TRIDENT-CBF

TRIDENT Correlator-Beamformer for the ngVLA
Preliminary Design Specification
TR-DS-000001

TRIDENT Document Number TR-DS-000001
File Type: ..... Preliminary Design Specification
Revision: 1
Author ..... Brent Carlson, Mike Pleasance
Date 2018-09-19
Status ..... RELEASE
Classification ..... Unrestricted

Prepared By	Name	Brent Carlson	Signature	X Brent Carlson Brent Carlson Signed by: Carlson, Brent
	Organisation	NRC		
Reviewed By	Name	Michael Rupen	Signature	X Michael P. Rupen Michael Rupen Signed by: Rupen, Michael
	Organisation	NRC		
	Name	Mike Pleasance	Signature	X Mike Pleasance Michael Pleasance Signed by: Pleasance, Michael
	Organisation	NRC		
Approved By	Name	Luc Simard	Signature	X Luc Simard Luc Simard Director, HAA-ATD Signed by: Simard, Luc
	Organisation	NRC		

National Research Council Canada | Conseil national de recherches Canada
Canada

Document History

Revision	Date	Changes/Notes	Author
A	2018-06-01	TOC only, release for review	B. Carlson
B	2018-08-31	Prelim DRAFT for initial NRC-internal review	B. Carlson
C	2018-08-30	Prelim DRAFT for external review	B. Carlson
D	2018-09-11	Prelim DRAFT: update to new requirements spec of 263 antennas, 10k km baselines, and 5/10 GHz 8+8b/4+4b sampling, among others.	B. Carlson
E	2018-09-18	DRAFT for final review. Flesh out some sections including thermal performance, use latest TALON-DX pricing from Mike. Fill in future power and cost estimates.	B. Carlson
1	2018-09-19	Rev. 1 release to NRAO.	B. Carlson

Table of Contents

List of Figures

List of Tables

Definitions, Acronyms, and Abbreviations

ADC – Analog to Digital Converter.

CBF – Correlator and Beamformer.

COTS – Commercial Off-the-Shelf.

CSP – Central Signal Processor.

DCLC – Direct Contact Liquid Cooling.

DDC – Digital Down Converter.

DDR – Double Data Rate.

EMI – Electro-Magnetic Interference.

EVLA – Expanded Very Large Array.

FinFET – “Fin” Field Effect Transistor. Also known as “Tri-gate” or “3D silicon”.

FPGA – Field Programmable Gate Array.

FS – Frequency Slice.

FSA – Frequency Slice Architecture.

FSP – Frequency Slice Processor.

HBM – High Bandwidth Memory.

HPBW – Half Power Beam Width.

HPS – Hard Processor System.

JM – Jones Matrix.

LRU – Line-Replaceable Unit.

M&C – Monitor and Control.

MBO – Mid-Board Optics.

MT/s – Mega Transfers per second.

MTBF – Mean Time Between Failures.

MTP – Brand of fiber connector containing x12 fibers.

ngVLA – Next Generation Very Large Array.

NRE – Non-Recurring Engineering (costs).

OSPPFB – Oversampling Poly-Phase Filterbank.

QSFP28 – Quad Small Form factor Pluggable, 28 Gbps/line (4 lines).

PDU – Power Distribution Unit.

RFI – Radio Frequency Interference.

SCFO – Sample Clock Frequency Offset.

SERDES – Serializer/Deserializer.

SFP – Small Form-actor Pluggable.

SKA1 Mid – Square Kilometer Array phase 1 Mid-frequency telescope array.

TALON – Name of the CBF for SKA1 Mid, “SKA1 Mid.CBF”.

TALON-DX – Name of the FPGA PCB assembly for SKA1 Mid CBF, providing the function/performance/cost basis for Trident-CBF.

TBC – To Be Confirmed.

TBD – To Be Defined.

Trident-CBF – The name of the proposed system described in this document.

VCC – Very Coarse Channelizer. The name of the one major processing step in the FSA.

VDIF – VLBI Data Interchange Format.

VLBI – Very Long Baseline Interferometry.

WIDAR – Wideband Interferometric Digital Architecture.

1 Introduction

This design specification defines the requirements, presents the design, and estimates the cost and power of a "Trident" correlator and beamformer ("Trident-CBF") system for the ngVLA. Key aspects of this system are as follows:

- Employs the NRC Frequency Slice Architecture (FSA), developed for the SKA1 Mid telescope array CBF (Mid.CBF), to provide a scalable high-performance system that meets and in many cases far exceeds ngVLA requirements, including 30 GHz/pol of bandwidth processing. Some corner-case requirements aren't technically met, but in practise, significant hardware savings are realized while providing considerable observing flexibility, compared with a design that strictly meets the corner cases (e.g. see TCBF-0007 in Table 4-1 on p.21).
- Poly-phase FX correlation for superior spectral performance, including ~80 dB of channel-to-channel isolation, seamless linear channel spacing across the digitized input bandwidth, tunable zoom window selection, and post-correlation channel averaging.
- Phase-delay beamforming for large numbers of beams on a ~40 km diameter restricted array aperture. True-delay beamforming for fewer beams to allow for an unlimited aperture for maximum tied-array output sensitivity.
- With antenna incoherent clocking [RD1] or SCFO (Sample Clock Frequency Offset sampling), spectral confinement of RFI, decorrelation of all digitizer clock frequency-dependent artefacts including those from interleaved sampling, and virtually unlimited SFDR[1] (Spurious Free Dynamic Range).
- Virtually 100% correlation and beamforming efficiency to facilitate maximum telescope sensitivity, even in an extreme RFI environment.
- Delay and phase tracking are performed virtually perfectly with no need for post-correlation or post-beamforming corrections.
- Beamforming Jones Matrix corrections for application of offset parabola polarization corrections.
- VLBI tied-array outputs in VLBI-standard VDIF format.
- Costed for 263 antenna, 30 GHz/pol capability but with fundamental physical architectural capable of handling up to 484 antennas with unlimited bandwidth in 10 GHz/pol swatches.
- Huge output bandwidth capability; an output switch, not included in this costing, is required to ingest and route these data to the archive.
- DCLC (Direct Contact Liquid Cooling) for maximum reliability and minimum total power use.

The estimated cost of the Trident-CBF system, using c. 2018 technology (Intel 14 nm FinFET FPGAs) is ~$130 M USD (2021 FPGA costs, 2018 dollars, see Section 13), including NRE and contingency. Power is estimated at ~1.5 MW. Considerable development has gone into these estimates largely leveraging the approximately 60 person-years of pre-construction effort for the SKA1 Mid.CBF system.

[1] Provided the antenna digitizer doesn't saturate.

If constructed in the ~2025 timeframe, the cost could reduce to a range of $108M to $121M (2018 dollars), based on FPGA cost multiplier factors of 0.5 and 0.8 and 900-1200 kW of power consumption based on power savings projections from Intel's FinFET devices available at that time.

2 Applicable and Reference Documents

2.1 Applicable Documents

The following documents at their indicated revision form part of this document to the extent specified herein.

Table 2-1 Applicable Documents

Ref No	Document/Drawing Number	Document Title	Issue Number
AD1	020.40.00.00.00-0001-SPE	ngVLA Central Signal Processor: Preliminary Technical Specifications, 09/10/2018.	08
AD2			

2.2 Reference Documents

The following documents provide useful reference information associated with this document. These documents are to be used for information only. Changes to the date and/or revision number do not make this document out of date.

Table 2-2 Reference Documents

Ref No	Document/Drawing Number	Document Title	Issue Number
RD1		Carlson, B.R., Incoherent clocking in coherent radio interferometers, *IEE Electronics Letters* (2018), 54 (14):909.	
RD2		Carlson, B., Gunaratne, T., Signal processing aspects of the sample clock frequency offset scheme for the SKA1 mid telescope array, 32nd URSI GASS, Montreal, 19-26 August 2017.	
RD3	SKA-TEL-CSP-00000069 311-000000-007	SKA1 CSP Mid Array Correlator and Central Beamformer Sub-element Signal Processing MATLAB Model (EB-7), 2018-08-16.	4
RD4		The TANGO Control System Manual	9.2

		http://ftp.esrf.fr/pub/cs/tango/tango_92.pdf	
RD5	SKA-TEL-CSP-00000066 311-000000-003	SKA1 CSP Mid Correlator and Beamformer Sub-element Detailed Design Document (EB-4a), 2018-07-26	2

3 System Overview

An overview block diagram of the Trident-CBF system is shown in Figure 3-1:

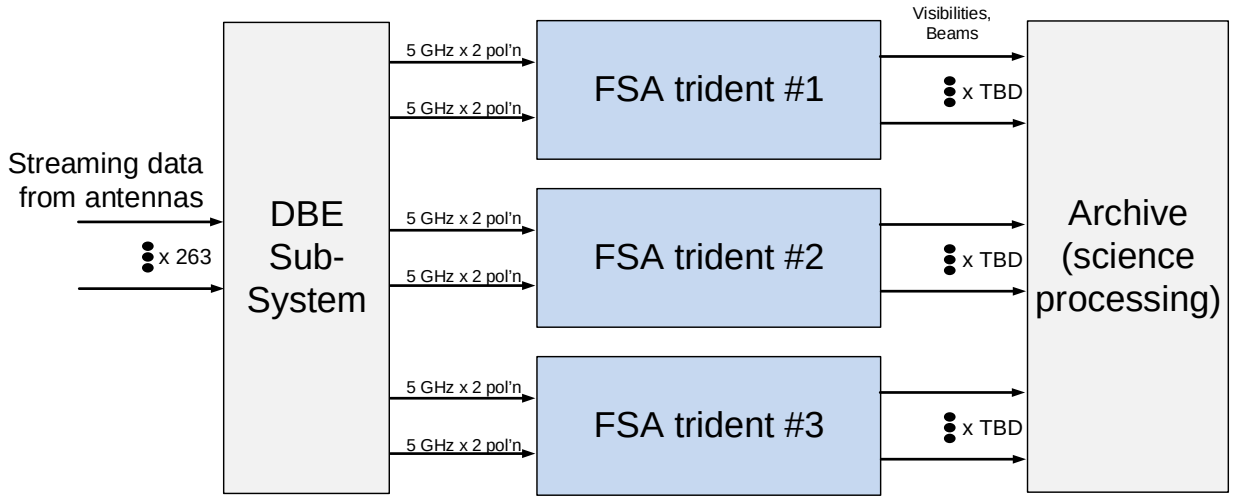

Figure 3-1 Overview block diagram of the Trident-CBF system.

In this architecture, three identical FSA tridents[2] are provided, each having 10 GHz/pol of ingest and correlation/beamforming processing capacity. The output from each trident feeds the archive via an external switch (not shown in the figure).

The fundamental data flow of *each* FSA trident is illustrated in Figure 3-2, with an associated Frequency Slice (FS) processing diagram shown in Figure 3-3. In these figures:

- The 10 GHz/pol (i.e. @ 4+4b sampling; 5 GHz/pol @ 8+8b sampling, with a separate FPGA bitstream for each such sampling mode) of ingest bandwidth is divided, via VCCs (Very Coarse Channelizers, one for each antenna), into up to 50 x 200 MHz (of processed bandwidth[3]) FSs, oversampled to 220 MHz of sampled bandwidth. Each FS is oversampled and overlaps with adjacent slices such that correlation and beamforming output data, when concatenated in frequency, are seamless across the input digitized bandwidths[4] with no undue anomalies in amplitude, phase, or sensitivity at the overlap boundaries. The VCCs additionally contain bulk coarse delay tracking[5] corrections, allowing each FSP to perform delay and phase corrections suited to its needs.

[2] Note that a "trident" refers to an individual 1/3rd of the entire system. "Trident-CBF" is the name of the entire system consisting of all 3 tridents.
[3] Processed bandwidth is defined as bandwidth, however processed, that is science-quality and delivered to the archive.
[4] The "digitized bandwidth" refers to the raw bandwidth out of a physical ADC digitizer, i.e. the antenna digitizers. When the term "sampled bandwidth" is used, it is similar but is the output from an entirely logic/digital process.
[5] Steps of many samples, updated relatively infrequently, using bulk memory external to the FPGA. By doing so, active coarse delay tracking in the FSPs is performed on-chip, greatly simplifying implementation and maximally distributing external memory bandwidth needs across VCCs and FSPs.

- Via a switch built into each VCC, each FS may be routed to any FSP (Frequency Slice Processor) within its trident, including routing the same FS to multiple FSPs, or not processing select FSs at all. This affords considerable flexibility in how bandwidth is processed.
- Each FSP performs one "Function Mode" (correlation and various beamforming modes) for all sub-arrays processed in the FSP. The Function Mode configuration of each FSP is independent of all other FSPs.
- Output correlation and beamformed data products from the FSPs are routed to the archive via an external switch.

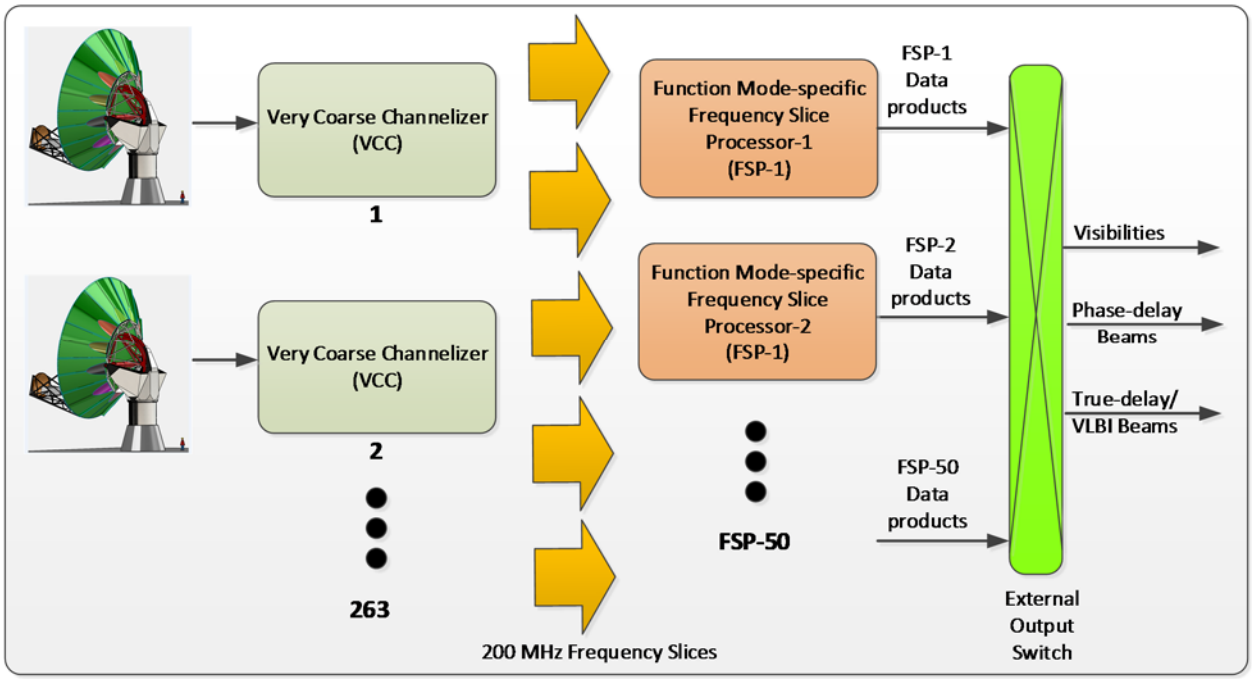

Figure 3-2 Fundamental data flow of each FSA trident.

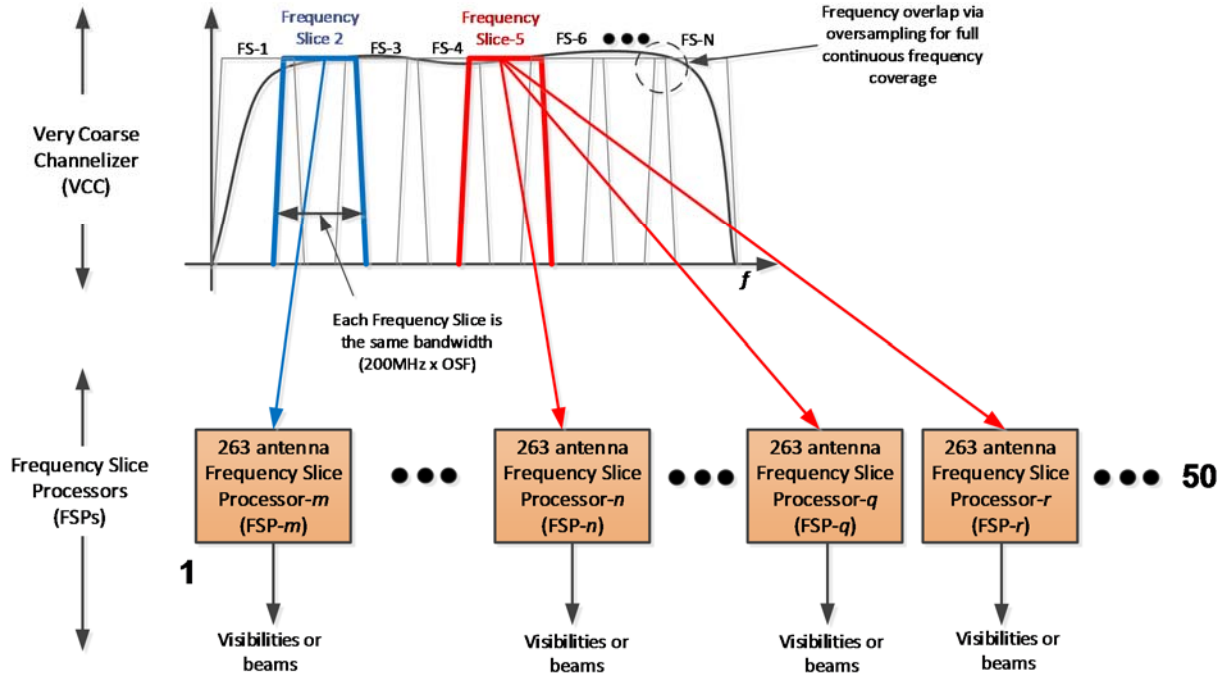

Figure 3-3 FSA Frequency Slice processing diagram for each trident.

Each FSP can be (FPGA-bitstream) independently configured to operate in exactly one "Function Mode", where these Function Modes have been specifically designed to meet ngVLA requirements [AD1], when the requisite FSPs are configured accordingly to meet required ngVLA Observing Modes [AD1]. ngVLA Trident-CBF Function Modes are briefly described below:

- **Correlation.** In this Function Mode, each sub-array performs normal channel resolution or zoom resolution cross-correlation including auto-correlation. At normal channel resolution, 16k channels across the FS sampled bandwidth of ~220 MHz are provided; at zoom resolution, for tunable filter sampled bandwidths of 220 MHz/2, /4, …, /64, the same number of channels are provided, thereby providing much higher spectral resolution. In both cases, post-correlation channel averaging can be used to reduce the delivered spectral channel resolution and therefore visibility output transport bandwidth.
- **VLBI.** In this Function Mode, Jones Matrix corrections are applied per antenna, followed by true-delay beamforming to form up to 4 dual-polarization beams with independent delay centers on the sky (within the beam of the antenna), each beam including any number of ngVLA antennas[6], with simultaneous correlation of 1k channel visibilities[7] for tied-array calibration solutions. After beamforming, Jones Matrix polarization corrections are applied to each beam, followed by tunable digital filters that can be used to select specific parts of the FS to ultimately be output as VLBI "beam-channels" at VLBI standard data rates in VDIF format to the archive.

[6] i.e. no beamforming aperture size restrictions.
[7] Note, with its own 1k critically sampled channelizer per antenna followed by correlation. i.e. visibilities are developed prior to beamforming for all antennas in the beamforming array.

- **Pulsar-True-Delay-4beam.** This is similar to the VLBI Function Mode except that after beamforming a 256-channel (up to 4k) oversampling channelizer is available on each beam with per-channel Jones Matrix corrections applied. Simultaneous correlation of 1k channel visibilities is also provided.
- **Pulsar-True-Delay-10beam-144ant.** Same as Pulsar-True-Delay-4beam, except 10 beams are produced from a pre-selected (hardwired -- i.e. specific antenna inputs into Trident-CBF) set of up to 144 antennas. Here, the number of antennas that can be beamformed is reduced to be able to increase the number of beams. Simultaneous correlation of 1k channel visibilities is also provided.
- **Pulsar-Phase-Delay-100beam.** Here, phase-delay beamforming is performed to produce a large number of beams very efficiently. As a consequence, the beamforming aperture diameter for any sub-array, to obtain required efficiency of each beam within the antenna primary beam, is restricted to ~40 km. Beamforming is performed on any number of antennas as long as the aperture diameter restriction is not exceeded. If beam offsets from boresight[8] are more restricted, the beamforming aperture can be increased accordingly. Here, no simultaneous correlation of 1k channel visibilities is performed in the FSP, relying on other FSPs being configured to provide correlation capability to develop tied-array calibration solutions.

Table 3-1 contains a list of all of the key capabilities of FSP beamforming Function Modes.

Figure 3-4 is a simplified overview signal flow diagram showing VCC processing and all of the Trident-CBF Function Modes. All processing is dual polarization. Note that each of the 3 tridents in the system:

- Provides processing for 10 GHz/pol of bandwidth.
- Provides 50 FSPs, each of which can be independently configured for each Function Mode.

[8] The main true-delay tracking delay center on the sky, typically in the center of the antenna primary beam.

Table 3-1 Trident-CBF FSP beamforming Function Modes

FSP Function Mode	Beam Steering	Aperture diameter	Beam offset from boresight	Nant	Beamforming Channelizer	Nbeams	JM per FS?	JM per ant?	JM per ch?	JM per bm?	Post-BF Channelizer	Post-BF JM?	Visibility Channelizer	Correlated vis per FS	Notes
VLBI	True-delay	no limit	Full antenna main BW	256	N/A	4/subarray (26 total)	Y	Y	N/A	Y	Multiple independently tunable DDCs, 1 for each beam-channel	Y: per-224 MHz beam	1k-CS on 1 beam	1k (across the 224 MHz OS FS)	Beam-channel output bandwidths of 1, 2, 4, 8, 16, 32, 64, 128, or 224 MHz.
Pulsar-True-Delay-4bm	True-delay	no limit	Full antenna main BW	256	N/A	4/subarray (26 total)	Y	Y	N/A	Y	256 OS	Y: per-OSchan	1k-CS on 1 beam	1k (across the ~220 MHz OS FS)	Post-beamforming channelizer could be any radix-2 up to perhaps 4k (bitstream compile-time-defined).
Pulsar-True-Delay-10bm-144ant	True-delay	no limit	Full antenna main BW	144	N/A	10/subarray (26 total)	Y	Y	N/A	Y	256 OS	Y: per-OSchan	1k-CS on 1 beam	1k (across the ~220 MHz OS FS)	144 antennas, selected from 2 broad pre-defined hardwired sub-arrays, with, for each group of 11 ants, 6 of 11 antenna selection. Post-BF channelizer could be any radix-2 to up
Pulsar-Phase-Delay-100bm	Phase-delay	40 km	Within antenna HPBW	BF-aperture limited	4k-CS	~100 tot	N	Y	Y	N	N/A	N	none	none	>40 km aperture possible: inversely proportional to beam offset from boresight. Simultaneous visibilities possible if only 10 beams are required.

