

Fast identification of continuous gravitational-wave signals

Iuri La Rosa

Thesis advisors:

Paola Leaci

Tania Regimbau

Department of Physics
Sapienza University of Rome
Université Savoie Mont Blanc
September 2022

Abstract

Continuous gravitational waves from non-axisymmetric spinning neutron stars are among the most interesting, although still undetected, targets of the Advanced LIGO-Virgo and KAGRA detectors. The search for this class of signals is difficult due to their expected weakness, and can be very computationally expensive when the source parameters are not known.

Stochastic gravitational-wave background searches use fast and consolidated cross-correlation techniques to search for either a stochastic background of gravitational waves or persistent gravitational waves in specific sky directions. Recent investigations have shown that stochastic directional searches have the ability to detect continuous waves as well, with less sensitivity than continuous waves searches, but with low computing requirements.

Using the algorithms from both the specific continuous-wave searches and the stochastic directional ones, in this thesis I present an approach that aims to improve the efficiency of the data analysis applied to the continuous-wave field.

Using the high parallel computing efficiency and computational power of the modern Graphics Processing Units (GPUs), I show how an all-sky search algorithm has been ported to GPUs. This work brought to a speed-up of all-sky searches for periodic signals by an order of magnitude, in terms of single device computational time. In addition, I describe a joint stochastic-background and continuous-wave pipeline chain. I managed to combine the robustness of the stochastic cross-correlation searches, used to quickly identify continuous-wave signals, to the accuracy of the continuous waves dedicated codes to properly follow up interesting candidates.

Contents

Introduction	vii
1 Sources of gravitational waves	1
1.1 The theory of gravitational waves	1
1.1.1 Geodesic equations and the curvature tensor	2
1.1.2 The Einstein's equations and their perturbative solutions . .	3
1.2 Compact objects	7
1.2.1 Degenerate fermions	8
1.2.2 Hydrostatic equilibrium and mass limit	9
1.2.3 Neutron stars	11
1.3 Sources of gravitational waves	17
1.3.1 Quadrupole formalism	17
1.3.2 Neutron star emission mechanisms	19
1.3.3 Other gravitational-wave sources	23
1.4 Detectors	25
1.4.1 Michelson configuration	25
1.4.2 Fabry-Perot Cavity	29
1.4.3 Main noise sources	30
2 Data analysis for Continuous Waves	31
2.1 General methods for data analysis	32
2.1.1 Filtering out noise	34
2.1.2 Searching for periodic signals	35
2.2 The FrequencyHough pipeline	37
2.2.1 The sky grid generation	45
2.2.2 The FrequencyHough transform	48
2.2.3 Search sensitivity	50
2.3 The GPGPU project	54
2.3.1 GPGPU with Tensorflow	56
2.3.2 GPU FrequencyHough	58
3 Continuous-wave searches with radiometer	67
3.1 The stochastic background search	68
3.1.1 Isotropic optimal filter	68
3.1.2 The radiometer method	71
3.1.3 Overlap reduction function	72
3.1.4 The sky map	74

3.1.5	Implementations of the radiometer method	76
3.1.6	Standard follow-up and upper limits	78
3.1.7	All-sky search	82
3.2	The fast identification project	83
3.2.1	Continuous-wave radiometer characterization	84
3.2.2	Hardware injection analysis with O2 and O3 data	89
3.2.3	Binaries	92
3.2.4	The radiometer continuous-wave follow-up	94
Conclusions		97
Bibliography		101

List of Figures

1.1	Gravitational-waves geodesic deviation.	7
1.2	Equilibrium configuration of degenerate fermions.	10
1.3	Comparison of equations of state.	14
1.4	Pulsar scheme.	15
1.5	$P-\dot{P}$ plot.	16
1.6	Rotating ellipsoid scheme.	19
1.7	Supernova modeled waveforms.	24
1.8	Chirp waveform and signal detected.	24
1.9	Michelson interferometer scheme.	26
1.10	Polarization and detector unit vectors.	27
1.11	Antenna pattern.	28
1.12	Simplified scheme of gravitational interferometer.	29
1.13	Sensitivity curves	30
2.1	Flowchart summarizing the FrequencyHough pipeline.	38
2.2	Scheme of the SFDB construction.	39
2.3	Example peakmap.	40
2.4	Example FrequencyHough map.	42
2.5	Example candidate selection.	44
2.6	Number of sky points for the all-sky analysis.	46
2.7	All-sky grid example.	47
2.8	Example Hough map.	48
2.9	Peakmap selection threshold.	51
2.10	CPU-GPU FLOPS comparison.	55
2.11	Joint computing with CPU and GPU.	56
2.12	Example TensorFlow code and its graph.	57
2.13	CPU-GPU comparison.	58
2.14	CPU-GPU cost-computing time comparison.	59
2.15	Hardware injections.	61
2.16	FrequencyHough CPU-GPU comparison.	62
2.17	FrequencyHough input-parameter spaces benchmarks.	63
2.18	Estimated time for O3 FrequencyHough production.	64
3.1	Scheme of the radiometer between two detectors.	71
3.2	Example of the overlap reduction function for the H,L,V detectors.	73
3.3	Scheme of the data folding process.	76
3.4	Narrowband radiometer search scheme	77

3.5	SNR plots from the O3 narrowband search	80
3.6	Example of HEALpix grids with various resolutions	82
3.7	HEALpix versus FrequencyHough grid comparison	83
3.8	Detection statistics comparison	86
3.9	Stochastic Narrowband.	87
3.10	Reproducibility	88
3.11	Analysis of a monochromatic disturbance	88
3.12	All-sky maps for software injections.	89
3.13	O2 and O3 hardware injections summary.	91
3.14	O3 checks summary	91
3.15	Binary spectrum example.	92
3.16	Test on simulation of binaries.	93
3.17	Example of FrequencyHough grid for PyStoch outliers	95
3.18	Follow up of a software injection.	96

Introduction

Gravitational-wave physics started its centenary story with the Theory of General Relativity by Albert Einstein (Einstein, 1918), in the second decade of the 20th century. Gravitational waves are perturbations of the space-time metric traveling at the speed of light. When they cross a region of space, their effect is to shrink the separation between two points, while relaxing the orthogonal direction, similarly to what happens when we compress a ring made of a flexible material. In the case of gravitational waves, the action is on the metric itself.

This effect is very tiny and, once theorized, for many years it has been impossible to have direct or indirect observation of their existence. The first breakthrough came from Russell A. Hulse and Joseph H. Taylor (Hulse et al., 1975; Taylor et al., 1976) which, at the turn of the mid-1970s, discovered the first binary system made by two neutron stars (one of them being a pulsar). They found that their orbit was decaying with a rate higher than their electromagnetic energy loss, compatible to the prediction of a gravitational-wave emission. This discovery granted them the Nobel prize in 1993.

In the meanwhile, to try to directly measure gravitational waves, between the late 1950s and the 1970s, the first gravitational-wave detectors were conceived as resonant solid-state antennas (Weber, 1960) bringing to the end of 1990s to the first network of gravitational-wave detectors under the IGEC collaboration (Astone, Babusci, et al., 2003). In present day, the search for gravitational waves revolves around ground based interferometric detectors. They are designed as Michelson interferometers with 3-4 km long arms, and, in the second decade of the 2000s, thanks to them, the first transient gravitational-wave signal has been detected. Thus, the gravitational-astronomy era began, precisely on September 14, 2015 when the first ever detection has been recorded (i.e. GW150914, B. P. Abbott et al., 2016) by the LIGO-Virgo collaboration during the first observational run O1 (Aasi et al., 2016). The kind of observed event was a coalescence of two black holes which, orbiting together, lost enough energy by gravitational emission to finally merge. The signature of the event was the waveform referred to as *chirp*.

In 2017, during the second observational run (O2) (B. P. Abbott et al., 2020a), two more detections represented new milestones in the gravitational-wave observations: the first signal seen by three detectors (the two of the LIGO collaboration in the USA and Virgo in Italy), still a coalescence of two black holes (GW170814, B. P. Abbott et al., 2017d), and the first merger of two neutron stars with an electromagnetic counterpart (GW170817, B. P. Abbott et al., 2017d). Hence, for the first time, a real signal could be studied both from the electromagnetic and gravitational point of view, giving birth to the multi-messenger astronomy.

With the third observational run (O3) (B. P. Abbott et al., 2019a, 2020b), the LIGO and Virgo collaborations released the third catalog of gravitational-wave detections (R. Abbott et al., 2021b), reaching 90 events and welcoming the fourth detector of the network, i.e. the KAGRA detector, which is located in Japan. With four operating gravitational-wave detectors we can have a very promising fourth observational run (O4) (B. P. Abbott et al., 2019a, 2020b).

However, several other gravitational-wave sources and emission mechanisms exist (Maggiore, 2018), which have yet to be detected due to their extremely faint expected amplitude and/or rarity of events. Few examples are bursts from supernova explosions, persistent signals coming from rotating neutron stars, and from a superposition of cosmological and astronomical phenomena forming a stochastic background of gravitational waves. Thanks to the detection of all these kinds of signals, we could probe large portions of the universe and shed light on a variety of phenomena in a way that has never been done before, potentially discovering also new physics scenarios, especially in view of the next generation of gravitational-wave detectors, such as Einstein Telescope (Michele Maggiore et al., 2020), Cosmic Explorer (Evans et al., 2021) and LISA (Danzmann et al., 2017).

In this PhD Thesis, I will present the work on data analysis searches for continuous gravitational waves coming from spinning neutron stars. In particular, from population studies, we expect in our galaxy approximately an order of a billion of unseen neutron stars (Bisnovatyi-Kogan, 1992). Hence, all-sky searches spanning the whole sky are of crucial importance. The current PhD project aimed at improving a few tools for this class of searches, acting on two different aspects.

First, the all-sky FrequencyHough continuous-wave search algorithm (Astone et al., 2014) has been improved, adapting it to the parallel computing efficiency of the modern Graphical Processing Units (GPUs), disclosing the outcomes in La Rosa et al., 2021. Afterwards, I have deployed the new version of the algorithm on a GPU cluster, and this optimized algorithm has been used in the LIGO-Virgo collaboration analysis of the O3 run data R. Abbott et al., 2022a, spanning almost one year between April 2019 and March 2020. Thanks to this work, I was able to dramatically reduce, up to one order of magnitude, the computation time for the all-sky search, thus giving a significant contribution to the analysis.

A further aspect has been to use and improve the stochastic gravitational-wave background search algorithm for continuous-wave signals. The analysis pipeline is based on the radiometer method (S W Ballmer, 2006; Mitra et al., 2008) to search for anisotropic signal distributions over the sky, and it has been used also to search for signals from point-like sources, such as neutron stars (B. P. Abbott et al., 2017b, 2019c; R. Abbott et al., 2021d, 2022c). The algorithm performs very efficiently for unmodeled analyses over a wide parameters space, but it cannot reach the sensitivity of dedicated continuous-waves analyses. My project consisted of a full characterization of the radiometer analysis applied to continuous waves to test its behavior in different situations (possible interference among several signals or with noise, isolated sources or located in binary systems, etc). Then, the adapted GPU version of the FrequencyHough algorithm has been used to rapidly follow-up outliers coming from the radiometer analysis, building a chain between the two pipelines which improves the sensitivity and robustness of the search for continuous waves using the radiometer method (scientific paper in development).

The Thesis is organized as follows:

- In chapter 1** the fundamentals of the theory of gravitational waves are summarized. The linear wave equation coming from the perturbative approach and its effects on the matter will be explained. The sources of major interest for this work, the neutron stars, are presented together with their structure, physical behavior and main gravitational-wave emission mechanisms. A brief description of the gravitational-wave interferometers will be also given.
- In chapter 2** we will discuss the used methods and algorithms to perform especially all-sky searches for periodic gravitational-wave signals. After a general introduction on few main features of the gravitational-wave analysis, the FrequencyHough pipeline will be presented in details. At last, the porting on GPU of the algorithm will be shown with the main achieved outcomes.
- In chapter 3** the use of the radiometer stochastic background algorithm for a fast identification of continuous-wave sources will be presented. After a detailed explanation on the search for stochastic gravitational-wave background, it will be explained how the analysis algorithms have been used for continuous-wave searches. It will follow the description of the in-depth characterization study and the developed follow-up procedure using the continuous-wave tools.

Chapter 1

Sources of gravitational waves

Gravitational-waves are space-time perturbations generated by non-stationary gravitational fields (Einstein, 1918). In general, the propagating field equation is difficult to solve, mainly due by the non-linearity of the gravitational field (Ferrari et al., 2020), but assuming that it can be considered as a small perturbation, an approximate solution can be derived as the equation of a wave traveling at the speed of light.

As we will see, unlike electromagnetic waves which come from dipolar oscillations, gravitational waves have a quadrupolar nature and the quadrupole oscillation of the sources are strictly related to their rotational asymmetry. Notably, the sources that are of major interest for this work are non-axisymmetric rotating neutron stars. Modeling them as ellipsoids rotating around one of their axis, the gravitational-wave emission from them is a continuous signal, with amplitude depending on their ellipticity, other than their moment of inertia (Ferrari et al., 2020). The ellipticity and the moment of inertia (that is, radius and mass) are in turn closely related to the internal structure of neutron stars. Hence, the detection of this class of gravitational-wave signal would reveal a lot of information about the physics of neutron stars.

In this chapter, after reviewing the basics of the General Relativity theory, the physics of the neutron stars will be described. The main gravitational-wave emission mechanisms will then be discussed, with a particular accent on non-axisymmetric rotating neutron stars. Other emission mechanisms will be summarized as well as other classes of sources of particular interest. The chapter will end with an overview of the basic principles of the interferometric detectors and their main noise sources.

1.1 The theory of gravitational waves

In the framework of General Relativity, the gravitational field is described as the metric tensor of a *curved* space-time geometry, usually referred to as $g_{\mu\nu}$, where $\mu, \nu = 0, 1, 2, 3$ indicate respectively the time and the three spatial dimensions. In absence of a gravitational field, the space-time is described by the *flat* Minkowski metric $\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1)$.

To derive the equation of the propagating wave, a perturbation method is used to solve the Einstein's equations, where the gravitational field is considered a small

perturbation $h_{\mu\nu}$ of the flat metric: $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$ (Ferrari et al., 2020).

1.1.1 Geodesic equations and the curvature tensor

In order to describe how a gravitational field acts on the motion of a particle, as seen by a frame different than the *locally inertial* one (i.e., the frame where the distance is always defined by the Minkowski metric, which can be always defined as states the *equivalence principle*), we need the equations of motion. In the locally inertial frame, the distance is defined as $ds^2 = \eta_{\mu\nu} d\bar{x}^\mu d\bar{x}^\nu$ ¹. Holding the flat metrics and hence the Newtonian mechanics, the equations of motion of a particle will be $d^2\bar{x}^\mu/d\tau^2 = 0$, where τ is the particle proper time.

Moving to another frame means that the coordinates x_μ will be expressed by a transformation $x_\mu = x^\mu(\bar{x}^\mu)$. Applying this transformation in the distance definition, and developing the calculations following the normal rules for differentials and derivatives, we will have the new distance (Ferrari et al., 2020)

$$ds^2 = \eta_{\alpha\beta} \frac{\partial \bar{x}^\alpha}{\partial x^\mu} dx^\mu \frac{\partial \bar{x}^\beta}{\partial x^\nu} dx^\nu = g_{\mu\nu} dx^\mu dx^\nu, \quad (1.1)$$

where the definition of the new metric tensor is then

$$g_{\mu\nu} = \frac{\partial \bar{x}^\alpha}{\partial x^\mu} \frac{\partial \bar{x}^\beta}{\partial x^\nu} \eta_{\alpha\beta}. \quad (1.2)$$

Following the same process for the equations of motion, we obtain the so-called *geodesic equations*, which describe the motion of a particle as seen by a generic frame:

$$\frac{d^2 x^\alpha}{d\tau^2} + \Gamma^\alpha_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} = 0. \quad (1.3)$$

The solutions $x^\mu(\tau)$ of these equations are the geodesics, i.e. trajectories in space-time. The key element in these equations is the second term on the left side, proportional to the elements $\Gamma^\alpha_{\mu\nu}$. They are called *affine connections* and contain the information of the space-time manifold curved by the gravitational field, describing how the geometry changes in a generic frame. It can be shown that $\Gamma^\alpha_{\mu\nu}$ has an explicit dependence on the metric tensor $g_{\mu\nu}$ in the formula

$$\Gamma^\alpha_{\mu\nu} = \frac{g^{\sigma\alpha}}{2} (\partial_\nu g_{\sigma\mu} + \partial_\mu g_{\nu\sigma} - \partial_\sigma g_{\mu\nu}), \quad (1.4)$$

and it contains only the first derivatives of the metric tensors. This is not enough to express in the most general way the information on the gravitational field. In fact, because of the equivalence principle, we can always define a locally inertial frame, where $g_{\mu\nu} \rightarrow \eta_{\mu\nu}$ and $\Gamma^\alpha_{\mu\nu} = 0$, not allowing to describe anymore the effect of the curved metric.

The tensor that mathematically describes in any reference frame the intrinsic curvature of a manifold given by a gravitational field is the curvature tensor, also called Riemann Tensor, defined as:

¹Here \bar{x} stands for the coordinates in the locally inertial frames, in order to distinguish them from the generic ones.

$$R^\alpha{}_{\beta\mu\nu} := \partial_\mu \Gamma^\alpha{}_{\beta\nu} - \partial_\nu \Gamma^\alpha{}_{\beta\mu} - \Gamma^\alpha{}_{\sigma\nu} \Gamma^\sigma{}_{\beta\mu} + \Gamma^\alpha{}_{\sigma\mu} \Gamma^\sigma{}_{\beta\nu}. \quad (1.5)$$

This tensor has several important properties, among which the most important is that, in a locally inertial frame it becomes

$$R^\alpha{}_{\beta\mu\nu} = \frac{1}{2} g^{\alpha\sigma} (\partial_\beta \partial_\mu g_{\sigma\nu} - \partial_\beta \partial_\nu g_{\sigma\mu} + \partial_\sigma \partial_\nu g_{\beta\mu} - \partial_\sigma \partial_\mu g_{\beta\nu}). \quad (1.6)$$

This means that even though the affine connections are 0 in the locally inertial frames, their derivatives are not, causing that the Riemann tensor gives information about the gravitational field even in this case, through the second derivatives of the curved metric².

Usually, aside the Riemann tensor $R^\mu{}_{\nu\alpha\beta}$, its contraction $R_{\mu\nu} := R^\alpha{}_{\nu\alpha\beta}$, called Ricci tensor, and the scalar $R = g^{\mu\nu} R_{\mu\nu}$, are used to write the Einstein field equations (Ferrari et al., 2020).

1.1.2 The Einstein's equations and their perturbative solutions

Unlike the electromagnetic field, the gravitational field is not linear. For example, an electromagnetic wave comes from oscillations of charged particles, but the carried four-momentum cannot be itself the source of the field. For the gravitational field, because the mass-energy relation in the Theory of General Relativity, both a mass distribution and an energy density can be sources of a gravitational field. This implies that the energy carried by gravitational waves has to appear as well in the field equations as a source. Hence, they cannot be linear.

While for the Newtonian gravitational field we can use the mass density ρ in the field equation $\partial_i g_i = -4\pi G \rho$, for a relativistic approach we need to use the stress-energy tensor $T_{\mu\nu}$ to describe an arbitrary energy-matter distribution, in an equation that has to reduce to the Newtonian case when $g_{\mu\nu} \rightarrow \eta_{\mu\nu}$. Then, the equation will have the following shape:

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}, \quad (1.7)$$

where G is the gravitation constant and c the speed of light in vacuum.

The tensor $G_{\mu\nu}$ is the Einstein tensor, defined as

$$G_{\mu\nu} = R_{\mu\nu} - \frac{g_{\mu\nu}}{2} R. \quad (1.8)$$

The Einstein's equations can also be written in terms of the Ricci tensor:

$$R_{\mu\nu} = \frac{8\pi G}{c^4} \left(T_{\mu\nu} - \frac{g_{\mu\nu}}{2} T \right), \quad (1.9)$$

²It can be useful to remark that in a flat space-time, $R^\mu{}_{\nu\alpha\beta} = 0$ in any reference frame.

³In general, $G_{\mu\nu} = R_{\mu\nu} - \frac{g_{\mu\nu}}{2} R + \lambda g_{\mu\nu}$, with the additional term proportional to the so called *cosmological constant* $\lambda \approx 1.11 \times 10^{-52} m^{-2}$. Being λ very small, it can be neglected since it does not play a role if not in cosmological scales, which are not treated in this thesis focused on astrophysical phenomena.

where the scalar $T = T^\rho{}_\rho$ is the trace of $T_{\mu\nu}$ ⁴

The Einstein's equations have the property to be *gauge invariant*, i.e. symmetric for an arbitrary coordinate transformation $x^{\mu'} = x^{\mu'}$ (x^μ).

Now let us consider a (constant) exact solution of Einstein's equations $\bar{g}_{\mu\nu}$, and let us assume that, at least in a certain reference frame, it exists a small perturbation $h_{\mu\nu}$ of the background metric. If, for instance, we choose $\bar{g}_{\mu\nu} = \eta_{\mu\nu}$, then we have that the total spacetime metric will be a small perturbation of the flat metric: $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$. We can define Einstein's equations for this metric and study the propagation of it as a wave with the perturbative approach.

Within this assumption, the stress-energy tensor and the Ricci tensor will be the sum of the unperturbed and the perturbed contributions. Neglecting terms of the order h^2 and above: $T_{\mu\nu} = T_{\mu\nu}^{(\eta)} + T_{\mu\nu}^{(h)}$, $R_{\mu\nu}(g) = R_{\mu\nu}(\eta) + R_{\mu\nu}(h)$. If we explicit the affine connections, since $\Gamma_{\mu\nu}^\alpha(\bar{g}_{\mu\nu} = \eta_{\mu\nu}) = 0$, we obtain

$$\begin{aligned}\Gamma_{\mu\nu}^\alpha(g) &= \Gamma_{\mu\nu}^\alpha(\eta) + \Gamma_{\mu\nu}^\alpha(h) + O(h^2) = \\ &= \Gamma_{\mu\nu}^\alpha(h) + O(h^2) = \\ &= \frac{1}{2}\eta^{\alpha\sigma}(\partial_\mu h_{\sigma\nu} + \partial_\nu h_{\sigma\mu} - \partial_\sigma h_{\mu\nu}) + O(h^2).\end{aligned}\tag{1.10}$$

In this case, the Einstein's equations will be

$$R_{\mu\nu}(g) = R_{\mu\nu}(h) = \frac{8\pi G}{c^4} \left(T_{\mu\nu}^{(h)} - \frac{\eta_{\mu\nu}}{2} T^{(h)} \right).\tag{1.11}$$

Note that already in this form the equations are linear, but we can show it in a more evident way if we explicit the Ricci tensor using the affine connections in Eq. 1.10:

$$\begin{aligned}R_{\mu\nu}(h) &= \partial_\alpha \Gamma_{\mu\nu}^\alpha - \partial_\nu \Gamma_{\mu\alpha}^\alpha = \\ &= -\frac{1}{2} [\Box_{flat} h_{\mu\nu} - (\partial_\lambda \partial_\mu h_\nu^\lambda + \partial_\lambda \partial_\nu h_\mu^\lambda - \partial_\mu \partial_\nu h_\lambda^\lambda)],\end{aligned}\tag{1.12}$$

where $\Box_{flat} h_{\mu\nu} = \eta^{\mu\nu} \partial_\mu \partial_\nu$ is the D'Alembertian in flat spacetime. If we equalize the second members of Eq.s 1.11 and 1.12, we have

$$-\frac{1}{2} [\Box_{flat} h_{\mu\nu} - (\partial_\lambda \partial_\mu h_\nu^\lambda + \partial_\lambda \partial_\nu h_\mu^\lambda - \partial_\mu \partial_\nu h_\lambda^\lambda)] = \frac{8\pi G}{c^4} \left(T_{\mu\nu}^{(h)} - \frac{\eta_{\mu\nu}}{2} T^{(h)} \right).\tag{1.13}$$

We know that, because of the Gauge symmetry, the solutions of these equations are not uniquely determined. If we choose a coordinate system where the so-called *harmonic gauge* condition is satisfied, that is $g^{\mu\nu} \Gamma_{\mu\nu}^\sigma = 0$, Eq. 1.13 can be simplified: it can be shown that in the weak field limit the harmonic gauge is equivalent to the relation $\partial_\mu h^\mu{}_\nu = \partial_\nu h^\mu{}_\mu/2$, and the second term on the left side of Eq. 1.13 will be equal to 0 (to hereafter $T = T^{(h)}$), i.e.:

$$\Box_{flat} h_{\mu\nu} = -\frac{16\pi G}{c^4} (T_{\mu\nu} - \frac{\eta_{\mu\nu}}{2} T_\lambda^\lambda).\tag{1.14}$$

⁴Note that in vacuum $T_{\mu\nu} = 0 \Rightarrow R_{\mu\nu} = 0$, and while Ricci tensor vanishes in vacuum, Riemann tensor does not, unless the gravitational field is null, or constant and uniform.

Keeping in mind equations 1.7 and 1.8, and defining the tensor $\tilde{h}_{\mu\nu} := h_{\mu\nu} - \eta_{\mu\nu}h^\lambda{}_\lambda$, we can write the two conditions as

$$\begin{cases} \square_{flat}\tilde{h}_{\mu\nu} = -\frac{16\pi G}{c^4}T_{\mu\nu} \\ \partial_\mu\tilde{h}^\mu{}_\nu = 0 \end{cases}. \quad (1.15)$$

Since in the vacuum we have $T_{\mu\nu} = 0$, then $\square_{flat}\tilde{h}_{\mu\nu} = 0$.

We will refer to this relations as *wave equation*, because it is very similar to the one describing an electromagnetic wave in the vacuum, so as a propagation of a linear superposition of plane waves. Their simplest solution will be a monochromatic plane wave $\tilde{h}_{\mu\nu} = \Re[A_{\mu\nu}e^{ik_\mu x^\mu}]$, where $A_{\mu\nu}$ is the polarization tensor and $k = (\omega/c, \mathbf{k})$ the wave vector.

If we replace this solution in the harmonic gauge condition, we will obtain $k_\mu A^\mu{}_\nu = 0$, which means that the wave vector and the polarization tensors are orthogonal and the wave is transverse. If we instead replace it in the wave equation, it will result $k_\mu k^\mu = 0 \Rightarrow \omega \equiv ck_0 = c\sqrt{k_i k_i}$, which means that the gravitational waves travel at the speed of light.

The harmonic gauge condition is not enough to determine the gauge uniquely, since it can be shown that from the 10 degrees of freedom of the (symmetric) metric tensor $h_{\mu\nu}$, it fixes only 4 of them, leaving 6 components still free. It can be shown that imposing the *transverse-traceless gauge* (TT-gauge) choice, where we have $h^{0\mu} = h_{ii} = \partial_j h_{ij} = 0$, the degrees of freedom are collapsed to only two, which are the polarization states of the gravitational waves. The metric tensor in the TT-gauge, which is usually referred to as, $h_{\mu\nu}^{TT}$ will then look like:

$$h_{\mu\nu}^{TT} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & h_+ & h_\times \\ 0 & 0 & h_\times & -h_+ \end{pmatrix}, \quad (1.16)$$

or, separating the two polarization states:

$$h_{\mu\nu}^{TT} = \Re[A_{\mu\nu}^+ e^{ikx} + A_{\mu\nu}^\times e^{ikx}], \quad (1.17)$$

with the two polarization tensors defined up to amplitude constants:

$$A_{\mu\nu}^+ \propto \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & +1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \quad A_{\mu\nu}^\times \propto \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & +1 \\ 0 & 0 & +1 & 0 \end{pmatrix}. \quad (1.18)$$

The reason of the $+$ and \times symbols used to label the two gravitational-wave polarization states can be explained by studying the effect of a metric perturbation on the geodesics, i.e., on the motion of particles. To do that we will study the effect of the relative motion of two particles. To understand what happens when a gravitational-waves signal passes between them, we have to define the concept of *geodesic deviation*.

Let us consider two particles moving freely in a gravitational field, respectively along the geodesics $x^\mu(\tau)$ and $x^\mu(\tau) + \delta x^\mu(\tau)$. If we define a set of geodesics $x^\mu(\tau, p)$,

with p a dimension that distinguishes the different geodesics of the set, we can write the separation between two geodesics as $\delta x^\mu = \partial x^\mu / \partial p$. The vector tangent to the geodesic is $u^\mu = \partial x^\mu / \partial \tau$, it follows that $\partial u^\mu / \partial p = \partial \delta x^\mu / \partial \tau$.

Another key element is the definition of the *covariant derivatives*. Since the derivative in general is not invariant for a coordinate transformation⁵, it is necessary to define a derivative of a four-vector V^μ which does not change in any reference frame. These are the covariant derivatives and they are defined as

$$D_\nu V^\mu := \partial_\nu V^\mu + V^\sigma \Gamma_{\sigma\nu}^\mu. \quad (1.19)$$

With these definitions, it is possible now to define the second covariant derivatives of the separation vector δx^μ along a fixed geodesic (i.e. $p = \text{const}$), called also geodesic deviation,

$$\frac{D^2}{D\tau^2} \delta x^\alpha = R^\alpha_{\beta\mu\nu} u^\beta u^\mu \delta x^\nu, \quad (1.20)$$

which describes the relative acceleration of nearby particles moving along close geodesics according to the curvature tensor.

Applying the geodesic deviation to nearby particles, and computing the results for a small perturbation, we have

$$\delta x_j \left(t - \frac{x}{c} \right) = \delta x_{jt=0} + \frac{1}{2} \eta^{ji} h_{ik}^{TT} \left(t - \frac{x}{c} \right) \delta x_{kt=0}. \quad (1.21)$$

Given the transverse nature of the gravitational waves, with a wave moving along the x axis, we will have $\delta x_1 = \text{const}$. The two other components, following equation from 1.16 to 1.17, will be

$$\begin{cases} \Delta \delta x^2 = (h_+ \overline{\delta x^2} + h_\times \overline{\delta x^3}) \\ \Delta \delta x^3 = (-h_+ \overline{\delta x^3} + h_\times \overline{\delta x^2}), \end{cases} \quad (1.22)$$

where, assuming that the amplitudes A are real functions, the two polarization states can be written as

$$h_{\left(\begin{smallmatrix} + \\ \times \end{smallmatrix}\right)} = 2\Re \left[A_{\left(\begin{smallmatrix} + \\ \times \end{smallmatrix}\right)} e^{i\omega \left(t - \frac{x}{c} \right)} \right] = 2A_{\left(\begin{smallmatrix} + \\ \times \end{smallmatrix}\right)} \cdot \cos \left[\omega \left(t - \frac{x}{c} \right) \right] \quad (1.23)$$

This is the periodic oscillation generated by a gravitational wave. If we imagine a ring in the vacuum and a linearly polarized plane wave, crossing perpendicularly the plane of the ring, it will cause a periodic distortion of the metric itself. The effect will be an orthogonal compression and dilation of the ring to an ellipse, with major semiaxes along the coordinate axis or rotated by 45 degrees, depending on whether the linear polarization is respectively $+$ or \times , explaining the reason of their names (see Fig. 1.1).

⁵In a curved manifold, the derivative of the basis vectors $\mathbf{e}_{(\mu)}$ is not in general zero, since $\partial_\beta \mathbf{e}_{(\mu)} = \Gamma^\mu_{\alpha\beta} \mathbf{e}_{(\mu)}$. Inverting the point of view, this means that the affine connections express how the basis vectors change when the coordinates are transformed




Figure 1.1: The effect of linearly polarized TT-gauge gravitational waves, crossing a ring of particles in vacuum along the x direction. Figure from Le Tiec et al., 2017.

1.2 Compact objects

For astrophysical phenomena, in order to have relativistic effects strong enough to produce detectable gravitational waves, it is necessary to have objects extremely compact. These objects come from the ending stage of the stars evolution, whose outcome is related to the progenitor mass: from stars with $M \lesssim 8M_\odot$ we expect to obtain white dwarfs, while neutron stars likely had progenitors with mass between $8M_\odot$ and $25 - 60M_\odot$ (Shapiro et al., 2008), while above the latter range the only possible result is a black hole.

White dwarfs are star cores which remain exposed for various reasons, usually either because they lose the outer layers dragged away by a companion object, either because the radiation pressure exceeds the gravitational self-attraction leaving around a planetary nebula (Carroll et al., 2017). Their radius is approximately the Earth radius and their mass is comparable to the Sun mass $M_\odot \approx 1.99 \times 10^{30} kg$. Given their high density, since the white dwarfs do not have anymore thermonuclear processes which can hold the pressure of the gravity of their mass, they remain stable because of the *degenerate condition* of their electrons.

Neutron stars come from a core collapse supernova where the energies and densities bring to a neutron synthesis not balanced anymore by their β decay. They are made then for the vast majority by neutrons, having a mass in the range $\simeq 1 - 1.4M_\odot$ and a radius of $\simeq 10 - 14$ km, and their equilibrium is driven by equations of state which in first approximation involve the degeneracy of neutrons.

In both cases there is a mass limit that breaks the equilibrium between gravity and fermion degeneracy, with a collapse of the star. As we will see in Sect. 1.2.2, in certain conditions the equilibrium of the degenerate matter is not able to find a stable configuration. This can happen when the core collapse of the parent star triggers, or because mass accretion, or during the collision of two compact massive objects. The result will be a collapse into a black hole, and a key parameter for

this behavior is the mass: as we will see in Sect. 1.2.2, the problem of the mass limit for compact objects is crucial for the study of their physics.

A black hole is a relativistic singularity theorized by the General Relativity, which is formed when an object reaches a density such that nothing can resist to the gravitational pressure and all its matter collapses in a material point. Following General Relativity, the physics of black-holes is fully described by a metric which includes as parameters only their mass, angular momentum and eventually electric charge, although the last quantity is usually assumed null. Their size is defined by the radius at which the escape velocity reaches the speed of light in the vacuum, the so-called *Schwarzschild radius* (Ferrari et al., 2020; Hartle, 2003):

$$R_s = 2 \frac{GM}{c^2}, \quad (1.24)$$

where G is the gravitational constant, M the mass of the object and c the speed of light in vacuum.

Gravitational waves coming from neutron stars are the main topic of this thesis and their emission mechanisms will be presented in Sec. 1.3.2.

Currently gravitational waves from white dwarfs cannot be detected from ground based detectors but white dwarfs binary systems produce a foreground that will be detected by future space interferometer LISA, limiting its sensitivity (Danzmann et al., 2017). In any case, even though we are not treating the physics of white dwarfs in this work, the concept of fermion degeneracy used to describe them plays a key role also in the existence of neutron stars, and this is why they have been mentioned.

1.2.1 Degenerate fermions

A property of all fermions (hence of electrons in white dwarfs and neutrons in neutron stars) is that the total wave function of a system of particles is antisymmetric, meaning that it changes sign for the exchange of two particles. This translates into the *Pauli exclusion principle*, that states that two or more fermions cannot occupy the same quantum state: if we had two fermions with the same state, e.g. the same atomic orbital and same spin, swapping them would not change the total wave function violating the antisymmetric condition.

