

EVLA Memo 148

WIDAR anti-aliasing and spectral dynamic range test

R.J. Sault

24 August 2010

Introduction

WIDAR, the correlator being commissioned as part of the EVLA upgrade, allows observing with multiple simultaneous subbands. Intrinsically WIDAR produces many channels within each subband. WIDAR aims to be a system with high spectral purity. In particular, WIDAR aims to produce low non-ideal in-band response of a spectral feature, low aliasing of a spectral feature from one subband to another, and good robustness to a variety of systematic errors and unwanted signals (e.g. sampler offsets and RFI). This memo aims to characterize some aspects of WIDAR's in-band dynamic range and its suppression of aliases from one subband to another.

In-band spectral response

At a low level, WIDAR can be thought of as an XF correlator. The most basic correlation operation is to measure 128 lags for a given polarization pair for a subband for a particular baseline. This results in 64 channels. The lags are uniformly weighted, and so the spectral point-spread response is a sinc function. For a signal much narrower than a WIDAR channel and placed at the center of a channel, the WIDAR response should be a delta function without sidelobes. Sidelobes will become apparent if the signals width becomes significant or if the signal is not aligned perfectly with channel center.

Anti-aliasing performance

When multiple subbands are observed simultaneously, the rejection of aliased responses of a spectral feature from one subband to another is important. This memo focuses on subbands that are strictly adjacent to each other. That is, the last channel of a subband and the first channel of the next subband are one channel width apart. In WIDAR, potential aliased responses are first substantially reduced by the digital filters that are used to select out each subband from the digitized data stream. In the finite-bit arithmetic used in these digital filters, and away from the subband edges, the attenuation in the bandstop parts of the subband filters is better than 60 dB (Figure 1).

A novel characteristic of the design of WIDAR is the so-called “fringe washing” technique which provides additional anti-aliasing between adjacent subbands. This aims to allow adjacent subbands to be “seamlessly stitched” together. Fringe washing also helps to improve the potential dynamic range of spectra because a number of other non-ideal responses also tend not to correlate between antennas.

The fringe washing technique introduces an antenna-specific frequency offset in the LO at the front end. This frequency offset, f_{shift} , is removed in the late stages of the correlation process. Aliasing causes a spectral response to be mirrored at the subband boundary. Because the f_{shift} frequency offset is removed after this mirroring happens and because the frequency offsets differ between antennas, the aliased responses of a signal will be at slightly different frequencies in different antennas. Consequently, provided the frequency difference of




Figure 1: The WIDAR digital filter response of a for a subband bandpass filter. This correctly represents the response after accounting for finite precision arithmetic.

the offsets between two antennas is significant enough, the aliased response will decorrelate away. This approach will work up to the band edges of adjacent subbands.

This fringe washing technique is closely related to the decorrelating effect that the natural astronomical fringe rate has on some undesired signals. In the tests of this memo, the astronomical fringe rate was made to be 0 by electronically tracking the north celestial pole. Consequently the fringe washing effect should be purely caused by the introduced f_{shift} frequencies rather than the natural source fringe rate.

Fringe washing also has some commonality with the traditional approach of Walsh switching. Walsh switching can be thought of as offsetting the LO frequency, with this offset then removed in the late stages of correlation.

It is worth noting that fringe washing does not suppress sidelobes (as distinct from aliases) from out-of-band signals. The digital filter response will attenuate the strength of an out-of-band signal. However the sidelobes of that attenuated signal will not be attenuated further if they fall in-band.

The test

This memo reports on some tests to measure both

- the spectral dynamic range achieved in-band when there is a strong spectral line, and
- the degree to which an aliased response is rejected in an adjacent subband.

The tests reported here were the fourth in a series of tests. They were performed on 10 May 2010. A CW tone of 0 dBm was transmitted from the Control Building. The tone is narrow (\sim Hertz, compared with the channel width used of 7.8 kHz), has been centered on a WIDAR channel and has locked to the site maser. For these tests the CW tone has been assumed to be an ideal. Consequently this signal could be fixed within a correlator channel to very high precision. The antennas were in the VLA D configuration, and were physically




Figure 2: The transmitter and antenna locations during the test.

stowed during the test observations. Figure 2 shows the antenna location and the approximate location of the transmitter.

