EVLA WIDAR Correlator:

Anti-Aliasing and Bin-Centered Tone Investigation

NRC-EVLA Memo# 032

Brent Carlson, October 7, 2010

ADDENDUM: March 24, 2011

ABSTRACT

This memo more thoroughly investigates WIDAR correlator anti-aliasing behaviour than previously performed. In addition, reported low-level artifacts surrounding bin-centered, timing-reference-locked tones are duplicated. Solutions, in terms of reasonably careful choice of antenna fshift parameters, and correlator chip integration times are found and tested. A more complete theory of anti-aliasing behaviour is also developed.

Introduction

For quite some time, there has been reported evidence that anti-aliasing behaviour observed when testing with the EVLA has not been observed to follow expected results as reported in [1]. Recently, more detailed testing has been performed at the EVLA with a system-timing locked, bin centered tone (Bob Sault's report of August 24, 2010), providing clear evidence that there are some unexplained artifacts in the results, in particular unexpected anti-aliasing attenuation, and low-level artifacts surrounding the test bin-centered tone.

This memo investigates these issues more thoroughly, resurrecting the software correlator simulator (used to originally investigate WIDAR signal processing) to run tests under well-known and controlled conditions. Hardware testing using Station Board "Delay Module Test Vectors" is also performed to ensure that the real hardware is behaving as the software correlator indicates it should.

During the development of this document and the testing that went into it, some new strange behaviour was encountered, the effects of which have been tested and incorporated. This document has therefore grown in scope to describe these findings, and some of the wording in this memo reflects that process.

Anti-Aliasing Principles

Anti-aliasing behaviour in the WIDAR correlator has not been found to strictly follow what the current theory [1] is, namely that aliased signals wash down as:

$$\frac{1}{2\pi(2\,fshiftX - 2\,fshiftY)T}$$

where fshiftX is the LO frequency offset in antenna X, and fshiftY is a different LO frequency offset in antenna Y. *T* is the integration time.

It seems that the anti-aliasing hits a "floor" of around -30 dB or so and that the magnitude of the differential fshifts don't seem to matter as long as they are non-zero (kHz or so).

I believe this is because the original theory is not complete and that there are one or more missing components.

Consider the cross-correlation of a broadband noise source, and another statistically independent broadband noise source. It is well known, in this case, that the amplitude of the correlation coefficient will integrate down proportional to $1/\sqrt{N}$, where N is the number of statistically independent samples correlated. An amplitude vs frequency plot of a simulation of this case is shown in Figure 1, for several decades of integration time.

The curves in Figure 1 follow the equation, resulting in the noise integrating down at the rate of ~-5 dB per decade of samples.

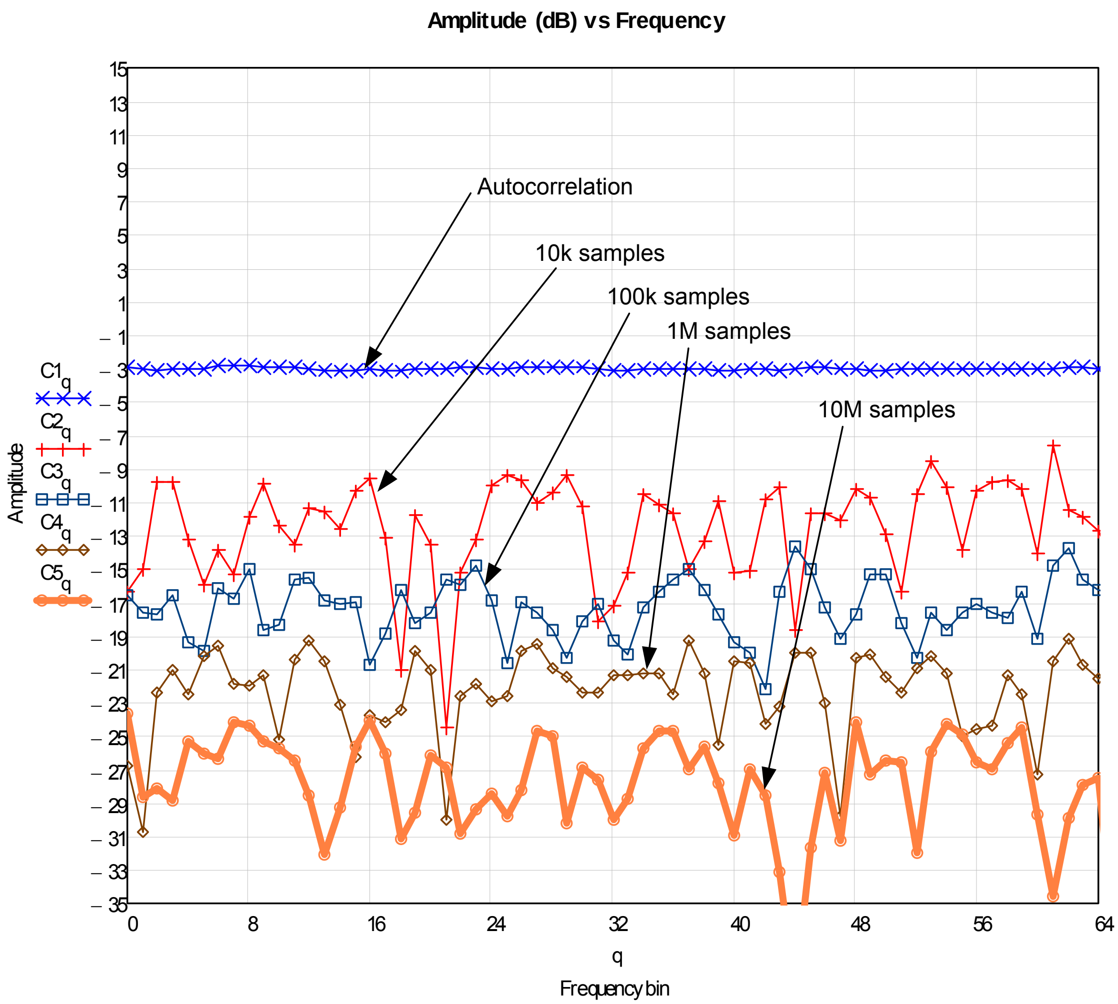

Figure 1 Uncorrelated noise spectrum showing decrease of ~5 dB per decade.

Now, consider the cross-correlation of a broadband noise source against an independent broadband noise source into which a tone of some amplitude, at an arbitrary frequency, has been injected. Everywhere the tone isn't the noise integrates down at -5 dB per decade as above. The place in the spectrum where the tone is placed also integrates down at this rate, however there is a "ghost" tone in the cross-power spectrum that is some level above the noise floor, and that has an uncertainty (noise) as any other place, as it has the same additive noise. A graph showing several decades of correlation is shown in Figure 2.

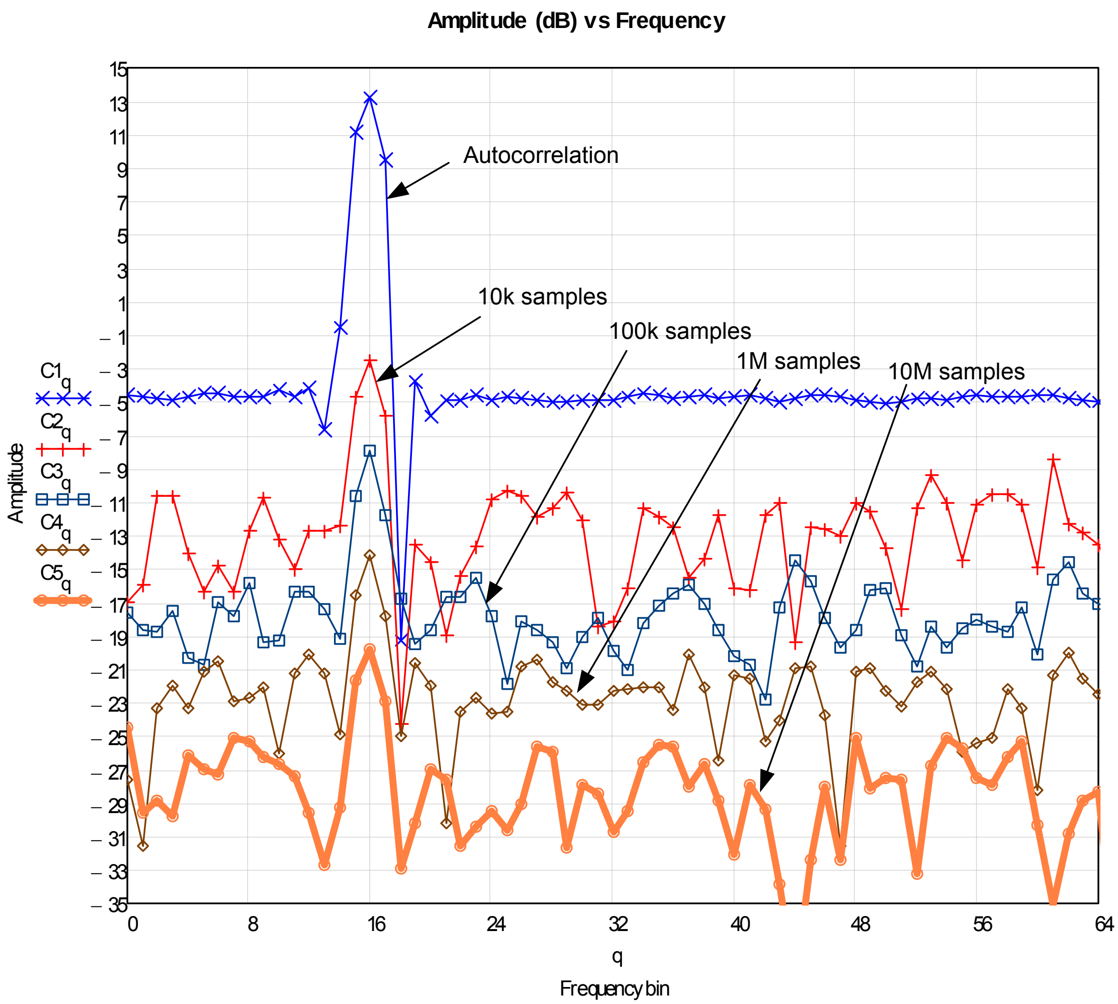

Figure 2 Uncorrelated noise with a tone injected into one antenna.

The ~-5 dB/decade reduction in the peak of the ghost tone is evident.

This test demonstrates that a tone injected into one signal, correlated against a statistically independent noise source, will see a "ghost" tone that integrates down at the same rate as the noise, $1/\sqrt{N}$, with an uncertainty in amplitude equal to the noise.

Next, consider 2 tones, the first tone injected into the first (X) antenna, and the second tone with a significant frequency difference injected into the second (Y) antenna. A plot highlighting this cross-correlation with only a 1M sample correlation performed is shown in Figure 3.

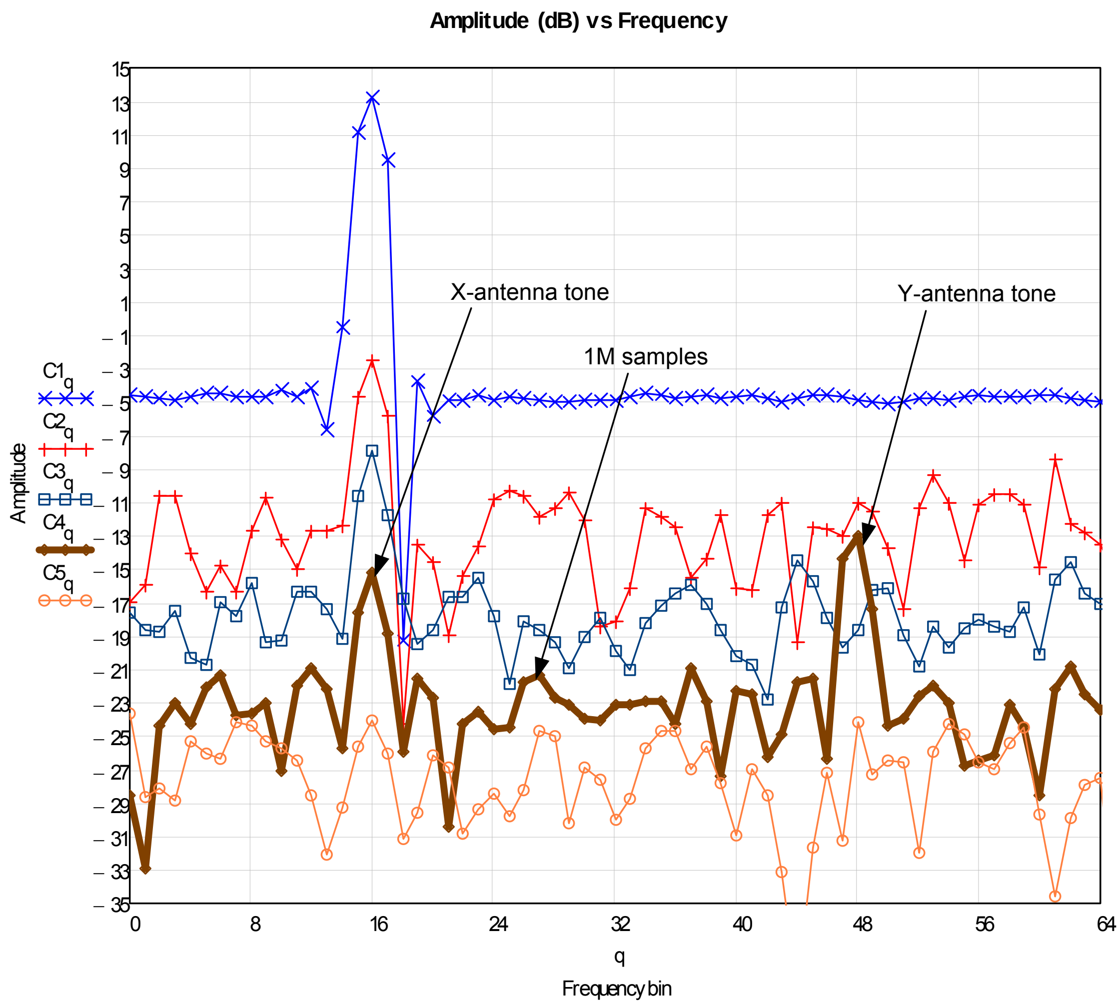

Figure 3 Uncorrelated noise with a different tone injected into each antenna.

Both tones are independent, and both integrate down with the noise independently.

Next, consider 2 tones, but at varying frequencies. If they are exactly the same frequency, the tones will correlate. As their frequencies change, they will correlate somewhat depending on their difference frequencies and depending on the integration time. Intuitively, the closer they are in frequency, the more they will correlate for a given integration time, and the farther apart they are in frequency, the less they will correlate until they are effectively independent of each other. The correlation of these tones is given by the well-known washing equation:

$$CC_{tones}=\mathrm{sinc}(2\pi(fx-fy)T)$$

Where fx is the tone frequency of the first antenna (X), fy is the tone frequency of the second antenna (Y), and T is the integration time.

The envelope of this function is simply $1/(2\pi(fx-fy)T)$, as long as the denominator is > 1.

The envelope and the function is shown in Figure 4 for a 1 second integration time, fx=100 Hz, and fy=200 Hz (i.e. a 100 Hz differential):

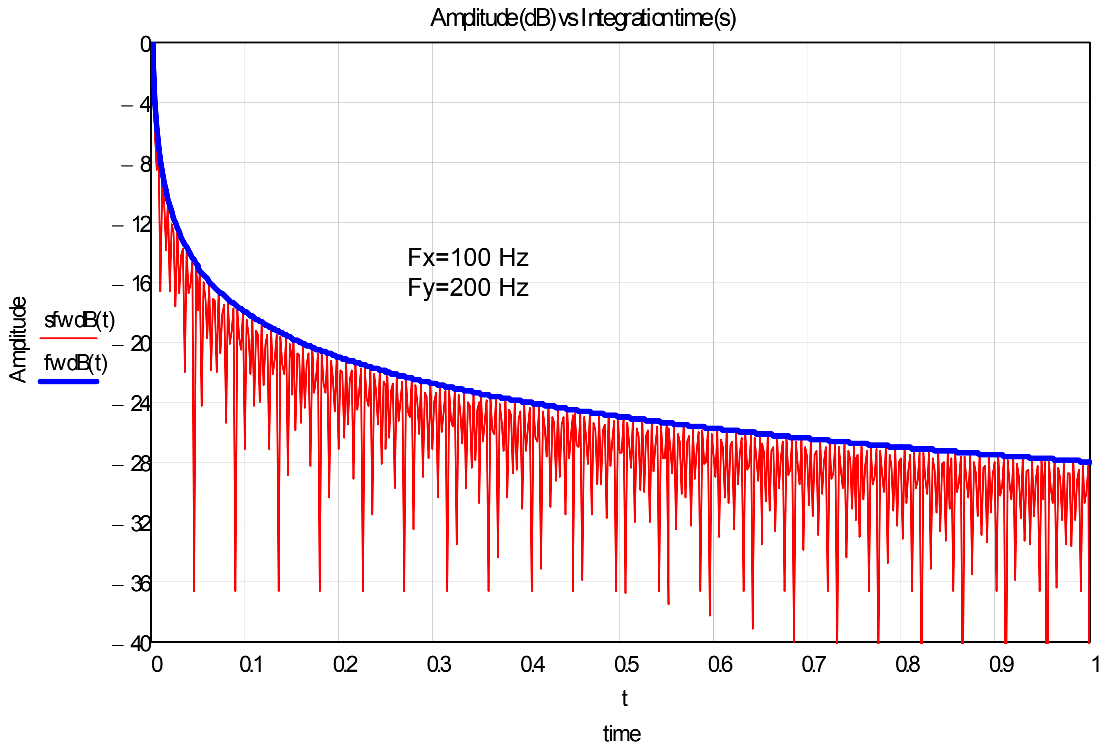

Figure 4 Sinc(x) washing function and envelope for 100 Hz frequency difference, and 1 second integration time.

If the difference between the tone frequencies is increased, say, by a factor of 10 (1 kHz, and 2 kHz tones), then the plot of Figure 5 indicates a much reduced correlation:

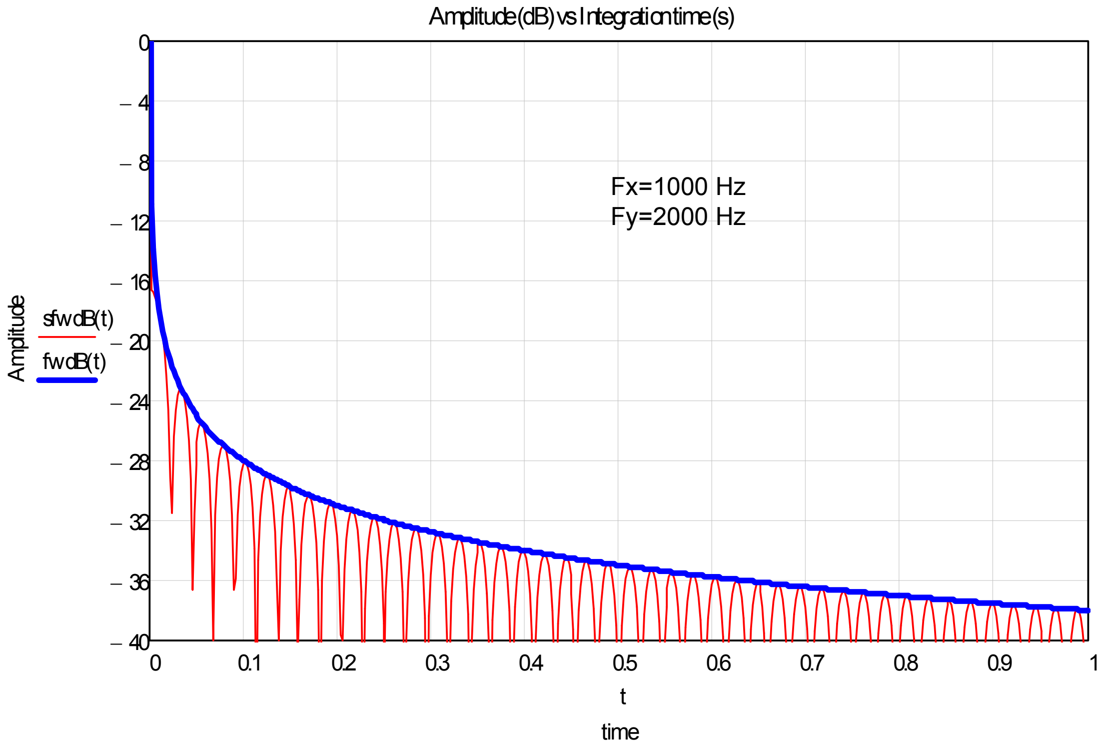

Figure 5 Sinc(x) washing function for a 1 kHz frequency difference.

