TOPICS IN BROADBAND GRAVITATIONAL-WAVE ASTRONOMY

by

Sydney J. Chamberlin

A Dissertation Submitted in
Partial Fulfillment of the
Requirements for the degree of

Doctor of Philosophy
in Physics

at
The University of Wisconsin–Milwaukee
May 2015

ABSTRACT

TOPICS IN BROADBAND GRAVITATIONAL-WAVE ASTRONOMY
The University of Wisconsin–Milwaukee, May 2015
Under the Supervision of Professor Jolien D. E. Creighton

The direct detection of gravitational waves promises to open a new observational window onto the universe, and a number of large scale efforts are underway worldwide to make such a detection a reality. In this work, we attack some of the current problems in gravitational-wave detection over a wide range of frequencies.

In the first part of this work, low frequency gravitational-wave detection is considered using pulsar timing arrays (PTAs). PTAs are a promising tool for probing the universe through gravitational radiation. Supermassive black hole binaries (SMBHBs), cosmic strings, relic gravitational waves from inflation, and first order phase transitions in the early universe are expected to contribute to a stochastic background of gravitational waves in the PTA frequency band of 1 nHz – 100 nHz. The detection of low-frequency stochastic backgrounds of gravitational waves in the PTA band is considered in the context of constructing an optimal cross-correlation statistic in the time domain. Also presented are some useful applications of this statistic, and discussion on its limitations in actual gravitational-wave searches.

Also considered are methods by which gravitational waves in the PTA frequency band can serve as a mechanism for testing general relativity (GR). In addition to providing a new paradigm for exploring the universe, the direct detection of gravitational waves will allow general relativity to be tested against other metric theories of gravity in the regime of strong gravitational fields. This work involves the analysis of the overlap reduction function (ORF), a geometrical factor that appears in the expected cross correlation of signals, for general metric theories of gravity. The ORF characterizes the loss of sensitivity due to detectors not being co-located or coaligned, and it is an important element in defining the optimal cross-correlation statistic. It is shown that PTA detection sensitivity increases for non-transverse gravitational waves. Additionally, the ORFs for a subset of the NANOGrav PTA are described, and are used to show that sensitivity to vector and longitudinal modes can increase dramatically for pulsar pairs with small angular

separations. Implications of these results are discussed.

In the second part of this work, the detection of gravitational-wave bursts in the 10 Hz – 1000 Hz frequency band is considered using ground-based laser interferometers. An excess power method for conducting unmodeled searches for gravitational-wave bursts is discussed, and its implementation into a search pipeline is described in detail. The performance of this pipeline is probed using software injections. Also discussed are potential applications of the ExcessPower pipeline to detector characterization efforts, which aim to improve interferometric searches by characterizing and mitigating non-Gaussian noise transients in the detectors.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

A great number of people have supported me on the path to this dissertation. First I would like to thank my advisor, Jolien Creighton, and the other faculty members of the Center for Gravitation, Cosmology and Astrophysics (CGCA) at UWM who have provided help and guidance on numerous occasions, and without whom this work would not have been possible. I would like to extend a special thanks to the postdoctoral researchers and graduate students of the CGCA for their friendship, laughter and support in this long process. I would like to thank Shane Larson for his career guidance and support. I would also like to thank my undergraduate research advisor at Utah State, Charles Torre, for introducing me to the beautiful interplay between gravity and geometry, and continuing to serve as a mentor throughout my career.

To Justin Ellis, Chiara Mingarelli, Chris Pankow, Adam Mercer, Sarah Caudill, Astrid Lamberts, Laleh Sadeghian, Laura Nuttall, Megan DeCesar, Joe Simon and Alex Urban: thank you for your continued help and support in research, for your friendship, and for helping me to gain confidence as a researcher.

To Maddie Wade, Les Wade, Tom Linz, and Whitney Linz: I am forever grateful to you for your friendship and laughter in completing this phase of our lives, and for helping me to make it through the hard times in grad school. I look forward to many more years of friendship, and will never forget the adventures we have had in our six years at UWM.

And finally, to my family, who have given me so much. To Kathryn Ricks for your many years of close friendship and support through hard times. To my wonderful fiancé, Brian Helf, for his love, help and encouragement in this process, and for supporting me in following my dreams. To my pup Lexi who spent hours curled around my feet while I wrote this dissertation, and who lifted my mood with many laughs when I was stressed. To my grandmother, who showed me what is important in life †. And to my parents, who taught me to look at the stars — and to reach for them.

Preface

Chapter 3 is based on the paper:

Time-domain Implementation of the Optimal Cross-Correlation Statistic for Stochastic Gravitational-Wave Background Searches in Pulsar Timing Data, S. J. Chamberlin, J. D. E. Creighton, P. B. Demorest, J. Ellis, L. R. Price, J. D. Romano, X. Siemens, Phys. Rev. D **91**, 044048 (2015).

I was the lead author on this paper, but J. Ellis developed the Section 3.1.1 and contributed to Section 3.1.2.

*

Chapter 4 is based on the paper:

Stochastic backgrounds in alternative theories of gravity: overlap reduction functions for pulsar timing arrays, S. J. Chamberlin, X. Siemens, Physical Review D **85**, 082001 (2012).

Chapter 1

Introduction

> "*We especially need imagination in science. It is not all mathematics, nor all logic, but it is somewhat beauty and poetry.*"
>
> — Maria Mitchell, Astronomer (1818-1889)

One of the predominant features of matter and its interactions is the interplay between a *field* and its *source*. This is easy to see in Maxwell's theory of electromagnetism; Maxwell's equations describe the interaction between electric and magnetic fields and their sources.

Many parallels can be drawn between Maxwell's electromagnetism and Einstein's theory of relativity. In fact, the disaccord between the theories of Maxwell and Newton in the late 19th century was one of the driving forces that led to the development of special relativity. To motivate the development of special, and consequently general relativity, it is useful to consider the parallels that can be drawn between Newtonian gravity and Einstein's relativity.

For Newton, the stage for all physical dynamics is a three-dimensional Euclidean space. Although this space is subject to the laws of spatial relativity, time is absolute, viewed the same by all observers. Using such a framework to describe dynamics in the universe, the gravitational field is described as a conservative vector field that can be obtained from a potential function $\vec{F} = -m\nabla\phi$ where m is the mass of a test particle[1] in the field. In the coordinate-free formulation of Newton's theory, the gravitational potential ϕ is related to the mass density ρ by the *field equation* $\nabla^2\phi = -4\pi G\rho$, where G is Newton's

[1] Test particles are not considered to be sources of the field.

gravitational constant. If at every point ϕ is bounded and ρ vanishes, ϕ is forced to be constant in space. It follows that the source of the gravitational field is the mass density ρ. A field equation therefore relates the field to its source, and characterizes the interplay between the two.

A simple generalization from Netwon's theory to Einstein's is not sensible, because the stage for physical interactions changes in general relativity. One of the key results of special relativity is that the laws of physics are equally valid in all inertial reference frames, and the universality of the speed of light as measured by any observer forces the *relativity of simultaneity*: there can be no absolute notion of simultaneity. Events that are simultaneous in one frame need not be simultaneous in another frame. In other words, time is not an absolute quantity in Einstein's universe.

In special relativity, physical events take place on a four-dimensional manifold with a flat Lorentzian metric; global inertial frames are a canonical feature of spacetime. General relativity, on the other hand, does not postulate a particular metric a priori. In this sense, general relativity is not canonically equipped with global inertial reference frames, and the existence of such frames depends on whether any gravitational field is present [1].

In Newton's theory of gravity, the source of the gravitational field is the mass density ρ. In general relativity, matter takes on the analogue *matter distribution*, and fields (*e.g.*, electromagnetic) may contribute to the gravitational field. The source of the gravitational field is therefore more complex than in Newton's theory. The object that generalizes the mass density ρ is called the *stress-energy* tensor.[2] The stress-energy tensor is a symmetric, rank $(0, 2)$ tensor, and contains all relevant information about mass density, energy density, momentum density, etc. It is now possible to relate the source of the gravitational field — the energy momentum tensor — to the field itself, which is dependent on the curvature of spacetime. The field equation is

$$R_{\alpha\beta} - \frac{1}{2}Rg_{\alpha\beta} = \frac{8\pi G}{c^4}T_{\alpha\beta}, \tag{1.0.1}$$

where $R_{\alpha\beta}$ is the Ricci tensor, R is the Ricci scalar (obtained through contraction with the metric $g_{\alpha\beta}$), G is the gravitational constant, c is the speed of light, and $T_{\alpha\beta}$ is the

[2]This is also frequently referred to as the energy-momentum or stress-energy-momentum tensor; the terms are used interchangeably.

stress-energy tensor. The left hand side of this equation is often simplified in appearance with the use of the *Einstein tensor*, which is defined as $G_{\alpha\beta} = R_{\alpha\beta} - \frac{1}{2}Rg_{\alpha\beta}$. As stated previously, a field equation describes the interaction between a field and its source; it describes the manner in which the source creates the field. This interdependence lays down the conceptual framework for the theory of general relativity.

1.1 Gravity and spacetime

The mathematical description of general relativity begins with the notion of coordinates and distances. Events in general relativity take place on a four-dimensional manifold with a metric tensor that may or may not be flat. To characterize the geometry taking place on this manifold, one can define the line element ds^2 that provides the distance between nearby points that are separated by coordinate intervals dx^α:

$$ds^2 = g_{\mu\nu}(x^\alpha)dx^\mu dx^\nu. \tag{1.1.1}$$

The metric $g_{\mu\nu}(x^\alpha)$ is a function that depends on the spacetime coordinates x^α, where α can take the values $\{0, 1, 2, 3\}$ (which correspond to a single time and three space coordinates, respectively). The metric completely characterizes the geometry of the manifold.

Other geometric quantities must be generalized from that of flat spacetime to accommodate curvature in a rigorous manner. Vectors can be described as directional derivatives: for a curve (or world line) $x^\alpha(t)$ parameterized by t passing through some point in the spacetime manifold, the derivative along that curve contains the components of a four-vector that lies in the tangent space to that point. One can therefore express the four-vector u^α as

$$u^\alpha = \frac{dx^\alpha}{dt}, \tag{1.1.2}$$

which not only satisfies the intuitive notion of a vector but also reduces to the definition used in flat space when the metric is flat. It is also worth pointing out here that general relativity is invariant under the symmetry group of all all possible coordinate transformations $x^\alpha \to x'^\alpha(x)$. Under this *gauge* symmetry, the metric transforms as

$$g_{\mu\nu} \to g'_{\mu\nu} = \frac{\partial x^\alpha}{\partial x'^\mu}\frac{\partial x^\beta}{\partial x'^\nu}g_{\alpha\beta}. \tag{1.1.3}$$

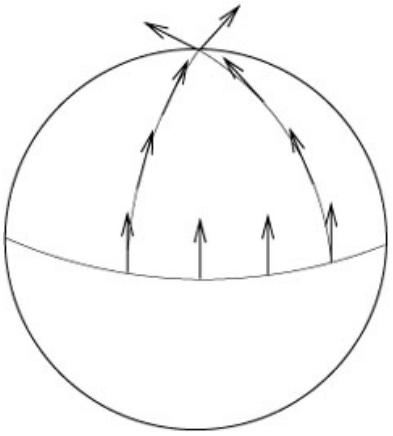

Figure 1: Vectors that are parallel-transported on a curved surface, such as a sphere, return rotated from their initial position. This leads to a very natural sort of definition for curvature.

1.1.1 Spacetime curvature

It is important to be rigorous in defining spacetime curvature since the human-intuitive view of curvature comes from imagining two-dimensional objects embedded in a three-dimensional space. This extrinsic notion of curvature turns out to be problematic for general relativity, since we are interested in the curvature of the spacetime manifold itself (which is not necessarily embedded in any higher dimensional space). To think about curvature in a way that makes sense without embedding, the concept of parallel transport is used. In a flat plane (or in flat spacetime), a vector that is parallel-transported around a closed path will always end up at its original position. On a curved surface, such as a sphere, the parallel-transported vector returns rotated from its initial position. This is shown in Fig. 1.

This example provides the motivation for a rigorous definition of curvature. Parallel-transporting vectors along a curve is related to the derivative of a vector field in the direction of the curve; given some notion of derivative, a vector is parallel-transported if its derivative along the curve is zero. Curvature is what occurs when successive differentiations on vector fields fail to commute.

To make this mathematically precise, one must define a derivative operator for vector fields on a manifold. This derivative, called the covariant derivative, is defined as

$$\nabla_\alpha v^\gamma = \frac{\partial v^\gamma}{\partial x^\alpha} + \Gamma^\gamma{}_{\alpha\beta} v^\beta \tag{1.1.4}$$

where $\Gamma^{\gamma}{}_{\alpha\beta}$ are the connection coefficients (or Christoffel symbol). In terms of the metric, the Christoffel symbol can be written as

$$\Gamma^{\gamma}{}_{\alpha\beta} = \frac{1}{2} g^{\gamma\delta} \left(\frac{\partial}{\partial x^{\alpha}} g_{\beta\delta} + \frac{\partial}{\partial x^{\beta}} g_{\delta\alpha} - \frac{\partial}{\partial x^{\delta}} g_{\alpha\beta} \right). \tag{1.1.5}$$

The equation of parallel transport for the components of a vector v along a curve with tangent vector u is then

$$u^{\mu} \nabla_{\mu} v^{\alpha} = u^{\mu} \frac{\partial v^{\alpha}}{\partial x^{\mu}} + \Gamma^{\alpha}{}_{\mu\nu} u^{\mu} v^{\nu} = 0. \tag{1.1.6}$$

One way to describe curvature is to relate it to a geodesic: a curve whose tangent vector is parallel-transported along itself. A curve $x^{\alpha}(t)$ with tangent vector u is a geodesic if u is parallel-transported along itself ($u = v$), i.e.

$$u^{\mu} \nabla_{\mu} u^{\alpha} = u^{\mu} \frac{\partial u^{\alpha}}{\partial x^{\mu}} + \Gamma^{\alpha}{}_{\mu\nu} u^{\mu} u^{\nu} = 0. \tag{1.1.7}$$

Mathematically, we can quantify curvature with the Riemann curvature tensor, which is defined by

$$R_{\alpha\beta\gamma}{}^{\delta} w^{\gamma} = - \left(\nabla_{\alpha} \nabla_{\beta} - \nabla_{\beta} \nabla_{\alpha} \right) w^{\delta} \tag{1.1.8}$$

for arbitrary w. This tensor completely characterizes the curvature of spacetime. It is often more useful, however, to work with contracted versions of the curvature tensor that appear in the Einstein field equations (Eq. (1.0.1)). The Ricci tensor $R_{\alpha\beta}$ is obtained from two contractions of the Riemann tensor,

$$R_{\alpha\beta} = R_{\alpha\gamma\beta}{}^{\gamma} \tag{1.1.9}$$

and the Ricci scalar R is simply the trace of the Ricci tensor:

$$R = g^{\alpha\beta} R_{\alpha\beta}. \tag{1.1.10}$$

For the study of gravitational radiation (or gravitational waves), one can expand the Einstein equations around the flat spacetime metric, and use the Riemann tensor to understand some of the properties of gravitational waves. This is what will be discussed next.

1.1.2 Linearized Gravity

When gravity is relatively weak, one can express the spacetime metric $g_{\alpha\beta}$ as the flat Minkowski metric $\eta_{\alpha\beta}$ along with a small perturbation $h_{\alpha\beta}$ (small meaning that Cartesian-like coordinates $|h_{\alpha\beta}| << 1$ in units where c =1),

$$g_{\alpha\beta} = \eta_{\alpha\beta} + h_{\alpha\beta}, \tag{1.1.11}$$

where the inverse metric is

$$g^{\alpha\beta} = \eta^{\alpha\beta} - h^{\alpha\beta}. \tag{1.1.12}$$

Note that in this weak-field regime, although Eq. (1.1.12) specifies the inverse metric, indices will be raised and lowered with the flat-metric $\eta_{\alpha\beta}$.

To understand the gravitational field in the weak-field limit, the Einstein field equations (Eq. (1.0.1)) must be solved to first order in the metric perturbation $h_{\alpha\beta}$. A few definitions are needed to proceed. To linear order in $h_{\alpha\beta}$, the Christoffel symbols may be written as

$$\Gamma^{\gamma}_{\alpha\beta} = \frac{1}{2}\eta^{\gamma\delta}\left(\partial_\alpha h_{\beta\delta} + \partial_\beta h_{\delta\alpha} - \partial_\delta h_{\alpha\beta}\right) \tag{1.1.13}$$

where we have simplified notation so that $\partial/\partial x^\mu = \partial_\mu$. Given the Christoffel symbols, the linearized Ricci tensor can be written to first order in h as

$$\begin{aligned} R_{\alpha\beta} &= \partial_\gamma \Gamma^{\gamma}{}_{\alpha\beta} - \partial_\alpha \Gamma^{\gamma}{}_{\gamma\beta} \\ &= \partial^\gamma \partial_\beta h_{\alpha\gamma} + \partial^\gamma \partial_\alpha h_{\beta\gamma} - \frac{1}{2}\partial^\gamma \partial_\gamma h_{\alpha\beta} - \frac{1}{2}\partial_\alpha \partial_\beta h \\ &= \frac{1}{2}\left(\partial^\gamma \partial_\beta h_{\alpha\gamma} + \partial^\gamma \partial_\alpha h_{\beta\gamma} - \partial^\gamma \partial_\gamma h_{\alpha\beta} - \partial_\alpha \partial_\beta h\right). \end{aligned} \tag{1.1.14}$$

By combining the Ricci scalar (which is just the trace of Eq. (1.1.14)) with the Ricci tensor, one can express the Einstein tensor as

$$G_{\alpha\beta} = \frac{1}{2}\left(\partial^\gamma \partial_\beta h_{\alpha\gamma} + \partial^\gamma \partial_\alpha h_{\beta\gamma} - \partial^\gamma \partial_\gamma h_{\alpha\beta} - \partial_\alpha \partial_\beta h - \eta_{\alpha\beta}\left(\partial^\gamma \partial^\delta h_{\gamma\delta} - \partial^\gamma \partial_\gamma h\right)\right), \tag{1.1.15}$$

and the resulting Einstein field equations are

$$\left(\partial^\gamma \partial_\beta h_{\alpha\gamma} + \partial^\gamma \partial_\alpha h_{\beta\gamma} - \partial^\gamma \partial_\gamma h_{\alpha\beta} - \partial_\alpha \partial_\beta h - \eta_{\alpha\beta}\left(\partial^\gamma \partial^\delta h_{\gamma\delta} - \partial^\gamma \partial_\gamma h\right)\right) = \frac{16\pi}{c^4} T_{\alpha\beta}. \tag{1.1.16}$$

In principle, one could stop at this point and seek out solutions to the field equations. However, the form of the equations in Eq. (1.1.16) are not particularly useful. It is helpful

to make a few choices that recast field equations into a more intuitive form. The first choice is to work with the *trace-reversed* metric, which is defined as

$$\overline{h}_{\alpha\beta} = h_{\alpha\beta} - \frac{1}{2}\eta_{\alpha\beta}h. \tag{1.1.17}$$

With this definition, the field equations take on a much simpler appearance,

$$-\partial^\gamma\partial_\gamma\bar{h}_{\alpha\beta} + \partial^\gamma\partial_\beta\bar{h}_{\alpha\gamma} + \partial^\gamma\partial_\alpha\bar{h}_{\beta\gamma} - \eta_{\alpha\beta}\partial^\gamma\partial^\delta\bar{h}_{\gamma\delta} = \frac{16\pi G}{c^4}T_{\alpha\beta}. \tag{1.1.18}$$

It is interesting to note that the first term on the left-hand side of Eq. (1.1.18) looks similar in form to the d'Alembertian wave operator (by definition the d'Alembertian is $\Box = \partial^\alpha\partial_\alpha$).

The second choice made to put the field equations in a more tractable form is to choose a specific gauge. There is a gauge freedom in general relativity which corresponds to the group of *diffeomorphisms*, which are (invertible) bijective functions that map one spacetime manifold to another in a smooth fashion. A complete discussion about diffeomorphisms and how they relate to gauge freedom goes beyond the scope of this dissertation (see the textbook by Wald, Ref. [1], for a detailed discussion on the topic), but essentially gauge freedom in linearized gravity means that for some vector field ξ^α, coordinates can be chosen so that

$$h_{\alpha\beta} \to h_{\alpha\beta} + \partial_\alpha\xi_\beta + \partial_\beta\xi_\alpha, \tag{1.1.19}$$

or in terms of the trace-reversed metric,

$$\bar{h}_{\alpha\beta} \to \bar{h}_{\alpha\beta} + \left(\partial_\alpha\xi_\beta + \partial_\beta\xi_\alpha - \eta_{\alpha\beta}\partial_\gamma\xi^\gamma\right). \tag{1.1.20}$$

Given the form of the field equations in Eq. (1.1.18), the best choice of coordinate transformation is one that will eliminate all terms in the Einstein tensor except those that look similar to the d'Alembertian wave operator; this will give the field equations an immediately useful (and somewhat physically intuitive) form.

In practice, this means making a coordinate choice that eliminates the divergence of the trace-reversed metric. This choice is in fact very similar to the Lorenz gauge choice made in electromagnetism. Mathematically, this means finding a solution to the equation

$$\partial^\beta\partial_\beta\xi_\alpha = -\partial^\beta\bar{h}_{\alpha\beta} \tag{1.1.21}$$

for ξ^α, which defines the choice of gauge

$$\partial^\beta \bar{h}_{\alpha\beta} = 0. \tag{1.1.22}$$

This gauge condition is referred to as the Lorenz gauge. With this gauge in place, the Einstein field equations (1.1.18) simplify dramatically to

$$\partial^\gamma \partial_\gamma \bar{h}_{\alpha\beta} = \Box \bar{h}_{\alpha\beta} = \frac{-16\pi G}{c^4} T_{\alpha\beta}. \tag{1.1.23}$$

In this form it is clear that the solutions to the field equations are wave solutions, where the stress-energy tensor serves as the source. The waves admitted by this equation are called *gravitational waves*, and they propagate at the speed of light c. In the next section, additional gauge freedom will be used to study the properties of gravitational waves and their effects on test masses.

1.2 Gravitational waves in the transverse-traceless gauge

To study the propagation of the gravitational waves admitted by Eq. (1.1.23), as well as the response of test masses to these waves, it is necessary to consider the field equations far from the source (where the stress-energy tensor vanishes), i.e. where

$$\Box \bar{h}_{\alpha\beta} = 0. \tag{1.2.1}$$

To make this problem as simple as possible, one can exploit additional gauge freedom that has not been used. This closely resembles the gauge freedom seen in electrodynamics. In that case, the choice of Lorenz gauge — $\partial_\alpha A^\alpha = 0$ — simplifies the equation of motion (which depends on the source current j^β) obtained from the Maxwell Lagrangian,

$$\partial_\alpha \left(\partial^\alpha A^\beta - \partial^\beta A^\alpha \right) = j^\beta \longrightarrow \Box A^\beta = j^\beta, \tag{1.2.2}$$

but there remains *additional* gauge freedom which allows

$$A_\alpha \longrightarrow A_\alpha - \partial_\alpha \phi \text{ (for } \Box\phi = 0). \tag{1.2.3}$$

The same additional gauge freedom is present in linearized gravity. Before proceeding, it is interesting to extend the electrodynamics analogy a bit further: when far from the source current j^β,

$$\Box A^\beta = 0 \tag{1.2.4}$$

and the additional gauge freedom of Eq. (1.2.3) (which appears via the scalar ϕ) can be used to set $A^0 = 0$. When $A^0 = 0$, the Lorenz gauge effectively represents a condition on transversality:

$$\partial_i A^i = 0. \tag{1.2.5}$$

For linearized gravity, the additional gauge freedom present means that just as in Eq. (1.2.3) the Lorenz gauge is still valid with another coordinate transformation

$$x^\alpha \longrightarrow x^\alpha + \xi^\alpha, \tag{1.2.6}$$

as long as $\Box\xi_\alpha = 0$. However if $\Box\xi_\alpha = 0$, then the quantity $\Box\xi_{\alpha\beta} = 0$ as well, where

$$\xi_{\alpha\beta} = \partial_\alpha\xi_\beta + \partial_\beta\xi_\alpha - \eta_{\alpha\beta}\partial_\gamma\xi^\gamma. \tag{1.2.7}$$

But Eq. (1.2.7) is precisely what appeared when making the coordinate transformation in Eq. (1.1.20), so the functions $\xi_{\alpha\beta}$ can be subtracted from the six independent components of $\bar{h}_{\alpha\beta}$ (that satisfy $\Box\bar{h}_{\alpha\beta} = 0$). Because the functions $\xi_{\alpha\beta}$ depend on the four functions ξ_α (and since the functions $\xi_{\alpha\beta}$ also satisfy $\Box\xi_{\alpha\beta} = 0$), the additional gauge freedom means that four conditions can be imposed on $\bar{h}_{\alpha\beta}$.

As in the electrodynamic analogue (Eq. (1.2.5)), these conditions can be chosen to impose transversality; if ξ^0 is specified so that the trace $\bar{h} = 0$, then $\bar{h}_{\alpha\beta} = h_{\alpha\beta}$, and the remaining three functions $\xi^i(x)$ can be chosen so that $h^{0i}(x) = 0$. In the Lorenz gauge with $\bar{h}_{\alpha\beta} = h_{\alpha\beta}$ and $\mu = 0$,

$$\partial^0 h_{00} + \partial_i h_{0i} = 0, \tag{1.2.8}$$

and with a fixed value $h_{0i} = 0$,

$$\partial^0 h_{00} = 0. \tag{1.2.9}$$

These choices characterize the *transverse-traceless* gauge (or TT gage), which can be summarized mathematically as

$$h^{0\alpha} = 0, \quad h^i_i = 0, \quad \partial^j h_{ij} = 0. \tag{1.2.10}$$

It is worth pointing out that these conditions are not valid near the source, i.e. for the field equations in Eq. (1.1.23). There is still the freedom to perform a transformation

satisfying $\Box\xi_{\alpha\beta} = 0$, but because $\Box\bar{h}_{\alpha\beta} \neq 0$ the components of $\bar{h}_{\alpha\beta}$ cannot be set to zero as in the TT gauge.

An important solution of the wave equation (1.2.1) is the *plane wave solution*,

$$h_{ij}^{TT}(x^\alpha) = \epsilon_{ij}(\vec{k})e^{ik_\alpha x^\alpha} \tag{1.2.11}$$

where $\epsilon_{ij}(\vec{k})$ is the *polarization tensor* and $k_\alpha = (\omega/c, \vec{k})$ is the wave vector satisfying $\omega/c = |\vec{k}|$. For a gravitational wave propagating in the z direction, this solution is of the form

$$h_{ij}^{TT}(x) = \begin{bmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{bmatrix}_{ij} \cos\left[\omega(t - z/c)\right], \tag{1.2.12}$$

where h_+ and $h_\times$ are the two possible gravitational-wave polarizations corresponding to the metric perturbation's remaining two degrees of freedom.

1.2.1 The effect of gravitational-waves on matter

A possible gravitational-wave coordinate system is shown in Fig. 2. Given the unit vectors $\hat{i}$, $\hat{j}$ and $\hat{k}$ lying along the x, y and z axes (blue arrows in Fig. 2) and the gravitational-wave coordinate system rotated by the polar and azimuthal angles θ and ϕ

$$\begin{aligned} \hat{i}' &= (\cos\theta\cos\phi, \cos\theta\sin\phi, -\sin\theta) \\ \hat{j}' &= (-\sin\phi, \cos\phi, 0) \\ \hat{k}' &= (\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta), \end{aligned} \tag{1.2.13}$$

the most general choice of coordinates is chosen by rotating about the gravitational wave's propagation axis with respect to the polarization angle ψ:

$$\begin{aligned} \hat{m} &= \hat{i}'\cos\psi + \hat{j}'\sin\psi \\ \hat{n} &= -\hat{i}'\sin\phi + \hat{j}'\cos\psi \\ \hat{\Omega} &= \hat{k}'. \end{aligned} \tag{1.2.14}$$

The coordinate system defined by $\hat{i}$, $\hat{j}$ and $\hat{k}$ is related to the coordinate system $\hat{m}$, $\hat{n}$ and $\hat{\Omega}$ by the angles θ and ϕ, and with this coordinate system in place it is possible to explicity define the polarization tensors.

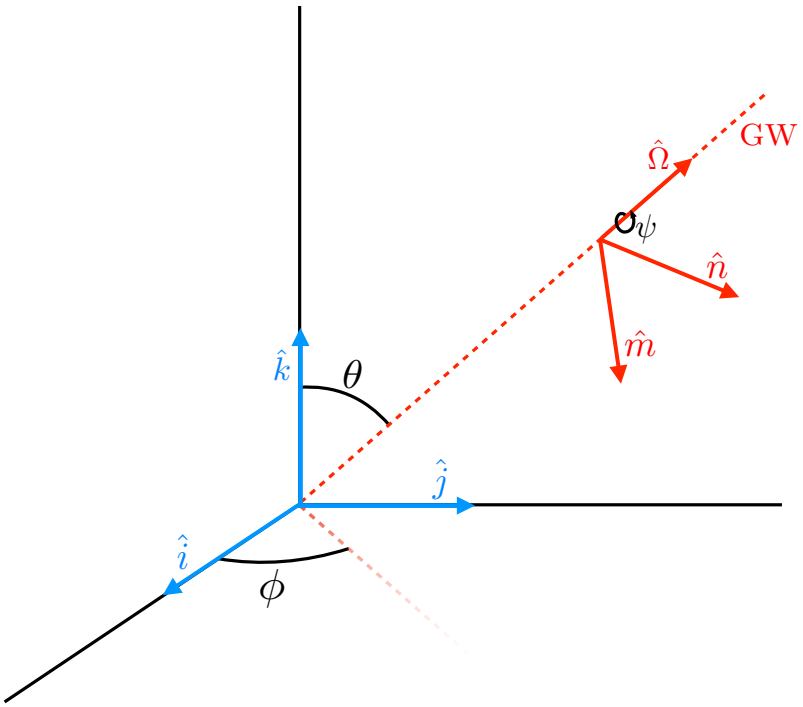

Figure 2: Illustration of a possible gravitational-wave coordinate system. The wave propagates in the direction $\hat{\Omega}$, and the vectors m and n, obtained through rotations of the angle ψ, represent the most general choice of coordinates.

These are

$$\begin{aligned} \epsilon_+ &= \hat{m} \otimes \hat{m} - \hat{n} \otimes \hat{n} \\ \epsilon_\times &= \hat{m} \otimes \hat{n} + \hat{n} \otimes \hat{m}. \end{aligned} \tag{1.2.15}$$

Their definition will be useful later on when characterizing the response of gravitational-wave detectors to gravitational waves.

It is worthwhile at this point to describe the manner in which gravitational waves are physically manifest. The equation of *geodesic deviation* can be used to study the effect of a gravitational wave on the relative motion of two freely falling particles. For two nearby freely falling bodies that are nearly "at rest" in a global coordinate system $\eta_{\alpha\beta}$, with deviation X^α,

$$\frac{d^2 X^\beta}{dt^2} = \sum_\alpha R_{\alpha 00}{}^\beta X^\beta, \tag{1.2.16}$$

and it can be shown that

$$R_{\alpha 00\beta} = \frac{1}{2} \frac{\partial h_{\alpha\beta}}{\partial^2 t}. \tag{1.2.17}$$

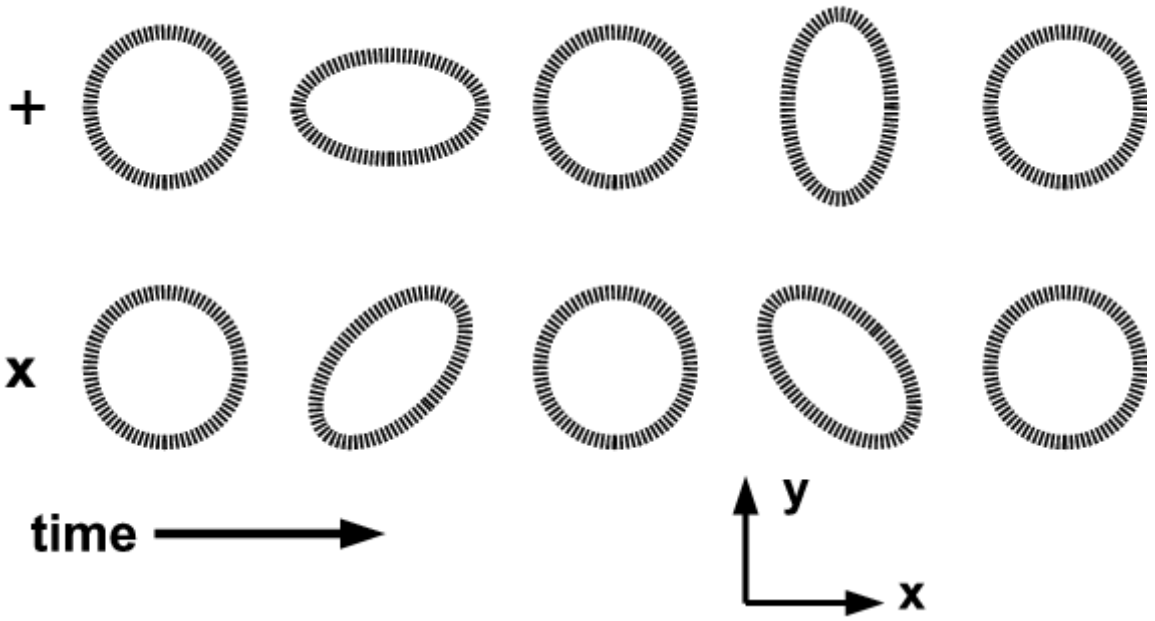

Figure 3: The effect of a plus-polarized (a) and cross-polarized (b) gravitational wave passing out of the page on a ring of test masses. Figure from Kalmus [2].

These two independent components of the Riemann tensor correspond to two physical polarization states for gravitational waves in general relativity. The outcome of this result is that the relative position of two test masses is measurable; a gravitational wave could be detected by watching the relative separation of two test masses over time.

Put differently, if a gravitational wave propagates in the $\hat{\Omega} = \hat{z}$ direction, the metric (within the TT gauge) takes on the form

$$ds^2 = -dt^2 + (1 + h_+)\, dx^2 + (1 - h_\times)\, dy^2 + dz^2. \tag{1.2.18}$$

If the gravitational wave is + polarized and two particles are located at positions $v^i = (v_1, 0, 0)$ and $w^i = (w_1, 0, 0)$ then at some instant of time the metric is

$$ds^2 = (1 + h_+(t - z))\,(v_1 - w_1)^2 \tag{1.2.19}$$

which means that the proper distance between the particles is

$$ds \approx (1 + h_+(t - z))\,(v_1 - w_1)\,. \tag{1.2.20}$$

The change in proper distance between the particles depends on the amplitude of the gravitational-wave. This concept forms the basis of gravitational wave detection efforts.

1.2.2 The production of gravitational waves

In the last section, the effect of gravitational waves on matter was considered in the far-zone, when the source of gravitational waves is distant. To gain a physical understanding of the sources of gravitational waves and connect far-field solutions of the Einstein equations to near-field solutions with sources present, solutions to the Einstein equations (1.0.1) are needed.

In general, it is sufficient to solve the linearized Einstein equations (1.1.23) for a specified source $T_{\alpha\beta}$. For an observer who is far away, the metric perturbation due to a slowly-moving source $T_{\alpha\beta}$ can be written as

$$\bar{h}^{\alpha\beta}(t,\vec{x}) = \frac{4G}{c^4}\int \frac{T^{\alpha\beta}(t-|\vec{x}-\vec{x}'|/c,\vec{x}')}{|\vec{x}-\vec{x}'|}d^3\vec{x}'. \tag{1.2.21}$$

Ultimately we are interested in the sources that generate the radiative gravitational field. To understand the dynamics of these sources, it is necessary to study both far and near-field solutions to Eq. (1.2.21).

In the far-field, the distance from the source to the field point r is much greater than the gravitational-wave wavelength (which is also much larger than the size of the source). In this case, the quantity $|\vec{x}-\vec{x}'| \simeq r$ appearing in Eq. (1.2.21) is approximately constant over the entire source and we can approximate $t-|\vec{x}-\vec{x}'|/c \approx t-r/c$ over the source. The metric perturbation can then be written as

$$\bar{h}^{\alpha\beta}(t,\vec{x}) = \frac{4G}{c^4 r}\int T^{\alpha\beta}(t-r/c,\vec{x}')d^3\vec{x}'. \tag{1.2.22}$$

This can be simplified further using the conservation law $\partial_\alpha T^{\alpha\beta}=0$. In this case, the spatial components of the trace-reversed metric perturbation are

$$\bar{h}^{ij}(t,\vec{x}) \simeq \frac{2G}{c^4 r}\frac{\partial^2}{\partial t^2}\int x'^i x'^j T^{00}(t-r/c,\vec{x}')d^3\vec{x}'. \tag{1.2.23}$$

The solutions to Eq. (1.2.23) are

$$\bar{h}^{ij}(t,\vec{x}) \simeq \frac{2G}{c^4 r}\ddot{I}^{ij}(t-r/c) \tag{1.2.24}$$

where we have defined the *quadrupole* tensor

$$I^{ij}(t) = \int x^i x^j T^{00}(t-r/c,\vec{x})d^3\vec{x}. \tag{1.2.25}$$

To obtain the solution in the TT-gauge, the transverse-projection operator $P_{ij} = \delta_{ij} - \hat{n}_i \hat{n}_j$ must be used. For radially traveling waves, where $\hat{n}^i = x^i/r$, the far-field solution is

$$h_{ij}^{TT}(t, \vec{x}) \simeq \frac{2G}{c^4 r} \ddot{I}_{ij}^{TT}(t - r/c) \tag{1.2.26}$$

and the quadrupole tensor is

$$I_{ij}^{TT} = P_{ik} I^{kl} P_{lj} - \frac{1}{2} P_{ij} P_{kl} I^{kl}. \tag{1.2.27}$$

Note that the quadrupole tensor is trace-free as it has been defined above.

In the near-field, the distance between the field point r and the source is much smaller than the gravitational-wave wavelength. Since the source is also moving slowly, we are in the Newtonian limit and the Newtonian potential can be used to understand the dynamics of the source. Ignoring internal stresses to the system and expanding $|\vec{x} - \vec{x}'|$ in powers of $1/r$, the Newtonian potential can be written as

$$\Phi(t, \vec{x}) = -G \left[\frac{M}{r} + \frac{D_i x^i}{r^3} + \frac{3}{2} \frac{\mathcal{I}_{ij} x^i x^j}{r^5} + \cdots \right] \tag{1.2.28}$$

where

$$\begin{aligned} M &= \int T^{00}(\vec{x}) d^3\vec{x} \\ D^i &= \int x^i T^{00}(\vec{x}) d^2\vec{x} \\ \mathcal{I}^{ij} &= \int \left(x^i x^j - \frac{1}{3} r^2 \delta^{ij} \right) T^{00}(\vec{x}) d^3\vec{x}. \end{aligned} \tag{1.2.29}$$

In the near-field, coordinates can be chosen so that the strain is

$$h_{ij}^{TT}(t, \vec{x}) \simeq \frac{2G}{c^4 r} \ddot{\mathcal{I}}_{ij}^{TT}(t - r/c) \tag{1.2.30}$$

which is identical to Eq. (1.2.30).

The significance of the near and far-field solutions is that accelerating masses produce gravitational waves. Furthermore, gravitational waves have no monopolar or dipolar components; only objects with a non-zero quadrupole moment (corresponding to the $\mathcal{I}$ term in Eq. (1.2.28)) will produce gravitational radiation.

1.3 Sources and resulting signals

In general, any system with a non-zero time-changing quadrupole moment ($\ddot{I}^{TT} \neq 0$) will produce gravitational waves. However, the effect of gravitational waves on matter is quite small, and modern gravitational wave detectors search for waves produced by the most massive and relativistic sources, which largely correspond to violent astrophysical events. In this section, we briefly touch on such sources and the types of gravitational waves they produce.

1.3.1 Continuous sources

Any astrophysical object that is rotating with axial asymmetry will emit gravitational waves. In the process of emitting gravitational waves, the object loses energy and slows down. The process of spinning down is quite slow, so the timescale over which gravitational radiation is emitted by the object is much longer than a typical observational timescale. The object is therefore called a continuous source of gravitational waves.

One of the most important such objects is an asymmetric or wobbling neutron star. The neutron star could have asymmetry due to precession, "mountains" or deformities in its crust, unstable fluid modes of oscillation, or other rotational instabilities [3, 4, 5, 6, 7, 8, 9]. If the neutron star is not axisymmetric, it will emit gravitational waves at a frequency that is twice its rotational frequency. Rotating axial asymmetric neutron stars are expected to produce gravitational waves with a frequency at the higher end of the detectable spectrum, $f \sim 1$ Hz – 1 kHz.

It is also possible for binary systems to emit continuously if their timescale of orbital decay is much longer than the observational timescale. More specifically, for a system to be considered as a continuous source of gravitational waves, its orbital frequency must remain relatively stable over the observational period. Because of this requirement, the frequencies of gravitational waves emitted by these systems are dependent upon the masses of the stars in the binary [7]. For example, the lowest frequency gravitational waves of $\sim 10^{-6}$ Hz emitted by such systems come from supermassive black hole binaries, while white dwarf binaries will produce gravitational waves in the $\sim 10^{-3}$ Hz frequency band.

Searches for continuous sources of gravitational waves exploit the fact that their signal is typically nearly monochromatic, and the signal can be well-modeled. One search that has been performed by the Initial Laser Interferometer Gravitational-wave Observatory (introduced in Chapter 5) involved observations of the Crab pulsar, a relatively young millisecond pulsar. Although no gravitational waves were found (see Abbott et al. 2008 [10]), upper limits indicate that less than 6% of the Crab's spin-down energy was gravitationally radiated. Future searches, with enhanced detectors and increased sensitivity, may be able to recover signals from the Crab pulsar.

1.3.2 Compact binary coalescence

When two compact objects such as neutron stars or black holes orbit each other, they lose energy due to the emission of gravitational waves. Over time, their orbit shrinks — while their frequency increases — until they merge with a characteristic chirp signal. The resulting black hole is highly deformed and reduces this deformity by emitting ringdown radiation. The final stages of this process are highly relativistic and the binary emits a vast amount of gravitational radiation in a short time.

The binaries that are most well-understood, both in terms of their expected waveform and expected event rate, are neutron star binaries (NS/NS). The existence of such binary systems has been confirmed through observations of pulsars in the galaxy; one example is the The Hulse-Taylor binary pulsar, discovered in 1974 [11]. Observations of the Hulse-Taylor binary (see Fig. 4) show that the system's orbital decay agrees strongly with general relativity's prediction of energy loss due to gravitational radiation (in fact, it won Hulse and Taylor the Nobel Prize in 1994). Unfortunately, the Hulse-Taylor binary will not coalesce for $\sim 300 \times 10^6$ years, but its existence (along with other more recently discovered binary pulsars) helps to constrain possible event rates for NS/NS binaries. Black hole/neutron star binaries (BH/NS) and black hole/black hole (BH/BH) binaries are another set of potential sources for gravitational wave emission.

As for the case of continuous sources, gravitational waves from coalescing binaries are typically very well-modeled. Post-Newtonian theory provides detail about the inspiral stage (when the orbit is decaying), and numerical relativity is able to describe the

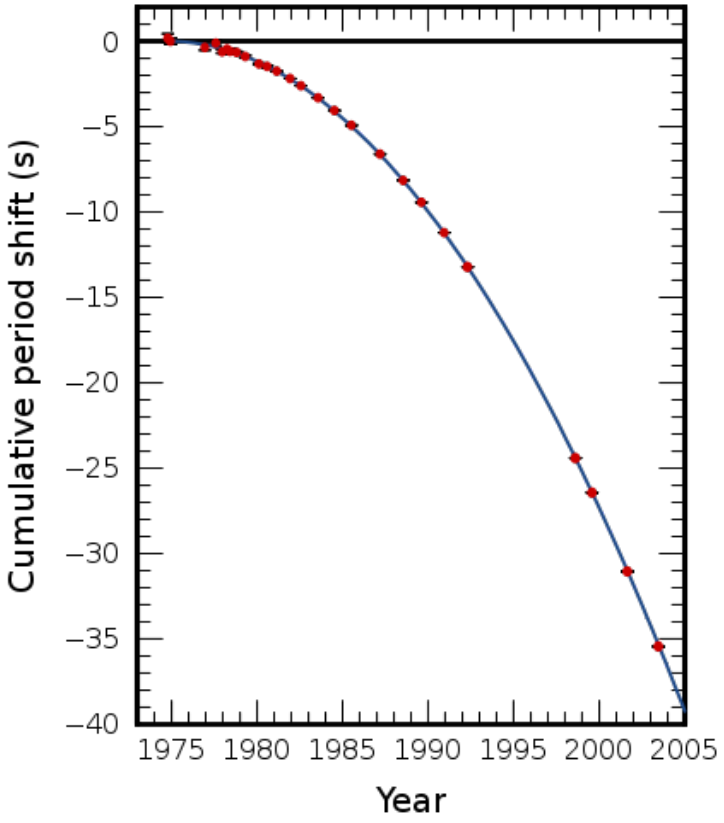

Figure 4: The measured orbital decay (red points) of the Hulse-Taylor binary pulsar is shown along with the prediction from general relativity (blue line). The data is in exceedingly good agreement with the prediction. Data in figure taken from [12].

late inspiral/merger stages (as the objects coalesce). This allows for matched filtering techniques to be used to search for gravitational waves from compact binary coalescence.

1.3.3 Sources of gravitational-wave bursts

The sources of transient gravitational waves, also known as "bursts" of gravitational waves, produce signals that are much shorter in duration than the typical observational timescale. One the most important sources of gravitational wave bursts is core collapse of massive stars.

For stars more than $\sim 8M_{\odot}$, core collapse signifies the end of the star's lifecycle. The outcome of collapse depends on the mass of the progenitor. Stars with $\lesssim 8M_{\odot}$ will form white dwarfs, supported by electron degeneracy pressure, at the end of their lives. If in a binary system, the dwarf could accrete enough matter from a companion to exceed the

Chandrasekhar limit, from which accretion-induced collapse ensues and a neutron star is born. Stars that are slightly more massive, $\gtrsim 8M_{\odot}$, will undergo stellar core collapse as electron degeneracy pressure can no longer counter gravitation. Depending on whether or not a supernova explosion follows, either a neutron star or black hole is born. Stars that are even more massive yet ($\sim 50M_{\odot}$) can collapse directly to black holes without producing a supernovae.

Although a spherically symmetric collapse would not result in gravitational radiation, the collapse of a rotating core causes it to flatten along the axis of rotation. The resulting aspherical collapse, as well as subsequent core oscillations, produces short bursts of gravitational radiation. Additionally, bar-mode instabilities, acoustic instabilities, and neutrino emission can also produce gravitational radiation.

A detailed gravitational wave signature for any one of these scenarios must be obtained through numerical models of stellar collapse. Such models do not tend to produce well-modeled waveforms, because the physical processes underlying collapse are quite complex and difficult to probe numerically. The rate of supernovae in the galaxy is estimated to be $1/25\,\mathrm{yr}^{-1}$ – $1/100\,\mathrm{yr}^{-1}$, potentially making gravitational-wave bursts rare.

1.3.4 Stochastic backgrounds

When the gravitational waves from a large number of independent, individually unresolvable sources overlap, a background of gravitational radiation is produced. If the many sources that contribute to this background are characterized by many different random times and frequencies, the central limit theorem suggests that the sum of the signals at any given time or frequency is a Gaussian random variable. The background is thus referred to as a *stochastic* background.

The sources of a stochastic background of gravitational waves range from cosmological (those exisiting in the early universe) to astrophysical (referring to recent astrophysical systems). A relic gravitational-wave background from inflation is expected to exist; other cosmological sources include cosmic strings and first-order phase transitions in the early universe. Astrophysical sources of a stochastic gravitational-wave background include neutron stars, white dwarf binaries and supermassive black hole binaries.

The stochastic background can be described in terms of its spectrum,

$$\Omega_{\rm gw}(f) = \frac{1}{\rho_c}\frac{d\rho_{\rm gw}}{d\log f}, \tag{1.3.1}$$

where $d\rho_{\rm gw}$ is the energy density of gravitational radiation contained the frequency range $(f, f + \Delta f)$ and ρ_c is the critical density needed to close the universe,

$$\rho_c = \frac{3c^2H_0^2}{8\pi G} \tag{1.3.2}$$

where H_0 is Hubble's constant. This spectrum will become important in Chapers 3 and 4 which involve searches for stochastic gravitational-wave backgrounds with pulsar timing arrays.

1.4 Gravitational-wave detectors

The first gravitational-wave detectors constructed were resonant mass (or bar) detectors that were designed to be oscillated by a passing gravitational wave. The famous Weber-bar experiment, conducted by Joseph Weber in the 1960s, involved 1.5 ton cylindrical aluminum bars that were sensitive to gravitational-waves with a relatively high frequency (1660 Hz). Although Weber's experiment did not lead to conclusive evidence of gravitational waves, it set forth some of the groundwork for modern gravitational-wave detectors.

Fig. 5 plots sensitivity curves for current and proposed gravitational-wave detectors, and illustrates the complementary nature of current gravitational-wave searches. The following subsections briefly describe some of the current and proposed gravitational-wave detectors.

