

**ngVLA Electronics Memo #11**  
**A SCREAM-Compatible ngVLA Pulsar Engine:**  
**Key Requirements Review and Option Trade-Off Study**

NOLAN DENMAN<sup>1</sup>

<sup>1</sup>*Central Development Laboratory, National Radio Astronomy Observatory*

Contents

1. Tasks Required of the Pulsar Engine	2
1.1. Input from the B&C Nodes	2
1.1.1. Data Itself	2
1.1.2. Metadata	3
1.1.3. Flags	4
1.2. Input from the M&C System	4
1.3. De-Dispersion	4
1.4. Detection	5
1.5. Folding	5
1.6. Output	6
1.6.1. Metadata	6
1.6.2. Phase Profiles	6
1.6.3. Offline Pulsar Search Data	7
2. Features and Extensions of the Pulsar Engine	7
2.1. Subarraying and Multiple Phase Reference Positions	7
2.2. Multiple Folding Threads	7
2.3. Corrections and Calibration	8
2.4. RFI Excision	8
2.5. I/O Handling	8
2.6. Higher-Order Accumulation	8
3. Potential Pulsar Engine Architectures	9
4. Potential Pulsar Engine Hardware	9
4.1. CPUs	9
4.2. GPUs	9
4.2.1. NVIDIA Ampere	9
4.2.2. Jetson AGX Xavier	9
4.3. FPGAs	10
4.3.1. General-Purpose FPGA	10
4.3.2. FPGA with High-Bandwidth Memory	10

4.3.3. TALON-DX	10
4.4. NRAO Custom ASIC	11
5. Evaluation	11
5.1. Metrics Used	11
5.2. Hardware Option Evaluation	12
5.3. Rationale for Selection of Hardware for Conceptual Design	13

**Note** This document discusses an ngVLA Pulsar Engine (P-Engine or PSE) design which is compatible with the SCREAM correlator design; the system presented may be tweaked for compatibility with any correlator which can supply channelized and beam-formed data in an appropriate format.

**Note** This document includes a description of selected key PSE requirements and a trade-off study of the architecture and hardware choices. For a description of the conceptual design for the PSE, please see [Denman et al. \(2021a\)](#).

## 1. TASKS REQUIRED OF THE PULSAR ENGINE

With the X-Engine or other portions of the correlator performing the synthesis imaging tasks, the Pulsar Engine provides the signal processing required for time-domain observations. The Pulsar Engine receives formed beams from the Beamforming and Channelization (B&C) nodes or their functional equivalent. Each beam has 2 polarizations with some total bandwidth, time- and frequency-resolution, and bit depth. These beams have been coarse- and fine-delay corrected and channelized into a number of frequency channels. They require further processing, which may be some combination of coherent de-dispersion, folding by some phase model, and other tasks as required. Figure 1 shows data and control flows within the PSE.

### 1.1. *Input from the B&C Nodes*

The Pulsar Engine receives channelized formed beams from the previous stage of the correlator, accompanied by a selection of metadata (potentially including a flag stream). Details of the interconnection such as routing and packet handling are not considered in this document; it is assumed to be an Ethernet switched network with sufficient bandwidth.

#### 1.1.1. *Data Itself*

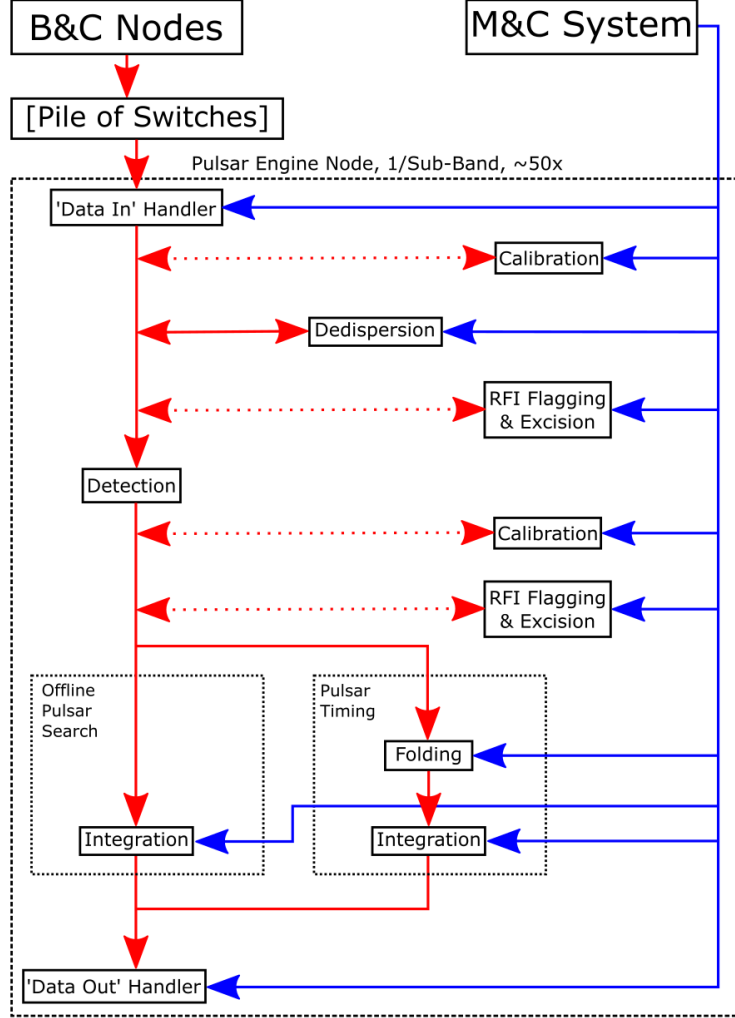
For the purposes of this discussion, the PSE’s initial input from the B&C nodes is assumed to be in 32+32 bit complex format ( $N_{bits} = 64$ ), with two polarization components ( $N_{pol} = 2$ ) in each of ten beams ( $N_{beams} = 10$ ). The bandwidth  $BW$  of one sub-band<sup>1</sup> is asserted to be approximately 200 MHz, and in modes in which the entire sub-band is processed it is divided into  $N_{chan} = 2^{7-14}$  (128-16384) channels. This suggests channel bandwidths of approximately 12 kHz to 1 MHz.

