



Universiteit
Leiden
The Netherlands

Gravitational wave detection and data analysis for pulsar timing arrays

Haasteren, R. van

Citation

Haasteren, R. van. (2011, October 11). *Gravitational wave detection and data analysis for pulsar timing arrays*. Retrieved from <https://hdl.handle.net/1887/17917>

Version: Corrected Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/17917>

Note: To cite this publication please use the final published version (if applicable).

4

Placing limits on the stochastic gravitational-wave background using European Pulsar Timing Array data

The most beautiful thing we can experience is the mysterious. It is the source of all true art and all science. He to whom this emotion is a stranger, who can no longer pause to wonder and stand rapt in awe, is as good as dead: his eyes are closed.

Albert Einstein

Abstract

Direct detection of low-frequency gravitational waves ($10^{-9} - 10^{-8}$ Hz) is the main goal of pulsar timing array (PTA) projects. One of the main targets for the PTAs is to measure the stochastic background of gravitational waves (GWB) whose characteristic strain is expected to approximately follow a power-law of the form $h_c(f) = A(f/\text{yr}^{-1})^\alpha$, where f is the gravitational-wave frequency. In this chapter we use the current data from the European PTA to determine an upper limit on the GWB amplitude A as a function of the unknown spectral slope α with a Bayesian inference method, by modelling the GWB as a random Gaussian process. For the case $\alpha = -2/3$, which is expected if the GWB is produced by supermassive black-hole binaries, we obtain a 95% confidence upper limit on A of 6×10^{-15} , which is 1.8 times lower than the 95% confidence GWB limit obtained by the Parkes PTA in 2006. Our approach to the data analysis incorporates the multi-telescope nature of the European PTA and thus can serve as a useful template for future intercontinental PTA collaborations.

This chapter is based on:

Placing limits on the stochastic gravitational-wave background using European Pulsar Timing Array data

R. van Haasteren et al.
MNRAS (2011), 414, 3117

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA

4.1 Introduction

The first direct detection of gravitational waves (GWs) would be of great importance to astrophysics and fundamental physics: it would confirm some key predictions of general relativity, and lay the foundation for observational gravitational-wave astronomy. Pulsar Timing Array (PTA) projects are collaborations which aim to detect low-frequency (10^{-9} — 10^{-8} Hz) extragalactic gravitational waves directly, by using a set of Galactic millisecond pulsars as nearly-perfect Einstein clocks (Foster & Backer, 1990). The basic idea is to exploit the fact that millisecond pulsars create pulse trains of exceptional regularity. GWs perturb space-time between the pulsars and the Earth, and this creates detectable deviations from the strict periodicity in the arrival times of the pulses (TOAs) (Estabrook & Wahlquist, 1975; Sazhin, 1978; Detweiler, 1979).

One of the main astrophysical targets of the PTAs is to measure the stochastic background of gravitational waves (GWB). This GWB is expected to be generated by a large number of black-hole binaries located at the centres of galaxies (Begelman et al., 1980; Phinney, 2001; Jaffe & Backer, 2003; Wyithe & Loeb, 2003; Sesana et al., 2008), by relic gravitational waves (Grishchuk, 2005), or, more speculatively, by oscillating cosmic-string loops (Damour & Vilenkin, 2005; Ölmez et al., 2010).

Currently, there are three independent PTA groups:

(i) the Australian-based programme PPTA, the Parkes Pulsar Timing Array, which uses data from the Parkes telescope (Hobbs et al., 2009; Verbiest et al., 2010), and archival Arecibo data.

(ii) the North-American based programme NANOGrav, North-American Nanohertz Observatory for Gravitational waves, which uses both the Green Bank Telescope (GBT), and the Arecibo radio telescope (Jenet, 2009).

(iii) and the European programme EPTA, European Pulsar Timing Array, which uses five different radio telescopes: the Lovell telescope near Manchester, United Kingdom, the Westerbork Synthesis Radio Telescope (WSRT) in the north of the Netherlands, the Effelsberg Telescope (EFF) near Bonn in Germany, the Nançay Radio Telescope (NRT) near Nançay in France, and the Sardinia Radio Telescope (SRT) in Sardinia, Italy¹.

It is likely that the first detection of GWs by a PTA will occur as a result of a joint effort of all current PTA projects: an International Pulsar Timing Array (IPTA; Hobbs et al., 2010). This will involve the combination of data from several different telescopes, each of them with its own specific hardware elements and software analysis tools. Combining data of different observatories is a challenging task,

¹The SRT is expected to become operational in 2011 (Tofani et al., 2008)

which requires extra care when dealing with the high quality data of modern observatories (Janssen, 2009).

In this chapter, we present a methodology on how to combine the data from several radio telescopes and use it in an optimal way to obtain the information on extragalactic gravitational waves. We use the data from three different radio telescopes located on the European continent, to place a new upper limit on the amplitude of the GWB. As part of our analysis, we obtain detailed information about the statistical properties of the individual pulse time series.

The calculation of upper limits on the GWB, based on pulsar timing, go as far back as the early 1990's (Stinebring et al., 1990; Kaspi et al., 1994; McHugh et al., 1996; Lommen, 2002). These analyses have been based on high quality datasets for single millisecond pulsars. The most stringent upper limits have been obtained recently by Jenet et al. (2006), who have used PPTA data and archival Arecibo data for several millisecond pulsars. Our dataset is different from that used by Jenet et al. (2006) since it includes only the pulse times of arrival measured by the EPTA telescopes, even though some of the pulsars are being timed by multiple PTA groups. The Bayesian inference method we use to obtain an upper limit on the GWB is also different from the methods used by all of the previous studies. Its potential advantages include the use of cross correlations between TOAs of different pulsars, and the simultaneous constraint on both the amplitude and spectral index of the GWB.

The outline of the chapter is as follows. In Section 4.2 we give a brief general overview of pulsar timing observations. In Section 4.3 we detail the observations from all of the EPTA telescopes which were used for our analysis. In Section 4.4 we compare the different data analysis methods that have thus far been developed for PTAs, and we motivate our choice for the Bayesian inference for placing a limit on the GWB. We outline the data analysis procedure in Section 4.5, after which, in Section 4.6, we present the upper limits on the amplitude of the GWB, and also the spectral analysis of the individual pulsar noises. Finally, in Section 4.7 we discuss the astrophysical implications of our results, and provide an outlook for the future.

4.2 EPTA data analysis

In this section we present a brief overview of the observations, instrumentation and data analysis used at the different EPTA observatories for transforming a series of measured pulses to a TOA.

The complete data reduction process that converts the incoming data stream from a radio telescope into one single TOA per observation, called “the pipeline”, is optimised by hand with much care and is observatory specific. The process can

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA

be described in five general steps, shown in Figure 4.1:

- 1) The incoming radio waves are received by the telescope.
- 2) The signal is converted from analog to digital, at a Nyquist sampled rate.
- 3) Data is (coherently) de-dispersed and, if possible, Stokes parameters are formed.
- 4) The de-dispersed timeseries are folded at the pulsar period, resulting in averaged pulse profiles. Typically a timespan containing several 10^5 pulses is used for each TOA.
- 5) A cross-correlation with a template pulse profile yields a TOA and associated uncertainty (Taylor, 1992).

Individual pulse amplitudes and pulse shapes are highly irregular, and pulse phases vary significantly from pulse to pulse (Cordes & Shannon, 2010). Therefore careful averaging (*folding*) has to be performed to obtain a single TOA. Furthermore, the interstellar medium (ISM) results in significant delays of the arrival time of the pulses over the receiver bandwidth. As a large bandwidth is required to reliably detect a pulse, accounting for the ISM is key for precision timing.

Differences in templates used, e.g. the use of integrated profiles versus analytic templates, all based on single-observatory data, and the difference in definition of the reference point in a template will result in offsets between data sets generated by different observatories. All extra offsets in our data will lead to information loss of other signals like the GWB. Therefore, using a common template for each pulsar at all observatories is desirable, and will be implemented in the near future.

The realisation of the five steps and therefore their output (the resulting TOA) might differ among observatories. Understanding and accounting for those differences is essential for the correct analysis and optimal combining of the EPTA data. A more detailed study on this subject is in preparation (Janssen et al. 2011).

