ngVLA Electronics Memo No. 4

Trident 2.0 Concept: A Minimum Delta Update to the Central Signal Processor Reference Design

Omar Yeste Ojeda

January 17, 2020

Executive Summary

This documents describes a modification to the reference design of the ngVLA Central Signal Processor in order to make it fully compliant with system requirements when connected to the Integrated Receivers and Digitizers. This proposal intends to keep changes to a minimum with respect to the reference design [1]. A newer proposal with more substantial differences is described in [2]. Additionally, there is a request for a slight modification to the frequency plan so that it nicely fits the bandwidth of Trident's frequency slices. The impacted subsystems are the Local Oscillator Reference and Timing subsystem and the Integrated Receivers and Digitizers. For the sake of brevity, I assume the reader is familiar with the ngVLA reference design [3].

The Integrated Receivers and Digitizers (IRD) [4] and the Central Signal Processor (CSP) [1] described in the ngVLA reference design are based on independent designs, neither of which were originally conceived for the ngVLA. As a result, their interfaces are not compatible with each other. The solution described in the reference design relies on inserting the Digital Back End (DBE) [5] as an adapter between both sub-systems. In this arrangement, the DBE must perform part of the CSP functionality in order to comply with preliminary system requirements, but the task division between the DBE and the CSP has not been clearly defined yet. A better solution is desired for the conceptual design of the CSP, which is the purpose of this document.

The proposed modification is summarized as follows:

1. Remove the threefold partition of the CSP.
2. Equip the CSP with 144 Frequency Slice Processors (FSP), which corresponds to 28 GHz of bandwidth.[1]

3-a. For those antennas within the maximum range of the unformatted data link from the CSP, remove the DBE and provide the CSP with 4 Very Coarse Channelizers (VCC) per antenna.

3-b. For the remaining antennas, keep the DBE at the antenna and substitute the corresponding VCC at the CSP with a simpler adapter for data buffering and bulk delay correction.

4. For the antennas with no DBE, aggregate the outputs of the 16 IRD modules at the antenna into 4 active outputs by using an optical matrix switch.

[1] In this document, bandwidth is always expressed on a per polarization basis.

Contents

1 Frequency Slice Bandwidth

The first parameter that requires harmonization is the frequency slice bandwidth. This is the information bandwidth contained in each of the frequency slices (or sub-bands) in which the digitized bandwidth by the ADC is split before it can be processed by the FSP. The reference design assumes an approximate bandwidth of 200 MHz, which is the bandwidth the FSP has been originally designed for.

Assumption 1 – *The optimum frequency slice bandwidth is 200 MHz. Higher values exceed the capabilities of a single FSP, while smaller values require more FSP hardware for the same overall bandwidth.*

The IRD modules can sample at two rates, 14 GS/s (giga-samples per second) and 7 GS/s. Using a frequency slice bandwidth of 200 MHz, those sampling rates result into 70 and 35 frequency slices, respectively. In this configuration the CSP would require 140 FSP to process 28 GHz of bandwidth.

The next possible value for the frequency slice bandwidth is 194.4 MHz, which produces 72 and 36 frequency slices at 14 and 7 GS/s, respectively. I prefer using this option for the straw-man design because the prime factorization of 72 and 36 allows a simpler (than that from 70 and 35) implementation of the Fast Fourier Transform (FFT) used in the frequency slice generation process. On the other hand, the number of FSP increases to 144 for 28 GHz of bandwidth.

Frequency slices are generated through an oversampled polyphase filter bank, where oversampling allows to attain high spectral selectivity while relaxing the prototype filter design. The oversampling factor is defined as the ratio between the sampling frequency of the frequency slices and the channel spacing or frequency slice bandwidth. In the reference design, this factor is assumed approximately 10/9, or 11% excess, which is the value in the original design of the FSP.

Assumption 2 – *The frequency slice oversampling factor must follow the expression* $\frac{M}{N}$*, where* M *is the number of frequency slices and* N *is any integer* $0 < N < M$*.*

The above assumption implies generating one sample for each of the M frequency slices every N input samples. This notably simplifies the filter bank design as no additional resampler is needed. For 36 frequency slices, I propose using an oversampling factor 9/8, or 12.5% excess, which is the valid value closest to 10/9.

