectral index of the bandpass calibrator be the same as that of the problem source?

J. van Gorkom: No, if UVLIN is used to reduce the calibrated data.

J. Uson: That is equivalent to stating that a linear baseline might have to be removed from the data and UVLIN does this in a way that is almost equivalent (except for the residual power in sidelobes) to removing the same linear baseline pixel by pixel in the final image cube (which can be done using the complementary task IMLIN also described in the reference by Cornwell *et al.*).

J. van Gorkom: The variation of the bandpass ripple with time is illustrated in an experiment in which 3C273 was continuously observed (the object was to detect low-level HI emission nearby). The interference-free edited data produced a spectral dynamic range of 5×10^3, with a significant ripple at that level.
A renormalization of the bandpass in 10-minute intervals resulted in a spectral dynamic range of 10^5. HI observations of CEN-A calibrated using the observation nearest in time (AIPS option DOBAND=2) gave a significantly worse result than that produced by a calibration using a bandpass table interpolated linearly in time between observations of the bandpass calibrator (this was done outside AIPS as this option has not been implemented in AIPS). This was determined by computing the rms noise of the channel images (after removal of the continuum) which was lower in the interpolated case.

K. Dwarakanath: It is not always possible to detect a change in the bandpass response over timescales as long as 40 min.

J. van Gorkom: The best way to eliminate the bandpass ripple is to use a nearby bandpass calibrator, observe it frequently (the optimal time needs to be determined) and to interpolate the bandpass solution in time.

The calibrated data are usually reduced using UVLIN (visibility-based subtraction of the continuum). Its main advantages are that it is very convenient to use, it does the correct thing, and, in addition, it gets rid of quite a few instrumental effects such as spectral index effects as well as the variation in the response of the primary beam with frequency. It also eliminates some time-dependent effects and, for example, it reduces the amplitude of the 3 MHz ripple.

Desiderata

(1) Hardware: replace the correlator and eliminate the waveguide. This is desirable for many other reasons in addition to improving the spectral dynamic range.

R. Sramek: What kind of a new correlator would you want?

J. van Gorkom: The "cadillac" version with 1 GHz bandwidth in two polarizations that could be synthesized in more than one piece if so wanted.

(2) Software: In the short term, it is necessary to develop display tools that can allow an easy way to track

and quantify changes in the bandpass shape. It is also important to incorporate the presently experimental version of UVMLN into standard AIPS

J. Uson: The present version works but does not conform to the AIPS conventions so it would have to be modified.

J. van Gorkom: In the long term it would be helpful to have flexibility in assigning weights to antennas, also selecting baselines and especially important is to include the possibility of interpolation of the bandpass response in AIPS.

P. Diamond: The AIPS program WTMOD allows assigning weights to antennas.

J. Uson: The decision should be made depending on the data and might not be uniform over a run if the visibility function is a strong function of position in the uv-plane.

P. Lilie: The 3 MHz ripple is always present. It is different on different antennas due to the different lengths of the 20 mm waveguide. It varies smoothly with time which might be due to changes in the termination of the coupler although this is hard to determine.

B. Clark: Phase changes could come from temperature changes.

P. Lilie: Introducing attenuators reduces the amplitude of the ripple.

P. Napier: Can we afford to keep the attenuators in place when the antennas are on D-array stations? This should be tested.

P. Lilie: Some of the antennas that have been rumored to have worse changes in their bandpass response seemed worse on tests that measured the standing-wave (for example antennas 13 and 21).

Further discussion of Bandpass correction

This subject was presented by J. Uson who started by showing some examples of bandpass ripples. The frequency varied in the range 0.5 - 5 MHz depending on the observation. Most cases would have been labeled "3 MHz ripple," especially if the observations should have been taken with bandwidths smaller than 3 MHz. In some cases, the ripples can be ascribed to RFI (as suitable editing seriously modified or eliminated them), and other causes of instability such as cross-talk, change in system temperature,...

