



ngVLA Memo # 109

Updated RFI Impact Estimates & Influence on the System Design

R. Selina, U. Rau, C. De Pree (NRAO)
04/24/2023

Abstract

As the ngVLA project has progressed through the conceptual design phase the RFI environment has been evolving, along with our understanding of these novel RFI sources and their impact on the system design. In particular, Low Earth Orbit (LEO) satellite constellations present an emerging source of RFI for ground and space-based radio astronomy, vehicular radar systems are now ubiquitous, and mobile communications systems are increasingly moving up in frequency. We summarize the ngVLA's overall strategy for RFI mitigation, associated studies that have been performed, and the expected impact on the ngVLA system design, observing strategies, and observing capabilities.

I Introduction

A number of ngVLA memos have been published over the conceptual design phase informing an overall strategy to RFI mitigation. In addition, a number of tests have been performed with the VLA to assess the impacts of emerging sources of RFI. We consider the findings of these memos and note any changes in assumptions and expectations since their publication, and summarize updated input for the preliminary design of the instrument and the estimated impact of RFI in operations.

Recently, Low Earth Orbit (LEO) satellite constellations have risen to prominence in astronomy, primarily due to their impact on optical telescopes, but clearly they present a new source of radio frequency emission as well. In the context of the ngVLA design, we must consider the evolving nature of environmental RFI, mitigation and avoidance strategies in the design, and the possible residual impact on observing capabilities and system performance.

The following memos related to RFI mitigation in the ngVLA era have been published since the start of the conceptual design phase. Their findings are summarized and updated where appropriate in the following sections.

Memo No.	Title	Authors	Date
ngVLA (Main) #38	Subarray Processing for Projection-Based RFI Mitigation in Radio Astronomical Interferometers	M. C. Burnett, B. D. Jeffs, R. A. Black, K. F. Warnick	03/12/18
ngVLA (Main) #48	ngVLA Radio Frequency Interference Forecast	A. Erickson	07/01/18
ngVLA (Main) #70	RFI Mitigation for the ngVLA: A Cost-Benefit Analysis	U. Rau, R. Selina, A. Erickson	12/10/19
ngVLA (Main) #71	RFI Mitigation in the ngVLA System Architecture	R. Selina, U. Rau, R. Hiriart, A. Erickson	02/06/20
ngVLA Electronics #8	Headroom, Dynamic Range, and Quantization Considerations	R. Selina, O. Y. Ojeda	01/20/21
020.10.25.00.00-0004-MEM	Antenna Requirements for Mitigation of LEO Constellations	R. Selina	01/22/21
ngVLA Computing #3	Software Requirements for RFI Management	R. Amestica, R. Hiriart, P. Brandt	08/04/21
EVLA Memo #222	Coordinated Starlink User Terminal Testing Near the Very Large Array	C. De Pree, U. Rao, R. Selina, B. Svoboda	04/21/23
EVLA Memo #223	SpaceX-VLA Alamo Pilot Testing	C. De Pree, U. Rao, R. Selina, B. Svoboda	04/21/23

2 Assessment & Updates to Prior ngVLA Memos

Before considering any new projections we will first revisit some of the conclusions of the ngVLA memos in Table 1. These memos all reflect sensible conclusions at the time of publication but much has been learned since the first publications (five years ago) and some underlying assumptions have changed.

ngVLA Memo 48 considered technical frontiers that may change the RFI environment on ngVLA timescales, and indeed concluded that the proliferation of satellites in low earth orbit in the near to

intermediate term, along with the broad adoption of vehicular radar systems and deployment of 5G+ base stations, would be important environmental changes on ngVLA timescales. The memo also concluded correctly that the biggest impact will be to the short baselines of the ngVLA core, necessitating care in managing local RFI sources on the Plains of San Agustin. However, the impact these new sources of RFI could have on observations were not well quantified at this time, and subsequent studies (discussed below) have placed better bounds on the impact, which is in general lower than initially assumed even without any new mitigation strategies. In particular, LEO satellites are believed to have a significantly lower impact, and 5G network implementations to be less likely to use high frequency band allocations (which are better suited for short-distance links) near rural ngVLA antenna locations. The associated prevalence of RFI over the 6 GHz to 50 GHz frequency span is now believed to be less.

ngVLA Memo 70 summarized an algorithmic cost-benefit analysis in terms of recovered time or spectrum by various mitigation techniques. Some of these techniques could be implemented in data post-processing after averaging while others would need to take place at high time and frequency resolution and/or in real time, providing an important input to the system architectural design and the location of RFI mitigation services in the design. This cost-benefit analysis by necessity had to make assumptions about the nature of the future RFI signals (building off Memo 48), the fraction of the array that might be impacted, and the fraction of observing time that the signals might impact. Some of these assumptions have been significantly revised since, especially as they pertain to LEO satellites and 5G systems as described above. For example, it was previously assumed that LEO satellite signals would be received at full strength continuously with time. It is now clear that with only one satellite per provider illuminating a cell at a time, the geometry predicts signal durations of approximately one minute before a handover in the constellation to a satellite at different sky coordinates. The memo concludes that with the application of existing post-processing techniques after data averaging, 20-35% of visibilities might be flagged and discarded, and the application of higher time and frequency resolution flagging could reduce this impact to 10%. These results align well with empirical data at the VLA for L and S band (1-2 and 2-4 GHz). Accounting for the changes in assumptions on Memo 48 will have an associated flow-down impact here. In particular, revised assumptions about LEOs and 5G systems means that the data loss to flagging in the 6 GHz to 50 GHz range with either current post-processing or high time and frequency resolution flagging techniques may therefore be overestimated.

ngVLA Memo 71 built on the cost-benefit analysis of Memo 70 and proposed the inclusion of key blocks in the system architecture for RFI mitigation. Building on the foundations laid in Memo 48 and 70, it is this memo that is most deserving of an update here, as some key design decisions have changed and the conceptual design baseline of the system architecture deviates in key ways from Memo 71. We introduce first the subsequent studies and memos before returning to updates to the system architecture presented in this memo (see Section 6).

