



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B



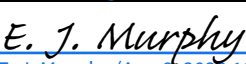



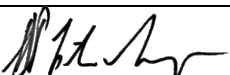
System Concept Options & Trade-Offs

020.10.25.00.00-0005-REP

Status: **Released**

PREPARED BY	ORGANIZATION	DATE
R. Selina et al.	Electronics Division, NRAO	2022-08-05

APPROVALS	ORGANIZATION	SIGNATURES
T. Kusel, System Engineer	ngVLA, NRAO	 Thomas Kusel (Aug 8, 2022 18:00 EDT)
R. Selina, Project Engineer	ngVLA, NRAO	 R. Selina (Aug 8, 2022 16:13 MDT)
E. Murphy, Project Scientist	ngVLA, NRAO	 E. J. Murphy (Aug 8, 2022 18:19 EDT)
W. Esterhuyse, Project Manager	ngVLA, NRAO	

RELEASED BY	ORGANIZATION	SIGNATURE
W. Esterhuyse, Project Manager	ngVLA, NRAO	



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

Change Record

Version	Date	Author	Affected Section(s)	Reason
01	2021-03-04	R. Selina	All	Started first draft outline.
02	2021-03-05	R. Selina	4.3.2, 4.6.3	Added sections recommended by TK.
03	2021-06-04	R. Selina	All	In-progress first draft. Circulated to EM, TK, RH, OO and BS for comment.
04	2021-06-15	R. Selina, O. Ojeda	4.4.5, 4.6, 4.7, 5	Updates to CSP section by OO. Updates to post processing discussion by RS. Added new sections on cryogenics, open trades.
05	2021-06-24	N. Denman	4.6	Updates to CSP section.
06	2021-10-12	R. Selina	3, 4.5	Updates to reference document numbers. Updates to antenna LO section based on LO trade study report by BS.
07	2021-12-28	R. Selina	All	Minor corrections and additions throughout.
08	2022-05-15	R. Selina	All	Updates to address RIDs from team review.
A	2022-06-01	R. Selina	3.2, 4.3.4	Final edits for release. Added reference to feed indexer geometry study.
B	2022-08-05	R. Selina	4.1.1, 4.1.2, 4.3.1, 4.3.2, 4.3.3, 4.3.4, 4.4.1, 4.4.3, 4.4.5, 4.5.1, 4.7.2.	Revising post Sys CDR to address RIDs.



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

Table of Contents

1	Introduction	4
1.1	<i>Purpose</i>	4
1.2	<i>Scope</i>	4
2	Overview of the System Concepts and Options	4
3	Related Documents and Drawings	5
3.1	<i>Applicable Documents.....</i>	5
3.2	<i>Reference Documents.....</i>	5
4	System Options.....	8
4.1	Array Configuration.....	8
4.1.1	Heterogeneous vs Homogeneous Array	8
4.1.2	Fixed Stations vs Reconfigurable Array.....	9
4.2	Array Calibration	11
4.2.1	Tropospheric Calibration Strategies	11
4.2.2	Instrumental Calibration	12
4.3	Antenna	13
4.3.1	Aperture Size.....	13
4.3.2	Field of View vs Sensitivity: Conic Optics & Shaped Optics	14
4.3.3	Symmetric Cassegrain vs Offset Gregorian Optics.....	17
4.3.4	Feed Indexer Geometry.....	19
4.3.5	Multi-band vs Single-band Operation.....	22
4.3.6	Altitude-Azimuth vs Equatorial Mounts & Rotation Axes.....	25
4.4	Antenna Electronics.....	26
4.4.1	Receiver Band Definition	26
4.4.2	Linear vs Circular Polarization Basis	29
4.4.3	Bit Depth	30
4.4.4	Downconversion Topology.....	32
4.4.5	Cryogenics System	33
4.5	Time & Frequency Reference Distribution	35
4.5.1	Antenna Local Oscillator Architectures.....	35
4.5.2	Frequency Reference Distribution	37
4.5.3	Timing Reference Distribution	38
4.6	Central Signal Processor.....	38
4.6.1	FX vs XF Correlator Architecture.....	38
4.6.2	Correlator-Beamformer Integration vs Independence.....	39
4.6.3	ASICs vs FPGA vs GPU Implementations.....	40
4.6.4	Phase-Delay vs True-Delay Beamformer	42
4.6.5	Additional Design Considerations.....	43
4.7	Post Processing System	45
4.7.1	Real-Time Processing vs Post-Processing.....	45
4.7.2	Low-Level vs High-Level Data Products	46
4.7.3	Observatory-Run vs Cloud Compute Resources.....	47
5	Baselined & Open Trades Summary	48
6	APPENDIX.....	50
6.1	Abbreviations and Acronyms.....	50



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

I Introduction

1.1 Purpose

This document presents a summary of the conceptual options and trade-offs applicable to the ngVLA system design. Over the course of the facility development phase and conceptual design phase, a large parameter space has been explored to ensure not only that the system supports the science and stakeholder requirements (AD01, AD02), but also that the design makes efficient use of resources and provides a flexible and extensible platform to adapt to future scientific interests and programs over the life of the instrument.

1.2 Scope

The scope of this document is the entire ngVLA facility, from the reception of external signals through to the storage of data products in the archive. The full operational model of the facility is reflected, from the preparation and submission of proposals, to the execution of an observation, and the delivery of data products to users.

This document responds to the trade space available after the derivation and baselining of the LI System Requirements. Decisions which are made as part of the Operations Concept (AD03), or otherwise precluded by the Level-0 Requirements and associated context documents, are not further considered in this report.

2 Overview of the System Concepts and Options

This document does not attempt to justify all trades. Rather, it provides an aggregation of the outcomes of the trades considered from 2015 through 2021 as the facility concept has matured, providing a reader a roadmap through the key technical decisions that informed the selected system concept. Detailed analysis can be found in supporting technical documentation and engineering memos referenced in the text.

This report will refer to the facility development phase, from 2015 through 2019, as the period where the facility concept was developed and formalized as the Reference Design presented to the decadal survey. A key milestone during this project phase was a Science and Technical Workshop held in Socorro in 2017, with broad participation from the user community. A number of the key trades described in this report were endorsed by our users at this workshop, and subsequently incorporated into the reference design that was the technical basis for the ngVLA Science Book performance estimates. The conceptual design phase refers to development in 2019 through 2021, building upon this facility concept towards the conceptual design baseline presently under review.

This document is roughly organized into chapters corresponding to the major sub-systems found in the system architecture (AD14). We start with the array as an integrated unit and consider its configuration and calibration (Sections 4.1, 4.2). We then progress to the antenna, antenna electronics, reference signals, central signal processor, and post-processing systems (Sections 4.3 through 4.7). Conceptual options and trades considered for each system are then presented in corresponding sub-sections. Where trades may impact other elements of the system, these relationships are noted.

The selected system conceptual design is documented in AD13.



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

3 Related Documents and Drawings

3.1 Applicable Documents

The following documents are applicable to this report to the extent specified. In the event of a conflict between the documents referenced herein and the content of this report, precedence is indicated in the table below as either “This doc”, indicating that this document takes precedence, or “Ref doc”, in which case the reference document takes precedence.

Reference No.	Document Title	Precedence	Rev/Doc. No.
AD01	L0 Science Requirements	Ref doc	020.10.15.00.00-0001-REQ
AD02	L0 Stakeholder Requirements	Ref doc	020.10.15.01.00-0001-REQ
AD03	Operations Concept	Ref doc	020.10.05.00.00-0002-PLA
AD04	L1 System Requirements	Ref doc	020.10.15.10.00-0003-REQ
AD05	L1 Environmental Specification	Ref doc	020.10.15.10.00-0001-SPE
AD06	L1 System EMC and RFI Mitigation Requirements	Ref doc	020.10.15.10.00-0002-REQ
AD07	L1 Safety Specification	Ref doc	020.80.00.00.00-0001-REQ
AD08	L1 Security Specification	Ref doc	020.80.00.00.00-0003-REQ
AD09	Reference Observing Program	Ref doc	020.10.15.05.10-0001-REP
AD10	Envelope Observing Program	Ref doc	020.10.15.05.10-0002-REP
AD11	Assembly, Integration and Verification Concept	Ref doc	020.10.05.00.00-0005-PLA
AD12	Commissioning and Science Validation Concept	Ref doc	020.10.05.00.00-0006-PLA
AD13	System Conceptual Design Description	Ref doc	020.10.20.00.00-0005-REP
AD14	Preliminary System Architecture	Ref doc	020.10.20.00.00-0002-DWG
AD15	Legacy Science Program	Ref doc	020.10.05.00.00-0004-PLA

3.2 Reference Documents

The following documents are referenced in the text and provide supporting context on the analyzed concepts and their associated trade-offs.

Reference No.	Document Title	Rev/Doc. No.
RD01	ngVLA Quantitative Exchange (Cost) Model Memo & Spreadsheet	NQXM V3.0, February 24, 2017
RD02	The Concept of a Reference Array for ngVLA	ngVLA Memo #4
RD03	Short Spacing Considerations for the ngVLA	ngVLA Memo #14
RD04	ngVLA Reference Design Development & Performance Estimates	ngVLA Memo #17
RD05	Summary of the Science Use Case Analysis	ngVLA Memo #18
RD06	The ngVLA Short Baseline Array	ngVLA Memo #43
RD07	System-level Evaluation of Aperture Size	ngVLA Antenna Memo #2



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

Reference No.	Document Title	Rev/Doc. No.
RD08	Fast Switching Phase Calibration at 3mm at the VLA Site	ngVLA Memo No. 1
RD09	Calibration Strategies for the Next Generation VLA	ngVLA Memo No. 2
RD10	Considerations for a Water Vapor Radiometer System	ngVLA Memo #10
RD11	Temporal and Spatial Tropospheric Phase Fluctuations at the VLA	ngVLA Memo #61
RD12	ngVLA Calibration Requirements	020.22.00.00.00-0001-REQ
RD13	Antenna Technical Requirements (L2)	020.25.00.00.00-0001-REQ
RD14	Water Vapor Radiometer: Design Description	020.45.00.00.00-0002-DSN
RD15	Antenna Optical Reference Design Report	020.25.01.00.00-0001-REP
RD16	EMSS Optical Design Report	EA-NGV-DR-05
RD17	System-level Cost Comparison of Offset and Symmetric Optics	ngVLA Antenna Memo #1
RD18	Antenna Optical Design Alternatives	ngVLA Antenna Memo #3
RD19	18m Antenna Conceptual Design Description	1021006-REP-21-000000-0001
RD20	Astrometrically Registered Simultaneous Observations of the 22 GHz H ₂ O and 43 GHz SiO Masers Towards R Leonis Minoris Using KVN and Source/ Frequency Phase Referencing	Dodson, R. et al (2014)
RD21	The Power of Simultaneous Multifrequency Observations for mm-VLBI: Astrometry up to 130 GHz with the KVN	Rioja, M. et al (2015)
RD22	The Power of (Near) Simultaneous Multi-Frequency Observations for mm-VLBI and Astrometry	Rioja, M. et al. Galaxies, 5, 9, (2017). doi:10.3390
RD23	Precise radio astrometry and new developments for the next generation of instruments	Rioja, M. & Dodson, R. The Astronomy and Astrophysics Review, 28, 6 (2020).
RD24	Observing Mode Framework	020.10.05.05.00-0005-PLA
RD25	Front End: Design Description	020.30.06.00.00-0006-DSN
RD26	ngVLA Receiver Cascade Analysis Tool	020.30.05.00.00-0004-GEN
RD27	Optimal Frequency Ranges for sub-microsecond precision pulsar timing	Lam, et al. (2018) https://arxiv.org/pdf/1710.02272.pdf
RD28	Headroom, Dynamic Range, and Quantization Considerations	ngVLA Electronics Memo #8
RD29	Antenna Requirements for LEO Satellite Mitigation	020.10.25.00.00-0004-MEM
RD30	ngVLA Radio Frequency Interference Forecast	ngVLA Memo #48
RD31	RFI Mitigation in the ngVLA System Architecture	ngVLA Memo #71
RD32	Integrated Receivers and Digitizers Design Description	020.30.15.00.00-0004-DSN
RD33	Downconversion and Digitization Methodology for the ngVLA	ngVLA Electronics Memo #1
RD34	An Integrated Receiver Concept for the ngVLA	ngVLA Memo #29
RD35	Central Signal Processor Design Description	020.40.00.00.00-0005-DSN



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

Reference No.	Document Title	Rev/Doc. No.
RD36	An Integrated Circuit for Radio Astronomy Correlators Supporting Large Arrays of Antennas	D'Addario & Wang (2016) https://doi.org/10.1142/S2251171716500021
RD37	Key Science Goals for the ngVLA	ngVLA Memo #19
RD38	Size of Computing Estimates for the ngVLA	ngVLA Computing Memo #4
RD39	System Architecture: Conceptual Design	020.10.20.00.00-0002-REP
RD40	EMSS Deformed Optical Analysis Report	EA-NGV-DR-06
RD41	VLBA Observing Manual: Frequency Bands and Performance	https://science.nrao.edu/facilities/vlba/docs/manuals/oss/bands-perf
RD42	Mid.CBF Design/Build-to-Cost: A Frequency Slice Approach – as submitted to the SKA Office	TALON SKAI MID.CBF Memo 0004
RD43	Trident 2.1 Concept: Updates to the CSP Reference Design	ngVLA Electronics Memo #5
RD44	A SCREAM-Compatible ngVLA Pulsar Engine: Key Requirements Review and Option Trade-Off Study	ngVLA Electronics Memo #11
RD45	A GPU based X-Engine for the MeerKAT Radio Telescope	Callanan (2020) Master's thesis, University of Cape Town
RD46	Trident 2.0 Concept: A Minimum Delta Update to the CSP Reference Design	ngVLA Electronics Memo #4
RD47	Cryogenics System Design Description	020.30.10.00.00-0007-DSN
RD48	Advanced Cryocoolers for the Next Generation VLA	ngVLA Memo #24
RD49	Thermoacoustic Stirling Cryocooler and Variable Speed Gifford McMahon Cryocooler Trade Study No. 2	ngVLA Electronics Memo #4
RD50	A SCREAM-Compatible ngVLA Cross-Correlation Engine: Key Requirements Review and Option Trade-Off Study	ngVLA Electronics Memo #10
RD51	ngVLA Antenna Local Oscillator Trade Study	020.30.35.00.00-0003-REP
RD52	An Optimal 18 m Shaped Offset Gregorian Reflector for the ngVLA Radio Telescope	Lehmensiek and de Villiers, IEEE Trans. A&P, Vol 69 No. 12, Dec. 2021
RD53	Front End Trade Study	ngVLA Electronics Memo #13 (<i>In prep.</i>)
RD54	Status Levels for New ngVLA PI Observing Modes	020.10.05.05.00-0004-PLA
RD55	Water Vapor Radiometer: Technical Requirements	020.45.00.00.00-0001-REQ
RD56	LO Reference & Timing Design Description	020.35.00.00.00-0004-DSN
RD57	Rotating FE Design: Benefits and Constraints	020.30.05.00.00-0007-REP
RD58	The Total Power Array Concept of Operations & System Level Requirements	020.27.00.00.00-0001-REQ



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

4 System Options

4.1 Array Configuration

4.1.1 Heterogeneous vs Homogeneous Array

Trades between heterogeneous and homogeneous array configurations have been considered as part of the facility concept development. This trade is closely related to the choice of aperture size, described in Section 4.3.1.

The science case development suggested that mapping speed in Band 6 could be an important capability, especially at lower angular resolutions that correspond to baselines from hundreds of meters to a few km. This suggested a smaller antenna aperture size in the core could be attractive scientifically, if it would not compromise total sensitivity and array cost. Early studies into atmospheric phasing strategies suggested a heterogeneous array with smaller apertures operating as a phase reference array could also be attractive (RD02).

Factors favoring a homogeneous array largely centered on development cost and maintenance cost, as multiple designs increases total development and construction effort (less learning curve gains), as well as requiring more spare parts inventory and maintenance team training. System complexity also increases downstream, with calibration and imaging complicated by the differences in field of view and primary beam response, including polarization patterns, when operated as a single array. Relative sensitivity differences between the two arrays further complicate data processing and quality assurance, as two different antennas means three different levels of baseline visibility weights and two different amplitude gain factors, complicating the display and statistics of gains and the heuristics for the automated identification of outliers.

The cost relationship between aperture sizes and the number of array elements that could be constructed within both construction and operations cost caps had to be established to inform the trade. In 2016-2017 a parametric cost and performance modeling tool was built (RD01) to explore the relationship between system cost, system performance, and a multitude of configuration parameters. The software system size and cost proved very sensitive to the number of apertures in the array and field of view in mapping cases. Operations cost constraints also proved limiting for array designs with many antennas of smaller aperture.

