

Probing the Universe with Gravitational Waves

Author: José Francisco Nuño Siles
Supervisor: Juan García-Bellido Capdevila

A thesis presented for the degree of
Doctor of Philosophy



Department of Theoretical Physics
Universidad Autónoma de Madrid
Spain

Contents

I. Acknowledgments	4
II. Prologue	5
III. Bayesian analysis of the spin distribution of LIGO/Virgo black holes	7
A. Journal & publication details	7
B. Authors	7
C. Abstract	7
D. Introduction	7
E. Spin observables from BBH events	9
1. The effective spin	9
2. The final spin	9
F. Bayesian Population Analysis	11
1. Likelihoods for LIGO/Virgo BBH events	11
2. Priors on the BBH population	13
G. Bayes Factors	13
H. Bayesian hierarchical modeling	16
I. Conclusions	16
IV. The False Alarms Induced by Gaussian Noise in Gravitational Wave Detectors	17
A. Journal & publication details	17
B. Authors	17
C. Abstract	17
D. Introduction	17
E. The false alarm rate of a template in a single detector	18
1. Probabilistic derivation of the FAR	20
2. Evaluation of the FAR of a template	22
3. NLO corrections to the FAR of a template	25
4. Dependence on the CBC template parameters	26
F. The false alarm rate of a template in a network of detectors	28
G. Application to GW events	30
1. Application to GW candidates in GWTC-3	31
H. Conclusions	35
I. Appendix: Study of the FAP for the bivariate complex Gaussian	36
V. Efficient Reduced Order Quadrature Construction Algorithms for Fast Gravitational Wave Inference	41
A. Journal & publication details	41
B. Authors	41
C. Abstract	41
D. Introduction	42
E. Theoretical framework	42
1. A primer on gravitational wave inference	43
2. Basics of Reduced Order Quadratures for Gravitational Wave inference	43
F. Efficient algorithm for ROQ computation	45
1. Reduced Order Basis	45
2. Empirical Interpolation Model	48
3. Construction of the ROQ	54
G. Code Validation	56
1. Basis Generation and comparison with other ROQ methods	56
2. Statistical tests	58
3. Speedup analysis	60
4. Application to GW events	63
H. Conclusions	65
I. Appendix: Fast way to update $\ \hat{A}^{-1}\ _F$ and $\sigma_{\text{EIM}}^{\text{tot}}$	66

VI. Analysis of a subsolar-mass compact binary candidate from the second observing run of Advanced LIGO	69
A. Journal & publication details	69
B. Authors	69
C. Abstract	69
D. Introduction	70
E. Significance of SSM170401	70
F. Properties of the source of SSM170401	71
1. Coherence Test	74
G. Discussion	75
H. Conclusions	77
VII. Primordial Black Hole clusters, phenomenology & implications	78
A. Journal & publication details	78
B. Authors	78
C. Abstract	78
D. Introduction	78
E. Methodology and initial mass function	79
F. Dynamics	80
G. BBH mergers	82
H. Merger count	83
I. Hierarchical mergers	84
J. Merger Rate	85
K. Mass distribution	86
L. Spin distribution	87
M. Residual eccentricity distribution	87
N. Conclusions	88
VIII. Gravitational Energy and Angular Momentum Loss in CHE	90
A. Introduction	90
B. Considerations & derivation	90
IX. Thesis Conclusion	93
References	93

I Acknowledgments

I want to thank my thesis supervisor, Dr. Juan García-Bellido, for believing in me from the beginning and encouraging me to make the correct decisions in the most critical of times, leaving aside emotions and rooting for my best fate. I would also like to thank Dr. Ester Ruiz Morales, a vital part of the team, for our endless talks about physics. Both of them and my colleague, soon-to-be PhD, Gonzalo Morrás, represent the pillars on which I stood throughout the thesis. Without them and their extraordinary abilities both as physicists but also as humans, I would not be who I have become. I am eternally grateful to have coincided in our paths.

I would also like to acknowledge monetary support from the research project CEX2020-001007-S-20-6 and FPI grant PRE2020-092571, the Centro de Excelencia Severo Ochoa Program SEV-2016-0597-20-2 and express my gratitude to all my professors and the staff at the Universidad Autónoma de Madrid (UAM) and the Centro Superior de Investigaciones Científicas (CSIC). I would like to make a specific mention to all the lecturers that represented an intellectual challenge during my studies at the University. It was thanks to them that I learned to be humble and appreciate being surrounded by people vastly more intelligent than myself. I would also like to thank my physics teacher from high school studies, Santiago Coronado, as he planted the seed that led me to study physics.

I would also like to acknowledge the computational resources utilized to carry out this work, including the Hydra cluster at the IFT, the LDAS cluster at Caltech, and the ARCCA Hawk Computing Centre at Cardiff. All of this numerical power would be useless without the proper code to take advantage of it, that is why I would like to thank all the developers whose open-source projects I have used throughout these years.

I would also like to mention my office colleague, also soon-to-be PhD, Santiago Jaraba, for the countless hours working together in the office and the big and vivid Cosmology group in the IFT for so many interesting Journal Clubs and seminars. Such a collaborative environment is the seedbed of the most brilliant ideas. Outside of my home research center, I would like to mention Dr. Rainer Spurzem for having welcomed me as a collaborator in the most important N-body simulations code and introducing me to the wider community in that field. I would also like to thank many collaborators in the LVK Collaboration for having guided me and answered my doubts, with particular mention to Dr. Sebastian Clesse and Dr. Christopher Berry.

This thesis is the most in-depth and curated piece of work I have ever written, and as such, I would like to dedicate it to my family, my mum Elvira, my dad Jose Antonio, and my aunt María Dolores, for their tireless understanding and attention. I would also like to dedicate it to my girlfriend, Iciar, and her sincere enjoyment when listening to the ideas behind this work. Lastly, I would like to add my friends, without whom the breaks between research sessions would not have been as entertaining and fruitful.

II Prologue

Gravitational wave (GW) astronomy has ushered in a new era of understanding the Cosmos. It offers a unique window into the most extreme events in the universe and deepens our understanding of its main actors, Black Holes. Beginning with an exploration of the origin and nature of the main GW source compact objects that includes two pieces of work, a population-informed study of the spin distribution of the binary black holes detected by the LIGO/Virgo collaboration in GWTC-1 and O3 Discovery Papers and a phenomenological investigation of dense Black Holes-only clusters with practical implications for future simulations. Following the cosmic trail to their Earthly arrival, we shift the focus to the detectors and present an analysis that characterizes the Gaussian noise intrinsic to any interferometer and the statistics to quantify the significance of any detection. We conclude the research included in this thesis with two studies focusing on parameter estimation, our best method of describing a given detector signal under the hypothesis of compact binary coalescence (CBC) origin. The first work emphasizes modern techniques used in the computation of the posterior distributions of the parameters, improving the current speed and flexibility of the algorithms. The second one presents an exhaustive analysis of an exotic sub-solar mass candidate from the second observing run of LIGO/Virgo. Together, these studies contribute to a narrative that not only refines our understanding of astrophysical phenomena but also poses intriguing questions about the nature of dark matter and the BH populations in our Universe, with a focus on the possibility of a primordial origin. Now we delve into the main results of all the research works included in this thesis.

In the first of the original works composing this thesis [1], we study the spin distribution of the underlying population of binary black holes detected by LIGO/Virgo, leveraging data from the GWTC-1 catalogue and the initial four binary black hole events from run O3. Using the Bayesian inference framework and four phenomenologically informed spin priors we find compelling population evidence for low spin magnitude and a preference for small and isotropic individual spins. These results support the idea that LIGO/Virgo black holes are primordial in origin and prompted us to pursue research in this line.

In the second article [2], we focus on gravitational wave interferometers, representing our best efforts to understand the arrival of a gravitational wave signal to Earth and its detection. We perform an in-depth theoretical study of stationary Gaussian noise, which represents an irreducible component of the background noise in any interferometer. We find that, even though it is uncorrelated in frequencies, there is an important correlation in time that depends on the template and the shape of the noise power spectral density (PSD).

Acknowledging this result, we compute from first principles the false alarm rate (FAR) of a gravitational wave template in Gaussian noise, which is defined as the number of occurrences per unit of time that the template's matched filter signal-to-noise ratio (SNR) with the noise goes over a threshold ρ . We find that this theoretical FAR that we call Gaussian FAR can be well approximated by the usual expression for uncorrelated noise if we replace the sampling rate with an effective sampling rate that depends on the parameters of the template, the noise PSD and the threshold ρ . This allows us to define a minimum SNR threshold that our catalogue of events needs to meet if we want to keep events generated from Gaussian noise below a certain FAR.

As a final result of this publication, we apply our theoretical formulas to the GWTC-3 events. Under the assumption that the multidimensional posterior distribution function (PDF) for the parameters as sampled via parameter estimation represents our best characterization of a given candidate, we can define a FAR per posterior sample. This is translated into a false alarm probability (FAP) for the candidate when considering the entire distribution. Using this, we find that Gaussian noise fluctuations could be the source of GW200308_173609 and GW200322_091133, rather than being true GW events.

The third published work [3] introduces new methods for gravitational wave data analysis, focusing on the development and optimisation of Reduced Order Quadrature (ROQ) techniques. ROQ methods can greatly reduce the computational cost of GW likelihood evaluations, and therefore speed up parameter estimation analyses. We improve upon previous ROQ construction algorithms allowing for more efficient bases in regions of parameter space that were previously challenging. In particular, we use singular value decomposition (SVD) methods to characterize the waveform space and choose a reduced-order basis close to optimal. We also propose improved algorithms for the selection of empirical interpolation nodes, greatly reducing the error added by the empirical interpolation model.

The main part in which I contributed was in the computational calculations and simulations using the constructed ROQ bases ranging from 4s to 256s for CBC waveforms including precession and higher order modes. We performed likelihood error tests and P-P tests to validate them. The likelihood tests consist of a comparison of the log-likelihoods evaluated using the standard waveform with those obtained using the ROQ approximation, both calculated over the same injected signals in Gaussian noise. The percent-percent or P-P tests represent a simulation-based validation procedure where, again, we inject several signals with random parameters inside the range of validity of our ROQ bases and perform parameter estimation analysis on them. We then plot the fraction of posterior distributions for which the injected (true) value of the parameters is found in a given confidence interval as a function of that same confidence interval [4, 5]. We expect this fraction to be drawn from a uniform distribution, which allows us to assign

a p-value to the test. We also study their speed-up both theoretically and empirically with very positive results. Finally, end-to-end parameter estimation analyses were also carried out on real events, including GW170817 with the waveform **IMRPhenomXPHM**, which represents the first time such a complex waveform has been used for such a long signal duration. With these tests and the comparison of the sampled posterior distributions to those obtained via standard methods, we validate our bases and their fidelity in producing precise and unbiased Parameter Estimations in real gravitational wave detector data.

The fourth article [6] delves into the intricacies of Bayesian parameter estimation and its specific application to a subsolar-mass gravitational wave candidate reported by Phukon et al. [7] from the second observing run of Advanced LIGO. This candidate has a reported signal-to-noise ratio (SNR) of 8.6 and a false alarm rate of 0.41 yr^{-1} which are too low to claim a clear gravitational wave origin. When improving on the search analysis by using more accurate waveforms, extending the frequency range from 45 Hz down to 20 Hz, and removing a prominent blip glitch in the strain around the time of the trigger, we find that the posterior distribution of the network SNR lies mostly below the search value, with the 90% confidence interval being $7.94^{+0.70}_{-1.05}$. We find that, under the CBC hypothesis, the secondary component has $m_2 = 0.76^{+0.50}_{-0.14} M_\odot$, with $m_2 < 1 M_\odot$ at 84% confidence level, suggesting an unexpectedly light neutron star or a black hole of primordial or exotic origin. The primary mass would be $m_1 = 4.71^{+1.57}_{-2.18} M_\odot$, likely in the hypothesized lower mass gap and the luminosity distance is measured to be $D_L = 124^{+82}_{-48} \text{ Mpc}$. We then perform a coherence test that compares the evidence between analysing the signal found in the interferometers in a coherent vs. incoherent way, obtaining a log Bayes factor of 4.96 ± 0.13 , which favours the existence of a coherent signal in the Hanford and Livingston interferometers.

With this paper, we set precedence and push the limits of current parameter estimation techniques towards the subsolar region of the parameter space by using real data for an SSM candidate with moderate SNR. The improved sensitivity of O4 and subsequent LIGO-Virgo-KAGRA observing runs could make it possible to observe similar signals, if present in the Cosmos, with a higher SNR and thus, more precise measurement of the parameters of the binary.

The final work [8] represents a general endeavour of employing direct N-body simulations to explore the evolution of black holes-only clusters. These are theorised to exist in our Universe by the primordial Black Holes paradigm, among other scenarios, and aim at explaining some of the dark matter and early Universe conundrums. With a focus on compact objects of primordial origin, we present the dynamical evolution of several mass and spatial distributions under the sole effect of Newtonian gravitational forces. The results highlight the critical role that relativistic black hole merger kicks play in the conceivable genealogical trees that can originate inside these clusters, abruptly diminishing the probabilities of higher-generation mergers that are widely assumed to be a viable route for producing some of the most massive events detected by the LIGO/Virgo interferometers. They also offer a glimpse into the rich and varied evolutionary dynamics of the clusters, their stability, the binary merger rate, and the spin distribution of the remnants.

On a final but complementary note in this topic, we also present here the derivation under the quadrupole emission approximation of the energy and angular momentum loss as well as the eccentricity and semimajor axes' evolution with respect to time for a two-body system following a hyperbolic trajectory under the influence of gravity. Although these results are partially derived already in [9], the novelty here is the accommodation of the final equations into their infinitesimal form for direct implementation in the state-of-the-art direct integration gravitational N-body code **Nbody6++GPU** [10] and any code alike.

Finally, it is worth mentioning two papers produced by the LVK Collaboration in which notable contributions were made during the development of this thesis. The first one is the GWTC-3 Catalog [11], in which numerous exploratory parameter estimation analyses were carried out on behalf of the Collaboration for selected candidates. The second concerns the official LVK search for subsolar-mass black hole binaries using the data acquired in the second part of Advanced LIGO and Advanced Virgo's third observing run [12]. The main contribution was, again, the exploratory parameter estimation runs performed on the most significant candidates, which planted the seed of other future work [13].

On the software side, it is worth mentioning a direct contribution to developing the latest upgrade of the widely implemented Bayesian GW framework **LAL** [14], mainly in the Python part.

III Bayesian analysis of the spin distribution of LIGO/Virgo black holes

A. Journal & publication details

Published in Physics of the Dark Universe 31 (2021) 100791 [1]

B. Authors

The authors in the same order as presented to the journal are

Juan García-Bellido^{a1}

Instituto de Física Teórica UAM-CSIC, Universidad Autonóma de Madrid, Cantoblanco, 28049 Madrid, Spain

José Francisco Nuño Siles^{a2}

Instituto de Física Teórica UAM-CSIC, Universidad Autonóma de Madrid, Cantoblanco, 28049 Madrid, Spain

Ester Ruiz Morales^{b3}

Departamento de Física, ETSIDI, Universidad Politécnica de Madrid, 28012 Madrid, Spain

C. Abstract

Gravitational wave detection from binary black hole (BBH) inspirals has become routine thanks to the LIGO/Virgo interferometers. The nature of these black holes remains uncertain. We study here the spin distributions of LIGO/Virgo black holes from the first catalogue GWTC-1 and the first four published BBH events from run O3. We compute the Bayes evidence for several independent priors: flat, isotropic, spin-aligned and anti-aligned. We find strong evidence for low spins in all of the cases, and significant evidence for small isotropic spins versus any other distribution. When considered as a homogeneous population of black holes, these results give support to the idea that LIGO/Virgo black holes are primordial.

D. Introduction

The regular detection with laser interferometers [15–19] of gravitational wave (GW) events from BBH mergers has opened a new window into the universe, and in particular to the exploration of the nature of black holes.

Before the first BBH detections by LIGO, stellar black holes with masses in the range $5 - 15 M_{\odot}$ had been detected as components of X-ray binaries, only a few intermediate mass black holes (IMBH) were known with masses above $100 M_{\odot}$, while supermassive black holes (SMBH) were known to exist at the centers of all galaxies. The origin of black holes in such a large range of masses remains a mystery, and one fascinating possibility is that part of these black holes are primordial in origin [20, 21]. In fact, soon after the first detection of black hole mergers by LIGO [22], there were claims of their primordial nature [23–25]. Since then, the best scenario consistent with all observational constraints so far [26] is that of spatially clustered and broad-mass distributed primordial black holes [21, 27].

As the number of binary black hole merger events detected with GW interferometers increases, a new population of black holes is emerging with unexpected properties in terms of their masses and spins. These properties significantly differ from previous black holes detected through X-rays, via stellar dynamics around SMBH at the center of galaxies, or via gravitational lensing effects for IMBH. When the range of masses and distances accessible by the GW interferometers improved, events like GW190425 and GW190521 appeared with BH masses in the lower and upper mass gaps, challenging existing astrophysical BH formation models. Moreover, events with small mass ratios $q \ll 1$, like GW190814, are also difficult to generate in stellar binary formation models, due to the expected mass transfer among the binary components, see however [28].

One of the most striking features of this new population of black holes detected in GW events is that they all seem to have a small spin. Although the individual spin of each of the black holes in the binary is poorly constrained,⁴

¹ juan.garciaellido@uam.es

² josef.nunno@estudiante.uam.es

³ ester.ruiz.morales@upm.es

⁴ It is worth pointing out, however, that the best determined spin for a single black hole is that of the massive companion of GW190814, with $S_1 < 0.07$ at 90% c.l., thanks to its particularly low mass ratio $m_2/m_1 = 0.11$.

a derived quantity called the *effective* spin, χ_{eff} , can be well inferred from the GW waveform. The O1+O2 events [15] plus the four run O3 published events [16–19] show that, in almost all the BBHs mergers events, the inferred χ_{eff} posteriors cluster around zero and are narrowly peaked. This single observation already constrains some BBH formation models, because even if it is possible to generate a single merger event with very small χ_{eff} in almost all physical scenarios of BBH formation, the fact that *the whole population* of BBH mergers have very small χ_{eff} cannot be explained by models involving high aligned spins of the two BH's in the binary.

Even before the first BBH merger detections, the effective spin had already been considered an optimal variable to discriminate among formation channels. By simulating populations of BBH mergers with different intrinsic spin values and spin alignments, the authors in [29, 30] showed that the effective spin could be relevant to distinguish the astrophysical environment in which these binary systems formed, since binaries evolved by dynamical interactions were expected to have their spins isotropically oriented, while those coming from the evolution of an isolated binary star system were more likely to have their spins highly aligned (see however [31] for updates on these models.).

The magnitude of the intrinsic spin of each black hole in the binary, though much poorly constrained from the data than the effective spin, is important to determine how these black holes were formed, and in particular to distinguish a possible stellar versus primordial origin. Stellar black holes are expected to have nonzero spin, due to conservation of angular momentum. On the contrary, since the size of their Schwarzschild radius is essentially identical to the size of the causal horizon at the moment of their formation, primordial black holes are all expected to have zero or very near zero spin at formation [32], although subsequent accretion could enhance it slightly [33].

On the other hand, it has been argued that the more massive LIGO/Virgo BH could originate from previous mergers [34, 35]. However, it is known that the spin distribution of second generation black holes is peaked very far from zero, near $S \sim 0.686$ [36], and this is in disagreement with the fact that most of the massive BH in LIGO/Virgo seem to have very low spin. In a careful analysis, Ref [37] showed that GWTC-1 BBH catalog is consistent with having no hierarchical mergers.

When the first detections of the O1 run became available, Farr et al. [38, 39] analysed the discriminating power of the χ_{eff} distribution with the first four GW LIGO/Virgo events, assuming equal masses ($q = m_2/m_1 = 1$) for all events and approximating the posterior χ_{eff} distributions by Gaussians. They compared the odds ratios for several models of spin (modulus and orientations) of the underlying BH population, finding a preference for either a population with an isotropic spin distribution or with low intrinsic spin values of the merging black holes. A more realistic analysis of the first six LIGO/Virgo mergers was made in [40], including Bayesian methods and taking into account the $q - \chi_{\text{eff}}$ correlations, reaching the general conclusion that highly spinning black holes were disfavoured against low spins, see also [41].

With the publication of the GWTC-1 catalog, the LIGO/Virgo collaboration (LVC) made a population analysis of the mass, redshift and spin properties of the ten BBH mergers detected in O1+O2 runs [42].⁵ Under different assumptions for the parameters of the population models, they observed a common trend of the inferred distribution for the BH spins to decrease with increasing spin magnitude, but were not able to place strong constraints on spin orientations, concluding that black hole spin measurements were not informative enough at that moment to discern between isotropic and aligned orientation distribution via χ_{eff} .

In this work, we improve such discriminating power by incorporating the projection along the total angular momentum of the spin of the final BH formed after the merger a_f as another variable of the BBH merger population analysis. The effective spin and the final spin can be measured independently, since χ_{eff} is inferred from the inspiral part of the waveform while a_f can be determined also from the ringdown part. We follow a multivariate approach keeping all correlations among the variables ($q, \chi_{\text{eff}}, a_f$) of each event, and show that this significantly improves the Bayesian evidence for an underlying population of black hole components with small spin magnitude and isotropic orientation.

There are still many uncertainties on the full spin (magnitude and orientation) distribution of the different BBH formation channels, since the computational models used to predict these properties depend on many assumptions about poorly understood environmental conditions related to the formation and evolution of the binaries. There are also many unknown aspects related to the mass, spatial distribution and dynamics of BBH of primordial origin. For this reason, instead of considering an specific astrophysical or primordial model for the mass and spin distributions, we have chosen general hypothesis on the underlying distributions of black hole spin magnitudes and orientations and compare the different hypothesis in a full Bayesian analysis.

In Section 1, we describe the spin variables and parameters that will be considered in our analysis of the BBH population. In Section 2, we describe the method employed in calculating the Bayesian evidence from the published parameter estimation samples, and define the population hypotheses together with their priors for the parameters

⁵ A hierarchical Bayesian analysis of spin distributions was also performed in [43]. Another Bayesian analysis, studying mass distributions and merger rates but not spins, was done in [44].

considered in our analysis. In Section 3 we use Bayesian methods to obtain the evidence for each hypothesis and the corresponding Bayes ratios. Section 4 is devoted to the hierarchical modelling method that allows us to infer a posterior distribution for the spin magnitude of the underlying BH population. In Section 5 we give our conclusions.

E. Spin observables from BBH events

In this section we will describe the observables relevant to the analysis of spins in LVC black hole binaries. The main observables that we will consider are the weighted-averaged effective projected spin, χ_{eff} , and the the projection along the total angular momentum of the final spin a_f . We could have also chosen the effective precession spin parameter, χ_p , but this quantity is much worse measured in the LVC events published so far and is mainly determined by the prior.

1. The effective spin

A derived physical quantity that can be measured very well by LIGO from the waveform templates is the mass-averaged (so called effective) projected spin,

$$\chi_{\text{eff}} = \frac{m_1 a_1 + m_2 a_2}{m_1 + m_2} = \frac{a_1 + q a_2}{1 + q}, \quad (1)$$

where $a_i \equiv (\vec{S}_i \cdot \vec{L})/(|\vec{L}| m_i^2)$ is the individual mass-weighted projection of each black hole spin onto the orbital angular momentum \vec{L} of the binary. Thus, χ_{eff} gives some information about the spin orientations of the inspiralling black holes w.r.t. the orbital angular momentum. Here $q = m_2/m_1$ is the binary mass ratio, in the range $q \in [0, 1]$.

A derived quantity that will be useful is the so-called “chirp mass”,

$$\mathcal{M}_c = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}} = m_1 \frac{q^{3/5}}{(1 + q)^{1/5}}, \quad (2)$$

which gives the mass ratio

$$q = \frac{3\alpha (2/3)^{2/3} + (9\alpha + \sqrt{81\alpha^2 - 12\alpha^3})^{2/3}}{3 (2/3)^{1/3} (9\alpha + \sqrt{81\alpha^2 - 12\alpha^3})^{1/3}}, \quad (3)$$

where $\alpha = (\mathcal{M}_c/m_1)^5$.

2. The final spin

The second best-measured spin-related quantity in LIGO/Virgo binaries is the projection along the total angular momentum of the spin of the final black hole after merging, a_f . We use the approximate expressions given in [45] for a_f in the particular case of spinning but non-precessing black hole binaries. We first define some quantities and then we assemble everything together.

First we need to define a_{tot} and a_{eff} in terms of the *projected* spins, $a_i \equiv (\vec{S}_i \cdot \vec{L})/(|\vec{L}| m_i^2)$,

$$a_{\text{tot}} = \frac{a_1 + q^2 a_2}{(1 + q)^2}, \quad a_{\text{eff}} = a_{\text{tot}} + \xi \nu (a_1 + a_2), \quad (4)$$

$$\text{with} \quad \nu(q) = \frac{q}{(1 + q)^2}, \quad \xi = 0.474046.$$

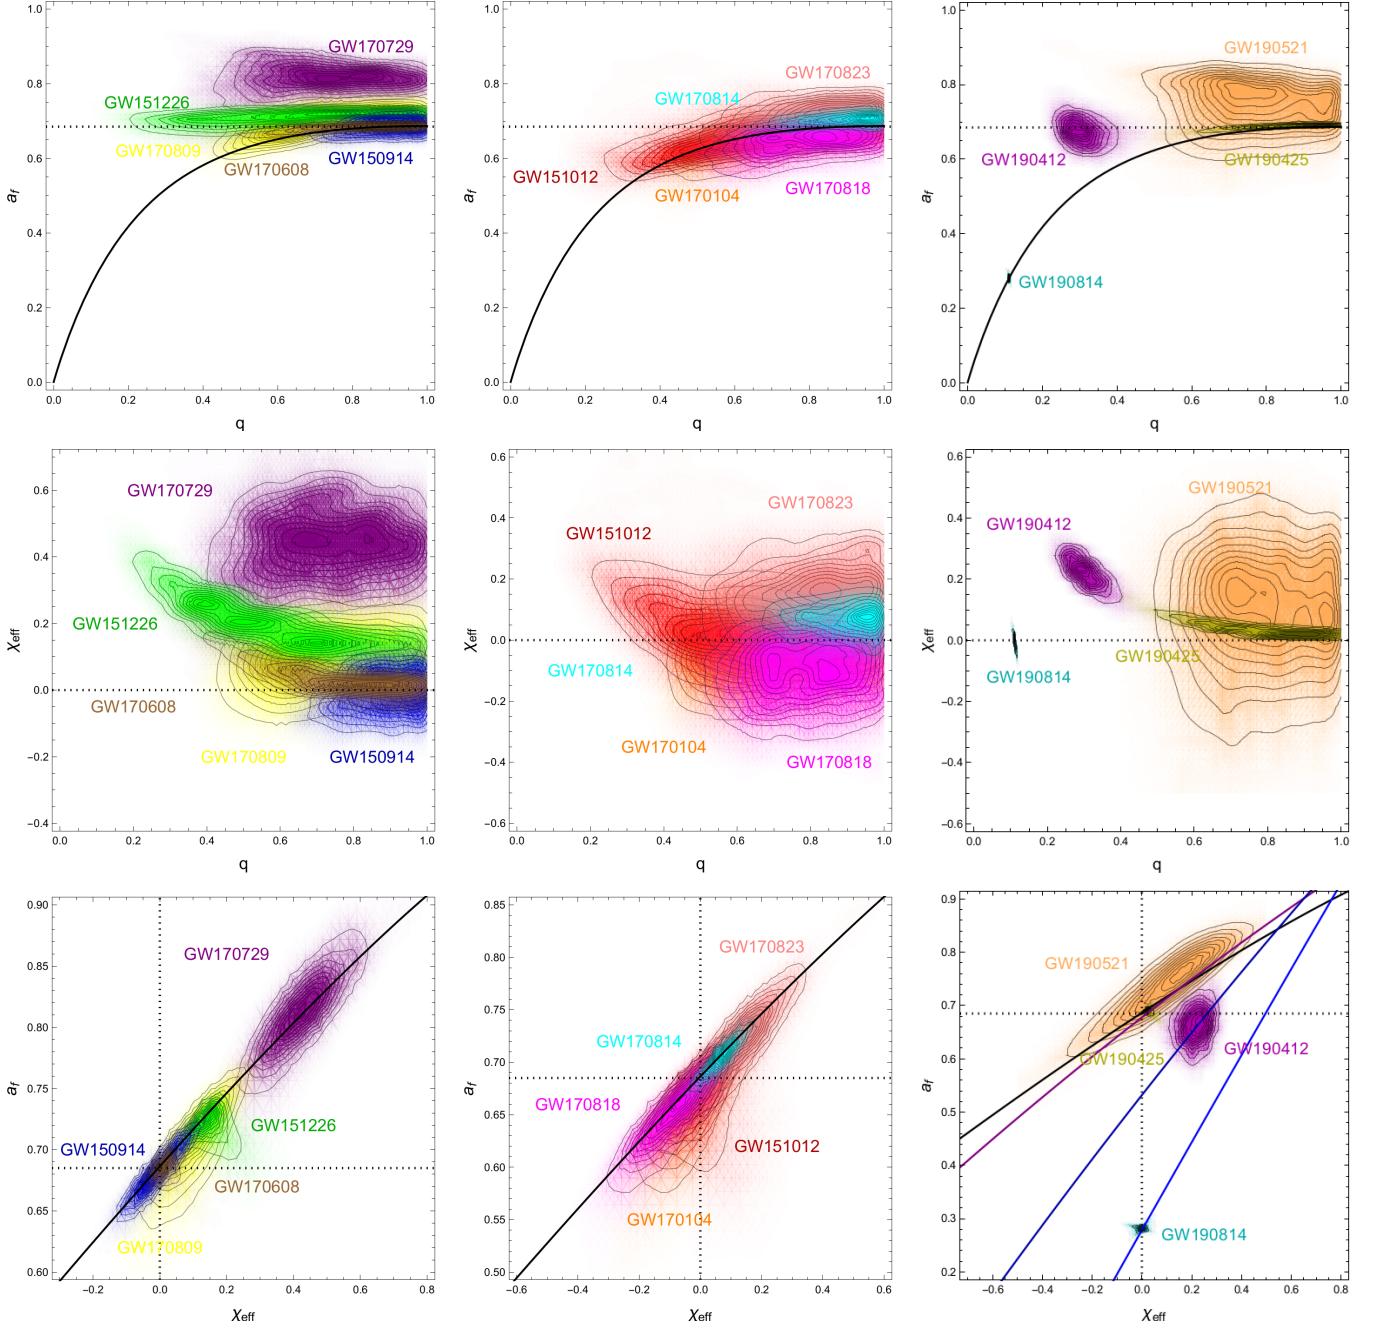


FIG. 1: The multidimensional LVC likelihoods projected on the planes (q, a_f) (top row), (q, χ_{eff}) (middle row) and (χ_{eff}, a_f) (bottom row) for the ten GWTC-1 events of LIGO/Virgo (left and middle column) plus the four runO3 events (right column). The dotted lines correspond to the fixed points $\chi_{\text{eff}} = 0$ and $a_f = 0.686$. The curves on the (χ_{eff}, a_f) plane correspond to $q = 1$ (black), $q = 0.75$ (purple), $q = 0.4$ (dark blue) and $q = 0.11$ (light blue), respectively.

-5.97723	3.39221	4.48865	-5.77101	-13.0459
35.1278	-72.9336	-86.0036	93.7371	200.975
-146.822	387.184	447.009	-467.383	-884.339
223.911	-648.502	-697.177	753.738	1166.89

TABLE I: The coefficients k_{ij} for $i \in [0, 3]$ and $j \in [0, 4]$.

We then have to define the energy, angular momentum and size of the Innermost Stable Circular Orbit (ISCO),

$$E_{\text{ISCO}}(a) = \sqrt{1 - 2/(3r_{\text{ISCO}}(a))}, \quad (5)$$

$$L_{\text{ISCO}}(a) = \frac{2}{3\sqrt{3}} \left(1 + 2\sqrt{3r_{\text{ISCO}}(a) - 2} \right), \quad (6)$$

$$r_{\text{ISCO}}(a) = 3 + Z_2 - \text{sign}(a) \sqrt{(3 - Z_1)(3 + Z_1 + 2Z_2)},$$

$$Z_1(a) = 1 + (1 - a^2)^{1/3} \left((1 + a)^{1/3} + (1 - a)^{1/3} \right), \quad (7)$$

$$Z_2(a) = \sqrt{3a^2 + Z_1^2}. \quad (8)$$

Then the final spin is given by (the coefficients k_{ij} can be found in Table I)

$$\begin{aligned} a_f &= a_{\text{tot}} + \nu \left(L_{\text{ISCO}}(a_{\text{eff}}) - 2a_{\text{tot}}(E_{\text{ISCO}}(a_{\text{eff}}) - 1) \right) \left(\right. \\ &\quad + \nu^2 (k_{00} + k_{01}a_{\text{eff}} + k_{02}a_{\text{eff}}^2 + k_{03}a_{\text{eff}}^3 + k_{04}a_{\text{eff}}^4) \\ &\quad + \nu^3 (k_{10} + k_{11}a_{\text{eff}} + k_{12}a_{\text{eff}}^2 + k_{13}a_{\text{eff}}^3 + k_{14}a_{\text{eff}}^4) \\ &\quad + \nu^4 (k_{20} + k_{21}a_{\text{eff}} + k_{22}a_{\text{eff}}^2 + k_{23}a_{\text{eff}}^3 + k_{24}a_{\text{eff}}^4) \\ &\quad \left. + \nu^5 (k_{30} + k_{31}a_{\text{eff}} + k_{32}a_{\text{eff}}^2 + k_{33}a_{\text{eff}}^3 + k_{34}a_{\text{eff}}^4) \right). \end{aligned}$$

F. Bayesian Population Analysis

The Bayes theorem relates the likelihood $\mathcal{L}(d_j|\theta)$ of the data d_j for a given set of parameters θ , with the posterior probability of the parameters given the data, $\mathcal{P}(\theta|d_j)$, via the prior knowledge about the parameters of the population model i , $\Pi_i(\theta)$,

$$\begin{aligned} \mathcal{P}_i(\theta|d_j) &= \frac{\mathcal{L}(d_j|\theta) \Pi_i(\theta)}{E(d_j)}, \\ E_{ij}(d) &= \iint d\theta \mathcal{L}(d_j|\theta) \Pi_i(\theta), \end{aligned} \quad (9)$$

where $E_{ij}(d)$ is the Bayesian evidence for the data d_j and the population model i . Here we will consider different priors, $\Pi_i(\theta)$, for the distribution of the parameters in each of the spin population models characterized by the hypothesis H_i . We first compute the multidimensional likelihood for each BBH event from the LIGO/Virgo published samples, and then calculate the priors associated with each population hypothesis to derive the evidence (9).

1. Likelihoods for LIGO/Virgo BBH events

We calculate the multidimensional likelihoods (i.e. including correlations) from the posteriors and priors provided by the LIGO public documentation page.⁶ We use the parameter estimation samples for each event, which are given in terms of the fundamental parameters $(m_1, m_2, s_1, s_2, \cos \theta_1, \cos \theta_2)$, and using the expressions of χ_{eff} in (1) and a_f

⁶ <https://dcc.ligo.org/cgi-bin/DocDB/DocumentDatabase>

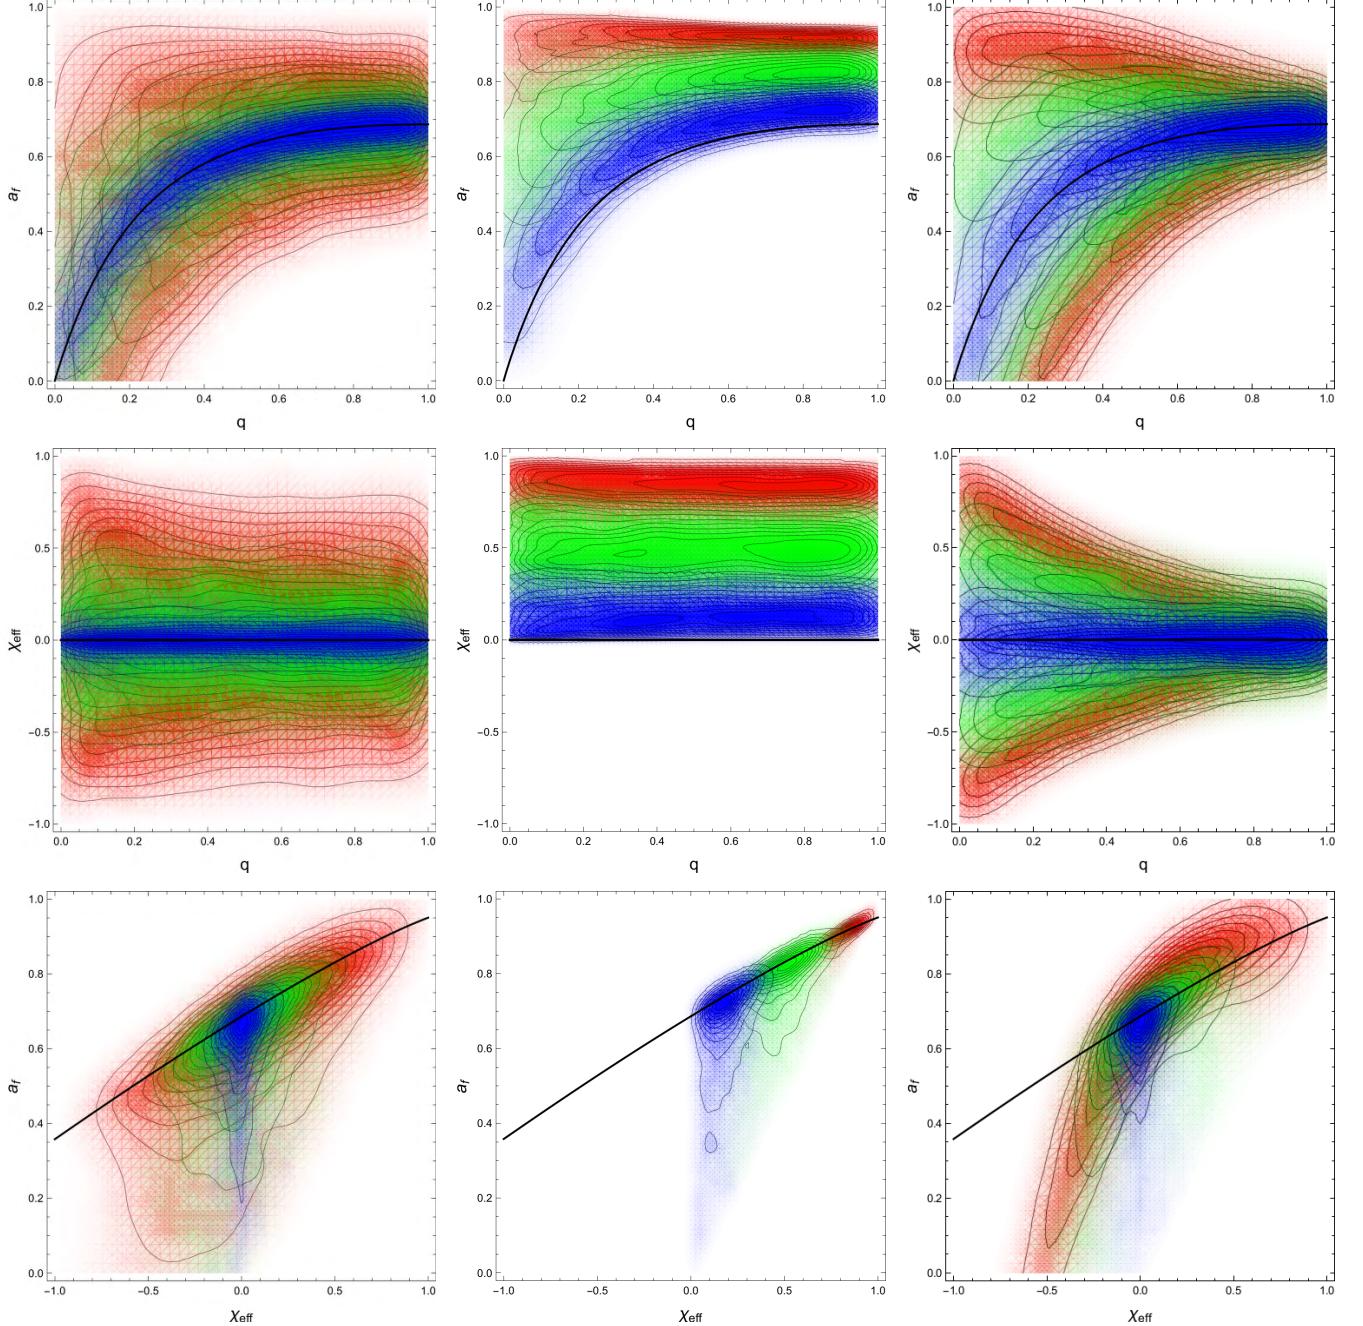


FIG. 2: The priors in the planes (q, a_f) (top), (q, χ_{eff}) (middle) and (χ_{eff}, a_f) (bottom) for the three hypothesis $H_{i=2,3,4}$: isotropic spin (left), aligned spin (middle) and anti-aligned spin (right), for spins centered at $\mu = 0, 0.5$ and 1 (blue, green and red resp.).

in (9) for each sample point, we construct the full multidimensional posterior and prior distributions for the derived parameters $\theta = \{q, \chi_{\text{eff}}, a_f\}$, marginalizing over the other parameters. Finally we generate the multidimensional likelihoods $\mathcal{L}_{\text{LVC}}(d_j|\theta)$ dividing the posteriors $\mathcal{P}_{\text{LVC}}(d_j|\theta)$ by their corresponding priors $\Pi_{\text{LVC}}(\theta)$. We have plotted the LVC likelihoods for the three parameters $\theta = \{q, \chi_{\text{eff}}, a_f\}$ in Fig. 1.

When using these likelihoods to infer population properties in our Bayesian analysis, we will not include possible selection effects in the $(\chi_{\text{eff}}, a_f, q)$ variables, like the observational bias towards positive χ_{eff} values described in [46], or possible effects from the fact that the waveform bank only considers aligned spins. We expect these effects to be small given the spin variables used, but should eventually be considered in a future analysis.

2. Priors on the BBH population

Once we have chosen to use the subset of parameters $\theta = \{q, \chi_{\text{eff}}, a_f\}$ for the Bayesian inference, we need to specify the different hypotheses on the spin distribution of the merging BBH population. We have chosen 5 different hypotheses H_i on the spin magnitude and orientation, which in turn depend on “hyperparameters” $\Lambda = \{\sigma, \mu\}$ that determine the width and the mean of the distribution of the spin magnitude.

These five basic Hypotheses are:

- H_0 : isotropic spin orientation, $\cos \theta_{LS_i} \in [-1, 1]$, flat prior on spin magnitude in the range $S_i \in [0, 1]$, flat prior on mass ratio, $q \in [0, 1]$. This is the “null” hypothesis.
- H_1 : isotropic spin orientation, $\cos \theta_{LS_i} \in [-1, 1]$, Gaussian prior on spin magnitude with $\mu = 0$ and $\sigma \in [0, 1]$, flat prior on mass ratio, $q \in [0, 1]$.
- H_2 : isotropic spin orientation, $\cos \theta_{LS_i} \in [-1, 1]$, Gaussian prior on spin magnitude with $\sigma = 0.2$ and $\mu \in [0, 1]$, flat prior on mass ratio, $q \in [0, 1]$.
- H_3 : aligned spin orientation, $\cos \theta_{LS_i}$ sampled from Gaussian centered at $+1$ and width 0.05 , Gaussian prior on spin magnitude for $\sigma = 0.2$ and $\mu \in [0, 1]$, flat prior on mass ratio, $q \in [0, 1]$.
- H_4 : anti-aligned spin orientation, $\cos \theta_{LS_i}$ sampled from Gaussian centered at ± 1 and width 0.05 , Gaussian prior on spin magnitude for $\sigma = 0.2$ and $\mu \in [0, 1]$, flat prior on mass ratio, $q \in [0, 1]$.

We generate the multidimensional prior probability distributions for each hypothesis $\Pi_i(\theta)$ from 10^5 random realizations in $(q, \cos \theta_i, \mu, \sigma)$, giving rise to the corresponding prior distributions in the $\theta = \{q, \chi_{\text{eff}}, a_f\}$ parameter space. To compute these distributions we have used the semianalytic expression of a_f given in equation (9) in terms of the fundamental parameters $(m_1, m_2, s_1, s_2, \cos \theta_1, \cos \theta_2)$.

Fig. 2 shows the projections on the planes (q, a_f) , (q, χ_{eff}) and (χ_{eff}, a_f) of the prior probability distributions $\Pi_i(q, \chi_{\text{eff}}, a_f)$ for the last three hypotheses $H_{i=2}$ (isotropic spins, left column), $H_{i=3}$ (aligned spins, central column) and $H_{i=4}$ (anti-aligned spins, right column). In each of these plots, we have chosen for illustrative purposes three different choices for the spin magnitude: Gaussian distributions centered at $\mu = 0, 0.5, 1$ (blue, green and red respectively) with a common width $\sigma = 0.2$. The black curves in the (q, a_f) and (q, χ_{eff}) planes correspond to the limit case of zero spin of the underlying black hole population, while the black curve in the (χ_{eff}, a_f) plane corresponds to the limiting case of $q = 1$.

These plots show that the different hypotheses on the underlying spin configurations populate very different areas of the $(q, \chi_{\text{eff}}, a_f)$ parameter space, and therefore, are taken as priors and integrated over the whole parameter space with the likelihood of each event shown in Fig. 1, we expect very different Bayesian evidence and therefore significantly informative Bayes ratios for the different hypotheses.

Comparing the distributions on the different planes we see that while the χ_{eff} distributions are flat with respect to q , a well-known degeneracy for random spins, the a_f distributions have a strong dependence on q . Models with high and aligned or anti-aligned spins simply cannot produce events with low a_f and q values, as found for event GW190814. Only populations with close to zero spin or high but isotropic spins can populate the lower left corner of the (q, a_f) plane. If the forthcoming O3a catalog includes more events of low q , we expect a much better determination of the spin of the underlying black hole population.

These five hypotheses give rise to multidimensional priors in the space of derived parameters $\theta = \{q, \chi_{\text{eff}}, a_f\}$, which can then be used to integrate them, together with the LVC likelihoods, to obtain the Bayesian evidence for each prior hypothesis H_i , which depend on $\Lambda = \{\sigma, \mu\}$,

$$E_{ij}(\Lambda) = \iint d^3\theta \mathcal{L}(h_j|\theta) \Pi_i(\theta|\Lambda). \quad (10)$$

Once we have obtained the LVC likelihoods from their parameter estimation samples, and computed the individual priors for our hypothesis, we can perform the 3D integration with Mathematica.

G. Bayes Factors

In order to evaluate the global significance of a prior hypothesis from the full BBH catalog, we compute the global Bayes factor for the whole BH population. For this, we will multiply the individual Bayes ratios, assuming that all

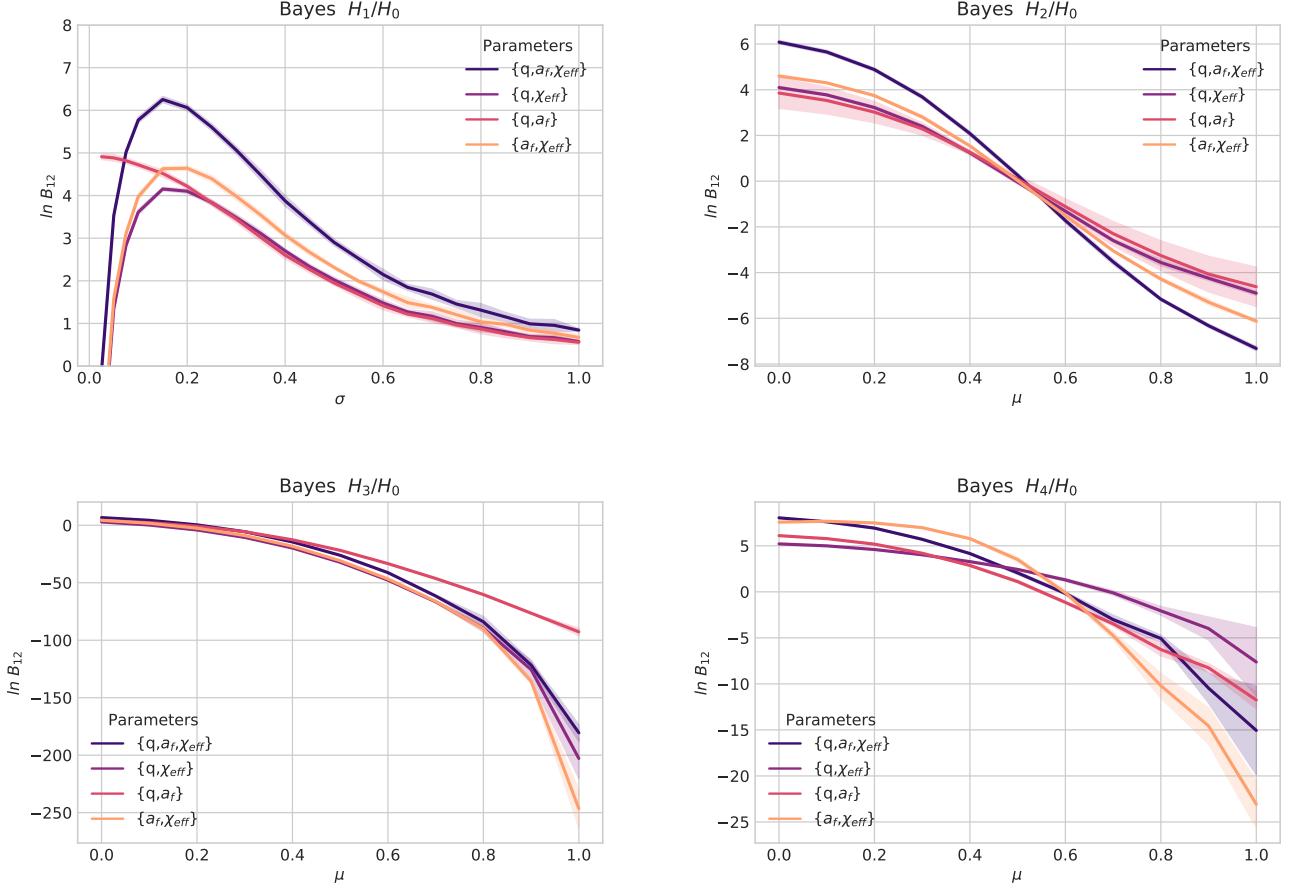


FIG. 3: The log of the Bayes ratios in the variables $(q, \chi_{\text{eff}}, a_f)$ for the four hypotheses: H_1 = isotropic spin centered at zero and width $\sigma \in [0, 1]$ (top left) and $H_{2,3,4}$ = isotropic spin (top right), aligned spin (bottom left) and anti-aligned spin (bottom right), for spin distributions centered at $\mu \in [0, 1]$, for $\sigma = 0.2$. We also study the sensitivity of the Bayes ratios to the evidence in the 2D planes (q, χ_{eff}) , (q, a_f) and (χ_{eff}, a_f) , with increasing values of the Bayes ratios. We find strong evidence $|\ln B_{12}| \geq 6$ at small values of the spin, for all four Hypotheses H_i versus the (null) “All Flat” Hypothesis H_0 , attaining the maximum value of $|\ln B_{12}| \geq 8$ for H_4 at $\mu \leq 0.01$. The events contributing to these Bayes factors are the GWTC-1 (10-event catalog) plus the four published run-O3 events from LIGO/Virgo Collaboration.

events in the catalog are independent,

$$\ln B_{12}^i(\Lambda) = \sum_{j=1}^N \ln \left(\frac{E_{ij}}{E_0} \right) (\Lambda), \quad (11)$$

for each hypothesis $H_i(\Lambda = \{\sigma, \mu\})$.

In practice what we do is a reweighting of the priors, using our population model hypothesis, $\Pi_i(\theta)$, versus the published LIGO/Virgo priors, $\Pi_{\text{LVC}}(\theta)$, which depend on the assumptions of the experiment on each event,

$$\begin{aligned} E_{ij}(\Lambda) &= \iint d^3\theta \mathcal{L}_{\text{LVC}}(d_j|\theta) \Pi_i(\theta|\Lambda) \\ &= E_{\text{LVC}} \iint d^3\theta \mathcal{P}_{\text{LVC}}(\theta|d_j) \frac{\Pi_i(\theta|\Lambda)}{\Pi_{\text{LVC}}(\theta)}. \end{aligned} \quad (12)$$

Since E_{LVC} is the same for all hypotheses, it factors out in the Bayes ratio (11). We can then perform the integral (12) in the full 3D parameter space $\theta = \{q, \chi_{\text{eff}}, a_f\}$ maintaining all the correlations in the multivariate priors and posteriors.

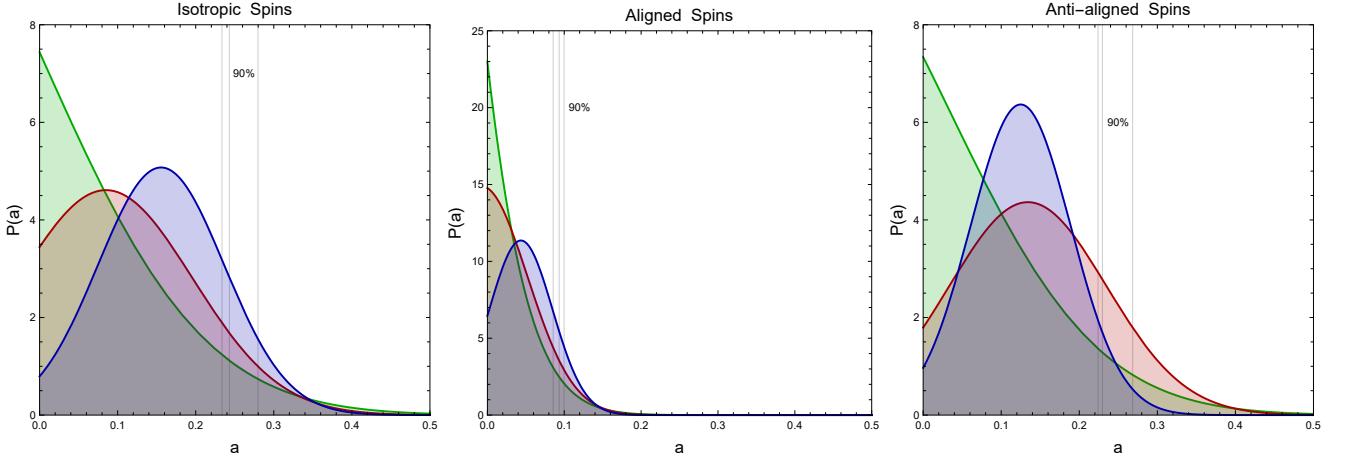


FIG. 4: The posterior PDFs $P(a)$ for the spin magnitude in the case of isotropic (left), aligned (center) and anti-aligned (right) hypothesis. We have computed the PDFs for three different spin widths, $\sigma = 0.05$ (blue), $\sigma = 0.1$ (red) and $\sigma = 0.25$ (green). The PDFs do not seem to depend much on the width σ of the distribution. In all cases, the PDFs peak close to zero spin, specially for $\sigma = 0.25$. The vertical gray lines correspond to the 90% c.l. limit, which is almost the same for all PDFs.

We show in Fig. 3 the Bayes ratios of the four different hypotheses $H_{i=1,2,3,4}$ with respect to the null hypothesis H_0 of flat distributions for mass ratios, spin magnitudes and random spin orientations. According to Jeffrey's scale [47], when $\ln B_{12} \geq 5$, hypothesis 1 is significantly more likely than hypothesis 2, see however [48].