Internal ngVLA Project Memo 020.10.25.00.00-0004-MEM considers the impact of LEO constellations on the antenna and antenna electronics design in particular. This was the first deeper dive into the impact of LEOs and its conclusions are relatively straight-forward. The memo makes conservative assumptions that there will be 5 independent constellation operators with 60,000 total satellites, and no coordination with the Observatory. These numbers were intended to represent an upper bound on such systems, and the conclusions should be considered pessimistic/conservative. Using SpaceX and OneWeb satellite design

parameters as representative, the memo presents a single dish received power spectral density model and concludes that:

- The interference from LEOs is 20dB below the input damage threshold of the LNAs, even when a satellite transmitting in the ngVLA antenna's direction directly transits the antenna primary beam.
- Saturation of the receiver, requiring the flagging of the full receiver band, could occur as much as 1.4% of the time in this conservative scenario when transmitting LEOs (in any direction) are in the primary beam or first few sidelobes. This is an upper bound on saturation given the lack of knowledge about the beam shape of the LEO satellites.
- Total Power observations would likely be infeasible in the downlink RF windows (10.7-12.7 GHz), in addition to the projected loss of 1.4% of time across the receiver band due to system saturation.
- Interferometric observation impacts are TBD but may be manageable with 26% of data in the 10.7 to 12.7 GHz downlink band flagged when transmitting satellites transit the near-in (first 15) sidelobes (where satellite power may match or exceed $\sim 10\%$ of T_{SYS} before any interferometric attenuation). In the remaining 74% of time, the interfering signal strength after integration is estimated to be comparable to the typical attenuation achieved through interferometric fringe rotation, suggesting these off-axis sources may not impact interferometric imaging for all but the shortest baselines. The effect on imaging sensitivity would need to be quantified, especially for partially attenuated RFI signals.
- Phased array impacts are TBD, but at least 26% of data in the 10.7 to 12.7 GHz band would be flagged and discarded in phased array modes as well. Phased array observations may remain feasible given the natural attenuation of off-axis signals in phased array modes (which generally goes as $1/N_{ANT}$!) especially if operators avoid directly illuminating the antenna sites. Further study of the off-axis suppression in this mode would be required before any concrete assessment could be made.

In large part due to the surprising conclusion that interferometric imaging in the transmit (downlink) band may remain feasible, but lacking any empirical data to support this conclusion, the associated memo was published as a project document rather than a public memo. Subsequent tests with the VLA [7,8] largely corroborate its conclusions as described later in this memo.

ngVLA Electronics Memo #8 builds upon the conclusions of 020.10.25.00.00-0004-MEM, adds an analysis for vehicular radar, and proposes revised system-level requirements for input dynamic range, dynamic range across system setups, and quantization levels. The resulting input dynamic range requirements are comparable to typical LNA specifications, requiring up to 29dB of linear dynamic range over 1 GHz spans when accounting for present RFI, future projected RFI, and supported solar observing capabilities. The associated effective number of bits (ENOB) required in digitization is approximately 7.35 bits. The resulting

¹ This is due to the phase of the interfering signal having a random phase imposed by entering through the antenna sidelobes in addition to having a different propagation phase from the astronomical source at the center of the phased beam. This incoherency adds with each antenna in the phased array. See EVLA Memo 106 for more details.

receiver and digitizer specifications are high-performant but do not suggest significant new R&D will be required.

The conclusions of Electronic Memo #8 are expected to be robust and continue to inform the analog and digital electronics design. The conclusions of 020.10.25.00.00-0004-MEM, the associated impact of LEO satellite constellations in imaging, projections of the impact on observing, and the need for RFI mitigation services in the ngVLA system architecture can be refined based on recent tests with the VLA which we consider in the following sections. We first introduce a framework for estimating the interferometric attenuation to LEO satellites.

3 Interferometric Attenuation Simulations for LEO Satellites

A radio interferometer tracks a fixed position on the sky by continuously adjusting the expected relative phase between all antenna pairs (or baselines) to be zero for that phase tracking center. Signals from any source of emission located away from the phase tracking center incur a non-zero sinusoidal relative phase which carries information about the source position at any instant in time. If this source is located far enough away from the phase center, or is moving relative to the sky, the non-zero and time-varying relative phase results in an attenuation of the correlated signal due to the non-zero time and frequency range over which the signals are averaged, per measurement.

Prior published estimates of interferometric RFI attenuation through fringe winding (Perley, Thompson [9,11]) are tailored to the attenuation of geostationary sources. This is a very sensible assumption for nearby interference sources such as cell phone towers or radiating equipment in the array, but an extension to this framework is necessary when considering fast moving LEO satellites that are moving in arbitrary (but known) directions. Here we reuse the model constructed for the VLA coordinated tests [7] to produce estimates of the impact on the ngVLA.

3.1 Single Baseline

In order to estimate the interferometric attenuation we can expect for fast moving sources, we start by analyzing a single baseline interferometer case. Interferometric attenuation occurs as a result of time averaging as well as bandwidth decorrelation, and we estimate the effect of each separately.

The amount of attenuation due to time averaging, per baseline, is given by $R_t = \text{sinc}(\pi v_f t)$

where v_f is the natural fringe frequency and t is the integration time. Since the relative phase between two antennas changes by 2π when w changes by one wavelength, a fringe frequency may be calculated as $v_f = dw/dt$, where u, v, w are related to the antenna locations L_X, L_Y, L_Z as per Eqn 2-30 of Thompson 1999 [12].

Allowing for arbitrary satellite motion with $\omega_H = dH/dt$ and $\omega_\delta = d\delta/dt$ we get

$$v_f = \omega_\delta v_s - \omega_H u_s \cos\delta_s$$

A fringe frequency due to Earth rotation alone (Eqn 2-32 of Thompson 1999 [12]) is given by $v_f = \omega_e u_\phi \cos\delta_\phi$ by setting $\omega_H = \omega_E$, $\omega_\delta = 0$ in the above equation, and where ϕ represents the phase center.

For our analysis, we are interested in the relative phase difference and fringe frequency between the phase center direction (H_ϕ, δ_ϕ) and the satellite location (H_s, δ_s) . This is given by

$$v_f = \omega_\delta v_s - \omega_H u_s \cos \delta_s + \omega_e u_\phi \cos \delta_\phi$$

For a satellite in a polar orbit, $\omega_H = 0$. For a satellite in an equatorial orbit, $\omega_\delta = 0$, and for a geostationary source of RFI, $\omega_H = 0$ and $\omega_\delta = 0$. A celestial source moving with the phase tracking center may be represented by $\omega_H = \omega_e$ and $\omega_\delta = 0$.

The amount of attenuation due to bandwidth decorrelation, per baseline, is given by

$$R_v = \text{sinc}(\pi\beta\tau_d)$$

where β is the observing bandwidth and τ_d is the difference in delay between the phase center direction and the source of interference (Eqn 18 of Thompson 1982 [11]).

For a single baseline, the combined attenuation is given by the product of the two terms: $R = R_v * R_t$. In our model, we calculate the attenuation for each baseline in the array, for one timestep, the observing channel bandwidth, and for a satellite's location, speed and direction of motion during that timestep. UTs are treated as stationary objects with a relative fringe rate given by the Earth rotation speed ($0.004^\circ/\text{sec}$).

