LWA Memo No. 154

Preliminary Design for the Digital Processing Subsystem of a Long Wavelength Array Station

L. D'Addario and R. Navarro
Jet Propulsion Laboratory, California Institute of Technology[1]
11 February 2009

I. Introduction and Summary

The digital processing (DP) subsystem of an LWA station (consisting of the DP1 and DP2 subsystems defined in [1]) is being developed at JPL. This document describes its preliminary design. Details are subject to change, and engineering documentation (outside this memo series) should be consulted for the latest revisions.

In some ways, this design departs from the preliminary component designs described in earlier LWA Memos. It is based on our current understanding of the project's requirements and it represents what we believe is a cost-effective approach.

System Highlights

- DP1 will process the signals from up to 260 dual-polarization antenna stands, and will include 4 beamformers, a wideband transient buffer and a narrow band transient buffer.
- Digitization of the signals will be performed at 196 Msps.
- Each signal will be adjusted in delay and amplitude and then all those of the same nominal polarization will be summed to form a beam that can track any position in the sky. Four such beam-forming systems will operate in parallel, all independently.
- DP2 will provide two independently-tuned digital receivers (DRXs) for each beam.
- COTS chassis infrastructures will be used to reduce electromechanical and thermal risks and development cost.

Future Expansion

- Extensive use of FPGAs allows the circuit boards to be reprogrammable, permitting future development of new beamforming and transient signal processing techniques. (However, the total amount of logic available for processing each antenna's signals is fixed, and the bandwidth available for combining the signals from different antennas is also fixed.)
- Modular nature will allow for arrays of fewer antennas, with proportional reduction in cost, and for adding more antennas at any time.

II. Top-Level Design

The top-level block diagram is given in Figure 1. We have used components that are readily available and with which we are familiar from other projects. Our choices include packaging in 14-slot ATCA chassis with high-speed backplanes; accomplishing most processing with Xilinx Virtex-5 series FPGAs; and implementing board-to-board data connections with the "RocketIO" serial channels available on these FPGAs (each supporting a practical user data rate up to about 2.5 Gb/s).

There are 512 analog input signals, consisting of a polarization pair from each of the 256 antenna stands. To minimize the need for board-to-board connections, all the stand-specific processing (DP1 processing) required for each signal pair is kept together. Thus, we have DP1

