Central Signal Processor Design Description

020.40.00.00.00-0005-DSN

Status: **RELEASED**

PREPARED BY	ORGANIZATION	DATE
O. Yeste Ojeda	Central Development Laboratory	2022-02-08
N. Denman	Central Development Laboratory	
W. Shillue	Central Development Laboratory	

APPROVALS	ORGANIZATION	SIGNATURES
W. Shillue, CSP & Timing IPT Lead	Project Office, ngVLA/NRAO	William Shillue William Shillue (Feb 10, 2022 09:05 EST)
R. Selina, Project Engineer	Project Engineer Dept., ngVLA/NRAO	R. Selina R. Selina (Feb 10, 2022 08:41 MST)
T. Küsel, System Engineer	Systems Dept., ngVLA/NRAO	Thomas Kusel Thomas Kusel (Feb 10, 2022 12:00 EST)
W. Esterhuyse, Antenna Project Manager	Project Office, ngVLA/NRAO	

RELEASED BY	ORGANIZATION	SIGNATURES
W. Esterhuyse, Antenna Project Manager	Project Office, ngVLA/NRAO	

Change Record

Version	Date	Author	Affected Section(s)	Reason
1	2018-07-05	O. Yeste Ojeda	All	Initial Draft. CBF part based on NRC's poster by M. Pleasance, B. Carlson, & M. Rupen (NRC).
2	2018-09-27	O. Yeste Ojeda	All	Update for ngVLA Reference Design Review. Incorporates feedback from Reference Design Workshop. CBF part based on NRC doc. No. TR-DS-000001, Rev. D.
3	2018-11-08	O. Yeste Ojeda	All	Minor adjustments for the internal review.
A	2019-07-31	A. Lear	All	Prepared PDF for signatures & release.
A.01	2021-10-31	O. Yeste Ojeda and N. Denman	All	Initial Draft.
A.02	2022-02-08	O. Yeste Ojeda	All	Incorporate feedback from internal review.
B	2022-02-09	A. Lear	All	Formatting, copy edits; prepared PDF for signatures and release.

Table of Contents

1 Introduction

1.1 Purpose and Scope

The purpose of this document is to define the design of the ngVLA Central Signal Processor (CSP) for the Conceptual Design phase of its development.

The design is driven by the requirements stated in [AD01] and the purpose of the design description is to define a design that can meet all the requirements stated in [AD01]. Compliance of the design to the requirements is defined in [AD02].

The design description is a holistic definition of the design, including performance, functional, mechanical, environmental, safety, reliability, availability and maintainability characteristics. The design should also show compliance to external interfaces in cases where the interfaces have a direct impact on the design.

2 Related Documents and Drawings

2.1 Applicable Documents

The following documents may not be directly referenced herein, but provide necessary context or supporting material.

Ref. No.	Document Title	Rev/Doc. No.
AD01	Central Signal Processor Requirements Specification	020.40.00.00.00-0001-REQ-B
AD02	Central Signal Processor Requirements Compliance	TBD
AD03	Integrated Receivers and Digitizers / Digital Back End Interface Specification	020.10.40.05.00-0002-ICD
AD04	Power Supply / Antenna Electronics Interface Specification	020.10.40.05.00-0006-ICD
AD05	Antenna Bins, Modules and Racks / Antenna Electronics Interface Specification	020.10.40.05.00-0040-ICD
AD06	Digital Back End / Monitor and Control System Interface Specification	020.10.40.05.00-0076-ICD
AD07	Antenna Fiber Distribution Interface Specification	020.10.40.05.00-0041-ICD
AD08	ngVLA Site Buildings Interface Specification	020.10.40.05.00-0095-ICD
AD09	Central Signal Processor / Monitor and Control System Interface Specification	020.10.40.05.00-0105-ICD
AD10	Central Signal Processor / Online Subsystem Interface Specification	020.10.40.05.00-0114-ICD
AD11	Central Signal Processor / Central Fiber Infrastructure Interface Specification	020.10.40.05.00-0119-ICD
AD12	Digital Back End / LO Reference and Timing Interface Specification	020.10.40.05.00-0122-ICD
AD13	Central Signal Processor / LO Reference and Timing Interface Specification	020.10.40.05.00-0123-ICD
AD14	ngVLA System Requirements	020.10.15.10.00-0003-REQ
AD15	ngVLA Integrated Receivers and Digitizers Design Description	020.30.15.00.00-0004-DSN
AD16	Monitor and Control Hardware Interface Layer: Reference Design Description	020.30.45.00.00-0004-DSN

Ref. No.	Document Title	Rev/Doc. No.
AD17	DC Power Supply: Reference Design Description	020.30.50.00.00-0002-DSN
AD18	ngVLA Safety: Risk Analysis Procedures	020.80.00.00.00-0002-PRO
AD19	L1 System Technical Budgets	020.10.25.00.00-0002-DSN

2.2 Reference Documents

The following documents are referenced within this text:

Ref. No.	Document Title	Rev/Doc. No.
RD01	Central Signal Processor: Preliminary Reference Design	020.40.00.00.00-0002-DSN
RD02	Trident Correlator-Beamformer for the ngVLA: Preliminary Design Specification	NRCC Doc. #TR-DS-000001
RD03	Trident 2.0 Concept: A Minimum Delta Update to the Central Signal Processor Reference Design	ngVLA Electronics Memo #4
RD04	Trident 2.1 Concept: Updates to the CSP Reference	ngVLA Electronics Memo #5
RD05	Experiments with Calibrated Digital Sideband-Separating Downconversion	M. A. Morgan and J. R. Fisher 2010 PASP 122 326
RD06	Interferometry and synthesis in radio astronomy	A. R. Thompson, J. M. Moran, and G. W. Swenson, 2017
RD07	A SCREAM-Compatible ngVLA Pulsar Engine: Key Requirements Review and Option Trade-Off Study	ngVLA Electronics Memo #11
RD08	A SCREAM-Compatible ngVLA Cross-Correlation Engine: Key Requirements Review and Option Trade-Off Study	ngVLA Electronics Memo #10
RD09	A GPU Based X-Engine for the MeerKAT Radio Telescope	G. M. Callanan (2020), Master's thesis, University of Cape Town
RD10	Digital Back End/Data Transmission System: Reference Design Description	020.30.25.00.00-0002-DSN
RD11	ngVLA Radio Frequency Interference Forecast	ngVLA Memo #48
RD12	Handbook of Pulsar Astronomy	D. Lorimer and M. Kramer, 2005
RD13	Arista 7800R3 Series Quick Look	Arista Networks, Inc. 2021
RD14	Arista 7800R3 Platform Architecture	Arista Networks, Inc. 2020
RD15	An Integrated Circuit for Radio Astronomy Correlators Supporting Large Arrays of Antennas	L. R. D'Addario and D. Wang, JAI 5.02, 2016

3 Subsystem Overview

3.1 High Level Description

The Central Signal Processor (CSP) is a heterogeneous subsystem consisting of various digital signal-processing (DSP) sub-elements that, together, convert the digitized voltage from each active receiver at the antenna into "raw" data products, such as uncalibrated visibilities, average pulse profiles, or beamformed digital voltage for data recorders. The different data products are generated through a set of CSP observing modes. The CSP design uses the concept of subarrays at its core and allows flexible allocation of resources to multiple subarrays without affecting one another's operation. Notwithstanding their flexible allocation, any new subarray may be limited in resources by previous subarrays depending on its operating mode. The CSP shall be populated with enough resources to satisfy CSP0018, Simultaneous Subarray Capabilities [AD01].

The CSP is composed of three sub-elements performing DSP tasks: the Digital Back End (DBE), the Sub-Band Processor (SBP), and the Pulsar Engine (PSE). Along with them, the CSP Switched Fabric (CSF) routes the data outputs from each sub-element to the next sub-element in the processing chain. This architecture, graphically represented in Figure 1, is the result of multiple iterative design processes and reveals itself as the most suitable for ngVLA system requirements [RD01, RD02, RD03, RD04].

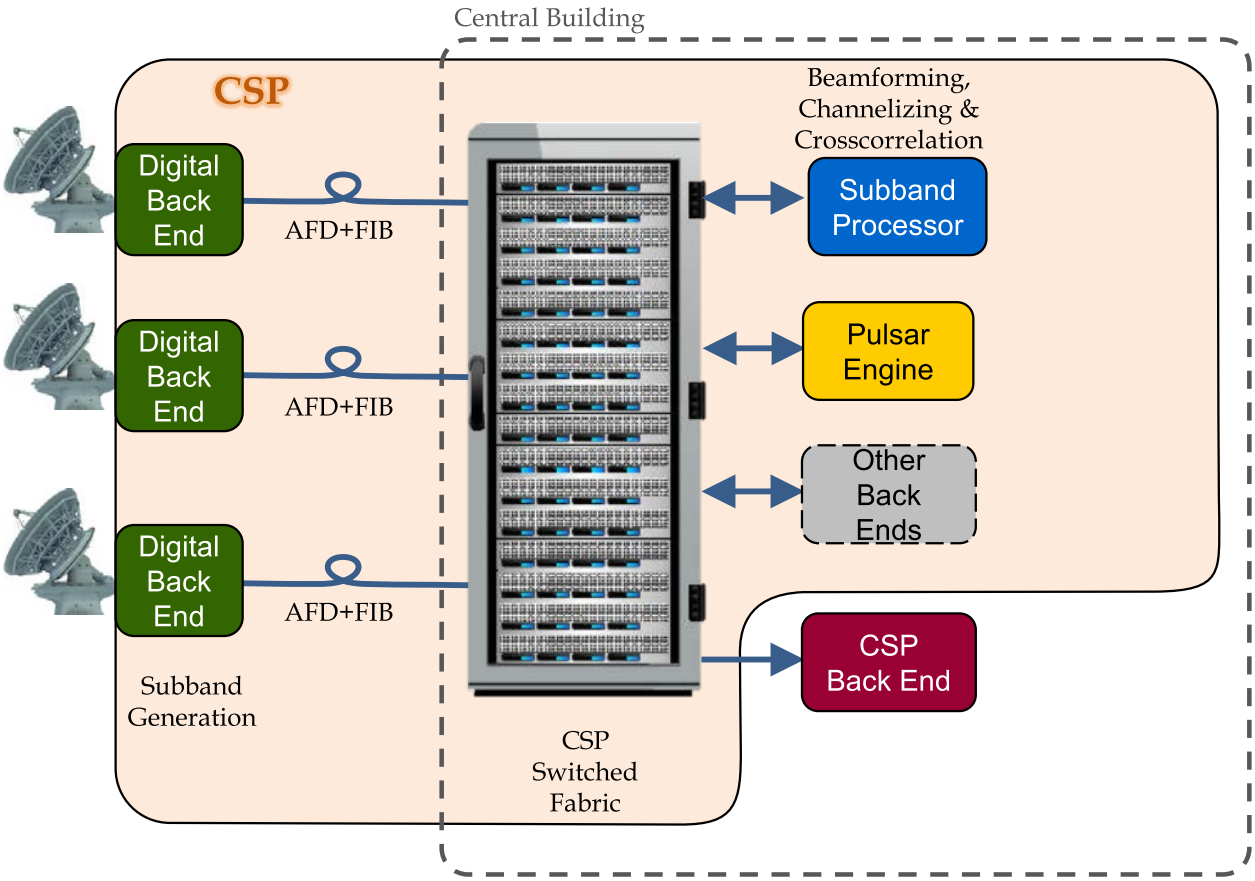

Figure 1: Central Signal Processor architecture.

The CSP follows an F-F-X architecture when operating in interferometric mode, and an F-B-F or F-F-B (where B stands for Beamforming) when operating in beamforming modes, depending on whether true

time delay or phase-shift beamforming is used, respectively. Regardless of the operation mode, the first F-Engine that splits the digitized voltage into frequency sub-bands is implemented by the DBE. The DBE is located at the antenna sites, with the option of moving some of the units into the same building as the rest of the CSP for the closest sites. This option is only available to antennas whose distance to the CSP building is shorter than the Integrated Receiver and Digitizer (IRD) maximum transmission range. Each DBE unit processes all the digitized bandwidth for one antenna, although one DBE unit may comprise a multiplicity of identical DBE modules (2 modules in the current design concept), each module processing the signal from different receivers. Therefore, the number of units scales with the number of antennas, but the size of each individual unit depends on the processed bandwidth.

After the DBE, the remaining DSP tasks depend on the observing mode and are carried out by the SBP. As the DBE, the SBP consists of a set of independent units, but in this case each unit processes a pair of sub-bands (or a single sub-band in high-res mode) for the whole array. Thus, the number of SBP units depends on the total instantaneous bandwidth the CSP must process, while the size of each individual unit depends on the number of antennas. In addition to the SBP, the ngVLA design of the CSP includes a custom back end, the Pulsar Engine (PSE), to perform DSP tasks specifically related to transient analysis, such as dedispersion and folding. The PSE receives beamformed data from the SBP and generates data products other than visibilities. The CSP output generated by either the SBP or the PSE is sent to the Computing and Software Subsystem (CSS) for further processing and archiving.

Except for the DBE, the other CSP sub-elements are housed in the central processing building. The data from the various antennas' DBE are routed to the proper SBP unit by the CSF. As required by the observing mode, the data from the SBP can be sent to the PSE via the CSF for additional processing. Ideally the CSS subsystem will connect to both sub-elements through the CSF as well. The alternative is to provide a separate network through which the SBP and the PSE can transmit their data products to the CSS. On the other hand, a communications network different from the CSF will carry the monitor and control data between the CSP and the CSS. The monitor and control data network is not designed with the capability necessary to sustain the CSP data output.

In addition to the DBE, the SBP, the PSE, and potentially the CSS, the CSF would also provide enough spigots to allow other custom back ends (e.g. SETI) to subscribe to sub-band data from the DBE, finely channelized[1] or beamformed data from the SBP, correlated data from the SBP (a.k.a. visibilities), or dedispersed and/or time folded data from the PSE. If the CSS is finally connected to the CSF, the output of those custom back ends could be sent to the CSS using the CSF, e.g., for archiving purposes. This assumes that the CSS is provided with the required resources.

3.2 Design Driving Requirements

A subset of the key requirements that drive the design is shown in Table 1 (next page). Most of these requirements become design drivers only when considered along with other requirements. They are collectively shown as one in the table as well as in the subsequent discussion.

[1] The data flow from the second F-Engine to the X-Engine, that is, phase-delay-corrected finely frequency channelized data from each antenna, is only available for SCREAM-based SBP units.

Parameter	Summary of Requirement	Reference
Subarray Capabilities	The CSP shall support subarray operation with combinations and capabilities equal to or greater than the functionality described in [AD14].	CSP0018
	The configuration of a subarray shall be completely independent of all other subarrays operating on different antenna subsets.	CSP0019
Beamforming Capabilities	The CSP shall support producing a minimum of 10 beams, with a goal of 50 beams, distributed over the active subarrays.	CSP0029
	The CSP shall be able to generate each beam with at least 8 GHz or the full instantaneous bandwidth of the band in use, whichever is less, with a goal of generating the full available bandwidth specified by CSP0008.	CSP0030
	The CSP shall generate cross-correlation products for one of the phase centers while operating in beamforming mode.	CSP0034

Table 1: Key Central Signal Processor requirements.

3.2.1 Subarray Independence and Simultaneous Subarray Capabilities

These two CSP requirement constitute the main driver of the CSP design. Subarray independence refers to the capability of the CSP to process the data from different subarrays with the same capabilities as if each subarray was the only one observing at any given time.

In the current design, each SBP unit, which processes one portion of the bandwidth for the whole array, can only operate in one functional mode. This limits the new subarrays to only the unused SBP units if they observe in a different functional mode from the subarrays currently observing. As a result, an active subarray can leave new subarrays depleted of SBP units or severely decrease their observation bandwidth. This limitation is implementation dependent. Section 5.6 describes an alternative architecture achieving full subarray independence.

The Simultaneous Subarray Capabilities defined in [AD14] allow the designer to establish a best-value trade-off between full subarray independence and single functional mode operation for the CSP. In addition, the discussion in the CSP Requirements Specification [AD01] under CSP0019 and CSP0058 differentiates between primary subarrays (operating in different subsets of antennas) and secondary subarrays (operating in the same subset of antennas as a primary subarray). Full subarray independence between primary subarrays is achievable by some CSP designs, but it is understood that some limitations must be imposed by a primary subarray on its related secondary subarrays, such as the RF Band and potentially the sub-band selection. On the other hand, many other configuration parameters of observations carried out through secondary subarrays may be configured independently of any other subarray, particularly its corresponding primary subarray. This includes the observation start time, as well as the time and frequency resolutions.

3.2.2 Beamforming Capabilities

The capabilities of the CSP in beamforming modes drive the computational requirements of the SBP, i.e., the amount of hardware resources it needs to perform all the required functions. As a reference, the

current TALON hardware, on which the SBP design is based, needs at least as much as twice the hardware to process the same amount of bandwidth in beamforming modes as compared to interferometric mode. And even so, the number of beams generated is less than the required 10, particularly if simultaneous visibilities are to be generated as well.

Similar conclusions have been obtained from the alternative SBP design (SCREAM, see Section 5.6). The SCREAM architecture uses separate hardware for the beamformer and the correlator. Initially, both sub-elements were based on ASIC technology for power efficiency. After evaluation of ngVLA requirements, it became evident that the beamformer is the dominant sub-element in terms of required hardware, justifying the use of ASICs. On the contrary, the correlator is small enough to be implemented using FPGA technology for a fraction of the cost.

3.3 Key Risks

The following key risks have identified in the design and development of the Central Signal Processor.

3.3.1 Lack of Canadian Support

While the DBE and the PSE are based on in-house development, the development of the SBP relies on a partnership with the National Research Council of Canada (NRC). The NRC serves Canadian interests as defined by the Canadian scientific community. There is a possibility that ngVLA does not obtain enough support to prioritize the ngVLA CSP over other similar projects, e.g., SKA or ALMA.

3.3.1.1 Mitigation Strategy

There is an ongoing research program at NRAO aiming at developing highly power-efficient devices for future radio astronomy interferometers. The Scalable, Reconfigurable, and Modular (SCREAM) project is currently at its early stages and does not show a competitive Technology Readiness Level as compared to this SBP design. Nonetheless, it can be used as an alternative SBP design in case the international partnership for the SBP does not materialize. Unfortunately, no low-risk alternative to the current SBP design has been identified. The SCREAM design is described in Section 5.6.

3.3.2 Global Chip Shortage

There is an ongoing crisis in which the demand for integrated circuits is greater than the supply, leading to major shortages of integrated circuits across the industry. The CSP design and development phase relies on the use of hardware prototyping platforms, as well as the fabrication of custom hardware designs. There is a high risk that these activities may be severely impacted by the global chip shortage, introducing delays of years in the development of the system. Current forecasts do not expect the CSP would be impacted beyond its design phase.

3.3.2.1 Mitigation Strategy

Not much can be done in this regard other than placing equipment purchasing orders as soon as possible to place high in the queue. This strategy is already in place, with some purchasing orders being placed more than one year before the equipment is going to be used. Hardware development activities should be brought forward as much as possible in the project, as those can be significantly impacted by the extended procurement periods.

3.4 Design Assumptions

This design of the Central Signal Processor is based on the following assumptions.

3.4.1 TALON Hardware Upgrade

The design of the SBP is based on the Frequency Slice Processor (FSP) of NRC's Frequency Slice Architecture (FSA) [RD02]. The FSA is the technology selected for the central signal processor of the future SKA1-Mid radio telescope. This design is based on the TALON hardware, which currently employs a Stratix-10 device based on Intel's 14 nm FinFET technology. This device does not incorporate enough memory and computing resources to satisfy the needs of the ngVLA CSP design described herein. Only a small delta over current capabilities has been incorporated under the assumption that a future technology node will enable such capabilities. For example, the sub-band bandwidth has been increased from 200 MHz to 218.75 MHz, or the data communications network is based on 400G technology, which is not supported by the current TALON hardware. This hardware upgrade assumption must be reassessed at PDR and verified during the development phase, contingent on the materialization of a Canadian participation in an international ngVLA partnership.

If this assumption is not met, there are two alternatives to cope with it. One is that the SCREAM architecture (Section 5.6) becomes the only viable option for the SBP design. The second option would be modifying the current CSP design to adapt the capabilities of the TALON hardware.

3.4.2 Digital Sideband Separation at the DBE

The CSP design assumes that the analog receiver response is smooth enough for a (reasonably long) Finite Impulse Response (FIR) filter approximation can be used to achieve the desired sideband rejection. This is related to a low calibration density, as described in [RD05]. More details on digital sideband separation can be found in Section 4.2.1.1.

If this assumption is not satisfied, it may imply that the DBE lacks the hardware resources to effectively perform the digital sideband separation. In that case, sideband separation can alternatively be carried out in the frequency domain at the SBP, after frequency channelization occurs. However, this requires the pertinent sub-bands from both sidebands to be sent to the same SBP unit to be processed. This imposes some constraints on the down-selection of sub-bands at the DBE prior to transmission. Moreover, each SBP unit processes only two sub-bands in interferometric standard resolution modes.[2] The CSP design should be modified in order to support digital sideband separation in high resolution and beamforming modes as well.

3.4.3 Technological Evolution of the CSF Switched Fabric

This design assumes that the technology used for the CSF allows a full non-blocking matrix switch topology in which all the CSP computing nodes can communicate to each other. This is something to be expected given the current pace at which networking technology evolves and the ever-increasing demand for higher bandwidth from many diverse industries.

However, current commercial solutions fall short in satisfying the connectivity needs of the CSP. A quick search across the market revealed that a crossbar switch with 2,000 400G ports is not available yet. Notwithstanding that, the imminent adoption of 112G transceivers by the industry, along with the first proof-of-concept CPO (Co-Packaged Optics) prototypes, make it likely that such a solution will be COTS available by the time the ngVLA project production phase begins.

[2] Refer to Section 4.2.1.4 for a definition of interferometric standard and high-resolution modes.

If this assumption is not satisfied, Section 5.4 describes how the CSF architecture can be partitioned in the bandwidth and the antenna domains, so that the port count decreases to the level of current COTS solutions. Unfortunately, such partitions are not exempt of disadvantages, among which is an increase in the DBE complexity and an overall increase of the CSF port count.