Satellite locations were chosen H and δ coordinates. We estimate satellite speeds for different altitudes as follows:

- At an altitude of 1150 km, the highest and slowest scenario, LEO satellites are completing a full revolution around the earth every ~6490 seconds (~108 min). Approximately 64.2 degrees of their orbit is above the local horizon, and they traverse that distance in 1157 sec (19.3 min), so their average angular rate of motion from the interferometer reference frame is $\omega_{sat} = 0.0027$ rad/sec ($0.155^\circ/\text{sec}$).
- The LEOs observed in the coordinated Starlink tests were in more common lower orbital shells at an altitude of 570km. These satellites are completing a full revolution around the earth every ~5755 seconds (~96 min). Approximately 46.76 degrees of their orbit is above the local horizon, and they traverse that distance in 748 sec (12.5 min), so their average angular rate of motion from the interferometer reference frame is $= 0.0042$ rad/sec ($0.24^\circ/\text{sec}$). Peak rate at Zenith can be approximately $0.76^\circ/\text{sec}$ and decreases towards the horizon, roughly consistent with the telemetry data.

3.2 Imaging Interferometer

The main emphasis of the modeling effort so far is on the single-baseline case to better understand the quantities of data flagged and to better connect the EVLA tests with a model that permits predictions for the ngVLA. However, we provide a first-order estimate of the impact on imaging by modifying the methodology used by Perley [9] to estimate the attenuation of a fixed interferer by an imaging interferometer using earth rotation synthesis. We will use the ngVLA SBA as a limiting case, given its short baselines and small number of apertures in the array.

Perley computes the interferometer attenuation factor for a fixed RFI source to be:

$$R(t) \sim \frac{1}{\sqrt{\pi}N_a} \sqrt{\frac{\lambda}{\omega_e t B_m \cos(\delta)}}$$

Where the distance travelled in the image by the geostationary interferer is $\omega_e t \cos(\delta)$. For a LEO satellite, which moves much faster than the phase center across the sky, we can approximate the distance traveled as $\omega_{sat} t$. We can then approximate the interferometric attenuation factor due to LEO satellite motion to be:

$$R(t) \sim \frac{1}{\sqrt{\pi}N_a} \sqrt{\frac{\lambda}{\omega_{sat} t B_m}}$$

For the SBA, the key parameters are $N_a=19$ and $B_m= 29\text{m}$. For $\omega_{sat}=0.25$ deg/sec and for $t = 1, 10, 100,$ and 1000 seconds, we can compute that R is approximately 9, 14, 29 and 34dB. For the ngVLA core, with $N_a=94$ and $B_m= 461\text{m}$, R(1000) is roughly 46dB. The equivalent computation for the VLA in D-config is 40dB, suggesting that the ngVLA core is no more sensitive to RFI than the VLA in its most compact configuration when used as an imaging array. Further tests and simulations would be required to confirm this point.

Including this attenuation factor when assessing the power spectral density levels calculated in 020.10.25.00.00-0004-MEM, this suggests that (1) all planned subarrays of the ngVLA, inclusive of the SBA, should be capable of imaging at least up to modest dynamic range over the LEO transmit bandwidth when satellite sidelobes are received through antenna sidelobes, and (2) the interferometric attenuation for the planned subarrays of the main array may be sufficient to image even in the case where the satellite is illuminating the ngVLA site but is received through an antenna sidelobe. These results should be considered 1st order approximations, but tests performed with the VLA in D-configuration should be considered representative of the expected impact on the performance of the ngVLA core subarray.

3.3 VLA Model / Estimates

Given the known geometry of the VLA configurations and available telemetry data from SpaceX, a model was built that estimates the visibility amplitude inclusive of interferometric attenuation of the RFI source (See VLA Memo 222). A conservative approach is taken here where (a) only the single baseline case is considered and (b) the amplitude is computed for a single integration ($t=1$) The resulting visibility amplitude is compared to the source flux and a fiducial residual imaging rms after averaging in the (u,v) plane. The goal with this more conservative analysis is to (a) be able to compare actual visibility amplitudes in a planned test and (b) identify impacts on actual imaging performance in subsequent tests.

This model is extended to the ngVLA in Section 5, but first we summarize the test results of the coordinated testing and ongoing imaging tests with the VLA as evidence of model validation.

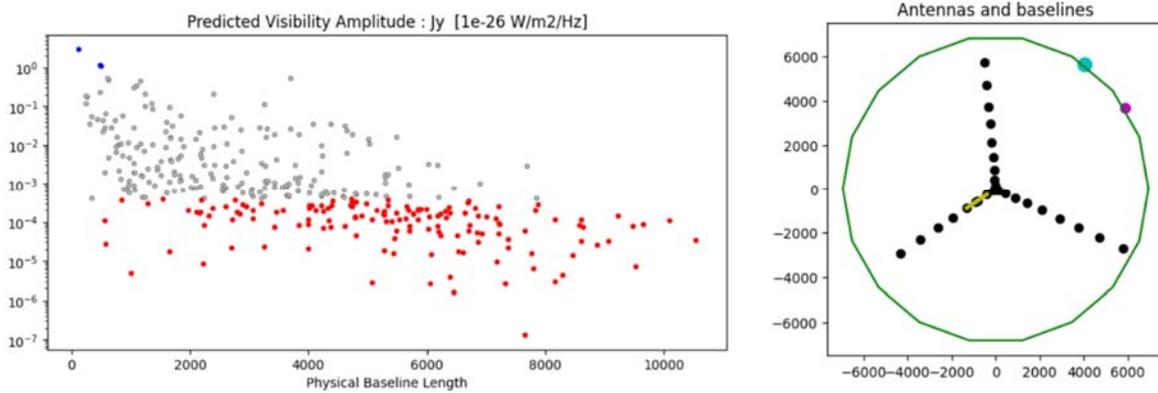


Figure 1 - (LEFT) Predicted visibility amplitude vs. baseline length. Blue: highest visibilities that correspond to the baselines marked in yellow on the right panel. Grey: visibilities that are below the source brightness (for this particular test) but still above the image noise level. Red: data points below a fiducial theoretical image rms level. (RIGHT) VLA array configuration, the azimuth locations of the phase tracking center (blue) and a satellite (magenta), and baselines for which the predicted satellite signal amplitude are strongest are shown in yellow.

4 VLA Testing Summary

The surprisingly positive conclusions of 020.10.25.00.00-0004-MEM, along with the recently deployed Starlink service in the vicinity of the VLA, provided a key opportunity to empirically validate (or refute) models of the impact of LEO satellites on interferometric imaging modes.

