

David Stanislaw Janiszyn

› Gravitational Waves from
Collapsing Domain Walls
in the Early Universe

2026



Gravitational Waves from Collapsing Domain Walls in the Early Universe

BACHELOR THESIS

submitted to the

Universität Münster

Institut für Theoretische Physik

by

David Stanislaw Janiszyn

August 2026

First examiner:

Jun.-Prof. Dr. Kai Schmitz

Second examiner:

Prof. Dr. Michael Klasen

Day of submission:

Friday 21st August, 2026

Abstract

Domain walls are among the sources capable of producing the stochastic gravitational wave background for which pulsar timing array collaborations have reported strong evidence. They are formed after a discrete symmetry is spontaneously broken and must collapse before dominating the energy budget of the Universe in order to be compatible with standard cosmology. Motivated by recent findings on the impact of including the collapse phase in the simulations of domain walls, this thesis analyses the NANOGrav 15-year data set with a corresponding gravitational wave spectrum template. Two scenarios are probed: one in which the domain walls are the only constituent of the signal, and one in which they contribute on top of an agnostic free spectrum capable of explaining the signal on its own. In the former, efficient gravitational wave emission is found to cease at $T_{\text{GW}} = 66.8^{+23.8}_{-17.5}$ MeV, safely before the onset of Big Bang nucleosynthesis, with a non-negligible energy density fraction of $\alpha_{\text{GW}} = 0.320^{+0.136}_{-0.096}$ stored in the domain walls, implying the necessity for corrections to the template, which assumes a radiation-dominated background. The deviations from a previous NANOGrav analysis are found to be consistent with the differences between the two templates. For the microscopic parameters, the domain wall tension is found to be $\sigma^{1/3} = 195.0^{+90.9}_{-62.0}$ TeV, providing access to otherwise unreachable energy scales, while the potential bias amounts to $\Delta V^{1/4} = 117.0^{+70.4}_{-44.0}$ MeV and is thus above the lower bound of $\sim 70\text{--}90$ MeV required to avoid an epoch of domain wall domination. In the latter scenario, the results depend strongly on the prior assumed for the free spectrum amplitudes. A log-uniform prior leaves a domain wall contribution preferred, whereas a linear-exponential prior suppresses it and yields conservative upper limits of $\alpha_{\text{GW}} < 0.221$ and $\sigma^{1/3} < 869$ TeV at the 95 % credible level. The resulting contribution ranges from dominant to negligible depending on this choice.

Declaration of Academic Integrity

I hereby confirm that this thesis, entitled *Gravitational Waves from Collapsing Domain Walls in the Early Universe*, is solely my own work and that I have used no sources or aids other than the ones stated. All passages in my thesis for which other sources, including electronic media, have been used, be it direct quotes or content references, have been acknowledged as such and the sources cited. I am aware that plagiarism is considered an act of deception which can result in sanction in accordance with the examination regulations.

I consent having my thesis cross-checked with other texts to identify possible similarities and to having it stored in a database for this purpose.

I confirm that I have not submitted the following thesis in part or whole as an examination paper before.

Münster, 21st August 2026

David Stanislaw Janiszyn

Contents

1	Introduction	1
2	Theoretical Background	3
2.1	Domain Walls	3
2.1.1	Cosmological Phase Transitions	3
2.1.2	Domain Wall Formation and Properties	5
2.1.3	Cosmological Evolution of Domain Walls	8
2.1.4	Biased Domain Walls	9
2.2	Gravitational Waves	12
2.2.1	Conventions and Notation	12
2.2.2	Derivation of the Wave Equation	14
2.2.3	Energy Carried by Gravitational Waves	17
2.2.4	Stochastic Gravitational Wave Background	19
2.3	Pulsar Timing	21
2.3.1	Millisecond Pulsars and Timing Residuals	21
2.3.2	Gravitational-Wave Signal in Pulsar Timing	23
2.3.3	The Hellings–Downs Correlation	25
3	Gravitational Waves from Domain Walls	29
3.1	Domain Walls as Gravitational Wave Sources	29
3.2	Modelling the Spectrum	31
3.2.1	State of the Art	31
3.2.2	Spectral Template of This Work	33
3.2.2.1	Simulation Setup	34
3.2.2.2	Spectral Parametrisation	35
3.2.2.3	Physical Parametrisation	36
3.2.3	Comparison to Previous Template	39

4	Methodology	43
4.1	Analysis Framework	43
4.1.1	PTA Likelihood	43
4.1.2	Bayesian Statistics	45
4.1.3	Markov Chain Monte Carlo Methods	46
4.2	Overview of the Performed Analyses	47
4.2.1	Analysis Setup	48
4.2.2	Signal Runs	48
4.2.3	Upper-Limit Runs	49
5	Results and Discussion	53
5.1	Signal Runs	54
5.1.1	Thermodynamic Parameters	54
5.1.2	Microscopic Parameters	58
5.1.3	Inclusion of Free Spectral Parameters	62
5.2	Upper-Limit Runs	65
5.2.1	Log-Uniform Prior	66
5.2.2	Linear-Exponential Prior	70
6	Conclusion and Outlook	75
A	PTArcade Model Files	79
	Bibliography	83

Introduction

The earliest picture we have of our Universe is the cosmic microwave background (CMB). It captures the state of the Universe roughly 380 000 years after the Big Bang and is considered one of the most important discoveries of modern cosmology, as its temperature anisotropies have enabled a precise determination of the parameters of the cosmological standard model, among them the spatial curvature and the densities of baryonic matter, dark matter, and dark energy [1]. Looking further back in time is not possible with light as the probed channel. The CMB originates from the very first moment photons were able to propagate freely through the Universe. Before, they were constantly scattered off free electrons, which only disappeared once the Universe had cooled down enough for protons and electrons to form neutral hydrogen [1]. Therefore, other channels of information have to be explored.

A promising candidate is the stochastic gravitational wave background (SGWB) in the nanohertz frequency band, for which pulsar timing array (PTA) collaborations have recently reported evidence [2–6]. Unlike photons, gravitational waves (GWs), which are wave-like solutions to the equations of general relativity (GR), propagate essentially unhindered since their production, carrying information from epochs inaccessible to electromagnetic observations [1]. The SGWB summarises the ‘background hum’ created by the superposition of GWs from a plethora of unresolved sources. While inspiralling supermassive black hole binaries (SMBHBs) are the most likely contributors to such a signal [7], the list of cosmological sources capable of producing a comparable background is long. Among these are the two-dimensional topological defects known as domain walls (DWs), which are theorised to arise when a discrete symmetry is spontaneously broken in the early Universe. When expanding to horizon size and entering an era of self-similar evolution, the so-called scaling regime, the DW energy density dilutes slower than that of matter and radiation, implying an epoch of DW domination, which is incompatible with the observations

of standard cosmology [8]. Thus, a common solution is to introduce a potential bias, softening the initially exact symmetry of the potential and driving the collapse of the DW network.

The goal of this thesis is to implement the state-of-the-art GW template from Notari et al. [9] for DWs into the Python-based GW analysis framework PTArcade [10, 11] and apply it to the NANOGrav 15-year (NG15) data set [2]. This is motivated by recent findings in DW research suggesting drastic changes to the GW spectrum through the inclusion of the collapse phase of a DW network in the underlying simulations [9, 12, 13]. With this recent GW template, two scenarios are probed: in ‘signal runs’ the DW network is the only source contributing to the SGWB, and in ‘upper-limit runs’ it sits on top of an otherwise perfectly explained SGWB signal, modelled by a free spectrum (FS). The latter is a source-agnostic component with a freely varying amplitude in each frequency bin of the data, capable of accounting for the observed signal. In both cases, the spectrum is described in two different ways. The thermodynamic parametrisation characterises the state of the Universe at the time when GW emission from DWs ceases, whereas the microscopic one is formulated directly in terms of the properties of the DWs themselves. For the FS amplitudes, two different prior choices are compared in order to assess their influence on the results.

This thesis is structured as follows. Chapter 2 introduces the theoretical foundations of this work, covering the formation, evolution, and collapse of DWs, the derivation of GWs from GR, and the pulsar timing technique used to measure the SGWB. Chapter 3 connects these topics: after establishing DWs as sources of GWs, the modelling of their GW spectrum is discussed, the template adopted in this work is presented, and it is compared to the one used in the previous analysis of the NG15 data set [14]. Chapter 4 describes the methodology, comprising the PTA likelihood, the relevant aspects of Bayesian statistics and Markov chain Monte Carlo sampling, together with an overview of the analyses performed. The results are presented and discussed in Chapter 5, separately for the signal runs and the upper-limit runs. Finally, Chapter 6 provides a summary and an outlook.

Theoretical Background

In this chapter, the necessary foundations for this thesis are established. In Sec. 2.1, the relevant physics of cosmological DWs in the early Universe, from their formation to their collapse, is introduced. Then, the derivation of gravitational waves from general relativity, alongside the energy carried by them as well as the SGWB are discussed in Sec. 2.2, followed by an introduction to pulsar timing, the technique used to detect the latter, in Sec. 2.3.

2.1 Domain Walls

In this section, we review domain walls. In Sec. 2.1.1, their origin in cosmological phase transitions is introduced, based on [15]. The remaining subsections follow [8] unless stated otherwise: Sec. 2.1.2 covers the formation of DWs and properties relevant for the subsequent analysis, Sec. 2.1.3 their evolution towards the scaling regime, and Sec. 2.1.4 the introduction of a potential bias leading to their collapse.

2.1.1 Cosmological Phase Transitions

One of the major goals of modern theoretical physics is the unification of the four fundamental interactions that govern the description of our Universe: the strong, electromagnetic, and weak forces, as well as gravitation. The approach of many theories is to seek the unification of those forces at high energies, which can be found in the early Universe, close to the big bang. Starting from there, with the passage of time, the Universe expanded and thus cooled down. In this process, each of the four fundamental forces obtained its own identity over the course of multiple cosmological *phase transitions* as the Universe cooled below a certain critical temperature. For example, the electroweak phase transition describes the

process in which the electromagnetic and weak forces adopted their own forms at a critical temperature of around 160 GeV [16] after being described by one joint symmetry group beforehand. In this context, a symmetry group describes the set of transformations under which the Lagrangian density $\mathcal{L}$ of a system is invariant. The breaking of such a symmetry leaves the considered system governed by a smaller group of symmetry rules.

In order to describe the process of symmetry breaking, a real scalar field ϕ , which pervades all space, is considered. A symmetric ground state of the field corresponds to a field value of $\phi = 0$. The development of the underlying effective potential during a second-order¹ phase transition is qualitatively visualised in Fig. 2.1 for the case of a breaking $\mathbb{Z}_2$ symmetry, under which the system is invariant with respect to inversion ($\phi \rightarrow -\phi$). Above the critical temperature T_C , the potential has a single minimum at $\phi = 0$, so that the ground state shares the symmetry of the potential. As the temperature falls below T_C , this minimum turns into a maximum and the potential develops two degenerate minima at $\phi \neq 0$, which no longer exhibit the inversion symmetry. The system settles into one of them and the symmetry is thus spontaneously broken.

Topological defects in the early Universe can arise due to such symmetry breaking. The formation process will be reviewed in the following Sec. 2.1.2. The list of typically considered cosmic defects is long: cosmic strings, monopoles, textures, etc. A rigorous classification of which type of defect arises is made by the homotopy groups $\pi_n(\mathcal{M})$ of the vacuum manifold $\mathcal{M}$. The vacuum manifold is the space of degenerate ground states left after breaking. Its structure is determined by the broken symmetry itself, such that the type of arising defect depends on the type of broken symmetry. Continuous symmetries such as a $U(1)$ symmetry may lead to the formation of cosmic strings, while monopoles can be generated if a $SO(3)$ symmetry is broken [17]. This thesis is centered around the type of cosmic defect named domain walls, which arise if a discrete symmetry, such as the already introduced $\mathbb{Z}_2$ symmetry, is broken.

¹The order of the phase transition is not essential for the formation of DWs. The necessary requirement is the existence of quasi-degenerate vacua after the transition [8]. The second-order case is shown here as an illustrative example.




Figure 2.1: Qualitative visualisation of the development of an effective potential during a second-order phase transition depending on the temperature relative to the critical temperature T_C . For temperatures above T_C , the potential has a single minimum at $\phi = 0$. As the temperature falls below T_C , this minimum turns into a local maximum and splits into two degenerate minima, marked by the white circles.

2.1.2 Domain Wall Formation and Properties

In order to illustrate the formation of DWs, the following toy model of a real scalar field ϕ is to be considered:

$$\mathcal{L} = -\frac{1}{2}\partial^\mu\phi\partial_\mu\phi - V(\phi), \quad (2.1)$$

$$V(\phi) = \frac{\lambda}{4}(\phi^2 - v^2)^2. \quad (2.2)$$

Here, λ is a self-coupling constant, while v denotes the vacuum expectation value (VEV) of the field after the phase transition. For a visual representation of the potential, Fig. 2.1 can be considered once again, where the blue and violet lines, representing the potential below the critical temperature T_C , qualitatively describe

the potential from Eq. (2.2). The VEVs $+v$ and $-v$ correspond to the scalar field values at the two minima, respectively.

As described above, the discrete symmetry of inversion-invariance is broken when the scalar field acquires an expectation value $\langle\phi\rangle = \pm v$. Due to the finiteness of the speed of light, regions of space that are far apart and thus not in causal contact choose the expectation value randomly and independently of each other. Therefore, different regions can acquire different field values. As these areas further grow with the expansion of the Universe, patches of different field values can come into contact and form steep field gradients in between. It is this transitional area between two distinct domains that is called a domain wall. Generally, the described formation process is referred to as the ‘Kibble mechanism’, first described by Kibble in 1976 [18].

For the simplified case of a static planar DW in Minkowski space oriented perpendicular to the z -axis, such that $\phi = \phi(z)$, the solution to the field equation $\frac{d^2\phi}{dz^2} - \frac{dV}{d\phi} = 0$ is given by

$$\phi(z) = v \tanh\left(\sqrt{\frac{\lambda}{2}}vz\right). \quad (2.3)$$

A visual representation of this solution is shown in Fig. 2.2. Here, the scalar field value tends to $\pm v$ for $z \rightarrow \pm\infty$, while the steep gradient mentioned above lies around $z = 0$. There, the scalar field passes through $\phi = 0$, corresponding to the local maximum of the potential in Fig. 2.1, so that energy is stored in the DW. The width of the transition area and thus the thickness δ of a DW can be characterised by the typical length scale in which $\phi(z)$ varies:

$$\delta \simeq \left(\sqrt{\frac{\lambda}{2}}v\right)^{-1}. \quad (2.4)$$

Another crucial property of a DW is its surface energy density σ , also referred to as the *domain wall tension*. As we will see throughout this work, this quantity is




Figure 2.2: Visualisation of the DW solution, obtained in Eq. (2.3), interpolating between the two VEVs. The typical thickness of a DW is given by $\delta \simeq \left(\sqrt{\frac{\lambda}{2}}v\right)^{-1}$.

essential for describing the GWs emitted from DWs. In order to obtain an expression for σ , the T_{00} component of the energy-momentum tensor

$$T_{\mu\nu} = \left(\frac{d\phi}{dz}\right)^2 \text{diag}(1, -1, -1, 0) \quad (2.5)$$

for the static solution introduced in Eq. (2.3) is integrated over the direction perpendicular to the DW:

$$\sigma = \int_{-\infty}^{\infty} T_{00} dz = \frac{4}{3} \sqrt{\frac{\lambda}{2}} v^3. \quad (2.6)$$

Therefore, the DW tension is directly set by the symmetry-breaking scale.

The value of σ is constrained by cosmological observations. As DWs can grow to horizon size, which will be considered thoroughly in the following section 2.1.3, their presence in the observable Universe would introduce large density perturbations in the CMB with an estimated magnitude of $\delta\rho/\rho \sim 10^{12} (\sigma/\text{TeV}^3)$. Observations of the CMB result in an upper bound for those given by $\delta\rho/\rho \leq \mathcal{O}(10^{-5})$. Following

this, the tension of stable DWs is restricted by what is commonly referred to as the Zel’dovich–Kobzarev–Okun (ZKO) bound, first introduced in [19]:

$$\sigma^{1/3} \lesssim \mathcal{O}(\text{MeV}). \quad (2.7)$$

2.1.3 Cosmological Evolution of Domain Walls

In order to describe a network of multiple DWs after their formation, the dynamics of the system have to be considered, which requires one to study the interplay of two forces: the tension force

$$p_T \sim \frac{\sigma}{R_{\text{DW}}}, \quad (2.8)$$

which depends on the DW tension σ and the typical curvature radius R_{DW} of the DW as well as the friction force

$$p_F \sim v_{\text{DW}} T^4 \quad (2.9)$$

with v_{DW} being the DW velocity and the temperature T . This friction term arises due to collisions of particles with the field describing the DWs, which is the reason the expression takes on the form shown: $v_{\text{DW}} T^4 \sim \Delta p \cdot n$, where $\Delta p \sim v_{\text{DW}} T$ is the typical momentum transfer of a particle, while $n \sim T^3$ is the number density of particles. Due to the heavy dependence of p_F on T , the term can be considered negligibly small for sufficiently late times and thus lower temperatures in the Universe. To be precise, this is true if DWs only interact with massive particle states, which is the case considered from now on. Hence, the dynamics of the DW network are now controlled by the tension force, which stretches them to horizon size.

The following period in the evolution of a DW network is called the *scaling regime*, which is an attractor solution [20], meaning that it is reached independently of the chosen initial conditions. During this epoch in the DW evolution, the number of DWs per Hubble volume is of order unity, and the typical length scales of the network are given by the Hubble radius $H^{-1} \sim t$. Here, t denotes the cosmic time. Thus, the energy density of the DW network is given by

$$\rho_{\text{DW}} \sim \frac{H^{-2} \sigma}{H^{-3}} \sim \frac{\sigma}{t}, \quad (2.10)$$

as presented in [21]. A conveniently used parametrisation is introduced via the area parameter $\mathcal{A}$, which is dimensionless and quantifies the number of DWs in one Hubble volume. It is defined as

$$\mathcal{A} = \frac{At}{a(t)V}, \quad (2.11)$$

where a represents the cosmic scale factor and A stands for the comoving DW area inside the comoving volume V . This results in the following expression for the DW energy density:

$$\rho_{\text{DW}} = \frac{\mathcal{A}\sigma}{t}. \quad (2.12)$$

$\mathcal{A}$ is consistently found to take on a constant value of order unity [13, 21, 22].

Comparing the dependency of the DW energy density on time, $\rho_{\text{DW}} \propto t^{-1}$, introduced in Eq. (2.10), to that of matter and radiation reveals a major problem for the theory of DWs. As the energy density of matter and radiation dilutes quicker, as $\rho_{\text{mat.}}, \rho_{\text{rad.}} \propto t^{-2}$ [23] in the respective eras of matter or radiation domination, eventually, an era of DW domination will be entered, where the scale factor evolves as $a(t) \propto t^2$. Because such a rapid expansion of the Universe is not compatible with the current cosmological model of our Universe, stable DWs are cosmologically unacceptable and thus must decay.

2.1.4 Biased Domain Walls

In order to circumvent the problem of stable DWs, presented in Sec. 2.1.3, a bias term V_{bias} is commonly introduced in the potential. One possible form is

$$V_{\text{bias}}(\phi) = \varepsilon v \phi \left(\frac{1}{3} \phi^2 - v^2 \right), \quad (2.13)$$

which is added to the exactly $\mathbb{Z}_2$ -symmetric potential from Eq. (2.2). ε is a dimensionless parameter that quantifies the strength of the bias term. The effect of the introduction of the bias term on the form of the potential is visualised in Fig. 2.3.

While the position of the two minima remains unchanged at the two VEVs $\pm v$, an energy difference is established:

$$\Delta V = V(-v) - V(+v) = \frac{4}{3}\varepsilon v^4. \quad (2.14)$$




Figure 2.3: Qualitative visualisation of the potential from Eq. (2.2) including the energy bias introduced in Eq. (2.13). The positions of the minima remain at $\phi = \pm v$, whereas the energetic difference between the minima is $\Delta V = \frac{4}{3}\varepsilon v^4$. Own image, based on [8].

Due to the scalar field's preference to occupy the energetically lowest state, DWs become unstable and collapse towards the energetically favoured vacuum, which corresponds to $\phi = v$ for the case visualised in Fig. 2.3. It is also worth noting that this approach to solving the DW overclosure problem comes at the cost of softening the discrete $\mathbb{Z}_2$ symmetry. Whereas it is exact in the absence of the bias term, it is only approximately fulfilled once the bias term is included. This poses a danger to the formation of DWs in the first place, as too large a potential bias would cause the scalar field to settle into the lower-energy vacuum everywhere. Thus, no distinct domains and no DWs in between would form. An upper limit for the energetic difference between the two minima ΔV , which prevents the occurrence of this scenario, can be deduced. For this, we consider the probability that a patch in

space occupies either of the two VEVs, p_+ for v and p_- for $-v$. The ratio of these two quantities is given by

$$\frac{p_-}{p_+} = \exp\left(-\frac{\Delta F}{T}\right) \simeq \exp\left(-\frac{\Delta V}{V_0}\right). \quad (2.15)$$

Here, $\Delta F = \Delta V \cdot \xi^3$ denotes the free energy difference between the two minima, with ξ being the correlation length at the time of the phase transition. The temperature is estimated as the Ginzburg temperature $T \simeq V_0 \cdot \xi^3$, where V_0 describes the height of the potential barrier between the two VEVs, as can be seen in Fig 2.3. Numerical simulations of percolation on a cubic lattice yield a critical probability of $p_c = 0.311$ [24]. Below this threshold, no infinite cluster of $\phi = -v$ vacuum regions can form. This results in the following condition:

$$\frac{\Delta V}{V_0} < \ln\left(\frac{1-p_c}{p_c}\right) \approx 0.795. \quad (2.16)$$

In order to ensure DW formation, this criterion for the potential bias must be satisfied.