The cross-correlation between the folded profile and the template yields an uncertainty of the TOA (Taylor, 1992). One would like this uncertainty to be solely due to the radiometer noise, i.e. the noise intrinsic to the measurement, but in practice the errors sometimes appear to have been systematically over- or underestimated. It is a common practice, which we follow here, to allow for an extra parameter to multiply these uncertainties for each pulsar-observatory-backend combination (Hobbs & Edwards, 2006). This extra multiplicative factor allows the TOA uncertainties to statistically account for the TOA scatter: the deviations of the strict periodicity of the pulses. This is clearly unsatisfactory, and in future timing experiments the origin of the predicted and measured TOA scatter will have to be thoroughly investigated.

4.2. EPTA DATA ANALYSIS




Figure 4.1: The processing pipeline for pulsar timing, step by step

	Telescope	WSRT	NRT	EFF
Equivalent dish size (m)		93.5	94.4	100
Centre observing frequencies (MHz)		1380	1398, 2048	1400
Observing bandwidth (MHz)		80	64/128	28-112
Obs. time per month per pulsar		1x30 min	4-6x60min	1x30 min
Pulsar backend		PuMaI	BON	EBPP
Dedispersion		incoherent	coherent	coherent
Used templates		integrated profiles	integrated profiles	analytic

Table 4.1: Details of the different EPTA observatories relevant for this chapter. The NRT observing bandwidth has doubled to 128 MHz in July 2009.

4.3 EPTA observations

4.3.1 Overview of the observatories

We have used pulsar timing observations of five radio pulsars, observed with three of the EPTA telescopes, to set a limit on the GWB. See Table 4.1, Fig. 4.2 and Appendix B of this chapter for an overview of the data sets used and the properties of each telescope. Each pulsar was observed on average once every month for 30 minutes at each telescope. Although additional observing frequencies are commonly used at WSRT and EFF, their respective 1380 and 1400 MHz observing bands have the best sensitivity and result in the highest precision TOAs. Therefore we have only used observations taken at those frequencies at WSRT and EFF for the analysis presented in this chapter. The data were either coherently de-dispersed (NRT and EFF) or incoherently de-dispersed (WSRT). The observations were folded and cross-correlated with an analytic template (EFF), or a high S/N, observatory specific, template (WSRT & NRT), to calculate one time-of-arrival (TOA) per observation. See e.g. Lazaridis et al. (2009) for a more complete description of the observing procedures and data analysis at the different observatories.

As discussed, any change to the pipeline or to the input of the pipeline can result in a difference in the calculated TOAs. We emphasise that it is essential to correctly identify these systematic effects and include them in the modelling of the TOAs. In our analysis, we have done this by introducing jumps between TOAs of the same pulsar anywhere the pipeline differs in some way.

Once the complete set of data for each pulsar is obtained, and corrected for global drifts by comparing to UTC, it is fit with the timing model. The timing model is a multi-parameter fit that represents our best knowledge of the many deterministic processes that influence the values of the TOAs. The timing residuals are then produced by subtracting the timing model, which is subsequently optimised by minimising these residuals through a least-squares fit. This was done using the pulsar timing package TEMPO2 (Hobbs et al., 2006).

4.3.2 Selection of data sets

The European observatories have been timing millisecond pulsars for many years, and potentially all of that data could be used in the calculation of an upper limit on the GWB. However, like Jenet et al. (2006) we choose to use only the data from the pulsars which perform best as ideal clocks, e.g. those with the highest precision TOAs and the most straight-forward noise characteristics.

TOA precision is not the only factor that determines the sensitivity to the GWB; other factors like the total timing baseline and the number of observations (i.e.

4.4. OVERVIEW OF DATA ANALYSIS METHODS

TOAs) affect this sensitivity as well. A great advantage of the EPTA data is that several pulsars have been monitored for a relatively long time: over 10 years. To determine which timing residuals (i.e. pulsar-observatory combinations) are most useful for GWB detection, we analyse each dataset separately. By doing this we can determine the sensitivity to the GWB of a set of TOAs: the lower the $3\text{-}\sigma$ upper limit $h_c^{\text{max}}(1\text{yr})$ we get using only a particular set of TOAs, the more sensitive that set of TOAs is to the GWB.

The timing residuals of the selected pulsars are shown in Figure 4.2. These five pulsars significantly outperform the other pulsars being timed by the EPTA in terms of how well they can limit the GWB amplitude: these five pulsars can each individually limit the GWB well below $h_c(1\text{yr}) = 10^{-13}$ for $\alpha = -2/3$, whereas other current EPTA datasets typically perform worse by a factor of several. Since there is such a difference between this set of five pulsars, and the other pulsars that have been observed by the EPTA, we do not expect to gain any significant sensitivity by including more pulsars that cannot meet this constraint. We therefore choose $h_c^{\text{max}}(1\text{yr}) \leq 10^{-13}$ with $\alpha = -2/3$ as a constraint for including a dataset in our calculation.

In addition to this constraint, we also demand that datasets that just barely satisfy $h_c^{\text{max}}(1\text{yr}) \leq 10^{-13}$ do not show prominent low-frequency (“red”) timing noise. Our criterion for presence of the latter is a peak in the posterior distribution which is inconsistent with zero amplitude for $\alpha \leq 0$.

4.4 Overview of data analysis methods

Over the years, different data analysis methods for PTAs have been developed. In this section we briefly discuss the differences between these methods, covering the merits and the weaknesses of each, and we motivate our choice for the Bayesian inference method.

Currently developed data analysis methods for isotropic GWB studies with PTAs can be divided in two classes:

- 1) methods designed to set limits on the GWB characteristic strain.
- 2) methods designed to detect the GWB signal, and measure its parameters.

In sections 4.4.1 and 4.4.2 we describe the differences between the two approaches. Among methods that belong to the first category are those of Stinebring et al. (1990); Kaspi et al. (1994); McHugh et al. (1996); Lommen (2002); Jenet et al. (2006). Methods that are primarily focused on detecting the GWB include those of Jenet et al. (2005); Anholm et al. (2009); van Haasteren et al. (2009); Yardley et al. (2011); see also chapter 2 of this thesis. As presented in section 4.6 of this chapter, the present-day data of the EPTA, is not sufficient to make a detection of the GWB,

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA




Figure 4.2: The timing residuals of all the pulsars used in the GWB limit calculation. The time in MJD is shown on the x-axis. On the left of the dash-dotted line we have placed a sample residual with an uncertainty of $1 \mu\text{s}$.

4.4. OVERVIEW OF DATA ANALYSIS METHODS

and this is likely to hold for the data of other PTAs as well. The primary goal of the data analysis method presented in this chapter is therefore to place a limit on the GWB.

Simultaneously with the characterisation of the GWB signal, the PTA data analysis method should also provide information about the timing model and the timing noise. Because of the quality of the PTA data, the parameters of the timing model and the timing noise are best determined from the same data set that we are analysing to obtain information about the GWB signal. We discuss how this is done in the different PTA data analysis methods in section 4.4.3.

4.4.1 Detecting the GWB

As described in chapter 2, the premise of GWB detection is that the GWB induces well-defined correlations between the timing residuals of different pulsars. The correlations between timing residuals of two pulsars induced by the GWB depend only on the separation angle between those two pulsars. This correlation curve is known as the Hellings & Downs curve (see Figure 1.3). Similar to the GWB signal which is expected to be very “red” (i.e. low frequency), the unknown timing noise for each pulsar is expected to have a very red component. Timing noise is poorly understood, and currently we only have empirical models. Therefore, we cannot reliably distinguish between the GWB or timing noise using data of one pulsar. The only unique characteristic of the GWB is apparent when considering data of many pulsars: a signal with identical statistical properties in the data of all pulsars, correlated according to the Hellings & Downs curve.

All current PTA GWB detection schemes rely on the unique feature of the GWB signal that it induces correlations between timing residuals of all pulsars. Different approaches to the detection problem include:

- 1) Measuring the correlations between all pulsar pairs. If GWB signal is present in the data, the correlations should match the Hellings & Downs curve (Jenet et al., 2005). Also, by considering the cross-power spectrum, the amplitude of the GWB can be determined, given a value for the spectral index of the power-law GWB signal (Demorest, 2007; Yardley et al., 2011).
- 2) Searching for a signal in the data of all pulsars that is correlated according to the Hellings & Downs curve. Such a search takes the correlations induced by the GWB as a given, and measures the parameters of the signal correlated according to the Hellings & Downs curve (Anholm et al., 2009; van Haasteren et al., 2009, see also chapter 2 of this thesis).

