

Pulsar Timing Arrays

Ryan S. Lynch

Department of Physics and Astronomy, West Virginia University, Morgantown, WV,
26506-6315, USA

Abstract.

Pulsars, and especially millisecond pulsars, can be used as precision clocks to explore a wide range of phenomena. An array of millisecond pulsars distributed across the sky can be used to detect low-frequency gravitational waves by searching for a correlated signature between pulsars, arising from the influence of gravitational waves at the Earth. These pulsar timing arrays complement other methods of gravitational wave detection in their frequency coverage and the physics they probe. I will introduce the basics of pulsar timing and discuss how it applies to the detection and study of the gravitational wave Universe. Pulsar timing array collaborations, current results, and future prospects will also be presented.

1. Pulsar Basics

A pulsar is a highly magnetized, rapidly rotating neutron star that emits beamed radio and high energy (X-ray and/or gamma-ray) radiation. A pulse of radiation is seen each time the beam crosses the line of sight of an observer, giving rise to a *pulse profile* that encodes information about the structure of the magnetosphere and the emission process (see Figure 1). Empirically, though there is substantial pulse-to-pulse variability, an average pulse profile emerges after ~ 100 – 1000 rotations. This pulse profile, and the rotation of the pulsar itself, are observed to be very stable over long timescales (though see [1] for some caveats). These two factors allow pulsars to be used as precision astronomical clocks, suitable for studying a wide range of phenomena in exquisite detail.

Pulsars are born with periods $P \sim 10$ s ms but quickly spin down due to the loss of rotational energy, to $P \sim 0.1$ – 1 s. Radio emission will continue for ~ 10 – 100 Myr, after which the pulsar will become a quiet neutron star. However, if the pulsar exists in a binary system, it may find new life as a millisecond pulsar (MSP). In this scenario, mass and angular momentum are accreted from a companion star, spinning the pulsar back up to periods $P < 10$ ms while simultaneously decreasing the magnetic field to $B \sim 10^8$ – 10^9 G. Due to their lower magnetic fields, MSPs spin down much slower than non-recycled pulsars, having lifetimes greater than 10 billion years. To date, ~ 260 MSPs have been discovered¹. Figure 1 shows the known pulsars as a function of period and spin-down.

2. Pulsar Timing

The extremely high rotational stability of pulsars, and especially MSPs, make them precision clocks, suitable for measuring minute deviations in the expected *time of arrival* (TOA) of a pulse. This is accomplished through phase coherent *pulsar timing*, which unambiguously accounts for

¹ For an up-to-date list, see the ATNF pulsar catalog (<http://www.atnf.csiro.au/people/pulsar/psrcat/>)






Figure 1. *Left:* A typical pulse profile, in this case of PSR J0741+6618, an MSP discovered in the Green Bank North Celestial Cap survey [2]. *Right:* The P - \dot{P} diagram, showing pulsar spin period vs. spin-down. Circles indicate binary systems, and lines of constant magnetic field and age are indicated by dotted and dashed-dotted lines, respectively. Two distinct populations are clearly visible: the non-recycled pulsars, with $P \sim 0.1\text{--}4\text{ s}$ and $\dot{P} \sim 10^{-17}\text{--}10^{-13}\text{ s s}^{-1}$, and the fully recycled millisecond pulsars with $P \sim 3\text{ ms}$ and $\dot{P} \sim 10^{-20}\text{ s s}^{-1}$, with partially recycled pulsars bridging the two populations. Data taken from the ATNF pulsar catalog.

every rotation of the pulsar over time-spans much greater than the star's rotational period. That is, the rotational phase of the pulsar at any time t (measured in an inertial reference frame with respect to the pulsar) is given by

$$\phi(t) = \phi(t_0) + f(t - t_0) + 1/2\dot{f}(t - t_0)^2 + \dots \quad (1)$$

where t_0 is a reference epoch, $f = d\phi/dt$ is the rotational frequency, and \dot{f} is the first time derivative of f . Higher order derivatives can sometimes be measured (particularly in young pulsars), but usually they are either too small (as is the case for most MSPs), or are contaminated by timing noise (see §3).

