

WILL PULSAR TIMING ARRAYS OBSERVE THE HELLINGS AND DOWNS CORRELATION CURVE?

Bruce Allen

MPI for Gravitational Physics, Callinstrasse 38, 30167 Hannover, Germany

Abstract

Pulsar timing arrays seek to detect gravitational waves (GW) by observing correlations which were predicted by Hellings and Downs in 1983. Here, we address the question, are these correlations going to be exactly what was predicted in 1983? The answer is no: interference between different GW sources creates a pattern of correlation that does not average to give the Hellings and Downs curve. We explain this effect, calculate the variance, and show that it is potentially observable.

1 Introduction

Thank you for inviting me to this wonderful place. The Vulcano workshops are famous not just for the interesting physics, but also for the friendly colleagues, the fantastic food, and the beautiful venue.¹ I am enjoying this a

¹This sentence remains correct under all 120 permutations of the adjectives.

lot, and really hope that you ask me to return for the next workshop in 2024.

My talk concerns pulsar timing arrays (PTAs), which are a way to detect low frequency (nHz) gravitational waves (GW). These waves have periods of years, meaning frequencies that lie far below the sensitivity band of detectors such as LIGO and LISA.

The results are described in detail in two arXiv preprints. The first of those is by me ¹⁾ and the second is work done in collaboration with Joe Romano ²⁾. All of the ideas and most of the figures in my talk come from those two preprints. So if some details are lacking, or if you want to learn more about this topic, please look there.




Figure 1: *The Hellings and Downs curve.*

This is the only talk about PTAs in this workshop, so I will spend some time introducing those. But I first want to show you the most important thing in my talk, which you may have seen before. This is the Hellings and Downs curve $\mu_u(\gamma)$, illustrated in Fig. 1. This function shows the average (hence the symbol “ μ ” for “mean”) correlation between the pulse arrival times (or pulsation frequency Doppler shifts) from two different pulsars, separated on the sky by angle γ as seen from Earth, induced by an unpolarized (hence the subscript “u”) isotropic GW stochastic background. For example, for two pulsars that are almost in the same direction on the sky (γ near 0°) you can see that, on the average, GWs induce correlated variations in the arrival times of the pulses. In contrast, for pulsars that are separated by about 90° on the

sky, on the average the effects of the GWs on the pulses is anticorrelated.

Observing this Hellings and Downs curve is important for PTAs. It plays the same role that the famous “chirp waveform” did for the first LIGO detection³⁾ of GWs. When we see this Hellings and Downs curve clearly, we can confidently proclaim “we have detected GWs”.

My talk addresses a simple question: should we expect to see *exactly* this curve? Or only something *close* to it? My conclusion: we will see this curve, but when enough good data is available, we will also see a certain deviation from it. While I can’t predict the sign of that deviation, I can predict its expected magnitude. Here, that expected (squared) deviation from the Hellings and Downs curve is called the variance, and is denoted by σ^2 .

2 Pulsar Timing Arrays

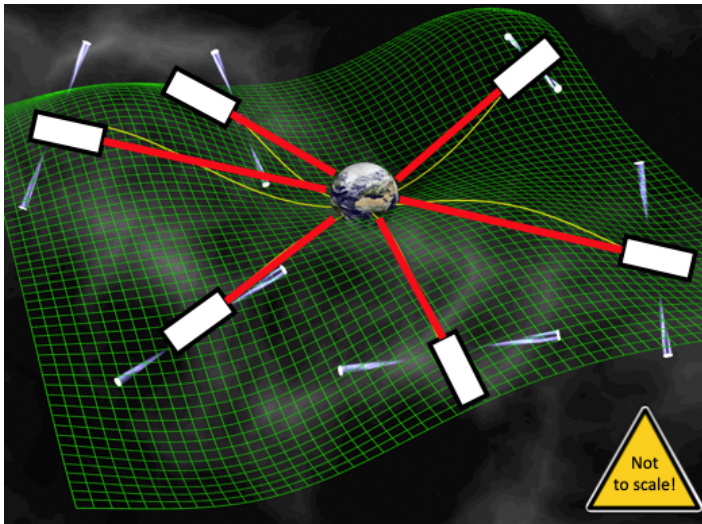


Figure 2: A pulsar timing array (PTA) made from six pulsars (a modified version of David J. Champion’s original illustration).