When the temperature of a gas of fermions of the same kind is very low, ultimately $T = 0$ K, the particles will occupy the lowest possible energy state, stacking up to higher energy when they will find a state already occupied, up to an energy (Carroll et al., 2017; Shapiro et al., 2008)

$$\varepsilon_F = \frac{\hbar}{2m_f} (3\pi^2 n_\nu)^{\frac{2}{3}}, \quad (1.25)$$

where m_ν and n_ν are respectively the mass of the fermions we are considering and their number per unit volume, $\hbar \approx 1.055 \times 10^{-34}$ m² kg/s is the reduced Planck constant. The quantity defined by Eq. 1.25 is referred to as *Fermi energy*, and the system we are considering is a *Fermi-Dirac condensate*, or a *gas of degenerate fermions*. In this condition, the pressure given by the momentum of the degenerate fermions can be enough to counterbalance the self gravity of the compact star.

It is important to remark that in real situations $T > 0$ K, hence not all the fermions will be in a degenerate state: some fermions will have thermal states with an energy $\varepsilon_T > \varepsilon_F$.

Nevertheless, writing down the thermal energy as $\varepsilon_T = \frac{3}{2} kT$, with $k \approx 1.38 \times 10^{-23} \text{ m}^2 \text{ kg s}^{-2} \text{ K}^{-1}$ the Boltzmann constant, we can define a condition for the gas of fermions to be still degenerate even with $T > 0$ K.

If we impose that the thermal energy has to be lower than the Fermi one, i.e. $\varepsilon_T < \varepsilon_F$ for a given T , replacing the expressions of the two energies in the latter relation we will have

$$3kT < \hbar (3\pi^2 n_\nu)^{\frac{2}{3}} / m_f, \quad (1.26)$$

which shows that even at $T > 0$ the degeneracy condition can be satisfied, depending on the mass and density of the fermions considered.

Using the definition of the Fermi momentum p_F (Carroll et al., 2017; Shapiro et al., 2008),

$$\epsilon_F = (p_F^2 c^2 + m_\nu^2 c^4)^{\frac{1}{2}} \quad (1.27)$$

we can express the number density n_ν of degenerate fermions, assuming an isotropic distribution of states in the momentum space of free particles (Carroll et al., 2017; Shapiro et al., 2008):

$$n_\nu = \int_0^{p_F} \frac{p^2}{\pi^2 \hbar^3} = \frac{p_F^3}{3\pi^2 \hbar^3}. \quad (1.28)$$

From the kinetic theory, the pressure will be

$$P_\nu = \int_0^{p_\nu} \frac{p^2 c^2}{\epsilon} n_\nu dp. \quad (1.29)$$

Replacing equations 1.27 and 1.28 in 1.29, the integral can be simplified using the non-relativistic limit ($E \approx mc^2 + p^2/2m$), yielding to the relation

$$P_\nu = \frac{1}{5} (3\pi^2)^{2/3} \left(\frac{\hbar}{m_\nu} \right) n_\nu^{5/3}. \quad (1.30)$$

Using the ultra-relativistic limit ($E \approx pc$) we obtain:

$$P_\nu = \frac{1}{4} (3\pi^2)^{1/3} (\hbar c) n_\nu^{4/3}. \quad (1.31)$$

These are the *equations of state* for a gas of completely degenerate fermions, in the non-relativistic and ultra-relativistic limit. They can be both written in the shape of a polytropic relation $P = K\rho^\gamma$, where γ depends on the model used to define the pressure-density relation of the system (Hartle, 2003; Steiner et al., 2010). A larger γ implies that the pressure increases more for a given increase of the energy density, and hence the equation of state is *stiffer*.

1.2.2 Hydrostatic equilibrium and mass limit

Another key element of a gas of free degenerate fermions is that there is a non-trivial relationship between mass and volume of a compact object in a general-relativistic treatment. To understand the structure and composition of a compact




Figure 1.2: Schematic plots from Shapiro et al., 2008 showing the equilibrium configuration of a gas of degenerate fermions for a given equation of state, as plots of mass M/M_\odot in function of central density ρ_c and star radius R . The turning points of stable-unstable intervals are shown with letters from A to E and the stable intervals for white dwarfs and neutron stars are specified.

star, the study of its hydrostatic equilibrium is the fundamental starting point, and a fully relativistic treatment is necessary for dense objects like neutron stars.

Let us consider a non-rotating spherically symmetric star of mass M . The so called Tolman-Oppenheimer-Volkoff (TOV) equation for the hydrostatic equilibrium (Oppenheimer et al., 1939; Tolman, 1939), solution of the Einstein's equation for a perfect fluid with spherical distribution, is:

$$\frac{dP}{dr} = -\frac{G}{r^2} \left(\rho(r) + \frac{P(r)}{c^2} \right) \left(m(r) + 4\pi r^3 \frac{P(r)}{c^2} \right) \left(1 - 2G \frac{m(r)}{rc^2} \right)^{-1}, \quad (1.32)$$

where r is the distance from the center and P , ρ and m are respectively the pressure, the density and the mass as functions of r .

Eq. 1.32 can be integrated with the following boundary conditions:

- $m(r=0) = 0$ and $\rho(r=0) = \rho_c$, with the latter being the free parameter central density;
- $m(r=R) = M$ and $\rho(r=R) = 0$, where M and R are respectively the star total mass and radius, which means that at the surface no source of pressure is left and all the mass of the star is therein included.

If ρ_c varies, we can have the family of all degenerate matter stars, which are in equilibrium for a given equation of state (Hartle, 2003; Shapiro et al., 2008).

Among all possible configurations, only a portion of them is stable. Working similarly to a configuration in energetic equilibrium, an unstable hydrostatic equilibrium will fall dynamically to a stable one at the first perturbation. The analysis of the stable configuration goes through the analysis of the stability of normal modes of vibration of the star. If for a given equation of state and central density ρ_c , there is at least one normal mode that grows indefinitely, then the star is unstable and it will break up or fall in a stable configuration (Hartle, 2003; Shapiro et al., 2008).

Fig. 1.2 shows the relationship of the star mass M versus ρ_c , and M versus the star radius R , for a fixed equation of state in a schematic way from Shapiro et al.,

2008. The points labeled from *A* to *E* highlight the starting and ending points of unstable configurations. The range where white dwarfs and neutron stars can exist in a stable configuration are shown as well.

It is important to remind that, in any general relativistic treatment, there is always a mass limit for a compact star above which it falls into a singularity. Even though, especially for neutron stars, this mass limit depends strongly on the equation of state and rotation rates, there are two general limits which strictly bound the maximum mass they can sustain (Kalogera et al., 1996; Rhoades Jr et al., 1974). The first comes straightforwardly from the Schwarzschild radius: any object with mass M smaller than its $R_s \leq 2GM/c^2$ cannot exist. The second comes from considerations about causality: the sound speed $c_s^2 = dP/d\rho$ inside the star has to be less than the light speed, which implies $R \lesssim 3GM/c^2$ (Lattimer, 2012; Lattimer and Prakash, 2004). From these considerations, a general limit for the maximum mass a neutron star can achieve is $M \sim 3M_\odot$.

As we have seen, the study of the condition of hydrostatic equilibrium of a compact star is very closely bound to the definition of its equation of state, that is the relationship between pressure and energy-mass density distribution (i.e., its gravitational attraction). While for white dwarfs the equation of state is well described by a gas of free degenerate electrons in a lattice of ionized atoms, with γ varying from 5/3 to 4/3 as the density increases (and hence the electron kinetic energy), for neutron stars the physics becomes much more complex both from the point of view of gravitation and particle physics.

While for white dwarfs a non-general relativistic treatment of the equation of state of a gas of degenerate electrons for the whole volume can be sufficient, for neutron stars the density reaches a level high enough to dramatically change the nature of the matter composing it. Usually, the integration of the *TOV* equation is split in several layers, with different fixed values of ρ and different equations of state that describe separately the matter in every layer. While the outer layers are well described by equations of state of perfect gas of degenerate fermions, in the inner layers, representing the vast majority of the mass of the star, the high density forces the matter to interact in a non trivial way (e.g nucleon-nucleon, meson-nucleon and QCD interactions, Lattimer and Prakash, 2004; Shapiro et al., 2008), giving significant contribution to the pressure-density relation, and hence bringing to several models of equation of state.

1.2.3 Neutron stars

As already stated, neutron stars are compact objects of mass around $1.5M_\odot$, radius of the order of $\simeq 10$ km and a density that at their center can reach values $\rho_c \gtrsim 10 \times \rho_n$, with ρ_n the typical nuclear equilibrium density of $\simeq 2.8 \times 10^{14}$ g/cm³ (corresponding to a nucleon density $n_n \simeq 1.6 \times 10^{38}$ cm⁻³) (Lattimer and Prakash, 2004; Shapiro et al., 2008). They are mainly composed by neutrons, but their structure is more complex and a complete physical description needs different layers with increasing density, different matter compositions and hence different equations of state.

Neutron stars are formed by the core-collapse supernova of a $\gtrsim 8M_\odot$ star in the latest stage of its life. The detailed mechanism of a supernova explosion, and then

the formation of a neutron star rather than a black hole are outside the purpose of this work, but key elements are the trigger of the collapse, due by the fact that the star core reached its mass limit and the role of neutrinos in driving away the gravitational energy loss by the collapse, giving a significant contribution to the explosive shockwave.

The collapse is triggered when the iron inner core of the progenitor reaches the Chandrasekhar limit. Because thermonuclear processes in the outer layers, structured as an onion-like structure made by elements progressively lighter (Burrows, 2000; Maggiore, 2018), the inner core steadily accretes mass (Burrows, 2000). Since no esothermic nuclear fusion reactions can happen for the iron nuclei, the only pressure holding the matter of the inner core against its self-gravity is the degeneracy energy of the electrons. When the iron core reaches M_{ch} , no equilibrium configurations can exist and the collapse is triggered (Burrows, 2000; Lattimer and Prakash, 2004; Shapiro et al., 2008).

When the core collapses to densities of $\sim 10^7 \text{ g cm}^{-3}$ (Carroll et al., 2017) the electrons become ultra-relativistic (Lattimer and Prakash, 2004) and, in the nuclei of the star core, the electron capture process is triggered $p^+ + e^- \rightarrow n + \nu_e$.

This phenomenon causes a neutron enrichment of the proto-neutron star and is the main source of neutrinos that will trigger the supernova explosion, eventually leaving then the star to its equilibrium configuration, where most of the matter is composed by free neutrons in a degeneracy state.

In this configuration they are stable and do not decay via a β decay ($n \rightarrow p^+ + e^- + \bar{\nu}_e$), because they are *already* in their minimum energy state allowed. Since the internal temperature is $\sim 10^9 \text{ K}$, thermally excited particles can still go through subsequent beta decays of non-degenerate neutrons and electron-proton combinations, eventually with the participation of another nucleon (the so-called URCA process, Lattimer and Prakash, 2004; Shapiro et al., 2008). After $\sim 10^5 \text{ s}$ the neutrino emission will be negligible and the neutron star will cool with a steady decreasing rate with emission of X photons (Lattimer and Prakash, 2004).

Two important features of neutron stars come from the fact that in the collapse of the parent star core, the angular momentum and magnetic field flux are expected to be conserved. This implies:

- a high rotation period of 1 - 10 s (reaching 10^{-3} s for objects that acquired angular momentum via mass accretion)
- magnetic fields of $\sim 10^{12} \text{ G}$ up to $\sim 10^{15} \text{ G}$

The very high magnetic field is responsible for periodic electromagnetic pulsations from most of the known neutron stars, which are then referred to as *pulsars*. This emission has been modeled as due to their magnetic fields, which — together with their rotation — let them act like beacons. The known pulsars are very interesting potential gravitational-waves sources (R. Abbott et al., 2021c).

It is expected that both in the early stage of their formation and later, various emission mechanisms will produce gravitational waves. A discussion is given in Sec. 1.3.2.

Structure and characteristics

On the basis of what is explained in Sec. 1.2.2, the description of the structure of neutron stars comes from the solution of its stable hydrostatic equilibrium configuration, with given equations of state that express the models of the matter constituting the object.

As we have seen, the central density plays an important role to define which neutron stars can be stable. At the same time, a fundamental hypothesis for any stable configuration is that the density is a function of the distance from the center, decreasing until it becomes 0 at the star surface. Two density values are very important to describe the neutron star composition, which are the neutron drip density $\rho_{drip} = 3 \times 10^{14} \text{ kg m}^{-3}$ and the nuclear density $\rho_{nucl} = 2.3 \times 10^{17} \text{ kg m}^{-3}$, which represent strong variations of the physical behavior of the matter inside the star. Because of that, neutron stars are described by concentric shells at various radii, where different equations of state describe the different physics (Lattimer and Prakash, 2004; Shapiro et al., 2008):

- A tiny atmosphere and envelope whose primary role is to determine the release of photons and thermal emission from the inner layers, with a negligible mass and a crust 1 - 2 km thick with density from 10^6 g cm^{-3} to ρ_{nucl} . From the outer layer of the crust to ρ_{drip} , the crust is made mainly by neutron rich nuclei starting from ^{56}Fe in the outer layers, to atoms having a number of nucleons $A \simeq 200$, made by 80 - 90% of neutrons in the inner part. This layer is made by a lattice of nuclei embedded in the electron degenerate gas. In this configuration, the neutrons in the nuclei do not decay because all the energy slots that a produced electron would occupy are already filled by the degenerate gas. When the neutron drip is reached, more and more neutrons will leak out of the nuclei lattice, overtaking the role of the electrons to determine the pressure of the system. They create a superfluid that interacts with the rest of the crust, and may generate rearrangements which contribute to change the moment of inertia of the star, causing a sudden increase of the rotational speed, the so-called *glitch*. Approaching the edge with the core, the nuclei are dissolved in the fluid of neutrons.
- The core represents $\simeq 99\%$ of the mass of the star and is usually described by an outer region made by nucleons, where the neutrons form a superfluid and protons a superconductor, plus electrons and muons, and an inner region composed probably by either strange-rich matter, condensate of π and K mesons, deconfined matter or other quantum chromodynamics modeling, where the strong interaction plays by far the dominant role at densities up to ~ 10 times that of typical atomic nuclei.

Representing almost all the mass of the star, the core determines the gravitational behavior of neutron stars and it is where the model of the equation of state becomes crucial. It is evident that the uncertainty in the description of matter at such high density in the neutron star's core is very significant.

Observations during the recent years have tried to reduce the space where equations of state can span, with an important role played by the detection of

the neutron star binary coalescence GW170817 (B. P. Abbott et al., 2017d, 2018). Following recent reviews (Lattimer, 2012, 2021), (Al-Mamun et al., 2021; Steiner et al., 2010), more and better estimations of masses and radii for several neutron stars are available from new detectors like the NICER X-ray telescope, in combination also with on-Earth nuclear experiments and the analysis of GW170817, which helped to constrain the space of models and key parameters such as maximum and minimum masses. An interesting example is a method developed in Steiner et al., 2010, based on a Bayesian analysis matching the available data on a parametrization of possible neutron star equation of states, modeling the external shells with known degenerate matter physics, up to density $\sim \rho_{nuc}$, and using two different polytropes to describe the core equation of state.

An example of the results of this kind of analysis is summarized in Fig. 1.3, where M - R relations coming from a set of different equation of states are compared to the estimation from Steiner et al., 2010. In the study of neutron star equation of states, in fact, it is useful to look at the M - R modeled curves produced by the modeled equation of states (Lattimer, 2012). In the left plot of Fig. 1.3 several curves are shown from four families of equation of states (Lattimer, 2012): purely strange quark matter neutron stars; hadronic matter with exotic physics (like condensate of K or π mesons, hyperons, quark deconfinement) that which, cannot reaching $2M_\odot$, are not able to describe the most massive observed neutron stars as observed they do not reach the value of $2M_\odot$ (Lattimer and Prakash, 2004); two other families of more standard hadronic matter neutron stars, one stiffer and one softer, showing that the radius is weakly dependent on the mass in the range between $\simeq 0.5 - 1.5M_\odot$ (Lattimer, 2012; Lattimer and Prakash, 2004). An interesting result from Steiner et al., 2010 is that the family of equation of states with the mass degeneracy at $R \simeq 12$ km is favoured by the available data available.

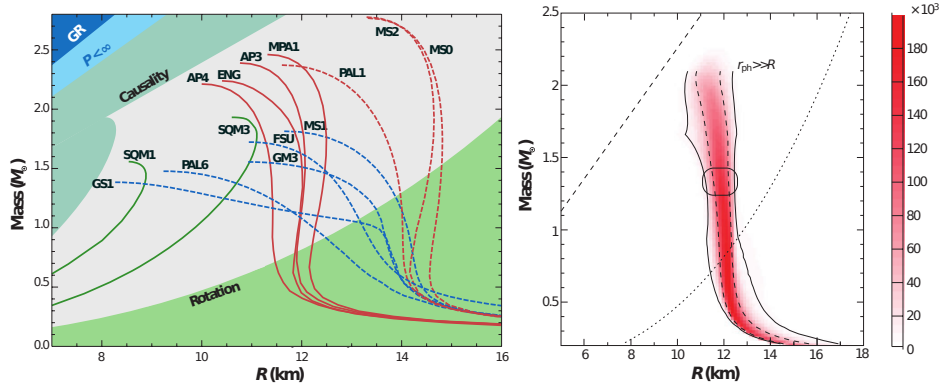


Figure 1.3: Plots showing recent results on equation of state models. Left plot: adapted from Lattimer, 2012, the M - R relations of several equations of state divided in four families: pure quark matter (green), soft equation of states with exotic matter models (dashed blue), hadronic equation of states with a mass distribution degeneracy at $\simeq 12$ km (red), softer hadronic equation of states with asymptotic behavior at higher radii (dashed red). Right plot: the results of the mentioned Bayesian analysis from Steiner et al., 2010 (image from Lattimer, 2012) as M - R posterior distribution, showing that the family of equation of states showing the vertical behavior at $\simeq 12$ km is favored with respect to the others with the data available so far.




Figure 1.4: Scheme of the pulsar emission due to the misaligned rotating magnetic field. Figure from Carroll et al., 2017.

Pulsars

Pulsars are neutron stars that emit electromagnetic pulsations, characterized by their strong magnetic field and fast rotation.

As previously stated, the emission mechanism of a pulsar is essentially explainable as a beacon of radio photons powered by the rotation of the magnetic field misaligned with respect to the rotation axis. The rapidly varying dipole field produces an induced electric field strong enough to drag charged particles from the surface of the star. The particles will follow the magnetic field lines producing a magnetosphere around the pulsar, co-rotating with it. At the distance $R_c = c/\omega_{rot}$, where ω_{rot} is the rotation pulsation of the pulsar, the speed of the particles will reach c and the field lines are open (Carroll et al., 2017); within this radius, defined as *light cylinder*, they are instead closed.

The electrons accelerated by the magnetic field travel along the curved field lines, spiraling around them and emitting gamma photons through Bremsstrahlung process, roughly parallel to their instantaneous motion direction. The emitted photons are energetic enough to decay in $e^+ - e^-$ pairs, which in turn are accelerated, triggering a cascade process which ultimately results in the narrow beacon of radio emission along the magnetic axis of the star (see Fig. 1.4).

It has been observed that the pulsars timing and pulsation profiles are generally very stable (excluding glitches and other transient phenomena). Their dynamics is strongly characterized by the rotation period P and, because of the steady loss of energy due to their pulsations, by the spindown rate \dot{P} . Hence, it is convenient to describe the population of known pulsars with P - \dot{P} plots, as shown in Fig. 1.5.

Pulsars usually emit in the radio band, with periods around 0.1-100 seconds and magnetic fields around $10^{10} - 10^{12}$ G. Few other groups with different characteristic




Figure 1.5: P - \dot{P} diagram for various classes of pulsar. Magnetars are shown in red squares, rotating radio transient (RRAT, pulsars which show unstable emission) in blue stars, X-ray dim isolated neutron stars (XDIN, characterized by their thermal X-ray emission) in green triangles. Expected magnetic field B and age in function of P and \dot{P} are also indicated. The green line is the so-called *death line*, threshold below which the particles dragged by the magnetic field are no longer energetic enough to produce the cascade pair production effect and the pulsar stops its emission (Steiner, 2020).

exist, especially for pulsars in binaries with mass accretion: the transfer of mass from the companion star provides additional angular momentum that accelerates the rotation up to periods of $P \sim 10^{-3}$ s. The in-falling particles, losing gravitational energy in the process, are much more energetic and produce X-ray emission instead of radio (Nagase, 1989). Millisecond radio pulsars can be found isolated if their original system has been disrupted (Bailes et al., 1997; Lyne et al., 1987).

Another class of pulsars, the magnetars (Mereghetti, 2008), have an emission rate that, for the observed rotational periods, can be explained with a magnetic field of various orders of magnitude higher than standard pulsars, up to $\sim 10^{16}$ G (see Fig. 1.5). . The high magnetic field makes these sources peculiar, with a persistent X-ray emission and a series of repeated gamma-ray bursts, as well as sudden giant flares in the gamma-ray band.

1.3 Sources of gravitational waves

Now that the relativistic explanation of the gravitational waves is given and having introduced the physical behavior of the sources of major interest for this work, we can clarify how the gravitational waves from them are generated.

Although many classes of sources exist, most of the astrophysical ones can be explained as an oscillation of the quadrupole moment of the system, for example an inspiraling binary system made by two compact objects which eventually ends in a coalescence, or the rotation of a non-axisymmetric neutron star, either isolated or in a binary system. The possibility to detect a gravitational-wave stochastic background of cosmological nature will be discussed as well in Sec. 1.3.3.

1.3.1 Quadrupole formalism

The quadrupole formalism is a very useful approach to compute waveforms emitted by a wide variety of gravitational-waves sources. Let us consider Eq. 1.15 and assume that the systems we are considering do not vary too fast, i.e., the space extension of the source is much smaller than the emitted wavelength, i.e., given a length ϵ , the emitted wavelength is $\lambda_G = 2\pi c/\omega_G \gg \epsilon$. This implies that the stress-energy tensor $T_{\mu\nu}$ is non-zero in a region with spatial coordinates $|x_i| < \epsilon$, and we are considering processes with typical velocities $\epsilon\omega_G \ll c$. This is called slow-motion approximation.

Defining the Fourier transform of a function $f(t, \mathbf{x})$ as

$$f(\omega, \mathbf{x}) = \int_{-\infty}^{+\infty} f(t, \mathbf{x}) e^{i\omega t} dt,$$

Eq. 1.15, in the frequency domain, can be written as follows (Ferrari et al., 2020):

$$\left(\nabla^2 + \frac{\omega^2}{c^2} \right) h_{\mu\nu}(\omega, \mathbf{x}) = -\frac{16\pi G}{c^4} T_{\mu\nu}(\omega, \mathbf{x}), \quad (1.33)$$

where $\nabla := \{\partial_1, \partial_2, \partial_3\}$ is the vectorial nabla operator whose components are the partial derivatives on the three spatial coordinates.

To find the gravitational-wave power emitted by a generic source, we have to solve Eq. 1.33 both inside and outside the source, given a certain stress-energy distribution described by $T_{\mu\nu}$. Starting outside the source, in the vacuum $T_{\mu\nu} = 0$, Eq. 1.33 becomes

$$\left(\nabla^2 + \frac{\omega^2}{c^2} \right) h_{\mu\nu}(\omega, \mathbf{x}) = 0. \quad (1.34)$$

The simplest solution for a signal coming from the source is a spherical wave expressed as

$$h_{\mu\nu}(\omega, r) = \frac{1}{r} A_{\mu\nu}(\omega) e^{i\omega r/c}. \quad (1.35)$$

The amplitude $A_{\mu\nu}$ is derived solving the wave equation inside the source, and it can be shown that, at first order of approximation, it is related to the stress-energy

tensor by the following integral relation over the source volume V :

$$A_{\mu\nu}(\omega) = \frac{4G}{c^4} \int_V T_{\mu\nu}(\omega, \mathbf{x}) d\mathbf{x}. \quad (1.36)$$

Replacing it in the outside solution and reverse-transforming to the time domain, the amplitude of emitted gravitational waves will be

$$h_{\mu\nu}(t, r) = \frac{4G}{c^4} \frac{1}{r} \int_V T_{\mu\nu} \left(t - \frac{r}{c}, \mathbf{x} d\mathbf{x} \right). \quad (1.37)$$

The integral over the spatial components of the stress-energy tensor can be expressed, via the multipolar expansion, as the second time derivative of the quadrupole moment tensor of the system Q^{ij} ⁶, i.e. (Ferrari et al., 2020):

$$\int_V T^{ij}(t, \mathbf{x}) d\mathbf{x} = \frac{1}{2} d_t^2 Q^{ij}(t), \quad (1.38)$$

with

$$Q^{ij} := \frac{1}{c^2} \int_V T^{00}(t, \mathbf{x}) x^i x^j d\mathbf{x}. \quad (1.39)$$

To obtain the physical degrees of freedom with respect to the observer, we project the metric and the quadrupole moment to the TT-gauge, with an operator dependant on the observer-source direction that returns the polarization states we can measure. The gravitational-wave solution will then be (Ferrari et al., 2020)

$$h_{ij}^{TT}(t, r) = \frac{2G}{c^4 r} d_t^2 Q_{ij}^{TT} \left(t - \frac{r}{c} \right), \quad (1.40)$$

where we note that emitted gravitational waves depend only on the second time derivative of the quadrupole moment of the system. It is important to underline that, for a stationary and spherical-symmetric energy-matter distribution, the quadrupole moment is constant, and hence $\ddot{Q} = d_t^2 Q = 0$, which implies no gravitational-wave emission. Even if the system is not rotating, as long as it does not violate its axisymmetry, the quadrupole moment will not change. When, instead, the system deviates from the rotational symmetry along its axis, $\ddot{Q} \neq 0$, gravitational radiation will be emitted. The quadrupole moment behavior of the source can be considered then as a measure of the asymmetry of the source.

Using the quadrupole moment of the source we can also derive the gravitational luminosity (Ferrari et al., 2020; Maggiore, 2008) as the total gravitational-wave energy E_{GW} emitted by the system:

$$L_{GW} := \int_{\Sigma} \frac{dE_{GW}}{dt dS} dS, \quad (1.41)$$

where the energy flux per unit surface S , $dE_{GW}/dt dS$ can be expressed in terms of the first derivative of the perturbation tensor h_{ij}^{TT} and hence, using the quadrupole formalism we have, (Ferrari et al., 2020; Maggiore, 2008)

⁶It can be shown that the mono- and di-polar terms are null (Ferrari et al., 2020) and we are neglecting higher order terms than the quadrupolar one.




Figure 1.6: Scheme of a rotating ellipsoid with three axes of different size and rotational axis parallel to one of them.

$$\begin{aligned} \frac{dE_{GW}}{dt dS} &= \frac{c^3}{32\pi G} \sum_{ij} \left(\dot{h}_{ij}^{TT}(t, r) \right)^2 = \\ &= \frac{G}{8\pi c^5 r^2} \sum_{ij} \ddot{Q}_{ij}^{TT} \left(t - \frac{r}{c} \right)^2. \end{aligned} \quad (1.42)$$

It can be shown that, dropping the TT gauge, the general quadrupole formula for the gravitational wave luminosity is (Ferrari et al., 2020; Maggiore, 2008):

$$L_{GW}(t) = \frac{G}{5c^5} \ddot{Q}_{ij}^2 \left(t - \frac{r}{c} \right) \quad (1.43)$$

1.3.2 Neutron star emission mechanisms

The main topic of this thesis is the search for continuous gravitational waves from non-axisymmetric rotating neutron stars. As shown in the previous section, gravitational-wave emission is closely related to spherical asymmetries of the system. For steadily rotating objects, if we schematize a neutron star as a rigid ellipsoid with a certain ellipticity ϵ , this is translated to cylindrical asymmetry with respect to the rotation axis.

Moreover, oscillations of the configuration of the neutron star structure can produce transient signals, according to its normal modes of oscillations, with characteristic frequencies and damping time which depend on the physical phenomena and forces that drive the oscillations, and on the equation of state of the matter involved in the process.

Asymmetrically rotating stars

Let us consider an ellipsoid with three semiaxes $\mathcal{A} \neq \mathcal{B} \neq \mathcal{C}$, rotating with angular velocity $\omega = (0, 0, \omega)$ along the axis \mathcal{C} . In the chosen reference frame, the

coordinate axis z is parallel to \mathcal{C} (see Fig. 1.6 for a scheme of the system). The equation describing the ellipsoid is:

$$\left(\frac{x_1}{\mathcal{A}}\right)^2 + \left(\frac{x_2}{\mathcal{B}}\right)^2 + \left(\frac{x_3}{\mathcal{C}}\right)^2 = 1. \quad (1.44)$$

For the sake of simplicity, let us consider the density ρ uniform, the total mass M and moment of inertia tensor

$$I_{ij} = \int_V \rho (r^2 \delta_{ij} - x_i x_j) d\mathbf{x}. \quad (1.45)$$

In the co-rotating frame, the inertia tensor can be expressed as (Ferrari et al., 2020)

$$I'_{ij} = \frac{M}{5} \begin{pmatrix} \mathcal{B}^2 + \mathcal{C}^2 & 0 & 0 \\ 0 & \mathcal{C}^2 + \mathcal{A}^2 & 0 \\ 0 & 0 & \mathcal{A}^2 + \mathcal{B}^2 \end{pmatrix} = \begin{pmatrix} I_{xx} & 0 & 0 \\ 0 & I_{yy} & 0 \\ 0 & 0 & I_{zz} \end{pmatrix}, \quad (1.46)$$

with $I_{xx,yy,zz}$ the principal moments of inertia.

The traceless quadrupole moment of the system will be

$$Q_{ij} = \int_V \rho x_i x_j d\mathbf{x} = - \left(I_{ij} - \frac{1}{3} \delta_{ij} I \right), \quad (1.47)$$

where $q = q_i^i$ and $I = I_i^i$.

With the configuration of Fig. 1.6, the rotation matrix is

$$R_{ij} = \begin{pmatrix} \cos(\omega t) & -\sin(\omega t) & 0 \\ \sin(\omega t) & \cos(\omega t) & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (1.48)$$

In the inertial frame, we have $I_{ij} = (R I' R^T)_{ij}$, and replacing it in Eq. 1.47, the time varying components of the quadrupole moment will be given by

$$Q_{ij} = \frac{I_{yy} - I_{xx}}{2} \begin{pmatrix} \cos(2\omega t) & \sin(2\omega t) & 0 \\ \sin(2\omega t) & -\cos(2\omega t) & 0 \\ 0 & 0 & 0 \end{pmatrix} = \frac{I_{yy} - I_{xx}}{2} \mathcal{A}_{ij}. \quad (1.49)$$

Note that in the last equation, the argument of the periodic functions \sin and \cos , in the expression of the time-varying quadrupole momentum, is twice the one of the rotational matrix. This means that the gravitational-wave frequency emitted will be twice the star rotational frequency. It can be shown though that, if the rotational axis does not coincide with any of the principal axes of the ellipsoid, the gravitational-wave frequency emitted will be the same of the rotational frequency (Ferrari et al., 2020).

As expected, if $I_{xx} = I_{yy} \Rightarrow \mathcal{A} = \mathcal{B}$, the axisymmetry of the rotating object will be preserved and there will not be any gravitational-wave emission.

We can express the degree of non-axisymmetry in terms of the ellipticity parameter (Ferrari et al., 2020):

$$\epsilon := 2 \frac{\mathcal{A} - \mathcal{B}}{\mathcal{A} + \mathcal{B}} \approx \frac{I_{yy} - I_{xx}}{I_{zz}} \quad (1.50)$$

$$\Rightarrow Q_{ij} = \frac{1}{2}\epsilon I_{zz} A_{ij}. \quad (1.51)$$

The gravitational waves emitted by a non-axisymmetric rotating object will be then (Ferrari et al., 2020)

$$h(t, r)_{ij} = -h_0 \cdot A_{ij}(t - r/c), \quad (1.52)$$

and the coefficient h_0 is

$$h_0 = \frac{4G}{c^4 r} I_{zz} \epsilon \omega^2. \quad (1.53)$$

This relation shows how the amplitude of the emitted wave depends on the moment of inertia of the star and its ellipticity. Moreover, it depends on the square of the rotational pulsation, so that a faster rotating star will emit more intensely and its amplitude will fade with the deceleration caused by the energy emission.

The emitted gravitational-wave luminosity from a rotating ellipsoids can be written as

$$L_{GW} = \frac{32}{5} \frac{G}{c^5} \omega^6 \epsilon^2 I_{zz}^2. \quad (1.54)$$

At first order approximation, for non-relativistic rotational speeds, the kinetic energy of a rotating object is $E = I\omega^2/2$, and its time derivative $\dot{E} = I\omega\dot{\omega}$. If we suppose a stationary rotating compact star with a given ellipticity ϵ , it will lose energy mainly via gravitational waves. Then, we will have $L_{GW} \approx -\dot{E}$, with a decreasing rotational frequency, implying that the system gravitational luminosity will fade according to the rotational energy loss.

Similarly to what happens for the electromagnetic emission in the case of pulsars, the frequency f - spindown \dot{f} relation strongly characterizes the physics of the neutron star. If we assume that the system is not accreting mass (gaining angular momentum from the inspiral matter and hence causing a positive \dot{f}), and is not losing energy through electromagnetic radiation, for the case of gravitational-wave emission it is possible to directly relate the observed signal dynamics with the neutron star ellipticity. Making explicit both members of the equation $L_{GW} \approx \dot{E}$ we have

$$\begin{aligned} \frac{32G}{5c^5} (2\pi f)^6 \epsilon^2 I^2 &\approx I^4 \pi^2 f |\dot{f}| \\ \Rightarrow \epsilon &\approx \sqrt{\frac{5c^5}{512\pi^4 G \frac{|\dot{f}|}{f^5 I}}} \end{aligned} \quad (1.55)$$

$$\Rightarrow h_0(f) \approx \sqrt{\frac{5G}{2c^3 r^2} I \frac{|\dot{f}|}{f}}. \quad (1.56)$$

If the compact star loses energy in other ways (e.g., the electromagnetic pulsations due to its magnetic field), then the ϵ estimation is an upper limit, which is related to the so-called *spin-down limit* (Lasky, 2015).

Normal modes

So far we have considered the neutron star as a rigid, stationary rotating body with a relationship between its structure, its moment of inertia and the maximum ellipticity sustainable for a given rotational speed.

In more realistic cases, neutron stars can be treated as spherical fluid distributions, with various kinds of perturbations that can be triggered by several dynamical phenomena, leading to a rich spectrum of normal modes of oscillation (Valeria Ferrari et al., 2003; Lasky, 2015; Maggiore, 2018), with characteristic frequencies and relaxation times. The physical processes able to excite these oscillation modes can arise from different origine: from the early stages of the formation of a neutron star during the core collapse or coalescence involving at least a neutron star, to rearrangements of the neutron star structure in equilibrium, which could lead for example to glitches or flares.

