

Requirements for 8-bit Processing in the Proposed WIDAR Correlator for the EVLA

NRC-EVLA Memo# 010

Brent Carlson, January 29, 2001

ABSTRACT

This short memo defines minor design modifications that are required to support 8-bit initial quantization and 8-bit correlation. It will be shown that 8-bit initial quantization can easily be incorporated with a minor design change to the FIR filter—but that it will come at the expense of only being able to process one-half the analog bandwidth with one-half the number of FIR filter taps as with 4-bit initial quantization. This is not a serious performance limitation since 8-bit initial quantization will normally only be required at lower frequencies where the bandwidth is fundamentally restricted anyway. This memo will also show that 8-bit correlation (after re-quantization) does not appear to be practical—but that 7-bit correlation is possible with insignificant design changes. Even in extreme interference cases, 7-bit correlation should yield close to 60 dB of spectral dynamic range [1]. Use of 7-bit correlation does reduce the number of available spectral channels by a factor of four since distributed arithmetic is employed.

Introduction

In previous analyses of VLA-site interference data [1], it seemed that 3-bit initial quantization and 4-bit correlation would be sufficient to ensure that ‘secondary’ interference products do not show up in the output cross-power spectrum. Recently, interference detection work at the VLA (private comm, Armendariz) seems to indicate that the narrowband interference from sources such as aircraft DME outside of the observing bands analyzed is considerably stronger than within the analyzed bands. Thus, it has become a requirement to design the correlator to handle data from 8-bit quantizers¹ to provide high performance in extreme interference cases. This document looks at what design changes are required to provide 8-bit capability.

Eight-bit Initial Quantization and FIR Filter Structures

In the WIDAR design, all wide-band data from the initial quantizer must pass through the FIR filters² from which the sub-bands are generated. Thus, to support 8-bit capability, the FIR filters must be able to process 8-bit data independent of whether or not there is 8-bit correlation. The Station Board will support 4-bit data paths with a de-multiplexing factor of 16. Eight-bit initial quantization will reduce the de-multiplexing factor by a

¹ It is recognized that ultimate 8-bit performance is strongly influenced by the performance of the 8-bit initial quantizer. The performance limitations of this real device are not considered in this memo.

² Not *strictly* required since the data equal to a sub-band bandwidth can pass through unchanged.



factor of 2 since one-half the filter is used for the MSN (Most Significant Nibble) and one-half is used for the LSN (Least Significant Nibble). This reduces the effective number of taps available by a factor of 2. The maximum analog bandwidth also drops by a factor of 2 (since the correlator operates at a constant sample rate) from 2 GHz to 1 GHz. A simplified block diagram of this reorganization is shown in Figure 1 for a FIR filter with (for clarity) a 4-bit de-multiplexing factor of 8.

