

Efficient searches for spinning compact binaries with advanced gravitational-wave observatories

Von der Fakultät für Mathematik und Physik
der Gottfried Wilhelm Leibniz Universität Hannover
zur Erlangung des Grades

Doktor der Naturwissenschaften
— Dr. rer. nat. —

genehmigte Dissertation von

M. Sc. Tito Dal Canton

geboren am 7 Februar 1983 in Asolo, TV (Italien)

2015

Referent

Prof. Dr. Bruce Allen
Albert-Einstein-Institut
Leibniz Universität Hannover
University of Wisconsin - Milwaukee

Korreferenten

Prof. Dr. Gianluca Maria Guidi
Università degli Studi di Urbino, Italien

Dr. Stephen Fairhurst
Cardiff University

Tag der Promotion

10 Juli 2015

LIGO document number

LIGO-P1500177

Abstract

Compact astrophysical objects (neutron stars and black holes) are the result of the evolution and death of massive stars. Their extreme properties can not be directly reproduced in a laboratory, but can be probed experimentally by observing the interaction between such objects in binary systems. Binaries of neutron stars, as well as the likely progenitors of binaries containing black holes, can be observed electromagnetically. General relativity predicts that all compact binaries also radiate energy via gravitational waves; the energy loss leads to the eventual coalescence of the two objects. A promising way to gain more information about compact objects is the direct observation of their gravitational radiation, which so far has not succeeded and represents a major challenge of modern experimental physics. However, advanced interferometric gravitational-wave detectors such as LIGO, Virgo and KAGRA will take data with unprecedented sensitivity starting from this year. Due to the weakness of the gravitational interaction and the rarity of compact binary coalescence events, analysis pipelines for detecting such events in interferometer data must ensure that the search sensitivity is as high as permitted by the fundamental noise of the detectors.

This thesis identifies possible limits to the sensitivity of such pipelines and proposes ways to overcome those limits. The main chapters (2 and 3) focus on detecting coalescing binaries of neutron stars and spinning black holes. We show the importance of black-hole spins, which have been neglected in past searches, on the sensitivity of the search pipeline. We demonstrate that including the effects of the spin component parallel to the orbital angular momentum is technically feasible and leads to a more sensitive *aligned-spin* pipeline. The exact gain depends on the distribution of spin magnitudes in nature. We show in chapter 3 that the aligned-spin pipeline can also detect a large fraction of coalescing neutron-star—black-hole binaries whose spins are *not* aligned with the orbital angular momentum, i.e. whose orbital planes precess. Chapter 4 compares the sensitivity and computational cost of two alternative ways to combine data from multiple gravitational-wave detectors in search pipelines for binary coalescence. We show that the two methods have comparable sensitivity, although the difference will become meaningful when several interferometric detectors will operate at the same time. Chapter 5 investigates how the detection statistic used to search for binary coalescences is affected by instrumental, non-astrophysical transient signals, which are commonly observed in interferometer data. We demonstrate that such transients can affect the pipeline for several minutes and we propose ways to counteract this effect.

Keywords: compact binaries, gravitational waves, data analysis

Kurzfassung

Kompakte astrophysikalische Körper (Neutronensterne und schwarze Löcher) sind die Endzustände der Evolution massereicher Sterne. Ihre extremen Eigenschaften können nicht direkt in Laboren reproduziert werden, aber durch Beobachtung der Wechselwirkung zwischen solchen Objekten in Binärsystemen experimentell überprüft werden. Paare von Neutronensternen, genauso wie die vermutlichen Vorstadien von Binärsystemen, die schwarze Löcher enthalten, können elektromagnetisch beobachtet werden. Die allgemeine Relativitätstheorie sagt voraus, dass alle kompakten Binärsysteme Energie in Form von Gravitationswellen abstrahlen; der Energie-Verlust führt dann schlussendlich zur Verschmelzung dieser Objekte. Eine vielversprechende Methode um Informationen über diese kompakten Objekte zu bekommen, ist die Gravitationsstrahlung direkt zu detektieren, welche bisher allerdings noch nicht erfolgreich war und daher eine große Herausforderung der aktuellen Experimentalphysik darstellt. Indes werden dieses Jahr fortschrittliche interferometrische Gravitationswellendetektoren wie LIGO, Virgo und KAGRA starten Daten mit unerreichter Empfindlichkeit aufzunehmen. Aufgrund der schwachen gravitativen Wechselwirkung und der Seltenheit von Verschmelzungen kompakter Binärsysteme müssen Analysemethoden, die solche Ereignisse in Interferometer-Daten detektieren sollen, gewährleisten können, dass die Suchempfindlichkeit so hoch ist wie durch das fundamentale Rauschen in den Detektoren zugelassen.

In dieser Arbeit werden mögliche Limits der Empfindlichkeit solcher Methoden identifiziert und Wege vorgeschlagen, um diese zu übertreffen. Die Hauptkapitel (2 und 3) konzentrieren sich auf die Detektion verschmelzender Paare von Neutronensternen und rotierenden schwarzen Löchern. Wir zeigen die Wichtigkeit den Spins schwarzer Löcher, welcher in vergangenen Suchen vernachlässigt wurde, für die Empfindlichkeit der Suchmethoden. Dann zeigen wir, dass die Berücksichtigung der Effekte von der Spinkomponente parallel zum Orbitaldrehimpuls technisch durchführbar ist und zu einer empfindlicheren Suchmethode führt. Der tatsächliche Nutzen hängt von der Verteilung der Drehimpulse schwarzer Löcher in der Natur ab. In Kapitel 3 zeigen wir, dass wir durch die Berücksichtigung des Parallel-Spins auch einen Großteil der Verschmelzungen von Binärsystemen aus Neutronensternen und schwarzen Löchern detektieren können deren Spins nicht parallel zum Orbitaldrehimpuls sind, d.h. deren Orbitalebene präzediert. Kapitel 4 vergleicht die Empfindlichkeit und den Rechenaufwand zweier verschiedener Wege, die Daten mehrerer Gravitationswellendetektoren in Suchmethoden für Verschmelzungen zu kombinieren. Wir zeigen, dass die beiden Methoden eine vergleichbare Empfindlichkeit haben, obgleich der Unterschied bedeutsam werden wird, wenn mehrere interferometrische Detektoren zur gleichen Zeit in Betrieb sein werden. Kapitel 5 untersucht, wie die Suche nach Verschmelzungen durch kurzlebige nicht-astrophysikalische Signale aus den Instrumenten, welche häufig in Interferometer-Daten beobachtet werden, beeinflusst wird. Wir zeigen, dass diese Signale die Suchmethoden für mehrere Minuten beeinflussen können und schlagen Wege vor, um diesem Effekt entgegenzuwirken.

Schlagworte: kompakte Binärsysteme, Gravitationswellen, Datenanalyse

Contents

1	Introduction	9
1.1	Stellar evolution	9
1.1.1	Formation and main sequence	9
1.1.2	Death as a white dwarf	10
1.1.3	Collapse to a neutron star	10
1.1.4	Collapse to a black hole	11
1.2	Compact binaries	13
1.2.1	Mass exchange	14
1.2.2	Supernova explosion	15
1.2.3	Early inspiral	15
1.2.4	Late inspiral, spins and precession	19
1.2.5	Coalescence	21
1.2.6	Standard evolutionary scenario	22
1.2.7	Coalescence rate	24
1.3	Studying compact binaries via gravitational radiation	25
1.3.1	Motivation	25
1.3.2	Experimental setup	26
1.3.3	Data model	28
1.3.4	Matched filtering and template banks	31
1.3.5	Signal-based vetoes	34
1.3.6	Multiple detectors: coincidence	35
1.3.7	Multiple detectors: coherent combination	37
1.3.8	Statistical significance of candidates	38
1.3.9	Results of past searches	39
2	Search pipeline for NSBH coalescence including aligned-spin effects	41
2.1	Overview	41
2.2	Simulation of the source population	43
2.3	Simulated data set	45
2.4	The PyCBC toolkit	46
2.5	Template bank construction and test	48
2.6	Pipeline configuration and execution	52
2.6.1	Matched filtering and trigger generation	57

2.6.2	Coincidence	58
2.6.3	Clustering	60
2.7	False-alarm background	62
2.8	Recovering simulated signals	65
2.9	Search sensitivity at fixed false-alarm rate	67
2.10	Conclusion	73
3	Effect of neglecting precession in NSBH search pipelines	75
3.1	Overview	75
3.2	Impact of precession on visibility of signals	77
3.3	Recovering precessing signals in realistic data	81
3.4	Search sensitivity at fixed false-alarm rate	82
3.5	Comparison with an idealized precessing search	84
3.6	Conclusion	88
4	Sensitivity of coincident and coherent CBC searches at fixed computational cost	91
4.1	Overview	91
4.2	Detector networks and template banks	92
4.3	False-alarm rate	95
4.3.1	Coincident search	95
4.3.2	Coherent search	99
4.4	Computational cost	100
4.5	Sensitivity	102
4.6	Hierarchical method	105
4.7	Conclusion	107
5	Modeling the effect of glitches on CBC searches	113
5.1	Overview	113
5.2	Glitch model and matched-filter approximation	114
5.2.1	Approximation I	115
5.2.2	Approximation II	116
5.2.3	Approximation III	117
5.3	Numerical simulations	118
5.4	Conclusion	124
	Bibliography	135
	Acknowledgments	137

Chapter 1

Introduction

This chapter gives an introduction to the topic of compact binaries and the observation of their coalescence via gravitational waves (GWs), providing the necessary context for the research presented in the next chapters. Section 1.1 is a quick summary of the evolution and final state of isolated stars and introduces compact objects. Section 1.2 reviews the current models and existing observations of compact binaries. Section 1.3 reviews the motivation and the experimental and data-analysis techniques behind the observation of compact-binary coalescence (CBC) through the associated GW emission.

1.1 Stellar evolution

Before introducing compact binaries, we quickly review the most important phases of stellar evolution and the basic properties of the resulting compact objects. For details, refer to classical textbooks such as [1].

1.1.1 Formation and main sequence

The formation of a star begins with the gravitational contraction of hydrogen in molecular clouds. Initially, the gas appears to form long and thin filament-like structures. Depending on their linear mass density, filaments may then undergo gravitational instability and fragment into protostellar cores. Further gravitational contraction increases the density and temperature of the gas in a core, eventually starting nuclear fusion reactions. These are typically exothermic and release high-energy photons. The resulting radiation pressure stops further gravitational contraction and the system enters a steady state known as a main sequence star.

Stars spend the majority of their life in the main sequence, fusing hydrogen into helium (through the proton-proton reaction or CNO cycle, depending on the mass of the star) as well as producing a smaller proportion of heavier elements. Heavier stars have typically higher reaction rates and thus shorter life. As helium accumulates at the core of the star, new fusion reactions may start converting it into heavier elements, preventing the gravitational collapse of the helium core. Thus, the interior of the star

may develop an onion-like structure with heavier elements accumulating towards its center. When the luminosity and surface temperature of stars are plotted one against the other (*Hertzsprung-Russell diagram*) stars form well-separated clusters, the most prominent one representing main-sequence stars.

Fusion reactions become less efficient as elements further down in the periodic table are formed, and they stop altogether as soon as they become endothermic or fuel is exhausted. Radiation pressure is then no longer able to contrast gravitational collapse. At this point the evolution of the star can take different paths, depending mainly on the initial mass and chemical composition and presumably on the intrinsic angular momentum.

1.1.2 Death as a white dwarf

If the initial mass of the star is sufficiently small, fusion leaves an inert C-O or O-Ne core. Such an object can no longer produce nuclear fusion, but its contraction is balanced by the *degeneracy pressure* induced by its electrons behaving as a Fermi gas. This equilibrium is realized at roughly the size of the Earth and a mass between fractions of the solar mass ($M_{\odot} \approx 1.99 \times 10^{30}$ kg) and $\approx 1.4M_{\odot}$ (Chandrasekhar limit) i.e. at an extremely large density. Although exothermic reactions no longer occur, its initial temperature still causes the object to radiate high-energy photons as a *white dwarf (WD)*, which cools down on a time scale longer than the expected duration of the universe.

1.1.3 Collapse to a neutron star

Initial masses below $\approx 20M_{\odot}$ result in either an oxygen-neon- magnesium or iron core whose gravitational collapse can no longer be halted by electron degeneracy. The object then undergoes electron capture, converting most proton-electron pairs into neutrons and collapsing violently on a time scale of seconds. The outer infalling material is thought to bounce on the newly formed neutron core, producing a *supernova explosion* which includes a burst of neutrinos from the neutronization process. The resulting collapsed core is essentially a degenerate, compact sphere of neutrons with density comparable to atomic nuclei, known as a *neutron star (NS)*. Conservation of the angular momentum of the collapsing core, or asymmetries in the supernova explosion [2], are thought to lead to an extremely fast rotation of the NS.

The observation of periodic radio pulses with excellent stability and periods between seconds and milliseconds (*pulsars*) suggests a rapidly rotating and extremely compact source emitting beamed radiation; spinning NSs with magnetic dipoles misaligned with the rotation axis are the best explanation [3]. Thousands of NSs have been discovered so far in this way [4]. Pulsars have also been observed in optical and γ bands [5, 6]. The emission mechanisms, however, are not yet firmly understood.

A review of NS masses can be found in [3]. Theoretical calculations of NS masses are difficult due to the unknown equation of state (EOS) of neutron matter. However, maintaining hydrostatic equilibrium in General Relativity (GR) requires an upper bound on the mass. Imposing the condition that the speed of sound does not exceed the speed of

light c (maximum stiffness) requires a maximum mass of $\approx 3M_\odot$ [7]. Measuring the mass of an isolated NS is also difficult. Fortunately however, when a NS is orbiting a companion star, observing the companion’s spectral lines allows one to constrain the Keplerian parameters and thus a combination of the masses (mass function). The constraint is improved if the NS is a pulsar and the Doppler shift of its pulses can be observed. When the NS is observed as a pulsar in a relativistic binary, and post-Keplerian parameters can be measured, the constraint is very tight and can lead to exquisite precision in the mass measurement. Precisely measured NS masses currently range from $1.2M_\odot$ to $2M_\odot$ with a narrow cluster at $1.35M_\odot$ for NSs in binaries [8, 9].

By definition, the rotation frequency of pulsars can be measured very precisely by timing the pulses; frequencies in the range $10^{-1} - 10^3$ Hz are typical. The intrinsic angular momentum (spin) \mathbf{S} of a NS can be characterized by the dimensionless vector

$$\boldsymbol{\chi} := \frac{c}{G} \frac{\mathbf{S}}{m^2} \quad (1.1)$$

with G and c the gravitational constant and speed of light in vacuum and m the NS mass. The dimensionless spin magnitude is limited by the break-up velocity of the NS to $\chi \lesssim 0.7$ [10]. Observed NSs are however far from this limit, with the fastest-spinning one at only $\chi \approx 0.4$ [11]; possible explanations for the lack of higher spins are r-mode instabilities and braking by GWs emission [12].

Reproducing the extreme conditions of NS matter in the laboratory is currently unfeasible. Thus its properties remain largely unknown; in particular, there are many possible models for its EOS. The EOS manifests itself macroscopically as a relation between the NS mass and radius and it also influences the NS deformability. EOS models can thus be tested via mass-vs-radius or deformability measurements. Unfortunately, obtaining useful measurements of the radius or deformation via electromagnetic observations has proved to be challenging, severely limiting our ability to constrain the available models.

1.1.4 Collapse to a black hole

Progenitors with mass above $\approx 20M_\odot$ undergo a collapse so violent that no known physical interaction can counter gravity and stop the contraction. GR predicts that this results in a *black hole (BH)*, a consequence of the Schwarzschild and Kerr solutions to Einstein’s equations [13]. In a static BH of mass m , matter and radiation are entirely confined within the *Schwarzschild radius*

$$r_s := 2\frac{G}{c^2}m. \quad (1.2)$$

The inside of a BH is effectively prevented from making any causal communication with the rest of the universe. Uniqueness theorems imply that a stationary BH is an extremely simple object, described only by its mass m , charge Q (expected to vanish for astrophysical BHs) and dimensionless spin $\boldsymbol{\chi}$ as in eq. (1.1). The assumption that the singularity of the Kerr solution is always covered by an horizon (*cosmic censorship*)

requires that BHs have $\chi \leq 1$; BHs can spin up to almost this limit by accreting matter [14]. Note that NSs (as well as non-compact objects) need not obey a fundamental spin limit.

Because of their very nature, observing BHs with electromagnetic radiation is difficult. There are currently no unambiguous observations of BHs. However, the gravitational interaction of a BH with the surrounding matter and radiation can be a very strong (albeit indirect) indication of its existence. Matter falling into a BH can form an accretion disk, where potential energy is converted very efficiently into heat causing the inner part of the disk to radiate X-rays. Observations of several X-ray binaries—notably Cygnus X-1 [15]—support this model and currently provide the best evidence for BHs produced by the collapse of massive stars. On the other hand, long-term tracking of stars in Sagittarius A* demonstrates the existence of a massive, dark and relatively small object at the galactic center, supporting the hypothesis that galaxies host *supermassive* BHs which did not directly result from the collapse of single stars. BHs could in principle also be observed as gravitational lenses, although no such observation is available yet [16].

Contrary to NSs, there are no fundamental limits to BHs masses; for instance, Sagittarius A* hosts a $4 \times 10^6 M_\odot$ BH [17, 18]. X-ray binaries whose Keplerian parameters can be constrained enable a measurement of their BH mass in a similar way as for NSs, although the typical precision is far from what is possible for NSs [3, 19]. Therefore, the distribution of BH birth masses is very uncertain. There is the possibility of a mass gap between NSs and BHs which could reflect different core-collapse scenarios [20]. It is not clear, however, whether this is a consequence of selection effects. This thesis will focus on BHs with masses between a few M_\odot and $\mathcal{O}(10M_\odot)$, i.e. those found in X-ray binaries, and will refer to them as *stellar-mass* BHs. An updated list of mass measurements for such objects is maintained at [8].

With important caveats, some of the available BH observations also enable spin measurements [3]. The X-ray spectrum of the accretion disk in a stellar-mass X-ray binary depends on the BH spin and can yield a spin measurement via the *continuum fitting* method [21]. However, this measurement can be influenced by the relative orientation of the spin and the disk [22]. The spin also induces a broadening of the iron $K\alpha$ emission line which is in some cases present in the spectrum of X-ray binaries, enabling a different spin measurement via spectroscopy [23]. A third method involves the interpretation of observed quasi-periodic oscillations in the X-ray flux in terms of the characteristic frequencies of the system [24, 25]. These methods now consistently indicate that a large fraction of BHs may be spinning close to the Kerr limit. Simulations modeling the formation of BHs in binaries also support this hypothesis [26]. Note however that some high-spin measurements have also been criticized [27] and alternative explanations of the data based on unexpected deviations from Kerr BHs have been proposed [28].

Mass and spin parameters for some of the observed stellar-mass BH candidates are plotted in figure 1.1.




Figure 1.1: Measured parameters for some of the observed stellar-mass BH candidates. Data from [8, 29–35].

1.2 Compact binaries

We have seen that, independently of the details, the generally accepted final outcome of a medium to high-mass star is always a *compact object* of relatively small radius (possibly tens of kilometers) and mass comparable to M_{\odot} or significantly larger. Bringing macroscopic amounts of matter to similar extreme regimes in a controlled laboratory experiment, let alone producing a BH, is currently impossible. Astrophysical compact objects are thus a unique and formidable opportunity to perform experiments with matter in such regimes. A promising way of extracting information from compact objects is observing their close interaction and nature already provides us with the necessary setup, as we shall see in this section.

Star formation often produces multiple-star systems [36]. Many studies estimate that a large fraction of stars are in binary systems, although these estimates are generally affected by selection biases. The evolutionary model described in the previous section generally also applies to each component of a binary, at least qualitatively, but its details can be greatly influenced by the presence of the companion star. The main parameter controlling the fate of the components and of the binary itself is the separation a between the components. In widely-separated binaries, whose a is much larger than the radii of the component stars, the two components essentially evolve independently and the only phenomenon which can dramatically affect the binary is the eventual collapse or explosion of the components. Close binaries, instead, enable qualitatively different and

interesting physical processes which greatly influence the fate of the components and the binary, in particular mass exchange, orbital decay by gravitational radiation and collision between the components. When both components of a close binary are compact objects we have a *compact binary*. Observing the distribution and behavior of such systems provides important data about compact objects and the processes leading to their formation.

The most important physical phenomena involving binaries, and the currently available observations and open issues, are summarized in the next subsections. For a full review of compact binary evolution, see [37].

1.2.1 Mass exchange

The shape of the stellar surfaces in a binary is largely determined by the effective potential in which they sit. This is usually approximated by the *Roche potential*, which includes the gravitational interaction of the components (assumed to be point masses corotating with the orbit) and the centrifugal force due to the orbital angular momentum (assuming a circular orbit) but neglects radiation pressure. The Roche potential has an equipotential surface containing both components, whose shape is composed of two *critical lobes* joined at the equilibrium saddle point along the separation between the components. If one of the components becomes larger than the corresponding lobe, matter is no longer bound to its surface and can be transferred to the other object or ejected from the system in different ways.

On the one hand, matter can simply be ejected from the system as a *fast wind* which interacts very little with the companion. The angular momentum loss associated with this fast-wind or *Jeans* mode can be calculated easily by assuming the components are point masses. The result is always a widening of the binary.

On the other hand, all stripped matter can fall from the donor star directly onto the companion. This *conservative accretion* mode preserves the total mass and orbital angular momentum. In principle, modeling the stream of matter requires numerical 3D hydrodynamics and possibly taking into account nuclear reactions. However, the same simplified point-mass model used for the *Jeans* mode shows that the binary can either shrink or widen, depending on the relative mass of the donor and accretor.

Matter stripped from the donor and falling towards the companion can fail to accrete completely, for instance if the Eddington luminosity is exceeded, and be ejected from the system (*isotropic reemission*). An observation compatible with this mode is the X-ray binary SS 433. The simple point-mass model can again be used to estimate the variation in orbital separation. The result depends strongly on the mass ratio and can be a dramatic shrinkage of the binary.

In fact, isotropic reemission can shrink the binary so much that the point-mass model is no longer valid and instabilities in mass transfer form. This can result in a *common envelope* engulfing the whole binary, with the accretor spiraling toward the core of the donor. In general, common envelope can result from any mechanism which dramatically reduces the orbital separation and brings one component into the other. The time scale of the common-envelope stage is expected to be years and the occurrence

of common-envelope events is thought to be rare, so its direct observation is unlikely. However, several close compact binaries ($a \approx R_\odot$) are known from observations—notable examples are the “Hulse-Taylor pulsar” B1913+16 [38] and the “double pulsar” J0737-3039 [39]—and we know that the formation of a compact object involves a progenitor with $r \gtrsim 100R_\odot$. Therefore, common envelopes appear to be a necessary and ubiquitous stage in the formation of close compact binaries. Unfortunately, studying their detailed formation, evolution and outcome requires a full 3D hydrodynamic treatment, possibly including nuclear reactions, which is computationally challenging. Simplified analytical models have been proposed based on balance of energy or angular momentum, but their parameters are poorly constrained and they are not consistent with all observed close binaries. As a result, common envelopes are still poorly understood and they are regarded as the most important open problem in the study of close binaries.

1.2.2 Supernova explosion

As a result of their independent evolution, or after accreting mass from the companion, components in a binary may explode as supernovae and turn into NSs or BHs. The time scale of the explosion is typically much shorter than the orbital period, so it is easy to model the phenomenon analytically. The result is an expression for the orbital parameters after the explosion in terms of the original ones, the mass ejected from the binary Δm and the change in velocity \mathbf{w} (*kick*) that the exploding object may receive as a result of asymmetries in the explosion.

The details of the explosion can greatly affect the fate of the binary. Some combinations of Δm and \mathbf{w} tend to increase the orbital radius and even disrupt the binary (e.g. $\Delta m \geq (m_1 + m_2)/2$ and $w \approx 0$). In other cases the binary can become tighter and its life can be significantly shortened, similarly to what happens with common envelopes. If the kick has a significant component orthogonal to the orbital plane, the orientation of the final orbit is affected and, if the objects have significant spins, this can then be important for precession effects, as we shall see later. Finally, large kicks can affect the spatial distribution of binaries by moving the systems away from the star-forming regions.

Unfortunately, birth kicks of NSs and BHs remain another unresolved problem in stellar evolution; at the moment there are observations and simulations supporting both large and small kicks and the symmetry of kicks during supernova explosions remains poorly understood.

1.2.3 Early inspiral

Systems surviving mass-exchange and supernova explosions as stellar-mass compact binaries are well approximated as detached point masses in Keplerian orbits, with separation comparable to the solar radius R_\odot . This evolutionary stage of compact binaries is thus fully driven by gravitation. In the Newtonian description of gravity, such systems would be in principle indefinitely stable. However, this picture is significantly modified by GR, which predicts that binary systems radiate energy and angular momentum as

GW, decaying with a spiral motion called *inspiral*. It is therefore important to incorporate this effect into the evolution of compact binaries; conversely, observing detached compact binaries enables probing GR effects.

GWs are a major prediction of GR. Their detailed derivation and properties can be found in classical text books (e.g. [40]) and are briefly summarized here. Start from Einstein's equations,

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu} \quad (1.3)$$

where $R_{\mu\nu}$ and R are the Ricci tensor and scalar describing the spacetime curvature, $g_{\mu\nu}$ is the metric tensor describing the distance between infinitesimally-close events and $T_{\mu\nu}$ is the tensor describing the local energy density and momentum density and flux. Assume a spacetime described by the following metric,

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \quad (1.4)$$

where $\eta_{\mu\nu}$ is the flat Minkowski metric and $h_{\mu\nu}$ is a small perturbation, i.e. it has the property $|h_{\mu\nu}| \ll 1$ at least in a particular reference frame and in a sufficiently restricted region of spacetime. Insert this perturbed metric into eq. (1.3), and neglect terms of second and higher orders in $h_{\mu\nu}$. The result is a linearized equation for $h_{\mu\nu}$. Under a particular gauge choice (*Lorenz* or *harmonic* gauge) this equation reads

$$\square \bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4}T_{\mu\nu} \quad (1.5)$$

where $\square := \partial_\mu \partial^\mu = -c^{-2} \partial_0^2 + \nabla^2$ is the d'Alembertian operator in flat space and $\bar{h}_{\mu\nu} := h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}\eta^{\rho\sigma}h_{\rho\sigma}$.

Far from the source $T_{\mu\nu}$ vanishes and the components of the perturbation admit wave-like solutions called *linearized GWs*. The propagation speed turns out to be c like for electromagnetic waves. In principle, $h_{\mu\nu}$ has 10 independent components, but they can always be reduced to two by fixing the remaining gauge freedom. In this so-called *transverse and traceless* gauge, the temporal components and the trace of the perturbation vanish. If we further restrict our attention to plane-wave solutions, which are excellent approximations to the GWs generated by astrophysical systems and observed on Earth, and let the wave vector define the z axis, the perturbation can be written as

$$\begin{aligned} h_{\mu\nu}^{TT}(t, z) &= \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+(t-z/c) & h_\times(t-z/c) & 0 \\ 0 & h_\times(t-z/c) & -h_+(t-z/c) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}_{\mu\nu} \\ &= \mathcal{E}_{\mu\nu}^+ h_+(t-z/c) + \mathcal{E}_{\mu\nu}^\times h_\times(t-z/c) \end{aligned} \quad (1.6)$$

The remaining h_+ and h_\times components represent physical degrees of freedom of the gravitational radiation, corresponding to the amplitudes of two independent polarization states, similarly to electromagnetic radiation. These states are described by the two

tensors $\mathcal{E}_{\mu\nu}^+$ and $\mathcal{E}_{\mu\nu}^\times$ and are commonly called *plus* and *cross* because of the geometry of their effect on test masses. In analogy with electromagnetic waves, each state can be transformed into the other by a rotation around the wave vector; however, the required angle is $\pi/4$ in the GW case and $\pi/2$ in the electromagnetic case. In our particular coordinate system, the only nonzero components of the perturbation are in the xy plane; in fact, the perturbation is always transverse to the wave vector, which is again a similarity with electromagnetism.

In linearized theory, GW emission from a source of finite size can be described, under the assumption that the gravitational field of the source is weak, using the Green's function and retarded potential. GWs produce a back-reaction on the source, subtracting energy and angular momentum from it and transporting them to infinity. The solution for an arbitrary source is complicated, but if the internal velocities of the source are non-relativistic, then the spatial components h_{ij}^{TT} of $h_{\mu\nu}^{TT}$ can be expressed in terms of the lowest multipole moments of the energy and linear-momentum density of the source. In particular, the radiation is dominated by the mass quadrupole moment,

$$h_{ij}^{TT}(t, \mathbf{x}) \approx \frac{2G}{c^4} \frac{1}{r} \ddot{Q}_{ij}^{TT}(t_{\text{ret}}) \quad (1.7)$$

where r is the distance from the source, $t_{\text{ret}} := t - r/c$ is the retarded time and Q_{ij}^{TT} is the transverse-traceless projection of the mass quadrupole tensor, which for a non-relativistic source of mass density $\rho(t, \mathbf{x})$ is

$$Q_{ij}^{TT}(t) \approx \left[\int_{\text{source}} d^3x \rho(t, \mathbf{x}) \left(x^i x^j - \frac{1}{3} r^2 \delta^{ij} \right) \right]^{TT}. \quad (1.8)$$

Note that there is no radiation associated with a time-varying mass monopole or dipole, meaning that an object undergoing a spherically-symmetric collapse or explosion does not radiate. On the other hand, a large, fast-rotating and non-axisymmetric mass distribution has a large time-varying quadrupole moment and is an efficient radiator. Perfect examples of such a system in nature are rapidly-rotating NSs deformed in a non-axisymmetric fashion and compact binaries after the common-envelope phase, which brings us back to the evolution of compact binaries.

The fundamental features of the gravitational radiation emitted by a compact binary can be derived by modeling the binary as two point masses m_1, m_2 in a non-relativistic circular Keplerian orbit [41]. In this regime the GW emission has a negligible effect on the trajectories of the masses on an orbital time scale. By calculating the quadrupole moment of the system and using it in eq. (1.7), the resulting GW polarization amplitudes for a distant observer are

$$\begin{aligned} h_+(t) &= A \frac{1 + \cos^2 \iota}{2} \cos(2\pi f_{\text{gw}} t_{\text{ret}} + 2\phi) \\ h_\times(t) &= A \cos \iota \sin(2\pi f_{\text{gw}} t_{\text{ret}} + 2\phi) \\ A &:= \frac{4}{r} \left(\frac{G}{c^2} \mathcal{M} \right)^{5/3} \left(\frac{\pi}{c} f_{\text{gw}} \right)^{2/3} \end{aligned} \quad (1.9)$$

where f_{gw} is twice the orbital frequency, $\mathcal{M} := (m_1 m_2)^{3/5} (m_1 + m_2)^{-1/5}$ is the *chirp mass* of the binary, ι is the angle between the orbital angular momentum and the line of sight from the source to the observer (*inclination*) and ϕ is a reference orbital phase. Due to GW radiation, the orbit slowly loses energy; by equating the total radiated power to the time derivative of the total energy of the Keplerian orbit, the evolution of the GW frequency is found to be

$$f_{\text{gw}}(t) = \frac{1}{\pi} \left(\frac{5}{256} \frac{1}{t_c - t} \right)^{3/8} \left(\frac{G}{c^3} \mathcal{M} \right)^{-5/8} \quad (1.10)$$

where t_c is the *coalescence time* i.e. the instant at which, under these approximations, frequency becomes formally infinite and the masses touch each other. Plugging the frequency evolution back into the polarization amplitudes gives

$$\begin{aligned} h_+(t) &= A' \frac{1 + \cos^2 \iota}{2} \cos \Phi(t) \\ h_\times(t) &= A' \cos \iota \sin \Phi(t) \\ A'(t) &:= \frac{1}{r} \left(\frac{G}{c^2} \mathcal{M} \right)^{5/4} \left(\frac{5}{c(t_c - t)} \right)^{1/4} \\ \Phi(t) &:= -2 \left(5 \frac{G}{c^3} \mathcal{M} \right)^{-5/8} (t_c - t)^{5/8} + \phi_c \end{aligned} \quad (1.11)$$

where ϕ_c is the phase at coalescence. Under these approximations, the evolution of the binary under GW emission is therefore an *adiabatic* or *Newtonian inspiral*: the objects perform a sequence of Keplerian orbits with slowly-decreasing orbital period and radius, while the emitted GW amplitude and frequency slowly increase. The GW signal is thus called a *chirp*. The polarization of the radiation depends on the inclination of the source: in the *face on* or *face off* cases ($\iota = \{0, \pi\}$) the polarization is circular, in the edge-on case ($\iota = \pi/2$) it is linear and for intermediate inclinations it is elliptic. Although we assumed a circular orbit, the description is easily extended to an eccentric orbit; the resulting signal contains potentially all harmonics of the orbital frequency and it can be shown that the GW back-reaction tends to reduce the eccentricity over time [42]. In any case, the orbital radius eventually shrinks to less than the size of the components, the orbital velocity becomes relativistic and the two objects merge. The Newtonian inspiral model is no longer an accurate description at that point.

GWs have never been directly observed on Earth. However, through long-term electromagnetic timing of pulsars in binary systems, several inspiraling NS binaries have been discovered and found to decay at a rate consistent with GW radiation. Notable examples are the already-mentioned close binaries B1913+16 and J0737-3039. Thanks to oscillations in their X-ray flux, inspiraling WD binaries have also been discovered, such as J0651+2844 [43]. The known inspiraling binaries currently represent the only firm observational evidence of the existence of GWs.

1.2.4 Late inspiral, spins and precession

The last few minutes of a compact binary decaying by GW radiation involve objects with masses of at least M_{\odot} orbiting at relativistic speeds and in a strong gravitational field, with separation comparable to their radius. Although the inspiral picture remains valid, the approximations made in deriving the Newtonian inspiral are no longer good and the amplitude and phase of the radiation require corrections. Such corrections can be computed via the *post-Newtonian* (*pN*) framework. For a full review of pN theory applied to compact binary inspiral, see [44]; the main concepts are summarized here.

The total energy $E(v)$ and outgoing GW power $P(v)$ of the binary system are written as power series of the parameter $\epsilon \propto v^2/c^2$ where v is the orbital velocity and the Newtonian inspiral corresponds to the leading (0th-order) term of the expansion. As done for the Newtonian inspiral, conservation of energy is assumed and the balance equation

$$\frac{d}{dt}E(v) = -P(v) \quad (1.12)$$

can then be used to derive the relationship between velocity and time and thus between GW phase and time. The final result is an expression of the GW phase as a power series and can be directly used to compute the GW signal as a function of time or frequency. pN corrections also introduce higher-order terms in the amplitude of the GW signal, beyond the leading quadrupole term. These give rise to multiple chirps in the time-frequency plane, at the orbital frequency as well as its higher harmonics. Depending on the region of interest in the parameter space, computing the phase to high pN order but neglecting corrections to the amplitude can be sufficient for detecting CBC signals. This approximation is called a *restricted* pN waveform [45] and will be used in chapters 2 and 3.

One source of corrections to the phase is the back-reaction of the outgoing radiation on the motion of the source and the non-linear scattering of the radiation off itself and off the background spacetime curvature produced by the binary. Phase correction terms associated with these effects have been computed at the 1, 1.5, 2, 2.5, 3 and 3.5 pN orders [44, 46, 47]. These terms contain a different combination of the component masses than the chirp mass, namely the *symmetric mass ratio*

$$0 < \eta := \frac{m_1 m_2}{M^2} \leq \frac{1}{4} \quad (1.13)$$

where $M := m_1 + m_2$ is the total mass of the binary. This additional dependency breaks the mass degeneracy present when only the Newtonian order is considered. Thus, if the signal is measured with sufficient precision, the individual masses can in principle be inferred.

More corrections to the phase appear when one or both components of the binary have non-vanishing spins. In fact, terms arise from the interaction between each spin and the orbital angular momentum (“spin-orbit” terms, $\chi_i \cdot \mathbf{L}$), from the mutual interaction of the two spins (“spin-spin” terms of type $\chi_1 \cdot \chi_2$), from the coupling of each spin to itself (“spin-squared” terms of type χ_i^2) etc. Spin contributions to the GW phase due

to these effects have been computed to the 2.5 pN order [44, 45, 48–51] and recently to 3.5 pN order [52]; the dominant effect is the spin-orbit interaction, entering already at 1.5 pN order. Spin also induces a Newtonian quadrupole moment; the coupling between this and the mass monopole then produces an effective 2 pN “quadrupole-monopole” correction to the phase [53]. The importance of spin corrections for a CBC search will be investigated in chapter 2.

Finally, the phase receives additional corrections when the orbital radius becomes comparable with the finite size of the components and tidal interaction becomes important [54–56]. The deformability of a NS is likely determined by its EOS, so observing the final cycles of the signal can in principle constrain properties of NS matter.

pN effects can not only introduce corrections to the orbital phase and GW amplitude, but also cause the orientation of the orbital angular momentum \mathbf{L} and the component spin vectors \mathbf{S}_1 , \mathbf{S}_2 to vary in time, i.e. to *precess* [57]. In fact, if one or both components of the binary have significant spins ($S_1 > 0$ and/or $S_2 > 0$) and the spins form a nonzero angle with \mathbf{L} , then a symmetry of the system is broken. Spin-orbit coupling then induces a time-varying orientation of the vectors \mathbf{S}_1 , \mathbf{S}_2 and \mathbf{L} with respect to an inertial frame, while the direction of the total angular momentum $\mathbf{J} := \mathbf{L} + \mathbf{S}_1 + \mathbf{S}_2$ remains constant except for small radiation-reaction effects. This is an entirely post-Newtonian phenomenon, analogous to the Lense-Thirring or gravitomagnetic effect.

Typically the motion of \mathbf{L} associated with precession is a simple rotation around \mathbf{J} with a time scale much longer than the orbital period; this regime is called *simple precession* and has two observable consequences on the emitted GW waveform. On the one hand, a secular correction to the orbital phase arises. On the other hand, the periodically-wobbling orbit means that the inclination and polarization angles associated with the emitted GW waveform are not constants but periodic functions of time. The resulting waveform acquires a characteristic amplitude and phase modulation. As we shall see in chapter 3, simple precession can have consequences in GW searches.

Precession can result in a short epoch characterized by a qualitatively different motion of \mathbf{L} if the spins are large and initially approximately anti-aligned with \mathbf{L} . In fact, \mathbf{L} slowly shrinks due to GW emission while S_1 and S_2 are constant. At some point, therefore, \mathbf{J} will transition from being dominated by \mathbf{L} to being dominated by $\mathbf{S}_1 + \mathbf{S}_2$. However, since \mathbf{J} must be approximately conserved, it can not suddenly change orientation; \mathbf{L} instead needs not be conserved and thus is free to rapidly reorient itself. This phenomenon is called *transitional precession* and leads to an episode of complicated three-dimensional tumbling of the binary. Once the vectors have been reoriented and the transition ends, simple precession is typically restored.

In practice the pN energy-balance equation (1.12) can be solved in different ways, producing different models (*approximants*) for inspiral waveforms (see e.g. [58]). The *SpinTaylorT2* and *SpinTaylorT4* approximants use eq. (1.12) to express respectively dt/dv or dv/dt as a power series in v . This is then combined with the relationship between orbital phase, velocity and time

$$\frac{d\phi}{dt} = \frac{v^3}{M} \tag{1.14}$$

and integrated numerically to produce the GW phase as a function of time. SpinTaylorT models also integrate the precession equations for the angular momenta and can thus simulate precessing pN waveforms. The *TaylorF2* approximant uses instead the stationary-phase approximation (SPA) to explicitly express restricted, non-precessing pN waveforms in the frequency domain as

$$\tilde{h}_{+,\times}(f) = A_{+,\times}(f) \cos[\Psi_{+,\times}(f)] \quad (1.15)$$

where $\Psi_{\cdot}(f)$ are power series in f and $\log f$. *SpinTaylorF2* uses the same approach but also includes single-spin precession [59]. These—and several other—waveform approximants are implemented in the LALSimulation software library maintained by the LIGO-Virgo collaboration [60]. Chapters 2 and 3 will use the SpinTaylorT2 model for simulating signals and TaylorF2 for analyzing the data.

