



UNIVERSITÀ DI PISA
TESI DI DOTTORATO IN FISICA
XXXII CICLO

**Characterisation and Mitigation of
Non-stationary Noise in Advanced
Gravitational Wave Detectors**

Author:
Francesco DI RENZO

Supervisor:
Dr. Giancarlo CELLA
Internal referee:
Prof. Francesco FIDECARO

*A thesis submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy*

in the

Virgo Collaboration
Department of Physics

UNIVERSITÀ DI PISA

*Abstract*Doctorate School in Physics
Department of Physics

Doctor of Philosophy

**Characterisation and Mitigation of Non-stationary Noise in Advanced Gravitational
Wave Detectors**

by Francesco DI RENZO

In December 2019, the LIGO Scientific Collaboration and the Virgo Collaboration have published the results of their first two joint observing runs (O1 and O2), describing the source properties of ten binary black hole (BBH) and one binary neutron star (BNS) events. Their validation has been made employing state of the art Data Analysis techniques. All of them rely, to some extent, on the assumption to know the statistical properties of the detector noise, from which the gravitational signals are extracted. Moreover, their performances are optimal, with respect to certain criteria, if the noise distribution is stationary and Gaussian. To this purpose, in this Thesis work we have studied several strategies aimed at the verification of the previous two hypotheses, and the characterisation of the detector noise. Once a specific noise feature, detrimental for gravitational wave searches, was found, we have proceeded to the investigation of its causes and some mitigation strategies. The techniques that we have implemented have been selected from many fields of research, like Digital Image Processing and state of the art Machine Learning. Two original contributions have been introduced. One consists of a new method for the identification of generic non-stationary noise, from the variations in the empirical distribution of the signal RMS value. The other is a wavelet-based, instantaneous causality statistic, specifically aimed at the study of transient noises. These aim to improve upon other existing strategies and have been applied for the investigation of specific noise issue in Advanced Virgo detector data.

Contents

Abstract	iii
Notation	xiii
Introduction	1
Outline of the Thesis	4
1 Gravitational waves: theory, detectors and detections	7
1.1 Gravitational wave theory overview	8
1.1.1 Linearised gravitational wave solutions	8
1.1.2 GWs in TT -gauge	12
1.1.3 Emission of gravitational waves	14
1.1.4 Effects of GWs on test masses	16
1.2 Sources of gravitational waves	18
1.2.1 Compact binary coalescence	21
1.2.2 Periodic sources	24
1.2.3 CCSNe and burst signals	26
1.2.4 Stochastic background	27
1.3 Gravitational wave detectors	29
1.3.1 Ground based interferometric detectors	30
1.3.2 Advanced interferometric detectors	33
1.3.3 Some sensitivity benchmarks: DARM, $h_{\text{rec}}(t)$, BNS range	38
1.3.4 A network of GW detectors	41
1.3.5 New detectors and next generations	43
1.4 Detections of the Advanced detectors	43
1.4.1 GW150914: the first detection	44
1.4.2 GWTC-1: an overview of current gravitational wave detections	49
2 Detector Noise	53
2.1 Prelude: statistical description of data	54
2.1.1 Probability spaces and random variables	54
2.1.2 Stochastic processes and time series	57
Stationary processes	57
Expectation values and distribution moments	58
Gaussian processes	59
Spectral representation of stationary processes	60
Cramèr representation and physical interpretation of the PSD	61

	Ergodic theorem	62
2.1.3	Some “finer-grain” notions of non-stationarity	63
	Locally stationary stochastic processes	64
	Stationary models and parametric notion of stationarity	65
2.1.4	Signal detection and the matched filter technique	67
2.1.5	Non-stationary noise and the “non-optimal” detection statistic	69
2.1.6	Time series analysis and spectral estimations	73
	Spectral estimations: the periodogram and the Bartlett’s method	74
	Welch’s method	76
	Median averaged spectrum	78
2.2	Noise sources in GW detectors	81
2.2.1	Noise origins and general classification	83
2.2.2	Fundamental noise sources and sensitivity	85
	The standard quantum limit	85
	Seismic noise	86
	Newtonian	86
	Thermal noise	87
2.3	Stationary noise	87
2.3.1	Effects of stationary spectral noise on search pipelines	87
2.3.2	Investigation of stationary noise	88
2.3.3	Coherence analysis	89
	Magnitude squared coherence	90
	Multiple-coherence	91
	Complex coherence and phase angle spectrum	93
2.3.4	Non-linear noise	95
2.4	Non-stationary noise	97
2.4.1	Glitches	98
2.4.2	Slow non-stationarities	99
2.4.3	Data Quality flags	101
2.4.4	Effects of vetoing in GW searches	103
3	Methods of investigation of non-stationary noise	105
3.1	Detection methods of non-stationary noise	106
3.1.1	Stationarity of the PSD: the Priestley-Subba Rao test	107
3.1.2	A BLRMS based test for noise stationarity	108
	Method of computation of the BLRMS	108
	Lines removal	109
	Glitches identification and removal	111
	Test on the empirical distribution function of the BLRMS	112
	Comments and choice of parameters	116
3.1.3	Other stationarity tests	118
	Variance of the detection statistics	118
	Discrete, orthogonal wavelet transform	119
3.2	Detection and investigation of glitches	120

3.3	Detection methods of non-Gaussianities	122
3.3.1	The Rayleigh test	123
3.3.2	Critical values and analysis of real data	124
3.4	Line tracker tool for frequency non-stationary noise	127
3.4.1	Image filtering	128
3.4.2	Frequencies series reconstruction	129
3.5	Investigating the correlation with the auxiliary channels	130
3.5.1	Cross-correlation analysis	131
3.5.2	Data transformation and non-linear methods	133
	Mutual Information	133
3.5.3	Multiple linear regression analysis	135
	Principal component regression	137
	Other regression methods: subset selection and shrinkage methods	140
	Residuals diagnostics and model selection	141
3.5.4	NonNA: Non-stationary Noise Analysis tools	144
3.5.5	Possible extensions: Machine Learning implementation	147
3.6	Wavelet coherence analysis of glitches	151
3.6.1	The continuous wavelet transform	152
	Algorithm implementation	155
	Wavelet power spectrum and significance level	156
	Averaging in the time-scale plane	158
3.6.2	Wavelet cross-spectrum and wavelet coherence	159
	Wavelet cross-spectrum and phase angle spectrum	159
	Wavelet coherence	161
	Wavelet cross-correlation	163
3.6.3	Instantaneous time delay through wavelet phase angle spectrum	164
	Phase angle spectrum between retarded signals	164
	Instantaneous time delay statistic	165
	Tests on simulated glitches	166
3.6.4	Discussion and applications	168
	On the possibility of a novel wavelet based causality test	169
	Analysis of “mystery glitches” and BNS range drops	170
3.7	On the concept of causality	173
3.7.1	Granger-Geweke causality	174
	Geweke formulation, limits and extensions	175
	Application to the study of interferometer fast unlocks	177
3.7.2	Convergent Cross Mapping	178
	The Takens’ reconstruction theorem	178
	Application to the study of the effects of squeezing	180
3.7.3	Some final thoughts about the application of causality studies to GW detector’s data	182

4	Noise artefacts mitigation	185
4.1	NNetFix: a Neural Network to Fix glitches on signals	186
4.1.1	The Multi-layer Perceptron and method description	186
4.1.2	Training by means of simulated waveforms	188
4.1.3	Results on injected and real BBH signals	189
4.2	Discussion and future developments	190
5	Conclusions	193
A	Examples of noise studies with Virgo data	195
A.1	Beam splitter control noise at 150 Hz	195
A.2	“Flat noise” investigations	198
A.3	The 83 Hz wandering line	202
B	Some notable distributions	207
B.1	Rayleigh	207
B.2	Gamma	208
	Bibliography	209
	LVC internal bibliography	233

List of Figures

1.1	GW polarisations in Einstein's GR	17
1.2	Strain amplitudes and frequency ranges for various GW sources	20
1.3	Mass space for compact binary coalescences	23
1.4	Expected rate of compact binary coalescences	24
1.5	Simplified optical scheme of the Advanced Virgo detector	34
1.6	Sensitivities of the Advanced LIGO and Advanced Virgo detectors	37
1.7	Prospects on Kagra sensitivity	38
1.8	Evolution of BNS range for the Advanced Detectors observing runs	41
1.9	Significance for the events detected during O1	45
1.10	Key results on GW150914	48
1.11	Sky localisation of the O1-O2 events.	50
1.12	Estimated component masses and final masses and spins for O1-O2 events . .	51
2.1	Welch's method for PSD estimation	76
2.2	Effects of padding in PSD estimations	77
2.3	Spectrogram of 600 seconds of Virgo data	80
2.4	Comparison of Welch's method and mean-median average method	82
2.5	Noise budget for Advanced Virgo	84
2.6	Depiction of the radiation pressure and shot noise	86
2.7	Spectral features in Advanced Virgo data	93
2.8	Example of study of modulated noise	96
2.9	Example of scattered light glitch in Advanced Virgo data	98
2.10	Median normalised spectrogram of the <i>Amplitude Spectral Density</i> (ASD) corresponding to figure 2.3.	99
2.11	Example of amplitude non-stationarity in Advanced Virgo data	100
2.12	Example of frequency non-stationary noise in Advanced Virgo data	101
3.1	ASD with and without spectral lines	110
3.2	Spectral lines removal before BLRMS computation	112
3.3	Glitch removal from BLRMS time series	113
3.4	Stationarity test: example of BLRMS non-stationarity	114
3.5	p -map of stationarity	115
3.6	Stationarity test: dependency on the sample length	117
3.7	Spectrogram of 10 minutes of data and corresponding Rayleigh Gaussianity test	125
3.8	Rayleigh-gram for the investigation of non-Gaussianities and non-stationarities	126
3.9	Line tracker for frequency non-stationarities	127

3.10	Line tracker with masking	129
3.11	Example of correlation analysis	132
3.12	Example of explained variance by a set of auxiliary channels	138
3.13	Flowchart of the preprocessing stage for the cross-correlation and regression analysis tools.	145
3.14	Flowchart of the cross-correlation analysis tool.	146
3.15	Flowchart of the regression analysis tool.	147
3.16	Output of the NonNA cross-correlation tool	148
3.17	Example of regression analysis with NonNA	149
3.18	Signals representations	151
3.19	Morlet wavelet	152
3.20	Scalogram of a glitch	155
3.21	Example of wavelet cross-spectrum	160
3.22	Example of wavelet coherence	162
3.23	Scalogram of a simulated glitch	167
3.24	Phase angle about the centre time of the glitch	168
3.25	Instantaneous time delay for simulated glitches	169
3.26	Application of WATERLOO to the study of some mystery glitches	171
3.27	Example of wavelet cross-spectrum	173
3.28	CCM causality	181
4.1	Saturation glitch in LIGO Livingston data in correspondence of GW170817	186
4.2	Performances of NNetFix reconstruction in terms of SNR error	189
4.3	Reconstruction of GW150914 by means of NNetFix	191
A.1	Beam splitter control noise at 150 Hz	196
A.2	Correlation between BS control mode and noise at 150 Hz	197
A.3	Effects of the flat noise on Virgo ASD	199
A.4	Application of the BRiSTOL stationarity test to the study of flat noise	200
A.5	Application of the Rayleigh test to the investigation of flat noise	201
A.6	ASD in correspondence of the 83 Hz wandering line	202
A.7	Wandering line of January 2020	204

List of Tables

3.1	Main input arguments to pass to the NonNA analysis tools.	144
-----	---	-----

Notation

Throughout the manuscript, we tried to consistently maintain the notation in use in the most recent LIGO and Virgo collaboration articles, like the *first Gravitational Wave Transient Catalogue* (GWTC-1) [1] or the *detector noise and signal extraction guide* [2]. This is also the one used, in general, in the reference textbook by Creighton and Anderson [3], members of the LIGO collaboration. Here we will summarise the most relevant conventions and notation choices. Further comments shall follow about the specific usage of symbols and notations in different chapters and contexts.

We use the symbol “:=” when a definition is introduced, and the *equivalence symbol* “ \equiv ” to denote an identity. When there is no ambiguity, we will simply denote one or the other by the equal sign “=”.

Coordinates on space-time are denoted by Greek letters, such as α, β, \dots or μ, ν, \dots , running from 0 to 3, with zero referring to the time coordinate. Sometimes it is useful to refer to the space and time components separately, so we use Latin letters from the middle of the alphabet, i, j, \dots , valued 1, 2 or 3, to stand for the space components alone. 3-dimensional spatial vectors are also sometimes indicated with bold-faced letters, like \mathbf{x} . The speed of light is always labelled by c , for this reason the partial derivative with respect to the zeroth space-time component is related to the one with respect to time by: $\partial_0 := c^{-1}\partial_t$. We make use of this convention in chapter 1. The metric signature convention is *space-like*, $(-, +, +, +)$, and all the other sign-conventions for General Relativity are those of Misner, Thorne and Wheeler [4]. Differently from [3, 4], we will not use bold sans-serif letter to represent tensor quantities, to which we prefer using indices inside equations.

Large part of the notation about random variables and stochastic processes (chapter 2 and 3) is adopted from [5] and [6]. Random variables are denoted with upper case letters, *e.g.* X . These take on values in some domain, and if we want to consider a particular observation of them (that is, X has been sampled and observed to have a particular value in the domain) then that *non-random value* is denoted by lower case, *e.g.* x . Probabilities over random variables are denoted as $P(X)$, or P_X , but a typical place we made use of the above upper case/lower case convention is when the outcome is “spelled out” as $P(X = x)$, or $P_X(x)$, representing the probability that the random variable X takes on the fixed value x . This distinction is made to emphasise the logical difference between the two concepts, or whether we are talking about actual experimental outcomes or the statistical properties of the corresponding estimators, although sometimes, when there is a clear distinction from the context, we will be sloppy and use lower case letters. This is the case for example when the variable is used as a subscript to label some other quantities. Expectation values are denoted both with $E[\dots]$ and with the “Quantum Mechanics” notation $\langle \dots \rangle$. We will usually assume the “ergodicity” of the processes under study, hence we will make no notation distinction in

the averages computed on their (inaccessible) ensembles and those computed with respect to time.

Following [3, 5], our conventions for the Fourier transform are as follows. For continuous quantities, the forward and inverse Fourier transforms are given by:

$$\tilde{x}(f) := \int_{-\infty}^{+\infty} x(t)e^{-2\pi i f t} dt, \quad x(t) := \int_{-\infty}^{+\infty} \tilde{x}(f)e^{2\pi i f t} df.$$

In most practical situations, the quantity $x(t)$ is sampled for a finite time T , at a *sampling frequency* f_S , so that we will have N values $x_n := x(n/f_S)$, where $n = 0, 1, \dots, N-1$, and $T = N/f_S$. Then, the discrete forward and inverse Fourier transforms are:

$$\tilde{x}_k := \frac{1}{f_S} \sum_{n=0}^{N-1} x_n e^{-2\pi i n k / N}, \quad x_n := \frac{1}{T} \sum_{k=0}^{N-1} \tilde{x}_k e^{2\pi i n k / N}.$$

This normalisation choice yields Fourier components \tilde{x}_k with the same units as the continuous Fourier transform $\tilde{x}(f)$. Consequently, the former can be interpreted as an approximation to the value of the latter at frequencies k/T : $\tilde{x}_k \approx \tilde{x}(k/T)$, for $0 \leq k \leq \lceil N/2 \rceil$ and $\tilde{x}_k \approx \tilde{x}((k-N)/T)$ for $\lceil N/2 \rceil \leq k \leq N$ (negative frequencies). Here the *ceiling brackets* (or *modulus*) “[...]” correspond to the smallest integer greater than or equal to what’s inside brackets. The DC component is $k = 0$, and $k = \lceil N/2 - 1 \rceil$ corresponds to the Nyquist-frequency. Analogously to [3], we prefer to work with “double sided” power spectral densities, defined as the the Fourier transform of the auto-correlation function R of the *stationary* stochastic process $X(t)$:

$$S(f) := \int_{-\infty}^{+\infty} R(\tau) e^{-2\pi i f \tau} d\tau$$

which is the quantity that is most convenient to use for mathematical calculations. However, the quantities that are directly measured by filtering procedures are the “one sided” (superscript “1s.”) spectral density functions:

$$\begin{aligned} S^{1s.}(f) &= 2 \int_0^{+\infty} R(\tau) e^{-2\pi i f \tau} d\tau \\ &= 4 \int_0^{+\infty} R(\tau) \cos(2\pi \tau f) d\tau = 2S(f) \end{aligned} \quad (1a)$$

defined for $0 < f < +\infty$. These are also the quantities reported in figures.

Different vectorial or matricial notations are sometimes implemented to replace indices and conveniently manipulate ordered sets of quantities. For example, in the context of signal analysis, the italic bold-faced letter $\mathbf{x} := \{x_n : n = 0, \dots, N-1\}$ is meant to collectively represent all the ordered data samples as a *vectorial quantity*: x_n with $n = 0, 1, \dots, N-1$. In the context of multiple linear regression (subsection 3.5.3), where the data series corresponding to p different channels are recorded, these will be collectively represented with the matrix $\mathbf{X} := (x_1, \dots, x_p)$. Other notations have been adopted, although often temporarily for few sections and calculations, and are promptly introduced and described in the main text.

Acronyms

AIC Akaike Information Criterion.	GW Gravitational Wave.
ANOVA analysis of variance.	GWTC-1 first Gravitational Wave Transient Catalogue.
AR Auto Regressive.	IMC Inpute Mode Cleaner.
ARMA Auto Regressive, Moving Average.	LSP Locally Stationary Processes.
ASD Amplitude Spectral Density.	LVC LIGO Scientific Collaboration and Virgo Collaboration.
BBH Binary Black Hole.	MA Moving Average.
BIC Bayesian Information Criterion.	MAD Median Absolute Deviation.
BLRMS Band-limited Root Mean Square.	MLP Multi-layer Perceptron.
BLUE Best Linear Unbiased Estimator.	MSE Mean Squared Error.
BNS Binary Neutron Star.	NN Neural Network.
BS beam splitter.	OLS Ordinary Least Squares.
c.d.f. cumulative distribution function.	p.d.f. probability density function.
CBC Compact Binary Coalescence.	PCA Principal Component Analysis.
CCM Convergent Cross-Mapping.	PR Power Recycling.
CLM Classical Linear Model.	PSD Power Spectral Density.
CNN Convolutional Neural Network.	PSR Priestley-Subba Rao test.
CW Continuous Waves.	rms Root Mean Square.
cWB coherent Wave Burst.	SGWB Stochastic Gravitational-wave Background.
DFT Discrete Fourier Transform.	SNR Signal-to-noise Ratio.
DGP data-generating process.	SR Signal Recycling.
DQ Data Quality.	SVD Singular Value Decomposition.
DQR Data Quality Report.	
EOB Effective One Body.	
fft Fast Fourier transform.	
GR General Relativity.	
GRB Gamma-Ray Burst.	

Physical Constants

Values in parentheses are the uncertainties (68% confidence range) in the last two digits.

Speed of Light	$c = 2.997\,924\,58 \times 10^8 \text{ m s}^{-1}$ (exact)
Newton's constant	$G = 6.674\,28(67) \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}$ $= 7.426 \times 10^{-28} \text{ m/kg c}^2 = 1477.1 \text{ m}/M_{\odot}c^2$
Hubble constant at present	$H_0 = 70.4(1.5) \text{ km s}^{-1} \text{ Mpc}^{-1} = 7.61 \times 10^{-27} \text{ m}^{-1}c$ $= [13.9(0.3) \text{ Gy}]^{-1}$
Solar mass	$M_{\odot} = 1.98847(7) \times 10^{30} \text{ kg} = 1477.1 \text{ m}/Gc^2$ $= 4.925 \times 10^{-6} \text{ s}/Gc^3$
Parsec	$\text{pc} = 30.857 \times 10^{12} \text{ m} = 3.2616 \text{ ly}$
Critical density at present	$\rho_c = 9.32(0.20) \times 10^{-27} \text{ kg m}^{-3}$
Age of the Universe	$t_0 = 13.73(0.16) \text{ Gy} = 1.30 \times 10^{26} \text{ m}/c$

Introduction

On the 14th of September 2015, the Advanced LIGO detectors [7] accomplished the first direct detection of a *Gravitational Wave* (GW) signal, GW150914 [8], originated from the merging of two stellar-mass black holes. This event provided further confirmation of the predictions of Einstein's *General Relativity* (GR) [9, 10] but, most importantly, it shed light on a class of astrophysical objects and phenomena in the Universe inaccessible with the means available at the time, that is, the electromagnetic radiation and the astroparticles. Further detections followed during this first observing run (O1), ended on the 19th of January 2016 [11]. In the next years, the detectors underwent some upgrades in preparation for the subsequent O2 observing run, accomplished from December 2016 to the 25th of August 2017. The Advanced Virgo detector [12] took part to this joining the LIGO detector *network*, contributing to the first three-detector observation of a *Binary Black Hole* (BBH) merger, GW170814 [13], and, in particular, to the sky localisation to just 32 square degrees (90% *C.L.*) of the first *Binary Neutron Star* (BNS) merger, GW170817 [14]. The unprecedented precision in the localisation of the source allowed the identification of its host galaxy, and a successful *multi-messenger* observation campaign across the electromagnetic spectrum [15]; no evidence for neutrinos or high energy cosmic ray particles was reported [16].

The results of the first two observing runs (O1 and O2) of the Advanced GW detectors were jointly published in December 2019 in the *first Gravitational Wave Transient Catalogue* (GWTC-1) [1]. This reports on the source properties of ten BBH and one BNS events, including the aforementioned ones, and represented the acknowledgement of GW observations as a new important means to learn about the Universe.

In order to validate an event, it is required the consistency of the data observed by multiple detectors [2]. This allows to suppress instrumental and environmental backgrounds, which represent a *noise* source for the identification of *signals* of astrophysical origin. State of the art Data Analysis techniques have been developed to make the best use of detectors data in order to distinguish these signals from the noise, and allow accurate detections and estimations of the source properties. Most of these algorithms, like PyCBC [17] and GstLAL library [18, 19], are based on *matched filters* [20, §17], which are basically noise-weighted correlations of the detector data with relativistic models of the expected GW signals from *Compact Binary Coalescences* (CBCs). There is also the *coherent Wave Burst* (cWB) algorithm, meant for *weakly modelled* or *unmodelled* short-duration transient signals [21]. For all of them, it is of fundamental importance to have an accurate knowledge of the properties of the noise. In particular, there is an assumption, which is almost ubiquitous, that is the *stationarity* of these properties. When this is not the case, most of the available techniques become inefficient, and the analysis results are doomed to suffer of larger statistical uncertainties. Moreover, if

these *non-stationarities* are ignored, and standard analysis performed, the corresponding results will be affected by systematic errors, and lack of reliability, due to an incorrect account for the *false alarm* probability, which might even lead to false detections.

It is part of the *detector characterisation* (*Detchar*) activities to check the “quality” of detectors data, such as the validity of the assumptions of the aforementioned analysis pipelines, including stationarity [22, 23]. After a negative result of this check, the culprit data should be *vetoed* from the analysis [LVC1, 24], and, from the detector point of view, the *causes* that may have produced them investigated thoroughly in order to avoid, or at least *mitigate*, them for the future.

Moreover, in the prospect of multi-messenger astronomy, it is of crucial importance to develop analysis tools able to report on the quality of the data with the lowest possible *latency*; this is essential for obtaining fast estimations of the *sky localisation* of the candidate events and other important parameters, like the type of CBC, whether a BBH or a BNS merger, or even an NS-BH signal,¹ recalling that only from the latter we are likely to expect some kind of counterpart, in order to send an *alert* to all the other observatories that detect electromagnetic radiation, neutrinos and high energy gamma rays. Indeed, the identification of some counterpart of a GW event provides a complementary insight into the physics that produced it, and represents the only way to obtain a comprehensive knowledge of the astrophysical sources and the related emission mechanisms [25].

This Thesis work has been carried out within the *Detchar* group of the Virgo Collaboration, and in close contact with the colleagues of the LIGO Scientific Collaboration. The main aim of it has been to develop adequate analysis tools for the investigation of non-stationary noise. Their purposes have slightly changed over time, in accordance with those that were the needs of the collaboration. From the end of O2, on the 25th of August 2017, and the beginning of O3, on the 1st of April 2019, the Advanced Virgo and LIGO detectors have undertaken a phase of upgrades, and a subsequent one of commissioning [26, 27]. In this period, for the interventions on the detectors and the corresponding tuning for achieving better and better performances, the noise was in general non-stationary, and it was assumed that no signal of gravitational origin was present. Hence, the focus was not the characterisation of the detector varying behaviour but rather to assess the influence of the changes implemented in certain parts of it, as well as those unintentional of environmental and technical origin. This was done *correlating* the detector output with the various signals that monitor its subsystems, and find whether and/or how modifications to any of them may reflect on the overall detector performances. In this stage, the most of the time has been spent in resuming and improving the analysis tool named *NonNA* (Non-stationary noise Analysis Tool) developed in 2008 by Dr. Gabriele Vajente [28, LVC2], of Caltech. Significant efforts have been put in making these analyses able to handle very large amounts of Virgo data,² in the least time possible, making an efficient use of the CPUs of the Virgo computing farm machines [LVC4]; this was motivated by all those “brute force” investigations, where, in the absence of clear direction where to look at, all the available information included in the data was analysed in

¹Not currently observed as of October 2019. Other type of not yet observed sources for transient signals are Core Collapse Supernovae and sudden and localized energy release in isolated neutron stars due, for example, to “star-quakes”.

²For reference, the data produced by a gravitational wave detector amounts to about 40 MB/s, and only the gravitational channel for Virgo in O1 is about 1.3 TB [LVC3].

search of correlations. Part of the studies with NonNA, and in particular those related to the environmental noise, have been described in [29].

Related to correlation, mostly from the point of view of the investigations of the underlying relations between signals, rather than for the logic behind it, there is the concept of *causation* [30–32]. Causality between *actions*, for example on detector subsystems or of external disturbances, and *effects* on the output has been studied in relation to Virgo data, following different approaches. This work has been partly carried out with the collaboration of Dr. Luca Rei, of INFN section of Genova. A general outline of the idea of causation, as well as three different approaches to test it, including an original one developed by the author, are described in this manuscript. To the best of our knowledge, this is to-date (October 2019) the first attempt to introduce this kind of analysis, recently applied in other fields like Climatology [33] and Neuroscience [34–36], to GW detectors data. This is an advanced topic, originating from the limitations of the standard, and well understood, analyses based on correlation and the need to dig deeper into the noise couplings inside the detectors. However, due to the complexity of the concept of causation itself and of the techniques involved, no significant insights in the detector non-stationary noise structure have been achieved so far with them. The algorithms that are described in this manuscript are aimed at presenting when the different causality tests are applicable, and when not, in trying to reproduce some coupling mechanisms understood and under control.

The 3rd observing run of the Advanced GW detectors started on the 1st of April 2019, whose first part (O3a) have lasted until the 30th of September of the same year. During this time the noise was, in general, more stationary (with a notable exception that we will analyse later) and, most importantly, the data was expected to contain signals of astrophysical origin. For the reasons explained before, it is of fundamental importance, for the validation of them and the corresponding source characterisation, to make some precise checks of the quality of the data for the purposes of the Data Analysis techniques exploited by the search pipelines. These checks, in correspondence of the trigger of a “candidate event”, are called *Data Quality Reports* (DQRs); some of them will be described in this manuscript. One of the original contributions of the author has been to provide a new analysis tool to asses stationarity of the data around the time of the trigger, which improves on the previously available methods. This has been done by means of time-frequency maps of where, in time and spectrum, the non-stationarities occurred, and relaxing part of the assumptions (*e.g.*, the Gaussianity of the data) of some of the other tests present in the DQRs. Moreover, this new test is attached with quantitative statements about the level of non-stationarity (*p*-values) in each time-frequency bin. The corresponding article, containing part of the material presented in this manuscript, is currently in preparation [37].

Besides some persisting periods of non-stationary data, often related to bad weather condition or insulated system malfunctions, which have been investigated with the aforementioned analysis tools, the most important manifestation of non-stationary noise were the so-called *glitches* [38, 39]. These are short duration ($\lesssim 1$ sec) bursts of excess power; they can mimic the waveform of some signals of gravitational origin, and for this reason it is important to find the actual mechanism that produced them, as commented before. For this purpose, we have develop a framework that specialises to the case of these short-duration

non-stationarities some of the concepts and the techniques discussed before, and which constitutes an improvement on the other available tools based on correlation and coherence.

The rate of these glitches is of one each ten-to-hundred of seconds, depending on its energy relative to the “noise background”. Hence, there are chances that some gravitational signals may overlap or fall in the proximity of a glitch. For this purpose, in collaboration with Prof. Marco Cavaglia and his team at Missouri University of Science & Technology (MS&T), it has been studied a “deglitching” algorithm, based on *artificial Neural Networks* (NNs), meant to reconstruct GW data corrupted by glitches, and mitigate therefore their effects on the searches for signals of Astrophysical origin. This is one of the “trending” fields of research inside the LIGO and Virgo collaborations, which will grow of importance in the near future with the increase in rate of expected detectable GW events.

Outline of the Thesis

The contents presented in this Thesis don’t follow the chronological order described above but a logical one, from a signal analysis perspective, in increasing level of complexity of the techniques and the algorithms involved. This choice has been made to emphasise the common thread that links each analysis tool described, the range of applicability of it together with its limitations, and the improvement carried by the one described next. The outline of the various chapter is as follows:

Chapter 1: This chapter describes the framework where GW research takes place. The first section recaps some fundamentals of GR, needed to introduce GWs. Here, particular relevance has been given to the derivation and discussion of the TT -gauge, through which the functioning of a simplified Michelson GW detector has been provided in section 1.3.1. An overview of the most promising sources is presented in section 1.2, with a particular focus on the ones that have already been observed, and which are the subject of section 1.3. The Advanced GW detector network is described in section 1.3, including an overview of the new coming detectors and of the planned next generation ones. Finally, and with very much pleasure, we review in section 1.4 the confirmed GW detections achieved during the first two observing runs and published in the GWTC-1. The search pipelines that have allowed this, and the methods adopted, will be briefly described as well.

Chapter 2: This chapter represents the “playground” where the noise studies of the author have taken place. It starts with a quite detailed prelude about stochastic processes and time series analysis. This is meant to fix the notation and, most importantly, it constitutes the *language* employed to describe the rest of the manuscript. Two sections of particular relevance, which motivates the importance of the following studies, are the one about the concept of stationarity of a stochastic process and that about “the non-optimal” matched filter; sections 2.1.3 and 2.1.5. The next section 2.2 contains an overview of the detectors noise, with its classification in fundamental, technical and environmental noises. A further classification is provided in the following two sections, where we present the investigation strategies of stationary noise (section 2.3)

and where we introduce the non-stationary one, with its subdivision in “glitches” and “slow non-stationarities” (section 2.4).

Chapter 3: Most of the work and the analysis implemented by the author are described in this chapter. The first part of it, comprising sections from 3.1 to 3.4, deals with the identification of specific noise features, like non-Gaussianities (section 3.3) and non-stationarities of different kinds: slower ones (3.1), glitches (3.2) and frequency non-stationarities (3.4). These are then related to the information provided by the detector auxiliary channels by means of the analysis techniques described in the following sections. Section 3.5 contains the description of the analyses based on cross-correlation and multiple linear regression, and some of their extensions. In particular, here we present the rebuilt of the NonNA tool. In the next section 3.6 we extend the analysis of non-stationarities by means of the *wavelets*, and we introduce the method of the wavelet coherence, which extends to a time-frequency map the previous analysis of correlation and coherence. In section 3.6.3 we define a new *causality* statistics that we have called *instantaneous time delay*. The concept of causality is further elaborated in section 3.7, where we present the notion of Granger-Geweke causality and that of Converging Cross-Mappings.

Chapter 4: This constitutes a sort of “epilogue” for this Thesis, and an outlook towards the most modern techniques in signal processing. In here, we present the glitch removal project named NNetFix (a NN to Fix glitches on signals) that the author is carrying on in collaboration with the LIGO MS&T group. This is meant to mitigate the effects of glitches in proximity or superimposed to GW signals. In section 4.1 we describe the algorithm, based on a *Multi-layer Perceptron* (MLP) architecture [40], and its application to LIGO data with simulated regions of “corrupted” data, which we aim to reconstruct. As anticipated before, the importance of this work will become even more relevant during the next observing runs, where, with the increased rate of expected events, the chances that some of them will occur in the vicinity of a glitch will be not negligible. Various algorithms similar to ours are currently under development within the LIGO and Virgo collaborations. In section 4.2 the state of the art of the most mature and documented ones, updated at October 2019, is compared to that of NNetFix for different kinds of signal and glitches.

Chapter 5: Conclusions about this work, with comments and perspectives for the future.

1 Gravitational waves: theory, detectors and detections

The existence of wave-like solutions (or *radiation*), propagating information through space and time at the speed of light,¹ is one of the most fascinating predictions of Einstein's theory of General Relativity (GR) [9, 10]. Indeed, thanks to the weakness of the *gravitational interaction*, the information carried by them is not appreciably altered in their travel from the source toward us.² This property makes Gravitational Waves (GWs) a reliable *messenger* to observe distant astrophysical objects, and probe regions opaque to photons, offering a new formidable window on the Universe, beyond that allowed with the electromagnetic spectrum and the cosmic rays [25].

As a drawback, despite every non-spherically symmetric accelerating energy/mass density can produce gravitational waves (ref. to (1.36)), due to the weakness of this interaction, only highly energetic astrophysical and cosmological events can produce waves actually detectable by us. Also, for the same reason, the detection of gravitational waves has been a tremendous challenge for almost sixty years³ and constant efforts are underway to improve the sensitivity of current detectors, and reach furthest sources and regions of space-time [26].

This chapter contains a brief introduction to the subject of Gravitational Wave research. In the first section, some basics of the GR framework are outlined, which are meant to be functional to the derivation of GWs as approximate solution to the linearised Einstein's equations. Then, the most relevant sources, and their effects on test masses, are presented. In section 1.3 the basic working principles of interferometric gravitational wave detectors is discussed, together with an overview of the present GW detectors and the next generation ones. For the material presented here, and in general for the entire GW community, the two reference textbooks on the subjects are those by Maggiore [49, 50] and the one by Creighton and Anderson [3], from which we tried to consistently adopt the notation. Lastly, in section 1.4.2 the first Gravitational Wave Transient Catalogue (GWTC-1) is presented, with an overview of the Advanced LIGO and Virgo discoveries during their first two observing run (O1-O2), updated to December 2018.

¹Or, more correctly, at the *speed of space-time* [41].

²This is certainly true in first approximation but there is at least an important exception, recently studied in literature, that is, the angular power spectrum of the astrophysical gravitational wave background of GWs [42].

³The milestones for this quest have been the 1957 Chapel Hill Conference [43], later known as GR1, and its follow-up, where Pirani [44] and his mentor Bondi [45] realised the physical nature of the gravitational radiation, and the first detection achieved by the two Advanced LIGO detectors in 2015 [8]. Some nicely comprehensive overview of the history of gravitational waves can be found in these review articles [46, 47] and in the book by Kennefick [48].

1.1 Gravitational wave theory overview

In the *geometric* approach pioneered by Minkowski [51], *space-time*, the structure on which all the physical phenomena take place, can be mathematically described as a 4-dimensional *Lorentzian manifold* (\mathcal{M}, g) [52, §2.5], that is, a topological space that resembles \mathbb{R}^4 near each point, equipped with a continuous and non-degenerate *metric*, $g_{\mu\nu}$, whose signature is, for our convention, of the type $(-, +, +, +)$ at every point of \mathcal{M} . This metric determines the geometry of space-time, as well as the *geodesic* paths of test masses and light beams. Both \mathcal{M} and $g_{\mu\nu}$ are *a priori* arbitrary, from this the origin of the name “General Relativity” [53]. The revolutionary intuition of Einstein’s GR was to link gravity to the curvature of space-time, which became then a dynamical entity. This idea is encapsulated in the set of non-linear, partial differential equations for the metric tensor, known as *Einstein’s field equations*:

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = \frac{8\pi G}{c^4}T_{\mu\nu} \quad (1.1)$$

where $R_{\mu\nu}$ and $R := R_{\mu\nu}g^{\mu\nu}$ are respectively the (components of the) *Ricci tensor* and *Ricci scalar*, computed from $g_{\mu\nu}$ and its mixed partial derivatives up to the second order. $T_{\mu\nu}$ is the stress-energy tensor of matter, which obeys the equations of motion $\nabla^\mu T_{\mu\nu} = 0$. The constant factor $8\pi G/c^4$ on the right-hand side can be found by requiring that the previous equations reduce to the usual form in the Newtonian limit of slow motions and weak gravitational fields involved [52, §4.2]. It is convenient to rewrite the previous equations in the form

$$R_{\mu\nu} = \frac{8\pi G}{c^4} \left(T_{\mu\nu} - \frac{1}{2}g_{\mu\nu}T \right) \quad (1.2)$$

obtained contracting the stress-energy tensor with the (inverse) metric, $T := T_{\mu\nu}g^{\mu\nu}$, and recalling the invariant relation $g_{\mu\nu}g^{\mu\nu} = \delta_\mu^\mu = 4$.

Despite their apparent simplicity, the previous equations cannot be generally integrated. The only known “exact solutions” are obtained imposing further assumptions, as for example some symmetry properties of the metric tensor. This is the case of the Schwarzschild [54] and Kerr [55] solutions, describing axisymmetric bodies in a surrounding empty space, and the Friedmann-Lemaître-Robertson-Walker solutions, for homogeneous and isotropic space-times that must contain a energy-momentum tensor of perfect fluid type [56]. For a comprehensive account of further exact solutions refer to [57, 58].

1.1.1 Linearised gravitational wave solutions

Besides the importance of the mentioned exact solutions in the global context of GWs production and propagation, what concerns us in this section is to derive them as approximate, *linearised* solution of the general Einstein’s equations (1.1).⁴ Indeed, in the situations of interest for the detection of GWs, the curvature of space-time can be typically *considered small* at least in a certain region of space-time (meaning that quadratic terms in $g_{\mu\nu}$ inside $R_{\mu\nu}$ do not significantly contribute to the equations of motion), and we can model the solutions to (1.1)

⁴For completeness, it should be noted that, as for the previous solutions, also GWs have been found as exact ones with the additional imposition of some particular symmetry [10, 58–60].

as being the *Minkowski metric* $\eta_{\mu\nu}$ plus a *small* perturbation term $h_{\mu\nu}$:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (1.3)$$

with $\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1)$ in Cartesian (inertial) coordinates ($x^0 = ct, x^1 = x, x^2 = y, x^3 = z$), and $|h_{\mu\nu}| \ll 1$.⁵ This condition requires both the curvature to be small, and in addition constrains the coordinate system to be approximately Cartesian [64].

The decomposition (1.3) is not unique for different choices of coordinates may give in general different forms for $h_{\mu\nu}$, although the space-time solution that they describe is the same. This is an example of *gauge symmetry*, where the system does not change when the underlying coordinate system is “shifted” by an infinitesimal amount [52, §7.1]. We will exploit this fact in a moment. For this freedom, equation (1.3) identifies a family of coordinate systems that are called *nearly Minkowskian* (or *Lorentzian*), which can be transformed into each other by means of global Lorentz transformations, provided the boost doesn’t spoil the condition $|h_{\mu\nu}| \ll 1$; $h_{\mu\nu}$ behave as a tensor under them.

To preserve the linear order, indices of tensor quantities are raised, by convention, by means of the flat metric $\eta^{\mu\nu}$, while the inverse linearised metric is:

$$g^{\mu\nu} = \eta^{\mu\nu} - h^{\mu\nu} + \mathcal{O}(h^2),$$

whose form is imposed by the condition: $(\eta_{\mu\rho} + h_{\mu\rho})g^{\rho\nu} = \delta_\mu^\nu + \mathcal{O}(h^2)$.

We want to find the equations of motion obeyed by the metric perturbation $h_{\mu\nu}$, which can be obtained expanding the Einstein’s equations to first order. We start by computing the linearised expression of the *connection coefficients* (or *Christoffel symbols*):

$$\begin{aligned} \Gamma_{\mu\nu}^\rho &:= \frac{1}{2}g^{\rho\lambda}(\partial_\mu g_{\nu\lambda} + \partial_\nu g_{\mu\lambda} - \partial_\lambda g_{\mu\nu}) \\ &= \frac{1}{2}\eta^{\rho\lambda}(\partial_\mu h_{\nu\lambda} + \partial_\nu h_{\mu\lambda} - \partial_\lambda h_{\mu\nu}) + \mathcal{O}(h^2), \end{aligned} \quad (1.4)$$

where the usual assumptions of *absence of torsion*, $\Gamma_{\nu\rho}^\mu = \Gamma_{\rho\nu}^\mu$, and *metric compatibility*, $\nabla_\rho g_{\mu\nu} := \partial_\rho g_{\mu\nu} - \Gamma_{\rho\mu}^\lambda g_{\lambda\nu} - \Gamma_{\rho\nu}^\lambda g_{\lambda\mu} = 0$, have been adopted [52, §3.2].

All the terms in the previous equations are already first-order quantities in the metric perturbation; thus, the only contributions to the *Riemann tensor* will come from their derivatives, while terms of the form Γ^2 won’t contribute:

$$\begin{aligned} R_{\mu\nu\rho\sigma} &:= g_{\mu\lambda}(\partial_\rho \Gamma_{\nu\sigma}^\lambda + \Gamma_{\alpha\rho}^\lambda \Gamma_{\nu\sigma}^\alpha - (\rho \leftrightarrow \sigma)) \\ &= \eta_{\mu\lambda} \partial_\rho \Gamma_{\nu\sigma}^\lambda - \eta_{\mu\lambda} \partial_\sigma \Gamma_{\nu\rho}^\lambda + \mathcal{O}(h^2) \\ &= \frac{1}{2}(\partial_\rho \partial_\nu h_{\mu\sigma} + \partial_\rho \partial_\mu h_{\nu\sigma} - (\rho \leftrightarrow \sigma)) + \mathcal{O}(h^2). \end{aligned} \quad (1.5)$$

⁵Equation (1.3) assumes that the background space-time is *flat*. A more general expansion can be done with respect to a *background metric* $\bar{g}_{\mu\nu}$: $g_{\mu\nu} = \bar{g}_{\mu\nu} + h_{\mu\nu}$, where the condition of $|h_{\mu\nu}|$ to be small is now referred to the diagonal elements of $\bar{g}_{\mu\nu}$ [49, §1.4]. In general this is true in all those reference frames where there is a clear separation of scales, either in space (with their wavelength) or in time (with frequencies), then the distinction of the “perturbation” with respect to the background can be done unambiguously. This concept is at the base of the works of Pirani [61] and Isaacson[62, 63], and the identification of GWs as a physical effect, carrying energy and information.

The Ricci tensor comes from contracting the Riemann tensor over one index in the first and one in the second pair:⁶

$$R_{\mu\nu} := R_{\mu\rho\nu\sigma}\eta^{\rho\sigma} = \frac{1}{2}\left(\partial_\sigma\partial_\mu h^\sigma{}_\nu + \partial_\sigma\partial_\nu h^\sigma{}_\mu - \partial_\mu\partial_\nu h - \square h_{\mu\nu}\right) + \mathcal{O}(h^2). \quad (1.6)$$

In the previous expression we have introduced the trace of the metric perturbation $h := \eta^{\mu\nu}h_{\mu\nu} = h^\mu{}_\mu$, and the *d'Alembertian* differential operator of flat space-time: $\square := \partial_\mu\partial^\mu = -\frac{1}{c^2}\frac{\partial^2}{\partial t^2} + \nabla^2$. Substituting (1.6) into (1.2) yields:⁷

$$\left(\square h_{\mu\nu} - \partial_\sigma\partial_\mu h^\sigma{}_\nu - \partial_\sigma\partial_\nu h^\sigma{}_\mu + \partial_\mu\partial_\nu h\right) = -\frac{16\pi G}{c^4}\left(T_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}T\right) + \mathcal{O}(h^2) \quad (1.7)$$

These are the *linearised Einstein's field equations*. As anticipated, in this form, they cannot yield unique solution because of the remaining *gauge freedom* to make coordinate transformations (and Physics must not depend on the choice of coordinates). Hence, in order to simplify the previous expression, it is convenient to exploit this freedom choosing a coordinate system satisfying the *harmonic* or *de Donder gauge condition* [65]:⁸

$$g^{\mu\nu}\nabla_\mu\nabla_\nu x^\rho = -g^{\mu\nu}\Gamma_{\mu\nu}^\rho = 0, \quad (1.8)$$

which is equivalent, for the metric perturbation, to demanding

$$\partial_\mu h^\mu{}_\nu = \frac{1}{2}\partial_\nu h \quad (1.9)$$

up to terms of the second order in $h_{\mu\nu}$, as it is immediate to prove contracting (1.4) with $\eta^{\mu\nu}$.

To find such a frame, we can consider a local change of coordinates, characterized by the vector field $\xi^\mu(x)$,

$$x^\mu \rightarrow x'^\mu = x^\mu + \xi^\mu(x), \quad (1.10)$$

which transforms the metric perturbation $h_{\mu\nu}$ as:

$$\begin{aligned} h_{\mu\nu}(x) &\rightarrow h'_{\mu\nu}(x') := g'_{\mu\nu}(x') - \eta_{\mu\nu} \\ &= h_{\mu\nu}(x) + \left(\frac{\partial x^\rho}{\partial x'^\mu}\frac{\partial x^\sigma}{\partial x'^\nu} - \delta_\mu^\rho\delta_\nu^\sigma\right)g_{\rho\sigma}(x) \\ &= h_{\mu\nu}(x) - (\partial_\mu\xi_\nu + \partial_\nu\xi_\mu) + \mathcal{O}((\partial\xi)^2), \end{aligned} \quad (1.11)$$

where in the second line we have substituted to the flat space-time metric $\eta_{\mu\nu}$ its expression in terms of $h_{\mu\nu}$ and $g_{\mu\nu}$. In order for $h'_{\mu\nu}$ to remain small, the partial derivatives $\partial_\mu\xi_\nu$ should be small at least of the same order of $h_{\mu\nu}$. Local changes of variables with such a property

⁶Here we adopt the Misner, Thorne and Wheeler convention [4] for the choice of which index to contract, as explained in the Notation.

⁷Notice that the stress-energy tensor already accounts for effects of first order in $h_{\mu\nu}$; for this reason it has been substituted the flat metric $\eta_{\mu\nu}$, in front of its trace, to the perturbed one $g_{\mu\nu}$. In any case, the same result can be obtained computing the linearised left-hand side of (1.1), as in [64, (2.60)].

⁸From (1.8), in this gauge each of the coordinate functions x^μ satisfies the curved space-time *d'Alembert's equation* (thus the alternative name *d'Alembert gauge*). This is a generalisation to space-time of the harmonic coordinates satisfying Laplace's equation in Riemannian geometry (*Laplace's gauge*). As regards the "potentials" $h_{\mu\nu}$, and in particular the trace-reversed ones $\bar{h}_{\mu\nu}$, the condition (1.21b) is also analogous to the *Lorentz gauge* for the electromagnetic potential vector A_μ : $\partial^\mu A_\mu = 0$.

are called *slowly varying diffeomorphisms*. From the previous equation, it is immediate to verify that the linearised Riemann tensor (1.5) is *invariant* under the effect of this gauge transformation, and also that:

$$\begin{aligned} \partial_\mu h^\mu{}_\nu - \frac{1}{2}\partial_\nu h &\rightarrow \partial'_\mu h'^\mu{}_\nu - \frac{1}{2}\partial'_\nu h' = \partial_\mu h^\mu{}_\nu - \frac{1}{2}\partial_\nu h - \partial^\mu (\partial_\nu \zeta_\mu + \partial_\mu \zeta_\nu) + \partial_\mu \partial_\nu \zeta^\mu \\ &= \partial_\mu h^\mu{}_\nu - \frac{1}{2}\partial_\nu h - \square \zeta_\nu \end{aligned} \quad (1.12)$$

Since the d'Alembertian operator is invertible, we can always solve the inhomogeneous wave equation $\square \zeta_\nu = \partial_\mu h^\mu{}_\nu - \frac{1}{2}\partial_\nu h$, without spoiling the slow diffeomorphism conditions, and choose the gauge parameter ζ^μ such that the right-hand side of the previous equation vanishes:

$$\partial'_\mu h'^\mu{}_\nu = \frac{1}{2}\partial'_\nu h' \quad (1.13)$$

satisfying the harmonic gauge condition (1.9). If we compute the derivative ∂'_ρ of the previous equation and we symmetrize with respect to ν and ρ , substituting the result (prime symbols are omitted)

$$\partial_\rho \partial_\mu h^\mu{}_\nu + \partial_\mu \partial_\nu h^\mu{}_\rho - \partial_\rho \partial_\nu h = 0$$

in equation (1.7), after renaming in a suitable way all the dummy indices, we have:

$$\square h_{\mu\nu} = -\frac{16\pi G}{c^4} \left(T_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu} T \right) + \mathcal{O}(h^2) \quad (1.14)$$

This choice of the gauge has reduced the linearised Einstein's equations to simple wave equations with source term given by the *trace reversed* stress-energy tensor: $T_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu} T$. Instead of having this form on the right-hand side, it is often customary to redefine it in terms of the *trace reversed metric perturbation*:

$$\bar{h}_{\mu\nu} := h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu} h \quad (1.15)$$

which gives to the linearised Einstein's field equations the particularly simple form:

$$\square \bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} T_{\mu\nu} + \mathcal{O}(h^2) \quad (1.16)$$

These equations are the starting point for computing GWs generation within the linearised theory. We will come back to this aspect in the next section.

Outside the source $T_{\mu\nu} = 0$, and (1.16) become homogeneous wave equations for each component:

$$\square h_{\mu\nu} \equiv \square \bar{h}_{\mu\nu} = 0. \quad (1.17)$$

The flat space-time d'Alembertian operator implies metric perturbations that travel at the speed of light, to which we attribute the meaning of Gravitational Waves. In this case, the most general vacuum solutions to $\square \bar{h}_{\mu\nu} = 0$ is given by the infinite summation of *monochromatic plane waves*

$$h_{\mu\nu}(x) = \int \frac{d^3\mathbf{k}}{(2\pi)^3} h_{\mu\nu}(\mathbf{k}) e^{ik^\rho x_\rho} + c.c. \quad (1.18)$$

where $k^\mu = (2\pi f/c, 2\pi f \hat{\Omega}/c)$ is the null, *wave four-vector*, and $\hat{\Omega}$ is the unit (three-dimensional) vector that identifies the direction of propagation of the wave. We can rewrite the differential of the wave vector as:

$$d^3\mathbf{k} = (2\pi/c)^3 f^2 df d^2\hat{\Omega},$$

where $d^2\hat{\Omega} := d \cos \theta d\phi$, and then rewrite (1.18) as:

$$h_{\mu\nu}(t, \mathbf{x}) = \frac{1}{c^3} \int_0^\infty df \int_{S^2} d^2\hat{\Omega} f^2 \tilde{h}_{\mu\nu}(f, \hat{\Omega}) e^{-2\pi i f(t - \hat{\Omega} \cdot \mathbf{x}/c)} + c.c. \quad (1.19)$$

The factors $1/c^3$ and f^2 can be reabsorbed in the definition of the amplitudes $\tilde{h}_{\mu\nu}(f, \hat{\Omega})$. Notice also that only “physical frequencies”, *i.e.* $f > 0$, enter in the previous equation; we can rewrite it in a more compact form, extending the definition of $\tilde{h}_{\mu\nu}(f, \hat{\Omega})$ to negative frequencies by defining:

$$\tilde{h}_{\mu\nu}(-f, \hat{\Omega}) := \tilde{h}_{\mu\nu}^*(f, \hat{\Omega}),$$

and then rewriting (1.19) as:

$$h_{\mu\nu}(t, \mathbf{x}) = \int_{-\infty}^{+\infty} df \int_{S^2} d^2\hat{\Omega} \tilde{h}_{\mu\nu}(f, \hat{\Omega}) e^{-2\pi i f(t - \hat{\Omega} \cdot \mathbf{x}/c)}. \quad (1.20)$$

1.1.2 GWs in TT -gauge

The harmonic gauge condition (1.9) doesn't completely fix the freedom we have to modify the components of the metric perturbation, and possibly to “gauge them away”. Indeed, this symmetric tensor $h_{\mu\nu}$ has 10 independent components, which are later reduced to 6 imposing the four constrain of the harmonic gauge (1.9). It is important then to convince ourselves about how many of them have a real physical meaning, and represent actual propagating degrees of freedom; refer to the comments in note 3.

Let us work with the trace reversed metric perturbation, and deduce the number of effective propagating degrees of freedom in vacuum ($T_{\mu\nu} = 0$). We can rewrite the gauge condition (1.9) and its transformation (1.12) for $\bar{h}_{\mu\nu}$ as:

$$\bar{h}_{\mu\nu} \rightarrow \bar{h}'_{\mu\nu} := \bar{h}_{\mu\nu} - (\partial_\mu \xi_\nu + \partial_\nu \xi_\mu) + \eta_{\mu\nu} \partial^\rho \xi_\rho \quad (1.21a)$$

$$\partial_\mu \bar{h}'^\mu{}_\nu = 0 \rightarrow \partial_\mu \bar{h}'^\mu{}_\nu = \partial_\mu \bar{h}^\mu{}_\nu - \square \xi_\nu = 0. \quad (1.21b)$$

Then, we are free to perform any further transformation satisfying $\square \xi_\nu = 0$ without spoiling the previous conditions.

For definiteness, we consider a plane wave solution of the form $\bar{h}_{\mu\nu}(t, x) = e_{\mu\nu} e^{ik_\rho x^\rho}$, where the *polarization tensor* $e_{\mu\nu}$ is a set of constants, with respect to x (but not with k), forming a symmetric tensor (10 components), and k_ρ is the null wave 4-vector introduced in the previous section. Putting it in the harmonic gauge condition we have

$$k^\mu e_{\mu\nu} = 0 \quad (1.22)$$

meaning that the polarization tensor must be *transverse* with respect to the propagation direction identified by \mathbf{k} (4 constraints).

The gauge vector field can be chosen of the form $\xi_\mu(x) = \chi_\mu e^{ik_\rho x^\rho}$, with χ_μ four constants (for fixed k_μ), which automatically satisfies the condition $\square \xi_\nu = 0$; the most general $\xi_\mu(x)$ is the infinite sum of such solution, similarly to (1.20). Substituting it into (1.21a), we have the transformation relation:

$$e'_{\mu\nu} = e_{\mu\nu} - ik_\mu \chi_\nu - ik_\nu \chi_\mu + i\eta_{\mu\nu} k^\rho \chi_\rho. \quad (1.23)$$

The residual gauge freedom, embodied in the four arbitrary constants χ_μ , can then be exploited choosing $e'_{\mu\nu}$ to be *traceless*:

$$e'^\mu{}_\mu = e^\mu{}_\mu - 2ik^\mu \chi_\mu = 0 \quad (1.24a)$$

that is, the constants in the gauge field chosen such that:

$$2ik^\mu \chi_\mu = e^\mu{}_\mu.$$

This constitutes one further constraint to the actual degrees of freedom of $e'_{\mu\nu}$, reducing the number of independent constants of χ_μ by one. In this *transverse-traceless gauge* (*TT*) the trace reversed metric coincides with the usual one: $\bar{h}_{\mu\nu} \equiv h_{\mu\nu}$.

Three more conditions can be then imposed. For example all the $e'_{0\mu}$ can be set to zero:

$$e'_{0\mu} = e_{0\mu} - ik_\mu \chi_0 - ik_0 \chi_\mu + i\eta_{0\mu} k^\rho \chi_\rho = 0. \quad (1.24b)$$

These correspond to three conditions, and not four, because $e'_{\mu\nu}$ must be transverse, and the previous four equations are constrained to fulfil

$$\begin{aligned} k^\mu e'_{0\mu} &= k^\mu e_{0\mu} - ik^2 \chi_0 - ik_0 k^\mu \chi_\mu + ik_0 k^\mu \chi_\mu \\ &= k^\mu e_{0\mu} = 0 \end{aligned}$$

meaning that $e'_{0\mu} = 0$ stands for just three independent conditions. An explicit solution for the two (1.24) is given by [64]:

$$\chi_\mu = \frac{3e_{\nu\rho} l^\nu l^\rho}{8ik_0^4} k_\mu + \frac{\eta^{\nu\rho} e_{\nu\rho}}{4ik_0^4} l_\mu + \frac{1}{2ik_0^2} e_{\mu\nu} l^\nu \quad (1.25)$$

with $l^\mu := (k_0, -\mathbf{k})$.

The symmetric polarisation tensor then has *two* independent elements, corresponding to two physical propagating degrees of freedom, which can't be cancelled out by any gauge choice. For example, a plane wave travelling in the $+z$ -direction has the *TT*-gauge expression:

$$h_{\mu\nu}^{TT}(t, z) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}_{\mu\nu} e^{-2\pi i f(t-z/c)}. \quad (1.26)$$

The two polarisation states can be taken to be

$$e_{\mu\nu}^+ := \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}_{\mu\nu}, \quad \text{and} \quad e_{\mu\nu}^\times := \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}_{\mu\nu} \quad (1.27)$$

respectively named “plus” and “cross” modes, and with amplitudes h_+ and h_\times :

$$h_{\mu\nu}^{TT}(t, z) = \sum_{A=+, \times} e_{\mu\nu}^A h_A e^{-2\pi i f(t-z/c)}. \quad (1.28)$$

The origin of these names for the two polarisation modes will be cleared in section 1.1.4. The non-tensorial part of the previous equation, $h_A(t, z) := h_A \exp(-2\pi f(t-z/c))$ is called the GW *strain*, and also the reason for that will become clear in the next section.

Although the previous procedure of passing to the TT -gauge has been derived for monochromatic waves, in fact, this can be extended to all the *radiative solutions*. Indeed, these can be generally expanded as *linear* superposition of monochromatic waves, as expressed by (1.20).

