

Simulation Tests to Quantify the Spectral Dynamic Range and Narrowband Interference Robustness of the WIDAR Correlator for the EVLA

NRC-EVLA Memo# 009

Brent Carlson, Peter Dewdney

November 1, 2000

ABSTRACT

In the first memo of this series [1], a number of simulations were performed with the primary goal of quantifying the performance of the proposed WIDAR correlator compared to an ideal correlator operating at the full bandwidth. These simulations demonstrated that the WIDAR correlator yields satisfactory results. It was also stated and demonstrated that the WIDAR correlator could increase the spectral dynamic range and interference robustness in the presence of powerful narrowband interference by causing quantizer-generated harmonics to decorrelate. This memo presents results of a more detailed investigation that more carefully quantifies these interference robustness characteristics. L-band interference spectra from the VLA site is analyzed to show that the correlator is more than capable of operating in its intended environment. Spectral dynamic range tests also yield the important conclusion that only a 3-bit, 8-level, initial quantizer should be required to yield interference-robust performance—even if the interference at the VLA site is an order of magnitude worse than it currently is. This is an important conclusion since 3-bit initial quantization will reduce the cost of the quantizers and the fiber-optic transmission system compared to a 4-bit system. The results of WIDAR correlation of a simulated L-band interference environment including a few +20 dB tones representing satellite interference into the “near-in” sidelobes of the antennas are also presented.

Background

An extensive discussion of the mechanism whereby the WIDAR correlator can cause narrowband harmonics to decorrelate is presented in [1] and is summarized here. Each quantizer sees the analog signal at a slightly different frequency imposed by offsetting the Local Oscillators in the antennas: when the frequency difference is removed in the correlator, narrowband harmonics—that are at multiples of the fundamental—decorrelate. Unfortunately, the degree of decorrelation is limited by the amplitudes of the harmonics of the digital mixer in the correlator: digital mixer harmonics stop the harmonics generated by the quantizer. A workaround to this problem involving some post-correlation fringe rotation was discovered and will be presented in this document.



The major flaw in the interference robustness investigation presented in [1], is that it did not consider the dynamic range limiting impact of intermodulation products that are generated by the quantizer from powerful narrowband interference. The simulations in this memo quantify this effect.

Simulation Results

A number of simulations were performed using a (software) WIDAR correlator and a full bandwidth correlator (without frequency shifting) where the interference consisted of a narrowband signal generated from a FIR filter and one or more additional tones added to a known quantity of Gaussian noise. The total power of the interference and noise is precisely known from the simulation parameters and so it is possible to precisely quantify the spectral dynamic range in terms of the ratio of total interference power to total power into the quantizer. The spectral dynamic range performance derived from these tests is summarized in Table 1.

	2 Signals			3 Signals			4 Signals		
	interf/total power			interf/total power			interf/total power		
	25%	50%	90%	25%	50%	90%	25%	50%	90%
Ideal/fullband 4-bit ¹	n.t.	n.t.	35	>50*	42.5	39	>50*	43.5	39
WIDAR 4-bit	>55*	>55*	>55	>50	42.5	39	>50*	43.5	39
WIDAR 3-bit	>50*	>50*	>50	>50	n.t.	38	n.t.	n.t.	n.t.

Table 1 Table summarizing the spectral dynamic range performance (in dB) of an ideal 4-bit (fullband) correlator, a WIDAR 4-bit correlator with 4-bit initial quantization, and a WIDAR 4-bit correlator with 3-bit initial quantization. In this case, 4-bit initial quantization is 15 levels (-7, -6, ..., -1, 0, 1, ..., 6, 7) with a threshold differential of 0.374σ and 3-bit initial quantization is initially 4-bit/15-level, but shifted by one bit to 8-levels (-4, -3, -2, -1, 0, 1, 2, 3). Quantities with an ‘*’ were derived from other tests within the table, and ‘n.t.’ means that no test was performed.

The WIDAR correlator clearly outperforms the fullband correlator when two narrowband signals are present because of the harmonic decorrelation effect². However, in the three signal and four signal cases, the WIDAR correlator does not perform any better strictly in terms of spectral dynamic range. This limitation arises from the fact that some intermodulation products of the fundamental frequencies correlate. For example, if a frequency shift of ϵ is imposed before quantization, then for three frequencies A, B, and C, the intermodulation product $(A+\epsilon)+(B+\epsilon)-(C+\epsilon)$ (which simplifies to $A+B-C+\epsilon$) will correlate once the frequency shift of ϵ is removed in the correlator’s fringe stopper (digital mixer). (A similar effect is possible with two narrowband signals if the quantizer

¹ This indicates the number of bits coming out of the initial quantizer.

² And incorporating the technique which overcomes the digital mixer harmonic correlation effect.



is too asymmetric and generates even harmonics.) This limitation is fundamental and can only be overcome by using more bits in the quantizer.

Despite the above fundamental limitation, the WIDAR technique still reduces the number of intermodulation products to only those that meet the criteria in the above example. Thus, a second metric that is useful in quantifying the performance of the WIDAR correlator is to consider the number of “secondary products” (intermodulation and harmonic products) appearing in the cross-power spectrum. In Table 2, a comparison is made of the number of secondary products for a WIDAR correlator and an ideal fullband 4-bit correlator. These results were obtained experimentally by counting the correlated intermodulation and harmonic products in plots of the simulation results.

	2 Signals			3 Signals			4 Signals		
	interf/total power			interf/total power			interf/total power		
	25%	50%	90%	25%	50%	90%	25%	50%	90%
Ideal/fullband 4-bit ³	n.t.	n.t.	15	0*	12	27	0*	31	46
WIDAR 4-bit	0*	0*	0	0	3	6	0*	10	16
WIDAR 3-bit	0*	0*	0	0	n.t.	7	n.t.	n.t.	n.t.

Table 2 Table summarizing the spectral dynamic range performance *in terms of number of secondary interference products* for an ideal 4-bit (fullband) correlator, a WIDAR 4-bit correlator with 4-bit initial quantization, and a WIDAR 4-bit correlator with 3-bit initial quantization. In this case, 4-bit initial quantization is 15 levels (-7, -6, ..., -1, 0, 1, ..., 6, 7) with a threshold differential of 0.374σ and 3-bit initial quantization is initially 4-bit/15-level, but shifted by one bit to 8 levels (-4, -3, -2, -1, 0, 1, 2, 3). Quantities with an ‘*’ were derived from other tests within the table, and ‘n.t.’ means that no test was performed.

Several plots of amplitude in dB versus frequency of some of the simulation results used in Tables 1 and 2 are shown in Figures 1 and 2. In Figure 1 plots are shown where the narrowband interference (i.e. the “signals”) power is 90% of the total power in the band. In all cases the WIDAR correlator outperforms the simple 4-bit correlator although the limitation of the correlation of intermodulation products is clear.

³ This indicates the number of bits coming out of the initial quantizer.