1.2.5 Coalescence

At some point the inspiral must terminate: the two objects fall towards each other with a *plunge* and then merge. This process involves full GR in its strong-field regime and is difficult to describe analytically from first principles. However, it is possible to study it by solving the full Einstein equations numerically. The so-called *3+1 formalism* is at the basis of modern numerical relativity (NR) codes [61–65].

In the binary black hole (BBH) case, it is found that the BHs simply merge into a single object with significant GW emission. The remaining object is thought to be an excited (deformed) BH which rapidly sheds its deformation by GW radiation and settles into a stationary Kerr BH. The resulting *ringdown* waveform can be described analytically via perturbation theory and consists of an exponentially-damped oscillation whose frequency depends on the parameters of the final BH.

When NSs are involved, however, the physics of the merger is greatly complicated by the presence of matter: tidal forces before the plunge may in fact disrupt one or both NSs, dispersing neutron-rich matter incoherently around the binary and completely invalidating the point-mass picture. Nevertheless, NR remains a powerful tool for studying NS mergers and predicting the resulting GW waveforms. Thanks to simulations, tidal disruption is thought to produce important electromagnetic emission from the system at different wavelengths, possibly an accretion disk around the remaining compact object and possibly a jet orthogonal to the disk due to relativistic beaming. For this reason, binary neutron star (BNS) and neutron-star–black-hole (NSBH) mergers are a strong candidate for explaining short gamma-ray bursts (GRBs), but despite a large number of such events have been observed, a firm proof of this connection is not yet available [66]. Decompressing neutron-rich matter released during NS disruption may undergo rapid neutron-capture processes (*r-processes*) resulting in abundant production of heavy elements [67]; thus, NS mergers could account for the observed abundance of neutron-rich heavy elements in the universe [68]. R-processes should also produce an infrared emission a few days after the merger (*kilonova*) which has indeed been observed [69]. BNS and NSBH mergers are also thought to produce an excited object; depending on

the progenitor masses and spins and on the NS EOS, this can be either directly a BH or an unstable NS above the maximum-NS-mass limit, which eventually also collapses to a BH [70–72].

Although NR simulations are invaluable for studying the merger and ringdown, they are computationally expensive, severely limiting the number of orbits that can be practically simulated before the merger. This makes it difficult to cross-check pN and NR results. Nevertheless, approximants which are able to fully simulate the inspiral, merger and ringdown parts of BBH waveforms have been developed. One approach, known as the *effective one-body (EOB)* model, combines pN theory, the self-force formalism and BH perturbation theory; the resulting waveform is calibrated to accurately match NR simulations [73, 74]. Another approach is to build a phenomenological model that smoothly connects the analytical pN inspiral to the analytical ringdown signal, producing the *IMRPhenom* models [75, 76]. Both methods have been recently extended to also include precession [77, 78]. Although EOB and IMRPhenom models are not built as solutions to Einstein’s equations, they enable computationally-efficient calculations of full inspiral-merger-ringdown (IMR) BBH waveforms which closely reproduce NR results. As is the case for pN approximants, implementations are available in the LALSimulation library [60]. Unfortunately, similar approximants taking into account matter effects are not available yet and therefore IMR models are not necessarily accurate for BNS and NSBH mergers.

1.2.6 Standard evolutionary scenario

The phenomena summarized in the previous subsections are part of a “standard” evolutionary scenario that emerged over the last few decades and is confirmed by many astronomical observations [37]. Its fundamental stages are visualized in figure 1.2 and summarized here.

Initially, two main-sequence massive stars are orbiting each other and well within their Roche lobes. Any eccentricity is rapidly dissipated by tidal interaction. The duration of this phase is determined by the evolution of the heavier star, which can take up to $\approx 10^7$ years to exhaust its hydrogen fuel. Thus a large number of binaries in this stage are expected in our galaxy.

As the heavier star runs out of hydrogen, a large helium core is left and the star expands rapidly, overflowing its Roche lobe. Mass starts spilling from the overflowing star onto the lighter one, which is still in the main sequence. This can lead to a stable mass transfer or to a common envelope, but in any case is expected to be a short process ($\approx 10^4$ years) so only $\mathcal{O}(10)$ such systems should be present in the galaxy. During this stage, the secondary star acquires a large angular momentum from the accretion.

Once the hydrogen envelope of the donor is lost, mass transfer stops. The binary now contains a dense helium object, observable as a Wolf-Rayet star, and a massive, rapidly-spinning companion. This phase can last for $\mathcal{O}(10^5)$ years so hundreds of these systems should be present in our galaxy.

If its initial mass is larger than $\approx 8M_{\odot}$, the helium-rich object eventually undergoes an electron-capture or core-collapse supernova and forms a compact object. Depending

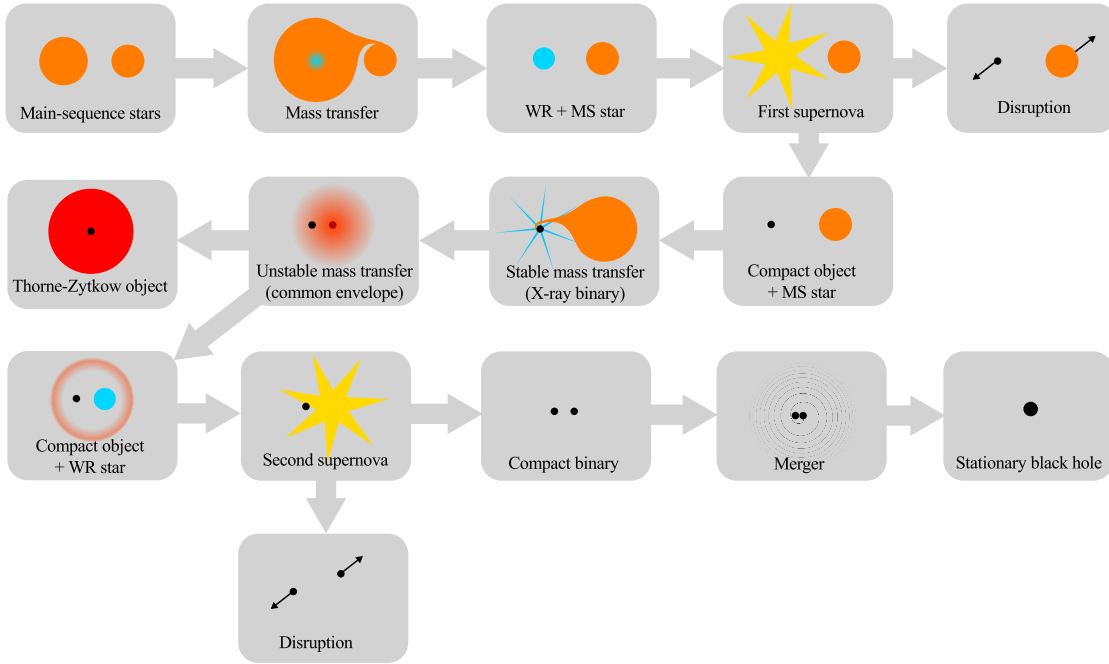


Figure 1.2: Standard evolutionary scenario for compact binaries.

on the kick associated with the supernova event, this can lead to the disruption of the binary.

Binaries surviving the kick from the first explosion contain a compact object and a rapidly-spinning, massive star, possibly in a very eccentric orbit. The non-compact star can be observed as a Be star, i.e. a hot B star with emission lines due to the presence of a ionized disk. Accretion of the Be stellar wind onto the compact object is thought to produce powerful X-ray emission and transfer a large amount of angular momentum to the compact object. This stage explains most Be/X-ray binaries observed in the galaxy. The duration of this stage is determined by the orbital parameters and it is in any case limited by the intrinsic evolution of the Be star.

The Be star eventually overflows its Roche lobe; for high mass-ratio systems, a common-envelope stage at this point is likely. This stage ends rapidly and forms either (i) a binary containing a compact object and a Wolf-Rayet star, surrounded by an expanding envelope or (ii) hypothetically, an exotic Thorne- Zytkow object. In the first case, the common envelope is expected to significantly shrink the orbit and align the spins of the objects.

Binaries containing compact objects and He stars eventually experience the second supernova explosion. The result can be two separated high-velocity compact objects or the formation of a compact binary; plenty of pulsar observations support both cases.

Compact binaries left from the previous stages evolve essentially by GW radiation and eventually coalesce, producing isolated compact objects. If one of the coalescing components is a NS, the merger can produce a short GRB. As mentioned above, pulsar

timing enabled the discovery of many compact binaries, including rapidly-inspiraling ones; however, there are no definitive observations of relativistic binaries involving BHs. Also, all observed inspiraling binaries will only merge in 10^6 years or more.

Note that the standard evolutionary scenario summarized here is only valid for close binaries and for a particular mass range. Different scenarios could also lead to the formation of compact binaries, for instance the capture of a compact object by a second one. Moreover, important parameters of the model remain poorly constrained, for instance the threshold mass for the formation of a BH from a main-sequence star, the BH birth mass and supernova kicks.

1.2.7 Coalescence rate

An important observable prediction of binary evolution models is the coalescence rate density, i.e. the number of coalescence events per unit of time and volume. This number can be computed for a particular model via population synthesis codes, which evolve large numbers of simulated binaries assuming particular distributions of initial parameters [79]. Current uncertainties in modeling common envelopes and supernova kicks imply a large uncertainty in theoretical estimates.

Existing observations of pulsars in binaries place important constraints on the galactic BNS coalescence rate. This can be estimated as $\mathcal{R} \approx N/(T_p + T_{gw})$ where N is the estimated number of binaries, T_p is the estimated characteristic age of their pulsars and T_{gw} is the time required for the systems to decay and coalesce due to GWs. The largest contribution to this estimate is given by the double pulsar J0737-3039, due to its short orbital period. Another way to estimate the BNS merger rate uses the observed formation rate of isolated galactic pulsars and the fraction of NS binaries where one of the components is a young pulsar. The two methods give consistent results of $\mathcal{O}(10^{-5})$ mergers per year in our galaxy.

As there are no direct observations of compact binaries with BHs, estimating their rate is more challenging. An upper limit is determined by the formation rate of their direct progenitors, high-mass X-ray binaries (HMXBs), which could be $\approx 10^{-3}$ to $\approx 10^{-5}$ yr^{-1} . Unfortunately, the fraction of HMXBs actually forming coalescing NSBH or BBH binaries is very uncertain. Cygnus X-1 has been used to set a lower limit on the Galactic NSBH rate of $\approx 10^{-9}$ yr^{-1} [80]. An upper limit of $\approx 6 \times 10^{-5}$ yr^{-1} on the NSBH rate can also be set by considering the abundance of heavy elements produced by r-processes when NSs are tidally disrupted [81].

Once an estimate of the galactic rate is available, this can be extrapolated to the local universe through a scale factor based on relative luminosity or star formation rate. This scale factor is currently estimated at ≈ 0.01 Mpc^{-3} . For an estimation of the rate density at cosmological scales, see [82].

Because of small statistics and selection effects, all these estimates are affected by uncertainties of at least a few factors. Better or independent measurement of the rate could therefore significantly improve our knowledge of binary evolution. In any case, current estimates consistently indicate that CBC events are extremely rare phenomena.

1.3 Studying compact binaries via gravitational radiation

1.3.1 Motivation

Due to the relative simplicity of compact objects, especially BHs, compact binaries merging on a time scale shorter than the Hubble time constitute excellent astrophysical laboratories for probing matter in extreme conditions and GR in the strong-field regime. As we have seen, electromagnetic observations of binaries already enabled important discoveries and significant understanding of compact objects.

In principle, however, the universe can also be observed through gravitational radiation. Due to the weakness of the gravitational interaction, this is far harder than using electromagnetism and in fact GWs have never been directly detected on Earth. Nevertheless, there are several reasons why detecting GWs from astrophysical sources, and in particular coalescing compact binaries, continues to be an active and important research topic in modern physics.

Several important predictions of GR have been confirmed experimentally. However, the properties of gravitational radiation still need to be tested (or at least confirmed) via a direct detection of GWs by an experiment on Earth. Producing detectable GWs in a controlled way, and thus implementing an analogy of Hertz’s experiment for gravitation, is unfortunately not feasible. GW detection attempts, therefore, must resort to the most violent astrophysical events as the signal source.

Compact binaries are naturally-occurring and efficient GW radiators; in fact, moments before coalescence, they are the loudest expected GW sources in their frequency band. In principle a “true” CBC waveform can be complicated by effects such as eccentricity, precession of the orbital plane, magnetic fields, tidal deformation or disruption; indeed, the exact solution to Einstein’s equation for the two-body problem is unknown even for two point particles. Nevertheless, compared to other GW sources like supernova explosions, CBC waveforms are simple quasi-monochromatic signals and their most important features are directly related to the physical parameters of the source. As we shall see shortly, this can be used to implement powerful data analysis techniques.

Confident direct observation of several CBC events would significantly improve our estimate of the binary coalescence rate. Moreover, even when simple pN approximations are used, CBC waveforms encode several important physical parameters of the source, notably the component masses and spins. Therefore, by observing several such signals and inferring the source parameters, the underlying population of compact binaries could be probed and our models of binary evolution could be refined [83]. Some of the open issues introduced in the previous sections—like the uncertain BHs mass distribution or the maximum mass of NSs—could be directly answered by GW detections. A simultaneous CBC and short-GRB detection would confirm the hypothesis that short GRBs are associated with BNS or NSBH mergers. In the case of BNS or NSBH, matter effects on the waveform would likely constrain the EOS of nuclear matter [84, 85]. Thanks to the cosmological redshift induced on their GW emission, coalescing binaries could in principle also enable a precise measurement of the Hubble constant [86, 87]. CBC detections may also highlight unexpected deviations from GR, if these affect the coefficients of the pN

expansion [88, 89]. Finally, many BBHs binaries may not be visible electromagnetically and their GW signal could be the only way to detect and study them.

Because GWs interact very weakly with matter, the entire universe is essentially transparent to them; the only limit to the detection and study of astrophysical sources is their distance. Once the detectability of GWs is proven, *gravitational astronomy* is thus expected to become an important complement to traditional electromagnetic observations.

Estimates of the CBC rate density can be used for predicting the rate of detections by existing and future GW detectors [90, 91]. Despite the large uncertainty, detections are expected to be relatively rare in the near future. Before gravitational astronomy can begin, therefore, extremely sensitive GW detectors must be built and the efficiency of the data analysis algorithms must be maximized; this thesis focuses on the latter problem.

1.3.2 Experimental setup

In order to describe the observable effect of gravitational radiation on matter and light we first need a suitable frame of reference representing our detector. This is usually chosen as the *Fermi normal coordinates*, where the world-line of the observer is a geodesic $x^\mu(t)$, spatial axes are defined by gyroscopes at the origin and time is defined by a clock also located at the origin. The curvature associated with GWs induces *tidal* effects on the separation between nearby particles in our detector. In fact, the separation vector δ^μ between nearby particles with four-velocity field v^μ satisfies the *geodesic deviation equation*

$$\frac{D^2 \delta^\mu}{d\tau^2} = R^\mu{}_{\nu\rho\sigma} v^\nu v^\rho \delta^\sigma \quad (1.16)$$

which, assuming the particles moves slowly and writing $R^\mu{}_{\nu\rho\sigma}$ in terms of the GW tensor $h_{\mu\nu}$, is approximately

$$\frac{d^2 \delta^\mu}{dt^2} = \frac{1}{2} \delta^\nu \frac{d^2}{dt^2} h^\mu{}_\nu. \quad (1.17)$$

The force field described by this equation can be visualized intuitively by considering a ring of particles at rest in our reference frame. Let the ring define the xy plane and consider a plane wave propagating along z , i.e. orthogonal to the ring. Since GWs are transverse, no effect is visible along the z axis. However, the wave induces an apparent distortion of the ring in the xy plane. The fractional displacement of the particles as a function of time and position in the xy plane is determined by the time-dependent polarization amplitudes $h_+(t)$, $h_\times(t)$. For this reason, the wave amplitude physically represents a dimensionless strain. A linearly-polarized wave stretches the ring into an ellipse oriented in a particular direction; the opposite polarization induces the same deformation but rotated by 45 degrees along z . Circularly-polarized radiation leads to an ellipsoidal deformation rotating in time.

Due to the weak coupling between GWs and matter and the extreme distance, the signal arriving on Earth from a CBC event is still very weak, with typical strain amplitude of order $\approx 10^{-21}$ or less. One way to achieve a sufficiently precise measurement




Figure 1.3: Simplified topology of an interferometric GW detector.

is optical interferometry with km-scale arm length [92, 93]. In fact, the GW strain can be measured by emitting two coherent and orthogonal light beams from the center of the above ring (directed along unit vectors ξ and ζ), reflecting them on two particles and having them interfere back at the origin. The differential deformation of the two optical paths produced by the wave, i.e. the strain, induces a phase shift between the two beams and thus an observable modulation of the intensity of the interference figure. Using eq. (1.17)—or an integration of the perturbed metric over null paths along the arms—the phase shift can be written as

$$\Delta\varphi \propto L(\xi^i \xi^j - \zeta^i \zeta^j) h_{ij} \quad (1.18)$$

where L is the arm length, assumed much smaller than the GW wavelength¹ [94]. This idea can be implemented with a Michelson interferometer, whose basic components are shown in figure 1.3. The two coherent beams are produced by reflecting monochromatic infrared laser light on a beam splitter. The beams are sent to two mirrors, reflected and recombined at the beam splitter. The phase shift $\Delta\varphi$ modulates the intensity of the light at the output port, which is read off by a photodiode. The photodiode signal is conditioned, digitized and stored for later analysis together with an accurate GPS time stamp.

Although this setup is in principle simple, when implemented on Earth it requires a few tricks to maximize its sensitivity—such as the use of mode cleaners, power recycling and Fabry-Perot cavities in the arms—which we will not describe in details here. An important point however is that ground-based interferometers are large and extremely complex instruments, very sensitive to a variety of complicated environmental and instrumental disturbances, such as local seismic motion, dust and gas in the optical paths,

¹For signals from stellar-mass compact binaries this is a good approximation.

electromagnetic interference, laser instability and even the modulation of the local gravitational potential by nearby moving masses. One of the major challenges in operating such instruments is therefore protecting the optical paths from these disturbances as much as possible. In particular, the mirrors and beam splitter are mechanically isolated from the ground and the arms are contained in vacuum chambers.

Nevertheless, fundamental noise sources always ultimately limit the sensitivity of the instrument. Ground-based interferometers are typically limited by seismic noise at frequencies below $\mathcal{O}(10)$ Hz. On the other hand, above $\mathcal{O}(10^3)$ Hz the GW wavelength becomes comparable to or shorter than the effective length of the interferometer arms and the response of the instrument drops, such that the limiting noise is due to the discreteness of the photons in the light beams (*shot noise*). Thus, ground-based interferometers are typically sensitive to GWs at frequencies between ≈ 10 and ≈ 1000 Hz. This includes the last minutes of BNS inspiral signals and the last seconds or tens of seconds of BBH signals with total mass up to $\approx 100M_\odot$. An important part of GW research focuses on improving technology in order to reduce the limits as much as possible.

Ground-based laser interferometers with km-long arms are the main technology currently used to search for stellar-mass CBCs (among other GW sources). Three interferometers are currently being upgraded from their first-generation design to an “advanced”, more sensitive configuration: the two American LIGO interferometers at Hanford and Livingston [95] and the Italian Virgo [96]. KAGRA is a brand new interferometer undergoing construction in Japan [97] and one more LIGO site is currently planned in India [98]. The only active instrument at the moment is the smaller and less sensitive GEO 600 interferometer in Germany [99], where some of the technology going into the larger instruments has been developed.

Although this thesis focuses on data from ground-based interferometers, it is worth noting that GWs can be sought in other ways too. A very promising technique is laser interferometry in space with $\mathcal{O}(10^6)$ m arm length, which led to the design of the eLISA mission [93, 100]. If successful, this method would enable probing a complementary frequency band where other kinds of astrophysical and cosmological GW sources could be detected, as well as some of the observed NS and WD binaries. Another promising technique is *pulsar timing arrays*, which uses radio telescopes to precisely monitor the arrival-time of pulses from several stable radio pulsars, thus implementing interferometry with astrophysical arm lengths [101]. Historically, the first attempt at detecting GWs used resonant-bar detectors [102]. Although cheap and simple, however, this technology is no longer considered sensitive enough to yield detections.

1.3.3 Data model

The data produced by an interferometric GW detector and subject to analysis consist of a time series $s(t)$ calibrated in units of dimensionless strain. The data receive contributions from the actual GW strain $h(t)$ we want to measure, as well as the noise of the detector $n(t)$, and can thus be modeled as $s(t) = n(t) + h(t)$. The weakness of GW signals implies that $n(t)$ typically dominates $s(t)$.




Figure 1.4: Noise spectral densities associated with past and future configurations of the LIGO interferometers. Data from [105–107].

$n(t)$ is usually described as a Gaussian stochastic process associated with a single-sided power spectral density (PSD) $S(f)$, defined by

$$\langle \tilde{n}(f)\tilde{n}^*(f') \rangle = \frac{1}{2}S(f)\delta(f - f') \quad (1.19)$$

where $\langle \cdot \rangle$ represents the expectation value over the ensemble of noise realizations. Because of the variability of the environment around the detector, this Gaussian component can only be considered stationary for a short time scale, and $S(f)$ typically slowly fluctuates over time by factors of a few. In practice it is therefore continuously estimated from the data using a modified Welch method [103, 104]. Figure 1.4 illustrates the typical PSD achieved by the LIGO detectors during the last scientific run [105], the target PSD of advanced LIGO in its final configuration (the so-called “zero-detuning, high-power” configuration) [106] and an intermediate PSD expected in the first run of advanced LIGO [107]. Chapters 2 and 3 shall focus on the latter curve.

The instrumental and environmental disturbances which affect interferometers also contribute to $n(t)$, typically via long-lasting, quasi-monochromatic signals (*lines*) and short transients lasting up to several seconds (*glitches*). Epochs of particularly unreliable or doubtful data are tagged with *category vetoes*, for instance based on observed correlation of the disturbances with auxiliary monitoring channels; epochs of a particular category can then be optionally excluded from the analysis [108–110]. Unfortunately, this does not guarantee completely clean data and many disturbances remain which are

not as easily modeled as the underlying quasi-stationary Gaussian component, representing a major challenge to the analysis of the data. We shall explore this problem for CBC searches in chapter 5.

$h(t)$ is the full astrophysical GW signal; in this thesis however we shall focus on CBC signals only. As the past sections outlined, the full CBC parameter space (spanned by the parameter vector $\boldsymbol{\theta}$) is highly-dimensional and poorly restricted by existing observations. The parameters are in principle:

- Coalescence time t_c and phase ϕ_c ;
- Component masses m_1, m_2 ;
- Dimensionless spin vectors at a reference time $\boldsymbol{\chi}_1, \boldsymbol{\chi}_2$;
- Distance r ;
- Sky position (right ascension α and declination δ , or unit vector $\hat{\boldsymbol{r}}$);
- Orientation of the total angular momentum (inclination ι and orientation on the plane of the sky ψ);
- Magnitude and orientation of the eccentricity vector.
- In the case of NS binaries, any parameter of the EOS model and possibly of magnetic field models.

Nevertheless, some of these quantities are usually neglected because they are likely small (e.g. eccentricity) or have a weak effect on the signal (e.g. magnetic fields). Others are partially or completely degenerate. The most important parameters to search over are the amplitude of the signal, the coalescence time and phase and the masses. Chapters 2 and 3 will show cases where this is no longer true and spins also need to be included. Those parameters which only affect the overall amplitude, phase and time shift of the signal are called *extrinsic*; those affecting the shape of the waveform, typically the masses and spins, are called *intrinsic*².

By combining eq. (1.6) and (1.18), one finds that the GW produces an observable strain at the detector given by a linear combination of the polarization amplitudes, i.e.

$$h(t; \boldsymbol{\theta}) = F_+(\alpha, \delta, \psi; t) h_+(t; \boldsymbol{\theta}) + F_\times(\alpha, \delta, \psi; t) h_\times(t; \boldsymbol{\theta}). \quad (1.20)$$

The *beam pattern* coefficients $F_{+,\times}(\cdot)$ are determined by the directional sensitivity of the interferometer, which is corotating with the Earth and thus time-dependent; in most cases however CBC signals spend only up to ≈ 10 minutes within the sensitive band of ground-based interferometers, such that the source as seen from Earth can be considered

²This distinction depends on the approximations made by the particular waveform model. For instance, sky location and orientation are extrinsic for restricted non-precessing pN waveforms, but they do affect the shape of the waveform when precession and higher-order modes are included.

fixed in the sky and $F_{+,\times}(\cdot)$ are almost time-independent. Assuming the restricted, non-precessing pN approximation, the signal at the detector can thus be modeled simply as [104]

$$h(t; \boldsymbol{\theta}) = A(t; \boldsymbol{\theta}) \cos \Phi(t; \boldsymbol{\theta}). \quad (1.21)$$

$A(t; \boldsymbol{\theta})$ is the slowly-increasing Newtonian amplitude given by

$$A(t; \boldsymbol{\theta}) := - \left(\frac{GM}{c^2 r_{\text{eff}}} \right) \left(\frac{5G}{c^3} \frac{\mathcal{M}}{t_c - t} \right)^{1/4} \quad (1.22)$$

with the *effective distance* defined as

$$r_{\text{eff}} := r \left[F_+^2(\alpha, \delta, \psi) \left(\frac{1 + \cos^2 \iota}{2} \right)^2 + F_\times^2(\alpha, \delta, \psi) \cos^2 \iota \right]^{-1/2}. \quad (1.23)$$

The effective distance combines the physical distance, orientation and sky location into a single degenerate parameter. It is in practice always larger than the physical distance, equality being achieved only when the source is face-on (or face-off) and directly overhead (or below) the detector. $\Phi(t; \boldsymbol{\theta})$ is the pN-corrected GW phase. As we shall see shortly, it is useful to express this signal in the frequency domain; for a restricted pN signal this can be done analytically with the TaylorF2 approximant

$$\tilde{h}(f; \boldsymbol{\theta}) \propto r_{\text{eff}}^{-1} \mathcal{M}^{5/6} f^{-7/6} e^{-i\Psi(f; \boldsymbol{\theta})}. \quad (1.24)$$

For an expression of $\Psi(f; \boldsymbol{\theta})$ at 3.5 pN order (2.5 in spins) see for instance [111].

1.3.4 Matched filtering and template banks

The problem of searching for CBC signals in the data can be cast as an hypothesis test, where the null hypothesis “ $h(t) = 0$ ” (absence of a signal) is compared to the signal hypothesis “ $h(t)$ contains an inspiral signal described by a parameter vector $\boldsymbol{\theta}$ ”. This comparison is formalized by computing the maximized likelihood ratio

$$\Lambda := \max_{\boldsymbol{\theta} \in \mathcal{S}} \Lambda(\boldsymbol{\theta}) = \max_{\boldsymbol{\theta} \in \mathcal{S}} \frac{P(s|h(\boldsymbol{\theta}))}{P(s|h=0)}. \quad (1.25)$$

Here $P(s|h=0)$ is the probability of observing the data segment in the absence of a signal, $P(s|h(\boldsymbol{\theta}))$ is the probability of the data when a CBC signal with parameters $\boldsymbol{\theta}$ is present and \mathcal{S} is the extent of CBC parameter space we are searching over. The assumption of stationary Gaussian noise allows one to write explicitly the probabilities associated with the two hypotheses and computing the (log) likelihood maximized over the amplitude of the signal,

$$\rho(\boldsymbol{\theta}) := \log \Lambda(\boldsymbol{\theta}) = \frac{4}{\sigma} \Re \left[\int_0^\infty \frac{\tilde{s}(f) \tilde{h}^*(f; \boldsymbol{\theta})}{S(f)} df \right]. \quad (1.26)$$

$\tilde{h}(f; \boldsymbol{\theta})$ is called a *template* waveform, $\rho(\boldsymbol{\theta})$ is the *signal-to-noise ratio* (SNR) associated with that template and

$$\sigma := \sqrt{4 \int_0^\infty \frac{|\tilde{h}(f; \boldsymbol{\theta})|^2}{S(f)} df} \quad (1.27)$$

is a normalization constant for the template quantifying how much it responds to the noise. The SNR represents the cross-correlation between the whitened data and the whitened, normalized template. The integrals can also be interpreted as the application of a filter to the data, where the filter is optimal (*matched*) in the sense that it maximizes $\rho(\boldsymbol{\theta})$ for a CBC signal described by exactly the same parameters $\boldsymbol{\theta}$. In practice the lower limit of the integrals is set to $f_L \gtrsim 10$ Hz due to the relatively poor low-frequency sensitivity and reliability of ground-based interferometers and the upper limit f_H is always at most the Nyquist frequency of the sampled data. In the case of pN templates, an even smaller f_H is chosen in order to “manually” truncate the pN inspiral before the merger. In past searches this was the frequency of the innermost stable circular orbit (ISCO) for a Schwarzschild BH of mass $M := m_1 + m_2$, which is $f_{\text{ISCO}} := c^3(6\sqrt{6}\pi GM)^{-1}$.

In order to evaluate Λ , we need to compute $\rho(\boldsymbol{\theta})$ over the parameter space \mathcal{S} . Maximization over the amplitude parameters is already achieved by using the matched filter. In addition, templates which differ only by a $\pi/2$ shift in coalescence phase are almost orthogonal, so the maximization over coalescence phase can also be done analytically by summing the squared SNRs of two templates with a $\pi/2$ phase difference. Furthermore, because a time translation of the template simply adds a term linear in frequency to the phase of $\tilde{h}(f; \boldsymbol{\lambda})$, computing the SNR over coalescence time can be done efficiently (albeit not analytically) by writing the integral in eq. (1.26) as an inverse Fourier transform and then using the fast Fourier transform (FFT) algorithm to compute it numerically. The SNR maximized over amplitude and coalescence phase is then

$$\rho(t_c, \boldsymbol{\lambda}) := \frac{|z|}{\sigma}, \quad z := 4 \int_0^\infty \frac{\tilde{s}(f)\tilde{h}^*(f; \boldsymbol{\lambda})}{S(f)} e^{-2\pi i f t_c} df \quad (1.28)$$

where $\boldsymbol{\lambda}$ is now the vector restricted to the intrinsic parameters (i.e. excluding the amplitude parameters and phase and time of coalescence) and the template now implicitly uses a coalescence time of 0. The complex matched-filter output z contains the SNR in its magnitude and the coalescence phase in its argument. The probability density for $\rho(t_c, \boldsymbol{\lambda})^2$ under the null hypothesis is a χ^2 distribution with 2 degrees of freedom (DOF), one for each choice of coalescence phase; in the presence of a signal, the distribution becomes non-central. An analytical or efficient maximization procedure over the remaining parameters $\boldsymbol{\lambda}$ has never been found, so the usual approach is to numerically evaluate $\rho(t_c, \boldsymbol{\lambda})$ over a discrete and finite set of points $\{\boldsymbol{\lambda}_i\}$ covering the parameter space \mathcal{S} , which is called a *template bank*. The algorithm invented for this calculation is called FindChirp [104]. In past searches, the computational cost of calculating $\rho(t_c, \boldsymbol{\lambda})$ for each template and data segment has been typically dominated by the FFTs.

When discussing template banks in the rest of this chapter and throughout this thesis, we shall encounter a few important quantities which we define here. The matched-filter

integral (eq. (1.26)) defines a noise-weighted inner product between two signals a and b , usually called *overlap*,

$$(a, b) := 4 \Re \left[\int_0^\infty \frac{\tilde{a}(f)\tilde{b}^*(f)}{S(f)} df \right] \quad (1.29)$$

and consequently the noise-weighted norm $\|a\| := \sqrt{(a, a)}$. With this definition, we can rewrite eq. (1.28) as

$$\rho(t_c, \boldsymbol{\lambda}) = \frac{\sqrt{(s, h_0(t_c, \boldsymbol{\lambda}))^2 + (s, h_{\pi/2}(t_c, \boldsymbol{\lambda}))^2}}{\|h_0(\boldsymbol{\lambda})\|}, \quad (1.30)$$

where h_0 and $h_{\pi/2}$ are the same template with a $\pi/2$ difference in coalescence phase; the SNR is thus the magnitude of the projection of the data on the normalized template. The *optimal SNR* of a CBC signal a is $\rho_{\text{opt}} := \|a\|$ and represents the expectation value of the SNR for a template exactly identical to the signal. Due to a variety of effects (discreteness of the template bank, approximations in the waveform model, calibration uncertainty etc) the optimal SNR is an idealization and the SNR observed in practice is always smaller. In order to characterize this loss, we introduce the *match* between two CBC waveforms a and b as the overlap between their normalized versions, maximized over their relative time and phase of coalescence:

$$0 \leq m(a, b) := \max_{t_c, \phi_c} \frac{(a, b)}{\|a\| \|b\|} \leq 1. \quad (1.31)$$

The match quantifies the similarity between a signal and a template (or between two templates) irrespective of their relative amplitudes, phases and time translations. As such, it measures the fraction of optimal SNR that can be recovered by an imperfect template. Importantly, it also defines a distance between points $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ in the intrinsic CBC parameter space

$$d(\boldsymbol{\lambda}, \boldsymbol{\mu}) := \sqrt{1 - m(h(\boldsymbol{\lambda}), h(\boldsymbol{\mu}))}. \quad (1.32)$$

For well-behaved parameter spaces (the intrinsic CBC space is such a case) the match falls off quadratically for nearby waveforms. This enables the introduction of a metric tensor g_{ij} on the parameter space, such that the match can be approximated as the quadratic form [112]

$$m(h(\boldsymbol{\lambda}), h(\boldsymbol{\lambda} + \boldsymbol{\delta})) \approx 1 - g_{ij}(\boldsymbol{\lambda}) \delta^i \delta^j \quad (1.33)$$

with the metric given by

$$g_{ij}(\boldsymbol{\lambda}) := -\frac{1}{2} \left. \frac{\partial^2 [m(h(\boldsymbol{\lambda}), h(\boldsymbol{\mu}))]}{\partial \mu_i \partial \mu_j} \right|_{\boldsymbol{\mu}=\boldsymbol{\lambda}}. \quad (1.34)$$

The metric depends on the derivatives of the template waveform with respect to its parameters and the noise PSD and provides an important description of the space in a condensed way; it has been computed analytically for various pN approximations, see e.g. [111, 113]. Given a CBC signal a and a template bank $\{\boldsymbol{\lambda}_i\}$, the maximum match

that can be obtained between a and any waveform $h(\boldsymbol{\lambda}_i)$ in the bank is called the *fitting factor* between the signal and the bank,

$$\varphi := \max_i m(a, h(\boldsymbol{\lambda}_i)) \quad (1.35)$$

The fitting factor represents the fraction of optimal SNR that the bank as a whole is able to observe and thus is a useful metric for comparing different template banks against one or more signals [114]. It will be used extensively in the following chapters.

A template bank is said to be *effectual* if its fitting factor with any signal in the search space of interest is equal to (or larger than) a given minimum value, typically 0.97. The criterion for constructing a template bank can thus be formalized as choosing the smallest possible effectual bank. This problem has been investigated extensively for parameter spaces of different dimensionality and various solutions have been proposed [112, 115–119]. This thesis considers two of them. *Stochastic* template placement works by randomly choosing a new template and adding it to the bank only if its fitting factor with the already-chosen templates is smaller than the requirement. This procedure has the advantage of simplicity and can be used to construct banks of arbitrarily complicated signals as long as their match can be computed reasonably efficiently, for instance using the metric approximation. The drawback is that it generates a sub-optimal bank which does not generally achieve the condition of minimum number of templates. *Geometric* template placement, on the other hand, identifies a coordinate system (typically in a higher-dimensional space) where the distance between points is coordinate-independent (i.e. the metric is explicitly flat) and places templates on a regular grid in those coordinates. This typically achieves a smaller number of templates, but it requires particular conditions on the parameter space. The number of templates required to cover a region \mathcal{S} of parameter space can be estimated via the metric as

$$N \approx \frac{1}{\Delta V} \int_{\mathcal{S}} \sqrt{\det [g_{ij}(\boldsymbol{\lambda})]} d^n \boldsymbol{\lambda} \quad (1.36)$$

where n is the dimensionality of the parameter space and ΔV is the proper (metric) volume covered by a single template. ΔV depends on the required minimum match and on how templates are distributed over \mathcal{S} [112, 116].

1.3.5 Signal-based vetoes

Eq. (1.26) and the resulting formalism stem from the assumption that the noise is Gaussian. We have seen however that the data contain not only Gaussian noise and (possibly) GWs, but frequently also glitches of environmental or instrumental origin, which are completely unrelated to astrophysics. Unfortunately, even if these signals are qualitatively different from CBC waveforms, their loudness can have a large impact on the SNR—as shall be investigated in detail in chapter 5—which would dramatically reduce the search sensitivity. This has been evident since the very first search and required the ad-hoc introduction of *signal-based vetoes*. These are functions of the data devised to

augment the SNR with an actual check of the consistency of the data with the template model, i.e. they represent goodness-of-fit tests.

A widely used veto in past searches (and in this thesis) is the reduced χ^2 veto [120]. The idea behind it is to split the frequency domain into p disjoint bands (indexed by $\ell = 1 \dots p$) and calculate the partial complex matched-filter output z_ℓ in each band. The statistic is then defined as

$$\chi_r^2 := \frac{p}{2p-2} \sum_{\ell=1}^p \left| z_\ell - \frac{z}{p} \right|^2. \quad (1.37)$$

Under *both* the noise and signal hypotheses, this quantity is described by a (normalized) χ^2 distribution, in particular its expectation value is 1. If however the data contain a disturbance, or a CBC signal with a large mismatch with the template, then the distribution becomes non-central. Thus χ_r^2 can be used to detect an inconsistency with both the noise and signal models. Intuitively, the meaning of eq. (1.37) is to check that the distribution of signal power in the time-frequency plane exactly matches that of the template; equivalently, it verifies that the partial SNRs all peak consistently at the same time as the actual SNR.

The χ^2 statistic has been used to augment the standard likelihood in an ad-hoc way by defining the so-called *reweighted SNR* [121]

$$\hat{\rho} := \begin{cases} \rho \left[\frac{1}{2} \left(1 + (\chi_r^2)^3 \right) \right]^{-1/6} & \text{if } \chi_r^2 > 1 \\ \rho & \text{otherwise} \end{cases} \quad (1.38)$$

This has proved to be a very effective method of protecting the search sensitivity from glitches [122] and we shall use it in chapters 2 and 3 as well. Note that other signal-based vetoes similar to the χ^2 test can be defined and used to weight the SNR as in eq. (1.38), but they are not used in this thesis [123].

Signal-based vetoes will generally respond to the unavoidable differences between true CBC signals and template waveforms. The non-centrality parameter of the χ^2 veto, for instance, depends on the match between the template and the signal [104]. If the models used to construct template waveforms are poor, such that large mismatches are present between true signals and the templates, signal-based vetoes could actually reduce the detectability of signals, i.e. they may be *unsafe*. It is therefore important to establish the safety of signal-based vetoes against uncertainties in CBC waveform modeling, which usually means allowing for a sufficiently large fitting factor between the template bank and the target signal population. We shall see cases where the χ^2 veto is not safe in chapters 2 and 3.

