

Simulation Tests of Phasing Subsystem Signal Processing in the WIDAR Correlator for the EVLA

NRC-EVLA Memo# 008

Brent Carlson, November 7, 2000

ABSTRACT

A crucial requirement of the correlator for the EVLA is the ability to produce a coherent sum of the real-time outputs of all of the antennas—the so-called “phased-VLA” output. The WIDAR design is well suited to performing this function on the wide bands involved since it breaks up the wide bands into narrower sub-bands that can be directly processed by digital phasing hardware. The WIDAR design has the additional benefit of processing the phased-VLA data in completely digital fashion—resulting in better performance and stability than methods which require an intermediate conversion to analog and then back to digital. In NRC-EVLA Memo# 001 [1] a basic concept for the phasing signal processing as well as a potential hardware architecture was presented. However, until now, there have not been any simulations of the signal processing to quantify the phased output performance or hardware word-size requirements. This memo presents the results of several simulations that have been performed to quantify phasing hardware performance and demonstrate that the signal processing operates as originally envisioned. This memo also defines a possible mechanism for feeding the phased output back into the correlator for real-time auto or crosscorrelation processing.

Background

The WIDAR correlator for the EVLA must be able to produce real-time coherent summed outputs from all antennas or sub-arrays of antennas. These outputs will be used when maximum single element sensitivity is required for activities such as pulsar searching and analysis and high sensitivity VLBI. In [1] the basic signal processing was described and it will be briefly repeated here. Figure 1 is a simplified block diagram of the phased-output signal processing. Phasing of each sub-band is performed by removing the frequency shift introduced by the Local Oscillator and the earth rotation fringe phase¹ with a complex digital mixer. Complex data from all antennas for each sub-band is then added before shifting the phase of the quadrature component with a -90° phase shifting Hilbert FIR filter. Complex addition before phase shifting minimizes the number of Hilbert FIR filters—saving cost and considerable circuit board area since these filters must use a large number of bits (~8).

¹ If desired—it is possible to remove the fringe phase by modifying the phase of the Local Oscillator, but since the digital mixer is available, it is probably the best place to do it.



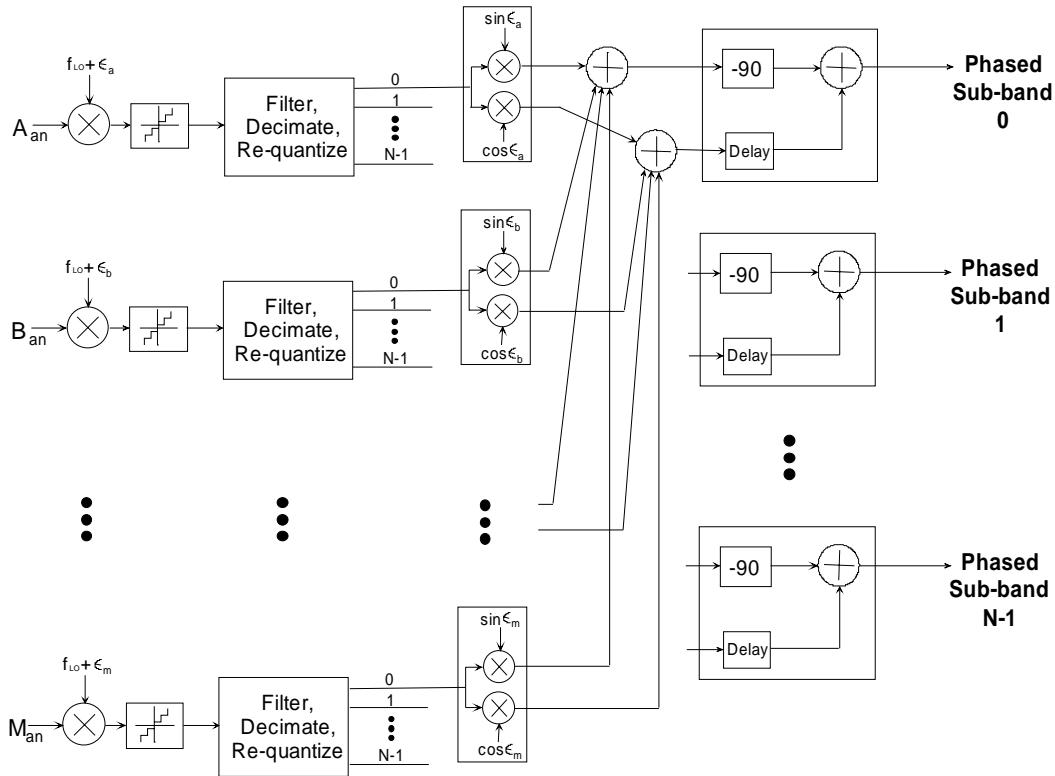


Figure 1 Simplified diagram of phasing signal processing. Each antenna’s Local Oscillator is offset in frequency by some small different ϵ as required by the WIDAR technique. The wideband signal is sampled and filtered where, for correlation, it goes to the downstream cross-correlator but, for phasing, goes to digital single sideband mixers. These mixers remove the small frequency shifts and, if desired, the earth-rotation fringe phase. Complex data from mixers from all (sub-array) antennas are added for each sub-band. Finally, the complex data is coherently added after the quadrature component is phase shifted by -90° with a digital Hilbert FIR filter (the in-phase data “Delay” matches the quadrature data delay through the FIR).

Simulation Methods

The WIDAR correlator simulator [1] was upgraded to include phasing simulation code that “tapped” into the data out of the sub-band FIRs after requantization. Two different complex mixer configurations were built and tested with different parameters. These are shown in Figure 2. Two different mixer configurations were tested because due to speed, cost, and circuit board area limitations, it is most likely necessary to include this mixer circuitry for ~4 antennas in a single (Xilinx Virtex) FPGA. Figure 2 (a) fundamentally has a smaller lookup table (LUT) than (b), but it must be followed by a multiplier. If the multiplier is not too large, it should be possible to implement in an FPGA. Figure 2 (b) has a larger LUT but requires no multiplier. This topology is normally chosen for this kind of function, but if the LUT is too large, it must be off (the FPGA) chip and, as mentioned, is not cost-effective to implement.