1.4.1 Pulsar timing arrays

Current gravitational-wave detection efforts encompass a wide rage of potential gravitational-wave frequencies and sources. At the relatively low-frequency end of the (detectable) gravitational-wave spectrum ($f \sim 10^{-9}$ Hz – 10^{-7} Hz), pulsar timing arrays (PTAs) exploit the high-precision timing of millisecond pulsars to search for gravitational waves.

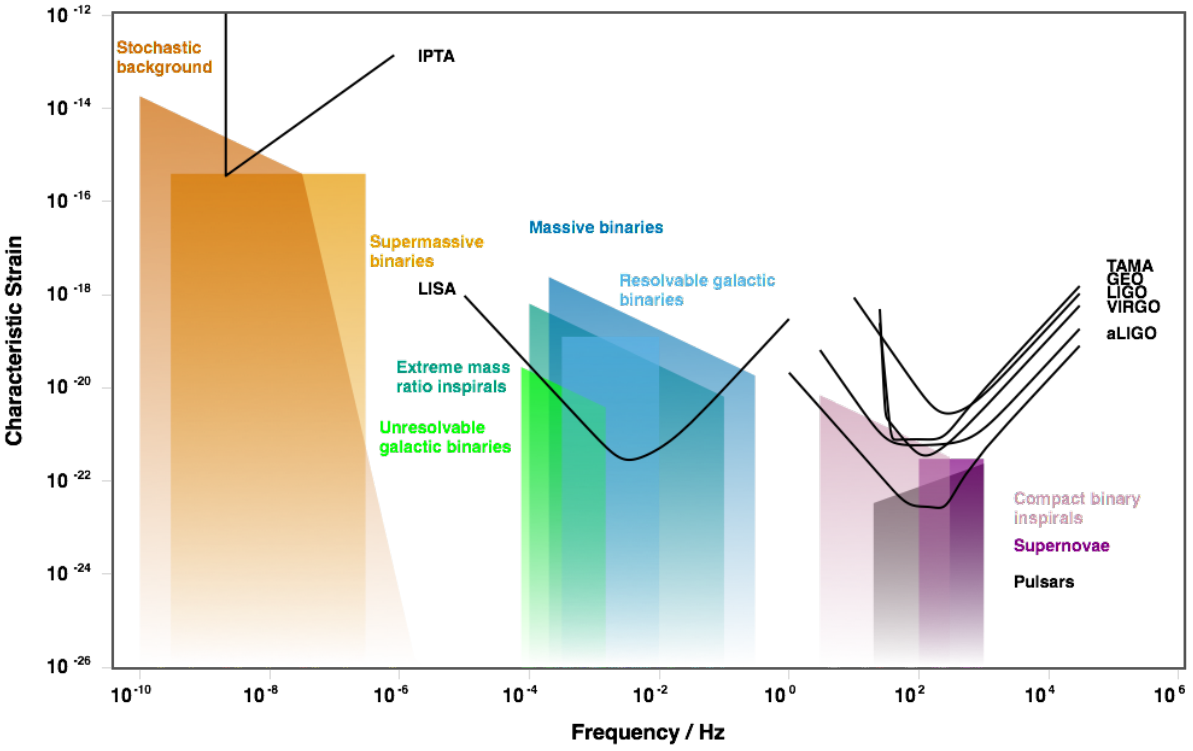

Figure 5: Gravitational-wave characteristic strain plotted as a function of gravitational-wave frequency for various detectors and sources. The anticipated source spectrums are indicated by the shaded regions below each curve. Note that this figure only illustrates a PTA sensitivity curve for an approximate combined IPTA data set and not current individual PTAs. Figure produced by Moore et al. [13].

A large-scale coordinated effort is underway across the globe to detect gravitational waves in the PTA frequency band. The PTAs contributing to this effort include the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) [14], the European Pulsar Timing Array (EPTA) [15], and the Parkes Pulsar Timing Array (PPTA) [16] in Australia. Together, these form the International Pulsar Timing Array (IPTA) [17].

Data for the current PTA experiments is collected at a variety of radio telescopes across the globe. The global nature of collaboration is particularly important for PTA experiments, because different telescopes probe different parts of the sky. Fig. 6 illustrates the global distribution of radio telescopes used by the IPTA.

The details surrounding PTAs and how they are used to detect gravitational waves are discussed in detail in Chapter 2.

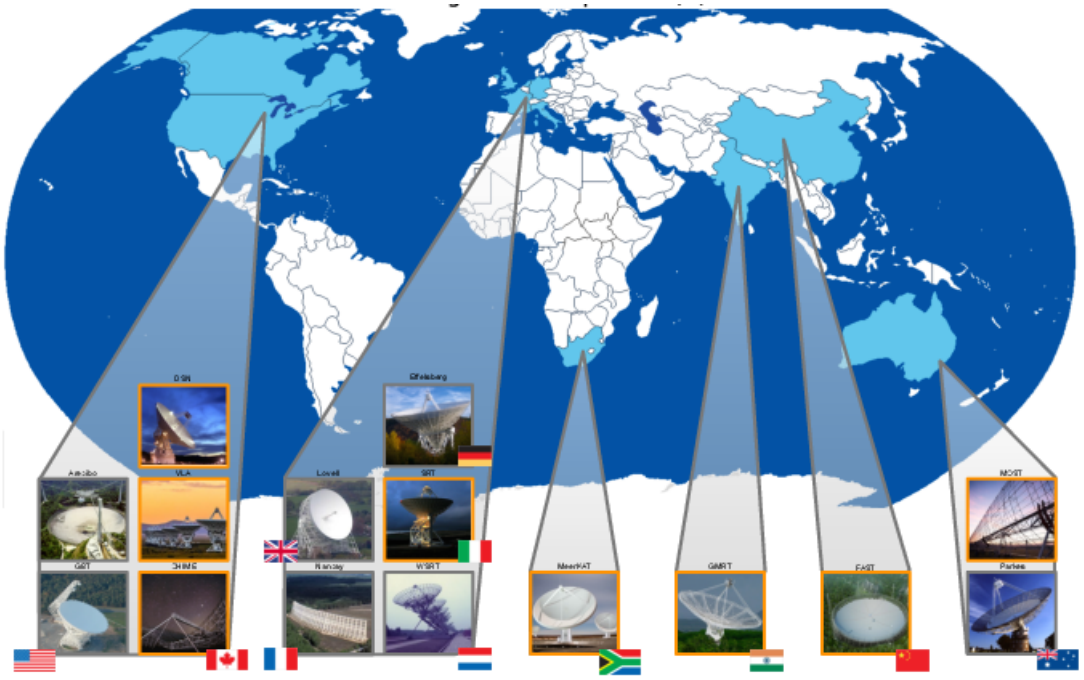

Figure 6: The data used by current PTA experiments is collected at radio telescopes across the globe, shown in this figure. Figure from the NANOGrav Collaboration.

1.4.2 Ground-based laser interferometers

At the relatively high-frequency end of the detectable gravitational wave spectrum ($f \sim$ 10 Hz – 1 kHz), kilometer-scale ground-based laser interferometers search for gravitational waves by measuring the differential changes in the positions of test masses at the ends of interferometer arms. Several ground-based laser interferometers have been constructed and taken data to search for gravitational waves.

The Laser Interferometer Gravitational-wave Observatory (LIGO) is comprised of two interferometers — one in Hanford, Washington (denoted as LHO) and one in Livingston, Louisiana (denoted as LLO) — with 4 km long arms each. A third detector, identical to the two in the United States, is planned for India.

The LIGO Scientific Collaboration (LSC) consists of the two LIGO detectors LHO and LLO, the GEO600 detector (a British-German interferometer with 600 m long arms) and the Virgo detector (a French-Italian interferometer with 3 km long arms). The LSC has completed several "science runs", or data taking periods. The S1-S5 science runs took place from 2002 to 2007. The configuration of the LIGO interferometers at that time is typically referred to as Initial LIGO.

Following the completion of the S5 science run, the two LIGO detectors and the Virgo detector were upgraded. The enhanced instruments, Enhanced LIGO and Virgo, were used to complete another science run (S6) which took place from July 2009-October 2010.

The two LIGO detectors, as well as the Virgo detector, are currently being upgraded for the Advanced LIGO and Advanced Virgo experiments. Advanced LIGO is anticipated to begin operating shortly after the publication of this dissertation.

The principles underlying gravitational-wave detection with ground-based laser interferometers will be presented in Chapter 5.

1.4.3 Space-based laser interferometers

Proposed space-based laser interferometers such as the Laser Interferometer Space Antenna (LISA) hold the potential to detect gravitational waves in the frequency band intermediate to the two discussed above, $f \sim 1$ mHz – 100 mHz. LISA and other space-based detectors will not be discussed in detail in this dissertation, but readers may consult the texts by Saulson [18] and Maggiore [8] for more details.

1.5 Dissertation Summary

The contents of this dissertation are based on my doctoral work, including two published papers. Chapter 2 introduces pulsar timing arrays and sets the foundation for the next two chapters. Chaper 3 presents a time-domain implementation of the optimal cross-correlation statistic for stochastic background searches in pulsar timing array data. Due to the irregular sampling typical of pulsar timing array data, as well as the use of a timing model to predict the times-of-arrival of radio pulses, time-domain methods are better suited for gravitational-wave data analysis of such data. This chapter presents a derivation of the optimal cross- correlation statistic starting from the likelihood function, a method to produce simulated stochastic background signals, and a rigorous derivation of the scaling laws for the signal-to-noise ratio of the cross-correlation statistic in two relevant pulsar timing array regimes: a weak signal limit where instrumental noise dominates over the gravitational-wave signal at all frequencies, and a second regime where the gravitational-wave signal dominates at the lowest frequencies.

In Chapter 4, stochastic background detection with pulsar timing arrays is studied in the context of alternative theories of gravity. The overlap reduction functions for pulsar timing arrays are determined and analyzed for *metric* theories of gravity (theories that obey the Einstein equivalence principle). It is shown that pulsar timing arrays have greater sensitivity to non-transverse gravitational waves than to transverse gravitational waves, a result that has interesting implications for developing gravitational-wave tests of general relativity. The physical origin of this enhanced sensitivity is discussed, and overlap reduction functions are calculated for some of the NANOGrav collaboration pulsars. It is shown that sensitivity to non-transverse polarization modes can increase dramatically for pulsar pairs with small angular separations.

Starting with Chapter 5, emphasis changes from gravitational-wave detection with pulsar timing arrays to detection with ground-based laser interferometers. Chapter 5 introduces the Laser Interferometer Gravitational-wave Observatory (LIGO) detectors and outlines relevant principles of gravitational-wave detection.

In Chapter 6, an excess power statistic is described to conduct searches for gravitational-wave bursts with LIGO. A data analysis pipeline constructed with the excess power statistic is applied to real and simulated data sets to determine its efficiency as a search tool. The excess power statistic turns out not only to be a tool for gravitational-waves searches, but also for detector characterization efforts which seek to mitigate problematic non-Gaussian noise transients in the LIGO detectors. Specific methods for performing these detector characterization tasks are discussed.

Chapter 7 summarizes the work outlined in this dissertation and describes planned future work.

Part I

Searches for stochastic gravitational wave backgrounds with pulsar timing arrays

Chapter 2

Pulsars as gravitational-wave detectors

> "*I switched on the high speed recorder and it came blip.... blip.... blip.... blip.... blip.... Clearly the same family, the same sort of stuff and that was great, that was really sweet. It finally scotched the little green men hypothesis cos it's highly unlikely there's two lots of little green men, opposite sides of the universe, both deciding to signal to a rather inconspicuous planet earth, at the same time, using a daft technique and a rather common place frequency. It has to be some new kind of star, not seen before, and that then cleared the way for us publishing, going public!*"
>
> — Jocelyn Bell Burnell, Astronomer

The pulsar is one of the most bizarre and interesting astrophysical objects known to exist in the universe. First discovered in 1967, pulsars — rapidly rotating neutron stars — are the most dense objects studied in astronomy that have not collapsed to form a black hole. Since their discovery, pulsars have been observed and characterized using pulsar timing experiments, and have given birth to a rich field of astronomy with applications ranging from the interstellar medium to plasma physics to gravitational-wave astronomy.

To explain how pulsars are used in gravitational-wave searches, it is useful to first describe some of their basic properties and illuminate the manner in which modern pulsar timing experiments are conducted. This will be done in the next section. The use of pulsars in gravitational-wave astronomy will be discussed in Sec. 2.2.

2.1 The pulsar as a stable clock

The first signals detected from a pulsar were discovered in 1967 by Jocelyn Bell-Burnell and Antony Hewish at Cambridge University [19]. The regularity of these pulses, combined with their distant astrophysical source, initially prompted the discoverers to dub the signal "LGM-1" (or "Little Green Men"). However, it was soon reasoned that these signals were originating from rotating neutron stars [20, 21, 22]. Decades of research on neutron stars has produced some understanding of pulsars and their properties. We will outline the basic scenario here, but readers are encouraged to consult the texts of Lorimer and Kramer [23] or Ostlie and Carroll [24] for a much more detailed discussion.

2.1.1 Pulsar formation and evolution

Neutron stars are born in supernovae, at the death of a massive main sequence star. The type of supernova that occurs depends on the mass of the progenitor. Core-collapse or type-II supernova occur for the most massive stars; such stars are generally believed to have masses of $M \gtrsim 8M_{\odot}$ – $10M_{\odot}$, although there is some evidence of stars as massive as $10M\odot$–$12M\odot$ avoiding supernova [25, 26, 27].

Stars that hold insufficient mass to produce a type-II supernova ($M \lesssim 8M\odot$) form white dwarfs at the end of their evolutionary cycle. White dwarfs are supported by electron degeneracy pressure, which is strong enough to counter the pull due to gravity on the star. However, there is a limit on how massive a white dwarf can be; if the dwarf exceeds a mass $\sim 1.4M\odot$ (also known as the *Chandrasekhar mass*), a type-I supernova can occur [28].

If the progenitor of the supernova is not massive enough to form a black hole, the remnant surviving the supernova is a rotating compact object composed of neutrons, with a canonical radius ~ 10 km, a mass $\gtrsim 1.4M_{\odot}$, and (surface) magnetic fields on the order $10^8 - 10^{15}$ Gauss. Neutron stars that are observed to emit electromagnetic radiation are referred to as pulsars.

The details surrounding the emission process of the pulsar are not completely understood and go beyond the scope of this dissertation (see [22, 29, 30, 31, 32, 23, 24] for more details), but in a basic model the pulsar contains a strong dipolar magnetic field

that is inclined relative to the rotational axis by some angle. As the pulsar rotates, its magnetic field varies in space and produces an electric field, which rips charged ions from the neutron star's polar regions [24]. These ions are accelerated relativistically and flow along the magnetic field lines, emitting a continuous spectrum of curvature radiation that forms a cone [32]. If the cone sweeps across the Earth's line of sight, the pulsar can be observed. Pulsars emit at a variety of frequencies [24], including radio-wavelengths that are observable at Earth via radio telecsopes.

As a result of energy lost due to a combination of electromagnetic radiation and particle emission [33, 34, 35], the pulsar gradually slows down in its rotation. Along with the pulsar's period P, the pulsar *spin-down* (or $\dot{P}$) can be determined with high levels of precision with modern pulsar timing experiments. The relationship between the pulsar's period and its derivative provide valuable information about the evolution of the pulsar. This information can be summarized with the $P - \dot{P}$ diagram, shown in Fig. 7. The $P - \dot{P}$ diagram shows two fairly distinct classes of pulsars: isolated pulsars, represented by single black dots, and pulsars in binary systems, represented by a dot in a circle. Interestingly, almost all of the MSPs are found in binary systems. Their separate location and grouping in the $P - \dot{P}$ diagram suggests that perhaps MSPs undergo a different evolutionary process than their slower isolated counterparts, which form the population of canonical (or typical) pulsars [23].

In fact, this is what is now believed for MSPs. The canonical pulsars are not only slower (with periods on the order of seconds) but also younger, with larger magnetic fields. Over the course of their life, as they emit electromagnetic radiation, these pulsars slowly evolve along lines of constant magnetic field to longer periods until they eventually land in the "pulsar graveyard" region in the $P - \dot{P}$ diagram. The MSPs are older and also have lower spin-down rates, which means that they are much more stable over time than pulsars in the canonical population.

Although MSP formation is still being studied in depth, the most favored current picture for their formation is based on "recycling" the pulsar and involves a binary system. If a supernova occurs for a star in a binary system, and the binary system is not disrupted (meaning the companion isn't ablated or ejected in the supernova), the resulting

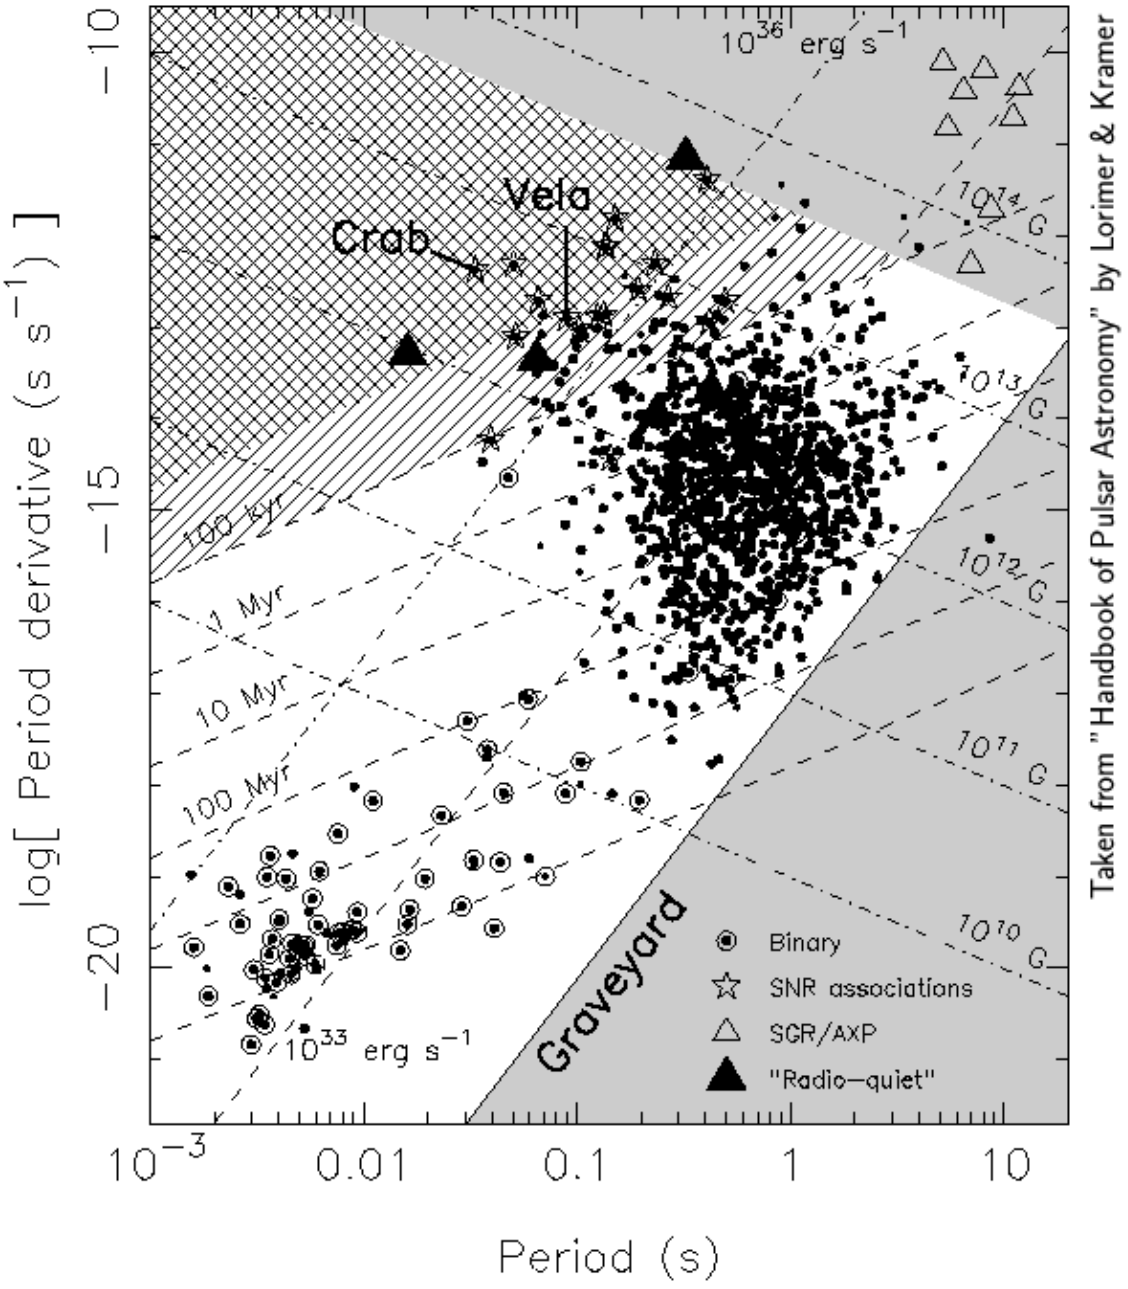

Figure 7: The $P - \dot{P}$ diagram illustrates the spin evolution for a population of pulsars. Single dots represent isolated pulsars; dots enclosed in circles correspond to pulsars in binary systems; and those represented by the star symbol indicate supernovae remnants. The three groups of dashed lines across the figure represent lines of constant characteristic age, magnetic field and spin-down luminosity. Figure 1.13 in Lorimer and Kramer [23].

neutron star may be observable as a pulsar. Over large timescales ($\sim 10^8$ yr) the pulsar loses energy and spins down until it no longer emits radiation. Because it is in a binary system, however, the now-quiet pulsar may be able to regain some of its lost spin by accreting matter from its companion. A massive enough companion will eventually overflow its Roche lobe, and accrete matter onto the pulsar. This accretion will also impart angular momentum that "spins-up" the pulsar. There is also evidence to suggest that the accretion of matter has the potential to "bury" the pulsar's intrinsic magnetic fields, which could explain why MSPs tend to have lower magnetic fields than their canonical counterparts [36, 37].

Because MSPs are so stable (relative to their younger canonical counterparts) they are the favored type of pulsar for gravitational-wave experiments. Before discussing the manner in which MSPs are used for gravitational-wave searches, basic elements of pulsar timing experiments will be reviewed.

2.1.2 Pulsar timing experiments

At the Earth, radio telescopes record the intensity of the pulsar's signal as a function of time. Individual pulsars are actually fairly weak radio sources, and require a process called *folding* to create a pulse profile; the folding process involves the coherent addition of hundreds or thousands of pulses together, and results in an integrated pulse profile for a pulsar. This process is illustrated in Fig. 8. Integrated pulse profiles are remarkably stable at a given frequency, despite the variation in individual pulses that comprise them.

The objective in pulsar timing experiments is to obtain the *timing residual* for the pulsar, which compares the pulse times-of-arrival (TOAs) to a theoretical timing model for the pulsar. The steps involved to produce timing residuals go beyond the scope of this document (see [39] for a more in depth analysis), but in rough detail the steps are as follows:

1. The pulsar's period is determined, either through an initial or previous observation.

2. Once the pulsar period is known, data with the correct phase are folded to obtain an integrated pulse profile.

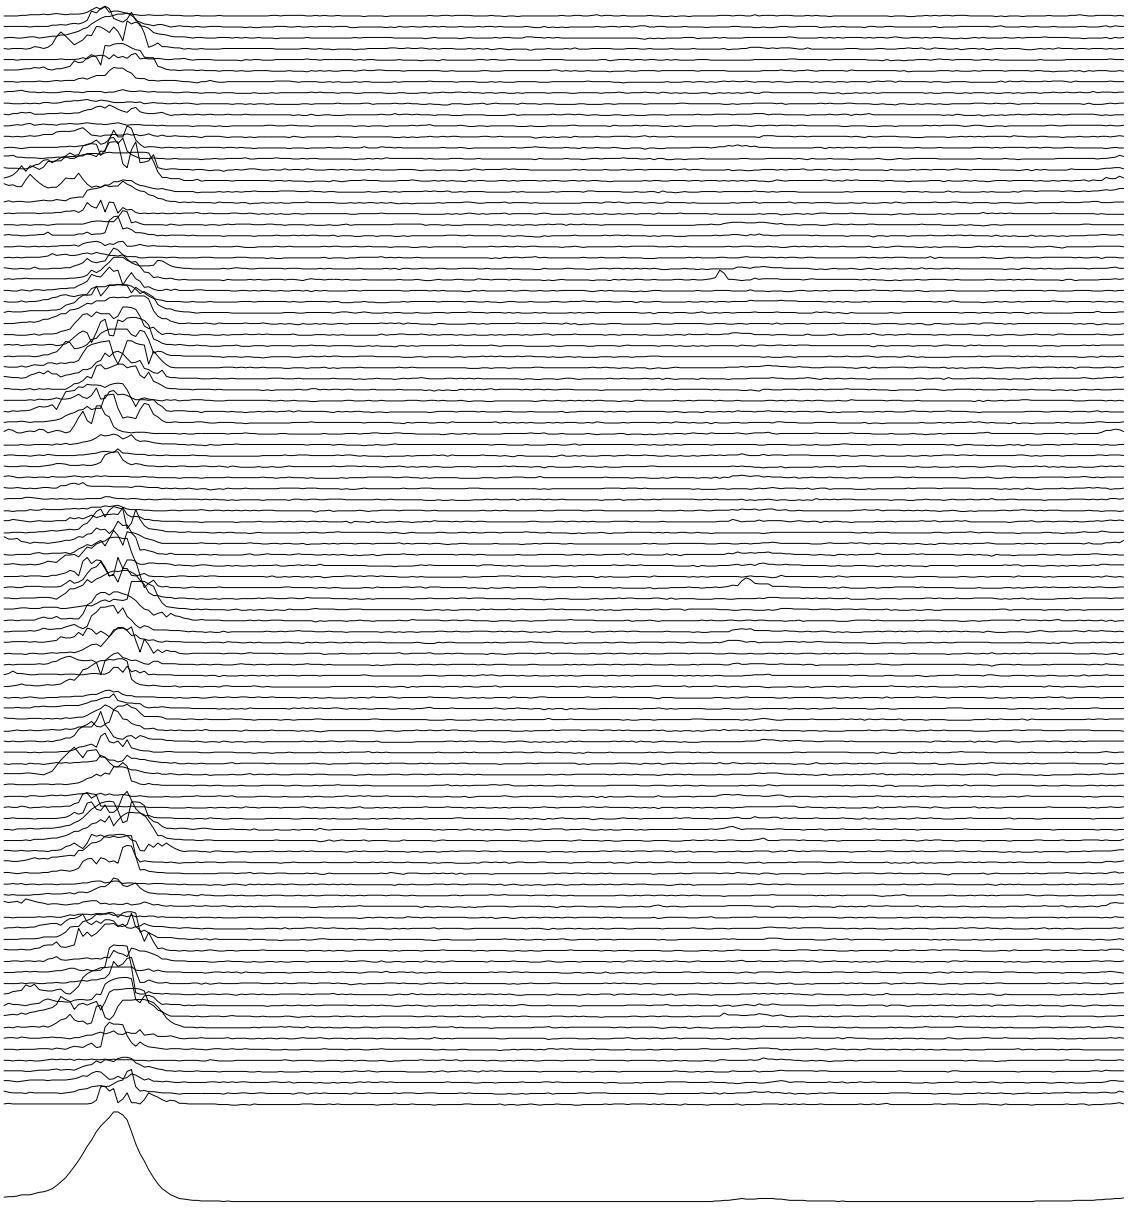

Figure 8: Individual pulses from the pulsar $B0950 + 08$ (253 ms) vary considerably in shape and intensity, but folding them results in a stable pulse profile that acts as a "fingerprint" of sorts for the pulsar. Figure 1 from Stairs [38].

3. The pulse profile is given a timestamp from the observatory atomic clock.

4. Additional timing parameters, such as the site-arrival-time and barycentric-arrival-time are determined and used so that measurements will correspond to an inertial reference frame.

5. The actual pulse TOAs are compared to arrival times produced by a timing model for the pulsar. The difference between the actual and expected TOAs defines the timing residuals.

This process is completed for many different pulsars, often with data from multiple radio telescopes, and leads to the term *pulsar timing array* (PTA) which describes this process for an entire array of pulsars across the sky. Current PTAs are accurate enough to predict pulse TOAs to within a nanosecond over several years [40, 41, 42].

A wealth of interesting science can be done using these timing residuals, but this dissertation involves detection efforts for gravitational waves. In the next section we will explain how the timing residuals are used to perform searches for gravitational waves.

2.2 Gravitational waves and the pulsar signal

The question of how a gravitational wave might effect the radio pulse traveling from a pulsar to the Earth was first probed in the late 1970s by Sazhin [43] and Detweiler [44]. In this section, we will determine explicitly the effect of the gravitational-wave signal on the signal from the pulsar, and relate that effect to the response of a PTA.

To begin, it is useful to make some choices surrounding the pulsar-Earth-GW geometry. Suppose that the pulsar is located in the direction of unit vector $\hat{p}$, pointing from Earth to the pulsar, and that a gravitational wave propagates in the direction $\hat{\Omega}$ (see Fig. 9). With this choice of coordinates, the polarization tensors defined in Sec. 1.2.1 of Chapter 1 can be used to describe the PTA response to gravitational waves.

This description begins with the *plane wave expansion*, which extends the plane wave solution of Chapter 1 (Eq. (1.2.11)). In terms of the geometry of Fig. 9, the plane wave

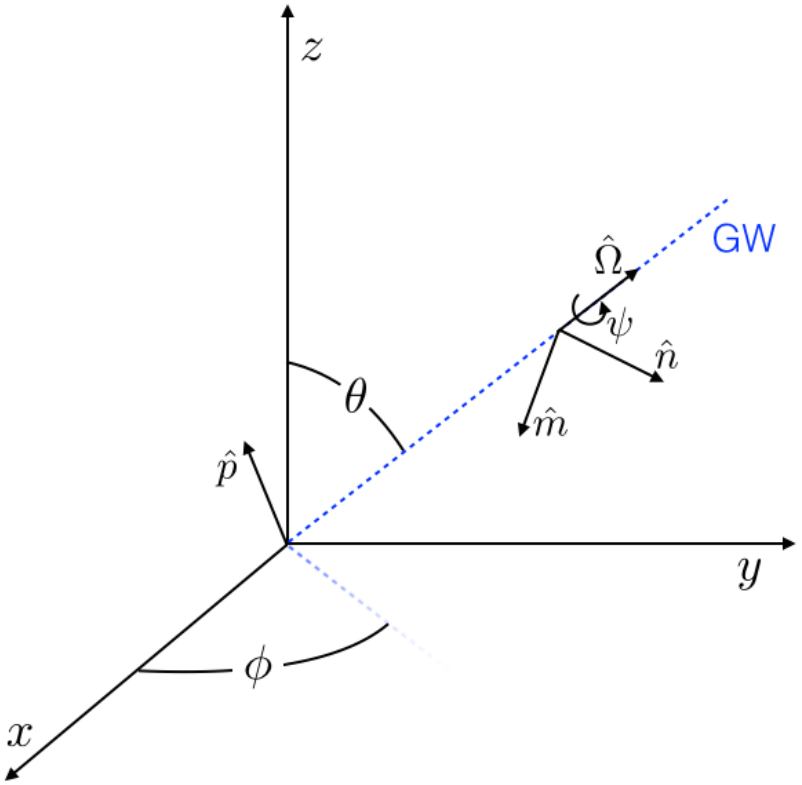

Figure 9: The pulsar-Earth-GW system, as visualized with the Earth at the origin. The gravitational wave propagates in the direction of the blue dashed line, and the vectors $\hat{\Omega}$, m and n defined in Sec. 1.2 are shown along with polar and azimuthal angles θ and ϕ. The unit vector $\hat{p}$ points from the Earth to the pulsar, and the angle ψ designates the polarization angle of the gravitational wave. For a stochastic gravitational wave background, all angles are averaged over many independent sources and ψ can be chosen to be zero.

expansion can be written as

$$h_{ij}(t,\vec{k}) = \sum_A \int_{-\infty}^{\infty} df \int_{S^2} d\Omega e^{2\pi i\, f(t-\hat{\Omega}\cdot\vec{k}/c)} h_A(f,\hat{\Omega}) \epsilon_{ij}^A(\hat{\Omega}) \tag{2.2.1}$$

where $A = \{+,\times\}$, f is the gravitational-wave frequency, $\vec{k}$ is the wave vector, $h_A(f,\hat{\Omega})$ is the gravitational-wave amplitude that depends on the wave's frequency, direction of propagation and polarization. Note that here we use the notation $\int_{S^2} d\Omega$ to represent the integral over the 2-sphere $\int d\theta \int d\phi$, and the hatless Ω appearing in the integrand should not be confused with the gravitational-wave direction $\hat{\Omega}$.

If the gravitational wave propagates in the z direction, then $\hat{\Omega} = \hat{z}$ and

$$\begin{aligned} h_{ij}(t, \hat{z}) &= \sum_A \int_{-\infty}^{\infty} df \int_{S^2} d\Omega e^{2\pi i f(t-z/c)} h_A(f, \hat{z}) \epsilon_{ij}^A(\hat{z}) \\ &= h_{ij}(t - z/c). \end{aligned} \tag{2.2.2}$$

The metric corresponding to this perturbation can be written as

$$\begin{aligned} g_{\alpha\beta} &= \eta_{\alpha\beta} + h_{\alpha\beta}(t - z/c) \\ &= \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 + h_+ & h_\times & 0 \\ 0 & h_\times & 1 - h_+ & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \end{aligned} \tag{2.2.3}$$

If the pulse from a pulsar is propagating in this spacetime background, the physical effect experienced by the pulse is a redshift:

$$z(t) = \frac{\nu_p - \nu_e}{\nu_e} \tag{2.2.4}$$

where ν_p is the pulse's emitted frequency and ν_e is the pulse's observed frequency on the Earth. The radio pulse from the pulsar must follow a null geodesic. This fact can be used to determine the gravitational-wave induced redshift in terms of the metric perturbation.

If a vector p^α is null in Minkowski space, in the perturbed spacetime the corresponding null vector k^α will be given by

$$k^\alpha = p^\alpha - \frac{1}{2}\eta^{\alpha\beta} h_{\beta\gamma} p^\gamma. \tag{2.2.5}$$

The flat-spacetime null vector that points from the pulsar to the Earth has components $p^\alpha = \nu(1, \hat{p})$, and these can be used to determine the *perturbed* null vector's components:

$$k^\alpha = \nu \begin{bmatrix} 1 \\ p_1(1 - h_+/2) - p_2 h_\times/2 \\ p_2(1 + h_+/2) - p_1 h_\times/2 \\ p_3 \end{bmatrix}. \tag{2.2.6}$$

The geodesic equation for the time component ($\alpha = 0$) of k^α is given by

$$\frac{dk^0}{d\lambda} = -\Gamma^0_{\alpha\beta} k^\alpha k^\beta, \tag{2.2.7}$$

where λ is an affine parameter for k^α and

$$\begin{aligned}\Gamma^0_{\alpha\beta} &= \frac{1}{2}g^{0\delta}\left(\partial_\alpha g_{\beta\delta} + \partial_\beta g_{\delta\alpha} - \partial_\delta g_{\alpha\beta}\right) = \frac{1}{2}\dot{g}_{\alpha\beta} \\ &= \frac{1}{2}\begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & \dot{h_+} & \dot{h_\times} & 0 \\ 0 & \dot{h_\times} & \dot{h_+} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}\end{aligned} \tag{2.2.8}$$

where the dot denotes differentiation with respect to time. The geodesic equation thus is

$$\begin{aligned}\frac{dk^0}{d\lambda} &= -\frac{1}{2}\dot{g}_{\alpha\beta}k^\alpha k^\beta \\ &= -\frac{1}{2}\left[\dot{g}_{11}(k^1)^2 - \dot{g}_{22}(k^2)^2\right] - \dot{g}_{12}k^1k^2 \\ &= -\frac{1}{2}\dot{h}_+\left[(k^1)^2 - (k^2)^2\right] + \dot{h}_\times k^1k^2\end{aligned} \tag{2.2.9}$$

The terms $(k^1)^2 - (k^2)^2$ and k^1k^2 evaluate to

$$(k^1)^2 - (k^2)^2 = \nu^2\left(p_1^2 - p_2^2\right) + \mathcal{O}(h) \tag{2.2.10}$$

and

$$k^1k^2 = \nu^2\left(p_1p_2\right) + \mathcal{O}(h) \tag{2.2.11}$$

respectively, and using the time component of $k^0 = \nu$, the geodesic equation becomes

$$-\frac{d\nu}{d\lambda} = \frac{1}{2}\dot{h}_+\nu^2\left(p_1^2 - p_2^2\right) + \dot{h}_\times\nu^2p_1p_2. \tag{2.2.12}$$

To turn this result into something integrable, it is necessary write the time derivatives on the right-hand side of Eq. (2.2.12) in terms of the affine parameter λ. Since $h_A = h_A(t - z/c)$,

$$\frac{dh_A}{d\lambda} = \frac{\partial h_A}{\partial t}\frac{dt}{d\lambda} + \frac{\partial h_A}{\partial z}\frac{dz}{d\lambda}. \tag{2.2.13}$$

From Eq. (2.2.6), it is clear that $dt/d\lambda = \nu$ and $dz/d\lambda = -\nu\,p_3$; furthermore $\partial h_A/\partial t = -\partial h_A/\partial z$. With these, Eq. (2.2.13) simplifies to

$$\frac{dh_A}{d\lambda} = \nu\left(1 + p_3\right)\dot{h}_A. \tag{2.2.14}$$

Replacing the time derivatives in Eq. (2.2.12) with this result, the geodesic equation becomes

$$-\frac{1}{\nu}\frac{d\nu}{d\lambda} = \frac{1}{2}\frac{p_1^2 - p_2^2}{1 + p_3}\frac{dh_+}{d\lambda} + \frac{p_1p_2}{1 + p_3}\frac{dh_\times}{d\lambda} \tag{2.2.15}$$

Integrating both sides of this expression,

$$
\begin{aligned}
-\int_{\lambda_p}^{\lambda_e} \frac{1}{\nu}\frac{d\nu}{d\lambda}d\lambda &= -\log\nu|_{\lambda_p}^{\lambda_e} = \log\left(\nu_p/\nu_e\right) \\
&= \frac{1}{2}\frac{p_1^2-p_2^2}{1+p_3}\int_{\lambda_p}^{\lambda_e}\frac{dh_+}{d\lambda}d\lambda + \frac{p_1p_2}{1+p_3}\int_{\lambda_p}^{\lambda_e}\frac{dh_\times}{d\lambda}d\lambda \\
&= \frac{1}{2}\frac{p_1^2-p_2^2}{1+p_3}\Delta h_+ + \frac{p_1p_2}{1+p_3}\Delta h_\times
\end{aligned}
\tag{2.2.16}
$$

where $\Delta h_A = h_A^e - h_A^p$ is the difference between the metric perturbation at the Earth when the pulse is received, and at the pulsar, when the pulse is emitted. Note that although the components of k^α are also functions of the affine parameter, terms of $\mathcal{O}(h)$ have been neglected. From the definition of redshift, $\log\left(\nu_p/\nu_e\right) = \log\left(1+z\right) \approx z$. This means that the gravitational-wave induced redshift of the radio pulse is

$$
z(t, \hat{\Omega}=\hat{z}) = \frac{1}{2}\frac{p_1^2-p_2^2}{1+p_3}\Delta h_+ + \frac{p_1p_2}{1+p_3}\Delta h_\times, \tag{2.2.17}
$$

when the gravitational wave propagates in the $\hat{\Omega}=\hat{z}$ direction.

In general, the gravitational wave can propagate in any direction and it is necessary to generalize this expression:

$$
z(t,\hat{\Omega}) = \frac{1}{2}\frac{p^i p^j}{1+\hat{\Omega}\cdot\hat{p}}\Delta h_{ij} \tag{2.2.18}
$$

where $\Delta h_{ij} = \Big[h_{ij}(t_e,\hat{\Omega}) - h_{ij}(t_p,\hat{\Omega})\Big]$ and the metric perturbation at each sky location takes the form

$$
h_{ij}(t,\hat{\Omega}) = \sum_A \int_{-\infty}^{\infty} e^{2\pi i f(t-\hat{\Omega}\cdot\vec{x})} \tag{2.2.19}
$$

These terms are typically referred to as the Earth-term and the pulsar-term, respectively; the Earth-term is correlated for all pulsars (regardless of their distances to the Earth), while the pulsar-term is not (the amplitude of this term is modulated by the direction of the pulsar with respect to the gravitational-wave source).

It is often useful to work with the gravitational-wave induced redshift in the frequency domain. By choosing a particular coordinate system where the Solar System barycenter is at the origin and the pulsar is some distance L away, we can write

$$
\begin{aligned}
t_p &= t_e - L = t - L, \\
\vec{x_e} &= 0, \\
\vec{x_p} &= L\hat{p}
\end{aligned}
\tag{2.2.20}
$$

where p denotes the position of the pulsar. The Fourier transform of Δh_{ij} is then

$$\Delta h_{ij} = \int_{-\infty}^{\infty} (1 - e^{-2\pi i f L(1+\hat{\Omega}\cdot\hat{p})}) \sum_A h_A(f,\hat{\Omega}) \epsilon_{ij}^A, \tag{2.2.21}$$

and Eq. (2.2.18) can be written in the Fourier domain as [45]

$$\tilde{z}(f,\hat{\Omega}) = \left(1 - e^{-2\pi i f L(1+\hat{\Omega}\cdot\hat{p}/c)}\right) \sum_A h_A(f,\hat{\Omega}) F^A(\hat{\Omega}) \tag{2.2.22}$$

where L is the distance to the pulsar from the Earth and we have defined *antenna patterns* that describe the detector response to the gravitational wave in terms of the pulsar-Earth-GW geometry:

$$F_A(\hat{\Omega}) = \frac{p^i \, \epsilon_{ij}^A(\hat{\Omega}) \, p^j}{2\left(1 + \hat{\Omega}\cdot\hat{p}\right)}. \tag{2.2.23}$$

As discussed in Sec. 2.1.2, the physical quantity determined at radio telescopes is the timing residual, or the difference between the actual and expected pulse TOAs at the Earth. The timing residual can be thought of as the sum of Doppler shifts due to the gravitational wave, and is thus defined as

$$r(t) = \int_0^t dt' \, z(t'). \tag{2.2.24}$$

It is this simple expression that makes pulsar timing experiments a viable tool for gravitational-wave astronomy since it is what relates the physically observed quantity to that which is affected by the gravitational wave.

Since the physical observable is the timing residual and not the redshift, it is useful to recast Eq. (2.2.22) in terms of the timing residual. Using the plane wave expansion (Eq. (2.2.1)), the frequency-domain expression for the timing residuals produced by a gravitational wave is

$$\tilde{r}(f,\hat{\Omega}) = \frac{1}{2\pi i f}\left(1 - e^{-2\pi i f L(1+\hat{\Omega}\cdot\hat{p}/c)}\right) \sum_A h_A(f,\hat{\Omega}) F^A(\hat{\Omega}), \tag{2.2.25}$$

2.3 Pulsar timing arrays and gravitational waves

The PTA gravitational-wave detection efforts discussed for the remainder of this dissertation focus on a particular type of gravitational-wave signal: a stochastic background, first discussed in Sec. 1.3 (and discussed in additional detail in Allen and Romano, Ref. [46]).

For pulsar timing arrays, the most likely source of a stochastic background of gravitational waves are supermassive black hole binaries (SMBBHs) that form following the merger of massive galaxies [47, 48, 49]. The superposition of gravitational waves from all SMBBH mergers forms a stochastic background of gravitational waves [50, 47, 51, 52, 53, 48, 54, 55]. Individual periodic signals [49, 56, 57, 58, 59] and bursts [60, 61] can also be produced by SMBBH systems. In addition, cosmic strings [62, 63, 64, 65], first order phase transitions in the early universe [66], and relic gravitational waves from inflation [67, 68] are potential sources of gravitational waves in the nanohertz band.

The backgrounds produced by these sources may be characterized by specifying how the gravitational-wave energy is distributed in frequency:

$$\Omega_{\rm gw}(f) = \frac{1}{\rho_c}\frac{d\rho_{\rm gw}}{d\log f}, \tag{2.3.1}$$

where $d\rho_{\rm gw}$ is the energy density of gravitational radiation contained the frequency range $(f, f + \Delta f)$ and ρ_c is the critical density needed to close the universe,

$$\rho_c = \frac{3c^2H_0^2}{8\pi G} \tag{2.3.2}$$

and H_0 is Hubble's constant,

$$H_0 = 100\,h\frac{\rm km}{\rm sMpc}, \tag{2.3.3}$$

which depends on a dimensionless constant h that accounts for the variation of H_0 appearing in the literature.

The gravitational-wave stochastic background is completely specified by $\Omega_{\rm gw}(f)$ if some assumptions are made:

i That the stochastic background is isotropic (although this condition has been probed in recent literature; see Mingarelli et al. [69] and Taylor et al. [70] for specific details);

ii That the stochastic background is unpolarized;

iii That the stochastic background is stationary;

iv That the stochastic background is Gaussian (this assumption is justified due to the central limit theorem, so long as there are a large number of stochastic sources distributed across the sky).

Given these assumptions and the plane-wave expansion, Eq. 2.2.1, the expectation value of the Fourier amplitudes $h_A(f,\hat{\Omega})$ can be expressed as

$$\langle h_A^*(f,\hat{\Omega})h_{A'}(f,\hat{\Omega}')\rangle = \delta^2(\hat{\Omega},\hat{\Omega}')\delta_{AA'}\delta(f-f')H(f) \tag{2.3.4}$$

where $\delta^2(\hat{\Omega},\hat{\Omega}')$ is the Dirac delta function on the two-sphere, $H(f)$ is some real, positive-valued function that satisfies $H(f) = H(-f)$ and $\langle\,,\rangle$ denote the expectation value. The function $H(f)$ can be related to $\Omega_{\rm gw}(f)$ by differentiating the plane-wave expansion (Eq. 2.2.1) with respect to time and determining the energy density of gravitational waves,

$$\rho_{\rm gw} = \frac{c^2}{32\pi G}\langle \dot{h}_{\alpha\beta}(t,\vec{k})\dot{h}^{\alpha\beta}(t,\vec{k})\rangle. \tag{2.3.5}$$

The frequency-dependent quantity that remains on the right-hand side of this expression determines $H(f)$ explicitly, and the expectation value of the Fourier amplitudes $h_A(f,\hat{\Omega})$ can then be fully expressed as

$$\langle h_A^*(f,\hat{\Omega})h_{A'}(f,\hat{\Omega}')\rangle = \frac{3H_0^2}{32\pi^3}\delta^2(\hat{\Omega},\hat{\Omega}')\delta_{AA'}\delta(f-f')|f|^{-3}\hat{\Omega}_{\rm gw}(|f|). \tag{2.3.6}$$

In the context of gravitational-wave detection, it is useful to relate the expectation value in Eq. 2.3.6 to the physical observable of PTAs, the timing residual. The stochastic background produces changes in the timing residuals of individual pulsars that are *correlated* between different pulsars, and Hellings and Downs first showed in 1983 [71] that the cross-correlation of the timing residuals from two pulsars I and J depends only on the angular separation ζ_{IJ} of the two pulsars:

$$\langle \tilde{r}_I^*(f)\tilde{r}_J(f')\rangle = \frac{H_0^2}{16\pi^4}\delta(f-f')|f|^{-5}\Omega_{\rm gw}(|f|)\chi_{IJ}(\zeta_{IJ}), \tag{2.3.7}$$

where $\chi_{IJ}(\zeta_{IJ})$ is the *Hellings-Downs curve* (sometimes called the Hellings-Downs coefficient) defined by

$$\chi_{IJ}(\zeta_{IJ}) = \frac{3}{2}\left[\frac{1}{3} + \frac{1-\cos\zeta_{IJ}}{2}\left[\ln\left(\frac{1-\cos\zeta_{IJ}}{2}\right) - \frac{1}{6}\right]\right] + \frac{1}{2}\delta_{IJ}. \tag{2.3.8}$$

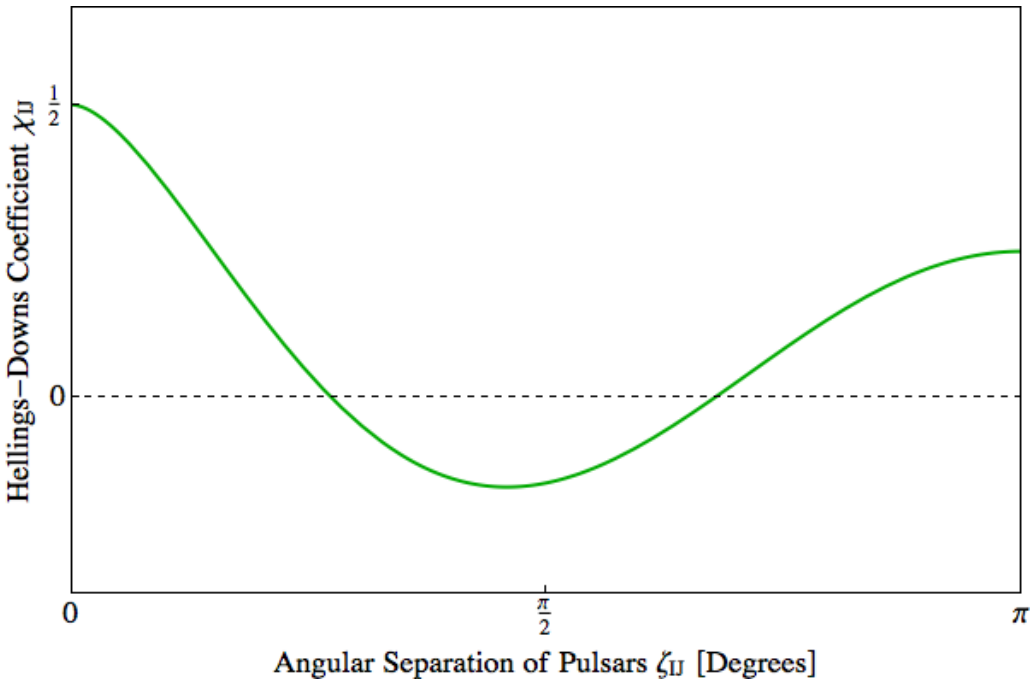

Figure 10: The Hellings and Downs curve, above, shows that the expected correlation χ_{IJ} for a pair of pulsars I and J depends only on their angular separation ζ_{IJ}.