A homogeneous array of small diameter apertures ($D < 15\text{m}$) and many elements proved infeasible due to the operations cost cap. While accounting for the pointing and surface accuracy required for Band 6 operation, sensitivity was maximized with apertures of approximately 22m, while survey speed favored 18m apertures and imaging fidelity favored 17m apertures once the lifecycle costs were considered. (RD07)

A small aperture core would require a heterogeneous array with larger apertures on the 1 to 1000 km scales to retain total system sensitivity while respecting the operations cost cap. The imaging fidelity losses at higher angular resolution would be significant due to the loss in the number of array elements compared to a heterogeneous array of 18m aperture.

The parametric modeling tool was used to prepare two options for the Science and Technology workshop held in 2017 (RD04). A homogenous array of 18m apertures was compared to a heterogeneous array with a compact core of 13m apertures and 26m apertures on 1-1000km baselines. These aperture sizes were selected to retain a comparable number of total apertures in the array, which is the primary driver of the operations cost estimate. Consensus from the user community was a clear preference for the



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

homogeneous 18m array when looking at a suitable suite of use cases (RD05). As a result of this feedback, a homogeneous main array of 18m apertures was adopted for the reference design.

Subsequent refinement of the Science Requirements identified the need to recover large scale structure beyond the capabilities of the main array. Baselines shorter than 11m would be required, and a total power capability would be desirable to achieve imaging fidelity targets. Based on the previously analyzed constraints on the main aperture size, only two viable options were available to meet this requirement: The inclusion of a large (50-100m) single dish antenna or a short baseline array and possibly a total power capability on a subset of the 18m antennas. Use of an existing 100m antenna (e.g., GBT) would meet our needs, but the lifetime of the instrument is uncertain and the overhaul costs appear prohibitive. The costs of new construction were compared with the parametric modeling tool for both the single-dish and short baseline array alternatives, and the inclusion of a short baseline array of 6m apertures is appreciably lower cost. Based on a combination of cost and calibration considerations, the short baseline array of 6m apertures has been adopted as the baseline solution to recover larger scale structure. The total power measurement capability on a subset of 18m antennas is still under review and the requirements and supporting concept [RD58] for the total power capability will be finalized for the system PDR.

Please consult the following materials for further information on this trade:

The ngVLA Quantitative Exchange Model	nqxm Ver3.0
The Concept of a Reference Array for ngVLA	ngVLA Memo #4
Short Spacing Considerations for the ngVLA	ngVLA Memo #14
ngVLA Reference Design Development & Performance Estimates	ngVLA Memo #17
Summary of the Science Use Case Analysis	ngVLA Memo #18
The ngVLA Short Baseline Array	ngVLA Memo #43
System-level Evaluation of Aperture Size	ngVLA Antenna Memo #2
The Total Power Array Concept of Operations & Requirements	020.27.00.00.00-0001-REQ

4.1.2 Fixed Stations vs Reconfigurable Array

Reconfiguration is a key feature of the VLA. The system collecting area can be placed on a narrow range of baselines, ensuring sensitivity to an appropriate range of angular scales on sky. E.g., VLA A-Config has baselines from 0.8km to 36.6km, with a ratio of 46 between B_{MAX} and B_{MIN} (Figure 1). The naturally weighted synthesized beam is still non-Gaussian, and still requires robust weighting for optimal imaging fidelity, but the associated losses in sensitivity compared to a naturally-weighted synthesized beam are manageable. The array is reconfigured three times a year, cycling through its available configurations over 1.3 years.

The VLBA by comparison uses fixed antenna locations. The obvious difference between the two facilities that leads to this are the vast distances between the VLBA antennas. Reconfiguration over rail for 20km distances is practical, but reconfiguration over 200km or 2000 km scales is not.

Given that ngVLA requirements lead to main array baselines extending out 100s of km from the core, with extended (VLB) baselines out to 8860 km, reconfiguration would only be practical for the antennas within ~30km from the core.



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

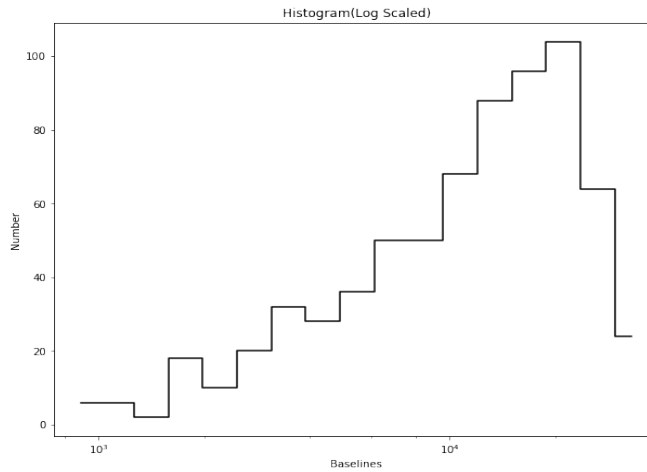


Figure 1 - Distribution of baselines in the VLA A-Configuration, in meters.

From a performance perspective, the cost of the additional foundations and reconfiguration system infrastructure can be compared to the additional collecting area that could be purchased with equivalent funds for a fixed station array. This analysis was performed with the parametric cost model, with the conclusion that reconfiguration capability has a comparable cost to 20% of the fixed collecting area. Should the difference in the geometric mean of the beam sculpting weights between the two scenarios exceed 20%, one might prefer a reconfigurable array (with lower beam sculpting weights) for that use case. (RD01).

However, the overall trade is more complicated. A use case that requires a range of angular scales would need to take observations in each configuration, which could take over a year, and some variable sources change in structure on these time scales. Beam shape for the extended baselines would be impacted in a mixed reconfiguration mode, reducing overall beam sculpting efficiency beyond the array core and partially offsetting the apparent benefits for reconfiguration when a suite of use cases are considered. Finally, reconfiguration is also a significant cost and complexity driver in operations. Offsetting these costs would require larger (and less) apertures in the array, negatively impacting imaging fidelity.

Ultimately, the benefit of multi-scale observations, recovering all scales in a single configuration and observation, was seen as a key feature by the science community, and the fixed station array with greater total collecting area was endorsed by the participants of the 2017 Science and Technical Workshop. This conclusion is captured in ngVLA Memo #17 and 18 (RD04, RD05) and adopted as a key feature of the facility concept presented as part of the Astro2020 reference design. This decision is retained for the conceptual design baseline.

Please consult the following materials for further information on this trade:

The ngVLA Quantitative Exchange Model
ngVLA Reference Design Development & Performance Estimates
Summary of the Science Use Case Analysis

nqxm Ver3.0
ngVLA Memo #17
ngVLA Memo #18



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

4.2 Array Calibration

4.2.1 Tropospheric Calibration Strategies

Tropospheric delay/phase calibration is a key calibration overhead for a cm to mm-wave array that can influence total observational efficiency, and the post-calibration residual errors can contribute to dynamic range limits in imaging. ngVLA has selected three complementary strategies for tropospheric calibration: Self-calibration, fast position switching, and water vapor radiometry. Each of these techniques will be supported in the design, providing improved observational efficiency while also mitigating the technical risks inherent in any one of the solutions.

Early in the facility concept development phase, a broad assessment of the available calibration strategies was performed and is summarized in the project memo series (RD02, RD08, RD09, RD10, RD11). In addition to the selected strategies, paired-element calibration was considered, along with a reference array (RD02) – a calibration array of smaller aperture antennas used solely for phase calibration.

Paired-element calibration was abandoned due to the high loss of effective sensitivity. Only half the array collecting area is observing the science target at a time, with effective losses comparable to fast position switching at high frequency and appreciably larger losses at lower frequencies or in stable atmospheric conditions.

The reference array strategy would have required extrapolating solutions from lower frequency receivers on the reference antenna to the science antenna, which is a significant technical risk. The reference antenna either has to be mounted to the 18m antenna mount, while still providing beam steering capabilities, or at a sufficient distance to avoid shadowing in most orientations with an independent mount. The former solution had technical complexities (large mass due to active position mechanisms, an asymmetric load on the mount, etc.) The latter solution has calibration residuals proportional to the baseline length. The extrapolation across frequency can be compromised by frequency dependent effects (e.g., ionospheric vs tropospheric contributions to delay) and errors in measurement are also multiplied at higher frequencies. The overall cost of the foundations, mounts and electronics was also cost prohibitive, leading to this concept being abandoned.

In-beam self-calibration is a computationally intensive but proven method of correcting for atmospheric effects when a sufficiently bright source (either the science target or a background source) is present in the primary beam. The highest dynamic range images from both the VLA and ALMA are generated using in-beam self-calibration techniques. While the sensitivity of the ngVLA will increase the availability of suitable calibrators, in-beam calibrators are still only infrequently present in-beam at high frequency, where the tropospheric effects on propagation are largest, so this technique cannot be relied upon as a sole method of tropospheric calibration.

The fast position switching is a reliable and proven strategy, employed on both the VLA and ALMA, but it does result in significant time spent off the science target. Although fast switching reduces the residual calibration error due to time variations in the atmosphere, it does not reduce the residual due to directional differences in the atmosphere. That is why subsequent self-calibration on the science target can further reduce the calibration error. The dynamic range that can be achieved in imaging by fast position switching is dictated by the total cycle time (from source to calibrator and back) and the structure function of the atmosphere. The antenna requirements include a slew and settling time restricted to 7 second for up to a 3-degree position change (RD12, RD13) in support of this calibration mode.

Position-switched astronomical calibration (effectively self-calibration on the astronomical calibrator, with solutions transferred to the science target) will generally be more effective on ngVLA than current arrays due to the system sensitivity which increases the density of astronomical sources that can be used for



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

calibration. The high angular resolution of the array will also resolve or partially resolve many calibrators, and suitable calibrator visibility models will need to be developed and included in the self-calibration algorithms to enable the derivation of complex gain solutions on all angular scales. Performing such a survey and generating these models will be a key commissioning activity. Accounting for these factors should provide a suitably strong calibrator within 2-degrees of a science target for most fields, and almost always within 4 degrees, at all bands. (RD09, RD12). The antenna slew and settle specifications and this dense grid of calibrators enable up to a 30-second cycle time (RD02, RD08) in position switching, albeit at a compromised observing efficiency of 50% to 35%. Only at a 3 minute cycling scale, which is too long at high frequency to freeze the phase fluctuations using position switching alone, is a 90% observing efficiency possible.

In an effort to provide a supplement to fast position switching, lengthening the position switching cycle, 22.2 GHz water vapor radiometers (WVRs) will be installed on each 18m antenna. These WVRs are broadband digital spectrometers, spanning the water line and up beyond 30 GHz to sample the oxygen line wings. The baseline implementation of the digital spectrometer WVR will utilize a standalone reflector and signal chain optimized for gain stability (RD14). However, an open trade study is exploring leveraging the Band 4 receiver and off-axis response of the antenna optics.

The WVR is intended to extend this astronomical calibration cycle, improving overall observing efficiency for a majority of cases and reducing post-calibration residual errors. It is the only solution that can provide sufficiently fast estimates of atmospheric delay to possibly enable the project to achieve the dynamic range requirements in the absence of self-calibration. The supporting calculations to demonstrate these findings are available in the Calibration Requirements (RD12). However, 22 GHz WVRs have not been demonstrated to operate reliably at the requisite precision to support these imaging dynamic range requirements, so further development of the associated calibration algorithms and redundant calibration strategies are warranted. Even if the WVR system fails to achieve the full desired specification, the inclusion of the WVR may enable the partial correction of the atmosphere in otherwise marginal conditions, enabling higher frequency observations (at relatively low dynamic range) in environmental conditions that would lead present arrays to be restricted to lower frequency observations.

Please consult the following materials for further information on this trade:

Fast Switching Phase Calibration At 3mm at the VLA Site	ngVLA Memo #1
Calibration Strategies for the Next Generation VLA	ngVLA Memo #2
The Concept of a Reference Array for ngVLA	ngVLA Memo #4
Considerations for a Water Vapor Radiometer System	ngVLA Memo #10
Temporal and Spatial Tropospheric Phase Fluctuations at the VLA (and Beyond) and Implications for Phase Calibration	ngVLA Memo #61
ngVLA Calibration Requirements	020.22.00.00.00-0001-REQ
Antenna Technical Requirements	020.25.00.00.00-0001-REQ
Water Vapor Radiometer: Technical Requirements	020.45.00.00.00-0001-REQ
Water Vapor Radiometer: Design Description	020.45.00.00.00-0002-DSN

4.2.2 Instrumental Calibration

The Operations Concept for ngVLA establishes a model where high-level data products (such as calibrated image cubes) are generated by automated pipelines, tailored to each standard observing mode, and maintained by the project. The progression of automation and validation of these observing modes is discussed in the RD54. Users have interaction points with the pipeline and data analysis tools to query or extract the desired components of the high-level products. This is distinct from the VLA, where users are responsible for both data reduction (generating the high level data products) and data analysis.



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

The approach to instrumental calibration is a critical decision to support the automated generation of these data products. Experience in developing automated pipelines for calibration, imaging, and data quality assurance for both ALMA and the VLA suggest that developing these pipelines for existing telescope architectures can result in highly complex heuristics and that the overall system design can be optimized to reduce this complexity (and cost) in the post-processing system. The VLA is highly flexible in the configuration of individual observations, but this flexibility amplifies the degree of complexity in the data product pipelines, and limits the degree of automation that is achievable.

ngVLA had two fundamental choices: (a) retain this degree of flexibility in configuration (b) eliminate degrees of freedom in standard modes, ideally without a loss in capability. A follow up trade is how much flexibility to permit in non-standard modes, and should the system design be fundamentally different for a telescope that delivers high-level data products.

While not strictly trades, these are central design decisions that should be explained in context. In order to reduce total system complexity, an approach to antenna configuration that reduces tunability has been selected as a guiding principle. This does not necessarily mean a less flexible or capable system, but it can increase up-front costs in some cases.

This consideration is most obvious in the design of the antenna electronics up through the digitizers. E.g., the system antenna electronics concept aims to digitize the full bandwidth of each receiver band. Arguably this leads to excess digitization infrastructure and correlator processed bandwidth in some scenarios, but it eliminates the need for fully tunable local oscillators, and associated changes in bandpass behavior that are unique to an individual tuning. Digitizing the full band is highly performant and provides extended data products that can increase archive data reuse, all while limiting the variability in the analog system configuration and supporting pipeline automation. This choice also reduces the burden on proposing observers specifying their technical setup.

Other areas of the system concept that are influenced by the top-down design for pipelined data products include the downconverter architecture, the elimination of mechanical RF switches, limited (ideally fixed) LO tunings, and fixed antenna positions. The data model and ancillary metadata streams also account for the need for observatory-level calibration observations and the need for tuning parameters for calibration and flagging algorithms to account for the conditions present during the observation (e.g., the presence of known RFI).

Please consult the following materials for further information on this trade:

ngVLA Operations Concept	020.10.05.00.00-0002-PLA
System Conceptual Design Description	020.10.20.00.00-0005-REP
ngVLA Calibration Requirements	020.22.00.00.00-0001-REQ
RFI Mitigation in the ngVLA System Architecture	ngVLA Memo #71
Preliminary System Architecture	020.10.20.00.00-0002-REP
Status Levels for New ngVLA PI Observing Modes	020.10.05.05.00-0004-PLA

4.3 Antenna

4.3.1 Aperture Size

A system-level assessment of the aperture size is considered in ngVLA Antenna Memo #2 (RD07). The analysis respected the programmatic constraints for construction cost and operations costs, and then optimized on three key performance parameters (KPPs): Sensitivity (ND^2), Survey Speed (N^2D^2), and Imaging Fidelity (N^2). The analysis was performed using the quantitative exchange model (RD01).



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

The key findings of the study are summarized in Figure 2, which shows the relative performance of array designs that respect the cost constraints to the three key performance parameters. Imaging fidelity is optimized at 17.5m, survey speed at 18m, and sensitivity at 22m. The sensitivity curve is relatively flat up to 18m and then descends rapidly for smaller aperture sizes. The degradation below 17m aperture is not based on constraints within the construction budget, but is driven by the annual operations budget constraint. The operations budget has a strong dependence on N for a number of parameters (e.g. number of maintenance technicians) and N^2 for others (such as the computing and data archive costs). While values from 17.5m to 22m can be justified based on this analysis, an 18m aperture size provides the best balance of the three criteria.




Figure 2 - Performance to KPPs for modeled array designs that account for both the construction and operations cost constraints. Values on each KPP are independently normalized to be fractions of maximum performance on the respective metric (RD07).

Please consult the following materials for further information on this trade:

System-Level Evaluation of Aperture Size
The ngVLA Quantitative Exchange Model

ngVLA Antenna Memo #2
nqxm Ver3.0

4.3.2 Field of View vs Sensitivity: Conic Optics & Shaped Optics

A key choice in the antenna optics is to use conic (parabolic) reflector segments or optical shaping. Natural conics have more even efficiency over the focal plane, making them well suited to multi-pixel receivers such as focal plane arrays (FPAs) or phased-array feeds (PAFs) depending on frequency. Conics also degrade more gracefully when deformed and are best suited for reflector homology techniques. However, the illumination efficiency achievable with conics is constrained by the feed pattern – a Gaussian feed pattern limits the combined illumination efficiency to approximately 75%, while an optimized feed pattern (with more of a ‘flat top’) can approach 80% efficiency.