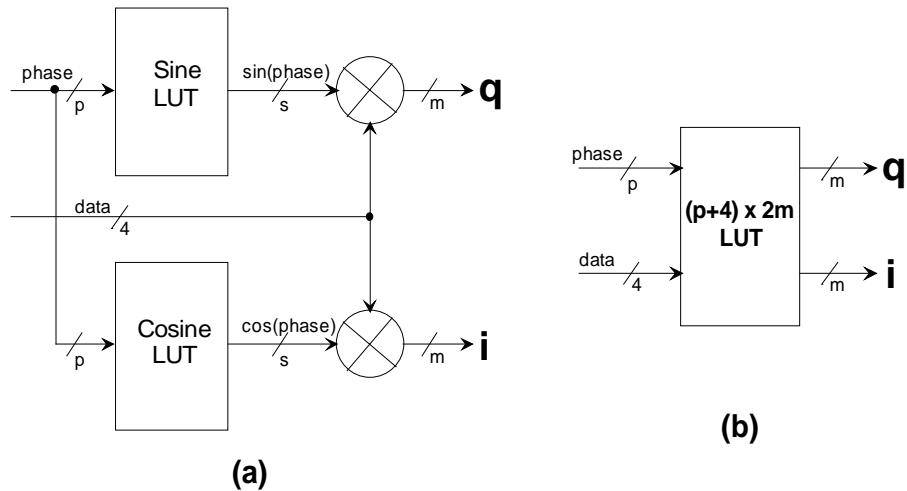


Figure 2 Alternative complex mixer architectures that are built into the simulator. In (a), p bits of phase are addresses into RAM lookup tables (LUTs) to generate s bits of sine and cosine quantities. This data is multiplied in a hardware multiplier by the 4 bits of sampled data. m bits of the in-phase and quadrature multiplier results are retained. In (b), phase and data are addresses of a larger RAM lookup table where the result is now the m -bits of in-phase and quadrature data.

Once the complex data is obtained from each X and Y station, it is added and one additional bit is retained. In the simulator, only two antennas are added and so retaining just one additional bit is sufficient. In the real system at this stage up to 40 antennas are added together and so one additional bit may not be sufficient (three are required for just noise). It is important to keep the number of bits here to a minimum ($< \sim 8$) so that the downstream Hilbert FIR does not have to be too large. To provide programmable flexibility, it may be a good idea to retain many more bits in the adder output, and then allow selection of an 8-bit window of data to go to the Hilbert FIR. In the worst case, there is only a small requantization loss and using the same window for the in-phase and quadrature data ensures that the Hilbert FIR and subsequent addition performs properly.

The Hilbert FIR is constructed in the simulator as an array of LUTs with one LUT for each tap. For efficiency, the input word size to the LUT should be ≤ 8 bits (although the simulator allows testing more or less than 8 bits) requiring a $256 \times N$ memory for each LUT. This is well within the Xilinx Virtex FPGA's capability for a reasonable number of taps. Since there is only one Hilbert FIR for every phased output (and there are going to be five phased outputs on every Phasing Board [1]), it is possible to dedicate a large FPGA to each FIR without a large cost impact. In the simulator, the output of the LUTs are simply added together—in the actual implementation this function would be performed with an adder tree. The FPGA must also contain the matching delay line that is $\frac{1}{2}$ the length of the FIR for the in-phase data. It is important that the Hilbert FIR not have any relative gain compared to the unfiltered in-phase data otherwise unwanted side-band cancellation will not occur.

Figure 3 is a block diagram of the phasing signal processing blocks along with auto-power spectrum plots of data at various points in the signal path. The test signal in this case is a strong tone and Gaussian noise. The increase in amplitude at the edges of the band is due to aliasing of the sub-band transition band (remember that this is the auto-

power spectrum of just one *sub*-band). The final output sees a decrease in the amplitude at the edge of the band due to the band edge roll-off from the Hilbert FIR filter. Also, note that the final output spectrum tone amplitude is double that of the individual inputs—this is the desired response of the system.

