EVLA Memo 160
More WIDAR spectral dynamic range tests

R.J. Sault

May 2, 2012

Introduction

This is a continuation of investigation of the spectral dynamic range achievable with the WIDAR correlator. Previous investigations (EVLA Memo 148) were limited by the so-called "splatter" phenomena, which is now understood and avoidable (EVLA Memo 157). This investigation is also stimulated by a desire to evaluate the "correlator chip beat artifact" seen in simulations (NRC-EVLA Memo 32) and to measure the isolation between WIDAR subbands.

The test

The tests performed were similar to those reported in EVLA Memo 147. Specifically observations were done with two adjacent subbands, each of 500 kHz, each configured to produce 64 channels. Only the RR and LL products were analyzed. The EVLA was in the C array. A test tone was transmitted from the Control Building (see Figure 1 for the physical layout). The test tone, at 8350 MHz, was locked to the site maser. The system was set to phase track a source less than an arcminute from the north celestial pole of date. Fringe and delay rates for this experiment were negligibly small.

These tests examined a modest range of parameter space. Two settings of $f_{\rm shift}$ frequencies were used for the nominally 28 antennas. The first set - the "slow $f_{\rm shift}$" - used frequencies of 10 Hz multiplied by 28 primes from 31 to 313 (ie 310 to 3130 Hz). The second set - the "fast $f_{\rm shift}$" - used 10 Hz multiplied by 28 primes from 283 to 3067 (ie 2830 to 30670 Hz). Antenna 1 used the highest $f_{\rm shift}$ frequency and antenna 28 the lowest.

The center of the observing bands were adjusted to place the fixed tone at different parts of the two subbands - see Figure 2. With channels numbered from 1 to 64 (channel 33 being the center channel of the subband) the observing frequencies were alternately configured to four different settings so that the CW tone was centered on:

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

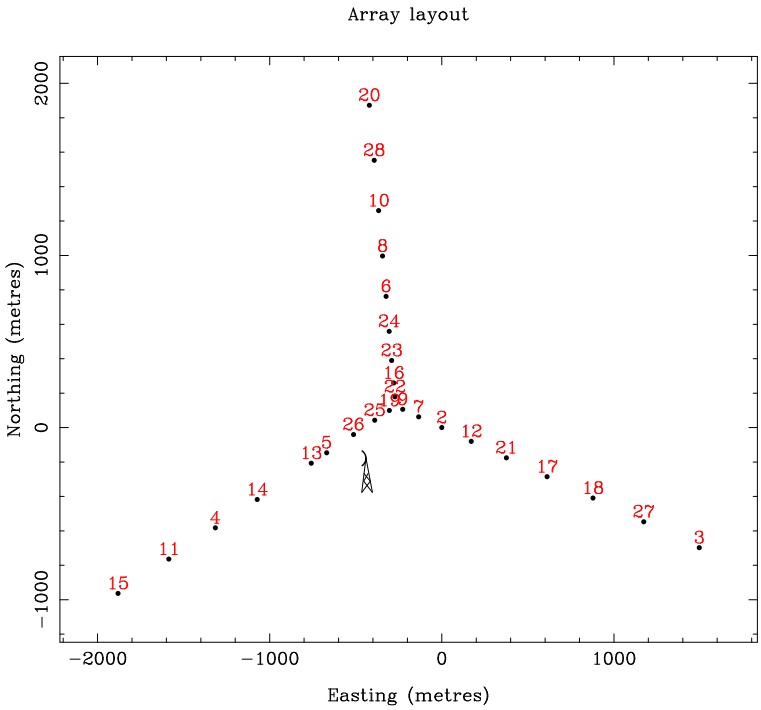

Figure 1: The transmitter and antenna locations during the test.

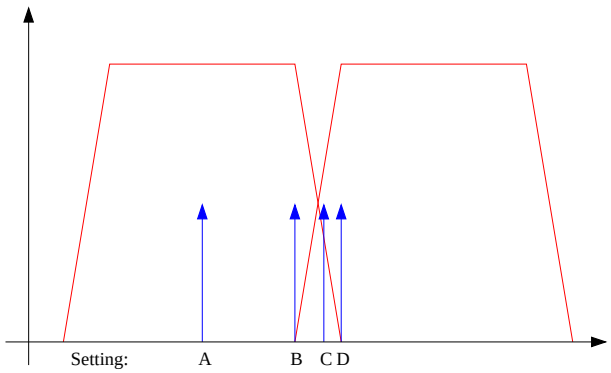

Figure 2: Representation of the location of tone within the two observing subbands. The frequency settings of the subbands were varied to place the fixed frequency tone at different locations within the passbands.

The location of the tone in the observing subbands is shown in Figure 2. During the observations, the re-quantizer gains were determined when the system was in frequency configuration A only. This was also done at each different correlator chip dump time setting (see below)[1]. Unfortunately in frequency configurations C and D, when the tone moves into the second subband, the power level in the second subband would have increased appreciably. As the re-quantizer gains were not re-determined, we believe significant clipping happened during re-quantization on some antennas. That is, the digital data saturates. This overloading was most apparent for antennas 5 and 25, i.e. the antennas where the largest power change would be expected. Consequently the data for some antennas were rendered bad or suspect for subband 2 in configurations C and D. These data have been discarded - they are not presented here.

NRC-EVLA memo 32 notes that there is a potential for a "correlator chip beat artifact" when the correlator chip dump period and $f_{\rm shift}$ frequencies are related. Each correlator chip dump period ends with 1 μs during which correlation is not performed and the transfer of the correlator chip accumulators is initiated. To characterize this effect, three different correlator chip dump times were used - $\tau_{\rm dump} = 400$, 401 and 390.625 μs.

The data were accumulated to 1 s (or 1.0025 s for $\tau_{\rm dump} = 401$ μs) before being written to disk.

Two power output levels of the transmitter where used: -13 dBm and 0 dBm. On baseline 5-26 (the antennas closest to the transmitter), this power level appeared as a $\sim 10^5$ and $\sim 10^4$ Jy line respectively. The tests were performed on 9 and 13 March 2012. Some follow-up observations, to clarify specific effects, were performed on 9 and 10 April.

The different combinations of parameters explored are summarised below.