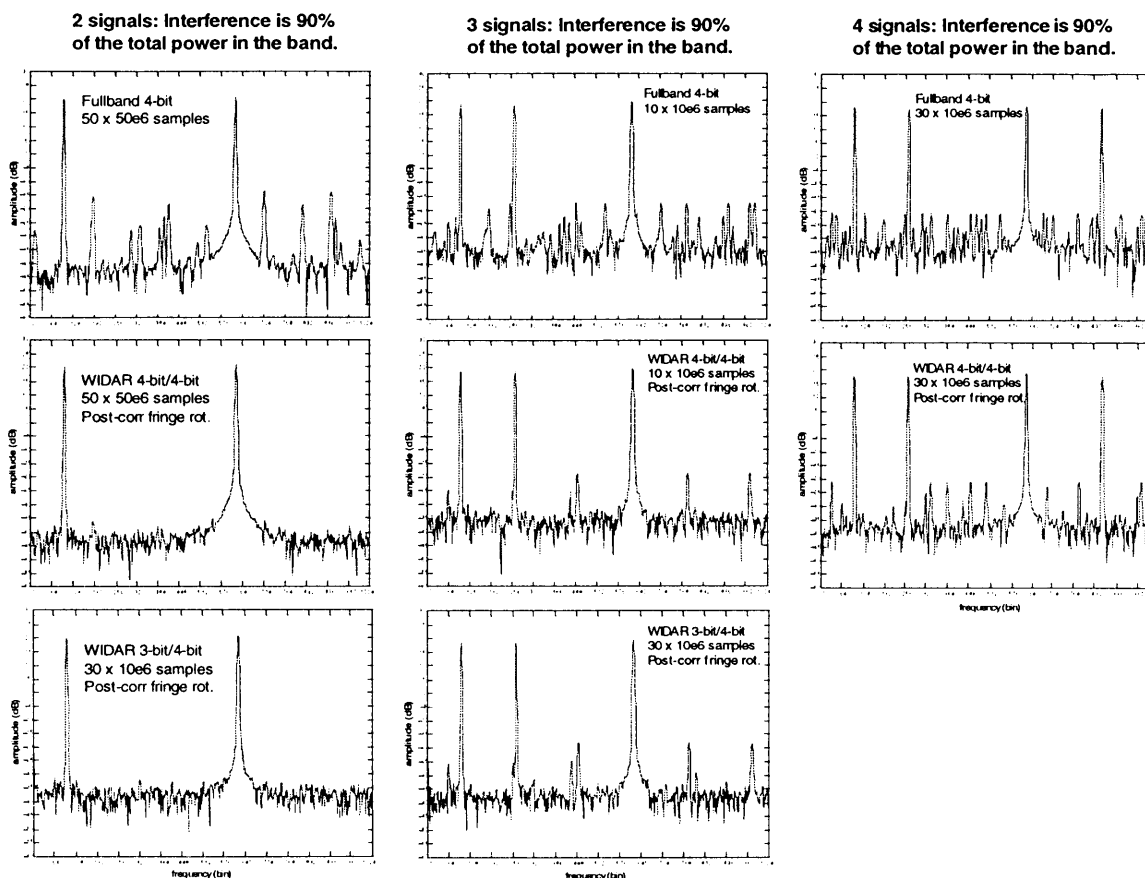


Figure 1 This figure contains several plots of simulation results that were used to fill in Tables 1 and 2. All plots are for the case where the total (narrowband) interference is 90% of the total power in the band. The fullband correlator is a straightforward 4-bit correlator with no fringe rotator. The two WIDAR correlators are with 4-bit (15-level) and 3-bit (15-level translated to 8 level *after* WALSH de-switching) initial quantization. The initial quantizer is a 4X interleaved, algorithmic A/D converter with 1 mV RMS decision level errors (which is representative of a worst-case quantizer). The WIDAR correlator uses a 5-level fringe rotator with post-correlation fringe rotation for harmonic suppression. The “skirted” interference line is resolved and has a power (in arbitrary units) of 124, the other narrowband (cw) lines have a power of 72 (each), and system noise power is 22, 29, and 37 for the 2, 3, and 4 signal cases respectively. Kaiser windowing was used with $\theta=5\pi$ because of the high spectral dynamic range of the signals. Note that not all of the integration times are equal in the above plots because of the considerable simulation time required for each one (i.e. integration times were roughly chosen to yield results deemed to be definitive).

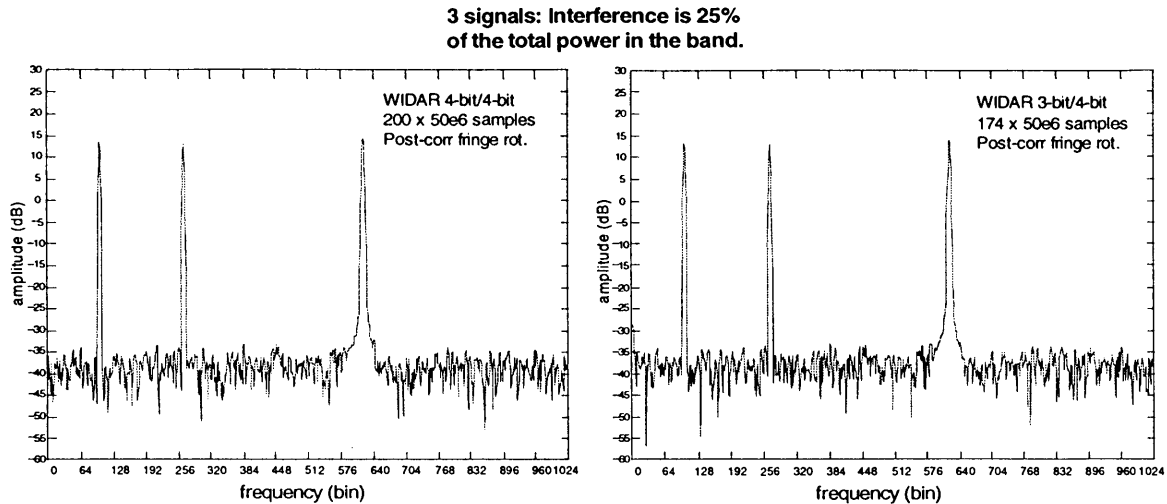


Figure 2 Plots of the cross-power spectrum with three “signals” where the interference (signal) is 25% of the total power in the band⁴. The plot on the left is with 4-bit (15-level) initial quantization, and the plot on the right is with 3-bit (15-level shifted to 8-level) initial quantization. Harmonic suppression with post-correlation fringe rotation was used (although it probably didn’t have to be in this case). These results indicate that close to 50 dB of spectral dynamic range can be obtained. It is unknown how much better the dynamic range is in this case since prohibitive simulation times would be required—as it is, each of the above plots correlated about 10^{10} samples (2.5 seconds of real-time data at 4 Gs/s).

3-bit or 4-bit Initial Quantization?