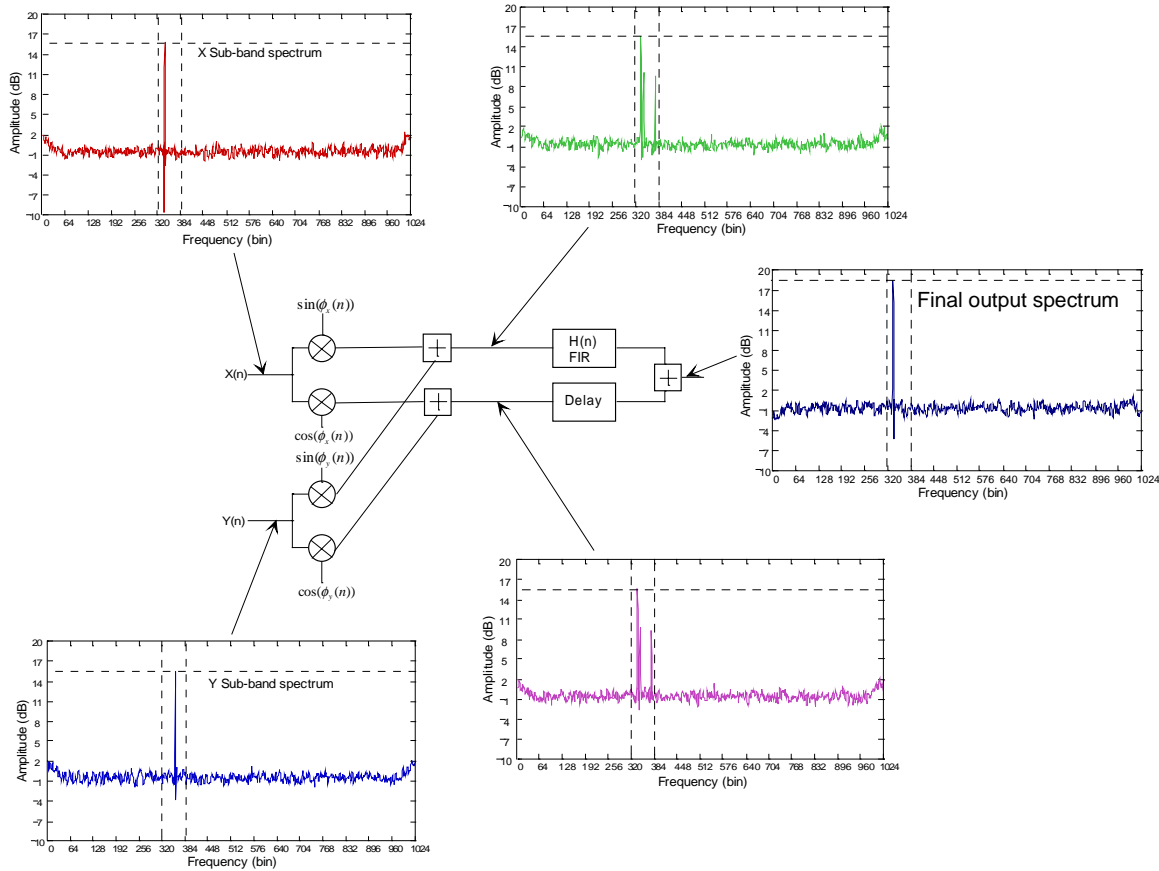


Figure 3 Phasing system signal processing block diagram where two antennas (X and Y) are phased together. The test signal is a single tone and Gaussian noise. Note that the final output tone amplitude is double (3 dB higher than) the individual inputs—as expected. The summed output of the complex mixers contains unwanted sidebands which are cancelled by the Hilbert (**H(n)**) FIR and final summation. The quantizer is not shown.

Complex Mixer Performance

A crucial component of the phasing system is the complex mixer. Since there is one mixer for every antenna and for every sub-band, it is important to optimize the design of the mixer so that hardware is minimized. Some work was done to try to find an optimized digital mixer function that would use the fewest number of phase bits and the fewest number of sine/cosine LUT output bits while still providing sufficient performance. Both of the configurations of Figure 2 were investigated. Figure 4 contains a number of plots of the mixer function amplitude versus phase and frequency for the configuration of Figure 2 (a) with different numbers of phase bits and output amplitude levels. The equation used to generate the sine/cosine output values is simply:

$$LUT_{sin_i} = \text{round} \left[A \cdot \frac{N_{levels}}{2} \sin \left(\frac{i}{N_{phase\ steps}} \cdot 2\pi \right) \right] \quad (1)$$

Where i is the phase step and ranges from $(0 \dots N_{phase\ steps} - 1)$, N_{levels} is the maximum number of amplitude levels, and A is the amplitude that is set to use all levels (determined empirically in this test by minimizing the mixer’s harmonic spectral content).




Figure 4 Plots of amplitude versus phase and frequency for several digital mixer functions for the LUT of Figure 2 (a). Of particular interest is (b): the size of the LUT as well as the subsequent multiplier (and multiplier result) should be manageable while providing decent performance of -30 dB harmonic attenuation. (c) is also interesting since there could be enough RAM on the Xilinx FPGA for a 256×4 LUT with the same multiplier size as (b).

Figure 5 contains two plots of the spectrum of the output of the LUT from Figure 2 (b). Figure 5 (a) is with 5 phase bits, and Figure 5 (b) is with 8 phase bits. The plots were obtained with the indicated number of phase bits, 256 output levels, and a pure sine-wave input (for data—so that only the effects of the quantized mixer phase could be seen). More phase levels yield lower harmonic amplitudes, but with larger memory requirements. With 5 phase bits and 4 data bits, an 8-bit output requires a 512×8 -bit RAM. If 8 phase bits are used, the harmonic amplitude drops to ~ -48 dB (as predicted) but a 4096×8 -bit LUT is required. It is not completely unreasonable to expect that

eight² 4096 x 8-bit LUTs could be implemented in an appropriate Xilinx Virtex-E device since these devices have many distributed RAM bits available.

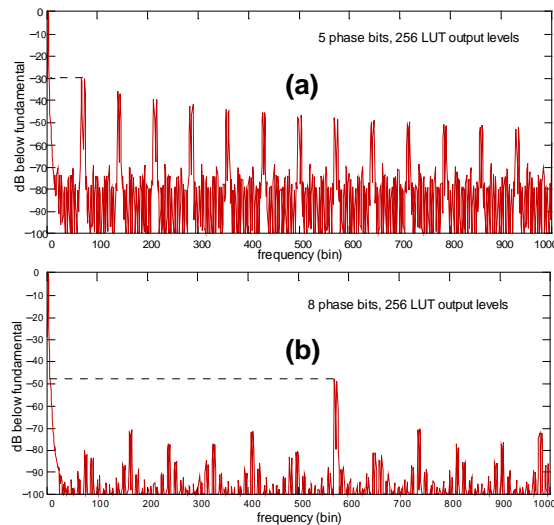


Figure 5 Spectral analysis of output of the LUT (Figure 2 (b) configuration) which contains the multiplication of the digital mixer function with a floating-point sine wave. (a) is with 5 phase bits and (b) is with 8 phase bits. In both cases, the output of the LUT is 8 bits.

It seems reasonable to be able to achieve a minimum of 30 dB of harmonic suppression in the complex mixer. More suppression can probably be achieved since the memory capacity of currently available FPGAs with on-chip SRAM seems sufficient.

A comparison of the performance of the two mixer topologies is shown in Figures 6 and 7. In Figure 6, the narrowband line is 7.5% of the total power in the *sub*-band and in Figure 7, the narrowband line is 63% of the total power in the *sub*-band. The top of each figure shows the auto-power spectrum with and without the narrowband line and the bottom of each figure shows a plot of the difference between the top spectra. In both figures, about 35 dB of spectral dynamic range is achieved even though the mixer harmonics are only 30 dB down from the fundamental. In the bottom plot of Figure 6, the harmonics of the mixer are not evident above the noise floor. A longer integration under similar conditions as Figure 6 (a) is shown in Figure 8, indicating that the mixer harmonic peaks do indeed appear. Thus, the spectral dynamic range is limited by mixer harmonic attenuation even in the presence of relatively weak signals³. The auto-power spectrum after requantization to 4 bits/15 levels is virtually identical to the plots shown.