As the introduction of the bias term is motivated by preventing the Universe from entering an era of DW domination, the potential bias must be chosen large enough to ensure that annihilation occurs before domination sets in. The concrete time scale for DW domination can be derived from the condition that the DW energy density ρ_{DW} equals the critical energy density, which is given in the assumed era of a radiation-dominated cosmology by $\rho_c = 3M_p^2/(4t^2)$. Thus, an epoch of DW domination arises at

$$t_{\text{dom}} = \frac{3M_p^2}{4\mathcal{A}\sigma}. \quad (2.17)$$

Here, M_p denotes the reduced Planck mass, $M_p = \sqrt{\hbar c/(8\pi G)} \simeq 2.435 \cdot 10^{18} \text{ GeV}$. Next, the time of DW annihilation can be estimated from the condition that it occurs once the volume pressure force $p_V \sim \Delta V$ due to the potential bias exceeds the DW tension force p_T , introduced in Eq. (2.8). As the typical curvature radius of a DW is given by $R_{\text{DW}} \simeq t/\mathcal{A}$ during the scaling regime, the DW annihilation time can be estimated to be

$$t_{\text{ann}} = C_{\text{ann}} \frac{\mathcal{A}\sigma}{\Delta V}, \quad (2.18)$$

where $C_{\text{ann}} \simeq 2\text{--}5$ is a dimensionless constant determined numerically. Combining these two expressions by demanding DW annihilation before the onset of DW domination, $t_{\text{ann}} < t_{\text{dom}}$, results in a lower bound for the potential bias term:

$$\Delta V > \frac{4C_{\text{ann}}\mathcal{A}^2\sigma^2}{3M_p^2}. \quad (2.19)$$

To summarise, this section has established an upper bound for the potential bias ΔV , given by Eq. (2.16), to ensure the formation of large-scale DWs, and a lower bound, quantified in Eq. (2.19), that prevents a cosmological era of DW domination.

2.2 Gravitational Waves

In this section, we study gravitational waves, a direct prediction of Albert Einstein's theory of general relativity [25]. The treatment follows [26] unless stated otherwise, while a more detailed description can be found in [27]. After introducing the notation and conventions in Sec. 2.2.1, the wave equation for GWs is derived from Einstein's field equations (EFEs) in Sec. 2.2.2, followed by an expression for the energy density carried by GWs in Sec. 2.2.3. Finally, the stochastic gravitational wave background is introduced in Sec. 2.2.4.

2.2.1 Conventions and Notation

Before we start, we introduce some notation and conventions of GR without going into too much detail about this theory. As usual, greek indices, such as μ and ν , denote space-time coordinates ranging from 0 to 3, where 0 denotes the time dimension and 1 to 3 denote the space dimensions, while Latin letters, such as i and j , range from 1 to 3 and purely describe spatial dimensions. The convention used for the metric of flat Minkowski space-time is

$$\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1) = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (2.20)$$

whereas we will describe a general metric with $g_{\mu\nu}$. Moreover, the Christoffel symbol is defined as

$$\Gamma_{\mu\nu}^{\rho} = \frac{1}{2}g^{\rho\sigma}(\partial_{\mu}g_{\sigma\nu} + \partial_{\nu}g_{\sigma\mu} - \partial_{\sigma}g_{\mu\nu}), \quad (2.21)$$

where we use the abbreviation

$$\partial_{\mu} \equiv \frac{\partial}{\partial x^{\mu}} = \left(\frac{1}{c}\partial_t, \partial_i \right), \quad (2.22)$$

with c denoting the speed of light. Later on, we will also make use of the following definitions:

$$\partial^{\mu} = \eta^{\mu\nu}\partial_{\nu}, \quad (2.23)$$

and

$$h = \eta^{\mu\nu}h_{\mu\nu}. \quad (2.24)$$

Following that, the Riemann tensor is given by

$$R^{\mu}_{\nu\rho\sigma} = \partial_{\rho}\Gamma^{\mu}_{\nu\sigma} - \partial_{\sigma}\Gamma^{\mu}_{\nu\rho} + \Gamma^{\mu}_{\alpha\rho}\Gamma^{\alpha}_{\nu\sigma} - \Gamma^{\mu}_{\alpha\sigma}\Gamma^{\alpha}_{\nu\rho}, \quad (2.25)$$

based on which the Ricci tensor is simply expressed as

$$R_{\mu\nu} = R^{\alpha}_{\mu\alpha\nu} \quad (2.26)$$

and the Ricci scalar as

$$R = g^{\mu\nu}R_{\mu\nu}. \quad (2.27)$$

This leads us to the most important set of equations of GR, around which the rest of it is centred: the Einstein field equations. They are given by

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu}. \quad (2.28)$$

Here, $G_{\mu\nu}$ is called the Einstein tensor, G is the gravitational constant, and c denotes the speed of light. On the right-hand side of the equation, the energy-momentum tensor $T_{\mu\nu}$ appears, which encodes the energy density, momentum density and stress of all matter and radiation fields present. The core principle of Einstein's theory of general relativity is often summarised by John Archibald Wheeler's statement that 'spacetime tells matter how to move; matter tells spacetime how to curve' [28].

The geometry of spacetime thus governs the motion of matter and radiation within it, while the distribution of energy and momentum in turn determines this geometry. The second part of this statement is encapsulated in Eq. (2.28).

2.2.2 Derivation of the Wave Equation

The goal of this section is to derive the wave equation for GWs by solving the EFEs from Eq. (2.28) for a small metric perturbation $|h_{\mu\nu}| \ll 1$ that causes a deviation from flat Minkowski space-time:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}. \quad (2.29)$$

Due to this weak field treatment of the perturbation, we can safely assume that the operations of index raising and lowering can be performed with the Minkowski metric. The inclusion of $h_{\mu\nu}$ causes the Einstein tensor, expanded up to linear order while neglecting higher order terms, to take the form of

$$G_{\mu\nu} = \frac{1}{2} (\partial_\mu \partial^\alpha h_{\alpha\nu} + \partial_\nu \partial^\alpha h_{\alpha\mu} - \partial_\mu \partial_\nu h - \square h_{\mu\nu} + \eta_{\mu\nu} \square h - \eta_{\mu\nu} \partial^\alpha \partial^\beta h_{\alpha\beta}), \quad (2.30)$$

where we make use of the flat space d'Alembertian operator $\square \equiv \eta^{\mu\nu} \partial_\mu \partial_\nu$. In an effort to make Eq. (2.30) easier to handle, we introduce the trace-reversed perturbation

$$\bar{h}_{\mu\nu} \equiv h_{\mu\nu} - \eta_{\mu\nu} \frac{h}{2}, \quad (2.31)$$

which further yields $\bar{h} = \eta^{\mu\nu} \bar{h}_{\mu\nu} = h - 2h = -h$ and $h_{\mu\nu} = \bar{h}_{\mu\nu} - \frac{1}{2} \eta_{\mu\nu} \bar{h}$. Thus, the Einstein tensor simplifies to:

$$G_{\mu\nu} = \frac{1}{2} (\partial_\mu \partial^\alpha \bar{h}_{\alpha\nu} + \partial_\nu \partial^\alpha \bar{h}_{\alpha\mu} - \square \bar{h}_{\mu\nu} - \eta_{\mu\nu} \partial^\alpha \partial^\beta \bar{h}_{\alpha\beta}). \quad (2.32)$$

This expression can be narrowed down even further by reducing the degrees of freedom, which is done by choosing an appropriate gauge. For this, we perform the coordinate transformation $x^\mu \rightarrow x^\mu + \xi^\mu$ which affects the introduced perturbation as $h_{\mu\nu} \rightarrow h_{\mu\nu} - (\partial_\mu \xi_\nu + \partial_\nu \xi_\mu)$. In order to keep the imposed weak field condition

$|h_{\mu\nu}| \ll 1$ fulfilled, we require that $\partial_\mu \xi_\nu$ is of the same order of magnitude as $|h_{\mu\nu}|$. This allows us to apply the Lorenz gauge

$$\partial^\mu \bar{h}_{\mu\nu} = 0, \quad (2.33)$$

which further simplifies the Einstein tensor to take the form of

$$G_{\mu\nu} = -\frac{1}{2} \square \bar{h}_{\mu\nu} \quad (2.34)$$

and reduces Eq. (2.28) to

$$\square \bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} T_{\mu\nu}. \quad (2.35)$$

In our case, we can set $T_{\mu\nu} = 0$, as we consider GWs in vacuum, far away from their source, such that Eq. (2.35) reduces to

$$\square \bar{h}_{\mu\nu} = 0. \quad (2.36)$$

This takes the form of a wave equation, whose solutions describe GWs propagating at the speed of light.

Starting from initially ten independent components of the symmetric 4×4 matrix $h_{\mu\nu}$, only six components remain independent after the four Lorenz gauge conditions have been imposed in Eq. (2.33). However, the gauge is not completely fixed. A further transformation $x^\mu \rightarrow x^\mu + \xi^\mu$ preserves $\partial^\mu \bar{h}_{\mu\nu} = 0$ under the condition that $\square \xi^\mu = 0$, which leaves four residual gauge functions at our disposal. These are used to impose the transverse-traceless (TT) gauge

$$h^{0\mu} = 0, \quad h_i^i = 0, \quad \partial^j h_{ij} = 0, \quad (2.37)$$

meaning that the perturbation is purely spatial, traceless and transverse to the direction of propagation. Since the trace vanishes, the trace-reversed and the original perturbation coincide, $\bar{h}_{\mu\nu} = h_{\mu\nu}$, and the bar may be dropped from here on. Of the six components remaining after the Lorenz gauge, only two are therefore of radiative nature.

Thus, a solution to Eq. (2.36) in a coordinate system chosen such that the direction of GW propagation $\vec{k}$ coincides with the z -axis can be given in the TT gauge by

$$h_{\mu\nu}(t, z) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \cos \left[\omega \left(t - \frac{z}{c} \right) \right]. \quad (2.38)$$

Here, the amplitudes of the two independent polarisation states, which correspond to the two radiative degrees of freedom, are given by h_+ and $h_\times$. The naming convention for the ‘plus’ (+) and ‘cross’ ($\times$) modes becomes apparent once we consider the effect of a passing GW on a ring of freely falling test masses, as visualised in Fig. 2.4. It should be stressed that in the TT gauge the coordinate positions of the test masses remain unchanged. The physical effect is a periodic modulation of the




Figure 2.4: Visualisation of the effect of the two GW polarisation states on a ring of test masses at different times t with respect to the period P of the GW. Reproduced from [29].

2.2.3 Energy Carried by Gravitational Waves

Richard Feynman proposed a simple thought experiment, known as the sticky bead argument, which addressed a long-standing debate about whether GWs are a genuine physical phenomenon or merely a coordinate artefact [30]. The setup consists of a rod with attached beads that are free to slide along it. We can now imagine a GW passing this construction, causing a displacement of the beads relative to the rod so that friction generates heat. The GW performs work on the system, from which it follows that GWs necessarily carry energy and are therefore physically real. Since energy is a source of space-time curvature, we can conclude that GWs should also influence the space-time through which they propagate. Motivated by this, an expression for the energy density of GWs is derived in the following, summarising the arguments presented in [27].

In order to account for the curvature sourced by a GW, we abandon the restriction to a flat background space-time, initially set up in Eq. (2.29). Instead, the background of the small perturbation $h_{\mu\nu}$ is given by a non-Minkowskian metric $\hat{g}_{\mu\nu}$:

$$g_{\mu\nu} = \hat{g}_{\mu\nu} + h_{\mu\nu}. \quad (2.39)$$

This raises the crucial question of how the two contributions can be distinguished, since a clear separation, such as in Eq. (2.29), is no longer possible. Without further specification, any redistribution between the two components would result in the same metric $g_{\mu\nu}$. Thus, we now characterise the two contributions to $g_{\mu\nu}$ via the typical scale of their variation. Concretely, we assume the background metric $\hat{g}_{\mu\nu}$ to vary on a scale L_B , while the perturbative contribution (the GW) varies on a scale λ . We separate the two regimes by imposing

$$\frac{\lambda}{2\pi} \ll L_B, \quad (2.40)$$

where $h_{\mu\nu}$ can be interpreted as a small ripple propagating on a smooth background geometry. Alternatively, the contribution of the GW can be identified as a rapidly varying perturbation with a frequency f in comparison to an almost static background, characterised by a frequency f_B :

$$f \gg f_B. \quad (2.41)$$

While both requirements can be used to derive the desired quantity, for simplicity, we will focus on Eq. (2.41). From here on, we consider the EFEs in the equivalent form

$$R_{\mu\nu} = \frac{8\pi G}{c^4} \left(T_{\mu\nu} - \frac{1}{2} g_{\mu\nu} T \right), \quad (2.42)$$

where T is the trace of the energy-momentum tensor $T_{\mu\nu}$. Next, we expand the Ricci tensor to quadratic order in $h_{\mu\nu}$, yielding

$$R_{\mu\nu} = \hat{R}_{\mu\nu} + R_{\mu\nu}^{(1)} + R_{\mu\nu}^{(2)} + \mathcal{O}(h^3). \quad (2.43)$$

Here, $\hat{R}_{\mu\nu}$ is constructed solely from $\hat{g}_{\mu\nu}$, while $R_{\mu\nu}^{(1)}$ and $R_{\mu\nu}^{(2)}$ collect all terms that are linear and quadratic in $h_{\mu\nu}$, respectively. By construction, $\hat{R}_{\mu\nu}$ contains only slowly varying modes, whereas terms linear in $h_{\mu\nu}$ consist purely of rapidly varying ones. However, terms quadratic in $h_{\mu\nu}$ contain a mixture of both, since the product of two high-frequency modes also gives rise to a slowly varying contribution. Eq. (2.42) can therefore be separated into a low- and a high-frequency part:

$$\hat{R}_{\mu\nu} = - [R_{\mu\nu}^{(2)}]^{\text{Low}} + \frac{8\pi G}{c^4} \left(T_{\mu\nu} - \frac{1}{2} g_{\mu\nu} T \right)^{\text{Low}}, \quad (2.44)$$

$$R_{\mu\nu}^{(1)} = - [R_{\mu\nu}^{(2)}]^{\text{High}} + \frac{8\pi G}{c^4} \left(T_{\mu\nu} - \frac{1}{2} g_{\mu\nu} T \right)^{\text{High}}. \quad (2.45)$$

Due to the separation of frequency scales, imposed in Eq. (2.41), we can now introduce an intermediate time scale τ , fulfilling the requirement $\frac{1}{f_B} \gg \tau \gg \frac{1}{f}$. Averaging Eq. (2.42) over τ removes all rapidly varying contributions while leaving the slowly varying ones unaffected, so that taking the average is equivalent to extracting the low-frequency part. In particular, this yields $\langle \hat{R}_{\mu\nu} \rangle = \hat{R}_{\mu\nu}$, and we obtain

$$\hat{R}_{\mu\nu} = - \langle R_{\mu\nu}^{(2)} \rangle + \frac{8\pi G}{c^4} \left\langle T_{\mu\nu} - \frac{1}{2} g_{\mu\nu} T \right\rangle. \quad (2.46)$$

Reversing the trace of Eq. (2.46) yields an equation of the form of the EFEs for the background metric,

$$\hat{R}_{\mu\nu} - \frac{1}{2} \hat{g}_{\mu\nu} \hat{R} = \frac{8\pi G}{c^4} (\langle T_{\mu\nu} \rangle + t_{\mu\nu}), \quad (2.47)$$

where the effective energy-momentum tensor of the GWs has been defined as

$$t_{\mu\nu} \equiv -\frac{c^4}{8\pi G} \left\langle R_{\mu\nu}^{(2)} - \frac{1}{2} \hat{g}_{\mu\nu} R^{(2)} \right\rangle, \quad R^{(2)} = \hat{g}^{\mu\nu} R_{\mu\nu}^{(2)}. \quad (2.48)$$

In Eq. (2.47), $t_{\mu\nu}$ enters the field equations for the background on the same footing as the energy-momentum tensor of matter. It is therefore interpreted as the energy-momentum tensor carried by the GWs themselves, which confirms the expectation formulated above that GWs source curvature of the space-time through which they propagate. The explicit evaluation of Eq. (2.48) is lengthy and thus briefly summarised here. Making use of the fact that the average of a total derivative vanishes, we obtain for GWs far away from their source, where the background can be taken as flat and the TT gauge can be imposed,

$$t_{\mu\nu} = \frac{c^4}{32\pi G} \langle \partial_\mu h_{\alpha\beta} \partial_\nu h^{\alpha\beta} \rangle, \quad (2.49)$$

so that the energy density carried by a GW is given by

$$\rho_{\text{GW}} = t^{00} = \frac{c^2}{32\pi G} \langle \dot{h}_{ab}^{\text{TT}} \dot{h}_{ab}^{\text{TT}} \rangle = \frac{c^2}{16\pi G} \langle \dot{h}_+^2 + \dot{h}_\times^2 \rangle, \quad (2.50)$$

where a dot denotes differentiation with respect to time. It should be emphasised that ρ_{GW} is meaningful only as an average over several wavelengths.

2.2.4 Stochastic Gravitational Wave Background

In this work, we will not consider individual GWs, such as those emitted by a resolvable black hole binary, but rather their superposition from a large number of unresolved sources. If many GWs of comparable frequencies reach us simultaneously, they cannot be discerned individually and add up to an incoherent ‘hum’, which is referred to as the *stochastic gravitational wave background*. Describing the signal as stochastic means that we cannot formulate a deterministic waveform with which the amplitude can be predicted at a given point in space and time. Instead, the SGWB is characterised by its statistical properties, in particular by its variance. In the following, the description of the SGWB will be introduced alongside a quantity of central importance, the energy-density spectrum $\Omega_{\text{GW}}(f)$, following the treatment in [26].

First, we define

$$\Omega_{\text{GW}}(f) \equiv \frac{1}{\rho_c} \frac{d\rho_{\text{GW}}}{d \ln f}. \quad (2.51)$$

In this way, we express the energy density contained in GWs as a spectrum with respect to the frequency f and relative to the closure density $\rho_c = 3H_0^2 c^2 / (8\pi G)$, which corresponds to a spatially flat Universe. H_0 denotes the present-day value of the Hubble constant. We choose a logarithmic frequency scale since the frequencies contributing to the SGWB span multiple orders of magnitude. In order to make use of the expression for the energy density ρ_{GW} , derived beforehand and given in Eq. (2.50), we formulate the perturbative metric as an expansion over plane waves in the TT gauge:

$$h_{ab}(t, \vec{x}) = \sum_{A=+, \times} \int_{-\infty}^{\infty} df \int_{S^2} d^2\Omega_{\hat{n}} h_A(f, \hat{n}) e^{2\pi i f(t + \hat{n} \cdot \vec{x}/c)} e_{ab}^A(\hat{n}). \quad (2.52)$$

Here, $h_A(f, \hat{n})$ denotes complex random fields, whose moments govern the statistical properties of the SGWB, $\hat{n}$ is a unit vector which is oriented towards the source of the GWs, and $e_{ab}^A(\hat{n})$ are the polarisation basis tensors, constructed from an orthonormal basis $\{\hat{l}, \hat{m}\}$ transverse to $\hat{n}$ as

$$e_{ab}^+(\hat{n}) = \hat{l}_a \hat{l}_b - \hat{m}_a \hat{m}_b, \quad e_{ab}^\times(\hat{n}) = \hat{l}_a \hat{m}_b + \hat{m}_a \hat{l}_b, \quad (2.53)$$

with

$$\hat{n} = \sin \theta \cos \phi \hat{x} + \sin \theta \sin \phi \hat{y} + \cos \theta \hat{z} \equiv \hat{r}, \quad (2.54)$$

$$\hat{l} = \cos \theta \cos \phi \hat{x} + \cos \theta \sin \phi \hat{y} - \sin \theta \hat{z} \equiv \hat{\theta}, \quad (2.55)$$

$$\hat{m} = -\sin \phi \hat{x} + \cos \phi \hat{y} \equiv \hat{\phi}. \quad (2.56)$$

Inserting Eq. (2.52) into Eq. (2.50) yields a double sum over polarisations and a double integral over frequencies and sky directions, so that the energy density is determined by the two-point function $\langle h_A(f, \hat{n}) h_{A'}^*(f', \hat{n}') \rangle$. Assuming the SGWB to be Gaussian, stationary, isotropic, spatially homogeneous, and unpolarised, this correlator reduces to

$$\langle h_A(f, \hat{n}) h_{A'}^*(f', \hat{n}') \rangle = \frac{S_h(f)}{16\pi} \delta(f - f') \delta_{AA'} \delta^{(2)}(\hat{n}, \hat{n}'), \quad (2.57)$$

where $S_h(f)$ denotes the one-sided power spectral density (PSD) of the SGWB Fourier modes. The delta functions collapse the double integrals, and with the identities $\int d\hat{\Omega} = 4\pi$ and $\sum_A e_{ab}^A e_A^{ab} = 4$, together with a restriction of the frequency integration to positive values, we obtain

$$\langle \dot{h}_{ab} \dot{h}^{ab} \rangle = 8\pi^2 \int_0^\infty df f^2 S_h(f). \quad (2.58)$$

Together with Eq. (2.50), this finally yields the energy density spectrum of the SGWB in terms of its PSD,

$$\Omega_{\text{GW}}(f) = \frac{2\pi^2}{3H_0^2} f^3 S_h(f). \quad (2.59)$$

This quantity is central to the remainder of this thesis, as the spectra predicted by the theoretical DW models will be tested against the observational data in terms of $\Omega_{\text{GW}}(f)$.

