32nd URSI GASS, Montreal, 19-26 August 2017

High-Performance Hardware Platform for the Square Kilomtre Array Mid Correlator & Beamformer

Michael Pleasance*(1), Heng Zhang(1), Brent Carlson(1), Ralph Webber(1), Dean Chalmers(1), Thushara Gunaratne(1), et.al. (1) (2)
(1) National Research Council Canada, National Science Infrastructure, P.O. Box 248, Penticton, BC, Canada V2A 6J9
(2) MDA Systems Ltd., 13800 Commerce Parkway, Richmond, BC, Canada V6V 2J3

Abstract

The 'TALON' architecture has been proposed to meet the unprecedented processing requirements and flexibility required for the Square Kilometre Array-Phase-1 (SKA1) Mid telescope, Correlator & Beamformer (Mid.CBF). The high-performance hardware platform of the TALON architecture incorporates two variants of TALON line-replaceable-units (LRUs); TALON-SX and TALON-MX. Each LRU features a single Intel Stratix-10 FPGA, 2 DDR4 DIMM modules, 4 100GE QSFP28 ports and 48 26 Gbps bi-directional optical channels that connect to a custom optical-backplane. Each LRU facilitates up to 7 TMAC/s of processing capability, 512 GB/s of memory bandwidth and 1.648 Tb/s of I/O capacity.

1. Introduction

Even in Phase 1, the Square Kilometre Array (SKA1) (www.skatelescope.org) will be the largest centimeter wavelength synthesis-array radio telescope in the world. The SKA1 'Mid-telescope' (SKA1_Mid) array is to be hosted in the Karoo region of South Africa. It consists of 133 SKA1 dish-antennas (15 m diameter) and 64 MeerKAT dish antennas (13.5 m diameter). In the SKA1_Mid, the frequency range 0.35 – 13.8 GHz is covered by 5 bands containing instantaneous bandwidths up to 5 GHz. The SKA1_Mid is required to facilitate a variety of imaging and non-imaging observation modes [1,2]. The imaging modes include continuum/spectral-line mode, zoom mode and pulsar-phase binning mode. The non-imaging modes include pulsar search (PSS) and fast radio bursts (FRB) search, pulsar timing (PST) and VLBI beamforming. It is also required to configure up to 16 sub-arrays by dynamically allocating the 197 antennas, where each sub-array would conduct any of the observations independently and many commensally [1,2].

The Central Signal Processor (Mid.CSP) is considered as a part of the 'brain' of the SKA1_Mid and consists of three sub-elements; (1) the Correlator and Beamformer (Mid.CBF) (2) Pulsar Search Engine and (3) Pulsar Timing Engine [1,2]. In particular, Mid.CBF is required to process the received signals in 'real-time' and produce visibilities for the imaging observations, pulsar-phase binning and form tied-array beams for VLBI observations, pulsar searches and pulsar timing observations [1,2]. Also, a transient buffer to capture real-time antenna data for segments of input band is required for FRB candidate search. Hence, in the SKA1_Mid.CBF, numerous signal processing operations such as re-sampling, delay-tracking, channelization, cross-correlation and beamforming [3] are required to be conducted at rates as high as 6 Giga-samples per-second (GSPS), yielding throughputs up to 200 Gb/s and occupying data storage capacities up to 128 GB.

In order to meet the tight requirements on performance, reconfigurability, power and cost of the Mid.CBF, the 'TALON' architecture has been proposed by Canada's National Research Council (NRC) (www.nrc-cnrc.gc.ca) and MDA Corporation (mdacorporation.com). The TALON architecture presents a complete solution that consists of hardware, firmware, software and mechanical designs to meet both functional and non-functional requirements of the SKA1_Mid.CBF. As a part of the hardware platform for TALON architecture, two variants of the TALON line-replaceable-units (LRUs) have been proposed to facilitate the required processing capability memory bandwidth and I/O capacity within the cost and power budgets. These LRUs are currently under development at the Dominion Radio Astrophysical Observatory (DRAO), Penticton, BC, Canada

This paper is organized as follows. An overview of the TALON architecture is presented in Section 2. In Section 3, details of the two TALON LRUs; TALON-MX and TALON-SX, are given. The summary of the current progress and future work are stated in Section 4.