PTAs are galactic-scale gravitational wave detectors. Fig. 2 is a schematic diagram of a PTA made from six pulsars.² The GW sources are not shown –

²This figure is not to scale. For example, typical PTA pulsars are thousands

the nearest ones are probably at distances that are five orders of magnitude larger than the Earth-pulsar separations.

To stress that this is like a six-arm LIGO detector, in the figure I have replaced the pulsars with lasers. Conceptually, one could also replace the pulsars with perfect clocks, which tick at a few hundred Hz rather than at typical laser frequencies of 10^{15} Hz. The idea is that when a GW goes flying by, it redshifts or blueshifts those clocks.³ Because the clock frequency is so low, there is no light or color, so it might be better to say “Doppler shift”. But I keep to the tradition of the literature, which uses “redshift” and “blueshift”.

The data stream from each pulsar is a redshift $Z = \Delta f/f$ as a function of time, where f is the mean pulsation frequency and Δf is the decrease in the frequency at time t . Typical PTA pulsars are observed every week or two for decades, so the time series consists of hundreds or thousands of redshift measurements.




Figure 3: *Redshift/blueshift of a pulsar signal induced by a GW that oscillates through three cycles with an oscillation period of five years.*

Here is an example. Suppose that a fixed-frequency GW, consisting of of years from Earth, but the GWs they detect, shown by the green ripples, should have wavelengths that are $O(10^2)$ times smaller.

³While pulsars are observed with terrestrial radio telescopes, their pulsation frequencies are then determined at the solar system barycenter, to remove effects of Earth’s motion and the gravitational effects of the Sun and planets.

three cycles with an oscillation period of five years, went flying by this detector. Fig. 3 shows the pattern of redshift/blueshift observed in one of these pulsars, in the absence of any noise. The maximum redshift corresponds to a lowering of the observed pulsar rotation frequency by three parts in 10^{15} . The maximum blueshift is the same fractional increase in the observed rotation frequency.

If you had a single perfect noise-free pulsar, then you could observe GWs simply by monitoring the pulse arrival times. The GW frequency is the frequency of this redshift/blueshift oscillation. The GW strain amplitude h is the maximum fractional frequency change, which in my example is $O(10^{-15})$. Because pulsars are monitored for decades with a timing precision of hundreds of nanoseconds, these small shifts are observable. However, because pulsars are not free of noise, PTAs must search for GWs by looking for a common signal, which appears the same in the different “pulsar arms”.

There are three active PTAs. The European Pulsar Timing Array (EPTA) currently monitors 42 pulsars ⁴⁾. The North American nanoHz Gravitational Wave Observatory (NANOGrav) currently monitors 66 pulsars ⁵⁾. The Parkes Pulsar Timing Array (PPTA) currently monitors 26 pulsars ⁶⁾. In all, the International Pulsar Timing Array (IPTA), which is an umbrella organization for all three PTAs, monitors 88 pulsars. (This number is smaller than you might have expected because many of the pulsars are common to two or more of the PTAs.)

The PTA data sets, which span several decades, show intriguing evidence for a stochastic background of GWs. A plausible source of these GWs is the slow orbital decay of supermassive black hole binaries ⁷⁾. We know that most galaxies have supermassive black holes at their centers. When galaxies merge, the black holes at their centers form binary systems, whose orbits decay due to interactions with other stars and with their environment. Those binaries eventually become close enough to orbit with periods of years or decades, emitting continuous gravitational waves at twice the orbital frequency. This would create a signal in the PTA band.

We expect that the closest of these supermassive black hole binaries is at a distance of order 50 Mpc, and there would be a much larger number of similar sources at greater distances, extending out to near the Hubble radius. The GW signals that these produce sum up to create a stochastic confusion-noise GW background, which has the statistical properties of a central-limit-theorem

Gaussian ensemble 8).