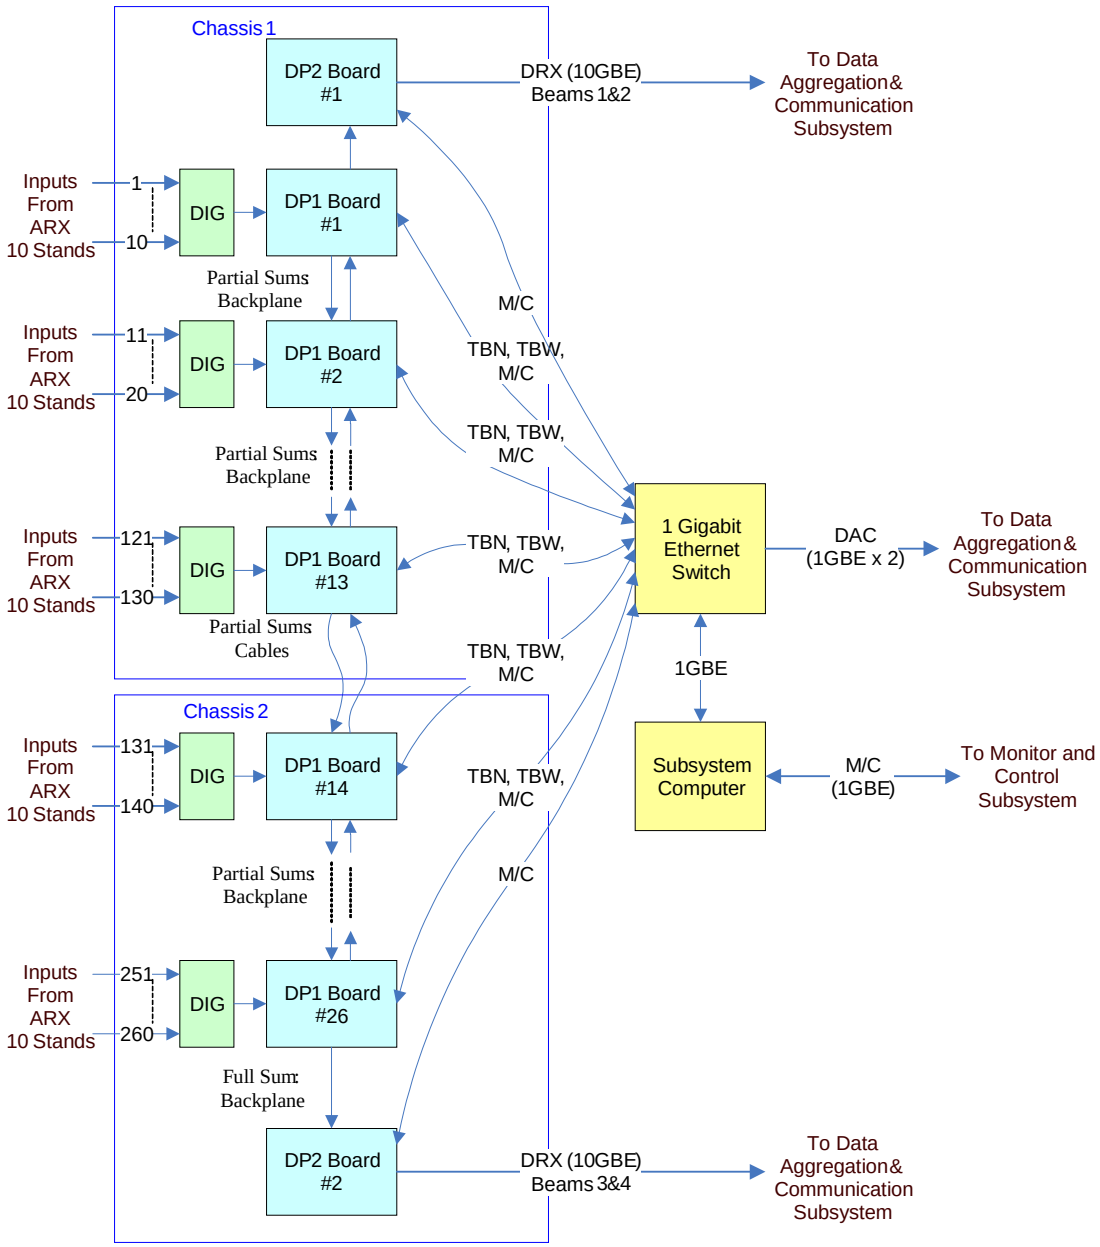

Figure 1: DP1-DP2 Subsystem Block Diagram

boards, each associated with a digitizer (DIG) board, containing all four beamformers (BFUs) and both transient buffers (TBN and TBW) for several stands, rather than separate boards for each function. We estimate that such processing for 10 stands can be accommodated on one 12.75in high by 11in deep ATCA board, requiring 26 boards to support 256 stands.

Each DIG board is constructed as a "rear transition module." It plugs into a connector at the back of the corresponding DP1 board, fitting within the ATCA chassis, and provides connectors for the input signals from the Analog Receiver (ARX) subsystem at the rear panel of the chassis. Thirteen of the DIG/DP1 board pairs are installed in each of two 13U high ATCA chassis, occupying a total of 45.5in vertically in a 19in wide rack.

As shown in Figure 1, the partial sums for combining the processed stand signals into beams are daisy chained through the DP1 boards with each board adding the signals from 10 stands. The end of the chain connects to a DP2 board (see Section IV) which is in the remaining slot of each 14-slot ATCA chassis. In order to make efficient use of the symmetrical bi-directional RocketIO ports of the FPGAs, the chains for two beams proceed in one direction (from DP1

board #1 through #26, ending in DP2 board #2) and those for the other two beams proceed in the opposite direction (DP1 #26 through #1 to DP2 #1). Each DP1 board can receive and send its partial sums for the daisy chain either through the ATCA backplane or through front panel connectors which accept Infiniband cables. The boards within a chassis are interconnected via the backplane. The 13th DP1 board in the first chassis connects to the first board in the second chassis (DP1 number 14) by a cable. The number of stands being combined can easily be increased or decreased in increments of 10 by adding or subtracting DP1 boards. (For more than 26 boards or 260 stands, the additional boards must be in one or more additional chassis inserted in the chain between Chassis 1 and Chassis 2.)