Variable parameter settings	
$f_{\rm shift}$ setting	310 - 3130 ("slow") and 2830 - 30670 ("fast") Hz
Frequency configuration	A, B, C and D
Correlator chip dump time, $\tau_{\rm dump}$	400, 401 and 390.625 μs
Transmitter power level	-13 dBm and 0 dBm

The results

The results show that there was a significant uncorrected phase rate in the observations. The apparent antenna phases varied much more rapidly than would be expected. Additionally it appears that there was significant, time-varying, phase slope across the observing subbands. This phase slope was inconsistent with multi-path variation: the phase slope across the subbands corresponded to a timing error of $\sim 10^{-5}$ s or a pathlength of about 3 km. We cannot offer a firm explanation for these phase errors. Given that this observation was run in a non-standard mode, we suspect we failed to appropriately drive the observing system. We believe the results in the follow sections appropriately account for this, and that the results are not compromised.

The resultant data showed a number of low-level artifacts which have varying origins in WIDAR. Below we consider each of these.

[1]Re-quantization is a re-normalization of the antenna data stream at the end of the "station" processing and before correlation. The gains in this normalization are determined on a subband basis.

Adjacent subband purity

With frequency configuration A, where the tone is at the center of the first subband, we have never observed any artifacts in the adjacent subband. The second subband was always consistent with noise at the expected nominal level. The rejection of the tone (and its potential artifacts) in the adjacent subband is observed to be better than 45 dB. This rejection of the tone is set by the subband filter shapes. When finite precision arithmetic is considered, the rejection of the subband filters is expected to be better than 60 dB (e.g. see EVLA Memo 148) away from the band edges. WIDAR shows excellent subband separation.

Given that no artifacts are see in the adjacent subband for frequency configuration A, it appears likely that the artifacts considered below originate late in the signal path (e.g. after the samplers and the first stage digital filters).

Spectral sidelobes cross subband boundaries

To generate a spectrum, WIDAR measures a finite set of lags and then Fourier transforms them with equal weight. Thus the natural spectral impulse-response function of WIDAR is a sinc function (at least approximately so). In most of the tests described here, the sinc function is not apparent: the tone was placed at the center of a channel, and so the nulls of the sinc function fall at all channels other than the channel containing the tone. However, inadvertently, some of the tests placed the tone somewhat offset from a channel center, and so the sinc function spectral sidelobes became apparent.

As described in many NRC-EVLA memos, an important architectural feature of WIDAR is the phase washing technique. This helps to beat down a number of aliases and artifacts. The phase washing technique, however, has no effect on spectral sidelobe responses. The isolation between subbands with respect to sidelobes is set by the subband filter responses alone. Note, because the attenuation provided by the subband filters is $\sim$60 dB except near the subband edges, the sidelobe response in adjacent subbands is generally negligible. The exception is for emission near the subband edges. Figure, 3 gives an example of an instant where sidelobes did cross subband boundaries, for baseline 5-26. Both an expected (model in red) and an observed spectrum (in green) are given. Frequency setting B was used (i.e. the tone was placed 3 channels away from the subband boundary). The expected spectrum is based on the measured tone strength of 13,100 Jy, the known frequency offset from the center of channel 61, as well as the known WIDAR subband filter response function. In interpreting the figure, channels 1-64 are from subband 1 whereas channels 65-128 are from subband 2. On the left half of the plot, the measurements almost perfectly overlie the model (to get such agreement, the full spectral PSF of WIDAR needed to be used - see Appendix A and following section). The right side shows modest agreement: the disparity may be because of imperfect understanding of the filter response.

WIDAR spectral point-spread function

Appendix A gives the spectral point-spread function (PSF) for a complex-valued, N-lag correlator such as WIDAR. The real part of the PSF is a sinc-like function. But somewhat surprisingly, the PSF is complex valued when N is even. For channel spacing $\Delta\nu$, the imaginary part of the spectral PSF is

$$\Im[h(\nu)] = \frac{1}{N}\sin(\pi\frac{\nu}{\Delta\nu}).$$

When the tone is placed at the center of a channel, the imaginary part is identically zero. However if the tone is offset, a imaginary response is apparent. This is readily seen in the observed data in those instances when the

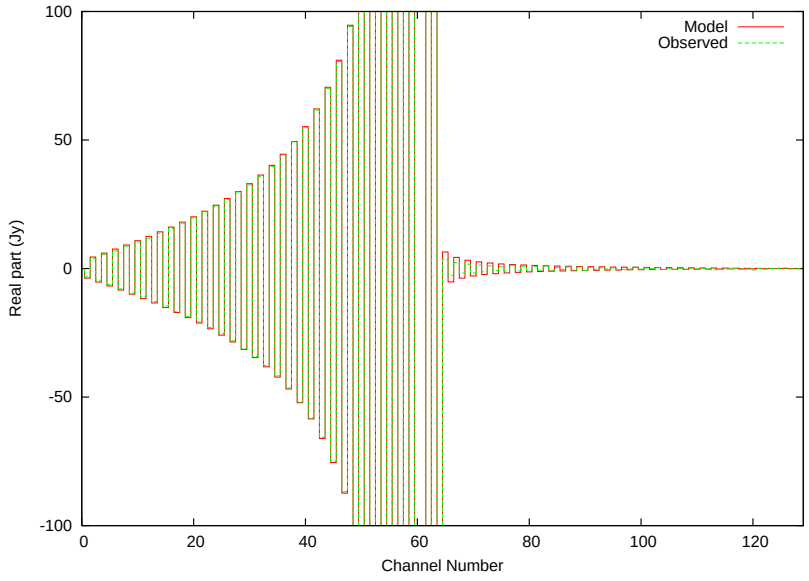

Figure 3: Example of sidelobes cross subband boundaries, for baseline 5-26. The observed data and the expected ('model') response are given. Channels 1-64 and 65-128 are in the first and second subbands respectively.

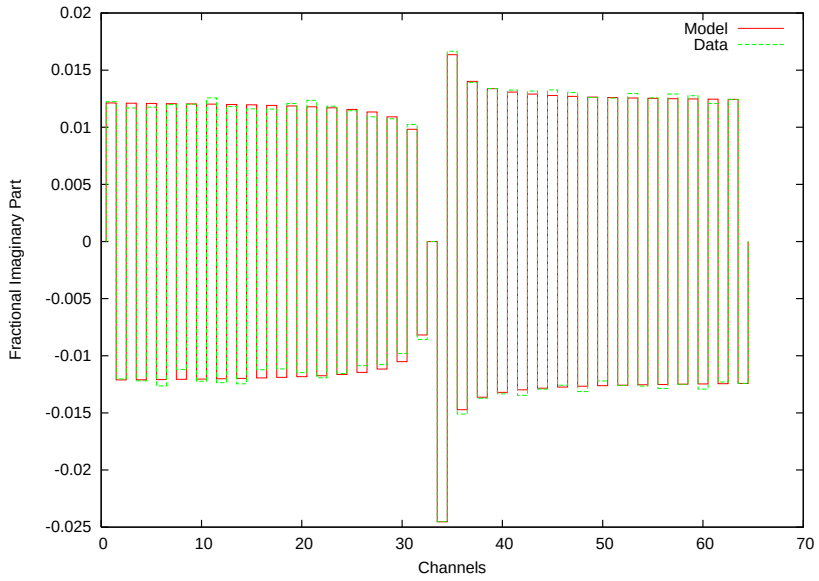

Figure 4: Imaginary part of the WIDAR spectral point-spread function for a tone offset half way between two channels. The tone is centered at channel 33.5. Self-calibration distorts the response pattern by enforcing channel 33 to have a phase of zero. The observed data and the expected result are given.

tone was not channel centered: it is of order a 1% effect. When the tone is not channel-centered, the imaginary part will oscillate between consecutive channels between a positive and negative value.