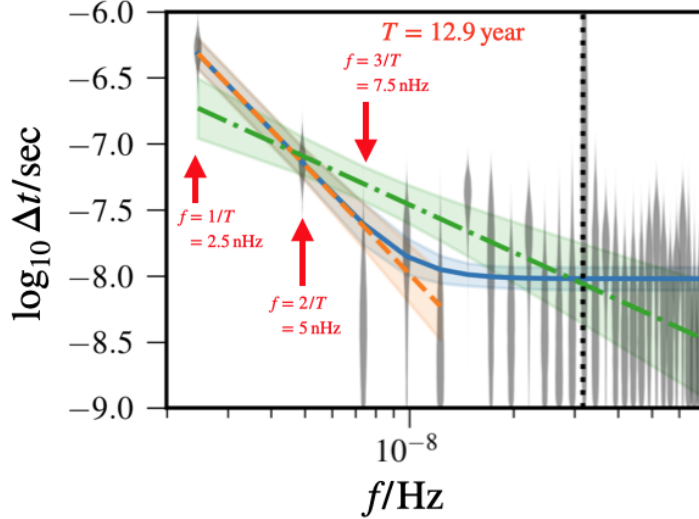


Figure 4: *The timing noise seen in the NANOGrav 12.5 year data set is consistent with a GW stochastic background. Reproduced from 5).*

Shown in Fig. 4 is $T = 12.9$ years of NANOGrav data 5); the data from other PTAs is similar 4, 6). This shows the result of Fourier-transforming the time series of pulsar timing residuals Δt , which are the time integrals of the redshift Z . The amplitudes are shown for frequency bins corresponding to the fundamental mode at $f = 1/T = 2.5$ nHz, $2/T = 5$ nHz and so on. This data, which only shows autocorrelations, has the sort of increasing amplitude at low frequencies which would be expected from a supermassive binary black hole background. The best fit power law (orange line) to the first five harmonics has a slope $\approx f^{-2.6}$. This is close to the $f^{-13/6}$ power law expected from supermassive black hole binaries. The characteristic strain amplitude $h_c \approx 1.9 \times 10^{-15}$ is also consistent with expectations for that source.

In my opinion, this is good evidence for a GW background. But there are other possible explanations. Perhaps, for example, all pulsars have some intrinsic source of rotation noise, with a power-law spectrum similar to that which would be produced by merging black hole binaries.

3 The Hellings and Downs curve $\mu_u(\gamma)$

Fortunately there is a simple way to understand if the source of these observed pulsar timing fluctuations is a GW background, and, if it is due to GWs, to gain confidence. This is to observe the pattern of correlation shown in Fig. 1. This pattern was first calculated by Ron Hellings and George Downs in 1983⁹, and is called the ‘‘Hellings and Downs curve’’. I want to describe how pulsar timing arrays might see this curve, which is described by the function

$$\mu_u(\gamma) = \frac{1}{4} + \frac{1}{12} \cos \gamma + \frac{1}{2} (1 - \cos \gamma) \log \left(\frac{1 - \cos \gamma}{2} \right). \quad (1)$$

The take-home message from my talk is that PTAs will not observe exactly this curve. However, they will observe something close to this curve, and we can predict and calculate the magnitude of the expected deviations away from it. This is what I mean by the variance in the Hellings and Downs correlation.

Here is what Hellings and Downs did in 1983. First, they placed a single distant unit-amplitude unpolarized GW source at a point Ω on the sky. (Here Ω is a unit-length three-dimensional vector.) Next, they wrote down the correlation ρ between the redshifts (or timing residuals) of two pulsars a and b , separated by an angle γ on the sky, which is

$$\begin{aligned} \rho &= \Re(F_a(\Omega)F_b^*(\Omega)) \\ &= F_a^+(\Omega)F_b^+(\Omega) + F_a^\times(\Omega)F_b^\times(\Omega). \end{aligned} \quad (2)$$

This is the product of the response of pulsar a to a GW point source at direction Ω with the response of pulsar b to that same source. The appearance of the real part, and the complex conjugate on the first line of Eq. (2), are because I use a complex polarization basis for the GWs. The real part of F is the response to the plus polarization, and (minus) the imaginary part is the response to the cross polarization: $F = F^+ - iF^\times$. The second line expresses this correlation in terms of these (real) linear polarizations. The functions F correspond to what are called ‘‘antenna patterns’’ in the context of LIGO or LISA. They depend upon the sky direction to the pulsar, as well as on the source direction Ω .