4 Central Signal Processor Design

4.1 Product Structure

4.1.1 Product Context

Figure 2 graphically represents the product context of the CSP housed in the central processing building, i.e., the SBP, the PSE, and the CSF. As seen in the figure, the DBE is treated in this regard as an independent subsystem because its location in the antenna pedestal results in a different product context. The product context of the DBE is shown in Figure 3 (next page).

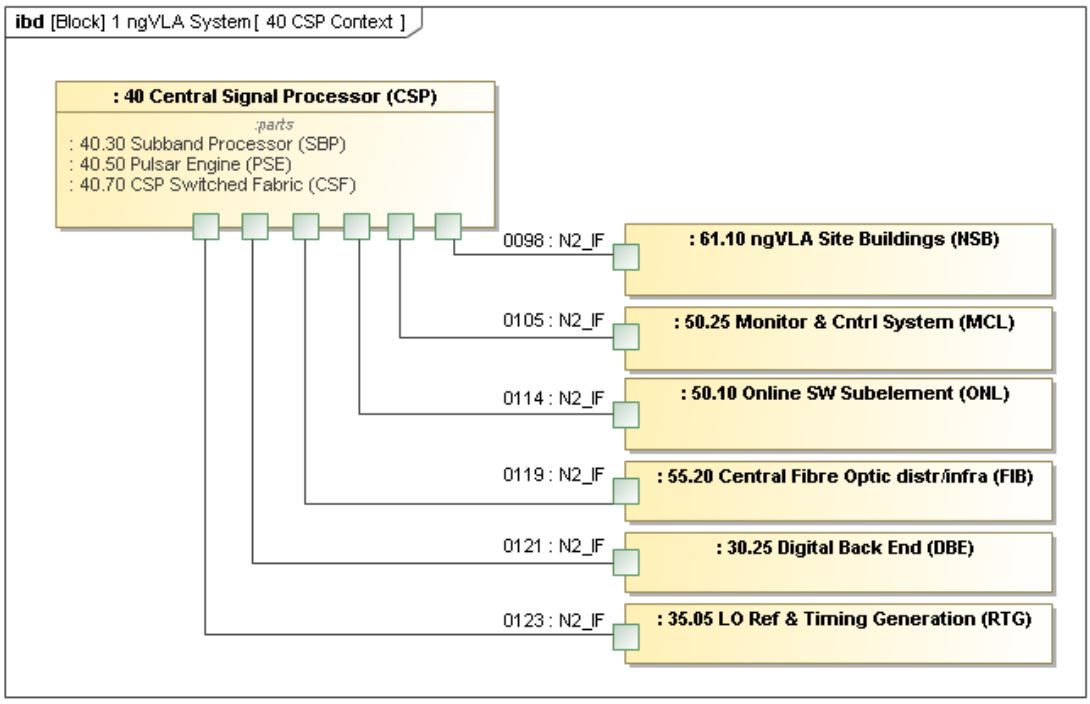

Figure 2: Central Signal Processor product context.

In this context, the CSP equipment is located at the central processing building within the ngVLA Site Buildings subsystem (NSB, CI number 020.61.10.00.00). The NSB subsystem provides the CSP with power, HVAC room cooling, physical space and RFI shielding [AD08]. At the antenna, the Bins, Modules & Racks sub-element (BMR, CI number 020.30.55.00.00) must provide the DBE with room, cooling and RFI shielding [AD05], while the DC Power Supply sub-element (PSU, CI number 020.30.50.00.00) provides it with power [AD04].

Data digitized at the antenna stations must be transmitted to the central processing building after being processed by the DBE. Data transmission is physically done through an interface between the DBE and the Antenna Fiber Optic subsystem (AFD, CI number 020.30.70.00.00) as per [AD07]. At the central

processor building, data from the antenna sites is received through an interface between the CSP and the Central Fiber Optic Distribution/Infrastructure subsystem (FIB, CI number 020.55.20.00.00) [AD11]. This interface connects the DBE to the CSP Switched Fabric (CSF, CI number 020.40.70.00.00) sub-element of the CSP.[3] Hence, in general, data transmitted from the DBE at the antenna is delivered to other CSP sub-elements in the central building by the AFD and the FIB subsystems.

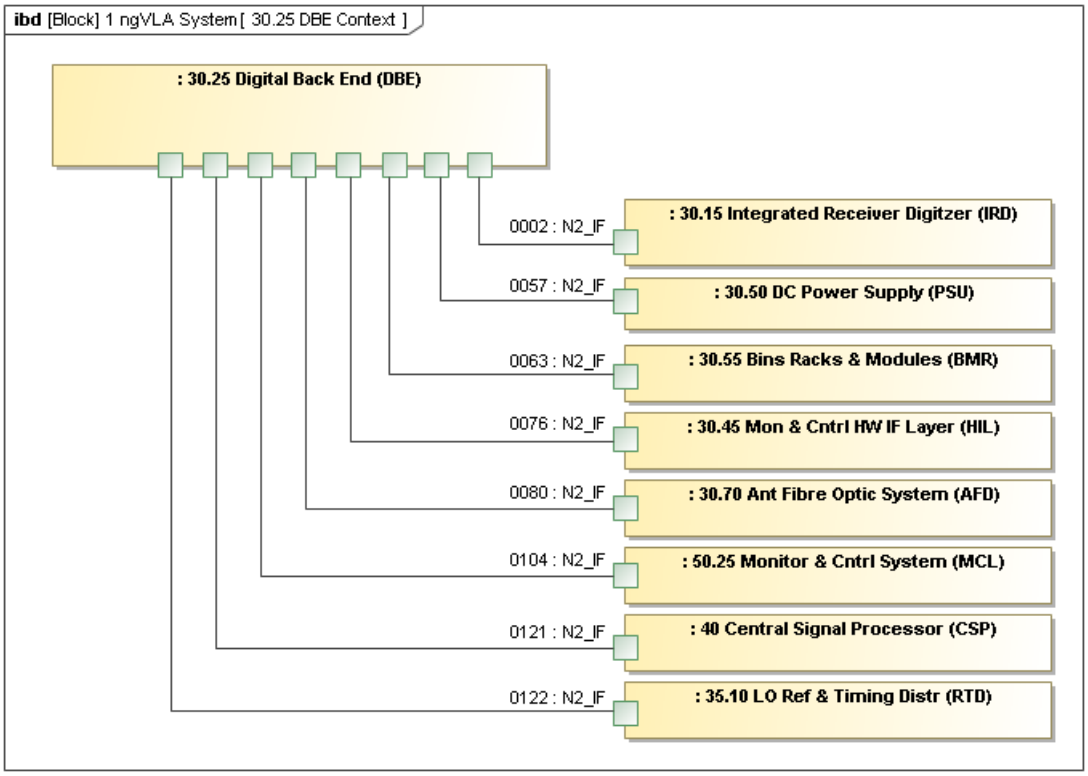

Figure 3: Digital Back End product context.

Time and frequency references needed by the CSP are obtained from its interfaces with the LO Reference and Timing Generation subsystem (RTG, CI number 020.35.05.00.00) [AD13] and the LO Reference and Timing Distribution subsystem (RTD, CI number 020.35.10.00.00) for its DBE sub-element [AD12].

Data products generated by the CSP are sent to the Computing and Software System (CSS, CI number 020.50.00.00.00), specifically its Online sub-element (ONL, CI number 020.50.10.00.00) [AD10]. The Monitoring & Control sub-element (MCL, CI number 020.50.25.00.00) of the CSS also interfaces with the CSP at various levels [AD06, AD09].

Finally, the Integrated Receiver Digitizer (IRD, CI number 020.30.15.00.00) generates the digital data that forms the CSP data input, which particularly received and processed by the DBE, whose common interface is described in [AD03].

[3] Notwithstanding that, the relocation of the nearest DBE units to the central signal building would be considered in the future if deemed beneficial.

4.1.2 Product Breakdown Structure

Consistent with the Product Context, the Digital Back End is separated from the Central Signal Processor in the Product Breakdown Structure (PBS), and falls under the Main Antenna System (18AS, CI number 020.12.00.00.00) instead. Considering that, the excerpt from the PBS relevant to the CSP (including the DBE) is shown in Figure 4.

020.12.00.00.00 (18AS) Main Antenna System
020.30.00.00.00 Antenna Electronics
020.30.25.00.00 (DBE) Digital Back End
020.30.25.10.00 D501 DBE Module
020.40.00.00.00 (CSP) Central Signal Processor
020.40.30.00.00 (SBP) Sub-Band Processor
020.40.30.10.00 SBP Line Type #1 (FSA-based)
020.40.30.40.00 SBP Local Monitor and Control Line
020.40.30.50.00 SBP Local Network Line Type #1
020.40.30.60.00 SBP Cooling System Line
020.40.50.00.00 (PSE) Pulsar Engine
020.40.50.10.00 PSE Sub-Band Processor Line
020.40.50.20.00 PSE Local Monitor and Control Line
020.40.70.00.00 (CSF) CSP Switched Fabric
020.40.70.10.00 CSF Chassis
020.40.70.20.00 CSF Line

Figure 4: Central Signal Processor product breakdown structure.

4.1.3 Block Diagram

The block diagram of the CSP, including major components as well as internal and external interfaces, is shown in Figure 5 (next page). As can be seen, the CSP external interfaces involve the DBE, which interfaces other antenna subsystems, and then the CSP in the central processing building interfaces other subsystems. The internal interfaces between the CSP sub-elements are performed via the CSF. In principle, there is no internal interface between the DBE and the PSE. In addition, both the SBP and the PSE transmit their output to the Online Sub-Element (ONL) through the CSF as specified in the CSP/ONL interface specification [AD10].

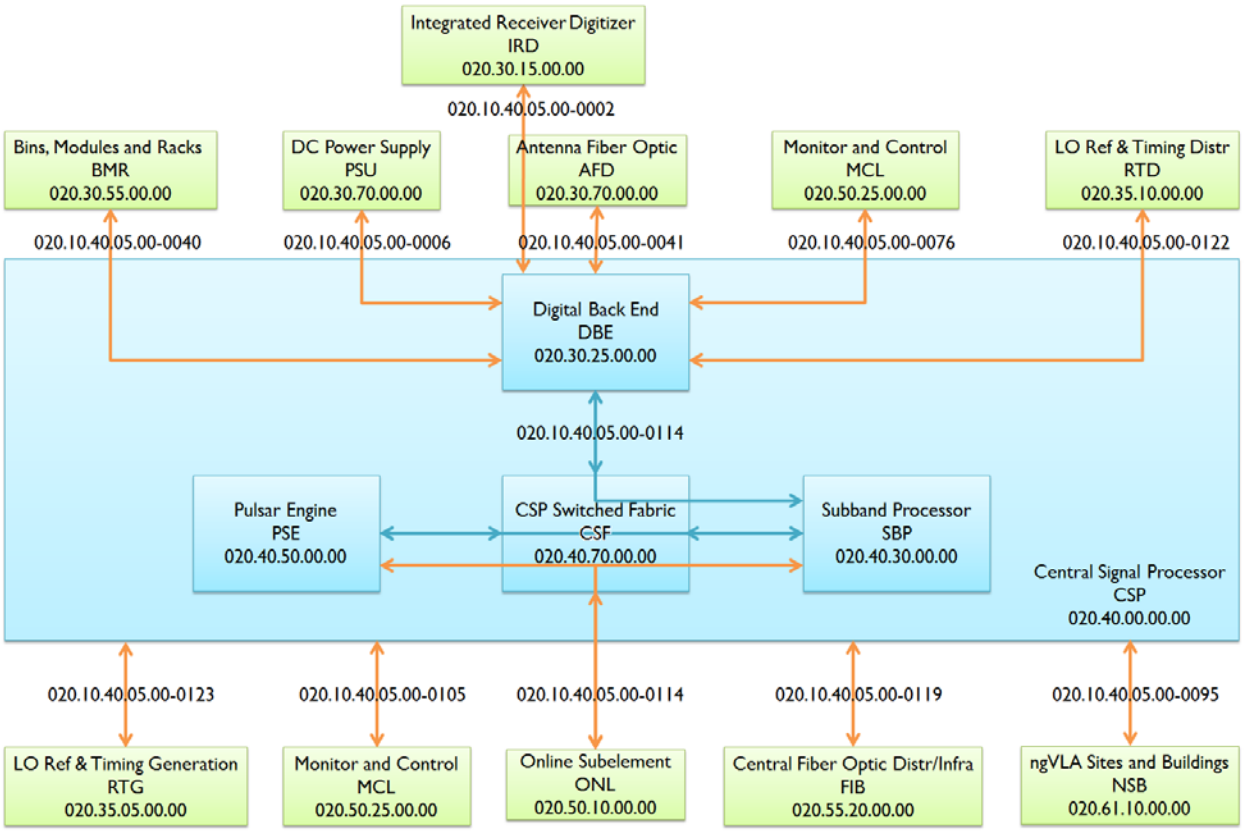

Figure 5: Central Signal Processor block diagram.

4.2 Product Design

The CSP ingests the digitizers' output at the antenna and produces low-level data products to be ingested by a high-performance computing system known as the CSP Back-End (CBE), which forms part of the Online Sub-Element (ONL). In addition to cross-correlation and auto-correlation capabilities, the CSP support further capabilities required of modern telescopes to enable VLBI and time-domain science. Specifically, the CSP operates in at least four different Observing Modes (OM), as per CSP requirements [AD01], depending on the desired data product:

- **Interferometric OM**: The CSP computes parallel and cross-polarization auto- and cross-correlation functions, per frequency channel, within a subarray. Spectral resolution and time averaging are independently configured across subarrays, and on a per-sub-band basis within each subarray, enabling spectral zoom windows for simultaneous spectral-line and continuum observations in the same subarray.
- **Pulsar Timing OM**: The CSP generates up to ten full bandwidth beams through beamforming, arbitrarily distributed across all subarrays. More beams are possible at reduced bandwidth. Each beam is independently dedispersed, detected, and folded according to given dispersion and timing models in order to generate an average pulse profile per frequency channel per beam.
- **Transient Search OM**: The CSP generates up to ten full bandwidth beams through beamforming, arbitrarily distributed across all subarrays. More beams are possible at reduced bandwidth. Each beam is converted to full Stokes parameters at a given time-frequency resolution, and finally sent to the

CBE for archiving and offline processing. The design includes the spigots necessary for an eventual transition to an Online Transient Search Back End.

- **VLBI OM**: Finally, this Observing Mode is thought to operate the ngVLA as one or multiple VLBI stations within a larger network. When operating in this mode, the CSP can generate up to ten full bandwidth beams arbitrarily distributed across all subarrays, although more beams are possible at reduced bandwidth. The resulting voltage stream is sent to the CBE for archiving in a VLBI standard formats.

These Observing Modes reveal the need for at least four sub-elements: a correlator, a beamformer, a Pulsar Engine (PSE) to compute pulsar profiles, and a communications network that interconnects these sub-elements, i.e., the CSP Switched Fabric (CSF). The current CSP design integrates the correlator and the beamformer into a single sub-element, i.e., the Sub-Band Processor (SBP). However, the alternative SCREAM architecture for the SBP (see Section 5.6) splits the SBP into a Beamformer and Channelizer (B&C) and a Cross-Correlation Engine (XE). Regardless of the architecture, both the SBP and the PSE are based on modular units that process a portion of the bandwidth (typically one or two sub-bands) for the whole array.

Thus, an additional sub-element, the Digital Back End (DBE), is needed to channelize the wideband data stream from the digitizer into such sub-bands. Other ancillary sub-elements, such as the local monitor and control and the temperature control sub-elements, are also critical for the CSP operation. However, these ancillary sub-elements are not considered key for the design or the cost of the CSP and will be deferred to future design reviews.

In the rest of this section, the functional, electronic, and mechanical design of the CSP and its sub-elements is fully described.

4.2.1 Functional Architecture

A high-level description of the CSP architecture has been given in Section 3.1. Briefly, the data flow and signal processing chain of the CSP are as follows: First, the DBE receives the digital data stream from the digitizers at the antenna, splits it into multiple frequency sub-bands, and selects and transmits those sub-bands that will be processed further at the central processor build by the SBP. The functional diagram of the DBE is shown in Figure 6 (next page). The diagram only represents the data flow of one receiver, so that multiple receivers are processed equally in parallel. Each IRD module consists of 2 receivers, one per polarization. Hence, the signal processing chain shown in the figure must be duplicated for each active IRD module to process both polarizations. After receiving the digitized data from the IRD modules, the DBE applies a calibration filter on the received signal in order to equalize the spectrum and further suppress the unwanted sideband from the down conversion process, effectively performing sideband separation digitally. Then, a frequency shifter coarsely tunes the digitized spectrum and removes any per-antenna LO-shift as needed. After that, the bandwidth is split into frequency sub-bands through an oversampled filter bank. Finally, the selected sub-band data streams undergo a requantization process and the resulting data packets are timestamped for its transmission to the SBP.

The sub-band processors receive data from all the antennas corresponding to one sub-band pair, or a single sub-band in high-res mode. Figure 7 depicts the functional diagram of the SBP in interferometric mode. The antenna data streams are first corrected for the bulk of the delay, which accounts for any variable network delay compensation, and then upsampled to a common reference time scale. This resampling process also compensates for any differences in sampling clocks at the various antenna sites (intentionally or unintentionally inserted, if known). Next, a time-variable phase correction is applied, before finely channelizing the sub-band bandwidth into narrow frequency channels. Alternatively, spectral zooming involves a digital down conversion and decimation process prior to the frequency channelization.

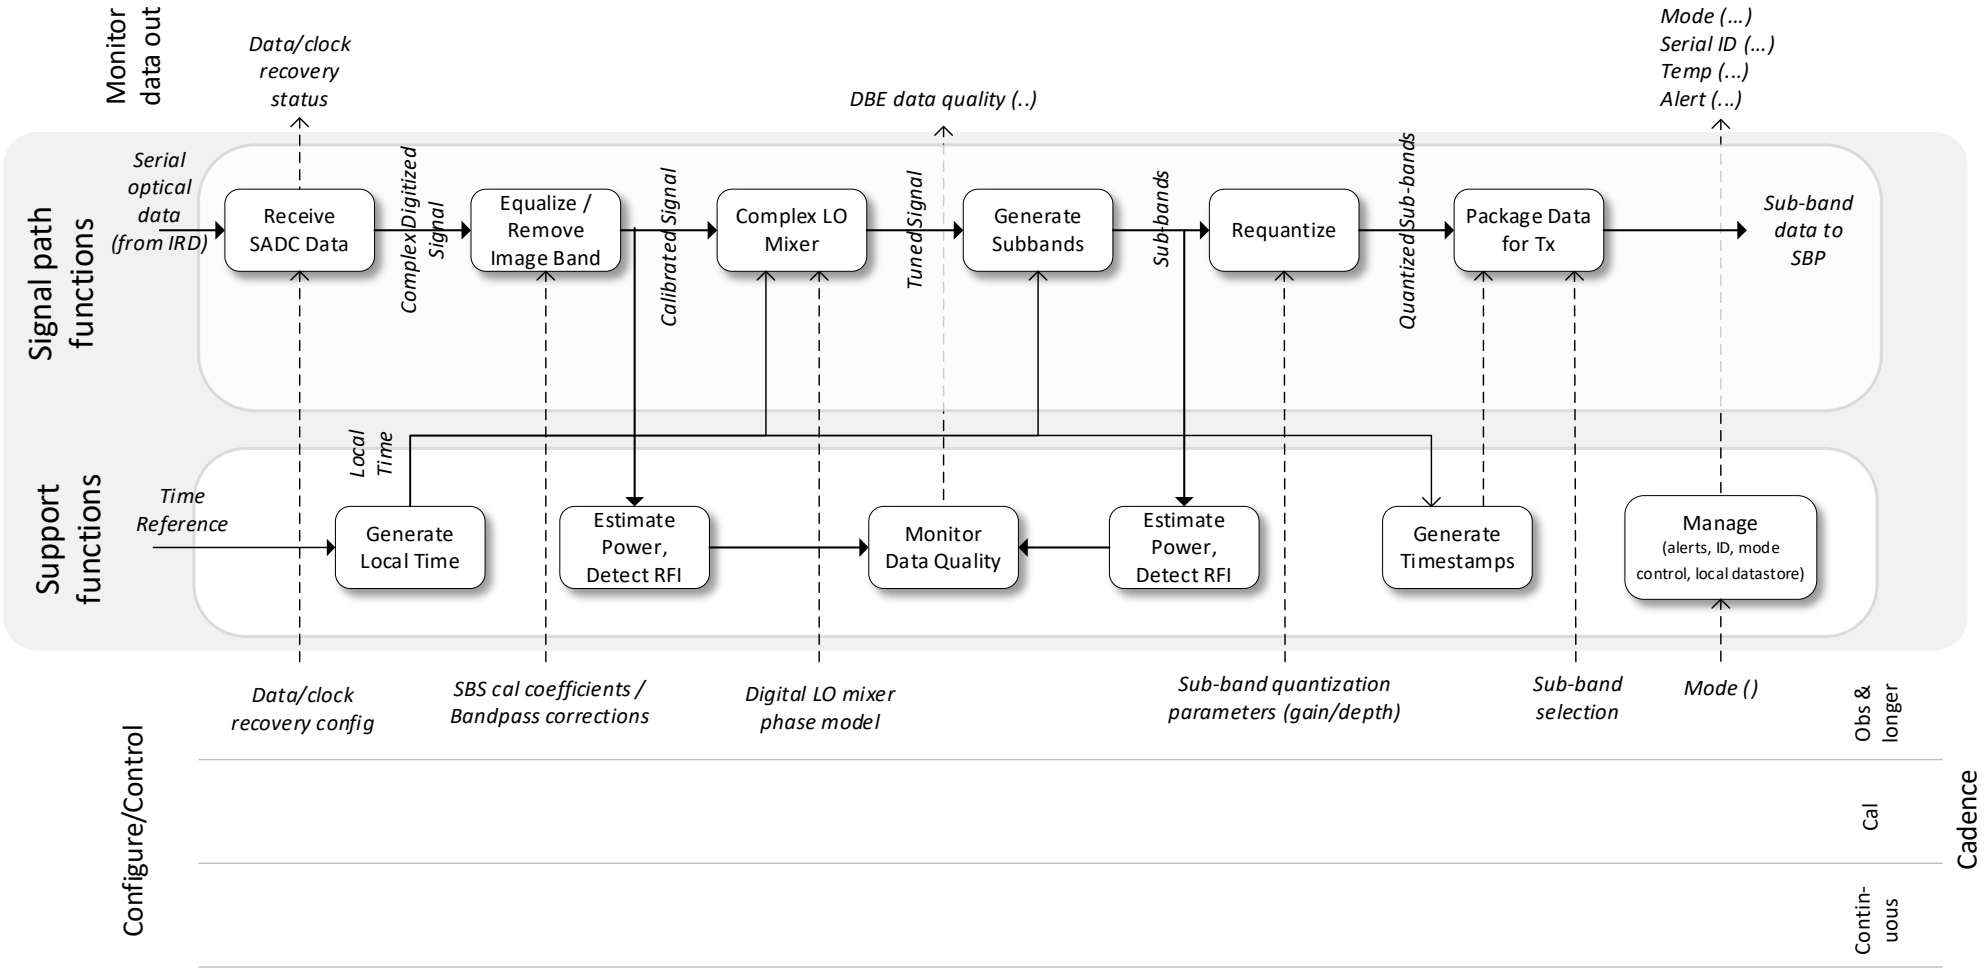

Figure 6: Functional overview of the Digital Back End.

