Maximum likelihood map-making with the Laser Interferometer Space Antenna

Carlo R. Contaldi,[1] Mauro Pieroni,[1] Arianna I. Renzini,[1, *] Giulia Cusin,[2, 3] Nikos
Karnesis,[4] Marco Peloso,[5] Angelo Ricciardone,[5] and Gianmassimo Tasinato[6]

(For the LISA Cosmology Working Group)

[1]*Blackett Laboratory, Imperial College London, SW7 2AZ, UK*
[2]*Astrophysics Department, University of Oxford, DWB, Keble Road, Oxford OX1 3RH, UK*
[3]*Universit de Genve, Dpartement de Physique Thorique and Centre for Astroparticle Physics,*
24 quai Ernest-Ansermet, CH-1211 Genve 4, Switzerland
[4]*Laboratoire Astroparticules et Cosmologie, Universit Paris Diderot,*
10 Rue Alice Domon et Lonie Duquet, 75013 Paris, France
[5]*INFN, Sezione di Padova, 35131 Padova, Italy*
[6]*Department of Physics, Swansea University, Swansea, SA2 8PP, UK*
(Dated: June 8, 2020)

Given the recent advances in gravitational-wave detection technologies, the detection and characterisation of gravitational-wave backgrounds (GWBs) with the Laser Interferometer Space Antenna (LISA) is a real possibility. To assess the abilities of the LISA satellite network to reconstruct anisotropies of different angular scales and in different directions on the sky, we develop a map-maker based on an optimal quadratic estimator. The resulting maps are maximum likelihood representations of the GWB intensity on the sky integrated over a broad range of frequencies. We test the algorithm by reconstructing known input maps with different input distributions and over different frequency ranges. We find that, in an optimal scenario of well understood noise and high frequency, high SNR signals, the maximum scales LISA may probe are $\ell_{\max} \lesssim 15$. The map-maker also allows to test the directional dependence of LISA noise, providing insight on the directional sky sensitivity we may expect.

I. INTRODUCTION

The Laser Interferometer Space Antenna (LISA) is an ESA led mission in collaboration with NASA planned to launch in the mid 2030s which will allow us to tune into a new, vast range of the gravitational-wave (GW) sky, at unprecedented depth and volume.

The LISA sensitivity curve spans roughly 4 orders of magnitude in frequency, from 10^{-5} Hz to 10^{-1} Hz [1], with maximum sensitivity around 10^{-3} Hz, and is such that we may observe loud GW events out to redshift $z = 6$ and beyond [2]. In this range, LISA is expected to measure a motley of signals with respective signal-to-noise ratios (SNRs) ranging from a few hundred to a few, such that data analysis will require careful component separation to disentangle the individual coherent sources. There will also be the need for tools to analyse the stochastic signals which build up incoherently in the time stream.

The collection of incoherent, unresolved signals in the data are typically referred to as the gravitational-wave background (GWB), which is considered stochastic in the limit where there is a statistically significant set of uncorrelated overlapping waves. There are several GW sources which will contribute to a stochastic background in the LISA band. These include galactic, extra-galactic or even primordial compact binaries, too distant or faint to be resolved [3], or a relic background from a burst of inflation or a phase transition at early times [4–9]. All backgrounds mentioned here are the focus of science objectives of the LISA mission [1]. The detection of a cosmological GWB in particular is considered to be the ultimate challenge in cosmology as GWs are the only signal expected to reach us from before recombination and the consequent generation of the cosmological microwave background (CMB), providing unique insight on the origins of the Universe. In this *paper* we present a method to extract the stochastic signal in the LISA data and how to solve for its directionality, effectively mapping it on the sky. This will allow us to assess the angular resolution of the LISA detector for a distribution of extended, unresolved sources, as a function of SNR.

The mapper we have developed is based on an optimal quadratic estimator which solves for maps iteratively in the pixel domain. Our approach is influenced by work done previously with LIGO data, both as a part of the

---

* Project coordinator and corresponding author
  arianna.renzini15@imperial.ac.uk

LIGO Virgo Collaboration [10–12] and as independent efforts [13, 14]. Much work has been done to assess the ability of LISA to detect and characterise GWBs [15–19], and also study the sky response of LISA for map-making purposes [20, 21]. Some studies have focused on how to distinguish detector noise from the galactic stochastic background, often referred to as a foreground when compared to other signals, which is expected to imprint a seasonal modulation on the time stream [22, 23]. These are particularly relevant here as we will probe the ability of LISA to map the galactic foreground in Section III by analysing an idealised signal with a Milky-Way–like sky distribution. Map-making with space-based detectors is discussed extensively in [24, 25], both in the low frequency limit and at high frequencies where a higher resolution may be achieved. However, these approaches reduce the mapping problem to a least-squares solution and solve for the spherical harmonic coefficients of the signal by looking at the individual cross-correlations between data-streams. Additionally, there is work in the literature which proposes solving for phase-coherent backgrounds, which we find is not possible, due to both the characteristics of the signal and the nature of the measurement. In this work, we present the first all-sky analysis based on the full likelihood of the data given a GWB intensity signal. The aim is to deliver the framework of the LISA SGWB map-making pipeline which will be progressively refined and updated until the data becomes available.

To test our procedure we will use injected signals from different input maps in order to verify the reconstruction of a known sky, with simple spectral dependence. The inputs will consist in stationary maps of strain intensity of varying amplitudes and with anisotropies given by a Gaussian random field. These are not to be regarded as realistic simulations of any particular GWB and are used solely for testing purposes.

This *paper* is organised as follows; in Section II we review the strain signal of GWBs and detail the LISA detector response to the intensity of the GWB on the sky. In Section III we introduce the quadratic estimator used to obtain maps of the GWB with LISA, and present results of a number of mapping tests we have run. Finally, we conclude in Section IV by giving an overview of the impact of our results and anticipate useful extensions of our algorithm to be explored in the future.

II. THE SIGNAL OF GRAVITATIONAL WAVE BACKGROUNDS

Typically the contribution of gravitational waves to the overall energy content of the Universe is parametrised by the dimensionless fractional energy density $\Omega_{\rm GW}$ [see e.g. 26]. The spectral dependence of this measure is given by the physical energy density of GWs per logarithmic frequency interval,

$$\Omega_{\rm GW}(f) = \frac{1}{\rho_c}\frac{d\rho_{\rm GW}}{d\ln f}\,, \qquad (1)$$

where ρ_c is the critical energy density. The fraction of this quantity produced by resolved GW sources is strictly marginal, since most of the signal is due to unresolved sources below the confusion limit. The collection of all unresolved GWs defines the GWB, which includes signals of both astrophysical and cosmological nature.