2.3 Pulsar Timing

Following the introduction of GWs in Sec. 2.2, a specific technique for detecting them is discussed in this section, which makes use of a class of neutron stars known as pulsars. First, millisecond pulsars and the concept of timing residuals are introduced in Sec. 2.3.1. It is then discussed in Sec. 2.3.2 how a pulsar is used to detect the imprint of a GW, before the Hellings–Downs correlation, the characteristic signature of the SGWB, is derived in Sec. 2.3.3. For this, we mainly follow the treatment presented in [26].

2.3.1 Millisecond Pulsars and Timing Residuals

The death of a star gives rise to a variety of compact remnants, depending primarily on the mass of its progenitor. Stars with initial masses between approximately $9 M_\odot$ and $25 M_\odot$, where $M_\odot$ denotes the solar mass, end their lives in a core-collapse supernova, leaving behind a neutron star [31]. The radius of such a neutron star lies between 10 km and 15 km and is thus drastically reduced compared to the radius of its progenitor. Due to the conservation of angular momentum and magnetic flux, such a collapse increases the star's rotational velocity and magnetic field strength by orders of magnitude. Charged particles are guided along the magnetic field

lines, leading to the emission of radio beams above the magnetic poles. As the axis of the magnetic field, and thus of radio emission, is generally not aligned with the rotational axis of the neutron star, the star acts as a ‘lighthouse’: the emitted beam sweeps across the sky once per rotation and can be perceived on Earth as a periodic pulse if it crosses our line of sight. Neutron stars that fulfil this criterion are called *pulsars*. For the measurement of the SGWB, we are interested in the class of so-called millisecond pulsars (MSPs), which are characterised by a pulse period of the order of 10^{-3} s and, most importantly, a remarkable rotational stability over long timescales.

This stability is what makes MSPs useful as precision clocks, as the arrival times of their pulses can be predicted over years. The process of measuring the time of arrival (TOA) t_{TOA} is lengthy and is thus only briefly summarised in the following. The incoming pulses are recorded over minutes to hours, corrected for the dispersion delay caused by the ionised interstellar medium, and folded with an estimate of the pulse period in order to obtain a profile with a high signal-to-noise ratio. At the end of this process, the folded profile is then cross-correlated with the pulsar-specific template profile, which represents the stable average pulse shape of the pulsar, and yields the arrival time t_{TOA} of the observation as the offset that best aligns the two.

These TOAs are described by a deterministic *timing model*, which relates the pulse emission time at the pulsar t_{PSR} to the measured time of arrival t_{TOA} :

$$t_{\text{PSR}} = t_{\text{TOA}} - \Delta_{\odot} - \Delta_{\text{ISM}} - \Delta_B. \quad (2.60)$$

Here, the additional terms describe various contributions that influence the TOA. $\Delta_{\odot}$ accounts for the transformation to the quasi-inertial frame of the Solar System Barycentre (Roemer, Einstein, and Shapiro delays within the Solar System), Δ_{ISM} includes the frequency-dependent propagation delays through the ionised interstellar medium and the solar wind, and, for pulsars in binary systems, Δ_B describes the analogous transformation from the binary barycentre to the pulsar. The central quantity of this section, the *timing residual* δt , is then given by the difference between the measured TOAs and those predicted by the timing model, whose parameters are determined by a least-squares fit to the measured TOAs. These timing residuals

contain every contribution not accounted for by the timing model and thus consist of noise and, most importantly for this work, the imprint of GWs.

2.3.2 Gravitational-Wave Signal in Pulsar Timing

Following the introduction of our ‘measurement setup’ in Sec. 2.3.1, we will now consider the effect of a GW on the measured TOAs. The proper path length for photons, travelling from a pulsar to our observatory, is perturbed by a GW, resulting in irregularities in the measured pulsar period. As a description for the effect of a GW, we introduce the line element

$$ds^2 = -(c dt)^2 + [\delta_{ab} + h_{ab}^{\text{TT}}(t, \vec{x})] dx^a dx^b, \quad (2.61)$$

which we can simplify by considering a photon propagating along the x -axis towards the observer at the origin, implying $ds^2 = 0$ and yielding

$$dx \approx -c \left\{ 1 - \frac{1}{2} h_{xx}^{\text{TT}}[t, \vec{x}(t)] \right\} dt. \quad (2.62)$$

Through integration, we obtain

$$\frac{L}{c} = t_{\text{obs}} - t_{\text{em}} - \frac{1}{2} \int_{t_{\text{em}}}^{t_{\text{obs}}} dt' h_{xx}^{\text{TT}}[t', \vec{x}(t')], \quad (2.63)$$

where L denotes the separation distance between a pulsar and Earth, and t_{obs} and t_{em} mark the time of observation and emission, respectively. Since h_{xx}^{TT} is small, we evaluate the integral along the unperturbed photon trajectory, using $t_{\text{obs}} \approx t_{\text{em}} + \frac{L}{c}$. Generalising from a photon travelling along the x -axis to an arbitrary pulsar direction $\hat{p}$, by replacing h_{xx}^{TT} with $p^a p^b h_{ab}^{\text{TT}}$, we obtain

$$t_{\text{obs}} = t_{\text{em}} + \frac{L}{c} + \frac{p^a p^b}{2} \int_{t_{\text{em}}}^{t_{\text{em}} + \frac{L}{c}} dt' h_{ab}^{\text{TT}}[t', (t_{\text{em}} + \frac{L}{c} - t') \hat{p}]. \quad (2.64)$$

Since the period of the pulsar P is known, the next step is to derive an expression for the deviation ΔP of the observed time from the expected one. We obtain an expression t'_{obs} for the TOA of the second pulse by shifting the time of emission by

one period: $t'_{\text{em}} = t_{\text{em}} + P$. The difference in the time of observation is then given as

$$t'_{\text{obs}} - t_{\text{obs}} = P + \Delta P, \quad (2.65)$$

with

$$\Delta P = \frac{p^a p^b}{2} \int_{t_{\text{em}}}^{t_{\text{em}} + \frac{L}{c}} dt' \{ h_{ab}^{\text{TT}}[t' + P, \vec{x}_0(t')] - h_{ab}^{\text{TT}}[t', \vec{x}_0(t')] \}. \quad (2.66)$$

Here, the abbreviation $\vec{x}_0(t') = (t_{\text{em}} + \frac{L}{c} - t') \hat{p}$ is used. Thus, the time difference between two subsequent pulses consists mainly of the pulse period of the order of milliseconds and a deviation induced by the GWs, which is only marginal as the period of the GWs of interest is of order months to decades. Based on this observation, we calculate the Taylor expansion of the first integrand term to first order, which yields

$$\frac{\Delta P}{P} = \frac{p^a p^b}{2} \int_{t_{\text{em}}}^{t_{\text{em}} + \frac{L}{c}} dt' \left[\frac{\partial}{\partial t'} h_{ab}^{\text{TT}}(t', \vec{x}) \right]_{\vec{x}=\vec{x}_0(t')}. \quad (2.67)$$

We use this quantity to define the redshift $z(t) = \frac{\Delta P}{P}$ of the pulse arrival rates due to the influence of a GW. For the case of a fiducial monochromatic GW propagating in the direction of $\hat{\Omega}$, given by

$$h_{ab}^{\text{TT}}(t', \vec{x}) = \mathcal{A}_{ab}(\hat{\Omega}) \cos \left[\omega_{\text{GW}} \left(t - \frac{\hat{\Omega} \cdot \vec{x}}{c} \right) \right], \quad (2.68)$$

where $\mathcal{A}_{ab}(\hat{\Omega})$ contains the amplitude and polarisation structure of the wave, the expression for the redshift is evaluated to

$$\begin{aligned} z(t, \hat{\Omega}) &= \frac{1}{2} \frac{p^a p^b}{(1 + \hat{\Omega} \cdot \hat{p})} \left[h_{ab}(t, \vec{x}_{\text{earth}}) - h_{ab} \left(t - \frac{L}{c}, \vec{x}_{\text{pulsar}} \right) \right] \\ &= \frac{1}{2} \frac{p^a p^b}{(1 + \hat{\Omega} \cdot \hat{p})} \Delta h_{ab}. \end{aligned} \quad (2.69)$$

Here, the time of observation is abbreviated as $t = t_{\text{obs}}$. Two terms contribute to the redshift: the Earth term, evaluated at the coordinate origin $\vec{x}_{\text{earth}} = 0$ (the Solar System Barycentre), and the pulsar term, evaluated at $\vec{x}_{\text{pulsar}} = L\hat{p}$ and retarded by the light travel time. While the pulsar term depends on the individual line of

sight and is therefore uncorrelated between different pulsars, the Earth term is common to all of them. It is this shared contribution that enables a correlation between the timing residuals, whose characteristic angular dependence is derived in Sec. 2.3.3.

2.3.3 The Hellings–Downs Correlation

If the GW-induced redshift given in Eq. (2.69) were measured for a single pulsar, the GW contribution could not be detected. A variety of noise contributions, such as intrinsic spin noise, unmodelled variations of the dispersion measure, and inaccuracies in the reference clock and the Solar System ephemeris, mimic the effect of a GW and make the GW contribution indistinguishable in the signal of an individual pulsar. This issue is circumvented by searching for the correlated signal between many, or at least two, pulsars across the sky. Such a network of pulsars is called a *pulsar timing array* (PTA). While the noise contributions are independent between different pulsars and therefore average out in the cross-correlation, the common Earth term gives rise to a correlated signal. In this section, the analytic form of the expected correlation is derived for an isotropic SGWB, known as the *Hellings–Downs correlation*, whose detection would provide decisive evidence for the SGWB. The treatment presented in the following is once again based on [26].

Performing a Fourier-transformation of the redshift from Eq. (2.69) yields the expression

$$\tilde{z}(f, \hat{\Omega}) = \left(e^{-2\pi i f L c^{-1} (1 + \hat{\Omega} \cdot \hat{p})} - 1 \right) \sum_A h_A(f, \hat{\Omega}) F^A(\hat{\Omega}), \quad (2.70)$$

where the GW antenna response pattern of mode A is denoted as $F^A(\hat{\Omega})$, which we define as

$$F^A(\hat{\Omega}) = \frac{1}{2} \frac{p^a p^b}{(1 + \hat{\Omega} \cdot \hat{p})} e_{ab}^A(\hat{\Omega}). \quad (2.71)$$

We set up the correlation between two pulsars i and j by including contributions of GWs from all directions $\hat{\Omega}$ on the sky and computing the expectation value $\langle \tilde{z}_i(f) \tilde{z}_j^*(f') \rangle$. The expectation value $\langle \dots \rangle$ denotes an average over the statistical ensemble of the SGWB, of which our Universe represents a single realisation. Under the assumption of a Gaussian, stationary, isotropic and unpolarised SGWB,

the result can be expressed in terms of the cross-power spectral density $S_s(f)_{ij}$ of the redshifts measured for pulsars i and j :

$$\begin{aligned} \langle \tilde{z}_i(f) \tilde{z}_j^*(f') \rangle &= \frac{1}{2} \delta(f - f') S_s(f)_{ij} \\ &= \frac{1}{2} \delta(f - f') S_h(f) \int_{S^2} \frac{d^2 \hat{\Omega}}{8\pi} \kappa_{ij}(f, \hat{\Omega}) \sum_{A=+, \times} F_i^A(\hat{\Omega}) F_j^A(\hat{\Omega}), \end{aligned} \quad (2.72)$$

with

$$\begin{aligned} \kappa_{ij}(f, \hat{\Omega}) &= \left[e^{-2\pi i f L_i c^{-1}(1 + \hat{\Omega} \cdot \hat{p}_i)} - 1 \right] \left[e^{2\pi i f L_j c^{-1}(1 + \hat{\Omega} \cdot \hat{p}_j)} - 1 \right] \\ &= 1 - e^{-2\pi i f L_i c^{-1}(1 + \hat{\Omega} \cdot \hat{p}_i)} - e^{2\pi i f L_j c^{-1}(1 + \hat{\Omega} \cdot \hat{p}_j)} \\ &\quad + e^{2\pi i f c^{-1} [L_j(1 + \hat{\Omega} \cdot \hat{p}_j) - L_i(1 + \hat{\Omega} \cdot \hat{p}_i)]}, \end{aligned} \quad (2.73)$$

which controls the rate at which the pulsar term of the measured signal spatially decorrelates. This expression can be simplified to

$$\kappa_{ij}(f, \hat{\Omega}) \approx \begin{cases} 2, & i = j, \\ 1, & i \neq j, \end{cases} \quad (2.74)$$

which can be seen by considering the expanded form in Eq. (2.73). The exponential terms contain the factor $f L c^{-1} \gg 1$, which takes on large values due to the large distances to the pulsars. These terms therefore oscillate rapidly and do not contribute significantly to the integral, leaving only the constant term, $\kappa_{ij} \approx 1$. For the auto-correlation of a single pulsar, $i = j$, the exponents in the fourth term cancel exactly, so that a second constant contribution survives and we obtain $\kappa_{ii} \approx 2$. In the further treatment, we restrict the calculation to the cross-correlation between distinct pulsars and thus $\kappa_{ij} \approx 1$, leaving the integral in Eq. (2.72) frequency-independent. The expression obtained for the cross-power spectral density is thus given by

$$\begin{aligned} S_s(f)_{ij} &= \frac{1}{2} S_h(f) \int_{S^2} \frac{d^2 \hat{\Omega}}{4\pi} \kappa_{ij}(f, \hat{\Omega}) \sum_{A=+, \times} F_i^A(\hat{\Omega}) F_j^A(\hat{\Omega}) \\ &= \frac{1}{2} S_h(f) \tilde{\Gamma}_{ij}, \end{aligned} \quad (2.75)$$

where the geometric integral is identified as the unnormalised *overlap reduction function* (ORF), $\tilde{\Gamma}_{ij}$. This quantity describes the correlation between the timing residuals of two pulsars induced by the SGWB, as a function of their angular separation on the sky. First derived by Hellings and Downs in 1983 [32], the term can be evaluated to

$$\begin{aligned}\tilde{\Gamma}_{ij} &= \int_{S^2} \frac{d^2\Omega_{\hat{n}}}{4\pi} \sum_{A=+, \times} F_i^A(\hat{n}) F_j^A(\hat{n}) \\ &= x_{ij} \ln(x_{ij}) - \frac{1}{6}x_{ij} + \frac{1}{3},\end{aligned}\tag{2.76}$$

where we define $x_{ij} = [1 - \cos(\theta_{ij})]/2$ and θ_{ij} quantifies the angular separation between the two pulsars i and j . Nowadays, the common expression found in the literature includes the case of auto-correlation $i = j$ and normalises the ORF to $\Gamma_{ii} = 1$, which results in

$$\Gamma_{ij} = \frac{3}{2}x_{ij} \ln(x_{ij}) - \frac{1}{4}x_{ij} + \frac{1}{2} + \frac{1}{2}\delta_{ij}.\tag{2.77}$$

Its characteristic shape, which can be seen in Fig. 2.5, is quadrupolar, arising from the quadrupolar antenna response of the Earth–pulsar systems to a GW. The maximum correlation of 0.5 is reached for two distinct pulsars at the same sky position, $\theta_{ij} \approx 0^\circ$, whereas the auto-correlation of a single pulsar is given by a value of 1. As the angular separation increases, the correlation decreases rapidly and enters a regime of anti-correlation between $\theta_{ij} \approx 49.3^\circ$ and $\theta_{ij} \approx 121.8^\circ$, reaching a minimum of -0.152 . At an angular separation of $\theta_{ij} \approx 180^\circ$, corresponding to pulsars on opposite sides of the sky, the correlation returns to a value of 0.25.

Lastly, the importance of the quadrupolar shape of the Hellings–Downs curve should be emphasised. The SGWB is not the only source that can induce correlations between the signals of different pulsars. Errors in the modelling of the Solar System ephemeris and slow drifts in the global timing standards do so as well. The decisive difference is that these deviations are correlated with a dipolar or monopolar pattern, respectively, whereas the characteristic quadrupolar correlation is induced only by the SGWB. Therefore, this quadrupolar signature is the clearest criterion for the detection of the SGWB.




Figure 2.5: Visual representation of the Hellings–Downs correlation, given by the normalised ORF of an isotropic SGWB, measured with a PTA. Reproduced from [26].

Gravitational Waves from Domain Walls

3

In the previous chapter, two subjects were introduced independently: the physics of cosmological domain walls in Sec. 2.1, and the theory of gravitational waves, including the stochastic gravitational wave background, in Sec. 2.2. The goal of this chapter is to bring the two together. In Sec. 3.1, it is shown, based on [8, 26], that domain walls act as a source of gravitational waves. Building on this, the expected spectrum of GWs is discussed in Sec. 3.2, where we briefly review the current state of the art in modelling GW spectra from DWs in Sec. 3.2.1, present the template analysed in this work in Sec. 3.2.2 and compare the resulting spectrum to that employed in a prior NANOGrav analysis [14] in Sec. 3.2.3.

3.1 Domain Walls as Gravitational Wave Sources

In contrast to electromagnetic waves, where dipole radiation dominates, gravitational waves are sourced at leading order by the second time derivative of the mass quadrupole moment of a system [26]. The two lower multipole orders do not contribute: the mass monopole is proportional to the total energy, whereas the mass dipole is proportional to the centre-of-mass position, whose second time derivative is given by the rate of change of the total momentum. Conservation of energy and momentum thus ensures that the relevant derivatives vanish. The metric perturbation sourced by the quadrupole contribution is given by the *quadrupole radiation formula* [26]

$$h_{ij} = \frac{2G}{c^4 r} \Lambda_{ij,kl} \ddot{Q}_{kl}, \quad (3.1)$$

where r denotes the distance to the source, $\Lambda_{ij,kl}$ projects the metric perturbation, given in the Lorenz gauge, into the TT gauge and the reduced quadrupole moment of the source's mass density $\rho(t, \vec{x})$ is given by

$$Q_{ij} = \int d^3x \rho(t, \vec{x}) \left(x_i x_j - \frac{1}{3} |\vec{x}|^2 \delta_{ij} \right). \quad (3.2)$$

Inserting Eq. (3.1) into Eq. (2.49) and integrating over a sphere, we obtain the GW power in the quadrupole approximation, given by

$$P_{\text{GW}} \equiv \frac{dE_{\text{GW}}}{dt} = \frac{G}{5c^5} \left\langle \ddot{Q}_{ij} \ddot{Q}^{ij} \right\rangle. \quad (3.3)$$

Since a network of DWs constitutes a time-dependent and anisotropic distribution of energy and momentum, it has a time-varying quadrupole moment and thus emits GWs. A common estimate of the power emitted in GWs is computed with the quadrupole moment of a DW, given by $Q_{ij} \sim M_{\text{DW}} t^2$, where its mass-energy is $M_{\text{DW}} \sim \sigma \mathcal{A} t^2$ [8]. Setting natural units, $G = c = 1$, we can estimate

$$P_{\text{GW}} \sim \frac{M_{\text{DW}}^2}{t^2}. \quad (3.4)$$

Based on this, the GW energy density can be estimated as

$$\rho_{\text{GW}} \sim \frac{P_{\text{GW}} t}{H^{-3}} \sim \frac{P_{\text{GW}} t}{t^3} \sim \mathcal{A}^2 \sigma^2. \quad (3.5)$$

Thus, we can conclude that the energy density is primarily determined by the DW tension σ and remains nearly constant with respect to time. It should be noted that for DWs, the quadrupole formula does not directly apply, since it is derived for GWs in the far field approximation, assuming compact sources, while DWs reach horizon size during the scaling regime [8]. The validity of this estimate is verified by an alternative, numerical computation of the GW spectrum from the transverse-traceless part of the stress-energy tensor of the scalar field [21, 33].

Having established that DWs can act as a source of GWs, we are now interested in the emission spectrum in order to analyse it with the measured data. The standard approach in the literature is to compute it numerically, using lattice simulations [9, 12, 21, 22]. We review the modelling of spectra in the following Sec. 3.2.

3.2 Modelling the Spectrum

When reviewing the literature regarding GW spectra resulting from DWs, a wide range of different approaches, foci and resulting spectra can be found. The aim of Sec. 3.2.1 is to give a broad overview of the various topics discussed and recent developments in the literature, in order to make a well-founded decision on which model to analyse. The chosen GW template is presented in Sec. 3.2.2 and compared to that used in the analysis of the NANOGrav 15-year data set [14], in order to evaluate the prospect of finding deviating results.

3.2.1 State of the Art

In this section, the current state of the art in the development of GW templates from DWs is summarised, based on the findings of [9, 12, 20, 22]. The commonly chosen potential in the context of DWs is the discrete symmetric $\mathbb{Z}_2$ potential, which was introduced in Eq. (2.2). In addition to that, potentials of cosine form, related to axion models, are also studied, e.g. in [9]. To avoid an epoch of DW domination, modern DW studies include some kind of annihilation mechanism, of which the most widely used in the literature is the potential bias presented in Sec. 2.1.4. Modifications to this regular scenario include a bias term whose amplitude is dependent on temperature, as seen in [9, 12], or a population bias, where the two vacua remain degenerate, but the initial field distribution favours one of them, so that the network decays without an explicit energy difference [34].

Regarding the methodological approach, current DW research heavily relies on numerical lattice field simulations, where the commonly used code is ‘CosmoLattice’ [35, 36]. It discretises the scalar field ϕ on a comoving grid and evolves its equation of motion in an expanding background. A more detailed description of the procedure employed in the chosen template will be given in Sec. 3.2.2. For the extraction of the GW spectrum, it is common to evolve the metric perturbation directly as an additional field on the lattice [9, 22], while other works such as [12] instead compute it via Green functions using the transverse-traceless part of the scalar field stress-energy tensor [33]. A more recent approach is to extract the equal-time correlator of the energy-momentum tensor and use it to predict the GW spectrum semi-analytically [20].