In a receiving antenna, optical shaping redirects rays from the edge of the main reflector towards the center of the subreflector. More rays are received near the boresight of the feed where the forward gain is highest. The illumination efficiency is improved in such cases, with illumination efficiency greater than 95% achievable (RD15). However, this gain in illumination efficiency is applicable at the secondary focus

only - the efficiency drops for off-axis feeds, so this solution is suitable for single-pixel receiver solutions only.

The trade-off in efficiency vs focal plane size is illustrated in Figure 3 - an unshaped system is contrasted with three different shaping functions. At the boresight, the simulated feed achieves approximately 68% efficiency when used with conic reflectors and set for a 15 dB edge taper. The curves are nearly flat though, so off-axis response is only marginally degraded compared to the on-axis response. The highest gain shown, shaping solution S3, achieves a sensitivity of 93% with the same feed, but the gain degrades rapidly, as evident by the steeper slope of the efficiency curve. Note that neither of these solutions is optimized, but the difference in performance is representative of the parameter space.




Figure 3 - illumination efficiency achieved with unshaped (conic) optics and three different shaping profiles using a real (simulated) feed with a 55-degree half angle. Delta represents an offset from the secondary, focus measured in feed apertures, for the three axes of the antenna optics. (RD16)

For an antenna tailored to single-pixel feeds, the primary impact at this level is clear: shaping improves illumination efficiency and overall system sensitivity. Holding sensitivity constant, the antenna aperture size can be decreased through shaping, increasing the field of view and survey speed. However, there are three secondary effects to consider:

- The optical shaping can raise the level of the first and second side lobes.
- The rate of change in gain and primary beam response due to a feed offset (e.g., due to gravitational deformation over elevation) is larger.
- The optical shaping effectively precludes the use of multi-pixel feeds in the future.

The importance of the sidelobe level varies by frequency. At low frequencies (<1 GHz) the beam is large and the sky density of sources is also high. A bright source in the 1st or 2nd sidelobe can dominate the flux in the field. High dynamic range imaging then requires imaging out to the 1st or 2nd null, which can drive the post-processing system sizing and complexity. At high frequencies, the beam is smaller and the sky

Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

density of bright sources drops, minimizing the concern. Given that ngVLA is optimized for operation above 10 GHz, this concern is not a design driver.

The rate of change in gain due to a feed offset is a concern for calibration. The feed indexer will need to be more precise to restrict the gain error after switching feeds, and the gain fluctuations with elevation angle will need to be corrected in the calibration pipeline. Finite element analysis of the antenna design suggests that the gravitational effect dominates and that thermal deformations have a minimal impact on feed displacement (RD40). Fortunately, the gravitation effect is expected to be deterministic and therefore suitable for calibration automation.

Considerations for future multi-pixel feeds are tied to the science use case analysis. As summarized in ngVLA Memo #18 (RD05), most use cases are sensitivity limited. There are use cases that would benefit from a larger field of view and associated improvements in mapping speed that could be achieved with FPAs or PAFs, but 50% of use cases are single antenna pointings, and 70% of fields can be imaged with 16 antenna pointings or less.

Fundamentally, ngVLA is a high-sensitivity follow-up instrument, not a mapping instrument, and the system design choices reflects this scientific emphasis. The contrast to SKA science is illustrated in Figure 4 which shows the field sizes that must be mapped for the key science use cases of the SKA and ngVLA. The SKA will map over 2000 times more area to achieve the key science, while ngVLA use cases are constrained to small patches of sky.

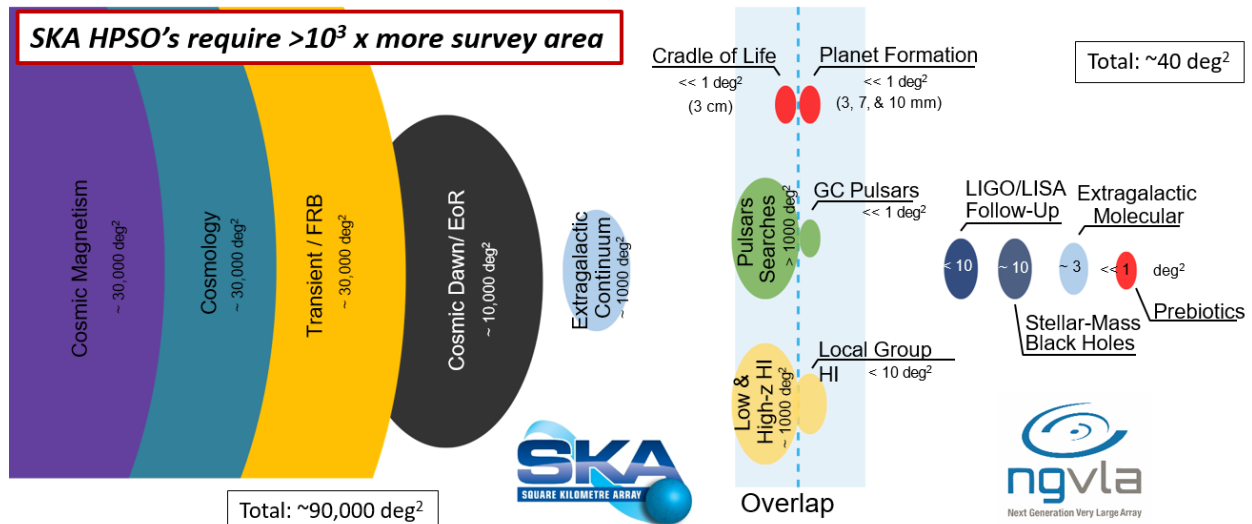


Figure 4 - Field of view requirements for SKA key science vs ngVLA key science goals.

Given these considerations, ngVLA has adopted a shaped optical system tailored to single pixel feeds. The system is optimized to maximize forward gain while reducing spillover temperature ($A_{\text{eff}}/T_{\text{sys}}$ being the key performance parameter), when averaged over the full elevation range at 30 GHz. The antenna optics selected for the preliminary design of the antenna, along with a tailored feed design, achieve an illumination efficiency of 96% while 1st and 2nd sidelobe levels are an acceptable -19dB and -25dB respectively. [RD15, RD16]

Please consult the following materials for further information on this trade:

ngVLA Reference Design Development & Performance Estimates
 Summary of the Science Use Case Analysis
 Antenna Optical Design Report

ngVLA Memo #17
 ngVLA Memo #18
 020.25.01.00.00-0001-REP



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

EMSS Optical Design Report
 An Optimal 18 m Shaped Offset Gregorian Reflector for the
 ngVLA Radio Telescope

EA-NGV-DR-05
 IEEE Trans. A&P, Dec. 2021

4.3.3 Symmetric Cassegrain vs Offset Gregorian Optics

A system-level comparison of offset and symmetric optics can be found in ngVLA Antenna Memo #1 (RD17). The analysis considered lifecycle cost and sensitivity as the key performance parameters. An aperture size of 18m is assumed for both designs, consistent with the analysis in Antenna Memo #2 (RD07) and the discussion in Section 4.3.1. Holding the aperture size constant significantly simplifies the analysis, and ensures comparable survey speed performance if the sensitivity of the two arrays can be matched.

The key performance differences between the offset and symmetric designs are captured in aperture blockage, scattered RF power, and the associated increase in spillover temperature for the symmetric design. The expected differences averaged over ngVLA operating frequencies are summarized in Table 1. An 18m symmetric Cassegrain antenna is expected to be 81% as efficient as an unblocked offset Gregorian equivalent. An additional 23% (1/0.81) more antennas are required with a symmetric Cassegrain array to achieve the same system-level point source sensitivity and survey speed.

Table 1 - Normalized Efficiency Factors for ngVLA Offset and Symmetric antenna implementations. Assumes an 18m aperture with a 3.5m subreflector. Spillover and system temperature point values given for 30 GHz operation (RD17).

Geometry	η_{BLOCK}	η_{SPILL}	T_{SYS}	η_{TOTAL}
Offset Gregorian	1.0	1.0	1.0	1.0
Symm. Cassegrain	0.87	0.99	0.94	0.81

Such an array would appear to perform well on imaging fidelity metrics, but the additional antennas will drive the total operations cost of the facility. An array design optimized for Cassegrain optics would likely need larger apertures to overcome this deficit, and would result in a comparable number of apertures to an optimized array with Gregorian optics. The net result is matching sensitivity and imaging fidelity, but a smaller field of view and slower survey speed.

Looking at the total lifecycle cost, the memo assesses the cost differences expected for offset Gregorian geometries, holding array sensitivity and aperture size constant, and concludes that:

- For the Cassegrain option to match the sensitivity of the Gregorian for the same nominal system construction cost, the antenna would have to cost 38% less.
- For the Cassegrain option to match the sensitivity of the Gregorian for the same nominal system lifecycle cost (construction and operations), the antenna would have to cost 76% less!
- The Cassegrain option would likely cost 28% less than the Gregorian equivalent, which will not offset either the construction cost total or the lifecycle cost difference.

The supporting cost analysis was performed using the quantitative exchange model (RD01). The reason for the large cost differences required is largely driven by the fact that the antenna mount represents only half the cost of outfitting an array element (i.e., if the antenna costs C, an outfitted array element inclusive of the antenna electronics and processing system costs 2C.). Operations costs scale roughly by the number of apertures in the array, and the Cassegrain option requires 23% more apertures which drives up the relative lifecycle cost for a symmetric Cassegrain array design.

While the project did not pursue a design and costing of an 18m symmetric Cassegrain antenna to project requirements, a comparison of the cost estimate for the selected 18m offset Gregorian concept to



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

parametric estimates for a Cassegrain equivalent suggests the real cost difference is likely less than projected in the memo, and closer to 15% (vs 28%) at ngVLA operating frequencies, further substantiating the decision to implement offset optics.

A broader study into the symmetric Cassegrain and offset Gregorian designs was performed and summarized in ngVLA Antenna Memo #3 (RD18). The analysis was multi-parameter, with four KPPs (life-cycle cost, sensitivity, survey speed, imaging performance), a number of ancillary criteria, and multiple feed and optical designs considered in an attempt to adequately explore the parameter space. Optical geometries, feed illumination angles, and focus and band selection mechanisms were considered as sets. The results of this study are summarized in Table 2.

Table 2 - Performance Summary of Offset Gregorian and Symmetric Cassegrain optics on selected metrics. (RD18)

Concept	Offset Gregorian Wide-angle feeds, fixed focus, linear translation mechanism	Symmetric Cassegrain Narrow-angle feeds, feed ring, focus-rotation mechanism
Life-Cycle Cost	Lower projected lifecycle costs, per analysis in RD01 and RD17.	Lower construction cost for the mount, but savings must be large (38% to 76%) to offset the lifecycle advantage of an offset Gregorian.
Sensitivity	Higher A_{EFF}/T_{SYS} of selected options.	Blockage and scattering reduce A_{EFF}/T_{SYS} . Requires 20% additional apertures at 18m, or 20m apertures which would reduce survey speed and require tighter pointing performance to support the imaging dynamic range requirements.
Survey Speed	Higher sensitivity at a given aperture size means fastest possible survey speed (since choice maximizes FoV and sensitivity)	
Imaging Performance	Excellent sidelobe, spillover, and cross-polarization performance.	Sidelobe level may be 3-6 dB higher, spillover also higher by a few K, but some choices could increase the number of apertures in the array, which has a positive influence on imaging performance.
Frequency Coverage	Excellent performance over full 1.2 GHz to 116 GHz range, due to an unblocked aperture.	The low frequency limit mandates a subreflector size (~3m) that is a significant fraction of the 18m primary aperture. The resulting blockage loss is applied to all frequencies.
Other Pros	Design continuity from the reference design, including feed, front-end and cryogenic system concepts.	More mechanically available concept, more designers and manufacturers, lower tooling costs for prototype.
Other Cons	Prototype and development costs are expected to be higher. Also, uncertainty in final production cost: no industry experience building comparable antennas on the scale envisioned.	Ring focus geometry adds inherent cross polarization. Required focus rotation mechanism (FRM) on apex is a high-maintenance item. Major redesign on front end and cryogenic system concepts.



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

Concept	Offset Gregorian Wide-angle feeds, fixed focus, linear translation mechanism	Symmetric Cassegrain Narrow-angle feeds, feed ring, focus-rotation mechanism
Final Ranking	Preferred option.	Should only be considered if the construction cost deltas exceed the projections in RD17, or the frequency coverage requirements change at the system level.

Please consult the following materials for further information on this trade:

The ngVLA Quantitative Exchange Model	nqxm Ver3.0
System-level Cost Comparison of Offset and Symmetric Optics	ngVLA Antenna Memo #1
Antenna Optical Design Alternatives	ngVLA Antenna Memo #3

4.3.4 Feed Indexer Geometry

As part of a broader study into the symmetric Cassegrain and offset Gregorian optics discussed in Section 4.3.3, a number of feed indexer geometries were considered. These trades are summarized in ngVLA Antenna Memo #3 (RD18).

The analysis was multi-parameter, with four KPPs (life-cycle cost, sensitivity, survey speed, imaging performance), a number of ancillary criteria, and multiple feed and optical designs considered in an attempt to adequately explore the parameter space. Optical geometries, feed illumination angles, and focus and band selection mechanisms were considered as sets. The following options were identified as the most viable:

- Symmetric Cassegrain optics with narrow angle feeds, a feed ring, and a focus/rotation mechanism.
- Offset Gregorian optics with wide angle feeds, fixed focus, and a linear translation mechanism.

As part of an integrated down select, the offset Gregorian concept was selected as described in Section 4.3.3. Building upon this conclusion, we will summarize the selection of the feed indexer geometry here.

The selected concept uses wide-angle low-gain feeds which are very compact. The combination of narrow-angle feeds and offset Gregorian optics was over constrained in an 18m aperture size, with no practical feed arrangement and feed selection mechanism that could work over the specified 1.2 to 116 GHz frequency range.

The compactness of the wide-angle feeds does offer performance benefits, since the high frequency feeds are sufficiently small to be cryogenically cooled, reducing receiver noise temperature. The compactness also offers lifecycle cost reductions, with the integration of multiple feeds into a limited number of cryostats (presently two) to improve overall efficiency and reduce the construction and electrical operation costs. Maintaining this compact feed arrangement was a priority in the selection of the feed indexer geometry.

The most common feed indexer arrangement selected by other projects for similar optics is a rotation stage for band selection. MeerKAT has such an arrangement with focus adjustment for the X-band receiver only. The SKA MID antenna also uses a rotation stage, and has no focus adjustment (at all frequencies). The turntable is mechanically simple and can be rotationally balanced, ensuring that the center of gravity does not move due to a band change.

A key difference between ngVLA and MeerKAT or SKA MID is the range of operating frequencies. ngVLA Bands 2 through 6 are expected to require focus adjustment to retain their sensitivity over the full range

Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

of environmental conditions. Focus adjustments are also likely to be made during an observation to compensate for gravitational sag for Bands 4 through 6. (RD40)

The need for focus adjustment on most bands complicates the down select. A turntable is only viable if either:

- The receivers are split into discrete cryostats with focus translation stages below each cryostat, or
- The focus translation stage is *below* the turntable, and moves the entire turntable.

The first option is unattractive since it requires multiple focus mechanisms and it increases the construction and operations cost for the multiple cryogenic cryostats. Mounting the turntable on a focus translation stage introduces a larger change in the center of gravity when the focus is adjusted and was deemed mechanically complex by antenna vendors who were consulted.

A two-axis linear translation stage was seen as appreciably mechanically simpler by the consulted antenna vendors, and could leverage a number of heavy duty precision actuators developed for two and three axis machinery. A downside of this design is that band changes do move the center of mass, but the overall cryostat package can be optimized to minimize the mass and the associated impact on antenna pointing and optical alignment. The compact packaging of the linear arrangement, and the elimination of counterbalance for rotational symmetry results in an appreciable reduction in mass compared to a turntable concept. Ancillary equipment (such as vacuum pumps) can also be moved off the moving platform, further reducing moving mass. Based on vendor input between this option and a focus stage below a turntable, it was ultimately selected as the preferred concept for the reference design.




Figure 5 – Left: Early concept sketch of the MeerKAT receiver turntable. Courtesy of General Dynamics and MeerKAT. Right: Offset Gregorian optical geometry with linear translation of the receivers for band selection and focus (Y and Z axes only). System presently translates +/- 650mm in Y and +/- 100mm in Z.

A third axis was considered to compensate for the movement of the secondary focus due to gravitational deformation. However, this additional axis would increase the overall cost, complexity and mass of the system. It would also add an additional position error that would need to be accounted for in the referenced pointing error budget. Elimination of moving parts is a guiding principle in the antenna design for lifecycle cost reasons, so a third axis would only be introduced if proven necessary to maintain overall aperture efficiency within specifications at high frequency, or as a way to reduce the lateral offset in the front end enclosure mass (in conjunction with a redesigned cryostat) in support of the pointing error specification.