In this architecture the amount of logic available for processing each stand's signals and the bandwidth of the interconnections is fixed, and this limits the number of beams. The current design supports four dual-polarization beams. An architecture in which each beam is processed in a separate set of boards (as in [1] and [4]) allows adding or deleting beams without changing the board designs, but this is achieved at a high cost in interconnections. For example, for four beams the latter architecture requires a total of 3104 board-to-board connections of 98 MHz bandwidth (each implemented as 4 LVDS lanes), whereas the architecture here requires 208 (each implemented as 2 RocketIO channels).[2]

The summed signals for each beam are processed by Digital Receivers (DRXs, see Sec. IV) in the DP2 boards. Each DP2 board sends the beamformer outputs directly to the Data Aggregation and Communication (DAC) subsystem via a 10 Gbps Ethernet link.

Meanwhile, the TBN or TBW outputs from all channels of a board are transmitted to an Ethernet switch (an off-the-shelf device), where they are made available to the DAC subsystem. For each board, 100 Mbps links are sufficient to support continuous TBN readout at a bandwidth of 100 kHz and 24 b/sample (complex), or to support TBW readout at a 1:1000 duty cycle and 12b/sample (real), allowing 50% Ethernet overhead. However, the aggregate TBN/TBW data from all boards requires two or more 1GB Ethernet links for transmission to the DAC subsystem. We assume that TBW readout and TBN streaming are not needed simultaneously. The same network is used for communication between the subsystem computer and all DP1 and DP2 boards, providing internal control and monitoring within the DP subsystem.

The subsystem computer is a general purpose server embedded in the DP1/DP2 subsystem. It is expected to be a 1U high rack-mounted unit in the same rack as the main processing chassis. It provides the interface to the station-level Monitor and Control (MC) subsystem, and it handles initialization and setup of the processing boards. It has access to the TBW and TBN data (although not at full bandwidth), so it can provide limited local processing of those data for diagnostic purposes. The full bandwidth of the TBW and TBN goes to the DAC subsystem, as do the filtered beamformer outputs from the DP2 boards.

[2] We have 26 DP1 boards through which the partial sums must be routed, and 8 signals (4 beams, 2 polarizations) from each board to the next, thus 26*8=208 connections. As shown in Fig. 6 of [4], the alternative architecture (for 4 beams) requires the same 8 partial sum connections among fewer (4) sets of beamforming boards, but it also requires 512 connections (256 stands, 2 polarizations) for the DP1 daisy chain from each functional group of boards to the next (DIG to 4xBFUs to TBW to TBN), for a total of 4*8 + 512*6 = 3104 connections. Generalizing, for B beams the totals are $52B$ and $512(B + 2) + 8B$, respectively. (Fig 6 of [4] shows 3 beams and omits the TBW and TBN modules.) These are counts of logical signal connections, irrespective of their representation (sample rate, bits per sample, and signaling rate per physical connection), which also affects the implementation cost.