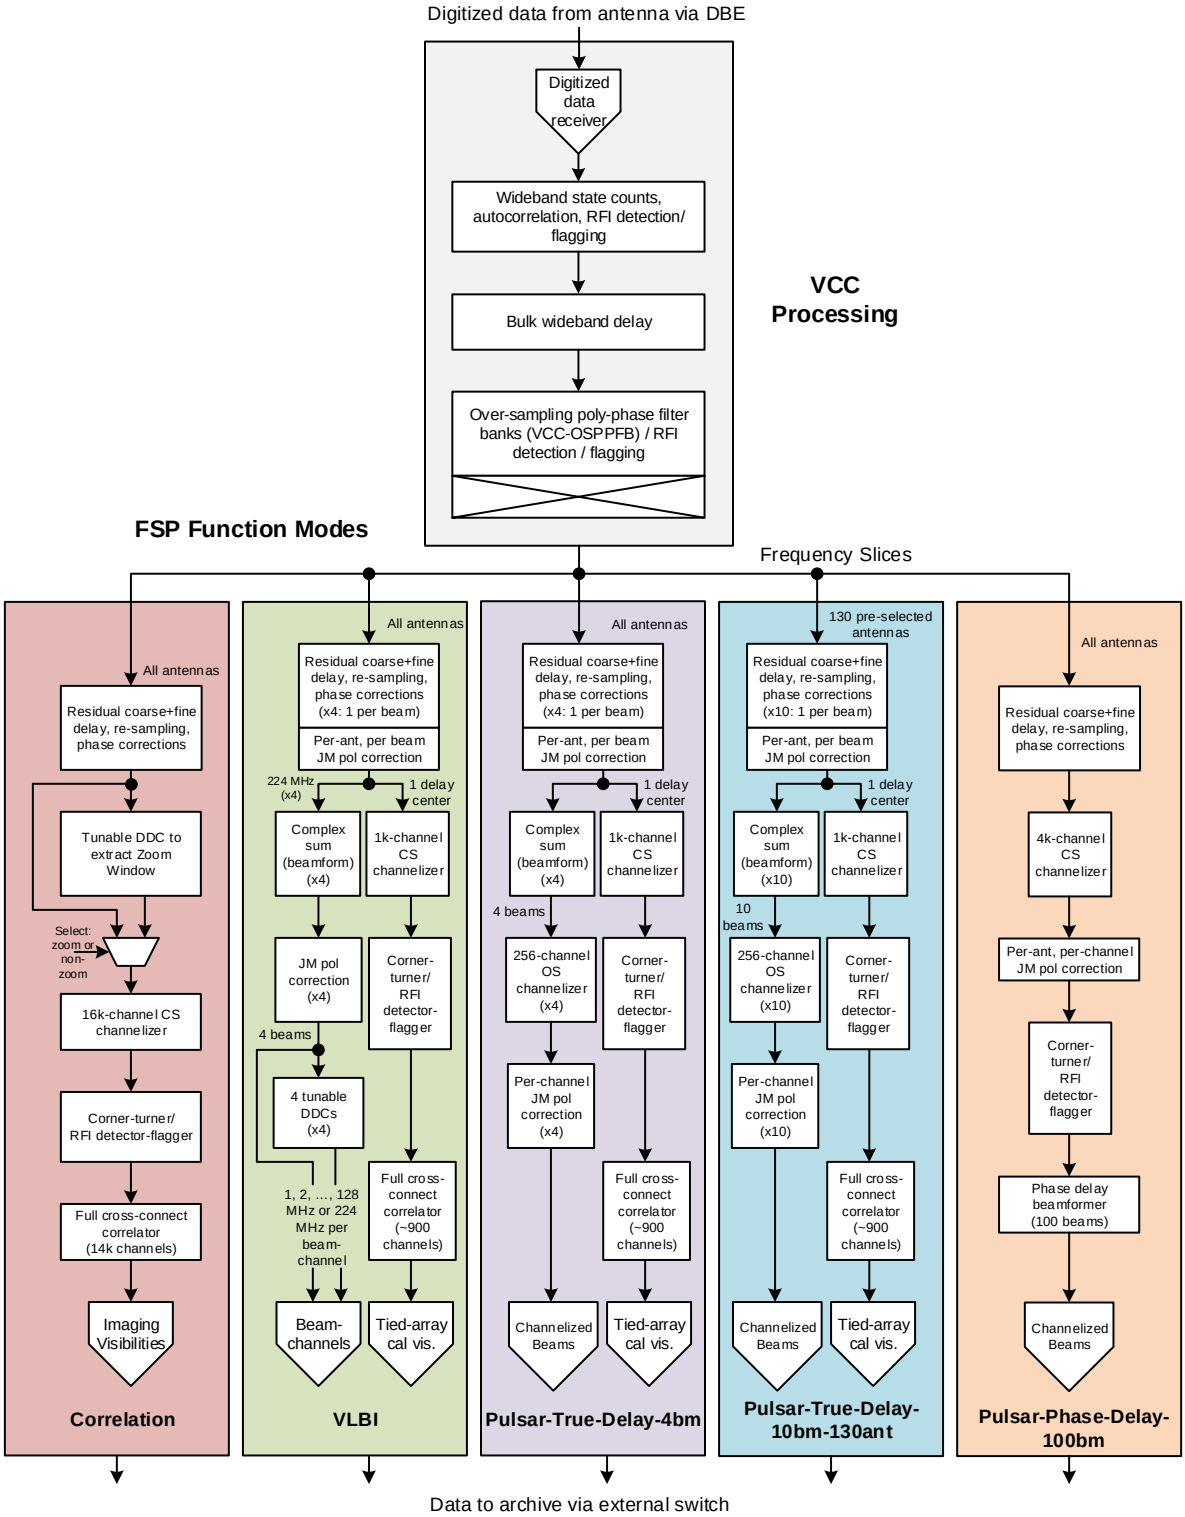

Figure 3-4 Trident-CBF simplified signal flow diagram.

4 System Requirements

This section develops requirements for the Trident-CBF system, with traceability to requirements as defined in [AD1]. Requirements are specifically defined in terms of the Trident-CBF FSA architecture outlined in Section 3, with further physical details in Section 5.3.

4.1 Functional Requirements

Table 4-1 Trident-CBF functional requirements

Trident-Req#	Requirement	Traceability [AD1]	Compliance/Comments
General			
TCBF-0001	There shall be 3 tridents in Trident-CBF, each of which ingests and processes 10 GHz/pol of single-sideband digitized data from each of at least 263 antennas via the DBE.	CSP0111 CSP0121 CSP0124 CSP0161	Compliant.
TCBF-0002	Trident-CBF shall configure its processing resources such that: • Data from each antenna digitizer is divided into multiple identical Frequency Slices whose collective contiguous processed bandwidth covers the entire digitized bandwidth and, • The processed bandwidth of each Frequency Slice is 200 MHz +/-2 MHz.	CSP0121 CSP0124 CSP0122 CSP0123 CSP0125	Compliant. Note on CSP0123: The intent in the requirement for 4 MHz tuning steps seems to be to provide contiguous bandwidth coverage across the full input digitized bandwidth.
TCBF-0003	Each trident of Trident-CBF shall contain 50 FSPs, each FSP processes any select Frequency Slice for each sub-array independently, and each FSP is independently configurable under software control to any of the following Function Modes: • Correlation, • VLBI, • Pulsar-True-Delay-4bm, • Pulsar-True-Delay-10bm-144antor, • Pulsar-Phase-Delay-100bm	CSP0121 CSP0124	Compliant.

TCBF-0004	The maximum data transport delay from any antenna to Trident-CBF shall be 250 msec.	CSP0141	Compliant. The bulk delay in the VCC is implemented with external DDR4 RAM with more than sufficient capacity.
TCBF-0005	Trident-CBF shall provide processing for a maximum ngVLA array diameter of 10,000 km.	CSP0142	Compliant. The bulk delay in the VCC is implemented with external DDR4 RAM with more than sufficient capacity.
TCBF-0006	The phase-center, for any sub-array, shall be within the smallest circle encompassing the sub-array.	CSP0143	Compliant.
TCBF-0007	Trident-CBF shall provide sub-arraying capability, for at least 10 independent simultaneous sub-arrays in each FSP for any Function Mode.	CSP0181 CSP0322 CSP0323 CSP0324	Partially compliant. Sub-arraying is independent in each FSP, for the configured Function Mode. The corner-case of each sub-array performing a different Observing Mode, each for the maximum aggregate bandwidth, cannot be met.
TCBF-0008	Trident-CBF shall automatically detect and flag RFI in time and frequency, at each frequency resolution available through each signal flow.	CSP0171 CSP0172	Compliant.
TCBF-0009	For all output correlation and beamforming data products, as appropriate, Trident-CBF shall calculate, and provide on request via monitor and control channels, all internally-determined and calculated gain correction factors to be applied in post-processing.	CSP0153	Compliant.
Function Mode: Correlation (for synthesis imaging)			
TCBF-0101	The normal (non-zoom) Correlation Function Mode shall provide:	CSP0222 CSP0221	Partially compliant.

	• 14,000 to 15,000 critically-sampled channels across the Frequency Slice processed bandwidth. • Post-correlation channel averaging factors of at least 2, 3, 4, 6, and 8. • Linearly-spaced channels, with or without channel averaging, across the processed bandwidth of the Frequency Slice and any number of adjacent Frequency Slices sourcing from the same digitizer. • Per-visibility integration times ranging from 2 msec to 10 sec, at least in steps of the minimum integration time, and such that integration times across all FSPs, across all tridents, that are part of the same sub-array are synchronized in time. • A phase-reference position on the sky (a.k.a the boresight delay center), which is within the antenna primary beam HPBW and independent of any other FSP.	CSP0225 CSP0231	14k to 15k channels provides a channel bandwidth of 14.2 kHz to 13.3 kHz. Channel averaging of 8 provides 114 kHz channel resolution; wider channels (likely) must be produced external to Trident-CBF. Larger channel averaging factors require further investigation. Linear channels across multiple adjacent Frequency Slices is not specifically in [AD1], however it is implied by CSP0222. Without the re-sampling function, it would be unwieldy or impossible to do so in a 2-stage oversampling poly-phase filterbank.
TCBF-0102	Correlation Function Mode, when correlating Zoom Windows, shall provide: • At least 1 tunable Zoom Window per antenna, with a tuning resolution of at least 10 kHz anywhere in the Frequency Slice processed bandwidth. • Zoom Window processed bandwidths of /2, /4, /8, /16, /32, /64 of a Frequency Slice processed bandwidth. • Number of channels, channel averaging, integration times, and independent phase-reference position identical to non-zoom visibilities. Linear channel	CSP0221 CSP0222 CSP0223 CSP0224 CSP0231	Compliant down to a channel width of 220 Hz (without averaging).

	spacing applies only to channels within the Zoom Window.		
TCBF-0103	In Correlation Function Mode, each sub-array in an FSP shall be independently configurable to be correlating zoom *or* non-zoom visibilities; if zoom, then each sub-array must have a single Zoom Window processed bandwidth, not dependent on any other sub-array's zoom or non-zoom configuration.	CSP0181	
Function Mode: VLBI			
TCBF-0201	In VLBI Function Mode, each FSP shall form at least 4 beams per-sub-array, on at least 1 VLBI sub-array, and each beam: • Has a contiguous, non-channelized, processed bandwidth of a Frequency Slice. • Is independently steerable to delay center on the sky anywhere within the HPBW of the antenna primary beam. • Can be formed from any select ngVLA antennas of any beamforming aperture within the sub-array.	CSP0521 CSP0522	Compliant. At least 10 simultaneous sub-arrays are possible, each one producing 4 beams, each beam independently steerable per sub-array.
TCBF-0202	In VLBI Function Mode, simultaneous with beamforming in each VLBI sub-array each FSP shall: • Produce 900 to 1000 supporting visibilities for at least 1 beam (i.e. delay center), across the Frequency Slice processed bandwidth. • Have a per-visibility integration times in the range of 20 msec to 5 sec, synchronized in time within each sub-array across all FSPs and tridents.	CSP0531 CSP0532 CSP0533	Compliant. [AD1] is TBD on CSP0532 and CSP0533; these are achievable values.
TCBF-0203	In VLBI Function Mode, the following Jones Matrix (JM) beam polarization corrections shall be applied:	CSP0541	Compliant. CSP0541 is TBD in [AD1].

	• Prior to beamforming, per-antenna, per-beam, on the Frequency Slice sampled bandwidth. • Post-beamforming, per-beam on the Frequency Slice sampled bandwidth (TBC)		
TCBF-0204	In VLBI Function Mode, at least 4+1 beam-channels shall be produced per sub-array with: • 4 beam-channels having at least 2 kHz tuning resolution anywhere within the Frequency Slice sampled bandwidth (of 224 MHz), with bandwidths of 1 MHz x 2^n (n=0, 1, …, 7). • 1 beam-channel fixed at the Frequency Slice sampled bandwidth of 224 MHz. • Real output. • Nyquist output sample rates. • Independently settable 2-bit, 4-bit, 8-bit, or 16-bit word size per beam-channel. • VDIF formatted output.	CSP0551 CSP0552 CSP0553 CSP0554	Compliant. In [AD1] the number of beam-channels output is TBD. The Trident-CBF design can provide up to an aggregate of 104+26 beam-channels per FSP.
Function Mode: Pulsar-True-Delay-4bm			
TCBF-0301	In Pulsar-True-Delay-4bm Function Mode each FSP shall produce at least 4 true delay beams arbitrarily allocated to any sub-array, and each beam: • Has a contiguous, non-channelized, processed bandwidth of a Frequency Slice. • Is independently steerable to delay center on the sky anywhere within the HPBW of the antenna primary beam. • Can be formed from any select ngVLA antennas of any beamforming aperture within the sub-array.	CSP0321 CSP0121 CSP0324 CSP0611	Partially compliant. Multiple FSPs allocated to the same Frequency Slices are required. The total beams x bandwidth product is 4 beams/FSP x 50 FSPs x 200 MHz/FS = 40 GHz. Thus, 50 beams at 0.8 GHz/beam can be formed. This is for the full array, any number of ngVLA antennas, any aperture size

	• Post-beamforming, before output, channelized to at least 128 oversampled channels across the Frequency Slice sampled bandwidth.		
TCBF-0302	In Pulsar-True-Delay-4bm Function Mode, the following Jones Matrix (JM) beam polarization corrections shall be applied: • Prior to beamforming, per-antenna, per-beam, on the Frequency Slice sampled bandwidth. • Post-beamforming, per-beam on each output oversampled channel independently.	CSP0331	Compliant.
TCBF-0303	In Pulsar-True-Delay-4bm Function Mode, simultaneous with beamforming in each Pulsar-True-Delay-4bm sub-array each FSP shall: • Produce 900 to 1000 supporting visibilities for at least 1 beam (i.e. delay center), across the Frequency Slice processed bandwidth. • Have a per-visibility integration times in the range of 20 msec to 5 sec, synchronized in time within each sub-array across all FSPs and tridents.	CSP0341	Compliant.
Function Mode: Pulsar-True-Delay-10bm-144ant			
TCBF-0401	In Pulsar-True-Delay-10bm-144ant Function Mode each FSP shall produce at least 10 true delay beams pre-allocated to a select 144 antennas, and each beam: • Has a contiguous, non-channelized, processed bandwidth of a Frequency Slice. • Is independently steerable to delay center on the sky anywhere within the HPBW of the antenna primary beam.	CSP0321 CSP0121 CSP0423 CSP0611	Partially compliant. CSP0423 requires 168 antennas but only 144 are believed possible. Multiple FSPs allocated to the same Frequency Slices are required. The total beams x bandwidth product is 10 beams/FSP x 50 FSPs x 200 MHz/FS =

	• Can be formed from any of the pre-selected ngVLA antennas of any beamforming aperture within the sub-array. • Post-beamforming, before output, channelized to at least 128 oversampled channels across the Frequency Slice sampled bandwidth.		100 GHz. Thus, 50 beams at 2 GHz/beam can be formed. The current design/resource utilization indicates that only 130 pre-selected antennas can provide this capability (CSP0423 requires 168).
TCBF-0402	In Pulsar-True-Delay-10bm-144ant Function Mode, the following Jones Matrix (JM) beam polarization corrections shall be applied: • Prior to beamforming, per-antenna, per-beam, on the Frequency Slice sampled bandwidth. • Post-beamforming, per-beam on each output oversampled channel independently.	CSP0331	Compliant.
TCBF-0403	In Pulsar-True-Delay-10bm-144ant Function Mode, simultaneous with beamforming in each Pulsar-True-Delay-10bm-144ant sub-array each FSP shall: • Produce 900 to 1000 supporting visibilities for at least 1 beam (i.e. delay center), across the Frequency Slice processed bandwidth. • Have a per-visibility integration times in the range of 20 msec to 5 sec, synchronized in time within each sub-array across all FSPs and tridents.	CSP0341	Compliant.
Function Mode: Pulsar-Phase-Delay-100bm			
TCBF-0501	In Pulsar-Phase-Delay-100bm Function Mode each FSP shall produce at least 100 phase-delay beams, and each beam: • Has a contiguous processed bandwidth of a Frequency Slice.	CSP0421 CSP0422 CSP0423 CSP0424 CSP0121	Compliant. This exceeds the CSP0421 requirement of 10 beams; if only 10 beams are produced, then

	• Is independently steerable to delay center on the sky anywhere within the HPBW of the antenna primary beam. • Can be formed from any up to 168 antennas with a beamforming aperture diameter of up to 30 km. • Channelized to 4096 critically-sampled channels across the Frequency Slice sampled bandwidth.		simultaneous supporting visibilities are possible.
TCBF-0502	In Pulsar-Phase-Delay-100bm Function Mode, the following Jones Matrix (JM) beam polarization corrections shall be applied: • Prior to beamforming, per-antenna, per-channel.	CSP0431	Compliant.

4.2 Interface Requirements

Table 4-2 Trident-CBF interface requirements

Trident-Req#	Requirement	Traceability [AD1]	Compliance/Comments
TCBF-1001	Each trident of Trident-CBF shall interface to each of 263 antennas as either: • 10 GHz/polarization at 4+4b complex per sample or, • 5 GHz/polarization at 8+8b complex per sample.	CSP0111 CSP2101 CSP2102 CSP2103	Compliant.
TCBF-1009	Trident-CBF shall automatically recover and require no on-sky calibration in the event of a packet loss or loss of link on any or all antenna data ingest pipelines.	CSP0144	Compliant. This requires that the ingest data includes an embedded 1PPS marker with a pre-determined and unchanging number of

			samples between 1PPS markers, which could be different for each antenna (i.e. to support SCFO sampling), and that synchronous counting/timing occurs in the antenna even though the antenna-CBF link sample data link has packet loss or has dropped.
TCBF-1002	Each of the 150 FSPs in Trident-CBF shall provide a data product output bandwidth of at least 40 Gbps.	CSP2201 CSP2202	Partially compliant. Trident-CBF require a 27 x 10/40G Ethernet switch for each FSP, which is currently not costed.
TCBF-1003	Trident-CBF shall receive and synchronize all operations to a 100 MHz and 1-PPS reference. Note: all antenna sampled inputs are frequency-synchronous with these references.	Section 3.4.6	Compliant.
TCBF-1004	Trident-CBF control and monitor interface shall be multi-mode fiber using TCP/IP over Ethernet.	Section 5.6	Compliant.
TCBF-1005	Trident-CBF shall provide an electrical power interface to 3-phase, 208 VAC (delta), 60 Hz power.	Section 5.4.1	Compliance expected.
TCBF-1006	Trident-CBF shall be water-based liquid cooled and the interface for such cooling shall be: • Per-rack water source and return at TBD L/s at a source temperature of TBD C per rack. • Fittings sizes TBD.	Section 5.4.2	Compliant.
TCBF-1007	Trident-CBF racks shall be standard 19” racks, each with maximum dimensions of 24” W x 42” D x 42U high.	TBD (trace to infrastructure ICD)	Compliant.

TCBF-1008	The complete Trident-CBF installation shall consist of TBD 19” racks.	TBD (trace to infrastructure ICD)	

4.3 Control and Monitor Requirements

Table 4-3 Trident-CBF control and monitor requirements

Trident-Req#	Requirement	Traceability [AD1]	Compliance/Comments
TCBF-1101	Trident-CBF shall monitor and report a set of parameters that allow for determination of its status.	CSP1001	Compliant.
TCBF-1102	Trident-CBF shall monitor and detect 99.9% of communications faults whilst operationally available.	CSP1001	Compliant.
TCBF-1103	Trident-CBF shall perform in-situ off-line (i.e. not operationally available) tests on a per-antenna and per-FSP basis to detect processing faults with 99.9% probability of detection.	CSP1001	Compliant. The idea here is to roll testing across the system so as to not have to take the entire system down to run such tests.

4.4 Performance Requirements

Table 4-4 Trident-CBF performance requirements

Trident-Req#	Requirement	Traceability [AD1]	Compliance/Comments
TCBF-1201	In FSP Correlation Function Mode normal and zoom visibilities shall: • Have a -3 dB +/-0.01 dB channel edge amplitude.	CSP0151 CSP0152 CSP0153 CSP0131	Compliance expected. Further analysis required to verify that CSP0151 -50 dB leakage is met.

	• Be monotonically decreasing in any region above -60 dB from the channel edge to the next adjacent channel center frequency. • Have an amplitude of -80 dB or better thereafter (i.e. reject band). • Have a post-gain factor correction amplitude that varies at most by +/-0.01 dB across the processed bandwidth across all Frequency Slice of a digitized bandwidth. • Provide at least 98.6% correlation efficiency where there is no more than 10% RFI power present in any particular channel.		
TCBF-1202	In FSP VLBI Function Mode simultaneously-produced visibilities shall: • Have visibility channel response: maximum +/-0.1 dB amplitude and +/-0.01 radian phase passband (channel edge-to-edge) ripple; -3 dB +/-0.2 dB channel edge amplitude; monotonically decreasing amplitude from the channel edge to the adjacent channel far edge; at least -60 dB reject-band amplitude anywhere greater than the 1 channel width away from the channel edge. • Have at least 98.6% correlation efficiency when in-channel RFI power is <10%.	CSP0131 CSP0151 CSP0152	Partially compliant. CSP0151: -50 dB total power leakage probably not achievable (requires more logic+memory), TBC.
TCBF-1203	Trident-CBF shall apply wavefront delay model corrections such that the coherence loss resulting from such corrections is <0.1%.	CSP0131	Compliant.
TCBF-1204	Trident-CBF shall apply wavefront phase model corrections such that the	CSP0131	Compliant.

	coherence loss resulting from such corrections is <0.1%.		
TCBF-1205	Trident-CBF shall apply delay and phase model corrections such that there is no discernable correlated phase or amplitude closure anomalies on any integration time scale above the expected noise.	TBD (missing requirement)	Compliant.
TCBF-1206	Trident-CBF shall apply delay and phase model corrections on true delay beams such that there is no discernible phase or amplitude variation due to such application over the expected noise.	TBD (missing requirement)	Compliant.