Let us start with the simplest configuration: a non-rotating neutron star modeled by a polytropic equation of state, without magnetic field. A class of important vibration modes are the so called *f-modes* or *fundamental modes*, which produce quadrupole variations emitting signals with frequencies (Maggiore, 2018)

$$f = (2.29 \pm 0.09) \left(\frac{M}{1.4M_{\odot}} \right)^{\frac{1}{2}} \left(\frac{R}{10km} \right)^{-\frac{3}{2}} \text{ kHz} \simeq 1.7 - 2.3 \text{ kHz}. \quad (1.57)$$

The damping time of this class of modes is $\tau \propto R^4/M^3$ (Maggiore, 2018) and, for typical neutron stars, with $M = 1.4M_{\odot}$ and $R = 10 - 14km$, is $\tau \simeq 0.1 - 0.3$ s.

If we take into account the rotation of a typical neutron star, another important class of vibration modes constitutes the so-called *r-modes* or *rotation modes*, due to the Coriolis force acting in the fluid inside a neutron star (Maggiore, 2018).

They can present an instability occurring when a r-mode, counter-rotating along a fixed axis (e.g. z) in the co-rotating frame (as it is a wave driven by a Coriolis force), can be seen instead as co-rotating by an observer at infinity, if the star angular velocity is high enough. A gravitational wave emission will drag away positive angular momentum as seen at infinity. In the co-rotating frame the r-mode has a negative angular momentum and the subtraction of a positive angular momentum via gravitational waves will increase its absolute value, consequently increasing the gravitational-wave radiation amplitude with a self-sustaining loop (Maggiore, 2018). This mechanism is called Chandrasekhar-Friedman-Schutz instability (Chandrasekhar, 1970; Friedman et al., 1978).

The expected emitted frequency is $\geq 4/3$ of the rotational frequency of the neutron star (Caride et al., 2019), while the expected growth and damping timescales and the gravitational-wave amplitude are widely uncertain (Caride et al., 2019): key factors are various forms of viscosity, which could limit the amplitude of the unstable oscillation, dampen it in a short time or even negating the process altogether (Caride et al., 2019; Maggiore, 2018).

It is worth to note that f-modes as well can show Chandrasekhar-Friedman-Schutz instabilities, even though only for rapidly rotating non-superfluid neutron stars and in this case the expected emitted frequencies are of the order of kHz (Zink et al., 2010).

1.3.3 Other gravitational-wave sources

Apart from the signals directly emitted by rotating neutron stars, several other kinds of possible sources of gravitational waves exist. Among those, we note the emission from inspiral binaries with a final coalescence, which is the only class of signals detected so far, with an event list at March 2020 of 90 detections (R. Abbott et al., 2021a), including black-hole binary systems (like the first detection, the already mentioned GW150914, B. P. Abbott et al., 2016) and systems with at least a neutron star (like GW170817, B. P. Abbott et al., 2017d). Their main features will be described in the next section, but there are other relevant sources of gravitational waves, which have yet to be detected :

- Bursts from core-collapse supernovae. As it has been explained in Sec. 1.2.3, when discussing the formation of neutron stars, this class of supernovae trigger when the mass of the object reaches the Chandrasekhar limit of $1.4M_{\odot}$ and the electron degeneracy cannot any more hold the gravitational pressure. The resulting dynamic of the collapse plus the neutrino-driven bounce and shockwave can have a highly varying quadrupole moment, with a burst of gravitational waves emission (Powell et al., 2017). Few waveform examples are shown in Fig. 1.7.
- Stochastic gravitational-wave background. It is a background of gravitational-wave emission that can come from the superposition of many unresolved astrophysical sources (Regimbau, 2011) or from cosmological phenomena (Maggiore, 2000).

The cosmological background would have great consequences for the cosmology research field, since it represents a probe on the stages of the Universe around $t \sim 10^{-43}s$ (Maggiore, 2000), much earlier than the decoupling of the Cosmic Microwave Background happened at $t \sim 380 \times 10^3$ years.

In fact, the decoupling of a fundamental interaction particle (e.g. photons, gravitons) from the primordial plasma, takes a snapshot of the Universe at the decoupling time t_{dec} , with a specific temperature T_{dec} . When the interaction is much weaker, like the gravitational one with respect to the electromagnetic one, the decoupling — and hence the respective gravitational-wave emission — happens at a far earlier time, with spectra able to give information on the Universe evolution at much higher energies far before the time of the Cosmic Microwave Background.

Although the spectrum of the cosmic gravitational-wave background is strongly model dependent, with inevitable implications from quantum gravity theories and beyond Standard Model particle physics, spectra coming from general consideration can be modeled as power-laws $\sim f^{\alpha}$, with α between 0 and 3, and cutoff frequencies at the GHz order (Maggiore, 2000).

More details on the search for stochastic gravitational waves background will be given in Chapter 2, since the analysis method used for the fast identification of continuous waves from non-axisymmetric rotating neutron stars, which is the topic of the current thesis, is normally used for this class of signals.

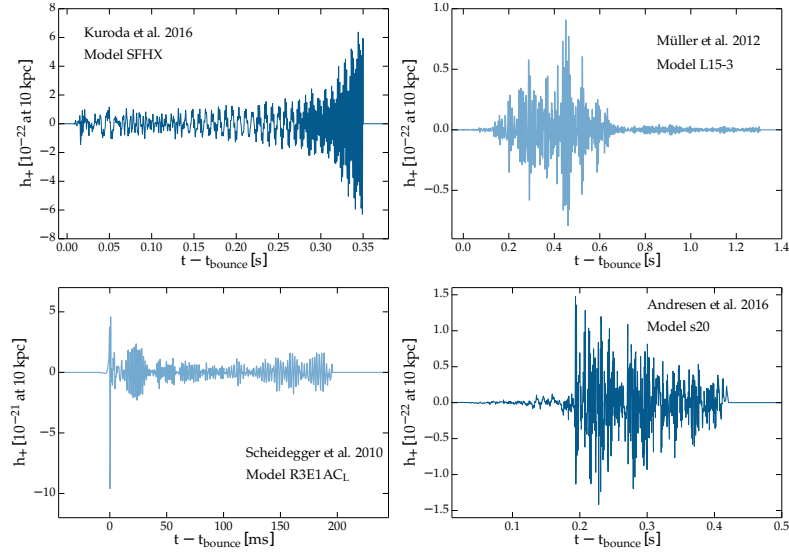


Figure 1.7: Simulated waveform using four models of core collapse supernova, with a $10kpc$ far $15M_{\odot}$ progenitor star ($20M_{\odot}$ for the bottom right plot) (Powell et al., 2017).

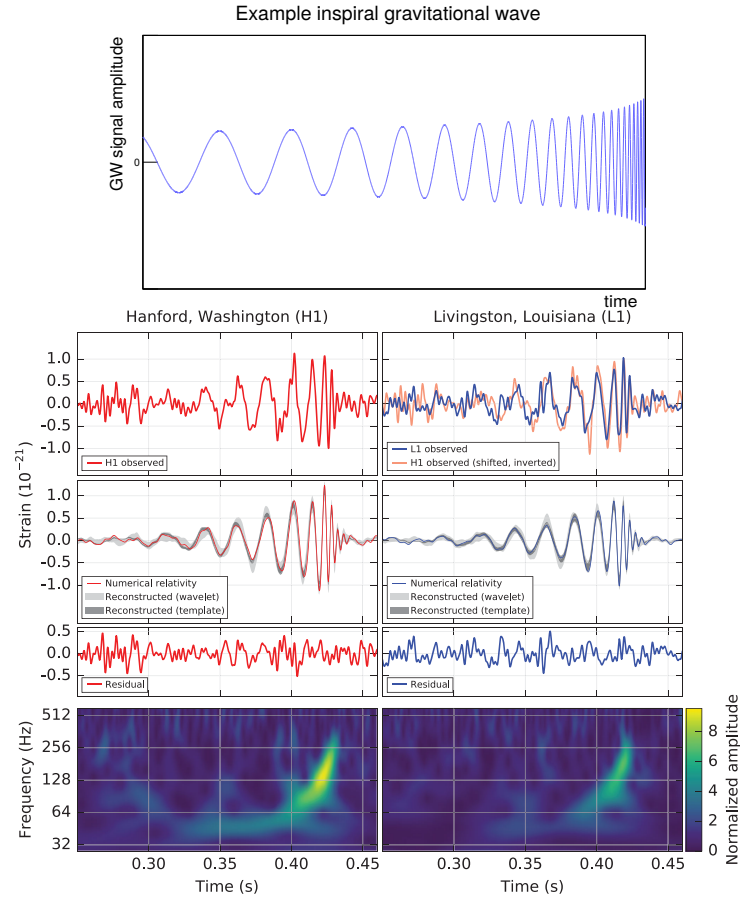


Figure 1.8: Top: evolution of the signal emitted by inspiral binaries, following the characteristic chirp behavior. Bottom: plots showing waveform and amplitude detected in the frequency-time plane for GW150914 (B. P. Abbott et al., 2016).

- Compact-binary coalescences. This class of signals, the only one detected so far (R. Abbott et al., 2021a), comes from the last stages of a binary system with compact objects (black holes and/or neutron stars). Due to a quadrupolar variation, the system loses energy resulting in the gravitational-wave emission. Because of the energy loss, the orbits of the system become closer, increasing the signal frequency as well as its amplitude, until the two objects coalesce together. The shape of the signal is the characteristic *chirp* waveform shown in the top panel of Fig. 1.8 (Ferrari et al., 2020), while in the bottom one the waveform and frequency-amplitude behavior of the GW150914 coalescence event (B. P. Abbott et al., 2016) are shown.

1.4 Detectors

The gravitational-wave detectors LIGO, Virgo and KAGRA (Aasi et al., 2015; Acernese et al., 2014; Somiya, 2012) are two-arms interferometric antennas which measure the effect of gravitational waves crossing them via the geodesic deviation caused on their arms.

The basic scheme is a Michelson interferometer, where a laser beam is split by a semi-transparent mirror (the *beam-splitter*) between two orthogonal arms. At the end of each arm, a mirror reflects back the beam to the splitter, where it is reconstructed via their respective interference. The interfered beam hits a photodiode, where the light power is measured.

The use of the light is fundamental because its travel speed does not change if it crosses any metric perturbations on its path. Considering the interferometer mirrors as test masses, we measure the geodesic deviation between them caused by the variation of the metric, i.e. the difference in distance as seen by a locally inertial frame as a variation of the path length crossed by the circulating laser.

1.4.1 Michelson configuration

Let us consider a plane gravitational wave $h_{\mu\nu}(t)$, linearly polarized along the $+$ polarization and incident perpendicularly on the plane of the interferometer. From now on, we will consider only metric perturbations in TT-gauge, so the superscript TT will be dropped. The test masses in ground based detectors are hung forming pendulums. Let us consider the test-mass n put at a distance $\mathbf{r}^{(n)} = (x_i^{(n)} - x_i^{CM})$ from the center of mass of the system.

Since the ground based interferometers have arms of the length of few kilometers, the period of a gravitational wave within the sensitivity range of the detectors will be always much higher than the time of a complete travel of a photon circulating in the detector. We can then safely use the small antenna approximation (J D Romano et al., 2017) to simplify the treatment.

As seen by an inertial frame, the geodesic deviation caused by the action of the incoming wave on the detector can be expressed as an acceleration field centered in the center of mass:

$$\frac{d^2}{dt^2} \delta x_{jtj} (t, \mathbf{r}^{(n)}) = a_j^{(n)} (t, \mathbf{r}^{(n)}) = \frac{1}{2c^2} \partial_t^2 h_{jk} r_k^{(n)}. \quad (1.58)$$




Figure 1.9: Simplified scheme of a ground based Michelson interferometer adapted from Pitkin et al., 2011. The beam splitter (m_1) is in the origin of the coordinate frame, the x and y axes are along the detector arms.

Taking a reference frame with origin on the beam splitter and the two axes along the arms (see Fig. 1.9), we can write the geodesic deviation, independently, for the three test masses of the detector and the two coordinates (the gravitational waves are always transverse with respect to their propagation). Defining $t_r^{(n)}$ the relaxing time and $\omega_p^{(n)}$ the characteristic pulsation of the pendulum n , we have (Maggiore, 2008; Saulson, 2017):

$$\begin{cases} \ddot{x}^{(n)} + \frac{\dot{x}^{(n)}}{t_r} + \omega_p^2 x^{(n)} = \frac{1}{2c^2} \left[\ddot{h}_{11} (x^{(n)} - x_{CM}) + \ddot{h}_{12} (y^{(n)} - y_{CM}) \right] \\ \ddot{y}^{(n)} + \frac{\dot{y}^{(n)}}{t_r} + \omega_p^2 y^{(n)} = \frac{1}{2c^2} \left[\ddot{h}_{21} (x^{(n)} - x_{CM}) + \ddot{h}_{22} (y^{(n)} - y_{CM}) \right] \end{cases} \quad (1.59)$$

Considering that the interferometer has both arms with same length L and, for the sake of simplicity, that the polarization axes coincide with the coordinate axes, then for the TT-gauge, we have $h_{11} = -h_{22}$; $h_{12} = h_{21} = 0$, and we can write the length difference caused by a gravitational wave as the differential equation

$$\frac{d^2}{dt^2} \Delta L + \frac{1}{t_r} \frac{d}{dt} \Delta L + \omega_p \Delta L = \frac{1}{c^2} \frac{d^2}{dt^2} h_{11} L. \quad (1.60)$$

With a difference ΔL for the length of the arms, the corresponding phase difference is $\Delta\phi = \omega_\gamma \Delta L / c$, where $\omega_\gamma / 2\pi$ is the frequency of the photons emitted by the laser. When a sinusoidal, linearly-polarized gravitational wave $h(t) = h_0 \cos(\omega_G t)$ crosses the interferometer, with amplitude h_0 and pulsation ω_G , the phase variation acquired by a photon during the $2L$ long trip from the beam splitter to the end-mirror and back will be (Maggiore, 2008; Saulson, 2017)

$$\Delta\phi \approx \frac{2\omega_\lambda}{c} h_0 \cos(\omega_G t). \quad (1.61)$$

When a signal crosses the interferometer, the photodiode will detect an oscillation of the light power at the frequency of the signal and with amplitude $\propto h_0$. Defining W_{in} the light power fed by the laser, the output power will be then $W_{out} = W_{in} \cos^2(\Delta\phi)$ and it is a direct measure of $h(t)$.

What is usually done for gravitational interferometers is to modulate the input light beam and the length of the arms in order to have an operating point of the instrument very close to the so-called *dark fringe*, that is, with $W_{out} = 0$ in absence of a gravitational-wave signal. In this way the signal-to-noise ratio is maximized and the detector is a *null instrument*, which is fundamental in a situation where we expect a very small fluctuation of power due by the very small amplitudes of gravitational-wave signals.

The antenna pattern

For the sake of simplicity, in the previous section we assumed a linearly polarized wave crossing face-on the interferometer. For a generic non polarized wave, coming from any direction with respect to the detector plane, the response of the detector depends on the polarization and incidence angles (Maggiore, 2008; J D Romano et al., 2017; Saulson, 2017).

We set the origin of the reference frame on the beam splitter, with z the axis orthogonal to the arms, according to the scheme in Fig. 1.10, let us define:

- \hat{l}_1 and \hat{l}_2 the unit vectors pointing along the two arms,
- \hat{r} the unit vector pointing to the sky, slanted with respect to the detector axis by polar angles (θ, ϕ) ,
- \hat{p} and \hat{q} unit vectors orthogonal to \hat{r} , defining the polarization angle ψ .

The antenna pattern for the two polarization states, given the detector and polarization basis vectors, are defined as (Maggiore, 2008; J D Romano et al., 2017; Saulson, 2017):




Figure 1.10: Scheme of an incoming wave, showing the polarization and detector basis used to define the antenna pattern of the interferometer.

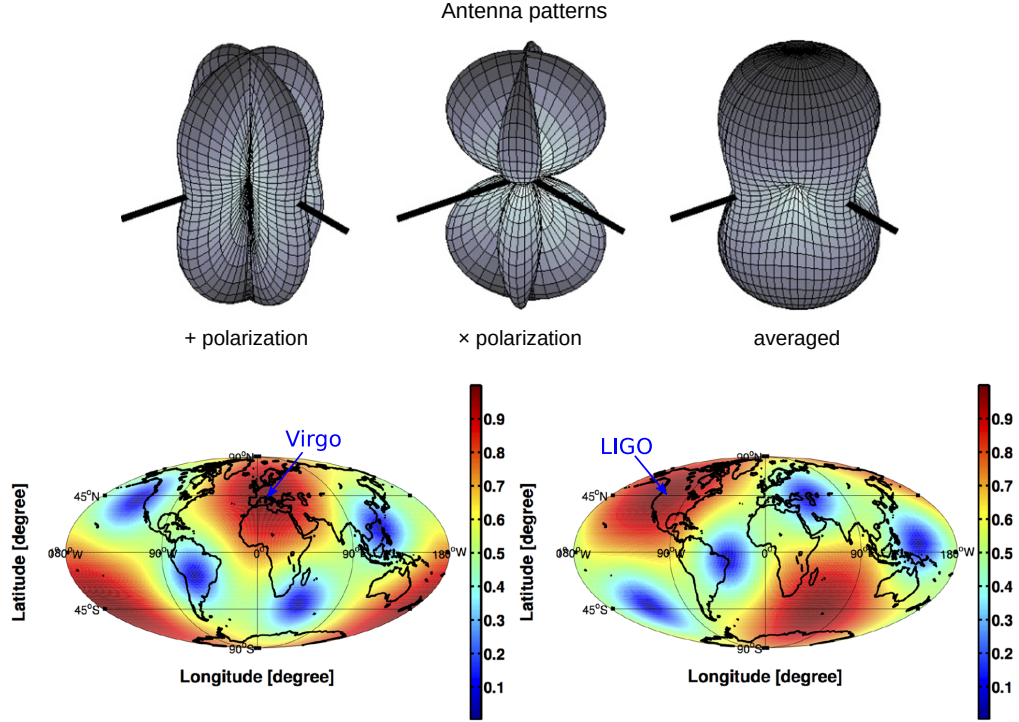


Figure 1.11: Top: antenna pattern for a Michelson interferometer, in the small antenna approximation, for +, \times and averaged polarization. The beam splitter is at the origin of the pattern, and the arms orientation are indicated by the solid black lines. Bottom: antenna patterns represented as colorbar plots on a Mollweide projection of the Earth, for LIGO and Virgo interferometers. The plot shows — at a fixed time — how a detector (at different positions on Earth) can respond to gravitational-wave signals, in function of the incoming direction in the sky J D Romano et al., 2017.

$$\begin{aligned} F^+(\theta, \phi, \psi) &= \frac{1}{2}(p_i p_j - q_i q_j)(l_1^i l_1^j - l_2^i l_2^j), \\ F^\times(\theta, \phi, \psi) &= \frac{1}{2}(p_i q_j + q_i q_j)(l_1^i l_1^j - l_2^i l_2^j). \end{aligned} \quad (1.62)$$

The detector output can then be written as $h(t) = F^+ h^+ + F^\times h^\times$, where the meaning of h^+ and h^\times is the same of those shown in Eq. 1.23.

The values of F^+ and F^\times are between 0 and 1 and tell us how the response of the detector is reduced by the misalignment between the incoming wave and the arms plane. It is important to remark that the direction of \hat{r} may change with time with respect to the detector frame, and the time dependence of the detector output includes the time variation of the sky and polarization angles.

For a network of interferometers on Earth, it is interesting also to relate their response to their position, since according to their distribution we are more or less sensitive to signals coming from certain directions. See Fig. 1.11 for an example given by the LIGO and Virgo detectors.




Figure 1.12: A scheme that shows the fundamental components of ground-based of second-generation gravitational interferometers. .

1.4.2 Fabry-Perot Cavity

Eq. 1.61 shows that the sensitivity of the detector depends on the length of the optical path $2L$. To increase it, each arm of the interferometer is usually equipped with a *Fabry-Perot cavity*, putting a semi-transparent mirror between the end mirror and the beam splitter: in this way the photons entering the cavity are reflected several times before being extracted and directed to the beam splitter, to be ultimately reconstructed with the light from the other arm. The cavities operate in resonant condition, where the light circulating inside will be in the form of standing waves (Maggiore, 2008; Saulson, 2017).

The characterizing parameter of a Fabry-Perot cavity is its *finesse* \mathcal{F} :

$$\mathcal{F} = \pi \frac{\sqrt{R_1 R_2}}{1 - R_1 R_2}, \quad (1.63)$$

with R_1 and R_2 being the reflectivity of the entrance and terminal mirror, respectively. The finesse defines how many times a photon can be bounced before being absorbed in the system because of imperfections in the reflectivity of the mirrors. The effective number of bounces is then $N_{eff} = 2\mathcal{F}/\pi$, while the effective arm length is $N_{eff}L$. The time before the photon is absorbed is $\tau = N_{eff}L/c$ and its inverse is the frequency cut-off of the cavity: $\omega_c = 1/2\pi\tau$. Hence, we obtain the effective cavity optical length (Maggiore, 2008; Saulson, 2017):

$$L_{opt} = \frac{2\mathcal{F}L}{\pi} \left[1 + \left(\frac{\omega_G}{\omega_c} \right)^2 \right]^{-\frac{1}{2}}. \quad (1.64)$$

This formula implies that, after the cut-off frequency of the cavity, the sensitivity of the detector degrades (Maggiore, 2008).

Two other important components are present in the ground-based detectors: a power recycling mirror set after the input laser source, which creates a cavity whose

purpose is to increase the light power circulating in the interferometer in order to reduce the noise; a signal recycling mirror at the output of the detector, right before the photodiode, which reflects back only the output light power, in order to increase its significance. A simplified scheme of a gravitational-wave detector, in Fig. 1.12, shows the described components.

1.4.3 Main noise sources

The sensitivity of a gravitational interferometer is determined by its noise sources. Some of the most important (Saulson, 2017) are shown in the left panel of Fig. 1.13 as design sensitivity of the LIGO detectors (Aasi et al., 2015):

- quantum uncertainty on the optical readout at the photodiode (i.e., the so-called *quantum noise*);
- seismic oscillations due to crustal activity on the test masses pendulums (seismic and Newtonian noise);
- thermal oscillations of the detector components (thermal and Brownian noise).

In the plot is shown also the noise by the scattering of the laser beam with the residual gas survived the vacuuming of the cavities.

In the right panel we find the estimation of the best sensitivity reached by the LIGO and Virgo detectors during the O3 run (Barsotti et al., 2018; Verkindt, 2021). The shape of the curves shows that in all three detectors exists a frequency region, roughly between 100 and 200 Hz, where the sensitivity is significantly better.

Usually, the sensitivity of the detectors is shown as the square root of the noise power spectral density $P_n(f)$, the so-called *strain sensitivity* $h_{strain}(f)$, with the dimensions of $\text{Hz}^{-1/2}$. A more detailed discussion on how these quantities are defined is given in Sect. 2.1.

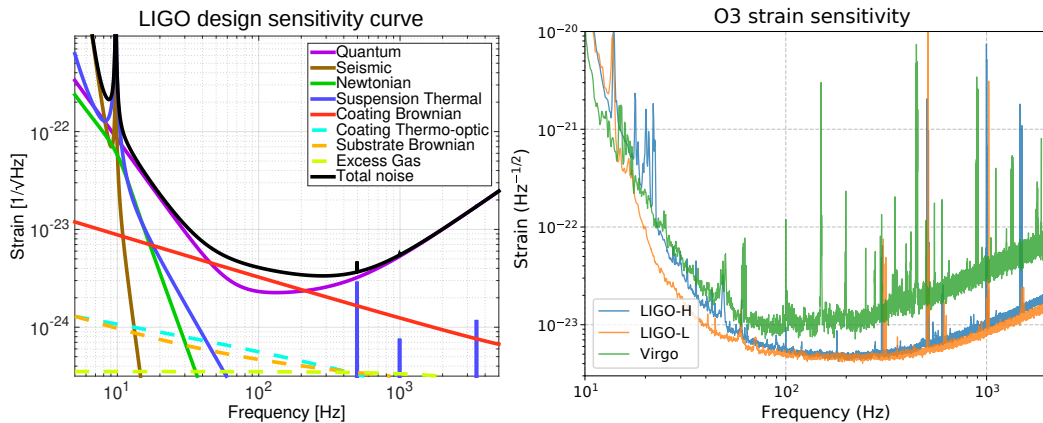


Figure 1.13: Left: LIGO design sensitivity curves, explicating the main noise sources (figure from Aasi et al., 2015). Right: plot of the O3 LIGO and Virgo best sensitivity curve estimates (Barsotti et al., 2018; Verkindt, 2021)

Chapter 2

Data analysis for Continuous Waves

One of the most interesting and promising classes of gravitational-wave sources is represented by Galactic non-axisymmetric rotating neutron stars. They emit periodic signals with frequency twice or equal their rotational frequency, and an amplitude that is expected to be below $\sim 10^{-26}$ - 10^{-25} (R. Abbott et al., 2021c, 2022a; Sieniawska et al., 2019). To give an element of comparison, the detected amplitudes from compact binary mergers reached amplitudes of the order of 10^{-22} - 10^{-21} , with sources hundreds of Mpc far from our galaxy.

In general, the data analysis in gravitational-wave searches plays a key role for any kind of source. In fact, given the faintness of the signals we observe, the noise fluctuations typically overcome the signals by orders of magnitudes. To be able to filter out the noise and isolate the signal from the detector data, a common approach has been to implement cross-correlation based filtering techniques.

Via the use of cross-correlation, the output of a detector is compared to a filter function that reproduces the expected signal shape, or, in case of stochastic background searches, to the output of another detector. The objective is to enhance correlated features (the modeled signals), discarding the supposedly uncorrelated noise. In the respective fields, it can be shown that these methods are the optimal way to define a detection statistics that, filtering out noise from the detectors, provide the best signal-to-noise ratio (SNR). For continuous waves, given the faintness of the sources, we need also to integrate the data over a long time period, typically of the order of months.

Since the efficiency of these algorithms depends strongly on the accuracy of the signal template taken from the modeled source, for continuous waves this method can be indeed applied when we know exactly the parameters of the emitting neutron star. Unfortunately, besides the ~ 3000 known pulsars (ATNF, 2021), it is expected from population studies that an order of a billion of undetected neutron stars populate the galaxy (Bisnovatyi-Kogan, 1992).

For unknown sources, it is impossible to derive the expected signal for a matched-filter analysis, so a wide set of possible templates should be generated. To generate a template we need:

- the intrinsic parameters of the expected emission, which -as we will see- are 2 in a first-order approximation (f and \dot{f});
- the Doppler effect due to the Earth rotation and orbital motion, which adds

to the parameters to take into account the two sky coordinates, in order to track down the observer-source direction;

- the 5 Keplerian orbital binary parameters (Dhurandhar et al., 2001; Leaci et al., 2017) for sources in binary systems.

For the search of yet unknown neutron stars, the combination between the lower amplitudes and the size of the parameter space makes this kind of analysis far more challenging. We are forced then to find a trade-off between the computational resources required (time, computation power, device memory capacity), to achieve a certain sensitivity, and their availability.

Summarizing then, searches for continuous waves can be broadly split into the following categories, which require a different data-analysis approach :

- *targeted* searches which look for known pulsars, whose parameters are known enough to run cross-correlation based coherent search algorithms (e.g. R. Abbott et al., 2021c);
- *directed* searches where a portion of the space parameter, notably a portion of the sky, is analyzed looking for unknown neutron stars where we expect to have a higher probability to find sources, like the Galactic Center (R. Abbott et al., 2022e; O J Piccinni et al., 2020) or supernova remnants (R. Abbott et al., 2021e);
- *all-sky* searches where no parameters are known and the search is performed spanning along all the sky directions (B. P. Abbott et al., 2017a, 2019b; R. Abbott et al., 2022a).
- other searches on more exotic sources, like dark matter particles (R. Abbott et al., 2022d) around spinning black holes, in particular boson clouds (R. Abbott et al., 2022b; Palomba et al., 2019).

2.1 General methods for data analysis

The data used for the analysis of gravitational-wave searches are usually in the form of strain data series, which come from the calibration process of the raw data of digitized electrical output from each detector readout system (Acernese et al., 2022; Sun et al., 2020).

It is important to remark that all fundamental methods for the analysis of gravitational waves start from the assumption that the noise follows a stationary Gaussian distribution over time. Because of this, after the calibration process, a series of cleaning procedures are used to minimize the non-stationarity features that are present in the noisy data (Vajente et al., 2020). In the O3 run a high rate of noise transients, usually called *glitches*, threatened to lower significantly the sensitivity of long-term searches like the ones for continuous waves and stochastic background. Hence, to exclude several glitches from the analysis, another cleaning procedure called *gating* has been used (Mata et al., 2021; Zweizig et al., 2021), consisting in vetoing the identified glitches setting to zero the time period interested

in the time series, with a smooth fade-out/in transitions of 0.25 seconds before and after.

At the end, the data we analyze have a form of a function of time $s(t) = h(t) + n(t)$, where $h(t)$ is the contribution given purely by the gravitational-wave signal and $n(t)$ is, in good approximation, the noise following a stationary Gaussian distribution.

We can characterize now the data coming from gravitational-wave detectors. Few key concepts are useful in signal analysis, which we will explained below:

Fourier transform. Given a function of time $s(t)$, its Fourier transform $\tilde{s}(f)$ is defined as (Saulson, 2017)

$$\tilde{s}(f) := \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{+\infty} s(t) e^{-i2\pi f t} dt. \quad (2.1)$$

The result of the Fourier transform is a complex function of frequency, containing all the information of the original time series. As we will see, the square modulus of the Fourier transform is proportional to the spectra of the time-dependent function.

If the transformed function is real, then the Fourier transform have the symmetry property $\tilde{s}(-f) = \tilde{s}^*(f)$.

Cross-correlation. In gravitational-wave analysis, the cross-correlation is widely used to measure how related two functions of the same variables are, as a function of a relative displacement. Taking for example two time series $s_1(t)$ and $s_2(t)$, it is defined as (Saulson, 2017) as

$$s_1 \star s_2(\tau) := \int_{-\infty}^{+\infty} s_1(t) s_2(t + \tau) dt, \quad (2.2)$$

with τ the offset between the two functions. If $s_1 = s_2 = s$ the self cross-correlation $s \star s(\tau)$ is called *auto-correlation*

Power spectral density. Another quantity to characterize the data in the frequency domain is their spectrum, defined as the Fourier transform of the autocorrelation of the time series (Saulson, 2017):

$$P(f) := \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{+\infty} s \star s(\tau) e^{-i2\pi f \tau} d\tau. \quad (2.3)$$

The power spectral density of a function is closely related with its *periodogram*, defined as (Saulson, 2017)

$$P_T = \frac{|\tilde{s}(f)|^2}{T}. \quad (2.4)$$

In the limit $T \rightarrow \infty$ converges to the power spectral density of $s(t)$. As we will see usually, to estimate the power spectral density with finite time series,

the data are split in several segments and the periodogram is averaged along each segment.

Since we are using real functions, $|\tilde{s}(f)|^2 = |\tilde{s}(-f)|^2$, we will consider only $f \geq 0$ and we define the single-sided power spectral density as $2P(f)$.

Assuming stationary noise, we can define the power spectral density of the noise $P_n(f)$ through the *ensemble average* of the noise components (Maggiore, 2008):

$$\langle \tilde{n}^*(f) \tilde{n}(f) \rangle = \delta(f - f') \frac{1}{2} P_n(f). \quad (2.5)$$

The ensemble average is the average computed over a statistical ensemble made by many realizations of the system. In practice we can consider the same detector at many different times as different realizations of the statistical *ensemble*¹, split the data in chunks of duration T_{coh} , computing the time-limited Fourier transforms from each chunk and average the periodogram over them.

In this case the approximation of the delta function for finite time, when $f = f'$, brings $\delta_t(0) = T_{coh}$ and we have

$$\langle |\tilde{n}(f)|^2 \rangle = \frac{1}{2} P_n(f) T \Rightarrow P_n(f) = 2 \langle |\tilde{n}|^2 \rangle \Delta f, \quad (2.6)$$

where the factor 2 comes from the fact that, being both noise and signal components of the output real functions, we are integrating only over positive frequencies and doubling the obtained value to have the *one-sided* spectral density.

This quantity has the dimensions of Hz^{-1} . The sensitivity of the detectors is usually expressed as strain sensitivity $h_{strain}(f) = \sqrt{P_n(f)}$ with dimensions $\text{Hz}^{-1/2}$.

2.1.1 Filtering out noise

Data analysis methods based on the use of the cross-correlation (Maggiore, 2008; Saulson, 2017) are widely used to extract various classes of signals from data whose noise level is higher than the expected gravitational-wave amplitude (for example, see B. P. Abbott et al., 2016; R. Abbott et al., 2021c,d).

We keep in mind that the data coming from the detector can be split in two independent contributions $s(t) = h(t) + n(t)$, from the signal and noise, respectively. The concept is that, cross-correlating $s(t)$ with a well defined function $H(t)$, from now on called *filter function*, the noise component $n(t)$ will not be correlated with the filter, not contributing to the cross-correlation integral; conversely, if the signal component matches with the filter, it will contribute positively. We can define then a detection statistic whose expectation value maximizes the signal-to-noise ratio SNR .

It can be shown that the integral in the time domain of the cross-correlation $s \star H(\tau)$ is equivalent to the following inverse Fourier transform (Maggiore, 2008; Moore et al., 2014; Zubakov et al., 1962)

¹An estimation on the minimum time needed to consider the detector as different realizations is ~ 50 s, derived in Joseph D Romano et al., 2015 computing the auto-correlation of the strain time series for a single detector, in the context of stochastic gravitational-wave background searches.

$$\begin{aligned}
Y(t) &= s \star H(\tau) = \int_{-\infty}^{+\infty} s(t)H(\tau - t)dt = \\
&= \int_{-\infty}^{+\infty} e^{i2\pi ft} \tilde{s}(f) \tilde{H}(f) df.
\end{aligned} \tag{2.7}$$

The optimum filter function \tilde{H} has the form:

$$\tilde{H}(f) = 2 \frac{\tilde{h}^*(f)}{P_n(f)}, \tag{2.8}$$

with signal template $\tilde{h}(f)$ generated from theory and $P_n(f)$ the noise spectral density defined in the previous section (Eq. 2.6)².

We have the so-called *matched filter* method, where the data are compared to a set of templates, in order to search for the one that gives an *SNR* above a certain statistical confidence. Coincidences from other detector network data are searched for to validate the detection.

The matched filtering method can be used for analyses where the signal can be well parametrized in the template, and whose parameters space is of an affordable size. This is the case of coalescence of compact objects (B. P. Abbott et al., 2016, 2017c, etc) or for known pulsars (R. Abbott et al., 2021c), but for blind searches of neutron stars other approaches need to be developed, as it will be better explained in Sec. 2.2.

Since the gravitational-wave signal is expected to be far lower than the instrumental noise, in the case of stochastic gravitational-wave background searches the data from two detectors are cross-correlated, exploiting the hypothesis that, being the signal by definition correlated in the two data streams, their respective noise is uncorrelated. As it will be shown in Sec. 3.1, with this approach, a specific optimum filter for the cross correlation of two detectors has to be derived. Thanks to the generality of the method, it will be explained also how the search for stochastic background searches, with anisotropic angular distribution, can be used also for point-like sources like rotating neutron stars.

