

Comparison of maximum-likelihood mapping methods for gravitational-wave backgrounds

Arianna I. Renzini^{*}

*LIGO Laboratory, California Institute of Technology, Pasadena, California 91125, USA
and Department of Physics, California Institute of Technology, Pasadena, California 91125, USA*

Joseph D. Romano

*Department of Physics and Astronomy, Texas Tech University,
Box 41051, Lubbock, Texas 79409-1051, USA*

Carlo R. Contaldi

Blackett Laboratory, Imperial College, London, SW7 2AZ, United Kingdom

Neil J. Cornish

*eXtreme Gravity Institute, Department of Physics, Montana State University,
Bozeman, Montana 59717, USA*

 (Received 15 July 2021; revised 4 December 2021; accepted 4 January 2022; published 13 January 2022)

Detection of a stochastic background of gravitational waves is likely to occur in the next few years. Beyond searches for the isotropic component of a stochastic gravitational-wave background, there have been various mapping methods proposed to target anisotropic backgrounds. Some of these methods have been applied to data taken by the Laser Interferometer Gravitational-wave Observatory (LIGO) and Virgo. Specifically, these directional searches have focused on mapping the *intensity* of the signal on the sky via maximum-likelihood solutions. We compare this intensity mapping approach to a previously proposed, but never employed, *amplitude-phase* mapping method to understand whether this latter approach may be employed in future searches. We build up our understanding of the differences between these two approaches by analyzing simple toy models of time-stream data, and we run mock-data mapping tests for the two methods. We find that the amplitude-phase method is only applicable to the case of a background which is phase coherent on large scales or, at the very least, has an intrinsic coherence scale that is larger than the resolution of the detector. Otherwise, the amplitude-phase mapping method leads to an overall loss of information, with respect to both phase and amplitude. Since we do not expect these phase-coherent properties to hold for any of the gravitational-wave background signals we hope to detect in the near future, we conclude that intensity mapping is the preferred method for such backgrounds.

DOI: [10.1103/PhysRevD.105.023519](https://doi.org/10.1103/PhysRevD.105.023519)

I. INTRODUCTION

Over the past two decades, several efforts have been made to understand the problem of gravitational wave (GW) map-making. The term “map-making” [1,2] is typically used when considering stochastic GW backgrounds (SGWBs), which are expected to be incoherent superpositions of GW signals arriving at a detector from all directions of the sky. These GWs are generated by a variety of underlying mechanisms, which include both compact or diffuse sources at both astrophysical or cosmological distances (see, e.g., [3,4]).

The primary map-making method adopted in searches for anisotropy in stochastic signals is a maximum-

likelihood approach that targets the overall background intensity [5], entirely discarding any time-domain phase information present in the data. This method, hereinafter referred to as *intensity mapping*, has been used on cross-correlated data from the Laser Interferometer Gravitational-wave Observatory (LIGO) detectors [6–8] and, more recently, also on data from the Virgo observatory [9]. We review the basics of this method in Sec. III. Intensity mapping is best suited for a background that is truly stochastic. In this case the signal is characterized by waves that have random time-domain phases but whose intensity varies as a function of frequency and, in principle, angular direction. This motivates the development of methods that compress the data by discarding the time-domain phase information of the signal. For noise-dominated detectors,

^{*}arenzini@caltech.edu

such as LIGO-Virgo, this is achieved by integrating the equal-time cross-correlation of the time stream of pairs of detectors. Each detector pair has a characteristic correlated sky response to the GW signal, which is time-dependent following the Earth's rotation. The set of responses determine the mapping capabilities of the detector array.

Other methods have been proposed that solve for the phase information as well, effectively mapping the amplitude and phase—or real and imaginary parts—of the two GW polarization modes on the sky [10–12]. We refer to these as *amplitude-phase mapping* in what follows. A clear motivation for preserving phase information would be the case of a diffuse, anisotropic, but *coherent* GW background. It is often argued that the primordial background generated by an inflationary epoch satisfies this condition [13,14]. This is due to the squeezing of modes induced by any period of superhorizon evolution. It has been shown, however, that, at accessible frequencies, any primordial coherence is wiped out by the effect of large-scale structure on the GWs as they propagate through the evolving universe [15,16]. Hence, reconstructing the phase information itself is effectively not necessary in stochastic analyses; however, there remains a question of whether it is actually *possible* to determine the phase information for a stochastic source. This question is what we set out to answer with this paper.

In this paper we compare the application and validity of intensity and amplitude-phase mapping methods for SGWB analyses. We do this by analyzing mock data, focusing on the cross-correlation of measurements made by a network of detectors. It is important to note, at this stage, that a key element in any stochastic analysis is the accurate modeling and estimation of detector noise. We do not focus on this aspect here since our aim is to clarify the applicability of the mapping methods with respect to the signal properties. Our analysis assumes the simplest possible noise component and artificially high signal-to-noise ratios (SNRs). Our conclusions, nonetheless, are independent of these choices.

This paper is organized as follows: in Sec. II, we review the strain description of the GW signal and the key assumptions made throughout the paper. Section III details the two map-making methods. In Sec. IV we discuss different measurement scenarios to build an understanding of the issues encountered when reconstructing an anisotropic SGWB using intensity and amplitude-phase mapping methods. In Sec. V we apply the two methods to mock datasets and provide a comparison between the mapping fidelity. We conclude with Sec. VI, where we discuss our findings and future prospects.

II. GW SIGNAL

The stochastic GW metric perturbations at time t and position vector \mathbf{x} may be written as an infinite superposition of plane waves having polarization A , having frequency f ,

and arriving from direction $\hat{\mathbf{n}} = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)$ [5]:

$$h_{ij}(t, \mathbf{x}) = \int_{-\infty}^{+\infty} df \int_{S^2} d\hat{\mathbf{n}} \sum_{A=+, \times} h_A(f, \hat{\mathbf{n}}) e_{ij}^A(\hat{\mathbf{n}}) e^{-i2\pi f(\hat{\mathbf{n}} \cdot \mathbf{x} + t)}, \quad (1)$$

where (θ, ϕ) are the standard angular coordinates on the 2-sphere and the spatial wave vector is written explicitly as $\mathbf{k} = -2\pi f \hat{\mathbf{n}}$. We set the speed of light $c = 1$ here for simplicity. The metric perturbations $h_{ij}(t, \mathbf{x})$ are real; hence a reality condition is imposed on the complex, frequency-domain modes, $h_A^*(f, \hat{\mathbf{n}}) = h_A(-f, \hat{\mathbf{n}})$. In the above expression, we choose the linear polarization basis $A = \{+, \times\}$, where the orthogonal polarization basis tensors e^A may be written as