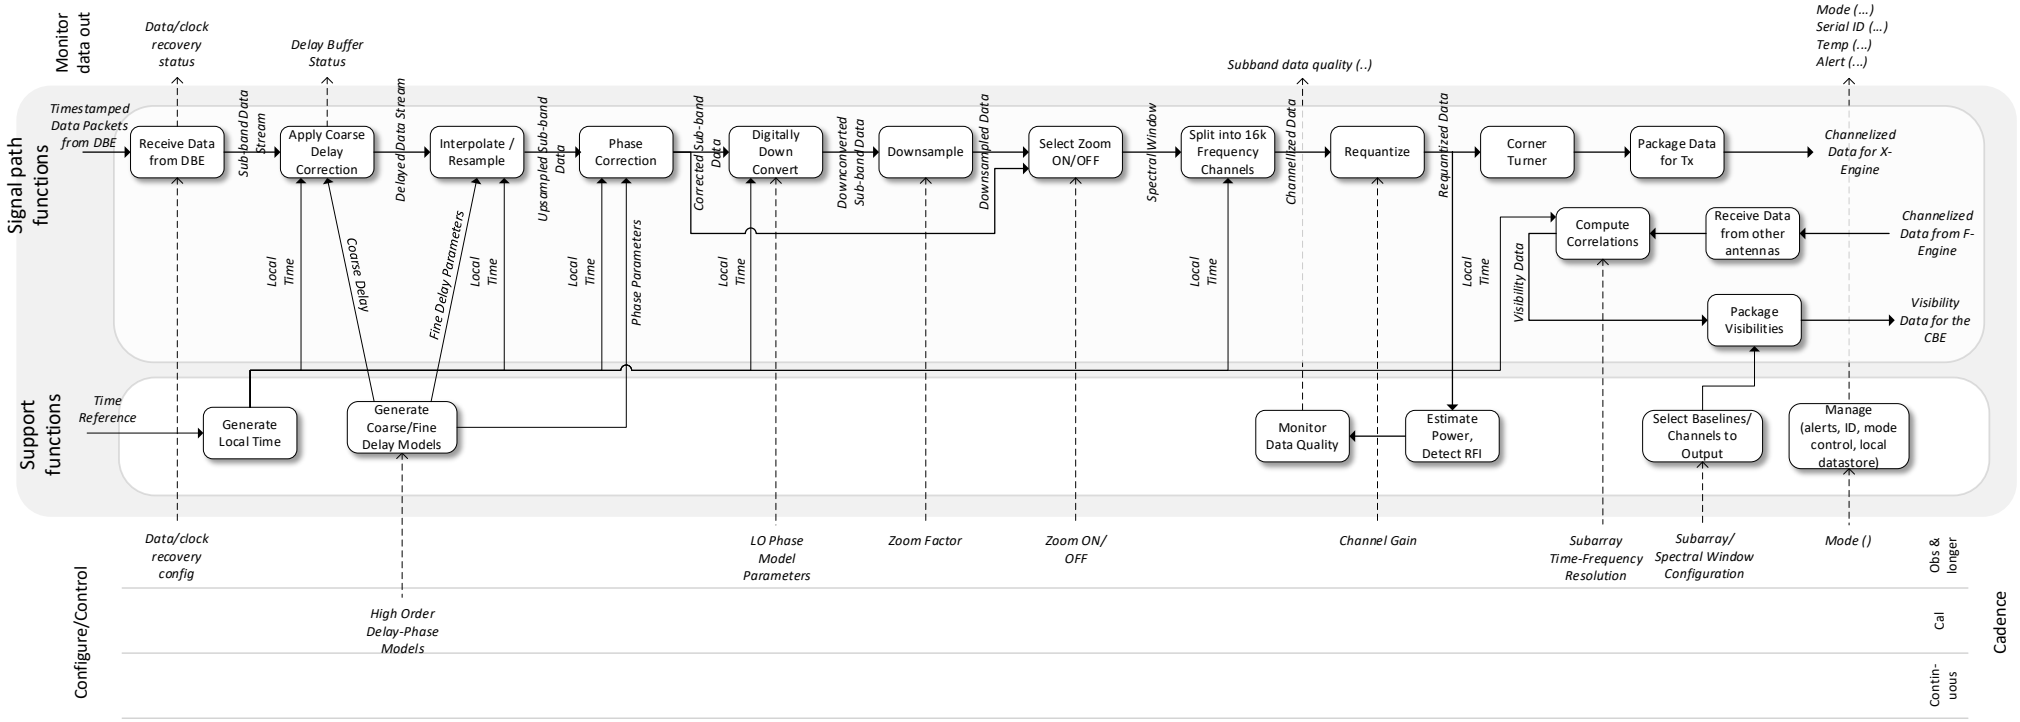

Figure 7: Functional overview of the Sub-Band Processor operating in interferometric mode.

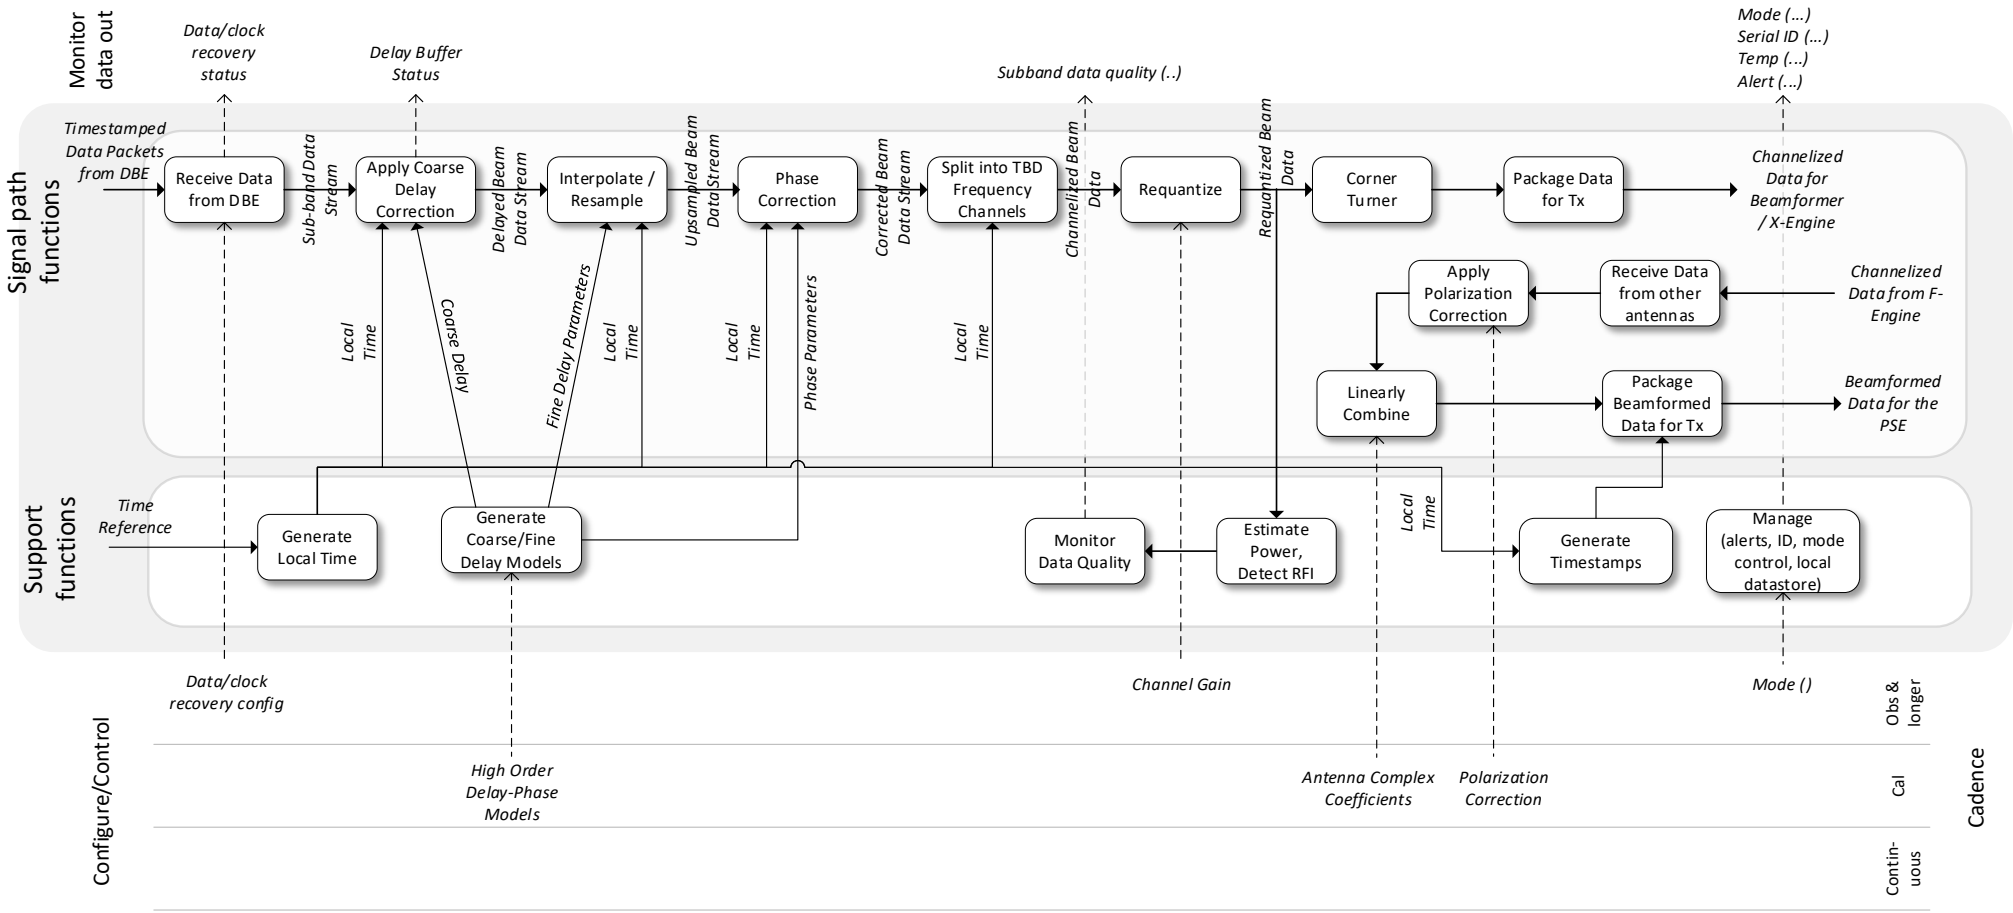

Figure 8: Functional overview of the Sub-Band Processor operating in beamforming mode.

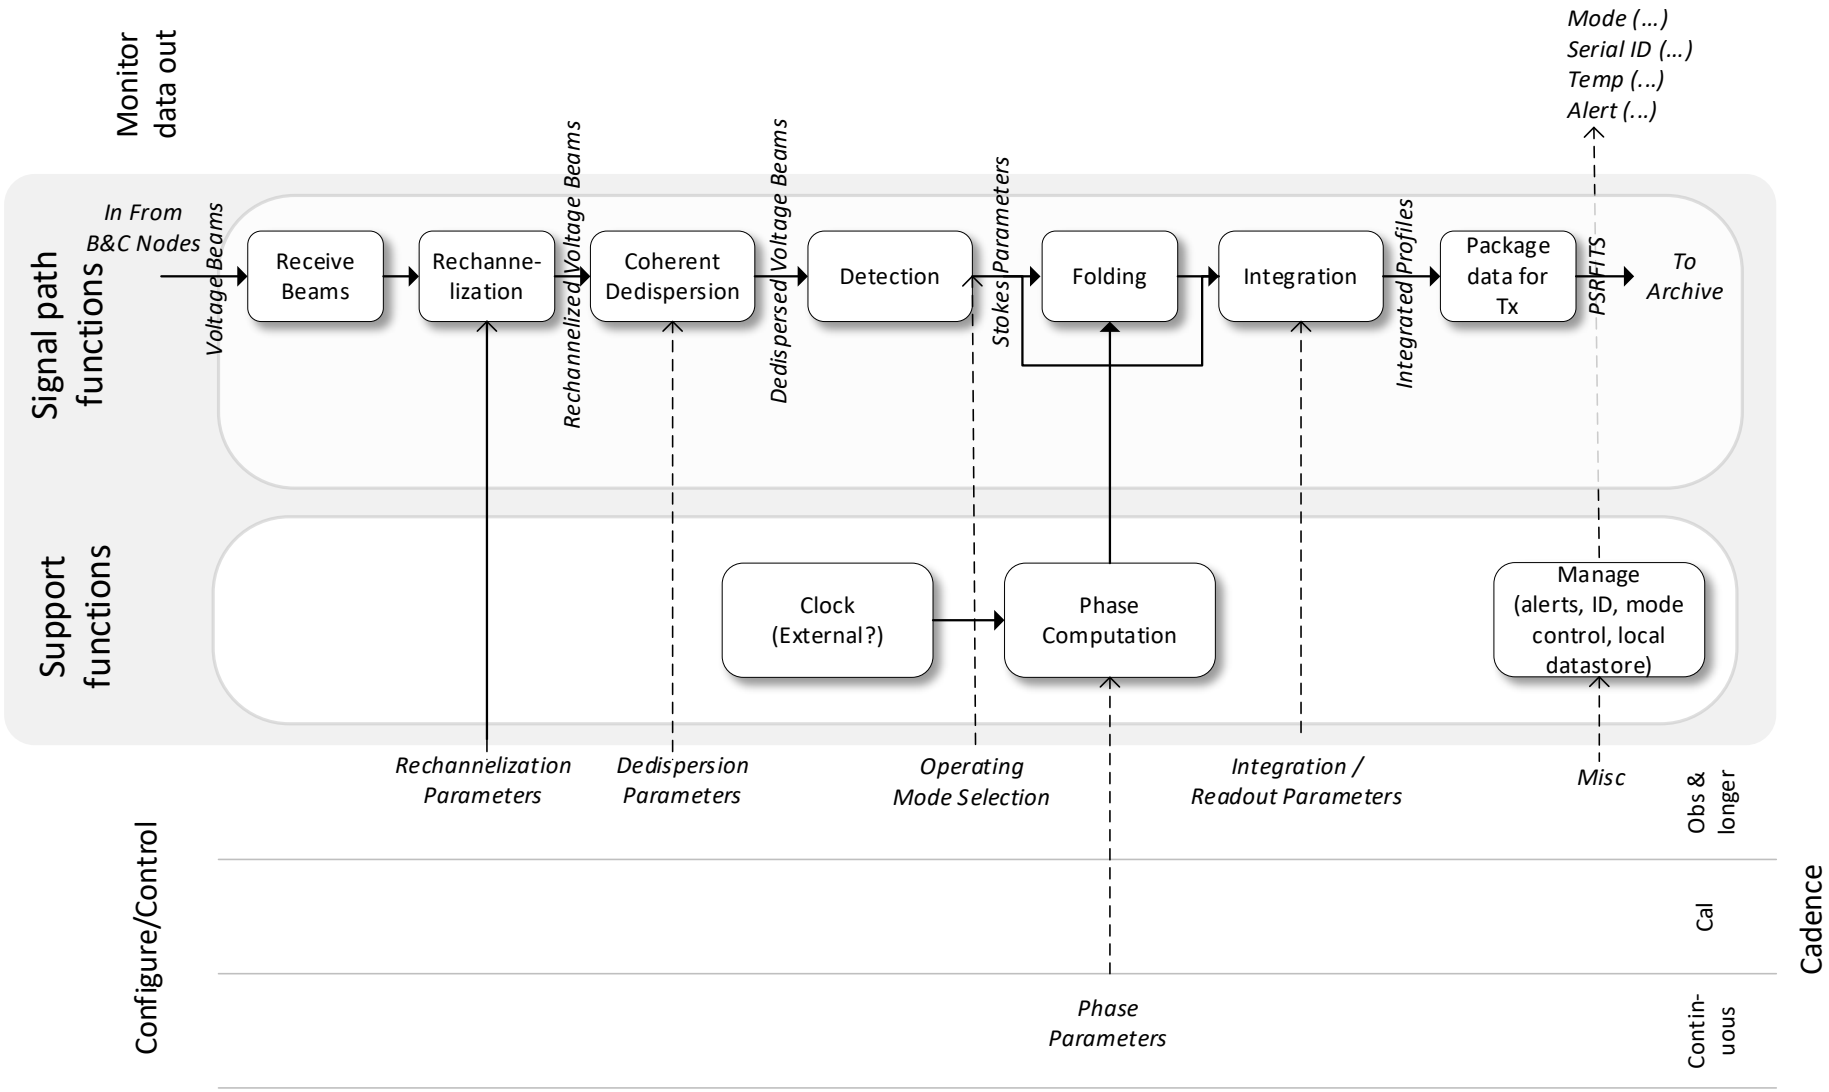

Figure 9: Functional overview of the Pulsar Engine.

The resulting data streams are then requantized, packetized, and transmitted internally[4] across the SBP computing nodes for cross-correlation. The correlator then performs a complex multiply-and-accumulate operation on a baseline pair basis according to the desired time resolution. Channel averaging may be carried out as well at the correlator. Finally, the uncalibrated visibilities are sent to the CBE for further processing and archiving.

In beamforming mode, the functional diagram of the SBP is slightly different, as represented in Figure 8. In this mode, which does not allow spectral zooming, the frequency channels from the antennas composing a subarray are linearly combined as opposed to cross-correlated. To meet ngVLA beamforming efficiency requirements [AD01], True Time Delay (TTD) beamforming is required given the maximum array aperture and spectral resolution. The maximum number of full bandwidth beams, 10, does not demand unreasonable hardware resources. While operating in beamforming mode, the SBP also generates concurrent visibilities for the phase center corresponding to one of the generated beams. The required additional functions are those of the interferometric mode (see Figure 7). However, some restrictions apply in terms of maximum frequency resolution (no spectral zoom modes available).

When operating in Pulsar Timing and Transient Search OMs, the SBP is also operating in beamforming mode. The difference now with respect to VLBI OM is that the beams formed are sent to the PSE for further processing prior to the CBE. The functional diagram of the PSE is included in Figure 9. After receiving the beamformed data streams from the SBP via the CSF, the PSE carries out a coherent dedispersion onto the data. This process may be preceded by a spectral rechannelization as needed. After dedispersion, Stokes parameters are generated through detection. Next, data can be folded according to a timing model and then integrated (Pulsar Timing OM), or simply accumulated in the absence of a timing model to reduce the output data rate as needed (Transient Search OM). In any case, the accumulated data is finally packetized and sent to the CBE for further processing and archiving.

In the following sections, each of the above functions is described in more detail, along with a specification of the responsible sub-element and pertinent Observing Mode.

4.2.1.1 Digital Sideband Separation/Coarse Bandpass Correction

The ngVLA down conversion scheme employs a direct conversion or homodyne receiver. Both the upper and lower sidebands can be retrieved from the in-phase and quadrature components digitized at the output of the analog receiver. The IRD is the subsystem responsible for the down conversion and digitization of the RF signal.

Due to nonidealities of the analog processing chain of the receiver, the amplitude and phase response of the in-phase and quadrature paths are not perfectly matched to each other. This creates some leakage between sidebands that must be compensated for at the DBE, as it is anticipated that the analog receiver alone cannot provide enough isolation to meet ngVLA dynamic range requirements.

The concept of digital sideband separation has been fully explained in the literature (see, e.g., [RD05]). The process consists of separately filtering both uncalibrated in-phase and quadrature components and then combining them together in order to form the calibrated components. A total of four digital filters are required. Two filters could be possible if no bandpass calibration is needed at this stage. Note that this calibration cannot recover the spectral gap around zero IF.

Digital sideband separation is performed as a first stage in the time domain. However, this process is normally carried out in the frequency domain, after fine channelization of the signal. In the CSP design, fine frequency channelization occurs in the last stages of the SBP processing, just before cross-correlation or beamforming. This creates a problem in the current design because both sidebands are needed for

[4] A SCREAM-based SBP transmits the narrow frequency channels to the X-engine via the CSF. That's why these data streams are accessible to a custom back end connected to the CSF.

sideband separation, but the SBP only has capability to process two sub-bands in interferometric standard resolution modes. High resolution and beamforming modes would then need a different solution.

Sideband separation could also be done in the frequency domain at the DBE, after sub-band channelization. However, it is expected that it can be more efficiently done in the time domain for two reasons: only 32 sub-bands are generated per receiver, limiting the frequency resolution of the calibration filter; and sub-band generation is carried out after a baseband LO Mixer finely tunes the digitized spectrum, which negatively affects the accuracy of the calibration process as both sidebands experience opposite frequency shifting at the mixer.

The number of FIR coefficients needed for the calibration filters is yet to be determined. This will depend on the amplitude and phase mismatch in the analog receiver paths of the in-phase and quadrature components. Such analog receivers are currently being designed.