It is important to note that Figures 6 and 7 show the autocorrelation spectral dynamic range. The mixer harmonics will not appear if the phased output is crosscorrelated with data from another antenna that does not contain such a digital mixer—as would be the case for phased-VLA VLBI.

² For four complex mixers within one 4-station mixer chip [1].

³ i.e. Those that have a smaller spectral dynamic range relative to the mixer dynamic range (mixer harmonic attenuation).

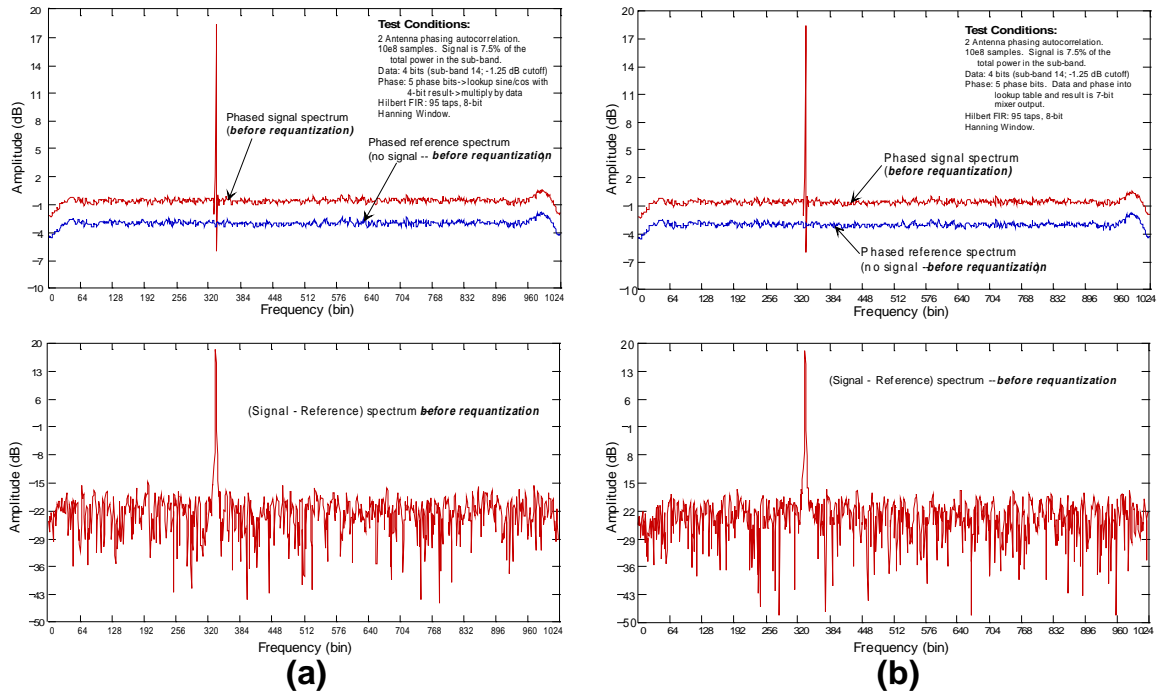


Figure 6 Auto-power spectrum of phased output—before requantization—with a narrowband test tone that has a power of ~7.5% of the total power in the sub-band. The top plots contain the raw spectrum with and without the tone, and the bottom contains the difference between the two. The raw reference spectrum amplitude (in the top plots) has been offset so it can be discerned from the signal spectrum. About 35 dB of spectral dynamic range is achieved, with the range limit imposed by the noise floor/integration time. (a) is for the mixer topology of Figure 2 (a) with 5 phase bits and 4 amplitude bits. (b) is for the mixer topology of Figure 2 (b) with 5 phase bits.

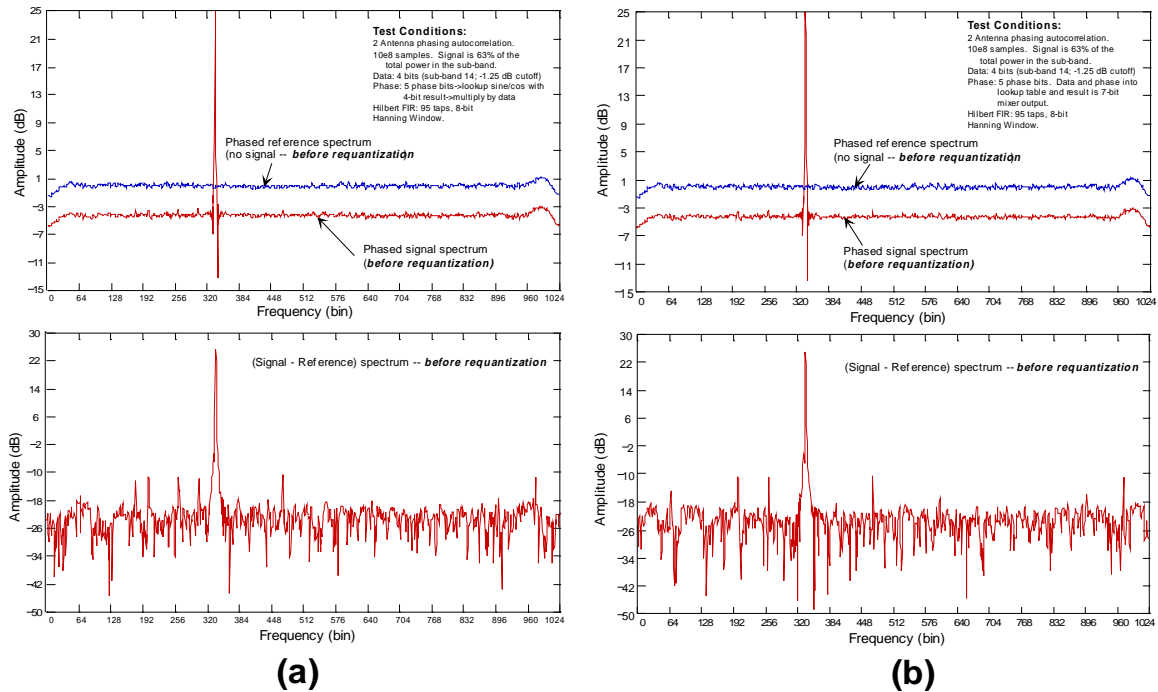


Figure 7 Same test as in Figure 6 only now the narrowband line is 63% of the total power in the sub-band. The harmonics generated by the coarse mixer function are now clearly visible in all plots.