Figure 4 gives the imaginary part of observed data when the tone was set to be half way between two channels. Figure 4 also gives the expected result. An added consideration that is apparent in this figure is that our data reduction process modifies the natural PSF. The data reduction performed self calibration using the peak channel of the line to determine antenna gains. Self calibration will subsume the complex-valued nature of the PSF into the antenna gains. This leads to a distortion of the simple oscillation of the imaginary part near the channel used for self-calibration. This effect is easily modeled: the resultant PSF with self calibration will be $h(\nu)/h(\delta\nu)$, where $\delta\nu$ is the offset frequency between the tone and the channel used for self calibration (i.e. $\delta\nu$ will always correspond to half a channel or less in our self-calibration approach). In Fig. 4, we have self calibrated to give a source flux density of 1 unit. In this case, the expected imaginary part will have an amplitude of

$$\frac{\pi}{N}\frac{\delta\nu}{\Delta\nu}$$

So for a tone offset of half a channel from channel center and number of lags $N = 128$, the expected imaginary response is 1.22% of the real response.

It is quite straightforward to eliminate this imaginary response from the WIDAR spectral PSF: it simply amounts to excluding the first lag value ("first" in the sense of Appendix A). This could be done in the correlator backend software just before the Fourier transform step. Alternatively it could be done by Hanning smoothing in the data reduction software (a simple filtering with channel weights [0.25,0.5,0.25]). Hanning smoothing will also cause the first lag to be zeroed.

Although readily measurable, are there practical consequences if this characteristic is not eliminated? These appear to be modest at best. The imaginary response will go hand-in-hand with the sinc-like response of the real part. If the sinc-like response is apparent, Hanning smoothing will most likely be required anyway. Because the response is the same on all baselines, the imaginary component will not introduce a closure error. It will introduce an imaging artifact. However this is appreciably attenuated in the real image.

As an example, an image was formed of a simulation of the effect: a VLA snapshot was simulated with the observing parameters of our test observation, but with a 1 Jy one-channel wide spectral line, which was also a point source The spectral line was placed midway between channels 33 and 34. Channel 5 was imaged and cleaned. Channel 5 is in the wings of the sinc-like response function. The expected and also the recovered flux density was 14.7 mJy. Although the simulation was noiseless, the RMS residual was 270 μJy/beam. This RMS residual is a result of the 12.2 mJy imaginary part that each visibility possessed, and which was "scattered" into the real image.

Although the effect of the characteristic appears to be modest, the method to remove it also appears simple and not prone to side effects. An equally good question to ask is whether there is a strong reason not to eliminate the characteristic.

Degraded SEFD: Correlator non-linearities and harmonics

The finite precision used within WIDAR (primarily the 4-bit representation of the data in the correlator cell, the 4-bit representation of the phase model and the 3-level phase rotator) all contribute to scattering power to different spectral frequencies. To first order, this scattered power does not correlate. When the signal is both strong (i.e. where self noise is a significant consideration) and spectrally rich, then the scattered power will also be spectrally rich.

During the analysis of the these tests, it was noted that the effective system temperature on some antennas at some channels was substantially more than normal. The effect was most pronounced for antennas nearest the transmitter. This increase was indeed like a heightened system temperature. For example, it affected all baselines including a particular antenna (it is not proportional to the correlated signal on a baseline, but rather it is an antenna property). It also integrated down over time, with the rms value goes down as the square root of the number of samples.

This behavior is expected: it can be thought of as self noise being scattered spectrally by WIDAR non-linearities. Higher $f_{\rm shift}$ frequencies are likely to lead to scattering the power over a larger frequency range. With a single test tone, the situation is complicated enough. When there are multiple, strong, tones, intermodulation products between these tones makes the situation a good deal more complicated. A detailed analysis has not been attempted here. See NRC-EVLA memos 1 and 9 for some analysis.

Figure 5 gives an approximate effective system equivalent flux density (SEFD) for antenna 5 for frequency configuration A. Note antenna 5 is the antenna closest to the transmitter. Three panels are given, corresponding to high power (line $\sim 10^5$ Jy), fast $f_{\rm shift}$; high power, slow $f_{\rm shift}$; and low power (line $\sim 10^4$ Jy), slow $f_{\rm shift}$. Because the system generates odd harmonics, we might expect increases in SEFD at offsets of $2n\nu_1$ from the main line. This corresponds to offsets of $6.8n$, $0.69n$ and $0.69n$ channels in Fig. 5a, 5b and 5c respectively. There is modest consistency with this, with SEFD increases corresponding to $n = \pm 2, \pm 4$ apparent (as is the alias of $n = \pm 10$ in Fig. 5a). The nominal SEFD of an EVLA antenna is about 300 Jy.

Although the degradation in SEFD can be appreciable, this is probably of little practical importance. Poor SEFD near a very strong line is usually not significant to the science. One exception is Zeeman experiments, where the noise in the V spectrum can be high and channel specific, but the line strength can be modest or low.

Scattering of power in this way is a basic characteristic of WIDAR. For practical observing, assuming appreciable changing fringe rotation, the power will be scattered more uniformly than in these test observations. This will make the average channel-to-channel variation in SEFD less pronounced than presented here.

Correlator chip dump beat artifact - In-band effect

The "correlator chip dump beat artifact" was noted in simulations described in NRC-EVLA Memo 32. It is caused by a beating between the correlator chip dump period, $\tau_{\rm dump}$, and the length of a "phase washing" cycle, $\tau_{\rm pw} = 1/(\nu_2 - \nu_1)$. At the end of each dump period, correlation is not performed (correlation is "blanked") for 1μs. When this blanking beats with the phase washing cycle, it is apparent that artifacts could build rather than being washed out. The dump period is configurable, and typically $\sim 400\mu$s. For a blanking time of 1μs and dump period of 400μs, the size of the error will not exceed 1 part in 400 – 0.25%.