It seems reasonable then, that there may be a cross-over between the 1/√N noise de-correlation curve and the sinc(x) washing function, depending on chosen parameters:

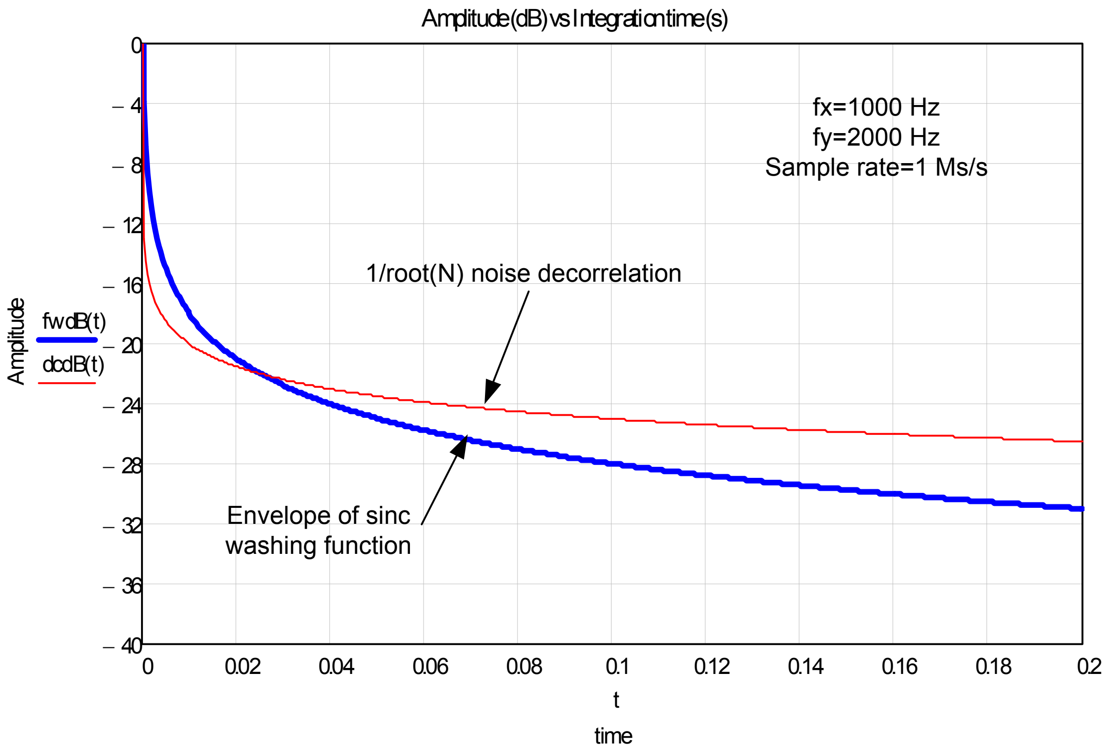

Figure 6 Comparison of sinc(x) washing function and 1/√N de-correlation, for select parameters.

In this case (fx=1000 Hz, fy=2000 Hz; sample rate is 1 Ms/s) after about 25 msec, the amplitude of the tones in the cross-correlation spectrum is dominated by 1/√N de-correlation and no further integration or larger frequency difference will matter; the tones are now independent.

This is a good news/bad news situation compared to the flawed original (sinc(x)) only theory. The good news is that for a given integration time and sample rate, the frequency difference (fshift difference) can be minimized so that it only does at least as well as 1/√N de-correlation, thereby minimizing bandwidth loss. Additionally, longer integration times (either incoherent—simple averaging, or coherent—via the FFT) will continue to beat down the aliasing levels as noise integrates down, provided the sinc(x) washing limit is well below the noise limit. The bad news, is that washing of the aliased signals is not as great as originally planned, but there is, unfortunately, no fundamental way to change this.

The graph of Figure 7 shows a comparison between the two curves for a 1 Ms/s sample rate (500 kHz bandwidth), with a frequency difference of 200 Hz.

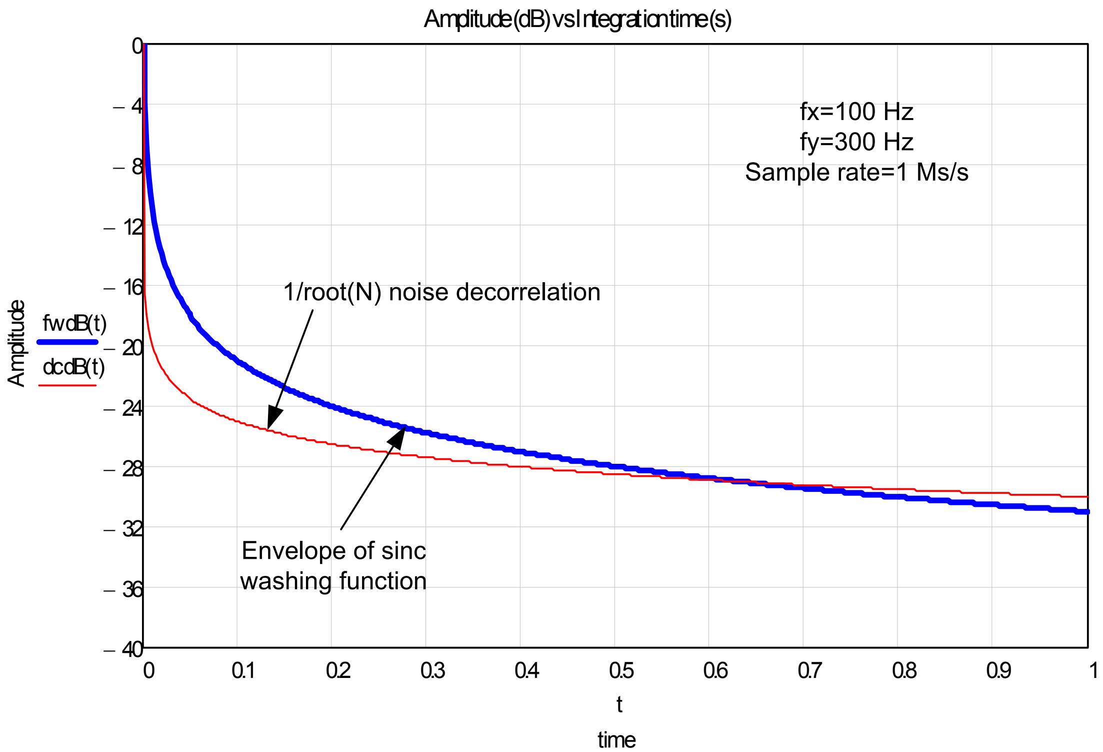

Figure 7 Comparison of sinc(x) washing function and 1/√N de-correlation for 1 Ms/s sample rate, and frequency difference of 200 Hz.

The 1/√N de-correlation always sets a limit, and so a "null" in the observed net washing/ de-correlation function does not exist, as has been ostensibly observed during sky testing.

Additionally, the amount of washing/de-correlation does not change with larger fshift differences for the same reason (provided the sinc(x) curve is below the 1/√N curve).

Anti-aliasing and Bin-Centered Tone Simulation Tests

Three tests (referred to as **Test #1**, **Test #2,** and **Test #3**) were performed to determine if anti-aliasing performance, in the software correlator simulation code originally developed to test correlator signal processing [1], is consistent with expected de-correlation/washing predictions as outlined in the previous section. Due to very long simulation times required to de-correlate down to the ~-35 dB level, a very strong tone was used. In the tests, the tone power was equal to ~1/3rd of the total power in the *wide* band (before 1/16 sub-band filtering). For both tests, the following conditions were used:

- The noise generator used a single-sideband mixer, 163-tap Hilbert FIR filter, with double-precision floating point for mixing and coefficients. Hilbert coefficients were Hanning windowed. This SSB mixer was used to add fshifts.
- 4-bit initial sampling, sigma ~=2.7.
- Correlator FIR filter was 1023-tap, 4-bit inputs, 12-bit LUT, and all integer operations afterwards. 1/16th band filtering was used.
- 4-bit re-quantizer, sigma ~=2.7.
- 4-bit, 15-level correlation with 3-level fringe rotation, operations performed exactly as in the actual correlator. For Test #1 a 16,384 lag (8192 channel) correlation was performed to allow discernment of the different frequency components that might be present. For Test #2, 2048 lags were used as the fshift frequencies are too small to resolve even with a much larger number of channels. Some tests were run with a double-precision floating-point fringe rotator as a comparison.

All tests were arranged according to the following diagram (fshift frequencies for Test #1 are shown in the figure). A single (after correlation) bin-centered (or nearly bin-centered[1]) tone was used, within ~500 kHz of the edge of the band. All operations are of course discrete-time, and a reference sample rate of 2.00 Gs/s was used. With a 1.00 GHz bandwidth, and 8192 channels across 1/16th of that bandwidth, it is 7.62939453125 kHz per point. The tone frequency, before filtering and decimation, was 0.031005859375 cycles/sample=62.0117188 MHz, or 488.28125 kHz before the edge of the sub-band 0-

[1] After correlation, FFT of the data without windowing indicated ~uniform ~-42 dB sidelobes spread across the band; the reason for this was unclear, and may be due to a slight bin un-centeredness, possibly due to numerical precision error, as frequencies were defined in “cycles/sample”, with resulting calculations for long simulations resulting in less and less precision in the fractional part of the tone phase calculation. Or, I didn’t get the bin-centered frequency calculations just right. Or, I’m mixed up on some fundamental principle. Therefore, Hanning windowing, with ~-45 dB sidelobes was used and believed to be sufficient, although a Kaiser window with -65 dB sidelobes might have been better (but looks complicated to implement, especially that nasty Bessel function).

to-sub-band 1 boundary. Test #3 is set up to verify hardware operation, and so its reference sample rate is 2.048 MHz, with the tone 1 MHz from the edge of the band.

For **Test #1** the X-antenna chosen fshift frequency is 151.1 kHz, and the Y-antenna fshift frequency is 31.1 kHz. These are large enough so the sinc(x) washing curve is below the noise de-correlation curve for all tests. They are also large enough so that spectral artifacts from them can be discerned with the chosen frequency resolution. A 16384-lag correlator was used to obtain high enough spectral resolution so that individual aliased tones could be seen.

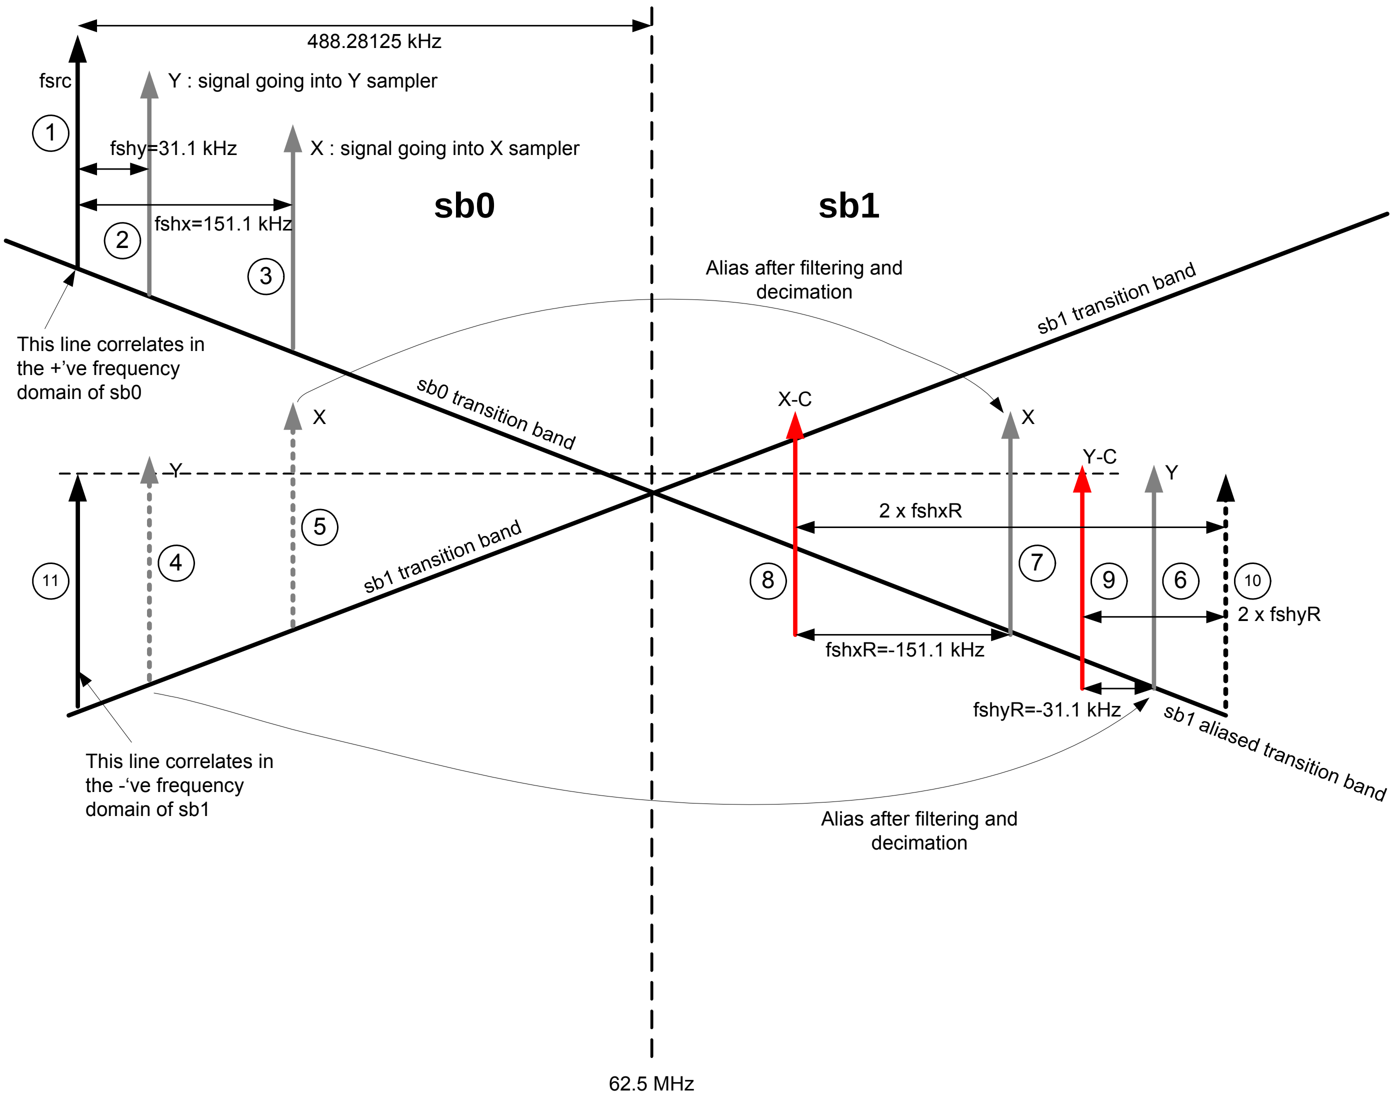

Figure 8 Anti-aliasing test signal setup.

For **Test #2** the X-antenna chosen fshift frequency is 1.511 kHz, and the Y-antenna fshift frequency is 0.311 kHz. These are small enough so the sinc(x) washing curve is well above the noise de-correlation curve for all tests. A 2048-lag correlator was used.

For **Test #3** various fshift frequencies were used to explore the effects of various ratios of fshift frequencies. A 2048-lag correlator was used.

The following notes are important to understand the above figure. Correlation of only sub-band 1 was performed[2], as it contained all required information to make dynamic range measurements. Numbers in the following list refer to "bubbled" numbers in the figure.

1. This is the original "sky" tone frequency, without regard to any fshift that may be imposed. It does not exist, except in sub-band 0 correlated output.

2. The Y-antenna fshift frequency puts the tone here, as seen by the Y antenna initial sampler.

3. The X-antenna fshift frequency puts the tone here, as seen by the X antenna initial sampler.

4. After sub-band filtering and decimation, the Y-antenna tone appears here in sub-band 1's negative frequency domain.

5. After sub-band filtering and decimation, the X-antenna tone appears here in sub-band 1's negative frequency domain.

6. After sub-band filtering and decimation, the Y-antenna tone aliases here into sub-band 1's positive frequency domain.

7. After sub-band filtering and decimation, the X-antenna tone aliases here into sub-band 1's positive frequency domain.

8. After fringe rotation, the X-antenna aliased tone appears here, 2 x fshiftX away from where the "virtual" reference tone (10) is. This is an artifact we want to de-correlate/wash out with time.

9. After fringe rotation, the Y-antenna aliased tone appears here, 2 x fshiftY away from where the "virtual" reference tone (10) is. This is an artifact we want to de-correlate/wash out with time.

10. This is the "virtual" reference tone. It doesn't actually exist anywhere, but would exist if no fshifts were introduced, and no fshift removal in the correlator were performed. i.e. without anti-aliasing, this is where the correlated aliased tone would appear.

11. This line correlates in the negative frequency domain of sub-band 1. To measure dynamic range (de-correlation of (8) and (9)), it sets our reference amplitude against which we measure low-level artifacts. Note that this sets a conservative

[2] sub-band 0 was correlated to ensure expected results were obtained, which they were (i.e. the tone in the expected place).

measurement, since due to the transition band slope, the actual aliased amplitudes (comparatively speaking) are set by the geometric mean of (6) and (7).

Test #1 Results

Test #1 was run for autocorrelation, and 3 decades of sample counts. The 1/√N noise de-correlation curve was well above the sinc(x) washing curve, as shown in the following figure:

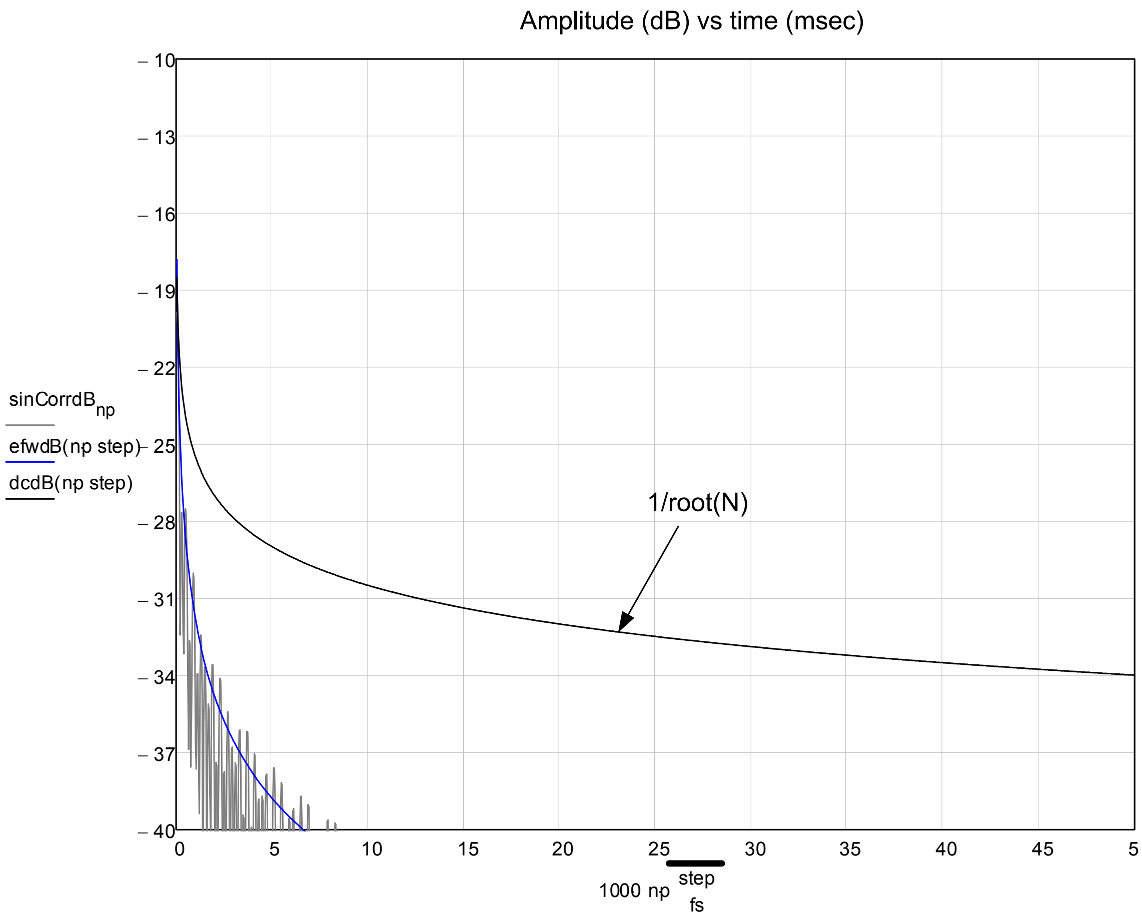

Figure 9 The 1/√N de-correlation curve for Test #1 limits anti-aliasing levels as it is well above the sinc(x) washing curve.

Results using a 3-level fringe rotator in the correlator are shown in the following figure.

Some points to note:

- The auto-correlation amplitude is not as high as the correlated tone (in the negative frequency domain) due to the fact that it is not bin centered. Auto-correlation was performed on X-antenna data without fringe rotation.
- There are definitely narrowband artifacts in addition to the aliased tones, coming from the 3-level fringe rotator's interaction with the strong tone. The following

figure shows the same results with floating-point fringe rotation and, except for one, they aren't there. However, these artifacts stay at or below the expected anti-aliasing levels, and integrate down with the noise.[3] One interesting artifact pointed out is (2fshiftX + 2fshiftY)/2, which also appears in the floating-point fringe rotation cross-correlation—is this of the form of a trigonometric identity stemming from sin(x)sin(y)?