The behaviour of the DW network during the scaling regime is typically monitored through the area parameter $\mathcal{A}$, whose value is consistently found to be of order unity. However, its concrete value varies noticeably across the different studies. While [12, 20] find it to be around $\mathcal{A} \sim 0.8$, [22] argues that discarding boundary effects of unphysical nature by suppressing infrared modes results in a genuine value for the area parameter of $\mathcal{A} \sim 1.2$. The value of [9] with $\mathcal{A} \sim 1$ lies right in between those two.

Furthermore, the most significant recent development in the field of DW simulation concerns the inclusion of the annihilation phase. While works such as [8, 21, 37] assume an instantaneous decay rather than including the annihilation phase in their simulations, [12] stresses the importance of doing so: the inclusion induces a shift of the peak frequency to lower values and raises the amplitude of the spectrum by a factor of $\mathcal{O}(100)$. Thus, according to current knowledge, the approximation of an instantaneous decay produces a GW spectrum that deviates significantly from the spectrum obtained when the annihilation phase is included. This is the core motivation for the development of a new GW template for DWs. Notably, the template used in the NG15 data analysis [14] is based on the treatment of Hiramatsu et al. [21] and thus relies on this approximation. A further comparison is presented in Sec. 3.2.3.

Among the works including the annihilation phase [9, 12, 22], the relation between the annihilation temperature T_{ann} and the bias amplitude ΔV is a current topic of discussion. The physically motivated expectation, following from the pressure balance $\sigma H_{\text{ann}} = \Delta V$, gives $T_{\text{ann}} \propto \Delta V^{1/2}$, which is the standard approximation applied by Cyr et al. [12] and Notari et al. [9]. Instead, Babichev et al. [38] reported $T_{\text{ann}} \propto \Delta V^{1/3}$ as a better match to the simulation results, supported by [22], implying an earlier annihilation and thus a smaller GW amplitude. This claim was tested in a recent study [34] using dedicated simulations with smaller bias sizes. There, a scaling of $T_{\text{ann}} \propto \Delta V^{0.59}$ was found, consistent with the original expectation, and the discrepancy with [38] was attributed to their networks not having reached the scaling regime before annihilation.

The GW template presented by Notari et al. [9] is adopted as the basis for the further analysis of this work. It includes the decay phase in the simulation, is consistent with annihilation occurring after the scaling regime has been reached, and yields

an area parameter roughly compatible with the findings on unphysical boundary effects reported in [22].

3.2.2 Spectral Template of This Work

In this section, the GW template by Notari et al. [9] is presented, considering how it was derived and its resulting parametrisation. It should be noted that the implemented potential slightly deviates from the one presented in Sec. 2.1.4 based on [8] in terms of the bias term, as it is realised in the form of

$$V = V_{\mathbb{Z}_2} + V_{\text{bias}} = \frac{\lambda}{4}(\phi^2 - v^2)^2 + qv\phi^3, \quad (3.6)$$

where $q \ll 1$ denotes a time-independent constant. While the form presented in [8] is constructed such that the minima remain located exactly at $\phi = \pm v$, the choice of Eq. (3.6) instead leaves the maximum of the potential at $\phi = 0$ unchanged, which is where the homogeneous mode of the field is initialised, so that no bias is introduced in the initial population of the two minima. The difference between the two minima is then given by

$$\Delta V = qv^4 \frac{(4\lambda^2 + 9q^2)^{\frac{3}{2}}}{4\lambda^3}. \quad (3.7)$$

The conventions used throughout their work are listed in the following. It is assumed that the expanding Universe is described by the Friedmann-Lemaître-Robertson-Walker (FLRW) metric [39],

$$ds^2 = c^2 dt^2 - a^2(t) \left(\frac{dr^2}{1 - kr^2} + r^2 d\Omega^2 \right), \quad (3.8)$$

where a is the cosmic scale factor, k the spatial curvature constant, and $d\Omega^2$ is defined as $d\Omega^2 \equiv d\theta^2 + \sin^2(\theta)d\phi^2$ ¹. Moreover, derivatives with respect to the conformal time η , defined through $d\eta = c dt/a(t)$, are abbreviated by a prime, $' \equiv d/d\eta$. The conformal Hubble rate is denoted by $\mathcal{H} \equiv a'/a$, such that the Hubble rate reads $H = a'/a^2 \equiv \mathcal{H}/a$. Natural units with $c = \hbar = 1$ are employed in the following.

¹Note that here, ϕ denotes the azimuthal angle and $d\Omega^2$ the solid angle element, whereas elsewhere ϕ refers to the scalar field and Ω to the density parameter.

3.2.2.1 Simulation Setup

The numerical work of simulating the DW network and extracting the GW spectrum was performed using a version of the code ‘CosmoLattice’ [35, 36] modified by the authors, which solves a discretised version of the equations of motion

$$\phi'' - \nabla^2 \phi + 2\mathcal{H}\phi' = -a^2 \partial_\phi V, \quad (3.9)$$

$$h''_{ij} - \nabla^2 h_{ij} + 2\mathcal{H}h'_{ij} = \frac{2}{M_p^2} (\partial_i \phi \partial_j \phi)^{\text{TT}}, \quad (3.10)$$

where h_{ij} denotes the metric perturbation in the TT gauge, sourced by the TT part of the anisotropic stress of the scalar field. The background is fixed to be radiation dominated at all times, which implicitly assumes the DWs to have a subdominant fraction of the total energy density. The system is discretised on a three-dimensional cubic lattice with $N = 3060$ points per dimension. Denoting by $\eta_{\Delta V}$ the conformal time at which the vacuum pressure equals the energy density of one Hubble-sized wall, $\sigma H(\eta_{\Delta V}) = \Delta V$, around which the network is expected to start deviating from the scaling regime, the accessible size of the bias term is constrained from two sides: the network has to reach the scaling regime before annihilation sets in, which requires a sufficiently large $\eta_{\Delta V}$, while a larger $\eta_{\Delta V}$ delays annihilation and thus demands longer simulations on larger lattices. The results presented in the following are therefore based on the single value $\eta_{\Delta V} = 14$. The spectrum is extracted at $\eta_f = 47$, after GW production has ceased to be efficient, and averaged over seven realisations of the initial conditions. The fit is restricted to wavenumbers $x \equiv k/(2\pi aH) \leq 50$, above which the spectrum is contaminated by lattice artefacts. Further details on the numerical implementation are provided in [9].

3.2.2.2 Spectral Parametrisation

The GW energy density spectrum at the final time of emission is parametrised as

$$\Omega_{\text{GW}}(f) = \Omega_{\text{GW}}(f_p) \times \mathcal{S}(f/f_p), \quad (3.11)$$

which separates the amplitude at the peak frequency f_p from the spectral shape function $\mathcal{S}(x)$, normalised to unity at $x_p \equiv k_p/(2\pi aH)$. Two shape functions are considered. The first is a singly broken power-law (BPL) of the following form,

$$\mathcal{S}(x) = \frac{(a+b)^c}{\left(b \left(\frac{x}{x_p}\right)^{-\frac{a}{c}} + a \left(\frac{x}{x_p}\right)^{\frac{b}{c}}\right)^c}, \quad (3.12)$$

where a denotes the spectral index for $x \ll x_p$ and b for $x \gg x_p$ respectively, while c quantifies the width around the maximum. The infrared slope is fixed to $a = 3$ by causality [40]. The second is a doubly BPL, which includes an additional point of steepening at a breaking wavenumber $x_b \gtrsim x_p$,

$$\mathcal{S}(x) = \frac{\alpha + \beta + (x_p/x_b)^{\beta+\gamma}}{\beta \left(\frac{x}{x_p}\right)^{-\alpha} + \alpha \left(\frac{x}{x_p}\right)^{\beta} + \left(\frac{x_b}{x_p}\right)^{-\beta} \left(\frac{x}{x_b}\right)^{\gamma}}, \quad (3.13)$$

with $\alpha = 3$, β and γ quantifying the slopes in the infrared, the intermediate and the ultraviolet regime, respectively. The amplitude is conveniently expressed through an efficiency parameter ϵ , which is defined relative to the normalised quadrupole [9] in the scaling regime evaluated for $\mathcal{A} = 1$,

$$\Omega_{\text{GW}}^{\text{quad}} \Big|_{\mathcal{A}=1} = \frac{3}{32\pi} \left(\frac{2\sigma H}{3H^2 M_p^2} \right)^2 \Big|_{\eta_f=47}. \quad (3.14)$$

Following a qualitative comparison of the data fits resulting from the two templates, which can be seen in Fig. 3.1, the doubly BPL from Eq. (3.13) was found to describe the numerical spectra better than the singly BPL from Eq. (3.12), in particular in the ultraviolet tail. Therefore, it is adopted in the following. The corresponding parameter values, resulting from the data fits, are listed in Tab. 3.1.

Table 3.1: Best-fit values for the parameters of the doubly broken power law template (3.13), taken from [9].

x_p	x_b	β	γ	ϵ
2.15 ± 0.19	6.2 ± 1.9	0.48 ± 0.15	1.79 ± 0.36	0.078 ± 0.006




Figure 3.1: Comparison of the data fits, resulting from the singly and doubly BPL, to the GW spectrum extracted from the simulation. Reproduced from [9].

3.2.2.3 Physical Parametrisation

In order to extract an interpretation regarding the nature of DWs, the peak frequency and amplitude of the template have to be related to physical quantities. The relation for the peak frequency is derived following [41].

A comoving mode k emitted at a temperature T is observed today at the frequency

$$f_0 = \frac{k}{2\pi a_0} = \frac{a(T)}{a_0} \cdot \frac{k}{2\pi a(T)}, \quad (3.15)$$

that is, the physical frequency at the time of emission, $k/(2\pi a)$, redshifted by the ratio of scale factors. The index 0 marks the values of a quantity on the present day.

The expression for the ratio of scale factors follows from the conservation of entropy, $g_{*,s}(T) T^3 a^3(T) = \text{const.}$, which yields

$$\frac{a(T)}{a_0} = \left(\frac{g_{*,s}(T_0)}{g_{*,s}(T)} \right)^{1/3} \left(\frac{T_0}{T} \right), \quad (3.16)$$

where $g_{*,s}$ denotes the effective number of entropy degrees of freedom and we use $T_0 \simeq 2.4 \cdot 10^{-13} \text{ GeV}$ as the photon temperature today. The frequency at emission is set by the Hubble rate at emission, which during radiation domination follows from the Friedmann equation $H^2 = \rho/(3M_p^2)$ with $\rho = (\pi^2/30)g_*(T)T^4$ as

$$H(T) = \frac{\pi}{\sqrt{90}} \frac{\sqrt{g_*(T)} T^2}{M_p}, \quad (3.17)$$

where g_* is the effective number of relativistic degrees of freedom. Since the spectral shape given in Eq. (3.13) is formulated in terms of the dimensionless wavenumber $x \equiv k/(2\pi aH)$, the peak of the spectrum corresponds to a physical wavenumber

$$\frac{k_p}{a} = 2\pi x_p H \quad (3.18)$$

at the time of emission, so that the factors of 2π in Eq. (3.15) cancel. Inserting Eqs. (3.16)–(3.18) into Eq. (3.15) and evaluating at $T = T_{\text{GW}}$, the temperature at which GW production ceases to be efficient, gives

$$f_p = x_p \frac{\pi}{\sqrt{90}} \frac{g_{*,s}^{1/3}(T_0) T_0}{M_p} \frac{\sqrt{g_*(T_{\text{GW}})}}{g_{*,s}^{1/3}(T_{\text{GW}})} T_{\text{GW}}. \quad (3.19)$$

Approximating $g_{*,s}(T) \simeq g_*(T)$, the remaining dependence on the degrees of freedom reduces to $g_*^{1/2-1/3} = g_*^{1/6}$, and normalising to representative values yields

$$f_p \simeq 24 x_p \text{ nHz} \left(\frac{g_*(T_{\text{GW}})}{100} \right)^{1/6} \left(\frac{T_{\text{GW}}}{150 \text{ MeV}} \right). \quad (3.20)$$

The amplitude of the signal is derived analogously. As stated in Eq. (3.14), the GW energy density at the time of emission is given by the quadrupole estimate, rescaled by an efficiency parameter ϵ determined from numerical simulations [8],

$$\Omega_{\text{GW}}(k_p)|_{\text{emit}} = \frac{3}{32\pi} \epsilon \alpha_{\text{GW}}^2, \quad \alpha_{\text{GW}} \equiv \frac{2\sigma H_{\text{GW}}}{3H_{\text{GW}}^2 M_p^2}, \quad (3.21)$$

where α_{GW} denotes the fraction of the total energy density stored in the DW network at T_{GW} . It follows from $\rho_{\text{DW}} = 2\mathcal{A}\sigma H$ together with the total energy density $\rho_{\text{tot}} = 3H^2 M_p^2$, where the area parameter is set to $\mathcal{A} = 1$. If the GW emission ceased during the era of radiation domination, the energy density of GWs at the present day can be estimated via

$$\Omega_{\text{GW}}(f_p)h^2 = \Omega_R h^2 \left(\frac{g_*(T_0)}{g_*(T_{\text{GW}})} \right)^{1/3} \Omega_{\text{GW}}(k_p)|_{\text{emit}}, \quad (3.22)$$

according to [21], where $\Omega_R h^2 \simeq 4.15 \times 10^{-5}$ denotes the present radiation density parameter with the reduced Hubble constant h and $g_*(T_0) = 3.36$ the corresponding number of relativistic degrees of freedom. Once again, inserting the obtained numerical values and normalising to representative values results in the following form,

$$\Omega_{\text{GW}}(f_p)h^2 \simeq 10^{-10} \epsilon \left(\frac{g_*(T_{\text{GW}})}{10.75} \right)^{-1/3} \left(\frac{\alpha_{\text{GW}}}{0.01} \right)^2. \quad (3.23)$$

As the shape parameters x_p , x_b , β and γ as well as the efficiency ϵ are fixed by the fit to the numerical spectra of [9], the spectrum defined by Eqs. (3.13), (3.20) and (3.23) depends on only two quantities: T_{GW} and α_{GW} . In the following, this parametrisation is referred to as *thermodynamic*, as it characterises the network purely through the state of the Universe at the time of emission, that is, the temperature at which GW production ceases and the fraction of the total energy density carried by the walls at that moment.

In order to obtain the *microscopic* parametrisation, which is instead given in terms of the underlying properties of the DWs themselves, both quantities are expressed through the DW tension σ and the energy difference ΔV between the two minima. The energy fraction, defined in Eq. (3.21), already depends on σ and on $H(T_{\text{GW}})$, so that only T_{GW} is left to be determined. To do so, the temperature $T_{\Delta V}$, at which the network starts deviating from the scaling regime, is obtained from the condition

of pressure balance $\sigma H = \Delta V$. Inserting the Hubble rate given in Eq. (3.17) and solving for the temperature yields

$$T_{\Delta V} = \left(\frac{\Delta V M_p}{\sigma \pi} \right)^{1/2} \left(\frac{90}{g_*(T_{\Delta V})} \right)^{1/4}, \quad (3.24)$$

where g_* has to be evaluated iteratively, as it depends on the temperature that is to be determined. Combining this with the delay of the end of GW emission, $T_{\text{GW}} \simeq 0.3 T_{\Delta V}$, found by Notari et al. [9], and normalising to representative values reproduces their result,

$$T_{\text{GW}} \simeq 133 \text{ MeV} \left(\frac{g_*(T_{\text{GW}})}{17.25} \right)^{-1/4} \left(\frac{\Delta V^{1/4}}{100 \text{ MeV}} \right)^2 \left(\frac{10^5 \text{ GeV}}{\sigma^{1/3}} \right)^{3/2}. \quad (3.25)$$

With this, we have introduced the two parametrisations that form the core of the analysis of this thesis.

3.2.3 Comparison to Previous Template

We now review the changes from the GW template employed in the NG15 data analysis [14] to the template considered in this thesis by Notari et al. [9]. The core equations of the former are given by

$$h^2 \Omega_{\text{GW}}(f) = \frac{3}{32\pi} \mathcal{D} \tilde{\epsilon} \alpha_*^2 \mathcal{S}(f/f_p), \quad (3.26)$$

$$f_p = 1.14 \text{ nHz} \left(\frac{10.75}{g_{*,s}} \right)^{1/3} \left(\frac{g_*}{10.75} \right)^{1/2} \left(\frac{T_*}{10 \text{ MeV}} \right), \quad (3.27)$$

$$\mathcal{S}(x) = \frac{(a+b)^c}{(b x^{-a/c} + a x^{b/c})^c}. \quad (3.28)$$

Here, the factor $\mathcal{D}$ accounts for the redshift of the GW amplitude and $\tilde{\epsilon}$ is an efficiency parameter fixed from numerical simulations. Quantities marked with a star in the index are evaluated at the temperature T_* , which is introduced in [14] as the temperature at which the DW network decays, fixing the peak frequency. For a

comparison of the two expressions for the peak frequency, we adopt the reference values of Eq. (3.20) in Eq. (3.27), yielding

$$\begin{aligned} f_p &= 1.14 \text{ nHz} \left(\frac{100}{g_{*,s}} \right)^{1/3} \left(\frac{10.75}{100} \right)^{1/3} \left(\frac{g_*}{100} \right)^{1/2} \left(\frac{100}{10.75} \right)^{1/2} \cdot 15 \left(\frac{T_*}{150 \text{ MeV}} \right) \\ &\approx 24.80 \text{ nHz} \left(\frac{100}{g_{*,s}} \right)^{1/3} \left(\frac{g_*}{100} \right)^{1/2} \left(\frac{T_*}{150 \text{ MeV}} \right). \end{aligned} \quad (3.29)$$

Comparing Eq. (3.29) to Eq. (3.20) and using $g_{*,s}(T) \simeq g_*(T)$, both expressions share the same dependence on g_* and T , and their numerical prefactors agree to within about 3.5 %. The only structural difference is the factor $x_p \simeq 2.15$ appearing in Eq. (3.20). According to Notari et al. [9], this factor arises because they compute the position of the peak rather than fixing it, as earlier analyses did, to the Hubble frequency at the temperature T_{GW} at which efficient GW production ceases. Its absence in Eq. (3.27) shows that the template of [14] belongs to the latter class of analyses. In addition, T_* is introduced there as the temperature at which the network decays, while it occupies the position of T_{GW} in Eq. (3.20), being the temperature at which efficient GW production ceases. The two temperatures are thus not distinguished, whereas Notari et al. [9] treat them separately.

Apart from that, the spectral shape $\mathcal{S}(x)$ used by Afzal et al. [14], given in Eq. (3.28), is a singly BPL, whereas Notari et al. [9] mainly study the doubly BPL, given in Eq. (3.13), following a comparison of those two. As can be seen in Fig. 3.1, the doubly BPL reproduces the spectrum, simulated by Notari et al. [9], more accurately. It follows the data more closely on the rising side of the peak and at the maximum itself. The deviation of the singly BPL becomes most noticeable in the ultraviolet tail. Therefore, the doubly BPL provides the better description of the simulated spectrum and is the form adopted in this work.

Analogous to Notari et al. [9], the resulting GW spectrum in [14], given in Eq. (3.26), is constructed from the peak amplitude at the time of emission, modelled by a spectral shape function. Notari et al. [9] use the explicitly calculated formula to account for the redshift, given in Eq. (3.23), whereas this contribution is stored in a factor $\mathcal{D}$ in Afzal et al. [14]. In the following, the definition of $\mathcal{D}$ given in [14] is rewritten into a form enabling a comparison with [9]. Inserting the present-day radiation density

parameter, $\Omega_R = \frac{\pi^2}{90} g_*(T_0) T_0^4 / (M_p^2 H_0^2)$ [42], and using $g_{*,s}(T) \simeq g_*(T)$ at T_{GW} , one obtains

$$\mathcal{D} = \frac{\pi^2}{90} \frac{T_0^4}{M_p^2 H_0^2} g_* \left(\frac{g_{*,s}^{\text{eq}}}{g_{*,s}} \right)^{4/3} \simeq \Omega_R \frac{(g_{*,s}^{\text{eq}})^{4/3}}{g_*(T_0)} g_*^{-1/3}(T_{\text{GW}}). \quad (3.30)$$

The corresponding redshift factor of Eq. (3.22) reads

$$\mathcal{D}_{\text{Notari}} = \Omega_R h^2 g_*^{1/3}(T_0) g_*^{-1/3}(T_{\text{GW}}), \quad (3.31)$$

so that both expressions share the same dependence $\propto g_*^{-1/3}(T_{\text{GW}})$. Apart from the factor h^2 , which simply reflects that $\mathcal{D}$ is defined for Ω_{GW} instead of $h^2 \Omega_{\text{GW}}$, the two prefactors are equivalent if $g_{*,s}^{\text{eq}} = g_*(T_0)$. A numerical evaluation yields $g_{*,s}^{\text{eq}}/g_*(T_0) \approx 1.16$, so that the prefactors differ by $(g_{*,s}^{\text{eq}}/g_*(T_0))^{4/3} \approx 1.22$. The two treatments of the redshift therefore deviate by about 22 %. Most of this deviation is absorbed by the numerical prefactor in Eq. (3.23), leaving a residual rounding error of about 2.5 % relative to the exact result obtained from the derivation.

