ersità degli Studi di Trieste – Corso di Astrofisica Nucleare e Subnucleare

A Gamma-Ray Pulsar Timing Array Constrains the Nanohertz Gravitational Wave Background

JOURNAL CLUB – 21.04.2023

https://arxiv.org/abs/2204.05226

Elisabeth-A. Ke

Abstract

After large galaxies merge, their central supermassive black holes are expected to form binary systems whose orbital motion generates a gravitational wave background (GWB) at nanohertz frequencies. Searches for this background utilize pulsar timing arrays, which perform long-term monitoring of millisecond pulsars (MSPs) at radio wavelengths. We use 12.5 years of Fermi Large Area Telescope data to form a gamma-ray pulsar timing array. Results from 35 bright gamma-ray pulsars place a 95% credible limit on the GWB characteristic strain of 1.0×10^{-14} at 1 yr^{-1}, which scales as the observing time span $t_{\mathrm{obs}}^{-13/6}$. This direct measurement provides an independent probe of the GWB while offering a check on radio noise models.

What are pulsars?

Pulsars are spinning neutron stars that emit beams of broadband radiation from radio to gamma-ray wavelengths that appear to pulse as they periodically sweep across the line of sight to Earth (*1*). Millisecond pulsars (MSPs) spin at hundreds of hertz and pulse with sufficient regularity to function as celestial clocks distributed across the sky and throughout the Galaxy.

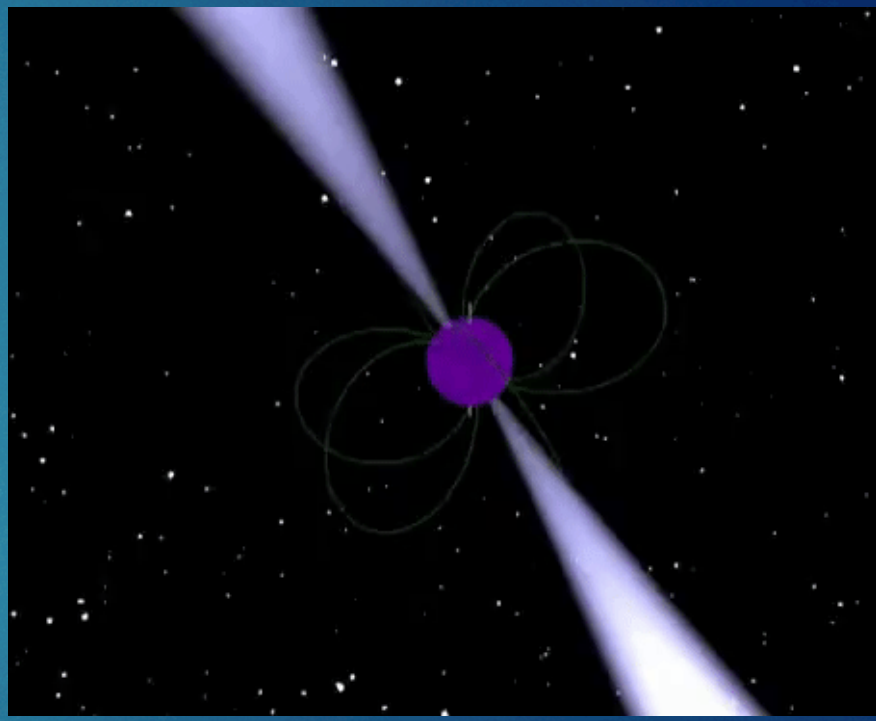

MSPs for GWB

- Long-term monitoring campaigns of ensembles of MSPs are used to search for **low-frequency GWs**, expected from supermassive black hole (SMBH) binaries that are predicted to exist at the **centers of galaxies** that have undergone mergers
- Because of this link between GW frequency and amplitude, the **superposition of GWs from many SMBH binaries** throughout the Universe is predicted to build up a GW background (GWB) with a characteristic GW strain:

$$h_c(f) = A_{\mathrm{gwb}} \left(\frac{f}{\mathrm{yr}^{-1}} \right)^{\alpha}.$$

GWB detection with MSPs

- This GWB can be detected with ensembles of MSPs—known as pulsar timing arrays (PTAs) by monitoring the times of arrival (TOAs) of the steady pulses from each pulsar, which arrive earlier or later than expected due to the spacetime perturbations
- The GWB is expected to be a sum of many individual sources, the induced TOA variations are random and differ for each pulsar, but have a common spectrum of power spectral densities:

$$P(f) = \frac{A_{\rm gwb}^2}{12\pi^2} \left(\frac{f}{\rm yr^{-1}}\right)^{-\Gamma} {\rm yr}^{-3},$$

where Γ=3-2 α, α spectral index

this functional form has more power at low frequencies so is referred to as a red spectrum.

Additionaly, for observations taken at an approximately fixed location (Earth), the GWB is expected to produce a signature quadrupolar pattern of TOA variations, known as the Hellings-Downs correlation.