1.1.3 Emission of gravitational waves

In linearised theory, the emission of GWs is described by equation (1.16), which can be solved, analogously to the case of Electrodynamics, with the *method of Green's function* [66, §12], obtaining a *retarded potential* solution

$$\begin{aligned} \bar{h}_{\mu\nu}(t, \mathbf{x}) &= -\frac{16\pi G}{c^4} \int d^4x' \left(-\frac{1}{4\pi|\mathbf{x}-\mathbf{x}'|} \delta(ct - |\mathbf{x}-\mathbf{x}'| - ct') \right) T_{\mu\nu}(x') \\ &= \frac{4G}{c^4} \int d^3\mathbf{x}' \frac{1}{|\mathbf{x}-\mathbf{x}'|} T_{\mu\nu} \left(t - \frac{|\mathbf{x}-\mathbf{x}'|}{c}, \mathbf{x}' \right) \end{aligned} \quad (1.29)$$

up to second order terms in h .⁹ In the first line, the term in parenthesis is the Green's function for the d'Alembertian operator in four dimensions, and the term $t - |\mathbf{x}-\mathbf{x}'|/c$ represents the *retarded time* in \mathbf{x} for a wave emitted by a source in \mathbf{x}' . The previous integrals extend over the past light-cone of the event (t, \mathbf{x}) .

We want to connect the results of these equations with the vacuum solutions obtained in (1.20). For the purposes of this section, we are particularly interested in the *far-field* limit of them, where the dynamics of the source is given, and we examine the GWs production at far distance from it. In general, GWs carry energy and momentum, which affect the motion of the source through a “back reaction” effect; then, a *near-field* solution is also needed in order to compute the equations of motion that describe the source dynamics. We will ignore this aspect in this treatment, referring to [3, 49] for the results.

In the far-field region the distance from the source, r , is much greater than the GW wavelength $\lambda = c/f$, which in turn is much greater than the characteristic size of the source R .

⁹These second order terms correspond to the GWs stress-energy tensor [49, §1.4].

This means that we can approximate:

$$|\mathbf{x} - \mathbf{x}'| \simeq r$$

meaning that there is no relative delay between waves emitted from two different sides of the source, and

$$\bar{h}_{\mu\nu}(t, \mathbf{x}) \simeq \frac{4G}{c^4 r} \int d^3\mathbf{x}' T_{\mu\nu}(t - r/c, \mathbf{x}'). \quad (1.30)$$

If we work in the TT -gauge of section 1.1.2, the only interesting components of the previous equation are the spatial ones:

$$\bar{h}_{ij}(t, \mathbf{x}) \simeq \frac{4G}{c^4 r} \int d^3\mathbf{x}' T_{ij}(t - r/c, \mathbf{x}').$$

It is convenient however to rewrite the components T^{ij} in terms of the T^{00} term only. From the time component of the matter equation of motion in the linearised theory, $\nabla_\mu T^{\mu 0} = \partial_\mu T^{\mu 0} + \mathcal{O}(h^2) = 0$, which is actually a conservation law to the second order in h , by differentiating with respect to time and multiplying by $x^i x^j$, we can obtain the *tensor virial theorem*:

$$0 = x^i x^j \frac{1}{c} \frac{\partial}{\partial t} \left(-\frac{1}{c} \frac{\partial T^{00}}{\partial t} + \frac{\partial T^{0k}}{\partial x^k} \right) = x^i x^j \left(-\frac{1}{c^2} \frac{\partial^2 T^{00}}{\partial t^2} + \frac{\partial^2 T^{k\ell}}{\partial x^k \partial x^\ell} \right)$$

whose integrated form is:

$$\int x^i x^j \frac{1}{c^2} \frac{\partial^2 T^{00}}{\partial t^2} d^3\mathbf{x} = \int x^i x^j \frac{\partial^2 T^{k\ell}}{\partial x^k \partial x^\ell} d^3\mathbf{x} = 2 \int T^{ij} d^3\mathbf{x}. \quad (1.31)$$

The integral in the last step has been evaluated by parts, where the surface terms obtained from the spatial derivatives vanish for a “bounded system” (*Gauss’s theorem*). The spatial part of the metric perturbation can then be written as:

$$\bar{h}_{ij}(t, \mathbf{x}) \simeq \frac{2G}{c^4 r} \frac{\partial^2}{\partial t^2} \int x'_i x'_j T^{00}(t - r/c, \mathbf{x}') d^3\mathbf{x}'. \quad (1.32)$$

The integral on the right-hand side is called the source *quadrupole tensor*:

$$I_{ij}(t - r/c) := \frac{1}{c^2} \int x'_i x'_j T^{00}(t - r/c, \mathbf{x}') d^3\mathbf{x}' \quad (1.33)$$

then:

$$\bar{h}_{ij}(t, \mathbf{x}) \simeq \frac{2G}{c^4 r} \ddot{I}_{ij}(t - r/c). \quad (1.34)$$

This is the *Einstein’s quadrupole formula* for the GW field. If we consider only linear terms in h , the stress-energy tensor includes only the matter contribution and $T^{00} = c^2 \rho'$, with ρ' the matter density of the source as measured by an observer at rest relative to the gauge coordinate frame (1.10); this differs from the *proper density* by corrections of order c^{-2} .

To project the right-hand side to TT -gauge, for radially propagating GWs, we need the *transverse projector operator*:

$$P_{ij} = \delta_{ij} - \hat{\Omega}_i \hat{\Omega}_j, \quad \text{where } \hat{\Omega}_i := x_i / r$$

which is used to obtain the transverse, traceless quadrupole tensor:

$$I_{ij}^{TT} := P_{ik} I^{k\ell} P_{\ell j} - \frac{1}{2} P_{ij} P_{k\ell} I^{k\ell}. \quad (1.35)$$

It is immediate to check that the previous quantity vanishes once projected onto \hat{r}_i , and that it has zero trace.

In conclusion, the far-field TT -gauge solution can be written in the compact form:

$$\bar{h}_{ij}^{TT}(t, \mathbf{x}) \simeq \frac{2G}{c^4 r} I_{ij}^{TT}(t - r/c). \quad (1.36)$$

In the presence of matter, the metric perturbation have more than just two degrees of freedom, but it can be shown that only the TT -part propagates as a wave; there is also a “Newtonian”, non-radiative part which instead obeys a “static” Poisson equation [64]. This justifies the application of the projector operator P_{ij} to the quadrupole moment in (1.35).

1.1.4 Effects of GWs on test masses

In this section we will derive how “test particles” respond to the passage of a GW, with particular attention to what concerns the working principle of interferometric detectors. Let’s start with a single particle, whose geodesic (*free fall*) trajectory is described in GR by

$$\frac{d^2 x^\mu}{d\tau^2} + \Gamma_{\nu\rho}^\mu \frac{dx^\nu}{d\tau} \frac{dx^\rho}{d\tau} = 0 \quad (1.37)$$

for any space-time metric $g_{\mu\nu}$ and coordinate system $\{x^\mu\}$. Here τ is the particle’s proper time, as measured by a clock attached to it all along the whole geodesic. From the latter, we can obtain a differential equation for the spatial coordinates x^i , differentiated with respect to the coordinate time $x^0 = ct$ instead of τ [64]:

$$\frac{1}{c^2} \frac{d^2 x^\mu}{dt^2} = - \left(\Gamma_{00}^i + 2\Gamma_{0j}^i \frac{v^j}{c} + \Gamma_{jk}^i \frac{v^j}{c} \frac{v^k}{c} \right) + \frac{v^i}{c} \left(\Gamma_{00}^0 + 2\Gamma_{0j}^0 \frac{v^j}{c} + \Gamma_{jk}^0 \frac{v^j}{c} \frac{v^k}{c} \right) \quad (1.38)$$

where $v^i := dx^i/dt$ is the i -th component of the coordinate velocity. If we now specialise to the linearised theory, described in TT -gauge, and to non-relativistic motions, $v^i \ll c$, whose a special case is that of a particle at rest in that specific frame, $v^i = 0$, we can neglect all the velocity dependent terms in the previous equation, and be left with:

$$\begin{aligned} \frac{d^2 x^i}{dt^2} &= -\Gamma_{00}^i + \mathcal{O}(v/c) \\ &\stackrel{(1.4)}{=} \frac{1}{2} \left(2\frac{1}{c} \partial_t h_{i0} - \partial_i h_{00} \right) + \mathcal{O}(v/c, h^2) \stackrel{(1.24)}{=} 0. \end{aligned} \quad (1.39)$$

This result implies that, in TT -gauge, coordinates of freely falling, non-relativist particles are left unchanged by the passage of a GW, at the considered order. This is a property already implicit in the equivalence principle.

Let us consider instead two neighbouring, test particles at rest ($v^i = 0$ for both). Their coordinates will be left unchanged as well, so lets put one at the origin of the coordinate system and the other at a coordinate distance L_x along the x -axis. We want to study the




Figure 1.1: Polarisation modes of GWs in Einstein’s GR, represented by the variation of the proper distance from the centre they induce on a circular array of test masses set at rest with respect to the gauge, and perpendicular to the waves direction of propagation (\odot symbol). The effect in time is represented by different shades of grey, which justifies the names given to a purely plus-polarized (left) and a purely cross-polarized GW (right). The lightest contour represents an argument of the right-hand side of equations (1.42) and (1.43) equals to $2n\pi$, the intermediate, circular contour to arguments $\pi/2 + n\pi$, and the darker one to $\pi + 2n\pi$, with n any integer.

propagation of a light beam between the two masses and derive the effects of the wave on it; for concreteness we will focus now on a $+$ -polarised wave travelling to the z -direction: $h_+^{TT}(t, z = 0)$. The *null* space-time interval of the light in the TT -frame will be (ref. to (1.27)):

$$0 = ds^2 = -c^2 dt^2 + (1 + h_+(t)) dx^2$$

from which we can extract an equation for dx :

$$dx = \pm \left(1 - \frac{1}{2} h_+(t) \right) c dt + \mathcal{O}(h^2).$$

Plus and minus signs are referred to a light ray travelling in the positive and negative x direction respectively. The latter equation can be integrated between the *fixed* coordinates of the two test masses, and the times t_0 and t_1 at which the light passes the two:

$$\begin{aligned} L_x &= \int_0^{L_x} dx = \int_{t_0}^{t_1} \left(1 - \frac{1}{2} h_+(t') \right) c dt' + \mathcal{O}(h^2) \\ &= c(t_1 - t_0) - \frac{c}{2} \int_{t_0}^{t_1} h_+(t') dt' + \mathcal{O}(h^2). \end{aligned} \quad (1.40)$$

In the last line, the second term makes evident the effect of the incoming GW. For simplicity, if we assume the amplitude of the wave not to change significantly in a time $t_1 - t_0$ (long wavelength approximation), the previous equation reduces to:

$$L_x \approx c(t_1 - t_0) \left(1 - \frac{1}{2} h_+(t) \right).$$

The term $c(t_1 - t_0) := L$ corresponds to the *proper distance* travelled by the light beam, which is different from the (constant) “coordinate distance” L_x . We can rewrite indeed:

$$\frac{c(t_1 - t_0) - L_x}{c(t_1 - t_0)} := \frac{\Delta L_x}{L} \approx \frac{1}{2} h_+(t) \quad (1.41)$$

which tells us that the relative separation, $\Delta L_x/L$, between the two test particles oscillates as the GW strain $h_+(t)$. For a monochromatic wave, as in (1.26):

$$\frac{\Delta L_x}{L} \approx \frac{1}{2} h_+ \cdot \cos(2\pi f t). \quad (1.42)$$

If the second mass was positioned on the y -axis, at coordinate L_y , a similar computation would have led, according to (1.27), to:

$$\frac{\Delta L_y}{L} \approx -\frac{1}{2} h_+(t) \stackrel{\text{monoc.}}{=} -\frac{1}{2} h_+ \cdot \cos(2\pi f t) \quad (1.43)$$

which differ from (1.42) by a global “−” sign. It is clear then how GWs act *tidally* on test masses, contracting the proper distance along one axis and stretching along the other. This interpretation will result useful in the next section, when we will discuss the working principle of interferometric GW detectors and the orders of magnitudes of the effect on them from various possible sources.

Repeating the previous computation for a circular array of test masses in the (x, y) plane, and for the other \times -polarisation, it can be shown that the resulting pattern of the effects of an incoming (“ \odot ” symbol, representing the z -axis) GW is that represented in figure 1.1, where different shades of grey correspond to different values of time. This motivates the origin of the names “plus and cross” given to the two polarisation modes in TT -gauge.

Although it was derived specifically in TT -gauge, the previous result is gauge independent (to the linear order in h),¹⁰ and in section 1.3 we will describe a detector configuration able to measure a gauge-independent observable related to the proper distance variation induced by GWs. Usually, many authors prove the previous statement directly computing the effect of GWs also in the so called “detector frame”, which is the natural one used in the laboratory [3, 5, 49, 64], while others derive it in a gauge-independent way, specialising only later to a specific coordinate frame [67–69].

1.2 Sources of gravitational waves

From equation (1.36) it is clear that GWs are emitted every time there is an accelerating (second time derivative), non-axisymmetric ($\ddot{I} \neq 0$), energy/mass density, and that this emission is weighted by the *tiny* factor G/c^4 . Let’s start with an order of magnitude estimate on the characteristic strain amplitude we may expect from a typical source, having in mind a coalescing binary but making some considerations easily generalisable to other systems. First of all, the (non constant part of the) quadrupole moment (1.33) of a system is approximately equal to the mass M of the part of the system that “moves”, times the square of the characteristic size R of it. Then, calling v the mean velocity of the moving part, and T the time scale for a mass to move from one side of the system to the other:

$$\frac{d^2 I}{dt^2} \approx \frac{M R^2}{T^2} \approx M v^2$$

¹⁰This is a consequence of the invariance of the Riemann tensor in the linearised theory.

and the quadruple formula can be approximated as:

$$h \approx \frac{2GM}{c^2 r} \left(\frac{v}{c}\right)^2. \quad (1.44)$$

The term $2GM/c^2$ is the *Schwartzschild radius* of the mass M , which for a celestial body of the size of the Sun corresponds to ~ 3 km. A reasonable value for the distance r of the source from us is of the order of some *Mega-parsecs* (Mpc) $\sim 3 \times 10^{19}$ km. Then, for a binary system, characterised by a mass corresponding to that of the Sun, distant 1 Mpc and moving at 10% the speed of light in the last stages of their coalescence,¹¹ the characteristic GW amplitude is about $h \sim 10^{-21}$. Looking at (1.41), the corresponding effect ΔL on a detector is of only 10^{-21} times its characteristic length L . For example, for a kilometre-scale, interferometric detector this would correspond to 10^{-18} m, that is one thousandth the size of the atomic nucleus. In the section 1.3 we will briefly describe how we can achieve such sensitivities with laser interferometric detectors.

As regards the possible sources, the previous estimate implies that GWs require very large masses and relativistic motions. Clearly the best candidates must be compact and dense, and therefore the archetypal ones are the coalescences of pairs of compact stars (*e.g.* neutron stars or black holes) in a tight binary system.

Similarly to the previous estimate on the strain amplitude, we can use *Kepler's third law* to make some previsions about the expected characteristic frequency of the waves emitted by various systems. With the same notation as above, we have:

$$f \sim \frac{1}{T} \sim \sqrt{\frac{GM}{R^3}} \sim 1/\sqrt{\rho} \quad (1.45)$$

with ρ a characteristic mean-density of the system. Let's consider for example a coalescing binary system composed of two neutron stars in a narrow orbit around each other. Their typical masses are $1.4 M_\odot$, while we can assume a value of ~ 100 km for the radius of their orbit. This results in a characteristic frequency for the emitted GWs of ~ 100 Hz.¹² At merger, the resulting object will have a total mass of $\sim 2.8 M_\odot$ and an estimated radius of ~ 10 km [70]. If we assume that a fraction of $\sim 0.1\%$ of its mass contribute to its dynamics, we obtain a frequency at merger of ~ 1 kHz. Analogously, a black-hole with a mass of $\sim 100 M_\odot$, and a Schwarzschild radius of ~ 300 km, will have an oscillation frequency $f \sim 100$ Hz. These are the main targets of ground-based, interferometric detectors, and, at December 2018, ten BBH and one BNS systems have been detected during the first two observing runs of the Advanced detectors. Further details on this will be presented in section 1.4.¹³

¹¹This may seem an unreasonably high value for the speed of a macroscopic system like a star, but it is indeed the actual one estimated for the coalescence of a BBH system. Refer to figure 1.10 and reference [8].

¹²Although the proposed numbers seem a little bit the result of a "cherry picking", these are indeed consistent with what actually measured for the BNS coalescence GW170817 [14].

¹³In the Introduction to this work, and also to this chapter, we have repeatedly compared GWs with electromagnetic ones. From the previous estimates we can notice an important difference between the two. Electromagnetic waves are usually generated at microscopic level by moving charges; therefore they have wavelengths smaller than the size of the emitting system, and so can be used to form an image of it. By contrast GWs are generated by the bulk dynamics of the source and corresponding wavelength, as described in the previous order of magnitude estimates, is typically comparable to or larger than the size of the radiating source. As a consequence, GWs cannot be used to resolve the image of the system, and only indirectly its structure [64].

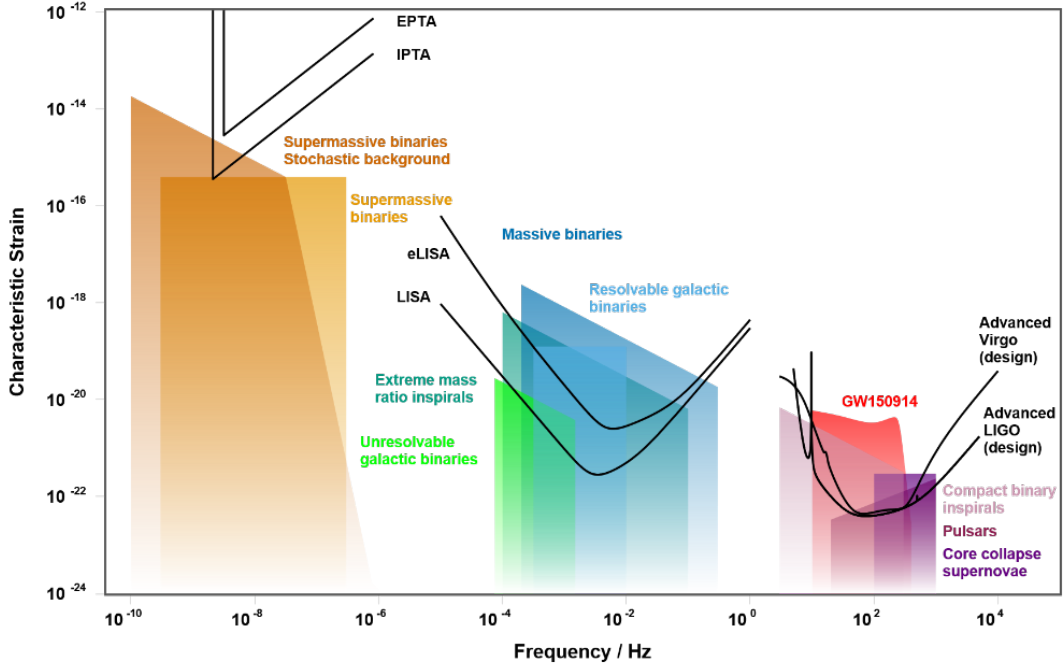


Figure 1.2: Strain amplitude and frequency ranges expected for various astronomical sources of GWs (colored boxes), and sensitivities of current and proposed detectors (black lines). The figure was generated using gwplotter. Ref. to [73] for the description and for the assumptions used in creating it.

Repeating the previous estimate for other kinds of sources, like massive ($10^4 \div 10^7 M_\odot$) and super-massive ($\gtrsim 10^7 M_\odot$) black-hole binary systems, allows to define a whole spectrum of possible sources and their characteristic strains. The results of more accurate calculations are reported in figure 1.2, together with the sensitivities of some of the current and proposed GW detectors.

Referring to that figure, a first way of classifying the possible sources is by the frequency band in which they produce GWs, or, equivalently, by the sensitivity range of the detectors able to detect them. Then, from right to left, in decreasing frequencies, we can distinguish an high frequency region, between 1 Hz and 10 kHz, which is the typical one of ground-based GW detectors and the emissions of stellar-mass objects. The low frequency band, with frequencies between 1 mHz and 1 Hz, is the characteristic one of space missions, like LISA [71], and GWs emitted by more massive stars. Pulsar timing experiments [72] are sensitive to gravitational waves in the very low frequency band, between 1 nHz and 1 mHz, the characteristic frequencies of emission of super-massive black-hole binaries.

An exception to this classification scheme, and also to the way the previous orders of magnitude estimates are thought, is the *cosmological stochastic background* of GWs, not shown in figure 1.2. There are strong theoretical reason to believe that a background of this kind, similar to the *cosmological microwave background radiation* (CMB), has been produced in the early Universe, and several possible mechanisms have been proposed [74–76].¹⁴ In general,

¹⁴For the Astrophysical sources of stochastic GW background, given by the incoherent superposition of a large number of the aforementioned astrophysical sources, refer to [77, 78].

all of them describe a GW spectrum spanning several orders of magnitude of frequencies. Further details about it will be discussed in section 1.2.4.

In the rest of this section, we will provide some further details about the most important sources of GWs for ground-based interferometric detectors, with particular attention to those already detected. A comprehensive and up to dated review on the subject can be found in the second volume of the monograph by Maggiore [50]. The physics of compact sources of GWs is comprehensively studied in [79]. A review of the science that can be learned from them is presented in the Living Review article [80].

Besides the previous classification, from a Data Analysis point of view, LIGO and Virgo scientists have defined four categories of signals, based on the method used to study them. These categories are: Continuous Gravitational Waves, Compact Binary Coalescences (CBCs), Stochastic Gravitational Waves, and Burst Gravitational Waves. Each one generates a unique “fingerprint” or characteristic signature in Advanced detectors data. We will proceed to review them.

1.2.1 Compact binary coalescence

From the previous estimates we have seen that any binary system loses energy in the form of gravitational radiation. However, among all of them, the most important sources of GWs are binaries made of *compact objects*, which are the most likely ones that can reach separations small enough and relative velocities large enough to produce detectable waves. These are typically systems formed by neutron stars (BNS), black holes (BBH), or mixed systems of neutron stars and black holes (NS-BH). These three classes are represented, as functions of the various masses of the objects involved, in figure 1.3. There is also a fourth class of objects, belonging to the so called “mass gap”, with masses above those of the observed neutron stars and below that of black holes, for which we don’t have evidence yet.

As a consequence of their emissions, the distance between the components decreases, producing an *inspiral* motion. The resulting GW signal will be a sinusoid whose amplitude and frequency increase with time, as the system gets tighter, and it is called a *chirp*. This phase can be studied, in first approximation, by means of adiabatic corrections to the Newtonian dynamics. An analytic approach to include further relativistic effects is, for example, that provided by the *post-Newtonian* (PN) approximation, in which corrections to the Newtonian emission are provided in increasing powers of the velocities of the objects, relative to that of light: v/c . This approach has been reviewed for example in [81, 82].

Next, when the separation of the two components becomes comparable with their size, the two objects “collide”, producing a *merger* phase characterised by a strong emission of GWs. The evolution of the orbit is no longer adiabatic, and this phase is typically studied by means of numerical simulations [83].

Finally, there is the *ringdown* phase, where the resulting object, typically a Kerr black hole, relaxes toward a stable state, emitting waves in the form of damped sinusoids. Their frequencies are those typical of the quasi-normal modes of the final black hole [84].¹⁵ These

¹⁵It should also be mention the *Effective One Body* (EOB) formalism, which consists of mapping, through the use of gauge-independent functions, the real two-body problem (two spinning masses orbiting around each other) onto an “effective one-body” problem: one spinless mass moving in some “effective” background metric, which is a

three phases of a CBC evolution are represented in the top panel of figure 1.10 for the first confirmed event of such type, GW150914 [8].

In the adiabatic Newtonian approximation for the circular orbit inspiral, the characteristic amplitude of the two GW polarizations is [5, §2]:

$$h(t) \simeq 2.6 \times 10^{-23} \left(\frac{\mathcal{M}}{M_\odot} \right)^{5/3} \left(\frac{f(t)}{100 \text{ Hz}} \right)^{2/3} \left(\frac{r}{100 \text{ Mpc}} \right)^{-1}, \quad (1.46)$$

with the *chirp mass* of the system $\mathcal{M} := \mu^{3/5} M^{2/5}$, which is the fundamental quantity that describes the GW emission at this level of approximation, and which is related to the total mass $M := m_1 + m_2$ and the reduced mass $\mu := m_1 m_2 / M$. The frequency of the waves f is twice that of the orbit, and r is the distance of the binary from us. The two polarisations can be expressed as:

$$h_+(t, \mathbf{x}) = -h(t - r/c) \frac{1 + \cos^2 \iota}{2} \cos \left(2\pi \int_{t_0}^{t-r/c} f(t') dt' + 2\phi_0 \right), \quad (1.47a)$$

$$h_\times(t, \mathbf{x}) = -h(t - r/c) \cos \iota \sin \left(2\pi \int_{t_0}^{t-r/c} f(t') dt' + 2\phi_0 \right), \quad (1.47b)$$

with ι the angle between the orbital angular momentum vector of the binary and the line of sight, and ϕ_0 the phase at the initial time t_0 when the observation starts.

The increase with time of the GW frequency is described by the *characteristic time* $\tau := f / \dot{f}$, equals to:

$$\tau \simeq 8.0 \text{ sec} \times \left(\frac{\mathcal{M}}{M_\odot} \right)^{-5/3} \left(\frac{f}{100 \text{ Hz}} \right)^{-8/3}. \quad (1.48)$$

Because ground-based detectors have a *sensitivity window* (further details in the next chapter) that starts from few tens of Hertz, signals from CBCs are observable as *transients* of excess power, whose typical time scales can be estimated from the previous equation. Assuming to start observing a coalescence at a frequency of 30 Hz, the characteristic time for a BBH ($\mathcal{M} \sim 10 \div 30$) is $\tau \sim 0.1 \div 1$ sec, while for a BNS ($\mathcal{M} \approx 1$) $\tau \approx 1$ minute. For the former, the most of the detectable GWs come from the merger phase as well as from the vibrational ringdown of the final BH. Instead, for the BNS coalescences the most of the GW emission happen during the inspiral phase.

GWs from BNS and NS-BH coalescences can provide information on the neutron stars equation of state [70, 86]. These phenomena are also candidates for triggering short *Gamma-Ray Bursts* (GRBs) [87]. The joint GW and GRB observation achieved on the 17th of August 2017 [14] (GW170817 and GRB170817a), accompanied by electromagnetic radiation emission in several other frequency bands, down to the low-frequency radio, has provided an unprecedented opportunity to study the mechanism of production of short GRBs, as well as a number of fundamental Physics concepts, like the possible evidence of higher dimensions of space-time, the propagation speed of GWs compared to that of light, and the polarisation modes predicted by GR [88, 89]. Also BHs are usually surrounded by matter (the *accretion disk*), which is expected to produce an electromagnetic counterpart at the time of the merger.

deformation of the Kerr one. This provides an accurate quasi-analytical description of the motion and radiation of a coalescing black-hole binary at all stages of its evolution [85].




Figure 1.3: Map of the four class of CBCs as functions of the masses m_1 and m_2 of the components involved, that is neutron stars, black holes, or objects belonging to the so called *mass gap*. Figure adapted from [26], licensed CC BY 3.0.

However, its amount, together with the typical distance at which BBH events are observed, make the detection of such a signal unlikely, and no evidence for that have been confirmed to-date [90].

Moreover, GWs from CBCs offer potential “standard sirenes”, independent of the “standard candles” represented by *type IA supernovae* [91], with which to measure cosmological distances. Indeed, the waveforms (1.47) can be used to determine the (luminosity) distance from the source [92]. If independent information are available on the possible location of the source, for example identifying its host galaxy from a joint GRB, the information on its redshift allows a determination of the Hubble constant H_0 [93]. This was the case for GW170817, which was used to make the first measurement of this kind [94].

The number of detectable CBCs depends on the number of coalescences per unit time in a certain volume of Universe. This is conveniently expressed by means of \mathcal{R} , the rate of event in a *Milky-Way Equivalent Galaxy* (MWEG) [95, §2]. This number can be predicted in two ways: *empirically*, based on the observations of compact binary systems that are thought to coalesce within a timescale comparable to the age of the Universe; and *theoretically*, based on models of binary star formation and evolution [95, §2]. Large uncertainties, however, arise on this number. For a comprehensive review refer to [96], while an up-to date estimate for BBHs, based on the detection accomplished by Advanced LIGO and Advanced Virgo, can be found in [97]. Multiplying \mathcal{R} by the estimated density of MWEGs, $\rho_{\text{MWEG}} \simeq 0.01 \text{ Mpc}^{-1}$, we can obtain a rate density of CBC events. This is shown in figure 1.4 for the different kids of CBCs, together with the upper limits on its values based on the data collectd by Virgo and LIGO during the S5 and S6 science runs.

Of course, the actual number of detected CBCs depends on the *sensitivity* of the detectors (and other factors that will be commented in section 1.3.4). We will describe this quantity and how to combine it with the rate density of CBCs in section 1.3.3.




Figure 1.4: CBC rate density estimates. The light grey regions represent the constraints obtained from the LIGO-Virgo S5/VSR1 data analysis, while the dark grey ones are those obtained with S6/VSR2-3 data, using the S5/VSR1 limits as priors. The lower (blue hatched) regions are the predictions based on population studies, as described in [98]. Figure adapted from [99].

1.2.2 Periodic sources

We call periodic those sources that radiate GWs at constant or nearly constant frequency, like radio pulsars. These compact stars are characterised by a spinning motion of up to several hundred revolutions per second. If they possess an *asymmetry* with respect to the rotational axis, hence $\ddot{I}_{ij} \neq 0$, from (1.44) they are expected to emit GWs. Such an asymmetry can arise for example if they have small “bumps” (~ 1 cm high, according to current theoretical models [100, 101]) on their surface, which could result from the turbulent environment characterizing a new born neutron star, frozen in the crust as it cools down in the first few seconds after its birth, or from “star-quakes”, fractures in the crust originating from changes in the surface stresses as the star’s spin rate slows. Global asymmetries can also result as a consequence of the strong inner magnetic fields that, if not aligned with the rotational axis, can produce pressures that distort the entire star. Refer to [102–104] for a review of the possible mechanisms of GW emission from pulsars.

Besides the mechanisms that can produce such an asymmetry, the consequent GW emission is well understood and can be computed, at a sufficient level of accuracy, entirely using Newtonian mechanics and Newtonian theory of gravitation, employing the so called weak-field, slow-motion, and small-stress approximation of GR [5, §2.5].

We can start from (1.36), approximating $\ddot{I} \approx (4\pi f)^2 \varepsilon I$, where ε is the ellipticity of the star, defined for example for a *triaxial* body rotating around the z -axis as:

$$\varepsilon := \frac{I_{xx} - I_{yy}}{I_{zz}}.$$

This factor is approximately $\varepsilon \approx 10^{-6}$ for a one centimetre “bump” on the neutron star crust.¹⁶ The dimensionless amplitude of the corresponding wave can then be approximated as:

$$\begin{aligned} h &\approx \frac{16\pi^2 G}{c^4} \frac{\varepsilon I f^2}{r} \\ &\simeq 10^{-26} \left(\frac{I}{10^{38} \text{ kg cm}^2} \right) \left(\frac{10 \text{ kpc}}{r} \right) \left(\frac{f}{100 \text{ Hz}} \right)^2 \frac{\varepsilon}{10^{-6}} \end{aligned} \quad (1.49)$$

where, in the second line, the typical values of the momentum of inertia I of a neutron star, the rotation frequency, and plausible values for its “ellipticity” ε , have been exposed [95, §2.4]. From the previous equation, we can notice that these sources are expected to produce comparatively weak GWs with respect to the “catastrophic” binary star mergers ($h \sim 10^{-26}$ vs. $h \sim 10^{-21}$). However, on their side they have the fact that this emission is *continuous*, hence the name *Continuous Waves* (CWs), and the detectors sensitivity can be significantly enhanced integrating their signal over time. The details on the related data analysis techniques are out of the scope of this dissertation, and we refer to the corresponding collaboration articles [105–107] or to the monograph [108].

Moreover, the number of possible CW sources is very large: $10^8 \div 10^9$ neutron stars in the Milky Way alone, on the basis of the supernova rate (about 1 every 50 years), the Milky Way age (about 13 billion years) and the assumption that most supernovae evolve in NS and not BH; further details in the next section. This, together with the simplicity in modelling their emission, makes periodic sources one of the preferred target of GW detectors.

From (1.49), the wave polarisations can be written as:

$$h_+(t, \mathbf{x}) = -h \frac{1 + \cos^2 \iota}{2} \cos 2\phi(t - r/c), \quad (1.50)$$

$$h_\times(t, \mathbf{x}) = -h \cos \iota \sin 2\phi(t - r/c) \quad (1.51)$$

with $\phi(t) = \phi_0 + 2\pi(f(t - t_0) + \frac{1}{2!}\dot{f}(t - t_0)^2 + \dots)$ the star’s rotational phase.

This signal is essentially monochromatic, up to correction of the order of the observation time over the spin-down time ($\dot{f} \sim 10^{-8} \text{ Hz s}^{-1}$). The rotational energy losses through the emission of GWs and electromagnetic radiation, causes a gradual decrease of the signal frequency. The actual signal frequency received at the detector is also modulated by the Doppler effect caused by the detector’s motion as the Earth rotates on its axis and revolves around the Sun.

Currently, an approximate number of 2500 neutron stars have been observed from their emission of electromagnetic radiation, and are monitored as candidate GW sources. No evidence of CW signals have been found so far. The most recent results on targeted searches by Advanced detectors have been published in [105, 106].

Without knowledge of the spin frequency, the spin-down rate and sky position, an enormous number of signal waveforms is possible, making searches extremely complex and computationally demanding. The updated results on an all-sky searches for insulated pulsar by the Advanced detector can be found in [107].

¹⁶For the well studied Crab pulsar, the upper limit on its ellipticity is $\varepsilon_{\text{max}}(\text{Crab}) \simeq 1.2 \times 10^{-4}$.

Recently, growing attention have been given to another “exotic” possible source of periodic waves. This is a “cloud” of ultralight bosons, such as QCD *axions* [109], surrounding a fast-spinning black hole. When the Compton wavelength of these bosons is comparable to the size of the black hole (so $m \lesssim 10^{-10} \text{ eV}/c^2$), these can bind to the black hole, forming a “gravitational atom”. Through the *superradiance process*, the number of axions occupying the bound levels grows exponentially, extracting energy and angular momentum from the black hole. This enables long-lasting, coherent and monochromatic GW emission from boson annihilation or from level transitions. Some constraints have been put on their masses and the GW emission mechanism. These can be found for example in [110, 111] or in [112] making use of Advanced detectors data.

1.2.3 CCSNe and burst signals

Compact objects, like neutron stars and black holes (of stellar masses), are formed in the gravitational collapse of a massive ($\gtrsim 8 M_\odot$) star, which leads to a *type II supernova* explosion, called a *core-collapse* supernova (CCSN) [113]. In this process, large amounts of mass ($1 \div 100 M_\odot$) flow in a compact region (hundreds to thousands of kilometres) at relativistic speeds ($v/c \gtrsim 0.1$). If this is not spherically symmetric, from (1.44), we expect to find an associated emission of GWs [114]. As in the case of CWs, also in this case the degree of asymmetry is the key aspect for determining the strength, and hence the detectability, of the associated GW emission. Contrarily to neutron stars, however, this asymmetry is not well modelled and predictable yet for supernovae. In general, we expect the relevant part of GW emission to be very fast, of the order of some tens of milliseconds, and the radiated energy $\lesssim 10^{-8} M_\odot c^2$ [115].¹⁷ This fact, together with the limited knowledge of the waveform emitted, partially reduce the prospects of detection mostly to our galaxy and the local group.

In the Milky Way, CCSN events are estimated by the Integral experiment to happen at a rate of one every 50 years, $25 \div 30$ in the local group, as deduced from the gamma rays produced by the radioactive decay of the rare isotope ^{26}Al [116]; the last one of such events was the famous SN 1987a in the Large Magellanic Cloud (51.4 kpc). Recently, the LIGO and Virgo Collaborations have published the constraints on the GW energy emitted during core-collapse for an optically targeted survey at distances up to 20 Mpc [117]. The results are $4.27 \times 10^{-4} M_\odot c^2$ and $1.28 \times 10^{-1} M_\odot c^2$ for emissions at 235 Hz and 1304 Hz respectively.

Because of our incomplete knowledge of the process of collapse and the emission mechanisms, GWs from these sources are typically described as unmodelled “bursts” of energy. Numerical models of stellar collapse are extremely challenging because of the complexity, and diversity, of the physical processes involved (electromagnetic, high-energy cosmic rays, neutrinos, etc.), that are important for such systems. This gives rise to a large variety of possible waveforms, depending on the particular scenario [118, 119]. For the same reason, research in this field is very active because the observation of the GWs from CCSN can potentially give a rare observation of the dynamics that are occurring in the core of the star, which is entirely obscured from electromagnetic observations.

¹⁷For comparison, the radiated energy in BBH events is $\mathcal{O}(1 M_\odot c^2)$.

Making an order of magnitude estimate, partially based on the ideas discussed at the beginning of section 1.2, one can obtain the dimensionless amplitude h of the emitted pulse of GWs:

$$h \simeq 1.4 \times 10^{-24} \left(\frac{\Delta E_{gw}}{10^{-8} M_{\odot} c^2} \right)^{1/2} \left(\frac{\tau}{1 \text{ ms}} \right)^{-1/2} \left(\frac{\Delta f_{gw}}{1 \text{ kHz}} \right)^{-1} \left(\frac{r}{15 \text{ Mpc}} \right)^{-1} \quad (1.52)$$

where ΔE_{gw} is the total gravitational energy carried away during the explosion, τ is the characteristic duration of the pulse and $\Delta f_{gw} \sim 200 \div 800$ Hz is its frequency bandwidth. The reference distance in the last term is that of the Virgo cluster of galaxies, while for a galactic supernova $r \sim 10$ kpc. Further details on this can be found in [5, §2.6].

Many other mechanisms have been proposed for the rapid emission of GWs. As for the case of CCSNe, these are generally described as unmodelled bursts of radiated gravitational energy. An example of a possible source is the radiation from *cosmic strings* [120]. These are one-dimensional topological defects that may have been formed in the early Universe during phase transitions due to symmetry breaking. Informally, they have been compared to the “cracks” (topological defects) that form when water freezes [121]. When two strings intersect each other at two points, or when a string crosses itself, they reconnect and an isolated cosmic loop is formed. These loops oscillate and radiate energy, mainly through GW emission. The strongest bursts of radiation are produced by *cusps* and *kinks* on the string, while also their fundamental vibration modes, for all the loops present in the Universe, are expected to superimpose into a *stochastic background* signal, as described in the next section. Their theory, as well as the various form their signals can take, has been reviewed in [122], while the constraints on their model parameters from the Advanced detector observations have been published in [121].

Burst searches are either all-sky or targeted. All-sky searches use multiple time shifts to target different points in the sky, and for this reason they are typically very resource demanding. Targeted searches use prior information from electromagnetic observations to aim at a specific sky direction, and they are usually more sensitive than the former.

1.2.4 Stochastic background

Whether weak or too distant to produce, individually, a signal strong enough to be detected, there are in the Universe a huge number of potential GW sources. However, the *incoherent* superposition of all of them can produce a “background” that could actually reach a detectable level. Differently from the signals described in the previous sections, such a background is not originating from a single *resolvable* source in the sky, but conversely it comes from everywhere. Having in mind the aforementioned astrophysical sources (CBCs, pulsars and supernovae), the uncertainties in the GW emission mechanisms and the parameters characterising each of them add up in this background, and the resulting signal is something that can’t be modelled as a deterministic one. A *stochastic model* is preferred, where the signal can be characterised only statistically, in terms of expectation values of the Fourier components of a plane-wave expansion of $h_{\mu\nu}$ (1.20), and for this reasons we refer to this as an *astrophysical Stochastic Gravitational-wave Background* (SGWB) [123].

After the discovery by the Advanced detectors of a population of merging black holes, we expect that, at some level, a stochastic background does exist, produced by all binary black hole (BBH) mergers over the history of the Universe [124]. This is hypothesised to become the principle observable in the not too distant future. Similarly, it is expected a background produced by BNS coalescences [125]. Other possible sources of an astrophysical background, of interest for ground-based detectors, are extragalactic supernovae [114] and non-axisymmetric highly spinning neutron stars, such as pulsars and magnetars, abundant in our galaxy [126, 127].

In addition to that, a background of gravitational waves is formed by all of the unresolvable galactic and extragalactic binaries in close orbit [128, 129]. This is expected to be relevant for future space-based detectors, like LISA; refer to figure 1.2.

Moreover, a *cosmological* SGWB could have been generated in the early Universe via a variety of possible mechanisms: amplification of primordial fluctuations in the Universe's geometry via inflation [130, 131], *phase transitions* as previously unified interactions separate [132, 133], a network of vibrating *cosmic strings* [134, 135], or the *condensation of a brane* from a higher dimensional space [64, 75]. Again, due to the large uncertainties in the proposed mechanisms and in the environment at the time of production, a stochastic model for the signal is preferred. Being the only known way to investigate the earliest moments after the Big Bang, the cosmological background is often referred to as the “Holy Grail” of GW astronomy [136].

The statistical characterisation of the SGWB resembles that of the detector noise, which is the main topic of the next chapter. This constitutes a complication, also due to the smallness of the expected amplitude of this background, that makes practically impossible to detect it using only one detector. We will come back to the detection technique for this background in section 2.1.4.

The magnitude of the SGWB is usually reported in terms of cosmological units of energy density per logarithmic frequency interval, $\Omega_{gw}(f)$, corresponding to its energy density ρ_{gw} normalised to the *critical density* for a closed Universe:

$$\Omega_{gw}(f) := \frac{1}{\rho_c} \frac{d\rho_{gw}}{d \log f}, \quad \rho_c = \frac{3c^2 H_0^2}{8\pi G} \simeq 1.90 \times 10^{-26} h_{100}^2 \text{ kg m}^{-3}$$

with $h_{100} = 0.70(2)$ the dimensionless Hubble parameter: $H_0 = h_{100} \times 100 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Due to the relatively large uncertainty on h_{100} , it is often preferable to quote the value of $h_{100}^2 \Omega_{gw}(f)$ in order to separate the knowledge on h_{100} from that on the SGWB itself. Following [5], its characteristic dimensionless amplitude can be expressed as:

$$h_c(f) = 2.8 \times 10^{-23} \left(\frac{h_{100}^2 \Omega_{gw}(f)}{5 \times 10^{-6}} \right)^{1/2} \left(\frac{f}{100 \text{ Hz}} \right)^{-1}. \quad (1.53)$$

The value at the denominator in the first parenthesis is the constraint on $h_{100}^2 \Omega_{gw}(f)$ from the observed abundances of light elements according to Big Bang Nucleosynthesis. This constitutes an upper limit on the expansion rate of the Universe at the time of Nucleosynthesis, which constrains the total energy density at that time, including the one of a cosmological SGWB [74, 137]. Currently, this is the most important broadband constraint, although

tighter ones are available for narrower frequency bands, as for example at extremely low frequencies from the observed anisotropies of the cosmic microwave background, and at high frequencies from ground-based interferometric detectors. The updated value of the latter upper limit on an isotropic SGWB have been published in [138], and it corresponds to

$$\Omega_{gw} < 6.0 \times 10^{-8}$$

at 95% credible level for a frequency-independent (flat) background and

$$\Omega_{gw}(f = 25 \text{ Hz}) < 4.8 \times 10^{-8}$$

for a background of CBCs. Directional limits for an anisotropic background are described in [139].

Last but most important, there are all those sources of GWs that have not been pondered yet, and which will (hopefully) surprise the researchers, yielding to new discoveries in the future.

1.3 Gravitational wave detectors

In section 1.1.4 the effects of GWs on test masses (at rest with respect to the TT -gauge) have been described as variations of their proper distance. For their ability to measure distance differences very accurately, laser interferometers have been soon realised to be the perfect instrument for detecting GWs. One of the pioneers of such an idea was the Nobel laureate Rai Weiss [140], while the key theoretical aspects of using interferometry for this scope were firstly studied by Gertsenshtein and Pustovoi [141, 142]. The (main) optical layout of a modern laser interferometric detector of GWs, like Advanced LIGO and Advanced Virgo, is shown in figure 1.5. This consists basically in a kilometre-scale, *Michelson interferometer* configuration, where the light coming from a laser source (left) is split in two beams by a semi-reflective mirror (*beam splitter*, at the centre of the figure) before travelling in the two (orthogonal) arms of the detector, and being reflected back by the *end-mirrors*, which are highly reflective test masses positioned at the end of each arm.¹⁸ The interference of the recombined light at the *asymmetric port* of the beam splitter (on the bottom of the picture) is used to measure variations in the relative distance of the end-mirrors and the beam splitter. In the next two sections we will describe the detection principle of laser interferometric detectors, as well as the modifications required to achieve the high sensitivities necessary to detect GWs.

A slight modification of the previous idea is that hypothesised for *space*-based detectors, like the LISA mission [71]. This will consist of three satellites forming an equilateral triangle, of about 2.5 million kilometre size, with one spacecraft at each vertex. The centre of it will be placed at the Lagrangian point $L1$ of the Sun-Earth system, trailing the Earth by about 20 degrees. Each spacecraft will have its own laser, whose phase is “locked” to match that of the others by means of laser beams exchanged by the three. This will allow to keep track

¹⁸The Michelson interferometer configuration is currently considered the optimal one for detecting GWs. Other possibilities, like the equilateral triangular shape proposed for the Einstein Telescope [143], have been studied for example in [144, 145].

of the relative distance between them; observing phase differences (at higher frequencies than those expected from the motion of the spacecrafts) in the frequency of the laser beam that returns to a spacecraft with respect to that of the local laser, will be a witness of the transit of a GW. Of course, differently from ground based detectors, no mirror reflections are involved in this mechanism, which would have suffered of significant diffraction loss in the propagation back and forth between different spacecrafts.

A similar strategy is exploited making use of *pulsars* instead of the “artificial” reference provided by lasers. These rapidly rotating, magnetised neutron stars are in fact extremely stable clocks. If we monitor simultaneously an *array* of them (from this the name *Pulsar Timing Array* to this technique), the time of arrival of each pulse can be correlated, and, accounting for their angular separation, the effects of an incoming GW observed as irregularities in this value [72, 146]. Currently, there are three major collaborations active in this project, which are PPTA, EPTA and NANOGRV [147].

Historically, large efforts have also been put into the development of *resonant mass* detectors, that is mechanical systems consisting of some tons of Aluminium alloy or Niobium, which are put into resonance by the passage of a GW. Contrarily to the previous ones, these are only sensitive to GWs with very specific frequencies, equal to their excitation ones, usually of the order of a kHz, and for this reason they are called “narrow band” detectors. Examples of such detectors are ALLEGRO [148], AURIGA [149], NAUTILUS [150], EXPLORER [151] and NIOBE [152]. A thorough description of their working principle, sensitivity and noise sources, is described in [49, §8], or in [142, §13] including an historic account on Weber’s pioneering experiments [153].

1.3.1 Ground based interferometric detectors

In this section we describe the main ideas behind the working principle of ground-based, interferometric detectors. The goal is to obtain, as a consequence of the effects of GWs on test masses, as derived in section 1.1.4, the phase shift induced on the recombined light at the antisymmetric port of the detector, in a simple Michelson configuration. This is the fundamental observable related to their working principle, and, as we will discuss momentarily, it is a *gauge independent* quantity. The following treatment has been readapted from [142, §2] and [49, §9].

Let’s consider an orthogonal arms, Michelson interferometer configuration, as in figure 1.5. We are working in *TT*-gauge, where the beam splitter is put at the origin of the coordinate system, while the position of the end-mirror that terminates the x -arm defines the point with coordinates $(L_x, 0)$, while that of the other end-mirror defines the point with coordinates $(0, L_y)$, along the y -arm. We recall that in this gauge the coordinates of masses at rest are left unchanged by the transit of a GW. Let’s start from equation (1.40); calling t_2 the time at which the light, reflected by the end-mirror at time t_1 , returns back to the beam splitter, we have, for the x -arm:

$$\int_{L_x}^0 dx = -c \int_{t_1}^{t_2} dt' \left(1 - \frac{1}{2} h_+(t') \right) + \mathcal{O}(h^2)$$

where the minus sign holds for the wave travelling towards the origin of the coordinates. Summing the implicit solution for $t_2 - t_1$ of the previous equation to (1.40), we obtain the back and forth time spent by the light into the detector x -arm:

$$\begin{aligned}
t_2 - t_0 &= \frac{2L_x}{c} + \frac{1}{2} \int_{t_0}^{t_2} dt' h_+(t') + \mathcal{O}(h^2) \\
&= \frac{2L_x}{c} + \frac{1}{2} \int_{t_0}^{t_0 + 2L_x/c} dt' h_+ \cdot \cos(2\pi f t') + \mathcal{O}(h^2) \\
&= \frac{2L_x}{c} + \frac{h_+}{4\pi f} \left(\sin(2\pi f(t_0 + 2L_x/c)) - \sin(2\pi f t_0) \right) + \mathcal{O}(h^2) \\
&= \frac{2L_x}{c} + \frac{L_x}{c} h_+(t_0 + L_x/c) \operatorname{sinc}(2\pi f L_x/c) + \mathcal{O}(h^2)
\end{aligned} \tag{1.54a}$$

In the second line, the “zeroth order” approximation $t_2 = t_0 + 2L_x/c + \mathcal{O}(h)$ has been used in the integration limit since the integrand is already to the linear order in $h_{\mu\nu}$. In the last, line the *cardinal sine* $\operatorname{sinc}(x) := \sin x/x$ has been introduced, after the application of the prostaphaeresis formula for the difference of sines. Again, the effect of the GW, with respect to the “unperturbed” propagation of the laser light, is visible in the second term in the last line, proportional to the arm (coordinate) length L_x and the GW strain $h_+(t_0 + L_x/c)$.

In the y -arm, a similar analysis leads to a difference in sign in the term proportional to h_+ , due to the nature of the $+$ -polarisation; ref. to (1.27):

$$t_2 - t_0 = \frac{2L_y}{c} - \frac{L_y}{c} h_+(t_0 + L_y/c) \operatorname{sinc}(2\pi f L_y/c) + \mathcal{O}(h^2). \tag{1.54b}$$

Notice a slight abuse of notation in the last two equations: t_2 and t_0 are, in general, different quantities, referred to each individual arm. However, if we consider the light recombined at the beam splitter at a common time $t_{(2)}$, we can compute at what previous times, $t_0^{(x)}$ and $t_0^{(y)}$ (different, in general), the light has left the beam splitter to travel into the two arms. Multiplying it by the laser angular frequency $\omega_L = k_L \cdot c$, with k_L the laser light *wave number*, we can express the effect of the GW as an additional phase to the “unperturbed” propagation; for the x -arm:

$$\Delta\phi_x(t) := k_L L_x h_+(t - L_x/c) \operatorname{sinc}(2\pi f L_x/c) + \mathcal{O}(h^2) \tag{1.55a}$$

where we made the substitution: $h_+(t_0^{(x)} + L_x/c) = h_+(t - 2L_x/c + L_x/c) + \mathcal{O}(h^2) = h_+(t - L_x/c) + \mathcal{O}(h^2)$. Analogously, for the y -arm:

$$\Delta\phi_y(t) = -k_L L_y h_+(t - L_y/c) \operatorname{sinc}(2\pi f L_y/c) + \mathcal{O}(h^2). \tag{1.55b}$$

In (1.54a) and (1.54b) notice also an “unperturbed” phase difference, for the possible case the two interferometer arms are not of equal length. It is convenient to write the corresponding contribution as:

$$2k_L L_x = k_L(L_x + L_y) + k_L(L_x - L_y) := 2k_L L + \phi_0$$

where $L := (L_x + L_y)/2$ is the average arm length, and the term ϕ_0 vanishes if they are equal. For the y -arm:

$$2k_L L_y = 2k_L L - \phi_0.$$

In general, the two arm lengths should be as close as possible, modulo odd multiples of half the laser wavelength, in order to cancel many *common mode* noises,¹⁹ except for a *microscopic difference* (smaller than half a laser wavelength), whose utility shall become clear in a moment.

The recombined light at the antisymmetric port of the beam splitter will be given by the superposition of two electromagnetic waves, with phase differences given by the previous terms. There is also an additional π difference, equivalent to a global “−” sign between the two fields, given by the fact that the light travelling in different arms has been reflected by opposite sides (of the same face) of the beam splitter; from this the name “asymmetric port”. This is a consequence of the *Stokes’ relations* for transmission and reflection, and the two kind of reflections are called *hard* (π phase difference, in the less dense medium) and *soft* (no phase difference, in the more dense medium); ref. to [66, §7.3] for further details. Let us assume to attribute this sign difference to the y -arm.

Then, besides common amplitude factors, the recombined wave will be given by the sum of the two phase terms:

$$e^{2ik_L L + i\phi_0 + i\Delta\phi_x(t)} - e^{2ik_L L - i\phi_0 + i\Delta\phi_y(t)}. \quad (1.56)$$

To simplify the previous expression, two approximations can be exploited. Firstly, in the long wavelength approximation, the two terms $h_+(t - L_x/c)$ and $h_+(t - L_y/c)$ in (1.55) are almost equal, up to corrections of the order of the ratio between the arm length difference (~ 0.1 m) and the GW wavelength ($100 \div 10^4$ km): negligible. Secondly, also the cardinal sines are approximately equal, and in particular equal to 1 up to corrections of the order of $\frac{1}{6}(fL/c)^2$, which for current kilometre-scale detectors is $\lesssim 10^{-6}$, negligible if multiplied to the already small GW amplitude. Under these approximations, the two (1.55) become one the opposite of the other, $\Delta\phi_x \approx -\Delta\phi_y := \Delta\phi(t)$, and (1.56) simplifies to:

$$e^{2ik_L L} \left(e^{i(\phi_0 + \Delta\phi(t))} - e^{-i(\phi_0 + \Delta\phi(t))} \right) = e^{2ik_L L} 2i \sin(\phi_0 + \Delta\phi(t)). \quad (1.57)$$

The effect of the wave is evident in the sine term. Comparing the “unperturbed” ϕ_0 term with $\Delta\phi(t)$, defined in (1.55), we see again that the effect of the GW is equivalent to a variation of the relative difference of arm lengths:

$$\frac{\Delta(L_x - L_y)}{L} \simeq h_+(t - L/c). \quad (1.58)$$

This variation can be measured as a time dependent power by the photodetector at the output of the beam-splitter, on the bottom of figure 1.5. Taking the square amplitude of equation (1.57):

$$\begin{aligned} P(t) &= P_0 \sin^2(\phi_0 + \Delta\phi(t)) \\ &= \frac{P_0}{2} (1 - \cos(2\phi_0 + 2\Delta\phi(t))) \end{aligned} \quad (1.59)$$

where P_0 is the power of the light injected by the laser into the interferometer. Looking again at the left-hand side of equation (1.57), this variation of power can also be understood as due

¹⁹What we want is to maximise the detector sensitivity to differential mode variations, typical of the transit of a GW, whilst minimising those of common mode, involving the average arm length [95, §3].

to the “beating” of the *acoustic sidebands* ($10 \div 10^4$ Hz) produced by the GW on the *carrier* laser frequency. As we will see in the next section, this idea is at the base of the working principle of the resonant cavity called *signal recycling*.

With (1.59) at hand, we can make some considerations about the optimal configuration, in terms of arm lengths, a detector should have to maximise its sensitivity. As anticipated, the two arms should be as similar as possible to cancel many common motion noises. However, from the previous equation we see that if they were chosen to be exactly equal, $\phi_0 = 0$, as it is the case for the classical Michelson interferometer working at *dark fringe*, the effect of $\Delta\phi(t) \propto h_+(t)$ would be only quadratic to $P(t)$. This would make impossible, in practice, to detect GWs, both for the smallness of h and for the sources of *noise* in P_0 and, inside $\Delta\phi(t)$, in k_L [95, §3]. Let’s take than $\phi_0 \neq 0$, and expand to linear order in $\Delta\phi(t)$:

$$1 - \cos(2\phi_0 + 2\Delta\phi(t)) = 1 - \cos(2\phi_0) + \sin(2\phi_0)\Delta\phi(t) + \mathcal{O}(h^2). \quad (1.60)$$

From the previous equation, it is then evident that, choosing a microscopic arm length difference, $\phi_0 = k_L(L_x - L_y) \neq k\pi$, we can obtain a time dependent variation of the laser light power at the detector output proportional to $\Delta\phi(t)$ and then to h . This is a tunable parameter in the interferometer readout, and it is typically chosen of the order of some tens of picometres.²⁰ With such a choice, the recombined light at the asymmetric port of the beam splitter is said to be at *grey fringe* [155], and this method is at the base of the *homodyne* or *DC detection principle* [156].

In this discussion we have shown how a simple Michelson interferometer can be considered as a prototype GW detector, acting as a transducer from (proper) length, or light travel time, differences to output optical power, corresponding to a certain *gauge invariant* number of detected photons. In the next section, we will briefly review the enhancements needed by actual interferometric detectors to achieve the sensitivities required to detect GWs. Then, in chapter 2 we will discuss how many non-gravitational effects may enter equation (1.59), representing a *noise* source for the detection of GW.

1.3.2 Advanced interferometric detectors

In the previous section, we have described an idealised Michelson GW detector, whose mirrors act as test masses to the passage of a GW. Of course, these masses are not in *free fall* and their movement is also affected by many other non-gravitational causes, which are in general several orders of magnitude more intense than the expected gravitational strain. To attenuate these effects, the mirrors of the actual interferometers are suspended by sets of pendulums in cascade, which have the role to filter out ground motion at frequencies larger than the proper ones of the pendulums. In Virgo, such devices are called *superattentors*, and consists of 10 metres high towers of multi-stage, inverted pendulums. LIGO detectors adopt similar solutions that mix active and passive insulation. With this expedient, detector mirrors are made to behave, at a certain level of approximation, as *free falling* test masses in

²⁰This *microscopic* arm difference should not be confused with the “macroscopic one”, much larger than the laser wavelength, called *Schnupp asymmetry* and equals to 23 cm in Advanced Virgo, 83 cm in the previous Virgo design and 27.8 cm in LIGO. This was at the base of the *radio-frequency* (RF) *read-out* scheme, and it is still used to “lock” part of the interferometer, like the *Fabry-Pérot cavities* through the *Pound-Draver-Hall* technique [154].




