

Approximating Simulated Stochastic Gravitational Wave Background with Broken Splines and Power Laws via MCMC Fitting

Author: Taylor Knapp¹

Mentors: Patrick Meyers², Arianna Renzi²

¹Brown University, Providence, RI, U.S.

²California Institute of Technology, Pasadena, CA, U.S.

(Dated: July 7, 2022)

In the search for gravitational waves (GWs), researchers have begun to investigate what makes up the stochastic gravitational wave background (SGWB). Detecting the SGWB may give insight into the decoupling of the gravitational field from the rest of the early evolving Universe and the composition of activity in the Universe at any given time. To better fit and understand SGWB, we develop a fitting pipeline based on the Metropolis-Hastings Markov Chain Monte Carlo (MCMC) algorithm using spline and power law fits. We introduce various proposals for probing the parameter space via spline and power law functions. For now, we look to match the power laws we inject into the simulated data from our pipeline with the MCMC-based fitter. In the future, we will work on understanding the physical meanings of our parameter posteriors for both the spline and power law fits.

I. INTRODUCTION

Gravitational waves are perturbations of the space-time manifold expressed by metric tensor $g_{\mu\nu}$ [6]. We are able to detect these fluctuations of space and time as strain, or change in length per unit length. This strain is detected by ground-based interferometers; two of such detectors are LIGO Hanford, WA and LIGO Livingston, LA. This style of interferometer involves two beams of light at the same wavelength propagated orthogonally to each other. These beams reflect off mirrors and coherently return to the source. When a gravitational wave passes, it strains the arms of the detector. This causes the light beams to move out of phase with one another, and so when the beams are recombined the resulting change in the interference pattern is evidence of a passing perturbation of spacetime.

Four primary sources of gravitational waves are coalescing binary systems, pulsars, supernovae, and stochastic gravitational wave backgrounds (SGWB) [1]. We know that coalescing binary systems appear as "chirps" with an uncertainty arising from the unknown number density of coalescences. These chirps are the only signals we have detected so far. Pulsars should appear as sine waves due to their periodic emission of gravitational waves. Supernovae are extremely challenging sources to understand since we have yet to detect or observe them, and parametrising their signal in order to include them in matched filtering searches is extremely unlikely. The fourth source, SGWBs, encompasses the unresolved gravitational wave sources. These unresolved sources include, for example, fluctuations from just after the Big Bang, as well as unresolved astrophysical sources like compact binary coalescences. "Stochastic" refers to a non-deterministic strain signal, either due to the generation process or detector limitations. We cross-correlate data from different detectors to try to detect the SGWB.

Understanding the SGWB will help researchers probe the Universe earlier than electromagnetic signals cur-

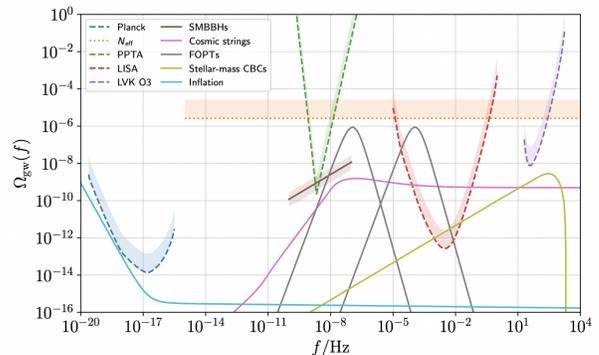


FIG. 1. Various GWB sources, their sensitivity, and their energy density. Figure from Ref [6].

rently allow [1]. Electromagnetic signals go back to about 400,000 years after the Big Bang, when scattering of particles decreased enough for photons to travel unimpeded. The SGWB could take us as far as 10^{-32} s after the Big Bang because GWs propagate through spacetime without the risk of scattering off particles [7]. For comparison, Planck Time is 5.39×10^{-44} s after the Big Bang. Thus, resolving the primordial background could help paint a clearer picture of the early Universe [6].