2.1.2 Searching for periodic signals

Unlike short-lasting signals, like bursts or compact-binary coalescences, the search for persistent periodic signals has to take into account the Doppler effect from the orbital and rotational Earth motions.

In the case of a monochromatic source, if we had an isolated detector freely traveling in space, the signal would appear as a (time limited) delta function in the power spectral density. If it generates data series with pure Gaussian noise, given

²The role of P_n in the denominator of the filter comes from the fact that the noise we are considering is not *white*, i.e., it does not have a flat distribution. Being not white, the filter weights the statistics using the power spectral density of the noise, in order to give less significance to the worst frequency bands. The *SNR* in this way is increased giving more significance not only to the bands where the source is expected to have the most spectral emission via the template $\tilde{h}(f)$, but also to the most sensitive frequency regions of the detectors (Mitra et al., 2008; Saulson, 2017)

an observation time T_{obs} , we expect the SNR scales proportional to $\sqrt{T_{obs}}$ (Saulson, 2017). It would be enough then to gather data for a sufficient amount of time to have a detection: with sufficient statistics we would identify as a signal a narrow single spectral peak above the noise fluctuations, at the frequency of the incoming wave .

In ground-based detectors the signal peak is instead modulated by the Doppler effect. If the frequency of the signal is f_0 , $\mathbf{v}(t)$ is the observer velocity vector at the time t and $\hat{\mathbf{\Omega}}(\mathbf{t})$ is the source direction unit vector, the received frequency is given by the relation (Saulson, 2017):

$$f = f_0 \left(+ \frac{(\mathbf{v} \cdot \hat{\mathbf{\Omega}})(t)}{c} \right) \quad (2.9)$$

and hence the Doppler spread δf_{dopp} can be defined as

$$\delta f_{dopp} = f_0 \frac{(\mathbf{v} \cdot \hat{\mathbf{\Omega}})(t)}{c} \quad (2.10)$$

The maximum Doppler spread caused by Earth rotation, for a detector located at the equator, is $\delta f_{max}/f_0 = v_{rot}/c \sim 10^{-6}$ Hz, while for the orbital motion we have the more significant value of $\delta f_{max}/f_0 = v_{orb}/c \sim 10^{-4}$ Hz (Saulson, 2017). In this case, the spectrum of a single monochromatic peak is modulated in a series of sub-peaks around the carrier frequency f_0 . If the modulation frequency is f_m , the number of sub-bands is approximatively equal to the modulation index δ , defined as (Saulson, 2017)

$$\delta := \frac{f_0}{f_m} \frac{(\mathbf{v} \cdot \hat{\mathbf{\Omega}})_{max}}{c}, \quad (2.11)$$

where $(\mathbf{v} \cdot \hat{\mathbf{\Omega}})_{max}$ is the maximum value of the product over the modulation cycle.

The power will be divided by each of the sub-peaks, which will have a spectral amplitude $\sim 1/\sqrt{\delta}$ times the amplitude of the single peak in the unmodulated case. For the daily modulation ($f_m = 1.2 \times 10^{-5}$ Hz) there are ≈ 10 sidebands, causing a factor 3 amplitude loss; with the annual modulation ($f_m = 3.2 \times 10^{-8}$ Hz) there are $\approx 10^6$ sidebands and the spectral peak will be reduced by a factor of 10^3 (Saulson, 2017).

This would decrease dramatically the chances to detect a signal and, because of that, methods to avoid or reduce the Doppler effect are adopted in the various classes of search. For example, in targeted or directed searches the Doppler effect is corrected at the time series level for each of the directions we want to inspect (R. Abbott et al., 2021c,e; Saulson, 2017). When passing to the frequency domain, the modulation spread for the given direction in this way is eliminated and the signal can be fully recovered in all its original features (Saulson, 2017).

When the source parameters are unknown, like in all-sky searches, a method adopted to reduce the computational cost of the search has been the *semicoherent* approach, which implies splitting evenly the time series in segments with a certain length. This length, far smaller than the observation time, is called the *coherence time* (T_{coh}) of the semicoherent search. The analysis is performed on each time segment and, afterwards, the detection statistics are properly integrated. Since, in

the case of semicoherent searches, the loss in sensitivity with respect to the coherent optimal search is $\propto \sqrt[4]{T_{\text{obs}}/T_{\text{coh}}}$ (Frasca et al., 2005), the objective is to minimize the loss in sensitivity, for example maximizing the coherence time.

The natural frequency resolution of semicoherent searches is given by the inverse of the coherence time, that is $\delta f = 1/T_{\text{coh}}$. As we will see better in Sect. 2.2, an approach to bypass the Doppler spectral spread is to properly choose a coherence time, in order to have a frequency resolution large enough to be able to contain in a single bin all the frequency modulation width generated in that time (Astone et al., 2014). Since the Doppler spread, as shown in Eq. 2.10 depends on the frequency, in Astone, Frasca, et al., 2005; Frasca et al., 2005 it has been shown that, to satisfy this condition when moving to the frequency domain, T_{coh} has to be lower than

$$T_{\text{max}}(f) = T_{\text{orb}} \sqrt{\frac{c}{4\pi^2 R_{\text{orb}} f}} \approx \frac{1.1 \times 10^5}{\sqrt{f}} \text{ s}, \quad (2.12)$$

where we have neglected the effect of the daily modulation, and T_{orb} and R_{orb} are respectively the period and the average radius of the orbital revolution.

It will be shown in Sect. 2.2 that with this method the frequency resolution will be still fine enough to see a difference in frequency between the Fourier transforms of different segments, due by the daily and yearly Doppler effects. Collecting the most significant peaks from each segment in the frequency domain, along the full observation time, we will then be able to correct the frequency of those peaks to remove the Doppler shift. In this way we can retrieve the expected frequency-time evolution for a given sky direction.

The Doppler effect represents an issue when it comes to improve the sensitivity of the searches, but at the same time it allows to retrieve the source locations. The gravitational-wave interferometers behave like antennas, that is they cannot be oriented like a telescope, but receive at the same time signals from all over the sky directions; the Doppler phase evolution of a signal is an important signature that allows us to effectively aim at a sky region when trying to resolve a source from the noise.

2.2 The FrequencyHough pipeline

As already mentioned, data analysis for all-sky blind searches requires computational resources at present day unachievable to cover the full parameter space with coherent methods. Because of that, hierarchical algorithms (referred to as pipelines) are developed using the semicoherent approach.

So far several independent pipelines have been developed, actively used in all-sky searches (R. Abbott et al., 2022a). One of them, the *FrequencyHough* pipeline (Astone et al., 2014), uses an adaptation of the so-called *Hough transform* (Hough, 1962) pattern recognition method; I have recently modified the algorithm in order to use modern graphical processing units (GPUs) and deployed for the O3 run analysis in the framework of this PhD project (see Sect. 2.3 and La Rosa et al., 2021 for more details).

The pipeline is composed by several steps explained below and summarized in the flowchart in Fig. 2.1 (Astone et al., 2014).

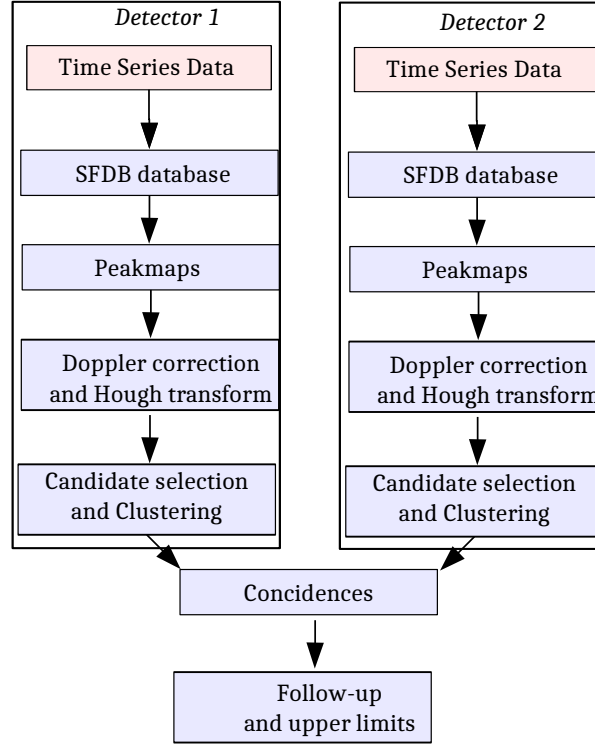


Figure 2.1: Flowchart showing the FrequencyHough search pipeline as exposed in Sect. 2.2, edited from (Astone et al., 2014).

Short Fourier transform database

The semicoherent pipeline starts from splitting the data time series with a certain time interval and compute the Short Fourier Transform (SFT) for each time segment. From the cleaned and calibrated frame files then a database of data in the frequency domain is generated using the Fast-Fourier Transform (FFT) algorithm with a certain segment duration, which we indicate with T_{coh} .

As we have seen in Sect. 2.1, the frequency resolution of a time-limited Fourier transform depends on the length of the time segment, following the relation $\delta f = 1/T_{coh}$. We have seen as well that the Doppler effect due to the Earth motion causes a spread of a spectral peak with a width δf_{dopp} depending on its frequency. For the FrequencyHough pipeline, the length of the segment is chosen in order to have a frequency resolution $\delta f_{dopp} < \delta f$. In this way a signal can be spread by the Doppler effect at most in two contiguous bins and the loss of spectral power is minimized.

For the searches performed so far, a database of four T_{coh} bands has been generated, i.e.:

$f(\text{Hz})$	$T_{coh}(s)$
10-128	8192
128-512	4096
512-1024	2048
1024-2048	1024

This choice is practical to ease the deploy of the search, but it causes a loss in its efficiency: we have defined a coherence time good for the maximum of the frequency interval of the dataset, for example, 4096 s for the band that ends at 512 Hz. But to analyze lower frequencies, let us say a small band of 10 Hz around 200 Hz, we could use a longer coherence time and have a better sensitivity, without worrying about a Doppler spread higher than our frequency resolution. To overcome this problem and increase the overall efficiency of all-sky searches, a new data infrastructure has been developed, called *Band Sampled Data* (BSD) (O. Piccinni et al., 2018). With the BSD data structure is possible to compute with a large adaptability the FFTs with the T_{coh} needed. Thanks to this, instead of defining only four different bands, we can define a higher number of shorter sub-bands with different coherence time, which will result then in general longer granting a gain in sensitivity. .

Either way, at the end of the SFT database generation process, the data of a several months long run will be organized in a series of Fourier transforms, each of them coming from a single time segment and hence identified by their initial time stamp t_i (see Fig. 2.2)

Peakmap

Once that the data in the frequency domain are organized in the way described above, we can start the search of the signal. Since we expect periodic signals, with a very small spindown for the typical coherence time of \sim minutes, we are looking for sharp peaks in the frequency domain.

«««< HEAD From each FFT $\tilde{h}_i(f)$, computing their squared modulus we have the periodogram $S_i(f) = |\tilde{h}_i(f)|^2$, where the index i runs over the N segments spanning the time period we are analyzing (e.g. the full observational run or a




Figure 2.2: Scheme of the SFDB construction: each segments of the time series is Fourier transformed and the set of short Fourier transforms constitutes the database.

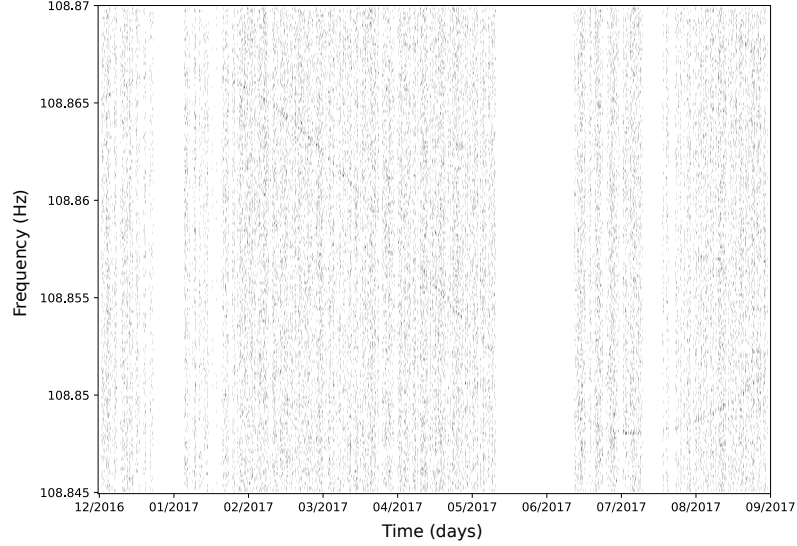


Figure 2.3: Example of a peakmap generated with O2 data with a pulsar injection. The frequency Doppler shift due to the yearly Earth motion is widely visible.

portion of it). The average spectrum $S_{AR_i}(f)$ is estimated with an auto-regressive method (Astone, Frasca, et al., 2005). The ratio

$$R_{ij} = \frac{S_i(f_j)}{S_{AR_i}(f_j)},$$

===== From each FFT $\tilde{h}_i(f)$, computing their squared modulus we have the periodogram $S_i(f) = |\tilde{h}_i(f)|^2$, where the index i runs over the N segments spanning the time period we are analyzing (e.g. the full observational run or a portion of it). An estimation of the average noise spectrum over the frequencies $S_{AR_i}(f)$ is derived with an auto-regressive method discussed in details in Astone, Frasca, et al., 2005. The principle is that the estimator of the spectrum should not be affected by high spectral peaks, and if the noise level varies, the estimator should be able to follow the noise variations.

The ratio

$$R_{ij} = \frac{S_i(f_j)}{S_{AR_i}(f_j)}, \quad (2.13)$$

»»»> 02d8ef2976e03f639d9fac6b6dbdce2d3f5f8091 represents a whitened estimation of the statistical significance of each frequency bin j from a given FFT.

Setting a threshold θ on the value of the ratio R_{ij} , we can select -from each FFT- the peaks above that value which are also local maxima. Since we are working with Fourier transforms discrete in frequencies, with a resolution $\delta f = 1/T_{coh}$, and since each FFT is uniquely identified by their time stamp, we will have in this way a collection of the most significant peaks in the frequency-time plane.

This collection is called *peakmap*, and it is usually represented as a Boolean matrix in the frequency-time coordinates, where each peak is identified as a 1 and the rest of the space is left as 0. To save memory usually it is expressed as a sparse matrix, that is, as a list of the coordinates of the peaks. If a periodic signal is present in the data, it will show a characteristic pattern in the peakmap, which

is analyzed by the FrequencyHough transform (see Figure 2.3 for an example of a peakmap with an injected signal).

It is worth to underline that if, from one side, if we selected all the spectral peaks above threshold, instead of only those which are local maxima, would imply a higher sensitivity because in this way we would consider peaks with a lower spectral amplitude than the maxima (see Sec. 2.2.3 for more details), but it would increase dramatically the computation and memory, since the number of peaks would be much higher. Moreover, using only the maxima, the robustness of the search against noise spectral disturbances is increased (Astone et al., 2014).

To improve the robustness of the search, the peaks collected can be weighted with values between 0 and 1, taking into account noise non-stationarity and the antenna pattern of the detector during the observation time. In this way, each peak, apart its time-frequency coordinates, has attributed a weight and will contribute accordingly to the final detection statistics.

Doppler Correction and FrequencyHough

Having identified a characteristic pattern in the frequency-time image represented by the peakmap, a pattern recognition algorithm like the Hough transform suits naturally for its identification.

In its original implementation, the Hough transform takes as input some images where a search for straight lines is done, returning as output an image of the portion of the analyzed space parameter, the so-called Hough map, where each point will have a certain value or number count, usually referred to as *peak amplitude* or simply *amplitude*. The values of the Hough map are related to the statistical significance that a line with certain parameters was present in the input image. Even though the Hough transform can be generalized in order to search not only for straight lines, but even a general functional form with a certain set of parameters (Ballard, 1981; Duda et al., 1972), the higher is the number of parameters of the search, the more its implementation can become complex.

Even though, with the proper frequency resolution, we are able to contain the peaks Doppler spread within a frequency bin, if we do not correct the yearly and daily frequency shifts, a signal in the peakmap will have a periodic pattern whose yearly Earth modulation is visible in Figure 2.3. This means that the two sky coordinates will be other search parameters, expanding further the computation power and the memory requirements to perform the search³. Since we expect that the source emission would follow a straight line in the frequency-time plane, it is indeed simpler to search for this pattern instead of a signal that is modulated both by the Earth rotation and orbital motion.

The choice done is then to correct the frequency shift of the peaks given a sky direction, in order to eliminate the Doppler effect. This is done for each sky direction $\hat{\Omega}$, according to the recorded velocity vector \mathbf{v}_j of the detector at the j th

³If the parameter space is 2-dimensional, the same will happen for the produced Hough map; with a 4-dimensional parameter space, where the two added parameters span over the full sky, the size of the output image becomes easily unmanageable.

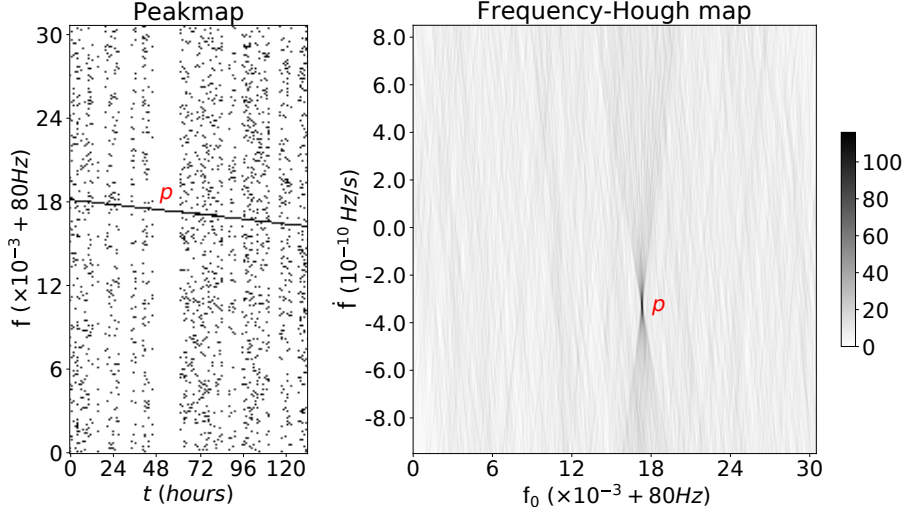


Figure 2.4: Example of a FrequencyHough map, from a portion of a peakmap taken from O2 real data and a straight line superimposed in order to simulate a continuous-wave signal. The parameters of the straight line labeled as p in the peakmap are retrieved as a point in the parameter space, where the lines coming from the transformation of the input space intersect together in the same area.

time, with the following equation:

$$f_{corr_j} = \frac{f_j}{1 + \hat{\Omega} \cdot \mathbf{v}_j}. \quad (2.14)$$

This is one of the most delicate parts of the pipeline, because an inaccurate correction of the Doppler effect can cause a severe loss in sensitivity. To search over the full sky, a proper grid of sky coordinates is defined in a way that, when we correct the peakmap, a loss of a signal is prevented (more details are given in Sec. 2.2.1).

Once the peakmap has been corrected for a given sky direction, it is possible to look for the characteristic pattern of an emitting neutron star. Assuming that the variation of the signal frequency is small enough, we can approximate, at first order on the frequency derivatives, the time law of the signal frequency as

$$f(t) = f_0 + \dot{f}(t - t_0), \quad (2.15)$$

where f_0 is the source frequency at the starting observation time t_0 and \dot{f} is the spindown term. With the above assumptions, the pattern of the signal in the Doppler corrected peakmaps will be a straight line, with parameters f_0 as the reference ordinate and \dot{f} as slope.

The adapted version of the Hough transform pattern recognition algorithm, the FrequencyHough transform, is well suited for the identification of continuous-waves signals. It searches for straight lines in the portion of the frequency-time space represented by the peakmap, returning an image of the frequency-spindown explored space parameter, so-called *FrequencyHough map*, where each pixel has an integer number count indicating the probability that a line with those parameters exists in the analyzed peakmap. An example of a FrequencyHough map generated from a

small peakmap coming from O2 data, with superimposed a straight line (to simulate a fake signal) is shown in Fig. 2.4. A detailed discussion on the Hough transform and its gravitational-wave implementation is given in Sect. 2.2.2.

Candidate selection

From the FrequencyHough map, the frequency-spindown bins with a high statistic significance are selected in order to be more thoroughly analyzed. These significant points of the parameter space are usually called candidates, or outliers.

The selection is done splitting the map in N frequency stripes and M spindown belts. This will identify frequency-spindown rectangular portions of the map, where the pixel with the maximum number count is selected and its amplitude and parameter space coordinates are recorded, together with the so called *critical ratio*, defined as (Astone et al., 2014)

$$CR = \frac{n - m(n)}{s(n)}. \quad (2.16)$$

Here n is the FrequencyHough map number count of the candidate selected, $m(n)$ the median of the amplitudes in the given map portion and $s(n)$ an estimator of the dispersion parameters, analogue to the standard deviation for the average in Gaussian distributions. It is defined as

$$s(n) = \frac{1}{C} m(|n| - m(n)), \quad (2.17)$$

where the normalization factor $C = 0.6745$ is defined in order to have $s(n) = \sigma_n$ if n follows a Gaussian distribution. The definition of CR with $m(n)$ and $s(n)$ comes indeed from the fact that because of non-Gaussian disturbances in the noise, we cannot use the average and standard deviation as estimators, introducing errors (Astone et al., 2014). If the noise is purely Gaussian, the critical ratio reduces to the distance from the average, in terms of sigma, of the pixel number count in the considered frequency-spindown portion of the Hough map. Being non perfectly Gaussian, we generalize the concept to have an estimation of the statistic significance of the selected candidate.

To have a schematic idea of the process for candidate selection, an example is given in Fig. 2.5 starting from the map from Fig. 2.4. The map is split in several rectangles identified by the frequency and spindown intervals we have chosen. The code goes through all rectangles, searching for the candidate with the maximum amplitude, and storing it in a database of outliers. The N and M parameters are chosen arbitrarily according to the amount of candidates we want to follow-up, for a total of $N_{cand} = N \times M$ candidates.

In a full sky analysis, a FrequencyHough map is computed for each Doppler corrected peakmap along all the points in the sky grid, and the selected candidates will be identified by their parameter space coordinates $\mathbf{c} = (\lambda, \beta, f, \dot{f})$, with λ, β being the ecliptic coordinates. To reduce the number of candidates to be followed up, a distance in the parameter-space is defined and candidates are clustered together if they are close enough.




Figure 2.5: Candidate selection scheme, which is shown for the example of Fig. 2.4. The matrix is split in vertical and horizontal stripes along the two dimensions. Each slot is delimited by the intervals in the chosen frequency and spindown values. When a slot is processed, a candidate is selected and the algorithm goes to the next one.

In particular, an Euclid distance is defined as

$$d = \sqrt{\delta_\lambda^2 + \delta_\beta^2 + \delta_f^2 + \delta_{\dot{f}}^2},$$

where the δ_X , with $X = \{\lambda, \beta, f, \dot{f}\}$ are the distances in terms of bins of the two candidates in the respective parameter-space dimension. When d is below a certain fixed value, for example 4 (Astone et al., 2014), the two candidates are considered to be coincident.

In general, in gravitational-wave searches, a network of detectors plays a crucial role to filter the noise and increase the search sensitivity. It is important then to find the coincidences between the sets of candidates coming from the detectors of the network. Similarly to what it is done when clustering the candidates of a single detector, supposing that noise disturbances can slightly alter the returned parameters of a signal, a coincidence is registered when two candidates coming from two detectors have a distance in the parameter space smaller than 3 bins (B. P. Abbott et al., 2017a, 2019b; R. Abbott et al., 2022a).

To reduce the cost of the follow up search, only a portion of the candidates are kept for the next steps of the analysis. A criterion is to set a threshold on the CR , which is directly related to the search sensitivity (see Sect. 2.2.3). Another way, which is adopted in the all sky analyses (B. P. Abbott et al., 2017a, 2019b; R. Abbott et al., 2022a), is to rank the candidates from each dataset according to their amplitude, and then the coincident candidate pairs for their parameters space distance. The strongest and closest pair of candidates will have a better rank. For each 0.1 Hz band of the frequency search interval, the first 8 candidates are selected and followed up (B. P. Abbott et al., 2017a, 2019b; R. Abbott et al., 2022a).

Follow up

The candidate follow-up consists in a more detailed analysis, following the same procedure as of the FrequencyHough transform, but refining the resolution of the parameter space around each selected outlier. The first step is the creation of a new peakmap, ± 3 bin wide around the frequency of the candidate and built with already Doppler corrected FFTs with a 10 times longer coherence time (R. Abbott et al., 2022a; Astone et al., 2014).

The peakmap is analyzed with a refined sky grid, which is ± 3 bin wide around the position of the candidate, where the residual Doppler effect is corrected. The FrequencyHough transform is performed as before, with a new candidate selection. For each of second-level candidates, another set of Doppler corrected FFTs is generated. Then, another peakmap is created in a 0.2 Hz range around the frequency outlier, and an histogram of the peaks is made by projecting the peakmap on the frequency axis. The follow-up candidates are vetoed according to the following criteria (R. Abbott et al., 2022a):

- considering the frequency of the candidate after removing the spindown effect and the Doppler correction, the candidate is discarded if it falls in a known instrumental disturbance or any of the hardware injections;
- if two coincident candidates, after the follow-up, exhibit a distance larger than 6 bins in the refined parameter space, they are removed;
- let us consider the 0.2 Hz peak histogram defined above and let us split it in sub-bands wide as ± 2 of the original frequency bins. In each sub-band the maximum values are identified and compared. If the sub-band where the candidate stays has a maximum which is the first or second higher among all the sub-bands, the candidate is kept, otherwise it is discarded
- let us consider again the above mentioned sub-bands of the 0.2 Hz peak histogram from the two detectors. Computing the critical ratios in those bands, if the CR of two coincident candidates, weighted by the detector noise level, differs by more than a factor of 5 (R. Abbott et al., 2022a), then candidates are discarded.

The surviving candidates undergo another iteration of the same procedure, but increasing further the coherence time by a factor of ten.

2.2.1 The sky grid generation

The construction of the sky grid is crucial for the accuracy and efficiency of all-sky searches. The rationale is that, due to the Earth Doppler effect, the frequency shift causes an uncertainty in the sky position. As we will see in this section, the proper construction of a grid depends on the frequency we analyze and on the coherence time we use.

The size of the sky grid is important for the management of the computational load of the FrequencyHough pipeline. For a typical analysis, the sky grid is computed on the last frequency of each of 5 Hz large peakmaps (1 Hz if $f < 128$ Hz) spanning




Figure 2.6: Number of sky points of the grid generated for the Doppler correction step. The frequency bands are obtained splitting the full search frequency range in bins with 1 Hz width for $f < 128$ Hz and with a width of 5 Hz for $f > 128$ Hz. The grid is always computed for the maximum frequency of the bin; the T_{coh} durations are the same shown in Tab. 2.2.

the search frequency range, and the number of sky points can go from thousands to hundreds of thousands points (Astone et al., 2014; La Rosa et al., 2021). In Fig. 2.6 we plot the number of sky points needed versus the maximum frequency of the analyzed peakmap. The four different datasets, with different T_{coh} , are shown separately with different colors. In the plot we also indicate, for each dataset, the total integrated number of points we need to process to complete the analysis of their respective frequency range. The quantity has been computed summing the size of the sky grids, generated for each sky map, along the full dataset.

By construction the analysis runs independently through all the sky points generated from a given peakmap. The deploy of the search over this wide parameter space represents the biggest technical challenge for this all-sky analysis. From one side, it is important to minimize the number of sky coordinates we have to inspect singularly, that is choosing the coarser sky resolution we can. On the other side we cannot allow a sky grid so coarse that could produce a miscorrection of the Doppler effect, causing a significant loss in sensitivity.

In fact, if the grid in the sky is too coarse, when we select a point of the grid, the performed Doppler correction will not be enough for a signal staying on the edge of the area defined by the selected point: in the corrected peakmap, it will not draw the expected straight line, but it will be still curved. On the other hand, if we refine the grid too much, we significantly increase the computation time needed to cover all the sky and, at the same time, we will have multiple candidates from different sky positions.

The formula for the natural sky grid, in terms of heliocentric ecliptic coordinates latitude β and longitude λ , is derived in the following way (Astone et al., 2014):

- considering two sources emitting at the same frequency f_0 , with same β and

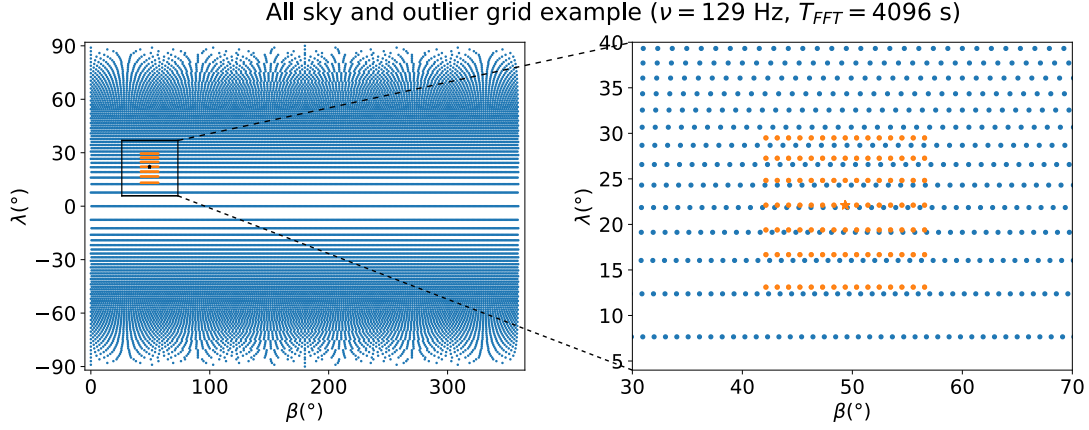


Figure 2.7: Example of a sky grid generated with $T_{coh} = 4096$ s and $f = 129$ Hz (blue dots). The red dots show a squared grid generated around a random position in the sky to simulate an outlier, with few degrees as length side, with $N_D = 56$ (Eq. 2.20). The plot in the right side is a zoom on the outlier grid.

a small separation in longitude $\delta\lambda$, the angular separation of the two signals translates in the detector in motion as a time delay $\Delta t \approx \delta\lambda/\omega_{orb}$, with ω_{orb} the Earth orbital angular speed⁴.

- The maximum Doppler shift the two signals can have is

$$\delta f_{max} = \frac{f_0 \omega_{orb} R_{orb} \cos \beta}{c} \delta \lambda, \quad (2.18)$$

where R_{orb} is the Earth orbit radius.

- Rearranging the previous equation in terms of $\delta\lambda$, and writing $\delta f = 1/T_{coh}$, the angular resolution along λ will be

$$\delta \lambda = \frac{c}{T_{coh} f_0 \omega_{orb} R_{orb} \cos \beta} = \frac{1}{N_D \cos \beta}, \quad (2.19)$$

where the so-called Doppler number N_D has been defined as (Astone et al., 2014)

$$N_D := f_0 T_{coh} \omega_{orb} \frac{R_{orb}}{c}. \quad (2.20)$$

- Following the same reasoning for two sources with this time same λ , but different β , the angular resolution along the latitude will be (Astone et al., 2014)

$$\delta \beta = \frac{1}{N_D \sin \beta}. \quad (2.21)$$

It is important to remark that, contrary to what happens for other sky tiling, such as the so-called *HEALpix* (Gorski et al., 2005), this grid is not uniform. As it is shown in Eqs. 2.19 and 2.21, the distance between two points along both directions depends on the latitude.

⁴For the sake of simplicity we are neglecting the Earth rotation.

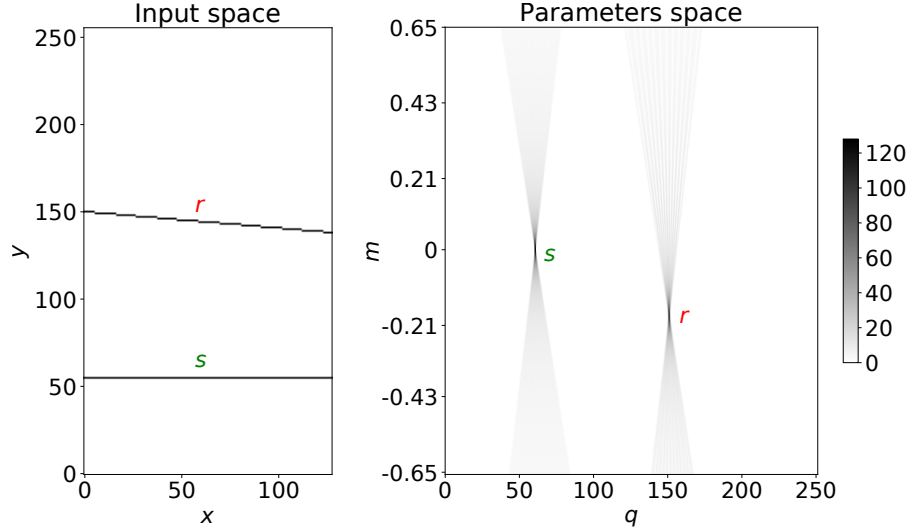


Figure 2.8: Example of a Hough map in its original implementation, having as input a picture with two straight tracks with different parameters, and as the output the Hough map showing their values in the parameter space.

For an all-sky analysis the grid can be computed covering all the directions, starting from three points: the poles $(\pm\pi/2, \lambda)$ and a point on the equator, e.g. $(0, 0)$. All other points follow using the distances given by Eqs. 2.19 and 2.21. Alternatively, as it is done for the follow up analysis of selected outliers, a grid can be built as a rectangle around a given sky point within a defined distance (see Figure. 2.7 for an example).

2.2.2 The FrequencyHough transform

The FrequencyHough transform is an implementation of the original Hough transform pattern recognition algorithm (Hough, 1962). The method has been conceived for the study of pictures of subatomic particle tracks in bubble chambers, and it was based on searching for straight lines in portions of the pictures small enough to have a curved track split, with good approximation, in line segments.

Considering as input an image where the line is formed by a series of co-linear black points on a white background, the process of the Hough transform consists in the definition of another image, called the Hough map, where the single black points with coordinates (x, y) are drawn as straight lines, whose slope and intercept are the input points coordinates. If we have a sequence of points forming a straight track with equation $y = mx + q$, they will be represented as a series of straight lines with varying intercept and slope. By simple Euclid geometry considerations, they will form a family of intersecting lines. The coordinates of the incidence point will be the values of the parameters (m, q) of the original line.

We can identify then the Hough transform space as the parameter space and the input one as the coordinates space; see Figure 2.8 for an example.

Even though the original version was identifying linear patterns in a 2D image, it is possible to generalize the Hough transform with any N -dimensional manifold as input space and an M -dimensional parameter space as output, searching for any

kind of curve with a set of M parameters (Ballard, 1981; Duda et al., 1972).