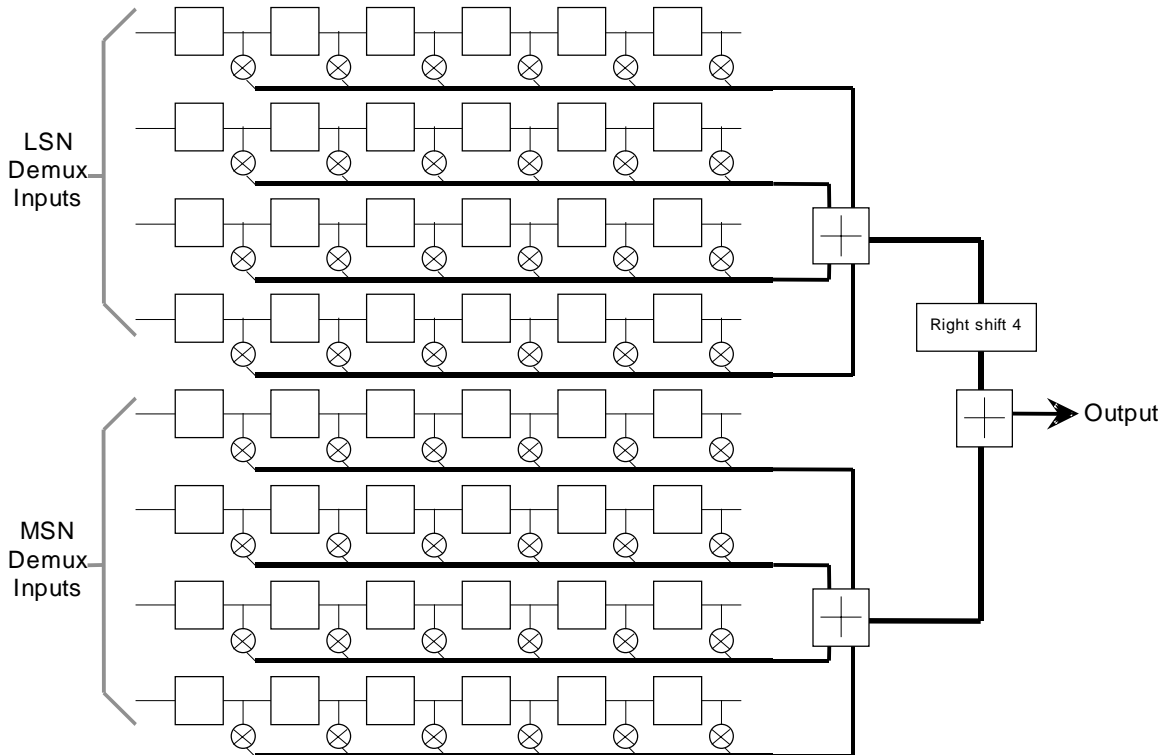


Figure 1 FIR filter structure to support 8-bit FIR filtering. In this figure a 4-bit de-multiplex/polyphase factor of 8 is shown resulting in an 8-bit de-multiplex factor of 4. A sum-of-products (SOP) is performed separately on the LSN (Least Significant Nibble) and the MSN data streams. The final SOP is generated by right-shifting the LSN by 4 bits before addition (or left shifting the MSN by 4 bits—requiring a larger adder).

The only additional requirement for the FIR filter to that described in [2] is to structure the adder tree so that the LSN and MSN SOPs can be performed separately, then shifting one of the SOPs before final addition. Generally, this requires simple organization and a programmable shift by 4—a negligible additional requirement.

FIR Distributed Arithmetic

Since the LSN and MSN SOPs are generated separately before final addition, it is necessary to ensure that the bit arithmetic works out properly. It would seem that there are two bit organizations where this is the case. The first one³ uses 8-bit 2’s complement data encoding. In this case, the 8-bit word is simply split in half, the LSN (without modification) is treated as unsigned data and the MSN is treated as 2’s complement data.

³ Source: Xilinx DSP seminar notes.

Since the product of the data and the tap coefficient is stored in a lookup table (LUT) in the FIR, this organization is easily supported. Two’s complement distributed data organization is shown in Table 1.

Address	MSN LUT	LSN LUT
0 0000	0	0
1 0001	1 x Constant	1 x Constant
2 0010	2 x Constant	2 x Constant
3 0011	3 x Constant	3 x Constant
4 0100	4 x Constant	4 x Constant
5 0101	5 x Constant	5 x Constant
6 0110	6 x Constant	6 x Constant
7 0111	7 x Constant	7 x Constant
8 1000	-8 x Constant	8 x Constant
9 1001	-7 x Constant	9 x Constant
A 1010	-6 x Constant	A x Constant
B 1011	-5 x Constant	B x Constant
C 1100	-4 x Constant	C x Constant
D 1101	-3 x Constant	D x Constant
E 1110	-2 x Constant	E x Constant
F 1111	-1 x Constant	F x Constant

Table 1 LSN and MSN LUT contents with 8-bit 2’s complement data. The 8-bit word is simply split in half, the LSN data is treated as unsigned and the MSN data is treated as 4-bit 2’s complement. Provided the 4-bit shift is performed before final addition as shown in Figure 1, the correct output is obtained. The ‘Constant’ in the table is the particular tap’s coefficient.

Example: If data is ‘-85’ and the ‘Constant’ is ‘+11’, is the correct answer obtained? -85 in 2’s complement is 0xAB (AB in hexadecimal). The LSN is 11 x B = 121; the MSN is -6 (i.e. A) x 11 = -66—shifted left by 4 bits becomes -1056. The final answer is 121 + (-1056) = -935—which is precisely -85 x 11.

The second method is to use *odd* data encoding with unsigned binary representation. This “naturally” comes out of an 8-bit A/D converter with the all zero/lowest output representing -255 and the all ones/maximum output representing +255. Once the 8-bit word is split in half, each half is treated in the same way: odd encoding with all zeros representing -15 and all ones representing +15. This representation is shown in Table 2.