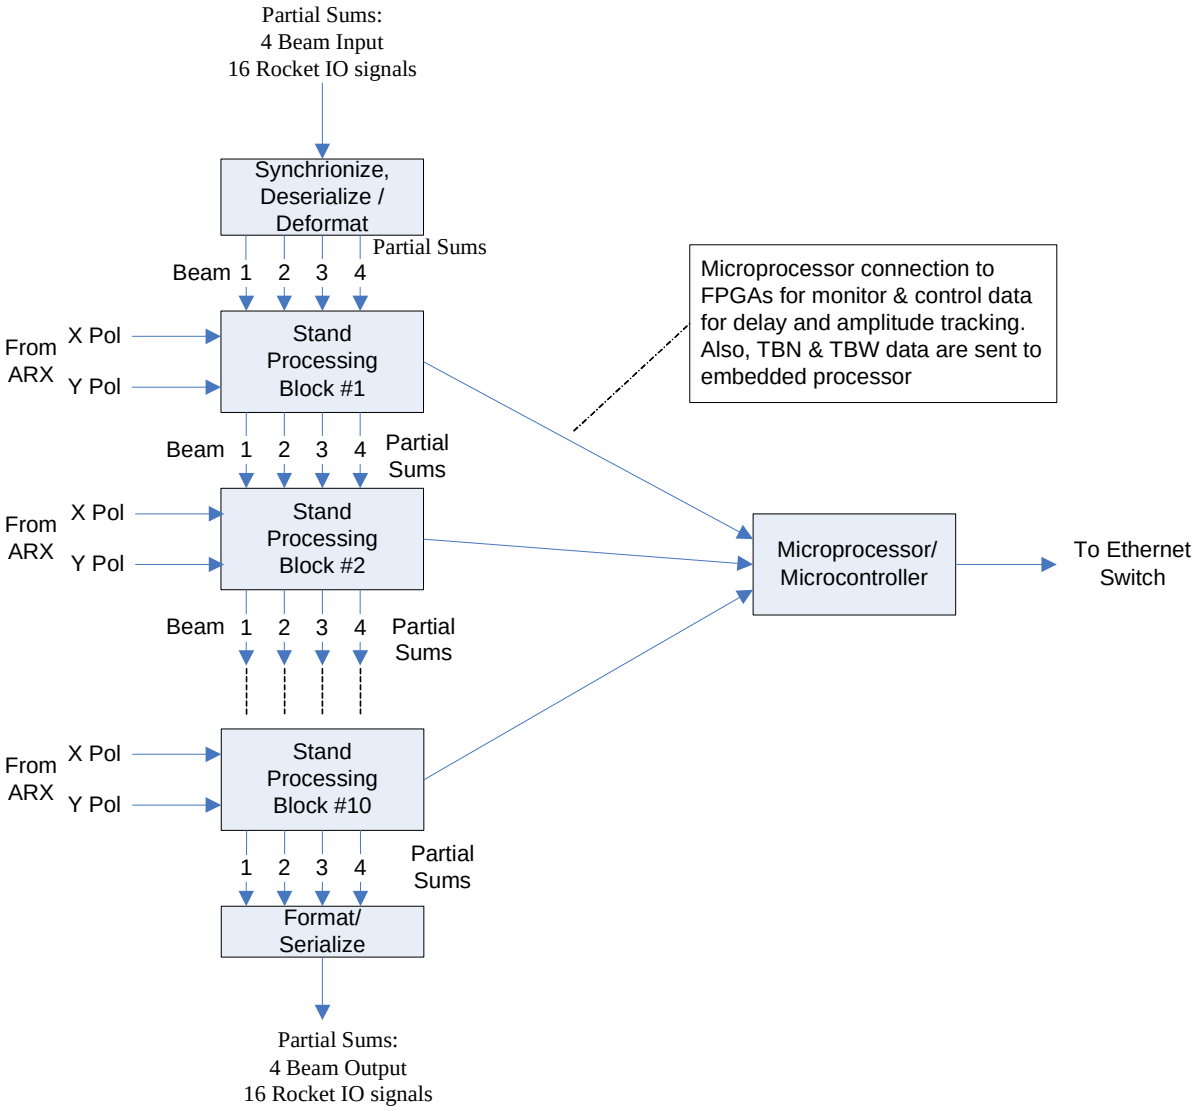

Figure 2: DP1 Board

III. Signal Processing For Each Antenna Stand (DP1)

The DP1 board's block diagram is given in Figure 2, and the processing for each stand is illustrated in Figures 3 and 4. In Figure 2, 8 partial sums (four dual-polarization beams) from the predecessor boards are received on 16 RocketIO lines. After deserializing, deformatting, and synchronizing, the partial sums pass through one processing block for each of the 10 stands handled by this board, and then are formatted and serialized for transmission to the successor boards. Of the four beam partial sums entering this board, two come from the board in the preceding slot in the chassis, and two come from the board in the next slot; similarly, two outputs go to the next slot and two to the preceding slot (see Fig. 1). The logic includes automatic correction for variations in the latency of these links. The board also includes a microprocessor to supervise the operation of the FPGAs, to implement the Ethernet link for the TBW and TBN outputs, and to provide a monitor/control interface to the subsystem computer.

The digitizer sampling clock runs at 196 MHz, and samples are processed at this rate throughout DP1. The digitizer resolution is 12 bits, but the partial sums are carried through the summing chain as 20 bit words so as to avoid truncation or overflow. This results in an aggregate rate of 31.36 Gbps for the 4 dual-polarization beams, and this is comfortably carried

