Pulsar Timing Array Requirements for the ngVLA

Next Generation VLA Memo 42

NANOGrav Collaboration

(Dated: April 5, 2018; Version 1.0)

1. SCIENCE WITH PULSAR TIMING ARRAYS

The recent detections of binary black hole and binary neutron star mergers by the LIGO-VIRGO collaboration provide spectacular confirmation of the existence of gravitational waves (GWs) at kilohertz frequencies, and challenges our concepts of stellar evolution and binary interactions. Our current understanding of the formation of galaxies and the history of mass assembly in the early universe requires the mergers of supermassive binary black holes, which produce GWs at much lower frequencies (in the microhertz to nanohertz range). GWs at such low frequencies (i.e., with wavelengths on the order of light-years) can be detected by a "pulsar timing array" (PTA) where a collection of stable pulsars are timed over a period of years to decades. Low frequency GWs would produce deviations in pulsar timing residuals that are *spatially correlated* in a quadrupolar pattern (e.g., Estabrook & Wahlquist 1975; Hellings & Downs 1983), enabling the identification of the GW signal against a background of confusing effects including pulsar timing noise, pulse propagation effects in the interstellar medium, and uncertainties in the position of the solar system barycenter.

The North American Nanohertz Observatory for Gravitational Waves (NANOGrav) is a collaboration of astronomers using the world's most sensitive pulsar timing facilities (Arecibo Observatory and the Green Bank Telescope) to construct a Galactic-scale detector for low-frequency GWs. We have recently released 11 years of pulsar timing data (Arzoumanian et al. 2018, see Figure 1 below), which place stringent limits on the stochastic background of GWs from supermassive binary black hole mergers, as well as meaningful constraints on various other (more speculative) source classes such as cosmic strings and phase transitions in the early universe that are expected to produce GWs in the low frequency band.

At present NANOGrav both competes and collaborates with its international partners, the European PTA and the Parkes PTA. As new telescopes (such as FAST and CHIME) come on line, we expect that collaboration under the umbrella of the International Pulsar Timing Array (IPTA) will play a more significant role. The leadership role of NANOGrav depends on continued access to telescope facilities for pulsar timing. In the next decade, we expect that the stochastic GW background will have been successfully detected, either by NANOGrav or as an IPTA effort. The scientific focus will shift to the measurement of the anisotropy of the stochastic GW