Finally, Hellings and Downs averaged the correlation ρ over all pairs of pulsars a and b separated by angle γ , assuming that these were uniformly distributed around the sky.⁴ We call this a ‘‘pulsar sky average’’. What they

⁴In fact Hellings and Downs fixed the pulsar positions and averaged over

found is exactly the function given in Eq. (1). That is to say, $\mu_u(\gamma) = \langle \rho \rangle$, where angle brackets denote the average over all pulsar pairs a, b separated by angle γ . People often call this a “quadrupole response pattern”, which is roughly correct. While it also includes higher modes, the quadrupole term is the largest one ²⁾.

The calculation corresponds to observational practice. Imagine that you are a PTA observer with access to many pulsars, distributed all around the sky. To determine the mean correlation at angle γ , you take all pairs that (within some tolerance window) are separated by angle γ , and average their correlations. This observational average can be directly compared to $h^2 \mu_u(\gamma)$, where h^2 is the characteristic squared amplitude of the GWs. (This factor of h^2 appears because the calculation that led to Eq. (1) is for a unit amplitude source of GWs.)

However, it is surely the case that our Universe is not like this simple model. As explained earlier, we expect that the Universe contains a very large number of GW point sources, rather than a single point source. So, what does the Hellings and Downs calculation and curve have to do with reality? Let us investigate this question.

4 The Hellings and Downs correlation for two GW sources

Imagine now that we have two GW sources, rather than one. For example, put the first source directly over the North pole (direction Ω_1) and the second source directly overhead us, here in Elba (direction Ω_2). We’re going to repeat the Hellings and Downs calculation, to find the average correlation of two pulsars separated by angle γ , under the influence of two GW sources.

The sources have complex waveforms $h_1(t)$ and $h_2(t)$; the real part is the plus polarization and the imaginary part is the cross polarization: $h = h^+ + ih^\times$. For simplicity, assume that both sources are unpolarized, which implies $\overline{h_1 h_1} = \overline{h_2 h_2} = 0$. (Here, the overline means “average over time”.) The real parts of these equation ensure that the average plus and cross intensities are equal, and the imaginary parts ensure that the product of the plus and cross amplitudes averages to zero. (If you look at the definition of the Stokes parameters, you’ll

source directions on the celestial sphere. From symmetry these are equivalent ¹⁰⁾.

see that this is what is meant by “unpolarized”.) The squared strain of the two GW sources are real quantities, given by $\overline{h_1 h_1^*} = |\overline{h_1}|^2$ and $\overline{h_2 h_2^*} = |\overline{h_2}|^2$ respectively.

The redshifts of pulsars a and b are obtained by summing the effects of the two GW sources:

$$\begin{aligned} Z_a &= \Re(h_1 F_a(\Omega_1) + h_2 F_a(\Omega_2)) \\ &= \frac{1}{2} (h_1 F_a(\Omega_1) + h_1^* F_a^*(\Omega_1) + h_2 F_a(\Omega_2) + h_2^* F_a^*(\Omega_2)) \\ Z_b &= \frac{1}{2} (h_1 F_b(\Omega_1) + h_1^* F_b^*(\Omega_1) + h_2 F_b(\Omega_2) + h_2^* F_b^*(\Omega_2)). \end{aligned} \quad (3)$$

Note that the final line is just the previous one with pulsar a replaced by pulsar b .

Now, we are going to consider two different possibilities for the behavior of these GW sources. In the first case, the two sources are not going to interfere with each other. This means that they are radiating at different frequencies: if we multiply their GW waveforms together and integrate over time, we get zero. Thus, $\overline{h_1 h_2} = \overline{h_1 h_2^*} = 0$. Later, I’ll consider the case where they do interfere.