2 Frequency Plan

The downconversion frequency plan described in the IRD reference design utilizes fixed Local Oscillator (LO) frequencies, except for slight LO offsets from antenna to antenna to be determined yet. The ratio between LO frequencies and the ADC sampling frequencies can remain constant, for instance, when such offsets are generated by offsetting the frequency reference centrally [6]. Nonetheless, coherent LO distribution is not the only alternative, and the frequency plan proposed next finds application as long as the ratio between LO frequencies and the frequency slice bandwidth is constant for all the antennas.[2] In this case, it becomes more efficient separating the LO frequencies

[2]Which most likely implies two things: 1) That sampling clocks and LO frequencies use the same reference; and 2) that any frequency offset is introduced in such reference, hence proportionally affecting both sampling clocks and LO frequencies equally.

RF Band	LO Frequency (GHz)	RF Range (GHz)	# of Frequency Slices	Frequency Slice Range
1	N/A	1.07 - 3.40	13	5 - 17
2	5.83	3.40 - 8.65	27	18 - 44
	11.67	8.65 - 12.35	19	45 - 63
3	14.58	12.15 - 17.40	27	63 - 89
	20.42	17.40 - 20.51	16	90 - 105
4	23.33	20.32 - 26.15	30	105 - 134
	29.17	26.15 - 31.99	30	135 - 164
	35.00	31.99 - 34.12	11	165 - 175
5	32.08	30.43 - 34.90	23	157 - 179
	37.92	34.90 - 40.74	30	180 - 209
	43.75	40.74 - 46.57	30	210 - 239
	49.58	46.57 - 50.65	21	240 - 260
6	75.83	69.90 - 81.57	60	360 - 419
	87.50	81.57 - 93.24	60	420 - 479
	99.17	93.24 - 104.90	60	480 - 539
	110.83	104.90 - 116.18	58	540 - 597

Table 1: Modified nominal frequency plan.

an integer number of frequency slices in order to avoid partially overlapped frequency slices from different IF signals. Using the numbers above, this implies a slight modification of the LO fundamental, $35/12 = 2.917$ GHz (equal to 15 frequency slices) in lieu of the 2.9 GHz proposed in the current frequency plan [7]. Table 1 summarizes the modified nominal frequency plan excluding frequency offsets.

3 Passive Optical Circuitry

One of the key components of the Trident design [8] is the optical flexible circuit which passively interconnects the VCC and FSP parts, and a single FSP internally. It reduces the networking power consumption as compared to active networking. In this straw-man design, I will use the following assumption:

Assumption 3 – *Passive optical point-to-point interconnection is the optimum solution to connect VCC and FSP parts.*

The above assumption derives from the fact that the FSP input interface is based on a point-to-point connection

scheme. From Table 1, it is clear that at most 60 frequency slices per IRD module are needed to fulfill the frequency plan. Nonetheless, the use of point-to-point connections is not exempt from constraints. First, even if the VCC hardware were capable of processing the output of more than one IRD module, hundreds of transceivers would be needed to output the corresponding frequency slices. For the Trident hardware, it is clear that at least 4 TALON boards [8] are required per antenna to support ngVLA Bands 5 and 6. Second, since 4 VCC can only serve 4 IRD modules, if one of the IRD modules required more than 60 FSP (its frequency slices are processed in multiple observing modes) another IRD module must be disconnected from the VCC. However, I do not consider this limitation will be very restrictive in practice. Also, hardwiring the VCC to the FSP part might make it more difficult to connect custom frequency back-end processors in the future. [3]

There is now the problem of how to connect 4 VCC units to 144 FSP. On the one hand, every VCC must be connected to at least 60 FSP. The current TALON hardware only devotes 56 high-speed serial transceivers to the Mid-Board Optical (MBO) modules. The other 8 transceivers, connected to the Quad Small Form-Factor Pluggable (QSFP) interfaces, are needed to receive the VCC input. Hence I can devise four possible solutions:

1. A new hardware is developed for the VCC part. For example, Xilinx's Virtex Ultra Scale+ family amply exceeds the number of transceivers required.
2. Maintaining the current TALON hardware would probably need to increase the frequency bandwidth, for instance 218.75 MHz (64 frequency slices at 14 GS/s), so that the number of frequency slices output by one VCC TALON board decreases. This contradicts previous Assumption 1.
3. Less likely, some slower serial transceivers might be available to populate the missing transceivers for the MBO interfaces. Note that 16-bit quantization yields 14 Gb/s per frequency slice, and perhaps a coarser quantization could be needed.
4. Whereas the sub-band generation and down-selection must be carried out at the antenna by the DBE for the antennas out of the IRD range, this solution can be applied to the entire array. This leverages the DBE development, which must be employed at the remote antennas in any case.