NRAO conducted two subsequent sets of tests that are relevant to the ngVLA design. The first were coordinated tests with SpaceX where telemetry data was provided to correlate observed results with constellation behavior and to intentionally illuminate the VLA site. [7] A second series of tests are passive with no coordination and instead aim to see the aggregated cumulative impact of constellation and ground station deployment in the vicinity of the array. [8]

The coordinated tests agree with the models in terms of large-scale behavioral expectations but fail to meet expectations at the detailed level. E.g., RFI amplitudes across the array and across time largely follow the distribution and pattern expected, but the specific baselines that are flagged do not consistently match predictions. This is likely due to approximations in the projected baseline and satellite vector computation with the coarse telemetry data, but this discrepancy will be further studied and debugged. These results confirm that even with LEO satellites illuminating the VLA site most baselines exhibit sufficient decorrelation to suppress the RFI level below the detrimental level, with only the shortest baselines impacted when the fringe rate drops due to the satellite trajectory. (EVLA Memo #222 [7])

The longer-term passive tests confirm that even with multiple downlinks operating within 30km of the site that the imaging performance is not significantly impacted in the downlink band for interferometric observations of at least 1 hr in length, with the LEO-generated RFI averaging down in a manner consistent with incoherent noise on these timescales. (EVLA Memo #223 [8]) For a significant fraction of baselines, the model predicts power levels below the single-visibility noise level but above the theoretical imaging sensitivity level, suggesting the presence of a noise floor when imaging in the presence of such RFI. Tests thus far, however, do not show any noise floor at least down to ~ 10 micro-Jy which is the level probed by the tests in EVLA Memo #223 [8].

5 Extrapolation of VLA Tests & Model to the ngVLA

When considering the applicability of the VLA/VLBA tests and the associated RFI attenuation model to the ngVLA, the following points are instructive:

- Detrimental emission thresholds for total power modes (e.g., ITU-R RA.769) are collecting area agnostic. Permissible emission levels are inversely proportional to T_{SYS} and receiver fractional bandwidth only. The detrimental emission levels are therefore comparable across all modern telescopes, as these parameters are largely maximized with present technology, and current telescopes are representative of future telescopes operating at the same frequencies.
- Interferometric modes provide a reduction in RFI sensitivity relative to the total power case, and single baseline sensitivity to RFI has an inverse proportionality to baseline length. This is due to the wider fringe spacing with shorter baselines (decreasing the fringe rate described in Section 3.1), and the fact that the interferometric attenuation is due to cycling through positive and negative lobes of the interferometric response function. ngVLA will have some baselines that are shorter than the VLA D-configuration in both the ngVLA Core and Short Baseline Array.
- Interferometric imaging RFI attenuation increases proportional to $1/N_a$ (See Section 3.2). This is based on the assumption that the phase of the measured RFI signal is randomized across baselines, compared to the desired astronomical signal. Importantly, this means that for interferometers with similar baseline distributions, more antennas leads to greater RFI attenuation. [9,10] A larger array like the ngVLA may be less sensitive to RFI than the VLA on this metric, though additional tests and/or simulations are required to prove this point.
- Interferometric attenuation of a coherent RFI signal within a single (u,v) -cell increases with time at a comparable or faster rate than the system noise drops through radiometric averaging. This is because of the cancellation of positive and negative response lobes of the sinc function across a given time window. The (u,v) -cell size limits the time duration across which the RFI signals can attenuate. However, since the time window across which response lobes are averaged varies across (u,v) -cells, we can also assume randomized amplitude and phase of the partially attenuated RFI across the gridded (u,v) -plane. With these assumptions of interferometric attenuation applying within each (u,v) -cell and a randomized result across (u,v) -cells, the detrimental emission levels for an RFI source may be considered time invariant across a full synthesis observation. [10]

When considered as a set, it is clear that VLA impact (or lack thereof) in testing should to first approximation be quite analogous to the expected impact on ngVLA performance. The sole outlying factor for interferometric modes is the baseline length distribution. The ngVLA will have shorter baselines than the VLA in both the Main Array Core and the Short Baseline Array, so these shorter baselines between neighboring antennas in the core may be more impacted, as would the baselines between antennas of the Short Baseline Array. Higher flagging rates may be expected on these baselines.

The other outstanding question to address is the impact on phased array modes. The beamforming process by definition introduces spatial filtering, and interfering signals will be received through sidelobes that will introduce random phase offsets that will further attenuate RFI, but a model and supporting test

must be constructed to assess the impact on KSG2 (pulsar observations near the galactic center) in particular.

5.1 Baseline Distribution Comparisons

To address the issue of baseline length distribution, a full set of plots are included in the appendix that compare the four main VLA configurations to subarray components of the ngVLA Main Array as well as the Short Baseline Array.

As is evident in Figures 7 and 8, the ngVLA appears to be minimally impacted on the scale of the spiral arms and mid-baselines, consistent with the VLA model and test results. Figure 9 shows individual baseline flagging within the core, and Figure 10 shows significant flagging within the SBA. We note that this is for the worst-case scenario of a satellite illuminating a cell at the ngVLA site, and is only applicable for the satellite transmit bandwidth. Coordinating with constellation operators to have them avoid cell illumination at radio astronomy sites will provide the most effective mitigation of this impact.

6 Updates to the ngVLA System Architecture

Given the prior studies and the evolution of the system design, we revisit the proposed RFI mitigation blocks in the ngVLA system architecture. ngVLA Memo #71 provides Figure 2 which proposes the inclusion of six blocks (A-F) in the architecture in addition to flagging in the data post-processing system. These blocks are:

- A - Outlier detection and masking or flagging on the time series voltage data in the antenna digital back end.
- B - Outlier detection and masking or flagging on the real-time spectrum (post F-engine in the correlator)
- C - Flagging and masking visibilities at high time/frequency resolution (post X-engine in the correlator, in the baseline/time/frequency space)
- D - Interference modeling and subtraction or nulling using matrix subspace projection, real-time imaging, or source location, in the correlator back end.
- E - RFI avoidance in the observation scheduler based on RFI metadata as a function of frequency and/or sky position.
- F - RFI database and manager to store RFI characteristics and metadata for reuse/retrieval. Analyze metadata to provide estimates of projected RFI for algorithm tuning in the current observation.

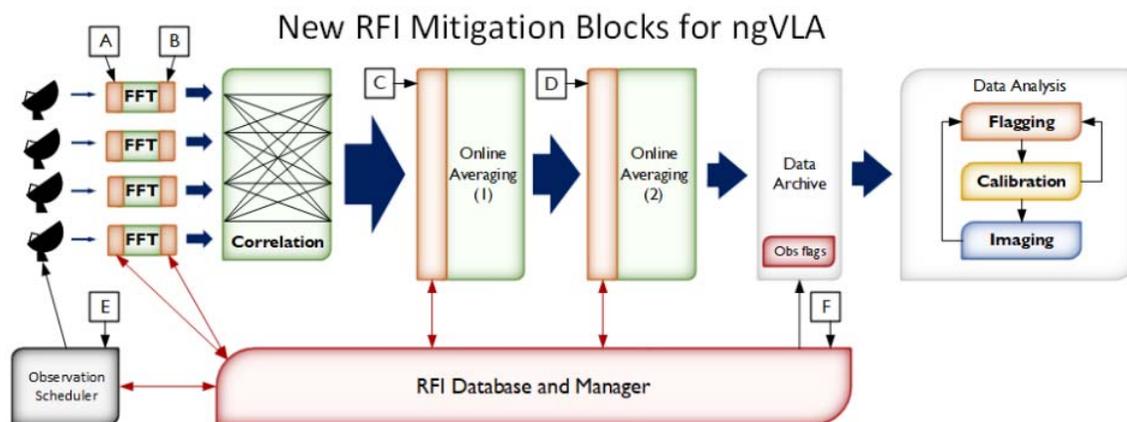


Figure 2 - RFI mitigation blocks in the system architecture as described in Memo 71.