Figure 1.5: A simplified scheme of the main optical layout of the Advanced Virgo detector, showing the power-recycling mirror (PRM), the beam-splitter (BS) that separates the light into the two beams that travel the West and North arms, formed by Fabry-Pérot cavities with input and end mirrors labelled WI and WE, and NI and NE respectively. Image credit Virgo public media gallery.

the horizontal (x, y) plane. Still, part of this *displacement noise* pass to the detector output, and its effect will be studied in the next chapter in relation to the detector sensitivity.

Other enhancements are needed to the detection principle of the simple Michelson configuration in order to achieve the required sensitivities for the expected GW strain of the sources described in section 1.2. We can introduce them starting with an order of magnitude estimate of the sensitivity of the Michelson interferometer and adding step by step the various modifications.

Equation (1.58) tells us that, in order to detect a GW, we should be able to measure relative length difference $\Delta(L_x - L_y)/L \sim h$. At the denominator, $L \simeq 3$ km for Virgo, and 4 km for the two LIGO detectors. At the nominator, the main scale of reference to which we can measure length differences is the wavelength of the laser: $\lambda_L := 2\pi/k_L$. For the main Nd:YAG laser used in Advanced detectors, this equals to 1064 nm. Comparing the ratio of the latter two quantities with the target $h \sim 10^{-21}$ of the most promising sources (ref. to section 1.2) reveals nine orders of magnitude of difference. However, with a sensitive photodetector, we can relate the smallest measurable distance not to the wavelength of the laser but to the number of photons detected by the photodiode in a time T . In turn, the maximum acquisition time for such photons should be of the same order of the period of the GW we aim to measure:

$$N_{\text{Photons}} = \frac{P_0}{hc/\lambda_L} T \sim \frac{P_0 \lambda_L}{hc f_{\text{gw}}}. \quad (1.61)$$

The uncertainty on the previous number is given by the Poisson statistic, $\Delta N_{\text{Photons}} \sim N_{\text{Photons}}^{1/2}$, and then in order to detect a change of length beyond the fluctuation given by

the chance (*shot noise*) we must have:

$$\Delta(L_x - L_y) \gtrsim \frac{N_{\text{Photons}}^{1/2}}{N_{\text{Photons}}} \lambda_L.$$

Putting all together:

$$h \sim \frac{\Delta(L_x - L_y)}{L} \gtrsim \frac{N_{\text{Photons}}^{-1/2} \lambda_L}{L} \sim 10^{-17}$$

for a 1 W laser power and $f_{\text{gw}} = 300$ Hz.

The previous sensitivity can be improved in two ways. Firstly, we can reduce the shot noise increasing the power of the laser. Unfortunately this doesn't come without consequences, since higher power implies an increment of the *radiation pressure noise*, besides the technical limitations in having an high power laser with the required stability; further details will be given in section 2.2.2. Secondly, one can think of increasing the arm length L . This solution, however, would be quite costly, and at most would produce a benefit of one order of magnitude for ground-based detectors of reasonable sizes. The neat solution to that was to implement *Fabry-Pérot resonant cavities* in the detector arms, as shown in figure 1.5.

A Fabry-Pérot cavity is an optical resonator that, when "locked on resonance", that is, to a length equals to an integer number of laser wavelengths, forces the light therein to be reflected back and forth, making many round trips, and increasing in such a way the "effective" optical path travelled. This behaviour is characterised by a parameter called the "finesse" of the cavity, approximately equal to:

$$\mathcal{F} \approx \frac{\pi \cdot \sqrt{r_{\text{ITM}}}}{1 - r_{\text{ITM}}} \quad (1.62)$$

where r_{ITM} is the *reflectivity* of the *Input Test Mirror* (ITM) of the cavity. This is a number close to 1, and typical value of the finesse was $50 \div 100$ for first generation detectors while it is about 450 for the Advanced ones [95, 157]. The effect of the wave on the phase shift in a Fabry-Pérot cavity can be computed to be [49, §9.3]:

$$\Delta\phi(t) = \frac{2\mathcal{F}}{\pi} k_L L h_+(t), \quad (1.63)$$

which shows an increment, if compared to (1.57) (after the simplification discussed at the end of the previous section), by a factor $G_{\text{FP}} := 2\mathcal{F}/\pi \sim 290$, called the cavity *optical gain*.²¹ Modern interferometric GW detectors exploit this idea adding to each arm an input mirror, with the aim of creating a Fabry-Pérot cavity in it; refer to figure 1.5. Then, repeating the computation of the previous section, and the argumentation at the beginning of this one for

²¹For simplicity, in the previous equation the frequency dependency of the optical gain has been omitted. This is due to the fact that, if the cavity length changes, the resonance condition gets lost and light power stored therein leaks away. Therefore, there is a characteristic time needed to "refill" it, and a characteristic *low-pass, cut-off* frequency, $f_c = c \cdot (4\mathcal{F}L)^{-1}$, above which the effect of the optical gain start to diminish [95, §3]. The characteristic "low-pass" behaviour, $G_{\text{FP}}(f) \sim \sqrt{1 + (f/f_c)^2}$, is clearly visible at high frequency (where the shot noise is dominant) in the sensitivity curves of the Advanced detectors, fig. 1.6 and 1.7. In Advanced Virgo $f_c \simeq 50$ Hz, while in LIGO $f_c \simeq 80$ Hz.

deriving the expected sensitivity to h , we have:

$$h \gtrsim \frac{N_{\text{Photons}}^{-1/2} \lambda_L}{G_{FP} L} \sim 10^{-20}.$$

The previous estimate is closer to the requirement $\sim 10^{-21}$, but not enough still. Then, similarly to the considerations about increasing the laser power, one may wonder why not to increase the finesse by one or two orders of magnitude. Unfortunately, besides the technical difficulties for a cavity some kilometres long, this would have the same consequences of the radiation pressure on the mirrors. As an alternative, one can instead put an additional mirror before the beam splitter. Indeed, if the two detector arms are (almost) equal, there is almost no light at the asymmetric port of the detector, arriving to the detection photodiode, as derived in the previous section; all the power is reflected back towards the laser source. Adding a new resonant cavity before the beam splitter, called *Power Recycling* (PR) cavity, has the twofold effect of avoiding high laser power going back towards the laser, while increasing the one circulating in the “enhanced” Michelson interferometer. The corresponding optical gain can be calculated as:

$$G_{PR} = \left(\frac{t_{PR}}{1 - r_{PR} \cdot r_{MICH}} \right)^2 \quad (1.64)$$

where t_{PR} and r_{PR} are the transmission and reflectivity of the PR mirror, respectively, and r_{MICH} is the reflectivity of the rest, the Fabry-Pérot enhanced, Michelson interferometer. This cavity increases the number of photons (1.61) by G_{PR} , which for the Advanced detectors is $\sqrt{G_{PR}} \sim 10$, hence:

$$h \gtrsim \frac{N_{\text{Photons}}^{-1/2} \lambda_L}{G_{FP} \sqrt{G_{PR}} L} \sim 10^{-21}.$$

These enhancements are those implemented in Advanced interferometric detectors to achieve the “reference” sensitivity of $h \sim 10^{-21}$. Additional improvements have been studied and partially implemented or scheduled for the next upgrades. One of these consists into inserting another mirror after the asymmetric port of the beam splitter, creating in such a way a resonance cavity for the light before reaching the detection photodiode, which is called *Signal Recycling* (SR) cavity. This cavity is (usually) *tuned* to be on resonance with the acoustic sidebands of the carrier laser light that are generated by the GW interaction with the detector, as reported in equation (1.55). In this way, the gravitational “signal” encoded in these sidebands is amplified with respect to the power of the “anti-resonant” carrier, as previously described for a Fabry-Pérot cavity. This is called SR cavity, and its net effect is a broadening of the response of the interferometer to GWs [158]. Currently (2019), it is installed in Advanced LIGO detectors, while it is foreseen to be added to Advanced Virgo before the next observing run (O4). An account of its effect is visible in figure 1.6 comparing, for example, the Advanced Virgo sensitivity achieved during O3a and that planned for O4, with SR active.

Quantum squeezing of light is an advanced Optics technique meant to improve the detector sensitivity by reducing the effects of quantum noise. These effects include the photon *shot noise*, already encountered as readout noise in (1.61), and the *quantum radiation pressure*

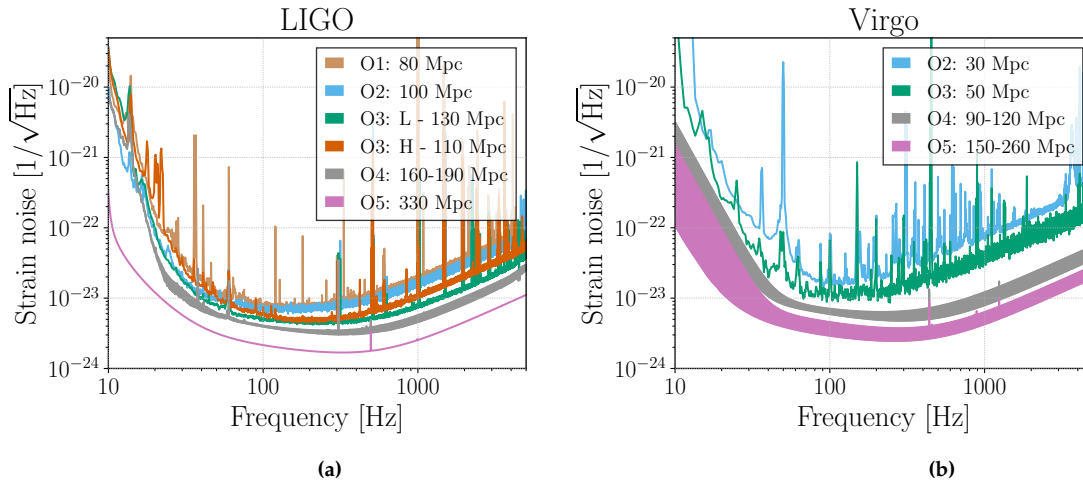


Figure 1.6: Sensitivities of the Advanced LIGO and Advanced Virgo detectors, expressed as the Amplitude Spectral Density (ASD) of the detector output assuming the absence of any GW signal. Estimates based on actual data are shown for science runs until O3a, whilst the bands shown for O4 and O5 are projections based on the scheduled improvements. Details about this quantity are described in section 1.3.3. Figure adapted from [26], licensed CC BY 4.0.

due to the variable number of photons impinging on test masses and transferring momentum to them; refer to figure 2.6. Together they constitute the so-called *standard quantum limit* to the sensitivity, which we will come back to discuss in section 2.2. To understand the origin of this noise, we recall, from the *Heisenberg uncertainty principle*, that there are unavoidable uncertainties in the conjugate variables amplitude and phase of the electro-magnetic field inside the detector. In a normal vacuum state, these fluctuations are uncorrelated. By using a crystal with non-linear optical properties, it is possible to prepare a special state of light where the fluctuations are mostly concentrated in one of the two variables, still without violating the uncertainty condition on their product [159, 160]. This technique has been pioneered in the GEO 600 detector [161] and used in the Advanced detectors to reduce the *shot noise* at frequency above $50 \div 80$ Hz by means of a “squeezed” vacuum state with phase uncertainty smaller than that of the normal vacuum [162, 163]. As a consequence, the radiation pressure noise, related to the higher amplitude uncertainty, has been observed to increase at low frequency. This constitutes, to some extent, a limit to the application of this technique to bypass the standard quantum limit. Radiation pressure can be reduced making use of *larger masses* for the interferometer mirrors, which however is quite a costly solution and has consequences also on all the suspension mechanisms. A neat work around to this has been studied and consists in the injection of a *frequency-dependent squeezed light*, with smaller amplitude uncertainty at low frequencies, in order to limit radiation pressure, and smaller phase noise at high frequency, to limit shot noise [164, 165]. This will constitute probably the major upgrade to the Advanced interferometers for their upcoming observing runs.

The sensitivities of Advanced LIGO and Advanced Virgo achieved during the first two observing runs (O1 and O2), and the first part of the third one (O3a), are reported in figure 1.6, together with the projections for the fourth and fifth runs. In particular, the improvements expected for the Virgo sensitivity due to the introduction of a signal recycling




Figure 1.7: Prospects on Kagra sensitivity, expressed by means of the projected detector output *Amplitude Spectral Density* (ASD) (2019 update). Figure taken from [26] licensed CC BY 4.0.

cavity are clearly visible at high frequency of fig. 1.6b [26].

In December 2019, during the second half of the third observing run (O3b), the Japanese detector KAGRA [166] is expected to join the Advanced detector network. The previsions about its sensitivities are reported in figure 1.7. The optical design and general infrastructure of KAGRA is similar to that of the other detectors, except for two significant improvements. First of all, the whole detector will be underground, in the Kamioka mountain site. This will reduce the micro-seismic noise on test masses, as the associated waves propagates mostly on surface, and the Newtonian gravity gradient noise, being smaller the relative changes of the mass density in the environment around the detector. The second improvement is that all the test masses will be cryogenically cooled to 20 K. This will have the advantage of reducing the thermal noise produced by their vibration. Further details about these and others noise sources will be described in the next chapter. These additions will also provide a testing ground for the development of the next generation of GW detectors, such as the European Einstein Telescope [143] or the American Cosmic Explorer [167].

What we discussed here is just a simplified sketch of the working principle of actual detectors. In reality, the whole infrastructure is a much more complex apparatus, made of many interacting subsystems [LVC5]. This brief description is only meant to introduce the fundamental concepts related to the quantity we shall describe in the next sections, or that we will refer to in the rest of this manuscript.

1.3.3 Some sensitivity benchmarks: DARM, $h_{\text{rec}}(t)$, BNS range

In the previous sections we have repeatedly made use of the term *sensitivity*, referring to, loosely speaking, “how good” a detector is able to measure the effects of GWs. In this section we will elaborate on this concept. We will present the sensitivity evolution of the Advanced detectors, and, in relation to that, their detections in section 1.4.

For the output of the simplified Michelson interferometer described in (1.59), a first way to represent its sensitivity is as the variation of *output* power for a given *input* variation,

equals in this case to the phase shift $\Delta\phi(t)$ produced by the wave, which in turn is proportional to the strain h :

$$\frac{\partial P}{\partial \Delta\phi} = \sin(2\phi_0) + \mathcal{O}(h).$$

Neglecting the contribution of the wave, this quantity is maximum for $\phi_0 = \pi/4 + k\pi$, which corresponds to a grey fringe condition where the response to the incoming GW is approximately linear. This “naive” working point however doesn’t take into account the possible noise sources entering in (1.59) [49, §9.3], like the aforementioned common mode noise. Noise will be the main subject of the next chapter, where we will describe its sources and how they dominate the detector output. Most importantly, we will proceed there with its statistical characterisation as a *stochastic process*. Hence, also the description of the detector sensitivity should be properly done statistically. Details on this will be formally presented in section 2.1.

Going back to the effects of GWs on the detector, we have described how they primarily affect the differential arm length (DARM), $\Delta(L_x - L_y)$. This corresponds to a degree of freedom that is constantly controlled by actuating differentially on the two end-masses in order to keep the interferometer *locked* on resonance.²² Hence, from the record of the corresponding control channel, named LSC_DARM, we can obtain a primary information about the possible presence of a GW signal and the general performance of the detector to this aim. This is indeed one of the main target for the studies that we will present, and in particular it was considered as the fundamental one during the commissioning phase in preparation to the O3 science run.

The information about the strain produced by the passing GW can be reconstructed from DARM by removing the contributions of the control signals and calibrating for the interferometer optical response transfer function. This is reconstructed from signal injections that simulate the effects of the passage of a GW with known waveform [168, 169]. The resulting time series corresponds to the *reconstructed value of the strain*, $h_{\text{rec}}(t)$, whose record is saved into the readout channel called Hrec_hoft in Virgo and CALIB_STRAIN in LIGO. This is the main input for GW searches and for the analysis described in section 1.4.

Again, due to the many noise sources that can enter into $h_{\text{rec}}(t)$, *a priori* this is not meant to represent an actual signal of astrophysical origin. No claim about the presence of a GW can be made before some statistical test is able to validate it. For this reason, what $h_{\text{rec}}(t)$ actually represents is the so called *strain equivalent noise*, $n(t)$, plus (hopefully) a true gravitational signal, whose presence is the objective of our inference: $h_{\text{rec}}(t) = n(t) + h(t)$. Some detection algorithms aimed at this will be described in section 1.4.

As we will discuss in the next chapter, the properties of the noise, $n(t)$, are conveniently described in the so called *frequency domain*, where specific noise sources have very characteristic signatures. Then, the detector, frequency dependent, *strain sensitivity* is defined as the *amplitude spectral density* of $n(t)$ (square root of its *one-sided* power spectral density, $\sqrt{S_n^{1s}(f)}$; refer to the conventions discussed in the Notation section), or equivalently of $h_{\text{rec}}(t)$ in the

²²As well as DARM, other important degrees of freedom, one for each independent resonant cavity in figure 1.5 [95, §5], are constantly controlled by the *longitudinal stability control* (LSC) system in order to keep the various detector cavities “locked” on resonance and the whole apparatus to work around its optimal configuration [12].

null hypothesis of the absence of any GW signals ($\tilde{h}_{\text{rec}}(f) = \tilde{n}(f)$), and it represents the minimum value of $\tilde{h}(f)$ detectable by our instrument.²³ Moreover, under the assumption of stationary and Gaussian noise, this quantity provides a complete statistical characterisation of the detector output, hence its sensitivity. We will justify this claim in section 2.1.2. The evolution of the strain sensitivities of the Advanced detectors are reported in figures 1.6 and 1.7.

The strain sensitivity represents, frequency by frequency, a figure of merit for the capabilities of the detector. Often, such a detail is not necessary, or it is just preferable to quote a single number rather than a full spectrum of values. A common benchmark for the “overall” detector sensitivity to GWs, which answers the recurrent question “how many coalescences of compact objects belonging to a certain population can we detect?”, is represented by the *detector range*. This is defined as the distance, averaged over polarisations and directions in the sky, at which a single detector can observe with *Signal-to-noise Ratio* (SNR) equals 8 the inspiral of compact objects with chirp mass \mathcal{M} [98, 170]:²⁴

$$\frac{d_{\text{range}}}{1 \text{ Mpc}} = \frac{1}{2.26} \times 1.95 \times 10^{-20} \left(\frac{\mathcal{M}}{M_{\odot}} \right)^{5/6} \sqrt{\int_{f_1}^{f_{\text{ISCO}}} \frac{f^{-7/3}}{S_n^{1s.}(f)} df} \quad (1.65)$$

The factor $1/2.26$ accounts for the average on polarisations and possible source directions. The dependency on $f^{-7/3}$ is characteristic of the inspiral phase in the adiabatic Newtonian approximation [170, §III], while f_{ISCO} is the frequency of the *innermost stable circular orbit*, where the inspiral ends and begins the merger, equals, to lowest order Post-Newtonian approximation, to [50, §4]:

$$f_{\text{ISCO}} \simeq 4.4 \text{ kHz} \left(\frac{M_{\odot}}{M} \right).$$

As a reference, it is convenient to consider the coalescence of two neutron stars, whose mass range is typically well clustered around $1.4 M_{\odot}$, instead of (stellar mass) black holes, which are known to possess masses spanning several decades of solar masses. In this case, we define, for $m_1 = m_2 = 1.4 M_{\odot}$, the so called *BNS range* substituting in the previous equation $\mathcal{M} \simeq 1.22 M_{\odot}$ and $f_{\text{ISCO}} \simeq 1.57 \text{ kHz}$. This value, together with the density rate of BNS coalescences discussed in section 1.2.1, provides an estimate of the number of events of this kind detectable by Advanced detectors [95, §2]. For example, with the median estimated density rate of figure 1.4 (dashed black line), $10^{-6} \text{ Mpc}^{-3} \text{ yr}^{-1}$ [98], and the Advanced Virgo target sensitivity of 60 Mpc for O3, we should be able to have an average of 0.9 BNS events per year. With the LIGO Livingston sensitivity (130 Mpc) this number equals 9.2 yr^{-1} . These estimates can be significantly improved operating with a network of detectors, as we will describe in the next section, which allows to increase the collective SNR of an event, and reduce the effects of non-stationary noise, as described in section 1.4 in relation to the analysis pipelines.

In figure 1.8 the evolution of the BNS range, for the current and projected GW detectors, is reported during their past and planned observation runs. Compare this values also with

²³This shouldn't be confused with the *upper limit* that the detector can establish on the strain amplitude of a certain GW signal, which, conversely, accepting the presence of that signal, corresponds to the maximum value of its amplitude that our instrument could have missed to detect.

²⁴From the matched filter technique that we shall describe in section 2.1.6, a value of 8 for the SNR in stationary and Gaussian noise is usually considered safe for calling a detection [171].




Figure 1.8: Timeline of the planned observing runs of the Advanced GW detectors, including Kagra [166], which should join Advanced LIGO and Advanced Virgo for the end of O3, and the third LIGO detector located in India [172], currently under construction. The sensitivity (range) of each detector for the various runs is expressed in terms of its past or projected BNS range (updated Sept. 2019). Picture taken from [26].

the detailed evolution of the strain sensitivities in figures 1.6 and 1.7.

1.3.4 A network of GW detectors

In the previous sections, we have mainly examined one single detector at a time. Instead, in the context of GW research primary importance has been addressed, by the collaborations involved, to the establishment of a *network* of detectors, by means of which to confidently identify and *triangulate* the origin in the sky of potential sources of GWs, especially transients. Indeed, interferometers are intrinsically non-directional. However, when a GW signal reaches the Earth it will arrive at the different detector sites at slightly different times, of the order of some milliseconds (the light travel time between the sites, which is *up to* 10 ms for the two LIGO detectors, and 26.39 and 27.20 ms for Virgo). The observed arrival time delay, and also the amplitude of the GW signal measured at each location, allows to localize the direction in the sky from which the signal has come. Using gravitational wave information alone, two detectors can constrain the direction to lie somewhere on a “broken ring” on the sky (refer to figure 1.11); it is necessary to have at least three detectors for localising it within the intersection of two of such rings, hence providing much better precision.²⁵ This information is of fundamental importance in order to alert the partner observatories to focus on that area of the sky in search for possible electromagnetic counterparts, as it was the case for GW170814 [13] or GW170817 [14, 15]. In the latter case, also a low signal amplitude in Virgo was useful to locate the source in those regions where the detector was less sensitive, which, combined with the *ring* identified by the two LIGO detectors alone, helped to reduce the network 90% credible sky localisation area to 28 square degrees.

²⁵A detailed analysis of the reconstruction of the sky position of a transient source, mainly based on the time delay of the signal received by the detector network, is presented in [173, §15].

In practice, also three detectors are not always sufficient for an accurate sky localisation. Indeed, the sky sensitivity, called “antenna pattern function”, of each interferometric detector doesn’t cover the entire sky but only a fraction $2/5$ of it, obtained averaging over all the directions. A figure of merit about this is the so-called *sky-coverage*, defined as the fraction of the sky over which the detector antenna pattern function is greater than half of its maximum value. This equals 33.6% for a single detector [80], which means it is able to efficiently detect GWs only from one third of the sky, and, most importantly, that there are regions where it is completely “blind”. Having multiple detectors increases the *network antenna pattern function*. For example, it can be shown that with the addition to the LIGO detectors of Virgo, the total sky-coverage rises to 71.8%. With the further addition of KAGRA (expected for the end of O3) this will be 73.5%. Refer to [174] for further details and for other means by which to characterise the network sensitivity of the planned and upcoming detectors.

Another fundamental aspect is that related to the rejection of instrumental artefacts and the avoidance of false detections. Indeed, some transient signals of *terrestrial origin*, for example from system malfunctions or environmental disturbances, may randomly mimic a waveform compatible with a certain GW model in the data from one detector. However, if a compatible waveform is not present, within the light travel time, in the other detectors that could have seen a similar signal, because acquiring data at the time of the putative event and with antenna pattern function sensitive to the direction identified by the first detector, then these events are discarded.²⁶ Coincidence of the signals detected by multiple interferometers is at the base of the analysis pipelines that we will describe in section 1.4. Also, the preliminary identification of instrumental artefacts potentially resembling a GW signal is of fundamental importance, and it will be thoroughly discussed in chapter 3.

Last but not least aspect of importance in the establishment of a network of GW detectors is related to their *duty cycle*. Transient periods of noise excesses or maintenance can let one detector temporary out of observing mode, reducing their effective number for GW searches [24]. This is quite a common situation, indeed maintaining the optimal working point outlined in section 1.3.2, with all the cavities locked on resonance, is sometimes difficult, especially in the presence of instrument malfunctions or external disturbances, like bad weather conditions. For reference, during O2 the individual duty cycle of the two LIGO detectors was approximately 60%, and about 45% that of the LIGO network, corresponding to 118 days of data suitable for coincident data analysis [1]. During August 2017, when Virgo joined O2, its individual duty cycle was above 80%, and the network one, with at least 2 detectors, reached about 86%.²⁷

All these reasons motivated the establishment of the current network of GW detectors, and the construction of newer ones, as we momentarily comment in the next section. Also, the causes that may reduce the detector duty cycle have been one of the main objective of

²⁶Only recently, thanks to the improved detector sensitivities, the confidence on the waveform to expect, both for the known astrophysical signals and the noise, and the analysis pipelines to detect the former and reject the latter, the LIGO and Virgo collaborations have opened to the possibility to consider also single detector candidate events, which of course require a dedicated validation procedure [175].

²⁷The duty cycle of the LIGO-Virgo network during O3 was of about 62.52% with two detectors, 23.21% with three, 8.90% with one, and only 5.37% of the time with no detector observing the sky. Notice that these are only preliminary estimates, further subject to quality cut (as we will discuss in section 2.4.4), gathered from the Virgo interferometer monitor [LVC6].

the studies that will be presented in this manuscript. We will come back to this subject in several occasions in the next sections and chapters.

1.3.5 New detectors and next generations

Following the motivations discussed in the previous section, in the next years, the current network of GW detectors will be enlarged with the already mentioned Japanese detector KAGRA [166], and a third LIGO detector located in India [172]. When all of them will be operating, the sky coverage will be almost complete and it is estimated that they will be able to achieve localisations of just $9 \div 12$ square degrees (median) [26]. Moreover, the most significant improvement will arguably be the possibility to have, almost always during science runs, at least three detectors observing in low noise conditions and able to provide accurate sky localisations.

Besides the increase of number of *second generation* detectors all over the world, there are under study improvements to upgrade further the current ones. For example, within the LIGO and Virgo Collaboration, there are plans to upgrade Advanced LIGO and Advanced Virgo detectors to A+3 and AdV+4 phases respectively, by around 2024 [26]. Their sensitivities are improved over the current version of detectors by roughly a factor of two. This is in part realized by reducing the coating thermal noise either from reducing the mechanical losses of the coating material or from implementing larger beam size. Also, broadband quantum noise reduction is expected by using a 300 metres filter cavity to generate frequency dependent squeezing [176], as described in section 1.3.2.

Completely new infrastructures have also been envisioned for the future, such as the Einstein Telescope [143] or the Cosmic Explorer [167, 177]. Besides state of the art technologies, for which Kagra will be a demonstrator, one of the major improvements in these detectors will be the increase in their arm length by about one order of magnitude ($10 \div 40$ km) with respect to LIGO and Virgo. This modifications will not change significantly the sensitivity band, except for some improvements down to 10 Hz for the reduced seismic noise, therefore the target sources for these detectors will be the same as for the second generation ones. However, their better sensitivity will allow to reach furthest and weakest sources [178].

A complementary point of view can be obtained with detectors based on different design, as for example the space-borne Laser Interferometer Space Antenna (LISA) [179], whose working principle has been briefly described at the beginning of section 1.3. As visible from figure 1.2, these detectors will open up to a new band of observation of the GW signals. Combined with that of ground based interferometric detectors, it will be possible that, for some sources, there will be multiband GW observations, with these detectors providing an “early warning” and sky localization for ground based ones [180]. Also, the data in this new band could provide additional information on system parameters [181] and new tests of general relativity [182].

1.4 Detections of the Advanced detectors

We conclude this chapter joining “sources” (more precisely, CBCs) and “detectors”, describing some of the fundamental aspects of the GW detections achieved by Advanced LIGO and

Advanced Virgo detectors as of December 2019. For brevity, we will mostly focus on the “first detection” event, GW150914 [8], slightly touching the general features of all the other detections achieved during the first two observing runs, and summarised in the first Gravitational Wave Transient Catalogue [1]. In particular, we will present those aspects, most closely related to the detector characterisation, of importance for assessing the significance of the events.

1.4.1 GW150914: the first detection

The first direct detection of GWs has been accomplished on the 14th of September 2015, at the very beginning of the first observing run of the Advanced LIGO detectors. In fact, the two detectors were still in “engineering mode” at the time of the event, meaning that they had finished the “commissioning” activities, and they were testing the behaviour of the apparatus; it was both surprising and rewarding the observation of such a strong signal right after the five years of upgrades spent to improve the sensitivity of the first generation of LIGO detectors. Then, the two detectors continued taking data for 38.6 days (~ 17 in coincidence), without modifications to the apparatus, in order to achieve the sufficient amount of information to trustworthy characterise the detectors and assess the significance of the event. Further details will follow momentarily.

This event comprised two major discoveries (merging stellar-mass black holes and *direct* GW observation) and two further confirmations for Einstein’s GR (GW prediction and CBC waveform), as summarised in the conclusions of the “Discovery paper” [8], which we will quote:

The LIGO detectors have observed gravitational waves from the merger of two stellar-mass black holes. The detected waveform matches the predictions of general relativity for the inspiral and merger of a pair of black holes and the ringdown of the resulting single black hole. These observations demonstrate the existence of binary stellar-mass black hole systems. This is the first direct detection of gravitational waves and the first observation of a binary black hole merger.

The result was worldwide announced by the *LIGO Scientific Collaboration and Virgo Collaboration* (LVC) on the 11th of February 2016. For this discovery Reiner Weiss, Barry Barish and Kip Thorne were awarded the 2017 Nobel Prize for Physics, “for decisive contributions to the LIGO detector and the observation of gravitational waves” [183].

A total of twelve Collaboration articles accompanied the one announcing the detection, and several more were published later, elaborating on various aspects ranging from the detector working principle and status at the time of the event [23, 184, 185], the assessment of its statistical significance [186], and the inference on black hole populations [187] and tests of GR [188]. Here we will mostly focus on those aspect in which the characterisation of the detector and its noise are most relevant, in particular for assessing the significance of the event. These are also important for the analysis that we will present in the rest of this work. A side aspect of fundamental importance, which however we will not have occasion to touch, is that of the detector calibration and its uncertainties, which allows to convert the intensity of




Figure 1.9: Search results for PyCBC (left) and GstLAL (right) pipelines, showing the histograms of the background distributions of the detection statistics (blue and black lines) together with the candidate events (orange squares). The second square from the right is “the boxing-day event”, GW151226 [192], while the third one, slightly above the background level, is GW151012, labelled at the time as a “LIGO-Virgo trigger” (LVT) [186] and only later promoted to an actual GW event after reanalysis [193]. Figure adapted from [8], licensed CC BY 3.0.

the laser light at the output photodiode (1.59), and several other controls, to the gravitational strain (1.47); the details have been published in the Collaboration paper [185].

At the time of this event, data from GW detectors were constantly monitored by two data analysis algorithms (*pipelines*) explicitly devoted to CBC searches, PyCBC [17] and GstLAL [18]. These use the *matched filter* technique, which we shall outline in section 2.1.4, to compare detectors data to templates representing the expected GR waveform for the coalescence of two BHs. These templates are obtained by two independent waveform models, SEOBNRv2_ROM_DoubleSpin [189] and IMRPhenomPv2 [190], both based on EOB and numerical relativity, and yielding consistent predictions. Once the best match has been found, within a huge template bank covering (the most of) the *parameter space* (masses, direction, distance, etc.),²⁸ the two pipelines compute the probability that this match would happen just by chance, from the random fluctuations of the detector noise. This part is evaluated in a slightly different way between the two pipelines, although they both use a detection statistic based on the SNR.

PyCBC makes use of a “reweighted” SNR $\hat{\rho}_c$ that evaluates the *consistency* of the signals of the two detectors in different frequency bands; this helps taking into account for random fluctuation of the signal, which may lower the consistency in certain frequency bands, or loud glitches, which may, conversely, have an high consistence in some bands but none in others.²⁹ The consistency is quantified by a *reduced* χ^2 -statistic, used to reweight the SNR. To calculate the *false alarm rate* given by the random fluctuation of the previous statistic, that is the probability to obtain a detection statistic as high or more if we only had noise, PyCBC uses the *time slide* method, shifting in time the data of one detector with respect to the other by an amount larger than the coherence time of a gravitational signal, and making in such a way the coincidence by two matches just due to the chance.

²⁸Further details can be found in the “Parameter estimation” paper [191].

²⁹Both these possibilities are clearly evident from the time-frequency plots for GW150914 and LVT151012, as reported in figure 10 and 12 of the Detecto Characterisation paper [23].

GstLAL instead, after the best matching template and the corresponding SNR has been found, computes the *residuals* of its subtraction to the detectors data. Then, it evaluates the likelihood-ratio statistic \mathcal{L} between the probability of finding those particular values of the SNR and the residuals in both detectors, and the probability of finding the same results in the case of noise only. The latter *background* is evaluated looking at those “triggers” (high SNR data) present in one detector only (and then assumed not to be of astrophysical origin; refer to the discussion in section 1.3.4), and computing the corresponding distribution function using Monte Carlo sampling methods. This gives the probability of having similar triggers due to noise and, similarly to the other pipeline, it is a measure of the significance of an event.³⁰

The results of the two pipelines for O1 are reported in figure 1.9. GW150914, corresponding to the rightmost orange square in the two plots, is the most significant event detected by both searches. The corresponding false alarm probabilities are $< 2 \times 10^{-7}$ and 1.4×10^{-11} for PyCBC and GstLAL respectively.³¹

Two comments should be made about the aforementioned methods. First of all, the matched filter technique used to compute the detection statistics is optimal on the assumptions that the detector “noise” is *stationary*, and the corresponding statistical properties are known. In reality, however, these properties are themselves estimated from the data (usually assuming *ergodicity*), up to certain distribution moments (usually the second, making the implicit assumption of *Gaussianity*), and known up to a certain limited *statistical* precision. Secondly, and even tightly related to the detector characterisation activities, there is the problem of *vetoing* those segments of data that exhibit “unusual behaviours” and, most importantly, we know with a high level of confidence that are not of astrophysical origin. As we will elaborate in section 2.4, these are usually related to *transient* (with different time scales) detector malfunctions, and often classified into *glitches* and slower *non-stationarities*. It is important to exclude them from the analyses, once we know for sure their origin, since otherwise they will introduce systematics in the noise properties estimation (related to the previous point) and also they would spoil the “background” estimates in figure 1.9. These aspects will be elaborated in the next chapters, where further details on the characterisation of the detector will be provided, as well as the statistical tools devoted to this.

We can report then the properties of GW150914 source, obtained from the previous data analysis pipelines, and arguably the most interesting part related to this event. The component masses, which are the most important parameter governing the dynamical evolution of the coalescence (in Newtonian adiabatic approximation (1.47)),³² are 35_{-3}^{+5} and 30_{-4}^{+3} solar masses *in the source frame* ($m_{1,2}^{\text{source}}$) at 90% credible level, obtained as an average between the

³⁰To elaborate a bit on the difference between the two pipelines, and the importance of having both of them, it should be mentioned the case of GW170818, reported in [1, §B.VI]. This was a triple detection event, where also Virgo contributed with its data, but the SNR in LIGO Hanford (and Virgo, although PyCBC didn’t consider it in the computation of the detection statistic) was not high enough to pass the consistency threshold required by PyCBC. Instead, the analysis of the residuals made with GstLAL (and using all three detectors) identified this as an actual event with a pretty decent (network) SNR of 11.3 (4.2 in Virgo, 4.1 in Hanford, therefore below the PyCBC threshold of 5.5, and 9.7 in Livingston).

³¹Actually, PyCBC was able to calculate only an upper limit on this value. Further details on this and also on the systematic errors on the value computed by GstLAL are reported in the corresponding Collaboration article [186].

³²To be more precise, what we “directly” measure is the *chirp mass* during the *inspiral* phase, as described in (1.47), and the *total mass* during the following *merger* and *ringdown* phases. From these, the values of the two masses are extracted.

results of the two waveforms, SEOBNR and IMRPhenom, in perfect agreement: refer to figure 1 from the corresponding “Parameter estimation” paper [191]. The mass of the resulting post-merger black-hole is $M_f = 62^{+4}_{-3} M_\odot$, implying that the missing 3.0 ± 0.5 solar masses were radiated away as GWs; all quantities referred to the source frame.

In fact, the masses that we measure from Earth are slightly larger (frequencies are smaller) as an effect of the redshift z at which the source is located. This, in turn, is related to the *luminosity distance* of the source D_L , which affects the amplitude of the detected wave (refer to (1.47), where this quantity was labelled with r). This is the second most important quantity, and was estimated to be 440^{+160}_{-180} Mpc ($z = 0.09^{+0.03}_{-0.04}$). Also the *inclination* of the source has a similar effect; instead of ι , which would be time-dependent for a precessing system, it is more convenient to report the inclination of the total angular momentum (which typically is approximately constant throughout the inspiral) with respect to the line of sight: θ_{JN} (at $f_{\text{ref}} = 20$ Hz). The two quantities D_L and θ_{JN} , which are strongly correlated, are usually reported on the same plot; refer to figure 2 in [191].

After the masses, the most important quantity describing the evolution of the coalescence are the spins of the two components. These are measured, for example, from the precessing IMRPhenom waveform model, and reported in figure 5 in [191]. The results favoured two components with small values of the spins, or larger values but anti-aligned with the orbital angular momentum.³³

The Numerical Relativity waveforms, for each phase of the coalescence (top cartoon), for the inferred source parameters are shown in figure 1.10 for the LIGO Hanford detector. The shaded grey region is the strain reconstruction using as a linear combination of sine-Gaussian wavelets [195, 196]. In the same picture, on the bottom, the Keplerian effective black hole separation, in units of Schwarzschild radii, and the effective relative velocity are reported.

We conclude the section reporting briefly also the results on GW150914 from the unmodelled searches, and described in the Collaboration article [196]. Three pipelines of this kind analysed the data in correspondence of the event: coherent Wave Burst (cWB) [21], which most notably identified for first the event, omicron-LALInferenceBurst (oLIB) [197], and BayesWave [195], which followed up on cWB trigger. All of them are meant as “excess power” detection algorithms, to identify “burst-like” signals. The first two are similar in many aspects; they both use a short duration, time-frequency representation based on some kind of *wavelet transform* to spot bins (or clusters of them) with transient excess power. Then, cWB classifies transients on the base of the shape of their time-frequency pattern; there is one class for known families of glitches [23], one for “chirp-like signals”, and one for everything else. Transients matching the first class are automatically discarded. The two algorithm then compute a measure of SNR based on the correlation between high power bins/clusters of

³³More specifically, it is the *dimensionless effective inspiral spin parameter*, χ_{eff} , the second most relevant factor to the evolution of the inspiral gravitational waveform [194], and it is the combination of the spins (and the masses) that we measure most accurately. This is the mass-weighted linear combination of the components of the black holes spins aligned with the orbital axis, $\chi_{\text{eff}} := (m_1 a_1 \cos \theta_{LS_1} + m_2 a_2 \cos \theta_{LS_2}) / (m_1 + m_2)$, and has values between -1 and $+1$. The individual dimensionless spin parameters are $\mathbf{a}_{1,2} := c \mathbf{S}_{1,2} / (G m_{1,2}^2)$. The results for O1 and O2 all favour values of χ_{eff} compatible with zero [1] (with possibly the exception of GW151226 and GW170729 [1, fig. 5]), meaning, from the previous equation, either small spins, or misaligned spins, or spins pointing towards the plane of the orbit. The part of the spin pointing to the orbital plane is expressed by the *dimensionless effective precession spin parameter*, χ_p , less constrained than the latter, as visible from [1, fig. 5] spanning almost all the allowed range of values.




Figure 1.10: Representation of some key results on GW150914. In the top figure the black hole horizons at various stages of their coalescence are shown. The middle part of the figure shows the strain reconstruction and its agreement with numerical relativity predictions. In the lower plot the evolution of the spacial separation and relative velocity of the two black holes is reported [8]. Figure readapted in fulfilment of license CC BY 3.0.

different detectors. The false alarm probability obtained for GW150914 was 2×10^{-6} for both of them.

BayesWave instead fits the detector signals with a variable number of sine-Gaussian wavelets. Then, comparing the coefficients of the fits in different detectors, it assesses the probability that what observed is a glitch (different coefficients) or the same signal (similar coefficients).³⁴ In such a way, the false alarm probability estimated for this event was 7×10^{-7} .

The reconstructed waveforms from cWB and BayesWave were then fitted with the waveform templates of EOBNR and IMRPhenom for black holes, yielding results consistent with those found with the previous template-based searches. An interesting comparison between the aforementioned five search pipelines can be found in figure 3 from [26], showing the estimated false alarm rates for the O1 events.

Unmodelled searches are more flexible than template-based ones, but also less sensitive, in general. The fluctuations of the noise in both detectors, in particular unclassified glitches, have indeed the effect to increase the false alarm probability, masking possible underlying astrophysical signals. For this reason, accurate detector characterisation and vetoing procedure are a critical aspects also for this kind of searches.

³⁴Very remarkably, in correspondence of GW170817, this technique was able to insulate the astrophysical signal in Livingston data from the part of it corresponding to a *saturation glitch* [198]; refer to figure 4.1. Further details on this will be discussed in the next section and in particular in chapter 4, when we will discuss about *deglitching* and artefact mitigation algorithms.

1.4.2 GWTC-1: an overview of current gravitational wave detections

Following GW150914, several more BBH coalescences have been observed during the first two observing runs of the Advanced detectors, plus the BNS event GW170817 [14]. The information acquired from them were collected in the masterpiece first Gravitational Wave Transient Catalogue (GWTC-1), which also contains the updated values of the O1 events parameters, and in the follow-up article reporting the inference about the population of observed stellar-mass black holes [97]. These two articles can be acknowledged as the initiation of gravitational-wave Astronomy, where the new messenger constituted by GWs is used to study the Universe.

In this conclusive section, we will briefly summarise the most important aspects presented in the two aforementioned articles.

The new black hole detections have been consistent with GW150914. One convenient and concise way of grouping them³⁵ is respect to the stellar-mass black hole population previously known from X-ray binaries, whose known most massive object is IC 10 X-1, $24 \div 33 M_{\odot}$ [199]. Reported BBH events, where both component masses are greater than the previous threshold, are GW170104 [200], GW170729, GW170809, GW170814, GW170818, GW170823, and of course GW150914. There were no evidence for the existence of such objects before the advent of GW astronomy. The Advanced Virgo detector officially joined the O2 run on the 1st of August 2017, although data from the 29th of July was also used for the reanalysis of the event GW170729, giving important contribution in the detection of the events of August 2017, and in particular to GW170814 [13] (as evident from figure 1 therein) and GW170818, with an $\text{SNR} > 4$.

On the lower mass side, there are GW151012, GW151226 and GW170608. These were also the closest events to us, which also justifies their small number with respect to the most massive ones (refer to (1.47)).

For the reasons explained in the previous section when talking about the search pipelines, having three detectors is of great importance in distinguishing coincident triggers due to noise fluctuations and actual events. Indeed, it shouldn't be surprising that 5 out of 7 black hole events detected in O2 occurred from the 29th of July to the end of the run. Also, as discussed in section 1.3.4, having three detectors constitute a major improvement for the sky localisation of the source. As visible from figure 1.11, the triple detection events are those with the smallest credible area in the sky; notice in particular GW170814, GW170817 and GW170818.

Very remarkably, on the 17th of August there was the discovery of the first BNS coalescence event, GW170817. This was the CBC with the smallest masses involved, $m_1 = 1.36 \div 1.60 M_{\odot}$ and $m_2 = 1.16 \div 1.36 M_{\odot}$, compatible with two neutron stars, and also the loudest event, in terms of SNR, ever detected. A total of ten Collaboration papers accompanied the one reporting the detection [14], with most of them signed also by the collaborations of the experiments that analysed the electromagnetic emission, and the cosmic rays. Indeed, the emission of GWs was followed by a short gamma-ray burst, GRB170817A, as observed by the Fermi-GBM and INTEGRAL telescopes 1.74 ± 0.05 sec after the merger of

³⁵Credit C.P.L. Barry.

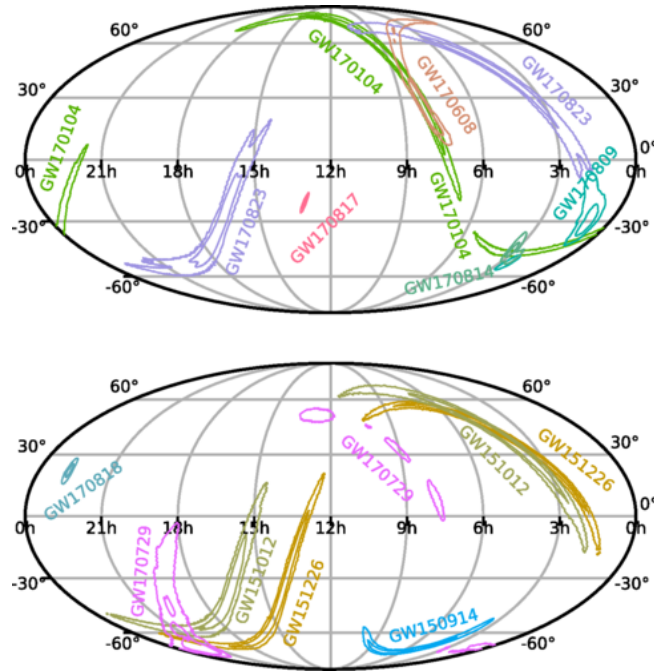


Figure 1.11: Sky localisation of the O1-O2 events, represented by the 90% credible region of the incoming GW signal, as inferred by the analysis pipelines and the triangulation method described in section 1.3.4. Figure readapted from [1], in fulfilment of license CC BY 3.0.

GW170817; this reinforcing the hypothesised of a BNS coalescence [201]. The “Multimesenger paper” [15] summarises the whole observation campaign, from GWs to the entire spectrum of the electromagnetic radiation, and even to neutrinos and cosmic ray particles (not observed), associated to this event.

From the point of view of the data, this event represented a major task for the team in charge of analysing it. Specifically, in the LIGO Livingston data, there was a glitch superimposed to the signal, few tens of milliseconds before the time of the merger. This usually comport the corresponding chunk of data to be automatically vetoed from the analysis. Refer to figure 2 from [14], also reported in chapter 4. However, the evidence of the signal in LIGO Hanford data, suggested the analysis team to throw away only the minimum amount of information from Livingston, just to remove the glitch and maintain a decent sky localisation (fundamental for the investigation of the electromagnetic counterpart). Fortunately, also Virgo contribute on this aspect, which, with its low SNR, pointed towards the region of the sky where it was less sensitive at that moment, hence improving upon the localisation obtained with the two LIGO detectors alone.³⁶

Fortunately, the glitch was short and the signal long, and this difference in the time evolution of the two allowed to clearly distinguish them and perform the subtraction of the former. This have been done by means of a model constructed out of the superposition of a certain number of sine-Gaussian wavelets [198]; with respect to this base, the components

³⁶The contribution of Virgo to the localisation of GW170817 is very effectively explained in the video realised by Leo Singer: Impact of Virgo on the localization of the binary neutron star merger GW170817.

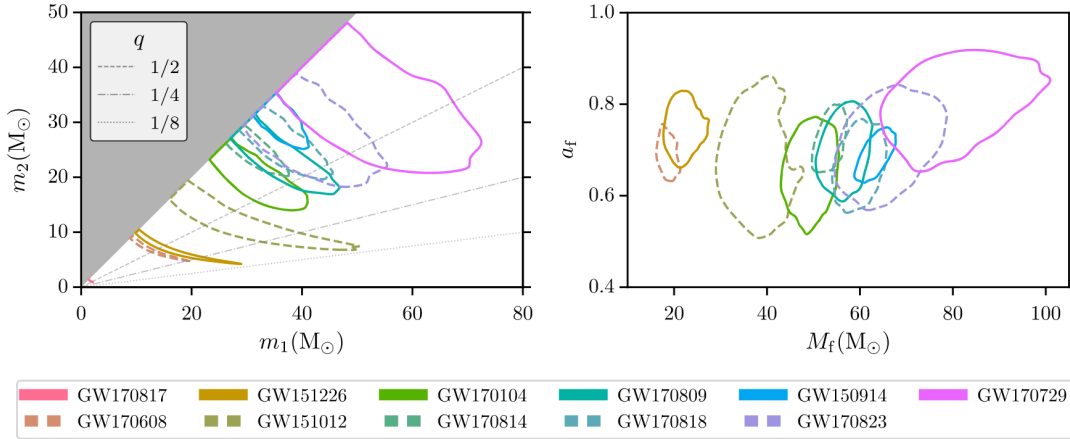


Figure 1.12: Estimated component masses (left plot) and final masses and spins for O1-O2 events. All the contours represent the 90% credible region. Figure readapted from [1], licensed CC BY 3.0.

corresponding to the impulsive glitch are very different from those of the long lasting inspiral, hence their separation has been feasible. If this signal had been a high-mass BBH, things would have been probably different.

Glitches are common features of the detector noise. Those with SNR (with respect to the “background, stationary noise”) between 5 and 8 are expected at a rate of one every ten seconds, although things may change significantly in different frequency bands and during different detector operation conditions. This means that it is very likely to find one superimposed with long BNS signal (GW170817 lasted about 100 seconds in the detector sensitivity band). Moreover, it is just a matter of time until we will find one overlapped with a BBH signal as well, for the increasing rate of events detected with the new sensitivities of the Advanced detectors.

These reasons motivated to develop some techniques to *remove* the possible glitch superimposed to a signal, without altering the properties of the latter, and affecting its parameter estimation. This will be the subject of chapter 4, and, as we will discuss, it is of fundamental importance for fast sky localisation and for sending rapid alerts for possible electromagnetic counterparts.

To conclude, we report in figure 1.12 some fundamental results from the analysis of the O1 and O2 events. On the left-hand side, the 90% credible regions for the component masses; we used the same convention of figure 1.3 that $m_1 > m_2$ (grey area). On the right-hand side of the same figure the estimated 90% credible regions for the spins of the final objects, $a_f := c\mathbf{S}_f / (GM_f^2)$, and the final masses M_f are reported. These values have been found perfectly consist from the two waveform models (the newest versions) described in the previous section.

Besides those described in this section, the release of the LIGO and Virgo data [202] allowed other research group to investigate the GW data and to perform their independent analyses. Some of them reported a few more confident BBH events, like [203] for O1 and [204] for O2.

2 Detector Noise

Data from GW detectors are dominated by *noise*. In practical terms, this is everything at the detector output that can produce an effect equivalent to a strain, in the sense of equation (1.59), but is not of astrophysical origin. For the most of the GW sources described in section 1.2, the noise is comparable or even more intense than the expected gravitational effect we aim to measure. For this reason, we don't simply "read" the GW signal at the detector output, but its "extraction" is accomplished by means of accurate modelling and statistical inference on the properties of both the noise and the GW signal itself.

This chapter, and the rest of this dissertation as well, will be focussed on detector noise. Some aspects of its characterisation have been already presented in section 1.3.3, where we introduced some figure of merit for its sensitivity and range. Moreover, the key ideas implemented in GW data analysis pipelines have been discussed in section 1.4 in relation to the events detected by Advanced LIGO and Advanced Virgo during their first two observing runs [1]. In that discussion, a central role was played by the correct evaluation of the SNR, either by means of the matched filter technique or by the unmodelled detection of excess of power with respect to the *noise background*, and the related false alarm probability. This was then used for assessing the credibility of an event and the estimates on its parameters. In this chapter we will provide further details on this, starting from detector noise description and characterisation, mostly from a statistical point of view; this have the twofold intent of constituting the statistical basis on which the detector is studied and characterised, in particular during the commissioning phase, and also serves as the workbench on which the extraction of the GW signal is performed.

In the first section we will introduce some concepts in stochastic processes theory and signal analysis, which will be used henceforth for the description and characterisation of the data. Although quite general and common to many fields of Physics, we decided to include this material in the main body of the text, instead of in an appendix, because it constitute the base and the "language" of all the material presented in this dissertation. In section 2.2.2 we will present the typical noise sources affecting a ground based interferometric detector. Part of them, the so called *fundamental noises*, like the *shot noise*, upon which we have based our estimates on the achievable sensitivity (equation (1.61)), are inherent to the detection principle described in section 1.3.1 and 1.3.2 [8, 12]. These can't be mitigated, unless by modifying the apparatus. Overall, these noises produce a main "background noise level", whose properties are assumed not to change with time (*i.e.*, they are "stationary"), as long as the detector is not altered by purpose, for example for some maintenance or calibration activities. This is represented by the curves, for the various detectors, in figures 1.6 (the O4 and O5 projections) and 1.7. Our analysis will mainly focus on the so called *technical* and *environmental* noise sources, described in the same section. These limitations come from, for

example, anthropogenic activity and weather conditions, or from systems malfunctions [22, 23]. Moreover, there is some occasional noise of *unknown origin*, whose causes and coupling mechanisms are a very active field of research for commissioners and the so called “noise hunters” [205]. Their manifestation is usually *transient*, with time scales that can vary a lot depending on the origin and the coupling mechanisms. This will be the subject of the third section, which will serve as an introduction for the third chapter.

In the last two sections, we will present the effects that both stationary and non-stationary noise may have on GW searches. We will describe the standard method for characterising the former, plus some extensions (sec. 2.3.3 and 2.3.3), and the usual “vetoing” procedure for dealing with the former.

2.1 Prelude: statistical description of data

Detector noise is modelled as a *stochastic* (or *random*) *process*. We consider a physical phenomenon random when we don’t, or we can’t, know everything about it, as for example the evolution of its status with time, or simply because we have convenience in such a modelisation.¹ Some of the sources of noise that we will consider are non-deterministic in nature, for example because originating from *quantum* processes, like the shot noise, which is the archetype of random processes. Very important is also the *thermal noise*, caused by the movement of the particles of a medium as a consequence of their temperature/energy; even if we knew the position and velocities of all of these particles, the description of their individual motion would be highly impractical, and we prefer to recover to statistical description of some “average” collective properties. Analogous considerations arise for seismic noise. Moreover, there are errors inherent to the measurement process that we are able to describe only statistically, and also their data are naturally modelled as random processes.

For the above reasons, the outcome of any experiment that we will consider, that is, the (real) quantities that we read at the output of the data acquisition systems, is a realization of a stochastic (or random) process. It can be described only statistically, by means of its *probability density function* (p.d.f.), and characterised with the values of some statistics computed from its realisations. Before digging deeper into more formal notions about stationarity, we shall clarify part of the formalism and the concepts that we will make use of in the next sections. This section is meant primarily to fix the notation and to provide a quick overview of basic concepts in stochastic process theory and time series analysis. All of the material presented here has been selected from some standard reference textbook in statistics and signal processing like [6, 206–208].

2.1.1 Probability spaces and random variables

The processes of measurement (*experiment*) of states of Nature that occur randomly (opposed to deterministically) is mathematically modelled by *probability spaces*. This mathematical construct due to Kolmogorov consists of different parts, which we shall describe shortly.

¹In the same sense, refer to the discussion in section 1.2.4 about why we have convenience into model the expected signal from a background of GW as a stochastic process as well.

The results of single execution of a certain experiment are commonly referred to as *outcomes*, ω . The (non-empty) set Ω of all the possible outcomes is called *sample space* or *observation set*. In general, it is of more practical use to deal with and characterise sets of outcomes, called *events*; these are represented by \mathcal{F} , the σ -algebra of subsets of Ω , consisting of a family of subsets closed with respect to countable union and complement with respect to Ω , representing all the possible outcome sequences and combination of them. The pair (Ω, \mathcal{F}) is called a *measurable* or *Borel space*, and it can be equipped with P , a *probability measure*, which is a way to assign to all members of \mathcal{F} a real valued number in the interval $[0, 1]$, satisfying the property of additivity with respect to \mathcal{F} , and which specifies the “likelihood” for each event to happen. Finally, the probability space, characterising a certain experiment, is the triplet (Ω, \mathcal{F}, P) .

When performing an experiment, it is assumed that the probability space corresponding to the particular system under study is fixed, and we randomly select a single outcome out of it: $\omega \in \Omega$. All the events in \mathcal{F} that contain the selected outcome ω are said to “have occurred”. This selection is done in such a way that if the experiment were to be repeated an infinite number of times, the relative frequencies of the occurrence of each of the events would correspond to the probabilities prescribed by P (*frequentist interpretation of P*).

As a classical example, for a die roll $\Omega = \{1, 2, 3, 4, 5, 6\}$, \mathcal{F} is composed by the 2^6 possible subsets (*power set*) of Ω , and for an event $A \in \mathcal{F}$, e.g. $A = \{2, 4, 6\}$, $P(A) = |A|/6$, where $|A|$ is the *cardinality* (number of elements) of the set A , in this case 3, and $P(A) = 1/2$. For $\Omega = \mathbb{R}^n$, we can define the σ -algebra \mathcal{F} generated by the open subsets of \mathbb{R}^n (or, equivalently, by the open balls), which is called the *Borel σ -algebra* and is denoted by $\mathcal{B}(\mathbb{R}^n)$ or just \mathcal{B}^n [209, §1]. Similarly, we can define a Borel algebra for any closed subsets of \mathbb{R}^n .