Appendix B analyzes the effect in some detail. As the appendix notes, in general the amplitude of the potential effect will be significantly smaller than 0.25%. The appendix gives a formula to calculate the relative amplitude of the error - ϵ. In addition to ϵ being small, the effect generally produces an artifact at negative frequencies[2]. The artifact is a mirror of the true visibility spectrum at negative frequencies, but shifted by the sum of the $f_{\rm shift}$ frequencies, $\nu_1 + \nu_2$. It is only when this (cyclic) shifting moves the artifact from negative to positive frequencies that the artifact will become noticeable to an observer. Thus the effect will only be significant when there is emission near the edges of the subband. This may be an issue for continuum observations where emission fills the subband.

[2]Recall that WIDAR produces a complex-valued lag spectrum. Fourier transforming this produces a visibility spectrum with independent positive and negative frequencies. Ideally the negative frequencies will be noise only. They are usually discarded before the observation is written to disk.

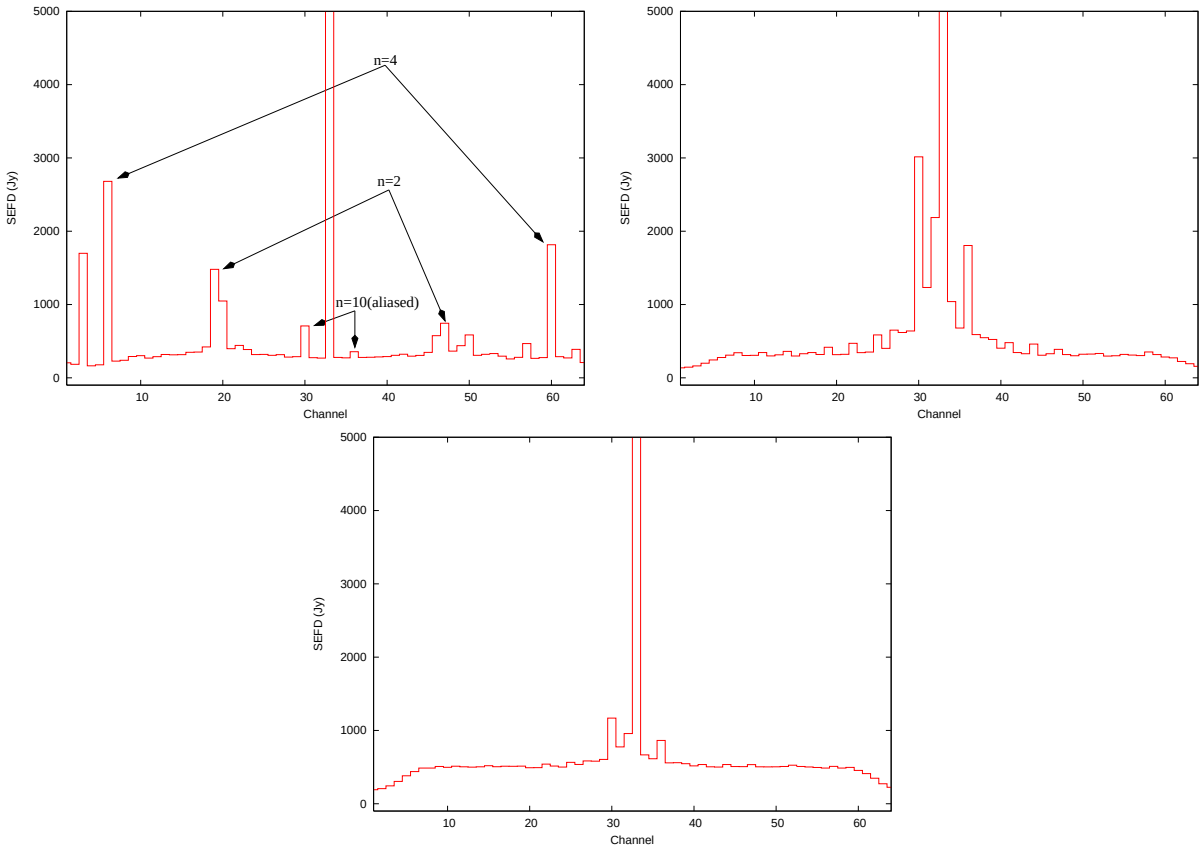

Figure 5: Modified SEFD as a function of channel for antenna 5. The three panels correspond to settings of high power, fast $f_{\rm shift}$; high power, slow $f_{\rm shift}$; and low power, slow $f_{\rm shift}$respectively.

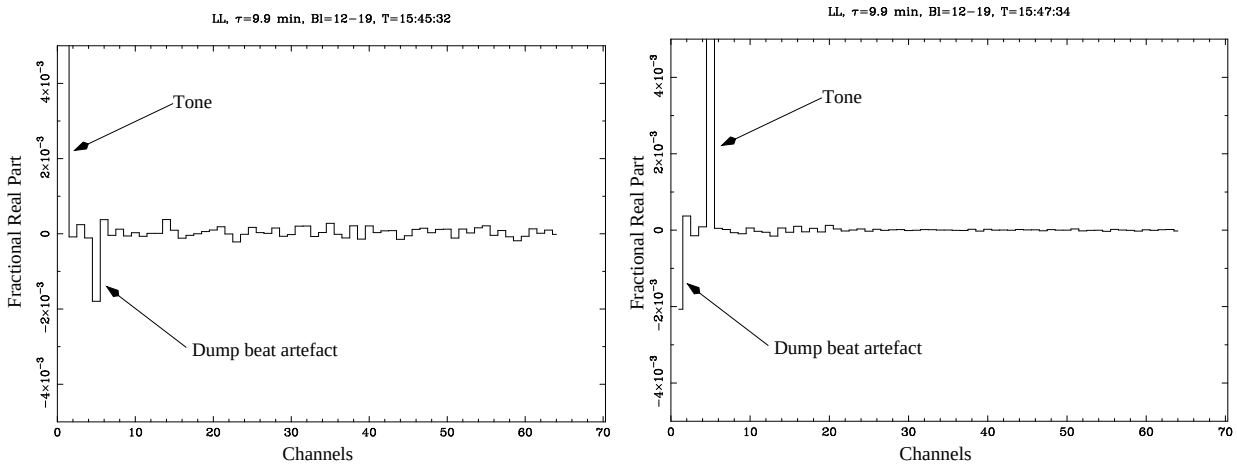

Figure 6: The correlator chip dump beat artifact: The left and right panels show the artifact for frequency configurations C and D for baseline 12-19.