The first step in pulsar timing is to transform a TOA measured at the telescope to an approximately inertial reference frame with respect to the pulsar, typically the Solar system barycenter. Then,

$$t = t_{\text{topo}} - t_{0,\text{topo}} + \Delta_{\text{clock}} - \Delta_{\text{DM}} + \Delta_{\text{R}\odot} + \Delta_{\text{E}\odot} + \Delta_{\text{S}\odot} + \Delta_{\text{R}} + \Delta_{\text{E}} + \Delta_{\text{S}}. \quad (2)$$

Here, topo indicates time measured in the topocentric frame of the observatory, and the Δ 's indicate various corrections and delays: Δ_{clock} accounts for differences between the observatory clock and a terrestrial time standard, Δ_{DM} corrects for a radio frequency-dependent dispersive delay as the signal travels through the ionized interstellar medium (ISM), Δ_{R} is the classical Römer delay that accounts for light travel time across an orbit, Δ_{E} is the Einstein delay, which takes into account time dilation due to orbital motion and the gravity of other bodies, and Δ_{S} is the Shapiro delay. These orbital delays arise from the orbit of the Earth and the influence of the Sun and other bodies in the Solar system (indicated by \odot) and from the orbit of the pulsar and its companion when found in a binary system.

Two points are worth emphasizing. First, the various delays encode information that reveals a wealth of science, such as information about the pulsar orbit. Second, many of these corrections

are time-dependent. The ISM is not static, so *dispersion measure* (DM, the column density of free electrons) must be measured at different epochs; clock corrections must be regularly updated; changes in the orbit of the pulsar owing to general relativistic effects or the presence of other bodies (planets or a third stellar companion) will manifest in the pulsar orbital terms.

The rotational, astrometric, and possible orbital parameters of a pulsar comprise a *timing model* that, along with the other corrections in Equation 2, predict the rotational phase of the pulsar at any given time as observed from the Solar system barycenter. The observed TOA of a pulse, or more precisely, a fiducial point on the pulse profile, can then be compared to this prediction. The difference between the observed and predicted arrival time is known as a *timing residual*. If we assume the fiducial point on the pulse profile always corresponds to the same rotational phase (i.e., that the profile does not drift or change), then an integer number of rotations must fit between any two TOAs. Model parameters are iteratively adjusted under this constraint to minimize the timing residuals. So long as the uncertainties in the model parameters do not propagate to a phase uncertainty $|\sigma_\phi| > 1$, phase coherence can be maintained, even though only a fraction of the total rotations of the pulsar are actually observed. This allows for extremely precise measurement of the model parameters, illustrating the power of pulsar timing.

3. The Limits and Challenges of Pulsar Timing

So far, we have not considered the influence of gravitational waves (GWs) on pulsar timing. Ultimately, we will be interested in measuring deviations in our timing residuals that are indicative of the influence of GWs on the Earth. Our sensitivity to GWs will depend in part on the uncertainty of our TOAs and our ability to model or otherwise account for so-called noise sources that will lead to non-zero residuals.

A TOA is calculated by cross-correlating a template and observed pulse profile. While this can be done in the time domain, far greater precision is achieved by working in the Fourier domain [3]. The TOA uncertainty is then $\sigma_{\text{TOA}} \sim \text{SNR}/\delta$, where SNR is the signal-to-noise of the observed profile and δ is the pulse width. It is common to obtain $\sigma_{\text{TOA}} \sim 1$ milliperiod, which for MSPs corresponds to $\sim \mu\text{s}$. However, for the best MSPs $\sigma_{\text{TOA}} \sim 10$ ns is possible.

Processes that lead to non-zero residuals can be broadly classified into two groups, depending on whether their power spectral density is flat or rises towards low frequency, i.e. *white* and *red noise* processes (see the contribution by Y. Wang in these proceedings). White noise processes include radiometer noise, pulse jitter, and interstellar scintillation, while red noise processes include timing noise, DM variations, and GWs. Each will be discussed in turn.