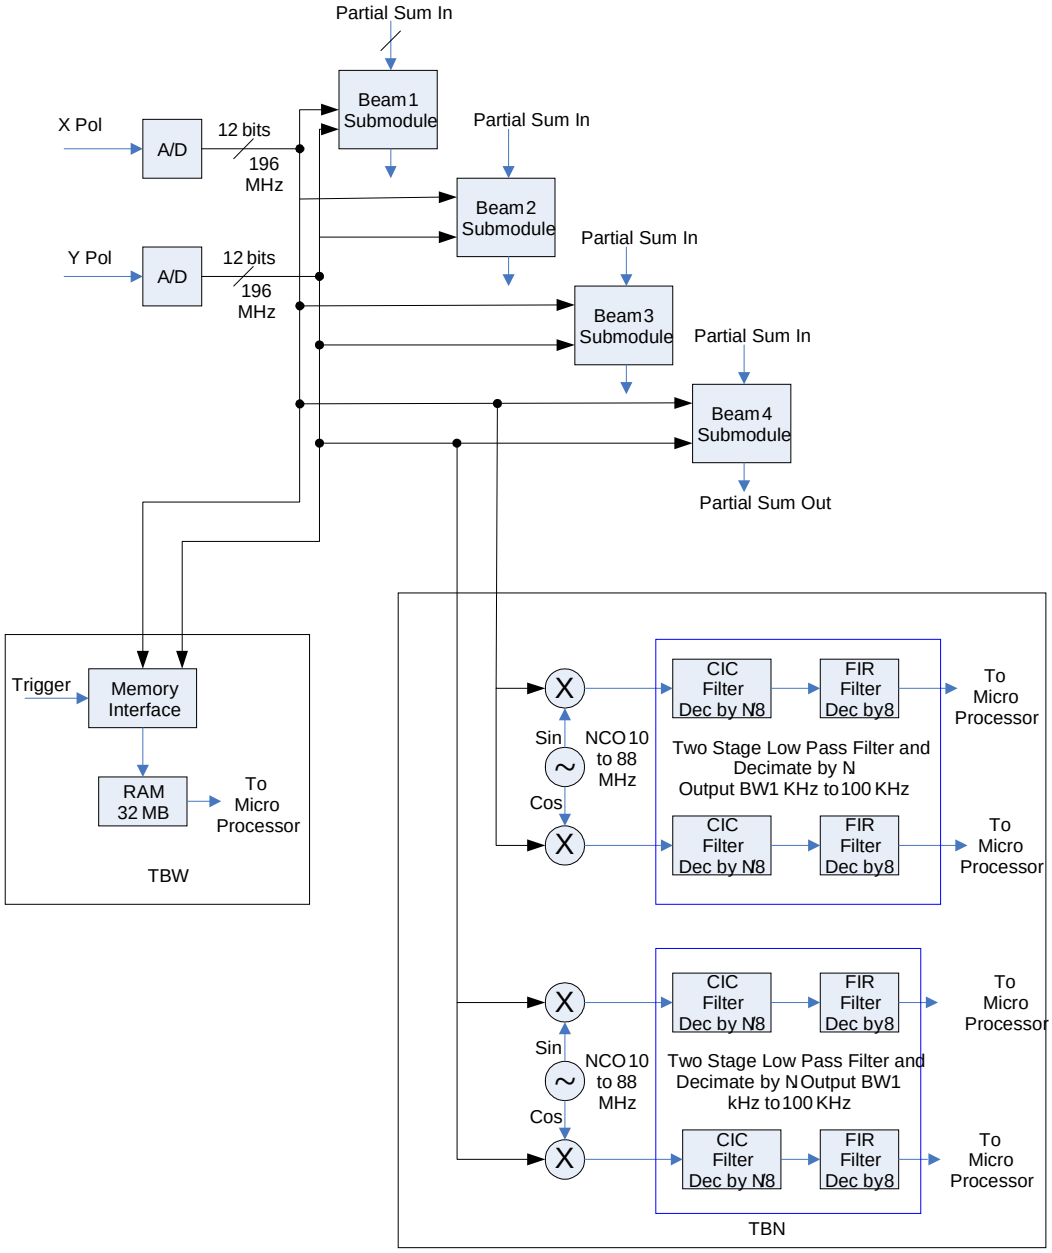

Figure 3: Stand Processing Block

by 16 RocketIO channels at 1.96 Gbps each.