On a per-sub-band, per-beam basis, a packet<sup>2</sup> of size

$$\left(\frac{N_{bits}}{64}\right) \left(\frac{N_{pol}}{2}\right) \left(\frac{N_{chan}}{256}\right) * 4096 \text{ bytes} \quad (1)$$

<sup>1</sup> A unit of signal bandwidth used to describe parallelism within the correlator.

<sup>2</sup> The term ‘packet’ refers here to a single time-interval’s observations for all beams and phase reference positions; it will consist of multiple network protocol packets.



**Figure 1.** An illustration of data and control flows within the Pulsar Engine. Operations marked with a dotted line may not be required in the final version.

will arrive at an average interval of

$$\left(\frac{200 \text{ MHz}}{BW}\right) \left(\frac{N_{chan}}{256}\right) * 1.3 \text{ microseconds} \quad (2)$$

for a per-sub-band per-beam data input rate of

$$\left(\frac{BW}{200 \text{ MHz}}\right) \left(\frac{N_{bits}}{64}\right) \left(\frac{N_{pol}}{2}\right) * 25.6 \text{ Gbps} \quad (3)$$

### 1.1.2. Metadata

In addition to the beam data itself, the B&C nodes will produce (or relay) metadata required by the Pulsar Engine or later processing stages.

Most immediately relevant to the Pulsar Engine's operation are metadata with timing and frequency information. Existing astronomical data interchange formats such as VDIF (Kettenis et al. 2014)

provide an example structure: in VDIF, time is specified as seconds since a reference epoch plus a frame counter within the second; the 24-bit counter supports packet-level timing up to a resolution of  $\approx 60$  nanoseconds.

Given the extremely high precision required for pulsar timing experiments (Demorest & Ransom 2018; NANOGrav Collaboration 2018), defining the arrival time and ensuring correct propagation through Pulsar Engine processing is vital to its success. Of particular importance is ensuring that the effects of de-dispersion and folding are consistently accounted for in the timing information relayed by the Pulsar Engine – this must be determined in consultation with downstream pulsar observers.

### 1.1.3. *Flags*

The topic of high-time-resolution Radio-Frequency Interference (RFI) mitigation within the ngVLA correlator is currently under examination (Rau et al. 2019; Selina et al. 2020a; Amestica et al. 2021). In the event that high-time-resolution flagging is implemented prior to the Pulsar Engine, the data may be accompanied by a set of flags of the same dimensions. These are considered separately from the general metadata due to their potentially large size.

1-bit flags for each of the input time-samples will increase the input rate by a few percent; higher-bit-depth flags have proportionally greater data transport requirements. These flags must be carefully propagated; the interaction of flagging and phase-folding has not yet been defined.

This is independent of, and possibly in addition to, any RFI excision which is implemented in the Pulsar Engine itself (§2.4).

## 1.2. *Input from the M&C System*

In addition to the data and metadata arriving from the B&C nodes, observing and processing parameters must be supplied by the ngVLA monitor and control (M&C) system (Koski et al. 2019). This will include the selection of observing and processing modes and their required parameters.

For pulsar folding modes, the PSE requires the reference epoch  $t_0$ , ephemeris  $\phi(t_0)$ , and phase model  $\phi(t - t_0)$  of the target (see §1.5). If de-dispersion is applied, the target dispersion measure must likewise be supplied.

The details of interaction between the PSE and M&C systems are not yet defined.

## 1.3. *De-Dispersion*

Corrections for interstellar dispersion are most easily applied in Fourier-space; this involves a Fourier transform, application of the transfer function, and then an inverse Fourier transform (van Straten 2003; Bassa et al. 2017). The use of overlap-save convolution (as in DSPSR (van Straten & Bailes 2011)) permits efficient application of the de-dispersion kernel, but this remains a resource-intensive process; a discussion of the computational and memory-bandwidth requirements follows.