Radiometer noise arises from the observing system, atmosphere, cosmic microwave background, and Galactic synchrotron emission. Radiometer noise can be reduced by increasing observing time and radio bandwidth, observing with larger telescopes, and by using low-noise receivers and amplifiers. Interstellar scintillation (ISS) refers to scattering due to inhomogeneities in the ISM, analogous to the twinkling of stars. ISS can be mitigated by increasing observing bandwidths and integration times. Pulse jitter is a property of the pulsar itself [e.g. 4]. Although the average pulse profile is observed to be very stable over time, individual pulses vary widely in shape. Because the observed profile is comprised of a finite number of pulses, one never achieves a perfectly stable pulse profile. Pulse jitter can only be overcome by increasing observing time, i.e. summing more pulses. Pulse jitter is currently the limiting factor in timing stability for a few pulsars, and it will become more important as more sensitive observing systems reduce the impact of radiometer noise and ISS [5].

Timing noise is a well known process in non-recycled pulsars owing to torque variations on the neutron star. The effect is usually large enough that it contaminates the measurement of higher-order rotational frequency derivatives. It is observed to correlate with rotational period, and so is not prominent in most MSPs, though it may become so as timing precision improves and the length of data sets grow [6]. DM variations are more common (see the contribution from

L. Levin in these proceedings). The ISM is turbulent, and both the pulsar and the Earth are in motion with respect to the ISM, so that over time the pulsar signal travels through a different column density of free electrons, i.e. different DMs. As such, Δ_{DM} in Equation 2 becomes $\Delta_{\text{DM}}(t)$. High-precision measurements of DM must be made for each observing epoch.

4. Using Pulsars to Detect Gravitational Waves

Assuming that we have accurately calculated all timing parameters and effectively mitigated all noise sources, we are left with timing residuals that may contain a measurable effect of GWs [7, 8]. The nature of the signature depends on the GW source. Continuous GWs emitted by a single supermassive binary black hole (SMBBH) will lead to sinusoidal deviation in timing residuals, while the superposition of all GW-emitting SMBBH in the Universe will have a stochastic signature (see the contribution by C. Moore in these proceedings). The goal is to detect and characterize this signature with a high degree of confidence. To do this, we should search for characteristics that are uniquely attributable to GWs, as opposed to other processes, namely the quadripolar signature of GWs.

GWs will influence the pulsars themselves, as well as the telescopes we use to observe them. The so-called *pulsar term* will not be correlated between different pulsars, and although it is recoverable in principle if the distance to the pulsar is accurately known, in practice it is not likely to be accessible using current analysis techniques. The *Earth-term*, however, will be correlated between different pulsars, and with a unique dependence on the angular separation between pulsar pairs owing to the quadripolar nature of GWs. This forms the basis for pulsar timing arrays (PTAs). Correlated signatures in timing residuals can arise from other processes, but not with this quadripolar nature. For example, an error in Δ_{clock} will have a monopolar signature, while an error in $\Delta_{\text{S}\odot}$ (due to an incorrect planetary mass) will have a dipolar signature.

The precise form of the correlation between pulsars is the famous Hellings-Downs curve, or overlap reduction function [9] (see Figure 2). Conceptually, it can be formed by simply cross-correlating timing residuals from pulsars distributed across the sky. The frequency of GWs that can be detected using this technique is bounded by the cadence of pulsar observations and the time-span of the timing data-set. This is currently $\sim 1/\text{years}$ – $1/\text{weeks}$, or $\sim 10^{-9}$ Hz– 10^{-7} Hz. At these frequencies, the most likely sources are believed to be SMBBH in the early stages of inspiral. Cosmic strings, if they exist, are also expected to emit at these frequencies. Inflationary GWs may also exist in this frequency range, but are likely to be well below PTA sensitivity. As such, PTAs are an excellent complement to CMB polarization experiments and space or ground based laser interferometers (see Figure 2).