$$e^+ = \hat{\boldsymbol{\theta}} \otimes \hat{\boldsymbol{\theta}} - \hat{\boldsymbol{\phi}} \otimes \hat{\boldsymbol{\phi}}, \quad (2)$$

$$e^\times = \hat{\boldsymbol{\theta}} \otimes \hat{\boldsymbol{\phi}} + \hat{\boldsymbol{\phi}} \otimes \hat{\boldsymbol{\theta}}, \quad (3)$$

where

$$\hat{\boldsymbol{\theta}} = (\cos \theta \cos \phi, \cos \theta \sin \phi, -\sin \theta), \quad (4)$$

$$\hat{\boldsymbol{\phi}} = (-\sin \phi, \cos \phi, 0) \quad (5)$$

are the standard unit vectors tangent to the sphere. Since $\{h_+, h_\times\}$ are complex valued, we can write them in terms of either their real and imaginary components or their amplitude and temporal phase,

$$h_A(f, \hat{\mathbf{n}}) = h_A^{\text{R}}(f, \hat{\mathbf{n}}) + ih_A^{\text{I}}(f, \hat{\mathbf{n}}) = \mathcal{A}_A(f, \hat{\mathbf{n}}) e^{i\varphi_A(f, \hat{\mathbf{n}})}, \quad (6)$$

all of which, in general, will depend on the frequency and direction of the waves.

For a stochastic background, the metric perturbations $h_{ij}(t, \mathbf{x})$ and hence the Fourier components $h_A(f, \hat{\mathbf{n}})$ are *random fields*, whose probability distributions define the statistical properties of the background. For the following discussion, we will assume that the background is (i) *Gaussian*, (ii) *stationary*, and (iii) *unpolarized*, which means that (i) the statistical properties of the random fields are completely characterized by their first- and second-order moments, (ii) there is no preferred origin of time (implying that random variables corresponding to different frequencies are statistically independent of one another), and (iii) the statistical properties of the background are invariant under rotations of the polarization tensors in the plane perpendicular to $\hat{\mathbf{n}}$ (implying statistically independent and equivalent $+$ and \times polarization components). From these assumptions it follows that

$$\langle h_A(f, \hat{\mathbf{n}}) \rangle = 0, \quad (7)$$

that $\{h_A^{\text{R}}, h_A^{\text{I}}\}$ are statistically independent variants drawn from the same Gaussian distribution with zero mean and half the variance of h_A , and that the phase φ_A is uniformly distributed between 0 and 2π and is statistically independent of the amplitude \mathcal{A}_A , which is Rayleigh distributed being the square root of the sum of squares of two statistically independent Gaussian random variables each with zero mean and equal variance.

We will further assume (iv) *that the background has no nontrivial phase coherence across the sky*. This means that the phase of the GWB signal coming from two different directions on the sky are statistically independent of one another¹:

$$\text{Cov}[e^{i\varphi_A(f, \hat{n})}, e^{i\varphi_{A'}(f', \hat{n}')}] \propto \delta_{AA'} \delta(f - f') \delta(\hat{n}, \hat{n}'). \quad (8)$$

Although primordial backgrounds, such as those generated during an epoch of inflation, may be phase coherent as modes reenter the horizon and begin to propagate, this coherence is lost through propagation effects [16]. It is therefore reasonable to assume that phase incoherence is a generic feature of *any* stochastic background that may, or may not, have angular correlations in amplitude. In other words, the intrinsic angular scale $\Delta\hat{n}$ over which the phases are correlated goes to 0, even though the correlation scale for the amplitude may be finite. Hence, any attempt to measure the GW phase using a detector with finite angular resolution will necessarily *average* the true phase over this angular resolution scale. This loss of information will degrade reconstruction of the amplitude if one tries to estimate it from the real and imaginary parts of the Fourier components. We will show this explicitly in Secs. IV and V.

Putting together all of the above results, we can write

$$\langle h_A(f, \hat{n}) h_A^*(f', \hat{n}') \rangle = \frac{1}{2} I(f, \hat{n}) \delta_{AA'} \delta(f - f') \delta(\hat{n}, \hat{n}'), \quad (9)$$

where

$$I(f, \hat{n}) \equiv \frac{2}{T} \sum_A \langle \mathcal{A}_A^2(f, \hat{n}) \rangle \quad (10)$$

defines the intensity of the GWB as a function of frequency f and direction \hat{n} . Here T is the total observation time. The angle brackets $\langle \dots \rangle$ denote *ensemble* averaging over the random amplitudes and phases of the Fourier coefficients of the metric perturbations at a fixed spatial location \mathbf{x} , under the assumption of time stationarity. In practice, this averaging is realized by averaging over all the available GW time-series data, assuming that the background is *ergodic*. The distribution of energy and matter in the universe, e.g., large-scale structure, which gives rise to

the GW background is *fixed* with respect to this averaging process.

The GW background may be *isotropic*, *anisotropic*, or *statistically isotropic* (i.e., invariant under arbitrary rotations of the sky) depending on the statistical properties of $I(f, \hat{n})$ with respect to the sky direction \hat{n} . For example, for both isotropic and anisotropic backgrounds, the intensity field is a *deterministic* quantity; it is independent of sky direction for a purely isotropic background—i.e., $I(f, \hat{n}) \equiv I(f)$, and has preferred directions for an anisotropic background. For a statistically isotropic background, the intensity is a random field, assumed here to be approximately Gaussian, whose mean is independent of sky direction

$$\langle I(f, \hat{n}) \rangle_{\Omega} \equiv I(f), \quad (11)$$

and whose quadratic expectation values depend only on the angular separation between two points on the sky as

$$\begin{aligned} \langle I(f, \hat{n}) I^*(f, \hat{n}') \rangle_{\Omega} &\equiv C(f, \hat{n} \cdot \hat{n}') \\ &= \sum_{\ell=0}^{\infty} \frac{2\ell+1}{4\pi} C_{\ell}(f) P_{\ell}(\hat{n} \cdot \hat{n}'), \end{aligned} \quad (12)$$

where P_{ℓ} is the Legendre polynomial of order ℓ . In the above expressions, $\langle \dots \rangle_{\Omega}$ denotes averaging over different GW universes [17] (e.g., over different realizations of large-scale structure) which are drawn from a rotationally invariant probability distribution described by the *angular power spectrum* $C_{\ell}(f)$. We note that there is a subtle difference if one imposes statistical isotropy on the h_A fields; for more details regarding this see [11].