A *random variable* is, loosely speaking, a quantity whose values depend on the random outcomes of an experiment, which allows for probabilities to be assigned to sets of its potential values. More precisely, given a probability space containing the observation set, (Ω, \mathcal{F}, P) , and a second measurable spaces (S, Σ) , a random variable X is an S -valued measurable function $X: \Omega \rightarrow S$ such that the event

$$\{\omega \in \Omega : X(\omega) \in A\}$$

belongs to \mathcal{F} for arbitrary $A \in \Sigma$. Often the values assumed by this variable are real, $S = \mathbb{R}$ equipped with its Borel σ -algebra $\Sigma = \mathcal{B}$, and the previous definition can be replaced with

$$\{\omega \in \Omega : X(\omega) \leq x\} \in \mathcal{F} \quad \text{for all } x \in \mathbb{R}.$$

In this case, we can define the *cumulative distribution function* P_X of the variable X as the function $P_X: \mathbb{R} \rightarrow [0, 1]$ for which

$$P_X(x) := P(X \leq x)$$

where $P(X \leq x)$ denotes the probability of the event $\{\omega \in \Omega : X(\omega) \leq x\}$. Similarly, we can define the *joint* cumulative distribution function of two (or more) random variables X and Y

defined on the same probability space (Ω, \mathcal{F}, P) to $(\mathbb{R}, \mathcal{B})$ as:

$$P_{X,Y}(x, y) := P(X \leq x, Y \leq y).$$

The variables are said to be *independent* if the events $\{\omega \in \Omega : X(\omega) \leq x\}$ and $\{\omega \in \Omega : Y(\omega) \leq y\}$ are such that, for all $x, y \in \mathbb{R}$, their *joint* cumulative distribution function satisfies

$$P_{X,Y}(x, y) := P(X \leq x, Y \leq y) = P(X \leq x) \cdot P(Y \leq y) := P_X(x) \cdot P_Y(y).$$

A *continuous* random variable X is defined as one with a continuous cumulative distribution function $P_X(x)$. If the derivative of the latter quantity exists everywhere we can define the p.d.f. of the variable X as:

$$p_X(x) := \frac{dP_X}{dx}, \quad (2.1a)$$

or implicitly as:

$$\int_{-\infty}^x p_X(x') dx' = P_X(x). \quad (2.1b)$$

As P_X is a non-decreasing function on \mathbb{R} , p_X is going to be a positive definite, real function. Similar definitions hold for the joint p.d.f. of two or more variables; if X and Y are continuous, and the mixed derivatives of $P_{X,Y}(x, y)$ exist, we can view them together as a *random vector* from Ω to \mathbb{R}^2 , and define the *joint probability distribution function*:

$$p_{X,Y}(x, y) := \frac{\partial^2 P_{X,Y}}{\partial x \partial y}(x, y), \quad P_{X,Y}(x, y) = \int_{-\infty}^x \int_{-\infty}^y p_{X,Y}(x', y') dx' dy'.$$

We can extend the above definition to random vectors of arbitrary finite dimensions.

We can use the distribution of a random variable to compute *expectation values*. For a continuous random variable X with p.d.f. p_x , its expectation value is defined as

$$E[X] := \int_{-\infty}^{+\infty} x p_X(x) dx = \int_{-\infty}^{+\infty} x dP_X(x).$$

The latter can be generalized to the variable $Y := f(X)$:

$$E[Y] = E[f(X)] = \int_{-\infty}^{+\infty} f(x) p_X(x) dx,$$

and in particular we can define the *n-th moment* of X about a value c as

$$\mu_n(c) := E[X^n] = \int_{-\infty}^{+\infty} (x - c)^n p_X(x) dx. \quad (2.2)$$

Particularly important are the *mean*, which is the first moment about the centre, $\mu := E[X]$, and the *variance*, which is the 2nd order moment about the mean, $\sigma^2 := \text{Var}[X] = E[(X - \mu)^2]$.

In the following, as it is common habit in GW literature, we will also adopt the ‘‘Quantum

Mechanics" notation of the angle brackets " $\langle \dots \rangle$ " to represent the expectation values, e.g. $\langle X \rangle \equiv E[X]$.

2.1.2 Stochastic processes and time series

If T is an ordered set, called the *index set*, a *stochastic* or *random process* is a family of (S -valued) random variables $\mathbf{X} := \{X_t(\omega), t \in T\}$ all defined on the same probability space (Ω, \mathcal{F}, P) to the measurable space (S, Σ) . This can be viewed as a function of both $t \in T$ and $\omega \in \Omega$, that is $X_t(\omega) \equiv X(t, \omega)$; in this sense, for any fixed sample point $\omega \in \Omega$, the function $X(\cdot, \omega): T \rightarrow S$ is called a *realization* (or *trajectory*) of the process \mathbf{X} , and the records $x(t)$ are the (real) observed values that it has assumed out of a whole collection of possible values, the "ensemble". Very often the variable $t \in T$ has the meaning of time, and consequently $x(t)$ is called a *time series*.² A very large amount of data from gravitational-wave detectors are represented in first instance as time series. For the reason that we will clarify in a moment, from now on, we will omit the dependency on ω and focus on that on the index variable t .

For a finite set $T = \{t_1, t_2, \dots, t_n\}$, $n \in \mathbb{N}$, the stochastic process is called a *random sequence* and its values are often denoted with the subscript corresponding to the time index: $x_i := x(t_i)$. Intuitively, this may be seen as a projection of the process onto a finite dimensional vector space.

For such a process, the *joint n -dimensional cumulative distribution function* is defined by:

$$P_{X_{t_1}(\omega), \dots, X_{t_n}(\omega)}(x_{t_1}, \dots, x_{t_n}) = P_{t_1, \dots, t_n}(\mathbf{x}) = P_{\mathbf{X}}(\mathbf{x}) := P(X_{t_i}(\omega) \leq x_i, i = 1, \dots, n). \quad (2.3)$$

The finite dimensional distribution of a stochastic process is then defined to be the set of all such joint distribution functions for all finite integer sets T of any size n .

Stationary processes

The previous definitions were necessary to introduce one of the leading concept of the analysis contained in this work, that is the *stationarity* of a certain stochastic process.

Intuitively, a stationary stochastic process is one whose statistical properties do not change over time (the index set). It is important to stress from now that this is a property of the process, and not of any finite (or infinite) realization of it, nor of the values assumed by this. More formally, stationarity (often called *strong* or *strict sense* stationarity) is a mathematical property that a stochastic process has when *all* of its finite-dimensional joint cumulative distribution functions (2.3) are shift-invariant along the index axis (usually, the time), that is if the set of points $\{t_1, \dots, t_n\} \in T$ is shifted by a constant amount τ (usually integer or real-valued):

$$P_{t_1+\tau, \dots, t_n+\tau}(\mathbf{x}) \equiv P_{t_1, \dots, t_n}(\mathbf{x}) \quad (2.4)$$

²Another frequently encountered meaning for the index set is that of *frequency*, and we will describe several example of this kind in what follows. Also, in many applications the index set is that of the spatial coordinate(s), as for example in a photographic image. We will make use of the latter interpretation in section 3.6.2, when we will face the problem of attributing the statistical significance of certain quantities in bidimensional time-frequency domain.

for any $n \in \mathbb{N}$, and $t_1, \dots, t_n, t_1 + \tau, \dots, t_n + \tau \in T$. From the previous definition trivially follows that all the sets of *i.i.d.* variables are stationary processes. This is the strongest and the most common definition of stationarity, though of little practical use since we don't have access to the full (*i.e.* arbitrary n) joint cumulative distribution of the process. We will discuss shortly other *weaker* definitions of stationarity.

Expectation values and distribution moments

To characterise the statistical properties of stochastic processes it is customary to consider their "simplest" numerical characteristics, embodied in their lowest order distribution moments; the following equation generalises (2.2) to the case of a random sequence:

$$\begin{aligned} \mu_{m_1, \dots, m_n} &:= E [X(t_1)^{m_1} \dots X(t_n)^{m_n}] \\ &= \int_{-\infty}^{+\infty} \dots \int_{-\infty}^{+\infty} x_1^{m_1} \dots x_n^{m_n} dP_{t_1, \dots, t_n}(x_1, \dots, x_n). \end{aligned} \quad (2.5)$$

The previous integral should be meant in the sense of *Riemann-Stieltjes* [210], accounting for the possibility of non-continuous variables X , *e.g.* with a discrete distribution, for which a probability density function (ref. to (2.6)) can't be defined, and for which the previous integral is not well defined in the usual sense of Riemann [5, §3.2].³

If the cumulative distribution function $P_{t_1, \dots, t_n}(x_1, \dots, x_n)$ is differentiable, its joint probability density function can be defined, analogously to (2.1), as:

$$p_{t_1, \dots, t_n}(x_1, \dots, x_n) := \frac{\partial^n P_{t_1, \dots, t_n}(x_1, \dots, x_n)}{\partial x_1 \dots \partial x_n}. \quad (2.6)$$

This can also be defined implicitly, as in (2.1b). Equation (2.7) becomes then an ordinary Riemann integral:

$$\begin{aligned} \mu_{m_1, \dots, m_n} &:= E [X(t_1)^{m_1} \dots X(t_n)^{m_n}] \\ &= \int_{-\infty}^{+\infty} \dots \int_{-\infty}^{+\infty} x_1^{m_1} \dots x_n^{m_n} p_{t_1, \dots, t_n}(x_1, \dots, x_n) dx_1 \dots dx_n. \end{aligned} \quad (2.7)$$

The most important distribution moments are the first- and the second-order ones, μ_1 and $\mu_{1,1}$, that is the *mean* and the *auto-correlation* functions:

$$\mu(t) := E[X(t)] \quad (2.8a)$$

and:

$$C(t, s) := E[X(t) X(s)] \quad (2.8b)$$

³More precisely, the Stieltjes integral generalises that of Riemann realizing the notion of integrating an "integrand" function with respect to another "integrating" one. It can be shown that one sufficient condition for this integral to exist is that the integrand is continuous and the integrating function is of bounded variation within the integrating interval, as it is in general the case for equation (2.7). This integral is linear in both the integrand and integrating functions, but differently from the Riemann integral is not additive.

with $t, s \in T$. Moreover, from the previous definitions, another useful quantity is the *auto-covariance* function:

$$\begin{aligned} \text{Cov}[X(t), X(s)] = K(t, s) &:= E[(X(t) - \mu(t))(X(s) - \mu(s))] = \\ &= C(t, s) - \mu(t)\mu(s) \end{aligned} \quad (2.9)$$

which equals the ordinary *variance*, $\text{Var}[X(t)] = \sigma_X^2$, when $t = s$.

From (2.4), if the process $\{X(t), t \in T\}$ is stationary, its mean value is constant:

$$E[X(t)] = \mu = \text{const.} \quad (2.10)$$

and its auto-correlation function $C(t, s)$ does not depend on the “epoch”, that is the global values of t or s , but only on their difference $\tau = t - s$:

$$C(t, s) = E[X(t)X(s)] = C(t - s, 0) := R(t - s) = R(\tau). \quad (2.11)$$

A normalized version of the previous quantity, with respect to lag $\tau = 0$, is called the *correlation coefficient* between $X(t)$ and $X(t + \tau)$:

$$\rho(\tau) = \frac{R(\tau)}{R(0)}. \quad (2.12)$$

This quantity measures the “similarity” between a realisation of $X(t)$ and an index shifted version of the same realisation. Intuitively, as $|\tau|$ increases we would expect the “memory” between $X(t)$ and $X(t + \tau)$ to decrease. A complete list of properties of $R(\tau)$, and their proofs, can be found for example in [6, §3.3]. For what concerns us, the most important thing to notice is that, from the definition, we can show that $R(t_i - t_j)$ is a positive semi-definite matrix in the indices i, j .

Then, a *weaker* condition of stationarity, often called *wide-sense stationarity* (or *weak-sense stationarity*), only requires the shift-invariance of the main statistical properties of the process, usually the first and second-order moments (2.8) (*second-order* or *covariance stationarity*). Strong stationarity does not imply weak stationarity, and vice versa.⁴

Gaussian processes

As we will comment in detail in the next sections, another very important property for data analysis purposes is *Gaussianity*. A stochastic process $\{X(t), t \in T\}$ is said to be *Gaussian* if its joint p.d.f. is a multivariate Gaussian for any $t_1, \dots, t_n \in T$:

$$p_{t_1, \dots, t_n}(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^n \det \mathbf{K}}} \exp \left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{K}^{-1}(\mathbf{x} - \boldsymbol{\mu}) \right] \quad (2.13)$$

⁴The reason why *strong* $\not\Rightarrow$ *weak* is quite intuitive, while the classical counterexample [211] for the opposite relation, *strong* \Rightarrow *weak*, is that of a process with a *i.i.d.* Cauchy distribution, which has no finite mean and second-order distribution moments, though its joint cumulative distribution function is always the same, and then it is strictly stationary. Extensions to the definition of weak-sense stationarity to this class of processes can be found in [212].

where the bold right-hand side quantities are expressed in vectorial and matricial form with respect to the index set $(\mathbf{x} := (x_1, \dots, x_n)^T, \mathbf{K}_{ij} := K(t_i, t_j))$, and to be well define it assumes the auto-covariance matrix to be non-degenerate: $\det \mathbf{K} \neq 0$.

As evident from (2.13), a key fact of Gaussian processes is that they can be completely characterised by their first second-order moments, and weak stationarity does indeed imply strong stationarity.

Gaussian random processes are very often encountered in physical contexts, and often may be mathematically predicted by the multidimensional *central limit theorem*. Also, it can be shown that if a Gaussian process undergoes a linear transformation, then the output will still be a Gaussian process. Moreover, a linear transformation of a weakly (strongly) stationary random process will form a weakly (strongly) stationary random process as well [213, §5.3].

In the remaining of this chapter we will motivate how GW detector noise can be considered, at least in first approximation, as stationary and Gaussian, and on the base of this we will describe the basic principles of the data analysis techniques used to study GWs. We will also clarify the limits of validity of the previous approximations, and discuss in which way they may influence the analysis.

Spectral representation of stationary processes

First of all, let us assume without much loss of generality that the stationary processes under consideration have zero mean, $E[X(t)] = 0$, as the general case can be easily recovered adding at the end the (known) constant mean. Two important theorems shall follow and will be useful for the characterisation of stationary stochastic processes.

One of the results of the renowned *Wiener-Khinchin theorem* [214] proves that a *necessary and sufficient* condition for $\rho(\tau)$ to be the correlation coefficient (2.12) of a continuous stochastic process⁵ $\{X(t), t \in T\}$ is that there exists a function $F(f)$, having the properties of a cumulative distribution function on $(-\infty, +\infty)$ (that is $F(-\infty) = 0$, $F(+\infty) = 1$, and $F(f)$ non decreasing) such that, for all τ , $\rho(\tau)$ may be expressed by the *Fourier-Stieltjes transform*:

$$\rho(\tau) = \int_{-\infty}^{+\infty} e^{2\pi i f \tau} dF(f). \quad (2.14)$$

Of course, the previous integral should be interpreted in the Riemann-Stieltjes sense. This result tells us that *any* stationary stochastic process can be characterised by an (infinite) sum of sine and cosine functions of frequency f , which can be interpreted as the *spectrum* of the signal, weighted by the amplitudes $dF(f)$; the integrating function $F(f)$ is therefore called the process *spectral distribution function*.

This provides a spectral representation of the auto-correlation function of a stationary stochastic process, without it to satisfy the usual assumptions for the Fourier integral or series, that is absolute integrability ($\int_{-\infty}^{+\infty} R(\tau) d\tau < \infty$). Indeed, the necessary part of the

⁵The analogue for discrete parameter, stochastic sequences of the Wiener-Khinchin theorem is due to Wold [215]. Assuming the elements of the index set to be equally spaced by $1/f_s$, with f_s the *sampling frequency* in the case of T representing time, the discrete index autocorrelation function $R(l/f_s) := R_l$ admits the *inverse Discrete Time Fourier-Stieltjes transform*: $R_l = \int_{-f_s/2}^{+f_s/2} e^{2\pi i f l / f_s} dP(f)$. Now the spectral function $P(f)$ is defined only for frequencies in the "Nyquist range" $-f_s/2 \leq f \leq +f_s/2$ [6, §4.8.3]. Detectors data are sampled at discrete times, so the latter is the form we will use in practice. However, it is more convenient for now to deal with continuous parameter processes.

Wiener-Khinchin theorem directly follows from the positive semi-definite property of $\rho(\tau)$, and the *Bochner's theorem* [216, §1.3], which ensures that any continuous function of such a kind have a Fourier-Stieltjes transform of the form (2.14). The continuity of $\rho(\tau)$ follows directly from the assumption that $X(t)$ is stochastically continuous [6, §4.8].

A similar definition (except for the normalisation) holds for the auto-correlation function:

$$R(\tau) = \int_{-\infty}^{+\infty} e^{2\pi i f \tau} dP(f), \quad (2.15)$$

with $P(f) = F(f) \cdot \sigma_X^2$ and $P(-\infty) = 0$, $P(+\infty) = \sigma_X^2$.

In the case when $P(f)$ (or $F(f)$) is differentiable everywhere, we have a purely continuous spectrum and we can rewrite the previous integral as a standard Riemann one, introducing the function $S(f)$,

$$dP(f) := S(f) df \quad (2.16)$$

called the *two-sided Power Spectral Density* (PSD) of the stationary stochastic process $\{X(t), t \in T\}$. Then equations (2.14, 2.15) can be immediately inverted giving an explicit expression for the power spectral density:

$$S(f) = \int_{-\infty}^{+\infty} e^{-2\pi i f \tau} R(\tau) d\tau. \quad (2.17)$$

The previous result, that any stationary stochastic process can be represented “as a Fourier(-Stieltjes) integral”, is one of the most important in the theory of stationary processes. It provides a “canonical” representation of them, with a physical and mathematical well-defined notion of power spectra. Moreover, if the process is also Gaussian, since $R(\tau)$ provides a complete description of its statistical properties so does $S(f)$.

In the context of GW detectors, the noise (hence the sensitivity, as for the discussion in section 1.3.3) is given, with some exceptions to be discussed in the next chapter, approximately by a superposition of stationary Gaussian processes, with very characteristic spectral signatures. It is therefore convenient to use this spectral representation to describe it, as we did with the familiar sensitivity curves shown in figures 1.6 and 1.7.

Cramèr representation and physical interpretation of the PSD

It can be shown that any continuous index, stationary stochastic process admit a *spectral* or *Cramèr representation* in terms of the form:⁶

$$X(t) = \int_{-\infty}^{+\infty} e^{2\pi i t f} dZ(f) \quad (2.18)$$

where $Z(f)$ is a stochastic process with $E[dZ(f)] = 0$ and *orthogonal* increments:

$$\text{Cov}(dZ(f), dZ(f')) = E[dZ^*(f) dZ(f')] = \delta(f - f') dP(f) \quad (2.19)$$

⁶The analogue to (2.18) for discrete index, stochastic process has the integration interval defined between $-f_s/2$ and $+f_s/2$ [6, §4.11]; refer also to note 5.

with $dP(f)$ the (non-normalised) spectral function introduced in (2.15). The previous *Fourier-Stieltjes* stochastic integral is a generalisation of the Fourier integral, where $Z(f)$ is itself a stochastic variable, and is therefore defined in the mean-square sense [6, §3.6]:

$$E \left[\left| X(t) - \int_{-\infty}^{+\infty} e^{2\pi i f t} dZ(f) \right|^2 \right] = 0.$$

Taking the complex conjugate of the representation in (2.18), evaluated at index $t + \tau$, from the expectation value of

$$E \left[X^{(*)}(t) X(t + \tau) \right] = \iint_{-\infty}^{+\infty} e^{2\pi i (f' - f)t} e^{2\pi i f' \tau} E \left[dZ^{*}(f) dZ(f') \right]$$

one can immediately recover the Wiener-Khinchin theorem (2.15).

Equation (2.18) provides the physical interpretation of $dP(f)$, that is the mean-square amplitude of the component in $X(t)$ with frequency f . More precisely, when $X(t)$ represents some physical process it may be shown that $P(+\infty) = \sigma_X^2$ is a measure of the average total power dissipated by the process, and $dP(f) (= S(f)df)$ represents the contribution to the total power from the components in $X(t)$ with frequency between $(f, f + df)$. The function $S(f)$ thus represents the distribution of the power density over frequency, thus the name *power spectral density* function.

Ergodic theorem

In many practical situations, we have no or little prior theoretical knowledge on the “underlying” statistical structure of the process under study, represented by its joint cumulative distribution function (2.3). A central problem in *statistical inference* is the extraction of this information from just one (or a finite number) of its realisations. In this respect, the practical value of stationary processes is that, under some further conditions very often satisfied in practice [5, §3.2], the expectation values (2.7), obtained by averaging over the whole space of experimental outcomes $\omega \in \Omega$, can “usually” be replaced by time averages of the same quantities using just one, “sufficiently long”, realisation of it. This is a consequence of the *ergodic theorem* (or *law of large numbers*). In the case of the lowest order moments, the mean and the auto-covariance function (2.8), we can define, with a suitable choice of the limits:

$$\mu(\omega) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T X(t, \omega) dt, \quad (2.20a)$$

$$R(\tau, \omega) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T X(t, \omega) X(t + \tau, \omega) dt \quad (2.20b)$$

and similarly for discrete index processes, with summations substituted to integrals; in this case, the existence of the first limit is provided by the *Birkhoff ergodic theorem* [217, §11.1]. In general, these quantities are themselves random variables, functions of the particular realisation chosen, $X(\cdot, \omega)$ (while the time dependency has been averaged away). If it turns out that they equal the values of the corresponding ensemble-averaged ones (2.8), independently on ω , then the process $\{X(t, \omega), t \in T\}$ is said to be *weakly ergodic*.

As for stationarity, there is a distinction between a “weak” version of the ergodic theorem, limited to moments up to a certain order, and a “strong” version of it, where all the ensemble-averaged statistical properties are deducible from the corresponding time averages. Thus, strong ergodicity implies weak ergodicity, but not conversely. Instead, if a Gaussian process is ergodic it is of course in both senses.

The idea at the base of this theorem is that any sufficiently large collection of random samples, regardless of what the individual samples are, must represent the average statistical properties of the entire process. Each sample function $X(\cdot, \omega)$ must then be representative of all the others in the sense described above, and it does not matter which particular realisation is used in the time-averaging calculations.

Of course, stationarity (to a certain level) is a necessary condition for an arbitrary random process to be ergodic, otherwise the previous limiting procedure will be time-dependent. Anyway, this is not also sufficient.

For the convergence of the sample (time-averaged) mean, in the limit $T \rightarrow \infty$, to the ensemble mean, it can be shown that the sufficient condition is that the integrated spectrum $P(f)$ is continuous at $f = 0$, that is it has “no jumps” at that frequency. Jumps in $P(f)$ at non-zero frequencies don’t cause the same difficulty. An alternative form of the same result is given by the *Slutsky’s theorem* [208, §12.1], which asserts that a weakly stationary stochastic process is mean ergodic if:

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T R(\tau) d\tau = 0. \quad (2.21)$$

The sufficient condition for convergence of the covariance is provided, for stationary Gaussian process, from the condition that $P(f)$ is continuous everywhere, that is, no delta functions appear in the auto-spectra corresponding to infinite spectral densities at discrete frequencies. Related to this, another class of process that can be said in advance to be covariance ergodic is the class of stationary *Gaussian Markov processes* [218]; a Markov process is a discrete process whose relationship to the past does not extend beyond the immediately preceding observation. The auto-correlation function of a stationary Gaussian Markov process may be shown to be of a simple exponential form. We will make use of a Markov process model in section 3.4 when describing an algorithm to track “drifting lines”.

Further details and the proof of the previous conditions for a process to be mean- or covariance-ergodic can be found in [217, §11.1]. A thorough discussion about the mathematical conditions under which processes have the ergodic property are presented in [219, §11].

The results discussed here about the conditions under which it is possible, from a single realisation, to estimate the mean and the covariance of process, are at the base of the most of the data analysis, in particular to what will be discussed in sections 2.1.4 and 2.1.6.

2.1.3 Some “finer-grain” notions of non-stationarity

Theory of stationary stochastic processes is well developed and is usually the basis for characterising the detector noise in most of the methods adopted in GW analyses. On the other hand, the assumption of stationarity is too restrictive in many practical situations, where

most measured signals exhibit significant changes over time, as it will be discussed in section 3.1, and data may require different analysis procedures and/or interpretations from those appropriate for stationary Gaussian random processes.

An appropriate general methodology does not exist for analysing the properties of all types of non-stationary random data from individual sample records. This is due partly to the fact that a non-stationary conclusion is a negative statement specifying only a lack of a property, rather than a positive statement defining the precise nature of the non-stationarity. It follows that special techniques must be developed for non-stationary data that apply only to limited classes of these data [213, §11].

Different approaches have been considered in literature to take into account time-varying behaviours. For example, *segmentation techniques* into stationary frames are common GW literature [220–222], where often *piece-wise stationary models*, such as time-varying *Auto Regressive* (AR) and *Auto Regressive, Moving Average* (ARMA) models, are used; it is often possible to force the data to be at least piecewise stationary for measurement and analysis purposes [213, §11]. An alternative approach revolves around classes of processes with desirable properties extending the stationary case; examples of these are the *Wiener* (also called *Brownian*) *processes*, and those allowing a *generalised Cramèr representation*. Refer to [5, §3.7] for an overview of these techniques in the context of GW data analysis. Moreover, the concept of a time-varying spectral representation of non-stationary processes were initially approached in [223], where the author developed the concept of “evolutionary spectrum”, a time dependent spectral function. A statistical test based on this concept will be discussed in section 3.1.1. Related to segmentation techniques and piece-wise stationary processes, of particular importance for many non-stationary noise manifestation in GW detectors is the class of *Locally Stationary Processes* (LSP) introduced by Silverman in [224].

In this section some extensions to the concept of non-stationary process, relevant for the analysis that shall be presented, will be discussed. In particular we will focus on LSPs and non-stationary process which can be model out of some stationary ones.

Locally stationary stochastic processes

The intuitive idea behind LSPs is that their statistical properties “change slowly” over time or, alternatively, that locally, at each time point, they are close to a stationary process but whose characteristics (covariances, parameters, etc.) are gradually changing as time evolves (refer to [225] or [226] for a complete review on the subject). The theory of LSPs is based on the principle that a non-stationary process can be locally approximated by a stationary one if the time variation of the model parameters is sufficiently smooth. This behaviour can be modelled qualitatively with a covariance equal to a stationary one multiplied by a time-dependent “sliding power factor”, which renormalizes the average instantaneous power to a representative local level. Besides being a natural generalization of stationary processes, another appealing feature of this definition is that it avoids time-varying parameters, as instead it is the case for the class of processes that will be described shortly. They have provided a sound statistical methodology for modelling data exhibiting non-stationary features without resorting to data transformations, trend removals and other related techniques.

In the context of GW detector data, we have experiments with finite durations. Moreover, we have some prior knowledge about a process containing the information of interest (the *signal*) with some characteristic time scale/duration, to which is superimposed (*additive noise model* assumption) a second process whose knowledge is not of interest for us (the *noise*). In both of this case we are usually not concerned with the possibility that stochastic process under study is not stationary “outside the observation window” or, more formally, with time scales longer than that of interest for our experiment and for the signal.

Also, in the context of noise studies, we have examples of noises that vary slowly with respect the characteristic times of the frequencies in the detector sensitivity band. For characterising this region of the spectrum (refer to its definition in section 2.1.2) we can assume that the underlying process is stationary. This is at the base of some of the *coherence analysis* described in section 2.3, used to study (quasi-)stationary noise features.

Stationary models and parametric notion of stationarity

The definitions of stationarity presented so far have been non-parametric, *i.e.*, it did not assume a model for the so called underlying *data-generating process* (DGP), and thus apply to any stochastic process. In this respect, *linear models* are very important because, as previously stated, if the underlying DGP is stationary and Gaussian, also the resulting data will be Gaussian, and under a suitable transformation also the stationary behaviour can be recovered. So, *linearity* can be regarded as the third most important property for stochastic processes, after stationarity and Gaussianity.

In general stochastic process modelling can have many forms and represent different processes. The simplest example of a non-stationary process originating from a stationary model is that of a *linear trend*:

$$Y(t) = a \cdot t + X(t), \quad (2.22)$$

where $X(t)$ is a stationary process. A generalised version of the previous trend to an arbitrary deterministic function $a(t)$ is:

$$Y(t) = a(t) + X(t).$$

It is immediate to check that the previous models have the effect to change the mean of the process $X(t)$ as a time dependent function. If this is the case, and if the function $a(t)$ is “sufficiently simple/common”, under a suitable definition of sample averages, one can try to reconstruct its functional form, and test the stationarity of the residual difference between this estimate and the observed process $Y(t)$.

Other more complex models include variance varying models (heteroscedastic), like:

$$Y(t) = a(t) \cdot X(t)$$

and frequency varying models, like

$$Y(t) = X(a(t))$$

Further details in [213, §11].

The previous models do not assume any form of “memory” (or just a deterministic one). Two of the most commonly used models in literature, that include memory information, are *Auto Regressive* (AR) and *Moving Average* (MA) models. In particular, we will make use of the former in section 3.7.1 when studying models to infer the *causality* relation between two processes. For a discrete index process, X_n , an $AR(p)$ model of order p can be expressed as

$$X_n = \varphi_0 + \sum_{i=1}^p \varphi_i X_{n-i} + \varepsilon_n \quad (2.23)$$

where the coefficients φ_n are some (constant) parameters measuring the influence of past values of X_n , and ε_n is an uncorrelated (*i.e. white*, in order to condense all the correlation structure in the coefficients φ_n) stationary process (often assumed to be also Gaussian, that is least informative among all the distribution with finite variance). The previous model is called *VAR* if the process X_n is instead *vector* valued. In terms of the *lag-operator*, $LX_n = X_{n-1}$, the previous model can be rewritten as:

$$(1 - L\varphi_1 - \dots - L^p\varphi_p) X_n := \Phi[L] X_n = \varepsilon_n.$$

The part inside the parenthesis on the left is called the *characteristic equation* of the model. We can find the roots of this equation, $\Phi[L] = 0$; it can be shown that the resulting process is weakly stationary if all of its roots lie inside the unit circle [227]. If 1 is a root of the equation then the stochastic process is said to be a difference stationary process, or unit root process, or integrated of order 1. This means that the process can be transformed into a weakly-stationary process by applying to it a certain type of transformation, called differencing. Unit root processes, and difference stationary processes generally, are interesting because they are non-stationary processes that can be easily transformed into weakly stationary processes. The simplest example for such a process is the following AR model:

$$X_n = X_{n-1} + \varepsilon_n.$$

This is also true if the multiplicity of the root 1 is greater, by applying a corresponding number of differentiations.

An $MA(q)$ process is mathematically defined as:

$$X_n = \varepsilon_n + \sum_{i=1}^q \vartheta_i \varepsilon_{n-i} \quad (2.24)$$

Contrary to the AR model, the finite MA models are always stationary. It can be shown that any $AR(p)$ model is *invertible*, that is it can always be rewritten in terms of an $MA(\infty)$ process, whereas for an $MA(q)$ process to be invertible, all the roots of its characteristic equation must lie outside the unit circle.

Stationary models will be exploited in chapter 3 for the characterisation of non-stationary noise. In particular, they will be adopted in section 3.4 for modelling the frequency evolution of “drifting lines”, in section 3.5.3 when studying *regression analysis* to characterise the correlations between the detector strain signal and other sensors data, and in section 3.7.1 for modelling the causal relation between two processes.

2.1.4 Signal detection and the matched filter technique

When searching for a GW signal, we usually adopt the *additive noise* model, assuming that our detectors data, $x(t)$, is represented by the superposition of the strain equivalent noise, $n(t)$, and the actual (unknown) gravitational signal, $s(t)$:⁷

$$x(t) = n(t) + s(t) \quad (2.25)$$

where the latter can be both a deterministic signal, as in the case of a CBC or a continuous-wave signal, or a stochastic one, like $n(t)$, as in the case of the SGWB. Moreover, it can be a completely unmodelled signal, as in burst searches, or one for which templates exist and are provided by IMRPhenom or SEOBNR [189, 190], for example. Although a wide variety of pipelines exists for investigating the previous classes of signals, all have fundamentally in common the assessment of whether the detector output is statistically different from our model for the noise. An accurate knowledge of it is therefore fundamental for all the GW data analysis searches, and motivates the importance of the Detector Characterisation activities.

For simplicity, in this section we will consider modelled searches only, and discuss the problem of finding the model $h(t)$ that best “match” the *unknown* signal $s(t)$. In CBC searches, most of the pipelines use the predictions of GR to construct the template $h(t)$, as it is done by the already mentioned IMRPhenom and SEOBNR waveform models. Then, $h(t)$ is considered a good model if the residuals of the subtraction to the data, $x(t) - h(t)$, is consistent with our model for the instrument noise, $n(t)$. More formally, the *likelihood* that the data $x(t)$ contains a possible signal $s(t)$, modelled by $h(t)$, is given by the *conditional probability* that $x(t) - h(t)$ is a realisation of the noise model $n(t)$ [2]. Using the vectorial “bold-face” notation for the data records in all the detectors,

$$\mathbf{x} := \{x_{nI} : n = 0, \dots, N - 1 \text{ and } I = H, L, V, \dots\} \quad (2.26)$$

and similarly for \mathbf{n} and \mathbf{h} , from (2.13), for Gaussian noise we have:

$$p(\mathbf{x}|\mathbf{h}) = \frac{1}{\sqrt{\det(2\pi\mathbf{K})}} \exp \left[-\frac{1}{2}(\mathbf{x} - \mathbf{h})^T \mathbf{K}^{-1}(\mathbf{x} - \mathbf{h}) \right] \quad (2.27)$$

where now \mathbf{K} is the noise covariance matrix in all the detectors, $I = H, L, V, \dots$ (Hanford, Livingston, Virgo, ...). Explicitly,

$$(\mathbf{x} - \mathbf{h})^T \mathbf{K}^{-1}(\mathbf{x} - \mathbf{h}) = \sum_{n,m} \sum_{I,J} (x_{nI} - h_{nI}) K_{nI,mJ}^{-1} (x_{mJ} - h_{mJ}). \quad (2.28)$$

If the noise is uncorrelated between the detectors, the covariance matrix is diagonal in the corresponding indices: $K_{nI,mJ} = \delta_{IJ} K_{nm}^I$. Currently, the only known potential correlated

⁷Notice from equation (2.25) two aspects related to the notation. From this section on, we will move the focus from the properties of the data, thought as a stochastic process, to what we can expect for their realisations. Hence, as stated in the Notation section, we will indicate it with lower-case letters, as customary in GW literature. Also, we adopted a different notation for the GW signal in this equation, labelled by $s(t)$, with respect to that in the previous chapter, $h(t)$. This is justified by their different meaning. In GW data analysis the actual signal is never known; what we aim to do is to provide the representation, $h(t)$, based on GR and the arguments provided in section 1.2 for example, that better describes it. This is indeed the main subject of this section.

noise sources, even for detectors several thousands kilometres apart, are *Schumann resonances*, that is low-frequency magnetic field excitations transmitting between the Earth's surface and the ionosphere [228, 229]. Their effects are well below today's detector sensitivities [138], but in the not so distant future they may influence in particular SGWB searches [230]. For this reason, they are monitored by magnetometers installed at each detector site, and some subtraction techniques have been developed, in order to mitigate their effects [231].

In addition to the above assumptions, if the noise is stationary, the covariance matrix will depend only on the relative difference between the two times, meaning that it is diagonal in the frequency domain: $K_{mn}^I = K^I(t_m, t_n) \stackrel{(2.11)}{=} R^I(t_m - t_n)$, whose Fourier transform provides the definition, according to (2.17), of the PSD of the I -th detector noise, $S^I(f)$. This allows to rewrite equation (2.28) in the popular form of GW literature as the *noise-weighted inner product* $(\mathbf{x} - \mathbf{h} | \mathbf{x} - \mathbf{h})$, where:

$$(\mathbf{a} | \mathbf{b}) := \int_{-\infty}^{+\infty} \frac{\tilde{a}(f)\tilde{b}^*(f) + \tilde{a}^*(f)\tilde{b}(f)}{S(f)} df \quad (2.29)$$

and the likelihood (2.27) becomes:

$$p(\mathbf{x} | \mathbf{h}) = \exp \left(-\frac{1}{2} \sum_I (\mathbf{x}_I - \mathbf{h}_I | \mathbf{x}_I - \mathbf{h}_I) + \int \log S^I(f) df \right). \quad (2.30)$$

Given the previous noise model for the (residuals of) the data, the statistical detection problem consists into distinguishing the two hypothesis:

H_0 : the "null hypothesis" of absence of a GW signal, $s(t) = 0$, and

H_1 : the "alternative hypothesis" of a GW signal at the detector output: $s(t) \neq 0$

on the base of whether the detector data \mathbf{x} are better described by the "noise only" model of hypothesis H_0 , or by the presence of a GW signal, modelled by $h(t)$. Several criteria exist to distinguish the previous two hypothesis: Bayes' theorem and the odds ratio [232], Von Neumann *minimax* criterion [233], and that based on the Neyman-Pearson lemma, of maximum detection probability at fixed false alarm rate [234]. As shown in [5], quite remarkably all of them return the same "optimal", in their respective sense, detection statistic. This is given by the *likelihood ratio* for the detector data in the two hypothesis, or equivalently, under the assumption of known stationary and Gaussian detector noise:

$$(\mathbf{x} | \mathbf{h}) \stackrel{(2.29)}{=} \int_{-\infty}^{+\infty} \frac{\tilde{x}(f)\tilde{h}^*(f) + c.c.}{S(f)} df. \quad (2.31)$$

This is called the *matched filter* detection statistic, which is basically a noise weighted correlation between the detector data and the putative GW signal. Notice also that in terms of the *whitened data* $x_w(t)$, whose Fourier transform is $\tilde{x}(f)/\sqrt{S(f)}$, the previous quantity is obtained by means of the *normalised filter* $\tilde{h}(f)/\sqrt{S(f)}$ of the GW signal template.

In the case of an SGWB, $h(t)$ is modelled as (a realisation of) a stochastic process, which is the same for each detector in the network. For this reason, the previous matched filter is

modified in such a way to match the output of the I -th detector with that of the J -th:⁸

$$\langle \mathbf{x}_I | \mathbf{x}_J \rangle = \int_{-\infty}^{+\infty} \frac{\tilde{x}_I(f) S_{IJ}(f) \tilde{x}_J^*(f) + c.c.}{S_I(f) S_J(f)} df \quad (2.32)$$

where $S_{IJ}(f) := \langle \tilde{x}_I^*(f) \tilde{x}_J(f) \rangle$ is the cross-PSD between the two detectors data. Assuming for $I \neq J$ the detector noises uncorrelated, the previous quantity represents the PSD of the gravitational signal, that is [235]:

$$\langle \tilde{x}_I^*(f) \tilde{x}_J(f) \rangle = \langle \tilde{h}_I^*(f) \tilde{h}_J(f) \rangle = \Gamma_{IJ}(f) S_h(f) \quad (2.33)$$

where $\Gamma_{IJ}(f)$ is called *overlap reduction function*, and represents the loss of *coherence* between the GW signals measured by the two detectors due to their separation and different orientation, and $S_h(f)$ is the PSD of the SGWB. Further details on the optimal detection techniques for this kind of signal can be found in [230] or in the seminal articles [236, 237].

The application of the matched filter technique and the corresponding modifications to equation (2.31) in the other case of interest, like unmodelled bursts or continuous waves, will not be presented in this manuscript, and we refer to their description in [3, 5, 50].

As we will discuss in the next section, it can be shown that (2.31) is the optimal linear filter for maximizing the SNR in the presence of additive stochastic noise. Also (2.32) is traditionally derived as the statistic maximising the SNR, *e.g.* in [75, 235]. Moreover, the considerations that we will make in the next section about the “non-optimal” filter apply to it as well.

To assess detection, the previous quantities should be compared to a threshold chosen on the base of the adopted optimality criterion, like the *false alarm probability* for the Neyman-Pearson criterion, or the value of the odds ratio for the Bayesian approach. This is also the base for *parameter estimation*, since the GW signal parameters θ are chosen as those that produce the model $h(t; \theta)$ (or $S_h(f)$ in the case of an SGWB) that maximise the previous quantities [3, §7.2].

2.1.5 Non-stationary noise and the “non-optimal” detection statistic

In (2.31) we have reported the expression for the “optimal” detection statistic, according to various criteria, and under the assumption that the noise is additive, stationary and Gaussian. In this section we will derive the optimality of its SNR in relation to a modified version of it. The latter can be the result, for example, of the failure of some of the aforementioned assumptions about the noise. For example, if the noise is not Gaussian, the likelihood (2.27) used to derive the expression of the filter, is no more valid, and so the expression in (2.31). Some authors have addressed the problem of deriving an alternative form of the matched filter statistic in the case of (usually small) deviations from the Gaussian distribution of the noise [238, 239]; *e.g.*, Student’s t -distribution [240]. Others, have proposed to *truncate* the detection statistic dropping those samples that make it not compatible with the previous assumption [241, 242].

⁸We have maintained the same notation of (2.29) since, when $I = J$, it reduces actually to $\langle \mathbf{x}_I | \mathbf{x}_I \rangle$ in the sense of that equation. For the formal derivation of (2.32), refer to [74] or [230].

Following [243], in this section we will consider the case of a “non-optimal” detection statistic arising from a *wrong* noise PSD, $S^w(f)$.⁹ This can be, for example, the consequence of the inherent statistical uncertainty in the estimation of the noise PSD (more details in section 2.1.6), or the result of non-stationarities in the data that make the standard estimation methods no more valid. From a theoretical point of view, in the latter case we can’t even apply the Wiener-Khinchin theorem discussed in section 2.1.2, and also the notion of a PSD is ill-defined. As mentioned in section 2.1.3, in this case one should refer to other definitions for the spectral content of a process, as for example evolutionary spectra or LSPs. To avoid complications, what we usually do in GW data analysis is to divide (from a formal point of view, at least) the detector noise in two subprocesses; a *background* noise, satisfying at a sufficient level of approximation the hypotheses of stationarity and Gaussianity, to which there are superimposed other noise transients. These are typically spurious excesses of power, called *glitches* [38], which of course are not compatible with the assumptions of the matched filter technique. If this is the case, and none of these glitches overlap with the signal we aim to detect (occurrence to be faced in chapter 4), we can refer to the PSD of the background noise in the expression of the matched filter, without any consequence but one: in practice, we don’t know in advance its value, which should be *estimated* from the recorded data. This problem will be discussed in the next section and, as we will see, the presence of glitches usually implies a bias in the PSD estimate. On the other hand, if the non-stationarity occurs with slower time scales, there is no clear distinction between this and the subprocess that constitutes what we have previously called “background”, and the previous separation of scales becomes ambiguous.

Let’s examine what are the effects that this incorrect value of the noise PSD on the matched filter SNR. For simplicity, let us assume that the GW signal has compact support in both time and frequency, and that the statistical properties of the (Gaussian) noise don’t change within the duration of the signal template $h(t)$. We can define our “wrong” PSD as:

$$S^w(f) = S^c(f) (1 + \epsilon(f)) \quad (2.34)$$

where $\epsilon(f)$ is the (frequency-dependent) fraction of power that the estimate has incorrectly taken into account. In practice, we want this quantity to be small with respect to one but, as it is written, it can be quite general and both positive and negative.

Let us also assume to have strain data (in the frequency domain) of the form:

$$\tilde{x}(f) = \rho \tilde{h}(f) + \tilde{n}(f) \quad (2.35)$$

where $\tilde{h}(f)$ is the template waveform and ρ is its SNR with respect to the “hypothetically” known noise; if we knew its spectral content, the normalisation of the template waveform

⁹Looking at (2.31), another factor that could make the matched filter non-optimal is the template waveform $\tilde{h}(f; \theta)$. This is a function of the source parameters θ and, of course, it is not possible to grid up the parameter space infinitely finely in order to cover exactly every possible value of them. It is important then to minimise the systematic error associated to a non-perfectly matching template, designing a “sufficiently fine” template bank of filters. This problem has been addressed for example in [244–246] and carefully reviewed in [3].

would have been such that

$$\int_{-\infty}^{+\infty} \frac{\tilde{h}(f)\tilde{h}^*(f) + c.c.}{S^c(f)} df = 2 \int_{-\infty}^{+\infty} \frac{|\tilde{h}(f)|^2}{S^c(f)} df := \int_{-\infty}^{+\infty} I(f) df = 1 \quad (2.36)$$

as we will derive in a moment. Let us stress the fact that the previous normalisation remains unknown as long as we have a wrong estimate of the noise PSD. Notice also that $I(f)$ satisfies the conditions for being a p.d.f. (refer to the definition in section 2.1.2); this interpretation will prove to be useful in what follows.

We can compute the normalised matched filter output with the wrong PSD estimate:

$$(\mathbf{x}|\mathbf{h})_w = \int_{-\infty}^{+\infty} \frac{\tilde{x}(f)\tilde{h}^*(f) + c.c.}{S^w(f)} df. \quad (2.37)$$

In the absence of a signal $\varrho = 0$, the expectation value of the previous quantity is clearly zero, and its variance can be computed as:

$$\begin{aligned} \sigma_w^2 &= \langle |(\mathbf{x}|\mathbf{h})_w|^2 \rangle = 2 \int \frac{\tilde{h}(f)\tilde{h}^*(f)}{S^w(f)^2} S^c(f) df \\ &= 2 \int \frac{|\tilde{h}(f)|^2}{S^c(f)} \frac{1}{(1 + \epsilon(f))^2} df = \int I(f) \frac{1}{(1 + \epsilon(f))^2} df. \end{aligned} \quad (2.38)$$

However, this quantity is unknown since it implicitly contains the correct value of the noise PSD at the time of the putative signal; we can (incorrectly) estimate it as:

$$\sigma_w^2 = 2 \int \frac{|\tilde{h}(f)|^2}{S^w(f)} df = \int I(f) \frac{1}{1 + \epsilon(f)} df.$$

Instead, if a signal is present, $\varrho \neq 0$, the expectation value of the previous quantity is not zero:

$$\langle (\mathbf{x}|\mathbf{h})_w \rangle = 2\varrho \int \frac{|\tilde{h}(f)|^2}{S^w(f)} df = \varrho \int I(f) \frac{1}{1 + \epsilon(f)} df. \quad (2.39)$$

Thus, if we use the previous “wrong” estimates of the noise PSD we have for the SNR of the wrong detection statistic:

$$\rho_w^2 = \frac{\langle (\mathbf{x}|\mathbf{h})_w \rangle^2}{\sigma_w^2} = \varrho^2 \frac{\left(\int I(f) \frac{1}{1 + \epsilon(f)} df \right)^2}{\int I(f) \frac{1}{(1 + \epsilon(f))^2} df} \quad (2.40)$$

which is clearly not equal to ϱ^2 , as we may have expected (except in the case $\epsilon = \text{const.}$).

On the other hand, if we had the correct estimate of the noise PSD, the detection statistic would be:

$$(\mathbf{x}|\mathbf{h})_c = \int_{-\infty}^{+\infty} \frac{\tilde{x}(f)\tilde{h}^*(f) + c.c.}{S^c(f)} df \quad (2.41)$$

whose expectation value is

$$\langle (\mathbf{x}|\mathbf{h})_c \rangle = 2\varrho \int \frac{|\tilde{h}(f)|^2}{S^c(f)} df = \varrho \int I(f) df = \varrho \quad (2.42)$$

and variance, in the absence of a signal ($\rho = 0$):

$$\sigma_c^2 = \langle |(\mathbf{x}|\mathbf{h})_c|^2 \rangle = 2 \int \frac{|\tilde{h}(f)|^2}{S^c(f)} df = \int I(f) df = 1. \quad (2.43)$$

Hence, the SNR in the correct case is, as expected:

$$\rho_c^2 = \rho^2.$$

The latter relation justifies the normalisation (2.36) chosen for the template $h(t)$.

We can then evaluate what is the loss of SNR that we have making use of the wrong PSD estimate, by means of the ratio of the correct estimate and the wrong one:

$$\frac{\rho_c^2}{\rho_w^2} = \frac{\int I(f) \frac{1}{(1+\epsilon(f))^2} df}{\left(\int I(f) \frac{1}{1+\epsilon(f)} df \right)^2}. \quad (2.44)$$

A couple of comments on the previous results shall follow. First of all, if $\epsilon(f)$ is frequency independent, the ratio in the previous equation is equal to one, and we have no loss of SNR due to a “scale” misestimate of the noise PSD. This is a property of the matched filter, which weights more those frequency components that are relatively less noisy, so it is insensitive to global (frequency independent) rescaling of the noise level. To put it in another way, in the language of Audio Engineering, the matched filter is most proficient to detect “colour differences” between the noise and the signal. For the same reason, the previous quantity is independent on the value of the theoretical SNR ρ . Second observation is that this ratio is always smaller than or equal to one, implying that (2.31), with the correct noise PSD at the denominator, is actually the optimal detection statistic that maximises the SNR. Indeed, if we had to interpret $I(f)$ as the p.d.f. for the random variable f , and we define the stochastic variable $z := (1 + \epsilon(f))^{-1}$, the previous result can be read:

$$\frac{\rho_c^2}{\rho_w^2} = \frac{\langle z^2 \rangle}{\langle z \rangle^2} \geq 1$$

which is greater than, or equal to one for the positiveness of the variance of a real random variable: $Var[z] = \langle z^2 \rangle - \langle z \rangle^2 \geq 0$.¹⁰ We can expand the previous result for small $\epsilon(f)$:

$$\begin{aligned} \frac{\rho_c^2}{\rho_w^2} &= \frac{\int I(f) (1 - 2\epsilon(f) + 3\epsilon(f)^2 + \mathcal{O}(\epsilon^3)) df}{\left(\int I(f) (1 - \epsilon(f) + \epsilon(f)^2 + \mathcal{O}(\epsilon^3)) df \right)^2} \\ &= \frac{1 - 2 \int I(f) \epsilon(f) df + 3 \int I(f) \epsilon(f)^2 df + \mathcal{O}(\epsilon^3)}{\left(1 - \int I(f) \epsilon(f) df + \int I(f) \epsilon(f)^2 df + \mathcal{O}(\epsilon^3) \right)^2} \\ &= (1 - 2 \int I \epsilon df + 3 \int I \epsilon^2 df + \mathcal{O}(\epsilon^3)) \\ &\quad \times \left(1 + 2 \int I \epsilon df - 2 \int I \epsilon^2 df + 3 \left(\int I \epsilon df \right)^2 + \mathcal{O}(\epsilon^3) \right) \end{aligned}$$

¹⁰This can also be seen as a special case of the *Cauchy-Schwarz inequality*. Indeed the expectation value of a square is always positive, $\langle (z - \langle z \rangle)^2 \rangle = \langle z^2 \rangle - \langle z \rangle^2 \geq 0$, and therefore $\langle z^2 \rangle \geq \langle z \rangle^2$, where the equality holds only for constant (non-stochastic) z , i.e. $\langle z \rangle = z$. In general, the larger the variations of z (or of $\epsilon(f)$), the bigger the ratio ρ_c^2 / ρ_w^2 , and the loss of SNR due to (2.34).

$$= 1 + \int I(f) \epsilon(f)^2 df - \left(\int I(f) \epsilon(f) df \right)^2 + \mathcal{O}(\epsilon^3) \geq 1 \quad (2.45)$$

The leading order correction to ρ_c^2 is then quadratic in $\epsilon(f)$, and of course it does not depend on its sign, that is whether the estimate error in the PSD is positive (less power) or negative (more power).

As commented by the authors of [243], for a typical PSD estimate by means of the Welch's method (refer to the next section) on M data segments, the relative error in this estimate is $\approx 0.7M^{-1/2}$ for each frequency bin, where the constant factor is obtained from numerical simulations and encompass the effects of the windowing and the overlap. In turn, from (2.45), the authors quote that this implies a relative error in the estimated SNR of order $0.5M^{-1}$ (although this results clearly depends on the waveform model, through the parameter $I(f)$ and, for comparison, the shape of the noise PSD).

2.1.6 Time series analysis and spectral estimations

As noted before, except for those cases where we have prior theoretical knowledge, we usually don't have access to the statistical structure of the stochastic process under study. What we have is just one, or few, finite realisation(s) of it, recorded in the form of time series, which we will assume to be evenly sampled at frequency f_S : $x_n := x(t_0 + n/f_S)$, with index $n = 0, \dots, N-1$. The *ergodic theorem*, when valid, provides an important recipe for reconstructing, from these values, some properties of the process. In particular, we are usually interested in the lowest order moments, the mean and the covariance, which comprise all the information needed to fully characterise a stationary Gaussian process. In this case, we are assuming to know that the data under study is stationary (no dependence on the *epoch* t_0) and Gaussian (described by (2.13)), with unknown parameters though. From (2.20) we can find the *estimators* ("hat" symbol) for the time averaged quantities related to the process. For example, the *sample mean*:

$$\hat{\mu} := \frac{1}{N} \sum_{n=0}^{N-1} x_n, \quad (2.46)$$

which is immediate to verify that, if considered itself a random variable of x_n , is an *unbiased estimator*, $E[\hat{\mu}] = \mu$, whose variance equals:

$$\text{Var}[\hat{\mu}] = \frac{1}{N^2} \sum_{n=0}^{N-1} \sum_{p=0}^{N-1} \text{Cov}(x_n, x_p) = \frac{1}{N} \sum_{l=-N}^N \left(1 - \frac{|l|}{N}\right) R_l$$

with $R_l := R(l/f_S) = K(n/f_S - p/f_S, 0) = K(n/f_S, p/f_S)$ the *auto-covariance sequence* of the stationary process, which therefore depends only on the "lag" $l = m - p$. This variance approaches zero as $N \rightarrow \infty$, making (2.46) a *consistent estimator*.

There are two definitions for the *sample correlation function*:

$$\hat{R}_l' := \frac{1}{N - |l|} \sum_{n=1}^{N-|l|} x_{n+|l|} x_n, \quad |l| < N \quad (2.47)$$

which is *unbiased*, $E[\hat{R}_l] = R_l$, and consistent, but has the major drawback of being not always positive semi-definite [6, §5.3]. This issue is avoided with the *biased* estimator [247, 248]:

$$\hat{R}_l := \frac{1}{N} \sum_{n=1}^{N-|l|} x_{n+|l|} x_n, \quad |l| < N. \quad (2.48)$$

This is the preferred choice of estimator (refer to the discussion in [6]), especially in the context of *spectral estimations*.

Spectral estimations: the periodogram and the Bartlett's method

From equation (2.17) we know that the PSD of a stationary stochastic process is defined as the Fourier transform of the auto-correlation function. Therefore, the most natural estimator for the PSD is the so called *periodogram* [249], given by the discrete Fourier transform of the previous auto-covariance sequence estimator:

$$\hat{S}(f_k) := \frac{1}{f_S} \sum_{l=-N}^N \hat{R}_l e^{-2\pi i k l / N} \stackrel{(2.48)}{=} \frac{1}{f_S N} \left| \sum_{n=0}^{N-1} x_n e^{-2\pi i k n / N} \right|^2 := \frac{f_S}{N} |\tilde{x}_k|^2, \quad (2.49)$$

with discrete frequencies $f_k := k/T$, $-\lceil N/2 \rceil \leq k \leq \lceil N/2 \rceil$.¹¹ Therefore, the periodogram is simply the squared modulus of the DFT of the data, normalised by $f_S/N = 1/T$. Its expectation value can be computed as:

$$\begin{aligned} E[\hat{S}_k] &= \frac{1}{f_S} \sum_{l=-N}^N E[\hat{R}_l] e^{-2\pi i k l / N} = \frac{1}{f_S} \sum_{l=-N}^N \frac{N-|l|}{N} R_l e^{-2\pi i k l / N} \\ &:= \frac{1}{f_S} \sum_{l=-\infty}^{+\infty} w_l^B(N) R_l e^{-2\pi i k l / N} \end{aligned} \quad (2.50)$$

where $w_l^B(N)$ is the *Bartlett's* (or *triangular*) *window* of length N [250]:¹²

$$w_l^B(N) = \begin{cases} \frac{N-|l|}{N} & |l| < N \\ 0 & \text{otherwise} \end{cases} \quad (2.51)$$

The discrete-time Fourier transform of the product $w_l^B(N) R_l$ in equation (2.50) is equals, from the Wiener-Khinchin theorem, to the convolution of their Fourier transforms, that is, the true PSD and the so called *Fejér kernel* [251] (*i.e.* the Fourier transform of the triangular window):¹³

$$E[\hat{S}_k] = \int_{-f_S/2}^{+f_S/2} S(f') F_N(f - f') df' \quad (2.52)$$

¹¹Refer to the convention on the *Discrete Fourier Transform* (DFT) in the Notation section, and the cyclic symmetry property, with respect to N , of this transform.

¹²This window is implicitly contained in the definition of the DFT \tilde{x}_k in (2.49). This equals indeed a "truncated" version of the discrete-time Fourier transform $\tilde{x}(f)$ by means of a rectangular window of length N . Being the periodogram the square modulus of the DFT of the data, this implies the convolution of two rectangular windows, hence a triangular or Bartlett window.

¹³The analogous computation for the unbiased estimator (2.47) would have led to an estimator for the PSD with expectation value given by the convolution of $S(f)$ with the *Dirichlet kernel* $D_N(f) := \sin(2\pi(N + 1/2)f) / \sin(2\pi f/2)$. Differently from the Fejér kernel, the latter is not always positive, and so is the PSD estimate [252].

with:

$$F_N(f) := \frac{1}{N} \frac{\sin^2(2\pi fN/2)}{\sin^2(2\pi f/2)}.$$

This term implies a bias in the PSD estimate,¹⁴ dependent on the finite sample size N , known as *spectral leakage*. The origin of this is clearly the “window function” (or *taper*) $w_n^B(N)$, which, for the case of the periodogram, is simply the Fourier transform of the triangular window. Therefore, the presence of this bias can be mitigated changing the window, that is, multiplying the data x_n by another windowing function $w_n(N)$ ($\tilde{w}_k(N) := \sum_{n=0}^{N-1} w_n(N) \exp(-2\pi i kn/N)$), before computing the DFT (*tapering*); details on this procedure and the window choice are thoroughly described in [250, §10.2.1].¹⁵

Another (major) problem of the periodogram estimate is that, asymptotically, the variance of the estimate is equal to the expected value squared:

$$\text{Var}[\hat{S}(f)] \approx S(f)^2, \quad \text{for } f \in (-f_s/2, f_s/2) \quad (2.53)$$

In particular this will never decrease to zero at large values of N , implying that the periodogram is not a consistent estimator of the true PSD.