In approach (1), one is primarily interested in the correlations that are induced by an unknown signal between all pulsars, whereas in approach (2) one is primarily interested in the parameters of the GWB signal. When attempting to present a re-

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA

liable GWB detection, an analysis that takes the first approach should yield results of the GWB-induced correlations between all pulsar pairs that are accurate enough to exclude the possibility that these correlations are induced by any other source. An analysis that takes the second approach should include other correlated signals in the analysis, such that a reliable distinction can be made between a GWB signal that is correlated according to the Hellings & Downs curve, and other signals that are not correlated according to the Hellings & Downs curve. Because methods that take the first approach rely on cross-power measurements of pulsar pairs, these methods are not suited for placing limits on the GWB with data of only a few pulsars.

4.4.2 Placing limits

Methods that aim to place a limit on the GWB amplitude all try to answer the same question: assuming there is a GWB-induced signal present in the observations, what is the largest amplitude this signal can have that is still consistent with the data. Again, there are two approaches to this problem:

1) The incoherent analysis, where one solely focuses on the power spectrum of the timing residuals of the pulsars, neglecting the GWB-induced correlations between pulsar pairs (Stinebring et al., 1990; Kaspi et al., 1994; McHugh et al., 1996; Lommen, 2002; Jenet et al., 2006; Demorest, 2007).

2) The coherent analysis, where one not only includes the certain power spectrum of the GWB-induced signal in the timing residuals, but also the correlations between residuals of pulsar pairs. This is essentially what is done in this chapter. We note that it is also possible to place a coherent GWB-limit by only considering the cross-power correlations between pulsar pairs (Demorest, 2007; Yardley et al., 2011). However, in that case, the power-spectrum information of the single-pulsar time series is neglected, which makes these methods suboptimal.

In the case where only observations of a single pulsar are available, the incoherent analysis should be as sensitive as the coherent analysis. However, when observations of multiple pulsars is available, the most stringent limits are produced when doing a coherent analysis.

The difference between a coherent limit-placing analysis, and a GWB detection scheme, is that it is not necessary in the limit-placing analysis to convincingly show that the GWB signal is correlated according to the Hellings & Downs curve. The Hellings & Downs correlations are then assumed, and used to place a stringent limit. This is what is done in this chapter: in our model we have not included sources that are similar to the GWB, but correlated with a different angular dependence than the GWB. To this date, the Bayesian inference presented in this chapter is the only published method to coherently analyse PTA data with the aim to place

4.5. BAYESIAN PTA DATA ANALYSIS

upper limits on the GWB without loss of information.

4.4.3 The timing model and timing noise

Both the timing model and the timing noise parameters should be determined from the same data as what is used to infer the parameters of the GWB. This has severe implications for the sensitivity to the GWB, since the timing model (especially the pulse period and pulse period derivative) and the timing noise parameters are covariant with the GWB. The way this is dealt with is different in Frequentist and Bayesian inference techniques.

Frequentist methods (e.g. Jenet et al., 2005; Yardley et al., 2011) to measure or place limits on the GWB parameters rely on estimates of the timing model and timing noise parameters. Typically, the best estimators for the timing model and timing noise parameters enter in the detection statistic for the GWB, and the uncertainty in the GWB parameters is determined by Monte Carlo simulations.

Bayesian methods like those presented in this chapter rely on marginalisation to deal with covariances. Instead of separately producing best estimators for the timing noise, the timing model, and the GWB parameters, all these contributions are present in the Bayesian posterior distribution. All parameters but those of the GWB are then marginalised over.

In Bayesian methods, the marginalisation process is computationally quite expensive. But once correctly carried out, the resulting marginalised posterior gives a very clear presentation of the analysis with all covariances correctly taken into account. In current frequentist methods, the point estimates for the GWB and the timing noise are relatively easily produced. But the Monte Carlo simulations required to take into account the covariances between the timing model, the timing noise, and the GWB parameters require a lot of computational resources, and optimality has yet to be proved (Yardley et al., 2011).

4.5 Bayesian PTA data analysis

The analysis presented in this chapter broadly follows the procedure introduced in van Haasteren et al. (2009, see also chapter 2). The Bayesian inference method of chapter 2 relies on creating the parametrised models of the timing residuals, and forming a probability distribution function (PDF) as a function of the model parameters. All known systematic contributions of known functional form should be included in the model. In the examples used in chapter 2 the model for the systematic errors included only the quadratic contribution to the TOAs from pulsar spindowns. The multi-telescope nature of the EPTA requires more complete mod-

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA

els for timing residuals than the one used in chapter 2. In this section we show how to build and implement these models in practice.

We first briefly review the inference method of chapter 2 in Section 4.5.1 and 4.5.2. We then present the extended model we use for the analysis of the TOAs in Section 4.5.3, after which we show how we handle TOAs coming from different observatories in Section 4.5.4.

4.5.1 Brief review of the Bayesian inference method

The set of TOAs from all pulsars forms the basic input used in the Bayesian data analysis. Many processes influence the measured TOA values; in this chapter we discriminate between deterministic processes, like quadratic spindown, and stochastic processes, like timing noise:

$$t_{(ai)}^{\text{obs}} = t_{(ai)}^{\text{det}} + \delta t_{(ai)}^{\text{stoch}}, \quad (4.1)$$

where $t_{(ai)}^{\text{obs}}$ represents the i -th TOA of pulsar a , $t_{(ai)}^{\text{det}}$ is the corresponding contribution to the TOA solely due to deterministic processes, and $\delta t_{(ai)}^{\text{stoch}}$ is the contribution due to stochastic processes.

The effects of deterministic processes are described by the set of model parameters $\vec{\eta}$: $t_{(ai)}^{\text{det}} = t_{(ai)}^{\text{det}}(\vec{\eta})$. As is done in chapter 2, we assume that the stochastic processes are Gaussian, though their spectra are not necessarily white. In such a model, the stochastic processes can be represented by the correlation matrix

$$\langle \delta t_{(ai)}^{\text{stoch}} \delta t_{(bj)}^{\text{stoch}} \rangle = C_{(ai)(bj)} = C_{(ai)(bj)}(\vec{\xi}), \quad (4.2)$$

where $\vec{\xi}$ are the model parameters.

The key distribution used in a Bayesian analysis is the likelihood function, the probability distribution of the data for a given model and its parameters. As described in chapter 2, for PTAs the likelihood takes the following form:

$$L(\vec{\theta}) = P(\vec{\delta t} | \vec{\theta}) = \frac{1}{\sqrt{(2\pi)^n \det C}} \exp \left[-\frac{1}{2} \sum_{(ai)(bj)} (\vec{t}_{(ai)}^{\text{obs}} - \vec{t}_{(ai)}^{\text{fit}}) C_{(ai)(bj)}^{-1} (\vec{t}_{(bj)}^{\text{obs}} - \vec{t}_{(bj)}^{\text{fit}}) \right], \quad (4.3)$$

where $\vec{\theta} = (\vec{\eta}, \vec{\xi})$, and $\vec{\delta t}$ is the difference between the observed TOAs, and the fitted TOAs. A Bayesian analysis assigns prior distributions $P_0(\vec{\theta})$ to the model parameters, and explores the parameter space of the posterior distribution (short-

handed simply as *the posterior*): $P(\vec{\theta} | \vec{\delta t}) = L(\vec{\theta})P_0(\vec{\theta})$.

4.5.2 Obtaining a marginalised posterior distribution

The posterior $P(\vec{\theta} | \vec{\delta t})$ contains information about all model parameters. We need to express the posterior as a function of only those parameters that represent the GWB. This process is called marginalisation, and consists of integrating over all other parameters. The resulting marginalised posterior is the posterior probability density of the GWB parameters.

Marginalisation of a posterior in a high-dimensional parameter space is non-trivial, and a direct numerical integration is prohibitively computationally expensive. As in chapter 2, we employ a mix of analytic integration and Markov Chain Monte Carlo (MCMC) methods to accomplish this. The marginalisation remains the computational bottleneck for the method's effectiveness, as the computational time scales with n^3 , with n the total number of TOAs to be analysed.