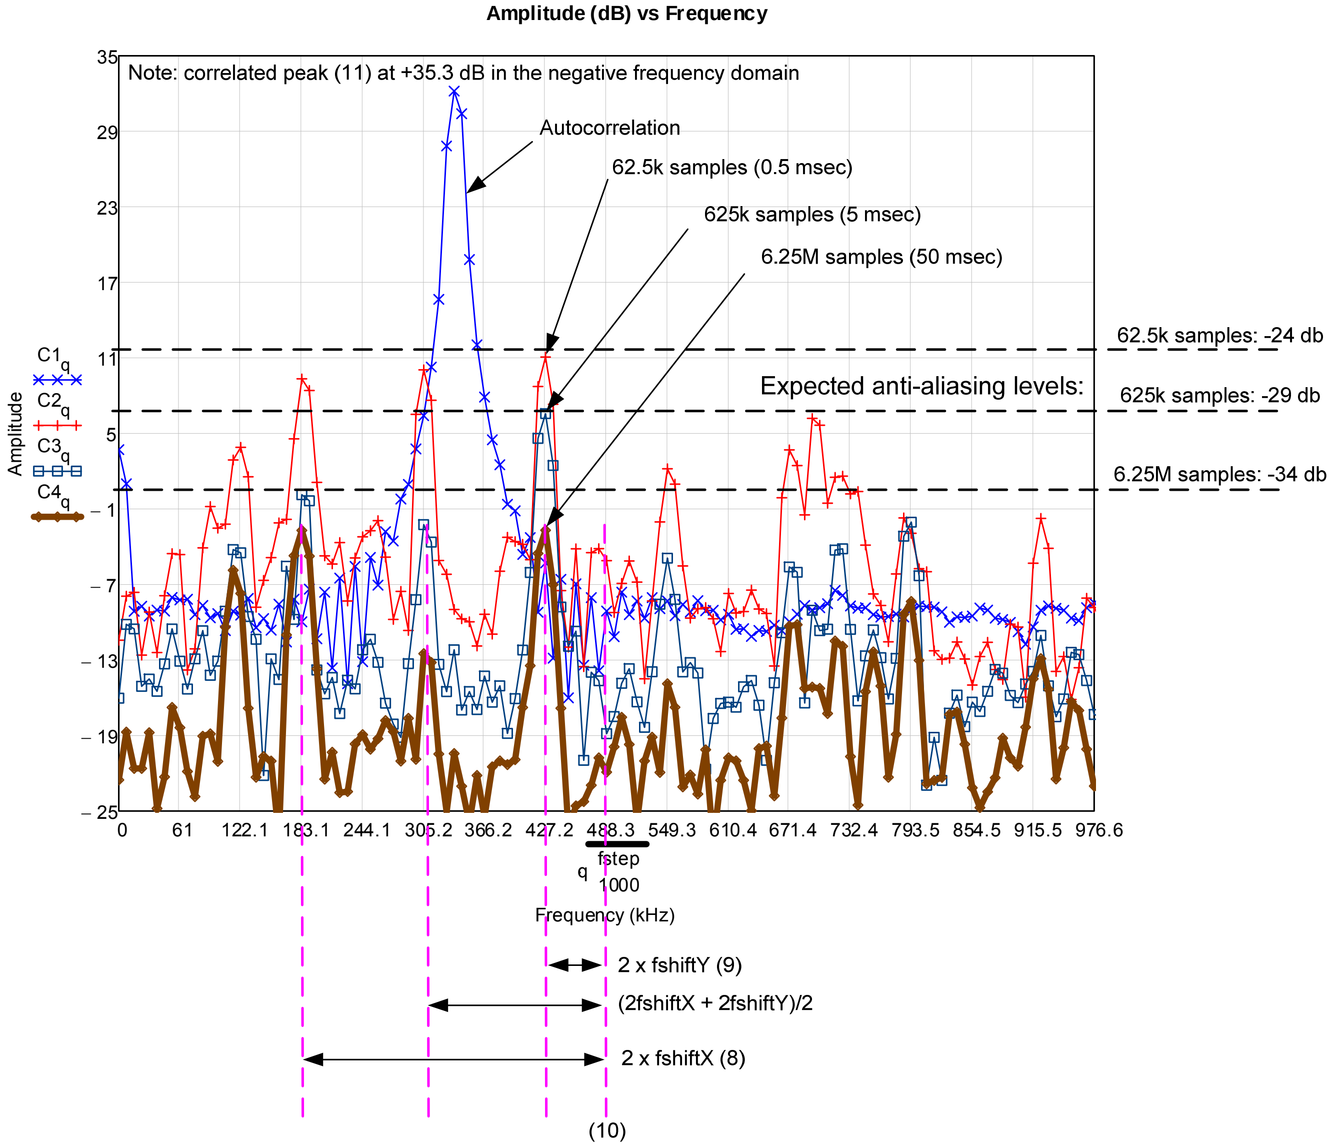

Figure 10 Test #1 anti-aliasing test results for 3 decades of "N" (samples) showing achieved and expected anti-aliasing. 3-level fringe rotation in the correlator.

- The figure clearly shows 1/√N anti-aliasing of roughly -5 dB per decade.
- The 6.25M sample (50 msec) correlation spectrum is highlighted in bold.

[3] If the tone is much stronger, some sampler-generated artifacts won't fundamentally decorrelate (as shown in detail in NRC-EVLA Memo# 009), but these are not evident in the auto-correlation or as amplitudes that are stable independent of integration time in the results.

The same test, with floating-point fringe rotators in the correlator is shown in Figure 11. Note the lack of most additional artifacts, but with roughly the same anti-aliasing results.

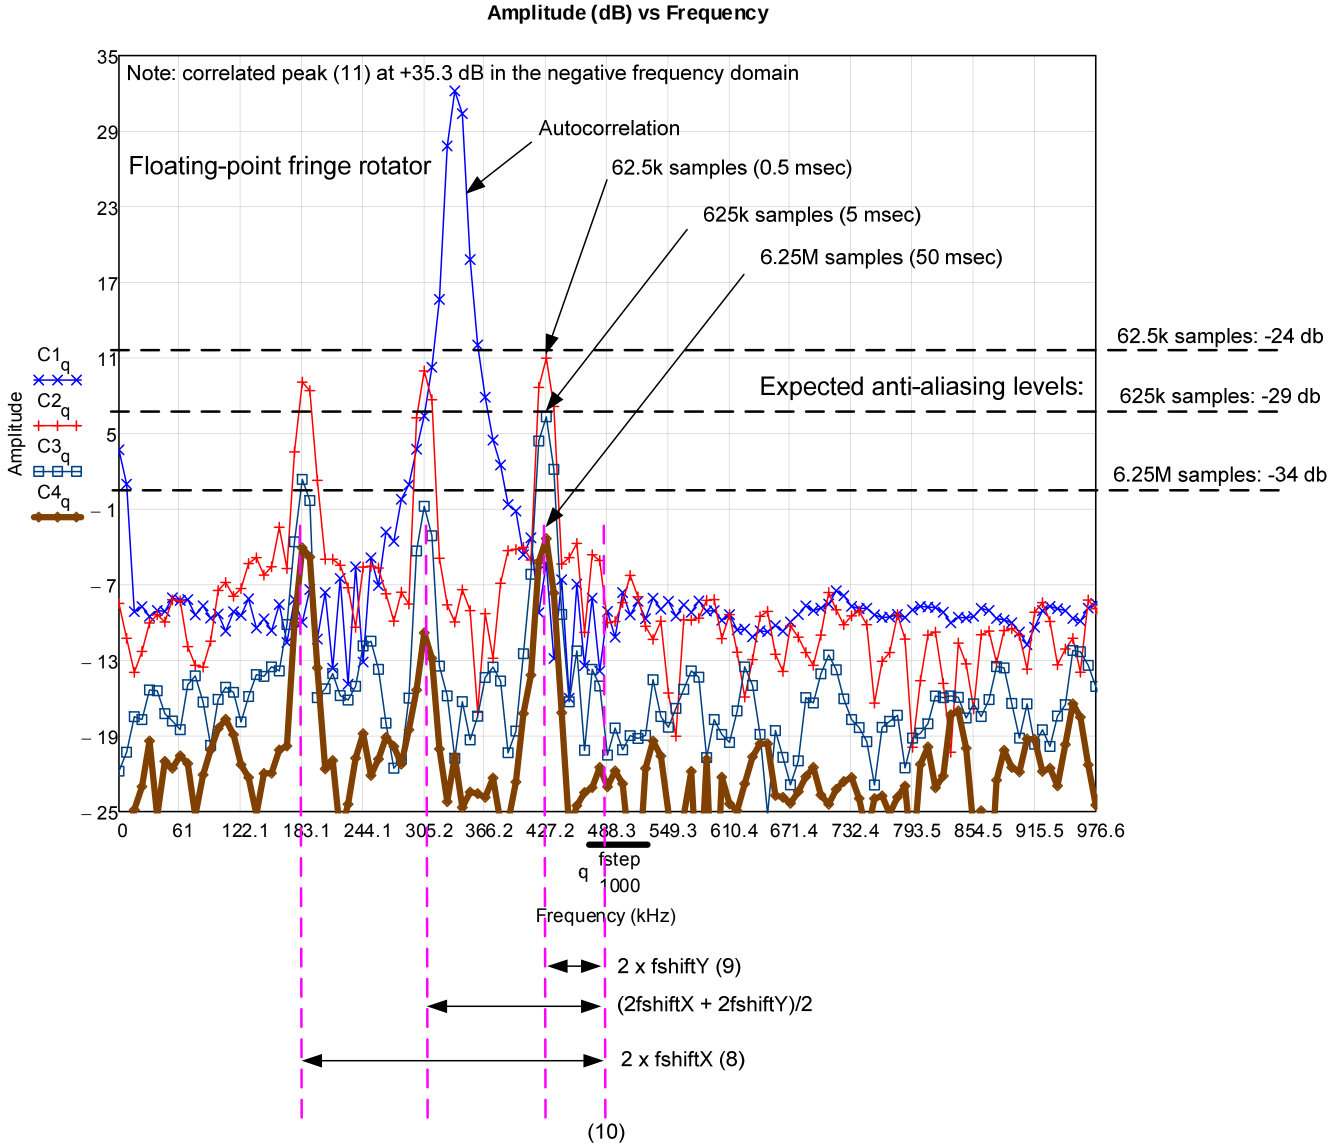

Figure 11 Test #1 anti-aliasing test results for 3 decades of "N" (samples) showing achieved and expected anti-aliasing. Floating-point fringe rotation in the correlator.

Figure 12 shows the first 128 channels of the negative frequency domain of sub-band 1 (i.e. 128 channels to the left of the above 2 figures), using a 3-level fringe rotator in the correlator. This is roughly what the top-end of sub-band 0 would look like.

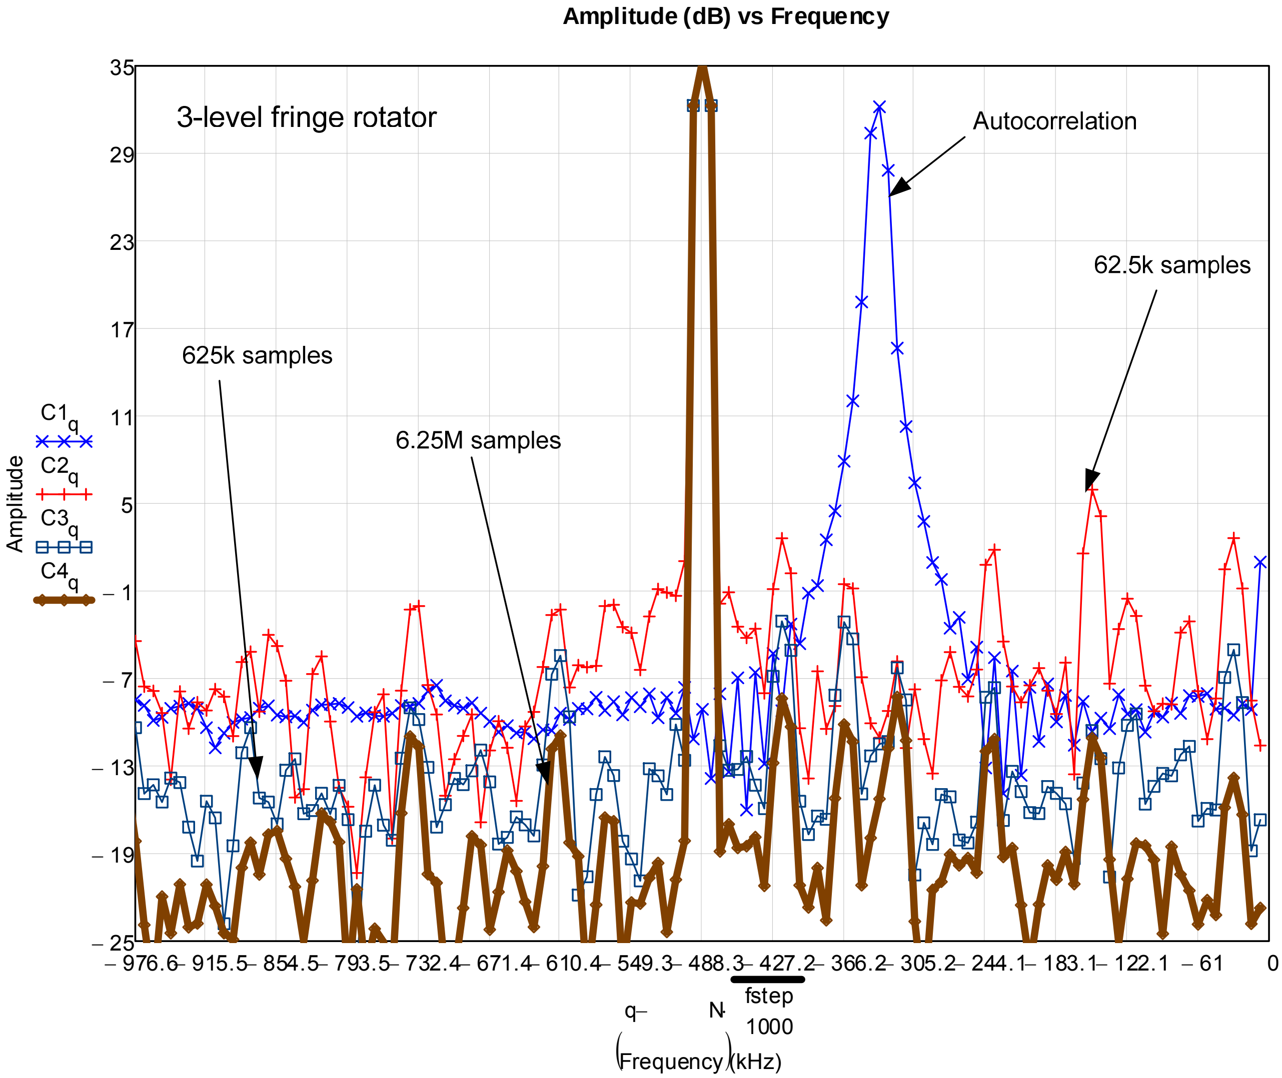

Figure 12 Test # 1. The first 128 channels of the negative frequency domain. 3-level fringe rotator in the correlator. The longest integration time spectrum is highlighted bold.

The bold spectrum in Figure 12 is for 6.25M samples of correlation, with artifacts at least 42 dB down from the tone peak, and indications that they all integrate down as 1/√N. The same spectrum with a floating point fringe rotator in the correlator is shown in Figure 13.

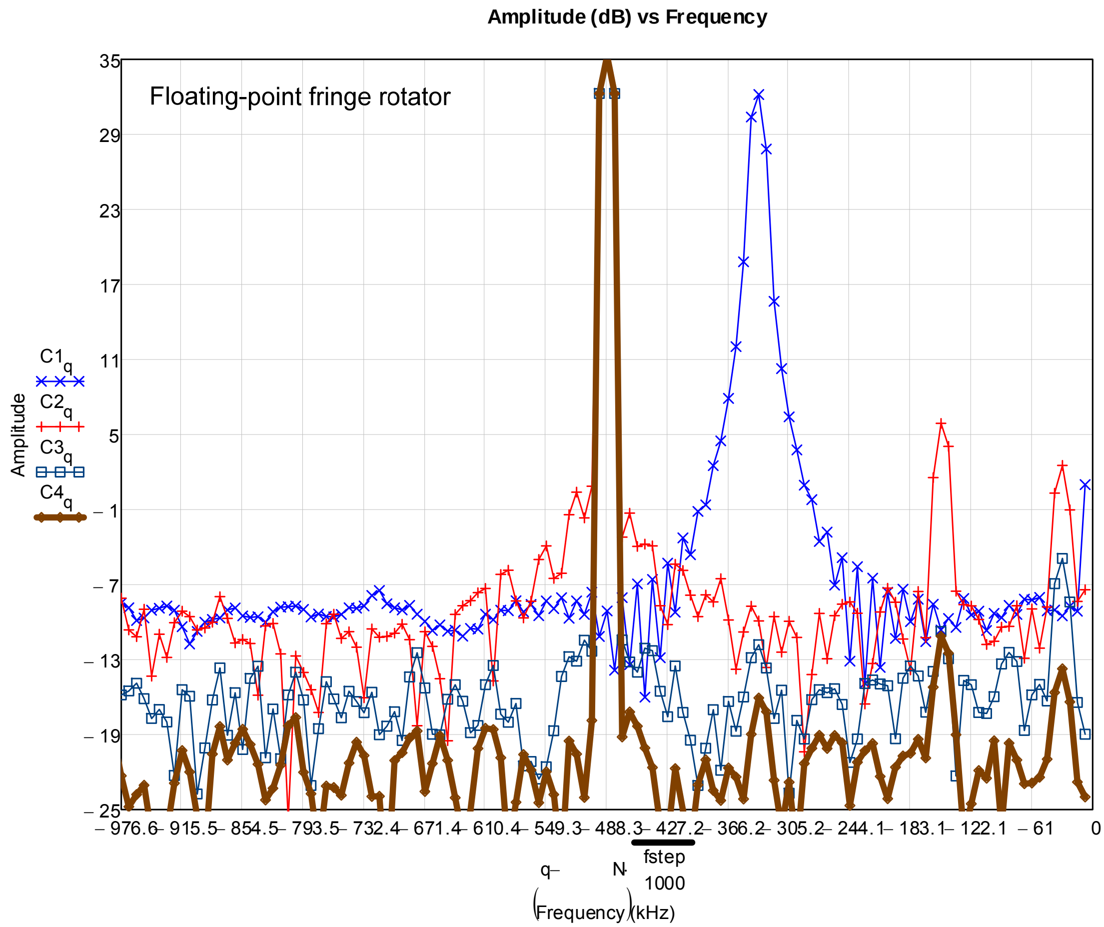

Figure 13 The first 128 channels of the negative frequency domain. Floating-point fringe rotator used in the correlator. There are still some artifacts, indicating they are likely quantizer-generated, but they integrate down as 1/√N. The longest integration time is highlighted bold.

Figure 14 is a continuum correlation, showing the first 256 channels of the negative frequency domain to show the transition-band slope of about 0.6 dB per ~240 kHz, indicating that using the correlated tone peak in the negative frequency domain as a reference in the anti-aliasing measurements is valid. For this test, fshift frequencies of $1/10^{th}$ of the above were used.

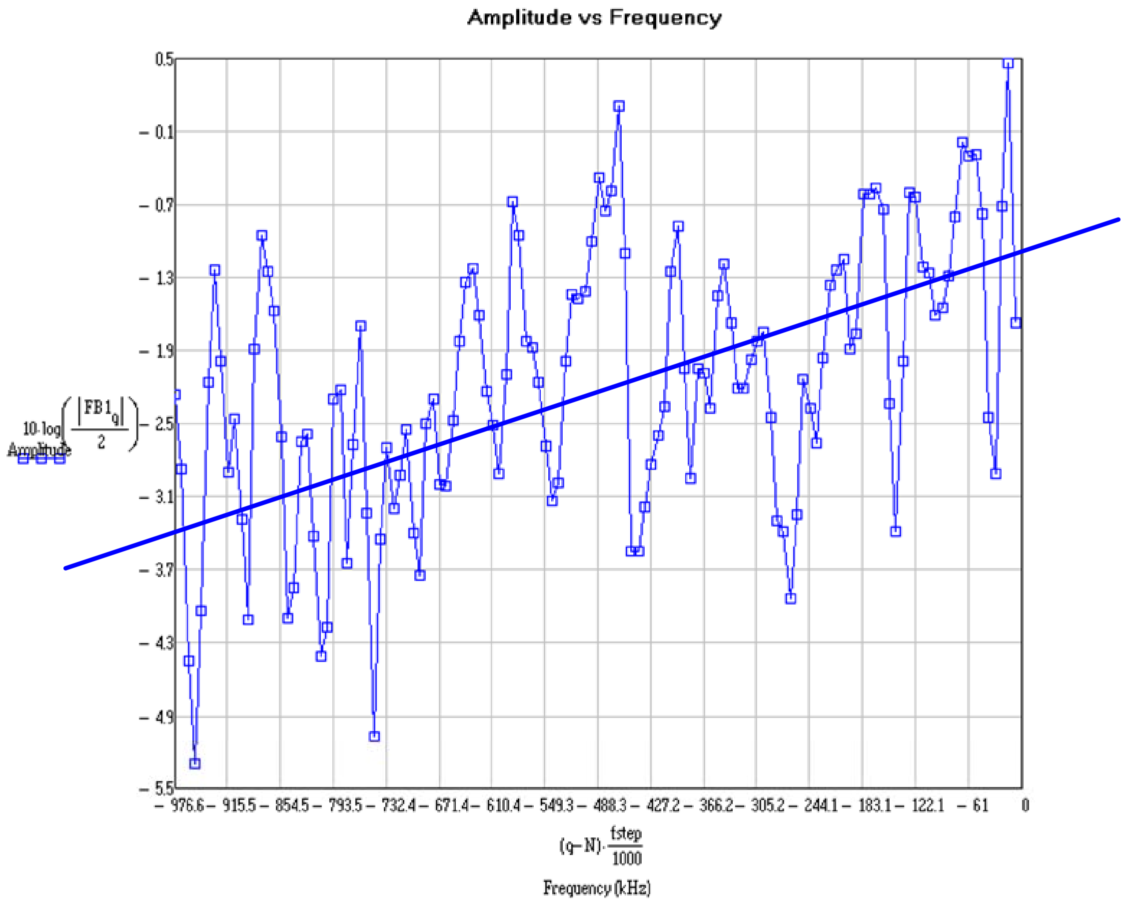

Figure 14 First 256 channels in the negative frequency domain, showing the transition-band slope of ~0.6 dB per ~240 kHz.