1.3.6 Multiple detectors: coincidence

Computing eq. (1.28) produces a discrete time series of SNRs for each template. Samples larger than a predetermined threshold $\bar{\rho}$ are recorded as *candidate events* or *triggers*. Because $\rho(t_c, \boldsymbol{\lambda})$ is a smooth function of t_c , several adjacent samples may produce correlated

triggers, especially if a loud signal or transient is present in the data. This correlation is reduced by only retaining the maximum trigger in a time window longer than the fluctuation time of $\rho(t_c, \boldsymbol{\lambda})$ [104]. After matched filtering and maximizing over t_c , the initial strain data are therefore reduced to a list of potentially interesting single-detector triggers, each tagged with the parameters of the template, estimates of coalescence time and phase and effective distance as well as the SNR and χ_r^2 .

Despite the discriminatory power of signal-based vetoes, an important part of CBC searches is ensuring that an event is observed in two or more independent instruments, possibly widely separated (*coincidence*). This dramatically reduces the false-alarm background and in particular its contribution from noise artifacts such as glitches. Coincidence is implemented by comparing candidate events across instruments and flagging events whose parameters (in particular coalescence time and masses) are consistent. The time-of-flight between detectors for a GW propagating at speed c must also be taken into account; for the Hanford and Livingston LIGO interferometers this is up to ≈ 10 ms. A widely-used coincidence algorithm, called *ethinca*, uses the metric of the parameter space to associate error ellipsoids to each event; coincidence is then achieved when the ellipsoids of two or more events overlap [124]. A simpler alternative—being considered for future searches—is called *exact match* and consists of using identical template banks for every detector, coincidence being defined by simply looking for events from identical templates in the detectors and with a compatible coalescence time. This has the advantage of being straightforwardly applicable to any parameter space, regardless of dimensionality and availability of a metric tensor. We shall use the latter method in this thesis. Coincident triggers are ranked by a quantity that somehow combines the single-detector likelihoods, typically the *combined reweighted SNR* defined as

$$\hat{\rho}_c := \sqrt{\sum_{d=1}^D \hat{\rho}_d^2} \quad (1.39)$$

where $d = 1 \dots D$ labels each detector [121].

Coincidence is simple and computationally cheap, but it has a number of drawbacks. In a network of detectors, triggers produced by CBC signals in different detectors exhibit a well-determined correlation between their SNRs, coalescence phases and coalescence times, while noise triggers do not necessarily obey the same constraint. With coincidence, however, single-detector SNRs are maximized independently over amplitude and coalescence phase at each detector and the difference in coalescence times is then simply accounted for by using a suitable coincidence window. For this reason, coincidence is called an *incoherent* combination of the single-detector statistics. The price of simplicity is therefore discarding information which could potentially improve the discrimination between noise and astrophysical signals. Another drawback of coincidence is that detection involves D independent thresholds and requires dealing with combinations of two, three or more detectors and it is not obvious how to optimally tune these parameters. Finally, as D increases, more DOF contribute to $\hat{\rho}_c$ and thus its probability density under the null hypothesis becomes wider, increasing the false-alarm probability for a given detection threshold.

1.3.7 Multiple detectors: coherent combination

When the gravitational radiation is modeled as a plane wave traveling at the speed of light, a precise relationship exists between the amplitude, arrival time and phase of signals observed by different detectors. This relationship is established by the relative positions and orientations of the detectors and contains information about the wave vector, i.e. the sky position of the source. The same approach that leads one to the single-detector SNR statistic can be used to construct a multi-detector statistic that effectively includes the plane-wave model into the matched filter [125–128]. This is called *coherent SNR* as opposed to the *incoherent* multi-detector SNR defined by eq. (1.39). Using a singular value decomposition, the coherent SNR can also be derived by considering the network as two independent detectors associated with the two polarizations of the GW signal [129].

The definition of noise-weighted inner product (eq. (1.29)) can be generalized to a network of detectors as

$$(a, b)_d := 4 \Re \left[\int_0^\infty \frac{\tilde{a}(f)\tilde{b}^*(f)}{S_d(f)} df \right] \quad (1.40)$$

where S_d is the noise PSD for detector d . The coherent SNR is then defined as

$$\rho_{\text{coh}}^2 := x_i B^{ij} x_j \quad (1.41)$$

with the following auxiliary definitions:

$$x_1 := \sum_{d=1}^D F_+^d D_d[(s_d, h_0)_d] \quad (1.42)$$

$$x_2 := \sum_{d=1}^D F_\times^d D_d[(s_d, h_0)_d] \quad (1.43)$$

$$x_3 := \sum_{d=1}^D F_+^d D_d[(s_d, h_{\pi/2})_d] \quad (1.44)$$

$$x_4 := \sum_{d=1}^D F_\times^d D_d[(s_d, h_{\pi/2})_d] \quad (1.45)$$

$$B^{ij} := \begin{pmatrix} \mathcal{B}_{++} & \mathcal{B}_{+\times} & 0 & 0 \\ \mathcal{B}_{+\times} & \mathcal{B}_{\times\times} & 0 & 0 \\ 0 & 0 & \mathcal{B}_{++} & \mathcal{B}_{+\times} \\ 0 & 0 & \mathcal{B}_{+\times} & \mathcal{B}_{\times\times} \end{pmatrix}^{-1} \quad (1.46)$$

$$\mathcal{B}_{ij} := \sum_{d=1}^D F_i^d F_j^d (h_0, h_0)_d. \quad (1.47)$$

In the above definitions, s_d is the data segment from detector d , h_0 and $h_{\pi/2}$ are the quadratures of the inspiral template, F_i^d are the beam pattern functions for detector

d and $D_d[\cdot]$ is the operator that shifts the coalescence time of a signal by an amount equal to the light-travel time between detector d and the geocenter, i.e. $\Delta t_d := \mathbf{r}_d \cdot \hat{\mathbf{r}}/c$ where, in Earth-centered coordinates, \mathbf{r}_d is the position of the detector and $\hat{\mathbf{r}}$ is the unit vector to the source. The definition of B^{ij} given here makes the approximation that the “0” and “ $\pi/2$ ” phases of the CBC signal are orthogonal, which is satisfied by signals completing many cycles within the sensitive bands of the detectors.

The coherent approach also enables the definition of additional signal-based vetoes, such as the *null SNR*, which check the consistency of signals with the plane-wave model. These, however, are not used in this thesis and their details can be found in the references above.

Note that all the above quantities, including ρ_{coh} and the signal-based vetoes, depend on the sky location of the source. If this is not known, ρ_{coh} must be computed over a template bank of sky locations in addition to the bank of intrinsic parameters. In comparison, the coincident method maximizes over the sky location in a computationally-trivial way. Chapter 4 shall perform a cost-benefit comparison of the two methods, taking into account the different sensitivities and computational costs.

1.3.8 Statistical significance of candidates

The final step in processing the candidate events is establishing their statistical significance, i.e. the probability that each event is generated under the assumption that the data contain *no* CBC signals. If we had a perfect model of our false-alarm background, this could simply be computed from the parameters of each trigger. The instability of the behavior of the detectors and the difficulty in modeling it accurately, however, require that this is somehow measured from the data.

The approach used in past searches is based on *time slides*, i.e. on shifting the data of one detector relative to the others multiple times and recalculating the list of coincident events [121]. This effectively removes GWs from the set of candidate events and enables sampling the false-alarm background many times; the variance in the estimation is nevertheless ultimately limited by the amount of data actually available [130]. Once the background sample is obtained, the significance of each “zero-lag” candidate—or equivalently the false-alarm rate (FAR) associated with it—can be estimated as the rank of that candidate with respect to the background triggers. A detection is then claimed if a given significance is achieved by some events. The method of time slides shall be used in chapters 2 and 3.

Although a detection has never been claimed, an important task in the case of a significant signal is estimating the parameters of the source. This inference problem is usually formalized in the Bayesian statistical framework and implemented, for instance, in the LALInference software [131]. Follow-up and parameter estimation of detected signals are outside the scope of this thesis.

1.3.9 Results of past searches

Data from first-generation interferometers have been used extensively to search for CBC signals. No plausible events were ever detected and the result was setting tighter and tighter upper limits on the CBC rate—typically as a function of the mass parameters—as the detector sensitivity improved.

[132] presents a search for BNS systems on data from the first LIGO scientific run. A similar search was also done using data from LIGO and TAMA300 [133]. One more BNS search was performed using LIGO’s second scientific run [134]. LIGO’s second run data were also used to search for coalescing primordial BBHs [135]. Data from the third LIGO run were used to search for precessing systems by introducing a “detection template family” designed specifically to capture the modulation of the signal [136]. [137, 138] searched for CBCs with total mass between $2M_{\odot}$ and $35M_{\odot}$ using data from LIGO’s fifth science run. In [139], BBH coalescences were searched in LIGO’s fifth science run using IMR templates for the first time, producing a rate upper limit of $2 \text{ Mpc}^{-3} \text{ yr}^{-1}$. [140] was the first search combining LIGO and Virgo data to search for systems with total mass between $2M_{\odot}$ and $35M_{\odot}$. In [122], total masses between $2M_{\odot}$ and $25M_{\odot}$ were searched in LIGO’s sixth science run and Virgo’s science runs 2 and 3; the resulting upper limits for BNS, NSBH and BBH are 1.3×10^{-4} , 3.1×10^{-5} and $6.4 \times 10^{-6} \text{ Mpc}^{-3} \text{ yr}^{-1}$. [141] presents a search for high-mass BBH coalescence in the latest LIGO-Virgo data, reporting an upper limit of $3.3 \times 10^{-7} \text{ Mpc}^{-3} \text{ yr}^{-1}$. All these searches used coincidence for combining multi-detector data. The upper limits on the CBC rate produced by these searches have been compared to predictions and found to be up to a factor of only a few away from optimistic predictions [83]. CBC events were also searched in association with GRBs detected electromagnetically, using the coherent SNR instead [142].

As visible in figure 1.4, advanced ground-based interferometers are expected to improve their strain sensitivity by roughly one order of magnitude, which translates to a potential factor of 10^3 increase in detection rate. In order to realize this potential, the data need to be analyzed with methods as close to optimal as possible. In the next chapters we shall investigate potential causes of sub-optimality in future CBC searches, propose ways to address them and estimate the resulting improvement.

Chapter 2

Search pipeline for NSBH coalescence including aligned-spin effects

2.1 Overview

We have seen in the introduction that most previous CBC searches have been carried out with the assumption that component spins are negligible. We have also seen that, according to recent research, BHs may have spin magnitudes close to the Kerr limit. Currently available pN waveform models and templates do have the possibility of including spin effects, at least to some pN order, without having to resort to phenomenological or ad-hoc waveforms. This raises the following questions: (i) does including spins in the parameter space of future searches actually increase their sensitivity? (ii) is the additional computational cost sustainable?

Spin effects have been considered in several previous studies. It was initially found that non-spinning 1.5 pN or 2 pN templates were effectual at detecting systems with spins aligned with the orbital angular momentum (i.e. non-precessing) but could not recover precessing signals [114, 143]. One of the previous CBC searches investigated spin effects on upper limits, although it did not use templates with spin [141]. One search used a “detection template family” designed to capture the modulation induced by strongly-precessing systems [136]. However, these templates included waveforms outside of the manifold of physical signals. Given that a signal consistency check such as the χ^2 veto was not available for those waveforms, the resulting false-alarm background was quite large. Another study attempted to capture the effect of precession, but again without signal-based vetoes [144]. The additional parameter freedom turned out to be extremely sensitive to noise artifacts, reducing the overall sensitivity instead of increasing it. No search attempted to capture the effect of spins aligned with the orbital angular momentum. More recently, the effect of spin on search sensitivity has been studied again via fitting-factor calculations, assuming stationary Gaussian noise, for BNS [111], NSBH [145] and generic low-mass binaries [146]. The latter studies demonstrated the

importance of including spin effects in a search. Nevertheless, a full search pipeline including pN spin effects, coincidence between detectors and robust signal-based vetoes like the χ^2 test has never been implemented and shown to be more sensitive than a non-spinning search.

In this chapter we address this problem for the particular case of NSBH binaries. As discussed in chapter 1, these systems are an important potential end product of some of the observed X-ray binaries, a candidate explanation for short GRBs and for the synthesis of a large fraction of heavy elements via r-processes. Tidal effects on the NS induced by the BH are also potentially useful for studying dense nuclear matter. NSBH mergers are expected to yield less GW detections than BNS or BBH, due to their relatively low estimated rate density [90, 91]. Because of the possibility of almost-extremal BH spins, we must thus ensure that future searches are able to detect these relatively rare signals without introducing unnecessary selection effects. Focusing first on the simpler case of non-precessing spins, we present here the first implementation and test of a prototype search for NSBH coalescences which includes the pN effect of the spin components parallel to the orbital angular momentum. We use this implementation to compare a “traditional” template bank which ignores spin effects altogether to a bank which includes aligned-spin effects only. We show that including aligned spins increases significantly the computational cost of the search. However, thanks to modern computing technologies and a new implementation of the matched-filtering code, the extra cost is sustainable. We also show that the expansion of the parameter space does increase the false-alarm rate of the pipeline, but (i) the χ^2 signal-based veto is able to mitigate the effect of instrumental artifacts even when spin is included and (ii) the additional sensitivity to strongly-spinning systems is able to offset the increased false-alarm background, to an extent which depends on the distribution of BH spins in the universe. In the worst case of non-spinning BHs, the reduction of sensitivity due to the larger background is minimal. Therefore, the aim of this chapter is to demonstrate that aligned-spin effects can and should be included in future NSBH searches.

The research presented in this chapter has been published in [147]. This chapter goes into a deeper level of detail than the published article, in particular providing greater detail on the use of online databases, on the false-alarm background of the pipeline and on the pipeline’s ability to recover simulated signals. The project is the result of a collaboration with several people. Most of the matched-filter engine and the fitting-factor calculation program has been implemented by Alex Nitz and Josh Willis. The code for constructing the template banks has been written by Ian Harry and already used in previous publications. Some of the fitting-factor calculations have been performed by Alex Nielsen. The MongoDB and Cassandra databases have been installed and maintained by Oliver Bock, who was also an important help in using them efficiently. Carsten Aulbert and Henning Fehrmann maintained the Atlas cluster, in particular the graphics processing units (GPUs) on which the matched-filtering jobs were executed. Karsten Wiesner implemented some of the early attempts at GPU acceleration of the matched-filter program. Duncan Brown, Tom Dent, Badri Krishnan, Andrew Lundgren, Alex Nielsen and Frank Ohme gave important advice and comments on the project

and the paper. The author (i) executed some of the fitting-factor calculations, (ii) contributed to developing and testing the matched-filter engine, (iii) implemented the infrastructure of the pipeline, the coincidence and clustering algorithms, the code for recovering the simulated signals and for estimating the sensitivity of the pipeline, (iv) set up the PostgreSQL database, (v) ran the search pipeline and followed up the most interesting background events, (vi) wrote a significant part of the paper and (vii) handled the publication process and communication with the anonymous referee.

The chapter is organized as follows. Section 2.2 describes the sources we are interested in and the CBC waveform model we use to simulate the signals. Section 2.3 describes the simulated data we use for testing the search pipeline. Section 2.4 presents a new software package for CBC searches which we contributed to developing and which our search is based on. In section 2.5 we describe a “traditional” template bank neglecting spin and one including aligned-spin effects and we compare their ability to recover the target sources. Section 2.6 describes our implementation of the full pipeline and the choice of its parameters. Section 2.7 investigates the features of the false-alarm background produced by the pipeline. In section 2.8 we compare the ability of the two banks to recover signals in the simulated data. Section 2.9 shows the sensitivity of the pipeline with the two banks and is the main result of this chapter. Finally, section 2.10 summarizes the conclusion.

2.2 Simulation of the source population

The parameters of NSBH binaries most relevant for GW detection comprise the masses and spins of the BH and NS. In this chapter we neglect any misalignment between the spins and \mathbf{L} , i.e. we assume that our NSBH binaries do not precess and focus instead on the spin components along \mathbf{L} . Therefore, we only need to consider the two masses m_{BH} and m_{NS} and the projections of the two dimensionless spins along \mathbf{L} , $\chi_{\text{BH}}^{\parallel}$ and $\chi_{\text{NS}}^{\parallel}$.

As we have seen in chapter 1, the mass distribution of stellar-mass BHs is not well known. Based on the results of [19], for our simulated signals we assume here a Gaussian distribution for m_{BH} with mean $7.8M_{\odot}$. However, since that value is appropriate for low-mass X-ray binaries, somewhat arbitrarily we adopt here a larger standard deviation of $3M_{\odot}$ and we truncate the resulting distribution to $3 \leq m_{\text{BH}}/M_{\odot} \leq 12$.

The population of spin parameters of stellar-mass BHs is also not well constrained. The available measurements now span the whole magnitude range allowed by the Kerr metric, as shown in figure 1.1. Population synthesis models for NSBH binaries suggest that a large fraction of BHs may actually be almost maximally spinning [26]. Assuming alignment between the BH spin and \mathbf{L} , we choose a uniform distribution for $\chi_{\text{BH}}^{\parallel}$ between ± 0.99 . Note that this assumes equal probability of spins aligned and anti-aligned with \mathbf{L} , but existing simulations suggest that anti-aligned systems are less likely [148–150]. In this chapter we shall thus consider a few different subsets of the uniform spin range and present results for each of them.

NS masses have been measured precisely thanks mainly to pulsar observations. The observed pulsar distribution is peaked at $m_{\text{NS}} \approx 1.4M_{\odot}$, with minimum mass $\approx 1M_{\odot}$ and maximum $\approx 2M_{\odot}$ [3]. In GR, an object can be in hydrostatic equilibrium only if

its mass is below a maximum value and for NSs this limit depends on the EOS and on the spin. NSs heavier than $\approx 3M_\odot$, however, are extremely unlikely [7]. Based on [9] we assume here a Gaussian distribution for m_{NS} with mean $1.35M_\odot$ and standard deviation $0.13M_\odot$ and we impose the hard limits $1 \leq m_{\text{NS}}/M_\odot \leq 2$.

By definition, pulsar timing also enables an exquisitely precise measurement of NS angular velocities. Unfortunately, because the EOS of self-gravitating neutron matter is not known, we have no handle on the moment of inertia of NSs and thus the angular momentum is not precisely known. In principle there is no fundamental limit to the spin as in the case of BHs. Nevertheless, beyond a maximum spin the NS will no longer be able to sustain the rotation and will break apart. This limit is thought to be $\chi_{\text{NS}} \approx 0.7$ for realistic EOS. The maximum observed NS spin is however only $\chi_{\text{NS}} \approx 0.4$ and when NSs in binaries are considered, the maximum observed spin—found in the double pulsar J0737-3039—is only $\chi_{\text{NS}} \approx 0.05$. A spin of 0.4 likely requires recycling the NS via accretion, which is unlikely to occur if the companion is a BH. We thus adopt a uniform distribution of $\chi_{\text{NS}}^{\parallel}$ between ± 0.05 .

The sky position and the orientation of the orbit are chosen to be uniform on the sphere. A physical choice for the distance distribution is one that leads to a uniform distribution of the binaries in a certain volume. However, this makes a large fraction of the signals too weak to be detectable, which can complicate the estimation of the sensitivity of the search. Moreover, binaries with a higher chirp mass are intrinsically louder and thus tend to dominate the sensitivity estimation. To compensate against these issues, we choose the distances in order to achieve a uniform distribution in the *chirp distance*, defined as

$$\mathcal{R} := r \left(\frac{\mathcal{M}_{\text{BNS}}}{\mathcal{M}} \right)^{5/6} \quad (2.1)$$

where $\mathcal{M}_{\text{BNS}} \approx 1.22M_\odot$ is the chirp mass of a canonical $1.4M_\odot + 1.4M_\odot$ BNS system. As described later, a uniform distribution in volume is restored *a posteriori* by applying appropriate weights to the signals when computing the sensitivity. In order to choose the limits of the chirp-distance distribution, we need to consider that below a certain minimum distance the sources are too rare and above a maximum distance they are frequent but too quiet to be detected. Based on this consideration we set the chirp distance uniform in the range [1, 160] Mpc. We demonstrate the safety of this choice later on.

The final moments of the binary can see either a rapid plunge of the whole NS into the BH, followed by a ringdown of the latter, or the disruption of the NS by tidal forces, spreading the NS matter around the remnant. The transition between these two scenarios depends mainly on the BH parameters and on the NS EOS [151]. Tidal disruption would likely attenuate the GW emission before merger and the resulting signal would have a different termination than the merger- ringdown expected for a BBH without matter effects. NSBH inspiral waveforms including such effects have been studied, but are not readily available at the time of this study [152, 153]. In this chapter we neglect the final part of the waveform, assuming that the majority of the SNR is given by the pN inspiral component. This should be a safe approximation across most of

the parameter space spanned by our simulated signals, but it may no longer be accurate for heavier systems ($M \gtrsim 15M_{\odot}$) whose merger phase may occur at frequencies within the most sensitive band of the detectors, i.e. ≈ 300 Hz or below (figure 1.4). We also neglect the effect of tidal deformation on the pN phasing, as this is unlikely to have a large effect on search pipelines [154].

We simulate our sources using the SpinTaylorT2 approximant from the LALSimulation library [60]. As described in the introduction, other pN approximants are available, for instance SpinTaylorT1 and SpinTaylorT4; however, SpinTaylorT2 is essentially the time-domain version of TaylorF2—which we use for our templates—and should therefore set aside issues related to approximant comparison, which are not in the scope of this study. We employ orbital pN phase corrections up to order 3.5 and aligned spin-orbit and spin-spin terms up to order 2.5. Amplitude corrections up to order 1.5 are also included. The starting frequency of the SpinTaylorT2 time-domain integration is in all cases set to 20 Hz, which is a safe choice because we expect interferometers to be much less sensitive here than at higher frequencies, at least in their early advanced configuration (figure 1.4), and because 20 Hz is well below the integration band of our matched filters, as described later. At the time of this work, the SpinTaylorT2 implementation in LALSimulation stops the time-domain integration of the waveform at the *minimum-energy circular orbit (MECO)* i.e. as soon as a minimum in the center-of-mass energy is discovered ($dE/dv = 0$). The MECO frequency depends on the mass and spin parameters and at fixed masses it increases with $\chi_{\text{BH}}^{\parallel}$, in some cases reaching several kHz for $\chi_{\text{BH}}^{\parallel} \approx 1$ or ceasing to exist altogether [155]. For such reasons, the integration is terminated before the MECO if an unphysical regime is reached, e.g. $v/c > 1$ or the frequency starts decreasing.

2.3 Simulated data set

In principle, our simulated signals could be added to stationary Gaussian noise with a specified noise PSD and this synthetic data set could be used to test the ability of our search pipeline to recover the signals. We have seen however that real interferometer data contain many kinds of noise artifacts, such as glitches and quasi-stationary lines, and the PSD itself is indeed time-dependent. Due to the sub-optimality of the matched-filter approach in the presence of non-stationary and non-Gaussian noise, when the parameter space of a search pipeline is enlarged to include spin, these artifacts may lead to a disproportionate increase in false alarms from the newly-included region of parameter space. It is not obvious that the existing signal-based vetoes should continue to work in the new region. In order to have a fair evaluation of the inclusion of spin in our search, we thus need a data set which realistically reproduces the conditions expected in data from advanced detectors, i.e. a realistic noise PSD as well as a broad variety of glitches and noise artifacts.

The LIGO collaboration constructed such a data set in order to facilitate comparisons like the one presented here (see e.g. [156]). We use this data set to test our pipeline. The set is generated starting from six months of real data from the sixth (last) scientific

run of the initial Hanford and Livingston LIGO detectors, in particular between GPS times 966388158 (Sat Aug 21 01:09:03 GMT 2010) and 971528001 (Tue Oct 19 12:53:06 GMT 2010). The noise PSD of these data is then measured and modified (*recolored*) to the “early advanced-LIGO” curve shown in figure 1.4. This produces six months of data with a PSD representative of the first runs of advanced LIGO, while preserving all kinds of noise artifacts already observed in initial LIGO. Because of this procedure, the result is called *recolored data*. In our test we use the recolored data segments flagged as `H1:CBC-MDC1_SCIENCE_EARLY_RECOLORED:2` and `L1:CBC-MDC1_SCIENCE_EARLY_RECOLORED:2` corresponding to epochs when the detectors were in “science” mode during their sixth scientific run.

There is obviously no guarantee that advanced detectors will be affected by the same population of artifacts as the initial ones, especially given the additional complexity introduced by the advanced design. Nevertheless, recoloring existing data is a simple method to extrapolate existing knowledge and the alternative would require complicated modeling of noise artifacts. Another caveat is that the calibration of initial-LIGO data is considered unreliable below ≈ 40 Hz and, as we shall see shortly, we use here a lower starting frequency for our matched filters. However, the recolored data only provide a proxy for the noise and we simply add our simulated signals to them, so the reliability of the calibration does not affect our study. Indeed, the poor stationarity of the recolored data at low frequency is likely an exaggeration with respect to advanced LIGO, which is designed to be reliable at those frequencies. Our result should therefore be conservative in terms of data quality at low frequency.

2.4 The PyCBC toolkit

CBC search pipelines for initial interferometers were implemented as a Python infrastructure known as *iHope*, which was part of LAL [60]. The most compute-intensive operations, i.e. the population of the template banks and the generation of single-detector triggers via the FindChirp algorithm, were instead implemented in C, respectively with the `lalapps_tmplbank` and `lalapps_inspiral` programs. The most expensive operations in this context were the FFTs used for computing the SNR and χ^2 statistics across coalescence time, because they were performed for each template and data segment. These used the highly-optimized FFTW library to achieve maximum performance [157].

Searches in advanced-interferometer data are likely to have even higher computational requirements than the past. On the one hand, the larger bandwidth of advanced interferometers leads to longer template waveforms and significantly larger template banks, which directly increases the number of FFTs that need to be performed in order to produce single-detector triggers. On the other hand—as is the case here—the development of more sophisticated waveform approximants extends the parameter space which can be explored and again leads to significantly larger template banks.

The last decade has seen an increase in computing power driven significantly by parallel computing architectures other than the traditional improvement in the performance of single central processing unit (CPU) cores. The introduction of multi-core and

many-core architectures, and in particular general-purpose computing on GPUs, has enabled massive acceleration of trivially-parallel algorithms [158]. In the last few years, LSC computing clusters have been equipped with various models of Nvidia GPUs. In order to meet the computational challenge of advanced detectors, it is important that the implementations of CBC search pipelines can leverage modern high-performance computing technologies, in particular multi-core CPUs and GPUs.

Since the dominant operation in past searches was the FFT in `lalapps_inspiral`, a reasonable attempt to accelerate this operation is by switching from the FFTW library to the Compute Unified Device Architecture (CUDA) FFT implementation, thus simply utilizing the GPU as an FFT accelerator. We quantify the improved performance by running the modified `lalapps_inspiral` on 2048 s of Gaussian noise. We use different SNR thresholds in order to evaluate the variable cost of the χ^2 calculation, which in `lalapps_inspiral` depends on the rate of single-detector triggers. This simple test shows that the GPU speeds up `lalapps_inspiral` to an extent which depends on how often the χ^2 test must be computed: when the SNR threshold is high, and the χ^2 is calculated rarely, we obtain a modest speed-up factor of ≈ 3.5 . On the other hand, when the χ^2 veto is calculated virtually all the time, the speed-up is as large as a factor of ≈ 20 . By profiling `lalapps_inspiral` we also observe, however, that the GPU is idle for a significant fraction of the time. This is because data need to be copied from the host to the GPU memory, Fourier-transformed, and then copied back to the host; unfortunately, data transfers between the host and the GPU are generally inefficient and should be avoided as much as possible. Data transfers can be minimized by moving the whole FindChirp algorithm to the GPU, such that the only transfers involve the preconditioned strain data and the final list of single-detector triggers. This should maximize the benefit of using the GPU and achieve a higher speed-up. Unfortunately, it also requires reimplementing and testing a non-trivial amount of code in the CUDA language. Moreover, other GPU and many-core architectures can potentially provide similar accelerations, for instance those based on OpenCL. A purely CPU-based implementation is also required in any case for systems with no GPUs. This would require maintaining several different implementations of the same algorithm.

In addition to this issue, the C implementations of `lalapps_tmpltbank` and `lalapps_inspiral` are not modular and flexible enough to easily accommodate changes that may be required with advanced detectors, for instance the introduction of templates with significantly higher-dimensional parameter spaces. This prompts a reimplementing of these tools with the Python language, which is a convenient, free, high-level language with a large community and wide adoption in the scientific computing world. The resulting toolkit, called PyCBC, is an evolution of the LAL-based software used in past CBC searches. PyCBC is developed primarily by the Syracuse LIGO group with contributions from us and is what we use here for implementing our NSBH pipeline.

PyCBC is able to benefit from GPU acceleration without having to maintain multiple implementations of the same algorithm. First of all, it uses the PyCUDA and PyOpenCL modules to wrap the architecture-specific details in a uniform Python interface. Should new compute architectures appear in the future, those could thus be added easily. On

the other hand, by introducing the concept of a *computing scheme*, PyCBC associates different implementations of a particular Python function to the different computing platforms; once the desired scheme is selected, the appropriate implementation of each function of the algorithm is called automatically and transparently. This means that the *same* Python script can be executed by different computing back-ends, for instance the CPU and a CUDA GPU (figure 2.1). Importantly, data transfers between the CPU and GPU are only performed when needed, i.e. when a data block residing on the CPU (GPU) is used by a function running on the GPU (CPU). This enables an efficient GPU implementation of FindChirp in simple Python. The same mechanism can be used to support different implementations of a particular feature—an important example being the FFT—allowing one to select the most efficient implementation depending on the particular CPU architecture.

Although PyCBC re-implements some of the functionalities of the old LAL-based software, it still relies on LAL where appropriate, e.g. for generating inspiral waveforms and for input and output to LAL file formats. This interfacing is achieved in a simple way via the SWIG framework. Finally, the Python language benefits from a large library of useful packages. These include scientific computing packages such as Numpy, Scipy and Matplotlib; file-format interfaces such as HDF5; and database interfaces such as SQLite. These technologies can then be used directly and easily in CBC pipelines.

2.5 Template bank construction and test

In order to carry out the comparison between neglecting and including aligned-spin effects, we construct two similar template banks which differ only in the extent of the covered NSBH parameter space: a non-spinning bank (NSB) and an aligned-spin bank (ASB). In both cases we use the restricted frequency-domain TaylorF2 approximant for our templates, similarly to what was done in past low-mass CBC searches. As we discussed already for our simulated signals, the use of inspiral-only waveforms is justified by the fact that, in the mass range we explore, the merger and ringdown happen at relatively high frequencies, where the early advanced detectors are not expected to have the maximum sensitivity. We include pN corrections to the orbital phase up to order 3.5. For the ASB, we include spin-orbit and spin-spin terms up to order 2.5¹. Since the low-frequency sensitivity of advanced LIGO is higher than initial LIGO, we use a starting frequency of 30 Hz for our templates, as opposed to 40 Hz in past searches; the early advanced LIGO configuration is not expected to have significant sensitivity and trustworthy calibration below 30 Hz. For template placement we assume a fixed termination of the templates at 1000 Hz, close to the largest ISCO frequency in our parameter space. This is necessary in order to have a constant metric across the whole parameter space and facilitate template placement². Note however that in generating

¹This was the highest implemented order at the time of the study; 3.5 pN spin terms are now implemented and can be used.

²PyCBC’s template bank generator code later included the ability to vary the termination frequency, but we do not explore this choice here.

File `inspiral.py`: platform-independent matched-filter program

```
scheme = CPUScheme()
template_bank = load_template_bank()
with scheme:
    data_segments = prepare_data()
    for params in template_bank:
        template = generate_template_waveform(params)
        for segment in data_segments:
            product = template * segment
            snr = abs(inverse_fft(product))
            triggers = find_triggers(snr)
            store(triggers)
```

File `waveform.py`: common interface to waveform generator

```
@schemed
def generate_template_waveform(params):
    # empty function definition
```

File `waveform_cpu.py`: CPU-specific waveform generator

```
def generate_template_waveform(params):
    compute waveform with Python/Numpy primitives
```

File `waveform_cuda.py`: CUDA-specific waveform generator

```
def generate_template_waveform(params):
    compute waveform with CUDA primitives
```

Figure 2.1: Simplified PyCBC pseudocode illustrating how the same implementation of the matched-filter algorithm can be executed on different platforms by using the concept of computing schemes. A different computing platform, e.g. a CUDA GPU, could be used by simply instantiating the appropriate scheme class in place of the `CPUScheme`; the functions within the `with` block pick the appropriate implementation based on the selected scheme. The selection is performed by the `schemed` Python decorator, which imports the actual function implementation from the appropriate file (in the case of `generate_template_waveform()`, either `waveform_cpu.py` or `waveform_cuda.py`) based on the global state of the scheme manager.

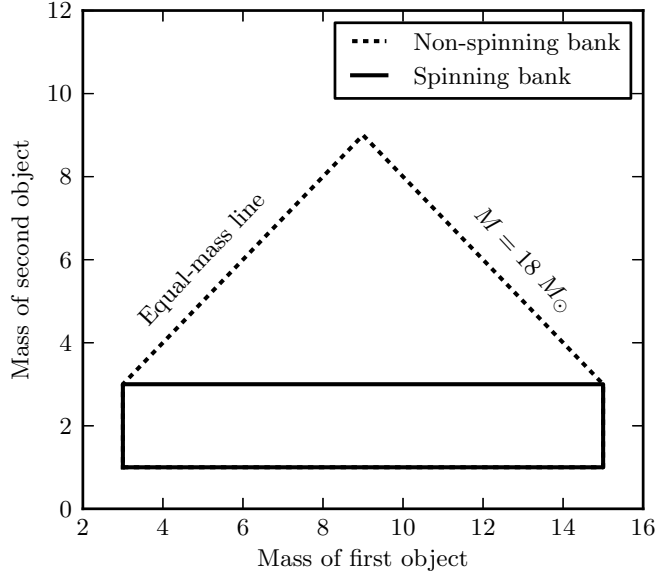


Figure 2.2: Mass boundaries used in our NSBH template banks.

the actual template waveforms for the matched-filter calculation we shall use instead the ISCO termination condition, consistent with previous CBC searches. Both banks are populated with PyCBC’s stochastic placement algorithm, described in [117].

The NSB uses a BH mass m_{BH} ranging from $3M_{\odot}$ to $15M_{\odot}$ and a NS mass m_{NS} ranging from $1M_{\odot}$ to the equal-mass boundary $m_{\text{BH}} = m_{\text{NS}}$ (figure 2.2). This mimics a “low-mass” search space similar to what used in past searches. Because of the degeneracy between spin and symmetric mass ratio [159], extending the mass range allows the NSB to recover spinning signals by introducing a bias in the recovered symmetric mass ratio. Both spins in the NSB are set to zero. With these settings, the NSB contains $\approx 2.8 \times 10^4$ templates; only $\approx 6\%$ of those are in the “extended” region $m_{\text{NS}} > 3M_{\odot}$.

The ASB also uses a m_{BH} range from $3M_{\odot}$ to $15M_{\odot}$, but m_{NS} is restricted instead to the expected NS masses, $[1, 3]M_{\odot}$ (figure 2.2). The parallel component of the dimensionless BH spin $\chi_{\text{BH}}^{\parallel}$ ranges from -1 to $+1$ and the parallel component of the NS spin $\chi_{\text{NS}}^{\parallel}$ ranges from -0.4 to $+0.4$. These settings result in $\approx 1.5 \times 10^5$ templates, a factor of ≈ 5 more than the NSB and a significant increase with respect to past searches. The distribution of templates in the ASB is shown in figure 2.3. The last plot shows the density in coordinates which make the metric approximately flat, i.e. the *chirp times* [160]

$$\begin{aligned} \tau_0 &:= \frac{5}{256\pi\eta f_0} (\pi G c^{-3} M f_0)^{-5/3} \\ \tau_3 &:= \frac{1}{128\pi\eta f_0} (\pi G c^{-3} M f_0)^{-2/3} \left[16\pi - \frac{1}{6} \chi_{\text{BH}}^{\parallel} (19\delta^2 + 113\delta + 94) \right] \end{aligned} \quad (2.2)$$




Figure 2.3: Density of templates in the aligned-spin bank. In the last plot, the black contour shows the extent of the zero-spin region of NSBH space. The dashed contour is the extra mass freedom ($m_{\text{NS}} > 3M_{\odot}$) included in the non-spinning bank. It can be seen that the region of highest density of the aligned-spin bank corresponds to small masses and extremal values of the BH spin.

with $\delta := (m_{\text{BH}} - m_{\text{NS}})/M$ and $f_0 := 20$ Hz a fiducial frequency. It can be seen that the majority of the templates are concentrated at the boundaries of maximum spin magnitude.

Before using the banks in the full search pipeline, it is useful to characterize their average ability to recover NSBH signals in stationary Gaussian noise. This can be done by numerically computing the fitting factor (eq. (1.35)) of each bank with simulated signals from our NSBH population (*bank simulation*). We use PyCBC’s bank simulation utility for this task.

We first consider the case of a “typical” NSBH signal with masses fixed at the center of our distributions ($m_{\text{BH}} = 7.8M_{\odot}$ and $m_{\text{NS}} = 1.35M_{\odot}$), no spin on the NS and variable $\chi_{\text{BH}}^{\parallel} \in [-1, 1]$. Figure 2.4 shows the fitting factor between such a signal and our NSB constructed assuming different noise models: the initial-LIGO, early-advanced-LIGO

and design-advanced-LIGO PSDs from figure 1.4 and an old advanced-LIGO model from Cutler and Flanagan [45]. It can be seen that, irrespective of the noise model, there is a region of small BH spin for which the NSB is able to satisfy the minimal-match requirement used for its construction. In other words, the NSB is effectual for BHs with small spin. This is because the effect of the small spin is compensated by the degeneracy between spin and symmetric mass ratio. As soon as the spin hits a critical value, however, the bias in η drives the parameters of the best-matching template outside of the mass range covered by the bank and the fitting factor quickly drops. The critical value ≈ 0.4 for positive spin corresponds to hitting the physical equal-mass boundary of the bank at $\eta = 1/4$. For negative spins the critical value is ≈ -0.4 and corresponds to hitting the artificial $m_{\text{NS}} > 1M_{\odot}$ boundary. In principle, the performance of the NSB at high anti-aligned spins could be improved by searching for lighter NSs. By also looking at the mass parameters of the best matching template as a function of $\chi_{\text{BH}}^{\parallel}$, we can confirm that the drop in fitting factor is due to hitting the $\eta = 1/4$ boundary at positive spin and the $m_{\text{NS}} = 1M_{\odot}$ boundary at negative spin (figure 2.5). When the NSB is replaced by the ASB, the fitting factor is always larger than 0.94 and no clear bias in the recovered parameters can be seen (figure 2.6). We can also see that the observed NS spin is essentially random, which suggests that $\chi_{\text{NS}}^{\parallel}$ effectively plays no role in our NSBH waveforms and can be safely ignored.