If a channel with central frequency  $\nu$  and width  $\Delta\nu$  is de-dispersed using overlap-save convolution with a dispersion of  $D$  and an FFT length of  $N_{FT}$ , the number of samples which must be discarded  $N_{discard}$  is

$$N_{discard} = \frac{D}{\Delta\nu} \left( \left( \nu - \frac{\Delta\nu}{2} \right)^{-2} - \left( \nu + \frac{\Delta\nu}{2} \right)^{-2} \right) \quad (4)$$

For large power-of-two values of  $N_{FT}$ , this requires approximately  $5N_{FT} \log_2(N_{FT})$  operations and produces  $N_{out} = N_{FT} - N_{discard}$  samples of de-dispersed output corresponding to a time period of

$\frac{N_{out}}{\Delta\nu}$ . The total number of operations per unit time per channel is therefore approximately

$$\text{ops} = \frac{10N_{FT} \log_2(N_{FT})}{N_{FT} - N_{discard}} \Delta\nu \quad (5)$$

For each FFT, the entire set of data must be read from working memory approximately twice (Hemmert & Underwood 2005). De-dispersion therefore requires  $4N_{FT}$  samples of size  $d$  to be written every  $\frac{N_{out}}{\Delta\nu}$ , for an approximate per-channel memory bandwidth of

$$\text{bw} = d \frac{4N_{FT}}{N_{FT} - N_{discard}} \Delta\nu \quad (6)$$

The most extreme case within the system specifications is an observation centered near 1.2 GHz with 1 MHz channel width and a dispersion measure of  $3000 \text{ pc cm}^{-3}$ . With  $N_{discard} \gtrsim 14,000$  and a choice of  $N_{FT} = 2^{16}$ , each FFT pair accesses  $2^{20}$  bytes ( $\approx 1 \text{ MB}$ ) of data and the system requires 0.2 GOPs of processing and 326 Mbps of memory bandwidth (assuming 32+32-bit samples).

In order to process 10 beams for a full 200 MHz sub-band, a node would therefore require 400 GFLOPs of processing capacity and 651 Gbps of memory bandwidth. Each set of FFTs accesses  $\approx 2 \text{ GB}$  of data, and will require a set of static de-dispersion kernels of comparable size.

A mitigating factor is that the most extreme cases are all in observing band 1, which has an overall bandwidth of only 2.3 GHz. High-DM operations in this band may be required to reduce the per-PSE-node signal bandwidth in order to permit the allocation of more resources per channel to de-dispersion. Additionally, the distribution of known pulsars' dispersion measures is weighted towards the lower end of the range considered; see Figure 2. This may significantly reduce the end-user impact of any performance trade-offs required for very-high-DM operation.

As an alternative, we might consider specific combinations to be 'out of range' and instead record the beamformed voltage data for these observations, with later computation done in offline analysis or SRDP pipelines. This would effectively cordon off a region of DM- $\nu$ - $\Delta\nu$  space as requiring additional end-user involvement to observe but would remove extreme de-dispersion resource requirements as a consideration in the PSE design.

#### 1.4. Detection

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

$$I = |E_x|^2 + |E_y|^2 \quad (7)$$

$$Q = |E_x|^2 - |E_y|^2 \quad (8)$$

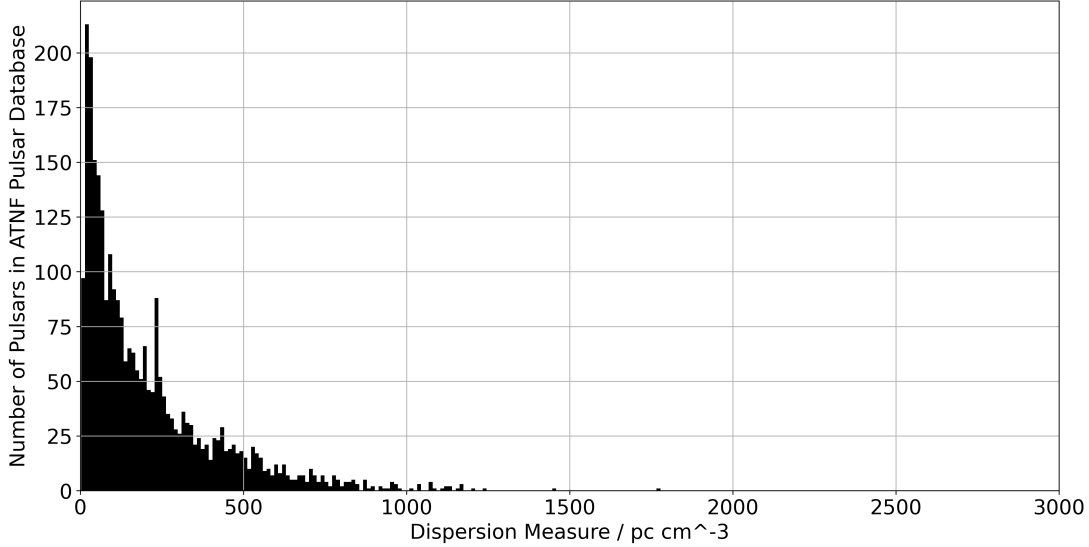
$$U = 2\text{Re}(E_x \overline{E_y}) \quad (9)$$

$$V = -2\text{Im}(E_x \overline{E_y}) \quad (10)$$

This conversion requires very few operations compared to de-dispersion, and it seems plausible that a dedicated processing sub-unit could perform this with great efficiency.

#### 1.5. Folding

The preliminary requirements envision the division of pulse periods ranging from 1 ms to 30 s into as many as 2048 phase bins. The fundamental procedure is to take the sample's arrival time  $t$



**Figure 2.** The distribution of dispersion measures for all pulsars in the ATNF Pulsar Catalogue (Manchester et al. 2005) as of February 2021.

and determine the pulse phase  $\phi(t)$ ; this maps to phase bin  $n$ . Each of the Stokes parameters are accumulated in the  $n$ th corresponding bin, and the  $n$ th counter is incremented. After the desired integration period, the parameters and counter are read out and zeroed.