Constraints on the GWB and Gamma-Ray PTAs

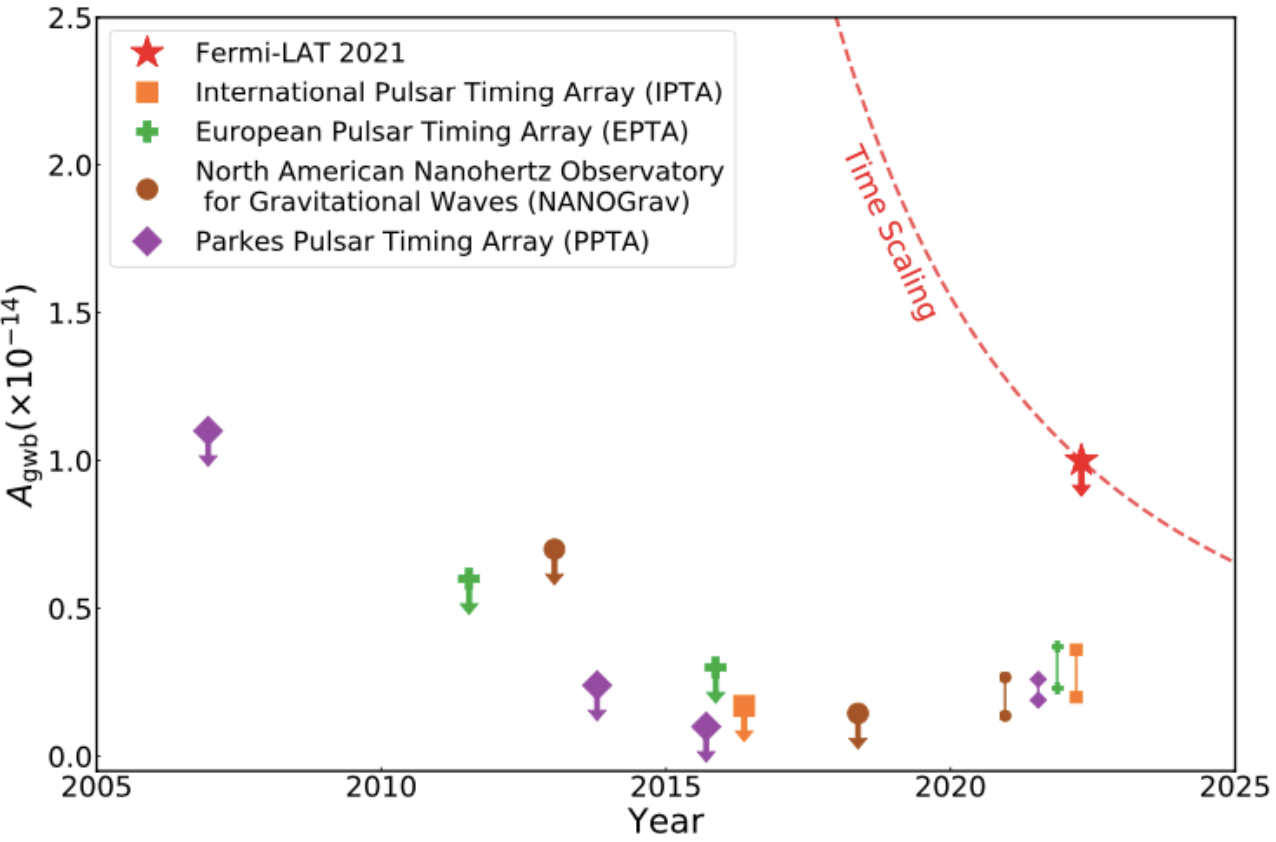

- The inferred constraints on the GWB amplitude at 1 yr-1 A_{GWB} are plotted a a **function of publication date and assume $\alpha = -2/3$**, as predicted for the superposition of GWs from merging BH
- The dashed red line indicates the **expected scaling** as the limit as a function of time

Possible alternative explanations for the signal

- Spin noise
- Frequency-dependent effect of radio propagation through plasma, including solar wind and the IISM. An estimate of the time delay due to dispersion is given by

$$\tau_{\mathrm{DM}} = 4.15\,\mathrm{ms} \times \left(\frac{\mathrm{DM}\,\mathrm{cm}^3}{\mathrm{pc}}\right) \times \left(\frac{\nu}{\mathrm{GHz}}\right)^{-2},$$

which can vary with time, due to relative motions of Earth and the pulsar

Correcting for this effect requires repeated measurements using multi-frequency radio observations and the introduction of many additional degrees of freedom to timing models.

A complimentary approach: Gamma-ray observations

Gamma-ray observations offer a potentially complementary approach: the much higher photon frequency means that the effects of the IISM and solar wind are negligible. The Large Area Telescope (LAT) (*22*), on the Fermi Gamma-ray Space Telescope, is sensitive to GeV gamma-ray photons emitted by MSPs. Its 2.4 steradian field-of-view performs a continuous survey, covering the full sky every 2 orbits (~3 hr). Its GPS clock records photon arrival times with <300 ns precision (*23*), enabling pulsar timing. Analyses of the LAT survey data have detected 127 (*21, 24*) of the over 400 known MSPs in the Milky Way. The large MSP sample, long observing span, and instrumental stability enable a gamma-ray pulsar timing array whose characterization of spin noise and a potential GWB signal is free from IISM effects.

A complimentary approach: Gamma-ray observations

- Observation of 35 gamma-ray MSPs: search of GWB through two different methods

Using the 35 brightest and most stable γ-ray MSPs and 12.5 yr of Fermi-LAT data, we searched for the GWB using two different techniques (*21*). First, we implemented a coherent photon-by-photon analysis which retains $<1\ \mu$s resolution. Second, for analysis with the established software used for radio data analysis, we directly measured TOAs from the LAT data (*25*). Because the TOA estimation procedure requires averaging up to one year of data, this method loses sensitivity to shorter-timescale signals, and only 29 of the 35 pulsars are suitable.