It is clear from figure 3 that low spin magnitudes are significantly preferred for the BBH population of LIGO/Virgo events, with Bayes ratios above $\ln B_{12} = 5$ for $\sigma, \mu < 0.2$, therefore we find strong evidence for small values of the spin, for all four Hypotheses H_i versus the "null" (All Flat) Hypothesis H_0 . In the case of Hypothesis H_1 , with zero spin and allowing for variable width, we find strong evidence ($\ln B_{12} = 6.2$) for relatively narrow spin distributions $\sigma = 0.15$. This low and narrow spin hypothesis represents the isotropic spin distribution that one would expect from an underlying population of primordial black holes. On the other hand, when we vary the spin magnitude as a Gaussian centered at μ for fixed width $\sigma = 0.2$, we find, in all three orientation Hypothesis $H_{2,3,4}$, that the maximum evidence occurs again for $\mu = 0$. Therefore we conclude that, whatever the orientation, there is very strong evidence for low spins in LIGO/Virgo BH.

Note that the aligned spin hypothesis H_3 has Bayes factors $\ln B_{12} < -5$ for spins $\mu > 0.25$, reaching large Bayes ratios $\ln B_{12} < -100$ for large spins $\mu > 0.8$, therefore our analysis of LVC catalog strongly *disfavours* a population of BBH with aligned spins and magnitudes greater than 0.25. Note also that hypothesis H_1 and H_4 , for isotropic and anti-aligned spins, have a similar behaviour in their Bayes factors at low spins, $\mu < 0.5$, reaching values above $\ln B_{12} \simeq 6$ for $\mu < 0.06$ in the isotropic case, and above $\ln B_{12} \simeq 8$ for $\mu < 0.02$ in the anti-aligned case, giving slightly higher significance for anti-aligned versus isotropic spins in the range $\mu < 0.5$. However, for large spins, $\mu > 0.5$, the anti-aligned hypothesis is much more strongly disfavoured than the isotropic one.

We also compute the Bayes factors for the three pairs (q, χ_{eff}) , (q, a_f) and (χ_{eff}, a_f) finding less evidence for low spins compared with the full analysis. The least significant is the Bayes ratio in the plane (q, χ_{eff}) , giving only mild preference for low spins. As we include information coming from a_f , as in (χ_{eff}, a_f) , (q, a_f) and $(q, \chi_{\text{eff}}, a_f)$, the evidence for low spin rises, in some cases up to $\ln B_{12} \sim 6$ or above, which is very strong evidence in favor of that hypothesis.

Historically, the emphasis has been focused on the importance of distinguishing between aligned and anti-aligned astrophysical models of BBH formation. Now that the evidence for small spins is so strong, the orientation of the spins becomes less relevant. It is much more interesting to characterize the posterior distributions of the whole population and to quantify the deviation from zero spins according to the various hypotheses H_i , which take into account orientation. This will be useful in order to characterize the possible origin of LVC BBH events from primordial black hole populations.

H. Bayesian hierarchical modeling

Given that the random isotropic spin hypothesis, H_1 , gives a value of $\ln B_{12} \sim 6$, strongly favouring low spin magnitudes $a < 0.2$ for the combined LVC BBH events, we would like to explore in detail the black hole population distribution of spin magnitudes.

In order to go forward in the information content with respect to the Bayesian approach of the previous section, we will use Bayesian hierarchical modeling, see e.g. [30], as a way to estimate the posterior distribution for spin magnitudes of the LVC black hole population. Here the spin magnitude a enters indirectly through the dependence of other measured quantities, like the effective spin χ_{eff} or the final spin a_f , on it. Then the posterior PDF for a is computed as

$$P_i(a) = \Pi(a) \prod_{j=1}^N E_{ij}(a), \quad (13)$$

where $E_{ij}(a)$ are the Bayesian evidences (10) computed from the LVC likelihoods with the spin-dependent priors $\Pi_i(q, \chi_{\text{eff}}, a_f | \Lambda)$, with $\Lambda = \{\sigma = 0.05, 0.1, 0.25, \mu = a\}$, and $\Pi(a)$ is assumed here to be flat in spin magnitude a , not to give any prior preference for any spin magnitude. We have plotted in Fig. 4 the posterior distributions for the three spin Hypothesis (isotropic, aligned and anti-aligned). In all cases, the preferred spin magnitude is below $a < 0.12$, within 50% c.l. For example, for isotropic spins, currently favoured by the data, we find $a < 0.25$ at 90% c.l. For aligned spins the PDF is even more strongly peaked around zero spin, with $a < 0.1$ at 90% c.l. All this suggests that the inclusion of the four published run O3 events, with increased sensitivity of the detectors, has revealed a property of the population of LIGO/Virgo black holes that was not present in previous studies [42].

I. Conclusions

The routine detection of BBH inspirals by LVC has opened the door to a detailed exploration of the nature of black holes and their populations. The online availability of the strain time-streams and the parameter estimation samples for each event allows for an independent analysis, opening the possibility to compute multivariate likelihoods for combinations of derived parameters like the mass ratio, or the effective spin.

It is then a matter of personal choice which parameters to use in order to infer properties of the populations of black holes detected by LVC. In this paper we have concentrated on just three parameters ($q, \chi_{\text{eff}}, a_f$) that we believe capture the essence of the spin nature of the population of LVC black holes.

We have put forward four different spin-magnitudes and spin-orientation prior hypothesis, $H_{i=1,\dots,4}$ (i.e. isotropic, aligned and anti-aligned), to compare with the (null) all-flat prior hypothesis, H_0 , and conducted a Bayesian analysis study to determine the goodness of a given spin-distribution hypothesis for the whole population of LIGO/Virgo black holes.

We find that all spin-orientation hypothesis have a larger Bayes factor for low spins. In some cases the log of the Bayes factor reaches values well above five, thus signalling a strong evidence in favour of low spins (below magnitude $a = 0.2$). Moreover, the largest Bayes factors are obtained by spin distributions peaked at zero spin (Bayes factors above 5 for widths $\mu < 0.2$), with very small width.

We also note that aligned spins are strongly disfavoured, specially for large spin magnitudes (with the log of the Bayes factors as low as -100), as would be expected from astrophysical black holes from isolated binaries. On the other hand, we find that LIGO/Virgo black hole population has a preference towards low spins with isotropic orientations, consistent with what one would expect from primordial black holes in clusters [24].

We have then computed the posterior PDF for the spin *magnitude* in the case of the three alternative spin hypothesis, for different spin widths ($\sigma = 0.05, 0.10, 0.25$). In all cases the spin distribution is peaked at very low values ($a \simeq 0.1, 0.05, 0.1$ for isotropic, aligned and anti-aligned cases, respectively), which clearly indicates a preference of the whole LIGO/Virgo BH population for low spins, irrespective of orientation.

Therefore, we conclude that only using the BH spin as a discriminator between the astrophysical versus primordial nature of LIGO/Virgo black holes, when considered as a homogeneous population of black holes, our analysis seems to suggest very strongly that they are consistent with being primordial.

IV The False Alarms Induced by Gaussian Noise in Gravitational Wave Detectors

A. Journal & publication details

Published in Physical Review D 107, 023027 (2023) [2]

B. Authors

The authors in the same order as presented to the journal are

Gonzalo Morrás^{a7}

Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco 28049 Madrid, Spain

José Francisco Nuño Siles^{a8}

Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco 28049 Madrid, Spain

Juan García-Bellido^{a9}

Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco 28049 Madrid, Spain

Ester Ruiz Morales^{b10}

Departamento de Física, ETSIDI, Universidad Politécnica de Madrid, 28012 Madrid, Spain. Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco 28049 Madrid, Spain

C. Abstract

Gaussian noise is an irreducible component of the background in gravitational wave (GW) detectors. Although stationary Gaussian noise is uncorrelated in frequencies, we show that there is an important correlation in time when looking at the matched filter signal to noise ratio (SNR) of a template, with a typical autocorrelation time that depends on the template and the shape of the noise power spectral density (PSD). Taking this correlation into account, we compute from first principles the false alarm rate (FAR) of a template in Gaussian noise, defined as the number of occurrences per unit time that the template's matched filter SNR goes over a threshold ρ . We find that the Gaussian FAR can be well approximated by the usual expression for uncorrelated noise, if we replace the sampling rate by an effective sampling rate that depends on the parameters of the template, the noise PSD and the threshold ρ . This results in a minimum SNR threshold that has to be demanded to a given GW trigger, if we want to keep events generated from Gaussian noise below a certain FAR. We extend the formalism to multiple detectors and to the analysis of GW events. We apply our method to the GW candidates added in the GWTC-3 catalog, and discuss the possibility that GW200308_173609 and GW200322_091133 could be generated by Gaussian noise fluctuations.

D. Introduction

A century after their theoretical derivation from General Relativity [49], Gravitational Waves (GWs) are now routinely detected by the laser interferometers of the LIGO-Virgo-KAGRA collaboration [50–52]. Their amplitude is so small that their detection above instrumental and environmental noise requires sophisticated pipelines [53–58], which look for signals in the data with various methods. These pipelines have to be designed to reject noise from very common non-Gaussian transient sources of noise (also known as glitches) [59], while being computationally efficient to search for events in a wide range of parameters within an affordable amount of time.

In the case of modeled searches for GWs from Compact Binary Coalescences (CBCs), templates from a predefined template bank are compared with the data at all times to find where a GW signal can be present. The likelihood that the observed data contains a GW signal is quantified by computing a pipeline-specific ranking statistic, defined in such a way that the larger its value the more it favors the signal hypothesis versus the noise hypothesis. If the detector noise were purely Gaussian, it can be proved that the optimal ranking statistic for a signal of known form would be the matched filter SNR [60]. However, the search pipelines that actually look for GWs use ranking statistics

⁷ gonzalo.morras@uam.es

⁸ josef.nunno@estudiante.uam.es

⁹ juan.garciaellido@uam.es

¹⁰ ester.ruiz.morales@upm.es

that, although based on the SNR, introduce corrections to consider the presence of non-Gaussian glitches which can give sizeable spurious SNR values. The corrections are usually based on signal consistency tests, a common example being the use of χ^2 [61] to weigh down the SNR.

In order to assign a significance to the candidate events in terms of their ranking statistic, the pipelines need to find the background distribution of the ranking statistic for the bank of templates. This is estimated in a data driven way, usually by running the search on the time-shifted strain of the different interferometers, so that coincidences become not physical and the triggers obtained this way represent an estimate of the background noise. The false alarm rate (FAR) of an event is then defined by the search pipeline as the rate of background triggers over the whole bank of templates with ranking statistic equal to or higher than the one observed for the event. Therefore, the FAR can give us an idea of how likely it is for noise to generate an event. Intuitively, for a total observation time T_{obs} , any trigger that has $\text{FAR} \geq 1/T_{\text{obs}}$ is compatible with being generated by noise, while $\text{FAR} \ll 1/T_{\text{obs}}$ disfavors the noise hypothesis.

In searches for GWs, the FAR estimates can differ several orders of magnitude among different pipelines [11], given that the FAR usually has an exponential dependence on the ranking statistic. Therefore, small variations in how the data is processed, what templates are used or what is looked for to rank the events in the different pipelines, can result in orders of magnitude discrepancies in the estimation of the FAR.

Moreover, the FAR does not contain any information about the foreground. To take this into account, together with the astrophysical prior knowledge, the p_{astro} , was introduced [62]. The rationale behind p_{astro} is to give the Bayesian probability that a candidate is from astrophysical origin under a model for the foreground rates $f(x, \vec{\theta})$ and background rates $b(x, \vec{\theta})$ that depend on the ranking statistic x and the template parameters $\vec{\theta}$. A threshold value of $p_{\text{astro}} > 0.5$ was required for any candidate event to be included in the GWTC-3 catalog [11]. The estimated expected contamination from events of terrestrial origin is $\sim 10\text{--}15\%$, or $\sim 4\text{--}6$ events. In the same fashion as the FAR, the p_{astro} for a given event can be very different between pipelines and presents large uncertainties, especially around $p_{\text{astro}} \sim 0.5$ [63].

As a consequence of the application of this threshold to enter the GWTC-3 Catalog, some events were accepted with FAR values greater than $1/T_{\text{obs}} \sim 2\text{yr}^{-1}$. One example is GW200322-091133 [11] with $\text{FAR} > 400\text{yr}^{-1}$, which, upon further investigation with Bayesian Parameter Estimation (PE), was found to have low SNR (≤ 8.5) and multimodal posterior distributions of its parameters. Since the likelihood used in PE is approximately proportional to $\exp(\text{SNR}^2/2)$, in events with small SNR the likelihood will not have a large enough peak so as to dominate the posterior, and there will be prior-dominated modes.

All these difficulties may prompt one to think that these candidate events with low SNR values might come from noise fluctuations. The noise and GW signal hypotheses are usually compared locally using the *Bayes factor* [64]. However, this number says nothing about how often we expect noise to generate a signal as “loud” as the observed one. This has motivated us to question whether we could aim to obtain a theoretical lower bound on the false alarm rate of an event, independently of all the complexities involved in the search pipelines. We start from the idea that Gaussian noise is always an irreducible component of the background in GW detectors [65, 66], and generates a rate of false alarms that could be calculated analytically. In the case in which non-Gaussianities are also present in the strain, more false alarms will be induced [67], as matches will occur more easily for a given template, thus making our estimate assuming only Gaussian noise a lower bound on their FAR, and thus an upper bound on their significance.

In this paper, we propose a new method to derive a local statistical measure of the significance of an event. The main idea will be to give a theoretical estimate of how often we would expect Gaussian noise colored with the local PSD to produce a fluctuation that matches a specific template with the same or higher SNR than the one observed. In Sec. IV E we develop the framework to compute the FAR for a given template in Gaussian noise from a single detector and study its dependence on different parameters for CBC templates. In Sec. IV F we extend the formalism to compute the FAR of a template when multiple detectors are online. In Section IV G we show how to apply our statistical method to events observed in the strain and in Sec. IV G 1 we use it on the O3b events included in GWTC-3. Finally in section IV H we present our conclusions.

E. The false alarm rate of a template in a single detector

In this section we want to determine, given a template $h(t)$, how much time of stationary Gaussian noise $n(t)$, from a given detector, we would have to look at, on average, to obtain a match with a signal to noise ratio (SNR) greater than some threshold ρ .

In general the noise will have zero mean, $\langle \tilde{n} \rangle = 0$, and assuming that it is stationary, the different Fourier modes are uncorrelated,

$$\langle \tilde{n}^*(f) \tilde{n}(f') \rangle = \frac{1}{2} S_n(f) \delta(f - f'), \quad (14)$$

which can be seen as the definition of the noise power spectral density (PSD) $S_n(f)$. If we assume that the noise is Gaussian, it is characterized completely by the fact that it has zero mean and a variance given in Eq. (14). Using the PSD we can define the following inner product,

$$\langle a, b \rangle = 4 \iint_{f_{\min}}^{f_{\max}} \frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n(f)} df, \quad (15)$$

where tildes denote Fourier transform. This inner product can be used to write down the usual definitions [68] of the optimal SNR:

$$\rho^{\text{opt}} = \sqrt{\langle h, h \rangle}, \quad (16)$$

and the matched filter SNR:

$$\rho^{\text{mf}} = \frac{\langle h, s \rangle}{\rho^{\text{opt}}}, \quad (17)$$

where $s(t)$ is the detector output strain, which in our case we will assume to be given by stationary Gaussian noise $n(t)$ with PSD S_n . Under this assumption, it can be proved that ρ^{mf} is a complex normal random variable (i.e. a Gaussian with unit dispersion, $\sigma = 1$) [69]:

$$p(\rho^{\text{mf}}) d\text{Re}\rho^{\text{mf}} d\text{Im}\rho^{\text{mf}} = \frac{1}{2\pi} e^{-\frac{1}{2}|\rho^{\text{mf}}|^2} d\text{Re}\rho^{\text{mf}} d\text{Im}\rho^{\text{mf}}. \quad (18)$$

and the real part of the matched filter SNR is the optimum quantity to rank the significance of events for a signal of known form under the assumption of Gaussian noise [60]. This quantity is very closely related to the likelihood ratio for the signal vs Gaussian noise hypotheses, which is the Bayes factor for a signal of known intrinsic parameters. However, it is common to be in the situation in which the global phase of the GW can be changed arbitrarily and does not contain any astrophysical information [70]. This is the case in a quasicircular compact binary coalescence, when we ignore higher order modes and precession. Even when including them, the global phase can typically be neglected since it is highly degenerate with other parameters such as polarization, location in the sky and the azimuthal angle separating the spin vectors of the component BHs. In these cases we will want to ignore the global phase of the GW in the search by taking as our ranking statistic the absolute value of the matched filter SNR:

$$|\rho^{\text{mf}}| = \sqrt{\text{Re}(\rho^{\text{mf}})^2 + \text{Im}(\rho^{\text{mf}})^2}, \quad (19)$$

which is invariant under global phase transformations $\tilde{h}(f) \rightarrow \tilde{h}(f)e^{i\phi_g}$. Indeed, the SNR usually used in searches is $|\rho^{\text{mf}}|$ [71] since it is equivalent to automatically finding the global phase ϕ_g of the GW that maximizes $\text{Re}(\rho^{\text{mf}})$. Because of this we will choose $|\rho^{\text{mf}}|$ as our ranking statistic in this paper. Defining $\tilde{h}(f)$ as the Fourier transform of the template $h(t)$, we can use the following property:

$$\mathcal{F}(h(t')) = \tilde{h}(f)e^{-2\pi if(t'-t)}, \quad (20)$$

and compute the matched filter signal to noise ratio, Eq.(17), at all times as

$$\rho^{\text{mf}}(t) = \frac{4}{\rho^{\text{opt}}} \iint_{f_{\min}}^{f_{\max}} df \frac{\tilde{h}^*(f)\tilde{n}(f)}{S_n(f)} e^{2\pi ift}. \quad (21)$$

where we assume that the strain only contains Gaussian noise. At any fixed point in time, $\rho^{\text{mf}}(t)$ of Eq. (21) will behave as a complex normal variable from Eq. (18) and the probability of obtaining a value of $|\rho^{\text{mf}}|$ greater than ρ will be:

$$\begin{aligned} P(|\rho^{\text{mf}}| > \rho) &= \frac{1}{2\pi} \int_0^{2\pi} d\arg(\rho^{\text{mf}}) \int_0^\infty |\rho^{\text{mf}}| d|\rho^{\text{mf}}| e^{-\frac{1}{2}|\rho^{\text{mf}}|^2} \\ &= e^{-\frac{1}{2}\rho^2}. \end{aligned} \quad (22)$$

A naive computation to estimate the rate of false alarms with $|\rho^{\text{mf}}| > \rho$ would be to multiply this probability by the number of trials per unit time, which in the case that different times were independent, would just be the sampling rate of the detector:

$$\text{FAR}_{\text{naive}} = \frac{1}{\Delta t_{\text{samp}}} e^{-\rho^2/2}. \quad (23)$$

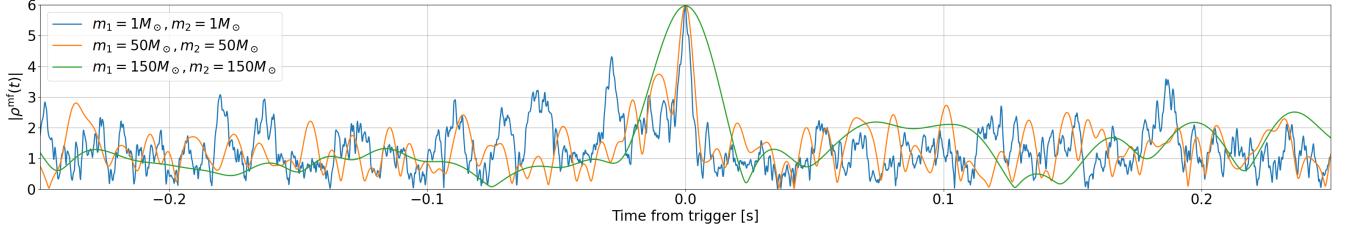


FIG. 5: Simulation of the modulus of the matched filter SNR $|\rho^{\text{mf}}(t)|$ for Gaussian noise generated using Advanced LIGO design sensitivity [50] and `IMRPhenomPv2` [72] templates of masses $m_1 = m_2 = 1M_\odot$, $m_1 = m_2 = 50M_\odot$ and $m_1 = m_2 = 150M_\odot$. Each template is matched with different noise realizations until we obtain a trigger of $|\rho^{\text{mf}}(t_{\text{trig}})| \sim 6$. We plot 0.5s around this trigger.

However, this would be incorrect because the value of $|\rho^{\text{mf}}(t)|$ at different times is correlated. The problem can be explicitly seen in Fig. 5, where we have generated Gaussian noise from Advanced LIGO at design sensitivity [50] and computed $|\rho^{\text{mf}}(t)|$ using Eq. (21) with `IMRPhenomPv2` [72] templates of the specified masses. Each template is matched with different noise realizations until we obtain a trigger of $|\rho^{\text{mf}}(t_{\text{trig}})| \sim 6$, which we show in Fig. 5. The correlation between different times manifests itself in the fact that $|\rho^{\text{mf}}(t)|$ is a smooth function, where the smoothing time scale will be related to the autocorrelation time, and we observe that it depends on the template mass. In particular, the larger the mass, the larger the autocorrelation time will be. This correlation of $|\rho^{\text{mf}}(t)|$ at different times has a direct effect on the False Alarm Rate (FAR), defined as the average time between peaks with $|\rho^{\text{mf}}| > \rho$, since the smoother the function $|\rho^{\text{mf}}(t)|$ is, the less peaks per second it will have, thus reducing the rate of false alarms. Assuming that the sampling rate of the detector is sufficiently fine to see $|\rho^{\text{mf}}(t)|$ as a smooth function, we will demonstrate in the rest of this section that the effect of the correlations will be to replace the sampling rate of the detector $1/\Delta t_{\text{samp}}$ in Eq. (23) by an effective sampling rate that depends on the template, the noise PSD and the threshold ρ .

1. Probabilistic derivation of the FAR

The autocorrelation of $\rho^{\text{mf}}(t)$ can be quantified by computing the covariance between the values of $\rho^{\text{mf}}(t)$ at different times, assuming that the strain only contains Gaussian noise:

$$\begin{aligned}
 \Gamma(t, t') &= \frac{1}{2} \langle \rho^{\text{mf}}(t) \rho^{\text{mf}}(t')^* \rangle = \\
 &= 8 \left\langle \int_{f_{\min}}^{f_{\max}} df \int_{f_{\min}}^{f_{\max}} df' \frac{\tilde{h}^*(f) \tilde{h}(f') \tilde{n}^*(f') \tilde{n}(f)}{S_n(f) S_n(f') (\rho^{\text{opt}})^2} e^{2\pi i (ft - f't')} \right\rangle \left(\right. \\
 &= 8 \int_{f_{\min}}^{f_{\max}} df \int_{f_{\min}}^{f_{\max}} df' \frac{\tilde{h}^*(f) \tilde{h}(f') \langle \tilde{n}^*(f') \tilde{n}(f) \rangle}{S_n(f) S_n(f') (\rho^{\text{opt}})^2} e^{2\pi i (ft - f't')} \\
 &= \Gamma(t - t') = \frac{4}{(\rho^{\text{opt}})^2} \int_{f_{\min}}^{f_{\max}} df \frac{|\tilde{h}(f)|^2}{S_n(f)} e^{2\pi i f(t - t')}, \tag{24}
 \end{aligned}$$

where we have used Eq. (14) and that $\langle \rho^{\text{mf}}(t) \rangle = 0$. We observe in Eq. (24) that for $t = t'$ we have $\Gamma(0) = 1$, as expected from the fact that $\rho^{\text{mf}}(t)$ is a complex normal variable at any specific point in time. In general $\Gamma(t - t')$ will be non-negligible for $t \neq t'$, so the value of the SNR at two different times will be correlated. If we consider the SNR at two different points separated by a time Δt , and define $\rho^{\text{mf}}(t) \equiv \rho_1^c$ and $\rho^{\text{mf}}(t + \Delta t) \equiv \rho_2^c$, from Eq. (24) we have that their joint probability distribution will be given by the following bivariate complex Gaussian:

$$p(\rho_1^c, \rho_2^c) = \frac{\exp \left\{ -\frac{|\rho_1^c|^2 + |\rho_2^c|^2 - 2\text{Re}(\Gamma(\Delta t) \rho_1^c \rho_2^c)}{2(1 - |\Gamma(\Delta t)|^2)} \right\}}{(2\pi)^2 (1 - |\Gamma(\Delta t)|^2)} \left(\right) \tag{25}$$

Using this expression we can compute the two-point false alarm probability (FAP_2), that is, the probability that either ρ_1 or ρ_2 are greater than some SNR threshold ρ ,

$$\text{FAP}_2 = P(\rho_1 > \rho \cup \rho_2 > \rho). \tag{26}$$

An in depth study of this quantity is made in appendix IV I, where we find expressions to compute it numerically and to analytically approximate it to arbitrary order. To understand how FAP_2 behaves, and to gain intuition on how the FAP of more variables will behave, it is interesting to discuss its limiting behaviours. When the separation between the two points is large ($\Delta t \rightarrow \infty$), the correlation between them vanishes ($|\Gamma(\Delta t)| \rightarrow 0$) meaning that FAP_2 becomes the FAP of two uncorrelated variables, that is, $\text{FAP}_2(|\Gamma(\Delta t)| = 0) = 2e^{-\rho^2/2} - e^{-\rho^2}$. As the points get closer together ($\Delta t \rightarrow 0$) the correlation increases ($|\Gamma(\Delta t)| \rightarrow 1$), and FAP_2 will decrease due to correlation effects until the correlation is maximal ($|\Gamma(\Delta t)| = 1$), when the two variables will behave as a single one and $\text{FAP}_2(|\Gamma(\Delta t)| = 1) = e^{-\rho^2/2}$.

In the real setup of a GW experiment, we are interested in determining the false alarm probability for N points separated by a sampling time Δt each. If we define $\rho_k \equiv |\rho^{\text{mf}}(t + k\Delta t)|$, this FAP is given by:

$$\begin{aligned} \text{FAP} &= P \left(\bigcup_{n=1}^N \rho_n > \rho \right) = 1 - P \left(\bigcap_{n=1}^N \rho_n < \rho \right) \left(\bigcap_{n=1}^N \rho_n < \rho \right) \\ &= 1 - P(\rho_1 < \rho) \prod_{k=2}^N P \left(\rho_k < \rho \mid \bigcap_{n=1}^{k-1} \rho_n < \rho \right), \end{aligned} \quad (27)$$

where $P(A|B)$ denotes the conditional probability of A given B and in the last equality we have used the multiplication rule of probability. To compute Eq. (27) we will thus need $P(\rho_2 < \rho | \rho_1 < \rho)$. This can be computed in terms of the FAP_2 defined in Eq. (26):

$$\begin{aligned} P(\rho_2 < \rho | \rho_1 < \rho) &= \frac{P(\rho_1 < \rho \cap \rho_2 < \rho)}{P(\rho_1 < \rho)} \\ &= \frac{1 - \text{FAP}_2(\rho, \Delta t)}{1 - e^{-\rho^2/2}} \\ &\approx 1 - (\text{FAP}_2(\rho, \Delta t) - e^{-\rho^2/2}), \end{aligned} \quad (28)$$

where in the last equality we have assumed that $e^{-\rho^2/2} \ll 1$ (which is true for $\rho \gtrsim 3$). In order to compute Eq. (27) we also need to calculate $P(\rho_k < \rho | \rho_1 < \rho \cap \rho_2 < \rho \cap \dots \cap \rho_{k-1} < \rho)$. We can determine this conditional probability in an approximate way by assuming that it depends only on the nearest neighbor, that is:

$$\begin{aligned} P(\rho_k < \rho | \rho_1 < \rho \cap \rho_2 < \rho \cap \dots \cap \rho_{k-1} < \rho) \\ \approx P(\rho_k < \rho | \rho_{k-1} < \rho) = P(\rho_2 < \rho | \rho_1 < \rho), \end{aligned} \quad (29)$$

where in the last equality we have just used the translation invariance of the problem. The Nearest Neighbor approximation of Eq. (29) will only be valid in the case in which the sampling time Δt is large enough such that second neighbor effects can be neglected, which could be taken into account by replacing the approximation of Eq. (29) by $P(\rho_3 < \rho | \rho_2 < \rho \cap \rho_1 < \rho)$.

Introducing Eqs. (28), (29) into Eq. (27) and assuming that $\text{FAP}_2 - e^{-\rho^2/2} \ll 1$, we have:

$$\begin{aligned} \text{FAP} &\approx 1 - (1 - e^{-\rho^2/2}) \left[1 - \left(\text{FAP}_2(\rho, \Delta t) - e^{-\rho^2/2} \right) \right]^{N-1} \\ &\approx 1 - \exp \left\{ -N \left[\text{FAP}_2(\rho, \Delta t) - e^{-\rho^2/2} \right] \right\} \left(\text{FAP}_2(\rho, \Delta t) - e^{-\rho^2/2} \right) \\ &\approx 1 - \exp \left\{ \left(\frac{T_{\text{obs}}}{\Delta t} \right) \left[\text{FAP}_2(\rho, \Delta t) - e^{-\rho^2/2} \right] \right\} \left(\text{FAP}_2(\rho, \Delta t) - e^{-\rho^2/2} \right), \end{aligned} \quad (30)$$

where T_{obs} is the observing time on which we are computing the FAP, which we assume to be long enough so that $N = T_{\text{obs}}/\Delta t \gg 1$.

To obtain a quantity that is independent of the observing time, we define the false alarm rate (FAR), which is the average number of false alarms per unit time. As we see in Fig. 5, the autocorrelation of the SNR has the effect of clustering its values in peaks. Though each peak of $|\rho^{\text{mf}}(t)|$ has many sample times over the threshold, which naively could count as false alarms, it is important to realize that each peak should be counted as a single false alarm, that is, we have to find the number of *uncorrelated* false alarms which are thus Poisson distributed. This is an important point, given that if each sample time that is over the SNR threshold ρ were counted as a false alarm, we would obtain the naive FAR of Eq. (23), since looking at individual points the probability is given by Eq. (22), and we would greatly overestimate the FAR.

By the definition of the FAR, the mean of the Poisson distribution describing the number of uncorrelated false alarms will be $\lambda = T_{\text{obs}}\text{FAR}$, assuming an observing time T_{obs} . Therefore, the probability of having k false alarms is:

$$p(k) = \frac{(T_{\text{obs}}\text{FAR})^k}{k!} e^{-T_{\text{obs}}\text{FAR}}. \quad (31)$$

Since the FAP is the probability of having one or more false alarms, it is given by:

$$\text{FAP} = \sum_{k=1}^{\infty} p(k) = 1 - p(0) = 1 - \exp\{-T_{\text{obs}}\text{FAR}\}. \quad (32)$$

By comparing Eq. (30) and Eq. (32), we immediately deduce the following relation between the FAR and the FAP:

$$\text{FAR}_2(\rho, \Delta t) = \frac{1}{\Delta t} \left[\text{FAP}_2(\rho, \Delta t) - e^{-\rho^2/2} \right], \quad (33)$$

where we add the subscript 2 to highlight that this FAR has been computed taking into account only nearest neighbors.

2. Evaluation of the FAR of a template

In order to further elaborate the expression of the FAR for a given template in Eq. (33), we need to study the $\text{FAP}_2(\rho, \Delta t)$ more in depth. In the case in which the detector has a high enough sampling rate, we can assume that $|\rho^{\text{mf}}(t)|$ is a continuous function, as is the case in Fig. 5. This will be a very good approximation in LIGO-Virgo, where the data is taken at a sampling rate of $1/\Delta t_{\text{samp}} = 16384$ Hz. In this case, instead of interpreting Δt as the sampling time of the detector, we leave it as a free parameter, as we imagine that the function $|\rho^{\text{mf}}(t)|$ can be resampled arbitrarily. We will want to make $\Delta t \rightarrow 0$, to obtain the result for when $|\rho^{\text{mf}}(t)|$ is continuously sampled, but if Δt is too small, the nearest neighbor approximation will stop being valid. The effect of the farther neighbors will be to reduce the number of effective trials. This compensates the increase in the number of sampling points in such a way that the exact FAR with all correlations taken into account will be smaller than the FAR from the nearest neighbor approximation, that is:

$$\text{FAR}(\rho, \Delta t_{\text{samp}}) \leq \text{FAR}_2(\rho, \Delta t_{\text{samp}}). \quad (34)$$

With this in mind, we approximate the FAR of Eq. (33) for $\Delta t \rightarrow 0$, which from Eq. (24) is equivalent to $|\Gamma(\Delta t)| \rightarrow 1$. We can do this by introducing in Eq. (33) the expression for FAP_2 of Eq. (88) found in Appendix. IV I, keeping only next to leading order terms in $1 - |\Gamma(\Delta t)|$ and assuming that $\rho^2 \gg 1$:

$$\text{FAR}_2 \approx \frac{e^{-\rho^2/2}}{\Delta t} \text{Erf} \left[\frac{\rho \sqrt{1 - |\Gamma(\Delta t)|}}{2} \right]. \quad (35)$$

Since we are interested in the limit $\Delta t \rightarrow 0$, we can substitute $\Gamma(\Delta t)$ by its Taylor expansion around $\Delta t = 0$, which using the definition in Eq. (24) will be given by:

$$\begin{aligned} \Gamma(\Delta t) &= \frac{4}{(\rho^{\text{opt}})^2} \iint_{f_{\text{min}}}^{f_{\text{max}}} df \frac{|\tilde{h}(f)|^2}{S_n(f)} e^{2\pi i f \Delta t} \\ &= \frac{4}{(\rho^{\text{opt}})^2} \iint_{f_{\text{min}}}^{f_{\text{max}}} df \frac{|\tilde{h}(f)|^2}{S_n(f)} \sum_{k=0}^{\infty} \frac{(2\pi i f \Delta t)^k}{k!} \\ &= \sum_{k=0}^{\infty} C_k \frac{(\Delta t)^k}{k!}, \end{aligned} \quad (36)$$

where C_k are real constants defined as

$$C_k = \frac{4}{(\rho^{\text{opt}})^2} \iint_{f_{\text{min}}}^{f_{\text{max}}} df (2\pi f)^k \frac{|\tilde{h}(f)|^2}{S_n(f)}. \quad (37)$$

To leading order in Δt , we then have that $|\Gamma(\Delta t)|$ will be given by:

$$|\Gamma(\Delta t)| = 1 - \frac{1}{2} (C_2 - C_1^2) (\Delta t)^2, \quad (38)$$

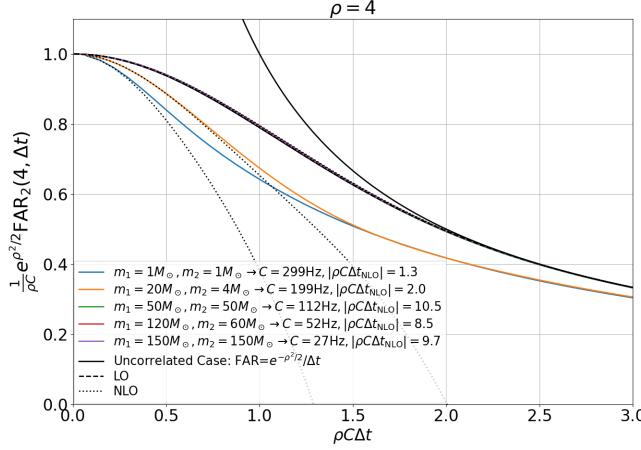


FIG. 6: The FAR_2 for IMRPhenomPv2 [72] templates of different masses, assuming Advanced LIGO at design sensitivity [50] and setting the threshold SNR $\rho = 4$. The FAR is computed using the exact expression (Eq. (33)), the leading-order expression of Eq. (39) and the NLO expression of Eq. (50), where the integrals in frequency are always computed between $f_{\min} = 20$ Hz and $f_{\max} = 2048$ Hz to mimic normal GW analysis. We normalize the FAR to its value at 0 separation and the time to make the LO approximation of cases appear the same. The uncorrelated case of Eq. (23) is also plotted.

where we have used that $C_0 = 1$. Substituting the expansion for $|\Gamma(\Delta t)|$ of Eq. (38) into Eq. (35) and keeping terms in Δt up to leading order, we obtain:

$$\text{FAR}_2(\rho, \Delta t) \approx \frac{e^{-\rho^2/2}}{\Delta t} \text{Erf} \left[\frac{\sqrt{\pi}}{2} \rho C \Delta t \right], \quad (39)$$

where for simplicity we have defined:

$$C \equiv \sqrt{\frac{C_2 - C_1^2}{2\pi}}, \quad (40)$$

which is always a real quantity, since $C_2 - C_1^2 \geq 0$.¹¹ This is a necessary condition given by the fact that $C_2 - C_1^2$ is the leading order coefficient in the Taylor expansion of $|\Gamma(\Delta t)|$ (see Eq. (38)) and we know that $|\Gamma(\Delta t)| \leq 1$.

From Eq. (39) we have that in the limit $\Delta t \rightarrow 0$:

$$\text{FAR}_2(\rho, 0) = C \rho e^{-\rho^2/2}. \quad (43)$$

The way to interpret the result of Eq. (43) is that even if we consider the separation between points to tend to 0, the FAR will not diverge, as we would have naively deduced from Eq. (23). The correlation between the neighboring points will regularize the FAR to the finite value of Eq. (43).

This can be seen in Fig. 6, where we show the FAR_2 for IMRPhenomPv2 [72] templates of different masses, assuming Advanced LIGO at design sensitivity [50]. The FAR is computed using the exact expression (Eq. (33)), the leading order (LO) expression of Eq. (39) and the next-to-leading order (NLO) expression of Eq. (50), which will be discussed in the next subsection. For the cases of large masses ($m_{1,2} = 50M_\odot$, $m_{1,2} = 150M_\odot$ and $m_1 = 120M_\odot$, $m_2 = 60M_\odot$),

¹¹ We can explicitly prove that $C_2 - C_1^2 \geq 0$ and gain some intuition on C , if we realize that

$$g(f) = \begin{cases} \frac{4}{(\rho^{\text{opt}})^2} \frac{|\tilde{h}(f)|^2}{S_n(f)} & f_{\min} < f < f_{\max} \\ 0 & \text{else} \end{cases} \quad (41)$$

can be interpreted as a probability distribution function, since it is always non-negative and it is normalized (i.e. $\int_{-\infty}^{\infty} g(f) df = 1$). Using this probability distribution function, we then observe that C is simply given by:

$$C = \sqrt{2\pi(\mathbb{E}_g[f^2] - \mathbb{E}_g[f]^2)} = \sqrt{2\pi \mathbb{E}_g[(f - \mathbb{E}_g[f])^2]} = \sqrt{2\pi} \sigma_f. \quad (42)$$

where $\mathbb{E}_g[X]$ denotes the expectation value of X in g , σ_f is the standard deviation of the frequency f in g , and from the second equality we explicitly see that the argument of the square root is always positive. From Eq. (42) we then observe that C will be directly related with the bandwidth, that is, how spread out in frequencies is $g(f)$. Therefore, the more broadband our detector and signals are, the larger C will be in general.

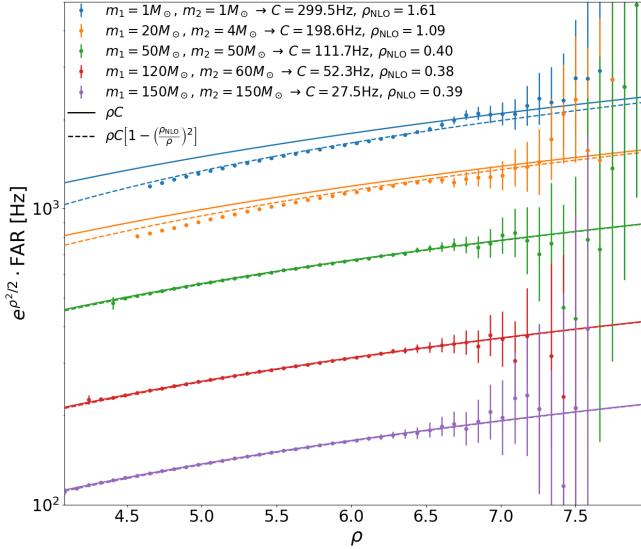


FIG. 7: Comparison between the simulated and predicted FAR for five different IMRPhenomPv2 templates. The simulation is done using 15 million chunks of 512 s of simulated Gaussian noise generated from Advanced LIGO at design sensitivity [50]. We directly compute the probability to have a trigger with $|\rho^{\text{mf}}| > \rho$ by performing matched filtering between $f_{\text{min}} = 20$ Hz and $f_{\text{max}} = 2048$ Hz with the corresponding GW template and dividing the number of chunks where we find a match with $|\rho^{\text{mf}}| > \rho$ by the total number of chunks analyzed. The error on the FAP is computed using the Wilson score 90% confidence interval [73]. Introducing this FAP in Eq. (45) (using $T_{\text{obs}} = 512$ s) we obtain the FAR plotted with dots, whose error bars represent the 90% confidence interval. For the theory curves, the corresponding values of C are computed with Eq. (40) and Eq. (37).

we have that the leading order expression gives an accurate representation of the exact result, as can be seen from the fact that the lines for the three high mass cases are on top of each other and on top of their corresponding LO and NLO approximations. This is no longer true for the low mass cases of $m_{1,2} = 1M_{\odot}$ and $m_1 = 20M_{\odot}, m_2 = 4M_{\odot}$, where the FAR decreases faster than expected at high values of Δt due to correlation tails at this high Δt . To describe this deviation from the LO result, we will have to take into account higher order corrections in Δt , which will be discussed in the next subsection.

The fewer trials we do, the smaller the FAR should be. Therefore the FAR is a monotonously decreasing function of Δt , and $\text{FAR}_2(\rho, \Delta t_{\text{samp}}) \leq \text{FAR}_2(\rho, 0)$, which can correctly be seen in Fig. 6. Using this together with Eq. (34) we obtain

$$\text{FAR}(\rho, \Delta t_{\text{samp}}) \leq \text{FAR}_2(\rho, 0) = \rho C e^{-\rho^2/2}. \quad (44)$$

We expect that the result of Eq. (44) will be a very tight upper bound, and thus a good approximation of the exact FAR in the case that the NLO corrections are small, since these are related with the length of the correlations and thus the importance of the next-to-near neighbors.

To study the validity of this result we will simulate the problem at hand. In particular, we will simulate the FAR by generating many chunks of simulated Gaussian noise from Advanced LIGO at design sensitivity [50] of duration $T_{\text{obs}} = 512$ s. We directly compute the probability to have a trigger with $|\rho^{\text{mf}}| > \rho$ by performing matched filtering on the noise using a GW template and dividing the number of chunks where we find a match with $|\rho^{\text{mf}}| > \rho$ by the total number of chunks analyzed. From this FAP we can obtain the FAR simply by inverting Eq. (32):

$$\text{FAR} = \frac{1}{T_{\text{obs}}} \log \left(\frac{1}{1 - \text{FAP}} \right) \quad (45)$$

In Fig. 7 we show the FAR computed in this way from the simulation of the FAP and multiplied by $e^{\rho^2/2}$ to extract the exponential decay behavior and make visualization easier. The matched filter is done with five different IMRPhenomPv2 templates with the same masses as the ones used in Fig. 6. We have observed that indeed, Eq. (44) is always satisfied and $\text{FAR}_2(\rho, 0)$ is an upper bound of $\text{FAR}_N(\rho, \Delta t_{\text{samp}})$ within the error. As was discussed previously,

this is a tight upper bound in the case in which the NLO corrections are small, deviating by less than 1 part in 1000 for the larger masses ($m_{1,2} = 50M_{\odot}$, $m_{1,2} = 150M_{\odot}$ and $m_1 = 120M_{\odot}, m_2 = 60M_{\odot}$). In the cases where the NLO corrections are important ($m_{1,2} = 1M_{\odot}$ and $m_1 = 20M_{\odot}, m_2 = 4M_{\odot}$) we can observe that even though Eq. (44) is still a good upper bound, it is not so tight any more. Nonetheless, the maximum relative error between the upper bound and the exact value always stays below 15% and decreases towards larger values of the SNR threshold ρ . We thus confirm that a good approximation of the FAR is:

$$\text{FAR} = C \rho e^{-\rho^2/2}. \quad (46)$$

Comparing this expression with the value of the naive FAR that we derived at the beginning in Eq. (23), we have that, as anticipated, the sampling time of the experiment is naturally replaced by an effective sampling time for which we can obtain the same result as for uncorrelated points. This effective sampling rate depends on the threshold ρ and on the template and noise PSD via the coefficient C :

$$\Delta t_{\text{eff}} = \frac{1}{\rho C} \quad (47)$$

Consistently computing corrections to this result, we would have to take into account the effect of next-to-leading order corrections. We do this in the next subsection.

3. NLO corrections to the FAR of a template

We will start by studying the next-to-leading-order (NLO) corrections to the expression for FAR_2 found in Eq. (39). For this we now substitute in Eq. (33) the expression for FAP_2 of Eq. (93) found in Appendix. IV I, keeping NLO terms in $1 - |\Gamma(\Delta t)|$ and assuming that $\rho^2 \gg 1$:

$$\text{FAR}_2 \approx \frac{e^{-\rho^2/2}}{\Delta t} \text{Erf} \left\{ \frac{\rho \sqrt{1 - |\Gamma(\Delta t)|}}{2} \left(\left(+ \frac{1 - |\Gamma(\Delta t)|}{4} \right) \right) \right\}. \quad (48)$$

And when considering the Taylor expansion of $|\Gamma(\Delta t)|$ we now keep up to quartic terms, that is:

$$\begin{aligned} |\Gamma(\Delta t)| = 1 - \frac{1}{2} (C_2 - C_1^2) (\Delta t)^2 \\ + \frac{1}{24} (C_4 - 4C_1 C_3 + 6C_1^2 C_2 - 3C_1^4) (\Delta t)^4, \end{aligned} \quad (49)$$

Introducing this Taylor expansion into Eq. (48) and keeping up to leading order terms, we have:

$$\text{FAR}_2(\rho, \Delta t) \approx \frac{e^{-\rho^2/2}}{\Delta t} \text{Erf} \left\{ \frac{\sqrt{\pi}}{2} \rho C \Delta t \left(\left(- \frac{(\Delta t)^2}{(\Delta t_{\text{NLO}})^2} \right) \right) \right\} \left(\right. \quad (50)$$

where we have introduced Δt_{NLO} as the characteristic time for which when $\Delta t \ll |\Delta t_{\text{NLO}}|$ we can neglect higher order effects. In terms of C_k , it will be given by:

$$(\Delta t_{\text{NLO}})^2 = \frac{24(C_2 - C_1^2)}{C_4 - 4C_1 C_3 - 3C_2^2 + 12C_1^2 C_2 - 6C_1^4}. \quad (51)$$

Looking again at Fig. 6 where the NLO FAR_2 of Eq. (50) is compared in with the LO expression (Eq. (39)) and with the exact expression (Eq. (33)), we can observe that the NLO corrections are not important for the high mass systems ($m_{1,2} = 50M_{\odot}$, $m_{1,2} = 150M_{\odot}$ and $m_1 = 120M_{\odot}, m_2 = 60M_{\odot}$), since $|\rho C \Delta t_{\text{NLO}}| \gg 1$. However, for the low mass cases of $m_{1,2} = 1M_{\odot}$ and $m_1 = 20M_{\odot}, m_2 = 4M_{\odot}$, which have $|\rho C \Delta t_{\text{NLO}}| \sim O(1)$, we can see that the higher order corrections in Δt are important. In these cases, the tails of the correlation are relatively longer, and so the FAR decreases faster than expected as a function of Δt , which is accurately described by the NLO corrections as long as $\Delta t \lesssim \Delta t_{\text{NLO}}$.

We also want to obtain a more accurate formula for the Gaussian FAR than the one in Eq. (46). To consistently compute corrections to the result of Eq. (46), we would have to take into account the effect of farther neighbors in Eq. (29). Nonetheless, doing this becomes very complicated rather quickly. Instead, a heuristic way to take into account the next to leading order corrections can be found by imposing that these preserve the same behavior as the leading order term of Eq. (39), which we have seen gives a very good description when higher orders can be neglected.

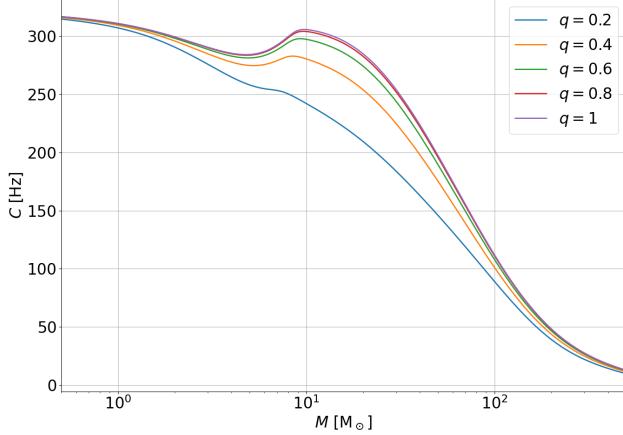


FIG. 8: The FAR prefactor C as a function of the CBC masses parameterized via the total mass of the binary $M = m_1 + m_2$ and the mass ratio $q = m_2/m_1$ and computed using the PSD of Advanced LIGO at design sensitivity [50] between $f_{\min} = 20$ Hz and $f_{\max} = 2048$ Hz. The waveform has been computed using `IMRPhenomPv2` with zero spin.

We can imagine that at $\Delta t_{\text{eff}}/2$ there will be a sampling point whose correlation we are neglecting when we resample $|\rho^{\text{mf}}(t)|$. We will then impose that the correlation $|\Gamma(\Delta t)|$ at this point has the same value as in the case where we only consider the leading order term in the Taylor expansion of Eq. (38):

$$\Gamma\left(\frac{\Delta t_{\text{eff}}}{2}\right) = 1 - \frac{\pi}{4\rho^2}. \quad (52)$$

Using the next to leading order expansion for $|\Gamma(\Delta t)|$ on the left hand side, we obtain:

$$1 - \frac{\pi}{4} (C\Delta t_{\text{eff}})^2 + \frac{\pi}{2} \rho_{\text{NLO}}^2 (C\Delta t_{\text{eff}})^4 = 1 - \frac{\pi}{4\rho^2}. \quad (53)$$

where for convenience we have defined ρ_{NLO} in the following way

$$\rho_{\text{NLO}} = \sqrt{\frac{(C_4 - 4C_1C_3 + 6C_1^2C_2 - 3C_1^4)}{48(C_2 - C_1^2)^2}}. \quad (54)$$

Solving Eq. (53) for Δt_{eff} , keeping only leading-order terms in ρ_{NLO}/ρ , we obtain:

$$\frac{1}{\Delta t_{\text{eff}}^{\text{NLO}}} = \rho C \left[1 - \left(\frac{\rho_{\text{NLO}}}{\rho} \right)^2 \right]. \quad (55)$$

This heuristic result is compared in Fig. 7 with the simulated value. Although we have to keep in mind that it has not been derived in a consistent way, we can observe that it closely follows the behavior of the deviations from Eq. (46) for the cases of $m_{1,2} = 1M_{\odot}$ and $m_1 = 20M_{\odot}, m_2 = 4M_{\odot}$ for which the corrections are important. Eq. (55) will thus be a useful model to understand how these deviations behave. As expected, the heuristic corrections of Eq. (55) make the FAR smaller than the upper bound of Eq. (46). Furthermore, we find that in this model the magnitude of the corrections is governed by ρ_{NLO} , Eq. (54), which is a parameter that characterizes how the correlation $|\Gamma(\Delta t)|$ deviates from a parabola around $\Delta t = 0$. From Eq. (55) we observe that when we increase the SNR threshold ρ , the magnitude of the correction decays as $(\rho_{\text{NLO}}/\rho)^2$, and so for $\rho \gtrsim 3\rho_{\text{NLO}}$, the relative error done when ignoring these corrections is smaller than $\sim 10\%$.

4. Dependence on the CBC template parameters

At a constant matched filter SNR, and neglecting higher order corrections ($\rho \gg \rho_{\text{NLO}}$), the False Alarm Rate of Eq. (46) will only depend on the signal via the multiplicative coefficient C defined in Eq. (40), which when multiplied

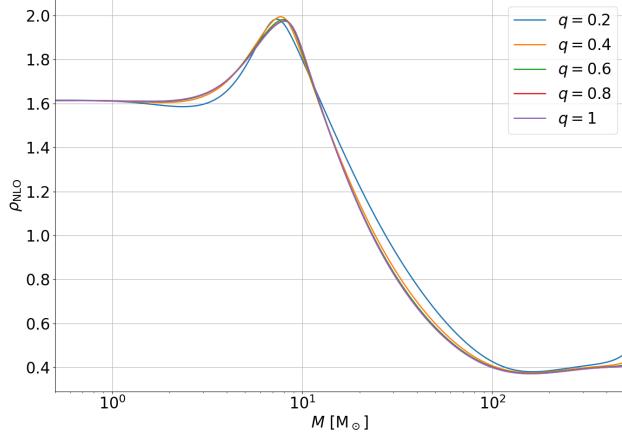


FIG. 9: The ρ_{NLO} as a function of the CBC masses parameterized via the total mass of the binary $M = m_1 + m_2$ and the mass ratio $q = m_2/m_1$ and computed using the PSD of Advanced LIGO at design sensitivity [50] between $f_{\min} = 20$ Hz and $f_{\max} = 2048$ Hz. The waveform has been computed using `IMRPhenomPv2` with zero spin.

by the SNR, gives us the effective sampling rate. Since the higher the effective sampling rate, the more false alarms we expect, we can study how much Gaussian noise background there is in different regions of the CBC parameter space by representing the coefficient C as a function of the CBC parameters. This is done in the Fig. 8, where we plot C as a function of the CBC component masses for the `IMRPhenomPv2` waveform with the spins set to 0.

The masses are parameterized via the total mass of the binary, $M = m_1 + m_2$, and the mass ratio, $q = m_2/m_1$, the leading order parameters that control the amplitude evolution of the waveform [74], which is the part that enters in the computation of C_k in Eq. (37). These results are robust with respect to the waveform choice since C_k depends only on the amplitude evolution which is not as sensitive to modeling uncertainties as quantities that depend on the phase evolution of the template [75].

In Fig. 8 we obtain the natural result that, as a general trend, the higher the mass, the smaller the FAR will be (at a constant ρ). This is because the characteristic frequency of the event will be smaller, and then the characteristic autocorrelation time of the matched filter SNR will be longer, meaning that the time between independent trials will be longer. On top of this general trend we observe a peak at around $M \sim 10M_\odot$, which will be due to events whose merger lies in the upper part of the most sensitive frequency range of the interferometer. Since during merger $|\tilde{h}(f)|^2 \propto f^{-4/3}$ instead of $|h(f)|^2 \propto f^{-7/3}$ as in the inspiral [74], this will make $g(f)$ (Eq. (41)) decay slower at larger frequencies where it is usually suppressed by the quantum shot noise ($S_n(f) \propto f^2$ [76] at high frequency). In this case where merger lies in the upper part of the most sensitive frequency range of the interferometer, the value of C will be larger because the band of frequencies that contribute will be larger. As a consequence of C being larger, the effective sampling rate will be larger, leading to more false alarms.

In Fig. 9, the parameter ρ_{NLO} giving the scale of the next to leading order corrections is shown. This quantity has a similar behavior as that of C , saturating at small masses where the merger is outside the sensitivity band, and generally decreasing at large masses whose merger happens at low frequency. It also has a peak at intermediate masses, corresponding to those systems that merge in the upper range of the frequency band that has the highest sensitivity. Note that in the case of ρ_{NLO} , this peak is more pronounced and towards smaller masses than in the case of C , which is due to the fact that in this range the value of ρ_{NLO} is dominated by the value of C_4 , which weighs more heavily higher frequencies than C_2 , see Eq. (37). The maximum of ρ_{NLO} is achieved in this peak around $M \sim 8M_\odot$, with a value of $\rho_{\text{NLO},\text{max}} \sim 2$. This means that if we go to $\rho \gtrsim 6$, the relative magnitude of the deviations from Eq. (46) will be smaller than $\sim 10\%$ for all CBC parameter range (see Eq. (55)). Therefore, as long as $\rho \gtrsim 6$ Eq. (46) will not only be an upper bound, but also a very good approximation of the FAR.