Without confirmation that the frequency slice bandwidth can be stretched further, or that the third option is even possible, I think that developing new hardware seems the best strategy because it allows maintaining the passive optical circuitry between the VCC and FSP parts.

On the other hand, it is evident that every individual FSP cannot be assigned exclusively to a single VCC. I propose the sharing scheme shown in Table 2. The scheme is represented graphically in Fig. 1. It allows any distribution of FSP, such as two VCC using 60 FSP each. When two VCC share a FSP, a passive optical power combiner is embedded in the optical circuitry, so the optical transmitter of the inactive VCC must be turned off.

4 Optical Matrix Switches

In practice, the first consequence of the configuration described in the previous section is that the VCC must be physically located near the FSP. Since the DBE is no longer needed, the active IRD modules connect directly to the corresponding VCC. As a result, that configuration is only feasible for those antennas within the range of the

[3] Stakeholder requirement STK2901 states: "It is desirable that commensal back-ends (e.g., RealFast) be considered in the design, and provisions/interfaces for future commensal back-ends be incorporated into the design. (i.e. not "designed out")."

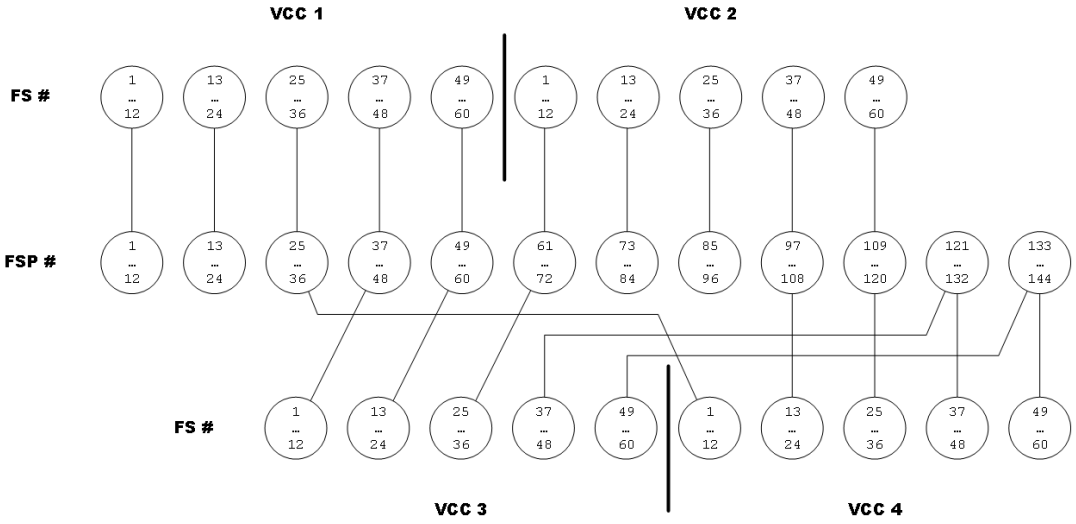

Figure 1: VCC-to-FSP wiring. Each node and wire represents 12 frequency slices.

unformatted serial data link, which is expected to be at least 40 km. This would be likely enough to cover the Core and the Plains sub-arrays. The IRD team will be responsible for developing and providing the IP block that the VCC needs for receiving unformatted data from the IRD modules.

The second practical consequence is that the 16 IRD modules at each antenna cannot be passively merged into four active connections at the VCC input. That would imply each IRD module can only communicate with a fixed set of 60 FSP. This constrains the observable frequency spectrum when sub-arrays operate in different observing modes, even when some FSP are still available. Hence, I propose using optical matrix switches (16x4) at the IRD output, so that any IRD module can be connected to any VCC, and hence, to any FSP. Installing this optical switch at the antenna should simplify the deployment of optical fiber.

FSP Range	1 - 60	61 - 120	37 - 72 , 121 - 144	25 - 36 , 97 - 144
VCC ID	1	2	3	4
1	1 - 24 (24)	–	37 - 60 (24)	25 - 36 (12)
2	–	73 - 96 (24)	61 - 72 (12)	97 - 120 (24)
3	37 - 60 (24)	61 - 72 (12)	–	121 - 144 (24)
4	25 - 36 (12)	97 - 120 (24)	121 - 144 (24)	–

Table 2: FSP assignment to VCC. The first row represents the FSP range assigned to the corresponding VCC indicated in the second row. The following rows show the FSP range shared by that same VCC in the second row and the VCC indicated by the first column follows. The number in parenthesis stands for the number of common FSP. For example, the FSP range 1-24 is exclusively assigned to VCC-1.