Another striking difference lies in the values of the efficiency factor obtained from the respective simulations. While Afzal et al. [14] use $\tilde{\epsilon} = 0.7$, taken from the simulations of Hiramatsu et al. [21], the efficiency factor of Notari et al. [9] is $\epsilon = 0.078$, differing by a factor of about 9. Since the efficiency enters the amplitude linearly and all remaining prefactors of the two templates agree to within the deviations discussed above, this difference directly influences the predicted GW amplitude. Evaluating the amplitude for identical parameter values, the template of [14] yields a signal larger by roughly an order of magnitude. This means that the reproduction of a given observed amplitude with the template used in this work requires a value of α_{GW} larger by a factor of about $\sqrt{9} = 3$ in comparison to α_* in [14], as the amplitude scales quadratically with α_{GW} . It should be kept in mind that the two quantities α_{GW} and α_* , while defined in the same way as the fraction of the total energy density stored in the DW network, are evaluated at different moments, as discussed above. Furthermore, it is important to stress that the simulations by Hiramatsu et al. [21] are performed without the bias term and do not include the collapse of the DWs. The coefficients are extracted from the scaling regime alone and extrapolated under the assumption of an instantaneous disappearance of the network. The efficiency factor of [14] therefore rests on an extrapolation from the scaling regime, whereas [9] determine it from simulations that follow the network through its collapse.

With all these differences in mind, the methodology and an overview of the performed analyses are presented in Chapter [4](#), before the results obtained by applying the new GW template to the data are presented and discussed in Chapter [5](#).

Methodology

4

In this chapter, we introduce the methodology underlying the analyses of this work. Sec. 4.1 presents the PTA analysis tool PTArcade [10, 11] and the statistical framework it builds on. Then, Sec. 4.2 specifies the analyses carried out.

4.1 Analysis Framework

PTArcade [10, 11] is a wrapper for the Python-based PTA analysis software packages Ceffyl [43] and ENTERPRISE [44, 45]. All analyses in this work are carried out in ENTERPRISE mode, evaluating the full PTA likelihood, and the description in the following sections refers to this mode throughout. We use it to probe PTA data for a contribution from DWs through Bayesian inference. This is done by implementing the GW spectrum $\Omega_{\text{GW}}(f; \vec{\theta})$ in Python, where $\vec{\theta}$ denotes the vector containing all free parameters the GW spectrum depends on. For a given model, PTArcade explores the parameter space by drawing samples according to the posterior, from which marginalised distributions and upper limits for the free parameters are obtained.

The following sections present in detail how this is achieved, mainly following the treatment in [11, 26]. In Sec. 4.1.1, the PTA likelihood, i.e. the probability of observing a given set of timing residuals for a given set of parameters, is constructed. This quantity is needed in Sec. 4.1.2, where it is embedded into the Bayesian framework from which the posterior distributions follow. Sec. 4.1.3 describes how these distributions are sampled numerically.

4.1.1 PTA Likelihood

As presented in Sec. 2.3.1, the central quantity measured by PTAs is the timing residual δt , which is given as the difference between the measured TOA and the one

predicted from the timing model. Such a discrepancy between the two quantities is mostly generated either by noise or a gravitational wave background. In PTArcade, the contributions to the timing residuals are modelled as

$$\vec{\delta t} = \vec{n} + \mathbf{F} \vec{a} + \mathbf{M} \vec{e}. \quad (4.1)$$

Here, $\vec{n}$ denotes a vector of remaining white noise contributions after subtraction of all known systematic influences, where each component corresponds to the white noise of one respective timing residual. White noise is modelled as a zero-mean normal random variable and is therefore fully characterised by its covariance matrix $\mathbf{N}$.

The second term, $\mathbf{F} \vec{a}$, contains time-correlated stochastic processes, such as the pulsar-intrinsic red noise and the contribution of the SGWB. Both are expanded in a Fourier basis with frequencies i/T_{obs} , where i denotes the index of the harmonic in the basis and T_{obs} quantifies the duration of the measurement, spanning all included timing residuals. A truncation is performed after N_f bins for each of the two individual contributions, such that the Fourier design matrix $\mathbf{F}$ holds the corresponding sine–cosine pairs evaluated at the observation times. The coefficients $\vec{a}$ are normally distributed with vanishing mean and covariance

$$[\phi]_{(ak)(bj)} = \delta_{ij} (\Gamma_{ab} \Phi_i + \delta_{ab} \varphi_{a,i}), \quad (4.2)$$

where a, b index the pulsars, i, j the frequency harmonics, and Γ_{ab} is the overlap reduction function of Sec. 2.3.3, reducing to the Hellings & Downs correlation for an isotropic, unpolarised background. The coefficients Φ_i carry the GWB signal and relate to the GW energy density via

$$h^2 \Omega_{\text{GW}}(f) \equiv \frac{h^2}{\rho_c} \frac{d\rho_{\text{GW}}}{d \ln f} = \frac{8\pi^4 f^5}{H_0^2/h^2} \frac{\Phi(f)}{\Delta f}, \quad (4.3)$$

with $\Delta f = 1/T_{\text{obs}}$ and $H_0 = h \times 100 \text{ km s}^{-1} \text{ Mpc}^{-1}$. PTArcade constructs $h^2 \Omega_{\text{GW}}$ directly from the spectrum given in the user model file. The remaining coefficients

$\varphi_{a,i}$ describe the intrinsic red noise of each pulsar, modelled as a power law with amplitude A_a and spectral index ξ_a :

$$\varphi_a(f) = \frac{A_a^2}{12\pi^2} \frac{1}{T_{\text{obs}}} \left(\frac{f}{\text{year}^{-1}} \right)^{-\xi_a} \text{year}^3. \quad (4.4)$$

The third term, $\mathbf{M} \vec{\epsilon}$, includes the imperfect knowledge of the m timing model parameters. Since these are determined by a fit prior to the formation of the residuals, they carry a small offset $\vec{\epsilon}$ from their true values, which in turn shifts the predicted TOAs. To linear order, this shift is governed by the derivatives of the TOAs with respect to the individual parameters, collected in the $N_{\text{TOA}} \times m$ design matrix $\mathbf{M}$ and evaluated at the best-fit point. Marginalising over all possible noise realisations yields the following expression for the likelihood of the parameter set $\vec{\theta}$ given the observed timing residuals:

$$p(\vec{\delta t} | \vec{\theta}) = \frac{\exp\left(-\frac{1}{2} \vec{\delta t}^T \mathbf{C}^{-1} \vec{\delta t}\right)}{\sqrt{\det(2\pi\mathbf{C})}}. \quad (4.5)$$

Here, $\mathbf{C} = \mathbf{N} + \mathbf{T}\mathbf{B}\mathbf{T}^T$ with $\mathbf{T} = [\mathbf{M}, \mathbf{F}]$, $\mathbf{B} = \text{diag}(\infty, \phi)$ and ∞ being the diagonal matrix of infinities, corresponding to the assumption of flat priors for the parameters of $\vec{\epsilon}$.

4.1.2 Bayesian Statistics

The most familiar interpretation of probability is the frequentist one, in which a probability corresponds to the frequency with which a certain outcome appears over many repetitions of an experiment [46]. The probability of a coin toss yielding heads is then given by the ratio of the number of heads N_{heads} to the total number of tosses N in the limit of infinitely many repetitions,

$$p(\text{heads}) = \lim_{N \rightarrow \infty} \frac{N_{\text{heads}}}{N}, \quad (4.6)$$

which yields $p = 1/2$ for a fair coin. Probability is thus a property of the experiment itself, and only the data are treated as random.

Such a definition is of limited use in cosmology and astrophysics, where the experiment cannot be repeated, e.g. a given black hole merger happens once, and we observe a single realisation of the Universe. What we want to quantify instead is how strongly a given data set supports a particular model or parameter value. This is the approach of the *Bayesian* framework, in which a probability expresses the degree of belief in a hypothesis, constructed on the basis of prior knowledge [46]. The core relation is Bayes' theorem, which, applied to the analysis of timing residuals $\vec{\delta t}$, is given by

$$p(\vec{\theta} | \vec{\delta t}) = \frac{p(\vec{\delta t} | \vec{\theta}) p(\vec{\theta})}{p(\vec{\delta t})}. \quad (4.7)$$

The central quantity we want to extract is the *posterior* $p(\vec{\theta} | \vec{\delta t})$, which is the probability of the parameters $\vec{\theta}$ describing the underlying GW spectrum given the measured residuals $\vec{\delta t}$. In addition to the PTA likelihood $p(\vec{\delta t} | \vec{\theta})$ of Eq. (4.5), Bayesian inference requires a *prior* $p(\vec{\theta})$, encoding the knowledge about the parameters before the measurement. The *evidence* is defined as

$$\mathcal{Z} \equiv p(\vec{\delta t}) = \int d\vec{\theta} p(\vec{\delta t} | \vec{\theta}) p(\vec{\theta}). \quad (4.8)$$

Since $\vec{\theta}$ is high-dimensional, statements about an individual parameter θ_i are obtained by marginalising over the remaining ones [26],

$$p(\theta_i | \vec{\delta t}) = \int d\vec{\theta}_{j \neq i} p(\vec{\theta} | \vec{\delta t}), \quad (4.9)$$

which is how the one-dimensional posteriors along the diagonal of the corner plots shown in Chapter 5 are constructed.

The integrals in Eqs. (4.8) and (4.9) are computationally expensive to calculate, given the typical high dimensionality of the considered parameter space. This is the reason that in PTArcade an evaluation via Markov chain Monte Carlo (MCMC) sampling is performed, which is introduced in Sec. 4.1.3.

4.1.3 Markov Chain Monte Carlo Methods

The basis of MCMC sampling is the ergodic theorem for Markov chains [26]. It states that a chain constructed such that the distribution of interest is its stationary

distribution will, after a sufficiently large number of iterative steps, sample from that distribution, independent of the initial conditions. This is exploited by the MCMC method, e.g. in the so-called Metropolis–Hastings algorithm, which proposes new points in parameter space and decides whether to move there based on the Metropolis–Hastings ratio

$$\text{MH} = \frac{p(\vec{y}) p(\vec{\delta}t \mid \vec{y})}{p(\vec{x}_{i-1}) p(\vec{\delta}t \mid \vec{x}_{i-1})} \times \frac{q(\vec{x}_{i-1} \mid \vec{y})}{q(\vec{y} \mid \vec{x}_{i-1})}. \quad (4.10)$$

Here, $\vec{x}_{i-1}$ denotes the current position of the chain, $\vec{y}$ the newly proposed point, and $q(\vec{y} \mid \vec{x}_{i-1})$ the proposal distribution from which $\vec{y}$ is drawn. The first factor is the ratio of prior times likelihood at the two points, i.e. the ratio of the corresponding posterior densities, in which the evidence $\mathcal{Z}$ cancels, making the posterior accessible without evaluating the integral in Eq. (4.8). The second factor corrects for an asymmetric proposal distribution and reduces to unity for symmetric q .

The proposal is accepted if $\text{MH} > u$, where $u \in U[0, 1]$ is drawn from a uniform distribution, in which case we set $\vec{x}_i = \vec{y}$. Otherwise the former point is kept, $\vec{x}_i = \vec{x}_{i-1}$. A proposal with $\text{MH} > 1$ is thus always taken, while one with $\text{MH} \leq 1$ is accepted with probability MH .

Starting from an initial point $\vec{x}_1$, the sampler first meanders through parameter space searching for regions of high probability. This initial phase, the so-called *burn-in*, heavily relies on the starting point and is discarded. The remaining samples are, in the limit of long chains, distributed according to the posterior itself, so that marginal distributions are obtained by arranging the samples in the parameters of interest in a histogram.

4.2 Overview of the Performed Analyses

The core of this thesis is to use the recently developed GW template of Notari et al. [9] to constrain the free parameters on which the DW-induced GW spectrum depends, using PTA data from the North American Nanohertz Observatory for Gravitational Waves. This is done through Bayesian inference on the template presented in Sec. 3.2.2, implemented into PTArcade [10, 11]. The realisation of this implementation in the form of a model file is discussed in Appendix A. The analyses performed

can be divided into two categories, which are distinguished by the role of the DW contribution within the total GW spectrum. Before introducing them, we summarise the technical setup in Sec. 4.2.1.

4.2.1 Analysis Setup

All of the results presented in Chapter 5 were obtained using PTArcade [10, 11] from runs evaluating the NG15 data set [2]. PTArcade was run in ENTERPRISE mode [44, 45], in which the analysis is carried out on the basis of timing residuals, using the full PTA likelihood of Eq. (4.5). Inter-pulsar correlations are included in the likelihood, meaning that the overlap reduction function is set to the Hellings–Downs correlation. The intrinsic red noise of each of the 67 included pulsars is modelled with 30 frequency components, while the GW background is described by the 14 lowest frequency bins. Posterior sampling is carried out with PTMCMCSampler [47] on the PALMA II cluster of the University of Münster. For each configuration, 20 independent chains are run in parallel and up to 2×10^6 samples per chain are drawn. The chains are automatically thinned by a factor of ten, while the first quarter of each chain is discarded as burn-in, and the remainder are merged. The number of samples varies between configurations, as they do not all require the same amount of sampling to reach convergence. Convergence is assessed via the split- $\hat{R}$ statistic, which is reported for each parameter. It quantifies the degree to which the chains have settled into a common stationary distribution by comparing the variance between the individual chains and their halves to the variance within them, approaching unity in the case of full convergence. Following Ref. [48], chains are regarded as converged only if $\hat{R} < 1.01$ for every parameter.

4.2.2 Signal Runs

In the ‘signal runs’ we assume that the entire observed spectrum originates from DWs. Both parametrisations introduced in Sec. 3.2.2.3 are probed, with two variants each. In the first variant, the spectral shape parameters are fixed to the best-fit values of Tab. 3.1, leaving T_{GW} and α_{GW} as the only free parameters in the thermodynamic parametrisation, and σ and ΔV in the microscopic one. For all four a log-uniform prior is adopted, i.e. a uniform prior on the logarithm of the parameter, reflecting that the quantities may span several orders of magnitude and that no decade within

the sampled range is preferred beforehand. In the microscopic case, we sample the rescaled variables

$$\theta_\sigma = \frac{1}{3} \log_{10}(\sigma/\text{GeV}^3) , \quad (4.11)$$

$$\theta_{\Delta V} = \frac{1}{4} \log_{10}(\Delta V/\text{GeV}^4) , \quad (4.12)$$

such that $10^{\theta_\sigma} = \sigma^{1/3}/\text{GeV}$ and $10^{\theta_{\Delta V}} = \Delta V^{1/4}/\text{GeV}$ both correspond to energy scales in GeV, making the two parameters directly comparable in magnitude. From the resulting constraints we aim to infer the state of the Universe and of the DW network at the time when efficient GW production ceases.

In the second variant, all parameters of Tab. 3.1 with a relative uncertainty larger than 20 %, namely x_b , β and γ , are included as free parameters. In contrast to the priors above, they follow Gaussian priors centred on their best-fit values with the standard deviation given by the corresponding uncertainties. This reflects the available prior knowledge from the fit to the simulated data by Notari et al. [9]. Through the inclusion of free spectral indices, we test whether the PTA data have any constraining power on the spectral shape.

4.2.3 Upper-Limit Runs

The second class of analyses, which we call ‘upper-limit runs’, is built on the same models, but adds a *free spectrum* (FS) component, which changes the underlying philosophy of the analysis. The SGWB is expected to contain, to some unknown degree, a contribution from inspiralling supermassive black hole binaries (SMBHBs) [7]. Further contributions of astrophysical or cosmological origin are plausible. The assumption underlying the signal runs is therefore an idealisation that will almost certainly not capture the full picture. We thus invert the logic: the FS is constructed such that it is capable of accounting for the entire PTA signal on its own, without reference to any specific source. Being agnostic to its particular constituents, it includes any additional contribution apart from the DWs.

The NG15 data set provides amplitude estimates for the first 14 harmonics $f_i = i/T_{\text{obs}}$, where $i = 1, \dots, 14$, of the inverse observational baseline $T_{\text{obs}} \approx 16.03 \text{ yr}$ [2]. Higher frequency bins are dominated by pulsar white noise and are therefore not included. Thus, we characterise the FS by 14 amplitude parameters, one per frequency bin.

The DW and FS contributions are then added and sampled in the same way as in the signal runs. This allows us to check what contribution DWs can have on top of an otherwise fully explained SGWB and whether its inclusion is favoured at all.

For the 14 amplitude parameters we compare two choices regarding the prior shape, as can be seen in Tab. 4.1. While the common log-uniform choice assigns equal weight to every decade of the sampled range, the alternative prior shape *LinearExp* is uniform in the amplitude itself and thus places most of its weight in the highest decade. A deviation between the two prior choices is expected, as the prior is included in Bayes' theorem calculating the posterior distribution, and we aim to qualitatively capture the form of this deviation.

In total, eight runs are performed. For each of the two parametrisations we carry out two signal runs: a two-parameter run in which the spectral shape parameters are fixed to their best-fit values in Tab. 3.1 and a five-parameter run in which x_b , β and γ are included into the list of free parameters as well. In addition to that, two upper-limit runs for each parametrisation are performed, differing in the prior shape of the free spectrum amplitudes. All free parameters as well as their prior distributions employed in this work are summarised in Tab. 4.1.

Table 4.1: Prior distributions used in the analyses. All spectral shape parameters are fixed at the best-fit values of Tab. 3.1, except for x_b , β and γ in the two five-parameter runs. For Gaussian priors, μ denotes the central value and s the standard deviation.

Parameter	Prior	Range / distribution parameters
<i>Thermodynamic parametrisation</i>		
$\log_{10}(T_{\text{GW}}/\text{GeV})$	Uniform	$[-3, 0]$
$\log_{10}(\alpha_{\text{GW}})$	Uniform	$[-4, 0]$
<i>Microscopic parametrisation</i>		
$\theta_\sigma = \frac{1}{3} \log_{10}(\sigma/\text{GeV}^3)$	Uniform	$[3, 7]$
$\theta_{\Delta V} = \frac{1}{4} \log_{10}(\Delta V/\text{GeV}^4)$	Uniform	$[-3, 1]$
<i>Spectral shape</i>		
x_b	Gaussian	$\mu = 6.2, s = 1.9$
β	Gaussian	$\mu = 0.48, s = 0.15$
γ	Gaussian	$\mu = 1.79, s = 0.36$
<i>Free spectrum</i>		
$\log_{10}[h^2\Omega_{\text{GW}}(f_i)], i = 1, \dots, 14$	Uniform	$[-14, -5]$
$\log_{10}[h^2\Omega_{\text{GW}}(f_i)], i = 1, \dots, 14$	LinearExp ^a	$[-14, -5]$

^a Uniform in $h^2\Omega_{\text{GW}}$ rather than in its logarithm, with a probability density $p(x) \propto \ln(10) 10^x$ [11] for $x = \log_{10}(h^2\Omega_{\text{GW}})$.

Results and Discussion

In this chapter, the results of the analyses outlined in Sec. 4.2 are presented and discussed. We visualise the obtained results in corner plots: along the diagonal, the one-dimensional posterior distribution of the parameter labelled below is shown, marginalised over all remaining parameters, whereas the panels below the diagonal display the posterior distributions of the corresponding parameter pairs, with contours enclosing the 68 % and 95 % credible regions [10].

Alongside these plots, PTArcade provides a summary of statistical properties for each parameter. First, the upper and lower bounds of the 68 % and 95 % credible intervals, given as the highest posterior density intervals (HPDI), are reported. Second, the Bayes estimator $\hat{\theta}$ together with its standard deviation is given. It is defined as the posterior mean

$$\hat{\theta} = \langle \theta \rangle = \int \theta p(\theta | \vec{\delta t}) d\theta, \quad (5.1)$$

which PTArcade evaluates as the mean over the merged samples after burn-in¹. Third, we report the maximum posterior (MAP), which is given as the parameter value at which the posterior density is at its maximum. These two estimators coincide only if the posterior is symmetric about its maximum, and may differ otherwise. Both are therefore quoted, and all values are collected in tables for each model configuration.

In the following, we first evaluate the results of the performed signal runs for the thermodynamic and microscopic parametrisations, with and without including the spectral parameters, in Sec. 5.1. Then, we turn to the analysis of the upper-limit runs in Sec. 5.2, where the influence of a varying prior shape is at the centre of focus.

¹See the implementation of `get_bayes_est` in the `chains_utils` module [49].

5.1 Signal Runs

In this section, we analyse the results of the performed signal runs, in which DWs are the only contribution to the GW spectrum. First, we discuss the results obtained for the thermodynamic parametrisation in Sec. 5.1.1. The same is done in Sec. 5.1.2 for the microscopic parametrisation. Lastly, we review the effect of including the spectral variables β, γ and x_b as free parameters in Sec. 5.1.3.

5.1.1 Thermodynamic Parameters

For the thermodynamic parametrisation, a uniform prior on $\log_{10}(T_{\text{GW}}/\text{GeV})$ is adopted over the range $[-3, 0]$, corresponding to temperatures between 1 and 1000 MeV. Since the peak frequency is entirely determined by T_{GW} , as can be seen in Eq. (3.20), this range is chosen such that it covers the temperatures for which the spectrum peaks in the nanohertz band probed by PTA searches. For $\log_{10}(\alpha_{\text{GW}})$, the prior is likewise chosen to be uniform. Its range is set to $[-4, 0]$, corresponding to 0.01 % to 100 % of the Universe's energy density stored in DWs at T_{GW} and thus includes the possibility of a negligibly small DW contribution as well as a DW-dominated Universe. At this point, we should note that the template of Notari et al. [9] assumes a radiation-dominated background, as it enters e.g. Eq. (3.22), and thus becomes less accurate the larger the contribution of the DWs to the total energy density is. This restriction is not imposed on the numerical realisation in PTArcade, so that the posterior of α_{GW} provides a consistency check on the applicability of the template. The resulting posterior distributions are shown in Fig. 5.1, whereas the corresponding statistical summary is listed in Tab. 5.1.