As an overall test of the performance of the two banks, we perform fitting-factor calculations by varying the signal parameters uniformly over our NSBH parameter space and assuming the early-advanced-LIGO noise model. The resulting worst-case fitting factor over the parameter space is shown in figure 2.7. In the case of the NSB, we again see that the bank performs well for sufficiently small BH spin and a large loss suddenly appears for higher spin magnitudes. The critical spin values depend on η ; for instance, for equal-mass systems there is no freedom to compensate any amount of positive $\chi_{\text{BH}}^{\parallel}$ and the fitting factor starts to drop already at $\chi_{\text{BH}}^{\parallel} \gtrsim 0$. We also see that the worst performance happens at large mass-ratio and extreme anti-aligned BH spin; this is the region of NSBH space where signals are shortest. The ASB is also performing poorly in this region, with fitting factors well below the minimum-match condition of 97%. The residual loss here is likely due to the fact that neither our templates nor the simulated SpinTaylorT2 signals include a physical termination of the signal. In fact, in this region the merger and ringdown happen in a very sensitive band of the early advanced LIGO noise PSD; the difference between the frequencies associated with ISCO and MECO is also the largest here. Nevertheless, this first calculation shows that including spin in the bank has a large potential for increasing the sensitivity.

2.6 Pipeline configuration and execution

The next step is assessing the behavior of the ASB with realistic noise and instrumental artifacts and evaluating the sensitivity of a full search pipeline based on it. At the time of this study, however, the existing pipeline does not allow one to perform coincidence with




Figure 2.4: Fitting factor of a typical signal from our NSBH population and our non-spinning bank, as a function of the BH spin component along the orbital angular momentum. Different plots assume different detector noise models (see text). In all cases, the non-spinning bank is unable to recover a significant fraction of the SNR when the magnitude of the BH spin is larger than ≈ 0.4 .




Figure 2.5: Bias in recovered mass parameters for a typical signal from our NSBH population observed with our non-spinning bank, as a function of the BH spin component along the orbital angular momentum. The dashed lines show the true mass parameters of the simulated signal. We can see that the drop in fitting factor at positive spin is associated with the $\eta = 1/4$ boundary (top right) while the drop at negative spin is due to the $m_{\text{NS}} = 1M_{\odot}$ boundary (bottom right).




Figure 2.6: Recovered mass and spin parameters for a typical signal from our NSBH population observed with our aligned-spin bank, as a function of the BH spin component along the orbital angular momentum. The dashed lines show the true parameters of the simulated signal. Although the fitting factor at large spin magnitude is slightly smaller than the required minimum value of 97%, there are neither clear biases in parameter recovery nor visible boundary effects as seen in the non-spinning bank. It can also be seen that the observed NS spin is essentially uniformly distributed over the range covered by the bank.




Figure 2.7: Performance of the non-spinning (left) and aligned-spin (right) banks over our NSBH parameter space. The color of each tile shows the worst-case fitting factor in the tile. Note that the color scales of the two banks are different. Although the fitting factor in the aligned-spin bank is still below 0.97 in some regions of the space, it is much higher than in the non-spinning bank when the spin magnitude is large.

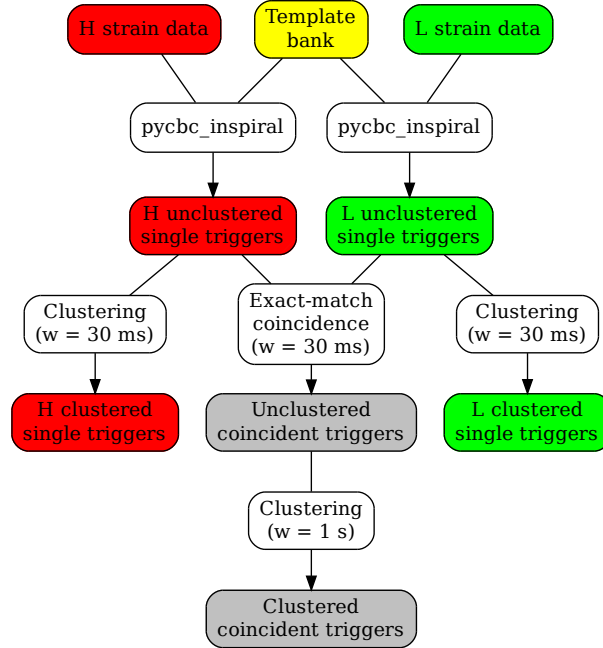


Figure 2.8: Flowchart of the search pipeline.

spinning templates in an easy way and it is not readily compatible with `pycbc_inspiral`, the implementation of the FindChirp algorithm provided by PyCBC. It is also not suited for experimentation with the additional degrees of freedom. We thus implement a simple prototype pipeline which can be configured to run on both the NSB and ASB with otherwise identical settings and on the same data and simulated signals. This enables a direct comparison of templates neglecting and excluding spin effects. Similarly to what was done in past CBC searches, we implement the pipeline in the Python language and we use HTCondor to split the analysis into many parallel jobs, which are executed on the Atlas computing cluster [161]. A flowchart of the pipeline is shown in figure 2.8.

2.6.1 Matched filtering and trigger generation

The first step of the analysis is breaking the strain data from the Hanford and Livingston interferometers into 2048 s segments, which are then analyzed independently by `pycbc_inspiral` to produce single-detector triggers. One such run is performed with the pristine data in order to produce the false-alarm background. A second run, with identical settings, is performed by first adding the simulated signals from the NSBH population to the data.

We configure the FindChirp algorithm in the following way. The noise PSD is estimated for each data segment by using the median-Welch method applied to segments with a duration of 64 s. The impulse response of the corresponding whitening filter

is then truncated to a maximum duration of 16 s. The SNR threshold for generating triggers is set to 5.9, which is higher than the typical 5.5 used in past searches. This choice is made in order to keep the number of resulting triggers reasonably small even when the large ASB is used and should have a small effect on the sensitivity. Triggers from each template are maximized over time by using a window of 1 s. The χ^2 statistic is then calculated for each trigger using 16 bands. We also apply a threshold of 5.9 on the reweighted SNR of the triggers (eq. (1.38)), which strongly reduces the number of triggers, especially in the presence of loud glitches.

Thanks to the flexibility of `pycbc_inspiral`, we are able to execute the jobs on the Nvidia C2050 GPUs currently available in the Atlas cluster; at the time of the study, this is able to accelerate the FindChirp algorithm by a factor which depends on the number of triggers generated for each data segment, with a maximum speedup of $\mathcal{O}(10)$ times with respect to running the same job on a single core of the node’s CPU when the data have good stationarity and Gaussian behavior.

2.6.2 Coincidence

The next step of the analysis is performing coincidence between the Hanford and Livingston detectors. We implement an exact-match coincidence algorithm: a coincident trigger is generated whenever both interferometers produce a trigger from the same template and with coalescence times within 15 ms. This value is based on the maximum light travel time between the Hanford and Livingston sites (≈ 10 ms) and a margin based on the typical uncertainty in coalescence time (a few milliseconds). Each coincidence is ranked via the combined reweighted SNR statistic defined in eq. (1.39) and already used in previous searches. Triggers from the run with no simulated NSBH signals are used to generate the false-alarm background via time slides. We perform 800 time shifts, all multiples of 5 s.

Although exact-match coincidence is simple and straightforwardly applicable to arbitrary parameter spaces, it requires storing and analyzing a very large number of single-detector triggers, especially when the ASB is used. In addition, triggers are produced and written independently by hundreds of jobs operating simultaneously on different regions of the search space. The access pattern of coincidence jobs, however, is quite different: each parallel coincidence jobs must quickly read all triggers for a particular template. The natural storage technology for triggers is thus an online database, which handles parallel read and write operations efficiently and intuitively. We test three different database technologies: PostgreSQL, MongoDB and Cassandra.

PostgreSQL is a well-established and widely-used relational database based on the structured query language (SQL) [162]. Our trigger storage problem maps very naturally to a relational model. We are able to efficiently implement coincidence and clustering by using a table for single-detector triggers, where each row is tagged with template and detector IDs (figure 2.9). This allows parallel coincidence jobs to rapidly query all triggers for a particular template and perform coincidence on them, independently from other templates. A second table is used for storing coincident triggers. A potential limitation of PostgreSQL is scalability to a much larger number of single-detector triggers,

```

-- template bank
CREATE TABLE templates (
    id SERIAL PRIMARY KEY,
    mass1 REAL,
    mass2 REAL,
    mchirp REAL,
    tau0 REAL,
    tau3 REAL,
    eta REAL,
    s1 REAL,
    s2 REAL
);

-- single-detector triggers
CREATE TABLE triggers (
    id SERIAL PRIMARY KEY,
    ifo TEXT,
    template INTEGER,
    coal_time DOUBLE PRECISION,
    coal_phase REAL,
    eff_distance DOUBLE PRECISION,
    snr REAL,
    rx2 REAL,
    newsnr REAL,
    FOREIGN KEY(template)
        REFERENCES templates(id)
);

-- coincident triggers
CREATE TABLE coinc_triggers (
    id SERIAL PRIMARY KEY,
    template INTEGER,
    lag REAL,
    coal_time_h DOUBLE PRECISION,
    snr_h REAL,
    rx2_h REAL,
    newsnr_h REAL,
    coal_phase_h REAL,
    eff_distance_h DOUBLE PRECISION,
    coal_time_l DOUBLE PRECISION,
    snr_l REAL,
    rx2_l REAL,
    newsnr_l REAL,
    coal_phase_l REAL,
    eff_distance_l DOUBLE PRECISION,
    FOREIGN KEY(template)
        REFERENCES templates(id)
);

-- coincident triggers
-- surviving clustering
CREATE TABLE coinc_clusters (
    trigger INTEGER,
    FOREIGN KEY(trigger)
        REFERENCES coinc_triggers(id)
);

```

Figure 2.9: SQL tables used for storing single and coincident triggers in our PostgreSQL implementation of the pipeline. In `coinc_triggers`, columns ending in “h” and “l” store quantities related to the Hanford and Livingston detectors respectively.

for instance for the significantly larger template banks expected with final advanced detectors. However, we do not experience such a limitation in our study.

MongoDB [163] is a so-called “NoSQL” database and is able to store *collections* of *documents*, where each document is an independent set of key-value associations. Thus, unlike relational databases which enforce a uniform row definition and normalization of the data, MongoDB provides a more flexible “schema-less” storage. Also, MongoDB is a distributed system which can scale to very large data sets simply by adding more physical nodes to its underlying cluster. Although it is possible to use such a schema for our problem, we find that MongoDB is not as convenient as a relational database and requires a significantly larger storage space than PostgreSQL. In addition, the structure of documents must typically be designed with a particular query in mind in order to be efficient, which in our case limits the flexibility of the system when triggers need to be accessed and explored in different ways. It is possible however that more recent

Collection <i>run</i>	
mc	Chirp mass
eta	Symmetric mass ratio
s1	BH spin along \mathbf{L}
s2	NS spin along \mathbf{L}
i	Detector
t	Coalescence time
snr	SNR
rx2	Reduced χ^2
newsnr	Reweighted SNR
phi	Coalescence phase
ed	Effective distance
c	Survives clustering

Collection <i>run_coinc</i>	
mc	Chirp mass
eta	Symmetric mass ratio
s1	BH spin along \mathbf{L}
s2	NS spin along \mathbf{L}
H	t Hanford coalescence time
	snr Hanford SNR
	rx2 Hanford reduced χ^2
L	t Livingston coalescence time
	snr Livingston SNR
	rx2 Livingston reduced χ^2
snr	Combined SNR
newsnr	Combined reweighted SNR
delay	Time shift between detectors
c	Survives clustering

Figure 2.10: Collections and document structure used for representing single and coincident triggers in our MongoDB implementation of the pipeline.

versions of MongoDB could address these limitations. The definitions of the collections and documents in our MongoDB implementation of the pipeline are given in figure 2.10.

Cassandra [164] is also a NoSQL distributed database but it uses a tabular data model and an SQL-like interface, although the set of possible queries is limited in such a way to prevent inefficient queries. This model enables an efficient implementation of coincidence, with performance similar to PostgreSQL. Similarly to MongoDB however, the data model has to be designed with a particular set of queries in mind in order to be efficient (figure 2.11). Again this complicates accessing the triggers with arbitrary queries. Therefore, the traditional relational model may be the most suited for our search pipeline. Testing these systems with a few orders of magnitude more triggers could enable a clearer decision on which one is more suitable.

2.6.3 Clustering

A fluctuation in the data, especially if significantly louder than the surrounding noise background, generally does not produce a single trigger but rather a cluster of triggers associated with different templates. Loud glitches, in particular, tend to excite the whole template bank and the resulting cluster can last fractions of a second or in some cases more. This means that triggers are correlated on a time scale of $\lesssim 1$ s. Since the FAR estimation procedure assumes that triggers are statistically independent, the correlation must be removed first. This is typically done by clustering or maximizing the triggers over the template bank.

We apply the following clustering algorithm to our coincident triggers: keep a trigger


```

CREATE TABLE templates (
    id bigint,
    mass1 float,
    mass2 float,
    mchirp float,
    eta float,
    spin1 float,
    spin2 float,
    tau0 float,
    tau3 float,
    duration float,
    PRIMARY KEY ((id))
);

CREATE TABLE single_triggers (
    ifo int,
    template_id bigint,
    coal_time bigint,
    coal_time_ns bigint,
    coal_phase float,
    effective_distance double,
    snr float,
    rx2 float,
    newsnr float,
    PRIMARY KEY ((template_id, ifo),
        coal_time, coal_time_ns)
);

CREATE TABLE coinc_triggers (
    id uuid PRIMARY KEY,
    template_id bigint,
    lag float,
    coal_time_h bigint,
    coal_time_ns_h bigint,
    snr_h float,
    rx2_h float,
    newsnr_h float,
    coal_phase_h float,
    effective_distance_h double,
    coal_time_l bigint,
    coal_time_ns_l bigint,
    snr_l float,
    rx2_l float,
    newsnr_l float,
    coal_phase_l float,
    effective_distance_l double,
    comb_snr float,
    comb_newsnr float
);

```

Figure 2.11: Tables used for storing single and coincident triggers in the Cassandra implementation of the pipeline. In `coinc_triggers`, columns ending in “h” and “l” store quantities related to the Hanford and Livingston detectors respectively. Clustered coincident triggers are stored in a separate table with the same columns as `coinc_triggers`.

if no other triggers with higher combined reweighted SNR exist within ± 0.5 s of that trigger. The clustering time window is chosen to be smaller than the typical time between false-alarm coincident triggers (in order to avoid clustering triggers produced by different fluctuations) but larger than the light travel time between the detectors. Note that past searches used a slightly different clustering algorithm based on a fixed and much longer window (several seconds). Because of the moving window, the algorithm used here is able to handle clusters lasting much more than the window itself.

In order to study the distribution of single-detector background, we also apply the same algorithm to single-detector triggers, but using a window of ± 15 ms only. This is in order to safely accommodate the much larger single-detector FAR, although a window of ± 0.5 s would likely produce similar results.

2.7 False-alarm background

We now look at the background of false alarms generated by our search pipeline with the two banks. When we switch from the NSB to the ASB, the proper volume of the searched parameter space increases by a factor which is approximately given by the relative increase in number of templates, i.e. ≈ 5 . In stationary Gaussian noise, we thus expect a similar increase in the rate of false alarms of our search pipeline. However, this is no longer guaranteed when noise artifacts are present; templates with certain spin values may respond unexpectedly to glitches and sudden changes in the noise characteristics. Thus it is important to investigate the background distribution associated with the two banks and verify that the inclusion of spin does not lead to a catastrophic increase in FAR.

Figures 2.12 and 2.13 show the single-detector false-alarm distribution for the Hanford and Livingston detectors. It is immediately evident that the SNR has a large tail in both detectors, which gets noticeably larger with the ASB. For both detectors, the SNR tail in the ASB is produced mostly by the high-mass, high anti-aligned spin region of the parameter space. Interestingly, these waveforms terminate at the ISCO, which is spin-independent. It is therefore remarkable that high-mass but positively aligned templates give a smaller contribution to the background. This suggests that the pN phasing may behave pathologically for high-mass, anti-aligned templates. Nevertheless, the χ^2 veto is quite effective at suppressing glitches and the distribution of reweighted SNR is much better behaved and approximately described by an exponential, i.e. $\exp(-k\hat{\rho})$. There seems to be no preferred region of parameter space producing high- $\hat{\rho}$ triggers in the ASB and the location of the loudest $\hat{\rho}$ appears to simply follow the template density. The relative FAR increase α at fixed $\hat{\rho}$ when going from the NSB to the ASB is roughly compatible with the value expected from the metric volume. Because of the almost exponential distribution of $\hat{\rho}$, the increase in reweighted SNR required to keep the FAR constant is only $\Delta\hat{\rho} \approx \log(\alpha)/k$. Still, even for $\hat{\rho}$ the tail is slightly larger than exponential, possibly because of the lack of category vetoes and thus the presence of a few data segments with exceptionally poor stationarity. The two background events with the highest reweighted SNR in the two instruments are given in table 2.1. Both

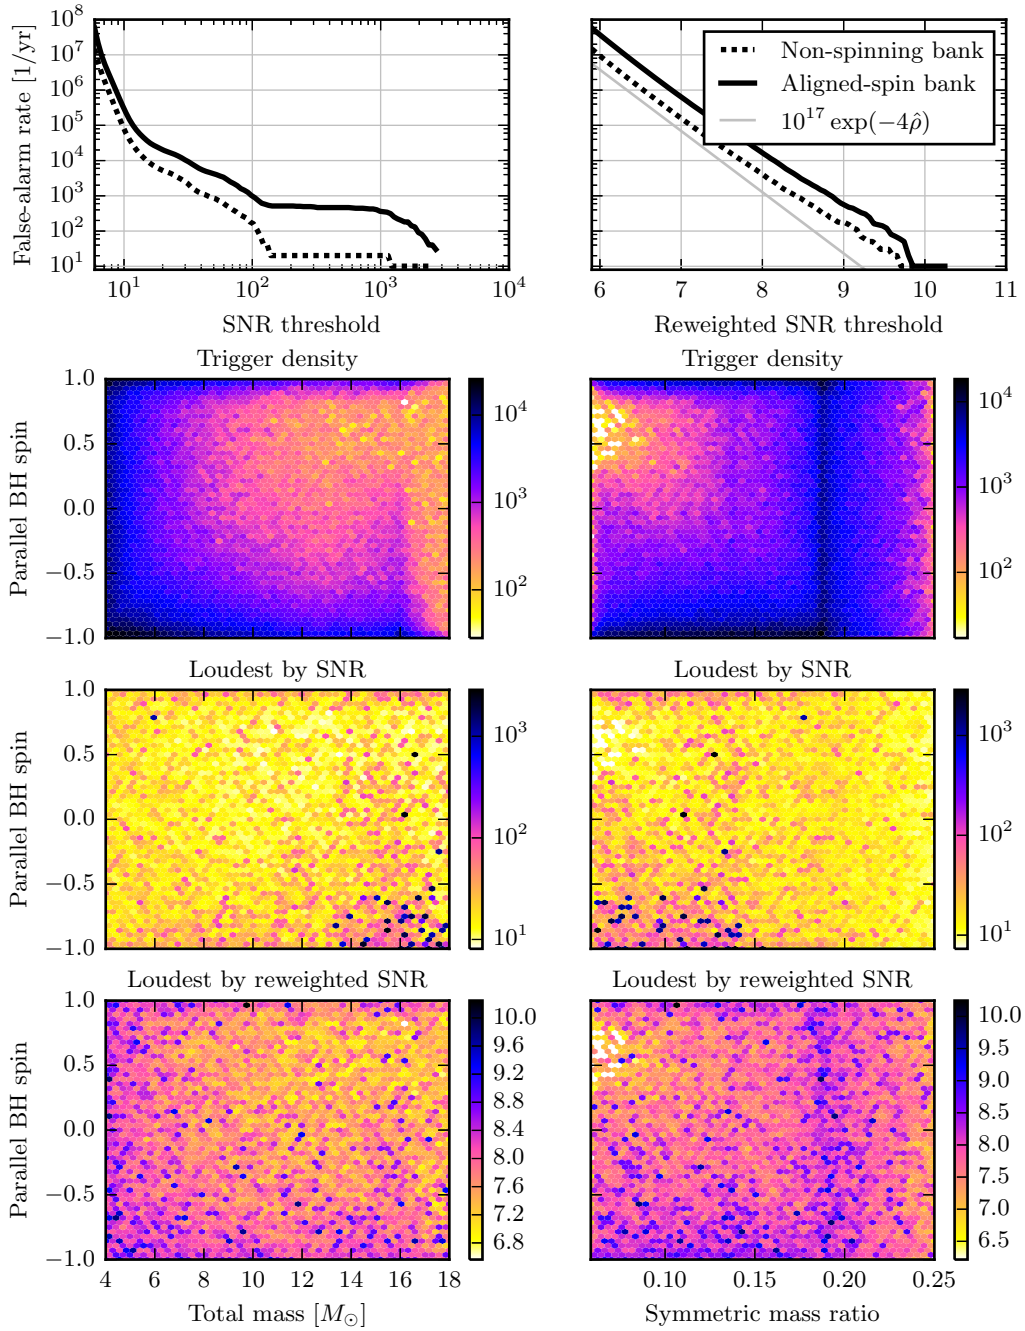


Figure 2.12: Distribution of single-detector false alarms in the Hanford interferometer. All heat maps refer to the aligned-spin bank. The reweighted SNR rate falls off almost exponentially, similarly to what expected in Gaussian noise (gray curve). The loudest triggers are not systematically associated with a particular part of the parameter space.

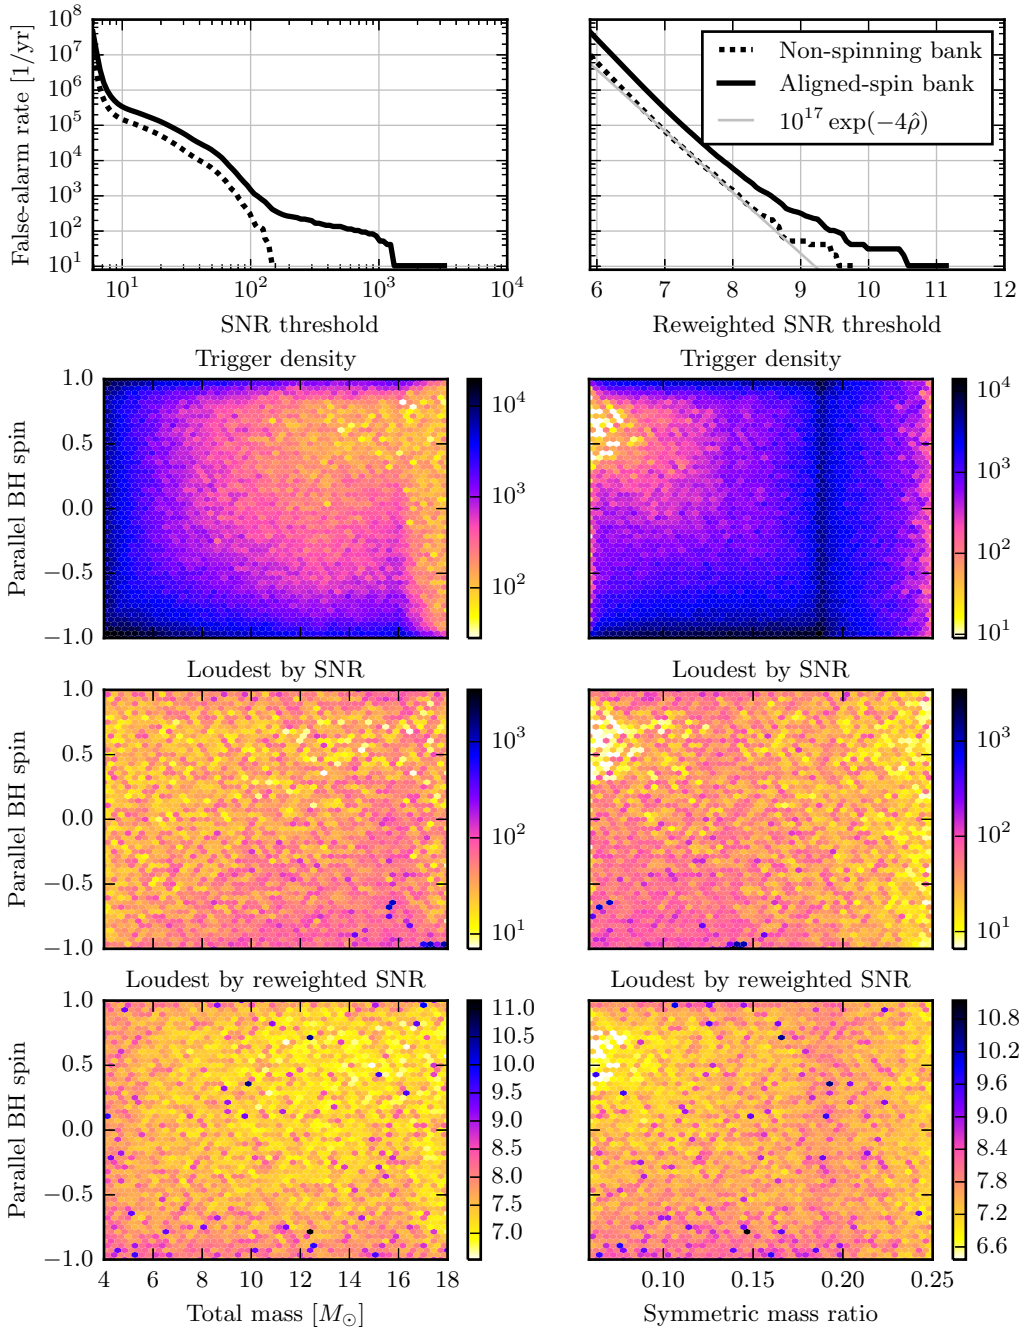


Figure 2.13: Distribution of single-detector false alarms in the Livingston interferometer. All heat maps refer to the aligned-spin bank. The reweighted SNR rate falls off almost exponentially, similarly to what expected in Gaussian noise (gray curve). The loudest triggers are not systematically associated with a particular part of the parameter space.

Interferometer	GPS end time	ρ	$\hat{\rho}$	m_{BH}	m_{NS}	$\chi_{\text{BH}}^{\parallel}$
Hanford	971 321 818.762	11.5	10.3	8.7	1.2	0.97
Livingston	967 609 107.681	36.8	11.1	10.1	2.2	-0.77

Table 2.1: Parameters of the single background triggers with the largest reweighted SNR in the aligned-spin bank.

triggers occur in the vicinity of very loud glitches (occurring 8 s before the trigger for Hanford and 5 s after the trigger for Livingston) and could possibly be related to them, although establishing this relation would require additional investigation.

In figure 2.14 we see the distribution of coincident false alarms from all the time-shifted samples. The effect of coincidence is dramatic and leads to a well-behaved distribution for both banks. The increase of the FAR at fixed threshold in combined reweighted SNR is given almost exactly by the relative increase in number of templates. The increase in threshold required to keep the FAR constant when going to the ASB is only ≈ 0.3 in the most interesting part of the curve ($\text{FAR} \lesssim 0.1/\text{yr}$). The distribution of coincident false alarms over the parameter space associated with the ASB follows what we see for the single-detector cases: their density matches the template density, the loudest combined SNR is located at the high-mass, high-aligned-spin corner and the loudest combined reweighted SNR is scattered according to the template density. The loudest single-detector triggers by reweighted SNR (table 2.1) have a combined reweighted SNR of ≈ 15.1 , but they do not survive coincidence and thus do not contribute to the tail of the coincident background. Had we chosen a larger number of time slides, we could have seen a slightly larger tail in the top plot of figure 2.14, but it is unlikely that it would have deviated significantly from the exponential part of the curve.

We can conclude that the inclusion of aligned-spin in the templates leads to an increase in false alarms which is well predicted by the increase in proper volume of the explored parameter space (cfr. eq. (1.36)). This is because the χ^2 test is effective at separating glitches from aligned-spinning waveforms, despite the presence of very poor data segments. Templates with high total mass and $\chi_{\text{BH}}^{\parallel} < -0.5$ respond more strongly than others to glitches and produce most of the loudest SNRs, but they are not entirely responsible for the tail of the reweighted SNR, at least in the region of the parameter space considered here and with the chosen number of χ^2 bins.

2.8 Recovering simulated signals

The runs containing simulated signals are used to estimate the ability of the pipeline to recover signals from different parts of the parameter space. We define a simulated signal as found if the search pipeline produces a coincident trigger with coalescence time within ± 0.5 s of the simulated value and chirp mass within $\pm 0.6M_{\odot}$ of the simulated value. If no trigger is present within this recovery window, the signal is missed. The window is based on the results of the bank simulations described earlier. Each found

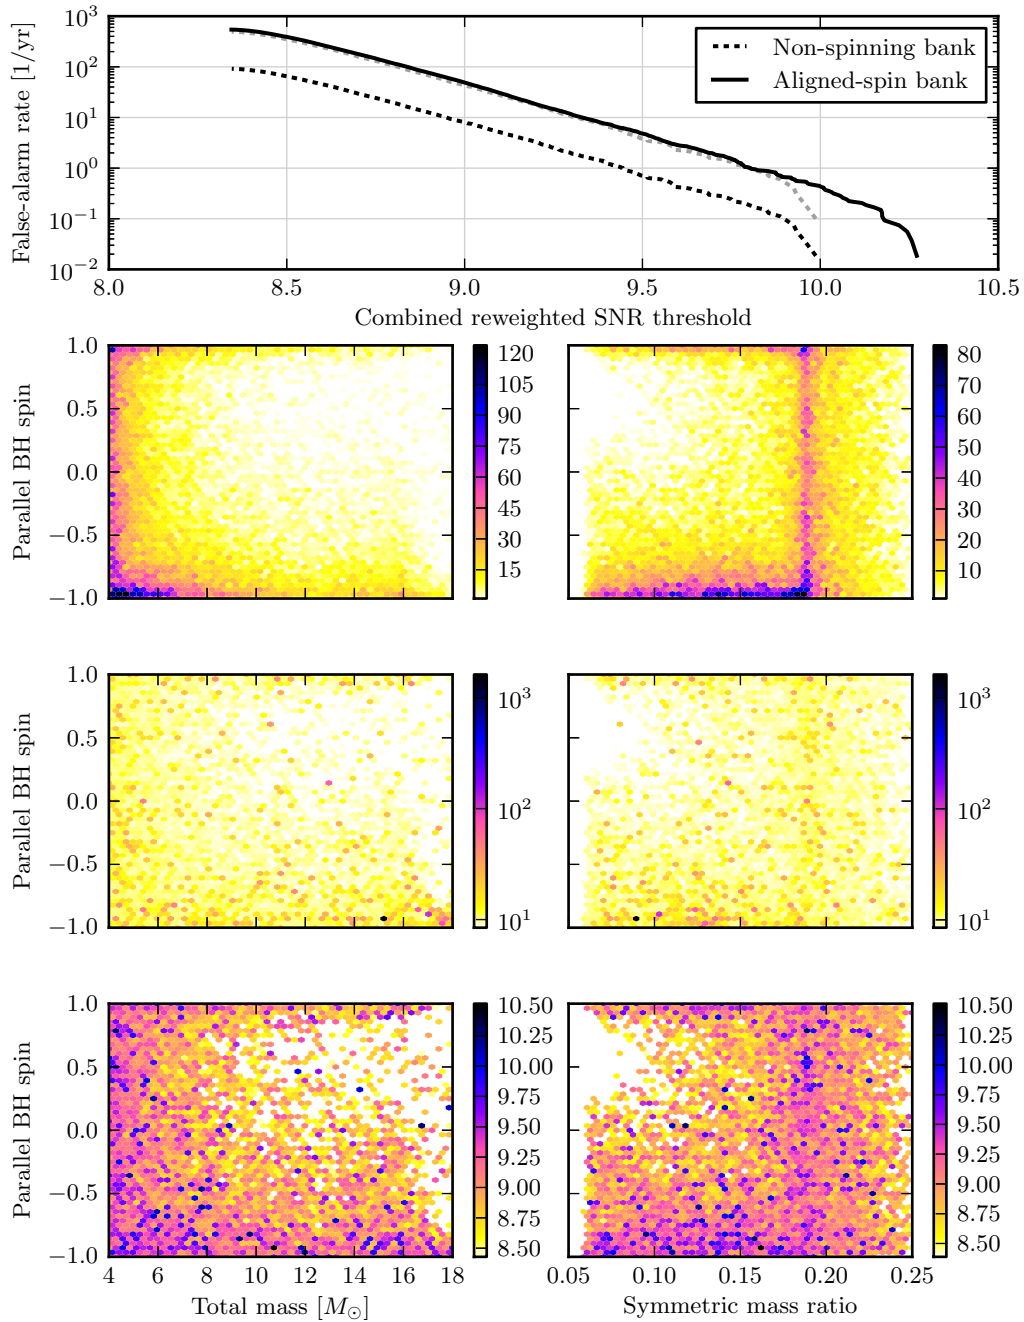


Figure 2.14: Top plot: coincident false-alarm rates for the two banks. The gray dotted curve is the non-spinning rate scaled by the relative number of templates in the ASB and NSB. Second row: density of false alarms over the ASB parameter space. Third and fourth rows: location of loud triggers (by combined SNR and reweighted SNR, in color) over the ASB parameter space. The coincident background of the ASB is distributed consistently with the metric density and falls off as expected from Gaussian noise.

signal is assigned a FAR by using the population of background coincident triggers (top plot of figure 2.14) to look up the FAR corresponding to the combined reweighted SNR of the trigger associated with the found signal.

A useful measure of the loudness or potential visibility of a signal is the optimal SNR ρ_{opt} defined in section 1.3.4. As a first check, we compare the optimal and observed single-detector SNRs signal-by-signal (figure 2.15). We can see that in the case of the NSB the observed SNR starts to drop for $|\chi_{\text{BH}}^{\parallel}| \gtrsim 0.4$, consistently with the earlier fitting-factor calculation. The ASB, instead, shows no clear dependence of the observed SNR on $\chi_{\text{BH}}^{\parallel}$ and it is able to recover a large fraction of the signal power even at high spin.

The reweighted SNR shows a different behavior due to the χ^2 veto and the discreteness of the templates. The optimal reweighted SNR is equal to ρ_{opt} , because the χ^2 for a perfectly matched signal with a vanishing noise realization is zero. However, even if a tiny mismatch exists, for a sufficiently large ρ the χ^2 becomes proportional to ρ^2 [104]. By combining this fact with eq. (1.38) we see that even a small mismatch leads to a significant reduction of $\hat{\rho}$ if ρ is sufficiently large; in fact, $\hat{\rho}$ appears to asymptote to ≈ 100 for $\rho \rightarrow \infty$. The large mismatch between highly-spinning signals and the NSB is thus exacerbated by the χ^2 veto and the resulting loss in $\hat{\rho}$ is even larger than the loss in ρ .

Signals found and missed by the NSB and ASB are shown in figure 2.16 as a function of the parallel BH spin and the *decisive* chirp effective distance, i.e. the largest chirp effective distance between the two detectors. It is evident that the NSB is unable to recover a large fraction of signals with spin, such that the critical range at which signals start to be missed drops significantly. This effect starts at $|\chi_{\text{BH}}^{\parallel}| \approx 0.4$, consistent with the fitting-factor calculation. For almost extremal positive spin, almost all signals are either missed or detected with a small significance. No drop in critical range is visible when the ASB is used. The range actually increases with $\chi_{\text{BH}}^{\parallel}$ due to the larger luminosity of highly-spinning and positively-aligned binaries.

As a check of the safety of our choice of distance distribution, we show in the left plots of figure 2.17 the detection efficiency (fraction of found signals) as a function of the chirp distance. We can see that our choice of distributing the chirp distances between 1 and 160 Mpc is safe, as the efficiency drops to almost zero at the end of the range. Figure 2.17 (right plots) shows the efficiency scaled by the chirp distance squared. This quantity represents the density of detections over the chirp distance assuming a source population uniformly distributed in volume.

2.9 Search sensitivity at fixed false-alarm rate

The final comparison of the sensitivity of the NSB and ASB requires the quantification of the number of detections made by each setup. For a population of coalescing binaries described by the number density $n(\boldsymbol{\theta})$, $\boldsymbol{\theta}$ being the full CBC parameter vector, the

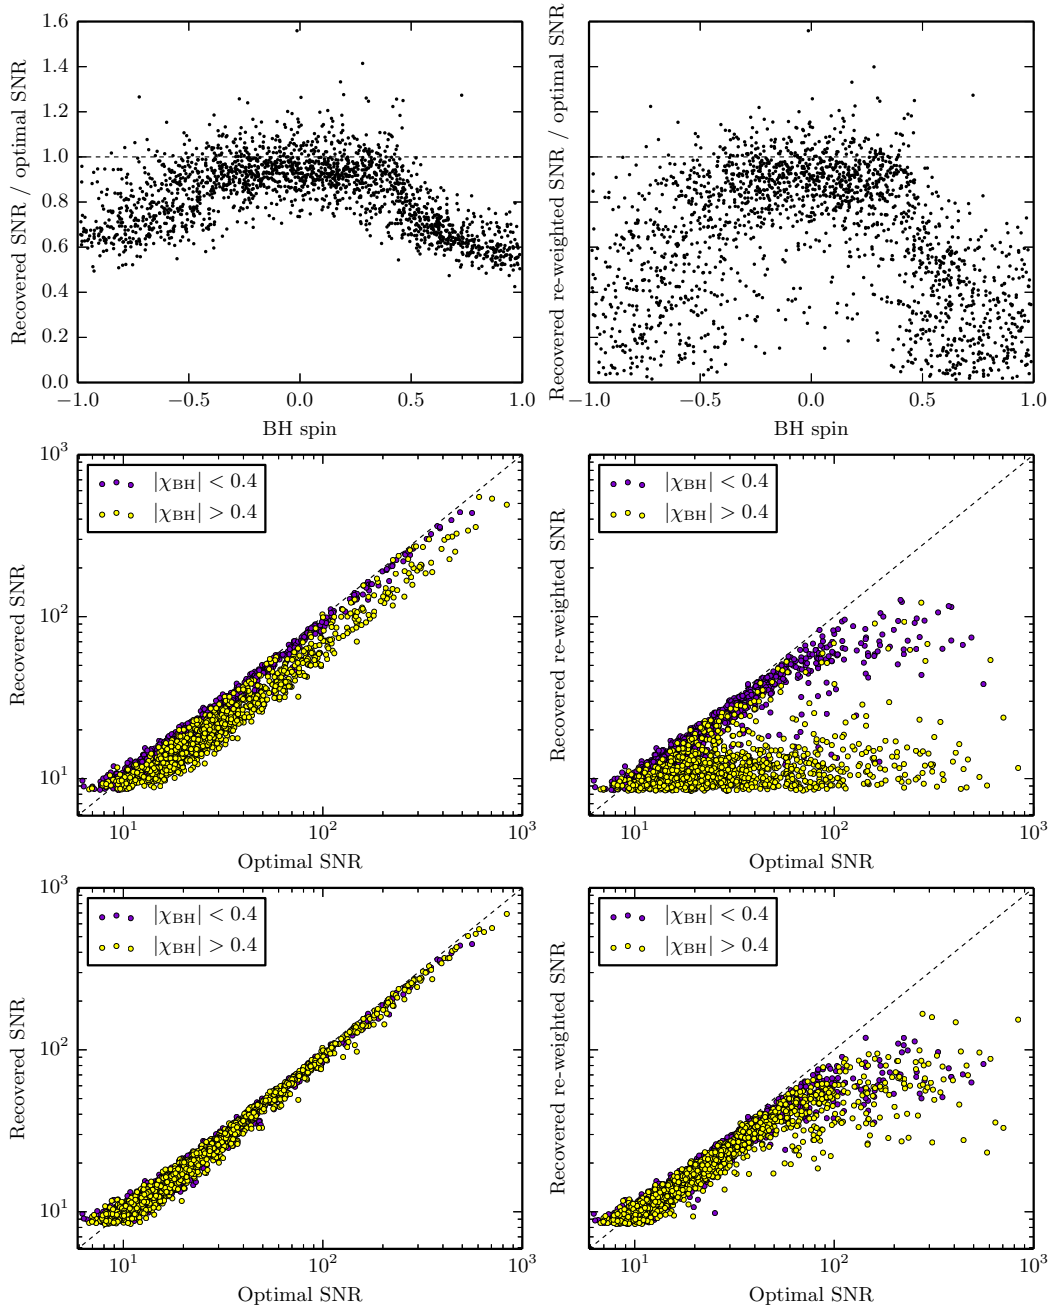


Figure 2.15: Ability of our banks to recover simulated signals in realistic noise. Top and middle: observed vs optimal combined SNR and reweighted SNR for the non-spinning bank. Bottom: same as middle, when using the aligned-spin bank. For spin magnitudes above 0.4, the non-spinning bank has a visible loss of SNR (consistent with the previous fitting-factor calculation) and an even larger loss of reweighted SNR. The aligned-spin bank, instead, performs well irrespective of the spin.