By multiplying the final two lines of Eq. (3) and then time averaging the product, you can calculate the correlation $\overline{Z_a Z_b}$ between pulsars a and b . It might sound complicated, but it’s simple – you can do it in your head. There are 16 possible terms in the product, but after taking the time average, 12 terms vanish. The only nonzero terms are when h_1 multiplies h_1^* or h_2 multiplies h_2^* , giving

$$\rho = \overline{Z_a Z_b} = \frac{1}{2} \overline{|h_1|^2} \Re(F_a(\Omega_1) F_b^*(\Omega_1)) + \frac{1}{2} \overline{|h_2|^2} \Re(F_a(\Omega_2) F_b^*(\Omega_2)). \quad (4)$$

These two terms look exactly the same as Eq. (2), which Hellings and Downs used in their 1983 calculation. So if we average Eq. (4) over all pulsar pairs separated by angle γ , then we get *exactly* the Hellings and Downs curve, $\langle \rho \rangle = \frac{1}{2} (|\overline{h_1}|^2 + |\overline{h_2}|^2) \mu_u(\gamma)$. It does not matter where the independent GW sources are located on the sky relative to each other: the pulsar average always gives exactly the Hellings and Downs curve.

Now, let us repeat the calculation for GW sources that interfere. For this, assume that $\overline{h_1 h_2} = 0$ but that $\overline{h_1 h_2^*}$ is real and positive. (These equations imply that the plus components of h_1 are uncorrelated with the cross components of h_2 and vice versa. However, they also imply that the plus components of

h_1 and h_2 are correlated with each other, and that their cross components have that same degree of correlation.) As before, we need to multiply the last two lines of Eq. (3) and average over time. Again, while this might sound complicated, you can do it in your head. In addition to the four terms that appeared for uncorrelated sources, we now get four additional terms where an h_1 multiplies the complex conjugate of h_2 , or vice versa. This gives

$$\rho = \overline{Z_a Z_b} = \frac{1}{2} \overline{|h_1|^2} \Re(F_a(\Omega_1) F_b^*(\Omega_1)) + \frac{1}{2} \overline{|h_2|^2} \Re(F_a(\Omega_2) F_b^*(\Omega_2)) + \frac{1}{2} \overline{h_1 h_2^*} \Re(F_a(\Omega_1) F_b^*(\Omega_2) + F_a(\Omega_2) F_b^*(\Omega_1)). \quad (5)$$

In contrast to the case of independent sources, the cross term, proportional to the (real) time average $\overline{h_1 h_2^*}$, is nonzero. If we average this correlation ρ over all pulsar pairs separated by angle γ , the first line, which is the same as in the independent source case, averages to give the Hellings and Downs curve, *but the second line does not*. If you average the second line over all pairs of pulsars separated by angle γ , it gives a different function of angle γ than the Hellings and Downs curve. In a minute, I'll show you what that function looks like.

That's the take-home message of my talk. After the various PTAs have observed enough pulsars, and determined the average correlation at angle γ , this interference term means that they won't observe exactly the Hellings and Downs curve. This is because our Universe contains many GW sources, with independent positions and GW phases, radiating in the lowest frequency bins. These generate a standing wave pattern whose pulsar average, in any representative universe, is not the Hellings and Downs curve.

5 Variance of the Hellings and Downs correlation for many GW sources

Let us now consider what happens when there are many GW sources radiating at the same frequency, so there is lots of interference. I'll denote the number of these sources by the integer N . So now, the response of each pulsar has N terms, and the time-averaged correlation, obtained as a product, has N^2 terms. If we label the sources by $j = 1, \dots, N$, then the pulsar-averaged correlation curve is

$$\langle \rho \rangle = \sum_j h_j^2 \mu_u(\gamma) + \sum_{j \neq k} h_j h_k \cos(\phi_j - \phi_k) \mu(\gamma, \beta_{jk}). \quad (6)$$

Here, the two-point function $\mu(\gamma, \beta)$ is shown in Fig. 5, h_j is the GW amplitude of the j 'th source, $\phi_j \in [0, 2\pi)$ is the GW phase of that source, and β_{jk} is the angle on the sky between sources j and k . (From here on, h_j is just a positive real number, whereas in Sec. 4, h_1 and h_2 denoted functions of time.)