Thus for the effect to be observable, two conditions must be met: the relative potential error, ϵ, must be detectable and the spectrum structure and $f_{\rm shift}$ frequencies must be such that the aretfact is shifted into positive frequencies. Both are quite stringent conditions. We investigate the potential to see this effect in the test observation.

Firstly, we find baselines which have relative amplitude of the error, ϵ, which are appreciable. Using the analysis from Appendix B, for a 28-antenna array, and using the $\tau_{\rm dump}$ of the tests, we find:

- For $\tau_{\rm dump}$ of 390.625 and 400μs, a total of 14 out of 1512 baselines (378 baselines for a 28-antenna array, for two dump times, for two $f_{\rm shift}$ settings) with appreciable potential error. These gave ϵ of 0.25%, which corresponds to the dump period being a multiple of the phase wash cycle period. No other baselines had potential error greater than 10^{-5}.
- For $\tau_{\rm dump}$ of 401μs, 15 out of 756 baselines had potential errors of $10^{-4} - 10^{-5}$. Other baselines had smaller potential errors.

Secondly, we see which frequency set-ups are conducive to seeing the artifact. For 3-cm observations, the observing band is in the lower sideband. $f_{\rm shift}$ frequencies are such that they increase the frequency of the first local oscillator. Taking these two characteristics together, the $f_{\rm shift}$ frequencies in a subband can be thought of as being negative. This means that frequency configuration C and D are conducive to seeing the beat effect.

A baseline conducive to seeing the effect is 12-19: this is a baseline where, for $\tau_{\rm dump} = 400\mu$s and the "fast" $f_{\rm shift}$ frequencies, there is an exact beat between the correlator chip dump time and the phase washing cycle ($\nu_{12} = 18670$ and $\nu_{19} = 11170$ Hz, $\nu_{12} + \nu_{19}$ is 3.8 channels). For frequency configurations C and D, where the tone is at channels 1 and 5 respectively of the second subband, the artifact is expected at channels 5 and 1, respectively, at near the 0.25% level. Figure 6 shows normalised plots of the observed spectra. The spectra show the artifact, but at the 0.2% level. The reduction in the amplitude of the effect from 0.25% to $\sim$0.2% is a result of the sinc response (the artifacts falls 0.2 channels from a channel center - the sinc pattern is seen faintly in the observed data) and some uncorrected phase errors.

To determine whether the correlator chip dump beat artifact is significant for continuum observing, we have

computed the expected potential errors, ϵ for standard wideband observing with WIDAR. In particular, we assume $\tau_{\rm dump} = 400\mu$s and $f_{\rm shift}$ frequencies between 75100 and 788300 Hz. We find no baselines with significant expected error.

Correlator chip dump beat artifact - Adjacent subband effect

NRC-EVLA Memo 32 first noted and investigated the correlator chip dump beat artifact. It did this while investigating the anti-alias response of WIDAR. Given the analysis of Appendix B, it is apparent that the artifact is a true sidelobe response of the system, albeit at a location that is usually unimportant - negative frequencies. However as negative frequencies in a subband correspond to positive frequencies in the adjacent subband, it is likely that the artifact will fall into the adjacent subbaand. We again note that phase washing does not beat down the response of a sidelobe falling in an adjacent subband - only the subband filter response will affect the apparent amplitude of a sidelobe response. At least for its primary response, the correlator chip dump beat artifact does not appear to be related to harmonics generated in the correlator cell phase rotator.

We wished to check the test observations to see if the dump beat artifact was present in an adjacent subband. For our observations, the most auspicious (or inauspicious?) settings to see the effect in the adjacent subband is frequency configuration B (configuration A has the artifact falling deep in the stop band of the subband filter response, and configurations C and D have the artifact falling in-band, not adjacent-band). The artifact will be apparent on baselines where the correlator chip dump time is a multiple of the phase wash cycle time. Again baseline 12-19 for $\tau_{\rm dump} = 400\mu$s is such a case. The expected artifact relative amplitude is $\epsilon = 2.5 \times 10^{-3}$ times the subband filter response of 1.9×10^{-3}. This gives an expected artifact at -53 dB. The expected sensitivity of the experiment reaches only -50 dB on this baseline. Thus it is expected that there is insufficient sensitivity to see this beat artifact. Indeed this was the case.

Aliased signal and self noise

An important design driver of WIDAR has been to provide good isolation between subbands even in the presence of strong signals. In a "conventional" system, isolation between adjacent subbands would be compromised by the aliasing of signals between subbands. In these systems, aliases are suppressed by the subband filter response, which may not be sufficiently good near the subband edges. The WIDAR concept of "phase washing" adds an extra level of rejection to aliased signals. See NRC-EVLA memos 1 and 32 for more details. NRC-EVLA memo 32 notes that, whereas phase washing is effective at reducing an aliased "signal", it has no effect on the aliased self noise of a strong signal.

As rules of thumb, we would expect the rms of the aliased signal to be of order

$$\sigma \sim \frac{rS}{4\pi|\nu_1 - \nu_2|T}$$

and the rms aliased self noise to be of order

$$\sigma \sim \frac{rS}{\sqrt{BT}},$$

where B is the width of a channel, T is the integration time, ν_1 and ν_2 are the $f_{\rm shift}$ frequencies, r is the subband filter response to the alias, and S is the signal strength. For our observations, for an $f_{\rm shift}$ frequency of ν and subband edge at ν_0, then the alias for a signal at $\nu_0 + \delta\nu$ is at $\nu_0 - \delta\nu + 2\nu$ (note when tuning near 8350 MHz, $f_{\rm shift}$ frequencies can be thought of as negative).

The data that would show the largest 'aliased signal' component is baseline 25-26, frequency configuration D, the "slow" $f_{\rm shift}$ frequencies and correlator chip dump time of 401 μs[3]. The alias should appear at channel 61 in the first subband. For this case, the subband filter rejection, r is about 1.01×10^{-2}, and the $f_{\rm shift}$ frequencies are 470 and 590 Hz. For a 1 s integration and signal strength of 10^5 Jy, the rules of thumb suggest an aliased signal strength of 0.6 Jy, and rms aliased self noise of 13 Jy.

Because the aliased signal averages down much faster than the aliased self noise, the rules of thumb suggest the aliased signal will not be detectable. We expect that there will be an enhanced "noise" level at channel 61 because of aliased self noise and that this should average down as $\sqrt{T}$.