As can be seen in Fig. 5.1, both parameters are well constrained, with unimodal and nearly symmetric posteriors. A slight asymmetry remains visible, which is supported by the asymmetric credible intervals around the Bayes estimator listed in Tab. 5.1. The split- $\hat{R}$ values lie below the threshold of 1.01 for both parameters, indicating that the chains have converged. Both posteriors close within the prior range, so that the constraints are not cut off by the prior boundaries.

The two-dimensional posterior in the bottom left panel shows a clear correlation between the two parameters, as the contour is visibly tilted and elongated, meaning that a change in one parameter can be compensated by a corresponding change in




Figure 5.1: Corner plot, which visualises the one- and two-dimensional posterior distributions for the thermodynamic parameters $\log_{10}(T_{\text{GW}}/\text{GeV})$ and $\log_{10}(\alpha_{\text{GW}})$.

the other without leaving the region of high posterior probability. To explain this, we first clarify the individual roles of each parameter in shaping the GW spectrum. While the peak frequency is set by T_{GW} alone, as can be seen in Eq. (3.20), the peak amplitude is governed almost entirely by α_{GW} , entering quadratically in Eq. (3.23). The remaining dependence of the amplitude on T_{GW} through $g_*(T_{\text{GW}})$ is weak, as it is less than a factor of two in $h^2\Omega_{\text{GW}}$ across the full prior range, whereas α_{GW} varies the amplitude by eight orders of magnitude over its prior. Both thermodynamic parameters thus govern one of the two spectral quantities, so that each spectral feature can, in principle, be adjusted independently. Therefore, the observed degeneracy does not arise because the spectrum depends on some effective combination of T_{GW} and α_{GW} . Instead, the degeneracy originates from the finite frequency band in which the data constrain the spectrum. The 95 % credible interval of T_{GW} corresponds to

Table 5.1: Statistical summary of the thermodynamic two-parameter signal run. All parameters satisfy $\hat{R} < 1.01$. The credible intervals are given as the highest posterior density intervals (HPDI). The maximum posterior is abbreviated as MAP. The lower block lists the corresponding values on a linear scale, which are obtained by transforming the sampled logarithmic quantities.

Parameter	68 % HPDI	95 % HPDI	Bayes estimator	MAP
$\log_{10}(T_{\text{GW}}/\text{GeV})$	$[-1.289, -1.033]$	$[-1.445, -0.922]$	-1.175 ± 0.132	-1.153
$\log_{10}(\alpha_{\text{GW}})$	$[-0.642, -0.344]$	$[-0.805, -0.188]$	-0.494 ± 0.154	-0.491
$T_{\text{GW}} / \text{MeV}$	$[51.4, 92.6]$	$[35.9, 119.6]$	$66.8^{+23.8}_{-17.5}$	70.3
α_{GW}	$[0.228, 0.452]$	$[0.157, 0.648]$	$0.320^{+0.136}_{-0.096}$	0.323

peak frequencies between roughly 9 and 32 nHz via Eq. (3.20), so that the peak lies within or close to the observed band throughout. Within this range, a shift of the peak can be partially compensated by adjusting the overall amplitude. For example, a shift of the peak frequency towards a higher value corresponds to a lower amplitude at a fixed frequency in the low-frequency tail of the spectrum, which can be compensated by shifting the peak amplitude and thus the entire spectrum towards higher values. However, this compensation is not exact. On the high-frequency side of the peak both changes act in the same direction, as shifting the peak frequency and amplitude towards higher values increases the predicted signal in this regime. The upper frequency bins therefore limit how far the two parameters can be traded against each other, which is reflected in the finite extent of the contour along the direction of the degeneracy.

Next, we consider the physical meaning of the resulting values for both parameters. Most notably, the resulting temperature allows the end of efficient GW emission to be placed within the thermal history of the Universe. With $T_{\text{GW}} = 66.8^{+23.8}_{-17.5}$ MeV, the GW emission from DWs ceases after the quantum chromodynamics (QCD) crossover, which lattice calculations place at $T_{\text{QCD}} = 156.5 \pm 1.5$ MeV [50], so that hadrons such as protons and neutrons have already formed out of a quark-gluon plasma.

Moreover, GW emission ceases safely before the onset of Big Bang nucleosynthesis (BBN) at $T \sim 1$ MeV [51]. This is essential for the viability of the scenario, since a DW network that survives into the BBN epoch and subsequently decays into Standard

Model particles would destroy already formed light elements, leaving the primordial abundances inconsistent with observations [8].

Furthermore, we consider the Bayes estimator of $\alpha_{\text{GW}} = 0.320^{+0.136}_{-0.096}$. This value implies that, in order to explain the observed GW spectrum entirely through a DW contribution, about one third of the energy density of the Universe has to be stored in DWs at the time at which efficient GW production ceases. While this does not correspond to an epoch of DW domination, the contribution of the DWs is non-negligible, especially under consideration of the upper bound of the 95 % credible interval, which reads the value $\alpha_{\text{GW}} = 0.648$, implying DW dominance. Their influence on the expansion rate is not accounted for in the template of Sec. 3.2.2, which assumes a radiation-dominated background, so that the necessary corrections are neglected in the spectral shape. A quantitative assessment of this effect would require a template derived for a non-negligible DW contribution, which is beyond the scope of this work.

As the thermodynamic parametrisation was also used in the NG15 data analysis [14], the same data set used in this work, a comparison can be performed, bearing in mind the differences between the GW templates pointed out in Sec. 3.2.3. To do so, the results provided in [14], given for what they call the DW-SM model, in which the decay of the DWs to standard model particles is considered, are listed in Tab. 5.2, together with the values obtained in this work with the GW template of Notari et al. [9]. The recovered temperature of this work lies systematically below that of [14], while the recovered α_{GW} lies above it. Both trends follow from the template differences discussed in Sec. 3.2.3. For the following comparison we consider the geometric means of the 68 % intervals, since the Bayes estimator values are not included in [14] and the geometric means differ by less than 3 % from the Bayes estimators for the values obtained in this work. The temperature of [14] exceeds the one obtained here by a factor of 2.5, in good agreement with the factor $x_p \simeq 2.15$ by which the peak frequency relations differ. Likewise, the value of α_{GW} obtained here exceeds theirs by a factor of 2.6, consistent with the factor of about 3 expected from the difference in the efficiency parameter. Therefore, the observed deviations between the two analyses are largely accounted for by the respective templates. The remaining deviation from the expected factors is plausible, as the two templates

differ in further respects, in particular in the spectral shape and in the rounding of the numerical prefactor in the GW template of Notari et al. [9].

Table 5.2: Comparison of the results for the thermodynamic parametrisation obtained in this work with those reported for the DW-SM model in the NG15 data analysis [14], in which the domain walls are likewise assumed to be the only source of GWs. It should be noted that T_* and α_* are evaluated at the decay of the network rather than at the end of efficient GW production, as discussed in Sec. 3.2.3. As [14] does not report the Bayes estimator, the geometric mean of the 68 % interval is quoted for comparison. For the values of this work, it agrees with the Bayes estimator to within 3 %.

Parameter	Reference	68 % CI	95 % CI	Geom. mean (68 %)
T / MeV	This work: T_{GW}	[51.4, 92.6]	[35.9, 119.6]	69.0
	Afzal et al. [14]: T_*	[110, 275]	[76, 505]	173.9
α	This work: α_{GW}	[0.228, 0.452]	[0.157, 0.648]	0.321
	Afzal et al. [14]: α_*	[0.080, 0.19]	[0.053, 0.35]	0.123

Overall, the thermodynamic parametrisation provides a straightforward access to quantities characterising the state of the Universe at the time when efficient GW production from DWs ceases, and allows conclusions to be drawn regarding the viability of the required scenario. Turning to the state of the radiating DW network itself, the results for the microscopic parametrisation are presented in the following Sec. 5.1.2.

5.1.2 Microscopic Parameters

The prior ranges for the microscopic parameters are listed in Tab. 4.1. For the DW tension, $\theta_\sigma \in [3, 7]$ covers energy scales between $\sigma^{1/3} = 1 \text{ TeV}$ and 10^4 TeV . The range is chosen such that the resulting GW amplitude can reach the level observed by PTAs. For $\theta_{\Delta V}$, the prior spans $[-3, 1]$, corresponding to $\Delta V^{1/4}$ between 1 MeV and 10 GeV. Together with the range chosen for θ_σ , T_{GW} spans many orders of magnitude via Eq. (3.25), so that the region relevant for frequencies obtained by PTA searches is contained. While the posterior distributions are visualised in Fig. 5.2, the statistical summary of the posterior values is given in Tab. 5.3.




Figure 5.2: Corner plot, which visualises the one- and two-dimensional posterior distributions for the microscopic parameters $\theta_\sigma = \frac{1}{3} \log_{10}(\sigma/\text{GeV}^3)$ and $\theta_{\Delta V} = \frac{1}{4} \log_{10}(\Delta V/\text{GeV}^4)$.

Considering the two one-dimensional posterior distributions in Fig. 5.2 for each individual parameter, both are well constrained and result in a unimodal, highly symmetric distribution. The two-dimensional posterior distribution in the bottom left panel points towards a precisely constrained parameter combination of σ and ΔV , matching the characteristic elongated and tilted elliptic form of a degenerate parameter constraint. Since the introduction of the microscopic parameters is essentially a reparametrisation of T_{GW} and α_{GW} in terms of σ and ΔV , the degeneracy observed in Sec. 5.1.1 cannot be expected to vanish. Instead, the reparametrisation preserves this effect, as can be seen by taking Eq. (3.25) into account, which links T_{GW} to σ and ΔV . The temperature does not depend on the individual values of the two parameters, but rather on the effective combination $T_{\text{GW}} \propto \sqrt{\Delta V/\sigma}$. This enables both of the parameters to occupy a variety of different values without

Table 5.3: Statistical summary of the microscopic two-parameter signal run. All parameters satisfy $\hat{R} < 1.01$. The credible intervals are given as the highest posterior density intervals (HPDI). The maximum posterior is abbreviated as MAP. The lower block lists the corresponding values on a linear scale, which are obtained by transforming the sampled logarithmic quantities.

Parameter	68 % HPDI	95 % HPDI	Bayes estimator	MAP
θ_σ	[5.146, 5.431]	[4.969, 5.586]	5.290 ± 0.166	5.294
$\theta_{\Delta V}$	[-1.108, -0.753]	[-1.331, -0.567]	-0.932 ± 0.205	-0.922
$\sigma^{1/3} / \text{TeV}$	[139.8, 270.1]	[93.2, 385.4]	$195.0^{+90.9}_{-62.0}$	196.9
$\Delta V^{1/4} / \text{MeV}$	[78.0, 176.5]	[46.7, 271.3]	$117.0^{+70.4}_{-44.0}$	119.6

altering the temperature, since a change can be compensated by an adjustment of the respective other parameter.

Before analysing the physical meaning of the retrieved statistical values, we can use the definition of T_{GW} in Eq. (3.25) and that of α_{GW} in Eq. (3.21) in order to run a consistency check with the values obtained in the previous section for the thermodynamic parametrisation, gathered in Tab. 5.1. Inserting the Bayes estimators of the microscopic parameters into these equations yields

$$T_{\text{GW, micro}} = 68.4 \text{ MeV} \quad \text{and} \quad \alpha_{\text{GW, micro}} = 0.330. \quad (5.2)$$

These values are in good agreement with those obtained in the thermodynamic run, $T_{\text{GW, thermo}} = 66.8 \text{ MeV}$ and $\alpha_{\text{GW, thermo}} = 0.320$, deviating by only 2.4 % and 3.1 %, respectively. The two parametrisations therefore describe the same physical scenario consistently, despite being sampled independently and with different prior distributions.

Turning to the physical meaning of the obtained values in Tab. 5.3, we first note that the DW tension, given with a Bayes estimator of $\sigma^{1/3} = 195.0^{+90.9}_{-62.0} \text{ TeV}$, violates the ZKO bound of Eq. (2.7). Given the prior setup, the posterior had no possibility of complying with this bound, since no value in the sampled range satisfies $\sigma^{1/3} \lesssim \mathcal{O}(\text{MeV})$. This range was chosen in order not to inflate the sampled volume with regions of parameter space far from being able to explain the observed GW

signal, which can be verified by a brief estimate. Demanding the spectral peak to fall inside the NG15 band restricts T_{GW} to roughly 8 to 107 MeV via Eq. (3.20). Solving the definition of α_{GW} in Eq. (3.21) for σ yields

$$\sigma = \frac{3}{2} \alpha_{\text{GW}} H_{\text{GW}} M_p^2, \quad (5.3)$$

such that even a modest DW contribution of $\alpha_{\text{GW}} = 0.01$ already requires tensions of $\sigma^{1/3} \sim 14$ to 84 TeV within this temperature range. A tension complying with the ZKO bound, $\sigma^{1/3} \sim 1$ MeV, would yield $\alpha_{\text{GW}} \sim 10^{-25}$ at $T_{\text{GW}} = 67$ MeV and thus a heavily suppressed GW amplitude relative to the observed signal. The violation of the ZKO bound is therefore a consequence of the amplitude of the measured SGWB. It poses no issue for the DW model considered here, since the potential bias renders the DWs unstable, causing them to annihilate.

Moreover, Eq. (2.6) links the DW tension directly to the scale of symmetry breaking, which can be solved for v , yielding $v = \left[3\sigma / (4\sqrt{\lambda/2}) \right]^{1/3}$. Although the self-coupling λ is not determined by the analysis, the resulting scale is largely insensitive to it, since $v \propto \lambda^{-1/6}$. Varying λ over three orders of magnitude changes v by less than a factor of three. With the Bayes estimator of $\sigma^{1/3} = 195.0^{+90.9}_{-62.0}$ TeV, we obtain $v \approx 200$ TeV for $\lambda = 1$ and $v \approx 430$ TeV for $\lambda = 0.01$. The symmetry breaking giving rise to the DW network therefore happens at a scale of several hundred TeV. To put this into perspective, we compare it to the maximum centre-of-mass energy of 13.6 TeV reached at the LHC [52], which is smaller by more than an order of magnitude. Thus, the energy scales probed here are far out of reach for present-day colliders and further motivate PTA observations as a probe of high-energy physics beyond the Standard Model.

In contrast, the obtained Bayes estimator value for the potential bias is given by $\Delta V^{1/4} = 117.0^{+70.4}_{-44.0}$ MeV, corresponding to an energy scale roughly six orders of magnitude below that of the DW tension. This is a remarkable finding, since it is precisely this comparatively minute energy difference between the two vacua that causes the DWs to be unstable and initiates their annihilation.

Moreover, the retrieved value can be compared to the lower limit of Eq. (2.19), which has to be satisfied in order to avoid an era of DW domination. Inserting $\mathcal{A} \simeq 1$,

the Bayes estimator of σ , and both ends of the estimated range of the numerical constant $C_{\text{ann}} \simeq 2\text{--}5$ yields

$$\Delta V_{\text{min}}^{1/4} = 70.5\text{--}88.7 \text{ MeV} . \quad (5.4)$$


The central value of $\Delta V^{1/4}$ therefore lies safely above this limit, so that no epoch of DW domination is entered within the framework of this model. However, the lower end of the 68 % credible interval, $\Delta V^{1/4} = 78.0 \text{ MeV}$, only satisfies the bound for the smallest values of C_{ann} , and would be excluded for larger values of C_{ann} . Hence, this tension can be classified as mild, since the Bayes estimator lies above the lower bound, while the lower limit of the 68 % credible interval still satisfies the bound for some part of the range of C_{ann} .

In summary, the microscopic parametrisation provides a complementary description of the same scenario analysed in Sec. 5.1.1, yielding access to the properties of a DW network capable of explaining the entire SGWB signal. Before leaving this assumption behind and analysing the results of the upper-limit runs in Sec. 5.2, we review the effect of including the three spectral parameters with the largest relative uncertainties as free parameters in Sec. 5.1.3.


5.1.3 Inclusion of Free Spectral Parameters

For the spectral parameters β , γ and x_b , included as free parameters in this supplementary analysis, the prior was chosen to be Gaussian, with the prior values set to the best-fit values obtained by Notari et al. [9], in order to leverage the prior knowledge gained through the data fit. The concretely chosen values are given in Tab. 4.1. The resulting corner plots for both parametrisations are shown in Fig. 5.3, while the statistical summaries for the thermodynamic and microscopic parameters, including the spectral indices and the chosen prior values, are tabulated in Tab. 5.4 and Tab. 5.5, respectively.

Considering the posterior distributions of the spectral parameters in Fig. 5.3, all three show the same behaviour: the one-dimensional posteriors reproduce the Gaussian priors that were initially set up, while the two-dimensional posteriors provide no sign of degeneracy with any other parameter, as reflected in their circular, non-elongated shape. This is confirmed quantitatively by Tab. 5.4 and Tab. 5.5: both



(a) Thermodynamic parametrisation.



(b) Microscopic parametrisation.

Figure 5.3: Corner plots of the five-parameter signal runs, visualising the one- and two-dimensional posterior distributions of the respective model parameters together with the spectral parameters x_b , β and γ .

Table 5.4: Statistical summary of the thermodynamic five-parameter signal run. All parameters satisfy $\hat{R} < 1.01$. The credible intervals are given as the highest posterior density intervals (HPDI). The maximum posterior is abbreviated as MAP. For the spectral parameters, the Gaussian priors of Tab. 4.1 are listed as $\mu \pm s$ for comparison. T_{GW} is given in GeV.

Parameter	Prior	68 % HPDI	95 % HPDI	Bayes estimator	MAP
$\log_{10}(T_{\text{GW}})$	$[-3, 0]$	$[-1.284, -1.021]$	$[-1.436, -0.896]$	-1.162 ± 0.136	-1.147
$\log_{10}(\alpha_{\text{GW}})$	$[-4, 0]$	$[-0.641, -0.346]$	$[-0.800, -0.188]$	-0.493 ± 0.152	-0.492
β	0.48 ± 0.15	$[0.319, 0.618]$	$[0.176, 0.762]$	0.470 ± 0.150	0.465
γ	1.79 ± 0.36	$[1.430, 2.149]$	$[1.081, 2.497]$	1.790 ± 0.361	1.789
x_b	6.2 ± 1.9	$[4.346, 8.093]$	$[2.562, 9.911]$	6.241 ± 1.877	6.190

Table 5.5: Statistical summary of the microscopic five-parameter signal run. All parameters satisfy $\hat{R} < 1.01$. The credible intervals are given as the highest posterior density intervals (HPDI). The maximum posterior is abbreviated as MAP. For the spectral parameters, the Gaussian priors of Tab. 4.1 are listed as $\mu \pm s$ for comparison.

Parameter	Prior	68 % HPDI	95 % HPDI	Bayes estimator	MAP
θ_σ	$[3, 7]$	$[5.152, 5.441]$	$[4.978, 5.605]$	5.301 ± 0.166	5.301
$\theta_{\Delta V}$	$[-3, 1]$	$[-1.100, -0.738]$	$[-1.320, -0.535]$	-0.915 ± 0.206	-0.914
β	0.48 ± 0.15	$[0.318, 0.616]$	$[0.172, 0.760]$	0.467 ± 0.150	0.465
γ	1.79 ± 0.36	$[1.432, 2.147]$	$[1.082, 2.494]$	1.789 ± 0.360	1.787
x_b	6.2 ± 1.9	$[4.364, 8.104]$	$[2.611, 9.877]$	6.254 ± 1.863	6.227

the Bayes estimators and their standard deviations agree with the corresponding prior values with a maximum deviation of 2.7 %, occurring for the central value of β in the microscopic parametrisation. Therefore, in the context of the considered DW model, the NG15 data provide no constraining power on the spectral shape of the GW spectrum beyond the peak in the spectrum; it is only sensitive to the spectral shape of the low-frequency tail. It should be noted that this statement is made within the informative Gaussian priors used here, which confine the spectral parameters to the neighbourhood of the values obtained from the numerical simulations. Whether the NG15 data would constrain the spectral shape at all under uninformative priors was not tested and would require a dedicated analysis.

A further point worth discussing is the influence that the inclusion of free spectral parameters has on the physical ones. Since we have seen that the NG15 data carry no information on the spectral shape, no drastic influence is expected. This expectation is confirmed by comparing the tabulated values for the thermodynamic parameters between Tab. 5.1 and Tab. 5.4, as well as for the microscopic parameters between Tab. 5.3 and Tab. 5.5. Neither the central values nor the uncertainties are largely affected: the Bayes estimators shift by a maximum of 0.1 standard deviations, and the widths of the 68 % credible intervals change by less than 3 %. Introducing three additional free spectral parameters therefore leaves the constraining power on the physical parameters essentially unchanged, demonstrating their robustness with respect to the uncertainty of the fit to the numerical simulations of Notari et al. [9].

Therefore, this supports the choice of fixing all spectral parameters in the two-parameter runs discussed in Sec. 5.1.1 and Sec. 5.1.2, as well as in the upper-limit runs presented in Sec. 5.2.

5.2 Upper-Limit Runs

In this section, we drop the assumption that the DW contribution has to account for the entire GW spectrum by including a free spectrum, as explained in Sec. 4.2.3. Since the posterior in Bayesian statistics depends not only on the data but also on the prior, two different prior choices for the 14 amplitude parameters of the FS are


compared, in order to evaluate their respective influence: the commonly applied log-uniform prior in Sec. 5.2.1 and the linear-exponential prior in Sec. 5.2.2.

5.2.1 Log-Uniform Prior


First, we analyse the results obtained for the log-uniform prior applied to the 14 FS amplitude parameters, sampled over the range $[-14, -5]$. This range is chosen to comfortably include the amplitudes recovered in the first 14 frequency bins of the NG15 data set, as can be seen in [53], so that no physically relevant value is excluded before sampling. With a log-uniform prior, every decade of the sampled range receives the same weight, corresponding to an uninformative search across the nine decades considered here. The resulting posterior distributions for the thermodynamic and microscopic parameters are directly compared to those obtained in the signal runs in Fig. 5.4, whereas the statistical values are compared in Tab. 5.6 for the thermodynamic and in Tab. 5.7 for the microscopic parameters.