For a thirty-second-long integration with half-microsecond time resolution, the minimum  $\dot{P}$  value required to move the pulse by half a resolution element is  $8 \cdot 10^{-9} \text{ss}^{-1}$ . This is an order of magnitude larger than the highest values I can find (J1808-2024,  $\dot{P} = 5.5 \cdot 10^{-10} \text{ss}^{-1}$  (Manchester et al. 2005)) which suggests that a first-order phase model like

$$\phi(t - t_0) = \phi(t_0) + \frac{t - t_0}{P} \quad (11)$$

will be sufficient within any single integration.

In this framework, use of a more general phase model (as in §2.6 of van Straten & Bailes (2011)) is left as an external responsibility; the first-order phase model must be supplied and updated at sufficient resolution to ensure it remains correct.

## 1.6. Output

**Note** Version B.04 of the ngVLA Preliminary System Requirements (Selina et al. 2020b) introduced the following constraint: “CON104: Maximum Data Rate: The maximum data rate from the correlator shall not exceed 132 GB/s.”

### 1.6.1. Metadata

In addition to metadata present in the input (§1.1.2), which may be relayed with the data, the PSE itself may generate metadata concerning observing parameters and results. These have not yet been specified; this will be informed by downstream requirements.

### 1.6.2. Phase Profiles

In the primary pulsar timing mode, the folded profiles are read out at intervals whose length is expected to be between 1 s and 10 s.

Avoiding overflow in the accumulation buffers is necessary; the usual radio astronomy model in which the data is effectively Gaussian random noise is not applicable as there will frequently be correlated structure. There are  $10^4 \approx 2^{13.3}$  1 ms periods in a 10 s integration, and the Stokes parameters are the sum of second powers, so for 32+32-bit input data a 64-bit intermediate sum will almost always be sufficient, but long-term accumulation will require larger registers to be entirely safe from overflow.

For data with a number of phase bins  $N_{bins}$  and a readout cadence of  $\Delta t$  on the order of seconds the total data output rate is a relatively modest

$$\left(\frac{1 \text{ second}}{\Delta t}\right) \left(\frac{N_{bins}}{2048}\right) \left(\frac{N_{bits}}{128}\right) \left(\frac{N_{Stokes}}{4}\right) * 1.0 \text{ Mbps} \quad (12)$$

for each beam and frequency channel; with a maximum  $N_{beams} = 10$  and  $N_{chan} = 16384$  this is approximately 172 Gbps. This is manageable in terms of node output, but may exceed the correlator’s overall output data rate limit (see **Note** above) if many sub-bands are simultaneously operating in a similar configuration; this should be incorporated into observation planning.

### 1.6.3. *Offline Pulsar Search Data*

The ‘offline pulsar search’ data product has been defined as the Stokes parameters incoherently integrated to a specified time-frequency resolution. As such, the accumulation is relatively simple but the output data rate may vary widely. For integration by a factor of  $k$  the per-sub-band output data rate is:

$$\left(\frac{1}{k}\right) \left(\frac{BW}{200 \text{ MHz}}\right) \left(\frac{N_{bits}}{128}\right) \left(\frac{N_{beams}}{10}\right) \left(\frac{N_{stokes}}{4}\right) * 1024 \text{ Gbps} \quad (13)$$

With an output data rate limit of 1056 Gbps for the correlator as a whole (see **Note** above), bounds on the value of  $k$  will be required to limit the output data bandwidth to a reasonable range.

## 2. FEATURES AND EXTENSIONS OF THE PULSAR ENGINE

### 2.1. *Subarraying and Multiple Phase Reference Positions*

As described in [Demorest & Ransom \(2018\)](#), the use of multiple subarrays with different main beam positions and the use of multiple phase reference positions within a single main beam are intended for very distinct use cases; ‘multiple beams within multiple subarrays’ is not considered a desirable operating condition. One subarray with multiple phase reference positions and a reduced number of other subarrays exists as a compromise which would require no additional resources.

The Pulsar Engine’s design is inherently parallel on a per-sub-band, per-beam, and per-phase-reference-position basis. With requirements of at most either ten phase reference positions ([Ojeda et al. 2019](#)) or ten subarrays ([NANOGrav Collaboration 2018](#)), for a maximum total bandwidth of 8.8 GHz, a system capable of managing 88 GHz of bandwidth-beam-reference product will satisfy current requirements. This may occur in a combination: for example, one subarray with eight phase reference positions and two additional single-phase subarrays. In the following document, and the companion document ([Denman et al. 2021a](#)), these are both referred to as ‘beams’ in a general sense.

### 2.2. *Multiple Folding Threads*

As specified in the science requirements (Murphy et al. 2019): “Timing multiple pulsars within a single primary beam is desirable. Support for five or more independent dedispersion and folding threads is desired.”

This is trivially satisfied by multicasting input streams to different pulsar nodes, each supplied with different timing information. It may be possible to substantially improve efficiency by supporting duplication of data streams at a later stage (for example, for globular cluster pulsars which may be de-dispersed identically but folded separately). This option would require substantial development; further work has been deferred until it is determined if it will be required.

### 2.3. Corrections and Calibration

As per the preliminary technical requirements (Ojeda et al. 2019), the Pulsar Engine may be required to apply calibration factors to the data. Per-receiver instrumental gains must be applied before beam formation, and therefore before the data is received by the Pulsar Engine. Corrections for quantization and lost data may be applied either in either the B&C or PSE.