4.5 Environmental Requirements

Table 4-5 Trident-CBF environmental requirements

Trident-Req#	Requirement	Traceability [AD1]	Compliance/Comments
TCBF-1301	The operating ambient temperature for Trident-CBF shall be in the range of 20 C to 30 C with a maximum rate of change of temperature of TBD C/Hr and humidity of TBD % relative.	CSP0811 CSP0812 CSP0813	Compliance expected. DCLC largely buffers electronics' dT/dt against room ambient temperature changes.
TCBF-1302	Trident-CBF shall self-protect and shutdown if a destructive over-temperature condition is reached.	Section 4.8.1	Compliant.
TCBF-1303	Trident-CBF packaging and installation shall be designed to survive, with no permanent damage, an earthquake event of up to 1 per year with up to 0.2 g peak acceleration in either the vertical or horizontal axis.	CSP0831	Compliance expected.

TCBF-1304	Trident-CBF shall operate up to an altitude of 2200 m above sea level.	CSP0841	Compliance expected.
TCBF-1305	The electrical power quality provided to Trident-CBF shall be such that Trident-CBF need not itself provide any battery backup or dirty power conditioning other than normally provided in COTS server AC-DC power supplies.	Section 5.4.1	Compliant.
TCBF-1306	At least 90% of Trident-CBF power shall be removed using water cooling.	TBD (trace to infrastructure ICD)	Compliance expected. This allows the water and air-cooling systems' requirements to be determined.
TCBF-1307	Each LRU in Trident-CBF shall, as a minimum, meet FCC Part 15, Sub-part J, Class A radiated and conducted emissions standards.	Section 5.4.3 Section 8.2	Compliant. Reasonable EMI design measures required.
TCBF-1308	The Trident-CBF installation shall comply with all relevant federal and State of New Mexico building and electrical codes.	CSP6001 CSP6002	Compliant. Assume meeting UL and U.S./Canada electrical codes is sufficient.

4.6 Reliability and Maintenance Requirements

Table 4-6 Trident-CBF reliability requirements

Trident-Req#	Requirement	Traceability [AD1]	Compliance/Comments
TCBF-1401	Trident-CBF shall have an inherent availability (Ai) of at least 99.9% while Operationally Capable. Note: Trident-CBF is Operationally Capable when it can produce 99% of required data products using at least 99% of the antennas.	CSP0901	Compliance TBD. Further analysis is required, but these numbers are based on a single TALON LRU failure.

4.7 Other Requirements

Table 4-7 Trident-CBF other requirements

Trident-Req#	Requirement	Traceability [AD1]	Compliance/Comments
TCBF-1501	Trident-CBF shall have a design-lifetime of at least 20 years.	CSP1101	Compliance expected.
TCBF-1502	The design of Trident-CBF shall minimize its life-cycle cost for 20 years of operation.	CSP1102	Non-compliant. Fundamentally impossible to do this; "minimal" is achievable, but beyond the current work scope to determine.
TCBF-1503	All materials used in the Trident-CBF installation shall be RoHS-compliant.	Section 7.2.5	Compliance expected.
TCBF-1504	All fasteners in Trident-CBF, with the exception of those in COTS units, shall be metric.	Section 8.3.1	Compliant.
TCBF-1505	Any painted coatings in Trident-CBF shall last without degradation for at least 20 years.	Section 8.3.2	Compliant.
TCBF-1506	Any unpainted surfaces in Trident-CBF shall be treated against corrosion and to provide necessary function for at least 20 years.	Section 8.3.3	Compliant.
TCBF-1507	All LRUs in Trident-CBF shall be equipped with nameplates which are visible after installation and contain the following information: • Name and model number. • Revision. • Serial number.	Section 8.3.5	Partially compliant. "Drawing number including revision" seems like an impossible requirement to meet since many such drawings could apply.

	• Manufacturing month and year. • Name of manufacturer.		Instead suggest “Name and model number”.

5 System Architecture

5.1 Simplified Overview

A simplified block diagram of *each* of the 3 tridents in Trident-CBF is shown in Figure 5-1. In this diagram:

- The "TALON LRU" is used for all signal processing functions in the VCC-part and the FSP-part. It contains 2 large FPGAs and substantial fiber I/O on 2 TALON-DX boards. Each LRU is identical: it is loaded with firmware and software needed to perform VCC or FSP functions, depending on where it is installed. See Section 6.1 for further information on the TALON LRU.
- A "VCC-UNIT" consists of 6 TALON LRUs and provides 10 GHz/pol (5 GHz/pol @ 8+8b/sample) VCC-part processing for 11 antennas (only one TALON-DX board of one of the LRUs is used). A VCC-UNIT is a convenient grouping of antennas for the purposes of VCC-UNIT-internal passive fiber routing via the "VCC-MESH" LRU, wherein each of 50 12-fiber MTP output connectors contains one Frequency Slice for 11 antennas. The VCC-UNIT grouping places no restrictions on sub-arraying or FS selection for any of the 11 antennas.
- Each "FSP-UNIT" consists 13 TALON LRUs, providing the required processing capacity for up to 286 antennas (263 is the requirement[9]). Each FPGA (of 26) in the FSP-UNIT accepts a FS from 11 antennas, providing the required processing for the configured Function Mode according to Figure 3-4. Within the FSP-UNIT the "FSP-MESH" LRU internal passive-fiber interconnect provides all mesh connectivity needed to cross-connect the data for correlation and beamforming. Each FSP-UNIT's FPGA outputs its data via its own 10/40G connector to the archive, consolidated into a single 40G output for the FSP-UNIT via a COTS 40G 27-port switch (not costed, but not a significant cost).
- The "VCC-FSP Passive Fiber Interconnect" block shown is not a separate block from the VCC-UNITs. It is the above-described VCC-MESH passive-fiber interconnects incorporated into each one. Thus, there is only point-to-point cabling from the VCC-UNITs to the FSP-UNITs. Each VCC-UNIT 12-fiber MTP output, consisting of one FS from each of 11 antennas, routes to exactly one FSP-UNIT MTP bulkhead input. In total, for a 263 antenna system, in *each* trident there are 50 (FSPs) x 26 (MTP 12-fiber cables per FSP) = 1300 12-fiber MTP cables.

[9] 286-antenna ingest input, but only 263 antenna correlation and beamforming capability. The number of VCC-UNITs can be reduced to 24 to provide 24x11=264-antenna ingest.

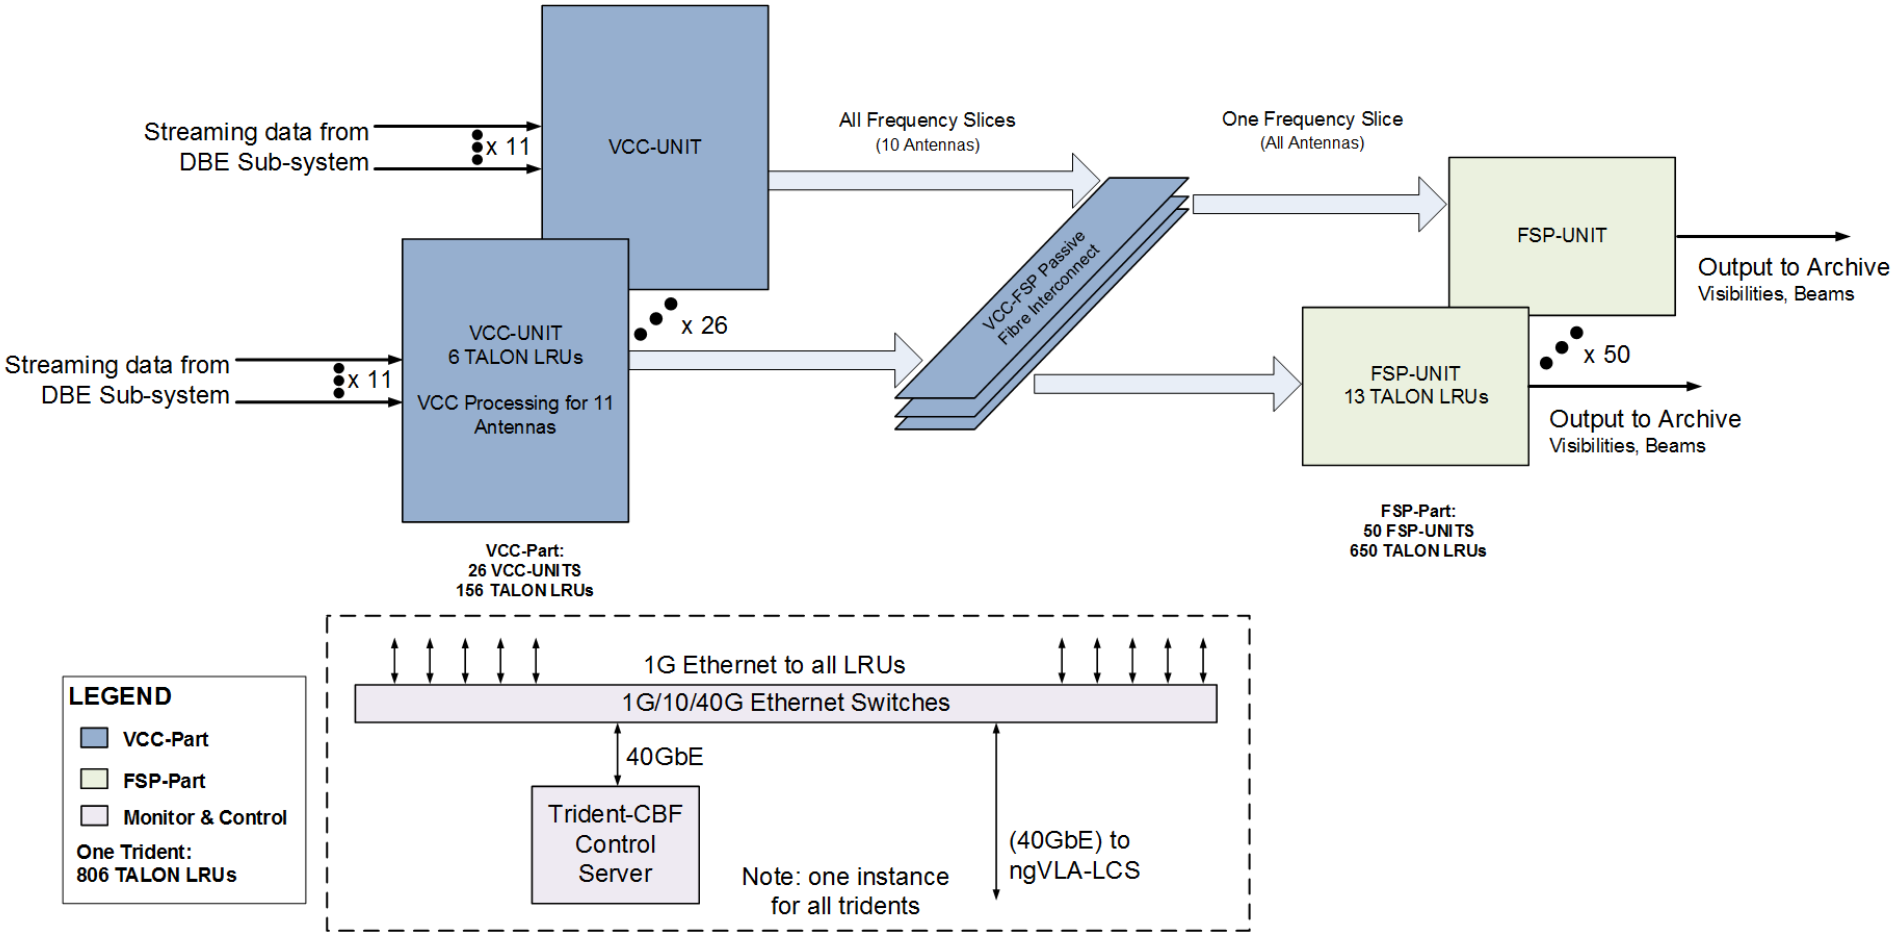

Figure 5-1 Simplified block diagram of each trident in Trident-CBF[10].

[10] 26 VCC-UNITs are shown, but this could be reduced to 24 to reduce cost and provide 24x11=264 antenna capacity, meeting the 263 antenna requirement. Another alternative if more FSP FPGAs are required for correlation/beamforming processing, is to keep each VCC-UNIT at 10 antennas and increase the number of FPGAs in each FSP to 28, providing up to 280 antenna capacity. In this case, there are 28 VCC-UNITs and each FSP-UNIT contains 14 TALON LRUs.

5.2 Signal Processing Architecture

The signal processing architecture is essentially a two-stage process, where the first stage is executed in the VCC-part and the second stage, specific to each FSP Function Mode, is executed in the FSP-part. This architecture allows separation of Band-specific signal processing (in the VCC-part) from Band-agnostic signal processing (in the FSP-part), which finds significant utility in the SKA1 Mid telescope with 6 different Bands, many of them with different sample rates, for which the FSA was developed. The utility of this architecture for ngVLA is allowing for ingest of either 5 GHz/polarization 8+8b samples or 10 GHz/polarization 4+4b samples, as well as considerable flexibility in how Trident-CBF resources are allocated to the signal processing tasks required for any given set of simultaneous observations, with the ability to optimize design for specific processing tasks with no disruption to other tasks (i.e. each FSP Function Mode design can be unique although many share the same signal processing block designs).

This section provides an overview description, rationale, and survey of the key signal processing blocks within the simplified signal processing flow of Figure 3-4.

5.2.1 VCC-part Processing

The first stage in the VCC consists of bulk coarse delay followed by an oversampling coarse poly-phase filterbank to produce the 200 MHz FSs.

The bulk coarse delay correction here is not fundamental, but rather simplifies delay tracking by adjusting/compensating for the bulk of the delay correction[11] every so often (10+ seconds) using high-capacity external DDR4 RAM, leaving only the small (O(10^3-to-10^4), TBC) number of residual coarse samples at the FS sample rate and the very fine delay, to be corrected in the FSP in the FSP FPGAs. The restriction is that for any particular antenna, all delay tracking centers on the sky must be somewhat constrained, normally within the HPBW of the antenna. This means that one would not be able to form a boresight beam and a beam with a large offset from it such as through a sidelobe, since the residual coarse delay in the FSP FPGA would exceed capacity (depending on distance to the array phase center).

The oversampling coarse poly-phase filterbank (VCC-OSPPFB) and FS bandwidth is chosen such that each FS is at a bandwidth and sample rate that is optimal for downstream direct-sample-rate FPGA processing, eliminating the complexities of a corner-turner before the next stage. For Trident-CBF, a 200 MHz bandwidth, oversampled to ~220 MHz is chosen, requiring an FPGA clock rate of ~450 MHz, something that is a reasonable performance target for the TALON-DX technology node. As FPGA performance increases, the FS bandwidth could be increased accordingly, although this is not thought to be cost-effective within the assumed ~2025 construction timeframe of Trident-CBF.

An example VCC-OSPPFB channel response, extracted from [RD3] is shown in Figure 5-2. In this case (for SKA1 Mid.CBF), the input sample stream is real, sampled at ~6 Gs/s, and the VCC-OSPPFB is a 30-channel 10/9 oversampling filterbank. For Trident-CBF, with 5 Gs/s complex (5 GHz digitized bandwidth) in, a different oversampling factor and some adjustment of the FS bandwidth will be needed, TBD.

[11] Encompassing wavefront as well as signal transmission delay from the antennas to Trident-CBF.

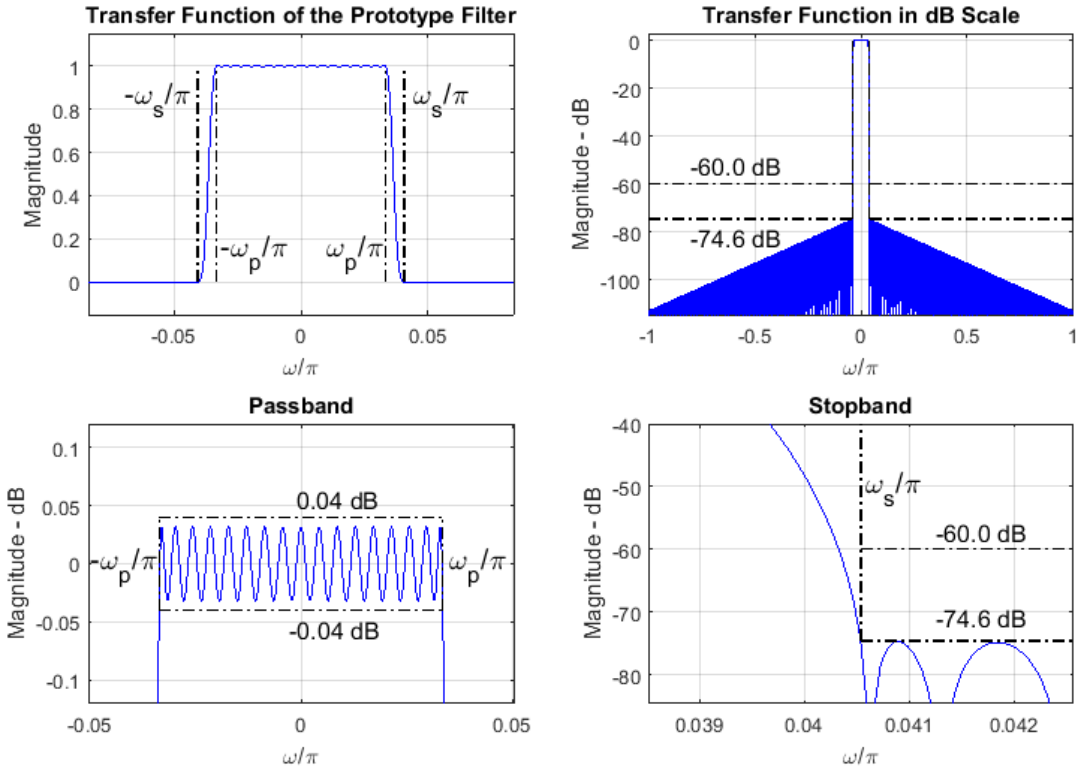

Figure 5-2 Example VCC-OSPPFB output channel (from SKA Mid.CBF, Band 5). (Upper-left: zoomed-in linear magnitude response; lower-left: zoomed-in magnitude response dB; upper-right: magnitude response dB; lower-right zoomed-in magnitude response transition band dB.) Not shown is that the phase response is linear (0) across the passband.

After the VCC-OSPPFB, as indicated in Figure 3-4, there is a circuit switch (internal to the FPGA) that allows any FS to be directed to any of the 50 outputs. This provides the system with considerable flexibility in allocating FSP resources to FSs.

5.2.2 FSP-Part Processing

Processing in each FSP is specific to the particular Function Mode that is currently configured (see Figure 3-4). However, there are some common components, and some insight into the operation and performance of each is provided in the following sub-sections.

5.2.2.1 Re-Sampler (and delay/phase tracker)

The Re-Sampler block performs the following functions:

- Performs residual coarse delay tracking correction (left over after bulk delay in the VCC-part), at the FS sampled rate. This is performed using on-FPGA memory, at the FS sample rate.
- Digitally re-sampling of the data from the raw Frequency Slice sample rate to the sample-rate conducive to providing linearly-spaced channels (requiring a post-re-sampling frequency shift) across multiple adjacent FSs and at channel sample rates that are conducive to distributed correlation and channel-averaging. As well, if the ngVLA employs SCFO sampling or incoherent

clocking, this re-sampler is essential to sample the data to a common rate prior to correlation and/or beamforming.

- Part and parcel with re-sampling is very-fine-delay tracking. Since re-sampling is a delay operation, then very-fine-delay tracking is simply a delay-model-varying offset into the Re-Sampler block.

The combination of re-sampling to a different sample rate and very-fine-delay tracking provides for virtually perfect very-fine-delay tracking with no discernible delay-tracking effects: since re-sampling is constantly and rapidly sweeping through many steps (typically 1024 or 2048) between +/-0.5 samples, the average of the filter re-sampler response is the average of the transfer function of the very-fine-delay tracking response. This is extracted from [RD3] and shown in Figure 5-3:

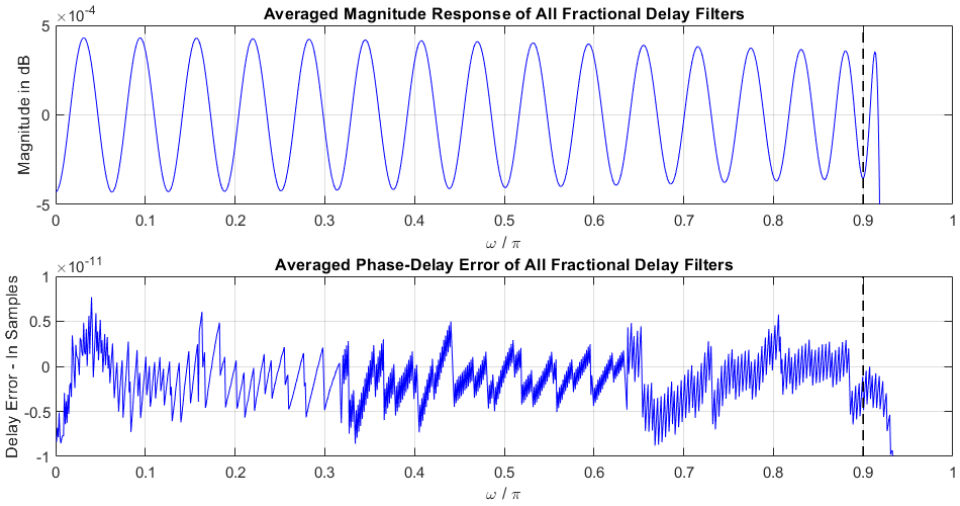

Figure 5-3 Effective average transfer function of the re-sampler/delay-tracker.

Furthermore, since there are a finite number of fine delay steps between +/-0.5 samples in the re-sampler, this imposes a sawtooth delay error waveform that is the magnitude of one step, e.g. 1/1024 samples for 1024 steps. This sawtooth modulates the input spectrum and produces a splatter of copies of the input spectrum. However, this effect is highly mitigated by scrambling the LSBit of the fractional delay so that the sawtooth modulation is turned into random noise. Results presented in [RD3] and Figure 7-2 show that all splatter products are gone and spread into low-level uncorrelated white noise.

The Re-Sampler block also performs a complex phase-rotation/mixing operation, to apply any required phase corrections to the data prior to further channelization, correlation, and beamforming. As above, such a mixing operation isn't perfect, even with 18-bit operations, and so scrambling the LSBit of phase before the required sin/cos LUT (look up table) turns splatter products that could correlate into uncorrelated white noise [RD3].

5.2.2.2 Tunable DDC

A tunable DDC (Digital Down Converter) is used for Zoom Window selection and VLBI beam-channel extraction from the 224 MHz beam. It is a complex decimating FIR filter, with a frequency shift to allow for Zoom Window tuning. A memory-efficient implementation from [RD3] is shown in Figure 5-4, and a typical response (/32 Zoom Window is shown in Figure 5-5).