Test #2 Results

The test was run with modified fshifts of 1.511 kHz and 0.311 kHz so that the sinc(x) washing function dominated. The washing and de-correlation curves along with several results obtained with the software correlator are shown in the following figure:

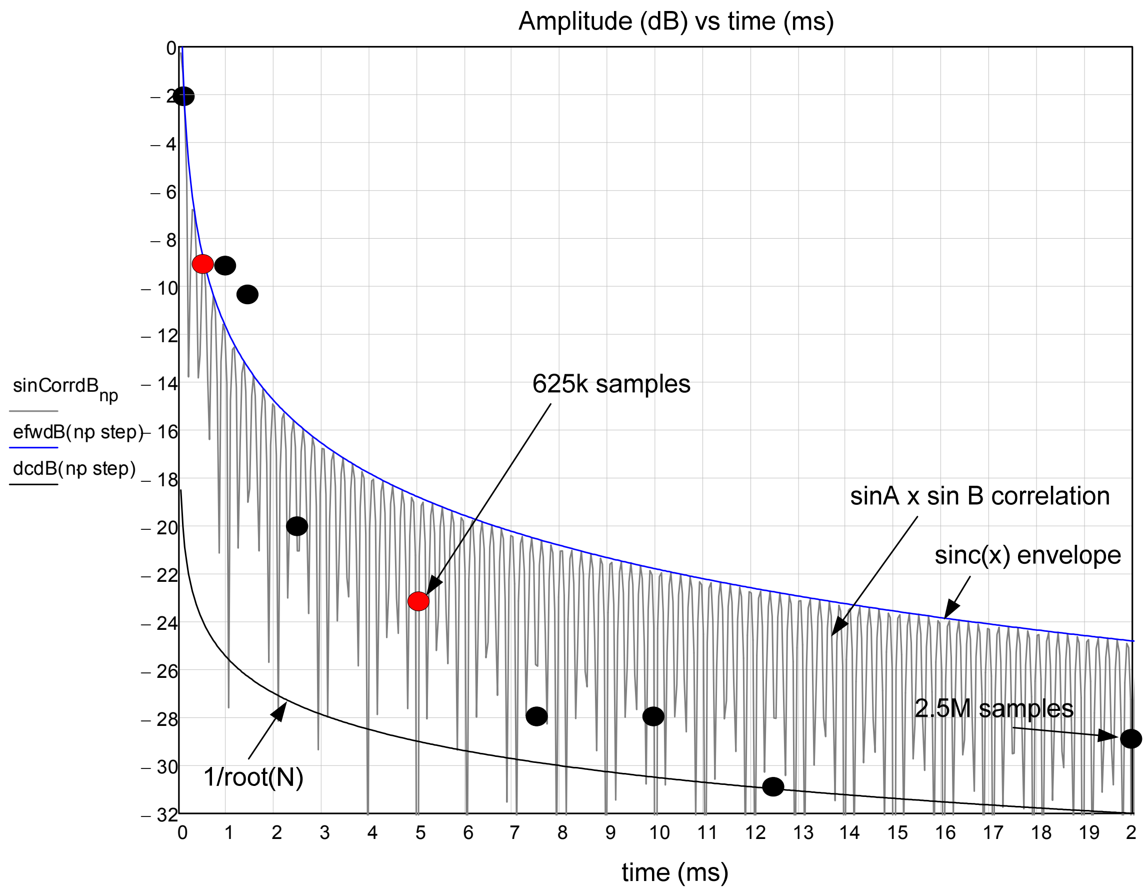

Figure 15 Test #2 sinc(x), 1/√N washing and de-correlation curves, and several correlation results. For this case, the sinc(x) function sets the anti-aliasing limit.

The actual "sinA x sinB" (i.e. sin[2π(f_{rest}-2fshiftX)·t] · sin[2π(f_{rest}-2fshiftY)·t]) calculated correlation (relative to $\sin^2(x)$ power=0.5) is shown in the figure. Note the large number of nulls that result when a tone with a relatively large f_{rest} frequency is offset by two relatively small frequencies and then correlated.

A second plot, with the last two points is shown in Figure 16:

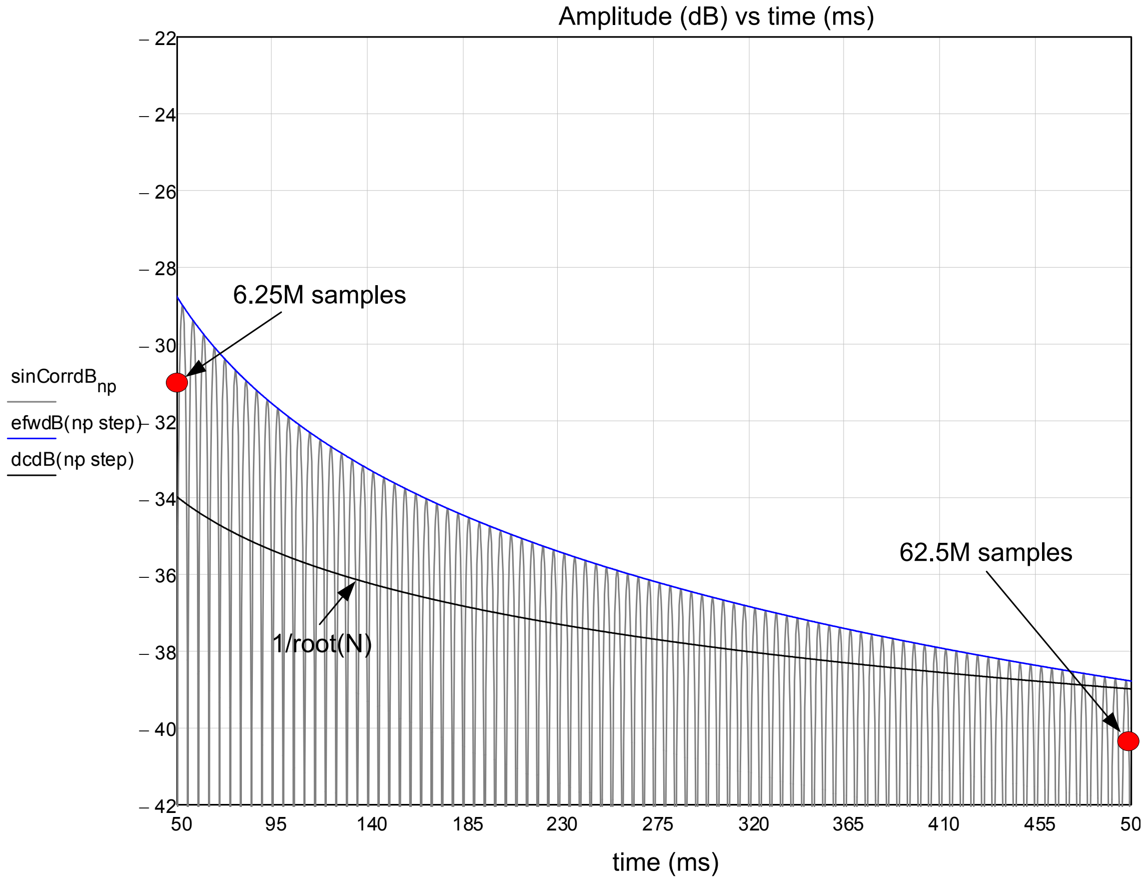

Figure 16 Test #2 final two data points, integrating down to ~-40 dB (the right-most time is 500 msec, not 50 msec—a Windows graphics copying error).

The lowest 128 channels of the amplitude vs frequency plot, with each spectral point ~61 kHz bandwidth (1024 channels across 62.5 MHz, Hanning windowing used) is shown in Figure 17, for 5 decades of correlation (corresponding to the RED points in Figures 15 and 16). The fshifts are so small relative to the spectral resolution obtained that there is only a single aliasing artifact seen.

Note in the plot that the noise is dropping at ~5 dB/decade, but the aliased artifact is dropping at ~10 dB per decade (also noticed by comparing the aliased peak to noise floor for the top curve and the bottom curve).

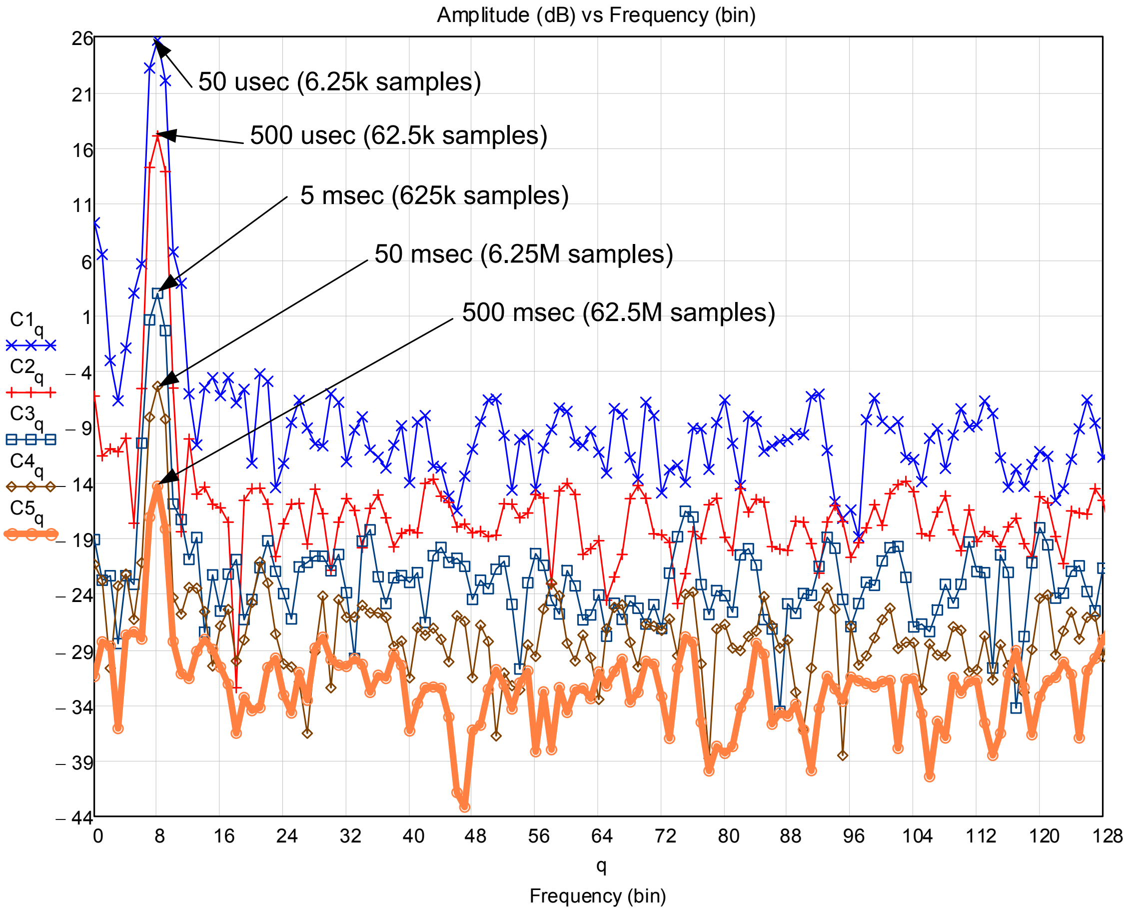

Figure 17 Test #2 amplitude vs frequency for 5 decades of correlation. The reference peak is at 26.2 dB. The 500 msec integration curve is highlighted bold.

Figure 18 is the same 5 data points, except looking at the 128 frequency channels just below DC. The correlated tone is the original un-aliased tone at the top of sub-band 0, but within the negative frequencies of sub-band 1 ((11) of Figure 8).

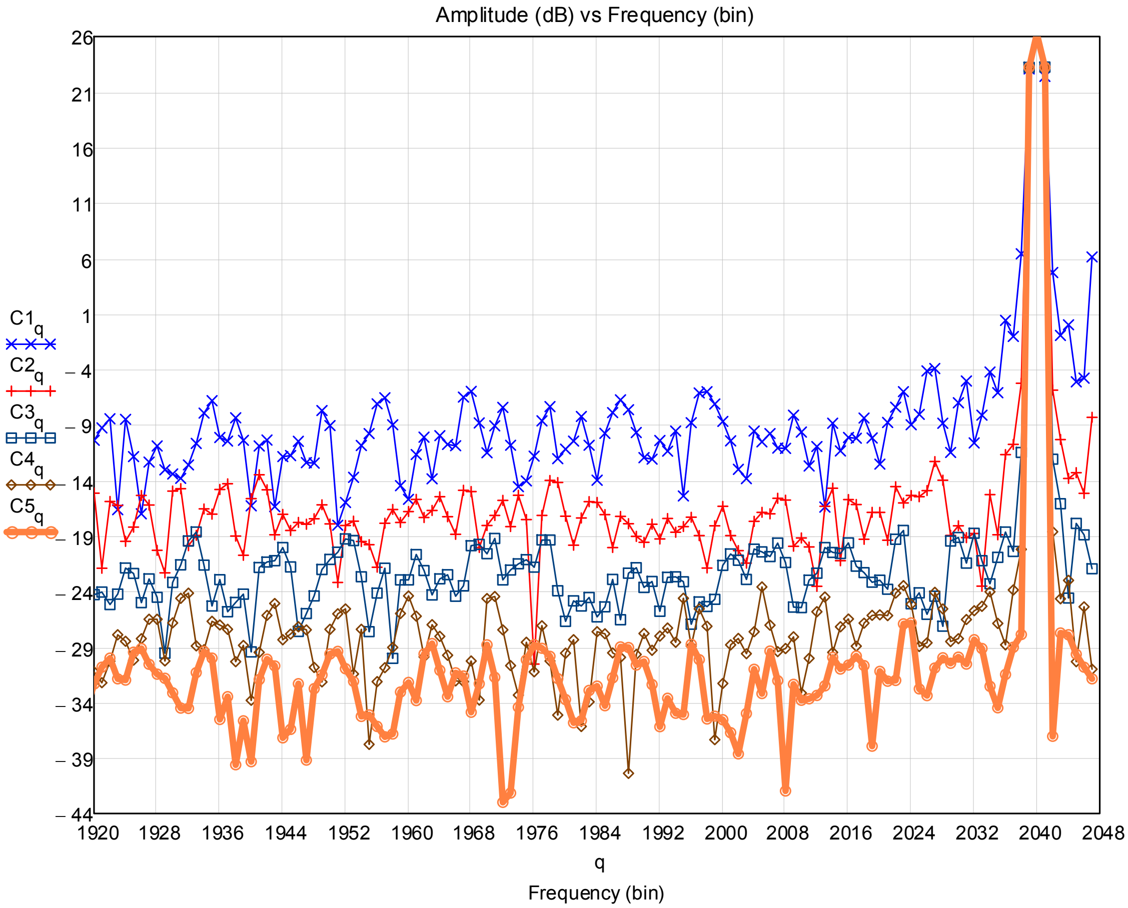

Figure 18 Test #2, 128 frequency channels just below DC. There are no apparent correlated artifacts around the tone, above ~-52 dB.

Test #3 Results

During hardware Delay Module Test Vector (DMTV) testing using simple fshift frequencies and fshift ratios, some unexpected behaviour in terms of anti-aliasing attenuation, and low-level artifacts around the tone were found. Thus, a series of tests were run to quantify these effects.

The conclusions that were found are as follows:

- Integer fshift ratios, whether even or odd, produce artifacts at low levels around a bin centered tone, and limit anti-aliasing at levels above what the sinc(x) washing function and 1/√N de-correlation curves predict. Odd ratios were the worst (refer to the following section), and confirm what Bob Sault saw during bin-centered on-the-sky testing. Even ratios produced artifacts ~10 dB lower than odd ratios. These artifacts appear to be at frequencies of 100's to 1000's of harmonics of fshift differences (i.e. fshiftX – fshiftY) away from the bin-centered tone. This may be a result of the interaction of the coarse mixer with the bin-centered tone.

- Non-integer fshift ratios, if the number, expressed in decimal form, has a repeating segment, results in similar behaviour, albeit at a lower level. For example for fshiftX=11 kHz, and fshiftY=6 kHz, the ratio calculates as 1.83333…, and artifacts appear. However, for fshift=11.1 kHz and 6.1 kHz, the ratio calculates as 1.81967213114754098…, and artifacts don't appear.

The following plots are examples of what has been found when testing a bin-centered tone with various fshifts and integration times. For these tests, the reference wideband sample rate is 2.048 GHz and $1/16^{th}$ sub-band filtering is performed with the tone originally at 127.0 MHz. The aliased tone shows up at 1 MHz into sub-band 2, and the unaliased, filter transition-band attenuated tone shows up at -1 MHz, and is used for aliasing attenuation comparisons.

Integer fshift ratios

Fshifts of 1 kHz, 500 Hz, and 10 kHz, 5 kHz were used in this test. The results are in the following plots.

The expected sinc(x) and 1/√N de-correlation curves are shown in Figure 19.

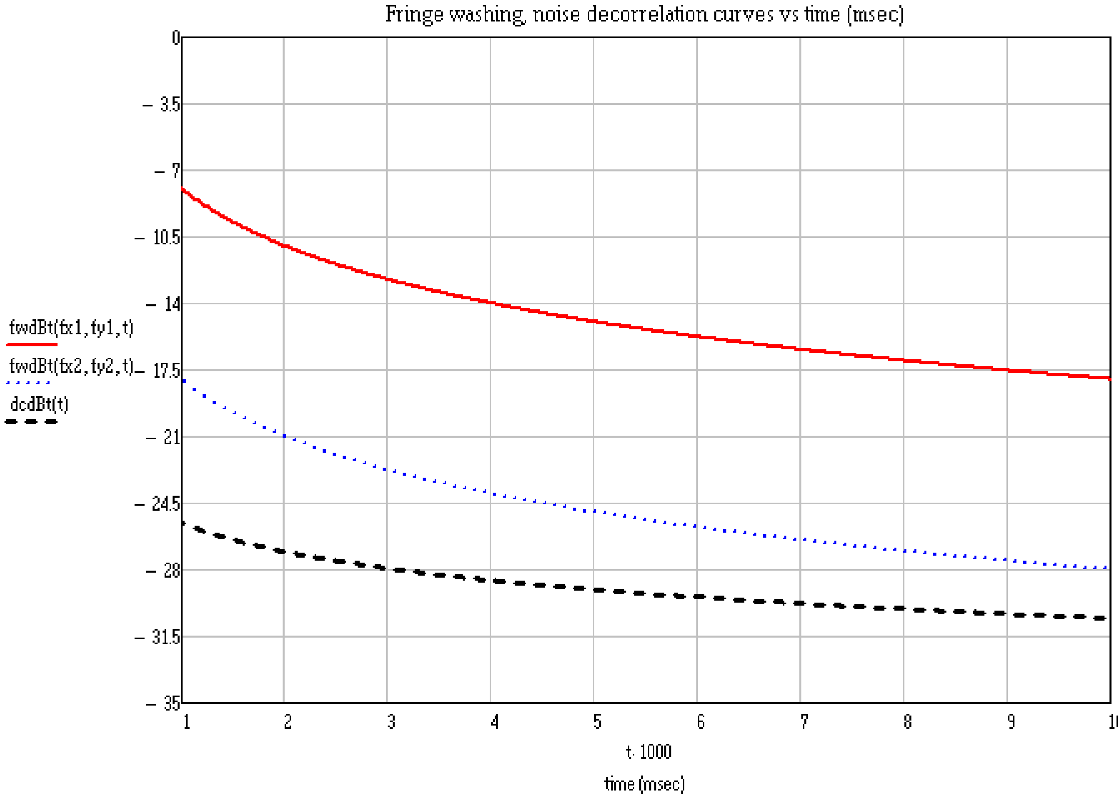

Figure 19 Expected sinc(x) washing and 1/√N de-correlation curves. The top (solid RED) curve is for fshiftX,Y=1 kHz, 500 Hz; the middle curve (dotted BLUE) is for fshiftX,Y=10 kHz, and the bottom (dashed BLACK) curve is the 1/√N noise curve. After 10 msec, the 500 Hz fshift difference predicts -17.5 dB attenuation, and the 5 kHz difference predicts -28 dB attenuation.

The correlation results for the aliased tone, showing the 64 frequency channels just above DC is shown in Figure 20.

The 5 kHz differential is worse than the 500 Hz. This effect was first seen when verifying hardware by loading up the Station Board with Delay Module Test Vectors, and later verified in software simulation.

Additionally, the correlated tone at -1 MHz (negative frequencies just below DC), shows "spreading" of the base of the tone, which gets worse with larger fshift difference. This is shown in Figure 21.