The recommendation that these blocks exist in the architecture still holds. What has changed is the feature set in each that should receive priority in implementation as part of the design and construction effort given the reduced prevalence of RFI expected in the 6-50 GHz bands. We consider the implementation of each block in the following subsections.

6.1 Block A: Flagging & Excision in the DBE

RFI flagging of the time-voltage series data in the DBE remains a useful mechanism to generate online system flags for downstream processing especially as it pertains to broadband RFI that is wider than the generated sub-bands processed by the sub-band processors in the CSP. The value of masking or excision has been reassessed though given the associated risks to scientific performance and available flexibility in requantization at the sub-band level.

In the reference design technical baseline, the time-voltage data was requantized to 4-bits in the DBE before data transmission at all bands. The generation of the sub-bands at these lower bit-depths presents a risk as spectral dynamic range drops in the requantization step, concurrent with a need for an *increase* in spectral dynamic range due to the lower integrated noise over a sub-band (the ratio of narrowband RFI power to the sub-band power is higher than the ratio of the RFI power to the full integrated baseband power by 12.4dB).

To mitigate this, we previously proposed excising high power RFI before the requantization step for data transmission. We now instead favor permitting the requantization level to be tunable at the sub-band level, and to use higher bit-depths by default. In the conceptual design, the DBE can requantize to 2/4/6/8/12/16 bits, with the bit-depth tunable per ~200 MHz sub-band. All antennas in a subarray would use the same bit-depth by sub-band to avoid more complex Van Vleck corrections with mixed modes.

Given this flexibility and the projected CSP sub-band processor processing resources, 8-bit quantization will be the default for each sub-band, reducing the risk of RFI saturating subbands. Lower frequency bands (Receiver Bands 1-3) can alternatively use 16-bit requantization to increase the spectral dynamic range while respecting network bandwidth limitations, fully supporting the system-level dynamic range requirements.

In the case of long-baseline sites with tighter network bandwidth limitations, either less bandwidth can be processed or RFI free sub-bands can be requantized to lower bit depths, providing the necessary flexibility across sub-arrays without the need for algorithmic development for RFI excision in the deployed DBE.

6.2 Block B: Flagging & Excision in Real-Time Spectrum in the CSP

RFI flagging at the channel level in the CSP sub-band processors (Block B) remains part of the design. No excision step is required here as this can always be performed post-correlation (in Block C) with no loss of information. We retain this detection and flagging functionality in the design in order to record information about RFI detected during the observation (in the RFI DB) to inform flagging and excision at later stages of the signal processing chain, while descopeing any excision functionality in this functional block.

6.3 Block C: Flagging & Excision in Visibilities at High-Resolution in the CSP

Block C proposed flagging in the time/frequency/baseline space before the multiply-accumulate (MAC) step. It also included optional blanking of the associated visibilities and their omission in the MAC as an experimental capability, while carrying forward the weights for each integrated visibility as a metadata stream. Flags from preceding stages of the signal processing chain, could also be applied in here.

Flagging and excision at the Block C stage in the CSP can be omitted in the initial deployment, though it is desirable that the relevant DSP resources be provided in the sub-band processors so this capability could be developed later in operations, if desired or necessary.

The lower priority ascribed to Block C is in large part due to the lower perceived risk of RFI across the 6-50 GHz frequency spans (e.g., lower likelihood of high-frequency 5G+ systems deployed near remote antenna sites). A secondary concern is the feasibility of updating and maintaining a capability that is largely experimental in FPGA firmware within the CSP. For this reason, flagging and excision of the visibilities post-correlation may be a preferable step for the development of this capability, though it will entail more restrictive time and frequency resolution. This simplifies the design of the CSP while adding complexity in the CBE, but combined with changes in Block D leads to a system with the most desirable set of characteristics available.

A decision to develop higher time-frequency resolution averaging in the CSP can be deferred until environmental measurements of RFI are made with the first deployments of the ngVLA antennas, or an adverse change in the RFI environment is measured in the periodic VLA RFI sweeps.

6.4 Block D: Flagging & Excision in the CBE

Interference modeling and subtraction through advanced interference excision or nulling techniques such as matrix subspace projection, real-time imaging, or source location is looking increasingly less necessary as we better understand the natural RFI attenuation offered by interferometric modes for orbiting RFI as described in the preceding sections of this memo.

The other system development is the inclusion of an active crossbar switch at the correlator output, enabling the inclusion of commensal back-ends in parallel to the CBE. A commensal back-end is a more suitable place to consider the inclusion of these experimental/advanced techniques at this time, especially real-time imaging techniques that are closely tied to FRB science cases, since the capability has (a) a strong interrelationship with these imaging cases and (b) the data streams must be parallel to establish the impact of these advanced techniques compared to standard data processing.

While we propose that these experimental techniques be removed from the CBE's prescribed functionality, we still suggest that a mitigation block may still be prudent to include in the Correlator Back End. This could be a natural place to implement semi-real-time flagging and excision algorithms, before further averaging over time or frequency for data processing and archival.

The supported maximum data rate from the CSP to the CBE is 132 GB/sec. Data must be averaged down over time to 8 GB/sec across all observing modes to avoid bottlenecks downstream, but one can consider flagging and excising data in the CBE at the available time-frequency resolution afforded by this CSP-CBE link. This is inherently lower time and frequency resolution than what could be implemented in the CSP directly (flagging and excision before the MAC), but it importantly can be developed iteratively and the code is more maintainable over the life of the instrument than a firmware implementation in the CSP resources.

At the max link rate of 132 GB/sec this corresponds to 33.1 GVis/sec (at 4 Bytes per visibility) or approximately 279k channels at 1 sec time resolution, equivalent to 72kHz channels across 20 GHz of processed bandwidth when generating full stokes products. Time-frequency resolution can of course be traded within the data rate constraint. While this does not achieve the μ sec and msec-scale flagging that is feasible within the CSP, it does permit the recovery of spectrum.

This approach does entail that a second averaging step be included in the CBE prior to data formatting in a science data model (SDM) inclusive of the observational metadata. This functionality may have other uses, such as the implementation of baseline-dependent averaging. It also ensures that the CSP runs at the highest possible rate which increases the utility of the data stream for commensal back end systems accessing the CSP data over the active crossbar switch.

While the computational resources for sub-space projection techniques are not required, this does entail corresponding computational resources for data flagging and averaging, in addition to other identified CBE functions. The detailed sizing of the CBE will be considered in follow up design activities of the computing and software IPT.