In all cases, the intensity $I(f, \hat{n})$ may be related to the fractional energy density parameter $\Omega_{\text{GW}}(f, \hat{n})$ via [5]

$$\Omega_{\text{GW}}(f, \hat{n}) = \frac{4\pi^2 f^3}{\rho_c G} I(f, \hat{n}), \quad (13)$$

which is the fundamental relation that allows one to connect GWB observations to the cosmological implications of the background. Integrating the above equation over direction on the sky yields

$$\Omega_{\text{GW}}(f) = \frac{4\pi^2 f^3}{\rho_c G} I(f), \quad (14)$$

which relates the monopole components of the fractional energy density parameter and the intensity of the background. Finally, we note that it is common in the literature to assume that the intensity factorizes as

$$I(f, \hat{n}) = E(f) I(\hat{n}), \quad (15)$$

where $E(f)$ and $I(\hat{n})$ encode the spectral and directional dependence of the background, respectively. We employ

¹We define the covariance of two complex variables A and B as $\text{Cov}(A, B) \equiv \langle AB^* \rangle - \langle A \rangle \langle B^* \rangle$.

this assumption throughout as it considerably simplifies map-making.

III. MAP-MAKING METHODS WITH GW DETECTOR DATA

In this section we present two proposed maximum-likelihood map-making methods for GWBs in a common formalism. First we review the intensity mapping method which has been previously presented in several references, e.g., [5,18,19], and applied consistently to LIGO data [7–9,20–22]; then, we lay out the amplitude-phase mapping method which follows the presentation in [12].

A. Intensity mapping

Intensity mapping, such as that used for map-making with the LIGO-Virgo detectors, works with the cross-correlation of time-coincident data directly to discard the autocorrelated noise terms which would otherwise dominate the calculation.

To start, we consider a set of detectors $i = \{1, \dots, N\}$. The data collected by detector i in a time segment τ may be considered as made up of separate signal and noise components, $d_i^\tau(t) = s_i^\tau(t) + n_i^\tau(t)$. We write the corresponding discrete Fourier transform as

$$\tilde{d}_i^\tau(f) = \tilde{s}_i^\tau(f) + \tilde{n}_i^\tau(f) \rightarrow \mathbf{d}^{\tau,f}, \quad (16)$$

where we have dropped the tilde and detector label in favor of the more concise boldface vector notation. As we are dealing with a real time-stream Fourier transformed into a discretized Fourier space, it is convenient to relabel frequencies as a discrete index, which picks out a single frequency mode in the transform. The signal component in the data is modeled as

$$\mathbf{s}^{\tau,f} = \int_{S^2} d\hat{\mathbf{n}} \sum_A \mathbf{R}_A^{\tau,f}(\hat{\mathbf{n}}) h_A^f(\hat{\mathbf{n}}), \quad (17)$$

where \mathbf{R} is the *response* function spanning detector space. Note that the h field has no τ dependence as it is assumed to be stationary. A pair of detectors i, j then observe $d_i^{\tau,f}, d_j^{\tau,f}$, respectively; these form a *baseline* and we can consider a correlated data vector as spanning the space of different baselines directly, $D_{ij}^{\tau,f} \equiv d_i^{\tau,f} d_j^{\tau,f*} \equiv \mathbf{D}^{\tau,f}$. We can write down the likelihood for the residuals of the cross-correlated data as

$$\mathcal{L} \propto \prod_{\tau,f,ij} \frac{1}{|\mathbf{C}_N|^{1/2}} e^{-\frac{1}{2}(\mathbf{D}-\mathbf{S})\mathbf{C}_N^{-1}(\mathbf{D}-\mathbf{S})^*}. \quad (18)$$

Note that we are using shorthand notation here, omitting τ and f everywhere in Eq. (18). As in [19], we take the signal model \mathbf{S} for the cross-correlation to be

$$\mathbf{S}^{\tau,f} = \int_{S^2} d\hat{\mathbf{n}} \mathbf{\Gamma}^{\tau,f}(\hat{\mathbf{n}}) I(\hat{\mathbf{n}}), \quad (19)$$

which may be derived directly from Eq. (9). Here $\mathbf{\Gamma}$ is the cross-correlated response vector to the GW intensity, obtained via the squared sum of the response terms above as

$$\mathbf{\Gamma}_{ij}^{\tau,f}(\hat{\mathbf{n}}) \equiv \sum_A \mathbf{R}_{Ai}^{\tau,f}(\hat{\mathbf{n}}) \mathbf{R}_{Aj}^{\tau,f*}(\hat{\mathbf{n}}). \quad (20)$$

The spectral dependence of I has been assumed factorizable here and absorbed into $\mathbf{\Gamma}$ for simplicity, assuming Eq. (15). The time dependence (encapsulated in τ) and spatial dependence of the response function are fundamental to map-making, as they define the scan strategy of the set of baselines, which sets the resolution of the measurement. More details on this may be found in [1,23–26].

The noise covariance matrix may be written in baseline space, $N_{ij}^{\tau,f} \equiv n_i^{\tau,f} n_j^{\tau,f*} \equiv \mathbf{N}^{\tau,f}$; hence the noise covariance matrix is

$$\mathbf{C}_N^{\tau,f} = \langle \mathbf{N}^{\tau,f} \otimes \mathbf{N}^{\tau,f*} \rangle \equiv \text{diag}(\mathbf{P}_i^{\tau,f} \mathbf{P}_j^{\tau,f*}), \quad (21)$$

assuming uncorrelated noise between detectors, where \mathbf{P}_i is the (two-sided) noise power spectrum in detector i . Maximizing the likelihood (18) above yields the mapping equation for GW intensity

$$I(\hat{\mathbf{n}}) = \mathbf{M}(\hat{\mathbf{n}}, \hat{\mathbf{n}}')^{-1} z(\hat{\mathbf{n}}'). \quad (22)$$

We refer to z as the *projection* of the dataset into pixel space, while \mathbf{M} is the *Fisher matrix* of the mapping problem. These are constructed from the following quantities calculated at the individual f and τ , as each of these constitutes an independent measurement according to the likelihood in Eq. (18),

$$\begin{aligned} z^{\tau,f}(\hat{\mathbf{n}}) &= \mathbf{\Gamma}^{\tau,f}(\hat{\mathbf{n}}) (\mathbf{C}_N^{\tau,f})^{-1} \mathbf{D}^{\tau,f}, \\ \mathbf{M}^{\tau,f}(\hat{\mathbf{n}}, \hat{\mathbf{n}}') &= \mathbf{\Gamma}^{\tau,f}(\hat{\mathbf{n}}) (\mathbf{C}_N^{\tau,f})^{-1} \mathbf{\Gamma}^{\tau,f}(\hat{\mathbf{n}}'), \end{aligned} \quad (23)$$

which are then summed over all times, frequencies, and baselines ij in the set to obtain the full (and most informative) projection and Fisher matrix,

$$z(\hat{\mathbf{n}}) = \sum_{\tau,f,ij} z^{\tau,f}(\hat{\mathbf{n}}), \quad \mathbf{M}(\hat{\mathbf{n}}, \hat{\mathbf{n}}') = \sum_{\tau,f,ij} \mathbf{M}^{\tau,f}(\hat{\mathbf{n}}, \hat{\mathbf{n}}'). \quad (24)$$

Note that the integration over frequencies here requires an assumption for the frequency dependence of the GWB, which we have transferred to the response function. One could also choose to *not* integrate over frequencies, and solve for maps mode by mode, avoiding this frequency modeling step.