This is indeed what has been found. Figure 7a shows the scalar average of the amplitude on subband 1. The increase in level at channel 61 is very apparent. The level for other channels is consistent with the expected thermal noise. Figure 7b shows how the estimated variance of the real and imaginary parts of the data of channel 61 goes down as the square root of the number of samples averaged. This plot is consistent with the excess in channel 61 being aliased self noise. It is not consistent with it being aliased signal.

We have analyzed other baselines and the other sets of data with different parameter settings. They are consistent with the above analysis. The observations had not 'detected' the aliased signal. The aliased self noise appears to be more significant for practical cases.

Generalized splatter

> *"The term numerologist is also used derogatorily for those perceived to place excess faith in numerical patterns"*

WIDAR uses a novel architectural approach - "phase washing" - to beat down a number of artifacts and aliases that would otherwise affect the data. Phase washingis very effective at producing high dynamic range spectra. However it is also important to chose the $f_{\rm shift}$ values with care. Poor choice of $f_{\rm shift}$ can lead to beat patterns and artifacts. EVLA Memo 157 analyzes the "splatter" artifact, where the harmonics of the correlator cell phase rotator phase stop the astronomical signal. Harmonics result from the three-level nature of the phase rotator as well as 4-bit quantization of the phase model that is used by the phase rotator. That memo gave the condition for splatter being when $m\nu_1 \approx n\nu_2$ for small, positive integers m and n.

The test observations of this memo showed splatter is a more general and pervasive effect. A "generalized" splatter was seen on a number of baselines, but at a level about a order of magnitude less than that reported in EVLA Memo 157. Given EVLA Memo 157, this splatter was unexpected. What EVLA Memo 157 failed to consider when investigating splatter was the effect of aliasing of the phase rotator harmonics. When aliasing is considered, the more general condition for splatter to occur is when

$$m\nu_1 + n\nu_2 = k\,BW.$$

Here m and n are non-zero (possibly negative) integers, k is a non-negative integer and BW is the bandwidth of the subband. Simple simulations of the WIDAR correlator cell architecture readily reproduce these generalized splatter artifacts.

In addition to this aliased splatter mechanism, we have also found instances where we cannot find simple relationships between the $f_{\rm shift}$ frequencies, bandwidths, etc yet artifacts which are otherwise qualitatively similar

[3]The data for correlator chip dump times of 401 μs should have better potential to show the aliased signal component because the total integration time is not an exact number of phase washing cycles.

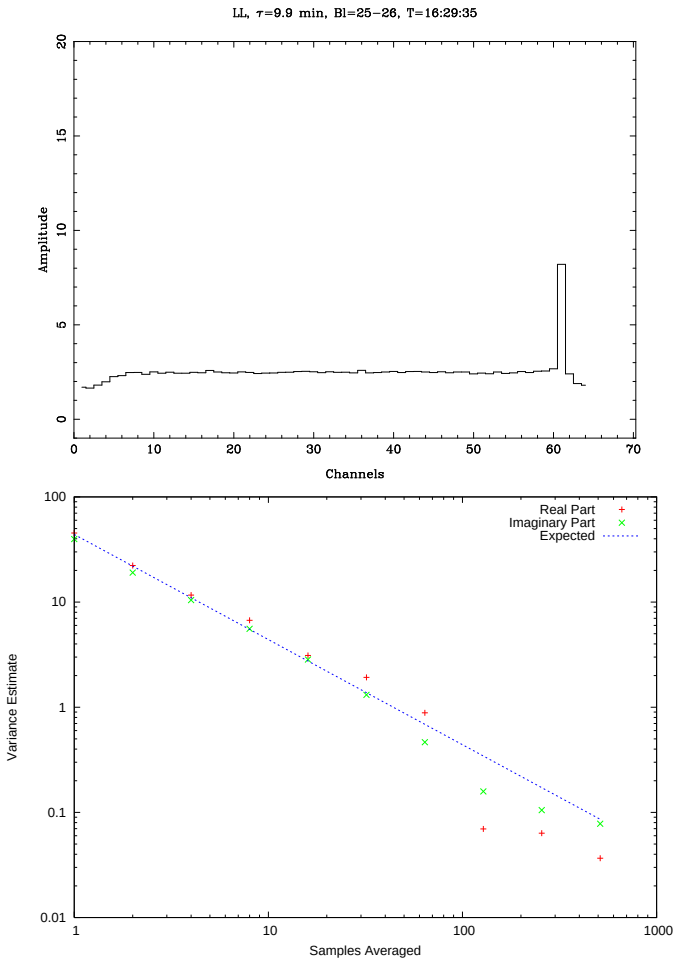

Figure 7: Aliased self noise. The upper panel shows the scalar amplitude average for subband 1. The increase in level at channel 61 corresponds to expected location of the alias from subband 2. The lower panel shows that the estimate of the variance of the real and imaginary parts of channel 61 average down as the square root of the number of samples - as would be expected from aliased self noise (but not phase-washed aliased signal).

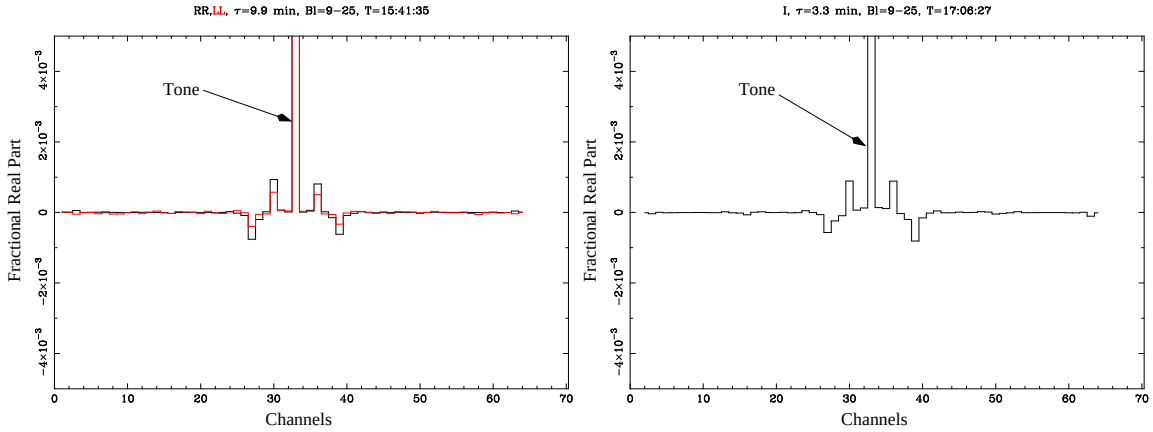

Figure 8: Generalized splatter: Left panel shows test data with satellite lines appearing as artifacts around the main tone. The right panel shows a simulation of this observation, which reproduces the artifacts.