Comparison between A_{GWB} measurements from each pulsar using three analysis methods

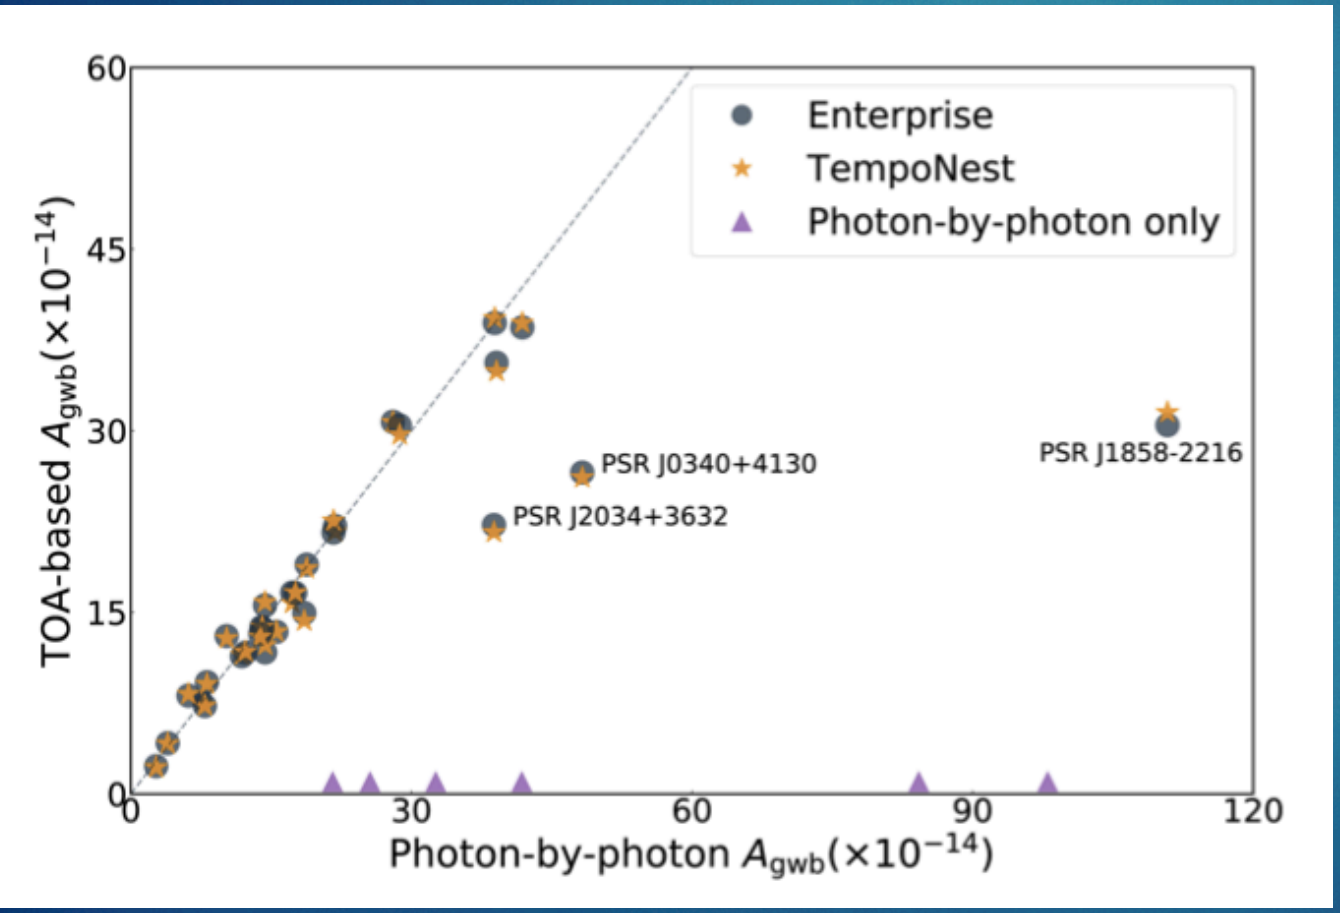

The dashed line indicates equality between the results of the TOA-based and photon-by-photon methods.

Advantages of using Fermi PTA Data

The Fermi PTA data have an essentially constant experimental setup: the data are almost uninterrupted and calibrations have been constant for the full 12.5 year dataset. Gamma-ray data are potentially less subject to astrophysical effects such as changes in the radio pulse shape (*21*). This stability is particularly useful for probing GWs with frequencies below 0.1 yr^{-1}. Such low frequencies are predicted to constrain the spectral shape of the GWB which contains information about the physical sources (*5*).