Determining a better and more general fit for the SGWB signal will also help us learn about the background itself [1]. As we add more time to our background detection survey, we will be able to resolve more features in the SGWB spectrum. The features in the SGWB spectrum will help us learn about the signals beyond the recognizable, precise events. As Allen [1] mentions, the more we record the signals on multiple detectors simultaneously, the better we will be able to transition the sources of the SGWB from "unresolved" to "resolved", helping contextualize our fitting parameters.

Additionally, developing better fits will help bound the

stochastic background signal. Narrowing down the frequency ranges where the SGWB signal is present will be helpful to deduce the components of this signal [9]. By fitting parameters to the models we develop during this project, we may be able to better constrain where to turn our attention in our GW searches. Since GW detection is a relatively new scientific development, interpreting as much of the data as we have now will only help us better understand what makes up the SGWB.

II. BACKGROUND

A. Power Law Spectrum

Most models for a GWB predict a power-law spectrum, which is given by:

$$\Omega_{\text{GW}}(f) = \Omega_{\text{ref}} \left(\frac{f}{f_{\text{ref}}} \right)^\alpha, \quad (1)$$

where $\Omega_{\text{GW}}(f)$ is the energy density per logarithmic frequency interval used to describe the isotropic stochastic background. This quantity can also be expressed as $\Omega_{\text{GW}}(f) = \frac{f}{\rho_c} \frac{d\rho_{\text{GW}}}{df}$ where ρ_c is the critical density and ρ_{GW} is energy density of gravitational waves in the infinitesimal frequency interval f to $f + df$ [1]. Fig. 1 provides a visualization for the energy densities expected from different sources across the frequency interval [6]. Ω_{ref} is the the amplitude at a reference frequency, f_{ref} . α is the spectral index. Both Ω_{ref} and α are constrained using strain data. Right now, we can fit various parameter combinations for different frequency ranges of our spectrum [4]. Fig. 2 shows the log-log profile of our expected GWB, with the solid red line indicating the GWB after subtracting events we can detect individually [8]. Since this is not linear, we cannot perfectly fit a power law to the solid red line. Fig. 2 is what we expect from third generation detectors. So far, we have tried to fit various regions of the frequency spectrum with individual power laws but aim to fit the entire profile as best as possible.

B. Alternative Fittings

Current alternative functional fittings for the SGWB are as follows:

- Power Law: $\Omega_{\text{GW}}(f) = \Omega_{\text{ref}} \left(\frac{f}{f_{\text{peak}}} \right)^\alpha$
- Broken Power law (BPL):

$$\Omega_{\text{GW}}(f) = \begin{cases} \Omega_{\text{peak}} \left(\frac{f}{f_{\text{peak}}} \right)^{\alpha_1} & \text{for } f \leq f_{\text{peak}} \\ \Omega_{\text{peak}} \left(\frac{f}{f_{\text{peak}}} \right)^{\alpha_2} & \text{for } f > f_{\text{peak}} \end{cases}$$
- Smooth BPL:

$$\Omega_{\text{GW}}(f) = \Omega_{\text{peak}} \left(\frac{f}{f_{\text{peak}}} \right)^{\alpha_1} \left[1 + \left(\frac{f}{f_{\text{peak}}} \right)^\Delta \right]^{(\alpha_2 - \alpha_1)/\Delta}$$

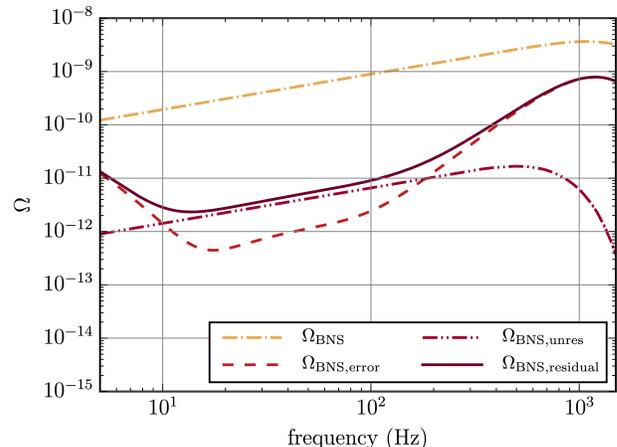


FIG. 2. The lines on this log-log plot correspond to our signal and errors. The solid red line is expected from our GWB search, which does not resemble a power law after subtracting the binary neutron stars (BNS). Figure from Ref. [8]