At this stage, it is necessary to pick a working resolution on the sky in order to explicitly carry out the calculations. Using the HEAPLIX PYTHON package to deal with pixelization, we can set a working resolution and translate sky direction \hat{n} to pixel p . The data model then becomes

$$s^{\tau,f} = \frac{4\pi}{N_{\text{pix}}} \sum_p \Gamma_p^{\tau,f} I_p. \quad (25)$$

The components of Eq. (23) can then be considered in the pixel domain, where $\hat{n} \rightarrow p$.

B. Amplitude-phase mapping

Let us now take a step back and reconsider the data model in Eq. (16), which we can write compactly as $\mathbf{d}^{\tau,f} = \mathbf{s}^{\tau,f} + \mathbf{n}^{\tau,f}$. Assuming zero-mean Gaussian noise in the strain, i.e., $\langle \mathbf{n}^{\tau,f} \rangle = 0$, Eq. (17) serves as our signal model. As $\{h_+, h_\times\}$ are complex valued, there are four independent (real) fields on the sky to estimate: $\{h_+^{\text{R}}, h_\times^{\text{R}}, h_+^{\text{I}}, h_\times^{\text{I}}\}$. We write down the likelihood in terms of the residuals in the strain as

$$\mathcal{L} \propto \prod_{\tau,f,i,j} \frac{1}{|C_n|^{1/2}} e^{-\frac{1}{2}(\mathbf{d}-s)^\dagger (C_n)^{-1} (\mathbf{d}-s)}, \quad (26)$$

assuming each time segment τ and frequency f in the dataset are statistically independent of one another, where we are again employing shorthand notation, using boldface font to represent a vector that now spans just the space of the detectors (and not baselines). C_n is the noise covariance $C_n = \mathbf{n} \otimes \mathbf{n}^*$. As the noise is modeled independently and is not part of the maximum likelihood estimation, it is possible to reduce the mapping problem to a closed-form χ^2 solution minimizing

$$\chi^2 = -\frac{1}{2} \sum_{\tau,f,i,j} (\mathbf{d}-s)^\dagger (C_n)^{-1} (\mathbf{d}-s). \quad (27)$$

We solve for the four h fields separately; hence it is useful to decompose the signal model into real and imaginary components simply as

$$s^{\tau,f} = \int_{S^2} d\hat{n} \sum_A \mathbf{R}_A^{\tau,f}(\hat{n}) [h_A^{\text{R}}(\hat{n}) + ih_A^{\text{I}}(\hat{n})], \quad (28)$$

where the spectral dependence of the stochastic field has been explicitly factored out of $h_A(\hat{n})$ and has been absorbed into the response term \mathbf{R} so as to keep track of fewer dependencies. Minimizing χ^2 with respect to each field yields the maximum likelihood solution

$$\begin{pmatrix} h_{A'}^{\text{R}} \\ h_{A'}^{\text{I}} \end{pmatrix}_{\hat{n}} = \begin{pmatrix} M_{AA'}^{\text{R}} & -M_{AA'}^{\text{I}} \\ M_{AA'}^{\text{I}} & M_{AA'}^{\text{R}} \end{pmatrix}_{\hat{n},\hat{n}'}^{-1} \begin{pmatrix} z_{A'}^{\text{R}} \\ z_{A'}^{\text{I}} \end{pmatrix}_{\hat{n}'}, \quad (29)$$

where

$$\begin{aligned} \begin{pmatrix} z_{A'}^{\text{R}} \\ z_{A'}^{\text{I}} \end{pmatrix}_{\hat{n}} &= \sum_{\tau,f,i,j} \begin{pmatrix} z_{A'}^{\text{R}} \\ z_{A'}^{\text{I}} \end{pmatrix}_{\hat{n}}^{\tau,f}, & \begin{pmatrix} M_{AA'}^{\text{R}} & -M_{AA'}^{\text{I}} \\ M_{AA'}^{\text{I}} & M_{AA'}^{\text{R}} \end{pmatrix}_{\hat{n},\hat{n}'} \\ &= \sum_{\tau,f,i,j} \begin{pmatrix} M_{AA'}^{\text{R}} & -M_{AA'}^{\text{I}} \\ M_{AA'}^{\text{I}} & M_{AA'}^{\text{R}} \end{pmatrix}_{\hat{n},\hat{n}'}^{\tau,f}. \end{aligned} \quad (30)$$

Each component is, explicitly,

$$\begin{aligned} z_A^{\text{R}}(\hat{n}) &= \mathbf{R}_A^{\text{R}}(\hat{n}) C_n^{-1} \mathbf{d}^{\text{R}} + \mathbf{R}_A^{\text{I}}(\hat{n}) C_n^{-1} \mathbf{d}^{\text{I}}, \\ z_A^{\text{I}}(\hat{n}) &= \mathbf{R}_A^{\text{R}}(\hat{n}) C_n^{-1} \mathbf{d}^{\text{I}} - \mathbf{R}_A^{\text{I}}(\hat{n}) C_n^{-1} \mathbf{d}^{\text{R}}, \\ M_{AA'}^{\text{R}}(\hat{n}, \hat{n}') &= \mathbf{R}_A^{\text{R}}(\hat{n}) C_n^{-1} \mathbf{R}_{A'}^{\text{R}}(\hat{n}') + \mathbf{R}_A^{\text{I}}(\hat{n}) C_n^{-1} \mathbf{R}_{A'}^{\text{I}}(\hat{n}'), \\ M_{AA'}^{\text{I}}(\hat{n}, \hat{n}') &= \mathbf{R}_A^{\text{R}}(\hat{n}) C_n^{-1} \mathbf{R}_{A'}^{\text{I}}(\hat{n}') - \mathbf{R}_A^{\text{I}}(\hat{n}) C_n^{-1} \mathbf{R}_{A'}^{\text{R}}(\hat{n}'). \end{aligned} \quad (31)$$

For the sake of conciseness we do not include the frequency and time labels in each of the terms above; these mirror those in Eq. (23). Note that the projection z and Fisher matrix M here are not the same as for intensity mapping—however, we have chosen to keep the same notation to draw the comparison between the two approaches. In the uncorrelated noise case the (noise) covariance matrix becomes diagonal in detector space,

$$(C_n)_{ij} = \delta_{ij} P_j, \quad (32)$$

so calculations in Eq. (31) simplify considerably. Furthermore, as in the intensity mapping approach described above, the integration over frequencies in Eq. (30) requires an assumption for the spectral shape of the signal. However, in the amplitude-phase case one must take extra care: when broadband integrates here one must assume a certain phase-coherence across modes, or else the phase introduced in Eq. (6) will (correctly) average to zero and yield null sky maps.