This is illustrated in Fig. 10, which plots the Hellings-Downs curve as a function of pulsar angular separation. The pulsar term in Eq. (2.2.25), proportional to $e^{-2\pi i f L(1+\hat{\Omega}\cdot\hat{p})}$, contributes to the expectation value in Eq. (2.3.7) only in the case of the same pulsar (i.e., when I = J), and averages to zero for different pulsars.

In many cases it is useful to refer not to $\Omega_{\rm gw}(f)$ but instead to the dimensionless gravitational wave amplitude $A_{\rm gw}$ (at reference frequency $f_{\rm 1yr} = {\rm yr}^{-1}$) which appears in the expression for the characteristic strain

$$h_c(f) = A_{\rm gw}\left(\frac{f}{f_{\rm 1yr}}\right)^{\alpha}. \tag{2.3.9}$$

The spectral index α depends on the astrophysical source of the background. For example, a stochastic background produced by supermassive black hole binary systems has $\alpha = -2/3$ [47, 48]. The amplitude $A_{\rm gw}$ is related to the strain spectral density $S_h(f)$ of the gravitational-wave background via:

$$S_h(f) = \frac{h_c^2(f)}{f}. \tag{2.3.10}$$

For one-sided power spectra, $S_h(f)$ and $A_{\rm gw}$ are related to $\Omega_{\rm gw}(f)$ by

$$S_h(f) = \frac{3H_0^2}{2\pi^2}\frac{\Omega_{\rm gw}(f)}{f^3}, \tag{2.3.11}$$

$$\Omega_{\rm gw}(f) = \frac{2\pi^2}{3H_0^2}A_{\rm gw}^2 f^2 \left(\frac{f}{f_{\rm 1yr}}\right)^{2\alpha}. \tag{2.3.12}$$

These expressions will be useful in the next chapter, which introduces the optimal-cross correlation statistic for PTAs in the time-domain.

Chapter 3

Optimal Strategies for Stochastic GW Detection with PTAs

> "*Life is not easy for any of us. But what of that? We must have perseverance and above all confidence in ourselves. We must believe that we are gifted for something, and that this thing, at whatever cost, must be attained.*"
>
> — Marie Curie, Physicist, Two-time Nobel Laureate

A number of data analysis techniques have been developed and implemented to search for isotropic stochastic backgrounds of gravitational waves in PTA data [44, 72, 73, 74, 75, 45, 76, 77, 78, 40, 79, 41, 58, 80, 81, 82]. More recently, these techniques have been generalized to searches for anisotropic backgrounds [69, 70, 83, 84]. Additionally, a range of data analysis methods have been developed to search for individual periodic sources that stand out over the stochastic background [49, 56, 57, 85, 59, 86, 87, 88, 89, 90, 91, 92, 93], bursts [94, 95, 96, 97, 98], and signals of unknown form [99].

In this chapter, we develop a practical time-domain implementation of the optimal cross-correlation statistic [45] that can be used to search for isotropic stochastic backgrounds. Readers may wish to consult Chapter 2 to review the effect of a gravitational wave on the pulsar-Earth-GW system, and the expected cross-correlations in the times-of-arrival of pulses from different pulsars arising from a stochastic background of gravitational waves.

In Section 3.1, we develop the formalism needed to implement the search for a stochastic background, including the timing model, and derive the optimal cross-correlation

statistic from the likelihood ratio. In Section 3.2, we develop a procedure for injecting simulated stochastic background signals into PTA data, and in Section 3.3, we describe the scaling laws that govern the expected signal-to-noise ratio of the cross-correlation statistic. We conclude in Section 3.4 with a discussion of the practicality of implementing the statistics introduced in this chapter for gravitational-wave searches. We work in units where $c = G = 1$.

3.1 The optimal cross-correlation statistic

3.1.1 Timing Model

In pulsar timing experiments the quantities that are directly measured are the times-of-arrival (TOAs) of radio pulses emitted from pulsars. These TOAs contain many terms of known functional form, including intrinsic pulsar parameters (pulsar period, spin-down, etc.), along with stochastic processes such as radiometer noise, pulse phase jitter, and possibly red noise either from interstellar medium (ISM) effects, intrinsic pulsar noise, and, potentially, a gravitational wave background.

Suppose that the TOAs for a pulsar are given by

$$\mathbf{t}^{\rm obs} = \mathbf{t}^{\rm det}(\boldsymbol{\xi}_{\rm true}) + \mathbf{n}, \tag{3.1.1}$$

where $\mathbf{t}^{\rm obs}$ are the $N_{\rm TOA}$ observed TOAs, $\mathbf{t}^{\rm det}$ are the deterministic modeled TOAs parameterized by $N_{\rm par}$ timing model parameters $\boldsymbol{\xi}_{\rm true}$, and $\mathbf{n}$ is the noise time series in the measurement which is assumed to be Gaussian with covariance matrix given by

$$\mathbf{N} = \langle \mathbf{n}\mathbf{n}^T \rangle = \mathbf{N}^{\rm white} + \mathbf{N}^{\rm red} \tag{3.1.2}$$

where the $N_{\rm TOA} \times N_{\rm TOA}$ matrtices $\mathbf{N}^{\rm white}$ and $\mathbf{N}^{\rm red}$ are the contributions to the covariance matrix from the white and red noise processes, respectively. We will discuss the exact form of this covariance matrix in the next section. Assuming that estimates of the true timing model parameters $\boldsymbol{\xi}_{\rm est}$ exist (either from information gained when discovering the pulsar or from past timing observations), we can form the pre-fit timing residuals as

$$\delta\mathbf{t}^{\rm pre} = \mathbf{t}^{\rm obs} - \mathbf{t}^{\rm det}(\boldsymbol{\xi}_{\rm est}) = \mathbf{t}^{\rm det}(\boldsymbol{\xi}_{\rm true}) + \mathbf{n} - \mathbf{t}^{\rm det}(\boldsymbol{\xi}_{\rm est}). \tag{3.1.3}$$

As mentioned above, we will assume that the initial estimates for our timing model parameters are correct to some linear offset $\boldsymbol{\xi}_{\rm est} = \boldsymbol{\xi}_{\rm true} + \delta\boldsymbol{\xi}$, for which the pre-fit residuals become

$$\delta\mathbf{t}^{\rm pre} = \mathbf{t}^{\rm det}(\boldsymbol{\xi}_{\rm true}) - \mathbf{t}^{\rm det}(\boldsymbol{\xi}_{\rm true} + \delta\boldsymbol{\xi}) + \mathbf{n}. \tag{3.1.4}$$

Expanding this solution around the true timing model parameters, we obtain

$$\begin{aligned} \delta\mathbf{t}^{\rm pre} &= -\left.\frac{\partial \mathbf{t}^{\rm det}}{\partial \boldsymbol{\xi}}\right|_{\boldsymbol{\xi}=\boldsymbol{\xi}_{\rm true}} \delta\boldsymbol{\xi} + \mathbf{n} + \mathcal{O}(\delta\boldsymbol{\xi}^2) \\ &\approx -\left.\frac{\partial \mathbf{t}^{\rm det}}{\partial \boldsymbol{\xi}}\right|_{\boldsymbol{\xi}=\boldsymbol{\xi}_{\rm true}} \delta\boldsymbol{\xi} + \mathbf{n} \\ &= \mathbf{M}\delta\boldsymbol{\xi} + \mathbf{n}, \end{aligned} \tag{3.1.5}$$

where $\mathbf{M}$ is an $N_{\rm TOA} \times N_{\rm par}$ matrix, commonly referred to as the *design matrix* [100, 101]. Here we have assumed that our initial estimate of the model parameters is sufficiently close to the true values so that we can approximate this as a linear system of equations in $\delta\boldsymbol{\xi}$. It is customary in standard pulsar timing analysis to obtain the best fit $\delta\boldsymbol{\xi}$ values through a generalized least-squares minimization of the pre-fit residuals. The function that we seek to minimize is (see [39])

$$\chi^2 = \frac{1}{2}(\delta\mathbf{t}^{\rm pre} - \mathbf{M}\delta\boldsymbol{\xi})^T\mathbf{N}^{-1}(\delta\mathbf{t}^{\rm pre} - \mathbf{M}\delta\boldsymbol{\xi}). \tag{3.1.6}$$

Minimizing this function with respect to the parameter offsets $\delta\boldsymbol{\xi}$ results in

$$\delta\boldsymbol{\xi}_{\rm best} = -\left(\mathbf{M}^T\mathbf{N}^{-1}\mathbf{M}\right)^{-1}\mathbf{M}^T\mathbf{N}^{-\mathbf{1}}\delta\mathbf{t}^{\rm pre}. \tag{3.1.7}$$

The post-fit residuals are then given by

$$\delta\mathbf{t}^{\rm post} \equiv \delta\mathbf{t}^{\rm pre} - \mathbf{M}\delta\boldsymbol{\xi}_{\rm best} = \mathbf{R}\delta\mathbf{t}^{\rm pre}, \tag{3.1.8}$$

where

$$\mathbf{R} = \mathbb{I} - \mathbf{M}\left(\mathbf{M}^T\mathbf{N}^{-1}\mathbf{M}\right)^{-1}\mathbf{M}^T\mathbf{N}^{-1} \tag{3.1.9}$$

is a an $N_{\rm TOA} \times N_{\rm TOA}$ oblique projection matrix that transforms pre-fit to post-fit residuals, and $\mathbb{I}$ is the identity matrix. All of the information about any noise source or stochastic gravitational-wave background is encoded in $\mathbf{N}$. However, in most cases we have no a priori knowledge of this covariance matrix and therefore assume that it is given by

$\mathbf{W} = \mathrm{diag}(\{\sigma_i^2\})$, where σ_i is the uncertainty of the ith TOA. Previous work [102] has used an iterative method to estimate the covariance matrix of the residuals and apply a generalized least squares fit. For this work we will only work with residuals that have been created using a weighted least squares fit. It should be noted that in standard pulsar timing packages such as TEMPO2 [39] this process must be iterated. In other words, the pre-fit residuals are formed with an initial guess of the parameters, and the chi-squared is then minimized to produce best estimates of the parameters. This may not be a good fit, however, as we have assumed that the pre-fit residuals are linear in the parameter offsets. Consequently, we form new parameter estimates from the best fit parameter offsets and iterate until the fit converges, with the reduced chi-squared serving as the goodness-of-fit parameter. For this reason, we must ensure that our timing model fit has converged prior to any gravitational-wave analysis.

3.1.2 Derivation of the optimal statistic

Likelihood function for a PTA

Much of the discussion in this section follows closely that of [82], with additional details included here. We begin by assuming that our PTA consists of M pulsars, each with some intrinsic noise $\mathbf{n}_I(t)$. Henceforth uppercase latin indices will label a pulsar and lowercase latin indices will label a particular TOA. Under the assumption that all intrinsic pulsar noise is Gaussian, we can write the full likelihood function for the PTA as

$$p(\mathbf{n}|\vec{\theta}) = \frac{1}{\sqrt{\det(2\pi\mathbf{\Sigma}_n)}} \exp\left(-\frac{1}{2}\mathbf{n}^T\mathbf{\Sigma}_n^{-1}\mathbf{n}\right), \tag{3.1.10}$$

where now we are using the full PTA noise time series that is just a concatenated length MN_{TOA} column vector

$$\mathbf{n} = \begin{bmatrix} \mathbf{n}_1 \\ \mathbf{n}_2 \\ \vdots \\ \mathbf{n}_M \end{bmatrix}, \tag{3.1.11}$$

$\mathbf{\Sigma}_n$ is the $MN_{\rm TOA} \times MN_{\rm TOA}$ covariance matrix and $\vec{\theta}$ is a set of parameters that characterize the noise. The covariance matrix for the noise is the block matrix

$$\mathbf{\Sigma}_n = \begin{bmatrix} \mathbf{N}_1 & \mathbf{X}_{12} & \dots & \mathbf{X}_{1M} \\ \mathbf{X}_{21} & \mathbf{N}_2 & \dots & \mathbf{X}_{2M} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{X}_{M1} & \mathbf{X}_{M2} & \dots & \mathbf{N}_M \end{bmatrix}, \tag{3.1.12}$$

where

$$\mathbf{N}_I = \langle \mathbf{n}_I \mathbf{n}_I^T \rangle, \tag{3.1.13}$$

$$\mathbf{X}_{IJ} = \langle \mathbf{n}_I \mathbf{n}_J^T \rangle \big|_{I \neq J}, \tag{3.1.14}$$

are the auto-covariance and cross-covariance matrices, respectively, for each set of noise vectors.

In general the autocorrelation matrices are defined via the Wiener–Khinchin theorem as

$$\begin{aligned} \mathbf{N}_I = \langle \mathbf{n}_I \mathbf{n}_I^T \rangle_{ij} = \int_0^\infty df e^{2\pi i f \tau_{ij}} \mathcal{P}_I(f) \\ + \mathcal{F}_I \mathbf{W}_I + \mathcal{Q}_I^2 \mathbb{I} \end{aligned} \tag{3.1.15}$$

where $\tau_{ij} = |t_i - t_j|$, $\mathcal{F}_I$ and $\mathcal{Q}_I$ are white noise parameters for pulsar I (usually denoted as EFAC and EQUAD, respectively), $\mathbb{I}$ is the identity matrix, and $\mathcal{P}_I(f)$ is a red noise power spectrum

$$\mathcal{P}_I(f) = \mathcal{P}_I^{\rm int}(f) + \mathcal{P}_g(f) \tag{3.1.16}$$

where

$$\mathcal{P}_I^{\rm int}(f) = \frac{A_I^2}{12\pi^2} \left(\frac{f}{f_{\rm 1yr}} \right)^{2\alpha_I} f^{-3} \tag{3.1.17}$$

is the intrinsic red noise in the pulsar parameterized by amplitude A_I and spectral index α_I, and

$$\mathcal{P}_g(f) = \frac{A_{\rm gw}^2}{12\pi^2} \left(\frac{f}{f_{\rm 1yr}} \right)^{2\alpha} f^{-3} \tag{3.1.18}$$

is the gravitational-wave background spectrum parameterized by the strain amplitude $A_{\rm gw}$ and spectral index α. In other words, the auto-covariance matrix of the noise in pulsar I consists of intrinsic white noise parameterized by $\{\mathcal{F}_I, \mathcal{Q}_I\}$ and red noise parameterized

by $\{A_I, \alpha_I, A_{\rm gw}, \gamma\}$. Notice that the gravitational-wave parameters do not have a pulsar label because they are common to all pulsars.

Similarly, the cross-covariance matrices are given by

$$X_{IJ} = \langle \mathbf{n}_I \mathbf{n}_J^T \rangle_{ij} = \chi_{IJ} \int_0^\infty df\, e^{2\pi i f \tau_{ij}} \mathcal{P}_g(f) \tag{3.1.19}$$

where χ_{IJ} are the Hellings and Downs coefficients for pulsar pair I, J defined in Eq. (2.3.8).

We now write the likelihood function for the timing residuals using Eqs. 3.1.5 and 3.1.10 as

$$p(\delta\mathbf{t}|\vec{\theta}, \delta\boldsymbol{\xi}) = \frac{\exp\left(-\frac{1}{2}(\delta\mathbf{t} - \mathbf{M}\delta\boldsymbol{\xi})^T \boldsymbol{\Sigma}_n^{-1} (\delta\mathbf{t} - \mathbf{M}\delta\boldsymbol{\xi})\right)}{\sqrt{\det(2\pi\boldsymbol{\Sigma}_n)}}, \tag{3.1.20}$$

where $\delta\mathbf{t}$ and $\delta\boldsymbol{\xi}$ are defined in an identical manner as $\mathbf{n}$ as the concatenated vector or residuals and timing parameters for each pulsar, respectively. Note that here we use $\delta\mathbf{t}$ instead of $\delta\mathbf{t}^{\rm pre}$ since this process can be thought of as another step in the iterative process of timing (where the post-fit residuals are formed from the previous set of pre-fit residuals); instead of minimizing chi-squared using $\mathbf{W}$ as the noise covariance, we now use the full noise covariance matrix $\boldsymbol{\Sigma}_n$ and the full PTA dataset to *maximize* the likelihood. In [81] it was shown that this likelihood can be maximized[1] analytically over the timing model parameters to give

$$p(\delta\mathbf{t}|\vec{\theta}) = \frac{\exp\left(-\frac{1}{2}\delta\mathbf{t}^T \mathbf{G}(\mathbf{G}^T\boldsymbol{\Sigma}_n\mathbf{G})^{-1}\mathbf{G}^T\delta\mathbf{t}\right)}{\sqrt{\det(2\pi\boldsymbol{\Sigma}_n)}}, \tag{3.1.21}$$

where $\mathbf{G}_I$ is an $N_{\rm TOA} \times (N_{\rm TOA} - N_{\rm par})$ matrix. The matrix $\mathbf{G}_I^T$ spans the null space of $\mathbf{M}_I$ and will project the data onto a subspace orthogonal to the linearized timing model. The full PTA G-matrix is then

$$\mathbf{G} = \begin{bmatrix} \mathbf{G}_1 & 0 & \dots & 0 \\ 0 & \mathbf{G}_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \mathbf{G}_M \end{bmatrix}. \tag{3.1.22}$$

[1] In [81], the authors actually *marginalize* the likelihood function over the pulsar timing parameters; however, when using uniform priors the resulting likelihood after maximizing or marginalizing only differs by a factor of $\det(\mathbf{M}^T\boldsymbol{\Sigma}_n\mathbf{M})$, so the data dependent part of the likelihood remains the same.

For the remainder of this chapter we will use the following notation

$$\mathbf{r}_I = \mathbf{G}_I^T \delta\mathbf{t}_I \tag{3.1.23}$$

$$\mathbf{P}_I = \mathbf{G}_I^T \mathbf{N}_I \mathbf{G}_I \tag{3.1.24}$$

$$\mathbf{S}_{IJ} = \mathbf{G}_I^T \mathbf{X}_{IJ} \mathbf{G}_J \tag{3.1.25}$$

$$\mathbf{\Sigma} = \mathbf{G}^T \mathbf{\Sigma}_n \mathbf{G}, \tag{3.1.26}$$

with the likelihood function written as

$$p(\mathbf{r}|\vec{\theta}) = \frac{1}{\sqrt{\det(2\pi\mathbf{\Sigma}_n)}} \exp\left(-\frac{1}{2}\mathbf{r}^T\mathbf{\Sigma}^{-1}\mathbf{r}\right). \tag{3.1.27}$$

Time-Domain Optimal Statistic

In [45] the authors presented the optimal cross-correlation statistic in both the frequency and time domains, with a focus on the frequency-domain implementation. The non-stationarity that arises from the timing model fit (Eq. (3.1.9)), along with the irregular sampling that is typical of realistic PTA data sets, however, make frequency-domain techniques unsuitable for PTA gravitational-wave data analysis. Therefore in this chapter we will focus on the time-domain implementation of the cross-correlation statistic. In [45] the time-domain derivation was done by constructing the likelihood ratio of a model that contained a stochastic gravitational-wave background and intrinsic noise to a model that contained only intrinsic noise. It was assumed that the amplitude of the intrinsic noise is much larger than the amplitude of the gravitational-wave background, and thus can be safely ignored in the auto-covariance matrices of the residuals. One can then perform an expansion of the log-likelihood ratio in powers of a small order parameter taken to represent the amplitude of the background. This assumption can lead to a significant bias in the recovered amplitude of the gravitational-wave background if the background is sufficiently large.

Fortunately it is possible to carry out a nearly identical derivation that takes into account a potential non-negligible contribution of the stochastic background to the auto-covariance terms. In [82] it was shown that it is possible to expand the covariance matrix $\mathbf{\Sigma}$ in a Taylor series expansion in the Hellings and Downs coefficients (as opposed to

an expansion in the amplitude of the background) to obtain a "first order" likelihood function. The log of this likelihood function can be written as

$$\ln p(\mathbf{r}|\vec{\theta}) \approx -\frac{1}{2}\left[\sum_{I=1}^{M}\left(\mathrm{tr}\ln\mathbf{P}_I + \mathbf{r}_I^T\mathbf{P}_I^{-1}\mathbf{r}_I\right) - \sum_{IJ}\mathbf{r}_I^T\mathbf{P}_J^{-1}\mathbf{S}_{IJ}\mathbf{P}_J^{-1}\mathbf{r}_J\right] \tag{3.1.28}$$

where $\sum_{IJ} = \sum_{I=1}^{M}\sum_{J<I}^{M}$ is a sum over all *unique* pulsar pairs. Let us now assume that we have done a single pulsar noise analysis [81, 103] on each pulsar so that we know $\mathbf{P}_I$, and consider the following log-likelihood ratio

$$\ln\,\Lambda = \ln p(\mathbf{r}|\vec{\theta}_{\rm gw}) - \ln p(\mathbf{r}|\vec{\theta}_{\rm noise}). \tag{3.1.29}$$

Here $\vec{\theta}_{\rm gw}$ are the parameters for a model with a *spatially correlated*[2] gravitational-wave background component along with uncorrelated red and white noise components, which include the gravitational-wave background present in the pulsar term, ISM noise, radiometer noise, jitter noise, etc. The parameters $\vec{\theta}_{\rm noise}$ are for a model with only *spatially uncorrelated* noise components. We treat the auto-covariance of each pulsar as a known measured quantity of the PTA data after the aforementioned noise analysis has been done. In this case, if we fix the spectral index to, say, the one corresponding to SMBBH backgrounds with a spectral index $\alpha = -2/3$, the only free parameter is the amplitude of the gravitational-wave background. Evaluating this log-likelihood ratio we have

$$\ln\,\Lambda = \frac{A_{\rm gw}^2}{2}\sum_{IJ}\mathbf{r}_I^T\mathbf{P}_J^{-1}\tilde{\mathbf{S}}_{IJ}\mathbf{P}_J^{-1}\mathbf{r}_J, \tag{3.1.30}$$

where we have used the amplitude-independent cross-correlation matrix $\tilde{\mathbf{S}}_{IJ}$ defined by

$$A_{\rm gw}^2\tilde{\mathbf{S}}_{IJ} = \langle\mathbf{r}_I\mathbf{r}_J^T\rangle = \mathbf{S}_{IJ}. \tag{3.1.31}$$

Notice that all terms that only include the auto-covariance matrices are cancelled by the noise model likelihood function. Note also that this expression is nearly identical to Eq. (75) of [45] with the caveat that now we are dealing exclusively with post-fit quantities and have allowed for a non-negligible contribution from the gravitational-wave background in the auto-covariance matrices. From Eq. (3.1.30) we define the optimal

[2]By spatially correlated we mean that the correlation is parameterized by the Hellings and Downs curve.

cross-correlation statistic for a PTA to be

$$\hat{A}^2 = \frac{\sum_{IJ} \mathbf{r}_I^T \mathbf{P}_I^{-1} \tilde{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \mathbf{r}_J}{\sum_{IJ} \mathrm{tr}\left[\mathbf{P}_I^{-1} \tilde{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \tilde{\mathbf{S}}_{JI}\right]}, \tag{3.1.32}$$

where the normalization factor

$$\mathcal{N} \equiv \left(\sum_{IJ} \mathrm{tr}\left[\mathbf{P}_I^{-1} \tilde{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \tilde{\mathbf{S}}_{JI}\right]\right)^{-1} \tag{3.1.33}$$

is chosen so that on average $\langle \hat{A}^2 \rangle = A_{\rm gw}^2$. This immediately follows from the observation that

$$\begin{aligned} \left\langle \sum_{IJ} \mathbf{r}_I^T \mathbf{P}_I^{-1} \tilde{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \mathbf{r}_J \right\rangle &= \sum_{IJ} \mathrm{tr}\left[\mathbf{P}_I^{-1} \tilde{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \mathbf{S}_{JI}\right] \\ &= A_{\rm gw}^2 \sum_{IJ} \mathrm{tr}\left[\mathbf{P}_I^{-1} \tilde{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \tilde{\mathbf{S}}_{JI}\right], \end{aligned} \tag{3.1.34}$$

where Eq. (3.1.31) was used in the second line.

In the absence of a cross-correlated signal (or if the signal is weak) the expectation value of $\hat{A}^2$ vanishes and its standard deviation is [45]

$$\sigma_0 = \left(\sum_{IJ} \mathrm{tr}\left[\mathbf{P}_I^{-1} \tilde{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \tilde{\mathbf{S}}_{JI}\right]\right)^{-1/2}, \tag{3.1.35}$$

so if in a particular realization we measure a value of the optimal statistic, the signal-to-noise ratio (SNR) for the *power in the cross-correlations* for that realization is

$$\hat{\rho} = \frac{\hat{A}^2}{\sigma_0} = \frac{\sum_{IJ} \mathbf{r}_I^T \mathbf{P}_I^{-1} \tilde{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \mathbf{r}_J}{\left(\sum_{IJ} \mathrm{tr}\left[\mathbf{P}_I^{-1} \tilde{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \tilde{\mathbf{S}}_{JI}\right]\right)^{1/2}}. \tag{3.1.36}$$

with an expectation value over all realizations of

$$\langle \rho \rangle = A_{\rm gw}^2 \left(\sum_{IJ} \mathrm{tr}\left[\mathbf{P}_I^{-1} \tilde{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \tilde{\mathbf{S}}_{JI}\right]\right)^{1/2}. \tag{3.1.37}$$

Note that this definition of the SNR measures the confidence (in standard deviations) with which we can reject the null hypothesis that there are no *spatially correlated* signals in our data. To clarify this a bit further we outline a standard frequentist hypothesis detection procedure:

1. Measure the optimal statistic value, $\hat{A}^2$ of Eq. (3.1.32), for our data set.

2. Compute the probability $p(\hat{A}^2 > \hat{A}^2_{\rm thresh}|A_{\rm gw} = 0)$, that is, the probability that our measurement of the optimal statistic, $\hat{A}^2$, is greater than some threshold value of the statistic, $\hat{A}^2_{\rm thresh}$, assuming that the null hypothesis, $A_{\rm gw} = 0$, is true.

3. If the aforementioned probability (sometimes called the p-value) is less than some value (this value is set to be a tolerable yet problem specific false-alarm probability (FAP)) then a detection is claimed.

Typically $\hat{A}^2_{\rm thresh}$ is given by

$$\alpha = \int_{-\infty}^{\hat{A}^2_{\rm thresh}} d\hat{A}^2 p(\hat{A}^2|A_{\rm gw} = 0), \tag{3.1.38}$$

where $\alpha = 1 - FAP$ and $p(\hat{A}^2|A_{\rm gw} = 0)$ is the probability distribution function of the optimal statistic given the null hypothesis. To a sufficiently good approximation, $p(\hat{A}^2|A_{\rm gw} = 0)$ can be described by a Gaussian distribution with zero mean and variance given by σ_0^2 (Eq. (3.1.35)), thus the probability $p(\hat{A}^2 > \hat{A}^2_{\rm thresh}|A_{\rm gw} = 0)$ can be expressed in terms of standard deviations away from the mean. For example, if the $\hat{A}^2$ that we measure is 3 standard deviations (i.e. 3-sigma) away from the mean (0 in this case) then this corresponds to a FAP of $\sim$0.003 meaning that we can rule out the null hypothesis with $\sim$99.7% confidence. Returning to Eq. (3.1.36) we see that the typical frequentist detection procedure mentioned above is contained in this definition of SNR. If we measure an SNR of 3, this carries the same meaning as the FAP above.

Figure 11 shows a histogram of the optimal statistic Eq. (3.1.32) in 10^4 simulations for PTA observations of $M = 36$ pulsars, with root-mean-squares (RMSs) $\sigma = 100$ ns, for an observational time $T = 5$ years, and a cadence $c = 20$ yr^{-1}. The black line shows the distribution of the statistic in the absence of a signal, and the gray curve shows the distribution in the presence of a signal with amplitude $A_{\rm gw} = 10^{-14}$ (using the methods described below in Section 3.2). The standard deviation of the distribution in the absence of a signal is $\sigma_0 = 1.08 \times 10^{-29}$. As shown in the figure, in the absence of a signal the distribution is not quite Gaussian, but using the true cumulative distribution of the simulations and the 3-σ Gaussian distribution threshold gives a FAP of ~ 0.006.

The optimal statistic in Eq. (3.1.32) has also been used to analyze the data sets produced for the International PTA Mock Data Challenge. In this challenge, the optimal

statistic was used to produce amplitude estimates for three closed data sets. These amplitudes were then compared to those from a first-order likelihood method (described in [82]). The amplitudes recovered using the optimal statistic were consistent with the first-order likelihood methods at the 95% level or better. Readers are encouraged to consult [68] for more details regarding the Mock Data Challenge and the results obtained using the optimal statistic.

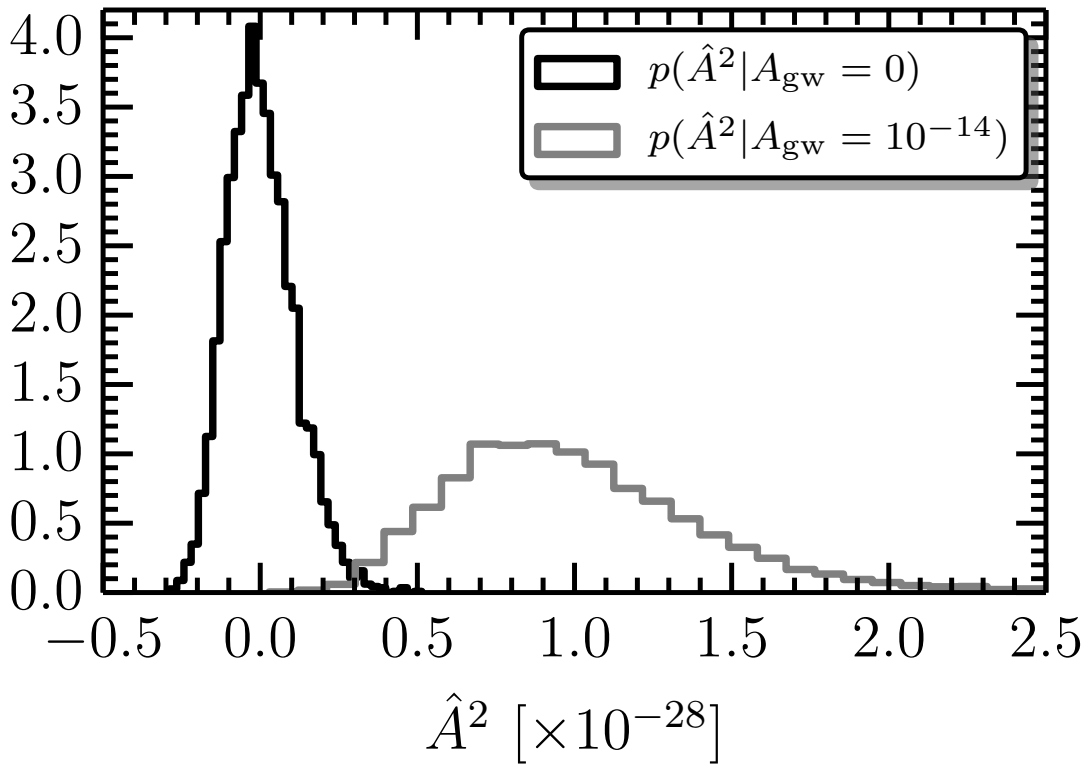

Figure 11: Histogram of the optimal statistic Eq. (3.1.32) in 10^4 simulations for a PTA consisting of $M = 36$ pulsars, all with RMSs $\sigma = 100$ ns, an observational time $T = 5$ years, and a cadence $c = 20$ yr^{-1}. We show the distribution of the statistic in the absence of a signal (black line), and the distribution in the presence of a signal with amplitude $A_{\rm gw} = 10^{-14}$ (gray line). The standard deviation of the distribution in the absence of a signal is $\sigma_0 = 1.08 \times 10^{-29}$.

3.2 Simulated signals

In this section we describe a software injection procedure that can be used to produce simulated stochastic background signals in PTA data. As we have shown, if a stochastic gravitational-wave background is present, the cross-correlation of timing residuals is given

by Eq. (2.3.7)

$$\langle \tilde{r}_I^*(f)\tilde{r}_J(f')\rangle = \frac{H_0^2}{16\pi^4}\delta(f-f')|f|^{-5}\Omega_{\rm gw}(f)\chi_{IJ}. \tag{3.2.1}$$

In the frequency domain it is possible to express the timing residuals as

$$r_I(f) = c(f)\sum_J H_{IJ}w_J(f), \tag{3.2.2}$$

where $w_I(f) = x_I(f) + iy_I(f)$ is a complex zero-mean white noise process, $c(f)$ is a real function that contains information about the spectral index and amplitude of the gravitational-wave spectrum (but does not depend on the pulsar pair), and H_{IJ} is a matrix that linearly combines the timing residuals in such a way as to simulate the expected spatial correlations in the signal so that the Hellings and Downs coefficients are recovered.

If the processes x_I and y_I are zero-mean unit-variance processes, i.e.

$$\langle x_I(f)\rangle = 0,\ \langle y_I(f)\rangle = 0 \tag{3.2.3}$$

and

$$\langle x_I^*(f)x_I(f')\rangle = 1,\ \langle y_I^*(f)y_I(f')\rangle = 1 \tag{3.2.4}$$

then $w_I(f)$ satisfies

$$\langle w_I^*(f)w_J(f')\rangle = \frac{2}{T}\delta(f-f')\delta_{IJ}, \tag{3.2.5}$$

where T is the length of observations, and we can use Eq. (3.2.1) to find $c(f)$ and H_{IJ}. Taking the ensemble average of Eq. (3.2.2) it is easy to show that

$$\langle \tilde{r}_I^*(f)\tilde{r_J}(f')\rangle = \frac{2}{T}c(f)c(f')\sum_K H_{IK}H_{KJ}\delta(f-f'), \tag{3.2.6}$$

which implies that

$$c^2(f)\sum_K H_{IK}H_{KJ} = \frac{TH_0^2}{32\pi^4}|f|^{-5}\Omega_{\rm gw}(f)\chi_{IJ}. \tag{3.2.7}$$

In matrix notation the equation above can be written as

$$c^2(f)\mathbf{H}\mathbf{H}^{\mathbf{T}} = \frac{TH_0^2}{32\pi^4}|f|^{-5}\Omega_{\rm gw}(f)\boldsymbol{\chi}. \tag{3.2.8}$$

Relating the functions of frequency on either side of Eq. (3.2.8), we readily identify the function $c(f)$ to be

$$c(f) = \left[\frac{TH_0^2}{32\pi^4}\Omega_{\rm gw}(f)|f|^{-5}\right]^{1/2}, \tag{3.2.9}$$

along with a condition for the matrix $\mathbf{H}$,

$$\mathbf{H}\mathbf{H}^{\mathbf{T}} = \boldsymbol{\chi} \tag{3.2.10}$$

which allows us to determine $\mathbf{H}$ given $\boldsymbol{\chi}$ via a Cholesky decomposition.

To construct simulated timing residuals one can: (1) start with M random complex frequency series $w_I(f)$, where M is the number of pulsars, (2) multiply these by $c(f)$, (3) find the Hellings and Downs coefficients for all pulsar pairs and construct the matrix $\boldsymbol{\chi}$, (4) perform a Cholesky decomposition of $\boldsymbol{\chi}$ to find $\mathbf{H}$, and (5) linearly combine the frequency series via Eq. (3.2.2) to find $r_I(f)$ for each pulsar. Finally, after inverse Fourier transforming the gravitational-wave residuals, they can be added to real or simulated TOA data that contains additional uncorrelated white and red noise components.

3.3 Scaling Laws for the Optimal Cross-correlation statistic

In [104] the authors considered a simple scenario where pulsar timing residuals have just two noise components, a gravitational-wave red noise piece and a white-noise piece, which are the same for all pulsars in the PTA, namely

$$\mathcal{P}_I(f) = \mathcal{P}_g(f) + 2\sigma^2\Delta t = bf^{-\gamma} + 2\sigma^2\Delta t. \tag{3.3.1}$$

Here all the frequency independent constants in Eq. (3.1.18) have been absorbed into the amplitude b, the index $\gamma = 3 - 2\alpha$ (recall that we are using one-sided power spectra in this chapter, in contrast to [104]), and the white noise RMS is denoted by σ.

In [104] it was shown that the SNR of the optimal cross-correlation scales in three different ways depending on the relative sizes of the gravitational-wave and white-noise components. Specifically the authors found scaling laws for the SNR in

(i) a weak signal regime where the white noise component of Eq. (3.3.1) is larger than the gravitational wave piece ($2\sigma_I^2\Delta t \gg bf^{-\gamma}$ at all relevant frequencies),

(ii) the opposite strong signal limit, where $2\sigma_I^2\Delta t \ll bf^{-\gamma}$ at all relevant frequencies, which turns out to be irrelevant for pulsar timing experiments, and,

(iii) an intermediate regime between the two cases where the gravitational wave power spectrum dominates at low frequencies, and the white noise dominates at high frequencies.

Additionally, they found that the latter regime is likely already relevant to current pulsar timing experiments. In this section we will review the scaling laws for the optimal statistic, and introduce an improved derivation of the scaling law for the intermediate regime.

To derive the scaling laws we begin with the expression for the expected SNR of the cross-correlation statistic,

$$\langle\rho\rangle = A_{\rm gw}^2 \left(\sum_{IJ} \mathrm{tr}\left[\mathbf{P}_I^{-1} \tilde{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \tilde{\mathbf{S}}_{JI} \right] \right)^{1/2}, \tag{3.3.2}$$

which can be written in the frequency domain as [45]

$$\langle\rho\rangle = \left(2T \sum_{IJ} \chi_{IJ}^2 \int_{f_L}^{f_H} df \, \frac{\mathcal{P}_g^2(f)}{\mathcal{P}_I(f)\mathcal{P}_J(f)} \right)^{1/2}. \tag{3.3.3}$$

Since we are assuming that all pulsars have the same noise characteristics we can write

$$\langle\rho\rangle = \left(\sum_{IJ} \chi_{IJ}^2 \right)^{1/2} \left(2T \int_{f_L}^{f_H} df \, \frac{b^2 f^{-2\gamma}}{\left(bf^{-\gamma} + 2\sigma^2\Delta t\right)^2} \right)^{1/2}. \tag{3.3.4}$$

In the weak signal regime, where $2\sigma_I^2 \Delta t \gg b f^{-\gamma}$ for all frequencies of interest, i.e., $f \in [f_L, f_H]$, the SNR is well approximated by

$$\langle\rho\rangle \approx \left(\sum_{IJ} \chi_{IJ}^2 \right)^{1/2} \frac{bcT^\gamma}{2\sigma^2\sqrt{\gamma - 1/2}}, \tag{3.3.5}$$

where $c = 1/\Delta t$ is the cadence, $f_L = 1/T$ and $f_H >> f_L$.

In the intermediate regime we cannot use this approximation because at low frequencies the power in the gravitational-wave backround is larger than the white noise level. Note that this happens when $bT^\gamma > 2\sigma^2\Delta t$, and the condition on the white noise RMS is

$$\sigma < \frac{A}{\pi f_{\rm 1yr}^\alpha} \sqrt{\frac{cT^\gamma}{24}}. \tag{3.3.6}$$

For pulsar timing experiment durations of $T = 5$ yr, cadence of $c = 20\ \mathrm{yr}^{-1}$, a background with amplitude $A = 10^{-15}$, and a spectral index like the one we expect for the SMBBH background ($\gamma = 13/3$), the pulsar timing array is in the weak signal limit only if the

pulsars have white noise RMSs greater than about 300 ns. There are already a handful of pulsars that are currently timed with better precisions than that (see, for example, [41]).

In this case the integral in Eq. (3.3.4) must evaluated generally. To do this, we write the integral as

$$\int_{f_L}^{f_H} df\, F(f) = \int_0^{f_H} df\, F(f) - \int_0^{f_L} df\, F(f) \tag{3.3.7}$$

where for convenience we have written

$$F(f) = \frac{bf^{-2\gamma}}{\left(bf^{-\gamma} + 2\sigma^2\Delta t\right)^2}. \tag{3.3.8}$$

The integrals on the right hand side of Eq. (3.3.7) have analytic solutions in terms of ordinary hypergeometric functions. To proceed, we evaluate the integral of $F(f)$ over a generic interval $[0, f_*]$ which yields

$$\int_0^{f_*} df F(f) = \frac{f_*}{\gamma}\left[\frac{1}{1+\frac{2\sigma^2\Delta t}{bf_*^{-\gamma}}} + (\gamma-1)\, G\left(\frac{-2\sigma^2\Delta t}{bf_*^{-\gamma}}\right)\right], \tag{3.3.9}$$

where $G(x) = {}_2F_1\left(1, \gamma^{-1}, 1+\gamma^{-1}, x\right)$. We can probe this solution in the context of Eq. (3.3.7) by replacing f_* with f_H or f_L.

For the second integral on the right hand side of Eq. (3.3.7) where $f_* = f_L = 1/T$, we have $(2\sigma^2\Delta t)/(bf_L^{-\gamma}) \ll 1$ and the hypergeometric function can be approximated to be unity

$${}_2F_1\left(1, \gamma^{-1}, 1+\gamma^{-1}, \frac{-2\sigma^2\Delta t}{bf_L^{-\gamma}}\right) \approx 1.$$

This simplifies Eq. (3.3.9) greatly, and the integral is easily evaluated as

$$\int_0^{f_L} df\, F(f) \approx \frac{1}{T}. \tag{3.3.10}$$

To evaluate the first integral in Eq. (3.3.7), we consider the case when $f_* = f_H$ in Eq. (3.3.9). In this case, since $(2\sigma^2\Delta t)/(bf_H^{-\gamma}) \gg 1$, the integral can be approximated as

$$\int_0^{f_H} df\, F(f) \approx \frac{f_H}{\gamma}\left[\frac{bf_H^{-\gamma}}{2\sigma^2\Delta t} + (\gamma-1)\, G\left(\frac{-2\sigma^2\Delta t}{bf_H^{-\gamma}}\right)\right], \tag{3.3.11}$$

We can then use standard identities relating the hypergeometric function to inverses of their arguments (see, for example, Eq. (15.8.2) in [105]). Using these identities along with

with Euler's reflection formula we obtain

$$\int_0^{f_H} df\, F(f) \approx \frac{f_H}{\gamma}\left\{\frac{bf_H^{-\gamma}}{2\sigma^2\Delta t} + (\gamma-1)\,\Gamma(\gamma^{-1}-1)\Gamma(2-\gamma^{-1})\Gamma(1+\gamma^{-1})\right.$$
$$\times\left[\frac{bf_H^{-\gamma}}{2\sigma^2\Delta t}\frac{\Gamma(2-\gamma^{-1})^{-1}}{\Gamma(\gamma^{-1})^2}\,{}_2F_1\left(1,1-\gamma^{-1},2-\gamma^{-1};\frac{-bf_H^{-\gamma}}{2\sigma^2\Delta t}\right)\right.$$
$$\left.\left.-\frac{1}{\Gamma(\gamma^{-1})}\left(\frac{bf_H^{-\gamma}}{2\sigma^2\Delta t}\right)^{1/\gamma}\,{}_2F_1\left(\gamma^{-1},0,\gamma^{-1};\frac{-bf_H^{-\gamma}}{2\sigma^2\Delta t}\right)\right]\right\}. \quad (3.3.12)$$

Since $bf_H^{-\gamma}/2\sigma^2\Delta t \ll 1$ both hypergeometric functions can be well approximated by unity. Additionally, since $bf_H^{-\gamma}/2\sigma^2\Delta t \ll \left(bf_H^{-\gamma}/2\sigma^2\Delta t\right)^{1/\gamma}$ for $\gamma > 1$, the last term in Eq. (3.3.12) dominates and the expression can be simplified to

$$\int_0^{f_H} df\, F(f) \approx \kappa(\gamma)\left(\frac{b}{2\sigma^2\Delta t}\right)^{1/\gamma} \quad (3.3.13)$$

with

$$\kappa(\gamma) = (1-\gamma^{-1})\frac{\pi/\gamma}{\sin(\pi/\gamma)}. \quad (3.3.14)$$

Putting the results of Eq. (3.3.10) and Eq. (3.3.13) together, we arrive at the solution to the original problem posed in Eq. (3.3.7):

$$\int_{f_L}^{f_H} df\, F(f) \approx \kappa(\gamma)\left(\frac{b}{2\sigma^2\Delta t}\right)^{(1/\gamma)} - \frac{1}{T}. \quad (3.3.15)$$

In terms of the cadence $c = 1/\Delta t$ the average value of the SNR is therefore given by

$$\langle\rho\rangle \approx \left(\sum_{IJ}\chi_{IJ}^2\right)^{1/2}\left[2T\left(\kappa(\gamma)\left(\frac{bc}{2\sigma^2}\right)^{(1/\gamma)} - \frac{1}{T}\right)\right]^{1/2}. \quad (3.3.16)$$

At late times,

$$\langle\rho\rangle \approx \left(\sum_{IJ}\chi_{IJ}^2\right)^{1/2}\left[2T\kappa(\gamma)\left(\frac{bc}{2\sigma^2}\right)^{(1/\gamma)}\right]^{1/2}$$
$$\propto M\left(\frac{cA_{\rm gw}^2}{2\sigma^2}\right)^{1/(2\gamma)} T^{1/2}. \quad (3.3.17)$$

In [104] the authors approximated the integral in a less accurate (albeit more pedagogical) way: they found the frequency $f_r = (bc/2\sigma^2)^{1/\gamma}$ at which the gravitational wave red noise equals the white noise, and assumed the integral was gravitational wave dominated at frequencies lower than f_r, and white noise dominated at frequencies higher than f_r. The integrals then become trivial. The result is the same as Eq. (3.3.16), but with a different value of the coefficient κ which was found to be $\kappa' = 2\gamma/(2\gamma-1)$. In the

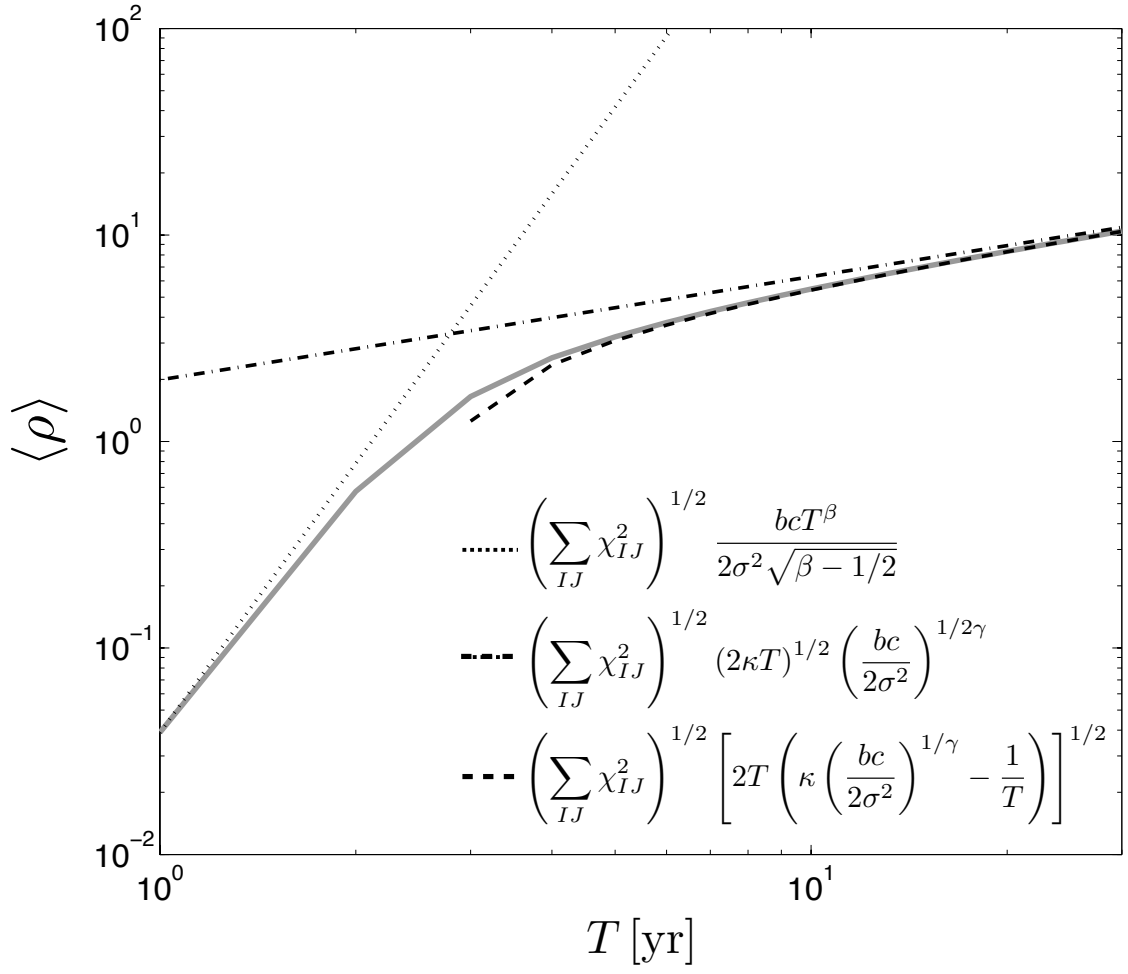

Figure 12: Average SNR versus time in years for PTA with 20 pulsars timed with a precision of $\sigma = 50$ ns and a gravitational-wave background produced by SMBBHs ($\gamma = 13/3$) with an amplitude $A_{\rm gw} = 10^{-15}$. The gray curve shows the SNR computed numerically using Eq. (3.3.2). The dotted curve shows SNR in the weak-signal limit, Eq. (3.3.5). The dashed-dot curve shows the SNR in the intermediate regime at late times, Eq. (3.3.17). The dashed curve shows the SNR calculated using Eq. (3.3.16).

approximation the integrand for the SNR is always over-estimated and the value of κ' is larger than what we have calculated for κ in this chapter.