At the conclusion of the antenna conceptual design studies the issue was revisited. Importantly, finite element analysis of the antennas demonstrated that:



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

- The two-axis positioner provided enough degrees of freedom to maintain overall efficiency within specification at all frequencies and elevation angles. Not only was a third-axis not required, it did not recover significant gain (<1%) when compensating for gravitational deformation.
- The pointing error budget allocation for the moving mass of the translation stage and a residual feed position error in the translation axis could be accommodated in the error budget.

The antenna conceptual design studies proved that the selected concept did not introduce unexpected risk or complexity into the design, and a two-axis linear positioner could be retained as the preferred concept for the preliminary design phase.

However, the project did revisit the radial arrangement before executing the antenna preliminary design contract, and noted the following observations should a radial turntable arrangement be adopted. Please refer to Figure 5 for definitions of the relevant coordinate axes.

- The feed boresight lines on all six bands must intersect at a single point, which we will call the feed boresight center. The feed boresight lines also have to lie in the same plane as the optical boresight from the secondary mirror, though this is a requirement for the linear translation scheme as well.
- The feed boresight center has to be precisely aligned with the axis of the rotation stage, to minimize any offset from the optical boresight that can degrade efficiency, pointing, and cross-polarization. This isn't an issue with a linear translator, because a y-axis offset can be easily taken out with a mechanical or sky calibration after installation of the front end. A radial band arrangement would require the entire cryostat/enclosure to be physically moved to take out any offset, unless proper alignment could be guaranteed by dead reckoning (e.g., with precision dowel pins).
- The axis of the rotation stage must lie within the optical plane of symmetry; i.e., to intersect the optical boresight from the secondary and be aligned along the local x-axis.
- The rotation center must be located such that the alignment of any feed phase center to the secondary focal point is possible within the z-axis translation limits. The feed phase centers don't necessarily have to be the same distance from the feed boresight center, but this reduces the change in center of gravity which is the most attractive feature of the radial concept.
- As previously noted, the rotation stage must be located between the front end and the z-axis stage, in order for focus adjustment at all rotation (band) settings. For a y-z translation stage, the installation order does not matter.

Given these observations, an updated list of pros (+) and cons (-) for the radial concept is as follows:

- A radial arrangement avoids the need for long throws in either y-direction for band selection, and reduces the shift in the load. One could also counterweight the rotating load to reduce or eliminate this center of gravity shift. (+) This could provide some relief in the antenna feed arm lateral (y-axis) stiffness (+), but the added mass would require improved stiffness in the x-axis to limit the losses due to gravitational sag (-).
- It is easier to avoid interference between adjacent beams on the lowest frequency receivers compared to an inline arrangement. (+)
- Because the lower frequency receivers tend to have greater lateral extent at the feed/thermal gap end, a radial arrangement could accommodate modularity while possibly reducing the overall thermal loading on the cryocooler. (+)
- Depending on the radius used and placement of individual receivers, there is the potential for mechanical interference between bands and/or cryostats. Given the feed volume constraints, a minimum radius may be 610mm, not including a space allocation for the downconverter-digitizer assembly. The overall package is larger than the linear translation option. (-)



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

- Phase center positioning error relative to the focal point is the product of the angular positioning accuracy (radians) and the phase center radius. So for a radius of 610 mm and a positioning error of +/- 0.25 mm, the required angular resolution is ~0.82 mrad, or 0.047 degrees. If a larger radius is needed to avoid mechanical interference, the requirements on positioning accuracy will go up accordingly. (-)
- For the IRD modules, available volume may be tighter and the mechanical/RF interface more complicated compared to present linear scheme. (-) The spacing between bands on the side opposite the feeds will be markedly closer, particularly if the phase center radius is small, making the available volume much less. A practical solution may require a larger radius for the rotation stage, and the overall mechanical complexity of IRD and LO system packaging and service will increase. (-)
- The total volume on the feed arm is likely higher, depending on the maximum radial extent of the cryostats/feeds, particularly in the z-direction (-). With a larger radius, the receiver platform extent is larger, with greater unbalanced loads on the positioner bearing, and greater rotational torque required. (-)
- The cable wrap design may be more complex than with linear translation. (-)

These findings were subsequently captured in RD57. In aggregate, this trade has a roughly balanced set of positive and negative features, with no clear preference for one design over the other. However, this does not account for the significant difference in the technical readiness and development level of the two approaches. As such, the existing linear translation approach, which has been developed throughout the reference design phase and validated by the antenna vendor, is retained as the preferred concept.

Please consult the following materials for further information on this trade:

Antenna Optical Design Alternatives
Antenna Conceptual Design Report
Rotating FE Design: Benefits and Constraints

ngVLA Antenna Memo #3
1021006-REP-21-000000-0001
020.30.05.00.00-0007-REP

4.3.5 Multi-band vs Single-band Operation

Some science use cases favor multi-band operation, and key VLBI use cases rely on multi-band approaches for calibration, such as S-X or X-Ka observations to improve atmospheric calibration for astrometry. A decision had to be made on multi-band or single-band operation as part of the broader down select of the antenna optics discussed in sections 4.3.3 and 4.3.4. Importantly, the preferred concept of wide-angle feeds with the offset Gregorian optics *precludes* multi-band observations without significant additional optical losses. The wide angle of the feeds cannot be intercepted with a dichroic mirror. If multi-band observations were required, this would likely lead to the selection of symmetric Cassegrain optics, with tradeoffs as described in Section 4.3.3.

Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B




Figure 6 - The ngVLA rev 7 Optics. Main reflector and subreflector in black, with subreflector extension shown in gray. Ray trace is at equal angular spacing propagating from the secondary focus to the subreflector, and is indicative of the aperture illumination pattern achieved through shaping. (RD16)

Figure 6 demonstrates the incompatibility of implementing a multi-frequency receiver system with the preferred optics. The subtended angle of the feeds is 110 degrees, with the feed positioned close to the subreflector. The beam expands rapidly and cannot be intercepted and redirected by a dichroic mirror.

Multi-band observations are especially common in VLBI as a strategy to calibrate instrumental and atmospheric effects for precision astrometry (RD22, 23). Importantly, multi-band observations are solutions to atmospheric calibration challenges, and are not directly required to support any of the identified ngVLA science cases. Some solar use cases and other transient science would be enhanced by multi-frequency observations, but a majority of use cases assume that the target is unchanging on the timescale of the observation, so frequency multiplexed observations (time shared in different bands) are an acceptable solution if the atmospheric calibration challenges can be overcome. Sub-arrays observing in different bands concurrently are also an acceptable scientific solution when the subarrays sample equivalent spatial scales.

Dual-band modes incur a sensitivity penalty by introducing additional system noise: The VLBA X-band SEFD rises 34% when the dichroic is in place (RD41). The relative difference in sensitivity between dual-band and single-band implementations significantly reduces the effective observing time difference between dual-band and time multiplexed modes, and the total observing time required across a diverse set of use cases is expected to be optimized with a single band solution. In the context of long baseline interferometry, this SEFD loss from a dichroic is also comparable to dedicating one of the three antennas in each LBA site to observe in a second receiver band, so a subarrayed solution should have comparable



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

SEFD, though calibration residuals would be expected to be larger given the separation between the LBA station antennas and the atmospheric phase structure function over the site.

Given the high emphasis on sensitivity in the science requirements, we have adopted the high efficiency single-band concept and offset Gregorian optics as the default solution, and will aim to replace the common dual-band modes as part of our observing mode and calibration strategy development. Some of the rationale for multi-band observations is explained below as well as the alternative calibration strategies selected for development for the ngVLA. The calibration strategies in support of each observing mode will be developed as part of the system preliminary design.

Multi-frequency observations are used in VLBI to overcome a number of propagation effects in the atmosphere due to the ionosphere at low frequency and the troposphere at high frequency, with the goal of increasing the coherence time in the observation. The techniques can also aim to mitigate the scarcity of calibrators at high frequency, and to tie reference frames together through relative position measurements at two frequencies (RD20).

Frequency phase transfer, or band-to-band gain transfer, uses observations at low frequency to calibrate for propagation effects (primarily due to the troposphere) at high frequency. This technique relies on the greater availability of calibrators at low frequency, since calibrator irradiance generally drops as a function of frequency in ngVLA operating frequencies. A secondary effect is the larger beam size, which increases the likelihood of an in-beam calibrator in some cases. The use of the additional source irradiance at low frequency can improve the SNR of the calibration observation, increase the coherence time in the observation if the calibrator is in-beam, and generally improve the achieved sensitivity as a function of time. However, a downside of this technique is that it can place additional amplitude and phase stability requirements on the signal chain. The method requires measuring the frequency dependence of gain in each receiver band with a strong source, then relies on this internal band gain being time-stable during the rest of the observation. Relative differences in amplitude and phase between the observed bands must also remain stable.

ngVLA will improve the availability of calibrators and lengthen the coherence time through other means. The limited availability of calibrators for mm-VLBI is in part driven by the limited sensitivity of existing arrays. ngVLA overcomes this issue through the number of antennas in the array, lower system temperatures, and significantly increased transmitted and processed bandwidth. The lower system equivalent flux density of the ngVLA will ensure that a suitable mm-wavelength calibrator is typically available within 2 degrees on sky and almost always within 4 degrees (RD09, RD12) at all frequencies, so fast switching phase calibration will be appreciably more effective on ngVLA than existing mm-wave long baseline arrays¹.

Coherence time between calibrator observations may also be increased through the use of water vapor radiometers (WVR). The WVR operates continuously in all bands and estimates the propagation delay from the troposphere on a 1 second cadence, with the intent of improving signal coherence on short timescales and extending coherence time from 10s of seconds to minutes or 10s of minutes depending on the observing band and the atmospheric conditions.

At low frequencies (<10 GHz) the ionosphere dominates propagation delays. Simultaneous multi-frequency observations (such as S-X observations) enable calibration of these ionospheric effects. Fortunately, ionospheric delays are relatively slow changing (compared to the troposphere) so frequency switching is a suitable solution to mitigate these effects for ngVLA (RD12). In addition, the ngVLA receiver band definition (discussed in Section 4.4.1) favors wideband feeds below 12 GHz. The Band 2 feed extends

¹ This statement is true when ngVLA is operated as a full array or subarray, but may not be true when operating in a VLBI-array with other facilities where the slowest antenna sets the switching cycle time.



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

from 3.4 GHz to 12.3 GHz, effectively supporting the calibration strategies that employ S-X modes today with a single receiver. The ngVLA solution will likely be higher performance in practice since changes in phase, delay or gain between the two receiver paths in a dual-band solution are eliminated, and the full bandwidth of the receiver is available, which improves delay solutions. The ngVLA Band 2 edges are also rather close to the VGOS array wideband design, so calibration strategies adopted for astrometry by others could be adapted to the ngVLA long baselines.

Returning to high frequencies (>10 GHz), the ngVLA system bandwidth can improve calibration solutions within a single receiver band. Up to 20 GHz of processed bandwidth is selectable (split into non-contiguous sub-bands) from across the full receiver band. This ensures that upper and lower band edges can be transmitted when solutions are improved by sampling widely spaced frequencies.

Finally, sub-arrayed calibration and observation strategies can replicate dual or tri-band observing modes. The inclusion of three antennas at each LBA site in the ngVLA configuration, with the antennas in a compact short-baseline arrangement within the site, lends itself to novel observation strategies not presently available with single antenna VLBA stations. The three antenna site can complete an all-sky survey faster in a geodetic use case, be used for paired-element calibration with one antenna dedicated to observing an astronomical calibrator, or be subarrayed for multi-frequency observing strategies depending on upon the scientific requirements for the observation.

In aggregate, the ngVLA calibration requirements (RD12) and observing mode definition (RD24) demonstrate that ngVLA can support the astrometric precision requirements with alternative techniques, while providing greater sensitivity to the majority of use cases.

4.3.6 Altitude-Azimuth vs Equatorial Mounts & Rotation Axes

Most modern arrays and single-dish telescopes are constructed with Altitude-Azimuth mounts. The performance of two-axis servo systems, when combined with a multi-term pointing model, improve upon the single-axis tracking performance achievable with an equatorial mount while being structurally simpler and more affordable to construct. However, the field rotation inherent in an Altitude-Azimuth mount can introduce challenges and does warrant brief consideration.

The preference for linear polarization feeds (as discussed in Section 4.4.2) when combined with Altitude-Azimuth mounts that are widely separated introduces a parallactic angle difference on sky. Effective calibration of this effect requires ‘full stokes’ imaging, where both parallel-hand and cross-hand correlation terms (i.e., XY and YX, in addition to XX and YY), which doubles the data rates from the correlator for most use cases. Given the significant cost difference for equatorial mounts, this trade is considered more closely associated with the feed polarization basis.

Multi-pixel feeds are also computationally complex with a rotating field introduced through earth rotation. The ASKAP telescope is informative, solving this issue with a rotation stage to maintain a constant field rotation across the array and across an observation. A rotation stage can also mitigate subtle polarization and imaging effects, such as measuring the intrinsic cross-polarization inherent in a feed ring geometry, or correcting for the effect of subreflector support structures on symmetric antennas.

The ngVLA concept aims to optimize the performance of single pixel feeds (Section 4.3.2) so the benefit of multi-pixel performance does not rank highly in any trade off. A rotation stage for calibration is only applicable to a symmetric antenna geometry, and independent trades favor the offset geometry (Section 4.3.3), so a rotation stage is not further considered in the system design.

Accepting the impact on the correlator data rates, data archive rates and compute capacity inherent in the combination of linear polarization with Alt-Azimuth mounts spread over continental-scale baselines is consistent with the overall ngVLA concept. These additional data products do have incremental value (e.g.



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

producing polarization products for science projects that only request Stokes I) that increases data reuse, and more generally, progresses the design towards the digital telescope ideal where digital signal processing replaces complexity previously mitigated through mechanical or RF means. Finally, ALMA also uses the combination of linear feeds with Alt-Azimuth mounts, so ngVLA may be able to leverage ongoing developments in polarization calibration techniques and their associated CASA implementation, increasing software reuse.

4.4 Antenna Electronics

4.4.1 Receiver Band Definition

The receiver band definition is a key conceptual choice for any antenna or array design. The evolution of the receiver band definition, supporting concept, and relevant metrics is discussed in detail in the Front End Design Description (RD25) and Front End Trade Study (RD53). These results are summarized here along with some ancillary information that informed on the wide-band feed designs for Bands 1 and 2.

Table 3 - ngVLA Receiver band definition. (RD25)

Band #	Cryostat	f_L (GHz)	f_M (GHz)	f_H (GHz)	$f_H: f_L$	BW (GHz)
1	A	1.2	2.01	3.5	2.91	2.3
2	B	3.4	6.56	12.3	3.62	8.9
3	B	12.3	15.9	20.5	1.67	8.2
4	B	20.5	26.4	34.0	1.66	13.5
5	B	30.5	39.2	50.5	1.66	20.0
6	B	70.0	90.1	116	1.66	46.0

The receiver band definition is summarized in Table 3. Bands 3 through 6 are waveguide bandwidth, while bands 1 and 2 are wideband designs. The receiver RF bandwidth is unique and largely non-overlapping except for Band 4 which overlaps with Band 5.

As mentioned elsewhere in this report, sensitivity is a driving requirement for the ngVLA. The distribution of use cases is roughly evenly split between continuum and spectral line use cases, so these are given roughly equal weight in any trade study. This is an important consideration for feed bandwidth selection, since at X-band and higher ($\sim > 10\text{GHz}$) waveguide orthomode transducers (OMTs) and low noise amplifiers (LNAs) lead to optimal noise performance, but limit the receiver band edge ratio to approximately 1.5:1 to 1.7:1 while constraining the degradation in performance at the band edges and avoiding the excitation of higher-order waveguide modes.

The sensitivity goals alone lead to a relatively traditional band definition above 12 GHz, with waveguide bandwidth receivers for bands 3 through 6. The overlap in Band 4 and 5 bandwidth is driven by continuum sensitivity requirements. Enabling thermal imaging at milli-arcsecond resolutions is a key science theme of the ngVLA. These thermal imaging use cases benefit from high continuum sensitivity around 30 GHz, where bandwidth can be maximized while avoiding atmospheric emission. As can be seen in Figure 7, atmospheric emission rises around the 22.235 GHz water line, drops for a span around 30 GHz, and then climbs rapidly as we approach the Oxygen feature around 60 GHz. Band 5's edges are constrained by the oxygen line emission at the upper end, and the waveguide bandwidth feasibility constraints on the low



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

end. Total system sensitivity in Band 5 continuum modes is then limited by the emission from the broad wings of the oxygen line.

