May 2019

Measuring Deterministic and Stochastic Gravitational Waves with Pulsar Timing Array Experiments

Kristina Islo
University of Wisconsin-Milwaukee

Recommended Citation

Islo, Kristina, "Measuring Deterministic and Stochastic Gravitational Waves with Pulsar Timing Array Experiments" (2019). *Theses and Dissertations*. 2078.
https://dc.uwm.edu/etd/2078

MEASURING DETERMINISTIC AND STOCHASTIC GRAVITATIONAL WAVES WITH PULSAR TIMING ARRAY EXPERIMENTS

by

Kristina Islo

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 2019

ABSTRACT

MEASURING DETERMINISTIC AND STOCHASTIC GRAVITATIONAL WAVES WITH PULSAR TIMING ARRAY EXPERIMENTS

by

Kristina Islo

The University of Wisconsin–Milwaukee, 2019
Under the Supervision of Professor Xavier Siemens

Pulsar timing arrays (PTAs) are uniquely poised to detect the nanohertz-frequency gravitational waves from supermassive black hole binaries (SMBHBs) formed during galaxy merger. Efforts are underway to observe three species of gravitational signal from these systems: the stochastic ensemble, individual, adiabatic binary inspirals, and bursts with memory. This dissertation discusses all three.

A typical Bayesian search for evidence of a stochastic gravitational wave background from the superposition of many unresolvable SMBHB inspirals requires weeks to months to deliver results. This is due in part to the inclusion of inter-pulsar spatial and temporal correlations induced in PTA data by such a signal. By integrating a simplified Bayesian search into an existing frequentist statistic, we are able to create a robust background amplitude estimator that requires minimal CPU time and does not compromise the key information gleaned from a full Bayesian analysis.

As PTA sensitivity increases, individual binary inspirals will rise above the stochastic background, promising information about local SMBHBs. PTAs, like the North American Nanohertz Observatory for Gravitational Waves, regularly conduct searches for these single-frequency sources, with the latest results coming from the 9- and 11-year data sets. Although detection is still in the future, results are already informing binary candidate properties.

Finally, an SMBHB coalescence is theoretically accessible to PTAs through bursts with memory, a purely General Relativistic phenomenon which imparts a permanent spacetime deformation and affects the coalescence signal amplitude at leading quadrupole order. Simulations parameterized by astrophysical observables from galaxy mergers out to z=3 predict the rates and signal-to-noise ratios for bursts occurring in the PTA-band. Extending the synthesized population to include less massive SMBHBs shows space-based interferometers may also observe this atypical signature.

for

Spencer and Ann

TABLE OF CONTENTS

LIST OF FIGURES

Acronyms

CW Continuous Wave

GBT Green Bank Telescope

GSMF Galaxy Stellar Mass Function

GW Gravitational Wave

GWB Gravitational Wave Background

HD Hellings and Downs

ISCO Innermost Stable Circular Orbit

LIGO Laser Interferometer Gravitational-Wave Observatory

LISA Laser Interferometer Space Antenna

MCMC Markov Chain Monte Carlo

MSP Millisecond Pulsar

NANOGrav North American Observatory for Nanohertz Gravitational Waves

ORF Overlap-reduction function

PTA Pulsar Timing Array

PSD Power Spectral Density

SMBHB Supermassive Black Hole Binary

SNR Signal-to-Noise Ratio

SSB Solar System Barycenter

SSE Solar System Ephemeris

TOA Time of Arrival

ACKNOWLEDGMENTS

My ability to engage in the research culminating into this document was fostered by numerous individuals, more than I can iterate here.

This dissertation is dedicated primarily to my family, for wholly supporting my curiosity and ambition. My brother for complying with my wishes to be called “Dr. Tina” well in advance – and regardless of – the completion of my degree.

To Muon and Kisa: thank you for nuzzling my laptop out of the way and curling up in my lap when you’ve understood me to be burnt out, for enthusiastically greeting me at the side door (you know, the one right by your food bowls) when I arrive home after an especially stressful day, and for respecting that I, personally, do not ever appreciate being bitten by tiny, needle teeth.

To Robert Wood for his outstanding teaching, rigorous assignments, and wit. I’ll remember you fondly as the man who fostered, from first principles, a literature major without a lick of background to a full-fledged physicist. Your instruction and uncompromising expectations quickly honed my scientific understanding to a caliber that rivaled my seasoned peers, and it couldn’t have come from a kinder, more invested individual. I’d be happy to know if you understand even a fraction of how much you have meant to me as a scientist and human being [1].

To Xavier Siemens and Luis Anchordoqui for their supervision and support.

To Sarah Vigeland for being an incredible friend and mentor. Thank you for believing in my abilities when I did not.

All the incredible scientists whom I befriended during my research. This includes, but is not limited to Megan Decesar, Laleh Shedanigan, and Danielle Berg. I’m a better woman and scientist for your influence.

To Tonia Klein for being one of my very best friends, and to whom I never tire of speaking. I am sure our office neighbors will miss the loud cackling.

To Glynn Lane, whose support buoyed me through difficult periods, and genuinely interesting conversations about programming provided distraction.

---

[1] generally overlapping categories

To Heidi Matera and Gracie Liebenstein for being above-and-beyond supportive department sisters. Thank you for grounding me and for showing me care when everything was overwhelming.

The many (then) NANOGrav postdocs who welcomed and mentored me: Sarah Burke-Spolaor, Joseph Simon, and Stephen Taylor.

Thank you to the individuals who made the department a safer and more welcoming space. David Kaplan for sharing adorable and inspiring cat videos, Joe Swiggum for his teatime conversation, and Michael Manske for unparalleled sensitivity to my struggles, laughing at my jokes, and reminding me to change the oil in my car.

Finally, to all the women whose curiosity has led them to physics; who grew up not obsessed with Star Trek; who are told they aren't legitimate scientists because so; who are discouraged from pursuing physics; who are accused of cheating by bewildered, sexist peers; who struggle alone in their otherness among self-congratulating cliques; who have no choice but to develop an acute social understanding in order to navigate academic culture intact; who work over time just to feel adequate; whose brilliance is often underestimated, buried or stolen; and who survive and persist in the face of violations by colleagues – those resilient individuals who all too often question their worth and right to exist in intellectual spaces: You are amazing.

"*Id like to ask the men here to consider idly, in some spare moment, whether by any chance they've been building any walls to keep the women out, or to keep them in their place, and what they may have lost by doing so.*"

— Ursula K Le Guin

PREFACE

The work presented in this dissertation stems from my participation in the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) collaboration.

Chapter 1

Introduction

1.1 Gravitational Wave Spectrum

Gravitational radiation offers a unique opportunity to observe “dark” phenomena in the Universe. Objects invisible to telescopes are illuminated through their propagated influence on spacetime curvature. In a manner independent of electromagnetic observations, gravity is able to communicate information in the form of spacetime perturbations traveling at the speed of light. These gravitational waves (GWs), like light, contain the signatures of various physical processes. With the addition of General Relativistic models and complementary electromagnetic observation, GWs can be used to further our understanding of astrophysics and cosmology.

Currently, detected GWs only inform the physics governing the inspiral and coalescence of some compact binary systems including stellar mass black holes and neutron stars. This is only a small slice of the types of potentially detectable binaries, however. An entire spectrum of GW frequencies exists – spanning an array of interesting sources (see Fig. 1).

Binaries with orbital periods longer than milliseconds emit radiation beyond the sensitivity of ground-based GW observatories, like the Laser Interferometer Gravitational-Wave Observatory (LIGO) and Virgo (Martynov et al. 2016). The development of GW detectors optimized to operate in lower-frequency bandwidths is necessary if we are to understand a broad range of binary configurations.

Seismic noise sets the minimum frequency detectable by terrestrial GW observatories – so, lower-band instruments must operate in space. The proposed Laser Interferometer Space Antenna (LISA) will cover frequencies within $10^{-4} - 1\,\mathrm{Hz}$, encapsulating signals emitted by extreme mass ratio compact binary inspirals and black hole binaries up with masses up to $10^8\,\mathrm{M}_\odot$ (Amaro-Seoane et al. 2017). The confusion background of signals from binaries in the Milky Way will cap the lower-frequency end of LISA's sensitivity.

Going beyond the solar system, Galactic-scale GW detectors will have access to the dynamics of supermassive black hole binaries (SMBHBs). Classic Michelson interferometers rely on measuring the passage of a GW by witnessing an increased or decreased displacement between two (effectively) freely-falling test masses, be they suspended mirrors or propulsion-stabilized spacecraft. For SMBHBs, this experimental method requires test-mass separations of thousands of parsecs. Thus, Galactic-scale detectors use existing astrophysical objects to probe these lowest GW frequencies.

1.2 Pulsar Timing Arrays

Pulsar timing arrays (PTAs) are GW observatories sensitive to frequencies in the $10^{-9} - 10^{-7}$ Hz range (Detweiler 1979; Sazhin 1978). PTAs consist of a network of galactic Millisecond Pulsars (MSPs): a class of rapidly rotating neutron stars with strongly beamed radio emission along its magnetic poles. If the spin and magnetic axes are misaligned, and some of the radio emission is projected onto the line-of-sight to Earth, MSPs appear as ultra-stable pulsing beacons, rivaling the precision of atomic clocks (Hartnett & Luiten 2011). MSPs exhibit stability over $\sim$ decades (see Cordes 2013) – a necessary condition when optimizing for long-wavelength GW detection.

By monitoring MSPs with radio telescopes, we can probe not only the physics which governs bodies of supranuclear density (e.g. Cromartie et al. 2019), but also the spacetime at and between the Earth and the pulsar. Importantly, a GW would alter the null geodesic along which these pulses travel, affecting the observed radio signal. PTAs aim to accurately measure the radio pulse times-of-arrival (TOAs). We develop a theoretical timing model which predicts pulse arrival and compare that against the measured TOAs. The difference between the two is the *timing residual.* A benefit of using MSPs is the

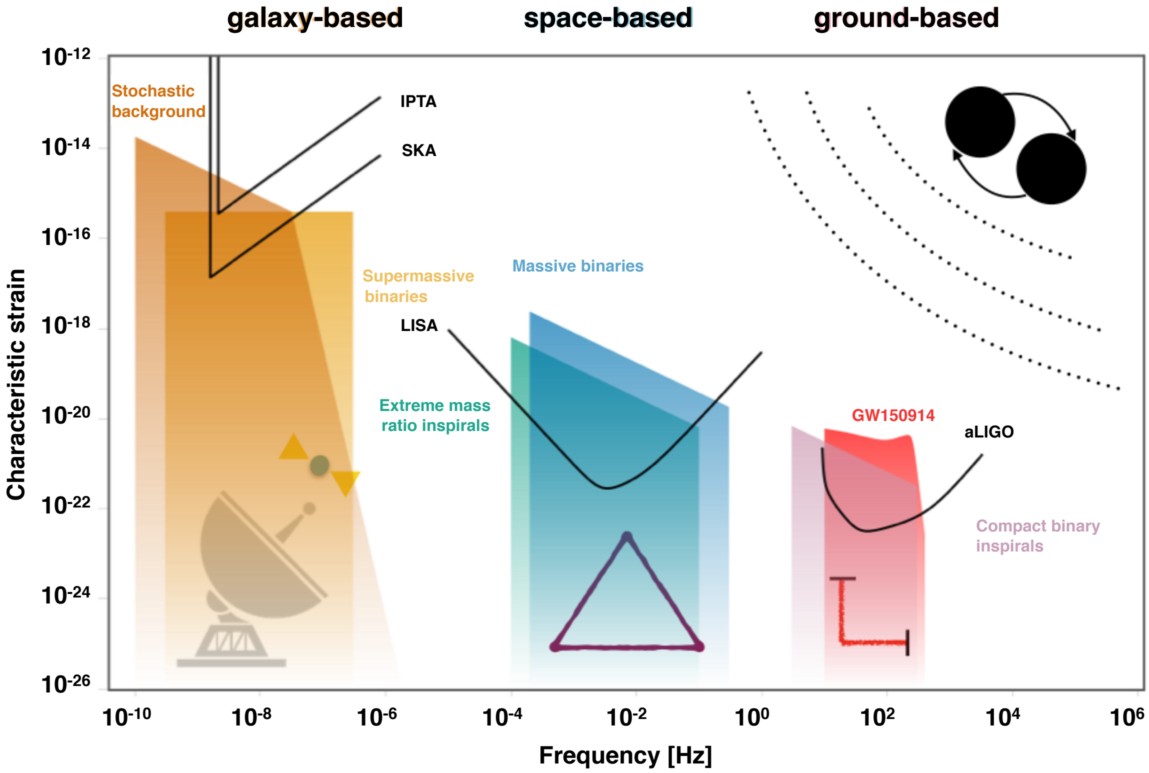

FIGURE 1: The gravitational wave (GW) spectrum with shaded colored regions corresponding to the different black hole binary sources. Sensitivity across the entire range requires various GW detectors operating within different frequency limits. Highest frequency phenomena like compact stellar mass black hole binaries are accessible to ground-based detectors, including aLIGO (advanced LIGO). Future space-based instruments like LISA will find sources emitting within $10^{-4} - 1\,\mathrm{Hz}$. Pulsar Timing arrays serve as a galactic-scale experiment receptive to the longest period binaries. Current and future sensitivity curves of these detectors are in black, reporting the strength of GW signal measurable by each instrument. Figure adapted from C. Moore, R. Cole, and C. Berry, with modifications by C.F. Mingarelli.

large number of pulses collected during each timing observation, e.g., we record $\sim 10^6$ pulses for the typical observation of a 10 ms pulsar. This allows us to *fold* the pulse data, stacking individual pulses to create an average pulse profile we can then measure to high precision. Current PTAs can accurately predict TOAs to a few tens of nanosecond precision.

Non-zero timing residuals of this order may be due to incorrect parameter estimates for the pulsar (for instance, pulse period, pulsar spin-down, astrometry, etc.), but, they may also be from a spacetime disturbances encountered during pulse propagation to Earth, such as gravitational wave. Depending on the considered source of GWs, we tailor a search to mine PTA data for evidence of temporal offsets consistent with the polarized stretch and squeeze induced by a GW.

Currently, there are three PTA experiments in operation: the North American Observatory for Gravitational Waves (NANOGrav; McLaughlin 2013), the European Pulsar Timing Array (EPTA; Desvignes et al. 2016), and the Parkes Pulsar Timing Array (PPTA; Hobbs 2013). Together these groups form the International Pulsar Timing Array (IPTA; Verbiest et al. 2016a) which periodically releases combined datasets (Verbiest et al. 2016b). Efforts are in place to continue these experiments as they reach sensitivities consistent with predicted GW signal levels. Dedicated telescope time ensures PTAs can not only continue to monitor well timed pulsars but embark on searches for more MSPs.

1.3 Supermassive Black Hole Binaries

A promising source of GWs in the PTA-band arises from a cosmological population of inspiraling supermassive black hole binaries with total mass $\sim 10^{7} - 10^{10}\,\mathrm{M}_{\odot}$ (Sesana et al. 2004; Sesana 2013; Burke-Spolaor et al. 2018). Binaries this massive are theorized to develop through accretion and repeated minor mergers of galaxies, each seeded with a central black hole. Since their formation is thought to be a direct consequence of the hierarchical clustering of large-scale structures, gravitational observation of these systems has widespread implications for cosmology.

During a galaxy merger, the resident black holes sink to the shared galactic nucleus where interactions with surrounding dark matter, gas and stars extract energy, shrinking the binary orbit until the black holes form a bound pair. If enough energy is siphoned by the environment, e.g., stellar scattering or accretion, the binary will reach sup-parsec separations at which point gravitational radiation from an accelerating quadrupole mass-moment dominates the orbital decay. Provided the black holes are sufficiently close, this GW emission will succeed in drawing the binary to coalescence within a Hubble time. The cosmic timeline of binary evolution strongly depends on galaxy morphology and contents; also, the expected rate and brightness of the signals from these systems strongly depend on the time spent undergoing various modes of energy loss (Kelley et al. 2017b).

There are well-publicized correlations between the mass of a SMBHB and observable parameters of its host galaxy, such as bulge mass or the velocity dispersion of stellar matter in the bulge (Gültekin et al. 2009), hinting that large galaxies and binaries are thought to grow symbiotically. These relations are sewn into population syntheses aimed at predicting the rate and strength of signals with respect to detector sensitivity from these sources, as further discussed in Chap. 2.

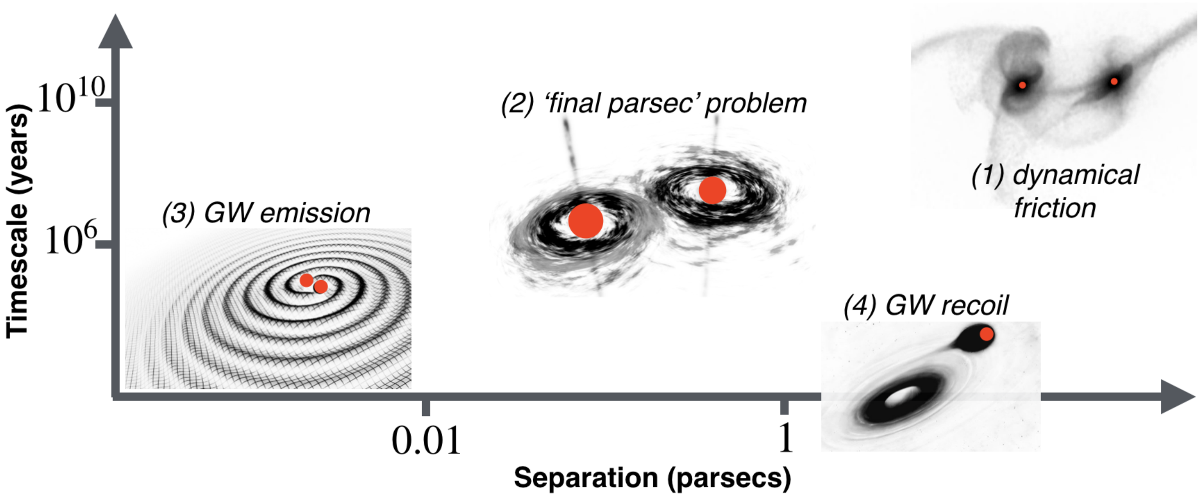

Figure 2: Schematic of the various mechanisms thought to contribute to SMBHB orbital decay over cosmological timescales. Depending on the available matter for coupling, a binary can be efficiently drawn to coalescence within a Hubble time or stall toward parsec separations. (Figure from Taylor et al. (2019); further credit therein).

1.4 SMBHBs in PTA data

SMBHBs imprint themselves in PTA data in two ways: individually and en masse. Historically, PTA data analyses focus on the stochastic GW background (GWB) formed by a population of inspiraling SMBHBs, as models predict that this signal is expected to be detected first (Rosado et al. 2015). Although the GWB contains the signatures of individual binaries, these properties are washed out in favor of the emergent properties describing the signal of the ensemble. Thus, our detection searches focus on extracting the statistical properties of the background, namely the power spectrum and the cross-power correlations of pulsar residuals.

A comprehensive derivation of the theoretical power spectral density (PSD) for a GW background can be found in numerous works (e.g. Allen & Romano 1999; Thrane

& Romano 2013). If we assume the energy density of the stochastic background suits a power-law description, and whose constituent gravitational waves collectively form an isotropic, unpolarized, stationary background of zero mean, then the strain power spectral density is

$$S_h(f) = \frac{A_{\rm gwb}^2}{12\pi^2}\left(\frac{f}{f_{\rm yr}}\right)^{-\gamma} f_{\rm yr}^3 \tag{1.4.1}$$

where $A_{\rm gwb}$ is the frequency-independent background amplitude, $f_{\rm yr} \equiv 1/(1\,{\rm yr})$, and γ dictates the spectral distribution of power. In the case of an ensemble of circular binaries evolving to increasingly smaller orbital separations by radiating energy through GW emission, $\gamma = 13/3$ (Phinney 2001).

The power spectral density due to the GWB within a single detector can be written as

$$P_h = \mathcal{R}(f) S_h(f) \tag{1.4.2}$$

with $\mathcal{R}(f)$ serving as the transfer function between GWB and the detector response power, i.e., the antenna pattern of the detector over polarizations and directions on the sky.

Any individual pulsar can, in theory, serve as a single detector, however, many MSPs show evidence for noise that can also be well described by a power-law with amplitude and spectral character similar to that of the background (Arzoumanian et al. 2018a). A network of pulsars provides addition means by which to increase signal-to-noise, introducing cross-correlated detector response terms.

The GWB-induced power shared by the residuals of pulsar a and pulsar b is

$$P_{h,ab} = \Gamma_{ab} S_h(f) \tag{1.4.3}$$

where Γ_{ab} is the *overlap reduction function* which scales the degree of cross-correlated power. An isotropic population of inspiraling SMBHBs induces unique quadrupolar cross-correlations in any pair of pulsar residuals. The cross-correlation coefficients depend only on the angle between the pulsar pairs and form the so-called Hellings and Downs curve (originally conceived in Hellings & Downs (1983) and rigorously detailed in Anholm et al. (2009a)).

Identifying not only the correct power spectrum but also measuring the proper cross-correlations is needed to claim a confident GWB detection with PTAs.

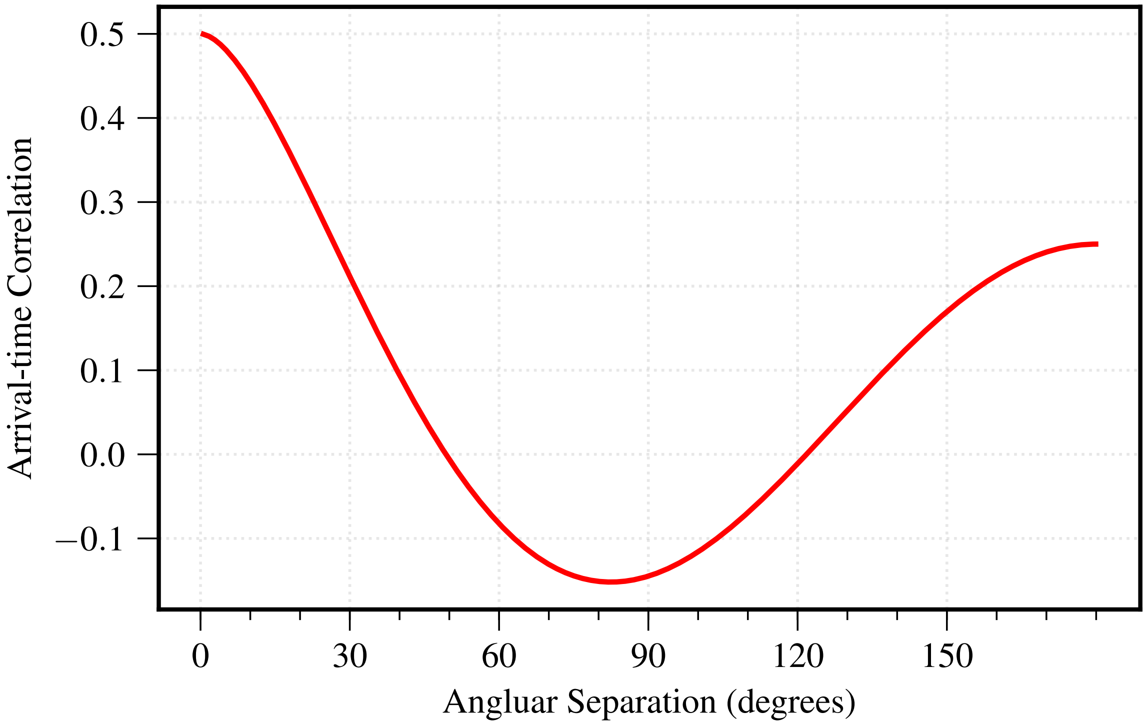

FIGURE 3: Hellings and Downs cross-correlation coefficients as a function of pulsar pair angular separation. Pulsar timing residuals expect to be related via this spread of power for gravitational waves from an isotropic distribution of inspiraling supermassive black hole binaries

We turn now to individually resolvable SMBHBs. PTAs are also sensitive to GWs emitted from nearby individual SMBHBs with periods on the order of months to years, total masses of $\sim 10^8 - 10^{10}\,\mathrm{M}_\odot$, and orbital separations of $\sim 10^{-3} - 10^{-1}\,\mathrm{pc}$, depending on the total mass of the binary. SMBHBs that are emitting in the PTA band have nearly-constant orbital frequencies, and hence the GWs from these sources are referred to as "continuous waves" (CWs). We treat these as a deterministic signals in our data with recoverable frequency, phase, and amplitude. PTAs may also be sensitive to bursts with memory. This is the DC component of the chirp signal emitted just prior to binary coalescence. In general, PTA sources evolve too slowly to directly detect the coalescence during the length of a typical PTA experiment.

Our target signals are present during the entirety of data collection. Consequently, PTA analyses use simultaneous modeling of interesting and nuisance sources of noise. This provides unique challenges when optimizing our detection analyses to accommodate all the available degrees of freedom. A boon, though, is that longer observation baselines grant increasing signal-to-noise to any astrophysical signal. This means there is no fundamental low-frequency limit – our sensitivity will continue to extend into lower-frequency territory without any additional instrumentation.

1.5 Dissertation Outline

In this dissertation, we explore aspects of all three types of GW signals described above. In the following chapter, we estimate detection prospects for bursts with memory using an observation-based simulation framework to construct SMBHB populations evolving to coalescence. This requires further our understanding of how this non-oscillatory GW signal appears in PTA data. We evaluate expected rates for various signal strengths and signal-to-noise ratios with respect to both current PTAs and LISA. Chapter 3 discusses NANOGrav's latest search for individual inspiraling SMBHBs in PTA data. We expound on methods for noise modeling and Bayesian procedures used in typical PTA data analysis. Our results are then used to inform astrophysical constraints of binary candidates. Lastly, Chapter 4 describes a new hybrid detection-frequentist statistic useful for discriminating the presence of a GWB power spectrum in PTA data without the need for comprehensive Bayesian analysis. Finally, we present a brief summary of results.

Part I

Deterministic Gravitational Wave Signals

Chapter 2

Prospects for Memory Detection with Low-Frequency Gravitational Wave Detectors

> “*All that you touch you Change.*
> *All that you Change Changes you.*
> *The only lasting truth is Change.*”
>
> — Octavia Butler, *Parable of the Sower*

This chapter is based on:
Prospects for Memory Detection with Low-Frequency Gravitational Wave Detectors
K.Islo et al. 2019
Submitted

2.1 Introduction

Non-linearities in Einstein’s field equations suggest that during the coalescence of compact objects, the oscillatory GW radiation is accompanied by a monotonically increasing or decreasing component of the strain. The result is an enduring change of gravitational potential in the wake of a propagating GW, a feature termed “memory” (Christodoulou 1991; Blanchet & Damour 1992; Thorne 1992). For example, a GW with memory passing through a system of two isolated, free-falling test masses would permanently stretch or compress the co-moving distance between them.