The length of the calibration FIR filter that can be implemented at the DBE has not been determined yet. This requires starting firmware development activities of the DBE. The targeted technology for the DBE hardware is state-of-the-art FPGA. Because of their finite length and finite precision in coefficient quantization, the calibration filters applied differ not only from the ideal response, but also from the calibration tables obtained by calibration of the receiver. The design must guarantee that these approximations are accurate enough to achieve the required sideband separation for the System (30 dB, with a goal of 40 dB).

Filter coefficients shall be stored by the DBE for faster reconfiguration purposes. They shall be loaded at bootup and upgradable from the standard monitor and control interface. They shall be under version control as well as the rest of the DBE firmware. A standard procedure to compute the sideband separation filter coefficients from calibration tables obtained through receiver characterization shall be determined and included within the DBE documentation package.

4.2.1.2 Baseband LO Mixer/Coarse LO Frequency Shift Correction

After digital sideband separation at the DBE, the baseband is frequency shifted prior to sub-band channelization. The objective is twofold: coarsely removing any per-antenna LO frequency shift intentionally introduced at the receiver LO, and tuning the sub-bands' frequency range, which is fixed in the subsequent filter bank. Notwithstanding that, fine frequency tuning is carried out at the SBP as part of the phase tracking process. The main goal of this mixer is to minimize any correlation losses due to sub-band frequency misalignment across antennas, which may be caused by a receiver's LO or sampling clock offsetting.

As regards the implementation of the mixer, a prospective design simply consists of a Numerically Controlled Oscillator (NCO), which generates a complex exponential that multiplies the incoming signal. The NCO shall be phase-dithered to minimize its spurious level.

4.2.1.3 Received Signal Analysis/Sub-Band Generation

Sub-Band data streams are generated by the DBE by means of a Polyphase Filter Bank (PFB) based on the Fast Fourier Transform (FFT) algorithm. Given the high sampling frequency of the digitizer, 7 GS/s [AD03], splitting the incoming data stream into multiple streams at a much more manageable sampling rate becomes very convenient and is inherent to the FFX architecture chosen for the ngVLA CSP.

One key parameter of the CSP design is the number of sub-bands generated at the DBE, or more specifically, its bandwidth. These are 32 sub-bands and 218.75 MHz, respectively. As described in Section 5.2, these figures are the result of an iterative process and the best trade-off between design complexity and current capabilities of the TALON hardware.

The PFB output is oversampled so that the extra bandwidth can be used to relax the design of the spectral response of a sub-band channels, as well as obtain much better performance than a critically sampled filter

bank, as required by ngVLA dynamic range requirements. The chosen oversampling factor is 8/7, which results in a sub-band sampling frequency of 250 MHz. This sampling frequency has many advantages in terms of timing and VLBI compatibility, as discussed in Section 5.3.

Finally, it is important to note that although the 32 sub-bands cover the full baseband, the IRD receiver design only guarantees minimum required performance for the central 5.8 GHz of the 7 GHz band. Therefore, only the 27 sub-bands around zero IF are normally selected for further processing. Although the DBE design does not prevent the selection of other sub-bands closer to the edges of the receiver frequency range.

4.2.1.4 Sub-Band Requantization

Sub-Bands that have been selected for further processing are requantized prior to transmission. The RMS amplitude of each sub-band is independently adjusted by the DBE prior to requantization. This process requires sending the RMS amplitude measurements through the monitor and control interface and receive each sub-band gain setting. The cadence of these measurements is similar to other sensors, in the order of 0.1s or less. The use of per-sub-band automatic gain control devices prior to quantization will be studied in future design phases.

The SBP allows two different resolution modes in interferometric OM: a standard resolution 8-bit mode and a high-resolution 16-bit mode. The SBP units using standard resolution can process two sub-bands in interferometric mode. This capability is not currently supported in beamforming modes.

In principle, the DBE will requantize sub-bands up to the maximum resolution allowed by the destination SBP unit. However, less resolution in the requantization processes is also supported by the DBE. This can prove useful when controlling the DBE output rate, particularly in those antenna stations connected to the central processing building via commercial infrastructure. See ngVLA Electronics Memo No. 5 [RD04, Table 3] for the DBE maximum output data rate as a function of the observing RF Band.

4.2.1.5 Delay Correction/Digital Resampling

Sub-Bands are timestamped at the DBE and sent to the SBP for further OM-dependent processing. In both beamforming and interferometric modes, the SBP must apply the delay corrections so that the signals from different antennas are received aligned in time at the beamformer and correlator input, respectively. This is done by keeping track, through calibration, of the antenna time reference in reference to the central time reference. Digital data streams are sampled and time stamped based on the antenna reference time. The ngVLA CSP must support different antenna time references at least for the outer stations in the array. It is desirable that the ngVLA be eVLBI capable as well. In addition, it is a goal that inner stations can use different sampling clock frequencies to increase the robustness against spurious signals related to the sampling clock. To enable all these features, the CSP must convert the sub-band data streams from different antennas to a common central reference time prior to cross-correlation or beamforming.

The conversion from antenna sampling rate to the central sampling time is done at the SBP through a FIR filter-based resampler. Since the SBP is in the central processing build, it has direct access to the central time reference signal from the RTG subsystem. Fine delay tracking is performed within the resampling process. The FIR filter-based resampler uses a collection of filters to apply a time-varying group delay correction to the incoming signal. The resampling effect is achieved by applying a linearly varying group delay to the periodically sampled input time series. Since the resampler continuously applies a delay correction to the processed sub-band, it is convenient that the applied delay model also accounts for the instrumental delay.

To minimize any aliasing effects, the sub-bands from different antennas must be up-sampled to a higher common sampling frequency. For convenience, 256 MHz is chosen as the sub-band sampling frequency at the correlator. This simplifies the SBP resources needed for VLBI OM, as 256 MHz is a standard VLBI

sampling frequency. The conversion from the sub-band input, nominally sampled at 250 MHz, can be achieved by applying a group delay that is decreased by 93.75 ps ([250 MHz]$^{-1}$ – [256 MHz]$^{-1}$) with every input sample. In order to do that, the incoming signal is passed through a time-variant FIR filter whose taps are updated with every input sample. The filter coefficients are chosen from a collection of FIR filters, each one with a different group delay. The maximum group delay variation required in the filter collection is just one sample. Whenever the next group delay to be applied goes out of range, the current input data stream at the FIR filter implementation is kept for another output sample, and the applied group delay is increased by one sample, which brings its value back to valid range. Equivalently, this up-sampling process can also be understood as an interpolation.

The current SBP design uses a collection of 1024 filters. This results into a delay correction error better than ±2 ps. Assuming the SBP resampler is provided with perfect delay models, the SNR loss for a narrow frequency channel due to the quantization of the delay correction can be computed according to equation (7.43) in [RD06]:

$$L_{\Delta\tau} = 1 - \left[\frac{\sin(\pi\nu\tau_0)}{\pi\nu\tau_0}\right]^2 \qquad (1)$$

where ν is the center IF of the frequency channel, and τ_0 is the quantization step of the delay correction. For sub-bands with 218.75 MHz information bandwidth (two sided) and about 4 ps delay quantization step as indicated above, the resulting SNR loss is approximately 0.000063%. The loss is smaller in continuum observations because the SNR loss is averaged across the sub-band spectrum. Such a small loss is not justified by CSP requirements, for which an overall correlation loss of up to 1% is allowed. Hence, future revisions of the design might increase the delay correction quantization step for resource optimization.

The delay correction is applied in two stages: coarse delay and fine delay corrections. The resampler can only apply small delay corrections. The bulk of the delay is applied at the SBP input data buffers with sample resolution. The SBP input buffers are needed to support widely variable delays of the communications network, particularly from commercial ones. Received packets with a timestamp not older than approximately 250 ms from the central time must be processed by the SBP. This implies buffering at least the last 250 ms of data. Data packets received out of order are chronologically arranged in the SBP based on their timestamps. Most of the delay correction is achieved through proper adjustment of the read pointer of these buffers.

The CSP receives current delay models on a low cadence through its monitor and control interface. The order of the delay models are high enough so that they are accurate for at least 10s intervals. These high-order delay models are valid from their applicability time to the validity time of the next received model. The CSP local monitor and control is responsible for distributing the high-order models across the pertinent SBP units where the real-time approximation by first-order delay models is done. The constant term of the first-order model is used to adjust the input buffer read pointer, while the delay drift is input to the filter selection algorithm. Continuity of the delay correction solution is guaranteed by the continuity across successive high-order models.

The objective of the resampler is to produce sub-band data streams that have been sampled not only at the same rate but also with the same time reference. To minimize the amount of data buffering required at the correlator input, or the beamformer input depending on the OM, the generation of delay-corrected data streams must keep certain synchronism across all devices within an SBP Unit. With this objective, the time retrieved by the CSP from the RTG subsystem is synchronously sent to every processing node within an SBP Unit using a similar approach as for SKA1-Mid and EVLA WIDAR correlators.

4.2.1.6 Sub-Band LO Mixer/Phase Correction

In addition to correcting for the delay, the SBP also applies the required phase correction for fringe stopping. The phase model generation process is very similar to the delay model generation, with real-time computation of first-order phase model occurring in the SBP. The generated instantaneous phase correction is fed into an NCO that produces a complex exponential that itself multiplies the delay-corrected data stream.

4.2.1.7 Spectral Zoom

CSP requirements call for a maximum interferometric frequency resolution of at least 1 kHz, with a goal of 400 Hz [AD01]. However, this resolution is not required for the whole processed bandwidth. The maximum number of frequency channels that must be supported is at least 240,000 channels, with a goal of supporting 2,000,000 frequency channels. Since this number is smaller than the number of channels resulting from splitting the desired processed bandwidth, 20 GHz, into channels of 1 kHz, it makes sense to enable the finest interferometric frequency resolutions though spectral zooming.

The maximum spectral resolution at which the whole processed bandwidth can be processed is 15.625 kHz. This resolution results from splitting one sub-band, sampled at 256 MHz, into 16,384 ($=2^{14}$) frequency channels (see Section 4.2.1.8). For 20 GHz of processed bandwidth, the overall number of channels at this resolution is 1,280,000 channels. For spectral resolutions narrower than 15.625 kHz (about 4 km/s at 1.2 GHz), the SBP operates in zoom mode and trades processed bandwidth in for spectral resolution.

In zoom mode, only a portion of the sub-band bandwidth is processed by the frequency channels that follow the zoom engine, while keeping the maximum number of frequency channels, 16,384. Therefore, the spectral resolution improvement is proportional to the bandwidth reduction or the zoom factor. Only zoom factors that are a power of 2, i.e., 2, 4, 8, and so on up to 64 can be selected. Hence, the maximum frequency resolution becomes slightly less than 250 Hz, meeting the CSP goal. The maximum bandwidth processed at this resolution is 312.5 MHz, for 1,280,000 channels, and assuming the same number of SBP Units as needed for processing 20 GHz of bandwidth in non-zoom modes.

The zoom engine consists of two main components, a complex NCO-based mixer to tune the portion of the bandwidth that will be selected, and decimating FIR filter that serves as anti-aliasing filter and downsampler. The filter coefficients are selected from a collection of FIR filters for each zoom factor. Both the phase and frequency of the NCO are free tuning parameters defining a first-order model for the phase. Note that the zoom engine oscillator and mixer can be integrated with the ones used for correcting the phase for resource optimization purposes.

Finally, the zoom mode is available only in Interferometric and VLBI OMs. In Interferometric OM, the zoom mode can be enabled on a per-sub-band and per-antenna basis, consistent with the required frequency resolution configurability. In VLBI OM, the zoom engine could be used after the beamformer to decrease the bandwidth per beam.

4.2.1.8 Frequency Channelization

The SBP includes a frequency channelizer that splits the input data stream into narrowband frequency channels. The input of the frequency channelizer are the delay-phase corrected sub-band data, and in beamforming OMs, the data could be beamformed or not.

In Interferometric OM, the frequency channelizer splits the incoming signal into 16,384 ($=2^{14}$) frequency channels. The resulting channel bandwidth is 15.625 kHz. As explained in Section 4.2.1.7, better spectral resolutions are achieved by trading in processed bandwidth via zoom modes. Coarser spectral resolutions are also necessary for controlling the output data rate while keeping radial smearing performance within acceptable levels. This spectral averaging process is implemented more efficiently at the correlator.

The number of frequency channels implemented for Pulsar Timing and Transient Search OMs still needs to be assessed by firmware development. These two OMs exceed the capabilities of the current technology node and rely on either a future hardware upgrade of the TALON hardware [RD02] or the SCREAM design (Section 5.6). The ideal number of channels used in Pulsar Timing OM is such that no rechannelization is required at the PSE. However, the need for simultaneous visibility generation makes uncertain the feasibility of this approach. If the number of frequency channels needs to be reduced, the beams would be rechannelized at the PSE as described in Section 4.2.1.13. Similar considerations apply to the Transient Search OM, where the requirement of simultaneous visibilities may impose analogous constrains to the achievable spectral resolution. The SCREAM design uses a multi-stream multi-resolution frequency channelizer, so that these constrains are not applicable.

The frequency channelizer is based on a PFB design optimized to meet the frequency selectivity and channels flatness requirement of the CSP [AD01]. In Interferometric OM, the frequency channels are critically sampled for optimal resource utilization. The same applies to the frequency channelizer used to produce simultaneous visibilities in beamforming OMs. On the contrary, the output of the frequency channelizer applied to the beamformed signal in Pulsar Timing and Transient Search OMs is oversampled. This is needed by the subsequent rechannelization and channel stitching processes performed by the PSE.

4.2.1.9 Corner Turner

The data produced by both the DBE and the SBP must undergo a "matrix transposition" or "corner turner" operation through which two-dimensional data is rearranged from one dimension into the other. This operation (Figure 10) is jointly carried out by both the digital logic and the communications network.

Each DBE unit processes the data from one antenna and produces a set of sub-bands. The DBE must separate the sub-band data into independent data streams for each sub-band. That allows the CSF to perform the corner turner operation. The data stream from one DBE unit is first demultiplexed into multiple sub-band data streams, and then each sub-band is multiplexed again, but this time with the data from other DBE units corresponding to the same sub-band. Thus, each data link from one DBE unit into the CSF contains multiple sub-bands corresponding to the same antenna but a different frequency band. In contrast, each data link out from the CSF to each SBP unit contains multiple sub-bands corresponding to the same frequency band but different antennas. Note that a single switch could perform the whole corner turner function if provided with enough throughput.

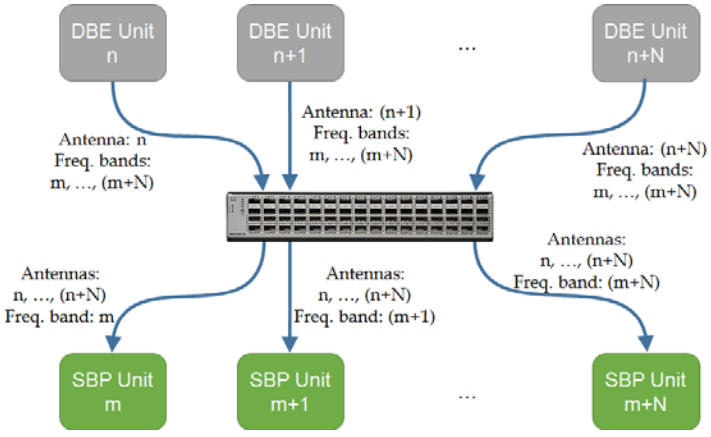

Figure 10: Corner turner operation between the DBE and the SBP.

The second corner turner operation is internal to each SBP unit. It occurs between the frequency channelizer output and the correlator or beamformer input. In this case, each frequency channelizer generates all the channels for one antenna, but every computing node of the correlator or the beamformer process a subset (or "bundle") of channels for all the antennas. Hence, the data at the frequency

channelizer output is rearranged in channel bundles, and each bundle from every antenna is sent to the same processing node for correlation or beamforming. In this case, each SBP Unit includes a passive fiber optic mesh, similar to the one shown in Figure 11, to interconnect all the processing nodes to one another.

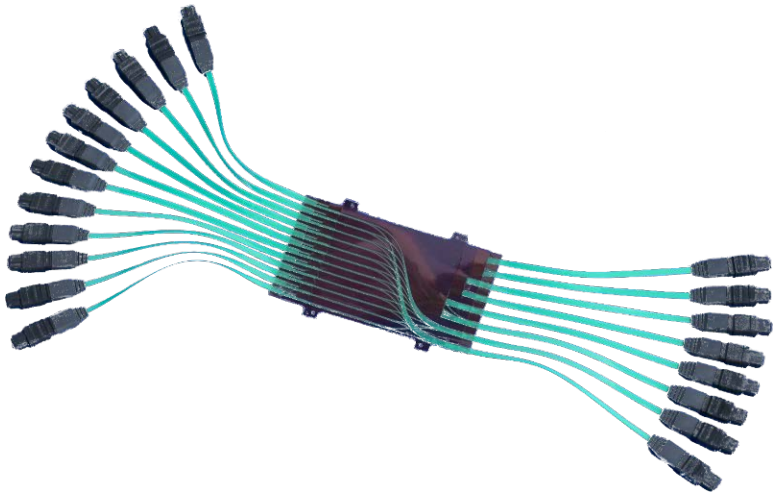

Figure 11: FlexPlane™ optical assembly from Molex.

4.2.1.10 Correlation/Spectral Averaging

Each SBP unit includes a correlator to compute all the auto and cross-correlation products for the whole subarray. The correlator is distributed across the same computing nodes that perform delay and phase tracking and the frequency channelization. Each processing node computes all the visibilities for a portion of the bandwidth, and hence all processing nodes within an SBP unit must be interconnected. The alternative SCREAM-based SBP design (see Section 5.6) follows a different approach in which the correlator does not share the same hardware with the per-antenna processing nodes (B&C part). In this approach, which uses more specialized hardware for power efficiency, the correlator connects to the B&C part via the CSF instead of the internal passive fiber optic mesh described in Section 4.2.1.9.

The correlator implementation is based on CMAC (Complex Multiply-Accumulate) units for estimating correlation products. These CMAC units multiply current data samples and add the result to previous accumulation results. The temporary accumulation results are stored in on-chip memory due to memory bandwidth requirements. The correlator also incorporates a small input data buffer for data alignment, as the data streams from different antennas are not received simultaneously. However, the time synchronization among processing nodes also guarantees that the required buffer size is kept at a minimum.

The SBP correlator not only allows integrating each baseline pair and frequency channel over time, but also multiple frequency channels can be added together in the processes. This is done by proper control of the address generator of the memory in which temporary results are stored.

For simplicity, the SBP design computes visibilities for all the baselines in the whole array, regardless of whether they make sense, e.g., when operating subarrays. Nonetheless, the CSP output data rate is not increased by this practice, as only the desired baselines are finally output. This contrasts to the SCREAM correlator (Section 5.6.2), where only the desired baselines are computed, at the cost of increased complexity of its controller. In return, each individual correlator nodes can process a higher bandwidth depending on the size of the subarray. The objective is to power off the unused nodes, or even reallocate them to increase the computational capabilities of the Pulsar Engine (PSE) as needed. With that objective in mind, the design of the SCREAM-based correlator targets the same hardware as the PSE.

The output of the correlator is finally packetized and send to the CSS, ideally via the CSF. Otherwise, an output switched fabric will be needed to gather and merge the various data streams flowing from the correlator nodes. The particular data format of the correlator output, e.g., FITS, will be determined in the pertinent ICD [AD10] and subsystem requirements [AD01].

4.2.1.11 Polarization Correction

All beamforming OMs allow the application of polarization correction coefficients in the form of Jones matrices, prior to beamforming on a per-antenna basis. Further work is required to determine if the required calibration must be done after frequency channelization, on a per-channel basis, or if calibration dynamic range requirements can be met by applying corrections on a per-sub-band basis, prior to the frequency channelizer. Polarization correction is always done by the SBP.

Zero-order (constant) models for the polarization calibration coefficient must be provided to the CSP by the MCL subsystem with fast enough cadency so that accuracy requirements are met. Like other calibration parameters, the models are applied by the CSP until a new set of parameters is received and applied.

4.2.1.12 Beamforming

Once the sub-band data streams have been delay- and phase-corrected, polarization calibrated and corrected to a common polarization angle, beamforming is a straightforward operation by which multiple antennas are linearly combined into a single data stream or "beam." Due to the reduced number of full-bandwidth beams required for the ngVLA (10 beams), the most suitable beamforming technique is True Time-Delay (TTD) beamforming. TTD beamforming requires one phase-delay tracker per beam, increasing the computational requirements, but does not impose constrains on the maximum array aperture or the maximum field of view of the beamformer.

Beamforming coefficients are complex and provided to the CSP through its monitor and control interface in a similar fashion as other calibration parameters. This is necessary for maximum control of the radiation pattern and nulling. Despite the complex nature of the beamforming coefficients, it is currently under study the possibility of applying only the coefficient amplitude as beamforming weighs, while incorporating the beamforming coefficient phase into the phase tracker model. This should allow minimization of the maximum data rate of the monitor and control interface.

Beamforming is carried out by the SBP. In the current TALON-based design, TTD beamforming is carried out through a distributed weighted added tree [RD02]. ngVLA beamforming requirements exceed the current hardware capabilities and rely on a TALON hardware upgrade (or a SCREAM-based SBP design) for their feasibility. For example, the maximum number of beams per subarrays is only 4 per SBP Unit, less than half the required number of beams. In addition, frequency channelization is only possible post-beamforming, which precludes any polarization correction on a per-channel per-antenna antenna basis. Per-antenna polarization corrections are still possible, but only at the sub-band frequency resolution. These limitations should be easily overcome by the next technology node.

The output of the beamformer is finally packetized and send to the PSE or the CSS, depending on the OM, ideally via the CSF. Otherwise, an output switched fabric will be needed to gather and merge the various data streams flowing from the beamforming nodes. Such an output network would connect the SBP, the PSE, and the CSS. The particular data format of the correlator output, e.g., VDIF, will be determined in the pertinent ICD [AD10] and subsystem requirements [AD01].

4.2.1.13 Rechannelization/Coherent Dedispersion

Propagation through the interstellar medium introduces a frequency-dependent delay which may be removed in the PSE [RD12]. The removal of this dispersion from the input signal is most efficient in Fourier space, which permits any desired rechannelization to occur at the same time.

