Using Pulsar Timing Arrays to Detect Gravitational Waves from Supermassive Black Hole Binaries

April 29th, 2025
Clementine Mitchell

Introduction

First discovered in 1967 by Jocelyn Bell Burnell and Antony Hewish, pulsars are rapidly spinning neutron stars that emit beams of electromagnetic radiation from their magnetic poles. As these beams sweep past Earth, they are observed as regular pulses of radio waves. Some pulsars, especially millisecond pulsars that have been spun up by accreting matter from a companion star, can be so regular that they rival the best atomic clocks on Earth.

This timing stability makes pulsars great tools for a wide variety of astronomical measurements. One such application is the use of a networks of pulsars – so-called Pulsar Timing Arrays (PTAs) – to detect gravitational waves (GWs). Gravitational waves are ripples in the “fabric of spacetime”, first predicted by Albert Einstein around 1915 as a consequence of his General Theory of Relativity.

While ground-based detectors, like LIGO and Virgo, have already made observations of gravitational waves from stellar-mass black hole and neutron star mergers, they are limited in the range of frequencies they can detect. In particular, the ultra-low-frequency gravitational waves emitted by supermassive black hole (SMBH) binaries are not detectable by these ground-based detectors. PTAs offer a complementary method that could detect this frequency band and thus give us more information about some of the more exotic interactions in our Universe.

In this paper, I will explain the basic physics behind pulsars and gravitational waves, I will describe how PTAs can be used to detect gravitational waves from supermassive black hole binaries, and I will mention what we might learn from these detections in the future.

Gravitational Waves

Einstein’s Relativity and the Prediction of Gravitational Waves

At the beginning of February in this class, we got an overview of Einstein’s relativity. Einstein’s General Theory of Relativity was formulated around 1915. In general relativity, gravity is not a force but the result of the curvature of spacetime caused by mass and energy. Spacetime itself can bend, stretch, and ripple. According to General Relativity, mass and energy tells spacetime how to curve, and this curvature tells the matter how to move. One consequence of the equations of General Relativity is the existence of gravitational waves. Gravitational waves are propagating disturbances in the curvature of spacetime, much like ripples spreading out on the surface of a pond.

Gravitational waves are generated by accelerating masses when mass distributions are non-spherical — for example, in a binary system where two massive objects orbit each other. As these objects move, they create changing quadrupole moments in the gravitational field, and these changes radiate away from the system at the speed of light as gravitational waves. Unlike Newtonian gravity, which acts instantaneously, general relativity predicts that changes in a mass distribution (like an orbiting binary system) produce waves that carry energy away from the system. This causes the system to slowly lose energy and spiral inward – a prediction confirmed by observations.

How We Detect Gravitational Waves Today

It was not until 2015 that LIGO (Laser Interferometer Gravitational-Wave Observatory) made the first confirmed observation of gravitational waves from a merger of two stellar-mass black holes. LIGO and similar observatories like Virgo detect gravitational waves by measuring tiny changes in the length of laser beams bouncing between mirrors separated by kilometers.

However, these ground-based detectors are sensitive to gravitational waves in the tens to thousands of hertz frequency range. Gravitational waves from SMBH binaries, on the other hand, have periods of years to decades, corresponding to nanohertz frequencies. Such low-frequency waves are undetectable with current interferometers because the wavelength is too large relative to the size of the detectors.

This is where PTAs come in. By observing the arrival times of pulsar signals across the sky over many years, astronomers can search for the characteristic patterns introduced by passing gravitational waves at nanohertz frequencies.

Pulsar Timing Arrays

How a Pulsar Timing Array Works

A Pulsar Timing Array is essentially a galaxy-scale gravitational wave detector. It consists of carefully timed observations of a network of millisecond pulsars spread across the sky. Because pulsars emit incredibly regular pulses, any deviation from the expected arrival time can hint at a perturbation in spacetime itself. [2]

When a gravitational wave passes between a pulsar and Earth, it stretches and squeezes spacetime along its path. This distortion leads to slight advances or delays in the arrival times of the pulses. By simultaneously monitoring many pulsars and looking for correlated patterns in the timing residuals (the difference between expected and observed pulse arrival times), astronomers can infer the presence of a gravitational wave background – or even detect individual SMBH binary systems.