Like the oscillatory component of a GW, memory is sourced by a changing time derivative of the system's mass multipoles, but only in part. It also grows through the cumulative history of GW emission. As such, the memory signal inherits the radiating system's evolving past: its strength at any time the result of the integrated history of the system. For an SMBHB undergoing a merger, the memory signal initially displays negligible growth corresponding to the slow time evolution of the binary's inspiral. Once the binary enters its most dynamic phase during coalescence, the system emits a 'burst' of memory signal which propagates outwards. Observations of SMBHB merger memory events would shed light on strong-field effects of General relativity, and provide information about SMBHB properties augmenting that obtained from the oscillatory components. More broadly, a memory observation would provide robust evidence of fundamental symmetries in GR (Strominger 2017).

In this chapter, we estimate the current and future potential to detect GW memory from SMBHB coalescence based on a simulation-suite of semi-analytic models for the SMBHB population. Importantly, the models are based in local observables for SMBHBs, and thus encompass only uncertainties from local mass functions, galaxy merger timescales, etc. rather than uncertainties in casting dark matter halo simulations to a binary supermassive black hole population (as has been the case for other similar simulations in the past (Enoki et al. 2004; Sesana et al. 2008; Ravi et al. 2012; Kulier et al. 2015)). This study is also novel from previous studies in that we expand our models to include lower black hole masses, down to $M_{\mathrm{BH}} \gtrsim 10^5$, and higher redshifts, such that the memory signals expected in the bands relevant to both PTAs and LISA can be explored. We also take into account the unknown decoupling radius for binary-host interactions. We look at the likely rate and SNR for memory events detectable within these frequency bands.

2.2 SMBHBs: A Viable Source of Memory

SMBHBs form during major mergers, and harden through repeated interactions with their galactic environment (Begelman et al. 1980; Volonteri et al. 2003). The effectiveness of the mechanisms by which these black hole systems are driven to merger is an open question in astrophysics, although many studies conclude these binaries will coalesce given sufficiently gas-rich galactic environments or continuous rates of stellar loss-cone refilling (Colpi 2014; Roedig et al. 2011; Khan et al. 2013; Vasiliev et al. 2015). SMBHBs that eventually merge are candidates for producing strong GW memory bursts.

The persisting spacetime offset induced by a passing GW memory signal is proportional to $\Delta E_{\rm rad}/D$, where $\Delta E_{\rm rad}$ is the total energy radiated in the form of GWs and D is the co-moving distance to the coalescence. For equal-mass SMBHBs, the energy available for the GW memory burst ranges from 5% to 10% of the total binary energy. The precise value depends on binary inclination and the degree of black hole spin-alignment (see Pollney & Reisswig 2011). For example, an optimally-oriented binary consisting of two $10^9\,\mathrm{M}_\odot$ black holes coalescing 1 Gpc away from Earth will emit a GW memory burst with amplitude $h_{\rm mem} \sim 10^{-15}$. per Kaplan: how does this change for various inclinations and orientations?

Population synthesis: To find SMBHB coalescence rates we use the population synthesis model described in Simon & Burke-Spolaor (2016). This simulation is constructed using empirical quantities that characterize galaxy mergers – namely, galaxy stellar mass function, galaxy pair fraction, and galaxy merger timescale. These are cast from galaxy pairs into inferred SMBHBs and mergers using the empirical relationship that has been found between host galaxy bulge mass and black hole mass (McConnell & Ma 2013; Kormendy & Ho 2013; Shankar et al. 2016). Our usage of the Simon & Burke-Spolaor (2016) prescription is as follows: In order to evaluate the number of bursts occurring per year, we start with the galaxy merger rate. The probability of a single galaxy completing a major merger is proportional to $f_{\rm pair}/T_{\rm mg}$, where $f_{\rm pair}$ is the fraction of galaxies in dynamically close pairs and $T_{\rm mg}$ is the merger timescale.

The probability of a galaxy of mass M becoming paired with another galaxy of mass qM, where $q < 1$, at a redshift z is,

$$\mathcal{R}(z, M, q) = \frac{\mathrm{d}f_{\rm pair}(z, q)}{\mathrm{d}q} \frac{1}{T_{\rm mg}(z, M)} \frac{\mathrm{d}t}{\mathrm{d}z}, \tag{2.2.1}$$

where $\mathrm{d}t/\,\mathrm{d}z$ is a cosmological term converting the proper time rate to a redshift rate. This differential merger rate along with the galaxy stellar mass function (GSMF) Φ, produces the differential number density of galaxy mergers,

$$\frac{\mathrm{d}^3 n}{\mathrm{d}z\,\mathrm{d}M\,\mathrm{d}q} = \Phi(z, M)\,\mathcal{R}(z, M, q). \tag{2.2.2}$$

In order to represent our merger rate as one of SMBHBs, we cast the host galaxy mass into a black hole (BH) mass using an $M_{\mathrm{BH}} - M_{\mathrm{bulge}}$ scaling relation. Then we can define the number of SMBHB mergers in an interval dt on Earth as,

$$\frac{\mathrm{d}^4 N}{\mathrm{d}t\,\mathrm{d}z\,\mathrm{d}M\,\mathrm{d}q} = \frac{\mathrm{d}^3 n}{\mathrm{d}z\,\mathrm{d}M\,\mathrm{d}q}\,\frac{\mathrm{d}V_c}{\mathrm{d}z}\,\frac{\mathrm{d}z}{\mathrm{d}t}, \tag{2.2.3}$$

where V_c is the co-moving volume at a redshift z. Note that this differs from the formulation of Simon & Burke-Spolaor: Instead of deriving binary coalescence rates, they calculated the number of binaries contributing to the strain in each GW frequency bin. We restrict the parameter space to include only what we know observationally. Specifically, we take the redshift to be $z < 3$, the primary galaxy mass to be $10^8\,\mathrm{M}_\odot < \mathrm{M} < 10^{12}\,\mathrm{M}_\odot$, and the mass ratio to be consistent with 'major mergers' with $0.25 < q < 1$.

Evolution of SMBHBs: Evolution of SMBHBs: At the early stages of galaxy merger, we expect dynamical friction to reduce the orbital angular momentum of the individual black holes until they sink to the center of the merger remnant, forming a SMBHB (Begelman et al. 1980; Volonteri et al. 2003). The dominant mode of energy loss below ∼10 pc binary separation is not yet understood, although many environmental interactions potentially contribute (Merritt & Milosavljević 2005). As such, it is unclear when the environment decouples from the binary, after which GW emission dominates. This has been the center of a growing debate in predicting the strength and spectral shape of the stochastic gravitational wave background in the nanohertz frequency band (e. g. Arzoumanian et al. 2018c; Taylor et al. 2017; Rasskazov & Merritt 2017; Kelley et al. 2017a).

Here, we incorporate varying environmental influence to predict GW memory bursts for populations with different evolutionary timescales. We introduce a parameter we term the *decoupling radius* a_d to be the orbital separation at which the binary's evolution becomes GW-dominated. SMBHBs embedded within sparse environments exhaust their environmental interactions earlier and have the potential to stall before reaching a regime where GW-radiation can drive the binary to merger (i.e., $a_d > 1\,\mathrm{pc}$). The opposite scenario involves a binary strongly coupled to its environment, undergoing extremely efficient orbital shrinking, and reaching coalescence quickly ($a_d \to 0$).

If we assume circular binaries, the time to coalescence can be expressed in terms of Keplerian parameters and chirp mass $\mathcal{M} = (M_1 M_2)^{3/5}/(M_1 + M_2)^{1/5}$

$$\tau_{\mathrm{GW}} = \frac{5\,c^5}{256\,G}\,\frac{a_d^4}{\mathcal{M}^{5/3} M_{\mathrm{tot}}^{4/3}}. \tag{2.2.4}$$

Establishing a_d is therefore akin to specifying the total time to binary coalescence, i.e., $t_{\mathrm{coal}} = t_{\mathrm{Gal}} + \tau_{\mathrm{GW}}$, where t_{Gal} is the time between galaxy merger and SMBHB formation. τ_{GW} contributes a difference in the redshift between SMBHB formation and coalescence – z_{merge} vs. z_{coal}. We must be careful to evaluate the last two factors in Eq. (2.2.3) at this lower redshift if we are to interpret it as a memory burst rate as seen at Earth.

Addition of lower-mass black holes: Previous simulations of SMBHB populations (such as that in Simon & Burke-Spolaor (2016)) only include progenitor galaxy mass greater than $10^{10}\,\mathrm{M}_\odot$ which corresponds to individual black hole mass of $\sim 10^7\,\mathrm{M}_\odot$. Lower galaxy masses are likely not relevant PTA-band sources, but $10^5 - 10^7\,\mathrm{M}_\odot$ black hole binaries can produce GW memory signals which enter the LISA band. We use results from the Galaxy And Mass Assembly survey as outlined in Wright et al. (2017) to estimate the distribution of these lower mass binaries at $z < 0.1$ along with those from the ULTRAVista survey reported in Ilbert et al. (2013) for higher redshifts.

Owing to their comparatively lower luminosity, observations do not constrain the distribution of these binaries well. As a comparison to the observations of the galaxy stellar mass function (GSMF) for these lower mass systems, we also use a GSMF fit to simulated binary populations as described in Barausse (2012) and Sesana et al. (2014).

2.3 GW Memory Signal Model

2.3.1 Memory signal from SMBHBs

Over the binary's lifetime memory undergoes a slow growth prior to merger, rapid accumulation of power during coalescence, and eventual saturation to a constant value at ringdown. Thus, in the time domain, the signature of a memory signal from a SMBHB can be approximated by a step-function centered at the moment of coalescence:

$$h_{+,\times}^{\rm mem}(t) = \Delta h_{+,\times}^{\rm mem}\Theta(t), \tag{2.3.1}$$

where Θ is the Heaviside-step function and $\Delta h_{+,\times}^{\rm mem}$ is the net difference between the spacetime perturbation long before and long after the passage of a gravitational wave with memory, i.e.,

$$\Delta h_{+,\times}^{\rm mem} = h(t \to +\infty) - h(t \to -\infty). \tag{2.3.2}$$

In the frequency domain,

$$\tilde{h}_{+,\times}^{\rm mem}(f) = \begin{cases} \dfrac{i\Delta h_{+,\times}^{\rm mem}}{2\pi f} & \text{for } 0 < f < f_c, \\ 0 & \text{for } f \geq f_c, \end{cases} \tag{2.3.3}$$

where f_c, the *cut-off frequency*, corresponds to twice the orbital frequency at coalescence. Frequencies larger than f_c do not contribute to the GW signal. We can approximate $f_c \sim \tau^{-1}$, where τ is the light crossing time of the merger remnant. τ is also the timescale for the rise of the memory signal during the merger. Eq. (2.3.3) is sufficient for PTA signal-to-noise ratio calculations, however, we include a minor correction since LISA may be able to resolve the time varying features of the memory signal between onset of coalescence and ringdown. From Favata (2009),

$$\tilde{h}_{+}^{\rm mem}(f) \approx i\frac{\Delta h_{+}^{\rm mem}}{2\pi f}\left[1 - \frac{\pi^2}{6}(\tau f)^2\right], \tag{2.3.4}$$

The magnitude of the spacetime offset $\Delta h_{\rm mem}$ is affected by black hole spin-alignment, with the maximally aligned spinning case exhibiting the strongest signal Pollney & Reisswig (2011). Higher-order spin effects should be incorporated in future simulations to properly reflect the saturated memory amplitude. We use the spin-averaged formula for circular binaries as given by Madison et al. (2017),

$$\Delta h_{+}^{\rm mem} \simeq \frac{(1-\sqrt{8}/3)}{24}\frac{G\mu}{c^2 D}\sin^2\left(17+\cos^2\iota\right), \tag{2.3.5}$$

$$\Delta h_{\times}^{\rm mem} = 0 \tag{2.3.6}$$

with dependence on additional binary parameters reduced mass μ, inclination angle ι, and co-moving distance D.

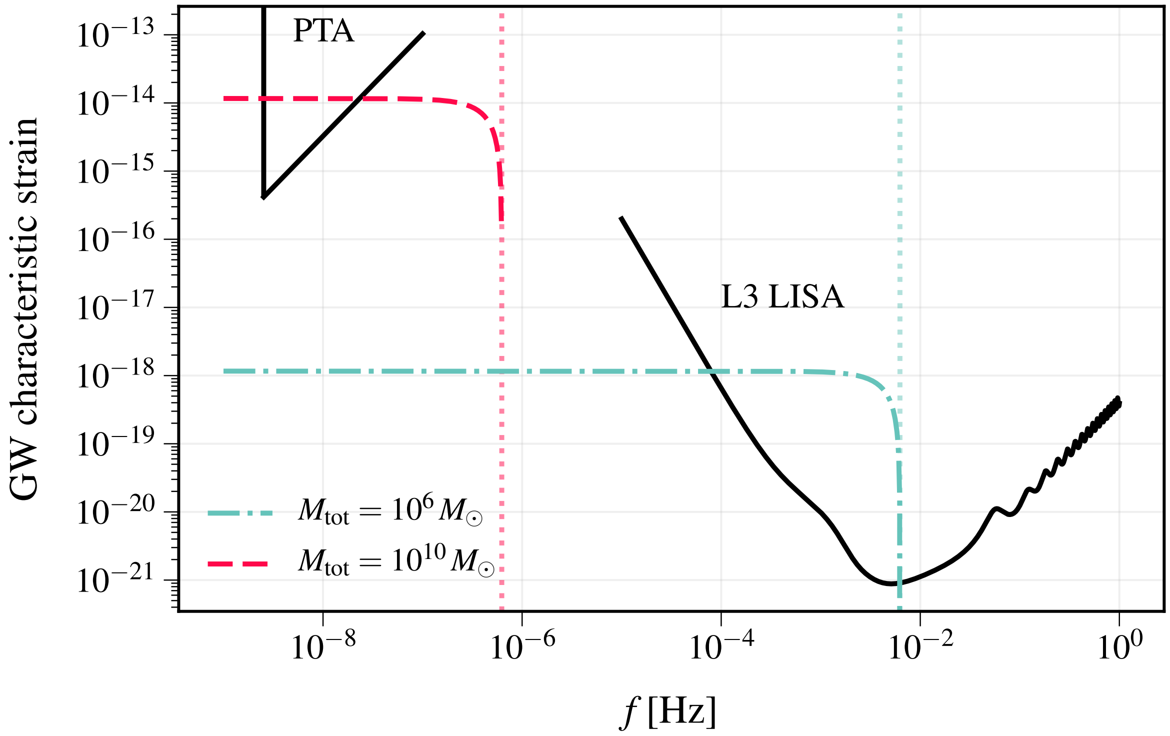

FIGURE 4: The low-frequency end of the gravitational wave spectrum. Estimated LISA sensitivity to individual black hole binary mergers as of the LS Mission Proposal (Amaro-Seoane et al. 2017) along with the sensitivity of a 12.5-year baseline PTA with biweekly observations, nanosecond root-mean-squared timing residuals, and only white noise (Moore 2015). Plotted are memory signals from binaries with total mass 10^{10} $M_\odot$ and 10^6 $M_\odot$ at $z = 1.0$ with mass ratio 0.25. (magenta and teal, respectively.) Include references here used to make curve

2.3.2 Signal-to-Noise Ratio

Previous studies expound on ideal strategies for searching and conducting analyses of memory signals (e.g. Braginskii & Thorne 1987). In this work, rather than discuss best practices, we evaluate the astrophysical motivation for implementing searches of GW memory signals from a cosmic population of SMBHBs, specifically focusing on the signal-to-noise for memory events within LISA and PTAs.

LISA

Proposed space-based interferometers like LISA would respond to a burst with memory as a permanent change to the proper distance between its free mirrors. In theory, the signal may be stored forever. Ground-based interferometers are not optimized for memory detection, requiring stacking of stellar-mass coalescence signals to accumulate SNR (Lasky et al. 2016) or waiting on new space-based observatories like the Big Bang Observatory to detect memory from GW150914-like events (Johnson et al. 2019). Here we calculate the SNR for memory from higher mass binaries, relevant for LISA.

The signal-to-noise ratio (SNR) is computed using an optimal match filter to extract a known signal model from noisy data as described in Flanagan & Hughes (1998). With Eq. (2.3.4) as our memory signal model, the average SNR is given by

$$\begin{aligned} \langle \rho^2 \rangle &= 4\,\mathrm{Re} \int_0^\infty \frac{\langle \tilde{h}_{\mathrm{mem}}(f) | \tilde{h}^*_{\mathrm{mem}}(f) \rangle}{S_f} df \qquad (2.3.7) \\ &= \frac{\Delta h^2_{\mathrm{mem}}}{\pi^2} \int_0^{f_c} \frac{\mathrm{d}f}{f^2} \left(1 - \frac{\pi^2}{6} (\tau f)^2 \right)^2 \frac{1}{S_n(f)}. \qquad (2.3.8) \end{aligned}$$

Fig. 5 shows the SNR for a memory burst produced by SMBHBs of total binary mass M_{tot} coalescing at redshift z. Expected SNR ranges from 100 to 10,000 with the highest SNR events coming from binaries at $z < 0.5$ and with $10^5\,\mathrm{M}_\odot < \mathrm{M}_{\mathrm{tot}} < 10^7\,\mathrm{M}_\odot$. Our approximated memory model only minimally diverges at higher redshift from the LISA SNRs reported in Favata (2009). Analysis of LISA mergers within the same range of mass and redshift, indicate that $M_{\mathrm{tot}} < 10^{4.2}\,\mathrm{M}_\odot$ coalescences will occur beyond the LISA frequency band (≥ 1 Hz), in which case, the memory signal will be the dominant coalescence signature. "Orphan memory" signals, as they are termed in McNeill et al. (2017), effectively increase the high-frequency limit of the detector, allowing access to previously forbidden binary masses and redshifts.

As LISA data analysis stands, the degree to which a memory burst will be distinguishable from any other cause of detector response relies on the co-incidentally modeled signals and sources of noise. Future work will hopefully elaborate on best practices for memory signal extraction in light of LISA's unique GW window.

Pulsar Timing Arrays

PTAs measure the times of arrival of radio pulses from an array of galactic millisecond pulsars. GW memory manifests in pulsar timing observations as an abrupt increase or decrease in pulsar rotational frequency. PTAs currently target their search to the

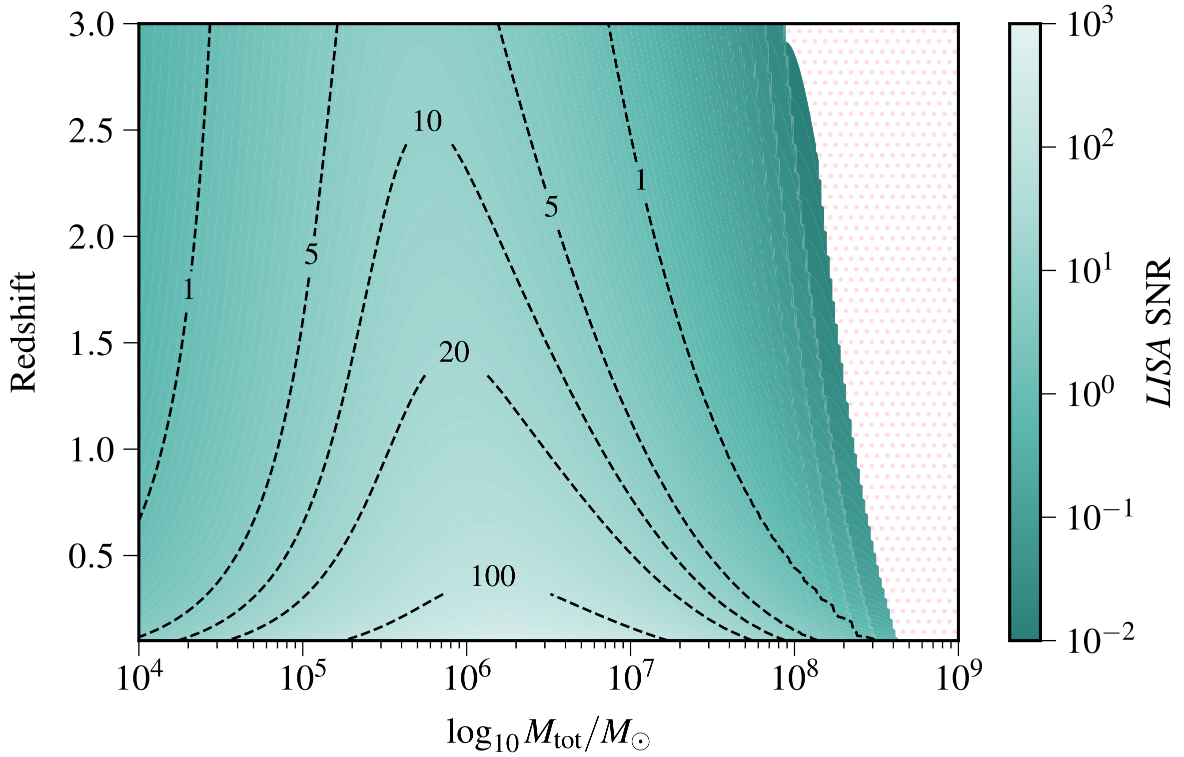

FIGURE 5: LISA signal-to-noise ratio for GW memory bursts from optimally-beamed (e.i., $\iota = \pm\pi/2$), equal-mass SMBHBs with $M_{\rm tot}$ between $10^4\,{\rm M}_\odot - 10^9\,{\rm M}_\odot$ and coalescence redshift between $0.1 - 3.0$. Dots fill the space in which the memory burst signal falls outside the LISA band due to strain amplitude below the noise amplitude or a cut-off frequency below the minimum band frequency. The hatched region encompasses binaries for which the memory signal will be observable, while the oscillatory merger signal will not.

stochastic background formed by the superposition of GW emission from a cosmological population of inspiraling SMBHBs, but, with longer baseline observations, will eventually be able to effectively marginalize over intrinsic pulsar noise to detect deterministic GW signal from an individual binary, which includes, in theory, GW memory.

GW memory manifests in pulsar timing observations as an abrupt increase or decrease in pulsar rotational frequency. A burst passing Earth will affect the timing residuals for all pulsars in the array in a correlated way. Leveraging the data from many pulsars is crucial when claiming a confident memory detection within PTA data as intrinsic noise in individual pulsars can mimic a burst. Below we derive the SNR of a memory event within a single pulsar's residuals – the SNR using an entire array is easy to estimate from this single pulsar derivation.

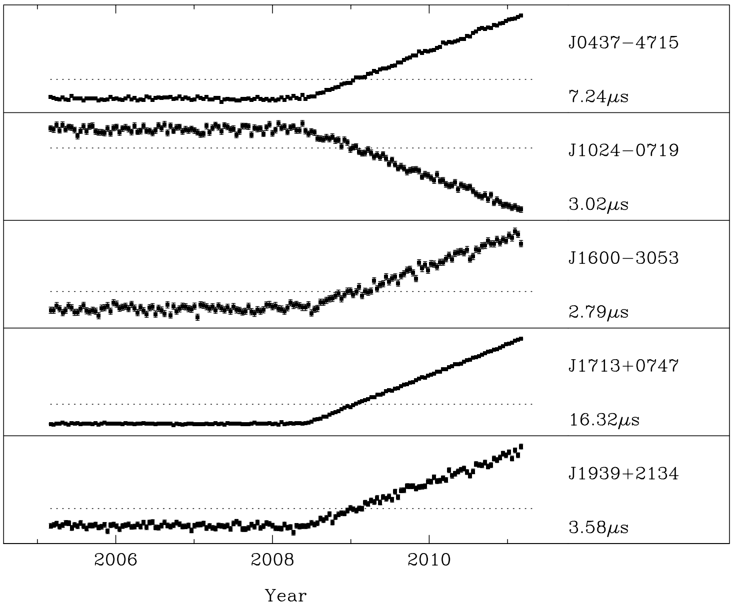

FIGURE 6: From Wang et al. (2015). Pre-fit pulsar timing residuals injected with a burst with memory. The time advance or lag of pulse arrival times depends on the black hole spin alignment of the source binary.

The timing residual for the jth pulsar accumulated during a length of observation t is computed by integrating the redshift of the signal over the time interval.

$$R_j(t,\hat{\Omega}) = \int_{t_0}^{t} dt'\, z_j(t',\hat{\Omega}), \tag{2.3.9}$$

where z is the redshift corresponding to the rate of arrival of signals from pulsar x and $\hat{\Omega}$ is the propagation direction of the GW contributing to the redshift modulation. Since the burst signal effectively alters the pulse redshift, we can write Eq. (2.3.9) as

$$\begin{aligned} R_j(t,\hat{\Omega}) &= B_j(\hat{\Omega}) \int_{t_0}^{t} dt' \Theta(t - t_{\mathrm{mem}}) & (2.3.10) \\ &= B_j(\hat{\Omega})(t - t_{\mathrm{mem}})\Theta(t - t_{\mathrm{mem}}). & (2.3.11) \end{aligned}$$

The Fourier transform of the integrated time-series residuals is

$$\tilde{R}_j(f) = -B_j \Delta h_{\mathrm{mem}} \frac{e^{ift_{\mathrm{mem}}}}{2\pi f^2}, \tag{2.3.12}$$

where B_j is a geometrical factor accounting for pulsar position, GW polarization and propagation direction. PTA SNR is similarly determined through a matched filter analysis. From van Haasteren & Levin (2010):

$$\langle \rho_j^2 \rangle = 4\,\mathrm{Re} \int_{1/T_{\mathrm{obs}}}^{\infty} \frac{|\langle \tilde{R}_j(f) | \tilde{R}_j^*(f) \rangle|}{S_r(f)}\, \mathrm{d}f, \tag{2.3.13}$$

where S_r is the power-spectral density of the timing residuals and $T_{\rm obs}$ is the total time span of pulsar observation. Assuming the residuals are white and Gaussian, the one-sided power-spectral density can be written as

$$S_r(f) = 2\sigma^2 \Delta t, \tag{2.3.14}$$

where σ is the white noise root-mean-square errors on the pulsar residuals and $1/\Delta t$ is the observation cadence (typically ~ 1/month). Inserting Eq. (2.3.12) and Eq. (2.3.14) into Eq. (2.3.13), and averaging over pulsar and source location gives

$$\mathrm{SNR} = \frac{\Delta h_{\rm mem} T_{\rm obs}^{3/2}}{6\,\pi \sigma \Delta t}.^{1} \tag{2.3.15}$$

If a GW memory wavefront were to strike the Earth, we can correlate the entire array's timing residuals over the burst event epoch to uncover frequency changes in every pulsar. Multiple pulsar observations enables more confident GW memory detection by increasing the observed SNR by a factor proportional to $\sqrt{N}_{\rm pulsars}$ (van Haasteren & Levin 2010). We incorporate Eq. (2.3.8) and Eq. (2.3.15) into the simulation outlined in Sec. 2.2 to repeatedly generate instances of an ensemble of coalescing SMBHBs, allowing us to arrive at GW memory event statistics for PTAs and LISA.