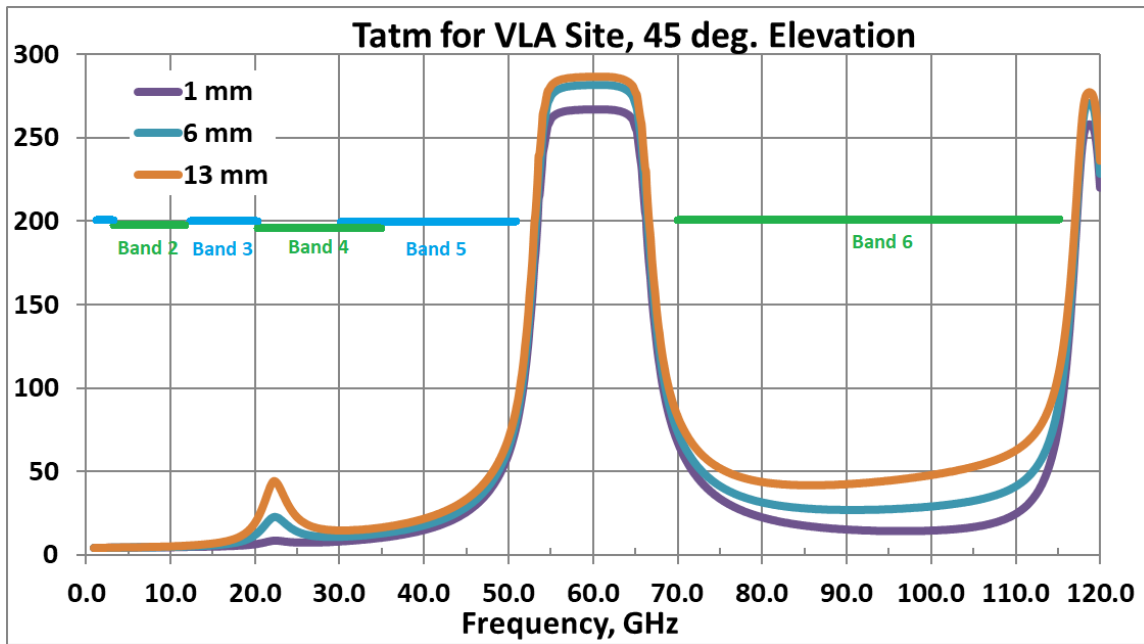


Figure 7 - Atmospheric contributions to system temperature (K) at the VLA site at 45-degree elevation for PWV column heights of 1mm, 6mm and 13mm. Band RF span (Table 3) shown for context. (RD26)

A Band 4 definition tailored to continuum observations would sample the widest bandwidth possible (while still achieving optimal receiver temperature) between the water line and the oxygen line, perhaps spanning from 24 GHz to 40 GHz. Such a receiver is difficult to accommodate while minimizing the total number of receivers in the design, but the ngVLA Band 4 definition overlaps with Band 5 to approximate this solution, especially in precision operating conditions with low precipitable water vapor.

The ultimate location of the band edges for Band 4 is driven by the practical bandwidth achievable in Band 2 as part of a total band definition, accepting that the band edge ratio of Band 3 will also be 1.67:1 for optimal noise performance. Bands 1 and 2 are driven by a combination of practical constraints and evolving science priorities at low frequency. We tackle each in turn.

A key constraint for ngVLA is the operations cost cap established in the stakeholder requirements (AD02). Operation and maintenance costs for the front end system roughly scale with the number of receivers in the antenna, so minimizing the total number of receivers is a supporting design goal. The VLA band definition, designed for optimized sensitivity without band gaps from 1-50 GHz has a total of 8 receiver bands. An extension of this definition to ngVLA operating frequencies system would result in a 9-band system which would drive both the construction and operations costs of the system.

Table 4 - Fraction of time used in each receiver band in the notional Envelope Observing Program. (AD10)

Receiver Band	1	2	3	4	5	6
EOP Time Fraction	7%	6%	14%	27%	18%	27%

The distribution of science use cases spans the full 1.2 to 116 GHz range, but is skewed to higher frequencies. The distribution of time by band in the envelope observing program (EOP, AD10) is given in Table 4. Bands 1 and 2 account for less time combined than any other receiver band. Accepting that a

Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

compromise in performance may be required to respect the operations cost cap, it is clear that this compromise should be made at lower frequencies by spanning more frequencies with a limited number of receivers.

The band edge between Band 1 and 2 is constrained by the downconverter-digitizer topology discussed further in 4.4.4. Direct sampling Band 1 with a common digitizer chip sets an upper limit to the first Nyquist zone at approximately 3.4 GHz, accounting for the antialiasing filter response. Band 2 then spans the gap between Band 1 and the lower edge of Band 3. A band edge ratio of less than 3:1 for Band 1 lends itself to an all metal feed design, while the 3.5:1 ratio of Band 2 favors dielectric loading to maintain illumination efficiency over the full frequency range of the receiver.

These practical constraints, combined with a science requirement to avoid key spectral lines at band edges, largely led to the current band definition. However, there are secondary science considerations that suggest wideband designs are attractive at these frequencies:

- A number of transient science use cases benefit from bandwidth, even at the expense of continuum sensitivity, and these transient science use cases are predominantly at low frequency (<10 GHz).
- The SKA MID array is expected to be a fantastic leading-edge instrument at frequencies below 10 GHz. Providing a different but complementary set of capabilities at low frequencies with the ngVLA improves the total suite of instruments available to the scientific community.

The importance of bandwidth in transient science is only recently becoming apparent. As an illustration, Figure 8 shows the uncertainty in time of arrival (TOA) for two selected pulsars as a function of bandwidth and receiver center frequency (RD27). Since the spectral index of pulsars vary, there is no ideal center frequency for a survey of many sources. What is common is that wider sampled bandwidth can improve the TOA uncertainty by providing a better constraint of the dispersion measure than would be feasible with a narrow bandwidth receiver, assuming comparable receiver noise and illumination efficiency.




Figure 8 – Left: TOA uncertainty as a function of observing center frequency and bandwidth for PSR J1909-3744 observed with the GBT. Solid contours here indicate TOA uncertainties of 2, 1, 0.5, 0.2, and 0.1 μ s, in order of increasing darkness and thickness. The minimum TOA uncertainty (black star) is $\sigma_{\text{TOA}}(\nu_0 = 8.1 \text{ GHz}, B = 14.8 \text{ GHz}) = 50 \text{ ns}$. The estimate given the GBT frequency coverage (black circle) is $\sigma_{\text{TOA}}(\nu_0 = 1.3 \text{ GHz}, B = 1.2 \text{ GHz}) = 59 \text{ ns}$. Assumes $A_{\text{eff}}/T_{\text{sys}}$ varies with center frequency, but not bandwidth. Right: TOA uncertainty as a function of observing center frequency and bandwidth for PSR J1903+0327 observed with AO. Solid contours here indicate TOA uncertainties of 200, 100, 50, 20, and 10 μ s, in order of increasing darkness and thickness. The minimum TOA uncertainty (black star) is $\sigma_{\text{TOA}}(\nu_0 = 9.8 \text{ GHz}, B = 13.2 \text{ GHz}) = 1 \mu$ s. The estimate given the AO frequency coverage (black circle) is $\sigma_{\text{TOA}}(\nu_0 = 1.8 \text{ GHz}, B = 1.2 \text{ GHz}) = 44 \mu$ s. Assumes $A_{\text{eff}}/T_{\text{sys}}$ varies with center frequency, but not bandwidth (RD27).



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

ngVLA is uniquely suited to incorporating wide bandwidth low frequency receivers since the system supports up to 20 GHz of processed bandwidth per polarization due to the broader system optimization for higher frequency observations. The frequency span of Band 2, processing the full 8.9 GHz of received bandwidth, is well suited to the key transient use cases (e.g., identifying pulsars at the Galactic center) and complementary to other arrays that will operate contemporaneously with the ngVLA.

As the antenna design has progressed, the mass of the front end enclosure has proven to have a significant influence on antenna pointing and surface performance with the preferred offset optical geometry. The reduced number of receiver bands, especially larger low frequency bands, helps minimize the front end mass and maintain the antenna pointing and surface accuracy required for high frequency observations. The chosen feed indexer geometry (Section 4.3.4) is also quite sensitive to total mass and receiver count. An optimized integrated offset Gregorian antenna design strongly favors a compact and low mass receiver configuration and supports the selected concept.

A final observation is that in the preferred system concept the Band 1 receiver has a dedicated cryostat and unique digitizer module (RD25, RD32). As the lowest frequency receiver, it is also inherently the largest, with a populated cryostat mass equivalent to about half of the combined Band 2-6 cryostat (RD25). The overall volume, mass and separate cryostat lead to a relatively higher construction cost for this receiver band, and incremental cost in the feed arm structure to accommodate the payload. The cryostat helium load also roughly doubles the cryogenic system electrical load within an antenna. However, this receiver is clearly required to support the full science requirements and is therefore retained in the baseline design.

Please consult the following materials for further information on this trade:

Front End: Design Description	020.30.05.00.00-0006-DSN
Antenna Optical Design Alternatives	ngVLA Antenna Memo #3
Front End Trade Study	ngVLA Electronics Memo #13

4.4.2 Linear vs Circular Polarization Basis

The decision to adopt a linear polarization basis was determined while settling on the band definition described in Section 4.4.1. The factors that led to this decision are summarized here.

Most astronomical sources are weakly linearly polarized, and these polarization attributes have traditionally been sampled with circularly polarized feeds. Circular polarization can simplify both the calibration process for a single observation and in some cases the data reduction process, since determining source intensity does not require generating cross-polarization correlations (LR/RL).

Due to these factors and institutional knowledge developed as part of VLA design and operation, circular polarization was assumed to be the preferred solution in the early design phase.

The first complexity is related to the wide bandwidth of Bands 1 and 2, with a band edge ratio of up to 3.5:1. Feed designs that operate over this wide bandwidth are intrinsically linearly polarized. Adopting a circular polarization basis would require conversion from linear to circular. This conversion is not desirable, since introducing the conversion prior to the LNA adds significant noise to the receiver noise figure, while conversion after the LNA (but still in the analog domain) limits the polarization dynamic range due to subtle complex gain (amplitude or phase) variations between the two LNAs over frequency and time.

For these reasons, the preference for Bands 1 and 2 is to retain the native linear polarization basis. Bands 3 through 6, with their waveguide RF trees, could be converted to circular polarization with a square



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

waveguide corrugated phase shifter and a square waveguide 45-degree transition before the orthomode transducer. This technique is employed today on the VLA very effectively.

However, phase shifters have only been demonstrated to work over ~1.5:1 bandwidth ratios, while the present design uses a 1.67:1 bandwidth ratio for the waveguide receivers. A waveguide conversion to circular polarization could require a redefinition of the band edges, the addition of a receiver band below 50 GHz, and a reduced frequency range in Band 6. A solution may exist using linear RF trees and hybrid couplers after the LNAs for Bands 3 and 4, but this implementation suffers from the same downside as the Band 1 and 2 post-LNA implementation.

Furthermore, the combination of the added path from the phase shifter, 45-degree transition, the longer RF path within a phased-matched OMT suitable for a circular polarization basis also introduce an additional ~3K noise to the receiver temperature compared to a linear orthomode transducer. The added mass and volume of these components also increase the cryostat size and mass, which has an impact on the antenna pointing and surface performance. These performance considerations favor linear feeds for Bands 3 through 6.

Linear polarization feeds, when combined with Altitude-Azimuth mounts that are widely separated, introduce a parallactic angle difference on sky. Effective calibration of this effect requires ‘full stokes’ imaging, where both parallel-hand and cross-hand correlation terms (i.e., XY and YX, in addition to XX and YY) are computed, which doubles the data rates from the correlator for most use cases. A circularly polarized feed basis results in a baseline-dependent phase offset in this scenario, but this can be removed through calibration without the need for the cross-hand correlation terms, possibly avoiding the need for these higher data rates. However, high dynamic range imaging and precision polarimetry still require full stokes imaging, even with circular feeds, due to the imperfect behavior of the 90-degree RF phase shifters over the receiver bandwidth, negating part of the perceived benefit of the circular polarization alternative. While these additional data products increase the compute load, they do have incremental value (e.g. producing polarization products for science projects that only request Stokes I) that increases the likelihood of archival data reuse, so their generation is a positive feature so long as the computational challenges are tractable (as discussed in Section 4.7.1).

A final consideration is harmonization of the calibration strategies for Bands 1 and 2 with Bands 3-6. Adopting a common linear-basis for all feeds simplifies the observing strategies and calibration pipeline, increasing reuse across all bands.

The decision to retain the linear polarization basis for Bands 3 through 6 was informed primarily by the bandwidth ratio limitation and the sensitivity improvement, with the harmonization of the calibration strategies as a secondary factor. A polarization calibration strategy based around this decision is also described in RD12.

Please consult the following materials for further information on this trade:

Front End: Design Description	020.30.03.00.00-0006-DSN
Calibration Requirements	020.22.00.00.00-0001-REQ

4.4.3 Bit Depth

The bit-depth of the analog to digital converter, and the bit-rate of the signal processor, both contribute to the overall efficiency of the system and also the dynamic range achievable within an observation.

The analysis of the required headroom required to prevent saturation of the digitizer is based on the VLA experience at low frequencies, and extended at higher frequencies to account for expected changes in the



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

RFI environment such as growing low earth orbit (LEO) satellite constellations and vehicular radar. The science requirements for observing bright sources, primarily the sun, are also considered in this analysis. These findings are documented in a number of memos (RD28, RD29) with ngVLA Antenna Electronics Memo #8 providing the required headroom for each band, accounting for the design bandwidth of each receiver band. The dynamic range requirements per band are summarized in Table 5.

Signal-to-Noise budgets were then built up for each band, to account for the noise of a real digitizer, the presence of RFI, the allowable gain slope across the digitized band, and the expected change in source irradiance within a single system setup. The resulting effective number of bits (ENOB) for Bands 1 through 5 is 7.35 bits. For Band 6, the requirement is an ENOB of 5.2 bits. (RD28) This ENOB definition is for the digitizer alone, and does not account for losses due to ADC clock jitter or LO jitter, which are accounted for separately in the system coherence budget. (AD04)

Table 5 - ngVLA Instantaneous dynamic range requirements for the ngVLA front ends. (RD28)

Frequency Range	Inst. Dynamic Range Required over 1 GHz	Inst. Dynamic Range Required, at specified quantization efficiency, over full band	Inst. Dynamic Range Required, at lower quantization efficiency, over full band
1.2 – 3.5 (Band 1)	26dB	26dB	30dB
3.5 – 12.3 (Band 2)	29dB	23dB	30dB
12.3 – 20.5 (Band 3)	29dB	23dB	30dB
20.5 – 34.0 (Band 4)	29dB	21dB	30dB
30.5 – 50.5 (Band 5)	29dB	20dB	30dB
70 – 116 GHz (Band 6)	15dB	6dB (20dB desired)	11dB (30dB desired)

Based on this analysis, the minimum bit-depth of the digitizer is 8 bits, with device noise constrained to an ENOB of 7.35. Such a device is expected to also be deployed in Band 6, assuming this proves to be the most cost effective solution.

The effective headroom of the system, employing an ADC of ENOB 7.35 is shown in Figure 9. The headroom and quantization efficiency achievable for an ideal 8-bit ADC, the specified ADC of ENOB 7.35, and the specified ADC inclusive of the permitted gain-slope of the preceding electronics are each shown. The efficiency inclusive of the gain slope is for the worst channel across the digitized sub-band assuming a worst-case passband gain structure. ADC saturation will be monitored and will generate online system data flags.

Please consult the following materials for further information on this trade:

- | | |
|--|----------------------------|
| Headroom, Dynamic Range, and Quantization Considerations | ngVLA Electronics Memo #8 |
| Antenna Requirements for LEO Satellite Mitigation | 020.10.25.00.00-0004-MEM-A |
| ngVLA Radio Frequency Interference Forecast | ngVLA Memo #48 |
| RFI Mitigation in the ngVLA System Architecture | ngVLA Memo #71 |

Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

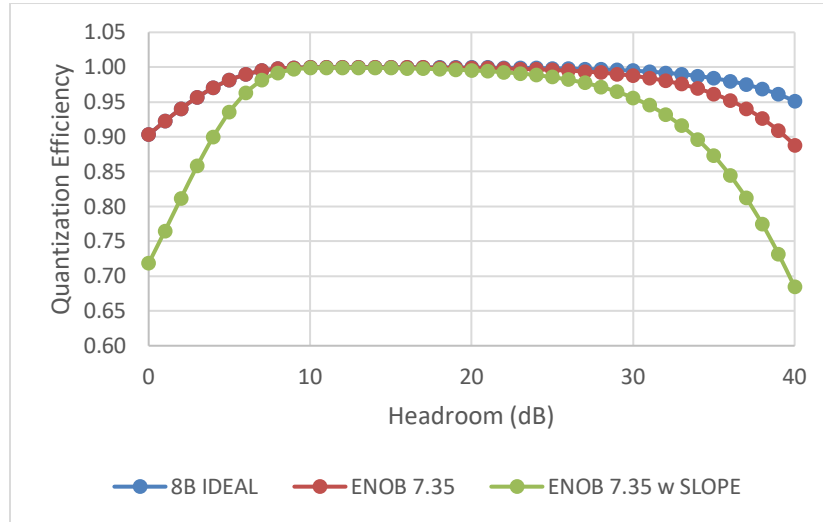


Figure 9 - Quantization efficiency of an 8-bit ADC for a Gaussian input as a function of its operating headroom, which is defined as the ADC's full-scale voltage relative to the input rms voltage. '8B IDEAL' represents for an ideal digitizer. 'ENOB 7.35' includes the additional noise such that the ADC's ENOB for a full-scale input tone results 7.35 bits. 'ENOB 7.35 w SLOPE' represents the minimum efficiency (worst channel) when the noise PSD is linear with an 8dB peak-to-peak variation across the digitized band.