Whilst GWBs of different origin will have differing spectral dependence [18], most will be unpolarised and incoherent in that the temporal phase of the signal is expected to be random. The randomisation of the temporal phase will either be a product of the generation mechanism of the background, a consequence of propagation through the inhomogenous universe, or both [27–32]. Beyond the isotropic value of the GWB over the entire sky we can include any anisotropy by adding a directional dependence $\Omega_{\rm GW}(f) \to \Omega_{\rm GW}(f, \hat{\boldsymbol{n}})$ where $\hat{\boldsymbol{n}}$ is the unit vector of a line-of-sight on the sky. Throughout this *paper* we will assume the frequency and sky dependence of the GWB are independent to good approximation, i.e.

$$\Omega_{\rm GW}(f, \hat{\boldsymbol{n}}) \approx H(f)\,\tilde{\Omega}_{\rm GW}(\hat{\boldsymbol{n}})\,; \qquad (2)$$

this applies both to the model of the incoming signal and the recipe for the signal's reconstruction. This is widely accepted to be a reasonable working assumption by the community, starting from [33]. Indeed, several models for GWBs predict this property to apply for the ensemble average of a stochastic GW signal. Specifically, the signal needs to be stochastic in both the time- and frequency-domain. It is worth noting that we may not be operating in the truly stochastic limit, however we will make this assumption throughout. In this paper, we further assume that the spectral shape of $\Omega_{\rm GW}$ is described by a power law with parameters (α, f_0) such that

$$\Omega_{\rm GW}(f, \hat{\boldsymbol{n}}) = \left(\frac{f}{f_0}\right)^{\alpha} \tilde{\Omega}_{\rm GW}(f_0, \hat{\boldsymbol{n}})\,, \qquad (3)$$

in allignment with both theory- and data- driven analyses of the stochastic GW signal [19, 34, 35]. Specifically for the monopole component of the stochastic background this assumption has been relaxed and a fitting technique has been tested using broken-power-law templates in the ground-based detector frequency band [36]. This fitting technique is very effective when tested on idealized signals, however the application to realistic data requires improvement as component separation will need to be taken into account. In the LISA case, a method to go beyond the single power-law assumption for the reconstruction of the monopole component $\Omega_{\rm GW}(f)$ has been proposed in [18]. We plan to explore the validity of (3) in the future, by analysing a realistic, time-domain generated GWB signal from the LISA Data Challenge (LDC).

FIG. 1. Illustrations of the TDI 1 (left panel) and TDI 1.5 configuration X (right panel). The three numbered grey circles represent the spacecrafts arranged in an equilateral triangle. The links are labelled $\boldsymbol{\ell}_i$ in the left panel for clarity, and are the same as the ones in the right panel.

A. Strain signal

We model the GW strain just outside the detector with the transverse, traceless tensor $h_{ij}(t, \boldsymbol{x})$ at time t and position $\boldsymbol{x}$. h_{ij} can be decomposed into the independent polarisation states h_+ and $h_\times$ and expanded using plane waves as

$$h_{ij}(t, \boldsymbol{x}) = \int_{-\infty}^{+\infty} df \int_{S^2} d\hat{\boldsymbol{n}} \sum_{P=+, \times} h_P(f, \hat{\boldsymbol{n}}) e_{ij}^P(\hat{\boldsymbol{n}}) e^{i2\pi f(\hat{\boldsymbol{n}} \cdot \boldsymbol{x} - t)},$$

$$(4)$$

where polarisation base tensors e^P are

$$e^+ = e_\theta \otimes e_\theta - e_\phi \otimes e_\phi, \qquad (5)$$

$$e^\times = e_\theta \otimes e_\phi + e_\phi \otimes e_\theta, \qquad (6)$$

with

$$e_\theta = (\cos\theta \cos\phi, \cos\theta \sin\phi, -\sin\theta),$$
$$e_\phi = (-\sin\phi, \cos\phi, 0), \qquad (7)$$

as in [35] and we are working in units where $c = 1$. For a stochastic GWB each Fourier mode will contribute independently to the overall signal, and h_+ and $h_\times$ are two random complex fields on the sky. We will assume the amplitudes are drawn from a Gaussian probability distribution. In the case of an astrophysical GWB this assumption holds if the signal is sourced by a sufficiently large number of independent events and all high signal-to-noise outliers have been subtracted from the detector time streams, such that the central limit theorem applies [37]. Under the Gaussian assumption, the statistical properties of the amplitudes are then characterised solely by the second order moments $\langle h_P(f, \hat{\boldsymbol{n}}) h_{P'}^\star(f', \hat{\boldsymbol{n}}') \rangle$, which, assuming statistical homogeneity, correspond to ensemble averages

$$\begin{pmatrix} \langle h_+ h_+'^\star \rangle & \langle h_+ h_\times'^\star \rangle \\ \langle h_\times h_+'^\star \rangle & \langle h_\times h_\times'^\star \rangle \end{pmatrix} = \frac{1}{2} \delta(\boldsymbol{n} - \boldsymbol{n}') \delta(f - f') \times$$
$$\begin{pmatrix} I(f, \hat{\boldsymbol{n}}) + Q(f, \hat{\boldsymbol{n}}) & U(f, \hat{\boldsymbol{n}}) - iV(f, \hat{\boldsymbol{n}}) \\ U(f, \hat{\boldsymbol{n}}) + iV(f, \hat{\boldsymbol{n}}) & I(f, \hat{\boldsymbol{n}}) - Q(f, \hat{\boldsymbol{n}}) \end{pmatrix}, \qquad (8)$$

where we have introduced the Stokes parameters I, the intensity, Q and U, giving the linear polarisation, and V, the circular polarisation. The four Stokes parameters completely describe the polarisation of the observed signal in analogy with electromagnetic Stokes parameters for the photon. The difference here is that whilst the electromagnetic Q and U Stokes parameters transform as spin-2 quantities with respect to rotations, their strain counterparts transform as spin-4 under rotations. In both cases the intensity I behaves as a scalar under rotations.

In the following we will restrict our analysis to the reconstruction of the GWB intensity $I(f, \hat{\boldsymbol{n}})$ only, corresponding to the combination

$$I(f, \hat{\boldsymbol{n}}) = \langle h_+(f, \hat{\boldsymbol{n}}) h_+^\star(f, \hat{\boldsymbol{n}}) \rangle + \langle h_\times(f, \hat{\boldsymbol{n}}) h_\times^\star(f, \hat{\boldsymbol{n}}) \rangle, \qquad (9)$$

note that this relates to the normalised logarithmic energy density $\Omega_{\mathrm{GW}}(f)$ as [33]

$$\Omega_{\mathrm{GW}}(f) = \frac{4\pi^2 f^3}{G\rho_c} I(f), \qquad (10)$$

such that the assumption in Eq. 3 carries down to the intensity

$$I(f, \hat{\boldsymbol{n}}) = \left(\frac{f}{f_0}\right)^{\alpha-3} I(f_0, \hat{\boldsymbol{n}}). \qquad (11)$$

This approximation is valid and sufficient for our calculations here, and has been assumed in multiple analyses of the stochastic background. Typically a value of $\alpha = 2/3$ is assumed for inspiral-dominated backgrounds [38], whereas $\alpha = 0$ may be used to describe scale-invariant cosmological backgrounds [9, 39].