Having established the validity of Eq. (46) to approximate the FAR, we can now use it to find what SNR threshold ρ would we need to set to discard all events with FAR higher than a given threshold FAR_{th} . To do this we have to invert Eq. (46), which can not be done exactly in terms of elementary functions, since it is a transcendental equation, but it can be done approximately in the limit that $\rho \gg 1$:

$$\rho = \sqrt{2 \log \frac{C}{\text{FAR}_{\text{th}}} + \log \left\{ 2 \left(\log \frac{C}{\text{FAR}_{\text{th}}} \right) 1 + \frac{1}{2 \log \frac{C}{\text{FAR}_{\text{th}}}} \right\}}, \quad (56)$$

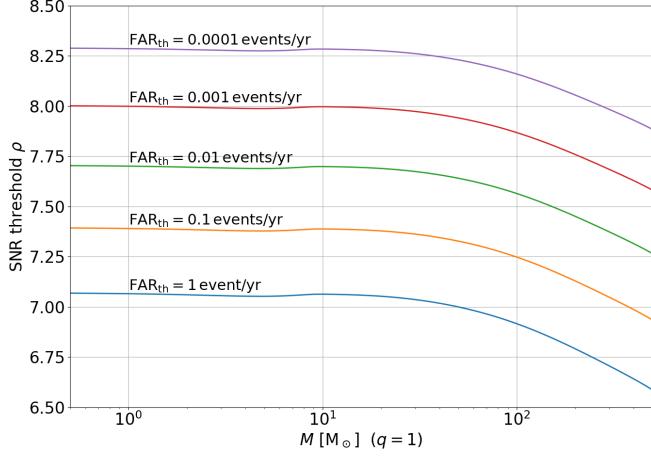


FIG. 10: SNR threshold ρ for different FAR thresholds as a function of the total mass of the binary M , assuming 0 spin and equal component masses ($q = 1$). ρ is plotted using Eq. (56), where the values of C are the same as the ones for the $q = 1$ curve of the left panel of Fig. 8. Direct comparison between the approximation of Eq. (56) and the ρ obtained by numerically inverting Eq. (46) shows that the maximum relative error made on ρ is of 2×10^{-5} for the values shown in this plot.

which gives ρ with a relative error of order $O(\log^2(\rho)/\rho^6)$. In Fig. 10 we have plotted this SNR threshold ρ for different FAR thresholds as a function of the total mass of the binary M , assuming equal component masses ($q = 1$). Even though from Fig. 8 we observe that the value of C depends strongly on M , when we introduce this C in Eq. (56), ρ depends to leading order on the square root of its logarithm and so has only a mild dependence on M as can be seen in Fig. 10. As a general trend, the higher M is, the smaller the SNR threshold ρ will have to be set to exclude false alarms at a given rate FAR_{th} , with the peak at $M \sim 10M_{\odot}$ that was observed in Fig. 8 now less prominent due to the logarithmic dependence. The dependence on FAR_{th} will also be mild, as ρ will also depend on the square root of the logarithm of this quantity. Because of this, the variation of an order of magnitude in FAR_{th} changes ρ by only a small amount. We observe that if we set $\rho = 8$, as is commonly done in the theoretical literature [77], we would be rejecting Gaussian noise false alarms with rates higher than $\text{FAR}_{\text{th}} \sim 10^{-3} \text{yr}^{-1}$.

F. The false alarm rate of a template in a network of detectors

In this section we want to determine how much time of stationary Gaussian noise $n_i(t)$ of the detectors in a network would we have to look at on average to obtain a match with a signal to noise ratio (SNR) greater than some threshold ρ , using a GW template for the two polarization $\{h_+(t), h_{\times}(t)\}$, which when projected in the i -th detector leaves a signal $h_i(t)$. For the problem to be well-posed we will have first to define what we mean by the SNR for multiple detectors. In the case we have more than one detector, the total optimal SNR $\rho_{\text{tot}}^{\text{opt}}$ is defined by summing the individual optimal SNRs (Eq. (16)) in quadrature, that is:

$$\rho_{\text{tot}}^{\text{opt}} = \sqrt{\sum_i \langle h_i, h_i \rangle_i} = \sqrt{\sum_i (\rho_i^{\text{opt}})^2}, \quad (57)$$

where $\langle \times, \times \rangle_i$ denotes the inner product (Eq. (15)) with the PSD $S_i(f)$ of the i -th detector. If $s_i(t)$ is the strain data in the i -th detector of the network, then the total matched filter SNR $\rho_{\text{tot}}^{\text{mf}}$ is defined as:

$$\rho_{\text{tot}}^{\text{mf}} = \frac{1}{\rho_{\text{tot}}^{\text{opt}}} \sum_i \langle h_i, s_i \rangle_i = \frac{1}{\rho_{\text{tot}}^{\text{opt}}} \sum_i \rho_i^{\text{opt}} \rho_i^{\text{mf}}, \quad (58)$$

which given that each ρ_i^{mf} is a complex normal variable, if there are no correlations between detectors, will also be a complex normal variable. As was the case for the single detector matched filter SNR, the real part of Eq. (58) will

be the optimal quantity to rank the triggers when the form of the signal is known. Nonetheless, as was discussed in Sec. IV E, in most cases of interest, the global phase of the GW can be changed arbitrarily and does not carry any information. Therefore we want to set $|\rho_{\text{tot}}^{\text{mf}}|$ as the ranking statistic, so that we get rid of the global phase while keeping the information contained in the relative phase and time of arrival of the GW in each detector, which will be related to the orientation and location of the detectors with respect to the direction and orientation of the GW source. The relative phase of the incoming GW in the different detectors is sometimes ignored in GW searches to reduce computational cost and can easily add single detector triggers [78], although methods to take it into account in a statistical way have recently been introduced [79]. The relative phases between detectors are ignored when using the incoherent SNR, which is obtained adding the absolute value of the single detector matched filter SNRs in quadrature:

$$\rho^{\text{inc}} = \sqrt{\sum_i |\tilde{h}_i^*(f) \tilde{n}_i(f)|^2}. \quad (59)$$

Nonetheless, in this paper this ranking statistic will not be used as a lot of information is lost with it. If we shift in time the signals in all detectors, they will change by the same factor ($\mathcal{F}(h_i(t')) = \mathcal{F}(h_i(t))e^{-2\pi i f(t' - t)}$), and then as in Eq. (21) we can compute the matched filter SNR of the signal at different times with Gaussian noise using the following expression:

$$\rho_{\text{tot}}^{\text{mf}}(t) = \frac{4}{(\rho_{\text{tot}}^{\text{opt}})^2} \iint_{f_{\text{min}}^{\text{min}}}^{f_{\text{max}}^{\text{max}}} df e^{2\pi i f t} \sum_i \frac{\tilde{h}_i^*(f) \tilde{n}_i(f)}{S_i(f)}, \quad (60)$$

where $S_i(f)$ is the noise PSD in the i -th detector. This quantity will also have correlations between different times that will affect the false alarm rate in a very similar way as in Sec. IV E. This correlation can be explicitly seen in Fig. 11, where in the top panel we have plotted a simulation similar to that of Fig. 5 for a random realization of the matched filter SNR for each detector in a Network formed by LIGO Livingston (L1), LIGO Hanford (H1) [50] and Virgo (V1) [51] at their design sensitivities. In the bottom panel we plot the sum of these single detector SNRs both in a coherent way (Eq. (58)) and incoherent way (Eq. (59)). We observe how these two are smooth functions and are thus autocorrelated in time. We also observe that the incoherent SNR is always above the coherent one (sometimes quite significantly), since it ignores the important information carried by the consistency of the GW phase in the different detectors.

In a similar way as in Eq. (24), we can quantify the autocorrelation in time of $\rho_{\text{tot}}^{\text{mf}}(t)$ by computing the covariance between different times:

$$\begin{aligned} \Gamma(t, t') &= \frac{1}{2} \langle \rho_{\text{tot}}^{\text{mf}}(t) \rho_{\text{tot}}^{\text{mf}}(t')^* \rangle = \\ &= \frac{8}{(\rho_{\text{tot}}^{\text{opt}})^2} \left\langle \iint_{f_{\text{min}}^{\text{min}}}^{f_{\text{max}}^{\text{max}}} df \iint_{f_{\text{min}}^{\text{min}}}^{f_{\text{max}}^{\text{max}}} df' e^{2\pi i (ft - f't')} \sum_i \sum_j \frac{\tilde{h}_i^*(f) \tilde{h}_j(f') \tilde{n}_j^*(f') \tilde{n}_i(f)}{S_i(f) S_j(f')} \right\rangle \left(\right. \\ &= \frac{8}{(\rho_{\text{tot}}^{\text{opt}})^2} \iint_{f_{\text{min}}^{\text{min}}}^{f_{\text{max}}^{\text{max}}} df \iint_{f_{\text{min}}^{\text{min}}}^{f_{\text{max}}^{\text{max}}} df' e^{2\pi i (ft - f't')} \sum_i \sum_j \frac{\tilde{h}_i^*(f) \tilde{h}_j(f') \langle \tilde{n}_j^*(f') \tilde{n}_i(f) \rangle}{S_i(f) S_j(f')} \\ &= \Gamma(t - t') = \frac{4}{(\rho_{\text{tot}}^{\text{opt}})^2} \iint_{f_{\text{min}}^{\text{min}}}^{f_{\text{max}}^{\text{max}}} df e^{2\pi i f(t - t')} \sum_i \frac{|\tilde{h}_i(f)|^2}{S_i(f)}, \end{aligned} \quad (61)$$

where we have used that when there is no correlation between the noise of different detectors, then $\langle \tilde{n}_j^*(f') \tilde{n}_i(f) \rangle = \frac{1}{2} S_i(f) \delta_{ij} \delta(f - f')$. What we observe in Eq. (61) is that in the many detector case we obtain the same formula of the covariance as in the single detector case of Eq. (24) if we do the following identification

$$\frac{1}{(\rho^{\text{opt}})^2} \frac{|\tilde{h}(f)|^2}{S_n(f)} \rightarrow \frac{1}{(\rho_{\text{tot}}^{\text{opt}})^2} \sum_i \frac{|\tilde{h}_i(f)|^2}{S_i(f)}. \quad (62)$$

Therefore the FAR will be given by the same expressions that were found in Sec. IV E for the single detector case doing the identification of Eq. (62). That is, an accurate upper bound approximation of the FAR is given by Eq. (46), with C given by the same formula of Eq. (40), but now using the following expression for C_k :

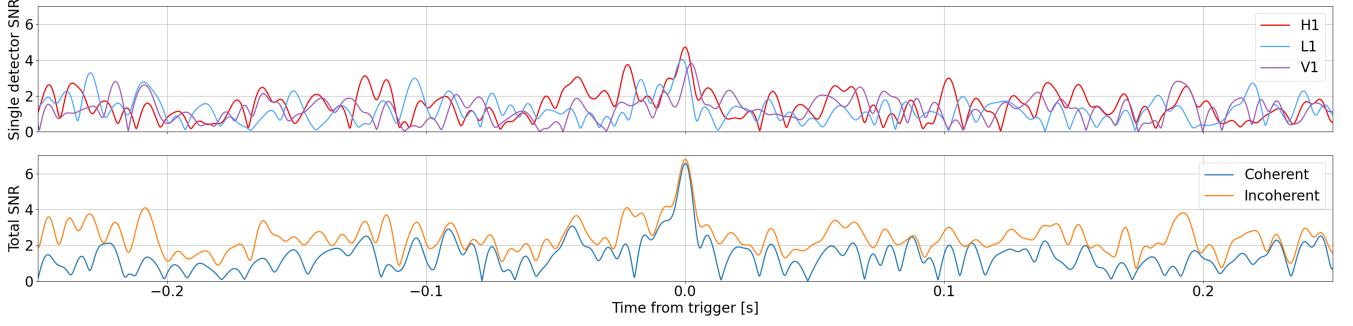


FIG. 11: Top panel: Simulation of the modulus of the single detector matched filter SNR $|\rho^{\text{mf}}(t)|$ for Gaussian noise generated of H1, L1 and V1 at design sensitivity. The match is performed using a spinless IMRPhenomPv2 [72] template of masses $m_1 = m_2 = 50M_{\odot}$, with extrinsic parameters right ascension 1.7rad, declination 1.7rad, polarization 0.2rad and a reference time $t_{\text{GPS}} = 1000000000s$. Bottom Panel: We show the result of adding the single detector SNRs of the top panel both coherently (Eq. (58)) and incoherently (Eq. (59)). To obtain the particular curves shown in this plot we generated random Gaussian noise in the three interferometers until we obtained a realization with a time at which $|\rho_{\text{tot}}^{\text{mf}}| > 6.5$ and we plot 0.5s around the maximum of this trigger.

$$C_k = \frac{4}{(\rho_{\text{tot}}^{\text{opt}})^2} \iint_{f_{\text{min}}}^{f_{\text{max}}} df (2\pi f)^k \sum_i \frac{|\tilde{h}_i(f)|^2}{S_i(f)}. \quad (63)$$

G. Application to GW events

So far we have discussed the FAR and the FAP for a predefined template given a threshold SNR ρ . However, in real settings what we observe is a fluctuation in the strain, that we do not know if it comes from a GW or from noise, and which we will generically call an event. This fluctuation can be interpreted under any template, each giving a different SNR. For a given template, the threshold SNR ρ to use in Eq. (46) for the FAR computation is the observed total matched filter SNR ($\rho = |\rho_{\text{tot}}^{\text{mf}}|$), since we want to know how likely it is to find SNRs equal to or larger than the one observed for the template. The problem will then be how to choose a template, given the observed strain, to determine the SNR and to compute the FAR using Eq. (46). The likelihood is the conditional probability of obtaining the observed strain given a GW signal with parameters $\vec{\theta}$. If we assume Gaussian noise, the likelihood takes the the following form [80]:

$$\begin{aligned} \mathcal{L}(\mathbf{s}|\vec{\theta}) &= \mathcal{N} \exp \left\{ -\frac{1}{2} \sum_i \langle s_i - h_i(\vec{\theta}), s_i - h_i(\vec{\theta}) \rangle_i \right\} \\ &\propto \exp \left\{ \rho_{\text{tot}}^{\text{opt}}(\vec{\theta}) \left(\text{Re} \left\{ \rho_{\text{tot}}^{\text{mf}}(\vec{\theta}, \mathbf{s}) \right\} - \frac{1}{2} \rho_{\text{tot}}^{\text{opt}}(\vec{\theta}) \right) \right\}, \end{aligned} \quad (64)$$

where \mathcal{N} is a normalization constant. Note that the likelihood will be larger for those templates that have the largest matched filter SNR and an optimum SNR such that $\rho_{\text{tot}}^{\text{opt}} = \text{Re}\{\rho_{\text{tot}}^{\text{mf}}\}$, which for GW templates can always be achieved by varying the distance to the source. We then have the expected result that, the more SNR a template has, the larger its Likelihood is and, therefore, the more likely it is to reproduce the observed strain.

However, when we associate a template with an event, we are interpreting the strain fluctuation in terms of a model, with underlying assumptions about the possible physics. The consistent way to take this into account is to think of the event as having a probability of being described by any template, with some priors on each template ¹². Because

¹² For example, even though the template that maximizes the SNR is the one that exactly reproduces the strain ($h_i(t) = s_i(t)$), this is usually a physically impossible GW template, and in this case, we will not consider it. We have that our prior probability for a template that can not be generated by GWs is 0.

we are characterizing a fluctuation observed in the data, we need to evolve our priors to find the probability of each template describing the specific strain. Therefore, what naturally arises is the need to employ Bayes' Theorem to determine the posterior probability $p(\vec{\theta}|\mathbf{s})$ of each template given the observed strain \mathbf{s} :

$$p(\vec{\theta}|\mathbf{s}) = \frac{\mathcal{L}(\mathbf{s}|\vec{\theta})\pi(\vec{\theta})}{\int d\vec{\theta}' \mathcal{L}(\mathbf{s}|\vec{\theta}')\pi(\vec{\theta}')} , \quad (65)$$

where $\pi(\vec{\theta})$ is the prior probability for each set of parameters and it is multiplied by the likelihood to give the posterior. The more SNR a template has, the larger its likelihood and the more weight it will be given in the posterior probability distribution. In Bayesian inference, the posterior $p(\vec{\theta}|\mathbf{s})$ is interpreted as the probability of the template given the strain. Therefore, the template corresponding to the maximum of the posterior probability distribution is the most likely template given the strain and our priors, while the maximum likelihood template is the template most likely to generate the observed strain. In general, these two templates will be different from each other, and they will have different FAPs when computed with Eqs. (32), (46), that we can call $\text{FAP}_{\max p}$ and $\text{FAP}_{\max \mathcal{L}}$ respectively. The most representative template when comparing to the LVK searches would correspond to the maximum likelihood sample, since the modeled searches performed by the LVK [11] deal with the unknown intrinsic parameters by setting up a template bank to cover a target parameter space, and then selecting the template which has the highest likelihood ratio for signal vs noise origin in a given segment of data which, in the Gaussian noise case, means the highest SNR sample. In practice, the FAR reported by LVK searches would be the FAR of this max likelihood template multiplied by the trial factor given by the number of independent templates within the search parameter space.

Another possibility to consider all the information contained in the posterior is to compute the FAP of the fluctuation. To do so, we combine the probability of each template describing the fluctuation given by the posterior, and the probability of each template to be generated by Gaussian noise with an SNR equal to or larger than the observed one, given by the FAP, see Eqs. (32), (46),

$$\begin{aligned} \text{FAP}_{\text{event}} &= \iint d\vec{\theta} p(\vec{\theta}|\mathbf{s}) \text{FAP}(\vec{\theta}, \mathbf{s}) \\ &= \int d\vec{\theta} p(\vec{\theta}|\mathbf{s}) \left(1 - \exp \left\{ \left(T_{\text{obs}} C(\vec{\theta}) \rho_{\text{tot}}^{\text{mf}}(\vec{\theta}, \mathbf{s}) e^{-\frac{1}{2} \rho_{\text{tot}}^{\text{mf}}(\vec{\theta}, \mathbf{s})^2} \right) \right\} \right) \end{aligned} \quad (66)$$

which will always be less than or equal to one, since the posterior $p(\vec{\theta}|\mathbf{s})$ is normalized, as can be seen in Eq. (65). The $\text{FAP}_{\text{event}}$ of Eq. (66) will now not only depend on a single template, but similarly to the Bayes Factor [64] it will take into account the distribution of the likelihood over the prior volume. Therefore, it can be seen as an effective way of considering the trial factor for the template that best matches the data over a parameter space.

In general, the normalization of the posterior, given by the evidence $\mathcal{Z} = \int d\vec{\theta} \mathcal{L}(\mathbf{s}|\vec{\theta})\pi(\vec{\theta})$, is extremely difficult to compute. However, even though the full posterior is unknown, one can use Monte Carlo methods to obtain independent samples from it, as done in Parameter Estimation Analysis [68]. In terms of these independent posterior samples, Eq. (66) can be approximated by:

$$\text{FAP}_{\text{event}} = \frac{1}{N_s} \sum_{i=1}^{N_s} \text{FAP}(\vec{\theta}_i, \mathbf{s}) . \quad (67)$$

where N_s is the number of samples, and the error of approximating the integral by a sum over independent posterior samples is given by:

$$\Delta \text{FAP}_{\text{event}} = \sqrt{\frac{1}{N_s(N_s-1)} \sum_{i=1}^{N_s} \left(\text{FAP}(\vec{\theta}_i, \mathbf{s}) - \text{FAP}_{\text{event}} \right)^2} . \quad (68)$$

1. Application to GW candidates in GWTC-3

As an application of the method previously outlined, we analyze the 35 CBC candidates included in the last gravitational wave transient catalog, GWTC-3 [11], detected during the second part of the third observing run (O3b). The Bayesian Parameter Estimation (PE) of these events has been performed by the LVK collaboration as described in Ref. [11] and the posterior samples obtained are publicly available in Ref. [82].

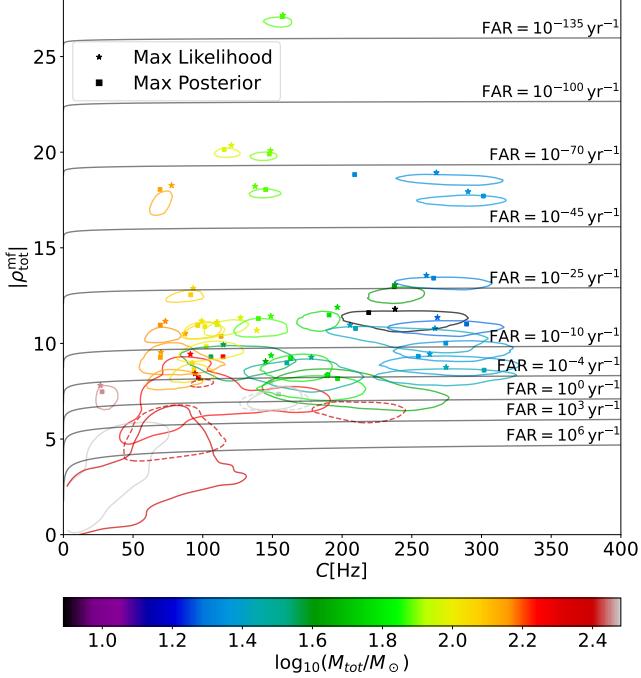


FIG. 12: In this plot we show the contours enclosing 90% of the IMRPhenomXPHM [81] posterior samples in the $(C, |\rho_{\text{tot}}^{\text{mf}}|)$ plane for all the O3b Catalog events. The value of C is computed using Eqs. (40), (63) with the local PSD around each event. The values of $(C, |\rho_{\text{tot}}^{\text{mf}}|)$ for the maximum likelihood and maximum posterior probability samples are marked with a star and square respectively. We also plot using dashed lines the contours of the two events GW200308_173609 (in grey) and GW200322_091133 (in red) after making the cut in the Likelihood as was done for GWTC-3. The color of the contours is given by the median total mass of the posterior samples of each event. We also plot lines of constant FAR as defined by Eq. (56).

We compute the value of C for each IMRPhenomXPHM [81] sample of every event in GWTC-3, using Eqs. (40), (63), where we use the local PSD around each event that is the same one employed in the PE, also available in Ref. [82]. In Fig. 12 we show the 90% credible intervals of $|\rho_{\text{tot}}^{\text{mf}}|$ and C , which are the contours enclosing 90% of the posterior samples in the $(C, |\rho_{\text{tot}}^{\text{mf}}|)$ plane. Since at first order the Gaussian FAR only depends on C and $|\rho_{\text{tot}}^{\text{mf}}|$, we can plot on top of Fig. 12 the contours of constant FAR using Eq. (56). We observe that for most of the events, almost all the samples are above a Gaussian FAR of 1 per year, meaning that we do not expect them to come from a Gaussian noise fluctuation. However, there are two notable exceptions which have almost no posterior support for templates with Gaussian FAR under 1 per year, which correspond to GW200308_173609 (grey) and GW200322_091133 (red), having only 4.16% and 0.71% of the posterior samples above this threshold respectively. These are the two events that were noticed in GWTC-3 to have multimodal posterior distributions, due to the likelihood not having a sufficiently large peak to dominate the posterior in all parameter space, which induces prior-dominated modes at large distances and high masses.

In GWTC-3, an ad hoc cut in the likelihood was made to get rid of these prior-dominated modes. For GW200308_173609 the samples with $\log\{\mathcal{L}/\mathcal{L}_0\} < 10$ are removed while for GW200322_091133 the samples with $\log\{\mathcal{L}/\mathcal{L}_0\} < 2$ are removed, where $\mathcal{L}_0 = \exp(-\sum_i(s_i, s_i)/2)$ is the likelihood of the data given no signal, i.e. substituting $h = 0$ in Eq. (64)[64]. We show with dashed lines the contour that encompasses in the $(C, |\rho_{\text{tot}}^{\text{mf}}|)$ plane 90% of the samples that remain after the ad hoc Likelihood cut. We observe that the result is to remove the lowest SNR samples (since the SNR and the Likelihood are intimately related) and it thus removes the posterior samples with the largest FAR. However, a large fraction of the remaining samples still have FARs larger than 1 per year, with 32.9% and 96.8% of them above this threshold for GW200308_173609 and GW200322_091133 respectively.

Looking only at the maximum likelihood sample of these two events (marked with a star in Fig. 12), they have large SNR values of 8.00 for GW200308_173609 and 8.42 for GW200322_091133, which makes them have a single template FAR_{max, L} of $4.7 \times 10^{-4} \text{ yr}^{-1}$ and $9.9 \times 10^{-6} \text{ yr}^{-1}$ respectively, without taking into account any trial factor due to the fact that the likelihood is maximized over a parameter space.

The Gaussian FAR that we have presented here is not directly comparable with the FAR computed by the LVK search pipelines, since they differ in methodology in various ways. The search pipelines make use of a template bank and a different ranking statistic from the bare SNR to take into account the presence of non-Gaussianities. The ranking statistic assigned to each trigger by the pipelines is the one maximized over all the template bank covering the parameter space of the search, with the background estimated by doing time-shifts in detector data. Another difference is that pipelines do not coherently sum the signal from all interferometers, as this would not allow marginalizing over the location in the sky, polarization and neither to work with single detector triggers, making the search computationally cost prohibitive. For this same reason, the template bank of the searches often use simplified waveform models, ignoring effects such as precession or Higher Order modes and do a coarser sampling of the parameter space than what is done in a Parameter Estimation.

In table II we present the most important parameters to quantify the significance of the events in GWTC-3, coming both from the LVK search and PE results and from our Gaussian FAR analysis. Looking at the rightmost column, we notice that there are several events with Gaussian FAPs (computed using Eq. (67) with $T_{\text{obs}} = 1\text{yr}$) that are of order 1. The highest FAPs come, as expected, from GW200308_173609 and GW200322_091133, which have FAPs of 0.97 and 0.99 respectively. After the Likelihood cut, the FAP of GW200308_173609 improves substantially, becoming 0.44. However, that's not the case for GW200322_091133, which keeps a very high FAP after the cut, with a value of 0.97 due to the fact that it has small SNR values in most of its posterior.

Since both GW200308_173609 and GW200322_091133 have a small subset of samples in their posteriors with larger SNRs and correspondingly small FARs, we can explore which samples have this larger significance by selecting only those that have a FAR below a 1yr^{-1} threshold. In Fig. 13 we show the distribution of some of the binary parameters using only those samples with FAR below a 1yr^{-1} . We observe that the parameters of the waveforms that satisfy this cut are very different from all other CBC observations [83], with both events having extremely large effective spin parameters χ_{eff} and with GW200322_091133 having a very extreme mass ratio for which waveform systematics might be important [81]. It's also noticeable that, due to the very low percentage of posterior samples with FAR below the 1yr^{-1} threshold in GW200322_091133, ($\sim 0.07\%$), the parameter space might be undersampled. In principle, both, the search [84] and the parameter estimation [82] should identify similar maximum likelihood points in the parameter space for a given trigger time. We can then compare the two template parameters' values as a sanity check. In the GW200308_173609 case, differences in the masses are not significant, with trigger masses of $(m_1, m_2) = (58.4, 41.3)M_{\odot}$ while the masses identified by the PE for the maximum likelihood template are $(m_1, m_2) = (64.2, 38.2)M_{\odot}$. We find larger discrepancies for the GW200322_091133, with trigger masses of $(m_1, m_2) = (56.0, 15.3)M_{\odot}$ while the masses identified by the PE are $(m_1, m_2) = (161.3, 7.8)M_{\odot}$. The calculation of p_{astro} depends crucially on the values of the masses and such an extreme mass ratio would definitely represent an outlier to the population. For both events, in the search and in the maximum likelihood of the PE, very large values of χ_{eff} are found, in contrast with the rest of the population of merging BH [83]. However, since the value of the spin is not taken into account for p_{astro} calculations [63], this does not downrank the event. Finally, for the case of GW200322_091133 we also find a substantial difference between the search SNR and the maximum likelihood SNR of the PE, being 9.0 and 8.4 respectively. Since the FAR and p_{astro} have an exponential dependence with the SNR, this difference would also downweight the event.

Since our method only gives a lower bound estimation on the FAP, it does not allow us to state that a candidate is indeed a gravitational wave event, but it can support the hypothesis of a noise origin. We can derive how likely Gaussian noise is to generate a signal, but we can not say anything about the possibility of non-Gaussianities mimicking it. With this in mind, in Fig. 14 we show how the GWTC-3 events are distributed in the Gaussian FAP and p_{astro} plane (the values are taken from Table II). We note that for all the 22 events with $p_{\text{astro}} > 0.9$, the Gaussian FAP also gives them low probability of generation from a Gaussian noise fluctuation, having all $\text{FAP} \leq 2 \times 10^{-4}$ and there is no inconsistency. However, for the 13 events with $0.5 < p_{\text{astro}} < 0.9$, results are mixed. The majority of these events (8/13) also have Gaussian FAP smaller than 10% and so we find that they are not likely to be generated from a Gaussian noise fluctuation. From the 13 events with $0.5 < p_{\text{astro}} < 0.9$ we have another 3 in the region of $10\% < \text{FAP} < 50\%$, which therefore have some non-negligible probability of being generated by Gaussian noise, although it is still more likely they are not. These 3 events correspond to GW191113_071753, GW200208_222617 and GW200220_061928, from which GW200208_222617 is the one with the largest Gaussian FAP ($\sim 31\%$) and also has a multimodal posterior distribution [11]. Finally, at $0.5 < p_{\text{astro}} < 0.9$ and $\text{FAP} > 50\%$ we have 2 points corresponding to GW200308_173609 and GW200322_091133 and which have already been discussed in detail as likely to be generated by a Gaussian fluctuation. It is also interesting to note that for all events with $\text{FAP} > 0.1$, the p_{astro} value quoted in GWTC-3 [11] is larger than 0.5 in only one of the pipelines, the others quoting significantly lower values.

Event	IFOs	Pipeline	Search SNR	Search $\log_{10}(\text{FAR} \cdot \text{yr})$	p_{astro}	PE SNR	Gaussian $\log_{10}(\text{FAR} \cdot \text{yr})$	Gaussian FAP _{event} ($T_{\text{obs}} = 1\text{yr}$)
GW191103_012549	HL	PyCBC-BBH	9.3	-0.34	0.94	$8.9_{-0.5}^{+0.3}$	$-6.40_{-1.08}^{+1.92}$	$(1.99 \pm 0.24) \times 10^{-4}$
GW191105_143521	HLV	PyCBC-broad	9.8	-1.92	> 0.99	$9.7_{-0.5}^{+0.3}$	$-9.44_{-1.32}^{+2.08}$	$(2.6 \pm 1.4) \times 10^{-6}$
GW191109_010717	HL	MBTA	15.2	-3.74	> 0.99	$17.3_{-0.5}^{+0.5}$	$-54.23_{-4.04}^{+3.34}$	$(1.96 \pm 0.56) \times 10^{-50}$
GW191113_071753	HLV	MBTA	9.2	1.41	0.68	$7.8_{-1.1}^{+0.6}$	$-2.40_{-2.14}^{+3.35}$	0.15724 ± 0.00085
GW191126_115259	HL	PyCBC-BBH	8.5	0.51	0.70	$8.3_{-0.5}^{+0.2}$	$-4.27_{-0.76}^{+1.68}$	$(4.05 \pm 0.12) \times 10^{-3}$
GW191127_050227	HLV	PyCBC-BBH	8.7	0.61	0.74	$9.1_{-0.6}^{+0.5}$	$-7.77_{-2.23}^{+2.29}$	$(2.35 \pm 0.54) \times 10^{-5}$
GW191129_134029	HL	GstLAL	13.3	< -5	> 0.99	$13.2_{-0.3}^{+0.2}$	$-26.58_{-1.21}^{+1.60}$	$(1.90 \pm 0.36) \times 10^{-25}$
GW191204_110529	HL	PyCBC-BBH	8.9	0.52	0.74	$8.8_{-0.6}^{+0.4}$	$-6.15_{-1.62}^{+2.28}$	$(4.68 \pm 0.36) \times 10^{-4}$
GW191204_171526	HL	PyCBC-broad	17.1	< -5	> 0.99	$17.5_{-0.2}^{+0.2}$	$-55.15_{-1.40}^{+1.80}$	$(1.7 \pm 1.4) \times 10^{-52}$
GW191215_223052	HLV	GstLAL	10.9	< -5	> 0.99	$11.2_{-0.4}^{+0.3}$	$-16.38_{-1.59}^{+2.00}$	$(9.5 \pm 2.7) \times 10^{-15}$
GW191216_213338	HV	GstLAL	18.6	< -5	> 0.99	$18.6_{-0.2}^{+0.2}$	$-63.74_{-1.47}^{+1.81}$	$(8.1 \pm 2.3) \times 10^{-62}$
GW191219_163120	HLV	PyCBC-broad	8.9	0.60	0.82	$9.1_{-0.8}^{+0.5}$	$-7.61_{-2.07}^{+3.00}$	$(2.29 \pm 0.33) \times 10^{-3}$
GW191222_033537	HL	GstLAL	12	< -5	> 0.99	$12.5_{-0.3}^{+0.2}$	$-23.29_{-1.12}^{+1.53}$	$(2.2 \pm 2.0) \times 10^{-21}$
GW191230_180458	HLV	PyCBC-BBH	9.9	-0.38	0.96	$10.5_{-0.4}^{+0.2}$	$-13.48_{-1.09}^{+1.74}$	$(3.6 \pm 3.5) \times 10^{-10}$
GW200112_155838	LV	GstLAL	17.6	< -5	> 0.99	$19.8_{-0.2}^{+0.1}$	$-74.28_{-1.17}^{+1.79}$	$(1.82 \pm 0.79) \times 10^{-72}$
GW200115_042309	HLV	GstLAL	11.5	< -5	> 0.99	$11.3_{-0.5}^{+0.3}$	$-16.69_{-1.51}^{+2.43}$	$(8.1 \pm 5.1) \times 10^{-14}$
GW200128_022011	HL	PyCBC-BBH	9.9	-2.37	> 0.99	$10.7_{-0.4}^{+0.3}$	$-14.16_{-1.38}^{+1.66}$	$(4.92 \pm 0.78) \times 10^{-13}$
GW200129_065458	HLV	GstLAL	26.5	< -5	> 0.99	$26.8_{-0.2}^{+0.2}$	$-144.95_{-2.21}^{+2.39}$	$(5.94 \pm 0.96) \times 10^{-143}$
GW200202_154313	HLV	GstLAL	11.3	< -5	> 0.99	$10.9_{-0.4}^{+0.2}$	$-14.63_{-1.05}^{+1.76}$	$(3.9 \pm 2.2) \times 10^{-11}$
GW200208_130117	HLV	PyCBC-BBH	10.8	-3.51	> 0.99	$10.9_{-0.4}^{+0.2}$	$-15.04_{-1.13}^{+1.96}$	$(3.4 \pm 2.0) \times 10^{-11}$
GW200208_222617	HLV	PyCBC-BBH	7.9	0.68	0.70	$7.4_{-2.0}^{+1.1}$	$-1.41_{-3.95}^{+5.28}$	0.31395 ± 0.00090
GW200209_085452	HLV	MBTA	9.7	1.08	0.97	$9.6_{-0.5}^{+0.3}$	$-9.67_{-1.38}^{+1.99}$	$(2.3 \pm 1.9) \times 10^{-6}$
GW200210_092254	HLV	PyCBC-BBH	8.9	0.89	0.54	$8.4_{-0.7}^{+0.5}$	$-4.66_{-1.84}^{+2.50}$	$(1.169 \pm 0.025) \times 10^{-2}$
GW200216_220804	HLV	GstLAL	9.4	-0.45	0.77	$8.2_{-0.5}^{+0.3}$	$-4.24_{-1.05}^{+1.72}$	$(2.948 \pm 0.095) \times 10^{-3}$
GW200219_094415	HLV	GstLAL	10.7	-3.00	> 0.99	$10.7_{-0.4}^{+0.3}$	$-14.45_{-1.33}^{+1.98}$	$(1.4 \pm 1.2) \times 10^{-11}$
GW200220_061928	HLV	PyCBC-BBH	7.5	0.83	0.62	$7.3_{-0.7}^{+0.4}$	$-1.66_{-1.13}^{+1.94}$	0.13003 ± 0.00070
GW200220_124850	HL	MBTA	8.2	-2.74	0.83	$8.5_{-0.5}^{+0.3}$	$-5.30_{-1.02}^{+1.74}$	$(5.41 \pm 0.47) \times 10^{-4}$
GW200224_222234	HLV	MBTA	19.0	< -5	> 0.99	$20.0_{-0.2}^{+0.2}$	$-75.77_{-1.41}^{+1.84}$	$(7.9 \pm 4.3) \times 10^{-74}$
GW200225_060421	HL	PyCBC-broad	12.3	< -5	> 0.99	$12.5_{-0.4}^{+0.3}$	$-23.14_{-1.57}^{+1.88}$	$(1.44 \pm 0.35) \times 10^{-21}$
GW200302_015811	HV	GstLAL	10.6	-0.96	0.91	$10.8_{-0.4}^{+0.3}$	$-14.76_{-1.64}^{+1.94}$	$(3.11 \pm 0.76) \times 10^{-13}$
GW200306_093714	HL	MBTA	8.5	2.61	0.81	$7.8_{-0.6}^{+0.3}$	$-2.46_{-1.19}^{+1.95}$	$(4.933 \pm 0.047) \times 10^{-2}$
GW200308_173609	HLV	PyCBC-BBH	8.0	0.38	0.86	$3.8_{-2.5}^{+3.1}$	$6.55_{-6.17}^{+1.91}$	0.96500 ± 0.00045
GW200308_173609*	-	-	-	-	-	$7.09_{-0.50}^{+0.47}$	$-0.90_{-3.31}^{+3.34}$	0.4366 ± 0.0040
GW200311_115853	HLV	GstLAL	17.7	< -5	> 0.99	$17.9_{-0.2}^{+0.1}$	$-58.41_{-1.12}^{+1.69}$	$(6.7 \pm 1.8) \times 10^{-57}$
GW200316_215756	HLV	GstLAL	10.1	< -5	> 0.99	$10.3_{-0.7}^{+0.4}$	$-12.24_{-1.81}^{+2.82}$	$(2.5 \pm 2.0) \times 10^{-8}$
GW200322_091133	HLV	MBTA	9.0	2.65	0.62	$2.5_{-1.7}^{+3.4}$	$8.00_{-5.35}^{+0.75}$	0.99327 ± 0.00021
GW200322_091133*	-	-	-	-	-	$5.3_{-0.9}^{+1.4}$	$9.15_{-7.37}^{+4.14}$	0.96870 ± 0.00096

TABLE II: In this table we report every candidate GW signal included in the O3b Catalog, as well as the detectors observing at the merger time of the events, the search pipeline in which it had the highest p_{astro} together with the Search estimated SNR, the Search FAR and the p_{astro} as calculated by that same pipeline. We also include the SNR as obtained by the LVK parameter estimation analysis, our Gaussian FAR and Gaussian FAP of the event assuming an observing time of 1yr. While the errors on the PE SNR and the Gaussian FAR represent the 90% credible intervals, for the Gaussian FAP they represent the uncertainty on the Monte Carlo integral used to compute it, given by Eq. (68). The events that have an asterisk and are in *italic*, correspond to the ones in which we have performed the ad hoc cut in the Likelihood.

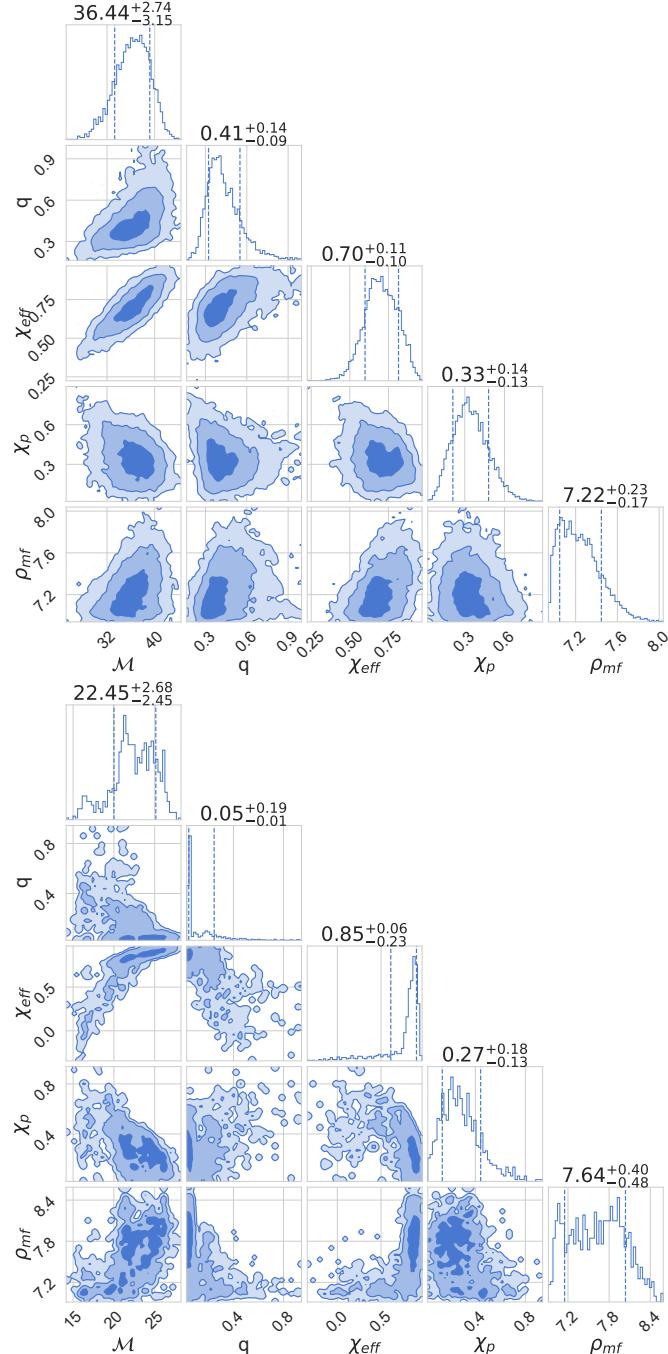


FIG. 13: Corner plots of selected parameters for the posterior samples with Gaussian FAR ≤ 1 per year. Top Panel: GW200308_173609, Bottom Panel: GW200322_091133

H. Conclusions

Understanding whether triggers in LIGO-Virgo detectors are from gravitational wave or noise origin is a hard task. For most of the events, the GW signal is expected to be extremely weak and in this paper we have explored the possibility of it being mimicked by the irreducible Gaussian noise in the gravitational wave detectors.

We have derived a mathematical framework for estimating the rate of false alarms induced by this Gaussian noise. Our main result is given in Eq. (46), which gives the rate at which the matched filter SNR of a specific template with the Gaussian noise of one (or multiple) GW detectors goes over a threshold ρ . The prefactor C multiplying the FAR depends on the specific template used for matched filtering. For CBC templates the most important parameter

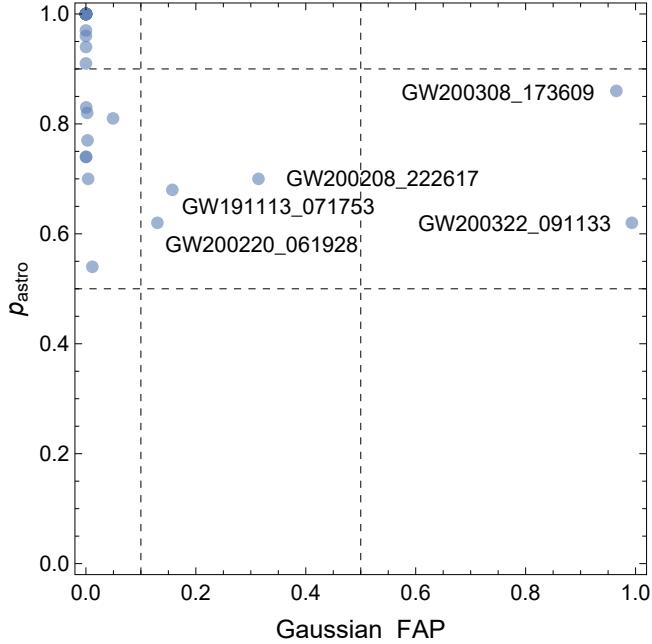


FIG. 14: The comparison in the Gaussian FAP and p_{astro} plane of the different GWTC-3 events from Table II. We have set cuts in FAP of 10% and 50%, as well as in $p_{\text{astro}} = 0.5$ and 0.9.

controlling the value of C is the total mass of the event, with C being significantly smaller for larger masses.

We have then studied how the Gaussian FAR of CBC templates behaves as a function of the threshold SNR, and gave an analytical expression for the minimum SNR needed for a given FAR threshold. We have also proposed a method to estimate the probability of Gaussian noise with the local PSD mimicking a given GW candidate in terms of a false alarm probability (66), using the samples from the Parameter Estimation analysis of such an event.

Finally, we have applied this formalism to the GW candidates that were added in the GWTC-3 catalog, obtaining a Gaussian FAR for each template in their PE posterior and a FAP for the events.

Summarizing these results, most of the samples of the events are clearly above the 1 yr^{-1} FAR threshold with event FAPs ranging from $\sim 10^{-143}$ to a more modest $\sim 10^{-1}$, assuming a reference observation time of one year. However, we find two clear outliers, GW200308_173609 and GW200322_091133, with event FAPs very close to one, signaling very high odds of Gaussian noise fluctuations mimicking them. We also explore the samples in their posterior that have single template $\text{FAR} < 1 \text{ yr}^{-1}$. These samples have very extreme parameter values with respect to the observed BBH population, and in the case of GW200322_091133 differ from those identified by the search.

We believe that the methods developed here may be useful in the future to further investigate GW triggers that are found in future LVK runs.

I. Appendix: Study of the FAP for the bivariate complex Gaussian

In this section we will study the FAP for the bivariate complex Gaussian (FAP_2) whose probability density function is given in Eq. (25). We will obtain Eq. (87) to numerically compute FAP_2 in an efficient and well behaved manner. We also obtain a prescription to analytically approximate the FAP_2 to arbitrary order in $1 - |\alpha|$ using Eq. (91). With this expansion we obtain the leading order and second order approximations of Eq. (88) and Eq. (93) respectively and shown in Fig. 15. As seen in Eq. (26), FAP_2 is given by the following expression:

$$\begin{aligned}
\text{FAP}_2 &= P(\rho_1 > \rho \cup \rho_2 > \rho) = 1 - P(\rho_1 < \rho \cap \rho_2 < \rho) \\
&= 1 - \frac{1}{(2\pi)^2(1-|\alpha|^2)} \int_0^{2\pi} d\theta_1 \int_0^{2\pi} d\theta_2 \int_0^\rho \rho_1 d\rho_1 \int_0^\rho \rho_2 d\rho_2 \exp \left\{ \left(\frac{\rho_1^2 + \rho_2^2 - 2|\alpha|\rho_1\rho_2 \cos(\theta_\alpha - \theta_1 + \theta_2)}{2(1-|\alpha|^2)} \right) \right\} \\
&= 1 - \frac{1}{2\pi(1-|\alpha|^2)} \int_0^\rho d\rho_1 \iint_0^\rho d\rho_2 \rho_1 \rho_2 \exp \left\{ -\frac{\rho_1^2 + \rho_2^2}{2(1-|\alpha|^2)} \right\} \int_0^{2\pi} d\theta \exp \left\{ \frac{|\alpha|\rho_1\rho_2}{1-|\alpha|^2} \cos \theta \right\} \\
&= 1 - \frac{1}{1-|\alpha|^2} \int_0^\rho d\rho_1 \iint_0^\rho d\rho_2 \rho_1 \rho_2 \exp \left\{ -\frac{\rho_1^2 + \rho_2^2}{2(1-|\alpha|^2)} \right\} I_0 \left\{ \frac{|\alpha|}{1-|\alpha|^2} \rho_1 \rho_2 \right\}, \tag{69}
\end{aligned}$$

where for notation simplicity we define $\alpha = \Gamma(\Delta t)$ and $I_n(z)$ is the modified Bessel function of the first kind [85]:

$$I_n(z) \equiv i^{-n} J_n(ix) = \frac{1}{\pi} \iint_0^\pi d\theta e^{z \cos \theta} \cos(n\theta) = \sum_{k=0}^{\infty} \frac{(\frac{1}{2}z)^{2k+n}}{k!(k+n)!} \quad (n \in \mathbb{Z}). \tag{70}$$

The integral of Eq. (69) can be further simplified by making the change of variables:

$$\rho_i = \sqrt{2(1-|\alpha|^2)u_i} \rightarrow d\rho_i = \sqrt{\frac{-|\alpha|^2}{2u_i}} \, , \tag{71}$$

which yields:

$$\text{FAP}_2 = 1 - (1-|\alpha|^2) \int_0^x du_1 \int_0^x du_2 I_0(2|\alpha|\sqrt{u_1 u_2}) e^{-(u_1+u_2)}. \tag{72}$$

where for notation simplicity we have defined:

$$x \equiv \frac{\rho^2}{2(1-|\alpha|^2)}. \tag{73}$$

From Eq. (70) we have that the Taylor series of $I_0(z)$ around $z = 0$ is given by:

$$I_0(z) = \sum_{k=0}^{\infty} \frac{z^{2k}}{2^{2k}(k!)^2} \rightarrow I_0(2|\alpha|\sqrt{u_1 u_2}) = \sum_{k=0}^{\infty} \frac{|\alpha|^{2k} u_1^k u_2^k}{(k!)^2}. \tag{74}$$

And substituting this expansion into Eq. (72) we obtain:

$$\text{FAP}_2 = 1 - (1-|\alpha|^2) \sum_{k=0}^{\infty} |\alpha|^{2k} \left[\frac{1}{k!} \iint_0^\pi u^k e^{-u} du \right]^2. \tag{75}$$

Since k is a natural number, the integral appearing in Eq. (75) is given by:

$$\frac{1}{k!} \iint_0^\pi u^k e^{-u} du = 1 - e^{-x} \sum_{n=0}^k \frac{x^n}{n!}, \tag{76}$$

Using this in Eq. (75), the FAP_2 will be given by:

$$\begin{aligned}
\text{FAP}_2 &= 1 - (1-|\alpha|^2) \sum_{k=0}^{\infty} |\alpha|^{2k} \left[1 - e^{-x} \sum_{n=0}^k \frac{x^n}{n!} \right]^2 \\
&= 1 - (1-|\alpha|^2) \left[\sum_{k=0}^{\infty} \left(|\alpha|^{2k} - 2e^{-x} \sum_{k=0}^{\infty} \sum_{n=0}^k \left(|\alpha|^{2k} \frac{x^n}{n!} + e^{-2x} \sum_{k=0}^{\infty} \sum_{n=0}^k \sum_{m=0}^k \left(|\alpha|^{2k} \frac{x^{n+m}}{n!m!} \right) \right) \right) \right]. \tag{77}
\end{aligned}$$

In the first sum of Eq. (77) we recognize a simple geometric series. Taking into account that $|\alpha|^2 < 1$, it will converge to the following expression:

$$\sum_{k=0}^{\infty} |\alpha|^{2k} = \frac{1}{1 - |\alpha|^2}. \quad (78)$$

The second sum of Eq. (77) can also be summed exactly by making some index manipulation:

$$\sum_{k=0}^{\infty} |\alpha|^{2k} \sum_{n=0}^k \frac{x^n}{n!} = \sum_{n=0}^{\infty} \frac{x^n}{n!} \sum_{k=n}^{\infty} |\alpha|^{2k} = \sum_{n=0}^{\infty} \frac{(|\alpha|^2 x)^n}{n!} \sum_{k=0}^{\infty} |\alpha|^{2k} = e^{|\alpha|^2 x} \frac{1}{1 - |\alpha|^2}. \quad (79)$$

Finally, the third sum of Eq. (77) can not be summed exactly, but it can be significantly simplified by making similar index manipulations:

$$\begin{aligned} \sum_{k=0}^{\infty} |\alpha|^{2k} \sum_{n=0}^k \sum_{m=0}^k \frac{x^{n+m}}{n!m!} &= \sum_{n=0}^{\infty} \sum_{m=0}^{\infty} \frac{x^{n+m}}{n!m!} \sum_{k=\max(n,m)}^{\infty} |\alpha|^{2k} = \sum_{k=0}^{\infty} |\alpha|^{2k} \sum_{n=0}^{\infty} \sum_{m=0}^{\infty} |\alpha|^{2\max(n,m)} \frac{x^{n+m}}{n!m!} \\ &= \frac{1}{1 - |\alpha|^2} \sum_{n=0}^{\infty} \sum_{m=0}^{\infty} |\alpha|^{2\max(n,m)} \frac{x^{n+m}}{n!m!}. \end{aligned} \quad (80)$$

Substituting the results of the sums of Eqs. (78), (79), (80) into Eq. (77), we obtain the following result:

$$\text{FAP}_2 = 2e^{-(1-|\alpha|^2)x} - e^{-2x} \sum_{n=0}^{\infty} \sum_{m=0}^{\infty} |\alpha|^{2\max(n,m)} \frac{x^{n+m}}{n!m!}. \quad (81)$$

To further simplify this expression we can change indices in the sum of Eq. (81), using $l = n - m$ and $k = \frac{1}{2}(n + m)$:

$$\sum_{n=0}^{\infty} \sum_{m=0}^{\infty} |\alpha|^{2\max(n,m)} \frac{x^{n+m}}{n!m!} = \sum_{l=-\infty}^{\infty} \sum_{k=|l|/2}^{\infty} |\alpha|^{2k+|l|} \frac{x^{2k}}{(k+\frac{l}{2})! (k-\frac{l}{2})!} = S_0 + 2 \sum_{l=1}^{\infty} S_l, \quad (82)$$

where we have used that $2 \max(n, m) = n + m + |n - m| = 2k + |l|$ and we have defined:

$$S_l = \sum_{k=l/2}^{\infty} |\alpha|^{2k+l} \frac{x^{2k}}{(k+\frac{l}{2})! (k-\frac{l}{2})!} = |\alpha|^l \underbrace{\sum_{k=0}^{\infty} \frac{(|\alpha|x)^{2k+l}}{(k+l)!k!}}_{I_l(2|\alpha|x)} = |\alpha|^l I_l(2|\alpha|x), \quad (83)$$

where we have identified the Taylor series of the modified Bessel function of the first kind of order l shown in Eq. (70). Using Eq. (83) and Eq. (82) we have that the FAP_2 of Eq. (81) will be given by:

$$\text{FAP}_2 = 2e^{-(1-|\alpha|^2)x} - e^{-2x} I_0(2|\alpha|x) + 2 \sum_{n=1}^{\infty} |\alpha|^n I_n(2|\alpha|x). \quad (84)$$

To compute the sum of modified bessel functions of the first kind, we can use their integral representation, shown in Eq. (70):

$$\begin{aligned} I_0(z) + 2 \sum_{n=1}^{\infty} |\alpha|^n I_n(z) &= \frac{1}{\pi} \iint_0^{\pi} d\theta e^{z \cos \theta} \left[1 + 2 \sum_{n=1}^{\infty} |\alpha|^n \cos(n\theta) \right] = \frac{1}{\pi} \iint_0^{\pi} d\theta e^{z \cos \theta} \left[1 + \sum_{n=1}^{\infty} (|\alpha| e^{i\theta})^n + (|\alpha| e^{-i\theta})^n \right] \\ &= \frac{1}{\pi} \int_0^{\pi} d\theta e^{z \cos \theta} \left[1 + \frac{|\alpha| e^{i\theta}}{1 - |\alpha| e^{i\theta}} + \frac{|\alpha| e^{-i\theta}}{1 - |\alpha| e^{-i\theta}} \right] = \frac{1}{\pi} \int_0^{\pi} d\theta e^{z \cos \theta} \frac{1 - |\alpha|^2}{1 - 2|\alpha| \cos \theta + |\alpha|^2}. \end{aligned} \quad (85)$$

And we have transformed the infinite sum in a definite integral of a relatively simple function. The integral can be expressed in a more simple and convenient way if we do the variable change $\theta = 2 \arctan\left(\frac{1-|\alpha|}{1+|\alpha|}u\right)$:

$$I_0(z) + 2 \sum_{n=1}^{\infty} |\alpha|^n I_n(z) = \frac{2}{\pi} e^z \int_0^{\infty} du \frac{1}{1+u^2} \exp \left\{ \left(2z \frac{(1-|\alpha|)^2 u^2}{(1+|\alpha|)^2 + (1-|\alpha|)^2 u^2} \right) \right\}. \quad (86)$$

Substituting this expression for the sum into Eq. (84) for the FAP_2 and using the fact that $z = 2|\alpha|x$, where x is defined in Eq. (73), we obtain:

$$\text{FAP}_2 = 2e^{-\rho^2/2} - \frac{2}{\pi} e^{-\rho^2/(1+|\alpha|)} \int_0^{\infty} du \frac{1}{1+u^2} \exp \left\{ \left(\frac{2|\alpha|(1-|\alpha|)\rho^2}{(1+|\alpha|)^3} \frac{u^2}{1 + (\frac{1-|\alpha|}{1+|\alpha|})^2 u^2} \right) \right\}. \quad (87)$$

The integral in this expression can not be analytically computed, but it can be numerically integrated as it is a well behaved one variable definite integral that does not suffer from divergences or accuracy problems due to large cancellations, as the previous integrals did. We can check that this formula has the correct limiting behavior if we realize that both when $|\alpha| = 0$ and when $|\alpha| = 1$, the argument of the exponential inside the integral of Eq. (87) vanishes and the value of the integral is $\pi/2$. Therefore in the case in which $|\alpha| = 0$, when there is no correlation, $\text{FAP}_2(|\alpha| = 0) = 2e^{-\rho^2/2} - e^{-\rho^2} = 1 - (1 - e^{-\rho^2/2})^2$ as is expected from two uncorrelated variables. In the opposite limit, when the correlation is maximal and $|\alpha| = 1$, FAP_2 coincides with the expected result in which the two variables behave as a single one, that is, $\text{FAP}_2(|\alpha| = 1) = e^{-\rho^2/2} = 1 - (1 - e^{-\rho^2/2})^1$.