The strength of the transmission did not add significantly to system temperature: the receiver electronics should be well within their linear regime. The observing control system was set to (electronically) observe within a few arcminutes of the north celestial pole of date. Two adjacent subbands were observed, with each subband being 500 kHz wide and correlated to produce 64 channels.

Earlier tests suggested some artifacts were present when the f_{shift} on different antennas were related by small odd integers. That is for a baseline involving antennas i and j , artifacts may be present when there are small odd integers m and n such that

$$m f_{\text{shift},i} = n f_{\text{shift},j}.$$

The EVLA observing system requires f_{shift} on all antennas to be some multiple of a base frequency, and so there will always be some m and n for which the above equation holds. The tests described here used two approaches to setting f_{shift} - the so-called “standard” and “non-standard” approaches:

- The standard approach sets f_{shift} equal to a base frequency of 100 Hz times the antenna number (e.g. $f_{\text{shift},1} = 100$ Hz, $f_{\text{shift},2} = 200$ Hz etc). This is the current normal approach for setting f_{shift} in the EVLA system. A comparatively small base frequency (100 Hz) has been used to ensure that the largest offset (2.8 kHz) is still a fraction of a channel spacing (7.8 kHz for these tests).
- The non-standard approach sets f_{shift} on the west arm antennas only so that the smallest values of m and n satisfying the equation above were comparatively large. Antennas on the north and east arm were set as per the standard approach. The largest f_{shift} was 3.8 kHz. Non-standard f_{shift} values were restricted to west arm antennas. These antennas were the closest to the control building and so received the strongest signal from the transmitter: they were the antennas of most interest in the test. The following table gives the f_{shift} value on these antennas.




Figure 3: Representation of the location of CW tone within the two observing subbands. The frequency settings of the subbands were varied to place the fixed frequency CW tone at different locations within the passbands.

Antenna	f_{shift} (Hz)
1	100
5	2500
17	3400
15	1500
24	600
19	3800
26	2600
16	2800

In this non-standard setting, given the antennas available for the tests, two pairs of antennas have the same f_{shift} values: antenna pairs 5 and 25 (2500 Hz) as well as 16 and 28 (2800 Hz). The response on these baselines provides an additional probe of the system.

The test investigated placing the CW tone at different points within the two subbands. One approach to this would be to have the observing system fixed and to change the frequency of the CW tone. Alternatively the frequency of the CW tone could be fixed, and the observing system changed. The latter was logistically more convenient and was the approach used. The frequency of the CW tone was set to 8350 MHz. With channels numbered from 1 to 64 (channel 33 being the center channel of the passband) the observing frequencies were alternately configured to four different settings so that the CW tone was centered on

- Setting A: subband 1 channel 33,
- Setting B: subband 1 channel 61,
- Setting C: subband 2 channel 1 or
- Setting D: subband 2 channel 5.

The location of the CW tone in the observing subbands is shown in Figure 3.

The results

Antenna 3 showed unique and poor behavior. This was later found to be caused by a hardware instability in its local oscillator system. It has been excluded. Antenna 2 was affected by low level spikes spaced at 125 kHz: it is not clear whether this was RFI or a correlator-based problem. Antenna 2 was also excluded from the




Figure 4: An example of a spectrum showing artifacts around the base of the tone. The apparent tone strength was 2600 and 1400 Jy in RR and LL correlations.

analysis. Antennas 6 and 10 were not included in the observation. Hence there are 24 good antennas. Various system problems meant that typically 20% to 25%, but sometimes 70% of data for a given baseline was flagged. The main problem were timing issues resulting in output data failing to be saved (the “identically zero data” problem). The observation cycled through the four different frequency settings over 12 minutes. This was first done with the “standard f_{shift} approach, and then the non-standard approach. The data were flagged and phase calibrated using the observation of the tone centered in the first subband. A nominal amplitude calibration was applied. There was no bandpass calibration applied.

