National Research Council Canada | Conseil national de recherches Canada

Canada

Trident Frequency Slice Architecture Correlator/Beamformer Reference Design for ngVLA

Mike Pleasance, Brent Carlson and Michael Rupen
National Research Council Canada (NRC), Herzberg Astronomy and Astrophysics (HAA), Penticton, BC, Canada

Trident Correlator/Beamformer Concept

The CBF ingests the voltage streams, digitized and packetized by the antennas, transmitted via the data transmission system, and possibly pre-processed by the DBE Sub-System. The CBF produces a number of low-level data products to be ingested by the Science Data Processor or Archive.

The CBF operates in up to five different simultaneous *Observing Modes*:

Synthesis Imaging: The CBF computes all spectral auto- and cross-correlation functions, including cross-polarizations within a sub-array. This includes normal fine-channel as well as zoom imaging.

(Sparse/Dense) Pulsar Timing: The CBF generates a number of channelized beams for further processing by the Pulsar Engine. There are two Pulsar Timing Observing Modes: Sparse and Dense. The difference between them is the number of beams that are generated per sub-array. Each of these two Observing Modes finds application for sparse or dense pulsar populations in the sky, which lead to different beamforming requirements.

Off-line Pulsar Search: In this Observing Mode the CSP outputs all four Stokes parameters at a given time-frequency resolution. In order to minimize telescope observation time, a great number of simultaneous beams is generated and subsequently stored for offline processing.

VLBI: The VLBI Observing Mode operates an ngVLA sub-array as a single VLBI station. Within this mode, several beams of varying bandwidth each can be generated and the resulting voltage streams stored in a VLBI standard format.

Table 1: Correlator/Beamformer Key Requirements

Parameter	Current CSP Requirement	Provided by NRC Trident
Maximum # of Antennas	256	260
Maximum # of Baselines	32896	33930
Instantaneous Imaging Bandwidth	20 GHz / pol.	Up to 30 GHz / pol.
Instantaneous Beamforming Bandwidth	20 GHz / pol.	Up to 30 GHz / pol.
Maximum aggregate bandwidth (simultaneous Observing Modes)	28 GHz / pol.	30 GHz / pol.
Maximum array diameter	1000 km	1000+ km
Channel Bandwidth	400 Hz to 2.5 MHz	210 Hz to 20 MHz
Maximum # of Channels	300,000	2,200,000
Maximum # of Pulsar Timing Beams	50 Beams @ 20 GHz / pol.	50 Beams @ 30 GHz / pol.
Maximum # of Pulsar Search Beams	91 Beams @ 20 GHz / pol.	100 Beams @ 30 GHz / pol.
Maximum # of VLBI Beams	2 Beams @ 20 GHz / pol.	4 Beams @ 30 GHz / pol.

Frequency Slice Architecture

The Frequency Slice Architecture splits the CBF into two parts. The VCC-Part provides Band-specific signal progressing to channelize wideband input streams into a number of narrower over-sampled Frequency Slices (FS). The bandwidth of an FS is common for all Bands making subsequent processing Band-agnostic. Each FS can be routed to one or more configurable Frequency Slice Processors (FSP) in the FSP-Part. Each FSP can be programmed to operate in one of the five Observing Modes described above and can perform that function for any number of sub-arrays in different Bands. The total number of FSPs dictates the aggregate bandwidth that can be processed simultaneously, but distribution of bandwidth across Observing Modes can vary by observation. This optimizes the amount of processing hardware required while providing the flexibility to assign processing resources to normal fine-channel imaging, zoom imaging or beamforming. The modular approach and clear division between the VCC-Part and FSP-Part allows implementation of future upgrades (such as new receivers or Observing Modes) with minimal impact to existing functionality.

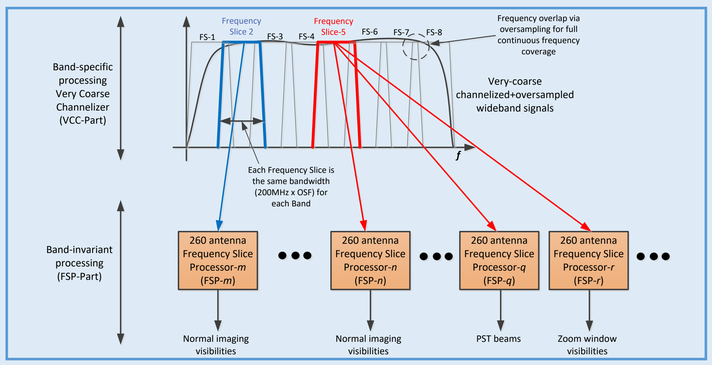

Figure 1: Frequency Slice Architecture showing channelization and distribution of Frequency Slices.