Figure 12 shows the average SNR versus time in years for PTA with 20 pulsars timed with a presicion of $\sigma = 50$ ns and a gravitational-wave background produced by SMBBHs ($\gamma = 13/3$) with an amplitude $A_{\rm gw} = 10^{-15}$. The gray curve shows the SNR computed numerically in the time domain using Eq. (3.3.2). For the timing model we have subtracted out a quadratic—i.e., we have fitted out a quadratic with the $\mathbf{R}$ projection matrices in the time domain. The dotted curve shows the average SNR as computed in the weak-signal limit using Eq. (3.3.5). The dashed-dot curve shows the SNR in the intermediate

regime at late times as calculated using Eq. (3.3.17). Finally, the dashed curve shows the SNR calculated using Eq. (3.3.16). At very early times the approximation is not valid: the first term in the square root is smaller than $1/T$ so the SNR is imaginary. At later times the predicted SNR is in excellent agreement with the time-domain numerical calculation. Note the remarkable accuracy with which the low frequency cutoff $f_L = 1/T$ approximates the effect of quadratic subtraction.

3.4 Summary

In this chapter, we have presented a time-domain implementation of the optimal cross-correlation statistic for stochastic gravitational-wave background searches using PTA data, originally presented in [45]. The derivation and implementation described here extends that of [45] by taking the timing model into account in a natural and statistically well-motivated way by including the linear timing model directly into the likelihood function, allowing for analytic maximization of the timing model parameters. The time-domain implementation also allows one to fully model the noise and naturally deal with non-stationarities and irregular sampling of the data, which cannot be modeled in the frequency domain.

An alternative approach for analyzing PTA data for stochastic gravitational-wave backgrounds is to use Bayesian inference, as described in [77, 40, 82, 106, 107]. In the Bayesian approach, one constructs the posterior probability distributions for the noise and gravitational-wave signal parameters via Bayes' theorem by specifying the likelihood function for the data given a set of model gravitational-wave and noise parameters and a prior distribution on the model parameters. By marginalizing over the model parameters, one also constructs the Bayesian evidence for various models, which allow for the construction of Bayes factors (ratio of Bayesian evidence) to determine which model is favored by the data.

While we believe that a Bayesian approach to the detection problem for stochastic backgrounds is preferred and indeed recommended, the frequentist cross-correlation statistic presented here has several advantages over the Bayesian approach. Firstly, the optimal statistic approach is computationally inexpensive as it involves only a single

function call (given a set of modeled noise parameters), while the Bayesian method must explore a very large dimensional space leading to millions of likelihood evaluations. For current data sets, the optimal statistic can be evaluated in seconds while the full Bayesian approach can take weeks to run on a super computer.

Furthermore, the SNR as defined in this work is a good approximation to the Bayes factor comparing a model for a correlated gravitational-wave background to a model for an uncorrelated intrinsic red noise source. Thus the computationally inexpensive optimal statistic has proven invaluable in large scale simulations and projections of detector sensitivity as it allows us to test many different signal models and pulsar observation scenarios with relative ease, while full Bayesian simulations on this scale are unfeasible. In addition, the relationship between the the optimal statistic SNR and the Bayes factors affords an analytically tractable environment from which to construct various scaling relations as shown in Figure 12.

The optimal statistic does have two major drawbacks that make it less desirable as a production-level detection statistic compared to the Bayes factor. Firstly, the point estimate of the amplitude of the gravitational-wave background depends on our ability to accurately model the total autocorrelated power for each pulsar. Typically this is done by modeling the noise for each pulsar *independently* and then including the maximum likelihood values in the auto-covariance matrices of the optimal statistic. If the signal is loud and the data does not contain any intrinsic red noise then this method is fairly robust and does not significantly bias results. However, if the signal is weak or there is other intrinsic red noise then this method will lead to biases. In low SNR scenarios the red noise due to the stochastic background may not be large enough to detect in an individual pulsar and will thus not enter the auto-covariance matrices used in the optimal statistic. This will lead to an inconsistency in the optimal statistic where it will still be able to detect cross-correlated power, but the point estimate of the amplitude will be biased low because the auto-covariance terms (from our single pulsar noise analysis) indicate that the red noise is very weak.

This problem does not arise in Bayesian analyses because the intrinsic pulsar noise and the stochastic background parameters are modeled *simultaneously*. This problem could be

ameliorated by performing the initial noise modeling over all pulsars simultaneously and including a correlated gravitational-wave background component. These noise estimates (which will include a common gravitational-wave background term in the auto-covariance) could then be input to the optimal statistic.

Despite these drawbacks, the optimal cross-correlation statistic serves as a proxy for a full Bayesian search when performing computationally-intensive simulations and will also serve as a very useful cross-check when making detection statements on future PTA data.

3.5 Relation to Demorest et al. Cross-Correlation Statistic

Here we show that the optimal statistic, although derived in a different manner, is identical to the cross-correlation statistic presented in [41]. In the notation used in this work, the cross-correlation coefficients can be written as

$$\rho_{IJ} = \frac{\mathbf{r}_I^T \mathbf{P}_I^{-1} \hat{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \mathbf{r}_J}{\mathrm{tr}\left[\mathbf{P}_I^{-1} \hat{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \hat{\mathbf{S}}_{JI}\right]}, \tag{3.5.1}$$

where $\hat{\mathbf{S}}_{IJ}$ is defined so that $A_{\rm gw}^2 \chi_{IJ} \hat{\mathbf{S}}_{IJ} = \mathbf{S}_{IJ}$. The uncertainty on the correlation coefficients is

$$\sigma_{IJ} = \left(\mathrm{tr}\left[\mathbf{P}_I^{-1} \hat{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \hat{\mathbf{S}}_{JI}\right]\right)^{-1/2}. \tag{3.5.2}$$

With these expressions we now have an estimate of the cross-correlation coefficients along with their uncertainty for each pulsar pair. Notice that only the spectral shape of the gravitational wave background is assumed. To determine an estimate of the gravitational wave background amplitude, the following chi-squared is minimized

$$\chi^2 = \sum_{IJ} \left(\frac{\rho_{IJ} - A_{\rm gw}^2 \chi_{IJ}^2}{\sigma_{IJ}}\right)^2. \tag{3.5.3}$$

The resulting best fit gravitational wave amplitude is

$$\hat{A}_{\rm gw}^2 = \sum_{IJ} \frac{\rho_{IJ} \chi_{IJ}}{\sigma_{IJ}^2} \Bigg/ \sum_{IJ} \frac{\chi_{IJ}^2}{\sigma_{IJ}^2}, \tag{3.5.4}$$

with variance

$$\sigma^2 = \left(\sum_{IJ} \frac{\chi_{IJ}^2}{\sigma_{IJ}^2}\right)^{-1}. \tag{3.5.5}$$

By using Eqs. 3.5.1 and 3.5.2 and by noting that $\chi_{IJ}\hat{\mathbf{S}}_{IJ} = \tilde{\mathbf{S}}_{IJ}$, we obtain

$$\hat{A}^2_{\rm gw} = \frac{\sum_{IJ} \mathbf{r}_I^T \mathbf{P}_I^{-1} \tilde{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \mathbf{r}_J}{\sum_{IJ} \mathrm{tr}\left[\mathbf{P}_I^{-1} \tilde{\mathbf{S}}_{IJ} \mathbf{P}_J^{-1} \tilde{\mathbf{S}}_{JI}\right]}, \tag{3.5.6}$$

which is identical to Eq. (3.1.32).

Chapter 4

Testing General Relativity with Pulsar Timing Arrays

> "*You look at science (or at least talk of it) as some sort of demoralising invention of man, something apart from real life, and which must be cautiously guarded and kept separate from everyday existence. But science and everyday life cannot and should not be separated.*"
>
> — Rosalind Franklin, chemist and x-ray crystallographer

General relativity is among the most successful theories of physics in the 20th century, passing all current weak-field, slow motion tests with flying colors. Progress in cosmology and high energy physics over the course of the last 50 years, however, has brought with it questions that may be unanswerable in the context of general relativity. The accelerated expansion of the universe, the dark matter problem, and inflation have led some authors to re-examine general relativity and attempt to modify it to explain some of these puzzles. Additionally, the incompatibility between general relativity and quantum field theory may be an indication that modifications to general relativity are necessary.

A number of alternative theories of gravity have been proposed to address some of these problems. Those which satisfy the Einstein Equivalence Principle are called metric theories of gravity. In these theories, the only gravitational fields that may influence matter are the components of the metric tensor $g_{\mu\nu}$. Additional fields play the role of generating spacetime curvature. Metric theories are grouped broadly into several categories: scalar tensor theories, in which a dynamical scalar field ϕ is present in addition to

the metric (see Refs. [108, 109, 110, 111, 112, 113, 114, 115, 116]); vector-tensor theories, which contain a dynamic gravitational four-vector field in addition to the metric (see Refs. [114, 117, 118, 116, 119]); and bimetric theories, which are characterized by "prior" geometry contained in dynamical scalar, vector or tensor fields (see Refs. [114, 116, 120]).

In a metric theory of gravity, the six independent components of the Riemann tensor provide up to six possible gravitational wave (GW) polarization states, four more than those allowed in general relativity. Detection of any extra GW polarization states would be fatal for general relativity. A non-detection could be used put constraints on the parameters of alternative theories of gravity.

Previous work on stochastic backgrounds of gravitational waves in the context of alternative theories of gravity has shown that three ground-based interferometers are sufficient to disentangle the polarization content of a general metric theory of gravity [121]. For pulsar timing arrays the form of the correlation between pulsar pairs as a function of pulsar pair angular separation depends on the polarization content of the theory [122]. Additionally it has been shown that pulsar timing arrays have a greater sensitivity to longitudinal and vector polarization modes than to transverse modes [122, 123].

It is also possible to investigate the problem of stochastic GW detection using PTAs in the context of the optimal statistic. In this chapter, the expected cross-correlations for pulsar timing arrays are determined for stochastic backgrounds of GWs for any metric theory of gravity. These cross-correlations are proportional to the so-called *overlap reduction function*, a function that characterizes the PTA response to a GW given the Earth-pulsar-GW geometry. It is the behavior of this function that describes the physical origin of the increased sensitivity to scalar-longitudinal and vector polarization modes.

This chapter is broken down as follows: In Section 4.1, the gravitational-wave induced redshift is determined for polarizations outside of general relativity. This will closely mirror the derivation of the redshift induced by relativistic gravitational waves, which was done in Section 2.2 of Chapter 2, although a slightly different sign convention is used in defining the redshift.

In Section 4.2 the optimal cross-correlation filter is written by maximizing the signal to noise for a pulsar pair, and the overlap reduction function for GWs of any metric theory

of gravity is expressed. In Section 4.3 the effect of GWs of various polarizations on the pulsar-Earth-GW system is probed, and the physical origin of the increased sensitivity to longitudinal and shear modes elucidated. This effect is most easily understood in the frequency domain. In Section 4.4, the form of the overlap reduction function for transverse vs. non-transverse GWs is discussed. The results of this analysis indicate that for the scalar-longitudinal and vector (shear) modes, the overlap reduction functions are frequency dependent in the ranges of frequencies and distances relevant to pulsar timing. This is not the case for the transverse tensor and breathing modes. In Section 4.5, overlap reduction functions are numerically determined for the a subset of NANOGrav pulsars. The resulting values of the overlap reduction function for non-transverse GWs show that sensitivity to the scalar-longitudinal and vector (shear) modes increases by at least an order of magnitude for nearby pulsar pairs for vector modes, and about four orders of magnitude for the longitudinal mode. Results from this chapter are summarized in Section 4.6. As in the last chapter, work is done in units where the speed of light $c = 1$.

Fig. 15 plots the antenna patterns for the various GW polarization modes in a system where the GW's direction of propagation is fixed and the pulsar's position is varied (see Appendix 4.A, Eqns. (4.A.20), (4.A.26), (4.A.24) and (4.A.12) for details), as is usually done in the literature.

In general relativity, for the frequency and distance ranges appropriate to pulsar timing experiments (i.e. for $f \gg 1/L$), the overlap reduction function $\Gamma(f)$ approaches a constant which is only a function of the angular separation between the two pulsars. This constant is proportional to the value of the Hellings-Downs curve for the angle between the pulsars [71, 45]. We will see that for longitudinal modes and for tensor modes the overlap reduction function is frequency dependent, even for $f \gg 1/L$, and is considerably larger than for the transverse modes. This indicates an increased sensitivity to such modes. To understand the physical origin of the increased sensitivity we first discuss the effect of GWs in the more simple case of a single pulsar-Earth baseline.

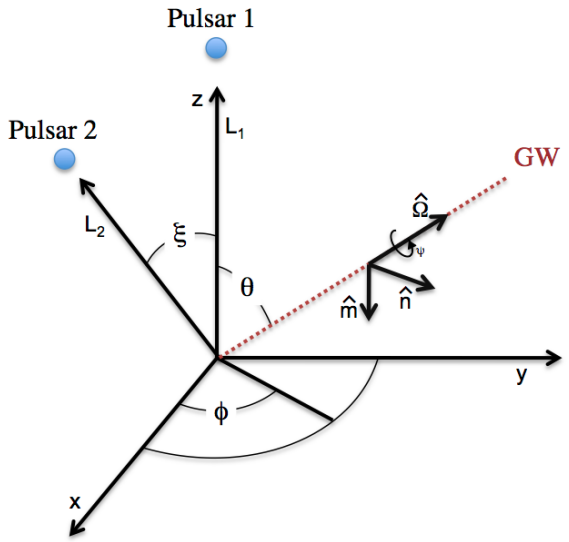

Figure 13: Pulsar positions are given with respect to the Solar System barycenter (located at the origin). Here θ and ϕ are the typical polar and azimuthal angles (as projected from the position of pulsar 1), and pulsar 1 and pulsar 2 are separated by angle ξ. A gravitational wave, characterized by polarization angle ψ, propagates along the $\hat{\Omega}$ direction.

4.1 Detecting gravitational waves with a pulsar timing array

The radio pulses from pulsars arrive at our radio telescopes at very steady rates. Pulsar timing experiments exploit this regularity. Fluctuations in the time of arrival of radio pulses, after all known effects have been accounted for, might be due to the presence of a GW background. If a GW is present the signal from the pulsar can be red-shifted (or blue-shifted). For a GW propagating in the direction $\hat{\Omega}$, the redshift of signals from a pulsar in the direction $\hat{p}$ is given by E q. (2.2.18)

$$z(t,\hat{\Omega}) = \frac{\hat{p}^i\hat{p}^j}{2\left(1+\hat{\Omega}\cdot\hat{p}\right)}[h_{ij}(t_p,\hat{\Omega}) - h_{ij}(t_e,\hat{\Omega})] \tag{4.1.1}$$

where h_{ij} is the metric perturbation and t_p, t_e represent the times the pulse was emitted at the pulsar and the time received at the Solar System barycenter, and we have defined

$$z(t,\hat{\Omega}) = \frac{\nu_e - \nu_p}{\nu_p}. \tag{4.1.2}$$

Note that this is the opposite of the sign convention normally used in the literature [44]. Modified gravity theories extend the possible polarization modes of GWs present in gen-

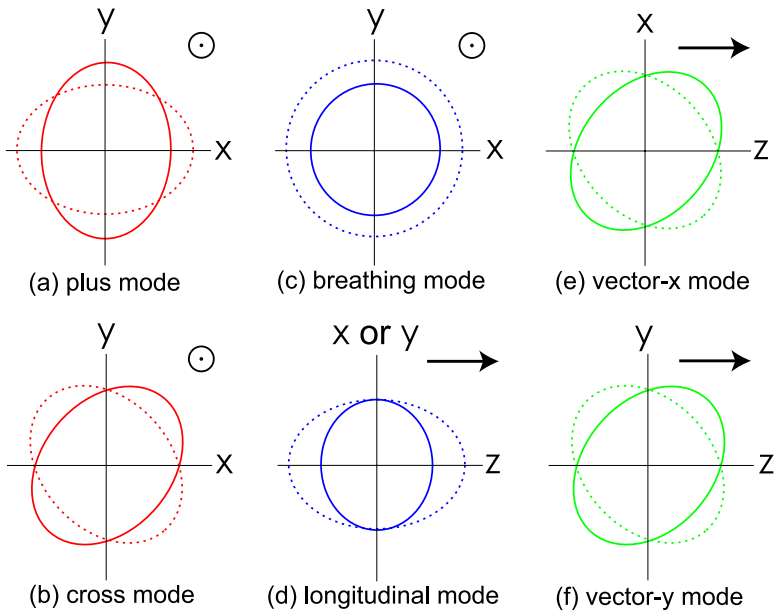

Figure 14: Motion of test masses in response to GWs of the six polarization modes. The plus (+), cross (×), and scalar-breathing (b) mode GWs are transverse, while the two vector modes (x, y) and the scalar-longitudinal (l) mode GWs are non-transverse. Figure reproduced from Nishizawa et al. [121] with permission.

eral relativity – the plus (+) and cross (×) modes– to a maximum of six possible modes. For the two pulsar–Earth system shown in Fig. 13, the GW coordinate system is given by

$$\begin{aligned} \hat{\Omega} &= (\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta) \\ \hat{m} &= (\sin\phi, -\cos\phi, 0) \\ \hat{n} &= (\cos\theta\cos\phi, \cos\theta\sin\phi, -\sin\theta) \end{aligned} \tag{4.1.3}$$

where, relative to [121], we have fixed the GW polarization angle $\psi = -\pi/2$ to agree with the conventions in [46]. From (4.1.3), the GW polarization tensors can be constructed [124, 122, 121, 125, 123]

$$\begin{aligned} \epsilon^{+}_{ij} &= \hat{m}\otimes\hat{m} - \hat{n}\otimes\hat{n}, \quad & \epsilon^{\times}_{ij} &= \hat{m}\otimes\hat{n} + \hat{n}\otimes\hat{m} \\ \epsilon^{b}_{ij} &= \hat{m}\otimes\hat{m} + \hat{n}\otimes\hat{n}, \quad & \epsilon^{l}_{ij} &= \hat{\Omega}\otimes\hat{\Omega} \\ \epsilon^{x}_{ij} &= \hat{m}\otimes\hat{\Omega} + \hat{\Omega}\otimes\hat{m}, \quad & \epsilon^{y}_{ij} &= \hat{n}\otimes\hat{\Omega} + \hat{\Omega}\otimes\hat{n} \end{aligned} \tag{4.1.4}$$

where $\otimes$ is the tensor product and $\hat{\Omega}$ is the direction of GW propagation. Here, x and y correspond to the vector (spin-1) polarization modes while b and l correspond to the scalar

(spin-0) breathing and longitudinal modes, respectively. The plus, cross and breathing modes are characterized by transverse GW propagation, while the longitudinal and vector (or shear) modes are non-transverse in nature (see Fig. 14).

Defining the antenna patterns as

$$F^A(\hat{\Omega}) = \epsilon^A_{ij}(\hat{\Omega})\frac{\hat{p}^i\hat{p}^j}{2(1+\hat{\Omega}\cdot\hat{p})}, \tag{4.1.5}$$

we can define the Fourier transform of (4.1.1) as [122, 45, 125]

$$\tilde{z}(f,\hat{\Omega}) = \left(e^{-2\pi i f L(1+\hat{\Omega}\cdot\hat{p})} - 1\right)\sum_A \tilde{h}_A(f,\hat{\Omega})F^A(\hat{\Omega}) \tag{4.1.6}$$

where the sum is over all possible GW polarizations: $A = +, \times, x, y, b, l$, and L is the distance to the pulsar.

The actual quantity measured in pulsar timing experiments is the timing residual, which is defined as the difference between the actual and expected time of arrival (TOA) of a pulse:

$$R(t) = \text{TOA}_{\text{ actual}} - \text{TOA}_{\text{ expected}}. \tag{4.1.7}$$

The expected TOA for a pulse is modeled and includes daily and yearly motion of the Earth, the position and proper motion of the pulsar, motion about a binary companion (if applicable), etc. The timing residual can be obtained by integrating the redshift in time [44].

In Fig. 15 we plot the antenna patterns for the various GW polarization modes in a system where the GW's direction of propagation is fixed and the pulsar's position is varied (see Appendix 4.A, Eqns. (4.A.20), (4.A.26), (4.A.24) and (4.A.12) for details), as is usually done in the literature.

4.2 GW detection statistic

In this section we introduce the optimal cross correlation statistic [46, 45] for stochastic background searches. The optimal cross-correlation statistic involves the calculation of the overlap reduction function, a geometrical factor that characterizes the loss of sensitivity due to detectors not being co-located or aligned. We will show how the overlap

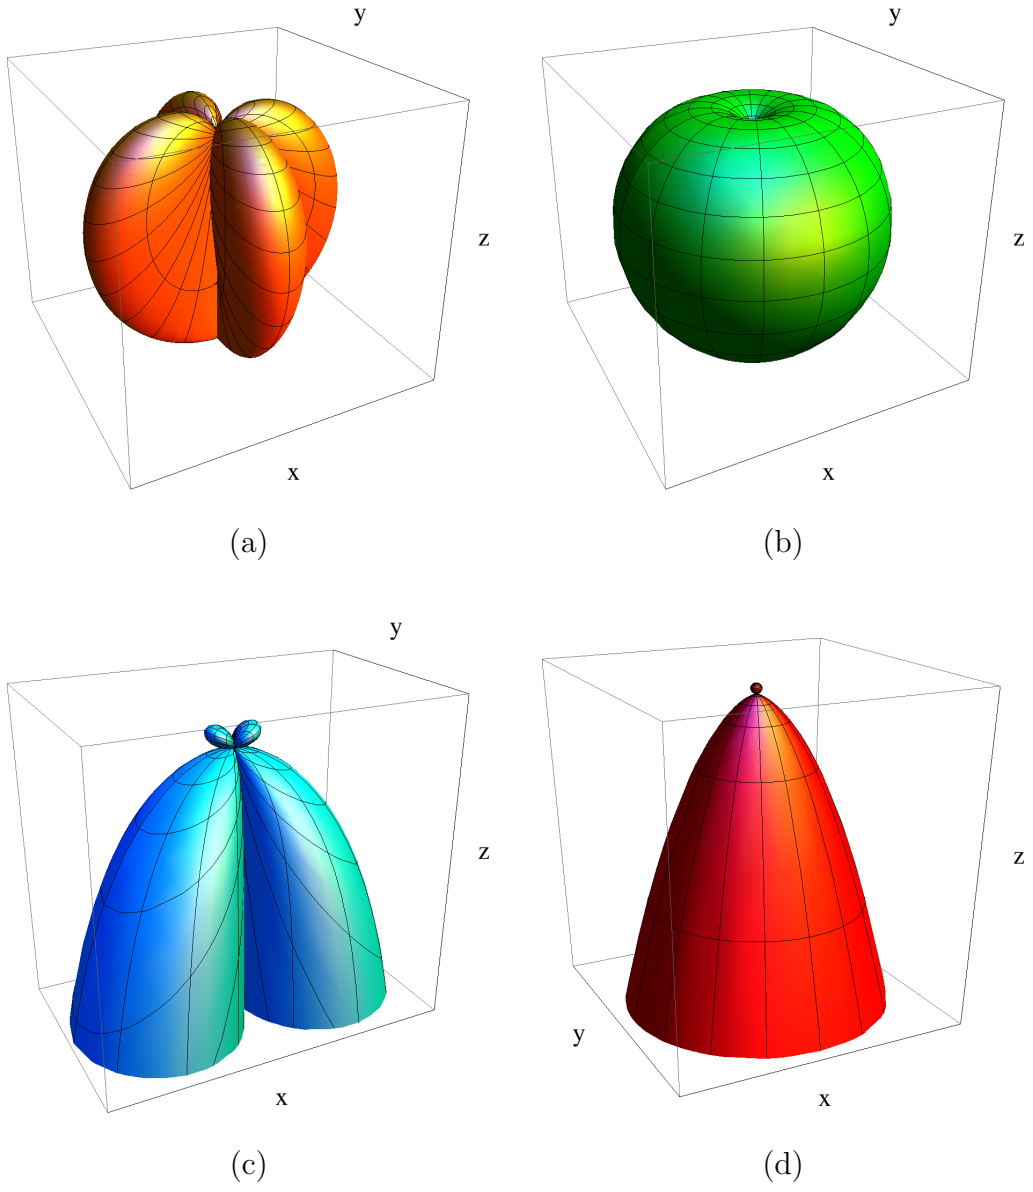

Figure 15: Antenna patterns (4.1.5) for plus/cross (a), breathing (b), vector-x/vector-y (c), and longitudinal (d) polarization modes. Note that the cross and vector-y modes are identical to plus and vector-x, respectively, but rotated by 45 degrees and thus do not appear separately here. In this figure, the GW propagates in the positive z-direction with the Earth at the origin, and the antenna pattern depends on the pulsar's direction, specified by polar angle θ and azimuthal angle ϕ_p. Exact expressions corresponding to each figure may be found in Appendix 4.A: (4.A.20) for the plus mode, (4.A.26) for the breathing mode, (4.A.24) for the vector-x mode, and (4.A.12) for the longitudinal mode. Note that fixing the GW propagation direction while allowing the pulsar position to change is analogous to fixing the pulsar position while allowing the direction of GW propagation to change (there is an inherent degeneracy in the GW polarization angle and the pulsar's azimuthal angle ϕ_p).

reduction function is computed for non-transverse modes. We follow the analysis (for General Relativity) of Allen and Romano [46].

The plane wave expansion for a GW perturbation propagating in the direction $\hat{\Omega}$ is given by [46]

$$h_{ij}(t,\vec{x}) = \sum_A \int_{-\infty}^{\infty} df \int_{S^2} d\Omega e^{2\pi i f(t-\hat{\Omega}\cdot\vec{x})} h_A(f,\hat{\Omega}) \epsilon_{ij}^A(\hat{\Omega}) \tag{4.2.1}$$

where i, j are spatial indices, the sum is over all six polarization states, and the Fourier amplitudes $h_A(f,\hat{\Omega})$ are complex functions satisfying $h_A(-f,\hat{\Omega}) = h_A^*(f,\hat{\Omega})$. A stochastic background of GWs is produced by a large number of weak, independent, unresolvable sources. The energy density of this background per unit logarithmic frequency is given by

$$\Omega_{\rm gw}(f) = \frac{1}{\rho_{\rm critical}} \frac{d\rho_{\rm gw}}{d\ln f} \tag{4.2.2}$$

where $d\rho_{\rm gw}$ is the energy density of the gravitational waves and $\rho_{\rm critical}$ is the critical energy density required to close the universe,

$$\rho_{\rm critical} = \frac{3H_0^2}{8\pi G} \tag{4.2.3}$$

where H_0 is the Hubble constant.

The characteristic strain spectrum, $h_c(f)$, is typically given a power-law dependence on frequency so that

$$h_c(f) = A\left(\frac{f}{\rm yr^{-1}}\right)^{\alpha}. \tag{4.2.4}$$

It may also be expressed in terms of the energy density of the background per unit logarithmic frequency, $\Omega_{\rm gw}(|f|)$:

$$h_c^2(f) = \frac{3H_0^2}{2\pi^2}\frac{1}{f^2}\Omega_{\rm gw}(|f|). \tag{4.2.5}$$

For an isotropic stochastic background of GWs, the signal appears in the data as correlated noise between measurements from different pulsars. The i^{th} data set is of the form

$$s_i(t) = z_i(t) + n_i(t) \tag{4.2.6}$$

where $z_i(t)$ corresponds to the unknown GW signal and $n_i(t)$ to noise (assumed in this case to be stationary and Gaussian). Because the signal is assumed to be much smaller

than the noise, the properties of the noise determine the variance. We can express these properties in the frequency domain as

$$\begin{aligned} \langle \tilde{n}_i(f) \rangle &= 0 \\ \langle \tilde{n}_i^*(f) \tilde{n_j}(f') \rangle &= \frac{1}{2} \delta(f - f') P_i(|f|) \end{aligned} \tag{4.2.7}$$

where we have introduced the one-sided noise power spectrum $P_i(|f|)$.

The cross-correlation statistic is defined as

$$S = \int_{-T/2}^{T/2} dt \int_{-T/2}^{T/2} dt' s_i(t) s_j(t') Q(t - t') \tag{4.2.8}$$

where $Q(t - t')$ is the filter function. The optimal filter is determined by maximizing the expected signal-to-noise ratio

$$\text{SNR} = \frac{\mu}{\sigma}. \tag{4.2.9}$$

Here μ is the mean $\langle S \rangle$ and σ is the square root of the variance $\sqrt{\langle S^2 \rangle - \langle S \rangle^2}$.

In the frequency domain (4.2.8) becomes

$$S = \int_{-\infty}^{\infty} df \int_{-\infty}^{\infty} df' \delta_T(f - f') \tilde{s}_i^*(f) \tilde{s}_j(f') \tilde{Q}(f'), \tag{4.2.10}$$

and the mean μ is

$$\mu = \int_{-\infty}^{\infty} df \int_{-\infty}^{\infty} df' \, \delta_T(f - f') \langle \tilde{z}_i^*(f) \tilde{z}_j(f') \rangle \tilde{Q}(f') \tag{4.2.11}$$

where δ_T is the finite time approximation to the delta function

$$\delta_T(f) = \frac{\sin \pi f t}{\pi f}.$$

The assumption that the background is unpolarized, isotropic, and stationary implies that the expectation value of the Fourier amplitudes $h_A(f, \hat{\Omega})$ must satisfy [46, 45]

$$\langle h_A^*(f, \hat{\Omega}) h_{A'}(f', \hat{\Omega}') \rangle = \frac{3H_0^2}{32\pi^3} \delta^2(\hat{\Omega}, \hat{\Omega}') \delta_{AA'} \times \delta(f - f') |f|^{-3} \Omega_{gw}(|f|) \tag{4.2.12}$$

where $\delta^2(\hat{\Omega}, \hat{\Omega}')$ is the covariant Dirac delta function on the two-sphere. With the demand (4.2.12) in place, the expectation value of the signals $z_i(f)$ may be written as

$$\langle \tilde{z}_i^*(f) \tilde{z}_j(f') \rangle = \frac{3H_0^2}{32\pi^3} \frac{1}{\beta} \delta(f - f') |f|^{-3} \times \Omega_{\text{gw}}(|f|) \Gamma(|f|).$$

Here β is a normalization factor and we define [45]

$$\Gamma(|f|) = \beta \sum_A \int_{S^2} d\Omega \, (e^{2\pi i f L_i (1+\hat{\Omega}\cdot\hat{p}_i)} - 1) \times (e^{-2\pi i f L_j (1+\hat{\Omega}\cdot\hat{p}_j)} - 1) F_i^A(\hat{\Omega}) F_j^A(\hat{\Omega})$$

where the sum is over all possible GW polarizations, and the exponential phase terms correspond to the pulsar term in the time domain.

The optimal filter is given by [46, 45]

$$\tilde{Q}(f) \propto \frac{\Omega_{\rm gw}(f)\Gamma(f)}{|f|^3 P_i(f) P_j(f)}, \tag{4.2.13}$$

where $P_i(f)$ and $P_j(f)$ are the power spectra for the ith and jth pulsar redshift time series that are being cross-correlated (see Eq. 4.2.8).

In general relativity, for the frequency and distance ranges appropriate to pulsar timing experiments (i.e. for $f \gg 1/L$), the overlap reduction function $\Gamma(f)$ approaches a constant which is only a function of the angular separation between the two pulsars. This constant is proportional to the value of the Hellings-Downs curve for the angle between the pulsars [71, 45]. We will see that for longitudinal modes and for tensor modes the overlap reduction function remains frequency dependent, even for $f \gg 1/L$, and is considerably larger than for the transverse modes. This indicates an increased sensitivity to such modes. To understand the physical origin of the increased sensitivity we first discuss the effect of GWs in the more simple case of a single pulsar-Earth baseline.

4.3 GW induced redshift on the pulsar-Earth system

In this section we will study the redshifts induced by GWs of different polarizations on the pulsar-Earth system. From (4.1.6), the redshift induced by this GW may be written as

$$\tilde{z}_A(f,\hat{\Omega}) = \left(e^{-2\pi i f L(1+\hat{\Omega}\cdot\hat{p})} - 1\right) \frac{p^i p^j}{2(1+\hat{\Omega}\cdot\hat{p})} \epsilon_{ij}^A(\hat{\Omega}) \tilde{h}_A. \tag{4.3.1}$$

The factor of $1/2(1+\hat{\Omega}\cdot\hat{p})$ comes from the relationship between the affine parameter λ and time t (see Eqs. (2.2.13) and (4.A.9)), and $\tilde{h}_A = \tilde{h}_A(f,\hat{\Omega})$.

In the region where the GW direction, $\hat{\Omega}$ and the pulsar direction, $\hat{p}$ are anti-parallel, Eq. (4.3.1) appears to become singular due to the $1+\hat{\Omega}\cdot\hat{p}$ term in the denominator (note that the derivative of h_A with respect to the affine parameter vanishes in this limit; see

Eq. (4.A.9)). There is in fact no divergence in the redshift induced. In this regime the exponential can be Taylor expanded and the $1 + \hat{\Omega} \cdot \hat{p}$ term in the denominator cancels.

A Taylor expansion of Eq. (4.3.1) can be performed in two cases. In the first, when $fL \ll 1$, the metric perturbation is the same at the pulsar and at the Earth. This case is often referred to as the long wavelength limit. In the second, when

$$1 + \hat{\Omega} \cdot \hat{p} \ll \frac{1}{fL},$$

the pulse's direction of propagation and the GW are nearly parallel (i.e. the GW is coming from a direction near the pulsar). In this case the metric perturbation at the pulsar when the pulse is emitted, and on Earth when the pulse is received, are also nearly the same. This is often described in the literature in terms of the pulse "surfing" the gravitational wave.

The surfing description, combined with Eq. (4.1.1), might lead one to incorrectly conclude that the effect of the GW should cancel in this case because the metric perturbations at the Earth and the pulsar are the same, despite the divergent $1/(1 + \hat{\Omega} \cdot \hat{p})$ term in the redshift. In fact, a delicate cancellation occurs with the divergent term in the denominator which is only manifest in the frequency domain. Let the pulse direction and the gravitational wave direction be nearly parallel so that $\hat{\Omega} \cdot \hat{p} = -1 + \delta$, where $\delta \ll 1$. Then as in [45, 125] we obtain

$$\tilde{z}_A(f, \hat{\Omega}) \sim -\pi i f L p^i p^j \epsilon_{ij}^A \tilde{h}_A. \tag{4.3.2}$$

The redshift is proportional to fL, but for finite δ increases only to the point where the argument of the exponential in (4.3.1) can no longer be Taylor expanded, at which point it becomes an oscillatory function of fL. Whether the redshift is finite in the $\delta \to 0$ limit depends on the projection term $p^i p^j \epsilon_{ij}^A h_A$. As we will see, the vanishing contribution for the tensor modes of general relativity occurs solely because of the transverse nature of these waves, and is unrelated to the "surfing" effect. For longitudinal modes the projection term does not vanish, and the increase in sensitivity to such modes originates from GWs that come from directions near the pulsar. To better understand this, we will look at the behavior of the redshifts induced by GWs of various modes.

The redshift for a longitudinal mode GW perturbation is

$$\tilde{z}_l(f,\hat{\Omega}) = \frac{\cos^2\theta}{2(1+\cos\theta)}(e^{-2\pi i f L(1+\cos\theta)}-1)\tilde{h}_l, \tag{4.3.3}$$

while the redshift for a plus mode GW perturbation is

$$\tilde{z}_+(f,\hat{\Omega}) = \frac{-\sin^2\theta}{2(1+\cos\theta)}(e^{-2\pi i f L(1+\cos\theta)}-1)\tilde{h}_+. \tag{4.3.4}$$

Here we note that the geometrical factor in the redshift for the transverse breathing mode differs from (4.3.4) only by a sign, and our analysis of (4.3.4) applies equally to the breathing mode.

In Fig. 16 we plot the geometrical and phase factor $|\tilde{z}(f,\hat{\Omega})/\tilde{h}|$ for both the +-mode and the longitudinal mode. We plot these for a value of fL in the long wavelength limit ($fL = 10^{-2}$), and for a value in the regime of pulsar timing experiments ($fL = 10$). In the regime of pulsar timing experiments the sensitivity is largest for GW directions near the pulsar $\theta \approx \pi$ for *both* polarizations. Although we do not show it here the same is true for all other polarization modes. In the long wavelength limit, $fL \ll 1$, the pulsar-Earth system is most sensitive to +-mode GWs coming from the equator, and longitudinal GWs from the poles.

As discussed above, these redshifts appear to become singular when $\theta \to \pi$, but the pulsar term may be Taylor expanded. Let $\theta = \pi - \delta$, where $\delta \ll 1$. Then

$$\tilde{z}_l(f,\hat{\Omega}) \sim \pi i f L(1-\delta^2)\tilde{h}_l \tag{4.3.5}$$

for the longitudinal case, while

$$\tilde{z}_+(f,\hat{\Omega}) \sim \pi i f L \delta^2 \tilde{h}_+ \tag{4.3.6}$$

for the plus mode. In the limit as $\delta \to 0$, $\tilde{z}_+$ vanishes while $\tilde{z}_l$ becomes proportional to fL. The vanishing redshift of $\tilde{z}_+$ is therefore due to the transverse nature of the mode, and does not occur for $\tilde{z}_l$, even though in both cases the pulse is "surfing" the GW. In the time domain, in the $\theta \approx \pi$ region, the redshift for both modes goes as

$$z_{l,+}(t,\hat{\Omega}) \propto L\dot{h}_{l,+}. \tag{4.3.7}$$

One may readily identify the right hand side of (4.3.7) as a velocity. The interpretation of this result is that, in this limit, the redshift is proportional to the relative velocity of

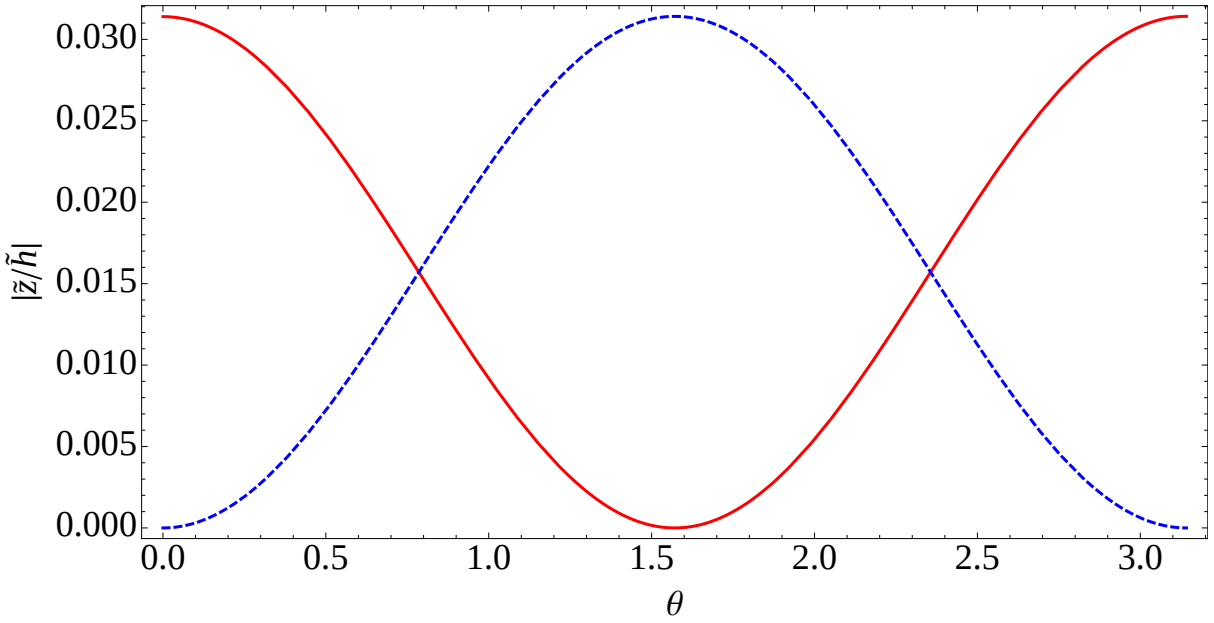

(a)

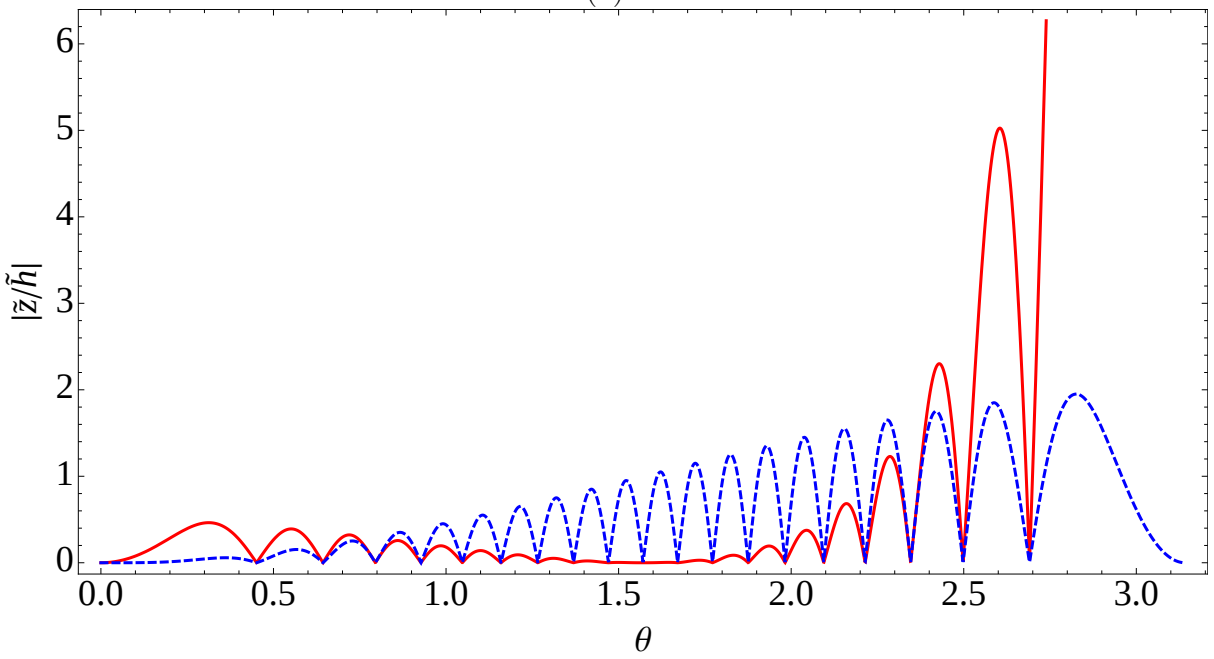

(b)

Figure 16: (color online) Plots of $|\tilde{z}(f,\hat{\Omega})/\tilde{h}|$ for the +-mode (dashed blue) and the longitudinal mode (solid red). We show these for $fL = 10^{-2}$ (a), a value of fL in the long wavelength limit, and (b) $fL = 10$, a value of fL typical of pulsar timing experiments. In the regime of pulsar timing experiments the sensitivity is largest for GW directions near the pulsar $\theta \approx \pi$ for both polarizations. In the longth wavelength limit, $fL \ll 1$, the pulsar-Earth system is most sensitive to +-mode GWs coming from the equator, and longitudinal GWs from the poles.

the pulsar-Earth system. The velocity of the pulsar when the pulse is emitted in this limit is approximately equal and opposite to the velocity of the Earth when the pulse is received.

An identical analysis for the shear GW modes produces analogous results. Starting from (4.1.6), the redshift for the vector-y mode goes as

$$\tilde{z}_y(f,\hat{\Omega}) = -\frac{\cos\theta\sin\theta}{(1+\cos\theta)}(e^{-2\pi i f L(1+\cos\theta)}-1)h_y. \tag{4.3.8}$$

The small δ expansion yields

$$\tilde{z}_y(f,\hat{\Omega}) \sim -2\pi i f L\delta\left(1-\frac{\delta^2}{2}\right)h_y. \tag{4.3.9}$$

Relative to the longitudinal mode the redshift of vector modes is smaller by a factor of δ and vanishes as $\delta \to 0$, but it is still larger than the transverse modes by a factor of $1/\delta$.

The same behavior is not present in other sky locations. If the GW propagates perpendicular to the pulsar-Earth line ($\theta = \pi/2 + \delta$), then up to second order in δ the redshifts

$$\tilde{z}_l = \frac{\delta^2}{2(1-\delta)}\left(e^{-2\pi i f L(1-\delta)}-1\right) \text{ (longitudinal)} \tag{4.3.10}$$

$$\tilde{z}_+ = \frac{-(1-\delta^2)}{2(1-\delta)}\left(e^{-2\pi i f L(1-\delta)}-1\right) \text{ (plus)} \tag{4.3.11}$$

$$\tilde{z}_y = \frac{\delta\,(1-\delta^2/2)}{(1-\delta)}\left(e^{-2\pi i f L(1-\delta)}-1\right) \text{ (shear)} \tag{4.3.12}$$

are obtained. In this case for small δ the exponential cannot be expanded unless $fL \ll 1$. For this sky location the redshift is always an oscillatory function of fL. The pulse comes across different phases of the GW as it propagates toward Earth.

To summarize, one can see that the surfing effect does not lead to a vanishing response of the pulsar-Earth system to GW waves coming from $\theta = \pi$. For the tensor and scalar-breathing modes, it is the transverse nature of GWs that is responsible for the vanishing response. For the scalar-longitudinal modes the response does not vanish—in fact, the response increases with both frequency and pulsar distance. For the vector modes the response does vanish, but more slowly than for the transverse modes. For *all* GW modes from directions near $\theta = \pi$, the redshift increases monotonically up to some limiting

frequency beyond which the Taylor series expansion of the pulsar term which leads to Eqs. (4.3.5) and (4.3.6) can no longer be performed.

We now discuss the implications of this effect on the overlap reduction functions.

4.4 Overlap reduction functions

As discussed in Section 4.2, the overlap reduction function for the two pulsars in Fig. 13 is equal to

$$\Gamma_A(|f|) = \frac{3}{4\pi} \sum_A \int_{S^2} d\Omega (e^{2\pi i f L_1 (1+\hat{\Omega}\cdot\hat{p}_1)} - 1)(e^{-2\pi i f L_2 (1+\hat{\Omega}\cdot\hat{p}_2)} - 1) F_1^A(\hat{\Omega}) F_2^A(\hat{\Omega}) \quad (4.4.1)$$

where all possible GW polarizations are allowed with equal amplitudes. It is advantageous to consider each term in the sum (4.4.1) separately since various gravity theories may have different polarization content [114, 121, 110, 111, 112, 113, 115, 116, 117, 118, 119, 120].

The overlap reduction function has a closed analytic form for transverse GWs. The overlap reduction function for the plus mode has been calculated by [71] and is given by

$$\Gamma_+(\xi) = 3\left[\frac{1}{3} + \frac{1-\cos\xi}{2}\left[\log\left(\frac{1-\cos\xi}{2}\right) - \frac{1}{6}\right]\right], \quad (4.4.2)$$

where ξ is the angular separation of the pulsars. For the scalar-breathing mode, a closed form is given by [122]:

$$\Gamma_b(\xi) = \frac{1}{4}(3 + \cos\xi)\,. \quad (4.4.3)$$

For the case of non-transverse GWs, the overlap reduction functions cannot be integrated analytically and we calculate them numerically.