As can be seen in Fig. 5.4, the effect of the FS is qualitatively the same for the thermodynamic and the microscopic parameters: the one-dimensional posterior distributions broaden and shift towards lower values. For the two-dimensional posteriors, the characteristic shape of the distribution is preserved, while the contour sweeps over a larger range of values, including the region obtained in the signal run.

This qualitative observation is supported by the numerical values listed in Tab. 5.6 and Tab. 5.7. For all parameters, both the Bayes estimator and the MAP are shifted towards lower values in the upper-limit runs compared to the signal runs. Taking the standard deviation of the signal-run posterior as a reference scale, the Bayes estimator shifts by 0.85 to 1.32 and the MAP by 1.43 to 1.82 of that width. In addition, the widths of the 68 % credible intervals increase by factors between 1.40 for $\log_{10}(\alpha_{\text{GW}})$ and 1.54 for $\log_{10}(T_{\text{GW}})$. We can also note that the Bayes estimator and the MAP deviate from each other notably for the first time, corresponding to a posterior distribution differing noticeably from a symmetric one. For $\log_{10}(T_{\text{GW}})$, the two differ by 0.056, compared to 0.022 in the corresponding signal run. This is qualitatively supported by the shape of the distributions in Fig. 5.4. However, all posteriors still close within the sampled prior range.



(a) Thermodynamic parametrisation.



(b) Microscopic parametrisation.

Figure 5.4: Comparison of the posterior distributions of the two-parameter signal runs with those of the corresponding upper-limit runs, in which an additional free spectrum (FS) component with a log-uniform prior is included.

Table 5.6: Comparison of the thermodynamic two-parameter signal run with the corresponding upper-limit run, in which an additional free spectrum with a log-uniform prior is included. All parameters satisfy $\hat{R} < 1.01$. The credible intervals are given as the highest posterior density intervals (HPDI). The maximum posterior is abbreviated as MAP. T_{GW} is given in GeV.

Parameter	Run	68 % HPDI	95 % HPDI	Bayes estimator	MAP
$\log_{10}(T_{\text{GW}})$	Signal	$[-1.289, -1.033]$	$[-1.445, -0.922]$	-1.175 ± 0.132	-1.153
	Upper limit	$[-1.516, -1.122]$	$[-1.672, -0.912]$	-1.312 ± 0.220	-1.368
$\log_{10}(\alpha_{\text{GW}})$	Signal	$[-0.642, -0.344]$	$[-0.805, -0.188]$	-0.494 ± 0.154	-0.491
	Upper limit	$[-0.921, -0.504]$	$[-1.062, -0.239]$	-0.698 ± 0.280	-0.771

Table 5.7: Comparison of the microscopic two-parameter signal run with the corresponding upper-limit run, in which an additional free spectrum with a log-uniform prior is included. All parameters satisfy $\hat{R} < 1.01$. The credible intervals are given as the highest posterior density intervals (HPDI). The maximum posterior is abbreviated as MAP.

Parameter	Run	68 % HPDI	95 % HPDI	Bayes estimator	MAP
θ_σ	Signal	$[5.146, 5.431]$	$[4.969, 5.586]$	5.290 ± 0.166	5.294
	Upper limit	$[4.892, 5.305]$	$[4.689, 5.629]$	5.098 ± 0.316	5.046
$\theta_{\Delta V}$	Signal	$[-1.108, -0.753]$	$[-1.331, -0.567]$	-0.932 ± 0.205	-0.922
	Upper limit	$[-1.420, -0.899]$	$[-1.679, -0.506]$	-1.107 ± 0.355	-1.216

These findings suggest that the inclusion of a DW contribution remains favoured, despite the addition of a FS which is capable of perfectly explaining the SGWB data on its own, since the parameter region mapped out by the signal runs is still roughly the one with the highest posterior density. In addition to the DW contribution, the FS assists in the explanation of the GW spectrum for slightly deviating parameter combinations, resulting in the mentioned broadening of the posterior. The shift towards lower values makes sense, as part of the signal is now explained by the FS, resulting in smaller necessary values of the DW parameters: a lower T_{GW} and α_{GW} mean that less energy was stored in DWs at the time when efficient GW emission ceased, and that this happened at cooler temperatures and thus later in the history of the Universe. The contribution missing to reach the signal strength of the measured SGWB is then filled in by the FS. The same holds true for the microscopic reparametrisation.


Since the deviations are, as noticed, small, the DW contribution remains dominant, such that the FS contributes a subdominant portion towards the explanation of the signal. The physical nature of this observation has to be taken with a grain of salt, as this result is likely to originate from the chosen prior. As stated, the latter was chosen to be log-uniform, such that the sampler scans nine decades of spectral amplitude for each of the 14 FS parameters, putting the same weight on each decade and thus spreading the probability mass over a huge, 14-dimensional parameter volume. Only a minute fraction of this volume contains amplitude combinations large enough to fully explain the GW spectrum via the FS alone, so that this possibility is strongly suppressed. This is a manifestation of Occam's razor, which is automatically included quantitatively as part of Bayesian inference [54, 55]: a model spreading its probability mass over a large parameter volume, of which only a small part is compatible with the data, is penalised relative to a simpler description. The much easier option is a fitting parameter combination in the 2-dimensional parameter subvolume constructed by the two thermodynamic or microscopic parameters, which is not subject to such a penalty. Thus, the DW contribution remains dominant, while the FS enables slight deviations from the signal-run results by compensating for those.

In the following Sec. 5.2.2, we investigate the influence that the differing choice of a linear-exponential prior has on the physical posteriors.


5.2.2 Linear-Exponential Prior

In this section, we review the results obtained for the upper-limit runs with a linear-exponential prior applied to the FS amplitude parameters. While the log-uniform prior assigns the same weight to every decade, the linear-exponential prior is uniform in the amplitude itself, corresponding to a probability density $p(x) \propto \ln(10) 10^x$ for $x = \log_{10}(h^2 \Omega_{\text{GW}})$ [11], and thus places most of its weight on the largest values. Analogous to Sec. 5.2.1, the prior range is chosen as $[-14, -5]$. The resulting posterior distributions are visualised in Fig. 5.5, while the statistical values, comparing the signal and upper-limit runs, are given in Tab. 5.8 for the thermodynamic and in Tab. 5.9 for the microscopic parameters.

The resulting posterior distributions for the linear-exponential prior in Fig. 5.5 differ drastically from those obtained with a log-uniform prior, depicted in Fig. 5.4. In the former, the one-dimensional posteriors do not form a pronounced peak, but are rather flat across the entire depicted range with only slight increases and decreases, and do not favour a distinct region. Instead, a clear trend towards a decreasing DW contribution can be discerned. While T_{GW} is largely flat over the whole parameter range, larger values of α_{GW} , corresponding to more energy being stored in DWs and thus a stronger contribution to the GW signal, are disfavoured. The same trend holds for the microscopic parameters, where smaller values of σ and larger values of ΔV are preferred, corresponding to less energy stored in the DW network and a to an earlier annihilation, respectively, and hence to a weaker GW contribution. From this we can conclude that the FS has taken over the explanation of the GW signal, leaving only a reduced contribution for the DWs. While a DW contribution on top of the FS thus remains feasible, the parameter configurations recovered in the signal runs now lie outside the 95 % credible region. This follows directly from the previous argument: since those parameters were obtained under the assumption that the DWs account for the entire spectrum, they are no longer favoured once the FS already provides the dominant contribution. Additionally, we should note that the patchy appearance of the two-dimensional contours is a consequence of the previously noted flatness: where the posterior density is nearly constant, small statistical fluctuations are sufficient to move the contour lines, so that the resulting islands and gaps reflect sampling noise rather than a structure of physical origin.



(a) Thermodynamic parametrisation.



(b) Microscopic parametrisation.

Figure 5.5: Comparison of the posterior distributions of the two-parameter signal runs with those of the corresponding upper-limit runs, in which an additional free spectrum (FS) component with a linear-exponential prior is included.

Table 5.8: Comparison of the thermodynamic two-parameter signal run with the corresponding upper-limit run, in which an additional free spectrum with a linear-exponential prior is included. All parameters satisfy $\hat{R} < 1.01$. The credible intervals are given as the highest posterior density intervals (HPDI). A dash replaces a bound that is not attained within the sampled prior range. The maximum posterior is abbreviated as MAP. T_{GW} is given in GeV.

Parameter	Run	68 % HPDI	95 % HPDI	Bayes estimator	MAP
$\log_{10}(T_{\text{GW}})$	Signal	$[-1.289, -1.033]$	$[-1.445, -0.922]$	-1.175 ± 0.132	-1.153
	Upper limit	$[-, -]$	$[-, -]$	-1.388 ± 0.881	-0.001
$\log_{10}(\alpha_{\text{GW}})$	Signal	$[-0.642, -0.344]$	$[-0.805, -0.188]$	-0.494 ± 0.154	-0.491
	Upper limit	$[-, -1.627]$	$[-, -0.656]$	-2.367 ± 0.990	-3.241

Table 5.9: Comparison of the microscopic two-parameter signal run with the corresponding upper-limit run, in which an additional free spectrum with a linear-exponential prior is included. All parameters satisfy $\hat{R} < 1.01$. The credible intervals are given as the highest posterior density intervals (HPDI). A dash replaces a bound that is not attained within the sampled prior range. The maximum posterior is abbreviated as MAP.

Parameter	Run	68 % HPDI	95 % HPDI	Bayes estimator	MAP
θ_σ	Signal	$[5.146, 5.431]$	$[4.969, 5.586]$	5.290 ± 0.166	5.294
	Upper limit	$[-, 4.657]$	$[-, 5.939]$	4.296 ± 0.884	3.458
$\theta_{\Delta V}$	Signal	$[-1.108, -0.753]$	$[-1.331, -0.567]$	-0.932 ± 0.205	-0.922
	Upper limit	$[-1.042, -]$	$[-2.532, -]$	-0.549 ± 1.070	1.000

The preceding qualitative findings are supported by the statistics in Tabs. 5.8 and 5.9. Five of the eight interval limits per credible level do not exist, since the posterior does not close within the prior range. No parameter is bounded towards values corresponding to a smaller DW contribution: neither α_{GW} nor θ_σ towards lower values, nor $\theta_{\Delta V}$ towards higher ones. T_{GW} , which sets the peak frequency, is not bounded at all. This is expected, as the DW spectrum becomes negligible compared to the FS contribution for small α_{GW} , so that the position of its peak is no longer constrained by the data.

Apart from that, the Bayes estimator and the MAP differ noticeably once more, both being shifted towards a minimal DW contribution. Since no region in parameter space is strongly preferred compared to the rest, both lose their meaning as parameter estimates. The MAP values of T_{GW} and $\theta_{\Delta V}$ coincide with the boundaries of their prior ranges, while the Bayes estimator, given as an average value, retains a dependence on the chosen prior range. In particular, this can be seen for $\log_{10}(T_{\text{GW}})$, where the Bayes estimator value of -1.388 ± 0.881 is close to -1.5 , the value that would be expected from a completely flat distribution across the prior range $[-3, 0]$. This indicates that the data provide essentially no information on this parameter. The remaining parameters deviate more strongly from their flat expectation, which reflects the trend towards a reduced DW contribution.

The results suggest that a DW contribution on top of a strong FS remains possible, since a broad range of parameter values is allowed, in particular those corresponding to a weak GW signal from DWs. The inclusion of DWs is, however, not particularly favoured in this model, as no distinct region of high posterior density can be singled out, meaning that the DW contribution is rather tolerated than required.

This finding can again be explained by the prior choice. Since most of the prior weight is placed on the largest FS amplitudes within the sampled range, the FS is driven towards large values, and this bias enters once per amplitude parameter. The volume argument of Sec. 5.2.1 thus applies in reverse, as configurations in which the FS amplitudes are small enough to result in a DW-dominated solution now occupy only a small fraction of the prior volume, whereas the region in which the FS accounts for the signal on its own is no longer penalised. Therefore, the FS governs the explanation of the GW signal, while the contribution of the DWs is reduced to be subdominant or entirely negligible.

Although no parameter determination is possible in this case, the run does yield conservative limits on the DW contribution under the assumption that other sources dominate the signal. For the 95 % credible interval, these are given as $\alpha_{\text{GW}} < 0.221$ and $\sigma^{1/3} < 869 \text{ TeV}$, together with a lower bound of $\Delta V^{1/4} > 2.9 \text{ MeV}$. These bounds are not independent of one another, as the thermodynamic and microscopic quantities are related through Eqs. (3.21) and (3.25), while they all express the same statement, namely that the DW contribution to the signal is favoured to be small. These bounds should be regarded as loose, as illustrated by the gap between the two credible levels, e.g. for $\sigma^{1/3}$, the 95 % limit exceeds the 68 % one by a factor of roughly 19, reflecting the flatness of the underlying posterior.

In summary, the upper-limit runs with a log-uniform prior favour a contribution from DWs, while the linear-exponential prior minimises it. Neither choice is more correct than the other, as they are merely based on different assumptions about the amplitude of the FS. The log-uniform run shows that DWs remain a feasible constituent of the GW spectrum, even though they are not required, whereas the linear-exponential run yields conservative upper limits under the assumption that other, unspecified sources dominate the spectrum. Thus, these results illustrate the broad range of roles DWs can play in the explanation of the SGWB, from negligible to dominant.

Conclusion and Outlook

In this bachelor's thesis, a recent gravitational wave (GW) template for domain walls (DWs) by Notari et al. [9], which includes the crucial phase of DW collapse and its dominant contribution to the GW spectrum [9, 12], was implemented into the Python-based framework PTArcade [10, 11] and applied to the NANOGrav 15-year (NG15) data set [2].

First, the scenario of the stochastic gravitational wave background (SGWB) being fully explained by the contribution of DWs was explored in the ‘signal runs’ for two different parametrisations of the spectrum. The thermodynamic parametrisation describes the GW spectrum through the temperature at which efficient GW emission ceases, T_{GW} , and the fraction of the Universe's energy density stored in DWs at this time, α_{GW} . The Bayes estimator of the former is given as $T_{\text{GW}} = 66.8^{+23.8}_{-17.5}$ MeV, placing the end of GW emission after the QCD crossover and well before the onset of big bang nucleosynthesis, which is required in order not to perturb the predicted primordial abundances. The Bayes estimator of $\alpha_{\text{GW}} = 0.320^{+0.136}_{-0.096}$ further implies that DWs contain a non-negligible portion of the Universe's energy budget at T_{GW} , and thus that a correction to the expansion rate would be required, which the template, assuming a radiation-dominated background, does not include. Moreover, these results were compared to former findings in the NG15 data analysis [14], which is based on a different GW template. The resulting deviations in both parameters can mainly be traced back to the differences in the peak frequency relation and the efficiency parameter, while the remaining discrepancy can be attributed to further variations, in particular the spectral shape and the rounding of a numerical prefactor.

Additionally, a dedicated microscopic parametrisation was used, describing the GW spectrum in terms of the DW tension σ and the difference between the potential minima ΔV induced by the potential bias. In order to explain the detected signal amplitude within this model, the DW tension, given with a Bayes estimator of $\sigma^{1/3} = 195.0^{+90.9}_{-62.0}$ TeV, has to violate the Zel’dovich–Kobzarev–Okun bound $\sigma^{1/3} \lesssim \mathcal{O}(\text{MeV})$ by eight orders of magnitude. This poses no issue here, as the potential bias destabilises the network and drives its collapse. Being directly linked to the scale of symmetry breaking via $\sigma^{1/3} \propto v$, the recovered tension provides access to physics at energies exceeding those achieved at the LHC by more than an order of magnitude. This scale is contrasted by the Bayes estimator of the potential bias, $\Delta V^{1/4} = 117.0^{+70.4}_{-44.0}$ MeV, which causes the DW decay despite its energy scale being six orders of magnitude below that of the tension, and implies annihilation before an era of DW domination. Since the lower end of the 68% credible interval falls below the corresponding bound for part of the range of the numerical constant C_{ann} , this scenario cannot be excluded entirely.

Furthermore, an analysis including the spectral parameters β , γ and x_b as free parameters was performed for both parametrisations. The resulting posteriors match the Gaussian priors to within 2.7 %, showing that the NG15 data do not shift the spectral parameters away from the values obtained from the numerical simulations. Apart from that, the physical parameters shift by a maximum of 0.1 standard deviations and their credible intervals broaden by less than 3 %, so that the results obtained above are robust with respect to the uncertainty of the spectral parameters.

Lastly, ‘upper-limit runs’ were performed, in which the DW contribution is supplemented by a free spectrum (FS), modelled via 14 amplitude parameters and capable of explaining the measured signal on its own. With a log-uniform prior on these amplitudes, the resulting posteriors deviate from those of the signal runs only slightly, upholding a preference towards the inclusion of DWs, while showing a shift towards a lower DW contribution and a broadening of the distributions. With a linear-exponential prior, which places most of the probability mass on the largest amplitudes, the DW posteriors become flat and no longer close within the prior range, implying a dominant FS contribution alongside a tolerated DW component.

Both outcomes follow from the way the prior volume of the 14-dimensional FS subspace is distributed relative to the two-dimensional DW subspace, and neither prior choice is more correct than the other. Overall, they illustrate that the role of DWs in the explanation of the SGWB can range from negligible to dominant, with the linear-exponential run yielding conservative upper limits of $\alpha_{\text{GW}} < 0.221$ and $\sigma^{1/3} < 869$ TeV at the 95 % credible level.

Turning towards an outlook on upcoming developments in this research field, the most immediate prospect lies in the analysis of future PTA data releases, such as the NANOGrav 20-year data set. A longer observation baseline extends the accessible band towards lower frequencies, while an increased number of pulsars improves the SGWB sensitivity [56]. Repeating the analysis performed in this thesis would allow the present results to be verified and could even tighten the constraints on the individual parameters. In particular, a tighter constraint on the potential bias would clarify whether an era of DW domination can be excluded for this model.

Apart from that, the results obtained in this thesis suggest a rather large value of α_{GW} , implying the need for corrections to the radiation-dominated background assumed in the derivation of the applied GW spectrum. Since the magnitude of this effect grows with α_{GW} , it is most relevant for the signal runs discussed here, whereas it becomes negligible in the upper-limit scenarios, where the DW contribution is smaller. While a quantitative assessment of the resulting effect on the template is beyond the scope of this work, accounting for a non-negligible DW contribution to the expansion rate would be a logical next step for future developments of DW spectra, since a strong DW contribution remains a viable possibility, as demonstrated here.

Another interesting prospect is the launch of the Laser Interferometer Space Antenna (LISA), which will probe the millihertz band of GWs [57]. Recent studies show that LISA could probe the ultraviolet tail of a GW signal from DWs peaking in the PTA band and extending into the millihertz region [58], or detect a signal peaking in the millihertz band itself, which corresponds to annihilation temperatures of 10^3 to 10^6 GeV and thus a region of the DW parameter space far above the one probed in this work [59]. Either way, LISA offers a complementary approach to the study of GWs from DWs.