Abstract

The Trident Frequency Slice Architecture Correlator/Beamformer (CBF) Reference Design describes a digital correlator/beamformer system that meets the science requirements of the ngVLA synthesis radio telescope – specifically, processing 28 GHz of aggregate bandwidth per polarization for 256 antennas. The reference design uses the Frequency Slice Architecture (FSA) developed by NRC which aims to optimize cost by reducing the processing hardware requirements while increasing modularity. Incredible flexibility is provided by delivering many independent Frequency Slice Processors (FSP) which can be allocated to continuum, spectral line (zoom), or beamforming work depending on the needs of an observation. The reference design implements an FSA CBF using NRC's TALON technology currently under development for the Square Kilometer Array Mid Frequency Telescope Correlator/Beamformer. The TALON technology is fiber-connected Intel Stratix 10 FPGA based signal processing boards in 2U (air-cooled) or 1U (liquid cooled) rack mount server boxes. While ngVLA will use future FPGA technology still in development, the reference design represents a low-risk solution using currently available technology that can be accurately costed. Cost, power consumption, and rack space requirements can be extrapolated to future technology nodes based on industry projections.

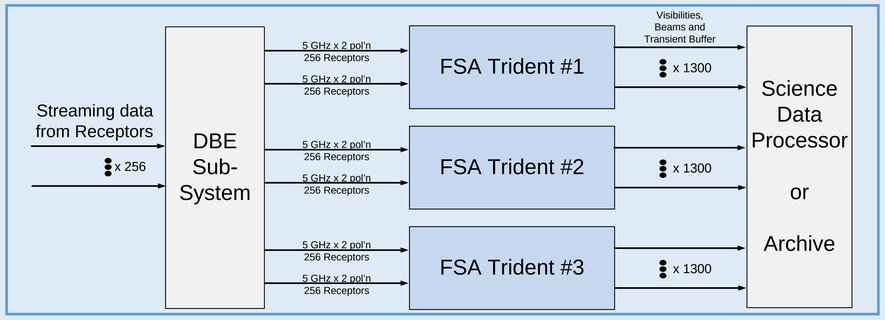

Figure 2: Top-level physical architecture of the Trident CBF for ngVLA

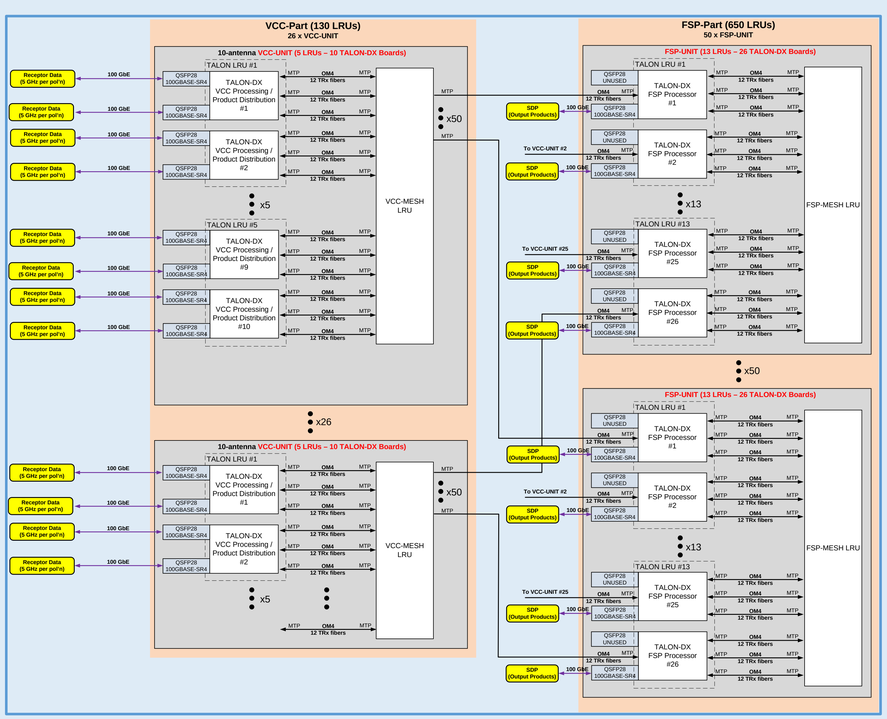

Figure 3: Physical architecture for one 10 GHz x 2 polarization trident

Further Information

For further information on the NRC Trident Frequency Slice Architecture CBF Reference Design for ngVLA, please contact: Mike Pleasance (Michael.Pleasance@nrc-cnrc.gc.ca) or Brent Carlson (Brent.Carlson@nrc-cnrc.gc.ca).