In practice, the small strain of GWs is expected to manifest in timing residuals at the level of ~ 10 ns, and it is technically challenging to obtain the necessary timing precision and theoretically challenging to model the relevant noise sources, while timing noise could pose a floor to the achievable residuals. Computationally, timing model parameters and noise processes are modeled simultaneously in a GW detection pipeline, which requires the inversion of large matrices and advanced analysis techniques. Nonetheless, several teams around the world are making excellent progress addressing all of these challenges.

5. PTA Collaborations

There are currently three major PTA collaborations. The North American Nanohertz Observatory for Gravitational Waves (NANOGrav) is made up of researchers at institutions in the United States and Canada [10]. NANOGrav primarily uses the 100-m Green Bank Telescope (GBT) and the 305-m Arecibo Observatory to time 42 MSPs with approximately monthly cadence. Arecibo is the most sensitive radio telescope in the world but has a limited field-of-view and can only see a subset of the pulsars. The GBT, which is fully steerable and can observe to a declination of $\delta \approx -46^\circ$, observes the MSPs that Arecibo cannot.




Figure 2. *Left:* The Hellings-Downs curve showing the expected angular correlation between pulsar timing residuals as a function of angular separation. The maximum correlation factor is 0.5, because only the Earth-term is correlated, while the pulsar-term is not. *Right:* A schematic representation of expected sources in the frequency range of PTAs and space and ground based interferometers. Approximate sensitivity curves are shown for various experiments. Figures courtesy of the NANOGrav collaboration.

The European Pulsar Timing Array (EPTA) [11] uses five telescopes: The 100-m Effelsberg Radio Telescope in Germany, the 76-m Lovell Telescope in England, the 64-m Sardinia Radio Telescope in Italy, the Westerbork Synthesis Radio Telescope (made up of 14, 25-m dishes) in the Netherlands, and the Nançay Radio Telescope in France, a Kraus-type design consisting of a flat primary and cylindrical secondary surface. The EPTA is able to use these telescopes to observe with a higher cadence than other PTAs. The LEAP (Large European Array for Pulsars) project is working to use all these telescopes as a phased array. When completed, the array will have a collecting area equivalent to a 194-m telescope.

The Parkes Pulsar Timing Array (PPTA) uses the 64-m Parkes Observatory in Australia [12]. As the southernmost telescope currently used for PTA observing, Parkes is able to see many southern hemisphere MSPs that are not visible to NANOGrav and EPTA telescopes, including PSR J0437–4715, the closest, brightest, and best timed MSP. PPTA members are primarily found in Australia and China.

All three major PTAs collaborate under the umbrella of the International Pulsar Timing Array (IPTA) [13]. The IPTA is a “collaboration of collaborations” with the goal of facilitating GW science through data and code sharing and greater cooperation. An IPTA data-set consisting of timing data from all three PTAs is currently being assembled. When complete, it will provide the greatest sensitivity to nanohertz-frequency GWs. But activities in all three PTAs and the IPTA are not limited to timing observations and data analysis. Surveys for new MSPs, instrumentation development, theoretical modeling of the nanohertz GW Universe, multi-messenger astronomy, and education and public outreach are all major projects in each PTA and the IPTA as a whole. Furthermore, PTA data sets have utility far beyond detecting GWs. Individual pulsars themselves are often excellent astrophysical laboratories, and there are currently IPTA-wide projects to develop a pulsar-based time standard (independent of terrestrial time standards) and to refine Solar system ephemerides.

6. Current Results and Future Prospects

The best limit on the amplitude of the stochastic GW background currently comes from the PPTA [14] (see Figure 3). These results are putting significant constraints on certain SMBBH merger models. The current best limit on the amplitude of single sources come from NANOGrav [15], and constrains the chirp mass of any SMBBHs in the Virgo cluster.