Corresponding author: Shami Chatterjee
shami.chatterjee@nanograv.org

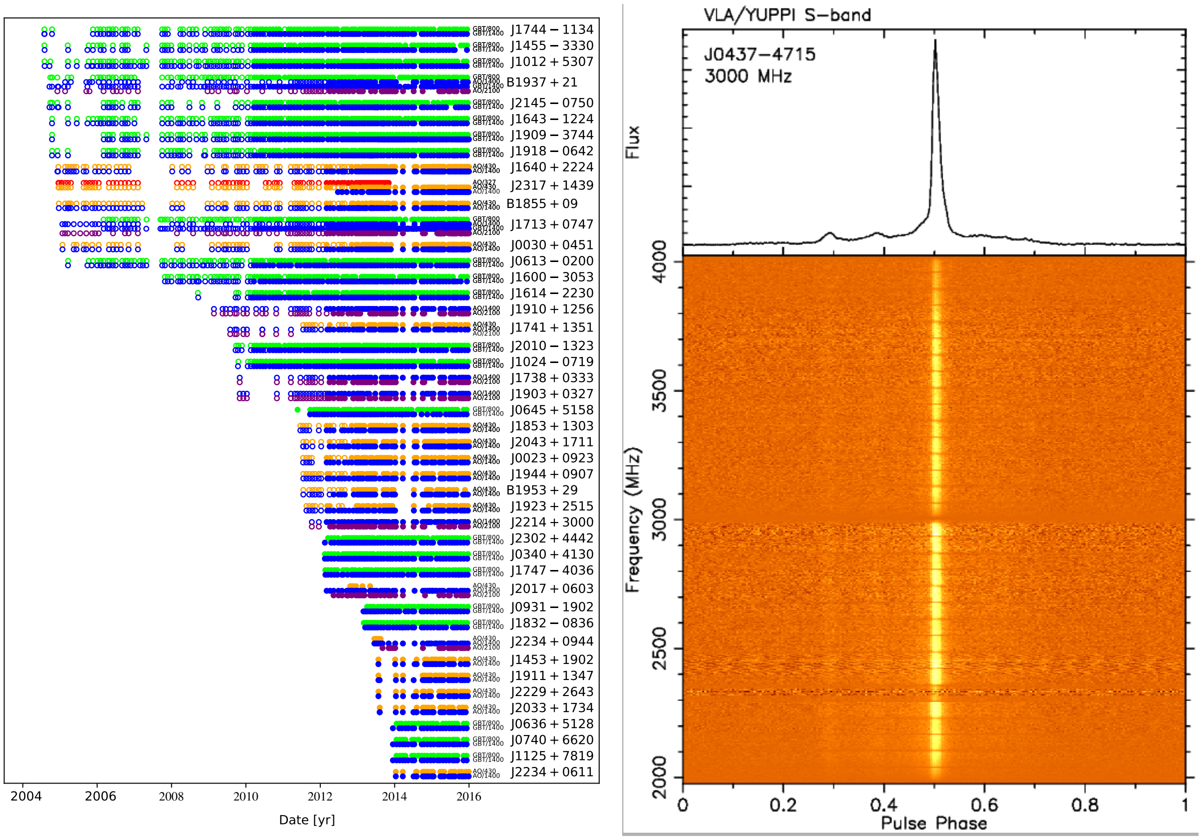

Figure 1. Left: The NANOGrav 11-year data release (Arzoumanian et al. 2018) includes pulse time-of-arrival measurements for 45 millisecond pulsars at Arecibo Observatory and the Green Bank Telescope, shown here sorted by the earliest included epoch. Colors indicate radio frequency band at either telescope: 327 MHz (red), 430 MHz (orange), 820 MHz (green), 1.4 GHz (blue), and 2.1 GHz (purple). Sensitivity to gravitational waves increases with steady growth in the number of pulsars being timed, and the expansion of the array demonstrated here exceeds our past projections.
Right: The pulse profile of PSR J0437–4715 observed over 2–4 GHz at the VLA, illustrating the use of phased arrays for pulsar timing and the utility of wider bandwidths in generating higher precision pulse times of arrival.

background, the detection of continuous waves and bursts with memory produced by individual GW sources, joint observations of their electromagnetic counterparts, as well as constraints on the predictions from more exotic physics.

In this memo we summarize the requirements and operational constraints such that the future ngVLA can play a leading role in PTA science. The technical and computational requirements are well within the current scope of the ngVLA project, and the operational requirements are not unreasonable.

2. OBSERVATIONAL PROGRAM

With its high sensitivity and broad spectral coverage, the ngVLA will be well-suited to pulsar timing observations. While the imaging capability of an array is not required for such observations, the capability of flexible sub-arrays provides an *advantage* for pulsar timing compared to large single dishes. Specifically: the average pulse profile of a pulsar over many rotations is very stable, but the individual single pulses in that average are subject to pulse-to-pulse variations, shot noise, and jitter. Pulse jitter sets a floor on the minimum number of pulses to be averaged in order to obtain an accurate pulse time of arrival at a given epoch, and thus, the required integration time does not continue to drop as available sensitivity increases. The current NANOGrav timing program (Arzoumanian et al. 2018) uses integration times of 0.25–0.5 hr per pulsar at Arecibo and Green Bank with cadences of weekly (for the best pulsars) to bi-weekly, and allocating similar integration times with a selected fraction of the ngVLA collecting area offers more flexibility and higher efficiency compared to steering the entire collecting area of a large single dish to a sequence of pulsar timing targets.

Thus we envisage two possible modes of PTA observation on the ngVLA: (1) A sole-user mode, where the array is split into ∼5 sub-arrays, each observing a pulsar of interest; or (2) A shared-user mode, where ∼20% of the ngVLA (i.e., comparable to the current GBT), or more, depending on the pulsar flux density, is phased up to observe one pulsar. As an aside, we note that sub-arraying reduces the net computational load, since baselines do not have to correlated across sub-arrays. Thus, as long as the correlator resources can be flexibly re-deployed, no extra computation resources will be required for shared-user or multi-target phased sub-array operation.

Frequency and Bandwidth: Pulsars are typically brighter at lower frequencies ($S_\nu \propto \nu^\alpha$; $\alpha \sim -1.6$) but receiver noise and sky background considerations typically lead to frequencies around 1 GHz for pulsar timing. The current NANOGrav program uses observations at both a low (typically $< 1\,\mathrm{GHz}$) and a high ($> 1\,\mathrm{GHz}$) frequency at Arecibo and GBT, as shown in Figure 1, in order to fit for the pulse dispersion measure (DM), which corresponds to the integrated line-of-sight electron column density ($\mathrm{DM} \equiv \int_0^s n_e\,ds$). For the very high precision timing required for PTAs, the time variability of the line of sight electron density requires contemporaneous measurement of the DM at each observation epoch. With wide enough available bandwidths, dual-frequency observations will not be required for DM estimation, leading to a large gain in efficiency. (See, e.g., Figure 1, right panel, for an illustration of VLA phased-array pulsar timing spanning 2–4 GHz.)