A computational shortcut can be used by analytically marginalising over the parameters of the timing model. As shown in chapter 2, this is possible provided that the parameters represent signals of known functional form. This condition is equivalent to the requirement that the timing residuals generated by the timing model are linear with respect to its parameters: $\delta t = d(\alpha - \hat{\alpha})$, where δt is the timing residual, d is a proportionality constant, $\hat{\alpha}$ is the best fit value for the model parameter, and α is the model parameter. While this is always true for quadratic spindown as considered explicitly in chapter 2, it is generally not true for other timing model parameters. However, when the deviations of the timing model parameters from their best-fit values are small, it is a good approximation that the residuals generated by the timing model are linear with respect to the deviations from their best-fit values: $\delta t \approx d(\alpha - \hat{\alpha})$.

Analytically marginalising over the timing model is therefore possible, and by doing so the number of parameters that must be integrated over numerically by the use of MCMC is reduced greatly. Dependent on the model we use to describe the statistics of the timing residuals, the number of parameters left to explore is then just several per pulsar/backend combination. The results of the analysis can be presented as a marginalised posterior as a function of any parameter in the model, provided that this parameter was present in the MCMC run.

4.5.3 Used model for the TOAs

We divide the actual parameterisation in 3 parts:

- a) The deterministic timing model.
- b) The gravitational-wave background.

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA

c) Other stochastic processes (e.g., timing noise).

In this section we discuss how we have taken these into account in our data analysis.

As a first step, the TOAs are processed using the software package TEMPO2, in order to determine the best-fit timing model. This procedure consists of the following steps:

1. TEMPO2 requires an initial guess α_{0i} for the timing model parameters α_i in order to find timing residuals (pre-fit timing residuals).
2. It then constructs an approximation to the timing model, in which the timing residuals depend linearly on $\alpha_i - \alpha_{0i}$.
3. It finds the best-fit α_i within this linear approximation, and uses those values to update the timing residuals using the full non linear timing model (post-fit timing residuals).
4. The newly obtained parameters and corresponding timing residuals are then judged by the person performing the model fitting, and if determined necessary the newly obtained parameters can act as the initial guess for a new fitting iteration. TEMPO2 also allows adjustment and fitting of α_i one by one.

Finding the timing solution with TEMPO2 is not fully algorithmic, but typically requires someone experienced with pulsar timing analysis, who approaches the TOAs fitting in several different ways, which ensures that phase coherence is maintained and that the relevant deterministic model parameters are included properly. Though this strategy works well in practice, we should remain conscious of the possibility that different solutions might be obtained by different observers, who may also choose to include additional model parameters.² In Appendix B of this chapter we present the timing solutions we found for the analysed pulsars. These are the values we used as our initial guess, α_{0i} . Note that these α_{0i} and their uncertainties, although created with TEMPO2 using the same datasets that we base our upper limit on, do not include our model for the red noise. The values and uncertainties we list in Appendix B of this chapter therefore do not represent our best estimates if we were to take into account the red timing noise. Although calculating these best estimates of α_i is reasonably straightforward, these estimates are not accessible in our MCMC because we have marginalised over these parameters analytically. The calculated upper limit on the GWB, however, does include all these effects, and therefore automatically incorporates the removal of power from the low-frequency GW signal by fitting for the timing model parameters and jumps.

In the above mentioned step 2 where the timing model is linearised, we have made an important simplification that we now describe in more detail. Since we

²Qualitatively, experienced observers are rightfully so very confident in their timing solutions. Quantitatively however, the only statistical tool currently available for observers to check whether the timing solution is reasonable is the reduced χ^2 statistic. But since the error bars obtained with the cross-correlation technique cannot be fully trusted, the same holds for the χ^2 statistic.

4.5. BAYESIAN PTA DATA ANALYSIS

take into account, and marginalise over, all timing model parameters present in our method, we are effectively working with the TOAs instead of just the timing residuals. However, the timing model has been linearised by TEMPO2 with respect to $\alpha_i - \alpha_{0i}$. This implies that we need to be sufficiently close to α_{0i} in the parameter space for this approximation to be valid, which means that the timing residuals derived with TEMPO2 need to be approved by the person fitting the data, before using these as inputs in the Bayesian inference method.

The stochastic component contributing to the TOAs is characterised as follows. Firstly, general relativity describes how the timing residuals of a pair of pulsars are correlated due to gravitational waves:

$$\zeta_{ab} = \frac{3}{2} \frac{1 - \cos \theta_{ab}}{2} \ln \left(\frac{1 - \cos \theta_{ab}}{2} \right) - \frac{1}{4} \frac{1 - \cos \theta_{ab}}{2} + \frac{1}{2} + \frac{1}{2} \delta_{ab}, \quad (4.4)$$

where θ_{ab} is the angle between pulsar a and pulsar b (Hellings & Downs, 1983). The GWB spectrum is parametrised as a power-law of the form (Maggiore, 2000; Phinney, 2001; Jaffe & Backer, 2003; Wyithe & Loeb, 2003; Sesana et al., 2008):

$$h_c = A \left(\frac{f}{\text{yr}^{-1}} \right)^\alpha, \quad (4.5)$$

where h_c is the characteristic strain as used in Jenet et al. (2006), A is the amplitude of the signal, and α is the spectral index. This then results in a correlation matrix for the GWB (see chapter 2):

$$C_{(ai)(bj)}^{\text{GW}} = \frac{-A^2 \zeta_{ab}}{(2\pi)^2 f_L^{2-2\alpha}} \left\{ \Gamma(-2 + 2\alpha) \cos(\pi\alpha) (f_L \tau)^{2-2\alpha} \sum_{n=0}^{\infty} (-1)^n \frac{(f_L \tau)^{2n}}{(2n)! (2n + 2\alpha - 2)} \right\}, \quad (4.6)$$

where, as in chapter 2, $\tau = 2\pi|t_i - t_j|$, and f_L is a cut-off frequency, set much lower than the lowest GW frequency we are sensitive to.

Secondly, the stochastic timing noise for each individual pulsar is split into three components:

1) Individual errors of TOA determination from the cross-correlation, represented by the TOA error bars. An extra free parameter, called the EFAC value, is commonly introduced by pulsar observers in order to account for possible miscalibration of the radiometer noise (Hobbs & Edwards, 2006); this parameter is a multiplier for all of the TOA error bars for a given pulsar.

2) An extra white noise component, independent of the error bars. This basi-

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA

cally acts as extra non-time-dependent noise, and the parameter is often called an EQUAD parameter.

3) Red noise, consisting of a power-law spectrum in the timing residuals. This component allows for structure in the timing residuals.

All three timing noise components are uncorrelated between the pulsars.

The resulting correlation matrices from components 1, 2, and 3, as derived in chapter 2, are given by:

$$\begin{aligned}
 C_{(ai)(bj)}^{\text{err}} &= E_a^2 \Delta t_{(ai)}^2 \delta_{ab} \delta_{ij} \\
 C_{(ai)(bj)}^{\text{WN}} &= N_a^2 \delta_{ab} \delta_{ij} \\
 C_{(ai)(bj)}^{\text{RN}} &= \frac{-R_a^2 \delta_{ab}}{(2\pi)^2 f_L^{2-2\alpha_a}} \left\{ \Gamma(-2 + 2\alpha_a) \cos(\pi\alpha_a) (f_L \tau)^{2-2\alpha_a} \right. \\
 &\quad \left. \sum_{n=0}^{\infty} (-1)^n \frac{(f_L \tau)^{2n}}{(2n)! (2n + 2\alpha - 2)} \right\}, \quad (4.7)
 \end{aligned}$$

where $C_{(ai)(bj)}^{\text{err}}$, $C_{(ai)(bj)}^{\text{WN}}$, and $C_{(ai)(bj)}^{\text{RN}}$ are the correlation matrices corresponding to the error bars, the extra white noise, and the red noise respectively, with a and b denoting the pulsar number, i and j denote the observation number, Δt is the TOA uncertainty (the error bar) as calculated in the pipeline, E_a is the scaling parameter of the error bars for the a 'th pulsar (the EFAC factor), N_a is the amplitude of the white noise, R_a is the amplitude of the red timing noise, α_a is the spectral index of the red noise spectrum of pulsar a , and τ is the time difference between two observations.