Gamma-ray constraints on different types of GWB sources

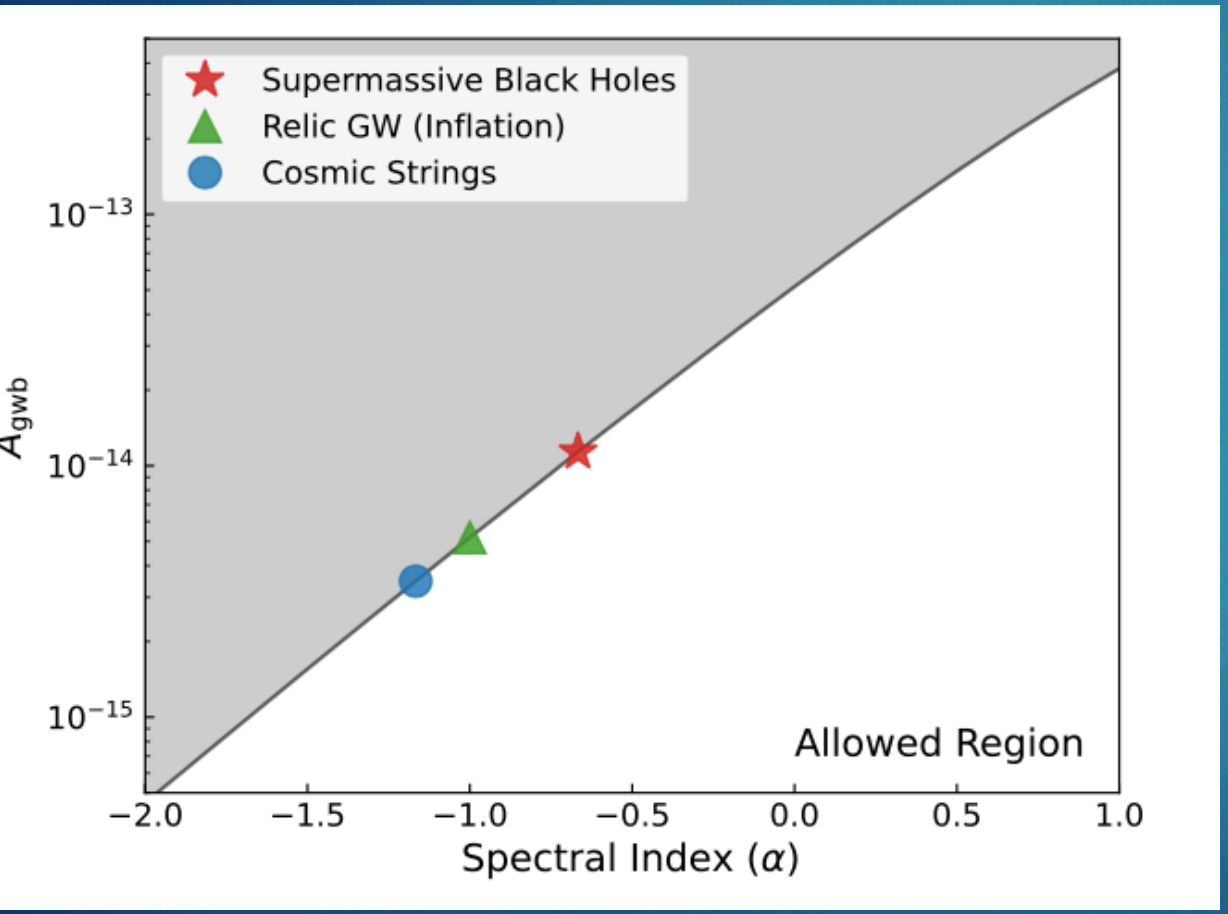

There may be other potential sources of power-law GWBs with different spectral indices α (Γ=3-2 α)

Materials and Methods

Pulsar Timing using Radio and Gamma-ray Observations

The spin phase at time t, $\phi(t)$, increases by 1 each time a pulsar rotates, and because the pulsation mechanism is fixed to the star, $\phi(t)$ can be inferred by observing pulses. Pulsar timing is the measurement of these pulse arrival times (TOAs) and comparison with a timing model that predicts $\phi(t, \lambda)$. The astrophysical parameters λ leave a characteristic imprint on the timing residuals between data and model: a spin frequency error $\delta\nu$ produces linear residuals, $\delta\dot{\nu}$ leaves quadratic residuals, a position error induces an annual sinusoid, etc. (*35*). The goal of pulsar timing is to measure λ and thus characterize a wide range of astrophysical phenomena and constrain fundamental physics (*2, 36–38*).

The Fermi-LAT, on the other hand, collects individual gamma rays. The arrival time of the ith photon, t_i, is recorded with ~300 ns precision, but aside from its energy, a photon carries no additional information: it could be from any pulsar phase ϕ or even from a background source, so these t_i cannot be interpreted as TOAs. In some cases, histograms in ϕ can be built up over a long enough time—hours for bright pulsars, years for the faintest—that a pulse profile emerges.

Pulsar Timing using Radio and Gamma-Ray Observations

For many of the faint MSPs used in this work, there is insufficient integration time to build up a gamma-ray pulse profile and estimate a TOA. For instance, resolving the annual sinusoidal residuals from a position error requires sampling of at least 2 TOAs per year, and preferably faster to avoid a systematic error. Instead, we can use a timing model to evaluate the phase $\phi(t_i, \lambda)$ at each individual photon time and then gauge the agreement of the resulting distribution with an assumed template $f(\phi)$ using the likelihood (Equation S1). Because the LAT has a broad, energy-dependent angular resolution, photons from different sources overlap and we must also account for the background, which we do by computing the probability weight $0 < w_i < 1$ that the ith photon originates from the pulsar (*40–42*). Using a normalized ($\int_0^1 f(\phi)\, d\phi = 1$) pulse profile model, the Poisson likelihood for the data, $\mathcal{L}_{\rm data}$, is

$$\log \mathcal{L}_{\rm data}(\lambda) = \sum_i \log \Big(w_i f(\phi[t_i, \lambda]) + (1 - w_i) \Big). \tag{S1}$$

By maximizing $\mathcal{L}_{\rm data}(\lambda)$, we obtain optimal estimates for parameters λ while preserving the <1 μs resolution of the LAT.