4.4.4 Downconversion Topology

The selected downconversion and digitization topology of the ngVLA is novel, employing single-stage downconversion to baseband, I-Q samplers in quadrature (for bands 2-6), with dedicated downconverter-digitizers for each receiver band. This architecture is quite different to the direct sampling architecture of the SKA, or the intermediate frequency and multi-stage conversion architecture of the VLA or ALMA. The factors that informed this trade are considered below.

Direct sampling is an attractive solution where the RF frequency sampled is within the analog input bandwidth of the digitizer. Direct sampling is used for ngVLA Band 1 (<3.5 GHz) for these reasons. This solution is not considered technically mature at higher frequencies for the bit depths required in Section 4.4.3. No production ADCs have been identified that operate above X-band, so a direct-sampled solution is not applicable for bands 3-6.

Multi-stage conversion with a fixed intermediate frequency band is the architecture of the VLA and ALMA. This architecture is best suited to connectorized components that are integrated into modules. This architecture has a high level of technical readiness, but suffers from a few negatives that impact its viability on the ngVLA:

- The equipment is bulky, occupying two 2m tall racks at the VLA. Given the offset Gregorian geometry, this equipment should be located on the feed arm for the ngVLA, but would violate volume and mass constraints of the antenna design. An assessment of the available space on the feed arm, and the sensitivity of the antenna pointing and surface performance to the mass supported by the feed arm, suggest this approach is not practical for the preferred offset Gregorian optics in an 18m aperture size that does not readily incorporate a receiver cabin.
- Alternatively, equipment could be partially located in the pedestal which introduces long electronic delays and gain slope in the intermediate frequency or baseband cable runs, and associated electronic delay management concerns. The length of the signal path will change with thermal fluctuations in the atmosphere and antenna enclosures, and introduce proportional phase changes in the digitized signal.



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

The ngVLA has tight phase and gain stability requirements, more comparable to ALMA than the VLA, to support the required imaging dynamic range of the system and ensure that the achievable sensitivity of the system is not dynamic range limited. Minimizing the total path length reduces the total phase change as a function of temperature, and is a desirable parameter in the design of the antenna electronics. The long path lengths associated with multi-stage conversion present a significant performance risk to the design.

- Multi-stage conversion requires a series of RF switches to select the respective receivers and route RF, IF and LO signals through the selected topology. The EVLA implementation includes a total of 13 connectorized packaged switches in each antenna. These switches are a frequent source of failures and introduce subtle amplitude fluctuations that impact system calibration and can limit the imaging dynamic range.
- ngVLA will be the first telescope operating at centimeter wavelengths that aims to deliver high-level data products with each standard observing mode. Improving the feasibility of system calibration is therefore a design driver for the electronics system design. The switching and tunability inherent in a multi-stage downconverter-digitizer topology conflicts with this goal.

The operations cost caps, and associated need for maintainability and reliability, favor reduced parts count and the elimination of mechanical switching components. Size and mass constraints at the secondary focus strongly favor integrated compact approaches. Component miniaturization and integration into MMICs for mass production would be highly desirable given both design constraints.

A downconverter-digitizer topology that minimizes electronic path lengths, eliminates the need for RF and LO switching, and is highly compact and thermally stabilized would seem to support the calibration and operations requirements, while respecting the size and mass constraints imposed by the antenna selection. The selected single-stage parallel down-conversion topology, employing a high degree of component integration, supports all these requirements.

While the selected concept is considered the optimal design solution, it does have risks that must be considered and mitigated as part of the design activities:

- The preferred implementation would use an ASIC digitizer/serializer/formatter to reduce the total size and power required, while also improving reliability with reduced parts count (RD32). This ASIC should be proven, on silicon, as soon as possible to retire the associated risk to changes in the packaging at the secondary focus. A fallback solution would rely on discrete components for the digitization stage, but may require appreciably more mass and volume, as well as up-sized power and environmental control system interfaces that could ripple through the antenna design.
- There is limited use of I-Q sampling architectures in radio astronomy so practical experience with the calibration of these devices must be developed. The system-level requirements for sideband separation have been derived and appear to be technically feasible, but proving this with an early prototype in a representative environment would largely mitigate this performance risk.

Please consult the following materials for further information on this trade:

Integrated Receivers and Digitizers Design Description	020.30.15.00.00-0004-DSN
Downconversion and Digitization Methodology for the ngVLA	ngVLA Electronics Memo #1
An Integrated Receiver Concept for the ngVLA	ngVLA Memo #29

4.4.5 Cryogenics System

The cryogenics system concept is closely related to both the receiver band definition (Section 4.4.1) and the feed indexer geometry (Section 4.3.4). The cryogenic cooling load is the largest component of the VLA electrical budget, as well as the most frequent antenna sub-system to receive corrective maintenance.



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

Given the close relationship of this system to the maintenance cost and maintenance model, it is a goal for ngVLA to improve both the electrical efficiency and MTBM by a factor of three compared to the VLA, with the associated detailed requirements captured in the MTBM and electrical budgets.

The VLA uses a Gifford-McMahon cycle system with a two-stage cryocooler per receiver band and three shared cryogenic compressors. The system is sized for a reasonable system cool down time, generally reaching operating temperatures within 8 hours after receiver service.

There are four Cryogenics related trades considered in the ngVLA system design: LNA operating temperatures, cryostat integration, Stirling vs GM cycle, and variable frequency operation.

The LNA operating temperature is a key design constraint established by the front end system. VLA operates on a two-stage system, with 70K and 15K stages. ALMA adds a third 4K stage for SIS mixers, with the LNA bands on the 15 K stage. At ngVLA operating frequencies, the LNA noise temperature contributions do not dominate the system temperature budget. Arguably, Band 1 InP LNAs could operate at 70K, while Bands 2-5 GaAs LNAs are well suited to two-stage cryocoolers. Band 6 approaches frequencies where a third stage could be considered, but the VLA site precipitable water vapor will lead to larger atmospheric contributions to system temperature than ALMA experiences, which offsets some of the relative gain of including a 3rd stage. At 15K, T_{LNA} is anticipated to be 26K within a T_{REC} of 49K at 90 GHz (RD25). This T_{REC} is a comparable value to T_{SKY} at 13mm of PWV and 45-deg elevation. Dropping the LNA physical temperature to 5K with a 3rd stage would reduce T_{LNA} by approximately 2K, or approximately a 2% reduction in T_{SYS} . This is not necessary to meet the system level requirements for system sensitivity. Given the added cost and complexity, and more importantly the opportunity to integrate the receivers in a single cryostat for the high frequency bands, a two-stage cooler is the preferred implementation.

Cryostat integration presents a clear opportunity to reduce the cryogenic system capital and maintenance costs. The receiver band definition and feed indexer geometry permit receiver integration into two cryostats, with Bands 2-6 located in a single cryostat serviced by one cryocooler. This reduces the total cryocooler count from six to two compared to a scaled VLA system, and the compressors can be similarly integrated into a single compressor. The reduced parts count dramatically improves the MTBF (three key components in a series reliability model vs nine components in the VLA case.) and reduces the capital cost as well. The integration also yields a marginal improvement in electrical efficiency, as the cooling load on the cryocooler has a roughly linear relationship with the size of the radiation shield. Integration of multiple receivers into a single cryostat reduces the total radiation shield area compared to discrete receivers (i.e., consider cube-shaped radiation shields; the integrated case can eliminate a cube face, or about a 17% reduction in radiation shield area.) We will assume two-stage cooling with two cryocoolers and cryostats for all further cryogenic trades discussed below.

The next trade concerns the cryogenic cycle with an emphasis on Stirling-cycle pulse tube coolers and GM based cryocoolers. The Stirling cycle systems are inherently more efficient and lower maintenance due to the elimination of Teflon rotating valves in the cryocooler and the recovery of energy in the return phase of the cycle. However, a Stirling cycle pulse tube has an elevation dependence on the cooling capacity, a higher capital cost, and requires that the compressor be mounted close to the cryocooler on the feed arm. Two-stage pulse tubes are also not commercially available in the correct capacity. A GM system is less efficient and has higher operating costs, but lower capital cost and a higher level of technical readiness. Both options were explored early in the facility concept development phase (RD48) with two studies performed for Stirling cycle systems with Northrop Grumman and RIX Industries during the conceptual design phase (RD47). While the second study with RIX in particular yielded a design that could meet the key performance requirements with the lowest lifecycle cost, a GM system is preferred at this time due to the high level of technical readiness, improved thermal stability over elevation, and the reduced



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

mass and volume demands on the antenna feed arm. The Stirling cycle system will be further developed in parallel leading up to the PDR to see if the solution can be suitably matured and risk retired before eliminating this option. (RD49)

Both the GM and Stirling cycle solutions can have their efficiency improved by implementing variable frequency operation. For a GM based system, this has the added benefit of extending the service interval (with corresponding improvements in MTBM) proportional to the average speed reduction, as the most frequent wear item are the Teflon seals in the cryocooler that cycle at the operating rate. Variable frequency drives do present an EMI risk, and any variable frequency system will require careful shielding and filtering which likely precludes the use of existing commercial products.

As mentioned above, the VLA system runs at a single speed. The system must therefore be sized not for steady-state operation, but for cool down, to achieve the required operating temperatures within a reasonable time frame, typically about 8 hours. This in effect means that the system is continuously operating at a faster rate than required, with an associated loss of efficiency. By adopting a scroll compressor (suitable to multi-speed operation) and variable frequency drives for the cryocoolers, the speed of each cryocooler can be cycled to match the present cooling demand. Early laboratory tests suggest a ~30% reduction in speed is feasible after the cryostat reaches the desired operating temperature, with corresponding improvements in MTBM and electrical efficiency (RD37). The tradeoff here is in the added complexity of the variable frequency drives, the associated loss of reliability from increased parts count, EMI risk, and the added capital expenditure. The ~30% reduction in speed results in an electrical savings of approximately \$3500 per cryocooler per year. With a projected capital expenditure of \$10K per VFD, the reduction in speed is sufficient to offset these costs within three years of operation. The increase in the failure rate from the VFDs must be offset by the lower failure rate of the cryocoolers for a cost-neutral total solution, and this requirement is reflected in the total MTBM allocation for the cryogenics system.

Please consult the following materials for further information on this trade:

Cryogenics System Design Description	020.30.10.00.00-0007-DSN
Thermoacoustic Stirling Cryocooler and Variable Speed	
Gifford McMahon Cryocooler Trade Study No. 2	ngVLA Electronics Memo #4
Advanced Cryocoolers for the Next Generation VLA	ngVLA Memo #24

4.5 Time & Frequency Reference Distribution

4.5.1 Antenna Local Oscillator Architectures

The Antenna local oscillator (LO) architecture is closely tied to the selected concept for downconversion (Section 4.4.4). The parallel downconverter topology permits both a traditional tunable synthesizer approach to generating the local oscillator signals as well as dedicated oscillators for each downconverter. A small degree of tunability is still required in the latter case to permit the application of per-antenna LO frequency offsets. These are necessary to provide continuous frequency coverage across each RF band, with no gaps at the system level (though this does require a dip in sensitivity). These frequency offsets also improve the suppression of image bands and spurious signals in the single stage downconverter-digitizers.

Four concepts for the antenna LO system were considered in an associated trade study (RD51). Options considered include:



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

- Option A: Discrete phase locked dielectric resonant oscillators (PLDROs) for each downconverter, with only limited tuning range.
- Option B: Shared (within the antenna) RF synthesizer LO references amongst downconverters.
- Option C: Photonic LO transmission from synthesizers in the pedestal.
- Option D: Photonic LO transmission from centrally located remote synthesizers.

A number of hybrid alternatives are also considered, such as discrete oscillators for bands 1-5 and tunable synthesizers dedicated to Band 6 (An Option A-B hybrid).

Based on a combination of performance factors, interface considerations, cost and technical risk, the preferred option for the conceptual technical baseline is Option A. This option has a high technical readiness, high performance, future extensibility, and acceptable size and mass on the feed arm. A simple block diagram for this concept is shown in Figure 10, and the frequency scheme is summarized in Table 6.

RF Band	Module	LO		DRO output frequency	multiplier value	harmonic compared to 2.9 GHz
		harmonic	(GHz)			
1	a					
2	a	2	5.8	5.8		2
	b	4	11.6	11.6		4
3	a	5	14.5	14.5		5
	b	7	20.3	20.3		7
4	a	8	23.2	23.2		8
	b	10	29	29		10
	c	12	34.8	34.8		12
5	a	11	31.9	31.9		11
	b	13	37.7	37.7		13
	c	15	43.5	21.75	2	7.5
	d	17	49.3	24.65	2	8.5
6	a	25	72.5	18.125	4	6.25
	b	27	78.3	19.575	4	6.75
	c	29	84.1	21.025	4	7.25
	d	31	89.9	22.475	4	7.75
	e	33	95.7	23.925	4	8.25
	f	35	101.5	25.375	4	8.75
	g	37	107.3	26.825	4	9.25
	h	39	113.1	28.275	4	9.75

Table 6 - Selection of PLDRO frequencies and multiplier values for each of the IRD modules. The fifth column is the PLDRO output frequency when locked to the reference. The items that are colored red indicate instances where the DRO may contain an internal doubler. Note that the resulting frequencies in column 5 are all direct multiples of 362.50 MHz (=2.9 GHz/8).

Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

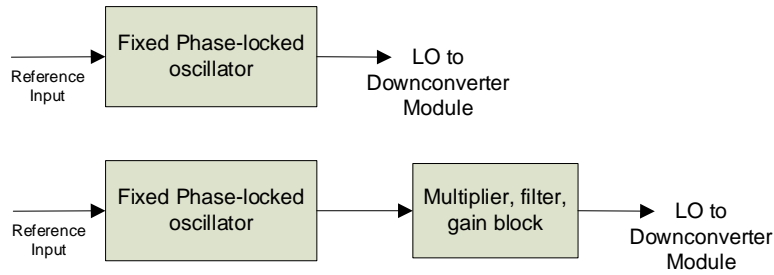


Figure 10 - LO Architecture A – Independent, fixed frequency oscillator model for each IRD downconverter module. A total of nineteen LO modules is needed with this approach. Per antenna frequency offsets are applied on the reference input frequency.

Development will also be pursued into the A-B hybrid and Option C leading up to the PDR, as possible future enhancements to the design. The A-B hybrid would permit the project to eliminate four of the downconverter-digitizer modules in Band 6, and half of the input bandwidth transmitted to the Digital Back End system, with significant system cost savings.

The photonic transmission development provides a risk reduction strategy for the mass and volume constraints at the front end enclosure, as well as a reduced parts count at the front end to improve maintainability and reliability. Finally, this photonic development is a necessary precursor to considering centrally located remote synthesizers or other architectural simplifications at the front end, though the centralized option clearly has higher technical risk given the transmission distances. The readiness of both development efforts will be assessed by the system PDR.

Please consult the following materials for further information on this trade:

ngVLA Antenna Local Oscillator Trade Study
LO Reference & Timing Design Description

020.30.35.00.00-0003-REP
020.35.00.00.00-0004-DSN

4.5.2 Frequency Reference Distribution

The selection of the Antenna LO architecture has a strong influence on the frequency reference distribution system. Given the maturity of the antenna design, the project has elected to complete the former trade first. The study of best technique for frequency transfer is still being conducted and will be completed prior to a downselect during the LO Reference and Timing subsystem conceptual design review (CDR).

The primary method being studied involves the distribution of a fixed frequency reference to each antenna for all antennas within approximately ~300km of the array core. A round-trip phase measurement is envisaged so that delay variations on the fiber path can be removed. Independent frequency references, locked to active hydrogen masers, would be provided for mid-baseline antennas and LBA sites outside the 300 km radius.

The LO frequency plan is currently arranged so that each LO frequency is a multiple of 2.9 GHz, and thus the frequency transfer design plan being studied involves transmission at this frequency or at a harmonic of this frequency.

In parallel, an additional design is being developed at NRC-Penticton. The main feature of this technique is that it absorbs the frequency reference transmission into already existing portions of the CSP and antenna-based data transmission system hardware.



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

Please consult the following materials for further information on this trade:

LO Reference & Timing Design Description

020.35.00.00.00-0004-DSN

4.5.3 Timing Reference Distribution

The concept for timing reference distribution will be advanced in parallel to the frequency reference distribution trade for the LO Reference and Timing subsystem CDR. The project is presently comparing the use of two techniques to provide ~nsec accuracy timing distribution to each antenna station.

- An optical hardware based approach utilizing direct lightwave modulation with two-way propagation and delay measurement.
- A White-Rabbit (an open-source synchronous Ethernet) solution.