These models, although simplistic and described by few parameters which require fitting, are not as general and generic as we would like [5]. Alternative functional approaches include *spline fitting*. Spline fitting utilizes smooth, piece-wise polynomials of different degrees to describe a curve. Parameters come in the form of coefficients of a polynomial expansion:

$$p_j(x) = a_0 + a_1x + a_2x^2 + \dots + a_nx^n \quad (2)$$

such that the a_i coefficients allow us to fit an n-degree polynomial to the curve segment j . This is advantageous where a single polynomial fit, such as attempting to use a single power law for the entire spectrum, fails. We will start with spline fitting to recover these parameters and their relationships to each other when constructing functional models for our data. In our SGWB analysis, we only require spline fit interpolations up to $n = 3$.

III. WESTLEY FITTING ALGORITHM

The Westley fitting algorithm utilizes a combination of single power laws and splines to interpolate a fit between an optimized number of knots. Westley is a Reversible Jump Metropolis-Hastings Markov Chain Monte Carlo algorithm. This means that a ratio of likelihoods between points drives the placement of nodes in either the spline or power law fits. The term "Reversible Jump" means that the MCMC can propose adding or removing parameters from a model instead of just probing the existing parameters. First, a guess is made for the placement of the set of nodes along a frequency interval. For now, this guess is made on a uniform prior over the local frequency bin. Next we calculate the likelihood of this node configuration, which is a function of the cross-correlation be-

tween detectors and the model we have injected. We are essentially proposing to move a node, interpolating between the nodes to calculate the model, then using that node-motivated model to calculate the likelihood of the data. We then propose to move the amplitude of a node, add another node, or remove a node, and then calculate the likelihood again. If the likelihood of the second node configuration is greater than the likelihood of the first node placement, the second node is kept. Otherwise, we throw out the second node configuration and keep the first node configuration. We repeat this process to form a chain, which should converge at a particular fit to the data.

A. Relevant Equations

An MCMC parameter probe such as Westley relies on probability formulations and statistics. In the following proposals, a set of equations is used to evaluate the likelihood of the proposal and the Hastings ratio, R . We define the acceptance probability for a node as:

$$P(m'|m) = \min \left(1, \frac{p(m') p(d|m') q(m|m')}{p(m) p(d|m) q(m'|m)} \right). \quad (3)$$

The lower-case p functions are standard probabilities while $q(m|m')$ is the ratio of probability to move from model m to model m' and vice versa. The R value, or Hastings ratio, is embedded into the acceptance probability and is expressed as:

$$R = \frac{q(m|m') p(m')}{q(m'|m) p(m)}. \quad (4)$$

The difference between the acceptance probability and the Hastings ratio lies in the probability of the data given the various models that are multiplied into the acceptance probability.

For now, the Hastings ratios are consistently $R=1$. However, in future Gaussian-based updates to the proposals, R will be more complicated, including a Gaussian exponential term. Future work will be deriving and implementing that expression.

B. Theoretical Proposals

Westley includes 4 proposals for node movement: birth, death, horizontal shift, vertical shift. These proposals create the changes to the nodes that are responsible for the acceptance or rejection of the nodes described above via changes in the likelihood. The fitter chooses a proposal by randomly generating a value between 0 and 1. Depending on the range of the random value, one of the proposals is chosen to alter a node.

1. Birth

The birth proposal turns "on" a node that was previously "off". First, a random index is chosen out of the "off" nodes which are the nodes set to "False". This node is then set to "True", thus turned "on". The newly birthed node, is set to the frequency and amplitude that it was last at before it was turned off. This is because the fitter saves the location of the nodes when they are turned off but just does not incorporate them into the fit calculations. The fit is then re-calibrated between the updated nodes and returned.