In general relativity the pulsar term can be excluded from the integral in Eq.(4.4.1) without any significant loss of optimality [45]. The reason for this is that the smallest frequencies that PTAs are sensitive to are $\sim 0.1\ \mathrm{yr}^{-1}$, and the closest PTA pulsar distances are $\sim$ 100 ly, so that $fL \gtrsim 10$. This is shown in Fig. 16, where we plot the overlap reduction functions $\Gamma(fL)$ with (solid curves) and without (horizontal dashed lines) the pulsar term for several pulsar separation angles ξ and GW polarization modes. The frequencies that PTAs are sensitive to are to the right of the vertical dashed line at $fL = 10$ in each plot. As seen in Fig. 16(a), $\Gamma_+(fL)$ is roughly independent of frequency

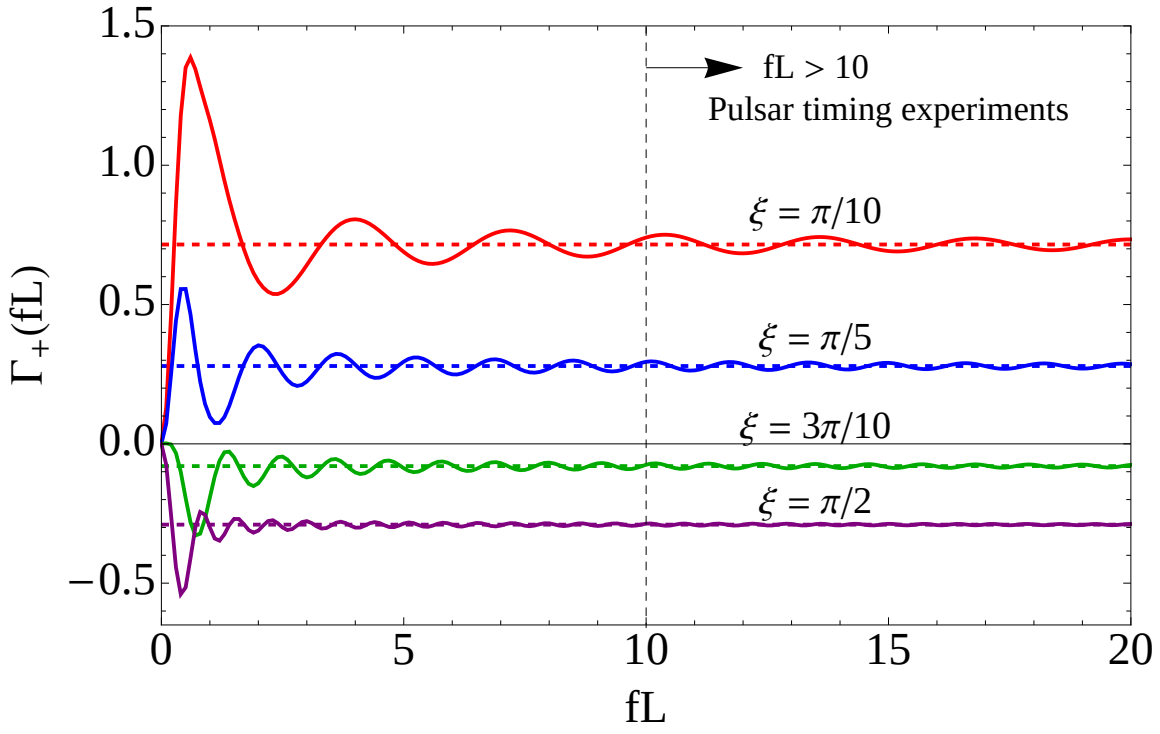
1.5
1.0
0.5
0.0
−0.5
0
5
10
15
20
fL
$\Gamma_+(fL)$
fL > 10
Pulsar timing experiments
$\xi = \pi/10$
$\xi = \pi/5$
$\xi = 3\pi/10$
$\xi = \pi/2$

(a)

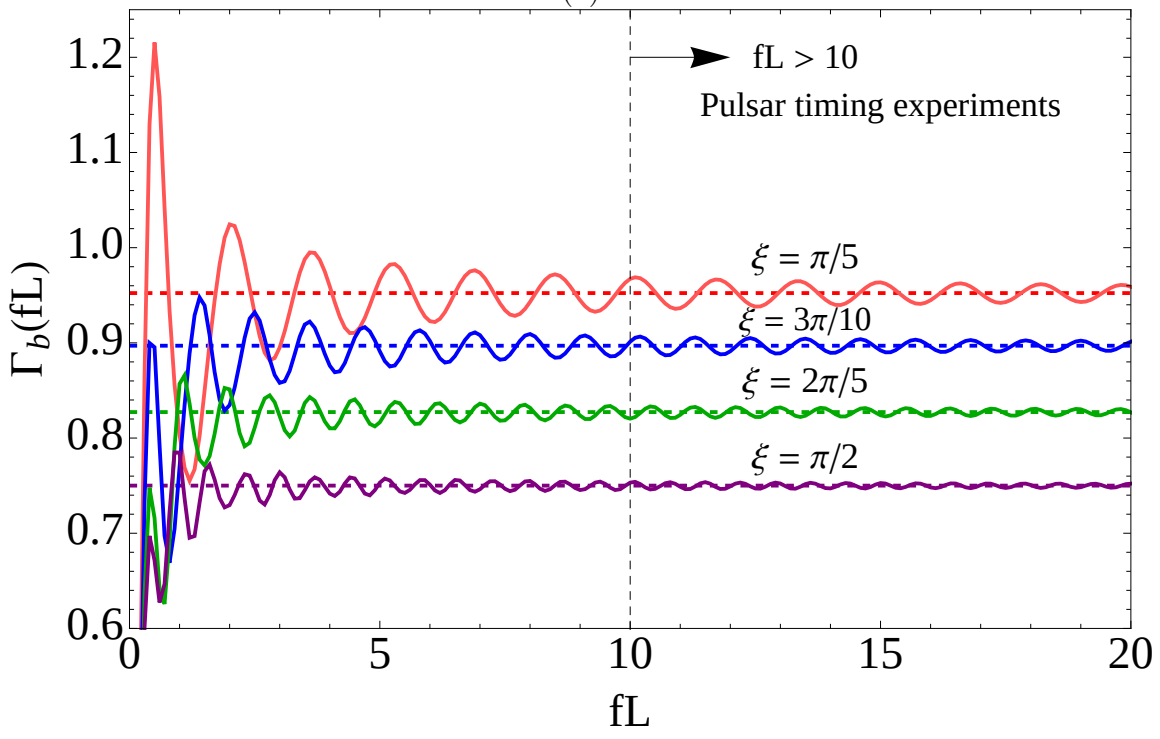
1.2
1.1
1.0
0.9
0.8
0.7
0.6
0
5
10
15
20
fL
$\Gamma_b(fL)$
fL > 10
Pulsar timing experiments
$\xi = \pi/5$
$\xi = 3\pi/10$
$\xi = 2\pi/5$
$\xi = \pi/2$

(b)

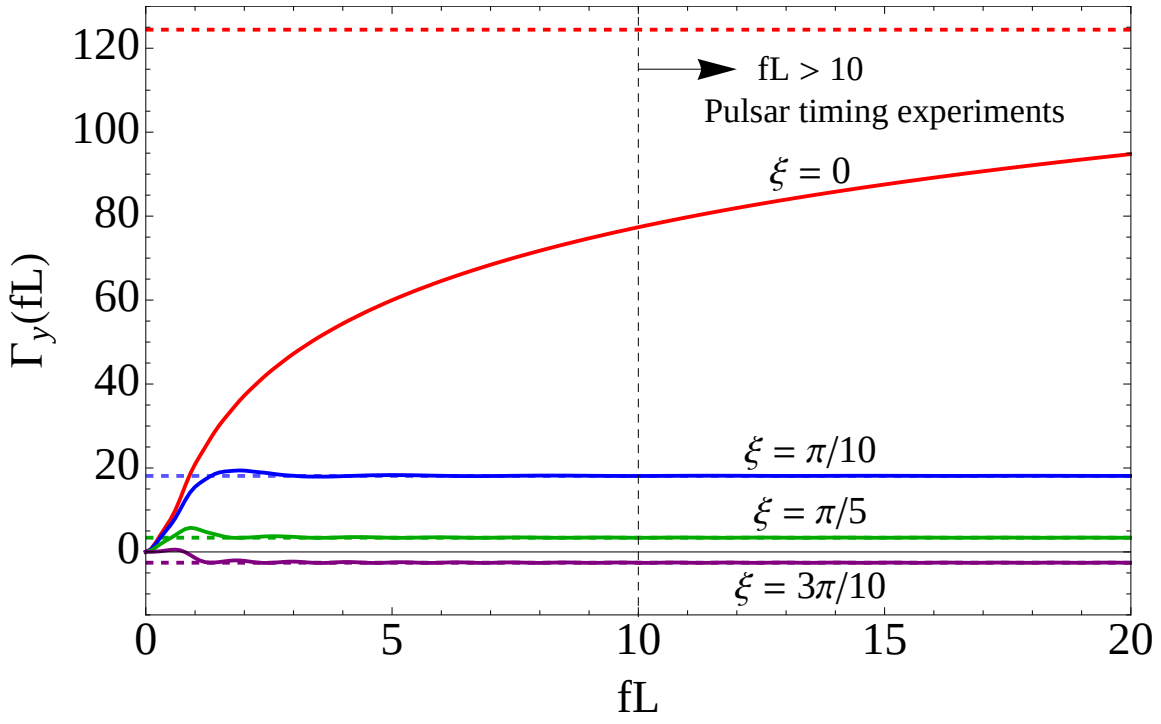

(c)

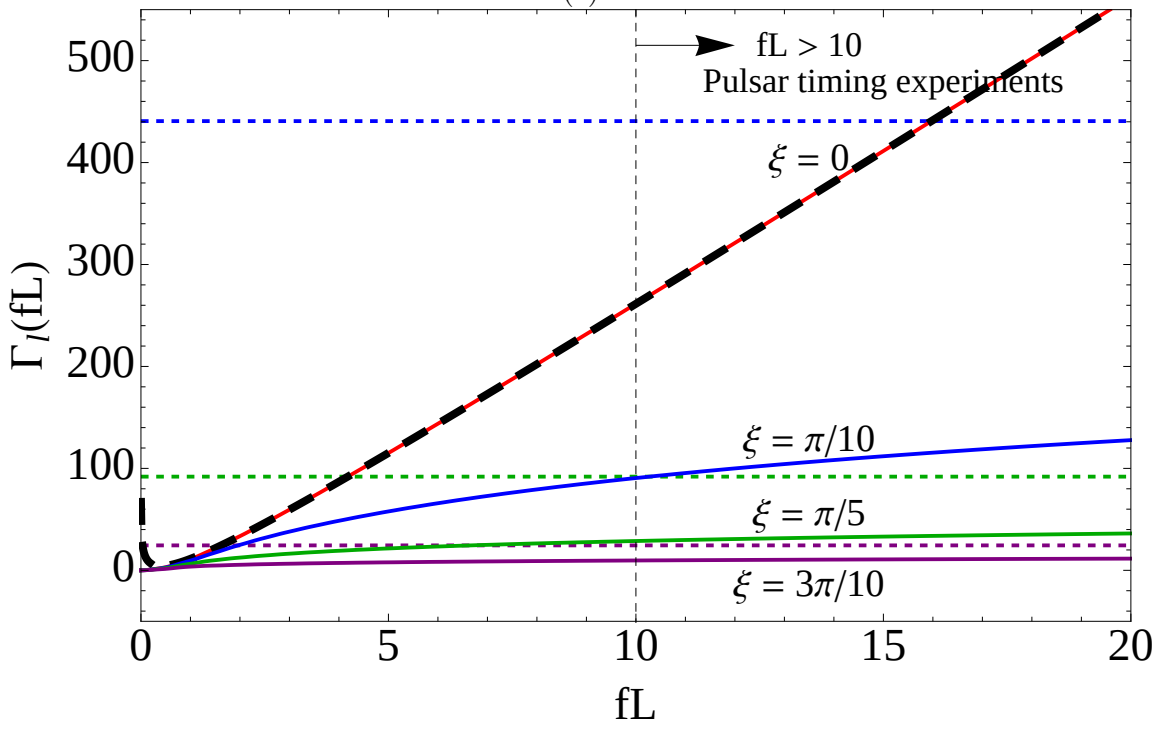

(d)

Figure 16: $\Gamma(fL)$ with (solid curves) and without (horizontal dashed lines) the pulsar term for the various polarization modes: plus (a), breathing (b), shear (c) and longitudinal (d). In the latter two modes, smaller pulsar separation angles are characterized by retained frequency dependence in $\Gamma(fL)$ in the range of frequencies relevant to pulsar timing experiments. Nearly all the non-transverse curves eventually converge, but at rather high values of $\Gamma(fL)$ relative to the transverse modes, indicating increased sensitivity to GWs with these polarizations. We have plotted the large limit approximation (4.4.5) (dashed black curve) along with $\Gamma_l(fL)$ in (d), which is in good agreement with the $\xi = 0$ curve.

over the range of frequencies relevant to pulsar timing experiments. The same is true for the scalar-breathing mode, which is shown in Fig. 16(b). It is worth pointing out that both $\Gamma_+(fL)$ and $\Gamma_b(fL)$ are normalized to unity for co-aligned pulsars. Note that the overlap reduction functions for all other modes are normalized with the same factor of $3/4\pi$ used in the +-mode.

In Fig. 16(c), we plot the overlap reduction function $\Gamma_y(fL)$ for the vector-y mode. Over the range of relevant frequencies, $\Gamma_y(fL)$ is frequency independent for most of the pulsar separation angles shown. For co-aligned pulsars, however, $\Gamma_y(fL)$ retains frequency dependence well into the range of pulsar timing frequencies, and takes on values an order of magnitude higher than those obtained by $\Gamma_+(fL)$ and $\Gamma_b(fL)$.

Similar behavior is shown in Fig. 16(d), where we have plotted the overlap reduction function for the scalar-longitudinal mode. Here $\Gamma_l(fL)$ retains frequency dependence throughout the relevant frequency range for each of the pulsar separation angles shown. For the case of co-aligned pulsars, $\Gamma_l(fL)$ diverges as $fL \to \infty$, and for separation angles that do converge $\Gamma_l(fL)$ takes on values that are at least an order of magnitude larger than those obtained by $\Gamma_+(fL)$ and $\Gamma_b(fL)$.

For co-located pulsars we can understand the behavior of the longitudinal mode analytically. In the problematic sky region ($\theta \approx \pi$), $\Gamma_l(fL)$ is proportional to the square of the redshift,

$$\Gamma_l(fL) \propto 2\pi \int_{-1}^{1} \left|\left(e^{-2\pi i fL(1+\cos\theta)} - 1\right)\right|^2 \frac{\cos^4\theta}{4(1+\cos\theta)^2} d(\cos\theta) \tag{4.4.4}$$

which may be evaluated analytically. In the limit of large fL ($fL \gg 1$),

$$\begin{aligned}\Gamma_l(fL) &= \pi\left\{37/6 - 4\gamma - 1/(\pi(fL)^2) + 4\ \mathrm{Ci}(4\pi fL) - 4\log(4\pi fL) + 2\pi fL\ \mathrm{Si}(4\pi fL)\right\} \\ &\sim (37/6 - 4\gamma)\,\pi - 4\pi\log(4\pi fL) + \pi^3 fL,\end{aligned} \tag{4.4.5}$$

where γ is Euler's constant. The overlap reduction function $\Gamma_l(fL)$ is roughly proportional to fL in this limit. Eq. (4.4.5) is shown along with the numerically integrated overlap reduction functions in Fig. 16(d) and, with the exception of the singular behavior near the origin (where the large fL approximation is not valid), agrees well with the numerical $\Gamma_l(fL)$ curve for co-aligned pulsars ($\xi = 0$).

PSR	Distance (kpc)	PSR	Distance (kpc)
J0030+0451	0.23	J1853+1303	1.60
J0218+4232	5.85	J1857+0943	0.70
J0613−0200	2.19	J1903+0327	6.45
J1012+5307	0.52	J1909−3744	0.55
J1024−0719	0.35	J1910+1256	1.95
J1455−3330	0.74	J1918−0642	1.40
J1600−3053	2.67	J1939+2134	3.58
J1640+2224	1.19	J1944+0907	1.28
J1643−1224	4.86	J1955+2908	5.39
J1713+0747	0.89	J2010−1323	1.29
J1738+0333	1.97	J2145−0750	0.50
J1744−1134	0.17	J2317+1439	1.89

Table 1: NANOGrav Pulsar Data

4.5 Overlap reduction functions for the NANOGrav pulsars

The NANOGrav PTA consists of 24 pulsars. The Australia Telescope National Facility (ATNF) data for the distances to these pulsars is given in Table 1 [126]. Using a simple numerical integration scheme, the overlap reduction function for each pulsar pair was computed. The main difference relative to the previous section is that we are including the effect of different pulsar distances. Results are given in Fig. 17 (a)–(d) and show that the calculated values of $\Gamma(f)$ are consistent with the more simple results discussed in Section 4.4 for the non-transverse modes for frequencies up to $\sim 10^{-9}$ Hz. Pulsar pairs with the smallest ($\xi \lesssim 12^{\circ}$) separation angles (starred curves in Fig. 17 (b), (d)) for non-transverse polarization modes are characterized by large values of the overlap reduction function and monotonic growth up to some limiting frequency. Pulsar pairs with larger ($\xi \gtrsim 12^{\circ}$) separation angles (un-starred curves in Fig. 17 (b), (d) and all curves in Fig. 17) do not display monotonic growth up to a limiting frequency, but still result in much larger values than those of the plus and cross modes. Fig. 17 shows that sensitivity is greater

for scalar-longitudinal and vector modes than for the tensor and scalar-breathing modes, and increases rapidly for pulsars that are nearly co-aligned in the sky.

Over the entire range of frequencies plotted for pulsar timing experiments (between $\sim 10^{-9}$ and $\sim 10^{-7}$ Hz), the overlap reduction functions are approximately constant. In practice, some optimality will be lost due to the fact that pulsar distances are known at best to only $\sim 10\%$ [127].

4.6 Discussion

Direct detection of GWs might be possible in the next decade using a pulsar timing array. A detection would provide a mechanism for testing various metric theories of gravity. To develop optimal detection strategies for stochastic backgrounds in alternative theories of gravity, we have computed overlap reduction functions for all six GW polarization modes, including four modes not present in general relativity.

We began by introducing the redshift induced by GWs of various polarizations, along with the polarization tensors unique to each mode. We then used the optimal detection statistic for an unpolarized, isotropic stochastic background of GWs, defined in Anholm et al. [45], to find the overlap reduction function, a geometric dependent quantity in the expression for the expected cross correlation.

We examined the redshifts induced by GWs of various polarizations on the pulsar-Earth system, and find that our results are consistent with those of Anholm et al. [45] and Tinto and Alves [125]: when the GWs are coming from roughly the same direction as the pulses from the pulsar, the induced redshift for *any* GW polarization mode is proportional to fL, the product of the GW frequency and the distance to the pulsar. When the GWs and the pulse direction are exactly parallel the redshift for the transverse and vector modes vanishes, but it is proportional to fL for the scalar-longitudinal mode.

We show that the vanishing contributions from the tensor, vector and scalar-breathing modes are not a result of the pulse surfing the GW. In fact, sensitivity to GWs coming from directions near the pulsar increases for *all* polarizations. It is the transverse nature of these modes that is responsible for the vanishing response. In this limit we also show that the redshift is proportional to the relative velocity of the pulsar-Earth system ($L\dot{h}$),

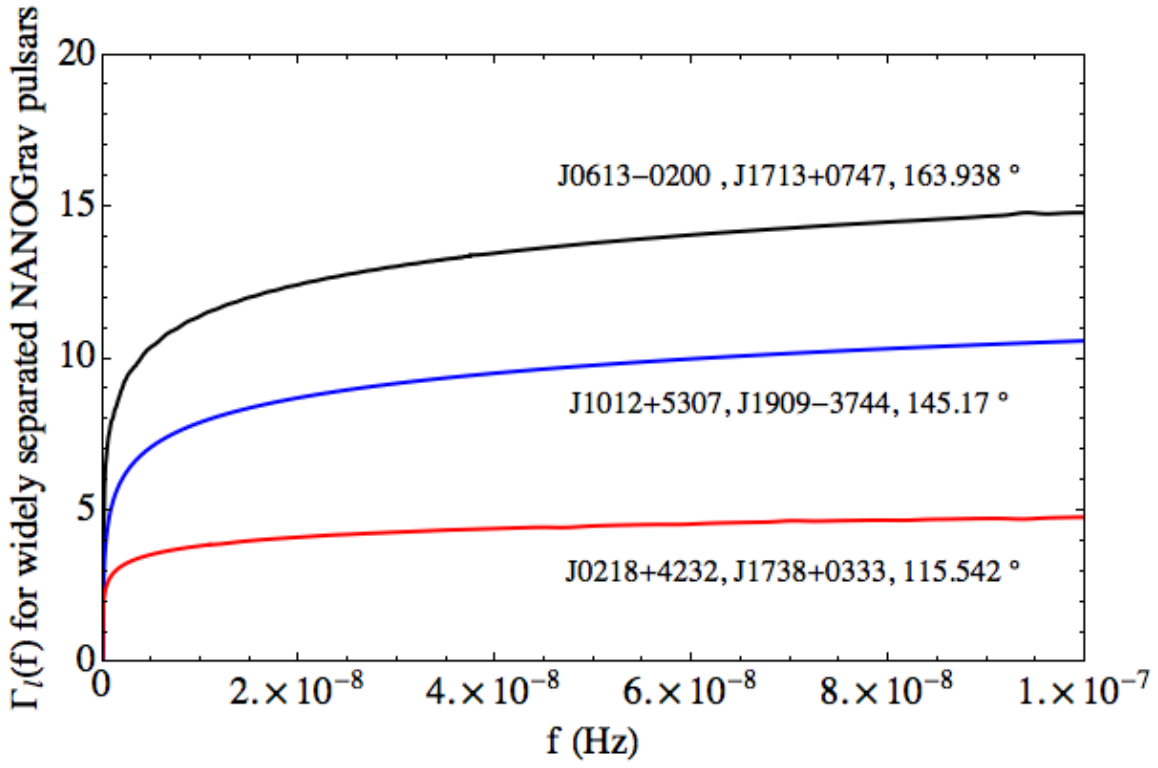
Γ_I(f) for widely separated NANOGrav pulsars
20
15
10
5
0
J0613−0200 , J1713+0747, 163.938 °
J1012+5307, J1909−3744, 145.17 °
J0218+4232, J1738+0333, 115.542 °
0
2.×10^−8
4.×10^−8
6.×10^−8
8.×10^−8
1.×10^−7
f (Hz)

(a)

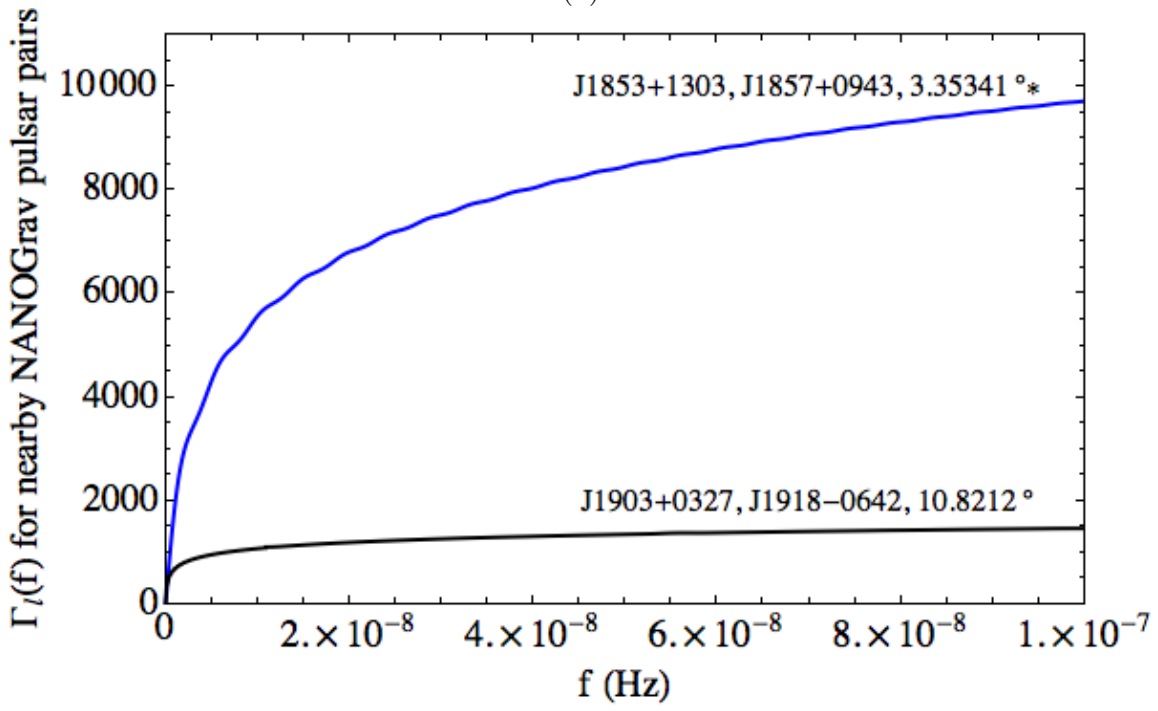
Γ_I(f) for nearby NANOGrav pulsar pairs
10000
8000
6000
4000
2000
0
J1853+1303, J1857+0943, 3.35341 °*
J1903+0327, J1918−0642, 10.8212 °
0
2.×10^−8
4.×10^−8
6.×10^−8
8.×10^−8
1.×10^−7
f (Hz)

(b)

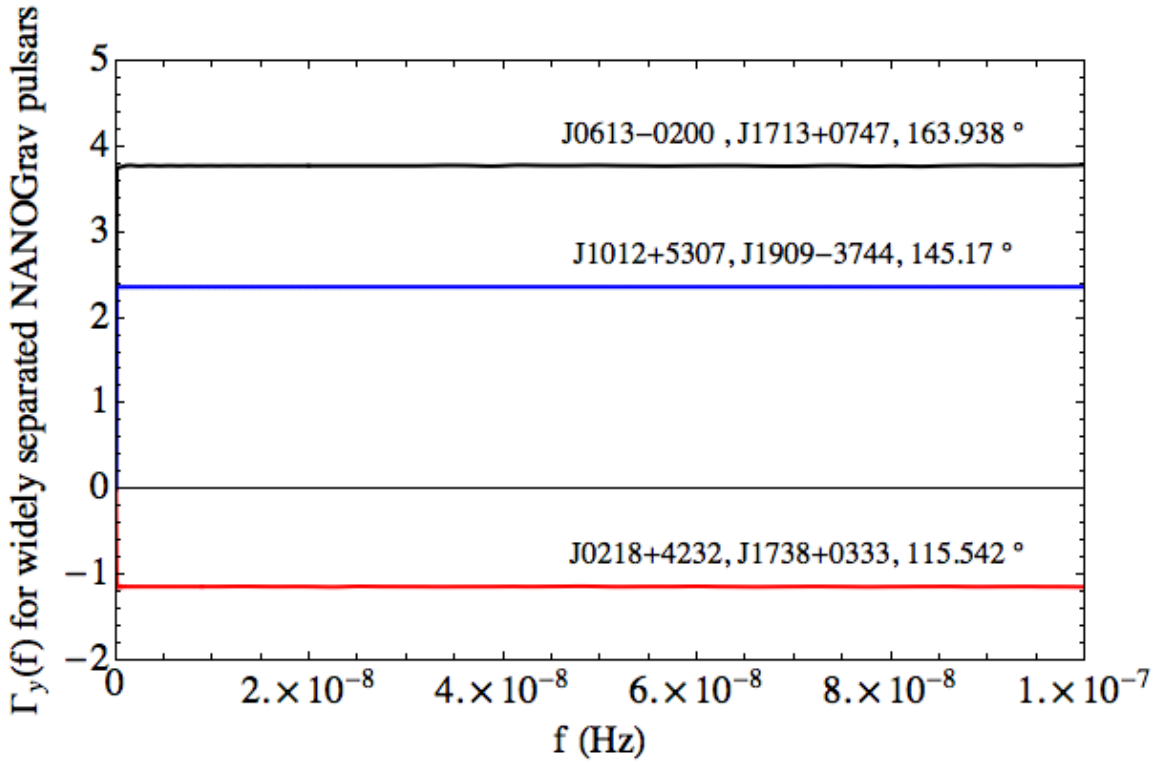

(c)

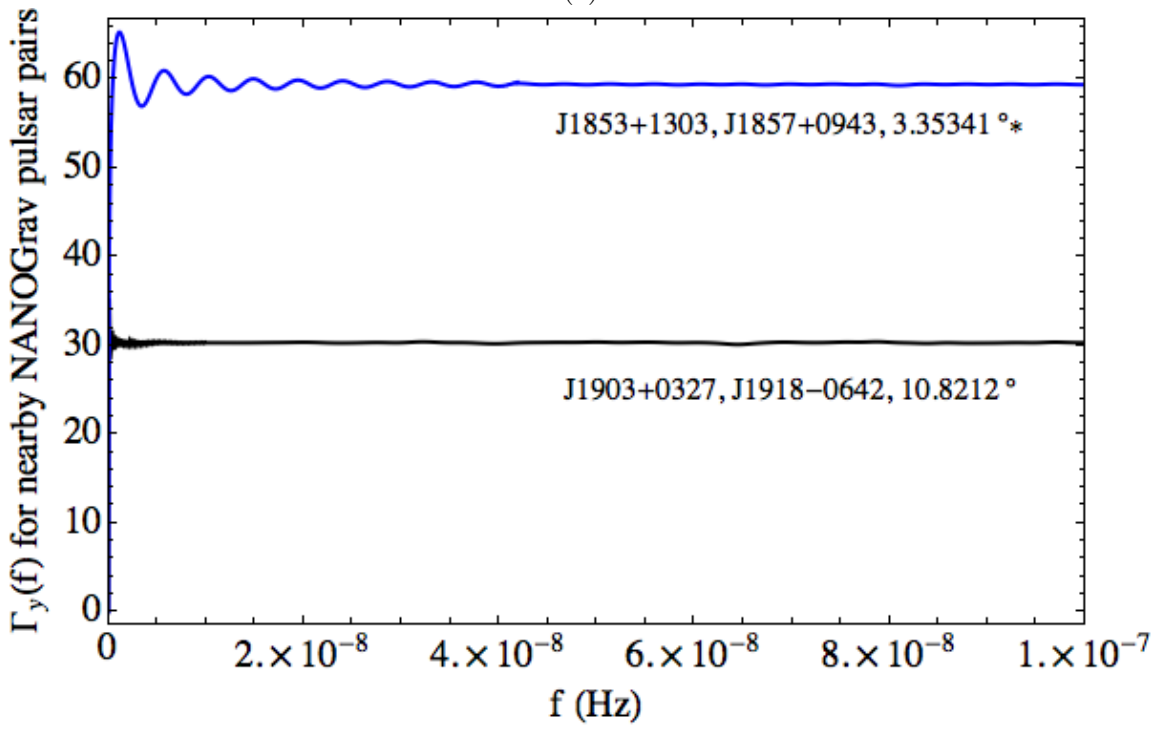

(d)

Figure 17: $\Gamma(f)$ for some of the NANOGrav pulsar pairs. Pulsar pairs, along with their angular separation in degrees, are shown with each curve. As f increases, $\Gamma(f)$ approaches a constant value. The asterisk indicates the NANOGrav pulsar pair with the smallest angular separation (~ 3.35 degrees). Note the larger values of the $\Gamma(f)$s for this pair.

which is the same when the pulse is emitted and when it is received.

We find that the overlap reduction functions for non-transverse GWs are characterized by frequency dependence that is significant for nearby pulsar pairs. The values of the overlap reduction function increase by up to one order of magnitude for the vector polarization modes and up to two orders of magnitude for the scalar-longitudinal mode. Pulsar timing arrays are significantly more sensitive to scalar-longitudinal and vector GW stochastic backgrounds.

Next, we used current pulsar distance and sky-location data from the ATNF pulsar catalog to calculate the overlap reduction functions for each pulsar pair in the NANOGrav pulsar timing array. Over the range of frequencies relevant to pulsar timing array experiments, these overlap reduction functions for all polarization modes are roughly constant for most pulsar pairs. For nearly co-aligned pulsars, the overlap reduction functions for scalar-longitudinal and vector modes exhibit marked frequency dependence and asymptote to much larger values than the overlap reduction functions for transverse modes. In fact for a pair separated but about 3° we find a sensitivity increase of about a factor of 10^4 for longitudinal modes.

The results discussed here may be compared to other recent work. Lee et al. [122] calculated the cross-correlation functions for stochastic GW backgrounds including all six GW polarizations, and found that the correlation functions for non-transverse GWs are frequency dependent, as well as an increased response in the cross-correlation to scalar-longitudinal GWs, in agreement with our results. This work was done in the context of the coherence statistic [122] for stochastic background detection, rather than the optimal statistic [45]. The coherence statistic is a measure of goodness of fit of the pulsar-pair cross-correlations to the Hellings-Downs curve. For non-transverse modes there is no Hellings-Downs curve because the overlap reduction functions remain frequency dependent for large fL. Lee et al. solved this problem by simulating GW backgrounds and finding effective background-dependent Hellings-Downs curves for these theories. In the context of the optimal statistic this is a non-issue: The frequency dependent overlap reduction functions can be used to construct the optimal filter in Eq. (4.2.13). This is identical to what is done for LIGO stochastic background optimal filter construction [46],

where the overlap reduction functions are also frequency dependent.

Alves and Tinto [123] have estimated antenna sensitivities to GWs of all six polarization modes by assuming a signal-to-noise ratio of 1 is achieved over 10 years of data and calculating the noise spectrum. Their results indicate an increase of two to three orders of magnitude in sensitivity to scalar-longitudinal mode GWs compared to that of plus and cross mode GWs. To explain this effect Alves and Tinto compare the effect of a tensor GW propagating orthogonally to the pulsar-Earth system, and a scalar-longitudinal GW propagating in a direction parallel to the pulse direction. They argue that the increased sensitivity to longitudinal GWs is due to the amount of time a longitudinal GW affects the pulsar-Earth radio link.

We have compared the effect of GW propagation from directions near the pulsar and orthogonal to the pulsar-Earth system for all polarization modes. For GW propagation directions parallel to the pulse direction we find that the redshift induced by a gravitational wave is large, and seemingly divergent when the GW and pulse directions are exactly parallel. This apparent divergence occurs for longitudinal, transverse, and shear modes alike. In that limit, however, the divergent term in the redshift that comes from the relationship between time and affine parameter derivatives cancels because the phase of the GW pulse when pulse is emitted is nearly equal to the phase of the GW when the pulse is received (see Eqs. (4.A.9), (4.3.1) and (4.3.2)). The redshift becomes proportional to the relative velocity of the pulsar-Earth system and a mode-dependent geometrical projection factor for *all* GW polarization modes. In this limit the relative velocity of the pulsar-Earth system is approximately equal when the pulse is emitted and received. For transverse and shear modes the projection factor vanishes when the GW and pulse directions become parallel. For longitudinal modes the geometrical factor goes to a constant, so that the pulsar-Earth system is very sensitive to GWs from directions near the pulsar. This is the physical origin of the increased sensitivity to scalar-longitudinal GWs.

Appendix 4.A Analog to Detweiler's equation for vector and scalar polarization modes

Here we show the derivation of the redshift induced by non-Einsteinian GW modes. We begin by considering the metric due to a longitudinal mode gravitational wave perturbation:

$$\begin{aligned} g_{ab} &= \eta_{ab} + h_{ab}(t-z) \\ &= \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1+h_L \end{pmatrix}. \end{aligned} \tag{4.A.1}$$

Given a null vector $s^a = \nu(1, -\alpha, -\beta, -\gamma)$ in Minkowski space (where α, β, γ are directional cosines) the corresponding perturbed null vector is given by

$$\begin{aligned} \sigma^a &= s^a - \frac{1}{2}\eta^{ab}h_{bc}s^c \\ &= \nu \begin{pmatrix} 1 \\ -\alpha \\ -\beta \\ -\gamma(1-\frac{h_L}{2}) \end{pmatrix}. \end{aligned} \tag{4.A.2}$$

From the geodesic equation, the t-component of σ^a must satisfy

$$\frac{d\sigma^t}{d\lambda} = -\Gamma^t_{ab}\sigma^a\sigma^b \tag{4.A.3}$$

where

$$\begin{aligned} \Gamma^t_{ab} &= \frac{1}{2}g^{tc}\left(\partial_a g_{bc} + \partial_b g_{ac} - \partial_c g_{ab}\right) \\ &= \frac{1}{2}\dot{g}_{ab}. \end{aligned} \tag{4.A.4}$$

Now we may write the geodesic equation as

$$\begin{aligned} \frac{d\sigma^t}{d\lambda} &= -\frac{1}{2}\dot{g}_{ab}\sigma^a\sigma^b \\ &= -\frac{1}{2}\dot{h}_L(\sigma^z)^2. \end{aligned} \tag{4.A.5}$$

To zeroth order in h_L,

$$\begin{aligned}(\sigma^z)^2 &= \nu^2\gamma^2\left(1+\frac{h_L}{2}\right)^2 \\ &\approx \nu^2\gamma^2 + O(h_L) \end{aligned} \tag{4.A.6}$$

allowing us to write the geodesic equation as

$$\frac{d\sigma^t}{d\lambda} = \frac{d\nu}{d\lambda} = -\frac{1}{2}\dot{h_L}\nu^2\gamma^2. \tag{4.A.7}$$

We now need to express the time derivative of the metric perturbation, $\dot{h_L}$, as a derivative of the affine parameter λ. Since $h_L = h_L(t-z)$, we may write

$$\begin{aligned}\frac{dh_L}{d\lambda} &= \frac{\partial h_L}{\partial t}\frac{dt}{d\lambda} + \frac{\partial h_L}{\partial z}\frac{dz}{d\lambda} \\ &= \frac{\partial h_L}{\partial t}\frac{dt}{d\lambda} - \frac{\partial h_L}{\partial t}\frac{dz}{d\lambda}. \end{aligned} \tag{4.A.8}$$

Identifying the relations $\frac{dt}{d\lambda} = \nu$ and $\frac{dz}{d\lambda} = -\nu\gamma$, we obtain the relation

$$\dot{h_L} = \frac{\partial h_L}{\partial t} = \frac{1}{\nu(1+\gamma)}\frac{dh_L}{d\lambda} \tag{4.A.9}$$

which makes the geodesic equation

$$\frac{d\nu}{d\lambda} = -\frac{1}{2}\dot{h_L}\nu^2\gamma^2 = -\frac{1}{2}\frac{\nu\gamma^2}{(1+\gamma)}\frac{dh_L}{d\lambda} \tag{4.A.10}$$

Integrating both sides, we obtain

$$\frac{\nu_e}{\nu_p} = \exp\left(-\frac{1}{2}\frac{\gamma^2}{(1+\gamma)}\Delta h_L\right) \tag{4.A.11}$$

where $\Delta h_l = h_l^e - h_l^p$. Expanding to first order in h_L, we may write

$$\frac{\nu_e - \nu_p}{\nu_p} \approx -\frac{1}{2}\frac{\gamma^2}{(1+\gamma)}\Delta h_L \tag{4.A.12}$$

$$= -\frac{\cos^2\theta}{2\,(1+\cos\theta)}\Delta h_L. \tag{4.A.13}$$

The derivation for vector modes is nearly identical to that of the longitudinal mode. For the sake of brevity we only detail the vector-y mode in the remainder of this document. For the vector-y mode, the metric perturbation takes the form

$$g_{ab} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & h_y \\ 0 & 0 & h_y & 1 \end{pmatrix}. \tag{4.A.14}$$

The null vector becomes

$$\sigma^a = \nu \begin{pmatrix} 1 \\ -\alpha \\ -\beta + \frac{h_y \gamma}{2} \\ \frac{h_y \beta}{2} - \gamma \end{pmatrix}. \tag{4.A.15}$$

Following the same algebraic steps used above, one obtains the geodesic equation

$$\frac{d\sigma^t}{d\lambda} = \frac{d\nu}{d\lambda} = -\dot{h_y}\nu^2\gamma\beta, \tag{4.A.16}$$

which leads to

$$\frac{d\nu}{d\lambda} = -\frac{\nu\gamma\beta}{(1+\gamma)}\frac{dh_y}{d\lambda}. \tag{4.A.17}$$

Integrating this expression and expanding the result to first order in Δh_y produces the result

$$\frac{\nu_e - \nu_p}{\nu_p} \approx -\frac{\beta\gamma}{(1+\gamma)}\Delta h_y \tag{4.A.18}$$

$$= -\frac{\sin 2\theta \sin\phi_p}{2\,(1+\cos\theta)}\Delta h_y. \tag{4.A.19}$$

where $\Delta h_y = h_y^e - h_y^p$.

For comparison, we also include the results for the plus, cross, vector-x, and breathing modes. For the plus mode, we obtain

$$\frac{\nu_e - \nu_p}{\nu_p} \approx -\frac{(\alpha^2 - \beta^2)}{2(1+\gamma)}\Delta h_+ \tag{4.A.20}$$

$$= -\frac{\sin^2\theta \cos 2\phi_p}{2\,(1+\cos\theta)}\Delta h_+; \tag{4.A.21}$$

for the cross mode,

$$\frac{\nu_e - \nu_p}{\nu_p} \approx -\frac{\alpha\beta}{(1+\gamma)}\Delta h_\times \tag{4.A.22}$$

$$= -\frac{\sin^2\theta \sin 2\phi_p}{2\,(1+\cos\theta)}\Delta h_\times; \tag{4.A.23}$$

for the vector-x mode,

$$\frac{\nu_e - \nu_p}{\nu_p} \approx -\frac{\alpha\gamma}{(1+\gamma)}\Delta h_x \tag{4.A.24}$$

$$= -\frac{\sin 2\theta \cos\phi_p}{2\,(1+\cos\theta)}\Delta h_x; \tag{4.A.25}$$

and for the breathing mode,

$$\frac{\nu_e - \nu_p}{\nu_p} \approx -\frac{(\alpha^2+\beta^2)}{2(1+\gamma)}\Delta h_b \tag{4.A.26}$$

$$= \frac{-\sin^2\theta}{2\,(1+\cos\theta)}\Delta h_b. \tag{4.A.27}$$

Here, $\Delta h_A = h_A^e - h_A^p$, and we can identify these expressions with Eq. (4.1.2).

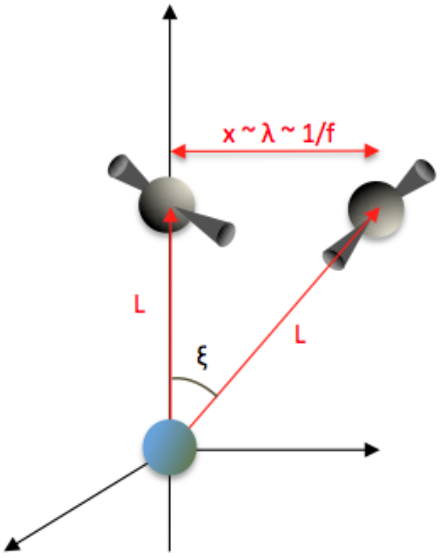

Figure 18: A system of two pulsars, distance L form the Earth, are shown along with their separation angle ξ and separation distance $x \approx L\xi$. When the GW is in the long wavelength limit, this separation distance is proportional to the GW wavelength.

Appendix 4.B Trends in $\Gamma_l(f)$ for nearby pulsar pairs

Consider a pair of pulsars separated by some *small* angle ξ and located approximately equidistant from the Earth so that $L_1 \sim L_2 \equiv L$.

As shown in Section 4.4, if the two pulsars are co-located the overlap reduction function $\Gamma_l(f) \sim \pi^3 f L$. We expect that if they are separated by a small angle the overlap reduction function will increase as $\pi^3 f L$ as though they were co-located, until the wavelength of the GW is comparable to the distance between the two pulsars. This happens when

$$\lambda \sim L\xi, \tag{4.B.1}$$

so that the value of fL where the behavior changes from the co-located case is

$$fL \sim 1/\xi. \tag{4.B.2}$$

For example, for the closest NANOGrav pulsar pair, separated by an angle $\xi \sim 3^\circ$ at a distance of $L \sim 1$ kpc, the frequency at which the linear growth of the overlap reduction function stops is

$$f \sim 10^{-9} \text{ Hz}. \tag{4.B.3}$$

The value of of the overlap reduction function where the behavior changes from the co-located case $\Gamma_l(\xi^{-1})$ is a poor estimate of the maximum value of $\Gamma_l(f)$, however, because after exiting the linear regime of Eq. (4.4.5), the overlap reduction functions continue to increase significantly before converging.

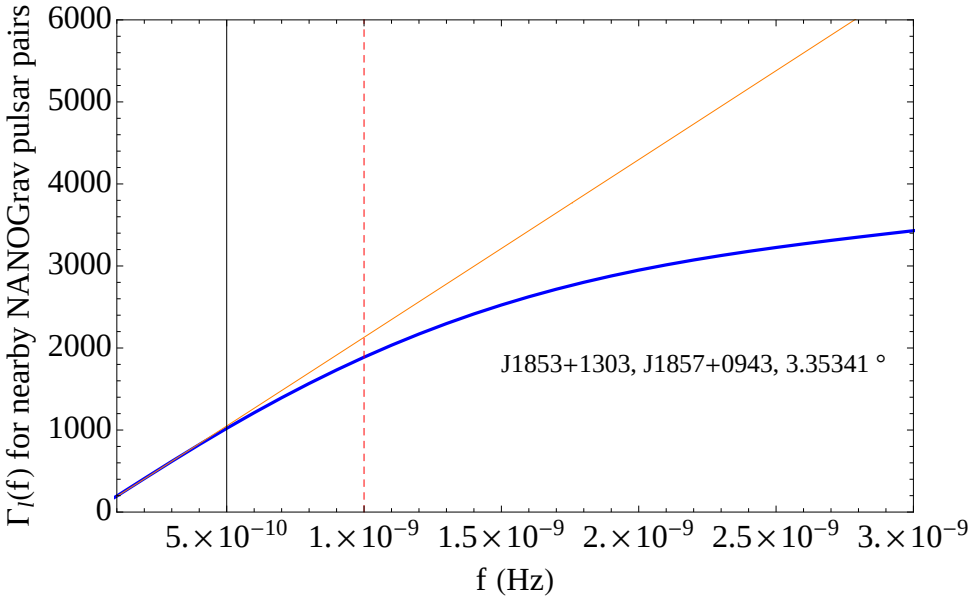

Figure 19: The NANOGrav pulsar pair $J1853+1303$, $J1857+0943$ has an angular separation of nearly 3°, with each pulsar approximately 1 kpc from the Earth. Using the estimate Eq. (4.B.2), the curve $\Gamma_l(f)$ should stop growing as $\sim \pi^3 fL$ near 10^{-9} Hz, which is shown as the red vertical dashed line. Note, however, that the curve does not converge onto constant values of $\Gamma_l(f)$ at this point; in fact the overlap reduction function continues to grow well past this point, but no longer linearly with fL (as indicated by orange solid line).

A closer look at the lower frequency portion of the plot Fig. 17(b), shown in Fig. 19, indicates that to order of magnitude this approximation is roughly valid.

Part II

Searching for bursts of gravitational waves with ground-based interferometers

Chapter 5

Laser interferometers as gravitational wave detectors

> *"Don't let others discourage you or tell you that you can't do it. In my day I was told women didn't go into chemistry. I saw no reason why we couldn't."*
>
> — Gertrude B. Elion, Biochemist, Nobel Laureate

In the previous chapters, gravitational-wave detection was considered in the context of timing extremely stable millisecond pulsars. A passing gravitational wave induces a correlated redshift on the signals from a set of pulsars, and determining this correlation forms the basis of PTA detection efforts. Laser interferometers, first mentioned in Chapter 1, also seek gravitational-wave signals on the basis of their induced effects, but in this case they seek to determine minuscule changes in length.

5.1 Laser interferometers

Interferometers are devices capable of determining very precise changes in length. The *Michelson* interferometer, a relatively simple interferometer illustrated in Fig. 20, uses a beam splitter to transmit monochromatic laser light down two orthogonal arms. Mirrors at the end of these arms reflect the incident beams in each arm.

When the light beams recombine at the beam splitter, part of the light — the part that constructively interferes — returns to the laser, in what is called the symmetric output direction. The other part — the part that destructively interferes — is sent to

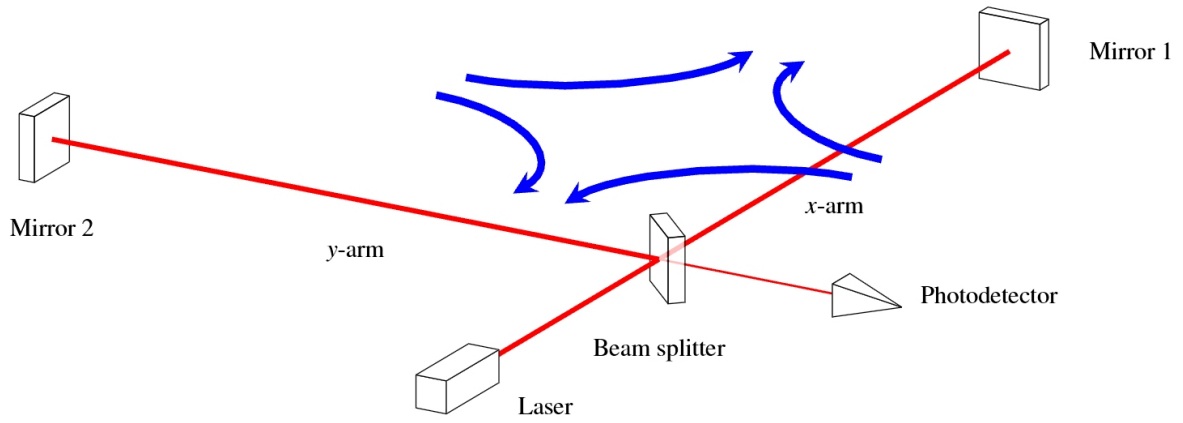

Figure 20: Schematic for a simple Michelson interferometer. Monochromatic (laser) light is transmitted to a beam splitter, which sends half the light down one arm (labeled here as the y-arm) and half down the other (labeled here as the x-arm). Mirrors at the end of each arm reflect the incident light, which recombines at the beam splitter and produces an interference pattern in the photodiode. Figure credit: European Space Agency

a photodiode in the antisymmetric output direction. Any difference in the path-length traversed by the beam (or the length of the interferometer arms) is manifest in light on the photodiode.