Imagine Earth at the center, surrounded by an array of pulsars at various points on the celestial sphere. A gravitational wave passing through this setup would systematically distort the distances between Earth and the pulsars, introducing timing deviations that depend on the relative positions of the pulsars to the wave's propagation direction.

The Hellings-Downs Correlation

One theoretical result in PTA science that continued to be brought up in the papers I read is the Hellings-Downs correlation, which was described in a 1983 paper by Hellings and Downs. This correlation function predicts the expected correlation between timing residuals from pairs of pulsars as a function of the angle between them on the sky, assuming an isotropic stochastic gravitational wave background.

The Hellings-Downs curve has a very distinctive shape: pulsar pairs separated by small angles show strong positive correlation, while those separated by about 90 degrees have a weaker correlation. As the separation approaches 180 degrees (antipodal pulsars), the correlation becomes negative. This pattern serves as evidence for the presence of gravitational waves and helps distinguish GW signals from other noise sources, such as intrinsic pulsar timing noise. I plotted the Hellings-Down Curve as the red curve in Figure 1 below.

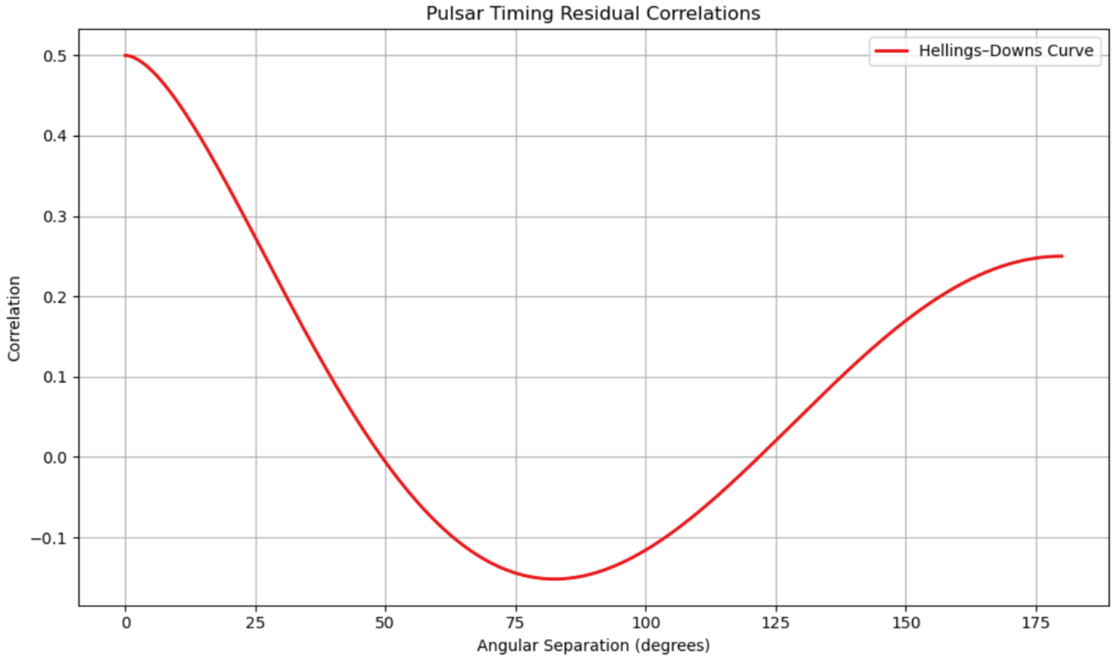

Figure 1: Hellings-Down Curve

Initially, I thought I could compare the Hellings-Down Curve to a model I built to show the impact of a SMBH binary on the timing residuals of various pulsars in a PTA to test how it was working. However, I realized during that process that the Hellings-Down correlations are only meaningful for a stochastic gravitational wave background (GWB), not for single events. The Hellings-Down Curve assumes an isotropic, stationary, Gaussian, and unpolarized background. That background is made of many weak, incoherent gravitational waves from different directions and frequencies – typically from thousands of supermassive black hole binaries across the universe [4]. Whereas the model I built is for a strong single source (like a nearby SMBH binary merger), which produces a coherent signal in all pulsars. This creates specific timing patterns depending on each pulsar's alignment to the source, but not the ensemble statistical structure the Hellings-Down curve describes.