2.4 Results

We begin by reporting the number of GW memory events per year with burst amplitude $h_{\rm mem}$ equal to or larger than a given value. We fix the decoupling radius to 0.001 pc, and present the average results recovered from 100 populations initialized with different GSMFs and $M_{\rm BH} - M_{\rm bulge}$ relations in Figure 7. Minor differences occur between the observed and simulated GSMFs (Ilbert (Ilbert et al. 2013) and Sesana (Sesana et al. 2014), respectively). In the most optimistic model (SMBHB masses determined from the Ilbert GSMF and McConnell & Ma (2013)), we expect bursts with $h_{\rm mem} \geq 10^{-16}$ to occur 0.03 times per 1000 years. Previous works estimate similar cumulative rates at this strain

---

[1] N.B. Henceforth we substitute $h_{\rm mem}$ in place of $\Delta h_{\rm mem}$. This is only a semantic change to be consistent with related literature.

amplitude (Ravi et al. 2015; Sesana et al. 2007) for assembled SMBHB populations. The selection bias incorporated into the $M_{\rm BH} - M_{\rm bulge}$ relation from Shankar et al. (2016) diminishes the median rate to less than once per 1000 years. In our combination of inputs, the chosen $M_{\rm BH} - M_{\rm bulge}$ relation proves the most influential (Fig. 7 includes McMa (McConnell & Ma 2013), KorHo (Kormendy & Ho 2013), and Shankar (Shankar et al. 2016)). GW memory detections (or lack thereof) could therefore have the potential to discriminate competing constraints on source parameter relations.

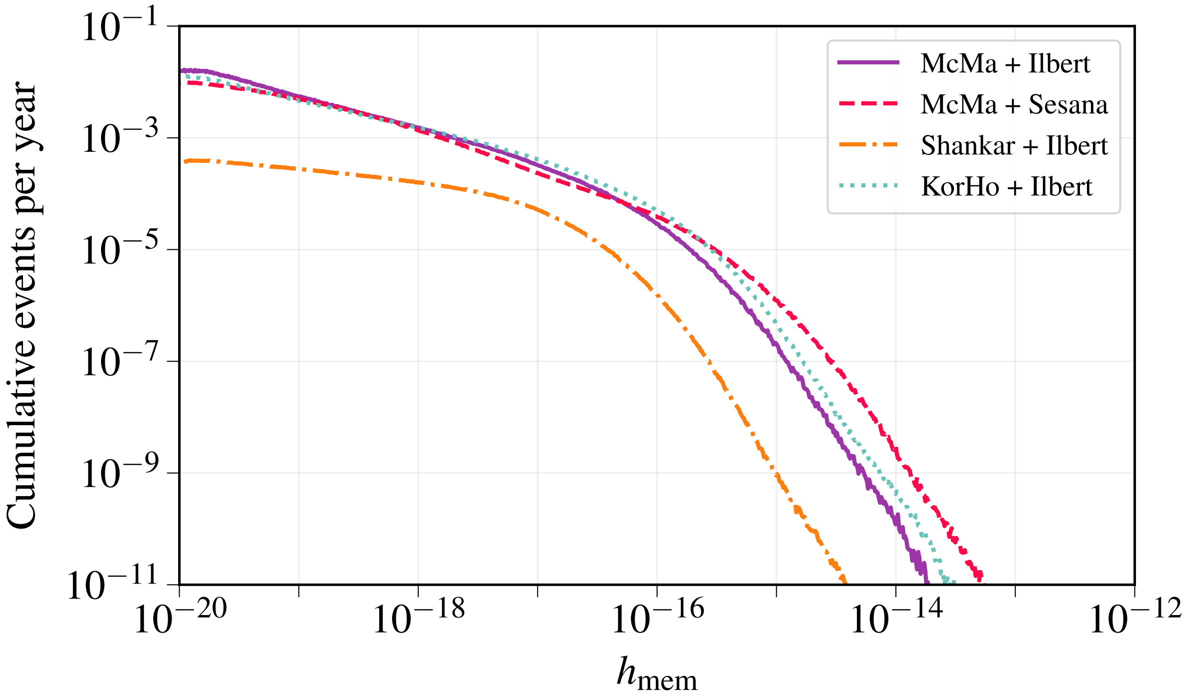

FIGURE 7: Median cumulative memory event rate for bursts with strain amplitude at or above given strain amplitude ($h_{\rm mem}$ from Eq. (2.3.5)) for 100 realizations of a simulated SMBHBs. We include results for populations generated using various $M_{\rm BH} - M_{\rm bulge}$ relations and GSMFs.

To establish the extent to which a more or less efficient environment can influence memory event rates, we vary the decoupling radius to between 0.001 pc and 1.0 pc. From Eq. (2.2.4), we see that increasing a_d yields longer times between galaxy merger and SMBHB coalescence. Decoupling at 1.0 pc returns a population where all but the most massive binaries have stalled. Binary stalling is most likely to occur following the merger of massive elliptical galaxies (Begelman et al. 1980). Therefore, due to the lack of low-mass elliptical galaxies ($M_{\rm Gal} < 10^{10}\,{\rm M}_\odot$), we fix the maximum decoupling radius for the binaries produced in $M_{\rm Gal} < 10^{10}\,{\rm M}_\odot$ to 1 mpc. The subset of stalling binaries

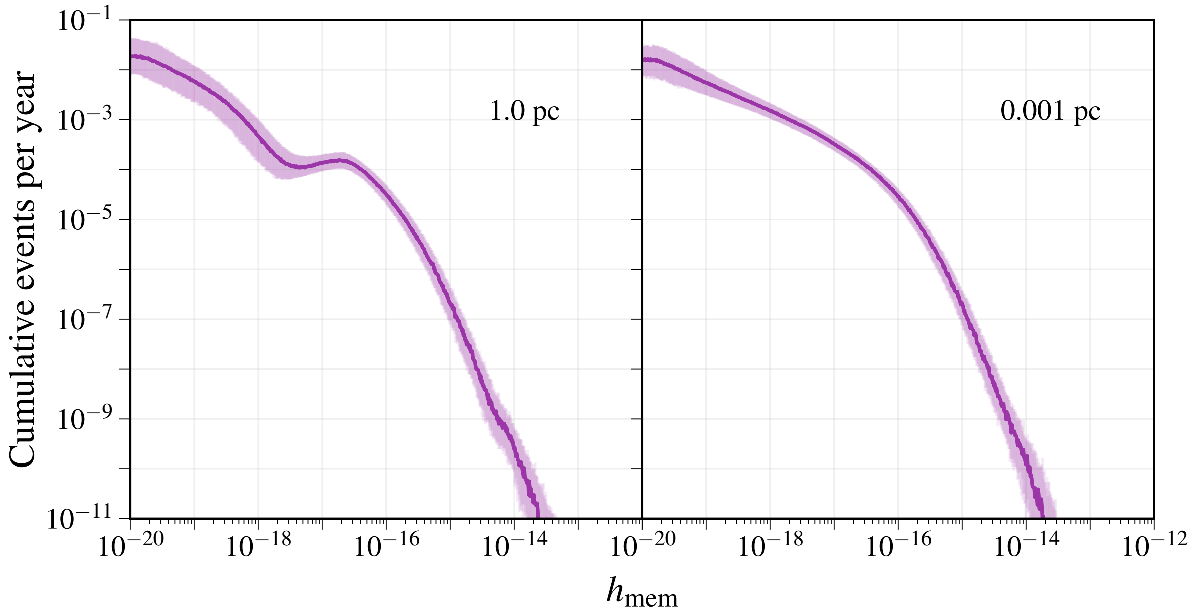

FIGURE 8: Cumulative memory event rate for bursts with strain amplitude at or above $h_{\rm mem}$ for 100 realizations of simulated SMBHBs. The shaded region is the $1-\sigma$ confidence interval from the distribution of simulated event rates using the McConnell and Ma $M_{\rm BH}-M_{\rm bulge}$ relation along with the Ilbert et al. GSMF.

creates a dip at $10^{-18} < h_{\rm mem} < 10^{-16}$ (as seen in Figure 8). However, that dip does not strongly impact the detectable population for either LISA or PTAs. Henceforth our simulated populations assume more efficiently driven environments.

We use Eq. (2.3.8) and Eq. (2.3.15) to calculate the SNRs associated with each memory burst event. Figure 9 shows the median number of burst events per year of given SNR specific to both a 12.5 year PTA observation and LISA for 100 realizations of a simulated SMBHB population using the most optimistic and pessimistic combination of $M_{\rm BH}-M_{\rm bulge}$ relation and GSMF.

Given the low rate of close, high-mass binary coalescences, PTAs systematically show poorer prospects. Memory events with SNR ≥ 5 occur $2.2\times10^{-9} - 9.8\times10^{-9}$ times per year in the optimistic model, and never in the pessimistic model.

Recent results indicate that even treating the pulsar lines of sight as independent lenses through time so that the entire array effectively has an observation baseline equal to the sum of individual baselines, constraints on event rates are far from arbitrating tension

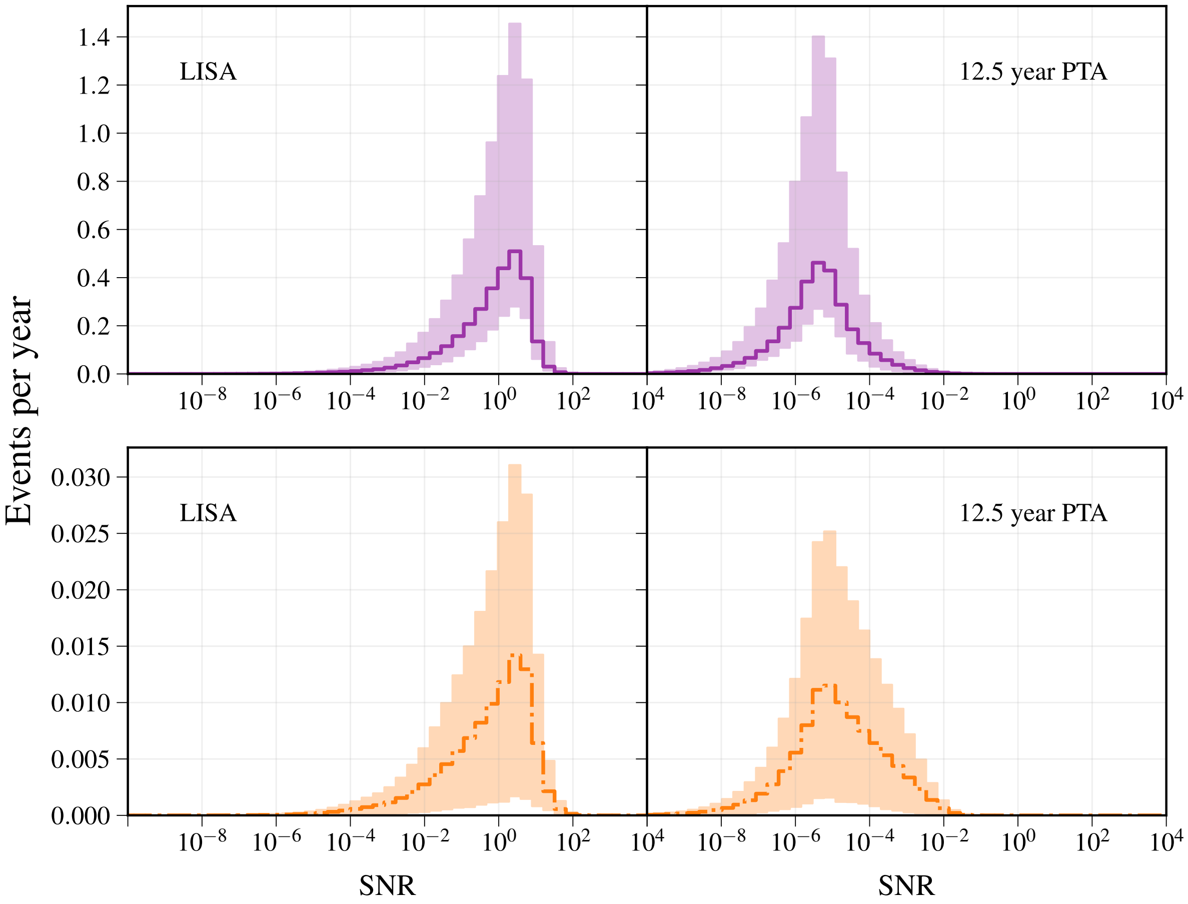

FIGURE 9: Median number of memory bursts per SNR bin from 100 realizations of a simulated population of SMBHBs with 1-σ confidence interval. All panels assume the Ilbert et al. GSMF. We use either the McConnell and Ma or Shankar $M_{\mathrm{BH}} - M_{\mathrm{bulge}}$ relation (upper and lower plots, respectively.)

between population models for these heavier SMBHBs. Stochastic searches currently (and will remain to) serve as the primary point of access to broad-stroke understanding of this population.

LISA demonstrates better prospects with SNR $= 5$ events occurring $0.21 - 1.27$ times per year in the optimistic model, and $0.001 - 0.03$ times per year in the pessimistic model. High-mass binaries have low cut-off frequencies, truncating the signal's power spectrum before entering the LISA band. Low-mass binaries, despite having more efficient environmentally driven decay, evolve too slowly to accumulate appreciable event counts. These two effects work to create a peak event rate at SNR ~ 1 in the upper left panel of Figure 9.

Figure 10 demonstrates the integrated event rate. From this figure, it is clear that given a fly time of 4 years, LISA has the potential to witness at least one memory burst with SNR > 1.

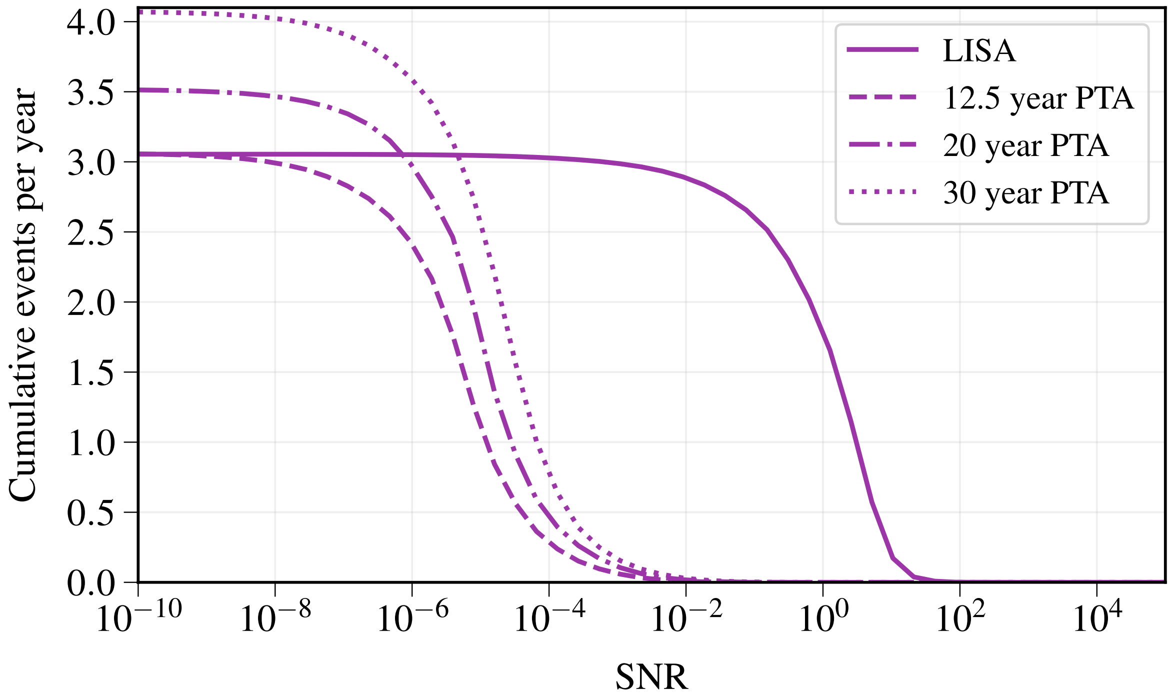

Figure 10: Cumulative number of memory bursts at or above specified LISA/PTA signal-to-noise ratio for same population as in Figure 9.

Figure 11 breaks down the demographics for the binary populations resulting in high SNR memory events with respect to LISA. Using the same average population as in Figure 9, we see that for SNR > 1 the most relevant binaries have total binary mass of $\sim 10^{5.5}\,\mathrm{M}_{\odot}$, coalescing below $z = 0.1$. Such a region of binary mass is relatively unprobed. The last few years have seen numerous searches for active galactic nuclei from the black holes in these local ($z < 1.5$), low-mass galaxies and future time-domain surveys, like the Large Synoptic Survey Telescope, may be able to constrain this population (Baldassare et al. 2018).

Memory detections from these sources would therefore complement our understanding of LISA's target population. As well, studies indicate discernible electromagnetic (EM) counterparts from these mergers (Tang et al. 2018). Paired with an EM trigger, memory signal searches could be conducted with restricted parameter priors to boost SNRs.

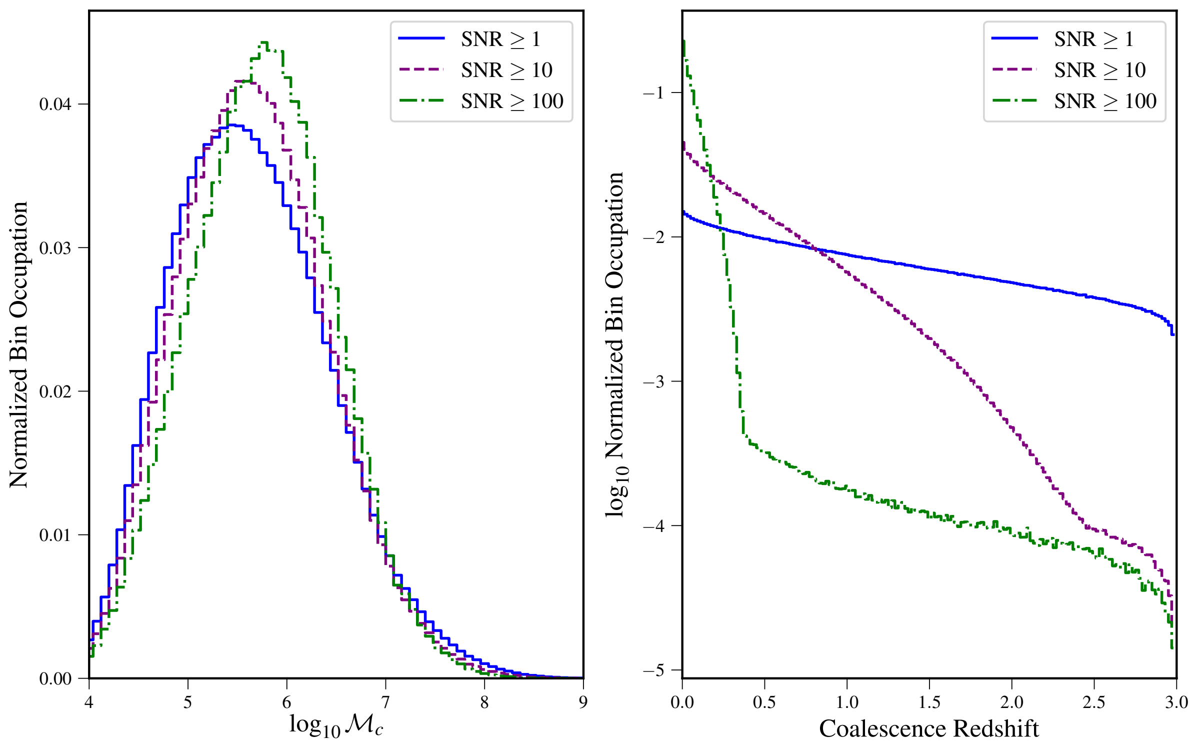

FIGURE 11: Distribution of total binary masses and coalescence redshifts from the upper two panels of Figure 9 yielding LISA SNRs between 1 and 100. Binaries of total mass $10^6\,\mathrm{M}_\odot$ and $z < 1$ form the most promising population of LISA memory sources.

2.5 Outlook and Conclusions

We evaluated the prospects for detecting a GW memory burst from source populations within reach of low-frequency GW detectors. We modified the initial inputs from population synthesis to include galaxy masses below $10^{10}\,\mathrm{M}_\odot$. This probes additional black hole binary masses down to $10^4\,\mathrm{M}_\odot$, and includes observable orphan coalescences from parent mergers outside the LISA band. While an orphan signal is unlikely, given the constraints of our models, it is a signal class worthy of future investigation.

The incorporation of the decoupling radius allowed us to take into account the efficiency of environmentally-driven orbital decay. Even assuming highly efficient environments, PTAs suffer from a sparser population of coalescing SMBHBs, making memory detection from these sources unlikely. This does not, however, preclude the possibility of detection memory from more exotic sources. Projections of PTA sensitivity at the advent of the LISA era strongly depend on the likelihood of PTA's access to increasingly sensitive radio telescopes, such as the planned ngVLA, SKA and DSA2000 (NANOGrav Collaboration 2018; Wang & Mohanty 2018; Hallinan et al. 2019). Detailed explorations of future PTA sensitivities are beyond the scope of this work, although SNRs ≥ 5 remain infrequent at below $1.4 \times 10^{-7}\,\mathrm{yr}^{-1}$ for even a 30 year PTA dataset. SNRs ≥ 5 can occur on average more than $0.5\,\mathrm{yr}^{-1}$ for LISA. If $10^6\,\mathrm{M}_\odot$ binaries are indeed driven to low separations before their orbital decay is dominated by GW emission, LISA has a compelling science case in GW memory.

Chapter 3

Continuous Waves from SMBHBs

> "*But now I'm not so sure I believe in beginnings and endings.*"
>
> — *Arrival,* Adapted from story by Ted Chiang

This chapter is based on:

The NANOGrav 11-Year Data Set: Limits on Gravitational Waves from Individual Supermassive Black Hole Binaries

Aggarwal et al. (The NANOGrav Collaboration, including K. Islo)

Submitted to ApJ

3.1 Introduction

In this chapter, we present the results of a search for GWs from individual circular SMBHBs performed on the NANOGrav 11-year data set (Arzoumanian et al. 2018b; hereafter NG11a). This is an extension of Arzoumanian et al. 2014 (hereafter NG5b), which performed a similar analysis on the NANOGrav 5-year data set. While conducting the search, we also considered the evolution of sensitivity to individual inspiral signals by simultaneously evaluating the 9-year data set (Arzoumanian et al. 2015; hereafter NG9a). This proved incredibly useful as we uncovered sources of unmodeled noise which spurred efforts to develop more advanced noise models and additional Bayesian diagnostic tools for identifying spurious detections. We place upper limits on the strain from individual sources across a range of gravitational wave frequencies and sky positions. Additionally, we establish minimum luminosity distances to potential SMBHBs.

3.2 The 11-year Data Set

A detailed description of the data set can be found in Arzoumanian et al. (2018a), and a summary in Aggarwal et al. (2018). Briefly, we made observations using two radio telescopes: the 100-m Robert C. Byrd Green Bank Telescope (GBT) of the Green Bank Observatory in Green Bank, West Virginia; and the 305-m William E. Gordon Telescope (Arecibo) of Arecibo Observatory in Arecibo, Puerto Rico.

We observed most pulsars once a month. In addition, we started a high-cadence observing campaign in 2013, in which we made weekly observations of two pulsars with GBT (PSR J1713+0747 and PSR J1909−3744) and five pulsars with Arecibo (PSR J0030+0451, PSR J1640+2224, PSR J1713+0747, PSR J2043+1711, and PSR J2317+1439). This high-cadence observing campaign was specifically designed to increase the sensitivity of our PTA to GWs from individual sources (Burt et al. 2011; Christy et al. 2014). In most cases, we observed pulsars at every epoch with two receivers at different frequencies with wide enough bandwidths to measure the pulse dispersion due to the interstellar medium (ISM).

For each pulsar, the observed TOAs were fit to a timing model that described the pulsar's spin period and spin period derivative, sky location, proper motion, and parallax. The timing model also included terms describing pulse dispersion along the line of sight. For those pulsars in binaries, the timing model also included five Keplerian parameters that described the binary orbit and post-Keplerian parameters that described relativistic binary effects if they improved the timing fit. In the GW analyses, we used a linearized timing model centered around the best-fit parameter values.

3.3 Data Analysis Methods

As in NG5b we dispatch Bayesian inference to obtain posterior distributions for our signal parameters. Here we discuss the signal model, PTA likelihood, and methods used in our analysis.

3.3.1 PTA Likelihood

We anticipate the residuals to be combination of offsets $\delta\boldsymbol{\xi}$ in the linearized timing model M, white noise $\mathbf{n}_{\text{white}}$, red noise encompassing any time-correlated signal not of cosmological origin $\mathbf{n}_{\text{red}}$, and the gravitational wave signal $\mathbf{s}$. Then we can write the residuals for a pulsar as

$$\delta\mathbf{t} = M\delta\boldsymbol{\xi} + \mathbf{n}_{\text{white}} + \mathbf{n}_{\text{red}} + \mathbf{s}. \tag{3.3.1}$$

If we parameterize the residuals in terms of white noise parameters $\vec{\theta}$, red noise parameters $\vec{\varphi}$, and gravitational wave parameters $\vec{\lambda}$, Bayes theorem states

$$p(\delta\boldsymbol{\xi}, \vec{\theta}, \vec{\varphi}, \vec{\lambda}|\delta\mathbf{t}) = \frac{p(\delta\mathbf{t}|\delta\boldsymbol{\xi}, \vec{\theta}, \vec{\varphi}, \vec{\lambda})p(\delta\boldsymbol{\xi}, \vec{\theta}, \vec{\varphi}, \vec{\lambda})}{p(\delta\mathbf{t})}. \tag{3.3.2}$$

where

$p(\delta\boldsymbol{\xi}, \vec{\theta}, \vec{\varphi}, \vec{\lambda}|\delta\mathbf{t})$: "posterior"; i.e., probability of parameter values generating residuals $\delta\mathbf{t}$

$p(\delta\mathbf{t}|\delta\boldsymbol{\xi}, \vec{\theta}, \vec{\varphi}, \vec{\lambda})$: "likelihood"; i.e., probability that residuals $\delta\mathbf{t}$ are drawn from a random distribution of parameter values

$p(\delta\boldsymbol{\xi}, \vec{\theta}, \vec{\varphi}, \vec{\lambda})$: "prior"; knowledge of parameter values a priori analysis

$p(\delta\mathbf{t})$: "evidence"

Specifically, we seek the posterior probability distributions for parameters of interest, namely GW parameter $\vec{\lambda}$. As previous analyses for individual SMBHB sources in the PTA band conclude a detection at this stage of the experiment is unlikely, we seek upper limit values of $\vec{\lambda}$. In that case, the evidence serves as a normalizing factor we can safely ignore.