Finally, as in the previous section, we can discretize the sky and transform sky direction \hat{n} to pixel p , to apply this method to data. The signal model becomes

$$s^{\tau,f} = \frac{4\pi}{N_{\text{pix}}} \sum_p \sum_A \mathbf{R}_{A,p}^{\tau,f} h_{A,p}, \quad (33)$$

and similarly the terms in Eq. (31) may be translated into the pixel domain.

IV. INSTRUCTIVE CONSIDERATIONS

Before discussing the application of the map-making methods laid out above, let us start with some useful considerations about the nature of the measurement of broadband, stochastic GWs from all sky directions with a set of interferometers. In particular, let us focus on the notion that a detector naturally low-pass filters the signal,

both in terms of temporal resolution and angular resolution; the subsequent digitization/sampling of the data needs to respect the maximum temporal and spatial frequencies present in the signal if aliasing of power is to be avoided.

It is useful to first break down the measurement into two components: the time or frequency-domain measurement, and the sky-dependent response. These two steps are inextricable in a real GW detector. However, to get an idea of the accessibility to each information component we can consider two simple scenarios separately: first, we can focus on the estimation of the amplitude and phase of a complex field in the frequency domain through time-domain measurements; then we can analyze the effect of a window-averaged response on the measurements of a complex field on the semicircle, in analogy with what happens with GW detectors observing the full sky.

When analyzing GW detector time streams that contain a measurement of a stationary stochastic field, say $h(t)$, we typically start by taking the Fourier transform (FT) of the data, as the frequency domain representation allows us to access information more efficiently. A way to see this is to note that the quadratic expectation value in the time and frequency domains are related by

$$\langle h(t)h(t') \rangle = C(t-t') \overset{\text{FT}}{\leftrightarrow} \langle \tilde{h}(f)\tilde{h}^*(f') \rangle = P(f)\delta(f-f'). \quad (34)$$

Thus, the Fourier transform of a stationary time series maps our measurement to the diagonal space for the autocorrelation of the field, where $P(f)$ is the (two-sided) power spectrum of the field. In other words, while in the time domain the correlation between the stationary field at different times depends (only) on the time difference $t-t'$, in the *conjugate space* (frequency space) the correlation between different frequency components is a delta function, and thus the frequency dependence is totally “compactified” in the power spectrum. We can either write an estimator for a real field, for example, $P(f)$, if we are interested only in the intensity of the signal, or we can aim to measure the complex field $\tilde{h}(f)$, to include both amplitude and phase information. In either case, we model the signal in the frequency domain and estimate it through time-domain measurements. Thus, in practice the measurement of the field is performed in the conjugate space of the model, over an observation time T , sampled at the finite time resolution Δt_d . According to the Nyquist-Shannon sampling theorem, given the Nyquist frequency of the measurement $f_{\text{Nyq}} \equiv 1/(2\Delta t_d)$, the field $\tilde{h}(f)$ may be completely determined as long as it is made up of modes with frequency content $1/\Delta t < f_{\text{Nyq}}$, where Δt is the time-coherence scale of the field. Hence, if the GWs in the time stream are at frequencies below f_{Nyq} , it is theoretically possible to fully reconstruct the Fourier coefficients $h_A(f)$, in the case of an isotropic background. If the target signal is

stochastic in nature, and the phase information is simply uninteresting, then the choice of method boils down to what is most computationally effective when dealing with noisy time streams and multiple detectors.

To understand the effects of the integrated beam of GW detectors, let us start by considering a complex and statistically isotropic field $\tilde{h}(\hat{n})$ on the sphere,

$$\tilde{h}(\hat{n}) = \mathcal{A}(\hat{n})e^{i\varphi(\hat{n})}, \quad (35)$$

where the phase φ has a certain functional dependence on direction. This is effectively a frequency-independent, unpolarized version of the GW strain in the Fourier domain given by Eq. (6). Note that, as in the relation $t \overset{\text{FT}}{\leftrightarrow} f$, there is a useful conjugate space to direction \hat{n} space,

$$\langle \tilde{h}(\hat{n})\tilde{h}^*(\hat{n}') \rangle = C(\hat{n} \cdot \hat{n}') \overset{\text{SHD}}{\leftrightarrow} \langle a_{\ell m'} a_{\ell' m'} \rangle = C_\ell \delta_{\ell\ell'} \delta_{mm'}, \quad (36)$$

where $a_{\ell m}$ are the spherical harmonic coefficients of $\tilde{h}(\hat{n})$ and C_ℓ is the angular power spectrum of \tilde{h} . The C_ℓ 's are related to $C(\hat{n} \cdot \hat{n}')$ via

$$C(\hat{n} \cdot \hat{n}') = \sum_{\ell=0}^{\infty} \frac{2\ell+1}{4\pi} C_\ell P_\ell(\hat{n} \cdot \hat{n}'), \quad (37)$$

where $P_\ell(x)$ are the Legendre polynomials. The spherical harmonic decomposition (SHD) introduced in Eq. (36) allows us to transform the field from directional space \hat{n} to degree ℓ and order m space, and, again, the correlation is diagonal in the latter. The measurement of \tilde{h} , however, does not occur in \hat{n} space directly, but in the time domain; the relation between these two spaces depends on the specific detector used to make the measurement. This is the crucial difference between measurements of a frequency-dependent GW observable and a direction-dependent one.

Let us now discuss an instructive example. Let us consider a case where we can fix the observation to a single value of ϕ on the sphere, such that we reduce the mapping problem to the estimation of a one-dimensional field on the semicircle. \tilde{h} is then parametrized solely by $0 \leq \theta \leq \pi$ on the semicircle, such that effectively we are observing $\tilde{h}(\theta, \phi = \text{const}) \equiv \tilde{h}(\theta) = \mathcal{A}(\theta)e^{i\varphi(\theta)}$. We assume that our detector makes a time-dependent measurement $r(t)$ of the field on the semicircle,

$$r(t) = \int_0^\pi d\theta R(t, \theta) \tilde{h}(\theta), \quad (38)$$

filtered by the detector response $R(t, \theta)$, which scans the semicircle as time goes by. Note here the similarity with the model for the GW strain in Eq. (17), relevant to the measurement of anisotropic GWBs.