In the real applications, we always work with digitized images, with discrete coordinates and parameter space having resolutions independent to each other. The Hough map will be then a 2D histogram, where each line falls in a series of pixels. When the lines cross the same pixel, that pixel will record how many of them entered in that parameter space region. The number count amplitude in the map will give the significance of a portion of the parameter space: keeping the example of the pictures of black tracks on a white background, if in the input image there are many colinear pixels on 1, many lines in the Hough map will intersect in the same parameter space pixel, giving a higher number count. The Hough transform in general works also with input images whose pixels do not have only binary values, but also grey scales. The degree of grey, between 0 and 1 represents the weight of that peak when it will be transformed. This means that a peak in the input space, instead of being directly transformed as a straight line of +1 in the Hough map, will undergo to a weighted transform, returning values between 0 and 1 according to the grey scale of the peak.

The implementation of the Hough transform for continuous gravitational-waves searches, the FrequencyHough, starts from the Doppler corrected peakmap, as we said in Sect. 2.2. This will be the input space of the transform: the frequency-time plane of the selected peaks. The output parameter space will be as well 2-dimensional: the frequency time-law of the emitting neutron star is approximated at first order as a straight line with equation

$$f = f_0 + \dot{f}(t - t_0), \quad (2.22)$$

where the *reference* frequency $f_0 = f(t = t_0)$ and the first frequency derivative \dot{f} , the spindown, are the parameters; t_0 is an arbitrary reference time that for the sake of simplicity will be set to 0 from now on.

Similarly to the original implementation, the FrequencyHough transforms each point in the peakmap with coordinates (t, f) in a straight line in the (f_0, \dot{f}) parameter space, following the relation (Antonucci et al., 2008)

$$\dot{f} = \frac{1}{t} (f_0 - f). \quad (2.23)$$

Since we are working with discrete quantities, a point in the peakmap will cover a frequency range given by $\Delta f = 1/T_{coh}$, and hence it will be naturally mapped by the Hough transform as a stripe between two parallel straight lines:

$$\dot{f}_{\pm} = \frac{1}{t} (f_0 - f \pm \Delta f/2), \quad (2.24)$$

keeping the resolution in frequency naturally given by the FFT we used to build the peakmap.

The precision on the parameters we want to recover can be increased enhancing the resolution of the parameter space, thus increasing the size of the Hough map, at the price of a bigger memory and computation load. For example about the frequencies, the resolution of the Hough map for the f_0 parameter is enhanced by a factor $\epsilon_f = 10$ with respect to the natural frequency resolution $\Delta f = 1/T_{coh}$:

$\Delta f_0 = \Delta f / \epsilon_f$. With this choice the signal pixel amplitude is improved by 8% (Antonucci et al., 2008), thanks to the fact that, with a higher resolution the signal is better localized, instead of being averaged on a wider range by the effect of the discretization. For the spindown, the natural resolution is $\Delta \dot{f} = \Delta f / T_{obs}$, where T_{obs} is the full observation time, and increasing it does not bring a significant improvement in precision and sensitivity. Especially with the enhancement in the frequency resolution, a differential method to implement the FrequencyHough transform has been developed (Antonucci et al., 2008), which in the latest versions of the algorithm has been updated to run as most as possible in parallel (See Sect. 2.3 for more details and La Rosa et al., 2021). The process can be summarized as follows:

- the map is created as a matrix with size given by the frequency range of the peakmap, the range of spindown we want to inspect and the resolutions Δf_0 and $\Delta \dot{f}$ (typically the size is on the order of $10^5 \times 10^2$ bins, for a memory occupation between 7GB to 0.5GB, depending on the frequency range analyzed ⁵);
- for each spindown bin, the values of the matrix elements corresponding to the relation $f_0 = f - \Delta f / 2 - t \dot{f}$ (Antonucci et al., 2008) are found and incremented by 1. Using the set of elements raised by one, the elements satisfying the relation $f_0 = f + \Delta f / 2 - t \dot{f}$ are decremented by 1. This is called *differential map* and it is the core of the FrequencyHough map computation. At this point each row of the map will be a series of positive and negative integers coming from the superimposition of the left (+1) and right (-1) edges of the stripes defined above;
- to have the final map, each row of the matrix is, in parallel, cumulatively integrated along the frequencies with efficient vectorial functions.

An example of FrequencyHough transform is shown in Figure 2.4, where a portion of a peakmap from the O2 run data has been analyzed with a superimposed straight line to simulate a signal.

2.2.3 Search sensitivity

The estimation of the sensitivity of the FrequencyHough pipeline comes from the two steps of the analysis where data are selected or rejected on the basis of a threshold: the peak amplitude threshold θ chosen to build the peakmap, and the CR threshold on the candidates selected from the Hough map. The choice on both thresholds is done looking for the best trade-off between the affordability of the search from the computational point of view, the increase of false-alarm rate and, most important, the loss in sensitivity.

⁵For higher frequencies, shorter FFTs are used, and hence the resolution both in frequency and spindown is larger, bringing to smaller Hough maps for the same parameter space.




Figure 2.9: Plot from Astone et al., 2014 showing the behavior of Φ (Eq. 2.29) versus θ , for various values of spectral amplitude λ . The red dashed lines are computed for the simple above threshold peaks selection, the blue continuous ones considering also the local maxima condition. The plot shows how the difference between the two conditions in terms of expected CR is faint, granting a higher robustness of the analysis and a reduced computational load with a small sensitivity loss. Moreover, it shows how the choice of $\theta = 2.5$, slightly after the maximum, does not cause a strong reduction in sensitivity in the expected situation where the spectral amplitudes are small.

Choice of θ

In the hypothesis of only random Gaussian noise from the data, the Hough map pixel amplitude follows a binomial probability distribution, with average number count $\langle n \rangle = Np_0$ and variance $\sigma_n^2 = Np_0(1 - p_0)$, where N is the number of FFTs used to create the peakmap. The probability p_0 , defined as the probability to have a local maximum above the threshold θ due to the noise can be expressed by (Astone et al., 2014):

$$p_0 = e^{-\theta} - e^{-2\theta} + \frac{1}{3}e^{-3\theta}. \quad (2.25)$$

Following the definition of critical ratio given in Eq. 2.16, in this case it can be written as

$$CR = \frac{n - Np_0}{\sqrt{Np_0(1 - p_0)}}. \quad (2.26)$$

Let us define then the measured spectral amplitude of a signal, in a FFT with coherence time T_{coh} , as

$$\lambda = \frac{4}{T_{coh}} \cdot \frac{|H(f)|^2}{S_n(f)}, \quad (2.27)$$

with $H(f)$ the Fourier transform of the segment and $S_n(f)$ the detector noise spectral density. Lastly, the probability to select a signal peak with spectral

amplitude λ is (Astone et al., 2014):

$$p_\lambda \approx p_0 + \frac{\lambda}{2} \theta (e^{-\theta} - e^{-2\theta} + e^{-3\theta}). \quad (2.28)$$

The choice of the best value of θ should maximize the expectation value of CR , which we will write as \overline{CR} , when a signal with spectral amplitude λ is present (Astone et al., 2014):

$$\overline{CR}(\theta, \lambda) = \frac{N(p_\lambda - p_0)}{\sqrt{Np_0(1 - p_0)}} = \sqrt{N}\Phi(\theta, \lambda). \quad (2.29)$$

Computing and plotting the value of the function Φ for various values of λ between 0.1 and 0.6, in function of θ (see Fig. 2.9 from Astone et al., 2014), we can see that for low spectral amplitudes, the expected CR changes according to $\theta \in [1, 3]$, and that $\theta = 2.5$ is a good value for the chosen threshold. Note that this threshold is slightly higher than the maxima of the function Φ . This causes a heavy peaks number reduction in the peakmap, reducing the computational cost of the Hough transform, with a negligible loss in sensitivity ($\sim 1\%$) (Astone et al., 2014).

It is important also to remind that we are selecting only the local maxima of the peaks in each FFT, instead of all the ones above the threshold. This causes a further sensitivity loss ($\sim 5\%$) with the advantage of a $\sim 9\%$ computational time reduction and, most importantly, gaining an increased robustness against disturbances, which could generate several peaks, thus affecting many neighbor frequency bins.

Choice of CR_{thr}

The second parameter that affects in a significant way the sensitivity of the search, is the threshold on the CR of the candidates we want to keep for the follow-up analysis, from now on indicated as CR_{thr} .

Typically, the number values of pixels in the FrequencyHough maps, in an analysis spanning the time scale of an observing run with duration 9-12 months, is $\geq 10^2$. Then, we can use the Gaussian approximation to the binomial distribution to express the probability to have the amplitude n in a given pixel:

$$P_{\theta, \lambda}(n) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(n-\mu)^2}{2\sigma^2}}, \quad (2.30)$$

with $\mu = Np$, $\sigma^2 = Np(1 - p)$ and where $p = p_0$ in the only noise assumption (see Eq. 2.25) and $p = p_\lambda$ in presence of a signal (Eq. 2.28).

Defining n_{thr} as the threshold on the number count that would be necessary to select a candidate, it can be shown that the false alarm probability in the no-signal hypothesis can be expressed as a function solely of CR_{thr} (Astone et al., 2014),

$$\begin{aligned} P_{fa}(n_{thr}) &:= \int_{n_{thr}}^{\infty} P_{\theta, 0}(n) dn = \\ &= \frac{1}{2} \operatorname{erfc} \left(\frac{n_{thr} - Np_0}{\sqrt{2Np_0(1 - p_0)}} \right) \\ &= \frac{1}{2} \operatorname{erfc} \left(\frac{CR_{thr}}{\sqrt{2}} \right), \end{aligned} \quad (2.31)$$

where erfc is the complementary error function. This equation shows also that candidate selection condition on the Hough amplitude translates to a condition on the CR : being N and p_0 fixed parameters, the maximum n in a portion of the FrequencyHough map will be also the maximum CR .

Then, as the candidates are selected with the highest CR in N_{cand} portions of the FrequencyHough map, in the no-signal hypothesis we have $P_{fa} = N_{cand}/N_{tot}$. In this case, N_{cand} is the number of selected candidates, while N_{tot} is the total number of points in the search parameter space. With this consideration, we can then write, inverting Eq. 2.31:

$$CR_{thr} = \sqrt{2} \text{erfc}^{-1}(2P_{fa}) = \sqrt{2} \text{erfc}^{-1}\left(2\frac{N_{cand}}{N_{tot}}\right), \quad (2.32)$$

and we can choose arbitrarily a CR_{thr} to have a certain level of false alarm probability, considering the availability of computing resources and time for the follow-up.

Sensitivity and upper limits

Let us consider a confidence level Γ . We can define as search sensitivity the minimum signal amplitude that would produce a candidate in a fraction $\geq \Gamma$ of a statistical *ensemble* of the experiment. As we will see, the sensitivity depends only on the thresholds chosen to contain the computational cost of the search.

Let us define the probability to select a candidate with a count $n > n_{thr}$, when a given spectral amplitude λ from a signal is present

$$P_{\lambda, n > n_{thr}} = \int_{n_{thr}}^{\infty} P_{\theta, \lambda}(n) dn, \quad (2.33)$$

where the probability density function in the integrand is given by Eq. 2.30. The value of the sensitivity is given by imposing

$$\begin{aligned} P_{\lambda, n > n_{thr}} &= \int_{n_{thr}}^{\infty} P_{\theta, \lambda}(n) dn = \\ &= \frac{1}{2} \text{erfc}\left(\frac{n_{thr} - Np_{\lambda}}{\sqrt{2Np_{\lambda}(1-p_{\lambda})}}\right) = \Gamma, \end{aligned} \quad (2.34)$$

where p_{λ} is given by Eq. 2.28.

Solving the equation in the small signal approximation, the minimum detectable strain amplitude, given the detector spectral power $S_n(f)$, is (Astone et al., 2014)

$$h_{0min} \approx \frac{4.02}{N^{1/4}\theta^{1/2}} \sqrt{\frac{S_n(f)}{T_{coh}}} \left(\frac{p_0(1-p_0)}{p_1^2}\right)^{\frac{1}{4}} \sqrt{CR_{thr} - \sqrt{2} \text{erfc}^{-1}(2\Gamma)}, \quad (2.35)$$

where $p_1 := e^{-\theta} - 2e^{-2\theta} + e^{-3\theta}$ and $N = T_{obs}/T_{coh}$ is again the number of FFTs used for the analysis. This value gives an estimation of the signal amplitude the search is able to detect.

For a 1 year long observational run, using $\Gamma = 0.95$, $\theta = 2.5$, $T_{coh} = 4096$ (used for all-sky analyses in the $[128 - 512]$ Hz frequency range) and $\sqrt{S_n} \sim 10^{-23}$, as

an order of magnitude value for the sensitivity curves for *O3* in this region (see Fig. 1.13 in Sec. 1.4), the latter equation becomes

$$h_{0_{min}} \sim 8 \times 10^{-26} \sqrt{CR_{thr} + 1.645}. \quad (2.36)$$

To complete the information on the minimum detectable h_0 , for the *O3* search, CR_{thr} has been chosen equal to 5 (R. Abbott et al., 2022a), returning $h_{0_{min}} \sim 1.5 \times 10^{-25}$ in the highest sensitive frequency region.

For the *O3* all sky search for isolated neutron stars, the definition of used upper limit has been Eq. 2.35 with, in place of CR_{thr} , the maximum CR retrieved from the candidates in each 1 Hz frequency band has been considered. For each band, the final upper limit is the worse between those computed separately for the LIGO-Hanford and LIGO-Livingston detectors (R. Abbott et al., 2022a)⁶. In the all sky analyses for previous runs (B. P. Abbott et al., 2017a, 2019b), the upper limits were computed directly through injections of simulated signals. The method explained above produces more conservative upper limits with a far lower computational effort (R. Abbott et al., 2022a).

2.3 The General Purpose GPU (GPGPU) computing project

In recent years, the so-called Moore's law on the number of transistors in integrated circuits started to show its limits. From one side, the physical limits are on the way to become too expensive to overcome, on the other side the computing paradigm has changed to face the need for higher and higher computational power in many fields, from science to industry, to many everyday life applications.

One of these fields is the videogame industry. Since the '70s of the last century, dedicated computer hardware were developed to be able to draw 2-dimensional scenes at a high frame rate. In the '90s, with the birth of 3D-graphics it became mandatory to demand all the rendering of 3-dimensional scenes to devices with a specific architecture: the Graphic Processing Units (GPU), structured as integrated circuits on a board connected to dedicated memory, with various interfaces and ports.

They have been created for the development of videogames, to solve the computing problem consisting on the management of different pixels on a screen, deciding what they should show to form images, taking into account the user interaction, the neighbor pixels and the physical parameters of the depicted image. Everything has to be done in real time, at the typical screen frame rate of 60 Hz. To be able to fulfill these functions, a high number of processors is necessary, which can take care independently of different tasks and, using dedicated memories, communicate each other in a very efficient way. This is the basic description of a GPU: a device created and developed to do efficiently computations with an extremely high degree of parallelism.

⁶Being the Virgo detector less sensitive (see Fig. 1.13), it has not been considered for the upper limit computation, as the data from it were used for validation purposes.

The ever-increasing demand of realistic graphics has brought to have a very fast technology evolution in the last 20 years, with, at present day, GPUs having from tens to hundreds of cores and memory capacity reaching fast order of gigabytes. This caused an explosion of the processing power in terms of floating point operations per seconds (see Fig. 2.10), with videogames able to reproduce on screen even very detailed physical simulations. It became natural then the idea to apply the use of GPUs in the most different fields, scientific research included, giving birth to the General Purpose GPU (GPGPU) computing.

A typical GPU has a number of cores far higher than any multi-core CPU system, but with single core clock rate far lower than the CPU ones. The high core number makes GPUs far more efficient in parallel computing, while CPUs are still the best option to perform sequential calculations. A good approach on many problems is then to split the tasks of a code, demanding serialized operations to the CPU and reserve the GPU only for well parallelizable problems (see Fig. 2.11).

The computing paradigm that exploits the features of the GPUs is called *stream processing* or *data flow*. The programs are modeled as graphs of operations (in general, not dependent to each other) where the data *flow* through. The key concept of this paradigm is the vectorial approach. The data are organized as N -dimensional arrays reproducing effectively the algebra of tensors, and the operations transform the data applying an operation over all the elements of the tensor. Operations and data are distributed over the core and memory topology of the GPU by specific functions.

The importance of the vectorial approach is crucial first of all from the logic point of view: we want to use the GPU on a well parallelizable problem, the most efficient approach then is to have an array of data and apply the same instruction on multiple data at once. Second of all, the devices we use are built to exploit this paradigm: the GPU hardware is optimized to work with an uniform stream of data and operations that are spread in parallel over many cores, introducing sequential instruction will result in sub-optimal bottlenecks.

Another key element is the memory management. Every memory transfer




Figure 2.10: Trend of GPU and CPU FLOPS capability in the last fifteen years, whose values are taking from specifications publicly released by manufacturers.

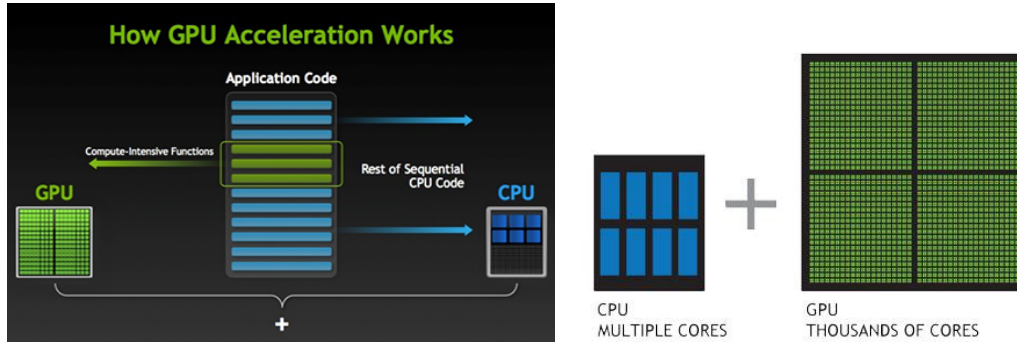


Figure 2.11: Scheme of joined CPU-GPU computation: the sequential part of a program is more efficiently when it is managed by the CPU, while the GPU manages the massive parallel computations (image from [nvidia.com](https://www.nvidia.com)).

introduces overhead, but while the bandwidth within the GPU is very large and the transfers are very fast and efficient, communicating with the storage or the system memory is highly inefficient. It is crucial then to load all the data possible at once, elaborate them and save the result at the very end of the process. If the data is much larger than the available memory, normally they are loaded in *batches* and the results are reconstructed after. Then it is worth to remark that, even though the raw power of the chipset is important, in data science the memory capacity can be decisive to define the efficiency of a code.

2.3.1 GPGPU with Tensorflow

Nowadays, thanks to the very low ratio between cost and computational power of the modern devices, the use of GPU for data science is an established standard, with devices developed specifically for the purpose, reaching thousands of cores and tens of GB of memory. Aside the hardware, also the software tools to develop codes using GPUs became more and more efficient and accessible, with the creation of dedicated libraries for the C and C++ programming languages. The two most important of them are CUDA (The NVIDIA Corporation, 2021) and OpenCL (The Khronos Group, 2021).

The *low-level* characteristics of C and C++ makes them among the best choices to optimize the efficiency of the developed codes, but on the other hand they require a time cost for the development, making the approach to GPU programming very steep. Even though the libraries help the task of the GPU resources management, it is still necessary some fine tuning that can result to be troublesome in the development process. To solve this issue, in 2015 Google released TensorFlow (The TensorFlow Authors, 2021) under the Apache 2.0 open source license: a new *high-level* framework for GPGPU programming, based on Python, but written in C++ using the CUDA libraries. TensorFlow was developed originally for machine learning, but as it will be shown in this Thesis, it is complete and general enough to be used with various purposes.

The APIs of TensorFlow consist in a set of CUDA-C++ functions covering a wide range of possible operations (tensorial products, convolutions, etc). With an approach and a syntax similar to the Scipy-Numpy libraries (Scipy developers, 2021),

```

1 import tensorflow as tf
2 N = 100
3
4 # Build a graph with:
5 # two matrices with random values;
6 matrix1 = tf.random_normal((N,N))
7 matrix2 = tf.random_normal((N,N))
8 # a matricial product between them.
9 product = tf.matmul(matrix1, matrix2)
10
11 # Launch the graph in a session,
12 sess = tf.Session()
13 # evaluate the tensor 'product'
14 sess.run(product)

```




Figure 2.12: Example TensorFlow code and its graph.

these functions can be called in a Python code, allowing an easier development and using an overlay of functions that, being written in a low-level language, are already well optimized.

Using the data-flow paradigm, a TensorFlow code builds a graph using tensors which go through a series of operations. The operations are the nodes of the graph, while the tensors, going from a node to another one, are the edges. Normally, the graph is first built and after executed by a specific call. However, in the recent versions 2.x, the so-called *eager execution* has been implemented to aid the development of new codes: the nodes are executed at the same time as they are added at the graph.

An example of a TensorFlow code, and the resulting graph, is shown in Fig. 2.12. The code produces two matrices with random values and then computes the matrix multiplication between them. It can represent a very useful benchmark to compare the performance and cost of different CPUs and GPUs on this kind of tasks.

In Fig. 2.13 the example code and an equivalent script written with Numpy have been executed, changing the parameter N (that is, the size of the generated matrices) from 3000 to 14000. The computing time as a function of the number of pixel is shown for the GPUs with TensorFlow using 32 and 64 bit data on a NVIDIA Tesla V100 (a GPU for data centers with 16 GB of memory and 5120 cores The NVIDIA Corporation, 2022). For the CPU side, the Numpy code, using 64 bit data with multithreading enabled, has been ran on an Intel Xeon E5-2650 v4, with 24 execution threads and a maximum clock of 2.9 GHz. The difference of more than one order of magnitude in the computing time for the 64 bit case shows

how the use of GPUs (in particular of the data center kind) can significantly boost the performances of numerical computations. The GPUs work natively on 32 bit data, so if it is possible to work on single precision, the performance improvement can be even larger.

To complete the picture, the same test has been ran on a set of different CPUs and GPUs, comparing performances and cost of the devices (see Fig. 2.14). Keeping in mind that the prices can vary significantly over time, the general trend that results from the benchmark is that with GPUs, when the problems to solve can be efficiently parallelized, is possible to run analyses one order of magnitude faster than on multi-core CPUs, with a fairly reduced cost. Moreover, with CPUs the performance on this kind of problems does not look strictly dependent on the architecture and computing capability; on the GPU side, even if desktop devices like the GTX series show similar execution times, the far larger memory availability of the data center series, like the Tesla, is a key factor when a large amount data has to be handled.

2.3.2 GPU FrequencyHough, details and performances

The GPU porting of the FrequencyHough algorithm and candidate selection (La Rosa et al., 2021) has the purpose of the application of GPGPU methods on the all-sky search for periodic gravitational waves from isolated neutron stars. The original version of the algorithm has been written in Matlab (Mathworks, 2021) and is based on the SNAG toolbox (Frasca et al., 2018), while the GPU-enabled version has been developed in Python, using the TensorFlow APIs updated up to

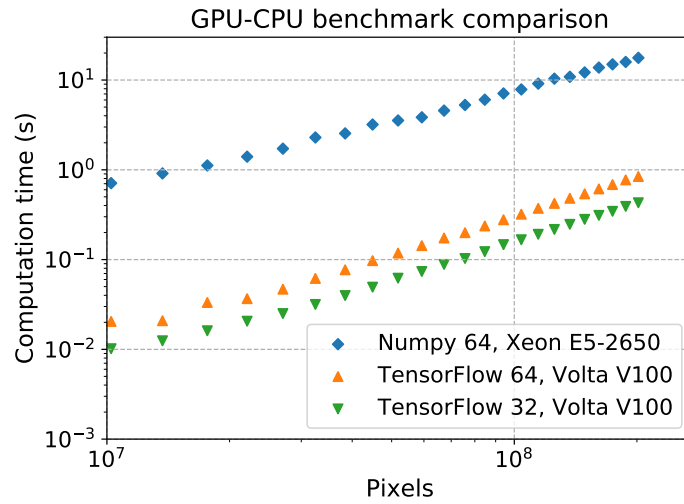


Figure 2.13: A benchmark based on the example code using two devices: an Intel Xeon E5-2650 v4 CPU system with 12 2.90 GHz execution threads, and a NVIDIA Volta V100 with 5120 cores at 1500 MHz. The test is based on the generation of two random 2-rank arrays and a matrix multiplication between them. The green and orange dots show the benchmark conducted, respectively, with 64 and 32 bit data with TensorFlow, the blue ones come from a test with 64 bit data with Numpy. The test with the CPUs has been performed keeping the multithreading enabled as by default for Numpy.




Figure 2.14: A set of tests using the example code on various different GPUs (left plot) and CPUs (right plot). For each device the benchmark of Fig. 2.13 has been performed and the best value of execution time, divided by the number of pixels has been registered. This has been done to take into account the possible different behavior between different systems and report a comparable quantity on the best performance reached by each used device. Aside to this, in the same frame, the market price at the time of writing is shown to give a general idea of the performance achievable in relation to the monetary cost.

the version 1.13.

The data used come from `.mat` files called *input files* containing Matlab structures with the peakmaps and various search parameters (observation time, frequency resolution, detector velocity for the Doppler correction, etc), coming from the previous steps of the search. The target of the project was to apply the new version of the Hough transform on all sky searches over the full detector sensitivity bandwidth, in order to produce and store the candidates for the follow-up analyses. Because the difficult integration of TensorFlow with Matlab, during the development of the new version, it resulted more efficient to write a code that, once loaded a peakmap, took care of all the steps of the analysis until the candidate selection step: sky grid generation and, for each sky point, Doppler correction, peakmap weighting, FrequencyHough transform and candidate selection.

The sky grid generation and peakmap manipulation parts are an exact transposition from the original Matlab code, using Python with Scipy libraries. The generation and analysis of the FrequencyHough map to produce the candidates has been implemented to run on GPU, being the core of the pipeline and the most computationally demanding part.

The original version was using a mix of vectorial operations and `for` loops, therefore, to improve the efficiency on GPUs, it was necessary first to fully vectorize the code⁷.

As said in Sect. 2.2.2, the FrequencyHough map is defined by the relation

$$f_H = f - \dot{f} \cdot t. \quad (2.37)$$

⁷Perform vectorial operations means to apply an operation on a complete array *at once*. Normally, in high level programming languages, to achieve this we need to use built-in functions optimized for this purpose. In this treatment, the term *vectorize* is used to mean that the optimization process goes through writing a code that uses only vectorial built-in functions of the framework used, in this case TensorFlow.

For each value of the spindown \dot{f} , a row of the Hough map is defined as the histogram of the values of the above relation (which in general are floating point values with single or double precision), binned in a number of columns $N_{col} = \Delta f / \delta f \cdot \epsilon_f$, where Δf is the frequency width of the peakmap, $\delta f = 1/T_{coh}$ the natural frequency resolution and ϵ_f the enhancement factor mentioned in section 2.2.2, typically set to 10.

In a typical all-sky search this implies maps with a number of columns from $\sim 2 \times 10^5$ to $\sim 5 \times 10^4$, depending on the T_{coh} used. With the natural choice on the spindown resolution, the number of rows of the map for 9 months of observation goes from ~ 1000 to ~ 250 , while it raises up to $\sim 1400 - 350$ for a 12 months long run like it was for O3. This implies maps that can reach 9 GB of memory occupation at single precision.

With these sizes, it was impossible to perform a full vectorization on both frequencies and spindown, which would have required a factor of 10^3 more of GPU memory. The solution have been to write a vectorized function of a single Hough map row, and then to use the built-in function `map_fn`, that maps a defined function over the values of a tensor, like in the pseudo-code shown below:

```

1 # Define a function that computes
2 # a single row of the matrix
3 def makeRow(i):
4     """
5     a series of instructions
6     """
7     return row
8
9 # The function is executed over an array of arguments
10 # (in this case, the range of indices between 0
11 # and the total number of spindown bins)
12 freqHoughMap = map_fn(makeRow, range(0,nSpindownBins)).

```

The `map_fn` function works similarly to the built-in Python `map` function, but it is able to distribute the calculation on the GPU topology automatically, applying an effective vectorization on a series of instructions having the logical shape of a `for` loop. This approach proved to be effective, as we will see in Sect. 2.3.2. An example of a FrequencyHough map produced for the analysis of a hardware injection on O2 data is shown in Fig. 2.15: the effect of the hardware simulated signal is evident in the characteristic double-cone shape of pixels with a higher number count. From the aliasing of the image it is possible to see how the rows of the map, computed independently, form the overall Hough map.

The candidate selection has been implemented in a very similar way: the candidates are selected over several frequency and spindown intervals, as explained in Sect. 2.2. Instead of going through the spindown intervals sequentially, with TensorFlow we can naturally parallelize the selection over them. Fixed the spindown portion, a function selects vectorially the candidates over the frequency sub-bands using the same criteria defined by the original Matlab code. The selection over the spindown intervals is done again with the use of the `map_fn` function. In this way,

since they run in parallel, the increase of computation time is negligible when we increase the number of spindown intervals and we can select more candidates over different spindowns, thus improving the search sensitivity.

The code execution then has been structured in this way:

- the peakmap, defined as a sparse matrix of the frequency-time coordinate array of the selected peaks, is loaded in the system memory and the sky grid is generated;
- a single sky point is selected, the peakmap frequencies are Doppler corrected and the weights are computed;
- the TensorFlow graph of the FrequencyHough map computation and candidate selection is built, with a placeholder for the input data;
- the graph is executed feeding the data into the placeholder. In this way there is only one memory transfer from the system to the GPU before the computation starts, and the overheads are reduced;
- the selected candidates from the map are saved;
- the same process is repeated for each sky point until the full grid has been covered with the loaded peakmap;
- once the peakmap is processed over the sky, another peakmap is loaded, a new sky grid is defined and the process is repeated for each peakmap of the analyzed frequency bandwidth.

Performances

The TensorFlow GPGPU approach on the FrequencyHough turned out to be very successful, with a speed-up of the analysis that can reach more than an order of magnitude with respect to the standard code, with same input and parameter




Figure 2.15: The GPU FrequencyHough algorithm applied to analyze a hardware injection (so-called `pulsar 8`) in O2 LIGO Hanford data, with parameters: $f_0 = 190.6373$ Hz, $\dot{f}_0 = -8.65 \times 10^{-9}$ Hz/s and ecliptic coordinates $\lambda = 351.39^\circ$, $\beta = -33.42^\circ$.




Figure 2.16: Detailed comparison between the estimated running time of the GPU- (left plot) and the CPU- (right plot) based FrequencyHough codes. The grey part of the benchmark is executed only once per job and contributes negligibly to the overall cost, while the colored parts of the bars constitute one iteration over the sky positions and are in the main loop. The TensorFlow involvement is split into two pieces: the graph building and variables initialization (orange) and the graph execution (red). The GPU code, after the graph creation, computes the FrequencyHough transform and the candidate selection at once, so it is not possible to split the two steps without the introduction of an overhead, caused by the fact that the graph is generated and run in two steps. It is to be remarked that the CPU code can run only on a single thread. Hence, in the right plot the performance are shown in terms of computing times using a single CPU core.

space. An estimation of the running time of the GPU code compared to the CPU one, running on single core, is shown in Fig. 2.16.

In the bar plots the execution time is shown for the four databases with different T_{coh} used in the analysis (see Sect. 2.2), with the subsequent analysis steps shown from the bottom to the top with different colors. The devices used are an Intel Xeon E5-2650 v4 CPU with 24 2.20 GHz execution threads and a NVIDIA Tesla V100 with 5120 cores at 1500 MHz and 16 GB of memory. The peakmaps used for the benchmarks were coming from the O3 analysis, that is with an observation time of $T_{obs} \sim 12$ months. The one with $T_{coh} = 8192$ s covers a frequency band of 1 Hz while the others are 5 Hz large. The spindown range was $-10^{-9} - 10^{-10}$ Hz/s, with the natural resolution (explained in Sect. 2.2.2) equal to $\delta\dot{f} = \delta f / T_{obs}$, i.e., the number of spindown bins for the four tests were respectively 566, 1284, 644, 322.

It is to remark that the comparison shown has been done between the new version running on a full GPU and the single core execution of the CPU code. The difference of device used prevents a detailed comparison, but it is evident how the FrequencyHough map computation becomes, from being computationally intensive on the CPU version, analogous to the computation of the peakmap weighting and

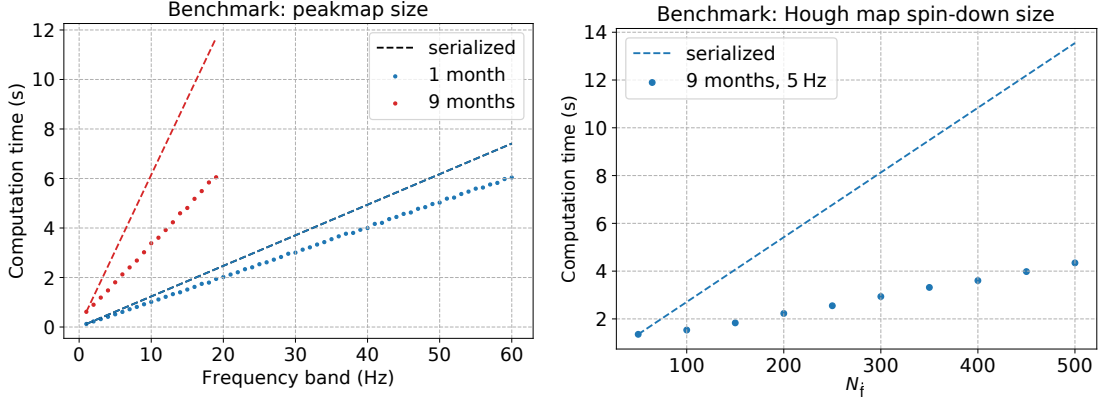


Figure 2.17: GPU tests for the FrequencyHough algorithm on a Tesla k20 (2496 cores at 706 MHz and 5 GB of memory) changing some key parameters: the size of the input (time interval–frequency range of the peakmap) and the size of the parameter space (number of spindown bins). The plots show the increasing efficiency of the GPU parallelism with a higher load on the memory and cores of the device. Left: computation time as a function of the frequency band covered by the peakmap. The measured computation times (dotted lines) are compared to those we would obtain by serializing the GPU code over 1 Hz peakmaps (dashed lines). The red plots show the test starting with a 9 months long dataset. It can be seen how when the memory of the device is full, with a 20 Hz large peakmap, the efficiency of the parallelization halves the computation time that we would need for a separate analysis of 20 Hz large peakmaps. The blue plot is the same comparison, but made with a peakmap covering only 1 month, with a 20% time gain. Note that when the memory is full the Hough computation time is the same. Right: computation time as function of the number of spindown steps of the Hough map, with $T_{obs} = 9$ months and $\Delta f = 5\text{Hz}$. Dotted and dashed lines have the same meaning as before. The plot shows a 68% time reduction with respect to the serialized case, proving that the vectorization with the TensorFlow `map_fn` function is successful.

Doppler correction, thanks to the GPU function, going from the $> 90\%$ to roughly the 50% of the total execution time.

To provide more insight on the performance and behavior of the GPU code, in Fig. 2.17 we show another set of benchmarks, performed varying either the input data size, either the parameter space size, i.e., the size of the Hough map.