The optimal frequency for pulsar timing observations has been investigated in detail by Lam et al. (2018), leading to the broad conclusion that the availability of broader bandwidths results in a higher center frequency being preferred for NANOGrav timing observations. As shown in Figure 2, higher frequencies and larger bandwidths can produce higher precision time-of-arrival (TOA) measurements than currently achieved. The effect is especially pronounced for pulsars at higher DMs, where propagation effects are increasingly important (Figure 2, right panel). As the NANOGrav PTA expands, newer pulsar discoveries are more likely to be fainter and more distant, so higher sensitivities and larger bandwidths become more important over time. However, access to frequencies down to 1–1.5 GHz remains essential.

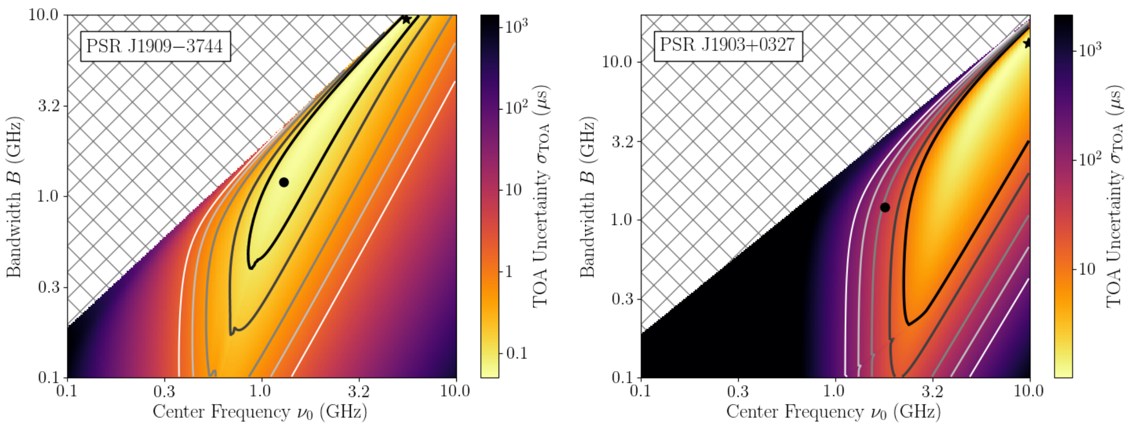

Figure 2. The uncertainty in pulse time of arrival measurements as a function of observing center frequency and bandwidth for two millisecond pulsars currently being observed by NANOGrav, adapted from our work in Lam et al. (2018). **Left:** PSR J1909–3744, which is one of the best-timed pulsars. Contours indicate TOA uncertainties of 2, 1, 0.5, 0.2, and 0.1 μs, in order of increasing darkness and thickness. The minimum TOA uncertainty (black star) is $\sigma_{\rm TOA}(\nu_0 = 5.5\ {\rm GHz}, B = 9.5\ {\rm GHz}) = 50$ ns and the estimate given the current frequency coverage (black circle) is $\sigma_{\rm TOA}(\nu_0 = 1.3\ {\rm GHz}, B = 1.2\ {\rm GHz}) = 60$ ns. **Right:** PSR J1903+0327, which is the pulsar with the highest dispersion measure (297.52 pc cm^{-3}) currently observed by NANOGrav. Contours indicate TOA uncertainties of 200, 100, 50, 20, and 10 μs, in order of increasing darkness and thickness. The minimum TOA uncertainty (black star) is $\sigma_{\rm TOA}(\nu_0 = 8.1\ {\rm GHz}, B = 10\ {\rm GHz}) = 1.0\ \mu$s (with 10 GHz being the limit of the plot, a fuller analysis shows the minimum at $\nu_0 = 9.8$ GHz, $B = 13.2$ GHz with a nearly similar amplitude) and the estimate given the current frequency coverage (black circle) is $\sigma_{\rm TOA}(\nu_0 = 1.8\ {\rm GHz}, B = 1.2\ {\rm GHz}) = 44\ \mu$s, with variable scattering being the largest unknown component. Higher observing frequencies with larger bandwidths potentially allow significant improvements in timing precision.

Observing cadence: NANOGrav currently observes ∼50 pulsars, with each pulsar in the array being observed (approximately) every two weeks. Beyond the detection of the stochastic GW background, the future PTA scientific program requires both an increasing number of pulsars in the array and an improved observing cadence.