2. The TALON Architecture

A simplified version of the TALON architecture for SKA1_Mid.CBF is shown in Figure 1. In general, the TALON-SX LRU is used for data handling intensive applications (i.e. PSS/PST beamformer) and the TALON-MX LRU is used for processing and memory bandwidth-intensive applications (i.e. the channelizer and the correlator). Note that the TALON architecture can support up to 200 antennas although there are only 197 in the array. It is estimated that about 1000 Talon LRUs are needed to implement the SKA1_Mid.CBF that can be hosted in 14 standard 19 inch wide and 89 inch high (i.e. 48 Rack Units) racks.

As shown in Figure 1, there are 25 F-ENGINEs and each F-ENGINE contains 16 TALON-MX LRUs and 8 TALON-SX LRUs. In SKA1_Mid bands 1, 2 and 3, each antenna generates two sample streams, corresponding to the two polarization components. These two streams are

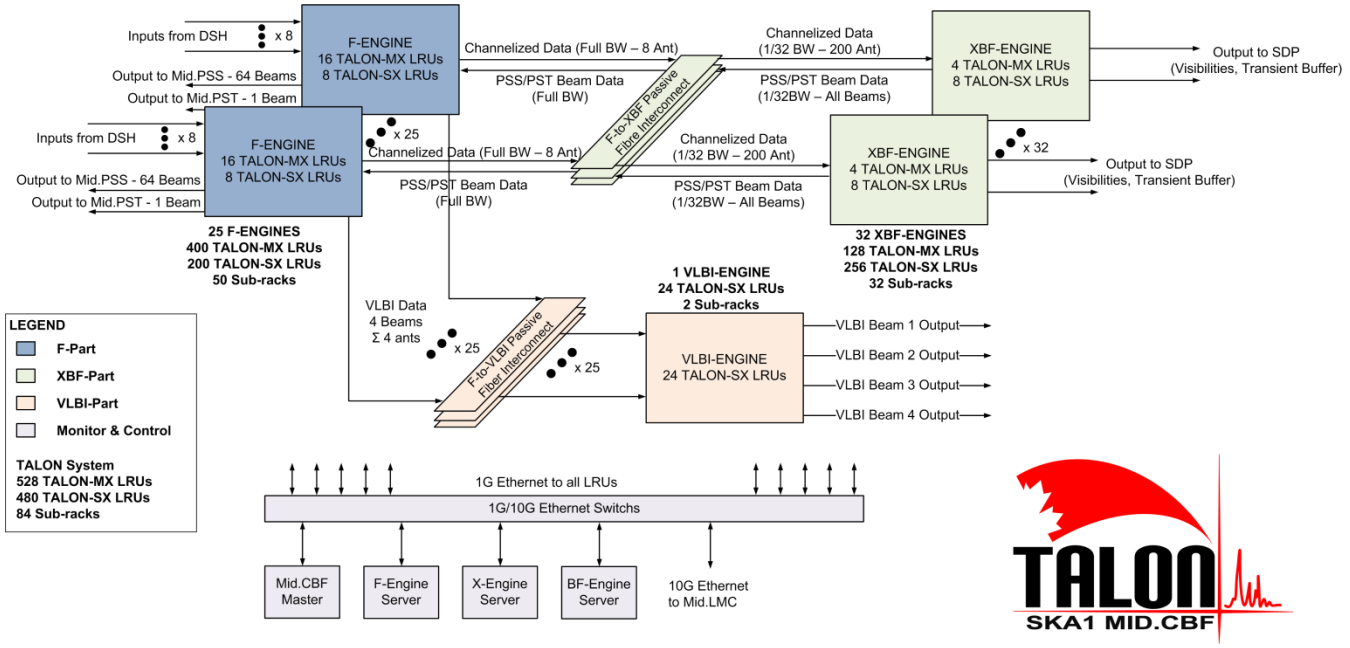

Figure 1. A simplified version of the TALON architecture for SKA1_Mid.CBF.

transmitted through a single 100GE (Giga-bit Ethernet) link to a TALON-MX LRU. For all observing modes except for VLBI beamforming, this LRU is configured to perform coarse and fine delay tracking, channelization for continuum/spectral-line mode and zoom mode imaging, or pulsar phase binning mode, band extraction for FRB search, channelization for PSS and PST beamforming. For VLBI beamforming this LRU is configured to steer the data to form up to 4 VLBI beams using coarse and fine delay corrections and to channelize the steered data in to standard VLBI bandwidths. In SKA1_Mid band 4, each antenna generates two sample streams corresponding to the two polarization components, whereas in SKA1_Mid band 5 each antenna generates four sample streams corresponding to two dual-polarization sub-bands. All two/four streams are transmitted through a single 100GE link to a TALON-MX LRU. Note that due to the increased bandwidth in band 4 and 5 two TALON-MX LRUs are required to perform the same processing.