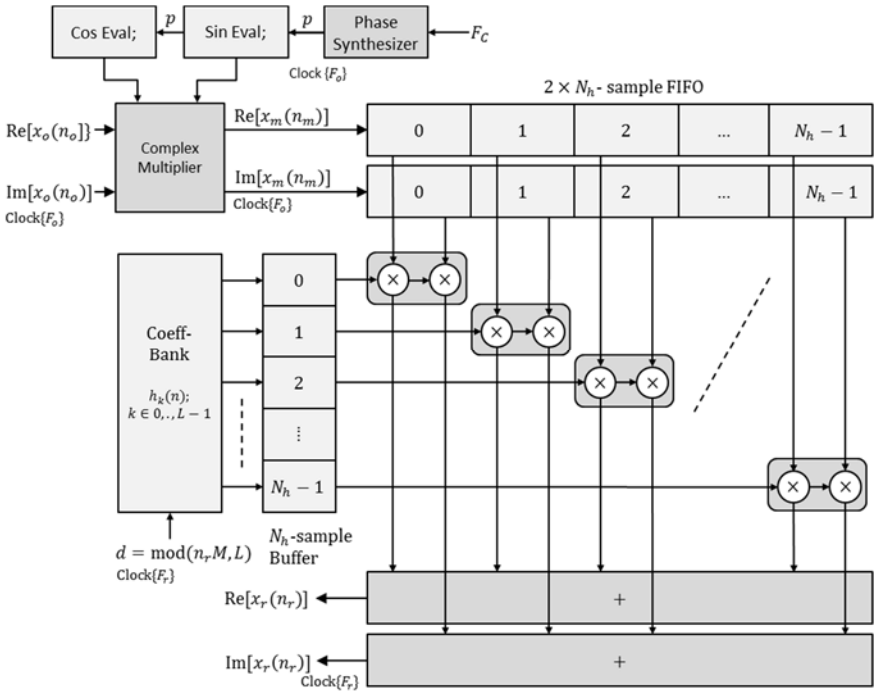

Figure 5-4 Tunable DDC efficient implementation block diagram.

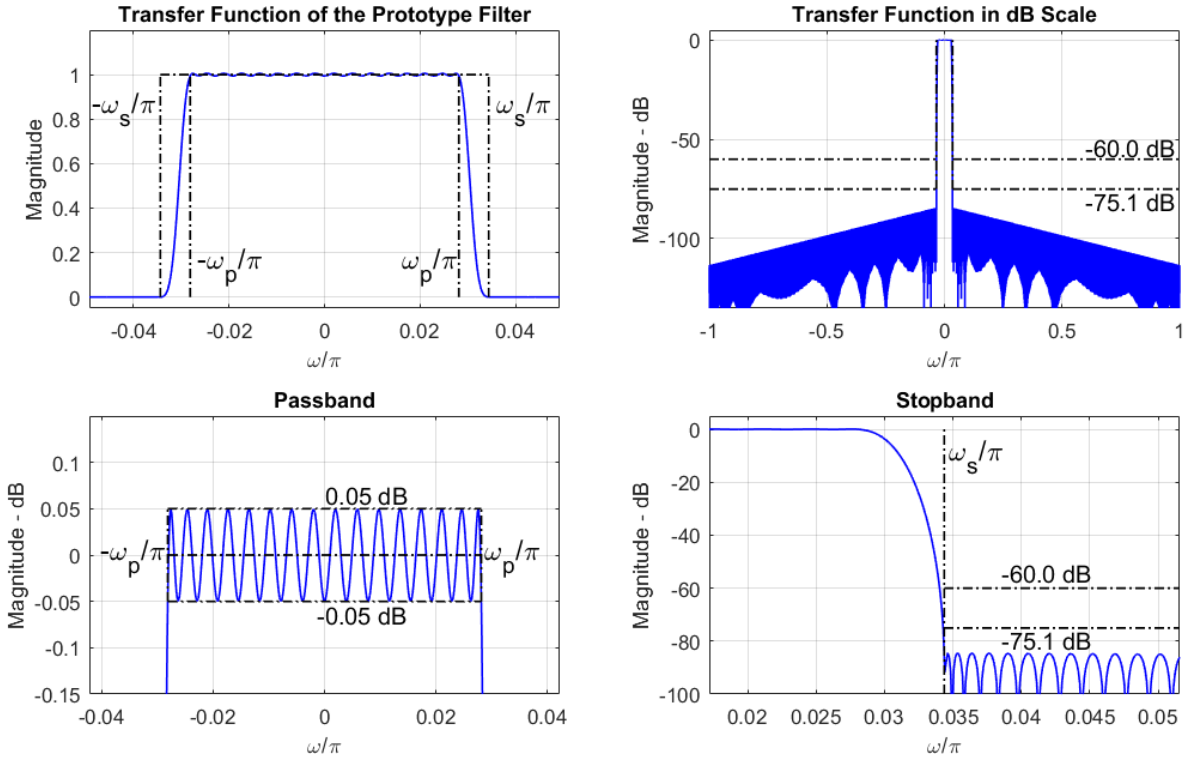

Figure 5-5 Zoom Window tunable DDC output for a FS/32 width.

5.2.2.3 16k Critically-sampled Imaging Channelizer

The imaging channelizer in the Correlation Function Mode, identical for zoom and normal (non-zoom) is crucial for providing required spectral resolution and channel-to-channel isolation. The channelizer filter response, extracted from [RD3], is shown in Figure 5-6.

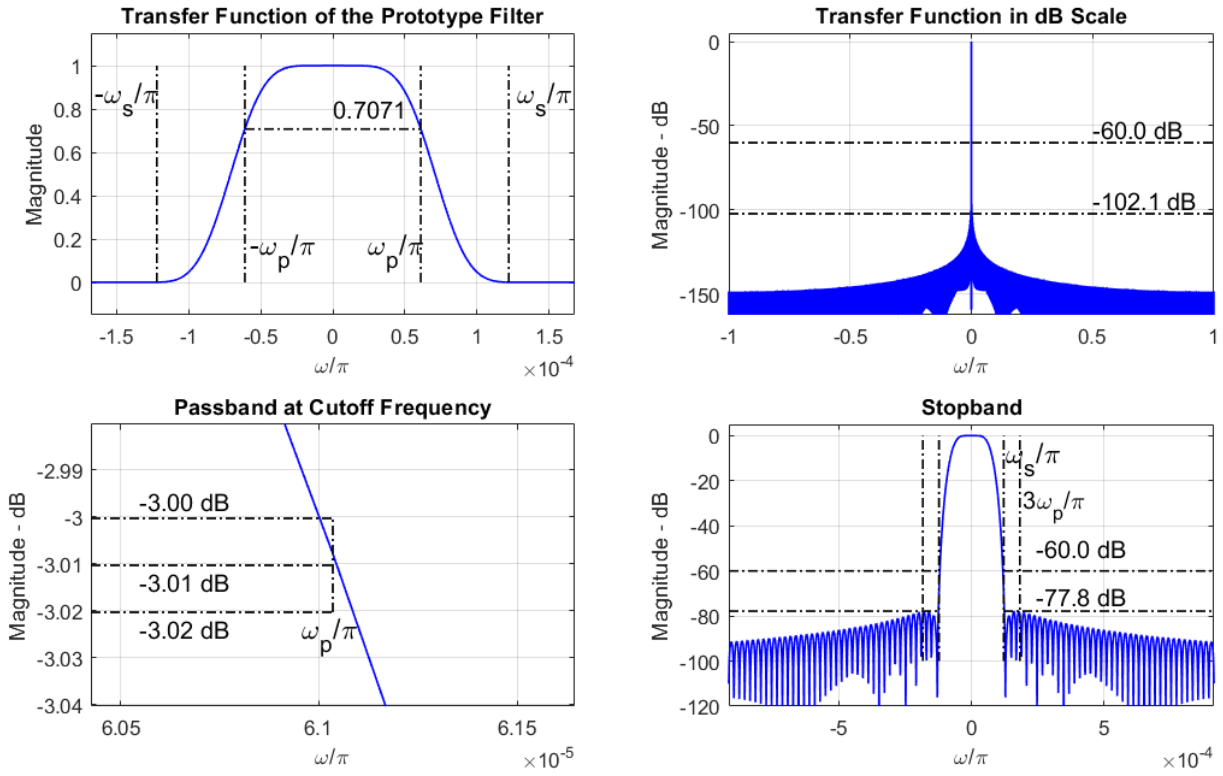

Figure 5-6 16k imaging channelizer response.

The total output channel response, of course, includes the effects of the VCC-OSPPFB, the Re-Sampler, the tunable DDC (for zoom), and the 16k channelizer. Such an end-to-end response, after application of per-channel gain corrections (provided by Trident-CBF to ngVLA-LCS) extracted from [RD3] is shown in Figure 5-7.

5.2.2.4 1k Critically-sampled "simultaneous visibilities" Channelizer

Several Function Modes require simultaneous correlation of ~1k visibilities in the FSP. The channel response of this channelizer, taken from [RD3], is shown in Figure 5-8.

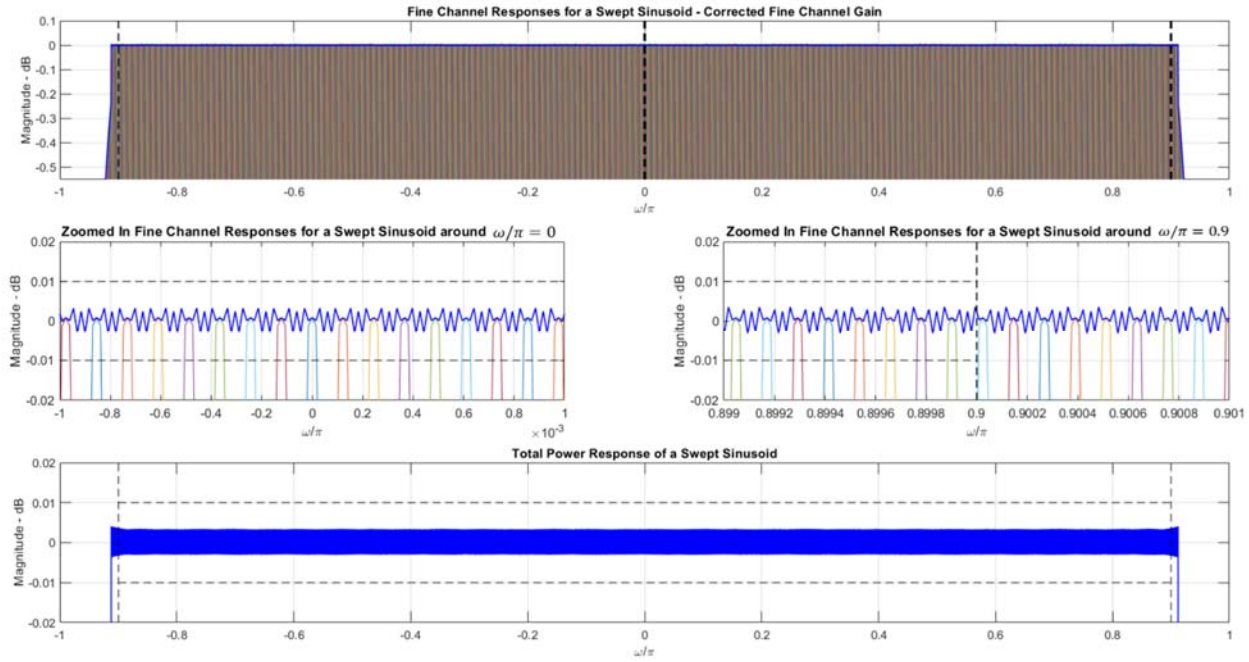

Figure 5-7 End-to-end channel response of the Zoom Window signal chain after application of channel gain correction factors.

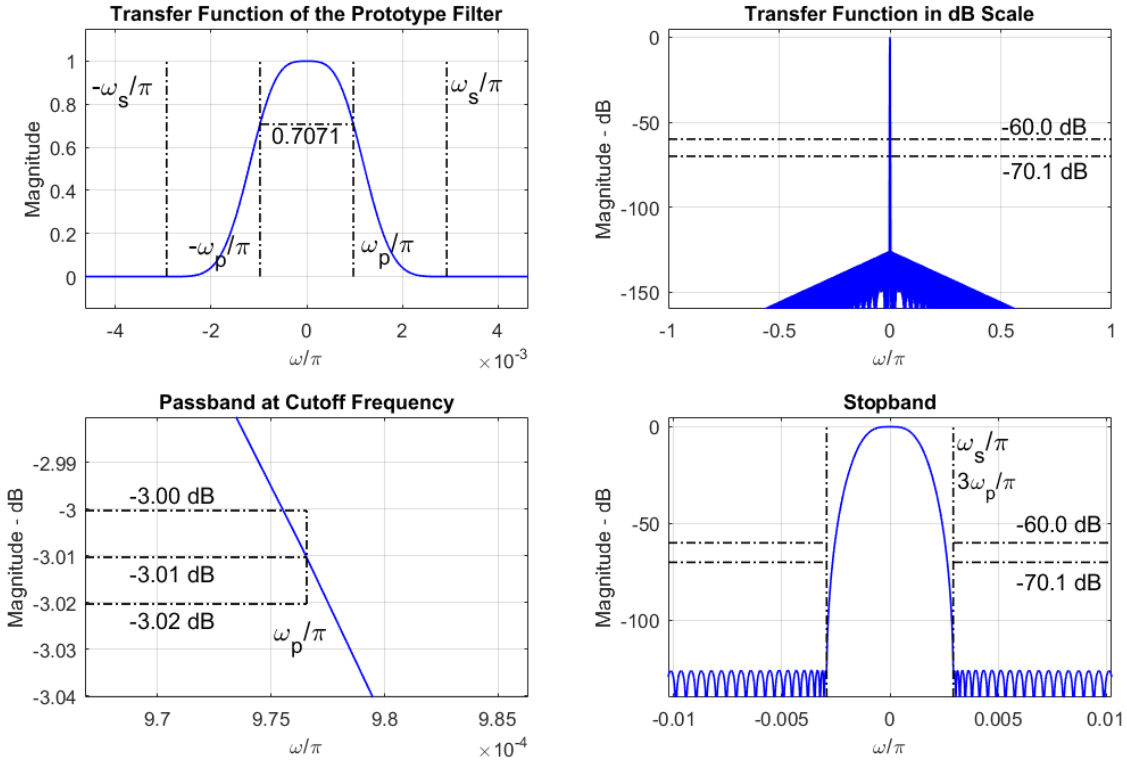

Figure 5-8 Channel response of the 1k channelizer for simultaneous visibilities with beamforming.

5.2.2.5 256-pnt Oversampling Channelizer Response

The response of the 256-pnt oversampling channelizer, performed post-beamforming to allow for fine-frequency Jones Matrix polarization corrections is something of a free parameter since it is not a cost/power driver. Typically, though, the response would be of a similar nature to that shown in Figure 5-2 or better.

5.2.2.6 RFI Detector-Flagger

The RFI detector-flagger, present at each accessible[12] channel resolution in each signal processing flow is extracted from [RD3] and shown in Figure 5-9. This block is a "gear-box" that can be configured ngVLA-LCS to detect and flag RFI with the desired effects.

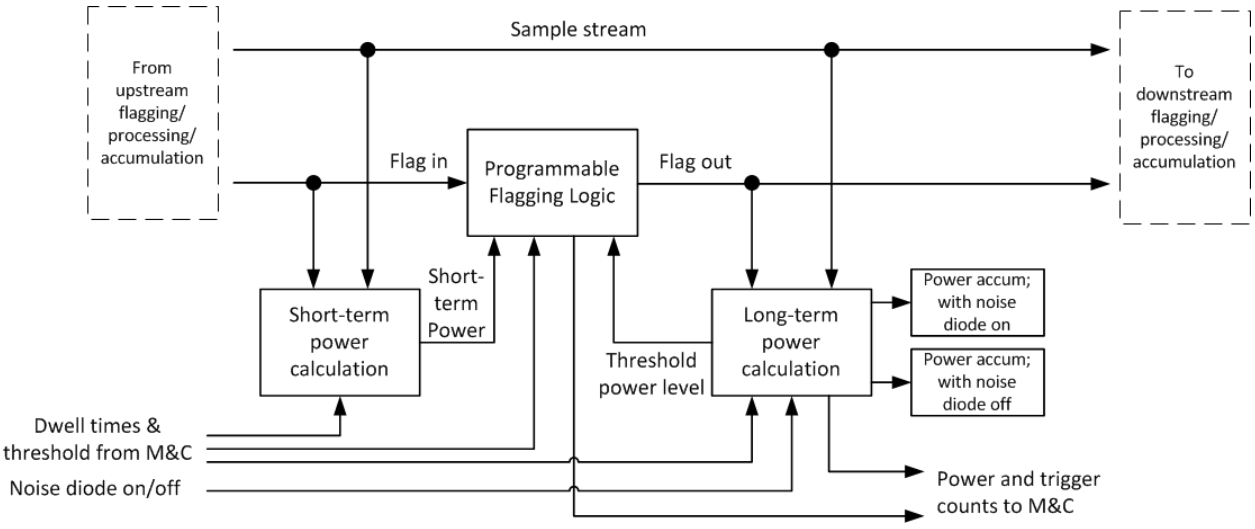

Figure 5-9 RFI detector-flagger.

Furthermore, built into the correlator gear-box is a higher time-resolution post-correlation RFI detector-flagger, if desired. This gear-box is shown in Figure 5-10.

[12] i.e. at the output of a filterbank and not, for example, internal to an FFT.

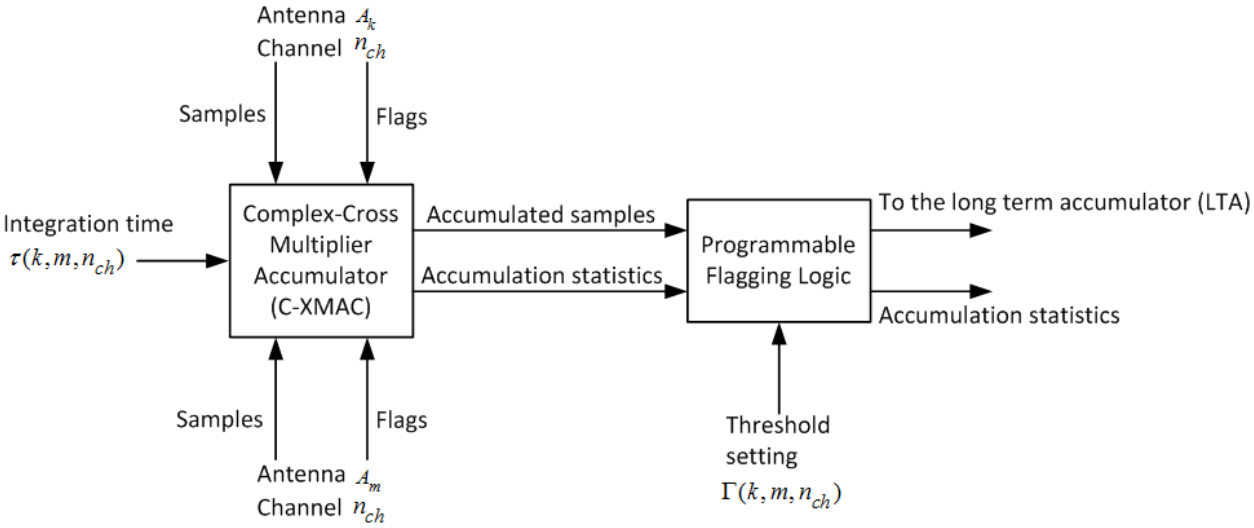

Figure 5-10 Post-correlation, high time-resolution RFI detector-flagger.

5.2.2.7 Other Signal Processing Blocks

The other signal processing blocks—correlation, phase-delay beamforming, Jones Matrix polarization corrections—occur in the FSPs and are straightforward operations with details not included here, rather see [RD5].

Correlation and beamforming are largely data management tasks with the key approaches described below:

- Correlation:
  - Full cross-connect design: time-multiplexed acquisition of all baselines for a portion of the channels in each FSP FPGA. This means that all timing must be synchronized for all sub-arrays within an FSP. Thus, there is a "kernel" integration time, and all sub-array integration times are an integer multiple of it.
  - Zoom and non-zoom handled simultaneously in different sub-arrays in the same FSP since zoom channel sample rates are integer divisions of non-zoom channel sample rates.
  - Separate designs, but same basic approach, for the 16k imaging correlator and the 1k correlator for visibilities in support beamforming. Both correlators perform 6+6b correlation. In the SKA1 Mid.CBF, 19503 baselines are correlated with 9+9b samples @ 744 channels per FSP FPGA (for imaging). In Trident-CBF, 34716 baselines @ ~570 channels per FSP FPGA must be processed and therefore Nbits/sample is reduced to 6+6b (which packs into an 18-bit multiplier, the key multiplier unit size in the FPGA), which should likely accommodate the increased number of baselines (TBC). Alternatively, more FPGAs could be deployed per FSP as noted for Figure 5-1.
  - Corner-turner integrated into the correlator design, after the mesh cross-connect.
- True-delay beamforming:

- The bulk of the beamforming task is performed with per-antenna, per-beam true-delay steering in the Re-Sampler blocks. After that, beamforming is a summing operation, with final beam sums and other post-beam operations performed after the mesh.

- Phase-delay beamforming:
  - Similar to correlation, after meshing, each FPGA in the FSP performs phase-delay corrections relative to boresight (set by Re-Sampler trued delay) for all antennas and a subset of the 4k channels for all 100 beams.

5.3 Physical Architecture

A simplified overview of the physical architecture of Trident-CBF was shown in Figure 5-1. This section delves into more detail of the key blocks, notably the VCC-UNIT, the FSP-UNIT, further details on the key LRUs, and hardware packaging and installation.

5.3.1 VCC-UNIT

A simplified diagram of the VCC-UNIT, and there are 26 such units in each trident of the system, is shown in Figure 5-11. Each TALON LRU consists of two TALON-DX boards[13], and each board contains a large high-performance Intel Stratix-10 FPGA. Each such FPGA performs the VCC-part processing (bulk coarse delay, VCC-OSPPFB, FS switch) for 2, dual-polarization sampled data streams.

Data is output from the FPGA onto fiber, each fiber containing one select Frequency Slice. Fibers for all FPGAs are passively cross-connected in the "VCC-MESH" optical circuit such that the output fibers are now arranged in MTP connectors ready for point-to-point connection to the 50 FSPs.