From the preceding tables and figures, it is clear that 3-bit, 8-level (translated from 4-bit 15-level) initial quantization performs as well as 4-bit 15-level initial quantization as far as spectral dynamic range is concerned. Further tests revealed that the 3-bit 8-level reduction in SNR is about 3.5% and the 4-bit 15-level reduction in SNR is about 1.5%—agreeing well with expected SNR losses. Other conclusions that have been determined through simulation testing are as follows:

- 3-bit 8-level spectral dynamic range is the same as 4-bit 15-level. If 3-bit 7-level is used, then the spectral dynamic range is worsened by ~6 dB.
- 3-bit 8-level encoding can either be 4-bit 15-level (7, 6, ... 1, 0, -1, -2, ... -6, -7) translated via a bit shift to 3-bit 8-level (3, 2, 1, 0, -1, -2, -3, -4) or just 3-bit 8-level with odd output encoding (7, 5, 3, 1, -1, -3, -5, -7).
- 4-bit 16-level with odd output encoding (13, 11, ... 1, -1, -3, ... -11, -13) improved the spectral dynamic range by about 4 dB compared to 4-bit 15-level if the same quantizer threshold steps (0.374σ) are used. Increasing this threshold step improved the dynamic range (Figure 3 (e))—presumably because it minimized clipping effects—although it may be due to optimization for two tones.
- These dynamic range effects are not due to a reduction in the ENOB (Effective Number Of Bits) since an ideal flash A/D converter was used in the simulations.

⁴ In this case the system noise power, in arbitrary units, is 806.

The conclusions seem to indicate that it is very important to use all available levels to achieve the maximum spectral dynamic range. Reducing the number of levels from the maximum reduces the spectral dynamic range a considerable amount. Figure 3 contains a number of cross-power spectra plots for several different initial quantizer configurations with two strong narrowband signals in the band. A simple, fullband correlator was used in all cases so that harmonics and intermodulation products did not decorrelate. The narrowband signals contain 90% of the total power in the band (same power relations as Figure 1).

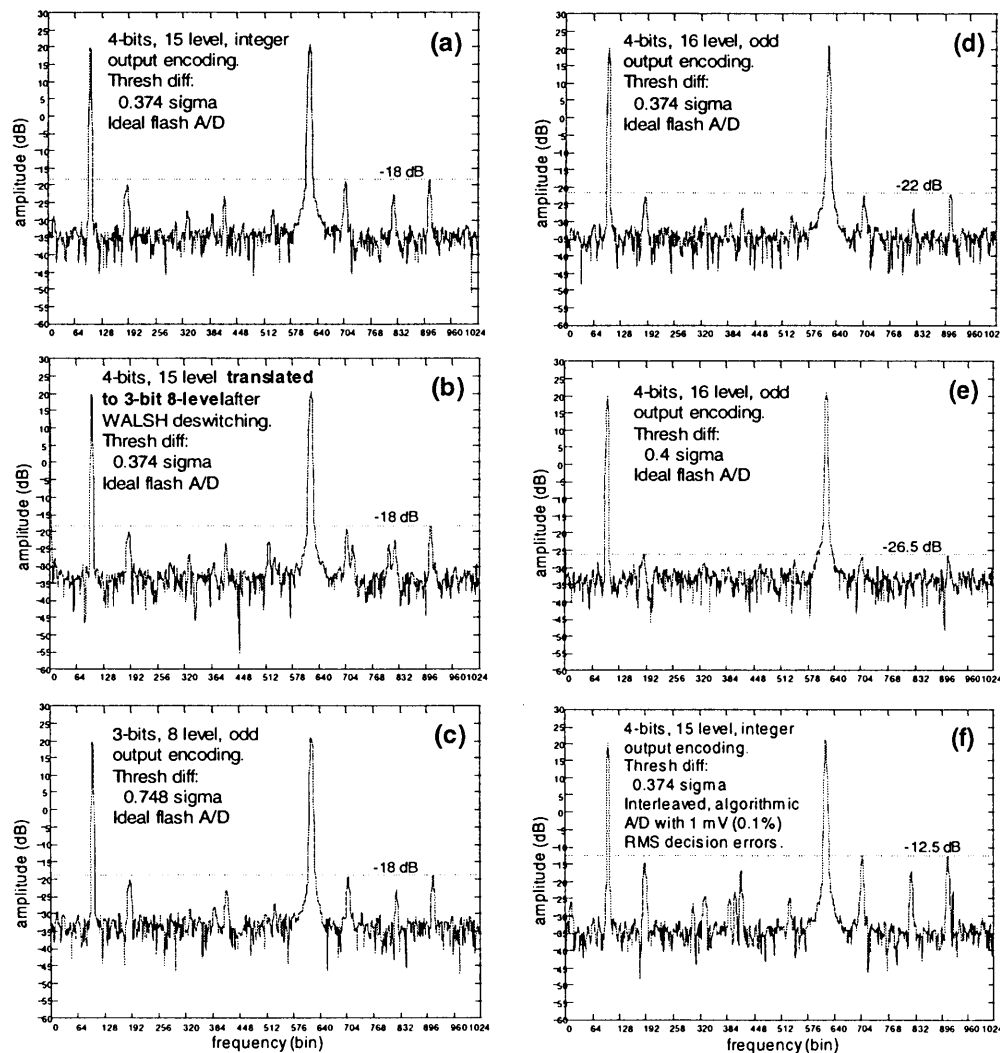


Figure 3 Fullband correlator cross-correlation plots used to demonstrate the performance of 4-bit and 3-bit initial quantization. Three-bit, 8-level quantization achieves the same spectral dynamic range as 4-bit 15-level quantization. In most cases, ideal flash A/D converters were used to factor out ENOB effects. A significant DC component is evident in (b)—due to the translation from 15-level to 8-level via a bit shift (on 2’s complement data) after WALSH de-switching. In (f), a simulated “real” A/D converter was used with results considerably worse than the ideal converter (a). Four-bit 16-level (e) achieves the best results if the threshold step is optimized—presumably to prevent clipping—although it may be an optimization for two tones in this case.

The above conclusions can be applied to the EVLA in one of two ways:



1. Use 3-bit, 8-level initial quantization and save money by only having to transmit 3 bits per sample via the Fiber Optic Transmission System (FOTS). To maximize spectral dynamic range and minimize A/D hardware, use 8-level odd encoding.
2. Use 4-bit, 16 level initial quantization with odd output encoding to maximize spectral dynamic range. This option requires a 4-bit FOTS but it ensures that significant spectral dynamic range margin is obtained. It is important to ensure that the A/D converter operate with close to 4 ENOBs in order to take full advantage of the increased number of levels.

Option 2. is suitable for use with a WIDAR correlator since the data must always be FIR filtered and requantized before correlation. That is, the 16-level odd output encoding (13, 11, ... 1, -1, -3, ... -11, -13) can be handled in the FIR multiplier lookup tables without any impact on the existing design⁵. This encoding is probably not suitable for a fullband 4-bit correlator because the size of the multiplier may become excessive⁶. There may be problems using the 15-level integer encoding (7, 6, ... 1, 0, -1, ... -6, -7) correlator chips for the wideband autocorrelator if 16 level odd encoding is used. (A 15-level correlator chip is contemplated so that the unused state (-8) can be used for data valid flagging.)