Simulating Timing Residuals from a SMBH Binary

To better understand how gravitational waves affect pulsar timing, I built a simple Python model that simulates the arrival time delays for pulsars under the influence of a gravitational wave emitted by a SMBH binary. The idea is to compute the timing residuals induced by a plane gravitational wave passing over the Earth-pulsar system.

In this model, the timing residuals depend on:

- The position of the pulsar relative to the gravitational wave propagation direction.
- The gravitational wave's amplitude and frequency.
- The time of observation.

Different pulsars show different residual patterns because they are at different angles relative to the wavefront. Some pulsars experience slight advances, and others slight delays, depending on how the spacetime stretching aligns with their line of sight. [1]

The very simple model that I simulated includes:

- A single GW source (the SMBH binary) with fixed frequency and amplitude.
- A set of pulsars randomly distributed across the sky.
- Timing residuals over 10 years caused by the GW as it passes Earth.
- A plot showing the timing residuals for each pulsar over time.

The timing residual $R(t)$ for a pulsar depends on:

- The GW amplitude: h_0
- The GW frequency: f_{gw}
- The angular position of the pulsar relative to the GW source

- The phase shift due to the distance between Earth and the pulsar

Residuals are computed using a bunch of math, which (to first order) is proportional to:

$$R(t) = \frac{1}{2}\left\{\frac{1+\cos\theta}{1-\cos\theta} h_0 \frac{\sin\left(2\pi f_{gw} t + \emptyset\right) - \sin\left(2\pi f_{gw}\left(t - L\,(1-\cos\theta)/c\right) + \emptyset\right)}{2\pi f_{gw}}\right\}$$

where:

- θ: angle between GW propagation and pulsar
- L: distance to the pulsar
- c: speed of light

I generated a synthetic dataset of pulsars randomly distributed across the sky and plotted the timing residuals over time for each one for a gravitational wave that is like one that would be generated by a SMBH binary.

Figure 2 shows an example plot from the toy model I built in Python. It shows the timing residuals for several pulsars over a 10-year observation span (each colored curve shows the timing residual for various randomly distributed pulsars), demonstrating the tiny but systematic timing changes introduced by the passing gravitational wave.