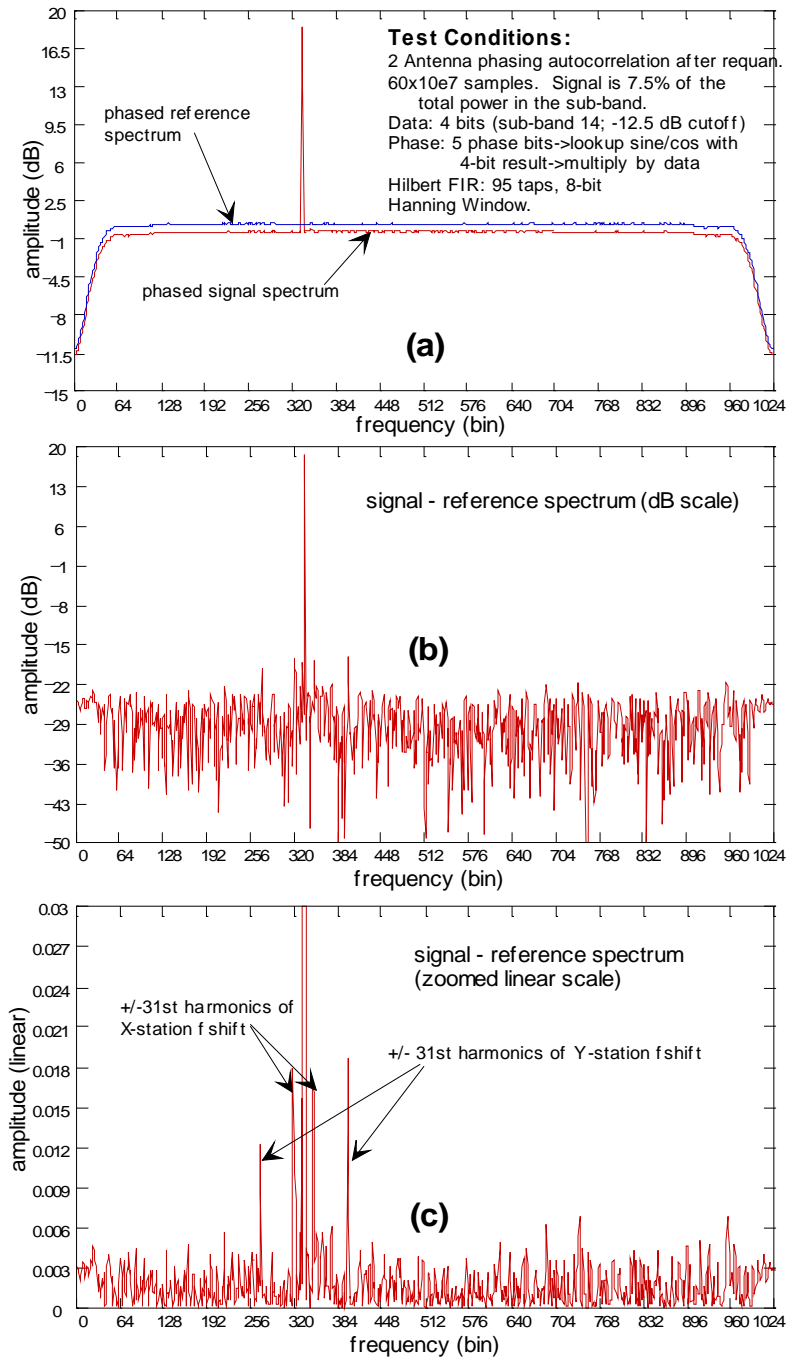


Figure 8 Longer integration under similar conditions used for Figure 6 (a) (a -12.5 dB cutoff sub-band filter was used and the frequency shift is smaller than in Figure 6). (a) is a raw plot of the auto-power spectra with and without the test tone. (b) is the difference between the two on a dB scale, and (c) is the difference on a magnified linear scale. Mixer harmonic-generated lines (31st harmonics—see Figure 4 (b)) are now readily apparent indicating that the spectral dynamic range is limited by the attenuation of the mixer harmonics. That is, mixer harmonic-generated lines will appear even in the presence of a narrowband signal with an auto-power spectral dynamic range that is less than the mixer harmonic attenuation. This indicates that, for robustness, maximum mixer harmonic attenuation should be a design goal.

Hilbert FIR

The Hilbert FIR filter introduces a broadband -90° phase shift to the quadrature data after complex mixing and summation of all antennas that are to be phased together. As mentioned previously, introducing the Hilbert FIR after complex summation means that only one Hilbert FIR is required for each sub-band phased output. Thus, it is not a major cost item but it is still necessary to study how FIR filter parameters—length and number of bits—affect its performance.

The equation used to generate the floating-point tap coefficients is well known and is as follows [2]:

$$h(n) = \begin{cases} \frac{2 \sin^2(\pi n/2)}{\pi n}, & n \neq 0 \\ 0, & n = 0 \end{cases} \quad (2)$$

Where $n=0$ at the center of the filter and takes on +/- values on either side of center. Note that every other tap is 0 and the coefficients are symmetrical about the center but with opposite sign. This equation has coefficients that guarantee that it has unity gain. (In the simulation tests that were run, the Hilbert coefficients were also Hanning windowed to reduce ringing effects on strong narrowband signals—although admittedly the exact benefit of doing this has not been properly quantified). These coefficients are scaled by the word size at the output of the complex mixer and antenna summer so that the data that comes out of the filter does not have any gain relative to the in-phase data that experiences only a matching FIR delay in the data path. This ensures that sufficient unwanted sideband cancellation occurs. Since the maximum value of $h(n)$ is about 0.64— or $(2/\pi)$ —and since no gain is allowed, the maximum use of the available number of levels is not realized. However, this does not seem to be a problem because testing with a 256 x 8-bit LUT for each tap yields results indiscernible from those which either have more address bits (larger word size out of mixer and summer) or more data bits (which, by itself, must have a compensating gain introduced before summation with the in-phase data). Since the nominal LUT seems sufficient, only the effect of the FIR filter length is investigated here.