The intrinsic coherence scale of the field, $\Delta\theta$, sets the input resolution of our example, inducing $\pi/\Delta\theta$ independent samples on the semicircle. We take the phase $\varphi(\theta)$ to be random at the intrinsic scale, such that the field has statistically independent phases at each sample. The amplitude \mathcal{A} is assumed to vary much more slowly than the phase, such that the scale over which the amplitude varies appreciably is much larger than $\Delta\theta$. We model the instantaneous response of the detector at time t_0 as a top-hat function with width $\Delta\theta_d$,

$$R(t_0, \theta) = \begin{cases} 1, & -\frac{\Delta\theta_d}{2} < \theta(t_0) < \frac{\Delta\theta_d}{2} \\ 0, & \text{otherwise} \end{cases}, \quad (39)$$

and the time dependence is imposed by the scanning strategy, i.e., the angular function $\theta(t)$. In this simple example we take $\theta(t) = \pi t/\tau$ where τ is the time period over the semicircle, such that consecutive measurements correspond to consecutive angular samples on the semicircle, and continuity between measurements at 0 and π is ensured. The measurement of \tilde{h} is then mediated by the detector-induced window width $\Delta\theta_d$, and the comparison between $\Delta\theta$ and $\Delta\theta_d$ will determine to what degree the field is resolvable. $\Delta\theta_d$ is the *angular resolution* of the measurement, and as long as $\Delta\theta_d$ is sufficiently less than $\Delta\theta$, the field is overresolved and its amplitude and phase information may be measured exactly. However, if $\Delta\theta_d > \Delta\theta$, there is loss of information as the window averages over the phases, which are randomly distributed along the semicircle. In the extreme case where $\Delta\theta_d \gg \Delta\theta$, the measurement becomes compatible with 0, and the amplitude of the field is entirely lost. Translating the measurement into the conjugate space, one may interpret this limit as an *angular* Nyquist frequency, $\ell_{\text{Nyq}} \equiv \pi/(2\Delta\theta_d)$, imposed by the detector. In this sense, the measurement only works when $\ell_{\text{Nyq}} > \ell_{\text{max}}$, where $\ell_{\text{max}} \sim \pi/\Delta\theta$ is the maximum angular frequency of the signal. Note the analogy then between $\Delta\theta_d$ and the sampling rate Δt_d described above. The useful estimator for this example is one that marginalizes over the phases and aims for the field intensity $C(\theta)$ directly, which is

$$\langle \tilde{h}(\theta)\tilde{h}(\theta') \rangle = C(\cos(\theta))\delta(\theta - \theta'), \quad (40)$$

obtained from Eq. (36) assuming total phase incoherence of $\tilde{h}(\theta)$. Equal-time measurements $r_1(t)$ and $r_2(t)$ made by two identical detectors 1 and 2 then satisfy

$$\langle r_1(t)r_2^*(t) \rangle = \int_0^\pi d\theta R_1(t, \theta)R_2^*(t, \theta)C(\cos(\theta)), \quad (41)$$

where the angle brackets introduce an expectation value similar to Eq. (9); i.e., they imply ensemble averaging over different data samples, and the Dirac delta in Eq. (40) has already been applied. Mapping then amounts to inverting

Eq. (41) to estimate $C(\theta)$. Hence to preserve the amplitude information it is necessary to cross-correlate (i.e., square) signals before averaging over the response window. Note that the spatial averaging introduced in Eq. (38) is analogous to what happens in a GW detector: the directional information is not directly accessible in the measurement, but rather needs to be reconstructed based on observation features, such as the time dependence of R .

In the simple case above the only relevant coherence scale of the field is that of the phase; however, in general it is also necessary to worry about the coherence scale of the intensity, $\Delta\theta_C$. Note this is directly related to the coherence scale of the amplitude, as in this case $C = \mathcal{A}^2$. When $\Delta\theta_C < \Delta\theta_d$, this will have an impact on the observation of C (or, equivalently, of \mathcal{A} , taking the square root of the intensity) similar to what occurs for φ . However, note that as the intensity is positive definite, this field will never average to zero but would rather approach zero from above in the limit of $\Delta\theta_C \ll \Delta\theta_d$.

V. MOCK-DATA MAPPING TESTS

Simplified mock-data mapping tests are presented here to illustrate the points made in the previous sections. We consider the two LIGO interferometers that have a well-known response function, illustrated, for example, in [18]. In these examples, the noise is taken to be white, Gaussian, stationary, and much weaker than the signal. This is to highlight the potential and shortcomings of the mapping algorithms based on the characteristics of the signal component only. As the noise is assumed to be known, it is not part of the estimation. In practice, all GW mapping attempts to date using ground-based detectors have relied on independent estimators for the noise, which is a valid approach when the signal is entirely subdominant and hence does not bias the noise estimation, at least on a segment-by-segment basis. This methodology needs to be revisited in the presence of competing signal and noise components—an example of this may be seen in [27] for the Laser Interferometer Space Antenna.

The reconstructions presented here are obtained following the data-handling recipe of [19], which imitates the steps performed on real LIGO data [28,29]. In both intensity and amplitude-phase mapping, the mock data are generated in segments that represent Fourier transformed one-minute segments of a detector array time stream. The one-minute timescale is chosen such that the response of the detectors may be considered constant on the sky throughout the segment. This not only sets a natural lower bound on the frequency range probed but also imposes a pixelization scheme: in order for the reconstruction to be complete, the sky response should vary smoothly from segment to segment.

The most challenging part of this procedure is the safe inversion of the Fisher matrix. This arises since the spatial sampling of the sky is suboptimal, as explored, for




FIG. 1. Set of input sky maps used in the mock-data tests described in the text. The top row shows injected maps for the four components of the h field, used for data injection and reconstruction with the amplitude-phase algorithm. The pixel phase is chosen randomly pixel to pixel. The bottom map corresponds to the total injected intensity I_{in} resulting from the combination of the maps in the top row. All these maps are injected at $N_{\text{side}}^{\text{in}} = 16$.

example, in [19]. However, as long as the noise is subdominant it is possible to pseudo-invert M discarding a minimal number of modes and recovering the sky signal perfectly, as shown below. The situation capsizes in the presence of noise: in that case, it is necessary to choose a cutoff on the singular values of the eigenmodes of M , to differentiate between the signal and noise modes. This is a highly nontrivial problem, so we do not discuss it here—see the discussion in [9] for example.

In all mapping examples presented here, the signal is modeled as a diffuse point source on the sky, with varying phase as a function of direction. This is obtained in practice by associating a random phase to each pixel of the amplitude sky map, which is then scanned to prepare the data segments. This is explicitly performed via Eq. (33) in the case of amplitude-phase mapping and Eq. (19) in the case of intensity mapping. The pixelization schemes used for data generation and map reconstruction need not match—in fact, one may not in general expect that the natural coherence scale of the field be comparable to that of the detector. The size of the unit pixel at injection, $4\pi/N_{\text{pix}}^{\text{in}}$, is analogous to $\Delta\theta$ in the example made in Sec. IV, while the size of the unit pixel at reconstruction, $4\pi/N_{\text{pix}}^{\text{out}}$, is analogous to $\Delta\theta_d$. For simplicity, we do not add a frequency-dependent phase term, although in reality it is present and may not be neglected in a real GW measurement; we discuss this further below.