We will briefly discuss the modification to the periodogram PSD estimate needed to overcome the previous issues, leading to how PSD are usually measured in GW data analysis. The first modification to the aforementioned method is due to Bartlett [253], and consists into dividing the whole data sample into M non-overlapping segments of N data: x_{n+iM} for $n = 0, \dots, N-1$ and $i = 1, \dots, M$. Then M periodograms are computed, as in equation (2.49), and the final Bartlett’s PSD estimate is computed as their mean:

$$\hat{S}_B(f) := \frac{1}{M} \sum_{i=1}^M \hat{S}^{(i)}(f). \quad (2.54)$$

In this way, it can be shown that the resulting estimator is still biased, as in (2.52), but now its variance decreases as $1/M$, making it a consistent estimator in the limit of large M .

This result masks a subtlety inherent to the averaging method. If the data are collected for a time T and evenly sampled at a frequency f_s , $x(t_n) = x(n/f_s)$ with $n = 0, \dots, N-1$ and $N = T \cdot f_s$, the usual selection of discrete frequency values for the components \tilde{x}_k is (refer to Notation section):

$$f_k := \frac{k}{T} = \frac{k f_s}{N}, \quad \text{for } k = 0, \dots, N-1.$$

These frequencies are unique only to $k = \lceil N/2 - 1 \rceil$, which corresponds to the *Nyquist frequency* associated to the sampling f_s , thus the frequencies range from f_s/N to $f_s/2$. Hence, in the Bartlett’s method, for a fixed number of sample data $N' = N \cdot M$, computing M averages reduces the variance of the estimate by a similar factor, but also the frequency resolution from N' points to $N = N'/M$. This is a trade-off one should take into account when choosing the parameters for the spectral estimation.

¹⁴Except in the case the process is *white*, that is the auto-correlation function is a delta: $K_{jk} = \delta_{jk}$, $S(f) = \text{const.}$.

¹⁵As observed in [2], windowing is of primary importance in GW data analysis, since an improper application of this method (or no windowing at all) can result in spurious phase correlations in the Fourier transform of the data; refer to figure 4 therein.




Figure 2.1: Representation of the basic ideas behind the Welch's method for PSD estimation. The recorded data sample (top of the figure) is divided into windowed, overlapping segments, which may also be zero-padded.

The previous choice of the sampling of the spectrum, in some sense, may not provide a good representation of the continuous PSD $S(f)$. The *zero padding* technique provides an efficient interpolation method, by means of the addition of intermediate frequencies, simply by increasing the sample length to $L > N$ points with the addition of zeroes before and after the actual data [250, §14.1.3]. It is important to note however that this doesn't provide any improvement in the frequency resolution; it is just an interpolation at more frequencies, as visible from the top-right frame of figure 2.2, where the padding (dashed black line) doesn't increase the resolution to which the 100 Hz *spectral line* is reconstructed with respect to the non-padded estimate (blue). The usefulness of this method is due to the fact that the computation of the discrete Fourier transform (an operation of typical *complexity* $\mathcal{O}(N^2)$) can be significantly speed up choosing L as the "next power of two" of N , and exploiting the increased speed of the *Fast Fourier transform* (fft) algorithm, which is of typical complexity $\mathcal{O}(N \log N)$, for example by means of the *Cooley-Tukey algorithm* [254]. This operation must be done *after* applying any windowing function to the signal, in order to avoid the spectral leakage due to the sharp transition between the data and the zero padding.

Welch's method

The most popular modification [255] to the simple periodogram PSD estimate is due to Welch [256]. This is based on the Bartlett's method, of averaging several periodogram estimates, with two important improvements. Firstly, as previously noted, one can multiply the data to an adequate window function, before computing the discrete Fourier transform, in order to reduce the spectral leakage (and marginally the variance):

$$\hat{S}^{(i)}(f) := \frac{1}{N f_S} \left| \sum_{n=0}^{N-1} w_n x_n^{(i)} e^{-2\pi i f n} \right|^2, \quad i = 1, \dots, M$$

$$\hat{S}_W(f) := \frac{1}{M} \sum_{i=1}^M \hat{S}^{(i)}(f) \quad (2.55)$$

where the window function is normalised such that:

$$\frac{1}{N} \sum_{n=1}^N w_n^2 = 1.$$

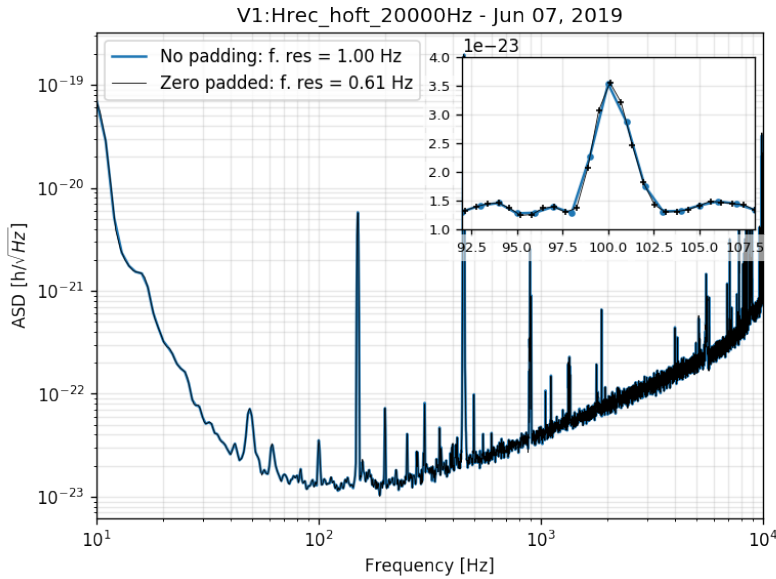


Figure 2.2: One-sided amplitude spectral density (ASD) of Hrec_hoft_20000Hz of June 7, 2019. The top right plot shows a zoom in of the region around 100 Hz.

Secondly, Welch proposed to use overlapping segments, by an amount that is usually taken around 50% of the segment data. The reason for that is to rebalance the degree by which all the samples in the various segments are weighted, in the estimation of the spectrum, for the effect of the windowing; refer to its depiction in figure 2.1. On the other hand, the use of overlapping segments implies correlations between them, which has the effect of increasing the variance and, usually, making more complex the computation of quantities related to them. The choice of the overlap is, in a certain sense, an optimization problem seeded in the choice of the window itself, and the final result thus depends on the shape of the particular window one aims to adopt. For the *Blackman window*, which is usually the preferred choice for reducing the leakage, the optimal overlap turns out to be 67%, as a generalization of Kaiser α -4 type windows [257]. For GW data analysis, the preferred choice is however 50% for the reason that we will discuss momentarily, in the next section.

The key ideas behind the Welch’s method are summarised in figure 2.1. The effects of the padding are represented in figure 2.2, where the ASD (square root of the PSD, $S(f)^{1/2}$) of Advanced Virgo strain data is estimated over 20 seconds of data, sampled at $f_s = 20$ kHz (4×10^5 data points), averaging (39 times) 50% overlapping segments of 1 second of duration (2×10^4 data per segment). The resulting frequency resolution is 1 Hz. The black line is obtained zero padding the previous segments to the next power of two: $20000 \rightarrow 32768 = 2^{15}$ data points.

To be precise, the quantity represented in figure 2.2, over positive frequencies only, is the so called “one-sided” ASD, whilst (the square root of) that defined in (2.17) is usually referred to as “two-sided” (PSD). For a real signal, the latter quantity is an even function of the frequency ($R(\tau)$ is real), so it is redundant to show both the contributions at positive and negative frequencies. Hence, they are usually shown one-sided spectra, where the power in

the negative frequency components is “folded” to the positive semiaxis:

$$S^{1s.}(f) = 2S^{2s.}(f) = \int_0^{+\infty} R(\tau)e^{-2\pi i f\tau}d\tau \quad (2.56a)$$

$$= 4 \int_0^{+\infty} R(\tau) \cos(2\pi f\tau)d\tau, \quad \text{for } f \geq 0. \quad (2.56b)$$

The two-sided spectral density is the most convenient one to use for calculations, and basically it is the only one used throughout this dissertation; we will therefore omit any superscript “2s.” for referring to it. Instead, the one-sided version of it is the quantity that is most convenient to report in graphics; therefore, in fig. 2.2 and all the others, when we say they represent the ASD or PSD, we actually mean the one-sided version of it. Moreover, the latter is sometimes the preferred version to Electronic Engineers since it can be directly measured by band-pass filtering(-squaring-averaging) procedures of the target signal; refer to its description in [213, §5.2.3].

Median averaged spectrum

The estimators provided by the Bartlett’s and Welch’s methods, (2.54) and (2.55), are sometimes referred to as “mean averaged” spectral estimates, for they are basically mean averages of the squared discrete Fourier transforms (2.49), computed in the various data segments.¹⁶

In the context of GW data analysis, for the application of the matched filter technique, we need to estimate what we have called in section 2.1.5 the *background noise* (i.e. the supposedly dominant part of it that is stationary and Gaussian) PSD. The problem with using the Welch’s method for power spectral estimation is that for detector noise containing significant excursions from its background behaviour, due to glitches or even to the transient GW signal we aim to detect, the mean used in (2.55) can significantly *bias* the estimate of this PSD.¹⁷ Indeed, typical GW detector data exhibits transient power excesses at a rate of one every 10 seconds for SNR ~ 7 glitches,¹⁸ and at about one every 10 minutes for loud SNR ~ 20 glitches [258].

To be more “robust” against these noise transients, the idea pioneered in [255, 259] has been to replace the mean by a *median*, which, as an estimator, is less sensitive to the tail of the distribution with respect to the *central tendency*. The following¹⁹

$$\hat{S}_{med}^{1/2}(f) = b^{-1} \text{median}[\hat{S}^{(1)}(f), \hat{S}^{(2)}(f), \dots, \hat{S}^{(M)}(f)]^{1/2} \quad (2.57)$$

¹⁶Notice also that, for the same reason, the corresponding estimator for the ASD, $\hat{S}(f)^{1/2} = \sqrt{f_s/N} \left(M^{-1} \sum_{i=1}^M |\tilde{x}^{(i)}(f)|^2 \right)^{1/2}$, is often referred to as the *Root Mean Square* (rms) ASD estimator.

¹⁷In reality, in the analysis of a potential GW transient signal, the PSD estimation is generally done “off-source”, that is, the stretch of data corresponding to the putative signal is *gated*, and the data therein not used for the estimate. The same gating procedure is also applied in correspondence of the loudest glitches.

¹⁸In this case, due to the high rate of this kind of non-stationarities, it is sometimes preferable to describe the detector noise, instead of a stationary Gaussian background plus transients, with a modified Gaussian distribution with “heavier tails”, like a *t*-distribution, as for example in [238].

¹⁹Notice that, since the square root is a monotonically increasing function, the square root of the median is equivalent to the median of the square roots of the values.

is called *median averaged ASD*,²⁰ and the factor b is a constant inserted in order to have consistent estimates, in term of expectation values, between this method and the usual mean averaged spectrum obtained with the Bartlett's or Welch's methods.

The normalisation factor b can be computed in the hypothesis of stationary and Gaussian noise. If the stochastic process corresponding to the strain $X(t)$ is Gaussian, this means that the real and imaginary parts of its Fourier transform (Cramèr representation, to be more precise; sec. 2.1.2), $Y := \Re \tilde{X}(f)$ and $Z := \Im \tilde{X}(f)$, are, at each frequency f , independent and identically distributed Gaussian random variables with zero mean and variance $\sigma^2(f)$.²¹ Ignoring, for now, the frequency dependency, this means that:

$$p_{Y,Z}(y,z) = p_Y(y) \cdot p_Z(z) = \frac{1}{2\pi\sigma^2} e^{-\frac{y^2+z^2}{2\sigma^2}}. \quad (2.58)$$

In the estimations of the PSD we are interested in the squared modulus of $\tilde{X}(f)$, $|\tilde{X}(f)|^2 = Y^2 + Z^2 := R^2$. We can then perform the change of variables, expressing the previous bivariate Gaussian distribution as a p.d.f. for the variable R :

$$\begin{aligned} p_R(r; \sigma^2) &:= \frac{d}{dr} P_R(r) = \frac{d}{dr} P_{Y,Z} \left(\sqrt{y^2 + z^2} = r \right) = \frac{d}{dr} \iint_0^{r=\sqrt{y^2+z^2}} \frac{1}{2\pi\sigma^2} e^{-\frac{y^2+z^2}{2\sigma^2}} dy dz \\ &= \frac{d}{dr} \int_0^{2\pi} \int_0^r \frac{1}{2\pi\sigma^2} e^{-\frac{r'^2}{2\sigma^2}} dr' r' d\phi = \frac{d}{dr} \int_0^r \frac{r'}{\sigma^2} e^{-\frac{r'^2}{2\sigma^2}} dr' \\ &= \frac{r}{\sigma^2} e^{-\frac{r^2}{2\sigma^2}}, \quad \text{for } r \geq 0. \end{aligned} \quad (2.59)$$

Some properties of this distribution are reported in Appendix B.1. The important thing to notice here is that, under the assumptions of stationarity and Gaussianity, the noise ASD is distributed, at each frequency, as a Rayleigh variable. Instead, the PSD is an exponential distributed variable with scale parameter $2\sigma^2$ ($R^2 \equiv S \sim \exp(-s/\beta)/\beta$, $\beta = 2\sigma^2$), also equivalent to a *Gamma* distribution with *shape parameter* $k = 1$ and *scale parameter* $\theta = 2\sigma^2$; further details in Appendix B.2. This property will be useful later.

Then, for non-overlapping segments of data, the rms ASD estimator (square root of (2.55)) provides an estimate of the rms of a Rayleigh distributed variable, while the *median* estimator (2.57) gives an estimate of its median. It can be shown that the expectation value of the first quantity, at each frequency, converges to the distribution rms $\sqrt{2}\sigma$ while the latter converges to the distribution median $(\log 4)^{1/2}\sigma$, for $M \rightarrow \infty$. Thus, asymptotically, the bias in the median ASD estimate is $b_\infty = (\log 2)^{1/2}$.

This is not exact for a finite number of averages ($M < \infty$) and if the segments are overlapping, since adjacent data segments will not have N data out of which to obtain M independent estimates of the ASD. In [255, §B] the authors discuss the value of this bias for finite values of M , for non-overlapping segments, obtained from the analytical expressions of the

²⁰It will be clear momentarily the reason why it is preferred to work with the ASD.

²¹Notice that the multiplication by a windowing function doesn't usually change this property. Notice also that for this section we have temporarily recovered the upper- lower-case distinction between random variables and the values assumed by their realisations.




Figure 2.3: Spectrogram representing the ASD of 600 seconds of Advanced Virgo strain data. Notice at 02:06:10 UTC a vertical line representing a loud glitch in the data.

biasses in the expectation values of the estimates (2.54) and (2.57) [260]:

$$b(M) = \begin{cases} \left(\sum_{i=0}^{M-1} \frac{(-1)^i}{m-1} \right)^{1/2} & \text{for } M \text{ odd} \\ b(M-1) & \text{for } M \text{ even} \end{cases} \quad (2.60)$$

As regards the overlap, in the case this is $\leq 50\%$, in [255] the authors suggest to compute two median spectra (2.57), one containing the “odd” i ’s, corresponding to M_{od} non-overlapping segments, and the other with the M_{ev} non-overlapping “even” i ’s. These are multiplied by the corresponding (finite M) bias factors, and their quadrature sum weighted by the number of segments in each of them:

$$\hat{S}_{mm}^{1/2}(f) := \left(\frac{M_{od} \hat{S}_{med}^{od}(f) + M_{ev} \hat{S}_{med}^{ev}(f)}{M} \right)^{1/2}. \quad (2.61)$$

This is known as *mean-median averaged* ASD estimate, and it is currently the preferred approach implemented in many GW data analysis pipelines.

To understand the difference between the standard Welch’s method and the mean-median averaged ASD, in figure 2.3 we reported the *spectrogram* of 600 seconds of Advanced Virgo data. This corresponds to a time-frequency map where at every time bin is shown (vertical axis) the ASD computed over the corresponding bin of data. The third dimension (colormap), represent the intensity of the ASD in the particular time-frequency bin. Spectrograms are in many circumstances the most immediate method to spot variations in the spectrum of a certain signal. For example, in this figure is visible a sharp vertical line, on the right-hand side of the map, at about 02:06:10 UTC. This is a quite loud glitch, not that infrequent over similar data lengths; many other fainter glitches are likely to be present in this data, although not evident from the spectrogram. Overall, the background noise appears

quite uniform, horizontally, during the entire stretch of data. In figure 2.4 the ASD is estimated for the first 300 seconds of figure 2.3 (no glitch), fig. 2.4a, and for the last 300 seconds, fig. 2.4b, where the glitch is present, using the two methods described: Welch’s method (solid blue line) and mean-median average (dashed black). In the plots on the bottom, the relative differences between the two estimates are reported. It is clearly visible their difference when the glitch is present. Also, the mean-median spectrum is almost left unchanged by the presence of the glitch, showing the robustness of this method against transients.

Let’s comment this result. The mean-median averaged method, being robust against noise transients, is able to provide a more trustworthy estimate of the background noise than the standard Welch’s method. Hence, its application should be preferred for the matched filter technique. This however doesn’t imply that this estimate comes without any potential bias or systematic error with respect to the “true” background noise PSD. Moreover, if the non-stationarities are slower, or the glitches occur at an high rate, the median is altered as well, and in general it can’t be identified any stationary background noise. Also, it is not unlikely, especially with the improved sensitivities of the advanced detectors, that some GW event occurs in the vicinity of a glitch/non-stationarity, as we will face in chapter 4.

What we have described here is a first technique for the mitigation of non-stationary noise effects in GW data analysis. The next chapter will be entirely devoted to other strategies aimed at this fundamental task.

2.2 Noise sources in GW detectors

As we derived for a simple Michelson interferometer, noise is everything at the detector output that can affect equation (1.59) but that is not of astrophysical origin. In section 1.3.1, we already started a classification of it on the base of how it can affect this equation and the assumptions on the base this was derived. We called *displacement noise* everything that can alter the position of the test masses, modifying for example $L_{x,y}$, and *phase* (or *timing*) *noise* what affects the phase of the light during its propagation between them. As we will discuss in this section, many noise sources can enter in the former class, producing a residual movement (after the suspension mechanisms and the corresponding filters) on the optical benches and the test masses. Some sources for the latter are instead: the scattering with the residual gas that makes the propagation of light not actually in a vacuum, the quantum vacuum fluctuations (briefly touched when talking about squeezing in section 1.3.2), and the light back scattering. Moreover, related to the readout principle, we discussed about the *measurement noise* relative to the fluctuating number of photons arriving at the detection photodiode (*shot noise*).

This section contains further details about noise sources in ground based interferometric detectors, particularly focussed on their origin. Part of the material presented here is taken from the Advanced detectors technical design reports [7, 12] and from some “standard” textbooks on the subject, like [95, 173, 261].

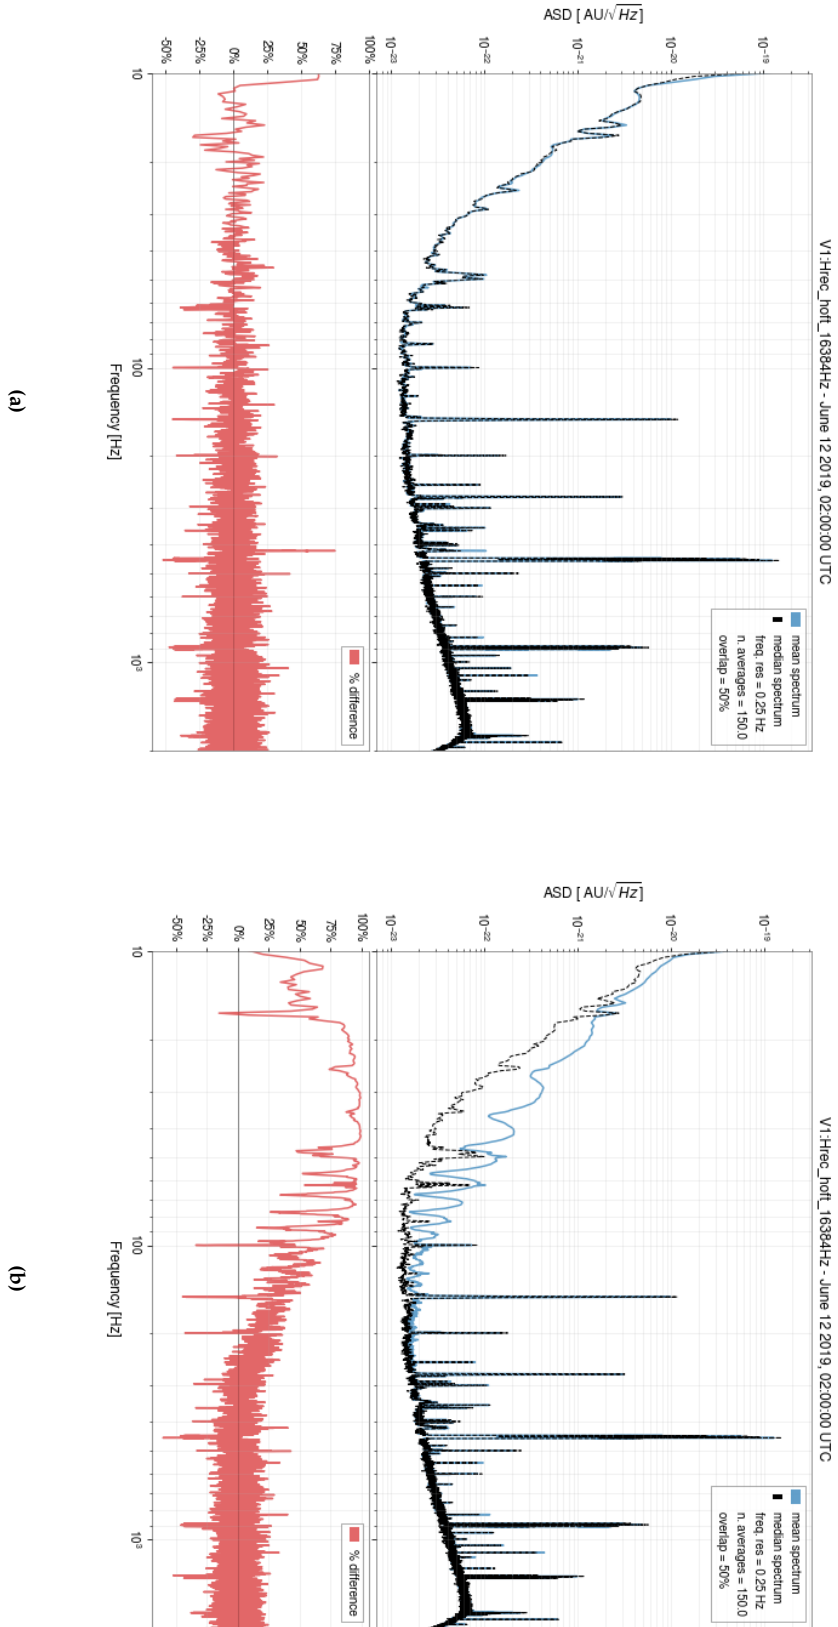


Figure 2.4: Comparison between the Welch's, "rms-averaged" ASD and the mean-median averaged one commonly used in GW data analysis. In figure 2.4a, the ASDs are estimated over a relatively quite stretch of data, corresponding to the left-hand side of figure 2.3, and their relative difference (bottom plot) is almost everywhere close to zero. In figure 2.4b the two estimates are computed in the presence of a loud glitch, as visible from the right-hand side of figure 2.3.

2.2.1 Noise origins and general classification

In “order of relevance” for the actual detector performances, the noise sources that limit the detector sensitivity can be classified on the base of their origin as [28]:

- *Fundamental noises*, arising from the physical limitation of the detection principle, as outlined in section 1.3.2, and its implementation. They are for example the *shot noise* (1.61), due to the quantum nature of light, or the *thermal noise*, from the vibrational motion of the various detector components related to their temperature, or the residual *seismic noise*, from the ground vibrations, not completely filtered out by the suspension systems. These are typically stationary, and their effects can be conveniently represented in the frequency domain. This is reported by the continuous lines in figure 2.5. As anticipated, in this domain most of them have a very distinguishable signature (or *color*), hence the convenience of this representation.
- *Technical noises*, coming from the actual implementation of the detector with components that are not optimal or that have imperfections that prevent them to work by design. Examples are the *residual gas* present in the optical cavities where the laser beams circulate, or the electrical and mechanical resonances due to the excitations of some components; the latter are responsible of some of the sharp spectral lines in figure 2.3 [262]. Moreover, this class also includes *control noise*, which is reintroduced in the system by the feedback control loops used to maintain the correct operating point of the various parts of the detector, as a consequence of non-optimal control filters or to the unavoidably noisy error signals or actuators. Most of the contributions of this family of noise (at least the stationary ones, representable by a PSD) is modelled by the dashed lines in figure 2.5.
- *Environmental noise*, in a sense, comprises all the causes of noise that are not included in the previous categories. Usually these are originated from the outside of the detector, and can affect several systems and sensors through different kind of couplings. In building the detector infrastructures, state-of-the-art hardware has been adopted to insulate the instrument from the local environment, but nothing is perfect and part of these disturbances are transmitted to the strain signal measured by the detectors. Some examples are *acoustic* and *seismic vibrations* and electro-magnetic fields, potentially of anthropogenic origin (people on site, traffic, airplanes) or weather conditions (wind, thunderstorms, sea-storms). In general their effects may change over the time, producing *non-stationary* noise, and they are not depicted in figure 2.5.

The superposition of all the previous noise sources is called *noise budget*. Of course, this is just a simplistic classification scheme; some ambiguity are possible and, most importantly, these classes are interconnected. An example is the sea and wind activity (environmental), which influences the ground motion vibration (seismic noise), and in turn can excite one of the characteristic resonance frequencies of some component of the apparatus.

Relevant for the statistical description provided in the previous section, and most importantly for the matched filter data analysis technique, is the fact that most of the contributions




Figure 2.5: Reference design sensitivity (solid black line), including all of the fundamental sources of noise, and the expected noise budget for Advanced Virgo (dashed black line). The reference values are for the Dual Recycling configuration (PR+SR) with 125 W of laser input power. Figure readapted from [12].

from the previous noises can be modelled as a stationary and Gaussian process. This constitutes the “background noise” we have talked about, for which it makes sense a representation in terms of PSD, and guarantees the validity of the matched filter technique.

The *stationarity* of the fundamental and (part of) the technical noises is motivated by the fact that, if the detector working point is not altered, by purpose for some realignment and calibration procedures or by some (usually slow) accidental drifts, these are expected to be fixed and under control.²² This is generally not true for environmental noise, which usually varies with different time scales.

The *Gaussianity* of the noise is usually invoked as a consequence of the *central limit theorem* [263]. If the number of individual contributions that compose a particular noise source, like the components thermal vibrations, is large, and there are no dominant contributions among them, or, to be more precise, the variances and cross-correlations of each single contribution are small compared with the total variance, then the so called *Lyapunov’s conditions* for this theorem are satisfied and we can indeed describe this noise as Gaussian. This is also true for shot noise; although this is inherently a Poissonian process, the number of photons is typically so high that the distribution of their number is practically indistinguishable from a Gaussian.²³ However, part of the technical noises, most noticeably the resonances of some components or some not-perfectly filtered calibration lines, are excluded from this criterion and their noises are likely to be non-Gaussian.

To have an idea of how different is the “noise budget” taking into account only the fundamental noises with respect to the actual one (in reality, limited to the differences in the PSD only) refer to figure 1.6, and compare the measured sensitivity curves during the first

²²Instrument malfunctions are a counterexample of technical noises that may manifest in a non-stationary manner, as for example as *glitches*.

²³The typical intensities of the *grey fringe* beam at the detection photodiode are of the order of some mW, and, from (1.61), the corresponding number of photons per unit time is: $N_{\text{Photons}}/T = P\lambda_L/h \sim 10^{30} \text{ s}^{-1}$, for $\lambda_L \simeq 1064 \text{ nm}$.

three observing runs and the projections for those of the next runs. In the latter, only the fundamental noises are taken into account. Besides a vertical offset, there are significant differences, mostly represented by “sharp” *spectral lines* and some broader noise features like “bumps”; a close-up example of them is visible in figure 2.7. The former represent quasi-periodic contributions to the noise. These are usually originating from the resonant excitations of the components of the apparatus, like the “wires” of the *superattenuators* and other mechanical parts, or from the coupling with a periodic signal, such as the 50 Hz power line or a cooling fan generating acoustic and magnetic noise with a frequency related to its angular speed [22]. Some of these line structures may originate from *non-linear* coupling mechanisms between two (or more) noise sources individually characterised by some sharp spectral features. The latter usually manifest in the form of “sidebands” of a central line, and they are visible in particular around the 50 Hz main line and its harmonic at 150 Hz. Other non-linear mechanisms may be at the origin of the bumps in the vicinity of the spectral lines, as in the case of the 50 Hz line. Others are manifestations of non-stationarities, like *scattered light* [264].

2.2.2 Fundamental noise sources and sensitivity

The most of the contributions to the detector sensitivity come from the fundamental noises. Looking at figure 2.5, the main one at low frequency is due to the seismic noise, characterised by the steep descent of the noise ASD. At high frequency, instead, the noise is dominated by the shot noise and the linear behaviour in f , characteristic of the Fabry-Pérot arm cavity; refer to note 21. In the centre part, many effects overlap, most noticeably those due to coating thermal noise. Relevant are also the “sharp”, high amplitude noise features, called *spectral lines*, at 50 Hz and at about 400 Hz. The former correspond to the frequency of the AC main power (60 Hz for LIGO detectors), while the latter correspond to the main vibrational of the mirror suspensions, called *violin mode*. In practice, also the harmonics of the previous spectral lines are visible in the detector sensitivity; see for example figures 1.6 or 2.5, and refer to [262] for a detailed analysis of them.

In this section we will provide some further details on the sources that produce it.

The standard quantum limit

Quantum noise is a consequence of the quantum nature of light used as the “sensing ruler” in the interferometric detectors. It comprises the *shot noise* (dominant at high frequencies) and the *radiation pressure* (dominant at low frequencies). We have already accounted for the former, deriving its characteristic amplitude and frequency dependence; refer to note 21.

The latter is the consequence of the force that the photons exert on the test masses, which, arriving at not equally spaced times, cause a fluctuating displacement of the mirrors. Contrarily to the shot noise, this effect is more important when the power increases since the force is larger as well. Its behaviour depends on the frequency as $1/f^2$, which is the reason why it is more important at low frequencies.

The effects of shot noise and radiation pressure are depicted in figure 2.6, and their combined effect is called the *standard quantum limit*.




Figure 2.6: Depiction of the radiation pressure and shot noise. Picture adapted from [95, Fig. 11.2].

The impact of radiation pressure can be reduced increasing the mass of the mirrors, as it was done in the upgrades between the first generation of detectors and the Advanced ones [95]. For mitigating the effects of shot noise at high frequency, we have already described the injection of squeezed light with smaller phase uncertainty with respect to the normal vacuum [162, 163]. For the future, it is planned to make use of frequency dependent squeezing in order to jointly reduce both of these effects [164, 165].

Seismic noise

Seismic noise is a consequence of the fact that our detectors are not in free fall but lie on ground, and hence they are subject to its motion. At frequencies around 0.1 Hz (*microseismic* peak) this is generated by the sea activity [265] (and the wind), whilst at frequencies of ~ 1 Hz this is mostly of anthropogenic origin [266].

This has a significant impact on the interferometer performance, and if it is not attenuated, it would dominate the sensitivity at low frequencies. To this purpose, in Virgo test masses are suspended to seven stages of inverted pendulums called *superattenuators*. A pendulum acts as a passive filter, attenuating the noise above its resonance frequency f_0 as $(f_0/f)^2$. In a LIGO a similar solution is obtained with four pendulums stages, plus some active control electronics.

The effect of these devices is to attenuate seismic noise at frequencies above 10 Hz to a level sufficient for the required sensitivities.

Newtonian

In the low frequency regime (< 10 Hz), density perturbations caused by a variation of the mass distribution induced by seismic waves in the ground and density fluctuations in the atmosphere, give rise to direct Newtonian couplings to the test masses of the interferometer [267]. This constitutes a “hard limit” to the low frequency sensitivity of detectors, as the gravitational field from these perturbations cannot be shielded. Nonetheless, some advanced subtraction techniques are under study to mitigate its effects by means of an array of auxiliary gravity sensors [268, 269].

Thermal noise

Thermal noise is a consequence of the vibrations of the atoms that constitute a system, as an effect of their temperature. This is described by the *fluctuation-dissipation theorem* [270], and its effect is visible in the “noise floor” of figure 2.5, from the coating Brownian noise, but it can also excite resonance modes of the mirrors and their suspensions, giving rise to spectral lines, with an energy proportional to the temperature of the system. As regards the former, the coating noise is distributed over the mirror surface where impinges the laser; by increasing the beam size, the noise vibrations can be mitigated *averaging* their value. This strategy has been adopted for the Advanced detectors, changing the geometry of the beams in the resonance cavities.

Another mitigation strategy is that implemented by KAGRA taking the test masses at cryogenic temperatures [271]. Moreover, for the same purpose, in the past new designs for the coating layers have been studied by means of *genetic algorithms* [272].

2.3 Stationary noise

In section 2.2.1, when we characterised the various sources of noise on the base of their origin, we have claimed that the fundamental noises, which are typically the best understood ones, can be considered stationary. Nonetheless, control noises, and the most of the technical noises in general, are essentially constant over “long enough” periods of time.²⁴ Under this assumption (to be checked, of course; refer to section 3.1), we know from the Wiener-Khinchin theorem discussed in section 2.1.2, that the corresponding processes can be conveniently described in the *frequency domain*, by means of their PSDs.

In this section we will firstly describe the impact of particular noise features, like spectral lines etc., on various search pipelines. Then we will present some of the techniques used for studying this noise. In general we will limit to study the second order moments only, that is the PSD, except in section 2.3.4, where we will describe some “higher order” spectral techniques.

2.3.1 Effects of stationary spectral noise on search pipelines

Regions in the sensitivity curve with high noise power, typically corresponding to spectral lines and other related structures, can affect searches for long duration signals, like continuous waves and SGWB. For the former, if the nearly monochromatic GW signal has characteristic frequency close to that of a spectral line, once reported to the Solar System barycentre and corrected for the Doppler shift from the motion of the Earth, this line can overlap with the putative signal. This causes periods of data where the corresponding search is almost blind to any underlying GW signal. The same effect is present in both targeted searches and in all-sky surveys.

Searches for SGWB look for correlated excess power in the signals from pairs of detectors; refer to eq. (2.32). Regions corresponding to spectral lines are set to zero in that equation in

²⁴We will clarify momentarily what we mean here with the statement “long enough”. Basically, it refers to the fact that some spectral analysis techniques can be used for such intervals of time.

order to avoid systematic errors in correspondence of lines that are typically non-Gaussian, and also because the coherence between the spectral components of the putative signals in the two detectors is usually negligible with respect to such high power regions. This reduces the search sensitivity by a factor equals to the fraction of frequency bins that have been removed, or “notched”, from the analysis for this reason.

For both CW and SGWB searches, the lists of known spectral lines and instrumental artefacts are created during the commissioning phases that precedes the observing run, and updated constantly. In Virgo, the analysis tool NoEMi has been developed for this purpose [273], and currently maintained by G. Hamming and O.J. Piccinni. Then, depending on the search, these lists are used to: 1) clean the data before performing the analysis, by removing the frequency bands corresponding to spectral artefacts, and interpolating them with Gaussian noise measured in the nearby frequency bins; 2) notch these frequencies, completely removing them from the analyses that are impacted by them; or 3) reject outliers that are clearly caused by the detector artefacts. These “countermeasures” are used to make the analysis pipelines to concentrate the computational resources only on regions of parameter space that are not degraded by those spectral features. If any of these searches returns a signal candidate which does not coincide with any known artefact, more detailed investigations are needed in order to trustworthily exclude the possible presence of other effects due to noise [262].

2.3.2 Investigation of stationary noise

Differently from fundamental noises, the contributions to the detector sensitivity of technical noises can, in some cases, be mitigated by means of specific hardware interventions (substitution of a noisy component or adjustments on the control loops) or with some refined software strategies [274, 275]. To this purpose, it is important to proceed first to their statistical characterisation, including the investigation of their potential causes.

Even before that, the “vulnerability” of the various parts that compose the detector are tested against known noise solicitations. This operation is called *noise injection*, and the analysis of its consequence on the detector is referred to as *linear noise projection* [276, LVC7].

When noise injection are not feasible or previous knowledge about the origin of the noise is not available, one may usually want to make a comprehensive search of the noise sources over the very large variety of sensor channels, typically $\mathcal{O}(10k)$ or even $\mathcal{O}(100k)$ [22, 23], that continuously monitor the status of the detector. The idea is that some of these channels may *witness* the noisy behaviour of the detector, as manifest in its ASD, and consequently on the BNS range, as described in section 1.3.3. This approach is usually referred to as “brute force” and can lead to significant discoveries and insight about what’s going on in the detector. We will also make use of this in section 3.5, when describing characterisation strategies for non-stationary noise.

2.3.3 Coherence analysis

In studying the relations between different detector signals, each modelled as a stochastic process, we will extend part of the notions already discussed in section 2.1.2. We start considering two *stationary* random processes $x(t)$ and $y(t)$. For the purpose of this section, it will be sufficient if they are just wide-sense (or covariance) stationary. In this case, since their means are constants (2.10), without loss of generality we can assume both to be zero: $E[x] = E[y] = 0$. We can define the following *correlation* functions:

$$R_x(\tau) = E[x(t)x(t+\tau)] \quad (2.62a)$$

$$R_y(\tau) = E[y(t)y(t+\tau)] \quad (2.62b)$$

$$R_{xy}(\tau) = E[x(t)y(t+\tau)] \quad (2.62c)$$

where the first two quantities are the *auto-correlation functions* already introduced in (2.11), while the latter is the *cross-correlation function* between the two processes. Now it is necessary to add the labels to declare which stochastic process the various quantities are referred to.

From the stationarity hypothesis, it follows that the auto-correlation functions are even functions of τ , while the cross-correlation function is neither odd nor even, but satisfies the relation

$$R_{xy}(\tau) = R_{yx}(-\tau). \quad (2.63)$$

If the previous quantities are absolutely integrable, namely,

$$\int_{-\infty}^{+\infty} |R(\tau)| d\tau < \infty$$

which will be always true for finite length records, we can define the *spectral density functions*:

$$S_x(f) = \int_{-\infty}^{+\infty} R_x(\tau) e^{-2\pi i \tau f} d\tau \quad (2.64a)$$

$$S_y(f) = \int_{-\infty}^{+\infty} R_y(\tau) e^{-2\pi i \tau f} d\tau \quad (2.64b)$$

$$S_{xy}(f) = \int_{-\infty}^{+\infty} R_{xy}(\tau) e^{-2\pi i \tau f} d\tau \quad (2.64c)$$

From the symmetry properties of stationary correlation functions it follows that:

$$S_x(-f) = S_x^*(f) = S_x(f) \quad (2.65a)$$

$$S_y(-f) = S_y^*(f) = S_y(f) \quad (2.65b)$$

$$S_{xy}(-f) = S_{xy}^*(f) = S_{yx}(f) \quad (2.65c)$$

thus the *auto-spectral* density functions $S_{xx}(f)$ and $S_{yy}(f)$ are positive definite, real-valued even functions of the frequency, while the *cross-spectral* density function $S_{xy}(f)$ is a complex-valued function of f .

We can separate real and imaginary part of $S_{xy}(f)$:

$$S_{xy}(f) = \int_{-\infty}^{+\infty} R_{xy}(\tau) e^{-2\pi i f \tau} d\tau = C_{xy}(f) - iQ_{xy}(f) \quad (2.66)$$

where $C_{xy}(f)$ is called the *coincident* spectral density function, or *co-spectrum*, and $Q_{xy}(f)$ is called the *quadrature* spectral density function, or *quad-spectrum*. Both functions are real-valued, the former being even while the latter odd function of f .²⁵ We will make use of this complex representation shortly; let's consider firstly the information provided by its amplitude.

Magnitude squared coherence

From the previous quantities, we can define the *magnitude squared coherence* (often simply called coherence) [277]:

$$C_{xy}^2(f) = \frac{|S_{xy}(f)|^2}{S_x(f) \cdot S_y(f)} \quad (2.67)$$

which is a measure of the amount of overlap in the spectral structures of the processes $x(t)$ and $y(t)$ [213].

As a consequence of the Cauchy-Schwarz inequality between the correlation functions (2.62), and hence between the spectral density functions (2.64), it follows immediately that $0 \leq C_{xy}^2(f) \leq 1$ for all f . Moreover, as in the case of ordinary correlation coefficient, the closest the previous quantity to 1, the more *linearly related* are the components at frequency f of the two signals. This claim will be motivated with the following example.

A remarkable property of the previous quantity is that it is left unchanged if we make *linear transformations* of the two processes. Indeed, if the signal $y(t)$ is just a "transformed version" of $x(t)$, obtained for example via a linear time-invariant filter of impulse response function $g(t)$,

$$y(t) = (g * x)(t) := \int_{-\infty}^{+\infty} g(\tau) x(t - \tau) d\tau, \quad (2.68)$$

it can be shown that their auto- and cross-correlation functions are

$$R_y(t) = g(t) * R_x(t) * g^*(-t), \quad \text{and} \quad R_{xy}(t) = (g * R_x)(t)$$

and their power spectral densities, obtained from the property of the convolution product under Fourier transform, are:

$$S_y(f) = |\tilde{g}(f)|^2 S_x(f), \quad \text{and} \quad S_{xy}(f) = \tilde{g}^*(f) \cdot S_x(f).$$

Substituting the previous result into (2.67) we obtain that the coherence of these signals is equal to one: $C_{xy}^2(f) = 1$.

In practice it is unlikely to find two channels that are linked by a simple linear transformation as the one in equation (2.68). Usually we may expect some more complex relation

²⁵From their definition, we can notice that $C_{xy} = |S_{xy}| \cos(2\pi f \tau)$ and $Q_{xy} = |S_{xy}| \sin(2\pi f \tau)$, hence they represent the in-phase and out-of-phase contributions to the total cross-spectral density.

between $x(t)$ and $y(t)$. For example, we can consider the simple additive noise model:

$$y(t) = (g * x)(t) + z(t), \quad (2.69)$$

where $z(t)$ is a further source of (stationary) noise, which we will assume for simplicity to be uncorrelated with both $x(t)$ and $y(t)$; the correlated case will be considered momentarily. Repeating the previous computation we find that the coherence of the two signals $y(t)$ and $x(t)$ is now:

$$C_{xy}^2(f) = 1 - \frac{S_z(f)}{S_x(f) + S_z(f)} \quad (2.70)$$

which shows a reduction in its value proportional to the ratio of the PSDs of the signal and the additional noise (that is, the spectral SNR), making it harder, in practice, to uncover the underlying relation (2.69) between the two signals $y(t)$ and $x(t)$.

In the practice of noise investigations, the quantities entering equation (2.67) are estimated, typically by means of the Welch's method described in section 2.1.6.

For investigating this kind of noise it has been developed for Virgo and LIGO by G. Vajente the analysis tool named Bruco, which computes the "brute force" coherence of a target channel (e.g. the strain signal) and a large set, $\mathcal{O}(1000)$, of auxiliary channels over time slices of some hundreds of seconds [LVC2, LVC8]. Then, in each frequency bin, the highest value of the coherence among all the auxiliary channels is selected, and the resulting curve constitutes a measure of the so-called "explained spectral content" of the target signal. A modification to this algorithm to take into account, simultaneously, the contributions from all the auxiliary channels, providing the *total explained spectral content* of the target, has been developed by the author, and will be discussed in the next section.

The noise characterisation method provided by coherence analyses has proven to be the reference technique for studying stationary noise, or, for example, for comparing the explained noise PSD in two different epochs (although in each of them the noise is assumed to be stationary for the whole time of the analysis); if new noise structures appear in one epoch or the other, it is interesting to understand what may have caused this change. This technique is also one of the checks produced in the Data Quality Report (DQR), when one candidate GW event is observed. If the coherence budget during the time of the putative event is different from what it used to be, this may be the manifestation of some unexpected noise contamination, and further analysis are demanded.

In the context of the aforementioned noise injections, a known signal, which should emulate a possible noise source, is used to excite a system, not necessarily the detector output. From its response, if a linear coupling between the injected signal and the output of the system is present, an estimation of the coherence and of the transfer function between the input and output signal is representative of the sensitivity of the system under study to that particular potential disturb.

Multiple-coherence

For the cases in which we want to study the linear spectral relations of one *target* signal, say $y(t)$, with a set of p *auxiliary* channels $X(t) := \{x_1(t), \dots, x_p(t)\}$, we can extend (2.67)

introducing the *multiple* (magnitude squared) *coherence* [278]:

$$C_{Xy}^2(f) := \frac{S_{Xy}^\dagger(f) S_{XX}^{-1}(f) S_{Xy}(f)}{S_y(f)} \quad (2.71)$$

where $S_{Xy}(f) := (S_{x_1y}(f), \dots, S_{x_py}(f))^T$ is the p -dimensional cross-spectral density vector between the x_i 's auxiliary channels and the target $y(t)$, and, analogously, $S_{XX}(f)$ is the p -by- p matrix the cross-spectral densities. The “dagger” symbol, “†”, means the transpose conjugate of the corresponding complex vector or matrix.

To estimate the *unique* contribution of the k -th channel alone to the previous equation, we can define the k -th partial coherence as:

$$C_{Xy}^{(k)2}(f) := \frac{C_{Xy}^2(f) - C_{X_ky}^2(f)}{1 - C_{X_ky}^2(f)} = \frac{R_{X_ky}(f) - R_{Xy}(f)}{R_{X_ky}(f)} \quad (2.72)$$

where $C_{X_ky}^2(f)$ is the multiple coherence (2.71) computed from all the channels in X but the k -th, and $R_{X_ky}(f) := 1 - C_{X_ky}^2(f)$ is the corresponding *residual*, or *unexplained* spectral content of $y(t)$ at frequency f . With the previous normalisation, it is immediate to show that this quantity is bounded by 0 and 1: $0 \leq C_{Xy}^{(k)2}(f) \leq 1$, where the former limit corresponds to a null contribution, either because $x_k(t)$ and $y(t)$ are uncorrelated or because its contribution is already completely encoded in $C_{X_ky}^2(f)$, and the latter to a zero residual after the inclusion of $x_k(t)$ (which however doesn't account for the *total* contribution of it but only to the original part).

Of course, the previous equations make sense as long as $S_{XX}(f)$ is *invertible*. If this is not the case, mathematically this means that the individual spectral contributions from the various channels can't be separated. In practical situations, when working with the estimators for the previous quantities, it is also important that $\hat{S}_{XX}(f)$ is not *ill-conditioned*. Further details will be discussed in section 3.5.3, when we will introduce the multiple linear regression technique. In practice, to avoid this issue we must select in advance auxiliary channels that are not coherent. An alternative to this could be represented by a *principal component* decomposition, as we will describe in the next chapter.

Unfortunately, signals with very similar spectral contents are quite common among all of those that typically monitor the detector, which makes this technique not suitable for “brute force” searches (unless the aforementioned decomposition in uncorrelated components is preposed to the analysis). This constitute a limit with respect to the more immediate simple coherence analysis implemented by Bruco. Moreover, the increased knowledge we can obtain with the latter extended analysis has proven to be no worth the significant augment of computational cost.

This technique has been used by the author to investigate the “rich” spectral structure in correspondence of the 150 Hz resonance in the spectrum of Advanced Virgo [LVC9]. In figure 2.7, the ASD of the LSC_DARM channel ($y(t)$) is shown in correspondence of this it. Besides the main sharp line, there are clearly visible another line at lower frequency and a “bump” centred at 158 Hz. We tried to investigate the relations of this with a selected set of auxiliary channels ($X(t)$) taken from the environmental sensors (ENV_) and the angular




Figure 2.7: Example of some spectral features in the Virgo LSC_DARM channel. From left to right, they are visible a sharp *spectral line* at 138 Hz, a strong *harmonic line* of the main AC power at 150 Hz, surrounded by at least two pairs of sidebands (± 0.5 and ± 1.3 Hz), and a wide *bump* centred at 158 Hz. The red line represents the values $\geq 50\%$ of the explained *multiple coherence* (2.71) with a set of auxiliary channel.

control sensors (ASC_) of the mirror suspensions. Values of the multiple coherence higher than 50% are represented by the continuous red line, equals to $C_{Xy}(f) \cdot \sqrt{S_y(f)}$. The vertical distance between the two curves is what we have called the (square root of the) unexplained spectral spectral content of the LSC_DARM signal with the described model: $R_{Xy}(f)$.

This linear technique was unable to explain the sidebands of the 150 Hz harmonic line, which are likely due to a *non-linear coupling* of the signal responsible to the central line (the mains current channel) with something else. We will discuss the analysis technique suitable for this case in section 2.3.4.

Complex coherence and phase angle spectrum

Usually, only the amplitude (or squared magnitude) of the coherence is used to characterise the relations between signals. Since it is a complex valued variable, like the cross-spectral density $S_{xy}(f)$, further insight can be obtained studying its *phase* value, as we are about to describe.

The complex valued cross-spectral density function (2.64c) can be represented in complex polar notation as:

$$S_{xy}(f) = |S_{xy}(f)| e^{-i\theta_{xy}(f)} \quad (2.73a)$$

where the *absolute value (magnitude)* and *phase angle* are determined, at each frequency, by

$$|S_{xy}(f)| = \sqrt{C_{xy}^2(f) + Q_{xy}^2(f)} \quad (2.73b)$$

$$\theta_{xy}(f) = \tan^{-1} \frac{Q_{xy}(f)}{C_{xy}(f)} \quad (2.73c)$$

The signs of $C_{xy}(f)$ and $Q_{xy}(f)$, and therefore the *quadrant* of the value of $\theta_{xy}(f)$, determines at each frequency f whether $y(t)$ “leads” $x(t)$ or vice versa [213]. We will clarify this statement immediately with an example.

Let us consider the simplest case of a signal $y(t)$ leading a second one $x(t)$, which is just a retarded copy of the former, plus some uncorrelated noise:

$$y(t) = x(t - \tau_0) + n(t), \quad \langle x(t) n(t') \rangle = 0$$

with $\tau_0 > 0$ their relative delay. This implies that:

$$\begin{aligned} S_{xy}(f) &= \int_{-\infty}^{+\infty} R_{xy}(\tau) e^{-2\pi i f \tau} d\tau = \int_{-\infty}^{+\infty} \langle x(t) y(t + \tau) \rangle e^{-2\pi i f \tau} d\tau \\ &= \int_{-\infty}^{+\infty} \langle x(t - \tau_0) y(t + \tau - \tau_0) \rangle e^{-2\pi i f \tau} d\tau \\ &= \int_{-\infty}^{+\infty} (\langle y(t) y(t + \tau - \tau_0) \rangle - \langle n(t) y(t + \tau - \tau_0) \rangle) e^{-2\pi i f \tau} d\tau \\ &= \int_{-\infty}^{+\infty} (\langle y(t) y(t + \tau) \rangle - \langle n(t) n(t + \tau) \rangle) e^{-2\pi i f (\tau + \tau_0)} d\tau \\ &= (S_y(f) - S_n(f)) e^{-2\pi i f \tau_0} = S_x(f) e^{-2\pi i f \tau_0} \end{aligned} \quad (2.74)$$

and therefore, from (2.73), the phase angle spectrum is $\theta_{xy}(f) = 2\pi f \tau_0$, which is a linear function of f , proportional to the delay between the two processes.

Normalising by the square root of the two signals PSDs, we obtain the *complex coherence*:

$$C_{xy}(f) = \left(1 - \frac{S_n(f)}{S_x(f) + S_n(f)} \right)^{1/2} e^{-2\pi i f \tau_0},$$

which improves on (2.70) for the presence of the phase information.

The *causal direction* is estimate to go from $y(t)$ to $x(t)$ if $\tau_0 > 0$ (opposite relation if negative). The notion of “causality” will be thoroughly discussed in the next chapter, as well as the application of this concept to the study of noise.

Of course, this example has little practical relevance, since it is usually not very interesting to study two stationary signals that are simply shifted copies of one another. If they even exist, they are usually already known to have such property. This however constitutes the base for a technique that we have developed to infer both short-time correlation and coherence, in a time-frequency representation of the signals, and their causal relation. Further details about this will be presented in section 3.6.2.

2.3.4 Non-linear noise

In the coherence method described so far, only linear relations between signals are taken into account. However, as visible from figure 2.7, some peculiar structures are present in GW detectors data as consequences of non-linear couplings between different noise processes. In this case, these structures constitute the two pairs of sidebands around the 150 Hz harmonic line, which are not explained by means of the linear coherence method of section 2.3.3. In general, a large variety of possible manifestations of non-linear couplings exists, whose shapes are quite not obvious *a priori*. The simplest, and best understood, example of this kind of noise is *bi-linear noise*, generated by the coupling of two noise sources that jointly affect a third signal. Usually this kind of noise is sub-dominant with respect to the linear one, which, however, in many circumstances is subtracted, making the non-linear one to stand out. Moreover, in the presence of strong spectral lines, the linear approximation is just not sufficient. In GW detectors, the main cause of bi-linear noise is due to the *upconversion* of the low frequency seismic noise, influencing for example the mirrors angular controls, which couples with some narrow-band noise processes, like power lines, calibration lines, violin modes and other narrow-band features in the noise spectrum [22, 274]. We will report here the theory and an example of investigation of this kind of noise.

Let's consider, for simplicity, the multiplicative mixing model of a narrow-band noise, manifesting approximatively as a spectral line of "carrier frequency" f_0 , represented by $e^{2\pi i f_0 t}$, and a generic second noise $X(t)$.²⁶ Their relative phase is unknown, and we can model it as a random variable: $\Theta \in [0, 2\pi)$. Once the experiment has been performed, and the realisation randomly drawn, the value of this phase becomes constant over time. If $X(t)$ is a wide sense stationary process, than also the multiplicative process:

$$Y(t) := X(t)e^{2\pi i f_0 t + i\Theta} \quad (2.75)$$

is (wide sense) stationary.²⁷ Its auto-correlation function becomes:

$$\begin{aligned} R_Y(\tau) &= \left\langle X(t)e^{2\pi i f_0 t + i\Theta} X(t + \tau)e^{-2\pi i f_0(t + \tau) - i\Theta} \right\rangle \\ &= \left\langle X(t) X(t + \tau)e^{-2\pi i f_0 \tau} \right\rangle = R_X(\tau)e^{-2\pi i f_0 \tau} \end{aligned} \quad (2.76)$$

whose Fourier transform yields the PSD:

$$S_Y(f) = \frac{1}{2} (S_X(f - f_0) + S_X(f + f_0)). \quad (2.77)$$

If also the process $X(t)$ is a narrow-band noise with characteristic frequency f_1 , the previous equation shows that the result of the multiplicative model (2.75) comprises two sidebands at frequency $\pm f_1$ with respect to the carrier f_0 , called *modulating frequency*. This is approximatively the situation in figure 2.7, where a carrier signal with $f_0 = 150$ Hz is modulated by

²⁶We recovered for this section the distinction between random variables, labelled with capitols letters, and their realisations in lower-case. This choice is of help in making clearer a point related to the extraction of the PSD of the resulting modulated process.

²⁷This is a consequence of the fact that the p.d.f.s of $2\pi f_0 t + \Theta$ and $2\pi f_0(t + \tau) + \Theta$ are the same, modulo 2π . Notice that without this random phase Θ , the process $Y(t) = X(t) \cos(2\pi f_0 t)$ is not stationary, and the usual notion of spectrum doesn't apply. Its formal definition is a cyclo-stationary process, whose statistical properties vary cyclically with time, and a specific definition of spectrum should be introduced [279].

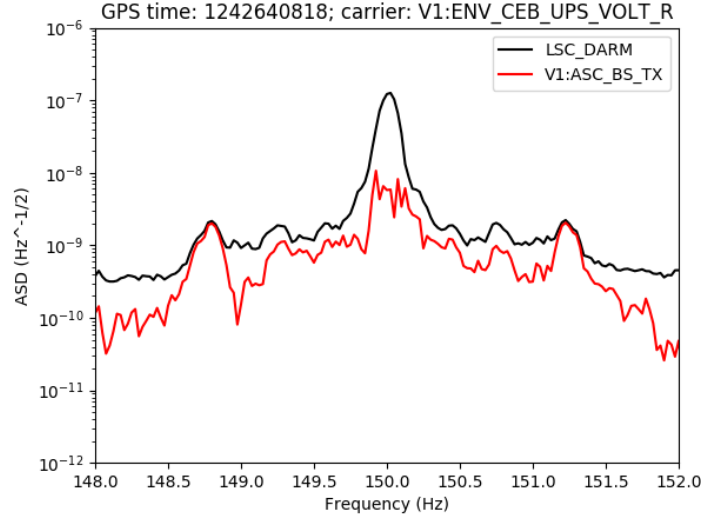


Figure 2.8: Modulated noise model analysis for the sidebands of the 150 Hz harmonic line. This view represents the detail of the central frequencies of figure 2.7. The carrier signal has been chosen to be ENV_CEB_UPS_VOLT_R, representing the main current power in the Virgo central building, while the modulating one, in the sense of (2.75), is the “pitch” angular control of the beam splitter ASC_BS_TX. The coherence of the resulting signal and LSC_DARM is represented by the red line.

two multiplicative noise processes, respectively with frequencies $f_1 \approx 0.5$ Hz and $f_2 \approx 1.3$ Hz. Hence, from a detector characterisation point of view, the task is to find some channels peaked at these frequencies.