Figure 5: The two-point function $\mu(\gamma, \beta)$. The cross-section at $\beta = 0$ is the Hellings-Downs curve: $\mu_u(\gamma) = \mu(\gamma, 0)$. An explicit formula for $\mu(\gamma, \beta)$ is derived in ¹.

Look carefully at Eq. (6). The first sum is the “diagonal” terms, where source j interferes with itself. The pulsar average of these gives exactly the Hellings and Downs curve. Then there are the “off-diagonal” terms, meaning the sum over $j \neq k$. These come from different sources interfering with each other. The product of the amplitudes of those two sources is multiplied by the cosine of the phase difference between the sources. These phases are indepen-

dent random numbers, different for each source, so that the value of the cosine is uniformly distributed in the interval $[-1, 1]$. Finally, this is multiplied by a function which I call the two-point function, which is illustrated in Fig. 5.

The function $\mu(\gamma, \beta)$ is a function of the angle γ between the two pulsars, and of the angle β between the two sources. It is defined by the pulsar average

$$\mu(\gamma, \beta_{jk}) = \langle F_a(\Omega_j) F_b^*(\Omega_k) \rangle . \quad (7)$$

Here β_{jk} is the angle between the sources: $\cos \beta_{jk} = \Omega_j \cdot \Omega_k$. The angle brackets mean “average over uniformly distributed pulsars a and b separated by angle γ on the sky”. You will recognize that this is precisely the pulsar average of the extra “interference” term that appeared in Eq. (5), when we looked at two interfering GW sources. (Note: after the pulsar average in Eq. (7), only the real part remains.)

The important thing about Eq. (6) is this. The first sum, the diagonal terms, gives us something proportional to the Hellings and Downs curve. But the second sum, the off-diagonal terms, adds up different cross-sections of the plot in Fig. 5, at values of $\beta \neq 0$. Those cross-sections are *not* proportional to the Hellings and Downs curve $\mu_u(\gamma)$. So, because of the interference between different GW sources, the pulsar-averaged correlation is *not* proportional to the Hellings and Downs curve.

6 Cosmic variance in the Hellings and Downs correlation

We can calculate the variance of the pulsar-averaged correlation from the Hellings and Downs curve. Start with Eq. (6), subtract the diagonal term, which *is* proportional to the Hellings and Downs curve, square the difference, and average over sources with independent random phases uniformly distributed on the sky. One obtains the cosmic variance ¹⁾

$$\begin{aligned} \sigma_{\text{cosmic}}^2(\gamma) &= \int_0^\pi d\beta \sin \beta \mu^2(\gamma, \beta) \\ &= -\frac{5}{48} + \frac{49}{432} \cos^2 \gamma - \frac{1}{6} (\cos^2 \gamma + 3) \log\left(\frac{1 - \cos \gamma}{2}\right) \log\left(\frac{1 + \cos \gamma}{2}\right) + \\ &\quad \frac{1}{12} (\cos \gamma - 1) (\cos \gamma + 3) \log\left(\frac{1 - \cos \gamma}{2}\right) + \\ &\quad \frac{1}{12} (\cos \gamma + 1) (\cos \gamma - 3) \log\left(\frac{1 + \cos \gamma}{2}\right) . \end{aligned} \quad (8)$$

This is the typical (squared) deviation away from the Hellings and Downs curve, for a universe filled with interfering GW sources, and is shown in Fig. 6 and Fig. 7. (Note: it might appear that there is no cosmic variance around $\gamma = 50^\circ$ and $\gamma = 130^\circ$. In fact, the variance there is small, but positive.)