For the first case, the Hough maps are computed with peakmap covering either 1 month or 9 months of data, changing the frequency range gradually from 1 Hz to the maximum allowed by the memory of the device used for the test (a Tesla k20 with 5 GB of memory), and keeping the same number of spindown bins. Then the time of the first execution has been projected as a straight line multiplying it for the increased size of the peakmap, to see the ratio of computing time between the actual computation with a N Hz large peakmap and what would be if we ran the FrequencyHough function on N small peakmaps, each large 1 Hz.

In this way we can have an idea on how the GPU parallelism gets more efficient when we use properly the device, filling the most amount possible of memory and cores. The result is that, for the 1 month case, the parallelism efficiency brings to a 20% time gain when the memory of the GPU is completely filled, and for the 9 month one reaches almost the 50%, even though we can cover a smaller range of

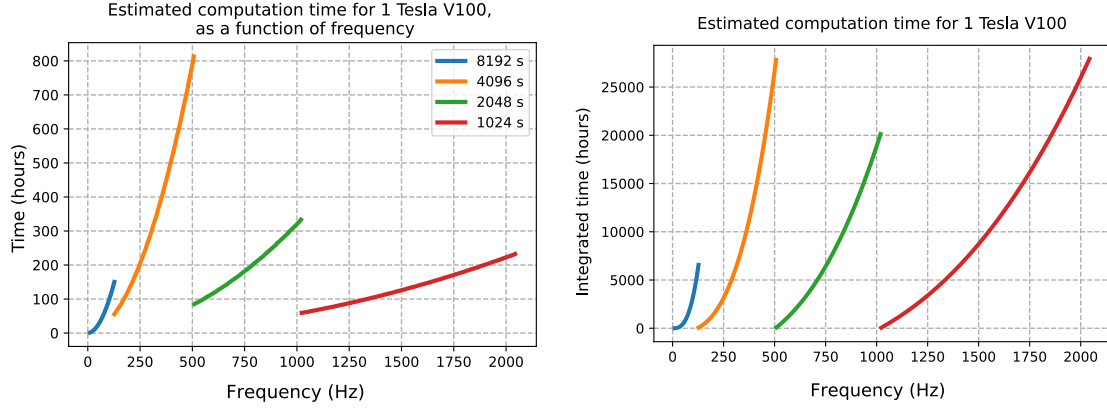


Figure 2.18: Plots showing the estimation on the time needed to complete the analysis with O3 data on the GPUs available in the CINECA Marconi 100 cluster. Left plot: the time in function of frequency needed to complete the analysis of a peakmap. Right plot: the integrated computation time to complete a full dataset.

frequencies before going out of memory. Moreover, in the two cases, the maximum computation time is the same, showing that the GPU occupancy is a key factor for the efficiency of the computation.

In the second test a single 5 Hz - 9 months large peakmap is loaded and only the number of spindown bins has been changed. In this way we can have information on the efficiency improvement of the code when we expand the parameter space, and also we can check if the pseudo-vectorization achieved using the `map_fn` function works well. With the same procedure described above, we compare the computation time increasing the number of spindown bins with the time of the first computation multiplied linearly. The results show that the parallelism achieved works very well, reaching a 68% increase in efficiency when the memory of the GPU is completely filled, and showing that the parameters space of the Hough transform can be expanded along the frequency derivative with a relatively low increase in computation time.

The results shown for the spindown dimension of the parameter space does not hold for the frequency dimension. As said in Sect. 2.2.2, the algorithm has an integration step, intrinsically sequential by definition, that in the new algorithm runs on the GPU. In this case, if we increase the resolution (that is, the number of bins) along the frequencies we have a loss in efficiency rather than a gain.

The deploy of the pipeline for O3 production

The new code has been tested and reviewed in order to be used within the LIGO-Virgo-KAGRA collaborations, showing that is able to produce the same set of candidates of the original code.

After the review process, it has been used for the O3 analysis. The search has been deployed on the CINECA Marconi 100 cluster (Cineca, 2021), providing 980 nodes with 4 NVIDIA Tesla V100 each.

A full set of benchmarks have been performed to have a precise estimation on the amount of computation hours needed to complete the search on single GPU

(we can call them *GPU-hours* as analogy to the core-hours quantity for the single core computation on multi-core systems). In Fig. 2.18 the result of this survey are shown as computation time in function of frequency, to have an idea on how much time is needed to complete a single peakmap that is 5 Hz wide (1 Hz in the range 10 - 128 Hz), and as integrated time in function of frequency to have the information on the time needed to complete the analysis of a full dataset.

The final values are summarized in the table below:

Dataset (T_{coh})	Time ($\times 10^3$ h)
8192 s	6.6 (~ 9 months)
4096 s	28 (~ 39 months)
2048 s	20 (~ 28 months)
1024 s	31 (~ 42 months)

The reason why the second band ($T_{coh} = 4096$ s, 128 - 512 Hz) takes the same time of the last one ($T_{coh} = 1024$ s, 1024 - 2048 Hz), despite being much shorter, is because the number of spindown bins was too high, increasing significantly the computation time, bringing often the GPU out of memory and slowing down, if not letting crash, the computation. Because of that, this band has been analyzed with the original code on CPUs, together with the first one ($T_{coh} = 8192$ s, 10 - 128 Hz), being smaller and more affordable for the CPU computation.

As part of my project, I took care then of the analysis and candidate selection of the peakmaps in the frequency range between 512 and 2048 Hz. In order to spread the computation accross the cluster in a homogeneous manner, the computation has been split in several jobs with the same duration ($\lesssim 24$ h). Given an input file containing the peakmap, a specific script was computing the expected computation time given the input data and search parameters. In the job submission then it was feeding the range of indexes of sky map that each single job had to process. Each job was producing an output file with the candidates list and all the output from the same input file were saved in the same folder.

Using approximatively the 25% of the resources of the cluster it was possible to perform the analysis in the selected range in few days, covering the full sky and the standard spindown range of -10^{-8} - 10^{-9} Hz/s. The results of the analysis have been published in a collaboration paper (R. Abbott et al., 2022a).

Chapter 3

The continuous-wave analysis with the radiometer method

As seen in the preceding chapters, the search for continuous waves from non-axisymmetric neutron stars, especially when we do not have any information on the source parameters, requires to span over a wide and fine-grained parameter space, demanding high computational resources.

In the recent years, the search for a stochastic background of gravitational waves with anisotropic distribution started to focus on point like sources, using the radiometer method in a pixel-wise search over the sky. Similarly to what happens with electromagnetic antennas, especially for radio and microwave frequencies, the distance between two detectors can be exploited to reconstruct a source via its phase mismatch, whose evolution depends on the source direction.

Thanks to the radiometer method, an unmodeled search can be performed over the full frequency sensitivity band of the detector and over the full sky with cross-correlation techniques. With a frequency resolution coarse enough, it is possible to neglect all the frequency variations, i.e. the spin derivatives and the Earth Doppler shift, keeping them confined in a single frequency bin.

In this way, narrowband maps of the sky can be built in a given frequency range, showing, pixel by pixel and for each frequency bin, detection statistics and *SNR*. This technique has been exploited in several searches (B. P. Abbott et al., 2017b, 2019c; R. Abbott et al., 2021d, 2022c) to search for continuous waves from non-axisymmetric rotating neutron stars.

Similarly to the FrequencyHough, presented in the preceding chapter, the radiometer search selects outliers from portions of the parameter space which show a high statistical significance. If they do not pass some verification steps, upper limits on the gravitational-wave amplitude, in function of frequency, are produced; otherwise, the statistical significance of the outliers is studied to deny or confirm a detection.

So far an in-depth characterization of the method for the continuous-wave search was missing as well as, very important, a dedicated follow-up procedure for the outliers. In the framework of my PhD project, I have performed this study and developed a pipeline to follow-up the radiometer all-sky candidates with an adapted version of the FrequencyHough.

This chapter will start with the description of the radiometer method and the

analyses performed for the search of continuous waves. Afterward, the characterization study and the follow-up method that will be applied in future searches, will be presented.

3.1 The stochastic background search

As mentioned in Sect. 1.3.3, a background of gravitational waves can be the result of two classes of sources: cosmological processes happened in the earliest stages of the Universe expansion; a coherent superposition of several unresolved astrophysical sources.

To characterize the stochastic gravitational-wave background we use the fractional energy density spectrum (Allen et al., 1999; Maggiore, 2000):

$$\Omega_{gw}(f) := \frac{8\pi G}{3H^2} \frac{d\rho_{gw}}{d\log f}, \quad (3.1)$$

where ρ_{gw} is the gravitational-wave energy density in function of frequency f , G is the gravitational constant. The Hubble constant, H_0 , is usually written as $H = h_{hubble} \times 100 \text{ km s}^{-1} \text{ Mpc}^{-1}$, with the values of h_{hubble} still widely uncertain between 0.50 and 0.75 (Maggiore, 2000 and references therein; for more recent results Aghanim et al., 2020; Denzel et al., 2021; Dietrich et al., 2020; Freedman, 2021).

Through this chapter we will focus on the searches for anisotropic stochastic gravitational-wave background based on the radiometer method (B. P. Abbott et al., 2017b, 2019c; R. Abbott et al., 2021d) and how it can be applied to point-like sources, but first we will see how the detection statistics is built starting from an isotropic distribution.

3.1.1 Isotropic optimal filter

The relation between Ω_{gw} and the gravitational-wave spectral density $P_h(f)$ is given by the formula (Allen et al., 1999; Maggiore, 2000)

$$\Omega_{gw}(f) = \frac{2\pi^2}{3H_0^2} f^3 P_h(f), \quad (3.2)$$

where $P_h(f)$ in the assumption of isotropic, unpolarized and stationary stochastic background, is defined through the ensemble average of the signal Fourier amplitudes (Allen et al., 1999; Maggiore, 2000):

$$\langle \tilde{h}_A^*(f, \hat{\Omega}) \tilde{h}_{A'}(f', \hat{\Omega}') \rangle = \delta(f - f') \frac{1}{4\pi} \delta^2(\hat{\Omega}, \hat{\Omega}') \frac{1}{2} P_h(f). \quad (3.3)$$

The index $A = \{+, \times\}$ in the formula stands for the two polarization states, $\hat{\Omega}$ and is a unit vector pointing at the celestial sphere, such that defining two angle coordinates (ϕ, θ) we have $\delta^2(\hat{\Omega}, \hat{\Omega}') = \delta(\phi - \phi') \delta(\cos\theta - \cos\theta')$.

It is worth to notice that the definition in Eq. 3.3 follows the same concept of the definition of the detector noise spectral density. The stochastic background

signal, in fact, in a single detector appears indistinguishable from the noise. The strategy to separate the signal contribution from the noise is to cross-correlate the output of a pair of detectors.

It can be shown that the optimal filtering detection statistics can be written as (Allen et al., 1999; S W Ballmer, 2006)

$$Y = \int_{-\infty}^{+\infty} df \int_{-\infty}^{+\infty} df' \delta_T(f - f') \tilde{s}_1^*(f) \tilde{Q}(f') \tilde{s}_2(f'), \quad (3.4)$$

where:

- $\tilde{s}_{\{1,2\}}$ are the Fourier transforms of the strain output time series from the two detectors;
- $\delta_T(f - f')$ is a finite-time approximation of the Dirac delta function, where T is the observation time;
- $Q(f)$ is the two-detector optimal filter, defined as the function that maximizes the SNR .

To define the optimal filter we need to express the $SNR := \langle Y \rangle / \sigma_Y$, where $\langle Y \rangle$ and $\sigma_Y^2 = \langle Y^2 \rangle - \langle Y \rangle^2$ are the expectation value and the variance of the cross-correlation statistics defined above. Reminding that $s_{\{1,2\}} = h_{\{1,2\}} + n_{\{1,2\}}$, the ensemble average in the hypothesis of a stationary gravitational-wave background and stationary and uncorrelated noise in the two detectors, will be:

$$\langle \tilde{s}_1^* \tilde{s}_2 \rangle = \langle \tilde{h}_1^* \tilde{h}_2 \rangle + \langle \tilde{h}_1^* \tilde{n}_2 \rangle + \langle \tilde{n}_1^* \tilde{h}_2 \rangle + \langle \tilde{n}_1^* \tilde{n}_2 \rangle \approx \langle \tilde{h}_1^* \tilde{h}_2 \rangle \quad (3.5)$$

The correlation terms of the signal with the noise from the detectors, $\langle \tilde{h}_1^* \tilde{n}_2 \rangle$ and $\langle \tilde{n}_1^* \tilde{h}_2 \rangle$, and the correlation between the two noise components, $\langle \tilde{n}_1^* \tilde{n}_2 \rangle$, with a long enough observation time are expected to be negligible, leaving only the signal term $\langle \tilde{h}_1^* \tilde{h}_2 \rangle$ (Allen et al., 1999).

Writing down the signal recorded by the detector $i = \{1, 2\}$, with antenna pattern F_i^A , as

$$\tilde{h}_i(f) = \sum_A \int d\hat{\Omega} h_A(f, \hat{\Omega}) e^{-i2\pi f \hat{\Omega} \cdot \mathbf{x}_i / c} F_i^A(\hat{\Omega}), \quad (3.6)$$

we can explicit the cross-correlation expectation value $\langle Y \rangle$:

$$\begin{aligned} \langle Y \rangle &= \int_{-\infty}^{+\infty} df \int_{-\infty}^{+\infty} df' \delta_T(f - f') \langle \tilde{h}_1^*(f) \tilde{h}_2(f') \rangle \tilde{Q}(f') = \\ &= \frac{T}{10} \int_{-\infty}^{+\infty} P_h(f) \gamma_{iso}(f) \tilde{Q}(f). \end{aligned} \quad (3.7)$$

Here we have introduced the normalized isotropic *overlap reduction function* γ_{iso} (Allen et al., 1999), which takes into account how the two detectors respond to an incoming signal when combined and is defined as:

$$\gamma_{iso} := \frac{5}{8\pi} \sum_A \int d\hat{\Omega} F_1^A(\hat{\Omega}) F_2^A(\hat{\Omega}) \quad (3.8)$$

In the case of an isotropic background search we cannot define an orientation between the baseline and the source, since the signal comes stochastically from all the directions. We then operate an average over the sky, and with respect to a situation with two coincident detectors, having a difference in orientation (hence we have misaligned antenna patterns $F_{\{1,2\}}$ (Allen et al., 1999)) and a spatial separation $\Delta \mathbf{x} = \mathbf{x}_1 - \mathbf{x}_2$, produces a reduction in sensitivity. Being normalized, the overlap function is limited between 0 and 1, expressing the fractional reduction of sensitivity when the detectors are not colocated and coaligned (in that case, $\gamma_{iso}(f) = 1$). Since it is a key element of the radiometer directional searches, and for the moment we are introducing the optimal filter in the isotropic case, a detailed discussion on the directional γ function is postponed to Sec. 3.1.3.

Having assumed that the noise from the detectors is much higher than the signal contribution, remembering the definition of the i th detector one-sided noise power spectral density (Eq. 2.5)

$$\langle \tilde{n}_i^*(f) \tilde{n}_i(f') \rangle = \frac{1}{2} \delta(f - f') P_i(f), \quad (3.9)$$

we can write the variance of the detection statistics as

$$\begin{aligned} \sigma_Y^2 &= \langle Y^2 \rangle - \langle Y \rangle^2 \approx \langle Y^2 \rangle \approx \\ &\approx \frac{T}{4} \int_{-\infty}^{+\infty} P_1(f) P_2(f) |\tilde{Q}(f)|^2 df. \end{aligned} \quad (3.10)$$

Defining the inner product

$$(A, B) := \int_{-\infty}^{+\infty} A^*(f) B(f) P_1(f) P_2(f) df, \quad (3.11)$$

we can write the squared SNR as

$$SNR^2 = \frac{\langle Y \rangle^2}{\sigma_Y^2} \approx T \frac{\left(\tilde{Q}, \frac{\gamma_{iso}(f) P_h(f)}{P_1(f) P_2(f)} \right)}{(\tilde{Q}, \tilde{Q})}. \quad (3.12)$$

The inner product of Eq. 3.11 behaves like a scalar product of vectors in a three-dimensional Euclidean space (Allen et al., 1999). Maximizing SNR^2 is analogous to find the maximum of the ratio $(\mathbf{Q} \cdot \mathbf{A})^2 / (\mathbf{Q} \cdot \mathbf{Q})$ for a generic vector \mathbf{A} . This happens when \mathbf{Q} and \mathbf{A} are parallel, and in our case it implies that the SNR is maximum when the filter function $\tilde{Q}(f)$ is

$$\tilde{Q}(f) = \lambda \frac{\gamma_{iso}(f) P_h(f)}{P_1(f) P_2(f)}, \quad (3.13)$$

with P_h representing our spectral model of the isotropic stochastic background, and with λ a normalization constant.

Due to the unavoidable long-term non stationarity of the detector, the data time series is split in segments that, in good approximation, satisfy the stationary noise hypothesis. On each segment j the detection statistics Y_j and variance $\sigma_{Y_j}^2$ are computed. The semicoherent cross-correlation estimator that maximizes the SNR is (S W Ballmer, 2006):

$$Y_{tot} = \frac{1}{\sigma_{tot}^2} \sum_j \frac{Y_j}{\sigma_{Y_j}^2}; \quad \frac{1}{\sigma_{tot}^2} = \sum_j \frac{1}{\sigma_{Y_j}^2}. \quad (3.14)$$




Figure 3.1: Scheme of the radiometer between two detectors from (Mitra et al., 2008). The plane wavefront coming from the direction $\hat{\Omega}$ crosses the two detectors of a baseline defined by the vector $\Delta\mathbf{x}(t)$ at different phases. Through the phase difference, with the radiometry technique, it is possible to reconstruct the source direction by the combination of the data streams from the two antennas.

3.1.2 The radiometer method

The radiometer search is a generalization of the isotropic cross-correlation optimal filter to search for an anisotropic distribution of stochastic gravitational-wave background. The concept of radiometry is to exploit the baseline distance $\Delta\mathbf{x}$ between two antennas to reconstruct the phase difference of an incoming signal from a given sky position (see Fig. 3.1). Delaying the data by the signal travel time between the detectors allows to effectively *point* the baseline toward a specific sky direction and extract the signal with the highest possible SNR via the use of a proper filter.

Instead of an isotropic distribution, we are considering the signal from a source having a certain angular power spectrum distribution over the sky $\mathcal{P}_A(\hat{\Omega}, f)$. This can be an anisotropic stochastic background, or a localized persistent source. We are assuming also that this distribution can be factorized in two elements, i.e., $\mathcal{P}_A(\hat{\Omega}, f) = \mathcal{P}_A(\hat{\Omega})P_h(f)$ ¹ (Mitra et al., 2008; Thrane et al., 2009). In this case, the gravitational-wave spectral density definition, with the above assumption, changes from Eq. 3.3 to (Mitra et al., 2008):

$$\langle \tilde{h}_A^*(f, \hat{\Omega}) \tilde{h}_{A'}(f', \hat{\Omega}') \rangle = \delta(f - f') \frac{1}{4\pi} \delta^2(\hat{\Omega}, \hat{\Omega}') \frac{1}{2} P_h(f) \mathcal{P}_A(\hat{\Omega}). \quad (3.15)$$

Our interest is to define a detection statistics for a selected sky position, in order to search for a point-like source as an emitting neutron star. We choose then $\mathcal{P}(\hat{\Omega}) = \delta^2(\hat{\Omega}, \hat{\Omega}')$: a two-dimensional delta function to describe a point source distribution. The estimator of Y and the variance will be (S W Ballmer, 2006; Mitra et al., 2008)

¹In the isotropic case we considered a uniform distribution $\mathcal{P}_A(\hat{\Omega}) = \text{const} = 1$.

$$\langle Y(\hat{\Omega}) \rangle = T \left(\tilde{Q}(f, \hat{\Omega}), \frac{\gamma(f, \hat{\Omega}) P_h(f)}{P_1(f) P_2(f)} \right); \quad \sigma_Y^2(\hat{\Omega}) \approx \frac{T}{4} \left(\tilde{Q}(f, \hat{\Omega}), \tilde{Q}(f, \hat{\Omega}) \right), \quad (3.16)$$

where the directional filter function is

$$\tilde{Q}(f, \hat{\Omega}) = \lambda \frac{\gamma^*(f, \hat{\Omega}) P_h(f)}{P_1(f) P_2(f)}. \quad (3.17)$$

The approach is similar to the continuous-wave blind searches described in Sec. 2.2. With the radiometer method and with the choice made on $\mathcal{P}(\hat{\Omega})$, it is possible to study the detector sensitivity bandwidth with a relatively refined resolution (usually, 1/32 Hz, B. P. Abbott et al., 2019c; R. Abbott et al., 2021d), pointing the analysis toward few selected sky coordinates (B. P. Abbott et al., 2017b, 2019c; R. Abbott et al., 2021d), or, via the definition of a grid in the sky, running an all-sky search in a very fast and efficient way (R. Abbott et al., 2022c; Ain, Suresh, et al., 2018). We refer usually to this kind of analysis as narrowband radiometer, as opposed to broadband searches which integrate the detection statistics over the full frequency range (e.g. R. Abbott et al., 2021d).

3.1.3 Overlap reduction function

The overlap reduction function in the radiometer analysis is crucially important because it includes the information of the phase shift between the two detectors of the baseline for any given direction.

In the preceding section we presented the best statistics for a given sky direction vector, not considering that the Earth based detectors follow the rotation and orbit motion of the planet. This means that the orientation of the single detector with respect to the source, as well as the orientation of the baseline between two detectors, will change with time.

In section 3.1 we assumed the stationarity of the detector noise to build the detection statistic, and since the assumption breaks up for coherence time above ~ 200 s (B. P. Abbott et al., 2019c; R. Abbott et al., 2021d), we use a semi-coherent approach in order to keep valid the stationarity condition. Here we add that, in the chosen time segment, also the detector parameters will be stationary: the antenna patterns $F_i^A(\hat{\Omega})$ and hence the source direction $\hat{\Omega}$, the baseline vector $\Delta \mathbf{x}$ and the overlap reduction function (Mitra et al., 2008).

Now the overlap factor depends on the frequency, on the direction $\hat{\Omega}$, but also on time, so it is different for different j segments. Keeping in mind that, with the semicoherent approach, time and frequency are discretized with resolution of T_{coh} and $1/T_{coh}$ respectively, the general expression for the overlap reduction function is (S W Ballmer, 2006)

$$\gamma_{ft}(\hat{\Omega}) = \frac{1}{2} \sum_A F_{1,ft}^A(\hat{\Omega}) F_{1,ft}^A(\hat{\Omega}) e^{i2\pi f \hat{\Omega} \cdot \Delta \mathbf{x} / c}, \quad (3.18)$$

which is the integrand of γ_{iso} defined in Eq. 3.8.

In general, the overlap reduction function is an object with a behavior difficult to handle and summarize in a plot. First of all, it is a complex quantity and it contributes to the detection statistics in a non trivial combination with the Fourier transform of the detector stream. Moreover, it changes with time and it depends also on frequency and sky coordinates. Since in our searches we integrate over several months (B. P. Abbott et al., 2019c; R. Abbott et al., 2021d), to show the magnitude of the effect of the overlap factor, it can be useful to compute an average over a sidereal day. In this way, with equatorial sky coordinates, we can collapse the dependency on the right ascension and time, keeping only the one on frequency and declination. Taking the absolute value of the result, the daily-averaged overlap factors is shown in Fig. 3.2 for the three pairs with the detectors LIGO Hanford (labeled as H), LIGO Livingston (L) and Virgo (V).

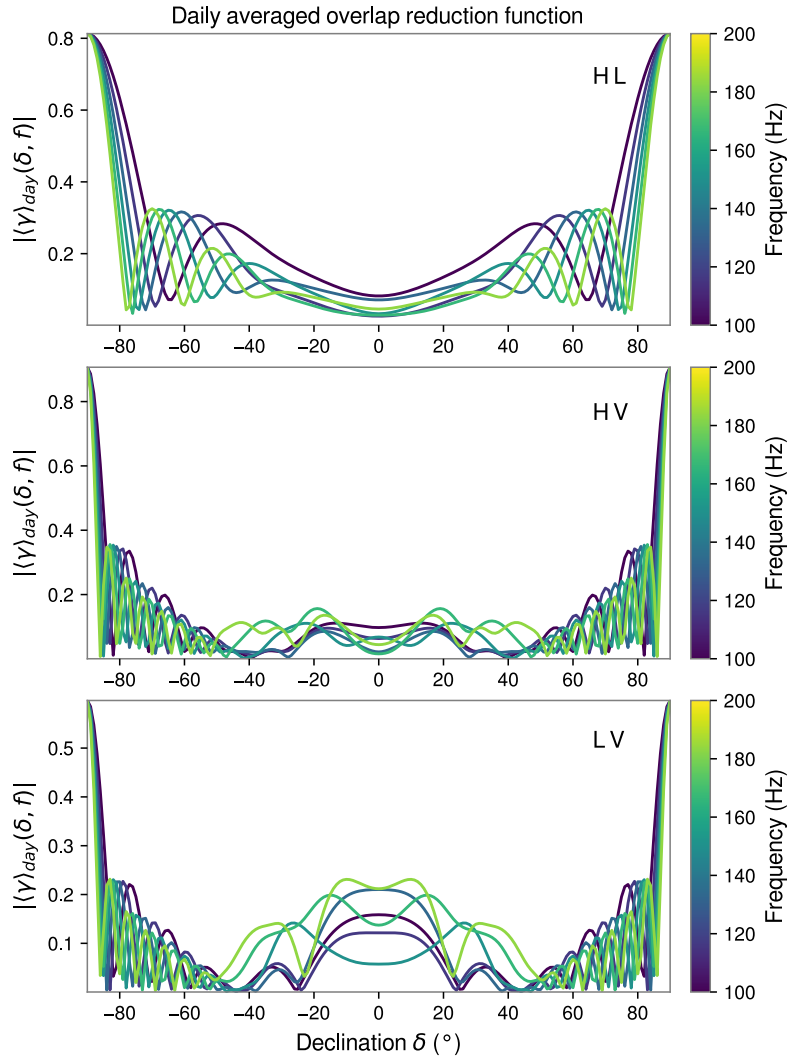


Figure 3.2: Contour plots showing, for the three pairs made by the LIGO-H, LIGO-L and Virgo detectors, the absolute value of the daily-averaged overlap reduction function $\langle\gamma\rangle_{\text{day}}$. The curves are shown as a function of the declination δ , for a range of frequencies between 100 Hz and 200 Hz. The highly convoluted shape of the function makes very difficult to foresee its interaction with the data Fourier transform when computing the cross-correlation statistics.

3.1.4 The sky map

Now that we have discussed all the elements of the optimal detection statistic, we are ready to define the narrowband sky map. In radiometry, the filter used for the cross-correlation creates the so-called *dirty map* (Mitra et al., 2008), which reproduces a sky distribution of the detection statistic influenced by the chosen filter. A signal will result smeared along the sky map (see Mitra et al., 2008 for a detailed discussion); to avoid this effect we create the *clean map*, starting from the dirty one and *deconvolving* it with the so-called *Fisher information* or *beam* matrix, which carries the information on the covariance of the detection statistic (Mitra et al., 2008; Thrane et al., 2009).

The dirty map is an array of the cross-correlation statistics, spanning the sky positions α , for each frequency bin f , summed over the short-Fourier transformed segments with timestamp t (Thrane et al., 2009):

$$X_{\alpha,f} = \frac{4}{T_{coh}} \sum_t \frac{H_f \gamma_{\alpha,ft}^*}{P_{1,ft} P_{2,ft}} \tilde{s}_{1,ft}^* \tilde{s}_{2,ft}, \quad (3.19)$$

the Fisher matrix, with same meaning for the indexes, is:

$$\Gamma_{\alpha\beta,f} = 4 \sum_t \frac{H_f^2}{P_{1,ft} P_{2,ft}} \gamma_{\alpha,ft}^* \gamma_{\beta,ft}. \quad (3.20)$$

The quantities in Eq. 3.19 and 3.20 are:

- $\tilde{s}_{i,ft}$ the Fourier transform amplitudes and $P_{i,ft}$ the power spectral density estimation for the time segment t and the frequency bin f , from the detector $i = \{1, 2\}$;
- $\gamma_{\alpha,ft}$ the overlap reduction function;
- T_{coh} the segment coherence time;
- $H_f = P_h(f)$ from now on will indicate the template function of the signal spectral model with $P_h(f)$ introduced in Sec. 3.1.2, but with discrete frequencies.

The deconvolution of the dirty map X_α with the Fisher matrix $\Gamma_{\alpha\beta}$ returns the detection statistics for the gravitational-wave power in each of the points of the map. The collection of them is the *clean map*:

$$Y_\alpha = (\Gamma_{\alpha\beta}^{-1}) X_\beta. \quad (3.21)$$

With the angular distribution we have chosen, for the radiometer search, the correlation between different pixels can be ignored (B. P. Abbott et al., 2017b, 2019c; R. Abbott et al., 2021d). Then, we will consider only the diagonal terms of the Fisher matrix, which takes the form of $\Gamma^{-1} = \sigma^2 \Rightarrow Y = X \sigma^2$. We can then write the detection statistics and standard deviation integrated over many segments, similarly to Eq. 3.14:

$$Y_{\alpha f, tot} = \frac{1}{\sigma_{\alpha f, tot}^2} \sum_j \frac{Y_{\alpha f, j}}{\sigma_{Y_{\alpha f, j}}^2}; \quad \frac{1}{\sigma_{\alpha f, tot}^2} = \sum_j \frac{1}{\sigma_{Y_{\alpha f, j}}^2}. \quad (3.22)$$

Equation 3.22 constitutes the final output of the radiometer method (Mitra et al., 2008; Thrane et al., 2009), from where outliers over a given threshold are selected and followed up with various methods (B. P. Abbott et al., 2017b, 2019c; R. Abbott et al., 2021d). If no detection has been claimed, upper limits are computed (as it will be later explained in Sec. 3.1.6).

From the study of the beam matrix in Mitra et al., 2008, it can be shown that the resolution of the radiometer is proportional to the sensitivity bandwidth and the distance between the two detectors $|\Delta \mathbf{x}|$. Assuming the same bandwidth for all the current ground-based interferometers, we have for the LIGO-H+LIGO-L baseline that a grid in the sky would have cells with a radius of ~ 0.1 rad (6°), implying ~ 1000 independent patches. For longer baselines, the resolution power is expected to be ~ 3 times higher.

Dataset combination

For the most recent narrowband radiometer searches, to improve the statistics obtained from the data, it was possible to exploit the baselines formed by the detectors LIGO-H, LIGO-L, Virgo, with the perspective to add the fourth detector KAGRA (Somiya, 2012) in the upcoming analyses. Moreover, like it has been done for O2 and O3 analyses, it is useful to combine the results coming from the previous runs with the most recent data (B. P. Abbott et al., 2019c; R. Abbott et al., 2021d). It can be shown that the combination of datasets coming from the same detector, but from different runs, and the combination of datasets coming from different detectors in the same period, can be treated in a similar way, assuming that in either cases the noise can be considered uncorrelated between the two datasets.

Let \mathcal{N} be the number of used datasets to obtain the combined dirty maps and Fisher matrices. It is enough to sum the statistics from the single dataset (B. P. Abbott et al., 2019c; R. Abbott et al., 2021d; Thrane et al., 2009):

$$X_\alpha^\mathcal{N} = \sum_k X_\alpha^k; \quad \Gamma_{\alpha\beta}^\mathcal{N} = \sum_k \Gamma_{\alpha\beta}^k. \quad (3.23)$$

The clean map will follow from Eq. 3.21: $Y_\alpha^\mathcal{N} = \Gamma_{\alpha\beta}^{\mathcal{N}-1} X_\beta^\mathcal{N}$, with $\sigma^\mathcal{N} = 1/\sqrt{\Gamma^\mathcal{N}}$ (B. P. Abbott et al., 2019c; R. Abbott et al., 2021d; S W Ballmer, 2006; Thrane et al., 2009).

Data folding

Let us take the equations for the dirty map and Fisher matrix (Eqs. 3.19 and 3.20) and let us define the following variables:

- the kernel $K = T_{coh} H \gamma$;
- the variance $\sigma^2 = T_{coh}^2 P_1 P_2 / 4$;
- the cross-correlated spectral density $C = \tilde{s}_1^* \tilde{s}_2$.

With these definitions, the sky map equations become

$$X_{\alpha,f} = \sum_t \frac{K_{\alpha,ft}}{\sigma_{ft}^2} C_{ft} \quad \Gamma_{\alpha\beta,f} = \sum_t \frac{K_{\alpha,ft}^* K_{\beta,ft}}{\sigma_{ft}^2}. \quad (3.24)$$




Figure 3.3: A scheme showing the concept of data folding from Goncharov et al., 2018. The $SNR(t, f)$ distribution over several days, with the folding, results to be collapsed in only one effective sidereal day.

It can be shown that the integration over time can be expressed in two parts: one over the sidereal days, along the full observation time, and one over the sidereal timestamp t_s of the segments, within a sidereal day with duration τ_s (Ain, Dalvi, et al., 2015). Defining the time label of a segment as $t = i\tau_s + t_s$, where i is the index of the day, the sum in Eq. 3.24 changes from \sum_t to $\sum_i \sum_{t_s}$. The dirty and variance maps can be then rewritten as follows:

$$\begin{aligned} X_\alpha &= \sum_{t_s=0}^{\tau_s} K_{ft_s,\alpha}^* \sum_i \frac{C_{f(i\tau_s+t_s)}}{\sigma_{f(i\tau_s+t_s)}^2} \\ \Gamma_{\alpha\beta} &= \sum_{t_s=0}^{\tau_s} K_{ft_s,\alpha}^* K_{ft_s,\beta} \sum_i \frac{1}{\sigma_{f(i\tau_s+t_s)}^2}. \end{aligned} \quad (3.25)$$

We note that, since the kernel contains only the overlap reduction function and the signal template, it has a daily periodic behavior. Thanks to this, we can exploit the linearity of the cross-correlation and variance statistics, compressing them in only one effective sidereal day containing the same information of several months of observation.

In Fig. 3.3 we see how the data coming from the segments are compressed into a sidereal day, keeping the binning in the frequency-time space.