Figure 9 contains several plots of the auto-power spectrum of the phased output with different FIR filter lengths, with 256 x 8-bit LUTs, and with two antennas “phased-up”. The test signal contains a number of tones—with symmetrical frequencies about the center of the band—and uncorrelated system noise. Most of the rolloff at the edges of the band is due to the sub-band FIR filter which, to minimize aliasing distortion, was chosen to have a sub-band boundary cutoff of -12.5 dB. In all cases, the mixer topology of Figure 2 (a) with 5 phase bits and 4 amplitude bits out of the sine/cosine LUT was used. Also, 4-bit/15-level sub-band data requantization was used.

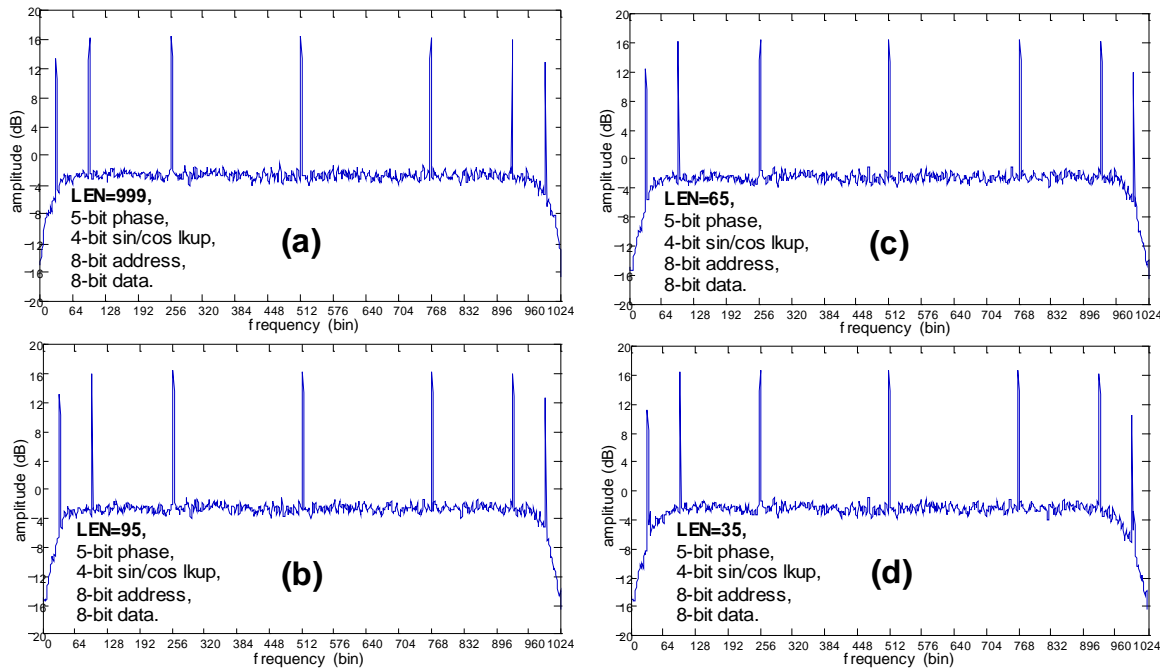


Figure 9 Auto-power spectra of the phased output with varying numbers of Hilbert FIR filter taps. In (a), 999 taps were used, and this output could be considered to be almost perfect (although because of the limited word size, all of the taps are not used). There is not much change from 95 (b) to 65 (c) taps, but in the 35-tap case (d), the band-edge tones suffer about 3 dB attenuation and have poor unwanted sideband attenuation.

With 256x8-bit LUTs, and using the knowledge that every other tap is zero, a 95-tap Hilbert FIR filter will require $256 \times 8 \times 95 / 2 = 97280$ RAM bits. A mid-range Xilinx Virtex FPGA (XCV600E) contains 294,912 “BlockRAM” bits requiring about 1/3rd of the FPGA for LUTs. This should leave plenty of chip logic for the delay lines, adder tree, and requantizer. Thus, it would seem that the implementation of the Hilbert FIR should not be a problem even in a mid-to-high capacity FPGA.

Sub-band Delay Tracking

As the station-based delay error (from digital delay tracking at the original sample rate) changes, each sub-band will see a residual phase slope and a changing phase [1][3]. This changing phase (a function of residual delay and sub-band) is tracked by offsetting the phase of the fringe rotators in the *correlator* and by offsetting the phase of the complex mixer in the phasing hardware. This function was tested in the simulator and the results of the test are shown in Figure 10. The test demonstrates that tracking the residual phase by offsetting the phase of the complex mixer accordingly operates as expected.