Figure 6: *The cosmic variance characterizes the expected deviation from the Hellings and Downs curve. (Plot is for a GW confusion-noise model, with $h^2 = 1$ and $h^4/h^4 = 1/2$, see ²).*

I want to give you an idea of the size of the cosmic variance. This is shown in Fig. 6. The solid black line is the Hellings and Downs curve. If you pick a random pair of pulsars separated by angle γ , they will have a correlation which lies (± 1 sigma) in between the two outer dashed lines. If you are a PTA observer, and you average over many pulsars on the sky separated by angle γ , you'll end up at the Hellings and Downs curve, plus or minus the amount shown by the dotted line, which is the (square root of the) cosmic variance of Eq. (8). This difference arises from the interference between GW sources radiating at the same frequency. That interference generates a standing wave pattern which doesn't average to give the Hellings and Downs curve.

7 How close can PTAs get to the Hellings and Downs curve?

Real pulsar timing arrays don't have access to an infinite set of pulsars, uniformly distributed on the sky. They only observe a finite number of pulsars. What happens is that as you add more pulsars to your array, you decrease the

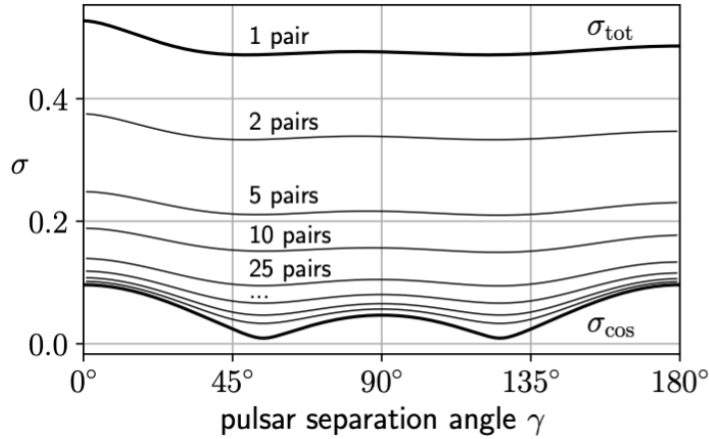


Figure 7: *The variance decreases as more pulsar pairs are added, but does not decrease to zero. Instead, it converges to the cosmic variance. (Plot is for a GW confusion-noise model, with $h^2 = 1$ and $h^4/h^4 = 1/2$, see ²⁾.)*

variance away from the Hellings and Downs curve. This is shown in Fig. 7. If you start with a single pair of pulsars at angle γ , then the variance is the top curve. As more pairs of pulsars are added, the variance decreases as shown, eventually converging to the cosmic variance, shown by the bottom curve.

There is another way to think about and to derive the cosmic variance ²⁾: it arises from the correlations between different pairs of pulsars. Once you have enough pulsar pairs, adding additional pairs at similar angular separations does not provide new information. So adding pairs does not reduce the variance to zero: there is always some remaining difference. At angles (say around $\gamma = 0^\circ$) where the cosmic variance is large, you only need a hundred pairs to get pretty close to the cosmic variance. In contrast, at angles (say around $\gamma = 50^\circ$) where the cosmic variance is small, thousands of pulsar pairs are required.

I'd like to illustrate the situation for four different PTAs ²⁾. This is shown in Fig. 8, where we have assumed that there is no timing noise, and no experimental noise of any kind. So this represents the absolute best-case limit of what might be achieved. The expected precision to which those PTAs, with their pulsar sky locations, can find the Hellings and Downs curve is represented by the distance between the “+” symbols. Let's look first at the PPTA, which