4.5.4 Combining datasets

The previous section gives a complete description of the model we use to analyse the TOAs of a single pulsar, observed with one telescope. That model does not yet account for the use of different observatories. In this section we explain what we do to accomplish this.

As discussed in Section 4.2, the reduced data products are (sometimes subtly) influenced by many different components of the reduction process. In order to account for slight offsets between TOAs, introduced by using slightly different reduction procedures on individual datasets, a calibration term needs to be introduced when merging TOAs from different observing systems. This extra calibration term takes the form of a ‘‘jump’’, an arbitrary phase offset between datasets, which is fit for simultaneously with other timing model parameters. We use the term dataset for any series of TOAs that can be analysed without a jump. In practice this is any series of TOAs, of the same pulsar, observed with the same hardware elements,

and processed with the same algorithms, at the same observing frequency. Here we combine 7 such datasets (those shown in Figure 4.2).

Jumps have been used routinely when combining data of different observatories and/or data recorders (e.g., Janssen, 2009). This allows us to find a single solution for the timing model of a pulsar timed by multiple observatories. However, the TOAs produced by pipelines at different observatories may have different statistical properties. In order to account for this, we allow the stochastic contributions in our model discussed in Section 4.5.3 to vary between datasets:

- 1) One timing model per pulsar (taken directly from TEMPO2)
- 2) Jumps between different datasets
- 3) A scaling factor for the error bars (EFAC) for each dataset
- 4) An extra white noise component (EQUAD) for each dataset
- 5) Power law red noise for each dataset

A major advantage of this approach is that it allows one to detect statistical differences between observatories that may be introduced by different algorithms/components, and then use this feedback to iteratively improve our datasets.

The analysis of the TOAs consists of two steps. In the first step TEMPO2 is used to find the timing solution for a single pulsar. This includes possible jumps between datasets. Once the timing solution is obtained, the results are passed on to the Bayesian inference method. The Bayesian method then analytically marginalises all parameters of the timing model, including jumps, while using MCMC to explore the rest of the parameter space.

4.6 Results

Now that we have developed the necessary framework to analyse the TOAs, we apply the Bayesian inference method to the observations. In the following subsections we explain in detail how we selected the five pulsars that we already mentioned in Section 4.3.2, and we present the GWB upper limit we are able to calculate using observations of those pulsars.

4.6.1 Selecting the most constraining datasets

For any pulsar, obtaining the timing solution and timing residuals is the first step after obtaining the TOAs. The timing residuals of the pulsars used in this chapter are shown in Figure 4.2, and the parameters of the timing model are shown in Appendix B of this chapter. The timing model also includes several jumps as some of these pulsars have been observed with several European telescopes. The timing solutions we find are quite consistent with the values already published in the literature. Given that we are solving for 56 parameters, it is to be expected that one

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA

or two parameters deviate at the $2\text{-}\sigma$ level. The only unexpected outlier we find is the proper motion in right ascension of J1713+0747, which deviates from Splaver et al. (2005) by over $5\text{-}\sigma$. Given that we are combining data of several telescopes, and that we do not take into account our red noise models in listing these timing solutions, we postpone exploring this difference to future work where the focus lies on investigating the statistics of the timing model parameters in the presence of red noise. Such an investigation is beyond the scope of this manuscript.

With the model of the systematic contributions in place, we first perform the analysis separately for each of the datasets and obtain the posterior probability distribution for their intrinsic noise parameters, specified in Equation (4.7) of the previous section. Note that at this stage of the analysis the contribution from a GWB is not yet included. We determine a marginalised posterior for each pulsar as a function of the following parameter combinations:

- 1) EFAC vs. EQUAD
- 2) Red noise amplitude vs. red noise spectral index

In both cases, the posterior is marginalised over all parameters but two, and the resulting 2-dimensional distribution is displayed as contours at the 1-, 2-, and 3- σ level (the regions where respectively 68%, 95%, and 99.7% of the volume of the posterior is enclosed).

As an example we consider the TOAs of pulsar J1713+0747, which consist of data taken with Effelsberg and Westerbork. Here we focus on the marginalised posterior distributions that represent information about the Effelsberg TOAs; these distributions and the residuals are shown in Figs 4.3 and 4.4. A traditional non-Bayesian analysis of the Effelsberg TOAs consists of a fit to the timing model with TEMPO2, which yields the best-fit parameters, the corresponding uncertainties, and a reduced χ^2 statistic. The reduced χ^2 is defined as:

$$\chi^2 = \frac{1}{n - m} \sum_{i=1}^n \frac{(t_i^{\text{obs}} - t_i^{\text{fit}})^2}{\epsilon^2 \sigma_i^2}, \quad (4.8)$$

where n is the number of observations, m is the number of free parameters in the least-squares fit, t_i^{obs} is the observed TOA, t_i^{fit} is best-fit value of the TOA, σ_i is the TOA uncertainty of t_i^{obs} , and ϵ is the EFAC value. It is usual practice to set the EFAC such that the reduced $\chi^2 = 1$, which is accomplished by: $\epsilon = \sqrt{\chi^2(\epsilon = 1)}$. For the J1713+0747 Effelsberg TOAs, we have $\chi^2(\epsilon = 1) = 18.9$ and therefore $\epsilon = 4.35$.

As can be seen from Figure 4.4, a non-zero red noise component is required to describe the TOAs. The EQUAD parameter is consistent with 0-amplitude according to Figure 4.3, while the EFAC parameter is significantly lower than what

a TEMPO2 $\sqrt{\chi^2}$ estimate would give. This tells us that no separate white-noise component is required to describe the TOAs: all the uncorrelated scatter can be assigned to the error bars on the TOAs. It is also of interest that in this case the EFAC parameter is much smaller, and indeed much closer to 1, than the more traditional estimator $\sqrt{\chi^2}$. The data is reasonably well-modelled by just the presence of red noise.

It is also worth noting that, due to practicalities having to do with hardware changes at the observatories, the TOAs of J1713+0747 end at an earlier epoch than the other 4 pulsars. Although in the future the inclusion of this data will obviously benefit the sensitivity to the GWB, we note that the GWB limit is not negatively effected by this lack of overlap of the TOAs between pulsars.

The analysis of the TOAs of the other pulsars yields similar, but slightly different results. As can be seen in Appendix B of this chapter, some pulsars do have non-negligible white noise, and some do appear to have EFAC values significantly different from 1. As of yet we do not have a complete explanation for the exact form of the marginalised posteriors.

We present the marginalised posterior as a function of the red noise parameters in an intuitive way: as pointed out in Section 4.3.2 we use the same units for the red noise amplitude and red noise spectral index as we use for the GWB parameters. For the analysis of TOAs of just one pulsar, the red noise can now be thought of as if it was generated solely by a GWB with a certain amplitude and spectral index. In this case, the marginalised posterior for the red noise parameters shows us what upper limit we are able to place on the GWB amplitude as a function of spectral index.

We choose a $3\text{-}\sigma$ threshold of $R_a \leq 10^{-13}$ at a spectral index of $\alpha_a = -2/3$. Based on the marginalised posteriors of all the EPTA pulsars, we can decide whether a particular dataset can put a constraint on the GWB lower than this or not. Using this threshold we include five pulsars in our final analysis. These five significantly outperform the other pulsars in terms of how well they can limit the GWB amplitude, and we do not expect to gain any significant sensitivity by including more pulsars in our current archival data sets. The residuals of the pulsars we use in our combined analysis are shown in Figure 4.2. More datasets will be added after some extensive and detailed recalibration procedure of existing datasets.

4.6.2 GWB upper limit

Now that we have selected our pulsars that can significantly contribute to a GWB limit, we are in the position to infer the amplitude and spectral index of the GWB. Our model of the combined data of the five pulsars we selected in Section 4.6.1 consists of all sources we included in the analysis for the single pulsars, and an

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA




Figure 4.3: The marginalised posterior of J1713+0747 (Effelsberg), as a function of the EFAC and EQUAD parameters. The contours are at the 1, 2, and 3- σ level, indicating a respective volume inside that region of 68%, 95%, and 99.7%.