Figure 2 shows results obtained with the amplitude-phase mapping algorithm, using the input maps from Fig. 1.

The top row of Fig. 2 shows the case when the input and output resolutions match exactly, $N_{\text{side}}^{\text{in}} \equiv N_{\text{side}}^{\text{out}} = 16$; hence the phases of the h_A fields in each pixel are perfectly recovered. However, the second and third rows of Fig. 2 show two cases where the output resolution is lowered from the input $N_{\text{side}}^{\text{in}} = 16$ to $N_{\text{side}}^{\text{out}} = 4$ and $N_{\text{side}}^{\text{out}} = 2$, respectively; here, the phase components are averaged out within neighboring pixels; hence the phase information is lost. This leads to an overall loss of intensity information as well, by the same reasoning described for the example in Sec. IV. Note that style choices have been made to underline the effect: the values chosen for the color bars in the intensity I plots in Fig. 2 match the correct values one would obtain when coarse graining the I_{in} intensity map to the relevant $N_{\text{side}}^{\text{out}}$. The coarse-graining exercise is performed with the `ud_grade` function of the HEALPY package [30].

This effect does not occur in the case of intensity mapping, since the phase information is not required to reconstruct the intensity on the sky. We have repeated the exercise above using the same input map as in Fig. 1 employing the intensity mapping algorithm, varying the values of $N_{\text{side}}^{\text{out}}$, and we find that the I_{out} maps agree with the `ud_graded` input maps within less than 1%. To quantitatively compare intensity and amplitude-phase mapping, we show in Fig. 3 the difference between the recovered maps with intensity mapping, $I_{\text{out},I\text{-map}}$, and amplitude-phase mapping, $I_{\text{out},\text{AP-map}}$, in the cases where $N_{\text{side}}^{\text{out}} < N_{\text{side}}^{\text{in}}$. In both cases, it is clear that most of the




FIG. 2. Set of sky map reconstruction examples with the amplitude-phase algorithm, using maps shown in Fig. 1 as input. The top row shows reconstructed maps with $N_{\text{side}}^{\text{out}} = 16$, which corresponds to the case when the intrinsic coherence scale of the field is equal to that of the detector. The second and third rows show reconstructions with $N_{\text{side}}^{\text{out}} = 4$ and $N_{\text{side}}^{\text{out}} = 2$, respectively. The bottom row shows the corresponding recovered intensity maps for the three different N_{side} cases. In the case when $N_{\text{side}}^{\text{out}} = N_{\text{side}}^{\text{in}}$, the recovered intensity matches exactly the one shown in Fig. 1; in the other cases, the intensity is degraded.

overall power is lost when performing amplitude-phase mapping, and in particular in the case where $N_{\text{side}}^{\text{out}} = 2$, the injected point source is hardly discernible from the fluctuations around 0 in the pixels.

Hence, we confirm that in the case of intensity mapping, the overall GW intensity is conserved, and the structure on the sky as well, as long as the coherence scale of the intensity is larger than that of the detector. In Fig. 4, we show an example where this is not the case: here, the injected signal is a very high resolution input map, $N_{\text{side}}^{\text{in}} = 32$, which is null everywhere apart from a tiny patch, mimicking a pointlike source, and is recovered at $N_{\text{side}}^{\text{out}} = 8$. Here indeed the signal from the (almost unique) incoming direction is averaged within the neighboring pixels, and hence the point value in that specific direction is much lower than the injection (by a factor of ~ 10). Note,

however, that the recovery is perfectly in line with a coarse graining of the input map: the monopole is conserved in the operation, and the correct value is recovered in the single “hot” pixel. Hence, in this case there is as little loss of information as possible. But this may present an issue when the signal is competing with a high level of noise.

Adding frequency-dependent phase terms to the injected maps has no effect on the intensity mapping results, as the data are first cross-correlated in the frequency domain and the phase term cancels out. However, in the amplitude-phase case this adds another level of complexity, as we need to keep track of the individual mode phases as well as those in each pixel. In fact, each frequency bin is independent, given the assumption of time stationarity, hence band-integrating over frequency bins would then have a similar effect as attempting map-making with a detector resolution




FIG. 3. A comparison between the recoveries of the same injected map I_{in} with the intensity I (first and third rows) and amplitude-phase AP (second and fourth rows) mapping algorithms, with the two different values of N_{side}^{out} as shown in Fig. 2. On the right, the difference maps $I_{in} - I_{out}$ are shown to quantitatively compare the performance of the algorithms. Note that the input map has been `ud_graded` to allow for this comparison; the actual input map is at $N_{side} = 16$ as shown in Fig. 1.




FIG. 4. Sky map injection (left panel) and reconstruction (right panel) with the intensity mapping algorithm, in a case where the intrinsic coherence scale of the signal intensity is considerably smaller than that of the detector.

that is larger than the intrinsic coherence scale of the field: the signal is degraded as the recovered amplitude averages to zero with the number of bins.

VI. DISCUSSION

We have taken a pedagogical approach in discussing the advantages and shortcomings of two different maximum-likelihood map-making techniques—intensity mapping and amplitude-phase mapping. The intent was to demystify these techniques and clarify the potential of detection and mapping of GWBs in general. Intensity mapping targets the GWB intensity as a function of direction. It constrains the average GW power on the sky and any anisotropies about this average. Versions of the method presented here have been applied to LIGO-Virgo data and have been proposed for LISA [1,2,12,31]. This method is best suited to astrophysical stochastic backgrounds, which carry no phase information, and hence solving for it would unnecessarily complicate the mapping procedure. On the other hand, amplitude-phase mapping allows one to do just that; however, we have shown that it is only possible to resolve the phase component if the intrinsic angular coherence scale of the signal is comparable to or larger than the angular resolution of the detector. When this is not the case, attempting to estimate the phase on the sky leads to loss of information, including the amplitude of the signal. Hence, this method is to be avoided in the case of stochastic backgrounds, where the intrinsic coherence scale is usually zero.

A possible exception to this rule may be the case of a stochastic background dominated by very few sources, such as that considered in [10]. This is a possibility for the stochastic background probed by pulsar timing arrays (PTAs), as the signal should be dominated by $\sim 10^2 - 10^3$ sources on the sky, such that a single source should dominate a pixel of $\sim 40 \text{ deg}^2$, which may be taken as the intrinsic coherence scale of the background. Even in this case, we are not quite there yet: achieving such an angular resolution with $\text{SNR} \sim 3$ would require a pulsar array with almost 10^4 pulsars [11], while current PTAs are monitoring $\sim 10^2$ pulsars, setting the present resolution to $\sim 400 \text{ deg}^2$.