A passing gravitational wave will change the relative length of the interferometer's arms. If the gravitational wave has amplitude $\sim h$ and propagates normal to the plane of an interferometer with arms of length ℓ, it will induce upon the arms a strain $h = \Delta\ell/\ell$.

The gravitational-wave amplitudes produced by astrophysical sources in the ground-based interferometer frequency band are expected to be quite small, on the order of $h \sim 10^{-22}$. To produce a measurable interference pattern, the change in arm length must be on the same order as the wavelength of the laser light, $\Delta\ell \sim \lambda_{\rm laser}$. For a typical laser wavelength $\lambda_{\rm laser} \sim 1\mu m$ and kilometer-scale arms, the interferometer would be sensitive to gravitational-wave amplitudes of

$$h = \frac{\Delta\ell}{\ell} \sim \frac{\lambda_{\rm laser}}{\ell} \sim \frac{10^{-6}m}{10^{3}m} = 10^{-9}. \tag{5.1.1}$$

This is not nearly sensitive enough. Longer arms would translate into improvements in sensitivity, but constructing interferometers on scales greater than $\ell \sim 1$ kilometer is highly impractical.

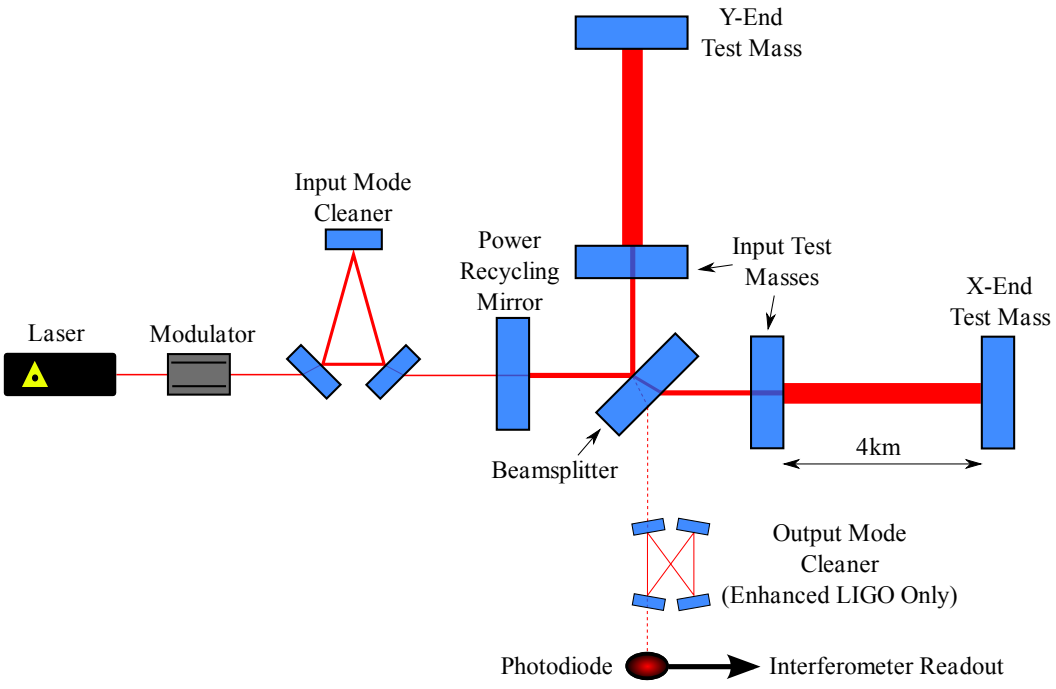

Figure 21: A simplified schematic for the Initial and Enhanced LIGO interferometers. Note that the output mode cleaner was not present in Initial LIGO. Figure from Smith et al. [128]

Instead, a few things are done to turn the Michelson interferometer into a practical gravitational-wave detector such as the LIGO interferometers. These features are described in the following subsections.

5.1.1 Fabry-Pérot cavities

As mentioned before, ground-based laser interferometers become more sensitive with longer arms. In theory the most optimal choice of arm length would be on the order of the gravitational-wave wavelength $\ell \sim \lambda_{\rm gw}$. Since the wavelengths of interest are typically $\sim$ 1000km, the optimal interferometer would need to have arms spanning thousands of kilometers. In practice, experimental and physical limitations restrict arm lengths to kilometer scales.

Fortunately, the optical path-length of the light can be extended to effectively create a longer interferometer arm. To do this, a Fabry-Pérot optical cavity is created in each arm of the interferometer. Each cavity consists of a pair of parallel, partially transmitting mirrors. In this cavity, laser light bounces back and forth many times, completing an

	Initial LIGO	Enhanced LIGO	Advanced LIGO
Laser Power (W)	10	35	200
Maximum Arm Cavity Power (kW)	10	100	800
Maximum Recycled Power (kW)	0.2	0.75	4

Table 2: Laser power properties for the Initial, Enhanced and Advanced LIGO.

effective path-length ℓ_{eff} that can greatly exceed the arm length ℓ.

5.1.2 Power recycling

Another strategy for improving interferometer sensitivity is to increase the power in the input light beam. One obvious way to do this is to construct the interferometer with a more powerful laser. Initial LIGO used a laser with power 10 W, and one of the modifications made for Enhanced LIGO was the addition of a more powerful 35 W laser. Advanced LIGO will utilize a 200 W laser. The maximum power achievable via upgraded lasers is limited, however.

There is another, more subtle method to increase the power that is input to the interferometer. This is called power recycling, and utilizes the light that would otherwise exit the interferometer though the symmetric output direction. To perform power recycling, a *power recycling mirror* (PRM) is constructed between the laser and the beam splitter to create an effective cavity. The placement of the PRM is shown in Fig. 21, which schematically illustrates Initial and Enhanced LIGO.

Light is emitted by the laser and propagates to the beam splitter, where some of the incident light is transmitted into the instrument, and some of the light (in fact, much of it) is reflected back toward the laser. With the addition of the PRM, the light reflected by the beam splitter can be be redirected back toward it. A standing wave can be built between the PRM and the beam splitter in this recycling cavity, effectively increasing the power that is incident on the interferometer. Table 2 compares laser power for the Initial, Enhanced and Advanced LIGO experiments along with the maximum effective power achieved through power recycling, and showcases the dramatic gains in power achieved

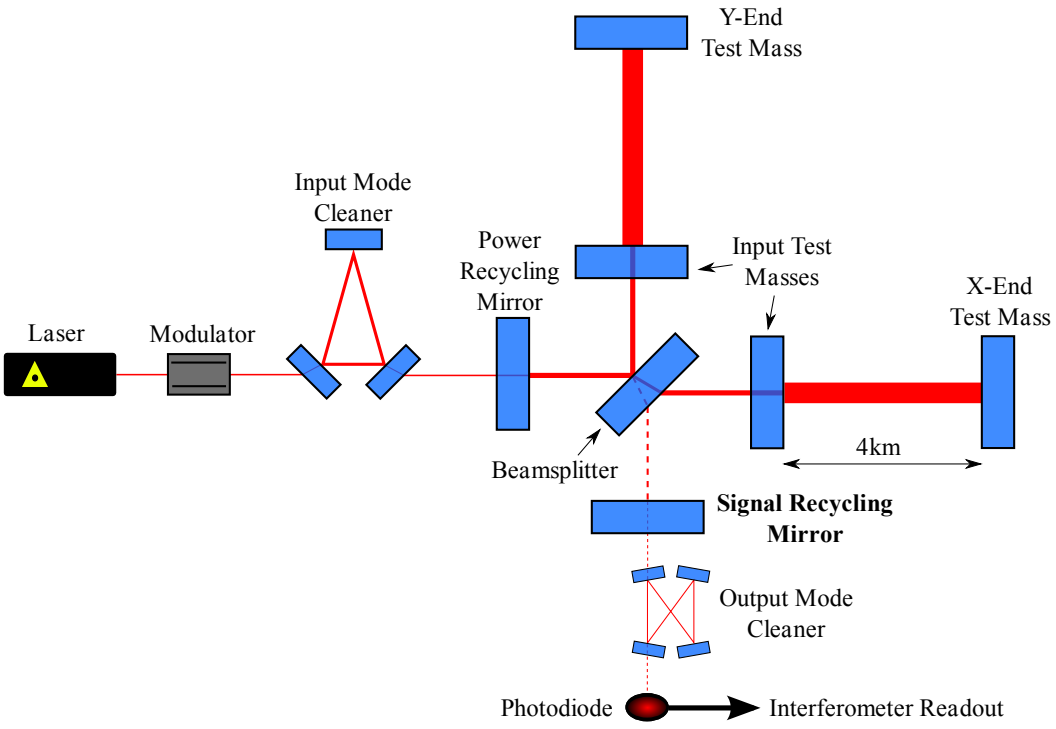

Figure 22: A simplified schematic for the Advanced LIGO interferometers. A key addition is the Signal Recycling Mirror. Figure from Smith et al. [128]

with the PRM.

5.1.3 Signal recycling

We have seen that Fabry-Pérot cavities and power recycling enhance the sensitivity of ground-based interferometers. There is yet another method to improve sensitivity even more; this method is called signal recycling, and involves the addition of a signal recycling mirror between the beam splitter and the output port of the interferometer (see Fig. 22, which diagrams Advanced LIGO's layout). The signal recycling mirror forms an optical cavity which allows gravitational-wave induced sidebands to be reflected back into the interferometer. Adjusting the position of the signal recycling mirror allows the interferometer to either be tuned to specific gravitational-wave sources (at the cost of decreased bandwidth) or to have an increased bandwidth.

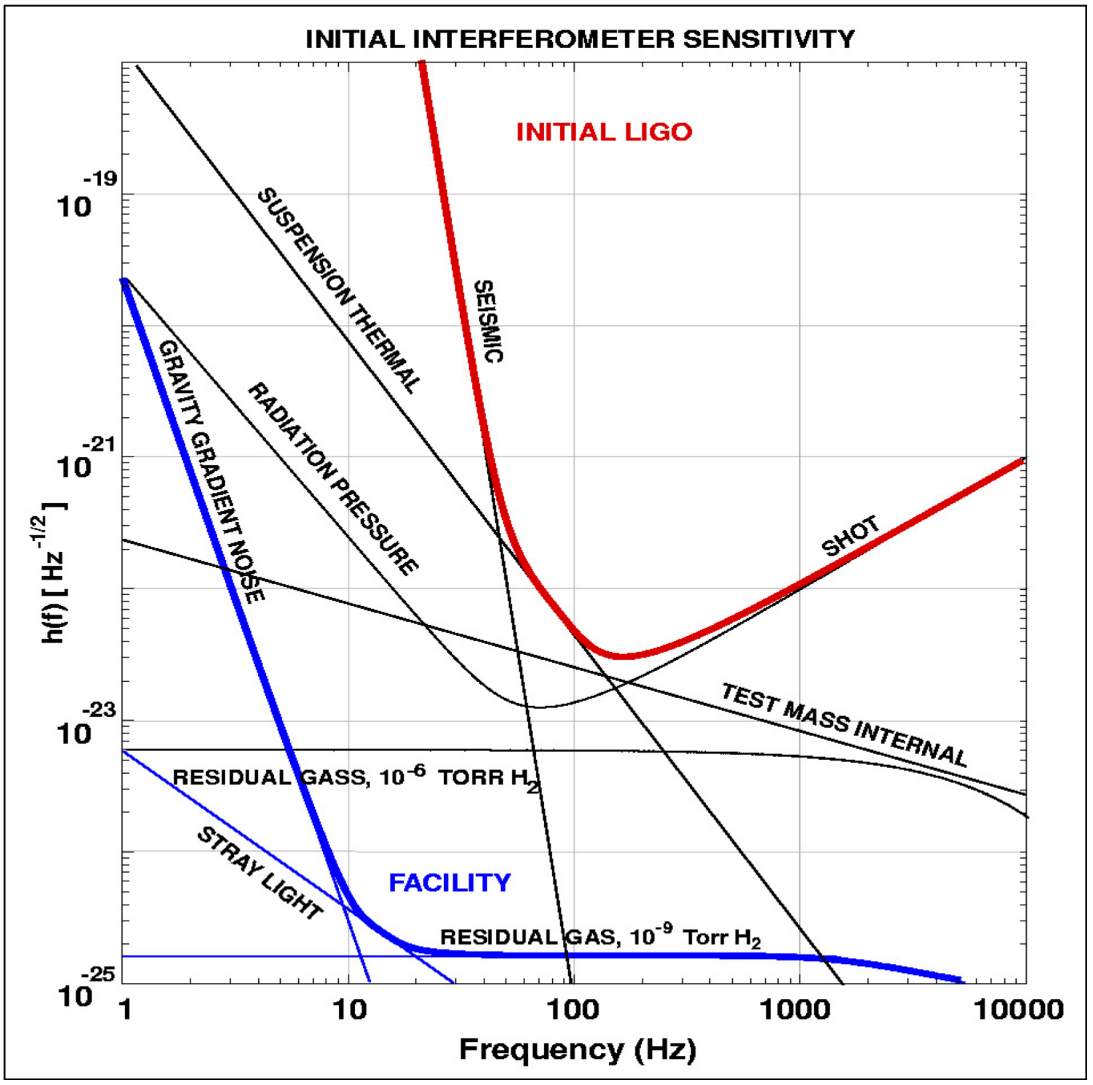

Figure 23: Various noise sources present in Initial LIGO are shown. The total sum of noise sources is shown in red; it is this curve that characterizes the overall sensitivity of the detector. Figure from Hughes et al. [129].

5.2 Sources of noise in interferometric detectors

In the last subsections, features designed to increase the sensitivity of ground-based laser interferometers were discussed. It is equally worthwhile to probe the noise inherent to these detectors since it is what limits their achievable sensitivity. Fig. 23 illustrates the primary sources of noise that limit the sensitivity in modern gravitational-wave interferometers.

5.2.1 Seismic noise

The Earth does not constitute a "quiet" laboratory. Geologic disturbances and the motion of people, aircraft and other human activity cause a substantial amount of vibration in the ground at any given location. Weather-related effects such as wind (and even tumbleweeds hitting the detector) can also contribute to seismic noise in the detectors. Seismic noise is what most limits detector sensitivities below ~ 40 Hz. To mitigate the effects from this noise, seismic isolation systems are employed. These systems typically involve spring-loaded platforms, which reduce motion above their resonant frequencies.

5.2.2 Shot noise

The light emanating from the laser in an interferometer comes in discrete packets or quanta of light. Measurements of the output power in the instrument really measure of the number of photons that arrive at the photodetector. Independently occurring discrete events, such as the arrival of a photon, are governed by the Poisson distribution. This means that for a large number of photons N arriving at the photodetector, there will be some fluctuations in the number of photons on the order of $N^{1/2}$; this will also result in fluctuations in the power observed at the photodetector. These fluctuations in power are known as shot noise, and this is what limits detector sensitivities at frequencies above a few hundred Hertz. More power in the laser — or more photons — will reduce the relative size of the fluctuations, but power recycling techniques are needed in addition to a powerful laser to mitigate shot noise sufficiently.

5.2.3 Radiation pressure

When a beam of photons A beam of photons that are incident on a mirror experience a change in momentum as they are reflected back. This results in the application of a radiation pressure force to the mirror. A beam of photons impinging upon a mirror thus produce radiation pressure noise within the instrument. If the radiation pressure noise were constant, it could very easily be corrected with the mirror controls. However, the fluctuations in the number of photons hitting the mirror result in fluctuating radiation pressure. Reducing laser power will reduce the radiation pressure, but this comes at

the cost of increasing shot noise. Together the radiation pressure and shot noise can be reduced to a point known as the "standard quantum limit". This is achieved when the amplitude of both noise sources is the same at some target frequency.

5.2.4 Thermal noise

Vibrations within the mirrors, suspensions and other parts of the instrument result in thermal noise within the detector. The most significant of these within the LIGO detectors arise from the pendulum suspension system for the mirrors and the internal vibration modes of the mirrors themselves. To mitigate such noise, suspension systems and optics are designed to have resonant frequencies far from the frequencies of interest in the detector (a few Hz for suspensions and kHz for the optics). The materials used in the detector are also chosen to have a high quality factor, which confines the noise they create to narrow bandwidths. Thermal noise sources are dominant in the detector from approximately $\sim$ 40 Hz–$\sim$ 100 Hz.

5.2.5 Gravity gradient noise

Another type of noise can arise due to gradients in gravity: fluctuations in the density of the atmosphere or ground due to seismic activity lead to fluctuations in the Earth's gravitational field, and this couples directly with thee test masses in the interferometer via Newtonian force. This noise cannot be easily isolated from the instrument, as is the case with seismic noise, because gravitational force cannot be screened. The only way to truly eliminate such noise would be to place the detector in space. However, it is not currently a dominant effect in the detectors; at low frequencies, gravity gradient noise is overwhelmed by seismic noise, and above a few Hz by thermal noise. Future generation detectors may attempt to mitigate some gravity gradient noise by moving mirrors underground or compensating for changes in the Earth's gravitational field.

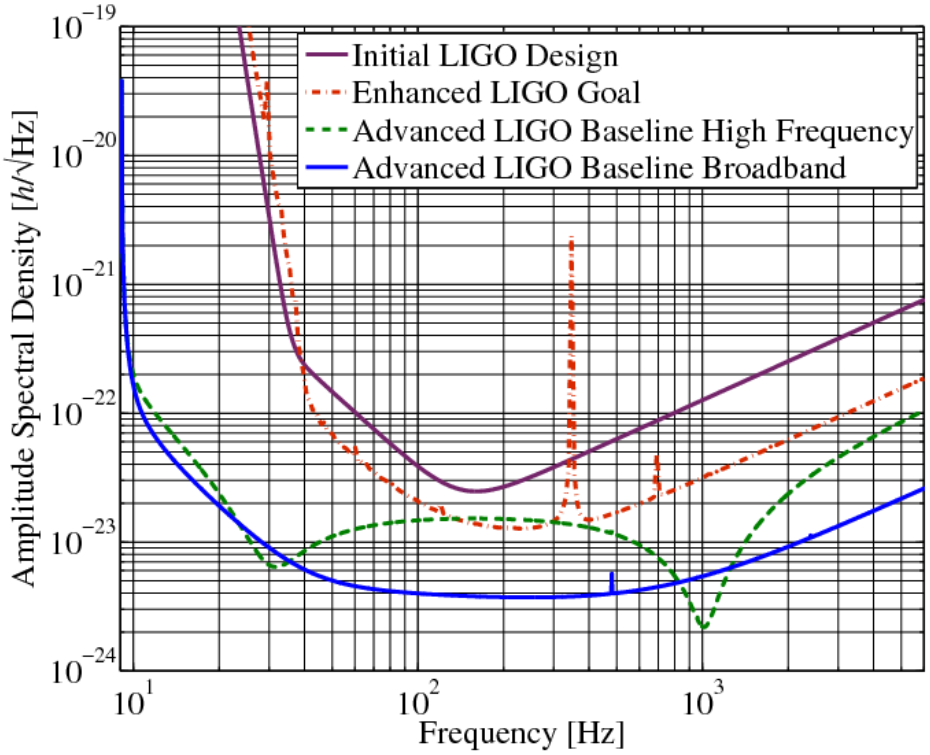

Figure 24: Sensitivity curves calculated for LIGO in strain amplitude spectral density. As mentioned in Sec. 5.1.3, Advanced LIGO will be tunable to specific sources via signal recycling; this is illustrated with the dashed green curve. The solid blue curve for Advanced LIGO corresponds to its baseline sensitivity, or to distances to which the coalescence of a binary system of $1.4M_{\odot}$ neutron stars could be detected by a single detector at 200 Mpc. Figure from Smith et al. [128].

5.3 Sensitivity of ground-based laser interferometers

The noise sources discussed in the last section provide the limits of what modern gravitation-wave detectors can achieve in terms of their sensitivity. In Fig. 24, sensitivity curves are shown for Initial, Enhanced, and Advanced LIGO, and demonstrate the gains in sensitivity that will be achieved with advanced era detectors. The broadband curves shown for initial LIGO, Enhanced LIGO and Advanced LIGO correspond to the distances at which the coalescence of a binary system of $1.4M_{\odot}$ neutron stars could be detected by a single detector. These are roughly 15, 30, and 200 Mpc, respectively.

Chapter 6

Use of the excess power statistic in LIGO burst searches

> "*My methods are really methods of working and thinking; this is why they have crept in everywhere anonymously.*"
>
> — Emmy Noether, Mathematician (1882-1935)

GW signals from the coalescence of intermediate black holes and other compact objects, being transient in nature, involve slightly different search algorithms than stochastic backgrounds of GWs. The underlying statistics, however, involve similar techniques. As with PTA detection experiments, the problem revolves around extracting a small signal (possibly with a known shape) from noisy data. The optimal statistic, discussed in Chapter 3, utilizes some concepts of matched filtering by constructing an optimal filter and using this filter to construct the statistic.

For compact binary coalescence events, if the shape of the expected signal (or waveform) can be well-modeled with post-Newtonian techniques or numerical relativity, matched filtering can be utilized to search for GW signals. In this case a template bank containing some large number of possible waveform filters is constructed, and the GW signal is extracted from the data by correlating each filter with the data.

For some astrophysical sources of gravitational waves, however, the physics underlying the GW emission process can be probed only by studying the highly non-linear self-interactions of the gravitational field, making it difficult (sometimes, impossible) to precisely model the gravitational wave waveform. Additionally, unknown signals from

unanticipated sources could exist. In these cases, the challenge is to extract a signal of unknown shape from the data.

One method that has been proposed and employed for this type of *unmodeled* search is the excess power method, which involves studying time-frequency decompositions of detector data for gravitational-wave bursts [130, 131, 7]. The excess power method effectively scans detectors' outputs for transients that are statistically significant relative to background noise. Blocks of time and frequency, or time-frequency tiles, are constructed based on knowledge of the signal's duration and frequency, and the total power within each tile is calculated. If a signal exists in the data, more power should be present than would exist from detector noise alone.

This chapter is laid out as follows. Sections 6.1 and 6.2 outline the methodology and implementation of a search pipeline based on the excess power statistic; in Section 6.1, the excess power statistic is developed from a frequentist point of view, and in Section 6.2, the steps used to implement the excess power statistic into an actual search pipeline — the ExcessPower pipeine — are outlined.

In the remainder of this chapter, efforts to characterize and improve the performance of ExcessPower are considered. Section 6.3 outlines the manner in which the performance of search pipelines can be probed, and presents two potential changes to the existing infrastructure of ExcessPower for improved funtionality. Section 6.4 outlines the manner in which the ExcessPower pipeline can be used to aid in detector characterization efforts for advanced generation detectors.

The derivation of the excess power statistic in this chapter follows closely the analysis of Anderson et al. [130] and Brady et al. [131], and the original implementation by Cannon [132].

6.1 Geometrical approach to the excess power statistic

To understand how the excess power statistic can be used in GW searches, it is useful to construct the statistic in an intuitive geometric fashion. In this section, the statistic will be defined for single-interferometer searches. The analysis for multiple detectors is not considered in this dissertation, and will be the subject of future work.

A single interferometric detector outputs stretches of discretely sampled data $\mathbf{s} = \{s_0, s_1, \ldots, s_{N-1}\}$ containing N data points. The collection of all such data stretches forms a N-dimensional vector space $\mathcal{V}$.

Assuming that the detector noise is stationary and Gaussian noise, the detector output may be expressed as $\mathbf{s}(t) = \mathbf{h}(t) + \mathbf{n}(t)$, where $h(t)$ represents the gravitational-wave signal and $n(t)$ represents the noise in the detector. This can also be written in index notation where each index denotes a data point, i.e.

$$s_i = h_i + n_i. \tag{6.1.1}$$

The statistical properties governing the noise in the data are fully described by a correlation matrix

$$R_{ij} = \langle n_i n_j \rangle = C_n(|i - j|\Delta t), \tag{6.1.2}$$

where $C_n(t)$ is the correlation function of the noise and Δt is the sampling time. From this definition, one can define an inner product on $\mathcal{V}$:

$$(\mathbf{a}, \mathbf{b}) = \sum_{i,j=0}^{N-1} a_i R_{ij}^{-1} b_j \tag{6.1.3}$$

which completely characterizes the geometry within $\mathcal{V}$.

As mentioned earlier, the excess power method works by constructing time-frequency tiles. Consider some time-frequency window defined by

$$\mathcal{W} = \{t_s, \delta t, f_s, \delta f\}. \tag{6.1.4}$$

Here f_s is the starting frequency of the window, δf is the bandwidth, t_s is a starting time and δt is a duration. To focus analysis on data that lie inside $\mathcal{W}$, a number of various methods can be used. The simplest way to do this is with geometry. The vectors in $\mathcal{V}$ are data stretches containing N points of data. Any data stretch that lies within the window $\mathcal{W}$ effectively lies on a *subspace* $\mathcal{V}_\mathcal{W}$ of $\mathcal{V}$.

Given a set of data (or vector within $\mathcal{V}$), one can then restrict work to stretches that lie in $\mathcal{W}$ using an orthogonal projection: any vector $\mathbf{s}$ can be broken into a piece $\mathbf{s}_{||}$ that lies parallel to $\mathcal{V}_\mathcal{W}$ and a piece $\mathbf{s}_\perp$ that is orthogonal to $\mathcal{V}_\mathcal{W}$, i.e.

$$\mathbf{s} = \mathbf{s}_{||} + \mathbf{s}_\perp. \tag{6.1.5}$$

It is useful to note at this point that the dimension of $\mathcal{V}_{\mathcal{W}}$ (over the field $\mathbb{R}$) is

$$\dim \mathcal{V}_{\mathcal{W}} = 2\,\delta t\,\delta f. \tag{6.1.6}$$

In practice, other methods can be used to choose data that lies within $\mathcal{W}$. A more data-intuitive method to do this utilizes the FFT: if data is produced in the time-domain, one can truncate the data to give it a needed duration, transform the data with a discrete Fourier transform, and then truncate the frequency values to lie within the necessary bandwidth.

Regardless of the projection method used, once vectors (or data) lying within the subspace $\mathcal{V}_{\mathcal{W}}$ have been obtained, they can be written as

$$s_J = \sum_{j=0}^{N-1} A^j_J\, s_j \tag{6.1.7}$$

where A^j_J is a real-valued $(\dim \mathcal{V}_{\mathcal{W}}) \times N = 2\,\delta t\,\delta f \times N$ matrix, s_j are real, and J labels indices that run from $0 \leq J \leq (2\,\delta t\,\delta f - 1)$.

Ultimately, the goal is to obtain a statistic that quantifies the amount of power in the data within the time-frequency window $\mathcal{W}$. This can now be done in a very straightforward way using Eq. 6.1.7:

$$\mathcal{E}_{\mathcal{W}}(\mathbf{s}) = \sum_{I,J=0}^{2\,\delta t\,\delta f-1} Q_{IJ}\, s_I\, s_J \tag{6.1.8}$$

is defined as the power statistic associated with $\mathcal{W}$, where Q_{IJ} is related to the correlation matrix R_{IJ} by

$$\sum_J Q_{IJ}\, R_{JK} = \delta_{IK} \tag{6.1.9}$$

and

$$R_{JK} = \langle s_J\, s_K \rangle = \sum_{j,k} A^j_J\, A^k_K\, R_{jk}. \tag{6.1.10}$$

This statistic describes the total power present in the data stream within $\mathcal{W}$, measured relative to the detector noise. Note that the correlation matrix R_{JK} depends on the acquired data (signal plus noise) since there is no way a priori to obtain the noise independently. In practice this is not a problem since gravitational-wave signals will be weak and small in number.

6.1.1 False alarm and detection probability

With a power statistic in hand, it is useful to understand the characteristics of this statistic in the presence of a signal. When a signal is *not* present in the data, the sum in Eq. 6.1.8 is over squares of independent, zero-mean, unit-variance Gaussian random variables (i.e., $\mathcal{E}(\mathbf{s}) = \mathcal{E}(\mathbf{n})$). By definition, this is just a χ^2 distribution with $D = \dim \mathcal{V}_{\mathcal{W}} = 2\delta t\, \delta f$ degrees of freedom. The cumulative probability for $\mathcal{E}$, $P(\mathcal{E} > \mathcal{E}^*)$, can be used to define the *false alarm probability*; this is

$$Q_0(\mathcal{E}^*) = P(\mathcal{E} > \mathcal{E}^*) = \frac{\Gamma(D/2, \mathcal{E}^*/2)}{\Gamma(D/2)}, \tag{6.1.11}$$

where $\Gamma(D/2, \mathcal{E}^*/2)$ is the incomplete Gamma function, $\Gamma(D/2)$ is the Gamma function, and $\mathcal{E}^*$ is the detection threshold.

When a signal *is* present, Eq. 6.1.8 contains noise *and* signal pieces, so that

$$\mathcal{E}(\mathbf{s}) = \mathcal{E}(\mathbf{n} + \mathbf{h}) = \sum_{I,J=0}^{2\,\delta t\,\delta f - 1} Q_{IJ}\,(n_I + h_I)\,(n_J + h_J)\,. \tag{6.1.12}$$

The amplitude of the signal is [1]

$$A^2 = \left(\mathbf{h}_{||}, \mathbf{h}_{||}\right) = \sum_{I,J=0}^{D-1} Q_{IJ}\, h_I\, h_J. \tag{6.1.13}$$

Here where we use the notation $||$ as a reminder that $\mathbf{h}$ lies within the subspace $\mathcal{V}_{\mathcal{W}}$.

It can be shown (see Ref. [130]) that the probability distribution for $\mathcal{E}$ is

$$\begin{aligned} p(\mathcal{E}|A, D/2) &= \sum_{n=0}^{\infty} \frac{e^{-A^2/2}(A^2/2)^n}{n!} \frac{e^{-\mathcal{E}/2}(\mathcal{E}/2)^{n+D/2-1}}{\Gamma(n + D/2)} \\ &= (1/2)e^{-(\mathcal{E}+A^2)/2} \left(\mathcal{E}^{1/2}/A\right)^{D/2-1} I_{V-1}\left(A\mathcal{E}^{1/2}\right), \end{aligned} \tag{6.1.14}$$

where the notation $I_n(x)$ designates a modified Bessel function of the first kind of order n. This distribution is a non-central χ^2 distribution with non-centrality parameter A^2. The *detection probability* can now be defined; this comes from the cumulative probability distribution for $\mathcal{E}$, and can be written as

$$Q_A(\mathcal{E}^*, A, D/2) = P(\mathcal{E} > \mathcal{E}^*|A, D/2) = \int_{\mathcal{E}^*}^{\infty} p(\mathcal{E}|A, D/2) d\mathcal{E}. \tag{6.1.15}$$

Eq. (6.1.15) represents the detection probability for a given threshold $\mathcal{E}^*$ and a given signal amplitude A.

[1] As explained in Ref. [130], A is the SNR that would be obtained by matched filtering, if matched filtering were possible, and if the signal h was contained within the relevant time-frequency window.

6.2 Implementation of the excess power statistic: `gstlal_excesspower`

Having laid out the theory governing the excess power statistic in Secs. sections 6.1 and 6.1.1 the details surrounding actual implementation in burst gravitational-wave searches can be described. The excess power statistic was first implemented by Cannon [132] as `lalapps_power` (fully detailed in [131]) and later restructured as `gstlal_excesspower` by Pankow [133], which is built on `gstreamer`, a low-latency framework for signal processing tools. The `gstlal_excesspower` program reads in gravitational-wave detector time series data from single or multiple detectors, and outputs a `sngl_burst` table that lists gravitational-wave signal candidates, along with parameters that characterize the candidates. Software injections can also be processed by `gstlal_excesspower`. This will be discussed in more detail in Sec. 6.3. Note that the description of the implementation in this section closely follows that of Cannon [132].

To describe the manner in which `gstlal_excesspower` applies the excess power statistic to data, consider a discretely-sampled time series of N samples, s_j (where $0 \leq j \leq N-1$), that has been sampled at a rate $1/\Delta t$.

6.2.1 Whitening

Given a subset or segment of detector data that has some duration and bandwidth, the excess power statistic quantifies the total power in that data segment relative to the detector noise. To calculate the excess power probability, the first procedure that must take place in the `gstlal_excesspower` pipeline is to whiten and normalize the input data so that it takes the form of stationary, white, unit-variance Gaussian noise in the absence of a signal.

Whitening (and much of the analysis that follows) takes place in the Fourier domain. Before proceeding, it is necessary to multiply s_j by a window function w_j that is tapered at its start and end to mitigate corruption that could arise from aperiodicity of the sample at its boundaries.

The data then undergoes a discrete Fourier transform. For a given frequency bin k,

the complex amplitude of the signal can be described as

$$\tilde{s}_k = \frac{\Delta t}{\sigma_w} \sum_{j=0}^{N-1} w_j \, s_j \, e^{-2\,\pi i\, j\, k/N} \tag{6.2.1}$$

where the mean of the tapered window samples is

$$\sigma_w = \left[\frac{1}{N} \sum_{j=0}^{N-1} w_j^2 \right]^{1/2}. \tag{6.2.2}$$

In principle the number of frequency bins can range from $0 \leq k \leq N-1$, but the frequency bins $[N/2] < k < N$ correspond to negative frequency components. Since the input data is entirely real, these are determined by the complex conjugates of the positive frequency components, $\tilde{s}_{-k} = \tilde{s}_k^*$. For the remaining frequency bins, each bin can be associated to some frequency

$$f_k = k\Delta f \tag{6.2.3}$$

where the size of the frequency bin Δf is determined by the duration $N\Delta t$, i.e., $\Delta f = 1/(N\Delta t)$.

To whiten, a power spectral density (PSD) is needed. In general, a one-sided PSD can be defined as[2]

$$\mathcal{P}_k = \Delta f \langle |\tilde{s}_k|^2 + |\tilde{s}_{N-k}|^2 \rangle, \tag{6.2.4}$$

where $\tilde{s}_{N-k}$ represent the negative frequency components and $\langle\,\rangle$ denotes the expectation value. Since we are working with real data, this can be simplified to

$$\mathcal{P}_k = 2\Delta f \langle |\tilde{s}_k|^2 \rangle,\ k = 1, \cdots, [(N-1)/2]. \tag{6.2.5}$$

In `gstlal_excesspower`, the PSD is estimated using the median power at each frequency for a number of overlapping segments. This convention is implemented to avoid any bias that would arise due to a gravitational-wave signal or large noise-transient in the data.

The goal of whitening is to make

$$\langle |\tilde{s}_k|^2 \rangle = 1 \tag{6.2.6}$$

[2]Here we can ignore P_0 and $P_{[N/2]}$ when N is even.

when no signal is present. This is accomplished by multiplying the Fourier amplitudes $\tilde{s}_k$ by an appropriate factor involving the PSD:

$$\hat{s}_k = \sqrt{\frac{2\Delta f}{\mathcal{P}_k}}\tilde{s}_k. \tag{6.2.7}$$

Note on the PSD and variance

It is useful to reflect on the relationship between the PSD and the variance of the time series $\langle s_j \rangle$ that would be obtained by inverse Fourier transforming the whitened frequency domain data. These quantities are related in the sense that the PSD effectively describes how the variance of the data is distributed over its frequency components.

Quantitatively, this relationship is described by

$$\begin{aligned} \langle s_j^2 \rangle &= \frac{1}{N^2\,\Delta t^2}\sum_{k=0}^{N-1}\sum_{k'=0}^{N-1}\langle \tilde{s}_k\,\tilde{s}_{k'}^*\rangle\, e^{2\pi i\, j\,(k-k')/N} \\ &= \frac{1}{2N\Delta t}\sum_{k=0}^{N-1}\mathcal{P}_k \end{aligned} \tag{6.2.8}$$

when the input time series corresponds to a stationary process. However, windowing a time series correlates its frequency components (because the Fourier transform of the series has effectively been convolved with that of the window). This means that $\langle \hat{s}_k\,\hat{s}_{k'}^*\rangle$ is not proportional to $\delta_{kk'}$, and is given instead by the quantity

$$\langle \hat{s}_k\,\hat{s}_{k'}^*\rangle = \frac{1}{\sigma_w^2}\sum_{j=0}^{N-1} w_j^2 e^{-2\pi i j(k-k')/N}. \tag{6.2.9}$$

Substituting Eq. (6.2.9) into Eq. (6.2.8), one can see that the variance of the (whitened) time series obtained by inverse Fourier transforming the whitened data becomes

$$\langle \mathfrak{s}_j^2 \rangle = \frac{1}{\Delta t^2 \sigma_w^2} w_j^2 \tag{6.2.10}$$

where here we have used $\mathfrak{s}$ to denote the difference between the un-whitened and whitened variances.

6.2.2 Channel Filters

After whitening the input time series in the frequency domain, digital filters are applied to the data. The filters turn a single frequency series into many frequency series (or channels), allowing `gstlal_excesspower` to function as a multi-resolution search pipeline.

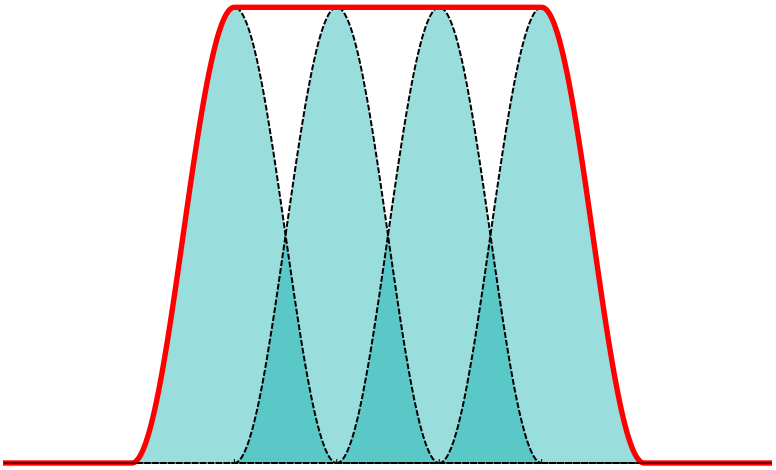

Figure 25: Illustration of the channel filter choice made in `gstlal_excesspower`. The Hann window (dashed lines filled in cyan) is chosen as the filter for the narrowest frequency channels, lying centered on each channel. When summed, adjacent Hann windows produce a Tukey window (solid red line) with a flat top and ends that taper as $\sin^2$. In this figure, 4 Hz filters are summed to produce a 16 Hz Tukey window. Note that overlap in the Hann windows need only be accounted for adjacent filters.

In the language of Sec. 6.1, it is the application of filters that partially projects the input time series onto the many different possible subspaces $s_{||}$.[3]

There are many possible ways to choose a channel filter. It turns out to be useful to choose the filter in such a way that summing narrow-band filters produces a filter for a wide-band channel. Suppose that the filter for some particular bandwidth B is denoted by $\tilde{H}_k(f_1, B)$ (where $f_1 \leq f_k < f_2 = f_1 + B$). In this case, if b is the bandwidth of the narrowest channel and $B = nb$, then the filter for the channel of bandwidth B is given by

$$\tilde{H}_k(f_1, B) = \sum_{m=0}^{n-1} \tilde{H}_k(f_1 + m\, b, b). \tag{6.2.11}$$

One choice of filter that satisfies Eq. 6.2.11 for the narrow-band filters is the *Hann* window, defined via

$$\tilde{H}_k(f_1, b) = \begin{cases} \sin^2\left[\pi/(2b)\,(f_k - f_1 + b/2)\right], & f_1 - b/2 \leq f_k \leq f_1 + 3b/2 \\ 0, & \text{otherwise.} \end{cases} \tag{6.2.12}$$

The sum of adjacent Hann windows (of equal bandwidth) that lie centered on their frequency channel produces a *Tukey* window that is flat over middle frequencies and

[3]The rest of this projection comes from choosing specific segments of time.

tapered at its ends. An example is shown in Fig. 25, which plots four 4 Hz wide Hann windows and their sum, which is a 16 Hz Tukey window.

To discuss the normalization of the filters, we define their "magnitude" to be

$$\left\{\tilde{X},\tilde{Y}\right\} = \sum_{k'=0}^{N-1}\sum_{k=0}^{N-1}(-1)^{(k-k')}\langle\hat{s}_k\hat{s}^*_{k'}\rangle\tilde{X}^*_k\tilde{Y}_{k'}. \tag{6.2.13}$$

The narrowest channels (of bandwidth b) are normalized so that they have a magnitude

$$\left\{\tilde{H}^*_k(f_1,b)\tilde{H}_{k'}(f_1,b)\right\} = \frac{b}{\Delta f}. \tag{6.2.14}$$

This choice of normalization will be important later on when considering the channel time series that are constructed from the filtered frequency-domain data.

It is worth pointing out that wide channels formed from summed from narrow channels will *not* have a magnitude of $nb/\Delta f$. This can be seen with a simple example: suppose that

$$\tilde{H}_k(f_1,2b) = \tilde{H}_k(f_1,b) + \tilde{H}_k(f_1+b,b). \tag{6.2.15}$$

Then the magnitude of the the filter $\tilde{H}_k(f_1,2b)$ can be found to be

$$\left\{\tilde{H}_k(f_1,2b),\tilde{H}_k(f_1,2b)\right\} = \frac{2b}{\Delta f} + 2\left\{\tilde{H}_k(f_1,b),\tilde{H}_k(f_1+b,b)\right\}. \tag{6.2.16}$$

In light of this result, the specific choice of the Hann and Tukey windows as filters becomes clear: it allows the approximation to be made that only adjacent channels have enough overlap for non-trivial inner products. In this approximation, the only cross-terms such as those appearing in Eq. (6.2.16) that need strict accounting for are the cross-terms from *adjacent* channels. In general, then, a filter spanning n channels can be given by

$$\tilde{H}_k(f_1,nb) = \sum_{i=0}^{n-1}\tilde{H}_k(f_1+ib,b), \tag{6.2.17}$$

and its magnitude is

$$\left\{\tilde{H}(f_1,nb),\tilde{H}(f_1,nb)\right\} = \frac{nb}{\Delta f} + 2\sum_{i=0}^{n-2}\left\{\tilde{H}(f_1+ib,b),\tilde{H}(f_1+(i+1)b,b)\right\}. \tag{6.2.18}$$

The quantity on the right-hand side of Eq. (6.2.18) will be denoted as $\mu^2(f_1,nb)$ in the remainder of this section.

6.2.3 Construction of Time-frequency Tiles

To construct time-frequency tiles, an inverse Fourier transform must be used to take the many channel-filtered frequency series back into the time domain. What began as a single time series is now many time series, one each for many different frequency channels. The inverse Fourier transform that describes these *channel* time series' can be written explicitly; for the channel that has bandwidth b and starts at frequency f_1, it is

$$z_j(f_1, b) = \frac{1}{N\,\Delta t} \sum_{k=0}^{N-1} \hat{s}_k \tilde{H}_k^*(f_1, b)\, e^{2\,\pi\, i\, j(k-k')/N}. \tag{6.2.19}$$

Following the inverse Fourier transform, the channel time series are handed to an undersampler, which resamples them. This is necessary because the channel time series are oversampled (for the narrowest frequency bands, the sample rate greatly exceeds the Nyquist rate). Oversampling is problematic for `gstlal_excesspower` because it means there are more samples per unit time than degrees of freedom per unit time in the channel time series: the tiles created at the original sampling rate are correlated to within the duration of the filter. The undersampler selects evenly spaced samples from each channel time series to yield appropriately sampled time series.

Ultimately, we will be interested in quantities that involve power, or the square of the channel time series. The mean square of Eq. 6.2.19 is

$$\langle z_j^2(f_1, b)\rangle = \frac{1}{N^2\Delta t^2} \sum_{k=0}^{N-1}\sum_{k'=0}^{N-1} \langle \hat{s}_k \hat{s}_{k'}^*\rangle \tilde{H}_k^*(f_1, b)\tilde{H}_{k'}(f_1, b) e^{2\pi i j(k-k')/N}. \tag{6.2.20}$$

This quantity is actually sample-dependent (meaning that it depends on the index j) due to the windowing that took place on the initial time series prior to being Fourier transformed. This can be corrected by applying another Tukey filter: by demanding that the middle part of the series is flat, the mean square of the series $z_j(f_1, b)$ should be independent of j when $j_1 \leq j < j_2$. In this case, the mean square can be written as

$$\langle z_j^2(f_1, b)\rangle = \frac{1}{N^2\Delta t^2} \sum_{k=0}^{N-1}\sum_{k'=0}^{N-1} (-1)^{(k-k')} \langle \hat{s}_k \hat{s}_{k'}^*\rangle \tilde{H}_k^*(f_1, b)\tilde{H}_{k'}(f_1, b) \tag{6.2.21}$$

which is just

$$\langle z_j^2(f_1, b)\rangle = \frac{1}{N^2\,\Delta t^2}\frac{b}{\Delta f} \tag{6.2.22}$$

thanks to the normalization choice made in Eq. (6.2.14). As in the frequency domain, a channel that spans n narrow channels can be written very simply with a sum

$$z_j(f_1, nb) = \sum_{k=0}^{N-1} \left(\sum_{m=0}^{n-1} \tilde{H}_k(f_1 + mb, b) \right) e^{2\pi ijk/N} \tag{6.2.23}$$

$$= \sum_{m=0}^{N-1} z_j(f_1 + mb, b) \tag{6.2.24}$$

and broader channel time series may be defined. In `gstlal_excesspower` this step is accomplished using an element known as the "matrix mixer", which combines samples along with the appropriate normalizations to achieve higher bandwidth resolutions.

The next step in the detection process involves the construction of time-frequency tiles. Before proceeding, it is helpful to express the number of degrees of freedom contained in a tile of bandwidth B and duration T; this is

$$d = 2BT. \tag{6.2.25}$$

Here the factor of two comes from the Nyquist-Shannon sampling theorem, which requires a signal with bandwidth B to be sampled at a rate $\geq 2B$.

Tiles are constructed by choosing a channel time series, and summing the squares of the samples for some duration of time in `gstlal_excesspower`'s "square-adder". A tile that spans the frequencies $f_1 \leq f < f_1 + B$ and times $t_1 \leq t < t_1 + T$ can thus be expressed using the samples from the time zeries $z_j(f_1, B)$. The use of the excess power statistic is considered next, along with the output of `gstlal_excesspower`.

6.2.4 Excess Power and Output

The excess power in the input time series is excess in the sense that it stands out relative to the pre-whitened data stream. The *whitened* energy contained in a tile that spans frequencies $f_1 \leq f < f_1 + B$ and times $t_1 \leq t < t_1 + T$ can be expressed as

$$\begin{aligned} E &= \frac{1}{\mu^2(f_1, B)} z(f_1, B) \cdot z(f_1, B) \\ &= \frac{1}{\mu^2(f_1, B)} \sum_{i=0}^{d-1} z^2_{j_1 + (i+\frac{1}{2})\Delta j}(f_1, B) \end{aligned} \tag{6.2.26}$$

where $j_1 = t_1/\Delta t$ and $\Delta j = T/(d\Delta t)$.

As we know from Sec. 6.1, in the absence of a signal E should be a χ^2 distributed random variable with d degrees of freedom. Once the energy of a time-frequency tile has been determined via Eq. 6.2.26, it is necessary to threshold on the energy to select "important" tiles; the threshold is the probability that a tile would be found with at least E from stationary Gaussian noise alone. The higher this probability is for a given tile, the less likely it is to correspond to a gravitational-wave burst.

Tiles that survive the thresholding process are recorded as *events* or *triggers* in a `sngl_burst` table along with several other parameters which include a confidence and SNR. The confidence is minus the natural logarithm of the probability of observing a tile with a whitened energy of E or greater in stationary Gaussian noise,

$$\text{confidence} = -\ln P(\geq E). \tag{6.2.27}$$

Larger confidence corresponds to a tile that one would be less likely to find in stationary Gaussian noise.

For a given event, the quantity that in `gstlal_excesspower` that is associated with "signal" is $E - d$; the expectation value of whitened energy is $\langle E \rangle = d$, and the quantity $E - d$ is the amount of whitened energy that exceeds what is expected from noise alone. In this case, the quantity

$$\eta = \frac{E - d}{d} \tag{6.2.28}$$

has $\langle \eta \rangle = 1$ and $\langle \eta^2 \rangle = 2/d$ for Gaussian noise, and in this sense can be thought of as a ratio involving signal and noise. The quantity that is conventionally stored as SNR by `gstlal_excesspower` is $\sqrt{\eta}$. However, the $\sqrt{\eta}$ is not a true estimator of the optimal SNR of a signal, which is defined as

$$\rho^2_{\text{opt}} = 4 \int_0^\infty \frac{|\tilde{h}(f)|^2}{\mathcal{S}(f)} df. \tag{6.2.29}$$

where $\mathcal{S}(f)$ is the noise spectral density. Note that this definition assumes a *one-sided* noise spectral density.