2. Death

This is the most straightforward proposal within the fitting. A random node index is selected from the list of active nodes. The chosen node is then set to "False" and removed from the list of active nodes. The fit is then re-calibrated between the remaining nodes and returned.

3. Horizontal Shift

A horizontal node shift corresponds to a change in the frequency position of the node. Based on the maximum number of nodes allowed (inputted as an argument into the fitter), the total frequency range is divided into even intervals. A random node is chosen corresponding to one of the segments of the overall frequency interval from 20-200 Hz. Then, the new frequency position of the node is chosen from a uniform prior over the frequency bin. The original amplitude of the node is kept. The fit is then re-calibrated between the updated nodes and returned.

4. Vertical Shift

A vertical node shift corresponds to a change in the amplitude or cross-correlation value of the node. First, a random node is chosen to be vertically perturbed. Then, we draw a normalized Gaussian about the initial vertical position of the node. Next, a random value within the first STD of the initial location is chosen and the node is moved to that vertical location. The original frequency position of the node is kept. The fit is then re-calibrated between the updated nodes and returned.

C. Updates on Proposals

I have been focusing on two primary updates to the proposals after writing the horizontal shift proposal. First, I am looking at generalizing the vertical shift proposal. Initially, that proposal only updated the amplitude of one node at a time. Now, every "on" node is perturbed about its initial position up to one STD. Adding

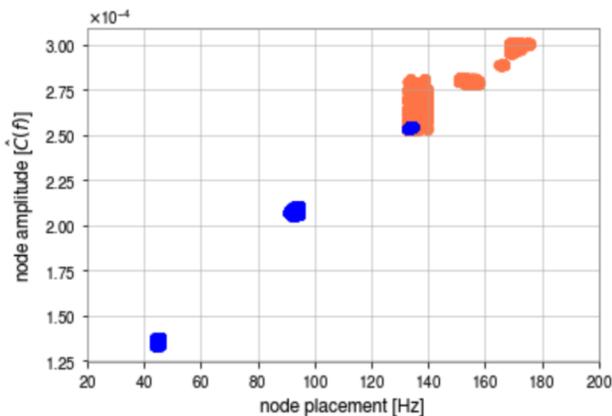


FIG. 3. Evidence of node shifting given updated proposals given by the point smearing. The Westley fitter tends to 2 nodes on average, which allows for a single power law. Blue corresponds to the first node and orange corresponds to the second node of the fitting.

in more movement allows the fitter to find the optimal node placements quicker and probe a larger expanse of the parameter space in each iteration of the fitter. So far, I have gotten the horizontal and vertical perturbations to work for various nodes, as shown in Fig. 3. The visible smearing about the first and second node locations for an injected single power law indicate the updated proposals are working for the fitter.

Next, I am working on making a more data-motivated birth proposal. Initially, the birth proposal would turn on a node where it had been turned off previously or initialized. This is a problem because the other nodes have likely moved since the node we are birthing was turned off. Thus, I am working on generating the birth proposal at the nearest data point to the last "on" location. This causes the newly birthed node to correspond directly to the data before it is moved around by the other proposals. An alternative approach is to place the newly birthed node as close to the neighboring "on" node as possible. This would place the node within the known fit bounds and the newly birthed node could be perturbed into place in following proposals. I plan to explore these both to determine which is a better approach, if there is any variation.

IV. SIMPLE MODELS

To test the effectiveness of the fitter in the context of SGWB, I injected both single and broken power laws to see if the known injection could be recovered. While the focus of my project is on the Westley fitter itself, I initially created a pipeline from the simulation of SGWB data through the pyGWB Python package. This pipeline returns a Baseline object that can be used to fit the energy density of the simulated data with the Westley

MCMC. The signal injection occurs at the beginning of the data simulation with pyGWB, where the injection is specified to the package and transformed into a time series and frequency spectrum automatically.