Noise Sources in PTA Data

	Radio		Gamma ray		
Noise Source	Impact	d.o.f.	Impact	d.o.f.	Note
White Noise					
Measurement	moderate	–	major	–	Sensitivity is major limiting factor for gamma rays.
RFI	minor	?	–	–	RFI varies widely between observing systems.
Calibration	minor	?	–	–	Affects certain pulsars/observing systems.
Jitter	moderate	10s	–	–	Jitter affects high signal-to-noise observations.
Red Noise					
DM variation	major	100s	–	–	DM(t) drives radio PTA observing strategies.
Solar wind	moderate	∼10s	–	–	Solar wind mitigation is poorly supported.
Scattering	moderate	100s	–	–	Affects some pulsars/low radio frequencies.
Pulse variability	moderate	0–10s	–	?	No gamma-ray MSP pulse profile changes known.
Discontinuities	moderate	10s	–	–	LAT data are continuous, not a general property.
Spin noise	major	10s–100s	major	10s	Fewer d.o.f. needed for less precise LAT data.

Why use Gamma-ray data?

Implications for gamma-ray analysis In summary, mitigating the noise sources for radio observations requires multi-frequency data, large fractional bandwidths, and homogeneous and regular monitoring, a substantial practical challenge. In contrast, gamma-ray data only require spin noise and Poisson noise models. This eases computational requirements and reduces systematic uncertainty. Due to continual all-sky monitoring, when a new MSP is discovered, archival LAT data can provide a full pulse timing history. The data span for each pulsar is uniform, ensuring that each pulsar is sensitive to the same spectrum of gravitational waves and enabling simple computational approaches.

Photon-by-photon GWB Analysis

Our photon-by-photon GWB analysis is implemented with maximum likelihood techniques. Degrees of freedom for spin noise or the GWB are incorporated into the log likelihood,

$$\log \mathcal{L} = \sum_i \log \left(w_i f(\phi[t_i, \lambda, \beta]) + (1 - w_i) \right) - 0.5 \sum_{kl} \mathcal{C}_{kl}^{-1} \beta_k \beta_l - 0.5 \det \mathcal{C}. \quad \text{(S2)}$$

The additional timing model parameters β are the coefficients of the Fourier transform of a potential noise signal in the data, such that

$$\phi(t, \lambda, \beta) = \phi(t, \lambda) + \sqrt{\frac{2}{t_{\rm obs}}} \nu^{-1} \sum_k \beta_{2k} \cos \left(2\pi k \frac{t}{t_{\rm obs}} \right) + \beta_{2k+1} \sin \left(2\pi k \frac{t}{t_{\rm obs}} \right), \quad \text{(S3)}$$

Timing properties of the 114 MSPs in the parent sample

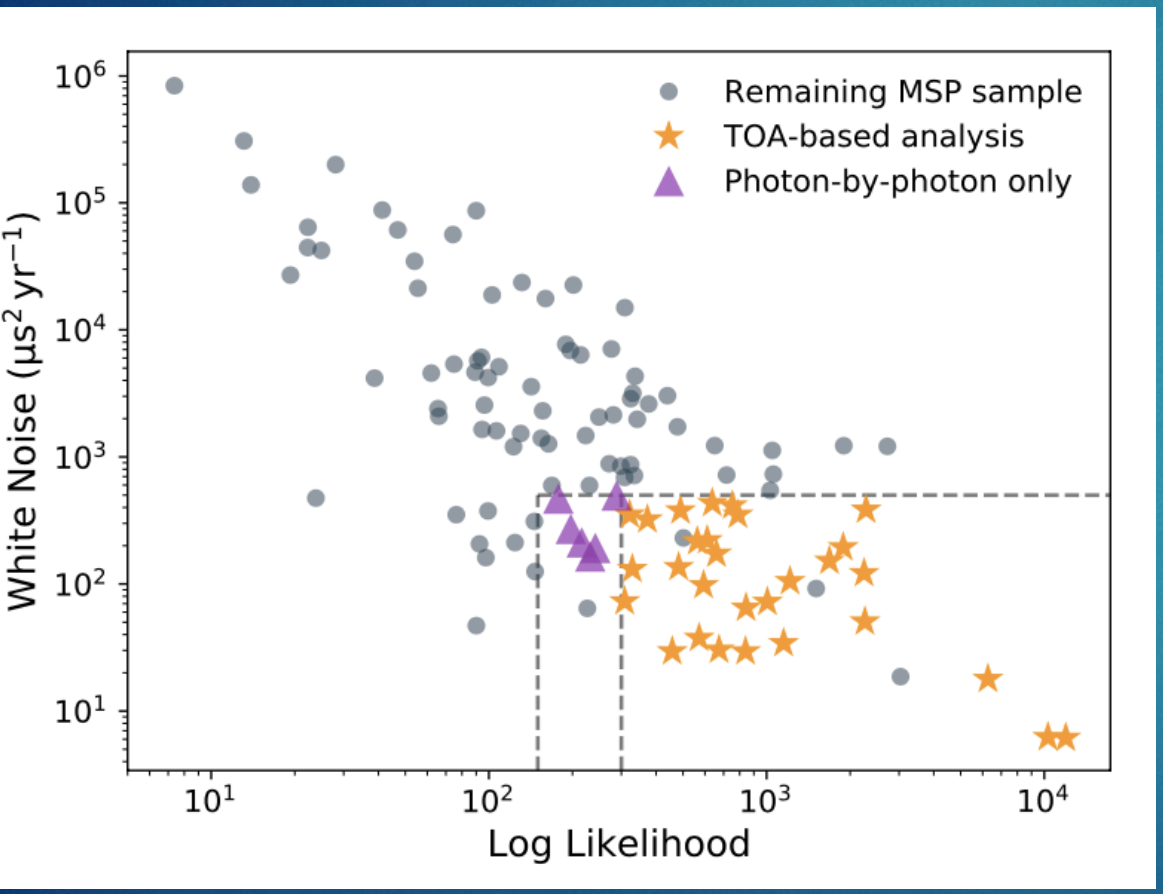

Only some MSPs are suitable for a GWB analysis:

- at a given intensity, a pulsar with a narrower pulse or faster spin frequency has better timing precision
- faint MSPs cannot constrain the GWB bu would increase computational complexity.

The single-pulsar log likelihoods as a function of A_{GWB} produced by the photon- by-photon method

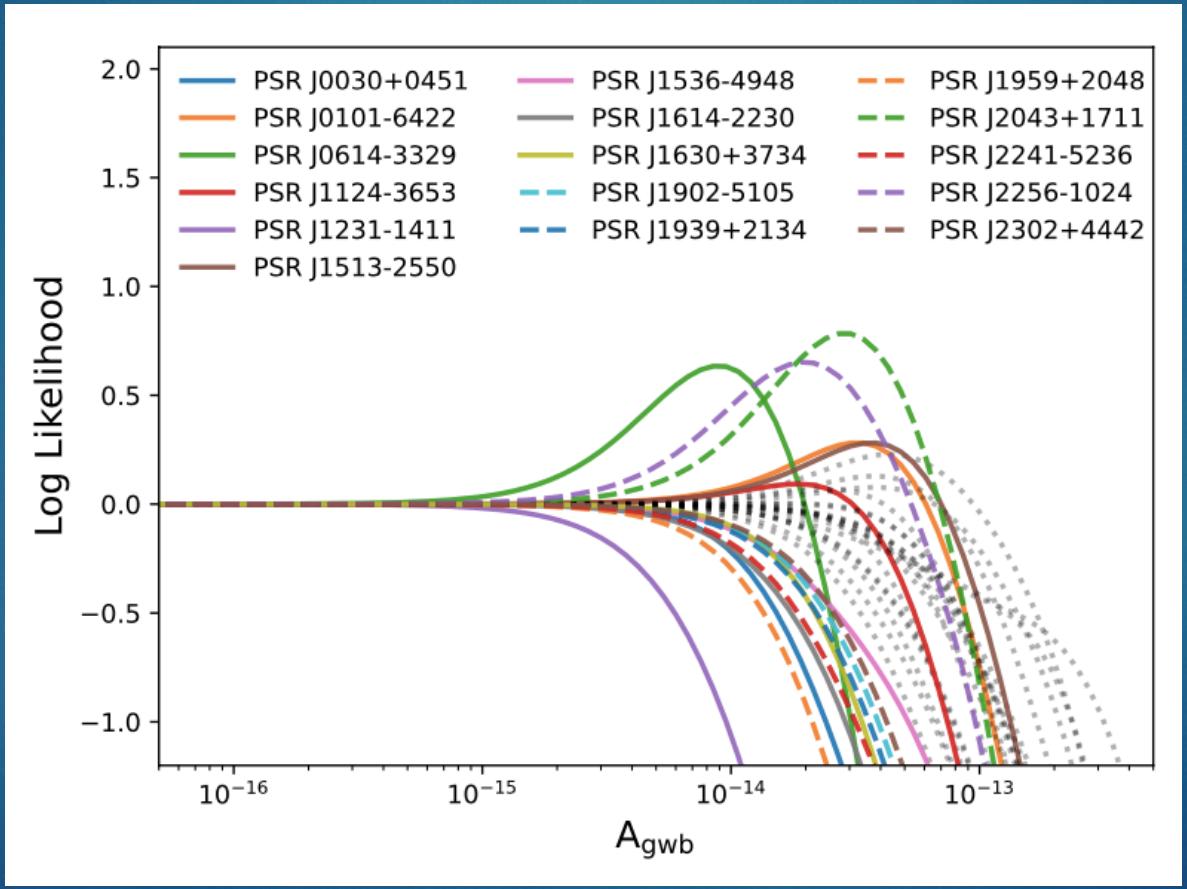

The single-pulsar log likelihoods with an additional, numerically marginalized intrinsic spin noise process