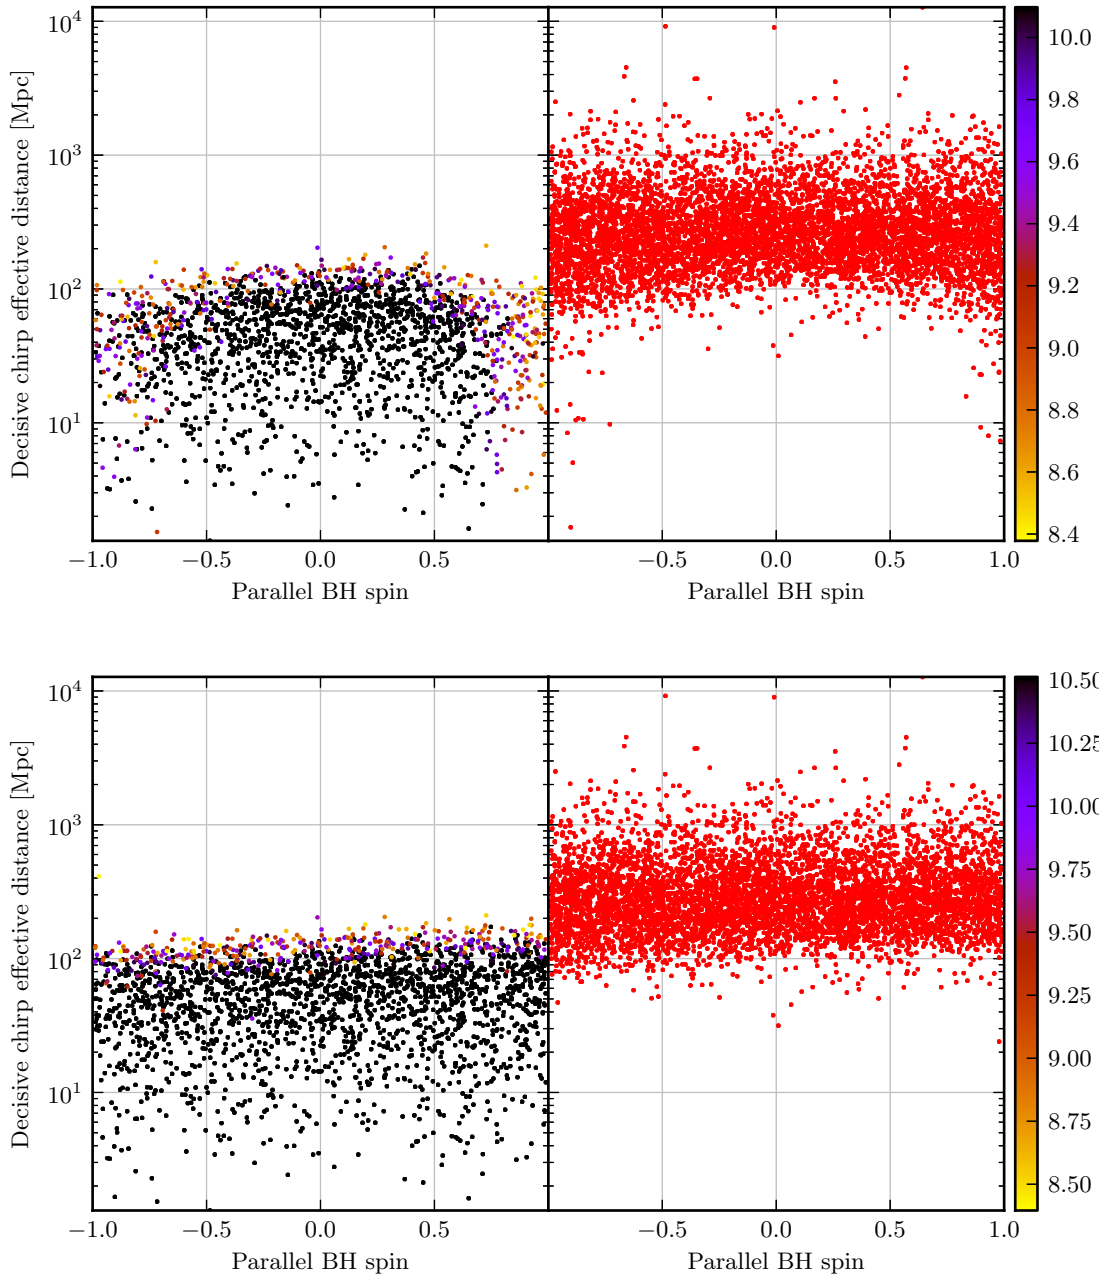


Figure 2.16: Simulated signals found (left) and missed (right) by the pipeline using the non-spinning bank (top) and aligned-spin bank (bottom). Found signals are colored with the observed combined reweighted SNR. Black found signals are louder than the loudest background trigger. With the non-spinning bank, the distance at which signals are confidently detected drops noticeably at large spin: there are only a few black points above $\chi_{\text{BH}}^{\parallel} \gtrsim 0.7$. With the aligned-spin bank, the distance is roughly constant.




Figure 2.17: Fraction of signals found as a function of the chirp distance (left) and density of detections over chirp distance (right). Top: non-spinning bank; bottom: aligned-spin bank. The lighter bands are the 68% binomial confidence intervals calculated from the number of found and missed signals in each distance bin. The simulated population correctly covers the whole range of the efficiency curve and most detections happen well below the maximum chirp distance of 160 Mpc used in the simulation.

number of signals detected with a FAR of ξ or less is

$$n^d(\xi) := \int_S p^d(\boldsymbol{\theta}, \xi) n(\boldsymbol{\theta}) d\boldsymbol{\theta} \quad (2.3)$$

where $p^d(\boldsymbol{\theta}, \xi)$ is the probability of detecting a signal with the given parameters at the given (or smaller) FAR. Our population of discrete simulated signals is effectively a Monte Carlo integration of eq. (2.3). Given that we need a large number of both found and missed signals in only two months of data, the number density we need to use is unrealistically large and the signals are distributed uniformly in chirp distance rather than volume. For this reason we use instead the following (somewhat unphysical) figure of merit,

$$V(\xi) := \frac{\sum_i p_i^d(\xi) \mathcal{R}_i^2}{\sum_i \mathcal{R}_i^2} \quad (2.4)$$

where $p_i^d(\xi)$ is 1 if the i -th signal is recovered with a FAR smaller than ξ and 0 otherwise, and \mathcal{R}_i is the chirp distance of the i -th signal (eq. (2.1)). The \mathcal{R}^2 weighting corrects our unphysical distance distribution such that $V(\xi)$ is proportional to $n^d(\xi)$ and can thus be used to compare the NSB and the ASB³. Although $V(\xi)$ represents simply the (weighted) fraction of detected signals, we sometimes refer to it improperly as the *sensitive volume* of the search pipeline, because one can define a (physical) sensitive volume from $n^d(\xi)$ in various ways. When $V(\xi)$ is plotted as a function of its argument, the curve is qualitatively equivalent to the receiver operating characteristic (ROC) curve typically used to characterize the performance of a signal detector [165]. Therefore, according to the Neyman-Pearson criterion, we need to look for the bank which produces the largest $V(\xi)$ at a fixed (sufficiently low) value of ξ .

This comparison is shown in figure 2.18 for different cuts on $|\chi_{\text{BH}}^{\parallel}|$. We compute 68% uncertainty intervals on $V(\xi)$ by constructing 100 realizations of this quantity using random selections of half of the background and half of the simulated signals, and taking the standard deviation of the result. We see that the relative improvement of the ASB depends strongly on the distribution of $\chi_{\text{BH}}^{\parallel}$. Assuming BHs can have any value of $\chi_{\text{BH}}^{\parallel}$ allowed by the Kerr limit, the ASB increases the sensitivity by 40 – 60% depending on the FAR. If we make the restriction $\chi_{\text{BH}}^{\parallel} > 0$, the improvement is slightly larger. If however we assume $\chi_{\text{BH}}^{\parallel} > 0.7$, the improvement can be a factor of $\mathcal{O}(10)$. This can be understood by considering the dramatic loss in reweighted SNR of the non-spinning search, which is due in part to the loss of SNR and in part to the poor χ^2 value of highly-spinning signals in the NSB. If the data were so well-behaved that the search could be carried out using just the combined SNR as the ranking statistic, therefore, it is possible that the improvement would be less dramatic. The same could happen with a different tuning of the χ^2 test, for instance with a different number of bins. Nevertheless, the bank simulations (figure 2.4 top-left and 2.7 left) suggest that in this extreme cases the improvement would still be a factor of $0.6^{-3} \approx 4.6$, ignoring the increase in false-alarm background. It is also worth noting that the improvement could be smaller if precession

³Another way to estimate the sensitivity is measuring the area under the curves in figure 2.17 (right).




Figure 2.18: Sensitivity of the pipeline with the non-spinning and aligned-spin banks at fixed FAR. The four plots assume NSBH systems with different limits on uniformly distributed BH spins. We can see that the aligned-spin bank leads to a more sensitive pipeline in almost all cases, with a large gain if most systems have nearly extremal-spin BHs. For weakly-spinning systems, the reduction in sensitivity is modest. This figure is the main result of the chapter.

is included in the simulated signals. In fact, this will be studied in the next chapter. For the case of weakly-spinning BHs, the ASB gives a slightly less-sensitive search, as can be expected from the larger background. Finally, although our statistics are insufficient for a ROC plot, we can do a back-of-the-envelope estimate of the relative sensitivity in the case $\chi_{\text{BH}}^{\parallel} = 0$. Using the background curve in figure 2.14 and a detection FAR of 1 yr^{-1} , corresponding to a threshold in combined reweighted SNR of $\bar{\rho} \approx 9.5$ for the non-spinning search, we can write

$$V(\xi) \approx \left(\frac{\bar{\rho}}{\bar{\rho} + 0.3} \right)^3 > 90\%. \quad (2.5)$$

In the unlikely worst case of exactly non-spinning BHs, therefore, the aligned-spin search would greatly increase the computational cost but only lose at most 10% of the detections.

2.10 Conclusion

In this study we demonstrate that the currently-available technology, waveform templates, signal-based vetoes and software implementation are sufficient to include aligned-spin effects in pipelines searching for NSBH coalescence in data from advanced GW detectors. We find that neglecting spin can lead to a large reduction in the number of detections if BHs have spins aligned with the orbital angular momentum and with magnitudes close to the Kerr limit. The sensitivity to such signals is restored by simply including the appropriate spin terms in the pN phasing of TaylorF2 templates and constructing a template bank which appropriately covers the new dimension of the search parameter space. If the universe is populated with a variety of spin magnitudes, neglecting spin would introduce an unnecessary selection bias. Even if BHs turn out *not* to spin significantly—a scenario in contrast with some of the existing observations and population-synthesis models—including spins leads to a small reduction in the sensitivity of the search.

The price to pay for taking spin into account is however a significant $\approx 5\times$ increase in computational cost of the search, which can nevertheless be met by modern implementations of the inspiral matched filter, such as PyCBC. Note that if we decided to neglect spin and raise instead the minimum match of the NSB from 97% to whatever is required to get a $5\times$ increase in number of templates, the sensitivity achieved to highly-spinning signals would still be much lower than what can be obtained by including spin. We argue instead that spin should be included and, should computing power be a limit, the minimum match of the ASB should be lowered appropriately in order to meet the available computing resources. More recent optimizations of PyCBC, however, are pushing the cost down and such a reduction in minimum match should not be necessary. A straightforward reduction in the computational cost of the pipeline can be achieved using the more efficient geometric placement method for constructing the ASB, which is possible with PyCBC.

Most templates in the ASB are located at $\chi_{\text{BH}}^{\parallel} < 0$, i.e. anti-aligned spins. We find that some of those templates are the most sensitive to instrumental glitches in the detector noise, although the $\hat{\rho}$ statistic is effective at separating glitches from astrophysical signals even in that region of the search space. Moreover, as will be discussed more in the next chapter, population synthesis models suggest that compact binaries with anti-aligned spins are the least likely. For these reasons, it could be argued that this part of the parameter space is the least worthwhile to include in the search. Excising anti-aligned templates would save a factor ≈ 2 in computational cost and make the search even more robust against instrumental artifacts in the data.

This study should be extended by considering the sensitivity of the search pipeline to generic precessing NSBH systems; this is the subject of the next chapter. Merger and ringdown effects, or tidal disruption of the NS, could be important in some parts of the parameter space. They should be included in the simulated signals and their effect should be investigated as well. Including merger and ringdown in the templates could enable the extension of the search space to higher BH masses. Further investigation, currently ongoing, is needed to tune the various parameters of the aligned-spin pipeline in order to maximize its sensitivity. It is also now possible to include higher-order spin terms into the phasing of our templates and signals, although it is unlikely that this would significantly increase the search sensitivity.

Chapter 3

Effect of neglecting precession in NSBH search pipelines

3.1 Overview

The previous chapter presented a search pipeline for NSBH binary coalescence which takes into account the effect of the BH spin component along the orbital angular momentum \mathbf{L} . It was shown that this can lead to a significant improvement if the spin magnitude is close to the Kerr limit.

BH spins and orbital angular momenta in astrophysical NSBH binaries, however, are most likely not exactly aligned. The standard evolutionary scenario for binaries (section 1.2.6) suggests that the common-envelope phase tends to align the angular momenta, but the kick associated with the second supernova explosion can then tilt the orbital plane, assuming the binary survives the event [166]. Nevertheless, kicks seem unlikely to be able to produce very large tilt angles. Simulations show that most coalescing NSBH binaries probably have tilts below 60 degrees with a significant fraction below 45 degrees [148–150]. Systems with anti-aligned spins are expected to be relatively rare. The physical interpretation of such results is that a large tilt requires a large kick component perpendicular to the initial orbital plane. This in turn implies a large kick magnitude, which makes the binary less likely to survive the kick. The actual distribution of tilt angles in nature is however currently unknown. Observations of X-ray binaries have so far been unable to provide clear measurements of the tilt. In fact, even measuring the spin magnitude often requires the assumption that the axis of the accretion disk is aligned with the BH spin [22].

The most important consequence of a tilted spin on the evolution of a NSBH binary is the precession of the orbital angular momentum, already introduced in chapter 1, which is analogous to the Lense-Thirring gravitomagnetic effect on a vector moving in a Kerr metric. Most systems radiating in the sensitive band of ground-based interferometers exhibit *simple precession*: the orientation of the total angular momentum \mathbf{J} remains constant to a good approximation while \mathbf{L} traces out a cone around it [57]. For NSBH binaries with massive and highly-spinning BHs, where $\mathbf{J} \approx \mathbf{S}_{\text{BH}}$, the opening angle of

the precession cone can be quite large. A distant inertial observer then sees a steady rotation of the orbital plane. The orientation of the orbital plane affects both the polarization amplitudes of the emitted radiation (via the inclination angle in eq. (1.9)) and the projection of the GW tensor on the detector (via the polarization angle and the beam pattern coefficients in eq. (1.20)). The result is that the precessing \mathbf{L} imprints a characteristic amplitude and phase modulation on the gravitational waveform measured by the detector. The time scale of one precession cycle is typically longer than the orbital period and this allows one to compute the modulation by averaging over several orbits [57]. Whereas the signal from a non-precessing circular orbit is dominated by a single chirp in the time-frequency plane, the modulation results in five chirps closely spaced in time and frequency, one corresponding to the carrier (at twice the orbital frequency) and four representing two symmetric modulation sidebands. The modulation is currently implemented by several waveform approximants in the LALSimulation library, notably SpinTaylorT2 and recently SPA and IMR approximants [59, 77, 78].

Existing studies attempted to include the effect of precession in a CBC search pipeline, but this did not result in any significant gain in sensitivity [114, 136, 144, 167–169]. This was mainly due to (i) the ability of precessing templates to pick up instrumental glitches and (ii) the lack of signal-based vetoes which could efficiently distinguish such spurious triggers from astrophysical signals. At the moment, it is not known how to include precession in a pipeline that can be practically run and achieve a satisfactory sensitivity. Given the scarce information available on spin tilts, it is therefore important to assess the effect of precession on the sensitivity of NSBH pipelines that neglect it, in particular on the non-spinning and aligned-spin pipelines presented in chapter 2. A first step towards this goal has been recently achieved by means of fitting-factor calculations and assuming the design sensitivity of the advanced LIGO detectors [145, 146].

In this chapter we extend the investigation carried out in chapter 2, and in previous studies, by presenting the effectualness of a non-spinning bank (NSB) and an aligned-spin bank (ASB) and the sensitivity of realistic non-spinning and aligned-spin search pipelines in the presence of precessing NSBH binaries. As done in chapter 2, we assume here the early- advanced-LIGO sensitivity curve. We first use bank simulations to compare the fitting factors of precessing NSBH signals in the NSB and the ASB and we study the variation of the fitting factor with the spin parameters (section 3.2). Then, we apply the same search pipelines described in chapter 2 to a population of such precessing signals and we show their sensitivity for various subsets of the spin parameter space, corresponding to different astrophysical expectations (sections 3.3 and 3.4). Finally, we use both the fitting factors and the results of the pipelines to estimate the loss of detections when aligned-spin effects are included in the pipeline, but precession is neglected (section 3.5). In section 3.6 we draw the conclusion from this study.

Results of this project are published in [170]. This chapter is a complete rewrite of the publication, but its contents do not differ significantly from the publication, except for an additional plot. Alex Nielsen and Andrew Lundgren gave important ideas, advice and comments on the project and the paper. Alex Nielsen in particular ran some of the fitting- factor calculations and wrote some sections of the article. Andrew Lundgren

wrote some of the code for handling precessing signals. Badri Krishnan, Tom Dent and Steve Privitera gave useful comments on the project and the paper. The author ran some of the fitting-factor calculations, ran the search pipeline, developed the estimations in section 3.5, wrote a major part of the paper and handled the publication process and communication with the anonymous referee.

3.2 Impact of precession on visibility of signals

In this section we perform bank simulations to study the effectualness of the NSB and ASB already used in chapter 2 and described in section 2.5, relaxing the assumption that our simulated NSBH signals have BH spins aligned with the orbital angular momentum. Enabling precession complicates the problem by introducing new degrees of freedom: the fitting factor of a signal can now depend on the tilt angle of the BH spin as well as the orientation of the total angular momentum \mathbf{J} with respect to the detector.

In order to simplify the problem and separate out the effect of the new parameters from the mass degrees of freedom, we begin with a population of NSBH binaries with masses fixed at $m_{\text{BH}} = 7.8M_{\odot}$, $m_{\text{NS}} = 1.35M_{\odot}$. The dimensionless BH spin magnitude χ_{BH} is distributed uniformly between 0 and 1 and the tilt angle ϑ_{BH} is determined by a uniform distribution of $\cos \vartheta_{\text{BH}}$ between ± 1 ¹. The orientation of \mathbf{J} is uniform on the unit sphere. The population is simulated using the SpinTaylorT2 approximant with the same settings are done in chapter 2.

Fitting factors for this first population against our banks are shown in figure 3.1 as a function of the main parameters of the dimensionless BH spin, namely its projection on the orbital angular momentum $\chi_{\text{BH}}^{\parallel}$ and the magnitude of its projection on the orbital plane

$$\chi_{\text{BH}}^{\perp} := \|\chi_{\text{BH}} - (\chi_{\text{BH}} \cdot \hat{\mathbf{L}})\hat{\mathbf{L}}\|. \quad (3.1)$$

By observing the fitting factors in the two banks, we can approximately separate the spin parameter spaces into three different regions.

The first region is defined by low values of the BH spin magnitude, roughly $\chi_{\text{BH}} \lesssim 0.4$. Not surprisingly, both banks work well here. The aligned-spin-orbit pN terms in the signal waveforms are small enough that the dephasing with respect to non-spinning templates can be compensated by a small bias in the symmetric mass ratio, allowing the NSB to be effectual, as we saw in section 2.5. Moreover, even if the spin is tilted, its magnitude is too small for precession to significantly modulate the signal.

The second region is defined by a small projection of the spin on the orbital plane but a large negative or positive spin component along \mathbf{L} , i.e. $\chi_{\text{BH}}^{\perp} \lesssim 0.5$ and either $\chi_{\text{BH}}^{\parallel} \gtrsim 0.4$ or $\chi_{\text{BH}}^{\parallel} \lesssim -0.5$. We can observe here a severe loss of the effectualness of the NSB, while the ASB performs significantly better. This happens because the pN spin-orbit terms

¹The resulting distribution of the spin vector is not uniform in the unit ball: sources are more likely to have a spin vector close to the center of the ball. This is appropriate if we assume that the magnitude and the orientation of the spin are determined independently by different physical processes (for instance, accretion on the BH and kick associated with the second supernova).

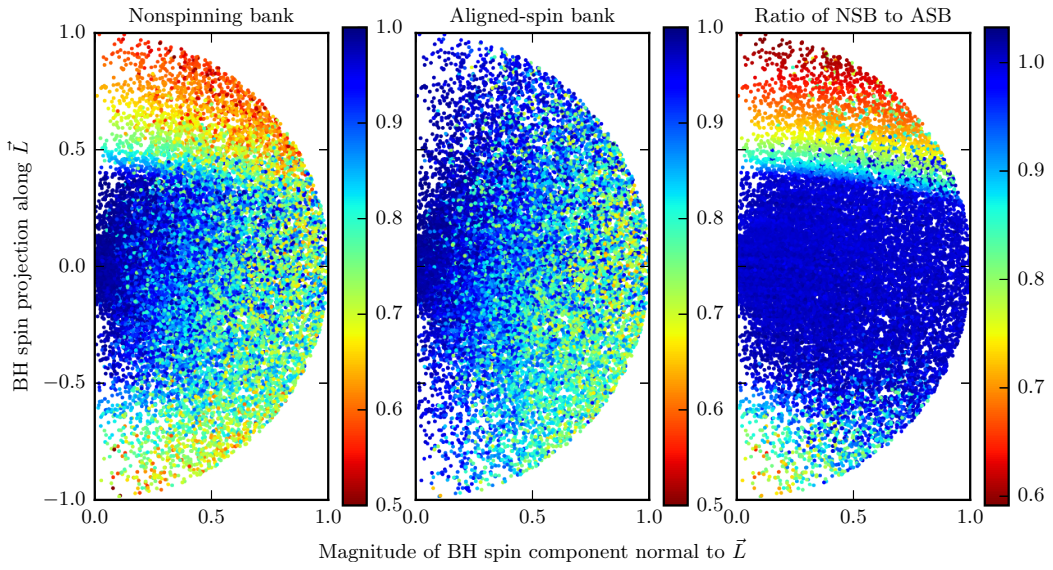


Figure 3.1: Fitting factors of our NSB, ASB and their ratio for 2×10^4 precessing NSBH binaries with $m_{\text{BH}} = 7.8M_{\odot}$, $m_{\text{NS}} = 1.35M_{\odot}$. The origin corresponds to sources with non-spinning BHs and the vertical axes correspond to the aligned-spin case. The NSB has a severe loss of effectualness when the spin projection on \vec{L} is large; this issue disappears with the ASB. However, neither bank is effectual for many systems having large orthogonal spin components.

in the waveform phase are large, causing a large dephasing with respect to non-spinning templates. Similarly to the effect found in chapter 2, this causes a bias in η so large that non-spinning templates with $\eta > 1/4$ would be required to recover positively aligned signals, and non-spinning templates with $m_{\text{NS}} < 1M_{\odot}$ would be required for negatively-aligned signals. However, the modulation induced by precession is still small in this region, so the ASB is effectual.

The third region represents the highly-precessing case and is roughly identified by $\chi_{\text{BH}}^{\perp} \gtrsim 0.5$. Two main features can be observed here: (i) the banks produce very similar fitting factors and (ii) there is a much larger variance in the fitting factor with respect to the other two regions, with some sources recovered well and others recovered poorly.

The similar behavior of the banks in the third region can be explained by the fact that pN spin-orbit terms in the signal phase are small here, so the NSB is able to compensate their effect with a small bias in η ; at the same time, the modulation induced by precession is the largest possible and neither bank is able to recover the power going into the modulation sidebands. This demonstrates how the effects of the aligned and orthogonal components of the spin are almost decoupled, as predicted for example in [171].

The spread in the fitting factor visible in the high-precession region is explained instead by the different possible orientations of \mathbf{J} relative to the detector, as demonstrated in figure 3.2. In fact, the orientation determines the fraction of signal power that goes into the modulation sidebands and that is lost when non-precessing templates are used. In particular, precessing binaries with \mathbf{J} pointed at or away from the detector (“parallel- \mathbf{J} ” systems) tend to look more like non-precessing systems than orthogonal- \mathbf{J} ones, so their waveforms have a smaller modulation and the fitting factor is larger [172]. We can see from figure 3.2 that large fitting factors are also possible for orthogonal- \mathbf{J} binaries with four particular orientations of \mathbf{J} . However, such systems produce relatively quiet signals at the detector, i.e. the maximum range at which they can be detected is lower than parallel- \mathbf{J} ones. As shown in the bottom plot, in fact, orthogonal- \mathbf{J} systems are generally quieter than parallel- \mathbf{J} ones at the same distance, because one polarization contains less signal power than the other. Thus, although detecting orthogonal- \mathbf{J} precessing systems with non-precessing templates is challenging, they are also the least likely to be detected even with an ideal precessing bank.

Having understood the main features of precessing signals for the fixed-mass case, we now extend the simulation to the whole range of masses considered in the construction of the ASB, i.e. $m_{\text{BH}} \in [3, 15]M_{\odot}$ and $m_{\text{NS}} \in [1, 3]M_{\odot}$ with uniform distributions. The result is shown in figure 3.3 and it is qualitatively consistent with the fixed-mass case, although the separation between the three regions is fuzzier. A visible difference is the much worse fitting factor between the NSB and some of the anti-aligned systems (lower region of first and last panels); these are heavy systems, with $m_{\text{BH}} \gtrsim 11M_{\odot}$. The poor performance of the NSB for those signals is not due to precession but again to the bias in η produced by attempting to recover spinning signals with non-spinning templates. Being anti-aligned, these systems introduce a negative bias, i.e. they would be well recovered by non-spinning templates associated with a lighter NS and a heavier BH.




Figure 3.2: Top plot: fitting factor of strongly-precessing NSBH binaries with the aligned-spin bank as a function of the orientation of the total angular momentum \mathbf{J} with respect to the detector (poles: \mathbf{J} parallel to the line of sight; equator: \mathbf{J} orthogonal to the line of sight). Almost identical fitting factors are obtained with the non-spinning bank. Bottom plot: optimal SNR of the same binaries at a fixed distance. Sources have $m_{\text{BH}} = 7.8M_{\odot}$, $m_{\text{NS}} = 1.35M_{\odot}$, $\chi_{\text{BH}}^{\perp} > 0.5$ and $|\chi_{\text{BH}}^{\parallel}| < 0.2$. We can see that the loudest systems (located close to the poles) tend to have fitting factors close to 100%, while systems with poor fitting factors are relatively quiet.




Figure 3.3: Fitting factors of our NSB, ASB and their ratio for 2×10^4 precessing NSBH binaries with $m_{\text{BH}} \in [3, 15]M_{\odot}$, $m_{\text{NS}} \in [1, 3]M_{\odot}$. The origin corresponds to sources with non-spinning BHs and the vertical axes correspond to the aligned-spin case. The result is very similar to the fixed-mass case shown in figure 3.1.

For systems with small enough masses (such as the earlier fixed-mass case) the biased masses are still within the mass range of the NSB and the bank is effectual. Signals with negatively-spinning BHs heavier than $\approx 11M_{\odot}$, however, “fall off” the $m_{\text{BH}} = 15M_{\odot}$ boundary of the bank; templates with larger m_{BH} would be required for the bank to be effectual.

3.3 Recovering precessing signals in realistic data

The sensitivity of a CBC search pipeline is determined primarily by the background of false alarms associated with the detector noise and by the effectualness of the template bank for the targeted signal population. We have seen in chapter 2 that the background increases in a predictable way when going from the NSB to the ASB and we have just investigated the effectualness of the banks for precessing binaries. Based on these results we expect the ASB to still perform at least as well as the NSB and to significantly outperform it for almost-aligned, high BH spins. However, the correct behavior of coincidence and signal-based vetoes remains to be verified when precession is important. Ultimately we also want to plot the sensitivity of the pipeline at fixed FAR (as done in figure 2.18) for different distributions of spin magnitude and tilt.

As the next step we thus add signals from a population of precessing NSBH binaries to the realistic simulated data described and used in chapter 2. We use again a fixed-mass

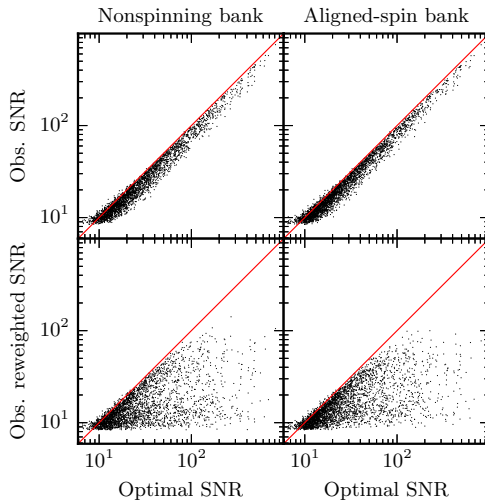


Figure 3.4: Combined SNR and reweighted SNR observed by the search pipelines for each simulated signal vs the optimal combined SNR of the signal.

population ($m_{\text{BH}} = 7.8M_{\odot}$, $m_{\text{NS}} = 1.35M_{\odot}$) and one where the masses are distributed as in chapter 2, namely: m_{BH} is a Gaussian with mean $7.8M_{\odot}$ and standard deviation $3M_{\odot}$, truncated between $3M_{\odot}$ and $12M_{\odot}$; m_{NS} is a Gaussian with mean $1.35M_{\odot}$, standard deviation $0.13M_{\odot}$ and truncated between $1M_{\odot}$ and $2M_{\odot}$. The BH spin is distributed as in section 3.2. All the other parameters, the waveform approximant and the configuration and background of the search pipelines are identical to what was described in chapter 2.

A signal-by-signal comparison between optimal combined SNR and observed combined SNR is shown in figure 3.4 for the fixed-mass population; results for the varying-mass signals are similar and are not shown. For many signals, both banks show a large loss of observed SNR and an even larger suppression of observed reweighted SNR. This is likely due to the χ^2 veto penalizing precessing signals. The advantage of the ASB is not clear from figure 3.4 because of the relatively scarce systems with large and aligned spins. The variation of the loss with the spin parameters is more clear in figure 3.5. Both the SNR and reweighted SNR show a consistent behavior with the fitting factor calculation: the ASB is much more effectual for large spin magnitude and tilt angle close to either 0 or 180 degrees, while both banks have similar performance for small magnitude or tilt angle around 90 degrees. We can thus be confident that the exact-match coincidence method and the χ^2 veto preserve the features of the banks observed with the bank simulations.

3.4 Search sensitivity at fixed false-alarm rate

We now present the sensitivity of the search pipelines with the NSB and the ASB at fixed FAR and for different cuts on the BH spin parameters. The figure of merit is the

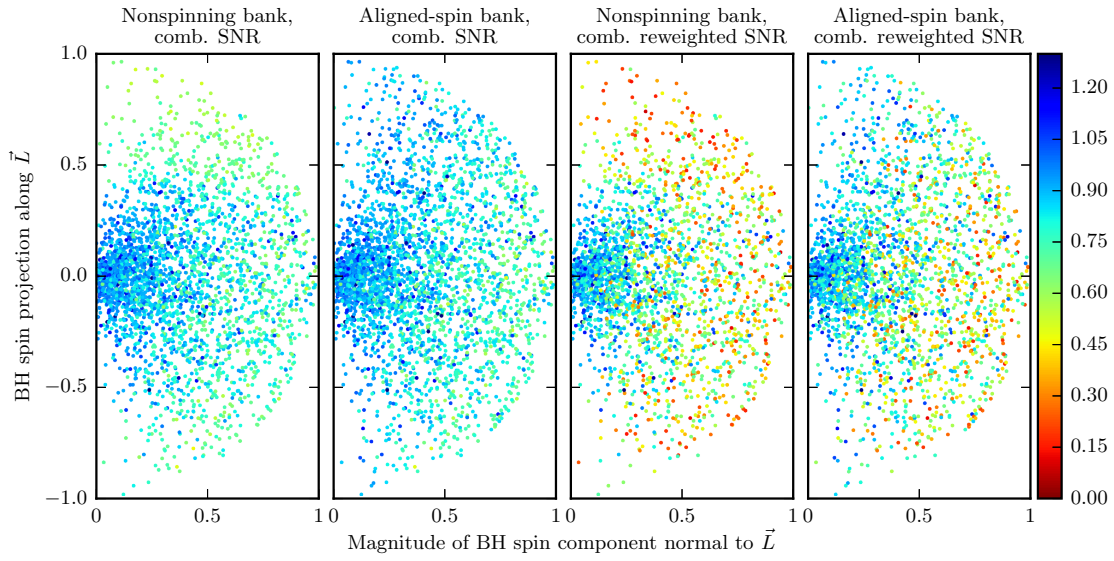


Figure 3.5: Fraction of optimal combined (reweighted) SNR observed by the search pipelines for each simulated signal as a function of the BH spin parameters. In the reweighted SNR plots, only systems with observed combined SNR smaller than 50 are shown, because louder systems are expected to have a relatively small observed combined reweighted SNR due to the discreteness of the bank.

same used in chapter 2, i.e. eq. (2.4); plotting this as a function of the FAR produces ROC curves. We compute 68% uncertainty intervals on this quantity by constructing 100 different curves using random selections of half of the background and half of the simulated signals, and taking the standard deviation of the result.

ROC curves associated with the fixed-mass population are shown in figure 3.6. The first plot shows all systems, regardless of their spin parameters: the ASB gives a slightly larger sensitivity, although the difference is compatible with the error interval. The second plot is the case of weak spin magnitudes and unrestricted tilts, for which the NSB is close to optimal. As we saw already in chapter 2, the ASB has a slightly lower sensitivity here due to the larger background. Plots 3, 4 and 5 assume respectively large spin magnitudes, large magnitudes and small tilt angles, and small tilt angles only. In these cases the advantage of the ASB is much clearer, giving a sensitivity between 50% and one order of magnitude larger. According to existing population synthesis studies and spin measurements, the most realistic scenario is between plots 4 and 5. From plot 6 we can see that, for the high-precession case, the banks perform again in a similar way. The ASB gives the worst sensitivity in this region. ROCs for the varying-mass population are shown in figure 3.7. The qualitative results are unchanged and we can thus extend our conclusions to a realistic mass distribution.

3.5 Comparison with an idealized precessing search

Although precessing waveform approximants are already available, further development is required to include precession in a search pipeline. The most important missing component is a full precessing bank (PB). Although this development is outside the scope of this thesis, as a final exercise we use the previous results to estimate the loss of detections produced by using the best search pipeline currently available, i.e. the one based on the ASB, and neglecting precession effects. This should help prioritizing the development of a fully precessing pipeline.

One way to estimate the detection rate of the ASB relative to an ideal PB is to use the fitting factors and optimal SNRs from our bank simulations. Consider one of our simulated systems. The spatial volume within which an identical system at arbitrary distance produces an SNR of at least ρ^{ref} in the ASB scales as

$$V_{\text{ASB}}^i \propto \left(\frac{\varphi_i \rho_i}{\rho^{\text{ref}}} \right)^3 \quad (3.2)$$

where i labels the system, φ_i is the fitting factor with the ASB and ρ_i is the optimal SNR at unit distance. The volume corresponding to the ideal PB can be estimated in the same way by simply assuming an average fitting factor $\varphi_{\text{PB}} = 0.985$ irrespective of the parameters of the signal. By summing over all the simulated signals and taking the ratio, we obtain the relative sensitivity

$$\mathcal{V} := \frac{V_{\text{ASB}}}{V_{\text{PB}}} = \frac{\sum_i (\varphi_i \rho_i)^3}{\varphi_{\text{PB}}^3 \sum_i \rho_i^3}. \quad (3.3)$$

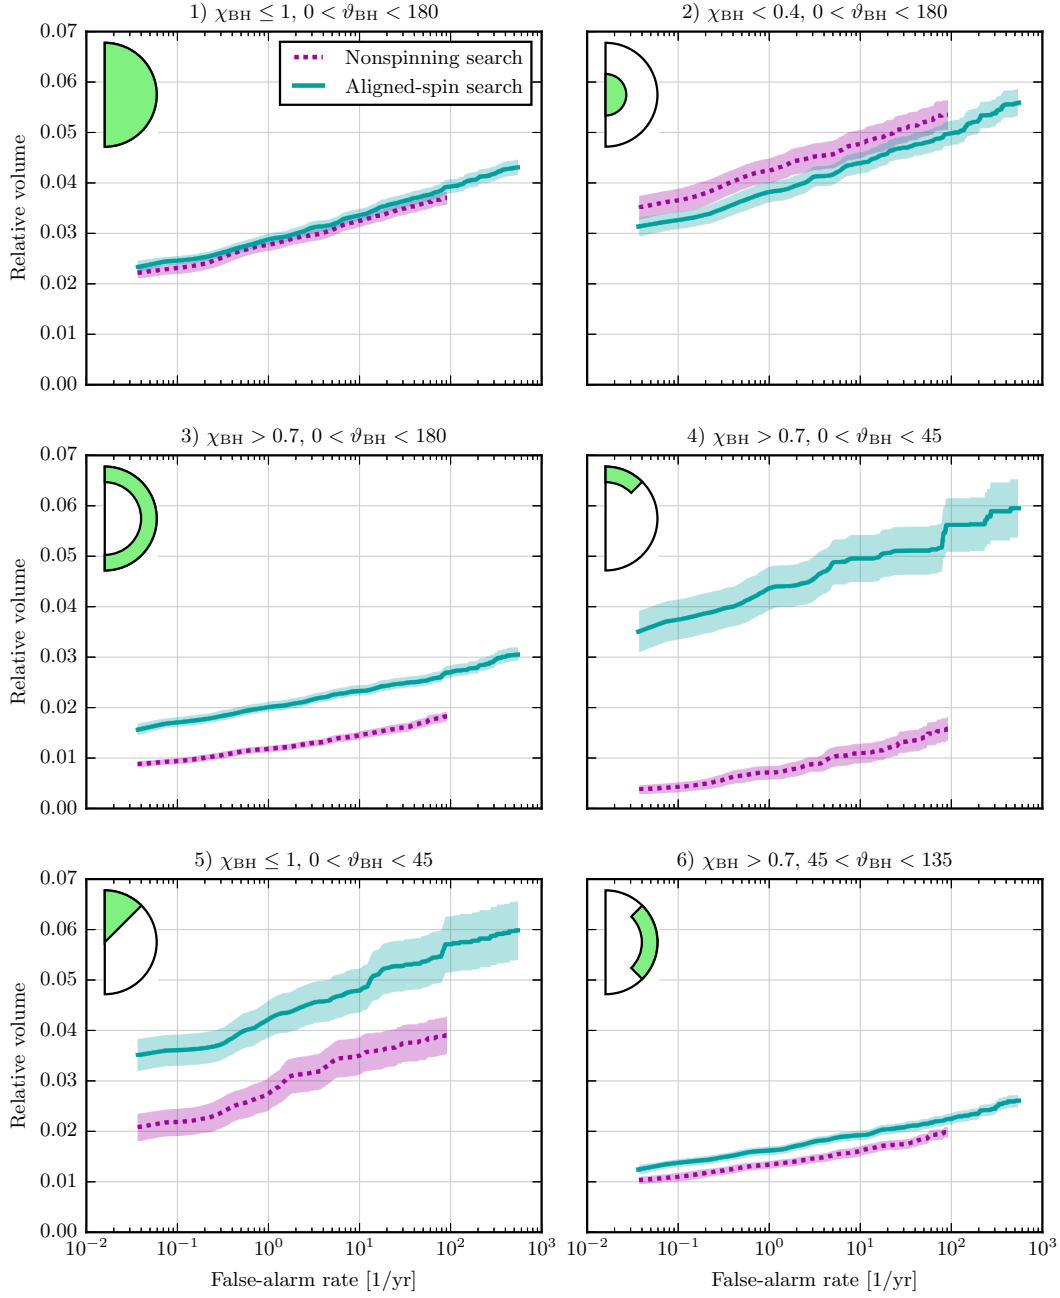


Figure 3.6: ROCs associated with the non-spinning and aligned-spin pipelines for NSBH binaries with $m_{\text{BH}} = 7.8M_{\odot}$, $m_{\text{NS}} = 1.35M_{\odot}$ and different constraints on the BH spin parameters (green areas in the insets). The aligned-spin pipeline is more sensitive in most cases, with large gains for astrophysically plausible spin distributions.