Figure 4.4: The marginalised posterior of J1713+0747 (Effelsberg), as a function of the power-law red noise parameters: the amplitude and the spectral index. The contours are at the 1, 2, and 3- σ level, indicating a respective volume inside that region of 68%, 95%, and 99.7%.

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA




Figure 4.5: The marginalised posterior distribution as a function of the GWB amplitude, and spectral index. The contours marked by 'van Haasteren et al. (2011)' are the results of this chapter at the $1\text{-}\sigma$ and $2\text{-}\sigma$ level, indicating a respective volume inside that region of 68%, and 95%. The vertical dash-dotted line at $\alpha = -2/3$ shows where we expect a GWB generated by supermassive black-hole binaries. The most recent published limits are shown as the three upper limit arrows pointing down, marked by 'Jenet et al. (2006)'.

extra source that corresponds to the GWB. As discussed in Section 4.5.3, the GWB source is a power-law correlated between pulsars as described by Equation (4.4).

As before, we use MCMC to sample the posterior distribution while analytically marginalising over the timing model; now the analytic marginalisation happens simultaneously for the timing models of the five pulsars. In Figure 4.5 we present the posterior, marginalised over all parameters except the GWB amplitude and spectral index. In the same figure we also show the PPTA published values of the GWB limit (Jenet et al., 2006). For the expected spectral index for a GWB generated by a large number of supermassive black-hole binaries, $\alpha = -2/3$, we find a 95% confidence GWB upper limit of $h_c(1\text{yr}) \leq 6 \times 10^{-15}$. This is smaller by a factor of 1.8 than the previously published PPTA limit.

As a cross-check with other codes, and to verify that we are definitely sensitive

4.7. IMPLICATIONS AND OUTLOOK

to the level of the limit we have calculated here, we perform an additional test. We use the TEMPO2 plug-in GWbkgd (Hobbs et al., 2009) to generate simulated timing residuals as produced by a GWB with an amplitude of $h_c(1\text{yr})$. We then create a new set of TOAs, consisting of the values of the simulated timing residuals added to the values of the observed TOAs of the five pulsars that we have analysed in this section. We then redo the whole analysis. Current PTAs aim to reach sensitivities in the order of $h_c(1\text{yr}) = 10^{-15}$ in the future (Jenet et al., 2005), which is over five times more sensitive than the limit we achieve here. In the case that the GWB just happens to be at the $2\text{-}\sigma$ level of our current limit, we demonstrate what such a fivefold increase in sensitivity could do for our ability to measure the GWB parameters by adding a signal of $h_c(1\text{yr}) = 30 \times 10^{-15}$ to our current TOAs. The result is shown in Figure 4.6. We find that the results are consistent with the input parameters of the simulated GWB, and that we can reliably detect a GWB in this case³. The values of the GWB parameters we have used to simulate the GWB lie within the $1\text{-}\sigma$ credible region of Figure 4.6.

4.7 Implications and outlook

The analysis performed in this chapter puts an upper limit on a GWB with a power-law characteristic strain spectrum $h_c = A(f/\text{yr}^{-1})^\alpha$. In the literature, upper limits are typically quoted for various values of α , where the considered α depends on the physics responsible for generation of the GWB. A useful feature of our approach is that we are able to measure α for a strong enough GWB (see chapter 2 for a discussion). The extra degree of freedom in our model, α , necessarily changes the interpretation of the posterior to some extent. We interpret the $2\text{-}\sigma$ contour in our plot of the marginalised posterior as the upper limit on the GWB as a function of α . Fixing α and re-evaluating the $2\text{-}\sigma$ limit based on the posterior for A only does not significantly alter our results.

In this section, we briefly discuss the implications of the new upper limits in the context of two different mechanisms for generation of the GWB: binaries of supermassive black holes, and cosmic strings. We also place the obtained limit in a context with respect to other PTA projects, and we discuss how we expect this limit to evolve in the near future.

³We note that, although such a detection is consistent with a GWB, we would need more pulsars to exclude the possibility that some other effect is causing the correlated signal we detect here.

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA




Figure 4.6: Same marginalised posterior distribution as in 4.5, but here we have injected the residuals of a simulated GWB with amplitude $h_c(1\text{yr}) = 30 \times 10^{-15}$ in the data.

4.7.1 Supermassive black hole binaries

Several authors discuss the characteristic strain spectrum generated by an ensemble of supermassive black holes (SMBHBs) distributed throughout the Universe (Begelman et al., 1980; Phinney, 2001; Jaffe & Backer, 2003; Wyithe & Loeb, 2003). They show that the characteristic strain spectrum generated by such black hole binaries can well be approximated by a power-law:

$$h_c = h_{1\text{yr}} \left(\frac{f}{\text{yr}^{-1}} \right)^{-2/3}, \quad (4.9)$$

where $h_{1\text{yr}}$ is a model-dependent constant. Though the form of the characteristic strain, the power-law, is quite general among the different SMBHB assembly models the authors use in their work, the parameterisations and assumptions about other physical quantities differ between all investigators. The predicted $h_{1\text{yr}}$ therefore differs depending on what SMBHB assembly scenarios were assumed.

Recently, Sesana et al. (2008) have extensively investigated a wide variety of assembly scenarios, including those considered in Jenet et al. (2006). For our purposes in this chapter, their most important result is an estimate of $h_{1\text{yr}}$ for all models⁴. In calculating this value, they take into account the uncertainties of the key model parameters for different scenarios, and come up with $h_{1\text{yr}} \approx 2 \times 10^{-16} - 4 \times 10^{-15}$. We are less than a factor of two away from this range, so we foresee that we can start to rule out some models in the near future.

Two more results of Sesana et al. (2008) are interesting with respect to the limit presented in this chapter. The first is that the frequency dependence of the GWB is expected to be steeper than a power-law $\propto f^{-2/3}$ for frequencies $f \gtrsim 10^{-8}$ Hz. The steepness depends on the chosen model. We have incorporated a varying spectral index α in our current analysis, and since we are not yet able to detect the GWB, we postpone a more thorough investigation of the exact dependence of h_c on f to later work with even better datasets. The second interesting result is that in the frequency range of $10^{-8} \text{ Hz} \leq f \leq 10^{-7} \text{ Hz}$, the GWB might be dominated by single sources. In that case, a search for just a certain characteristic strain spectrum is not appropriate, and we note that further investigation is required in this regard.

4.7.2 Cosmic strings

Several authors have suggested that oscillating cosmic string loops will produce gravitational waves (Vilenkin, 1981; Damour & Vilenkin, 2005; Ölmez et al.,

⁴The model for the GWB that Sesana et al. (2008) use is a broken power-law. Their $h_{1\text{yr}}$ therefore has a slightly different meaning, and our quoted value should be taken as a crude approximation.

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA

2010). Damour & Vilenkin (2005) have used a semi-analytical approach to derive the characteristic strain h_c of the GWB generated by cosmic strings:

$$h_c(f) = 1.6 \times 10^{-14} c^{1/2} p^{-1/2} \epsilon_{\text{eff}}^{-1/6} \times (h/0.65)^{7/6} \left(\frac{G\mu}{10^{-6}} \right)^{1/3} \left(\frac{f}{\text{yr}^{-1}} \right)^{-7/6}, \quad (4.10)$$

where μ is the string tension, G is Newton's constant, c is the average number of cusps per loop oscillation, p is the reconnection probability, ϵ_{eff} is the loop length scale factor, and h is the Hubble constant in units of $100 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Usually, the dimensionless combination $G\mu$ is used to characterise the string tension. Theoretical predictions of string tensions are $10^{-11} \leq G\mu \leq 10^{-6}$ (Damour & Vilenkin, 2005).

From the above expression for the characteristic strain generated by cosmic strings, we see that this is again a power-law, but now with $\alpha = -7/6$. Using a standard model assumption that $c = 1$, the facts that p and ϵ_{eff} are less than one, and that h is expected to be greater than 0.65, we can safely use our derived upper limit on h_c for $\alpha = -7/6$ to limit the string tension: $G\mu \leq 4.0 \times 10^{-9}$. This already places interesting constraints on the theoretical models, and in a few years the EPTA will be able to place much tighter restrictions in the case of a non-detection of a GWB: with only a factor of five decrease of the upper limit, we would be able to completely exclude the $10^{-11} \leq G\mu \leq 10^{-6}$ range mentioned in Damour & Vilenkin (2005).