6.5 Block E & F: Scheduler, RFI DB & RFI Manager

Avoidance is still a useful strategy for some types of RFI and this high-level functionality should be retained in the scheduler. Avoiding the sun, moon, and geostationary satellites are key functionality. However, the avoidance of LEO satellites is no longer a concern given the single dish analysis [5] and the VLA test results [7, 8]. Should LEO satellites transit the main beam or near-in sidelobes of the antenna, the associated data will simply be flagged and discarded. Anticipated data flagging rates should simply be included in the associated observation preparation tools when observing in the satellite downlink bands. The smaller set

of objects to be avoided will simplify scheduling in practice and may eliminate the need for regular ephemeris updates from external databases.

Retaining the interfaces from the ngVLA scheduler to external services seems increasingly prudent as the NRDZ concept continues to gain support. Spectrum coordination seems more plausible on ngVLA timescales, and Observational Status Sharing (OSS) systems are being actively developed. The RFI manager and database concept therefore remains unchanged, as recording the nature of interfering signals permits an a priori estimate of RFI present in the data stream of an observation that can be used to establish observation time and to tune offline RFI flagging algorithms.

7 Conclusions & Discussion

Based on a survey of the ngVLA memos on the topic of RFI mitigation, subsequent tests with the EVLA, and an evolving understanding of new spectrum-deploying technologies, we conclude that the risk of emerging RFI sources does not present a significant change from the present environment at the VLA and VLBA sites for interferometric operating modes. Satellite constellations in low earth orbit will impact total power observations but their impact on interferometric operating modes appears to not be as drastic as initially thought. The impact on phased array modes needs to be further studied.

Given the reduced impact associated with RFI risks, the more advanced RFI flagging and excision features incorporated into the system architecture can be scaled back on first deployment, though retaining them in the architecture provides mitigation options for changes in the operating environment over the life of the instrument and evolving science use cases.

Acknowledgements: Thank you to Matthew Schiller, Omar Ojeda, and Rafael Hiriart for fruitful discussions, feedback and corrections to the memo content and supporting calculations.

8 References

- [1] Erickson, A. “ngVLA Radio Frequency Interference Forecast” ngVLA Memo #28, July 2018.
- [2] Rau, U., et al. “RFI Mitigation for the ngVLA: A Cost-Benefit Analysis” ngVLA Memo #70, December 2019.
- [3] Selina, R., et al. “RFI Mitigation in the ngVLA System Architecture” ngVLA Memo #71, February 2020.
- [4] Selina, R., Ojeda O. Y. “Headroom, Dynamic Range, and Quantization Considerations” ngVLA Electronics Memo #8, January 2021.
- [5] Selina, R. “Antenna Requirements for Mitigation of LEO Constellations” ngVLA Project Internal Memorandum 020.10.25.00.00-0004-MEM.
- [6] Amestica, R., Hiriart R., Brandt, P. “Software Requirements for RFI Management” ngVLA Computing Memo #3, August 2021.

[7] De Pree, C. et al. “Coordinated Starlink User Terminal Testing Near the Very Large Array” EVLA Memo #222, April 2023.

[8] De Pree, C. et al. “SpaceX-VLA Alamo Pilot Testing” EVLA Memo #223, April 2023.

[9] Perley, R. “Attenuation of Radio Frequency Interference by Interferometric Fringe Rotation” EVLA Memo #49, November 2002.

[10] Perley, R. “RFI Emission Limits for Equipment at the EVLA Site” EVLA Memo #106, September 2006.

[11] Thompson, A. R. “The Response of a Radio-Astronomy Synthesis Array to Interfering Signals” IEEE Transactions on Antennas and Propagation, Vol. AP-30, No. 3, May 1982, pp. 450 - 456.

[12] Thompson, A. R., in Synthesis Imaging in Radio Astronomy II, ed. Taylor, G. B., Carilli, C. L., Perley, R. A. 1999, ASP Conference Series, 11-36, 1999

9 Appendix A – Attenuation by VLA Configuration

The following plots show the predicted visibility amplitudes and interferometric attenuation for a point in time in the telemetry data from the coordinated Starlink tests [7] with various VLA Configurations. The points in blue are the strongest (above the source brightness for our test) and correspond to the baselines marked in yellow on the lower left panel. The grey points are visibilities that are below the source brightness (for this particular test) but still above the image noise level. The red points represent the data points below a fiducial theoretical image rms level. The lower right plot shows the baseline distribution histogram with the same color coding.

Due to approximations in geometry in the current implementation of the attenuation modeling software, we believe that the aggregate behavior across baseline lengths is predicted accurately and corresponds to test results, but single baseline expectations have not yet been validated. The model will be further studied, debugged and validated.

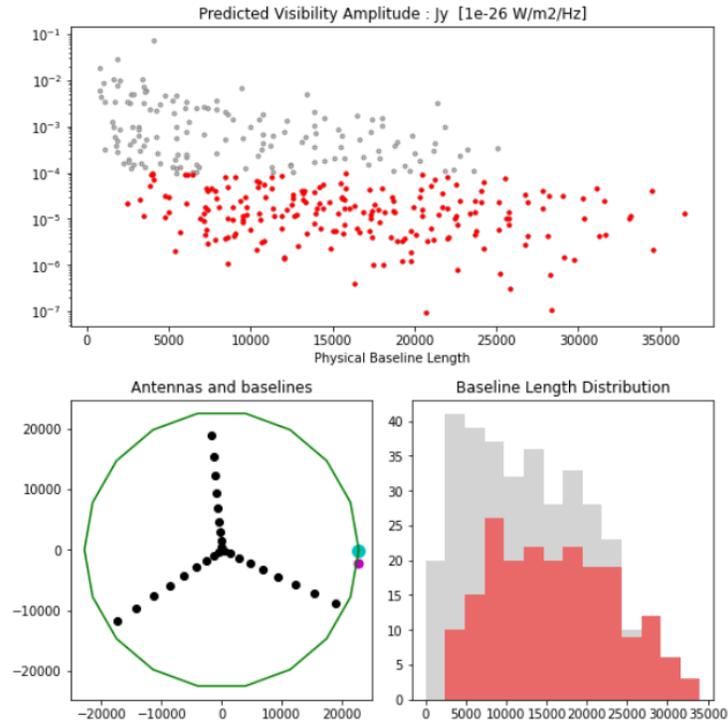


Figure 3 - Projected visibility amplitudes for a simulated LEO interference source observed with the VLA in A-Configuration.