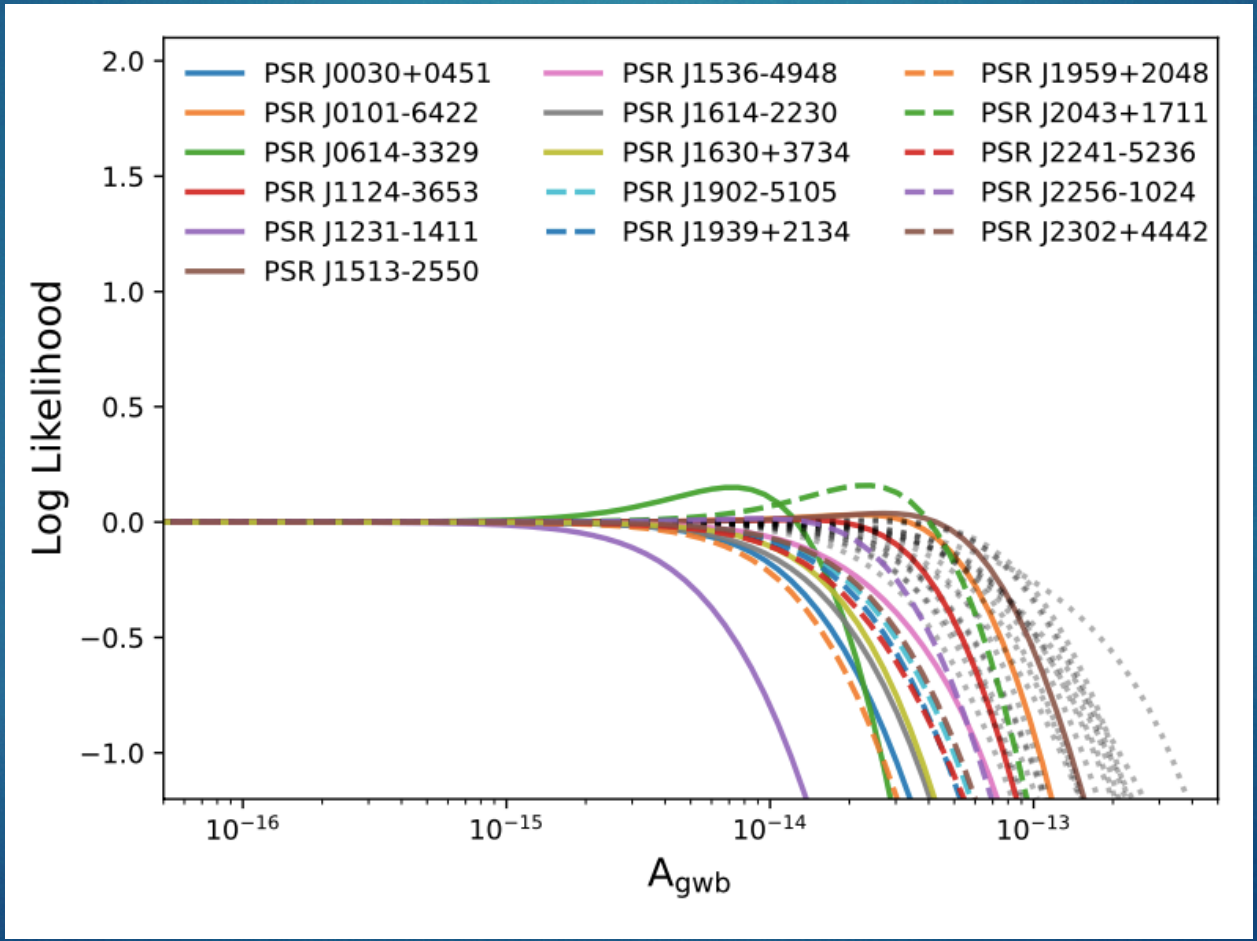

TOA-based GWB analysis

The likelihood based-method described above takes advantage of the time resolution of the LAT data and, by construction, avoids potential systematic errors from reducing the full photon data to TOAs. However, TOA-based methods are computationally efficient, well-tested, and commonly used by radio PTAs, so we have implemented a parallel TOA-based analysis for comparison.

Results

Comparison of methods The two codes ENTERPRISE and TEMPONEST provided single pulsar limits which were consistent with each other, so we compared TOA-based and photon-by-photon approaches. Aside from computational aspects—the TOA-based methods are sampled, while the photon-by-photon method is analytic—there are two primary differences: the photon-based approach avoids the assumption of Gaussianity on TOA uncertainties, and it retains sensitivity to all timescales. We expect the photon-based approach to generally be more precise due to these advantages. The agreement between these two methods provides us with confidence in our photon-by-photon approach. The few exceptions (discussed below) stem from the fundamental differences in the methods, and they have little impact on the final GWB limits because none of these pulsars contribute strongly to the sensitivity of the timing array.