In each F-ENGINE, channelized data from all TALON-MX LRUs are distributed to the eight TALON-SX LRUs. This is the first stage of the processing bandwidth distribution, with each TALON-SX LRU handling $1/8^{th}$ of the bandwidth for all eight antennas handled by the F-ENGINE. The TALON-SX LRUs also apply instrumentation polarization corrections to PSS and PST channelized data. From the TALON-SX LRUs in the F-ENGINE, channelized data for cross-correlation and PSS/PST beamforming are distributed to XBF-ENGINES via the F-to-XBF passive optical fiber interconnect. The channelized data for VLBI beamforming is fed to the F-to-VLBI passive optical fiber interconnect to be delivered to the VLBI-ENGINE, where VLBI beamforming is performed using 24 TALON-SX LRUs.

A single XBF-ENGINE receives channelized data from all antennas containing $1/32^{nd}$ of bandwidth via the fiber interconnect. In each XBF-ENGINE, the cross-correlations for continuum/spectral-line mode and zoom mode imaging or pulsar phase binning mode or calibration solutions for VLBI beamforming mode are carried out by four TALON-MX LRUs. Also in each XBF-ENGINE, both PSS and PST beamforming and post beamforming instrumentation polarization corrections for PST beams are carried out in the 8 TALON-SX LRUs.

From XBF-ENGINES, the cross-correlated visibilities for imaging along with the transient buffer data for FRB search are packaged and sent via 100GE links to the Science Data Processor (SDP) element of the SKA1_Mid telescope. However, the PSS and PST beams are fed back through the fiber interconnect to F-ENGINES in order to recombine the segments of bandwidth processed by different XBF-ENGINES before being sent to the PSS-Engine and the PST-Engine, respectively. SKA1_Mid.CBF generates up to 16 PST beams in Bands 1 – 3, or 8 PST beams in Bands 4 and 5. Each PST beam (bandwidth up to 2.5GHz in Band 5) is sent to the PST-Engine via two 100GE links. SKA1_Mid.CBF generates 1500 PSS beams (~300 MHz bandwidth each) and connects to the PSS-Engine via 750 10GE links.

In line with the Telescope Manager (TM) element, the monitor and control functions of the Mid.CBF are handled using the TANGO control protocol [6]. The Mid.CBF Master, a Linux server, facilitates management of processing in the LRUs with delay polynomial evaluations, gain settings and fault monitoring and alarm setting. These servers communicate with processors embedded in each of the FPGAs in the LRUs. The monitor and control network is a combination of 1GE links to each LRU and 10GE links between the Mid.CBF

Master server and each processing rack top-of-rack switch.

The Mid.CBF is expected to consume about 400 kW. Due to its high compute/power density, high thermal-efficiency liquid cooling system is proposed for the Mid.CBF system.

3. TALON LRUs

There are 2 variants of the TALON LRUs; TALON-SX LRU and TALON-MX LRU. The TALON-SX LRU contains an Intel Stratix-10 SX2800 FPGA [4], whereas the TALON-MX LRU contains an Intel Stratix-10 MX2100 FPGA [5]. The main difference between the SX2800 FPGA and the MX2100 FPGA is that the MX2100 FPGA contains the High Bandwidth Memory (HBM) blocks within the FPGA package. The on-chip HBMs provide extremely-high memory bandwidth up to 512 GB/s for each DDR stack and the total memory capacity up to 16 GB. Also, the two FPGAs are different in terms of the number of Logic Elements (LEs), DSP multipliers and internal memory blocks. The SX2800 FPGA has 2.8 million LEs, 11,520 18×19-bits multipliers and 11,721 M20k memory blocks [4] where the MX2100 FPGA has 2 million LEs, 7,488 18×19-bits multipliers and 6,501 M20k memory blocks [5]. However, both the SX2800 FPGA and the MX2100 FPGA contains 64 26G-SERDES transceivers [4,5].