As seen in Sec. IV E of the main text, we are interested in obtaining an approximation in the limit in which the correlation is large and thus $|\alpha| \rightarrow 1$. However, we will take into account that the SNR threshold ρ can be large in such a way that $(1-|\alpha|)\rho^2$ can be of order $O(1)$. In this case, an upper bound approximation for the FAP_2 is obtained in the following way:

$$\begin{aligned} \text{FAP}_2 &\approx 2e^{-\rho^2/2} - \frac{2}{\pi} e^{-\rho^2/(1+|\alpha|)} \int_0^{\infty} du \frac{1}{1+u^2} \exp \left\{ -\frac{2|\alpha|(1-|\alpha|)\rho^2}{(1+|\alpha|)^3} u^2 \right\} \left(\right. \\ &= e^{-\rho^2/2} \left[\left(-\exp \left\{ \left(\frac{1}{2} \left(\frac{1-|\alpha|}{1+|\alpha|} \right)^3 \rho^2 \right) \right\} \text{Erfc} \left\{ \rho \sqrt{\frac{2|\alpha|(1-|\alpha|)}{(1+|\alpha|)^3}} \right\} \right] \\ &\approx e^{-\rho^2/2} \left[1 + \text{Erf} \left\{ \frac{\rho \sqrt{1-|\alpha|}}{2} \right\} \right], \end{aligned} \quad (88)$$

where we have used that [85]:

$$\frac{2}{\pi} \int_0^{\infty} \frac{du}{1+u^2} e^{-\eta^2 u^2} = e^{-\eta^2} \text{Erfc}(\eta), \quad (89)$$

and where $\text{Erf}(z)$ and $\text{Erfc}(z)$ are the error function and the complementary error function respectively. Eq. (88) can be taken to be as the leading order term in an expansion in $1-|\alpha|$ of the FAP_2 . To analyze higher order terms it will be convenient to introduce two new variables:

$$\eta = \rho \sqrt{\frac{2|\alpha|(1-|\alpha|)}{(1+|\alpha|)^3}}, \quad (90a)$$

$$\epsilon = \frac{1-|\alpha|}{1+|\alpha|}. \quad (90b)$$

In the regime we are interested, η is of order $O(1)$, while $\epsilon \ll 1$. Using these variables we have:

$$\begin{aligned}
\text{FAP}_2 &= e^{-\rho^2/2} \left[2 - e^{-\frac{1-|\alpha|}{2(1+|\alpha|)}\rho^2} \frac{2}{\pi} \int_0^\infty \frac{du}{1+u^2} \exp \left\{ \left(\frac{\eta^2 u^2}{1+\epsilon^2 u^2} \right) \right\} \right] \\
&= e^{-\rho^2/2} \left[2 - e^{-\frac{1-|\alpha|}{2(1+|\alpha|)}\rho^2} \frac{2}{\pi} \int_0^\infty \frac{du}{1+u^2} e^{-\eta^2 u^2} \exp \left\{ \left(\frac{\epsilon^2 \eta^2 u^4}{1+\epsilon^2 u^2} \right) \right\} \right] \left(\right. \\
&= e^{-\rho^2/2} \left[2 - e^{-\frac{1-|\alpha|}{2(1+|\alpha|)}\rho^2} \frac{2}{\pi} \int_0^\infty \frac{du}{1+u^2} e^{-\eta^2 u^2} \sum_{n=0}^\infty \frac{1}{n!} \left(\frac{\epsilon^2 \eta^2 u^4}{1+\epsilon^2 u^2} \right)^n \right]. \tag{91}
\end{aligned}$$

If we truncate the sum at n -th order, we obtain an upper bound approximation that is accurate to order $(\eta\epsilon)^{2n}$ and that has correct limiting behavior when $\epsilon \rightarrow 0$, when $\epsilon = 1$, when $\eta = 0$ and when $\eta \rightarrow \infty$. Since we want only the first order correction, we can keep terms up to $n = 1$ and integrate, obtaining:

$$\begin{aligned}
\text{FAP}_2 &\approx e^{-\rho^2/2} \left[2 - e^{-\frac{1-|\alpha|}{2(1+|\alpha|)}\rho^2} \frac{2}{\pi} \int_0^\infty \frac{du}{1+u^2} e^{-\eta^2 u^2} \left(1 + \frac{\epsilon^2 \eta^2 u^4}{1+\epsilon^2 u^2} \right) \right] \left(\right. \\
&= e^{-\rho^2/2} \left[2 - e^{-\frac{1-|\alpha|}{2(1+|\alpha|)}\rho^2} \left(\left(1 + \frac{\epsilon^2 \eta^2}{1-\epsilon^2} \right) \left(\eta^2 \text{Erfc}(\eta) + \frac{\eta}{\sqrt{\pi}} - \frac{\eta^2}{\epsilon(1-\epsilon^2)} e^{\eta^2/\epsilon^2} \text{Erfc} \left(\frac{\eta}{\epsilon} \right) \right) \right) \right] \left(\right. \\
&\approx e^{-\rho^2/2} \left[2 - e^{-\frac{1-|\alpha|}{2(1+|\alpha|)}\rho^2} \left(\left(1 + \epsilon^2 \eta^2 \right) \left(\eta^2 \text{Erfc}(\eta) - \epsilon^2 \frac{\eta}{\sqrt{\pi}} \left(1 - \frac{1}{2\eta^2} \right) \right) \right) \right]. \tag{92}
\end{aligned}$$

We can express this result in terms of the correlation $|\alpha|$ and the SNR threshold ρ substituting the expressions for η and ϵ of Eq. (90). To be consistent in the approximation, we keep the two first orders in $1 - |\alpha|$, assuming that $(1 - |\alpha|)\rho^2$ is of order $O(1)$. Doing this we obtain:

$$\begin{aligned}
\text{FAP}_2 &\approx e^{-\rho^2/2} \left[1 + \text{Erf} \left\{ \frac{1}{2} \rho \sqrt{1-|\alpha|} \left(\left(1 + \frac{1-|\alpha|}{4} - \frac{(1-|\alpha|)^2}{32} \right) \right\} \left(\frac{(1-|\alpha|)^{3/2}}{4\sqrt{\pi}\rho} e^{-\frac{1}{4}(1-|\alpha|)\rho^2} \left(1 - \frac{(1-|\alpha|)\rho^2}{2} \right) \right] \left(\right. \\
&\approx e^{-\rho^2/2} \left[1 + \text{Erf} \left\{ \frac{1}{2} \rho \sqrt{1-|\alpha|} \left(1 + \frac{1-|\alpha|}{4} \left(1 - \frac{1}{\rho^2} \right) + \frac{3(1-|\alpha|)^2}{32} \right) \right\} \right] \left(\right. \tag{93}
\end{aligned}$$

where for simplicity of the final result, in the last step we have introduced all the corrections inside the argument of the error function in a way that is consistent with the order of the approximation. We check that ignoring the higher order corrections in $1 - |\alpha|$, we recover the leading order expression of (88).

In Fig. 15 we show the relative error, between the exact FAP_2 computed using Eq. (87) and the approximations of Eq. (88) (left panel) and Eq. (93) (right panel), as a function of the correlation $|\alpha|$ and the SNR threshold ρ . We observe that the leading order approximation (left panel), already gives an accurate description of the FAP_2 , having sub-percent accuracy for $\rho \gtrsim 5$ and reproducing the exact result as $|\alpha| \rightarrow 1$. On the right hand panel we can see the effect of introducing the higher order correction, we observe that the description is now much improved, reaching an accuracy better than 1 part in 10000 for $\rho \gtrsim 4$ and describing much better the limit $|\alpha| \rightarrow 1$. If we wanted to approximate the FAP_2 to higher precision, we could take into account more terms in the sum of Eq. (91) and analytically integrate them using Eq. (89).

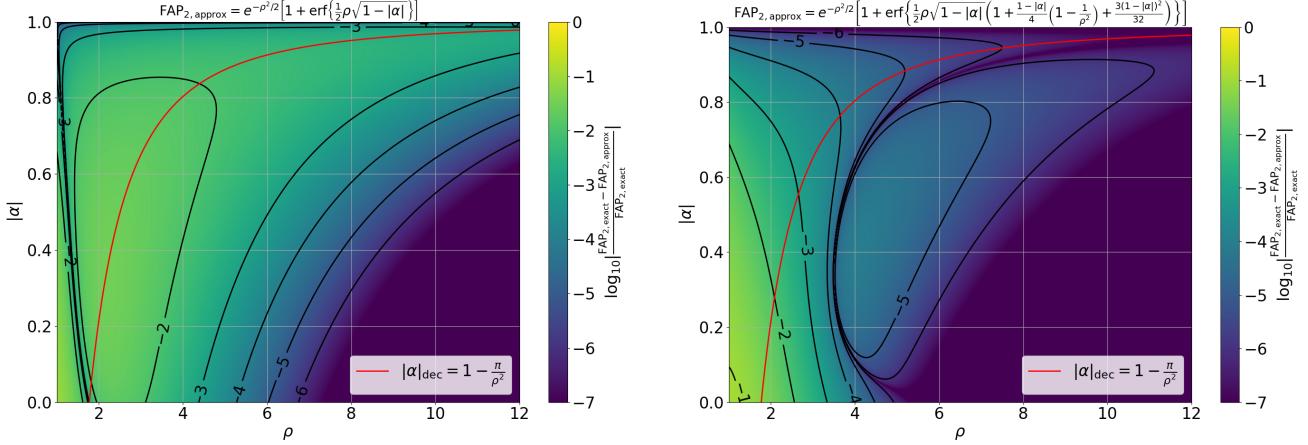


FIG. 15: Base 10 logarithm of the relative error between the exact value of FAP_2 , computed using Eq. (87) and the approximations proposed in Eq. (88) (left panel) and Eq. (93) (right panel), as a function of the correlation $|\alpha|$ and the SNR threshold ρ . We also show the with a red line the value of $|\alpha|$ at the decoupling time, to get an idea of the region where we are interested in having a good approximation as a function of ρ .

V Efficient Reduced Order Quadrature Construction Algorithms for Fast Gravitational Wave Inference

A. Journal & publication details

Published in Physical Review D 108, 123025 (2023) [3]

B. Authors

The authors in the same order as presented to the journal are

Gonzalo Morrás^{a13}

Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco 28049 Madrid, Spain

José Francisco Nuño Siles^{a14}

Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco 28049 Madrid, Spain

Juan García-Bellido^{a15}

Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco 28049 Madrid, Spain

C. Abstract

Reduced Order Quadrature (ROQ) methods can greatly reduce the computational cost of Gravitational Wave (GW) likelihood evaluations, and therefore greatly speed up parameter estimation analyses, which is a vital part to maximize the science output of advanced GW detectors. In this paper, we do an in-depth study of ROQ techniques applied to GW data analysis and present novel algorithms to enhance different aspects of the ROQ bases construction. We improve upon previous ROQ construction algorithms allowing for more efficient bases in regions of parameter space that were previously challenging. In particular, we use singular value decomposition (SVD) methods to characterize the waveform space and choose a reduced order basis close to optimal and also propose improved methods for empirical interpolation node selection, greatly reducing the error added by the empirical interpolation model. To demonstrate the effectiveness of our algorithms, we construct multiple ROQ bases ranging in duration from 4s to 256s for compact

¹³ gonzalo.morras@uam.es

¹⁴ josef.nunno@estudiante.uam.es

¹⁵ juan.garciabellido@uam.es

binary coalescence (CBC) waveforms including precession and higher order modes. We validate these bases by performing likelihood error tests and P-P tests and explore the speed up they induce both theoretically and empirically with positive results. Furthermore, we conduct end-to-end parameter estimation analyses on several confirmed GW events, showing the validity of our approach in real GW data.

D. Introduction

Gravitational wave (GW) astronomy has been made possible in recent years by ground-based observatories like LIGO [86], Virgo [87], and KAGRA [88], revolutionizing our understanding of the Universe by enabling the direct detection of GW signals emitted during extreme cosmic phenomena such as the mergers of binary black holes, binary neutron stars, and neutron star-black hole binaries. With the continuous improvement in sensitivity of current detectors [89] and the advent of next-generation detectors, including projects like the Einstein Telescope [90], Cosmic Explorer [91], LISA [92–94], we anticipate a dramatic increase in the number of GW candidates detected. For maximum science outputs, a parameter estimation (PE) for each candidate will have to be performed. With standard PE methods [95], this can be prohibitively computationally expensive, especially as we reduce the frequency from which we can detect gravitational waves and the duration of the signals becomes much longer [69].

To fully exploit the enhanced sensitivity of these advanced detectors, it is essential to use accurate waveform models that incorporate important physical effects such as precession or higher-order modes [96]. However, the computational challenge of calculating the likelihood of such signals poses a significant bottleneck in the analysis pipeline. Traditional likelihood calculations can be computationally intensive, particularly for long-duration waveforms. Several methods have been explored in the literature to reduce this computational burden, such as multi-banding [97], heterodyned likelihood [98, 99], likelihood-free approaches [100, 101], Reduced Order Quadrature methods [102–107] and others [108–110].

In this work, we will focus on the ROQ method, which is one of the most promising approaches to fast GW likelihood evaluations, due to its ability to achieve very large speed-ups while maintaining high accuracy and being able to accommodate the effects of precession and higher-order modes [104, 105]. ROQ methods exploit the fact that for a given parameter range, the corresponding GW waveforms span only a small subspace of the vector space of all possible signals. By constructing reduced bases that capture the essential information of the templates, ROQ techniques provide an efficient representation that enables fast likelihood evaluations. The ROQ has a *start-up* cost associated with the offline basis building stage, which needs to be performed in advance only once per waveform model and parameter space. However, since for typical PE analyses we have to compute more waveforms than what is needed to construct the ROQ and a basis can be used to perform multiple PEs, this start-up cost quickly pays off.

This paper presents several algorithms for ROQ construction, which offer some key advantages over existing methods. They are specifically designed to tackle the challenges of speed in the basis construction and accuracy in GW likelihood evaluation while maximizing the ROQ speedup. As we will see, these algorithms have the ability to handle complex waveform models in parameter ranges that were intractable with existing procedures.

The paper is organized as follows. In Sec. V E, we introduce the basic theoretical framework, including a discussion on GW inference as well as on the basics of ROQ. In Sec. V F, we describe the ROQ algorithms we introduce in depth, going through the construction of the reduced order basis, the choice of empirical interpolation model and how to construct a ROQ with a set tolerance for a given parameter space. In Sec. V G we present several bases created for two phenomenological waveform models, `IMRPhenomPv2` [111] and `IMRPhenomXPHM` [112], and test their speed and accuracy. We further test the ROQ by performing parameter estimation analyses on three confirmed GW events. In Sec. V H we finally conclude. We relegate some of the more convoluted numerical methods used by our algorithms to the Appendices.

The methods introduced in this paper have been implemented in a `python` code named `EigROQ`, which is publicly available at <https://github.com/gmorras/EigROQ>.

E. Theoretical framework

In this section we will briefly describe the basic theoretical framework to contextualize the rest of the paper. In Sec. V E 1 we give a very brief overview on the basics of GW parameter estimation while on Sec. V E 2 we summarize the basics of the ROQ rule. For more details, we refer the reader to Refs. [95, 104].

1. A primer on gravitational wave inference

GW inference refers to the modern scientific discipline taking care among other things, of computing the posterior probability distribution of the GW model parameters $\vec{\theta}$ that best fit the data, using Bayes Theorem

$$p(\vec{\theta}|d) = \frac{\mathcal{L}(d|\vec{\theta})\pi(\vec{\theta})}{\mathcal{Z}}. \quad (94)$$

In this equation, there are several objects that enter the calculation. The first, $\pi(\vec{\theta})$ refers to the prior employed, from the nature of the event, which throughout this paper will always be a CBC to the distributions describing the parameters of the binary. Next, the likelihood function $\mathcal{L}(d|\vec{\theta})$ of the data given the parameters $\vec{\theta}$ and the evidence \mathcal{Z} representing the probability of the data given the model.

The likelihood is the most computationally expensive part of estimating the posterior. Given a CBC signal without eccentricity, there are 15 different parameters to fit that enter the likelihood computation. The typical gravitational-wave astronomy likelihood is based on the hypothesis that only Gaussian noise is present in the detector and deviations from it are the result of a GW signal. In such case, the likelihood can up to a normalization constant be expressed as [80]

$$\begin{aligned} \log \mathcal{L}(d|\vec{\theta}) &= -\frac{1}{2}(d - h(\vec{\theta}), d - h(\vec{\theta})) \\ &= -\frac{1}{2}(d, d) + (d, h(\vec{\theta})) - \frac{1}{2}(h(\vec{\theta}), h(\vec{\theta})), \end{aligned} \quad (95)$$

where $h(\vec{\theta})$ represents, in this specific case, the CBC waveform with parameters $\vec{\theta}$ used to fit the data d . The overlap integral (\cdot, \cdot) is defined as

$$(d, h(\vec{\theta})) = 4\Delta f \mathcal{R} \sum_{j=1}^L \frac{\tilde{d}^*(f_j) \tilde{h}(f_j; \vec{\theta})}{S(f_j)}, \quad (96)$$

, where $S_n(f)$ is the detector's noise power spectral density (PSD) and $\tilde{a}(f)$, denotes the Fourier transform of $a(t)$. Since the data of GW detectors are discretely sampled, we will have discrete Fourier transforms having a frequency spacing $\Delta f = 1/T$, with T being the observation time. For a frequency window $(f_{\text{high}} - f_{\text{low}})$ there will be $L = \text{int}[(f_{\text{high}} - f_{\text{low}})T]$ terms in the sum of Eq. (96).¹⁶ Repeatedly computing the overlap integrals in Eq. (95) is the bottleneck in gravitational waves inference, and the main part we aim to speed up in this paper.

2. Basics of Reduced Order Quadratures for Gravitational Wave inference

The parameters $\vec{\theta}$ of the GW signal $h(\vec{\theta})$ we are fitting to the data (Eq. (95)) can be split on intrinsic and extrinsic parameters. The extrinsic parameters are common to all transient GW sources and they are the sky location, usually measured with right ascension α and declination δ , the polarization ψ , luminosity distance d_L and a reference time of arrival of the signal t_c .¹⁷ The intrinsic parameters are related to the source of the GW and are generically referred to as $\vec{\lambda}$. For a quasi-circular CBC they are comprised of the 2 component masses m_1 and m_2 , 3 components per BH spin vector \vec{s}_i , the inclination angle ι and the coalescence phase ϕ_c . For CBCs with at least one neutron star (NS) $\vec{\lambda}$ can also contain a tidal deformability parameter Λ per NS in the binary [113], as well as any other matter effect information included in the model. If we break the assumption of quasi-circular orbits, the eccentricity e would also have to be taken into account in the intrinsic parameters $\vec{\lambda}$ [114].

We assume that the signal $h(t, \vec{\theta})$ is short enough to ignore the dependence of the detector antenna patterns $F_{+, \times}$ with time and the time-varying Doppler shift due to motion of the detector with respect to the solar system barycenter [115]. In practice, the signal will have to last less than a few hours, to be able to ignore the effects of Earth's rotation. Then, in the frequency domain, the GW signal can be written as:

¹⁶ Here $\text{int}[x]$ refers to taking the integer part of x .

¹⁷ We use t_c because, for the CBC case, the reference time of arrival for the signal is usually given by the coalescence time at the geocenter.

$$\begin{aligned} \tilde{h}(f, \vec{\theta}) &= e^{-i2\pi f t_c} \frac{1}{d_L} \left(F_+(\alpha, \delta, \psi) \tilde{h}_+(f, \vec{\lambda}) \right. \\ &\quad \left. + F_\times(\alpha, \delta, \psi) \tilde{h}_\times(f, \vec{\lambda}) \right) \left(\right. \\ &\equiv e^{-i2\pi f t_c} \tilde{h}(f, \vec{\Lambda}) \end{aligned} \quad (97)$$

The main idea of the ROQ is to represent the GW waveform model $\tilde{h}(f_i; \vec{\theta})$ and its modulus squared $|\tilde{h}(f_i; \vec{\theta})|^2$ in terms of an empirical interpolant each, which is described in more detail in Sec. V F. For now, we assume that they can be approximated to arbitrary precision as:

$$\tilde{h}(f_i; \vec{\Lambda}) \approx \sum_{j=1}^{N_L} B_j(f_i) \tilde{h}(F_j; \vec{\Lambda}) \quad (98a)$$

$$|\tilde{h}(f_i; \vec{\Lambda})|^2 \approx \sum_{k=1}^{N_Q} C_k(f_i) |\tilde{h}(\mathcal{F}_k; \vec{\Lambda})|^2, \quad (98b)$$

where the main focus of this paper is to find the optimal values of the interpolation nodes $\{F_j\}_{j=1}^{N_L}$ and $\{\mathcal{F}_k\}_{k=1}^{N_Q}$ and of the ‘‘bases’’ $B_j(f_i)$ and $C_k(f_i)$ such that we minimize the required number of elements ($N_L + N_Q$) entering Eq. (98) while respecting a given specified precision.

If we input Eq. (97) into Eq. (95) and use the approximation for the GW waveform $\tilde{h}(f_i; \vec{\theta})$ and its modulus squared $|\tilde{h}(f_i; \vec{\theta})|^2$ of Eq. (98), we can represent the likelihood as

$$\log \mathcal{L}(d|\vec{\theta}) \approx -\frac{1}{2}(d, d) + (d, h(\vec{\theta}))_{\text{ROQ}} - \frac{1}{2}(h(\vec{\theta}), h(\vec{\theta}))_{\text{ROQ}}, \quad (99)$$

where the term $-\frac{1}{2}(d, d) \equiv \log \mathcal{L}_{\text{noise}}$ is a constant that depends only on the data and cancels with the evidence \mathcal{Z} when we compute the posterior probability distribution using Bayes theorem (Eq. (94)). In Eq. (99) we have also implicitly defined the quantities:

$$(d, h(\vec{\theta}))_{\text{ROQ}} \equiv \mathcal{R} \sum_{j=1}^{N_L} w_j(t_c) \tilde{h}(F_j; \vec{\Lambda}) \quad (100a)$$

$$(h(\vec{\theta}), h(\vec{\theta}))_{\text{ROQ}} \equiv \sum_{k=1}^{N_Q} \psi_k |\tilde{h}(\mathcal{F}_k; \vec{\Lambda})|^2, \quad (100b)$$

which approximates the corresponding overlap integrals appearing in the Likelihood calculation of Eq (95). In Eq. (100) we have introduced the linear and quadratic ROQ weights, $w_j(t_c)$ and ψ_k , defined as:

$$w_j(t_c) \equiv 4\Delta f \sum_{i=1}^L \frac{\tilde{d}^*(f_i) B_j(f_i)}{S(f_i)} e^{-i2\pi f_i t_c} \quad (101a)$$

$$\psi_k \equiv 4\Delta f \sum_{i=1}^L \frac{C_k(f_i)}{S(f_i)}. \quad (101b)$$

Before starting PE analysis on an event, the weights have to be computed for the observed data strain $\tilde{d}(f)$ and the corresponding PSD ($S(f)$). Since the linear weights are smooth functions of time, they are usually evaluated in a discrete set of times N_t and are interpolated for the PE analysis [104]. The spacing between time samples is usually of the order of the expected resolution in t_c , which for CBC signals can be as small as 0.1ms, and for the typical t_c prior, which is uniform in ± 0.1 s around trigger time, this equates to $N_t \sim O(10^3)$. Therefore, at the beginning of the analysis, we have to perform $N_t N_L + N_Q$ full overlaps, as prescribed in Eq. (101), and the startup cost of the ROQ is $O((N_t N_L + N_Q)L)$.

Once the weights have been initialized, computing the ROQ likelihood will only require $N_L + N_Q$ terms to estimate the overlap integrals (Eq. (100)), compared to the L terms in the full overlap integrals. We can therefore expect a speed-up in the likelihood computation of $O(L/(N_L + N_Q))$ when using the ROQ rule. In GW astronomy, typical CBC PE analyses require $O(10^8 - 10^9)$ likelihood evaluations, which dominate the computational cost required to sample the posterior of Eq. (94). In most applications the startup cost of the ROQ is negligible compared to the sampling time and the ROQ will greatly speed up the whole analysis. The likelihood speedup is further explored in Sec. V G 3.

The biggest overhead when using the ROQ rule is in constructing the ROQ basis (Eq. (98)), since to explore typical CBC parameter spaces we need $O(10^6 - 10^7)$ random waveforms. With the methods outlined in this paper, we also aim to reduce the computational time of the basis generation, allowing us to handle complex waveform models in parameter ranges that were intractable with existing procedures. In practice, for the CBC case, we train the ROQ on the h_+ polarization, varying only the values of the intrinsic parameters $\vec{\lambda}$, defined in Eq. (97). The same ROQ basis is valid for both polarizations since they can be jointly decomposed in spherical harmonics of spin weight -2 , ${}_{-2}Y_{lm}$ as [96]

$$h_+ - ih_\times = \sum_{l=2}^{\infty} \sum_{m=-l}^l {}_{-2}Y_{lm}(\iota, \phi_c) h_{lm} \quad (102)$$

where the inclination ι and coalescence phase ϕ_c are also being sampled.

F. Efficient algorithm for ROQ computation

1. Reduced Order Basis

We generate N templates from the waveform model we are trying to approximate:

$$\{h_A(x), A = 1, \dots, N\}, \quad (103)$$

where, in GW astronomy, x can be either frequency f or time t . We can define the matrix of inner products between templates as

$$M_{AB} = \langle h_A, h_B \rangle. \quad (104)$$

In this context, the inner product is usually defined as:

$$\langle h_A, h_B \rangle = \int_{\text{low}}^{f_{\text{high}}} \tilde{h}_A^*(f) \tilde{h}_B(f) df \quad (105)$$

although we could also use a reference PSD $S_n(f)$ to give different weights at different frequencies to the integrand, as in Eq. (96). Since M_{AB} is a matrix of inner products, it is hermitian and positive semi-definite, and therefore can always be diagonalized as

$$M_{AB} = \sum_{C=1}^N E_{AC} \lambda_C E_{BC}^*, \quad (106)$$

where $\lambda_C \geq 0$ are the eigenvalues and E_{AB} is a unitary matrix whose columns are the orthonormal eigenvectors

$$\sum_{C=1}^N E_{CA}^* E_{CB} = \delta_{AB}. \quad (107)$$

In the waveform space we can then define the eigenvectors with $\lambda_A \neq 0$ as:

$$e_A(x) = \frac{1}{\sqrt{\lambda_A}} \sum_{C=1}^N h_C(x) E_{CA}. \quad (108)$$

It can be proven that these are an orthonormal set of vectors under $\langle \cdot, \cdot \rangle$. That is:

$$\begin{aligned} \langle e_A, e_B \rangle &= \left\langle \frac{1}{\sqrt{\lambda_A}} \sum_{C=1}^N h_C(x) E_{CA}, \frac{1}{\sqrt{\lambda_B}} \sum_{D=1}^N h_D(x) E_{DB} \right\rangle \left(\begin{array}{c} \\ \\ \\ \end{array} \right) \\ &= \frac{1}{\sqrt{\lambda_A \lambda_B}} \sum_{C=1}^N \sum_{D=1}^N E_{CA}^* E_{DB} \underbrace{\langle h_C, h_D \rangle}_{M_{CD}} \left(\begin{array}{c} \\ \\ \\ \end{array} \right) \\ &= \frac{1}{\sqrt{\lambda_A \lambda_B}} \sum_{C=1}^N E_{CA}^* \sum_{D=1}^N M_{CD} E_{DB} \left(\begin{array}{c} \\ \\ \\ \end{array} \right) \end{aligned} \quad (109)$$

$$\begin{aligned} &= \sqrt{\underbrace{\lambda_B}_{\lambda_A} \sum_{C=1}^N E_{CA}^* E_{CB}} = \delta_{AB}. \left(\begin{array}{c} \\ \\ \\ \end{array} \right) \end{aligned} \quad (110)$$

We can also prove that the inner product between one of the waveforms used to compute M_{AB} and a given eigenvector will be given by:

$$\begin{aligned} \langle h_A, e_B \rangle &= \left\langle h_A, \frac{1}{\sqrt{\lambda_B}} \sum_{C=1}^N h_C(x) E_{CB} \right\rangle \left(\begin{array}{c} \\ \\ \\ \end{array} \right) \\ &= \frac{1}{\sqrt{\lambda_B}} \sum_{C=1}^N \underbrace{\langle h_A, h_C \rangle}_{M_{AC}} E_{CD} \left(\begin{array}{c} \\ \\ \\ \end{array} \right) \\ &= \frac{1}{\sqrt{\lambda_B}} \sum_{C=1}^N M_{AC} E_{CD} = \sqrt{\lambda_B} E_{AB}. \left(\begin{array}{c} \\ \\ \\ \end{array} \right) \end{aligned} \quad (111)$$

We can define our reduced order basis (ROB) as a subset of $n < N$ elements of $\{e_A\}_{A=1}^N$, which we will learn how to optimally select later. To represent the waveform h_A in terms of this ROB $\{e_a\}_{a=1}^n$, we project h_A using the orthonormality property of the ROB:

$$h_A^{\text{ROB}}(x) = \sum_{b=1}^n \langle e_b, h_A \rangle e_b(x) = \sum_{b=1}^n \sqrt{\lambda_b} E_{Ab}^* e_b(x). \quad (112)$$

We can compute the representation error of projecting h_A as:

$$\begin{aligned} \sigma_{\text{ROB}, A} &= \|h_A - h_A^{\text{ROB}}\|^2 = \langle h_A - h_A^{\text{ROB}}, h_A - h_A^{\text{ROB}} \rangle \\ &= \left\langle h_A - \sum_{b=1}^n \langle e_b, h_A \rangle e_b, h_A - \sum_{c=1}^n \langle e_c, h_A \rangle e_c \right\rangle \left(\begin{array}{c} \\ \\ \\ \end{array} \right) \\ &= \langle h_A, h_A \rangle - \sum_{b=1}^n |\langle e_b, h_A \rangle|^2 = \langle h_A, h_A \rangle - \sum_{b=1}^n \lambda_b |E_{Ab}|^2. \end{aligned} \quad (113)$$

Ideally, to construct a ROB we would take a very large number of templates $\{h_A\}_{A=1}^N$, that capture most of the variability of the waveform in the parameter space of interest, compute the matrix M_{AB} as in Eq. (104), diagonalize it and, to construct our ROB, pick the minimum number of eigenvectors $\{e_a\}_{a=1}^n$ such that the ROB error of Eq. (113) is smaller than a specified tolerance. Unfortunately, this cannot be done in practice, since the number of random templates needed to fully span the typical parameter spaces for GW applications is of order $O(10^7)$. Using the fact that M_{AB} is hermitian, we need $N(N-1)/2$ complex numbers to store the off-diagonal elements, and N real numbers for the diagonal elements. Assuming that each real number is stored with n_B Bytes, the memory required to store

M_{AB} is¹⁸:

$$\text{Memory}(M_{AB}) = N^2 n_B = 80 \text{ GB} \left(\frac{N}{10^5} \right)^2 \left(\frac{n_B}{8 \text{ B}} \right). \quad (114)$$

Therefore, in current computers, examining more than a few tens of thousands of waveforms at a time is unfeasible, and we will not be able to analyze the entire parameter space at once. Motivated by this issue, we have developed a multi-step approach summarised in Algorithm 1. We construct a first ROB for a set tolerance with random waveforms. Then, we calculate its orthogonal space and obtain the corresponding ROB which we add to the original ROB. We repeat this process iteratively, reducing the tolerance at every step. The equivalent to the matrix M_{AB} of Eq. (104) for the orthogonal space to the basis $\{e_a\}_{a=1}^n$ is:

$$\begin{aligned} M_{AB}^{\text{ROB}} &= h_A - h_A^{\text{ROB}}, h_B - h_B^{\text{ROB}} \\ &= \left\langle h_A - \sum_{c=1}^n \langle e_c, h_A \rangle e_c, h_B - \sum_{d=1}^n \langle e_d, h_B \rangle e_d \right\rangle \left(\begin{array}{c} \\ \end{array} \right) \\ &= \langle h_A, h_B \rangle - \sum_{c=1}^n \langle h_A, e_c \rangle \langle e_c, h_B \rangle. \end{aligned} \quad (115)$$

Algorithm 1 Construction of reduced order basis

- 1: **Input:** Maximum number of waveforms selected N , tolerances of each step $[\sigma_0, \dots, \sigma_s]$, maximum number of waveforms computed per step $[N_{\text{lim},1}, \dots, N_{\text{lim},s}]$
 - 2: Generate N waveforms $\{h_A\}_{A=1}^N$
 - 3: Compute the matrix $M_{AB} = \langle h_A, h_B \rangle$
 - 4: Diagonalize M_{AB} to obtain eigenvalues λ_A and eigenvectors E_{AB}
 - 5: Input $\{\sigma_0, \{h_A\}_{A=1}^N, \lambda_A, E_{AB}\}$ in Algorithm 2 to obtain initial ROQ basis $\{e_i\}_{i=1}^{n_0}$
 - 6: **for** $j = 1 \rightarrow s$ **do**
 - 7: **repeat**
 - 8: Generate $N_{\text{lim},j}$ waveforms $\{h_A\}_{A=1}^{N_{\text{lim},j}}$ and compute their ROB error $\sigma_{\text{ROB},A}$
 - 9: Select the N waveforms $\{h_A\}_{A=1}^N$ with largest σ_{ROB}
 - 10: Save the minimum value of σ_{ROB} for the selected waveforms: $\sigma_{\text{ROB,min}}$
 - 11: $M_{AB}^{\text{ROB}} = \langle h_A, h_B \rangle - \sum_{c=1}^{n_{j-1}} \langle h_A, e_c \rangle \langle e_c, h_B \rangle$
 - 12: Diagonalize M_{AB}^{ROB} and obtain eigenvalues λ_A and eigenvectors E_{AB}
 - 13: Input $\{\sigma_j, \{h_A - h_A^{\text{ROB}}\}_{A=1}^N, \lambda_A, E_{AB}\}$ in Algorithm 2 to obtain next ROQ basis elements $\{e_i\}_{i=n_{j-1}+1}^{n_j}$
 - 14: **until** $\sigma_{\text{ROB,min}} < \sigma_s$
 - 15: **end for**
 - 16: **Output:** ROB $\{e_i\}_{i=1}^n$
-

We observe that the same formulas and reasoning of Eqs. (104-114) apply to the space orthogonal to the ROB if we make the identification $h_A \rightarrow h_A - h_A^{\text{ROB}}$. To find the minimum number of elements that have to be added to the ROB to reduce the error below the set tolerance σ , we use Algorithm 2, where we iteratively subtract the contribution of the eigenvalue that produces the largest drop in any $\sigma_{\text{ROB},A}$, according to Eq. (113), until $\sigma_{\text{ROB},A} < \sigma$ for all A .

¹⁸ 1GB = 10^9 Bytes = $8 \cdot 10^9$ bits

Algorithm 2 Selection of Eigenvectors

1: **Input:** Tolerance σ , waveforms $\{h_A\}_{A=1}^N$, eigenvalues λ_A and eigenvectors E_{AB} of the matrix $M_{AB} = \langle h_A, h_B \rangle$

2: Initialize σ_A : $\{\sigma_A = \langle h_A, h_A \rangle\}_{A=1}^N$

3: Compute the maximum contribution of each eigenvector
 $\{\delta\sigma_{A,\max} = \lambda_A \max_B |E_{BA}|^2\}_{A=1}^N$

4: Find order of $\delta\sigma_{A,\max}$: $\{B_n\}_{n=1}^N = \text{argsort}(\delta\sigma_{A,\max})$

5: $n = N$

6: **repeat**

7: Compute current error $\{\sigma_A \leftarrow \sigma_A - \lambda_{B_n} |E_{AB_n}|^2\}_{A=1}^N$

8: $n \leftarrow n - 1$

9: **until** $\sigma_A < \sigma \forall A = 1, \dots, N$

10: **Output:** Eigenvectors in waveform domain

$$\left\{ e_k(x) = \frac{1}{\sqrt{\lambda_{B_k}}} \sum_{A=1}^N h_A(x) E_{AB_k} \right\}_{k=n}^N$$

The process of diagonalizing the matrix M_{AB} of Eq. (104) and finding the eigenvalues in the waveform domain using Eq. (108) is equivalent to performing Singular Value Decomposition (SVD) on a set of waveforms $\{h_A\}_{A=1}^N$, which has been previously used in the literature for the Reduced Order Modeling (ROM) of GW waveforms (See Refs. [116, 117]). However, we follow the procedure outlined in this paper since it has a few numerical advantages. Namely, if we have waveforms with a number of sampling points M , storing them will require $2MNn_B$ bytes, which in the usual case that $M \gg N$, will be much larger than the memory needed to store M_{AB} (Eq. (114)) and we will be even more limited in the number of waveforms we can analyze at once. Moreover, if we are studying the ROB of the space orthogonal to $\{e_a\}_{a=1}^n$, our algorithm is equivalent to computing the SVD of the orthogonal part of the waveforms $\{h_A - h_A^{\text{ROB}}\}_{A=1}^N$. Finding this orthogonal part is, in general, a computationally expensive process that can be avoided if M_{AB}^{ROB} is obtained using Eq. (115). Since we are going to select $n_{\text{new}} \ll N$ eigenvectors of M_{AB}^{ROB} , we can just compute the orthogonal projection of their corresponding eigenvectors in the waveform domain at the end of the algorithm.

2. Empirical Interpolation Model

Writing a given template in the form of Eq. (112) will not save computational cost, since one needs the full waveform $h_A(x)$ to compute the inner product $\langle h_A, e_B \rangle$. To avoid this, we approximate the inner products $\langle h(\vec{\lambda}), e_i \rangle$ by some coefficients $c_i(\vec{\lambda})$ that will in general be functions of the parameters of the waveform $\vec{\lambda}$ (e.g. for a CBC this would be masses, spins, inclination and coalescence phase). The approximate waveform can then be written as:

$$I_n[h](x, \vec{\lambda}) = \sum_{i=1}^n c_i(\vec{\lambda}) e_i(x). \quad (116)$$

We force the approximation to be exact at some interpolation nodes $\{X_j\}_{j=1}^m$

$$I_n[h](X_j, \vec{\lambda}) = h(X_j, \vec{\lambda}) = \sum_{i=1}^n c_i(\vec{\lambda}) e_i(X_j). \quad (117)$$

This is what we define as an interpolant. If we identify the matrix

$$A_{ij} = e_j(X_i), \quad (118)$$

and take the number of interpolation nodes m to be equal to the number of basis elements n , then \hat{A} is a square matrix which we construct by choosing the interpolation nodes $\{X_j\}_{j=1}^n$. Assuming that we construct \hat{A} to be invertible, we can solve Eq. (117) for $c_i(\vec{\lambda})$ in the following way:

$$c_i(\vec{\lambda}) = \sum_{j=1}^n (\hat{A}^{-1})_{ij} h(X_j, \vec{\lambda}). \quad (119)$$

We therefore observe that the value of $c_i(\vec{\lambda})$ will just be a linear combination of the values of the waveform at the different interpolation nodes $\{X_j\}_{j=1}^n$. In practice, the functions $h(x)$ and the ROB elements $\{e_i(x)\}_{i=1}^n$ are discretely sampled in a set of points $\{x_i\}_{i=1}^M$, and we can define the matrix:

$$\hat{V} \equiv [\vec{e}_1, \dots, \vec{e}_n] \in \mathbb{C}^{M \times n}, \quad (120)$$

where $\vec{e}_A = e_A(\vec{x}) \in \mathbb{C}^M$. From Eq. (118), we observe that the matrix \hat{A} can be written in terms of \hat{V} as:

$$\hat{A} = \hat{P}^\dagger \hat{V} \in \mathbb{C}^{n \times n}, \quad (121)$$

where the matrix $\hat{P} \in \mathbb{C}^{M \times n}$ is a projector that selects the rows of \hat{V} corresponding to the interpolation nodes. That is:

$$P_{\alpha j} = \delta_{\alpha \beta_j} \quad (122)$$

with $\{\beta_j\}_{j=1}^n$ the indices of the interpolation nodes (i.e. $x_{\beta_j} = X_j$). In terms of these matrices, the empirical interpolation model (EIM) can be written as:

$$I_n[\vec{h}] = \hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger \vec{h}. \quad (123)$$

which is an interpolant because $\hat{P}^\dagger I_n[\vec{h}] = \hat{P}^\dagger \vec{h}$. In terms of the matrix \hat{V} , the ROB representation of \vec{h} is given by

$$\vec{h}^{\text{ROB}} = \hat{V} \hat{V}^\dagger \vec{h}. \quad (124)$$

Note that even though the basis elements \vec{e}_A are orthonormal, and therefore $\hat{V}^\dagger \hat{V} = \mathbb{1}_{n \times n}$, since the matrices are not square, we have that in general $\hat{V} \hat{V}^\dagger \neq \mathbb{1}_{M \times M}$. From Eqs. (123, 124) we can explicitly see that the EIM acting on a waveform in the ROB space will have no effect. That is:

$$\begin{aligned} I_n[\vec{h}^{\text{ROB}}] &= \hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger (\hat{V} \hat{V}^\dagger \vec{h}) = \hat{V}(\hat{P}^\dagger \hat{V})^{-1} (\hat{P}^\dagger \hat{V}) \hat{V}^\dagger \vec{h} \\ &= \hat{V} \hat{V}^\dagger \vec{h} = \vec{h}^{\text{ROB}} \end{aligned} \quad (125)$$

This can be used to relate the representation error of the EIM with the representation error of the ROB. Computing the modulus of the difference between the exact waveform and its EIM representation we obtain:

$$\begin{aligned} \sigma_{\text{EIM}}(\vec{h}) &= \|\vec{h} - I_n[\vec{h}]\|_2^2 = \left\| \mathbb{1} - \hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger \right\|_2^2 \|\vec{h}\|_2^2 \\ &= \left\| \mathbb{1} - \hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger \right\|_2^2 \left\| \vec{h} - \vec{h}^{\text{ROB}} \right\|_2^2 \\ &\leq \|\mathbb{1} - \hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger\|_2^2 \underbrace{\|\vec{h} - \vec{h}^{\text{ROB}}\|_2^2}_{\sigma_{\text{ROB}}(\vec{h})}, \end{aligned} \quad (126)$$

where $\|\cdot\|_2$ denotes the matrix 2-norm, which is given by:

$$\|\hat{M}\|_2 = \max_{\vec{x} \neq 0} \frac{\|\hat{M}\vec{x}\|}{\|\vec{x}\|} = \sqrt{\lambda_{\max}(\hat{M}^\dagger \hat{M})} = \sqrt{\lambda_{\max}(\hat{M} \hat{M}^\dagger)}, \quad (127)$$

where $\|\vec{x}\|$ is the usual vector norm and $\lambda_{\max}(\hat{M}^\dagger \hat{M})$ denotes the maximum eigenvalue of $\hat{M}^\dagger \hat{M}$. Since $\hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger$ is idempotent, that is $(\hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger)^2 = \hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger$, and it is different from $\mathbb{0}$ or the identity $\mathbb{1}$, it follows that [118]:

$$\|\mathbb{1} - \hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger\|_2 = \|\hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger\|_2. \quad (128)$$

Furthermore, since $\hat{V}^\dagger \hat{V} = \mathbb{1}_{n \times n}$ and $\hat{P}^\dagger \hat{P} = \mathbb{1}_{n \times n}$, from the definition in Eq. (127) of the matrix 2-norm, we have that

$$\|\hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger\|_2 = \|\hat{P}^\dagger \hat{V}\|_2^{-1}. \quad (129)$$

Substituting in Eq. (126)

$$\sigma_{\text{EIM}}(\vec{h}) \leq \|(\hat{P}^\dagger \hat{V})^{-1}\|_2^2 \sigma_{\text{ROB}}(\vec{h}) = \|\hat{A}^{-1}\|_2^2 \sigma_{\text{ROB}}(\vec{h}). \quad (130)$$

Using the definition of the matrix 2-norm of Eq. (127), we have that

$$\begin{aligned} \|\hat{A}^{-1}\|_2^2 &= \lambda_{\max} \left((\hat{A}^{-1})^\dagger \hat{A}^{-1} \right) = \lambda_{\max} \left((\hat{A}^\dagger)^{-1} \hat{A}^{-1} \right) \left(\right. \\ &= \lambda_{\max} \left((\hat{A} \hat{A}^\dagger)^{-1} \right) \left. \right) = \frac{1}{\lambda_{\min}(\hat{A} \hat{A}^\dagger)}, \end{aligned} \quad (131)$$

and we can rewrite Eq. (130) as

$$\sigma_{\text{EIM}}(\vec{h}) \leq \frac{\sigma_{\text{ROB}}(\vec{h})}{\lambda_{\min}(\hat{A} \hat{A}^\dagger)}. \quad (132)$$

Therefore, given a maximum error of the ROB, the error of the EIM model is bounded from above by Eq. (132). To make this bound as stringent as possible, we could maximize the smallest eigenvalue of $\hat{A} \hat{A}^\dagger$. Using the definition of \hat{A} from Eq. (121) we can write

$$(\hat{A} \hat{A}^\dagger)_{ij} = \sum_{k=1}^n e_k(X_i) e_k^*(X_j) = \langle \vec{v}_j, \vec{v}_i \rangle, \quad (133)$$

where we have defined the vectors $\{(\vec{v}_i)_k = e_k(X_i) | k = 1, \dots, n\}_{i=1}^M$ as the rows of \hat{V} corresponding to the interpolation nodes X_i . We then observe that $\hat{A} \hat{A}^\dagger$ is the same as the scalar product between the corresponding selected rows of \hat{V} .

If the vectors \vec{v}_i were orthonormal, we would obtain that $(\hat{A} \hat{A}^\dagger)_{ij} = \delta_{ij}$, and therefore $\lambda_{\min}(\hat{A} \hat{A}^\dagger) = 1$ and the EIM would not introduce additional error over the ROB. Selecting n orthonormal rows of \hat{V} is in general not possible, however, we can try to minimize the EIM error by picking rows which are as close to orthogonal as possible using algorithm 3.

Algorithm 3 Selection of interpolation nodes

- 1: **Input:** Evaluated basis $\{\vec{e}_i\}_{i=1}^n$
 - 2: Define row vectors: $\{\vec{v}_\alpha = \{e_i(x_\alpha)\}_{i=1}^n\}_{\alpha=1}^M$
 - 3: Initialize ortonormal base of columns: OB = $\{\vec{w}_i\}_{i=1}^0$
 - 4: Initialize the norm of the orthogonal part of \vec{v}_α to OB: $\{N_\alpha = |\vec{v}_\alpha|^2\}_{\alpha=1}^M$
 - 5: **for** $j = 1 \rightarrow n$ **do**
 - 6: Choose vector with largest N_α : $\beta_j = \text{argmax}(N_\alpha)$
 - 7: Append \vec{v}_{β_j} to OB using Gram-Schmidt
 - 8: Update N_α : $\{N_\alpha \leftarrow N_\alpha - |\langle \vec{w}_j, \vec{v}_\alpha \rangle|\}_{\alpha=1}^M$
 - 9: **end for**
 - 10: **Output:** EIM interpolation nodes $\{\beta_i\}_{i=1}^n$
-

We observe that Algorithm 3 is equivalent to picking the EIM nodes that maximize the determinant of $\hat{A} \hat{A}^\dagger$, since

$$\det(\hat{A} \hat{A}^\dagger) = \det(\hat{A}) \det(\hat{A}^\dagger) = |\det(\hat{A})|^2 = \prod_{j=1}^n |\langle \vec{w}_j, \vec{v}_{\beta_j} \rangle|^2. \quad (134)$$

Algorithm 3 does not directly maximize the minimum eigenvalue of $\hat{A} \hat{A}^\dagger$. However, based on the expression for the determinant of $\hat{A} \hat{A}^\dagger$

$$\det(\hat{A} \hat{A}^\dagger) = \prod_{j=1}^n \lambda_j, \quad (135)$$

to maximize it, the values of the individual eigenvalues have to be large, and thus, the output of the algorithm is near to the minimum of $\|\hat{A}^{-1}\|_2^2$. When compared to the greedy algorithm typically used in the literature (e.g. Refs. [102–105]) to compute the interpolation nodes, we observe a superior performance of algorithm 3, as we will later discuss in relation to figure 16.

If we wanted to create an EIM with a tolerance smaller than σ , from Eq. (132) we could in principle just construct a ROB with a tolerance better than $\lambda_{\min}(\hat{A}\hat{A}^\dagger)\sigma$. However, in real settings, we observe that Eq. (132) is a loose upper bound on the EIM error, and we can obtain an EIM with a tolerance better than σ using fewer basis elements.

Instead of bounding $\sigma_{\text{EIM}}(\vec{h})$ using the inequality of Eq. (126), we can refine this expression by doing:

$$\begin{aligned}\sigma_{\text{EIM}}(\vec{h}) &= \left[\mathbb{1} - \hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger \right] (\vec{h} - \vec{h}^{\text{ROB}})^2 \\ &= \vec{h} - \vec{h}^{\text{ROB}}^2 + \hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger (\vec{h} - \vec{h}^{\text{ROB}})^2 \\ &= \sigma_{\text{ROB}}(\vec{h}) + (\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger (\vec{h} - \vec{h}^{\text{ROB}})^2,\end{aligned}\quad (136)$$

where we have used that $\hat{V}^\dagger \hat{V} = \mathbb{1}$ and that the EIM projects the waveform onto the ROB, and therefore $\langle \hat{V}(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger (\vec{h} - \vec{h}^{\text{ROB}}), \vec{h} - \vec{h}^{\text{ROB}} \rangle = 0$. From Eq. (136) we have that the EIM error is always larger than or equal to the ROB error. We also observe that for the bound of Eq. (132) to be saturated we need $\hat{P}^\dagger(\vec{h} - \vec{h}^{\text{ROB}})$ to be the eigenvector of $\hat{A}^\dagger \hat{A}$ with the maximum eigenvalue, which is extremely unlikely in general. To explore this we assume that $\vec{h} - \vec{h}^{\text{ROB}} \equiv \delta \vec{h}$ is a random variable, such that:

$$\mathbb{E}[\delta h_\alpha^* \delta h_\beta] = c_\alpha \delta_{\alpha\beta}, \quad (137)$$

where $\mathbb{E}[\cdot]$ denotes the expected value (i.e. the average over random waveform realizations). Using Eq. (137), we compute the expected value of σ_{EIM} as

$$\begin{aligned}\mathbb{E}[\sigma_{\text{EIM}}] &= \sum_{\alpha=1}^M \mathbb{E}[\delta h_\alpha^* \delta h_\alpha] \\ &+ \sum_{k=1}^n \sum_{q=1}^n \sum_{l=1}^n \left((A^{-1})_{lk}^* (A^{-1})_{lq} \mathbb{E}[\delta h_{\beta_k}^* \delta h_{\beta_q}] \right) \\ &= \sum_{\alpha=1}^M c_\alpha + \sum_{k=1}^n \sum_{l=1}^n c_{\beta_k} |(A^{-1})_{lk}|^2 \\ &= \sum_{\alpha=1}^M c_\alpha + \hat{A}^{-1} \frac{2}{F},\end{aligned}\quad (138)$$

where \hat{A} is the matrix

$$\tilde{A}_{kl} = \frac{1}{\sqrt{c_{\beta_k}}} A_{kl} = \frac{1}{\sqrt{c_{\beta_k}}} e_l(x_{\beta_k}). \quad (139)$$

Such that $(\tilde{A}^{-1})_{lk} = \sqrt{c_{\beta_k}} (A^{-1})_{lk}$ and $\|\cdot\|_F$ is the Frobenius norm, defined as:

$$\|\hat{M}\|_F = \sqrt{\sum_{k=1}^n \sum_{l=1}^n |M_{kl}|^2} = \sqrt{\text{Tr}\{\hat{M}^\dagger \hat{M}\}} \quad (140)$$

Therefore, to optimize the EIM such that the expected value of σ_{EIM} is minimum, we want to minimize the value of the Frobenius norm of \hat{A}^{-1} . Using the properties of the trace we can rewrite it as:

$$\hat{\hat{A}}^{-1} \Big|_F = \sqrt{\sum_{k=1}^n \frac{1}{\lambda_k(\hat{\hat{A}}^\dagger \hat{\hat{A}})}} \quad (141)$$

To minimize the Frobenius norm of $\hat{\hat{A}}$ we can start from the EIM given by Algorithm. 3 and allow the interpolation nodes to “walk” in the direction of diminishing $\hat{\hat{A}}^{-1} \Big|_F$, as outlined in Algorithm 4.