A. Single Power Law

As a base case, I injected the following single power law into the pipeline:

$$\Omega_{GW}(f) = \Omega_{ref} \left(\frac{f}{f_{ref}} \right)^\alpha, \quad (5)$$

where $\Omega_{ref} = 5 \times 10^{-5}$, $f_{ref} = 10$ Hz, and $\alpha = 2/3$. Fig. 4 shows the injection (orange) and the simulated data by the pyGWB package detector simulator. Following simulation within the pipeline, the cross-correlation or energy density is returned as shown in Fig. 5. The orange corresponds to the simulated noise, the green with the injected single power law, and the blue with the overall combined data. The green power law is what we wish to recover with our fitting. The fitting is only provided the blue and orange data and noise from Fig. 5.

We then fit the data and noise with Westley. When probing the returned data, we see the Westley fitter favors 2 nodes to fit the data, which corresponds to a single power law. Visually, these posterior fits have found the injected signal. Fig. 6 shows the spectral indices for this fit. Since we injected $2/3$, we expect to get back this same value if the fits are as good as they appear in Fig. 6. Calculating the spectral index is a simple computation using the node locations:

$$\alpha = \frac{\log_{10} \Omega_2 - \log_{10} \Omega_1}{\log_{10} f_2 - \log_{10} f_1}, \quad (6)$$

where Ω_i and f_i reflect the positions of nodes $i = 1, 2$. Fig. 6 shows the histogram of the spectral indices for the returned fits. The estimated spectral index is $\alpha = 0.66 \pm 0.03$ where 0.03 is the standard deviation of the sampled spectral index. We observe that the average is close to the injected value. We expect this value to refine and the histogram to appear more solidly Gaussian with more iterations of the Westley fitter.

B. Broken Power Law

As a slightly more complex case, I injected a broken power law (BPL) into the pipeline:

$$\Omega_{GW}(f) = \begin{cases} \Omega_{peak} \left(\frac{f}{f_{peak}} \right)^{\alpha_1} & \text{for } f \leq f_{peak} \\ \Omega_{peak} \left(\frac{f}{f_{peak}} \right)^{\alpha_2} & \text{for } f > f_{peak} \end{cases} \quad (7)$$

with $\Omega_{peak} = 5 \times 10^{-5}$, $f_{peak} = 45$ Hz, $\alpha_1 = 6$, and $\alpha_2 = 2/3$. Applying our Westley fitter yields the posterior fits shown in Fig. 7. The BPL fits do not

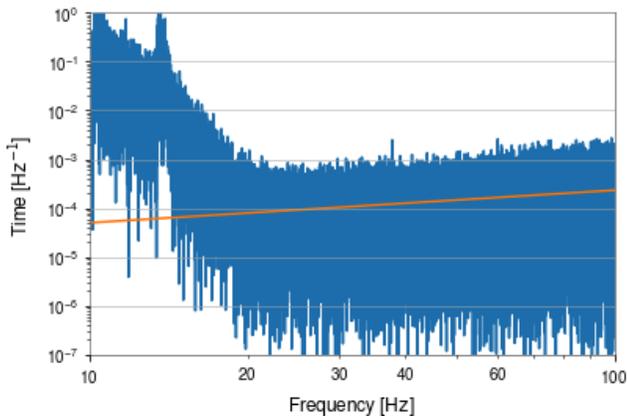


FIG. 4. Injected single power law (orange) over the simulated data (blue) using the pyGWB detector strain simulator. This plot reflects the data from one of the simulated detectors. Two sets such as this are required for the pipeline to generate the cross-correlation function.