B. Detector Response

The antenna is comprised of three spacecrafts which will be positioned on three heliocentric yearly orbits in

FIG. 2. Normalised auto-correlated response of TDI channel X, $\boldsymbol{A}_{XX}$, at time $t = 0$ and in the Solar System Baycentre (SSB) reference frame, at frequencies $f = 10^{-4}$ Hz, $f = 10^{-1}$ Hz, $f = 5 \times 10^{-1}$ Hz from left to right respectively.

such a way that they will maintain an equilateral triangle configuration throughout the whole duration of the mission (4+ years) [1]. The spacecrafts will specify a detector plane at a 60° angle with the ecliptic consistently throughout the motion of the constellation, which will appear to rotate on the detector plane with period 1 year. Each pair of spacecrafts will share two laser links, which we can represent as two opposite vectors along the virtual detector arm. The arm-length will be time dependent during the mission, however in this work, for the sake of simplicity, we assume equal and stationary arms. Consequently, we consider equal and opposite links along each arm. In this paper we will use the trajectories as defined in [40], which correspond to analytical elliptical orbits in the Solar System Barycentre (SSB) frame. The starting position at zero time is chosen such that $\boldsymbol{\ell}_2$ in Fig. 1 is aligned with the negative y axis of the frame, and the two remaining arms are determined accordingly.

GW detection with LISA will rely on time-delay interferometry (TDI), which involves time-shifting and linearly combining independent Doppler measurements [41]. This will occur in post-processing and will require accurate modelling of the response of each link. The study of the optimal TDI combinations, or channels, for GW detection is ongoing [42]. Given 6 links it is possible to construct three correlated channels, referred to as X, Y, Z, or two independent channels, typically referred to as A, E channels, plus a *Sagnac* channel T [43]. The independent channel basis is found simply by diagonalising the correlated basis and rotating into the diagonal TDI space. As the measurements are time-delayed, this rotation will never be perfect and it adds a degree of complication which may be hard to constrain, hence we find it is preferable to work in the fully correlated basis, with the caveats it entails. Specifically, the rotation will only cancel out the correlations assuming the three armlengths are equal, and that the noise in each spacecraft is identical. While these assumptions are made in this work, the method developed here is general and may be easily extended to accommodate a more complete noise model and detector response. The simple Michelson-Morley-style combination is referred to as TDI 1 and it features three channels, each of which is centered around a spacecraft such that X is centered around spacecraft

1, Y around spacecraft 2, and Z around 3. A similar, slightly more sophisticated combination designed to minimise the spacecraft breathing noise is referred to as TDI 1.5; both are illustrated in Figure 1. Throughout this paper we will adopt the TDI 1.5 configuration.

The time stream $s_C(t, \boldsymbol{x_i})$ measured by a single 1.5 TDI channel C at time t and ith spacecraft position $\boldsymbol{x_i}$ can be Fourier expanded between t and $t + \Delta t$ to yield the signal as a function of frequency,

$$s_C^\tau(f) = \sum_P \int_{S^2} d\hat{\boldsymbol{n}} \, R_C^P(f, \hat{\boldsymbol{n}}; \tau) \, h^P(f, \hat{\boldsymbol{n}}), \qquad (12)$$

where τ is the time segment label, R_C^P is the polarisation response function for the TDI channel C and $h^P(f, \hat{\boldsymbol{n}})$ is the incoming GW strain decomposed into its polarisation components. We will now drop the τ label for simplicity. The response functions for TDIs 1.5, $\{X, Y, Z\}$ are derived from the *LDC Manual* which is an internal document. In the case of TDI 1.5 X the measured strain may be written as

$$s_X(f) = -2i \sin(a) \, e^{-ia} \left[e^{-ia} \left(y_{\boldsymbol{\ell}_1} - y_{-\boldsymbol{\ell}_3} \right) + y_{-\boldsymbol{\ell}_1} - y_{\boldsymbol{\ell}_3} \right], \tag{13}$$

where the $y_{\hat{\boldsymbol{\ell}}_i}$ terms correspond to the strain measured along the oriented link $\boldsymbol{\ell}_i$ and the exponential terms correspond to the phase shifts in Fourier space required to build the TDI 1.5 channel. Here $a = 2\pi f L$ for the sake of conciseness. The response of each link is

$$y_{\boldsymbol{\ell}_i} = -\frac{ia}{2} \sum_P \int_{S^2} Q_{\hat{\boldsymbol{\ell}}}^P h^P \, e^{-\frac{ia}{L} \hat{\boldsymbol{n}} \cdot \boldsymbol{x}_i} \operatorname{sinc}(b) \, e^{-ib}, \tag{14}$$

where $b = \frac{a}{2}(1 - \hat{\boldsymbol{n}} \cdot \hat{\boldsymbol{\ell}}_i)$ and the response of a single arm $\boldsymbol{\ell}$ to P- polarised modes, $Q_{\hat{\boldsymbol{\ell}}}^P$, is simply the contraction of the arm tensor with the polarisation basis element,

$$Q_{\hat{\boldsymbol{\ell}}}^P = \epsilon^P : \hat{\boldsymbol{\ell}} \otimes \hat{\boldsymbol{\ell}}. \tag{15}$$

Plugging (14) into (13) we recover the sky response

$$R_X^P = \frac{ia}{2} \left(1 - e^{-2ia} \right) e^{-\frac{ia}{L} \hat{\boldsymbol{n}} \cdot \vec{x}_1} \left[Q_{\hat{\boldsymbol{\ell}}_1}^P \mathcal{T}(+\hat{\boldsymbol{\ell}}_1) - Q_{\hat{\boldsymbol{\ell}}_3}^P \mathcal{T}(-\hat{\boldsymbol{\ell}}_3) \right] \tag{16}$$

where the transfer function $\mathcal{T}$ to each sky direction $\hat{\boldsymbol{n}}$ is

$$\mathcal{T}(+\hat{\boldsymbol{\ell}}) = \text{sinc}\left(\frac{a}{2}(1 - \hat{\boldsymbol{n}} \cdot \hat{\boldsymbol{\ell}})\right) e^{-\frac{ia}{2}(3+\hat{\boldsymbol{\ell}}\cdot\hat{\boldsymbol{n}})} +$$
$$+ \text{sinc}\left(\frac{a}{2}(1 + \hat{\boldsymbol{n}} \cdot \hat{\boldsymbol{\ell}})\right) e^{-\frac{ia}{2}(1+\hat{\boldsymbol{\ell}}\cdot\hat{\boldsymbol{n}})}, \quad (17)$$

assuming a permanent equal arm configuration, such that $\boldsymbol{x}_i = \boldsymbol{x}_{i-1} + \boldsymbol{\ell}_{i-1}$ where i obeys cyclic permutations. One can then derive R_Y^P, R_Z^P by permuting the arms.