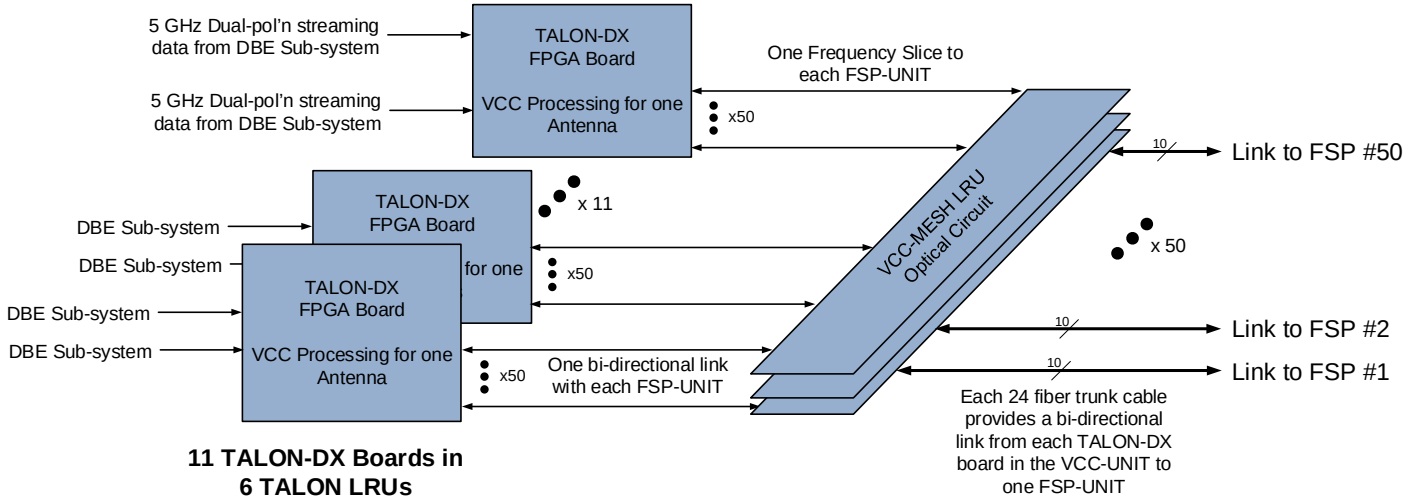

Figure 5-11 Simplified diagram of the VCC-UNIT.

A more detailed look at VCC-UNIT connectivity is shown in Figure 5-12. In this figure:

[13] Since only 11 TALON-DX boards are required, one LRU could contain only one board. However, doing so would break the uniformity of there being one TALON LRU, used everywhere.

- All connections to ngVLA antennas via the DBE are made using 100GBASE-SR4 fiber-optic modules plugged into the TALON-DX board dual QSFP28 cages via the TALON LRU front panel.
- All connections to the VCC-MESH are via the TALON-DX board's FCI LEAP 12 x 26G fiber Mid-Board Optical modules (MBOs), also via the TALON LRU front panel. Each TALON-DX has 5 of these: 4 contain 12 TRx fibers each and 1 contains dual 100G capability (i.e. using only 8 TRx fibers). All-tolled, including the QSFP28 cages, there are 64 x 26G transceiver channels connecting to the FPGA.
- The VCC-MESH LRU is a 1U-packaged box with MTP bulkhead connectors on the front and back panels, and an internal fiber passive cross-connect—specific to the needs of the VCC-UNIT and hidden from access—on the inside. The connectivity of this mesh is shown in Figure 5-13; refer to Section 5.3.4 for details of this COTS passive fiber mesh technology.
- Although the diagram shows bi-directional connectivity through the VCC-MESH, only the outgoing direction contains data, namely Frequency Slices en-route to the FSP-UNITs.

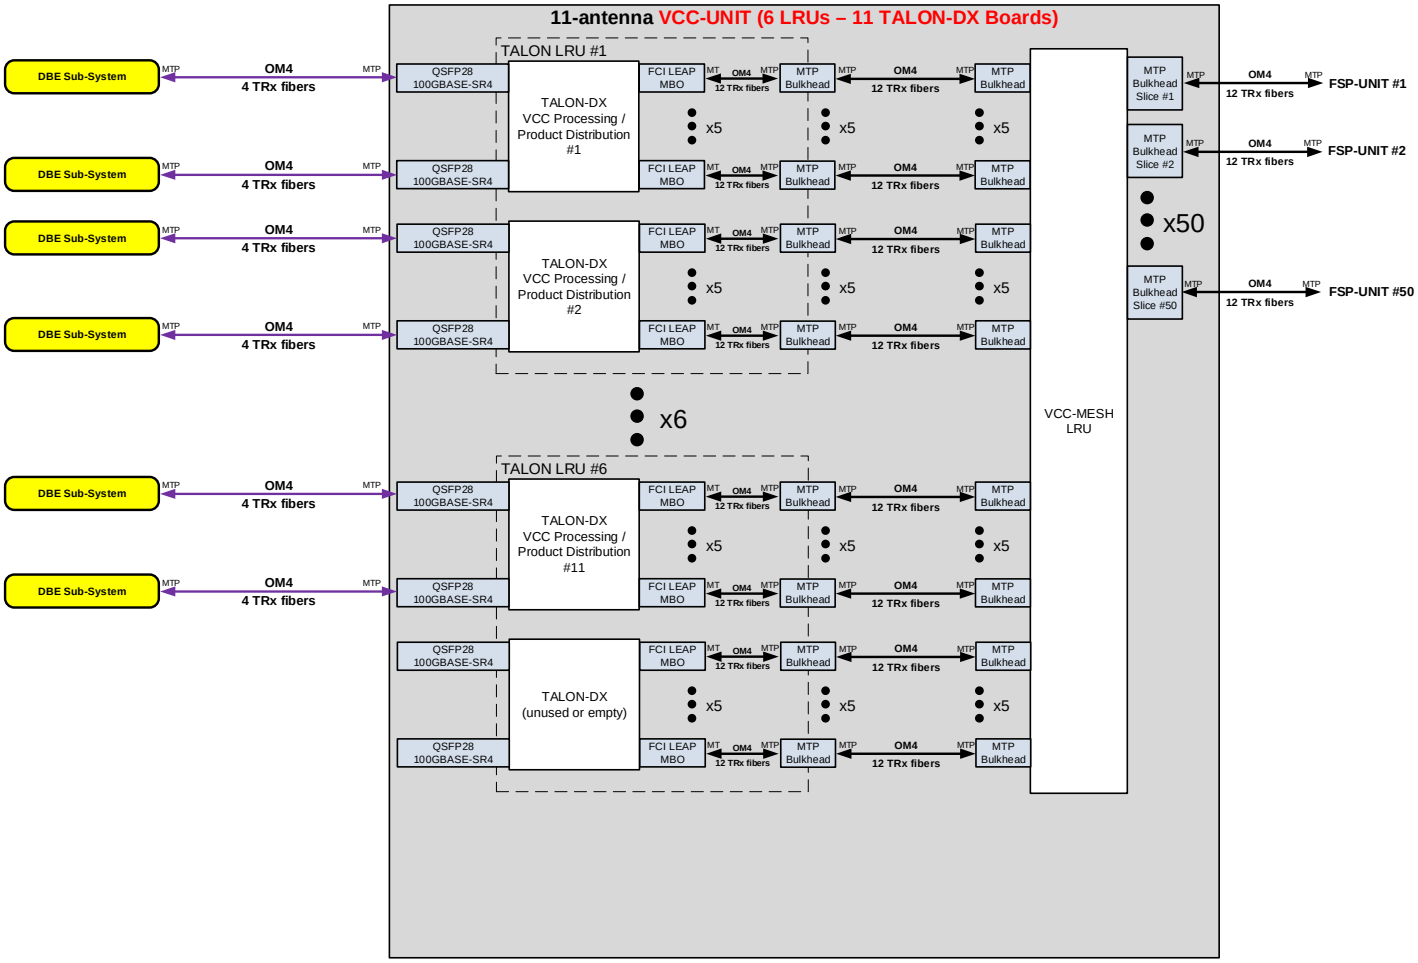

Figure 5-12 VCC-UNIT internal connectivity to TALON-DX boards and the VCC-MESH.

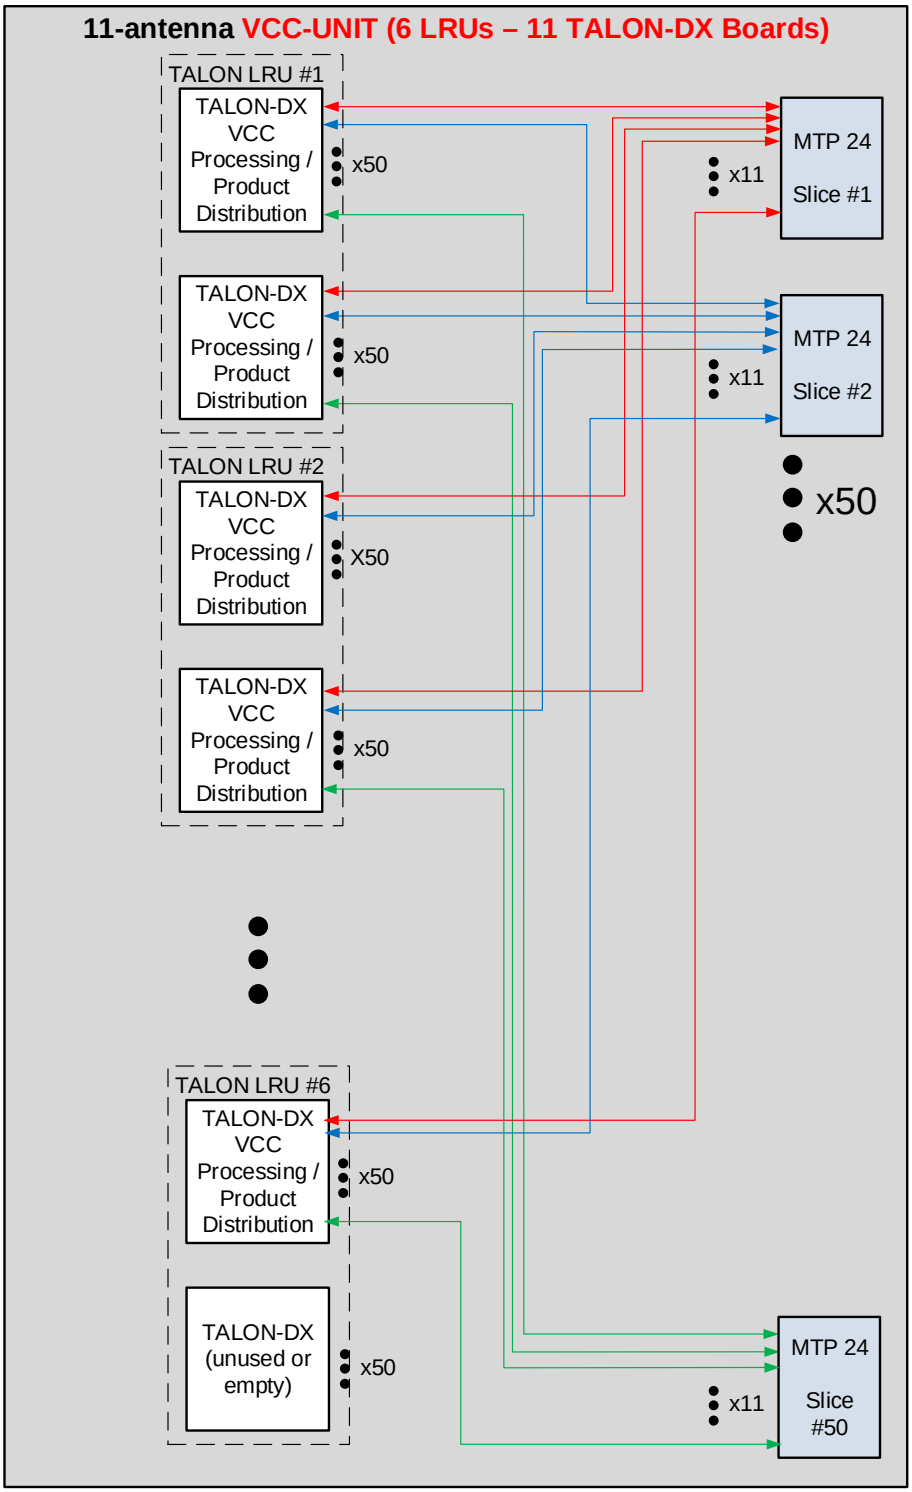

Figure 5-13 VCC-MESH LRU connectivity.

5.3.2 FSP-UNIT

Each FSP-UNIT in each trident is identical. It consists of 13 TALON LRUs (and therefore 26 TALON-DX boards), with direct plug-in inputs from VCC-UNITs (i.e. antenna data in), and a passive fiber FSP-MESH to provide the final mesh connections to organize the data back into the TALON-DX FPGAs for correlation and beamforming. A simplified diagram of the VCC-UNIT is shown in Figure 5-14:

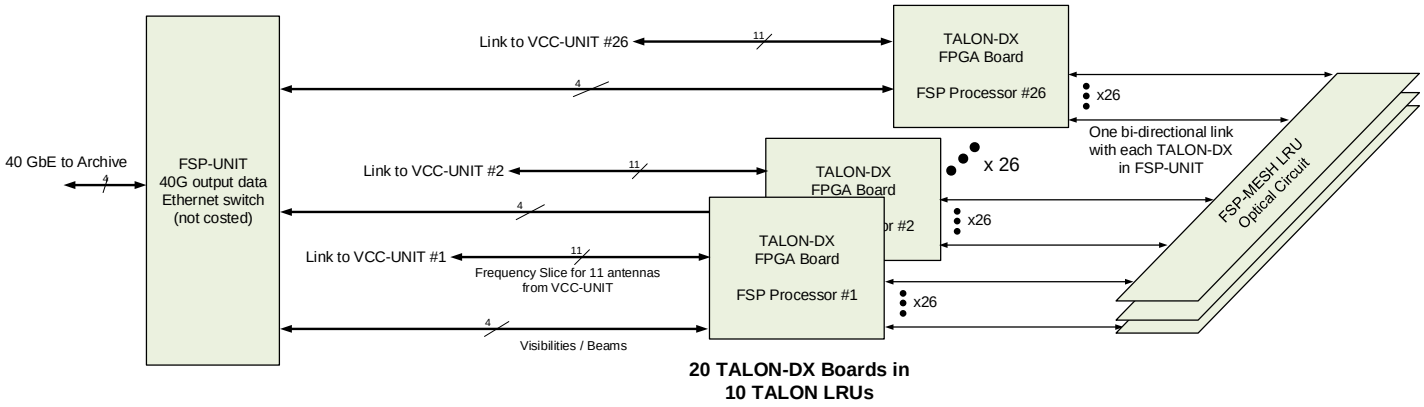

Figure 5-14 FSP-UNIT simplified block diagram.

A more detailed look at FSP-UNIT connectivity is shown in Figure 5-15. In this figure:

- All meshing cross-connectivity is provided by the FSP-MESH LRU, shown in Figure 5-16. As with the VCC-MESH, it is a 1U passive box with MTP bulkheads on the front panel and all fiber cross-connect details hidden internally. Thus, to the installer/maintainer, all connections here (and indeed the entire system) are via MTP fiber cables and connectors.
- Only one of the 40G QSFP28 cages[14] per TALON-DX board is used to output data to a COTS 40G Ethernet switch (not costed), which consolidates data for transmission via as single 40G link to the archive. There is much more bandwidth if needed.
- Although bi-directional connections are shown to VCC-UNITs, Frequency Slice data is only flowing into the FSP-UNIT.

[14] These cages support 100G-BASE-xR4 modules, plug compatible with 40GBASE-xR4.

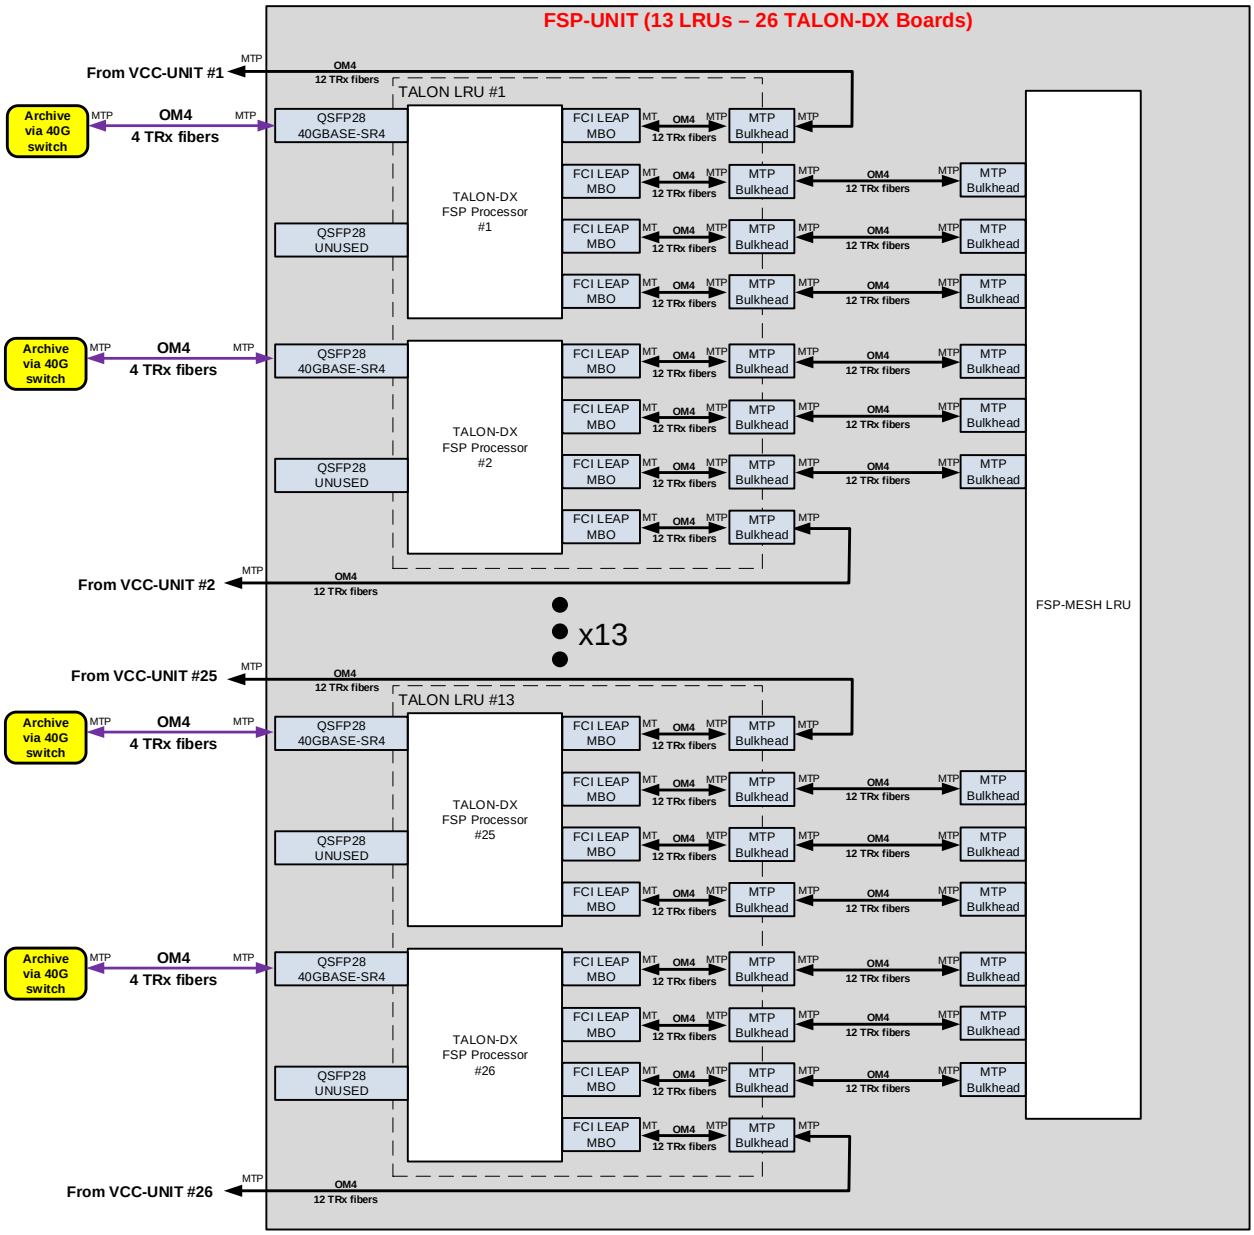

Figure 5-15 FSP-UNIT internal connectivity to TALON-DX boards and the VCC-UNIT(s).

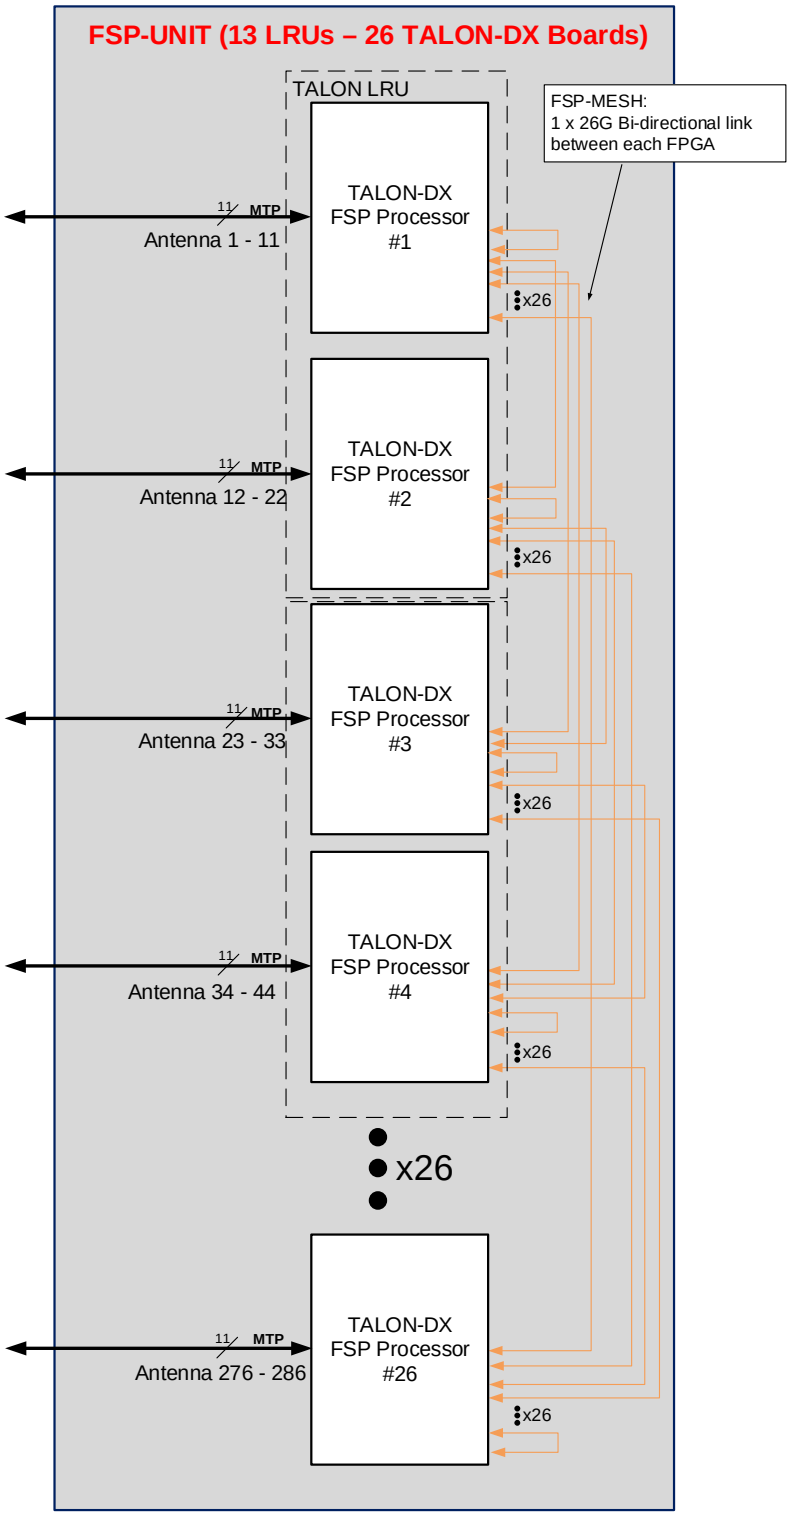

Figure 5-16 FSP-MESH internal connectivity between TALON-DX boards. Although the FSP-UNIT can ingest up to 286 inputs, correlation and beamforming is only provided for the required 263.