The deduced antenna phases were relatively stable, although the amplitude was less so. This is expected. A subsequently “tone extractor” run with WIDAR was performed to measure stability of the tone. It found that the phase of the tone, relative to the WIDAR model of the expected tone, fluctuated by of order 5 to 10° over a few seconds. This shows that the CW tone and the observing system are locked to good precision. The fluctuation in phase may be instrumental or may be a multipathing effect.

Spectral dynamic range results

Low level artifacts were noted around the base of the tone signal in the spectra of a number of baselines. An example is shown in Figure 4. This spectrum is for frequency setting A and baseline 4-20 (The f_{shift} frequencies were 400 and 2000 Hz, giving an f_{shift} ratio of 5 – see below). To quantify the artifacts, we have used the rms value of the data in the three frequency channels above and below of the tone as a measure of spectral impurity. This rms is normalized by the tone amplitude for that baseline. This fractional spectral impurity is the reciprocal of a spectral dynamic range measure. In computing spectral impurity, the data for each baseline and polarization were (vector) averaged in time: the artifacts were noted to be phase coherent with the tone, and so after phase calibrating the data on the tone averaging in this way reduces the thermal noise level but not the artifacts. This measure of spectral impurity was noted to be a significant function of the ratio of the f_{shift} frequency of the antennas. There was no significant difference between polarizations or between frequency settings A, B and D. Nor was there significant difference between the “standard” and “non-standard” ways of setting the f_{shift} frequency - this shows that the effect is tied to the f_{shift} setting rather than antennas in some unusual way. Figure 5 gives a plot of fractional spectral impurity as a function of f_{shift} ratio. The points plotted are all above




Figure 5: A plot of fractional spectral impurity against f_{shift} ratio for frequency configurations A, B and D.

the thermal noise limit. This figure includes frequency settings A, B and D, standard and non-standard f_{shift} data as well as RR and LL polarizations. Baselines for which f_{shift} was the same for both antennas are not shown in this figure. It is clear from this figure that spectral purity is significantly poorer where the f_{shift} ratio is exactly an odd integer. This is particularly so when the f_{shift} ratio is 3, 5 or 7 but also for larger odd integers: 9, 11, 13 and 17. As is apparent from Figure 5, artifacts at the 1-2% level are present on some baselines. As noted above, these artifacts are phase coherent with the tone and so do not average down with a long integration.

The artifacts in frequency settings A, B and D were never noted to cross the subband boundary. Indeed sometimes they were noted to stop abruptly at the subband boundary. This implies that they are correlator-based errors that are generated within the subband processing. They do not represent a linear response to the tone.

Frequency setting C behaves quite differently to A, B and D [recall that frequency setting C is when the tone is placed on a channel adjacent to a subband boundary]. For setting C, the artifacts were noted to cross the subband boundary and the size of the artifacts are about a factor of 3 larger - up to 6%. Additionally dependence on the f_{shift} ratio, if any, is much less clear. The behavior in this case is different from the other frequency settings. Figure 6 shows the fractional spectral impurity for setting C as a function of f_{shift} ratio.

Alias suppression results

The suppression of the alias should be the product of the digital filter response and the fringe washing. Specifically for digital filter response r , integration time T and f_{shift} frequencies of the two antennas of ν_1 and ν_2 , the alias attenuation should be better than

$$\frac{r}{2\pi|\nu_1 - \nu_2|T}.$$

The expected suppression offered by the digital filter response is given by Figure 1: it is approximately 60, 20, 5 and 20 dB for configurations A through D respectively. The fringe washing decorrelation is a function of integration time T : because the rms thermal noise level decreases as \sqrt{T} and fringe washing as T , the largest alias-to-noise ratio will be for the smallest T . The best approach to detecting aliases is to incoherently average the visibility data (i.e. average the visibility amplitudes) at the correlator dump rate. For the tests described here, the correlator dump time was set to 1 s. For differential f_{shift} frequencies of a hundred to a few thousand Hertz, the resultant fringe washing suppression in a 1-s correlator dump should exceed approximately 30 to 40 dB.