According to Demorest & Ransom (2018), pre-summation calibration must be an option; this would take place in the B&C nodes.

In order to correct for per-beam (but not per-antenna) effects, it may be necessary to apply Jones matrix corrections to the pre-detection data in the PSE (Ojeda 2018). These are assumed to be supplied from other calibration stages or external input.

### 2.4. RFI Excision

Automated RFI excision is possible in the Pulsar Engine, provided it is sufficiently frequency-parallel to be implemented within each PSE node. At the cadence of PSE operations, computationally intensive RFI excision algorithms would contribute substantially to the overall system load; this requires careful evaluation.

The detection of RFI might be indicated in a set of flags accompanying the data, either new or merged with pre-existing flags. The interaction of flagging with folding and integration has not yet been defined.

### 2.5. I/O Handling

The above discussion has treated the transfer of data between the B&C nodes and the PSE as perfectly reliable and time-order-preserving. Should this not be the case, it will be necessary to ensure proper handling of errors. Some of this may be handled by the connection protocol; for other cases data buffering and data validity checks (as well as appropriate responses should those checks fail) will be required. The precise nature of these will be informed by experimental tests of candidate data transmission systems’ reliability.

Consultation with downstream pulsar observers will be required to determine the appropriate interactions of flagging, folding, missing or invalid data, and time-integration.

### 2.6. Higher-Order Accumulation

In addition to the Stokes parameters, which are second powers of the voltage measurements, we may optionally compute and accumulate the third and fourth powers of the voltages. These cumulants are described (see van Straten & Tiburzi (2017)) as primarily useful for RFI excision, data quality checks, and instrument calibration rather than routine observation.



This extension would multiply the intermediate storage and output bandwidth requirements, and require some additional computation. Evaluation of the potential usefulness of this feature is required before further development is warranted.

### 3. POTENTIAL PULSAR ENGINE ARCHITECTURES

As the processing tasks for the Pulsar Engine are parallel on a per-sub-band and per-beam basis, the PSE architectures considered invariably take the form of some number of units each of which handles a given beam-bandwidth product. With a switched network distributing input data to the PSE, multicasting and redirection of input are trivial and the interconnection is not detailed further.

An exception is the Frequency Slice Architecture (FSA) of the reference design, which describes a set of PSE nodes connected directly to a set of Frequency Slice Processors; the PSE output is then connected to a switched network. This architecture relies on the selection of the FSA for the correlator as a whole, and additionally on the selection of the TALON-DX board as the hardware for the Frequency Slice Processors (FSPs).

### 4. POTENTIAL PULSAR ENGINE HARDWARE

#### 4.1. CPUs

General-purpose Central Processing Units (CPUs) provide a powerful and flexible platform for many types of scientific computing, but are not well-optimized for the types of extremely large parallel computations the Pulsar Engine requires. If selected, CPUs would require host machines and external Network Interface Cards (NICs).

#### 4.2. GPUs

Modern Graphics Processing Units (GPUs) are heavily optimized for large-scale parallel operations on limited-bit-depth input. A number of types, sizes, and formats of GPU are available; two of the most promising options are described below.

##### 4.2.1. NVIDIA Ampere

The NVIDIA Ampere GPU architecture features a large number of Streaming Microprocessors, each of which combines caches, schedulers, parallel computing modules, and tensor cores (NVIDIA Corporation 2020a, 2021). These GPUs are capable of extremely high rates of processing; the primary difficulty with a GPU-based Pulsar Engine is with data transport to and from the GPU. GPUs are almost all limited to PCIe interfaces with relatively low bandwidth; for the computations required, they have difficulty receiving data quickly enough to occupy their processing capability.

An additional consideration for GPU hardware selection is the choice of specific model and form factor. They typically take the form of individual GPUs hosted in servers and connected by PCIe to separate NICs. One alternative is multi-GPU self-hosting units like the NVIDIA HGX A100 (NVIDIA Corporation 2020c) or DGX A100 (NVIDIA Corporation 2020b) which support high-speed internal networking via the NVLINK protocol. PCIe GPUs offer increased flexibility in networking configuration and power-bandwidth-cost choices, while the latter make inter-GPU data transport an order of magnitude faster.

##### 4.2.2. Jetson AGX Xavier

A low-power GPU optimized for mobile and autonomous computing, the NVIDIA Jetson AGX Xavier module (NVIDIA Corporation 2020d,f) hosts eight Volta-architecture Streaming Multiprocessors (NVIDIA Corporation 2017) equipped with tensor cores capable of up to approximately 1.6 TFLOPs of FP32 computation.

Although the Jetson AGX Xavier is described in marketing materials as having an x8 PCIe Gen4 connection, in practice only one lane is accessible as an endpoint, limiting its data transfer bandwidth severely (NVIDIA Corporation 2020e).

### 4.3. FPGAs

With a range of configurations and options, Field-Programmable Gate Arrays (FPGAs) offer a number of possible solutions; a few are elaborated upon below.

#### 4.3.1. General-Purpose FPGA

Current top-of-the-line all-purpose FPGAs such as the Xilinx Virtex UltraScale+ (Xilinx Inc 2021a,b) and Intel Stratix 10 (Intel Corporation 2020a, 2021) families combine large numbers of high-bandwidth interfaces and a wealth of programmable logic resources.

These FPGAs are available in a number of hosts; these include pre-built standalone development boards like the Xilinx VCU 118 (Xilinx Inc 2018a) and PCIe cards with integrated network interfaces like the BittWare XUP-VV8 and 520N (BittWare Inc 2021b,c). If selected, PCIe FPGAs would require host machines.