5.3.3 TALON LRU

The TALON LRU, containing two independent TALON-DX boards, performs all signal processing functions. There is one LRU, loaded with firmware (FPGA code) and software, which is able to take on VCC or FSP processing tasks depending on where it is located in the system and how it is configured.

This LRU design is based on extensive design and prototyping work done for the SKA1 Mid.CBF system. In that system, air-cooling is a requirement, and so the design of the LRU (enclosure, fans, heat sink) is for air cooling, but the basic design of the board and internal layout was built to be transportable to DCLC (Direct Contact Liquid Cooling)[15].

Each TALON-DX board is independent and has the following key capabilities:

- Intel SX280 Stratix-10 14 nm FinFET FPGA. This device has ~11k 18-bit multipliers and extensive logic and on-chip memory operating typically and comfortably at 500 MHz, and 64 x 26G SERDES transceiver channels—all of which are available as I/O. The device also contains an on-chip HPS (Hard Processing System) based on a quad-core ARM A53 CPU.
- 5 FCI LEAP Mid Board Optical (MBO) modules connecting to FPGA SERDES; 4 contain 12 transmit/receive SERDES at up to 26G each; 1 contains 8 transmit/receiver SERDES connecting to FPGA 2x100G Ethernet cores—they may be used for 100G Ethernet, or individual SERDES channels. The number of fiber I/O places a limit on the number of FPGAs in the FSP, and therefore the number of antennas. At 11 antennas per FPGA, the limit is (3 x 12 fibers + 1 x 8 fibers + 1 x 4 fibers [100G cage]) x 11 = 528 antennas.
- 2 x QSFP28 cages, each supporting COTS plug-in 100G (or 40G) fiber transceivers and connecting to the other 2x100G Ethernet cores in the FPGA.
- 4 x 72-bit, 2400 MT/s DDR4 DIMM modules. Three of these are accessible to the FPGA fabric (two of which may be as large as 256 GB each), one is accessible to the HPS cores.
- 100 Mbps fiber module I/O via 6 SFP cages including Clock Data Recovery (CDR) PLLs for Timecode distribution across the system. Timecode is a 100 MHz encoded clock containing the 1PSS and 32-bit 1 second timestamp.
- Dual 1 Gb/s Ethernet I/O for M&C, via SFP cages.
- High-capacity micro-SD card to hold FPGA bitstreams and HPS software.

The TALON LRU is powered by dual 1+1 redundant COTS 800 W AC-DC hot-pluggable power supply, providing each TALON-DX board with its +12 VDC mains supply. If desired, a plug-compatible power supply is available with 48 VDC mains supply. In the 1U DCLC design, 1+1 redundancy of this power supply is likely not possible due to height constraints.

A rendering of the air-cooled 2U TALON LRU for the SKA1 Mid.CBF is shown in Figure 5-17. The 100G QSFP28 cages are located along the bottom, with the MTP optical breakouts (for MBOs and Timecode) in a consolidated bulkhead on the right side. A DCLC 1U LRU is similar, but without the need for such a large air-flow cross-section, and clearly with some re-arrangement of the MTP optical bulkheads. See

[15] Indeed, the initial concept for the LRU was to use DCLC, but the infrastructure in the South African SKA1 Mid facility cannot support it.

Section 6.1 for further detailed technical and prototyping information on this 2U air-cooled LRU for SKA1 Mid.CBF.

Figure 5-17 Rendering of the 2U air-cooled 19" rack-mount TALON LRU for SKA Mid.CBF. The 1U DCLC LRU would be ½ the height. Note that this LRU has no front-rail mounting ears because of fiber routing space limitations; instead, it is mounted on custom slide rails in a 19" rack, and quick-lock secured to the rack at the back

5.3.4 Passive Fiber Meshes (VCC-MESH, FSP-MESH)

The passive fiber meshes in the VCC-UNITs and FSP-UNITs provide all of the necessary fiber cross-connect routing for the system. These are based on the Molex FlexPlane optical circuit, shown in Figure 5-18, packaged in a COTS 1U LRU and broken-out into MTP bulkhead connectors, as shown in Figure 5-19.

The required optical cross-connects for the VCC-UNIT (Figure 5-13) and FSP-UNIT (Figure 5-15) are custom, but procured—turn-key—through a COTS procurement process. All of the complex cross-connect details are contained and hiding in these mesh LRUs, leaving system installation and maintenance dealing only with point-to-point MTP-MTP connections within and across racks in the system.

An additional "TC-MESH" for Timecode distribution is also required, but not discussed in detail here.

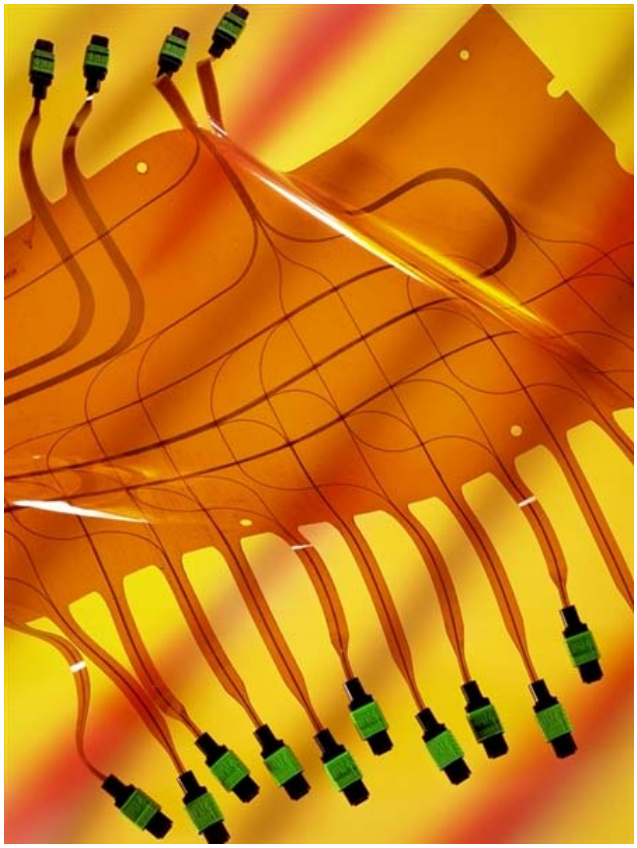

Figure 5-18 Molex FlexPlane optical circuit—a passive optical cross-connect.

Figure 5-19 Leviton 5R1UD-S12 1U optical circuit enclosure.

5.3.5 Timecode Distribution

An identical approach to timing and synchronization across the Trident-CBF system is taken as for SKA1 Mid.CBF (details found in [RD5]), itself derived from a similar approach used in the EVLA WIDAR correlator. Here, a single Timecode fiber signal, derived from the 100 MHz and 1PPS reference input and therefore synchronized to the ngVLA timing reference, is distributed to all TALON-DX boards in the system using a hierarchical approach, each trident separately and identically. Distribution components include the 6 Timecode fiber receivers/repeaters in each TALON-DX board and a TC-MESH optical circuit LRU contained in each rack.

The Timecode signal is a 100 Mbaud signal formatted as follows:

<pre-amble><1PPS mark> **<32-bit 1 sec count><32-bit msg><CRC-8>** <pre-amble>…

The bold 32-bit fields including the CRC-8 are scrambled with a known scramble code whose sequence starts at the 1PPS mark, to ensure DC-balance and that the receiver Clock Data Recovery (CDR) unit in the TALON-DX board doesn't see long strings of 1's or 0's. An example sequence is:

…10101010_0_00110001110101010101…_010101011100101001…_01011001_01010101…,

with the second 0 in the first double-0 of the sequence being the <1PPS mark>. Further work is required to develop the exact Timecode distribution design for Trident-CBF; SKA1 Mid.CBF distributes Timecode to a total of 27 racks—therefore a similar design, for each trident independently[16], is applicable.

5.3.6 Hardware Packaging and Installation

The entire Trident-CBF system is contained in COTS 19" server racks. Each such rack, with a notional population shown in Figure 5-20, contains:

- Two FSP-UNITs requiring 15U each. 13U are required for the TALON DCLC LRUs, 1U for FSP-MESH optical circuit LRU, and 1U for the 40G Ethernet output switch.
- One VCC-UNIT requiring 7U. 6U for the 6 TALON DCLC LRUs providing 11 antenna, single-trident, processing capability, and 1U for the VCC-MESH optical circuit.
- 1U for the TC-MESH optical circuit for Timecode distribution in the rack.
- 2U for a COTS 10G/1G Ethernet switches for monitor and control into the TALON-DX FPGA HPS processors.
- 2U for the COTS CoolIT CHx40, 40 kW[17] heat exchanger (https://www.coolitsystems.com/coolant-distribution-units/). Estimated power dissipation of each TALON LRU, based on extensive prototyping testing for SKA1 Mid.CBF, is ~600 W; with 32 TALON LRUs (with one TALON-DX board missing or powered off), the heat exchanger must handle ~19.3 kW.

[16] But still with cross-trident synchronization to meet requirements for data product timestamping synchronization.
[17] At 30 C supply water temperature.

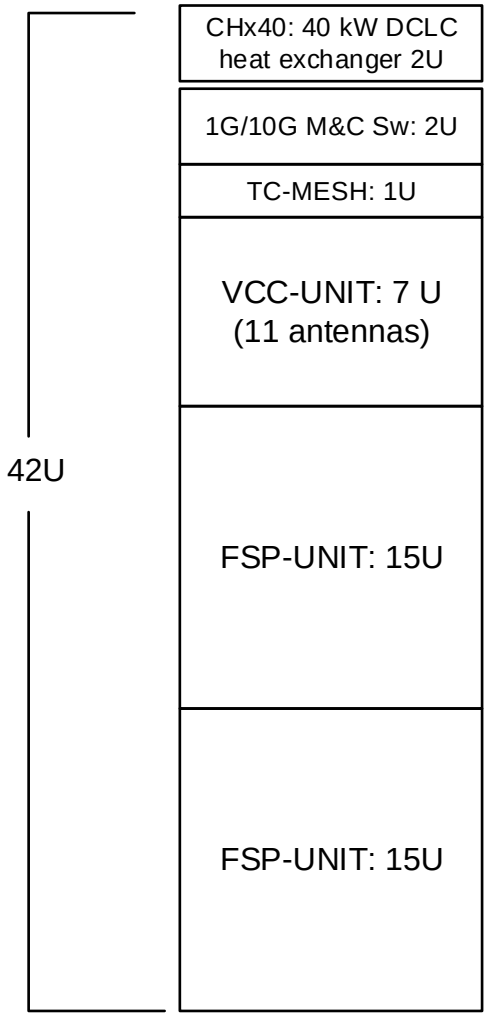

Figure 5-20 Trident-CBF notional 42U 19" server rack population. If required, a 48U server rack could be used since there is no spare space available.

- At the rear, and AC-AC COTS Power Distribution Unit, as well as a DCLC manifold to provide connectivity between the CHx40 unit and the DCLC heat exchangers on the TALON-DX boards.

For *each* trident, with 50 FSPs and 263 antennas, 26 such racks are required, with 24 such racks containing only the VCC-UNIT and potentially a redundant FSP, if desired.

The total number of racks for all 3 tridents is therefore 78, plus 1 or 2 air-cooled racks for central monitor and control and networking. Thus, space for 80 racks, 78 at ~19.3 kW and 2 at ≤ ~10 kW is required for a Trident-CBF based on this generation of technology. See Section 13 for forecasting of power and space to ~2025 technology.

See Section 6.4 for further details on rack technology.

5.4 Software/Control and Monitor Architecture

Details of the software/M&C architecture for Trident-CBF are TBD, however, such details have been worked out extensively for the SKA1 Mid.CBF design. A summary of key aspects is as follows:

- The TANGO software communications and abstraction infrastructure [RD4] is used throughout, including on the TALON-DX FPGA HPS processors.
- Hard real-time embedded processing, such as delay and phase tracking model updates in the FPGA fabric, occurs in the TALON-DX FPGA HPS processors.
- Firmware IP blocks in the FPGA connect to the HPS processors via a standard bus (Avalon or AXI). All such IP blocks' read/write/control registers are accessed by software running on the HPS processors, by the Trident-CBF Control Server, and indeed any network node with available access, via device drivers that are "TANGO devices".
- The FPGAs' ability for "partial bitstream re-configuration" is used extensively:
    - The FPGA I/O ring and HPS processors are always active. This means that M&C intelligence, high-speed serial connectivity, and Timecode connectivity across the system is always maintained.
    - When a new FSP Function Mode is configured, only the internal partial reconfiguration region of the FPGA is bitstream-configured. This approach means minimal time to change Function Modes, typically a few seconds. Similarly for the VCC FPGA 10 GHz/pol 4+4b or 5 GHz/pl 8+8b configurations.
- The HPS processors run real-time Linux. Boot time from the micro-SD card is a few seconds, and the processors are always active independent of FPGA bitstream.
- The central Trident-CBF Control Server translates high-level configuration commands from the ngVLA-LCS to internal VCC-UNIT and FSP-UNIT M&C, via the 10G/1G Ethernet network provided.

Whilst the above M&C approach is used for SKA1 Mid.CBF, for Trident-CBF any standardized M&C infrastructure required by the ngVLA may be used, simply requiring the same hierarchy of processing.

6 Technologies

6.1 TALON-DX Board and TALON LRU

A simplified block diagram of the TALON-DX board is shown in Figure 6-1, with a layout diagram shown in Figure 6-2. The key capabilities of this board are described in Section 5.3.3.

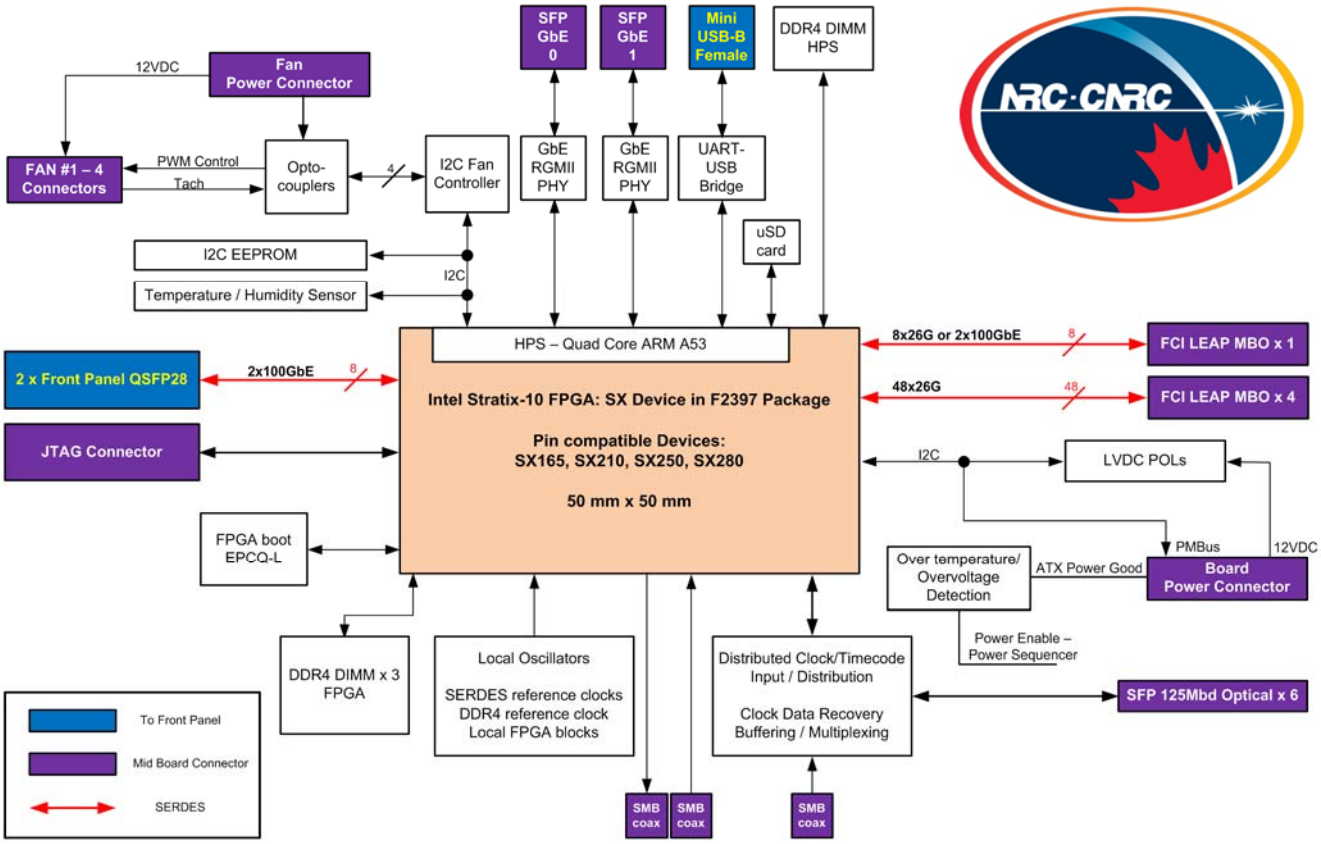

Figure 6-1 TALON-DX board simplified block diagram.

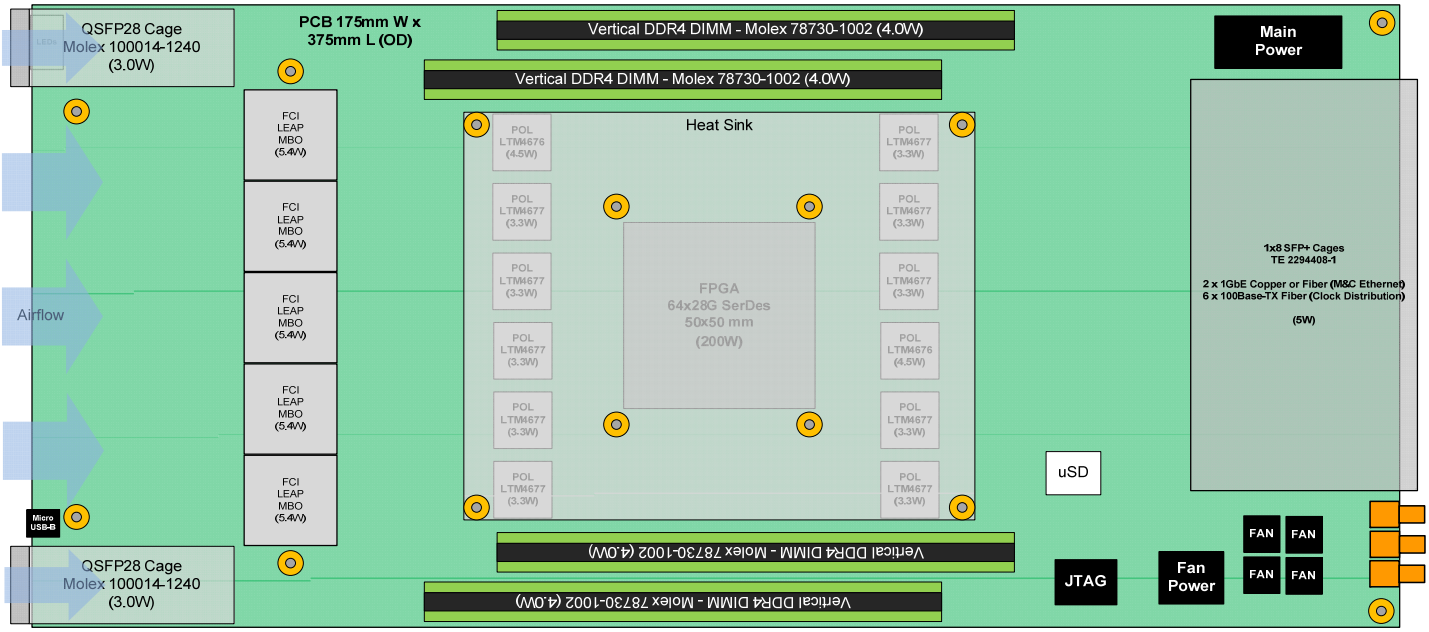

Figure 6-2 TALON-DX board layout diagram. The location of a stacked-fin heatsink for air cooling is indicated; for DCLC, a liquid cooling plate would be used instead, allowing for a 1U packaging height.

As mentioned previously, significant development and prototyping of this board, including HPS processor boot/software development, has been done for SKA1 Mid.CBF and all of the key risks (SERDES 26G performance, DDR4 2400 MT/s performance) have been retired, using engineering prototypes of the FPGA. As of this writing (Sept. 2018) the NRC team is gearing up for procurement of several more prototypes of the 2nd spin of the board to include bug fixes and updates for the board to be used for SKA1 Mid antenna post-digitization signal processing (digital sub-band selection). Also, more extensive performance and margin testing will be performed on these models.

The following figures contain photos of the TALON-DX board and LRU prototypes.