Algorithm 4 Selection of interpolation nodes to minimize target function of the EIM $F(\cdot)$

```

1: Input: Maximum number of rounds  $N_{\text{rounds}}$ , initial interpolation nodes  $\vec{\beta}$ , function to be minimized  $F(\vec{\beta})$ .
2: for  $j = 1 \rightarrow N_{\text{rounds}}$  do
3:   for  $k = 1 \rightarrow n$  do
4:     for  $\delta\beta$  in  $[-1, 1]$  do
5:       Copy interpolation nodes:  $\vec{\beta}' = \vec{\beta}$ 
6:       repeat
7:         Test new EIM:  $\beta'_k \leftarrow \beta'_k + \delta\beta$ 
8:         if  $F(\vec{\beta}') \leq F(\vec{\beta})$  then
9:           Update reference EIM:  $\vec{\beta} \leftarrow \vec{\beta}'$ 
10:          end if
11:          until  $F(\vec{\beta}') > F(\vec{\beta})$ 
12:        end for
13:      end for
14:      if  $\{\beta_i\}_{i=1}^n$  didn't change this iteration then
15:        break for loop
16:      end if
17:    end for
18: Output: EIM interpolation nodes  $\{\beta_i\}_{i=1}^n$ 

```

The time complexity of Algorithm. 4 is $O(N_{\text{rounds}}nN_F)$, where N_F denotes the number of operations required to compute $F(\vec{\beta})$. Given that our target function is $F(\vec{\beta}) = \hat{\hat{A}}^{-1} \Big|_F$, one could naively expect that, based on the size $n \times n$ of the matrix $\hat{\hat{A}}$, directly inverting it would take $O(n^3)$ operations, and therefore the time complexity of Algorithm 4 would be $O(N_{\text{rounds}}n^4)$. This can be computationally very expensive even if $n \ll M$. However, updating the value of $\|\hat{\hat{A}}^{-1}\|_F$ when only one row of the matrix changes, can be done in $O(n^2)$ by following the procedure of Appendix VI, and we can implement Algorithm 4 with target function $F(\vec{\beta}) = \hat{\hat{A}}^{-1} \Big|_F$ in a way that takes $O(N_{\text{rounds}}n^3)$ operations.

Even though Algorithm 4 is considerably better than the greedy algorithms used in the literature, as we will later discuss in relation to figure 16, it can still be improved by training the EIM directly on the waveform data. For this purpose, we assume that we have an initial ROB $\{\vec{e}_i\}_{i=1}^n$ with a corresponding EIM that can be computed with e.g. algorithm 4. We want to update this EIM to better fit a training set of waveforms $\{h_A\}_{A=1}^N$. We first generate a ROB for the part of the training set orthogonal to the initial ROB $(\vec{h} - \vec{h}^{\text{ROB}})$, which can be done by diagonalizing the matrix of Eq. (115). Analogously to Eq. (112) we can write:

$$\vec{h}_A - \vec{h}_A^{\text{ROB}} = \sum_{B=1}^N \sqrt{\lambda_B} E_{AB}^* \vec{u}_B, \quad (142)$$

where λ_B and E_{AB} are the eigenvalues and eigenvectors of the matrix M_{AB}^{ROB} defined in Eq. (115) and \vec{u}_B represent the eigenvectors in the waveform domain. Substituting Eq. (142) in the expression for σ_{EIM} derived in Eq. (136), we obtain

$$\begin{aligned}
\sigma_{\text{EIM},A} &= \vec{h}_A - I_n[\vec{h}_A]^2 \\
&= \sum_{B=1}^N \sqrt{\int_B E_{AB}^* \vec{u}_B}^2 + \sum_{B=1}^N \sqrt{\int_B E_{AB}^* (\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger \vec{u}_B}^2 \\
&= \sum_{B=1}^N \int_B |E_{AB}|^2 + \sum_{B=1}^N \sum_{C=1}^N \sqrt{\int_B \lambda_C} E_{AB}^* E_{AC} \langle \vec{w}_C, \vec{w}_B \rangle,
\end{aligned} \tag{143}$$

where we have defined

$$\vec{w}_B = (\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger \vec{u}_B \tag{144}$$

From Eq. (143), we can compute the sum of all the EIM errors of the waveforms in the training set. That is:

$$\begin{aligned}
\sigma_{\text{EIM}}^{\text{tot}} &= \sum_{A=1}^N \sigma_{\text{EIM},A} = \sum_{B=1}^N \int_B (1 + \langle \vec{w}_B, \vec{w}_B \rangle) \\
&= \sum_{B=1}^N \int_B \left(1 + (\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger \vec{u}_B \right)^2 \\
&\approx \sum_{B=1}^{n_\lambda} \int_B \left(1 + (\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger \vec{u}_B \right)^2
\end{aligned} \tag{145}$$

Where we have used that E_{AB} is unitary and that the matrix M_{AB}^{ROB} will usually have a small number of large eigenvalues, with the rest of the eigenvalues close to 0. Therefore, we can truncate the sum to be made only over the largest n_λ eigenvalues and obtain a very good approximation of $\sigma_{\text{EIM}}^{\text{tot}}$.

Algorithm 5 Selection of interpolation nodes trained on a set of waveforms $\{h_A\}_{A=1}^N$

- 1: **Input:** Evaluated basis $\{\vec{e}_i\}_{i=1}^n$, maximum number of rounds N_{rounds} , n_λ eigenvalues λ_B and eigenvectors in waveform domain \vec{u}_B of the matrix $M_{AB}^{\text{ROB}} = \langle \vec{h}_A - \vec{h}_A^{\text{ROB}}, \vec{h}_B - \vec{h}_B^{\text{ROB}} \rangle$.
 - 2: Compute weights: $c_\alpha = \sum_{B=1}^{n_\lambda} \lambda_B |u_{B,\alpha}|^2$
 - 3: Compute weighted basis: $\left\{ \{w_i(x_\alpha)\}_{\alpha=1}^M \right\}_{i=1}^n = \left\{ \left\{ \frac{1}{\sqrt{c_\alpha}} e_i(x_\alpha) \right\}_{\alpha=1}^M \right\}_{i=1}^n$
 - 4: Get initial EIM $\vec{\beta}$ inputting $\{\vec{w}_i\}_{i=1}^n$ in Algorithm 3
 - 5: Update $\vec{\beta}$ using Algorithm 4 with maximum rounds N_{rounds} and target function $F(\vec{\beta}) = \|\hat{A}^{-1}\|_F$, where $\tilde{A}_{ij} = w_j(x_{\beta_i})$
 - 6: Update $\vec{\beta}$ again with Algorithm 4 with maximum rounds N_{rounds} and target function $F(\vec{\beta}) = \sigma_{\text{EIM}}^{\text{tot}}$, where $\sigma_{\text{EIM}}^{\text{tot}} = \sum_{B=1}^{n_\lambda} \lambda_B \left(1 + \sum_{i=1}^n \sum_{j=1}^n (\hat{A}^{-1})_{ij} u_{B,\beta_j} \right)^2$
and $A_{ij} = e_j(x_{\beta_i})$
 - 7: **Output:** EIM interpolation nodes $\{\beta_i\}_{i=1}^n$
-

To minimize the value of $\sigma_{\text{EIM}}^{\text{tot}}$, we follow Algorithm 5, in which we start with an EIM and perform walks around the initial solution in the direction of diminishing $\sigma_{\text{EIM}}^{\text{tot}}$. For the initial solution, we will use the EIM generated by Algorithm 4 with target function $F(\vec{\beta}) = \|\hat{A}^{-1}\|_F$. Since we want to fit $\{h_A\}_{A=1}^N$, following Eq. (137), the weights c_α of Eq. (139) are

$$\begin{aligned}
c_\alpha &= \mathbb{E} [\delta h_\alpha^* \delta h_\alpha] = \frac{1}{N} \sum_{A=1}^N |h_{A,\alpha} - h_{A,\alpha}^{\text{ROB}}|^2 \\
&= \frac{1}{N} \sum_{A=1}^N \sum_{B=1}^N \sum_{C=1}^N \sqrt{\lambda_B \lambda_C} E_{AB}^* E_{AC} E_{CB} u_{C,\alpha}^* u_{B,\alpha} \\
&= \frac{1}{N} \sum_{B=1}^N \lambda_B |u_{B,\alpha}|^2 \approx \frac{1}{N} \sum_{B=1}^{n_\lambda} \lambda_B |u_{B,\alpha}|^2,
\end{aligned} \tag{146}$$

where we have once again used that E_{AB} is unitary and that the sum can be approximated by taking only the largest n_λ eigenvalues. In algorithm 4, using $\sigma_{\text{EIM}}^{\text{tot}}$ as target function, the value of $\sigma_{\text{EIM}}^{\text{tot}}$ can be efficiently updated with $O(nn_\lambda)$ operations, as described in appendix VI. Therefore, the algorithm 4 to walk around an initial solution minimizing $\sigma_{\text{EIM}}^{\text{tot}}$ will require $O(N_{\text{rounds}} n^2 n_\lambda)$ operations.

In Figure 16 we show for the 256s `IMRPhenomPv2` ROB listed in Table III a comparison between algorithms 3 4 5 proposed in this paper, the usual greedy algorithm used in the literature and the lower bound imposed by the ROB error. We show only the analysis for the 256s `IMRPhenomPv2` basis of Table III, but we find similar results for all the other cases in Tables III IV. In the upper panel of Figure 16 we show the fraction of points with an EIM error larger than a tolerance σ as a function of σ . Comparing the methods we observe that the *Training* one (algorithm 5) outperforms the others, which is expected since it has been trained on the waveform data to reduce the EIM error. The worst performer is the *Greedy* method since it induces the largest EIM error in all cases tested. We also observe that the *Frobenius* method, which uses algorithm 4 to minimize $\|\hat{A}^{-1}\|_F$ induces the smallest EIM error among the algorithms that do not train on waveforms, which could make it more robust against overfitting.

In the lower panel of Figure 16 we show the ratio between the EIM and the ROB error for the same methods and test samples as in the upper panel. We observe that this ratio is in the range $1 \leq \sigma_{\text{EIM}}/\sigma_{\text{ROB}} \leq \|\hat{A}^{-1}\|_2^2$, as was derived in Eqs. (130,136). In general, we observe that the EIM errors obtained with the different methods are always considerably below the upper limit imposed by Eq. (130) ($\sigma_{\text{EIM}}/\sigma_{\text{ROB}} \ll \|\hat{A}^{-1}\|_2^2$). This is expected since to saturate this upper bound we need $\hat{P}^\dagger(\vec{h} - \vec{h}^{\text{ROB}})$ to be the eigenvector of $\hat{A}^\dagger \hat{A}$ with the maximum eigenvalue, which is hard to get in practice. We also observe that the *Training* method is almost optimal since most samples are close to the lower bound of $\sigma_{\text{EIM}}/\sigma_{\text{ROB}} \geq 1$. In contrast, most of the samples for the methods that do not involve training on waveform data, concentrate at values of $\sigma_{\text{EIM}}/\sigma_{\text{ROB}} \gtrsim 10^3$. This is probably because when we train on the waveform data, we are selecting an EIM that avoids coincidences between $\hat{P}^\dagger(\vec{h} - \vec{h}^{\text{ROB}})$ and eigenvectors of $\hat{A}^\dagger \hat{A}$ with large eigenvalues.

3. Construction of the ROQ

In this section, we describe how we use the methods of sections VF 1 and VF 2 to create, in an efficient way, an EIM that fits a waveform model over a parameter space with a tolerance better than σ .

We obtain an initial ROB $\{\vec{e}_i\}_{i=1}^n$ using Algorithm 1 and construct its corresponding EIM with Algorithm 5, where the set of training waveforms is the $\{h_A\}_{A=1}^N$, selected in the last step of Algorithm 1. We add elements to this initial ROB following a similar philosophy to that of Algorithm 1, in which we generate N_{lim} random waveforms, compute their EIM error σ_{EIM} , and select the N waveforms with largest EIM error for further study. Again, we want N to be as large as allowed by the memory (see Eq. (114)). We then compute the matrix M_{AB}^{ROB} for the N selected waveforms, find its Eigenvalues λ_B and compute the $n_\lambda < N$ most relevant eigenvectors in the waveform domain $\{\vec{u}_B\}_{B=1}^{n_\lambda}$, where the value of n_λ is again limited by the memory of the system. We iteratively select the eigenvector with the largest contribution to the EIM error, add it to the ROB and construct a new EIM with Algorithm 5 until all N waveforms are fitted with a tolerance better than the required one. The process is summarized in Algorithm. 6.

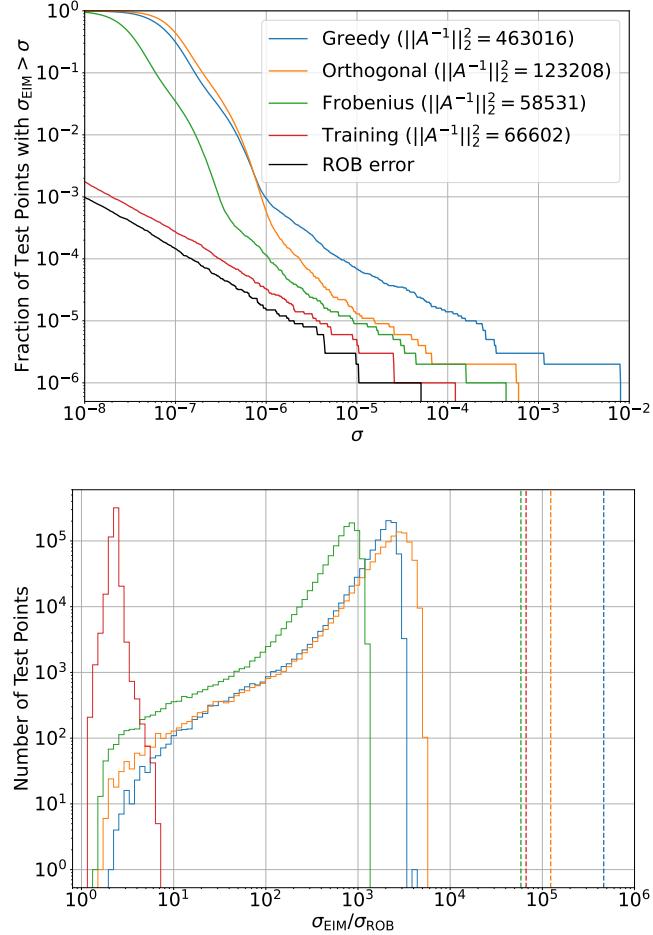


FIG. 16: Comparison of methods to compute the EIM for the 256s IMRPhenomPv2 ROB of Table III. We test the different EIMs on the same 10^6 samples randomly drawn from the parameter space over which the ROB is generated (see Table III). The *Greedy* method is the one outlined in [103], the *Orthogonal* method stands for algorithm 3, the *Frobenius* method corresponds to using algorithm 4 to minimize $\|\hat{A}^{-1}\|_F$ and the *Training* method is the one used to construct the EIM of Table III with algorithm 5. Upper panel: Fraction of samples with an EIM error larger than a tolerance σ as a function of σ . For comparison purposes, we also show the distribution of the ROB error. Lower panel: Histogram of the ratio between the EIM error and the ROB error for the same methods and test samples as in the upper panel. The vertical dashed lines represent an upper bound, defined by the value of $\|\hat{A}^{-1}\|_2^2$ for each method.

Algorithm 6 Enrich ROB to construct an EIM under tolerance

1: **Input:** Initial ROB $\{\vec{e}_i\}_{i=1}^n$ and EIM $\{\beta_i\}_{i=1}^n$, maximum number of waveforms selected N , tolerance σ , maximum number of waveforms computed N_{lim} , maximum number of eigenvectors used n_λ

2: **repeat**

3: Generate $N_{\text{lim},j}$ waveforms $\{h_A\}_{A=1}^{N_{\text{lim},j}}$ and compute their EIM error $\sigma_{\text{EIM},A}$

4: Select the N waveforms $\{h_A\}_{A=1}^N$ with largest σ_{EIM}

5: Save the minimum value of σ_{EIM} for the selected waveforms: $\sigma_{\text{EIM},\text{min}}$

6: $M_{AB}^{\text{ROB}} = \langle h_A, h_B \rangle - \sum_{c=1}^{n_j-1} \langle h_A, e_c \rangle \langle e_c, h_B \rangle$

7: Diagonalize M_{AB}^{ROB} and obtain eigenvalues λ_A and eigenvectors E_{AB}

8: Compute the n_λ normalized eigenvectors in waveform domain with largest $\delta\sigma_{A,\text{max}} = \lambda_A \max_B |E_{AB}|^2$: $\{\vec{u}_A\}_{A=1}^{n_\lambda}$

9: **repeat**

10: Compute the maximum contribution of each eigenvector to $\sigma_{\text{EIM},A}$:

11: $\{\delta\sigma_{A,\text{max}}^{\text{EIM}} = \lambda_A (1 + \|(\hat{P}^\dagger \hat{V})^{-1} \hat{P}^\dagger \vec{u}_B\|^2) \max_B |E_{BA}|^2\}_{A=1}^{n_\lambda}$

12: Find largest $\delta\sigma_{A,\text{max}}^{\text{EIM}}$: $A_{\text{sel}} = \operatorname{argmax}_A (\delta\sigma_{A,\text{max}}^{\text{EIM}})$

13: Add the corresponding eigenvector to the ROB: $\{\vec{e}_i\}_{i=1}^n \leftarrow \{\vec{e}_i\}_{i=1}^n \cup \{\vec{u}_{A_{\text{sel}}}\}$

14: Remove the selected eigenvector from the eigenvector list: $\{\vec{u}_A\}_{A=1}^{n_\lambda} \leftarrow \{\vec{u}_A\}_{A=1}^{n_\lambda} \setminus \{\vec{u}_{A_{\text{sel}}}\}$

15: Input $\{\vec{e}_i\}_{i=1}^n$, N_{rounds} , $\{\vec{u}_A\}_{A=1}^{n_\lambda}$ and their corresponding eigenvalues $\{\lambda_A\}_{A=1}^{n_\lambda}$ into Algorithm. 5 to obtain a new EIM $\{\beta_i\}_{i=1}^n$.

16: Find new error of selected waveforms $\{\sigma_{\text{EIM},A}^{\text{new}}\}_{A=1}^N$

17: **until** $\max_A \sigma_{\text{EIM},A}^{\text{new}} \leq \sigma$

18: **Output:** ROB $\{e_i\}_{i=1}^n$

G. Code Validation

In this section, we aim to quantify and assess the validity of the ROQ basis obtained using the algorithm described in VF. For that matter, we would like to evaluate the accuracy of the different basis in reconstructing the original waveform as well as the speed up gained. First, in Sec. VG 1 we describe the bases to be tested and compare them with examples found in the literature, in Sec. VG 2 we show the results of two statistical tests for the various bases, in Sec. VG 3 we comment on the theoretical and empirical speedups using the ROQ, and finally in Sec. VG 4 we compare the results of doing a parameter estimation analysis with the standard and the ROQ likelihoods.

1. Basis Generation and comparison with other ROQ methods

In this section, we describe how we generate the bases that will be used for testing and parameter estimation. We construct bases for both `IMRPhenomPv2` [111] and `IMRPhenomXPHM` [112]. Both waveform models take into account the effects of spin precession and `IMRPhenomXPHM` also includes higher order mode GW emission.

For `IMRPhenomPv2` we generate the bases listed in Table III, covering a chirp mass (\mathcal{M}) range between $0.95M_\odot$ and $45M_\odot$. Given that integration is performed from a low-frequency cutoff of 20Hz, we find bases duration ranging from 256s to 4s. For `IMRPhenomXPHM` we generate the bases listed in Table III, with chirp masses ranging between $2.18M_\odot$ and $110M_\odot$ and corresponding durations between 64s and 4s from 20Hz. We show in Fig 17 an example of an `IMRPhenomXPHM` waveform and its corresponding empirical interpolant. More specifically, the upper panel shows the real part of the plus polarization $\operatorname{Re}(\tilde{h}_+)$ and the lower panel, its square $|\tilde{h}_+|^2$ in the frequency domain. The corresponding interpolation nodes and empirical interpolant are shown to visually confirm the goodness of the fit to the original waveform. The parameters of the template are shown in the caption of Figure 17 and are selected so that the quadratic EIM error is equal to the median quadratic EIM error over the testing set of waveforms of the basis of Table IV covering $\mathcal{M} \in [10, 15]M_\odot$. We can observe how both the linear and quadratic parts have a complicated dependence on frequency, coming from the interference of the higher order modes with the main (2,2) mode. This is the principal reason for the larger number of linear and specially quadratic elements when comparing the basis of `IMRPhenomPv2` and `IMRPhenomXPHM`.

The 4s and 8s basis of `IMRPhenomPv2` and `IMRPhenomXPHM` are directly comparable with those published in Ref [105] computed using `PyROQ`, since they cover the exact same parameter space and frequency range. We observe that the number of basis elements in `PyROQ` and `EigROQ` is generally similar and we expect it to be smaller than that of a

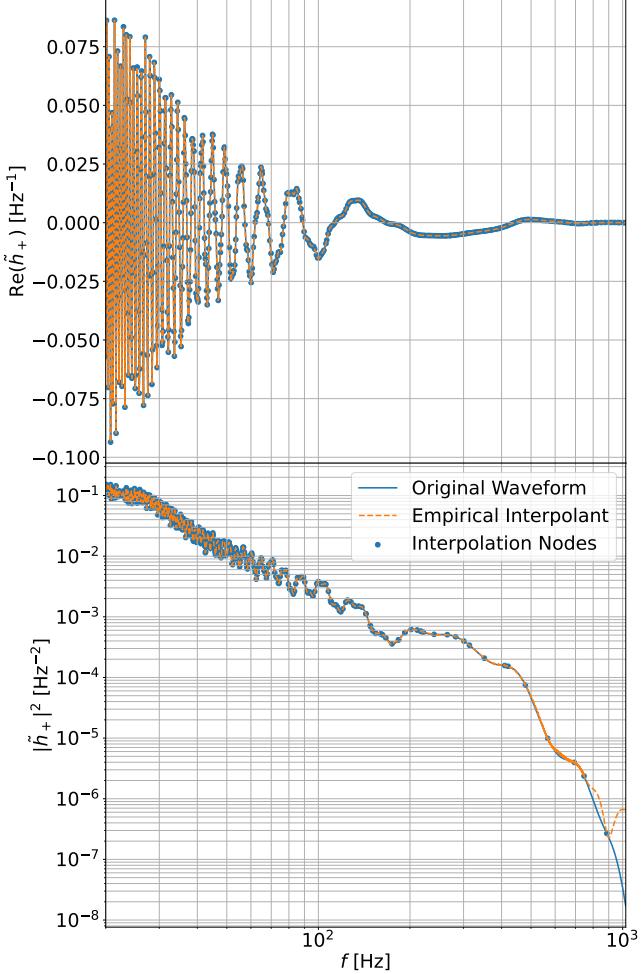


FIG. 17: Example of an `IMRPhenomXPHM` template and its empirical interpolant. In the upper panel, we show the real part of the plus polarization of the template $\text{Re}(\tilde{h}_+)$ as a function of the frequency and in the lower panel we depict its square $|\tilde{h}_+|^2$. We superimpose in each panel the corresponding interpolation nodes and empirical interpolants as defined in (117). The template shown has $\mathcal{M} = 13.6M_\odot$, $q = 2.61$, $\chi_{\text{eff}} = -0.011$ [119, 120], $\chi_p = 0.208$ [121] and inclination angle $\iota = 61.6^\circ$. Using the ROQ basis of Table IV covering $\mathcal{M} \in [10, 15]M_\odot$, we have linear and quadratic EIM errors of $\sigma_{\text{linear}}^{\text{EIM}} = 1.26 \cdot 10^{-9}$ and $\sigma_{\text{linear}}^{\text{EIM}} = 6.06 \cdot 10^{-8}$ respectively.

comparable basis constructed with `GreedyCPP`. However, the number of test points over the set tolerance is about an order of magnitude smaller in our bases than in `PyROQ`'s ones.¹⁹ We attribute this improvement to the way we approach the minimization in the error of the Empirical Interpolant. In the `PyROQ` algorithm, it is implicitly assumed that once a template is below the tolerance it will remain like this throughout the computation, which would be true if the EIM error were monotonically decreasing. This, however, is not true in general as adding new templates to the base, can deteriorate the fit, and in particular it can bring some of the waveforms which were under the tolerance, back over tolerance. The fact that this is happening can be explicitly seen in Ref [105] because the maximum EIM error in the training set is over the tolerance. To alleviate this problem, we simultaneously use the N waveforms with initially more EIM error, even if some of them are already below tolerance.

We have also extended the parameter space of the ROQ bases with respect to those computed by `PyROQ` in Ref [105], with durations up to 256s for `IMRPhenomPv2` and 64s for `IMRPhenomXPHM`. Doing this in `PyROQ` is computationally challenging since finding the template with the largest associated EIM error requires the recomputation of the waveforms in the training set many times. With our methods, this is no longer the case as we only need to compute any given

¹⁹ Note that while in Ref [105] the bases use 10^6 points for testing, we use 10^7 points.

Freq. range (Hz)		Δf (Hz)	$M_c(M_\odot)$		Basis size		$\sigma_{EI,max}$ Linear	$\sigma_{EI > 10^{-5}}$ Linear	Likelihood Speedup	
Min	Max		Min	Max	Linear	Quadratic			Theoretical	Empirical
20	1024	1/4	12.3	45	242	194	1.00×10^{-3}	1.09×10^{-4}	31	19
20	1024	1/8	7.93	14.76	369	294	4.91×10^{-4}	1.46×10^{-4}	55	31
20	2048	1/16	5.14	9.52	493	389	6.85×10^{-4}	5.72×10^{-4}	110	59
20	2048	1/32	3.35	6.17	631	438	6.88×10^{-4}	5.83×10^{-4}	98	75
20	2048	1/64	2.18	4.02	848	407	1.51×10^{-3}	5.71×10^{-4}	103	71
20	4096	1/128	1.42	2.60	1315	306	6.4×10^{-4}	2.46×10^{-3}	83	50
20	4096	1/256	0.95	1.72	2196	300	1.43×10^{-4}	6.32×10^{-5}	69	28

TABLE III: Summary of the reduced bases constructed with `EigROQ` for the `IMRPhenomPv2` waveform model. We limit the mass ratio $1 \leq q \leq 8$, the magnitudes of the two spins $-0.8 \leq \chi_i \leq 0.8$ for $i \in [1, 2]$, and the full range for the spin angles $(0, 0) \leq (\theta_J, \alpha_0) \leq (\pi, 2\pi)$. For the first base ($\Delta f = 0.25\text{Hz}$) we extend the coverage in spins to $-0.88 \leq \chi_i \leq 0.88$. For the creation of all the basis, we run `EigROQ` with the same configuration. In algorithm 1 we set the maximum number of waveform selected $N = 20000$, tolerances of each step $\sigma_i = [10^{-1}, 10^{-3}, 10^{-5}]$ and maximum number of waveforms computed per step $N_{\text{lim},i} = [10^6, 3.16 \cdot 10^6]$, and in algorithm 6 we set $N = 10^7$, $\sigma = 10^{-5}$, $N_{\text{lim}} = 10^7$ and the maximum number of eigenvectors used $n_\lambda = 5000$, except for the 256s basis where we set $n_\lambda = 4000$ due to memory limitations. The basis are tested on 10^7 randomly generated waveforms in the same parameter space as the training was done on. The “Theoretical” speedup has been computed with Eq. (149) while the “Empirical” speedup is the median and 90% credible interval of the corresponding points in the upper panel of Figure 20.

waveform once. This allows more complex case studies to be feasible.

2. Statistical tests

In this section, we perform 2 different statistical tests to check the faithfulness of the ROQ basis in gravitational waves inference, a likelihood test and a P-P test.

The likelihood test consists of a comparison of the log-likelihood ratios evaluated using the standard waveform with those obtained using the ROQ approximation. The log-likelihood ratio is defined as the ratio between the likelihood of Eq. (95) and the likelihood of the noise hypothesis ($h = 0$), that is

$$\log \mathcal{L}_{\text{ratio}}(d|\vec{\theta}) = \log \frac{\mathcal{L}(d|h(\vec{\theta}))}{\mathcal{L}(d|0)} = (d, h(\vec{\theta})) - \frac{1}{2}(h(\vec{\theta}), h(\vec{\theta})). \quad (147)$$

This quantity, which is just the likelihood of Eq. (95) removing the constant part that only depends on the data, is what we will be referring to as the log-likelihood throughout the rest of the text. The log-likelihood is the crucial quantity used in estimating the parameters of a given GW event, which is the ultimate end for which the ROQ is created. We perform likelihood tests on the `IMRPhenomPv2` and `IMRPhenomXPHM` bases described in table III and table IV respectively, and show the results on Figure 18. To obtain the difference in the log-likelihood, we create a random realization of Gaussian noise and inject a waveform calculated using the corresponding approximant. The injected waveforms’ parameters are randomly sampled from uniform distributions whose boundaries are the respective ROQs’ ranges of validity. We use a fixed distance of 100Mpc and randomly sample the incoming direction of the GW from a uniform distribution in the sky. We then compute the standard log-likelihood and the ROQ log-likelihood using the same injected waveform and compare them. What we plot is the relative difference between both logarithms for a total of $1.5 \cdot 10^5$ realizations. We see the maximum discrepancy lies below 0.1 for every case considered here, and the bulk of the samples lie below 10^{-3} .

Given that the likelihood is the only signal-dependent quantity that enters the computation of the posterior (Eq. (94)), as long as the ROQ and standard likelihoods agree reasonably well, we can expect the PE posteriors with and without the ROQ to be virtually the same. According to Wilks theorem [123] in the frequentist and large sample size limits, the quantity $-2\log\{\mathcal{L}/\mathcal{L}_{\text{max}}\}$ is distributed as a χ^2 with a number of degrees of freedom equal to the number of parameters being fitted by the PE. In the case of a CBC, we need 15 parameters to fully characterize the binary, although, since the azimuthal spin angles and phase of coalescence are usually so poorly constrained, in

Frequency range (Hz)	Δf (Hz)	$M_c(M_\odot)$	Basis size		Test set $\sigma_{\text{EI,max}}$		Test set $\sigma_{\text{EI}} > 10^{-4}$		Likelihood Speedup		
			Min	Max	Linear	Quadratic	Linear	Quadratic	Linear	Quadratic	Th.
20 1024	1/4	55 110	303	195	3.67×10^{-2}	2.47×10^{-2}	119	86	8.1	$3.2_{-0.6}^{+1.2}$	$1.4_{-0.3}^{+0.3}$
20 1024	1/4	35 66	339	192	6.95×10^{-2}	2.47×10^{-2}	115	64	7.6	$4.5_{-1.0}^{+1.7}$	$1.7_{-0.3}^{+0.5}$
20 1024	1/4	26 42	328	204	9.57×10^{-3}	1.04×10^{-2}	84	21	7.6	$6.1_{-1.1}^{+1.8}$	$2.2_{-0.5}^{+0.6}$
20 1024	1/4	18 33	348	201	1.80×10^{-2}	1.32×10^{-3}	70	19	7.3	$7.8_{-1.7}^{+0.4}$	$2.6_{-0.6}^{+0.7}$
20 1024	1/4	12 20	371	264	1.18×10^{-2}	1.03×10^{-3}	67	16	6.3	$7.5_{-1.6}^{+0.3}$	$3.1_{-0.6}^{+0.7}$
20 1024	1/8	10 15	491	386	4.32×10^{-3}	4.39×10^{-4}	50	6	9.2	$11.1_{-0.4}^{+0.3}$	$4.3_{-1.0}^{+1.2}$
20 1024	1/8	8.6 11.8	505	435	9.33×10^{-3}	1.96×10^{-4}	48	3	8.5	$10.5_{-0.9}^{+0.3}$	$4.8_{-1.0}^{+0.8}$
20 2048	1/16	5.1 9.6	868	942	2.95×10^{-3}	2.38×10^{-3}	56	11	17.9	$24.6_{-4.6}^{+2.3}$	$4.8_{-0.8}^{+1.2}$
20 2048	1/32	3.35 6.17	1539	1826	9.62×10^{-4}	2.53×10^{-4}	46	1	19.3	$27.6_{-0.8}^{+1.1}$	$4.6_{-0.9}^{+1.8}$
20 2048	1/64	2.18 4.02	2924	3636	6.37×10^{-4}	2.68×10^{-4}	19	7	19.8	$28.6_{-0.5}^{+0.7}$	$4.2_{-0.7}^{+1.7}$

TABLE IV: Summary of the reduced bases constructed with `EigROQ` for the `IMRPhenomXPHM` waveform model. We limit the mass ratio $1 \leq q \leq 4$, the magnitudes of the two spins $-0.8 \leq \chi_i \leq 0.8$ for $i \in [1, 2]$, and the full range for the spin angles $(0, 0) \leq (\theta_J, \alpha_0) \leq (\pi, 2\pi)$. For the creation of all the basis, we run `EigROQ` with the same configuration. In algorithm 1 we set the maximum number of waveform selected $N = 20000$, tolerances of each step $\sigma_i = [10^{-2}, 10^{-3}, 10^{-4}]$ and maximum number of waveforms computed per step $N_{\text{lim},i} = [10^6, 3.16 \cdot 10^6]$, and in algorithm 6 we set $N = 10^7$, $\sigma = 10^{-4}$, $N_{\text{lim}} = 10^7$ and the maximum number of eigenvectors used $n_\lambda = 5000$. The basis are tested on 10^7 randomly generated waveforms in the same parameter space as the training was done on. The “Theoretical” speedup has been computed with Eq. (149) while the “Empirical” speedup is the median and 90% credible interval of the corresponding points in the lower panel of Figure 20. For the empirical speedups, we show the values both without (Emp.) and with (MB) the `IMRPhenomXPHM` multibanding option enabled [122].

most cases the effective number of parameters is reduced to 12. We then expect $\log\{\mathcal{L}/\mathcal{L}_{\text{max}}\} = -5.7_{-4.8}^{+3.1}$, which is in accordance with most of the GW observations, specially those with high signal to noise ratio (SNR). Under the same model, the standard deviation of $\log \mathcal{L}$ is $\sigma_{\mathcal{L}} \sim \sqrt{N_{\text{eff}}/2} \sim \sqrt{6} \sim 2.4$, where N_{eff} is the effective number of parameters. Therefore, as long as the difference between the logarithm of the standard and the ROQ likelihoods is much smaller than $\sigma_{\mathcal{L}} \sim 2.4$, we expect the posteriors to be similar.

From Eq. (147), we observe that the likelihood ratio of a GW signal will approximately be given by $\log \mathcal{L} \sim \rho^2/2$, where ρ is the SNR. Therefore, the condition that $\Delta \log \mathcal{L} \ll 2.4$ can be translated into a condition on the SNR:

$$\rho \ll 2 \sqrt{\frac{\log \mathcal{L}}{\Delta \log \mathcal{L}}}, \quad (148)$$

which can be used to interpret figure 18 in terms of up to which SNR we can trust the posteriors obtained when using the corresponding ROQ. If we want the ROQ to be valid for the analysis of larger SNRs, we can always decrease the tolerance σ with which we generate it, at the expense of having more basis elements.

The second of the tests is the percent-percent (P-P) plot [4, 5]. P-P plots have been widely used in the literature [124] to validate codes that perform Bayesian parameter estimation (PE). Therefore, we use the P-P plots to directly test the ROQ’s faithfulness in its intended use. In this specific case, to make the P-P plots shown in Fig. 19, we use the posteriors pdfs resulting from performing PE on 200 injections. The PE analyses are done using the ROQ likelihood and the `dynesty` [125] sampler within the `Bilby` [126] framework and the injections use the same waveform model for which the corresponding ROQ was constructed. The priors of the PE and the distribution from which the injections are drawn are the same and coincide with the parameter space in which each ROQ basis has been constructed. For the extrinsic parameters we put priors which are uniform in the sky and in comoving volume, going to a maximum distance tailored for each chirp mass range to have detectable signals.

In the P-P plots of Fig. 19 we show the fraction of posterior pdfs for which the injected value of the parameter is found in a given confidence interval as a function of that same confidence interval. We expect the fraction of injected parameter values that fall into a particular confidence interval of the posterior pdfs to be drawn from a uniform distribution. We can thus assign a p-value quantifying such claim [5], individually for each binary parameter and jointly for all the parameters. The p-values are shown in the legends of figure 19. For all the PP-plots shown, the

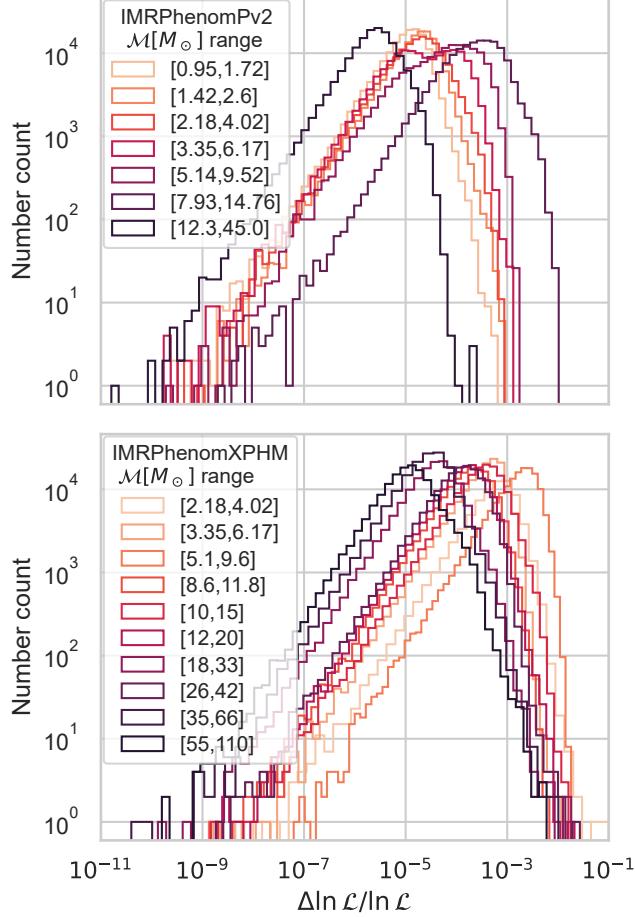


FIG. 18: Likelihood error tests for various \mathcal{M} ranges. Specifically, we plot $\Delta \ln \mathcal{L} / \ln \mathcal{L}$, that is, the fractional error of the $\ln \mathcal{L}$ when calculated with and without the ROQ. Upper pannel: IMRPhenomPv2. Lower pannel: IMRPhenomXPHM.

cdfs of the majority of the parameters fall within the $3 - \sigma$ regions, leading to p-values that are consistent with a uniform distribution. The combined p-values lie between 0.49 and 0.89, indicating that the posterior pdfs produced using these ROQs are well-calibrated.

3. Speedup analysis

The main purpose of the ROQ is to accelerate the computation of the GW likelihood. To test how good it is in this regard we perform a series of speed-up trials shown in figure 20. There are two quantities which we evaluate for the benchmarking test, the waveform and the Gaussian log-likelihood described in Eq (95). The tests consist in timing several calculations of both quantities for the standard case and the ROQ case. The sets of parameters used as inputs are drawn from uniform distributions with boundaries based on the range of validity of the corresponding ROQ basis. The ratio between the time for the standard method and the ROQ is what we call the empirical speedup, where we use the term empirical because we perform the actual likelihood and waveform computations using `python` [127] and the `Bilby` [126] framework. For IMRPhenomXPHM waveform speedups, we disable the default multibanding [122], which is used to speed up the full waveform computation by reducing the number of frequencies the model is evaluated at, and then interpolates between them. Therefore, we disable this to test if the model is linear with the number of frequencies at which it is evaluated. However, for the likelihood test, we compute the speedups both without and with multibanding enabled, to explore real-world speedup gains.

In figure 20 we differentiate the speedups using triangles for the waveform, squares for the log-likelihood and in the IMRPhenomXPHM case, circles for the log-likelihood with multibanding enabled. We can also compare with the

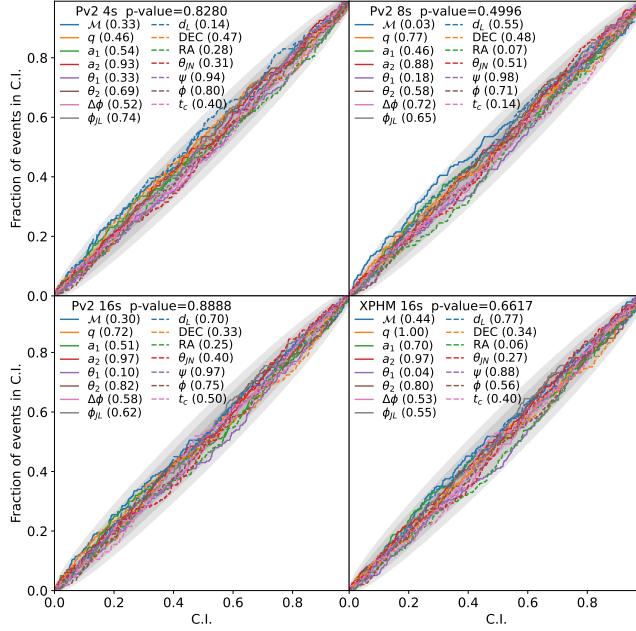


FIG. 19: P-P plots performed with different ROQ basis as stated in each subplot's legend. We show here the result of 200 injections being drawn from the corresponding ROQ-compatible prior, as stated in tables [III,IV]. The contours in grey delimit the 1σ , 2σ and 3σ regions. We plot a line for every parameter that uniquely characterizes a given CBC with consistent colors and styles across subplots. The lines represent the cumulative fraction of events.

theoretical speed-ups that are plotted as histogram-like bars in the figure. There are two kinds of bars, the solid ones represent the quantity

$$\text{Theoretical Speedup} = \frac{L}{N_L + N_Q} \quad (149)$$

where L is the number of frequencies for the waveform evaluation in the standard computation and N_L and N_Q are the frequency nodes for the linear and quadratic ROQ bases without factoring out repeated frequencies. This is the theoretical speedup that is usually attributed to the ROQ in the literature [104]. The dashed bars are the same quantity as in Eq. (149) when the frequencies belonging to both the linear and the quadratic interpolation list of frequency nodes are just considered once, thus the notation L/N_{LUQ} . In the ROQ likelihood we need to call the waveform model only once at the frequencies defined by $\{f_i\}_{i=1}^{N_{\text{LUQ}}} = \{F_j\}_{j=1}^{N_L} \cup \{\mathcal{F}_k\}_{k=1}^{N_Q}$, as is done in `Bilby`. Therefore, L/N_{LUQ} will be the theoretical speedup of the waveform evaluation if we assume that its computation time is proportional to the number of sampling points. For the `IMRPhenomXPHM` case, the difference between $N_L + N_Q$ and N_{LUQ} can be significant since there are many repeated interpolation nodes at low frequencies. The reason is that in the low-frequency region, the amplitude is larger and the waveform oscillates more rapidly than in the high-frequency part. Consequently, the interpolation nodes tend to concentrate at low frequencies leading some of them to coincide in the linear and quadratic ROQ. This behavior can be seen in figure 17.

For \mathcal{M} smaller than $\sim 20M_\odot$, we see that the waveform speedups are constant in the entire \mathcal{M} range of a given basis and are always close to the theoretical value of L/N_{LUQ} . This is in agreement with our expectations, since the `IMRPhenom` models describe the inspiral in a way that the computation time is linear with the number of sampling points, and their implementation in `LALsimulation` [128] being tested is written efficiently in `C` [129], with minimal overheads. In the case of large \mathcal{M} , above $\sim 20M_\odot$, the waveforms start being dominated by the merger and ringdown, the last two phases of a CBC, which are harder to model, and the speed-up of the `IMRPhenom` models is smaller than the theoretical expectation. This can be due to the waveform generation stopping above the ringdown frequency, meaning that the model is evaluated at fewer frequency points for high mass signals. Furthermore, when the waveform uses sufficiently few frequency points, fixed-costs associated with calculating post-Newtonian and phenomenological parameters of the model become important. Therefore, as \mathcal{M} increases, the trend of the waveform speedup is to decrease until a value of $\mathcal{O}(1)$ is reached and we have no speed up at all.

For the `IMRPhenomXPHM` likelihood speedups, we show both the results with and without disabling the default

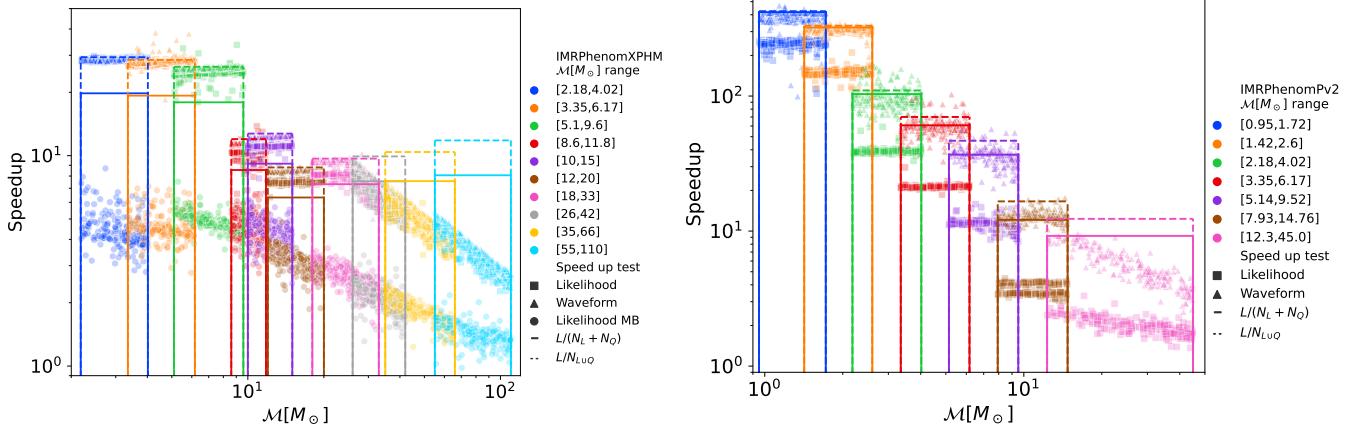


FIG. 20: Speed up factor for the `IMRPhenomXPHM` (upper panel) and `IMRPhenomPv2` (lower panel) waveforms in the different regions in chirp mass where the ROQ has been computed. We can differentiate theoretical and empirical speedups. The empirical speedups are calculated as the ratio between the time spent in computing the waveform without the ROQ and with it and are plotted as triangles. Squares are obtained in the same way but employing the Likelihood. In the `IMRPhenomXPHM` case, we include the likelihood speedups with the multibanding option enabled (circles) and disabled(squares). The theoretical speedups are drawn as bars. The dashed bars represent the speedup when array frequency duplications are accounted for while solid bars don't.

multibanding [122], which is used in the standard likelihood to speed up the full waveform computation. We observe that without multibanding `IMRPhenomXPHM` has a likelihood speedup very close to the theoretical value. This is due to the fact that the computation time of the likelihood is dominated by the waveform evaluation, and the `Bilby` implementation of the ROQ likelihood only generates the waveform once at the frequencies $\{f_i\}_{i=1}^{N_{L\cup Q}} = \{F_j\}_{j=1}^{N_L} \cup \{\mathcal{F}_k\}_{k=1}^{N_Q}$. However, when one includes the multibanding option, the `IMRPhenomXPHM` is already internally being evaluated in fewer frequency points, and therefore the speedup can be significantly lower than the expected one, although it still reaches median values that can be as large as 5, and which will be noticeable in PE applications. Looking at the targeted bases that are introduced in Table V, we observe that in this case, the speedup over the standard multibanded case can be even larger, reaching a value of $29.2^{+1.4}_{-4.6}$ for the base targeted at GW170817 [130].

In the `IMRPhenomPv2` case, we observe that the likelihood speedup is significantly below the waveform speedup and therefore, also below the theoretical speedup. To understand this discrepancy, we note that for the standard likelihood case, the computation time is dominated by evaluating the waveform in all the required frequencies and computing the overlap integrals of Eq. (95), both of which will be proportional to L . However, for the ROQ likelihood, the time to compute the waveform and overlap integrals is significantly reduced since they are proportional to $N_L + N_Q \ll L$. Given the fact that `IMRPhenomPv2` is much faster to generate than `IMRPhenomXPHM`, the computation time starts to be dominated by fixed-cost operations, which for example include computing the parameters of the waveform models, finding the detector responses as well as possible overheads.

To further explore this hypothesis, we model the computation time of the likelihood as a coefficient multiplying the number of frequencies being evaluated plus a constant term which represents the fixed-cost operations. Since for `IMRPhenomPv2`, $N_{L\cup Q} \sim N_L + N_Q$, we have,

$$T = A \cdot L + B \quad (150)$$

$$T_{ROQ} = a \cdot (N_L + N_Q) + b. \quad (151)$$

To compute the speedup, we divide Eq. (150) by Eq. (151), obtaining

$$f(L, N_L, N_Q; B, a, b) = \frac{L + B}{a(N_L + N_Q) + b}, \quad (152)$$

where we have divided all the coefficients by A , which is not expected to be 0. In figure 21 we show the ratio between the empirical and theoretical likelihood speedups, together with the best fit of our model in Eq. (152). We observe very good agreement between the model and the data. From the fitted values of B , a and b , also displayed in the plot, we can substantiate our hypothesis that the fixed-cost operations in the ROQ likelihood is making the empirical speedup of the `IMRPhenomPv2` smaller than the theoretical value. We find a value of $a = 1.00 \pm 0.16$, and therefore,

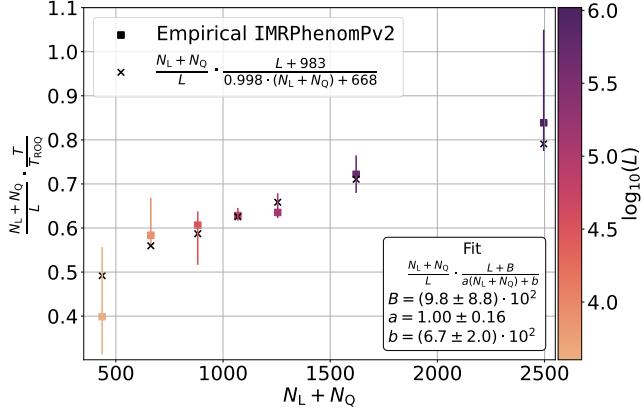


FIG. 21: Ratio between empirical speedup and the theoretical speedup of Eq. (149), plotted as a function of the total elements of the ROQ basis ($N_L + N_Q$) for `IMRPhenomPv2`. The colour of the error bars encodes the logarithm of the number of frequencies where the waveform is evaluated in the standard computation $\log \mathcal{L}$. In the bottom right box, we show the functional form we fit, which comes from Eq. (152), as well as the 1σ uncertainty for the three fitted parameters. We also plot as black crosses the results obtained evaluating the best fit in the data points.

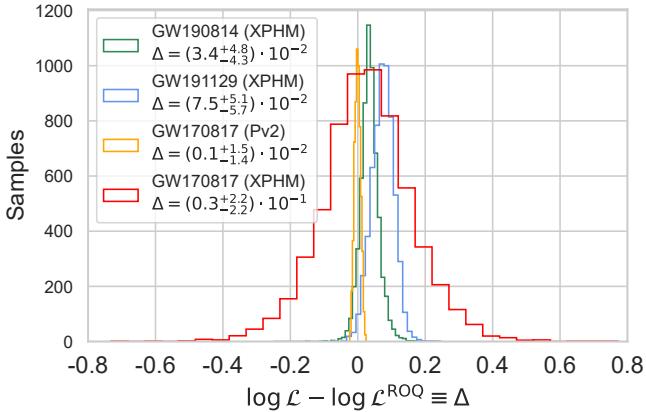


FIG. 22: Difference between the logarithm of the standard Likelihood and the logarithm of the ROQ Likelihood for the three events analysed.

from Eq. (152), we observe that if the coefficients B and b describing the fixed-costs were 0, we would recover the theoretical speedup result. However, since we find a value of $b = (6.7 \pm 2.0) \cdot 10^2$, the `IMRPhenomPv2` speedup is reduced with respect to the theoretical unless we have a very large number of basis elements such that $a \cdot (N_L + N_Q) \gg b$.

4. Application to GW events

We now perform four PE analyses [95] on three confirmed GW events using the ROQ approximation. More specifically, we use the `IMRPhenomXPHM` 16s basis described in table IV for the GW191129_134029 [11] event and the `IMRPhenomPv2` 256s basis of table III for the GW170817 [130] event. For the other two PE analyses of GW190814 [18] and GW170817 with `IMRPhenomXPHM`, in a similar spirit to Refs [106, 107], we construct targeted ROQ bases with narrow \mathcal{M} ranges, listed in Table V. These bases are centered on the search \mathcal{M} value and have a narrow width tuned to be larger than the expected chirp mass resolution. Note that the bases have been generated using a factor of 10 times fewer waveforms than that of Tables III IV, since the parameter space they cover is smaller.

The analyses use the ROQ likelihood and the `dynesty` [125] sampler within version 2.1.0 of `Bilby` [126] and the version 5.1.0. of `LALSimulation`. The PSDs employed were estimated using `BayesWave` [131, 132] and are those used by the LVK collaboration for the public analysis of the events. We also include the effects of calibration

Freq. range (Hz)		Δf (Hz)	$M_c(M_\odot)$		Basis size		Likelihood Speedup		
Min	Max		Min	Max	Linear	Quadratic	Th.	Emp.	MB
20	2048	1/16	6.2	6.6	1090	816	17.0	$21.6^{+3.8}_{-3.6}$	$4.8^{+1.1}_{-0.6}$
20	2048	1/256	1.195	1.200	1392	2007	152.7	$151.8^{+4.5}_{-4.1}$	$29.2^{+1.4}_{-4.6}$

TABLE V: Focused `IMRPhenomXPHM` bases for GW190814 ($\Delta f = 1/16$ Hz) and GW170817 ($\Delta f = 1/256$ Hz). We limit the magnitudes of the two spins $-0.8 \leq \chi_i \leq 0.8$ for $i \in [1, 2]$, and the full range for the spin angles $(0, 0) \leq (\theta_J, \alpha_0) \leq (\pi, 2\pi)$. For the GW190814 we limit the mass ratio $q \leq 16$ while for GW170817 we limit it $q \leq 4$. For the creation of the two basis, we run `EigROQ` with the same configuration. In algorithm 1 we set the maximum number of waveform selected $N = 20000$, tolerances of each step $\sigma_i = [10^{-2}, 10^{-3}, 10^{-4}]$ and maximum number of waveforms computed per step $N_{\text{lim},i} = [10^5, 3.16 \cdot 10^5]$, and in algorithm 6 we set $N = 10^6$, $\sigma = 10^{-4}$, $N_{\text{lim}} = 10^7$ and the maximum number of eigenvectors used $n_\lambda = 5000$. The “Theoretical” speedup has been computed with Eq. (149) while the “Empirical” speedup is the median and 90% credible interval of the corresponding points in the lower panel of Figure 20. For the empirical speedups, we show the values both without (Emp.) and with (MB) the `IMRPhenomXPHM` multibanding option enabled [122].

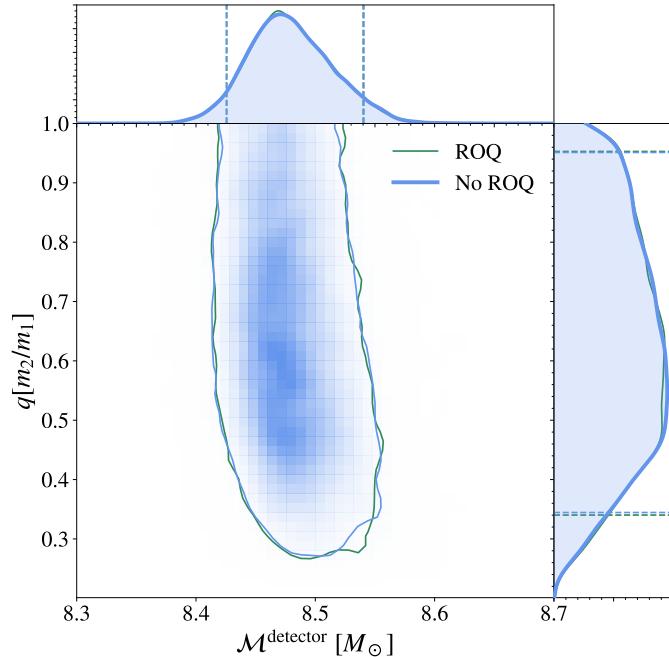


FIG. 23: Posterior distributions for the mass ratio and \mathcal{M} in the detector frame for the ROQ and non-ROQ analysis of GW191129_134029. The 90% credible regions are indicated by the solid contour in the joint distribution, and by the dashed vertical and horizontal lines in the marginalized distributions.

uncertainties [133–135] in the phase and the amplitude.