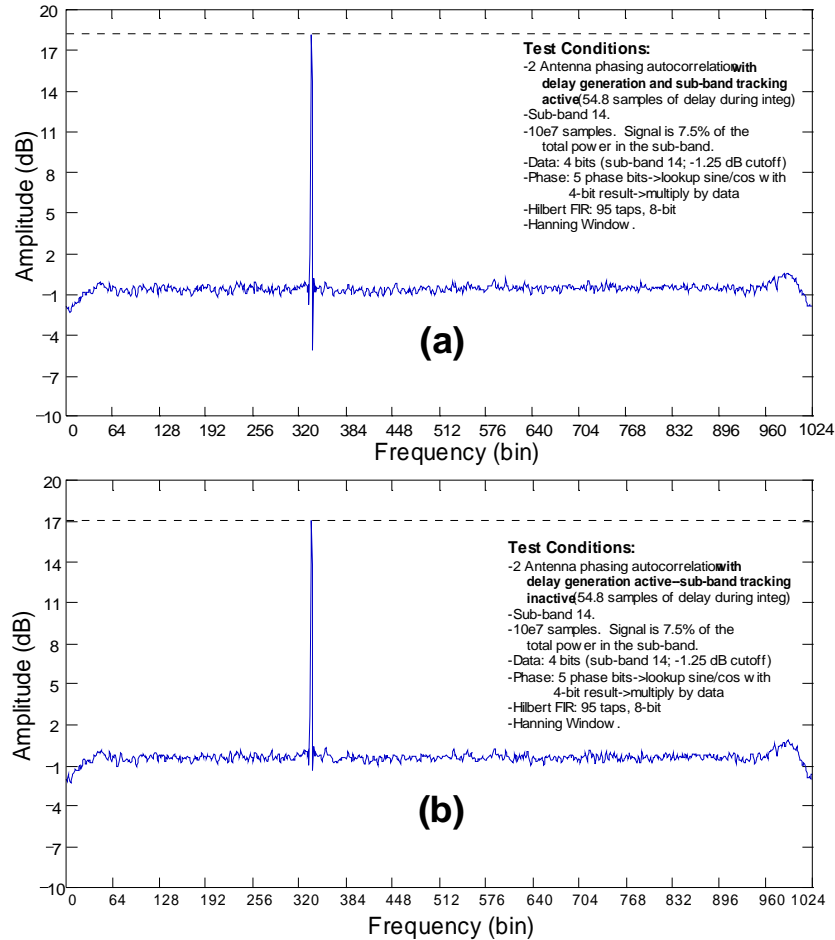


Figure 10 Auto-power spectrum plots of phasing signal processing tests with residual delay phase tracking. The test signal consists of a tone and broadband uncorrelated noise. Sub-band 14 was used so that there would be a large phase swing as the station delay error oscillated between ± 0.5 samples. The test in (a) is with delay/phase tracking in the digital mixers active and the test in (b) is with a delay error in the signal, but *no tracking* in the complex mixer. The tone amplitude in (b) is down about 1.5 dB—equivalent to about a 30% coherence loss.

Cross-correlation Testing

One of the main uses of the phased-VLA output is to be able to use it as a high sensitivity VLBI station. Thus, it is useful to look at the crosscorrelated phased output to study what artifacts may be present from upstream digital processing. The simulator was modified so that it could crosscorrelate two phased outputs—where each output phased-up the data from two antennas. The nominal phasing complex mixer topology of Figure 2 (a) was used with 5 phase bits, 4 LUT output bits, and 4 sub-band data bits. The cross-correlator was a 4-bit correlator with a 5-level fringe rotator. A residual phase rate was introduced into the phasing complex mixers⁴ which was taken out in the correlator. This is equivalent to the actual behaviour using real antennas if the phase center of each group of

⁴ i.e. the complex mixer frequencies were offset so that they did not completely stop the frequency shift.

antennas is different—although in the nominal case, the second “group of antennas” is some other VLBI antenna. Figure 11 contains a number of plots of auto-power and cross-power spectra for different input signals as noted. An important conclusion here is that digital mixer-generated harmonics of narrowband signals crosscorrelate and the cross-power spectral dynamic range is limited by the mixer harmonic attenuation. This is indeed a well-known effect in VLBI and is why XF VLBI correlator chips employ a coarse digital mixer in one data path instead of both. These mixer-generated harmonics will not appear in the cross-power spectrum if the second antenna does not include digital mixer circuitry—normally the case where the phased-VLA is used for VLBI.