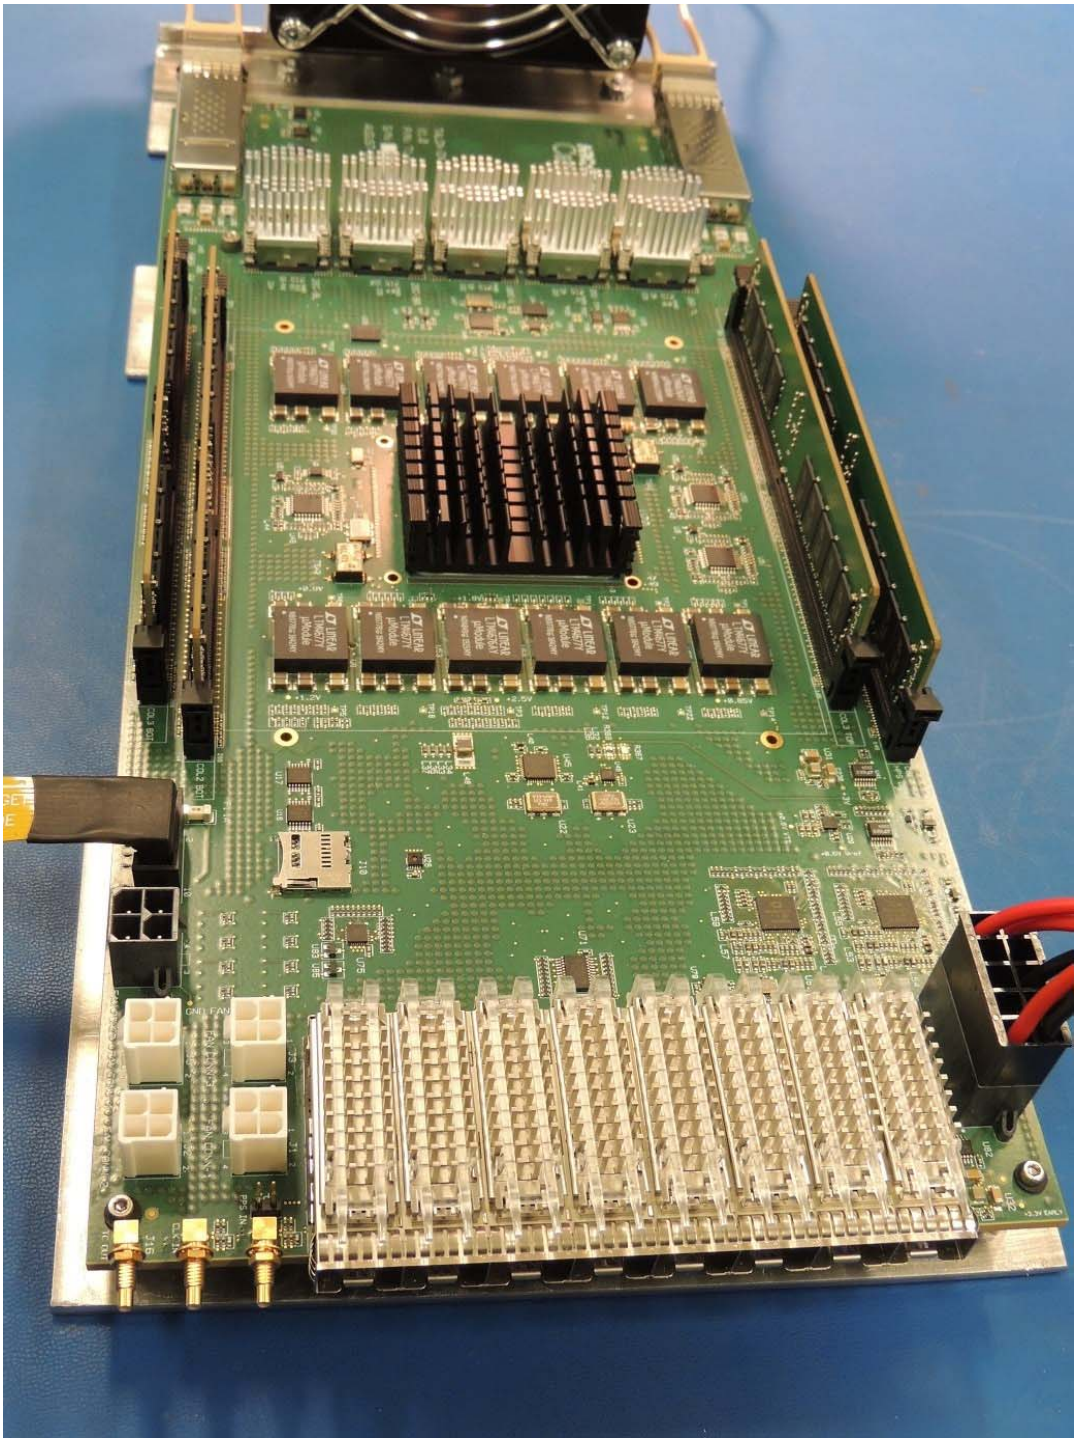

Figure 6-3 TALON-DX prototype board under test.

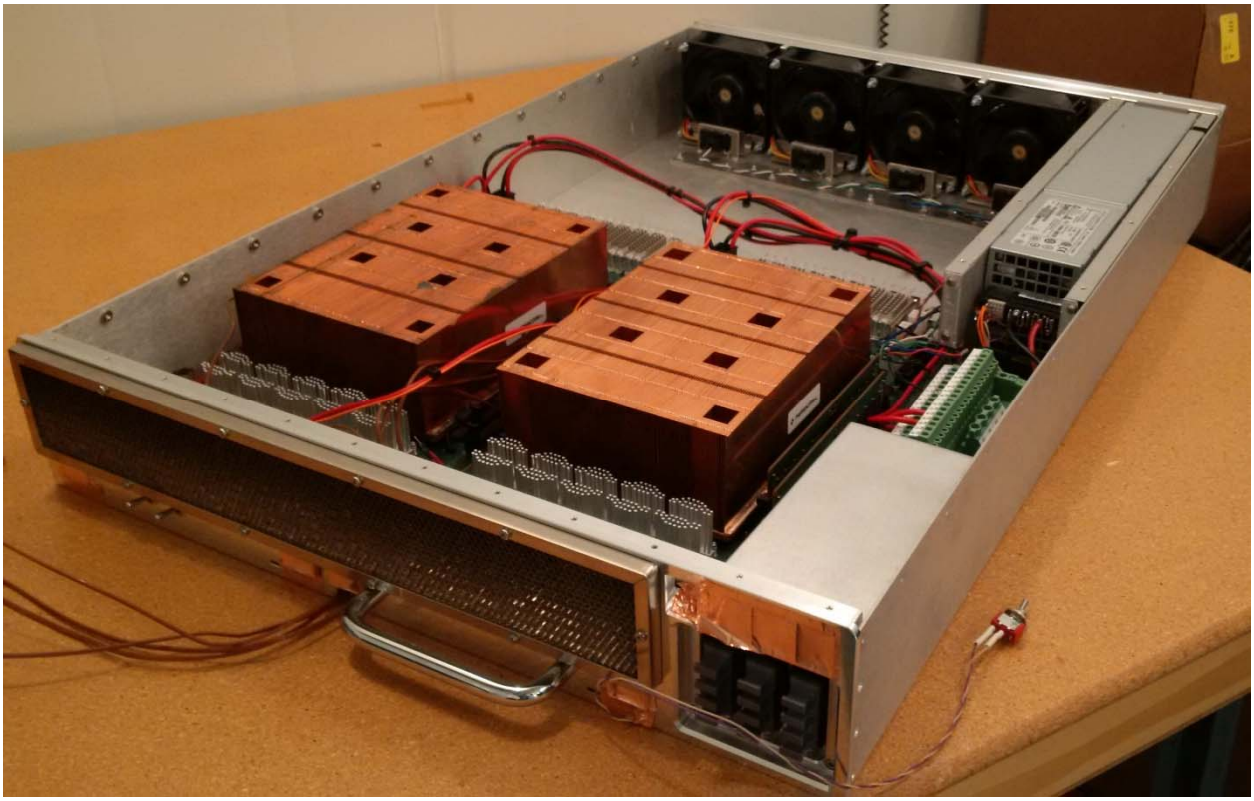

Figure 6-4 2U air-cooled TALON LRU prototype under test. The large orange boxes are copper stacked-fin heatsinks for air cooling. These would be replaced by liquid cooling plates for a liquid-cooled 1U DCLC design.

6.2 Passive Fiber Meshes

See Section 5.3.4.

6.3 Direct Contact Liquid Cooling (DCLC)

There are several flavours of liquid cooling solutions used throughout the industry. In DCLC, a cooling plate—through which cooling water-based fluid circulates—is thermally attached to the main heat generating devices on the board. For the TALON LRU, this means thermal attachment to the FPGA, MBOs, and LVDC power supplies. Thermal attachment to the DDR4 DIMMs is possible, but instead could be a small amount of air cooling—provided by ~2 1U fans at the rear of the LRU, TBC.

A leading player in the design, manufacture, and deployment of DCLC solutions is CoolIT (https://www.coolitsystems.com/). This manufacturer provides custom liquid cooling plate designs and associated COTS components. These include the passive cold plate, rack manifold, and top-of-rack liquid-to-liquid heat exchanger between the rack liquid loop and the facility liquid loop (the latter dumping heat to the outside via a liquid-to-air heat exchanger (radiator). Photos of key components in a DCLC solution are shown in the following figures.

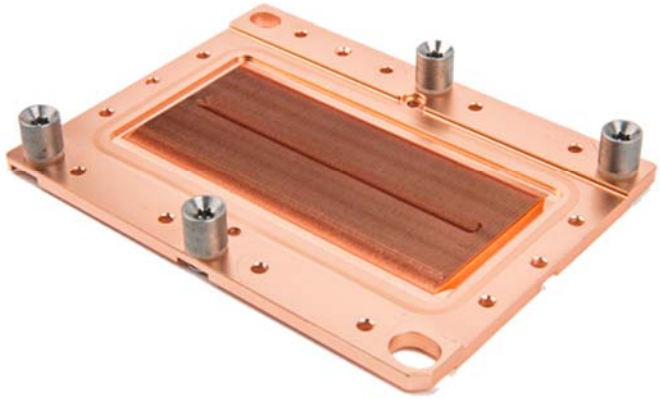

Figure 6-5 CooIT DCLC cold plate example showing the interior micro-fins for heat exchanging. For a DCLC TALON-DX, a customized plate like this would cover the FPGA, MBOs, and LVDC power supplies.

Figure 6-6 View looking down the stainless-steel manifold, located at the rear of the rack. The manifold provides the interface between the cold plate attached to the FPGA (etc.) and the CHx40 heat exchanger. The manifold has a supply and return side, providing 2 hoses to each TALON-DX board (i.e. 4 per LRU).

Figure 6-7 CoolIT CHx40 liquid-to-liquid heat exchanger. Dual 1+1 redundant circulating pumps are located in this unit. If a failure is detected, all of the LRUs in the rack must be powered down, but with dry quick connect/disconnects, this unit can be quickly replaced. It also has an Ethernet network connection to allow its status to be continuously monitored.

The CoolIT solution is very robust and is used extensively throughout the industry, including the CHIME GPU correlator, located at the NRC-Penticton site. All liquid connections, including the CHx40 connections to the rack manifolds and the (overhead) facility liquid lines, are via dry quick-connect/disconnect lines. Nevertheless, each TALON LRU will contain humidity/water sensors to detect leaks (CoolIT also has an integrated solution to detect manifold and other leaks) to take mitigating action if a leak is detected. As well, a welded screened room to house Trident-CBF is highly recommended so that a major leak won't cause catastrophic room failure.

6.4 Other Rack Components

Each of the Trident-CBF racks is a COTS 19" rack with the TALON LRUs, optical circuit LRUs, Ethernet switches, and DCLC components installed. Additionally, each rack contains a COTS AC-AC Power Distribution Unit (PDU), located at the back of the rack. The PDU is Ethernet-controllable to monitor power and control each output; a switch in the Trident-CBF Control Server rack would provide access to each of these PDUs.

7 Performance

7.1 Signal Processing

Extensive performance modelling and testing of the SKA1 Mid.CBF signal processing chains have been undertaken, with results presented in [RD3]. Since the Trident-CBF signal processing chains are based on this design, then these results are directly relevant, with key excerpts presented in the following figures.

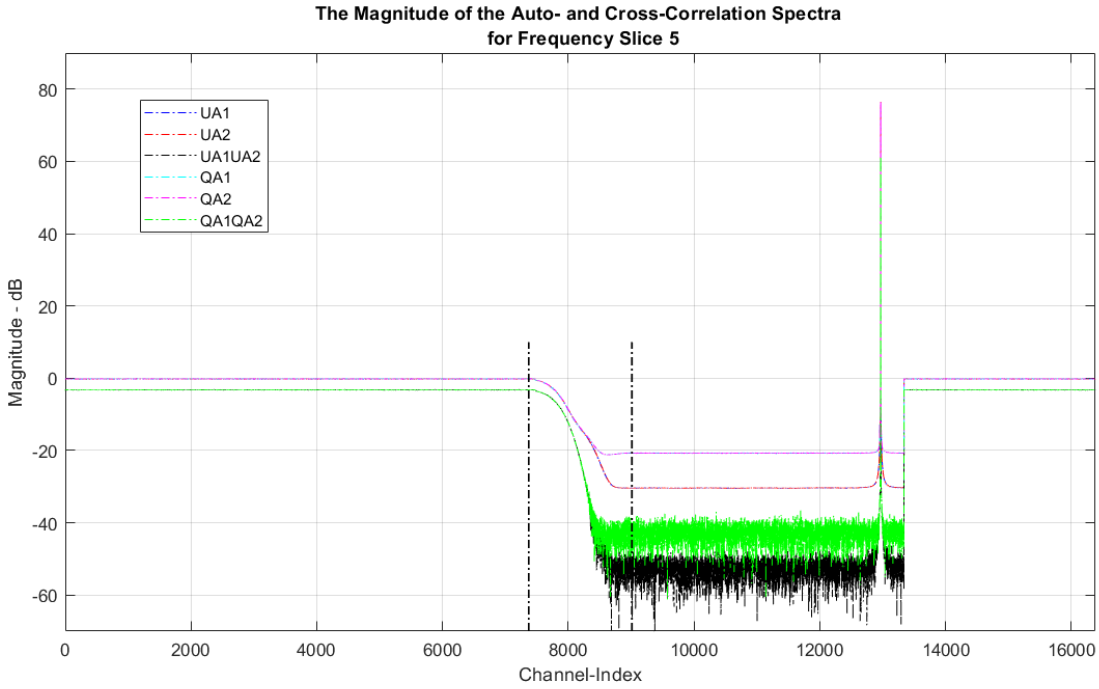

Figure 7-1 End-to-end imaging correlation result with a strong in-FS RFI tone (at 30 dB power relative to the total noise power into an 8-bit digitizer) in the Frequency Slice. Here, the green trace is the quantized (Q) cross-correlation result (i.e. representative of the actual implementation) with ~99% correlation efficiency, per-channel gain/ripple corrections applied, and showing over 100 dB of SFDR (U=floating point/un-quantized model) with some RFI amplitude reduction due to earth-rotation phase wrapping. Note that in this plot the edges of the FS are near the center, and there is a deliberate uncorrelated portion of the FS from channel index ~8500 to ~13500.

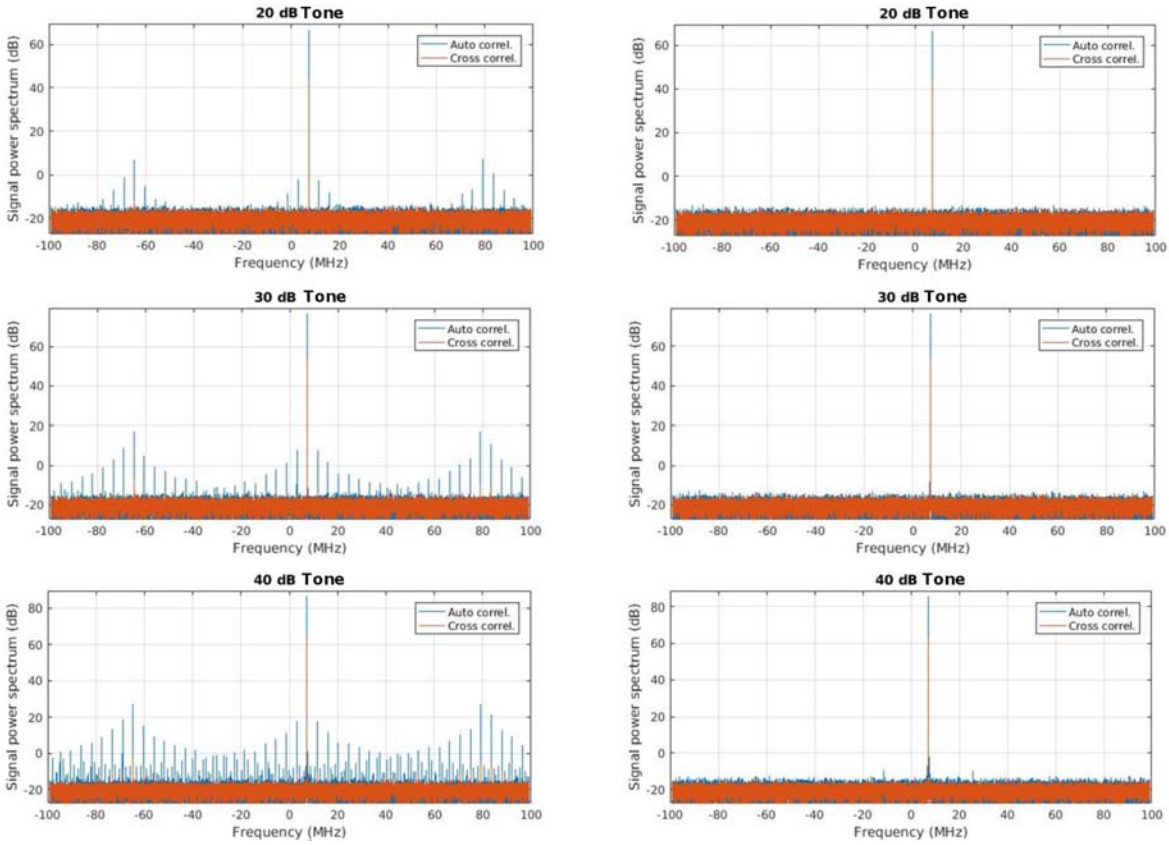

Figure 7-2 Auto-correlation and cross-correlation noise spectra for digitizer input signals containing a single tone of power 20 dB (top), 30 dB (middle) and 40 dB (bottom) without (left) and with (right) scrambling (dithering) in the Re-Sampler phase/delay correction.

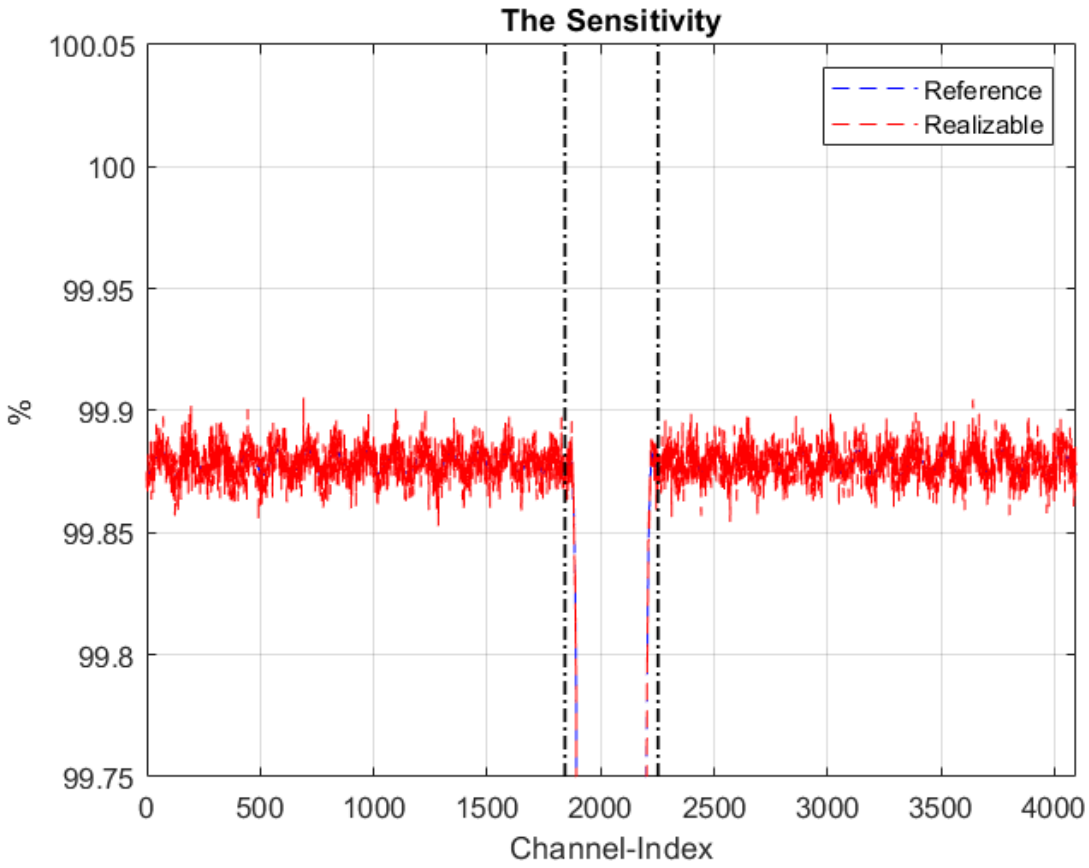

Figure 7-3 Sensitivity across the sampled (channelized) bandwidth for phase-delay beamforming ("realizable"=quantized realized signal processing chain). Here for SKA1 Mid.CBF, the bandwidth is ~330 MHz and the aperture is 20 km, with one antenna on either side of it. For the ngVLA this corresponds to a ~40 km aperture, a 220 MHz FS bandwidth, and less beam offset from boresight due to the narrower primary beam of the ngVLA antenna. The ripple across the band is uncorrected and due to the VCC-OSPPFB response (see Figure 5-2).

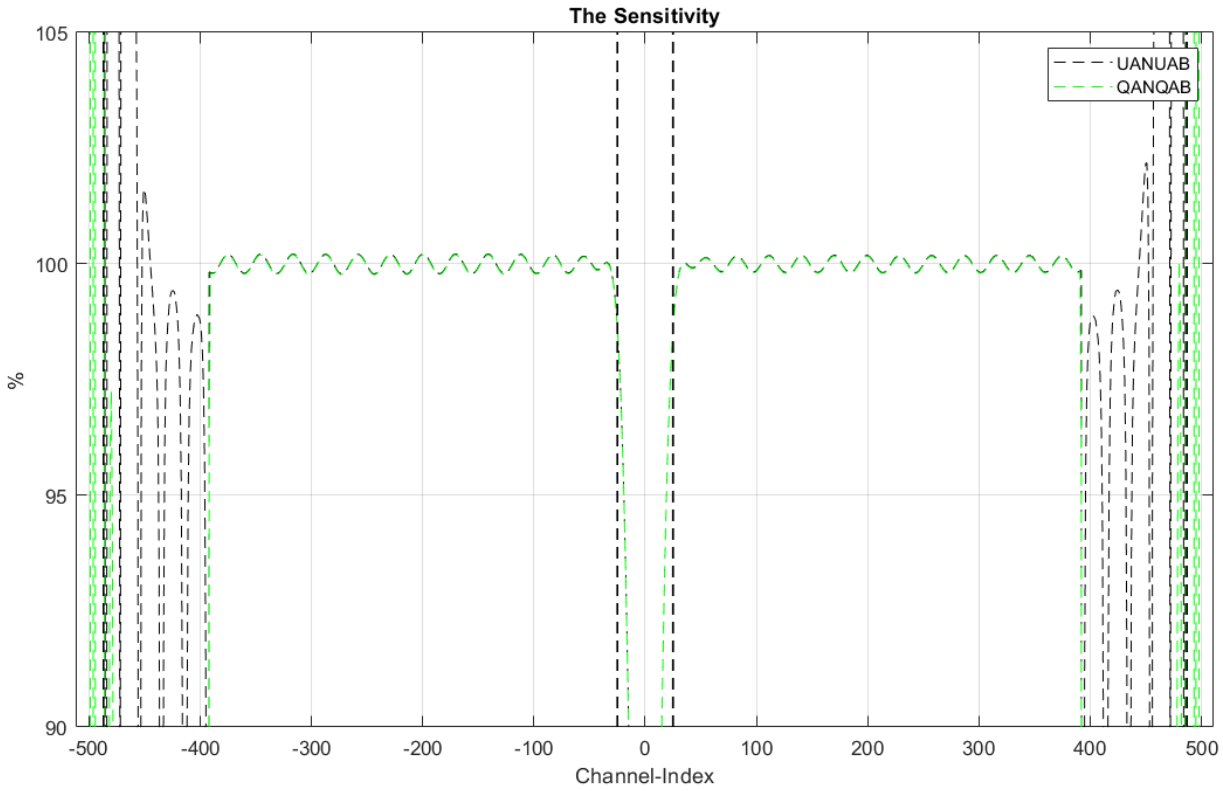

Figure 7-4 Sensitivity of true-delay beamformed real output for 2 antennas across the FS sampled bandwidth of 224 MHz, shifted by ~10 MHz in frequency (hence the asymmetry of the incoherent regions). Here, +/- frequencies of the spectrum of the FS are shown and ripple is due to the VCC-OSPPFB. The apparent >100% sensitivity is due to a modelling anomaly in the sensitivity calculation relative to a single antenna receiving the same signal.