The likelihood for a single pulsar is a weighted multivariate Gaussian distribution, i.e.,

$$p(\delta\mathbf{t}|\delta\boldsymbol{\xi}, \vec{\theta}, \vec{\varphi}, \vec{\lambda}) = \frac{\exp\left[-\frac{1}{2}(\delta\mathbf{t} - \mathbf{s} - M\delta\boldsymbol{\xi})^T C^{-1}(\delta\mathbf{t} - \mathbf{s} - M\delta\boldsymbol{\xi})\right]}{\sqrt{(2\pi)^{N_{\text{TOA}}}\det C}} \tag{3.3.3}$$

with correlations between the residuals encapsulated in C. It is possible to separate the white and red components of noise so that

$$C = C_{\text{white}} + C_{\text{red}}. \tag{3.3.4}$$

Depending on the complexity of parameter covariance, the likelihood can be expensive to compute numerically. It is crucial that our models are concise enough to minimize the dimensionality of the problem, but not so concise as to compromise recovered information in the posteriors.

3.3.2 Signal and noise models

White noise

We used the same white noise model as NG5b, which has three parameters: EFAC, EQUAD, and ECORR. The EFAC parameter scales the TOA uncertainties, and the EQUAD parameter adds white noise in quadrature. The ECORR parameter describes additional white noise added in quadrature that is correlated within the same observing epoch across radio frequencies, such as pulse jitter (Dolch et al. 2014; Lam et al. 2017a). We used the improved implementation of ECORR described in NG11b.

Red noise

We use *red noise* to describe any time-correlated noise process with an higher ensemble-average power at lower frequencies. There are many possible sources of red noise in pulsar timing residuals including spin noise, variations in pulse shape, pulsar mode changes, and errors in modeling pulse dispersion from the ISM (Cordes 2013; Lam et al. 2017a; Jones et al. 2017). For our core individual SMBHB search, we decompose the red noise in each pulsar through a Fourier analysis as follows

$$\mathbf{n}_{\text{red}} = \sum_{j=1}^{N_{\text{mode}}} \left[a_j \sin\left(\frac{2\pi j t}{T}\right) + b_j \cos\left(\frac{2\pi j t}{T}\right) \right] = F\mathbf{a} \tag{3.3.5}$$

where we divide the noise spectrum into 30 bins spaced linearly between $f = 1/T_{\text{obs}}$ and $f = 30/T_{\text{obs}}$, where T_{obs} is the total observation time for that particular pulsar.

If we assume the pulsar's red noise does not vary measurably across observation time, then the Fourier coefficients $\mathbf{a}$ may be treated as spanning an orthogonal basis of sines and cosines. These enter the likelihood through the red noise piece of the covariance matrix

$$C_{\rm red} = \langle F\mathbf{a}(F\mathbf{a})^T \rangle = F \langle \mathbf{a}\mathbf{a}^T \rangle F^T \tag{3.3.6}$$

where $\langle \mathbf{a}\mathbf{a}^T \rangle$ is a diagonal matrix containing the red noise power at every Fourier frequency bin, i.e.,

$$\langle \mathbf{a}\mathbf{a}^T \rangle_{ij} = \varphi_i \delta_{ij}. \tag{3.3.7}$$

Our standard red noise model is a power-law that can be written as

$$P(f) = A_{\rm red}^2 \left(\frac{f}{f_{\rm yr}} \right)^{-\gamma} \tag{3.3.8}$$

where $f_{\rm yr} \equiv 1/(1 \text{ yr})$, $A_{\rm red}$ is the amplitude, and γ is the spectral index.

Then

$$\varphi_i^2 = \frac{1}{T_{\rm obs}} \frac{1}{12\pi f_{\rm yr}^3} P(f). \tag{3.3.9}$$

Ultimately, we marginalize over the individual Fourier coefficients and focus on recovering posterior distributions for parameters $A_{\rm red}$ and γ.

Gravitational wave signal

Finally, we include the gravitational wave signal from an individual SMBHB. Consider a GW source whose location in equatorial coordinates is given by declination δ and right ascension α. It is convenient to write the sky position in terms of the polar angle θ and azimuthal angle ϕ, which are related to δ and α by $\theta = \pi/2 - \delta$ and $\phi = \alpha$. The emitted GWs can be written in terms of two polarizations:

$$h_{ab}(t, \hat{\Omega}) = e_{ab}^{+}(\hat{\Omega}) \; h_{+}(t, \hat{\Omega}) + e_{ab}^{\times}(\hat{\Omega}) \; h_{\times}(t, \hat{\Omega}) \,, \tag{3.3.10}$$

where $\hat{\Omega}$ is a unit vector from the GW source to the Solar System barycenter (SSB), $h_{+,\times}$ are the polarization amplitudes, and $e_{ab}^{+,\times}$ are the polarization tensors. The polarization tensors can be written in the SSB frame as (Wahlquist 1987)

$$
\begin{aligned}
e^{+}_{ab}(\hat{\Omega}) &= \hat{m}_a\, \hat{m}_b - \hat{n}_a\, \hat{n}_b\,, && (3.3.11)\\
e^{\times}_{ab}(\hat{\Omega}) &= \hat{m}_a\, \hat{n}_b + \hat{n}_a\, \hat{m}_b\,, && (3.3.12)
\end{aligned}
$$

where

$$
\begin{aligned}
\hat{\Omega} &= -\sin\theta\cos\phi\,\hat{x} - \sin\theta\sin\phi\,\hat{y} - \cos\theta\,\hat{z}\,, && (3.3.13)\\
\hat{m} &= -\sin\phi\,\hat{x} + \cos\phi\,\hat{y}\,, && (3.3.14)\\
\hat{n} &= -\cos\theta\cos\phi\,\hat{x} - \cos\theta\sin\phi\,\hat{y} + \sin\theta\,\hat{z}\,. && (3.3.15)
\end{aligned}
$$

The response of a pulsar to the source is described by the antenna pattern functions F^{+} and $F^{\times}$ (Sesana & Vecchio 2010; Ellis et al. 2012; Taylor et al. 2016),

$$
\begin{aligned}
F^{+}(\hat{\Omega}) &= \frac{1}{2}\frac{(\hat{m}\cdot\hat{p})^2 - (\hat{n}\cdot\hat{p})^2}{1+\hat{\Omega}\cdot\hat{p}}\,, && (3.3.16)\\
F^{\times}(\hat{\Omega}) &= \frac{(\hat{m}\cdot\hat{p})(\hat{n}\cdot\hat{p})}{1+\hat{\Omega}\cdot\hat{p}}\,, && (3.3.17)
\end{aligned}
$$

where $\hat{p}$ is a unit vector pointing from Earth to the pulsar.

The effect of a GW on a pulsar's residuals can be written as

$$
s(t,\hat{\Omega}) = F^{+}(\hat{\Omega})\,\Delta s_{+}(t) + F^{\times}(\hat{\Omega})\,\Delta s_{\times}(t)\,, \qquad (3.3.18)
$$

where $\Delta s_{+,\times}$ is the difference between the signal induced at Earth and at the pulsar (the so-called "Earth term" and "pulsar term"),

$$
\Delta s_{+,\times}(t) = s_{+,\times}(t_p) - s_{+,\times}(t)\,, \qquad (3.3.19)
$$

where t is the time at which the GW passes the SSB and t_p is the time at which it passes the pulsar. From geometry, we can relate t and t_p by

$$
t_p = t - L(1+\hat{\Omega}\cdot\hat{p})\,, \qquad (3.3.20)
$$

where L is the distance to the pulsar.

For a circular binary, at zeroth post-Newtonian (0-PN) order, $s_{+,\times}$ is given by

(Wahlquist 1987; Lee et al. 2011; Corbin & Cornish 2010)

$$\begin{aligned} s_+(t) &= \frac{\mathcal{M}^{5/3}}{d_L\,\omega(t)^{1/3}}\left[-\sin 2\Phi(t)\,\left(1+\cos^2 i\right)\,\cos 2\psi\right. \\ &\quad \left.-2\cos 2\Phi(t)\,\cos i\,\sin 2\psi\right]\,, \qquad (3.3.21) \\ s_\times(t) &= \frac{\mathcal{M}^{5/3}}{d_L\,\omega(t)^{1/3}}\left[-\sin 2\Phi(t)\,\left(1+\cos^2 i\right)\,\sin 2\psi\right. \\ &\quad \left.+2\cos 2\Phi(t)\,\cos i\,\cos 2\psi\right]\,, \qquad (3.3.22) \end{aligned}$$

where i is the inclination angle of the SMBHB, ψ is the GW polarization angle, d_L is the luminosity distance to the source, and $\mathcal{M} \equiv (m_1 m_2)^{3/5}/(m_1+m_2)^{1/5}$ is a combination of the black hole masses m_1 and m_2 called the "chirp mass." Note that the variable $\mathcal{M}$ is the observed redshifted value, which is related to the rest-frame value $\mathcal{M}_r$ according to

$$\mathcal{M}_r = \frac{\mathcal{M}}{1+z} \qquad (3.3.23)$$

Currently PTAs are only sensitive to sources in the local Universe for which $(1+z) \approx 1$.

For a circular binary, the orbital angular frequency is related to the GW frequency by $\omega_0 = \pi f_{\rm gw}$, where $\omega_0 = \omega(t_0)$. For our search, we defined the reference time t_0 as 31 December 2015 (MJD 57387), which corresponded to the last day data were taken for the 11-year data set. The orbital phase and frequency of the SMBHB are given by (NG5b)

$$\begin{aligned} \Phi(t) &= \Phi_0 + \frac{1}{32}\mathcal{M}^{-5/3}\left[\omega_0^{-5/3} - \omega(t)^{-5/3}\right]\,, \qquad (3.3.24) \\ \omega(t) &= \omega_0\left(1 - \frac{256}{5}\mathcal{M}^{5/3}\omega_0^{8/3}t\right)^{-3/8}\,, \qquad (3.3.25) \end{aligned}$$

where Φ_0 and ω_0 are the initial orbital phase and frequency, respectively. Unlike in NG5b, we used the full expression for $\omega(t)$ in our signal model rather than treating the GW frequency at Earth as a constant, as high-chirp-mass binaries will evolve significantly over the timescale of our observations.

3.3.3 Bayesian methods and software

The Bayesian inference procedure used followed closely that of NG5b, with the addition of the BayesEphem model for the uncertainty in the SSB introduced in NG11b. Pulsar timing uses a Solar System ephemeris (SSE) to transform from individual observatories' reference frames to an inertial reference frame centered at the SSB. We used DE436

(Folkner & Park 2016) to perform this transformation, plus the BayesEphem model, which parameterizes uncertainty in the outer planets' masses, Jupiter's orbit, and the rotation rate about the ecliptic pole.

We implemented the likelihood and priors and performed the searches using NANOGrav's new software package `enterprise`[1]. We confirmed the accuracy of this package by also performing some searches using the software package `PAL2`[2], which has been used for previous NANOGrav GW searches. Both packages used the Markov Chain Monte Carlo (MCMC) sampler `PTMCMCSampler`[3] to explore the parameter space.

For detection and upper-limit runs, we described the Earth-term contribution to the GW signal by eight parameters:

$$\lambda_0 = \{\theta, \phi, \Phi_0, \psi, i, \mathcal{M}, f_{\rm gw}, h_0\}\ . \tag{3.3.26}$$

We have reparameterized Eqs. (3.3.21) – (3.3.22) in terms of the characteristic strain h_0 rather than d_L, where

$$h_0 = \frac{2\mathcal{M}^{5/3}(\pi f_{\rm gw})^{2/3}}{d_L}\,. \tag{3.3.27}$$

We used log-uniform priors on h_0 for detection analyses, and a uniform prior on h_0 to compute upper limits on the strain. For both types of analyses, we searched over $\log_{10} h_0 \in [-18, -11]$.

We used isotropic priors on the sky position of the source (θ, ϕ), source inclination angle i, GW polarization angle ψ, and GW phase Φ_0. We searched over $\log_{10}\mathcal{M}$ with a uniform prior $\log_{10}(\mathcal{M}/{\rm M}_\odot) \in [7, 10]$. For high $f_{\rm gw}$, we truncated the prior on $\log_{10}\mathcal{M}$ to account for the fact that high-chirp-mass systems will have merged before emitting high-frequency GWs. Assuming binaries merge when the orbital frequency is equal to the innermost stable circular orbit (ISCO) frequency, $\mathcal{M}$ must satisfy

$$\mathcal{M} \leq \frac{1}{6^{3/2}\pi f_{\rm gw}}\left[\frac{q}{(1+q)^2}\right]^{3/5}\,, \tag{3.3.28}$$

where q is the mass ratio. For our analyses, we used the chirp-mass cutoff with $q = 1$. This change to the prior on $\mathcal{M}$ only affected $f_{\rm gw} \geq 191.3\,$nHz.

---

[1] https://github.com/nanograv/enterprise

[2] https://github.com/jellis18/PAL2

[3] https://github.com/jellis18/PTMCMCSampler

We performed searches at fixed values of f_{gw}. The minimum GW frequency was set by the total observation time, $f_{\mathrm{gw}} = 1/(11.4 \text{ yrs}) = 2.8\,\mathrm{nHz}$. The maximum GW frequency was set by the observing cadence. Because of the high-cadence observing campaign, the 11-year data set can detect GWs with frequencies up to 826.7 nHz; however, the data are not very sensitive at high frequencies. Also, we do not expect to find any SMBHBs with orbital periods of weeks because high-chirp-mass systems would have already merged before emitting at those frequencies, and low-chirp-mass systems would be evolving through the PTA band very quickly at that point. Therefore, we only searched for GWs with frequencies up to 317.8 nHz, which corresponded to the high-frequency-cut-off adopted in NG5b.

The pulsar-term contributions to the GW signal used the pulsar distances to compute the light-travel-time between when the GW passed the pulsars and when it passed the SSB (see Eq. (3.3.20)). We used a Gaussian prior on the distances with the measured mean and uncertainty from Verbiest et al. (2012); for the pulsars not included in that paper, we used a mean of 1 kpc and error of 20%. The phase at the pulsar can be written as

$$\Phi(t) = \Phi_0 + \Phi_p + \frac{1}{32}\mathcal{M}^{-5/3}\left[\omega(t_{p,0})^{-5/3} - \omega(t_p)^{-5/3}\right] , \tag{3.3.29}$$

where Φ_p is the phase difference between Earth and the pulsar. The pulsar phase parameters Φ_p can be computed from the pulsar distances and chirp mass as

$$\Phi_p = \frac{1}{32}\mathcal{M}^{-5/3}\left[\omega_0^{-5/3} - \omega(t_{p,0})^{-5/3}\right] ; \tag{3.3.30}$$

however, in most cases the pulsar distance uncertainties ($\Delta L \sim 10 - 100$ pc) are significantly greater than the GW wavelengths ($\lambda_{\mathrm{gw}} \sim 0.1 - 10$ pc), and so the phase differences between the Earth terms and pulsar terms are effectively random. Therefore, following the approach of Corbin & Cornish (2010), we treated Φ_p as an independent parameter with a uniform prior $\Phi_p \in [0, 2\pi]$.

We fixed the white noise parameters to their best-fit values, as determined from noise analyses performed on individual pulsars. In the GW analyses, we simultaneously searched over the individual pulsars' red noise using the power-law model with uniform priors on $\log_{10} A_{\mathrm{red}} \in [-20, -11]$ and $\gamma \in [0, 7]$. In order to burn-in the red noise and

BayesEphem parameters efficiently, we introduced jump proposals that drew proposed samples from empirical distributions based on the posteriors from an initial Bayesian analysis with only the pulsars' red noise and BayesEphem (i.e., excluding a GW signal).

We computed Bayes factors for the presence of a GW signal using the Savage-Dickey formula (Dickey 1971),

$$\mathcal{B}_{10} \equiv \frac{\text{evidence}[\mathcal{H}_1]}{\text{evidence}[\mathcal{H}_0]} = \frac{p(h_0 = 0|\mathcal{H}_1)}{p(h_0 = 0|\mathcal{D}, \mathcal{H}_1)}\,, \tag{3.3.31}$$

where $\mathcal{H}_1$ is the model with a GW signal plus individual pulsar red noise, $\mathcal{H}_0$ is the model with only individual pulsar red noise, and $\mathcal{D}$ the data. $p(h_0 = 0|\mathcal{H}_1)$ is the prior volume at $h_0 = 0$, and $p(h_0 = 0|\mathcal{D}, \mathcal{H}_1)$ is the posterior volume at $h_0 = 0$. We were able to use the Savage-Dickey formula because $\mathcal{H}_1$ and $\mathcal{H}_0$ are nested models, i.e., $\mathcal{H}_0$ is $\mathcal{H}_1 : h_0 = 0$. We approximated $p(h_0 = 0|\mathcal{D}, \mathcal{H}_1)$ as the fraction of samples in the lowest-amplitude bin of a histogram of h_0. We computed the Bayes factor for a range of bin sizes, and reported the mean as $\mathcal{B}_{10}$ and the standard deviation as its error.

For upper limits, following the approach of NG11b, we computed the standard error as

$$\sigma = \frac{\sqrt{x(1-x)/N_s}}{p(h_0 = h_0^{95\%}|\mathcal{D})}\,, \tag{3.3.32}$$

where $x = 0.95$ and N_s is the number of effective samples in the chain. This definition of σ is the error in the computed 95% upper limit due to using a finite number of samples. We estimated the number of effective samples by dividing the total number of samples by the autocorrelation chain length, which is a measure of how far apart two samples in the chain must be in order to be statistically independent.

3.4 Results

In this section, we report the results of both detection and upper limit analyses of the NANOGrav 11-year data set for GWs from individual circular SMBHBs. We used the data to place upper limits as a function of frequency and sky location and to compare upper limits from the 11-year data set to those from the 5- and 9-year data sets. We also briefly discuss a new Bayesian technique to determine how much each pulsar in a

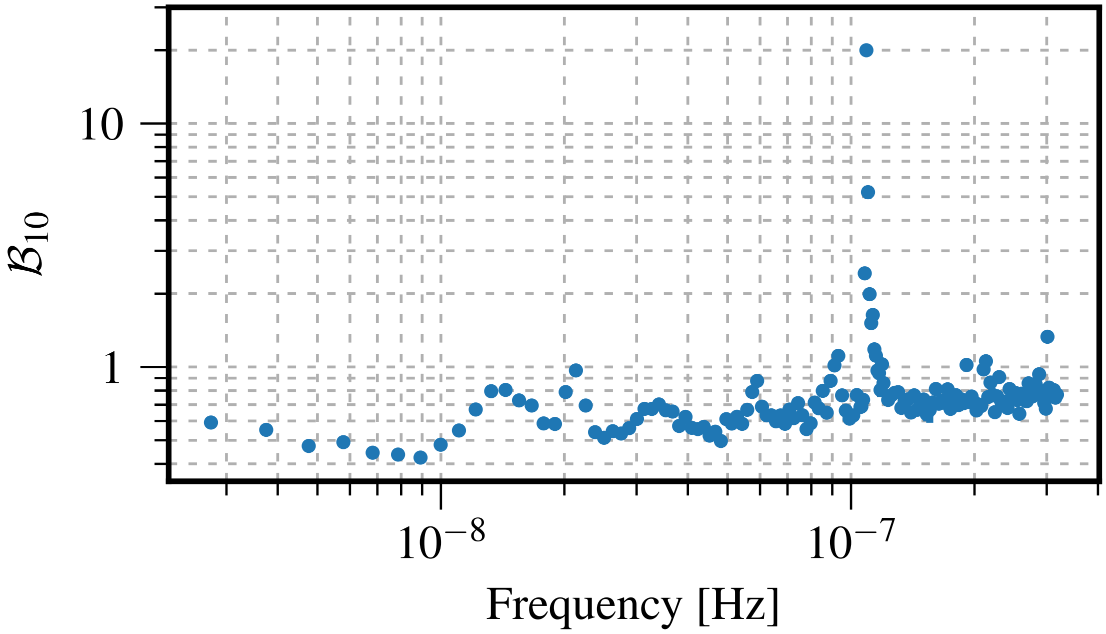

FIGURE 12: Bayes factors for a GW signal from an individual circular SMBHB as a function of GW frequency in the NANOGrav 11-year data set. We found no strong evidence for GWs in our data. The highest Bayes factor was at $f_{\rm gw} = 109\,{\rm nHz}$, for which $\mathcal{B}_{10} = 20.0(9)$. For all other frequencies searched, the Bayes factors were close to 1.

PTA contributes to a common signal in order to diagnose spurious signals. Following the approach of NG11b, our analyses of the 11-year data set only used the 34 pulsars which had been observed for at least three years. Our analyses of the 5- and 9-year data sets used the same subset of pulsars that were used in the corresponding analyses for the GWB (NG5a, Arzoumanian et al. 2016), which included 17 and 18 pulsars, respectively.

3.4.1 Detection analyses

We performed detection searches for GWs from individual circular SMBHBs on the 11-year data set. Fig. 12 shows the Bayes factors for each frequency, marginalized over the sky location. We did not find strong evidence for GWs in the 11-year data set. The largest Bayes factor was at $f_{\rm gw} = 109\,{\rm nHz}$, for which $\mathcal{B}_{10} = 20.0(9)$. For all other frequencies, the Bayes factors were between $\mathcal{B}_{10} = 0.43(1)$ and $\mathcal{B}_{10} = 1.31(4)$, indicating no evidence of GWs in the data.

Although the detection search at $f_{\mathrm{gw}} = 109\,\mathrm{nHz}$ found a higher Bayes factor than any of the other values of f_{gw}, we emphasize that the Bayes factor is not high enough to claim a detection. A Bayes factor of 20 means 20:1 odds for the presence of a GW signal; similarly, frequentist analysis yields a signal-to-noise ratio (SNR) of 1.2. Neither of these metrics supports the claim that the data shows evidence of GWs. Furthermore, as we discuss in more detail in Sec. 3.4.3, we determined that most of the evidence for this signal was in the residuals of a single pulsar, J1713+0747, whereas a true GW signal should be seen in many pulsars.

3.4.2 Upper limit analyses

As we did not find strong evidence for GWs from individual circular SMBHBs in the 11-year data set, we placed upper limits on the GW strain. Fig. 13 shows the sky-averaged 95% upper limit on the GW strain amplitude. At the most sensitive frequency of 8 nHz, we placed a 95% upper limit on the strain of approximately $h_0 < 7.3(3) \times 10^{-15}$. We also show the strain upper limits from the 5- and 9-year data sets for comparison. There was an improvement of about a factor of two between the 5-year and 9-year data sets, and more than a factor of two between the 9-year and 11-year data sets. Our upper limit based on the 11-year data set was about 1.4 times lower than that of $h_0 < 10^{-14}$ set by the EPTA based on observations of 6 pulsars observed for up to 17.7 years (Babak et al. 2016; Desvignes et al. 2016).

We note that there is an increase in the strain upper limit from the 9-year data set creating a “bump” in the spectrum centered around $f_{\mathrm{gw}} = 15\,\mathrm{nHz}$. At first glance, this loss of sensitivity indicates the presence of some noise that could be a GW signal. We examined the posterior distributions at this frequency and saw our posterior distributions found maximum likelihood for a source located nearest PSR J0613−0200 (see Fig. 14).

However, there is not a significant Bayes factor at this frequency ($\mathcal{B}_{10} = 1.42(3)$). Furthermore, this “bump” in the spectrum is not present in the 11-year data set. A true GW signal would only increase in significance with the passage of time as a longer baseline would further diminish down any other noise at this frequency. As discussed in more detail in Sec. 3.4.4, this increase in the strain upper limit is due to an unmodeled

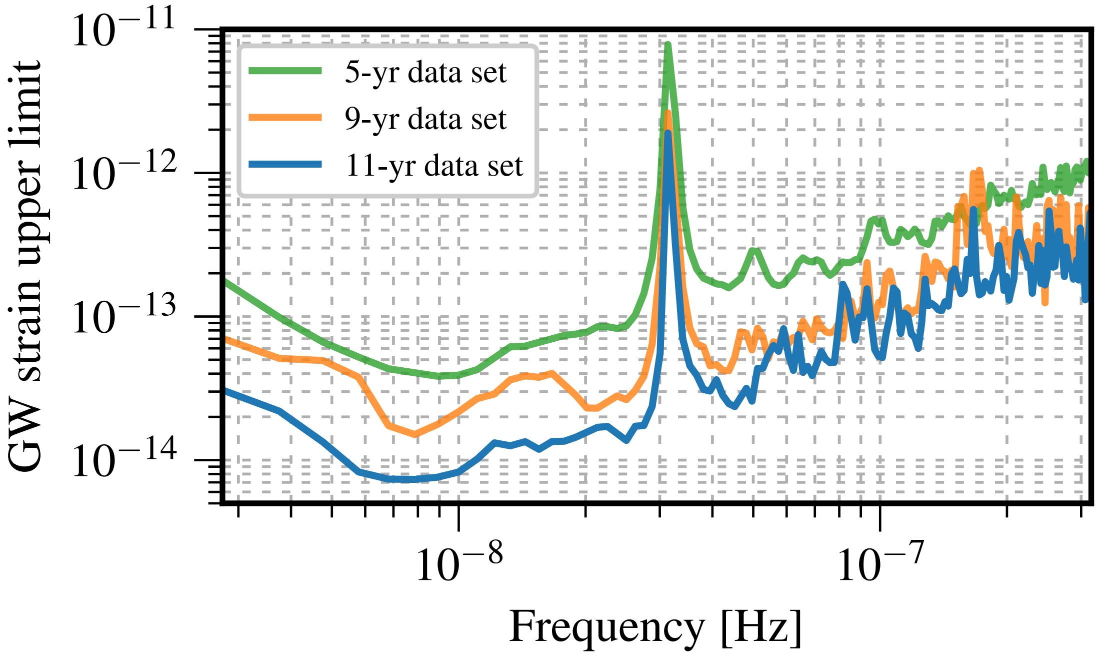

FIGURE 13: Sky-averaged 95% upper limit on the GW strain amplitude from a circular SMBHB as a function of GW frequency from the NANOGrav 5-year data set (teal), 9-year data set (purple), and 11-year data set (orange). These analyses used BAYESEPHEM to parameterize uncertainty in the SSB. The data were most sensitive at $f_{\mathrm{gw}} = 8\,\mathrm{nHz}$, with a strain upper limit of approximately $h_0 < 1.51(7) \times 10^{-14}$ from the 9-year data set, and $h_0 < 7.3(3) \times 10^{-15}$ from the 11-year data set.

signal in a single pulsar, PSR J0613−0200.