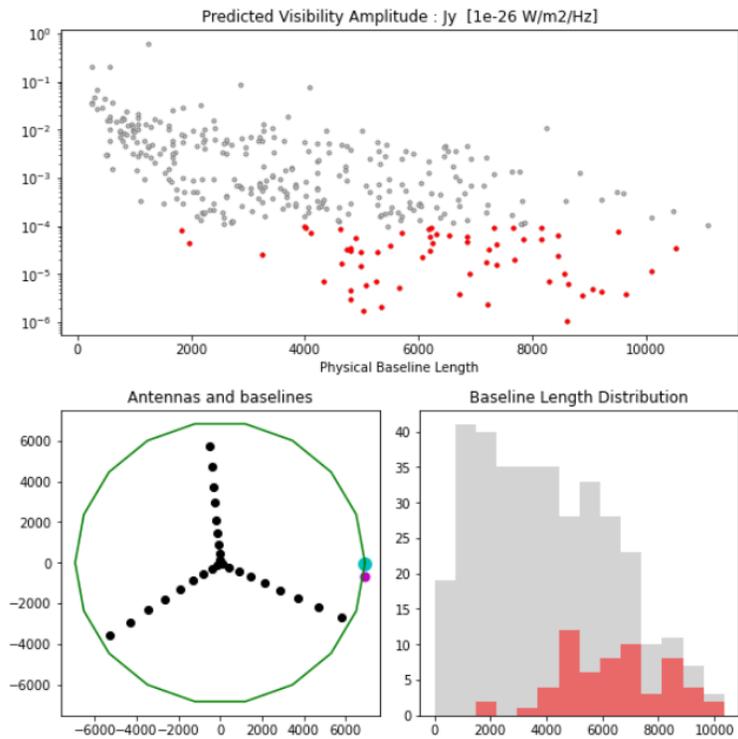


Figure 4 - Projected visibility amplitudes for a simulated LEO interference source observed with the VLA in B-Configuration.

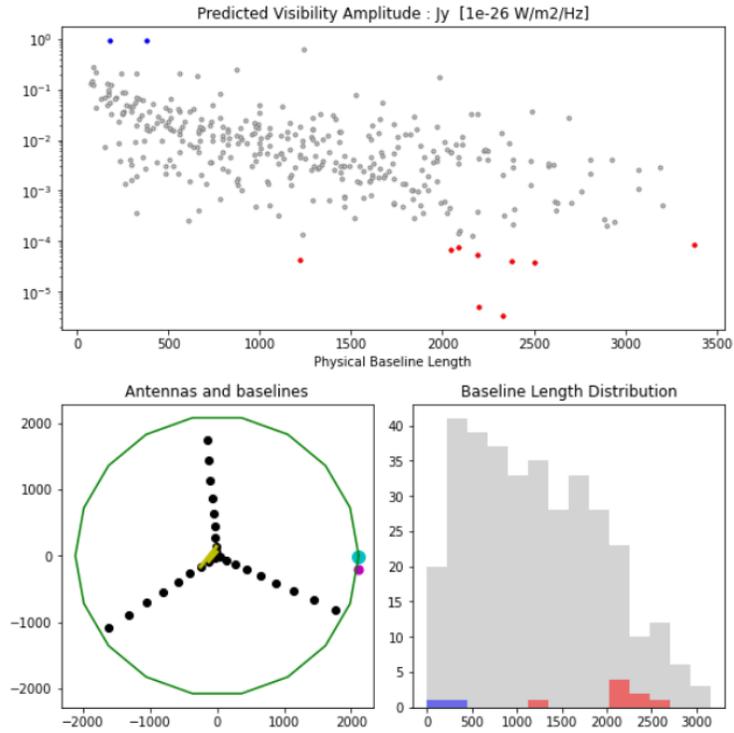


Figure 5 - Projected visibility amplitudes for a simulated LEO interference source observed with the VLA in C-Configuration.

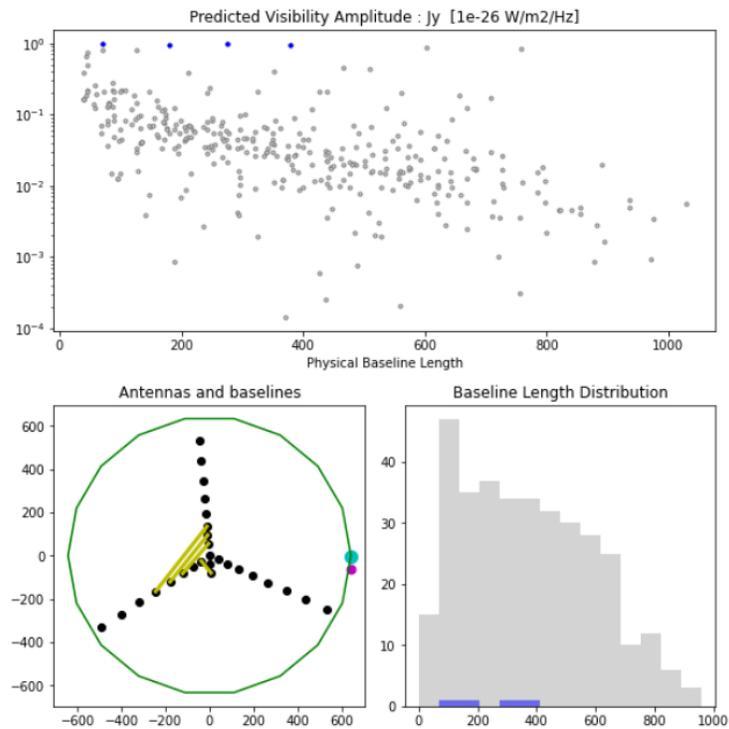


Figure 6 - Projected visibility amplitudes for a simulated LEO interference source observed with the VLA in D-Configuration

10 Appendix B – Attenuation by ngVLA Subarray

The following plots show the predicted visibility amplitudes and interferometric attenuation for a point in time in the telemetry data from the coordinated Starlink tests [7] for various ngVLA subarrays of the Main Array and the Short Baseline Array. The points in blue are the strongest (above the source brightness for our test) and correspond to the baselines marked in yellow on the lower left panel. The grey points are visibilities that are below the source brightness (for this particular test) but still above the image noise level. The red points represent the data points below a fiducial theoretical image rms level. The lower right plot shows the baseline distribution histogram with the same color coding.

Due to approximations in geometry in the current implementation of the attenuation modeling software, we believe that the aggregate behavior across baseline lengths is predicted accurately and corresponds to test results, but single baseline expectations have not yet been validated. The model will be further studied, debugged and validated.