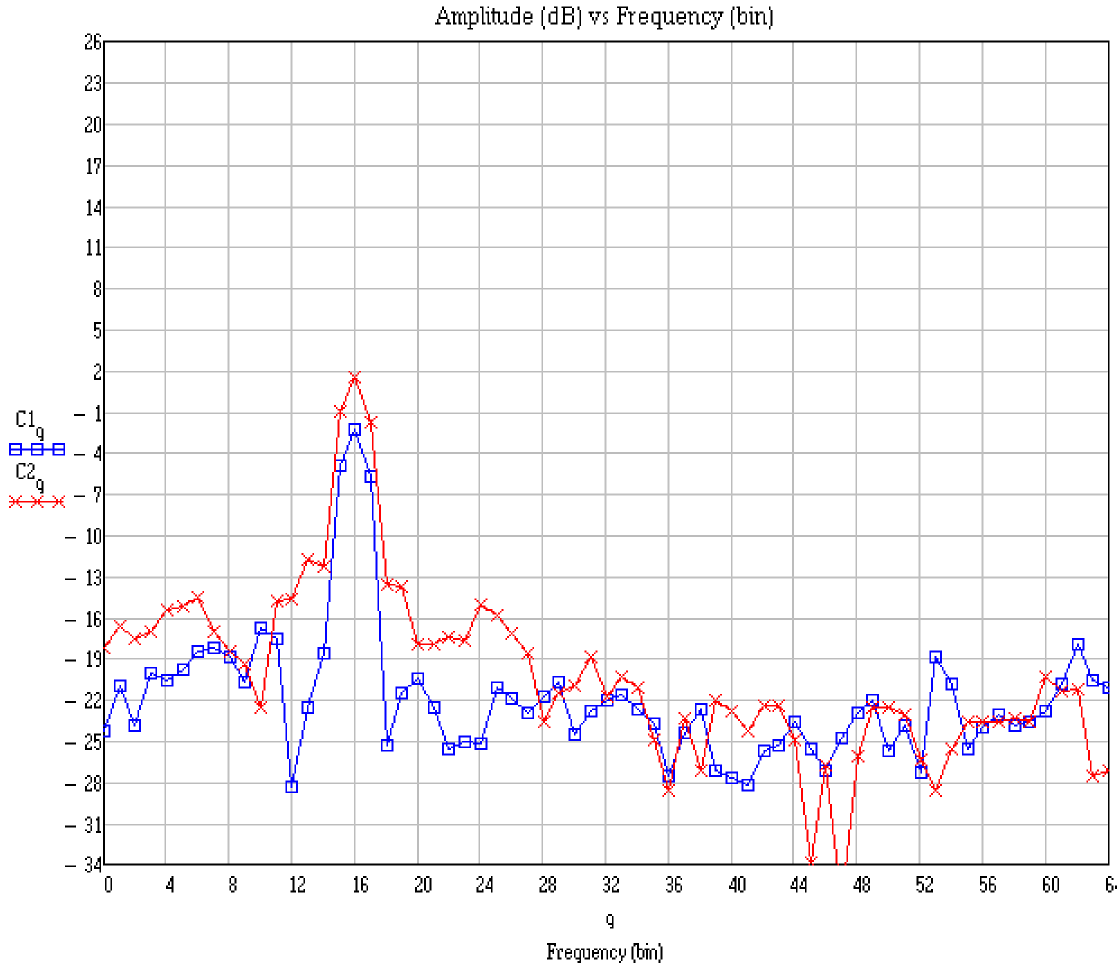

Figure 20 Aliased tone for fshift differential of 500 Hz (□ BLUE curve) and 5 kHz (x RED curve), 10 msec integration time. The 500 Hz differential has an attenuation of ~-28.5 dB, which is much better than predicted, due to it possibly being on a null of the sinc(x) function. However, the 5 kHz differential has an attenuation—worse at that—of ~-24 dB, and the prediction is -28 dB.

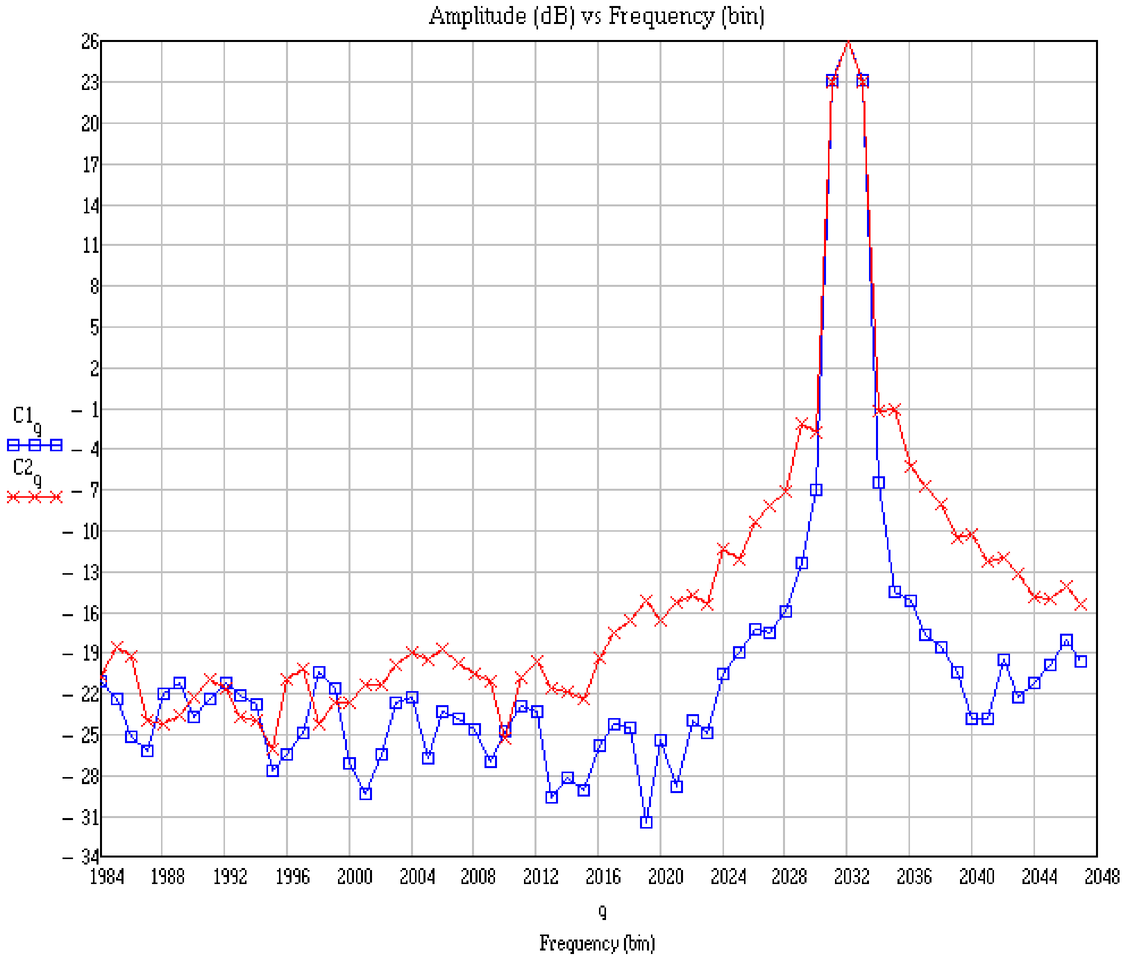

Figure 21 Correlated tone in the negative frequencies of sub-band 2. The spreading of the tone with fshift difference of 5 kHz (10 kHz, 5 kHz, x RED curve) is much worse than 500 Hz (□ BLUE curve).

Non-integer fshift ratios

The tests were re-run with fshift frequencies adjusted to be non-integer, and for an integration time of 250 msec.

The expected sinc(x) and 1/√N de-correlation curves are shown in Figure 22:

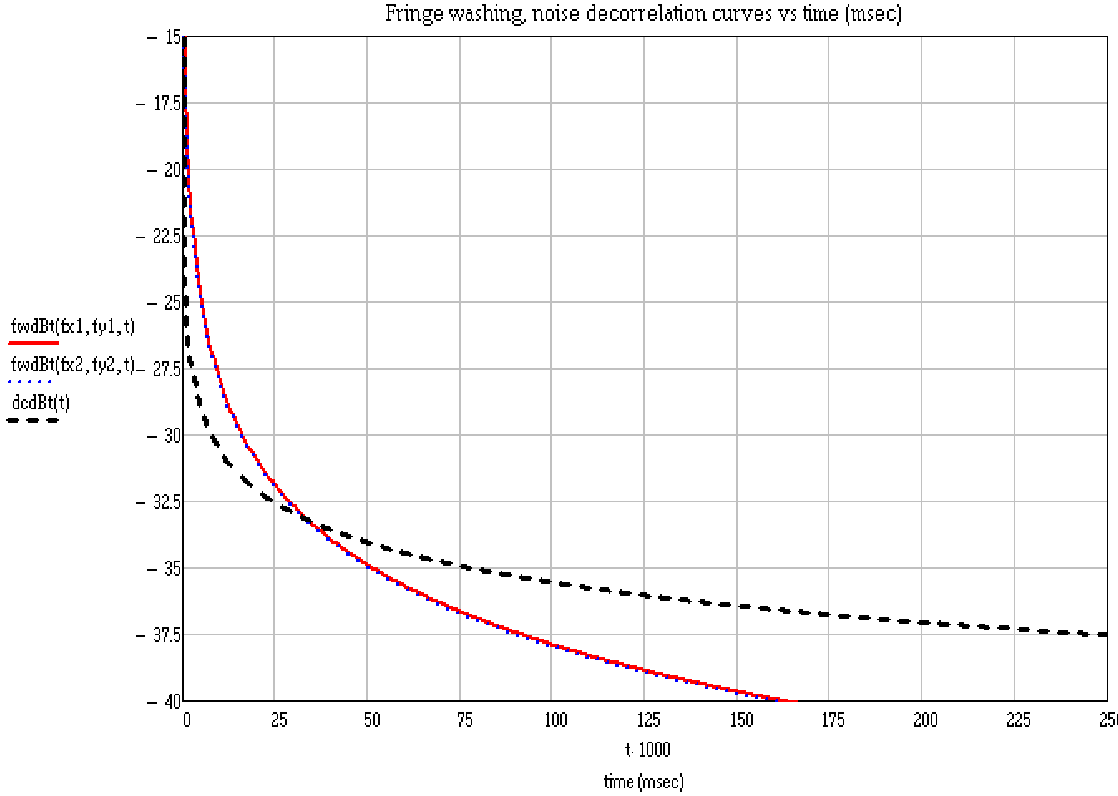

Figure 22 1/√N de-correlation (dotted BLACK curve), and sinc(x) washing curves for fshifts of 11 and 6 kHz (and 11.1 and 6.1 kHz). Noise de-correlation dominates at 250 msec integration time.

The aliased tone plot for the two cases is shown in Figure 23:

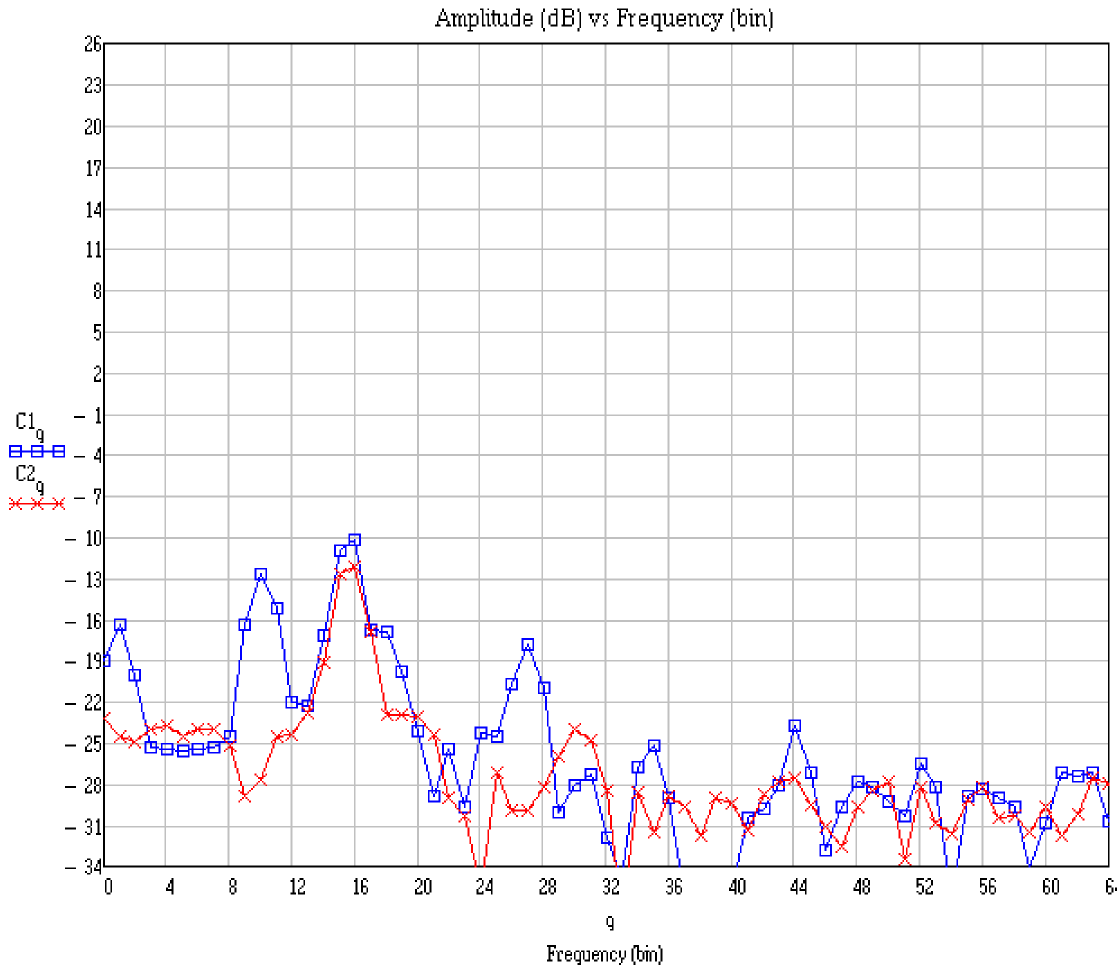

Figure 23 Aliased tone spectral plot for the 11 kHz, 6 kHz case (□ BLUE curve), and the 11.1 kHz, 6.1 kHz case (x RED curve), 250 msec integration time. The 11.1/6.1 case achieves the noise limit of ~-38 dB, whereas the 11/6 case achieves only -36 dB, with additional artifacts. Although, it is hard to tell without a longer integration if either case would continue to integrate down.

The correlated tone spectrum for this case is shown in Figure 24, clearly showing artifacts for the 11/6 kHz case compared to the 11.1/6.1 kHz case.

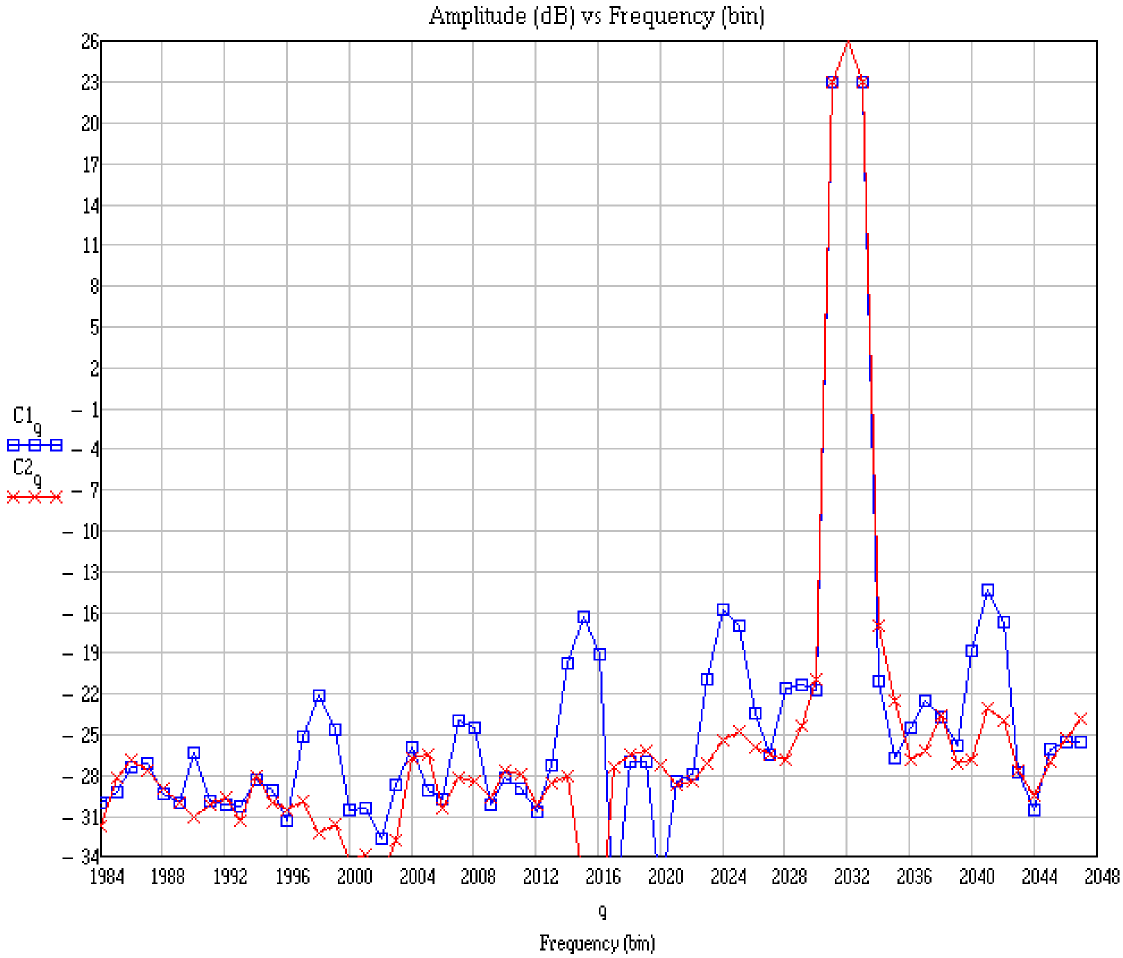

Figure 24 Correlated tone spectra for the 11/6 kHz case (□ BLUE curve), and the 11.1/6.1 kHz case (x RED curve) for a 250 msec integration time. In the 11/6 case, additional spectral artifacts at ~-40 dB level are at (what looks like) ~+/-100 and 200 times the fshift difference. Looking at Figure 17 and Figure 18 results though, indicate that better than -40 dB of aliasing attenuation and artifacts better than -52 dB are obtained, with a longer integration time.

More tests could be run, specifically looking at whether this is a side-effect of the combination of a very strong tone that is bin centered with an integer fshift ratio etc. Suffice to say that the evidence indicates that fshift ratios that are non-integer, and that when calculated result in "non-repeating+infinite[4]" decimal places is best advised. Further sky testing should be done, however, to integrate deeper and test this tentative conclusion.

Fshift algorithm

A simple fshift algorithm of:

$$fshift_n = n \cdot f1 + f_\varepsilon$$

[4] Or, at least out to ~> 6 decimal places.

where "n" is the antenna number, was investigated for f1 = 2.3 kHz, and f_ε = 70 Hz. A partial fshift ratio matrix is shown below:

Fr =

	2	3	4	5	6	7	8	9	10
0	0	0	0	0	0	0	0	0	0
1	0.507494647	0.340028694	0.25566343	0.204840104	0.170872386	0.146567718	0.128316188	0.114106885	0.102730819
2	1	0.670014347	0.50377562	0.403630078	0.336697909	0.288806432	0.252842447	0.224843524	0.202427395
3	1.492505353	1	0.75188781	0.602420052	0.502523432	0.431045145	0.377368706	0.335580164	0.302123971
4	1.985010707	1.329985653	1	0.801210026	0.668348955	0.573283859	0.501894965	0.446316803	0.401820546
5	2.47751606	1.659971306	1.24811219	1	0.834174477	0.715522573	0.626421224	0.557053442	0.501517122
6	2.970021413	1.989956958	1.49622438	1.198789974	1	0.857761286	0.750947482	0.667790082	0.601213697
7	3.462526767	2.319942611	1.74433657	1.397579948	1.165825523	1	0.875473741	0.778526721	0.700910273
8	3.95503212	2.649928264	1.992448759	1.596369922	1.331651045	1.142238714	1	0.889263361	0.800606849
9	4.447537473	2.979913917	2.240560949	1.795159896	1.497476568	1.284477427	1.124526259	1	0.900303424
10	4.940042827	3.30989957	2.488673139	1.99394987	1.663302091	1.426716141	1.249052518	1.110736639	1
11	5.43254818	3.639885222	2.736785329	2.192739844	1.829127614	1.568954855	1.373578776	1.221473279	1.099696576
12	5.925053533	3.969870875	2.984897519	2.391529818	1.994953136	1.711193568	1.498105035	1.332209918	1.199393151
13	6.417558887	4.299856528	3.233009709	2.590319793	2.160778659	1.853432282	1.622631294	1.442946558	1.299089727
14	6.91006424	4.629842181	3.481121899	2.789109767	2.326604182	1.995670996	1.747157553	1.553683197	1.398786303
15	7.402569593	4.959827834	3.729234088	2.987899741	2.492429704	2.137909709	1.871683812	1.664419836	1.498482878
16	7.895074946	5.289813486	3.977346278	3.186689715	2.658255227	2.280148423	1.99621007	1.775156476	1.598179454
17	8.3875803	5.619799139	4.225458468	3.385479689	2.82408075	2.422387137	2.120736329	1.885893115	1.697876029
18	8.880085653	5.949784792	4.473570658	3.584269663	2.989906273	2.56462585	2.245262588	1.996629754	1.797572605
19	9.372591006	6.279770445	4.721682848	3.783059637	3.155731795	2.706864564	2.369788847	2.107366394	1.897269181
20	9.86509636	6.609756098	4.969795038	3.981849611	3.321557318	2.849103278	2.494315106	2.218103033	1.996965756
21	10.357601713	6.93974175	5.217907228	4.180639585	3.487382841	2.991341991	2.618841364	2.328839673	2.096662332
22	10.850107066	7.269727403	5.466019417	4.379429559	3.653208363	3.133580705	2.743367623	2.439576312	2.196358908
23	11.34261242	7.599713056	5.714131607	4.578219533	3.819033886	3.275819419	2.867893882	2.550312951	2.296055483
24	11.835117773	7.929698709	5.962243797	4.777009507	3.984859409	3.418058132	2.992420141	2.661049591	2.395752059
25	12.327623126	8.259684362	6.210355987	4.975799481	4.150684932	3.560296846	3.1169464	2.77178623	2.495448635
26	12.82012848	8.589670014	6.458468177	5.174589455	4.316510454	3.70253556	3.241472658	2.88252287	2.59514521
27	13.312633833	8.919655667	6.706580367	5.37337943	4.482335977	3.844774273	3.365998917	2.993259509	...
28									

Figure 25 Partial table showing fshift ratios for a simple n · f1 + f_ε algorithm yields, by inspection, non-repeating decimal numbers for f1 = 2.3 kHz, and f_ε = 70 Hz. Many other parameters are likely possible and I can imagine a more sophisticated algorithm being developed to optimally choose these numbers for maximum anti-aliasing attenuation and minimum bandwidth loss, bearing in mind the inchoherent (averaging) time, coherent (FFT) integration time, and sub-band sample rate.