3.1.5 Implementations of the radiometer method

The radiometer narrowband search has been implemented in the context of the stochastic gravitational wave background search using the `stochastic.m` (The LIGO, Virgo, KAGRA collaborations, 2022) and `PyStoch` (Ain, Suresh, et al., 2018) pipelines. The analysis can be separated in two steps: a pre-processing one, where intermediate files containing the noise power spectral density estimation and the cross-correlation statistics are generated; a post-processing step, where dirty map

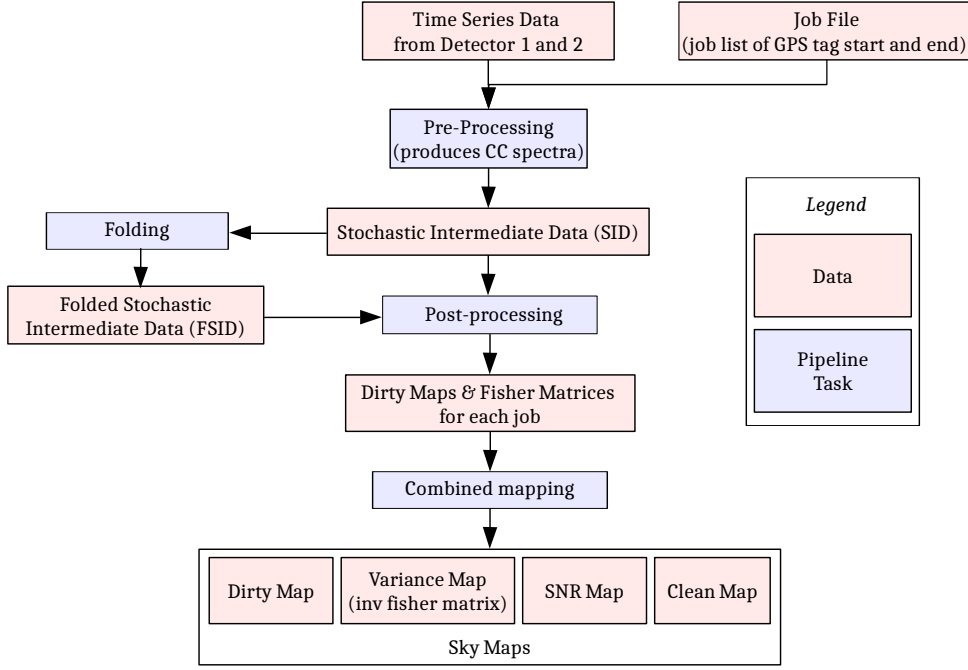


Figure 3.4: Flowchart showing the narrowband radiometer search pipeline as exposed in Sect 3.1.5, edited from Ain, Dalvi, et al., 2015.

and Fisher matrix are calculated and integrated, producing the final clean map, variance and SNR .

In the next sections we describe the most updated procedures used for the O3 narrowband search (R. Abbott et al., 2021d). Since we have shown how the results from multiple baselines can be easily combined, we will consider the analysis of data for only one baseline of two detectors.

A flowchart summarizing the steps of the narrowband radiometer search is shown in Fig. 3.4; details on the main steps of the search are given below.

Pre-processing The strain time series from the two detectors are down-sampled from the default sampling frequency of 16 kHz to 4 kHz and filtered with a high pass filter above 11 Hz, to reduce the contamination from low-frequency noise (Thrane et al., 2009). From the two time series we exclude all segments that cannot be used for the analysis, for example because the condition of Gaussianity or stationarity of the noise is violated, or due to instrumental artifacts or signal injections present in the data (see Sect. 3.2.2). The data then are split in 50% overlapping windowed segments of 192 s (B. P. Abbott et al., 2019c; R. Abbott et al., 2021d).

The portions of the data which have at least three contiguous segments are identified. From these selected time periods, the segments are Fourier transformed and the frequency bins are averaged together in order to bring the resolution from the natural $1/192$ Hz to $1/32$ Hz (Thrane et al., 2009). This gives the SFT_t dataset for the semicoherent analysis, where t refers to the timestamp of the analyzed segment.

Taken a series of three contiguous SFT_t , the cross-correlation spectral density

C is calculated using the middle one; the noise power spectral density P_n is instead estimated using the segments immediately before and after (Thrane et al., 2009). Lastly, the overlap factor for each analyzed time segment is computed.

This step of the analysis is carried out by the `stochastic.m` code, which takes the data and the search parameters and produces the intermediate files used for the post-processing step. Each contiguous segments interval compounds one job, and separated output files are generated, for a total number of jobs around 1000 usually ran on a cluster (for O3 it has been used the LIGO-Caltech cluster, LIGO Caltech, 2021).

Post-processing The post-processing step takes the data from the intermediate files, either folded or not, and computes the sky maps. The sky maps from each jobs are then integrated in order to have the final clean, variance and SNR maps as a function of frequency. All frequencies containing known instrumental lines or sharp noise features are vetoed.

The analysis can be carried out as a targeted narrowband search, computing Y and σ_Y only for a given set of sky coordinates, or an all sky search where all the maps for each frequency bins are computed.

Even if `stochastic.m` is able to run the targeted post-processing portion of the analysis, to save a significant amount of computation time we have used the newly developed and very efficient `PyStoch` tool (R. Abbott et al., 2021d), which considers by default the folded data (and can also be used for all-sky searches).

3.1.6 Standard follow-up and upper limits

In the first three observing run searches, a set of three promising targets has been chosen for the radiometer narrowband analysis: *Scorpius X-1*, *SN1987A* and the *Galactic Center* (B. P. Abbott et al., 2017b, 2019c; R. Abbott et al., 2021d). For each of these directions, the SNR versus frequency values have been stored, like in the plots in Fig. 3.5, for the three baselines combination with the LIGO-H, LIGO-L and Virgo detectors. All the frequency bins having an $|SNR| \geq 4^2$ were selected.

This procedure is repeated by introducing a delay of 1 second between the time series from the two detectors. In this way we are disrupting any correlation between the signal components of the data and we have a good approximation of an only-noise dataset, which is useful to compare with the results of the real *zero-lag* analysis.

The outliers selected from the analysis are then compared with known instrumental lines, in order to search for eventual sub-peak that would have survived

²The reason why also *negative* $SNRs$ are considered will be more clear in Chapter 3: an effect showing an excess power in a certain frequency, but not following the expected phase evolution for a given target, can appear as a high negative SNR . These frequencies are considered as outliers, in order to inspect the possibilities that high negative $SNRs$ are off-position signals or, in general, signals that are not correctly recovered by the pipeline.

the vetoing step, and also with the time-shifted SNR versus f plots. If in the time-shifted plot there is a $SNR \geq 4$ in the same or in an adjacent bin of the outlier, then probably that frequency is affected by a noise disturbance and will be discarded. For survival candidates, we study the evolution of the SNR during the observation time. Recording the cumulative sum of the detection statistics over the time segments, we check the proportionality with \sqrt{T} that would be signature of a persistent gravitational-wave signal.

Moreover, the statistical significance for each frequency bin is computed according to the following description:

Bin combination The frequency bins of Y and σ_Y are combined together to avoid spectral peaks whose power is split in two bins by Doppler modulations. Selecting the frequency bin f_j , and the number N of neighbors we want to combine together, starting from the frequency $j = N$, all bins between f_{j-N} and f_{j+N} are combined summing the detection statistics Y_j and squared-summing the standard deviation σ_{Y_j} .

By default, only the first neighbors are combined, because with the frequency resolution commonly used in the narrowband radiometer searches (1/32 Hz, B. P. Abbott et al., 2019c; R. Abbott et al., 2021d), we do not expect the signal to be spread over multiple bins. If we expect other frequency modulations, like for sources in binary systems, or a significant spindown for the observation time, a wider range of neighbors is combined to include all the modulation in one bin, depending on the source parameters (see Sect. 3.2.3 for more details on the binary systems expected Doppler modulation)

At the end, we will have a new SNR versus f plot with combined bins having width $N/32$ Hz, and a range between $f_{start+N}$ and f_{end-N} , with the labels *start* and *end* indicating the chosen analysis bandwidth.

p -values When we combine bins, we go from a situation where the different frequencies were independent to each other, to SNR distributions with a high degree of correlation. To give a statistical significance to our results, we need to simulate a high number of random Gaussian instances ($\gtrsim 1000$) of the estimator Y_j for each frequency bin, with standard deviation equal to the one obtained from the data for that bin. The simulated Y are combined, summing together the same number of bins of the measured ones according to the source expected frequency modulation. In this way we have combined-bin SNR distributions for the noise-only hypothesis.

From all noise realizations, the highest SNR for each bin is selected to form the maximum SNR versus frequency distribution. The difference between these two distributions gives the significance of each frequency bin as p -values (B. P. Abbott et al., 2017b).

In absence of a detection, upper limits are computed starting from the definition of the expected estimator μ_Y (B. P. Abbott et al., 2017b; Messenger, 2010; Whelan et al., 2014), given the gravitational wave signal amplitudes $A^+ = \frac{1}{2}h_0(1 + \cos^2\iota)$,




Figure 3.5: SNR versus frequency for the O3 targeted narrowband search. The search has been performed with the `stochastic.m` and `PyStoch` pipelines over the full O3 data, calibration C-01, sub-sampled at 4096 Hz split in segments with $T_{coh} = 192$ s. The frequency range is between 20 and 1720 Hz, with a coarse grained resolution of $1/32$ Hz. The data from the detector LIGO-H, LIGO-L and Virgo have been used to produce the combined detection statistics and variance, as explained in Sect. 3.1.4. Known instrumental lines have been removed from the final SNR results.

$A^\times = h_0 \cos \iota$, with ι the inclination angle:

$$\mu_Y = \frac{\sum_j (A^{+2} F_{1j}^+ F_{2j}^+ + A^{\times 2} F_{1j}^\times F_{2j}^\times) (F_{1j}^+ F_{2j}^+ + F_{1j}^\times F_{2j}^\times)}{\sum_j (F_{1j}^+ F_{2j}^+ + F_{1j}^\times F_{2j}^\times)^2}. \quad (3.26)$$

Note that the antenna patterns $F_{\{1,2\}j}^{\{+, \times\}}$ for the segment j , defined in Eq. 1.62, have an explicit dependency on the polarization angle ψ . The sum Σ_j is carried over all segments used for the analysis.

Since we do not know the source inclination angle and polarization, a set of values of μ_Y are computed for several values, for example 50, of both $\cos \iota$ and ψ , and with the Bayesian approach the upper limits are calculated marginalizing over those values. Moreover, to take into account the uncertainty and systematic errors in the calibration process (Acernese et al., 2022; Sun et al., 2020), another marginalization is done over the unknown calibration correction factor l . The prior probability distribution is assumed to be Gaussian, with mean 1 and standard deviation given by the estimated calibration uncertainty σ_l (B. P. Abbott et al., 2017b; Whelan et al., 2014).

The likelihood function that will be marginalized is

$$p(Y|h_0, \iota, \psi, l) = \exp \left\{ -\frac{1}{2} \left(\frac{l}{\sigma_l} \right)^2 - \frac{1}{2} \left[\frac{Y(l+1) - \mu_Y}{\sigma_Y(l+1)} \right]^2 \right\}, \quad (3.27)$$

and the marginalized posterior for each frequency³ is (B. P. Abbott et al., 2019c; Whelan et al., 2014)

$$p(h_0|Y) = \int_{-1}^1 d(\cos \iota) \int_{-\pi/4}^{\pi/4} d\psi \int_{-1}^{\infty} dl p(h_0, \iota, \psi, l). \quad (3.28)$$

The upper limits are defined integrating the posterior function up to the value h_0^{UL} that returns the chosen confidence of 90% (B. P. Abbott et al., 2017b):

$$0.9 = \int_0^{h_0^{UL}} p(h_0|Y, \sigma_Y). \quad (3.29)$$

To compute the upper limit for the combined data from multiple datasets (baselines or run, as stated in Sect. 3.1.4), assuming that the detectors are independent, as well as their calibration factors $l_{\mathcal{N}}$ (having given uncertainty $\sigma_{l_{\mathcal{N}}}$), we multiply the single likelihood functions (Whelan et al., 2014):

$$p(\vec{Y}|h_0, \iota, \psi, \vec{l}) = \prod_{\mathcal{N}} p(Y_{\mathcal{N}}|h_0, \iota, \psi, l_{\mathcal{N}}), \quad (3.30)$$

with \mathcal{N} running over the combined datasets

The results of the O3 targeted narrowband analysis, which I have actively contributed to, are shown in R. Abbott et al., 2021d.




Figure 3.6: An illustration of the HEALpix grid from Gorski et al., 2005. It shows the meaning of the N_{side} parameter, where the side of the first grid is divided in equispaced segments to generate the grids with $N_{side} = 2, 4, 8$.

3.1.7 All-sky search

The all-sky search for persistent sources, using the radiometer method, has been performed with the `PyStoch` pipeline on O3 data (R. Abbott et al., 2022c). Exploiting the reduced input size thanks to the folding, the pipeline is able to load all needed data at once from the intermediate file (coming from the pre-processing step described in Sect. 3.1.5). With those data, it builds efficiently the full sky maps for each frequency bin using a single cluster node in the timescale of ~ 1 hour.

For this search, an equispaced spherical grid, the so-called HEALpix (Gorski et al., 2005), has been used.

The HEALpix grid is defined by the sole parameter N_{side} , which indicates the number of divisions on the side of a cell in the base pixel resolution (see Fig. 3.6), and each pixel is identified by a unique index that can be directly translated to angle sky coordinates. For the O3 search $N_{side} = 16$ was chosen, which corresponds to 3072 pixels with a resolution of $\sim 3.67^\circ$. This resolution is a factor of 3 more refined to the one computed in Mitra et al., 2008 for the natural resolution of the radiometer method using the LIGO-H plus LIGO-L baseline, and it is comparable with the one with baselines involving Virgo and KAGRA (see Sect. 3.1.2).

This grid is very different from the one used in the FrequencyHough pipeline described in Sect. 2.2.1. A comparison is shown in Fig. 3.7, where the two spherical grids are projected on a plane: the main difference is that the distance between the center of the cells of the grid in the HEALpix does not depend on the latitude, as it does in the FrequencyHough grid. As it will be shown in Chapter 3, this difference

³The frequency dependency is omitted in the prior and posterior equations to lighten the notation, since we are considering a signal coming from a monochromatic source. Nevertheless, the computation of the upper limits is carried on for each frequency bin in the analyzed range, with the same estimator μ_Y and with their respective Y_f and σ_{Y_f} .




Figure 3.7: A comparison of the HEALpix grid with $N_{side} = 16$, 3072 pixels and a FrequencyHough grid with comparable size. It is important to remark the different sky tiling for the follow-up analysis of candidates from the all-sky radiometer search.

has been addressed when developing the continuous-wave dedicated follow-up for the radiometer candidates.

The candidate selection method consisted into comparing the SNR results with the SNR of the time-shifted analysis, similarly to what has been done for the directed narrowband analysis (see Sect. 3.1.5).

Since the analysis is performed over ~ 3000 points (rather than 3), and we will have a far higher production of outliers, the following procedure has been chosen (R. Abbott et al., 2022c):

- the SNR distributions in the zero-lag and time-shifted analyses are projected on the frequency dimension, taking the maximum SNR from each map;
- defining frequency bands of 10 Hz, a histogram of the $SNR_{max}(f)$ distributions for each band is computed;
- for each 10 Hz band, the 99th percentile of the time-shifted SNR_{max} histogram is identified, then a threshold SNR_{thr} is defined averaging three neighboring bands;
- any of the original frequency-sky bins falling in each sub-band, with $SNR_{max} \geq SNR_{thr}$, is stored for the follow-up study.

The treatment of the follow-up for the candidates coming from the all-sky radiometer analysis will be described in the next chapter.

3.2 The fast identification project

As we have seen in the previous section, with gravitational-wave radiometry we can create a map of the cross-correlated filter statistics over the full sky, with a very high efficiency and relatively low computational resources, with the trade-off of a lower reachable sensitivity.

Thanks to the speed of the radiometer algorithm, it is possible though to have quickly a set of outliers to analyze, but in order to be able to claim a detection we need a dedicated follow-up.

The project presented here then consists of two steps: the first has been focused on the characterization and usage of the stochastic analysis pipelines. The objective is to study how to reduce the parameter space to frequency-sky regions with a high SNR , selecting candidates to be further analyzed. The second one consists in the development of a follow-up of those candidate parameter-space regions, using the FrequencyHough pipeline to analyze them in order to reach a higher sensitivity.

For the first part, the code used has been `PyStoch`, which takes advantage of data folding and allows us to run the all-sky, all-frequency analysis in a short time, building SNR sky maps for each frequency bin within a defined range and resolution. On these maps we can select the pixel areas satisfying some relevant conditions, such as SNR above a certain threshold, or a stricter criterion as the one presented in R. Abbott et al., 2022c (see Sect. 3.1.7), and investigate the outliers with a more thorough analysis with continuous-wave tools.

The continuous-wave follow-up consists in building customized peakmap for each candidate, for the full observation time and in a small frequency range around the outlier frequency, and performing the FrequencyHough transform directed to a certain number of sky points in a region around the candidate position.

This section will show in first instance the results of an in-depth set of tests performed to characterize the behavior of the radiometer analysis when used for the continuous-wave searches. Subsequently, the follow-up method will be explained.

3.2.1 Continuous-wave radiometer characterization

The first step of the project has been dedicated to a set of tests aiming to have more information on the behavior of the radiometer search for monochromatic point-like sources.

The study involved tests on software generated data, using both simulated noise and real data where fake signals have been injected. To keep the computation simpler and faster, in first instance, a dataset of the so-called *frame files*, containing the data time series, has been generated with simulated white gaussian noise, with a sampling frequency of 512 Hz (that means with a band with maximal frequency of 256 Hz) and $\sqrt{S_h} = 4 \times 10^{-24} \text{ Hz}^{-1/2}$.

On this simulated noise, in turn, various fake signals have been injected. For both the noise and the pulsar signals generation, the tool `lalapps_MakeFakeData_v5` has been used, which is part of the LALsuite software (LIGO Scientific Collaboration, 2018). With this tool it is possible to simulate signals without noise. The feature has been proven useful to perform tests on real data, where frame files with only simulated signals were generated with the same time intervals and sampling frequency of the data files. Afterward, using the `PyCBC` Python libraries (Nitz et al., 2022), the two time series were summed together element by element. The result was finally stored for the radiometer analysis.

For most of the tests, the pair of detectors LIGO-Hanford and LIGO-Livingston was used both for the noise simulation and for the injections on real data, because they have similar sensitivity.

With `lalapps_Makefakefata_v5` we could simulate at the same time a certain number of isolated and/or binary pulsars with the complete set of parameters:

- signal amplitude h_0 , frequency and spin derivatives;
- equatorial sky coordinates in radians;
- source inclination angle and signal polarization angle;
- reference phase and time;
- if needed, orbital parameters.

In all of our tests, the spin derivatives have been set to 0, as well as the inclination and polarization angles. Moreover, to stay in a fairly stable region of the detectors, the simulations and tests have been limited to the [100-200] Hz frequency range.

The data have been analyzed equivalently with the narrowband targeted mode of `PyStoch` and `stochastic.m`, holding the general conventions used for the narrowband stochastic background searches (see Sect 3.1 and B. P. Abbott et al., 2017b, 2019c; R. Abbott et al., 2021d), such that the coherence time $T_{coh} = 192$ s and the frequency resolution $\delta f = 1/32$ Hz.

Detection statistics

The first test with simulated signals aimed to study the trend of the output detection statistics Y , when varying the amplitude h_0 of the signals injected.

To achieve this, a reasonable number of 10 signals from isolated pulsars have been simulated, keeping for all the same parameters except for the amplitudes, which were varied uniformly in a defined range, and the frequencies, varied as well uniformly in order to have the results in only one plot. The separation between the injections has been set with a wide random variation to avoid interferences between the different signals.

The same procedure was followed both for injections on the simulated noise mentioned above and on O2 real data (B. P. Abbott et al., 2019c). To have comparable observation time, the time duration of the dataset with simulated noise has been chosen equal to 3 months, similar to the integrated time of the segments used for the O2 search (B. P. Abbott et al., 2019c).

In the real data case, the injected amplitudes spans the range 7×10^{-26} - 5×10^{-25} ; in the simulated noise case, because the lower noise spectral amplitude, it was possible to reduce the range to 6×10^{-26} - 1.8×10^{-25} .

The results are shown in Fig. 3.8. For a general comparison of the obtained values with the theoretical ones, Eq. 3.26⁴ has been used for each amplitude value to compute the expected Y and it has been superimposed in the plot.

⁴Reproduced here for the reader's convenience:

$$\mu_Y = \frac{\sum_j (A^{+2} F_{1j}^+ F_{2j}^+ + A^{\times 2} F_{1j}^\times F_{2j}^\times) (F_{1j}^+ F_{2j}^+ + F_{1j}^\times F_{2j}^\times)}{\sum_j (F_{1j}^+ F_{2j}^+ + F_{1j}^\times F_{2j}^\times)^2}.$$

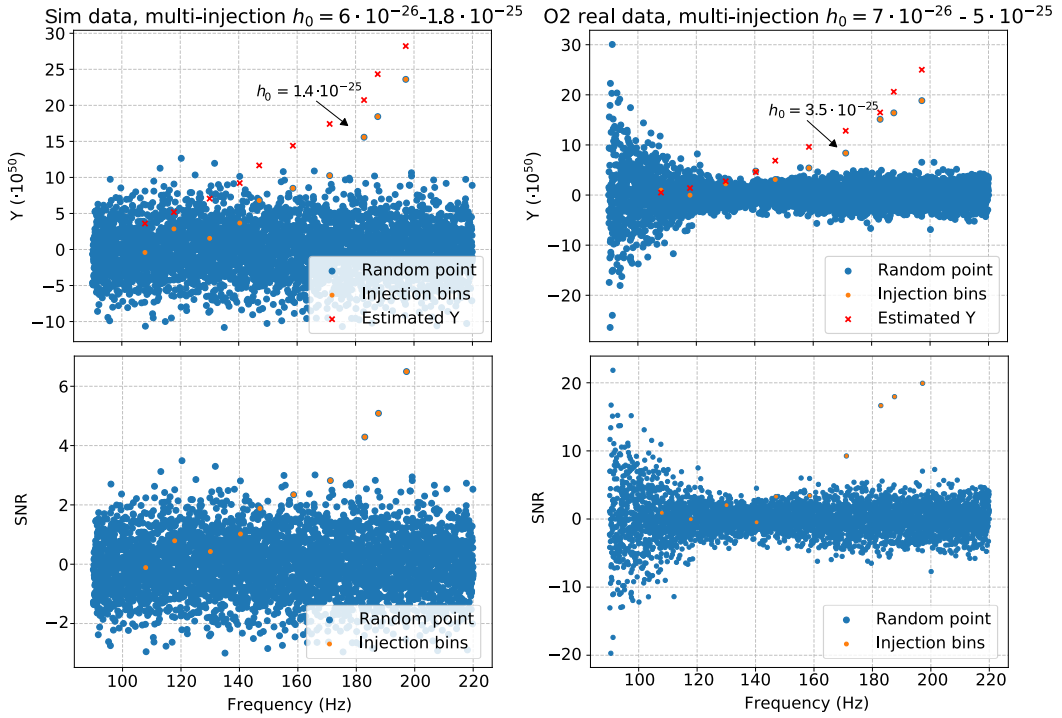


Figure 3.8: Comparison between the estimated and retrieved clean maps statistics Y , analyzing injections of isolated pulsars at various amplitudes on simulated noise (left plots) and real data (right plots). The output of the radiometer analysis is shown in blue, while the frequency bins of the injected signals are highlighted in orange; superimposed in red the theoretical estimation from Eq. 3.26. In the two plots below, the SNR values are plotted with same color codes. The first signal in the $Y(f)$ plots with an SNR above 4 (respectively 4.2 and 9.3) is indicated with an arrow. The values of Y are defined up to a normalization constant, that has not been applied in the computation of μ_Y , this is why the expectation values are systematically higher. Nevertheless, the purpose of this test was to study the trend of the detection statistics and not the actual values, since the determining quantity for the identification of a signal is the SNR .

Multi-injections test

Here we explore the analysis made for the targeted narrowband stochastic search. Four isolated pulsars were injected on simulated noise and on O2 real data with the same amplitude ($h_0 = 2 \times 10^{-25}$ for the test on simulated noise, $h_0 = 4 \times 10^{-25}$ for the software injections on real data) at random well separated frequencies.

These tests, whose results are shown in Fig. 3.9, revealed that noise fluctuation on real data can affect significantly the chances to detect a signal even with some months of gathered statistics, especially in a more realistic situation with much fainter signals.

While the simulated noise shows a fairly stable SNR distribution for the injected amplitudes, the case on real data shows significant differences in the retrieved SNR . Having lower level of noise, as well as a longer observational period to achieve a higher level of statistics, would stabilize the results obtained by the analysis.




Figure 3.9: Results of a test analysis on four signal injections at different sky positions on simulated noise (left) and real data (right). The signals have the same amplitudes.

To investigate further the behavior when many signals are present, a set of three different tests have been performed, based on the generation of a single instance of simulated gaussian noise, and injecting in turn three different sets of signals:

- four neutron stars at different coordinates (the ones of Sco X-1 and SN1987a and two other random sky points) with amplitude of $h_0 = 2 \times 10^{-25}$;
- the same four sources plus six more neutron stars at different coordinates, with same amplitude;
- the same starting four sources, but with the amplitude lowered to $h_0 = 1.5 \times 10^{-25}$.

The results are shown in Fig. 3.10 and they are fairly enlightening in showing how radiometer analyses respond in a situation where many signals are loud enough to appear with a high SNR .

First of all, we see again the fluctuation of the SNR even if the amplitudes of the injected signals are the same. Secondly, lowering the amplitudes, the retrieved SNR is scaled accordingly, with some differences probably due to the different behavior of the overlap reduction function in the different coordinates and frequencies. The last important consideration is that, when increasing significantly the number of loud signals, the SNR of the original four fake signals remains the same, showing that the assumption of independent sky map pixels and frequency bins hold, and the signals do not interfere together.

Given the loudness of the injected signals, however, some weird behaviours appear, often shown as a high negative SNR especially in bins where there are injections and we are not analyzing the respective sky coordinates. This is what we expect to happen when a loud feature is present in the data, but does not follow the expected phase evolution for the chosen coordinates (see Fig. 3.11): the feature will result anti-correlated with the expected template for a given sky position, returning a negative detection statistics.

Sometimes, when analyzing a certain source from a given sky position, signals which are injected in other sky coordinates even exhibit a high positive SNR in their respective frequency bins. The explanation is similar to what happens for




Figure 3.10: Results of different tests on: 10 isolated pulsars with the same amplitude and random different coordinates (black squares), a subset of those pulsars (colored squares), the same subset but with lower amplitude (colored triangles).




Figure 3.11: Directed analysis for two sky directions (i.e., the coordinates of Sco X-1 and SN1987a) for simulated noise with an injection of a sinusoid function with fixed frequency. The output of the radiometer analysis is a high negative SNR , showing the presence of a feature in the data high enough to be seen by the cross correlation statistics, but being anti-correlated with respect to the expected phase evolution for the positions we analyzed.

loud instrumental lines: the involved frequencies show high SNR independently of what point in the sky we are analyzing. However, this is not a source of concern, since the simulated amplitudes on these tests were much higher than the ones we expect from real sources.

An example of how a monochromatic disturbance behaves in the radiometer search is shown in Fig. 3.11. On the Gaussian noise time series we have injected a sinusoid function, with no modulations, to simulate any kind of source that is not following the phase evolution expected for the targeted direction. The data then were analyzed towards the usual directions Sco X-1 and SN1987a, and the results show in both cases a high negative SNR corresponding to the injected sinusoid frequency. In this way it was possible to verify the interpretation of the strong negative $SNRs$ that appear in the analyses and in the tests done.

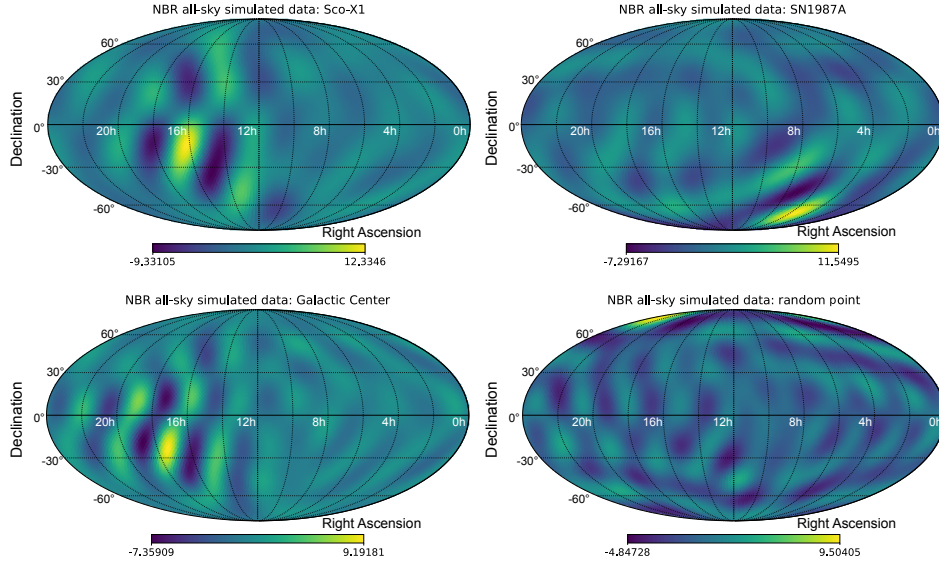


Figure 3.12: Selection of sky maps at the frequencies of some software injections. The parameters are the same of the previous tests based on the O3 narrowband search (see Sect. 3.2.1 and Fig. 3.9)

All-sky test

So far we have always considered targeted analyses because it was useful in the study of the SNR versus frequencies, as we selected only one pixel in the narrowband sky maps.

Using the `PyStoch` code, however, we can compute and save all narrowband maps at once in a single output file. Then, going through the maps for each frequency bin, we can select sky positions with a higher statistical significance, like SNR above a certain threshold or more robust criterion (see Sect. 3.1.7 and R. Abbott et al., 2022c for more details on the outlier selection for all-sky analysis performed with `PyStoch`).

In Figure 3.12 we show as an example the narrowband maps at frequencies corresponding to the four injections on simulated noise. The analysis process is the same of the tests shown in Sect. 3.2.1. This time, instead of using the targeted mode of the code, the full sky maps have been generated for the full search frequency range. The maps at the frequencies of the four software injections have been selected and plotted.

3.2.2 Hardware injection analysis with O2 and O3 data

In this section we summarize the results of an investigations on O2 and O3 LIGO Hanford-Livingston data to test the behavior of the radiometer search with an analysis aiming at retrieving the hardware injections on real data. The hardware injections are continuous signals injected directly in the detector during the observational runs, and not via software simulations. The list of injections with their parameters is shown in Tab. 3.1.

The analyses set up is the same as the one used for O2 and O3 narrowband

Table 3.1: O2 and O3 hardware injection parameters (LIGO collaboration, 2017, 2019)

#	h_0 (O2)	h_0 (O3)	f_0 (Hz)	\dot{f} (Hz/s)	Dec ($^\circ$)	RA ($^\circ$)
0	9.8×10^{-25}	6.1×10^{-26}	265.58	-4.15×10^{-12}	-56.22	71.55
1	4.4×10^{-24}	5.5×10^{-25}	848.94	-3×10^{-10}	-29.45	37.39
2	1.7×10^{-24}	7.6×10^{-26}	575.16	-1.37×10^{-13}	3.44	215.26
3	8.2×10^{-25}	1.3×10^{-25}	108.86	-1.46×10^{-17}	-33.44	178.37
4	4.6×10^{-24}	1.1×10^{-24}	1391.01	-2.54×10^{-08}	-12.47	279.99
5	1.5×10^{-24}	4.0×10^{-25}	52.81	-4.03×10^{-18}	-83.84	302.63
6	8.9×10^{-25}	3.8×10^{-25}	145.50	-6.73×10^{-09}	-65.42	358.75
7	4.6×10^{-24}	1.7×10^{-25}	1220.44	-1.12×10^{-09}	-20.45	223.43
8	1.1×10^{-24}	1.3×10^{-25}	190.17	-8.65×10^{-09}	-33.42	351.39
9	3.0×10^{-24}	1.3×10^{-25}	763.85	-1.45×10^{-17}	75.69	198.89
10	4.0×10^{-24}	6.3×10^{-25}	26.33	-8.5×10^{-11}	42.88	221.56
11	5.6×10^{-24}	3.2×10^{-25}	31.42	-5.07×10^{-13}	-58.27	285.10
12	3.3×10^{-24}	2.6×10^{-25}	37.85	-6.25×10^{-09}	-16.97	331.85
13	2.2×10^{-22}	2.6×10^{-24}	12.43	-1×10^{-11}	14.32	14.32
14	3.4×10^{-24}	1.8×10^{-24}	1991.09	-1×10^{-12}	-14.32	300.80
15	\times	8.7×10^{-25}	2991.09	-1×10^{-12}	-14.32	300.80
16	\times	1.6×10^{-25}	234.57	0	-15.64	19.98
17	\times	8.4×10^{-26}	890.12	0	-15.64	109.98

searches (B. P. Abbott et al., 2019c; R. Abbott et al., 2021d). Regarding the data, for the O2 test the standard calibrated and cleaned (Cahillane et al., 2017; P. Covas et al., 2018) strain time series have been used. For the O3 search, the peculiar noise behavior faced during the run, already mentioned in Sect. 2.1, brought to several datasets with different cleaning and filtering processes applied, and it has been difficult to find the one that could have fitted properly the analysis.

After several attempts, the choice has been to use the cleanest possible dataset, where, apart for the standard calibration and cleaning procedures, also the gating (Mata et al., 2021; Zweizig et al., 2021) and an additional filtering for instrumental noise below 60 Hz has been applied. This dataset was used with the FrequencyHough pipeline to retrieve successfully a subset of the available hardware injections.

The analysis on O2 data has been successful and most of the hardware injected signals (with parameters shown in Tab. 3.1) are retrieved with correct parameters. The analysis on O3 data, on the other hand, does not show in the search frequency interval (20-1726 Hz) any of the injections from the list shown in Tab. 3.1.

To establish whether the problem involves the data or the search parameters, another test has been done on O3 data, injecting via software the set of pulsars with exactly the same parameters from the list of candidates shown in Tab. 3.1, plus a set of four pulsars at different frequencies between 100 and 200 Hz, all located in a single random sky position and amplitudes of $h_0 = 1, 1.22, 1.38, 1.5 \times 10^{-24}$ ⁵.

The data files with the software injections were prepared generating a set of

⁵Like all the other simulate signals used for characterization tests, spindown, inclination angle and polarization angles are set to 0

frames, with same starting time, duration and sample frequency as the O3 real frames, and contained only the signals of the fake pulsars via `lalapps_MakeFakeData_v5`. Then, the software injections and the O3 datasets have been summed together using `PyCBC`, like for the O2 tests shown in the above sections. The resulting frames have been saved and analyzed.

Results

The plots in this section (Fig. 3.13 and Fig. 3.14) show all the SNR versus f results from the hardware injections study (13 pulsars in different sky positions for O2, 15 for O3 in the search frequency range).

Similarly to what has been done before, for each analysis the different sky points are independent to each other and they are plotted in the same figure with different color codes. The colors indicate at which frequency we expect to see the injection

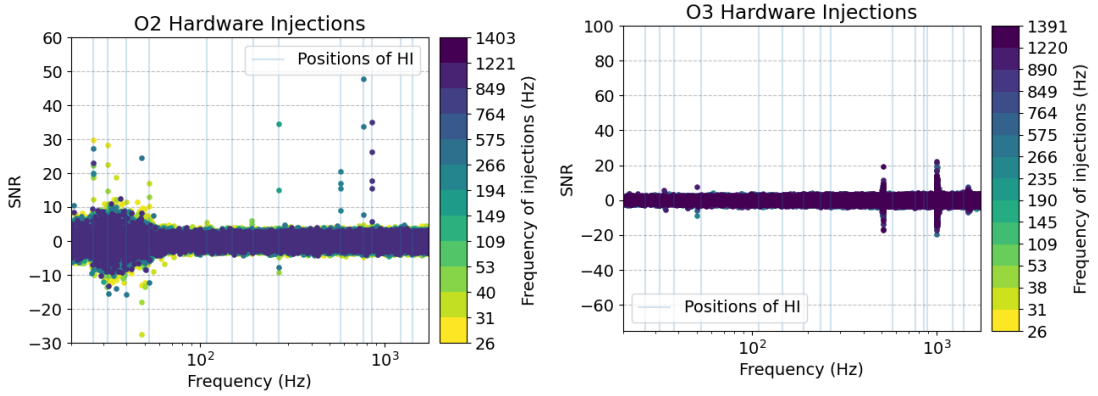


Figure 3.13: O2 and O3 hardware injections summary. Most of the O2 ones appear very sharp, while the O3 plots show only instrumental lines.