A signal composed of an incoherent superposition of many components, as in the case that we are focused on here, will vanish when averaged in time. To observe an incoherent GWB we therefore need to consider the integration of the square of the signal. Making use of all three channels we can write the TDI strain vector as $\boldsymbol{s} = (s_X, s_Y, s_Z)$ and construct the quadratic strain tensor $\boldsymbol{S}(f)$,

$$\boldsymbol{S}(f) = \boldsymbol{s}(f) \otimes \boldsymbol{s}^{\star}(f). \quad (18)$$

We expect its ensemble average to yield

$$\langle \boldsymbol{S}(f) \rangle = \int_{S^2} d\hat{\boldsymbol{n}}\, \boldsymbol{A}(f, \hat{\boldsymbol{n}})\, I(f, \hat{\boldsymbol{n}}), \quad (19)$$

having expanded in the polarisation bases and inserting the relation between the second order moments of the strain and the Stokes parameters [33]. Note that we consider a GWB which is stochastic also in polarisation; as such the Q, U, and V Stokes parameters average to zero. The quadratic response tensor $\boldsymbol{A}$ is simply constructed with the linear response TDI vectors $\boldsymbol{R}^P = (R_X, R_Y, R_Z)^P$

$$\boldsymbol{A}^{\tau}(f, \hat{\boldsymbol{n}}) = \boldsymbol{R}^+ \otimes \boldsymbol{R}^{+\star} + \boldsymbol{R}^{\times} \otimes \boldsymbol{R}^{\times\star}. \quad (20)$$

We have recovered the τ label to point out the time-dependence of the detector response, which is implicitly expressed in Eqs. (13-20) through the arm and positions vectors. Note also the directional and frequency dependence of $\boldsymbol{A}$, a sample of which is given in Fig. 2. Here the continuous sky response has been discretised to produce the Mollweide projection in pixel space, $\boldsymbol{A}(\hat{\boldsymbol{n}}) \to \boldsymbol{A}_p$, where p labels a pixel on the sky, and it is shown in the SSB reference frame. The projections and pixel calculations are carried out using the HealPix package [44]. As in the case of LIGO, the sky response gives rise to an inhomogeneous and non-compact antenna pattern. However, in the case of the LISA TDI channels this pattern is also significantly dependent on the frequency probed, as may be observed in the different panels of Fig. 2. It is straightforward to show that in the low frequency limit the spectral dependence of the quadratic response tensor $\boldsymbol{A}$ scales as f^4 [18], then starts deviating above 10^{-2} Hz, as may be observed in Fig. 3. The constant pattern and flat trend will set the limit to the resolution of LISA at low frequency, while this substantial deviation and the high-ℓ pattern will prove crucial to obtain higher resolution at high frequency.


FIG. 3. Sky-integrated auto-correlated and cross-correlated responses, A_{XX} and A_{XY}, across the frequency spectrum. Note that the slopes in the low frequency limit scale as $\propto f^4$.

For the simulations and study that follows we fix the coordinate system and polarisation basis to the SSB frame as it is convenient for the description of both the spacecraft trajectories and the injected signal's directional dependence.

Although we have developed a map–maker that solves for a map directly in the pixel domain, we can use it to assess the ability of the cross-correlated TDI channels to reconstruct the angular power spectrum of a background defined as

$$C_\ell^{\text{GW}} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \left| \int_{S^2} \frac{d\hat{\boldsymbol{n}}}{4\pi} Y_{\ell m}(\hat{\boldsymbol{n}})\, I(\hat{\boldsymbol{n}}) \right|^2. \quad (21)$$

It is then informative to decompose the detector response in spherical harmonic space, similarly to [35], and assess the sensitivity to C_ℓ^{GW} induced by the isotropised spherical harmonic response,

$$\boldsymbol{A}_\ell = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \left| \int_{S^2} \frac{d\hat{\boldsymbol{n}}}{4\pi} Y_{\ell m}(\hat{\boldsymbol{n}})\, \boldsymbol{A}(\hat{\boldsymbol{n}}) \right|^2. \quad (22)$$

Studying the relative sensitivity of a single TDI channel to angular scales ℓ on the sky, the angular resolution of the LISA TDI 1.5 configuration is predicted to be constrained as $\ell \lesssim 10$ [15, 19]. This resolution will strongly depend on the frequency interval examined, such that for lower frequencies the cutoff will be much more severe. This discussion is expanded in [45].

III. MAXIMUM LIKELIHOOD MAPS

A. Mapping method

In this section we describe the method we have developed to extract maximum likelihood maps of the GWB

from three LISA TDI 1.5 data streams. We start with some remarks on analogies and differences with CMB analysis.

The final parts of the analysis of LISA data for the purpose of constraining GWBs will be split into two main stages. The first will be an estimate of a map in a way that is similar to the estimation of CMB maps from timestream data. This first compression step, from time (frequency) domain to sky coordinates, will be useful for identifying systematics in the data and for separating out different signals contributing to the GWB. The second stage will be a further compression to an isotropised statistic i.e. the angular power spectrum of the (separated) signal. This is similar to the power spectrum estimation step in CMB taking the maps to a final C_ℓ estimate.

In CMB observations made with incoherent (bolometric) detectors, the first stage uses a χ^2 estimate to obtain a maximum likelihood map [46]. This is possible because the signal component of the data does not vanish in the ensemble mean

$$\langle d \rangle \equiv \langle s + n \rangle = m \,, \tag{23}$$

where, generically, d, s, and n are the time-domain data, signal, and noise with $\langle n \rangle = 0$ and model $m \equiv \langle s \rangle$. This allows the definition of a residual $d - m$ whose covariance is just the noise covariance $N \equiv \langle nn^\dagger \rangle$. This means we can estimate the signal by considering simply a Gaussian likelihood for the residual and, since the noise is usually considered to be independent of the signal, the normalisation of the likelihood is constant and we only need to minimize the χ^2 part

$$\chi^2(d|m) = (d - m)^\dagger N^{-1} (d - m) \,, \tag{24}$$

to obtain a maximum likelihood solution for m using the closed-form solution to $\partial \chi^2 / \partial m = 0$.

This is not possible for data obtained from noise correlated, coherent detectors such as LISA. In this case both signal and noise components in the time-domain vanish in the ensemble mean. Transforming to the Fourier domain does not alleviate this problem unless the background itself is phase-coherent or if the signal is a single point source on the sky. In this case no residual can be defined since $\langle d - m \rangle = 0$ and we need to use the full likelihood of the data with both χ^2 and normalisation depending on the signal being estimated. There is no closed-form solution for the maximum likelihood in this case and iterative methods such as quadratic estimators must be used to estimate the signal [47]. This is analogous to the *second* stage in CMB analysis where the map signal component also vanishes in the ensemble mean limit and therefore the full data likelihood is used to estimate the signal. For LISA *both* stages will require the use of a quadratic likelihood estimator and we define a method for the first stage below.

Note that a χ^2 estimate is possible for LIGO style detectors where the noise can be assumed to be uncorrelated between detectors [12–14, 48]. In that case we can

define a likelihood for a residual in the cross-correlated data where the signal component ensemble mean does not vanish but the noise does. In the case of LISA TDI $\{X, Y, Z\}$ channels, the noise will be correlated between channels and this approach is not possible.