Single pulsar limits on A_{GWB} for 35 pulsars in the sample

Pulsar	Cadence TOA/yr	Noise model (favored)	TN $\times 10^{-14}$	ENT. $\times 10^{-14}$	Photon $\times 10^{-14}$	Photon+RN $\times 10^{-14}$
PSR J0030+0451	2	None	7.54	7.77	7.61	8.44
PSR J0034−0534	2	None	13.40	13.39	15.56	18.00
PSR J0101−6422	2	None	18.63	18.94	18.80	22.16
PSR J0102+4839	2	None	39.29	38.90	38.90	38.76
PSR J0312−0921*	–	–	–	–	21.57	27.86
PSR J0340+4130	2	None	26.13	26.54	48.26	58.21
PSR J0418+6635*	–	–	–	–	32.59	36.41
PSR J0533+6759	1.5	None	21.66	22.14	21.83	26.23
PSR J0613−0200	2	None	22.57	21.59	21.66	25.80
PSR J0614−3329	2	None	4.15	4.20	3.94	4.36
PSR J0740+6620	1.5	None	15.76	16.62	17.24	20.55
PSR J1124−3653	2	None	15.84	15.58	14.34	16.77
PSR J1231−1411	2	None	2.19	2.30	2.71	3.54
PSR J1513−2550*	–	–	–	–	25.57	43.22
PSR J1514−4946	2	None	38.89	38.54	41.85	34.90
PSR J1536−4948	2	None	12.30	11.69	14.36	15.38
PSR J1543−5149*	–	–	–	–	98.02	1356.47
PSR J1614−2230	2	None	9.08	9.23	8.14	9.76
PSR J1625−0021	1.5	None	30.70	30.74	28.01	31.07
PSR J1630+3734	2	None	7.28	7.27	7.91	9.08
PSR J1741+1351*	–	–	–	–	84.24	120.30
PSR J1810+1744	2	None	14.31	14.95	18.54	21.35
PSR J1816+4510	2	None	34.91	35.61	39.09	41.31
PSR J1858−2216	2	None	31.51	30.46	110.83	1417.63
PSR J1902−5105	2	None	11.50	11.38	11.87	15.06
PSR J1908+2105*	–	–	–	–	41.80	47.97
PSR J1939+2134	1.5	None	12.86	13.04	10.24	12.99
PSR J1959+2048	2	WN	8.25	8.12	6.15	7.84
PSR J2017+0603	2	None	16.63	16.60	17.59	20.25
PSR J2034+3632	1.5	None	21.62	22.21	38.82	74.75
PSR J2043+1711	2	None	13.38	13.80	13.96	15.07
PSR J2214+3000	2	None	29.67	30.44	28.74	38.60
PSR J2241−5236	2	WN	13.80	13.66	14.05	16.39
PSR J2256−1024	1.5	None	12.98	13.23	13.84	13.31
PSR J2302+4442	2	None	11.74	11.67	12.26	14.82

- These results use TEMPONEST (TN in column 4), ENTERPRISE (ENT. in column 5) and the photon-by-photon method (columns 6 and 7)
- Pulsars with only photon-by-photon limits are indicated with an asterisk
- Data for PSR J1959+2048 and PSR J2241−5236 favor a model with white noise, while all others favor no additional noise
- Most pulsars can be analyzed with a 2 yr −1 (182 day) cadence, while six pulsars require longer integrations (1.5 yr−1, 243 day) to produce reliable TOAs.

95% credible upper limits on $A_{GWB}/10^{-14}$ from the combined samples

Subset	ENTERPRISE	ENTERPRISE with RN	Photon	Photon with RN
Best 2	1.89	1.94	1.84	2.02
Best 3	1.71	1.74	1.50	1.66
Best 9	1.15	1.19	1.06	1.16
Full 29	1.12	1.04	1.14	1.06
Full 35	–	–	1.14	1.05
Full 29 w/HD	1.06	–	–	–

The pulsars corresponding to each subset are **ranked by their single-pulsar GWB upper limits**. The "Full" rows indicate the total sample for the two methods, 29 pulsars common to TOA-based and photon-by-photon, and 35 to photon-by-photon only

Combined limits and scaling

For both methods, the limits steadily improved with additional pulsars. In the case of the TOA-based approach, the limit also improved with the inclusion of the spatial correlation information predicted for a GWB (*8*). The two full-array limits were nearly identical, while the tightest constraint was $A_{\rm gwb} < 1.0 \times 10^{-14}$.

As discussed above, the limit in the photon-by-photon case can be degraded by including pulsars whose $A_{\rm gwb}$ posteriors peak at $A_{\rm gwb} > 0$, which can happen even in the absence of a GWB signal due to statistical fluctuations. For completeness, however, we considered limits computed when removing two MSPs with the strongest signals, PSR J2043+1711 and PSR J2256−1024 (which exhibits modest orbital period variations). This is justifiable if the log likelihood peak is due to some deficiency in the timing solution. When removing these two pulsars, we obtained limits of 9.8×10^{-15} and 10.3×10^{-15}.

Combined limits and scaling

Table S6: **95% credible upper limits on $A_{\rm gwb}/10^{-14}$ from the combined samples using BAYESEPHEM.** Here, we compare results obtained with ENTERPRISE, in both cases using spatial correlations (*8*) but with and without BAYESEPHEM.

Subset	ENTERPRISE	ENTERPRISE with BAYESEPHEM
Best 2	1.90	1.81
Best 3	1.70	1.64
Best 9	1.17	1.16
Full 29	1.06	1.08

Comparison to radio measurements

Table S8: **Comparison between measured radio PTA spin noise amplitudes and 95% Fermi PTA upper limits.** The spin noise power spectra are given in $s^2\,yr^{-1}$ evaluated at $f = 1/yr$. Fermi limits are computed similarly to GWB amplitude limits except using the value of Γ measured in each row rather than $\Gamma = 13/3$. The last columns give the ratio of the Fermi upper limit to the radio PTA power.

Pulsar	PTA	Γ	P(f) PTA	P(f) Fermi	Ratio
PSR J0030+0451	NANOGrav	6.3	9.0×10^{-18}	8.3×10^{-18}	0.9
PSR J0613−0200	NANOGrav	2.1	1.5×10^{-14}	2.0×10^{-13}	13
PSR J0613−0200	PPTA	4.2	2.5×10^{-16}	5.4×10^{-15}	21
PSR J1939+2134	NANOGrav	3.3	9.8×10^{-15}	6.6×10^{-15}	0.7
PSR J1939+2134	PPTA	5.4	1.8×10^{-16}	1.1×10^{-16}	0.6

The Fermi upper limits for PSR J0030+0451 and PSR J1939+2134 were lower than the measured radio values and statistically incompatible. This could be evidence for uncorrected IISM/solar wind effects leaking into the spin noise estimation in the radio data.