Many pre-built FPGA host boards feature one or more ANSI/VITA 57.1 FPGA Mezzanine Connector (FMC) or ANSI/VITA 57.4 FMC+ interfaces, which enable high-speed data transport via a standard for interchangeable modules. Currently-available Commercial Off-The-Shelf (COTS) adapters provide up to 6x 100GbE-capable QSFP28 connectors per FMC+ array (HiTech Global LLC 2021).

#### 4.3.2. FPGA with High-Bandwidth Memory

The families of FPGAs described in §4.3.1 also include models which feature in-package High-Bandwidth Memory (HBM); examples include the Intel Stratix 10 ‘MX’ line (Intel Corporation 2020b) and the Xilinx Virtex UltraScale+ VU37P (Xilinx Inc 2021a,b). These offer similar processing and expansion capabilities but an order of magnitude more local memory bandwidth.

As with the general-purpose FPGAs, these are available as pre-built standalone development boards like the Xilinx VCU 128 (Xilinx Inc 2018b) and PCIe cards with integrated network interfaces like the BittWare XUP-VVH and 520N-MX (BittWare Inc 2021d,a). If selected, PCIe FPGAs would require host machines.

#### 4.3.3. TALON-DX

Featured in the ngVLA Reference Design (Ojeda 2018) and in previous minimally-modified derived designs (Ojeda 2020a,b) the TALON-DX board hosts one of a selection of Intel Stratix 10 SX-variant FPGAs along with large quantities of DDR4 memory and a number of high-speed interfaces (Pleasant et al. 2017; Carlson & Pleasant 2018).

The TALON-DX has a great deal of similarity to notional NRAO-designed ngVLA-specific FPGA host boards, save for its use of Leap On-Board Transceivers<sup>3</sup> (OBTs) to interface with a planned

<sup>3</sup> Referred to in Carlson & Pleasant (2018) as the FCI Leap Mid-Board Optic (MBO).

passive optical interconnect. For the purposes of this discussion, it is considered an instance of the General-Purpose FPGA design described in §4.3.1 which has committed to a specific FPGA package and optical interface.

#### 4.4. NRAO Custom ASIC

An Application-Specific Integrated Circuit (ASIC) designed explicitly for the ngVLA Pulsar Engine would require an initial design process comparable to that for a dedicated FPGA solution on an NRAO-designed board, followed by an in-depth verification and fabrication process. Performance could be precisely as required, and power consumption would be substantially reduced relative to an FPGA implementation.

The potential advantages of this solution are primarily the power and space savings, which must be balanced against the cost of implementing and manufacturing an ASIC.

## 5. EVALUATION

### 5.1. Metrics Used

The metrics used to assess the choice of hardware employed in the conceptual design description (Denman et al. 2021a) are presented and briefly described in Table 1.

The extremely parallel structure of the Pulsar Engine means that almost any computing hardware is capable of fulfilling the minimum requirements, if deployed in sufficient quantity. The evaluation of potential hardware choices is therefore an optimization; the capabilities of each hardware option inform the required quantity of nodes, and therefore the Hardware Cost and Power Consumption. The Hardware Cost metric represents the cost of enough of a given choice of hardware to satisfy the requirements of the Pulsar Engine, rather than a per-unit price. This evaluation models the lifecycle cost as predominantly NRE plus Hardware Cost plus Power Consumption. The ‘maturity’ of a proposed solution is included as a consideration to represent our confidence in (and the probability of) deploying a stable and effective Pulsar Engine based upon a specific choice of hardware.

Potentially also a consideration, but not currently included due to high uncertainty, is the difficulty of sourcing a given hardware choice independent of price.

(Hardware Cost)	Estimated cost of purchasing the hardware itself
(Non-Recurring Engineering)	Design effort required to implement Pulsar Engine
(Power Consumption)	Cost of power consumed during operation
Maturity	Previous use of this hardware for a similar purpose

**Table 1.** Metrics used to evaluate hardware options. For those in parentheses, lower/less is better; for others, higher/more is better.

For the purposes of balancing hardware cost and NRE against power consumption, we have used the value of \$0.15 per kWh or \$1.31 per watt-year from Ojeda (2018) and a twenty-year hardware refreshment cycle; the results in an additional cost-of-power of  $\approx$ \$26 per node per watt of additional power consumption. This is intended to capture both the direct cost of additional power and the cost of cooling required to dissipate the additional waste heat generated.

## 5.2. Hardware Option Evaluation

Table 2 presents a qualitative evaluation of the hardware options described in §4 against the metrics in Table 1. A more detailed description of each option’s evaluation follows, with quantitative comparisons where appropriate.

Hardware		Ref	(HW Cost)	(NRE)	(Power)	Maturity
CPU	Generic	§4.1	DQ	-	DQ	-
GPU	Ampere	§4.2.1	High	Med	High	Med
GPU	Jetson	§4.2.2	Med	Med	Med	Med
FPGA	General	§4.3.1	Med	Med	Med	Med
FPGA	HBM	§4.3.2	Low	Med	Low	Med
FPGA	TALON	§4.3.3	Med	Med	Med	Med
ASIC	Custom	§4.4	?	High	Low	Low

**Table 2.** Qualitative evaluation of hardware options from §4 against criteria from §5.1 For those metrics in parentheses, lower/less is better; for others, higher/more is better. ‘DQ’ indicates that performance regarding a particular metric is considered disqualifyingly bad; ‘?’ indicates that evaluation of a metric is ongoing.