Firstly, we define the likelihood for the data given a signal *intensity*. We decompose the data TDI vector as $\boldsymbol{d}_f^\tau = \boldsymbol{R}h + \boldsymbol{n}$, where $\boldsymbol{n}$ is the noise TDI vector, $\boldsymbol{R}$ is the linear response TDI vector and there is an implicit sum over polarisations. Assuming the noise is zeromean and Gaussian with covariance $\boldsymbol{N}_f = \boldsymbol{n} \otimes \boldsymbol{n}$ and the signal component is also Gaussian with covariance $\boldsymbol{C}_f^\tau = \boldsymbol{A}\tilde{I} + \boldsymbol{N}$, the likelihood $\mathcal{L}$ of a pixel map $\tilde{I}$ given the TDI data $\boldsymbol{d}$ is

$$\mathcal{L} \propto \frac{1}{|\boldsymbol{C}|^{1/2}} e^{-\frac{1}{2} \boldsymbol{d} \, \boldsymbol{C}^{-1} \boldsymbol{d}^\star} \,, \tag{25}$$

Following [47] we find the solution which maximises $\mathcal{L}$,

$$\tilde{I}_p = \mathcal{F}_{pp'}^{-1} \cdot \mathrm{Tr} \sum_{\tau, f} \left[\boldsymbol{C}^{-1} \frac{\partial \boldsymbol{C}}{\partial I_{p'}} \boldsymbol{C}^{-1} \left(\boldsymbol{D} - \boldsymbol{N} \right) \right] \,, \tag{26}$$

$$\mathcal{F}_{pp'} = \mathrm{Tr} \sum_{\tau, f} \left[\boldsymbol{C}^{-1} \frac{\partial \boldsymbol{C}}{\partial I_p} \boldsymbol{C}^{-1} \frac{\partial \boldsymbol{C}}{\partial I_{p'}} \right] \,, \tag{27}$$

where we have explicitly written down the pixel indices for clarity. $\boldsymbol{D} = \boldsymbol{d} \otimes \boldsymbol{d}^\star$ is constructed thanks to the permuting property of the trace. $\mathcal{F}$ is the Fisher information matrix, whereas the trace in Equation (26) is referred to as the gradient term.

This approach is general and may be used for both a broad- or a narrow- band analysis, and there is no specific limit on the time-discretisation as long as the Gaussian ansatz is not violated. In the application considered here, the trace is taken over the three TDI channels, and the sum is over all frequencies in the FFT and all observation times, to maximise sky coverage. We can also consider integrating over short frequency intervals or single frequencies to probe the frequency dependence of the GWB, within the limits imposed by the conditioning of the Fisher matrix. We employ an iterative scheme to reach the maximum likelihood estimate for I, starting with an initial guess $\tilde{I}_{\mathrm{in}}$ to plug into $\boldsymbol{C}$, get a first estimate of $\tilde{I}_p$, then repeat until convergence. In the analysis here we assume to perfectly know the noise model $\boldsymbol{N}$, and expect that when this method will be applied to real LISA data it will be possible to rely on an independent noise estimation which informs the gradient term. However, it will be possible yet expensive to extend this estimator to include the noise as a free parameter, and solve for both noise and signal components directly.

In case the signal is non-Gaussian, the maximum–likelihood maps derived below remain unchanged, but the interpretation of their covariance would be affected by the presence of higher order cumulants in the underlying probability densities. See e.g. [49] for an analysis of detection methods of non-Gaussianity in a GWB induced

by short-duration signals, and [50] for a study of how to use higher-order cumulants to characterize its properties.

It is important to stress that in this approach we assume that the intensity scales simply across the spectrum by some function $E(f)$. In reality this will not be the case and the model will have to account for a number of sky signals, each having a distinct spectral dependence. For a broad band analysis such as that carried out here one would have to solve for multiple maps. Alternatively, an analysis where the data is separated into narrow frequency bands, such as one of the approaches in [14], could deal with multiple signals through a later component separation stage.

B. Mapping tests

Here we present the results of the application of the method detailed above to simulated data of LISA TDI 1.5 $\{X, Y, Z\}$ channels. These serve as a proof of concept of the recipe, which is the first stepping stone towards testing it on mock LISA time streams. We generate correlated data in the Fourier domain as $\boldsymbol{D}_f^\tau = \boldsymbol{A}I_{\text{in}} + \boldsymbol{N}_{\text{in}}$ over the period of one year divided in two-day segments for different frequency bands. This FFT time-scale determines the pixelisation resolution, as we consider the sky response to be stationary throughout the whole time segment, hence it is chosen to be sufficiently short to allow for a high N_{pix}, but sufficiently long so as not to exclude the lower frequencies. Each segment is heavily down-sampled in frequency to $\mathbb{N}_f = 200$ samples in order to lighten the computational load as much as possible without losing too much resolution in frequency, which would lead to a loss of information on the sky due to the coupling between f and $\hat{\boldsymbol{n}}$ via $\boldsymbol{A}$. $\boldsymbol{A}I_{\text{in}}$ is the signal component of the data contracted over pixels and $\boldsymbol{N}_{\text{in}}$ is a Gaussian realisation of the noise model $\boldsymbol{N}$. The resulting data streams for the three channels are then plugged into Eq. (26), and the Fisher matrix is iteratively computed, where the covariance matrix $\boldsymbol{C}$ is the full covariance of the data and simply $\frac{\partial \boldsymbol{C}}{\partial I_p} = \boldsymbol{A}_p$. All the terms in Eqs. (26 - 27) are complex matrices and the reality condition imposed by the data is recovered when integrating over frequencies.

The nominal SNR of the signal component of the data is estimated by integrating over observation time and frequency range as follows:

$$\text{SNR} = \sqrt{\Delta T_{\text{obs}} \int_{f_{\text{min}}}^{f_{\text{max}}} df \left(\frac{AI_{\text{in}}}{N}\right)^2}, \qquad (28)$$

where A is the instantaneous response of channel X and N is the appropriate noise model.
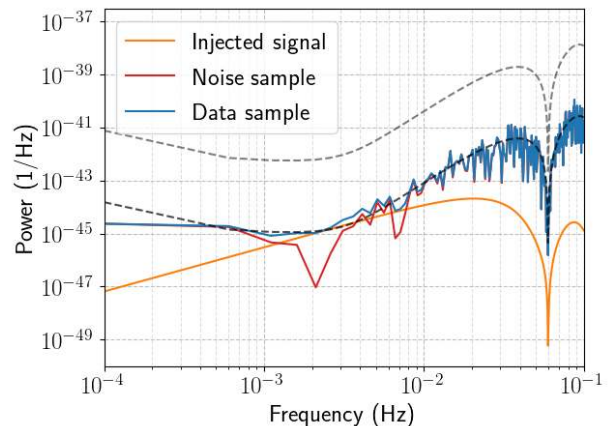

FIG. 4. Data power sample from the autocorrelated X channel, decomposed into injected signal (orange line) and simulated noise (red line) components. The Injected signal has and SNR of 0.6, and is evidently buried in the noise. The sample represents the data accumulated over a two-day period and then FFTed. The dashed grey lines are the noise model which informs the noise generation; the darker line is renormalised by a factor which accounts for the lower sampling rate, while the lighter line is the true instantaneous noise curve used to calculate the SNR.

1. Signal generation

The signal component in the data $\boldsymbol{A}I_{\text{in}}$ is generated by simply scanning an input map for the GWB intensity on the sky $I_{\text{in}}^p(f_0)$ fixed at a reference frequency f_0 with the quadratic response of our TDI configuration directly in pixel space,

$$(\boldsymbol{A}I_{\text{in}})_f^\tau = \frac{4\pi}{N_{\text{pix}}} \sum_p \boldsymbol{A}_p^\tau(f) I_{\text{in}}^p(f_0), \qquad (29)$$

where we include the spectral dependence of the signal $E(f)$ in the response $\boldsymbol{A}$ for convenience. This is a completely deterministic calculation which relies on the assumptions that the signal is truly stochastic and obeys Eq. 2, and that the overall integration over time and frequencies satisfies the ensemble average limit, such that Eq. 8 holds. All the maps presented here are reconstructed with number of pixels $N_{\text{pix}}^{\text{out}} = 768$ which corresponds to a HealPix $N_{\text{side}} = 8$, while the injected maps are over-resolved with $N_{\text{pix}}^{\text{in}} = 3072$ which corresponds to a HealPix $N_{\text{side}} = 16$.