Example: If data is -7 (01111100 binary; 0x7C hex) and the ‘Constant’ is +15, is the correct answer obtained? The LSN is 1100 or +9 x +15 = 135; the MSN is 0111 or -1 x +15 = -15—shifted left by 4 bits becomes -240. The final answer is 135 + -240 = -105—which is precisely -7 x 15.

It is important to note that the output of every LUT is 2’s complement and so the 4-bit shift before final addition must perform a proper sign extension.

Address	MSN LUT	LSN LUT
0 0000	-15 x Constant	-15 x Constant
1 0001	-13 x Constant	-13 x Constant
2 0010	-11 x Constant	-11 x Constant
3 0011	-9 x Constant	-9 x Constant
4 0100	-7 x Constant	-7 x Constant
5 0101	-5 x Constant	-5 x Constant
6 0110	-3 x Constant	-3 x Constant
7 0111	-1 x Constant	-1 x Constant
8 1000	+1 x Constant	+1 x Constant
9 1001	+3 x Constant	+3 x Constant
A 1010	+5 x Constant	+5 x Constant
B 1011	+7 x Constant	+7 x Constant
C 1100	+9 x Constant	+9 x Constant
D 1101	+11 x Constant	+11 x Constant
E 1110	+13 x Constant	+13 x Constant
F 1111	+15 x Constant	+15 x Constant

Table 2 MSN and LSN LUT contents with odd data encoding and unsigned binary representation. Here, the 8-bit odd-encoded word is split in half and each half is treated identically. The final 4 bit shift before adding the MSN and LSN SOPs is required as shown in Figure 1.

Additional FIR Structure Notes

In the memo on 2-stage FIR filtering [2], the FIR filter architecture that is presented (Figure 6 of the memo) contains 2:1 data selection MUXes that can be used to allow different FIR filter operating modes. The FIR’s flexibility can be improved by allowing *each* selection MUX to be individually programmed. This will permit all of the FIR’s taps to be used—in both 4-bit and 8-bit applications—if the original sampled analog bandwidth is less than the maximum. That is if, at maximum bandwidth, the de-multiplex factor is 16, then de-multiplex factors of 8, 4, and 2 will be able to fully utilize all taps rather than leaving taps unused. Of course, each FIR can still only handle one sampled analog band.

8-bit (Or, Nearly 8-bit) Correlation

After FIR filtering, it is possible to re-quantize the data to 4 bits as the filter design currently allows. However, for 8-bit correlation, it is necessary to re-quantize to 8 bits. Because of the way the sub-band data is rearranged by the “Sub-band Distributor Backplane” and carried by the cables to the Baseline Boards, it is not possible to directly use the 8-bit re-quantized data out of the FIR filter. It is thus necessary to split up the 8-bit re-quantized data into two 4-bit chunks for transport and correlation. Additionally, since the correlator contains 4-bit multipliers, it is necessary to obtain four, 4-bit products to achieve 8-bit correlation. This means that 8-bit correlation will reduce the number of spectral channels that are available by a factor of four. However, a mitigating factor is that it is possible to configure each sub-band for correlation with 8 bits or the normal 4

bits. Those sub-bands that contain extreme interference are correlated with 8 bits, and those without extreme interference are correlated, as normal, with 4 bits.

The need to perform four 4-bit correlations within a sub-band to generate the 8-bit result is similar to full polarization operation. Thus, it seems feasible to use an entire Station Board for one sampled analog band: one-half of the board generates the re-quantized MSN and one-half of the board generates the LSN. Two Station Boards are therefore required to be able to do 8-bit full polarization⁴. When the data for a particular sub-band shows up at the Baseline Board, the station receiver chip⁵ will have access to all MSN and LSN data for up to two baseband pairs (remember that, in this case, each Station Board is one polarization). Four 128 complex-lag cross-correlators can be used to do all of the 4 correlations to effect 8-bit correlation. A simplified correlator data path diagram showing how to do this is shown in Figure 2.