Coherent dedispersion via the overlap-save method requires an FFT, after which the time-domain dedispersion kernel (and a window function, if desired) is applied. An inverse FFT then produces a dedispersed time series from which the uncontaminated center portion is extracted for further processing. The FFT length must be no less than that of the discarded samples, and should be significantly greater to permit efficient operation. In the most extreme case, an observation centered near 1.2GHz with a 1MHz channel width and a dispersion measure of 3000 pc cm^{-3}, the overlap-save convolution must discard ~14,000 samples; an FFT length of 16k results in only ~2,000 uncontaminated samples, while a 64k FFT length yields ~50,000 samples after only ~4.6 times as many operations.

The selection of the dispersion measure for a given beam is asserted to be a part of the observation configuration, as is the choice of a window function. The resulting kernels are recalculated only when required, which is expected to be infrequent (approximately on the scale of beam pointing selection).

Consistent with other implementations of overlap-save dedispersion [RD12] the PSE will use a uniform kernel (and FFT) length within a single sub-band; this greatly simplifies dedispersion operation with a negligible impact on efficiency. In the most extreme case (parameters as above) this results in an ~8.7% drop in computational efficiency across the 1200–1418.75 MHz sub-band; in most cases the loss in efficiency is negligible.

4.2.1.14 Detection/Folding/Integration

The conversion of the two beam-polarizations' complex electric field data (E_x, E_y) to real Stokes parameters (I, Q, U, V) requires the computation of the quantities:

$$\begin{aligned} I &= |E_x|^2 + |E_y|^2 \\ Q &= |E_x|^2 - |E_y|^2 \\ U &= 2Re\left(E_x\overline{E_y}\right) \\ V &= -2Im(E_x\overline{E_y}) \end{aligned} \tag{2}$$

Once detected, the Stokes parameters will be integrated either according to an observer-provided timing model (pulsar observing mode) or to a desired time resolution (transient search observing mode). Prior to this stage, the processing required by both observing modes is the same.

The preliminary pulsar observing mode requirements envision the division of pulse periods ranging from 1ms to 30s into as many as 2048 phase bins. The fundamental procedure is to take the sample's arrival time t and determine the pulse phase $\Phi(t)$; this maps to phase bin n. Each of the Stokes parameters are accumulated in the nth corresponding bin, and the nth counter is incremented. After the desired integration period, the parameters and counter are read out and zeroed.

The pulsar ephemeris and timing model are asserted to be the observer's responsibility; the PSE should be provided with a low-order model at a cadence sufficient to maintain phase alignment. In the case of highly relativistic pulsar binaries, it is possible that Doppler accelerations could cause large effective phase drifts during a sufficiently long integration; this will introduce limitations on the total possible integration length for these observations.

In the case of transient search mode data, which is simply the Stokes parameters incoherently integrated to a specified time-frequency resolution, the folding step is eliminated. The choice of time-resolution will be informed by the observer, but may be constrained by data transfer limitations.

4.2.1.15 RFI Detection/RFI Excision

RFI is expected to disrupt ngVLA operations with increasing severity across the lifetime of the observatory [RD11]. Cellular networks and satellite downlinks are foreseen as the main sources of RFI in interferometric data because such signals can be visible to many antennas and can perturb the

measurement data for long periods of time. These signals can be detected and flagged by the post-correlation processing pipeline. The CSP will support RFI management downstream through producing power spectra.

However, there are also many signals that are expected to affect science data that have high energy but low time occupancy, such as vehicular radar and aviation transponders [RD10]. These signals are not easily flagged in correlated data. Detecting the signals in a high-time-resolution and antenna-specific domain can potentially improve the robustness of the system against this elusive type of RFI.

Despite the potential benefits, current CSP requirements do not require RFI detection, flagging, or excision at any of the CSP sub-elements. This is due to the lack of enough research data that could be used to develop meaningful requirements. Any prospective RFI mitigation algorithm to be implemented within the CSP signal processing chain must first prove its ability to operate effectively within the ngVLA radio environment, which itself needs an extensive characterization first. Therefore, more research needs to be conducted, and RFI detection and excision capabilities may be included in future revisions of the CSP design. As a side note, current SBP firmware incorporates simple RFI detection algorithms.

4.2.2 Electronic Design

As indicated in the general subsystem overview, the CSP is made of heterogeneous electronics components, each one with its own independent electronic design. The following sections visit each of these designs for the main CSP sub-elements, i.e., the DBE, the SBP, the PSE, and the CSF, while providing a general description of its electronic components. Note that no hardware design activities have been carried out at this stage of the project. The electronic design described for the SBE is based on NRC-developed TALON hardware, which is the chosen solution for future SKA1-Mid correlator and has been proposed for the ALMA correlator upgrade as well.

4.2.2.1 DBE Electronic Design

Each DBE unit, itself heavily shielded against RFI, is further enclosed in a shielded rack in the antenna pedestal according to the specifications in the applicable ICD [AD05]. Inside, the optical streams from all IRD modules feed into quad-channel fiber-optic transceivers hosted on printed circuit boards (PCB). The transceivers utilize four wavelengths on a single fiber, one color per lane. They modulate the optical carrier at 28 GBd using four-level pulse-amplitude modulation (PAM4), achieving a serial rate of 56 Gb/s per lane. Physically, the transceivers conform to the QSFP28 form-factor.

Figure 12 (next page) shows a block diagram of the hardware structure of a DBE module. Data stream processing resides on a high-performance PCB. The main component of the processing PCB is a large FPGA with enough high-speed transceivers for the required I/O capabilities, and enough logical resources to perform all the functional tasks. This includes clock recovery, data stream alignment, digital signal processing, and Ethernet framing. The need for more than one FPGA per DBE unit is still under evaluation and must be verified during the firmware development.

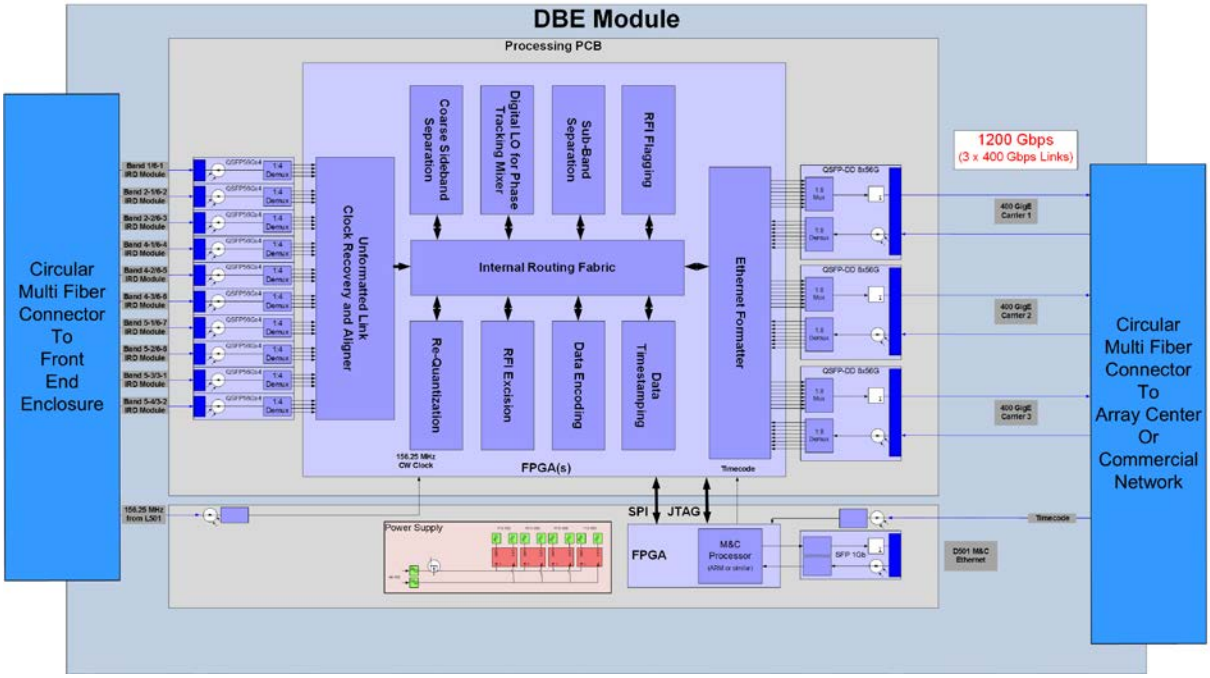

Figure 12: Block diagram of one DBE module. Adapted from [RD10]. Original drawing credit: J. Jackson et al.

Other considerations regarding the number of connected IRD inputs can also affect the design. The processing PCB includes multiple QSFP modules for connecting the IRD modules to the FPGA. In the current concept, 2 DBE modules are needed per antenna, each one connected to 10 IRD modules. The first DBE module connects to Band 6 and Band 3 receivers, while the second DBE module connects to Band 1, Band 2, Band 4, and Band 5 receivers.[5] Three additional QSFP-DD modules are needed to handle the DBE data output. For 8-bit quantization, 20 GHz of dual polarization bandwidth, and a sub-band oversampling factor of 8/7, the output data rate of the DBE is approximately 731 Gb/s. The three 400G data links provide enough room for communication protocol overhead, as well as some extra capacity for increased resolution in case of severe RFI, or for future CSP upgrades (more instantaneous bandwidth). It may be possible that the DBE data output interface must be adapted for the remote antennas connected through commercial fiber.

Each DBE Module will contain a separate PCB to provide all necessary M&C functionality. The M&C board will interface to the antenna M&C Ethernet switch via standard SFP modules. It will contain a high-performance microprocessor as defined in the MCL documentation [AD16]. A high-performance microprocessor is necessary in this application because this board will be tasked with acquiring and analyzing multiple data samples for diagnostic and system setup purposes. It will run an embedded variant of the Linux operating system.

The last PCB in the module will be a high-efficiency power supply that converts the antenna's –48 VDC to the voltages for hardware in this module. It will be consistent with the power supply designs described in the PSU documentation [AD17]. The power supply PCB will be tailored to provide the unique low-voltage, high-current supply voltages needed by the components in this module.

[5] In this scheme, only one DBE module is active at any given time, allowing to put the inactive DBE module in a low-power state. An alternative connection scheme would distribute the computational load between the 2 DBE modules, so that the FPGA hardware resources are minimized. In such a scheme, half the receivers of each RF Band is connected to each DBE module, so that each module only processes half the digitized bandwidth.

4.2.2.2 SBP Electronic Design

The Sub-Band Processor is based on NRC's Frequency Slice Architecture (FSA), which has been chosen for the SKA1-Mid correlator and proposed for the ALMA correlator upgrade. This design is based on high-performance FPGA technology used by NRC's design of the TALON hardware [RD02]. Upgrading current technology is strongly recommended to fully meet ngVLA SBP requirements (see Section 3.4.1).

The SBP consists of a set of SBP Units, each one processing, depending on observing mode:

- Two sub-bands per antenna, for all the antennas, in standard resolution (up to 8-bit quantization) interferometric mode.
- One sub-band per antenna, for all the antennas, in high resolution (up to 16-bit quantization) interferometric mode.
- One sub-band per antenna, for all the antennas, in any of the beamforming modes.

Beamforming modes can use up to 16-bit quantization as they only process one sub-band per antenna. Quantization is mainly constrained by the I/O capacity.

In the current design, each SBP Unit consists of 12 Line Replaceable Units (LRUs), where each LRU includes 2 high-performance FPGAs (14-nm Intel Stratix 10), hence amounting to 24 FPGAs per SBP Unit. A general block diagram of the TALON FPGA board, which the SBP is based on, is shown in Figure 13.

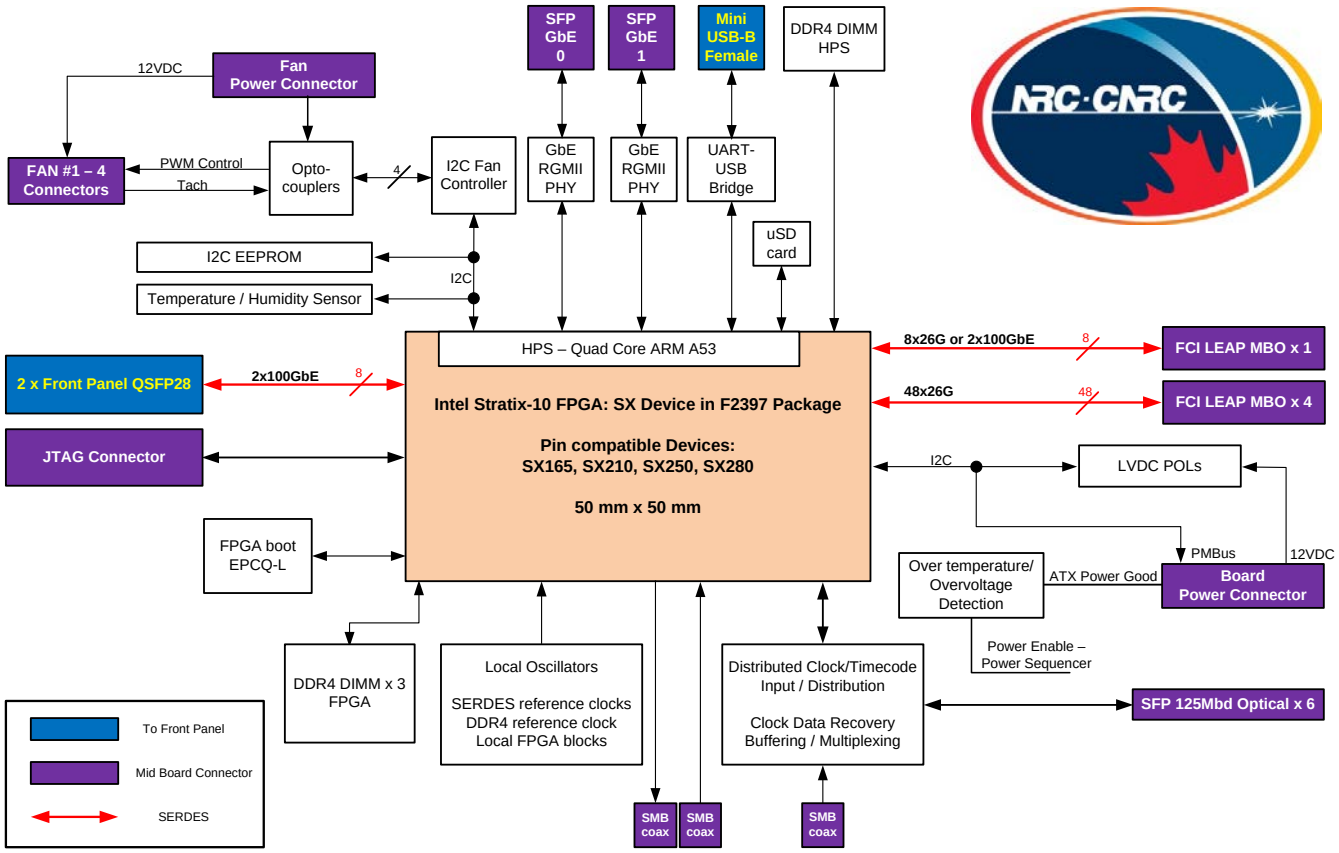

Figure 13: Block diagram of the TALON FPGA board. 24 FPGA boards, arranged in 12 LRUs, compose an SBP unit. Taken from [RD02]. Credit: B. Carlson et al.

Data input and output is through the QSFP modules on the front panel. Each FPGA receives and processes sub-band data of up to 11 antennas, for a maximum of 264 antennas satisfying ngVLA requirements. More antennas would be possible at the cost of additional LRUs per SBP Unit.

Internally, all the FPGA boards connect to each other through the passive optical mesh described in Section 4.2.1.9, which itself connects to each board through multiple Amphenol FCI Leap® On-board Transceiver (OBT) system, as shown in Figure 14. The TALON board includes additional optical interfaces for monitor and control, and for time synchronization.

Figure 14: An On-Board Transceiver system is used to interconnect each SBP unit's FPGA to one another through a passive optical mesh. Taken from fibersystems.com. Credit: Amphenol FSI.

In addition to the FPGA and the optical modules, other components of the FPGA board are DDR4 SDRAM modules for external memory, power system and thermal control components (fans, heat sinks).

4.2.2.3 PSE Electronic Design

The current PSE design consists of a set of LRU "nodes," each of which independently performs the processing for a single sub-band. This processing takes place on an FPGA equipped with High-Bandwidth Memory (HBM) to handle the memory-bandwidth-intensive process of coherent dedispersion; the FPGA host board also contains external memory modules, board management components, an M&C interface, thermal management equipment, and related environmental sensors. Primary I/O connectors are either mounted directly on the FPGA host board or are attached through an industry-standard ANSI/VITI 57.4 FPGA Mezzanine Connector Plus (FMC+).

The PSE nodes connect to the CSP switched fabric (CSF) through a set of 100GbE- or 400GbE-capable connectors on the front panel; these may be QSFP cages with separate COTS optical transceiver modules or a board-integrated optical connector (e.g. MTP or LC), depending on future evaluation of each option's life-cycle cost. Due to the highly frequency-parallel nature of the pulsar processing tasks, dedicated interconnections are not required within the PSE; each node receives beamformed voltage data and outputs the appropriate products for downstream processing independent of any other node.

4.2.2.4 CSF Electronic Design

The CSP switched fabric (CSF) design is based on 400G technology and COTS Ethernet switches to route the high-speed data flowing between the different CSP elements, and the CSP output towards the Online

Sub-Element (ONL) of the CSS. The CSF uses a scalable architecture to meet current and future networking needs of the ngVLA CSP. It relies on commercial high-density non-blocking modular systems consisting of a multiplicity of Ethernet switches, or line cards, such as the one represented in Figure 15.

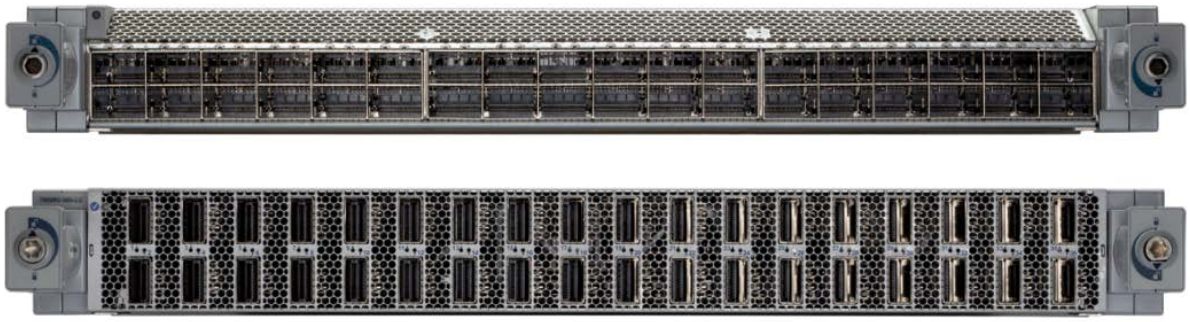

Figure 15: Arista 7800R3 400G line cards. Taken from [RD13]. Credit: Arista Networks, Inc.

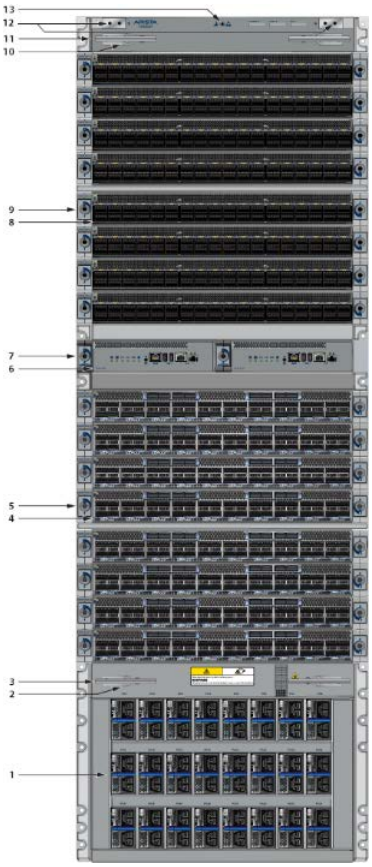

Each of these Ethernet switches is based on ASIC technology for maximum performance. Typically, each line card includes multiple devices that perform as packet processors, with the number of devices depending on the type and number of input ports. For example[6], Arista's Jericho2 IC features 96 50G PAM4 transceivers that process 12 400G I/O ports, resulting in 4.8 Tb/s of processing capabilities [RD14]. These devices are connected to each other through internal mid- and backplanes in a mesh topology, to scale the switching capabilities to the desired size. In addition to the 96 transceivers serving the front panel, the Jericho2 IC incorporates 112 transceivers for internal connectivity, as well as 8 GB of on-chip HBM memory for data buffering.

To meet CSP connectivity requirements, the network switches are commonly mounted on a special chassis. Chassis design and layout are a key aspect that enables high performance. Some fabric modules that interconnect the installed line card modules are directly behind them and oriented orthogonally [RD14]. This design reduces trace lengths between system elements and enables high-speed signals to operate more efficiently with high signal integrity thanks to the shorter lengths. It eliminates the need for additional passive mid- or backplane, allowing the direct connection of line cards to fabric modules improving airflow. Figure 16 shows an example of a fully populated chassis. The resulting system provides full matrix switch capabilities with 576 400G ports, by means of 16 line cards with 36 400G ports each, for an overall maximum full-duplex throughput of 460 Tb/s (230 Tb/s transmit and 230 Tb/s receive).

Figure 16: Arista DCS-7816 chassis (fully populated). Credit: Arista Networks, Inc. 1. Power supplies. 2. Line card and Supervisor extraction tool tether. 3. Extraction tool. 4. Line cards. 5. Line card lock. 6. Supervisor modules. 7. Supervisor lock. 8. Line cards. 9. Line card lock. 10. Line card and Supervisor extraction tool tether. 11. Extraction tool. 12. Grounding locations. 13. ESD attach point. Taken from www.arista.com. Credit: Arista Networks, Inc.

[6] This CSP design document uses commercial networking solutions from Arista Networks, Inc. as an example of technological capabilities. This should not be interpreted as the result of a proper vendor selection, which will be conducted in future stages of the design.

4.2.3 Mechanical Design

This section provides a general overview of the mechanical design of the main components of the CSP.