Trident Physical Architecture

The Trident Correlator/Beamformer for ngVLA is implemented as three x 10 GHz per polarization 'trident' sub-systems (Figure 2). Antenna data is received by each trident from the DBE Sub-System via two 5 GHz x 2 polarization streams. Data products generated by the FSPs are output on up to 3900 100GbE links for further processing or archiving.

Each trident contains a VCC-Part and a FSP-Part. The VCC-Part contains one FPGA for each antenna and performs very coarse channelization to generate 25 Frequency Slices (FS) for each of the 5 GHz x 2 polarization input streams (50 FS in total for 10 GHz). Each VCC FPGA can direct any FS to one or more Frequency Slice Processors (FSP-UNIT) in the FSP-Part using one of 50 point-to-point output fibers. Output fibers from the 10 FPGAs in a VCC-UNIT are grouped together with a passive optical circuit so one FS from all 10 FPGAs can be transported to an FSP using a 12-fiber MTP cable.

The FSP-Part (of each trident) contains 50 FSP-UNITs which can each be configured to perform one of the five Observing Modes. An FSP-UNIT contains 26 FPGAs and a passive optical circuit that implements a mesh providing 26 Gbps communication links between each pair of the FPGAs. Each FPGA receives an input of one Frequency Slice from 10 antennas (one VCC-UNIT). This provides a total input capacity for 260 antennas into the FSP-UNIT. In general, the processing in each FPGA is performed in two stages. First, per-antenna processing is performed on each of the ten antennas received by the FPGA. Data is then distributed to the other FPGAs based on bandwidth or beams via the FSP-MESH so that each FPGA gets a portion of the data for all antennas for correlation or beamforming. Each FPGA in the FSP-UNIT has a 25GbE or 100GbE output to the Archive/SDP.

Figure 3 shows the physical architecture of one 10 GHz x 2 polarization trident. Each trident requires 780 1U TALON LRUs, with each LRU containing two TALON-DX FPGA based signal processing boards plus support equipment (power supplies, liquid cooling plates, fans, etc.). Each TALON-DX board contains an Intel Stratix 10 FPGA (SX family device equipped with an embedded quad-core ARM A53 processor), two QSFP28 cages for 100GbE interfaces, 54 high-speed optical SERDES via five FCI Leap mid-board optical modules and four DDR4 DIMM modules (Figure 4).

Each trident requires 21 processing racks plus one rack for monitor and control servers. The entire ngVLA CBF requires 2340 TALON LRUs, 66 racks and will consume ~1500 kW of power.

As FPGA technology advances, a new version of the TALON-DX signal processing board will be designed and prototyped. Moving to next generation FPGAs will decrease power consumption (up to 30%) and rack space requirements. It will also provide devices with High Bandwidth Memory (HBM) which will significantly increase memory bandwidth over the existing DDR4 based solution.

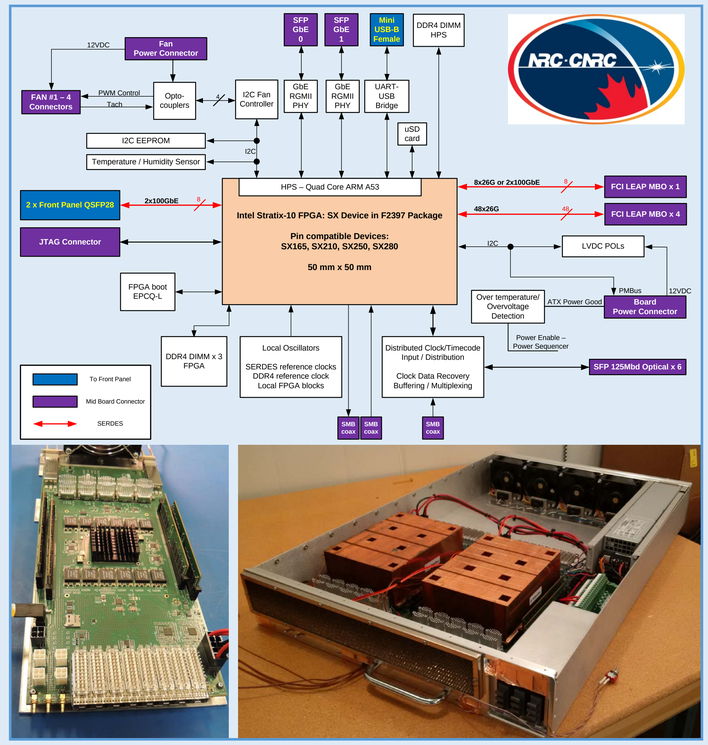

Figure 4: TALON-DX block diagram (top), prototype TALON-DX board (left) and prototype 2U TALON LRU (right)