Later in this document, analysis of the interference environment at the VLA site indicates that a 3-bit 8-level initial quantizer should be more than adequate for operation in the existing and (probably) future interference environments. Thus, it would seem that option 1. is the most cost-effective configuration.

Extending Harmonic Decorrelation

Frequency shifting required for sub-band transition band anti-aliasing in the WIDAR correlator causes narrowband harmonics to decorrelate. Unfortunately, the decorrelation effect is limited by the amplitude of the harmonics of the digital mixer function in the correlator: digital mixer harmonics fringe-stop narrowband signal harmonics [1]. A 3-level digital mixer/fringe stopper has a 3rd harmonic amplitude of ~ -8 dB and thus, the amount of 3rd harmonic decorrelation that occurs is limited to -8 dB. A 5-level digital mixer function extends the 3rd harmonic decorrelation to ~ -15 dB, although it requires some more silicon in the correlator chip. A 7-level mixer would extend this to ~ -21 dB, but it would require a prohibitive increase in silicon in the correlator chip.

A way to work around this limitation was discovered and it operates by offsetting the correlator's fringe stoppers so that there is a residual fringe rate that is removed after correlation. Since the post-correlation fringe stopper uses many bits with virtually no harmonics, the residual harmonics from the correlator output will decorrelate within the *incoherent* integration time. It is expected that this function could be performed within the correlator chip readout circuitry (and Baseline Board phase generator circuitry) so that front-end and back-end software need not deal with it or even know it exists. Furthermore, it was found that the residual fringe rate need only be generated by putting a

⁵ Although this has not yet been tested in the simulator.

⁶ Although this assertion would have to be tested in more detail since there may be some multiplier table trick that can be played.

varying *discrete* phase offset into the fringe stopper, rather than a phase offset and rate. This means that there will be no fringe smearing within the hardware integration time or subsequent *incoherent* LTA (Long-Term Accumulator) integration time. If decorrelation is required within a longer *coherent* integration time, some random discrete phase offset function versus time—rather than a simple phase “stairstep” versus time—will be required. Discrete phase offset insertion and post-correlation removal was used in the tests shown in Figure 1 to virtually eliminate the harmonics present at bins 191 and 704 in the WIDAR correlator. Figure 4 shows a number of plots of cross-power spectra that clearly demonstrates the sort of attenuation that can be achieved.

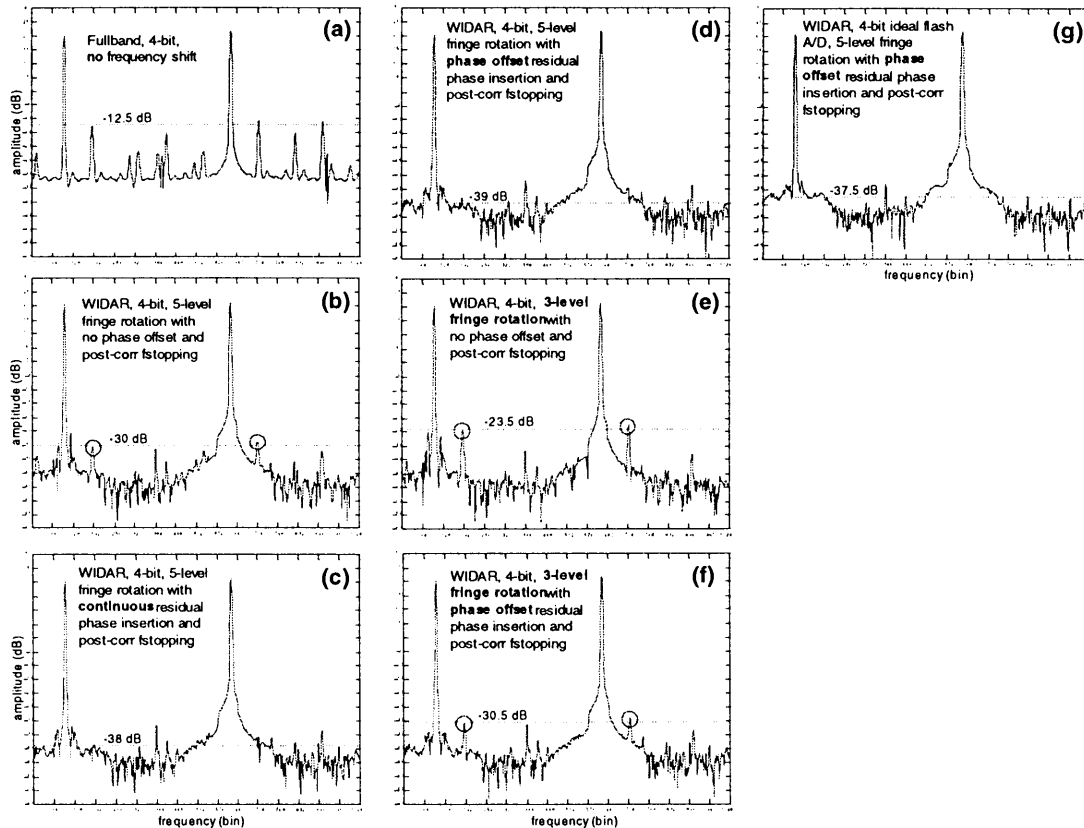


Figure 4 Several cross-power spectrum plots showing the effect of continuous and discrete phase-offset insertion with post-correlation removal. (a) is a fullband correlator with no frequency shifting and thus no harmonic decorrelation. (b) is a 4-bit WIDAR correlator with 5-level fringe stopping and no phase insertion and removal. (c) shows that harmonics at bins 191 and 704 decorrelate with *continuous* residual phase insertion and removal. (d) shows that *discrete* residual phase offset insertion and removal achieves the same result as continuous phase. (e) is WIDAR correlation with 3-level fringe rotation. (f) is WIDAR correlation with 3-level fringe rotation, phase offset insertion, and post-correlation removal. In this case the attenuation of the harmonics is not as much as expected—requiring further investigation. In all cases there are 30 hardware dumps of 5×10^6 samples each with an X-station phase offset of 0.05381 cycles/dump and a Y-station phase offset of -0.0328 cycles/dump. In this extreme situation the narrowband lines contain 99.5% of the total power in the band—resulting in some signal distortion. Distortion in the right-most narrowband line from the requantizer is evident and is probably the same reason for the correlated “plateau” surrounding the left-most CW signal. (g) is a WIDAR correlation with an ideal flash A/D converter for the initial quantizer, indicating that the distortion effects are not due to non-ideal quantizer behaviour.

An additional interesting bit of information can be determined from Figure 4 by comparing (a) and (b). The suppression of the 3rd harmonic near frequency bin 191 using a 5-level fringe rotator is about -17.5 dB. From [1], it was determined that the 3rd harmonic of a 5-level fringe rotator is about -14.5 dB down from the fundamental. This discrepancy can be explained as being due to the reduction in the 3rd harmonic amplitude from phase dithering due to quantized phase differencing. Spectral analysis indicates that phase dithering spreads the harmonic power across many more frequencies than if there is no phase dithering. Since the total harmonic power is the same in both cases [2], it is reasonable to conclude that the 3rd harmonic power has been reduced. This also appears to be the case for 3-level fringe rotation plotted in Figure 4 (e).