Except where explicitly stated, the I_{in} maps scanned in the simulations below are random Gaussian realisations of an $\ell^2 C_\ell$ flat power spectrum with $\ell_{\text{max}} = 20$, chosen uniquely for testing purposes. This means that we are assuming that the monopole C_0 sets the size of the anisotropies which scale as a fixed power law $C_\ell = C_0/\ell$. In a more realistic modeling of the signal, one should rather introduce two separate scales: a monopole C_0, which would set the size of the SNR as in 28, and a typical scale of anisotropies C_1. These two scales are simply

FIG. 5. An example input map from the simulations (left panel) to be compared to the final output maps obtained integrating with different frequency cutoffs, $f_{\max} = 0.1$ Hz (central panel) and $f_{\max} = 0.01$ Hz (right panel). These highlight the different resolutions the LISA channels have in different ranges of frequency.


FIG. 6. Transfer functions T_ℓ for the average reconstructed C_ℓs obtained with different frequency cutoffs (left panel) and different spectral shapes, $\alpha = 3$ and $\alpha = 0$, both in the high frequency case $f_{\max} = 0.1$ Hz (right panel). Each simulation set consists of 50 maps, each a different realisation of the same C_ℓ input. There appears to be a clear one-to-one relation between the resolution $\ell_{\max}$ of the instrument and the frequency cutoff. Conversely, there is an average difference of 5% between the transfer functions obtained with different spectral weightings, however this does not affect the resolution cutoff.

assumed to be the same here. We explore the ability of LISA to reconstruct these input maps by changing the frequency integration range, varying the spectral parameter α, and toggling the sky-integrated amplitude $I(f_0)$ with respect to the noise level set by N.

2. Noise generation

LISA TDI noise is an active area of research, and there are ongoing studies on which TDI configuration will yield the lowest noise measurements. However these go beyond the scope of our project, as these model the orbits and breathing modes of the constellation more realistically. In the simplified equal-arm scenario considered here, to good approximation the noise in each auto-correlation of TDI channels may be described by the same power spectrum, S_{CC}, and similarly the noise in each cross-correlation may be described by S_{CD}. The 3×3 model TDI noise correlation matrix N is then completely described by S_{CC} on the diagonal and S_{CD} on the off-diagonal terms. The expressions for these power spectra

used in this work are

$$S_{CC} = 16 \sin^2 a \left(S_{\text{int}} + (3 + \cos 2a) S_{\text{acc}} \right), \tag{30}$$

$$S_{CD} = -8 \sin^2 a \, \cos a \left(S_{\text{int}} + 4 S_{\text{acc}} \right), \tag{31}$$

as presented in [19], where S_{int} and S_{acc} are the interferometer and acceleration noise-components respectively.

The noise realisation N_{in} must respect the correlations imposed by N, hence the noise is first generated linearly in the noise-diagonal space, and then rotated back into correlated noise space. This is achieved by generating a random 3-vector in noise-diagonal space and rotating it into the TDI noise vector $n_{\text{in}} = (n_X, n_Y, n_Z)_{\text{in}}$ using the eigenvector matrix of N. In the limit of equilateral configuration and identical noise at the vertices that we are considering here, the noise-diagonal space is precisely the space of the $\{A, E, T\}$ channels, and the same rotation as described in [43] is employed to transform from one to the other. N_{in} is then simply the outer product of n_{in}. An example data segment is provided in Fig. 4.

FIG. 7. Input (left panel) and reconstructed output (central panel) maps in a low SNR case, integrating over a frequency range of $[10^{-4}, 10^{-1}]$ Hz. The input map is an ℓC_ℓ flat gaussian realisation with $\ell_{\max} = 10$. The output map is a heavily smoothed version of the input, with some loss of power due to the strong conditioning required. The SNR of this particular reconstruction is shown in the right panel.

3. Results

We have run a variety of simulations to extensively probe the ability of the TDI configurations $\{X, Y, Z\}$ to reconstruct an anisotropic GWB, both in high and low signal-to-noise ratio (SNR) scenarios. For all the maps discussed below the inversion problem is strongly ill-conditioned, hence the Fisher matrix is inverted using the singular value decomposition technique as in similar work done with LIGO data [13, 14]. Specifically, the condition number of the Fisher matrices in the cases presented below, which are all of dimension $N_{\text{pix}}^{\text{out}} \times N_{\text{pix}}^{\text{out}}$, is of order 10^{18}. In the case of high SNR the output maps mildly depend on the conditioning imposed at pseudo-inversion as the signal will dominate the information, whereas in the low SNR scenario the output map is extremely dependent on the conditioning. In these tests, the choice made was to monitor the monopole level of the output map and gauge the conditioning such that it would match the input. This technique may be employed also with real data, where the measurement of the monopole can be done independently of the inversion problem.

To test the limits of the geometric set-up of LISA, we have generated data with anisotropic, flat spectrum backgrounds ($\alpha = 3$) of effectively infinite SNR and reconstructed output maps truncating the frequency integration at different values of $f_{\max}$. In this extremely high signal scenario, the solution converges after a single iteration. As may be seen in the example maps in Fig. 5, the input map is reconstructed remarkably differently in the case of $f_{\max} = 10^{-1}$ Hz and $f_{\max} = 10^{-2}$ Hz. The higher angular modes are well preserved when allowing the reconstructor to integrate up to higher frequencies, where there is finer structure in the response pattern, whereas they are aliased into lower modes when integrating only over the lower frequencies. To study this trend we have run sets of 50 analogous simulations when $f_{\max} = 10^{-3}$ Hz, 10^{-2} Hz, 5×10^{-2} Hz, 10^{-1} Hz and calculated the average output C_ℓs in each set. We also compute the transfer functions $T_\ell = C_\ell / C_\ell^{\text{in}}$. The comparison between transfer functions with different fre-


FIG. 8. Sky distribution of the noise in the SSB frame. Note the imprinted 6-fold symmetry of the orbit, given by the three concentric orbits of the spacecrafts.

quency cutoffs may be seen in the left panel of Fig. 6.

Different spectral shapes are reconstructed slightly differently, as may be seen in the right panel of Fig. 6 where the high frequency $\alpha = 3$ set is compared with a high frequency $\alpha = 0$ set of simulations. The C_ℓ reconstruction and resulting transfer functions differ at certain modes, on average by 5%, but the cutoff $\ell_{\max}$ appears to be the same, hence we conclude the frequency weighting in the signal and reconstruction do not have a strong impact on the resolution of the reconstructed maps.