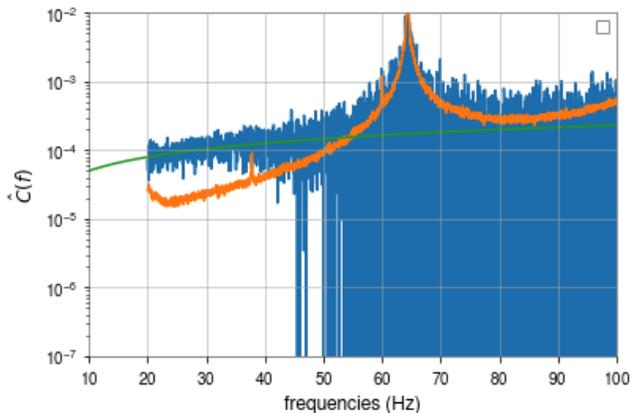


FIG. 5. The cross-correlation as a function of frequency after the pyGWB pipeline incorporating simulated data from Hanford and Livingston detectors. The green is the same injected signal shown in Fig. 4. The simulated detector noise is in orange and the simulated cross-correlation data is in blue. The orange and blue are fed into the Westley fitter to determine a best fit.

appear as neat as the SPL fits, with BPL posterior turnovers between 130-160 Hz, evident in Fig. 7.

An interesting feature of the BPL fits is the node placement tendencies. The fitter favors 4 nodes when only fitting with power laws, drawing a power law between each node that translates to a linear segment in log-log space. In Fig. 7, the first and second nodes find the minimum frequency of 20 Hz and the knee of the injected BPL at 45 Hz. We expect this. The second 2 nodes are a bit more unexpected because we only anticipate 3 total nodes. We need one before the peak, one at the peak, and one after the peak to get the two spectral indices. The placements of the second 2 nodes indicate a

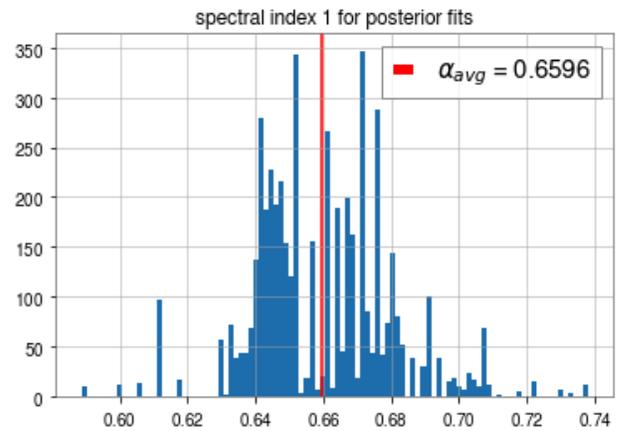


FIG. 6. Histogram of spectral indices calculated for 2 node fits in the single power law case. Our injected signal has $\alpha = 2/3$. The average spectral value for our posterior is in red.

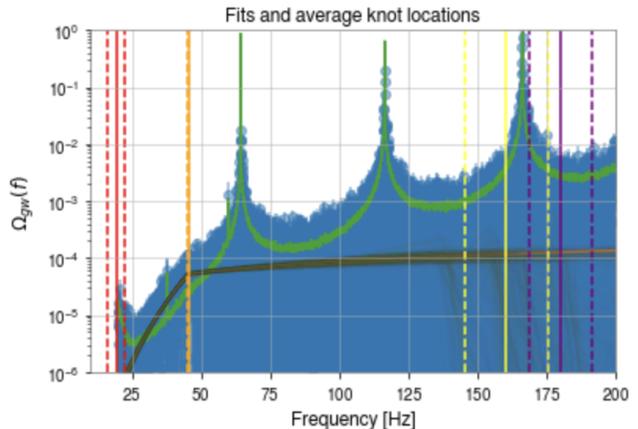


FIG. 7. The posterior fits of an injected BPL. The dark orange band reflects the posterior fits from Westley of the injected BPL. Green is noise and blue is simulated data. Vertical solid lines correspond to the average node placement for the BPL, which tends to 4 node fits. The vertical dotted lines of the same color indicate one STD from the average node value. Red is the first node, orange is the second, yellow is the third, and purple is the fourth.

turnover at higher frequencies, which does not follow our injected BPL. The first STD of the third and fourth node placement, indicated by the dotted lines about the average value solid line, are much wider spaced than the first two nodes. Perhaps running more iterations of these fits will help more patterns emerge for the locations of these second two nodes in the power law fits of the injected BPL. In doing this, we may be able to discover why the fits diverge at higher frequencies with an injected BPL.