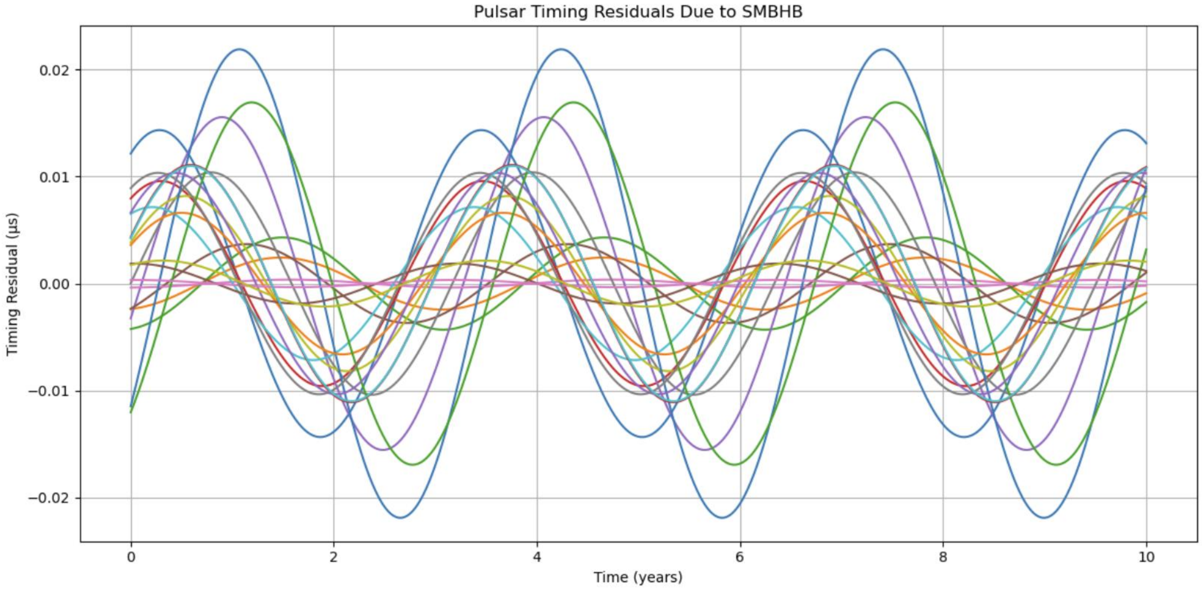

Figure 2: Model of Timing Residuals of Various Pulsars Due to a Gravitational Wave Modeled to Reflect that Caused by an SMBH Binary

<u>What We Could Learn from PTA Observations</u>

If PTAs successfully detect gravitational waves from SMBH binaries, it would open up a new realm of gravitational wave astronomy. We could:

- Measure the population and merger rate of SMBH binaries throughout the universe.
- Gain insights into the processes that drive galaxy mergers and the growth of supermassive black holes.
- Test aspects of General Relativity in a new, low-frequency regime.
- Potentially identify individual SMBH binaries, enabling multi-messenger studies combining radio and electromagnetic observations.

Furthermore, detecting the gravitational wave background itself – the "hum" produced by the ensemble of all SMBH binaries – would give us a glimpse into the merger history of galaxies across cosmic time.

The Future of Pulsar Timing Arrays

The field of PTA research is entering a very exciting phase. In 2023, the NANOGrav collaboration and other PTA groups around the world announced hints of a common-spectrum process across many pulsars, consistent with the early stages of a gravitational wave background detection. Continued observations, improved timing precision, and the addition of new pulsars to the arrays will strengthen the data and potentially lead to a definitive detection. [3]

New projects like the planned Square Kilometer Array (SKA) will dramatically enhance the sensitivity of PTAs, making it possible not only to detect the gravitational wave background but also to resolve individual sources. The field of the detection of low-frequency gravitational waves seems like it is just getting started.

Conclusion

Summary of What I Learned

Through this project, I gained an understanding of how pulsar timing arrays can be used as a method for detecting gravitational waves. Pulsars are cosmic clocks, allowing us to perform precision experiments on scales far beyond what we can achieve on Earth.

Reading the papers by Babak and Sesana (2017) and Maiorano et al. (2021), and Taylor et al. (2019) helped me understand not only the basic mechanisms behind PTA gravitational wave detection but also the broader astrophysical context. In particular, I am interested in the way that PTAs offer a unique opportunity to detect supermassive black hole binaries, which I think will give us a nice understanding of galaxy mergers in general, especially in the early Universe when galaxy mergers were more common (due to the smaller size of the Universe).

Final Thought

It was cool to see how simple ideas – pulsars as clocks, spacetime ripples introducing timing shifts, and the correlations between different pulsars – together into a powerful method for probing some of the universe’s most massive and mysterious objects.

Bibliography

[1] *Pulsar Timing Array Based Search for Supermassive Black Hole Binaries* – Babak and Sesana (2017)

[2] *Principles of Gravitational-Wave Detection with Pulsar Timing Arrays* – Maiorano et al. (2021)

[3] *Supermassive Black-hole Demographics & Environments With Pulsar Timing Arrays* – S. R. Taylor, S. Burke-Spolaor, et al. (2019)

[4] *The Predictor of Pulsar Timing* – Hellings and Down (1983)