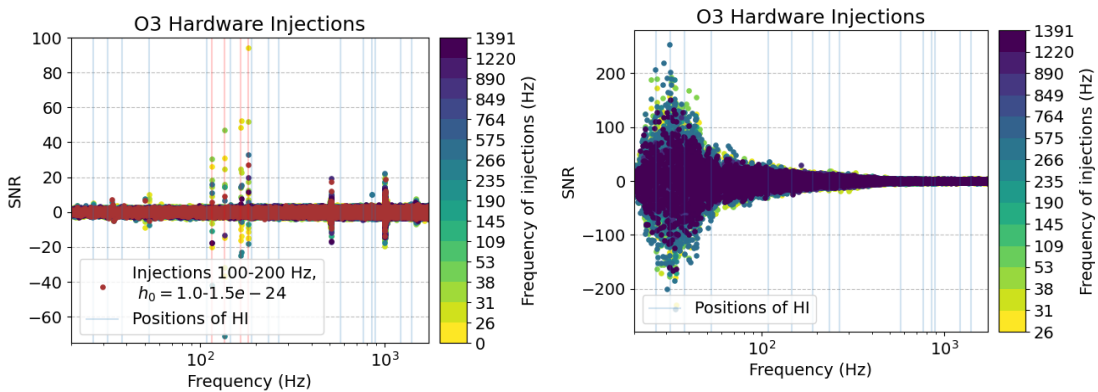


Figure 3.14: Left: plot of the results of the analysis with the four software injections, superimposed in red over the rest of the hardware injections analysis results. Four red line highlighting their frequency are also shown. The software injections were intentionally very loud in O3, to show up with a high SNR in every other sky direction we analyzed. Right: analysis on calibrated and cleaned, non-gated frames. The extremely high disturbance prevents to retrieve any hardware injection.

at the respective sky position. In this way it is possible to see, in only one plot for each test, how many injections have been spotted.

From the comparison study between O2 and O3 it appears that either the injected amplitudes were too low for the radiometer analysis or the noise disturbances in the O3 data prevent to retrieve correctly the hardware injections, and is giving troubles also for software injections.

A last test has been performed using the parameter file and executable used for the O2 test, but with O3 frames with no gating procedure. We chose to select the same kind of frames used for the O2 analysis, that is just with standard calibration and cleaning process. Unfortunately this test showed that this dataset cannot be used for this study, as shown in Figure 3.14: a very strong disturbance populates the frequencies up to ~ 500 H, making impossible to retrieve any hardware injections.

Looking forward to the upcoming fourth observational run (O4), the expected behavior of hardware injections needs to be thoroughly studied.

3.2.3 Binaries

So far, we have considered only isolated neutron stars in our tests. Since, of the ~ 3000 observed pulsars, roughly half of them is located in binary systems (ATNF, 2021), there is the concrete possibility that an unknown source would orbit with a companion star. It is important then to analyze what happens when we want to retrieve binary signals with the radiometer search.

As shown in Leaci et al., 2017, the spectrum energy of an emitting neutron star in a binary system is spread over a certain frequency range by the orbital Doppler effect, in a characteristic double-horn shape. In Fig. 3.15 we report an example of how a simulated binary signal appears, with a modulation of ~ 37 mHz, after the application of a series of filters to enhance the spectral amplitude distribution as a function of frequency (see Leaci et al., 2017 for more details).

The maximum Doppler shift caused by the orbital motion depends on the source Keplerian parameters Leaci et al., 2017:




Figure 3.15: Figure from Leaci et al., 2017 showing the output of a couple of *ad hoc* built filters applied to the power spectral peak frequencies, expressed as W_w/W .




Figure 3.16: Results from the test using simulated signals in binary systems. Two of the retrieved signals are zoomed as an example, showing the frequency profile of the SNR of the injected signal. As expected, we show in the zoomed panels the Doppler shift of four bins coming from the chosen parameters and the double-horn feature is visible.

$$\Delta f_M = \frac{2\pi a_p}{T_{orb}(1-e)}, \quad (3.31)$$

where we have: $a_p = a \cdot \sin \iota / c$ as the projected semimajor axis, with a being the orbit semimajor axis and ι the inclination angle of the orbit; T_{orb} as the orbital period and e the orbital eccentricity.

Using this formula we can compute the maximum range of frequencies over which a signal will be spread by the Doppler effect. Exploiting that, a set of four injection has been generated choosing the orbital parameters to reproduce a Doppler spread of four times the typical radiometer search bin of $1/32$ Hz.

Four well distanced frequencies and sky points have been chosen for the signal simulation study. After that, the parameters a_p and T_{orb} have been chosen to obtain a signal spread over four frequency bins for each of them (for the sake of simplicity we consider circular orbits, $e = 0$). To limit the choice of the two parameters to a reasonable range, values of $a_p \in [1, 3]$ s and $T_{orb} \in [10, 48]$ h (Leaci et al., 2017) have been considered. The signals have been simulated and injected on O2 real data, together with three isolated signals for comparison.

From the tests done, we see that the power of the binary signals is distributed over the expected number of bins, causing a reduction of the SNR we would obtain with respect to signals with the same amplitude, but coming from isolated sources.

Even though this implies that it is more difficult to detect a source in a binary system, looking at the plot in Fig. 3.16, we note that given the shape of the signal profile, the signal power is not evenly spread and that holds also for the SNR . Even if we have spread the binary signals over four bins, the power of the higher horn of the profile still reaches $SNRs$ similar to those of the isolated signals we have injected, which had only half the amplitude.

Nevertheless, to properly see the injected signals above the noise fluctuations, we have injected amplitudes much higher than the ones we expect to detect. This implies that in real situations it will be difficult to see the full profile of a signal.

Another consideration regards the possibility that a source would have orbital parameters producing a maximum modulation that, with the used frequency resolution, will result in a signal that would split only over two frequency bins. In this case it might be difficult to distinguish the binary signal from an isolated one as the shift of two frequency bins could be caused also by the spindown or the Earth Doppler effect. Increasing the frequency resolution would help to resolve the modulation profile, but it will result in a lowered retrieved SNR because the signal will be spread over more bins.

3.2.4 The radiometer continuous-wave follow-up

Once properly understood how the radiometer search works in various situations, we can move to the final step of the project, regarding the follow-up of the outliers selected from the all-sky search.

As explained in Sect. 3.1.7, the HEALpix grid, in the analysis with PyStoch, does not match with the grid defined in the FrequencyHough pipeline: not only the shape is different, but since the FrequencyHough grid resolution depends on the frequency, the resolutions of the two grids can have a big mismatch.

Taking the coordinates of a PyStoch candidate, a HEALpix cell is identified and needs to be completely covered by the follow-up analysis: to avoid to miss a signal, we need to be sure that we are able to fully cover the identified cell with the proper FrequencyHough sky grid. The procedure to achieve this is the following:

- for each radiometer outlier, the coordinates of the center of the cell are used to retrieve the first HEALpix grid neighbors, by using the HEALpy (The HEALpy contributors, 2022) Python library. In this way, as it is shown in Fig. 3.17 by the blue dots, we can define a safe area containing the signal, covering the HEALpix cell of the outlier plus a portion of all neighboring cells around it;
- using HEALpy, the maximum distance between the given point and their first neighbors is computed, defining a circle that covers the above defined area where we want to search the signal with the FrequencyHough;
- with a function of the SNAG toolbox (Frasca et al., 2018), the FrequencyHough grid for a given frequency, centered in a sky point and with a given angular distance can be created, as shown in Sect. 2.2.1. Using the radius obtained using HEALpy, the grid is created for the frequency of the outlier.

An example of this procedure is shown in Fig. 3.17.




Figure 3.17: Example of FrequencyHough grid for two example PyStoch outliers at frequency 216.16 (left) and 129.02 (right). The HEALpix cell we have chosen is identified by a black dot, the first neighbors are shown in blue and the maximum distance between the outlier and its neighbors has been used to draw a red ellipse around the outlier, which covers the area we have to analyze. The FrequencyHough grid, generated with the source parameters and completely covering the found area is plotted with red dots.

Now we can launch the FrequencyHough analysis for every radiometer candidate, over its own sky grid and on a given portion of the peakmap around the candidate frequency (let us say of 1 Hz). The range of spindown can be arbitrarily chosen, but the natural choice is to keep the range used for the O2 and O3 all-sky analyses (see Sect. 2.3.2 and R. Abbott et al., 2022a). An example is shown in Fig. 3.18.

Following the same procedure explained in Sect. 2.2, we can reach the sensitivity of standard all sky analysis in the portion of the parameter space we selected.

The number of produced outliers in the all-sky radiometer analysis (see Sect. 3.1.7 and R. Abbott et al., 2022c) is around 500. Building the sky grids with the procedure described above for each outlier, the total number of searched sky points would be $\sim 8 \times 10^5$. Keeping the same data structure used for the O3 analysis, as shown in Sect. 2.2, we have an estimated total time for the FrequencyHough selection of candidates from the radiometers outliers of $\sim 10^3$ GPU hours (~ 1.5 months), that is roughly the 1.3% of the time needed to perform the full FrequencyHough all-sky analysis.

Nevertheless, the selection of 1 Hz portion from the pre-defined peakmaps with bandwidth of 5 Hz ($1 \text{ Hz for } 10 \leq f \leq 128 \text{ Hz}$) results in a strong inefficiency. A very useful tool is in this case the Band-Sampled-Data (BSD) data system shown in Sect. 2.2 (O J Piccinni et al., 2020). Using the BSD we can build customized peakmaps, not only with the needed frequency range, but also defining a proper coherence time specific for each outlier frequency, as opposite to the fixed coherence time broadly used for the FrequencyHough all-sky analysis. The possibility to increase the coherence time of the short-Fourier transforms used to build the peakmaps will grant an improvement in sensitivity. The implementation of the BSD GPU-FrequencyHough code will be reported in an upcoming paper.




Figure 3.18: Example of an all sky search on a software injection on real data. The signal is from an isolated source with no spindown, frequency 133.34 Hz, equatorial coordinates $(120^\circ, -15^\circ)$ and amplitude $h_0 = 3.5 \times 10^{-25}$. In the top plot we show the sky map generated by *PyStoch* at the injection frequency bin: the candidate appears as brighter regions of the sky around the source coordinates. In the bottom panels we plot the maps generated by the FrequencyHough transform, applied to peakmaps created from the O2 LIGO-H and LIGO-L detectors time series with the simulated signal injected, Doppler-corrected for the source coordinates.

Conclusions

The content of the current thesis shows that, even with well established data analysis algorithms, there is still large space for improvement, both developing new methods and using new programming techniques to have better performances, computing efficiency and increased search sensitivity.

Thanks to the GPGPU approach, it has been possible to use GPUs for all-sky continuous-wave searches, with a programming high-level effort that focused only on the code efficiency via the vectorization of the data. In this way, it has been possible to successfully exploit the parallel computing in which the GPUs are highly optimized, and achieve a very significant analysis speed-up, which could reach computing times of one order of magnitude smaller than the standard search (B. P. Abbott et al., 2019b; R. Abbott et al., 2022a; La Rosa et al., 2021). Thanks to this porting, it has been possible to run the analysis for the 75% of the frequency detector sensitivity interval for the O3 search in a couple of weeks rather than in a few months with a comparable amount of computing nodes, provided the availability of a modern GPU cluster like the Cineca Marconi100 (Cineca, 2021).

The same approach could be in principle applied to other search algorithms, given that it is expected that the trend of exponential growth of the capabilities of GPUs will continue, following also the growing interest in GPU computing across several different scientific and industrial fields. Porting of pipelines on GPUs, if they are adaptable to the vectorial data flow paradigm, is already showing to be very promising in various searches, and in future it can help the physics of gravitational waves on many different sectors, especially in data analysis.

One of the main elements that make the all-sky searches for rotating neutron stars difficult, from the computational point of view, is the large search parameter space size. Even with semi-coherent (B. P. Abbott et al., 2017a) approaches, this will be still challenging as — due to the Doppler effect — we are forced to analyze one by one all sky directions that are defined by a grid covering the whole sky. Since the number of points of the grid is very high, this brings to analyses that take $\sim 10^5$ single-device computation hours to be completed, even with the code running on GPUs, as it has been shown in Sect. 2.3.2.

We have seen, however, that in the framework of the stochastic gravitational-wave background searches, the radiometer method can run unmodeled analyses very efficiently over the full sky, behaving like a tomography of the sky and producing *SNR* maps through the search frequency band. This method has been used in recent years to search for point-like sources, and produced upper limits for persistent signals from neutron stars, that are approximately one order of magnitude higher than those we can obtain with dedicated continuous-wave searches.

Since the radiometer search may produce a large set of outliers, which have never been followed up so far with specific continuous-wave procedures, in future it will be very useful to apply the follow-up method that is described in Sect. 3.2.4 to all-sky radiometer outliers (R. Abbott et al., 2022c). The outcomes will be published in an upcoming paper. Thanks to the radiometer analysis, we can identify very quickly portions of the parameter space of interest. Due to the optimized GPU FrequencyHough algorithm, *ad hoc* analyses for the radiometer outliers can be performed. Given a candidate frequency, we can create customized peakmaps centered on it in the BSD framework (O. Piccinni et al., 2018). Defining a proper grid around the candidate sky coordinates allows one to explore a reduced portion of the sky producing a set of FrequencyHough candidates, which are in turn followed-up using the procedure described in Sect. 2.2. In this way we are able to verify thoroughly the selected candidates, and eventually put upper limits, which would be more stringent than those obtained by the sole radiometer search, which are compatible with those of all-sky continuous-wave analysis, but with a reduced delay and computational effort.

In particular, for the upcoming fourth observational run of the LIGO-Virgo-KAGRA collaborations, it will be possible to have quick results on wide portions of the parameters space via the radiometer search, far before the all-sky continuous-wave analysis is completed. Furthermore, with the combination of the two projects I have carried on, the search parameters space can be expanded without a significant increase of computational cost. An example regards to take into account more frequency derivatives, either enlarging the related search range or adding the second order spindown parameter.

Another promising direction that has been opened by the present work is the study of neutron stars in binary systems, for which a dedicated follow up needs to be tailored. However, it has been shown that the radiometer search can reproduce the signature of an emitting binary system, assuming to have good sensitivity and frequency resolution.

The achievements obtained with the thesis projects, together with the good prospects of further improvements in the near future, both for the algorithms and detector sensitivity, might give a timely contribution to the first detection of continuous gravitational-wave signals.

Acknowledgements

I want to thank the people in the LIGO and Virgo collaborations for all the things I have learned, all the answers to the many questions and doubts, for the data fundamental for all our work and for the many tools, codes, algorithms developed to analyze them, without which this thesis could not exist.

I thank also the Rome Virgo people, I do not see most of them for years and still I feel part of the group as the time I left the G23. The group that physically accompanied me for years at the LAPP, even with several lockdowns, curfews, endless telework, cannot be acknowledge enough. The LAPP Virgo group, two “generations” o PhD students, my office mates, colleagues and ex-colleagues, ex-flatmates, the Spanish and Latin American crew, the Italian crew plus the naturalized Italians, thanks to everyone, fellow sailors even during the biggest turmoils.

A special thanks goes to my supervisors. They went over the normal duties of a thesis advisor and tirelessly supported me in a PhD that, perhaps, has been one of the most adventurous, crossing a global pandemy, few accidents and, basically, life. Thanks a lot.

My last thanks goes to the persons who always been there to support me and always will be, to the persons randomly appeared in the last years in Annecy and joined the sailing in many different ways, and to the persons that on the way decided to disembark.

Bibliography

- Aasi, J et al. (2015). “Advanced ligo”. In: *Classical and quantum gravity* 32.7, p. 074001.
- (2016). “Prospects for observing and localizing gravitational-wave transients with Advanced LIGO and Advanced Virgo”. In: *Living Reviews in Relativity* 19.
- Abbott, B P et al. (2016). “Observation of gravitational waves from a binary black hole merger”. In: *Physical review letters* 116.6, p. 061102.
- (2017a). “All-sky search for periodic gravitational waves in the O1 LIGO data”. In: *Physical Review D* 96.6, p. 062002.
- (2017b). “Directional limits on persistent gravitational waves from Advanced LIGO’s first observing run”. In: *Physical review letters* 118.12, p. 121102.
- (2017c). “GW170814: A three-detector observation of gravitational waves from a binary black hole coalescence”. In: *Physical Review Letters* 119.14, p. 141101.
- (2017d). “GW170817: observation of gravitational waves from a binary neutron star inspiral”. In: *Physical Review Letters* 119.16, p. 161101.
- (2018). “GW170817: Measurements of neutron star radii and equation of state”. In: *Physical review letters* 121.16, p. 161101.
- (2019a). *Advanced LIGO, Advanced Virgo and KAGRA observing run plans*. LIGO Scientific Collaboration and Virgo Collaboration and KAGRA Collaboration.
- (2019b). “All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO O2 data”. In: *Physical Review D* 100.2, p. 024004.
- (2019c). “Directional limits on persistent gravitational waves using data from Advanced LIGO’s first two observing runs”. In: *Physical Review D* 100.6, p. 062001.
- (2020a). “Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA”. In: *Living reviews in relativity* 23.1, pp. 1–69.
- (2020b). “Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA”. In: *Living Reviews in Relativity* 23.1, p. 3.
- Abbott, R et al. (2021a). *Gravitational Wave Open Science Center: Event List*. Online; accessed 30-March-2022. LIGO Scientific Collaboration and Virgo Collaboration. URL: <https://www.gw-openscience.org/eventapi/html/allevvents/>.
- (2021b). *GWTC-3, a third catalog of gravitational wave detections*. Online; accessed 30-March-2022. LIGO Scientific Collaboration and Virgo Collaboration. URL: <https://www.ligo.org/science/Publication-03bCatalog/>.

- Abbott, R et al. (2021c). “Narrowband searches for continuous and long-duration transient gravitational waves from known pulsars in the LIGO-Virgo third observing run”. In: *arXiv preprint arXiv:2112.10990*.
- (2021d). “Search for anisotropic gravitational-wave backgrounds using data from Advanced LIGO and Advanced Virgo’s first three observing runs”. In: *Physical Review D* 104.2, p. 022005.
- (2021e). “Searches for continuous gravitational waves from young supernova remnants in the early third observing run of Advanced LIGO and Virgo”. In: *The Astrophysical Journal* 921.1, p. 80.
- (2022a). “All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data”. In: *arXiv preprint arXiv:2201.00697*.
- (2022b). “All-sky search for gravitational wave emission from scalar boson clouds around spinning black holes in LIGO O3 data”. In: *Physical Review D* 105.10, p. 102001.
- (2022c). “All-sky, all-frequency directional search for persistent gravitational waves from Advanced LIGO’s and Advanced Virgo’s first three observing runs”. In: *Physical Review D* 105.12, p. 122001.
- (2022d). “Constraints on dark photon dark matter using data from LIGO’s and Virgo’s third observing run”. In: *Physical review D* 105.6, p. 063030.
- (2022e). “Search for continuous gravitational wave emission from the Milky Way center in O3 LIGO–Virgo data”. In: *arXiv preprint arXiv:2204.04523*.
- Acernese, F et al. (2014). “Advanced Virgo: a second-generation interferometric gravitational wave detector”. In: *Classical and Quantum Gravity* 32.2, p. 024001.
- (2022). “Calibration of advanced Virgo and reconstruction of the detector strain $h(t)$ during the observing run O3”. In: *Classical and Quantum Gravity* 39.4, p. 045006.
- Aghanim, Nabila et al. (2020). “Planck 2018 results-VI. Cosmological parameters”. In: *Astronomy & Astrophysics* 641, A6.
- Ain, A, P Dalvi, et al. (2015). “Fast gravitational wave radiometry using data folding”. In: *Physical Review D* 92.2, p. 022003.
- Ain, A, J Suresh, et al. (2018). “Very fast stochastic gravitational wave background map making using folded data”. In: *Physical Review D* 98.2, p. 024001.
- Allen, Bruce et al. (1999). “Detecting a stochastic background of gravitational radiation: Signal processing strategies and sensitivities”. In: *Physical Review D* 59.10, p. 102001.
- Antonucci, F et al. (2008). “Detection of periodic gravitational wave sources by Hough transform in the f versus plane”. In: *Classical and Quantum Gravity* 25.18, p. 184015.
- Astone, P et al. (2014). “Method for all-sky searches of continuous gravitational wave signals using the frequency-Hough transform”. In: *Physical Review D* 90.4, p. 042002.
- Astone, P, D Babusci, et al. (2003). “Methods and results of the IGEC search for burst gravitational waves in the years 1997–2000”. In: *Physical Review D* 68.2, p. 022001.

- Astone, P, S Frasca, et al. (2005). “The short FFT database and the peak map for the hierarchical search of periodic sources”. In: *Classical and Quantum Gravity* 22.18, S1197.
- ATNF (2021). *The ATNF Pulsar Catalogue*. URL: <http://www.atnf.csiro.au/people/pulsar/psrcat/>.
- Bailes, M et al. (1997). “Discovery of four isolated millisecond pulsars”. In: *The Astrophysical Journal* 481.1, p. 386.
- Ballard, D H (1981). “Generalizing the Hough transform to detect arbitrary shapes”. In: *Pattern recognition* 13.2, pp. 111–122.
- Ballmer, S W (2006). “A radiometer for stochastic gravitational waves”. In: *Classical and Quantum Gravity* 23.8, S179.
- Barsotti, Lisa et al. (2018). *Updated Advanced LIGO sensitivity design curve*. URL: <https://dcc.ligo.org/LIGO-T1800044/public>.
- Bisnovatyi-Kogan, G S (1992). “The Neutron Star Population in the Galaxy”. In: *Symposium-International Astronomical Union*. Vol. 149. Cambridge University Press, pp. 379–382.
- Burrows, A (2000). “Supernova explosions in the universe”. In: *Nature* 403.6771, pp. 727–733.
- Cahillane, Craig et al. (2017). “Calibration uncertainty for Advanced LIGO’s first and second observing runs”. In: *Physical Review D* 96.10, p. 102001.
- Caride, Santiago et al. (2019). “How to search for gravitational waves from r-modes of known pulsars”. In: *Physical Review D* 100.6, p. 064013.
- Carroll, B W et al. (2017). *An introduction to modern astrophysics*. Cambridge University Press.
- Chandrasekhar, S (1970). “Solutions of two problems in the theory of gravitational radiation”. In: *Physical Review Letters* 24.11, p. 611.
- Cineca (2021). *Marconi100*. URL: www.hpc.cineca.it/hardware/marconi100.
- Covas, PB et al. (2018). “Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO”. In: *Physical Review D* 97.8, p. 082002.
- Danzmann, Karsten et al. (2017). “Laser interferometer space antenna”. In: *arXiv preprint arXiv:1702.00786*.
- Denzel, Philipp et al. (2021). “The Hubble constant from eight time-delay galaxy lenses”. In: *Monthly Notices of the Royal Astronomical Society* 501.1, pp. 784–801.
- Dhurandhar, Sanjeev V et al. (2001). “Searching for continuous gravitational wave sources in binary systems”. In: *Physical Review D* 63.12, p. 122001.
- Dietrich, Tim et al. (2020). “Multimessenger constraints on the neutron-star equation of state and the Hubble constant”. In: *Science* 370.6523, pp. 1450–1453.
- Duda, R O et al. (1972). “Use of the Hough transformation to detect lines and curves in pictures”. In: *Communications of the ACM* 15.1, pp. 11–15.
- Einstein, A (1918). “Näherungsweise integration der feldgleichungen der gravitation”. In: *Albert Einstein: Akademie-Vorträge: Sitzungsberichte der Preußischen Akademie der Wissenschaften 1914–1932*, pp. 99–108.
- Evans, Matthew et al. (2021). “A horizon study for cosmic explorer: science, observatories, and community”. In: *arXiv preprint arXiv:2109.09882*.

- Ferrari, V et al. (2020). *General Relativity and Its Applications: Black Holes, Compact Stars and Gravitational Waves*. CRC Press.
- Ferrari, Valeria et al. (2003). “Gravitational waves from newly born, hot neutron stars”. In: *Monthly Notices of the Royal Astronomical Society* 342.2, pp. 629–638.
- Frasca, S et al. (2018). *SNAG*. URL: grwavs.f.roma1.infn.it/snag.
- Frasca, S et al. (2005). “Evaluation of sensitivity and computing power for the Virgo hierarchical search for periodic sources”. In: *Classical and Quantum Gravity* 22.18, S1013.
- Freedman, Wendy L (2021). “Measurements of the Hubble constant: tensions in perspective”. In: *The Astrophysical Journal* 919.1, p. 16.
- Friedman, John L et al. (1978). “Secular instability of rotating Newtonian stars”. In: *The Astrophysical Journal* 222, pp. 281–296.
- Goncharov, Boris et al. (2018). “All-sky radiometer for narrowband gravitational waves using folded data”. In: *Physical Review D* 98.6, p. 064018.
- Gorski, K M et al. (2005). “HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere”. In: *The Astrophysical Journal* 622.2, p. 759.
- Hartle, J B (2003). *Gravity: an introduction to Einstein’s general relativity*. Addison-Wesley, Pearson.
- Hough, P V C (1962). *Method and means for recognizing complex patterns*. Tech. rep.
- Hulse, Russell A et al. (1975). “Discovery of a pulsar in a binary system”. In: *The Astrophysical Journal* 195, pp. L51–L53.
- Kalogera, Vassiliki et al. (1996). “The maximum mass of a neutron star”. In: *The Astrophysical Journal* 470.1, p. L61.
- La Rosa, Iuri et al. (2021). “Continuous Gravitational-Wave Data Analysis with General Purpose Computing on Graphic Processing Units”. In: *Universe* 7.7, p. 218.
- Lasky, Paul D (2015). “Gravitational waves from neutron stars: a review”. In: *Publications of the Astronomical Society of Australia* 32.
- Lattimer, J M (2012). “The nuclear equation of state and neutron star masses”. In: *Annual Review of Nuclear and Particle Science* 62, pp. 485–515.
- (2021). “Neutron Stars and the Nuclear Matter Equation of State”. In: *Annual Review of Nuclear and Particle Science* 71, pp. 433–464.
- Lattimer, J M and M Prakash (2004). “The physics of neutron stars”. In: *science* 304.5670, pp. 536–542.
- Le Tiec, A et al. (2017). *An Overview of Gravitational Waves: Theory, Sources and Detection*. World Scientific.
- Leaci, P et al. (2017). “Novel directed search strategy to detect continuous gravitational waves from neutron stars in low-and high-eccentricity binary systems”. In: *Physical Review D* 95.12, p. 122001.
- LIGO Caltech (2021). *CIT*. URL: www.ligo.caltech.edu.
- LIGO collaboration (2017). *Selected pulsar injection parameters for injection epoch O2*. URL: https://www.gw-openscience.org/O2_injection_params/.
- (2019). *Selected pulsar injection parameters for injection epoch O3*. URL: https://www.gw-openscience.org/O3/O3April1_injection_parameters/.
- LIGO Scientific Collaboration (2018). *LIGO Algorithm Library - LALSuite*. free software (GPL). DOI: [10.7935/GT1W-FZ16](https://doi.org/10.7935/GT1W-FZ16).

- Lyne, AG et al. (1987). “The discovery of a millisecond pulsar in the globular cluster M28”. In: *Nature* 328.6129, pp. 399–401.
- Maggiore, M (2000). “Gravitational wave experiments and early universe cosmology”. In: *Physics Reports* 331, pp. 283–367.
- (2008). *Gravitational waves: Volume 1: Theory and experiments*. Vol. 1. Oxford university press.
- (2018). *Gravitational Waves: Volume 2: Astrophysics and Cosmology*. Oxford University Press.
- Maggiore, Michele et al. (2020). “Science case for the Einstein telescope”. In: *Journal of Cosmology and Astroparticle Physics* 2020.03, p. 050.
- Al-Mamun, M et al. (2021). “Combining electromagnetic and gravitational-wave constraints on neutron-star masses and radii”. In: *Physical Review Letters* 126.6, p. 061101.
- Mata, Andrew et al. (2021). *Applying gating to stochastic searches in O3*. Tech. rep. URL: <https://dcc.ligo.org/LIGO-P2000546/public>.
- Mathworks (2021). *MATLAB*. URL: www.mathworks.com/products/matlab.
- Mereghetti, Sandro (2008). “The strongest cosmic magnets: soft gamma-ray repeaters and anomalous X-ray pulsars”. In: *The Astronomy and Astrophysics Review* 15.4, pp. 225–287.
- Messenger, C (2010). *Understanding the sensitivity of the stochastic radiometer analysis in terms of the strain tensor amplitude*. Tech. rep.
- Mitra, S et al. (2008). “Gravitational wave radiometry: Mapping a stochastic gravitational wave background”. In: *Physical Review D* 77.4, p. 042002.
- Moore, Christopher J et al. (2014). “Gravitational-wave sensitivity curves”. In: *Classical and Quantum Gravity* 32.1, p. 015014.
- Nagase, Fumiaki (1989). “Accretion-powered X-ray pulsars”. In: *Publications of the Astronomical Society of Japan* 41, pp. 1–79.
- Nitz, Alex et al. (June 2022). *gwastro/pycbc: v2.0.4 release of PyCBC*. Version v2.0.4. DOI: [10.5281/zenodo.6646669](https://doi.org/10.5281/zenodo.6646669). URL: <https://doi.org/10.5281/zenodo.6646669>.
- Oppenheimer, J R et al. (1939). “On massive neutron cores”. In: *Physical Review* 55.4, p. 374.
- Palomba, Cristiano et al. (2019). “Direct constraints on the ultralight boson mass from searches of continuous gravitational waves”. In: *Physical review letters* 123.17, p. 171101.
- Piccinni, O J et al. (2020). “Directed search for continuous gravitational-wave signals from the Galactic Center in the Advanced LIGO second observing run”. In: *Physical Review D* 101.8, p. 082004.
- Piccinni, OJ et al. (2018). “A new data analysis framework for the search of continuous gravitational wave signals”. In: *Classical and Quantum Gravity* 36.1, p. 015008.
- Pitkin, M. et al. (2011). “Gravitational wave detection by interferometry (ground and space)”. In: *Living Reviews in Relativity* 14.1, p. 5.
- Powell, J et al. (2017). “Inferring the core-collapse supernova explosion mechanism with three-dimensional gravitational-wave simulations”. In: *arXiv preprint arXiv:1709.00955*.

- Regimbau, Tania (2011). “The astrophysical gravitational wave stochastic background”. In: *Research in Astronomy and Astrophysics* 11.4, p. 369.
- Rhoades Jr, Clifford E et al. (1974). “Maximum mass of a neutron star”. In: *Physical Review Letters* 32.6, p. 324.
- Romano, J D et al. (2017). “Detection methods for stochastic gravitational-wave backgrounds: a unified treatment”. In: *Living Reviews in Relativity* 20.1, p. 2.
- Romano, Joseph D et al. (2015). “Phase-coherent mapping of gravitational-wave backgrounds using ground-based laser interferometers”. In: *Physical Review D* 92.4, p. 042003.
- Saulson, P R (2017). *Fundamentals of interferometric gravitational wave detectors*. World Scientific.
- Scipy developers (2021). *Scipy*. URL: www.scipy.org.
- Shapiro, S L et al. (2008). *Black holes, white dwarfs, and neutron stars: The physics of compact objects*. John Wiley & Sons.
- Sieniawska, M et al. (2019). “Continuous gravitational waves from neutron stars: current status and prospects”. In: *Universe* 5.11, p. 217.
- Somiya, K (2012). “Detector configuration of KAGRA—the Japanese cryogenic gravitational-wave detector”. In: *Classical and Quantum Gravity* 29.12, p. 124007.
- Steiner, A W (2020). *P pdot diagram*. URL: <https://neutronstars.utk.edu/code/nstar-plot/ppdot.html#p-pdot-diagram>.
- Steiner, A W et al. (2010). “The equation of state from observed masses and radii of neutron stars”. In: *The Astrophysical Journal* 722.1, p. 33.
- Sun, L et al. (2020). “Characterization of systematic error in Advanced LIGO calibration”. In: *Classical and Quantum Gravity* 37.22, p. 225008.
- Taylor, JH et al. (1976). “Further observations of the binary pulsar PSR 1913+ 16”. In: *The Astrophysical Journal* 206, pp. L53–L58.
- The HEALpy contributors (2022). *HEALpy*. URL: <https://healpy.readthedocs.io/en/latest/>.
- The Khronos Group (2021). *OpenCL*. URL: www.khronos.org/opencvl.
- The LIGO, Virgo, KAGRA collaborations (2022). *stochastic.m*. URL: <https://svn.ligo.caltech.edu/svn/matapps>.
- The NVIDIA Corporation (2021). *CUDA*. URL: www.developer.nvidia.com/cuda-toolkit.
- (2022). *Tesla V100 datasheet*. URL: <https://images.nvidia.com/content/technologies/volta/pdf/volta-v100-datasheet-update-us-1165301-r5.pdf>.
- The TensorFlow Authors (2021). *TensorFlow*. URL: www.tensorflow.org.
- Thrane, Eric et al. (2009). “Probing the anisotropies of a stochastic gravitational-wave background using a network of ground-based laser interferometers”. In: *Physical Review D* 80.12, p. 122002.
- Tolman, R C (1939). “Static solutions of Einstein’s field equations for spheres of fluid”. In: *Physical Review* 55.4, p. 364.
- Vajente, G et al. (2020). “Machine-learning nonstationary noise out of gravitational-wave detectors”. In: *Physical Review D* 101.4, p. 042003.
- Verkindt, D. (Oct. 2021). *Sensitivity of the Advanced Virgo detector during O2 and O3 runs*. URL: <https://tds.virgo-gw.eu/?content=3&r=19594>.

- Weber, Joseph (1960). “Detection and generation of gravitational waves”. In: *Physical Review* 117.1, p. 306.
- Whelan, John T et al. (2014). “Treatment of calibration uncertainty in multi-baseline cross-correlation searches for gravitational waves”. In: *Journal of Physics: Conference Series*. Vol. 484. 1. IOP Publishing, p. 012027.
- Zink, Burkhard et al. (2010). “Frequency band of the f-mode Chandrasekhar-Friedman-Schutz instability”. In: *Physical Review D* 81.8, p. 084055.
- Zubakov, LA et al. (1962). *Extraction of signals from noise*. Prentice-Hall, Inc., New Jersey.
- Zweizig, John et al. (2021). *Information on self-gating of $h(t)$ used in O3 continuous-wave and stochastic searches*. Tech. rep. URL: <https://dcc.ligo.org/T2000384/public>.