Observations with high spectral dynamic range require high precision bandpass correction. It is not always sufficient to interpolate in time. For example, a test run at L-band observed several sequences containing 1226+023 (this is 3C273 used as bandpass calibrator for all sources in this run), 3C286, 3C273 and 0826-373. The observation was made on May 27, 1993 using two pairs of 6.25 MHz bandwidth spectra (15 channels each spectrum, after on-line hanning-smoothing). The sources were assumed featureless in frequency and the resulting spectra had no clear fixed-frequency ripple although some frequency could be found for each one of the individual spectra (in the range indicated above). No proper interpolation was done as the data were reduced using AIPS. In general, nearest-neighbor bandpass calibration was used (AIPS DOBAND=2) with some exceptions discussed below.
The peak departure from a featureless spectrum was at the relative level of 0.1% for both 3C286 and 0826-373. For 3C273 it was 0.06% whereas for 1226+023 it was 3×10^{-5} (this is consistent with the thermal noise for this observation). This shows that, in this case, some interpolation in time (which was of course fairly successful for 3C273) was not sufficient when the observed source was at a different position. The much better correction achieved on 1226+023 (bandpass table entries computed as "scan" averages, then used to correct all entries for that scan) is consistent with J. van Gorkom's result using proper interpolation in time. Presumably, such interpolation on 3C273 (using the 1226+023 data) would have improved its bandpass correction to a level close to that achieved for 1226+023.

Furthermore, for 3C286 and 0826-373 there was no significant difference between spectra obtained with options DOBAND=2 and DOBAND=1 (one single entry in the bandpass table with the average of all the data on 1226+023). This indicates that, for this database, the variation in the bandpass with position in the sky, or rather position with respect to the ground, was much larger than the change due to the passage of time. The L-band system suffers from a highly variable contribution to the system temperature due to ground pick-up and it seems likely that this affects the bandpass response at this level by changing the operating point of the receiver systems (the ALC gain corrections).

Hopefully, interpolation in time will be adequate under most optimal or good circumstances. This might require nighttime observing for the most demanding projects. The required software tools must be developed to do proper interpolation. This requires sensible weights attached to each solution. These could be computed from theory (T_{SYS}, t_{INT}, signal strength in each baseline,...), or from the data themselves (scatter).

In addition, it is possible that using the strongest possible source as a bandpass calibrator might not be the best strategy to reach the highest dynamic range. At some bands, such strategy might more than double the system temperature. It is not known at this time whether the change induced in the ALC gains might allow the channel to channel stability to be preserved at the 10^{-5} level, although that seems unlikely. It would seem that choosing a more moderate bandpass calibrator and appropriately increasing the time that it is observed might lead to a better bandpass correction. This should probably be tested.

Bandpass interpolation: One workable approach

It would seem that the following approach might produce useful and correct results:

(1) Compute baseline-based bandpass response. This could be done by averaging the data over a suitable time interval, T, perhaps the duration of a scan. The rms noise for each channel of each entry should be derived from the scatter of the data, perhaps in a robust way. These $\sigma_{CHANNEL}$ can be used to determine one σ for the bandpass. This could perhaps be the median value to avoid biasing due to RFI.

(2) Compute antenna-based bandpass response. This would then be done by weighted inversion of the baseline-based bandpasses. The associated errors have to be computed by correct propagation of the input uncertainties through the weighted inversion process. It should be possible to exclude from the inversion those baselines susceptible to cross-talk using a uv-range threshold or a neighborhood criterion. It should also be possible to reject some of the baselines with RFI contamination, high σ, or other reasons to be determined by the user.

(3) Interpolate an n-th degree polynomial in time to each channel for each set of antenna-based bandpasses. This procedure might not preserve the normalization so the interpolated bandpasses should perhaps be normalized. However, the antenna-based bandpasses are not applied to the data during calibration but the corresponding products for each antenna pair are. Consistency with the principle that led to the initial normalization of the antenna based spectra would suggest that the normalization should be restored after these products are computed. Although this is expected to be a small effect under most circumstances, the method should be tested for systematic errors once the necessary software is in place.