PTArcade Model Files

The code given in the following is the PTArcade model file for the microscopic two-parameter model without free spectral parameters, chosen as a representative example. Generally, the model files are structured as follows. First, the required packages are imported, among them `models_utils`, which provides constants such as the reduced Planck mass and functions such as `g_rho`, returning $g_\rho := g_*$ at a given temperature. The model name is then defined and a SMBHB contribution is explicitly excluded. Next, the model parameters are defined. In the version shown here, all shape parameters are fixed, while priors for the free parameters θ_σ and $\theta_{\Delta V}$ are set up by choosing the shape and the range. In the thermodynamic parametrisation these two are replaced by $\log_{10}(T_{\text{GW}})$ and $\log_{10}(\alpha_{\text{GW}})$. When including β , γ and x_b as free parameters, they are set up with a Gaussian prior, where the first argument is the central value and the second one the standard deviation, both obtained from the best-fit values of Notari et al. [9]. Then, all relevant equations presented in Sec. 3.2.2 are implemented. Special care is required for the calculation of T_{GW} , since it not only depends on σ and ΔV , but also implicitly on itself through g_* . This is taken care of through a fixed-point iteration, which converges within a few steps since g_* varies only mildly across the sampled temperature range. In the thermodynamic parametrisation, neither T_{GW} nor α_{GW} has to be computed since they are the parameters that are directly sampled, whereas the calculation of the peak and break frequency remains the same. The final spectrum is returned in the format $h^2\Omega_{\text{GW}}$.

```

1 import ptarcade.models_utils as aux
2 import numpy as np
3
4 name = 'dw_micro_2par'
5 smbhb = False

```

```

6
7 # Fixed shape parameters
8 BETA          = 0.48
9 GAMMA         = 1.79
10 X_P           = 2.15
11 F_B_OVER_F_P = 6.2/2.15
12 EPSILON       = 0.078
13
14
15 parameters = {
16     'theta_sigma': aux.prior("Uniform", 3, 7),
17     'theta_deltaV': aux.prior("Uniform", -3, 1),
18 }
19
20 def S_shape(f, f_p, f_b, beta, gamma):
21     # Spectral shape, eq. (3.13)
22     numerator = 3 + beta + (f_p / f_b)**(beta + gamma)
23     denominator = beta * (f / f_p)**(-3) + 3 * (f / f_p)**
24         beta + (f_p / f_b)**beta * (f / f_b)**gamma
25     return numerator/denominator
26
27 def spectrum(f, theta_sigma, theta_deltaV):
28     sigma_term = 10**theta_sigma
29     deltaV_term = 10**theta_deltaV
30
31     # Temperature at stop of GW emission, eq. (3.25).
32     # Implicit through g_*(T_GW), solved by fixed-point
33     iteration.
34     T_GW = 133.0
35     for _ in range(5):
36         g_star = aux.g_rho(T_GW * 10**(-3))
37         T_GW = 133.0 * (g_star / 17.25)**(-0.25) * (10 *
38             deltaV_term)**2 * (10**5 / sigma_term)**1.5

```

```

37
38     g_star    = aux.g_rho(T_GW * 10**(-3))
39     T_GW_GeV = T_GW * 10**(-3)
40
41     # Peak and break frequency today, eq. (3.20)
42     f_p = 24 * X_P * (g_star / 100)**(1/6) * (T_GW / 150)
         * 1e-9
43     f_b = F_B_OVER_F_P * f_p
44
45     # Energy fraction stored in DWs, eq. (3.21)
46     H_GW    = np.sqrt((np.pi**2 / 90.0) * g_star *
         T_GW_GeV**4) / aux.M_pl
47     alpha_GW = 2.0 * sigma_term**3 / (3.0 * H_GW * aux.
         M_pl**2)
48
49     # Peak amplitude, eq. (3.23)
50     omega_p_h2 = 1e-10 * EPSILON * (g_star / 10.75)
         **(-1/3) * (alpha_GW / 0.01)**2
51
52     return omega_p_h2 * S_shape(f, f_p, f_b, BETA, GAMMA)

```

For the upper-limit runs, the parameter list includes 14 additional amplitude parameters, one per frequency bin. The spectrum function is extended by a code block, where each frequency is assigned to its corresponding bin, the associated amplitude is read off, and the result is added to the DW spectrum. Frequencies outside the first 14 bins receive no free spectrum contribution. The two upper-limit runs differ only in the prior shape passed to `aux.prior`, being `Uniform` for the log-uniform and `LinearExp` for the linear-exponential case.

```

1  N_FS          = 14
2
3  # ...
4
5  for i in range(1, N_FS + 1):

```

```

6     parameters[f'log10_fs_{i}'] = aux.prior("LinearExp",
7         -14, -5)
8
9 def spectrum(f, theta_sigma, theta_deltaV,
10             log10_fs_1, log10_fs_2, ..., log10_fs_14):
11
12     # ... DW part ...
13
14     # FS part
15     log10_fs = np.array([log10_fs_1, log10_fs_2,
16                          log10_fs_3, log10_fs_4,
17                          log10_fs_5, log10_fs_6,
18                          log10_fs_7, log10_fs_8,
19                          log10_fs_9, log10_fs_10,
20                          log10_fs_11, log10_fs_12,
21                          log10_fs_13, log10_fs_14])
22
23     fs      = 10**log10_fs
24     i_bin   = np.rint(f / np.min(f)).astype(int) # bin
25             index of each frequency
26     omega_fs = np.zeros_like(f)
27     in_range = (i_bin >= 1) & (i_bin <= N_FS) # keep only
28             frequencies in bins 1 to 14
29     omega_fs[in_range] = fs[i_bin[in_range] - 1] #
30             frequencies outside receive no FS contribution
31     return omega_p_h2 * S_shape(f, f_p, f_b, BETA, GAMMA)
32         + omega_fs

```

Bibliography

- [1] Kai Schmitz. ‘Modern Cosmology, an Amuse-Gueule’. In: *Advances in Cosmology*. Springer International Publishing, 2022, pp. 37–70. ISBN: 9783031056253. DOI: [10.1007/978-3-031-05625-3_3](https://doi.org/10.1007/978-3-031-05625-3_3). URL: http://dx.doi.org/10.1007/978-3-031-05625-3_3.
- [2] Gabriella Agazie et al. ‘The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background’. In: *The Astrophysical Journal Letters* 951.1 (2023), p. L8. DOI: [10.3847/2041-8213/acdac6](https://doi.org/10.3847/2041-8213/acdac6).
- [3] J. Antoniadis et al. ‘The second data release from the European Pulsar Timing Array: III. Search for gravitational wave signals’. In: *Astronomy and Astrophysics* 678 (Oct. 2023), A50. ISSN: 1432-0746. DOI: [10.1051/0004-6361/202346844](https://doi.org/10.1051/0004-6361/202346844). URL: <http://dx.doi.org/10.1051/0004-6361/202346844>.
- [4] Daniel J. Reardon et al. ‘Search for an Isotropic Gravitational-wave Background with the Parkes Pulsar Timing Array’. In: *The Astrophysical Journal Letters* 951.1 (2023), p. L6. ISSN: 2041-8213. DOI: [10.3847/2041-8213/acdd02](https://doi.org/10.3847/2041-8213/acdd02). URL: <http://dx.doi.org/10.3847/2041-8213/acdd02>.
- [5] Heng Xu et al. ‘Searching for the Nano-Hertz Stochastic Gravitational Wave Background with the Chinese Pulsar Timing Array Data Release I’. In: *Research in Astronomy and Astrophysics* 23.7 (2023), p. 075024. ISSN: 2397-6209. DOI: [10.1088/1674-4527/acdfa5](https://doi.org/10.1088/1674-4527/acdfa5). URL: <http://dx.doi.org/10.1088/1674-4527/acdfa5>.
- [6] Matthew T Miles et al. ‘The MeerKAT Pulsar Timing Array: the first search for gravitational waves with the MeerKAT radio telescope’. In: *Monthly Notices of the Royal Astronomical Society* 536.2 (Dec. 2024), pp. 1489–1500. ISSN: 1365-2966. DOI: [10.1093/mnras/stae2571](https://doi.org/10.1093/mnras/stae2571). URL: <http://dx.doi.org/10.1093/mnras/stae2571>.

- [7] Gabriella Agazie et al. ‘The NANOGrav 15 yr Data Set: Bayesian Limits on Gravitational Waves from Individual Supermassive Black Hole Binaries’. In: *The Astrophysical Journal Letters* 951.2 (2023), p. L50. ISSN: 2041-8213. DOI: [10.3847/2041-8213/ace18a](https://doi.org/10.3847/2041-8213/ace18a). URL: <http://dx.doi.org/10.3847/2041-8213/ace18a>.
- [8] Ken’ichi Saikawa. ‘A Review of Gravitational Waves from Cosmic Domain Walls’. In: *Universe* 3.2 (May 2017), p. 40. ISSN: 2218-1997. DOI: [10.3390/universe3020040](https://doi.org/10.3390/universe3020040). URL: <http://dx.doi.org/10.3390/universe3020040>.
- [9] Alessio Notari, Fabrizio Rompineve, and Francisco Torrenti. ‘The spectrum of gravitational waves from annihilating domain walls’. In: *Journal of Cosmology and Astroparticle Physics* 2025.07 (2025), p. 049. ISSN: 1475-7516. DOI: [10.1088/1475-7516/2025/07/049](https://doi.org/10.1088/1475-7516/2025/07/049). URL: <http://dx.doi.org/10.1088/1475-7516/2025/07/049>.
- [10] Andrea Mitridate. ‘PTArcade’. In: (Apr. 2023). DOI: [10.5281/zenodo.7876430](https://doi.org/10.5281/zenodo.7876430).
- [11] Andrea Mitridate et al. ‘PTArcade’. In: (June 2023). arXiv: [2306.16377](https://arxiv.org/abs/2306.16377) [hep-ph].
- [12] Bryce Cyr, Steven Cotterill, and Richard Battye. *Near-Peak Spectrum of Gravitational Waves from Collapsing Domain Walls*. 2025. arXiv: [2504.02076](https://arxiv.org/abs/2504.02076) [astro-ph.CO]. URL: <https://arxiv.org/abs/2504.02076>.
- [13] Ricardo Z. Ferreira et al. *Collapsing Domain Wall Networks: Impact on Pulsar Timing Arrays and Primordial Black Holes*. 2024. arXiv: [2401.14331](https://arxiv.org/abs/2401.14331) [astro-ph.CO]. URL: <https://arxiv.org/abs/2401.14331>.
- [14] Adeela Afzal et al. ‘The NANOGrav 15 yr Data Set: Search for Signals from New Physics’. In: *The Astrophysical Journal Letters* 951.1 (2023), p. L11. ISSN: 2041-8213. DOI: [10.3847/2041-8213/acdc91](https://doi.org/10.3847/2041-8213/acdc91). URL: <http://dx.doi.org/10.3847/2041-8213/acdc91>.
- [15] Alejandro Gangui. *Topological Defects in Cosmology*. 2001. arXiv: [astro-ph/0110285](https://arxiv.org/abs/astro-ph/0110285) [astro-ph]. URL: <https://arxiv.org/abs/astro-ph/0110285>.
- [16] Michela D’Onofrio and Kari Rummukainen. ‘Standard model cross-over on the lattice’. In: *Physical Review D* 93.2 (Jan. 2016). ISSN: 2470-0029. DOI: [10.1103/PhysRevD.93.025003](https://doi.org/10.1103/PhysRevD.93.025003). URL: <http://dx.doi.org/10.1103/PhysRevD.93.025003>.

- [17] Giacomo Ferrante. *Cosmological consequences of spontaneous symmetry breaking*. 2025. arXiv: 2512.20474 [hep-ph]. URL: <https://arxiv.org/abs/2512.20474>.
- [18] T W B Kibble. ‘Topology of cosmic domains and strings’. In: *Journal of Physics A: Mathematical and General* 9.8 (Aug. 1976), p. 1387. DOI: 10.1088/0305-4470/9/8/029. URL: <https://doi.org/10.1088/0305-4470/9/8/029>.
- [19] Ya. B. Zel’dovich, I. Yu. Kobzarev, and L. B. Okun. ‘Cosmological consequences of a spontaneous breakdown of a discrete symmetry’. In: *Soviet Physics JETP* 40.1 (1974), pp. 1–5.
- [20] Simone Blasi et al. *Domain walls in the scaling regime: Equal Time Correlator and Gravitational Waves*. 2025. arXiv: 2511.16649 [hep-ph]. URL: <https://arxiv.org/abs/2511.16649>.
- [21] Takashi Hiramatsu, Masahiro Kawasaki, and Ken’ichi Saikawa. ‘On the estimation of gravitational wave spectrum from cosmic domain walls’. In: *Journal of Cosmology and Astroparticle Physics* 2014.02 (Feb. 2014), pp. 031–031. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2014/02/031. URL: <http://dx.doi.org/10.1088/1475-7516/2014/02/031>.
- [22] I. Dankovsky et al. ‘Cosmic domain walls on a lattice: Illusive effects of initial conditions’. In: *Physical Review D* 112.12 (Dec. 2025). ISSN: 2470-0029. DOI: 10.1103/19gp-7vp5. URL: <http://dx.doi.org/10.1103/19gp-7vp5>.
- [23] John A. Peacock. ‘Acceleration of the Universe’. In: *arXiv preprint* (2000). Sec. 2.3, Eqs. (7)–(8), p. 4. eprint: astro-ph/0009491. URL: <https://arxiv.org/abs/astro-ph/0009491>.
- [24] D. Stauffer. ‘Scaling theory of percolation clusters’. In: *Physics Reports* 54.1 (1979), pp. 1–74. DOI: 10.1016/0370-1573(79)90060-7.
- [25] Albert Einstein. ‘Die Grundlage der allgemeinen Relativitätstheorie’. In: *Annalen der Physik* 354.7 (1916), pp. 769–822. DOI: 10.1002/andp.19163540702.
- [26] Stephen R. Taylor. *The Nanohertz Gravitational Wave Astronomer*. 2021. arXiv: 2105.13270 [astro-ph.HE]. URL: <https://arxiv.org/abs/2105.13270>.
- [27] Michele Maggiore. *Gravitational Waves. Vol. 1: Theory and Experiments*. Oxford University Press, 2007. ISBN: 978-0-19-171766-6, 978-0-19-852074-0. DOI: 10.1093/acprof:oso/9780198570745.001.0001.

- [28] John Archibald Wheeler and Kenneth Ford. *Geons, Black Holes, and Quantum Foam: A Life in Physics*. New York: W. W. Norton, 1998. ISBN: 978-0-393-04642-7.
- [29] Nigel T. Bishop and Luciano Rezzolla. ‘Extraction of Gravitational Waves in Numerical Relativity’. In: *Living Reviews in Relativity* 19.1 (2016), p. 2. DOI: [10.1007/s41114-016-0001-9](https://doi.org/10.1007/s41114-016-0001-9). arXiv: [1606.02532](https://arxiv.org/abs/1606.02532) [gr-qc].
- [30] Marco Di Mauro, Salvatore Esposito, and Adele Naddeo. *Towards detecting gravitational waves: a contribution by Richard Feynman*. 2021. arXiv: [2111.00330](https://arxiv.org/abs/2111.00330) [physics.hist-ph]. URL: <https://arxiv.org/abs/2111.00330>.
- [31] A. Heger et al. ‘How Massive Single Stars End Their Life’. In: *The Astrophysical Journal* 591.1 (2003), pp. 288–300. ISSN: 1538-4357. DOI: [10.1086/375341](https://doi.org/10.1086/375341). URL: <http://dx.doi.org/10.1086/375341>.
- [32] R. W. Hellings and G. S. Downs. ‘Upper Limits on the Isotropic Gravitational Radiation Background from Pulsar Timing Analysis’. In: *The Astrophysical Journal Letters* 265 (1983), pp. L39–L42. DOI: [10.1086/183954](https://doi.org/10.1086/183954).
- [33] Jean-François Dufaux et al. ‘Theory and numerics of gravitational waves from preheating after inflation’. In: *Physical Review D* 76.12 (Dec. 2007). ISSN: 1550-2368. DOI: [10.1103/physrevd.76.123517](https://doi.org/10.1103/physrevd.76.123517). URL: <http://dx.doi.org/10.1103/PhysRevD.76.123517>.
- [34] Davide Barbini et al. *Biased Domain Wall Networks and their Gravitational Waves*. 2026. arXiv: [2607.18107](https://arxiv.org/abs/2607.18107) [astro-ph.CO]. URL: <https://arxiv.org/abs/2607.18107>.
- [35] Daniel G. Figueroa et al. ‘CosmoLattice: A modern code for lattice simulations of scalar and gauge field dynamics in an expanding universe’. In: *Computer Physics Communications* 283 (Feb. 2023), p. 108586. ISSN: 0010-4655. DOI: [10.1016/j.cpc.2022.108586](https://doi.org/10.1016/j.cpc.2022.108586). URL: <http://dx.doi.org/10.1016/j.cpc.2022.108586>.
- [36] Daniel G. Figueroa et al. ‘The art of simulating the early universe. Part I. Integration techniques and canonical cases’. In: *Journal of Cosmology and Astroparticle Physics* 2021.04 (Apr. 2021), p. 035. ISSN: 1475-7516. DOI: [10.1088/1475-7516/2021/04/035](https://doi.org/10.1088/1475-7516/2021/04/035). URL: <http://dx.doi.org/10.1088/1475-7516/2021/04/035>.

- [37] Naoya Kitajima et al. *Gravitational Waves from Domain Wall Collapse, and Application to Nanohertz Signals with QCD-coupled Axions*. 2023. arXiv: 2306.17146 [hep-ph]. URL: <https://arxiv.org/abs/2306.17146>.
- [38] E. Babichev et al. *Biased domain walls: faster annihilation, weaker gravitational waves*. 2025. arXiv: 2504.07902 [hep-ph]. URL: <https://arxiv.org/abs/2504.07902>.
- [39] Fulvio Melia. ‘The Friedmann–Lemaître–Robertson–Walker metric and the principle of equivalence’. In: *Z. Naturforsch. A* 78.6 (2023), pp. 525–533. DOI: 10.1515/zna-2022-0307.
- [40] Chiara Caprini et al. ‘General properties of the gravitational wave spectrum from phase transitions’. In: *Phys. Rev. D* 79 (8 Apr. 2009), p. 083519. DOI: 10.1103/PhysRevD.79.083519. URL: <https://link.aps.org/doi/10.1103/PhysRevD.79.083519>.
- [41] Gabriele Franciolini, Davide Racco, and Fabrizio Rompineve. *Footprints of the QCD Crossover on Cosmological Gravitational Waves at Pulsar Timing Arrays*. 2024. arXiv: 2306.17136 [astro-ph.CO]. URL: <https://arxiv.org/abs/2306.17136>.
- [42] Yann Gouttenoire, Nicholas Leister, and Pedro Schwaller. *Gravitational Waves from Black Hole Reheating: The Scalar-Induced Component*. 2026. arXiv: 2605.21474 [hep-ph]. URL: <https://arxiv.org/abs/2605.21474>.
- [43] William G Lamb, Stephen R Taylor, and Rutger van Haasteren. ‘Rapid refitting techniques for Bayesian spectral characterization of the gravitational wave background using pulsar timing arrays’. In: *Physical Review D* 108.10 (2023), p. 103019.
- [44] Justin A. Ellis et al. *ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite*. Zenodo. Sept. 2020. DOI: 10.5281/zenodo.4059815. URL: <https://doi.org/10.5281/zenodo.4059815>.
- [45] Stephen R. Taylor et al. *enterprise_extensions*. Version 2.3.3. Zenodo, 2021. DOI: 10.5281/zenodo.15132505. URL: https://github.com/nanograv/enterprise_extensions.
- [46] Roberto Trotta. *Bayesian Methods in Cosmology*. 2017. arXiv: 1701.01467 [astro-ph.CO]. URL: <https://arxiv.org/abs/1701.01467>.

- [47] Justin Ellis and Rutger van Haasteren. *jellis18/PTMCMCSampler: Official Release*. 2017. DOI: [10.5281/zenodo.1037579](https://doi.org/10.5281/zenodo.1037579).
- [48] Aki Vehtari et al. ‘Rank-Normalization, Folding, and Localization: An Improved $\hat{R}$ for Assessing Convergence of MCMC (with Discussion)’ . In: *Bayesian Analysis* 16.2 (2021). ISSN: 1936-0975. DOI: [10.1214/20-ba1221](https://doi.org/10.1214/20-ba1221). URL: <http://dx.doi.org/10.1214/20-BA1221>.
- [49] Andrea Mitridate and David Wright. *PTArcade documentation: chains_utils*. https://andrea-mitridate.github.io/PTArcade/reference/ptarcade/chains_utils/. Accessed: 2026-08-01. 2023.
- [50] A. Bazavov et al. ‘Chiral crossover in QCD at zero and non-zero chemical potentials’. In: *Physics Letters B* 795 (Aug. 2019), pp. 15–21. ISSN: 0370-2693. DOI: [10.1016/j.physletb.2019.05.013](https://doi.org/10.1016/j.physletb.2019.05.013). URL: <http://dx.doi.org/10.1016/j.physletb.2019.05.013>.
- [51] Keith A. Olive. *Primordial Big Bang Nucleosynthesis*. 1999. arXiv: [astro-ph/9901231](https://arxiv.org/abs/astro-ph/9901231) [astro-ph]. URL: <https://arxiv.org/abs/astro-ph/9901231>.
- [52] CERN. *The third run of the Large Hadron Collider has successfully started*. July 5, 2022. URL: <https://home.cern/news/news/cern/third-run-large-hadron-collider-has-successfully-started> (visited on 08/02/2026).
- [53] Gabriella Agazie et al. *The NANOGrav 15 yr Data Set: Piecewise Power-Law Reconstruction of the Gravitational-Wave Background*. 2026. arXiv: [2601.09481](https://arxiv.org/abs/2601.09481) [astro-ph.HE]. URL: <https://arxiv.org/abs/2601.09481>.
- [54] David J. C. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge: Cambridge University Press, 2003. ISBN: 978-0-521-64298-9.
- [55] Roberto Trotta. ‘Bayes in the sky: Bayesian inference and model selection in cosmology’. In: *Contemporary Physics* 49.2 (Mar. 2008), pp. 71–104. ISSN: 1366-5812. DOI: [10.1080/00107510802066753](https://doi.org/10.1080/00107510802066753). URL: <http://dx.doi.org/10.1080/00107510802066753>.
- [56] Xavier Siemens et al. ‘The stochastic background: scaling laws and time to detection for pulsar timing arrays’. In: *Classical and Quantum Gravity* 30.22 (Nov. 2013), p. 224015. ISSN: 1361-6382. DOI: [10.1088/0264-9381/30/22/224015](https://doi.org/10.1088/0264-9381/30/22/224015). URL: <http://dx.doi.org/10.1088/0264-9381/30/22/224015>.

-
- [57] Pau Amaro-Seoane et al. *Laser Interferometer Space Antenna*. 2017. arXiv: 1702.00786 [astro-ph.IM]. URL: <https://arxiv.org/abs/1702.00786>.
 - [58] Satyabrata Datta and Rome Samanta. *PTA-Compatible Domain Walls at LISA and Taiji: Bayesian Reconstruction and Multiband Inference*. 2026. arXiv: 2606.09713 [hep-ph]. URL: <https://arxiv.org/abs/2606.09713>.
 - [59] Satyabrata Datta and Rome Samanta. *Mapping Domain-Wall Bayesian Reconstruction with LISA*. 2026. arXiv: 2606.15387 [astro-ph.HE]. URL: <https://arxiv.org/abs/2606.15387>.

Acknowledgement

I would like to express my gratitude to Kai Schmitz for providing me with the opportunity to write this thesis in his working group, for the many fruitful discussions and for his invaluable guidance. My co-supervisor, Alessandra Amosso, provided me with numerous helpful suggestions and much appreciated support, for which I am very grateful. Thanks also go to all members of the PCM group for the welcoming atmosphere and the enjoyable moments along the way.

The numerical computations for this thesis were performed on the HPC cluster PALMA II of the University of Münster, funded by the DFG (INST 211/667-1).

To Franziska, thank you for listening to explanations of domain walls you never asked for, as well as for your warmth and patience throughout the work on this thesis. Finally, my deepest gratitude goes to my parents. Without your continuous support, encouragement and love, pursuing my bachelor's degree would never have been possible.