Non-bin centered tone tests

A sub-test was run to determine if artifacts and anti-aliasing with integer fshift ratios also applied to a non-bin-centered tone. In this case, the strong tone is "repetitively sampled" at different points along its sine curve as the sample rate has no relation to the tone frequency. It was found that indeed, anti-aliasing functioned as expected (i.e. worked) with an fshiftX of 10 kHz and an fshiftY of 5 kHz. Figure 26 shows the drop in aliased tone amplitude versus integration time for integration times up to 250 msec. The 1/√N and sinc(x) de-correlation curves for this case are the same as Figure 22.

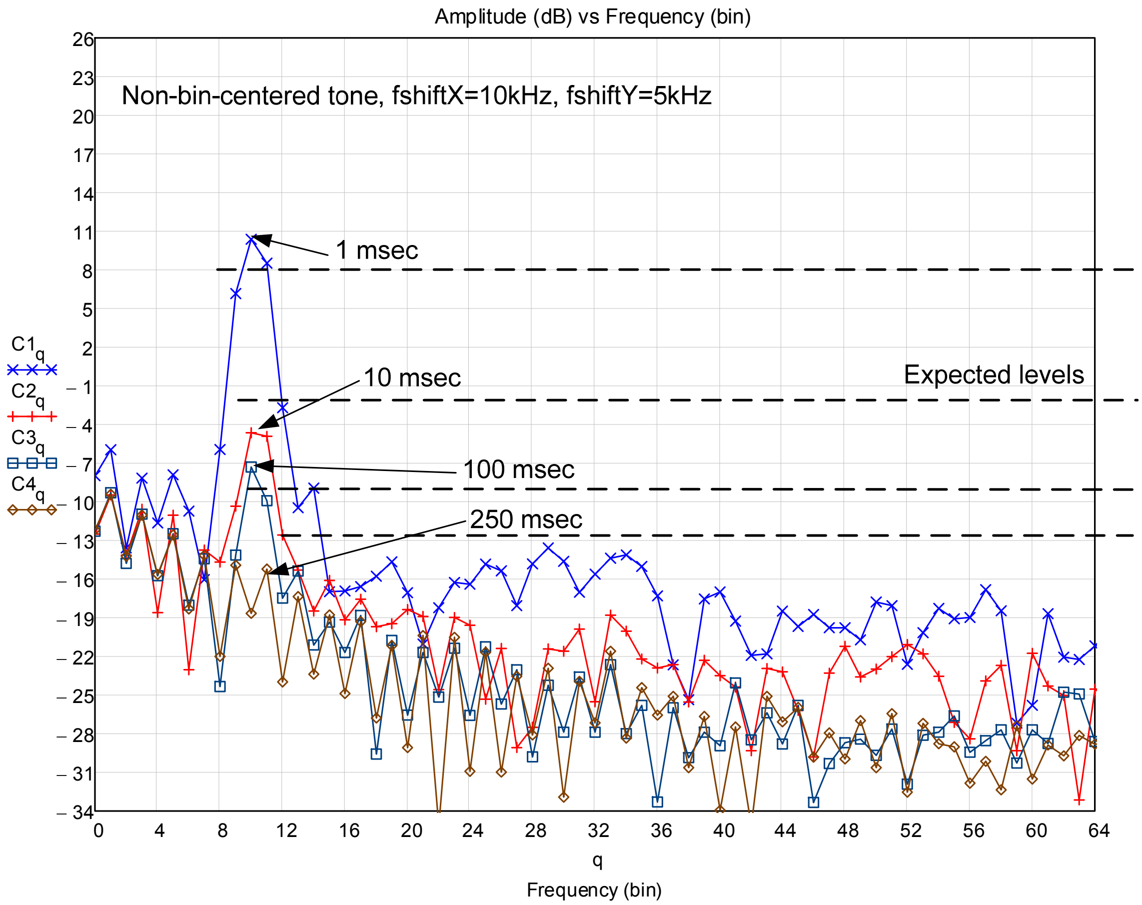

Figure 26 Aliased tone amplitude vs frequency for 4 integration times showing expected anti-aliasing levels, with fshiftX=10 kHz and fshiftY=5 kHz, for a non-bin-centered tone. The ringing in the signal amplitude is due to ringing from Hanning windowing originating from the correlated tone in the adjacent negative frequencies.

For this case it was not possible to determine if artifacts (spreading) around the base of the correlated tone were present, due to sidelobes from the non-bin-centered tone (i.e. the sidelobes are about the same size as the artifacts).

For an odd ratio of 3:1, similar anti-aliasing results were obtained. Figure 27 shows a 1 msec and 10 msec integration time with fshiftX=15 kHz, and fshiftY=5 kHz.

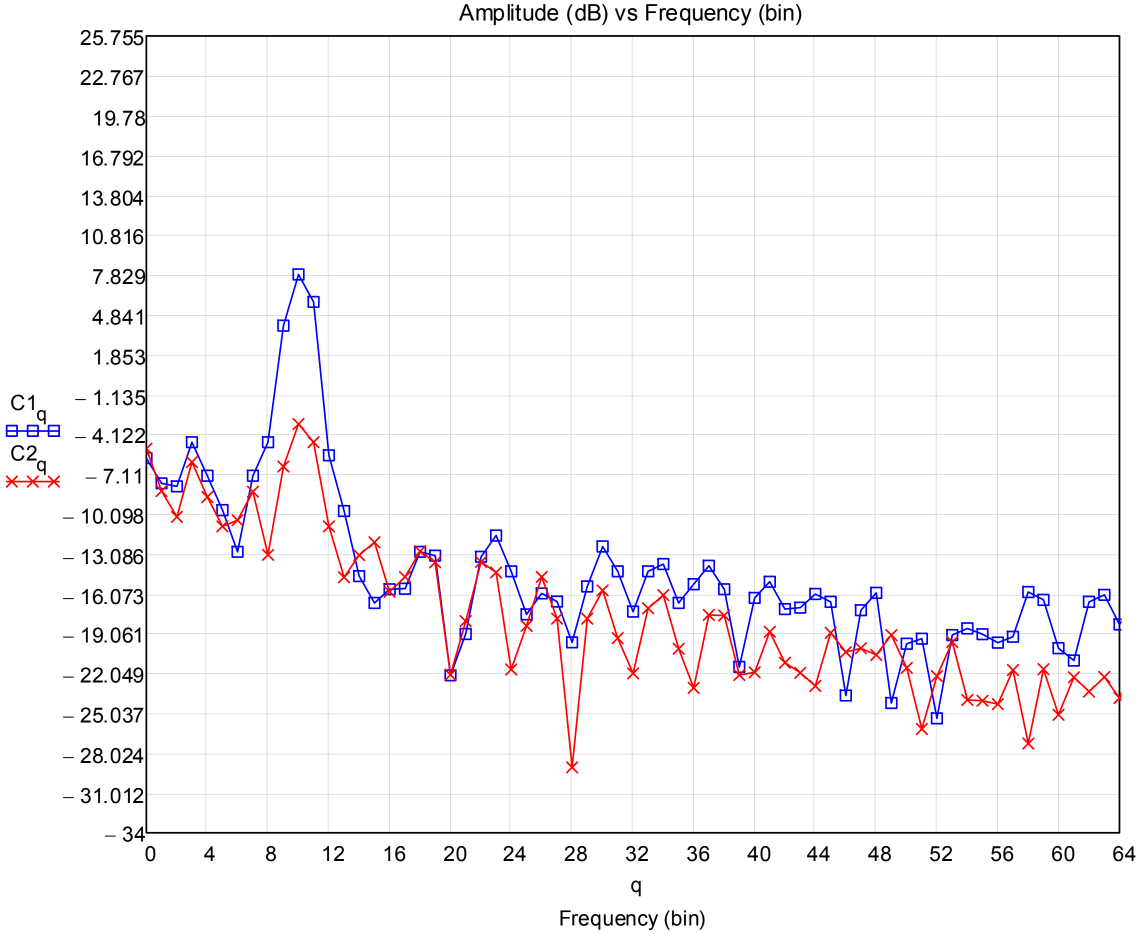

Figure 27 Aliased tone de-correlation for fshiftX=15 kHz, fshiftY=10 kHz, for 1 msec (□ BLUE curve), and 10 msec (x RED curve) integration times. The expected level for the top curve is ~6 dB, and for the bottom curve is ~-5 dB.

For this test, the differential fshift was 10 kHz, double the 10k / 5k test, and this allowed us to see that spreading of the base of the tone does indeed occur. This spreading is shown in Figure 28:

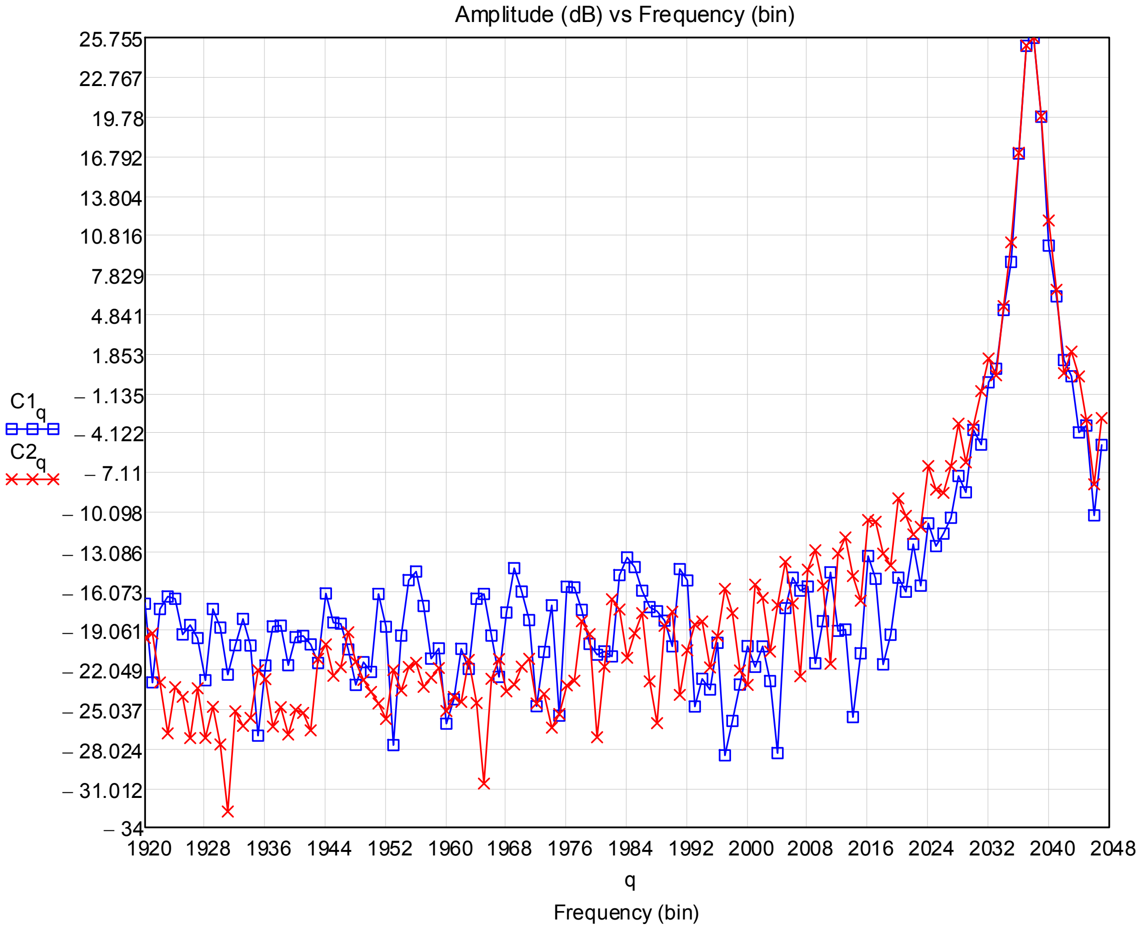

Figure 28 Correlated non-bin-centered tone in the negative frequencies of sub-band 2 for fshiftX=10 kHz, fshiftY=5 kHz (□ BLUE curve), and fshiftX=15 kHz, fshiftY=5 kHz (x RED curve) for a 10 msec integration time. The spreading of the RED curve over and above Hanning window sidelobe spreading is clear, indicating that false artifacts are present.

Further tests could be run to determine if spreading occurs only for integer fshift ratios, or extends to “repeating decimal” ratios as determined for bin-centered tones.

Tentative Conclusions

1. Anti-aliasing is primarily determined by two curves; the sinc(x) washing curve, which depends on the fshift differential, and by the 1/√N noise de-correlation curve, which purely depends on integration time and sample rate.

2. Depending on fshift frequencies, either curve may dominate and so expected anti-aliasing levels must take this into account. Fundamentally, the 1/√N curve forms a limit, and so it is never possible for anti-aliasing to hit complete "nulls" in the sinc(x) curve.

3. For bin-centered tones (and for a real system this means tones that are frequency bin centered AND frequency locked to the observing system timing reference), with integer or "repeating+infinite decimal place" fshift ratios, there are false spectral artifacts produced around the base of the tone. Additionally, for these fshift ratios, anti-aliasing does not produce the expected attenuation—it reaches a limit which could be well above the sinc(x) or 1/√N curves, depending on fshift differential, and integration time.

4. For non-bin-centered tones, there does not appear to be an effect on anti-aliasing depending on fshift ratio. However, there do appear to be artifacts that show up at low levels around the base of the tone for integer or repeating decimal place fshift ratios.

For maximum signal integrity non-integer and non-repeating+infinite decimal fshift ratios should be used. A simple fshift calculation algorithm is likely all that is necessary, with some care in choice of parameters, to achieve the desired result.

Delay Module Test Vector Test

Tests were run to verify that the results obtained in the software correlation investigation were also obtained in real correlator hardware.

Each Station Board has two Delay Modules, and each Delay Module has a large enough memory capacity to hold 250 msec worth of pre-filtered data. The Delay Module is loaded with computer-generated data samples (vectors) and is then run in “Delay Module Test Vector” mode (DMTV). In this mode, data is output continuously but repeats every 250 msec. Delay Module output data goes to the downstream digital filters and the correlator just like it would normally do if connected to an antenna.

A test was run with a tone near the top of sub-band 0, at 127 MHz, 1 MHz below the edge of the band, identical to the setup for the previously described “Test #3”. For this test, fshiftX=11.1 kHz, and fshiftY=6.1 kHz, with the tone amplitude similar in magnitude as for the software-only test.

The expected aliasing attenuation curves are as follows:

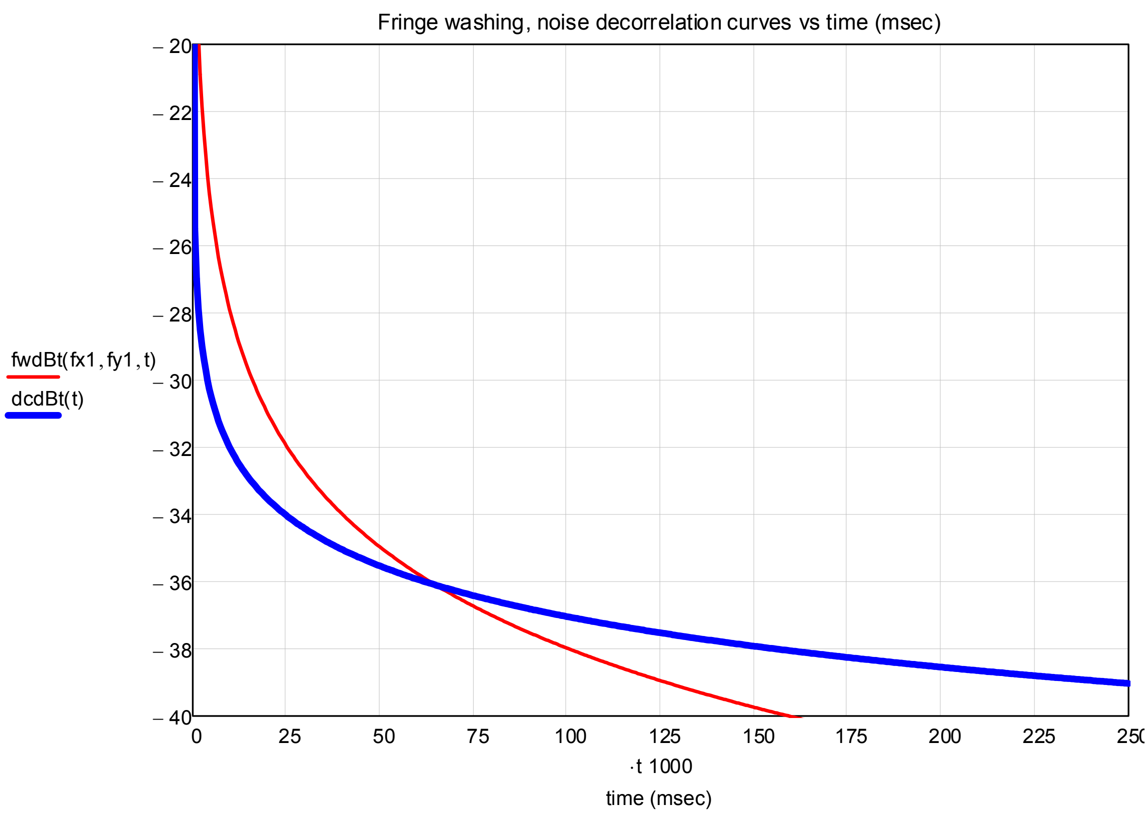

Figure 29 Expected sinc(x) washing (RED) and 1/√N noise de-correlation (bold BLUE) curves for fshiftX=11.1 kHz, fshiftY=6.1 kHz, and 128 MHz bandwidth.

The aliased component spectrum is shown in Figure 30 for 10 msec and 250 msec integration times:

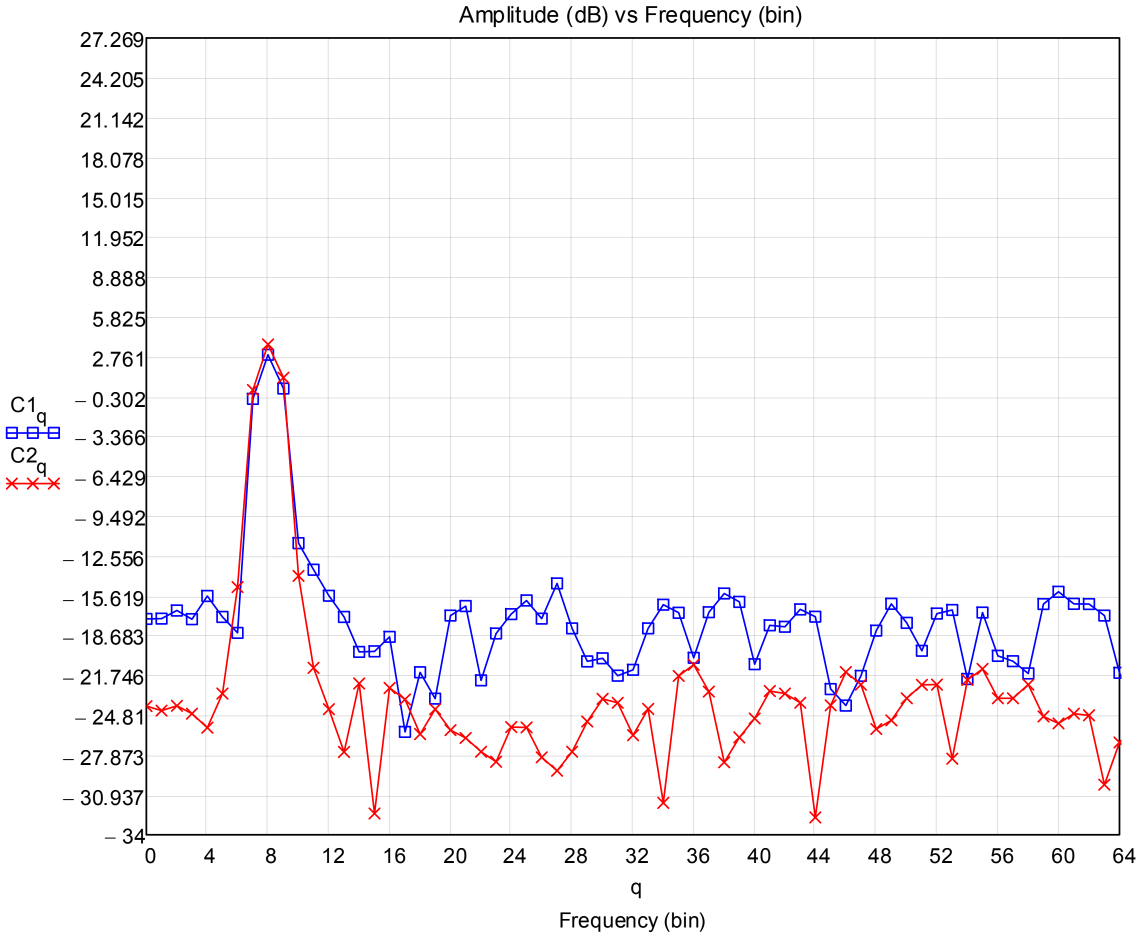

Figure 30 Aliased signal spectra obtained from hardware filtering and correlation of 250 msec of Delay Module Test Vectors (DMTV) for 10 msec (□ BLUE) and 250 msec (x RED) integration times. This is for fshiftX=11.1 kHz, and fshiftY=6.1 kHz. The peak of the correlated negative frequency tone is at 27.3 dB, thus achieving ~25 dB attenuation, far short of the expected -29 dB and -39 dB respectively.