Both TALON LRUs form identical 156.2 mm high, 35.5 mm wide and 340 mm deep modules that plug into a standard 4U-high 19 inch shelf. In addition to the FPGA, each LRU contains 2 DDR4 DIMM modules, 4 100GE QSFP28 ports, 4 26G 12-Channel bi-directional FCI LEAP Mid-Board Optical (MBO) modules, a 48-Channel bi-directional Molex HBMT optical backplane connector, as well as other support devices. A functional block diagram of both TALON LRUs is shown in Figure 2.

There are 4 100GE QSFP28 ports located at the front plate of both LRUs. Each port can be configured as a single 40/100GE link or four 10/25GE links. These ports are intended to interface with the DISH, SDP and PSS/PST Engines.

Each TALON LRU contains 2 DDR4 DIMM modules each supporting a 64-bit wide transfer rate of 2400 MT/s. With density up to 128 GB per module, the DDR4 DIMM modules provide the memory density and bandwidth for capturing transient buffer data up to 30 seconds. The DDR4 DIMM modules are also used by the FPGA on-chip processor to store the delay polynomials and various signal-path gain settings. Dual RJ45 1GE ports, which are located at the front plate, provide the interface to the monitor and control subsystem of the LRU. Component layout diagram of a TALON LRU is shown in Figure 3.

Each of the TALON LRU supports several configuration options for the FPGA. These include JTAG and Intel

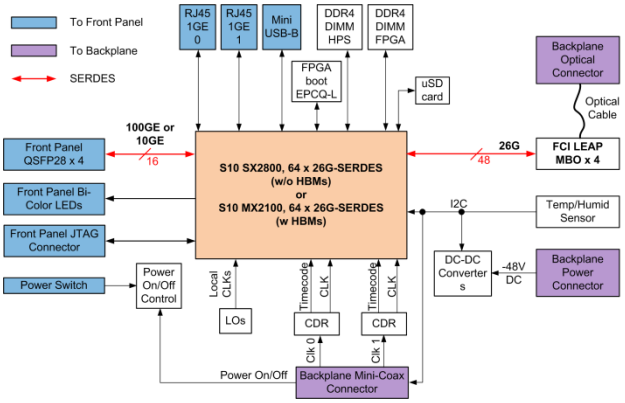

Figure 2. Functional block diagram of the TALON LRUs.

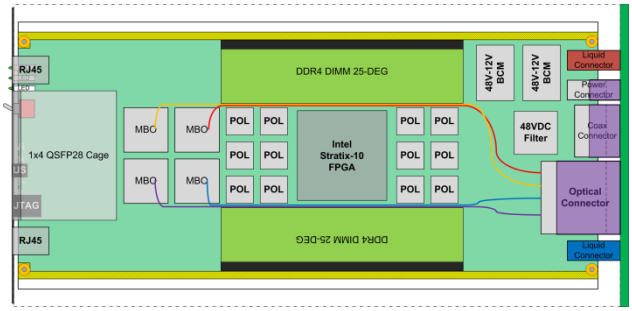

Figure 3 Component layout diagram of the TALON LRU.

QSPI configuration, along with full or partial-reconfiguration from the embedded processor in the FPGA.

The backplane 48 VDC connector delivers power to the TALON LRUs. The power consumption of the TALON LRU can be as high as 380 W. A liquid-cooling system is proposed for efficiently removing heat from the LRUs in order to keep the junction temperature in the FPGA below 85° C. As shown in Figure 4, the assembly of the TALON LRU includes the TALON PCB assembly, the cold plate, the rear cover, front panel/cover and the back cover. All the components with high power dissipations are placed on the top side of the PCB such that the cold plate is in direct thermal contact with them. The de-coupling capacitors for the power supplies are placed on the back side of the PCB. For those, heat dissipates through the PCB itself.

The cold plate contains an embedded cold pipe. Cooling liquid flows through the cold pipe and removes the heat from the PCB. Figure 5 shows the cold plate with embedded cold pipe.

The Mid.CBF is required to comply with the Electromagnetic interference (EMI) standard CISPR – 22 (Class-A) [7]. Hence, as shown in Figure 4, the PCB assembly is shielded by the front panel, back-plate, rear cover and the cold-plate from the top. Additional shielding is provided by the assembly rack.