Figure 6: A plot of fractional spectral impurity against f_{shift} ratio for frequency setting C.

For frequency setting A, the tests did not detect an alias. The data in the tone-free subband were consistent with the expected thermal noise level. The noise level integrated down as expected. The best the test could conclude is that alias rejection was better than 44 dB.

For frequency configurations B and D, an alias response was detected on many baselines. The expected alias suppression is 50 to 60 dB (a digital filter contribution of 20 dB and fringe washing suppression of 30 to 40 dB). The upper panel of Figure 7 shows an example of the detected alias. The alias suppression that was achieved varied from about 27 to 40 dB, with suppression of 35 dB being typical. This suppression is about 20 dB poorer than might be expected.

The following characteristics of the detected alias can be noted:

- The alias suppression on a particular baseline *should* depend on the difference in f_{shift} of the antennas. We did *not* find such a relationship.
- The alias suppression should improve in proportion to integration time. Averaging data up to 16 s does show a roughly proportional decrease in the alias response. The alias response becomes difficult to detect beyond integration times of 16 s.
- No significant difference in alias suppression was noted between the standard and non-standard way of setting the f_{shift} frequencies. This is as expected.

The level of artifacts in frequency setting C prevented the alias being detected (or are the artifacts the alias?).

The following table summarizes the expected and achieved alias suppression in the four frequency settings.

Frequency Setting	Expected Suppression (dB)	Achieved Suppression (dB)
Setting A	90-100	>44
Setting B	50-60	~ 35
Setting C	35-45	?
Setting D	50-60	~ 35

An additional check on the performance of alias suppression is available on those baselines where there is not fringe washing. Without fringe washing, alias suppression is set purely by the digital filter response. For frequency




Figure 7: Upper panel: Baseline 7-17 showing an aliased response. The thermal noise floor is near -50 dB. Artifacts approximately 15 dB above this noise floor can be seen around the CW tone. The aliased response is about 35 dB below the tone. Lower panel: Baseline 5-25. f_{shift} for antennas 5 and 25 were set to the same frequencies, so there is no fringe washing in reducing the alias response.

settings B and D, alias suppression of about 20 dB is expected. The lower panel of Figure 7 shows the tone and alias as detected on baseline 5-25 for the non-standard f_{shift} observation. For this observation, f_{shift} for both these antennas was set to 2500 Hz. The figure shows that the alias is suppressed by the expected 20 dB.

Conclusions

This memo notes some test observations where some characteristics of the spectral response of WIDAR that are not currently understood:

- Spectral artifacts are present which can be 1-2% of the spectral line peak. These artifacts are related in some way to the ratio of the f_{shift} frequencies of the antennas: artifacts are most significant when the ratio of the f_{shift} frequencies are small odd integers.
- Significant artifacts were seen when there was signal in the channel at the edge of a subband. Understanding this effect is important to allow “seamless stitching” between subbands.
- Although the fringe washing technique was shown to help suppression the alias of a spectral line into an adjacent subband, the suppression was ~ 20 dB less effective than expected.

Tests and analysis are on-going to understand these results. It is clear that at least some of these effects are related to finite precision characteristics within the correlator. Having said this, the results should be taken with a significant caveat: the effects have not been shown to affect standard observing. Certainly an observation near the north celestial pole is likely to accentuate a number of errors: some error effects will not be decorrelated by the astronomical fringe rate or delay rate of the source. Indeed this is why the test observation was framed in this way. Additionally, these test observations used a tone which was coherent with the telescope frequency standard: the tone was locked to the site maser. Consequently it is possible that there are unforeseen correlations between the signal and system that will not be relevant to observations of an astronomical source.