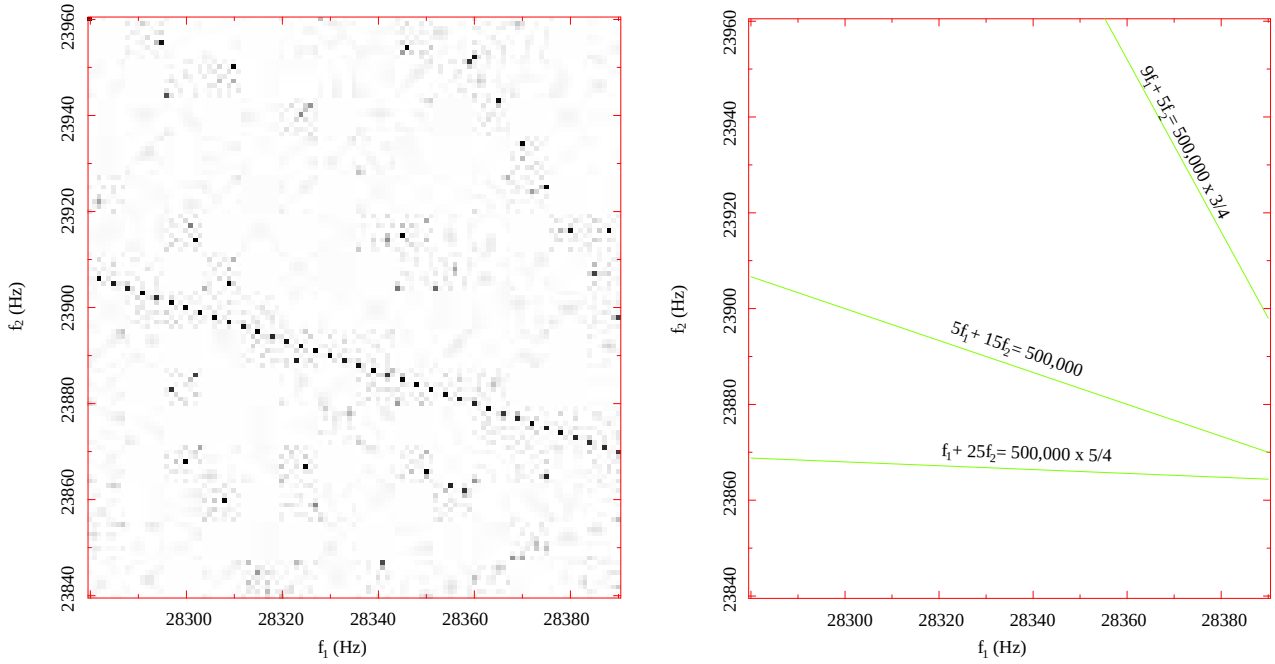

Figure 9: $f_{\rm shift}$ space: The left panel shows values of $f_{\rm shift}$ where artifacts are apparent. The right panel gives the numerology of some lines of significant artifacts that are apparent in the left panel.

to "splatter" are seen. Simulations also show these artifacts. The only finite-precision effects needed to for them to appear in the simulations is the three-level phase rotator and quantized phase models. As with "classical" splatter, the artifacts vanish if the simulations use either a high precision phase rotator or or a high precision phase model.

Figure 8 shows an example of an observed and simulated artifact (frequency configuration A). This corresponds to $\nu_1 = 21790$ Hz and $\nu_2 = 5630$ Hz. This was one of the instances where no "numerology" was discovered between the $f_{\rm shift}$ frequencies and/or the bandwidth.

Given that simulations reproduce the effects well, it is instructive to use simulations to explore $f_{\rm shift}$ space. Figure 9 shows an example of a 110×120 Hz portion of $f_{\rm shift}$ space. The left panel (grayscale plot) shows the severity of the artifacts for a particular $f_{\rm shift}$ pair, with black representing a relative error of 1×10^{-3}. The right panel shows the numerology for some paths of large artifacts apparent in the left panel. It is clear that, in addition to the paths of poor artifacts associated with some particular numerology, there are many points in $f_{\rm shift}$ space where artifacts will be detectable. However, most $f_{\rm shift}$ choices lead to low artifact levels. The root mean square relative artifact level in the left panel is 9×10^{-5}. Simulations of other parts of $f_{\rm shift}$ space were found to have a similar rms relative artifact level.

The relative amplitude of these generalized splatter artifacts that we have noted (test observations and simulations) were always 1×10^{-3} or less. This relative amplitude is set by the $f_{\rm shift}$ setting. These artifacts are phase coherent with the source, and so may not average down with time. But this assumes that the natural fringe rate is 0. However the range of $f_{\rm shift}$ for particular artifacts is very narrow: the artifacts require quite specific values of $f_{\rm shift}$. For sources with natural fringe rates, these artifacts will be averaged down with fringe rotation. Given this, plus their intrinsic weakness, their significant variation between baselines, and their narrowness with $f_{\rm shift}$ frequency, we do not believe they are an important error source.

If these artifacts are considered significant, NRC-EVLA memo 9 notes an approach to reducing them further. That memo suggests that a small ($\sim$ Hz) residual phase rate be left uncorrected by the correlator cell phase rotator. This residual phase rate can then be corrected in the back-end software using a software (high precision) phase rotator. This approach would wash out the artifacts at the residual phase rate.

Conclusions

We have taken data which would be taxing to any system: we have an observation with of a line with equivalent strength of 10^4 to 10^5 Jy. Placing the line at the north celestial pole prevents fringe rotation from washing out some artifact responses. Taken in this light, WIDAR has performed exceptionally well. The artifacts that we identify are small and not generally not significant. WIDAR has exceptional spectral dynamic range.

The effects identified and analyzed do not appear to warrant further action. None of the effects noted here would prevent 1:10,000 spectral dynamic range, and probably significantly better. The possible only exception to this is the complex-valued nature of the WIDAR spectral point-spread function. Although this does not appear to be a significant issue in practise, it is easily remedied.

Appendix A: WIDAR spectral point-spread function

For a pure lag correlator where the lags are uniformly weighted, it is widely stated that the spectral point-spread function (PSF) is a sinc function. It is usually appreciated that this is only a good approximation, and that the exact spectral PSF is somewhat more complicated. This PSF, however, does have a comparatively simple form.

For WIDAR, the measured complex-valued lag function for a spectral tone at ν_0 is

$$\exp(i2\pi\nu_0 k\Delta\tau), \qquad k \in [-N/2, N/2-1].$$

Here N is an power of two[4]. For our observations, there are 64 visibility channels produced from 128 complex-valued lags.