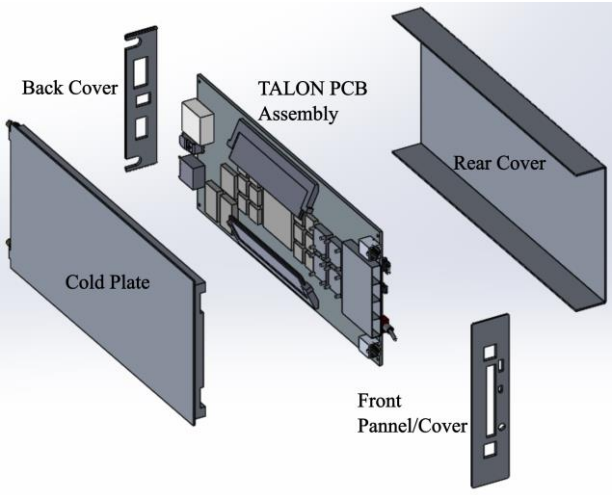

Figure 4. Assembly of a TALON LRU with the cover and the cold-plate.

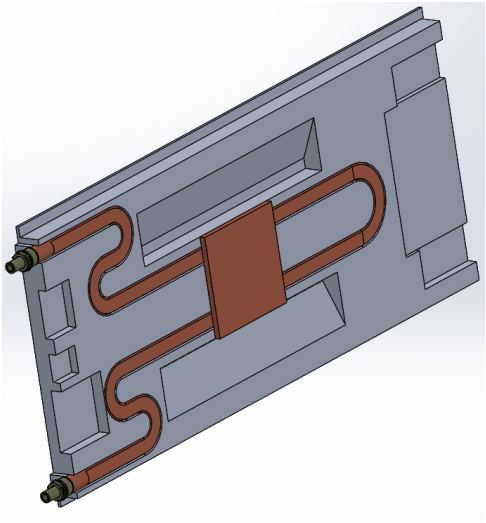

Figure 5. Cold plate with an embedded cold pipe.

4. Summary and Future Work

The Canadian NRC and MDA collaboration has proposed the TALON architecture for the SKA1_Mid.CBF. This architecture features hardware, firmware, software and mechanical designs to meet the functional and non-functional requirements. The hardware design for the TALON architecture mainly consists of two LRUs; TALON-SX that contains a Stratix-10 SX2800 FPGA and TALON-MX that contains a Stratix-10 MX2100 FPGA with HBM memory blocks. In addition, each of these LRUs contains 2 DDR4 DIMM modules, 4 100GE QSFP28 ports and 48 26 Gbps bi-directional optical channels that connect to a custom optical-backplane. Up to 7 TMAC/s of processing capability, 512 GB/s of memory bandwidth and 1.648 Tb/s of I/O capacity can be facilitated by each LRU.

The PCBs for the two LRUs are currently being designed at the DRAO. The engineering samples of Stratix-10 SX2800 FPGA are currently available from Intel. The first prototypes of the TALON-SX LRU are expected to be delivered by the middle of 2017.

5. Acknowledgements

We would like to acknowledge our counterparts at CSIRO Astronomy (www.csiro.au/en/Research/Astronomy) and ASTRON (www.astron.nl) for their valuable critical reviews of the TALON architecture.

6. References

1. P. E. Dewdney, SKA1 SYSTEM BASELINE DESIGN (SKA-TEL-SKO-DD-001), Rev.-01, 2013-03-12. (Available online https://www.skatelescope.org/key-documents/)

2. P. E. Dewdney, SKA1 SYSTEM BASELINE V2 DESCRIPTION (SKA-TEL-SKO-0000308), Rev.-01, 2015-11-04. (Available online https://www.skatelescope.org/key-documents/)

3. A. R. Thompson, J. M. Moran, and G. W. Swenson, *Interferometry and Synthesis in Radio Astronomy*, Ed.-02. New York: Wiley-VCH, 2001.

4. Stratix 10 GX/SX Family Overview Table, Altera (Gen-1032-1.1). (Available online https://www.altera.com/content/dam/altera-www/global/en_US/pdfs/literature/ pt/stratix-10-mx-product-table.pdf)

5. Stratix 10 MX Family Overview Table, Altera (Gen-1023-1.3). (Available online https://www.altera.com/content/dam/altera-www/global/en_US/pdfs/literature/pt/stratix-10-product-table.pdf)

6. The TANGO Control System Manual, Ver. 9.2, January 2016. (Available online http://www.tango-controls.org/)

7. EN 55022:2010 Information technology equipment – Radio disturbance characteristics – Limits and methods of measurement (Available Online http://www.rfemcdevelopment.eu/en/en-55022-2010)