Figure 3 shows the processing of the signals from each stand. After digitization (12 b, 196 MHz), the samples are distributed to 6 parallel processes: 4 beamformers, the wideband transient buffer (TBW), and the narrow band transient buffer (TBN). The TBW consists of a memory controller (implemented in the same FPGA as the other processing) that records all samples to a 32 MB RAM (implemented in a separate DRAM device) beginning when a trigger is asserted and ending when the RAM is full. The memory controller and RAM are shared among several processing blocks because it is assumed that the same trigger signal will be used for the whole array. The interface can provide some flexibility in recording time: 57 Msec is available if the full 12b-wide samples are recorded, or proportionally longer if fewer bits of each sample are recorded. Meanwhile, the samples also go to the TBN where they are digitally downconverted, low-pass filtered, decimated to the reduced Nyquist frequency, and made available to the microprocessor for output via the Ethernet link. (Figure 3 shows a separate NCO

as the local oscillator for each downconverter, but a single NCO per FPGA will suffice because all channels of the array will normally be tuned to the same TBN frequency.)

Figure 4 shows one beamformer module. Each signal of the polarization pair is first delayed to compensate for the array's geometrical delay in the desired beam direction, as well as for the cable delay and latencies within the DP subsystem. Delay is set to the nearest sample using a FIFO, and then an interpolating FIR filter is used to provide fractional-sample delay. The FIR filter can also provide a small amount of dispersion correction, if needed. Next each pair of samples is multiplied by a 2-by-2 matrix. The diagonal elements of the matrix, whose values are normally between 0.5 and 1, provide amplitude weighting as part of the beam forming process. The range of the weight is limited so as to avoid carrying additional bits after the multiplication. It is assumed that further amplitude adjustment, if needed, can be provided within the analog processing to within a factor of 2 (6 dB) of the desired value. However, the weight can also be set to zero, for example to remove a malfunctioning antenna from the beam. The off-diagonal elements of the matrix, nominally zero, provide for a simple transformation of the signal polarizations. This is intended to help align the polarizations of all stands, if necessary. It cannot correct a misalignment that varies with frequency, since the same transformation is applied to every sample pair.

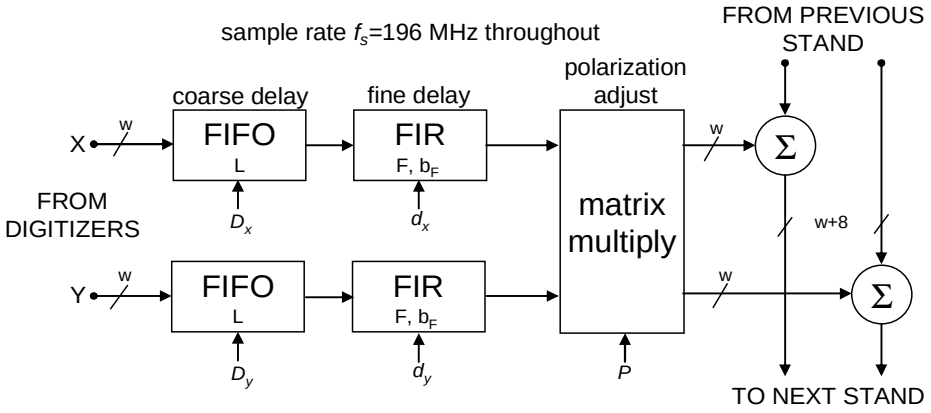

Figure 4: Beam Submodule. Each incoming sample is *w*=12 bits wide, and the partial sums are *w*+8=20 bits wide. D_x, D_y, d_x, and d_y are delay control numbers; *P* is a 2x2 coefficient matrix; and *L*, *F*, b_F are filter design parameters.

The processed samples are then added to the partial sums from the previous stand, and the new partial sums are sent to the processor of the next stand. Processing for the previous or subsequent stand may be within the same FPGA, in another FPGA on the same board, or on another board.

A single Xilinx XC5VSX50T FPGA is expected to provide enough processing for 2 stands, so that 5 of them are needed to handle the 10 stands of each board. Although the amount of processing available for each stand is fixed by the size of this FPGA, the nature of that processing is re-programmable, and if requirements change it could be made very different from that shown in Figures 3 and 4.