The first event we discuss is GW191129_134029 [11, 136]. This is an event with $\mathcal{M}^{\text{detector}} = 8.48^{+0.06}_{-0.05} M_\odot$ so we can use the 16 seconds `IMRPhenomXPHM` ROQ basis. It has a relatively big median network SNR of 13.1, allowing us to put tight constraints on the parameters and better see if any differences arise between the ROQ and the standard posterior. We perform two `Bilby` runs with the exact same configuration, one using the standard GW likelihood and the other using the ROQ likelihood.

In figure 22 we show the difference between the logarithm of the standard and the ROQ likelihoods, for the posterior samples of the PE with the ROQ likelihood. This difference corresponds to the ROQ error in the log-likelihood. We find a 90% c.l. error of $\Delta \log \mathcal{L} = 0.075^{+0.051}_{-0.057}$. Since $\Delta \log \mathcal{L} \ll 1$, we expect the posteriors with and without the ROQ to be almost the same. Using that the log likelihood of this event is $\log \mathcal{L} = 84.2^{+2.9}_{-4.1}$, the fractional error in the ROQ

log-likelihood computation is $\delta_{\mathcal{L}} = (9.1^{+6.3}_{-6.8}) \cdot 10^{-4}$.²⁰ The distribution of errors is centered at a positive value, as one would expect if the waveform model were a good representation of reality since any error in the ROQ modelization of the waveform would push it away from the true GW and thus, to a lower likelihood value.

In figure 23 we corroborate that indeed the posteriors with and without the ROQ are similar by showing the corresponding distributions for the detector frame chirp mass \mathcal{M} and the mass-ratio q . We find a Jensen-Shannon Divergence (JSD) [137] of $1.3 \cdot 10^{-4}$ and $1.9 \cdot 10^{-4}$ respectively, robustly assessing the similarity between the distributions with and without the ROQ approximation.

The second event we analyze is GW190814 [18, 136]. This event was measured to have a chirp mass of $\mathcal{M} = 6.42^{+0.02}_{-0.02} M_{\odot}$ and a very unequal mass ratio of $0.11^{+0.01}_{-0.01}$, which is below the mass ratios of $q > 0.25$ explored in the bases of Table IV. Therefore we create a targeted ROQ base with 16 seconds of duration, and chirp mass range from $6.2 M_{\odot}$ to $6.6 M_{\odot}$ for the **IMRPhenomXPHM** waveform. In figure 22 we show the ROQ log-likelihood errors of the posterior samples of the PE performed using this targeted basis. We have that $\Delta \log \mathcal{L} = 0.034^{+0.048}_{-0.043}$ which is similar in magnitude to that of GW191129_134029. Again, since $\Delta \log \mathcal{L} \ll 1$, we expect the posteriors with and without the ROQ to be almost the same. However, for this event, the log-likelihood is larger, at $\log \mathcal{L} = 310.3^{+3.1}_{-5.0}$, and therefore the relative error in the ROQ log-likelihood computation is smaller, at $\delta_{\mathcal{L}} = (1.1^{+1.5}_{-1.4}) \cdot 10^{-4}$.

The last GW event we analyze is GW170817 [130], the event with the largest Network SNR (~ 33) ever detected. It was identified as a binary neutron star with $\mathcal{M} = 1.1976^{+0.0004}_{-0.0002}$ [138] and we use it to probe the longest of our **IMRPhenomPv2** bases with 256s in duration as well as a targeted ROQ using **IMRPhenomXPHM** for such long signals. For our analysis, we make use of the public strain data after noise subtraction [139]. In figure 22 we show the ROQ log-likelihood errors of the posterior samples of both PEs. For both cases, we do not expect the ROQ error to significantly impact the posterior, since $\Delta \log \mathcal{L} \ll 2.4$. The **IMRPhenomPv2** PE has an order of magnitude smaller ROQ error than the **IMRPhenomXPHM** case. This is most likely the result of the **IMRPhenomPv2** basis being constructed with a tolerance $\sigma = 10^{-5}$, which is an order of magnitude smaller than the tolerance $\sigma = 10^{-4}$ used in the **IMRPhenomXPHM** case. In the **IMRPhenomPv2** case, the log-likelihood is $536.1^{+3.2}_{-4.3}$ and the corresponding fractional error is $\delta_{\mathcal{L}} = (-0.1^{+2.8}_{-2.6}) \cdot 10^{-5}$. In the **IMRPhenomXPHM** case, we find a larger likelihood of $538.1^{+4.3}_{-5.1}$, which is expected since the higher order modes give more freedom to the waveform to fit the data. The corresponding fractional error is $\delta_{\mathcal{L}} = (0.5^{+4.1}_{-4.0}) \cdot 10^{-4}$. Comparing the Bayes Factors of both PE runs, adjusted to have the same priors, we find $\log \mathcal{B} = 1.1 \pm 0.3$ in favour of **IMRPhenomXPHM**, which can be taken as evidence for Higher Order Modes in the signal. This highlights the importance of considering all physical effects of the waveform. To further make this point, we show in figure 24 how the addition of the Higher Order Modes improves the determination of the mass ratio and the inclination angle θ_{JN} , even for this low mass CBC for which the Higher Order Modes are harder to measure in LIGO-Virgo [96].

H. Conclusions

In this paper, we have explored in-depth Reduced Order Quadrature (ROQ) methods applied to GW data analysis and have presented novel algorithms to improve different aspects of the ROQ bases construction. ROQ methods offer a significant advantage by reducing the computational burden associated with likelihood evaluations, especially for long-duration waveforms, and therefore can greatly speed up parameter estimation analyses. Existing procedures for constructing ROQ bases encounter challenges in approximating waveforms that include complicated features such as precession or Higher Order Modes. We present algorithms to address these limitations by making use of SVD methods to characterize the waveform space and choose a reduced order basis close to optimal. We also propose improved methods to select the empirical interpolation nodes, greatly reducing the error induced by the empirical interpolation model.

We have demonstrated the effectiveness of our algorithm by constructing multiple ROQ bases for the **IMRPhenomPv2** and **IMRPhenomXPHM** waveforms, ranging in duration from 4s to 256s. These bases have been subjected to various tests, including likelihood error tests and P-P tests, to validate their accuracy and trustworthiness for data analysis applications. The speedup of these bases has also been empirically explored, confirming that ROQ methods provide close to the expected reduction in computational time compared to traditional likelihood calculations.

Furthermore, we have performed end-to-end parameter estimation analyses on several confirmed GW events. The results provide compelling evidence of the algorithm's ability to generate ROQ bases that accurately represent complex waveform models over both broad and targeted parameter spaces. By directly comparing the posterior distributions using the ROQ and standard methods and understanding the log-likelihood error distributions, we validate that our bases can straightforwardly be incorporated into current pipelines to produce precise and unbiased Parameter Estimations in real gravitational wave detector data.

²⁰ We define the fractional error in the ROQ log-likelihood computation as $\delta_{\mathcal{L}} = \Delta \log \mathcal{L} / \log \mathcal{L}$

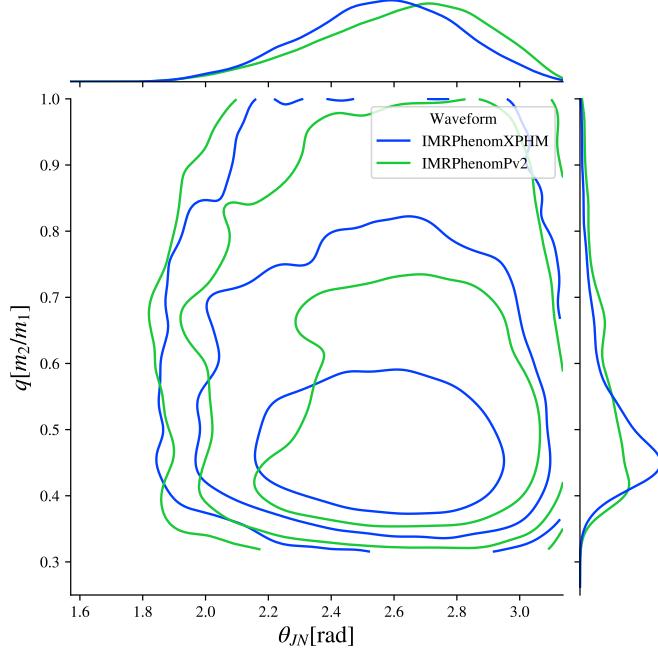


FIG. 24: Posterior distributions for the mass ratio q and the inclination angle θ_{JN} for the ROQ analysis of GW170817. In blue we plot the IMRPhenomXPHM run and in green IMRPhenomPv2. Three contours per run delimit the 1σ (68.3% C.L.), 2σ (95.4% C.L.) and 3σ (99.7% C.L.) credible regions in the joint $q - \theta_{JN}$ distribution. Note that the non-continuous behaviour of the contours near the border is an artefact of the Gaussian kernel employed in the drawing. This is expected whenever the parameter is bounded and presents many samples close to the border.

In conclusion, the algorithms introduced in this paper represent a step forward in the quest to efficiently exploit the capabilities of advanced gravitational wave detectors. We improve upon previous ROQ construction algorithms allowing for more efficient bases in regions of parameter space that were previously inaccessible. As gravitational wave astronomy continues to evolve, and the number of events detected per year continues to grow, having fast and accurate techniques to perform Parameter Estimation will undoubtedly play a vital role in maximizing the scientific potential of future observatories and advancing our knowledge of the Universe.

I. Appendix: Fast way to update $\|\hat{A}^{-1}\|_F$ and $\sigma_{\text{EIM}}^{\text{tot}}$

In this section we assume that we have the inverse of the matrix $\bar{A}_{ij} = e_j(X_i)$ and its Frobenius norm $\|\hat{A}^{-1}\|_F$, defined in Eq. (140), and we want to compute the inverse and Frobenius norm of the inverse of the matrix A_{ij} , defined as:

$$A_{ij} = \begin{cases} q_j(x_{\beta_i}) & i \neq k \\ e_k(x_{\beta'_k}) & i = k \end{cases} \quad (153)$$

which is just the result of changing the row k of \bar{A}_{ij} . We then use the fact that, from the properties of the inverse \bar{A}_{ij} , we have:

$$(\hat{A}\hat{A}^{-1})_{ij} = \begin{cases} q_j & i \neq k \\ \sum_{l=1}^n e_l(x_{\beta'_k})(\hat{A}^{-1})_{lj} \equiv c_j & i = k \end{cases} \quad (154)$$

Since the matrix of Eq. (154) has such a simple structure, it can be analytically inverted as:

$$((\hat{A}\hat{A}^{-1})^{-1})_{ij} = \begin{cases} \frac{d_{ij}}{c_j} & i \neq k \\ \frac{c_j}{c_k} & i = k, j \neq k \\ \frac{1}{c_k} & i = j = k \end{cases} \quad (155)$$

And we can use that $\hat{A}^{-1} = \hat{A}^{-1}(\hat{A}\hat{A}^{-1})^{-1}$ to show that:

$$(\hat{A}^{-1})_{ij} = \begin{cases} (\hat{A}^{-1})_{ij} - \frac{c_j}{c_k}(\hat{A}^{-1})_{ik} & j \neq k \\ \frac{1}{c_k}(\hat{A}^{-1})_{ik} & j = k \end{cases} \quad (156)$$

We observe that this way of computing the inverse will require $O(n^2)$ for computing c_j with Eq. (154) and also $O(n^2)$ operations to update each element of \hat{A}^{-1} using Eq. (156). So the total number of operations will be $O(n^2)$, much smaller than the $O(n^3)$ required to directly invert the matrix.

Using this expression for the updated inverse we can find a way to update also the Frobenius norm of the inverse, which is given by:

$$\begin{aligned} \|\hat{A}^{-1}\|_F^2 &= \sum_{i,j=1}^n |(\hat{A}^{-1})_{ij}|^2 \\ &= \sum_{i,j=1}^n (\hat{A}^{-1})_{ij} - \frac{c_j}{c_k}(\hat{A}^{-1})_{ik}^2 + \sum_{i=1}^n \frac{1}{c_k}(\hat{A}^{-1})_{ik}^2 \\ &= \|\hat{A}^{-1}\|_F^2 + \frac{1}{|c_k|^2} \left(1 + \sum_{j=1}^n |c_j|^2 \right) \left[\sum_{i=1}^n |(\hat{A}^{-1})_{ik}|^2 \right] \\ &\quad - 2\operatorname{Re} \left(\frac{1}{c_k} \sum_{j=1}^n \left[\sum_{i=1}^n (\hat{A}^{-1})_{ik}(\hat{A}^{-1})_{ij}^* \right] c_j \right), \end{aligned} \quad (157)$$

where we can precompute with $O(n^2)$ operations the factors in square brackets that only depend on \hat{A} for each row k which we will change, and afterwards, updating the Frobenius norm will only need $O(n)$ operations on top of the $O(n^2)$ operations needed to compute c_j for each new row q we want to test. Since with Eq. (157) we do not need to update the inverse each time that we want to update its Frobenius norm, we can avoid the $O(n^2)$ memory allocations that are needed in Eq. (156).

We will now also look for a method to rapidly compute $\sigma_{\text{EIM}}^{\text{tot}}$. We assume that we have the value computed for an EIM whose variables we denote with a bar over them:

$$\bar{\sigma}_{\text{EIM}}^{\text{tot}} = \sum_{B=1}^{n_\lambda} \left(\lambda_B + \hat{A}^{-1} \bar{v}_B^2 \right), \quad (158)$$

where we have defined

$$\bar{v}_{B,i} = \sqrt{\lambda_B} u_{B,\beta_i}. \quad (159)$$

When we change the k 'th interpolation node of the EIM from β_k to β'_k , this becomes:

$$v_{B,i} = \begin{cases} \sqrt{\lambda_B} u_{B,\beta_i} & i \neq k \\ \sqrt{\lambda_B} u_{B,\beta'_k} & i = k \end{cases} \quad (160)$$

And the value of multiplying \hat{A} by \bar{v}_B will change to:

$$\begin{aligned}
(\hat{A}^{-1} \vec{v}_B)_i &= \sum_{j=1}^n (\hat{A}^{-1})_{ij} v_{B,j} \\
&= \sum_{j=1}^n \left(\hat{A}_{ij}^{-1} - \frac{c_j}{c_k} \hat{A}_{ik}^{-1} \right) v_{B,j} + \frac{1}{c_k} \hat{A}_{ik}^{-1} v_{B,k} \\
&= \underbrace{\sum_{j=1}^n \hat{A}_{ij}^{-1} \bar{v}_{B,j}}_{\bar{\Omega}_{B,i}} + \underbrace{\left[\frac{1}{c_k} v_{B,k} - \sum_{i=1}^n c_j \bar{v}_{B,j} \right]}_{\Theta_B} \underbrace{\hat{A}_{ik}^{-1}}_{\bar{\Gamma}_i},
\end{aligned} \tag{161}$$

where we have used the updated value of \hat{A}^{-1} computed in Eq. (156) and we put bar over the variables that do not depend on the value of the new interpolation node. Using Eq. (161), $\bar{\sigma}_{\text{EIM}}^{\text{tot}}$ becomes:

$$\begin{aligned}
\sigma_{\text{EIM}}^{\text{tot}} &= \sum_{B=1}^{n_\lambda} \lambda_B + \sum_{i=1}^n \left(\bar{\Omega}_{B,i} + \Theta_B \bar{\Gamma}_i^2 \right) \left(\bar{\Omega}_{B,i} + \Theta_B \bar{\Gamma}_i^2 \right)^* \\
&= \bar{\sigma}_{\text{EIM}}^{\text{tot}} + \sum_{B=1}^{n_\lambda} |\Theta_B|^2 \left[\sum_{i=1}^n |\bar{\Gamma}_i|^2 \right] \left[\sum_{i=1}^n |\bar{\Gamma}_i|^2 \right]^* \\
&\quad + 2\text{Re} \left\{ \sum_{B=1}^{n_\lambda} \Theta_B \left[\sum_{i=1}^n \bar{\Omega}_{B,i}^* \bar{\Gamma}_i \right] \right\}.
\end{aligned} \tag{162}$$

In general we will assume that $n_\lambda \gg n$. For each row k that we change, we can precompute with $O(nn_\lambda)$ operations all the factors in square brackets that will stay constant. Afterwards, the computational complexity of updating the value of $\sigma_{\text{EIM}}^{\text{tot}}$ will require $O(nn_\lambda)$ operations for computing Θ_B and only $O(n_\lambda)$ additional operations to evaluate Eq. (162).

VI Analysis of a subsolar-mass compact binary candidate from the second observing run of Advanced LIGO

A. Journal & publication details

Physics of the Dark Universe, Volume 42, December 2023 [6]

B. Authors

The authors in the same order as presented to the journal are

Gonzalo Morrás

Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco 28049 Madrid, Spain

José Francisco Nuño Siles

Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco 28049 Madrid, Spain

Juan García-Bellido

Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco 28049 Madrid, Spain

Ester Ruiz Morales

Departamento de Física, ETSIDI, Universidad Politécnica de Madrid, 28012 Madrid, Spain Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco 28049 Madrid, Spain

Alexis Menéndez-Vázquez

Institut de Física d'Altes Energies (IFAE), Barcelona Institute of Science and Technology, E-08193 Barcelona, Spain

Christos Karathanasis

Institut de Física d'Altes Energies (IFAE), Barcelona Institute of Science and Technology, E-08193 Barcelona, Spain

Katarina Martinovic

Theoretical Particle Physics and Cosmology Group, Physics Department, King's College London, University of London, Strand, London WC2R 2LS, UK

Khun Sang Phukon

Nikhef - National Institute for Subatomic Physics, Science Park, 1098 XG Amsterdam, The Netherlands. Institute for High-Energy Physics, University of Amsterdam, Science Park, 1098 XG Amsterdam, The Netherlands. Institute for Gravitational and Subatomic Physics, Utrecht University, Princetonplein 1, 3584 CC Utrecht, The Netherlands. School of Physics and Astronomy and Institute for Gravitational Wave Astronomy, University of Birmingham, Edgbaston, Birmingham, B15 9TT, United Kingdom

Sebastien Clesse

Service de Physique Théorique, Université Libre de Bruxelles (ULB), Boulevard du Triomphe, CP225, B-1050 Brussels, Belgium

C. Abstract

We perform an exhaustive follow-up analysis of a subsolar-mass (SSM) gravitational wave (GW) candidate reported by Phukon et al. from the second observing run of Advanced LIGO. This candidate has a reported signal-to-noise ratio (SNR) of 8.6 and false alarm rate of 0.41 yr which are too low to claim a clear gravitational-wave origin. When improving on the search by using more accurate waveforms, extending the frequency range from 45 Hz down to 20 Hz, and removing a prominent blip glitch, we find that the posterior distribution of the network SNR lies mostly below the search value, with the 90% confidence interval being $7.94^{+0.70}_{-1.05}$. Assuming that the origin of the signal is a compact binary coalescence (CBC), the secondary component is $m_2 = 0.76^{+0.50}_{-0.14} M_\odot$, with $m_2 < 1 M_\odot$ at 84% confidence level, suggesting an unexpectedly light neutron star or a black hole of primordial or exotic origin. The primary mass would be $m_1 = 4.71^{+1.57}_{-2.18} M_\odot$, likely in the hypothesized lower mass gap and the luminosity distance is measured to be $D_L = 124^{+82}_{-48} \text{ Mpc}$. We then probe the CBC origin hypothesis by performing the signal coherence tests, obtaining a log Bayes factor of 4.96 ± 0.13 for the coherent vs. incoherent hypothesis. We demonstrate the capability of performing a parameter estimation follow-up on real data for an SSM candidate with moderate SNR. The improved sensitivity of O4 and subsequent LIGO-Virgo-KAGRA observing runs could make it possible to observe similar signals, if present, with a higher SNR and more precise measurement of the parameters of the binary.

D. Introduction

The development of gravitational wave astronomy, with about 90 compact binary coalescence (CBC) events detected so far [11, 15, 22, 140–142] by the LIGO-Virgo-KAGRA (LVK) collaboration [143], is driving a true revolution in astrophysics and cosmology. As the number of detected events grows with successive observing catalogs, the range of the inferred component masses has extended to previously unexplored regions, with black holes (BH) found [19, 144] in the pair-instability mass gap [145] and in the hypothesized lower mass gap [18]. The frequency range of the LIGO [86] and Virgo [87] detectors also makes them sensitive to CBC signals in which one of the compact objects has a mass below $1M_{\odot}$. The detection of a subsolar mass (SSM) compact object would be of utmost interest since it would require either modification of the standard astrophysical evolution and collapse of ordinary matter or a new formation mechanism operating in the Universe, such as primordial black holes (PBHs) [146–150] or SSM objects originated by dark matter with exotic properties [151–163].

Before the advent of GW astronomy, the only way to detect SSM black holes was via X-ray binaries [164] or microlensing [165]. At present, some hints of the existence of such light black holes come from microlensing events toward the bulge [166], from Andromeda [167] and lensed quasars [168, 169], although the mass and the abundance of the lenses remain uncertain. Complementary to these astrophysical searches, GW signals of CBC with at least one subsolar component have been searched for in the first (O1), second (O2) and third (O3) observing runs of LVK [170–174], without finding compelling evidence for a clear detection. However, a further search in the O2 data for SSM black holes with low mass ratios [7] and the latest O3b SSM search results from LVK [174] have reported several potential candidates with a false alarm rate smaller than 2 yr^{-1} .

In this work, we follow up on the O2 search reported in [7], using the standard parameter estimation (PE) methods to further investigate the candidates reported in Table I. Given that PE on these long GW signals is extremely time consuming, we have focused on the third candidate of the table, which is the lowest FAR trigger found in coincidence by both LIGO Hanford and LIGO Livingston interferometers, which allows more confident rejection against terrestrial noise. This candidate was observed on April 1st 2017 and we will refer to it here as SSM170401.

In this analysis we have extended the frequency range of the search from 45 Hz down to 20 Hz. We have also improved upon the `TaylorF2` [175] waveform used in the template bank of search, by using for PE the more accurate waveforms `IMRPhenomPv2` [111] and `IMRPhenomXPHM` [112] that include the merger and ringdown phases, as well as spin precession and, in the case of `IMRPhenomXPHM`, higher order modes. We have also inspected the quality of the data and discussed the impact of a prominent glitch removal using standard tools such as `BayesWave` [131, 132, 176]. Finally, assuming that the origin of the candidate is a BBH merger, the PE allows us to infer the component masses, spins, distance and sky location, as well as the posterior probability of having an SSM component of the hypothetical source for SSM170401.

In the following sections, we describe in detail different aspects of this analysis that reveal the peculiarities and difficulties of doing PE on this type of candidate, as well as the necessary analysis tools in preparation for a possible future significant candidate, given the increase of sensitivity expected in the O4 run.

E. Significance of SSM170401

The candidate was found in data taken on April 1st, 2017, 01:43:34 UTC during the O2 LIGO-Virgo observing run. It was not reported by any of the LVK searches, both generic [15] and SSM specific [177], but it was found in a dedicated search for SSM mergers in asymmetric binaries using the `GstLAL` pipeline [7]. The search reported detector frame masses of $4.897 M_{\odot}$ and $0.7795 M_{\odot}$, with a false-alarm-rate (FAR) of 0.4134 yr^{-1} and a combined network signal-to-noise ratio (SNR) of ~ 8.67 . Given that the time of O2 coincident data suitable for observation is $T_{\text{obs}} = 118$ days [15], the false alarm probability (FAP) of this candidate, according to the search is:

$$\text{FAP} = 1 - \exp\{-\text{FAR} \cdot T_{\text{obs}}\} = 0.12. \quad (163)$$

The interpretation of this FAP is that the search would produce a higher-ranked candidate in 12% of trials over data containing only noise.

We can also estimate an upper bound for the probability of this signal coming from a CBC merger with an SSM component, given the upper limits on event rates obtained from the O3 SSM searches. [173, 174]. In Ref. [174] the 90% C.L. constraints on the merger rate \mathcal{R}_{90} of SSM binaries are reported in the (m_1, m_2) plane, assuming null results of these searches. For the median values of the component masses of the source of SSM170411 (see Table VI), we find $\mathcal{R}_{90} \sim 2 \times 10^2 \text{ Gpc}^{-3} \text{ yr}^{-1}$. Moreover, in the search where the signal was identified [7], the volume-time surveyed for these same masses is reported to be $\langle VT \rangle \sim 3 \times 10^{-3} \text{ Gpc}^3 \text{ yr}$. Since the arrival of GWs from binary mergers to the

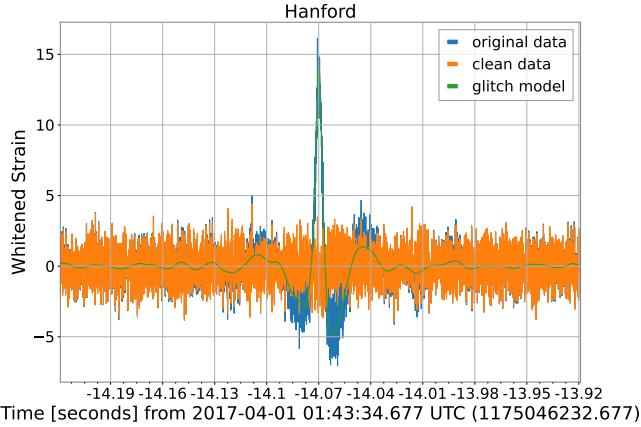


FIG. 25: Figure showing the Hanford original whitened strain $\tilde{h}^{\text{whitened}}(f) = \tilde{h}(f)/\sqrt{S_n(f)}$, the whitened glitch model and the whitened clean data after subtracting the glitch. Times are shown relative to the trigger time.

detectors is Poisson distributed, with an expected number of events $\mu = \mathcal{R}\langle VT \rangle$, the probability of finding n events would be:

$$P(n) = \frac{\mu^n}{n!} e^{-\mu}. \quad (164)$$

Using the values previously mentioned for \mathcal{R}_{90} and $\langle VT \rangle$, the upper bound on the expected number of events is $\mu_{90} \sim 0.6$ at 90% C.L. and the corresponding upper bound on the probability for the search in Ref. [7] to have found one or more events would be smaller than $1 - P_{90}(0) \sim 0.45$. Therefore, the results of O3 do not particularly constrain the possibility that SSM170401 could come from a real SSM merger.

The strain in Hanford presents a glitch 14 s before coalescence, as shown in Fig. 25. The search presented in Ref. [7] uses templates starting at 45 Hz. The loudest template, in this case, is only 10 s long and so should be unaffected by the glitch. However, PE was performed with templates starting at 20 Hz, which are roughly 100 s long for the component masses discussed. In this situation the glitch must be removed. Using **BayesWave** [131, 132], we model excess power in the detectors as a sum of sine-Gaussian wavelets. We fit for the glitch and the PSD of the Gaussian noise component simultaneously. We ignore the modelling of the signal due to the extremely low coherent energy per frequency bin deposited in the detectors in the 0.3 s duration of the glitch by such low mass sources, even more when the subtraction is done ~ 14 s before coalescence. The same procedure is used routinely by the LVK collaboration in the main GW catalog [11].

F. Properties of the source of SSM170401

To obtain the properties of the potential source of SSM170401 we interpret the signal as coming from the coalescence of two compact objects. We infer the CBC parameters of the signal using a Bayesian analysis of the data from LIGO Livingston and LIGO Hanford, following the methodology outlined in Appendix B of Ref. [15]. In analysing the data, we fit two different waveform models: **IMRPhenomPv2** [111] and **IMRPhenomXPHM** [112], the latter including higher order modes. Comparing the PE analyses using the two waveforms models, we find that their posterior distributions are consistent with each other, noting that both of them take into account precessing spins.

We use a low-frequency cutoff of 20 Hz in both detectors for the likelihood evaluation and choose uninformative and wide priors. The primary tool used for sampling the posterior distribution is the **LALInference** Markov Chain Monte Carlo implementation as described in [180]. The power spectral density used in the calculations of the likelihood is estimated using **BayesWave** [131, 132]. The study uses the O2 open access data [181] with a sampling frequency of 4096 Hz; however the likelihood is integrated up to 1600 Hz.

The best fit CBC template has ~ 3000 cycles in the detector, allowing us to constrain with relatively high accuracy the source properties of SSM170401 in spite of the low SNR [182]. The estimated parameters are reported in Table VI. The marginalized posterior for the absolute value of the matched filter SNR is $7.98^{+0.62}_{-1.03}$ for **IMRPhenomPv2** and $7.94^{+0.70}_{-1.05}$ for **IMRPhenomXPHM**. The median value of the SNR is lower than that found by the search, which was 8.67. However, these two quantities are not directly comparable. The SNR from the search is obtained by maximizing the

Parameter	IMRPhenomPv2	IMRPhenomXPHM
Signal to Noise Ratio	$7.98^{+0.62}_{-1.03}$	$7.94^{+0.70}_{-1.05}$
Primary mass (M_\odot)	$4.65^{+1.21}_{-2.15}$	$4.71^{+1.57}_{-2.18}$
Secondary mass (M_\odot)	$0.77^{+0.50}_{-0.12}$	$0.76^{+0.50}_{-0.14}$
Primary spin magnitude	$0.32^{+0.47}_{-0.26}$	$0.36^{+0.46}_{-0.30}$
Secondary spin magnitude	$0.48^{+0.46}_{-0.43}$	$0.47^{+0.46}_{-0.42}$
Total mass (M_\odot)	$5.42^{+1.10}_{-1.65}$	$5.47^{+1.43}_{-1.68}$
Mass ratio ($m_2/m_1 \leq 1$)	$0.17^{+0.34}_{-0.05}$	$0.16^{+0.34}_{-0.06}$
χ_{eff} [119, 120]	$-0.06^{+0.17}_{-0.32}$	$-0.05^{+0.22}_{-0.35}$
χ_p [121]	$0.28^{+0.34}_{-0.21}$	$0.33^{+0.33}_{-0.26}$
Luminosity Distance (Mpc)	119^{+82}_{-48}	124^{+82}_{-48}
Redshift	$0.028^{+0.018}_{-0.010}$	$0.028^{+0.017}_{-0.011}$
Ra ($^\circ$)	-2^{+34}_{-35}	-1^{+34}_{-37}
Dec ($^\circ$)	47^{+14}_{-26}	46^{+14}_{-29}
Final mass (M_\odot)	$5.34^{+1.11}_{-1.70}$	$5.40^{+1.45}_{-1.73}$
Final spin	$0.39^{+0.24}_{-0.07}$	$0.42^{+0.22}_{-0.10}$
$P(m_2 < 1 M_\odot)$	85%	84%

TABLE VI: Parameters of the source of SSM170401. All masses are in the source frame. We assume *Planck15* Cosmology [178]. The statistical uncertainty of all the parameters is quantified by the equal-tailed 90% credible intervals about the median of the marginalized one-dimensional posteriors. Right ascension (Ra) and declination (Dec) are measured in the International Celestial Reference System (ICRS) [179].

ranking statistic over a discrete template bank [7, 183–185], while the quoted SNR from the PE is the median value over the samples. Since the ranking statistic and the SNR are closely related, the SNR that is more comparable to that of the search would be the maximum SNR as found by the PE. The values of this maximum PE SNR are 9.09 for IMRPhenomPv2 and 9.18 for IMRPhenomXPHM. These values are slightly larger than that of the search, which is consistent with what would happen if the signal was astrophysical. However, this is also expected in the noise case due to the larger parameter space that allows more flexibility for the PE analysis to fit the data. We also notice the maximum value of the SNR to be larger for IMRPhenomXPHM than for IMRPhenomPv2. In a similar way, this is expected for an astrophysical signal but also for noise, since the waveform includes Higher Order Modes and thus has more flexibility to fit the data.

The source is then compatible with a compact binary system having an unequal mass ratio $q = 0.17^{+0.34}_{-0.05}$ (all uncertainties are quoted at 90% C.L.), a source frame primary mass $m_1 = 4.65^{+1.21}_{-2.15} M_\odot$ and a source frame secondary mass $m_2 = 0.77^{+0.50}_{-0.12} M_\odot$ as shown in Fig. 26. The marginalised posterior distribution for the secondary mass favors a mass lower than $1 M_\odot$ (85% C.L.). Using the IMRPhenomXPHM waveform, we find almost identical results, with a mass lower than $1 M_\odot$ at 84% C.L.

The left panel of Fig. 27 shows the posterior distributions for the magnitude and tilt angle of the individual spins, measured at a reference frequency of 20 Hz. All pixels in this plot have an equal prior probability. The spin of the secondary BH is largely unconstrained, as expected for very unequal masses, while the primary spin shows a preference for small spin magnitudes ($a_1 = 0.32^{+0.47}_{-0.26}$), where the posterior samples with large primary spin tend to have it misaligned with the orbital angular momentum. As can be seen in the right panel of Fig. 27, this leads to a χ_{eff} compatible with zero ($\chi_{\text{eff}} = -0.05^{+0.22}_{-0.35}$) and an uninformative posterior in χ_p ($\chi_p = 0.33^{+0.33}_{-0.26}$).

The luminosity distance and inclination angle θ_{JN} posterior distributions are shown together in the left panel of Fig. 28, since these two quantities are correlated. We find a luminosity distance of $d_L = 119^{+82}_{-48}$ Mpc. We identify a bimodal distribution for θ_{JN} due to the fact that we can not distinguish whether the system is being observed face-on ($\theta_{JN} \sim 0$) or face-away ($\theta_{JN} \sim \pi$), but it being edge-on ($\theta_{JN} \sim \pi/2$) is disfavoured. In the face-on(away) configuration, the effects of precession and higher order modes in the signal are suppressed [96, 186, 187], as is the

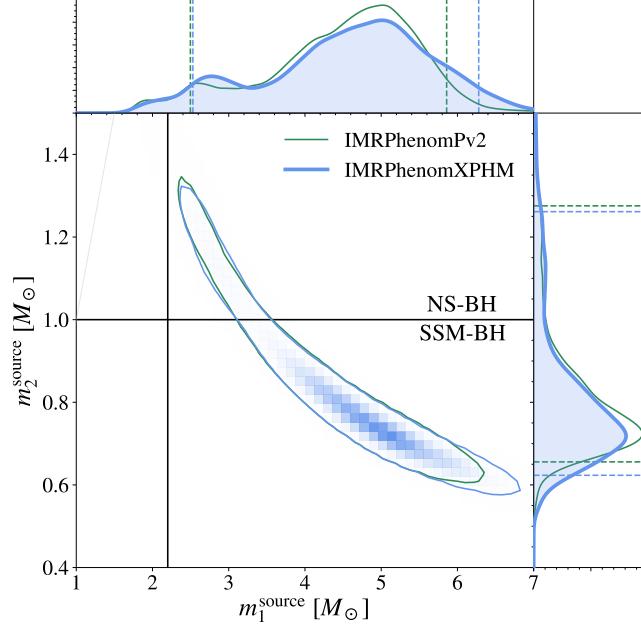


FIG. 26: Posterior distributions for the primary and secondary mass in the source frame. The 90% credible regions are indicated by the solid contour in the joint distribution, and by the dashed vertical and horizontal lines in the marginalized distributions. We also paint a vertical line for the upper bound in the mass of any given Neutron Star and a horizontal one for the subsolar mass threshold.

case here.

In the right panel of Fig. 28, we show the posterior distribution of the location in the sky of the event. This sky map looks abnormal when compared with the typical ones of the events detected exclusively by Hanford and Livingston [11, 15, 141, 142]. When the trigger is seen in two detectors only, most of the information for the sky localisation comes from the time delay between the observation of the signal in both interferometers. This time delay will be given by:

$$\Delta t_{L-H} = \frac{\vec{d}_{H-L} \cdot \hat{n}}{c} = \frac{d_{H-L}}{c} \cos \theta, \quad (165)$$

where \vec{d}_{H-L} is the position vector of Livingston with respect to Hanford and \hat{n} is the direction in the sky of the source. We observe in Eq. (165) that the time delay only constrains the inclination angle with respect to \vec{d}_{H-L} , but leaves the azimuthal angle completely unconstrained. This has as a result a ring-like shape in the sky maps usually observed. However, when the source direction corresponds to $\theta = 0, \pi$, that is, the line joining both detectors, the ring will collapse to have a blob like shape in the sky. As can be seen in the top panel of Figure 29, this is what is happening for the source of SSM170401, since the time delay between Livingston and Hanford is close to the maximum light travel time. In the bottom panel of Figure 29 we also show the posterior distribution of the network antenna pattern \mathcal{F} [188], defined as:

$$\mathcal{F} = \sqrt{\frac{F_{+,H1}^2 + F_{\times,H1}^2 + F_{+,L1}^2 + F_{\times,L1}^2}{2}} \quad (166)$$

where $F_{+(x),D}$ are the $+$ (\times) antenna patterns of detector D . In this plot we observe that the event is coming from a region in the sky where the network antenna pattern is significantly smaller than 1, peaking at $\mathcal{F} \sim 0.5$. This means that the sensitivity in this direction will be half of that of the most sensitive direction ($\mathcal{F} \sim 1$) located on top of the continental US and its antipodes [189]. Since the direction joining both LIGO detectors has smaller sensitivity, the distance up to which LIGO can detect astrophysical signals is also smaller. This leads to a lower expected event rate coming from that direction, which is the reason sky maps like the one of Fig. 28 are uncommon.

Even though the posterior PDFs for the parameters of the source of SSM170401 seem to have converged to a well-defined distribution that differs from the prior, it is known that GW signals can be mimicked by gaussian noise [2] or non-gaussian transients, specially given the relatively low SNR and high FAR.

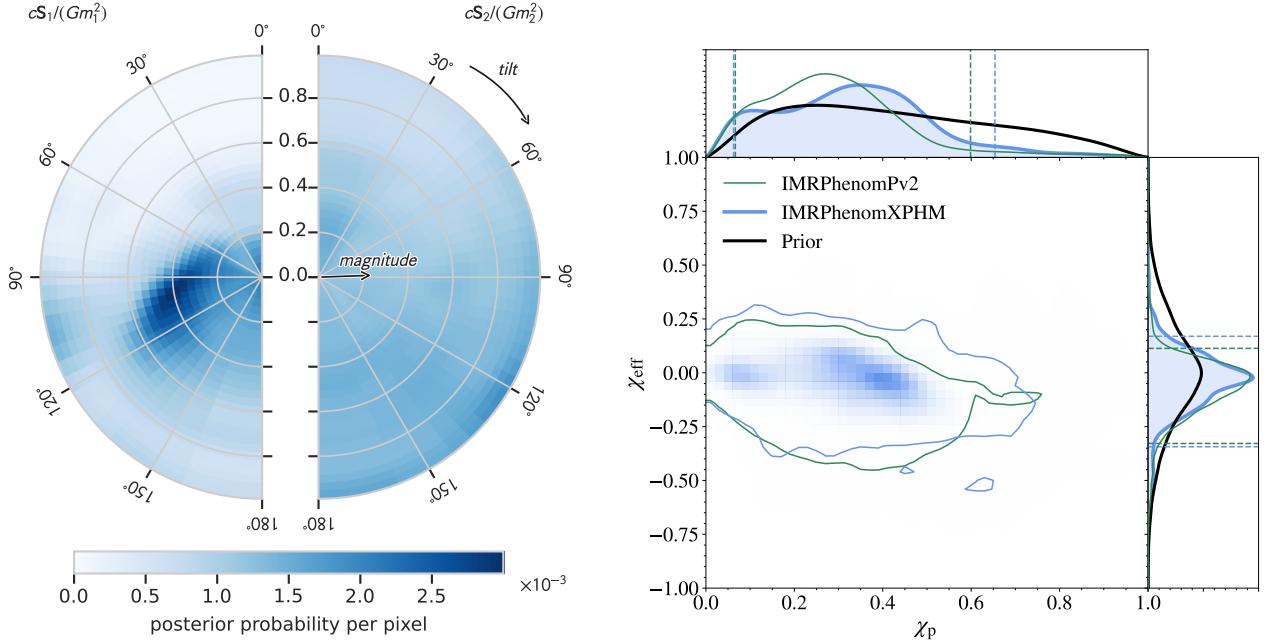


FIG. 27: Left: posterior distribution for the individual spins of the source of SSM170401, according to the `IMRPhenomXPHM` waveform models. The radial coordinate in the plot denotes the dimensionless spin magnitude, while the angle denotes the spin tilt, defined as the angle between the spin and the orbital angular momentum of the binary at a reference frequency of 20 Hz. A tilt of 0° indicates that the spin is aligned with the orbital angular momentum. A nonzero magnitude and a tilt away from 0° and 180° imply a precessing orbital plane. All bins have an equal prior probability. Right: posterior distributions for the effective spin and effective in-plane spin parameters. The black lines in the right panel show the prior distributions for the effective spin parameters. The 90% credible regions are indicated by the solid contour in the joint distribution, and by dashed vertical and horizontal lines in the marginalized distributions. The large density for tilts close to 90° leads to non-zero values for χ_p and low values for χ_{eff} .

1. Coherence Test

To test the compatibility of the SSM170401 with a GW coming from a CBC, we perform the coherence test proposed in Ref. [190]. The idea of this test is to perform Bayesian PE using the data from all the detectors together to calculate the evidence Z_{coh} for a coherent CBC signal and compare this with the evidence Z_{inc} for incoherent CBC signals. The incoherent evidence is defined as the product of the CBC signal evidences obtained performing PE individually in each detector. These incoherent CBC signals are used to represent noise in the detectors that can be picked up by CBC templates. The coherent versus incoherent hypothesis Bayes factor is then defined as

$$\mathcal{B}_{\text{coh,inc}} = \frac{Z_{\text{coh}}}{Z_{\text{inc}}} = \frac{Z_{\text{coh}}}{\prod_{i=1}^N Z^{(i)}} = \frac{\mathcal{B}_{\text{coh}}}{\prod_{i=1}^N \mathcal{B}^{(i)}}, \quad (167)$$

where we have used that for each interferometer, the Bayes factor of the signal versus noise hypothesis is defined as $\mathcal{B}^{(i)} = Z^{(i)}/Z_{\text{noise}}^{(i)}$, while in the coherent analysis it is defined as $\mathcal{B}_{\text{coh}} = Z_{\text{coh}}/\prod_{i=1}^N Z_{\text{noise}}^{(i)}$. We use the same priors for coherent and single-detector analyses. To get a more reliable estimate of the evidence, in the PE we use nested sampling, particularly the `Dynesty` sampler [125] as implemented in `Bilby` [126]. The computational cost of performing the coherence test will be very large since it requires us to perform three separate PEs with the costly nested sampling. To make this analysis feasible we employ Reduced Order Quadrature (ROQ) methods [104], which greatly speed up the computation time of the likelihood, specially for long signals like SSM170401. The analysis is done using only the `IMRPhenomPv2` waveform, since we have seen that it gives consistent results with `IMRPhenomXPHM`

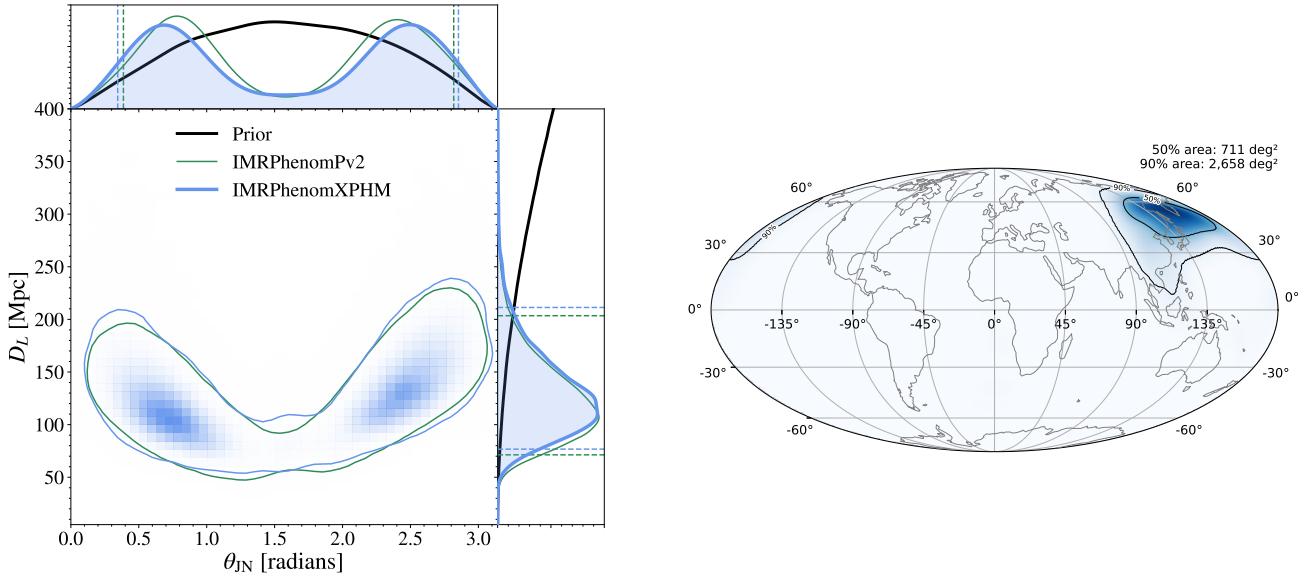


FIG. 28: Left: posterior distributions for the luminosity distance and the inclination angle of the source of SSM170401, according to the `IMRPhenomXPHM` and `IMRPhenomPv2` waveform models. The inclination angle indicates the angle between the line of sight and the total angular momentum of the binary. For nonprecessing binaries, this is equal to the angle between the orbital angular momentum and the line of sight. The solid lines and the central contour denote 90% credible regions. Right: sky position of the event as evaluated from the Greenwich meridian according to the `IMRPhenomXPHM` waveform model.

$\log \mathcal{B}_{H1L1}$	$\log \mathcal{B}_{H1}$	$\log \mathcal{B}_{L1}$	$\log \mathcal{B}_{\text{coh,inc}}$
7.00 ± 0.10	1.56 ± 0.07	0.48 ± 0.06	4.96 ± 0.13

TABLE VII: Natural logarithm of the Bayes factors of the signal versus noise hypotheses obtained from the PE in the data of Hanford-Livingston $\log \mathcal{B}_{H1L1}$, only Hanford ($\log \mathcal{B}_{H1}$), only Livingston ($\log \mathcal{B}_{L1}$) and the natural logarithm of the Bayes factor of the coherent versus incoherent hypothesis
 $\log \mathcal{B}_{\text{coh,inc}} = \log \mathcal{B}_{H1L1} - \log \mathcal{B}_{H1} - \log \mathcal{B}_{L1}$.

while leading to much greater ROQ speedups [105]. We obtain the Bayes factors listed in Table. VII. We find a value for $\log \mathcal{B}_{\text{coh,inc}} = 4.96 \pm 0.13$, strongly favoring the coherent hypothesis over the incoherent hypothesis [191].

This result, however, cannot be used to update the statistical significance of the candidate since we have not run the coherence test over the background triggers of the search. While it is unlikely that a randomly selected noise candidate would give such a large value of $\log \mathcal{B}_{\text{coh,inc}}$ [192], we note that the search ranking promotes candidates with parameters consistent between different detectors [193], thus it may be less surprising that a highly ranked candidate has large $\log \mathcal{B}_{\text{coh,inc}}$. In Fig. 30 we also show the posterior distributions of (\mathcal{M}, q) obtained performing PE in each detector individually and coherently in both of them. We observe that the 2D contours are compatible with each other, having larger areas in the single-detector analyses. This behavior is what is expected if the signal in both detectors were generated by the GWs coming from a single CBC [190].

G. Discussion

To discuss the possible source of SSM170401, assuming it is a CBC, we have divided the (m_1, m_2) in four regions, according to the SSM threshold ($m_2 = 1 M_\odot$) and the maximum allowed mass of a NS ($m_1 = 2.2 M_\odot$) [194–196]. We observe that the full 90% credible region of the posterior lies in the region of $m_1 > 2.2 M_\odot$, excluding the NS origin of the primary component. We find that 16% of the posterior distribution lies in the region of $m_2 > 1 M_\odot$, which would point to a likely NS origin, although a light black hole cannot be excluded. However, the most probable region, representing 84% of the posterior, would imply a mass of the secondary component below $1 M_\odot$. It is thus

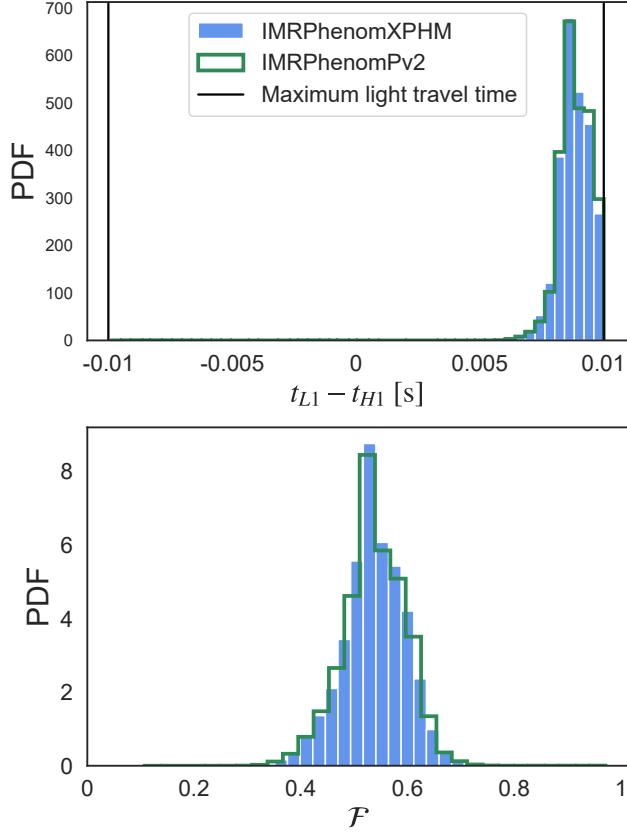


FIG. 29: Upper panel: posterior distribution of the time delay between the arrival of the signal in LIGO Livingston and LIGO Hanford. Lower panel: posterior distribution of the network antenna pattern \mathcal{F} , computed using Eq. (166).

interesting to explore what could be the origin and nature of a possible SSM object.

The first possibility to consider for an SSM component would be a neutron star. Neutron stars have relatively well-determined masses from observations of binary systems, including pulsars or X-ray binaries involving an accreting neutron star from a companion. Their measured masses are above $1.2 M_{\odot}$ [197], further confirmed by the observation of GW170817 [198]. However, there is a recent claim [199] for a neutron star of mass $0.77^{+0.20}_{-0.17} M_{\odot}$, although it has been argued [200] that such a small mass for a neutron star probably requires a strange QCD equation of state. Therefore, the neutron star interpretation of a possible SSM component cannot be excluded, although it is disfavored by the bulk of observational data.

Another possibility is PBHs formed by the gravitational collapse of large inhomogeneities in the early Universe which are already considered as a possible explanation of LVK GW detections, see e.g. [1, 23–27, 41, 44, 201–205]. Depending on the model, they may explain anything from a tiny fraction of Dark Matter to its entirety. PBHs have been the main motivation to conduct searches of SSM black holes in the LVK data [7, 172, 173, 177, 206–208], in particular, the extended subsolar search with low-mass ratios in O2 which reported SSM170401 as a possible candidate [7]. If some of the observed binary coalescences are indeed due to PBHs, they must have a relatively extended mass distribution that would have been imprinted by the thermal history of the Universe [26, 209]. This would lead to a peak in the mass distribution around a solar mass which is naturally produced at the QCD transition [26, 209–214], and the source of SSM170401 could be an example of a subsolar PBH around the QCD-induced peak. The spin posterior is quite broad and the spin is compatible with zero, although a slight preference for a primary spin around 0.3 is observed. In this case, the non-zero but relatively low spin of the primary component may have been acquired by matter accretion, previous mergers or hyperbolic encounters [1, 33, 215].

Alternatively, in scenarios with complex and dissipative particle Dark Matter, SSM black holes could form through the cooling and gravitational collapse of Dark Matter halos [157]. This model was constrained by the LVK data in [12, 173, 216]. Furthermore, it could be that the secondary component of the source of SSM170401 is a boson star,

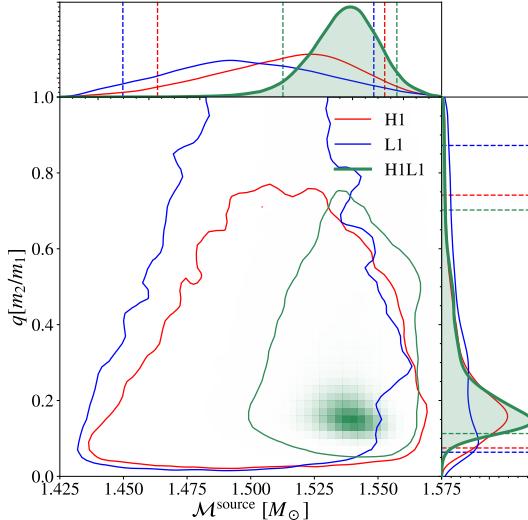


FIG. 30: Posterior distributions for the mass ratio and the source frame chirp mass (\mathcal{M}) for the PE analysis performed with the waveform `IMRPhenomPv2` when considering both detectors together (H1L1) and individually.

a hypothetical horizonless compact object formed by an ultralight bosonic field. If the mass of the bosonic particle is larger than $10^{-10} \text{ eV}/c^2$, the boson star can have subsolar mass [217]. Whether a merger with an SSM boson star component could produce a signal similar to the SSM170401 trigger, though, remains to be investigated.

Finally, we note that the primary component mass of the hypothetical source of SSM170401 would preferably lie in the hypothesized low mass gap between 2.5 and $5M_{\odot}$ (64% C.L.). However, this is not unique, since other candidates with components possibly in this lower mass gap have been observed, namely GW190814 [18] and GW200210_092254 [11].

H. Conclusions

We have performed an in-depth investigation of the most significant double-detector candidate reported in [7] in an SSM search over O2 data. We have removed a prominent blip glitch 14 s before coalescence in the data and estimated the parameters of the possible CBC source using a low frequency cutoff of 20 Hz. Parameter Estimation runs were performed using `LALInference` with two different waveforms `IMRPhenomPv2` and `IMRPhenomXPHM`, where the latter includes effects from higher order modes. The source parameters obtained by both PE runs show good agreement with each other and with the parameters of the template that triggered the search. We find a median network SNR of $7.94^{+0.70}_{-1.05}$ (90% credible interval), which is lower than the SNR of 8.6 obtained in the search. However, the search SNR is more closely related to the maximum SNR, which we find to be higher in the PE, where it reaches values of 9.09 for `IMRPhenomPv2` and 9.18 for `IMRPhenomXPHM`. The secondary mass is $m_2 = 0.76^{+0.50}_{-0.14} M_{\odot}$ (90% credible interval), with 84% confidence of being below one solar mass. For the location in the sky posterior, we find an atypical distribution when compared with the usual Hanford-Livingston events detected thus far, which can be explained if it were a GW coming from the direction joining the two LIGO interferometers.