Our sensitivity to individual sources varied significantly with the angular position of the source due to having a finite number of pulsars distributed unevenly across the sky. Fig. 15 shows the 95% upper limit on the GW strain for $f_{\mathrm{gw}} = 8\,\mathrm{nHz}$ as a function of sky position, plotted in equatorial coordinates. The upper limit varies from $h_0 < 2.0(1) \times 10^{-15}$ at the most sensitive sky location to $h_0 < 1.34(4) \times 10^{-14}$ at the least sensitive sky location.

3.4.3 “Dropout” analyses

In our searches of the NANOGrav 9-yr and 11-yr data sets, we found two low-SNR signals. We introduced a new type of analysis that used “dropout” parameters to determine how much each individual pulsar contributed to these signals. The dropout method tests the robustness of the correlations in the signal by determining whether evidence for the

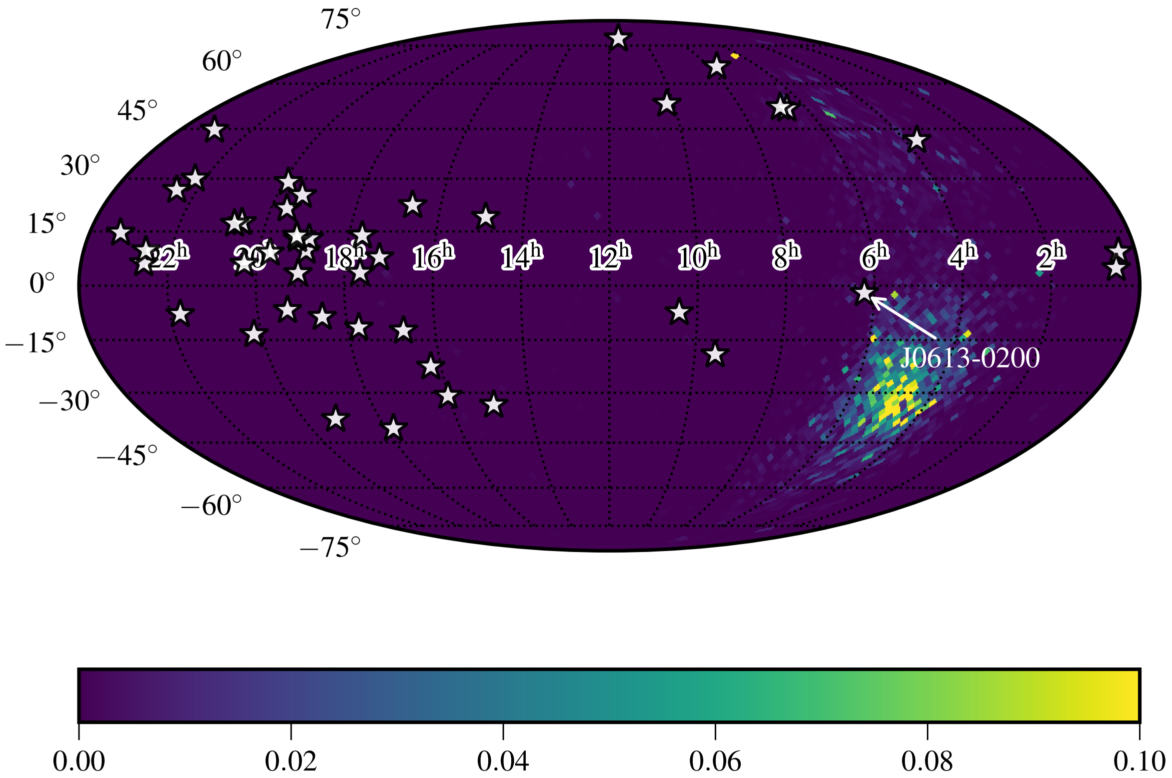

FIGURE 14: Normalized sky position posterior distribution at $f_{\mathrm{gw}} = 15\,\mathrm{nHz}$. Color bar corresponds the percentage of MCMC samples in a particular sky position. Maximum likelihoods were found in a region nearest PSR J0613−0200 indicating a potential signal originating in that direction.

signal comes from correlations between multiple pulsars, or it only originates from a single pulsar. It is similar to the dropout technique in neural networks, where units are randomly dropped during training in order to strengthen the network (Srivastava et al. 2014). This method is also similar to jackknife resampling (Efron & Stein 1981); however, in jackknifing, samples are removed in order to estimate the bias in parameter estimation, whereas in dropout analyses the parameter values are held fixed, and the dropout parameters indicate how much each pulsar is biasing the parameter estimation. An upcoming paper will further describe and develop this method (Vigeland et al. 2019)

In a dropout analysis, the GW parameters were held fixed at the values that maximized the likelihood, and dropout parameters k_a were introduced into the signal model. If k_a was above a threshold, then the ath pulsar's residuals included the contribution from the GW; otherwise, that pulsar's residuals did not. Thus at each iteration of the MCMC, the GW was present in only a subset of the pulsars' residuals. The posteriors of

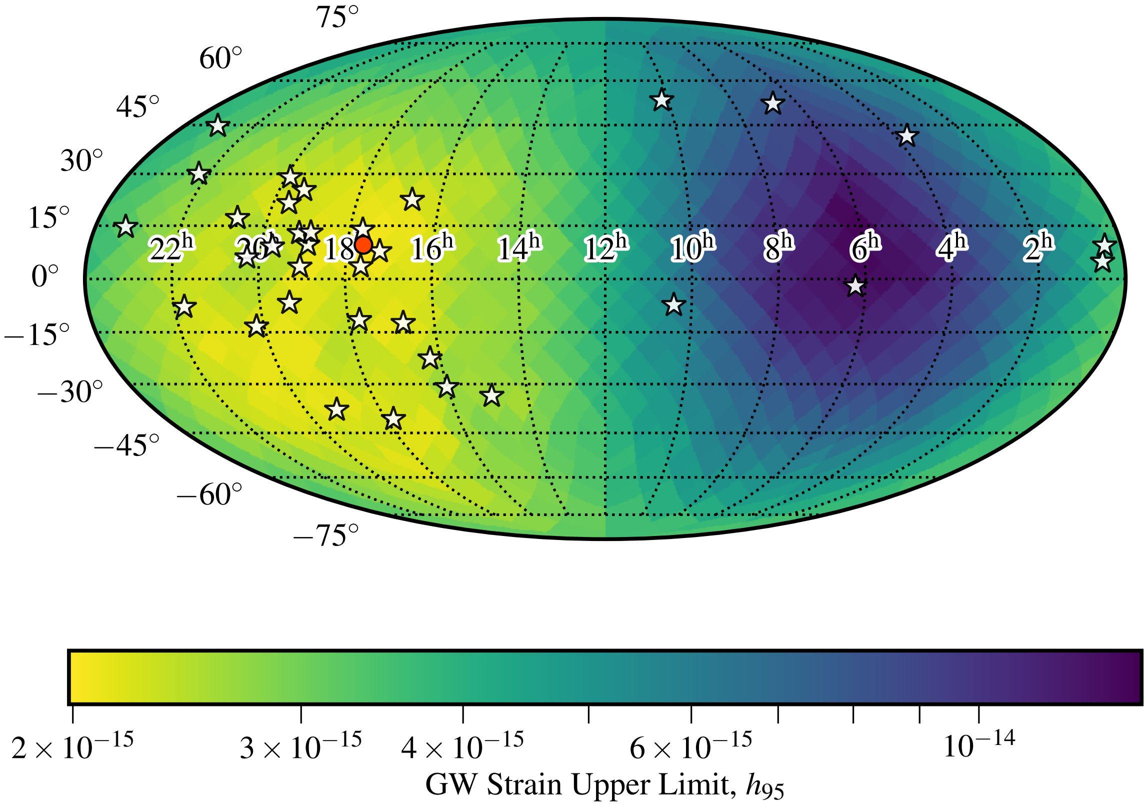

FIGURE 15: The 95% upper limit on the GW strain amplitude from a circular SMBHB with $f_{\rm gw} = 8\,$nHz as a function of sky position from an analysis of the 11-year data set plotted in equatorial coordinates using the Mollweide projection. We used the DE436 ephemeris model with BAYESEPHEM to model uncertainty in the SSB. The positions of pulsars in our array are indicated by stars, and the most sensitive sky location is indicated by a red circle. The 95% upper limit ranged from $2.0(1) \times 10^{-15}$ at our most sensitive sky location to $1.34(4) \times 10^{-14}$ at our least sensitive sky location.

the dropout parameters indicated whether the data preferred for the GW to be present in each pulsar.

We performed two dropout analyses. The first was on the 9-yr data set at $f_{\rm gw} = 15\,$nHz. The analysis of the 9-year data set found an increase in the 95% strain upper limit at $f_{\rm gw} = 15\,$nHz compared to the upper limits at neighboring frequencies. Furthermore, as shown in Fig. 16, we found that the strain upper limit decreased significantly when PSR J0613−0200 was removed from the 9-year data set. However, there was very little difference in the Bayes factor: $\mathcal{B}_{10} = 1.42(3)$ with all pulsars, and $\mathcal{B}_{10} = 1.25(3)$ excluding PSR J0613−0200. Fig. 17 shows the results of a dropout analysis. We fixed the GW signal parameters to the best-fit values from a detection analysis including all pulsars and only

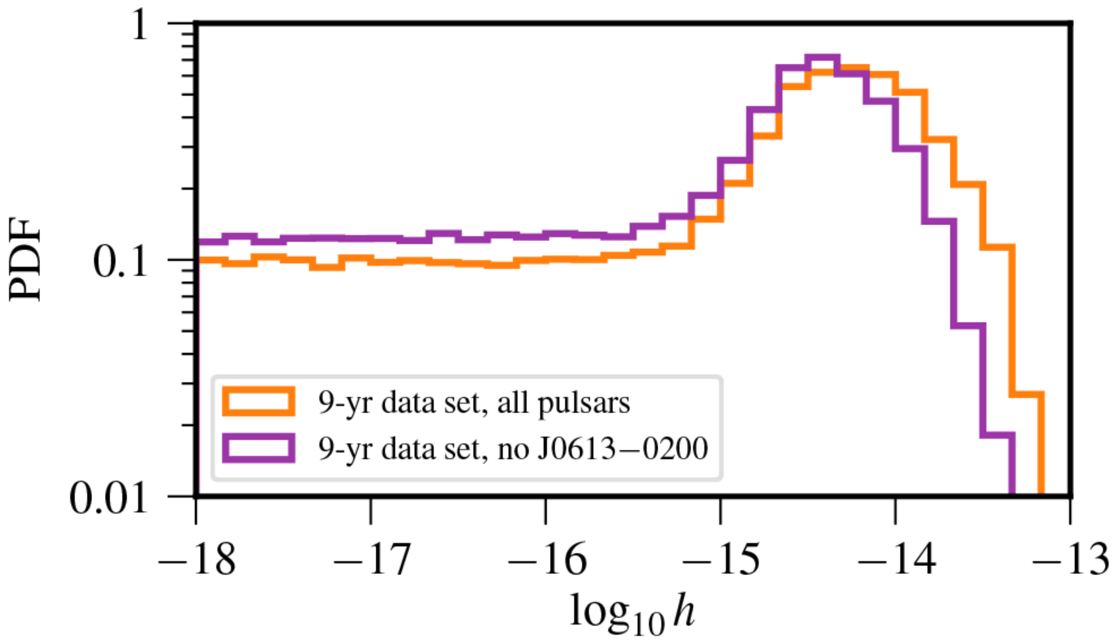

FIGURE 16: Comparison between a search at $f_{\rm gw} = 15\,$nHz performed on the 9-yr data set with all pulsars (orange) and excluding PSR J0613−0200 (purple). There was very little difference between the Bayes factors ($\mathcal{B}_{10} = 1.42(3)$ with all pulsars, and $\mathcal{B}_{10} = 1.25(3)$ excluding PSR J0613−0200), but there was a significant difference in the 95% strain upper limit. We found an upper limit of $4.1(2) \times 10^{-14}$ with all pulsars, compared with $3.2(3) \times 10^{-14}$ without PSR J0613−0200.

allowed the dropout parameters to vary. We set $k_{\rm threshold} = 1/2$, so that there was an equal prior probability of the signal being included or excluded in the model for each pulsar's residuals. PSR J0613−0200 had the largest Bayes factor while all other pulsars had Bayes factors of order 1, from which we concluded that the increase in the strain upper limit at $f_{\rm gw} = 15\,$nHz was caused by an unmodeled non-GW signal in PSR J0613−0200.

3.4.4 Spurious spectral features in our data:PSR J0613−0200

Luckily, the dropout analysis provides a means to prevent erroneous detections. It is more efficient, however, to anticipate such problems so that rote upper limits analyses can be made robust. This has spurred efforts to deploy more bespoke noise models for individual pulsars in our analyses.

Typically, the power-law is sufficient to capture the pulsar's these features, but if certain frequencies have more power than anticipated, that can be mistakenly absorbed

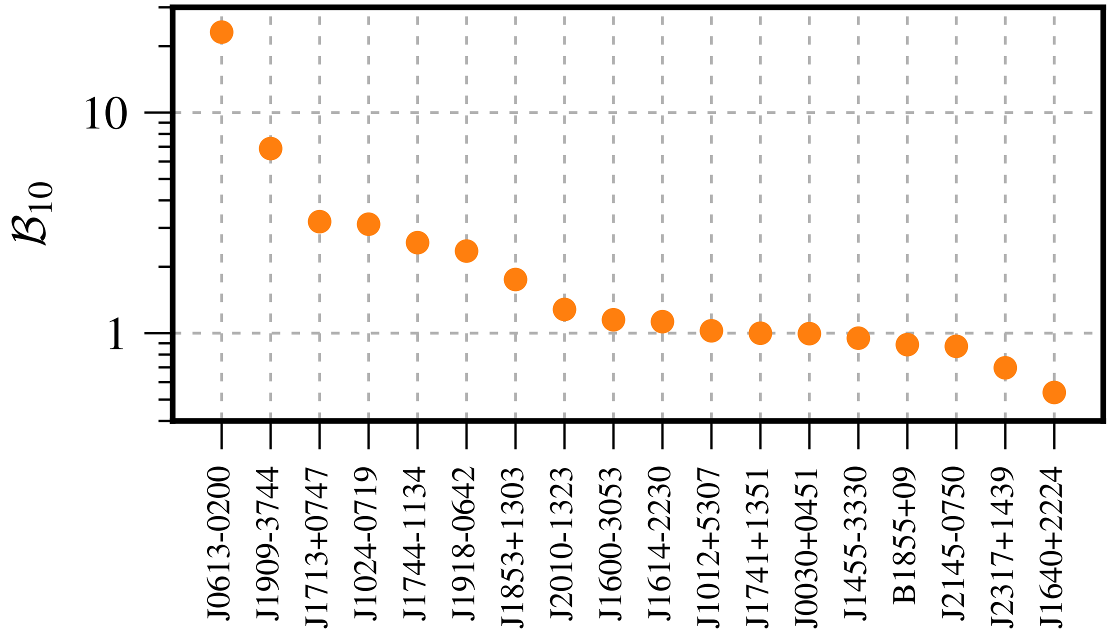

FIGURE 17: Bayes factors for the presence of a GW signal in each pulsar's residuals, from an analysis of the 9-yr data set with $f_{\rm gw} = 15\,{\rm nHz}$. The GW parameters are fixed to the maximum-likelihood values, as determined from a detection analysis. PSR J0613−0200 had the largest Bayes factor for the signal, with $\mathcal{B}_{10} = 23.2(5)$, indicating that PSR J0613−0200 was the primary source of this signal.

by our GW model. One way to combat this is through the use of a free-spectral noise model, which each frequnecy bin is treated independently treats the red noise power spectral density as a series of power-laws broken with varying amplitude and spectral index in every Fourier frequency bin. This technique is the so-called "sledgehammer" method because it catches most of a pulsar's intrinsic red noise, however, it adds $N_{\rm modes} \times 2$ parameters to the model, due to the Occam penalty associated with the extra parameters. which reduces the efficiency of our search. Another possibility is a modified power-law called the t-process. This red noise power spectral density is written as

$$P_i = \alpha_i A_{\rm red}^2 \left(\frac{f_i}{f_{\rm yr}}\right)^{-\gamma} \tag{3.4.1}$$

where P_i is the power in a Fourier frequency bin centered around a frequency f_i, and α_i is a deviation parameter. The t-process allows for deviations from the fiducial power-law, with a student's t-distribution as a prior on α_i. The t-process also increases the dimensionality of the searched parameter space, like the free spectral model, but this is partially offset by the effectively reduced prior volume you have to sample.

Fig. 18 shows the results of running the canonical red noise models on the residuals of PSR J0613−0200. The top panel shows the results from the 9-year data set, and the bottom panel shows the results from the 11-year data set. For the power-law model, we plot the spectrum using the maximum-likelihood values for $A_{\rm red}$ and γ. For the free spectral model, we plot either the median and 68% confidence interval or the 95% upper limit, depending on whether the posteriors of the bin amplitudes are consistent with zero. For the t-process model, we plot the spectrum using the maximum-likelihood values for $A_{\rm red}$ and γ, and the median values for the deviation parameters α. We identified an excess of noise at a frequency of about 15 nHz in both the 9-year and 11-year data sets using the free spectral model and t-process model. This kind of sharp spectral feature is not adequately captured by the power-law, allowing it to be mistakenly identified as a GW signal by our search.

3.5 Limits on Astrophysical Properties of Nearby SMBHBs

In this section, we discuss what we can infer about the astrophysical properties of nearby SMBHBs from our limits on the GW strain. We used the 95% upper limits on the GW strain to place 95% *lower* limits on the distance to SMBHBs using Eq. (3.3.27) for a given chirp mass. Fig. 19 shows the 95% lower limit on the distances to individual SMBHBs as a function of sky position, plotted in equatorial coordinates, for sources with $\mathcal{M} = 10^9\,M_\odot$ and $f_{\rm gw} = 8\,\mathrm{nHz}$. The limits on the luminosity distance varied by a factor of 7 between the most-sensitive and least-sensitive sky locations. At the most-sensitive sky location, we found $d_L > 120\,\mathrm{Mpc}$ for SMBHBs with $\mathcal{M} = 10^9\,\mathrm{M}_\odot$ and $d_L > 5.5\,\mathrm{Gpc}$ for SMBHBs with $\mathcal{M} = 10^{10}\,\mathrm{M}_\odot$.

Fig. 20 shows the limits on the chirp masses of any SMBHBs in the nearby Virgo Cluster, which is at a distance of 16.5 Mpc. We found that there are no SMBHBs in the Virgo Cluster with $\mathcal{M} > 1.6(1) \times 10^9\,M_\odot$ emitting GWs in the PTA band. Furthermore, there are no SMBHBs with $\mathcal{M} > 3.8(1) \times 10^8\,M_\odot$ emitting GWs with $f_{\rm gw} = 9\,\mathrm{nHz}$. These chirp-mass limits imply that none of the galaxies NGC 4472 (estimated black hole mass of $2.5 \times 10^9\,M_\odot$; Rusli et al. 2013), NGC 4486 (estimated black hole mass of $6.6 \times 10^9\,M_\odot$;

Gebhardt et al. 2011), or NGC 4649 (estimated black hole mass of $4.5 \times 10^9\, M_\odot$; Shen & Gebhardt 2010) could contain binaries emitting GWs in this frequency range.

In order to assess how likely we were to have detected an SMBHB given our current sensitivity, we compared our strain upper limit curves to simulations of nearby SMBHBs. A similar technique was introduced in Babak et al. (2016) to estimate the detection probability from the strain upper limit curve. We generated populations of SMBHBs using the technique in Mingarelli et al. (2017), which is based on galaxies in the Two Micron All-Sky Survey (2MASS; Skrutskie et al. 2006) and merger rates from the Illustris cosmological simulation project (Genel et al. 2014; Rodriguez-Gomez et al. 2015). We estimated the number of detectable sources as the number lying above our sky-averaged 95% strain upper limit curve. Fig. 21 shows the loudest GW sources for a sample realization plotted alongside our 95% strain upper limit curve. We show both the sky-averaged strain upper limit curve (solid, blue line) and the strain upper limit curve at the most-sensitive sky location (dashed, red line). For this particular simulation, none of the sources were above the sky-averaged strain upper limit curve; therefore, we concluded there were no detectable sources in this particular realization. Out of 75,000 realizations of the local Universe, 34 contained a source that lay above the sky-averaged strain upper limit curve (i.e., 0.045% of realizations contained an observable SMBHB), from which we concluded that our non-detection was unsurprising given our current sensitivity. We point out, though, that our sensitivity varies significantly with sky location, and therefore some sources that are below the sky-averaged strain upper limit curve may be detectable depending on their sky locations. In our simulations, we found that a GW source lay above the strain upper limit curve at the most-sensitive sky location in 918 realizations (1.22%).

3.6 Summary and Conclusions

We searched the NANOGrav 11-year data set for GWs from individual circular SMBHBs. As we found no strong evidence for GWs in our data, we placed limits on the GW strain. We determined that the 11-year data set was most sensitive to $f_{\rm gw} = 8\,\mathrm{nHz}$, for which the sky-averaged strain upper limit was $h_0 < 7.3(3) \times 10^{-15}$. We produced sky maps of the GW strain upper limit at $f_{\rm gw} = 8\,{\rm nHz}$. At the most sensitive sky location, we placed a strain upper limit of $h_0 < 2.0(1) \times 10^{-15}$. These results are the first limits on GWs from individual sources to be robust to uncertainties in the SSE due to the incorporation of BayesEphem in our model.

We introduced a new detection technique that uses "dropout" parameters to determine the significance of a common signal in each individual pulsar. We applied this technique to two low-SNR signals found in the 9-year and 11-year data sets and identified the pulsars contributing the most to these signals. This technique is currently being used within NANOGrav in other GW searches, and a methods paper developing this technique is underway. Determining the physical processes causing these low-SNR signals is beyond the scope of this paper. Advanced noise analyses of all the pulsars in the NANOGrav PTA are underway, using more complicated models for the red noise and incorporating models for time-variations in the dispersion measure, and the methods and results will be the subject of an upcoming paper.

We used our strain upper limits to place lower limits on the luminosity distance to individual SMBHBs. At the most sensitive sky location, we placed a limit of $d_L > 120\,{\rm Mpc}$ for $\mathcal{M} = 10^9\,{\rm M}_\odot$ and $d_L > 5.5\,{\rm Gpc}$ for $\mathcal{M} = 10^{10}\,{\rm M}_\odot$. Our non-detection of GWs was not surprising given our current sensitivity limits. We generated simulated populations of nearby SMBHBs using the method introduced in Mingarelli et al. (2017) and found that only 34 out of 75,000 realizations of the local Universe contained an SMBHB whose GW strain lay above our sky-averaged 95% upper limit curve. These simulations also supported the conclusion that the two low-SNR signals found in the 9-year and 11-year data sets were not GW signals.

Although we have not yet made a positive detection of GWs from individual SMBHBs, the NANOGrav PTA is sensitive enough to place interesting limits on such sources. Based on our non-detection of GWs, we have determined that there are no SMBHBs in the Virgo Cluster with $\mathcal{M} > 1.6(1) \times 10^9\,{\rm M}_\odot$ emitting GWs in the PTA band. Furthermore, our sensitivity to GWs from individual SMBHBs will continue to improve as we increase our observation times, add MSPs to our array, and develop improved pulsar noise models.

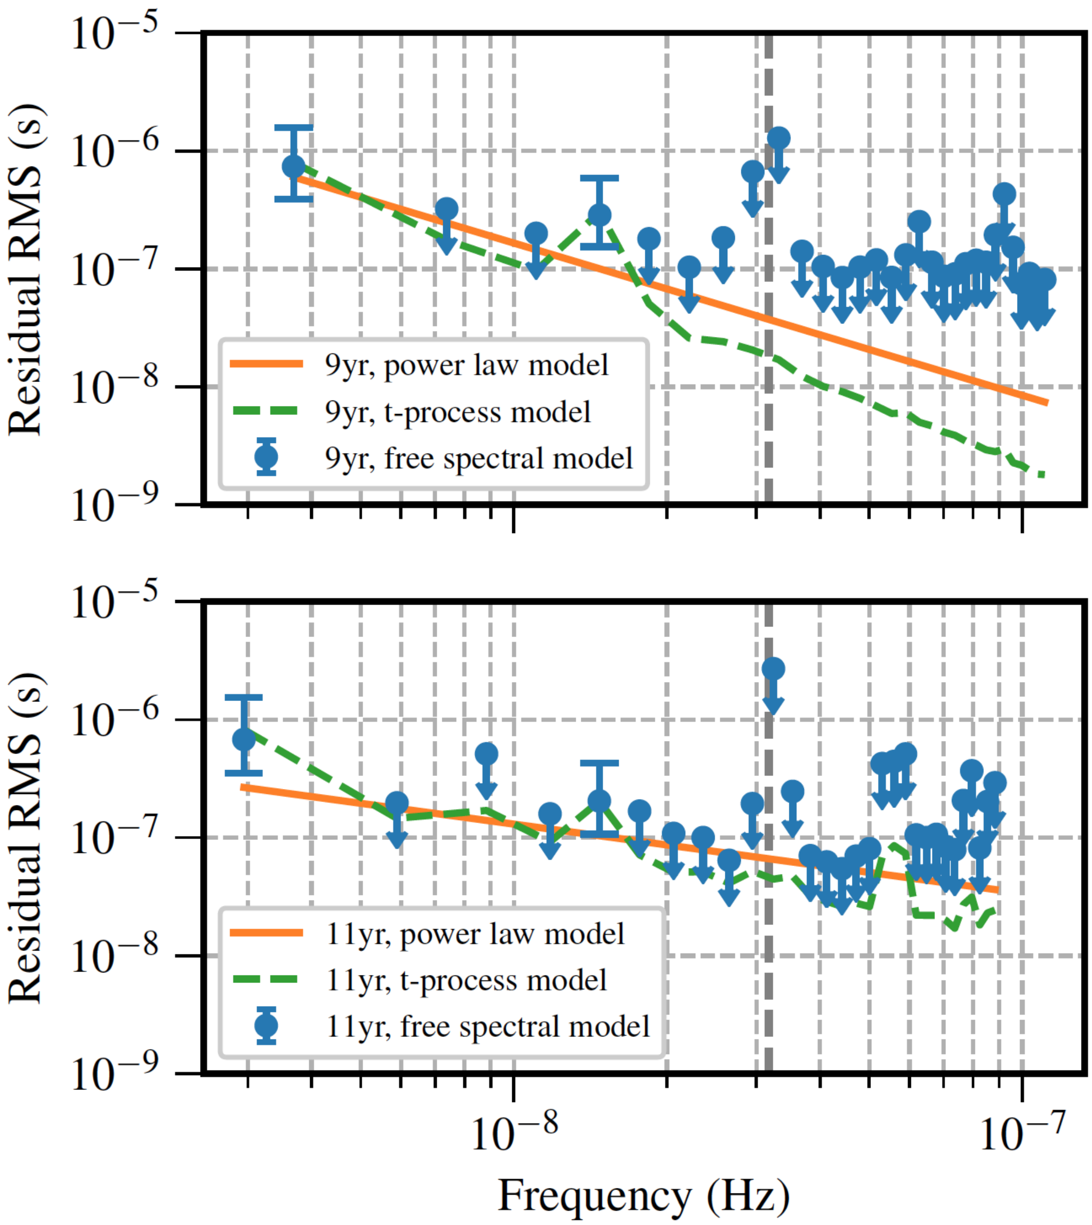

Figure 18: Red noise spectrum of PSR J0613−0200based on the 9-year data set (top panel) and 11-year data set (bottom panel). The vertical dashed line is at $f = 1/(1\mathrm{yr})$; excess noise at this frequency is caused by fitting for the pulsars position and proper motion. We show the results of three different red noise models: a power-law model (solid orange line), a free spectral model (blue points), and a t-process model (dashed green line). We found an excess of noise at a frequency of about 15 nHz in both the 9-year and 11-year data sets using the free spectral model and t-process model, which coincides with the bump in the 9-year sky-averaged CW upper limit spectrum.