The Interference Environment at the VLA Site

A representative sample of L-band⁷ interference spectra was obtained from the interference protection group at the VLA (Raul Armendariz). This is so-called “SYSQUICK” data but obtained in ASCII format to facilitate quantitative analysis. A discussion of interference monitoring and SYSQUICK data is available from the VLA web site⁸. Briefly, the L-band SYSQUICK data is obtained by pointing the antennas near the celestial north pole and scanning across the wide band, 1.5625 MHz at a time with 256 spectral channels in the correlator and with a 10 second integration time. (Additionally, reading “between the lines” indicates that several integrations and baselines are incoherently averaged to yield the final output spectrum—resulting in a “less noisy” noise floor.) The spectral resolution is about 6 kHz, and the scan covers a bandwidth of 540 MHz from 1220 MHz to 1760 MHz. A plot of some L-band SYSQUICK data is shown in Figure 5.

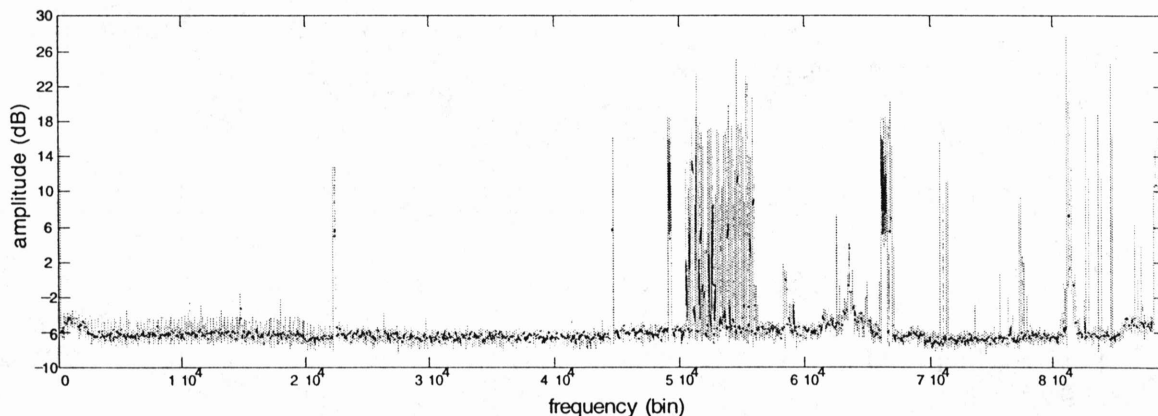


Figure 5 Plot of L-band SYSQUICK data obtained from the interference protection group at the VLA. This covers the frequency range from 1220 MHz to 1760 MHz and there are about ~ 88000 spectral points (bins) in the plot. Data has been cross-correlated and integrated for 10 seconds with a bin bandwidth of about 6 kHz.

At first glance, the data in Figure 5 would seem to indicate that the narrowband interference is extremely powerful compared to the system noise power. As we’ve seen

⁷ A survey of interference in several bands indicates that L-band is the worst case.

⁸ <http://www.aoc.nrao.edu/vla/html/rfi.shtml>

from previous tests in this document, it is the percentage of narrowband interference power to total power *into the initial quantizer* that determines whether secondary harmonics and intermodulation products splatter across the band. Since the SYSQUICK data has been cross-correlated and integrated for 10 seconds, it is necessary to account for the decorrelation of independent system noise in order to be able to determine what the initial quantizer is subjected to.

A MathCad program was written to read in the SYSQUICK data and determine what the percentage of the interference power to the total power in the band is. To do this, the independent noise power before integration must be determined. The attenuation of the independent noise from cross-correlation and integration is:

$$\text{noise atten (dB)} = 10 \cdot \log \left(\sqrt{\frac{\# \text{ freq bins}}{2 \cdot T \cdot BW}} \right) \quad (1)$$

Where T is the integration time and BW is the total bandwidth in Hz⁹. For a bandwidth of 540 MHz, 88992 frequency bins in the SYSQUICK data, and a 10 second integration time, the noise attenuation is -25.4 dB. If this is taken into account in Figure 5, the noise floor should be at about 19 dB to eliminate the noise decorrelation effect. This is also the same (relative) level that the noise floor would be in the auto-power spectrum. Thus, the “true” interference spectrum would have the narrowband interference only about a maximum of 10 dB above the noise floor. This conclusion is roughly consistent with the “W8 L-band RFI monitor” data from the VLA site that is acquired with a spectrum analyser using a cooled receiver on a separate interference detection antenna.

Analysis of the SYSQUICK data of Figure 5 using the MathCad program and taking into account the noise decorrelation effect indicates that only about 0.7% of the total power in the band is due to the narrowband interference. This of course **assumes** that the interference has not decorrelated within the integration time—which should be the case since the antennas are pointing near the celestial north pole. It also **assumes** that in the regions where there is no narrowband interference (e.g. bins 30000 to 40000), there is no correlation. As a check to ensure that equation (1) and the analysis method are correct, the same method was applied to simulation data with 3 narrowband signals (the data shown in Figure 2), with 4 narrowband signals (with data contained in Table 1), and with many narrowband signals and 16384 spectral points. Hanning windowing was used to minimize the main lobe spread¹⁰ and provide decent—although not perfect—ringing suppression. In the 3 narrowband signals case, it is known a priori that 25% of the total power is interference power—the analysis indicates that 27% is interference. In the 4 narrowband signals case, 50% of the total power is interference power—the analysis indicates that 55% is interference. In the many narrowband signal case, 12.3% of the total power is interference power—the analysis indicates that 13.8% is interference. Thus, the method yields good results, giving confidence that the analysis of VLA interference is correct. The MathCad program with VLA SYSQUICK data, 3 tone (1024 spectral channels) data, and multi-tone data are in the appendix.

⁹ This equation assumes Nyquist sampling.

¹⁰ Which, if too wide, skews the results.

Based on the simulation results contained in Table 1, it is probably safe to state that it is possible for the correlator to handle¹¹ any conditions whereby the interference power is less than 25% of the total power into the initial quantizer. Analysis of VLA interference data indicates that this provides about 16 dB (factor of 40) of margin in the design of the system—**assuming** the statistics of the interference environment across 1 GHz (at L-band) is the same as across the 540 MHz that was analyzed. In the event that the interference environment at the VLA site exceeds this margin, the WIDAR design eliminates quantizer splatter for one or two strong interferers, and reduces the splatter if there are three or more strong interferers.

WIDAR Correlator with Simulated VLA Interference

It was suggested (Napier, Perley—private comm.) that it would be interesting to simulate the interference environment at the VLA but with “a couple” of strong tones ~20 dB above the others to represent interference from satellites into the “near-in” sidelobes of the antennas. Several simulations were run with about 40 tones¹² to simulate interference, and the WIDAR correlator was configured for the maximum number of spectral channels of 16384 to approach the SYSQUICK spectral resolution (at least as far as the human eye can tell). These results are contained in Figures 6 through 9.