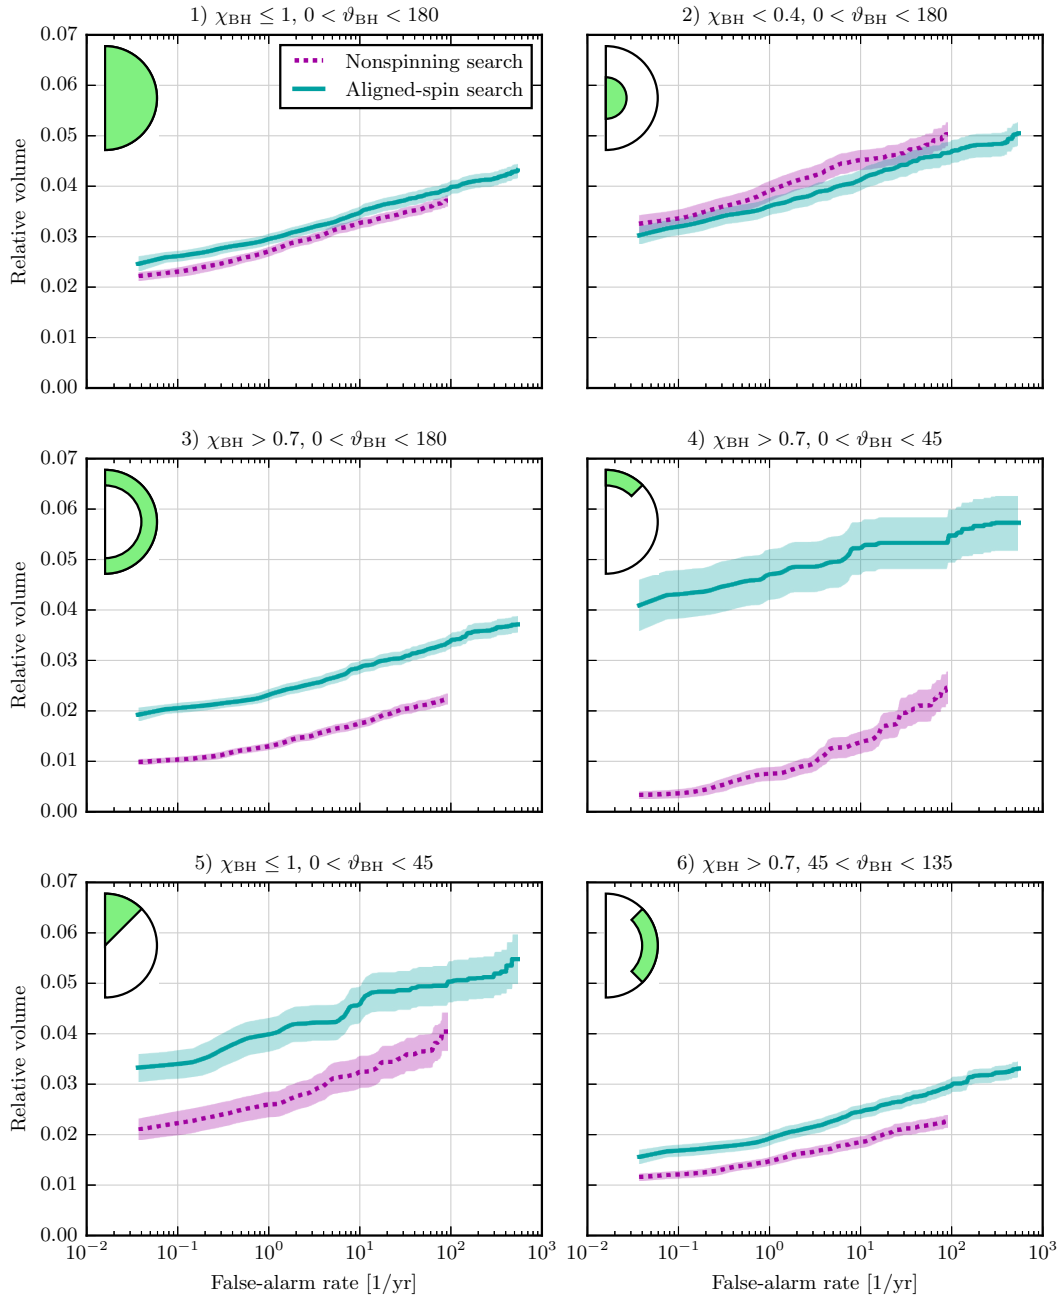


Figure 3.7: ROCs associated with the non-spinning and aligned-spin pipelines for NSBH binaries with varying masses (see text) and different constraints on the BH spin parameters (green areas in the insets). Even with a realistic mass distribution, the aligned-spin pipeline is more sensitive for most spin distributions, with large gains for astrophysically plausible ones.

Note that we are ignoring the increased background of the precessing pipeline—which could be considerable—so \mathcal{V} is a lower limit. Coincidence and signal-based vetoes are also obviously ignored. With this caveat in mind, using the fixed-mass (“FM”) or varying-mass (“VM”) bank simulations we obtain

$$\mathcal{V}_{\text{FM}}^{\text{all}} \approx 81\% \quad (3.4)$$

$$\mathcal{V}_{\text{FM}}^{\text{HP}} \approx 61\% \quad (3.5)$$

$$\mathcal{V}_{\text{VM}}^{\text{all}} \approx 86\% \quad (3.6)$$

$$\mathcal{V}_{\text{VM}}^{\text{HP}} \approx 71\% \quad (3.7)$$

where the superscripts “all” and “HP” indicate respectively the case of no restriction on the spin parameters and the high-precession case $\chi_{\text{BH}} > 0.7$, $45 < \vartheta_{\text{BH}} < 135$. It appears from this estimation that the ASB is able to detect a large number of systems, even when precession is important. This can be explained by considering that the brightest systems are parallel- \mathbf{J} ones and the ASB is quite effectual at recovering them, as shown in figure 3.2.

The results of running the search pipeline allow us to estimate the loss of the ASB vs PB in a different way, if we make one more approximation (other than neglecting the larger background of the PB). This is that the number of detections with the PB, averaged over the possible orientations of \mathbf{J} , does not depend on χ_{BH}^{\perp} . In other words, we assume that an ensemble of precessing binaries with $\hat{\mathbf{J}}$ uniformly distributed on the sphere is not brighter than a similar ensemble of non-precessing binaries. This is justified because precession merely distributes the radiated power more evenly across the sky—making orthogonal- \mathbf{J} systems more detectable and parallel- \mathbf{J} ones less detectable—but it does not make the signals longer nor does it increase their radiated power² [172]. Under these assumptions, the fraction of generic binaries detected by the precessing pipeline should be equal to the fraction of weakly-precessing binaries detected by the aligned-spin pipeline. Since the fraction of detected sources is our definition of ROC, we can then use the ROC of the ASB associated with weakly-precessing sources (say, $\chi_{\text{BH}}^{\perp} < 0.4$) as a proxy for the ROC of the precessing pipeline with no restrictions on the spin parameters. We can thus estimate the relative sensitivity of the ASB as

$$\mathcal{W}_{\text{FM,VM}}^{\text{all,HP}}(\xi) := \frac{W_{\text{FM,VM}}^{\text{all,FM}}(\xi)}{W_{\text{FM,VM}}^{\text{LP}}(\xi)} \quad (3.8)$$

where $W_M^S(\xi)$ is the ROC associated with the aligned-spin pipeline assuming the particular cut S of the spin parameters (“all”: no restriction; “LP”: low precession, $\chi_{\text{BH}}^{\perp} < 0.4$; “HP”: high precession, $\chi_{\text{BH}} > 0.7$ and $45 < \vartheta_{\text{BH}} < 135$) and mass distribution M (“FM”:

²Since precession produces modulation sidebands around the fundamental chirp, in principle it could boost the detectability of the signal by spreading its power over a frequency band where the detector is more sensitive. However, because the modulation is slow with respect to the instantaneous period of the chirp, the sidebands are too close to the main chirp for this effect to be important.

fixed; “VM”: varying) and ξ is the FAR at which the ROC is evaluated. The result is

$$78\% < \mathcal{W}_{\text{FM}}^{\text{all}} < 83\% \quad (3.9)$$

$$41\% < \mathcal{W}_{\text{FM}}^{\text{HP}} < 50\% \quad (3.10)$$

$$80\% < \mathcal{W}_{\text{VM}}^{\text{all}} < 87\% \quad (3.11)$$

$$50\% < \mathcal{W}_{\text{VM}}^{\text{HP}} < 68\% \quad (3.12)$$

where the ranges include the different possible FARs. These estimates do include the effect of coincidence and χ^2 veto. In fact the high-precision estimates are smaller than eq. (3.5) and (3.7), suggesting that the veto is penalizing strongly-modulated signals.

3.6 Conclusion

In this chapter we extend the results of chapter 2 by relaxing the assumption that BH spins are aligned with the orbital angular momenta in our simulated sources. We present the effectualness of our non-spinning and aligned-spin banks as a function of the spin parameters and we show how the sensitivity of the pipelines used in the previous chapter depends on different assumptions about spins, using realistic synthetic data.

We find again that the NSB is not effectual when the projection of the dimensionless spin on the orbital angular momentum is larger than ≈ 0.4 or smaller than ≈ -0.5 , i.e. for strongly-spinning and weakly precessing binaries. When precession is strong, including aligned-spin effects does not improve the effectualness of the bank. The resulting SNR loss depends on the orientation of the total angular momentum and can be as large as $\approx 40\%$ for cases where \mathbf{J} is roughly orthogonal to the line of sight. Nevertheless, both banks are effectual or close to effectual for precessing systems with \mathbf{J} pointing either at or away from the detector, which are also the most likely to be detected based on their intrinsic luminosity.

Assuming a universe composed of NSBH binaries with weak and tilted BHs spins, the aligned-spin pipeline is slightly less sensitive than the non-spinning one because of the larger false-alarm background. With no restriction on spin parameters, both pipelines have very similar sensitivity. However, when spins are large and almost aligned, which is a likely scenario based on existing simulations and measurements, the aligned-spin search can be roughly one order of magnitude more sensitive. The improvement is also significant assuming either large spin magnitude only or small tilt only. In the unlikely case of a vast majority of strongly-precessing binaries, both pipelines perform similarly again and the aligned-spin one has its worst sensitivity.

Using these results, we estimate that developing a generic precessing NSBH search pipeline could further increase the sensitivity by a few tens of percent or, under less realistic assumptions, possibly double it. This assumes that the FAR does not increase significantly when precessing templates are used, such that the detection threshold can be kept unchanged. The development of a precessing template bank is needed to check this assumption and also to estimate the increase in computational cost.

Figure 2.4 as well as existing studies [145, 146] suggest that our conclusions are likely to remain valid for the final sensitivity curve of advanced LIGO interferometers. Precessing IMR waveforms are now available or in final development [77, 78]; these could be a more appropriate model for the high-mass, anti-aligned systems in our population, but we do not expect the inclusion of IMR to change our results significantly. Given the large uncertainty about BH masses, the availability of computationally efficient IMR templates could enable the extension of our search to much higher masses, which may reduce the importance of including spin. This must be investigated. We also neglect tidal interaction and disruption of the NSs, which could significantly change the signals depending on the BH parameters [151]. Based on [154] however we expect tidal interaction to have a small effect on the detectability of signals. Moreover, our choice of terminating templates at ISCO should make the pipeline robust against tidal disruption. Nevertheless, this is also an effect that must be checked with further studies.

We thus conclude with the recommendation of using an aligned-spin pipeline for future NSBH searches, until a fully precessing pipeline is technically feasible, its robustness against instrumental noise artifacts and its larger sensitivity are both demonstrated and its computational requirements can be satisfied.

Chapter 4

Sensitivity of coincident and coherent CBC searches at fixed computational cost

4.1 Overview

Every past matched-filter search for CBC events with unknown coalescence time and sky position used coincidence as the method for combining data from different detectors (eq. (1.39)). We have seen in chapter 1 that data can also be combined coherently by defining a coherent matched filter (eq. (1.41)) which has a number of advantages. This method has been used for searches triggered by electromagnetic GRB observations, but not for all-sky, all-time searches [128, 142]. Comparisons between the coherent and coincident methods have been done in the past. In [173], for instance, the comparison was done for two interferometers at the same site and with the same orientation. The result suggests that the coherent method has a significantly larger sensitivity. In [174] the study was extended to two geographically separated detectors and to a more realistic definition of coincidence, showing that in this case the coherent methods perform only slightly better. Although the quantitative improvement depends on the details of the analysis and detector network, the coherent statistic seems to always be more sensitive at a fixed FAR.

With the larger bandwidth of advanced interferometers, and the availability of more developed waveform models such as the ones used in chapters 2 and 3, the number of templates and thus the computational cost of CBC searches is expected to grow significantly. A complete comparison between coincident and coherent searches, thus, should take into account the computational cost of both methods. As we saw earlier, the cost of coincident pipelines is essentially dominated by the FFT algorithm. The cost of a fully coherent search can be more complicated and has been estimated for initial detector networks, but not for advanced detectors [175, 176]. A study of the relative sensitivities of the two methods at fixed FAR *and* computational cost has never been published. It is also possible that a *hierarchical* (semicoherent) method combining the

two strategies could produce an interesting increase in sensitivity at a cost comparable to the coincident method. Such a hierarchical method has been partially investigated, but the associated cost has not been studied in detail [177]. The coherent matched filter engine used for triggered GRB searches is already using a hierarchical method [128].

In this chapter, using a semi-analytical approach, we estimate the false-alarm background, the sensitivity and the computational cost associated with coincident and coherent CBC search pipelines for different networks of advanced interferometers. We then fix the FAR and the cost and estimate the relative sensitivity of the two methods. We repeat the estimation for a hierarchical method, where a coincident detection with a low threshold is followed by a coherent follow-up of the surviving candidates. We make the simplifying assumption that the detectors produce stationary Gaussian noise and neglect the effect and cost of computing signal-based vetoes and null combinations of the data; in other words, we compare the methods under ideal data quality. We also neglect the cost of estimating the false-alarm background for establishing the significance of candidates.

The research presented in this chapter has not been published yet. The original idea for the project, important guidance and comments came from Drew Keppel. Drew also developed the code for placing templates over the sky, based on his own previous research, and provided some of the estimations of the computing cost. The author constructed the template banks, developed the various models, implemented the code for simulating the false alarms of the coincident and coherent methods, ran the simulations on the Atlas cluster and implemented the code for combining the models and producing the results. The author is grateful to Bruce Allen, Tom Dent, Shaon Ghosh and Reinhard Prix for useful comments and discussion.

The chapter is organized in the following way. Section 4.2 describes our models of the detectors and the template banks. Section 4.3 presents our models of the false alarms. In section 4.4 we estimate the computational cost of the two methods. Section 4.5 combines the previous calculations into the final comparison. Section 4.6 describes the hierarchical semicoherent method and shows the associated reduction in cost with respect to the fully coherent method. Conclusions are drawn in section 4.7. In our notation, $d = 1 \dots D$ indexes the detectors in the network, T is the duration of the analyzed data segment and N is the number of data samples. The subscripts “coi” and “coh” label quantities associated with the coincident and coherent methods respectively. We refer to a χ^2 distribution with n DOF as $\chi^2(n)$.

4.2 Detector networks and template banks

We consider here the following interferometers: LIGO Hanford (H), LIGO Livingston (L), Virgo (V), KAGRA (K) and LIGO India (I). The location and orientation of LIGO India are not established at the time of this study and we choose arbitrary values used already in [178]. For KAGRA’s parameters, see [179]. Although the first detector network to begin taking new data in 2015 is the HL combination, based on [174] we expect it to yield very similar sensitivity with both coincident and coherent methods. For simplic-

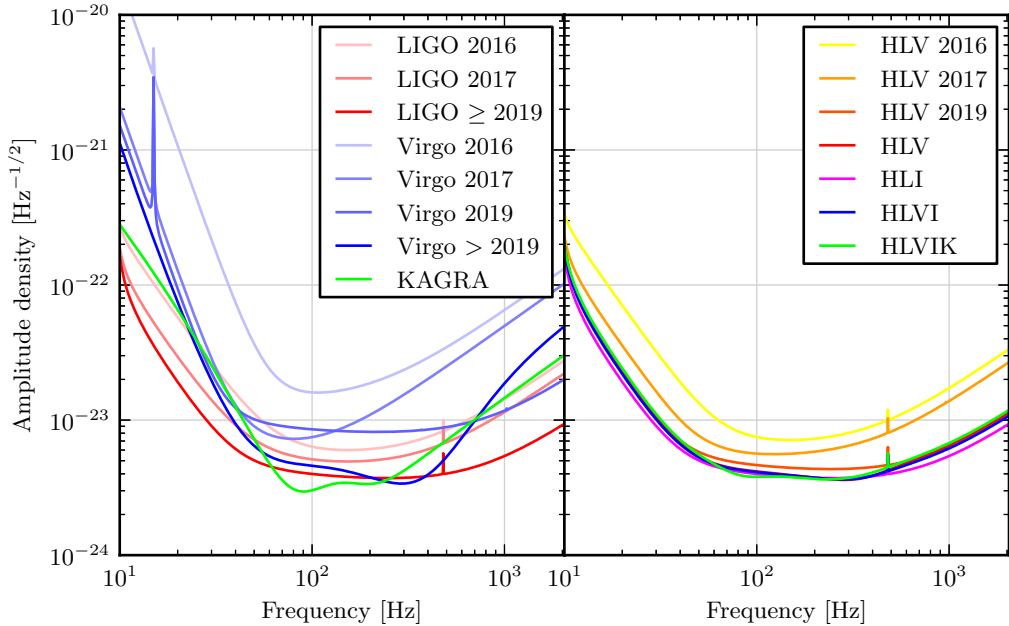


Figure 4.1: Noise models assumed for individual detectors (left) and network sensitivities obtained by harmonic-averaging the single detector curves (right).

ity we thus only consider networks of three or more detectors, specifically the following combinations: HLV, HLI, HLVI and HLVIK. The evolution of the sensitivity curves of different detectors towards their design configuration could produce a significant variation of the size of the template banks and thus noticeably affect the computational cost of the pipelines. Therefore, for the HLV network we consider four sensitivity scenarios tracing the plausible evolution towards the final design curves: HLV 2016 (“mid” sensitivity curves from [180] for HL, “early” sensitivity curve from [180] for V), HLV 2017 (“late” sensitivity for HL, “mid” for V), HLV 2019 (design sensitivity for HL, “late” sensitivity for V) and final HLV configuration (design sensitivities for every detector). For design sensitivities of LIGO, Virgo and KAGRA we use the models implemented in LALSimulation [60] and we assume identical curves for H, L and I. The amplitude spectral densities associated with the different noise models are shown in figure 4.1.

An important aspect affecting both the computational cost and the sensitivity of a pipeline is the template bank, because the number of templates that need to be searched over is directly related to the maximum mismatch μ used for constructing the bank. For simplicity and speed, in this study we neglect the spin of the components of the binary and only search over the component masses. As done in the previous chapters, we assume here an exact-match method for the coincident search, so a single bank covering the component masses is associated with each network. A coherent search, on the other hand, needs to search over both the masses and the sky location, i.e. 4 parameters instead

Network	Noise curves	Template count		k_m	k_s
		Mass	Sky		
HLV	2016	58288	1641	1749	49.2
HLV	2017	67062	1405	2012	42.2
HLV	2019	84650	2595	2540	77.9
HLV	final	78063	2723	2342	81.7
HLI	final	86001	4274	2580	128
HLVI	final	81040	6552	2431	197
HLVIK	final	73008	8006	2190	240

Table 4.1: Empirical estimates of the template-bank parameters k_m and k_s for different detector networks.

of 2. It has been found that the mass and sky parameters are weakly correlated, so the full coherent template bank can be split into the Cartesian product of a mass bank and a sky bank [181]. For each detector network we can thus employ a common mass bank for the two methods and an additional sky bank for the coherent method.

The number of templates required to cover a certain region of an n -dimensional parameter space with a maximum mismatch μ scales as $\mu^{-n/2}$ [116]. We thus model the number of mass templates as

$$M := k_m/\mu \quad (4.1)$$

where k_m is a proportionality constant which depends on the detector network. We fix this constant empirically by constructing template banks for a reference 3% maximum mismatch and counting the templates. The construction is performed using the `lalapps_tmplbank` program of LALSuite. For simplicity we use pN inspiral templates with non-spinning phase terms up to 2 pN order only. For each detector network, single-detector noise PSDs are combined into a multi-detector curve by taking their harmonic mean, as suggested in [181], and the result is used to build the bank. The resulting network sensitivities are shown in figure 4.1. The frequency limits for the metric calculation are fixed at 20 Hz and 2048 Hz. We explore individual component masses between $1M_\odot$ and $25M_\odot$ and with a total mass below $25M_\odot$. The resulting bank sizes and k_m estimates are shown in table 4.1. k_m appears to be fairly independent from the detector network, with the exception of HLV 2016 which has a visibly narrower network PSD.

The size of the sky bank given a maximum mismatch μ is modeled in a similar way,

$$S := k_s/\mu. \quad (4.2)$$

k_s is estimated, as before, by constructing sky banks for $\mu = 0.03$. Such banks are built using the stochastic placement method and the analytical metric associated with the coherent SNR [181, 182]. The same frequency limits as for the mass banks are used. The result is shown in table 4.1. Networks containing several widely separated detectors clearly require many more sky templates than smaller networks, as can be expected from

their increased angular resolution [183, 184]. Two example sky banks are shown in figure 4.2.

4.3 False-alarm rate

As explained in section 2.9, the comparison of the sensitivity of two detection methods is typically done at a fixed FAR ξ . Requiring a sufficiently low ξ in order to claim a detection implies a sufficiently high threshold that must be crossed by the detection statistic and thus reduces the number of astrophysical signals that can be detected. One step in estimating the sensitivity of a search method is therefore establishing the relation between the FAR of the method and the threshold on its detection statistic. In this section we model the relation between ξ and the thresholds on the coincident SNR ($\bar{\rho}_{\text{coi}}$) and on the coherent SNR ($\bar{\rho}_{\text{coh}}$). This model will allow us to obtain $\bar{\rho}_{\text{coi}}$ and $\bar{\rho}_{\text{coh}}$ given ξ . The most important approximation we make here is that the detector noise is stationary and Gaussian, which allows us to work with simple χ^2 distributions. A second approximation is that the false alarms of both the coincident and coherent methods are Poisson-distributed with rate $\xi(\bar{\rho})$ (where $\bar{\rho}$ can be either $\bar{\rho}_{\text{coi}}$ or $\bar{\rho}_{\text{coh}}$). This allows us to write the false-alarm probability (FAP) in a data segment of duration T as

$$p_{\text{fa}}^{\text{tot}}(\bar{\rho}) = 1 - e^{-T\xi(\bar{\rho})} \quad (4.3)$$

and the “tot” superscript indicates that we are considering false alarms from the whole template bank and data segment. In order to obtain the relation between ξ and $\bar{\rho}_{\text{coi}}$ ($\bar{\rho}_{\text{coh}}$) we then need to model the relation between $p_{\text{fa}}^{\text{tot}}$ and $\bar{\rho}_{\text{coi}}$ ($\bar{\rho}_{\text{coh}}$) which we do in the next subsections.

4.3.1 Coincident search

In a coincident search, the single-detector SNRs ρ_d for a given template are first maximized over a time window \mathcal{W} of length W samples, producing the statistic

$$\varrho_d^2 := \max_{k \in \mathcal{W}} \{\rho_d^2[k]\} \quad (4.4)$$

where k is the sample index of the SNR time series. We choose here a window of 50 ms—slightly larger than the maximum light travel time through the Earth—although in a realistic search \mathcal{W} can be several seconds long. The network statistic of the coincident method is then the incoherent SNR

$$\rho_{\text{coi}}^2 := \sum_{d=1}^D \varrho_d^2. \quad (4.5)$$

Because the ρ_d^2 are all distributed as $\chi^2(2)$ under stationary Gaussian noise, if we simply summed them in quadrature—without the maximization— ρ_{coi}^2 would be distributed as $\chi^2(2D)$ and we could directly relate the FAR to the threshold $\bar{\rho}_{\text{coi}}$. However, (i) coincidence is conditional on two or more detectors crossing the single-detector threshold first

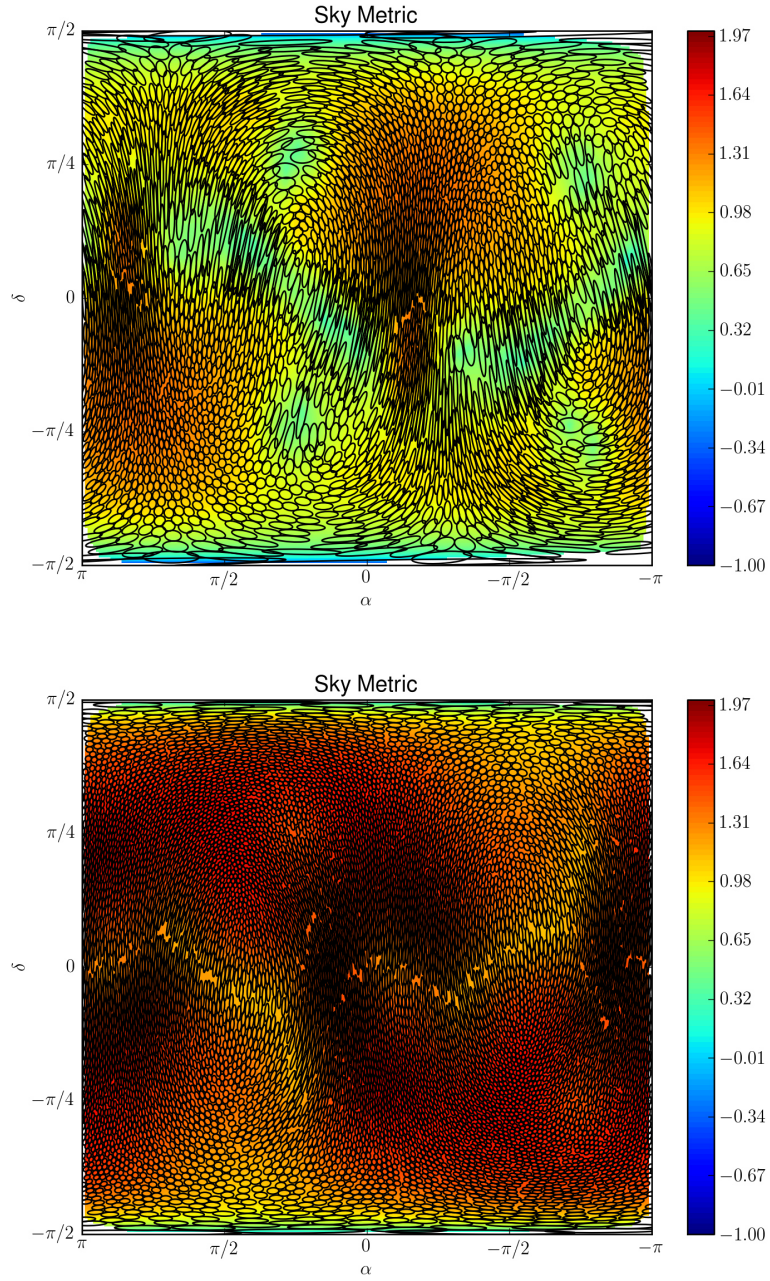


Figure 4.2: Sky template banks obtained for the HLV network (top) and the HLVK network (bottom). α and δ are the right ascension and declination at a fixed sidereal time. Each ellipse is the 1.5% mismatch contour for the associated template. The color shows the metric density in \log_{10} scale.

and (ii) ϱ_d^2 does not quite follow a $\chi^2(2)$ distribution. We deal with the first issue by assuming that the single thresholds are small enough that they can be neglected and we simply have a discrete time series of ρ_{coi} at each coalescence time sample. In order to solve the second issue, we compute first the distribution of ϱ_d^2 (momentarily dropping the d subscript for simplicity) and then use that to compute the distribution of ρ_{coi}^2 .

ϱ is smaller than a threshold $\bar{\varrho}$ if and only if all samples of ρ in the maximization window \mathcal{W} are also smaller than $\bar{\varrho}$. Therefore we can write

$$P(\varrho < \bar{\varrho}) = P(\varrho^2 < \bar{\varrho}^2) = P(\rho^2[k] < \bar{\varrho}^2 \forall k \in \mathcal{W}). \quad (4.6)$$

We now make the approximation that all samples of ρ in the maximization window are statistically independent, which allows us to write

$$P(\varrho < \bar{\varrho}) = \prod_{k \in \mathcal{W}} P(\rho^2[k] < \bar{\varrho}^2) \quad (4.7)$$

and then we use the cumulative distribution function of a $\chi^2(2)$ variable to arrive at

$$P(\varrho < \bar{\varrho}) = \left(1 - e^{-\bar{\varrho}^2/2}\right)^W. \quad (4.8)$$

The probability density of ϱ^2 is then

$$f(\varrho^2) = \left. \frac{d}{d\varrho^2} P(\varrho^2 < \bar{\varrho}^2) \right|_{\bar{\varrho}^2 = \varrho^2} = \frac{W}{2} \left(1 - e^{-\varrho^2/2}\right)^{W-1} e^{-\varrho^2/2}. \quad (4.9)$$

The mean and the variance of this distribution are

$$\mathbb{E}[\varrho^2] = 2 \sum_{w=1}^W \frac{1}{w}, \quad \text{Var}[\varrho^2] = \frac{2}{3} [\pi^2 - 6\psi'(1+W)] \quad (4.10)$$

where $\psi(x) = \Gamma'(x)/\Gamma(x)$ is the digamma function. In practice, single-detector SNRs have an autocorrelation with a typical width of the order of milliseconds for advanced-detector noise curves (figure 4.3) and thus the samples in \mathcal{W} are certainly not independent. Therefore, W must be replaced by an effective window size $W_{\text{eff}} < W$ taking into account a finite decorrelation time. As the autocorrelation level needed to make two samples statistically independent is not obvious, in practice W_{eff} is a free parameter of our model and is fixed via Monte Carlo simulations later on.

We are now ready to extend our result to a network of D detectors and compute the FAP associated with the coincident SNR,

$$p_{\text{fa}}(\bar{\rho}_{\text{coi}}) = P(\rho_{\text{coi}}^2 > \bar{\rho}_{\text{coi}}^2) = \int H\left(-\bar{\rho}_{\text{coi}}^2 + \sum_{d=1}^D \varrho_d^2\right) \prod_{d=1}^D f(\varrho_d^2) d\varrho_d^2 \quad (4.11)$$

where the integral is over the positive D -dimensional real set and $H(\cdot)$ is the Heaviside step function defining the integration limits. Using eq. (4.9), and assuming that W_{eff} is the same for each detector,

$$\prod_{d=1}^D f(\varrho_d^2) = \left(\frac{W_{\text{eff}}}{2}\right)^D \exp\left(-\sum_{d=1}^D \varrho_d^2/2\right) \times \prod_{d=1}^D \left(1 - e^{-\varrho_d^2/2}\right)^{W_{\text{eff}}-1}. \quad (4.12)$$




Figure 4.3: Duration of the autocorrelation of single-detector squared-SNR time series in stationary Gaussian noise as a function of the total mass of the template. Dashed and solid curves show respectively the 20% and 50% autocorrelation.

We now notice that the last product is always between 0 and 1 and in particular very close to 1 for $\varrho_d^2 \gg 1$. By dropping it we thus obtain an upper bound for $p_{\text{fa}}(\bar{\rho}_{\text{coi}})$, which however should be a reasonable first approximation in the limit $\bar{\rho}_{\text{coi}} \gg 1$:

$$\begin{aligned}
 p_{\text{fa}}(\bar{\rho}_{\text{coi}}) &< \left(\frac{W_{\text{eff}}}{2}\right)^D \int H\left(-\bar{\rho}_{\text{coi}}^2 + \sum_{d=1}^D \varrho_d^2\right) \exp\left(-\sum_{d=1}^D \varrho_d^2/2\right) d\varrho_d^2 \\
 &= W_{\text{eff}}^D \int_{\bar{\rho}_{\text{coi}}^2}^{\infty} dx_1 e^{-x_1} \int_{\bar{\rho}_{\text{coi}}^2 - x_1}^{\infty} dx_2 e^{-x_2} \dots \\
 &= W_{\text{eff}}^D \Phi(\bar{\rho}_{\text{coi}}^2, 2D)
 \end{aligned} \tag{4.13}$$

with $\Phi(\cdot, 2D)$ being the survival function, i.e. one minus the cumulative distribution function, of the $\chi^2(2D)$ distribution. We thus see that, to a first approximation, the maximization of single SNRs simply introduces a trials factor W_{eff}^D into an otherwise unmodified χ^2 distribution. For $W_{\text{eff}} = 1$, corresponding to no maximization, we correctly recover a $\chi^2(2D)$ distribution. For $D = 1$ we recover eq. (4.8)—or rather, its complement—in the approximation $\bar{\rho}_{\text{coi}} \gg 1$.

The previous results refer to the FAP for a single sample of the detection statistic and a single template. Since we are searching over coalescence time and component masses, the overall FAP $p_{\text{fa}}^{\text{tot}}$ is enhanced by a trials factor proportional to the proper

volume of the search parameter space, i.e. to the quantity

$$\int \sqrt{\det[g_{ij}]} dt_c dm_1 dm_2 \propto T k_m \quad (4.14)$$

where g_{ij} is the metric tensor associated with the parameter space spanning the masses and the coalescence time. Therefore we write the overall FAP, in the usual limit $\bar{\rho}_{\text{coi}} \gg 1$, as

$$p_{\text{fa}}^{\text{tot}}(\bar{\rho}_{\text{coi}}) \approx \gamma_{\text{coi}} T k_m W_{\text{eff}}^D \Phi(\bar{\rho}_{\text{coi}}^2, 2D) \quad (4.15)$$

where γ_{coi} is a proportionality factor accounting for the fact that samples of ρ_{coi} at different coalescence times and templates are not completely statistically independent. We can finally set eq. (4.15) equal to (4.3) and expand the exponential to first order, obtaining the relation between FAR and threshold

$$\Phi(\bar{\rho}_{\text{coi}}^2, 2D) \approx \frac{\xi}{\gamma_{\text{coi}} k_m W_{\text{eff}}^D} \quad (4.16)$$

valid in the regime $T\xi \ll 1$. The free parameters W_{eff} and γ_{coi} are empirically estimated by performing Monte Carlo simulations. We simulate the distribution of the coincident SNR from stationary Gaussian noise for different detector networks and perform a global fit of eq. (4.16) to the measured trigger rates. This results in $W_{\text{eff}} \approx 11.2$ samples and $\gamma_{\text{coi}} \approx 3.47 \times 10^4$. By naively considering all samples and templates statistically independent, we would obtain instead $\gamma_{\text{coi}} \approx 10^5$. At a typical sample rate of 4096 Hz, the estimated value of W_{eff} corresponds to a decorrelation time of ≈ 3 ms, i.e. an autocorrelation of a few tens of percent according to figure 4.3.

4.3.2 Coherent search

Under the approximation of stationary Gaussian noise, the FAP associated with a single template and a single sample of the coherent SNR is simply given by a $\chi^2(4)$ distribution, independently from the number of detectors [127, 128, 185]:

$$p_{\text{fa}}(\bar{\rho}_{\text{coh}}) = \Phi(\bar{\rho}_{\text{coh}}^2, 4). \quad (4.17)$$

The trials factor for going from $p_{\text{fa}}(\bar{\rho}_{\text{coh}})$ to $p_{\text{fa}}^{\text{tot}}(\bar{\rho}_{\text{coh}})$ corresponds now to a 5-dimensional search volume (coalescence time, component masses, right ascension and declination) such that

$$p_{\text{fa}}^{\text{tot}}(\bar{\rho}_{\text{coh}}) \approx \gamma_{\text{coh}} T k_m k_s \Phi(\bar{\rho}_{\text{coh}}^2, 4) \quad (4.18)$$

where γ_{coh} takes into account the fact that nearby coalescence times, mass templates and sky templates are not statistically independent. Again, by setting eq. (4.18) equal to eq. (4.3) we obtain the relationship between FAR and threshold for the coherent method

$$\Phi(\bar{\rho}_{\text{coh}}^2, 4) \approx \frac{\xi}{\gamma_{\text{coh}} k_m k_s} \quad (4.19)$$

valid for $T\xi \ll 1$. This is similar to eq. (4.16) with k_s now playing the role of W_{eff} . The only free parameter left is thus γ_{coh} , which again we estimate by simulating the false




Figure 4.4: Relationship between the FAR and the threshold on coincident SNR (solid lines) or coherent SNR (dashed lines) according to our model.

alarms associated with the coherent SNR in Gaussian noise and fitting eq. (4.19) to their empirical rate. The result is $\gamma_{\text{coh}} \approx 2.05 \times 10^5$.

The relation between the coincident and coherent thresholds and the FAR is shown in figure 4.4 for the various detector networks. At a given FAR, the coherent method has a consistently lower threshold which is fairly independent from the detector network, while $\bar{\rho}_{\text{coi}}$ increases significantly with the number of detectors. This can be explained by the different DOF of the two detection statistics. The small increase of $\bar{\rho}_{\text{coh}}$ for larger networks can be attributed to the variation of k_s with D , i.e. to the metric volume associated with the sky parameters.

4.4 Computational cost

We now estimate the computational cost associated with the coincident and coherent methods. We define the cost as the number of floating-point operations (FLOPs) required to calculate the detection statistic (either coincident or coherent SNR) for the whole data segment and template bank and generate the candidate triggers from the statistic.