If the spectrum of the multiplicative process $X(f)$ is more “broad-band”, this results in the characteristic, quasi symmetric “bump” of excess power, centred around the spectral line. This is also partially visible around the 150 Hz line, and even more clearly around the 50 Hz line (whose central contribution is subtracted, though): figure 2.2.

To study this kind of noise, B. Patricelli and G. Cella have developed a data analysis tool called MONET (MODulated NoisE Tool) [LVC10], which computes the coherence between a target signal and a “synthetic” one, constructed multiplying a signal representing the carrier (or even a representative sinusoidal one) with a set of modulating auxiliaries. In figure 2.8, we reported the analysis of the same central noise feature of figure 2.7 making use of this tool, and documented in [LVC11]. The carrier channel has been chosen to be one measuring the main power current in the Virgo central building (ENV_CEB_UPS_VOLT_R), representing the “culprit” of the 150 Hz line in LSC_DARM, whilst the set of modulating channels are taken among some of the alignment control channels used for the mirror suspensions of Virgo.²⁸ As visible from the red curve, the sideband at ± 1.3 Hz is explained with the model (2.75) using ASC_BS_TX channel as modulator. Only part of the other sideband is explained with this method and the previous choice of signals.

²⁸A similar result is obtained choosing, as a carrier, the “artificial” signal constituted by a sinusoid of frequency $f_0 = 150$ Hz. Of course, using actual signals may shed light on not evident manifestation of this multiplicative coupling in other regions of the spectrum. This is the preferred strategy when we have prior inform, for example from previous application of the coherence method and the identification of the potential “carrier” signal. Refer to figure 2.7 and the corresponding discussion.

2.4 Non-stationary noise

The analysis methods presented in the previous sections rely on the assumption that the processes under study are (wide-sense) stationary, and ergodic, during the observed period of time. Concepts like PSD, the coherence function (2.67), and the matched filter technique are not defined for non-stationary signals. However, as a matter of fact, GW detectors data present non-stationarities, whose extent and influence must be understood and taken into account. Many attempts to extend these analysis to non-stationary signals can be found in the specialized literature [280–282]. We will discuss about the characterisation of non-stationary noise, as well as the implementation of some of these strategies in the next chapter. In this section we will just introduce some of the “phenomenology” with which non-stationary noises typically manifest in GW detectors data, and the methods commonly adopted to cope with it.

We already presented the “mathematical” definition of (non-)stationarity in section 2.1. Here we will elaborate further with an “operative” classification of the kind of non-stationarity that may be present in our data. In GW community we commonly divide non-stationarities into *fast transients*, colloquially referred to as *glitches* [38, 39], manifesting as short duration “bursts” of excess power, and *slow noise transients*. There is no net threshold between the two classes. From an “experimentalist” point of view, one can choose as a typical time-scale that of variation of the position of the mirrors and the optical benches that constitute the interferometer.²⁹ These components are suspended to the chains of pendulums that constitute the *superattenuators*, whose cutoff frequency is at about $\mathcal{O}(1)$ Hz, as visible from the steep drop off in the sensitivity curves: figure 2.5. Then, according to this criterion, one can classify as a glitch the transient noise that has typical duration shorter than, or comparable to, a few seconds, and slow non-stationarities everything occurring over longer time scales.

There is a second viewpoint, which yields to a similar conclusion, and is that of the “GW data analyst”. We already made use of it in section 2.1, when talking about the matched filter technique and spectral estimations. This is based on the effects that noise non-stationarities produce on the search pipelines for various sources. Glitches are those “fast” noise transients that could be misleadingly confused for GW transient signals, like CBCs and bursts: *false positive*. Slower non-stationarities has an indirect effect on these searches, like the mis-estimation of the background noise PSD, which results in a lower value for the detection statistic SNR (refer to the discussion in section 2.1.5): *false negative*, and systematics in the parameter estimation with the matched filter technique.³⁰ This sets a scale of reference that is again of the order ~ 1 sec, where the typical time of BBH events only has been taken into account since the longer BNS chirp has a characteristic waveform and “shape” in a time-frequency map that no other noise sources are known with similar features; see eq. (1.48). This shades importance on another aspect, that is, the classification of glitches on the base of their morphology. We will come back to this aspect later.

²⁹Perhaps D. Hoak should be credited by the author for this idea, since the first person who introduced me to this concept, with all of the consequences about “where to look for” when doing noise characterisation.

³⁰The systematics in parameter estimations due to slow non-stationarities are visible, for example, for GW170814 in figure 3 of [13], where the 90% credible area obtained by fast sky localisation and that obtained after noise removal and improved detector calibration are reported, and only partially compatible.

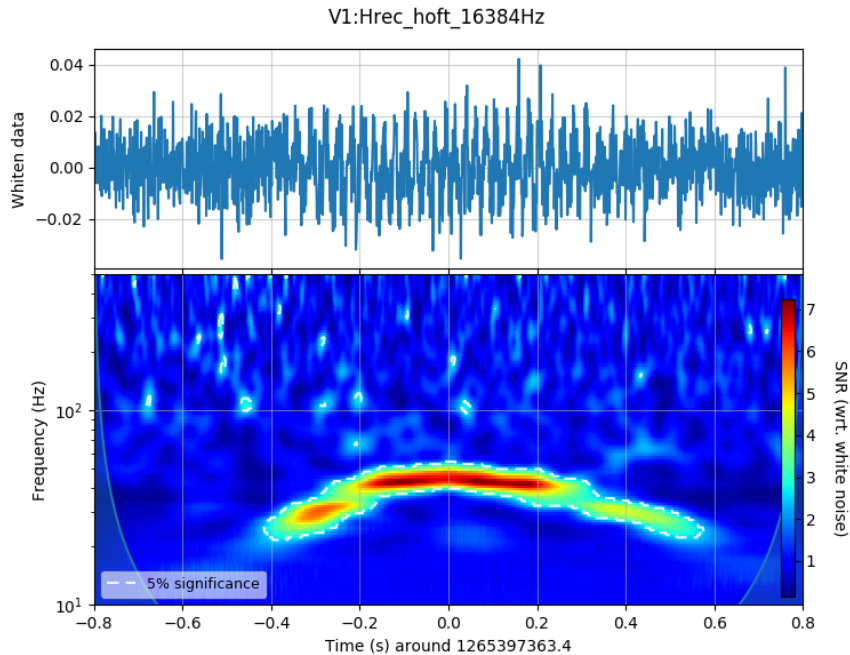


Figure 2.9: Characteristic example of *scattered light glitch* in Advanced Virgo data. The top plot represents its *whitened* time series, where the “signature” of this kind of noise is visible oscillations with increasing frequency towards the centre of the figure. The bottom plot is a time-frequency map representing the *wavelet transform* of the data. Its definition, choice of normalisation and significance levels (dashed white contours) will be presented in section 3.6.2.

2.4.1 Glitches

Probably “the most famous glitch” in GW literature is that affecting LIGO Livingston data during GW170817 event [14], and reported in figure 4.1. This high-amplitude transient originated from a short-duration excess noise in the DARM control loop, which is transferred then to the calibrated strain signal sensing the GW [198]. The morphology of this glitch, in a time-frequency representation called *Q-scan*, which will be described in the next chapter, is clearly different from that of the BNS signal upon which it is superimposed. Actually, it is different from the most (all, to the best of our knowledge) of the astrophysical waveform predicted by GR. This difference is exploited as a method for distinguishing, at least in first analysis, candidate events from “known” classes of glitches, from their shape [283]. For example, as discussed in section 1.4.1, the cWB data analysis pipeline [21] implement it in its search for coherent signals between the detectors, which don’t resemble the shape of known glitches. Moreover, this is implemented in the DQRs as a preliminary classification based on a *Convolutional Neural Network* (CNN) model to generate glitch categorisation confidences for candidate events.

Fast glitches, with time scales of the order of 1 msec, are usually caused by the electronics in the sensing chain, while slower ones can be caused by optics misalignments, or external disturbance (ground motion induced by anthropic activity, thunderstorms, etc.). A specific manifestation of the previous disturbances is through *scattered light* glitches [22, 23]. This is a non-linear effect caused by some unfiltered excess ground motion or acoustic noise coupling

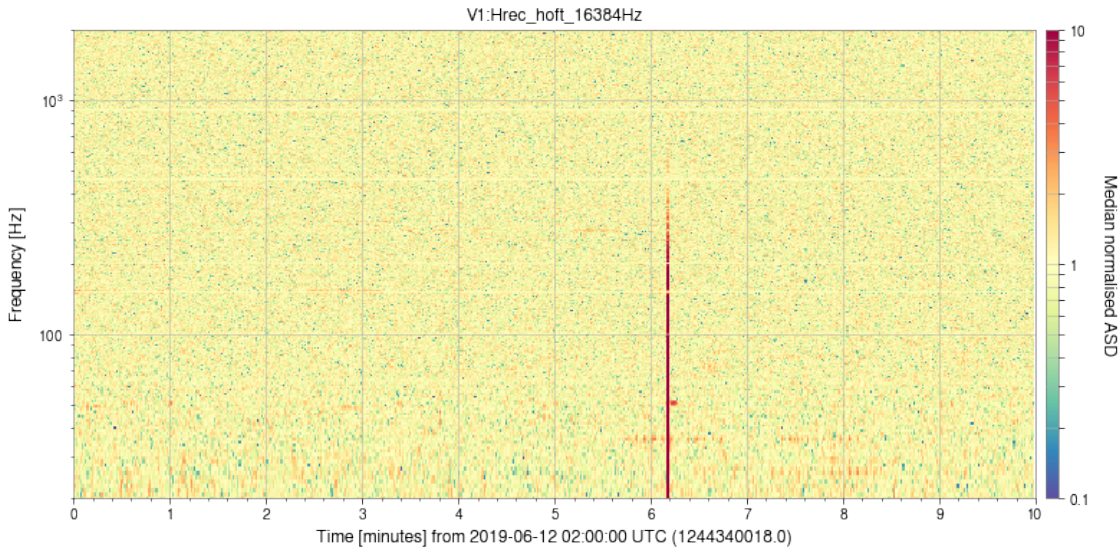


Figure 2.10: Median normalised spectrogram of the ASD corresponding to figure 2.3.

in to the detector, which produce the movement and misalignment of some of the mirrors and optical benches. This, in turn, causes part of the laser light to be scattered off moving surfaces and coupled back to the main beam [258]. The resulting glitches are characterised by arch-like shapes, as visible in figure 2.9. It is important in this case to understand the part of the detector which are more vulnerable to the disturbances causing these glitches. This is done by means of the noise injections described in section 2.3.2.

More details on the characterisation of glitches will be discussed in the next chapter. We will report in section 2.4.3 what are the common, and immediate, “countermeasures” to cope with glitches, and their consequence on GW searches.

2.4.2 Slow non-stationarities

Slow non-stationary noise can be usually associated with (external) environmental disturbances, operation cycle of some apparatus or tidal fluctuations [22, 23]. Its characterization goes then through identifying the typical time-scales of the variations, the spectral shape, including the characteristic frequencies (if any), and how sudden is its onset and/or fade off. These operations are conveniently done inspecting the signal spectrogram, already introduced in figure 2.3. A useful way to make the non stationarities to stand out is by normalising it, along its time axis, with the value of its mean-median ASD (2.57). The idea behind this strategy is the same already discussed for that modified spectral estimate. An example of this kind of map is shown in figure 2.10, where the glitch at 2:06:10 UTC is clearly highlighted. Moreover, we can notice that the spectral lines, in particular that at about 470 Hz, are highly stable: uniform yellow colour. Some other weaker non-stationarities stand out from this map at low frequency; these were not evident in the classical spectrogram representation in figure 2.3. We will discuss a method for efficiently detecting them in section 3.1.2.

Usually, different noise sources can affect the target at different characteristic frequencies; for example, magnetic noise associated with the electric main is centred in a band

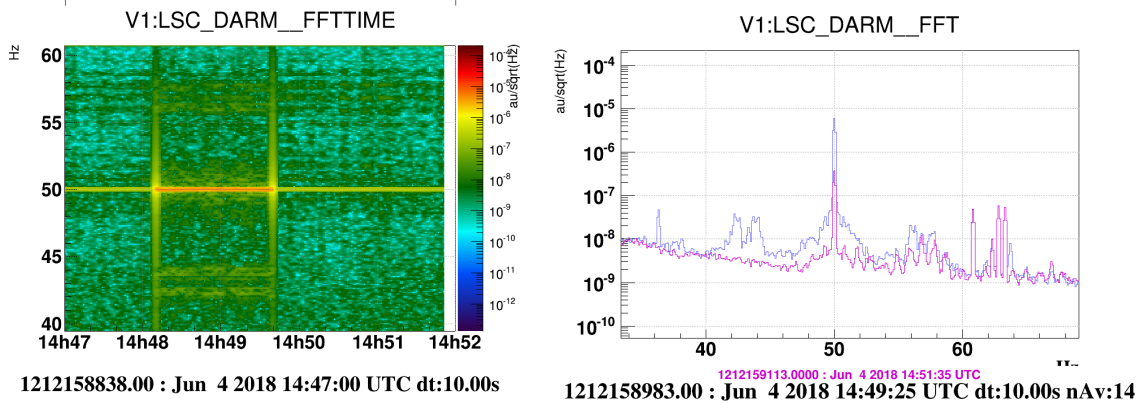


Figure 2.11: Example of amplitude non-stationarity in a frequency band around 50 Hz of the LSC_DARM channel. The left-hand side plot shows its spectrogram, while the right-hand side one its amplitudes at two different times, corresponding to during the noisy phase and after it. Figures obtained with `Virgo dataDisplay [LVC13]`.

around 50 Hz (see figure 2.11), *scattered light* has typical noise contributions below some tens of Hertz, glitches due to the misalignment of the *pre-mode cleaner* have typical frequencies at 1 kHz and above, etc. If these noises involve mostly variations of the power intensity in the specific frequency band, as in the example in figure 2.11, we talk about *amplitude non-stationarities*. It is useful to characterise them by means of their *Band-limited Root Mean Square* (BLRMS), that is, the time series corresponding to the total energy of the signal in a specific frequency band. This approach in Virgo was pioneered by G. Vajente with the NonStatMoni analysis tool [28, LVC12]. It has been rebuilt and improved with some modification that will be described in the next chapter 3.5.4, and now it constitutes the “pre-processing” stage of some of the tools developed by the author.

The non-stationarities can also affect the characteristic frequency of some spectral features. For example, particular spectral lines have proven to changing their frequency over time, as can be seen in figure 2.12; this kind of noise is therefore colloquially referred to as *drifting* or *wandering line*, and is a particular kind of spectral noise [22].³¹ These frequency variation can be due to resonant modes of the mechanical components that are in the optical path of the beams, which may change over time. We would like to investigate these features relating the characteristic frequencies of these lines, for example corresponding to their maximum values changing over time, with the information provided by the auxiliary channels. In the example of figure 2.12, it is shown that the evolution of the frequencies of the line in LSC_DARM was correlated with the temperature measured at the ring heater in correspondence of the West-end mirror (ENV_TCS_WE_RH_TE), and at the *filter-7* of the corresponding suspension (ENV_WE_F7_TE1) [LVC14].

Frequency non-stationary noise has been one of the main issues in June 2017, right before Virgo joined the O2 science run [LVC15]. Also during O3 some wandering lines were present, and in particular the one at reference frequency 83 Hz that has been the subject of extensive investigations by the author [LVC16], which will be described in appendix A.3.

³¹Usually, we prefer the adjective “drifting” when the origin of the line is somehow known, and the varying frequencies assumed by the line are driven by it. Instead, we call “wandering” lines those that appear to wander in the spectrum without an apparent reason. We discuss both cases in section 3.4, and elaborate the study of an example of the latter in appendix A.3.

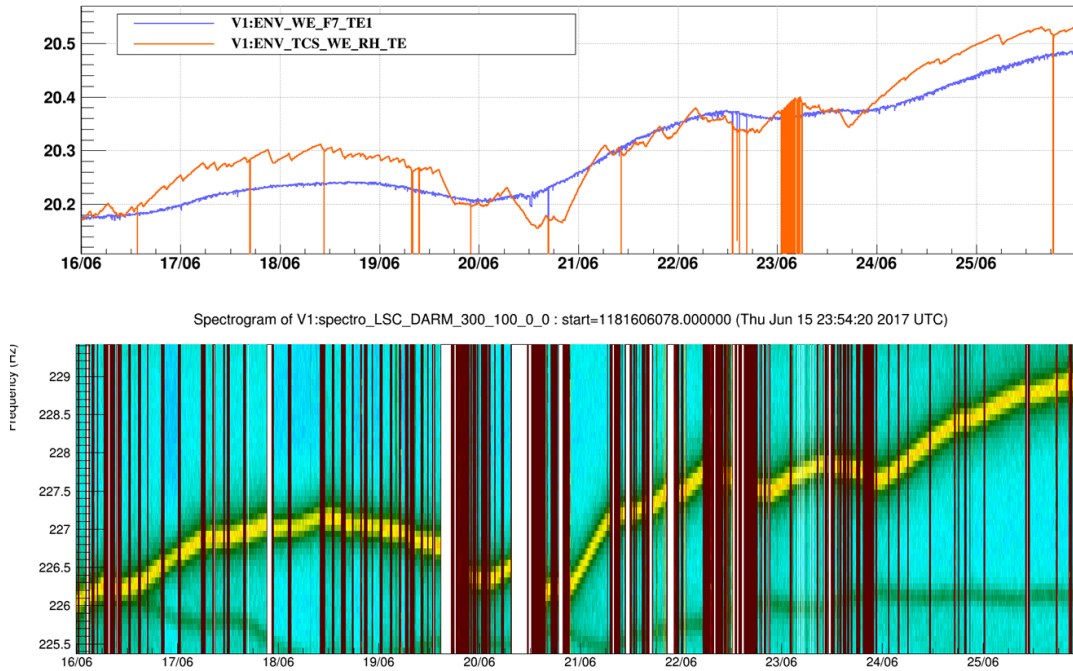


Figure 2.12: Example of frequency non-stationary noise in Advanced Virgo data during the pre-O2 phases. Bottom plot: spectrogram of a “drifting” or “wandering” line (a fainter one is also visible in green in the lower part of the spectrogram) in the Advanced Virgo LSC_DARM channel. In the top plot two thermometer channels exhibiting a similar behaviour. Image taken from [LVC14], courtesy of I. Fiori.

For this purpose, we have developed a tool that extract the information about these lines, and correlates it with that available from the auxiliary channels.³² We will describe the “line tracker” in section 3.4, and the tools developed for correlation analyses in 3.5. Also, the aforementioned NoEMi tool [273] can be useful to this purpose, especially for very slow variations with time, for the way it is conceived to the purpose of continuous wave searches. Indeed, it has demonstrated not to be sensitive to variations frequency taking place in few hours,³³ hence the relevance of our tool as a useful complement.

Mixed situations, where non-stationary noise induces both frequency and amplitude variations, are not uncommon and usually imply much complex underlying coupling mechanisms. “Ad hoc” applications of the available analysis tools, including the ones that we shall describe, can be performed.

2.4.3 Data Quality flags

Non-stationarities are unavoidably present in detectors data, despite the efforts in trying to prevent them by intervening on the systems and the controls that are more likely to produce them. Removing them by performing some sort of subtraction is not a simple task, as we will comment in chapter 4, although various groups have started investigating new approaches

³²The first attempts to this kind of study have been successfully conducted by S. Koley, in the case of the aforementioned noise during the pre-O2 phases [LVC14, LVC15]. Based on this, a similar analysis tool has been developed, in more recent times, by B. Swinkels [LVC17]. This differs from ours for the way the time series of the frequency maxima is extracted. A comparison of the two methods will be discussed in section 3.4.

³³We will further comment this in appendix A.3, in relation to the study of a specific noise line in O3 Virgo data.

based on state of the art techniques in Machine Learning [284, LVC18, 285]. Hence, the most immediate intervention to mitigate their effect is to exclude the corrupted data segments from the analysis, imposing *Data Quality (DQ) flags* that constitutes a sort of “veto”. In the “Advanced Virgo DQ model”, two kinds of vetos exist [LVC1], which we are going to describe.

“Interferometer Status Flags” are mostly generated by the automatic processes that monitor the status of the subsystems and the control loops of the interferometer. A combination of them is used to determine the segments of data where the interferometer is working in nominal, stable conditions, and then producing valid data for GW searches or, to be more precise, good for evaluation by the next level of DQ. This process is also supervised by the operator and the run coordinator in charge, also ensuring that there are no human interventions or other “known” interferences occurring on or near the instrument. If the previous checks are passed, the interferometer is said to be in “Science mode” (Virgo terminology) or “Observing mode” (LIGO). This kind of flag is produced in real-time [22, 286].

Of course, the interferometer status is independent on the particular GW search and pipeline used (but not the opposite). To this purpose, there is another level of DQ flags, marking specific conditions of the instrument, or subparts of it, which are likely to affect the quality of the searches. Since periods of noisy data will affect each type of analysis differently, these are separately defined for each of them, like those for CBC searches, bursts, CW and stochastic backgrounds; refer to section 1.2 for the description of these sources and to 2.1.4 for the Data Analysis basis of some of them (also mentioned in 1.4 in the case of CBCs), or to [17] for the details on the gating procedure adopted by PyCBC. Contrarily to the previous flags, these are only partially produced online. Some of them also require longer analysis times, and are added later.

The first, and arguably most severe flag of this kind, is the so called “CAT1” flag. Failing the corresponding check indicate some *known* critical issue in a detector component, which is not operating in its nominal conditions and compromising the quality of the data, or at least of a part of it. This corresponds to a major issue for GW searches, and it is almost identical for all the sources, and pipeline independent.

Two less severe flags are the so-called “CAT2” and “CAT3” flags. These differ among different kind of searches and pipelines. The former is referred to some issues in the strain channels, with *known* origin, that are likely to affect the corresponding analysis. This may include bad weather conditions and high micro-seismic activity, visible in the higher level of noise in the low frequency region and, most importantly, in the higher presence of scattered light glitches entering the sensitivity band, as described in section 2.4.1, plus other effects (like the ones described in appendix A.1). The latter flag is usually referred to some statistical coupling between the strain channel data and that of some other auxiliary sensors, which are likely to affect the analyses but in not a fully understood way.

Of course, the previous flags are used in a cumulative way: what doesn’t pass CAT1 automatically does the same for CAT2 and CAT3. However, having different definitions of them for the various groups allows them to be independent, as well as the data at their disposal for searches. This constitutes an efficient usage of all the available information gathered by our detectors, although it comes without saying that form a Commissioning

and Detector Characterisation point of view, we should work for maximising the quality and amount of data for all of them. This point will be elaborated further in the next section.

2.4.4 Effects of vetoing in GW searches

The previous DQ flags have proven to be of fundamental importance, in terms of the benefits they provide to the search sensitivity, especially for transient searches. For example, one search for compact binary coalescences had as much as a 90% improvement in its sensitivity in parts of O1 [24]. This however comes at the price of “throwing away” some potentially valuable data, as we will elaborate in this section. Moreover, for the previous reason some noise-subtraction and deglitching procedures are under study by different groups (including ours), as we will describe in chapter 4.

The discussion presented in this chapter had the purpose of convincing why having a good strain sensitivity curve is important but not sufficient. In order to maximize the chances of detection, it is mandatory that the detector has a good amount of full functioning operating time. This quantity has been called the detector’s *duty cycle*, as introduced in section 1.3.4. During the Advanced Virgo O2 run this was about the 86% of the total time, and about 60% for the two LIGO detectors [1]. However, such amount must be slightly reduced, to the purpose of the Data Analysis, by the further application of the DQ flags described before. In the case of Virgo, this resulted in 80% of the time with valuable data for GW searches.

For the most of the analysis however, an even more important aspect than the previous one is the actual coincident time between two, or possibly three detectors, for achieving good sky localisations by triangulations. We have already touched this aspect in section 1.3.4 when describing the importance of having a *network* of multiple GW detectors constantly monitoring the sky. This coincidence time reduces further the total amount of valuable data for GW searches.

For transient searches, the previous aspect has mostly relevance for the total number of events detected during an observing run. For persistent signals, like CWs and SGWB, where the detection statistics take advantage of the integration time, it is important to avoid also spectral features containing dominant non-Gaussian components, which would introduce systematics in the analysis if not otherwise corrected. We have already presented the characterisation of this kind of noise in section 2.3.3, while the detection methods for identifying this particular issue will be described in section 3.3.1.

As a result, for the isotropic SGWB search in the Advanced LIGO first observing run, 35% of the time series data T (“delta-sigma” cut [LVC19], on the variance of the detection statistic described in (2.32)) and 21% of the frequency domain observing band Δf were removed from the analysis [287]. The sensitivity to this background can be estimated, in cosmological units of energy density Ω_{gw} , from [74] as:

$$\Omega_{\text{gw}} \propto \frac{S_n(f)}{\sqrt{T \Delta f}}.$$

The previous vetoes have then implied a reduction of sensitivity of about 28%. In the second observing run these numbers were significantly reduced to 16% T and 15% Δf , thanks to the improvements to the detectors [138]. However, this issue come back more severe in the

first part of O3, where due to non-stationarities in the data a relevant amount ($\sim 57\%$) of the times were discarded from the analysis [LVC20]. This motivates the very high importance that Detector Characterisation studies aimed at non-stationarities identification, and the investigation of their causes, have to this kind of search.

Continuous GW searches share the same dependence on the observation time T but, especially for targeted searches, where the emission of known pulsars is investigated, they are mostly sensitive to narrower frequency bands [108]. The effects of non-Gaussianities, namely those associated to spectral lines, have been previously discussed in the context of stationary noise in section 2.3.1.

3 Methods of investigation of non-stationary noise

Statistical Data Analysis is easier when the underlying process that generates the data is stationary. If this is the case, the Ergodic Theorem (section 2.1.2) is often taken into account, allowing to substitute ensemble averages with time ones, and in this way defining the estimators for the statistical properties of interest [6, 213]. As presented in section 2.1.4, this is a common practice in Gravitational Wave (GW) data analysis, mostly based on some fashion of the matched filter technique, whose filter function is made assuming to know the statistical properties of the detector noise from which to extract the astrophysical signal. It is crucial for the validity of this technique and the consequent detection significance and parameter estimation that the hypothesis underlying the noise distribution are satisfied. In the previous chapter, we have discussed when this may not be the case and the consequences on various kind of GW searches; sections 2.3 and 2.4.

To this purpose, within the LIGO Scientific Collaboration and Virgo Collaboration (LVC), the Detector Characterisation group has the role of checking the validity of the previous assumptions, and the investigation of the causes that may have led these to break down. This has to be done before any inference on the possible presence of a GW signal in the data is made.

In this chapter, we will describe the methods developed to identify non-stationarities and non-Gaussianities in the data, and then the analysis strategies to investigate their causes. In section 3.1 we will present some tests for non-stationarity, including an original one developed by the author [LVC21, LVC22], which improves upon other test known in the literature and commonly used in GW searches. Since the beginning of O3b (the second half of the third Advanced detectors observing run), this test has become part of the standard checks present in the Data Quality Report (DQR) for the preliminary validation/retraction of candidate events. Connected to that, we will also present in section 3.3 what is known in the LIGO-Virgo community as the “Rayleigh test” for Gaussinity [LVC23].¹ We will provide an overview of the method, plus some proposals of improvements [LVC24].

Once the non-stationary and non-Gaussian behaviour of the data has been identified, the next step is to try to correlate it with the signals from the thousands of auxiliary sensors that monitor the detector and its environment. This analysis will be described in section 3.5 and 3.5.3, where we will present the *rebuilt* and extension of the analysis tool called NonNA, developed by G. Vajente [28, LVC2].

¹Although its base version is now standard in many GW Data Analysis libraries, like `GWpy` [288], and in the DQR, it was surprisingly hard to find any documentation for it. Apparently, the first to introduce it were S. Finn, G. Gonzalez and P. Sutton in [289]. Some reference to this in a collaboration article are also present in [290].

In section 3.6 we will present a possible extension of the previous analysis, based on correlation and coherence, to the case of glitches. This will make use of the *wavelet transform* representation of the data [291]. We will also introduce an original “instantaneous” causality statistics based on the phase angle spectrum of the *wavelet coherence* between two signals. While all the methods based on correlations and coherences are simply measures of the similarity in the time or frequency domain between two signals, the latter is an attempt to infer whether a *causation* relation between them is present. This concept will be elaborated in section 3.7, where we will present the application of three methods for the assessment of causality on GW detectors data.

3.1 Detection methods of non-stationary noise

In order to apply the standard data analysis techniques described in section 2.1.4, we must ensure that the data (the noise component) that we are analysing is stationary. Or, to be more precise, we have to test whether we have reasons to *reject* the null-hypothesis that the available data sample is representative of (a realisation of) a stationary stochastic process, as defined in section 2.1.2. Actually, this implicitly assumes the “zeroth order assumption” that any given sample record will effectively reflect the (non-)stationary character of the random process in question [213].

In signal analysis literature, many tests are available to investigate specific (non-)stationary models, like those described in section 2.1.3. Some examples are the Augmented Dickey-Fuller (ADF) test for *unit roots* [292] or those based on AR and MA models discussed in section 2.1.3. In general, these are too restrictive for the analysis of interferometers data, and only seldom used in practice, except than for checking specific hypotheses about the noise behaviour. Moreover, due to the very rich frequency structure of the detectors data, and to the presence of processes that are known to have influence only on restricted bands of the spectrum, it is most convenient to characterise the non-stationarities in a time-frequency map. We have already discussed how an immediate information about this comes from the inspection of the spectrogram, representing the signal energy content in each time-frequency bin. In particular, the median normalised version of it, shown in figure 2.10, highlights the deviations from the “central” tendency of its distribution. Basically, what we want is to quantify this “deviation”, formalising it in a statistical hypothesis test.

In this section, we will describe some of the statistical tests currently adopted in GW Data Analysis to check stationarity. For the reasons previously explained, all of them rely on some statistics based on the estimated PSD, or related spectral representations, of the data over short periods of time. Loosely speaking, the tests aim to verify whether this estimates (don’t) remain constant over time. In this sense, they can be considered as test for “covariance” or “wide-sense” stationarity, according to their definitions of section 2.1.2.

All the concepts that we are going to discuss momentarily are closely related to the treatment on spectral estimation of section 2.1.6. The basic ideas behind the tests are already present there, at least in some crude form.

3.1.1 Stationarity of the PSD: the Priestley-Subba Rao test

The *Priestley-Subba Rao test* (PSR), from the authors of the concept of *evolutionary spectra* [223], is based on the time-varying Fourier spectrum of a signal [293]. Without introducing the formal definition of processes described by evolutionary spectra, for our purpose, we can take the series of PSD estimates (2.49), over various data segments, to be an effective representation of it.

If the process that generated the data sample is stationary, the previous estimate shouldn't vary with time, except for its inherent statistical uncertainty, as reported in (2.53). Then, the PSR test aims to reject this hypothesis studying "how non-constant" that function of time is. It does this by looking at the test statistics given by the logarithm of the previous estimator:²

$$Y^{(i)}(f) = \log \hat{S}^{(i)}(f), \quad i = 1, \dots, M,$$

with i the "temporal" index that labels the segment where the PSD has been estimated. Here the logarithm acts as a "variance stabiliser transformation" [294], allowing to focus on changes in the mean structure of Y . Asymptotically, we may expect it to be a consistent estimator of the logarithmic PSD: $E[Y^{(i)}(f)] \approx \log S^{(i)}(f)$, and the variance of $Y^{(i)}(f)$ can be proved to be approximately constant [295].³ These actions allow us to write $Y^{(i)}(f)$ as a linear model with constant variance and test the constancy of $S^{(i)}(f)$ using a standard one-way *analysis of variance* (ANOVA) on the residuals: F -test statistics.

This method has three advantages: firstly, as requested, it tests the stationarity of the PSD estimate frequency-by-frequency; it doesn't make any assumption on the data distribution,⁴ which is useful if we want to test Gaussianity independently; and, lastly, the ANOVA provides p -values, on the base of which we can decide to reject the null hypothesis of stationarity or not. The disadvantage is that some of the assumptions that the authors, and the references therein, make are valid only asymptotically (or not thoroughly justified at all).⁵ Some tests on simulated models, taken from 2.1.3 (and partially already included in [293]) have shown good performances, but, on real interferometer data, at fixed data sample size its sensitivity has proven to be worst than that of the test that we will discuss momentarily.

²Notice that the $Y^{(i)}(f)$ are meant as stochastic variables, hence we have indicated them with upper-case letters. Refer to the convention discussed in the Notation section.

³As already discussed in section 2.1.6, and in particular for the periodogram estimator (2.52), the standard deviation of the estimate is proportional (or even equal) to its mean: *fixed relative error* estimator. It is often desirable to "disentangle" the two quantities. In the case of an exponential distributed variable S (refer to the comments after equation (2.59)), with scale parameter $2\sigma^2$, also equal to its mean and standard deviation, the variable $Y \equiv \log S = \log 2\sigma^2 + \frac{1}{2\sigma^2}(S - 2\sigma^2) + \mathcal{O}((S - 2\sigma^2)^2)$ has expectation value $E[Y] \approx \log 2\sigma^2$ and variance $\text{Var}[Y] \approx 1$, where a Taylor expansion of S with respect to $2\sigma^2$ has been performed (*delta method* [296]), and all the moments of order two or higher have been omitted. This is justified by the fact that $E[S^n] = n!(2\sigma^2)^n$, as characteristic of an exponential variable, and $\sigma^2 \ll 1$ (as typical of h_{rec} or, in general, after a suitable rescaling).

⁴Although it ignores any distribution moment beyond those of order two, as implicit in the fact that it is based on the PSD estimates.

⁵Specifically, this is a consequence of the fact that this method actually implies a χ^2 test on the linear model residuals for $Y^{(i)}(f)$. A comparison of this with the Kolmogorov-Smirnov test that will be discussed in the next section, including the evaluation of their sensitivity in relation to the sample size, can be found in [297].

3.1.2 A BLRMS based test for noise stationarity

We will describe here the statistical test developed by the author to check the stationarity of the data. This is based on the empirical distribution of the signal BLRMS, whose stationarity is verified, in each frequency band, comparing their empirical distributions in two neighbouring data segments, by means of a two-sample Kolmogorov-Smirnov test [298].

First of all, we proceed to describe the method for computing this BLRMS, which is based on that developed by G. Vajente in [LVC12], and recently improved in collaboration with him [LVC21]. The content of this section has been taken from an article in preparation by the author [37].

Method of computation of the BLRMS

The BLRMS is (the square root of) the total power of a signal in a limited frequency band:⁶

$$BLRMS(t; [f^{\min}, f^{\max}]) = \sqrt{\int_{f^{\min}}^{f^{\max}} S(f; t) df} \quad (3.1)$$

where $S(f; t)$ is the signal PSD referred to the time t [223]. In practice, without recurring to the theory of evolutionary spectra, we will be interested in its estimate from the value of the PSD in the i -th segment: $\hat{S}^{(i)}(f)$, defined in (2.49).

There are two reasons to prefer, in the context of GW detectors data and non-stationary noise analysis, the previous quantity with respect to the usual ASD estimate. Firstly, as already noted and clearly visible in figure 2.7, the spectral properties of the data are very rich but also very diverse. There are regions with characteristic features, like lines, bumps and structures due to non-linear couplings, where we would need a better frequency resolution, and regions where the spectrum is *flatter*. For the latter, averaging over a certain frequency range has a variance reduction effect similar to that obtained with the Welch method (equation (2.55)), averaging in time. In practice, if we don't need to resolve detailed frequency structures, we can average on fewer time segments and more frequency bins to obtain a similar variance for the corresponding estimate.⁷

Secondly, we know from the discussion in sections 2.1.3 and 2.3 that many kinds of disturbances affect specific frequency bands only.⁸ If we know a band division of the spectrum, for example provided by somebody in charge of the commissioning of the detector, where different kind noise sources have their support, we can use this to compute the BLRMSs. This is more practical than retaining one time series $\hat{S}_i(f)$ per every frequency in the spectrum. Moreover, it has the advantage that once a non-stationarity stand out in a specific band known to be affected by a particular disturbance, the person who is investigating the noise is immediately triggered about the possible cause. For example, the band around 30 Hz in LIGO Hanford was known to have some disturbances related to scattered light at the beginning of O3, the 60 Hz main line also has some glitches in LIGO Livingston, while the same

⁶With this normalisation, the BLRMS has the same units of the signal. If it has to be interpreted as an "average" ASD, one can divide (3.1) by the frequency interval $f^{\max} - f^{\min}$.

⁷This is certainly true for a white noise process. For coloured ones, and more specifically in the presence of spectral lines, the latter will dominate the estimate and the consequent variance. We will discuss shortly the modification implemented to deal with them.

⁸With the exception of fast electric glitches, which are broadband.

happens in Virgo around 150 Hz; refer to the noise study in appendix A.1 for an overview of the latter. In the case this prior information is not available, one can provide a “sound” band division of the spectrum.

However, these averages have the drawback of “masking” the frequency non-stationarity of a wandering line, like that in figure 2.12, if this is entirely contained within a single band. For this reason, the BLRMS method is best suited for studying amplitude non-stationarities; we will describe how to attack wandering lines in section 3.4.

As described in [LVC7], there are two methods to compute the BLRMS. One operates the frequency selection directly on the PSD, and the other simply consist in a band-pass filter applied to the signal, whose details can be found in [213, §5.2.3]. For few bands, the latter is usually faster but less precise in the frequency selection. This issue is particularly relevant in the vicinity of spectral lines, as we will discuss shortly, therefore we have preferred the direct, frequency-domain method.

Then, given a set of n bands (which can be pretty general, overlapping, etc.), we proceed to compute the corresponding n time series:

$$BLRMS(t; b), \quad \text{with} \quad b \in \left\{ [f_1^{\min}, f_1^{\max}], \dots, [f_n^{\min}, f_n^{\max}] \right\}.$$

Three scales are involved in computing this by means of the Welch’s method. With the same notation of section 2.1.6, we have: *i*) the *segment* length N , in number of data points, used to compute the fft, which gives the actual frequency resolution, *ii*) the *chunk* length $N' = M \cdot (N - \text{overlap})$, where M segments are averaged to reduce the variance, which, including the overlap, gives the actual time resolution with which our BLRMS estimates are computed, and *iii*) the *total* data sample, which, divided by N' gives the number of independent chunks over which the BLRMS is computed. A trade-off between the previous quantities should be decided in advance, before performing the analysis, in order to have an adequate time-frequency sensitivity to the problem under study.

Lines removal

As previously noted, if a spectral line is contained in the band where we are about to compute the BLRMS, this is going to dominate the estimate. This is sometimes undesirable, especially if one is mostly interested in testing the stationarity of the “noise floor”, since lines are already usually *notched* in most of the search pipelines. Also, we will describe in appendix A.1 an example of noise study where this strategy was necessary. One could then proceed to define *ad hoc* bands, very narrow around the frequencies of interest of the particular line or broader and without lines at all. This is often not practical nor dynamical enough.

For this purpose, we have inserted in the method for computing the BLRMS the possibility to identify spectral lines and remove them automatically. The method is similar to that developed for NoEMi for the identification of spectral lines, and documented in [299, LVC25], but optimised for the opposite scope. We report here its steps:

1. Computation of the “flattened PSD” (without lines):

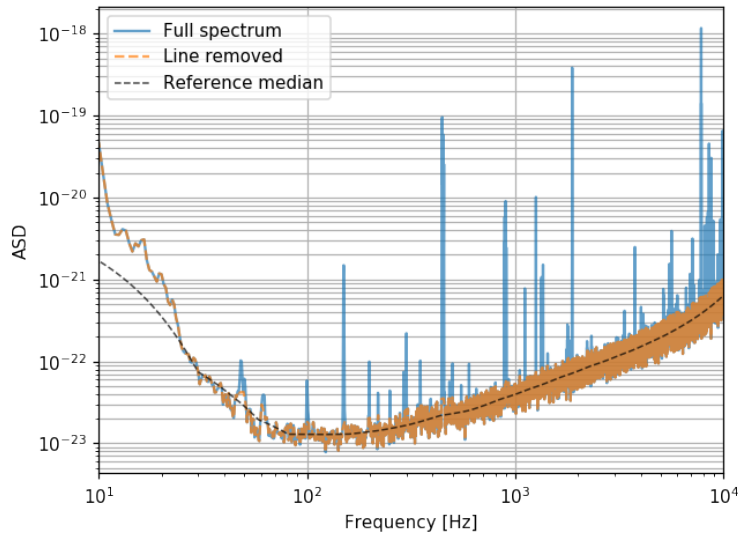


Figure 3.1: Effects of lines removal on the ASD of Virgo strain channel. The blue line represents the ASD with all the lines, while the orange line is the “flattened” ASD with them removed. The dashed lines is a reference median for the application the method.

- (a) A first PSD estimate is done over one or more chunks. If the spectral lines are reasonably stable (which is usually the case), we can increase the number of data (chunks) to include in this estimate in order to have better frequency resolution and less variance. Otherwise, we can repeat this computation and that in the next steps for every chunks, with a consequent higher computational cost due to more iterations;
- (b) In order to reduce the contribution of strong lines in the computation of average values, we define the logarithm of the previous quantity. The reason is the same already mentioned for the PSR test [293], and elaborated in note 3. We divide then the spectrum in a certain number of bands, containing several frequency bins, and in each of them we compute the median over the frequencies;
- (c) The resulting log-PSD medians are linearly interpolated, attributing their values to the central frequency of each band;⁹
- (d) The *Median Absolute Deviation* (MAD) [301] between the interpolated median values and the actual log-PSD are computed over the entire spectrum. Those values of the log-PSD that differ more than a fixed number of MADs (usually 3 or more) from the interpolated values are removed and the corresponding values linearly interpolated;

⁹Actually, this choice is quite arbitrary. However, we have verified that this doesn’t really change the result as long as the previous band division is done in a suitable way. Typically, a linear division is done until a reference frequency of ~ 100 Hz, than a logarithmic one. An alternative approach could have been that implemented by N. Cornish and T. Littenberg in the *BayesWave* search pipeline [195, 300], where, given a set of pivotal frequencies, the noise floor is estimated by means of a spline on them.

- (e) For the resulting interpolated PSD, steps from (1b) are repeated for a fixed number of iterations (usually two or three are sufficient) or according to some other criterion, like the *Mean Squared Error* (MSE) with respect to the interpolated median values. This will constitute the *flattened PSD interpolant function*.
2. Lines removal: if just one chunk has been used for point (1a), proceed to compute the BLRMS on the flattened PSD obtained at point (1e), and go ahead with the other chunks. Otherwise:
 - (a) In each of the chunks used for the PSD estimation at point 1a, estimate the PSD;
 - (b) Compare the logarithm of the previous estimate with that of the interpolant PSD function obtained in (1e). As before, remove those values exceeding a certain threshold, in terms of MADs, and interpolate their values;
 - (c) In each chunk, proceed to compute the BLRMS on these interpolated PSDs.

An example of application of the previous algorithm is reported in figure 3.1; the yellow curve represents the flattened spectrum, where the spectral lines have been removed. The BLRMSs for the Virgo strain data in six bands around the 50 Hz electric mains line are reported in figure 3.2. Far from this frequency, the effect of line removal is absent or minimal, whilst close to it this becomes relevant, as visible from the plots in the central row. Notice also that by removing the energy content of the main spectral line, the visibility of the glitches is enhanced, as they mostly contribute to the increase of the power level in the noise floor surrounding the spectral line.

Glitches identification and removal

Since the BLRMS represents a time series of (the square root of) the energy in a certain frequency band, it is also suitable for identifying glitches. These are clearly visible in figure 3.2 as “sharp”, *i.e.* fast, excesses of power. Many other methods are in use for that within the GW community; some of them will be described in section 3.2. Thanks to them, glitches are already identified and cross-checked, and DQ flags attached to the corresponding data segments; refer to the discussion in section 2.4.3. Instead, very few tools are available for investigating slower non-stationarities. For this reason, we developed a “glitch identification tool”, which, in the same way of the previous line removal algorithm, acts to remove the glitches from the BLRMS’. We did this because we want our stationarity test to be sensitive only (preferably) to slow non-stationarities.¹⁰

For the identification of glitches we have developed a “rolling” *thresholding MAD* algorithm:¹¹

1. The “rolling” median for a certain duration of BLRMS data is computed, and progressively shifted forward in time.¹² At the edges, a “mirror padding” is applied, adding a time reversed copy of the first or last data;

¹⁰An example of study where the difference between working with the BLRMS with the glitches and that with them subtracted will be presented in appendix A.1.

¹¹The term “rolling median” has been borrowed from the same name method in the Pandas Python library [302].

¹²To this purpose, it is important to notice that the complexity of the median computation through the `quickselect` algorithm is $\mathcal{O}(n)$ [303].

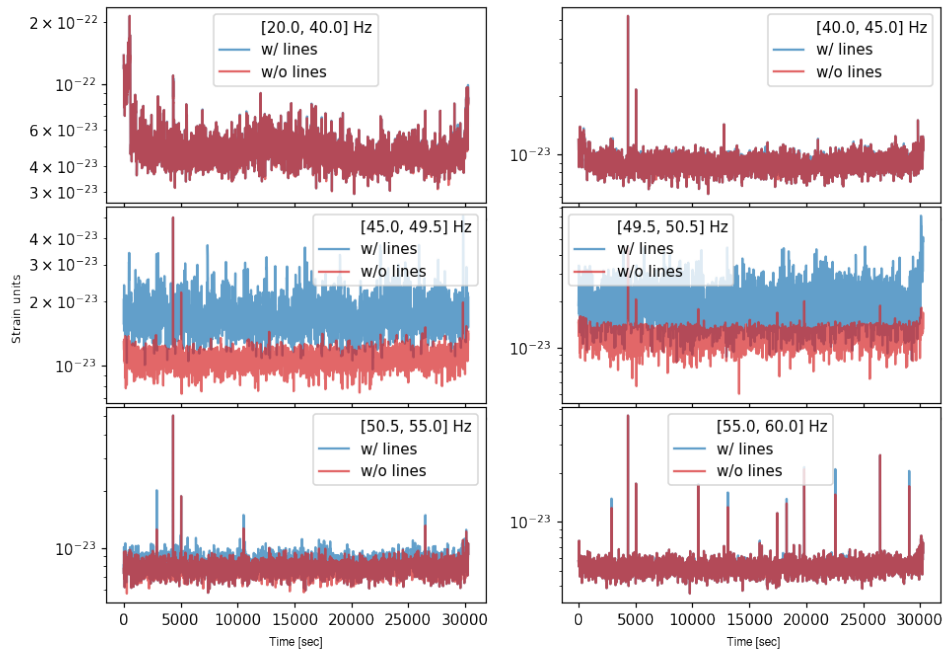


Figure 3.2: Effects of line removal for the BLRMS computed in different frequency bands of Virgo strain data. Notice the reduction in those values close to the 50 Hz main line: plots in the second row. Notice also the increased visibility of glitches, which mostly contribute to an increased noise level in the noise floor surrounding the spectral lines.

2. A *modified* version of the MAD is computed, adding an optional increased weight value to the last data points. This solution has no effect with glitches, but is useful in order not to “trim” the possible inset of a slower transient. Also, the previously found “outliers” are excluded from this computation, in order not to bias it;
3. Data exceeding a fixed threshold based on the previous modified MAD are identified as glitches and optionally discarded (holes in the data) or interpolated. For the purpose of the stationarity test that we will describe shortly, we removed the data corresponding to the glitch, in order not to alter the distribution of the BLRMSs. In other circumstances, evenly sampled data are preferable and the interpolation solution should be adopted.

An example of application of this algorithm is reported in figure 3.3. Some loud glitches are clearly visible in the BLRMS in the top plot, while the effect of their removal is evident from the plot on the bottom.

The next step is to infer whether the time series reported in the previous figure is compatible with the hypothesis of being generated by a stationary process.

Test on the empirical distribution function of the BLRMS

From these BLRMS', we wanted to develop a test sensitive to a large variety of non-stationarities: slow trends, seasonal effects, the models described in section 2.1.3, and possibly also glitches. Some of them are already visible in figure 3.3. Recalling the definition of strict sense




Figure 3.3: Effects of glitch removal in a BLRMS time series. In the top plot, the blue line represent the value of the this function, where some glitches are clearly visible. The yellow line is the moving median, computed over 5 minutes of data, while the shadowed band is that corresponding to 7 times the value of the modified MAD. In the bottom plot is visible the effect of removal of these glitches.

stationarity (2.4), it seems natural to make use of a statistic that compares the “empirical” *cumulative distribution function* (c.d.f.) of the BLRMS time series at two different epochs, but also that is able to provide the information about when the possible non-stationarity appeared.

For this reason, we firstly divided the BLRMS data sample in *intervals* containing a number ~ 100 of points. More details about the choice of this number will be discussed in the next section. Then, in each of the intervals, we compute the *empirical* c.d.f., which, for a set of n *i.i.d.* random variables $\{X_i\}$, is defined as:

$$F_n(x) = \frac{1}{n} \sum_{i=1}^n I(X_i \leq x) \quad (3.2)$$

where $I(\dots)$ is the *indicator* or *characteristic function*, equals to 1 if its argument is true, zero otherwise. The previous quantity counts the proportion of the sample points below level x . To compare it between two different (*neighbouring*) intervals, with n and n' samples respectively (which may also differ), we can use the metric provided by the *two-samples Kolmogorov-Smirnov test* [304, 305], equals to the maximum “vertical” distance between the two functions

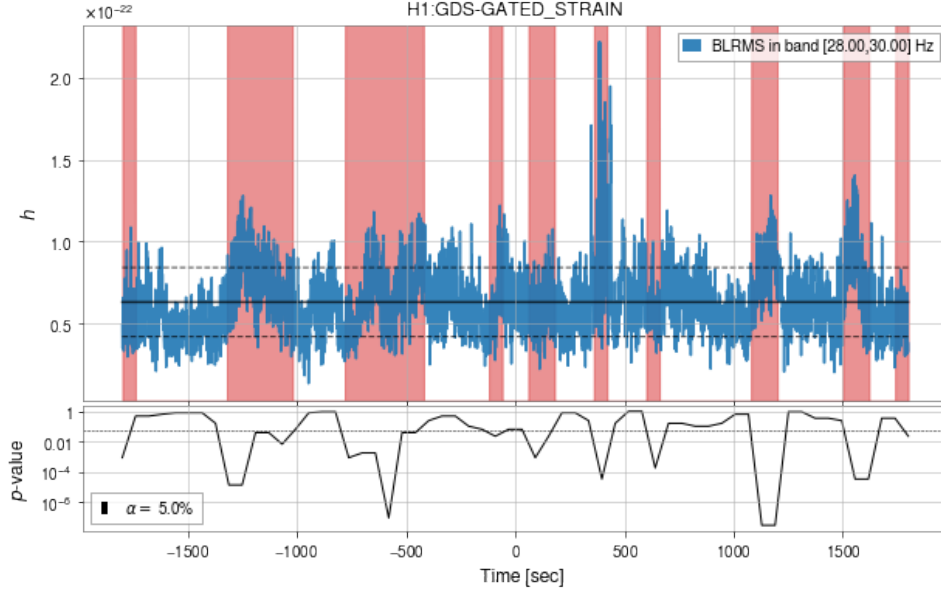


Figure 3.4: The top plot shows an example of BLRMS of LIGO Hanford strain data. The shadowed red regions mark those intervals where the stationarity test has been rejected, according to the p -values reported in the bottom plot, and the chosen significance level $\alpha = 5\%$.

(supremum norm, or $\|\dots\|_\infty$ metric):

$$KS_{nn'} = \sqrt{\frac{nn'}{n+n'}} \sup_x |F_n(x) - F_{n'}(x)|. \quad (3.3)$$

The normalisation factor has been introduced in such a way that, asymptotically, the previous function approaches the *Kolmogorov distribution*, which is independent on the kind of distribution under test (*Kolmogorov Theorem*). This will also be useful when making inference by means of the previous test statistic, as we will discuss momentarily. The critical values of (3.3) have been derived numerically by Massey [311], and are provided by many Data Analysis software libraries, like Python’s `scipy.signal` [312], which we have made use of.

Analogously to the supremum norm, another possibility was to use the L^2 -norm, given by the squared absolute difference between the empirical c.d.f.s; this would have led to the *Cramér-von Mises test* [306, 307], whose generalisation to the two-samples case is due to Anderson [308]: $CvM_{nn'} = \frac{nn'}{(n+n')^2} \left[\sum_i^n (F_n(x_i) - F_{n'}(x_i))^2 + \sum_j^{n'} (F_n(x'_j) - F_{n'}(x'_j))^2 \right]$.¹³

Then, for every BLRMS interval, two p -values are computed, according to the chosen test statistic, for the *null hypothesis* of stationarity of its empirical c.d.f., with respect to those of the two neighbouring segments. These represent the probability of obtaining more “extreme” values for the test statistic (3.3) assuming valid the null hypothesis. The *maximum* of these two values is compared to a *significance level* α , chosen before performing the test (typical values are 5% or 1%); if this value is lower than the threshold, we *reject the stationary hypothesis at a significance level* α . The choice of the maximum has been made in order to avoid

¹³Also the L^1 [309] and L^p metric versions of the previous test statistic exist [310], although they have not been examined yet for the purposes of our stationarity test.




Figure 3.5: Time-frequency map of the p -values of the stationarity test for each time interval and frequency band the BLRMSs have been computed into. Darker shades of red denote those regions where the null hypothesis of stationarity should be rejected with a significance corresponding to the values reported in the colour-bar.

to consider as non-stationary both the interval before (or after) the actual problem and that containing it. This may create some ambiguity in the unlikely case there is a discontinuity falling exactly in between two intervals; we will discuss how to remedy to it in the next section.

In figure 3.4 the BLRMS, with glitch removed, of the LIGO Hanford strain channel in the band $[28, 30]$ Hz is reported for one hour of data. The sampling frequency of this series is 1 Hz, and it has been divided into 90 seconds intervals, *overlapping* by 30 seconds. In this way, the test resolution is 1 min, although it is able to resolve non-stationarity both slower and faster than even 1 second, since they all change the point-wise BLRMS value, and than its distribution in the interval. In the bottom plot the p -values for the stationarity hypothesis are reported, together with a reference significance level $\alpha = 5\%$. Then, those intervals exceeding this threshold have been marked in red in the top plot; in there, the hypothesis of stationarity should be rejected with the (previously) chosen significance level.

Non-stationarities in this region was a known issue in LIGO Hanford data for the first part of O3, which was likely related to scattered light.

In figure 3.5 the time-frequency map of the p -values is reported for the previous test, performed on each time interval and frequency band onto which the BLRMS has been computed. The darker the colour, the most likely the null hypothesis is to be rejected.¹⁴ Again,

¹⁴In figure 3.5 no significance level has been fixed. Then, in performing the test, one should choose the desired α , identifying from the colour-bar the corresponding shade of red, and then reject in the p -map the null hypothesis everywhere is present a darker colour. In the DQR version of this test, this operation is simplified with a value of

the band around 30 Hz presents evidences of non-stationarity, as well as that at 300 Hz. It is also visible, about 1500 seconds after the central time, a non-stationary range of frequency between 500 and 3000 Hz. Another application of this test will be discussed in the context of a specific noise study in appendix A.2 (notice there the double colour-map, to highlight the regions where the test is passed and where it should be rejected according to a significance level fixed in advance; refer to the comment in note 14.) [LVC26].

This analysis test has been named BRiSTOL (Band RMS Stationarity Test toOL) by the author, and presented in [LVC22, LVC24]. The material discussed in this section constitutes also the body of an article in preparation by the author on this new method, to be submitted in 2020 [37]. Since the beginning of O3b in November 2019, it has become part of the DQR produced in low latency in correspondence of every candidate GW event, and used to assess the stationarity of the data in correspondence of it. Moreover, it has also been used for commissioning purposes to monitor the changes on the detector when some adjustment operations were performed on it. Refer for example to [LVC27] or to the example described in appendix A.2.

Comments and choice of parameters

This test has all the properties demanded at the beginning of this section: time-frequency representation of the inferred non-stationarities, independence on the underlying distribution of the data, quantitative statements about the level of rejection of the null hypothesis (p -values). Since, in principle, every kind of non-stationarity, measurable by the BLRMS, alters its distribution, this test is sensitive to a wide variety of them; not just glitches, trends, etc. From the computational side, it is highly parallelisable, since every test is performed independently on the intervals and the bands. The complexity is due to the BLRMS computation, which is basically $\mathcal{O}(N \log N)$ like the fft, and repeated application of the sorting algorithm (usually $\mathcal{O}(N)$) to obtain the empirical distributions.¹⁵

Some parameters and method choices enters in this algorithm, which we shall justify. First of all, they apply all the choices of parameters and time-scales in the computation of the BLRMSs, with their consequences on the time-frequency resolution, and estimate variance; refer to the discussion at the beginning of this section and in 2.1.6. Moreover, the number of BLRMS estimates entering the test, hence the length of the intervals, influences the test outcome. For the Kolmogorov-Smirnov test a number $n \gtrsim 50$ of samples is desirable [313], although some authors have advertised its application also in smaller sample cases [314].¹⁶ what we have observed is that, with less samples, the test tends to be more conservative in rejecting the stationarity hypothesis, and less sensitive to smaller variations in the distribution parameters.

$\alpha = 1\%$ chosen in advance and a two-colour colour-bar palette. This is also shown in another example of application of this tool in appendix A.2.

¹⁵The BRiSTOL tool that implements it on LIGO and Virgo data requires about 2 minutes to read and analyse one hour of data with a typical parameter configuration.

¹⁶To be precise, Lilliefors in [314] and [315] makes explicit evaluations, on the small sample case, referring to Gaussian and Exponential distributions, respectively, with unknown parameters. This may not be the case for other generic distributions but, under the assumption of stationary and Gaussian noise, the PSD estimator is indeed distributed, at each frequency, as an exponential variable, as derived in section 2.1.6.