4.2.3.1 DBE Mechanical Design

Each DBE Module consists of high-speed digital components, a high-performance microprocessor, and switching power supplies. All these devices will contribute to high levels of RFI that are in-band to the ngVLA receiver band and must be attenuated.

Figure 17 shows a prototype module design based on the Advanced RFI Containment System (ARCS) module system designed for the ngVLA project [RD10]. This enclosure is heavily shielded to prevent interference with the telescope from the electronics inside. The DBE PCBs will attach to heatsinks inside this module. Two multi-fiber connectors and a filtered Cannon power connector allow signaling and power across the module's RFI shielding.

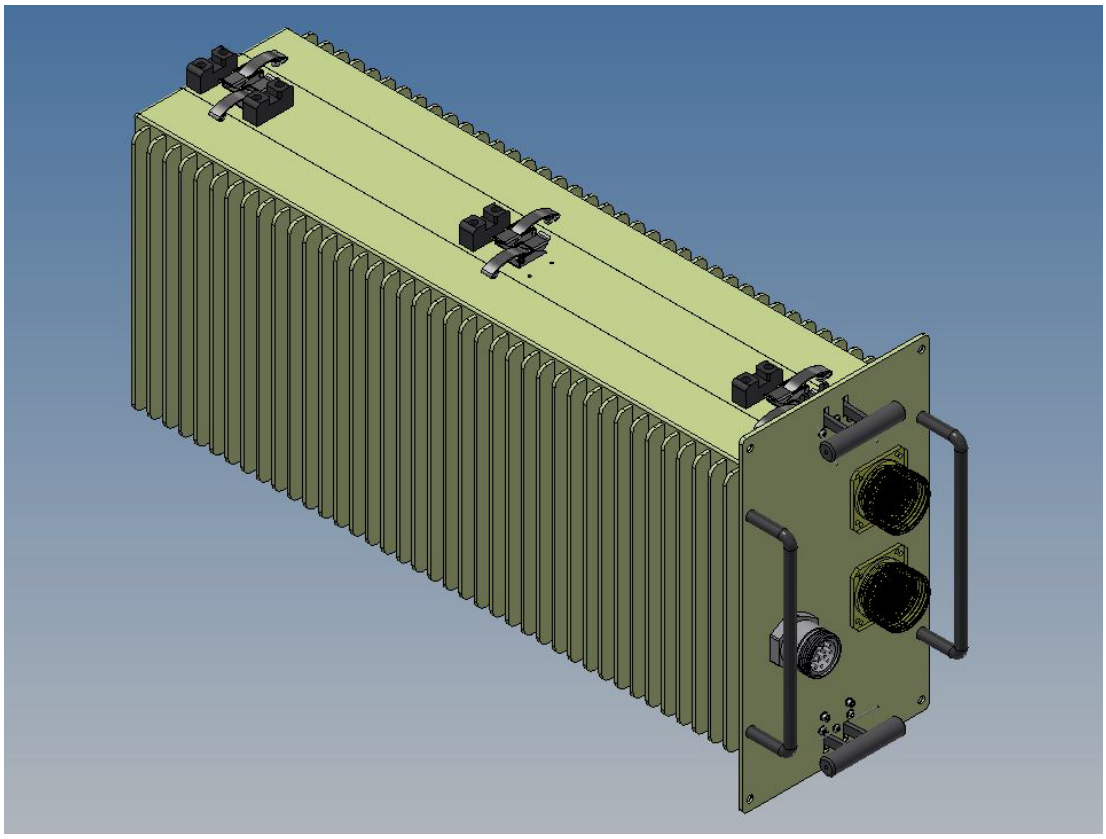

Figure 17: 3D rendering of the Advanced RFI Containment System module. Taken from [RD10]. Credit: J. Jackson et al.

Module dimensions and aspect ratio are for demonstration purposes and will depend on the space available in the pedestal area enclosure available in the selected antenna design. In the current design concept, two DBE modules are connected to all the IRD modules. A very preliminary power consumption estimate would be about 300W per DBE Module, 600W per antenna unit.

4.2.3.2 SBP Mechanical Design

The SBP design is based on NRC's TALON FPGA board design. A picture of an early prototype board is shown in Figure 18. This board will be upgraded to a newer FPGA technology node.

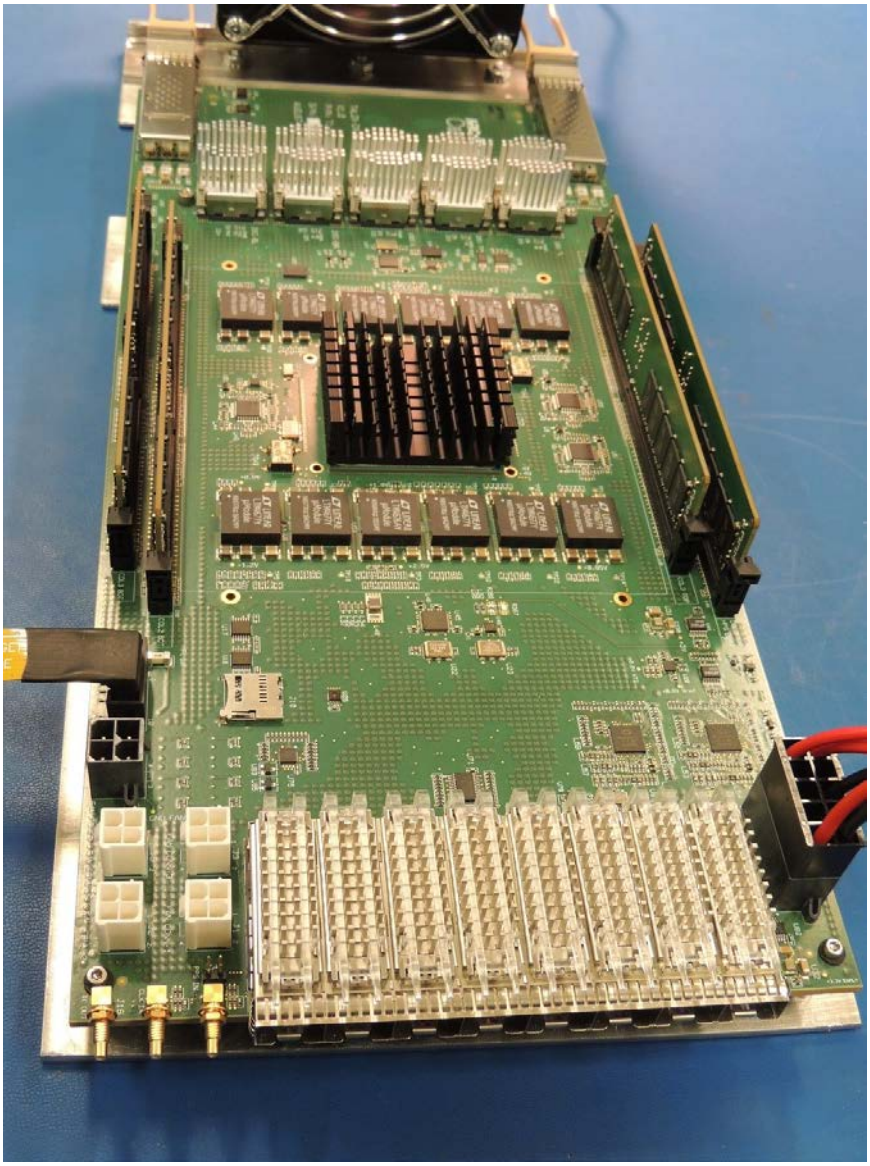

Figure 18: TALON-DX prototype board under test. Taken from [RD02]. Credit: B. Carlson et al.

Each two of these boards are installed inside an LRU cage designed to fit into any standard 19-inch rack. A picture of a 2U prototype LRU is shown in Figure 19 (next page). As can be seen, the tall FPGA heat sinks dominate the LRU volume. These are required as each TALON board power consumption estimates based on extensive testing for SKA1-Mid are about 300 W. Some power reductions down to 200-250 W can be expected after a technology upgrade.

The proposed cooling system for the SBP design is based on Direct Contact Liquid Cooling (DCLC). When using DCDL, a cooling plate through which cooling water-based fluid circulates is thermally attached to the main heat generating devices on the board, i.e., the FPGA, power supplies, optical transceivers, and even the DIMM memory modules if needed. A Coolant Distribution Unit (CDU) installed at the top of

each rack performs the heat exchange between the rack liquid loop and the facility (NSB) liquid loop. An example of such a CDU is shown in Figure 20.

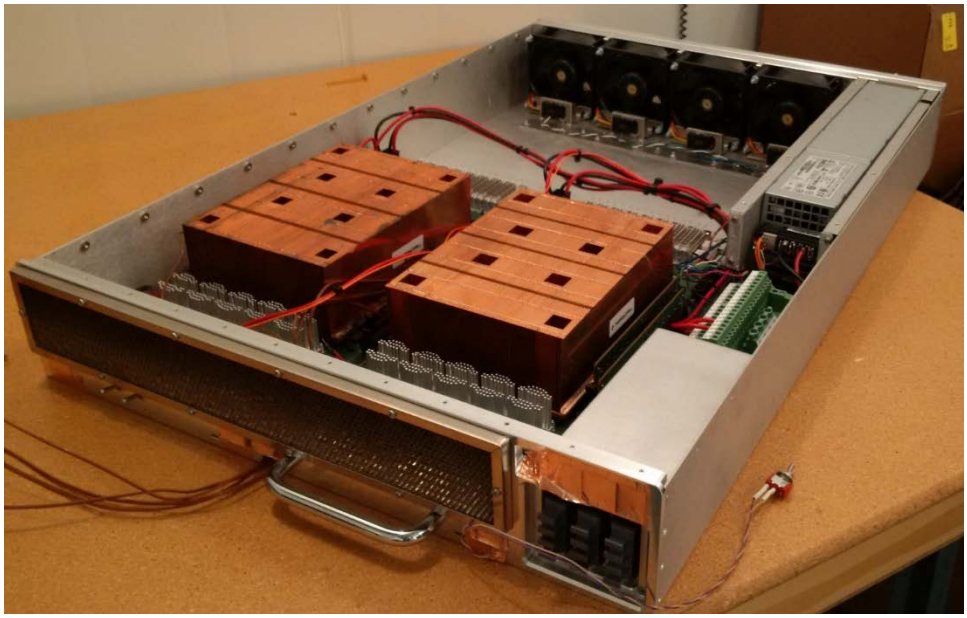

Figure 19: 2U air-cooled TALON LRU prototype under test. The large orange boxes are copper stacked-fin heatsinks for air cooling. These would be replaced by liquid cooling plates for a liquid-cooled 1U DCLC design. Taken from [RD02]. Credit: B. Carlson et al.

Figure 20: An instance of a 40-kW DLDC heat exchanger 1U DCLC design. Taken from [RD02]. Credit: B. Carlson et al.

By using DCLC, the physical size of the SBP can be reduced by about a factor of two, as each FPGA LRU can now be installed into a 1U line, as opposed to the 2U line shown in Figure 19. Assuming standard 42U 19-inch electronics racks, a tentative rack distribution is shown in Figure 21 (next page). Within each rack, 2 SBP Units (or Frequency Slice Processor Units, using FSA terminology [RD02]) can be installed, each one taking 13U height: 12U for the 12 FPGA LRUs, and the additional 1U for the passive fiber optic mesh to interconnect them. In addition to the SBP Units, 2U are needed for the CDU, 1U for an M&C Ethernet switch, and 1U for the timecode distribution optical circuit. At 600 W per LRU, the total power consumption of both SBP Units is about 15kW. Hence, a 40-kW capacity heat exchanger should suffice

for the whole rack. The remaining 12U are spare and can be used to install either PSE or CSF electronics. If needed, a higher capacity heat exchanger can be installed, after the final rack configuration is fully defined.

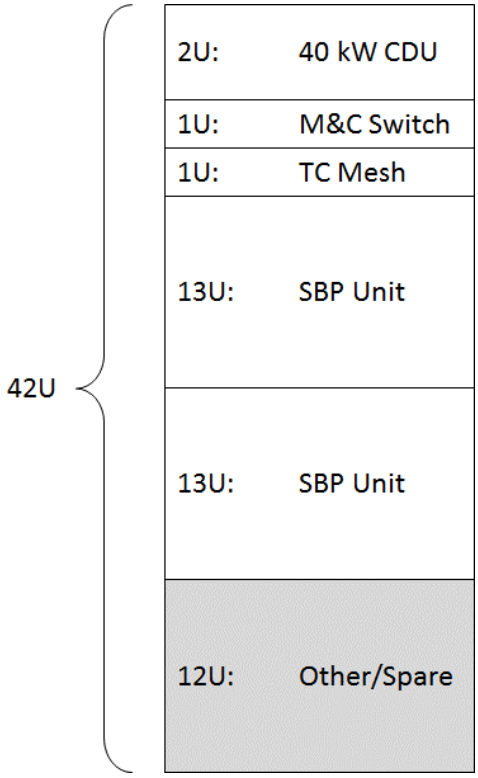

Figure 21: Tentative SBP rack distribution within a 42U rack. The 12U spare can be used to accommodate PSE equipment.

The number of SBP units that needed to satisfy ngVLA requirements has been thoroughly studied in different iterations of the SBP design [RD02], [RD03], [RD04]. As per CSP requirements, the current CSP design criterion is to install as many SBP units as needed to process 20 GHz of bandwidth in (standard precision) Interferometric OM, simultaneous to 8 GHz of bandwidth, overall, in any combination of beamforming OMs, i.e., Pulsar Timing OM, Transient Search OM, or VLBI OM.

Since two sub-bands per antenna are processed by each SBP unit in Interferometric OM, each sub-band carrying 218.75 MHz of bandwidth (Section 4.2.1.3), a minimum of 46 SBP units are needed for processing 20 GHz of bandwidth in Interferometric OM. Because the TALON-based SBP units can simultaneously operate in a single OM only, additional SBP units must be installed to provide the CSP with simultaneous beamforming capabilities. This contrasts to a SCREAM based design, in which the OM is independently defined on a per-sub-band, per-antenna basis (Section 5.6). Hence, no additional SBP units are needed, which results into a huge reduction in hardware. For the TALON-based SBP design, as one sub-band per antenna is processed in all beamforming modes, 37 additional SBP units must be procured for 8 GHz of beamforming bandwidth. Overall, a TALON-based SBP design needs at least 83 SBP units to satisfy CSP subarray operation requirements [AD01]. This is a significant reduction over the reference design [RD01], which employed 150 SBP units, thanks to the increase in the sub-band bandwidth and, most importantly, the processing of two sub-bands in Interferometric OM.

In summary, the SBP sub-element of the CSP will need 42 standard 42U 19-inch racks. There should be enough room in those same racks to install the PSE as well (more than 500U in spare space across all the racks). The total estimated power consumption of the SBP is estimated in approximately 600 kW, which does not account for other ancillary subsystems such as cooling and local M&C.

4.2.3.3 PSE Mechanical Design

The 50 nodes of the PSE will occupy, depending on the final node density, between 25 and 100 vertical units of space ("U") in standard 19-inch equipment racks; these need not necessarily be contiguous. These racks are assumed to be of a four-post design, with appropriate provisions for cable routing. The PSE LRUs will mount to sliding rails to permit easy installation and removal.

Each PSE LRU will connect to a system circulating liquid coolant through a pair of quick-connect couplers, to provide direct cooling to the FPGA. The coolant pressure must be sufficient to ensure a flow rate capable of maintaining the desired operating temperature of the LRU when operating at full capacity, but not high enough to endanger the equipment. Forced-air ventilation will aid in cooling those components of the LRU which are not in thermal contact with the liquid coolant.

The power consumption of a single PSE node is estimated at 250W with current technologies, with some potential for further reduction with future technological developments; the PSE as a whole will therefore consume approximately 12.5 kW of electrical power, exclusive of cooling.

4.2.3.4 CSF Mechanical Design

To meet CSP connectivity needs, the CSF must provide around 2000 400G ports. Assuming a high-density organization, with at least 36 ports per individual switch, or line cards, the CSF would be populated with 64 line cards offering 2304 ports. This would be enough for the baseline ngVLA design needs, but additionally it would incorporate enough spare ports to connect custom commensal back ends or additional antenna inputs for VLBI real time correlation, such as the GBO.

The CSF uses special chassis for internal midplane and backplane connectivity. For example, taking the Arista 7800R3 as a reference [RD14], the 64 line cards could be installed in four DSC-7816 chassis, with overall dimensions 60H x 70W x 40D inches (about four 31U 19-inch racks). Besides the line cards, the chassis provides housing for 8 supervisory modules, 48 internal fabric modules, and 96 power supplies. Figure 22 shows a conceptual image of the CSF.

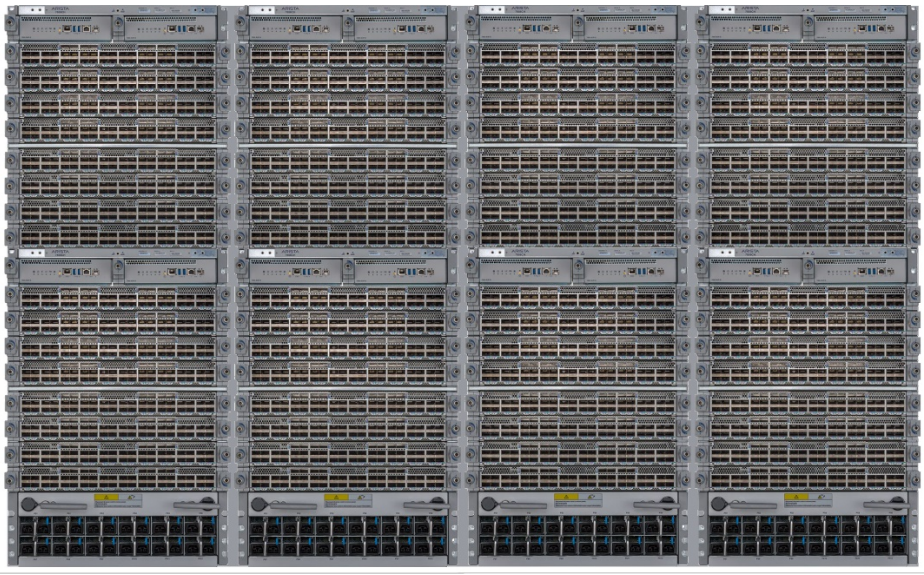

Figure 22: Rendered image of the CSF concept. A single switched matrix network connects all CSP elements based on 400G technology. 64 line cards are distributed in 4 columns with 16 line cards each. Supervisory and fabric modules are included as well. Adapted from www.arista.com. Credit: Arista Networks, Inc.

Being COTS modules, the CSF chassis and modules use air cooling with front-to-rear air flow. The estimated maximum power consumption would be around 160 kW, at 40 kW per line card stack. When

added to the other CSP sub-elements, the overall CSP power consumption yields slightly less than 1 MW, as summarized in Table 2.

Location	CSP Unit	No. of Units	Power/Unit	Total Power
Antenna	DBE Unit	263	600 W	157.8 kW
			Total Antenna	157.8 kW
Central Building	SBP Unit	83	7200 W	597.6 kW
	PSE Node	50	250 W	12.5 kW
	CSF Line Card	64	2500 W	160.0 kW
			Total Central Building	770.1 kW
			Total CSP	***927.9 kW***

Table 2: Preliminary CSP power consumption budget.

4.3 Performance Budgets

Most of the CSP functionality is well constrained within one of its processing subsystems, i.e., the DBE, the SBP, or the PSE. Nonetheless, there are two notable exceptions: the correlation loss, sometimes referred to as the digital efficiency of the correlator, and the beamforming efficiency. Note that the PSE does not contribute to either of these, as both beamforming and correlation occur prior to the PSE in the signal processing chain. In the following, a performance budget for both figures of merit is developed to help develop Level 3 requirements of the DBE, as well as the SBP.

4.3.1 Correlation Loss

CSP requirement CSP0009 states that the correlation loss attributable to the CSP shall be less than 1% [AD01]. This does not include the loss due to digitizer quantization. Since most computations carried out by the CSP use fixed-point arithmetic, they incur a correlation loss through the addition of quantization noise.

Table 3 summarizes the preliminary correlation loss budget within the CSP to meet CSP0009. This budget may be revised in the future during the firmware development phase.

Functional Operation	CSP Subsystem	Loss
Digital Sideband Separation	DBE	0.1%
Baseband LO Mixer	DBE	0.1%
Sub-Band Generation	DBE	0.1%
Sub-Band Requantization	DBE	0.1%
	Total DBE	0.4%
Resampling	SBP	0.1%
Phase Correction	SBP	0.1%
Spectral Zoom	SBP	0.1%
Frequency Channelization	SBP	0.1%
Corner Turner	SBP	0.1%
Correlator	SBP	0.05%
Output Format	SBP	0.05%
	Total SBP	0.6%
	Total Reserved	0.0%
	Total CSP	**1.0%**

Table 3: CSP correlation loss budget.

The intention is not to establish hard constraints that could end up driving the design, but to provide some preliminary figures that could be used as a reference during the design. It is understood that the target performance can be readily satisfied without too much complication. The preliminary CSP correlation loss budget allocates 0.4% loss to the DBE and 0.6% loss to the SBP, with no room for additional losses. Hence, any deviation from these number should be coordinated across subsystems. The correlation loss must be computed against an ideal system for the same inputs. For example, such inputs include, but are not limited to, any set of parameters generated externally to the CSP. It does not include, however, any quantization or transformation of such input parameters performed by the CSP. In addition, the correlation loss must be computed in the relevant case of weak signal.

At the DBE, the filter coefficients, sine/cosine look-up tables, and other fixed stored parameters are all quantized and hence introduced some correlation loss. In additional, multiple truncations of the data word length need to happen internally to keep the number of bits under control. In principle, at least 16-bit precision should be possible, which should produce small correlation losses for normal input levels. The loss attributable to sub-band requantization assumes standard precision, i.e., 8-bit quantization, and optimal sub-band gain adjustment. Obviously, the correlation loss target cannot be guaranteed in presence of strong RFI (without imposing some constrains on the RFI) or for sub-bands requantized with a smaller number of bits, for instance, when the output data rate has to be reduced.

Similarly at the SBP, there filter coefficients and other fixed parameters that must be quantized, but also some other input parameters such as phase/delay model parameters or calibration parameters are used to compute corrections that need to be quantized before their application. The quantization at the corner turner prior to the correlator employs at least the 8-bit standard resolution, which should suffice to achieve the allocated correlation loss. Both the correlator and the conversion to an output format, probably floating point, are expected to meet a smaller loss target without much difficulty.