Please consult the following materials for further information on this trade:

LO Reference & Timing Design Description

020.35.00.00.00-0004-DSN

4.6 Central Signal Processor

4.6.1 FX vs XF Correlator Architecture

NRAO has institutional experience with both FX and XF correlator architectures: the VLA WIDAR and ALMA correlator (in Frequency Division Mode) are both FFX architectures, while the VLBA DiFX software correlator uses an FX approach (as the name suggests).

A practical correlator supporting the required 20 GHz of processed bandwidth will necessarily include an initial F-stage before the cross-correlator; this limits the possible architectures to FFX and FFX. This first F-stage is a polyphase filter bank (PFB) used to generate sub-bands which can be processed by later-stage processors at clock rates suitable for a massively-parallel processing architecture. An additional opportunity presented by the first F-stage is the generation of duplicate sub-band data streams that can interface with future commensal back-end instruments (such as a commensal SETI or FRB search instrument) by including a crossover switch between the first F-stage and the subsequent FX or XF sub-band processing stage.

Fundamentally both FX and XF architectures are possible, but each presents different requirements for the preceding F-stage. An FX architecture can accommodate a wider sub-band bandwidth, while an XF architecture will require narrower sub-bands to respect data rate limits on the number of lags that must be generated by the X engine.

A key differentiator between the two in the ngVLA context is the ability to leverage development work from partners – for example, an FX architecture is consistent with the FSA architecture developed at NRCC for the SKA-MID telescope and under consideration for the ALMA correlator upgrade. With no inherent feasibility differences between the two architectures or explicit design requirements which would mandate one or the other, adopting an FX architecture would enable broader collaboration opportunities with other next-generation signal processor initiatives including SKA-MID and the ALMA correlator upgrade.

An additional motivation for selecting an FX architecture for the sub-band processor is its compatibility with alternative back-ends, such as the pulsar engine or dedicated SETI and FRB search instruments. Although these back-ends could be required to process the sub-band data immediately after the first F-engine, it would be much more efficient to duplicate the data after the delay correction and further frequency channelization carried out at the FX sub-band processors. An XF architecture would require



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

each back-end to implement a separate F-engine where one was required. The FX architecture may additionally re-use its F-engine for phase-delay beamforming (see Section 4.6.4 for more details), while an XF architecture would require an additional F-engine specifically for this purpose.

An FX architecture is also much more efficient at implementing the extended window functions needed to satisfy the ngVLA's channel purity requirements. For example, it can employ polyphase DFT (FFT-based) filter banks in which the spectral decimation (the removal of highly overlapped frequency channels to decrease information redundancy) is actually performed in the time domain prior to the frequency transform. As a result, the redundant frequency channels are never computed, which results in significant computational savings. In contrast, with an XF architecture, the number of lags required for the desired channel response increases significantly, with no ability to avoid redundant computations or reduce the required lag range at the X-engine.

For all the reasons above, the FX architecture is considered the most suitable for the CSP sub-band processors, and an overall FFX architecture has been adopted for the ngVLA CSP.

Please consult the following materials for further information on this trade:

Central Signal Processor Design Description

020.40.00.00.00-0005-DSN

4.6.2 Correlator-Beamformer Integration vs Independence

The ngVLA is presently carrying two sub-band processor designs forward into the CSP conceptual design review. Both designs employ an FX architecture for interferometric mode operation, consistent with the decision presented in Section 4.6.1, but differ in a key architectural choice: integrating correlator and beamformer modes into a single type of sub-band processor or retaining distinct hardware for each mode. A down-select between these two alternatives is expected at the central signal processor sub-system conceptual design review.

The Trident Frequency Slice Architecture (FSA) is a mature FPGA-based design that uses a common sub-band processor platform (the TALON-DX board) as either a beamformer or a correlator for all antennas. It has been proposed for the ngVLA CSP after its successful adoption by the SKA1-Mid telescope. This design includes 28 GHz of sub-band processing to accommodate the requisite subarray commensality requirements of 20 GHz of processed bandwidth in interferometric modes concurrent with 8 GHz of bandwidth in the beamformed modes. This integrated solution makes use of reprogrammable FPGAs, supporting the sub-array commensality described in Section 4.6.5.1 with 1.4x the sub-band processors required to support full capabilities in a single mode (RD43).

The SCREAM (SCalable REconfigurable And Modular) architecture features an independent X-engine and combined F-engine/beamformer subelement that is available concurrently in all subarrays. SCREAM's F-engine/beamformer is based on a fine-grained design in which each antenna-sub-band (or sub-band pair in normal resolution mode) has corresponding hardware which is not shared with any other antenna's device in the same sub-band processor. This allows the F-engine/beamformer to achieve full subarray independence, as well as complete commensality of interferometric and beamformed modes within a single sub-band processor. The X-engine is likewise being designed with full subarray commensality as one of its core capabilities. As a result, the SCREAM design supports different modes for subarrays within a single sub-band. Subarrays are supported in all six operation modes at full specification with the same number of sub-band processors required to fully support a single mode.

The choice between integrated and separated correlator-beamformer designs has significant consequences outside of multimode operation. For example, the FSA's correlator stage is distributed across the FPGA devices which form every sub-band processor. The FSA further specifies that these devices communicate through a passive optical network with a mesh topology. The resulting distributed



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

X-engine is sized for the worst case (i.e., the entire array), and is continuously computing (and then discarding) the cross-correlations of antennas assigned to different subarrays. This strategy does not impose any restrictions on the spectral or frequency resolutions that are independently configured for each subarray; however, it is not possible to put any distributed elements in a low-power state even if the subarray configuration would suggest fewer units are required.

In contrast, nodes in the SCREAM design's X-engine can receive a varying total bandwidth to ensure optimal resource utilization for different sizes of subarrays. As the number of visibilities can decrease dramatically depending on the subarray configuration (while the total input bandwidth remains constant), a significant number of computing elements may be set in a low-power state when not required. This flexibility does, however, require an actively-switched network to properly route data from the dedicated F-engine/beamformer hardware to the allocated X-engine resources.

An additional benefit of SCREAM's separated beamformer-correlator architecture is that beamformed data streams may be themselves correlated against other receivers or subarrays. Although this capability is not currently part of the ngVLA design requirements, and it is therefore not under active development, it remains a potential future expansion and further illustrates the potential of a flexible architecture with distinct correlation and beamforming components.

Please consult the following materials for further information on this trade:

Central Signal Processor Design Description

020.40.00.00.00-0005-DSN

4.6.3 ASICs vs FPGA vs GPU Implementations

The ngVLA CSP design consists of heterogeneous subelements performing the varied DSP tasks required to support system requirements. This includes the digital backend (DBE); data transportation network; pulsar engine; and per-sub-band channelization, correlation, and beamforming (integrated or separated). Each of these subelements 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, SMA), FPGAs (VLA, SKA-MID, SMA) 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 for small volumes, while the reduced power draw of ASICs 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 sixteen active ADCs (thirty-two in some down-conversion schemes), each at 56-Gbps, while transmitting sub-band data streams at an overall data rate of almost 1 Tbps. While some



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

modern FPGAs can support these data rates with a single device, a GPU-based system would require several devices (and additional networking hardware) to support the same overall bandwidth.

For the switched fabric, a solution featuring Commercial Off-The-Shelf (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. RD44 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 Commercial Off-The-Shelf (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.

The SCREAM B&C nodes have data-interchange requirements which 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 10 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; 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 their 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 (RD45) will be closely monitored until a final design downselection is made. As detailed in RD50, 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.

Please consult the following materials for further information on this trade:

- | | |
|--|----------------------------|
| Central Signal Processor Design Description | 020.40.00.00.00-0005-DSN |
| A SCREAM-Compatible ngVLA Pulsar Engine: Key Requirements Review and Option Trade-Off Study | ngVLA Electronics Memo #11 |
| A SCREAM-Compatible ngVLA Cross-Correlation Engine: Key Requirements Review and Option Trade-Off Study | ngVLA Electronics Memo #10 |



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

4.6.4 Phase-Delay vs True-Delay Beamformer

The phased-array mode of the correlator calls for beamforming capabilities, which may be obtained through two different methods:

A “true time delay” beamformer applies an explicit time delay to the incoming voltage stream prior to channelization. This delay correction process has been routinely used in radio astronomy interferometers to track a specific reference position on the sky. The main differences are that a beamformer usually tracks multiple reference positions at once, pointing one beam formed at each, and the signals from the antennas are linearly combined instead of cross-correlated. The delay tracking process would be carried out at the central building, by the sub-band processors. Given the ngVLA beamforming efficiency requirements, this time delay must be applied at a sub-sample resolution determined by the sub-band bandwidth and oversampling factor. Sub-sample correction is a computationally expensive process, but this approach supports phasing over large apertures and bandwidths while incurring in small beamforming efficiency losses.

A “phase delay” beamformer is computationally inexpensive, and can therefore generate a great number of beams; however, the beamforming efficiency degrades as the beams deviate from the reference position and as the aperture of array increases. The cause of this inefficiency is that for a channel of nonzero width, the phase (a linear function of frequency) is approximated as a step function, which is only correct for the frequency at the channel center. The beamforming efficiency will depend on the error introduced, which is related to the channel width and the slope of the true phase function. The sub-band processors minimize the slope to be applied by tracking the delay of a reference position in the sky. The beamforming error is therefore related to the residual delay between each beam’s pointing direction and the reference direction in the sky. As a result, the phase approximation error component related to the slope of the phase function depends on the beam’s offset from the reference position and the aperture of the array.

The greatest aperture size that must be phased for the ngVLA is driven by the VLBI and pulsar timing use cases. Phasing over the full aperture of the main array is required in the LI system requirements, with a beamformer efficiency greater than 95% of the SNR achievable with an ideal beamformer (AD04). These specifications can only be supported with a true delay beamformer, so a true delay beamformed mode will therefore be necessary for both the VLBI and pulsar timing modes.

For a pulsar search mode that has many (hundreds) of beams, a phase delay mode would also be required to efficiently use the CSP resources and generate the requisite number of beams. ngVLA requirements (AD04) for pulsar search focus on small fields of view and only require 10 beams (with a goal of 50). With the limited number of beams required, a true time delay beamformer can support all the required beamforming modes, and a separate phase delay mode should not be required. The higher efficiency, and therefore sensitivity, of a true time delay mode is also more responsive to the ngVLA science goals for pulsar search, which focus on the detection of weak pulsars that are not detected by other instruments.

It is important to note that the goal of 50 beams at reduced bandwidth may be enabled in the CSP architecture by requiring multicast capabilities at the switched network. The same sub-band can then be processed at multiple sub-band processors, each generating the full 10 beams. In order to offset the increased number of sub-band processors this sub-band requires, the maximum bandwidth of each beam must be reduced proportionately. Trading in bandwidth for beams also reduces pressure on subsequent systems, as the overall data output from the CSP in pulsar search mode is proportional to the product of the bandwidth and number of beams.

Please consult the following materials for further information on this trade:

Central Signal Processor Design Description
ngVLA System Requirements

020.40.00.00.00-0005-DSN
020.10.15.10.00-0003-REQ



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

4.6.5 Additional Design Considerations

4.6.5.1 Subarray Commensality and Independence

The ngVLA is expected to operate in subarrays for a majority of observing time, and subarray use is central to the operations concept for the facility. The operations concept prescribes that maintenance, test and commissioning activities be primarily conducted within subarrays, with no regularly scheduled “maintenance days” where the full array is dedicated to maintenance and testing.

The degree of independence between concurrent subarrays can be a design driver for the central signal processor and online software system, so settling the degree of independence amongst commensal subarrays is a key design decision that informs the architecture of the central signal processor.

An assessment for the ngVLA identified that the minimum operating mode commensality required would support full capabilities in the interferometric imaging mode concurrent with 8 GHz of processing in a beamformed mode. This combination supports the reference observing program and envelope observing program, ensuring efficiency in science operations without infrequently-used redundant processing resources. This reduces the number of subarray processors needed from six times what is required for single-mode full capabilities (20 GHz) to just 1.4 times. This factor is also dependent on the trade-off between integrated or independent correlator and beamformer modes, which is considered in Section 4.6.2. A similar analysis was performed for SKA (RD42) and incorporated in to the SKA-MID CSP design. The supporting design for ngVLA is documented in RD35.

An operational concern to consider is the degree of independence in the software configuration of each subarray. Maintenance and commissioning activities will favor supporting a high degree of independence for testing purposes, but this independence can also add design complexity and cost, so once again this must be approached as a best-value study.

We note that subarray software independence is highly desirable for sub-scale testing and necessary to support a continuous integration software development and test approach. A maintenance and development model where software is tested on the array rapidly, within days or weeks of development and not months or years, can significantly increase the productivity of the software development team and avoid the need for managing multiple long-lived software branches as part of the development and release process. Therefore, an ability to deploy all or most of the software stack independently for a subarray is highly desirable. Conversely, many issues do not appear in real-time software systems until tested at full scale, so the facility will always need to reserve time for commissioning with the full array.

Acknowledging the value, we will aim for full software stack independence as a design goal for the central signal processor. As a goal (rather than a requirement) the design team will aim to implement this functionality but will carefully consider otherwise interesting parameter space that meets the sub-system requirements and lifecycle cost targets. This factor will be considered in more detail at the CSP conceptual and preliminary design reviews.

Please consult the following materials for further information on this trade:

Central Signal Processor Design Description	020.40.00.00.00-0005-DSN
Reference Observing Program	020.10.15.05.10-0001-REP
Envelope Observing Program	020.10.15.05.10-0002-REP



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

4.6.5.2 Digital Back-End at the Antenna vs Central System

As described in Sections 4.6.1 and 4.6.5.1, the CSP follows an FFX architecture in which the Digital Back-End (DBE) implements the first F-engine to generate sub-bands and the sub-band processors carry out subsequent FX processing on a per sub-band basis.

One of the first questions raised in the design was the location of the DBE - at the antenna or in the central facility along with the rest of the CSP. Locating the DBE at the antenna is mandatory for the remote stations, per system requirements, as their sub-bands must be generated and down selected before being transmitted to a central facility in order to reduce operational costs of the network infrastructure. Consequently, the question of DBE location only applies to the nearby antennas, which are connected through facility-owned fiber. The main advantages of locating the DBE in the central building are ease of maintenance and the possibility of using a power-efficient passive optical mesh between the DBE and the sub-band processors. Any benefits from an integrated design in which a piece of equipment processed several signals coming from multiple antennas was discarded due to the complexity of the DBE and current technological limits. Moreover, the use of a passive network was finally abandoned (see Section 4.6.5.3).

In contrast, locating the DBE at the antenna has its own advantages. For example, it reduces signal integrity risks from a link between the digitizers and the DBE that could extend for several kilometers, it allows timestamping to occur closer to the digitizer, and it also allows increasing the resolution of the digitizer in future upgrades without necessarily having to increase the fiber capabilities. However, no significant decrease in the maximum output data rate is obtained.

None the above considerations have been sufficient to make the decision obvious. Since the remote stations require their DBEs be located at the antenna, it has been decided to extend this to the closer antennas as well in order to minimize the development effort.

Please consult the following materials for further information on this trade:

- | | |
|----------------------------|---------------------------|
| Trident 2.0 Concept (RD46) | ngVLA Electronics Memo #4 |
| Trident 2.1 Concept (RD43) | ngVLA Electronics Memo #5 |

4.6.5.3 Active vs Passive Networking

The CSP reference design is based on the TRIDENT FSA, which utilizes a point-to-point passive fiber optic network to connect the first F-engine and the remaining FX processors. This solution minimizes initial hardware cost and power consumption of the network. However, this approach does not scale well to a correlator of the ngVLA’s size, as the number of transceivers required at each F-engine device exceeds what is possible with the technology employed. As a result, the FSA correlator is split into three smaller correlators (hence “trident”), each processing 10 GHz of bandwidth; this introduces hardware redundancies which negate any power reduction benefit.

For its conceptual design, the CSP employs an active switched fabric to distribute the sub-band data generated at the antennas to any desired sub-band processor. Prior to this, a significant effort was made (see RD46) to eliminate hardware redundancies from the Trident design while retaining the advantages of passive networking. That solution was ultimately abandoned because it lacked the required extensibility to support future commensal back-ends, imposed limitations on the routing flexibility necessary to completely satisfy commensality requirements, and presented compatibility issues with remote antennas that require sub-band generation at the antenna location.

Please consult the following materials for further information on this trade:

- | | |
|----------------------------|---------------------------|
| Trident 2.0 Concept (RD46) | ngVLA Electronics Memo #4 |
| Trident 2.1 Concept (RD43) | ngVLA Electronics Memo #5 |



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

4.7 Post Processing System

4.7.1 Real-Time Processing vs Post-Processing

The selection of a near real-time processing system vs a post-processing system with reprocessing capability is a key architectural choice for the array. In this trade we focus on the processing of science data products for delivery to observers and retention in the science data archive. Intermediate data products such as ‘Quick Look’ images or data quality assurance products (QA0) may require a tailored implementation since the requirements differ, with a key functional requirement to rapidly calibrate and produce rough images to inform subsequent observations.