The final quantity recorded by `gstlal_excesspower` is an amplitude, which called the root sum square strain or h_{rss}. The h_{rss} characterizes the total gravitational-wave

amplitude at the Earth and is given by

$$h_{rss}^2 = \int \left[|h_+(t)|^2 + |h_\times(t)|^2\right] dt. \tag{6.2.30}$$

The amplitude that is actually received at the detector is modulated by the detector antenna patterns F_+ and F_x, i.e.

$$h_{\rm det}(t) = F_+ h_+(t) + F_\times h_\times(t), \tag{6.2.31}$$

and one can thus define the quantity

$$\begin{aligned} h_{\rm rss,\,det}^2 &= \int |h_{\rm det}(t)|^2 dt \\ &= \int \left[F_+^2 h_+^2(t) + F_\times^2 h_\times^2(t)\right] dt, \end{aligned} \tag{6.2.32}$$

which represents the root sum square strain at the detector. This quantity is related to the optimal value of the SNR $\rho_{\rm opt}$ defined in Eq. (6.2.29).

The method conventionally used by `gstlal_excesspower` to estimate the $h_{\rm rss,\,det}$ is to discretize the integral in Eq. 6.2.29 under the assumption that the noise power spectral density is approximately constant over the bandwidth of interest. In this case,

$$h_{\rm rss,\,det}^2 = \frac{\rho^2}{2(f_H - f_L)} \sum_{f_L}^{f_H} \mathcal{S}(f)\Delta f = \frac{\rho^2}{2}\langle \mathcal{S}(f)\rangle \tag{6.2.33}$$

where ρ is an estimator for the optimal SNR, Eq. (6.2.29). It turns out that this choice is not necessarily optimal. A modified definition is presented in the next section.

Events recorded by `gstlal_excesspower` are stored in a `sngl_burst` table, a file of .xml format that is structured to contain tables and various parameters for each table. In many cases, the number of events recovered is quite large. While it is possible that some events correspond to gravitational-wave bursts, noise in the detector is often to blame for the large number of events. To reduce the event rate and assist in parameter determination, a clustering process is used to combine overlapping tiles.

Given a pair of overlapping tiles, the *cluster tile* is the smallest tile that contains the original two tiles. It is constructed to inherit the $h_{\rm rss}$, stored SNR and confidence of the most significant contributor; the most significant contributor is the tile whose boundaries are the SNR^2-weighted (η^2-weighted) average boundaries of the two contributing tiles.

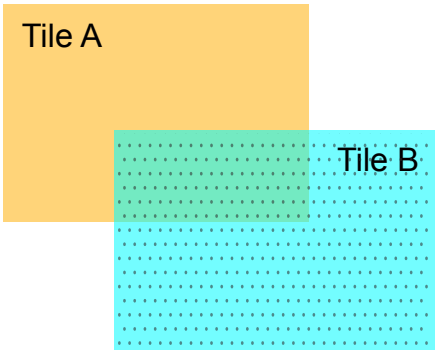

sngl_burst quantity	tile A	tile B
sngl_burst.confidence	C_A	C_A
sngl_burst.amplitude	h_A	h_B
sngl_burst.snr	$\sqrt{\eta_A}$	$\sqrt{\eta_B}$
sngl_burst.peak_time	t_A	t_B
sngl_burst.peak_frequency	f_A	f_B

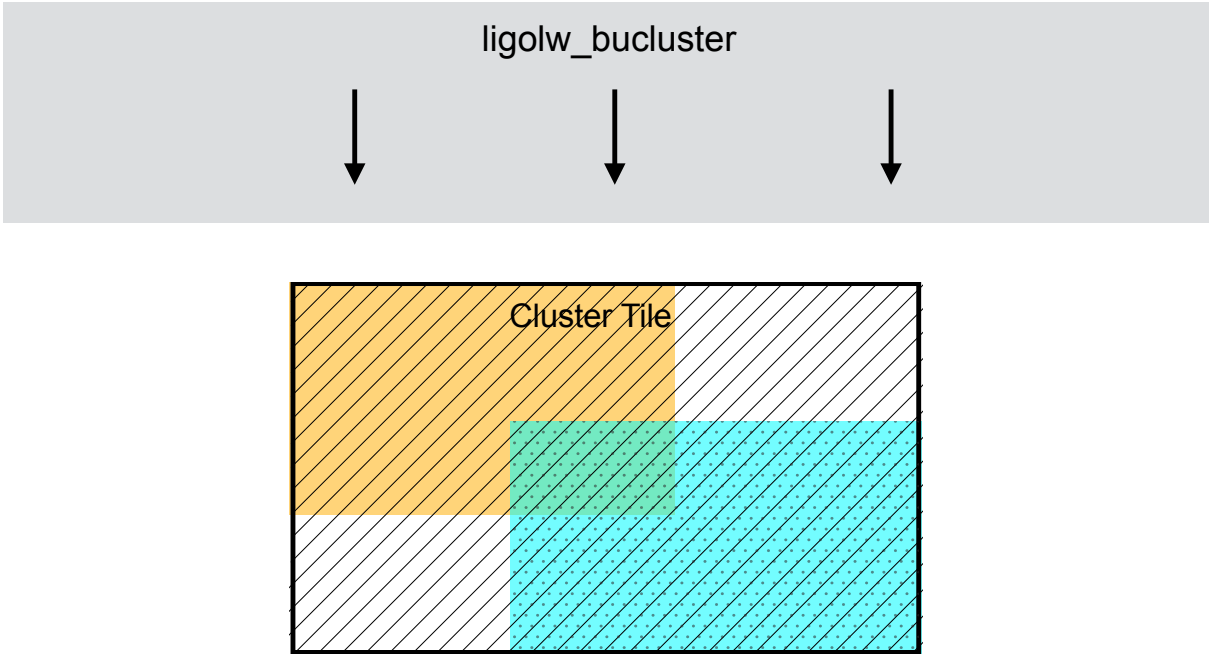

sngl_burst quantity	cluster tile
sngl_burst.ms_confidence	C_B
sngl_burst.ms_hrss	h_B
sngl_burst.ms_snr	$\sqrt{\eta_B}$
sngl_burst.peak_time	$(\eta_A t_A + \eta_B t_B) \,/\, (\eta_A + \eta_B)$
sngl_burst.peak_frequency	$(\eta_A f_A + \eta_B f_B) \,/\, (\eta_A + \eta_B)$
sngl_burst.amplitude	$h_A + h_B$
sngl_burst.snr	$\sqrt{\eta_A + \eta_B}$

Figure 26: Overlapping tiles are clustered using `ligolw_bucluster`. Given two overlapping tiles A (orange) and B (cyan), the algorithm creates a new cluster tile (shaded region), which is the smallest tile that contains both original tiles. The most significant contributor is chosen from the tile whose boundaries are the SNR^2-weighted average boundaries of the contributing tiles. In this schematic, tile B is assumed to be the most signifiant contributor, and the quantities stored in the `sngl_burst` table are adopted from this tile. In addition to these parameters, several new "most significant" quantities (e.g., *ms_snr*, *ms_hrss*, etc.) are added to the `sngl_burst` table.

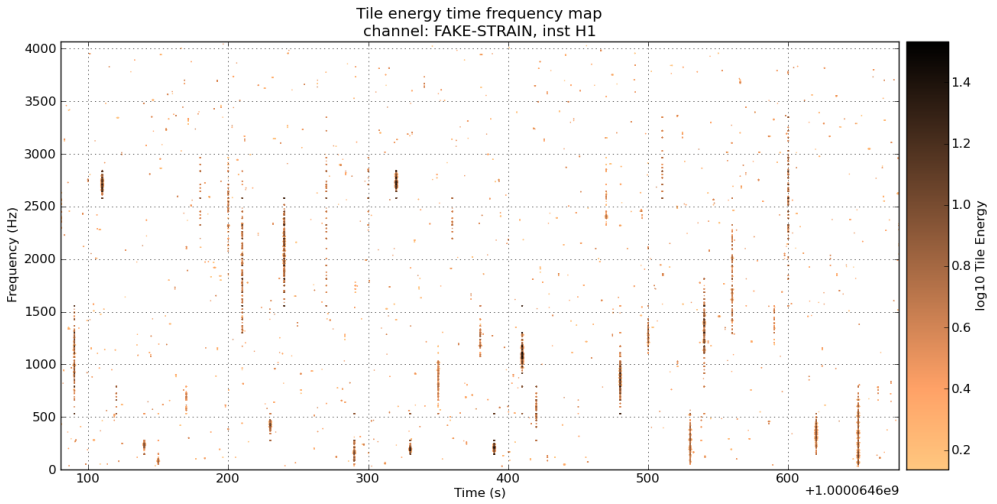

(a)

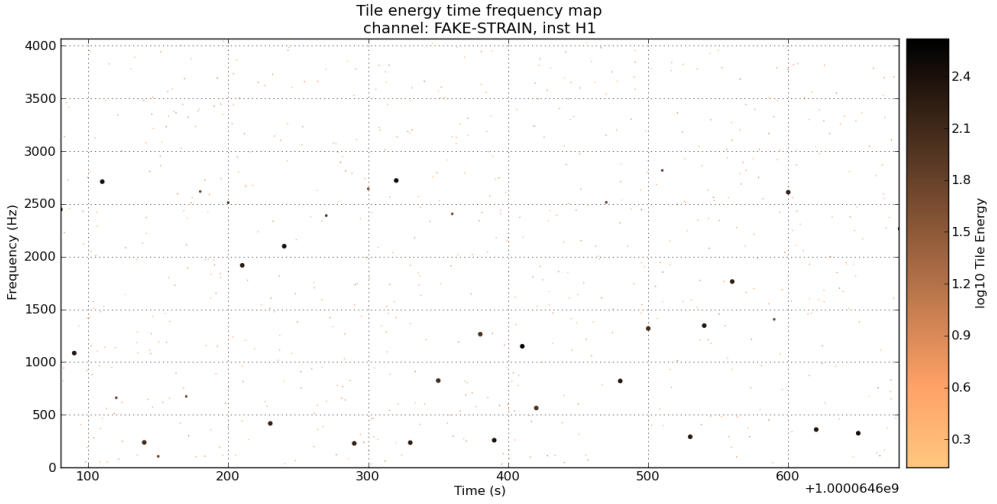

(b)

Figure 27: Clustering greatly reduces the event rate, as seen in these before-clustering (top) and after-clustering (bottom) plots. Here round markers are used to indicate the location of tiles, and the markers are colored by tile energy.

The most significant $h_{\rm rss}$, SNR and confidence adopted by the most significant contributor are stored in the `sngl_burst` table in separate columns from the clustered tile's $h_{\rm rss}$ and SNR, which are determined by summing the $h_{\rm rss}$ and SNR (in quadrature) of the original tiles. It is worth pointing out that tile overlap is *not* considered in this process, which results in over counting these parameters; this is needed to continue carrying information needed to compute the SNR^2-weighted peak time and frequencies. A most significant bandwidth and duration are also determined by choosing the smallest band and interval,

respectively, containing the originals.

Fig. 26 shows a schematic of the clustering process, which is completed using `ligolw_bucluster`, and shows how the contents of the `sngl_burst` table change according to the clustering algorithm.

The effect that clustering has is significant on the number of events stored in `sngl_burst`; this is illustrated in Fig. 27 which plots time-frequency tiles that are stored as `sngl_burst` events before and after clustering.

6.3 Signal detection in a single-interferometer

As mentioned in the last section, a typical search for burst cgravitational-wave signals with `gstlal_excesspower` results in a very large number of events in the `sngl_burst` table. For single-interferometer searches, part of the gravitational-wave problem involves determining which recorded events should be called "detections". One way to tackle this problem is to threshold on some statistic. In this case, the statistic for an event recorded by ExcessPower is compared to the threshold value; if the statistic exceeds the threshold value, a detection is said to occur. If not, a non-detection occurs.

Suppose that some signals are injected into data and processed by a search pipeline, outputting a `sngl_burst` table conaining events. Some of these events will (ideally) correspond to injected signals, but some will also correspond to noise. To determine whether a given event is a gravitational-wave signal or noise, some sort of threshold must be used; events that pass the threshold test can be referred to as detections, while events that fail the threshold test are referred to as non-detections. Several outcomes are possible for a given event:

1. The event passes the threshold test, and corresponds to an injected signal. In this case, the result is a *true-positive*.

2. The event fails the threshold test and is labeled a non-detection, but actually corresponds to an injected signal. In this case the result is a *false-negative*.

3. The event fails the threshold test and corresponds to noise; the event is labeled a non-detection. This result is called a *true-negative*.

4. The event passes the threshold test, but corresponds to noise. This result is referred to as a *false-positive*.

The choice of threshold obviously has a large impact on the outcome of a signal injection search. Setting the threshold relatively low will result in more signals being detected, but comes at the price of many false alarms. Setting the threshold relatively high reduces the number of false alarms, but results in more missed signals. To design an optimal search, it is useful to answer the question, "how many false-positives can be tolerated to reduce the number of false-negatives that occur?"

The Receiver Operating Characteristic (ROC) curve can be used to probe this question. This curve plots the percent of signals reported as positive against the percent of non-signals erroneously reported as positive, for various threshold values. Put differently, the ROC curve plots *efficiency* versus the *false-alarm probability*. This curve thus quantifies the tradeoff of true-positive versus false-positive results.

The ROC curve can be compared to probability density functions for the target population (gravitational-wave signals) and the non-target population (non-signals). This is shown in Fig. 28; sample signal and non-signal probability density functions are shown with a particular choice of threshold (vertical line) next to the point represented in the corresponding ROC curve. Each point on the ROC curve corresponds to a specific threshold choice (or a single possible trade-off between true and false positives). Points can be compared to what would be achieved through pure guessing (efficiency = FAP). The ROC curve thus serves as an excellent indicator of performance.

One way to illustrate the ROC curve is to perform a pipeline efficiency test. Pipeline efficiency tests aim to identify how well a given pipeline can recover signals with known parameters. To do this, signals with the desired parameters are injected into appropriate background noise and analyzed by the pipeline. Roughly speaking, an efficient pipeline is one that recovers the injected signals well. In the remainder of this section, the efficiency of `gstlal_excesspower` is probed using injections, and a ROC curve is produced.

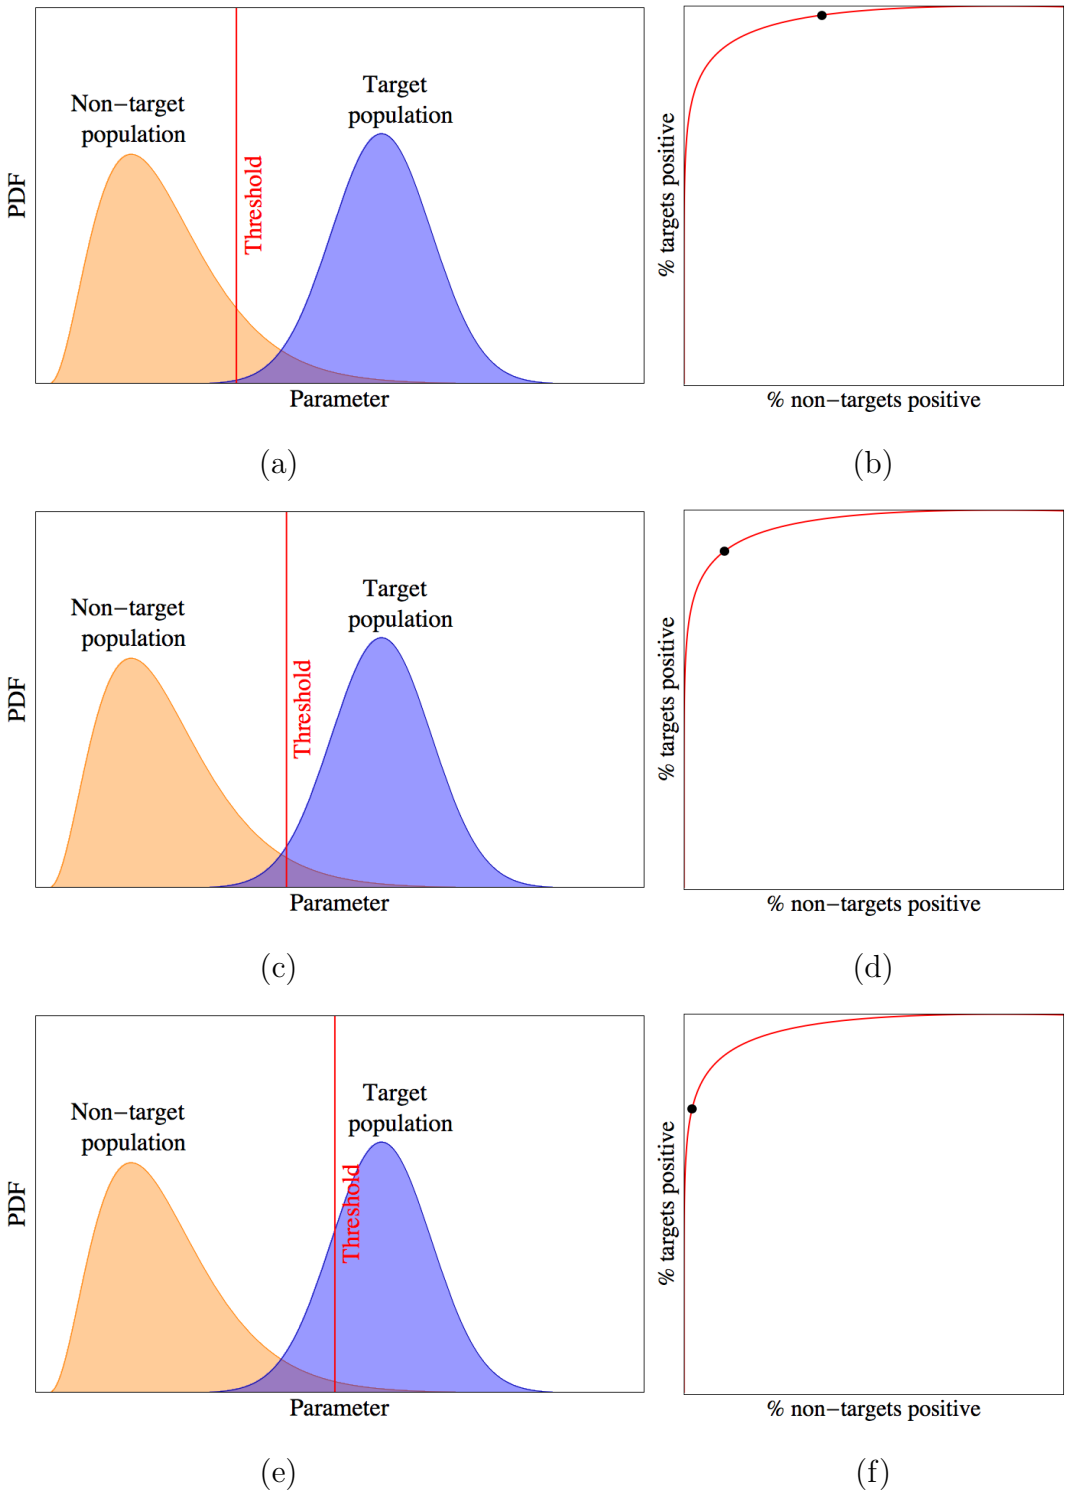

Figure 28: Relationship between target (signal) and non-target (non-signal) pdfs, parameter threshold and ROC curves. The left figure in each panel displays pdfs for non-target (orange) and target populations (blue) along with a parameter threshold (red vertical line). The right figure in each panel displays the ROC curve. Each choice of threshold specifies a point on the ROC curve. In (a), where the threshold is fairly low, more target signals are found in (b), but the false-alarm is higher; in (e), the threshold is relatively high, resulting in a lower false-alarm rate in (f) but fewer target signals recovered. The center panel lies between these two extremes.

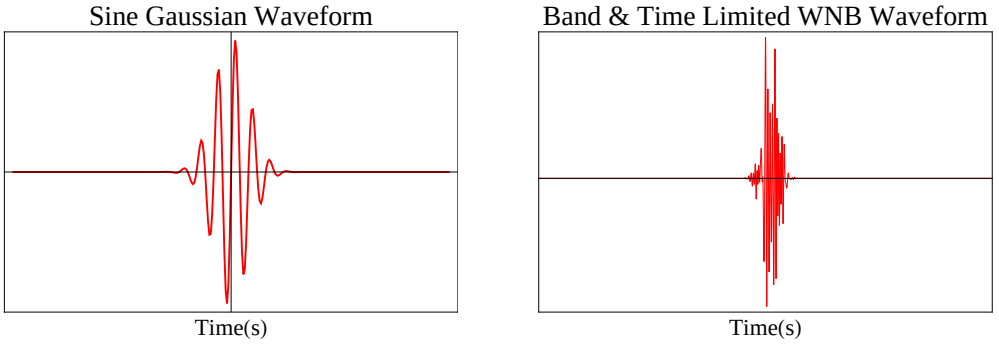

(a) (b)

Figure 29: Sample sine Gaussian (a) and band-time limited white noise burst (b) signals are shown in the time domain, with signal amplitude along the vertical axis.

6.3.1 Injected waveforms

To study the performance of `gstlal_excesspower` as a gravitational-wave burst pipeline, a set of mock gravitational-wave burst signals were injected into one week of simulated aLIGO data. Two signal waveforms were used: sine-Gaussians, which are sine signals modulated by Gaussian envelopes,

$$h_{\mathrm{SG}}(t) = h_0 \sin(2\pi f_0 t) e^{-t^2/\tau^2}, \tag{6.3.1}$$

and band-and-time limited white noise bursts (BTLWNBs). Note that the sine-Gaussian waveform is characterized by a frequency and Q-number where $Q = \sqrt{2}\pi\tau f_0$, while the BTLWNBs are characterized by a frequency band and duration. Two sample waveforms are illustrated in Fig. 29). These two families of signals have the advantage that they can be used as surrogates for nearly any practical signal (including non-astrophysical signals, discussed later on in Sec. 6.4).

	Sine Gaussian	BTLWNB
Number of injections	20137	20137
Central frequency	40 Hz – 4000 Hz	50 Hz – 1500 Hz
Bandwidth range	N/A	50 Hz – 1500 Hz
Q number	3-9	N/A
SNR (ρ_{opt}) distribution	Uniform, $\mu = 3 - 50$	Uniform, $\mu = 3 - 50$

Table 3: Parameters for the ETG mock data challenge injection waveforms. A total of 20000 injections were created, and injected into the data at 1/30s. Note that signals were scaled to have SNRs with the values shown above, based on the optimal SNR defined in Eq. (6.2.29)

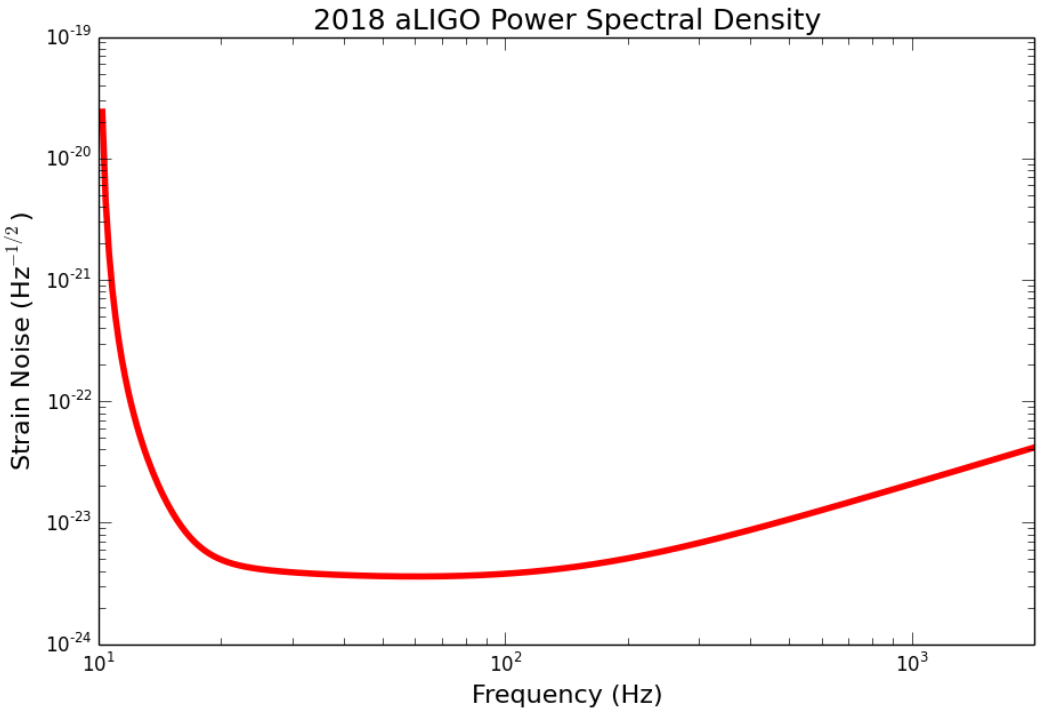

Figure 30: Simulated 2018 aLIGO power spectral density used in the ETG mock data challenge. The simulated PSD provides the noise power spectrum for aLIGO under the high-power broad-band signal recycling (no detuning of the signal recycling cavity), and includes thermal and quantum noise.

The injections used in this study comprise an ongoing event trigger generator (ETG) mock data challenge (to be published in McIver et al. [134]) and were established to have

the parameters shown in Table 3, where the SNR is defined via Eq (6.2.29). The injections were created as elements of a `sim_burst` table within an injection file; to process the injections plus noise within `gstlal_excesspower`, the injections were added to "FAKE-STRAIN" channel gravitational-wave frame files containing the estimated 2018 aLIGO power spectral density. A plot of this power spectral density is shown in Fig. 30. Because the highest frequency component of any single injection did not exceed 4000 Hz, the data were processed at a sampling rate of 8192 Hz.

6.3.2 Determination of signal amplitude

There is a subtle problem with the method used by `gstlal_excesspower` to estimate recovered signal amplitude $h_{\mathrm{rss,\,det}}$. This method, shown in Eq. (6.2.33), involves an assumption that the PSD is approximately constant over the bandwidth of interest. In practice, this is often true for signals of relatively high frequency where the PSD is fairly flat (see Fig. 30), but fails to approximate the $h_{\mathrm{rss,\,det}}$ well for signals with bandwidths over which the PSD is changing more rapidly. In this section, we derive an alternative expression to estimate the $h_{\mathrm{rss,\,det}}$ that takes this factor into consideration.

The $h_{\mathrm{rss,\,det}}$ was defined in Eq. (6.2.32) as

$$\begin{aligned} h^2_{\mathrm{rss,\,det}} &= \int_{-\infty}^{\infty} |\tilde{h}_{\mathrm{det}}(t)|^2 dt = \int_{-\infty}^{\infty} |\tilde{h}_{\mathrm{det}}(f)|^2 df \\ &= 2\int_{0}^{\infty} |\tilde{h}_{\mathrm{det}}(f)|^2 df \end{aligned} \tag{6.3.2}$$

where the last equality on the first line comes from Parseval's theorem. With actual data, the range of frequencies probed is not $0 \leq f \leq \infty$ but some finite bandwidth $f_L \leq f \leq f_H$, and the integral in Eq. (6.3.2) can be replaced by

$$h^2_{\mathrm{rss,\,det}} = 2\int_{f_L}^{f_H} |\tilde{h}_{\mathrm{det}}(f)|^2 df \tag{6.3.3}$$

If the signal doesn't change too much over the bandwidth of interest (i.e., if $\tilde{h}(f)$ is roughly constant over the bandwidth), the product $|\tilde{h}_{\mathrm{det}}(f)|^2$ can be written approximately as

$$|\tilde{h}_{\mathrm{det}}(f)|^2 \approx H^2 \tag{6.3.4}$$

where H is some constant. In this case, the integral in Eq. (6.3.3) can be evaluated

exactly using the Fundamental Theorem of Calculus, and the $h_{\mathrm{rss,\,det}}$ is given by

$$h^2_{\mathrm{rss,\,det}} = 2\int_{f_L}^{f_H} |\tilde{h}_{\mathrm{det}}(f)|^2 df = 2H^2(f_H - f_L). \tag{6.3.5}$$

Under the approximation used in Eq. (6.3.4), the expression for the SNR given by Eq. (6.2.29) can be rewritten as

$$\begin{aligned} \rho^2 &\approx 4H^2 \int_{f_L}^{f_H} \frac{1}{\mathcal{S}(f)} df \\ &= 2\frac{h^2_{\mathrm{rss}}}{(f_H - f_L)} \int_{f_L}^{f_H} \frac{1}{\mathcal{S}(f)} df. \end{aligned} \tag{6.3.6}$$

which can be simplified to

$$\rho^2 = 2h^2_{\mathrm{rss}} \left\langle \frac{1}{\mathcal{S}(f)} \right\rangle. \tag{6.3.7}$$

Solving for $h_{\mathrm{rss,\,det}}$ in Eq. (6.3.7) yields

$$h^2_{\mathrm{rss}} = \frac{\rho^2}{2} \frac{1}{\langle 1/\mathcal{S}(f) \rangle}, \tag{6.3.8}$$

which, although similar in form to Eq. (6.2.33), is different; in general,

$$\langle X(f) \rangle \neq \frac{1}{\langle 1/X(f) \rangle} \tag{6.3.9}$$

for an arbitrary function $X(f)$. The value of ρ^2 in Eq. (6.3.8) is an estimator for the optimal SNR in Eq. (6.2.29) and has been chosen to be $\rho^2 = E - d$. This choice will be discussed in the following section.

The amplitude definition in Eq. (6.3.8) more effectively approximates the $h_{\mathrm{rss,\,det}}$ in regions where the PSD is changing. For example, consider the sine Gaussian signal with central frequency ~ 50 Hz that was injected into simulated aLIGO data and recovered by `gstlal_excesspower`. Table 4 compares the injected $h_{\mathrm{rss,\,det}}$ value with the $h_{\mathrm{rss,\,det}}$ values obtained by Eq. (6.2.33) and Eq. (6.3.8). It is clear that Eq. (6.3.8) estimates the injected amplitude much more closely than Eq. (6.2.33).

Injected $h_{\mathrm{rss,\,det}}$	$h_{\mathrm{rss,\,det}} \propto 1/\langle 1/\mathcal{S}(f) \rangle$	$h_{\mathrm{rss,\,det}} \propto \langle \mathcal{S}(f) \rangle$
3.6×10^{-21}	4.2×10^{-21}	4.1×10^{-18}

Table 4: A signal with known injected $h_{\mathrm{rss,\,det}}$ (given in the first column) was recovered by `gstlal_excesspower`. Its $h_{\mathrm{rss,\,det}}$ was then estimated using the recovered SNR and Eq. (6.3.8 and Eq. (6.2.33), respectively.

6.3.3 Recovered signal parameters for the one week ETG mock data challenge

Post-processing of the outputs from `gstlal_excesspower` involves several steps beyond handing frame files to the pipeline handler. First, `sngl_burst` events — or triggers — are clustered to reduce the event rate. The file containing the clustered triggers is then conglomerated with the `sim_burst` table containing injections, and injections are matched to corresponding `sngl_burst` events by event identification information that is stored in both `sngl_burst` and `sim_burst` tables. This step thus crucially identifies which events in `sngl_burst` match an injected signal exactly. The results from one week of data are summarized in figures of merit (FOMs). Of 40274 injections, 33974 — or about 85% — were recovered. This is shown in Fig. 31, which plots the frequencies of found injections (black +) and missed injections (red x) versus time.

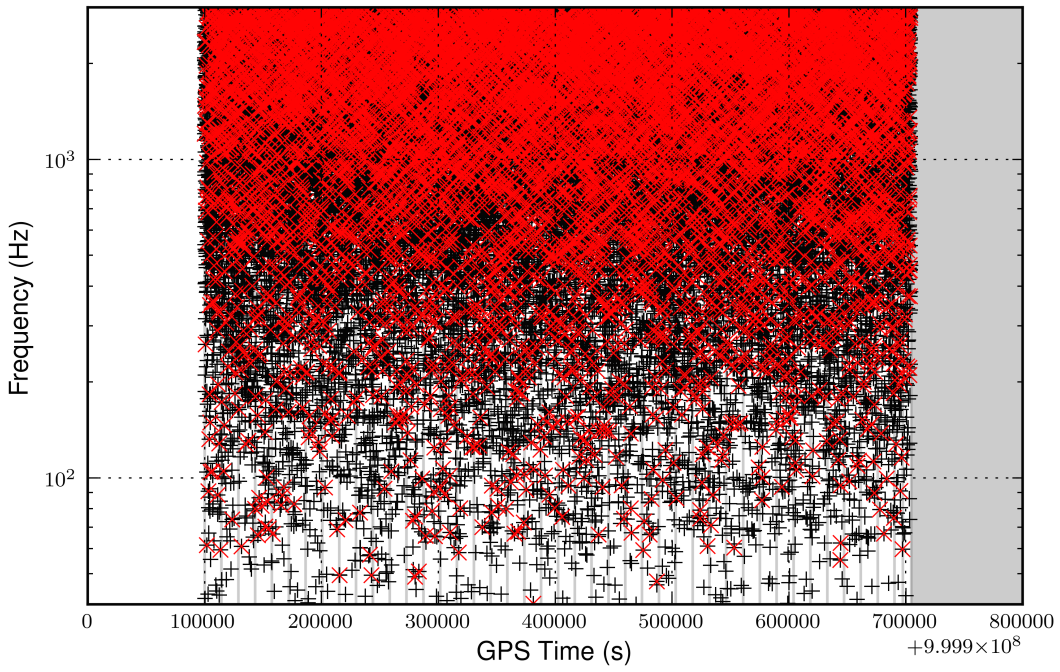

Figure 31: Found (black +) and missed (red ×) injections for the one week ETG mock data challenge as a function of GPS time.

The recovered $h_{\mathrm{rss,\,det}}$, determined using the estimate in Eq. (6.3.8), is plotted against the injected $h_{\mathrm{rss,\,det}}$ in Fig. 32 and indicates that the distribution follows a shape roughly expected.

Fig. 33 shows recovered most significant SNR ($\sqrt{\eta}$) versus injected SNR (ρ_{opt}), which was defined in Eq. (6.2.28). The SNRs $\sqrt{\eta}$ recovered using the most significant value of Eq. (6.2.28) significantly underestimate the injected SNR ρ_{opt}. To optimally estimate ρ_{opt}, it appears that a quantity other than the most significant SNR stored in the `sngl_burst` table should be chosen or calculated.

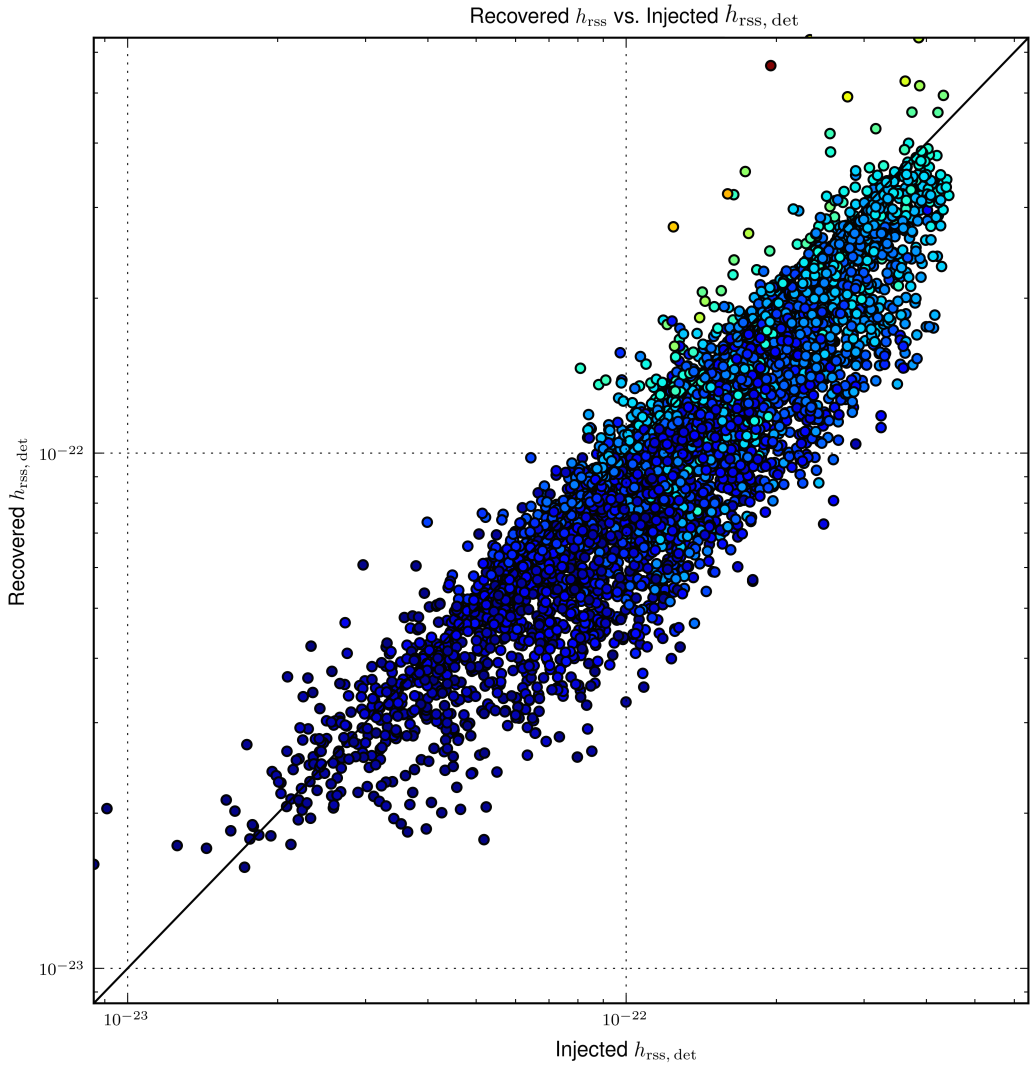

Figure 32: The estimated recovered $h_{\mathrm{rss,\,det}}$ is plotted as a function of injected $h_{\mathrm{rss,\,det}}$ for injections found in the one week ETG mock data challenge. The diagonal line represents values at which the recovered $h_{\mathrm{rss,\,det}}$ exactly equals injected $h_{\mathrm{rss,\,det}}$. The recovered $h_{\mathrm{rss,\,det}}$ values were estimated using Eq. (6.3.8); some of the recovered values slightly underestimate the $h_{\mathrm{rss,\,det}}$, but the distribution has the form roughly expected.

There is an alternative way to think about the SNR. If the data measured is expressed in the frequency domain as $\tilde{s}(f) = \tilde{h}(f)+\tilde{n}(f)$, where $\tilde{h}(f)$ is the gravitational-wave signal and $\tilde{n}(f)$ is the noise, then the energy in the signal can be written as

$$\begin{aligned} E &= 4\int_0^\infty \frac{|\tilde{s}(f)|^2}{\mathcal{S}(f)}df \\ &= 4\int_0^\infty \frac{|\tilde{h}(f)+\tilde{n}(f)|^2}{\mathcal{S}(f)}df \end{aligned} \tag{6.3.10}$$

Ignoring any cross-terms of the form $\tilde{h}(f)\tilde{n}(f)$ (since $\langle\tilde{n}(f)\rangle = 0$), we can then write

$$\begin{aligned} \langle E\rangle &= 4\int_0^\infty \frac{|\tilde{h}(f)|^2}{\mathcal{S}(f)}df + 4\int_0^\infty \frac{\langle|\tilde{n}(f)|^2\rangle}{\mathcal{S}(f)}df \\ &= \rho^2 + d. \end{aligned} \tag{6.3.11}$$

The first term on the right-hand side of in Eq. (6.3.11) is the optimal SNR $\rho_{\rm opt}$. The second term can be be thought of as the power in the noise alone. The power in the noise alone is distributed as a χ^2 distribution with d degrees of freedom, so this term averages

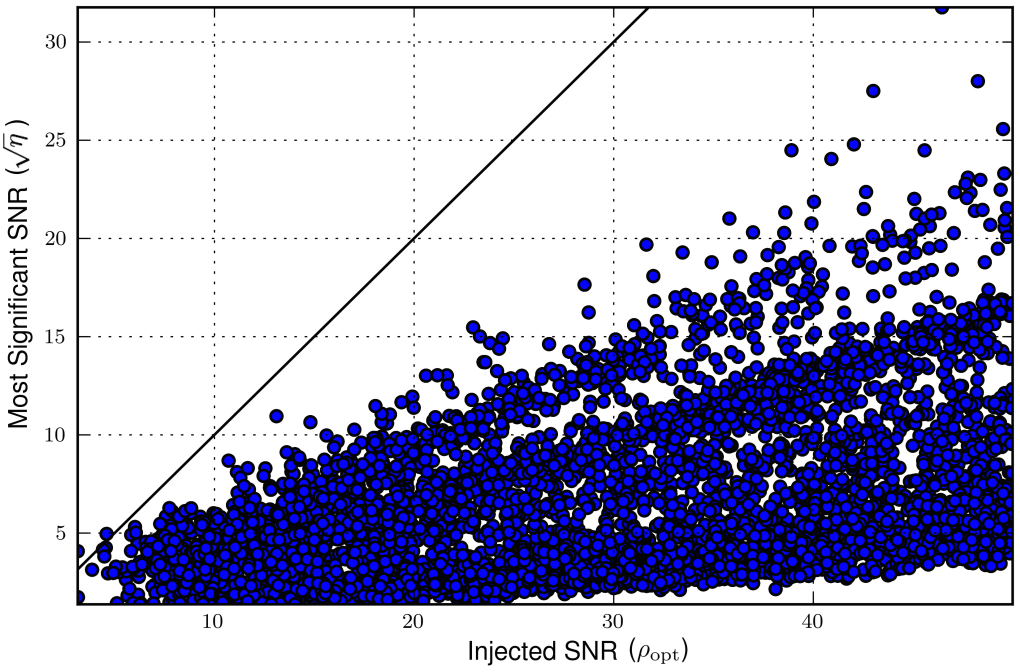

Figure 33: Recovered most significant SNR, defined by $\sqrt{\eta}$ in Eq. (6.2.28), is plotted against the injected SNR $\rho_{\rm opt}$. The solid black line indicates optimal recovery when recovered SNR equals injected SNR. When recovered using η, the injected values of the SNR $\rho_{\rm opt}$ are significantly underestimated. This reflects the idea that Eq (6.2.28) should not be used as an estimator for the SNR.

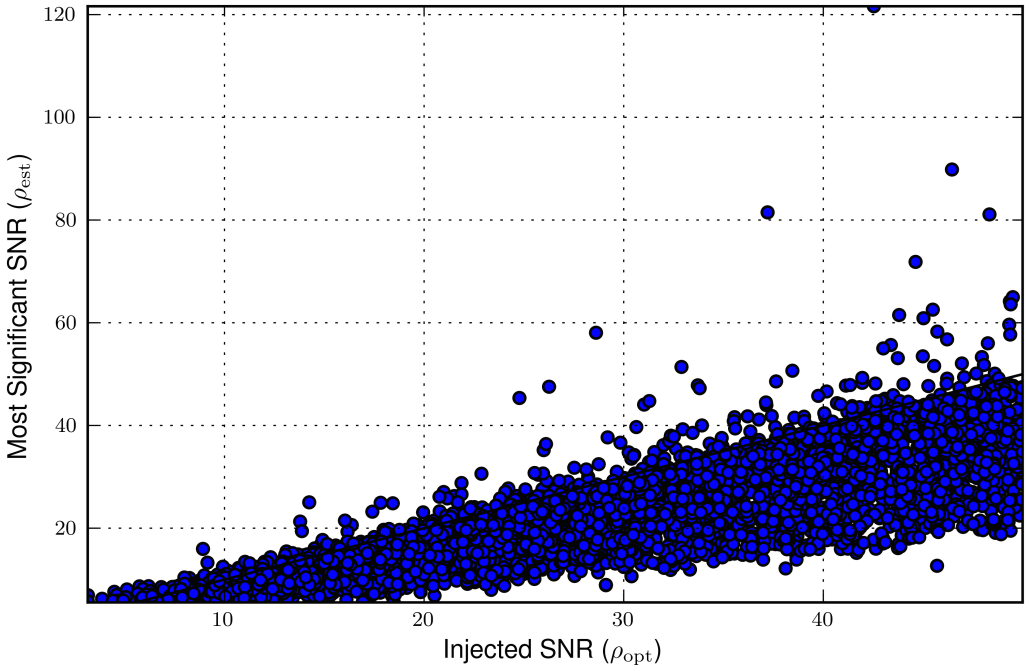

Figure 34: Recovered most significant SNR, defined by $\rho_{\rm est} = \sqrt{E-d}$ in Eq. (6.2.28), is plotted against the injected SNR $\rho_{\rm opt}$. The solid black line indicates points where recovered SNR $\rho_{\rm est}$ equals injected SNR $\rho_{\rm opt}$ exactly. Closeness of points to the solid black line illustrate that Eq (6.3.12) is a better estimator for the optimal SNR $\rho_{\rm opt}$ than $\sqrt{\eta}$.

to d. The optimal SNR can thus be approximated as

$$\rho_{\rm est} \approx \sqrt{E-d} \tag{6.3.12}$$

which differs from the convention in Eq. (6.2.28). The definition in Eq. (6.3.12) proves to be more optimal, as seen in Fig. 34, which forms an improved distribution of values.

6.3.4 ROC curve

A ROC curve was determined for `gstlal_excesspower` using the results just described. As discussed in Sec. 6.3, the ROC curve quantifies the tradeoff between efficiency and FAP for various threshold choices. In `gstlal_excesspower`, confidence is used as the threshold. Each event recovered by `gstlal_excesspower` stores a value of confidence in the `sngl_burst` table; this value is determined using a percent point function, which is effectively an inverse distribution function. In the absence of any signal, the power in the noise should be distributed as a χ^2 distribution with N degrees of freedom.

The percent point function takes the tile energy and degrees of freedom corresponding to an event and returns a value of confidence.

The data that were processed and produced results in Sec. 6.3.3 consisted of injections along with simulated 2018 aLIGO noise. In addition to this data, `gstlal_excesspower` was used to process the noise alone. Events that were recovered from noise alone were used to determine the FAP. Since each noise event had an associated confidence value, each value of the FAP uniquely specifies a confidence interval. Efficiency was determined by finding the fraction of injections possessing a confidence that exceeded each confidence interval. The ROC curve, shown in Fig. 35, plots efficiency versus FAP and has an expected shape.

It is worth pointing out that the ROC curve depends significantly upon the injections used in the study, as well as the distribution of the noise. The noise used to construct the ROC curve in Fig. 35 contains estimated quantum and thermal noise for the 2018 aLIGO detector, but it is likely that real detector noise will contain other components, including non-Gaussian noise components, that would modify the ROC curve presented here.

Several extensions of this work would be useful. It would be interesting to plot the ROC curve for different injection families, which could yield insight into the efficiency of recovering signals of various morphology. The ROC curve could also be reproduced for different simulated noise, and could help determine which components of noise affect gravitational-wave searches the most.

6.4 Detector characterization

Initial LIGO's sixth science run (S6) took place from July 2009 to August 2010, with data collected by the Livingston (L1), Hanford (H1), GEO600 and Virgo(V1) detectors. The data were searched for gravitational waves from a variety of sources, including unmodeled bursts [135, 136]. Searches such as these are limited by the ability of detection algorithms to distinguish astrophyscially induced signals from transient noise bursts in the detectors.

Transient noise bursts — also known as glitches — can arise from many potential sources, including environmental, mechanical, or electronic components of the detectors,

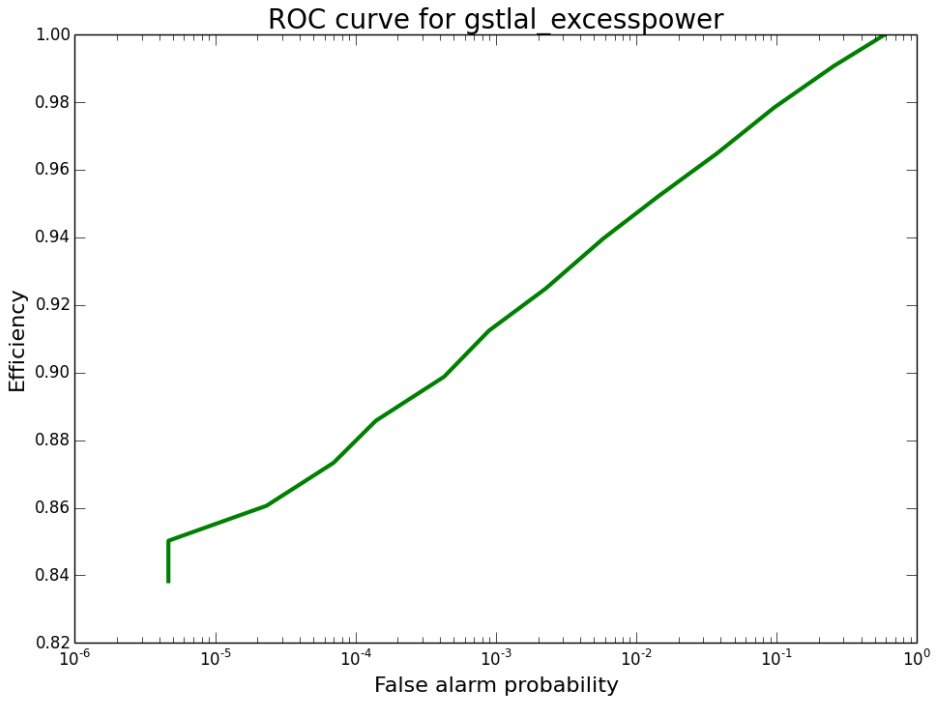

Figure 35: A ROC curve of efficiency versus FAP produced for `gstlal_excesspower`, shown on a semilog scale.

and can appear in both the gravitational-wave strain channel as well as the hundreds of *auxiliary channels* comprising various detector subsystems. In addition to potentially mimicking gravitational-wave burst signals, noise transients can cause loss of lock in the detectors, resulting in less observational time and poorer data quality. There is thus a large demand for *detector characterization* efforts, which seek to characterize the noise inherent to the LIGO detectors and develop data quality tools.