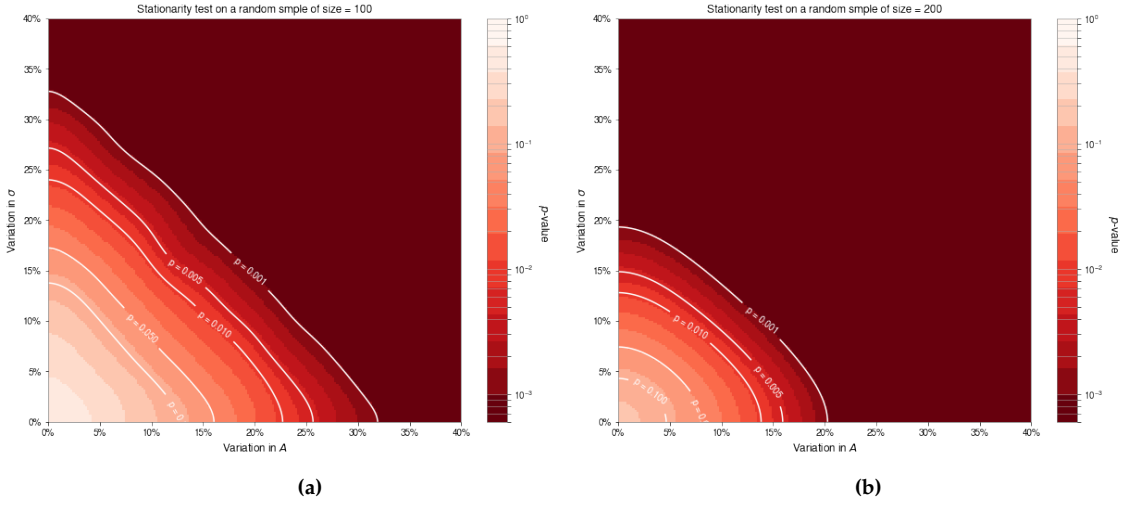


Figure 3.6: Dependency of the test sensitivity to parameters variations, according to model (3.4), for different interval size n . The axes represent percentage variations on the corresponding parameters. In darker regions the null hypothesis is rejected with significance corresponding to the contour.

We have tested the dependency on the sample size n generating our dataset from a “shifted” Rayleigh distribution,

$$p_X(x; \sigma, A) = \theta(x - A) \cdot \frac{x - A}{\sigma^2} e^{-(x-A)^2/2\sigma^2} \quad (3.4)$$

where A represents a constant offset in the values of the variable X , and $\theta(\cdot)$ is the *Heaviside step function* (equals 1 if the argument is positive, zero otherwise). Fixed the reference value $\sigma = A$, we have made these parameters to change smoothly, and run the test comparing the empirical c.d.f.s of the data from the reference distribution and its shifted and scaled versions. In figure 3.6 the p -values as functions of the relative variations in the shape and shift parameters, σ and A , have been reported averaged and smoothed over a large number ($\mathcal{O}(10^4)$) of sample sets. The left-hand side figure is referred to $n = 100$, and, at a significance level $\alpha = 5\%$, it shows that the test is able to identify variations in both the parameters smaller than 15%. For $n = 200$, it was able to detect variations smaller than 10% at the same significance level. We have repeated the same test for other continuous distributions and models, obtaining similar dependency on the length n .

From the previous observations, one may want a test able to spot even small variations in the underlying distribution parameters. This can be done increasing n . However, at the same time this implies a reduced time resolution. To overcome this, we have included in our test the possibility to use overlapping intervals (typically $\sim 20 \div 30\%$ of n). In this way, the test time resolution is usually given by the non-overlapping part of it, while its sensitivity to distribution variations is given by the total n .

We have verified the last claim on the null hypothesis. We have generated, from the same p.d.f. (Rayleigh), a set of 10^3 samples of duration n , and we have estimated the empirical distribution of (some of) their Kolmogorov-Smirnov test statistics (3.3). Then, we have

constructed a similar set of samples of duration $n' > n$, where only n of them were independent, and analogously estimated the empirical p.d.f. for their test statistic. Notice that in the first case the normalisation constant in front of (3.3) was $\sqrt{n/2}$, while in the latter case it was $\sqrt{n'/2}$. Then, we have tested that, again with a Kolmogorov-Smirnov test, the two distributions were equal. Up to values of the overlap $\sim 30\%$, this was indeed the case, with a tendency of the test to be more conservative as this overlap increases. This is somewhat a desirable property, since an increased number of samples has proven to be more able to spot differences in the distribution parameters (fig. 3.6) and the overlap has reduced the number of false-positives.

Lastly, let's comment on the choice of the test statistic. The Kolmogorov-Smirnov test was motivated by the fact that it is moderately sensitive to all characteristics of the distributions, including location, dispersion and shape [313]. Other tests have been tried as well, like the two-samples Anderson-Darling test [316, 317] or the Cucconi test [318]. These have proven to be usually more sensitive to tails in the distributions, like glitches, and, on simulated data, more powerful in detecting even small changes in the distribution parameters for small sample sizes. The Kolmogorov-Smirnov test is more conservative in these cases. However, for the typical parameter configuration in the test performed on LIGO-Virgo data, on BLRMS intervals of $60 + 30$ seconds, with glitch removed, both the Kolmogorov-Smirnov and the Anderson-Darling tests have proven to produce similar results.¹⁷ Anyway, the BRiSTOL tool developed for Virgo implements natively both of the options, and can be run either ways [LVC22].

3.1.3 Other stationarity tests

Before proceeding to discuss Gaussianity tests we report here a couple more stationarity tests (among the many available), in use within the LIGO-Virgo collaboration. These have been selected because involving concepts and ideas already encountered in this dissertation, or because useful complement of the two previous tests.

Variance of the detection statistics

In a recent work, S. Mozzon, L. Nuttal and A. Lundgren have presented a very effective test to determine the influence of non-stationarities (but also of non-Gaussianity) on the matched filter statistic of equation (2.31) [LVC28]. We report here a brief description of their method, which is entirely based on what already discussed in section 2.1.5, as an interesting example that provides also an immediate way to account for the effects of non-stationarity on CBC searches.

As shown in section 2.1.5, the optimal detection statistic, under the assumptions of stationary and Gaussian noise, whose PSD has not being misestimated, is provided by equation (2.41). Its variance should be equal to one, while it is different, in general, if some of the previous assumptions is not valid. The case for Gaussian noise but wrong PSD estimate has been faced in equations (2.43) and (2.38). The idea of the authors is then to monitor the value

¹⁷After the initial testing, we ended up preferring the Kolmogorov-Smirnov test over the Anderson-Darling one because the python implementation of the latter, by means of the `scipy` library `v0.19`, tended to present some bugs in the extrapolated p -values. This is a known issue, which will be possibly fixed in new versions of that library.

of the detection statistic in time, and check if its variations are compatible with 1, as under the assumption of stationary and Gaussian noise.

This test is very effective in the sense that, for a signal model, which for a typical CBC inspiral can be something like $|h(f)| \propto f^{-7/8}$ [103], the loss of SNR (2.44) due to PSD misestimate can be directly evaluated, quantifying how much of the potential signal, under favourable detector condition, we are losing. This method however doesn't take into account (at least, not directly) the potential effect of non-Gaussianity, which may alter the previous variance in a non-obvious way.¹⁸ It is also an "integrated test" over the whole frequency spectrum (which is not an issue for a modelled search, as long as the model has been previously specified, but maybe it is not optimal for noise investigation purposes), and it doesn't consist in a formal hypothesis test; the threshold proposed by the authors to reject the null hypothesis of stationarity is somewhat arbitrary, although quite effective.

Discrete, orthogonal wavelet transform

We should mention here another interesting time-frequency method to assess stationarity, which has been firstly presented in a LIGO-Virgo internal note by N. Cornish [LVC29], and recently discussed also in the collaboration article [2]. This is based on a *discrete, orthogonal wavelet transform* of the signal [319], which, similarly to the short-time Fourier transform, is a way of mapping it onto a *time-scale* space, where the latter has a similar role as the frequency for stationary signals. More details on the wavelet method, which we will use for the study of glitches, can be found in section 3.6.

Instead of the spectrogram representation of figure 2.3, the author suggests firstly to *whiten* the data, that is, divide their discrete Fourier transform by an estimate of their ASD and then transform back to the time domain [320]:¹⁹

$$x(t) \xrightarrow{\text{fft}} \tilde{x}(f) \rightarrow \tilde{x}_w(f) := \frac{\tilde{x}(f)}{\sqrt{S_n(f)}} \xrightarrow{\text{fft}^{-1}} \tilde{x}_w(t). \quad (3.5)$$

The effect of the previous transformation is that of making, overall, the whiten-signal spectrum "more uniform", or, in the time domain, to make its auto-correlation function more similar to a Dirac δ -function: $R(\tau) \propto \delta(\tau)$. In this way, the onset of a transient signal or excess noise (a non-stationarity) is made more evident in the transformed time series; refer for example to figure 2 from [2]. The matched filter (2.31) itself embodies this operation, and we already commented how it can be thought as a "noise whitened" cross-correlation, which makes the resulting SNR optimal. Lastly, it turns out to be particularly convenient to refer to a standard "white" background to assess the significance of other noise features standing out from it, as we will see shortly and as we will elaborate in section 3.6.

The discrete wavelet transform is then applied to the whiten data, producing the so called *scalogram*, which is the equivalent of the *spectrogram* but with wavelet scales on the vertical axis instead of frequencies, and where at each time-scale bin the amplitude of the corresponding wavelet is reported. Of course, in order to help the intuition that we have for the

¹⁸Nonetheless, in practical situations the expectation is of course that its value should increase. This is the case for example of glitches.

¹⁹This operation has been already described in section 2.1.4 in relation to the matched filter technique. In many aspects, it is also analogous to the median normalisation for the spectrogram reported in figure 2.10.

latter, the scales are usually converted to “equivalent frequencies”, following the method of Mayers *et al.* [321]. The squared amplitude of each of the previous coefficients represents a measure of the signal power at the corresponding time-scale (or frequency).

However, as commented for the periodogram PSD estimate in 2.1.6, these coefficients has a natural statistical error whose variance is equal to the estimate itself.²⁰ Then, some sort of averaging procedure, as for the Welch’s method in (2.55), is needed in order to better distinguish the statistical fluctuations of the estimate from other kind of non-stationarities. To this purpose, the author computes the scale-averaged power, summing the squared amplitude of the wavelet coefficient in each time bin, completely losing any frequency dependency of this power; refer to figures 7-9 from [2]. Maybe, an intermediate *smoothing* between the “fine grain-high variance” of the wavelet coefficients, and the scale-averaged spectrum could have been useful to maintain both the informations with a decent estimate variance. We will discuss about a procedure for that in section 3.6.

To quantify the presence of non-stationarities, the author considers the null hypothesis of a stationary and Gaussian process, for which the previous scale-averaged spectrum is distributed, in each time bin, as a χ^2 with a number of degrees of freedom N_s equals to the scales/frequencies over which it has been averaged.²¹ He proceeds then comparing the measured values with the theoretical ones by means of a (one-sample) Anderson-Darling test [322].

It is interesting to note how conceptually this method and the one discussed in section 3.1.2 are very similar for the way they attack the problem of assessing stationarity. Nonetheless, it is relevant the different ways they aim to solve them: discrete orthogonal wavelets instead of short-time Fourier transforms, scale-averaging in place of time and frequency band averaging, one-sample Anderson-Darling test of a reference χ^2 distribution instead of a generic two-samples Kolmogorov-Smirnov test.

3.2 Detection and investigation of glitches

Although the BLRMS method is in principle suitable also to identify glitches, more specific algorithms have been developed within LIGO and Virgo. Some of them have the primary aim to detect transient GW signals, and distinguish their morphology from noise artefacts. The most common, and arguably the reference, method to search for such transient excesses of power is *Omicron*, by F. Robinet [LVC30]. This is based on the search pipeline called *Q*-pipeline or *Omega*, described in S. Chatterji Ph.D. Thesis [323]. Similar to the previous wavelet transform method, it *project* the whitened data stream onto a template bank of sine-Gaussian wavelets, finding time-frequency localised excesses of power. This is used to produce “triggers” in quasi real-time, to which is attributed the central time and frequency (and *Q*-factor, that is the ratio of the central frequency and the characteristic wavelet bandwidth) of the wavelet they are referred to. Neighbouring triggers in the wavelet space are usually clustered together, resulting in just a single one for the whole power excess. Moreover, an

²⁰We will face the same issue when discussing about glitches and wavelet-coherence in section 3.6.

²¹This distribution is also equivalent to a $\Gamma(N_s/2, 2)$ distribution, as already found for the Bartlett’s PSD estimate in section 2.1.6, for $\sigma^2 = 1$. The latter is a consequence of the whitening, which makes the Fourier transformed (complex) time series to have unit variance.

SNR is associated to them on the base of the expectation of a similar wavelet coefficient due to white noise. This is similar to the way p -values are computed for noise excesses in the previous discrete orthogonal wavelet method.

Other matched filter based search pipelines, which in low-latency inspect the strain data stream for transient GW signals, produce a list of triggers that turn out not to be coincident, both in time and parameter space, with those of the other detectors. Or, to be more precise, whose occurrence has low significance, according to the criteria described for PyCBC and GstLAL in section 1.4. These triggers are usually removed from the search, and, interpreted as glitches, used for Data Quality purposes. An example of Data Analysis pipeline that produce this kind of triggers is the “live version” of the already mentioned PyCBC [193], that is PyCBC `live`. Another search pipeline that produces in very low latency the information about excess power triggers is the Multi-Band Template Analysis (MBTA) pipeline [324]. This is a template based, CBC search pipeline, similar to PyCBC and GstLAL, which uses the expedient of dividing the matched filter in two frequency bands, low and high with respect of the pivotal frequency of 100 Hz, to speed up the production of the results. The corresponding templates $h(t)$ in each band are indeed shorter, as discussed in section 1.2, and consequently the matched filter is less computationally demanding, allowing for lower latency results.

Both PyCBC and MBTA triggers are available, together with those from Omicron, in the Virgo Interferometer Monitoring (VIM) web page for noise studies [LVC6]. Similar information, on the LIGO side, with the exception of MBTA, is reported in the detectors status summary pages [LVC31].

Once these glitch triggers have been generated, it is interesting to find correlations with the other “auxiliary channels” in the detector, in order to trustworthily exclude the astrophysical origin and, possibly, to identify the causes that may have generated them. This operation is analogous to the search for coincident triggers in the strain channels of two different detectors, to which is attributed a probability of being of astrophysical origin, as described in section 1.4.

One of the reference data analysis tools to investigate this is the LIGO *Hierarchical Veto* (HVeto) algorithm, by J. Smith *et al.* [325]. This compares the triggers found in a target (*e.g.* the strain or DARM) channel with those in other auxiliary channels. The statistical significance of coincident glitches is evaluated, and those channels that appear “to glitch” at the same time of the target are marked as possible causes, and further investigations are demanded. A similar tool, customary used in Virgo analysis is the *Use Percentage Veto* (UPV) tool, firstly developed by T. Isogai [326]. This computes the coincidence significance of a glitch trigger in the target channel and in the auxiliary ones by means of the percentage of time they “glitch together” in various frequency bands.

It is important to stress from now that *coincidence* and *correlation* are not logically related to *causation*. We will elaborate on the latter in section 3.7. For now it is important to have in mind that if similar features, like glitches, are found in two or more channels, this is not a proof of a causation relation. Instead, this should be just the starting point for further investigations, as for example to which part of the detector to start looking at, and acting on it try to solve to occurrence of the particular noise feature. The previous considerations

will be relevant also in the next sections when discussing about other tools for correlation analysis.

Another important aspect that we should mention is that of “unsafe channels”. These are usually signals used for the reconstruction of the calibrated strain channel, h_{rec} , or with known relationships with DARM, as for example other longitudinal control channels. It is not surprising to find correlation between them, but this is usually of little interest when looking for the causes of the glitches in a set of them.²² Moreover, the excess of power corresponding to a transient GW signal can be simultaneously present in some of them. If these were not previously “flagged” and excluded from the analysis, this would have misleadingly lead HVeto or UPV to consider this signal as a glitch. The identification of safe and unsafe channels is part of the Detector Characterisation activities, and it usually done by means of noise injections, as described in section 2.3.

3.3 Detection methods of non-Gaussianities

Once some reliable test have convinced us about the stationarity of the data, we can proceed by applying the Ergodic Theorem, computing statistics and expectation values in all those frames where this hypothesis has not been rejected. The next step is to test Gaussianity, which is fundamental for the application of the matched filter technique in the form discussed in section 2.1.4. Notice that most of the Gaussianity tests are automatically rejected if the underlying process is still Gaussian but not stationary, just for the way test statistics are computed by means of time averages. For the same reason, a Gaussianity test may be rejected both because the data are non-Gaussian but also because they are non-stationary.

We have repeatedly insisted on the importance of testing stationarity and Gaussianity separately. Then, before proceed to describe the tests for the latter, let us comment on why this should be the case. Of course, from the point of view of the process properties, these are clearly very different aspects, with usually different origins. In the spirit of the Detector Characterisation, these are studied with various approaches, as described in section 2.4.4. From a practical point of view, the immediate “countermeasure” to mitigate the effects of non-Gaussianities and non-stationarities on GW searches is to apply *notches* to the frequencies, and *cuts* to the data segments. This however has the drawback of throwing away information simply because we are not able to analyse it with standard methods without introducing systematics. Quite remarkably, some techniques have been recently developed, and some others are currently under study, able to perform the *subtraction* of particular noise features, like spectral lines [274, 327, LVC32], non-linear sidebands [328] and other kinds of broadband, bi-linear couplings [275], like for example the oscillations, or “jitter”, of the pre-stabilised laser (PSL) [329]. Refer to their description and characterisation in section 2.3.3. In general, all of these techniques assume a certain level of stationarity, either for the durations needed to estimate the signals transfer functions or for the underlying noise model, which however may generate a non-stationary process, as in the case of [274] and [328]. Instead, Gaussianity is not always required; this motivates our choice of studying the two aspects separately.

²²It is interesting instead the opposite situation, when, for example, a glitch or a particular noise feature is present in h_{rec} but not in DARM.

We proceed then describing how to test Gaussianity of GW detectors data. For the same reasons discussed at the beginning of section 3.1.2, a good test for detector characterisation purposes should be able to investigate this property frequency-by-frequency. Hence, the popular one-sample Kolmogorov-Smirnov test [298, 304], for data with known mean and variance, or the Lilliefors test if these are unknown [314], or the Anderson-Darling test [316], just to mention a few of the most common, are not particularly useful if applied to the whole data. Any small or “narrow” non-Gaussianity may be averaged away with these methods.²³ One can then think to apply the previous tests to each frequency component of the signal DFT. This is a possibility, but the one we are about to describe yields a similar result and is computationally less expensive.

3.3.1 The Rayleigh test

The common way the hypothesis of Gaussianity is tested in LIGO and Virgo is by means of the so called *Rayleigh test* [LVC23, 289]. Refer to note 1 for a brief account of its origin. Recalling the theory of spectral estimation discussed in section 2.1.6, if a stochastic process is stationary, it possess a well defined PSD, which we can estimate with the various techniques described in that section. In particular, if the data is Gaussian, the Bartlett’s method ASD (2.54) can be considered an estimator for the rms of a variable $|\tilde{X}(f)|$ distributed like a Rayleigh with scale parameter σ . Its rms is $\sqrt{2}\sigma$. Instead, the *median-averaged ASD* (2.57), is an estimator for the median of this variable: $\sqrt{\log 4}\sigma$. Hence, if the Gaussianity hypothesis is valid (as well as stationarity), one can compute, from a single fft of the signal (complexity $\mathcal{O}(N \log N)$), the previous two quantities and then define the statistic:

$$Ry_1 := \frac{\text{rms}(|\tilde{X}(f)|)}{\text{median}(|\tilde{X}(f)|)} \stackrel{\text{Gauss.}}{=} \frac{1}{\sqrt{\log 2}} \simeq 1.20 \quad (3.6)$$

besides normalisation factors due to the finite averaging size M [330], as in (2.60). Most noticeably, this doesn’t depend on the distribution scale parameter σ , *i.e.* it is the same at all frequencies. Another nice property of the previous quantity is that the denominator, the median, is a robust measure of the central tendency of the distribution, which is barely influenced by the possible presence of non-stationarities, while the numerator, the rms, is highly influenced by them. This makes the test also very proficient in detecting non-stationarities. In the next section we will investigate its performances both on non-Gaussianities and non-stationarities.

Let us mention a couple more forms of the Rayleigh test statistic before discussing their performances and the method we used to derive their critical values for proper hypothesis testing.²⁴ Exploiting the fact that most of the Rayleigh distribution properties are factorisable by the scale parameter σ , as shown in appendix B.1, one can define many other test statistics independent on σ (but not on M). For example, in [290] the authors suggest as a figure of merit to compare the estimated 95th and 99.7th percentiles with their theoretical value for

²³This is a consequence of the central limit theorem. To be more specific, the Anderson-Darling test is likely to identify them as “tails” in the distribution, and then reject the null hypothesis. At this point however there is no information about “where” these non-Gaussianities are, and one could be tempted to discard the whole data.

²⁴To this purpose, I’m thankful to N. Sorrentino, who greatly helped me in performing these tests, and for the related discussion about which statistic was the most suitable to use for Detector Characterisation.

a Gaussian distributed variable of known (previously estimated) parameter σ . Although, they don't formally quote any critical value to assess rejection, the plot shown in figure 15 therein, and the corresponding discussion, would suggest to reject Gaussianity everywhere below 100 Hz, which is obviously a bit too extreme. Presumably, this is a consequence of the high uncertainties on the estimation of these quite high percentiles, where there are few point for moderate values of M , and a consequent high tendency to oscillate.

The original version of the test, implemented at LIGO in the `Rayleigh monitor` tool, is instead [289]:

$$Ry_2 := \frac{\text{std}(|\tilde{X}(f)|^2)}{\text{mean}(|\tilde{X}(f)|^2)} \stackrel{\text{Gauss.}}{=} \frac{2\sigma^2}{2\sigma^2} = 1 \quad (3.7)$$

except for the usual corrective factors for finite M [331]. The previous values are referred to a $\Gamma(1, 2\sigma^2)$ distribution, like that in the arguments of the standard deviation (`std`) and mean functions under the hypothesis of Gaussianity. Notice that in this formulation of the Rayleigh test, it is evident its interpretation as a measure of the dispersion of the stochastic variable $|\tilde{X}(f)|^2$ with respect to its expected value [332]. Moreover, an equivalent form of the previous quantity is:

$$Ry_2 = \sqrt{\frac{E[|\tilde{X}(f)|^4] - (E[|\tilde{X}(f)|^2])^2}{(E[|\tilde{X}(f)|^2])^2}} = \sqrt{\frac{E[|\tilde{X}(f)|^4]}{(E[|\tilde{X}(f)|^2])^2} - 1} \stackrel{(2.2)}{=} \sqrt{\frac{\mu_4(0)}{\mu_2(0)^2} - 1}$$

where $\mu_i(0)$ is the i -th distribution moment about zero for the variable $|\tilde{X}(f)|$; refer to its definition in section 2.1.2. This makes the fraction inside the square root similar to a *kurtosis*, $\mu_4(\mu)/\mu_2(\mu)^2$, which is a measure of "tail propensity", that is the presence of "outliers" in the tails of the distribution [333].

Then, if the fluctuations of the measured noise PSD are smaller than those expected for a Gaussian noise, as for the case of the *coherent* amplitude oscillation of a spectral line, the values of the Rayleigh statistic are expected to be smaller than 1. Conversely, *incoherent* variations, as it is typical for glitches, are expected to produce values for Ry_2 higher than 1. We will verify this momentarily.

Similar to the previous one, let's introduce also the Virgo version of the Rayleigh test [LVC23], regularly produced for the DQRs:

$$Ry_3 := \frac{\text{std}(|\tilde{X}(f)|)}{\text{mean}(|\tilde{X}(f)|)} \stackrel{\text{Gauss.}}{=} \sqrt{\frac{2 - \pi/2}{\pi/2}} \simeq 0.52. \quad (3.8)$$

All the comments made about the interpretations of Ry_2 in terms of dispersion and kurtosis for the variable $|\tilde{X}(f)|^2$, also apply to Ry_3 for the variable $|\tilde{X}(f)|$. Since the two are linked by a monotonic transformation, which preserves the percentiles, the two test statistics Ry_2 and Ry_3 are in fact equivalent.

3.3.2 Critical values and analysis of real data

We have commented that many possible versions of the previous Gaussianity test can be built exploiting the properties of the Rayleigh distribution. However, for brevity, we report

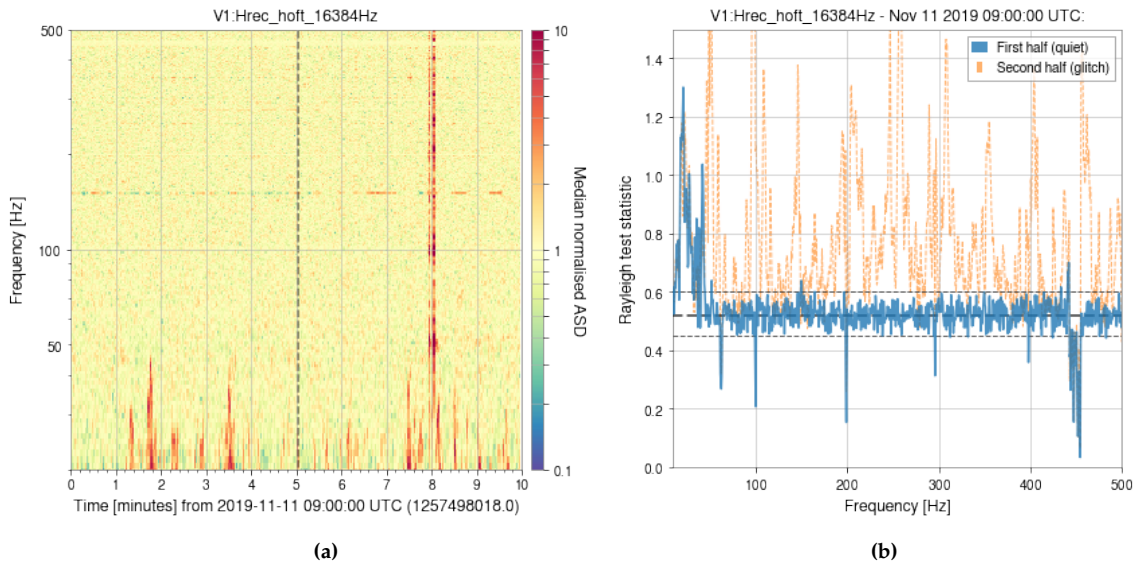


Figure 3.7: Left: median normalised spectrogram of a stretch of 600 seconds of Virgo strain data, where a loud glitch is present at about 8 minutes from the beginning, as well as some low frequency non-stationarities. Right: Gaussianity test by means of the Ry_3 test statistic, for the two halves of the data; fft length of 2 seconds, corresponding to $M = 150$ averages.

here only the results about the tests we made on Ry_1 and $Ry_{2,3}$, to assess their performances, which made us to prefer the latter.

As it should be the case for every proper hypothesis test, firstly we constructed a *table of critical values*²⁵ for these three test statistics. We did this by generating a large ($\mathcal{O}(10^5)$) number of Gaussian data samples (null hypothesis), and computing, for various M from 5 to 1000, their critical values from 10^{-5} to $1 - 10^{-5}$. These turned out to be smoothly distributed on the $(M, \text{percentile})$ plane. Then we proceeded to construct the cubic interpolant (and extrapolant) function, which was used later to find the corresponding critical value for any M and percentile.

Applying the previous test statistics to real data, it turned out that Ry_1 was very good at spotting non-stationarities in the data, but it was not the same in making “non-Gaussianities” to stand out with an high significance from the statistical noise of the null hypothesis. This is probably a consequence of using the rms at the numerator of (3.6), which is too sensitive to even random fluctuations in the data, which in turn broadens the distribution of this test statistic. The other statistics, Ry_2 and Ry_3 , instead turned out to be quite good at detecting both non-Gaussianities and non-stationarities, as opposite outliers in their distributions according to the null hypothesis of (stationary and) Gaussian data.

In figure 3.7 we report the analysis of a stretch of 600 seconds of Virgo data, containing some interesting noise features. First of all, it obviously contains the spectral lines typical of Virgo strain sensitivity curve, fig. 2.2; it also contains some low frequency noise, likely due to bad weather conditions and scattered light, and a loud glitch on the second half of the data, about 8 minutes after the beginning. Figure 3.7a contains a median normalised

²⁵Which in modern programming language is actually an instance of the class `griddata` of the `scipy.signal` Python library [312].

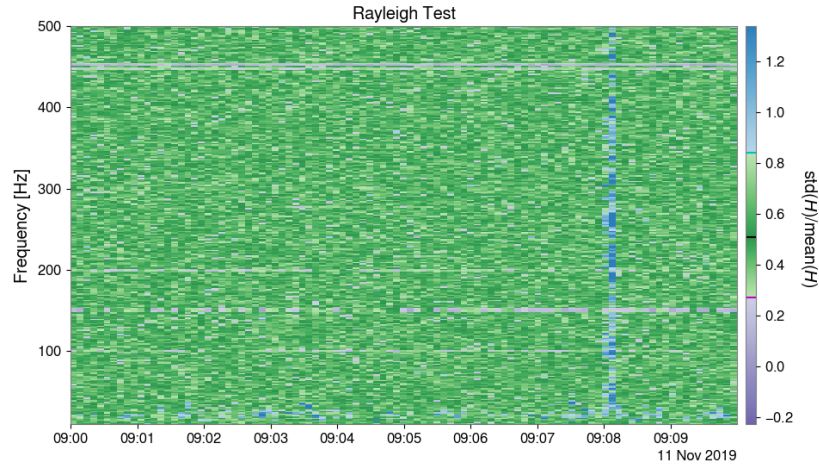


Figure 3.8: Rayleigh-gram corresponding to figure 3.7 obtained by means of Ry_3 test statistic. Red and blue lines on the colour-bar represent the .5 and 99.5th percentiles for the null-hypothesis distribution of Ry_3 . Image obtained with RAGoUT.

spectrogram where the previous non-stationary features stand out in darker shades of red. In figure 3.7b the results for the Gaussianity test by means of Ry_3 are shown for the first and second half of the data. The ffts have been computed on segments of 2 seconds of duration, which implies $M = 150$ averages. For the latter number, the corresponding critical values for the 0.5, 50 and 99.5th percentiles, implying a two-sided hypothesis test of significance $\alpha = 1\%$, are represented with dashed lines.

First of all, the test on the second half of the data, where there is a loud glitch, is completely out of the critical region for the null hypothesis of Gaussianity. This is the same effect already noted for the Welch’s ASD estimate in figure 2.4b. The test for the first half instead behave more or less as expected. It is altered at low frequency for the presence of non-stationary noise; it correctly identifies as non-Gaussian the calibration line at ~ 60 Hz, the 100 and 200 Hz main harmonic lines, and the wide region from 443 to 454 Hz, corresponding to the violin modes for the large mirrors suspensions, as measured by T. Hardwick and Y. Michimura in the pre O2 commissioning phase [LVC33]. Interestingly, the 150 Hz region behave, overall, as incoherent noise ($Ry_3 > 0.52$), and not like the other harmonic lines. This is likely to be a consequence of the non-linear noise couplings in this region, already described in section 3.5.

To have another perspective on the effects of non-stationarities and the presence of non-Gaussianities, we can make use of the so called *Rayleigh-gram*, representing a “short-time” version of the previous test, as reported in figure 3.8 for ffts of 0.5 seconds and $M = 10$. The corresponding critical values for a 1% significance two-sided Gaussianity test are reported on the colour-bar; blue regions are those corresponding to incoherent noise, like glitches, while purple ones are those typical of spectral lines and coherent variations.

This tests have been implemented in the Virgo Detchar tool RAGoUT (RAYleigh GaUSSianity Test), developed by the author in collaboration with N. Sorrentino [LVC24]. Another example of application of this in the context of a noise study is described in appendix A.2.

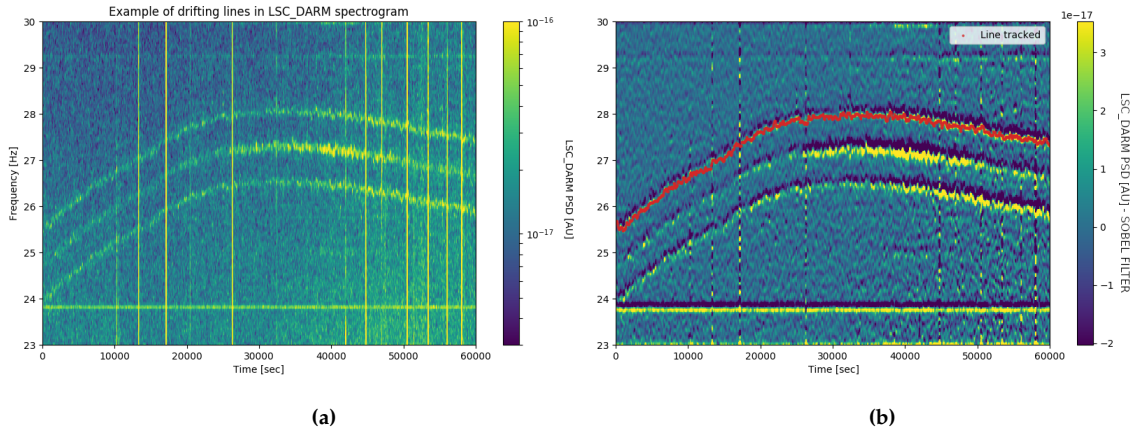


Figure 3.9: Example of a “triplet” of wandering lines in Advanced Virgo LSC_DARM data of August 2018 (during the upgrade phase after O2). Left: spectrogram with these lines represented by the curves ranging between 24 and 28 Hz. Other noise features are visible as well, like glitches (vertical lines) and a (horizontal) spectral line at 23.8 Hz. Right: application of a Sobel vertical transform, which enhance the presence of the wandering lines and removes the glitches. The red dots constitute the frequencies time series corresponding to the tracking of the line.

3.4 Line tracker tool for frequency non-stationary noise

In the previous sections we have presented several techniques to identify, mainly, amplitude non-stationary noise, according to the definition we gave in section 2.4. We want to focus now on “frequency non-stationary” noise, like the *drifting line* visible in figure 2.12. One possible strategy to attack this problem and identify its causes is to *track* the varying frequencies assumed by this line with time, and relate them to the information provided by some auxiliary channel [LVC15]. Indeed, as in the case of figure 2.12 (top plot), in some circumstances these frequencies are linearly correlated with the amplitudes of other auxiliary sensor channels, like the two thermometers in the top plot of that figure, which witness a similar non-stationarity behaviour, hence providing a hint about the possible causes that make this line to drift.

Let’s start from describing the algorithm to extract the information about the varying frequencies of this line. This has been implemented in the `LineTracker` pre-processing stage of the analysis tools that we are going to describe in section 3.5 [LVC34]. The starting point is the spectrogram of the stretch of data containing that particular feature. From this, what we track the frequencies of the local maxima of this line, and record them into a time series. If we had only that feature, that is, a single wondering line in a certain frequency band, everything we need is to find the maximum of the PSD amplitude in that interval, in each time bin. This was the basic idea behind the first implementation of a similar algorithm in 2017 by S. Koley [LVC15]. We improved upon this in the way and for the reasons that will be clarified momentarily.

Unfortunately having a single, “clean” wandering line is almost never the case with real data, especially after the improvement of the detectors and the removal of all the principal and most evident causes of these lines, as it has been done after the pre-O2 phase described

in section 2.4.2. Indeed, for example in O3 or in the commissioning phases that preceded it, the typical wandering lines had intensities not much greater than the noise background, and very often they overlap with other (stationary) spectral lines and glitches. An example of this is reported in figure 3.9a, comprising of a “triplet” of wandering lines in Virgo DARM signal, overlapping with many glitches (vertical lines). The direct application of the simple tracking of the maxima of the PSD would have probably led in this case to get stuck on the 23.8 Hz line, or jumping between one and the other of these lines, or to be influenced by the presence of the glitches. An even more complex scenario was present (before the application of the Hough transform to line masking) in figure 3.10 and in the noise study described in appendix A.3.

3.4.1 Image filtering

To make the wandering lines to stand out with respect to other noise features, we borrowed some *imaging techniques* from the field of Digital Image Processing [334, 335]. Referring again to figure 3.9a, for example, one can notice that the disturbance produced by the vertical lines due to glitches can be *removed* by applying some vertical differentiating operator. A typical choice, popular for *edge detection* purposes, is the application of a “vertical” *Sobel filter* [336]. This acts on $S(f; t)$, at each point of the time-frequency map, *image*, computing the convolution of it with a kernel, representing an approximate gradient, or a vertical derivative, for our purposes. Its effect is to emphasise regions of *high spatial frequency* that correspond to significant variations, *edges*, in the value of $S(f; t)$, *e.g.* glitches.²⁶ Typical dimensions for this kernel are 3×3 , or other odd integers, which makes its application quite inexpensive in computational cost. For example, defining the “spectrogram matrix” $S := \{S(f; t)\}$, its Sobel vertical transform is [335]:

$$\mathbf{G}_y := \begin{pmatrix} +1 & +2 & +1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{pmatrix} * S$$

An example of its application is visible in figure 3.9b, where the vertical lines have been considerably attenuated by this operation.

In case the wandering line had overlapped also with some (stationary) spectral line, like the one at 22.8 Hz in figure 3.9, instead of proceeding with an horizontal Sobel filter, which would have removed significant portions of the wandering line too, we usually prefer to make use of a *Hough transform* for detecting straight lines [337]. This allows to identify the time and frequency regions corresponding to stationary spectral lines (horizontal) and glitches (vertical). An adequate tuning of the transform parameters, such as the slope and the length, are necessary to avoid the wandering line.²⁷ Then we have proceeded to mask the previous region and identified the wandering line in the rest of the image.

²⁶In this interpretation, we consider an image as a realisation of a stochastic process whose *index set* is given by the tensor product of two spatial dimensions; recall its definition in section 2.1.2. In this case, the spatial frequency can be defined from the Wiener-Khinchin theorem analogously to what done in (2.15) for the times.

²⁷To this purpose, we have made use of the `OpenCV` image processing library and its Python implementation through the package `cv2`, and the function `HoughLinesP`.




Figure 3.10: Example of line tracking with a preliminary identification, by means of a Hough transform, of other noise features. These consist for example in the spectral lines at 87 Hz and between 92 and 94 Hz, and in some thin vertical lines, corresponding to glitches, and some broader ones, in correspondence of interferometer unlocks or high noise periods. The corresponding regions are interpolated by means of the median ASD over all the image. With this preprocessing stage, the line extraction algorithm was able to correctly track the wandering line also in those region corresponding to the previous noise features.

Other techniques we have made use of, for the purpose of making the wandering line more easily detectable, include the Prewitt transform [338] and the *Canny's edge detection* algorithm [334, 335], which consists into a series of *Gaussian filter*, some passages of the Sobel operator, and *hysteresis* to suppress all those lines that are weak and not connected to strong ones.

An example of masking after the application of the Hough transform is reported in 3.10, together with the frequencies series reconstruction that we shall describe momentarily. This line corresponds to a well known issue in Advanced Virgo data, present, to the best of our knowledge, since June 2018 [LVC35]. We will cover the full story of the investigation related to this line in appendix A.3.

3.4.2 Frequencies series reconstruction

Once we have pre-processed the spectrogram image to make the wandering line more evident, we proceed to reconstruct the time series of the frequencies corresponding to the maxima of the PSD at each time. We do that by modelling the evolution of the frequencies as a *Markov process*, already introduced in section 2.1.3, where the value at a certain time depends only on the previous one, through a fixed thought unknown probability function. Part of the correct reconstruction of the frequencies time series consists into the adequate modelling of this function.

The algorithm starts identifying the frequency at the maximum PSD in the first time bin (*i.e.* column of the spectrogram). Then, the next frequency value is chosen as the next PSD

maximum, weighted by the certain *transition probability* given by the previous model. This should reflect the unknown distribution governing this Markov process, and the choice of the metric has the role of down-weighting too wide transitions to other noise features. This is useful in the presence of other spectral features, like the line in figure 3.9, where there is the risk to jump from one line to the other, instead of following the path along a single wandering line. We made use of both a Gaussian and an uniform p.d.f. models, and iterated the frequencies series reconstruction (usually 2 ÷ 5 times) to find the best values for their variances.

The first iteration is usually done with a pretty loose value of the variance. Then, step after step, the frequencies of the weighted maxima are reconstructed through the whole line, as visible in figures 3.9 and 3.10. In the second iteration of the algorithm, the sample variance of the previously reconstructed frequencies series, which is usually quite high, is substituted to the variance of the model p.d.f. for the Markov process. This iteration usually results in a smaller value of the frequencies series variance, which corresponds to a reduction of the random noise fluctuations of the frequencies time series.

Termination criteria for the previous algorithm depend in general on how net is the wandering line with respect to other noise features. This fact, together with the other parameters that should be tuned on the base of the frequency region and expected variations of the line, makes this algorithm not suitable at present for automatic line identifications. It should also be mentioned that in particularly noisy situations, the previous filtering techniques and line reconstruction algorithm are prone to failed to correctly resolve the wandering line from the rest of the noise.

Within the Virgo Detchar group there is another line tracking algorithm under development by B. Swinkels [LVC17]. This is based on a different strategy that consists into the manual selection, by means of graphical user interface, of the *bounding region* where to look for the wandering line. In there, it performs the search for the maxima of the PSD, as described before. From a computational point of view, this is clearly faster than ours, but it is more sensitive to the presence of other spectral features, which should be manually excluded from the bounding region for the search. This is not necessary with the algorithm previously described that is able (at least in not too pathologic situations) to automatically remove all the glitches (vertical lines) and spectral lines (horizontal) before attempting the search and reconstruction of the frequency evolution of the line, thought as a Markov process. Several example of this, in noisy contexts and other overlapping features have been reported in [LVC16] and will be described in appendix A.3.

3.5 Investigating the correlation with the auxiliary channels

Having identified some segments where the data can't be considered stationary, or some regions of the spectrum where non-Gaussianities are present, is usually just the beginning of the Detector Characterisation activity. As regards the latter, we have already presented in section 2.3.3 how to relate these features to the information provided by the auxiliary channels by means of the coherence method, and its extensions. For non-stationary noise we can

apply a similar strategy but in the “time domain”. This can be done investigating the *cross-correlation* of a target channel, namely the strain or the BNS range described in section 1.3.3, with the time series corresponding to any of the hundred of thousands of auxiliary channels that monitor the interferometer subsystems and their environment. The target can also be the frequencies series of a wondering line or some transformed version of the data, as we will discuss momentarily.

The cross-correlation technique will be presented in the next subsection, while its extensions to non linear terms and to the multiple linear regression will be discussed in sections 3.5.2 and 3.5.3. In section 3.5.4 the non-stationary Noise Analysis tool NonNA that the author have implemented for Virgo will be described [LVC34]. Its application to the study of environmental noise (refer to section 2.2) has been described in [29].

3.5.1 Cross-correlation analysis

The idea behind the *auto-correlation* function, as a measure of similarity between a time series and a “time-shifted” version of it, has already been introduced in section 2.1.2. In section 2.3.3 we have briefly touched the cross-correlation, which extends the previous concept to two different time series. Assuming both of them to be sampled at the same frequency, x_t and y_t , with $t = 1, \dots, N$, we have introduced in (2.47) the (*unbiased*) estimator for their cross-covariance:²⁸

$$\hat{R}_{xy}(\tau) = \frac{1}{N - |\tau|} \sum_{t=1}^{N-|\tau|} (x_t - \hat{\mu}_x)(y_{t+\tau} - \hat{\mu}_y), \quad \tau = 0, \pm 1, \dots, \pm(N-1) \quad (3.9)$$

where $\hat{\mu}_x = (N)^{-1} \sum_{t=1}^N x_t$, and similarly for $\hat{\mu}_y$, is the *sample mean*. The previous summation goes from $t = 1$ to $(N - \tau)$ if $\tau \geq 0$, or from $t = 1 - \tau$ to N if $\tau < 0$. We will be particularly interested in making use of the “normalized” version of the previous estimator, whose values range from -1 to 1 , that is the *sample cross-correlation function*:

$$r_{xy}(\tau) = \frac{\hat{R}_{xy}(\tau)}{s_x s_y}$$

where the *sample standard deviations* of the series are obtained from $s_x = \sqrt{\hat{R}_{xx}(0)}$. Notice that unlike the auto-correlation function, which assumes its maximum at zero lag, this is not in general true for the cross-correlation; remember the example of two retarded signals in section 2.3.3. Nonetheless, it is convenient to define the *sample Pearson’s correlation coefficient* as the value at zero lag of the previous function:

$$r_{xy} := r_{xy}(0) = \frac{\hat{R}_{xy}(0)}{s_x s_y} = \frac{\sum_{t=1}^N (x_t - \hat{\mu}_x)(y_t - \hat{\mu}_y)}{\left(\sum_{t=1}^N (x_t - \hat{\mu}_x)^2 \sum_{t=1}^N (y_t - \hat{\mu}_y)^2 \right)^{1/2}}. \quad (3.10)$$

²⁸Previously we have usually labelled the time index with a lower case-letter from the middle of the alphabet, like in x_i or x_n . In this section we will have to deal with multiple indices referring to different kind of quantities, hence we decided to stress the meaning of temporal index by using the letter “ t ”. Also, notice that we prefer now to make use of the *unbiased estimator* for the cross-covariance rather than the biased one used for spectral estimations in section 2.1.6.




Figure 3.11: Example of cross-correlation study, where the Virgo LSC_DARM BLRMS in band $[40, 47]$ Hz is correlated to the time series corresponding to the (maximum) intensities measured by a magnetometer in the Virgo Central Building, ENV_CEB_MAG_V_max. In correspondence of this a malfunctioning chiller was eventually found to be the origin of the glitches, represented by spikes, visible in both channels [LVC36].

This quantity measures the *linear relationship* between x_t and y_t , indeed, similarly to the coherence, it is left unchanged by linear transformations of the two signals.²⁹ The sign of r_{xy} represents the “direction” of the linear relationship, that is, whether to high (low) values of the first signal correspond also high (low) values of the second one. The absolute value $|r_{xy}|$ gives the strength of this relationship, hence the similarity of the two signals, in the time domain.

The previous considerations can be applied to the study of non-stationarities. For example, in figure 3.11 we have reported the BLRMS (without glitch removed) of the Virgo DARM signal in the band $[40, 47]$ Hz (blue line) in relation to the signal corresponding to the magnetic field measured by a magnetometer sensor in the Virgo Central Building (ENV_CEB_MAG_V, red line). Several “spikes” in the former were coincident to those in the latter; the resulting Pearson’s coefficient was $r_{xy} \simeq 0.65$. Although this doesn’t constitute a proof of a causal link between the two, it was a clear hint to look for something electric, activating at the times of the spikes, located in the Virgo Central Building. The “culprit” was eventually found in a malfunctioning chiller in the Laser Lab of the Central Building [LVC36].

Situations similar to this one suggest the importance of a Detector Characterisation tool that, in the presence of a particular non-stationarity, like the previous glitches or some “drops” in the channel measuring the BNS range, executes a “cross-correlation search” among all the auxiliary channels, or in general large subsets of $\mathcal{O}(1000)$ of them. As previously mentioned, this kind of approach is colloquially referred to as “brute force”. The tools developed by the author to this purpose will be momentarily described in section 3.5.4 [LVC34].

²⁹This claim can be easily verified in the frequency domain, as discussed in section 2.3.3, than transforming back to times.

3.5.2 Data transformation and non-linear methods

As for the case of the BLRMS, which allows to focus on the variations of power in a particular frequency region, also for the auxiliary channels the possibility of making transformations has been taken into account. One example is indeed to apply to them the same BLRMS transformation of the target channel. This is the case, for example, if we find some noise features in a specific band of the target and we want to investigate if this is present in the same band in other channels.

We will describe here some common transformations implemented in the tools described in section 3.5.4 as a *pre-processing* stage for the signals from the auxiliary channels before executing on them cross-correlation or regression analyses. These are implemented in the bottom passages of the flowchart in figure 3.13.

First of all, the various time series to be compared with the previous method must all be sampled at the same frequency. So, a typical pre-processing step is that constituted by the *downsampling* of the various auxiliary channels time series to the frequency of the target (for example a BLRMS or the BNS range) or to a reference output frequency. This is done by some iterations of an anti-aliasing low-pass filter and a decimation operation, which consists in selecting only an equally spaced portion of the data sample. In case the sampling frequency of the original time series and the desired output are not integer multiple, then also a *resampling* operation is needed.

The most popular transformation is the *standardisation* of the data set, which converts all the time series to have zero mean and unit variance. This has computational relevance, for many numerical algorithms, when they have to compare data with very different variation scales, as for example the strain $\sim 10^{-23}$, and others sensors like magnetometers ~ 1 nT or powers measured by photodiodes ~ 0.1 mW. In particular, for the correct interpretation of the *principal components* that we discuss momentarily, it is mandatory that the signals are all converted to the same scale.

Another operation we have made use of is the computation of the derivative of certain signals. For example, for step constant signals it could be relevant to the target not their absolute value but only the instant when they change status. This is for example the case of *counter* channels, which counts some transitions or some other happening.

Analogous to the previous idea is that to introduce in this *brute force* search for correlations also non-linear terms. Indeed, as we previously mentioned, the cross-correlation is a measure of the linear relationship between two time series. If X is a Gaussian variable, $Y \equiv X^2$ is the classical example of a stochastic variables that is un-correlated with the former, but not independent. For this reason, we have included in the pre-processing stage also the possibility to compute powers of the auxiliary channels, and in this way to account for non-linear terms, up to a specified order, usually 2nd or 3rd.

Mutual Information

A popular, consistent way to account for general, non-linear, dependencies between two signals is to consider, instead of their cross-covariance, their *mutual information* [339]. This quantity measures the reduction of *uncertainty*, or *entropy*, about a stochastic variable Y after

observing X . It is defined as:³⁰

$$\begin{aligned} I(X, Y) &:= E \left[\log \left(\frac{p_{X,Y}(x, y)}{p_X(x) p_Y(y)} \right) \right] \\ &= \iint p_{X,Y}(x, y) \left(\log p_{X,Y}(x, y) - \log (p_X(x) p_Y(y)) \right) dx dy. \end{aligned} \quad (3.11)$$

The term within parenthesis in the last integral is the *Kullback-Leibler divergence* (also called *relative entropy*) between the joint p.d.f. of X and Y and the product of the individual p.d.f.s, which is a measure of how the first is different from the product of the latter [341]. From the *Jensen's inequality* [342], which relates the value of a convex function of an integral to the integral of the convex function, it can be shown that the previous quantity is always greater than, or equal to, zero [339].³¹ It is useful to understand what is the difference between it and the cross-covariance rewriting the latter as:

$$\begin{aligned} Cov(X, Y) &:= E[XY] - E[X] E[Y] \\ &= \iint xy p_{X,Y}(x, y) dx dy - \left(\int x p_X(x) dx \right) \left(\int y p_Y(y) dy \right) \\ &= \iint xy (p_{X,Y}(x, y) - p_X(x) p_Y(y)) dx dy. \end{aligned} \quad (3.12)$$

From the previous expressions we can interpret the mutual information as the value of the Kullback-Leibler divergence weighted by the joint p.d.f.. The cross-covariance is instead the “usual” difference between the joint p.d.f. and the product of the individual p.d.f.s weighted by the values of the variables, xy . In this sense, mutual information is more general and can measure *non-monotonic* relationships³² and other more non-linear relationships between the variables.

Equation (3.11) can also be extended to the case of n auxiliary variables: *joint mutual information*, and *marginal* or *incremental mutual information* (or *redundancy*), which is the difference in I when considering n auxiliary variables and $n - 1$.

Although quite promising, when applied to the search for relationships between non-stationary noise features in GW detectors data, the previous method have shown two main limitations that have made us to prefer the usual cross-correlation analysis. Firstly, to be applicable to sampled values from continuous random variables, we need to evaluate their empirical p.d.f.s. This is done partitioning the observation space into M bins, and similarly to what described for the Kolmogorov-Smirnov method in section 3.1.2, define $p(x_i) = n_i/N$, where n_i is the number of observations, out of the total N , falling into the i -th bin. Similarly for two variables: $p(x_i, y_j) = n_{ij}/N$, with M^2 bins. It can be shown that increasing the number of bins augment the value of the entropy. Conversely, reducing their number reduces the variance of the estimation but has the drawback of making it less accurate, and with higher

³⁰In terms of the Information Theory's concept of *entropy*, for two discrete random variables X, Y , whose sample space is $\Omega = \{x_i, i = 1, \dots, n\}$, $H(X) = -\sum_{i=1}^n P(x_i) \log P(x_i)$, the mutual information can be rewritten as: $I(X, Y) = H(X) + H(Y) - H(X, Y)$ [340].

³¹If $f(x)$ is a convex function, and $g(x)$ a non-negative Lebesgue integrable function, Jensen's inequality reads: $f(\int g(x) dx) \geq \int f(g(x)) dx$. A sketch of the proof of the relation $I(X, Y) \geq 0$ can be found, for the discrete variable case, in [343, p. 75].

³²For this reason, it is also more general than the *Spearman's correlation coefficient*, which measures the Pearson's cross-correlation between the “ordered” (or *ranked*) version of the values of the two time series [344].

bias [345]; refer to the definition in note 30. The dependency on the number of bins makes the results to vary significantly, with a consequent lack of credibility in them.³³

The previous issue is a consequence of a more profound one that this method usually has when applied to the study of non-stationary signals. This is the interpretation to give to any estimate of a p.d.f., or to “the” p.d.f. itself, for a non-stationary signal.³⁴

These considerations have led us to prefer the “standard” cross-covariance analysis, with the possible introduction of non-linear terms or transformed versions of the signals, instead of the Mutual Information. However, we will come across this concept again in section 3.7, where it will be fundamental for the clarification of one of the possible interpretations to give to the concept of *causality* between different signals. Moreover, it will be the criteria by means of which to represent the signals in the application of the *convergent cross mapping* causality test 3.7.2.

3.5.3 Multiple linear regression analysis

The quantities described so far are valuable for “one-to-one” comparisons between time series. Often, greater contribution to the noise may come from more complex relations, involving more than one auxiliary signal at the same time. This aspect, similar to what implemented for the multiple coherence in (2.71), can be modelled with the statistical method provided by the *multivariate linear regression* analysis.

As in the previous case, let us assume to have a *target* time series y_t and we want to investigate its relationships with a set of n *auxiliary signals* $x_{1t}, x_{2t}, \dots, x_{nt}$. These can be the outputs provided by the auxiliary channels or their transformed versions, as discussed in the previous section, like their squares $x_{kt} = x_{kt}^2$, cross-products $x_{kt} = x_{lt} \cdot x_{mt}$, or lagged versions of them $x_{kt} = x_{lt-\tau}$. The linear regression analysis consists into modelling the target as a *linear combination* of the auxiliary signals:

$$y_t := \hat{y}_t + e_t = p_0 + p_1 x_{1t} + \dots + p_n x_{nt} + e_t \quad \text{for } t = 1, \dots, N \quad (3.13)$$

where the p_i 's are the $n + 1$ real *model coefficients*, assumed not to change over the time, and $e_t = y_t - \hat{y}_t$ is an *error term*, or *residual*, introduced to account for the possibly non perfectly linear relationship between y_t and the x_{it} 's. In practice this term accounts for the addition “noise” our model has not being able to take into account; we'd want it to contain the least residual amount of information in order to consider the model predictive. These consideration will be formalised and elaborated momentarily, starting from equation (3.14).

³³From a computational point of view, for the mutual information we have used the function `mutual_info_classif` from the Python library `sklearn.feature_selection`. We have also confirmed its prediction, and most importantly the annoying dependency on M , with an “hand written” version of it. It should be mentioned that other techniques, less dependent on the previous choice, have recently (2018 – 2019) started to be investigated. Refer to [346] for an account about some of them. Unfortunately, none of these seem to be already implemented in common Python libraries, and have not been taken into account so far by the author.

³⁴To be precise, also the estimation of the cross-covariance between two non-stationary signals can be said to be not fully legitimated operation, since the Ergodic Theorem is not valid. This however can be factually interpreted as a measure of their time similarity, without recurring to any inference about their underlying statistical structures, as explained at the beginning of this section. On the other hand, Mutual Information and the estimated p.d.f. require many more assumptions, and the proof of facts has shown that the corresponding uncertainty is too large to be of any practical use.

It is convenient to collect the auxiliary signals x_{kt} , together with the constant term in (3.13), in a $(n + 1)$ -dimensional row vector:

$$\mathbf{x}_t := (1, x_{1t}, \dots, x_{nt}), \quad \text{for } t = 1, \dots, N$$

and the regression coefficients as the column vector $\mathbf{p} := (p_0, p_1, \dots, p_n)^T$. In such a way, \hat{y}_t can be rewritten as an inner product: $\hat{y}_t := \mathbf{x}_t \mathbf{p}$ for $t = 1, \dots, N$. Even more compactly, we can stack the various elements of the series corresponding to different times into *column vectors*:

$$\mathbf{y} := \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{pmatrix}, \quad X := \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \vdots \\ \mathbf{x}_N \end{pmatrix} = \begin{pmatrix} 1 & x_{11} & \dots & x_{1n} \\ 1 & x_{21} & & x_{2n} \\ \vdots & & \ddots & \vdots \\ 1 & x_{N1} & \dots & x_{Nn} \end{pmatrix}, \quad \text{and } \mathbf{e} := \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_N \end{pmatrix}$$

and rewrite (3.13) as:

$$\mathbf{y} = \hat{\mathbf{y}} + \mathbf{e} = X\mathbf{p} + \mathbf{e}.$$

Next, the goal of the regression is to find the coefficients p_i 's that make the *fit*, that is the matching between the estimate $\hat{\mathbf{y}}$ and the target \mathbf{y} , best according to some optimality criterion. For this purpose, it can be shown, from the *Gauss-Markov theorem* [347], that under the *Classical Linear Model* (CLM) assumptions the *Ordinary Least Squares* (OLS) estimator $\hat{\mathbf{p}}$ provide many desirable statistical properties, and in particular $\hat{\mathbf{p}}$ is BLUE [348]:

- Best** (minimum variance among all the other linear estimators)
- Linear** (as a function of the data)
- Unbiased** ($E[\hat{\mathbf{p}}] = \mathbf{p}$)
- Estimator** (of \mathbf{p}).