4.7.3 Comparison with other PTA projects, and prospects

The EPTA data we have analysed in this chapter consists of only a small subset of pulsars from the complete ensemble of pulsars that is observed by the EPTA. The observatories of the EPTA have been continuously upgrading their systems and their data processing pipelines. The limit on the GWB produced by EPTA data will therefore improve substantially in the near future. But even though the limit $h_c(1\text{yr}) \leq 6 \times 10^{-15}$ we present here is significantly lower than the limit published by Jenet et al. (2006), it is worth placing this value into context.

Because of the redness of the GWB signal, the rms of the timing residuals induced by the GWB increases sharply with the duration of the PTA experiment. For an spectral index $\alpha = -2/3$, the rms scales as $\sigma^{GWB} \propto T^{5/3}$, where σ^{GWB} is the rms of the GWB-induced timing residuals, and T is the duration of the experiment. Since PTA observations have been rapidly improving over the past several years, we conclude that the Jenet et al. (2006) results are not representative for the PPTA data quality. Also the greater timespan of high-quality data makes the PPTA much

4.8. CONCLUSION AND DISCUSSION

more sensitive to the GWB now than it was in 2006.

The sensitivity to the GWB should be compared between the three projects PPTA, NANOGrav, and EPTA, by performing the same analysis on their respective datasets: there may still be noise contributions to the TOAs that can be mitigated, and the effect on GWB sensitivity due to the differences between data reduction pipelines is not completely understood. We note here that all three projects have their strengths. The PPTA was the first to organise itself and has been timing a full array of 20 pulsars regularly for over five years, and it has access to the southern hemisphere pulsars. NANOGrav has the largest telescopes at their disposal, and the EPTA has many telescopes which, next to their regular observing schedule, will be coherently combined in the near future (Ferdman et al., 2010).

Most likely, the first detection of GWs by a PTA will occur as a result of a joint IPTA effort. We use the graphs of chapter 2 to create rule of thumb that sensitivity of a PTA to the GWB with a fixed T scales as \sqrt{n} , with n the total number of TOAs. Because the GWB signal increases in strength with time as $\sigma^{GWB} \propto T^{5/3}$, we can be optimistic about prospects for detection of the GWB by PTAs. Given that we have only used five pulsars, most of which were observed for five years, and that there are three currently organised PTAs, we can hope to reach a sensitivity to the GWB by the IPTA in five years an order of magnitude greater than we have accomplished in this chapter. Although this prediction is probably optimistic due to our selection of best pulsars in this chapter, and because it neglects the mostly unknown level of red timing noise which is expected to be present in all millisecond pulsars, we conclude that the sensitivity of PTAs to the GWB will greatly increase in the near future.

4.8 Conclusion and discussion

In this chapter we have developed the methodology on how to handle combined PTA datasets of several telescopes and how to robustly calculate a corresponding upper limit on the GWB. Our Bayesian approach has handled in a straightforward way different data sets of varying duration, regularity, and quality. The current upper limit on the GWB, calculated with EPTA data, is $h_c \leq 6 \times 10^{-15}$ in the case of $\alpha = -2/3$, as predicted for a GWB created by an ensemble of supermassive BH binaries. More generally, the analysis has resulted in a marginalised posterior as a function of the parameters of the GWB: the GWB amplitude and the spectral index.

Due to hardware and software upgrades at the EPTA observatories, and due to the ever increasing time baseline of the data, we expect the sensitivity to increase greatly over the next few years. Especially the combination of the EPTA data sets

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA

with the data of the other PTA projects seems promising.

The raw telescope data must first undergo careful reduction and modelling before it can be analysed by the Bayesian inference method. We have provided some discussion of these processes and have motivated our choice of model for the TOAs. As part of our analysis, we have studied the probability distribution of the pulsar noise parameters, and highlighted the crucial importance of precise characterisation of the red component of pulsar timing noise.

Appendix B

Here we show the timing solutions of all datasets used in this chapter, combined with the posterior distributions for the timing noise.

4.8. CONCLUSION AND DISCUSSION

Pulsar name	J0613–0200	J1012+5307	J1713+0747
Fit and data set			
Telescopes used	NRT	NRT	EFF & WSRT
MJD range	53367 - 55012	53443 - 55030	51426 - 54637
Number of TOAs	280	107	195
Rms timing residual (ns)	912	769	396
Reduced χ^2 value	1.00	1.00	1.13
Epoch	54189	54236	53031
Measured Quantities			
Right ascension, α (J2000)	06:13:43.97385(4)	10:12:33.43241(10)	17:13:49.530782(3)
Declination, δ (J2000)	–02:00:47.0720(12)	+53:07:02.665(2)	+07:47:37.52343(8)
Pulse freq., ν (s^{-1})	326.600562095168(13)	190.26783448248(14)	218.811840486637(30)
Derivative of pulse freq., $\dot{\nu}$ (s^{-2})	$-1.02281(3)\times 10^{-15}$	$-6.1998(4)\times 10^{-16}$	$-4.0836(2)\times 10^{-16}$
PM in RA, μ_α ($mas\ yr^{-1}$)	1.90(4)	3.17(7)	5.017(12)
PM in DEC, μ_δ ($mas\ yr^{-1}$)	–10.31(9)	–24.96(9)	–3.96(3)
Parallax, π (mas)	—	—	0.915(7)
Dispersion measure, DM ($cm^{-3}pc$)	38.77700	9.0176	15.9907
Binary model	DD	ELL1	DD
Orbital period, P_b (d)	1.19851257534(5)	0.60462272322(4)	67.8253309255(20)
Derivative of orbital period, \dot{P}_b	—	—	—
Epoch of periastron, T_0 (MJD)	54189.019(6)	—	53014.9592(7)
Projected sm. axis of orbit, x (lt-s)	1.09144417(8)	0.58181742(13)	32.34242015(7)
Longitude of periastron, ω_0 (deg)	47.1(1.6)	—	176.2109(12)
Orbital eccentricity, e	$5.47(15)\times 10^{-6}$	—	$7.49312(13)\times 10^{-5}$
Time of ascending node (MJD)	—	54236.2078302(3)	—
EPS1 (ϵ_1), $e\ sin\ \omega$	—	$1.18(5)\times 10^{-5}$	—
EPS2 (ϵ_2), $e\ cos\ \omega$	—	$2.20(5)\times 10^{-5}$	—
Sine of inclination angle, $\sin i$	—	—	—
Companion mass, M_c (M_\odot)	—	—	—
Assumptions			
Clock correction procedure	TT(TAI)		
Solar system ephemeris model	DE405		

Table 4.2: The timing solutions for the pulsars used in this chapter before applying the Bayesian inference method. These solutions are determined using TEMPO2, which uses the International Celestial Reference System and Barycentric Coordinate Time. As a result this timing model must be modified before being used with an observing system that inputs TEMPO format parameters. See Hobbs et al. (2006) for more information. Note that the figures in parentheses are the nominal $1-\sigma$ TEMPO2 uncertainties, with EFACs included, and therefore do not include the red noise model. In the GWB limit calculation these respective parameters are marginalised over. Also, the dispersion measure quoted here results from combining these observations with EPTA data of other frequencies. These DM values are used in the dedispersion, but we didn't include all observations in our GWB analysis. We therefore have not fit for the DM here, and an error estimate cannot be