In Figure 6, sub-band FIR filters with transition band cutoffs of -1.25 dB are used. The SNR degradation region— $\sim\pm 3$ MHz¹³ about the sub-band boundaries with a degradation >20%—contains many spectral points and is evident in the plot¹⁴. In Figure 7, sub-band FIR filters have transition band cutoffs of -12.5 dB. The SNR degradation region is $\sim\pm 0.5$ MHz, does not contain as many spectral points, and is not as evident in this Figure. In Figure 8, a continuum correlation component of $\rho=0.02$ (and with -12.5 dB cutoff filters) has been added to illustrate the amplitude and phase performance with a correlated continuum signal. In this figure, the roll-off at the edges of the wide band is from the software digital SSB mixer in the noise generator. This roll-off appears here and not in the previous two plots because in the noise generator the mixer operates on common (source) signals, but not on independently generated system noise. Finally, Figure 9 uses the same simulation parameters as Figure 7, but with one additional satellite interference tone for an interference power to total power ratio of 17%.

¹¹ i.e. not create additional quantizer-generated interference across the band.

¹² Tones here are pure sine-waves—trying to generate many very narrow tones with FIR filtering in the noise generator was found to be too computationally intensive.

¹³ That is, if one takes the full scale bandwidth as 2 GHz.

¹⁴ Or, at least was evident in the original MathCad plot before inclusion in this document!

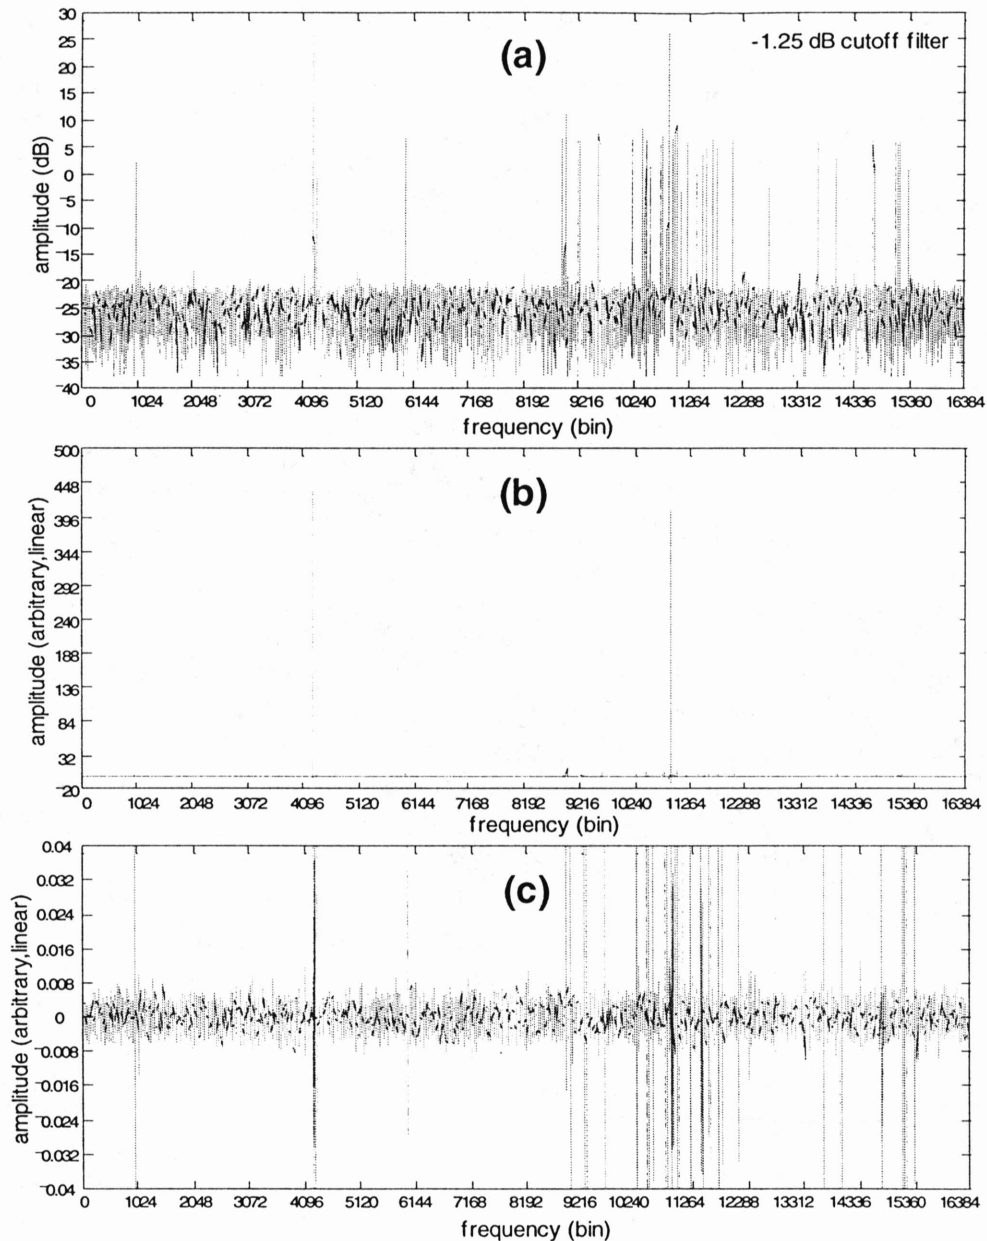


Figure 6 WIDAR correlation of simulated interference environment with 2 tones that are 20 dB higher than the other “nominal” tones to simulate satellite interference in the near-in sidelobes of the antennas. This simulation uses a 4-bit, 15-level initial quantizer (interleaved, algorithmic A/D) followed by bit shift translation to 3 bits, 8 levels. There are 16384 spectral points and the sub-band FIR filters have a **-1.25 dB cutoff** at the sub-band boundary. Hanning windowing using window method #3 [1] is used. Five-level fringe stopping with residual discrete phase offsets and post-correlation removal is used. The total system power (arbitrary units) is 100, the total “nominal” tone power is 1.86, and each of the strong tones has a power of 6.1. The interference power is thus ~12% of the total power in the band. 100 dumps of 10^7 samples each were integrated. In the figure, (a) is a log (dB) plot, (b) is a full-scale linear plot, and (c) is a magnified scale linear plot of the real part of the cross-power spectrum. The additional noise at the sub-band boundaries is evident in (a) and (c).