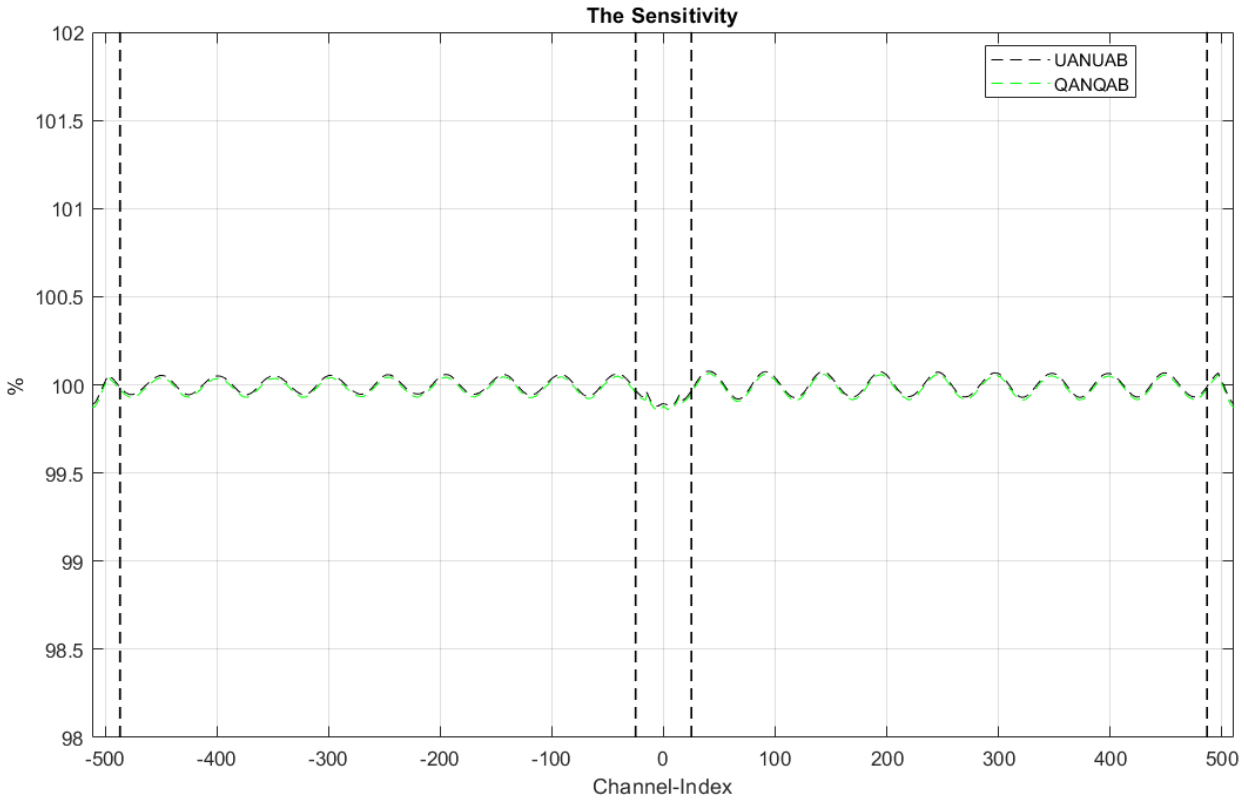

Figure 7-5 Sensitivity of a true-delay 128 MHz VLBI beam-channel output. The signal is real and aliasing occurs at the band edges (noted with dashed vertical lines) due to the beam-channel tunable DDC finite transition bands. Ripple is due to the VCC-OSPPFB response.

7.2 Thermal

Extensive thermal modelling of the air-cooled TALON-LRU for SKA1 Mid.CBF has been performed and it has been found that acceptable thermal operation at the expected power up to Tamb=27 °C is obtained.

However, for Trident-CBF with DCLC, thermal performance is much better since a) only ~20 kW/rack needs to be cooled, easily within the capacity of the CHx40 heat exchanger and b) DCLC heat transfer is much more efficient than air. The CHx40 operates with facility water supply temperatures in the range 2-45 °C with a cooling capacity of 40 kW at 30 °C supply water temperature. Cooling capacity curves for the CHx40 are shown in Figure 7-6:

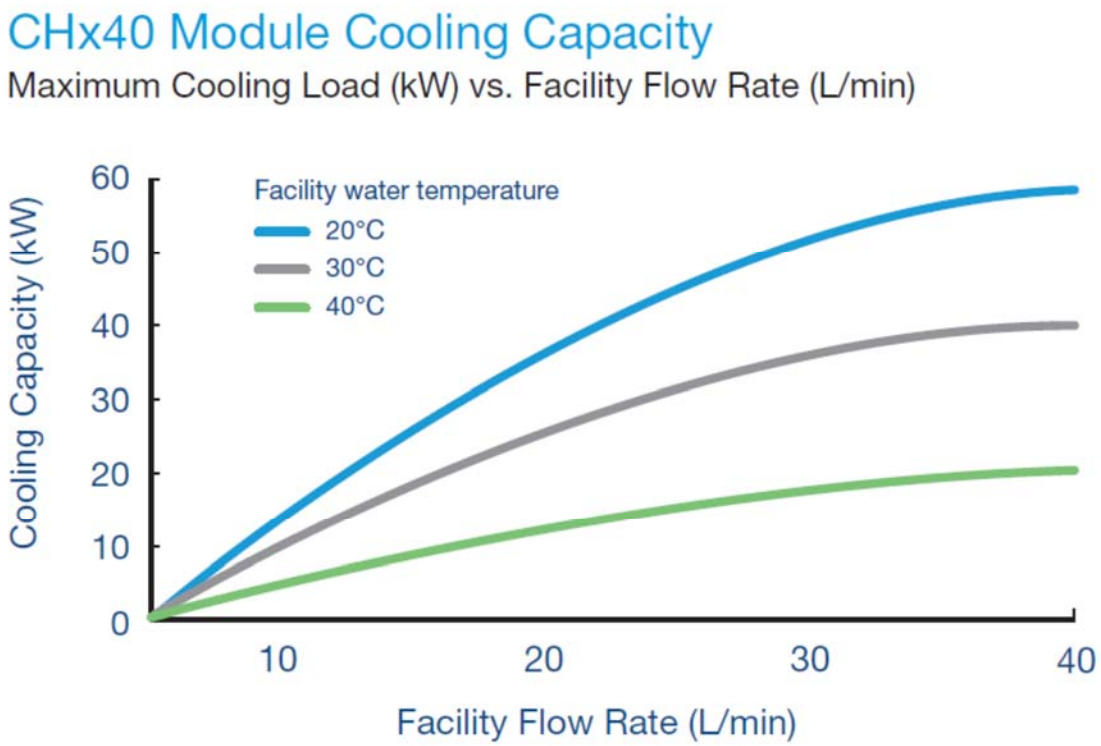

Figure 7-6 Cooling capacity curves for the CoolIT CHx40. Trident-CBF is expected to require 20 kW of cooling capacity per rack which, at a facility water temperature of 40 °C, requires a flow rate of ~40 L/min. The CHx40's power consumption is ~120W.

If the CHx40 turns out to not have sufficient capability, the CHx80, with up to 80 kW of cooling capacity can be used instead. It requires 4U of rack space and its cooling capacity curves are shown in Figure 7-7:

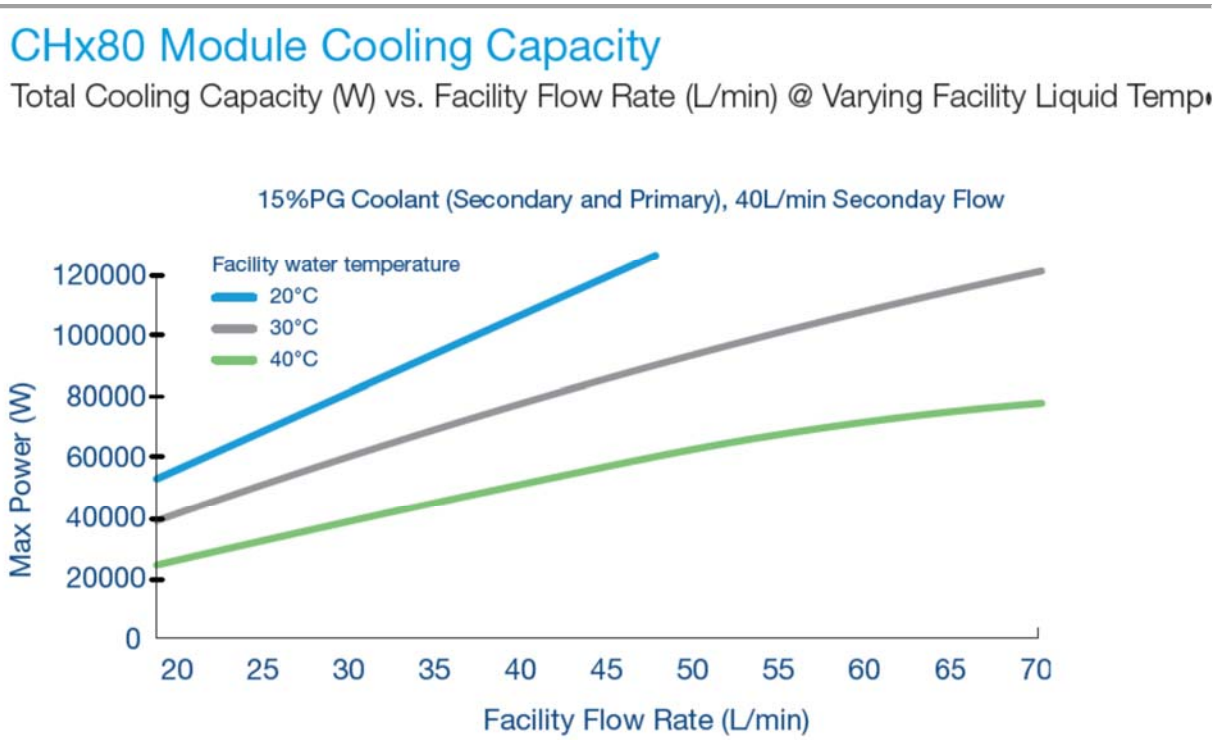

Figure 7-7 CHx80 cooling capacity curves. The CHx80's power consumption is ~625W.

8 Reliability, Availability, Maintainability

8.1 Reliability

From SKA1 Mid.CBF studies reported in [RD5], the MTBF of the TALON-DX assembled board, calculated using the parts-count method according MIL-HDBK-217FN2 is 5 years. In practise, MIL-217 is very conservative and an MTBF of at least 10 years is more likely.

8.2 Availability

Availability requirement TCBF-1401 is met assuming a single TALON LRU fails at any given time. As well, the Intel Stratix-10 FPGAs on the TALON-DX boards contain integral single SEU detection and correction and multiple SEU detection. Given the size of the devices, the number in the system, and the VLA-site 7000 foot location, SEUs due to cosmic rays will occur on a regular basis, sometimes needing FPGA bitstream reboots to correct.

8.3 Maintainability

All of the following items are hot-swap replaceable:

- TALON LRUs. There is one LRU type, loaded with hardware, firmware, and software to take on the required function (VCC or FSP), depending on where it is located. The LRU contains on-board high-capacity non-volatile memory that holds all firmware (FPGA bitstreams) and software—updates to these are pushed-out to the LRUs over the monitor and control network. TALON LRUs have dry quick connect/disconnect DCLC cooling lines at the rear-of-rack manifold (Figure 6-6).
- AC-DC power supply in the TALON LRU. This can be replaced without removing the TALON LRU from the rack.
- VCC-MESH, FSP-MESH, and TC-MESH. These are passive LRUs so shouldn't fail. Removing/replacing the VCC or FSP-MESH boxes necessarily takes down an entire VCC-UNIT, and therefore 11 antennas for one trident, or FSP-UNIT.
- Rack COTS Ethernet switches.

The following items are replaceable, but doing so requires powering-down the rack they are in.

- CoolIT CHx40 top-of-rack heat exchanger. This 2U module has dry quick connect/disconnect fittings. Replacement is very quick, typically ~0.5 hrs or so. This unit has internal 1+1 redundant circulating pumps and network monitoring and so normally it is possible to replace a unit before it fails.
- Rack AC PDU.

As with the EVLA WIDAR correlator, all communications paths between TALON LRUs are continuously tested with embedded error checking, however, off-line tests need to be performed periodically to ensure FPGA-internal calculations are correct. In the Trident-CBF system it is not necessary to take the entire system off-line to do so, FSP-UNITs can be tested individually, as can select portions of VCC-UNITs (doing so takes off-line one or more antennas of a trident).

9 Environmental

The Trident-CBF system is design to be located in a clean, ground-benign, office environment. Since DCLC cooling is employed there should be little to no dust build-up in TALON LRUs, however a dust contamination environment similar to the EVLA WIDAR correlator should be assumed.

With DCLC, there is always the possibility of a liquid cooling leak. The rack local loop (from the TALON LRUs to the manifold, to the CHx40) holds several liters of water and is at relatively low pressure. However, the main liquid loop (CHx40s in all racks to the external heat exchanger), is much higher volume flow and pressure water. In any case, a room construction that would withstand a leak of many 10s of liters of water is prudent, e.g. if in an EMI screened room, welded rather than bolted construction.

Since TALON LRUs, and in fact all equipment in Trident-CBF, are enclosed in individual 19” rack-mount boxes, no special room ESD protection measures are required.

Total estimated power consumption of the Trident-CBF system, implemented using TALON-DX Stratix-10 technology is ~1.5 MW (Section 5.3.6).

10 Safety

All equipment, including the TALON LRU, meet industry standards for electrical safety. The TALON LRU uses a COTS hot-pluggable AC-DC power supply, which is certified to all relevant safety standards.

11 Manufacturing and Procurement Plan

All Trident-CBF PBS items are COTS or COTS-specified, except the TALON LRU. However, for it, a COTS procurement model is applicable as follows:

- The TALON-DX board/assembly manufacturer is tooled-up by the developer to build the board, but also to assemble the entire LRU, load it with firmware and software, and run tests to verify that is operates according to specification. The manufacturer also is tooled and capable of supporting a full RFQ, quote, purchase order, deliver, and defect/repair RMA cycle, complete with standard and extended warranty support.
- For each deployment milestone, required COTS and TALON LRU items are purchased. TALON LRUs are purchased in a COTS-manner from the manufacturer through a normal procurement cycle.

12 Testing and Deployment Plan

This plan will be based on ngVLA project-driven testing and deployment methods and milestones, TBD.

13 Development, Construction, and Operational Cost

The total cost of Trident-CBF is based on the very-well developed SKA1 Mid.CBF cost model, with the following points of note:

- TALON-DX assembled board pricing is based on CY2021production in 2018 dollars and includes actual vendor quotes for all aspects of it, including manufacturing.
- TALON LRU pricing includes production costs of the all components, assembly, and testing. Even though Trident-CBF uses DCLC, the cost of this is judged to be approximately the same as air cooling. Same for rack components.
- TALON LRU pricing is for the volume of SKA1 Mid.CBF (359 LRUs + spares), whereas Trident-CBF contains ~2382 LRUs[18], thus LRU pricing for Trident-CBF is somewhat conservative.
- Although Trident-CBF has some different FPGA firmware blocks, it is largely the same as for SKA1 Mid.CBF. Thus, it is judged that overall the NRE of Trident-CBF will be approximately the same as SKA1 Mid.CBF.
- The complete SKA1 Mid.CBF NRE (including management, system engineering, testing, and deployment) is included in the cost, i.e. it is not assumed that Trident-CBF NRE can be reduced because much of the work has already been done.

Based on these points, the SKA1 Mid.CBF cost model was updated to increase the number of TALON LRUs, optical cables, mesh boxes, Ethernet switches, racks etc. to that required for Trident-CBF. A summary of the results are shown in Figure 13-1 and Figure 13-2. Based on FPGA cost multiplier factors of 0.5 and 0.8 for future 2025 technology, the procurement and production cost of ~$91.8M[19], could reduce to ~$73M-$84M, for a total cost of ~$108M-$121M respectively (2018 dollars).

From Section 5.3.6, total Trident-CBF power, not including cooling infrastructure, is ~1.5 MW. However, this is for 14 nm FinFET 2018 technology. Based on forecasts from Intel, in ~2025 power per operation should be 0.6 to 0.8 of this, and so total system power could be in the 900-1200 kW range.

[18] For the 6 TALON LRU per VCC-UNIT and 13 TALON LRU per FSP-UNIT case. All TALON LRUs in the VCC-UNITs are fully populated with TALON-DX boards. Some cost could be shaved by having an LRU with only one TALON-DX board.

[19] The most significant costs are the FPGAs in the TALON LRU.

WBS Category Definition	WBS Category	Labour PDs									Total PDs	Labour PD%	Labour Cost	Non-Labour Cost	Travel Cost	Contingency	TOTAL COST
		PM:Sr $1,100	PM:Int $900	Eng:Sr $1,000	Eng:Int $900	Eng:Jr $750	Sci:Sr $1,050	Sci:Int $850	Con $950	Admin $500							
Management	MGT	825	0	660	550	0	0	0	209	825	3069	16.0%	$2,673,550	$0	$261,000	$293,455	$3,228,005
System Engineering	SE	0	0	917	766	0	0	0	0	0	1683	8.8%	$1,606,367	$0	$220,400	$182,677	$2,009,443
Product Design	PD	0	0	0	0	0	0	0	0	0	0	0.0%	$0	$0	$0	$0	$0
Hardware Development	HW	0	0	0	606	0	0	0	0	0	606	3.2%	$545,738	$449,893	$69,600	$228,319	$1,293,550
Firmware Development	FW	0	0	0	4937	0	0	0	0	0	4937	25.8%	$4,443,683	$360,000	$208,800	$1,282,535	$6,295,018
Software Development	SW	0	0	0	3986	0	0	0	0	0	3986	20.8%	$3,587,704	$240,000	$139,200	$731,197	$4,698,101
Sub-System Integration Infrastructure	SII	0	0	88	461	0	0	0	0	0	549	2.9%	$502,975	$291,668	$34,400	$118,332	$947,376
Integration & Test (a.k.a. Development Testing)	I&T	0	0	193	1348	0	0	99	0	0	1639	8.6%	$1,489,400	$0	$413,800	$420,456	$2,323,656
Acceptance Test for Verification	AT	0	0	88	807	0	0	0	0	0	895	4.7%	$814,000	$116,006	$116,000	$83,680	$1,129,686
Shipping and Installation	S&I	0	0	0	128	0	0	0	0	293	422	2.2%	$262,167	$160,000	$60,200	$42,589	$524,956
Procurement and Production Hardware Cost	PHW	0	110	0	0	231	0	0	0	110	451	2.4%	$327,250	$91,763,078	$46,400	$9,205,200	$101,341,928
Warranty	WTY	41	6	102	692	12	0	5	10	61	929	4.8%	$828,392	$4,669,032	$83,710	$631,618	$6,212,752
TOTALS		**866**	**116**	**2047**	**14282**	**243**	**0**	**104**	**219**	**1290**	**19166**	**100.0%**	**$17,081,224**	**$98,049,677**	**$1,653,510**	**$13,220,059**	**$130,004,469**

Figure 13-1 Trident-CBF cost summary, based on the SKA1 Mid.CBF cost model, 2018 dollars. This is for 6 TALON LRUs per VCC-UNIT, 24 VCC-UNITs, and 13 TALON LRUs per FSP-UNIT.

WBS Category Definition	WBS Category	Labour PDs									Total PDs	Labour PD%	Labour Cost	Non-Labour Cost	Travel Cost	Contingency	TOTAL COST
		PM:Sr $1,100	PM:Int $900	Eng:Sr $1,000	Eng:Int $900	Eng:Jr $750	Sci:Sr $1,050	Sci:Int $850	Con $950	Admin $500							
Management	MGT	825	0	660	550	0	0	0	209	825	3069	16.0%	$2,673,550	$0	$261,000	$293,455	$3,228,005
System Engineering	SE	0	0	917	766	0	0	0	0	0	1683	8.8%	$1,606,367	$0	$220,400	$182,677	$2,009,443
Product Design	PD	0	0	0	0	0	0	0	0	0	0	0.0%	$0	$0	$0	$0	$0
Hardware Development	HW	0	0	0	606	0	0	0	0	0	606	3.2%	$545,738	$449,893	$69,600	$228,319	$1,293,550
Firmware Development	FW	0	0	0	4937	0	0	0	0	0	4937	25.8%	$4,443,683	$360,000	$208,800	$1,282,535	$6,295,018
Software Development	SW	0	0	0	3986	0	0	0	0	0	3986	20.8%	$3,587,704	$240,000	$139,200	$731,197	$4,698,101
Sub-System Integration Infrastructure	SII	0	0	88	461	0	0	0	0	0	549	2.9%	$502,975	$291,668	$34,400	$118,332	$947,376
Integration & Test (a.k.a. Development Testing)	I&T	0	0	193	1348	0	0	99	0	0	1639	8.6%	$1,489,400	$0	$413,800	$420,456	$2,323,656
Acceptance Test for Verification	AT	0	0	88	807	0	0	0	0	0	895	4.7%	$814,000	$116,006	$116,000	$83,680	$1,129,686
Shipping and Installation	S&I	0	0	0	128	0	0	0	0	293	422	2.2%	$262,167	$160,000	$60,200	$42,589	$524,956
Procurement and Production Hardware Cost	PHW	0	110	0	0	231	0	0	0	110	451	2.4%	$327,250	$96,647,568	$46,400	$9,693,649	$106,714,867
Warranty	WTY	41	6	102	692	12	0	5	10	61	929	4.8%	$828,392	$4,913,257	$83,710	$656,040	$6,481,399
TOTALS		**866**	**116**	**2047**	**14282**	**243**	**0**	**104**	**219**	**1290**	**19166**	**100.0%**	**$17,081,224**	**$103,178,392**	**$1,653,510**	**$13,732,930**	**$135,646,056**

Figure 13-2 Trident-CBF cost summary (2018 dollars) for the 5 TALON LRUs per VCC-UNIT, 28 VCC-UNITs, and 14 TALON LRUs per FSP-UNIT option mentioned as a note to Figure 5-1.