The compatibility of SSM170401 with a CBC origin has been further tested by performing the signal coherence tests of Ref [190], obtaining a log Bayes factor of 4.96 ± 0.13 for the coherent vs incoherent hypothesis. Furthermore, we observe that the (\mathcal{M}, q) posteriors of each independent IFO converge to mutually compatible contours. These tests provide significant support in favor of a coherent signal, which generally is not expected if it were generated by noise fluctuations [192]. We also checked that the O3 limits on the SSM merger rate [12] do not put a significant constraint on the probability of this candidate being astrophysical ($P_{90} \lesssim 0.45$). Therefore, we do not find compelling arguments against a possible CBC origin of SSM170401.

Finally, even if most of the m_2 posterior support is in the SSM region, there is still a 16% probability of m_2 being over $1 M_{\odot}$, which does not allow us to convincingly exclude a NS origin. Candidates with a higher SNR would have better measurements on their parameters, allowing for more confident discrimination between sub- and super- solar masses [218]. Therefore, the data from future planned LIGO-Virgo-Kagra runs with improved sensitivity [143], O4 and O5, offer a great opportunity for discovering CBC mergers with SSM components, if they are out there in the Cosmos.

VII Primordial Black Hole clusters, phenomenology & implications

A. Journal & publication details

Submitted to Physics of the Dark Universe for consideration. [8]

B. Authors

The authors in the same order as presented to the journal are

José Francisco Nuño Siles^{a21}

Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco 28049 Madrid, Spain

Juan García-Bellido^{a22}

Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco 28049 Madrid, Spain

C. Abstract

We present direct N-body simulations of black-hole-only clusters with up to $2 \cdot 10^4$ compact objects, zero natal spin and no primordial binaries as predicted by various primordial black hole (PBH) Dark Matter models. The clusters' evolution is computed using NBODY6++GPU, including the effects of the tidal field of the galaxy, the kicks of black hole mergers and orbit-averaged energy loss by gravitational radiation of binaries. We investigate clusters with four initial mass distributions, three of which attempt to model a generic PBH scenario using a lognormal mass distribution and a fourth one that can be directly linked to a monochromatic PBH scenario when accretion is considered. More specifically, we dive into the clusters' internal dynamics, describing their expansion and evaporation, along with the resultant binary black hole mergers. We also compare several simulations with and without black hole merger kicks and find modelling implications for the probability of hierarchical mergers.

D. Introduction

Black holes (BHs) have long captured the attention of astrophysicists due to their elusive and mysterious nature. Recent progress in observational astronomy has revolutionized our ability to study them across various scales, from detailed images of supermassive BHs at the centres of galaxies [219] to the detection of microlensing events caused by BHs passing in front of background sources [220]. More relevant to this work, the detection of gravitational waves emitted by merging binary black holes (BBHs) in Earth-based interferometers [11] has improved our understanding of their formation mechanisms and population statistics while, at the same time, giving rise to new fundamental questions [16, 18, 19, 203], like the nature of Dark Matter [21, 23, 24], and the thermal evolution of the early Universe [26].

In this context, the investigation of black hole-only clusters takes on particular significance as a possible alternative origin capable of shedding light onto some observations. Theoretical considerations, supported by numerical simulations, have demonstrated the plausibility of forming such clusters. For instance, BHs created during the radiation-dominated era are naturally clustered [221, 222] if they derive from large non-Gaussian tails [223–225]. Other phenomena such as the appearance of closed domain walls [226] and their collapse [227] can lead to their clustering too. We should also mention the possibility of the existence of black-hole-only clusters originating via stellar evolution [228] for which our analysis and conclusions remain valid with the appropriate time delay.

Despite the progress in the field, our comprehension of the dynamics of black hole-only clusters remains limited [229]. The gravitational interactions between thousands of individual BHs together with the influence of the surrounding environment, present hard challenges for any attempt at theoretical modelling given the vast range of scales needed to be considered. However, these complexities also offer a unique opportunity for N-body simulations to excel.

N-body simulations have proven instrumental in understanding a diverse range of complex astrophysical phenomena. Using numerical methods to integrate the gravitational equations of motion for all of the individual particles, we obtain

²¹ jose.nunno@uam.es

²² juan.garcia@uam.es

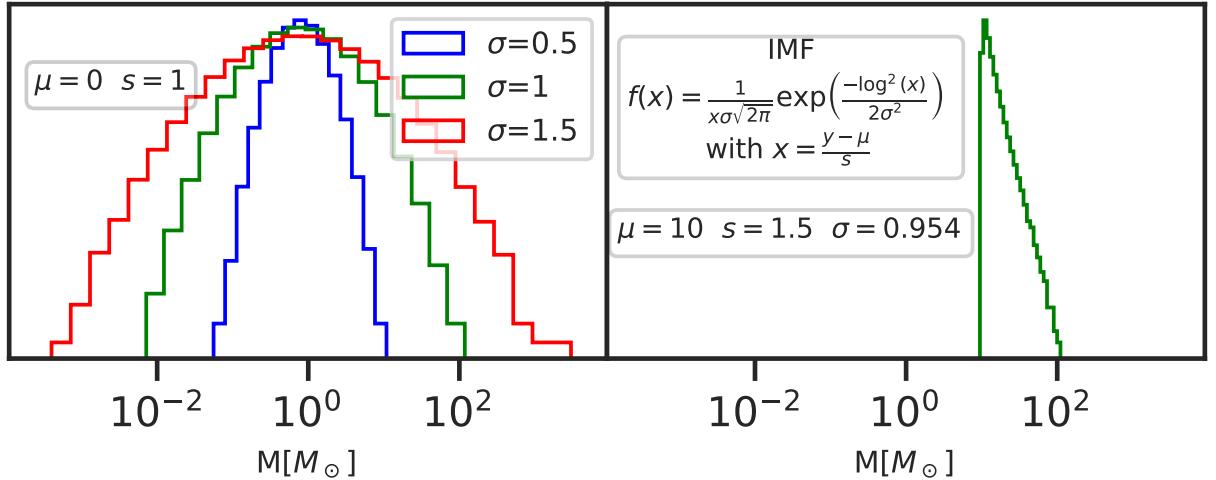


FIG. 31: Initial mass function of the clusters. Those based on a typical PBH mass function (σ_i) are in the left panel, while the monochromatic with accretion IMF ($M\&A$) is in the right one. We also write the general formula for all the IMFs and the values for the specific parameters to reproduce our results.

a detailed description of the dynamics governing the evolution of such complex systems [230]. A complete introduction to the theoretical foundations of the field of N-body simulations can be found in Ref. [231] and a shorter, but less recent one, in [232].

In this paper, we present direct N-body simulations of black hole-only dense clusters, focusing on systems with up to $2 \cdot 10^4$ compact objects. Our simulations incorporate physical effects such as the tidal field of a given host galaxy, black hole relativistic merger kicks, and orbit-averaged energy loss due to gravitational radiation from binary systems. By utilizing the NBODY6++GPU code [10], we accurately model the evolution of all the BHs for a Hubble time. The text is structured as follows, in the first section VII E we describe the initial conditions in detail, in the second section we analyze the cluster dynamics, trying to understand why the clusters dissolve and expand. The following section is focused on binary black hole mergers and the parameters characterizing these events. In the last section, we summarize our findings and conclude the work.

The main body of this work consists of the description of the physical variables of interest regarding such hypothetical celestial bodies and their behaviour in a Hubble time. We find promising hints of results that could be linked with current observations but would need more in-depth research before satisfying conclusions could be drawn. While there are possibly many more findings in the data, we don't embark on such extensive endeavour and treat this work as a comprehensive proof of concept for future works to be carried out.

E. Methodology and initial mass function

To evolve our BHs-only clusters, we made use of the direct numerical integrator NBODY6++GPU. The models we evolve consist of N black holes, where N ranges from $\mathcal{O}(10^3)$ to $\mathcal{O}(2 \cdot 10^4)$, drawn randomly from 4 different initial mass functions plotted in Fig 31. The exact numbers can be found in Table VIII. The first one, a log-normal distribution with $\mu = 10$, $s = 1.5$ and $\sigma = 0.954$ is depicted in the right panel and models, for instance, a monochromatic mass function when accretion is considered and thus the masses of the BHs can increase. The other three represent an approximation to the PBH Thermal model [26, 233] composed of a wide mass function with three widths for the log-normal $\sigma = \{0.5, 1, 1.5\}$. See [234] for a recent review of the model. We will refer to them as $\{M\&A, \sigma_{0.5}, \sigma_1, \sigma_{1.5}\}$ respectively. We set the initial fraction of binaries as well as the individual spins of the BHs to zero according to theoretical expectations. The BHs are then spatially distributed in such a way that they follow a Plummer distribution [235] with a Plummer Radius $r_p = 10$ pc. Based on the assumption that they could reside in the halo of a galaxy similar to our own, these clusters are themselves immersed in a central gravitational potential with orbital radius $R_c = 34$ kpc and central mass $M = 4.37 \times 10^{10} M_\odot$ throughout the entire evolution. This is just a point mass approximation which leads to a circular movement of period $T = 2.81$ Gyr. In Table VIII we explicitly write key statistics for the individual clusters.

For these models to be of any real physical interest Primordial Black Holes (PBHs) should have been created in

ID M&A	$M_{\text{total}}[M_{\odot}]$	$M_{\text{max}}[M_{\odot}]$	$M_{\text{max}}[M_{\odot}]$	$R_{\text{HM}}[\text{pc}]$	$R_{\text{HM}}[\text{pc}]$	ID $\sigma_{0.5}$	$M_{\text{total}}[M_{\odot}]$	$M_{\text{max}}[M_{\odot}]$	$M_{\text{max}}[M_{\odot}]$	$R_{\text{HM}}[\text{pc}]$	$R_{\text{HM}}[\text{pc}]$
	$t = 0$	$t = 0$	$t = T_U$	$t = 0$	$t = T_U$		$t = 0$	$t = 0$	$t = T_U$	$t = 0$	$t = T_U$
1295	16140	46.54	46.54	2.20	27.04	1520	1389	5.03	5.03	9.72	12.25
2570	32166	66.10	97.80	2.38	27.17	3345	3007	4.50	5.08	3.23	15.55
4046	50088	63.64	91.98	4.27	26.74	5480	5009	6.67	6.67	0.77	16.21
5199	64280	69.98	74.80	2.20	28.39	7678	7008	5.69	7.52	1.92	15.15
8077	100116	78.27	78.27	12.53	29.50	9937	9001	5.48	5.48	1.25	15.85
8922	110196	76.15	81.52	13.31	32.01	12201	11001	5.00	7.77	2.90	15.29
10372	128337	104.26	104.26	6.20	27.54	14366	13005	5.42	8.22	2.65	16.04
10535	130198	68.50	74.35	9.72	28.20	16428	14912	7.72	7.72	0.48	16.04
16159	200392	57.75	113.59	4.44	27.24	18776	17007	5.40	6.51	2.29	15.00
20738	256346	134.10	145.70	1.37	26.65	20866	18918	5.31	7.36	3.91	15.35
ID σ_1	$M_{\text{total}}[M_{\odot}]$	$M_{\text{max}}[M_{\odot}]$	$M_{\text{max}}[M_{\odot}]$	$R_{\text{HM}}[\text{pc}]$	$R_{\text{HM}}[\text{pc}]$	ID $\sigma_{1.5}$	$M_{\text{total}}[M_{\odot}]$	$M_{\text{max}}[M_{\odot}]$	$M_{\text{max}}[M_{\odot}]$	$R_{\text{HM}}[\text{pc}]$	$R_{\text{HM}}[\text{pc}]$
$t = 0$	$t = 0$	$t = T_U$	$t = 0$	$t = T_U$	$t = 0$	$t = 0$	$t = 0$	$t = 0$	$t = T_U$	$t = 0$	$t = T_U$
1505	2003	32.07	32.07	1.93	27.07	1220	3020	63.67	63.67	3.37	58.92
3090	4004	18.56	18.56	7.52	22.21	3423	8010	105.74	105.74	7.15	41.00
5288	7007	34.14	34.14	7.60	21.97	5258	13014	541.18	541.18	1.95	40.15
7507	10001	25.82	34.29	3.95	24.77	8025	20021	157.23	157.23	0.41	41.62
10663	14000	46.77	46.77	6.64	26.96	10011	24187	179.83	280.36	1.45	60.33
12834	17004	32.81	58.06	4.41	23.91	12409	30007	114.51	114.51	0.07	55.98
15136	20000	58.31	58.31	0.47	25.47	14691	38027	731.69	1004.13	1.40	37.81
17554	23008	45.68	64.90	0.82	25.62	17182	43013	554.27	554.27	5.59	46.36
20590	27008	51.60	76.41	4.70	26.64	20261	49021	657.47	784.02	3.52	41.41

TABLE VIII: In this table we show the initial conditions for the clusters studied in the paper. The ID also corresponds to the initial number of BHs. We also provide the total mass of the cluster, the maximum mass of any single BH in the cluster at $t=0$ as well as after we have evolved it for the age of the Universe and the radius that encircles half of the total mass of the cluster at both times. The reasons for the maximum mass differing at the beginning and end of the evolution are due to mergers retained in the clusters or the escape of the most massive body.

large numbers at the beginning of the Universe and close enough to each other so that they would have formed clusters so dense as to overcome the Hubble expansion. Several mechanisms could lead to such scenarios [227, 236]. The idea of them making up a significant fraction of the Dark Matter would add relevance to the study but is not essential. To understand the implications of such a scenario, we can estimate the number of clusters needed to get the observed DM mass. Given the clusters' total masses range is $\mathcal{O}(10^3 - 10^5) M_{\odot}$, we would need around $\mathcal{O}(10^6 - 10^9)$ clusters to conform the entirety of the dark matter in our galaxy, where we have used a total DM mass of $1.2 \times 10^{12} M_{\odot}$ [237]. Given the nature of the clusters' components, direct detection is intrinsically difficult; therefore, our approach will focus on identifying potential footprints they may leave behind.

F. Dynamics

The dynamics governing self-gravitating systems are extremely non-linear, thus, trying to infer general qualitative principles is the most we aim to do in this paper. In this section, we will try to shed some light on the understanding of the passage of time for these systems. We will do so by following various global physical variables that would serve as a proxy for the whole cluster.

From the point of view of the spatial extension of the clusters with time, we can study global quantities such as the Lagrangian radii and the core radii as defined in [231]. We plot in the left panel of Fig 33 the Lagrangian radii for 50% of the mass of the different clusters. This quantity is also called the half-mass radius. The tendency is for the half-mass radius to grow with time, with a steeper slope at the beginning of the simulation and flattening towards

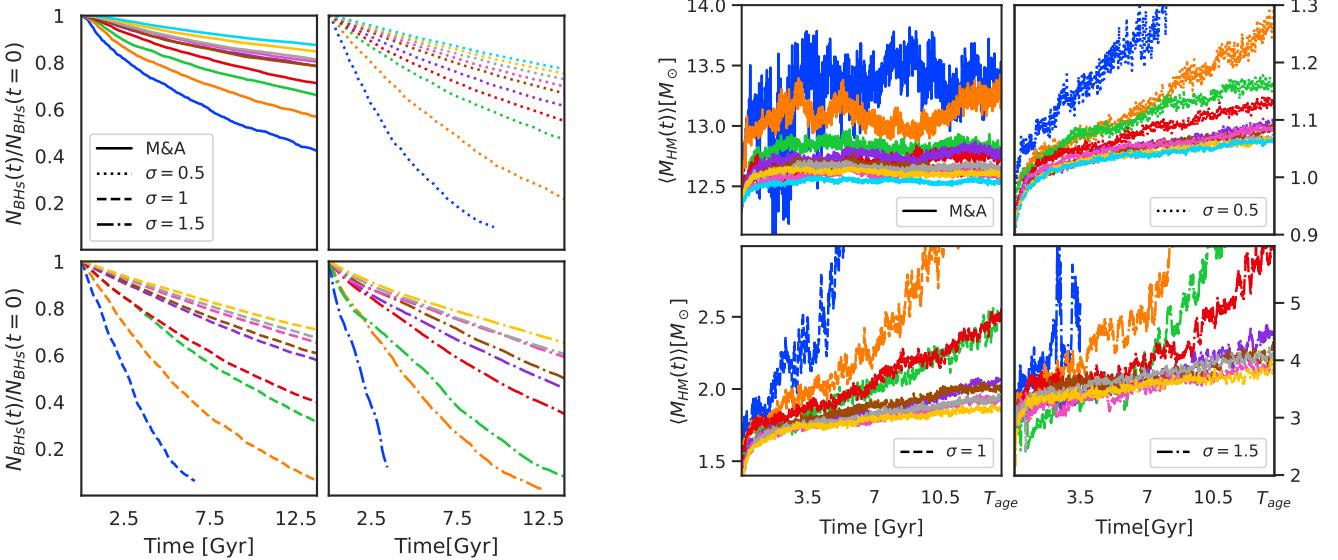


FIG. 32: In the left panel of this figure we plot the number of BHs for the various clusters as a function of time normalized to the initial number of BHs. The dark blue lines represent the lightest cluster for a given IMF increasing upwards with no overlapping between the lines in the last ~ 7 Gyrs of evolution. In the right panel, we plot the average mass of BHs for the various clusters inside the half-mass radius as a function of time.

the end, albeit always increasing. In the case of the core radius, which is displayed for all the clusters in the right panel of Fig 33, we observe that it remains constant, with a slight upward trend throughout the evolution. This translates into a faster expansion of the outer layers, converging towards an almost non-expanding core. Additionally, it is interesting to note that most clusters exhibit a similar core radius regardless of the initial number of BHs.

Now we analyse the stability of these clusters. When carrying out the numerical evolution, it is evident from the beginning that BHs are constantly escaping the combined gravitational influence of the cluster and the host galaxy so that the total number being evolved by the code gets reduced with time. In NBODY6++GPU, BHs are removed from the cluster evolution once they have reached twice the tidal radius [238] as calculated using the mass of the cluster at the specific time. This will be our definition of an escaper (single or binary).

These BHs acquire the needed escape velocity after one or several close encounters with other BHs. This process happens so often that we can consider all clusters to be metastable, that is, most of them don't release enough BHs to dissolve in the age of the Universe, but they will eventually do so. In the left panel of Fig 32, we show the number of BHs left in the clusters as a function of time, normalized to the initial number of BHs of each cluster so that the various runs are comparable. We see that some of the lightest clusters of the σ_i type dissolve completely in the age of the Universe and, in general, the larger the σ for the IMF of the cluster, the shorter the expected lifespan. The most stable set of initial conditions seems to be that of *M&A*. To explain this, we have to understand the main mechanism behind the clusters' dissolution and the difference in their behaviour as a function of the IMF. We can imagine a simplistic scenario in which we just consider 2-body interactions as the main driver of the evaporation of the clusters. This is probably a good approximation as 3+ body encounters are suppressed based on their scattering cross-section. Now, when considering the different IMFs, we see that the main difference across the various σ is the mass ratio distribution for random pairs. The larger the σ is, the more extreme mass ratios exist, and the expected mass ratio goes further away from 1. In the *M&A* case, most of the pairs concentrate close to equal mass ratios. This led us to the conclusion that the more extreme the mass ratios are, the more slingshots of the lighter BHs occur which is translated into more BH escapers. This is also in agreement with the fact that the lighter BHs escape earlier, leaving the heavier ones in the cluster. This can directly be appreciated in the right panel of Fig 32 where we plot the average mass of the BHs inside the half-mass radius. This quantity is strictly increasing in the σ_i cases with a slope that steepens in the last stages of the cluster life due to the wider range of masses present in those clusters. Given the almost monochromatic nature of the *M&A* IMFs, the average mass tends to stay constant with a variance that decreases inversely with the initial BH number.

If a sufficient number of these clusters existed in our Universe, single black hole escapers could potentially be detected from Earth through microlensing or via stellar disruption events. Those that escape as binaries could also be detected in the final stages of the inspiral, as well as the merger and the ringdown in Earth-based gravitational

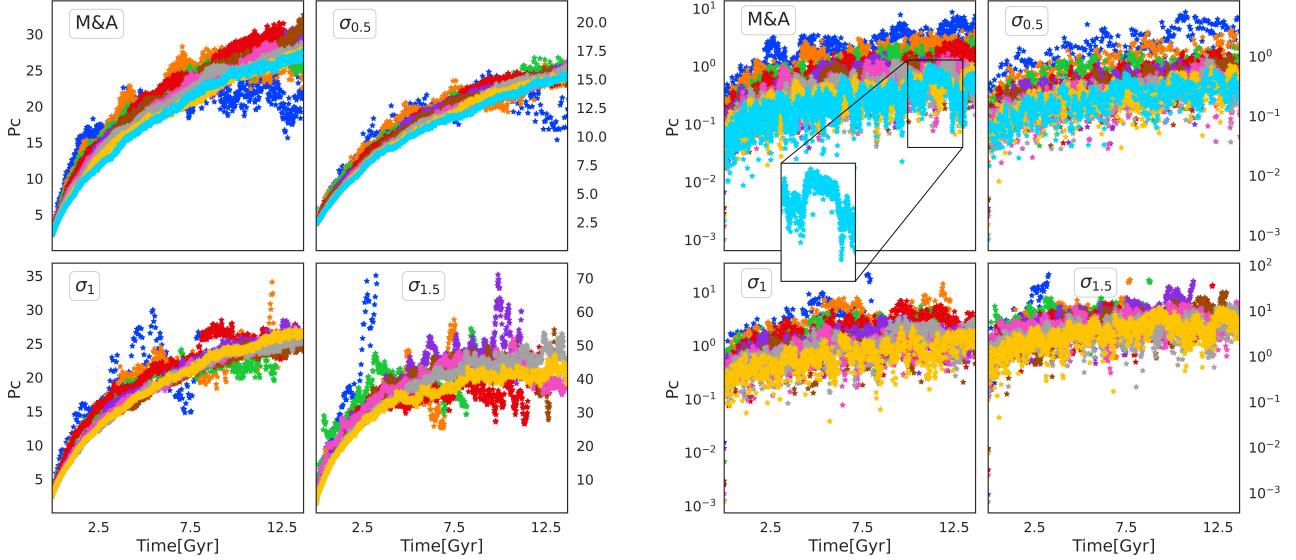


FIG. 33: In this figure we plot the Lagrangian radii for 50% of the mass and the core radii for the various clusters as defined in [231]. Most lines overlap cluster-wise, so we don't add a legend to distinguish them only pointing out that the lighter clusters are also those showing the largest variance. For the core radii, we include a zooming window to what looks like possible oscillations for the heaviest of the *M&A* clusters, although it could also be just shot noise in the calculation.

waves interferometers. We will call these binary escapers, off-cluster mergers if the initial conditions regarding their orbits' parameters once they have escaped from the gravitational pull of the cluster are such that they coalesce in the age of the Universe. To calculate the coalescence time, we assume orbit shrinking via gravitational wave emission and use the formula as derived in [239] in post-processing.

Starting with the single BH escapers, we first acknowledge the fact that the distributions of masses are very similar to the IMF of the clusters with a slight skew towards the lighter BHs due to usual gravitational mass segregation. This means that enough observations of these rogue compact objects would paint a good picture of the progenitor clusters' distributions.

Regarding binary escapers, of which off-cluster mergers represent a subset, we find a correlation in the distribution of semi-major axes with the initial BH number. The larger the initial number of BHs in the cluster, the tighter the binaries that escape (smaller semi-major axis) and vice-versa. This fact can be visually spotted in figure 34. This is expected as the binaries that may survive in a less dense cluster based on the rate of interactions, may not do so in denser environments. In other words, the binding energy needed for a binary to prevail and escape the cluster grows with the density. Regarding their eccentricities (e_0), we see an excess close to maximum eccentricity. This is just the result of the many interactions the binaries need to endure before escaping the cluster.

Results also show that the absolute number of binary escapers increases with the initial BH number within the same type of IMF. Going even further, it also seems to correlate with the stability of the clusters surveyed. The more stable the cluster type, the more binaries escape the gravitational influence with the order from more to less stable being $\{M\&A, \sigma_{0.5}, \sigma_1, \sigma_{1.5}\}$.

G. BBH mergers

During the numerical evolution of the clusters, the centres of two (or more) BHs might get close enough due to the gravitational interactions that their event horizons would merge and the code would start evolving them as a single entity, thus declaring a coalescence. These events occur inside the clusters [240, 241], mainly induced by interactions where 3 or more compact objects are involved [242], that is, binary-single interactions or binary-binary interactions. We refer to them as in-cluster mergers. The second type of merger, as described in Section VII F, occurs when two

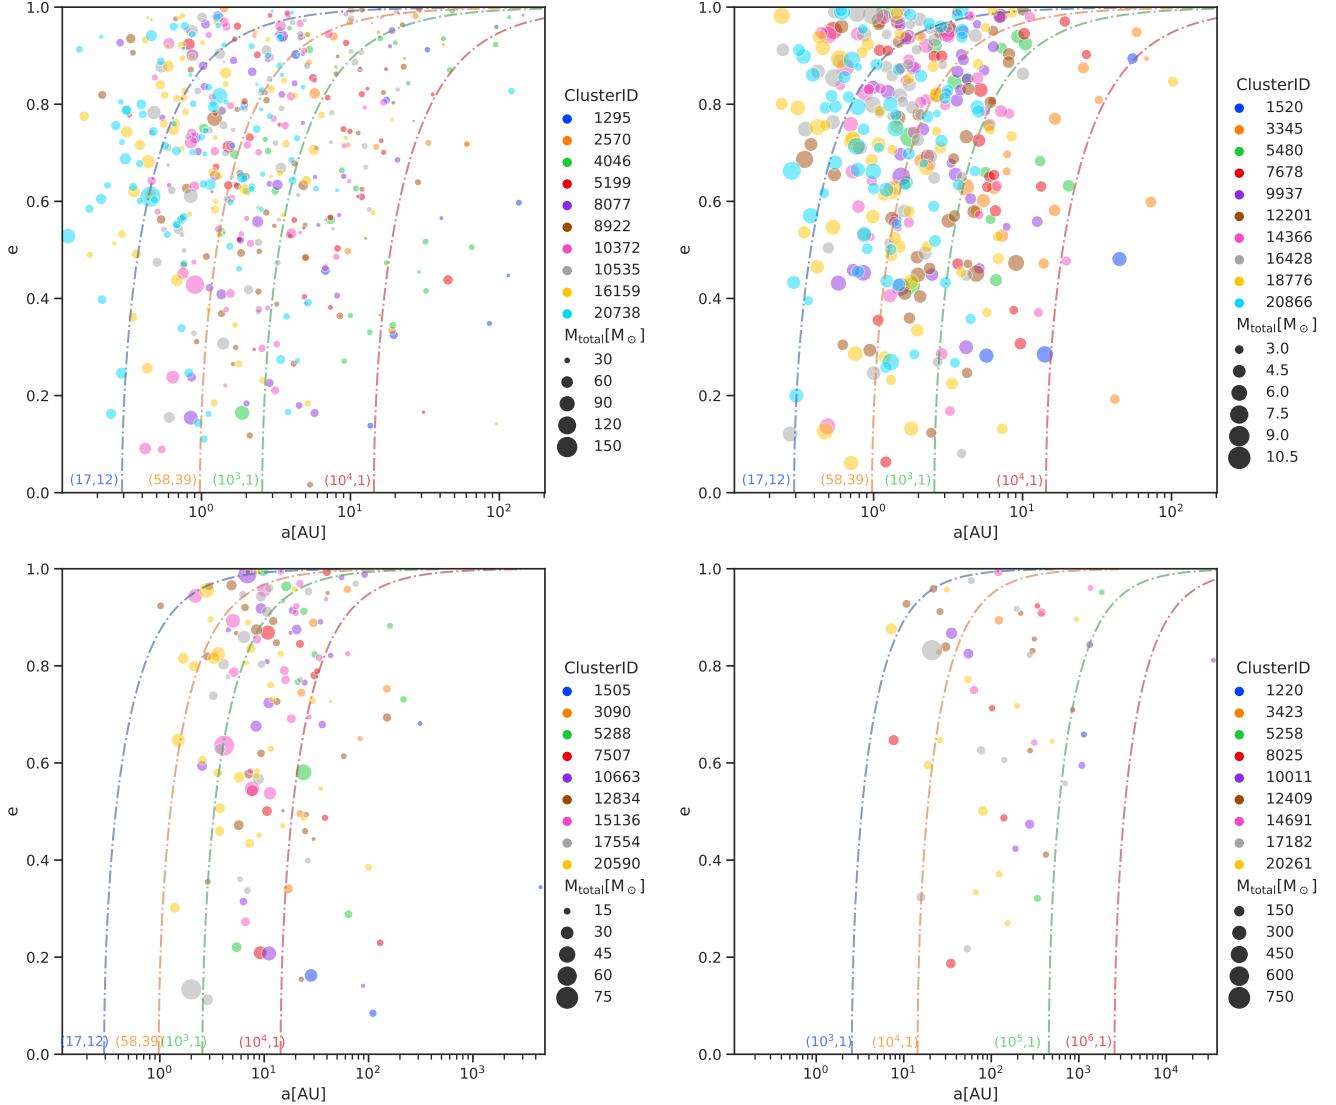


FIG. 34: In this figure we plot the initial semimajor axis and eccentricity (a_0, e_0) of all the BHs that escape the gravitational pull of the clusters within the age of the Universe. These are the binaries logged by the Nbody code as escapers. In order from left to right and top to bottom, we are showing the clusters $\{M \& A, \sigma_{0.5}, \sigma_1, \sigma_{1.5}\}$. The dashed lines represent the values of (a_0, e_0) for which $\tau_{\text{merger}}(m_1, m_2, a_0, e_0) = T_{\text{Hubble}}$ with $(m_1, m_2)[M_\odot]$ being the pairs of number in the base of the plot beside each line. These calculations, however, does not take into account the time delay from the beginning of the simulation until the binary escapes the cluster which can be quite substantial.

The width of the dots represents the total mass of the binary defined as $M_{\text{Total}} = m_1 + m_2$.

BHs leave the cluster as a binary bound by their mutual gravitational interaction. Due to gravitational radiation and under the assumption that nothing else ever interacts with these pairs, they will inevitably end up merging. If the time of merger is less than a Hubble time, we refer to them as off-cluster mergers.

H. Merger count

Understanding the difference between the various initial conditions can be very useful in assessing the underlying mechanisms of the mergers. In the right panel of Fig 35 we show four histograms with the number of mergers per cluster. Discussing first the in-cluster mergers, we find that among the σ_i clusters there is a trend in the number of mergers inverse to the width (σ) of the distribution. In the $\sigma_{0.5}$ and σ_1 cases, a maximum in the merger count seems

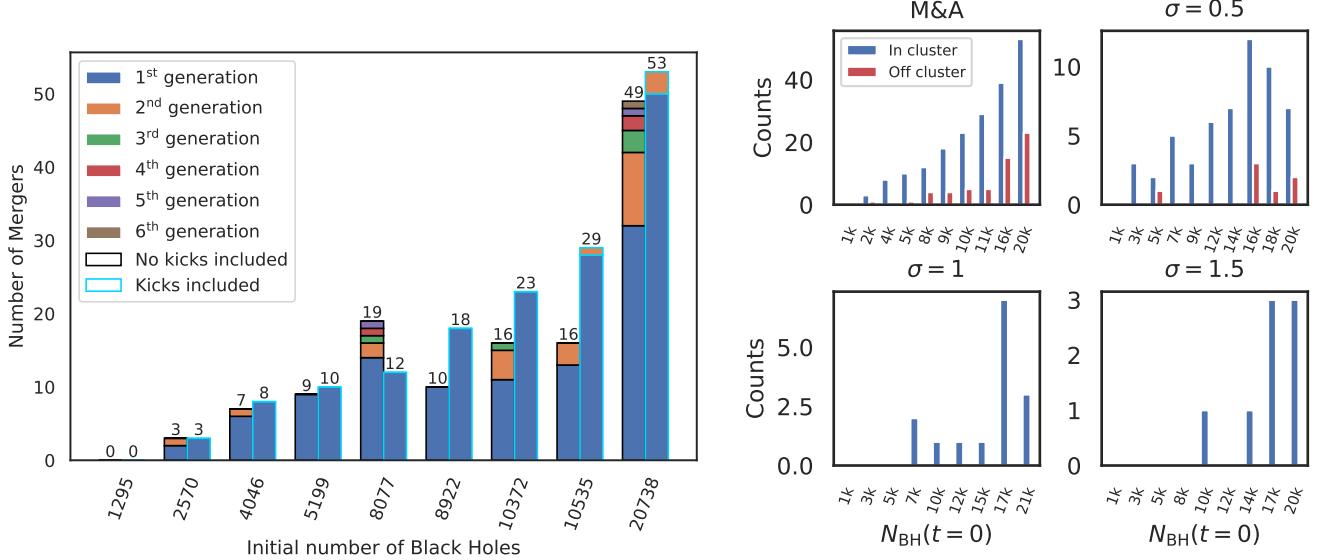


FIG. 35: In the left panel we show a histogram of the number of in-cluster mergers for the $M\&A$ cluster type. The bars with the blue border represent the simulations used throughout the paper, where BH merger kicks are taken into account. The bars with the black border are those same simulations when the kicks are not considered. In the right panel, we plot a histogram of the number of mergers per cluster type. We differentiate, in blue the in-cluster and in red the off-cluster mergers.

to arise for $N \sim \mathcal{O}(1.7 \cdot 10^4)$, although it could be just a statistical anomaly. The total number count for comparable initial N_{BH} lies below that of the $M\&A$ cluster type. This can be explained based on the lower stability of the clusters themselves due to the larger mass range present. The large number of mergers for the $M\&A$ cluster type has as a result also the existence of second-generation BHs, which we discuss in sec VIII. For this last case, there is a clear upward-sloping tendency in the number count both for in and off-cluster mergers.

We observe a lower number count in the off-cluster merger statistics. This is expected as the delay between the start of the simulations and the ejection of the binary from the clusters can be significant, thus making the expected merger time for the binary escapers way longer than Hubble time. This comes together with the fact that gravitational wave emission is a slow mechanism for shrinking the orbit of the binary in comparison with the interactions governing the in-cluster mergers. Comparing cluster types, we find again fewer counts for the σ_i with no off-cluster mergers happening in any $\sigma_{1,1.5}$ clusters. This means that no binary escaper ever merged in the age of the Universe. One of the reasons for this is that the distribution in the semi-major axis of these two cluster types lies at higher values always, which is generally equivalent to larger merger times. Another reason is that the average delay time before escaping is larger for $\sigma_{1,1.5}$ than for the other two clusters. It can also be explained by the fact that there are fewer binary escapers for those two clusters, as can be seen in figure 34.

I. Hierarchical mergers

The product of any merger is another BH with the mass equal to the sum of the masses and spins calculated using [243]. These remnants may or may not stay in the cluster as the velocity imprinted on them [244–246] can be very high. Those which not leave the cluster are subject to the possibility of being involved in another merger. To understand how common this phenomenon is, we show a histogram in the left panel of Fig 35 with the total number of BH mergers identified in the lifetime of the clusters for the $M\&A$ cluster type. There are two different bars per Cluster ID. The ones whose border is black represent simulations in which BHs' natal kicks are not taken into account while bars with a blue border are the normal simulations that do include such a physical phenomenon and thus represent a more complete description of reality. The initial conditions are the same for both cases. We find that when kicks are not present, the probability that a remnant stays in the cluster is high as the velocity after the merger is obtained based on linear momentum conservation. However, when kicks are included, relativistic effects imprint a velocity in the remnant that is usually much larger than the escape velocity of the cluster. This has as a consequence

a sharp drop in the probability of hierarchical mergers. We can see this reflected in the fact that hierarchical mergers of up to 6th generation can happen in most of the clusters if kicks are ignored but it's only for N larger than 10535 BHs that we find hierarchical mergers when kicks are considered. What is more, anything beyond 2nd generation is not on the cards, as the kick velocity is even larger in the case where progenitor spin is non-zero [247]. This has broader consequences in the distribution of other binary parameters, such as the final spin. As can be seen in Fig. 38 a limited number of BH remnants have $a_f > 0.7$ and they are all produced via second-generation mergers. A variation in the statistics of such products would lead to critical differences in the detected population of BHs via for example gravitational wave interferometers.

This result is of great importance as it highlights the significant impact that an occasionally overlooked physical effect can have on observational probes. Any simulation lacking the modelling of BH merger kicks is doomed to overestimate the hierarchical merger rate and the skew to larger remnant spins.

J. Merger Rate

Cluster ID <i>M&A</i>	Redshift last merger	Rate z=0 (events/yr/Gpc ³)	Cluster ID $\sigma_{0.5}$	Redshift last merger	Rate z=0 (events/yr/Gpc ³)
2570	5.89	0.00	3345	1.71	0.00
4046	2.79	0.00	5480	0.03	263.69
5199	0.18	1.36	7678	2.67	0.00
8077	0.44	2.72	9937	0.83	80.93
8922	1.77	0.00	12201	1.65	0.00
10372	0.43	1.36	14366	0.33	77.48
10535	0.04	5.45	16428	0.60	29.55
16159	0.18	6.81	18776	0.34	45.33
20738	0.05	14.98	20866	0.44	18.34

TABLE IX: We present the local merger rate and redshift of the last identified in-cluster BH merger for every cluster with at least one merger within the age of the Universe. We focus on the $\{M&A, \sigma_{0.5}\}$ cluster types due to their larger merger counts. The rate is calculated using all the in-cluster mergers with $z < 1$ to approximate the local merger rate, assuming that all of the Dark Matter is contained within such dense clusters.

Following the merger history of the various clusters throughout their evolution could give us clues about possible probes of these models, as currently, our best chance to infer the existence of these objects is via gravitational waves. For that purpose, we plot in Fig 36 the in-cluster and off-cluster merger rates as a function of time for the cluster types $\{M&A, \sigma_{0.5}\}$ based on their larger merger statistics compared to the other two. The first thing we observe is there is a pronounced peak for the in-cluster merger rate at a time $T \sim \{0.5, 3\}$ Gyrs or $z \sim \{9, 2\}$ respectively for the two cluster types and a distribution for the off-cluster merger rate very broad with time. In Table IX we include the estimated local merger rate of the different clusters assuming they individually comprise all of the dark matter in our Universe. This way of presenting the rates for individual clusters rather than cumulatively for all clusters within the same type facilitates comparison with the rate we could estimate if we observed any one of these clusters in our Universe. Given the limited statistics available in some cases, these estimates are susceptible to shot noise for the single clusters.

From the in-cluster mergers perspective, current Earth-based interferometers are on the verge of being able to detect events coming from as far as $z = 2$, with examples in O3b going as high as $z = 1.18^{+0.73}_{-0.53}$ in the case of GW190403_051519 [11]. Current star formation models [248] predict a peak in the merger rate corresponding to the peak in the star formation rate ($z \sim 1.8 =$ Madau-Dickinson 3.5 Gyr after BB) plus a time delay of about a Gyr due to the time of BH collapse, binary formation and subsequent merger ($z \sim 1.1 - 0.9$). Finding an excess beyond this redshift would possibly hint towards one of the models of dense clusters of PBH analysed here.

In the case of off-cluster mergers, the rate is smaller, and instead of increasing, their redshift dependence decreases towards larger redshifts. This characteristic could be utilized to distinguish between their primordial and astrophysical origins.

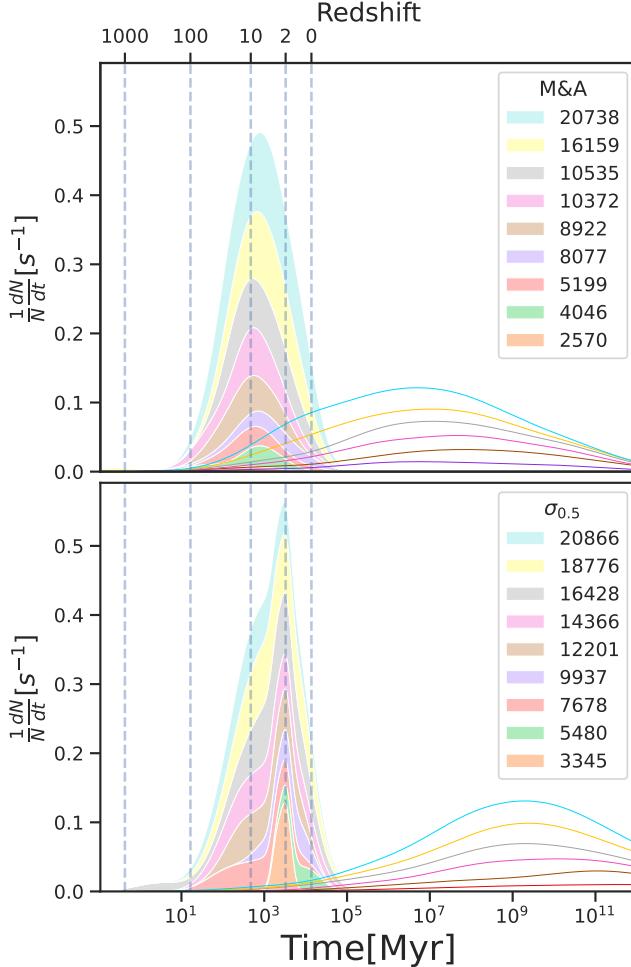


FIG. 36: We show the in-cluster and binary escapers merger rate for the cases of $M\&A$ and $\sigma_{0.5}$, which have the highest number of mergers, allowing for more accurate kernel density plots of the underlying merger rate. The binary escapers' merger rates are represented as solid lines with similar colours as the filled density plots depicting the in-cluster merger rates. We observe a general trend where the in-cluster mergers peak around 0.54 Gyr after the Big Bang for $M\&A$ clusters and 3.3 Gyr for $\sigma_{0.5}$ clusters, well before a Hubble time, while the binary escapers' merger rate extends significantly beyond it.

K. Mass distribution

Another approach to analyzing BBHs is through their mass distribution, which is highly dependent on the original distribution of masses due to the strong suppression of hierarchical mergers in realistic simulations. In Fig 37 we present a scatter plot of the mergers in the $(M_1, M_2)[M_\odot]$ plane. We can differentiate the two cases: in-cluster mergers directly identified by the code and depicted as circles in the figure and off-cluster mergers that we evolve in post-processing. A noticeable clustering around equal mass ratios is observed, deviating from the expected binary distribution one would obtain by randomly drawing pairs from the initial mass function. This suggests a genuine preference for $q = 1$ embedded in the mechanisms responsible for these mergers. This preference implies that the binary formation cross-section has a maximum for equal mass ratios, while unequal ratios result in more hyperbolic encounters. The BBH merger with the minimum mass ratio identified sits around ~ 0.2 in the $\sigma_{1.5}$, not shown in the plot, which is above the minimum mass ratio detected in a confident BBH which is ~ 0.1 for GW190814.

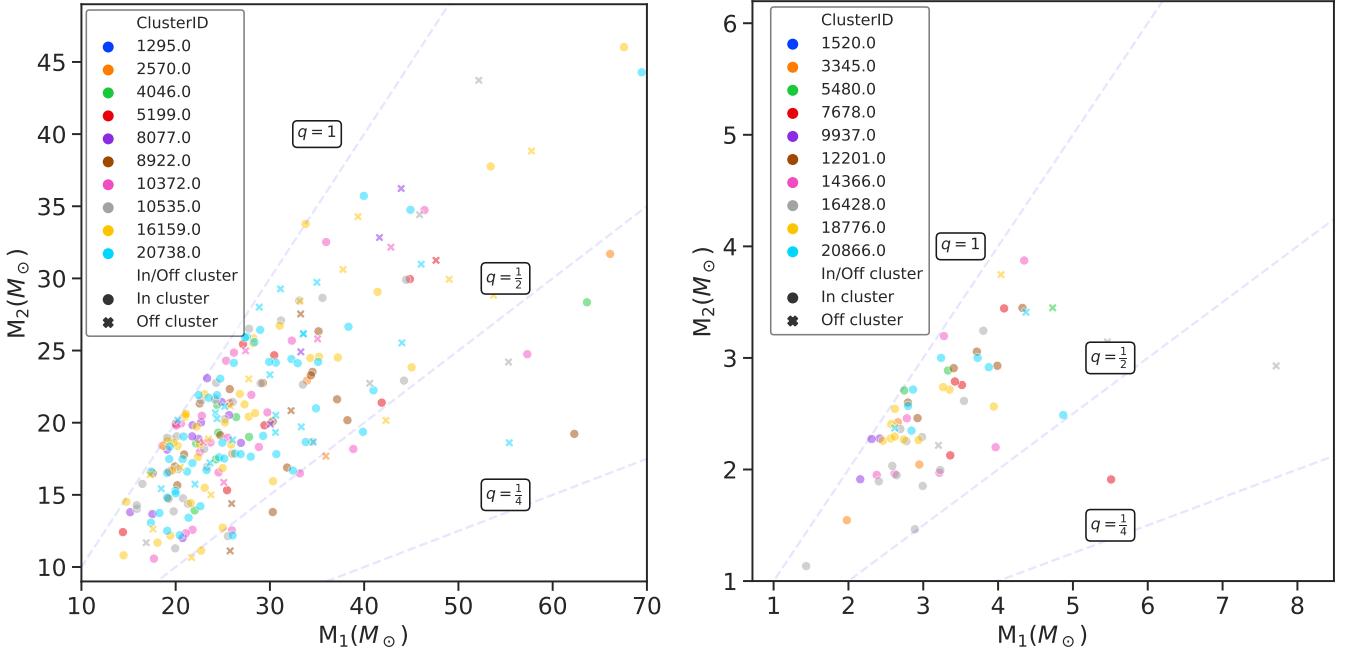


FIG. 37: In this figure, we present the distribution of the masses of the BHs that merge within the age of the Universe. They are categorized into two groups: those merging inside the clusters, identified as collisions by the code (in-cluster), and those accounted for as binary escapers and evolved outside of the simulation (off-cluster). Some clusters don't produce any merger and thus, not every color in the legend is represented in the plots. The left plot corresponds to the $M\&A$ type and the right plot represents the $\sigma_{0.5}$ models.

L. Spin distribution

Initially, the spins of all the BHs are set to 0 inspired by the isotropic collapse of the primordial overdensities predicted by most PBH formation theories. This assumption, nonetheless, may also apply to stellar BHs as the efficiency of the angular momentum transport from the spin of the progenitor star is still under debate (e.g. [249–251]). The various BHs can acquire spin via merging with other BHs. The final spin of the remnant is calculated using [243], which depends also on the progenitors' spins and masses. It is also possible that due to the multiple close encounters in dense clusters, the spin of the BH population is induced stochastically, with a final spin distribution peaked at zero with a dispersion of about $\sigma_s = 0.2$ [215]. We did not consider in this work the induced spin due to close hyperbolic encounters in the distribution of spins of the BH population.

In Fig 38 we present the distribution for the absolute magnitude of the spin of the remnant BHs. The left panel illustrates the distribution for in-cluster mergers, while the right panel depicts the same for off-cluster mergers. We observe that the bulk of the distribution is similar for both cases, as the initial spins for all BHs are 0, and thus, the final spin depends entirely on the initial masses. The most frequent value for the final spin is around $a_f = 0.68$ as you need to reach very unequal mass ratios $q < 0.4$ to go below $a_f = 0.6$. There are however, some extreme spins, sitting around $a_f = 0.9$ due to 2nd generation mergers where one of the progenitors already had a non-negligible spin. They represent a 4% of all the BBHs that coalesce while representing a larger fraction of the total number of off-cluster mergers. The smallest final spin sits around $a_f = 0.4$, corresponding to an extreme mass ratio of $q = 0.2$ in the $\sigma_{1.5}$ type of clusters and the largest one lays very close to $a_f = 0.9$, coming from a second generation merger and high total mass. This distribution aligns with general observational results and our expectations that most, if not all, BHs in our universe are Kerr BHs, as there are many mechanisms to gain spin but very few to lose it.

M. Residual eccentricity distribution

Another parameter crucial for understanding the nature of BBHs is the residual eccentricity of the binaries. Looking at fig 36, we can differentiate the in-cluster and the off-cluster mergers as two very different cases. The in-cluster mergers are the result of complex dynamical processes where more than two bodies are usually involved. The main

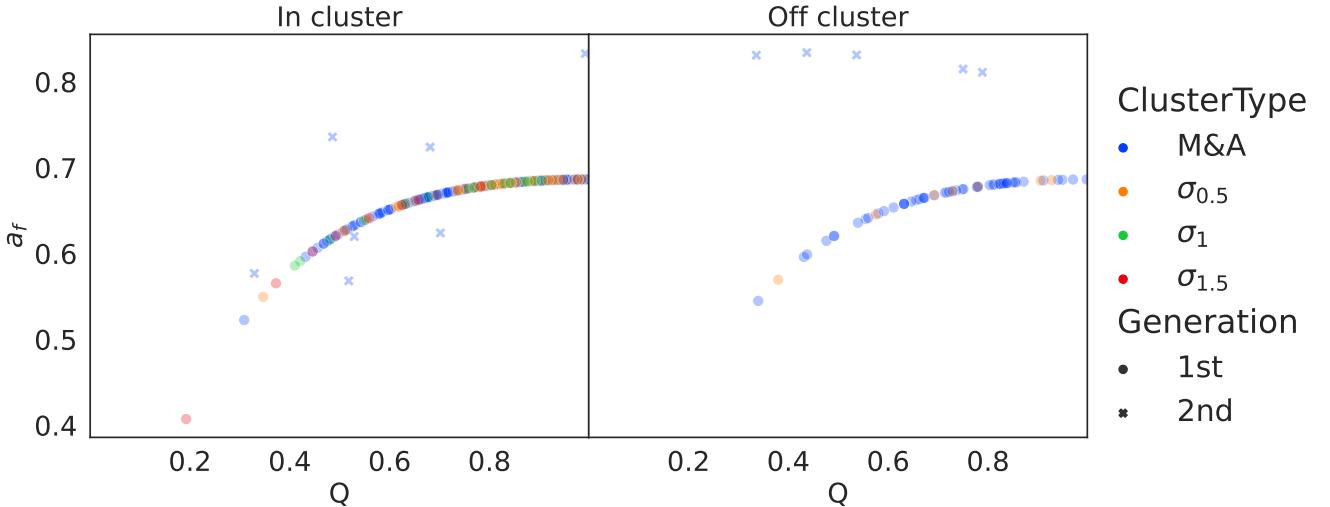


FIG. 38: In this plot we show the value for the final spins of the remnant BHs as a function of the binary mass ratios of the progenitors. The simulations assume zero spins as the initial condition. We distinguish between mergers happening inside the cluster and outside and also between 1st and 2nd generation mergers. Spins are calculated using [243] either by the code itself or in post-processing.

process by which the BHs end up merging is not gravitational wave emission, and thus, we expect to find highly eccentric signals as eccentricity cannot be radiated away. This could be a distinct feature if we were to detect an eccentric merger at a sufficiently high redshift as we would not expect in-cluster mergers to happen before the peak star formation rate era. This kind of highly eccentric merger, however, presents challenges on its own to be measured (see [252]). Given the lack of a significant inspiral phase in these events, in-cluster mergers would likely manifest in Earth-based gravitational wave interferometers as short bursts, of which one example might be GW190521 [19]. This comes with various challenges as the very low number of detectable cycles leaves the floor open to numerous alternative hypotheses [253–255].

The other case, the off-cluster mergers, follow a completely different route and the only mechanism by which they end up merging is via orbit shrinking through GW emission. As a result, eccentricity is radiated away [239] before the binary reaches the detectability threshold in current Earth-based interferometers. To investigate the validity of such a hypothesis, we plot in fig 39 the residual eccentricity for the off-cluster mergers, that is, the eccentricity that, based on quadrupolar GW emission, the binaries would have at the moment of the merger. Most of the cases that would concern us ($z < 2$) present a maximum residual eccentricity of $e \sim 10^{-4}$, which is hardly detectable given current sensitivity and the fact that it might be confused with spin effects [256]. The detectability threshold lies at around $e \sim 10^{-3}$ for a high enough SNR event, justifying the use of the quasi-circular approximation.

N. Conclusions

In this paper, we have reported the results of our study of the phenomenology of BHs-only clusters with different initial mass functions based on both astrophysical and cosmological assumptions for their origin.

From a dynamics point of view, we have encountered differences in the evaporation as well as the expansion rate of the clusters. The main distinction between cluster types can be understood via their stability, which seems to depend highly on the mass ratios of the typical encounters occurring inside of them and driving the energy exchange between the BHs. In this regard, we find that single BH escapers' distributions resemble very closely the original IMF except for a larger skew towards smaller masses and the total number of binary escapers seems to be correlated again with the stability criteria so that, the more stable the cluster is, the more binaries can fly from its gravitational potential well.

Regarding BBH coalescences, we distinguished between in-cluster and off-cluster mergers due to their significant differences in orbital parameters and potential detectability from Earth. On the one hand, in-cluster mergers are produced via 3+ body encounters inside the densest regions of the clusters, with a peak in the merger rate at a relatively high redshift, generally above the detectability horizon of current interferometers, although this is highly model-dependent. These mergers are also characterized by very high residual eccentricity and few detectable cycles

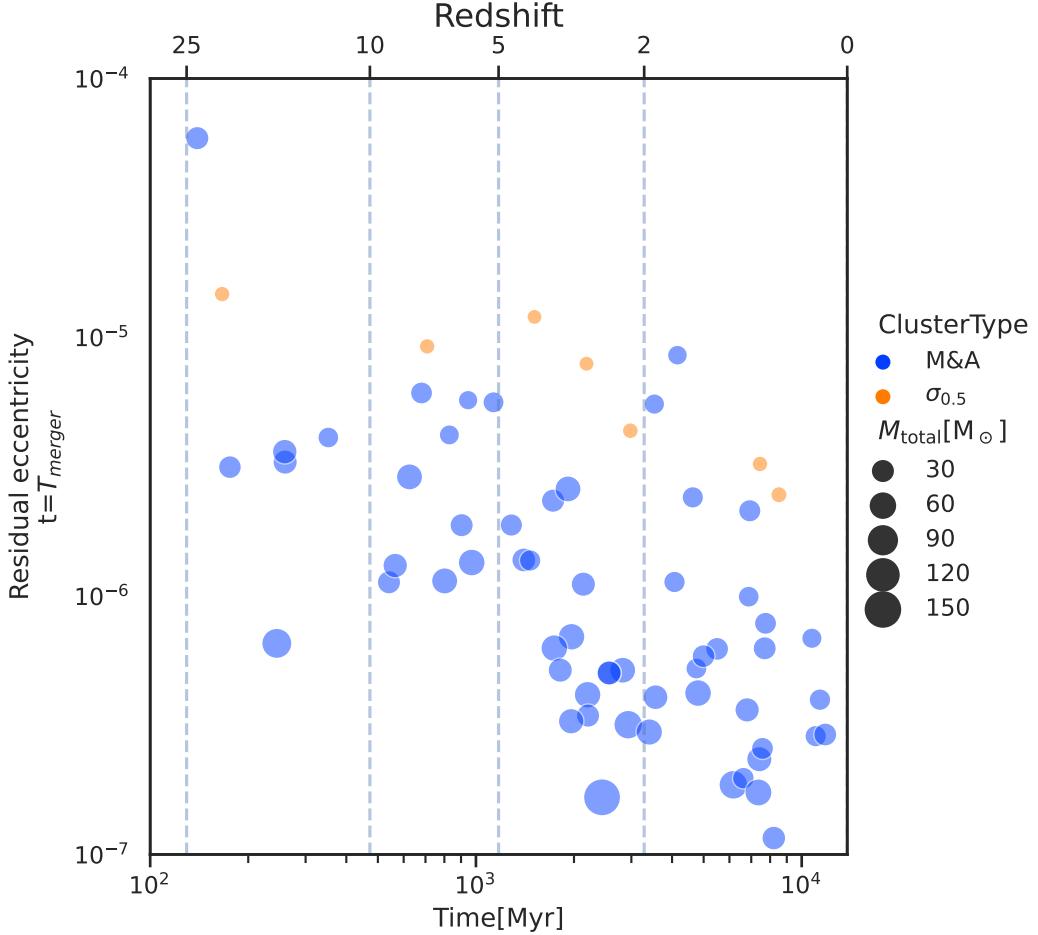


FIG. 39: In this figure, we display the residual eccentricity of the off-cluster mergers. We define residual eccentricity as the eccentricity the binary would possess at the time of merger, after it has been radiated away throughout the inspiral phase. This residual eccentricity could potentially be detected from ground-based interferometers.

from Earth. On the other hand, off-cluster mergers occur outside of the cluster due to gravitational wave emission. This mechanism leads to a very low residual eccentricity ($e < 10^{-4}$), leading to quasi-circular BBH coalescence as observed from Earth. The masses of the BBH mergers depend very much on the IMF although we find a trend favouring mass ratios close to 1 independent of the IMF. Given our zero-natal spin assumption, most of the final spins depend on the progenitor masses with a very low probability of non-zero progenitor spin (< 4%).