An example of map reconstruction in the presence of high noise is shown in Fig. 7. The map in input here is a ℓC_ℓ flat Gaussian realisation with maximum angular scale $\ell_{\max} = 10$ to make the visual comparison easier. The spectral shape of the signal is $\alpha = 2/3$, which together with the power spectrum is typically associated with an astrophysical inspiral-dominated GWB [38] which traces the large scale structure [51]. The signal component has an SNR of 0.6 and is thus buried under the noise, as may be seen in the specific data sample in Fig. 4. As may be observed in Fig. 7, the presence of loud noise has a similar effect as having a lower frequency cutoff, as the noise dominates at higher frequencies. Additionally, severe conditioning is required to cut out the

FIG. 9. Sky inputs (left panels) and reconstructed outputs (right panels) for a Milky Way - like GWB distribution. The signal is injected in the range $[10^{-4}, 5 \times 10^{-3}]$ Hz. The top row is an extremely high SNR case, whereas the bottom row is an SNR 52, hence it emerges above the noise around 10^{-3} Hz. Both results are obtained over a full year of integration. These Mollweide projections are presented in log scale.

noisier high modes, and multiple iterations are necessary to reach convergence. Specifically, the final conditioning cuts 97% of the total eigenvalues of the Fisher information matrix. The value of the monopole of the reconstructed map however does not depend strongly on the conditioning and is in good agreement with that in input. Thus the resulting heavily conditioned output map appears to be a low-ℓ smoothed version of the initial map with some loss of power. We also recover the noise distribution on the sky $\mathcal{N}$, shown in Fig. 8, as the diagonal of the inverse Hermitian square root of the converged Fisher matrix. This is an incomplete representation of the noise as the off-diagonal correlations are ignored, however it already shows that there is more noise off the ecliptic, as expected.

Finally, an example of reconstruction of a galaxy-like sky signal both with and without noise over a full year of integration is shown in Fig. 9. The input signal has spectral shape $\alpha = 2/3$ in the frequency range $[10^{-4}, 5 \times 10^{-3}]$ Hz, as suggested in [52], and the directional dependence is that of the Milky Way's white dwarf binary (WDB) population from [53]. The high SNR case is extremely unrealistic but provides a useful visual test of what the best resolution of LISA can be for a confused GWB at the frequency range where the signal from the compact

galactic WDB is dominant. By relaxing the conditioning and including the higher modes in the Fisher matrix before inverting, the resolution can be pushed slightly higher however the pixels appear noisier, as may be observed in the top right panel of Fig. 9. The low SNR case proves, as in the simulation discussed above, that strong conditioning is required to recover the input distribution and that this leads to the exclusion of higher modes in the Fisher matrix. The level of the noise in this case is the same as in the one above. The monopole of the signal is very well recovered in both cases.

IV. CONCLUSIONS

We have presented a mapping algorithm for the LISA GW detector, which is treated as 3 correlated TDI channels. The method involves an iterative estimator which calculates the Fisher information matrix and gradient term by tracing over the channels, observation time, and frequencies efficiently. The tests discussed above clearly show there is a one-to-one relationship between the resolution $\ell_{\max}$ of the instrument and the frequency cutoff in the trace, and this is not particularly dependent on the input signal's spectral shape nor the spectral weight-

ing assumed in the reconstruction. This frequency cutoff will naturally depend on both the nature of the signal and the level of the noise. As may be read off of Fig. 6, the maximum angular resolution for LISA TDI 1.5 channels is $\ell_{\max} \sim 15$ in the best case scenario when the signal is loud and clear over the noise up to $f_{\max} = 10^{-1}$ Hz. A possible candidate for this type of signal is a stochastic background of astrophysical origin from stellar mass black hole and neutron star binaries [54]. The outlook is a lot less optimistic for lower frequency signals which are strong at $f < 10^{-2}$, for which $\ell_{\max} < 7$. An intermediate signal, which should peak around $f \sim 5 \times 10^{-2}$ [52], is the galactic background from binary white dwarfs. Considering this type of signal to be described by a simple power law in frequency, we have tested its reconstruction on the sky and have found that the strongest limit is set by the strong noise modes at high frequency. The reconstruction of low SNR signals is similarly limited, even if they are present at high frequencies, due to the shape and nature of the noise.

In the tests presented above we assume that the size of the monopole sets both the SNR and the anisotropy level of the injected signal. However a more realistic modeling would see the introduction of two separate scales, a monopole C_0 and a typical anisotropy scale C_1. This would allow to test separately the detectability of background components with similar monopole values but different levels of anisotropy. For example, astrophysical extra-galactic backgrounds are expected to have a level of anisotropy $\delta\Omega/\bar{\Omega} = 10^{-3}$ whereas cosmological backgrounds should present CMB-like ansotropies of $\delta \sim 10^{-5}$ [39, 51]. Hence, cosmological backgrounds from the early universe are expected to have a low SNR and a highly suppressed scale invariant anisotropy spectrum. These include inflationary backgrounds which seed primordial black holes [55], highly non-Gaussian backgrounds [56, 57], post-inflationary backgrounds due to strong first order phase transitions [9] or cosmic defect networks [58]; see [59] for a full review. Map reconstruction for these background components will be quite challenging. The most promising candidate remains an astrophysical galactic background, hence we have carried out a preliminary study here. The LDC has produced mock time-domain data for a galactic background made up of thousands of superposed white dwarf binary inspirals which we plan to send through our pipeline soon.

The results obtained in this paper have been derived under the ideal assumption of a perfectly known noise PSD model. A natural expansion of this method would be to parametrise a given uncertainty in the noise model and performing parameter estimation on the noise and signal parameters simultaneously. This can happen by fitting noise parameters for the acceleration and interferometer components as seen in Equ.s (30–31) for each cross-correlated channel or in each individual TDI channel. A technique for simultaneous reconstruction of the monopole component of the signal and a simplified mode of the noise power spectra, parametrised by acceleration and optical metrology system contributions only, has been implemented in [22, 23], and more recently in [17, 18]. In addition, it has been shown that, under specific assumptions, a similar technique can also be efficient in the presence of idealised foregrounds [60].

Another significant limit of this study is the assumption that the GWB intensity $I(f, \hat{\boldsymbol{n}})$ can be reduced to two independent components: a simple spectral shape and a constant pattern on the sky. We plan to explore the validity of this assumption in future work, studying specific GWB signals in the LISA band and analysing the time-domain generated LDC data containing anisotropic backgrounds. The results of this mapping algorithm may also be used to assess the detectability of different types of anisotropic backgrounds in the LISA band by studying the effective sky sensitivity of the detector and comparing it to the expected levels of anisotropy. This will be investigated further in [45].

V. ACKNOWLEDGEMENTS

[1] P. Amaro-Seoane, H. Audley, S. Babak, J. Baker, E. Barausse, P. Bender, E. Berti, P. Binetruy, M. Born, D. Bortoluzzi, et al. (2017).

[2] N. Tamanini, C. Caprini, E. Barausse, A. Sesana, A. Klein, and A. Petiteau, JCAP **1604**, 002 (2016).

[3] T. Regimbau, Research in Astronomy and Astrophysics **11**, 369 (2011), ISSN 16744527, arXiv:1101.2762v3, URL http://stacks.iop.org/1674-4527/11/i=4/a=001.

[4] L. Grishchuk, Soviet Journal of Experimental and Theoretical Physics **40**, 409 (1975).

[5] M. Maggiore, Phys. Rept. **331**, 283 (2000), gr-qc/9909001.

[6] C. J. Hogan, Monthly Notices of the Royal Astronomical Society **218**, 629 (1986), URL http://dx.doi.org/10.1093/mnras/218.4.629.

[7] R. A. Battye, R. R. Caldwell, and E. P. S. Shellard (1997), arXiv:9706013.