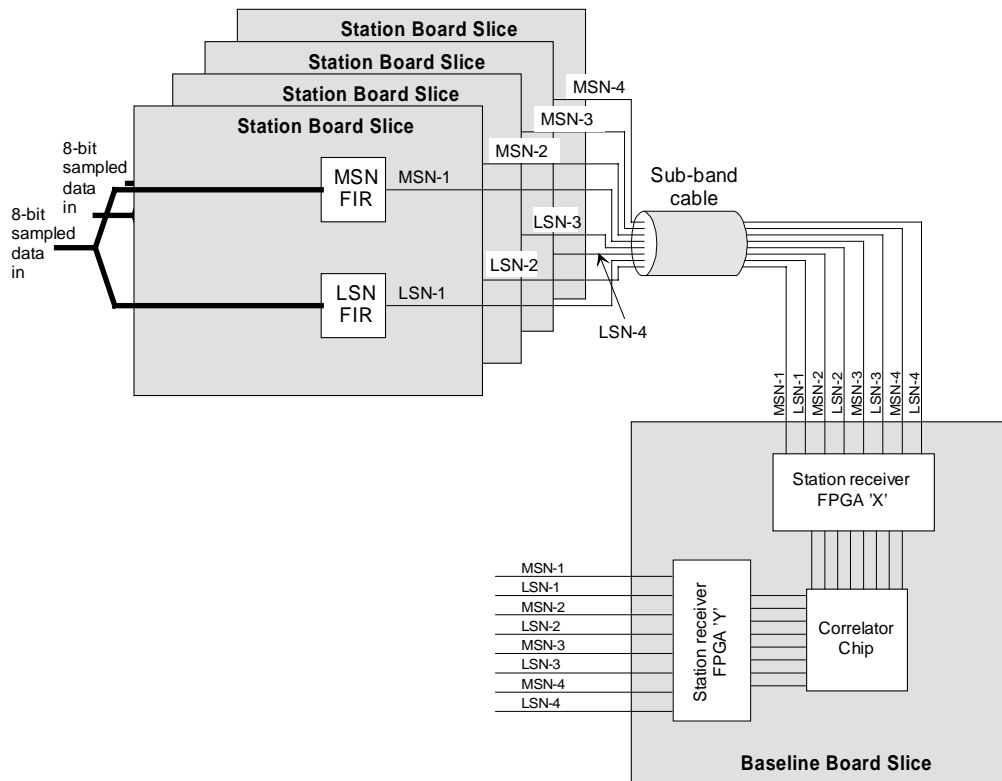


Figure 2 Simplified correlator data path diagram showing a single sub-band slice of a station input (4 Station Boards) and a Baseline Board with a correlator chip slice. The Station receiver FPGAs are provided with all of the data for two baseband pairs of 8-bit data. One correlator chip could perform all of the four 8-bit cross-correlations for one polarization pair since there are four sets of four 128 complex-lag correlators.

It is important to note from Figure 2 that at some point the same 8-bit data must be fed into the R and L FIR filter banks on the Station Board. This is a function that could be

⁴ Alternatively, the number of useable FIRs could be reduced by a factor of two, but this would require that the MSN and LSN SOPs reside on different chips. Final addition then becomes problematic—especially at the speeds being contemplated.

⁵ i.e. the re-circulation controller FPGA that gets the data ready for the correlator chip.

done at the antenna before the FOTS but it may be more robust to do it on the Station Board after coarse delay compensation. Also, both the MSN and LSN FIRs for a particular sub-band operate identically—they are only different when it comes to selecting whether the upper or lower 4 bits of the 8-bit re-quantized data is output from the chip.

In the current design concept, one state of the 4-bit data that is transported from the Station Board to the Baseline Board is used to indicate “data invalid” (i.e. for pulsar gating). Also, the 4-bit data contains embedded synchronization and identification information for downstream processing [3]. If 8-bit correlation (or, as we shall see, 7-bit correlation) is used, then there is no spare state to indicate “data invalid”. This means that the gating operation will not be available with more than 4-bit re-quantization. The way to handle this is to design the correlator chip with a full 4-bit multiplier and a *separate* data invalid line. Thus, with 4-bit re-quantization, gating is allowed and the separate data invalid line is invoked when the invalid state or embedded synchronization is recognized. With 7-bit re-quantization, gating is not allowed and the separate data invalid line is invoked only when embedded synchronization is recognized.

Correlator Distributed Arithmetic

With four 4-bit multipliers it is possible, in principle, to generate all of the necessary products to achieve 8-bit correlation. As we shall see, however, it does not appear to be practical to do this. It would seem that there are two encoding schemes that will allow 8-bit correlation to be achieved with 4-bit multipliers. These are analogous to those covered in the previous section on 8-bit FIR filtering: if the data is 2’s complement encoded, then the LSN is treated as an unsigned quantity and the MSN is treated as a 2’s complement number; if the data is odd encoded with the lowest level all zeros and the highest level all ones then the 4-bit multiplier must use a similar encoding.