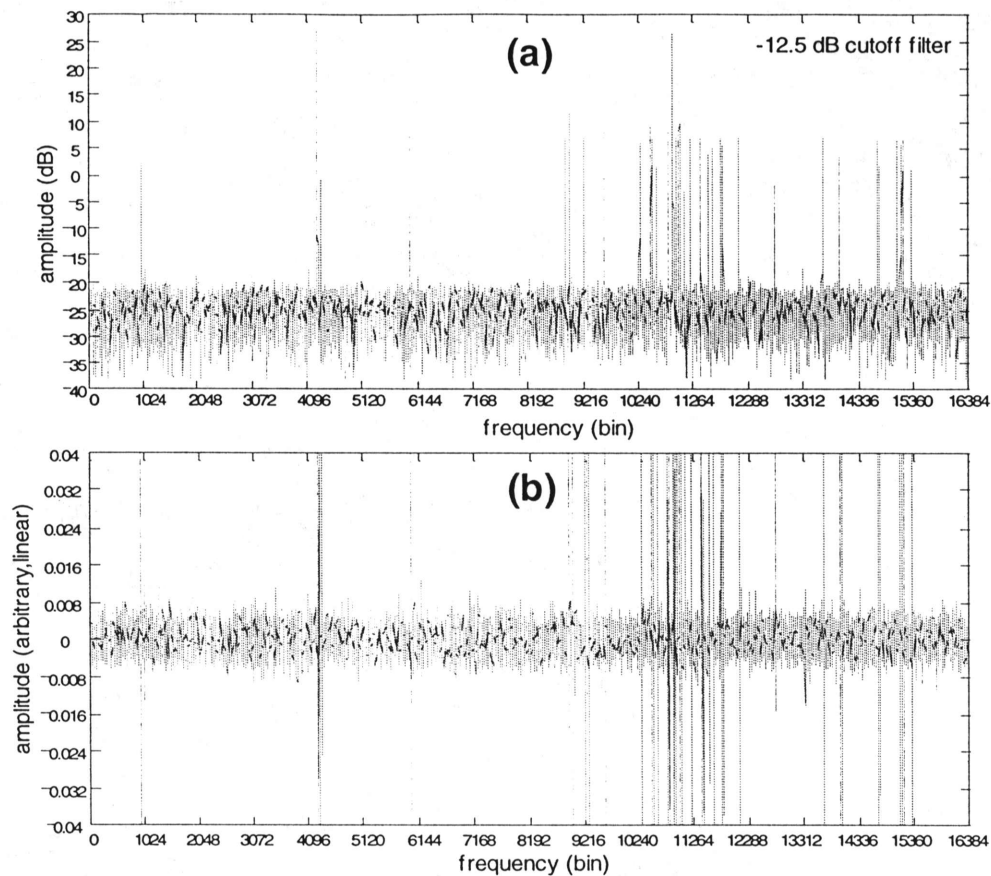


Figure 7 WIDAR correlation of simulated interference environment with 2 tones that are 20 dB higher than the other “nominal” tones to simulate satellite interference in the near-in sidelobes of the antennas. This simulation uses a 4-bit, 15-level initial quantizer (interleaved, algorithmic A/D) followed by bit shift translation to 3 bits, 8 levels. There are 16384 spectral points and the sub-band FIR filters have a **-12.5 dB cutoff** at the sub-band boundary to reduce noise near the sub-band boundary. Hanning windowing using window method #3 [1] is used providing reasonable, but not perfect ringing suppression. Five-level fringe stopping with residual discrete phase offsets and post-correlation removal is used. The total system power (arbitrary units) is 100, the total “nominal” tone power is 1.86, and each of the strong tones has a power of 6.1. The interference power is thus ~12% of the total power in the band. 100 dumps of 10^7 samples each were integrated. In the figure, (a) is a log (dB) plot and (b) is a zoomed-scale linear plot of the real part of the cross-power spectrum. The additional noise at the sub-band boundaries does not appear to be as evident as in Figure 6 due to the smaller region of SNR degradation.

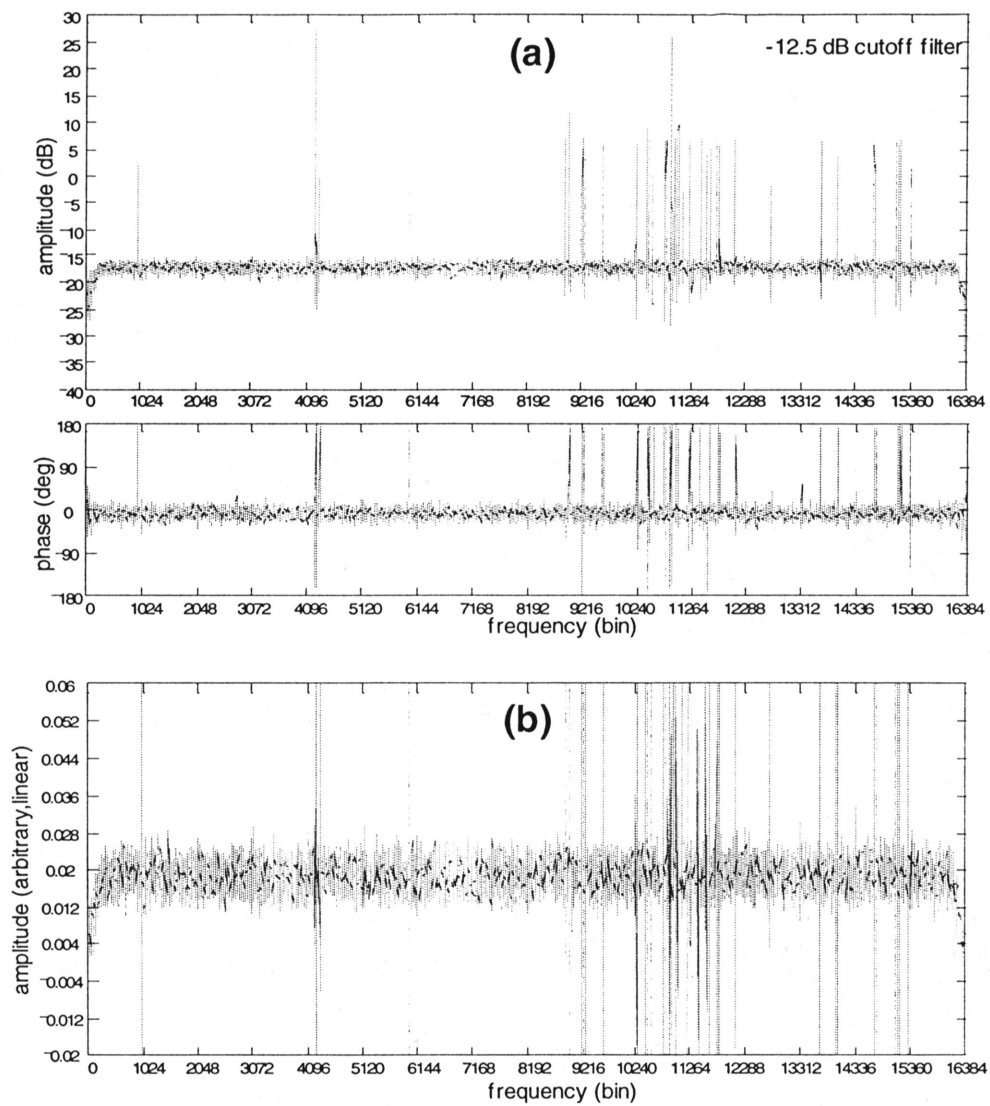


Figure 8 WIDAR correlation with identical conditions to those in Figure 7, but with a correlated continuum component of $\rho \approx 0.02$. In (a) the amplitude (dB) versus frequency and phase (deg) versus frequency is plotted. In (b), the amplitude is plotted on a magnified linear scale for the real part of the cross-power spectrum. Large phase discontinuities where there are interference tones are from ringing due to the high dynamic range and relatively poor ringing suppression provided by Hanning windowing.