As a first step, both statistics require the calculation of the single-detector complex matched-filter output. A common practice in frequency-domain inspiral searches is filtering twice as many samples as the desired output size and then ignoring the first and

Operation	Cost	Factor
Template generation	$\alpha_{\text{tmp}}N$	$\alpha_{\text{tmp}} = 10$
Frequency-domain correlation	$\alpha_{\times}N$	$\alpha_{\times} = 6$
Inverse FFT	$2\alpha_{\text{FFT}}N \log_2 2N$	$\alpha_{\text{FFT}} = 4$
SNR calculation	$\alpha_{\text{abs}}N$	$\alpha_{\text{abs}} = 3$
Trigger generation	$\alpha_{\text{trg}}N$	$\alpha_{\text{trg}} = 1$

Table 4.2: Approximations made in our computational cost estimate.

last quarters of the output time series. This avoids issues arising from the length of the impulse response of the filters and the FFT wrap-around [104]. Therefore, analyzing N samples requires processing N complex frequency bins. In particular, the following steps are required. (i) Computation of the pN template waveform at each frequency bin. This amounts to $\mathcal{O}(10)$ FLOPs per sample for a 3.5 pN template using the TaylorF2 waveform model. (ii) Correlation of data and template in the frequency domain. This is just a complex multiplication, so 6 FLOPs per sample per detector are required. (iii) Inverse Fourier transform. This can be done efficiently using a complex-to-complex FFT of size $2N$, adding $\mathcal{O}(2N \log_2 2N)$ FLOPs per data segment for each detector.

After the basic operations listed above, the coincident method proceeds by calculating the squared magnitude of single-detector SNRs (3 FLOPs per sample per detector), maximizing the SNR over the window \mathcal{W} (which only involves comparisons and thus we neglect) and summing the SNRs across detectors ($D - 1$ FLOPs per sample). Finally, generating the list of triggers from the resulting time series adds a small contribution to the total cost, which we assume to be just one operation per sample. The total FLOP count per data segment for the coincident method is then

$$C_{\text{coi}} = M_{\text{coi}}N [\alpha_{\text{tmp}} + D(\alpha_{\times} + 2\alpha_{\text{FFT}} \log_2 2N + \alpha_{\text{abs}} + 1) - 1 + \alpha_{\text{trg}}]. \quad (4.20)$$

The values we take for the α . constants are summarized in table 4.2. From eq. (4.20) we see that $C_{\text{coi}} \approx 2M_{\text{coi}}ND\alpha_{\text{FFT}} \log_2 2N \approx 100 M_{\text{coi}}ND$, i.e. the FFT dominates the total cost, as we already mentioned in chapters 1 and 2.

The first operation of the coherent method, after the common matched-filtering steps discussed above, is time-shifting the matched-filter outputs to account for the different signal arrival times. This can be done via an interpolating filter with N_{ts} taps, requiring $2N_{\text{ts}}$ FLOPs per sample, per sky location and per detector ($N_{\text{ts}} = 0$ corresponding to nearest-neighbor interpolation). The next step is calculating x_i (eq. (1.42)), requiring $8D - 4$ FLOPs per sample and sky location. The coherent SNR, eq. (1.41), can then be formed efficiently with 15 more FLOPs per sample and sky location. In fact, we can write $\rho_{\text{coh}}^2 = y_i y^i$ where $y_i := x^j \mathcal{L}_{ji}$ and \mathcal{L}_{ij} is the Cholesky decomposition of the matrix

B^{ij} defined in eq. (1.46). The only non-null elements of \mathcal{L}_{ji} are

$$\mathcal{L}_{11} = \mathcal{L}_{33} = \sqrt{\frac{\mathcal{B}_{\times\times}}{\mathcal{B}_{++}\mathcal{B}_{\times\times} - \mathcal{B}_{+\times}^2}} \quad (4.21)$$

$$\mathcal{L}_{21} = \mathcal{L}_{43} = -\sqrt{\frac{\mathcal{B}_{+\times}^2/\mathcal{B}_{\times\times}}{\mathcal{B}_{++}\mathcal{B}_{\times\times} - \mathcal{B}_{+\times}^2}} \quad (4.22)$$

$$\mathcal{L}_{22} = \mathcal{L}_{44} = \sqrt{\frac{\mathcal{B}_{++} - \mathcal{B}_{+\times}^2/\mathcal{B}_{\times\times}}{\mathcal{B}_{++}\mathcal{B}_{\times\times} - \mathcal{B}_{+\times}^2}} \quad (4.23)$$

where \mathcal{B}_{++} , $\mathcal{B}_{+\times}$ and $\mathcal{B}_{\times\times}$ are defined in eq. (1.47). Finally, the trigger generation must now be performed for each sky location, so α_{trg} gets a factor S . In total,

$$C_{\text{coh}} = M_{\text{coh}}N [\alpha_{\text{tmp}} + D(\alpha_{\times} + 2\alpha_{\text{FFT}} \log_2 2N) + S(2DN_{\text{ts}} + 8D + 11 + \alpha_{\text{trg}})]. \quad (4.24)$$

Note that if $\mu_{\text{coh}} \lesssim 0.03$ then $S \gtrsim 10^3$ (table 4.1) and thus, depending on how the sky bank is constructed, the cost of the FFT can be easily overwhelmed by the calculation of the coherent SNR for all sky locations.

We can now set eq. (4.20) and (4.24) equal to the same cost C and express the number of templates in terms of the maximum mismatches μ_{coi} , μ_{coh} via eq. (4.1) and (4.2). This establishes the following relation between μ_{coi} and μ_{coh} :

$$\mu_{\text{coi}} = \mu_{\text{coh}} \frac{\frac{\alpha_{\text{tmp}}}{D} + \alpha_{\times} + 2\alpha_{\text{FFT}} \log_2 2N + \alpha_{\text{abs}} + \alpha_{\text{trg}}}{\frac{\alpha_{\text{tmp}}}{D} + \alpha_{\times} + 2\alpha_{\text{FFT}} \log_2 2N + \frac{k_s}{\mu_{\text{coh}}} \left(2N_{\text{ts}} + \frac{\alpha_{\text{trg}} + 11}{D} + 8 \right)}. \quad (4.25)$$

When μ_{coh} is very small, eq. (4.25) simplifies to $\mu_{\text{coi}} \propto \mu_{\text{coh}}^2$, which gives an idea of the significant difference between the computational cost of the two methods when the sky is covered with many templates. Figure 4.5 shows a comparison between μ_{coi} and μ_{coh} as a function of the fixed computational cost required to analyze one second of data. As can be seen, the two quantities scale with a different power and μ_{coi} is always smaller by orders of magnitude.

4.5 Sensitivity

Given a detection statistic ρ (either coincident or coherent SNR) and a threshold $\bar{\rho}$, a particular binary is visible on average if

$$\text{E}[\rho^2] \geq \bar{\rho}^2 \quad (4.26)$$

where the expectation value is taken over the ensemble of noise realizations. In stationary Gaussian noise, ρ^2 has a non-central $\chi^2(n)$ distribution with non-centrality $(\lambda\rho_x/r)^2$, such that

$$\text{E}[\rho^2] = \left(\frac{\lambda\rho_x}{r} \right)^2 + n \quad (4.27)$$

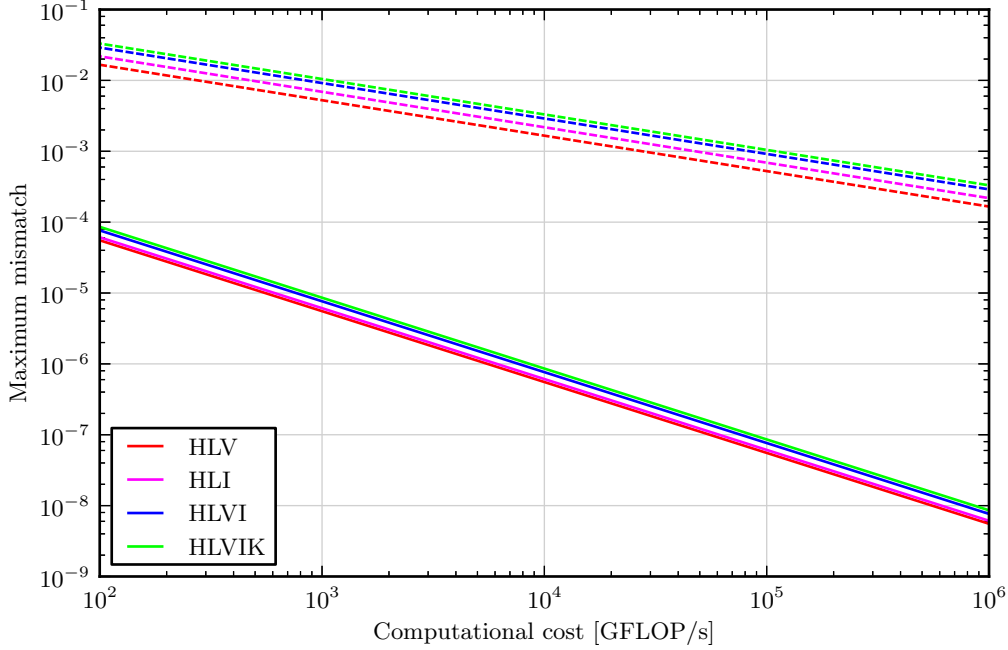


Figure 4.5: Maximum mismatch of the coincident method (solid lines) and coherent method (dashed lines) for a given computational cost.

where ρ_x is the optimal SNR for the binary at unit distance, r is the distance and $0 \leq \lambda \leq 1$ represents the SNR loss introduced by the particular detection method for the binary we are considering. By combining eq. (4.26) and (4.27) we can define the visible range

$$r_v := \frac{\lambda \rho_x}{(\bar{\rho}^2 - n)^{1/2}}. \quad (4.28)$$

r_v is a function of the component masses, the sky location $\hat{\mathbf{r}}$ and the orientation of the orbit $\hat{\mathbf{J}}$, as well as the geometry of the network and the sensitivity curves of the detectors. We can first of all remove the dependence on $\hat{\mathbf{r}}$ by integrating over the sky up to a distance $r_v(\hat{\mathbf{r}})$, which defines a sensitive volume,

$$V := \int_{S^2} \left(\int_0^{r_v(\hat{\mathbf{r}})} r^2 dr \right) d\hat{\mathbf{r}} = \frac{1}{3} \int_{S^2} r_v^3(\hat{\mathbf{r}}) d\hat{\mathbf{r}}. \quad (4.29)$$

We can further average the volume over masses and orientation, obtaining

$$\langle V \rangle = \frac{1}{3(\bar{\rho}^2 - n)^{3/2}} \left\langle \int_{S^2} \lambda^3 \rho_x^3 d\hat{\mathbf{r}} \right\rangle. \quad (4.30)$$

Although in principle both λ and ρ_x depend on the parameters of the binary, in practice the template banks are so fine that λ^3 varies much more rapidly with the parameters

than ρ_x^3 . We can thus replace λ^3 with its average over the parameters and take it out of the integral,

$$\langle V \rangle \approx \frac{\langle \lambda^3 \rangle}{3(\bar{\rho}^2 - n)^{3/2}} \left\langle \int_{S^2} \rho_x^3 d\hat{r} \right\rangle. \quad (4.31)$$

We now note that despite the important differences in how they are calculated, both ρ_{coi}^2 and ρ_{coh}^2 measure the total signal power projected on the network. In other words, the optimal coincident and coherent SNRs are both equal to ρ_x/r . By taking the ratio of the average sensitive volumes of the coherent and coincident methods, we thus get rid of the integrals and obtain a simple expression for the gain in average sensitive volume,

$$\mathcal{V} := \frac{\langle V_{\text{coh}} \rangle}{\langle V_{\text{coi}} \rangle} = \left(\frac{\bar{\rho}_{\text{coi}}^2 - n_{\text{coi}}}{\bar{\rho}_{\text{coh}}^2 - n_{\text{coh}}} \right)^{3/2} \frac{\langle \lambda_{\text{coh}}^3 \rangle}{\langle \lambda_{\text{coi}}^3 \rangle} = \left(\frac{\bar{\rho}_{\text{coi}}^2 - 2D}{\bar{\rho}_{\text{coh}}^2 - 4} \right)^{3/2} \frac{\langle \lambda_{\text{coh}}^3 \rangle}{\langle \lambda_{\text{coi}}^3 \rangle} = \mathcal{V}_{\text{fa}}(\xi) \mathcal{V}_{\text{loss}}(C) \quad (4.32)$$

where in the last step we have defined the factor containing the thresholds as $\mathcal{V}_{\text{fa}}(\xi)$ and the factor containing the losses as $\mathcal{V}_{\text{loss}}(C)$. Assuming a population of coalescing binaries with uniform distribution in volume, \mathcal{V} gives directly the improvement in number of detections. In the limit of infinite computational cost, we can afford large template banks and very precise interpolation, so the SNR losses vanish and we obtain

$$\lim_{C \rightarrow \infty} \mathcal{V} = \mathcal{V}_{\text{fa}}(\xi) = \left(\frac{\bar{\rho}_{\text{coi}}^2 - 2D}{\bar{\rho}_{\text{coh}}^2 - 4} \right)^{3/2}. \quad (4.33)$$

This limit depends purely on our false-alarm model.

The SNR loss of the coincident method λ_{coi} is determined entirely by the mismatch between the signal and the closest mass template. We approximate the mass bank as a 2-dimensional hexagonal lattice, which is the optimal covering scheme [116]. By considering that the mismatch plays the role of the squared metric distance (eq. (1.32)), and computing the distribution of the distance from the center of a regular hexagon over the whole hexagon, one can show that

$$\langle \lambda_{\text{coi}}^3 \rangle = 1 - \frac{5}{4}\mu_{\text{coi}} + \frac{7}{10}\mu_{\text{coi}}^2 - \frac{83}{560}\mu_{\text{coi}}^3. \quad (4.34)$$

The coherent search, however, also loses SNR due to the finite mismatch with the sky bank. Since we take the mass and sky banks to be independent, and both use the same maximum mismatch, we make the approximation that the mismatch with the sky bank introduces an additional factor identical to (4.34). A small loss also occurs when the single-detector matched-filter-output time series are interpolated to synchronize the arrival times at the different detectors, giving an average factor λ_{ts} . Thus

$$\langle \lambda_{\text{coh}}^3 \rangle = \left(1 - \frac{5}{4}\mu_{\text{coh}} + \frac{7}{10}\mu_{\text{coh}}^2 - \frac{83}{560}\mu_{\text{coh}}^3 \right)^2 \lambda_{\text{ts}}^3. \quad (4.35)$$

In section 4.4 we saw that the condition of fixed computational cost requires $\mu_{\text{coi}} < \mu_{\text{coh}}$. Since $0 < \lambda_{\text{ts}} < 1$, this implies that $0 < \langle \lambda_{\text{coh}}^3 \rangle < \langle \lambda_{\text{coi}}^3 \rangle < 1$, i.e. $\mathcal{V}_{\text{loss}}(C) < 1$: at fixed computational cost, the coherent method has always a larger SNR loss.

We can now combine eq. (4.16), (4.19), (4.25), (4.34) and (4.35) with (4.32) to numerically evaluate \mathcal{V} as a function of the computational cost for different detector configurations. Several parameters need to be chosen for this evaluation; our choice is $\xi = 10^{-4} \text{ yr}^{-1}$, $N = 2^{22}$, $T = 1024 \text{ s}$ (4096 samples/s), $N_{\text{ts}} = 0$ and $\lambda_{\text{ts}} = \sqrt{0.99}$. The resulting \mathcal{V} curves are shown in the top plot of figure 4.6 as a function of the computational cost for analyzing one second of data. The advantage of the coherent method increases monotonically with the computational cost as its SNR loss is reduced and eventually approaches \mathcal{V}_{fa} at very large cost. The limit \mathcal{V}_{fa} as a function of the FAR is shown in the bottom plot of figure 4.6. Increasing the number of detectors seems to enhance the advantage of the coherent method if the available computational power is high enough to enable a realistically small μ_{coh} , since the false alarms of the coincident method are increased by the additional DOF introduced by each new detector, while the coherent SNR has only 4 DOF irrespective of D . If the computational power is too restricted, however, the 4-dimensional parameter space can not be explored with enough resolution and the SNR loss of the coherent method is too large for the method to be efficient. It also appears that the most important parameter explaining the variation of \mathcal{V}_{fa} with the detector network is just the number of detectors rather than the details of the particular network (such as the constants k_m and k_s). We also see that \mathcal{V}_{fa} slowly decreases with decreasing FAR, suggesting that strong signals are detected equally well by both methods, as one would expect. Note however that, as stated before, our false-alarm model for the coincident search gives only an upper limit at small thresholds. Hence, \mathcal{V}_{fa} may be overestimated towards large ξ . Interestingly, imposing a fixed FAR always favors the coherent method (i.e., $\mathcal{V}_{\text{fa}} > 1$) while fixing the computational cost always penalizes it ($\mathcal{V}_{\text{loss}} < 1$).

4.6 Hierarchical method

In section 4.4 it was shown that, when μ_{coh} is very small, forming the combination of the matched-filter outputs at each sky location easily becomes the dominant cost of a fully coherent search. From figure 4.6 we also saw that the coherent search starts to be more sensitive when μ_{coh} is a few tens of percent or less, depending on the number of detectors. In this regime, the calculation of ρ_{coh} over the sky is expensive. However, if the single-detector SNRs are small enough, it may not be necessary to calculate ρ_{coh} on the whole sky in order to rule out a possible signal. Also note that the operations required for computing ρ_{coi} are almost a subset of those required for ρ_{coh} . Thus, we investigate next whether it is possible to reduce the cost of the coherent search by combining the coincident and coherent methods in a hierarchical, or semicoherent, way. By this we mean (i) computing first each sample of the coincident SNR for a given template, (ii) comparing it to a threshold $\bar{\rho}_{\text{h}}$ and (iii) calculating the coherent SNR at each sky location for the coalescence times of the surviving samples only, by reusing the matched-filter outputs already calculated for step (i).

Let α_{h} be the average fraction of samples that pass the first test under the null hypothesis and thus require the calculation of the full coherent SNR. This is just the FAP




Figure 4.6: Sensitivity of the coherent search relative to the coincident search as a function of the computational cost for a FAR of 10^{-4} years (top) and as a function of the FAR at infinite cost (bottom). The triangles mark the points where the template banks have the conventional 3% maximum mismatch in the coincident search (∇) or in the coherent one (Δ). For all networks the coherent method is more sensitive past a critical value of the computational cost. The asymptotic gain is larger for larger networks.

of the coincident method for a single template and sample, which we already determined in eq. (4.13),

$$\alpha_h \approx W_{\text{eff}}^D \Phi(\bar{\rho}_h^2, 2D). \quad (4.36)$$

This is a good approximation in the limit $\bar{\rho}_h \gg 1$ and an upper limit otherwise. The average cost of the hierarchical method is then no larger than

$$\begin{aligned} C_h &= M_{\text{coh}} N [\alpha_{\text{tmp}} + D(\alpha_{\times} + 2\alpha_{\text{FFT}} \log_2 2N + \alpha_{\text{abs}} + 1) \\ &\quad + \alpha_h S (2DN_{\text{ts}} + 8D + 11 + \alpha_{\text{trg}})] \\ &=: M_{\text{coh}} N (C_{\text{MF}} + \alpha_h C_{\text{sky}}) \end{aligned} \quad (4.37)$$

where C_{MF} is the cost of the matched filtering and $C_{\text{sky}} \propto S$ the cost of constructing the coherent SNR for the whole sky. Since we are looking at the regime in which $C_{\text{MF}} \ll C_{\text{sky}}$, the relative cost of the hierarchical vs fully coherent search (shown in figure 4.7 for $\mu_{\text{coh}} = 0.03$) is approximately

$$\frac{C_h}{C_{\text{coh}}} \approx \frac{C_{\text{MF}}}{C_{\text{sky}}} + \alpha_h. \quad (4.38)$$

For $\mu_{\text{coh}} = 0.03$, an interesting cost reduction is achievable with $\alpha_h \approx 10^{-3}$, corresponding to $\bar{\rho}_h$ between 6 and 8 depending on the number of detectors. This happens to be in the regime where eq. (4.36) is not a good approximation but rather an upper limit, so in practice the cost will be even less than predicted (although always at least as large as the $C_{\text{MF}}/C_{\text{sky}}$ asymptote). A higher $\bar{\rho}_h$ brings no further reduction because the cost is eventually saturated by the matched filter operation. In fact, as we want the coincident stage to have a negligible effect on the sensitivity at some given FAR, we also have the condition $\bar{\rho}_h < \bar{\rho}_{\text{coh}}$. Thus a sensible choice is to set $\bar{\rho}_h$ just at the turnover of the curves in figure 4.7, for instance by solving the equation $\Phi(\bar{\rho}_h^2, 2D) = C_{\text{MF}}/2C_{\text{sky}}W_{\text{eff}}^D$ for $\bar{\rho}_h$.

Taking for simplicity $\bar{\rho}_h$ such that $\alpha_h = 10^{-3}$, we now repeat the comparison done in the previous section and compute the sensitivity improvement of the hierarchical method over the coincident one at fixed computational cost. The results for different networks are displayed in figure 4.8 as a function of the cost. We can see that the result is the same as for the fully coherent method, but the curves are shifted to a computational cost roughly two orders of magnitude smaller. For 3-detector networks, the two methods have a very similar sensitivity already at a cost corresponding to the canonical choice $\mu_{\text{coi}} \approx 0.03$, and for networks of 4 or 5 detectors there is already a noticeable improvement.

4.7 Conclusion

Through simplified analytical models and simulations, we investigate the sensitivity and computational cost of CBC searches with advanced interferometers based on a coincident network statistic and on a coherent one. We assume stationary Gaussian noise and we neglect the cost of computing the signal-based vetoes and null streams, as well as the estimation of the false-alarm background. Consistent with past studies, we find




Figure 4.7: Upper limit on the average cost of the hierarchical method, relative to the fully coherent one, as a function of the pre-threshold on the coincident SNR and for a coherent maximum mismatch of 3%. The shaded lines show the fraction α_h of samples requiring the calculation of the coherent SNR. For moderately large values of the pre-threshold, the cost of the hierarchical search is dominated by the matched filtering, which is much smaller than the cost of the fully coherent search.




Figure 4.8: Sensitivity of the hierarchical search relative to the coincident one as a function of the computational cost, for a FAR of 10^{-4} years and $\alpha_h = 10^{-3}$. The triangles mark the points where the coincident template bank (∇) or the coherent one (Δ) are built using the conventional 3% maximum mismatch. The hierarchical search outperforms the coincident one at a much smaller computational cost than the fully-coherent one.

that a fully coherent search is more sensitive, provided that the computational resources are enough to enable sufficiently fine mass and sky template banks. Because the DOF of the coincident statistic increase with the number of detectors, while the DOF of the coherent statistic are network-independent, the coherent method produces a larger gain in sensitivity for larger detector networks. For networks of three detectors, the resulting gain in detection rate is below 10%: the coincident method is not very far from the optimal performance of the coherent one. On the other hand, networks of five or more detectors have a more interesting gain of at least a few tens of percent. A hierarchical (or semicoherent) method, where the computation of the coherent SNR is conditional on crossing a coincident threshold first, should be able to achieve the maximum sensitivity gain of the fully coherent method at a computational cost comparable to that of conventional coincident searches. This could be important for networks of four or five advanced interferometers. In the presence of more realistic, non-stationary noise the advantage of the coherent search could be even larger, provided that the additional consistency tests that the method enables (such as the null SNR statistic) are employed¹. However this would likely complicate our simple semianalytical treatment of the problem.

Neglecting the cost of vetoes and background estimation is a good approximation for the coincident search, but may not be so for the coherent one. Background estimation via time slides, in particular, may require filtering the data multiple times with different time shifts between detectors. This would multiply the coherent cost by a factor determined by the SNR of the loudest zero-lag triggers, which could be one order of magnitude or more. The development of alternative methods of background estimation may however lift this requirement.

In this investigation we consider for simplicity non-spinning inspiral templates with phase terms up to 2 pN order. As shown in chapters 2 and 3, however, non-spinning terms up to order 3.5 and aligned-spin terms up to at least 2.5 order can and should be included in future searches. Similarly to mass parameters, we expect aligned-spin effects to be somewhat orthogonal to the sky parameters. Moreover, the effective dimensionality of a bank of aligned-spin pN waveforms is not much larger than the non-spinning case. Based on these considerations, we do not expect a dramatic deviation from our result when higher-order and spinning pN templates are used (except, obviously, for a larger overall computational cost). However, results could be significantly different if *precessing* waveforms are used, since those introduce an additional dependence on sky location.

Our accounting of the computational cost only considers arithmetic operations and ignores the cost of memory access. In some cases, however, the cost of the FFT algorithm could have a significant contribution from memory input/output operations. Computing the coherent SNR from the matched filter outputs, on the other hand, should have a high enough number of arithmetic operations to make the cost of memory access negligible. If our assumption of the FFT cost is underestimated because of this effect, the matched filter would give a bigger contribution to the total cost of the fully coherent and hierarchical methods. Thus, the gain of the (semi)coherent methods at a given cost

¹In the absence of vetoes and with non-stationary noise, in fact, the coherent method may actually be *less* efficient than coincidence [186].

would be slightly larger.

An open question is whether the additional sensitivity of the coherent method can also be obtained with the coincident method without the need for an explicit search over the sky and the simultaneous filtering of the data from different detectors. A route towards improving the sensitivity of the coincident method could be to fold the observed difference in SNR, coalescence times and phases at different detectors into the ranking of coincident triggers. This investigation is ongoing and we leave the presentation of the results to future work.

Chapter 5

Modeling the effect of glitches on CBC searches

5.1 Overview

Chapter 1 highlighted the issue of noise transients (*glitches*) affecting the results of searches for CBC signals. Despite the improved sensitivity of advanced interferometers, such signals are also expected in data from future observing runs. Because of the increase in sensitivity of advanced interferometers at low frequencies (10–40 Hz) CBC templates will become significantly longer than they were in the past. For instance, an NSBH signal ($m_{\text{BH}} = 10M_{\odot}$, $m_{\text{NS}} = 1.4M_{\odot}$, no spins) starting at 10 Hz would be visible for more than 200 s in advanced LIGO; the same signal would have been visible for about 6 s only in initial LIGO. The introduction of much longer templates raises a new concern: short glitches could produce spurious CBC triggers with a significant delay with respect to the central time of the glitch and certainly much longer than the glitch itself. Therefore, the association between a trigger and a glitch is no longer guaranteed by their vicinity in time, but must take into account a potentially large delay. Moreover, the delay is in principle a function of the template parameters and the glitch waveform.

This chapter addresses this issue in a semi-analytical way by studying the effect of short, isolated, non-CBC transients on the matched filter operation at the core of CBC searches. We use a simple sine-Gaussian model to describe glitch waveforms, parametrized by a central frequency and a quality factor. CBC templates are simplified as Newtonian (0 pN) inspiral-only waveforms. Using three different approximations, appropriate in different regions of the glitch parameter space, we are then able to estimate the average SNR time series produced by the CBC matched filter as a function of the parameters of the glitch and the template. By maximizing the time series we obtain the time and SNR of triggers produced by the glitch. We compare these to numerical simulations to check the accuracy of each approximation.

The research presented in this chapter is published in [187] and the contents of the chapter are mostly unchanged. The idea for the project, major guidance and comments are from Andrew Lundgren. Swetha Bhagwat and Sanjeev Dhurandhar developed Ap-

proximation I and some simulations. The author developed Approximations II and III, wrote and ran the code for the numerical simulations, wrote the article and handled the publication and communication with the anonymous referees. Useful comments and help came from Bruce Allen, Tom Dent, Drew Keppel, Badri Krishnan, Alex Nielsen and Chris Pankow.

The chapter is organized as follows. Section 5.2 introduces our models of glitch and template waveforms, formalizes the problem and presents the approximations to the matched-filter integral. In section 5.3 the approximations are tested against numerical evaluations of the matched-filter integral. The response of a full template bank to a glitch is also described there. Conclusions are drawn in section 5.4.

5.2 Glitch model and matched-filter approximation

The problem we address is calculating the average response of the CBC matched filter (eq. (1.28)) to a transient signal which reasonably represents glitches. Once the SNR time series is available, we want to obtain the value of its maximum ρ^{\max} and the corresponding time t_c^{\max} , because these are the most important parameters of the trigger produced by the glitch for a given template. Although we must assume a particular noise PSD in order to evaluate the matched-filter integral, we ignore the actual noise realization and only include the glitch waveform in our data model; thus, our result represents an average over the ensemble of Gaussian noise realizations. Moreover, we only consider the case of an isolated glitch; the far less likely case of a glitch overlapping with an actual CBC signal is not studied here.

A realistic model of every class of glitch waveforms observed in initial LIGO data is a complicated endeavor and requires either an understanding of the physical origin of glitches or a large number of observations of each glitch class. Such requirements are not always met. A detailed model, moreover, is likely unnecessary if we are only interested in the maximum of the inspiral SNR. Therefore, we do not go into such detail here. A common feature of many glitches, however, is their approximate localizability in either time, frequency, or both [188]. A suitable general model, which is also very simple, is therefore represented by monochromatic signals whose amplitude is modulated with a Gaussian envelope, i.e. *sine-Gaussian* waveforms. This is also the signal basis used by burst detection pipelines (such as Omega [189]) routinely used for investigating the quality of data from GW detectors. Our glitch model is therefore

$$s(t) = A \exp\left(-\frac{(t-t_0)^2}{\tau^2}\right) \cos(2\pi f_0 t + \phi_0) \quad (5.1)$$

where A is the overall amplitude of the glitch, t_0 and f_0 are the central time and frequency, $\tau := Q/2\pi f_0$ is the time duration, Q is the dimensionless quality factor and ϕ_0 is the phase at $t = 0$. t_0 can be arbitrarily set to zero, such that t_c^{\max} represents the time delay between the glitch and the CBC trigger. We can also set $A = 1$ since ρ^{\max} scales linearly with the amplitude of the glitch. The frequency- domain expression of

the glitch is then

$$\tilde{s}(f) = \frac{\sqrt{\pi}}{2} \tau \exp(-\pi^2 \tau^2 (|f| - f_0)^2 + i\phi_0) \left[1 + \exp\left(-Q^2 \frac{|f|}{f_0} - 2i\phi_0\right) \right] \quad (5.2)$$

for frequencies $f > 0$ and the complex conjugate of that for $f < 0$. The magnitude of the second exponential is always at most 1. If the glitch completes at least a few oscillations ($Q > 1$) or its center frequency is not much larger than the low-frequency cutoff of the matched filter ($f_0 \approx f_L$) then the magnitude of the exponential is actually much smaller than 1. We make this approximation and neglect the factor in square brackets altogether. The SNR is then no longer affected by the phase of the glitch and we can set $\phi_0 = 0$, so we are left with only two parameters for our glitches: f_0 and Q .

Assuming a pN frequency-domain inspiral template as in eq. (1.24), the matched-filter integral for a data segment containing only our glitch reads

$$z(t_c) = 2\sqrt{\pi}\tau h_0 \int_{f_L}^{f_H} U(f) \exp[-\pi^2 \tau^2 (f - f_0)^2 + i\Psi(f; t_c)] df \quad (5.3)$$

where $U(f) := f^{-7/6} S^{-1}(f)$, $S(f)$ is the noise PSD, h_0 contains all frequency-independent factors of the template amplitude and we have neglected the parameter vector of the template for simplicity. Although chapter 2 demonstrated the importance of including spin and high-order pN terms in the phase Ψ , we are mainly concerned here with the largest contribution to the duration of the inspiral. It is therefore sufficient to only consider a Newtonian (0 pN) template, which only depends on the chirp mass \mathcal{M} and greatly simplifies the problem:

$$\Psi(f; t_c) = 2\pi f t_c - \phi_c - \frac{\pi}{4} + \frac{3}{128} (\zeta f)^{-5/3} \quad (5.4)$$

with $\zeta := \pi G \mathcal{M} / c^3$.

5.2.1 Approximation I

One way to evaluate eq. (5.3) is via the stationary-phase approximation (SPA). In fact, for the Newtonian chirp we have $d\Psi/df = 0$ at a frequency

$$f_s := \left(\frac{5}{256\pi}\right)^{3/8} \zeta^{-5/8} t_c^{-3/8} = f_0 \left(\frac{\tau_0}{t_c}\right)^{3/8} \quad (5.5)$$

with τ_0 being the Newtonian chirp time [160] evaluated using f_0 as the fiducial frequency. This allows us to approximate the phase as

$$\Psi(f; t_c) \approx \Psi(f_s) + \beta (f - f_s)^2 \quad (5.6)$$

$$\beta := \frac{5}{96} f_s^{-11/3} \zeta^{-5/3}. \quad (5.7)$$

We now notice that, in eq. (5.3), the function $U(f)$ varies more slowly with $(f - f_s)$ than the rest of the integrand. We thus evaluate it at f_s and bring it out of the integral, which becomes the integral of the product of two Gaussians:

$$z(t_c) \approx 2\sqrt{\pi}\tau U(f_s) \int_{f_L}^{f_H} \exp\left(i\frac{(f - f_s)^2}{2\sigma_f^2} - \frac{(f - f_0)^2}{2\sigma_{sg}^2}\right) df \quad (5.8)$$

where we have omitted the global phase factors and the template amplitude h_0 , since they disappear once the SNR is calculated from $z(t_c)$. The standard deviations of the two Gaussians are

$$\sigma_f := \left(\frac{3f_0}{16\pi\tau_0}\right)^{1/2} \left(\frac{f_s}{f_0}\right)^{11/6} \quad (5.9)$$

$$\sigma_{sg} := \frac{1}{\sqrt{2\pi}\tau}. \quad (5.10)$$

The integral can be solved via Cauchy's theorem and we can then use the resulting $z(t_c)$ and eq. (1.28) to obtain the SNR as a function of t_c ,

$$\rho(t_c) \approx \frac{1}{\sigma_0} U(f_s) \left(1 + \frac{\sigma_{sg}^4}{\sigma_f^4}\right)^{-1/4} \exp\left(-\frac{(f_0 - f_s)^2}{2\sigma_f^2} \frac{\sigma_{sg}^2/\sigma_f^2}{1 + \sigma_{sg}^4/\sigma_f^4}\right) \quad (5.11)$$

with $\sigma_0^2 := \int_{f_L}^{f_H} f^{-7/3} S^{-1}(f) df$ being the normalization of the template.

Despite our approximations, eq. (5.11) is still a pretty complicated function of the coalescence time through f_s and σ_f . In particular, the fact that the coalescence time appears in the argument of $S(f_s)$ means that the shape of the function depends on the noise PSD. For a smooth PSD with a single global minimum, $\rho(t_c)$ turns out to be a fairly simple peak-shaped function. Nevertheless, ρ^{\max} and t_c^{\max} must be found via numerical maximization of $\rho(t_c)$.

5.2.2 Approximation II

The SPA can be applied in a different way to give another approximation, valid for small Q . We still use eq. (5.6), but this time we make the assumption that the integral in eq. (5.3) is dominated by a narrow band around f_s . We can therefore consider most of the terms in the integrand as constant (evaluated at f_s) and take them out of the integral, which reduces to

$$\int_{f_L}^{f_H} \exp[i\beta(f - f_s)^2] df \approx \sqrt{\frac{\pi}{\beta}}. \quad (5.12)$$

In this way we obtain the following SNR time series:

$$\rho(t_c) \approx \pi\sigma_0^{-1}\tau U(f_s)\beta^{-1/2} \exp[-\pi^2\tau^2(f_s - f_0)^2] \quad (5.13)$$

For the Newtonian inspiral in particular, using eq. (5.5) and (5.7),

$$\rho(t_c) \approx \frac{\sqrt{6}\pi^{3/4}}{5^{1/4}}\sigma_0^{-1}\tau\zeta^{5/12}S^{-1}\left[f_0(\tau_0/t_c)^{3/8}\right]t_c^{-1/4}\exp\left[-\frac{Q^2}{4}\left((\tau_0/t_c)^{3/8}-1\right)^2\right]. \quad (5.14)$$

Although this second approximation is somewhat simpler than eq. (5.11), it still depends on the shape of the noise PSD. If we assume the PSD is a well-behaved function with no narrow peaks and one global minimum, however, we can see that eq. (5.14) is the product of two peaks in t_c with different time scales: one given by the exponential (with a time scale determined by Q) and one given by the other terms containing t_c (with a time scale determined by the bandwidth of the detector). In order to get a better understanding of this shape, we can try to make these time scales very different by going to the limit $Q \gg 1$. In this limit the exponential becomes more peaked than the other term, so the maximum SNR is approximately given by the condition $f_s = f_0$ and we obtain

$$\rho^{\max} \approx \pi\tau\sigma_0^{-1}U(f_0)\beta^{-1/2} \quad (5.15)$$

$$t_c^{\max} \approx \tau_0. \quad (5.16)$$

This suggests that, when Q is sufficiently large, the delay between the glitch and the resulting trigger is just the time taken by the inspiral signal to go from the center frequency of the glitch to coalescence. This result can be expected from intuition by visualizing the time-frequency support of the inspiral signal and the sine-Gaussian signal. Note however that, since this approximation is expected to be valid for small Q , this limit is not necessarily meaningful and the result needs to be checked against numerical simulations.

5.2.3 Approximation III

The $Q \gg 1$ case is probably the least useful in practice, because glitches typically do not last more than a few cycles. However, this limit enables one more approximation for $\rho(t_c)$ which is simpler than the others. In fact, a large Q implies that the glitch is very narrow in the frequency domain. Thus, the integrand in eq. (5.3) is no longer dominated by the neighborhood of f_s , but by the narrow peak of the sine-Gaussian centered on f_0 . As before, the function $U(f)$ varies slowly with frequency and can be regarded as constant. We can also extend the integration limits to the full real axis because they are effectively already limited by the bandwidth of the glitch. Finally, we still approximate the phase of the template to second order, but this time around f_0 , which means that we also need the linear term:

$$\Psi(f; t_c) \approx \Psi(f_0) + \alpha(f - f_0) + \beta(f - f_0)^2. \quad (5.17)$$

For the Newtonian chirp the coefficients are

$$\alpha = 2\pi t_c - \frac{5}{128}f_0^{-8/3}\zeta^{-5/3} = 2\pi(t_c - \tau_0) \quad (5.18)$$

$$\beta = \frac{5}{96}f_0^{-11/3}\zeta^{-5/3}. \quad (5.19)$$

By switching to the new integration variable $x := \sqrt{\pi^2\tau^2 - i\beta}(f - f_0)$ and dropping global phase factors,

$$z(t_c) \approx 2\sqrt{\pi}h_0\tau(\pi^2\tau^2 - i\beta)^{-1/2}U(f_0) \int_{-\infty}^{+\infty} \exp\left(-x^2 + \frac{i\alpha x}{\sqrt{\pi^2\tau^2 - i\beta}}\right) dx \quad (5.20)$$

which is just a Gaussian integral and readily gives

$$z(t_c) \approx 2\pi h_0\tau(\pi^2\tau^2 - i\beta)^{-1/2}U(f_0) \exp\left(-\frac{\alpha^2}{4(\pi^2\tau^2 - i\beta)}\right) \quad (5.21)$$

and the SNR time series is

$$\rho(t_c) \approx \sigma_0^{-1}U(f_0) \left(1 + \frac{\beta^2}{\pi^4\tau^4}\right)^{-1/4} \exp\left(-\frac{\alpha^2}{4\pi^2\tau^2(1 + \frac{\beta^2}{\pi^4\tau^4})}\right). \quad (5.22)$$

Assuming a Newtonian chirp, α is the only function of t_c . The SNR time series is then a simple Gaussian pulse, with maximum given by

$$\rho^{\max} \approx \sigma_0^{-1}U(f_0) \left(1 + \frac{\beta^2}{\pi^4\tau^4}\right)^{-1/4} \quad (5.23)$$

$$t_c^{\max} \approx \tau_0. \quad (5.24)$$

We thus see that Approximation II taken in the limit $Q \gg 1$ did indeed give us the correct result for t_c^{\max} , but not quite the same ρ^{\max} . We expect Approximation III to be correct in this case.