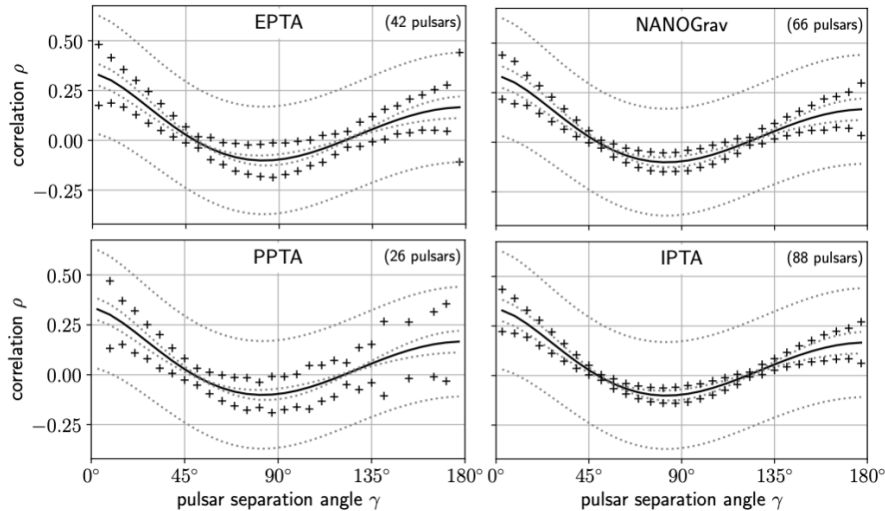


Figure 8: *Best-case variance for current PTAs, using $30 \times 6^\circ$ angular bins, and assuming noise-free measurements, plotted with “+” symbols. This assumes a GW Gaussian ensemble with a binary inspiral spectrum, and plots timing residual correlations with $h^2/h^2 = 0.4$ and $h^2 = 1$, see ²⁾.*

has the smallest number of pulsars, and can form 861 distinct pairs. Since there are 30 angular bins, on the average this is only 29 pairs per bin. Note that some bins, for example the $0^\circ - 6^\circ$ bin, are empty, because the PPTA does not have any pulsar pairs separated by an angle under 6° . So while PPTA does get well below the single-pair variance, shown by the outer dotted lines, it does not get really close to the cosmic variance, shown by the inner dotted lines. In contrast, the IPTA, which has 88 pulsars, can form 3828 distinct pairs, so the average bin contains about 128 pairs. You can see from the crosses that it can, in principle, get much closer to the cosmic variance than the PPTA.

The good news is that these predicted deviations away from the Hellings and Downs curve are not enough to prevent one from recognizing it, and from announcing a confident GW detection. However these deviations are also interesting, because they are a fundamental prediction. If our Universe matches the Hellings and Downs curve much more closely than predicted by the cosmic variance, or if it differs from that curve by much more than the cosmic vari-

ance, then this implies that our Universe is not dominated by many weak GW sources forming a confusion-noise background.

8 Conclusion

I am fairly sure that the Hellings and Downs correlations will be confidently detected in the coming decade, and hope that the organizers of this workshop will invite me to provide updates over that time. In the longer term, as the Square Kilometer Array (SKA) discovers more pulsars and they are timed with greater precision, I am also confident that the cosmic variance will be measured, and will be found to agree with these predictions. If the observed deviations from the Hellings and Downs curve are much smaller or larger than I have predicted, then it means that our Universe does not have a GW background which is described by a Gaussian ensemble, as would be expected from many supermassive black hole binaries, radiating incoherently.

9 Acknowledgments

Parts of this work ²⁾ were done with my friend and colleague Joe Romano.

References

1. B. Allen, arXiv:2205.05637, *to appear in Phys. Rev. D* (2022).
2. B. Allen and J. D. Romano, arXiv:2208.07230 (2022).
3. LIGO/Virgo Scientific Collab., *Phys. Rev. Lett.* **116**, 061102 (2016).
4. S. Chen *et al.*, *Mon. Notices Royal Astron. Soc.* **508**, 4970 (2021).
5. Z. Arzoumanian *et al.*, *Astrophys. J. Lett.* **905**, L34 (2020).
6. B. Goncharov *et al.*, *Astrophys. J. Lett.* **917**, L19 (2021).
7. C. M. F. Mingarelli *et al.*, *Nat. Astron* **1**, 886 (2017).
8. B. Allen and J. D. Romano, *Phys. Rev. D* **59**, 102001 (1999).
9. R. W. Hellings and G. S. Downs, *Astrophys. J. Lett.* **265**, L39 (1983).
10. N. J. Cornish and A. Sesana, *Class Quantum Gravity* **30**, 224005 (2013).