4.3.2 Beamforming Efficiency

CSP requirement CSP0032 states that the SNR loss attributable to the CSP in beamforming OMs shall be less than 5% [AD01]. Similarly to the correlation loss budget, most computations carried out by the CSP use finite precision arithmetic and hence produce an SNR loss in the form of added quantization noise. Table 4 shows a preliminary SNR loss budget for the CSP when operating in beamforming OMs.

Functional Operation	CSP Subsystem	Loss
Digital Sideband Separation	DBE	0.1%
Baseband LO Mixer	DBE	0.1%
Sub-Band Generation	DBE	0.1%
Sub-Band Requantization	DBE	0.1%
	Total DBE	0.4%
Resampling	SBP	0.1%
Phase Correction	SBP	0.1%
Frequency Channelization	SBP	0.1%
Polarization Correction	SBP	0.1%
Beamforming	SBP	0.1%
Output Format	SBP	0.1%
	Total SBP	0.6%
	Total Reserved	4.0%
	Total CSP	**5.0%**

Table 4: CSP beamforming SNR loss budget.

The CSP SNR must be measured at the SBP output and compared against what an ideal system with the same inputs would produce. The input parameters are inclusive of, but not limited to, phase and delay models, beamforming weights, and polarization correction parameters, among others. The signal processing chain at the DBE does not change from Interferometry OM to beamforming modes. Hence, the more restrictive interferometric requirements developed for correlation loss apply in this case as well.

Something similar occurs with the SNR loss caused by the SBP, where the more stringent interferometric requirements apply as well, under the assumption that the same DSP building blocks are used in both interferometric and beamforming firmware development. The same level of accuracy should be easily attained in beamforming exclusive tasks, such as polarization correction and the beamforming or linear combination data streams from different antennas. The output formatting loss has been relaxed assuming 8-bit quantization is used at the beamformer output.

As a result, 0.4% SNR loss is allocated to the DBE, 0.6% to the SBP, and 4.0% is reserved for future use, since only the additive noise has been considered in Table 4. For example, the additional SNR loss due to the phase noise introduced by the CSP, which may not be attributable to a specific DSP task, would make use of the extra 4.0% headroom. This budget may be reviewed as needed during the firmware development phase.

4.4 Environmental Protection

All CSP equipment shall be installed in environmentally controlled facilities or racks. Normal local dust contamination levels shall be assumed in the design. Racks supporting the CSP electronics shall be compliant with seismic protection requirements. For the equipment installed at the antenna, additional vibration protection requirements shall be considered part of normal operations. All CSP equipment shall be affixed to its respective (LRU) enclosures, which provide it with needed ESD protection. Fastening shall be compliant with normal operating and transportation conditions.

CSP equipment uses both air cooling and DCLC. The applicable ICDs will define the maximum (room) temperature at which CSP equipment shall operate under specifications. The hosting facility is responsible for monitoring and maintaining operational environmental conditions. DCLC is only used in the central processing building. The CSP room shall be designed and constructed to withstand low-probability coolant leaks on the order of several gallons (specific volume will be specified in the corresponding ICD [AD08]).

4.5 RFI, EMC, and Lightning Protection

RFI shielding to the CSP equipment will be provided by the hosting facility. Special care shall be put in the design of the DBE, which shall be installed at the antenna, by designing an enclosure that shall provide any additional RFI shielding as specified in the pertinent ICD [AD05]. All CSP equipment shall comply with applicable EMC industry standards. Lightning protection shall be provided by the facility at which the CSP equipment is installed.

4.6 Power Supply and Distribution

Each DBE Unit receives its power from the antenna PSU, which supplies –48 VDC to its high efficiency power supply. All components of the SBP, the PSE and the CSF (hosted by the central processing building) are mounted on standard 19-inch racks, each one including a COTS AC–AC Power Distribution Unit (PDU) at the back of the rack. Distribution of power to the racks is performed by the NSB facility, according to the specifications in [AD08].

4.7 Reliability, Availability, and Maintainability

A preliminary reliability (FMECA) analysis of the CSP shall be carried out in the CSP PDR phase, and updated and finalized by CDR. As side node, the reliability of the TALON-based SKA1-Mid correlator and beamformer has been formally assessed. The study reports an MTBF of the TALON-DX assembled board of 5 years, calculated using the parts-count method in accordance with MIL-HDBK-217FN2 [RD02].

Similarly, the availability of the CSP shall be formally assessed in the same fashion. The reported results of the study will be fed back into the design to meet CSP availability requirements, developed according to the system availability budget [AD19]. A TALON-based SBP design would exhibit an inherent availability greater than 99.9% [RD02], assuming a single TALON LRU fails at any given time. Additionally, the FPGA on the TALON board contain integral single SEU detection and correction, as well as multiple SEU detection. Similar capabilities will be pursued in all FPGA devices used in the CSP. Given the size of the devices, the number in the system, and the ngVLA site elevation requirements of 2500 m, SEUs due to cosmic rays will occur on a regular basis, sometimes needing FPGA bitstream reboots to correct.

Maintainability of the CSP will be formally analyzed in the PDR and CDR phases. As per CSP requirements, the CSP design is modularized into Line Replaceable Units (LRUs) to facilitate site maintenance. In addition, all CSP sub-elements incorporate self-diagnosis functions to identify faults based on recorded monitor data. The firmware and software of any individual LRU can be updated remotely via the appropriate monitor and control network connection, as defined in [AD06] and [AD09].

Table 5 summarizes other important maintainability aspects of each item in the CSP section of the PBS.

CI number	Item	Maintainability
020.30.25.10.00	D501 DBE Module	This item shall be replaceable on-site. Replacing it requires powering down the rack.
020.40.30.10.00	SBP Line Type #1 (FSA-based)	TALON-based LRUs are hot swappable. Their power supply is also hot-swap replaceable without removing the LRU from the rack. The LRU has dry quick connect/disconnect DCLC cooling lines at the rear-of-rack manifold [RD02].
020.40.30.40.00	SBP Local Monitor and Control Line	This item shall be hot swappable.
020.40.30.50.00	SBP Local Network Line Type #1	This item is passive and should not fail. Replacing it takes down the entire SBP Unit. [RD02]
020.40.30.60.00	SBP Cooling System Line	This item includes dry quick connect/disconnect fittings. Replacement is quick, typically 30 min. It includes redundant circulating pumps and network monitoring to allow replacement before failure [RD02].
020.40.50.10.00	PSE Sub-band Processor Line	This item shall be hot swappable.
020.40.50.20.00	PSE Local Monitor and Control Line	This item shall be hot swappable.
020.40.70.10.00	CSF Chassis	These items are hot swappable components, with redundant supervisor, power, fabric, and cooling modules.
020.40.70.20.00	CSF Line	

Table 5: Maintainability of the main CSP items in the PBS.

If necessary, the SBP and the PSE may include a redundant unit to allow cyclical testing of individual units without a negative impact on the availability of the system.

4.8 Manufacturability

Many CSP items are COTS available except for LRUs containing custom PCBs, e.g., those hosting FPGA devices. The complete list of non-COTS items is:

- TALON-based LRU
- DBE Module, which includes three custom PCBs, as described in Section
- PSE processing LRU

For the TALON-based LRU, a COTS procurement model is applicable as follows [RD02]:

1. 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 it 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.
2. 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.

A similar COTS procurement model will be used for the other non-COTS items in the list above.

4.9 Safety Analysis

All CSP equipment shall meet industry standards for electrical safety. All CSP COTS components shall be certified to all relevant safety standards. A formal hazard analysis according to the standards and procedures defined in [AD18] shall be done and included in the CSP documentation for PDR.

4.10 Technology Readiness Assessment

Table 6 provides a technology readiness assessment for the main CSP components captured in the product breakdown structure included in Section 4.1.2.

CI number	Item	TRL	Motivation
020.30.25.10.00	D501 DBE Module	3	This item is based on technology that has been proven in other applications. The concept is fully developed and ready for initial development.
020.40.30.10.00	SBP Line Type #1 (FSA-based)	7	The TALON-based hardware has been extensively tested and integrated with other subsystems in relevant simulated environments the frame of SKA1-Mid telescope. It is ready to be tested in operational environments.
020.40.30.40.00	SBP Local Monitor and Control Line	N/A	This is a COTS item.
020.40.30.50.00	SBP Local Network Line Type #1	N/A	This is a COTS item.
020.40.30.60.00	SBP Cooling System Line	N/A	This is a COTS item.
020.40.50.10.00	PSE Sub-Band Processor Line	3	This item is based on technology that has been proven in other applications. The concept is fully developed, and ready for initial development.

CI number	Item	TRL	Motivation
020.40.50.20.00	PSE Local Monitor and Control Line	N/A	This is a COTS item.
020.40.70.10.00	CSF Chassis	N/A	This is a COTS item.
020.40.70.20.00	CSF Line	N/A	This is a COTS item.

Table 6: Technology readiness assessment for the main CSP items in the PBS.

5 Appendix A: Trade Studies

5.1 Digital Sideband Separation

One of the first decisions in the design of the signal chain is where to apply the calibration coefficients that correct the passband ripple and bring the sideband rejection to the level in System Requirements [AD14]. There are basically three options:

1. Apply the correction on the digitizer data stream at the DBE
2. On the sub-band data stream, at the SBP (time domain) or at the DBE (frequency domain)
3. On the frequency channel data stream, at the SBP

In principle, the best accuracy can be achieved if the correction is applied after fine frequency channelization at the SBP. However, this turns problematic in some CSP observing modes for which the computational capabilities required for the CSP only allow processing one sub-band. Without the image sub-band, the compensation for sideband separation is not possible.

The second alternative, applying calibration coefficients on the sub-band data stream, consists of two different alternatives: in the time domain or in the frequency domain. In the frequency domain, it would be very simple by combining at the DBE each pair of sub-bands to generate calibrated ones. However, given the coarse frequency resolution of the sub-band channelizer, and the frequency shifter prior to the sub-band channelizer, it is not clear that the achievable accuracy can provide the necessary rejection level.

Another possibility is performing the sideband separation through time-domain filters on the sub-band data streams. Doing this at the DBE is quite complex, because of the number of filters needed. A better option would be to implement the filters at the SBP, but again this is not possible unless both sub-bands are processed at the SBP.

The best trade-off between complexity and accuracy can be achieved through FIR filters acting on the input data streams. This solution has the advantage of being implemented before the baseband digital LO mixer, so the accuracy is only limited by the filter approximation error. For all the above reasons, this is the solution chosen for the design.

5.2 Number of Sub-Band Channels per Baseband

One key parameter of the CSP design is the number of sub-bands generated from each receiver digital output, which consists of IQ samples at a rate of 7 GS/s. The number of sub-bands determines the sub-band bandwidth, which itself determines the minimum computational capabilities required from each SBP unit, which must process at least one sub-band in all observing modes.

The reference design used a sub-band bandwidth of 200 MHz per sub-band, which results in 35 sub-bands for the two-sided receiver bandwidth. However, it is more convenient from the implementation point of view of the PFB that the number of sub-band be a power of 2. Hence, the current CSP design employs 32 sub-bands per receiver, yielding 218.75 MHz of bandwidth for each sub-band. This increase assumes that a future upgrade of the TALON hardware will have no problem in handling a slight increase in the input data rate.

The different options for the number of sub-bands, their bandwidth, and the implications on the CSP design have been studied in the technical reports [RD01], [RD03], [RD04]. The chosen solution is the result of an iterative process after taking into considering many factors.

5.3 Sub-Band Sampling Frequency

Once the number of sub-bands is determined, the most convenient oversampling factor from an implementation point of view is to set it as N/M, where N is the number of sub-bands, and $M < N$ is any integer greater than zero. The rationale can be found in [RD04] and assumes a PFB implementation of the sub-band channelizer.

Given 32 sub-bands, M must be as close as possible to N, but still give some margin to ease the filter bank design. A good trade-off between the increased data rate and the length of the filter response is given for oversampling factor more than 110%, which is first achieved by $M = 29$. The resulting sampling frequency becomes approximately 241,379,310.345 Hz.

For timing purposes, a more convenient sub-band sampling frequency is such that an integer number of samples are produced every second. This is achieved by the next possible value of $M = 28$, which is hence the chosen solution. As a result, the sub-band sampling frequency is now exactly 250 MHz, and the oversampling factor 8/7.

Another factor considered in the selection of the oversampling factor is that the resulting sub-band sampling frequency results slightly smaller than a standard VLBI sampling frequency, 256 MHz in this case. The additional benefit from this is that the SBP can upsample the sub-band data stream to this standard frequency, which not only facilitates the implementation of the VLBI Observing Modes, but also maximizes the true information bandwidth in this OM.

5.4 CSF Architecture

Different architectures have been studied for the ngVLA CSP Switched Fabric. The following sections describe three different alternatives: a single crossbar switch, partitioned in the bandwidth domain, and partitioned in the antenna domain. The single crossbar switch is the preferred solution for its simplicity and optimization of resources. However, current COTS technology cannot satisfy the connectivity requirements required by a single switched matrix for the ngVLA CSP, although it is very likely that these needs will be fulfilled in time for the production phase of the CSP. Thus, the partitioned CSF architectures are considered only as a contingency plan in case the preferred choice is not available.

In this trade study, 400G technology is assumed for the CSF. This technology is already commercially available and, although it will be replaced in the short term by 800G as the most advanced COTS option, that should be deemed as an opportunity more than a risk, thanks to the probable cost decrease it will bring. In any case, the employed technology only changes the quantitative results of the study, but it remains essentially the same for the qualitative ones.

5.4.1 Single Switch Matrix

In this architecture, there a single switch to which all the CSP elements connect to, as well as the Online Sub-Element (ONL) of the Computing and Software System (CSS) for receiving the data products generated by the CSP. Thus, the used network topology is a star network, which allows full flexibility in bandwidth allocation between every pair of connected elements. This is the preferred architecture as it minimizes the number of line cards (individual switches) needed, and hence the size, power, and cost of the CSF.

The number of 400G connections (ports) required by the CSF are as follows. For the DBE subsystem, at three 400G connections per antenna, and 263 antennas, 789 400G ports are needed for the whole array. Regarding the SBP subsystem, each TALON board receives 11 sub-bands (high-resolution quantization) from 11 antennas. At a sub-band sampling rate of 250 MHz (see Section 5.3 for details), each TALON board receives 176 Gbps from the CSF. This may allow sharing a single 400G port between the two TALON boards in each SBP LRU by using breakout cables. Hence, at 12 LRUs per SBP Unit, and a CSP

populated with 83 SBP Units (see Section 4.2.3.2), the number of CSF ports needed to fully connect the SBP is 996 400G ports. Concerning the PSE, a conservative estimate would be that each of its 50 nodes would need one full 400G port from the CSF. Potentially, just three 400G ports would be needed for connecting the ONL and the CSP, as the output data rate from the CSP shall not exceed 132 GB [AD14]. As a result, the minimum number of CSF ports amounts to 1838.

The proposed solution entails populating the CSF with 64 line cards, each one containing 36 400G ports, for a total of 2304 ports. That would be sufficient to serve all the CSP communication needs, as well as providing enough spare ports to allow additional antenna inputs in the future (e.g., GBO), or additional custom back ends in the future (for example, a SETI back end). 64 line cards may be installed in four vertical stacks of 16 lines each within a single chassis, and consume around 100 kW, depending on the final solution selected. Technological advances currently being prototypes, such as 112G transceivers, may be available in the near future for higher port density and would allow reducing the CSF size to half.

5.4.2 Bandwidth Domain Partition

Current COTS solutions cannot meet the CSF requirements as described in the single switch matrix architecture, although this is likely to change in the near future before the ngVLA enters its production phase.

A CSF partitioned in the bandwidth domain is presented next as a contingency plan. This solution is less capable than the fully connected matrix and makes the design of the DBE more complex. In this CSF architecture, the CSF consists of three independent matrix switches, each one processing one-third of the bandwidth for the whole array. This architecture is reminiscent of the TRIDENT architecture of the CSP reference design [RD02], with the main difference being the use of active switches to connect the DBE, the SBP, and the PSE subsystems, instead of passive fiber optic circuits that scale poorly with more systems added to the network.

A threefold partition is used because the DBE is not partitioned in the bandwidth domain but in the antenna domain (one DBE unit processes all the bandwidth for one antenna), and this partition nicely fits the number of 400G data links coming out from the DBE. Therefore, each of these DBE output links is connected to each of the three CSF networks. An internal crossbar switch, possibly within the FPGA firmware, allows routing any of the sub-band data streams generated by the DBE to any of the CSF networks, and thus, to any of the SBP Units. This scheme can become more complex if separate FPGA devices are required per DBE Unit.

In contrast to the DBE, both the SBP and the PSE subsystems are naturally partitioned in the bandwidth domain, with each SBP unit or PSE node processing a portion of the bandwidth, typically one or two sub-bands, for the whole array. Hence, simply one-third of the SBP units and PSE nodes connect to each of the three CSF networks.

As regards the ONL subsystem, only a few SBP units can potentially generate enough data to saturate the ONL input data link. Thus, full connectivity capabilities are needed from each of the CSF networks towards the ONL subsystem.

As a result, the size of each individual CSF network becomes as follows. For the DBE, one 400G port is required per unit, for a total of 263 ports. For the SBP, 12 LRUs per SBP unit and 28 SBP units (one-third) make a total of 336 ports. 17 additional ports are required for one third of the PSE nodes, and three more ports for the ONL. Overall, at least 619 400G ports are required in each of the three individual CSF networks. This is roughly one-third of the ports required in the single switch matrix architecture, which is obtained at the cost of routing flexibility and DBE complexity.

5.4.3 Antenna Domain Partition

The partition of the CSF in the bandwidth domain alone would not be enough for enabling current COTS solutions. However, the risk that this situation extends to the ngVLA production phase is very low. For example, just considering the upgrade from 56 Gb/s per lane transceiver technology to 112 Gb/s per lane, already available in some commercial devices, would allow doubling the current maximum number of 400G ports.

Partitioning the CSF in the antenna is not desirable because it increases the overall number of ports required. This increase only affects the PSE, whose processing nodes must now connect to all the CSF subnetworks serving different sets of antennas. Splitting the CSF in such a way that one subnetwork connects the DBE to the SBP, while another subnetwork connects the SBP, the PSE, and the ONL to each other, is even worse as the number of ports doubled in this case are those of the SBP, much more numerous than the PSE's.

In order to keep the increase of overall ports to a minimum, we will just consider the previous case of a threefold partition of the CSF in the bandwidth domain, followed by a twofold partition of each of subnetwork in the antenna domain. The result is six independent CSF networks, each distributing the data corresponding to one-third of the processed bandwidth and half the antennas, as shown in Figure 23.

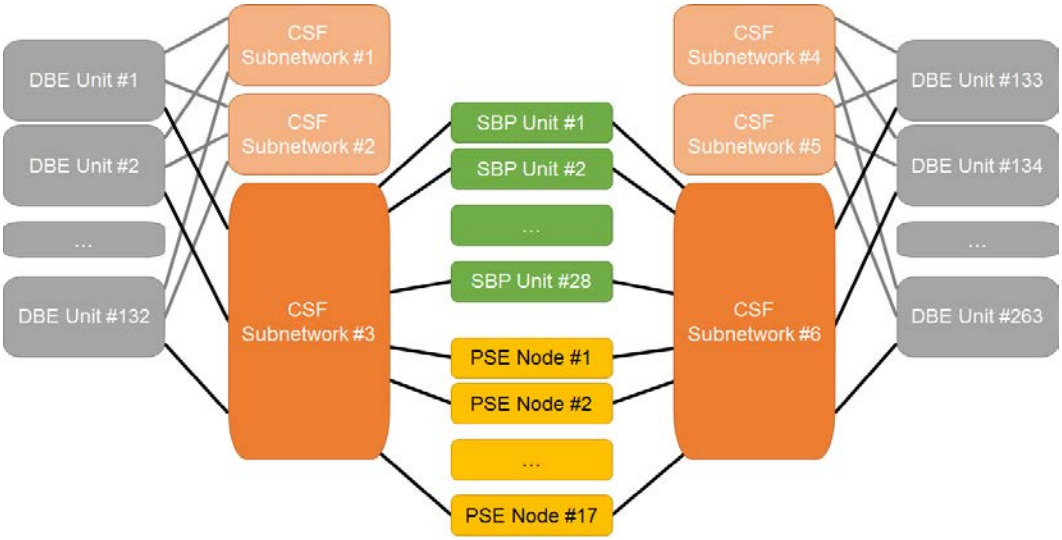

Figure 23: CSF architecture partitioned in the bandwidth and the antenna domain. Six independent CSF subnetworks transport the data streams from one third the bandwidth and one half the antennas. Each subnetwork must provide at least 320 400G ports. The overall number of ports increases from 1838 in the single switch matrix architecture, to YY port in this architecture.

In this architecture, each DBE distributes one third of the bandwidth through each subnetwork. But now, each subnetwork only receives data from half the DBE units. Thus, the number of ports of each subnetwork used by the DBE is 132. On the SBP side, each LRU only processes the data from 22 antennas (11 antennas per FPGA). As a result, six LRUs of every SBP unit connect to one subnetwork to process the data from 132 DBE units, and the other six LRUs of each SBP unit connect to a different subnetwork through which it receives the data from the other half of antennas. Therefore, to support 28 SBP units, with six connected LRUs each, a CSF subnetwork needs 168 ports. As regards the PSE, 17 nodes are needed to process one-third of the bandwidth. However, they must connect to the two subnetworks each SBP unit is connected to, because they need to receive data from any LRU in an SBP unit. Something similar happens for the ONL, which will need three additional ports from each of the six subnetworks. Overall, each subnetwork need to be populated with at least 320 ports, which is readily possible with

COTS equipment. Excluding any spare inputs for future capabilities, the overall number of ports increases to 1920, about 4%, with respect to the single switch matrix CSF architecture.

5.5 CSP Technology

The ngVLA CSP design consists of heterogeneous sub-elements performing the varied DSP tasks required to support system requirements. This includes the digital backend (DBE), the CSP Switched Fabric (CSF), the Sub-Band Processor (SBP), and the Pulsar Engine (PSE), among others. Each of these sub-elements has gone through its own technology review process, and the preferred hardware platform for each—GPU, FPGA, or ASIC—depends on the characteristics of the tasks required.