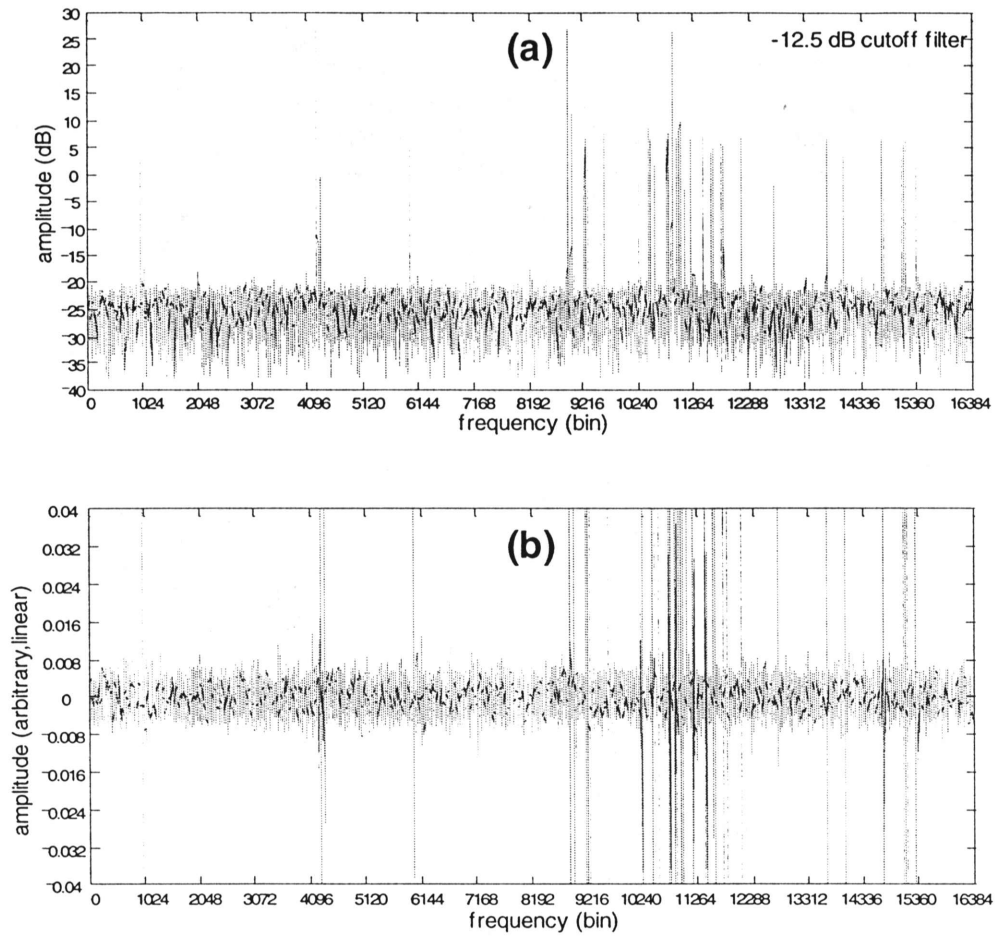


Figure 9 WIDAR correlation with identical conditions to those of Figure 7 but with three strong satellite interference tones for a interference to total power ratio of 17%. (a) is a plot of amplitude (dB) versus frequency, and (b) is a linear scale plot of zoomed-in amplitude versus frequency of the real part of the complex cross-power spectrum.

Conclusions

This memo has investigated the interference robustness of the WIDAR correlator. Interference robustness¹⁵ in this investigation is defined as the ability of the correlator to obtain the cross-power spectrum in the presence of powerful narrowband interference sources but without any secondary harmonic or intermodulation products. Eliminating these secondary products minimizes spectral contamination and aids post-correlation interference mitigation. Frequency shifting that is necessary for sub-band transition band anti-aliasing has the positive side effect of causing quantizer-generated harmonics and many intermodulation products to decorrelate. However, there is a fundamental limitation that prevents some intermodulation products from decorrelating because they are shifted in frequency precisely by the same amount that is removed in the correlator. This condition only appears if there are three or more narrowband interference signals. Nevertheless, the WIDAR correlator generally significantly reduces the number of secondary products in extreme high-power narrowband interference cases.

A method was discovered and tested that potentially overcomes the limitation in the ability of the correlator to eliminate narrowband *harmonics*. This limitation is due to the harmonics of the coarse digital mixer (fringe rotator) in the correlator “fringe-stopping” the quantizer-generated harmonics of the interference. This method involves inserting a changing phase offset into the correlator’s fringe rotator and removing it with a large number of bits after correlation. It is expected that if this function is realized, it can be implemented in Baseline Board hardware so that front-end and back-end hardware need not even know of its existence.

A number of tests were performed to quantify the performance of the initial quantizer and it was found that maximum spectral dynamic range (minimum secondary product amplitude) is obtained when all available levels, afforded by a given number of bits, are used. For example, it was found that the spectral dynamic range of 3-bit/8-level quantization was equivalent to 4-bit/15-level quantization—although the sensitivity loss for 3-bit/8-level was ~3.5% and ~1.5% for 4-bit/15-level. It was also found to be important to ensure that the effective number of bits (ENOB) was as close to the actual number of bits as possible—implying that a well designed 3-bit/8-level precision quantizer is better than a 4-bit quantizer with significant decision level errors.

A survey of available interference data from the VLA web site indicates that the worst-case interference environment is at L-band. Some “SYSQUICK” L-band interference data was obtained from the interference protection group at the VLA in ASCII format so that some quantitative analysis could be performed. After accounting for the drop in the noise floor due to cross-correlation of independent system noise, it was found that the interference to total power ratio into the initial quantizer was ~1%. This is well within 3-bit/8-level initial quantization capability and affords a margin of a minimum of about 16 dB in the interference robustness of the system. These results indicate that 3-bit/8-level

¹⁵ Interference robustness also includes the correlator’s ability to be immune to the amplitude modulating effects of time variable interference. This was discussed in [1] but was not discussed in this memo.

initial quantization should be sufficient for the system which will reduce the cost of the Fiber Optic Transmission System over that of 4-bit initial quantization. Some WIDAR correlations of a simulated L-band interference environment with a few tones 20 dB stronger than seen in the SYSQUICK data were performed. These tones represent those that might come from satellite carriers into the near-in sidelobes of the antennas. These simulations yielded satisfactory results.

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, An Analysis of the Effects of Phase Dithering in a Lag-based Fringe-Stopping XF Correlator, NRC-EVLA Memo# 002, May 26, 2000