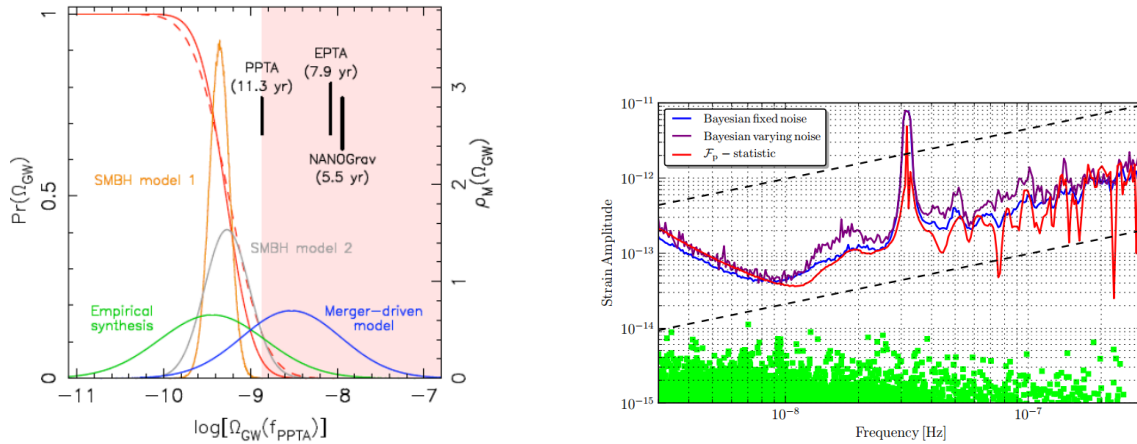


Figure 3. The current best results from PTAs. *Left:* The PPTA 11.3-year 95% confidence upper limit on the density of GWs, Ω_{GW} due to a stochastic background, from [14]. Probability density functions for the predictions of various models for SMBBH inspirals are also shown, with the PPTA results excluding the shaded region. Earlier NANOGrav and EPTA results are also shown. *Right:* NANOGrav 95% confidence upper limits on strain amplitude due to individual, continuous GW sources from [15]. Both a Bayesian (blue and purple lines) and frequentist (red line) analysis are shown, along with the amplitude of the strongest sources in 1000 Monte Carlo realizations based on [16]. Dashed lines indicate the expected strain amplitude for a source in Virgo with chirp mass of $10^9 M_\odot$ and $10^{10} M_\odot$.

These limits are expected to improve substantially over the next 5–10 years as the length of the data sets increase and new pulsars are added to the arrays. The upcoming 9-year data release from NANOGrav and the combined IPTA data release will improve upon the PPTA result. Based on current estimates, a GW detection is likely in as little as three years, and almost certainly by 2022. The biggest variable is the unknown amplitude of the stochastic GW background, which will likely be the first detection, though the red noise properties of the data set (mostly arising from possible timing noise) are also important.

Detection is merely the first step in studying the low-frequency GW Universe. The frequency dependence of the stochastic background will shed light on the merger history of galaxies. PTAs will also contribute to multi-messenger astronomy: they are sensitive to SMBBHs with separations $\sim 10^{-1}$ – 10^{-4} pc, and electromagnetic counterparts at these stages of merger include circumbinary disk emission, dual nuclei, and periodic changes in radio jets, among others. However, alternative explanations for some of these phenomena are available, so a PTA single-source detection can provide supporting evidence for a SMBBH interpretation.

New telescopes will significantly contribute to PTA science. The Five Hundred Meter Aperture Spherical Telescope being built in China will provide Arecibo-like sensitivity over a wider range of declinations, significantly increasing the timing precision that can be achieved for many pulsars. The increased sensitivity will also allow pulsars that are currently too weak to achieve the requisite timing precision to be added to the arrays. No current or planned telescopes will be able to match the sensitivity expected from the Square Kilometer Array, which will time hundreds of pulsars with sub-microsecond precision. This will greatly increasing PTA sensitivity, especially to single sources as this is currently dominated by the two or three most precise MSPs. In the meantime, new receivers and back-ends will improve upon the telescopes that are already in use.

Over the next several years and decades, PTAs will provide insight into a new and unique regime of the GW Universe, nicely complementing ground and space based interferometers.

7. Additional Resources

The Handbook of Pulsar Astronomy [17] provides an excellent introduction to pulsar timing and many other areas of pulsar astronomy. More information about all areas of PTA science can be found in the focus issue of *Classical and Quantum Gravity* (2013, Volume 30, Issue 22).