**CPUs:** Intended for flexible general-purpose computing, CPUs are poorly suited to the Pulsar Engine’s large-scale parallel processing. The vast number of CPUs required and corresponding power consumption remove this option from consideration.

**NVIDIA Ampere GPUs:** Although more than capable of the required processing, and featuring very large quantities of on-die HBM, GPUs have relatively restricted input data bandwidth; this inflates the quantity (and therefore cost) of the hardware required to perform the specified pulsar processing.

Future developments in the GPU field may alter this conclusion; the introduction of user-accessible high-speed external interfaces should prompt a re-evaluation of this option.

**NVIDIA Jetson AGX Xavier GPUs:** As with other GPUs, I/O limitations require an extremely large number of nodes to receive the incoming beams. This, in turn, results in increased initial hardware costs and power consumption significantly greater than other solutions, removing this option from consideration.

**General-Purpose FPGAs:** Modern general-purpose FPGAs are, at the chip level, well-supplied with transceivers supporting extremely high input and output data rates and with flexible computational resources. They are not, however, supplied with memory which is both capacious and fast enough to support the required buffering and processing; typically, they feature distributed block RAM totaling less than a gigabyte and potentially massive but comparatively slow external memory interfaces. These are therefore a sub-optimal solution for the ngVLA Pulsar Engine as specified – the sub-bands would need to be divided into smaller portions to be processed, increasing the hardware

cost and power consumption of the system.

**HBM FPGAs:** The addition of multi-gigabyte HBM to a powerful general-purpose FPGA produces a system which is extremely well-aligned with the requirements for ngVLA pulsar processing. Several commercial standalone and PCIe host boards exist which appear fully capable of pulsar processing for a full ten-beam sub-band, and which have retail prices which are not significantly higher than equivalent non-HBM FPGA hardware. As a modification of general FPGA-based PSE plans, this permits much more efficient use of resources and therefore significantly lower hardware and power costs.

**TALON-DX:** The TALON-DX board features an Intel Stratix 10 SX FPGA with significant programmable logic and data transfer resources; it is effectively a general-purpose FPGA host board, although designed for a large-scale radio beamforming task with different parameters than ngVLA. The total available I/O is similar to that envisioned for an ngVLA-custom FPGA host, but is divided between multiple form factors. Due to the incorporation of a passive optical mesh into the TRIDENT design, the TALON board features only two 100 GbE QSFP28 connectors, with the majority of its I/O taking the form of specialized mid-board optical connections.

It additionally lacks HBM or similar high-speed RAM; it is unclear if the three on-board DDR4 DIMMs would provide sufficient memory bandwidth for one node to handle processing for an entire sub-band. If this option is to be considered further, it must be after additional testing to ensure that it is capable of the required processing.

**NRAO Custom ASIC:** As detailed in [D’Addario & Wang \(2016\)](#), the power consumption of an ASIC may be significantly less than that of comparable general-purpose hardware. This comes with extremely high development and verification costs, the precise extent of which is presently unclear. Implementation would likely also require the ASICs to be hosted by FPGA data transfer and control units, requiring their own hardware selection and development. At present, it is not believed that the power savings relative to an FPGA-based solution would justify the additional costs.

### 5.3. *Rationale for Selection of Hardware for Conceptual Design*

The design option currently under exploration ([Denman et al. 2021a](#)) is a board hosting an HBM FPGA with additional network interfaces hosted either directly or via FMC+ modules. This exposes the extremely high data transfer bandwidth which is required for the ngVLA, while minimizing the quantity of hardware required for pulsar operations.

An additional benefit is that when combined with the recommendations of the accompanying X-Engine trade study (an AI-optimized FPGA on a similar host board, see [Denman et al. \(2021b\)](#)) it may be feasible to select package-compatible HBM and AI FPGAs which permit a unified host board, resulting in a significant reduction in design costs. The release of future FPGAs featuring both HBM and accelerator cores (likely, given current industry trends) would potentially permit convergence on a single hardware platform for all sub-band processing.

## REFERENCES

- |  |  |
|--|--|
| <p>Amestica, R., Hiriart, R., &amp; Brandt, P. 2021, Software Requirements for RFI Management, Tech. Rep. C3, National Radio Astronomy Observatory</p> | <p>Bassa, C. G., Pleunis, Z., &amp; Hessels, J. W. T. 2017, <i>Astronomy and Computing</i>, 18, 40</p> |
|--|--|