We also conducted a comparison between the *M&A* clusters with and without considering BH merger kicks, finding a strong dependence on the existence and number of hierarchical mergers. This finding provides valuable insight into the accuracy of numerical simulations. Predictions regarding the occurrence of these types of mergers cannot be relied upon if merger kicks are not incorporated.

In summary, this work aims to elucidate the phenomenology of complex non-linear systems represented by gravitationally-bound BH-only clusters. We hint at some results as promising candidates for unexplained observations and expand the world of possibilities in numerical research.

VIII Gravitational Energy and Angular Momentum Loss in CHE

A. Introduction

Close encounters and their accurate description represent one of the main challenges in modern gravitational direct N-body simulation codes. The correct and efficient integration of such phenomena is critical for the comprehension of the phenomenology of dense stellar environments. What is more, given the distances and velocities involved in the most intense of these interactions, radiative corrections start playing a non-negligible role. Until now, most of the efforts had been dedicated to the two-body closed-orbit case as already implemented in the reference N-body code **NBODY6++GPU** [10]. In the following section, we will derive analytic formulas describing the gravitational wave emission of an open orbit encounter between two given bodies and the quadrupolar approximation. The format and similarity with the closed-orbit case made them easily and straightforwardly implementable in the reference N-body code as an extension when the eccentricity parameter is above 1.

B. Considerations & derivation

The equation of motion for a body in a hyperbolic trajectory around another is

$$r(\phi) = \frac{a(e^2 - 1)}{1 + e \cos(\phi - \phi_0)} \quad (168)$$

First, we are going to use the quadrupole approximation [69]; that is, the fact that the gravitational wave emission is dominated by the quadrupole moment of the mass distribution of the system. For that matter, the energy and angular momentum loss can be written as

$$\begin{aligned} \frac{dE}{dt} &= -\frac{G}{45c^5} \langle \ddot{Q}_{ij} \ddot{Q}_{ij} \rangle \\ \frac{dL^i}{dt} &= -\frac{2G}{45c^5} \epsilon^{ijk} \langle \ddot{Q}_{ja} \ddot{Q}_{ka} \rangle \end{aligned} \quad (169)$$

We assume the orbit lays in the z plane and thus we only need to keep the $i = 3$ component. Using the quadrupole moment given by [9]

$$Q_{ij} = r^2 \mu \begin{pmatrix} -1 + 3 \cos^2 \phi & 3 \cos \phi \sin \phi & 0 \\ 3 \cos \phi \sin \phi & -1 + 3 \sin^2 \phi & 0 \\ 0 & 0 & -1 \end{pmatrix} \quad (170)$$

with $\mu = \frac{M_1 M_2}{M_1 + M_2}$ and $M = M_1 + M_2$ we calculate such derivatives

$$\begin{aligned} \frac{dE}{dt} &= -\frac{4G^4 M^3 \mu^2 (1 + e \cos(\psi))^4}{15|a|^5 c^5 (-1 + e^2)^5} f(\psi) \\ f(\psi) &= (24 + 13e^2 + 11e^2 \cos(2\psi) + 48e \cos(\psi)) \end{aligned} \quad (171)$$

where we have defined $\psi = \phi - \phi_0$ and

$$\begin{aligned} \frac{dL^z}{dt} &= -\frac{4G(|a|GM)^{5/2} \mu^2 (1 + e \cos(\psi))^3}{5a^6 c^5 (-1 + e^2)^{7/2}} g(\psi) \\ g(\psi) &= (8 + e^2 + 3e(e \cos(2\psi) + 4 \cos(\psi))) \end{aligned} \quad (172)$$

Now using that for hyperbolic encounters

$$\begin{aligned} e^2 &= 1 + \frac{L^2}{|a|GM\mu^2} \\ E &= \frac{GM\mu}{2|a|} \end{aligned} \quad (173)$$

where $a < 0$ so that we are consistent with the elliptical case. We can write the variation with time of the eccentricity and the semimajor axes, as needed for the evolution of the orbit

$$\begin{aligned}\frac{da}{dt} &= -\frac{8G^3M^2\mu(1+e\cos(\psi))^4}{15|a|^3c^5(-1+e^2)^5}f(\psi) \\ \frac{de}{dt} &= -\frac{2G^3M^2\mu}{15a^4c^5e(-1+e^2)^4}(1+e\cos(\psi))^3(e(116+6e^2+ \\ &\quad +11e\cos(3\psi)+2(26+9e^2)\cos(2\psi))+ \\ &\quad +(72+109e^2)\cos(\psi))\end{aligned}\quad (174)$$

With $f(\psi)$ as defined in eq (171). Given the needs of our program, we need to write all variables as functions of $\{a, e, r, \mu, M\}$, thus, we can use eq (168)

$$\psi = \arccos\left(\frac{1}{e}\left(\frac{q(e^2-1)}{r}-1\right)\right)\quad (175)$$

We could implement already these equations. However, it would not go in the same spirit as the Peters and Mathews approximation [257]. Thus, as suggested also by Rainer, we should calculate the orbit averaged energy and angular momentum emission. To do so, we need to integrate from $\psi = -\phi_0$ to $\psi = \phi_0$ with

$$\phi_0 = \cos^{-1}\left(-\frac{1}{e}\right)\quad (176)$$

and divide the emission by the typical crossing time $T = 2\pi\sqrt{|a|^3/GM}$. The result for the energy emission is the following

$$\begin{aligned}\delta E &= \iint_{-\phi_0}^{\phi_0} \frac{dE}{dt} \frac{dt}{d\psi} d\psi \\ \frac{\delta E}{T} &= -\frac{G^4 M m_1^2 m_2^2}{45\pi|a|^5 c^5 (e^2-1)^{7/2}} \bar{f}(e) \\ \bar{f}(e) &= (673e^2 + 602) \sqrt{e^2 - 1} + 3(37e^4 + 292e^2 + 96) \sec^{-1}(-e)\end{aligned}\quad (177)$$

Having used

$$\frac{dt}{d\psi} = \frac{[|a|(e^2-1)]^{3/2}}{(GM)^{1/2}} \frac{1}{(1+e\cos\psi)^2}\quad (178)$$

Now we do the same with the angular momentum

$$\begin{aligned}\delta L^z &= \iint_{-\phi_0}^{\phi_0} \frac{dL^z}{dt} \frac{dt}{d\psi} d\psi \\ \frac{\delta L^z}{T} &= -\frac{4m_1^2 m_2^2 M^{1/2} G^{7/2}}{5\pi|a|^{7/2} c^5 (e^2-1)^{2}} \bar{g}(e) \\ \bar{g}(e) &= (2e^2 + 13) \sqrt{e^2 - 1} + (7e^2 + 8) \sec^{-1}(-e)\end{aligned}\quad (179)$$

For the semimajor axis and eccentricity, we have

$$\begin{aligned}\frac{\delta a}{T} &= -\frac{2G^3 m_1 m_2 M}{45\pi|a|^3 c^5 (e^2-1)^{7/2}} \bar{f}(e) \\ \frac{\delta e}{T} &= -\frac{G^3 M m_1 m_2}{45\pi a^4 c^5 e (e^2-1)^{5/2}} ((134 + 1069e^2 + 72e^4) \sqrt{(e^2-1)} \\ &\quad + 3e^2 (304 + 121e^2) \sec^{-1}(-e))\end{aligned}\quad (180)$$

We see that in the limit $e \rightarrow 1$ the expressions are equivalent to those of [257] and thus, these formulas represent their analytic continuation. Now, for the implementation in Nbody6++GPU, we will make them symmetric for efficiency purposes

Defining the following mass coefficients

$$f_c = \frac{m_1 m_2 M}{c^5} \quad M = m_1 + m_2, \quad (181)$$

$$\{F_0, F_2\}$$

$$F_0 = 1 + \frac{73}{24}e^2 + \frac{37}{96}e^4 \quad (182)$$

$$F_2 = 1 + \frac{7}{8}e^2 \quad (183)$$

and $\{G_0, G_1, G_2\}$ similar to the $e < 1$ case

$$\begin{aligned} G_0 = & \frac{1}{\pi} \sec^{-1}(-e) F_0 + \\ & + \frac{1}{24\pi} (e^2 - 1)^{0.5} \left(\frac{301}{6} + \frac{673}{12} e^2 \right) \left(\right. \end{aligned} \quad (184)$$

$$\begin{aligned} G_1 = & \frac{1}{8\pi} \left(\frac{67}{36} + \frac{1069}{72} e^2 + e^4 \right) \sqrt{(e^2 - 1)} \\ & + \frac{1}{\pi} e^2 \left(\frac{19}{12} + \frac{121}{192} e^2 \right) \left(\sec^{-1}(-e) \right. \end{aligned} \quad (185)$$

$$G_2 = \frac{1}{8\pi} \left(\frac{1}{2} e^2 + 13 \right) \sqrt{e^2 - 1} + \frac{1}{\pi} \sec^{-1}(-e) F_2 \quad (186)$$

we write the final formulas entering the actual Fortran code

$$\frac{dE}{dt} = \frac{32}{5} f_c \frac{m_1 m_2}{a^5} (e^2 - 1)^{-3.5} G_0 \quad (187)$$

$$\frac{dJ}{dt} = \frac{32}{5} f_c \frac{m_1 m_2}{M^{1/2} a^{7/2}} (e^2 - 1)^{-2} G_2 \quad (188)$$

$$\frac{da}{dt} = \frac{64}{5} f_c \frac{1}{a^3} (e^2 - 1)^{-3.5} G_0 \quad (189)$$

$$\frac{de}{dt} = \frac{64}{5} f_c \frac{1}{e} \frac{1}{a^4} (e^2 - 1)^{-2.5} G_1 \quad (190)$$

IX Thesis Conclusion

As this thesis ends, it is evident that gravitational wave astronomy has opened up new avenues for understanding our Cosmos. Through the various research works presented here, we've gained insights into many aspects of the production and detection of gravitational waves, with deep dives into the mechanics of the detectors, the analysis techniques employed in the characterisation of the signals and new proposals regarding the potential origins of their main sources, binary black holes. The primordial black hole paradigm has been a main repeating point throughout the thesis. It has guided our work and many of the discussions surrounding it. Apart from these works, we have also contributed to the LIGO-Virgo-KAGRA Collaboration, adding our bit to one of the largest scientific collaborations in the history of humankind.

We have witnessed the birth of a new era, one where gravitational waves act as messengers from the depths of space, carrying information about celestial phenomena that were once thought to be beyond our reach. This thesis represents nothing more than a first crude and wide glimpse at this vast and yet-to-be-explored new window to the Universe, with plenty of more gravitational wave observations to come with the improvements of current interferometers and the outright construction of new and more advanced apparatus [91, 258]. As we stand here, it is worth reflecting on the journey that led to this discovery starting with the first derivation of the linearized equations of motion of General Relativity in 1916 by Albert Einstein himself. At that time, the reality of these quadrupolar ripples in the fabric of spacetime was doubtful and later on, its detectability was thought to be impossible. Yet, humanity's relentless pursuit of knowledge defied expectations and right now we detect events every week with ever-increasing accuracy. This is one of the many examples of the importance of humility in the face of discovery.

As we continue the scientific endeavour to expand human understanding, we need to accept that human knowledge will continue to advance faster than our expectations and imagination. And this, at least for me, is sufficient reason to embrace curiosity and pursue learning as a guiding principle.

-
- [1] J. García-Bellido, J. F. Nuño Siles, and E. Ruiz Morales, *Phys. Dark Univ.* **31**, 100791 (2021), arXiv:2010.13811 [astro-ph.CO].
- [2] G. Morras, J. F. N. n. Siles, J. Garcia-Bellido, and E. R. Morales, *Phys. Rev. D* **107**, 023027 (2023), arXiv:2209.05475 [gr-qc].
- [3] G. Morras, J. F. N. Siles, and J. Garcia-Bellido, *Phys. Rev. D* **108**, 123025 (2023), arXiv:2307.16610 [gr-qc].
- [4] S. R. Cook, A. Gelman, and D. B. Rubin, *Journal of Computational and Graphical Statistics* **15**, 675 (2006), <https://doi.org/10.1198/106186006X136976>.
- [5] S. Talts, M. Betancourt, D. Simpson, A. Vehtari, and A. Gelman, *arXiv e-prints* , arXiv:1804.06788 (2018), arXiv:1804.06788 [stat.ME].
- [6] G. Morras *et al.*, *Phys. Dark Univ.* **42**, 101285 (2023), arXiv:2301.11619 [gr-qc].
- [7] K. S. Phukon, G. Baltus, S. Caudill, S. Clesse, A. Depasse, M. Fays, H. Fong, S. J. Kapadia, R. Magee, and A. J. Tanasijczuk, (2021), arXiv:2105.11449 [astro-ph.CO].
- [8] J. F. N. Siles and J. G.-B. Capdevila, "Primordial black hole clusters, phenomenology & implications," (2024), arXiv:2405.06391.
- [9] J. García-Bellido and S. Nesseris, *Physics of the Dark Universe* **21**, 61 (2018).
- [10] Wang, Long and Spurzem, Rainer and Aarseth, Sverre and Nitadori, Keigo and Berczik, Peter and Kouwenhoven, M. B. N. and Naab, Thorsten, *Monthly Notices of the Royal Astronomical Society* **450** (2015), 10.1093/mnras/stv817.
- [11] R. Abbott, T. D. Abbott, *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), *Phys. Rev. X* **13**, 041039 (2023).
- [12] R. Abbott *et al.* (LIGO Scientific, Virgo, KAGRA), "Search for subsolar-mass black hole binaries in the second part of advanced ligo's and advanced virgo's third observing run," (2022).
- [13] M. Prunier, G. Morras, J. F. N. Siles, S. Clesse, J. García-Bellido, and E. R. Morales, "Analysis of the subsolar-mass black hole candidate ssm200308 from the second part of the third observing run of advanced ligo-virgo," (2023), arXiv:2311.16085 [gr-qc].
- [14] LIGO Scientific Collaboration, "LIGO Algorithm Library - LALSuite," free software (GPL) (2018).
- [15] B. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. X* **9**, 031040 (2019), arXiv:1811.12907 [astro-ph.HE].
- [16] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, and et al., *The Astrophysical Journal* **892**, L3 (2020).
- [17] R. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. D* **102**, 043015 (2020), arXiv:2004.08342 [astro-ph.HE].
- [18] B. Abbott *et al.*, *The Astrophysical Journal* **896**, L44 (2020).
- [19] R. Abbott, T. D. Abbott, and others. (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **125**, 101102 (2020).
- [20] S. Clesse and J. García-Bellido, *Phys. Rev. D* **92**, 023524 (2015), arXiv:1501.07565 [astro-ph.CO].

- [21] J. García-Bellido, *Proceedings, 11th International LISA Symposium: Zurich, Switzerland, September 5-9, 2016*, J. Phys. Conf. Ser. **840**, 012032 (2017), arXiv:1702.08275 [astro-ph.CO].
- [22] B. P. Abbott *et al.* (LIGO Scientific, Virgo), Phys. Rev. Lett. **116**, 061102 (2016), arXiv:1602.03837 [gr-qc].
- [23] S. Bird, I. Cholis, J. B. Muñoz, Y. Ali-Haimoud, M. Kamionkowski, E. D. Kovetz, A. Raccanelli, and A. G. Riess, Phys. Rev. Lett. **116**, 201301 (2016), arXiv:1603.00464 [astro-ph.CO].
- [24] S. Clesse and J. García-Bellido, Phys. Dark Univ. **15**, 142 (2017), arXiv:1603.05234 [astro-ph.CO].
- [25] M. Sasaki, T. Suyama, T. Tanaka, and S. Yokoyama, Phys. Rev. Lett. **117**, 061101 (2016), arXiv:1603.08338 [astro-ph.CO].
- [26] B. Carr, S. Clesse, J. García-Bellido, and F. Kuhnel, Phys. Dark Univ. **31**, 100755 (2021), arXiv:1906.08217 [astro-ph.CO].
- [27] S. Clesse and J. García-Bellido, Phys. Dark Univ. **22**, 137 (2018), arXiv:1711.10458 [astro-ph.CO].
- [28] M. Zevin, M. Spera, C. P. L. Berry, and V. Kalogera, Astrophys. J. Lett. **899**, L1 (2020), arXiv:2006.14573 [astro-ph.HE].
- [29] S. Vitale, R. Lynch, R. Sturani, and P. Graff, Class. Quant. Grav. **34**, 03LT01 (2017), arXiv:1503.04307 [gr-qc].
- [30] C. Talbot and E. Thrane, Phys. Rev. D **96**, 023012 (2017), arXiv:1704.08370 [astro-ph.HE].
- [31] K. Belczynski *et al.*, Astron. Astrophys. **636**, A104 (2020), arXiv:1706.07053 [astro-ph.HE].
- [32] V. De Luca, V. Desjacques, G. Franciolini, A. Malhotra, and A. Riotto, JCAP **1905**, 018 (2019), arXiv:1903.01179 [astro-ph.CO].
- [33] V. De Luca, G. Franciolini, P. Pani, and A. Riotto, JCAP **2004**, 052 (2020), arXiv:2003.02778 [astro-ph.CO].
- [34] D. Gerosa and E. Berti, Phys. Rev. **D95**, 124046 (2017), arXiv:1703.06223 [gr-qc].
- [35] M. Fishbach, D. E. Holz, and B. Farr, Astrophys. J. Lett. **840**, L24 (2017), arXiv:1703.06869 [astro-ph.HE].
- [36] E. Barausse and L. Rezzolla, Astrophys. J. Lett. **704**, L40 (2009), arXiv:0904.2577 [gr-qc].
- [37] C. Kimball, C. Talbot, C. P. L. Berry, M. Carney, M. Zevin, E. Thrane, and V. Kalogera, Astrophys. J. **900**, 177 (2020), arXiv:2005.00023 [astro-ph.HE].
- [38] W. M. Farr, S. Stevenson, M. Coleman Miller, I. Mandel, B. Farr, and A. Vecchio, Nature **548**, 426 (2017), arXiv:1706.01385 [astro-ph.HE].
- [39] B. Farr, D. E. Holz, and W. M. Farr, Astrophys. J. Lett. **854**, L9 (2018), arXiv:1709.07896 [astro-ph.HE].
- [40] V. Tiwari, S. Fairhurst, and M. Hannam, Astrophys. J. **868**, 140 (2018), arXiv:1809.01401 [gr-qc].
- [41] N. Fernandez and S. Profumo, JCAP **08**, 022 (2019), arXiv:1905.13019 [astro-ph.HE].
- [42] B. Abbott *et al.* (LIGO Scientific, Virgo), Astrophys. J. Lett. **882**, L24 (2019), arXiv:1811.12940 [astro-ph.HE].
- [43] S. Miller, T. A. Callister, and W. Farr, Astrophys. J. **895**, 128 (2020), arXiv:2001.06051 [astro-ph.HE].
- [44] A. Hall, A. D. Gow, and C. T. Byrnes, Phys. Rev. D **102**, 123524 (2020).
- [45] F. Hofmann, E. Barausse, and L. Rezzolla, Astrophys. J. Lett. **825**, L19 (2016), arXiv:1605.01938 [gr-qc].
- [46] K. K. Y. Ng, S. Vitale, A. Zimmerman, K. Chatzioannou, D. Gerosa, and C.-J. Haster, Phys. Rev. **D98**, 083007 (2018), arXiv:1805.03046 [gr-qc].
- [47] H. Jeffreys, *The Theory of Probability*, Oxford U.P. (1939).
- [48] S. Nesseris and J. García-Bellido, JCAP **1308**, 036 (2013), arXiv:1210.7652 [astro-ph.CO].
- [49] A. Einstein, Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin), 688 (1916).
- [50] J. Aasi and others. (LIGO Scientific Collaboration), Classical and Quantum Gravity **32**, 074001 (2015).
- [51] F. Acernese and others. (Virgo Collaboration), Classical and Quantum Gravity **32**, 024001 (2014).
- [52] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatsumi, and H. Yamamoto (KAGRA), Phys. Rev. D **88**, 043007 (2013), arXiv:1306.6747 [gr-qc].
- [53] Q. Chu *et al.*, Phys. Rev. D **105**, 024023 (2022), arXiv:2011.06787 [gr-qc].
- [54] S. A. Usman *et al.*, Class. Quant. Grav. **33**, 215004 (2016), arXiv:1508.02357 [gr-qc].
- [55] S. A. Usman *et al.*, Class. Quant. Grav. **33**, 215004 (2016), arXiv:1508.02357 [gr-qc].
- [56] S. Sachdev *et al.*, (2019), arXiv:1901.08580 [gr-qc].
- [57] F. Aubin *et al.*, Class. Quant. Grav. **38**, 095004 (2021), arXiv:2012.11512 [gr-qc].
- [58] M. Drago *et al.*, (2020), arXiv:2006.12604 [gr-qc].
- [59] D. Davis *et al.* (LIGO), Class. Quant. Grav. **38**, 135014 (2021), arXiv:2101.11673 [astro-ph.IM].
- [60] C. W. Helstrom, *Statistical theory of signal detection* (Pergamon, Oxford, 1968).
- [61] B. Allen, Phys. Rev. D **71**, 062001 (2005), arXiv:gr-qc/0405045.
- [62] W. M. Farr, J. R. Gair, I. Mandel, and C. Cutler, Physical Review D **91** (2015), 10.1103/physrevd.91.023005.
- [63] N. Andres *et al.*, Class. Quant. Grav. **39**, 055002 (2022), arXiv:2110.10997 [gr-qc].
- [64] J. Veitch and A. Vecchio, Phys. Rev. D **78**, 022001 (2008), arXiv:0801.4313 [gr-qc].
- [65] B. P. Abbott *et al.*, Phys. Rev. D **93**, 112004 (2016), [Addendum: Phys. Rev. D **97**, 059901 (2018)], arXiv:1604.00439 [astro-ph.IM].
- [66] B. P. Abbott *et al.*, Classical and Quantum Gravity **37**, 055002 (2020).
- [67] B. P. Abbott *et al.*, Classical and Quantum Gravity **33**, 134001 (2016).
- [68] E. Thrane and C. Talbot, Publications of the Astronomical Society of Australia **36** (2019), 10.1017/pasa.2019.2.
- [69] M. Maggiore, *Gravitational Waves. Vol. 1: Theory and Experiments*, Oxford Master Series in Physics (Oxford University Press, 2007) p. 572.
- [70] J. Veitch *et al.*, Phys. Rev. D **91**, 042003 (2015), arXiv:1409.7215 [gr-qc].
- [71] G. S. Davies, T. Dent, M. Tápai, I. Harry, C. McIsaac, and A. H. Nitz, Phys. Rev. D **102**, 022004 (2020).
- [72] S. Khan, K. Chatzioannou, M. Hannam, and F. Ohme, Physical Review D **100** (2019), 10.1103/physrevd.100.024059.
- [73] E. B. Wilson, Journal of the American Statistical Association **22**, 209 (1927).
- [74] P. Ajith *et al.*, Phys. Rev. Lett. **106**, 241101 (2011), arXiv:0909.2867 [gr-qc].

- [75] C. Cutler *et al.*, Phys. Rev. Lett. **70**, 2984 (1993), arXiv:astro-ph/9208005.
- [76] H. J. Kimble, Y. Levin, A. B. Matsko, K. S. Thorne, and S. P. Vyatchanin, Phys. Rev. D **65**, 022002 (2001).
- [77] H.-Y. Chen, D. E. Holz, J. Miller, M. Evans, S. Vitale, and J. Creighton, Class. Quant. Grav. **38**, 055010 (2021), arXiv:1709.08079 [astro-ph.CO].
- [78] S. A. Usman *et al.*, Class. Quant. Grav. **33**, 215004 (2016), arXiv:1508.02357 [gr-qc].
- [79] A. H. Nitz, T. Dent, T. Dal Canton, S. Fairhurst, and D. A. Brown, Astrophys. J. **849**, 118 (2017), arXiv:1705.01513 [gr-qc].
- [80] L. S. Finn, Phys. Rev. D **46**, 5236 (1992), arXiv:gr-qc/9209010.
- [81] G. Pratten *et al.*, Phys. Rev. D **103**, 104056 (2021), arXiv:2004.06503 [gr-qc].
- [82] L. S. Collaboration, V. Collaboration, and K. Collaboration, “GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run — Parameter estimation data release,” (2021).
- [83] R. Abbott, T. D. Abbott, *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), Phys. Rev. X **13**, 011048 (2023).
- [84] L. S. Collaboration, V. Collaboration, and K. Collaboration, “GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run — Candidate data release,” (2021).
- [85] M. Abramowitz and I. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables* (Dover Publications, Incorporated, 1974).
- [86] T. L. S. Collaboration, J. Aasi, B. P. Abbott, R. Abbott, *et al.*, Classical and Quantum Gravity **32**, 074001 (2015).
- [87] F. Acernese, M. Agathos, *et al.*, Classical and Quantum Gravity **32**, 024001 (2014).
- [88] T. Akutsu *et al.* (KAGRA), Nature Astron. **3**, 35 (2019), arXiv:1811.08079 [gr-qc].
- [89] B. P. A. others, Living Reviews in Relativity **23** (2020), 10.1007/s41114-020-00026-9.
- [90] M. Maggiore, C. V. D. Broeck, N. Bartolo, E. Belgacem, D. Bertacca, M. A. Bizouard, M. Branchesi, S. Clesse, S. Foffa, J. García-Bellido, S. Grimm, J. Harms, T. Hinderer, S. Matarrese, C. Palomba, M. Peloso, A. Ricciardone, and M. Sakellariadou, Journal of Cosmology and Astroparticle Physics **2020**, 050 (2020).
- [91] M. Evans, R. X. Adhikari, C. Afle, S. W. Ballmer, S. Biscoveanu, S. Borhanian, D. A. Brown, Y. Chen, R. Eisenstein, A. Gruson, A. Gupta, E. D. Hall, R. Huxford, B. Kamai, R. Kashyap, J. S. Kissel, K. Kuns, P. Landry, A. Lenon, G. Lovelace, L. McCuller, K. K. Y. Ng, A. H. Nitz, J. Read, B. S. Sathyaprakash, D. H. Shoemaker, B. J. J. Slagmolen, J. R. Smith, V. Srivastava, L. Sun, S. Vitale, and R. Weiss, “A horizon study for cosmic explorer: Science, observatories, and community,” (2021), arXiv:2109.09882 [astro-ph.IM].
- [92] A. Klein, E. Barausse, A. Sesana, A. Petiteau, E. Berti, S. Babak, J. Gair, S. Aoudia, I. Hinder, F. Ohme, and B. Wardell, Physical Review D **93** (2015), 10.1103/PhysRevD.93.024003.
- [93] N. Tamanini, C. Caprini, E. Barausse, A. Sesana, A. Klein, and A. Petiteau, Journal of Cosmology and Astroparticle Physics **2016**, 002 (2016).
- [94] S. Babak, J. Gair, A. Sesana, E. Barausse, C. F. Sopuerta, C. P. Berry, E. Berti, P. Amaro-Seoane, A. Petiteau, and A. Klein, Physical Review D **95** (2017), 10.1103/physrevd.95.103012.
- [95] E. Thrane and C. Talbot, Publications of the Astronomical Society of Australia **36** (2019), 10.1017/pasa.2019.2.
- [96] C. Mills and S. Fairhurst, Phys. Rev. D **103**, 024042 (2021), arXiv:2007.04313 [gr-qc].
- [97] S. Morisaki, Phys. Rev. D **104**, 044062 (2021).
- [98] B. Zackay, L. Dai, and T. Venumadhav, (2018), arXiv:1806.08792 [astro-ph.IM].
- [99] N. J. Cornish, Phys. Rev. D **104**, 104054 (2021), arXiv:2109.02728 [gr-qc].
- [100] A. J. K. Chua and M. Vallisneri, Phys. Rev. Lett. **124**, 041102 (2020), arXiv:1909.05966 [gr-qc].
- [101] S. R. Green, C. Simpson, and J. Gair, Phys. Rev. D **102**, 104057 (2020), arXiv:2002.07656 [astro-ph.IM].
- [102] H. Antil, S. E. Field, F. Herrmann, R. H. Nochetto, and M. Tiglio, J. Sci. Comput. **57**, 604 (2013), arXiv:1210.0577 [cs.NA].
- [103] P. Canizares, S. E. Field, J. R. Gair, and M. Tiglio, Phys. Rev. D **87**, 124005 (2013), arXiv:1304.0462 [gr-qc].
- [104] R. Smith, S. E. Field, K. Blackburn, C.-J. Haster, M. Pürer, V. Raymond, and P. Schmidt, Phys. Rev. D **94**, 044031 (2016), arXiv:1604.08253 [gr-qc].
- [105] H. Qi and V. Raymond, Phys. Rev. D **104**, 063031 (2021), arXiv:2009.13812 [gr-qc].
- [106] S. Morisaki and V. Raymond, Phys. Rev. D **102**, 104020 (2020), arXiv:2007.09108 [gr-qc].
- [107] S. Morisaki, R. Smith, L. Tsukada, S. Sachdev, S. Stevenson, C. Talbot, and A. Zimmerman, Phys. Rev. D **108**, 123040 (2023).
- [108] C. Pankow, P. Brady, E. Ochsner, and R. O’Shaughnessy, Physical Review D **92** (2015), 10.1103/physrevd.92.023002.
- [109] J. Lange, R. O’Shaughnessy, and M. Rizzo, (2018), arXiv:1805.10457 [gr-qc].
- [110] L. Pathak, A. Reza, and A. S. Sengupta, Phys. Rev. D **108**, 064055 (2023).
- [111] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürer, Phys. Rev. Lett. **113**, 151101 (2014).
- [112] G. Pratten, C. García-Quirós, M. Colleoni, A. Ramos-Buades, H. Estellés, M. Mateu-Lucena, R. Jaume, M. Haney, D. Keitel, J. E. Thompson, and et al., Physical Review D **103** (2021), 10.1103/physrevd.103.104056.
- [113] T. Dietrich, A. Samajdar, S. Khan, N. K. Johnson-McDaniel, R. Dudi, and W. Tichy, Phys. Rev. D **100**, 044003 (2019), arXiv:1905.06011 [gr-qc].
- [114] D. Chiaramello and A. Nagar, Phys. Rev. D **101**, 101501 (2020), arXiv:2001.11736 [gr-qc].
- [115] P. Jaranowski, A. Krolak, and B. F. Schutz, Phys. Rev. D **58**, 063001 (1998), arXiv:gr-qc/9804014.
- [116] M. Pürer, Class. Quant. Grav. **31**, 195010 (2014), arXiv:1402.4146 [gr-qc].
- [117] R. Cotesa, S. Marsat, and M. Pürer, Phys. Rev. D **101**, 124040 (2020), arXiv:2003.12079 [gr-qc].

- [118] D. B. Szyld, *Numerical Algorithms* **42**, 309 (2006).
- [119] P. Ajith, M. Hannam, S. Husa, Y. Chen, B. Brügmann, N. Dorband, D. Müller, F. Ohme, D. Pollney, C. Reisswig, L. Santamaría, and J. Seiler, *Phys. Rev. Lett.* **106**, 241101 (2011).
- [120] L. Santamaría, F. Ohme, P. Ajith, B. Brügmann, N. Dorband, M. Hannam, S. Husa, P. Mösta, D. Pollney, C. Reisswig, E. L. Robinson, J. Seiler, and B. Krishnan, *Phys. Rev. D* **82**, 064016 (2010).
- [121] P. Schmidt, F. Ohme, and M. Hannam, *Phys. Rev. D* **91**, 024043 (2015).
- [122] C. García-Quirós, S. Husa, M. Mateu-Lucena, and A. Borchers, *Class. Quant. Grav.* **38**, 015006 (2021), arXiv:2001.10897 [gr-qc].
- [123] S. S. Wilks, *The Annals of Mathematical Statistics* **9**, 60 (1938).
- [124] I. Romero-Shaw, C. Talbot, S. Biscoveanu, V. d'Emilio, G. Ashton, C. Berry, S. Coughlin, S. Galaudage, C. Hoy, M. Huebner, K. S. Phukon, M. Pitkin, M. Rizzo, N. Sarin, R. Smith, S. Stevenson, A. Vajpeyi, M. Arene, K. Athar, and L. Xiao, *Monthly Notices of the Royal Astronomical Society* **499** (2020), 10.1093/mnras/staa2850.
- [125] J. S. Speagle, *Monthly Notices of the Royal Astronomical Society* **493**, 3132 (2020).
- [126] G. Ashton *et al.*, *Astrophys. J. Suppl.* **241**, 27 (2019), arXiv:1811.02042 [astro-ph.IM].
- [127] G. van Rossum, in *Proceedings of the 2007 USENIX Annual Technical Conference, Santa Clara, CA, USA, June 17-22, 2007*, edited by J. Chase and S. Seshan (USENIX, 2007).
- [128] LIGO Scientific Collaboration, “LIGO Algorithm Library - LALSuite,” free software (GPL) (2018).
- [129] L. Prechelt, *Computer* **33**, 23 (2000).
- [130] B. Abbott *et al.*, *Physical Review Letters* **119** (2017), 10.1103/physrevlett.119.161101.
- [131] N. J. Cornish and T. B. Littenberg, *Classical and Quantum Gravity* **32**, 135012 (2015).
- [132] T. B. Littenberg and N. J. Cornish, *Phys. Rev. D* **91**, 084034 (2015).
- [133] C. Cahillane, J. Betzwieser, D. A. Brown, E. Goetz, E. D. Hall, K. Izumi, S. Kandhasamy, S. Karki, J. S. Kissel, G. Mendell, R. L. Savage, D. Tuyenbayev, A. Urban, A. Viets, M. Wade, and A. J. Weinstein, *Physical Review D* **96** (2017), 10.1103/physrevd.96.102001.
- [134] F. Acernese *et al.* (The Virgo Collaboration), *Classical and Quantum Gravity* **39**, 045006 (2022).
- [135] L. Sun, E. Goetz, J. S. Kissel, J. Betzwieser, S. Karki, A. Viets, M. Wade, D. Bhattacharjee, V. Bossilkov, P. B. Covas, L. E. H. Datrier, R. Gray, S. Kandhasamy, Y. K. Lecoeuche, G. Mendell, T. Mistry, E. Payne, R. L. Savage, A. J. Weinstein, S. Aston, A. Buikema, C. Cahillane, J. C. Driggers, S. E. Dwyer, R. Kumar, and A. Urban, *Classical and Quantum Gravity* **37**, 225008 (2020).
- [136] R. Abbott *et al.* (LIGO Scientific, VIRGO, KAGRA), (2023), arXiv:2302.03676 [gr-qc].
- [137] J. Lin, *IEEE Transactions on Information Theory* **37**, 145 (1991).
- [138] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. X* **9**, 011001 (2019), arXiv:1805.11579 [gr-qc].
- [139] LIGO Open Science Center (LOSC), “Data release for event GW170817,” (2017).
- [140] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. X* **6**, 041015 (2016), [Erratum: *Phys. Rev. X* 8, 039903 (2018)], arXiv:1606.04856 [gr-qc].
- [141] R. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. X* **11**, 021053 (2021), arXiv:2010.14527 [gr-qc].
- [142] R. Abbott, T. D. Abbott, *et al.* (The LIGO Scientific Collaboration and the Virgo Collaboration), *Phys. Rev. D* **109**, 022001 (2024).
- [143] B. Abbott, R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, V. Adya, C. Affeldt, M. Agathos, K. Agatsuma, N. Aggarwal, O. Aguiar, L. Aiello, A. Ain, P. Ajith, T. Akutsu, G. Allen, and A. Allocca, *Living Reviews in Relativity* **23** (2020), 10.1007/s41114-020-00026-9.
- [144] B. Abbott *et al.*, *The Astrophysical Journal* **900**, L13 (2020).
- [145] S. E. Woosley, *Astrophys. J.* **836**, 244 (2017), arXiv:1608.08939 [astro-ph.HE].
- [146] Y. B. Zel'dovich and I. D. Novikov, *Soviet Astronomy* **10**, 602 (1967).
- [147] S. Hawking, *Mon. Not. Roy. Astron. Soc.* **152**, 75 (1971).
- [148] B. J. Carr and S. W. Hawking, *Mon. Not. Roy. Astron. Soc.* **168**, 399 (1974).
- [149] G. F. Chapline, *Nature* **253**, 251 (1975).
- [150] J. García-Bellido, A. D. Linde, and D. Wands, *Phys. Rev. D* **54**, 6040 (1996), arXiv:astro-ph/9605094.
- [151] C. Kouvaris and P. Tinyakov, *Phys. Rev. D* **83**, 083512 (2011), arXiv:1012.2039 [astro-ph.HE].
- [152] A. de Lavallaz and M. Fairbairn, *Phys. Rev. D* **81**, 123521 (2010), arXiv:1004.0629 [astro-ph.GA].
- [153] J. Bramante and T. Linden, *Phys. Rev. Lett.* **113**, 191301 (2014), arXiv:1405.1031 [astro-ph.HE].
- [154] J. Bramante and F. Elahi, *Phys. Rev. D* **91**, 115001 (2015), arXiv:1504.04019 [hep-ph].
- [155] J. Bramante, T. Linden, and Y.-D. Tsai, *Phys. Rev. D* **97**, 055016 (2018), arXiv:1706.00001 [hep-ph].
- [156] C. Kouvaris, P. Tinyakov, and M. H. G. Tytgat, *Phys. Rev. Lett.* **121**, 221102 (2018), arXiv:1804.06740 [astro-ph.HE].
- [157] S. Shandera, D. Jeong, and H. S. G. Gebhardt, *Phys. Rev. Lett.* **120**, 241102 (2018), arXiv:1802.08206 [astro-ph.CO].
- [158] J. H. Chang, D. Egana-Ugrinovic, R. Essig, and C. Kouvaris, *JCAP* **03**, 036 (2019), arXiv:1812.07000 [hep-ph].
- [159] M. A. Latif, A. Lupi, D. R. G. Schleicher, G. D'Amico, P. Panci, and S. Bovino, *Mon. Not. Roy. Astron. Soc.* **485**, 3352 (2019), arXiv:1812.03104 [astro-ph.CO].
- [160] B. Dasgupta, R. Laha, and A. Ray, *Phys. Rev. Lett.* **126**, 141105 (2021), arXiv:2009.01825 [astro-ph.HE].
- [161] J. Gurian, M. Ryan, S. Schon, D. Jeong, and S. Shandera, *Astrophys. J. Lett.* **939**, L12 (2022), arXiv:2209.00064 [astro-ph.CO].
- [162] M. Ryan and D. Radice, *Phys. Rev. D* **105**, 115034 (2022), arXiv:2201.05626 [astro-ph.HE].
- [163] M. Hippert, E. Dillingham, H. Tan, D. Curtin, J. Noronha-Hostler, and N. Yunes, (2022), arXiv:2211.08590 [astro-ph.HE].

- [164] J. M. Corral-Santana, J. Casares, T. Munoz-Darias, F. E. Bauer, I. G. Martinez-Pais, and D. M. Russell, *Astron. Astrophys.* **587**, A61 (2016), arXiv:1510.08869 [astro-ph.HE].
- [165] B. Paczynski, *Astrophys. J.* **304**, 1 (1986).
- [166] L. Wyrzykowski and I. Mandel, *Astron. Astrophys.* **636**, A20 (2020), arXiv:1904.07789 [astro-ph.SR].
- [167] H. Niikura *et al.*, *Nature Astron.* **3**, 524 (2019), arXiv:1701.02151 [astro-ph.CO].
- [168] M. R. S. Hawkins, *Astron. Astrophys.* **643**, A10 (2020), arXiv:2010.15007 [astro-ph.CO].
- [169] M. R. S. Hawkins, *Mon. Not. Roy. Astron. Soc.* **512**, 5706 (2022), arXiv:2204.09143 [astro-ph.CO].
- [170] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. Lett.* **121**, 231103 (2018), arXiv:1808.04771 [astro-ph.CO].
- [171] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. Lett.* **123**, 161102 (2019), arXiv:1904.08976 [astro-ph.CO].
- [172] A. H. Nitz and Y.-F. Wang, arXiv (2021), arXiv:2102.00868 [astro-ph.HE].
- [173] R. Abbott *et al.* (LIGO Scientific, VIRGO, KAGRA), *Phys. Rev. Lett.* **129**, 061104 (2022), arXiv:2109.12197 [astro-ph.CO].
- [174] R. Abbott *et al.* (LIGO Scientific, VIRGO, KAGRA), (2022), arXiv:2212.01477 [astro-ph.HE].
- [175] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 084043 (2009), arXiv:0907.0700 [gr-qc].
- [176] D. Davis, T. B. Littenberg, I. M. Romero-Shaw, M. Millhouse, J. McIver, F. D. Renzo, and G. Ashton, *Classical and Quantum Gravity* **39**, 245013 (2022).
- [177] B. P. Abbott, R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. Adya, C. Affeldt, *et al.*, *Physical review letters* **123**, 161102 (2019).
- [178] Planck Collaboration, *A&A* **594**, A13 (2016).
- [179] E. F. Arias, P. Charlot, M. Feissel, and J. F. Lestrade, *Astronomy and Astrophysics* **303**, 604 (1995).
- [180] J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff, S. Vitale, B. Aylott, K. Blackburn, N. Christensen, M. Coughlin, W. Del Pozzo, F. Feroz, J. Gair, C.-J. Haster, V. Kalogera, T. Littenberg, I. Mandel, R. O'Shaughnessy, M. Pitkin, C. Rodriguez, C. Röver, T. Sidery, R. Smith, M. Van Der Sluys, A. Vecchio, W. Vousden, and L. Wade, *Phys. Rev. D* **91**, 042003 (2015).
- [181] R. Abbott *et al.* (LIGO Scientific, Virgo), *SoftwareX* **13**, 100658 (2021), arXiv:1912.11716 [gr-qc].
- [182] C. Cutler *et al.*, *Phys. Rev. Lett.* **70**, 2984 (1993), arXiv:astro-ph/9208005.
- [183] C. Messick, K. Blackburn, P. Brady, P. Brockill, K. Cannon, R. Cariou, S. Caudill, S. J. Chamberlin, J. D. Creighton, R. Everett, *et al.*, *Physical Review D* **95**, 042001 (2017).
- [184] K. Cannon, R. Cariou, A. Chapman, M. Crispin-Ortuzar, N. Fotopoulos, M. Frei, C. Hanna, E. Kara, D. Keppel, L. Liao, *et al.*, *The Astrophysical Journal* **748**, 136 (2012).
- [185] K. Cannon, S. Caudill, C. Chan, B. Cousins, J. D. Creighton, B. Ewing, H. Fong, P. Godwin, C. Hanna, S. Hooper, R. Huxford, R. Magee, D. Meacher, C. Messick, S. Morisaki, D. Mukherjee, H. Ohta, A. Pace, S. Privitera, I. de Ruiter, S. Sachdev, L. Singer, D. Singh, R. Tapia, L. Tsukada, D. Tsuna, T. Tsutsui, K. Ueno, A. Viets, L. Wade, and M. Wade, *SoftwareX* **14**, 100680 (2021).
- [186] N. V. Krishnendu and F. Ohme, *Phys. Rev. D* **105**, 064012 (2022), arXiv:2110.00766 [gr-qc].
- [187] J. Calderón Bustillo, P. Laguna, and D. Shoemaker, *Phys. Rev. D* **95**, 104038 (2017).
- [188] S. Klimenko, G. Vedovato, M. Drago, G. Mazzolo, G. Mitselmakher, C. Pankow, G. Prodi, V. Re, F. Salemi, and I. Yakushin, *Physical Review D* **83** (2011), 10.1103/physrevd.83.102001.
- [189] N. Andersson, J. Baker, K. Belczynski, S. Bernuzzi, E. Berti, L. Cadonati, P. Cerdá-Durán, J. Clark, M. Favata, L. S. Finn, and et al., *Classical and Quantum Gravity* **30**, 193002 (2013).
- [190] J. Veitch and A. Vecchio, *Phys. Rev. D* **81**, 062003 (2010), arXiv:0911.3820 [astro-ph.CO].
- [191] R. Trotta, *Contemp. Phys.* **49**, 71 (2008), arXiv:0803.4089 [astro-ph].
- [192] M. Isi, R. Smith, S. Vitale, T. J. Massinger, J. Kanner, and A. Vajpeyi, *Phys. Rev. D* **98**, 042007 (2018), arXiv:1803.09783 [gr-qc].
- [193] K. Cannon, C. Hanna, and J. Peoples, arXiv (2015), arXiv:1504.04632 [astro-ph.IM].
- [194] B. Margalit and B. D. Metzger, *Astrophys. J. Lett.* **850**, L19 (2017), arXiv:1710.05938 [astro-ph.HE].
- [195] M. Ruiz, S. L. Shapiro, and A. Tsokaros, *Phys. Rev. D* **97**, 021501 (2018), arXiv:1711.00473 [astro-ph.HE].
- [196] H. T. Cromartie *et al.* (NANOGrav), *Nature Astron.* **4**, 72 (2019), arXiv:1904.06759 [astro-ph.HE].
- [197] J. M. Lattimer, *Universe* **5** (2019).
- [198] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. Lett.* **119**, 161101 (2017), arXiv:1710.05832 [gr-qc].
- [199] V. Doroshenko, V. Suleimanov, G. Pühlhofer, and A. Santangelo, *Nature Astronomy* **6**, 1444 (2022).
- [200] H. Sotani, K. Iida, K. Oyamatsu, and A. Ohnishi, *PTEP* **2014**, 051E01 (2014), arXiv:1401.0161 [astro-ph.HE].
- [201] K. Jedamzik, *JCAP* **09**, 022 (2020), arXiv:2006.11172 [astro-ph.CO].
- [202] K. Jedamzik, *Phys. Rev. Lett.* **126**, 051302 (2021), arXiv:2007.03565 [astro-ph.CO].
- [203] S. Clesse and J. García-Bellido, *Physics of the Dark Universe* **38**, 101111 (2022).
- [204] G. Hütsi, M. Raidal, V. Vaskonen, and H. Veermäe, *JCAP* **03**, 068 (2021), arXiv:2012.02786 [astro-ph.CO].
- [205] G. Franciolini, V. Baibhav, V. De Luca, K. K. Y. Ng, K. W. K. Wong, E. Berti, P. Pani, A. Riotto, and S. Vitale, (2021), arXiv:2105.03349 [gr-qc].
- [206] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. Lett.* **121**, 231103 (2018), arXiv:1808.04771 [astro-ph.CO].
- [207] A. H. Nitz and Y.-F. Wang, *Phys. Rev. Lett.* **126**, 021103 (2021), arXiv:2007.03583 [astro-ph.HE].
- [208] A. H. Nitz and Y.-F. Wang, *Phys. Rev. D* **106**, 023024 (2022).
- [209] C. T. Byrnes, M. Hindmarsh, S. Young, and M. R. S. Hawkins, *JCAP* **1808**, 041 (2018), arXiv:1801.06138 [astro-ph.CO].
- [210] K. Jedamzik, *Phys. Rev. D* **55**, 5871 (1997), arXiv:astro-ph/9605152.

- [211] J. C. Niemeyer and K. Jedamzik, Phys. Rev. Lett. **80**, 5481 (1998), arXiv:astro-ph/9709072.
- [212] J. Garcia-Bellido, H. Murayama, and G. White, JCAP **12**, 023 (2021), arXiv:2104.04778 [hep-ph].
- [213] A. Escrivà, E. Bagui, and S. Clesse, Journal of Cosmology and Astroparticle Physics **2023**, 004 (2023).
- [214] G. Franciolini, I. Musco, P. Pani, and A. Urbano, Phys. Rev. D **106**, 123526 (2022).
- [215] S. Jaraba and J. García-Bellido, Phys. Dark Univ. **34**, 100882 (2021), arXiv:2106.01436 [gr-qc].
- [216] D. Singh, M. Ryan, R. Magee, T. Akhter, S. Shandera, D. Jeong, and C. Hanna, Phys. Rev. D **104**, 044015 (2021), arXiv:2009.05209 [astro-ph.CO].
- [217] S. L. Liebling and C. Palenzuela, Living Rev. Rel. **15**, 6 (2012), arXiv:1202.5809 [gr-qc].
- [218] N. E. Wolfe, S. Vitale, and C. Talbot, (2023), arXiv:2305.19907 [astro-ph.HE].
- [219] T. E. H. T. Collaboration, K. Akiyama, A. Alberdi, *et al.*, The Astrophysical Journal Letters **875**, L1 (2019).
- [220] Wyrzykowski, L., Kruszyńska, K., Rybicki, K. A., Holl, B., Lecœur-Taïbi, I., Mowlavi, N., Nienartowicz, K., Jevardat de Fombelle, G., Rimoldini, L., Audard, M., Garcia-Lario, P., Gavras, P., Evans, D. W., Hodgkin, S. T., and Eyer, L., Astronomy & Astrophysics **674**, A23 (2023).
- [221] J. R. Chisholm, Physical Review D **73** (2006), 10.1103/physrevd.73.083504.
- [222] J. R. Chisholm, Physical Review D **84** (2011), 10.1103/physrevd.84.124031.
- [223] J. M. Ezquiaga, J. García-Bellido, and V. Vennin, JCAP **03**, 029 (2020), arXiv:1912.05399 [astro-ph.CO].
- [224] J. M. Ezquiaga, J. García-Bellido, and V. Vennin, Phys. Rev. Lett. **130**, 121003 (2023), arXiv:2207.06317 [astro-ph.CO].
- [225] C. Animali and V. Vennin, (2024), arXiv:2402.08642 [astro-ph.CO].
- [226] K. M. Belotsky, V. I. Dokuchaev, Y. N. Eroshenko, E. A. Esipova, M. Y. Khlopov, L. A. Khromykh, A. A. Kirillov, V. V. Nikulin, S. G. Rubin, and I. V. Sadvkovsky, The European Physical Journal C **79** (2019), 10.1140/epjc/s10052-019-6741-4.
- [227] S. G. Rubin, A. S. Sakharov, and M. Y. Khlopov, Journal of Experimental and Theoretical Physics **92**, 921 (2001).
- [228] M. Gieles, D. Erkal, F. Antonini, E. Balbinot, and J. Peñarrubia, Nature Astronomy **5**, 957 (2021).
- [229] M. Trashorras, J. García-Bellido, and S. Nesseris, Universe **7**, 18 (2021), arXiv:2006.15018 [astro-ph.CO].
- [230] D. Inman and Y. Ali-Haïmoud, Phys. Rev. D **100**, 083528 (2019), arXiv:1907.08129 [astro-ph.CO].
- [231] S. J. Aarseth, *Gravitational N-Body Simulations: Tools and Algorithms*, Cambridge Monographs on Mathematical Physics (Cambridge University Press, 2003).
- [232] R. Spurzem, Journal of Computational and Applied Mathematics **109**, 407 (1999).
- [233] E. Bagui *et al.* (LISA Cosmology Working Group), (2023), arXiv:2310.19857 [astro-ph.CO].
- [234] B. Carr, S. Clesse, J. García-Bellido, M. Hawkins, and F. Kuhnel, Phys. Rept. **1054**, 1 (2024), arXiv:2306.03903 [astro-ph.CO].
- [235] P. H.C., Monthly Notices of the Royal Astronomical Society **71**, 460 (1911).
- [236] M. Y. Khlopov, S. G. Rubin, and A. S. Sakharov, “Strong primordial inhomogeneities and galaxy formation,” (2002).
- [237] G. Fragione and A. Loeb, New Astronomy **55**, 32 (2017).
- [238] S. von Hoerner, Astrophys. J. **125**, 451 (1957).
- [239] P. C. Peters, Phys. Rev. **136**, B1224 (1964).
- [240] S. Sigurdsson and L. Hernquist, Nature (London) **364**, 423 (1993).
- [241] S. F. P. Zwart and S. L. W. McMillan, The Astrophysical Journal **576**, 899 (2002).
- [242] D. C. Heggie, Monthly Notices of the Royal Astronomical Society **173**, 729 (1975).
- [243] X. Jiménez-Forteza, D. Keitel, S. Husa, M. Hannam, S. Khan, and M. Pürrer, Physical Review D **95** (2017), 10.1103/physrevd.95.064024.
- [244] C. O. Lousto, Y. Zlochower, M. Dotti, and M. Volonteri, Phys. Rev. D **85**, 084015 (2012).
- [245] C. O. Lousto, M. Campanelli, Y. Zlochower, and H. Nakano, Classical and Quantum Gravity **27**, 114006 (2010).
- [246] M. Campanelli, C. O. Lousto, Y. Zlochower, and D. Merritt, Phys. Rev. Lett. **98**, 231102 (2007).
- [247] M. Campanelli, C. O. Lousto, Y. Zlochower, and D. Merritt, Physical Review Letters **98** (2007), 10.1103/physrevlett.98.231102.
- [248] P. Madau and M. Dickinson, Annual Review of Astronomy and Astrophysics **52**, 415 (2014).
- [249] M. Cantiello, C. Mankovich, L. Bildsten, J. Christensen-Dalsgaard, and B. Paxton, Astrophys. J. **788**, 93 (2014), arXiv:1405.1419 [astro-ph.SR].
- [250] A. Maeder and G. Meynet, Annual Review of Astronomy and Astrophysics **38**, 143 (2000), <https://doi.org/10.1146/annurev.astro.38.1.143>.
- [251] J. Fuller, A. L. Piro, and A. S. Jermyn, Monthly Notices of the Royal Astronomical Society **485**, 3661 (2019), <https://academic.oup.com/mnras/article-pdf/485/3/3661/28230562/stz514.pdf>.
- [252] J. C. Bustillo, N. Sanchis-Gual, A. Torres-Forné, and J. A. Font, Phys. Rev. Lett. **126**, 201101 (2021).
- [253] J. C. Bustillo, N. Sanchis-Gual, A. Torres-Forné, J. A. Font, A. Vajpeyi, R. Smith, C. Herdeiro, E. Radu, and S. H. W. Leong, Phys. Rev. Lett. **126**, 081101 (2021).
- [254] V. Gayathri, J. Healy, J. Lange, B. O’Brien, M. Szczepanczyk, I. Bartos, M. Campanelli, S. Klimenko, C. Lousto, and R. O’Shaughnessy, “Eccentricity estimate for black hole mergers with numerical relativity simulations,” (2022), arXiv:2009.05461 [astro-ph.HE].
- [255] I. Romero-Shaw, P. D. Lasky, E. Thrane, and J. C. Bustillo, The Astrophysical Journal Letters **903**, L5 (2020).
- [256] I. M. Romero-Shaw, D. Gerosa, and N. Loutrel, Monthly Notices of the Royal Astronomical Society **519**, 5352 (2023).
- [257] P. C. Peters, Phys. Rev. **136**, B1224 (1964).
- [258] M. Maggiore, C. V. D. Broeck, N. Bartolo, E. Belgacem, D. Bertacca, M. A. Bizouard, M. Branchesi, S. Clesse, S. Foffa, J. García-Bellido, S. Grimm, J. Harms, T. Hinderer, S. Matarrese, C. Palomba, M. Peloso, A. Ricciardone, and M. Sakellariadou, Journal of Cosmology and Astroparticle Physics **2020**, 050–050 (2020).