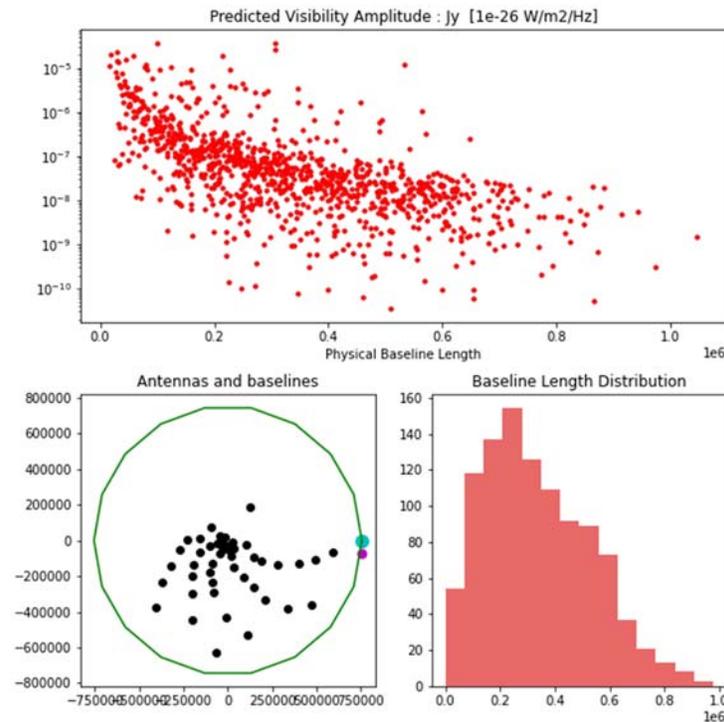


Figure 7 - Projected visibility amplitudes for a simulated LEO interference source observed with the mid-baselines of the ngVLA Main Array

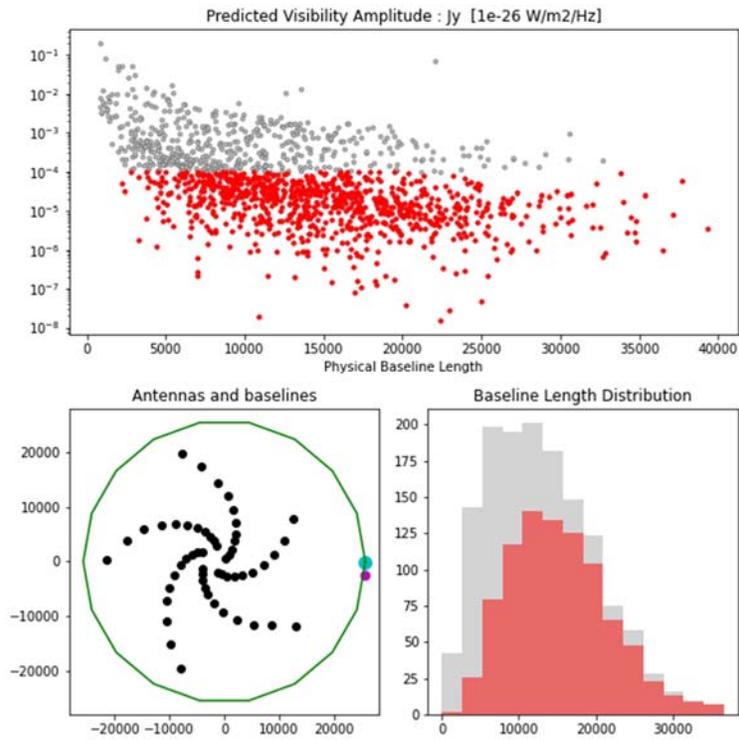


Figure 8 - Projected visibility amplitudes for a simulated LEO interference source observed with the spiral arms of the ngVLA Main Array.

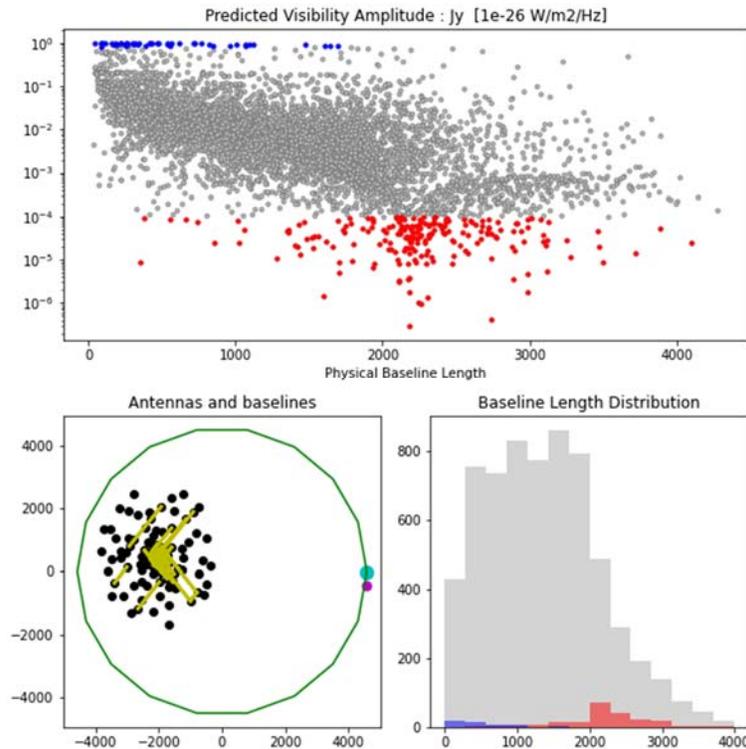


Figure 9 - Projected visibility amplitudes for a simulated LEO interference source observed with the ngVLA Core.

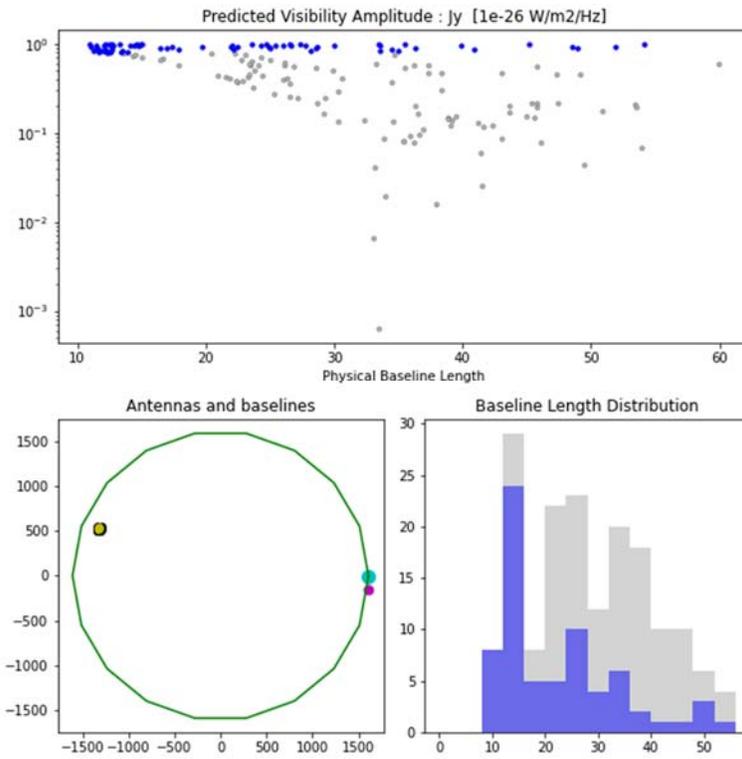


Figure 10 - Projected visibility amplitudes for a simulated LEO interference source observed with the ngVLA Short Baseline Array.