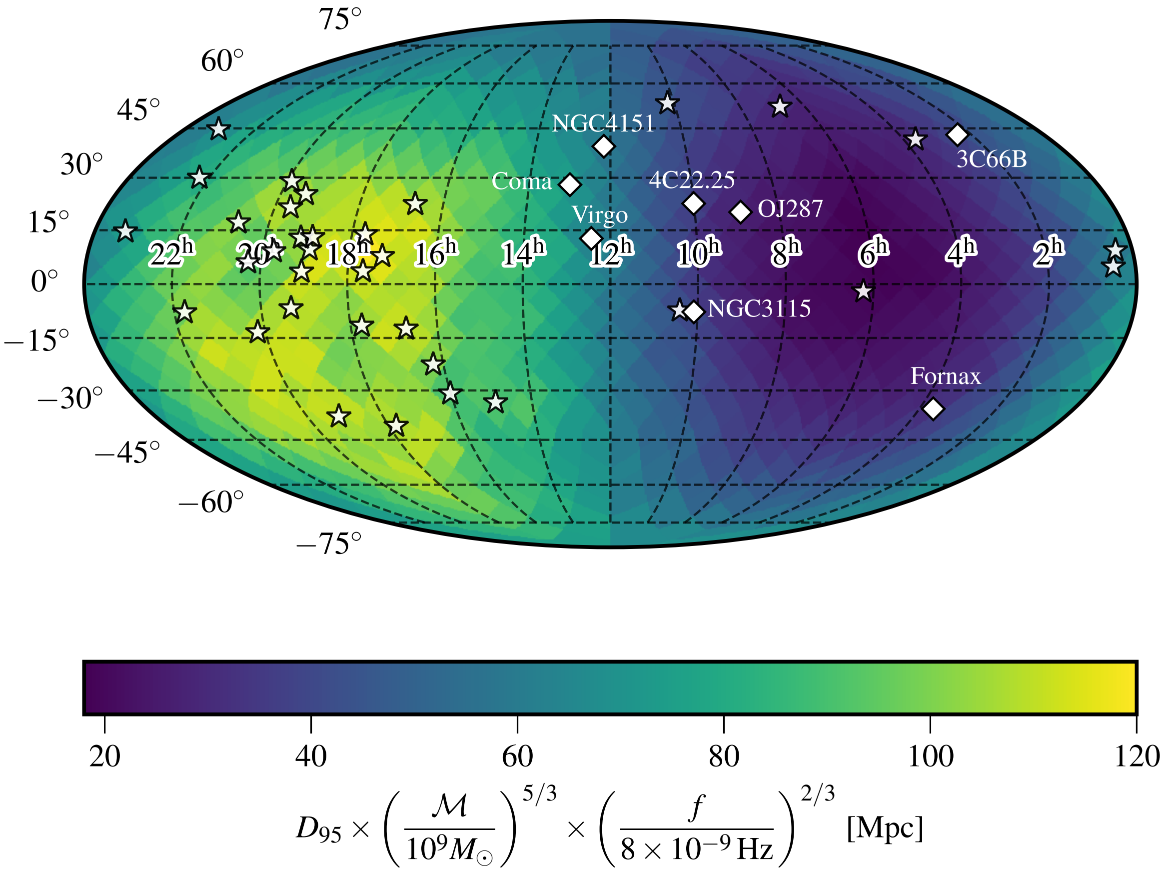

FIGURE 19: The 95% lower limit on the distance to individual SMBHBs with $\mathcal{M} = 10^9\, M_\odot$ and $f_{\rm gw} = 8\,\text{nHz}$ as a function of sky position based on an analysis of the 11-year data set plotted in equatorial coordinates using the Mollweide projection. The stars indicate the positions of pulsars in our array, and the diamonds indicate the positions of known SMBHB candidates or galaxy clusters that may contain SMBHBs. At our most-sensitive sky location, we place a limit of $d_L > 120\,\text{Mpc}$ for SMBHBs with $\mathcal{M} = 10^9\, M_\odot$, and $d_L > 5.5\,\text{Gpc}$ for SMBHBs with $\mathcal{M} = 10^{10}\, M_\odot$.

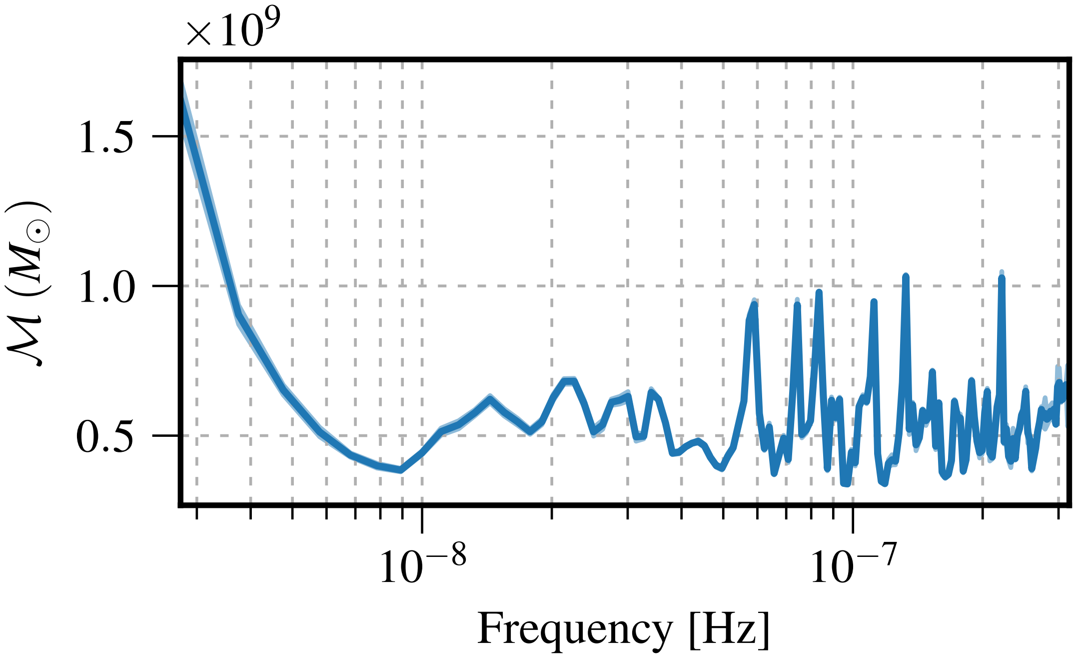

FIGURE 20: The 95% upper limit on the chirp mass of any SMBHBs in the Virgo Cluster as a function of GW frequency. We found that there are no SMBHBs in the Virgo Cluster with $\mathcal{M} > 1.6(1) \times 10^9\, M_\odot$ emitting GWs in this frequency band. At $f_{\rm gw} = 9\,$nHz, we placed an upper limit of $3.8(1) \times 10^8\, M_\odot$.

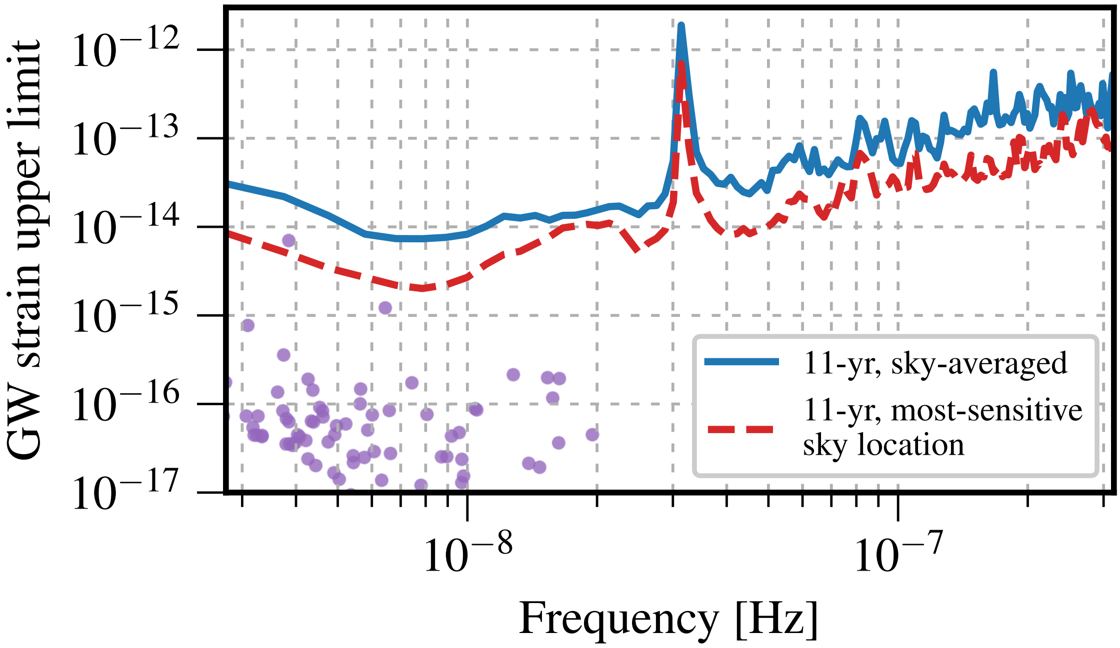

FIGURE 21: GW frequency and strain for the loudest GW sources for a sample realization of the local Universe plotted alongside our 95% strain upper limit curve. This simulation used the method from Mingarelli et al. (2017) to determine the number of SMBHBs emitting GWs in the PTA band based on galaxies in 2MASS (Skrutskie et al. 2006), using merger rates from the Illustris cosmological simulation project (Genel et al. 2014; Rodriguez-Gomez et al. 2015). For this realization, there are 87 SMBHBs – none of them lie above the sky-averaged strain upper limit curve, and there is one source that lies above the strain upper limit curve at the most-sensitive sky location. This source could be detectable depending on its sky location.

Part II

Stochastic Gravitational Wave Signals

Chapter 4

Noise-Marginalized Optimal Statistic

> "*But there is a limit to thinking about even a small piece of something monumental. You still see the shadow of the whole rearing up behind you, and you become lost in your thoughts in part from the panic of realizing the size of that imagined leviathan.*"
>
> — Jeff VanderMeer, *Annihilation*

This chapter is based on:

Noise-marginalized optimal statistic: A robust hybrid frequentist-Bayesian statistic for the stochastic gravitational-wave background in pulsar timing arrays

Sarah J. Vigeland, Kristina Islo, Stephen R. Taylor, Justin A. Ellis.

2018, Phys. Rev. D, 98, 044003

4.1 Introduction

PTAs primarily use Bayesian data analysis methods to search pulsar timing residuals for evidence of gravitational waves and infer signal parameters (van Haasteren et al. 2009; Lentati et al. 2013). Bayesian inference is a powerful tool because it properly accounts for degeneracies between parameters and incorporates all sources of uncertainty into the analysis. However, running a full Bayesian analysis is computationally intensive, particularly when searching for evidence of Hellings and Downs (HD) spatial correlations – the "smoking gun" of the GWB.

The significance of the GWB can also be assessed using the optimal statistic, a frequentist estimator for the GWB amplitude (Anholm et al. 2009b; Demorest et al. 2013b; Chamberlin et al. 2015). Not only does it provide an independent detection procedure, complementing a more robust Bayesian analysis, but it requires significantly less time to compute. In particular, the optimal statistic produces results for a given spatial correlation function within seconds; a full Bayesian analysis including correlations has to run for many weeks on a supercomputing cluster.

However, when pulsars have significant red noise, the optimal statistic gives biased results due to the strong covariance between the individual red noise parameters and the GWB amplitude. Many individual pulsars show evidence for red noise (Lam et al. 2017b; Arzoumanian et al. 2018b), and uncertainty in the position of the Solar System barycenter leads to a common red process in all pulsars (Arzoumanian et al. 2018d).

Here we present a technique for improving the accuracy of the optimal statistic by including an additional step: marginalizing over the individual pulsars' red noise parameters using the posterior distributions from a full Bayesian analysis of all the pulsars. Instead of a single estimate for GWB amplitude $A_{\rm gw}$ we generate a distribution of values from which we extract the mean. This hybrid approach produces a more precise estimate of $A_{\rm gw}$ and its uncertainty, while requiring only a few minutes more than the standard optimal statistic.

Furthermore, the same Bayesian analysis drawn upon by the noise marginalization can be used to compute the optimal statistic for any choice of spatial correlations simply by changing the overlap-reduction function (ORF). For example, clock errors lead to a common red signal with monopole spatial correlations (Hobbs et al. 2012), while uncertainty in the SSB produces dipole spatial correlations (Champion et al. 2010). This technique is used to perform the frequentist searches for common red signals with Hellings and Downs correlations coefficients, monopole, and dipole spatial correlations in the NANOGrav 11-year data set (Arzoumanian et al. 2018d).

4.2 Noise-marginalized optimal statistic

4.2.1 Derivation

The optimal statistic can be derived by analytically maximizing the PTA likelihood function in the weak-signal regime (Anholm et al. 2009b; Chamberlin et al. 2015). It is constructed from the timing residuals $\delta\mathbf{t}$, which can be written as

$$\delta\mathbf{t} = M\epsilon + F\mathbf{a} + U\mathbf{j} + \mathbf{n}\,. \tag{4.2.1}$$

The term $M\epsilon$ describes the contributions to the residuals from perturbations to the timing model. The term $U\mathbf{j}$ describes noise that is correlated for observations made at the same time at different frequencies and uncorrelated over different observing epochs (ECORR), while $\mathbf{n}$ describes uncorrelated white noise from TOA measurement uncertainties. [1]

The term $F\mathbf{a}$ describes red noise, including both red noise intrinsic to the pulsar and a common red noise signal common to all pulsars (such as a GW signal). We model the red noise as a Fourier series,

$$F\mathbf{a} = \sum_{j=1}^{N} \left[a_j \sin\left(\frac{2\pi jt}{T}\right) + b_j \cos\left(\frac{2\pi jt}{T}\right) \right] , \tag{4.2.2}$$

where N is the number of Fourier modes used (typically $N = 30$) and T is the span of the observations.

The optimal statistic is constructed from the auto-covariance and cross-covariance matrices C_a and S_{ab},

$$\begin{aligned} C_a &= \left\langle \delta\mathbf{t}_a \delta\mathbf{t}_a^T \right\rangle , & (4.2.3)\\ S_{ab} &= \left\langle \delta\mathbf{t}_a \delta\mathbf{t}_b^T \right\rangle\Big|_{a\neq b} , & (4.2.4) \end{aligned}$$

where $\delta\mathbf{t}_a$ is a vector of the residuals of the ath pulsar in the PTA. For the GWB with power spectral density (PSD) $P_{\rm gw}(f)$ and overlap reduction function (ORF) Γ_{ab}, the cross-covariance matrices are

$$S_{ab} = F_a\, \phi_{ab}^{\rm gw}\, F_b^T , \tag{4.2.5}$$

where

$$\phi_{ab}^{\rm gw} = \Gamma_{ab}\, P_{\rm gw}(f)\,. \tag{4.2.6}$$

---

[1] see Section 3.1 in Arzoumanian et al. (2014) for details

The ORF is given by the Hellings and Downs curve (Hellings & Downs 1983),

$$\begin{aligned}\Gamma_{ab} &= \frac{1}{2}\left[1-\frac{1}{2}\left(\frac{1-\cos\theta_{ab}}{2}\right)\right. \\ &\left.+3\left(\frac{1-\cos\theta_{ab}}{2}\right)\ln\left(\frac{1-\cos\theta_{ab}}{2}\right)\right]\,, \end{aligned} \tag{4.2.7}$$

where θ_{ab} is the angle between the pulsars. We model the PSD of the GWB as a power law:

$$P_{\mathrm{gw}}(f) = \frac{A_{\mathrm{gw}}^2}{12\pi^2}\left(\frac{f}{f_{\mathrm{yr}}}\right)^{-\gamma}\,, \tag{4.2.8}$$

where $\gamma = 13/3$ assuming SMBHBs evolve solely due to GW emission and $f_{\mathrm{yr}} \equiv 1/(1\,\mathrm{yr})$. The optimal statistic $\hat{A}^2$ is given by

$$\hat{A}^2 = \frac{\sum_{ab}\boldsymbol{\delta}\mathbf{t}_a^T C_a^{-1}\tilde{S}_{ab}C_b^{-1}\boldsymbol{\delta}\mathbf{t}_b}{\sum_{ab}\mathrm{Tr}\left(C_a^{-1}\tilde{S}_{ab}C_b^{-1}\tilde{S}_{ba}\right)}\,, \tag{4.2.9}$$

where $\tilde{S}_{ab}$ is the amplitude-independent cross-correlation matrix,

$$A_{\mathrm{gw}}^2\tilde{S}_{ab} = S_{ab}\,. \tag{4.2.10}$$

This definition of the optimal statistic ensures that $\langle\hat{A}^2\rangle = A_{\mathrm{gw}}^2$. If $A_{\mathrm{gw}} = 0$, the variance of the optimal statistic is

$$\sigma_0 = \left[\sum_{ab}\mathrm{Tr}\left(C_a^{-1}\tilde{S}_{ab}C_b^{-1}\tilde{S}_{ba}\right)\right]^{-1/2}\,. \tag{4.2.11}$$

For a measured value of $\hat{A}^2$, the significance of $\hat{A}^2 \neq 0$ is given by the signal-to-noise ratio (SNR)

$$\rho = \frac{\sum_{ab}\boldsymbol{\delta}\mathbf{t}_a^T C_a^{-1}\tilde{S}_{ab}C_b^{-1}\boldsymbol{\delta}\mathbf{t}_b}{\left[\sum_{ab}\mathrm{Tr}\left(C_a^{-1}\tilde{S}_{ab}C_b^{-1}\tilde{S}_{ba}\right)\right]^{1/2}}\,. \tag{4.2.12}$$

4.2.2 Results for NANOGrav-like PTA

When constructing the residuals $\delta\mathbf{t}_a$, we typically fix the red noise parameters to the values that maximize the single-pulsar likelihood. However, this leads to a bias in the optimal statistic because the individual red noise and common red noise parameters are highly covariant, the optimal statistic estimates computed using fixed red noise parameters are lower than the true value of A_{gw}^2. In this section, we compare three techniques for computing the optimal statistic. First, we fix the individual pulsars' red noise parameters to the maximum-likelihood values from individual Bayesian pulsar noise analyses. Second, we fix the pulsars' red noise parameters to the values that jointly maximize the likelihood for a Bayesian analysis of all of the pulsars in our PTA that searches over the pulsars' red noise parameters and a common red process. For the noise-marginalized method, we draw values of the pulsars' red noise parameters from the posteriors generated by the common Bayesian analysis.

We use these methods to compute the optimal statistic for simulated "NANOGrav-like" data sets consisting of 18 MSPs with observation times, sky positions, and noise properties matching the 18 longest-observed pulsars in the NANOGrav 11-year data set (Arzoumanian et al. 2018b). We include white noise for all pulsars, plus red noise parametrized as a power law,

$$P_a(f) = \frac{A_{\rm red}^2}{12\pi^2}\left(\frac{f}{f_{\rm yr}}\right)^{-\gamma}, \tag{4.2.13}$$

for those pulsars that show evidence of red noise (see Table 1 for more details). We use the PTA data analysis package PAL2 (Ellis & van Haasteren 2017) to perform the noise analyses and compute the optimal statistic.

Fig. 22 shows the fixed-noise and noise-marginalized optimal statistic for a simulation with a GWB with $A_{\rm gw} = 5 \times 10^{-15}$. For this particular realization of the GWB, the fixed-noise analysis using the individual noise results gives $\hat{A}^2 = 6.6 \times 10^{-30}$ with SNR = 2.4, and the fixed-noise analysis using the common noise results gives $\hat{A}^2 = 2.6 \times 10^{-29}$ with SNR = 6.0. The noise-marginalized analysis gives $\hat{A}^2 = (2.5 \pm 0.1) \times 10^{-29}$ with SNR $= 4.8 \pm 0.8$. The value of $\hat{A}^2$ from the fixed-noise analysis using the individual noise results is significantly lower than the injected level of the GWB, while the values of $\hat{A}^2$ from the fixed-noise analysis using the common noise results and the noise-marginalized analysis are in good agreement with each other and the injected value. The fixed-noise analysis using the individual noise results also gives a significantly lower SNR than the other two.

In Fig. 23 we show the optimal statistic for 300 different realizations of a GWB with $A_{\rm gw} = 5 \times 10^{-15}$ computed using the three techniques described above. For the noise-marginalized analysis, we plot the mean values of $\hat{A}^2$ and ρ. Using the noise values

Table 1: Pulsar parameters used in simulated PTA data sets.

Pulsar	T_{obs} (yrs)	σ_w (μs)	A_{red}	γ_{red}
J0030+0451	11.0	0.339	-13.93	3.56
J0613−0200	11.0	0.281	-13.14	1.22
J1012+5307	11.0	0.320	-12.79	1.51
J1024−0719	6.0	0.421	–	–
J1455−3330	11.0	0.773	–	–
J1600−3053	8.0	0.146	–	–
J1614−2230	7.0	0.261	–	–
J1640+2224	11.0	0.202	–	–
J1713+0747	11.0	0.093	-14.14	1.58
J1741+1351	6.0	0.106	–	–
J1744−1134	11.0	0.096	–	–
B1855+09	11.0	0.218	-13.75	3.54
J1853+1303	7.0	0.215	–	–
J1909−3744	11.0	0.034	-13.84	1.74
J1918−0642	11.0	0.342	–	–
J2010−1323	7.0	0.413	–	–
J2145−0750	11.0	0.281	-12.69	1.30
J2317+1439	11.0	0.160	–	–

from individual noise analyses systematically underestimates the strength of the GWB; while using the noise values from a common noise analysis more accurately recovers the injected value. The fixed-noise analysis using the individual noise results finds $\hat{A}^2 = (7.9 \pm 6.8) \times 10^{-30}$ and $\rho = 2.3 \pm 1.5$, averaging over realizations of the GWB. The fixed-noise and noise-marginalized analyses using the common noise results both give $\hat{A}^2 = (2.4 \pm 1.2) \times 10^{-29}$ and $\rho = 4.1 \pm 1.7$.

The fixed-noise and noise-marginalized analyses using the common noise results give the same results for $A_{\mathrm{gw}} = 5 \times 10^{-15}$, but there are differences between them when analyzing data sets containing smaller injected values of $A_{\rm gw}$. In Fig. 24 we show a P–P plot of the cumulative fraction of simulations for which the injected $A^2_{\rm gw}$ lies within a given confidence interval of the measured $\hat{A}^2$. The confidence interval of $\hat{A}^2$ is determined assuming $\hat{A}^2$ follows a Gaussian distribution, with mean and variance $\sigma^2_{\hat{A}^2}$ taken from the distribution for $\hat{A}^2$ found from our 300 realizations of the GWB (i.e., the top panel of Fig. 23). If $\hat{A}^2$ has a Gaussian distribution centered around $A^2_{\rm gw}$, the curves should lie along a straight line with slope equal to unity (the dotted, diagonal lines in Fig. 24).

We compare the three methods for computing the optimal statistic for simulations with $A_{\rm gw} = 5 \times 10^{-15}$, $A_{\rm gw} = 10^{-15}$, and $A_{\rm gw} = 5 \times 10^{-16}$. The fixed-noise optimal statistic using the individual noise results underestimates $\hat{A}^2$ (Fig. 24, left panel). The fixed-noise optimal statistic using the common noise results recovers $\hat{A}^2$ well for large values of $A_{\rm gw}$, but for small values it also underestimates $\hat{A}^2$ (Fig. 24, middle panel). The noise-marginalized optimal statistic provides the most accurate estimate of $\hat{A}^2$ over the range of $A_{\rm gw}$ considered here (Fig. 24, right panel).

4.3 Monopole and Dipole Spatial Correlations

The optimal statistic is particularly well-suited for comparing models of spatial correlations since the ORF enters simply into Eq. (4.2.6). Tiburzi et al. (2016) demonstrated how the optimal statistic can be altered to fit for multiple spatial correlations at once in order to mitigate common noise sources such as clock error and ephemeris error. Here we take a different approach – rather than simultaneously fitting for signals with different spatial correlations, we look at how well we can distinguish between different spatial correlations by computing the optimal statistic with monopole and dipole spatial correlations for the same simulated data sets as in the previous section. For a monopole signal, the ORF becomes simply $\Gamma_{ab} = 1$, while for a dipole signal, the ORF becomes $\Gamma_{ab} = \cos\theta_{ab}$.

Our ability to discriminate different spatial correlations depends on the strength of the GWB and the angular separations between pulsar pairs, θ_{ab}. We can determine the overlap between ORFs corresponding to different spatial correlations by computing the "match statistic" (Cornish & Sampson 2016),

$$\bar{M} = \frac{\sum_{a,b\neq a} \Gamma_{ab}\Gamma'_{ab}}{\sqrt{\left(\sum_{a,b\neq a} \Gamma_{ab}\Gamma_{ab}\right)\left(\sum_{a,b\neq a} \Gamma'_{ab}\Gamma'_{ab}\right)}}, \tag{4.3.1}$$

where Γ and Γ' are two different ORFs. For the 18 pulsars used in these simulations, the match statistic for monopole and HD correlations is $\bar{M} = 0.264$, and the match statistic for dipole and HD correlations is $\bar{M} = 0.337$. These match statistics describe a fundamental limit on our ability to identify the spatial correlations of a common red signal as HD rather than monopole or dipole that depends only on the number of pulsars in our PTA and their sky positions.