A straw-man future observing program might involve ∼200 pulsars distributed over the sky. If each pulsar can be timed with a 20% sub-array of the ngVLA for ∼0.5 hr every week, the total NANOGrav observing program may require ∼20 hr/week of the full array. However, such estimates require two important caveats. (1) The PTA observing program needs to be sustained for years—see, e.g., our recent data release spanning 11 years (Arzoumanian et al. 2018). (2) As the PTA expands, future pulsar additions to the array may be disproportionately fainter, requiring larger integration times. Instead, it appears more likely that NANOGrav will rely on the ngVLA to time the most critical and faintest pulsars, with the remainder being timed at partner facilities.

Pulsar Search: The PTA requirements for the ngVLA described here do *not* include the capability of blind, full field-of-view pulsar searches, which impose enormous computational loads. However, sensitivity to GWs increases with the number of pulsars, and requires a distribution of sources all over the sky for a good GW detector response function. It is likely that the discovery of new pulsars will be dominated by single dish telescopes and by hybrid methods, where multi-wavelength imaging is used to identify candidate compact sources that are then searched for pulsations. (A pulsar that is suitable for timing as part of a PTA will be easily identified as a compact source in radio sky survey images, or potentially via high-energy counterparts.)

3. TECHNICAL REQUIREMENTS FOR PTA OBSERVATIONS

Given the more general programmatic issues described above, here we summarize the technical requirements for the ngVLA in order to enable PTA observations. We note that at present, none of these requirements represent a "tall pole" or driver for the ngVLA specifications.

Phased Sub-Arrays: Multiple independent sub-arrays are required, up to $\sim$10, with each sub-array providing an independent phased array beam. It is preferable if each sub-array can maintain phase coherence over the duration of a pulsar timing observation, possibly using real-time self-calibration strategies with in-field calibrator sources.

Frequency coverage and Bandwidth: As described above and investigated further in Lam et al. (2018), wide-band frequency coverage is required and coverage down to 1–1.5 GHz is essential. We note that some pulsars with large and time-variable DMs may be suited to simultaneous dual-band observations using two independent sub-arrays, a capability of the ngVLA that is simply not available at any single-dish facility.

Correlator and Computation needs: Correlator specifications for pulsar observations are described in another ngVLA memo by Demorest et al., but in broad outline, phased array beams that can be sampled at 50 μs with 0.5 MHz channels are sufficient for PTA requirements. The output of the phased array beam will be coherently de-dispersed (i.e., a digital filter will be applied to remove the known average pulse dispersion as a function of frequency) in real time, before sampling. Given that cross-correlation for imaging will not be routinely required (and in any case, dishes observing different fields will never be correlated against each other), PTA observations using sub-arrays will be far less computationally challenging than full field-of-view, full-resolution imaging observations.

Polarization Calibration: Emission from millisecond pulsars is typically polarized, with linear polarizations at the few–50% level (e.g., Yan et al. 2011). The phased array beams will thus require polarization calibration (and more importantly, polarization stability) at the 5–10% level, preferably better.

Clock Stability: The long-term clock stability requirement for pulsar timing observations is currently met by tying observatory masers to GPS time. On the short term, clock stability requirements for pulsar timing are exceeded by the requirements for high-frequency VLBI.

Data Management and Curation: PTA observations are a long-term enterprise, and the scientific value of the data set increases with time baseline as wider ranges of GW frequencies are probed to higher sensitivities. NRAO has an exemplary track record of incorporating infrastructure for data management over decade-long timescales at the VLA, but pulsar observations (e.g. at the GBT) have typically not been included in these plans. For the ngVLA, the volume of PTA observations is expected to be dwarfed by full-field visibility storage requirements, and a long term archival and curation plan is both essential and straightforward.

4. SUMMARY

Pulsar timing arrays have the potential to open a new window on the gravitational wave spectrum, probing nanohertz emission from the supermassive binary black hole mergers that accompany mass assembly in the universe, as well as other, more exotic, sources. The ngVLA will play a key role in such PTA observations, as long as certain key requirements are met: most importantly, independently phased sub-arrays and frequency coverage down to 1–1.5 GHz. We have outlined the other requirements and constraints in this memo, and find that none of them pose a significant obstacle given the existing specifications of the ngVLA.

The NANOGrav Physics Frontiers Center is supported by the National Science Foundation award number 1430284.

REFERENCES

Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2018, arXiv:1801.01837

Estabrook, F. B., & Wahlquist, H. D. 1975, General Relativity and Gravitation, 6, 439

Hellings, R. W., & Downs, G. S. 1983, ApJL, 265, L39

Lam, M. T., McLaughlin, M. A., Cordes, J. M., Chatterjee, S., & Lazio, T. J. W. 2018, arXiv:1710.02272

Yan, W. M., Manchester, R. N., van Straten, W., et al. 2011, MNRAS, 414, 2087