Hence, this remains a target for future observatories, such as the Square Kilometre Array.

We have provided a useful investigation into the relationship between the time-frequency and the pixel-angular frequency domains. Time samples t are conjugate to frequencies f , while pixels p which correspond to sky locations \hat{n} are conjugate to the angular “frequency” scale l . However, our sky observations are effectively carried out in the time domain. A detector naturally low-pass filters the data, in terms of both temporal resolution and angular resolution; however, while the temporal resolution is well-known and controlled by the experimenter, the angular resolution is determined by the geometry and motion of the detector array. Furthermore, the digitization and sampling of the data needs to respect the maximum temporal and spatial frequencies, if aliasing of power is to be avoided.

The first detection of a GWB is drawing near, whether it be with PTAs, possibly by the NANOGrav Collaboration [32], or with the LIGO-Virgo-Kagra ground-based interferometer network. Currently, searches are focusing on the isotropic background, as astrophysical backgrounds are not expected to be highly anisotropic [17,33]. However, understanding how to handle any anisotropy to confirm or rule out this hypothesis remains an essential task of our collaborations. Ideally, the isotropic GWB could be estimated as the monopole of GWB maps (as in Ref. [21], for example), and not under the strict assumption of zero anisotropy.

ACKNOWLEDGMENTS

A. I. R. acknowledges the support of the National Science Foundation and the LIGO Laboratory. J. D. R. acknowledges support from NASA Grant No. 80NSSC19K0318, NSF Physics Frontiers Center Grants No. PFC-1430284 and No. PFC-2020265, and start-up funds from Texas Tech University. C. R. C. acknowledges support by Science and Technology Facilities Council consolidated Grant No. ST/P000762/1. N. J. C. appreciates the support of NASA LISA foundation Science Grant No. 80NSSC19K0320, NSF Grant No. PHY1912053, and NSF Physics Frontiers Center Grants No. PFC-1430284 and No. PFC-2020265.

[1] N. J. Cornish, *Classical Quantum Gravity* **18**, 4277 (2001).
 [2] N. J. Cornish, *Classical Quantum Gravity* **19**, 1279 (2002).
 [3] T. Regimbau, *Res. Astron. Astrophys.* **11**, 369 (2011).
 [4] C. Caprini and D. G. Figueroa, *Classical Quantum Gravity* **35**, 163001 (2018).
 [5] B. Allen and A. C. Ottewill, *Phys. Rev. D* **56**, 545 (1997).

[6] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. Lett.* **107**, 271102 (2011).
 [7] B. P. Abbott, R. Abbott, T. D. Abbott *et al.* (LIGO Scientific, Virgo Collaborations), *Phys. Rev. Lett.* **118**, 121102 (2017).
 [8] B. P. Abbott *et al.* (LIGO Scientific, Virgo Collaborations), *Phys. Rev. D* **100**, 062001 (2019).

- [9] R. Abbott *et al.* (LIGO Scientific, Virgo, KAGRA Collaborations), *Phys. Rev. D* **104**, 022005 (2021).
- [10] N. J. Cornish and R. van Haasteren, [arXiv:1406.4511](https://arxiv.org/abs/1406.4511).
- [11] J. Gair, J. D. Romano, S. Taylor, and C. M. F. Mingarelli, *Phys. Rev. D* **90**, 082001 (2014),
- [12] J. D. Romano, S. R. Taylor, N. J. Cornish, J. Gair, C. M. F. Mingarelli, and R. van Haasteren, *Phys. Rev. D* **92**, 042003 (2015),
- [13] L. Grishchuk, *Sov. Phys. JETP* **40**, 409 (1975).
- [14] C. R. Contaldi and J. A. Magueijo, *Phys. Rev. D* **98**, 043523 (2018).
- [15] N. Bartolo, V. Domcke, D. G. Figueroa, J. García-Bellido, M. Peloso, M. Pieroni, A. Ricciardone, M. Sakellariadou, L. Sorbo, and G. Tasinato, *J. Cosmol. Astropart. Phys.* **11** (2018) 034.
- [16] A. Margalit, C. R. Contaldi, and M. Pieroni, *Phys. Rev. D* **102**, 083506 (2020).
- [17] A. C. Jenkins, J. D. Romano, and M. Sakellariadou, *Phys. Rev. D* **100**, 083501 (2019).
- [18] J. D. Romano and N. J. Cornish, *Living Rev. Relativity* **20**, 2 (2017), ISSN 14338351.
- [19] A. I. Renzini and C. R. Contaldi, *Mon. Not. R. Astron. Soc.* **481**, 4650 (2018).
- [20] A. I. Renzini and C. R. Contaldi, *Phys. Rev. Lett.* **122**, 081102 (2019).
- [21] A. I. Renzini and C. R. Contaldi, *Phys. Rev. D* **100**, 063527 (2019).
- [22] B. P. Abbott, R. Abbott, T. D. Abbott *et al.* (LIGO Scientific, Virgo Collaborations), *Phys. Rev. Lett.* **118**, 121101 (2017); **119**, 029901(E) (2017).
- [23] A. Taruya and H. Kudoh, *Phys. Rev. D* **72**, 104015 (2005).
- [24] S. W. Ballmer, *Classical Quantum Gravity* **23**, S179 (2006), ISSN 0264-9381.
- [25] A. Taruya, *Phys. Rev. D* **74**, 104022 (2006).
- [26] S. Mitra, S. Dhurandhar, T. Souradeep, A. Lazzarini, V. Mandic, S. Bose, and S. Ballmer, *Phys. Rev. D* **77**, 042002 (2008).
- [27] C. R. Contaldi, M. Pieroni, A. I. Renzini, G. Cusin, N. Karnesis, M. Peloso, A. Ricciardone, and G. Tasinato (LISA Cosmology Working Group), *Phys. Rev. D* **102**, 043502 (2020).
- [28] B. P. Abbott, R. Abbott, T. D. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J.* **659**, 918 (2007).
- [29] B. P. Abbott, R. Abbott, T. D. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Nature (London)* **460**, 990 EP (2009).
- [30] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, *Astrophys. J.* **622**, 759 (2005).
- [31] S. Banagiri, A. Criswell, T. Kuan, V. Mandic, J. D. Romano, and S. R. Taylor, *Mon. Not. R. Astron. Soc.* **507**, 5451 (2021).
- [32] Z. Arzoumanian *et al.* (NANOGrav Collaboration), *Astrophys. J. Lett.* **905**, L34 (2020).
- [33] G. Cusin, C. Pitrou, and J.-P. Uzan, *Phys. Rev. D* **96**, 103019 (2017).