5.3 Numerical simulations

At this point we have different estimations for ρ^{\max} and t_c^{\max} and we can visualize and compare them as a function of the main parameters involved: f_0 , Q and the mass of the Newtonian template. In order to check the accuracy of the approximations, we now simulate the response of the matched filter to a noiseless sine-Gaussian and find the maximum of the SNR numerically. Because we are concerned with detectors sensitive at low frequencies and long inspiral templates, we take advanced LIGO's final design sensitivity as the model for $S(f)$. For simplicity, and to avoid numerical issues, we set the amplitude of the glitch to $A = 1$ and multiply $S(f)$ by 10^{48} : thus, the SNRs we obtain refer to sine-Gaussians with amplitude $A = 10^{-24}$. We explore the glitch parameter ranges given by $10 \text{ Hz} \leq f_0 \leq \min(1 \text{ kHz}, f_{\text{ISCO}})$ and $1 \leq Q \leq 1000$. We test equal-mass templates with total mass $M = \{2.8M_\odot, 5M_\odot, 10M_\odot\}$.

Approximation I produces the results shown in figure 5.1. It correctly reproduces both ρ^{\max} and t_c^{\max} given by the simulations, as long as f_0 , Q and mass are sufficiently small. We find empirically that the accuracy decreases roughly as f_0^2QM and the 5% accuracy for ρ^{\max} is at $f_0^2QM/M_\odot \approx 5 \times 10^7 \text{ s}^{-2}$. This is a promising result because




Figure 5.1: Comparison between numerical simulations (color) and Approximation I (black contours). Left: trigger SNR; right: trigger delay; dashed lines: ISCO frequency. Approximation I is accurate almost everywhere, except at large f_0 , Q and mass.

we expect most glitches to last up to $\mathcal{O}(10)$ cycles; moreover, at high frequency or mass the time delay is small and thus the problem we are investigating is less of a concern. We also see in figure 5.1 that neglecting the exponential term in eq. (5.2) does not seem to have major effects at $Q \approx 1$, but it might be responsible for the wrong prediction of t_c^{\max} at $Q \approx 1$, $f_0 \approx f_{\text{ISCO}}$ and large mass.

Results for Approximation II are shown in figure 5.2. This approximation has excellent performance for t_c^{\max} across all the explored parameter space, remarkably even for large Q . However, as expected, it does not predict ρ^{\max} correctly for $Q \gg 1$. As Q increases, in fact, the sine-Gaussian peak in the frequency domain becomes smaller and smaller until the integral is no longer dominated by the region around f_s . Nevertheless, this is a useful approximation which works well in the most interesting regime, although it is not significantly simpler than Approximation I.

Approximation III's results are shown in figure 5.3. Consistently with its assumption, Approximation III works well in the high-mass, high- Q region where the other approximations begin to break down. It is also the one that gives the simplest result. However, at $Q \lesssim 10$ this approximation fails to reproduce the important dependence of t_c^{\max} on Q . Although Approximation III is not very useful for describing short glitches, it can be used to understand the effect of very long-lived and quasi-monochromatic disturbances; in particular it clearly shows how the triggers no longer depend on Q and the main factor determining ρ^{\max} is just the noise PSD.

Although we only consider Newtonian templates, search pipelines need to use higher-order pN terms to achieve a sufficient effectualness. Thus, we check next the accuracy of Approximation I against numerical simulations with a 3.5 pN template. We consider a single case with equal mass and $M = 5M_\odot$. The result is shown in figure 5.4. We find that the accuracy is worse but still within a few percent both for ρ^{\max} and t_c^{\max} . Moreover, the variation of the accuracy with the parameters is retained. A full exploration taking into account chirp mass, mass ratio and possibly aligned spin components could highlight other cases where our approximations break down, but we do not go into such detail here.

So far we considered the case of a single template. In general, a sufficiently strong glitch excites every template in the bank covering the search space, so a cluster of spurious triggers is produced. Our approximations allow us to model the shape of this cluster and in particular describe the correlation between the SNR and time of the triggers. As an example, figure 5.5 shows such a cluster when a simplified bank with uniform distribution of templates in M is assumed. It can be seen that low-mass templates produce the triggers with the largest delay and smallest SNR. Triggers from the lightest templates can be delayed by several minutes with respect to the glitch; such a delay is much longer than the duration of the glitch itself. Vetoing every trigger within such a long cluster leads to a large loss of live time. On the other hand, only vetoing triggers in the close vicinity of the glitch is not sufficient. The approximations introduced here can be used to regulate the vetoing procedure template-by-template, based on the estimates of the time delay and possibly SNR.




Figure 5.2: Comparison between numerical simulations (color) and Approximation II (black contours). Left: trigger SNR; right: trigger delay; dashed lines: ISCO frequency. This approximation predicts the delay correctly almost everywhere, but fails for the SNR at large Q , as expected.




Figure 5.3: Comparison between numerical simulations (color) and Approximation III (black contours). Left: trigger SNR; right: trigger delay; dashed lines: ISCO frequency. The approximation agrees with the numerical result at $Q \gg 1$, as expected.




Figure 5.4: Comparison between numerical simulations with a 3.5 pN template (color) and Approximation I (black contours). Left: trigger SNR; right: trigger delay; dashed lines: ISCO frequency. Despite neglecting post-Newtonian corrections, the approximation still gives a reasonable prediction of the SNR and delay in its region of validity.




Figure 5.5: Triggers generated by a sine-Gaussian glitch ($t_0 = 0$, $f_0 = 20$ Hz, $Q = 20$) exciting a simplified template bank of 50 3.5 pN equal-mass waveforms with M uniformly spaced between $2M_\odot$ and $20M_\odot$. The glitch produces a cluster of triggers lasting as long as the longest template, with high-mass triggers coming first and with the largest SNRs.

5.4 Conclusion

This chapter investigates the response of low-mass inspiral templates to short instrumental glitches commonly found in interferometer data. It is an initial step in understanding spurious inspiral triggers produced by short instrumental transients. By employing a simplified sine-Gaussian glitch model, parametrized by a central frequency and a quality factor, we are able to find three approximations to the matched-filter integral, which allow us to estimate the average SNR and the time of inspiral triggers produced by glitches. In other words, these approximations map the parameters describing the glitch and the template to the SNR and time of the resulting trigger. We test the approximations via numerical simulations of noiseless sine-Gaussian glitches and identify their region of validity. We find that the approximations are appropriate in different regions of the parameter space describing the glitches and the templates. If each approximation is used in its region of validity, the full space of physically interesting glitches and templates is covered effectively.

This study demonstrates how glitches in advanced interferometers can trigger long inspiral templates several minutes after the occurrence of the glitch. This represents a potential issue for future inspiral searches, because simply vetoing inspiral triggers in the vicinity of a short glitch will not be sufficient. On the other hand, vetoing every trigger within the duration of the template could lead to a significant reduction of the live time of the search. This requires a new way of dealing with short glitches. Assuming the parameters of the glitch can be estimated first (for instance via an event trigger generator such as Omega [189] or by analyzing auxiliary channels with known coupling to the strain signal [190]) the approximations introduced here can be used to calculate the expected time of the spurious trigger on a template-by-template basis. If indeed a trigger is found near the expected time, it can be either vetoed or down-ranked. An alternative approach to deal with delayed triggers is excising the short portion of the data affected by the glitch (*gating*) as part of the conditioning step prior to the matched filter. However, this must be done in a way that does not also introduce spurious triggers. We reserve the exploration of both approaches to a future study.

We neglect here the signal-based vetoes described in section 1.3.5. We have seen in chapter 2 that these are very effective at protecting against glitches. This is particularly true for long inspiral templates which are very different from sine-Gaussian waveforms. The problem described here may thus be strongly suppressed by the usage of ranking statistics such as eq. (1.38). Nevertheless, it is possible that the two loudest single-detector background events described in chapter 2 are the result of nearby loud glitches, despite the use of the χ^2 veto. A detailed investigation of the response of signal-based vetoes to sine-Gaussian waveforms is therefore a necessary follow-up of this study. In particular, it could indicate optimal choices for the parameters associated with the signal-based vetoes, such as the number of bins (and their boundaries) for the χ^2 test.

Bibliography

- [1] A. Weiss, W. Hillebrandt, H. C. Thomas, and H. Ritter, *Cox and Giuli's Principles of Stellar Structure*, 2nd ed. (Cambridge Scientific Publishers, 2004).
- [2] H. Spruit and E. S. Phinney, *Nature* **393**, 139 (1998), arXiv:astro-ph/9803201 [astro-ph] .
- [3] M. C. Miller and J. M. Miller, *Phys.Rept.* **548**, 1 (2014), arXiv:1408.4145 [astro-ph.HE] .
- [4] “ATNF pulsar catalogue,” <http://www.atnf.csiro.au/research/pulsar/psrcat/>.
- [5] A. Shearer and A. Golden, in *Neutron Stars, Pulsars, and Supernova Remnants*, edited by W. Becker, H. Lesch, and J. Trümper (2002) p. 44, astro-ph/0208579 .
- [6] P. A. Caraveo, *Ann.Rev.Astron.Astrophys.* **52**, 211 (2014), arXiv:1312.2913 [astro-ph.HE] .
- [7] N. Chamel, P. Haensel, J. Zdunik, and A. Fantina, *Int.J.Mod.Phys.* **E22**, 1330018 (2013), arXiv:1307.3995 [astro-ph.HE] .
- [8] “stellarcollapse.org,” <http://www.stellarcollapse.org>.
- [9] B. Kiziltan, A. Kottas, M. De Yoreo, and S. E. Thorsett, *Astrophys.J.* **778**, 66 (2013), arXiv:1309.6635 [astro-ph.SR] .
- [10] K.-W. Lo and L.-M. Lin, *Astrophys.J.* **728**, 12 (2011), arXiv:1011.3563 [astro-ph.HE] .
- [11] J. W. Hessels, S. M. Ransom, I. H. Stairs, P. C. C. Freire, V. M. Kaspi, *et al.*, *Science* **311**, 1901 (2006), arXiv:astro-ph/0601337 [astro-ph] .
- [12] D. Chakrabarty, E. H. Morgan, M. P. Muno, D. K. Galloway, R. Wijnands, *et al.*, *Nature* **424**, 42 (2003), arXiv:astro-ph/0307029 [astro-ph] .
- [13] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation* (W H Freeman and Company, 1973).

- [14] K. S. Thorne, *The Astrophysical Journal* **191**, 507 (1974).
- [15] J. A. Orosz, J. E. McClintock, J. P. Aufdenberg, R. A. Remillard, M. J. Reid, R. Narayan, and L. Gou, *The Astrophysical Journal* **742**, 84 (2011).
- [16] V. Bozza, *General Relativity and Gravitation* **42**, 2269 (2010).
- [17] A. M. Ghez, S. Salim, N. N. Weinberg, J. R. Lu, *et al.*, *The Astrophysical Journal* **689**, 1044 (2008).
- [18] S. Gillessen, F. Eisenhauer, S. Trippe, T. Alexander, R. Genzel, F. Martins, and T. Ott, *The Astrophysical Journal* **692**, 1075 (2009).
- [19] F. Özel, D. Psaltis, R. Narayan, and J. E. McClintock, *Astrophys.J.* **725**, 1918 (2010), arXiv:1006.2834 [astro-ph.GA] .
- [20] W. M. Farr, N. Sravan, A. Cantrell, L. Kreidberg, C. D. Bailyn, I. Mandel, and V. Kalogera, *The Astrophysical Journal* **741**, 103 (2011).
- [21] J. E. McClintock, R. Narayan, and J. F. Steiner, (2013), arXiv:1303.1583 [astro-ph.HE] .
- [22] P. C. Fragile, *The Astrophysical Journal* **706**, L246 (2009), arXiv:0910.5721 [astro-ph.HE] .
- [23] C. S. Reynolds, *Space Sci.Rev.* **183**, 277 (2014), arXiv:1302.3260 [astro-ph.HE] .
- [24] M. A. Abramowicz and W. Kluzniak, *Astron.Astrophys.* **374**, L19 (2001), arXiv:astro-ph/0105077 [astro-ph] .
- [25] R. Wijnands, M. van der Klis, J. Homan, D. Chakrabarty, C. B. Markwardt, *et al.*, *Nature* **424**, 44 (2003), arXiv:astro-ph/0307123 [astro-ph] .
- [26] R. W. O’Shaughnessy, J. Kaplan, V. Kalogera, and K. Belczynski, *Astrophys.J.* **632**, 1035 (2005), arXiv:astro-ph/0503219 [astro-ph] .
- [27] S. Yamada, K. Makishima, Y. Uehara, K. Nakazawa, H. Takahashi, T. Dotani, Y. Ueda, K. Ebisawa, A. Kubota, and P. Gandhi, *The Astrophysical Journal Letters* **707**, L109 (2009).
- [28] C. Bambi, *The European Physical Journal C* **75**, 22 (2015), 10.1140/epjc/s10052-014-3248-x.
- [29] L. Gou, J. E. McClintock, J. F. Steiner, R. Narayan, A. G. Cantrell, C. D. Bailyn, and J. A. Orosz, *The Astrophysical Journal Letters* **718**, L122 (2010).
- [30] L. Gou, J. E. McClintock, R. A. Remillard, J. F. Steiner, M. J. Reid, *et al.*, (2013), arXiv:1308.4760 [astro-ph.HE] .

- [31] W. R. Morningstar, J. M. Miller, R. C. Reis, and K. Ebisawa, *Astrophys.J.* **784**, L18 (2014), arXiv:1401.1794 [astro-ph.HE] .
- [32] W. R. Morningstar and J. M. Miller, *Astrophys.J.* **793**, L33 (2014), arXiv:1408.7028 [astro-ph.HE] .
- [33] J. Miller, M. Parker, F. Fuerst, M. Bachetti, F. Harrison, *et al.*, *Astrophys.J.* **775**, L45 (2013), arXiv:1308.4669 [astro-ph.HE] .
- [34] J. M. Miller, A. C. Fabian, R. Wijnands, C. S. Reynolds, *et al.*, *The Astrophysical Journal Letters* **570**, L69 (2002).
- [35] J. F. Steiner, J. E. McClintock, J. A. Orosz, R. A. Remillard, C. D. Bailyn, *et al.*, *Astrophys.J.* **793**, L29 (2014), arXiv:1402.0148 [astro-ph.HE] .
- [36] G. Duchêne and A. Kraus, *Annual Review of Astronomy and Astrophysics* **51**, 269 (2013), arXiv:1303.3028 [astro-ph.SR] .
- [37] K. A. Postnov and L. R. Yungelson, *Living Rev. Relativity* **17** (2014).
- [38] J. M. Weisberg, D. J. Nice, and J. H. Taylor, *The Astrophysical Journal* **722**, 1030 (2010), arXiv:1011.0718 [astro-ph.GA] .
- [39] M. Burgay, N. D’Amico, A. Possenti, R. Manchester, A. Lyne, *et al.*, *Nature* **426**, 531 (2003), arXiv:astro-ph/0312071 [astro-ph] .
- [40] M. Maggiore, *Gravitational Waves*, Vol. 1 (Oxford University Press, 2008).
- [41] P. C. Peters and J. Mathews, *Phys. Rev.* **131**, 435 (1963).
- [42] P. C. Peters, *Phys. Rev.* **136**, B1224 (1964).
- [43] J. Hermes, M. Kilic, W. R. Brown, D. Winget, C. A. Prieto, *et al.*, *The Astrophysical Journal Letters* **757**, L21 (2012), arXiv:1208.5051 [astro-ph.SR] .
- [44] L. Blanchet, *Living Rev. Relativity* **17** (2014).
- [45] C. Cutler and E. E. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
- [46] L. Blanchet, T. Damour, B. R. Iyer, C. M. Will, and A. G. Wiseman, *Phys. Rev. Lett.* **74**, 3515 (1995).
- [47] L. Blanchet, T. Damour, G. Esposito-Farèse, and B. R. Iyer, *Phys. Rev. Lett.* **93**, 091101 (2004).
- [48] K. Arun, A. Buonanno, G. Faye, and E. Ochsner, *Phys. Rev. D* **79**, 104023 (2009), arXiv:0810.5336 [gr-qc] .
- [49] L. E. Kidder, *Phys. Rev. D* **52**, 821 (1995).

- [50] L. Blanchet, A. Buonanno, and G. Faye, *Phys. Rev. D* **74**, 104034 (2006).
- [51] B. Mikóczi, M. Vasúth, and L. A. Gergely, *Phys. Rev. D* **71**, 124043 (2005).
- [52] S. Marsat, A. Bohe, G. Faye, and L. Blanchet, *Class. Quant. Grav.* **30**, 055007 (2013), arXiv:1210.4143 [gr-qc] .
- [53] E. Poisson, *Phys. Rev. D* **57**, 5287 (1998).
- [54] D. Lai and A. G. Wiseman, *Phys. Rev. D* **54**, 3958 (1996).
- [55] J. Vines, E. E. Flanagan, and T. Hinderer, *Phys. Rev. D* **83**, 084051 (2011).
- [56] J. E. Vines and E. E. Flanagan, *Phys. Rev. D* **88**, 024046 (2013).
- [57] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, *Phys. Rev. D* **49**, 6274 (1994).
- [58] A. H. Nitz, A. Lundgren, D. A. Brown, E. Ochsner, D. Keppel, and I. W. Harry, *Phys. Rev. D* **88**, 124039 (2013).
- [59] A. Lundgren and R. O’Shaughnessy, *Phys. Rev. D* **89**, 044021 (2014), arXiv:1304.3332 [gr-qc] .
- [60] “The LIGO algorithms library,” <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>.
- [61] M. Alcubierre, *Introduction to 3+1 Numerical Relativity* (Oxford University Press, UK, 2008).
- [62] L. Lehner and F. Pretorius, *Annual Review of Astronomy and Astrophysics* **52**, 661 (2014), arXiv:1405.4840 [astro-ph.HE] .
- [63] U. Sperhake, (2014), arXiv:1411.3997 [gr-qc] .
- [64] F. Löffler, J. Faber, E. Bentivegna, T. Bode, *et al.*, *Class. Quantum Grav.* **29**, 115001 (2012), arXiv:1111.3344 [gr-qc] .
- [65] “Spectral einstein code,” <http://www.black-holes.org/SpEC.html>.
- [66] E. Nakar, *Phys.Rept.* **442**, 166 (2007), arXiv:astro-ph/0701748 [ASTRO-PH] .
- [67] S. Goriely, A. Bauswein, and H.-T. Janka, *The Astrophysical Journal* **738**, L32 (2011), arXiv:1107.0899 [astro-ph.SR] .
- [68] O. Just, A. Bauswein, R. A. Pulpillo, S. Goriely, and H. T. Janka, (2014), arXiv:1406.2687 [astro-ph.SR] .
- [69] N. Tanvir, A. Levan, A. Fruchter, J. Hjorth, K. Wiersema, *et al.*, *Nature* **500**, 547 (2013), arXiv:1306.4971 [astro-ph.HE] .

- [70] L. Baiotti, B. Giacomazzo, and L. Rezzolla, *Phys. Rev. D* **78**, 084033 (2008).
- [71] Y. Sekiguchi, K. Kiuchi, K. Kyutoku, and M. Shibata, *Phys. Rev. Lett.* **107**, 051102 (2011).
- [72] K. Hotokezaka, K. Kyutoku, H. Okawa, M. Shibata, and K. Kiuchi, *Phys. Rev. D* **83**, 124008 (2011).
- [73] Y. Pan, A. Buonanno, M. Boyle, L. T. Buchman, L. E. Kidder, *et al.*, *Phys. Rev. D* **84**, 124052 (2011), arXiv:1106.1021 [gr-qc] .
- [74] A. Taracchini, A. Buonanno, Y. Pan, T. Hinderer, *et al.*, *Phys. Rev. D* **89**, 061502 (2014).
- [75] P. Ajith, S. Babak, Y. Chen, M. Hewitson, B. Krishnan, *et al.*, *Class.Quant.Grav.* **24**, S689 (2007), arXiv:0704.3764 [gr-qc] .
- [76] L. Santamaría, F. Ohme, P. Ajith, B. Brügmann, *et al.*, *Phys. Rev. D* **82**, 064016 (2010).
- [77] Y. Pan, A. Buonanno, A. Taracchini, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **89**, 084006 (2014).
- [78] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, *Phys. Rev. Lett.* **113**, 151101 (2014).
- [79] M. Dominik, K. Belczynski, C. Fryer, D. Holz, E. Berti, *et al.*, *The Astrophysical Journal* **759**, 52 (2012), arXiv:1202.4901 [astro-ph.HE] .
- [80] K. Belczynski, T. Bulik, and C. Bailyn, *The Astrophysical Journal* **742**, L2 (2011), arXiv:1107.4106 [astro-ph.GA] .
- [81] A. Bauswein, R. Ardevol Pulpillo, H. Janka, and S. Goriely, *The Astrophysical Journal* **795**, L9 (2014), arXiv:1408.1783 [astro-ph.SR] .
- [82] M. Dominik, K. Belczynski, C. Fryer, D. E. Holz, E. Berti, *et al.*, *The Astrophysical Journal* **779**, 72 (2013), arXiv:1308.1546 [astro-ph.HE] .
- [83] K. Belczynski and M. Dominik, (2012), arXiv:1208.0358 [astro-ph.HE] .
- [84] B. D. Lackey, K. Kyutoku, M. Shibata, P. R. Brady, and J. L. Friedman, *Phys. Rev. D* **85**, 044061 (2012).
- [85] T. Damour, A. Nagar, and L. Villain, *Phys. Rev. D* **85**, 123007 (2012).
- [86] B. F. Schutz, *Nature* **323**, 310 (1986).
- [87] W. Del Pozzo, *Phys. Rev. D* **86**, 043011 (2012).

- [88] C. K. Mishra, K. G. Arun, B. R. Iyer, and B. S. Sathyaprakash, *Phys. Rev. D* **82**, 064010 (2010).
- [89] M. Agathos, W. Del Pozzo, T. G. F. Li, C. Van Den Broeck, J. Veitch, and S. Vitale, *Phys. Rev. D* **89**, 082001 (2014).
- [90] J. Abadie *et al.*, *Classical and Quantum Gravity* **27**, 173001 (2010).
- [91] M. Dominik, E. Berti, R. O’Shaughnessy, I. Mandel, K. Belczynski, *et al.*, (2014), arXiv:1405.7016 [astro-ph.HE] .
- [92] P. Saulson, *Fundamentals of Interferometric Gravitational Wave Detectors* (World Scientific, 1994).
- [93] M. Pitkin, S. Reid, S. Rowan, and J. Hough, *Living Rev. Relativity* **14** (2011).
- [94] L. S. Finn, *Phys. Rev. D* **79**, 022002 (2009).
- [95] G. M. Harry (LIGO Scientific Collaboration), *Classical and Quantum Gravity* **27**, 084006 (2010).
- [96] F. Acernese *et al.* (The Virgo Collaboration), *Classical and Quantum Gravity* **32**, 024001 (2015), arXiv:1408.3978 [gr-qc] .
- [97] K. Somiya (KAGRA Collaboration), *Classical and Quantum Gravity* **29**, 124007 (2012), arXiv:1111.7185 [gr-qc] .
- [98] B. Iyer *et al.*, *LIGO-India, Proposal of the Consortium for Indian Initiative in Gravitational-wave Observations (IndIGO)*, Tech. Rep. LIGO-M1100296 (2011).
- [99] C. Affeldt, K. Danzmann, K. L. Dooley, H. Grote, *et al.*, *Classical and Quantum Gravity* **31**, 224002 (2014).
- [100] P. Amaro-Seoane, S. Aoudia, S. Babak, P. Binetruy, E. Berti, *et al.*, *GW Notes* **6**, 4 (2013), arXiv:1201.3621 [astro-ph.CO] .
- [101] G. Hobbs, (2010), arXiv:1006.3969 [astro-ph.SR] .
- [102] J. Weber, *Phys. Rev.* **117**, 306 (1960).
- [103] D. Percival and A. Walden, *Spectral Analysis for Physical Applications: Multitaper and Conventional Univariate Techniques* (Cambridge University Press, Cambridge, UK, 1993).
- [104] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, *Phys. Rev. D* **85**, 122006 (2012), arXiv:gr-qc/0509116 [gr-qc] .
- [105] J. Abadie *et al.* (Virgo Collaboration, LIGO Scientific Collaboration), (2012), arXiv:1203.2674 [gr-qc] .

- [106] D. Shoemaker, *Advanced LIGO anticipated sensitivity curves*, Tech. Rep. LIGO-T0900288-v3 (2010).
- [107] L. Barsotti and P. Fritschel, *Early aLIGO Configurations: example scenarios toward design sensitivity*, Tech. Rep. LIGO-T1200307-v4 (2014).
- [108] J. Slutsky, L. Blackburn, D. A. Brown, L. Cadonati, *et al.*, *Classical and Quantum Gravity* **27**, 165023 (2010).
- [109] N. Christensen, the LIGO Scientific Collaboration, and the Virgo Collaboration, *Classical and Quantum Gravity* **27**, 194010 (2010).
- [110] J. Aasi *et al.*, *Classical and Quantum Gravity* **29**, 155002 (2012).
- [111] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz, *Phys. Rev. D* **86**, 084017 (2012).
- [112] B. J. Owen, *Phys. Rev. D* **53**, 6749 (1996).
- [113] D. Keppel, A. P. Lundgren, B. J. Owen, and H. Zhu, *Phys. Rev. D* **88**, 063002 (2013).
- [114] T. A. Apostolatos, *Phys. Rev. D* **52**, 605 (1995).
- [115] T. Cokelaer, *Phys. Rev. D* **76**, 102004 (2007).
- [116] R. Prix, *Classical and Quantum Gravity* **24**, S481 (2007), arXiv:0707.0428 [gr-qc].
- [117] I. W. Harry, B. Allen, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 104014 (2009).
- [118] C. Messenger, R. Prix, and M. A. Papa, *Phys. Rev. D* **79**, 104017 (2009).
- [119] G. M. Manca and M. Vallisneri, *Phys. Rev. D* **81**, 024004 (2010).
- [120] B. Allen, *Phys. Rev. D* **71**, 062001 (2005), arXiv:gr-qc/0405045 [gr-qc].
- [121] S. Babak, R. Biswas, P. R. Brady, D. A. Brown, *et al.*, *Phys. Rev. D* **87**, 024033 (2013).
- [122] J. Abadie *et al.*, *Phys. Rev. D* **85**, 082002 (2012).
- [123] C. Hanna, *Searching for gravitational waves from binary systems in non-stationary data*, Ph.D. thesis, Louisiana State University (2008).
- [124] C. A. K. Robinson, B. S. Sathyaprakash, and A. S. Sengupta, *Phys. Rev. D* **78**, 062002 (2008).
- [125] S. Bose *et al.*, *Int. J. Mod. Phys. D* **09**, 325 (2000).
- [126] L. S. Finn, *Phys. Rev. D* **63**, 102001 (2001).

- [127] C. Cutler and B. F. Schutz, *Phys. Rev. D* **72**, 063006 (2005).
- [128] I. W. Harry and S. Fairhurst, *Phys. Rev. D* **83**, 084002 (2011).
- [129] K. Haris and A. Pai, *Phys. Rev. D* **90**, 022003 (2014).
- [130] M. Was, M.-A. Bizouard, V. Brisson, F. Cavalier, *et al.*, *Classical and Quantum Gravity* **27**, 015005 (2010), arXiv:0906.2120 [gr-qc] .
- [131] J. Veitch, V. Raymond, B. Farr, W. Farr, *et al.*, *Phys. Rev. D* **91**, 042003 (2015).
- [132] B. Abbott *et al.*, *Phys. Rev. D* **69**, 122001 (2004).
- [133] B. Abbott *et al.*, *Phys. Rev. D* **73**, 102002 (2006).
- [134] B. Abbott *et al.*, *Phys. Rev. D* **72**, 082001 (2005).
- [135] B. Abbott *et al.*, *Phys. Rev. D* **72**, 082002 (2005).
- [136] B. Abbott *et al.*, *Phys. Rev. D* **78**, 042002 (2008).
- [137] B. P. Abbott *et al.*, *Phys. Rev. D* **79**, 122001 (2009).
- [138] B. P. Abbott *et al.*, *Phys. Rev. D* **80**, 047101 (2009).
- [139] J. Abadie *et al.*, *Phys. Rev. D* **83**, 122005 (2011).
- [140] J. Abadie *et al.*, *Phys. Rev. D* **82**, 102001 (2010).
- [141] J. Aasi *et al.*, *Phys. Rev. D* **87**, 022002 (2013).
- [142] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *The Astrophysical Journal* **715**, 1453 (2010).
- [143] T. A. Apostolatos, *Phys. Rev. D* **54**, 2421 (1996).
- [144] C. Van Den Broeck, D. A. Brown, T. Cokelaer, I. Harry, G. Jones, *et al.*, *Phys. Rev. D* **80**, 024009 (2009), arXiv:0904.1715 [gr-qc] .
- [145] I. W. Harry, A. H. Nitz, D. A. Brown, A. P. Lundgren, E. Ochsner, and D. Keppel, *Phys. Rev. D* **89**, 024010 (2014).
- [146] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, N. Mazumder, and A. J. Weinstein, *Phys. Rev. D* **89**, 084041 (2014).
- [147] T. Dal Canton, A. H. Nitz, A. P. Lundgren, A. B. Nielsen, *et al.*, *Phys. Rev. D* **90**, 082004 (2014), arXiv:1405.6731 [gr-qc] .
- [148] P. Grandclément, M. Ihm, V. Kalogera, and K. Belczynski, *Phys. Rev. D* **69**, 102002 (2004), arXiv:gr-qc/0312084 [gr-qc] .

- [149] K. Belczynski, R. E. Taam, E. Rantsiou, and M. van der Sluys, *The Astrophysical Journal* **682**, 474 (2008), arXiv:astro-ph/0703131 [ASTRO-PH] .
- [150] T. Fragos, M. Tremmel, E. Rantsiou, and K. Belczynski, *The Astrophysical Journal Letters* **719**, L79 (2010), arXiv:1001.1107 [astro-ph.HE] .
- [151] F. Foucart, *Phys. Rev. D* **86**, 124007 (2012), arXiv:1207.6304 [astro-ph.HE] .
- [152] B. D. Lackey, K. Kyutoku, M. Shibata, P. R. Brady, and J. L. Friedman, *Phys. Rev. D* **89**, 043009 (2014).
- [153] F. Pannarale, E. Berti, K. Kyutoku, and M. Shibata, *Phys. Rev. D* **88**, 084011 (2013).
- [154] F. Pannarale, L. Rezzolla, F. Ohme, and J. S. Read, *Phys. Rev. D* **84**, 104017 (2011), arXiv:1103.3526 [astro-ph.HE] .
- [155] M. Cabero, A. P. Lundgren, and A. B. Nielsen, (in preparation).
- [156] J. Aasi *et al.*, *Classical and Quantum Gravity* **31**, 115004 (2014).
- [157] M. Frigo and S. G. Johnson, *Proceedings of the IEEE* **93**, 216 (2005), special issue on “Program Generation, Optimization, and Platform Adaptation”.
- [158] B. Allen, B. Knispel, J. M. Cordes, J. S. Deneva, *et al.*, *The Astrophysical Journal* **773**, 91 (2013).
- [159] E. Baird, S. Fairhurst, M. Hannam, and P. Murphy, *Phys. Rev. D* **87**, 024035 (2013).
- [160] A. S. Sengupta, S. Dhurandhar, and A. Lazzarini, *Phys. Rev. D* **67**, 082004 (2003), arXiv:gr-qc/0301025 [gr-qc] .
- [161] C. Aulbert and H. Fehrmann, *Max-Planck-Gesellschaft Jahrbuch 2009* (2009).
- [162] “PostgreSQL,” <http://www.postgresql.org>.
- [163] “MongoDB,” <http://www.mongodb.org>.
- [164] “Apache Cassandra,” <http://cassandra.apache.org>.
- [165] C. W. Helstrom, *Statistical theory of signal detection* (New York, Pergamon Press, 1968).
- [166] V. Kalogera, *Astrophys.J.* **541**, 319 (2000), arXiv:astro-ph/9911417 [astro-ph] .
- [167] A. Buonanno, Y.-b. Chen, and M. Vallisneri, *Phys. Rev. D* **67**, 104025 (2003), arXiv:gr-qc/0211087 [gr-qc] .

- [168] D. Fazi, *Development of a physical-template search for gravitational waves from spinning compact-object binaries with LIGO*, Ph.D. thesis, Università di Bologna (2009).
- [169] I. Harry and S. Fairhurst, *Class. Quant. Grav.* **28**, 134008 (2011), arXiv:1101.1459 [gr-qc] .
- [170] T. Dal Canton, A. P. Lundgren, and A. B. Nielsen, *Phys. Rev. D* **91**, 062010 (2015), arXiv:1411.6815 [gr-qc] .
- [171] P. Schmidt, F. Ohme, and M. Hannam, *Phys. Rev. D* **91**, 024043 (2015), arXiv:1408.1810 [gr-qc] .
- [172] D. A. Brown, A. Lundgren, and R. O’Shaughnessy, *Phys. Rev. D* **86**, 064020 (2012), arXiv:1203.6060 [gr-qc] .
- [173] H. Mukhopadhyay *et al.*, *Phys. Rev. D* **74**, 083005 (2006), arXiv:gr-qc/0608103 [gr-qc] .
- [174] H. Mukhopadhyay *et al.*, *Phys. Rev. D* **80**, 123019 (2009), arXiv:0910.4302 [gr-qc] .
- [175] A. Pai *et al.*, *Phys. Rev. D* **64**, 042004 (2001), arXiv:gr-qc/0009078 .
- [176] A. Pai *et al.*, *Class. Quant. Grav.* **19**, 1477 (2002), arXiv:gr-qc/0110041 [gr-qc] .
- [177] S. Bose *et al.*, *Class. Quant. Grav.* **28**, 134009 (2011), arXiv:1104.2650 [astro-ph.IM] .
- [178] B. F. Schutz, *Class. Quant. Grav.* **28**, 125023 (2011), arXiv:1102.5421 .
- [179] “KAGRA location,” <http://gwdoc.icrr.u-tokyo.ac.jp/cgi-bin/DocDB/ShowDocument?docid=1699>.
- [180] J. Aasi *et al.* (LIGO Scientific, VIRGO), (2013), arXiv:1304.0670 [gr-qc] .
- [181] D. Keppel, (2013), arXiv:1307.4158 [gr-qc] .
- [182] D. Keppel, *Phys. Rev. D* **86**, 123010 (2012).
- [183] L. Wen, X. Fan, and Y. Chen, *Journal of Physics: Conference Series* **122**, 012038 (2008).
- [184] S. Fairhurst, *Class. Quant. Grav.* **28**, 105021 (2011), arXiv:1010.6192 [gr-qc] .
- [185] R. Prix, *Phys. Rev. D* **75**, 023004 (2007).
- [186] J. D. E. Creighton, *Phys. Rev. D* **60**, 021101 (1999), arXiv:gr-qc/9901075 [gr-qc] .
- [187] T. Dal Canton, S. Bhagwat, S. V. Dhurandhar, and A. Lundgren, *Classical and Quantum Gravity* **31**, 015016 (2014), arXiv:1304.0008 [gr-qc] .

- [188] L. Blackburn, L. Cadonati, S. Caride, S. Caudill, S. Chatterji, *et al.*, *Class. Quant. Grav.* **25**, 184004 (2008), arXiv:0804.0800 [gr-qc] .
- [189] S. K. Chatterji, *The search for gravitational wave bursts in data from the second LIGO science run*, Ph.D. thesis, Massachusetts Institute of Technology. Dept. of Physics (2005).
- [190] P. Ajith, M. Hewitson, J. R. Smith, H. Grote, S. Hild, and K. A. Strain, *Phys. Rev. D* **76**, 042004 (2007).

Acknowledgments

First of all, I am deeply grateful to Bruce Allen and Badri Krishnan for giving me the opportunity to work in such a stimulating and active environment as the AEI Hannover and for their guidance during the development of my projects. I also thank other people at the AEI and in the LIGO collaboration, in particular: Swetha Bhagwat, Sukanta Bose, Duncan Brown, Collin Capano, Gergely Debreczeni, Tom Dent, Sanjeev Dhurandhar, Steve Fairhurst, Shaon Ghosh, Gianluca Guidi, Ian Harry, David Keitel, Drew Keppel, Andy Lundgren, Lars Nieder, Alex Nielsen, Alex Nitz, Frank Ohme, Steve Privitera, Reinhard Prix, Karsten Wiesner and Josh Willis. Discussing with them and receiving their comments has been very useful for my projects and the development of this thesis. I thank Carsten Aulbert, Oliver Bock and Henning Fehrmann for their help with the Atlas cluster, its GPUs, MongoDB and Cassandra.

I am also grateful to all the people I met at the AEI in the last few years for the chats, the trips and the nice time we had together, which made my experience at the AEI really pleasant: Alex, Andy, Avi, Colin, David, Evan, Francesco, Gian Mario, Giulio, Holger, Heinz-Bernd, Irene, Jing, Karl, Lars, Miriam, Miroslav, Nat, Natalia, Pablo, Paola, Tom and the several visitors we had during my time here.

I thank all the friends who have been an important part of my Hannoverian life in the last years, in particular Giulio and Pablo, who gave a big contribution to my interest in moving here, and Sasan for pushing me to practice German. I deeply thank my wife, my brother and sister and my parents for their constant presence, love, advice and support in all my adventures.

This thesis has been developed entirely on Debian GNU/Linux systems.

Curriculum vitae

Tito Dal Canton

Date of birth: February 7, 1983

Place of birth: Asolo (TV), Italy

2012-02 — present

Member of the LIGO scientific collaboration

2012-01 — present

PhD student position

Max-Planck-Institut für Gravitationsphysik

Leibniz Universität Hannover

2009-04 — 2011-12

Senior research worker position

New York State Psychiatric Institute

Columbia University Medical Center

2005-10 — 2008-07

Master of science (laurea specialistica) in physics

Università degli Studi di Padova

Thesis work at Laboratori Nazionali di Legnaro, INFN

2002-10 — 2005-07

Bachelor of science (laurea triennale) in physics

Università degli Studi di Padova

Thesis work at Laboratori Nazionali di Legnaro, INFN

1997-09 — 2002-09

High school (liceo scientifico)

Liceo Ginnasio Statale G. B. Brocchi, Bassano del Grappa

List of publications

T. Dal Canton, A. P. Lundgren and A. B. Nielsen

Impact of precession on aligned-spin searches for neutron-star-black-hole binaries

Physical Review D **91** 062010 (2015)

T. Dal Canton, A. H. Nitz, A. P. Lundgren, A. B. Nielsen et al.

Implementing a search for aligned-spin neutron star-black hole systems with advanced ground based gravitational wave detectors

Physical Review D **90** 082004 (2014)

T. Dal Canton, S. Bhagwat, S. Dhurandhar and A. P. Lundgren

Effect of sine-Gaussian glitches on searches for binary coalescence

Classical and Quantum Gravity **31** 015016 (2014)

M. Ciszak, F. Marino, A. Ortolan and T. Dal Canton

Identification of gravitational wave signals from chaotic astrophysical systems through phase space and attractor properties

Physical Review D **80**, 043002 (2009)

P. Leaci, T. Dal Canton, A. Ortolan and G. A. Prodi

Antenna pattern of DUAL detectors of gravitational waves and its exploitation in a network of advanced interferometers

Classical and Quantum Gravity **25** 195018 (2008)

(Full-collaboration LIGO papers are not listed)