Modern correlators have been built on each platform: ASICs (VLA, ALMA), FPGAs (VLA, SKA-Mid), and GPUs (CHIME, HERA). In broad terms, modern GPUs offer extremely powerful vector and matrix processing but have relatively limited I/O; FPGAs have a broad range of flexible capabilities and may feature exceptionally high-bandwidth I/O resources; and ASIC solutions consume very little power but require extensive development efforts. Particularly for the ASIC-to-FPGA comparison, the number of units to be produced is a dominant factor in the overall cost comparison: FPGAs are more cost-effective in small volumes, while the reduced cost per unit of ASICs makes them more suitable for large volumes. In addition, ASIC's much lower power draw can lead to significant lifecycle cost savings for an instrument with an operating life of 20 years.

The details of this trade are requirement- and design-dependent, so absent a reference architecture for the central signal processor, an assessment cannot be conclusive. Contextually, ASIC development costs declined through 2019 for all but the most cutting-edge processes, as ASIC development has become more common. However, as of 2021, the ASIC design and fabrication market is oversubscribed, with significant backlogs for both designers and wafer fabricators. The lifecycle cost comparison will depend on the evolution of these market trends and on the process nodes available in FPGA and ASIC architectures at the time of construction.

The low volume of DBE units required (a few hundred for the entire array) strongly favors FPGAs over ASICs, while the limited I/O of GPUs removes them from consideration. The DBE must process in real-time the output of at least 16 active ADCs (32 in some down-conversion schemes), each at 56 Gbps, while transmitting sub-band data streams at an overall data rate of almost 1 Tb/s. While some modern FPGAs can support these data rates with a single device, a GPU-based system would require many devices (and additional networking hardware) to support the same overall throughput.

For the switched fabric, a solution featuring COTS switches (themselves based on ASICs) is by far the most cost-effective and efficient option. A number of vendors offer turn-key solutions satisfying current and future ngVLA needs; a competitive evaluation process will determine the most suitable choice.

Both FPGAs and GPUs have been considered for the design of the Pulsar Engine. As with the DBE, ASIC development costs are not justified given the small volume of units needed. [RD07] examines the various trade-offs driving the pulsar engine design and concludes that FPGA devices with on-chip High Bandwidth Memory (HBM) are the best choice for the pulsar engine's design. This is primarily due to their excellent I/O and internal memory bandwidth, supported by pre-existing signal processing designs which demonstrate a low design risk.

The downselection between CSP designs is planned to be part of the central signal processor conceptual design review. The selection of the TRIDENT FSA design for the CSP would entail the choice of the FPGA-based TALON-DX board for sub-band processing; although the reference design calls for a passive optical network, an active switched network using ASIC-based COTS hardware would offer considerable additional capabilities. The SCREAM design separates the beamforming and channelization (B&C) stage from the correlation (X) stage, permitting different technologies to be used for each (see Section 5.6).

The SCREAM B&C nodes have data interchange requirements that limit the hardware choices to FPGAs or ASICs; the larger volume required makes an ASIC design more feasible. This is because the SCREAM architecture achieves full subarray independence by performing per-antenna, per-sub-band signal processing in separate devices for each antenna, an approach that would be highly inefficient if implemented with FPGAs. In contrast, each individual FPGA device of the FSA processes data from a set of 11 antennas, which is the main motivator for the single-mode operation of its sub-band processors. In this regard, the ASIC design considers the use of commodity manufacturing processes and clock rates that minimize power consumption but at the same time should not lead to unaffordable total development costs.

An X-engine based on ASICs which could also be reconfigured to support pulsar engine functions was examined [RD08]; this was an attractive option, as it would permit unused X-engine resources to be transferred to pulsar operations when in beamforming modes. However, poor functional overlap and the additional design and manufacture costs it would require made this option untenable. A GPU-based X-engine design was also considered, but its limited I/O and requirement for supplemental network hardware made the system's hardware cost and overall power consumption unreasonably high. Notwithstanding that, GPU-based solutions, such as those under development for MeerKAT [RD09], will be closely monitored until a final design downselection is made. As detailed in [RD08], an X-engine based on AI-optimized FPGAs was determined to be the best choice; these offer the I/O and processing required with relatively low power consumption.

5.6 SCREAM Design for the Sub-Band Processor

The proposed Sub-Band Processor subsystem design relies on the participation of the Canadian NRC for its design and development. As NRC is currently occupied with the design and development of the SKA1-Mid Correlator and Beamformer, a lack of their availability is one of the biggest risks in the development of the ngVLA CSP.

The SCREAM (SCalable, REconfigurable And Modular) architecture is an ongoing NRAO project that aims to capture the orders-of-magnitude improvements that ASIC technology can offer for future radio astronomy interferometers. Although development is currently in its early stages, the SCREAM project presents a compelling alternative design for the SBP subsystem in the event that NRC's participation in developing the ngVLA CSP does not materialize.

One key difference in the SCREAM design is the SBP consists of two parts: the B&C (Beamformer and Channelizer), which takes the role of the B- and F-Engines, and the X-Engine. The following sections describe the differences in the mechanical and electrical design of a SCREAM-based SBP, as pertaining to each of these subsystems.

5.6.1 Beamformer and Channelizer

The B&C subsystem performs the functions of the SBP that relate to the B- (beamforming) and the F- (channelization) Engines. These have been described in Sections 4.2.1.5 through 4.2.1.8. There is no significant difference in how delay and phase tracking are performed; this is also based on resampling the incoming data stream and an NCO for the phase rotation. The same applies to the zoom engine. However, the frequency channelization is much more flexible in the SCREAM design, using a multistream, multiresolution PFB-based frequency channelizer. The TALON-based design uses a fixed frequency resolution (the finest one) at the F-Engine, then multiple frequency channels are averaged at the X-Engine to achieve the desired frequency resolution. The frequency resolution is also fixed for beamforming modes. On the contrary, SCREAM's B&C design uses a much more flexible frequency channelizer, where the frequency resolution can be selected at the F-Engine. In addition, the frequency channelizer can also process multiple data streams (e.g., phase centers in beamforming modes) by trading in frequency

resolution for number of streams. The increased flexibility is a natural reaction to the lack of reprogrammability in an ASIC-based design, as compared to an FPGA-based one.

Another key difference of the SCREAM-based design is that it has been developed with full subarray independence at its core. As a result, each B&C unit can operate in multiple observing modes (independently set on a per-antenna basis), removing the need for additional units for simultaneous operation in multiple observing modes. Full subarray independence demands that all antennas must be processed independently from each other. The simplest way of achieving this is by using separate hardware for each antenna, which naturally leads to ASIC technology as the most suitable for the SCREAM architecture. To illustrate the above conclusion, consider a big radio interferometer such as the ngVLA. The ngVLA design targets 20 GHz of bandwidth and 263 antennas. Assuming two sub-bands from the same antenna are processed by each device, and 218.75 MHz per sub-band, the number of devices (B&C nodes) in the B&C part would amount to 12,098 nodes (46 sub-band pairs, 263 antennas).

An FPGA-based design with so many devices would be orders of magnitude less efficient. Indeed, FPGAs are most efficient when as much functionality as possible is integrated in a single FPGA. For example, the FPGA-based SBP design assumes the data streams from 11 antennas are processed at each FPGA, as well as a distributed X-Engine module. The cost of this integration is the lack of freedom to process each antenna independently, which ultimately results in having to almost double the required hardware in order to support simultaneous beamforming. By keeping a modular design, i.e., one B&C node per antenna and independent F- and X-Engines, the SCREAM architecture avoids these limitations and minimizes power consumption.

A consequence of using separate devices for each antenna is that it becomes impractical to connect all B&C nodes in an SBP unit to each other for beamforming. Instead of a fully connected mesh topology, which scales poorly with the number of antennas, the SCREAM design uses a ring topology through which the beamforming data is sent from one node to the adjacent ones, using both chip-to-chip and board-to-board communications. A star topology (or hub-spoke) is also under study. This topology would use the CSF to communicate the beamforming data among nodes. The advantage is better support for subarrays in beamforming modes, but at the cost of higher complexity in the management of data transport delays.

In the current B&C board concept, 16 nodes are packed on a single board, as shown in Figure 24 (next page). A QSFP-DD connector is used to connect the B&C board to the CSF, hence for receiving data from the DBE, as well as outputting the data generated by the B&C nodes. These are either beamforming products or phase-delay corrected antenna data streams ready for cross-correlation at the SCREAM's X-Engine.

The B&C board also includes an FPGA device acting as communications gateway for the B&C node and providing housekeeping functionality for the whole board. This device also includes enough external memory to provide the board with the required data buffering capability. Thus, the coarse delay correction is performed by the FPGA. All B&C nodes are connected to adjacent nodes in a ring topology which is used for beamforming, as described above. 2 mid-board optical modules allow extending the beamforming ring to multiple B&C boards.

Each B&C LRU would consist of 2 B&C boards and fit on a standard 1U, air-cooled, 19-inch enclosure. The power estimate per LRU is between 100-200W. For processing the whole array, each B&C Unit will need 9 LRUs, providing 25 spare nodes that can be used in the future for processing other antenna inputs, e.g., eVLBI. Since each antenna is independently processed within each B&C Unit, only 46 Units (as opposed to 83 units in the FSA proposed design) are needed to process 20 GHz of bandwidth in either interferometric or beamforming OMs. The total power consumption estimate for the B&C part is approximately 50–100kW.

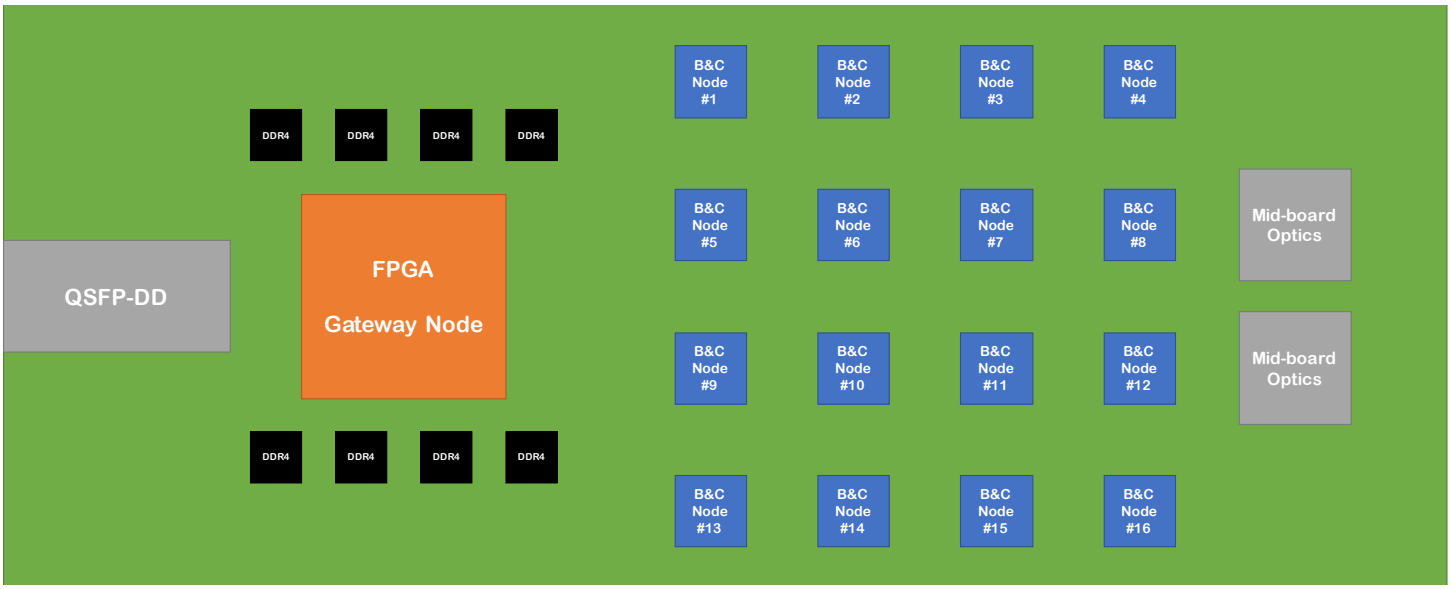

Figure 24: Preliminary B&C board layout featuring 16 nodes per board and other main processing and data communications components. Two boards would fit into a 19-inch enclosure constituting a B&C LRU.

The 46 SBP units could be installed in 12 standard 42U 19-inch racks, each one populated with 4 B&C units, at 9U per unit.

5.6.2 Correlator or X-Engine

The key changes in the SCREAM X-Engine design are to perform the frequency-axis distribution of the correlation task in the CSF (rather than within each SBP) and to arrange computation to more efficiently occupy the selected processing hardware. The latter is achieved by not computing visibilities which will be immediately discarded and by reducing the factor by which the correlation task is partitioned. The effect of these changes is to reduce the number of FPGAs dedicated to the correlation task from ~2000 in the FSA design to ~400. The initial hardware cost, infrastructure requirements, and power consumption decrease proportionately.

Each SCREAM XE node consists of a single FPGA with associated high-bandwidth I/O interfaces, external memory, M&C connections, power supply, and other required support hardware. Although general-purpose FPGAs with capabilities similar to that currently specified in the FSA's TALON module could be used, variants with emphasis on low-bit-depth computational assets (marketed as "AI FPGAs") offer a significant cost–performance advantage. A number of appropriate COTS options are currently in production, although substantially higher efficiency (and support for liquid cooling) could be achieved by pursuing a modified COTS or other ngVLA-specific design. Additional networking capacity (above that currently available) would permit large efficiency gains when operating with many subarrays.

It is possible to construct a SCREAM XE LRU which shares the structure of the TALON design: a SCREAM XE LRU would consist of two SCREAM XE nodes and be contained in a standard 1U 19-inch enclosure which supports liquid cooling. Assuming the same power consumption as a TALON-based LRU (600W) yields an overall power consumption of around 120kW for the SCREAM XE. When combined with the estimated power consumption of the SCREAM B&C, this results in approximately 200kW for a SCREAM-based SBP, about one third of the TALON-based design.

While both use similar FPGA technology, SCREAM's XE consumes just about one-third the power of a hypothetical TALON-based design, after excluding the extra SBP units required to support simultaneous beamforming operation. This suggests that the TALON F-Engine requires twice the power of the X-

Engine, justifying its choice as the first element for power optimization through ASIC implementation. Due to the separation of the X-Engine from the F-Engine, it may be possible to further optimize the X-Engine design in the future—for example, extrapolating the results in [RD15] (based on an older 32nm IBM process), a SCREAM XE could consume as little as 13kW.

Alternately, using FPGA technology for the SCREAM XE reveals synergies with the PSE subsystem. If a compatible hardware platform is used for both subsystems, a single PSE-XE subsystem can be developed in which each node is reprogrammed to function as either an XE or PSE node as required. This has particular relevance as the XE must be sized to correlate the entire array while the PSE would normally be employed by select subarrays operating in beamforming modes. This implies that some portion of the XE nodes are necessarily underutilized when the PSE is active. As a result, a joint FPGA-based PSE-XE subsystem would reduce the size, power, complexity, and cost of independent FPGA-based PSE and XE designs.

A disadvantage of the independent F- and X-Engine design is the increase in the number of independent connections to the CSF. Assuming 8-bit quantization, the overall data rate flowing from the B&C part to the XE is approximately 170 Tb/s, which requires at least 425 400G ports on each side, i.e., 850 ports for the whole SBP. This is less than the current port count of the SBP, but only because the current design nearly doubles the number of SBP units in order to fully support simultaneous beamforming capabilities.

Furthermore, the number of ports required in the SCREAM design will necessarily be greater than this due to inefficiencies in the partitioning of the system. For example, at 16 B&C nodes per board, the output data rate of a B&C board will be around 224 Gb/s, much lower than the theoretical 400 Gb/s available. The resulting number of ports just for the B&C part therefore increases to 782, from the 425 above. This must be considered in future revisions of a SCREAM-based SBP design, particularly those examining the partition of its architecture.

Mechanically, the SCREAM XE could be installed in six standard 19-inch 42U racks, assuming 36 XE LRUs per rack. When combined with the 12 racks populated with the B&C LRUs, a SCREAM-based SBP would use a total of 18 racks, in addition to the ~2 racks required for the pulsar engine.

6 Appendix B: Abbreviations and Acronyms

Acronym	Description
AC	Alternating Current
AD	Applicable Document
AFD	Antenna Fiber Optic System
ALMA	Atacama Large Millimeter Array
ARCS	Advanced RFI Containment System
ASIC	Application Specific IC
B&C	Beamformer and Channelizer
BMR	Bins, Modules, & Racks
CBE	CSP Back End
CDU	Coolant Distribution Unit
CHIME	Canadian Hydrogen Intensity Mapping Experiment
CMAC	Complex Multiply-Accumulate
COTS	Commercial Off-the-Shelf
CSP	Central Signal Processor
CSS	Computing and Software System
DBE	Digital Back End
DC	Direct Current
DCLC	Direct Contact Liquid Cooling
DDR4	Double Data Rate 4
DIMM	Dual In-line Memory Module
DSP	Digital Signal Processing
EMC	Electromagnetic Compatibility
ESD	Electrostatic Discharge
EVLA	Expanded VLA
eVLBI	Real-Time VLBI
FFT	Fast Fourier Transform
FinFET	Fin Field-Effect Transistor
FIR	Finite Impulse Response
FITS	Flexible Image Transport System
FMECA	Failure Mode, Effects, and Criticality Analysis
FPGA	Field Programmable Gate Array
FSA	Frequency Slice Architecture
FSP	Frequency Slice Processor
Gb/s	Gigabits per second
GBd	Gigabaud
GPU	Graphics Processing Unit
GS/s	Giga samples per second
HBM	High-Bandwidth Memory
HERA	Hydrogen Epoch of Reionization Array
HVAC	Heating, Ventilation and Air Conditioning
FIB	Central Fiber Optic Distribution/Infrastructure
IC	Integrated Circuit
IF	Intermediate Frequency
I/O	Input/Output
IQ	In-phase and Quadrature

Acronym	Description
IRD	Integrated Receivers and Digitizers
LO	Local Oscillator
LRU	Line Replaceable Unit
M&C	Monitor and Control
MCL	Monitoring and Control System
ngVLA	Next Generation VLA
NCO	Numerically Controlled Oscillator
NRAO	National Radio Astronomy Observatory
NRC	National Research Council
NSB	ngVLA Site Buildings
OBT	On-Board Transceiver
OM	Observing Mode
ONL	Online Sub-element
PAM	Pulse Amplitude Modulation
PBS	Product Breakdown Structure
PCB	Printed Circuit Board
PDR	Preliminary Design Review
PDU	Power Distribution Unit
PFB	Polyphase Filter Bank
PSE	Pulsar Engine
PSU	DC Power Supply System
QSFP	Quad SFP
QSFP-DD	QSFP Double Density
RAM	Random-Access Memory
RF	Radio Frequency
RFI	RF Interference
RMS	Root Mean Square
RTD	LO Reference and Timing Distribution
RTG	LO Reference and Timing Generation
SBP	Sub-Band Processor
SCREAM	Scalable, Reconfigurable, and Modular
SDRAM	Synchronous Dynamic RAM
SETI	Search for Extraterrestrial Intelligence
SEU	Single-Event Upset
SFP	Small Form-factor Pluggable
SKA	Square Kilometre Array
Tb/s	Terabits per second
TBD	To Be Defined/Determined
TRL	Technology Readiness Level
TTD	True Time-Delay
VDIF	VLBI Data Interchange Format
VLA	Karl G. Jansky Very Large Array
VLBI	Very Long Baseline Interferometry
XE	X-Engine (Cross-correlation Engine)

020.40.00.00.00-0005-DSN-B-CSP_DESIGN_DESCRIPTION-CoDR

Final Audit Report 2022-02-10

Created:	2022-02-10
By:	Anne Lear (alear@nrao.edu)
Status:	Signed
Transaction ID:	CBJCHBCAABAAD9AD89IFJql0SLoDQfR6bsq5vr2EunhB

"020.40.00.00.00-0005-DSN-B-CSP_DESIGN_DESCRIPTION-CoDR" History

Document created by Anne Lear (alear@nrao.edu)
2022-02-10 - 3:35:48 AM GMT- IP address: 75.161.242.227

Document emailed to William Shillue (bshillue@nrao.edu) for signature
2022-02-10 - 3:36:55 AM GMT

Email viewed by William Shillue (bshillue@nrao.edu)
2022-02-10 - 2:05:03 PM GMT- IP address: 216.30.188.197

Document e-signed by William Shillue (bshillue@nrao.edu)
Signature Date: 2022-02-10 - 2:05:26 PM GMT - Time Source: server- IP address: 216.30.188.197

Document emailed to R. Selina (rselina@nrao.edu) for signature
2022-02-10 - 2:05:29 PM GMT

Email viewed by R. Selina (rselina@nrao.edu)
2022-02-10 - 2:26:58 PM GMT- IP address: 75.161.226.40

Document e-signed by R. Selina (rselina@nrao.edu)
Signature Date: 2022-02-10 - 3:41:54 PM GMT - Time Source: server- IP address: 75.161.226.40

Document emailed to Thomas Kusel (tkusel@nrao.edu) for signature
2022-02-10 - 3:41:56 PM GMT

Email viewed by Thomas Kusel (tkusel@nrao.edu)
2022-02-10 - 4:59:57 PM GMT- IP address: 71.62.226.171

Document e-signed by Thomas Kusel (tkusel@nrao.edu)
Signature Date: 2022-02-10 - 5:00:28 PM GMT - Time Source: server- IP address: 71.62.226.171

Document emailed to Willem Esterhuyse (westerhu@nrao.edu) for signature

2022-02-10 - 5:00:31 PM GMT

Email viewed by Willem Esterhuyse (westerhu@nrao.edu)

2022-02-10 - 5:32:55 PM GMT- IP address: 105.224.252.103

Document e-signed by Willem Esterhuyse (westerhu@nrao.edu)

Signature Date: 2022-02-10 - 5:44:21 PM GMT - Time Source: server- IP address: 105.224.252.103

Agreement completed.

2022-02-10 - 5:44:21 PM GMT