IV. Digital Receivers (DP2)

The final sum of all stands for each beam is sent to a DP2 board. This board contains the digital receivers (DRXs) specified in [1]. The logic of an elementary DRX is illustrated in Figure 5. For each polarization pair of signals of one beam, a DRX selects a portion of the 10 to 88 MHz signal band by digital downconversion using a tunable local oscillator (numerically controlled oscillator, NCO), sine/cosine mixers, and low-pass filters. The bandwidth of the selected region is adjustable from 0.25 MHz to 8.0 MHz in factors of 2. The signals are

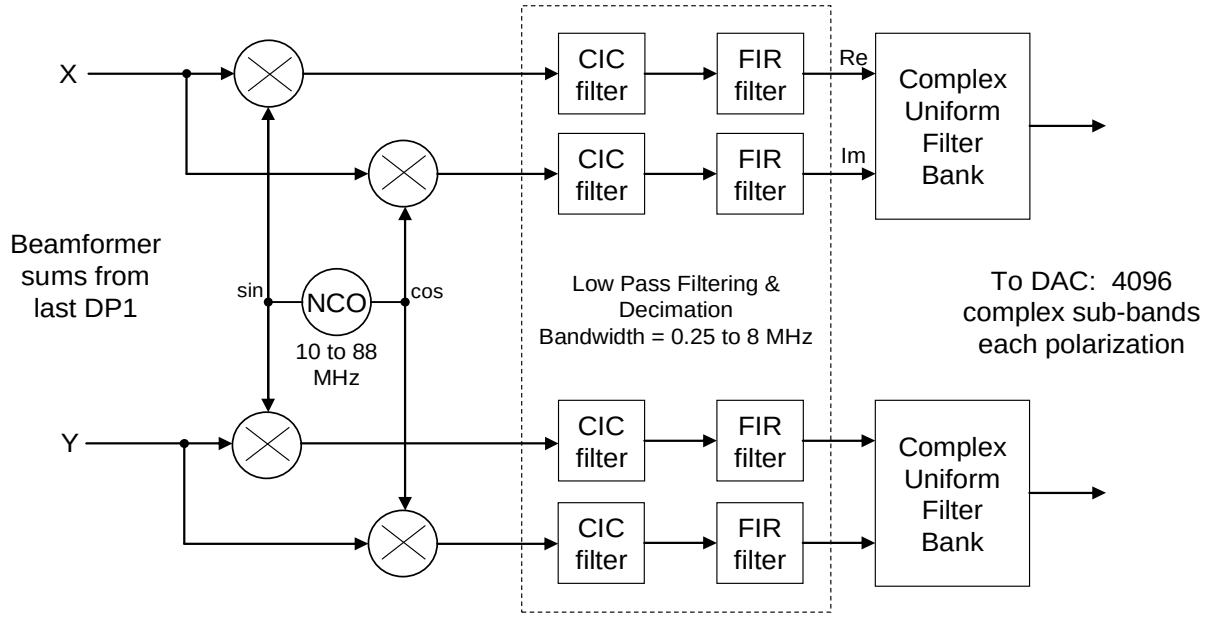

Figure 5: Digital Receiver (DRX) on the DP2 board. Each board includes 4 of these, one for each of (2 beams) × (2 tunings).

decimated to the new Nyquist rate and then processed by a filter bank of length 4096, producing subchannels with bandwidths from 61 Hz (at 0.25 MHz channel bandwidth) to 2 kHz (at 8 MHz).

Two independently-tunable DRXs are provided for each beam.

In the digital downconverters of the DP2 boards and also those of the TBN processors in the DP1 boards, the low pass filters are implemented efficiently by breaking them into several stages. The first few stages are simple cascaded integrator-comb filters [5], which do not require multipliers; the last stage, where the sample rate is slower because of partial decimation, is an FIR filter.

The hardware of the DP2 board is intended to be identical to DP1. They differ only in their FPGA and microprocessor programming, and the fact that DP2 has no associated DIG board. Whereas the signal processing required in DP2 by the current specifications is considerably less than that in DP1, it would be possible to reduce the construction cost of DP2 by designing a different board with fewer or smaller FPGAs. However, keeping them identical reduces the requirement for spare boards (we recommend two spares per station of each board type), reduces development cost, and allows for the addition of new capabilities in the future.

REFERENCES

[1] Steve Ellingson, "Long Wavelength Array Station Architecture, ver 1.0." LWA Memo 119, November 11, 2007.

[2] Steve Ellingson, "Polarization Processing at LWA Stations." LWA Memo 106, October 28, 2007.

[3] Steve Ellingson, "ADC Sample Rate and Preliminary Design for a Full-RF ADC Post-Processor, Ver.0.1." LWA Memo 101, September 11, 2007.

[4] Steve Ellingson, "BFU and DP1 Preliminary Design." LWA Memo 108, November 4, 2007.

[5] Eugene B. Hogenauer, "An economical class of digital filters for decimation and interpolation." *IEEE Transactions on Acoustics, Speech and Signal Processing* 29 (2): 155-162 (April 1981).