- BittWare Inc. 2021a, 520N-MX Hardware Reference Guide, Tech. Rep. NT107-0519, BittWare Inc
- 2021b, 520N PCIe FPGA Board Datasheet, Tech. rep., BittWare Inc
- 2021c, XUP-VV8 PCIe FPGA Board Datasheet, Tech. rep., BittWare Inc
- 2021d, XUP-VVH Hardware Reference Guide, Tech. rep., BittWare Inc
- Carlson, B., & Pleasance, M. 2018, TRIDENT Correlator-Beamformer for the ngVLA: Preliminary Design Specification, Tech. Rep. TR-DS-000001, National Research Council Canada
- D’Addario, L. R., & Wang, D. 2016, *Journal of Astronomical Instrumentation*, 5, 1650002
- Demorest, P., & Ransom, S. 2018, ngVLA Time-domain Correlator Considerations, Tech. rep., National Radio Astronomy Observatory
- Denman, N., et al. 2021a, A SCREAM-Compatible ngVLA Pulsar Engine: Design Description, Tech. Rep. TBA, National Radio Astronomy Observatory
- 2021b, ngVLA Electronics Memo #10: A SCREAM-Compatible ngVLA Cross-Correlation Engine: Key Requirements Review and Option Trade-Off Study, Tech. Rep. E10, National Radio Astronomy Observatory
- Hemmert, K., & Underwood, K. 2005, in 13th Annual IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM’05), 171–180
- HiTech Global LLC. 2021, 6-Port QSFP28 (6x100G) FMC+ Module, <http://www.hitechglobal.com/FMCModules/x6QSFP28.htm>
- Intel Corporation. 2020a, Intel Stratix 10 GX/SX Device Overview, Tech. Rep. S10-OVERVIEW, Intel Corporation
- 2020b, Intel Stratix 10 MX (DRAM System-in-Package) Device Overview, Tech. Rep. S10-MX-OVERVIEW, Intel Corporation
- 2021, Intel Stratix 10 Device Datasheet, Tech. Rep. S10-DATASHEET, Intel Corporation
- Kettenis, M., Phillips, C., Sekido, M., & Whitney, A. 2014, VLBI Data Interchange Format (VDIF) Specification, Tech. Rep. 1.1.1, VDIF Task Force
- Koski, W., Baca, J., & Durand, S. 2019, ngVLA Monitor & Control Hardware Interface Layer: Reference Design Description, Tech. Rep. 020.30.45.00.00-0004-DSN-A, National Radio Astronomy Observatory
- Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *AJ*, 129, 1993
- Murphy, E., et al. 2019, ngVLA: Science Requirements, Tech. Rep. 020.10.15.00.00-0001-REQ-B, National Radio Astronomy Observatory
- NANOGrav Collaboration. 2018, Pulsar Timing Array Requirements for the ngVLA, Tech. Rep. 42, National Radio Astronomy Observatory
- NVIDIA Corporation. 2017, NVIDIA Tesla V100 GPU Architecture, Tech. Rep. WP-08608-001, NVIDIA Corporation
- 2020a, NVIDIA A100 Tensor Core GPU Architecture, Tech. rep., NVIDIA Corporation
- 2020b, NVIDIA DGX A100 Datasheet, Tech. rep., NVIDIA Corporation
- 2020c, NVIDIA HGX A100 Datasheet, Tech. rep., NVIDIA Corporation
- 2020d, NVIDIA Jetson AGX Xavier Developer Kit Carrier Board Specification, Tech. Rep. SP-09778-001, NVIDIA Corporation
- 2020e, NVIDIA Jetson AGX Xavier Series OEM Product Design Guide, Tech. Rep. DG-09840-001, NVIDIA Corporation
- 2020f, NVIDIA Jetson AGX Xavier Series System-on-Module Data Sheet, Tech. Rep. DS-09654-001, NVIDIA Corporation
- 2021, NVIDIA Ampere GA102 GPU Architecture, Tech. rep., NVIDIA Corporation
- Ojeda, O. 2018, ngVLA Central Signal Processor: Preliminary Reference Design, Tech. Rep. 020.40.00.00.00-0002-DSN-03, National Radio Astronomy Observatory
- 2020a, Trident 2.0 Concept: A Minimum Delta Update to the Central Signal Processor Reference Design, Tech. Rep. E04, National Radio Astronomy Observatory
- 2020b, Trident 2.1 Concept: Updates to the CSP Reference, Tech. Rep. E05, National Radio Astronomy Observatory
- Ojeda, O., Lacasse, R., Selina, R., et al. 2019, ngVLA Central Signal Processor: Preliminary Technical Requirements, Tech. Rep. 020.40.00.00.00-0001-REQ-A, National Radio Astronomy Observatory



- Pleasance, M., Zhang, H., Carlson, B., et al. 2017, in 2017 XXXIInd General Assembly and Scientific Symposium of the International Union of Radio Science (URSI GASS), 1–4
- Rau, U., Selina, R., & Erickson, A. 2019, RFI Mitigation for the ngVLA: A Cost-Benefit Analysis, Tech. Rep. 70, National Radio Astronomy Observatory
- Selina, R., Rau, U., Hiriart, R., & Erickson, A. 2020a, RFI Mitigation for the ngVLA: A Cost-Benefit Analysis, Tech. Rep. 71, National Radio Astronomy Observatory
- Selina, R., et al. 2020b, ngVLA: System Requirements, Tech. Rep. 020.10.15.10.00-0003-REQ, National Radio Astronomy Observatory
- van Straten, W. 2003, PhD thesis, Swinburne University Of Technology
- van Straten, W., & Bailes, M. 2011, PASA, 28, 1
- van Straten, W., & Tiburzi, C. 2017, ApJ, 835, 293
- Xilinx Inc. 2018a, VCU118 Evaluation Board User Guide, Tech. Rep. UG1224, Xilinx Inc
- . 2018b, VCU128 Evaluation Board User Guide, Tech. Rep. UG1302, Xilinx Inc
- . 2021a, UltraScale Architecture and Product Data Sheet: Overview, Tech. Rep. DS890, Xilinx Inc
- . 2021b, UltraScale+ FPGAs: Product Tables and Product Selection Guide, Tech. Rep. XMP103, Xilinx Inc