In the first case, the multiplier in the correlator chip is a 4-bit 2’s complement multiplier (a Baugh-Wooley multiplier being one such efficient implementation). This multiplier generates the 2’s complement products complete with sign extension. To be able to perform 2’s complement x unsigned multiplication (e.g. $MSN_X \times LSN_Y$) would require modification of this multiplier—probably an impossible task without an increase in power and cost. Thus, the LSN must use only the lower 3 bits with the 4th bit forced to zero—so that it is taken as unsigned by the multiplier, and the MSN must use the next 4 bits as is⁶. This means that 7-bit multiplication can be performed. Seven-bit correlation will yield an additional 18 dB of dynamic range over 4-bit correlation which, in worst-case interference environments⁷, should yield close to 60 dB of dynamic range [1]. The distributed correlation equation (where each term is one 4-bit complex multiplier-accumulator output) is:

$$\langle r_7 \rangle = \frac{64}{N} \sum_{i=0}^{N-1} B_{3456i_X} B_{3456i_Y} e^{j\hat{\phi}_i} + \frac{8}{N} \sum_{i=0}^{N-1} B_{3456i_X} B_{0120i_Y} e^{j\hat{\phi}_i} + \frac{8}{N} \sum_{i=0}^{N-1} B_{0120i_X} B_{3456i_Y} e^{j\hat{\phi}_i} + \frac{1}{N} \sum_{i=0}^{N-1} B_{0120i_X} B_{0120i_Y} e^{j\hat{\phi}_i}$$

⁶ Refer to Table 1. In this case, the “Constant” is the other station’s data.

⁷ i.e. 90% of the total power is interference.

In the second case (to support full 8-bit multiplication with odd encoding), the 4-bit multiplier⁸ consists of an array of 16 XOR gates (since a 4 x 4 multiplier can be broken down into 16, ± 1 multipliers), followed by conversion to 2's complement, followed by a 3-stage adder. This design, and the resulting larger output word, would seem to make 4-bit multiplication with odd encoding unwieldy. (Recall that 4-bit correlation with distributed 2-bit odd-encoded correlators is entirely feasible because the 2-bit multiplier is very simple and each term is accumulated separately—not true for the 4-bit multiplier under consideration.)

Conclusions

This memo has demonstrated that 8-bit initial quantization with (optional) 7-bit correlation is entirely feasible and comes with the small price of some enhancements to the design of the FIR filter chip and the correlator chip. The enhancements to the FIR chip include more general programming of data-path control selectors compared to that defined in [2], division of the adder tree so as to allow proper addition of MSN and LSN sum-of-products, and 7-bit re-quantization with 3-bit LSN and 4-bit MSN selection. The enhancements to the correlator chip are full 4-bit multiplication and a separate “data invalid” line rather than using one data state for data invalid. The Station Board may⁹ need to be able to duplicate data in the R and L baseband data paths, and reprogramming or reconfiguration of the Fine Delay Controller is necessary to take into account the different word size. Eight-bit initial quantization restricts the input analog bandwidth to 1 GHz and the number of taps in the FIR filters are reduced by a factor of two. Seven-bit correlation reduces the number of spectral channels by a factor of four, and does not allow the use of pulsar gating. The analog bandwidth restriction is not seen as problematic since 8-bit initial quantization will probably only be used at lower frequencies where the interference is the worst and where the analog bandwidth is fundamentally limited anyway. The reduction in the number of spectral channels with 7-bit correlation may be problematic, but a per-sub-band selection of 4-bit or 7-bit correlation should help to mitigate this effect. Loss of pulsar gating with 7-bit correlation is unfortunate, but phase-binning capability should adequately compensate for this loss.

References

- [1] Carlson, B., Dewdney, P. Simulation Tests to Quantify the Spectral Dynamic Range and Narrowband Interference Robustness of the WIDAR Correlator for the EVLA, NRC-EVLA Memo# 009, November 1, 2000.
- [2] Carlson, B., A Closer Look at 2-Stage Digital Filtering in the Proposed WIDAR Correlator for the EVLA, NRC-EVLA Memo# 003, June 29, 2000.
- [3] 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.

⁸ At least a first go at a design anyway.

⁹ Since this function *could* be performed before the Station Board.