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA

Pulsar name	J1744–1134	J1909–3744
Fit and data set		
Telescopes used	EFF & NRT	NRT
MJD range	51239 - 55001	53366 - 55127
Number of TOAs	159	113
Rms timing residual (ns)	444	134
Reduced χ^2 value	1.05	1.00
Epoch	53120	54247
Measured Quantities		
Right ascension, α (J2000)	17:44:29.391592(7)	19:09:47.437982(5)
Declination, δ (J2000)	-11:34:54.5762(6)	-37:44:14.3176(2)
Pulse freq., ν (s^{-1})	245.426119777227(4)	339.31568732355(1)
Derivative of pulse freq., $\dot{\nu}$ (s^{-2})	$-5.3817(4)\times 10^{-16}$	$-1.614853(8)\times 10^{-15}$
PM in RA, μ_α ($mas\ yr^{-1}$)	18.817(10)	-9.490(11)
PM in DEC, μ_δ ($mas\ yr^{-1}$)	-9.30(6)	-35.89(4)
Parallax, π (mas)	2.602(10)	1.01(7)
Dispersion measure, DM ($cm^{-3}pc$)	3.13632	10.37877
Binary model	—	ELL1
Orbital period, P_b (d)	—	1.53349947490(6)
Derivative of orbital period, \dot{P}_b	—	$3.5(5)\times 10^{-13}$
Epoch of periastron, T_0 (MJD)	—	—
Projected sm. axis of orbit, x (lt-s)	—	1.89799108(11)
Longitude of periastron, ω_0 (deg)	—	—
Orbital eccentricity, e	—	—
Time of ascending node (MJD)	—	54247.169903748(15)
EPS1 (ϵ_1), $e \sin \omega$	—	$6.4(5.5)\times 10^{-8}$
EPS2 (ϵ_2), $e \cos \omega$	—	$-3(3)\times 10^{-8}$
Sine of inclination angle, $\sin i$	—	0.9980(3)
Companion mass, M_c (M_\odot)	—	0.208(7)
Assumptions		
Clock correction procedure	—	TT(TAI)
Solar system ephemeris model	—	DE405

Table 4.3: Same as table 4.2.

4.8. CONCLUSION AND DISCUSSION

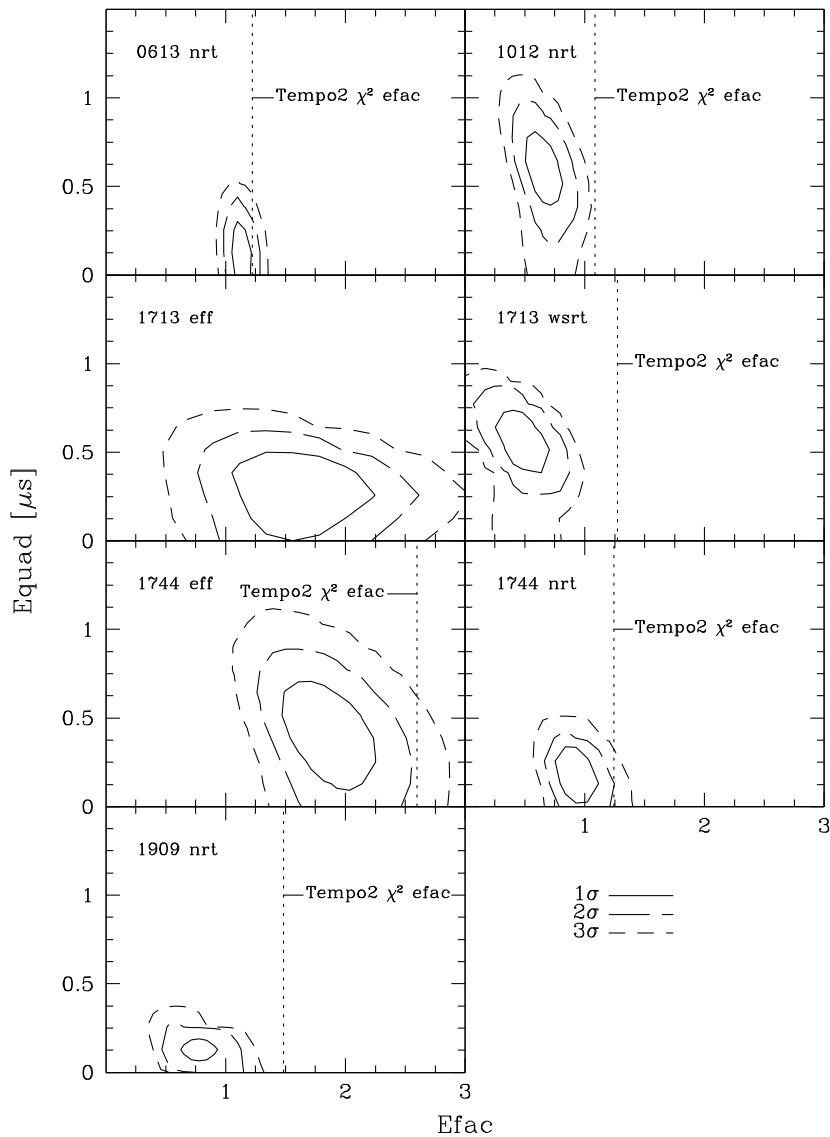


Figure 4.7: The marginalised posteriors of all datasets, as a function of the EFAC and EQUAD parameters. The contours are at the 1, 2, and 3- σ level, indicating a respective volume inside that region of 68%, 95%, and 99.7%. For the J1713-0747 posterior, the TEMPO2 χ^2 estimate is not shown because it has the off-scale value of 4.4.

CHAPTER 4. LIMITING THE GRAVITATIONAL-WAVE BACKGROUND WITH EPTA DATA

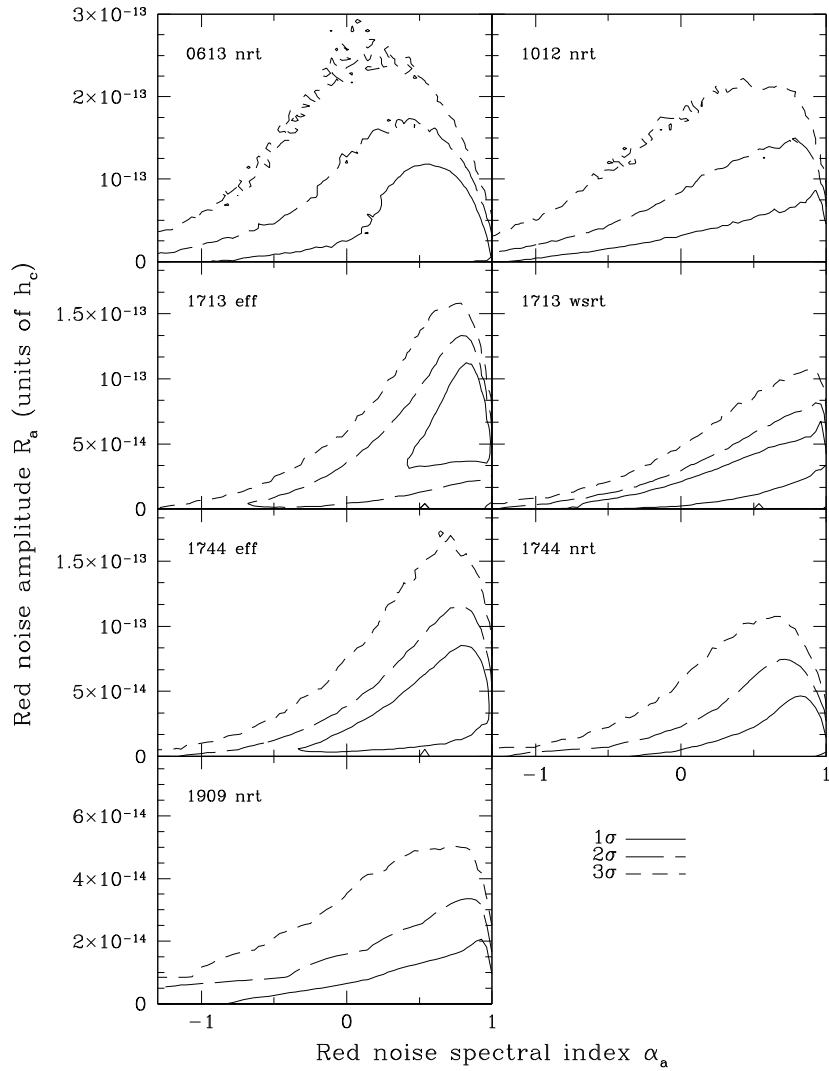


Figure 4.8: The marginalised posterior of all datasets, as a function of the power-law red noise parameters: the amplitude and the spectral index. The contours are at the 1, 2, and 3- σ level, indicating a respective volume inside that region of 68%, 95%, and 99.7%. The more negative the value of α , the steeper the power-law spectrum, with the spectrum approaching a white spectrum at the right of the plot. We also note that the amplitude of the red noise cannot be trivially scaled linearly to an rms value of the timing residuals.