Fig. 25 shows the noise-marginalized mean value of $\hat{A}^2$ and the mean SNR computed assuming monopole, dipole, and HD spatial correlations for 300 simulated data sets. Using a monopole or dipole ORF gives a lower value for the mean optimal statistic and mean SNR compared to the HD ORF. Using HD spatial correlations gives $\hat{A}^2 = (2.4\pm1.2)\times10^{-29}$, while using monopole spatial correlations gives $\hat{A}^2 = (2.5\pm3.1)\times10^{-30}$, and dipole spatial correlations gives $\hat{A}^2 = (5.3\pm4.2)\times10^{-30}$. We find a noise-marginalized mean SNR above 1.0 in 97% of our simulated data sets using the HD ORF, and in 50% and 68% of our simulated data sets using the monopole and dipole ORFs, respectively. The mean SNR using the HD ORF, averaged over realizations of the GWB, is 4.1, and we find an SNR greater than this using the monopole and dipole ORFs in just 3% and 3.5% of our simulations, respectively.

This overlap between the monopole, dipole, and HD ORFs also means that a common red process that does not have HD correlations may be confused for a GWB. Fig. 26 shows the results of 300 simulations containing a stochastic signal with dipole spatial correlations. Although a dipole signal has been injected, the HD ORF gives a mean SNR greater than 5 in 82% of the simulations. However, both the monopole and HD ORFs give smaller values of the mean $\hat{A}^2$ and mean SNR compared to the dipole ORF. Furthermore, there are no simulations for which the mean SNR with HD ORF is greater than the mean SNR with dipole ORF. This demonstrates the importance of comparing the SNR from different spatial correlations when determining the type of spatial correlations present.

4.4 Sky Scrambles

The significance of spatial correlations can also be tested with “sky scrambles,” where the ORF is altered in order to simulate changing the pulsars’ positions (Cornish & Sampson 2016; Taylor et al. 2017). The scrambled ORFs are required to have small values of $\bar{M}$ with the true ORF and each other so that they form a nearly-orthogonal set. This ensures that the distribution of $\hat{A}^2$ computed using the scrambled ORFs forms the null hypothesis distribution. Taylor et al. (2017) showed how sky scrambles affect the Bayes’ factor for simulated data sets. We performed a similar analysis using frequentist methods.

We generated 725 scrambled ORFs using a Monte Carlo algorithm. We required the scrambled ORFs to have $\bar{M} < 0.2$ with respect to the true ORF and each other. This threshold was chosen to be comparable to the match statistics between the HD ORF with monopole and dipole ORFs given in Sec. 4.3. We did not choose a smaller threshold because significantly more time would have been needed to generate 725 scrambled ORFs. For each simulation, we computed the noise-marginalized mean optimal statistic and mean SNR for each scrambled ORF, and compared them to the values found using the true ORF.

Fig. 27 shows the results of a sky scramble analysis for a sample data set with $A_{\rm gw} = 5 \times 10^{-15}$. For this particular realization of the GWB, none of the 725 scrambled ORFs resulted in a mean SNR greater than the mean SNR using the true ORF ($p < 0.0014$). In Fig. 28, we plot the distribution of p-values of the 725 sky scrambles for 300 realizations of the GWB. For a GWB with $A_{\rm gw} = 5 \times 10^{-15}$, 95% of the simulations have $p \leq 0.05$, and 74% of the simulations have $p \leq 0.003$. For a GWB with $A_{\rm gw} = 10^{-15}$, 76% of the simulations have $p \leq 0.05$, and 39% have $p \leq 0.003$. This shows that for smaller values of $A_{\rm gw}$, there is a greater chance that noise fluctuations will appear to have the spatial correlations of the GWB.

4.5 Conclusion

The definitive signature of a GWB in PTA data is spatial correlations described by the HD curve. Searching for these using a fully Bayesian approach is computationally expensive, requiring many weeks on a super-computing cluster. In contrast, the optimal statistic can be computed in seconds. Above, we introduced an improved method for computing the optimal statistic, which uses the output from a Bayesian analysis for individual and common red signals to marginalize the optimal statistic over the individual pulsars' red noises. As shown in Sec. 4.2, the noise-marginalized optimal statistic more accurately recovers the GWB amplitude than the fixed-noise optimal statistic, which underestimates the GWB amplitude when significant red noise is present in some pulsars.

Although the noise-marginalized optimal statistic requires computing the optimal statistic thousands of times, it is still many orders of magnitude faster than a Bayesian search. Furthermore, the results from a single Bayesian analysis, which are needed to marginalize over the red noise parameters, can be used to compute the optimal statistic for many different spatial correlations. In Sec. 4.3 we use the noise-marginalized optimal statistic to compare the strength of monopole, dipole, and HD correlations in simulated PTA data with a GWB. In Sec. 4.4 we use the noise-marginalized optimal statistic to perform sky scramble analyses, where we compare the mean SNR computed using the true ORF to the mean SNR computed using scrambled ORFs and measure the significance of HD spatial correlations through the p-value.

The primary strength of the optimal statistic is how quickly it can be computed. This is useful for analyses where the significance of many spatial correlations is compared, as with the sky scrambles. It also makes the optimal statistic a valuable tool for analyzing simulations where many realizations of the GWB are compared. The noise marginalization technique described in this paper is key to being able to accurately measure the GWB with the optimal statistic for real PTAs and realistic PTA simulations, for which red noise is significant.

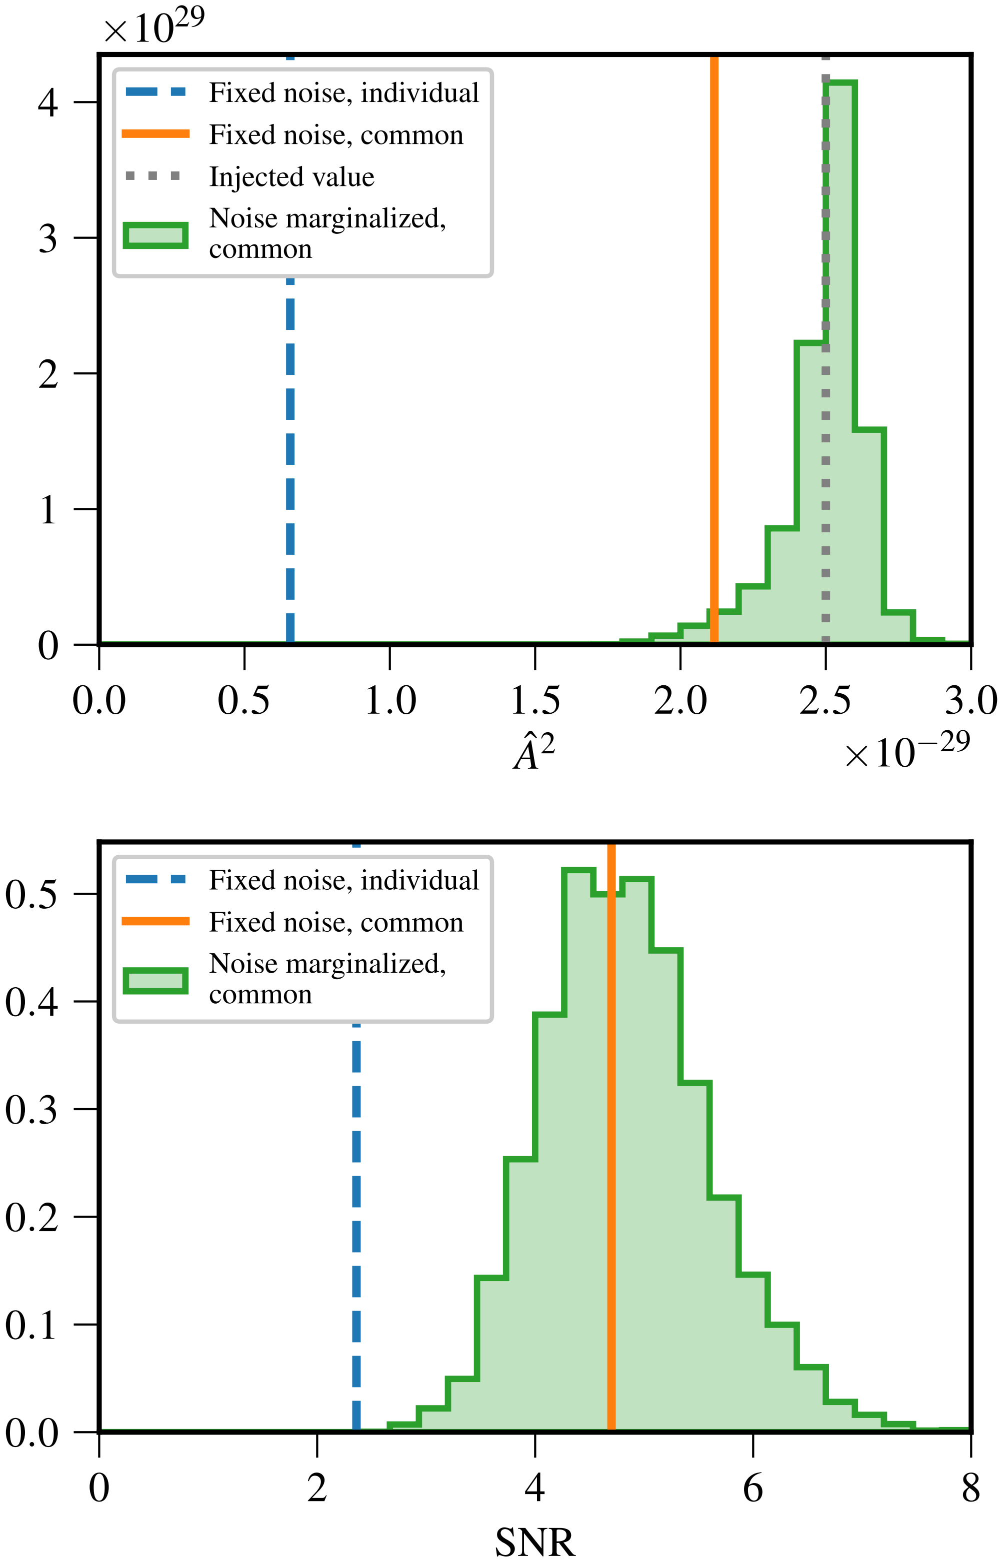

FIGURE 22: Optimal statistic for a simulated PTA data set containing a GWB with $A_{\mathrm{gw}} = 5\times10^{-15}$. The fixed-noise analysis using the individual noise values (dashed blue lines) systematically underestimates $\hat{A}^2$, while the fixed-noise analysis using the common noise values (solid orange lines) and the noise-marginalized analysis (green histograms) more accurately recover A_{gw}.

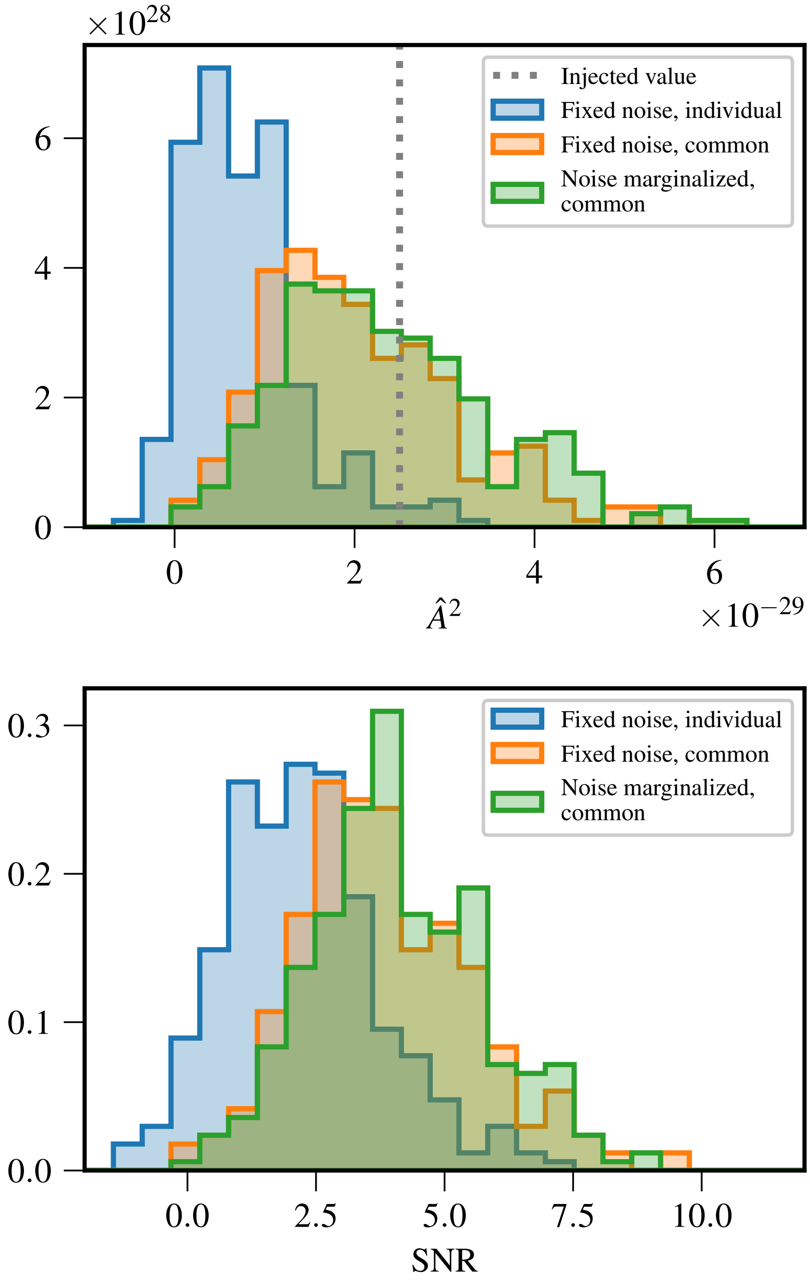

FIGURE 23: Optimal statistic and SNR for 300 simulated data sets containing a GWB with $A_{\mathrm{gw}} = 5 \times 10^{-15}$. The fixed-noise analysis using the individual noise values (blue) systematically underestimates $\hat{A}^2$, while both the fixed-noise analysis using the common noise values (orange) and the noise-marginalized analysis (green) accurately recover A_{gw}.

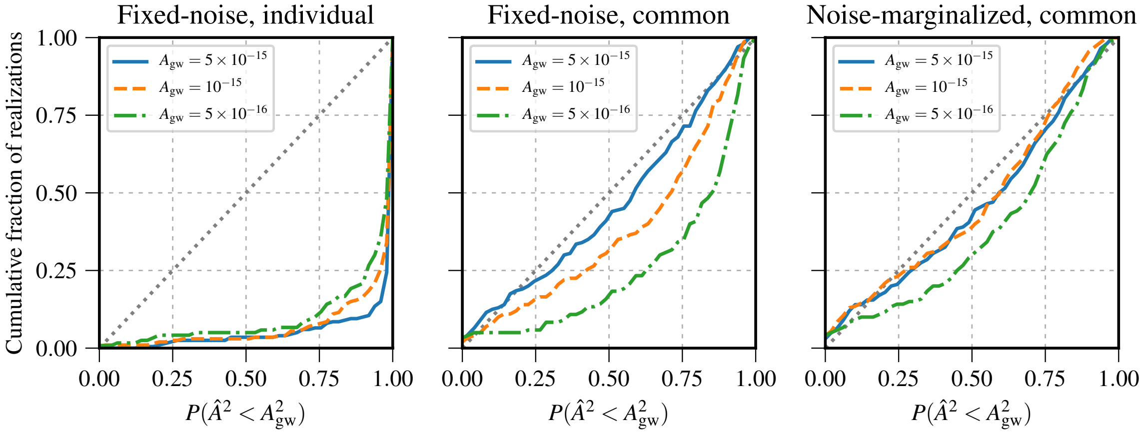

FIGURE 24: P–P plot showing the cumulative fraction of simulations for which $A_{\rm gw}^2$ lies within a given confidence interval of the measured $\hat{A}^2$. The probability distribution of $\hat{A}^2$ is assumed to be a Gaussian with variance $\sigma_{\hat{A}^2}^2$. The fixed-noise optimal statistic using the individual and common noise results both give biased values of $\hat{A}^2$, particularly for small values of $A_{\rm gw}$, while the noise-marginalized optimal statistic gives more accurate values of $\hat{A}^2$ over a large range of injected values of $A_{\rm gw}$.

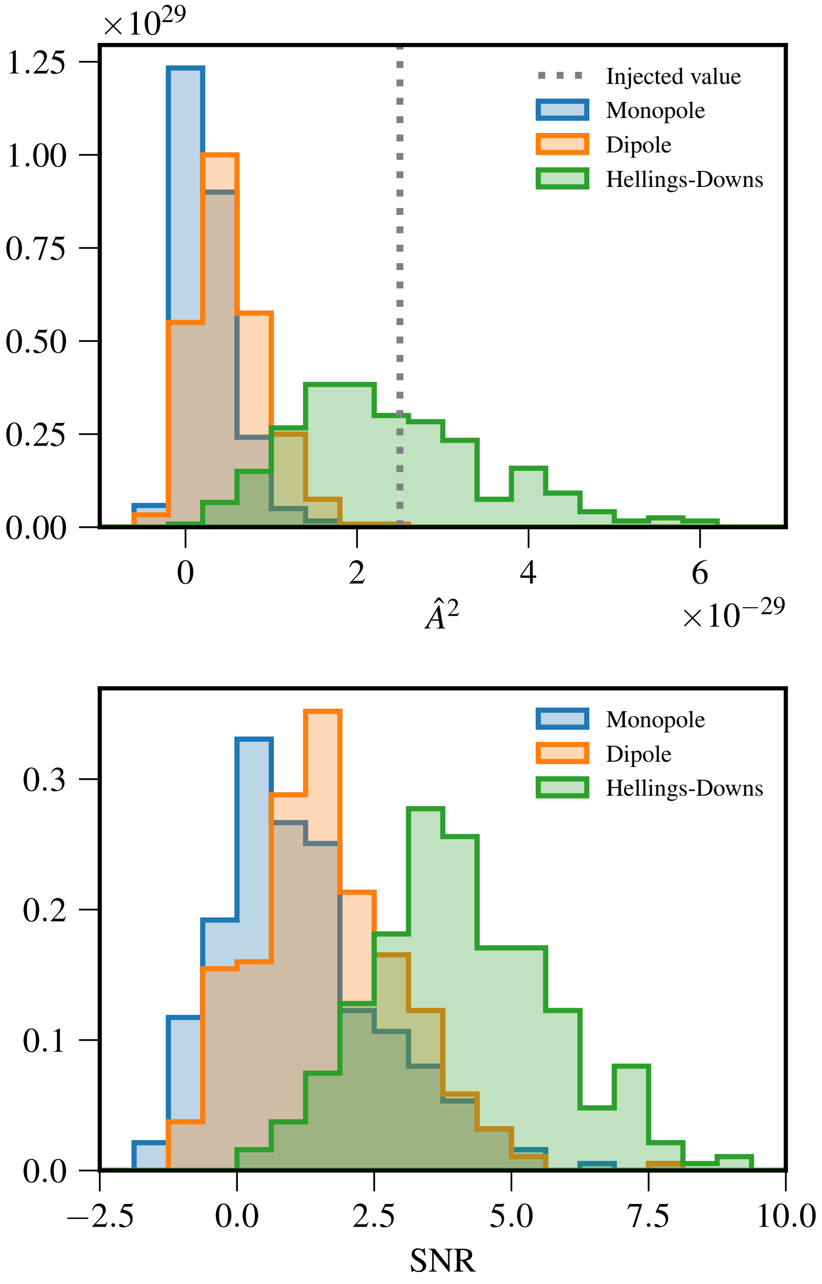

FIGURE 25: Noise-marginalized mean optimal statistic and mean SNR for 300 simulated data sets containing an injected GWB with $A_{\rm gw} = 5 \times 10^{-15}$. We compare the values of the mean optimal statistic and the SNR found using monopole (blue), dipole (orange), and HD (green) spatial correlations. The dashed vertical line indicates the injected value, $\hat{A}^2 = 2.5 \times 10^{-29}$.

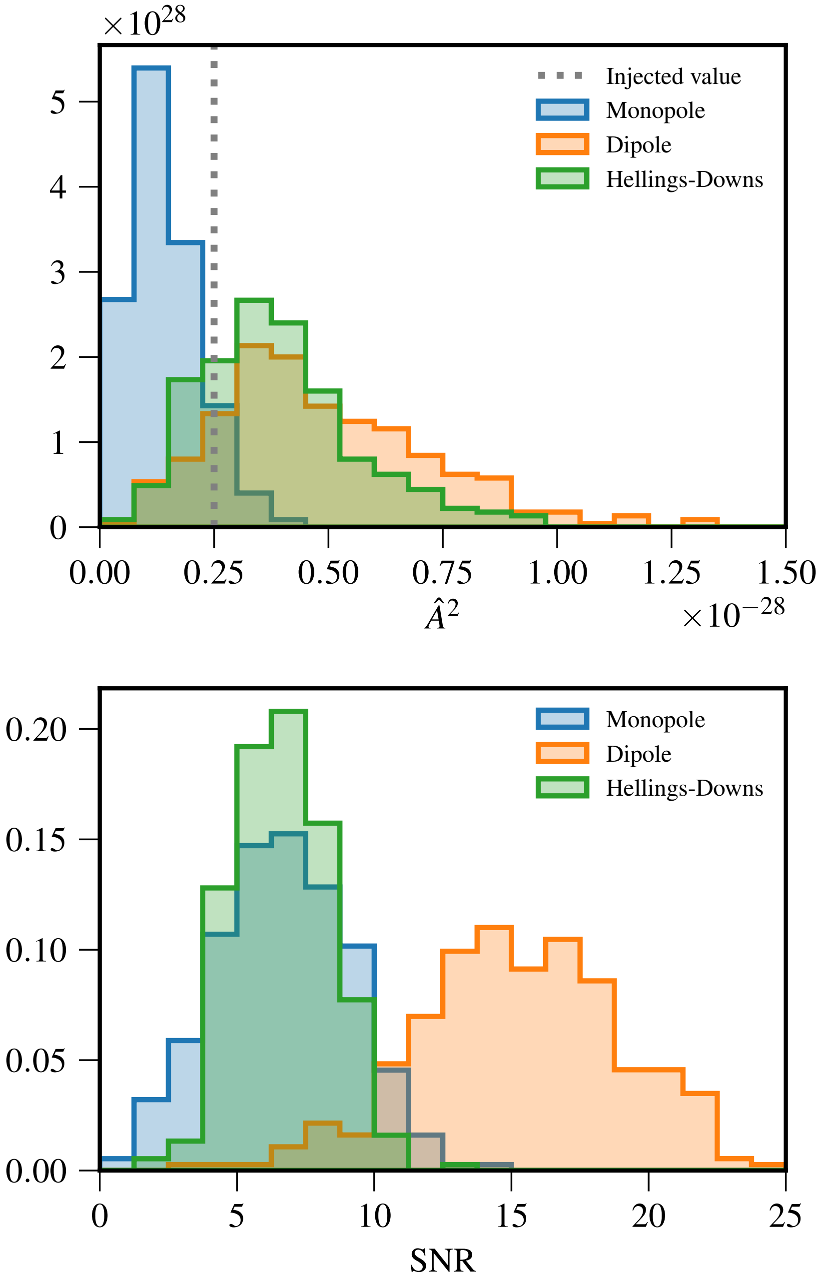

FIGURE 26: Noise-marginalized mean optimal statistic and mean SNR for 300 simulated data sets containing an injected stochastic signal with dipole spatial correlations and $A = 5 \times 10^{-15}$. We compare the values of the optimal statistic and mean SNR found using monopole (blue), dipole (orange), and HD (green) spatial correlations. The dashed vertical line indicates the injected value, $\hat{A}^2 = 2.5 \times 10^{-29}$.

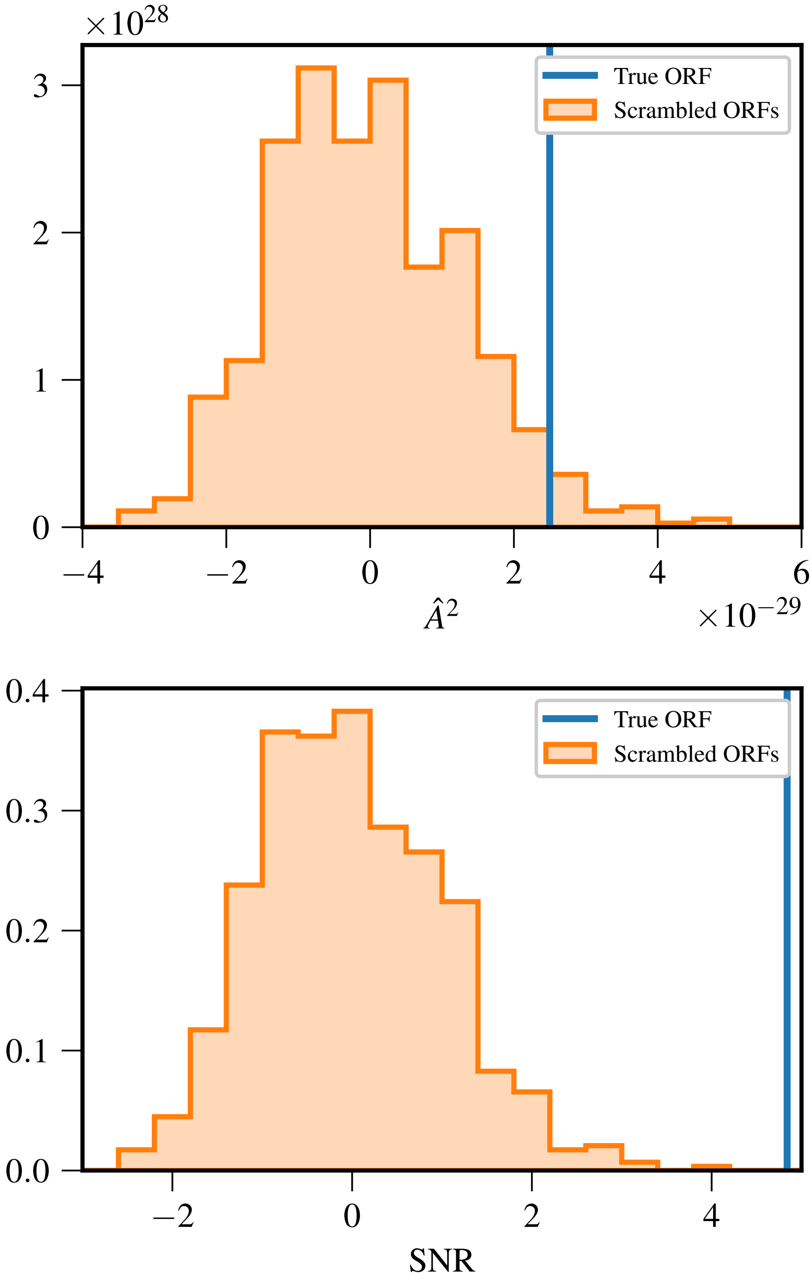

FIGURE 27: Comparison between the noise-marginalized mean optimal statistic and mean SNR with and without sky scrambles for a simulated data set containing a GWB with $A_{\mathrm{gw}} = 5 \times 10^{-15}$. None of the 725 scrambled ORFs gave a mean SNR as large as the mean SNR using the true ORF ($p < 0.0014$).