V. NEXT STEPS

Before we can begin interpreting the physical meaning of our fit parameters, we need to run the Westley fitter for more iterations. Patterns in the spectral indices and node placements are emerging at a million iterations, but more iterations will better reinforce the fits. Furthermore, by running the fitter for longer, we can thin out the results more. Thinning removes a certain percentage of the posterior fits, helping eliminate the outliers of the fits. Increasing the thinning from 10 to >100 will reduce the number of fits received but hopefully increase the quality. This is something that can be probed following the proposal refining previously discussed.

Once the parameters have converged or give a nearly Gaussian posterior, we can begin to physically interpret these values. The benefits of using power law fits are we know the physical correspondence of the parameter values and can connect them with our injected values. While this only holds for our simulated data, we can use these trends on real data to interpret the fits we

get back. Alternatively, spline fits give more parameters and flexibility, but we do not yet know the physical correspondence. This would be an interesting avenue of investigation going forward using higher iteration runs. However, the fits systematically utilize a higher number of nodes in spline fits, so this would be something to confirm and investigate constraining before moving on to interpreting.

VI. ACKNOWLEDGEMENTS

I'd like to thank Patrick Meyers and Arianna Renzini for their tremendous support in this project. Also, I'd like to thank the other LIGO Summer Undergraduate Research Fellows (SURFs) for their collaborative spirit and encouraging teamwork. I also gratefully acknowledge the support from the National Science Foundation Research Experience for Undergraduates (NSF REU) program, the California Institute of Technology, and the LIGO Summer Undergraduate Research Fellowship.

-
- [1] Bruce Allen. The stochastic gravity-wave background: sources and detection. In Relativistic Gravitation and Gravitational Radiation, Proceedings of the Les Houches School of Physics, held in Les Houches, Haute Savoie, volume 26, pages 373–418, 1997.
 - [2] Thomas Callister, Letizia Sammut, Shi Qiu, Ilya Mandel, and Eric Thrane. Limits of astrophysics with gravitational-wave backgrounds. Physical Review X, 6(3):031018, 2016.
 - [3] Eanna E Flanagan. Sensitivity of the laser interferometer gravitational wave observatory to a stochastic background, and its dependence on the detector orientations. Physical Review D, 48(6):2389, 1993.
 - [4] Vuk Mandic, Eric Thrane, Stefanos Giampanis, and Tania Regimbau. Parameter estimation in searches for the stochastic gravitational-wave background. Physical review letters, 109(17):171102, 2012.
 - [5] Katarina Martinovic, Patrick M Meyers, Mairi Sakellariadou, and Nelson Christensen. Simultaneous estimation of astrophysical and cosmological stochastic gravitational-wave backgrounds with terrestrial detectors. Physical Review D, 103(4):043023, 2021.
 - [6] Arianna I Renzini, Boris Goncharov, Alexander C Jenkins, and Patrick M Meyers. Stochastic gravitational-wave backgrounds: Current detection efforts and future prospects. Galaxies, 10(1):34, 2022.
 - [7] Joseph D Romano, Neil Cornish, et al. Detection methods for stochastic gravitational-wave backgrounds: a unified treatment. Living reviews in relativity, 20(1):1–223, 2017.
 - [8] Surabhi Sachdev, Tania Regimbau, and BS Sathyaprakash. Subtracting compact binary foreground sources to reveal primordial gravitational-wave backgrounds. Physical Review D, 102(2):024051, 2020.
 - [9] LIGO Scientific, BP Abbott, R Abbott, TD Abbott, S Abraham, F Acernese, K Ackley, C Adams, VB Adya, C Affeldt, et al. Search for the isotropic stochastic background using data from advanced ligo's second observing run. Physical Review D, 100(6):061101, 2019.
 - [10] Ashish Sharma and Jan Harms. Searching for cosmological gravitational-wave backgrounds with third-generation detectors in the presence of an astrophysical foreground. Physical Review D, 102(6):063009, 2020.