5 Remote antennas

Herein, remote antennas refers to those for which transmitting the IRD raw data directly to the CSP is not practical. Instead, data must be formatted and processed prior to transmission. For these antennas, the DBE must process IRD raw data and generate frequency slices to avoid transmitting the whole digitized bandwidth.

Assumption 4 – *The data transport network between remote antennas and the CSP is a switched network.*

The above assumption implies that any DBE can route any frequency slice to any FSP at will. Therefore, the IRD modules can be hardwired to the DBE, as well as the Data Transmission System (DTS) receiver to the FSP. Because most of the VCC functionality is now performed by the DBE, I suggest removing all VCC corresponding to remote antennas. In addition, the DBE must support flexible quantization of the frequency slices, anywhere between 4 bits (minimum requirements) and 16 bits, which is the maximum supported by the FSP.

Assumption 5 – *The TALON hardware does not allow the FSP to implement the long data buffers needed for compensating the transport delay from the remote antennas.*

From conversations with NRC, coarse delay corrections are applied by the VCC as the FSP cannot perform this task due to limited memory resources. Whereas transport latency is quite variable for the remote antennas, a signal buffer is required at the CSP end, which needs to be implemented by additional hardware other than the FSP. I suggest developing new hardware to carry out this task, as the TALON hardware seems quite an expensive option for this simple task. A first cut design would be based on a small FPGA or System on Chip (SoC), which would receive data related to a specific frequency slice and subset of antennas from the switched network, reorder data packets through an external DRAM-based data buffer, and feed the buffer output into the corresponding TALON-FPGA's MBO input interface. The number of antennas processed per each of these FGPA, and hence the total number of FPGA required, would depend on the selected device. Note that each TALON board processes a subset of 10 antennas, which must be considered when partitioning the system. Coarse delay corrections can readily be applied by this additional hardware as well. [4]

6 Conclusion

This proposal is an attempt to address the IRD-Trident incompatibility without having to modify their reference design significantly. It relies on several key assumptions requiring confirmation. Additionally, any of the choices made are open to discussion, and better alternatives could be possible. As compared to the reference design, the elimination of either the DBE or the VCC part for each antenna should be enough compensation for the extra cost of the added elements such as the optical matrix switches or the DTS receivers.

7 Acknowledgments

I would like to thank Christophe Jacques, Matthew Morgan, Rob Selina and William Shillue, for their invaluable help in conceiving this solution.

[4] The alternative, applying coarse delay corrections through the DBE at the antennas seems much more complex to me, if not impractical.

References

[1] O. Yeste Ojeda. *Central Signal Processor: Preliminary Reference Design*. Tech. rep. 020.40.00.00.00-0002-DSN-A. ngVLA Reference Design, 2019. URL: https://ngvla.nrao.edu.

[2] O. Yeste Ojeda. *ngVLA Electronics Memo No. 5. Trident 2.1 Concept: Updates to the Reference Design*. Tech. rep. 2020. URL: https://ngvla.nrao.edu.

[3] R. Selina et al. *System Reference Design*. Tech. rep. 020.10.20.00.00-0001-REP-B. ngVLA Reference Design, 2019. URL: https://ngvla.nrao.edu.

[4] M. Morgan and S. Durand. *Integrated Downconverters and Digitizers Design Description*. Tech. rep. 020.30.15.00.00-0002-DSN-A. ngVLA Reference Design, 2019. URL: https://ngvla.nrao.edu.

[5] J. Jackson et al. *Digital Back End/Data Transmission System: Reference Design Description*. Tech. rep. 020.30.25.00.00-0002-DSN-A. ngVLA Reference Design, 2019. URL: https://ngvla.nrao.edu.

[6] W. Shillue. *Local Oscillator Reference and Timing Design Description*. Tech. rep. 020.35.00.00.00-0002-DSN-A. ngVLA Reference Design, 2019. URL: https://ngvla.nrao.edu.

[7] M. Morgan. "IRD and Front-end Interfaces". ngVLA Antenna Electronics Workshop. Aug. 2019.

[8] B. Carlson and M. Pleasance. "Trident Correlator-Beamformer for the ngVLA: Preliminary Design Specification". 2019. URL: https://ngvla.nrao.edu.