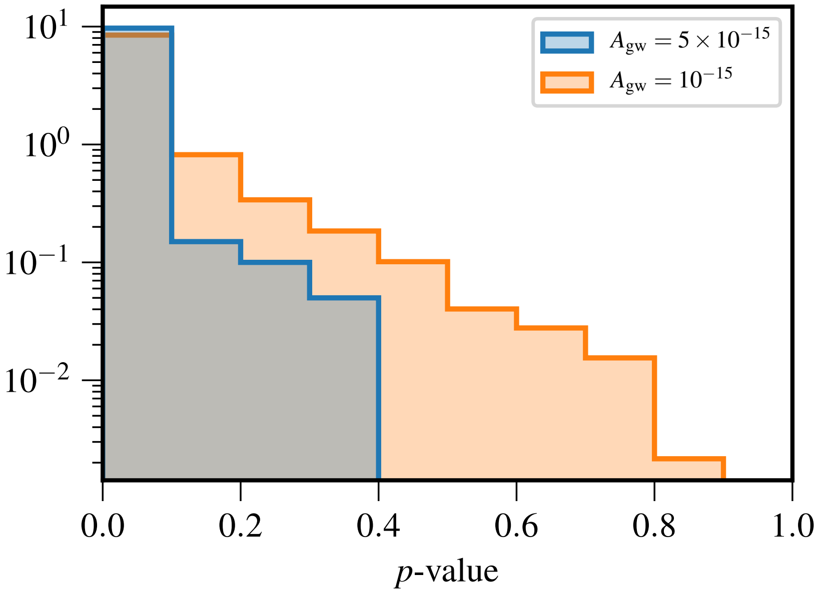

FIGURE 28: Distribution of p-values for the noise-marginalized optimal statistic mean SNR using the true ORF compared to 725 sky scrambles from 300 realizations of the GWB. We show results for simulations with $A_{\rm gw} = 5 \times 10^{-15}$ (blue) and $A_{\rm gw} = 10^{-15}$ (orange).

Chapter 5

Conclusions and Outlook

> "*I have seen the dark universe yawning,*
> *Where the black planets roll without aim;*
> *Where they roll in their horror unheeded,*
> *Without knowledge or lustre or name.*"
>
> — H. P. Lovecraft, *Nemesis*

Here we summarize the essential themes and key findings in the above manuscript.

5.1 Incorporating and Imposing SMBHB Constraints

PTAs and LISA are uniquely poised to detect ultra-low-frequency GW signals. This includes GW bursts with memory, which can be treated as a step-function of the measured strain in the time domain, as discussed in Chap. I. Using a model of the memory signal amplitude derived from General Relativity, we simulated a population of coalescing SMBHBs and compared their GW memory emission to current low-frequency detector sensitivities. We modified the initial inputs from population synthesis to include galaxy masses below $10^{10}\,\mathrm{M}_\odot$ and redshifts out to $z = 3$. This expanded the initial parameter space to include sources in the LISA band. We constructed a galaxy mass-, redshift-, and mass ratio-dependent SMBHB merger rate. Parameterized in this way, we describe the expected rate of detectable events through astrophysical observables.

The incorporation of the decoupling radius allowed us to take into account the efficiency of environmentally-driven orbital decay. Even assuming highly efficient environments, PTAs suffer from a sparser population of coalescing SMBHBs, making memory detection from these sources unlikely. This does not, however, preclude the possibility of detection memory from more exotic sources.

SNRs > 5 can occur on average upwards of $0.5\,\mathrm{yr}^{-1}$ for LISA. This assumes minimal selection bias in the $M_\bullet - M_{\mathrm{bulge}}$ relation used to populate simulated galaxy mergers. The largest SNRs originate from cosmologically local $10^6\,\mathrm{M}_\odot$ binaries. These exhibit evidence of material-rich nuclear cores; strong environment coupling drives these binaries to the GW-emission regime quickly. As a result, these events occur frequently enough to be detectable during LISA's mission.

In Chap. I we were able to use PTA limits of continuous waves to place contraints on nearby candidate binaries. Importantly, we established 95% lower luminosity limits across the sky as a function of chirp mass and GW frequency. At our most sensitive sky location, we found $d_L > 120$ Mpc for $\mathcal{M} = 10^9\,\mathrm{M}_\odot$ and $d_L > 5.5$ Gpc for $\mathcal{M} = 10^{10}\,\mathrm{M}_\odot$ around our most sensitive GW frequency.

If we include the pulsar-term in our Bayesian continuous wave analysis, we can effectively construct a history of GW frequency evolution. Since this depends on the chirp mass of the emitting binary, we can set upper limits on candidate binary chirp mass by searching for evidence of GW frequency evolution. We created a targeted search in the direction of one potential binary candidate: the Virgo cluster. Setting the sky position and distance from electromagnetic observations, our search reconstructed posterior distributions for frequencies in the PTA-band. We showed a maximum chirp mass of $\mathcal{M} > 1.6(1) \times 10^9\, M_\odot$ for the lowest sensitive frequency bin.

5.2 Red Noise in PTA Data

We encountered instances where spectral peculiarities intrinsic to individual pulsars complicated estimates for the strength of both deterministic and stochastic gravitational waves in PTA data. Gravitational waves from a single inspiraling SMBHB appear as excess noise with fixed amplitude and frequency. PTAs are most sensitive to these sources in a frequency range consistent with time-correlated phenomena we term red noise.

While conducting a continuous analysis of NANOGrav's 9-year data set, we found a maximum likelihood value SMBHB candidate near PSR J0613−0200 emitting at $f_{\mathrm{gw}} =$

15 nHz. Although the calculated Bayes factor was uninteresting at ~ 1, the 95% upper limit worsened by roughly a factor of 2 near this frequency when compared to previous upper limit analyses.

A free-spectral analysis of the PSR J0613−0200ś residuals showed an uptick in power at $f_{\rm gw} = 15\,$nHz. When fit to a power-law with fixed amplitude and slope, this frequency saturated the allowable fit, the excess thus leaking into a continuous wave signal. This inspired the development of further diagnostic tools for discriminating between GW signals and unmodeled noise.

The t-process increases the flexibility of a typical red-noise power law. Per frequency bin, we included a deviation parameter which adequately accounts for abrupt increases or decreases in residual power across the PTA-band. Constraining the deviation parameter prior to a t-distribution allowed for accurate modeling without suffering the Occam penalty of an inflated Bayesian prior volume. We demonstrated in Sec. 3.4 how incorporating the t-process adequately mitigates the “bump” to within expected values from scaling laws Moore et al. (2015).

In addition to red noise, we consider the antenna response generated from an $f_{\rm gw} = 15\,$nHz SMBHB in the direction of PSR J0613−0200. An individual source detected at Earth will induce quadrupolar spatial inter-pulsar correlations. Consequently, we must tackle simultaneous model selection across the entire PTA. We do this with a “dropout” analysis. In the full Bayesian search we append to our noise model a GW signal with a multiplier between 0 and 1, corresponding to the GW being absent or present in a subset of pulsars. This effectively serves as model selection in that the posterior for the multiplier indicates which state the data prefer. We evaluated the evidence for the signal appearing each pulsars’ residuals and found the betting odds for PSR J0613−0200 being the culprit against all other pulsars to be $> 2:1$.

In the case of the stochastic GWB, we expect the signal to emerge as red noise shared among all the pulsars in the PTA. Therefore, removing degeneracies between intrinsic and common red noise is of primary focus. One way we combat this confusion is by developing advanced statistics which incorporate intrinsic pulsar noise, including the noise-marginalized optimal statistic. This statistic provides an alternative method for searching

PTA data for evidence of a GWB using the same PTA likelihood as used for Bayesian parameter inference. Importantly, the noise-marginalized optimal statistic delivers quick GWB amplitude estimates without significant differences from a full Bayesian analysis.

Iterations of the optimal statistic take seconds to compute. This makes it trivial to conduct rapid estimates of the amplitude and SNR for a common-red process with various spatial correlations among pulsar timing residuals. Scrambling pulsar positions and finding orthogonal ORFs effectively generated a null distribution against which we evaluated evidence for any specific spatial correlation. Results indicate this estimator prescription has limited ability to arbitrate tension between monopole, dipole and HD correlations to the degree that a Bayesian model selection procedure could. This is primarily true of weak GWBs. As PTA data sets lengthen and more sources of noise are identified and modeled within our likelihood, deploying Bayesian inference will become increasingly cumbersome for rote analyses. Frequentist statistics like these will be necessary to rapidly and accurately report evidence of a GWB in PTA data.

Bibliography

Aggarwal, K., Arzoumanian, Z., Baker, P. T., et al. 2018, arXiv e-prints, arXiv:1812.11585 3, 3.2

Allen, B., & Romano, J. D. 1999, Phys. Rev. D, 59, 102001 1.4

Amaro-Seoane, P., Audley, H., Babak, S., et al. 2017, arXiv e-prints, arXiv:1702.00786 (document), 1.1, 4

Anholm, M., Ballmer, S., Creighton, J. D. E., Price, L. R., & Siemens, X. 2009a, Phys. Rev. D, 79, 084030 1.4

Anholm, M., Ballmer, S., Creighton, J. D. E., Price, L. R., & Siemens, X. 2009b, Phys. Rev. D, 79, 084030 4.1, 4.2.1

Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2014, ApJ, 794, 141 3.1, 3.3, 3.3.2, 3.3.2, 3.3.2, 3.3.3, 3.3.3, 1

—. 2015, ApJ, 813, 65 3.1

—. 2016, ApJ, 821, 13 3.4

—. 2018a, The Astrophysical Journal Supplement Series, 235, 37 1.4, 3.2

—. 2018b, ApJS, 235, 37 3.1, 4.1, 4.2.2

Arzoumanian, Z., Baker, P. T., Brazier, A., et al. 2018c, ApJ, 859, 47 2.2

—. 2018d, ArXiv e-prints, arXiv:1801.02617 3.3.2, 3.3.3, 3.3.3, 3.4, 4.1

Babak, S., Petiteau, A., Sesana, A., et al. 2016, MNRAS, 455, 1665 3.4.2, 3.5

Baldassare, V. F., Geha, M., & Greene, J. 2018, ApJ, 868, 152 2.4

Barausse, E. 2012, MNRAS, 423, 2533 2.2

Begelman, M. C., Blandford, R. D., & Rees, M. J. 1980, Nature, 287, 307 2.2, 2.2, 2.4

Blanchet, L., & Damour, T. 1992, Phys. Rev. D, 46, 4304 2.1

Braginskii, V. B., & Thorne, K. S. 1987, Nature, 327, 123 2.3.2

Burke-Spolaor, S., Taylor, S. R., Charisi, M., et al. 2018, arXiv e-prints, arXiv:1811.08826 1.3

Burt, B. J., Lommen, A. N., & Finn, L. S. 2011, ApJ, 730, 17 3.2

Chamberlin, S. J., Creighton, J. D. E., Siemens, X., et al. 2015, Phys. Rev. D, 91, 044048 4.1, 4.2.1

Champion, D. J., Hobbs, G. B., Manchester, R. N., et al. 2010, ApJ, 720, L201 4.1

Christodoulou, D. 1991, Phys. Rev. Lett., 67, 1486 2.1

Christy, B., Anella, R., Lommen, A., et al. 2014, ApJ, 794, 163 3.2

Colpi, M. 2014, Space Sci. Rev., 183, 189 2.2

Corbin, V., & Cornish, N. J. 2010, ArXiv e-prints, arXiv:1008.1782 3.3.2, 3.3.3

Cordes, J. M. 2013, Classical and Quantum Gravity, 30, 224002 1.2, 3.3.2

Cornish, N. J., & Sampson, L. 2016, Phys. Rev. D, 93, 104047 4.3, 4.4

Cromartie, H. T., Fonseca, E., Ransom, S. M., et al. 2019, arXiv e-prints, arXiv:1904.06759 1.2

Demorest, P. B., Ferdman, R. D., Gonzalez, M. E., et al. 2013a, ApJ, 762, 94 3.4

—. 2013b, ApJ, 762, 94 4.1

Desvignes, G., Caballero, R. N., Lentati, L., et al. 2016, MNRAS, 458, 3341 1.2, 3.4.2

Detweiler, S. 1979, ApJ, 234, 1100 1.2

Dickey, J. M. 1971, The Annals of Mathematical Statistics, 42, 204 3.3.3

Dolch, T., Lam, M. T., Cordes, J., et al. 2014, ApJ, 794, 21 3.3.2

Efron, B., & Stein, C. 1981, Ann. Statist., 9, 586 3.4.3

Ellis, J., & van Haasteren, R. 2017, jellis18/PAL2: PAL2, doi:10.5281/zenodo.251456 4.2.2

Ellis, J. A., Siemens, X., & Creighton, J. D. E. 2012, ApJ, 756, 175 3.3.2

Enoki, M., Inoue, K. T., Nagashima, M., & Sugiyama, N. 2004, ApJ, 615, 19 2.1

Favata, M. 2009, ApJ, 696, L159 2.3.1, 2.3.2

Flanagan, E. E., & Hughes, S. A. 1998, Phys. Rev. D, 57, 4566 2.3.2

Folkner, W. M., & Park, R. S. 2016, JPL planetary and Lunar ephemeris DE436, Tech. rep., Jet Propulsion Laboratory, Pasadena, CA, `https://naif.jpl.nasa.gov/pub/naif/JUNO/kernels/spk/de436s.bsp.lbl` 3.3.3

Gebhardt, K., Adams, J., Richstone, D., et al. 2011, ApJ, 729, 119 3.5

Genel, S., Vogelsberger, M., Springel, V., et al. 2014, MNRAS, 445, 175 3.5, 21

Gültekin, K., Richstone, D. O., Gebhardt, K., et al. 2009, ApJ, 698, 198 1.3

Hallinan, G., Ravi, V., Weinreb, S., et al. 2019, The DSA-2000: A Radio Survey Camera for The Next Decade, `http://www.astro.caltech.edu/~vikram/DSA_2000__Specifications.pdf` 2.5

Hartnett, J. G., & Luiten, A. N. 2011, Reviews of Modern Physics, 83, 1 1.2

Hellings, R. W., & Downs, G. S. 1983, ApJ, 265, L39 1.4, 4.2.1

Hobbs, G. 2013, Classical and Quantum Gravity, 30, 224007 1.2

Hobbs, G., Coles, W., Manchester, R. N., et al. 2012, MNRAS, 427, 2780 4.1

Ilbert, O., McCracken, H. J., Le Fèvre, O., et al. 2013, A&A, 556, A55 2.2, 2.4

Johnson, A. D., Kapadia, S. J., Osborne, A., Hixon, A., & Kennefick, D. 2019, Phys. Rev. D, 99, 044045 2.3.2

Jones, M. L., McLaughlin, M. A., Lam, M. T., et al. 2017, ApJ, 841, 125 3.3.2

Kelley, L. Z., Blecha, L., & Hernquist, L. 2017a, MNRAS, 464, 3131 2.2

Kelley, L. Z., Blecha, L., Hernquist, L., Sesana, A., & Taylor, S. R. 2017b, MNRAS, 471, 4508 1.3

Khan, F. M., Holley-Bockelmann, K., Berczik, P., & Just, A. 2013, ApJ, 773, 100 2.2

Kormendy, J., & Ho, L. C. 2013, ARA&A, 51, 511 2.2, 2.4

Kulier, A., Ostriker, J. P., Natarajan, P., Lackner, C. N., & Cen, R. 2015, ApJ, 799, 178 2.1

Lam, M. T., Cordes, J. M., Chatterjee, S., et al. 2017a, ApJ, 834, 35 3.3.2, 3.3.2

—. 2017b, ApJ, 834, 35 4.1

Lasky, P. D., Thrane, E., Levin, Y., Blackman, J., & Chen, Y. 2016, Phys. Rev. Lett., 117, 061102 2.3.2

Lee, K. J., Wex, N., Kramer, M., et al. 2011, MNRAS, 414, 3251 3.3.2

Lentati, L., Alexander, P., Hobson, M. P., et al. 2013, Phys. Rev. D, 87, 104021 4.1

Madison, D. R., Chernoff, D. F., & Cordes, J. M. 2017, Phys. Rev. D, 96, 123016 2.3.1

Martynov, D. V., Hall, E. D., Abbott, B. P., et al. 2016, Phys. Rev. D, 93, 112004 1.1

McConnell, N. J., & Ma, C.-P. 2013, ApJ, 764, 184 2.2, 2.4

McLaughlin, M. A. 2013, Classical and Quantum Gravity, 30, 224008 1.2

McNeill, L. O., Thrane, E., & Lasky, P. D. 2017, Phys. Rev. Lett., 118, 181103 2.3.2

Merritt, D., & Milosavljević, M. 2005, Living Reviews in Relativity, 8, astro-ph/0410364 2.2

Mingarelli, C. M. F., Lazio, T. J. W., Sesana, A., et al. 2017, Nature Astronomy, 1, 886 (document), 3.5, 3.6, 21

Moore, C. J. 2015, in Journal of Physics Conference Series, Vol. 610, Journal of Physics Conference Series, 012018 (document), 4

Moore, C. J., Taylor, S. R., & Gair, J. R. 2015, Classical and Quantum Gravity, 32, 055004 5.2

NANOGrav Collaboration. 2018, arXiv e-prints, arXiv:1810.06594 2.5

Phinney, E. S. 2001, arXiv e-prints, astro 1.4

Pollney, D., & Reisswig, C. 2011, The Astrophysical Journal Letters, 732, L13 2.2, 2.3.1

Rasskazov, A., & Merritt, D. 2017, The Astrophysical Journal, 837, 135 2.2

Ravi, V., Wyithe, J. S. B., Hobbs, G., et al. 2012, ApJ, 761, 84 2.1

Ravi, V., Wyithe, J. S. B., Shannon, R. M., & Hobbs, G. 2015, MNRAS, 447, 2772 2.4

Rodriguez-Gomez, V., Genel, S., Vogelsberger, M., et al. 2015, MNRAS, 449, 49 3.5, 21

Roedig, C., Dotti, M., Sesana, A., Cuadra, J., & Colpi, M. 2011, MNRAS, 415, 3033 2.2

Rosado, P. A., Sesana, A., & Gair, J. 2015, MNRAS, 451, 2417 1.4

Rusli, S. P., Thomas, J., Saglia, R. P., et al. 2013, AJ, 146, 45 3.5

Sazhin, M. V. 1978, Soviet Ast., 22, 36 1.2

Sesana, A. 2013, Classical and Quantum Gravity, 30, 224014 1.3

Sesana, A., Barausse, E., Dotti, M., & Rossi, E. M. 2014, ApJ, 794, 104 2.2, 2.4

Sesana, A., Haardt, F., Madau, P., & Volonteri, M. 2004, ApJ, 611, 623 1.3

Sesana, A., & Vecchio, A. 2010, Phys. Rev. D, 81, 104008 3.3.2

Sesana, A., Vecchio, A., & Colacino, C. N. 2008, MNRAS, 390, 192 2.1

Sesana, A., Volonteri, M., & Haardt, F. 2007, MNRAS, 377, 1711 2.4

Shankar, F., Bernardi, M., Sheth, R. K., et al. 2016, MNRAS, 460, 3119 2.2, 2.4

Shen, J., & Gebhardt, K. 2010, ApJ, 711, 484 3.5

Simon, J., & Burke-Spolaor, S. 2016, ApJ, 826, 11 2.2, 2.2

Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, AJ, 131, 1163 (document), 3.5, 21

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. 2014, Journal of Machine Learning Research, 15, 1929 3.4.3

Strominger, A. 2017, arXiv:1703.05448 2.1

Tang, Y., Haiman, Z., & MacFadyen, A. 2018, MNRAS, 476, 2249 2.4

Taylor, S. R., Huerta, E. A., Gair, J. R., & McWilliams, S. T. 2016, ApJ, 817, 70 3.3.2

Taylor, S. R., Lentati, L., Babak, S., et al. 2017, Phys. Rev. D, 95, 042002 4.4

Taylor, S. R., Simon, J., & Sampson, L. 2017, Phys. Rev. Lett., 118, 181102 2.2

Taylor, S. R., Burke-Spolaor, S., Baker, P. T., et al. 2019, arXiv e-prints, arXiv:1903.08183 2

Thorne, K. S. 1992, Phys. Rev. D, 45, 520 2.1

Thrane, E., & Romano, J. D. 2013, Phys. Rev. D, 88, 124032 1.4

Tiburzi, C., Hobbs, G., Kerr, M., et al. 2016, MNRAS, 455, 4339 4.3

van Haasteren, R., & Levin, Y. 2010, MNRAS, 401, 2372 2.3.2, 2.3.2

van Haasteren, R., Levin, Y., McDonald, P., & Lu, T. 2009, MNRAS, 395, 1005 4.1

Vasiliev, E., Antonini, F., & Merritt, D. 2015, ApJ, 810, 49 2.2

Verbiest, J. P. W., Weisberg, J. M., Chael, A. A., Lee, K. J., & Lorimer, D. R. 2012, ApJ, 755, 39 3.3.3

Verbiest, J. P. W., Lentati, L., Hobbs, G., et al. 2016a, MNRAS, 458, 1267 1.2

—. 2016b, MNRAS, 458, 1267 1.2

Vigeland, S. J., Vallisneri, M., & Taylor, S. R. 2019, in preparation 3.4.3

Volonteri, M., Haardt, F., & Madau, P. 2003, ApJ, 582, 559 2.2, 2.2

Wahlquist, H. 1987, General Relativity and Gravitation, 19, 1101 3.3.2, 3.3.2

Wang, J. B., Hobbs, G., Coles, W., et al. 2015, MNRAS, 446, 1657 6

Wang, Y., & Mohanty, S. D. 2018, Journal of Physics: Conference Series, 957, 012003 2.5

Wright, A. H., Robotham, A. S. G., Driver, S. P., et al. 2017, MNRAS, 470, 283 2.2

CIRRICULUM VITAE

Kristina Islo

Education

Ph.D., Physics, University of Wisconsin Milwaukee, WI, USA **May 2019**

Dissertation Advisor: Prof. Xavier Siemens

Dissertation Title: *Measuring Deterministic and Stochastic Gravitational Waves with Pulsar Timing Array Experiments*

B.S., Physics, University of Wisconsin Milwaukee, WI, USA **June 2012**

Degree conferred with honor.

Employment

Research and Teaching Assistant, UWM, Milwaukee WI, USA **September 2012–May 2019**

Honors and Awards

Dean's List, UWM	**Fall 2006–Spring 2011**
Sophmore Honors, UWM	**2007**
Phi Beta Kappa Honors Society, UWM	**Spring 2011 –Present**
Ed Zore Honors College Scholarship	**2008 –2009**
Lawrence Baldassaro Honors College Scholarship	**2009**
James Martin Award Scholarship	**June 2009**

Teaching Experience

Teaching Assistant, University of Wisconsin Milwaukee

Phys. 209, Physics I **Fall 2011, Spring 2012, Fall 2011, Spring 2013, Fall 2014**

Phys. 210, Physics II **Spring 2014, Fall 2015, Fall 2016**

Phys. 110, Physics for the Health Professions **Fall 2013, Fall 2015**

Phys. 120, General Physics I **Summer 2013**

PROFESSIONAL ACTIVITIES, OUTREACH, AND SERVICE

Member, NANOGrav **2010–Present**

Member, Pierre Auger Collaboration **2010–2016**

Conference organizer

NANOGrav Fall 2017 Science Meeting, Lafayette College **October 2017**

Outreach

Invited panelist, Medical College of Wisconsin **February 2018**

Women in Science: Tales from the Lab

Invited panelist, Conference for Undergraduate Women in Physics **January 2017**

University of Wisconsin Madison

Coffeeshop Astrophysics public lecture series speaker and volunteer **2015–2019**

Approximately every four months

Invited speaker, Milwaukee Underground Science Society **August 2017**

The Summer Solar Eclipse Countdown!

NANOGrav exhibit representative **September 2016, 2017, 2018**

Milwaukee Maker Faire

Publications in Progress

13. The NANOGrav collaboration [including **Islo, K.**] (2019) *The NANOGrav 11-year Data set: Constraints on Bursts with Memory*

Submitted Publications

12. **Islo, K.**, Siemens, X., Simon, J., Burke-Spoloar, S., (2018) *Prospects for Memory Detection with Low-Frequency Gravitational Wave Detectors*

Collaboration Publications

11. *The NANOGrav 11-year Data Set: Pulsar Timing Constraints on the Stochastic Gravitational-Wave Background* Astrophysical Journal, **859**, 47 [arXiv:1801.02617]

Refereed Publications

10. Vigeland, S., **Islo, K.**, Taylor, S. R., Ellis, J. A., *Noise-marginalized optimal statistic: A robust hybrid frequentist-Bayesian statistic for the stochastic gravitational-wave background in pulsar timing arrays Physical Review D* **98***, 044003 [arXiv:1805.12188]*

9. *The NANOGrav collaboration [including* **Islo, K.***] (2018) The NANOGrav 11-year Data set: Constraints on Continuous Wave from Individual Supermassive Black Hole Binaries. [arXiv:1812.11585]*

Contributed Talks & Posters

8. "Prospects for Detecting Gravitational Wave Memory with Low-Frequency Gravitational Wave Detectors"

   American Physical Society Meeting April 2019

7. "Prospects for Detecting Gravitational Wave Memory with Low-Frequency Gravitational Wave Detectors"

American Astronomical Society Meeting January 2019

6. "Prospects for Detecting Gravitational Wave Memory with Low-Frequency Gravitational Wave Detectors"
   NANOGrav meeting, Green Bank Observatory October 2018

5. "Prospects for Detecting Gravitational Wave Memory with Low-Frequency Gravitational Wave Detectors"
   28th Midwest Relativity Meeting October 2018

4. "Prospects for Detecting Gravitational Wave Memory with Pulsar Timing Array Experiments"
   27th Midwest Relativity Meeting October 2017

3. "Gravitational Wave Memory Source Populations"
   NANOGrav meeting, McGill University, Montreal October 2015

2. "Measuring Gravitational Wave Memory with Pulsar Timing Experiment"
   CGCA seminar, UWM March 2016

1. "Generating a Memory Background from Coalescing Binary Black Holes"
   International Pulsar Timing Array (IPTA) conference July 2015