Clearly in the above figure, anti-aliasing hits a limit at ~-25 dB and does not achieve the expected -29 dB and -39 dB levels for 10 msec and 250 msec integration times. -25 dB is the expected level for just 5 msec of integration time.

This result is, to say the least, completely unexpected, given the investigations and results obtained in the previous section. As a result a test was run where a copy of the test vectors loaded into the Delay Module was run through the software correlator. The copy is identical, except for format: the Delay Module hardware works on 8-bit offset binary and the s/w correlator works on 8-bit 2's complement. A simple conversion between the two formats was performed. Both the hardware and software correlators implement the FIR filter mathematically (8-bit input, 12-but LUT output, integer adder tree) the same except for different bandpass widths and number of taps (the h/w correlator has 256 taps, $1/8^{th}$ sub-band, and the software has 1023 taps, $1/16^{th}$ sub-band). The correlation part is also the same, 4-bit re-quantization, lag-based differential phase, 3-level phase rotation.

The software correlator was set as in the previous sections, for 1/16th sub-band bandwidth (64 MHz). This is ½ of the bandwidth of the hardware setup and so the aliased tone is as expected 1 MHz from the bottom of sub-band 2 in the software correlator (rather than sub-band 1 in the hardware correlator). The noise limit curve in the software correlator is higher as only ½ the number of samples are correlated (setting an expected anti-aliasing attenuation of ~-37.5 dB at 250 msec, rather than -39 dB).

The results of software correlation on identical vectors as produced the hardware result of Figure 30 are shown in Figure 31:

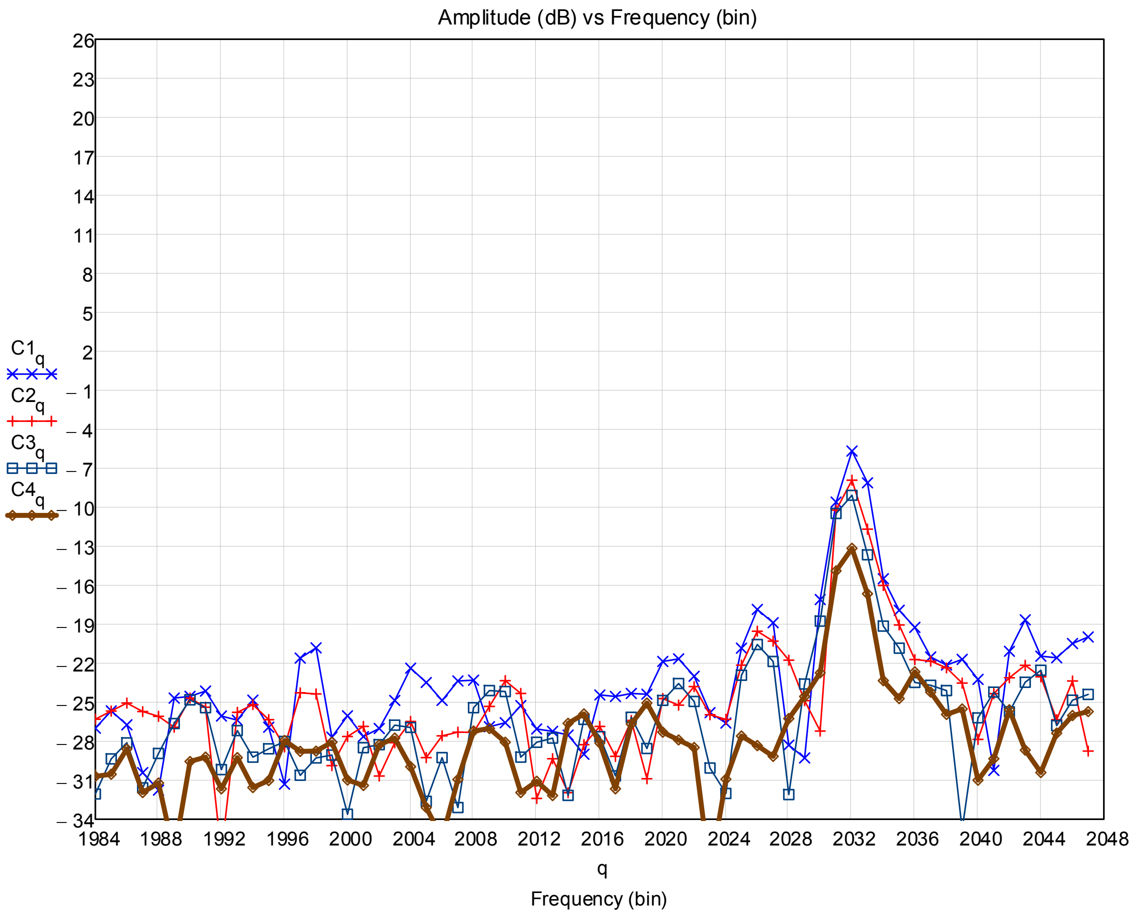

Figure 31 Aliased signal spectra from software correlation of Delay Module test vectors used in hardware correlation. The curves are from top to bottom, for 25 msec, 50 msec, 100 msec, and 250 msec integration times. The peak of the tone in the negative frequencies just below DC is at +26.01 dB. This is for fshiftX=11.1 kHz and fshiftY=6.1 kHz. The frequency sense is flipped compared to software-only and hardware-only results, due to a mismatch in the fshift sign convention used in the different signal generator implementations. -39 dB of attenuation is achieved at 250 msec integration time, slightly better than the expected ~-37.5 dB.

Further software correlation testing was performed, using the identical FIR coefficients (and therefore filter shape and decimation factor), and the same result was obtained.

Clearly, the software correlator achieves the expected result, and it is evident that the aliased signal amplitude drops with integration time. The fact that the hardware correlator does not achieve the same result, given ostensibly identical inputs, indicates that something, at a very low level, is possibly amiss in the hardware signal processing chain.

It is not yet known if this is due to loading of test vectors, operation of the Delay Module in test vector mode causing, or something downstream in the normal signal processing chain. Note that these tests, h/w and s/w, were run with no delay model, and with constant fshift frequencies that return to zero phase every 10 msec (so the hardware PHASEMODels are always the same). This test also independently checks that the fshift frequencies are set as indicated because the s/w correlator program and user (Brent) is independent from the signal generator and user (Dave).

Problem Found

After extensive investigation, a cause for the faulty anti-aliasing behaviour observed was found in hardware.

It turns out that the correlator chip [2] dump time, T_{dump} coupled with the fact that there is a 1 micro-second blanking time[5] for each dump, beats with the fshift differential (and likely fshift frequency sum) in a kind of "inverse" strobing fashion to produce correlation at ½ the differential frequency (and its odd harmonics if 3-level phase rotation is turned on). Indeed, it was found that the correlation of two independent tones, frequency locked to system timing, let alone aliasing, produced a correlation where there should have been no correlation!

Figure 32 shows two hardware correlations. The BLUE (□ points) trace is with fshift1=131.1 kHz and fshift2=351.1 kHz, with a 2fshift difference of 440 kHz. The RED (x bold points) is with fshift1=131.1 kHz, and fshift2=353.8 kHz, for a 2fshift difference of 445.4 kHz. For both cases, T_{dump}=200 μsec (i.e. f_{dump}=5 kHz). The difference is clear, the RED trace showing no signs of the 3 artifacts in BLUE, indeed artifacts are completely wiped out.

[5] Blanking occurs at the same time, for all lags of the cross-correlation, however long the lag length and whether or not recirculation is used. This is not the same behaviour as repetitive data valid blanking.

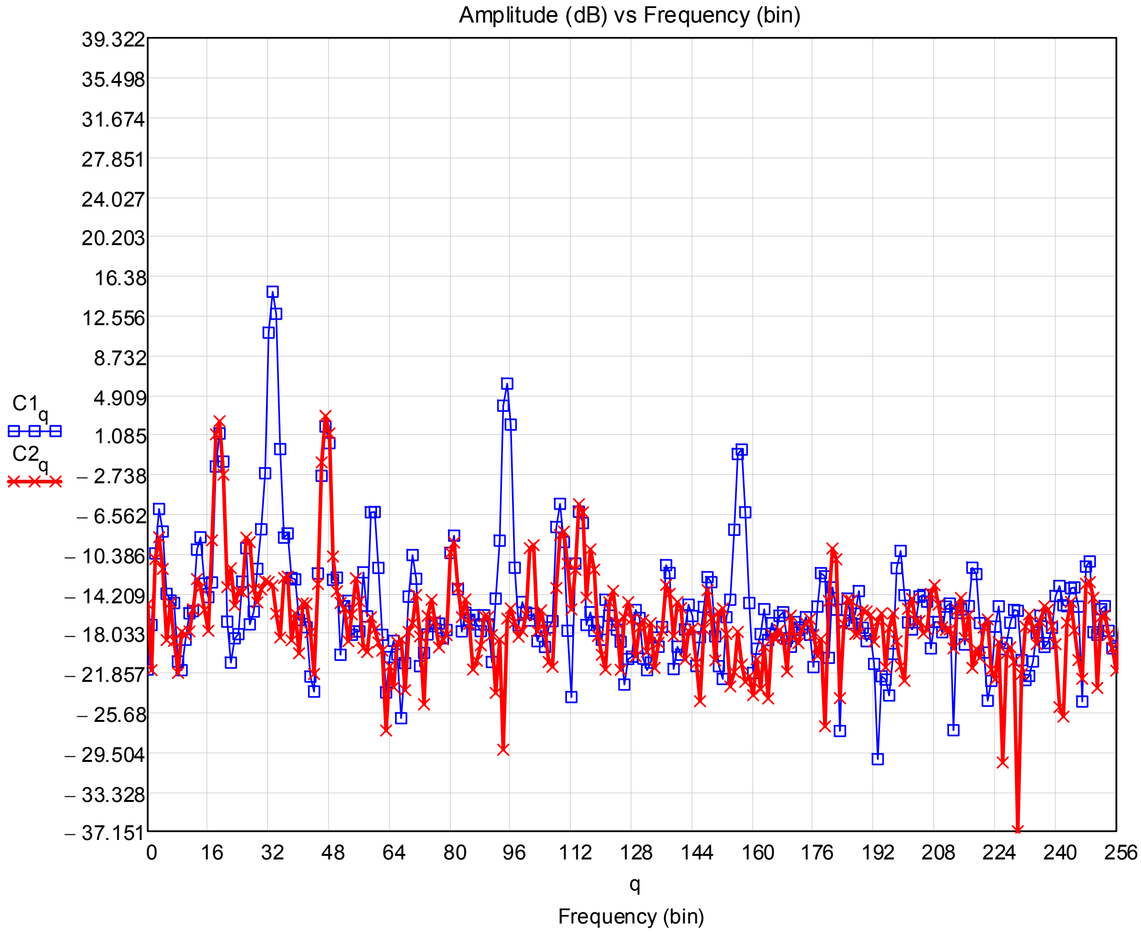

Figure 32 BLUE trace has an fshift differential of 220 kHz, an integer multiple of the correlator dump frequency of 5 kHz. The RED (bold) trace has a 2fshift differential of 445.4 kHz, which is not an integer multiple of the correlator dump frequency. The artifacts are now gone, leaving the primary aliased tones (left 2 peaks), at expected levels.

Normally, any in-band tones are not frequency locked to the system timing reference and this beating behaviour would not occur. However, for any tones, or even continuum signals that are aliased, their frequency difference or sum *is* locked to the system timing reference and so this beating behaviour will occur as long as f_{dump} has the integer relation.

Fshift and Correlator Integration Time Relationships

If there is no integer relation between $2fshift_{diff}$ and f_{dump}, and, $2fshift_{diff}$ % f_{dump} is non-zero and has only 1 digit in the tenths location, then there should ostensibly be a 10 dB reduction in the artifacts as beating at the same time in the $2fshift_{diff}$ waveform will happen only every 10 f_{dump} cycles. If there are 2 digits, one in the tenth and one in the hundredth location, then it follows that a 20 dB reduction is expected as beating will occur only every 100 f_{dump} cyles, 3 digits should be 30 dB and so on.

Further investigation, however, finds that this is not the case. It is only necessary that $2fshift_{diff}$ % f_{dump} is non-zero to virtually wipe out the beating artifact.

Consider the following diagram. It shows a sine-wave, with exactly 7 cycles in the plot, the duration of which we will say is equivalent to our (LTA or CBE) integration time. The plot contains some x-axis ticks, which we will say is equivalent to the correlator chip dump blanking time, spaced apart by the correlator chip integration time T_{dump}. In the integration time, there are 5 dump intervals, and so the sine frequency to fdump ratio is 7/5 = 1.4; i.e. only 1 digit in the $2fshift_{diff}$ % f_{dump} equation.

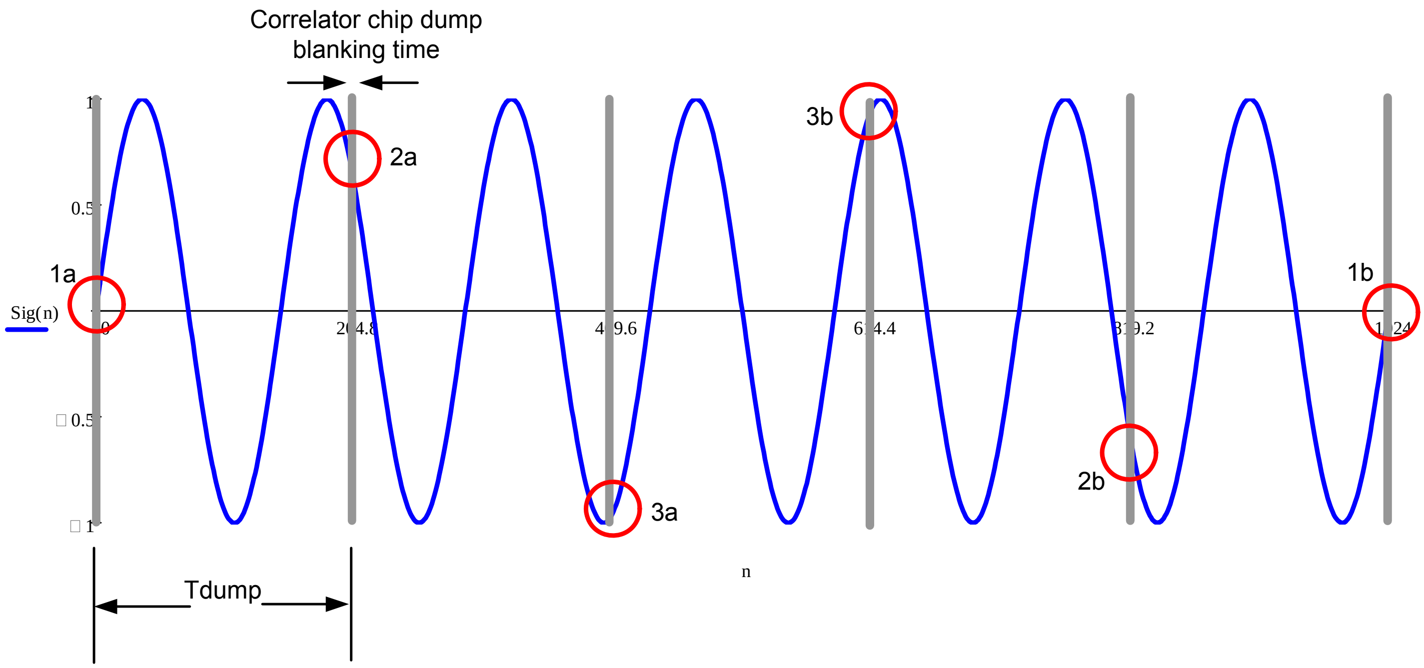

Figure 33 $2fshift_{diff}$ % f_{dump} example. 5 dumps in 7 sine wave cycles for a ratio of 1.4.

Now, when the correlator chip is blanking, integration does not occur, and so *if* there is another part of the sine waveform where integration *does* occur, but is of opposite sign, then a net correlation will be present. We can see that for this scenario, this is not the case. Point 1a "balances" point 1b, point 2a balances point 2b, and point 3a balances point 3b. Correlator chip blanking is perfectly balanced, and so no false correlation or artifacts will occur.

Considering points 2a and 2b, it is clear that point 2a occurs at a phase of the signal slightly less than 180^o (i.e. phase % 360^o), and point 2b occurs at a phase of the signal slightly more than 180^o; they are balanced about 180^o, or about 0.5 cycles.

Generally then:

$$\phi(n) = ([n \times f_r - \mathrm{trunc}(n \times f_r)] - 0.5) \ \% \ 0.5$$

where, n is a particular correlator chip dump, f_r is the ratio of $2fshift_{diff}$ to f_{dump} and $\phi(n)$ is the phase of the sine wave relative to 0.5 cycles.

If:

$$\sum_{n=0}^{N-1} \phi(n) = 0$$

then within the total integration time consisting of N correlator chip dumps, blanking is perfectly balanced, and beating artifacts are completely wiped out.

It turns out that for ratios of 1.2, 1.3, 1.4, 1.6, 1.7, and 1.8, only 6 dumps (5 dump intervals) are required for perfect cancellation. Ratios of 1.1 and 1.9 require 11 dumps (10 dump intervals) for perfect cancellation, and a ratio of 1.5 requires only 3 dumps (2 dump intervals) for perfect cancellation as each dump is exactly 1.5 $2fshift_{diff}$ cycles apart.

A plot showing $\phi(n)$ for the above example ($f_r = 1.4$) is shown in the following figure.

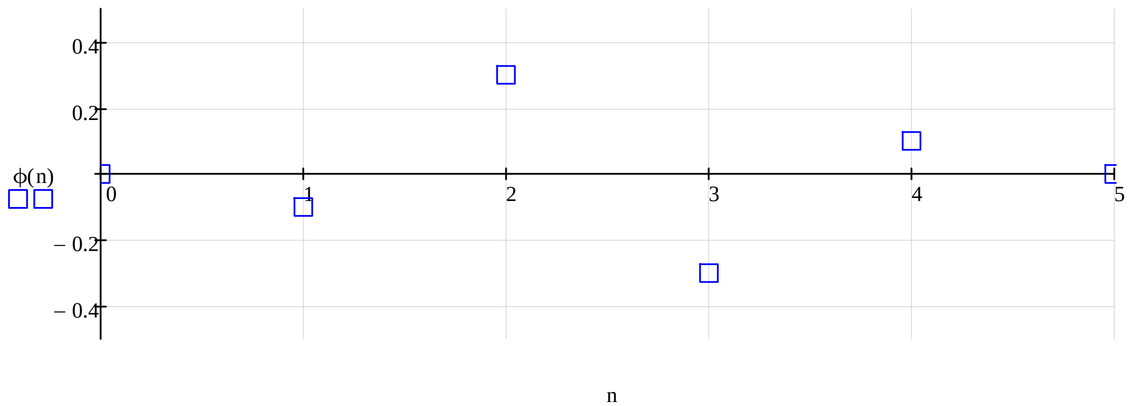

Figure 34 F(n) for the example of $f_r = 1.4$. With 6 correlator chip dumps (5 dump intervals), perfect cancellation of beating artifacts occurs.

If the $2fshift_{diff}$ to f_{dump} ratio f_r has exactly 2 significant digits to the right of the decimal place, then 101 dumps (100 dump intervals), generally, are required to obtain perfect cancellation to wipe out the artifacts completely.

A plot of $\phi(n)$ for 100 dump intervals for a f_r ratio of 1.37 is shown in the plot below:

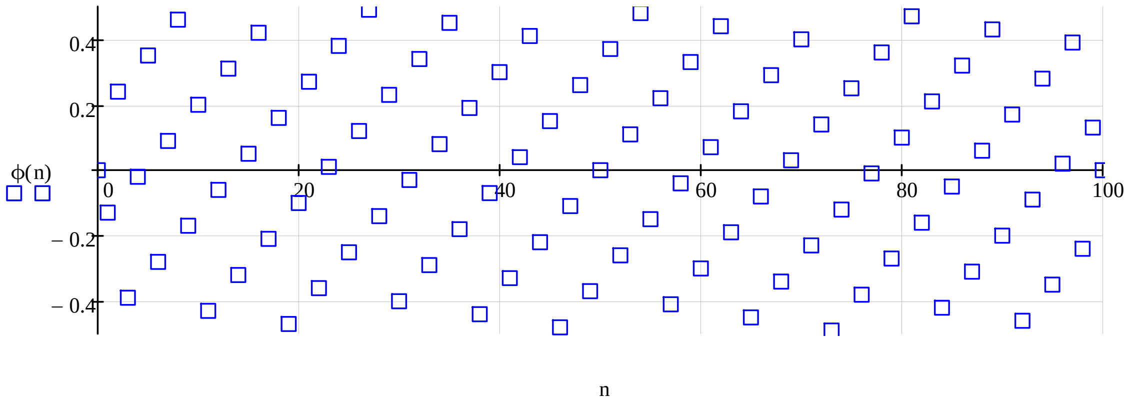

Figure 35 φ(n) for f_r=1.37, and 100 dump intervals.