[8] A. Vilenkin and E. P. S. Shellard, *Cosmic Strings and Other Topological Defects* (Cambridge University Press, 2000).

[9] C. Caprini and D. G. Figueroa, Class. Quant. Grav. **35**, 163001 (2018).

[10] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, and R. Adhikari, Phys. Rev. Lett. **271102**, 1 (2011).

[11] B. P. Abbott, R. Abbott, T. D. Abbott, et al. (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. Lett. **118**, 121102 (2017).

[12] B. P. Abbott, R. Abbott, T. D. Abbott, et al. (LIGO Scientific Collaboration and Virgo) (2019), arXiv:1903.08844.

[13] A. I. Renzini and C. R. Contaldi, Phys. Rev. Lett. **122**, 081102 (2019).

[14] A. Renzini and C. Contaldi, Phys. Rev. **D100**, 063527 (2019).

[15] N. J. Cornish, Class. Quant. Grav. (2001), 0105374v1.

[16] N. Cornish and T. Robson, J. Phys. Conf. Ser. **840**, 012024 (2017), 1703.09858.

[17] N. Karnesis, A. Petiteau, and M. Lilley (2019), 1906.09027.

[18] C. Caprini, D. G. Figueroa, R. Flauger, G. Nardini, M. Peloso, M. Pieroni, A. Ricciardone, and G. Tasinato, JCAP **11**, 017 (2019), 1906.09244.

[19] T. L. Smith and R. R. Caldwell, Phys. Rev. D **100**, 104055 (2019), URL https://link.aps.org/doi/10.1103/PhysRevD.100.104055.

[20] H. Kudoh and A. Taruya, Phys. Rev. D **71**, 024025 (2005), gr-qc/0411017.

[21] J. D. Romano and N. J. Cornish, Living Reviews in Relativity **20**, 1 (2017), ISSN 14338351, 1608.06889.

[22] M. R. Adams and N. J. Cornish, Phys. Rev. D **82**, 022002 (2010), URL https://link.aps.org/doi/10.1103/PhysRevD.82.022002.

[23] M. R. Adams and N. J. Cornish, Phys. Rev. D **89**, 022001 (2014), URL https://link.aps.org/doi/10.1103/PhysRevD.89.022001.

[24] A. Taruya and H. Kudoh, Phys. Rev. D **72**, 104015 (2005), gr-qc/0507114.

[25] A. Taruya, Phys. Rev. D **74**, 104022 (2006), gr-qc/0607080.

[26] B. Allen and J. D. Romano, Phys. Rev. D **59**, 102001 (1999), ISSN 0556-2821, 9710117, URL http://arxiv.org/abs/gr-qc/9710117{%}5Cnhttp://link.aps.org/doi/10.1103/PhysRevD.59.102001.

[27] G. Cusin, C. Pitrou, and J.-P. Uzan, Phys. Rev. D **97**, 123527 (2018), 1711.11345.

[28] C. Conneely, A. H. Jaffe, and C. M. Mingarelli, Mon. Not. Roy. Astron. Soc. **487**, 562 (2019), 1808.05920.

[29] N. Bartolo, V. De Luca, G. Franciolini, A. Lewis, M. Peloso, and A. Riotto, Phys. Rev. Lett. **122**, 211301 (2019), 1810.12218.

[30] N. Bartolo, D. Bertacca, S. Matarrese, M. Peloso, A. Ricciardone, A. Riotto, and G. Tasinato (2019), 1912.09433.

[31] N. Bartolo, D. Bertacca, S. Matarrese, M. Peloso, A. Ricciardone, A. Riotto, and G. Tasinato, Phys. Rev. D **100**, 121501 (2019), 1908.00527.

[32] A. Margalit, C. R. Contaldi, and M. Pieroni (2020), 2004.01727.

[33] B. Allen and A. C. Ottewill, Phys. Rev. D **56**, 545 (1996).

[34] N. Christensen, Rept. Prog. Phys. **82**, 016903 (2019).

[35] A. I. Renzini and C. R. Contaldi, Mon. Not. Roy. Astron. Soc. **481**, 4650 (2018), 1806.11360.

[36] S. Kuroyanagi, T. Chiba, and T. Takahashi, JCAP **11**, 038 (2018), 1807.00786.

[37] Y. B. Ginat, V. Desjacques, R. Reischke, and H. B. Perets (2019), 1910.04587.

[38] A. Sesana, A. Vecchio, and C. N. Colacino, Monthly Notices of the Royal Astronomical Society **390**, 192 (2008).

[39] C. R. Contaldi, Phys. Lett. **B771**, 9 (2017).

[40] *LISA documents*, https://www.cosmos.esa.int/web/lisa/lisa-documents.

[41] M. Tinto and S. V. Dhurandhar, Living Rev. Rel. **8**, 4 (2005), gr-qc/0409034.

[42] M. Muratore, D. Vetrugno, and S. Vitale (2020), 2001.11221.

[43] T. A. Prince, M. Tinto, S. L. Larson, and J. Armstrong, Phys. Rev. D **66**, 122002 (2002), gr-qc/0209039.

[44] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, Astrophys. J. **622**, 759 (2005), astro-ph/0409513.

[45] *Characterization and detectability of stochastic gravitational wave background anisotropies with LISA*, CosWG project (2020), in preparation.

[46] J. Borrill, *Madcap - the microwave anisotropy dataset computational analysis package* (1999), astro-ph/9911389.

[47] J. R. Bond, A. H. Jaffe, and L. Knox, Phys. Rev. **D57**, 2117 (1998), astro-ph/9708203.

[48] B. P. Abbott, R. Abbott, T. D. Abbott, et al. (LIGO Scientific Collaboration and Virgo), Phys. Rev. Lett. **118**, 121101 (2017), [Erratum: Phys. Rev. Lett.119,no.2,029901(2017)].

[49] S. Drasco and E. E. Flanagan, Phys. Rev. **D67**, 082003 (2003), gr-qc/0210032.

[50] N. Seto, Phys. Rev. **D80**, 043003 (2009), 0908.0228.

[51] G. Cusin, I. Dvorkin, C. Pitrou, and J.-P. Uzan, Phys. Rev. Lett. **120**, 231101 (2018), 1803.03236.

[52] G. Nelemans, L. Yungelson, and S. F. Portegies Zwart, Astron. Astrophys. **375**, 890 (2001), astro-ph/0105221.

[53] V. Korol, E. M. Rossi, P. J. Groot, G. Nelemans, S. Toonen, and A. G. A. Brown, Mon. Not. Roy. Astron. Soc.

470, 1894 (2017), 1703.02555.

[54] Z.-C. Chen, F. Huang, and Q.-G. Huang, Astrophys. J. **871**, 97 (2019), 1809.10360.

[55] N. Bartolo et al., JCAP **12**, 026 (2016), 1610.06481.

[56] A. Ricciardone and G. Tasinato, JCAP **1802**, 011 (2018), 1711.02635.

[57] E. Dimastrogiovanni, M. Fasiello, and G. Tasinato, Phys. Rev. Lett. **124**, 061302 (2020), 1906.07204.

[58] P. Auclair et al., JCAP **04**, 034 (2020), 1909.00819.

[59] C. Caprini et al., JCAP **03**, 024 (2020), 1910.13125.

[60] M. Pieroni and E. Barausse (2020), 2004.01135.