References

- [1] Lyne A, Hobbs G, Kramer M, Stairs I and Stappers B 2010 *Science* **329** 408–412 (*Preprint* [1006.5184](#))
- [2] Stovall K, Lynch R S, Ransom S M, Archibald A M, Banaszak S, Biwer C M, Boyles J, Dartez L P, Day D, Ford A J, Flanigan J, Garcia A, Hessels J W T, Hinojosa J, Jenet F A, Kaplan D L, Karako-Argaman C, Kaspi V M, Kondratiev V I, Leake S, Lorimer D R, Lunsford G, Martinez J G, Mata A, McLaughlin M A, Roberts M S E, Rohr M D, Siemens X, Stairs I H, van Leeuwen J, Walker A N and Wells B L 2014 *ApJ* **791** 67 (*Preprint* [1406.5214](#))
- [3] Taylor J H 1992 *Royal Society of London Philosophical Transactions Series A* **341** 117–134
- [4] Liu K, Keane E F, Lee K J, Kramer M, Cordes J M and Purver M B 2012 *MNRAS* **420** 361–368 (*Preprint* [1110.4759](#))
- [5] Shannon R M, Osłowski S, Dai S, Bailes M, Hobbs G, Manchester R N, van Straten W, Raithel C A, Ravi V, Toomey L, Bhat N D R, Burke-Spolaor S, Coles W A, Keith M J, Kerr M, Levin Y, Sarkissian J M, Wang J B, Wen L and Zhu X J 2014 *MNRAS* **443** 1463–1481 (*Preprint* [1406.4716](#))
- [6] Shannon R M and Cordes J M 2010 *ApJ* **725** 1607–1619 (*Preprint* [1010.4794](#))
- [7] Foster R S and Backer D C 1990 *ApJ* **361** 300–308
- [8] Jenet F A, Hobbs G B, Lee K J and Manchester R N 2005 *ApJL* **625** L123–L126 (*Preprint* [astro-ph/0504458](#))
- [9] Hellings R W and Downs G S 1983 *ApJL* **265** L39–L42
- [10] McLaughlin M A 2013 *Classical and Quantum Gravity* **30** 224008 (*Preprint* [1310.0758](#))
- [11] Kramer M and Champion D J 2013 *Classical and Quantum Gravity* **30** 224009
- [12] Hobbs G 2013 *Classical and Quantum Gravity* **30** 224007 (*Preprint* [1307.2629](#))
- [13] Manchester R N and IPTA 2013 *Classical and Quantum Gravity* **30** 224010 (*Preprint* [1309.7392](#))
- [14] Shannon R M, Ravi V, Coles W A, Hobbs G, Keith M J, Manchester R N, Wyithe J S B, Bailes M, Bhat N D R, Burke-Spolaor S, Khoo J, Levin Y, Osłowski S, Sarkissian J M, van Straten W, Verbiest J P W and Wang J B 2013 *Science* **342** 334–337 (*Preprint* [1310.4569](#))
- [15] Arzoumanian Z, Brazier A, Burke-Spolaor S, Chamberlin S J, Chatterjee S, Cordes J M, Demorest P B, Deng X, Dolch T, Ellis J A, Ferdman R D, Garver-Daniels N, Jenet F, Jones G, Kaspi V M, Koop M, Lam M, Lazio T J W, Lommen A N, Lorimer D R, Luo J, Lynch R S, Madison D R, McLaughlin M, McWilliams S T, Nice D J, Palliyaguru N, Pennucci T T, Ransom S M, Sesana A, Siemens X, Stairs I H, Stinebring D R, Stovall K, Swiggum J, Vallisneri M, van Haasteren R, Wang Y and Zhu W W 2014 *ArXiv e-prints* (*Preprint* [1404.1267](#))
- [16] Sesana A 2013 *MNRAS* **433** L1–L5 (*Preprint* [1211.5375](#))
- [17] Lorimer D R and Kramer M 2012 *Handbook of Pulsar Astronomy*