6.4.1 `gstlal_excesspower` and detector characterization

Because the excess power method is suitable for searching any type of data for a signal of unknown shape, it can be utilized not only to perform astrophysically motivated searches for gravitational-wave bursts, but also to aid in detector characterization efforts. To avoid confusion between glitches and a possible gravitational-wave signal, detector characterization efforts are typically limited to the *auxiliary* channels: channels that contain data from sensors monitoring instrumental and environmental variables such as seismic motion, laser beam alignment, mirror suspension and control, etc.

It is important to note that while the auxiliary channels do not contain gravitational-wave signals, couplings between the auxiliary channels and the gravitational-wave strain channel (the channel that *should* contain any gravitational-wave signal present in the data) have been observed, and it is possible for noise transients in the auxiliary channels to manifest in the gravitational-wave strain channel where gravitational-wave signals are expected [137]. The complexity of the many subsystems and data channels is high, and much is still unknown about the couplings between various channels and glitches.

Noise transients in the LIGO detectors can be characterized by a variety of parameters. In general, it is useful to group noise transients into populations based on their "loudness" or SNR, central frequency, and spectral morphology. Spectrograms, which visually display the frequency spectrum as a function of time (or some other parameter), are often used to identify noise transients and the populations to which they belong.

Existing gravitational-wave search pipelines are affected differently by noise transients. Searches that use matched filtering tend to be affected by noise transients that mimic the waveform templates. Unmodeled searches, such as `gstlal_excesspower`, are affected by the frequency of noise transients occurring; despite demanding coincidence in events between different detectors, glitches occur frequently enough that coincidence alone cannot eliminate the problem.

Beyond demanding coincidence in data between different detectors, two methods (discussed in detail in McIver et al. [137]) can be used to mitigate noise transients. The first involves identifying the source of various noise transients. In practice, such sources are difficult to determine and can range from environmental events to equipment malfunction. The properties or characteristics of a given noise transient become are very useful in determining its origin.

The second method to mitigate noise transients is to develop *data quality flags* which are used to veto segments of data that are "suspicious". This is necessary because even when the source of a given noise transient has been identified, it is not always possible to correct it. For example, if a noise transient is caused by a malfunctioning piece of hardware, it may not be possible (or may take a long time) to fix the hardware. There are several categories of data quality (DQ) flags:

- *Category 1.* These DQ flags veto segments of data that were taken when the detector had a serious problem or was not operating near nominal sensitivity.

- *Category 2.* These flags veto segments of data that were taken during a time with a well-understood detector problem. For instance, if a particular subsystem is known to cause noise transients that couple into the gravitational-wave strain channel, then segments of data displaying that problem receive a category 2 veto.

- *Category 3.* These DQ flags veto intervals that are characterized by a problem that is not yet completely understood.

The outcome of a DQ flag depends on the particular search method being used, but in general results in the removal of suspect data segments.

6.4.2 Low-frequency mock data challenge

With the era of Advanced LIGO about to begin, data quality studies are in high demand. There is a need to develop new vetoing tools and identify the sources of problematic noise transients, as discussed in the last section. To help fulfill this need, efforts to test and "tune" `gstlal_excesspower` for detector characterization are underway.

A first proposed task is to complete an efficiency test similar to that discussed in Sec. 6.3, but with injections designed to mimic the types of noise transients currently appearing in auxiliary channel data, and with a power spectral density appropriate to the channels chosen. A set of proposed injection parameters has been developed, based on the seismic and suspension subsystems of the detector; these parameters are shown in Table 5. The most notable difference between the detector characterization injection parameters and the astrophysical injection parameters in Sec. 6.3 is the variation in signal frequency; the injections in Table 5 are meant to represent the low-frequency noise transients that appear int he seismic and suspension auxiliary channels.

Future work will involve analyzing these injections with `gstlal_excesspower`, and making needed changes within the pipeline to tune it as an optimal search tool for low-frequency transients. It can then be used to identify sources of problematic noise transients and to develop new DQ tools for the Advanced LIGO era.

	Sine Gaussian	BTLWNB
Number of injections	10000	1000
Central frequency	0.10 Hz – 10 Hz	0.10 Hz – 10 Hz
Bandwidth range	N/A	0.75 Hz – 5.00 Hz
Q number	3-9	N/A
SNR (ρ_{opt}) distribution	Gaussian, $\mu = 4$, $\sigma = 3.5$	Gaussian, $\mu = 4$, $\sigma = 3.5$

Table 5: Parameters for the low-frequency ETG mock data challenge injection waveforms.

Chapter 7

Summary

"*Organize, agitate, educate, must be our war cry.*"

— Susan B. Anthony (1820-1906)

Throughout the past century, a number of dramatic scientific discoveries have expanded human knowledge of the universe far beyond Earth's boundaries. The direct detection of gravitational waves, anticipated within the next several years, promises to add to this legacy of science. Gravitational-wave astronomy will completely change the paradigm of modern astronomy, and holds the potential to unlock hidden mysteries of the cosmos. This dissertation has presented a series of works that probe gravitational-wave detection efforts across a wide range of frequencies.

Part I focused on efforts to detect a stochastic background of gravitational waves with PTAs. In Chapter 3, the optimal detection statistic for PTAs was presented in the time-domain. The methodology presented takes into account two challenges in PTA data analysis: the irregular sampling typical of PTA data, and the use of a timing model to predict the times-of-arrival of radio pulses. As a result, the time-domain optimal statistic is better suited to gravitational-wave data analysis than its frequency domain counterpart. This statistic is a robust tool for making quick estimates of gravitational-wave amplitude, creating software injections, and probing scaling laws for the SNR in PTA experiments. The statistic does have drawbacks that make it an incomplete substitute for a full Bayesian analysis, however.

In Chapter 4, isotropic stochastic gravitational-wave background detection with PTAs

was probed in the context of the optimal detection statistic and metric theories of gravity. Although general relativity is considered one of the most successful theories in physics, problems in cosmology and high energy physics have spurred the development of other viable theories of gravity. To robustly test these theories against general relativity, gravity must be tested in a strong-field regime. Gravitational waves provide a means to do so.

While allowing for any possible gravitational wave polarization mode, the expected PTA cross-correlations for a stochastic background were determined. The expected cross-correlations are proportional to a function called the overlap reduction function, a factor that characterizes the response of the PTA to gravitational waves based on the pulsar-Earth geometry. The overlap reduction function has a different form for each gravitational-wave polarization and turns out to be frequency-dependent for non-transverse gravitational waves.

In numerically evaluating the overlap reduction functions for each possible gravitational-wave polarization, it was found that for "nearby" pulsar pairs — pulsar pairs with small angular separations on the sky — PTAs have increased sensitivity to the scalar-longitudinal and vector polarization modes that are present in some alternative theories of gravity. In fact, for nearby pulsar pairs PTAs are more sensitive to non-transverse gravitational waves than transverse gravitational waves by several orders of magnitude. This has interesting implications for testing general relativity, since three of the four additional polarization modes possible in a metric theory of gravity correspond to non-transverse gravitational waves; this work suggests that PTAs are particularly well-suited to testing general relativity. In Section 7.1, I outline a few directions in which this work can be extended to perform tests of general relativity.

In part II, gravitational-wave detection was probed in the context of ground-based laser interferometers. In Chapter 6, an excess power statistic was described as a tool for use in gravitational-wave burst searches. The implementation of this statistic into the `gstlal_excesspower` pipeline was laid out in detail. An effort to improve the `gstlal_excesspower` vis-a-vis signal amplitude recovery was presented, and used to estimate the amplitude of signals recovered in one week of simulated data containing injections. The results of the one week ETG mock data challenge presented in Chapter 6

also provided an opportunity to test a new method for estimating the optimal SNR, which was shown to perform significantly better at recovering the optimal SNR than the conventional method used by `gstlal_excesspower`.

The efficiency of `gstlal_excesspower` was probed at the end of Chapter 6 by producing a ROC curve, which plots efficiency against the false alarm probability. A future objective includes producing families of ROC curves for various parameters, which will illustrate how different parameters affect the performance of the pipeline.

The potential of `gstlal_excesspower` as a tool for detector characterization was discussed at the end of Chapter 6, along with some of the methods currently used in detector characterization to mitigate noise transients that pose a problem in gravitational-wave burst searches.

7.1 Future directions motivated by Chapter 4

The work described in Chapter 4 shows that for stochastic backgrounds of gravitational waves, PTAs receive an enhanced response to non-transverse gravitational-wave polarizations predicted by some theories of gravity. It is worth pointing out that this effect is unique to PTA experiments. Similar work probed in the context of ground-based laser interferometers indicates that they have approximately the same response to all six possible gravitational-wave polarizations [121]. This suggests that PTAs are uniquely well-suited to searches for non-transverse gravitational waves.

Several future lines of research, based on this fact, are planned and/or ongoing.

7.1.1 Using the Deep Space Network to search for non-transverse gravitational waves

Precision pulsar timing capabilities for the Deep Space Network (DSN) are just being finalized [138], with pulsar backends being installed in Australia (to follow in the U.S., shortly) and observations anticipated to begin within a year. In addition to the advantage of having antennas distributed across the globe (which improves PTA sensitivity), the DSN will offer high-cadence observations that lead to dramatic improvements in

gravitational wave detectability with PTAs on 5-10 year timescales. The new DSN pulsar backend system will also have the ability to remove dispersive effects caused by the interstellar medium. This makes it an excellent tool for pulsar timing experiments.

With its new pulsar framework in place, the DSN could be utilized for a novel pulsar timing experiment to test general relativity. Because PTAs are more sensitive to the non-transverse gravitational waves produced in some metric theories of gravity than to those of general relativity (for nearby pulsar pairs), globular cluster (GC) pulsars could be timed specifically to search for non-transverse gravitational waves. In recent years, large numbers of millisecond pulsars with small (many $< 1^\circ$) angular separations have been discovered within GCs. In 2005, over 100 pulsars had been discovered in GCs [139] with 21 millisecond pulsars in the cluster Terzan 5 alone [140], and the last decade has seen the discovery of even more.

One observational advantage afforded by the new DSN pulsar backend lies in its ability to make high cadence observations. Recent results in the literature [104], derived completely in Chapter 3, have probed which observational parameters — the number of pulsars, the cadence of observation, or the white-noise RMS values of the pulsars — lead to the greatest gains in recovered signal-to-noise (SNR). Over large timescales ($T \gtrsim 5 - 10$ yrs), the best strategy to detect a stochastic background of gravitational waves is to increase the number of pulsars in the array. On shorter observational timescales, however, higher observational cadence lends itself to dramatic rises in the SNR (for PTAs with the same number of pulsars). This can be seen in Fig. 12. Additionally, the large number of millisecond pulsars in GC suggests that there need not be a trade-off between the number of pulsars in the PTA and the cadence; with the DSN and GCs, both parameters are likely to contribute to high values of recovered SNR.

To take advantage of gains in SNR due to high cadence, one could established a short (~ 3 yr) GC pulsar timing experiment with the DSN, using GCs with known millisecond pulsar populations. One advantage to PTA experiments is that sensitivity only improves over the observational timescale. Such an experiment could result in the direct detection of a gravitational-wave signal; in the absence of a direct detection, constraining upper

limits could be determined. It is worth pointing out, however, that the proposed experiment has one challenge; the close proximity of GC pulsars means that signals from these pulsars would likely have some correlated noise (in particular, correlations in the effects wrought by the interstellar medium). Extra care would need to be used in searching for a stochastic background from such signals to mitigate the challenges arising from correlated noise due to the interstellar medium.

7.1.2 Generalizing current detection pipelines to search for alternative-theory polarizations

A number of stochastic gravitational-wave pipelines for PTAs are currently under development, including one designed on the principles described in Chapter 3. The optimal cross-correlation statistic-based pipeline serves as an excellent proxy to robust Bayesian searches, and is dramatically less computationally expensive than Bayesian methods.

To generalize this pipeline so that searches for gravitational waves of all six polarizations can take place, the Hellings-Downs curve (that describes the expected correlation for general relativity) must be replaced by the more general overlap reduction functions described in Chapter 4. For transverse gravitational-wave polarizations, these functions depend only on the angular separation of the pulsars and can be described in a closed analytic form; adding these functions to existing pipeline frameworks is trivial. For the case of non-transverse gravitational-wave polarizations, however, the overlap reduction function depends not only on the angular separation of the pulsars, but also on the pulsar-Earth distances and the gravitational-wave frequency.

While pulsar distances are generally poorly constrained, observations of GCs provide more optimistic measurements. The real challenge in incorporating the non-transverse overlap reduction functions into existing pipeline framework comes from their frequency dependence, which must be integrated out. This can be done numerically (at the expense of computational speed), but recent work by Mingarelli et al. [69, 70] probing anisotropic stochastic backgrounds suggests that there may be another way. In the framework Mingarelli et al. devised to study anisotropic stochastic backgrounds, the Hellings-Downs

curve is generalized to an overlap reduction function with a spherical harmonic decomposition. It is possible that this framework may yield more a more tractable form of the non-transverse overlap reduction functions than numerical integration, although it is not immediately clear what the computational cost of doing so might be.

Numerous efforts are also underway to establish pipelines that utilize Bayesian analysis, including work by Stephen Taylor, Rutger van Haasteren and Justin Ellis at the Jet Propulsion Laboratory. I have agreed to co-lead an effort with Stephen Taylor to constrain alternative-theory polarization states with PTAs using developing Bayesian tools. The frequentist cross-correlation statistic pipeline will serve as an excellent proxy to these Bayesian tools, and will help us to make robust constraints on PTA data.

7.1.3 Extracting and/or Disentangling Polarization Content from Observations

As a last step in utilizing current PTA pipelines to test general relativity, I intend to develop a framework to extract polarization content from PTA observations. Current frameworks, such as that laid out in the frequentist optimal cross-correlation statistic pipeline, typically produce an estimate of the stochastic gravitational wave background and a SNR value. While these estimates can be used to place upper limits on gravitational-wave content in the data, they do not provide any information about what gravitational-wave polarizations are present. Disentangling the polarization content is essential to placing stringent constraints on various metric theories of gravity.

Nishizawa et al. (2009) has investigated the feasibility of extracting polarization content from gravitational-wave observations with ground-based laser interferometers [121]. In their approach, they express the optimal detection statistic as a linear combination of contributions from three classes of gravitational-wave polarizations (see Fig. 14): two *tensor* polarization modes $(+, \times)$, two *vector* modes (x, y) and two *scalar* modes (b, l). The constant coefficients multiplying each of the three terms in the linear combination depend on the gravitational-wave source spectrum, which (in principle) can vary from one metric theory to another. To extract the polarization content from this linear combination, the constant coefficients multiplying each term must be determined. This requires

at least three ground-based detectors, although more than three may also be used.

For the case of PTAs, each pulsar-Earth link is effectively the arm of an interferometer, and the framework of Nishiawa et al. could be applied to the optimal cross-correlation statistic for PTAs to extract polarization information from PTA observations. Since there are many more pulsar-Earth pairs than ground-based interferometer arms, it may be possible to optimize the analysis so that certain combinations of pulsars are most sensitive to each class of polarization.

The software injection procedure discussed in Sec. 3.2 could be used to test the method discussed above in PTA data; signals of various gravitational-wave polarization could be injected into data (real or simulated), and the linear combination described above determined. By solving for the coefficients in this linear combination, it should be possible to determine how efficiently content can be extracted.

As a last part of this problem, the polarization content that is extracted (which depends on the source spectrum that produces it) must be related to the source spectra that are possible in various metric theories of gravity. That is, different metric theories of gravity yield different source spectra, so once presented with polarization content, additional work needs to be done to determine what parameters in a given metric theory can be constrained.

Bibliography

[1] R. M. Wald, *General Relativity* (The University of Chicago Press, 1984).

[2] P. Kalmus, *Search for Gravitational Wave Bursts from Soft Gamma Repeaters*, PhD thesis, PhD Thesis, 2009, 2009.

[3] D. Lai and S. Shapiro, (1994).

[4] D. Lai and S. L. Shapiro, Class.Quant.Grav. **442**, 259 (1995).

[5] K. C. B. New, G. Chanmugam, W. W. Johnson, and J. E. Tohline, Class.Quant.Grav. **450**, 757 (1995).

[6] L. Bildsten, Class.Quant.Grav. **501**, L89 (1998).

[7] J. D. E. Creighton and W. G. Anderson, *Gravitational-wave physics and astronomy: an introduction to theory, experiment and data analysis* (Wiley-VCH, Weinheim, 2011).

[8] M. Maggiore, *Gravitational Waves: Theory and experiments*Gravitational Waves (Oxford University Press, 2007).

[9] B. Haskell *et al.*, arXiv.org (2014), 1407.8254v1.

[10] LIGO Scientific, B. Abbott *et al.*, Astrophys.J. **683**, L45 (2008), 0805.4758.

[11] R. A. Hulse and J. H. Taylor, Astrophys. J. Lett. **195**, L51 (1975).

[12] J. M. Weisberg and J. H. Taylor, ASP Conf.Ser. **328**, 25 (2005), astro-ph/0407149.

[13] C. J. Moore, R. H. Cole, and C. P. L. Berry, Class.Quant.Grav. **32**, 015014 (2015).

[14] F. Jenet *et al.*, arXiv:0909.1058 (2009), 0909.1058.

[15] G. H. Janssen *et al.*, European Pulsar Timing Array, in *40 Years of Pulsars: Millisecond Pulsars, Magnetars and More*, edited by C. Bassa, Z. Wang, A. Cumming, & V. M. Kaspi, , American Institute of Physics Conference Series Vol. 983, pp. 633–635, 2008.

[16] R. N. Manchester, The Parkes Pulsar Timing Array Project, in *40 Years of Pulsars: Millisecond Pulsars, Magnetars and More*, edited by C. Bassa, Z. Wang, A. Cumming, & V. M. Kaspi, , American Institute of Physics Conference Series Vol. 983, pp. 584–592, 2008, 0710.5026.

[17] G. Hobbs *et al.*, Classical and Quantum Gravity **27**, 084013 (2010), 0911.5206.

[18] P. R. Paulson, *Fundamentals of Interferometric Gravitational Wave Detectors* (World Scientific, 1994).

[19] A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins, Nature **217**, 709 (1968).

[20] T. Gold, Nature **218**, 731 (1968).

[21] V. L. Ginzburg, Soviet Physics Uspekhi **14**, 83 (2007).

[22] P. A. STURROCK, Nature **227**, 465 (1970).

[23] D. Lorimer and M. Kramer, *Handbook of Pulsar Astronomy*, Cambridge Observing Handbooks for Research Astronomers Vol. 4, 1st ed. (Cambridge University Press, Cambridge, U.K.; New York, U.S.A, 2005).

[24] D. A. Ostlie and B. W. Carroll, *An Introduction to Modern Stellar Astrophysics* (Pearson Education, Inc., 2013).

[25] D. Kazanas and D. N. Schramm, Nature **262**, 671 (1976).

[26] C. L. Fryer, Classical and Quantum Gravity **20**, S73 (2003).

[27] E. Chaisson and S. McMillan, *Astronomy: A Beginners Guide to the Universe* (Pearson Education, Inc., 2013).

[28] S. Chandrasekhar, apj **74**, 81 (1931).

[29] P. A. STURROCK, Astrophys.J. **164**, 529 (1971).

[30] P. R. Chaudhuri, J.Phys. **A4**, 508 (1971).

[31] V. L. Ginzburg and V. V. Zheleznyakov, Ann.Rev.Astron.Astrophys. **13**, 511 (1975).

[32] Y. Ochelkov and V. Usov, Astrophysics and Space Science **69**, 439 (1980).

[33] I. Contopoulos, (2006).

[34] I. Contopoulos and A. Spitkovsky, Class.Quant.Grav. **643**, 1139 (2006).

[35] O. Hamil, arXiv.org (2015), 1503.09110.

[36] W. C. G. Ho, arXiv.org (2011), 1102.4870v1.

[37] A. R. Choudhuri and S. Konar, arXiv.org (2003), astro-ph/0311031v1.

[38] I. H. Stairs, Living Rev Relativity (2003).

[39] G. B. Hobbs, R. T. Edwards, and R. N. Manchester, Mon. Not. R. Astron. Soc. **369**, 655 (2006), arXiv:astro-ph/0603381.

[40] R. van Haasteren *et al.*, Mon. Not. R. Astron. Soc. **414**, 3117 (2011), 1103.0576.

[41] P. B. Demorest *et al.*, Astrophys. J. **762**, 94 (2013).

[42] R. N. Manchester *et al.*, Publications of the Astronomical Society of Australia **30**, 17 (2013), 1210.6130.

[43] M. V. Sazhin, Soviet Astron. **22**, 36 (1978).

[44] S. Detweiler, Astrophys. J. **234**, 1100 (1979).

[45] M. Anholm, S. Ballmer, J. D. E. Creighton, L. R. Price, and X. Siemens, Phys. Rev. D **79**, 084030 (2009), 0809.0701.

[46] B. Allen and J. D. Romano, Phys. Rev. D **59**, 102001 (1999).

[47] A. H. Jaffe and D. C. Backer, Astrophys. J. **583**, 616 (2003), arXiv:astro-ph/0210148.

[48] A. Sesana, A. Vecchio, and C. N. Colacino, Mon. Not. R. Astron. Soc. **390**, 192 (2008), 0804.4476.

[49] A. Sesana, A. Vecchio, and M. Volonteri, Mon. Not. R. Astron. Soc. **394**, 2255 (2009), 0809.3412.

[50] A. N. Lommen and D. C. Backer, Astrophys. J. **562**, 297 (2001), arXiv:astro-ph/0107470.

[51] J. S. B. Wyithe and A. Loeb, Astrophys. J. **590**, 691 (2003), arXiv:astro-ph/0211556.

[52] M. Volonteri, F. Haardt, and P. Madau, Astrophys. J. **582**, 559 (2003), arXiv:astro-ph/0207276.

[53] M. Enoki, K. T. Inoue, M. Nagashima, and N. Sugiyama, Astrophys. J. **615**, 19 (2004), arXiv:astro-ph/0404389.

[54] A. Sesana, arXiv:1211.5375 [astro-ph.CO] (2012).

[55] S. T. McWilliams, J. P. Ostriker, and F. Pretorius, arXiv:1211.4590 [astro-ph.CO] (2012).

[56] A. Sesana and A. Vecchio, Phys. Rev. D **81**, 104008 (2010), 1003.0677.

[57] C. Roedig and A. Sesana, arXiv:1111.3742 [astro-ph.CO] (2011).

[58] V. Ravi *et al.*, arXiv: 1210.3854 [astro-ph.CO] (2012).

[59] C. M. F. Mingarelli, K. Grover, T. Sidery, R. J. E. Smith, and A. Vecchio, Phys. Rev. Lett. **109**, 081104 (2012), 1207.5645.

[60] R. van Haasteren and Y. Levin, Mon. Not. R. Astron. Soc. **401**, 2372 (2010), 0909.0954.

[61] J. Cordes and F. Jenet, Astrophys. J. **752**, 54 (2012).

[62] S. Olmez, V. Mandic, and X. Siemens, Phys. Rev. D **81**, 104028 (2010), 1004.0890.

[63] L. Sousa and P. P. Avelino, Phys. Rev. D **88**, 023516 (2013), 1304.2445.

[64] K. Miyamoto and K. Nakayama, Journal of Cosmology and Astroparticle Physics **7**, 12 (2013), 1212.6687.

[65] S. Kuroyanagi, K. Miyamoto, T. Sekiguchi, K. Takahashi, and J. Silk, Phys. Rev. D **87**, 023522 (2013), 1210.2829.

[66] C. Caprini, R. Durrer, and X. Siemens, Phys. Rev. D **82**, 063511 (2010), 1007.1218.

[67] A. A. Starobinsky, JETP Lett. **30**, 682 (1979).

[68] W. Zhao, Y. Zhang, X.-P. You, and Z.-H. Zhu, Phys. Rev. D **87**, 124012 (2013), 1303.6718.

[69] C. M. F. Mingarelli, T. Sidery, I. Mandel, and A. Vecchio, Phys. Rev. D **88**, 062005 (2013), 1306.5394.

[70] S. R. Taylor and J. R. Gair, Phys. Rev. D **88**, 084001 (2013), 1306.5395.

[71] R. W. Hellings and G. S. Downs, Astrophys. J. Lett. **265**, L39 (1983).

[72] D. R. Stinebring, M. F. Ryba, J. H. Taylor, and R. W. Romani, Phys. Rev. Lett. **65**, 285 (1990).

[73] A. N. Lommen, New Limits on Gravitational Radiation using Pulsars, in *Neutron Stars, Pulsars, and Supernova Remnants*, edited by W. Becker, H. Lesch, & J. Trümper, pp. 114–+, 2002, arXiv:astro-ph/0208572.

[74] F. A. Jenet, G. B. Hobbs, K. J. Lee, and R. N. Manchester, Astrophys. J. Lett. **625**, L123 (2005), arXiv:astro-ph/0504458.

[75] F. A. Jenet *et al.*, Astrophys. J. **653**, 1571 (2006), arXiv:astro-ph/0609013.

[76] R. van Haasteren, Y. Levin, P. McDonald, and T. Lu, Mon. Not. R. Astron. Soc. **395**, 1005 (2009), 0809.0791.

[77] R. van Haasteren, Y. Levin, P. McDonald, and T. Lu, Mon. Not. R. Astron. Soc. **395**, 1005 (2009), 0809.0791.

[78] D. R. B. Yardley *et al.*, Mon. Not. R. Astron. Soc. **414**, 1777 (2011), 1102.2230.

[79] J. M. Cordes and R. M. Shannon, Astrophys. J. **750**, 89 (2012), 1106.4047.

[80] L. Boyle and U.-L. Pen, Phys. Rev. D **86**, 124028 (2012), 1010.4337.

[81] R. van Haasteren and Y. Levin, Mon. Not. R. Astron. Soc. **428**, 1147 (2013), 1202.5932.

[82] J. Ellis, X. Siemens, and R. van Haasteren, arXiv:1302.1903 [astro-ph.IM] (2013).

[83] J. Gair, J. D. Romano, S. Taylor, and C. M. F. Mingarelli, Phys. Rev. D **90**, 082001 (2014).

[84] C. M. F. Mingarelli and T. Sidery, Phys. Rev. D **90**, 062011 (2014).

[85] J. A. Ellis, X. Siemens, and J. D. E. Creighton, Astrophys. J. **756**, 175 (2012), 1204.4218.

[86] A. Petiteau, S. Babak, A. Sesana, and M. de Araújo, Phys. Rev. D **87**, 064036 (2013).

[87] J. A. Ellis, F. A. Jenet, and M. A. McLaughlin, Astrophys. J. **753**, 96 (2012), 1202.0808.

[88] M.-L. Tong *et al.*, Chinese Physics Letters **30**, 100402 (2013), 1306.6719.

[89] S. Taylor, J. Ellis, and J. Gair, arXiv:1406.5224 [gr-qc] (2014).

[90] Z. Arzoumanian *et al.*, arXiv:1404.1267 (2014).

[91] Y. Wang, S. D. Mohanty, and F. A. Jenet, arXiv:1406.5496 (2014).

[92] S. Yi *et al.*, arXiv:1409.2296 [astro-ph.SR] (2014).

[93] X.-J. Zhu *et al.*, Mon. Not. R. Astron. Soc. **444**, 3709 (2014), 1408.5129.

[94] L. S. Finn and A. N. Lommen, Astrophys. J. **718**, 1400 (2010), 1004.3499.

[95] J. Cordes and F. Jenet, Astrophys.J. **752**, 54 (2012).

[96] M. Pitkin, Mon. Not. R. Astron. Soc. **425**, 2688 (2012), 1201.3573.

[97] D. R. Madison, J. M. Cordes, and S. Chatterjee, Astrophys. J. **788**, 141 (2014), 1404.5682.

[98] X. Deng, Phys. Rev. D **90**, 024020 (2014), 1404.0663.

[99] C. Cutler, S. Burke-Spolaor, M. Vallisneri, J. Lazio, and W. Majid, Phys. Rev. D **89**, 042003 (2014), 1309.2581.

[100] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical recipes in C (2nd ed.): the art of scientific computing* (Cambridge University Press, New York, NY, USA, 1992).

[101] P. B. Demorest, *Measuring the gravitational wave background using precision pulsar timing*, PhD thesis, University of California, Berkeley, 2007.

[102] W. Coles, G. Hobbs, D. J. Champion, R. N. Manchester, and J. P. W. Verbiest, Mon. Not. R. Astron. Soc. **418**, 561 (2011), 1107.5366.

[103] J. A. Ellis *et al.*, in preparation (2014).

[104] X. Siemens, J. Ellis, F. Jenet, and J. D. Romano, Classical and Quantum Gravity **30**, 224015 (2013).

[105] F. W. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark, *NIST Handbook of Mathematical Functions*, 1st ed. (Cambridge University Press, New York, NY, USA, 2010).

[106] L. Lentati, P. Alexander, M. P. Hobson, S. Taylor, and S. T. Balan, arXiv:1210.3578 [astro-ph.IM] (2012).

[107] S. R. Taylor, J. R. Gair, and L. Lentati, arXiv:1210.3489 [astro-ph.IM] (2012).

[108] S. Nojiri and S. D. Odintsov, arXiv.org (2006), hep-th/0601213v5.

[109] S. Nojiri and S. D. Odintsov, Class.Quant.Grav. **505**, 59 (2011).

[110] F. S. N. Lobo, arXiv.org (2008), 0807.1640v1.

[111] M. E. S. Alves, O. D. Miranda, and J. C. N. de Araujo, arXiv.org (2009), 0908.0861.

[112] S. Capozziello and M. Francaviglia, arXiv.org (2007), 0706.1146.

[113] A. De Felice and S. Tsujikawa, arXiv.org (2010), 1002.4928.

[114] C. M. Will, *Theory and Experiment in Gravitational Physics* (Cambridge University Press, New York, NY, USA, 1985).

[115] M. Brunetti, E. Coccia, V. Fafone, and F. Fucito, Class.Quant.Grav. **D59**, 044027 (1999).

[116] T. Clifton, P. G. Ferreira, A. Padilla, and C. Skordis, arXiv.org (2011), 1106.2476.

[117] E. Sagi, Class.Quant.Grav. **D81**, 064031 (2010).

[118] T. Clifton, M. Bañados, and C. Skordis, arXiv.org (2010), 1006.5619.

[119] C. Skordis, arXiv.org (2009), 0903.3602.

[120] M. Milgrom, arXiv.org (2009), 0912.0790.

[121] A. Nishizawa, A. Taruya, K. Hayama, S. Kawamura, and M.-a. Sakagami, Physical Review D **79** (2009).

[122] K. J. Lee, F. A. Jenet, and R. H. Price, Astrophys. J. **685**, 1304 (2008).

[123] M. E. d. S. Alves and M. Tinto, arXiv.org (2011), 1102.4824.

[124] D. M. Eardley, D. L. Lee, A. P. Lightman, R. V. Wagoner, and C. M. Will, Phys. Rev. Lett. **30**, 884 (1973).

[125] M. Tinto and M. E. da Silva Alves, Phys.Rev. **D82**, 122003 (2010).

[126] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, Astron.J. **129**, 1993 (2005). 149

[127] J. M. Cordes and T. J. W. Lazio, arXiv.org (2002), astro-ph/0207156.

[128] J. R. Smith and F. t. L. S. Collaboration, arXiv.org (2009), 0902.0381.

[129] S. A. Hughes, Ann. Phys. **303**, 142 (2002).

[130] W. G. Anderson, P. R. Brady, J. D. E. Creighton, and E. E. Flanagan, arXiv.org (2000), gr-qc/0008066v1.

[131] P. Brady, D. Brown, K. Cannon, and S. Ray-Majumder, Preprint , 1 (2007).

[132] K. Cannon, lalapps_power [software], 2006.

[133] C. Pankow, gstlal_excesspower [software], 2011.

[134] J. McIver, F. t. L. S. Collaboration, and f. t. V. Collaboration, In Progress (2015).

[135] V. Predoi, the LIGO Scientific Collaboration, V. Collaboration, K. Hurley, and the IPN, Journal of Physics: Conference Series **363**, 012034 (2012).

[136] L. S. Collaboration, t. V. Collaboration, J. Abadie, and Others, arXiv.org (2012), 1202.2788v3.

[137] J. McIver, F. t. L. S. Collaboration, and F. t. V. Collaboration, arXiv.org (2012), 1204.2497v1.

[138] W. A. Majid *et al.*, Developing Precision Pulsar Timing Capability for the DSN, in *American Astronomical Society Meeting Abstracts #223*, , American Astronomical Society Meeting Abstracts Vol. 223, p. 153.16, 2014.

[139] F. Camilo and F. A. Rasio, arXiv.org (2005), astro-ph/0501226v1.

[140] S. M. Ransom *et al.*, arXiv.org (2005), astro-ph/0501230.

Sydney J. Chamberlin

EDUCATION

8/2009–present **Ph.D Student in Physics**
Center for Gravitation, Cosmology and Astrophysics, UWM, Milwaukee WI

8/2004–6/2009 **B.S. in Physics with Honors**
Utah State University, Logan, UT

1/2007–6/2009 **B.S. in Mathematics**
Utah State University, Logan, UT

RESEARCH EXPERIENCE

06/2010–present **Graduate Research Assistant**
Center for Gravitation, Cosmology and Astrophysics, UWM, Milwaukee WI
Advisor: Dr. Jolien D. E. Creighton, Dr. Xavier Siemens

1/2008–5/2009 **Undergaduate Research Assistant**
Dept. of Physics, Utah State University, Logan, UT
Advisor: Dr. Charles Torre

6/2007–8/2007 **Summer Intern**
Lunar and Planetary Institute, NASA Johnson Space Center, Houston, TX
Advisor: Dr. Roy Christoffersen

6/2007–8/2007 **Summer Exchange Student**
Institute for Diagnostic and Interventional Radiology, Fredrich Schiller University, Jen
Advisor: Dr. Jan Sedlacik

RESEARCH INTERESTS

Gravitational wave data analysis, gravitational wave searches with pulsar timing arrays, gravitational wave searches with the Laser Interferometer Gravitational-wave Observatory (LIGO), tests of General Relativity, metric theories of gravity, compact objects, relativistic astrophysics, physics education

TEACHING EXPERIENCE

1/2015–5/2015 Physics 210 (calculus-based introductory physics II) Discussion TA
Dept. of Physics, University of Wisconsin-Milwaukee

8/2014–12/2014 Physics 610 (graduate-level course) Instructor
Dept. of Physics, University of Wisconsin-Milwaukee

1/2014–5/2014 Physics 215 (calculus-based introductory physics II lab) Lab TA
Dept. of Physics, University of Wisconsin-Milwaukee

1/2013–5/2013 Physics 215 (calculus-based introductory physics II lab) Lab TA
Dept. of Physics, University of Wisconsin-Milwaukee

1/2010–5/2010 Physics 108 (introductory physics lab for non-science majors) Lab TA
Dept. of Physics, University of Wisconsin-Milwaukee

8/2009–12/2009 Physics 108 (introductory physics lab for non-science majors) Lab TA
Dept. of Physics, University of Wisconsin-Milwaukee

1/2009–5/2009 Physics 2215 (calculus based introductory physics I lab) Lab TA
Dept. of Physics, Utah State University

AWARDS & HONORS

2010-2015	Wisconsin Space Grant Consortium Graduate Fellowship
2013	James B. Hartle Award, 10th Amaldi Conference
2011-2012	Distinguished Graduate Fellowship, University of Wisconsin-Milwaukee
2009	Utah State University College of Science *Undergraduate Researcher of the Year*
2008	Rhodes Scholarship Interview Finalist
2008	Barry M. Goldwater Scholarship Honorable Mention
2008-2009	Willard L. Eccles Undergraduate Research Fellowship, Utah State University
2007	Carl A. Spaatz Award, Civil Air Patrol, United States Air Force Auxiliary
2004–2008	Utah State University Presidential Scholarship

OUTREACH AND VOLUNTEER ACTIVITIES

CoffeeShop Astrophysics: Designed, obtained funding for, and implemented a Milwaukee public lecture series, *CoffeeShop Astrophysics*, aimed at bringing topics from astronomy and physics to the public. Developed partnership with a local coffee shop, *Anodyne Coffee Roasting Co.*, where lectures are held. Coordinate space and volunteers, manage grant funding, prepare lectures and visual aids, and present material to the public (ongoing, 2014-2015).

Hawley Environmental School: Worked with two fifth grade classes ($\sim$ 30 students each) to teach children about the the life cycle of stars and other astrophysics topics at the K-5 Hawley Environmental School. Students participated in hands-on "Make your own play-dough pulsar" activity and learned about objects currently studied in astronomy. (April 2013, April 2014)

University of Wisconsin-Waukesha STEM Day: Presented "An Astro Safari", a guide to the structure and contents of our Universe, to 6-8th grade students at the University of Wisconsin-Waukesha *STEM Day* along with Sarah Caudill (UWM) and Megan DeCesar (UWM). (March 2014)

Nicolet High School Science Fair: mServed as judge of physics projects at the Nicolet High School Science Fair. (April 2014)

Milwaukee Public School District Science Fair: Judged student science fair projects in the Milwaukee Public School District Science Fair. (May 2013)

Badger State Science Fair: Served as a *Scientific Process Judge* in the Badger State Science Fair, which selects students to take part in the Intel International Science and Engineering Fair. (March 2013)

COMPASS Scholarship Program, UWM: Served as graduate mentor for undergraduate science scholars: discussed career options and graduate school applications with students (Fall 2012, Spring 2013).

UWM Astronomy Club: Attended UWM Astronomy Club meetings and field trips to local observatories and planetariums. (Fall 2011 – Fall 2012)

Civil Air Patrol: Twelve years of experience with extensive leadership training, including attendance at the week-long Cadet Officer School (Maxwell Air Force Base, 2003) and leadership of a local unit of Civil Air Patrol youth for a year (2002-2003). Additionally, organized and orchestrated a week-long summer camp for Utah Civil Air Patrol youth, and spent significant time as an aerospace instructor, teaching topics on the basic physics of flight and model rocketry. (August 1998 – July 2010)

ORGANIZATIONS

2014–present	Member of IRG (International Society for General Relativity and Gravitation)
2012–present	Member of the American Physical Society
2011–present	Member of AAS (American Astronomical Society)
2010–present	Member of IPTA (International Pulsar Timing Array)
2010–present	Member of NANOGrav (North American Nanohertz Observatory for Gravitational Waves)
2010–present	Member of LIGO (Laser Interferometer Gravitational-wave Observatory)
2007–present	Member of Phi Sigma Sigma National Honor Society
2007–present	Member of the Spaatz Association

SELECTED CONFERENCES AND SEMINARS

"Gravity and the unseen sky, *or*, how to probe General Relativity with dead stars." Jet Propulsion Laboratory, Pasadena, California, USA. 18 February 2015. (invited talk)

"Gravity and the unseen sky, *or*, how to probe General Relativity with dead stars." California Institute of Technology, Pasadena, California, USA. 17 February 2015. (invited talk)

"Gravity and the unseen sky, *or*, how to listen to the universe with dead stars." Joint Institute for Nuclear Astrophysics – Center for the Evolution of the Elements, Michigan State University, East Lansing, MI, USA. 21 January 2015. (invited talk)

"Gravity and the unseen sky, *or*, how to listen to the universe with black holes." Utah State University, Logan, Utah, USA. 19 November 2013. (invited talk)

"Listening to the universe with pulsars: gravitational wave data analysis with pulsar timing arrays." University of Wisconsin-Milwaukee Astronomy Club, Milwaukee, Wisconsin, USA. 14 October 2013. (invited talk)

"Searches for a stochastic gravitational wave background with pulsar timing arrays." Albert Einstein Institute, Max Planck Institute, Golm-Potsdam, Germany. 3 October 2013. (invited talk)

"Studying the universe with gravity." University of Wisconsin-La Crosse, La Crosse, Wisconsin, USA. 5 December 2012. (invited talk)

"Probing alternative theories of gravity with gravitational wave detection from pulsar timing arrays." 2012 International Pulsar Timing Array Student Workshop, University of Sydney, Sydney, Australia. 20 June 2012. (invited talk)

"Developments in ExcessPower: Applications of the ExcessPower Pipeline to Burst Searches and Detector Characterization." LIGO/Virgo Scientific Collaboration Meeting, Stanford University, Stanford, USA. 26 August 2014. *(poster presentation)*

"The Quest to Characterize Non-Gaussian Noise Components in Interferometric Gravitational Wave Data." 24th Wisconsin Space Conference, BioPharmaceutical Technology Center, Fitchburg, Wisconsin, USA. 15 August 2014. (contributed talk)

"Applications of the ExcessPower Data Analysis Pipeline to Gravitational Wave Detection Efforts." 23rd Midwest Relativity Meeting, Milwaukee, Wisconsin, USA. 26 October 2013. (contributed talk)

"When Galaxies Collide: The Search for Low-frequency Gravitational Wave Backgrounds in the Universe." 23rd Wisconsin Space Conference, Milwaukee, Wisconsin, USA. 14 August 2013. (contributed talk)

"Searches for a Stochastic Gravitational Wave Background with Pulsar Timing Arrays: A Data Analysis Pipeline." 20th International Conference on General Relativity and Gravitation/10th Amaldi Conference on Gravitational Waves, University of Warsaw, Poland. 10 July 2013. (contributed talk) ***Won James B. Hartle Award for talk***

"Searches for a Stochastic Gravitational Wave Background with Pulsar Timing Arrays: A Data Analysis Pipeline." 221st American Astronomical Society Meeting, Long Beach, California, USA. 7 January 2013. (contributed talk)

"Operating the Stochastic Gravitational Wave Background Pipeline." NANOGrav Collaboration Fall Meeting, Oberlin College, Oberlin, Ohio. 24 October 2012. (contributed talk)

"Searches for a Stochastic Gravitational Wave Background with Pulsar Timing Arrays: A Data Analysis Pipeline." 22nd Midwest Relativity Meeting, University of Chicago, Chicago, Illinois, USA. 28 August 2012. (contributed talk)

"Testing General Relativity with Pulsar Timing Arrays." 22nd Wisconsin Space Conference, University of Wisconsin-Whitewater, Whitewater, Wisconsin, USA. 16 August 2012. (contributed talk)

"Probing Alternative Theories of Gravity with Gravitational Wave Detection from Pulsar Timing Arrays." Phenomenology 2012 Symposium, University of Pittsburgh, Pittsburgh, Pennsylvania, USA. 7 May 2012. (contributed talk)

PUBLICATIONS

Z. Arzoumanian, A. Brazier, S. Burke-Spolaor, **S. J. Chamberlin**, et al, "NANOGrav Constraints on Gravitational Wave Bursts with Memory", arXiv:1501.05343 [astro-ph.GA] (2015)

S. J. Chamberlin , J. D. E. Creighton, P. B. Demorest, J. Ellis, L. Price, J. D. Romano, X. Siemens, "Time-domain Implementation of the Optimal Cross-correlation Statistic for Stochastic Gravitational Wave Background Searches in Pulsar Timing Data", Physical Review D **91**, 044048 (2015).

Z. Arzoumanian, A. Brazier, S. Burke-Spolaor, **S. J. Chamberlin**, et al, "NANOGrav Limits on Gravitational Waves from Supermassive Black Hole Binaries in Cicular Orbits", The Astrophysical Journal **794**, 141 (2014).

J. A. Ellis, "A Bayesian analysis pipeline for continuous GW sources in the PTA band," *Classical and Quantum Gravity* **30** (2013) no. 22, 224004.

J. Ellis, X. Siemens, **S. Chamberlin**, "Results of the First IPTA Closed Mock Data Challenge", arXiv:1210.5274 [astro-ph.IM] (2012).

P. B. Demorest, R. D. Ferdman, M. E. Gonzalez, D. Nice, S. Ransom, I. H. Stairs, Z. Arzoumanian, A. Brazier, S. Burke-Spolaor, **S. J. Chamberlin**, et al., "Limits on the Stochastic Gravitational Wave Background from the North American Nanohertz Observatory for Gravitational Waves", The Astrophysical Journal **762**, 94 (2013).

S. J. Chamberlin and X. Siemens, "Stochastic Backgrounds in Alternative Theories of Gravity: Overlap Reduction Functions for Pulsar Timing Arrays", Physical Review D **85**, 082001 (2012).

J. Aasi et al. (the LIGO Scientific Collaboration), "Searches for continuous gravitational waves from nine young supernova remnants", arXiv:1412.5942 [astro-ph.HE] (2014).

J. Aasi et al. (the LIGO Scientific Collaboration, the Virgo Collaboration), "A directed search for gravitational waves from Scorpius X-1 with initial LIGO", arXiv:1412.0605 [gr-qc] (2014).

J. Aasi et al. (the LIGO Scientific Collaboration, the Virgo Collaboration), "Narrow-band search of continuous gravitational-wave signals from Crab and Vela pulsars in Virgo VSR4 data", arXiv:1410.8310 [astro-ph.IM] (2014).

M. G. Aartsen et al. (the IceCube Collaboration, the LIGO Scientific Collaboration, the Virgo Collaboration), "Multimessenger Search for Sources of Gravitational Waves and High-Energy Neutrinos: Results for Initial LIGO-Virgo and IceCube", arXiv:1407.1042 [astro-ph.HE] (2014).

J. Aasi et al. (the LIGO Scientific Collaboration, the Virgo Collaboration), "Narrow-band search of continuous gravitational-wave signals from Crab and Vela pulsars in Virgo VSR4 data", arXiv:1410.8310 [astro-ph.IM] (2014).

J. Aasi et al. (the LIGO Scientific Collaboration, the Virgo Collaboration), "First all-sky search for continuous gravitational waves from unknown sources in binary systems", Physical Review D **90**, 062010 (2014).

J. Aasi et al. (the LIGO Scientific Collaboration, the Virgo Collaboration), "Improved upper limits on the stochastic gravitational-wave background from 2009-2010 LIGO and Virgo data", arXiv:1406.4556 [gr-qc] (2014).

J. Aasi et al. (the LIGO Scientific Collaboration, the Virgo Collaboration), "Methods and results of a search for gravitational waves associated with gamma-ray bursts using the GEO600, LIGO, and Virgo detectors", arXiv:1405.1053 [astro-ph.HE] (2014).

J. Aasi et al. (the LIGO Scientific Collaboration, the Virgo Collaboration), "Search for gravitational radiation from intermediate mass black hole binaries in data from the second LIGO-Virgo joint science run ", Physical Review D **89**, 122003 (2014).

J. Aasi et al. (the LIGO Scientific Collaboration, the Virgo Collaboration), "Search for gravitational waves associated with gamma-ray bursts detected by the InterPlanetary Network", Physical Review Letters **113**, 011102 (2014).

J. Aasi et al. (the LIGO Scientific Collaboration), "Implementation of an F-statistic all-sky search for continuous gravitational waves in Virgo VSR1 data", submitted to Classical and Quantum Gravity (2014).

J. Aasi et al. (the LIGO Scientific Collaboration, the Virgo Collaboration), "Search for gravitational wave ringdowns from perturbed intermediate mass black holes in LIGO-Virgo data from 2005-2010", Physical Review D **89**, 102006 (2014).

J. Aasi et al. (the LIGO Scientific Collaboration, the Virgo Collaboration), "The NINJA-2 project: Detecting and characterizing gravitational waveforms modelled using numerical binary black hole simulations", Classical and Quantum Gravity **31**, 115004 (2014).

S. J. Chamberlin, "Algebraic Computing Tools in General Relativity: Energy-Momentum Tensors and Exact Solutions to the Einstein Field Equations," Undergraduate Thesis Published in Undergraduate Honors Theses, Utah State University Digital Commons: pp. 1-93, (2009).

R. Christoffersen and **S. Chamberlin**, "Space Plasma Processing of Lunar Dust: Modeling of Radiation-damaged Rim Widths on Lunar Soil Minerals", LPI Contributions **1415**, 2092 (2008).

S. Chamberlin, R. Christoffersen, and L. P. Keller, "Space Plasma Ion Processing of the Lunar Soil: Modeling of Radiation-damaged Rim Widths on Lunar Grains", in Lunar and Planetary Institute Science Conference Abstracts, vol. 39 of Lunar and Planetary Institute Science Conference Abstracts, p. 2302 (2008).