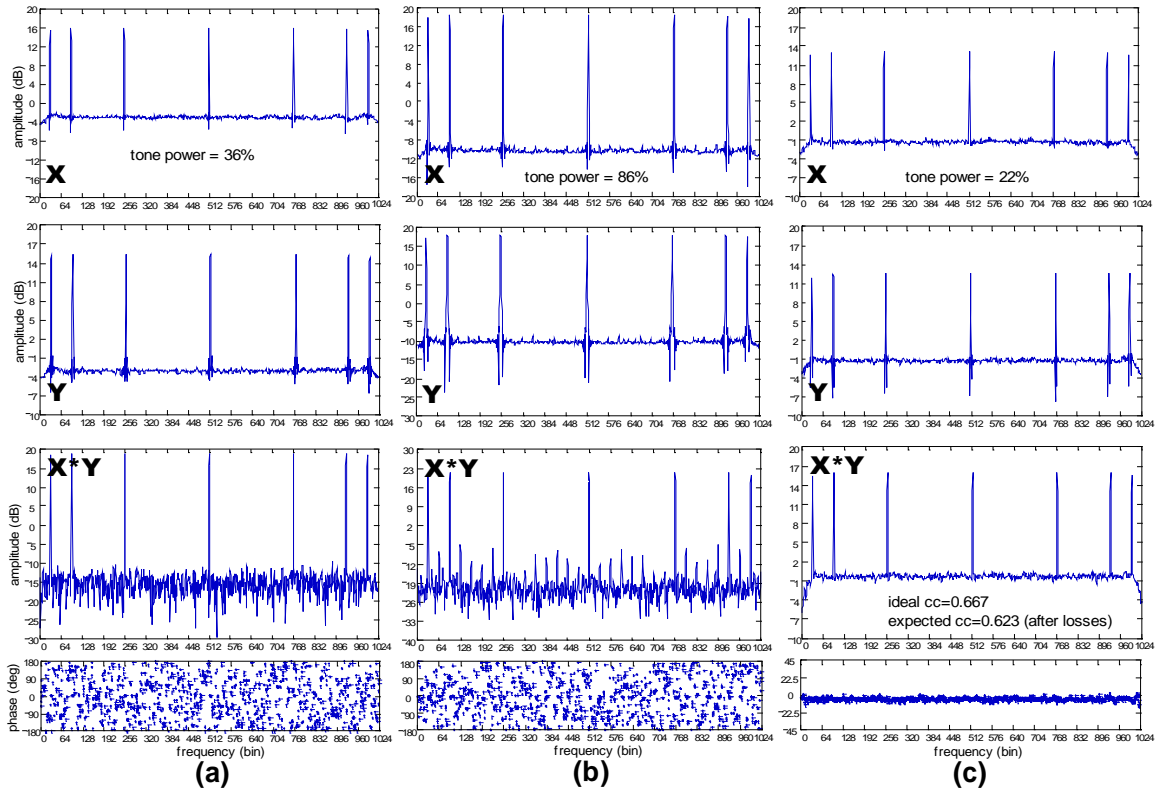


Figure 11 A number of plots of auto-power and cross-power spectra for three different test cases, (a), (b), and (c). In (a) the auto-power (“X” and “Y” plots) and cross-power (“X*Y”) plot show no signs of mixer harmonic tones—although a longer integration should reveal them. In (b), the auto-power and cross-power mixer harmonic tones are clearly evident demonstrating that these do not somehow decorrelate during cross-correlation. However, these mixer-generated tones will not correlate if the second antenna does not include the digital mixer function (as would normally be the case if the phased output is used for VLBI). In (c), a continuum source signal was added which, on a single antenna basis, should yield a correlation coefficient of 0.5. Since two antennas are phased-up and then crosscorrelated, the expected continuum correlation coefficient is 0.667 (since the source noise adds *coherently* and the system noise adds *incoherently*). The expected correlation coefficient after three 4-bit quantization losses and a 5-level fringe rotation loss is 0.623 (~93% efficiency). A WIDAR correlation⁵ (to eliminate the amplitude effects of the tones) of the sub-band data yielded the expected correlation coefficient. In all cases, the mixer topology of Figure 2 (a) with nominal parameters was used along with a 95-tap Hilbert FIR utilizing 256 x 8-bit LUTs. The sub-band FIR filter cutoff was -1.25 dB.

⁵ In software, anything is possible! Although, this could be possible in the implemented system if desired.

Feeding the Phased Output into the Correlator

In the action items section of [4], it was requested that NRC investigate how phased output data could be fed back into the correlator so that autocorrelation and, if possible, crosscorrelation processing could be performed. This section aims to address this action item.

In [5] an extensive investigation shows how a VLBA correlator could fit within the planned EVLA correlator. Part of the investigation established the need for the correlator to have VLBI capability, and be able to process analog sampled bands from VLBI antennas with less than the maximum 2 GHz bandwidth (in appropriate decrements). That investigation, and the minor design changes coming from it to support a VLBA correlator, all work to support auto or crosscorrelation processing of phased output data: phased output data is simply fed back into the correlator via a (dedicated or switchable) Station Board. The phased output data (and its associated TIMECODE) will be delayed a small amount by the phasing digital circuitry—the lion's share of which is from the Hilbert FIR filter delay. This delay can easily be compensated for with the geometrical delay memory on the Station Board.

If 16 polarized sub-bands are phased-up⁶, then all of them can be accommodated in the correlator with one Station Board for autocorrelation processing. (It will be necessary to design a Station Board daughter board [1] specifically for the purpose of getting phased data into the Station Board.) If it is desired to crosscorrelate one phased sub-array with another phased sub-array, then one Station Board can be used as long as the 'X' and 'Y' sub-array data is assigned to different polarized inputs (because sub-bands within the same polarization cannot be correlated). More Station Boards would be required to be able to crosscorrelate more than two phased sub-arrays. Currently though, there is no need or plan to crosscorrelate sub-arrays.

To summarize, with phased data fed into *one* Station Board, it will be possible to:

- Perform sub-band autocorrelations on phased sub-band data for up to 16 polarized sub-bands.
- Crosscorrelate phased data with antenna data fed into other Station Boards as would be the case with real-time VLBI.
- Crosscorrelate two phased sub-arrays. If all polarization products are required, then only 1/2 of the 16 phased sub-bands can be accommodated.

Conclusions

This memo has presented simulation results that demonstrate that the phasing system signal processing for the WIDAR EVLA correlator operates as originally envisioned. Various aspects of the digital single sideband mixers such as complex mixer word size

⁶ Requiring 32 Phasing Boards—the original plan [1] was to deliver 8.

and Hilbert FIR filter length were investigated. The important conclusion was reached that the autocorrelation spectral dynamic range is limited by digital mixer harmonic attenuation even if relatively weak narrowband signals are present. A nominal digital mixer with -30 dB harmonic attenuation was simulated to support this conclusion. More mixer harmonic attenuation can be obtained with a more precise representation using more bits—with an 8-bit mixer, better than -48 dB harmonic attenuation occurs. It may be possible to implement 8-bit mixers, given the size of currently available FPGAs. Certainly, 8-bit mixers should be a design goal. Cross-correlation tests between two phased data streams also verify that this dynamic range limit is present and that the mixer harmonic-generated signals do not decorrelate. This is a well-known effect that influences the design of XF VLBI correlators. If the phased-output is cross-correlated with an antenna that *does not* include such a digital mixer in its data path, then the spectral dynamic range limit is not imposed. Finally, conclusions were drawn as to how it is relatively simple to feed phased outputs back into the correlator for autocorrelation or crosscorrelation processing. Nominally, it is necessary to design a daughter module to plug into the Station Board to facilitate data entry, and one Station Board will accommodate up to 16 phased (polarized) sub-bands. This will allow autocorrelation, real-time VLBI correlation where one element is the phased-VLA, or single interferometer phased sub-array crosscorrelation. If more sub-bands are required, or more phased sub-arrays are to be correlated, then more Station Boards would be required.

References

- [1] Carlson, B., 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] Crochiere, R.E., & Rabiner, T.R., 1983, Multirate Digital Signal Processing, Englewood Cliffs: Prentice-Hall
- [3] Carlson, B., Simulation Tests of Sub-Sample Delay Tracking in the Proposed WIDAR Correlator for the Expanded Very Large Array, NRC-EVLA Memo# 007, October 3, 2000
- [4] Carlson, B., Summary of Discussions Held During the July 10-14, 2000 Workweek in Socorro Regarding the EVLA-WIDAR Correlator, NRC-EVLA Memo# 005—DRAFT, August 22, 2000
- [5] Carlson, B., Two Correlators for the Price of One: How a VLBA Correlator Could Fit Within the Proposed 40-Station WIDAR EVLA Correlator, NRC-EVLA Memo# 006, September 28, 2000