Of course, generally it is likely not possible to get the precise number of dump intervals to completely wipe out any beating false correlation. Paradoxically, the fewer digits in the f_r ratio to the right of the decimal place, the more easily (fewer correlator chip integration times) it will be to obtain perfect cancellation for a given integration time. However, generally it is true that for 1 digit, multiples of 10 correlator chip integration times are required, for 2 digits, multiples of 100, and for 3 digits, multiples of 1000 etc.

If we assume an arbitrary number of f_r digits, what is the artifact attenuation, if there isn't an exact multiple of 1000 correlator chip integration times within a given final (LTA or CBE) integration time? Generally, as long as the number of f_r digits is ≥ 1 the correlation is:

$$\frac{1}{N}\sum_{i=0}^{N-1}\sin(\phi(i))$$

where N is the number of correlator chip integration times.

It turns that this is the same form as the sinc(x) function, with envelope of 1/N, where N is the number of correlator chip integration times within a final (LTA or CBE) integration time.

So, as long as the number of f_r digits is ≥ 1, without worrying about how many f_r digits there are, the minimum artifact amplitude will be down by 10*log(1/N), for N correlator chip integration times. If there are fewer f_r digits, there are more nulls in the function. To hit a null requires only that the number of correlator chip integration times is an integer multiple of 10, 100, or 1000 for number of f_r digits 1, 2, and 3 respectively.

Figure 36 shows the above equation, and the 1/N envelope for f_r = 1.7 for the first 100 correlator chip integration times.

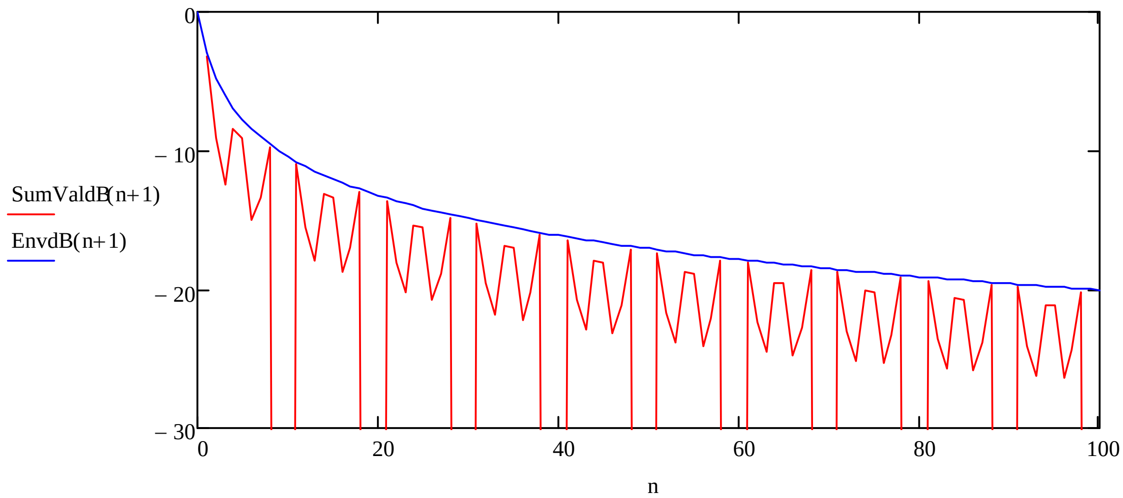

Figure 36 1/N envelope and exact artifact amplitude versus the number of correlator chip integration times.

Choices of Fshift Frequencies and Correlator Chip Integration Times

There are now 3 criteria that need to be met (and a 4^{th} is desirable) to ensure that there are no (or at least highly attenuated) spectral artifacts in the data, and that anti-aliasing functions as expected:

1. The fshift ratio (fshift_x / fshift_y) must be non-integer, and a non-repeating+infinite decimal value.

2. $2fshift_{diff}$ % f_{dump} must be non-zero.

3. $2fshift_{sum}$ % f_{dump} must be non-zero[6].

4. If possible the number of correlator chip integration intervals in an LTA or CBE integration time should be a multiple of 1000 as 3-digit f_r ratios are normally more prevalent and easily obtained than 1 or 2 digit ratios.

The simple equation:

fshift = fbase + f1 * n + f2

has been found to be sufficient to find fshifts that meet criteria 1., given well chosen parameters. (Example: fbase=13 kHz, f1=3 kHz, and f2=100 Hz, limited to a sinc(x) attenuation of ~38 dB in a 1 second integration time.)

[6] Only $2fshift_{diff}$ has been empirically tested; it is assumed that $2fshift_{sum}$ has a similar effect; $2fshift_{diff}$ has been found to be the more difficult relation to obtain, and once obtained, $2fshift_{sum}$ essentially comes for free.

With these frequencies (and trying many others with 100 Hz fshift tuning resolution), however, and using "standard" correlator chip integration times[7] of 200 μsec, 250 μsec, 312.5 μsec, and 400 μsec, it was *not* found to be possible to also meet criteria 2. on all baselines.

It was found that either the correlator chip integration time had to be adjusted by +/-1 μsec (i.e. 1 μsec resolution), or fshift tuning resolution had to be improved to 10 Hz to achieve all 3 criteria.

If fshift tuning resolution is changed to 10 Hz, and still using standard correlator chip integration times, values of fbase=0 Hz, f1=3.11 kHz, and f2=100 Hz met all criteria, as was found to generally be possible with many other values (Example: fbase=3.13 kHz, f1=0.37 kHz, and f2=0.11 kHz for maximum fshift differential of ~10 kHz on 27 antennas for acceptable bandwidth loss when operating at minimum bandwidths—good enough for ~28 dB of sinc(x) anti-aliasing attenuation, limited to ~23 dB of 1/√N attenuation at 62.5 ks/s in a 1 second integration time).

[7] i.e. correlator chip integration times such that there are an integer number of each in 10 msec.

Conclusions

This memo has investigated anti-aliasing and bin-centered tone behaviour in the EVLA WIDAR correlator. Some new information has come to light, namely the following:

1. Anti-aliasing is fundamentally limited by either the sinc(x) washing curve, or the $1/\sqrt{N}$ noise de-correlation curve, whichever is greater. Thus, one can never expect anti-aliasing to hit a “null” and be completely wiped out, and anti-aliasing performance is bandwidth and integration time dependent.

2. Fshift frequency choice is important to ensure there are no low-level artifacts around bin-centered tones. In particular, the fshiftX/fshiftY ratio must have a non-repeating+infinite decimal part. The fundamental cause of these artifacts was not investigated or explained.

3. Correlator chip integration times, due to the fact that there is a 1 μsec blanking time for every dump, must be chosen so that they do not beat with fshift frequencies. This can be accomplished with either (or both) 10 Hz fshift tuning resolution with carefully chosen frequencies according to a simple equation, or correlator chip integration times with 1 usec timing resolution that are “non-standard”.

4. Spectral artifact amplitudes from correlator chip integration times beating with fshift frequencies, generally decay as $1/N_{dumps}$, but specifically can be completely wiped out if the LTA or CBE integration time is a multiple of 1000 correlator chip integration times.

Further telescope testing with the H-maser-locked bin-centered tone, and with narrowband astronomical sources is recommended to determine if the solutions found in this memo function as described.

Finally, there is some probability that these solutions may solve the problem of the “correlator wobble” conundrum, as the mechanisms found are such that they could produce this undesired effect. Further testing with careful choice of fshift frequencies and correlator chip integration times, using continuum sources and 100 millisecond LTA integration times should be performed to see if this is indeed the case.

Appendix

Some thoughts regarding Bob Sault's August 24, 2010 report, "WIDAR anti-aliasing and spectral dynamic range test"

In particular, regarding Figure 7 of the report, the upper panel indicates only about -15 dB of anti-aliasing suppression (-35 dB total, with -20 dB attributed to the filter transition band cutoff). This is far short of the (now) expected 1/√N de-correlation of -30 dB (1 second integration, 10^6 samples).

There are some things to note in the figure. Comparing the sub-band boundary level in the upper panel to the lower panel shows ~-24 dB of suppression, but this is false economy as both sub-bands' spectra are virtually obliterated by boxcar windowing of non-bin centered tone sidelobes, also indicating that the tone is quite far off from being bin-centered. Thus a fairly large fshift differential is likely at play (also, the "Aliased response" tone in the lower panel is clearly well off bin center). Even though the fringe rotators were reported on for the lower panel, the coarse nature of their operation, and the very few lags used (128), effectively means they were off.

It is highly likely that the large disparity between the expected (1/√N) response and the achieved response could be attributed to a combination of the integration time beating with fshift differences, to the integer fshift ratio effect, and less so to the very effect noted for in Figure 8 in this document, namely the steepness of the transition band and the shifting of the tones to other than expected parts of the transition band.

Bin-centered tone testing

Bob Sault's report also found that when there were odd ratios of fshifts, bin-centered tones had sidelobes around them to the level of a few percent. Software correlator simulation tests did, indeed, confirm that this is the case (in addition to the lower sidelobes found for other fshift ratios).

The following figure shows the bin-centered tone test of Figure 8, showing the upper 128 channels of sub-band 0 cross-correlation. FshiftX=15 kHz, and fshiftY=3 kHz, for a 5-to-1 ratio. Hanning windowing was used, and each frequency point is 61.035 kHz. The sidelobes around the main peak at ~-18 dB level are clearly evident. Each of the sidelobes has a well-defined phase, indicating that they will not de-correlate with further integration time and show up as (false) real signals.

N.B.: In this test, the tone power was significantly lower (~10 dB) than tests discussed previously.

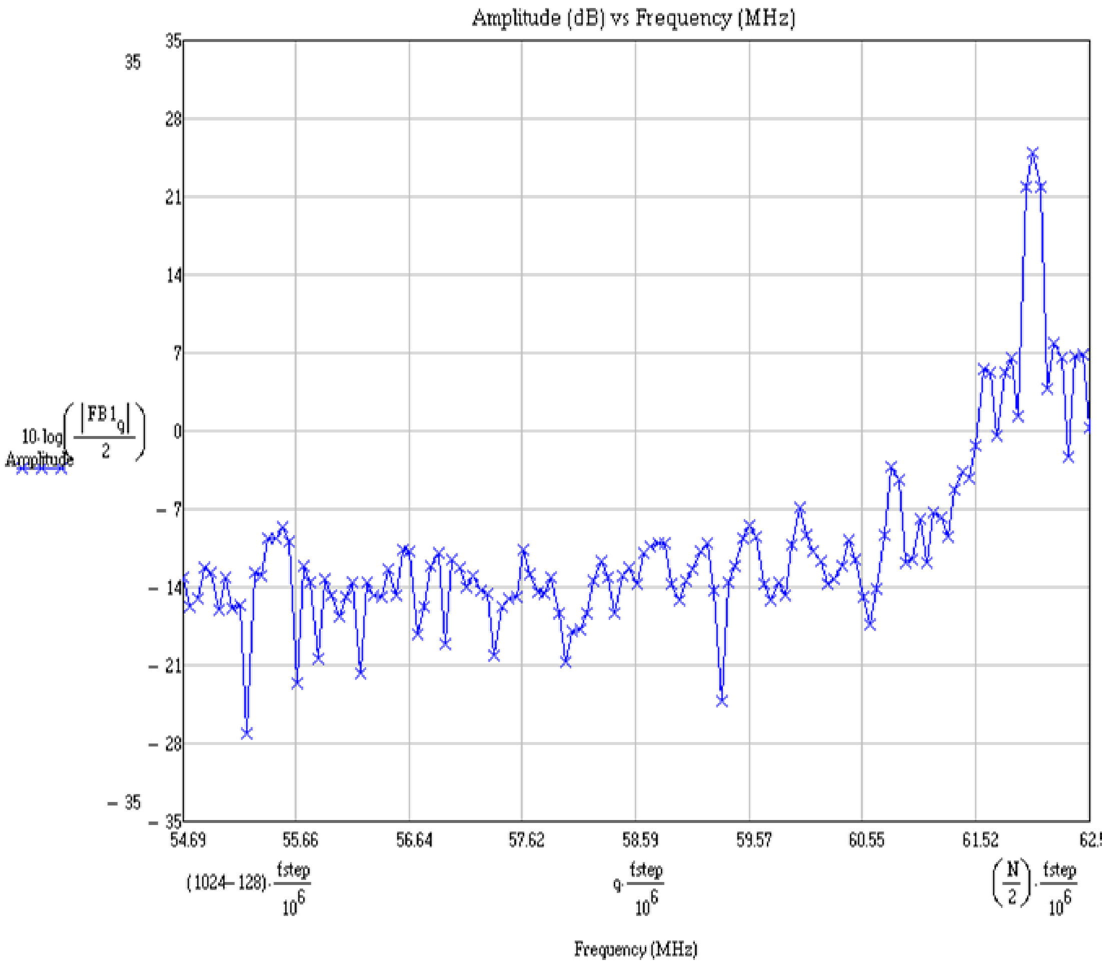

Figure 37 Bin-centered tone with an fshift ratio of 5-to-1. Sidelobes at ~-18 dB level are evident.

A small offset was added to each fshift to produce the improvement (fshiftX=51.1 kHz, fshiftY=11.1 kHz) shown in the following figure[8], although this ratio, when calculated has a value of 4.603636…, and so should contain sidelobes at the ~-40 dB level (which it appears to show).

[8] The drop in the noise floor is due to a longer (6.25 Msample) correlation, rather than fshift frequency changes.

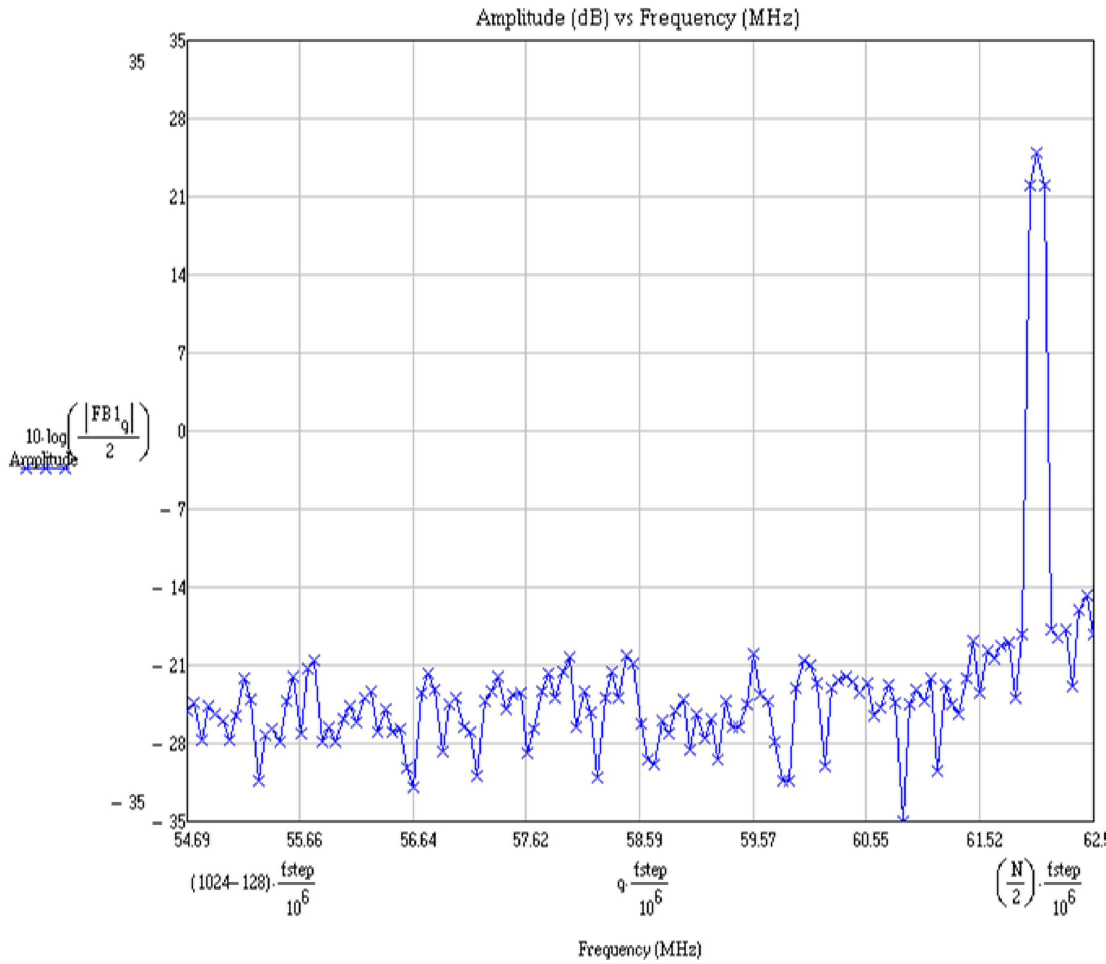

Figure 38 5-to-1 bin-centered sidelobes are greatly once a small offset was added to each fshift. In this case, fshiftX=51.1 kHz, fshiftY=11.1 kHz. It would seem though that sidelobes are appearing at <-40 dB level (51.1/11.1=4.6036036…).

References

[1] Carlson, Brent, A Proposed WIDAR Correlator for the Expansion Very Large Array Project: Discussion of Capabilities, Implementation, and Signal Processing, NRC-EVLA Memo# 001, May 18, 2000.

[2] Carlson, Brent, REQUIREMENTS AND FUNCTIONAL SPECIFICATION: EVLA Correlator Chip, RFS Document A25082N0000, Revision 2.5, January 20, 2010.

ADDENDUM: MARCH 24, 2011 – Recirculation Phase Effects

This addendum includes information on the effects of using recirculation (either static or dynamic) for correlation, when phase is passed serially through the recirculation memory.

The effect has to do with limited memory width in transporting phase information thru the recirculation memory. For all 8 streams to undergo recirculation in the Recirc FPGA, phase for each two streams is sampled at fs/8, serialized into one bit stream through the memory, and then reconstructed at the other end (refer to Recirc RFS A25090N0000 Rev. 3.1a (or higher), page 49, RSSR0_1 register description). If the phase rate is too high, it results in phase jitter, which reduces correlated coherence. (If only 4 streams need to undergo recirculation, then phase can go thru the memory at full resolution and this does not apply. The Recirc GUI allows for this setting—see the Baseline Board User Manual, A25080N0001, Figure 5-1, bullet "S" Recirc GUI description.) It should be noted that for the technology of the time, it would likely be impossible from a board routing perspective and certainly very expensive (~$1M) to provide enough memory width so as not to have this effect.

This description is derived from an email, "Recirculation phase sampling", of January 28, 2011. Tests also indicated that jitter from two different fshifts would beat, and result in a periodic coherence vs lag function, which resulted in false artifacts showing up in the cross-power spectrum. To attenuate these spectral artifacts requires that the phase-jitter-induced coherence loss be very low.

Phase jitter is:

$$\phi_{jitter} = 2\pi \cdot \frac{8}{fs} \cdot fshift$$

Where *fs* is the sub-band sample rate. Note that *fshift* includes the earth-rotation fringe phase rate, which for EVLA baselines and ν_{LO} = 50 GHz is normally very much smaller than the antenna-dependent LO offset. Coherence is then:

$$\eta = \frac{\sin(\phi_{jitter})}{\phi_{jitter}}$$

For example, for a 78 kHz *fshift*, and *fs*=2 Ms/s, $\eta \approx 50\%$. A plot of magnitude vs lag, for a section of lags with (RED) and without (BLUE) this effect is shown below:

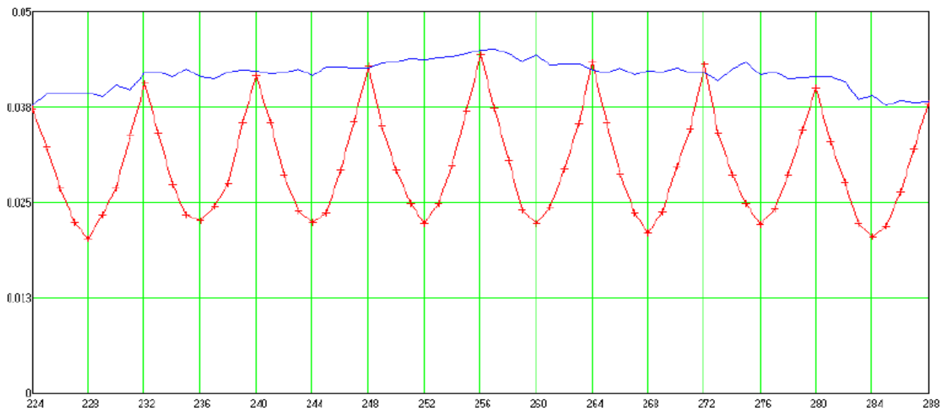

This data was taken with the EVLA on a 6.7 GHz line, with and without recirculation. Spectral artifacts show up at about -11 dB (RED), but no artifacts show up for the BLUE data.

If, for the above example, fshift is reduced to 1 kHz, the coherence *loss* is reduced to ~0.01%, or the artifacts should be down another ~-37 dB to ~-48 dB from the main signal. (I recall that Ken Sowinski reports subsequent testing on this line with reduced fshift did indeed wipe out the spectral artifact, although I can't find the specific email with this note.)

There is also a thought that to introduce uniform phase jitter, 8/fs should not nicely divide into 1/fshift, where "nicely" means no finite number of decimal places in the result. For the fs=2 Ms/s and fshift=78 kHz case, 1/((8/fs)*fshift) = 3.2051282051282…

Further testing is required to determine how much phase jitter is allowed, if the rules for fshift and correlator dumping, established in the conclusions (page 48) of this memo are followed. i.e. if coherence vs lag is not regularly repeating as in the above plot—resulting in spectral artifacts, then more coherence loss and a larger range of fshift is likely allowed.