In this context, “real-time” will be used to refer to systems that process the visibilities once, generating high level data products such as image cubes, and subsequently discard the visibilities. Such systems must be sized, at a minimum, to match the average throughput of the central signal processor. The SKA Science Data Processor is an example of such an architecture.

“Post-processing” will be used to refer to systems that record the low-level data products (e.g., visibilities) for the life of the instrument, and generate high-level data products asynchronously. Post processing can be accomplished on either observatory-managed resources or user resources, depending on the observing model. ALMA and the VLA are examples of each model, with ALMA relying more on central observatory-managed resources for the generation of higher level data products.

The post-processing approach is more familiar to the project team, given VLA and ALMA experience, and is expected to provide a greater degree of reuse of existing post-processing software that would reduce both the technical risk and total development effort. The approach also enables future reprocessing, providing corrected or enhanced data products and improving the “science multiplier” of the archive. For these reasons, post-processing is the preferred and default approach.

However, a key risk must be assessed to determine if post processing is feasible: the average data rate generated by the central signal processor must be feasible to store, both technically and economically within the constraints of the operations budget.

This risk was initially determined to be manageable, and the preference for post-processing defensible, based on trades conducted with the ngVLA quantitative exchange model. The model was updated with a set of use cases that approximate the Reference Observing Program (AD09), scaled to use all available observing time. The data rates and processing complexity vary dramatically amongst use cases, with ratios of 1000s between the simplest and most demanding use cases, so a sensible and defensible set of use cases is a key input.

The model suggested an average of 9000 giga-visibilities would be generated per hour, for a total of 13.2 PB of data generated and stored each month. Processing this data would require 9800 core-hours per telescope observing hour, using a reference 2016 single-core processor.

The archiving costs would be prohibitive in 2016 USD (using 2016 technology), totaling \$67M per year. However, with a conservative assumption about the rate of decreased cost of storage in the future, assuming the cost of storage drops by half every 36 months, the reduces the annual data storage cost to less than \$900K as the system reaches full operations in 2035. With these inputs and assumptions, the annual archive costs would peak within 2 to 3 years after the commencement of full operations, suggesting the model is supportable for the life of the instrument.

Of course, these results are highly dependent on the continuation of trends in the cost of data storage, so these trends should be monitored throughout the system design phase. The cost curve for disk-based storage appeared to be flattening from 2016 to 2020, raising concerns. However, NAND-based solid-



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

state drives are now cost competitive with disk storage and are exhibiting a 30% cost reduction per year (twice the former disk rate), suggesting the model assumptions remain defensible for the near future.

The early results from studies with the quantitative exchange model were subsequently refined in a much more sophisticated compute scaling study and memo (RD48) using the reference observing program (ROP) as a representative data set for analysis. This model is based on the measured single core performance of existing code implemented in CASA, scaled to ngVLA data rates with appropriate parallelization efficiency factors for each step in the data reduction process. The recorded data rates and corrections applied in calibration and imaging are tailored to the requirements of each use case in the reference observing program, and are weighted to reflect a full observation calendar. This refined model suggests the system will generate 20.1 PB of data each month, and will require a 50 PFLOPs/sec system to satisfy a full observing calendar based on the referenced observing program use case distribution. The data processing requirements are highly skewed, with use cases representing 10% of observing time in the ROP requiring 44 PFLOPs/sec of capacity, and the remaining 90% of time processed with a 6 PFLOPs/sec system. A 60 PFLOPs/sec compute system is adopted as both the baseline design and a constraint for the construction project. I.e., the ngVLA project will deliver a 60 PFLOPs/sec computing system within the construction project scope. The operations team will need to either treat this as a finite resource that is allocated by priority (much like observing time on the instrument) or secure additional computational resources outside the scope of the construction project. The 60 PFLOPs/sec figure was adopted based on the 50 PFLOPs/sec estimate from the study, combined with a 20% allocation for reprocessing.

Please consult the following materials for further information on this trade:

- | | |
|--|-------------------------|
| Size of Computing Estimates for the ngVLA | ngVLA Computing Memo #2 |
| The ngVLA Quantitative Exchange Model | nqxm Ver3.0 |
| ngVLA Reference Design Development & Performance Estimates | ngVLA Memo #17 |

4.7.2 Low-Level vs High-Level Data Products

The decision to generate only low-level data products vs low-level and high-level data products with observatory resources is closely tied to the computational complexity associated with data reduction. For the VLA and ALMA, the data volume (GB-scale) associated with a typical observation can be transferred over a broadband internet connection. The post-processing complexity varies with each use case, but typical use cases can be processed with single high-performance desktop computers or small computer clusters such as those found at research universities.

The data sets expected for typical ngVLA observations (100 TB scale) are not practical to transfer over broadband to most users, and the compute requirements (PFLOP-scale) exceed the projected capacities available to a majority of users, even those associated with research universities and computing research centers when ngVLA approaches full operations.

Given these practicalities, the generation of high-level data products must be an observatory-performed function. This need is captured in the ngVLA Operations Concept (AD03) and flows to the Stakeholder requirements and subsequently System requirements.

The computing capacity required within the system design to support the generation of these high-level data products is analyzed in the compute scaling study and memo (RD48). The use of observatory-run resources vs cloud resources to provide this capability are considered in Section 4.7.3.

In addition to the standard data deliver model, there are special use cases associated with Large and Legacy Projects (AD15). In such projects, sophisticated users may create their own high-level data products with custom pipelines operating on cloud services or research clusters, necessitating the transfer of low-level data to these user-provided computing systems for processing. Such use cases will be accommodated in



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

the data transmission model, preferably by providing sufficient network bandwidth for data transmission to these remote data processing centers, or relying on shipped storage (e.g., AWS Snowballs) in the most demanding scenarios.

Please consult the following materials for further information on this trade:

ngVLA Operations Concept	020.10.05.00.00-0002-PLA
Size of Computing Estimates for the ngVLA	ngVLA Computing Memo #4
The ngVLA Quantitative Exchange Model	nqxm Ver3.0
ngVLA Reference Design Development & Performance Estimates	ngVLA Memo #17

4.7.3 Observatory-Run vs Cloud Compute Resources

As discussed in Section 4.7.1, the computing resources required for the processing of high-level data products are on the scale of 60 PFLOP/sec. As of summer 2021, this would be the 7th largest super computer in the world.

ngVLA will not need this degree of computational capacity until it approaches full operations in 2035, as the computing needs will scale with the array construction and the deployment of observing modes. Capacity on this scale is not expected to be a commodity by 2035, but should present a much more tractable challenge than it would if deployed today.

As described in the Computing Sizing Memo, the computational load can vary by 1000x between use cases, so a scope contingency approach (i.e., deferment of key operating modes and the most challenging science products until later in operations) is also feasible while retaining most functionality, should trends in increasing capacity per dollar not continue through the design and construction phase.

The computational capacity cannot be provided by users on this scale, but the method of delivery for observatory-provided resources still permits a trade. In particular, the observatory could chose to expend capital for an observatory-run computing center, could collaborate with an existing advanced computing center (e.g., the Texas Advanced Computing Center, TACC), or could outsource the computing system to a cloud-based service.

The project position is that the progression of technology and the relative costs of these three options are not knowable 10 to 15 years in advance of full operations – there is too much uncertainty in the projections based on current capital and service costs. The project will develop a post processing software system that can run on a massively parallel computing architecture, and can defer the trade study of the hardware selection to a future date. This will provide flexibility regardless of market trends, and the project will endeavor to make a final decision on the compute system model prior to the system CDR/FDR.

Please consult the following materials for further information on this trade:

Size of Computing Estimates for the ngVLA	ngVLA Computing Memo #4
---	-------------------------



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

5 Baselined & Open Trades Summary

The following conceptual decisions are made as part of the system conceptual design technical baseline:

- 5.1. Heterogeneous vs Homogeneous Array: The system will employ a heterogeneous configuration with 18m main array and long baseline array elements, 6m short baseline array elements, and 18m total power elements.
- 5.2. Fixed Stations vs Reconfigurable Array: The system will employ fixed antenna stations (no reconfiguration capability).
- 5.3. Tropospheric Calibration Strategies: The antenna will be designed to support both fast-switching and water vapor radiometry for tropospheric calibration.
- 5.4. Instrumental Calibration: The system will implement a finite and repeatable set of system hardware configurations to support instrumental calibration by a post-processing pipeline.
- 5.5. Aperture Size: The main array and long baseline array will incorporate 18m apertures.
- 5.6. Field of View vs Sensitivity: The antennas will incorporate aggressively shaped optics to optimize sensitivity (G/T_{SYS}) with single pixel feeds.
- 5.7. Feed Indexer Geometry: The system will use a two-axis linear positioner for band selection and focus.
- 5.8. Multi-band vs Single-band Operation: The system will permit only single-band operation within a sub-array in order to optimize system sensitivity.
- 5.9. Altitude-Azimuth vs Equatorial Mounts: The antenna will include an altitude-azimuth mount with no field rotation stage.
- 5.10. Receiver Band Definition: The system will adopt a 6-band receiver definition, with wideband receivers below 12 GHz and waveguide bandwidth receivers above 12 GHz.
- 5.11. Linear vs Circular Polarization Basis: The system will adopt a linear polarization basis at all bands.
- 5.12. Bit Depth: The system will adopt a bit depth of 8 bits, with an ENOB of 7.2 bits, at all bands.
- 5.13. Downconversion Topology: The system will adopt an integrated (MMIC) approach to downconversion, with single-stage downconversion to baseband with in-phase and quadrature digitizer pairs for sideband separation. A total of 19 downconverters and 20 digitizers are incorporated in parallel to span the RF range of the system.
- 5.14. Cryogenics System: The system will adopt a variable speed GM system integrated in to two cryostats.
- 5.15. Antenna LO Architecture: The system will have 19 fixed frequency local oscillators, with one LO per downconverter. All will be locked to a common reference frequency distributed to the antenna.
- 5.16. Subarray Commensality and Independence: The system will incorporate a central signal processor architecture that shares processing resources amongst operating modes, providing an efficient use of resources but limiting the maximum achievable subarray independence.
- 5.17. Correlator-Beamformer Integration vs Independence: The baseline solution will employ an integrated correlator-beamformer that shares processing resources between modes.



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

- 5.18. FX vs XF Correlator Architecture: An FFX correlator architecture is adopted.
- 5.19. Phase-Delay vs True-Delay Beamformer: A true-delay beamformer will be provided for all phased array modes.
- 5.20. Real-Time Processing vs Post-Processing: The system will deploy an asynchronous post-processing system sized for total average throughput.
- 5.21. Low-Level vs High Level Data Products: The system will produce and record both low-level and high-level data products using observatory provided resources.

The following trade studies remain open as the project proceeds into the preliminary design phase:

- 5.22. WVR System: Parallel concepts are being advanced using the Band 4 receiver and a stand-alone receiver/dish. A trade between these two concepts will be assessed by the system PDR. A change from the baseline stand-alone design would require a change control board action.
- 5.23. Cryogenics System: An alternative concept using Stirling cycle pulse tubes will be advanced in parallel with the baseline design. This development project will be assessed at the system PDR. A change from the baseline GM design would require a change control board action.
- 5.24. Commercial Digitizers: An alternative concept employing Xilinx RFSoc devices in place of the Serial ADC ASIC is being explored as a way to reduce development effort and technical risk. A change from the baseline SADC design would require a change control board action.
- 5.25. Antenna LO System: Development towards a tunable synthesizer for Band 6 will be pursued, with the goal of eliminating 4 downconverter-digitizer and LO modules. Direct photonic delivery will also be developed as a risk mitigation for volume/mass constraints at the front end, and to reduce parts count for reliability. These development projects will be assessed at the system PDR. A change from the baseline design would require a change control board action.
- 5.26. Time & Frequency Reference Distribution: Multiple methods for distributing time and frequency references are under consideration. A decision will be made at an LO reference and timing subsystem CDR.
- 5.27. Correlator-Beamformer Integration vs Independence: An alternative design that employs a separate correlator and beamformer capabilities within the CSP will be advanced to a higher degree of technical readiness in parallel with the baseline design. A final decision between the two designs will be made at the CSP subsystem CDR.
- 5.28. ASIC vs FPGA vs GPU implementations: Associated decisions in technology implementation are tied to the correlator-beamformer integration vs independence trade. A decision between the various architectures will be made at the CSP subsystem CDR.
- 5.29. Observatory-Run vs Cloud Computing: TBD between the two options, or working with another FFRDC that specializes in advanced computing systems. A down-select will be made no later than the system FDR.



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

6 APPENDIX

6.1 Abbreviations and Acronyms

Acronym	Description
AD	Applicable Document
ALMA	Atacama Large Millimeter/submillimeter Array
AST	Division of Astronomical Sciences (NSF)
BW	Bandwidth
CDL	Central Development Laboratory
CDR	Conceptual Design Review
CSIRO	Commonwealth Scientific and Industrial Research Organization
CW	Continuous Wave (Sine wave of fixed frequency and amplitude)
EIRP	Effective Isotropic Radiated Power
EMC	Electro-Magnetic Compatibility
ENOB	Effective Number of Bits
FFRDC	Federally-Funded Research and Development Centers
FOV	Field of View
FRM	Focus Rotation Mechanism
FWHM	Full Width Half Max
HPC	High Performance Computing
HVAC	Heating, Ventilation, & Air Conditioning
IF	Intermediate Frequency
KPP	Key Performance Parameters
KSG	Key Science Goals
LEED	Leadership in Energy and Environmental Design
LNA	Low Noise Amplifier
LO	Local Oscillator
MREFC	Major Research Equipment and Facilities Construction (NSF)
MTBF	Mean Time Between Failure
MTTF	Mean Time To Failure
NES	Near Earth Sensing
ngVLA	Next Generation VLA
NRAO	National Radio Astronomy Observatory
NSF	National Science Foundation
OMT	Ortho-mode Transducer
PFB	Polyphase Filter Bank
PLL	Phase Locked Loop
PSD	Power Spectral Density
PWV	Precipitable Water Vapor
RD	Reference Document
RFI	Radio Frequency Interference
rms	Root Mean Square
RSS	Root of Sum of Squares
RTP	Round Trip Phase
SAC	Science Advisory Council
SEFD	System Equivalent Flux Density



Title: System Concept Options and Trade-Offs	Author: Selina et. al.	Date: 2022-08-05
NRAO Doc. #: 020.10.25.00.00-0005-REP		Version: B

SKA	Square Kilometer Array
SWG	Science Working Group
SNR	Signal to Noise Ratio
SRDP	Science Ready Data Products
TBC	To Be Confirmed
TBD	To Be Determined
VLA	Jansky Very Large Array
VLBI	Very Long Baseline Interferometry
WVR	Water Vapor Radiometer












020.10.25.00.00-0005-REP-SYS_TRADES

Final Audit Report


2022-08-09

Created:	2022-08-05
By:	Alicia Kuhn (akuhn@nrao.edu)
Status:	Signed
Transaction ID:	CBJCHBCAABAAY_GEqdOFU47sVbA1qMNdqzcWgYG5Dh0o

"020.10.25.00.00-0005-REP-SYS_TRADES" History

-  Document created by Alicia Kuhn (akuhn@nrao.edu)
2022-08-05 - 6:06:47 PM GMT
-  Document emailed to Thomas Kusel (tkusel@nrao.edu) for signature
2022-08-05 - 6:08:07 PM GMT
-  Email viewed by Thomas Kusel (tkusel@nrao.edu)
2022-08-08 - 9:59:59 PM GMT
-  Document e-signed by Thomas Kusel (tkusel@nrao.edu)
Signature Date: 2022-08-08 - 10:00:27 PM GMT - Time Source: server
-  Document emailed to rselina@nrao.edu for signature
2022-08-08 - 10:00:29 PM GMT
-  Email viewed by rselina@nrao.edu
2022-08-08 - 10:11:43 PM GMT
-  Signer rselina@nrao.edu entered name at signing as R. Selina
2022-08-08 - 10:13:37 PM GMT
-  Document e-signed by R. Selina (rselina@nrao.edu)
Signature Date: 2022-08-08 - 10:13:39 PM GMT - Time Source: server
-  Document emailed to emurphy@nrao.edu for signature
2022-08-08 - 10:13:41 PM GMT
-  Email viewed by emurphy@nrao.edu
2022-08-08 - 10:18:25 PM GMT
-  Signer emurphy@nrao.edu entered name at signing as E. J. Murphy
2022-08-08 - 10:19:31 PM GMT




 Document e-signed by E. J. Murphy (emurphy@nrao.edu)


Signature Date: 2022-08-08 - 10:19:32 PM GMT - Time Source: server

 Document emailed to westerhu@nrao.edu for signature


2022-08-08 - 10:19:35 PM GMT

 Email viewed by westerhu@nrao.edu

2022-08-09 - 8:29:50 AM GMT

 Signer westerhu@nrao.edu entered name at signing as Willem Esterhuysen

2022-08-09 - 8:30:19 AM GMT

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

Signature Date: 2022-08-09 - 8:30:21 AM GMT - Time Source: server

 Agreement completed.

2022-08-09 - 8:30:21 AM GMT