A closed form for the Fourier transform of this lag spectrum is readily derived. The spectral PSF is:

$$h(\nu) = \frac{1}{N}\exp\big(i\frac{\pi}{N}\frac{\nu}{\Delta\nu}\big)\,\frac{\sin(\pi\frac{\nu}{\Delta\nu})}{\sin(\frac{\pi}{N}\frac{\nu}{\Delta\nu})}$$

where $\Delta\nu$ is the channel spacing. Near the center of the response, where

$$\frac{\pi}{N}\frac{\nu}{\Delta\nu} \ll 1,$$

using small angle approximations, the real part of $h(\nu)$ reduces to the sinc function. This is as expected. Somewhat surprisingly, $h(\nu)$ is complex-valued. The imaginary part of $h(\nu)$ has a simple form:

$$\Im[h(\nu)] = \frac{1}{N}\sin(\pi\frac{\nu}{\Delta\nu}).$$

The complex-valued nature of the PSF is a consequence of the even number of lags used. In the set of lags, it is apparent that every negative lag except one has a partnering positive lag. The combined response of the partnered lags do not lead to a imaginary response. The imaginary part of the PSF is a result of the one lag without a partner – the first lag value in the notation of this appendix. An odd number of lags (lags that are symmetric about zero lag) would have a purely real-valued PSF.

Appendix B: The correlator chip dump beat artifact

The WIDAR correlator uses short-term (on chip) and long-term accumulators. When the transfer between the accumulators is being initiated, there is a brief blanking period when no accumulation is performed. The blanking period is $\tau_{\rm blank} = 1\mu$s. The current standard dump time for the correlator chip accumulators is $\tau_{\rm dump} = 400\mu$s.

The blanking results in small (insignificant) loss of signal. The correlator chip beat artifact results when the dump time, and so the blanking period, happens to be synchronized with difference of the $f_{\rm shift}$ frequencies on a baseline. This is an effect first noted and partially analyzed in NRC-EVLA memo 32. That memo was most concerned with the effect this had on alias suppression.

[4]The basic lag architecture of WIDAR tends to produce lags that are a power of 2. However it has the flexibility with recirculation and board stacking to generate N which is not a power or two or indeed even.

To analyze this effect, we will initially ignore fringe rotation (i.e. we will consider a source at the north pole). Consider the lag function measured by the correlator:

$$\ell(\tau_1,\tau_2) = \left\langle \cos\Big((\mathbf{\Omega}+\omega_1)(t+\tau_1)\Big)\cos\Big((\mathbf{\Omega}+\omega_2)(t+\tau_2)\Big)\exp\Big(i\big(\omega_1(t+\tau_1)-\omega_2(t+\tau_2)\big)\Big)\right\rangle$$

where ω_1 and ω_2 are the $f_{\rm shift}$ angular frequencies and $\mathbf{\Omega}$ is the angular frequency of the source emission. The latter is a stochastic process varying around ω. Expanding and ignoring the terms in $\mathbf{\Omega}t$ (these have an expected value of 0), then

$$\ell(\tau_1,\tau_2) = \left\langle \exp\Big(i\big(2(\omega_1-\omega_2)t + 2\omega_1\tau_1 - 2\omega_2\tau_2 - \mathbf{\Omega}(\tau_2-\tau_1)\big)\Big)\right\rangle + \exp\Big(i\tau\omega\Big).$$

Here $\tau = \tau_2 - \tau_1$. The second term is the desired one: it is the term that the correlator is intended to measure. Generally the first term will have an expected value of 0 - it will 'wash out'. However if there is blanking which is synchronous with $2(\omega_1-\omega_2)$, the first term will not completely wash out. The first term, which we represent as $\ell_{\rm a}(\tau)$, is the lag function of the correlator chip dump beat artifact.

Consider the case where the correlator chip dump period is a multiple of the phase washing cycle time, and assume that the $\tau_{\rm blank}(\omega_1-\omega_2) << 1$[5]. Taking into account the details of the blanking and the geometry of the WIDAR delay ladders, and ignoring some arbitrary phase terms, the first term in the above equation has an expected value of

$$\ell_{\rm a}(\tau) = \begin{cases} a\exp\big(-i\tau(\omega+\omega_1+\omega_2)\big) & \text{for even lags;} \\ a\exp\big(-i\tau(\omega+\omega_1+\omega_2)-(\omega_1-\omega_2)\Delta\tau\big) & \text{for odd lags.}\end{cases}$$

Here a is $\tau_{\rm blank}/\tau_{\rm dump}$ (i.e. the fraction of time that the correlator is blanked)[6]. The lag function is sampled at increments of $\Delta\tau$ (i.e. the sampling period). Note the additional phase term in the case for odd lags. This is a consequence of the geometry of the WIDAR delay ladder. However the phase is near zero and can be ignored to first order.

Recall that WIDAR measures a complex-valued lag function and the negative frequencies of the Fourier transform of the lag function are expected to be noise. The lag function $\ell_{\rm a}(\tau)$ represents an artifact at negative frequencies. Generally negative frequencies are ignored. If the true source visibility spectrum is $V(\nu)$, then the expected value of the artifact will have a spectrum of $aV(-\nu-\nu_1-\nu_2)$. There is potential, however, for the artifact to affect positive frequencies: when $\nu+\nu_1+\nu_2$ is larger that the bandwidth, or less than zero, then the negative frequencies can wrap around into positive frequencies. Thus if there is emission at the very edge of a subband (within $\Delta = \nu_1+\nu_2$ of the edge), then this can wrap into positive frequencies. Observing with emission near the subband edge is probably most relevant for continuum science.

The above analysis has focused on a correlator chip dump period that is a multiple of the phase washing cycle time. When this is not the case, the response is a sum of the response in each correlator chip dump. In particular, the relative amplitude of the effect, ϵ, will be the average of the individual dump responses, i.e.

$$\epsilon = \frac{a}{N}\sum_{n=0}^{N-1}\cos\big(2(\omega_1-\omega_2)n\tau_{\rm dump}\big)$$

where N is the number of correlator chip dumps in a correlator integration and a is (as before) the ratio of the blank to dump periods[7].

[5]This assumption is not necessarily good for wideband observations where the $f_{\rm shift}$ frequencies can be large. The assumption, however, leads to a conservative estimate. When the $f_{\rm shift}$ frequencies are large, there is some additional beating down of the response within a blanking period.

[6]Note the desired response is also reduced by a factor of $1-a$

[7]This equation for ϵ is closely related to equations on pages 39 and 40 of NRC-EVLA Memo 32.