The CLM assumptions can be put in the form that the expectation values of the regression errors must satisfy:

$$E[e_t] = 0, \quad \text{and } E[e_t e_{t'}] = \delta_{tt'} \sigma^2 \quad (3.14)$$

The first assumption is automatically achieved with the inclusion of the constant term p_0 in the regression model (3.13). The last one implies that the errors must have equal variance (*homoscedasticity*) and has to be *independent*; these are usually the most subtle aspects of the analysis, and must be carefully taken into account when building the model (3.13) and in order to make reliable inference on the regression results. We will discuss this in more details in the next subsection.

Given the CLM assumptions, an overall goodness of fit measure is given by the *error sum of squares* (ESS):

$$ESS := \mathbf{e}^T \mathbf{e} = (\mathbf{y} - X\mathbf{p})^T (\mathbf{y} - X\mathbf{p}) = \sum_{t=1}^N (y_t - \hat{y}_t)^2 \quad (3.15)$$

and the OLS estimator $\hat{\mathbf{p}}$ can be found minimizing this quantity:

$$\frac{\partial ESS}{\partial \mathbf{p}} = -2X^T \mathbf{y} + 2X^T X \mathbf{p} \quad (3.16)$$

which, equating to zero, gives the unique minimum of the positive definite quadratic form (3.15):

$$\hat{\mathbf{p}} = \left(X^T X \right)^{-1} X^T \mathbf{y}. \quad (3.17)$$

From (3.14), it is immediate to verify that this estimator is unbiased and its covariance is given by:

$$\text{Cov}(\hat{\mathbf{p}}) = E \left[(\hat{\mathbf{p}} - \mathbf{p})(\hat{\mathbf{p}} - \mathbf{p})^T \right] = \sigma^2 \left(X^T X \right)^{-1}. \quad (3.18)$$

An unbiased estimate of the error variance σ^2 is the *mean squared error*:

$$\text{MSE} = \frac{1}{N - n - 1} \sum_{t=1}^N e_t^2. \quad (3.19)$$

If, besides the CLM assumptions, the regression errors are also Gaussian distributed, $\hat{\mathbf{p}}$ will also be Gaussian. In this case, reliable *t*- and *F*-tests can be carried out on the coefficient estimates to assess predictor significance, and confidence intervals can be constructed to describe estimator variance. Also, $\hat{\mathbf{p}}$ achieves the Cramèr-Rao lower bound, becoming an *efficient estimator* [349].

Principal component regression

From the previous equations, it is evident a potential problem, already anticipated in the context of multiple coherence (2.71). If the auxiliary channels have zero means (as provided by a suitable pre-processing transformation) $X^T X / N = \hat{R}_X(0)$ is the zero lag (estimator of the) cross-covariance matrix of these signals; compare it to (3.9). In the presence of collinearity among them, implying that one can be linearly predicted from the others, the previous quantity is singular (*rank deficient*) and the OLS solution (3.17) can't be found. In practice, also in the case of *quasi*-collinearity, meaning $\det \hat{R}_X(0) \approx 0$ although not strictly zero, the previous estimate has a very large variance (3.18); this makes the OLS estimate very prone to large variations from small changes in the data *X* (*ill-conditioning*), and in fact of little practical use. More details will be provided momentarily. In order to avoid this situation, which is very likely to occur with large number of auxiliary channels, we have to modify the standard approach to regression, implementing some algebraic "countermeasures" and approximations.

For the covariance matrix is real and symmetric, we can apply the *spectral theorem* [350, §6.3]:

$$\hat{R}_X(0) \cdot N = X^T X = V \Lambda V^T \quad (3.20)$$

where $\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n)$ is a diagonal matrix of eigenvalues, corresponding, besides the constant factor N , to the variances of the data in the orthogonal directions given by the columns of $V := (\mathbf{v}_1, \dots, \mathbf{v}_n)$, which constitute a basis of orthonormal eigenvectors.³⁵ Then, every data point x_t can be reconstructed as a linear combination of these eigenvectors. If the eigenvalues are ordered such that $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq 0$, the previous decomposition

³⁵For this operation is a projection of the original data onto directions which maximise the variance, it is fundamental for the consequent interpretation that the various signals have been previously standardised to unit variance, as discussed in section 3.5.2.




Figure 3.12: Explained variance, normalised to 1, for a set of Virgo environmental channel signals (ENV). The dashed line represents the 95% of the variance, and the corresponding number of independent components. In this example, with just about 70 independent components we were able to explain 95% of the variance of 140 Virgo environmental channels.

is known as *Principal Component Analysis* (PCA). Then, if we retain only the first $m < n$ eigenvectors v_k , it can be shown that every x_t can be approximated with a mean squared error equals to $\sum_{k=m+1}^n \lambda_k$ [351]. To understand this in a practical situation, in figure 3.12 the explained variance $\sum_{k=1}^m \lambda_k$ for the signals from a set of Virgo Environmental sensors is reported as a function of their number. It is interesting to note that just $m \sim 70$ of them, out of about 140, are able to explain 95% of their variance, and with $m \gtrsim 100$ their covariance matrix (3.20) became singular. Indeed, it is not unreasonable to expect, with the thousands of sensors that monitor the apparatus, that some collinearity shows up.

This gives a recipe to get rid of collinear signals, corresponding to null eigenvalues, and find the OLS solution to (3.16) as [350, §7.3]:

$$\hat{p} = (V\Lambda^+V^T)X^T\mathbf{y}$$

where $\Lambda^+ = \text{diag}(1/\lambda_1, \dots, 1/\lambda_m, 0, \dots, 0)$ is the *pseudo-inverse* of Λ , made with the inverses of all the m non-zero eigenvalues, and zeros in the $(n - m)$ remaining positions. This solution is therefore insensitive to collinearity issues. Furthermore, the previous expression can be simplified introducing the *Singular Value Decomposition* (SVD) of the data, $X = U\Sigma V^T$, where $\Sigma^T\Sigma = \Lambda$ and the columns of U are given by the eigenvectors of XX^T :

$$\hat{p} = (V\Sigma^+U^T)\mathbf{y}. \quad (3.21)$$

In the previous form the OLS solution becomes quite fast to compute numerically, even for large data sets, for efficient algorithms exist for computing the SVD of X without having to form the product X^TX [40, 352].³⁶

³⁶Let's report some considerations on the computational cost of the proposed solutions. OLS, where the decomposition of X^TX is usually done via the *Cholesky algorithm* [353], requires $n^3 + Nn^2/2$ operations and, depending

This strategy can be also applied to the case of quasi-collinearity of the auxiliary signals, that is, when some $\lambda_k \simeq 0$, although not strictly equals to zero. In this situation the problem of finding the OLS estimate is said to be *ill-conditioned*, being the *condition number* of $X^T X$ very large: $\kappa = \lambda_{\max}/\lambda_{\min}$. Consequently, the solution $\hat{\mathbf{p}}$ will be affected by a correspondingly high variance (3.18). To avoid this, it is often preferable to get rid of some of the smallest eigenvalues, equating a number $n - m$ of them to zero. This implies a reduction in the covariance of the estimate $\hat{\mathbf{p}}_m$ with respect to the full OLS $\hat{\mathbf{p}}$, which we can compute to be [348, §12.2]:

$$\text{Cov}(\hat{\mathbf{p}}) - \text{Cov}(\hat{\mathbf{p}}_m) = \sigma^2 \sum_{k=m+1}^n \frac{\mathbf{v}_k \mathbf{v}_k^T}{\lambda_k} \succeq 0$$

where \mathbf{v}_k is the k -th column of V , and the symbol “ \succeq ” indicates a *positive semi-definite* matrix. The price to be paid is that this estimator is no longer unbiased, being constrained to be perpendicular to the last $(n - m)$ eigenvectors \mathbf{v}_k . Its bias is usually very small, and can be quantified by the approximation error $\sum_{k=m+1}^n \lambda_k$. This implies that, by choosing to exclude only the very smallest λ_k 's, the loss due to the bias is usually less significant than the gain in the covariance reduction, and than the estimator $\hat{\mathbf{p}}_m$ can still be *more efficient* than $\hat{\mathbf{p}}$. The problem of finding the “adequate amount” of principal components to discard from the analysis is very common in statistics and *Machine Learning*, and it is usually referred as “the bias-variance *trade-off*” [40], or *dilemma*, sometimes.

The previous technique, called PCA regression, natively allows to adjust towards which direction make the bias-variance trade-off to lean. Moreover, the transformation from the “space of the auxiliary channels” to their principal components directions can give a deeper insight about what are the underlying uncorrelated mechanisms that are generating the noise. The *dimensionality reduction*, accomplished discarding some of the lowest energetic principal components, allows also to express the results in terms of a small number of uncorrelated contributions, instead of a possibly large sum of correlated ones.

However, it must be noted that all the previous considerations aim at improving the coefficients estimate, reducing its covariance, disregarding the overall regression prediction, as can be quantified by (3.15). In general, neglecting part of the data increase the *ESS* (3.15), and, as pointed out by some authors [351], even the lowest variance components of $\hat{R}_X(0)$ can play a crucial role in predicting the target \mathbf{y} . To overcome this issue, finding a trade off between a low covariance estimate and a predictive regression model, some authors have proposed to use a *supervised Principal Component Regression* algorithm [40, 354, 355], which we have implemented in the analysis tools we are going to describe in section 3.5.4. Instead of zeroing all the smallest eigenvalues of the cross-covariance matrix, we can remove a subset of them up to a fixed number or up to a fixed energy content $\sum_k \lambda_k$, according to some criterion such as the smallest cross-correlation of the corresponding eigenvectors with the target signal \mathbf{y} , or the smallest value of the *ESS* obtained in such a way. For our purposes, we have chosen the latter criterion, and a total amount of discarded principal components energy of $5 \div 10\%$.

All of these choices constitute part of the model construction. In order to compare different models, we will introduce in the next subsection several goodness of fit benchmarks.

on the relative size of N and n , it can be numerically unstable. The same is for Lasso regression. SVD, through *QR decomposition*, requires Nn^2 operation and is usually more stable and faster [40].

Other regression methods: subset selection and shrinkage methods

From figure 3.12 and the previous considerations, it should be clear that standard regression analysis is not suitable for “brute force” investigations. This approach is generally doomed to result in a singular cross-correlation matrix, or in an ill-conditioned model estimation. PCA regression is just one method to overcome this difficulty. Some others have been investigated in GW literature, and many more in Statistics and Machine Learning ones, which we are going to briefly describe.

To the best of our knowledge, the first attempt to correlate noise features in the strain signal with the auxiliary channels, was implemented by G. Vajente in the original version of the NonNA analysis tool [LVC7], which used to run on Virgo data back in 2013. The rebuilt and update of the code was part of the work of the author and it will be described in section 3.5.4. This method implemented a *backward-stepwise selection* [40]. It started with the full set of channels by means of which we want to model the target, and, step by step, the least contributing one to the estimation error reduction, or some other criterion, was discarded. This process was iterated up to a fixed number of remaining channels, or until all of them was ranked. Needless to say that this was computationally very expensive, hence suitable only for quite a moderate number of channels. In general, without some sort of pre-selection, the first iterations resulted in singular cross-covariance matrices, and a random channel was discarded. For these reasons this approach was soon abandoned in the new version of the code for Advanced Virgo.

In literature, a common solution to the previous problem can be provided by the so called *shrinkage methods*, which can be characterised by the following *Lagrangian form* for the model parameters estimator [40]:

$$\hat{\mathbf{p}} = \underset{\mathbf{p}}{\operatorname{argmin}} \left\{ \|\mathbf{y} - X\mathbf{p}\|^2 + \lambda \sum_{j=1}^p |p_j|^q \right\}$$

for $q, \lambda \geq 0$. The first term inside braces is identical to (3.15). The additional one is a “penalty factor” that weights the number of non-zero parameters for the model by means of the *complexity parameter* λ ; the larger its value, the more the shrinkage of the parameters to values close to zero. For $q = 2$, corresponding to an ℓ_2 penalty, we have the so called *Ridge regression*, while for $q = 1$, ℓ_1 penalty, we have the *least absolute shrinkage and selection operator* or *Lasso regression*, currently implemented in the same-name algorithm in LIGO [356]. This algorithm is mainly used to relate the variations of BNS range of the two LIGO detectors to a small ($5 \div 10$) subset of auxiliary channels by arbitrarily varying the parameter λ . As well as PCA regression, these regularisation methods introduce a bias in the parameter estimation, but posses a smaller *MSE* with respect to the OLS solution.

A thorough comparison of the previous methods, including PCA regression, can be found in [40], by some of the authors and developers of these techniques. Since the ℓ_1 penalty has the effect of “truncating” the smallest coefficients to zero, Lasso is generally preferred over Ridge when the solution is believed to have sparse contributions from the auxiliary channels, and in general not strong multi-collinearity (which may or may not be the case depending on the channels selection in noise studies). Instead, the ℓ_2 penalty does a proportional shrinkage making to prefer Ridge regression n case of known multi-collinear

channels. Sometimes, a weighted mixture of the two is desired, which is called *elastic net*. Refer to [357] for an account about this, and also for some interesting considerations about which shrinkage method to prefer, depending on the context and the prior information. PCA regression does an “hard truncation”, like Lasso, but on the least energetic principal components instead of on the channels themselves.

In the context of noise characterisation, it is interesting to notice that PCA regression and Lasso regression (and “brute force” correlation) give substantially a complementary information, being the former more focussed on the underlying uncorrelated noise contributions and the latter on those from a main subset of individual channels. From a commissioner’s point of view, maybe the first is more suitable to answer the question “what kind of noise is affecting the detector?” while the latter to “where to intervene?”, since PCAs are not referred to physical sensors and places in the detector, whilst the latter sensors channels are.

Residuals diagnostics and model selection

Before attempting to interpret any regression result and try to make inference from it, it is necessary to verify the model assumptions, such as the CLM assumption described before. If the model residuals $e_t = y_t - \hat{y}_t$ depart markedly from this, the estimate obtained in such a way doesn’t benefit of the Gauss-Markov theorem, and the regression is unlikely to perform well, either in explaining variable relationships or in predicting the target channel.³⁷

Many statistical tests have been developed to assess the CLM assumption, and we have implemented some of them in our regression analysis tool. To verify the Gaussianity of the residuals, our algorithm performs some of the most common tests, including Shapiro-Wilk [358], Kolmogorov-Smirnov [359] and Jarque-Bera’s [360] tests. Although not necessary for the Gauss-Markov theorem, fulfilling the latter condition is of great help for testing the CLM assumption and, most of all, for constructing confidence intervals and hypothesis test for the prediction significance using standard techniques (refer to the discussion in the previous section). As previously mentioned, despite being convenient, the interpretation of having Gaussian and uncorrelated residuals is that all of the information in the target signal has been correctly captured by the regression model, leaving to the residuals only the inherent stochasticity of the processes (and possibly a sum of all the other omitted variables, which approach the conditions for the central limit theorem [263]).

Independence of the residuals can be visually investigated plotting their auto-correlation function, $R_e(\tau)$, and checking that for any lag $\tau \neq 0$ this is smaller than 1 in absolute value. The Durbin-Watson statistic is the traditional test for the presence of first-order auto-correlation, and suitable p -values can be computed for it [361]. Both of these checks are implemented in our tool, and reported in the example in figure 3.17.

Unequal variance of the residuals, or *heteroscedasticity*, can be tested by means of the Engle’s ARCH test for Auto-Regressive Conditional Heteroscedasticity [362]. It is meant to test the null hypothesis that the squared residual e_t^2 can’t be predicted by a linear combination

³⁷In modern Machine Learning applications, this seems to be no more a concern, as long as the estimation variance and the *overfitting* issue are taken under control. Indeed, the primary criteria on the base of which to judge the goodness of a model, estimated on an observed set of data (namely the *training set*), is its *ESS* computed on a different *test* set of previously unobserved data, and other related quantities to be introduced in equation (3.22). We will come back to discuss this aspect related to the predictivity of a model momentarily.

of its lagged values plus a white noise term, $w_t: e_t^2 = a_0 + a_1 e_{t-1}^2 + \dots + a_L e_{t-L}^2 + w_t$; refer to the discussion in section 2.1.3. In practice, this test makes a second regression analysis on the squared residuals, verifying this CLM assumption.

As a consequence of the violation of the last two assumptions, one may usually obtain less efficient estimators; this implies underestimating their confidence intervals and making overly optimistic claims of accuracy. In this case, many countermeasures have been developed, like the adoption of the *generalized* least square instead of the OLS method, and variance stabilizing transforms [348, §5-6]. Instead of implementing these solutions in our algorithm, the occurrence of this kind of violations has been interpreted as a consequence of model misspecification, and hence we have consequently provided to modify the model, rather than the analysis, adding additional terms, couplings or lagged variables to the regression model under consideration. Indeed, the presence of correlated residuals is usually a consequence of the exclusion of some relevant auto-correlated variable from the model, while heteroscedasticity is often the result of interactions between model variables and omitted variables.

A standard measure of goodness of fit of the adopted regression model is the *coefficient of multiple determination*, or colloquially “ R^2 coefficient”, defined as the ratio of the *regression sum of squares* (RSS) and the *total sum of squares* (TSS):³⁸

$$R^2 := \frac{\text{RSS}}{\text{TSS}} = 1 - \frac{\text{ESS}}{\text{TSS}} = \frac{(X\hat{p})^T(X\hat{p})}{\mathbf{y}^T\mathbf{y}} = 1 - \frac{\mathbf{e}^T\mathbf{e}}{\mathbf{y}^T\mathbf{y}}. \quad (3.22)$$

This number corresponds to the fraction of the variability of the target \mathbf{y} that can be explained by the regression model $X\hat{p}$; the closer to 1 the better, but taking into account not to commit the *overfitting* error (the “curse of dimensionality” [363]), including in the model a number of variables greater than, or comparable to, the number of observations: $n \gtrsim N$. For this reason it is often preferable to make use of the *adjusted determination coefficient* R_{adj}^2 , defined as:

$$R_{\text{adj}}^2 = 1 - (1 - R^2) \frac{N - 1}{N - n - 1} \quad (3.23)$$

which is always less than R^2 , and identical to (or smaller than) zero if the number of variables is equal to (respectively smaller than) the number of observations. This quantity accomplish the principle of *parsimony*, and ease of interpretability, in the selection of the variables to include in the model. In the case of PCA regression, this metric can be used as the criteria by means of which to discard certain least energetic components.

To formally infer whether the proposed model is effective in representing the target, one can test the following hypotheses:

$$\begin{aligned} H_0 : & \quad p_1 = \dots = p_n = 0 && \text{(no linear relationships)} \\ H_1 : & \quad \exists p_k \neq 0 && \text{(at least one witness variable)} \end{aligned}$$

³⁸Beware that many authors use “E” for explained and “R” for residuals, inverting their definitions of ESS and RSS with respect to ours.

From equation (3.19) we can notice that, under the assumption that the residuals are $e_k \sim N(0, \sigma^2)$, the *ESS* is distributed like a χ_{N-n-1}^2 , and the test statistic

$$\frac{MSR}{MSE} = \frac{RSS}{ESS} \frac{N-n-1}{n} \sim F_{n, N-n-1}$$

that is, it is distributed as an *F*-variable with n and $N-n-1$ degrees of freedom [348, §3.4]. Therefore, if the previous quantity is greater than F_α , the value corresponding to the $1-\alpha$ distribution percentile of *F*, we can reject the null hypothesis (H_0) with a level of significance α , and assess that our model has been predictive. Similarly, we can make *t*-test to assess whether the inclusion of a particular variable is relevant or not for the model:

$$\hat{p}_k / \sqrt{Cov(\hat{\mathbf{p}})_{kk}} \sim t_{N-n-1} \quad (3.24)$$

where $Cov(\hat{\mathbf{p}})_{kk}$ is the (k, k) -element of the covariance matrix of $\hat{\mathbf{p}}$ (3.18). From the previous equation we can also define the *confidence intervals*:

$$C.I.(p_k)_\alpha = \left[\hat{p}_k \pm t_{N-n-1, \alpha/2} \cdot \sqrt{Cov(\hat{\mathbf{p}})_{kk}} \right]$$

where $t_{N-n-1, \alpha/2}$ is the critical value corresponding to the upper $\alpha/2$ of the cumulative *t*-distribution with $N-k-1$ degrees of freedom. Similar results can be found for the confidence interval for \mathbf{y} [348, §3.8].

With the previous quantities and tests at our disposal, we can now face the problem of *model building*. An adequate model (3.13) should accomplish the criteria of *necessity*, *parsimony*, *sufficiency* and *stability*, meaning that all the variable included in the model should contribute to the prediction and no additional variable should improve and/or change the coefficient estimate significantly.

If we have several models among which to choose, assumed all to produce normally distributed residuals $e_i \sim N(0, \sigma^2)$, we can compute the *likelihood function* of their parameters:

$$\mathcal{L}(\mathbf{p}, \sigma^2 | \mathbf{y}) = (2\pi \sigma^2)^{-N/2} \exp \left[-\frac{1}{2\sigma^2} (\mathbf{y} - X\mathbf{p})^T (\mathbf{y} - X\mathbf{p}) \right].$$

Substituting the parameters estimate $\hat{\mathbf{p}}$ and *MSE* for each of them, we can compare different models on the base of the value assumed by their likelihood, $\hat{\mathcal{L}}$; we will opt for the one with the highest estimated likelihood.

Other criteria related with the likelihood function are the *Akaike Information Criterion* (AIC) [364] and the *Bayesian Information Criterion* (BIC) [365]. The former provides a relative estimate of the information loss between different models, while the latter is based on Bayesian arguments. Both of them introduce a cost for models with large number of parameters, usually more severe for BIC, penalizing their complexity:

$$AIC = 2n - 2 \log \hat{\mathcal{L}} \quad (3.25)$$

$$BIC = n \log(N) - 2 \log \hat{\mathcal{L}} \quad (3.26)$$

In this case, a model with lower AIC or BIC is preferable.

Parameter name	Description
target	Name of the channel to choose as a target for the analysis
aux name spec.	Tags of the auxiliary channel names to include in the analysis
excluded spec.	Tags of the aux. channel names to exclude from the analysis
gps start	Start time, in gps or UTC units
duration, gps end	Duration in seconds or end time in gps or UTC units
output frequency	Downsampling frequency for studying slow variations
order	Maximum power of the signals to include in the analysis
band(s)	Frequencies where to compute the BLRMS(s)

Table 3.1: Main input arguments to pass to the NonNA analysis tools.

3.5.4 NonNA: Non-stationary Noise Analysis tools

The correlation and regression methods described in the previous sections have been implemented in the rebuilt of the Detector Characterisation tool NonNA, Non-stationary Noise Analysis tool [LVC34]. Two versions of it are available; one that performs a cross-correlation analysis and the other a PCA regression. These have become a reference tool for noise investigations in Virgo, such as those related to environmental noise [29], and their results are collected in the corresponding web area [LVC37]. They have helped understanding the origin of various glitches during O3, like those in figure 3.11 or those related to the *pre-stabilised laser* and the *pre-mode cleaner* [LVC38, LVC39], of slow non-stationarities, for example related to the weather conditions [LVC40], and of various BNS range drops [LVC41]. The cross-correlation tool has also been routinely used for the investigation of the “infamous” 83 Hz wandering line, whose overview is presented in appendix A.3. In this section we describe their working principle, comprising the implementation and summary of the previous theoretical concepts, as well as their usual parameter configurations.

In figure 3.13 we have described a schematic flowchart of the preprocessing stage common to both the cross-correlation and regression analyses. The main configuration parameters for our tools are listed in table 3.1, which we are going to comment.

Commonly chosen targets are the BNS range, whose drops and variations are typical noise features it is important to understand, the BLRMS’ of LSC_DARM or the strain channel, and also wandering lines, tracked with the method described in section 3.4. The auxiliary channels can be chosen among the about 80k sensor channels that monitor the detector. These can also undergo some of the transformations described in section 3.5.2. In the case their BLRMS’ is desired, their acquisition is from the Virgo *raw* data frame, with channels at their full sampling frequencies. Otherwise they are read from the *trend* data frame, with all the channels sampled at 1 Hz, which is usually sufficient for slow non-stationary noise investigations. Typical choices for them are the environmental channels and the alignment control channels: ENV_, ASC_ and LSC_. A list of about 400 “standard” channels, customary used for Detector Characterization purposes, is used by default. Some channels with known correlations with the targets, strain and DARM, are excluded from the analysis; these include all the channels related to the detection photodiode or used in the reconstruction of $h(t)$.

It is important to note that the operation of reading hours or even few days of data, in particular from the raw frame, is quite memory demanding. Large efforts have been put into developing efficient acquisition and pre-processing algorithms for these analyses.

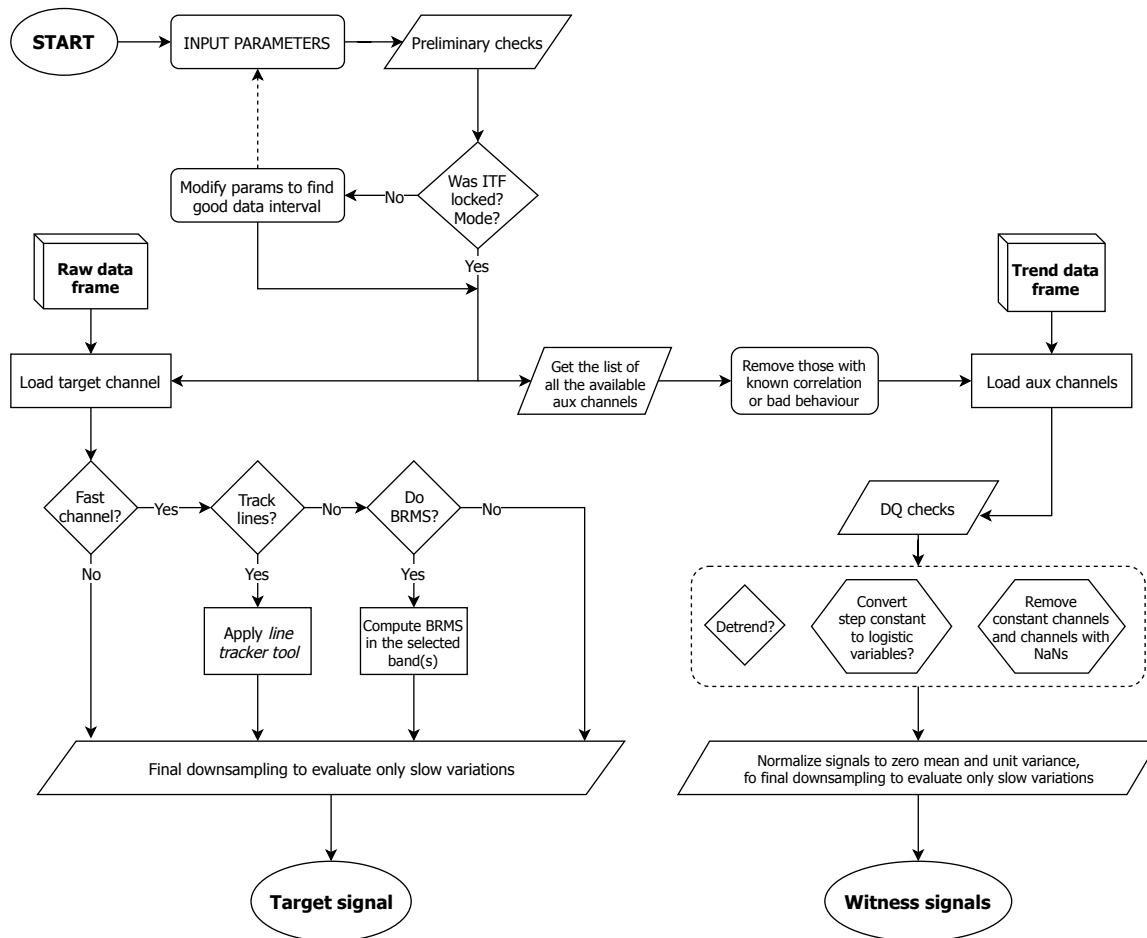


Figure 3.13: Flowchart of the preprocessing stage for the cross-correlation and regression analysis tools.

This was done implementing a parallel structure in the code, which exploits multiprocessing computing over the up to eight cores of the Virgo farm machines [LVC4]. Some limitations are still present when reading from the raw data frame and computing the BLRMS' of large lists of auxiliary channels, while for the single target channel the maximum data length can be up to a week.

Usually, as a first step, our algorithm checks if the interferometer was locked, or in a specified status as described in section 2.4.3, at the selected start time for the analysis and for all the specified duration. If this is not the case, the start time is moved to the beginning of the next interferometer locking and/or the duration is reduced up to the next unlocking. For long duration correlation analyses, this option can be disabled, provided the data segments corresponding to no data are filled in a suitable way.³⁹

The auxiliary channels can be specified by their names or making use of some “wild-cards”, like in “ENV_*”, representing all the *environmental* channels. Similar operation is done for all the channels that we want to exclude from the analysis, for example because we already know of their correlation with the target or because not reliable, for example

³⁹Usually with non-numeric values (NaNs), which are automatically ignored by the analysis.

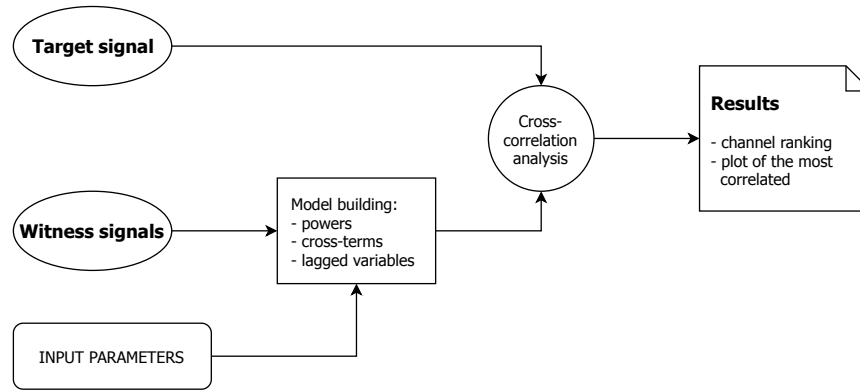


Figure 3.14: Flowchart of the cross-correlation analysis tool.

for some malfunctioning or undesired behaviour. All the available channels at the specified analysis time matching with the specifications provided are selected, and all the ones whose names matching the exclusion specs excluded. Usually, this list comprises about 40k channels for the “brute force” cross-correlation analysis, while for the regression analysis the starting model usually comprises few hundreds of channels. Next, there is the pre-processing phase described before, at the end of which all the channels should match the desired output frequency, the same as the target.

The pre-processed signals are then correlated by means of the cross-correlation and the regression analysis tools, whose flowcharts are represented in figure 3.14 and 3.15. The output of the former is shown in figure 3.16. It consists into an html page with a table containing, for every frequency band analysed, the names of the most correlated auxiliary channels, ranked on the base of (the absolute value of) their Pearson’s correlation coefficient with the target. The darker the colour the higher (the absolute value of) the correlation.⁴⁰ The name of each channel is a link to the corresponding plot of its time series, along with that of the target, analogously to what shown in figure 3.11. In the same page we also put the link to a *log file* containing a longer list of correlated channels. The results of the analysis are available at the corresponding Virgo Data Analysis web page [LVC37].

The regression tool is more complex. Many model parameters are meant to change through iteration of the algorithm, as described in the previous sections. For example, the *glitch removal* algorithm, described in section 3.1.2, is applied to the target channel as an additional pre-processing. This can be done with or without setting from the beginning the “outlier threshold”. Its value can be decided by the algorithm itself on the base of the analysis of the residuals and the value of the (adjusted) determination coefficient (3.23). For example, it can start with none of the outliers removed, and it lowers the threshold until finding a maximum for R_{adj}^2 (or a minimum of AIC or BIC (3.25-3.26), in case of Gaussian residuals).⁴¹ This choice has been made to account for the possibility that some of the auxiliary channels is able to explain the “outliers” (mostly glitches) in the target; if not, these are removed and the focus is given to the other non-stationarities.

⁴⁰This design has been inspired by that of Bruco, for studying the “brute force” coherence, already described in section 2.3.3.

⁴¹Recall that R_{adj}^2 depends on the ratio of the parameters in the model and the number of observations; so, removing too many outliers is penalized.

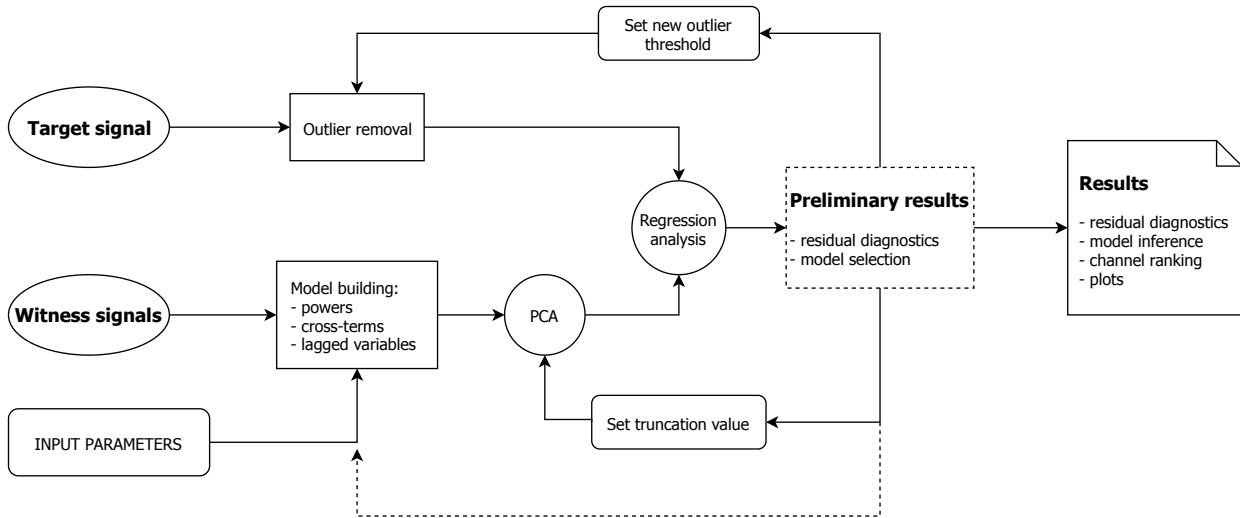


Figure 3.15: Flowchart of the regression analysis tool.

Also, the choice of which eigenvector remove in the truncated PCA is computed for a fixed amount of energy, $\sum_k \lambda_k$, in order to maximize the R_{adj}^2 (or minimize AIC and BIC). Empirically, it turned out that for $\sum_k \lambda_k \simeq 1\%$, the condition number was $\kappa \simeq 20 \div 80$, and none of the lowest eigenvalue components influenced significantly the regression determination coefficient. Other recursive operations (dashed line in figure 3.15) can be implemented in order to reduce the model complexity, reducing the dimensionality of the witness channel space, and to make the interpretation of the results easier.

The results of the regression tool are a *log-file*, containing all the relevant results for making statistical inference as described in section 3.5.3, and the plots of the target prediction and the residual diagnostics.⁴² In figure 3.17 we report the regression analysis of the amplitude non-stationarity already shown in figure 2.11, consisting in an excess of power in the band around 50 Hz of the channel LSC_DARM. A regressive model comprising 100 environmental sensor channels was able to reproduce quite accurately the shape of this noise feature, with an $R_{\text{adj}}^2 \simeq 84\%$. Those that contributed the most, on the base of their *t*-statistic were some magnetometer sensors in the Virgo Injection Electronic Room and in the West End Building.

3.5.5 Possible extensions: Machine Learning implementation

The previous correlation analyses have been tested during the commissioning phases that preceded the third observing run of the Advanced detectors (O3), and in particular during the *commissioning runs* C10 and C11, of August and October 2018 respectively [LVC43, LVC44]. The main purpose was to correctly model non-stationarities in a main target, related to the reconstructed strain, by means of the auxiliary channels, providing also the statistical quantification of their relationship. Related to this, it was also interesting to explore the

⁴²Also, some attempts have been made to try to represent the principal components in a convenient format. The idea was to exploit the Virgo channel names convention [LVC42], `v1:SUBSYSTEM_LOCATION_SENSOR_extras`, in order to identify which subsystems, locations, and kind of sensors were witnessing particular noise components. Unfortunately, this convention is not comprehensively adopted, and the simple inspection of the names that resulted most correlated provided substantially an analogous information.

NonNA cross-correlation analysis tool

by Francesco Di Rienzo
Version: 1.0, of 7/03/2019

This tool investigates the non-stationary behavior of a target signal correlating it with a set of auxiliary channels. The key statistics is the [Pearson correlation coefficient](#). The auxiliary channels are then ranked on the basis of this.

In case of necessity: [send me an email!](#)

Configuration parameters

Process id: test_2019-11-12_10-40-49
Operator: direnzo

Command string

NonNA_tool/s/nonna_corre.py -t hrec_hofc_20000Hz -b [[7,10],[10,15],[15,17],[17,29],[20,40],[40,45],[45,55],[55,60]] -n [ENV_*,ASC_*] -B 19-11-11-12 -d 19-11-11-18 -f -.02
Type nonna_corre.py -h for displaying the help string

Parameters

Target_Hec: hofc_20000Hz
Freq_bands (60 bins): [[7, 10,0], [10,0, 15,0], [15,0, 17,0], [17,0, 29,0], [20,0, 40,0], [40,0, 45,0], [45,0, 55,0], [55,0, 60,0]] Hz
Aux_name_specs: [ENV_*, ASC_*] (1597)
Ops_start: 1257510204.0
Ops_end: 1257510204.0
Operation: 20.0 seconds (in LNS)
Output fs: 0.02 Hz

Analysis results

Analysis duration: 0 hours, 7 minutes, and 27.95 seconds

This is the log file with the analysis summary: [test_2019-11-12_10-40-49.log](#)

Ranking

Channel ranking: from left to right the channels are ranked on the base of their correlation with the target (0rms). Cell colors represent these values.

Frequency band [Hz]	Aux channels correlation													
[7, 10,0]	ENV_DT_F4_TE2 -0.17	ENV_DT_F0_TE2 -0.14	ENV_WEB_RF_6MHz_dfreq_FS_mean -0.12	ENV_CEB_S_TE8 0.12	ENV_WEB_NI_TE9 -0.12	ENV_MCB_IPS_CURR_T_min -0.12	ENV_DT_F0_TEL -0.12	ENV_TGS_HWS_HR_NI_TE -0.12	ENV_WEB_NI_TE10 -0.12	ENV_PP_LINK_ACC_Z_mean 0.12				
[10,0, 15,0]	ASC_B4_QD1_Sum_rms 0.28	ASC_B3_QD1_Sum_rms 0.28	ASC_B3_QD2_Sum_rms 0.28	ASC_B5_QD2_Sum_min -0.27	ENV_WEB_SEIS_W_min -0.27	ASC_B2_QD1_Sum_max 0.26	ASC_B4_QD2_Sum_rms 0.26	ASC_B2_QD2_Sum_max 0.26	ASC_B3_QD1_Sum_mean 0.26	ASC_PP_X_CORR_rms 0.26				
[15,0, 17,0]	ASC_B5_QD2_Sum_min -0.27	ASC_B2_QD1_Sum_max 0.27	ASC_B4_QD1_Sum_rms 0.27	ASC_B2_QD1_Sum_mean 0.27	ASC_B3_QD1_Sum_rms 0.27	ASC_B2_QD2_Sum_min 0.27	ASC_Adj_ellipsoid_time_min 0.27	ASC_B5_QD2_Sum_min 0.27	ASC_B4_QD1_Sum_min 0.27	ASC_B2_QD2_Sum_min 0.26				
[17,0, 29,0]	ENV_WEB_SEIS_W_min -0.41	ENV_WEB_SEIS_W_min 0.41	ENV_WEB_SEIS_W_min 0.36	ENV_CEB_SEIS_V_mean 0.22	ASC_B4_QD1_V_max 0.21	ASC_B4_QD2_H_min 0.21	ASC_Adj_ellipsoid_time_mean 0.19	ASC_Adj_ellipsoid_time_max 0.19	ASC_B4_QD1_V_min 0.21	ASC_Adj_Time_FS_min 0.21				
[20,0, 40,0]	ENV_WEB_SEIS_W_min 0.34	ENV_WEB_SEIS_W_min -0.34	ENV_WEB_SEIS_W_min 0.31	ASC_B4_QD1_V_max 0.30	ASC_Adj_Time_FS_max 0.19	ASC_Adj_ellipsoid_time_mean 0.19	ENV_MCB_IPS_VOILT_T_min 0.15	ENV_MCB_IPS_VOILT_T_max 0.15	ASC_B4_QD1_V_min 0.18	ENV_CEB_S_TE7 -0.18				
[40,0, 45,0]	ENV_WEB_SEIS_W_min 0.26	ENV_WEB_SEIS_W_min 0.25	ENV_WEB_SEIS_W_min -0.22	ENV_WEB_MFC_mean 0.16	ENV_WEB_MFC_mean 0.16	ENV_WEB_MFC_mean 0.15	ENV_WEB_MFC_max -0.15	ENV_WEB_MFC_max -0.15	ENV_WEB_SEIS_N_min 0.15	ENV_WEB_RF_6MHz_min_FS_min 0.15				
[45,0, 55,0]	ENV_B4_MFC_mean -0.68	ENV_B3_MFC_mean -0.40	ASC_B4_QD1_V_mean 0.28	ENV_WEB_MFC_mean 0.26	ENV_WEB_MFC_mean 0.24	ASC_B4_QD1_H_mean -0.24	ASC_B5_QD1_V_mean -0.22	ASC_B4_QD2_V_6MHz_Q_mean -0.19	ASC_B4_QD2_H_mean -0.18	ASC_B4_QD2_V_6MHz_Q_min -0.18				

Figure 3.16: Example of web page output of the NonNA cross-correlation analysis tool. The table in shades of blue represents the list of the most correlated channels at each frequency band where the target BLRMS has been computed into.




Figure 3.17: Example of regression analysis with the NonNA tool referred to the amplitude non-stationarity of figure 2.11, whose BLRMS is reported in the top left plot. In the same plot is shown (red) its best prediction and 90% confidence interval. In the left bottom plot the residuals between target and prediction is shown in orange. On the right, two tests about Gaussianity and the autocorrelation of the residuals are reported.

possibility to *predict future, unobserved values* of the target by means of the estimated \hat{p} with the past values of X and y . This is the typical interest of regression analysis in the context of Machine Learning [366]; refer also to note 37.

Unfortunately, this “extrapolation process” applied to the context of non-stationary noise resulted in very poor performances, as measured by (3.15) or (3.22), except for those cases where the non-stationarities exhibited some recurring pattern, and therefore it could have been described by a *stationary model*, in the interpretation given in section 2.1.3. Also “interpolation”, that is, randomly deselecting some of the observed data, and testing the algorithm performance on them, produced results strongly dependent on their choice; this can be clearly understood from figure 3.17. In the language of Machine Learning, our regression models poorly performed in *generalising* to new data [40].

The main reason for that is in the assumption that the model coefficients p in (3.13) were fixed, that is *time invariant*. For generic non-stationary noise studies, there is no need for this to be true, and consequently the possibility to make predictions on the base of their estimated values at a certain time.

Nonetheless, it is interesting to report a couple of situations where regression models, similar to the one discussed in dissertation, have proven to contribute to understand specific noise features in GW data, although not regarding non-stationary noise. These comprise in particular the structures around the 60 Hz spectral line due to the electric mains in LIGO detectors (similar to those around the 50 Hz line studied in section 2.3.3) and its harmonics. In [274] the authors develop a regression model similar to (3.13) except for the inclusion of a certain number (~ 10000) of past and future values of the auxiliary channels (*Volterra series* [367]), representing their correlation structure with the target. Then, they make use of a neat trick to reduce the complexity of the inversion of a normal equation with that many parameters, transforming the model to a mixed time-frequency representation by means of a Wilson-Daubechies-Meyer transform (WDM) [368]. In such a way, making use of some witness channels related to the electric current monitors, they have been able to predict and filter out the spectral lines corresponding to 60 Hz and its harmonics. Moreover, exploiting a paradigm similar to the one discussed in section 2.3.4, they were able to predict the sidebands of these lines, as consequences of the up-conversion of low frequency disturbances non-linearly coupled with the previous resonances.

Another very recent example of application of a similar technique is that presented in [328].⁴³ The authors attack the same problem faced in the previous work by means of a *Deep Neural Network* (DNN) [40] to model the Volterra series that represents the relation between the target and a certain number of auxiliary channels, together with their past values. This technique constitutes a modern approach to solve non-linear problems, and indeed it can be proficiently applied to the subtraction of non-linear noise couplings in GW data, as shown by the authors. We will describe an example of application of neural networks in the next chapter, about the *deglitching* project of GW data [LVC18, 285].

⁴³Beware to a subtle difference in terminology with respect to what we have described in this manuscript. The authors indeed refer to *non-stationarities* in the data, rather than in the process generating them. For example, the sidebands described in section 2.3.4 and originating from the *non-linear* coupling of some noise of characteristic frequency ~ 0.1 Hz and the mains (either 50 or 60 Hz) are considered by them as non-stationarities in the observed data. To this purpose they clarify: “The distinction between a non-linear and a non-stationary noise coupling is simply a matter of time scales or frequencies.”

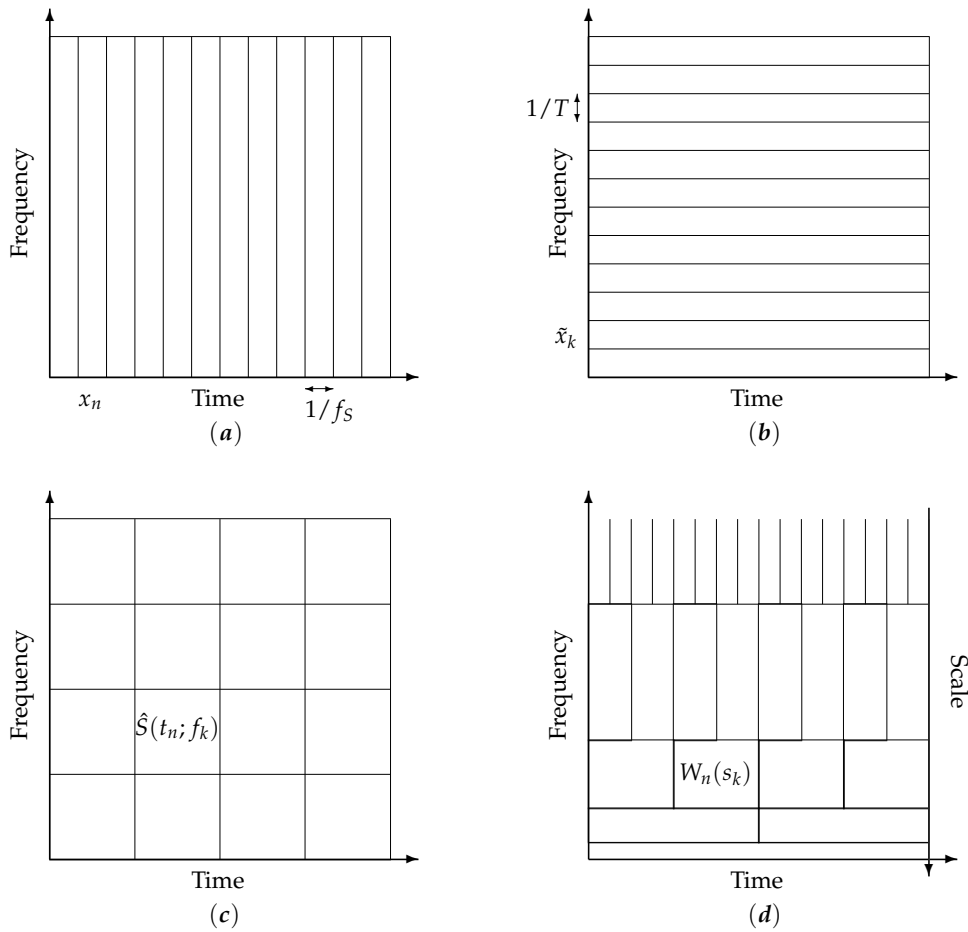


Figure 3.18: Different representations of a signal: (a) time series, (b) Fourier series, (c) short-time Fourier transform, by means of which spectrograms are computed, and (d) wavelet transform, with *octaves* frequency division (refer to equation (3.35)), yielding a scalogram representation of the signal.

3.6 Wavelet coherence analysis of glitches

In the previous section and in 2.3.3 we have presented two methods for the study of detector noise. The former, suitable for non-stationary noise, is completely implemented in the time domain representation of the signals, while the latter, for stationary noise, is completely in the Fourier domain. It was somewhat natural then to find an extension to the concepts of *correlation* and *coherence* to accommodate a time-frequency representation, for those noise features that exhibit both peculiar time and frequency characteristics.

This framework can be provided by the *wavelet transform* [291, 369], which naturally embodies the concept of a time-frequency representation, and, most importantly, is also *multi resolution*. This aspect is relevant for the fact that low frequency signals require long duration windows to be estimated with a sufficient frequency precision, while high frequency signals are better isolated in time using short duration windows. This is known as *time-frequency localization problem*, and is a consequence of the *Gabor uncertainty principle* [370] (analogue of the Quantum Mechanic's *Heisenberg principle*). The standard "short-time" Fourier transform, using windowed data, is *inaccurate* to this purpose, imposing at all frequencies a resolution

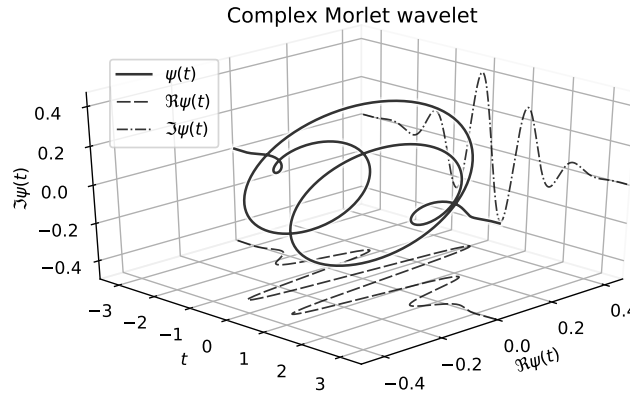


Figure 3.19: Complex Morlet(-Gabor) wavelet for $\omega = 2\pi$, together with the projections of its real and imaginary parts, represented by the dashed and dot-dashed lines respectively.

$1/T$, and at all times a resolution $1/f_S$ [371]. Refer to figure 3.18 for a depiction of the signals representations encountered so far and their resolution in the time and/or frequency domain.

The implementation of a wavelet representation is well established in GW research. Glitches, which are transient noise features with compact support and characteristic patterns in a time-frequency map [23, 258], are naturally represented in this way, as well as transient GW signals. Pinto *et al.* [372] have shown that many classes of glitches can be reconstructed by a small sum of *continuous* Morlet(-Gabor) wavelet components [373], interpretable as “glitch atoms”. To this purpose, we have already mentioned the Omega pipeline [323] and Omicron [LVC30], which use this kind of representation. Also the BayesWave algorithm [195] uses the same modellistion to “fit” glitches and GW signals. There are also examples of studies by means of *discrete wavelets*, which constitute an *orthogonal basis* on which to expand and reconstruct the signals. The aforementioned cWB [21] and the stationarity test described in section 3.1.3 [2, LVC29] exploit this strategy.

In this section we will formalise the definition of wavelet transform and we will introduce the concepts of wavelet coherence [374] and wavelet correlation [375, 376]. The short-time version of the phase delay introduced in section 2.3.3 will be described in the last section. This constitutes a new attempt to make inference on the “instantaneous causality” between signals, whose concept will be clarified in section 3.7.

3.6.1 The continuous wavelet transform

In this treatment we will consider *continuous wavelet transforms* only, but many considerations apply to discrete orthogonal ones as well. Loosely speaking, a *wavelet* is a *finite energy* function representing a “brief oscillation”, with a limited time support and some kind of periodicity assimilable to a “frequency” [377]. More formally, there are some conditions for a generic function $\psi(\eta)$ of the dimensionless time parameter η to be considered an “admissible wavelet” [319]. Basically, this should be *square integrable* ($\psi(\eta) \in L^2(\mathbb{R})$), in order to admit some kind of Fourier transform, and satisfy an *energy conservation* condition analogue to a generalised version of the *Parseval’s Theorem*, which ensures that no information gets

lost in the transform corresponding to the convolution of the wavelet with the signal. From the former, we can conveniently choose the function to be normalised to have unit energy, $\int_{-\infty}^{+\infty} |\psi(\eta)|^2 d\eta = 1$; this, in some sense, expresses the fact that the function $\psi(\eta)$ should be “localised”, with respect to the metric given by the L^2 -norm. The second requirement can be imposed by checking the ability of $\psi(\eta)$ to “resolve the identity”, which can be shown to be equivalent to demanding that its Fourier transform⁴⁴ satisfies $\int_{-\infty}^{+\infty} |\tilde{\psi}(\omega)|^2 / |\omega| d\omega < +\infty$, which reduces to $\tilde{\psi}(0) = 0$ for sufficiently fast decay with ω [378, 379]. In the time domain, the latter is equivalent to $\int_{-\infty}^{+\infty} \psi(\eta) d\eta = 0$, which actually means that $\psi(\eta)$ should be a “rapidly oscillating” function, in order for its average to be equal to zero.

The “archetypical” example, and a natural choice after the discussion at the beginning of this section, is the *complex Morlet(-Gabor) wavelet* [373]:⁴⁵

$$\psi(\eta) = \pi^{-1/4} e^{i\omega_0\eta} e^{-\eta^2/2} \quad (3.27)$$

representing a monochromatic wave of (non-dimensional) *central angular frequency* ω_0 , windowed by a Gaussian function that accomplishes its localisation in time. Since ω_0 is a free parameter, the previous equation represents in fact a family of “candidate wavelets”. Its value determines the number of oscillations in a *wave packet*; the higher the number of oscillations the better the frequency resolution, in spite of the temporal one [378]. Also, this factor should be chosen in such a way to make $\psi(\eta)$ satisfy the admissibility condition of vanishing average. A common choice consists in making the second peak of the function equal to half of the first: $\omega_0 = \pi\sqrt{2\ln 2} \simeq 5.336$ [382].⁴⁶ Historically, a very popular choice is $\omega_0 = 6$, which makes it even closer to the admissibility condition $\tilde{\psi}(0) = 0$ [291, 319]. In our analysis, we have found particularly convenient, for its interpretability, to choose $\omega_0 = 2\pi$, which slightly emphasises the “frequency resolution” with respect to the former options; the reason for that will be clarified in a moment, and the admissibility condition is not spoiled either. In general, this is a tunable parameter to choose on the base of the characteristic duration and bandwidth of the signal to analyse. Some criteria about how to choose it will be discussed after its interpretation will be clarified. Namely, lower values of ω_0 ($\sim 2\pi$) are more suitable for shorter, impulsive signals, while higher ones are better for longer lasting signals and better frequency resolutions.

Moreover, the choice of a complex wavelet function is useful in order to return information about both amplitude and phase and is better adapted, with respect to real one, for capturing oscillatory behaviours [291]. We will exploit this aspect in section 3.6.3.

In figure 3.19 we reported the 3D plot of the complex Morlet(-Gabor) wavelet, for $\omega_0 = 2\pi$, together with the projections of its real and imaginary parts, represented by the dashed and dot-dashed lines respectively. As commented before, higher values of ω_0 correspond to more oscillations of this function (the real and imaginary parts of it) and a better frequency resolution, in spite of a worst time localisation.

⁴⁴The existence of the Fourier transform $\tilde{\psi}(\omega)$ is ensured by the first condition: $\psi(\eta) \in L^2(\mathbb{R})$.

⁴⁵The normalisation choice by means of the factor $\pi^{-1/4}$ for the Morlet wavelet is the popular one adopted in [291, 380, 381], which makes its squared norm simply equal to $\exp(-\omega_0^2)$.

⁴⁶Given the Fourier transform of (3.27), $\tilde{\psi}(\omega) = \pi^{-1/4} \exp(-(\omega - \omega_0)^2/2)$, this implies $\tilde{\psi}(0) \simeq 7 \times 10^{-7} \approx 0$, which is close enough to zero-mean for most practical purposes. In general, in literature are considered admissible Morlet wavelet with parameter choice $\omega_0 > 5$.

The *wavelet transform* (or coefficient) of a signal $x(t)$ is given by its convolution with a *scaled* and time-shifted version of the prototype “mother” (or *analysing*) wavelet, $\psi(\eta, s)$, acting as a time-scale projection kernel:

$$W(t, s) = \frac{1}{\sqrt{|s|}} \int_{-\infty}^{+\infty} \psi^* \left(\frac{\eta - t}{s} \right) x(\eta) d\eta \quad (3.28a)$$

$$:= \int_{-\infty}^{+\infty} \psi^*(\eta - t, s) x(\eta) d\eta \quad (3.28b)$$

The *scale parameter* s represents here the analogue to a frequency f .⁴⁷ The normalization by $1/\sqrt{|s|}$ maintains the energy invariance of the “daughter wavelets” as a function of scale.

For the Morlet wavelet (3.27):

$$\begin{aligned} \psi(\eta - t, s) &:= \pi^{-1/4} \frac{1}{\sqrt{|s|}} \psi \left(\frac{t - \eta}{s} \right) \\ &= \pi^{-1/4} \frac{1}{\sqrt{|s|}} \exp \left(-\frac{(t - \eta)^2}{2s^2} \right) \exp (i\omega_0(t - \eta)/s) \end{aligned} \quad (3.29)$$

where, choosing $\omega_0 = 2\pi$, makes evident the relation between scales and frequencies, $s \equiv 1/f$:

$$= \pi^{-1/4} \sqrt{f} \exp \left(-\frac{(t - \eta)^2}{2} f^2 \right) \exp (2\pi i(t - \eta)f) \quad (3.30)$$

representing in fact, the Gaussian *envelope* about t of a monochromatic plane wave of frequency f ; the former localise the wavelet in time and the latter determines its frequency. When the mother wavelet can be interpreted as a windowed sinusoid, as in this case, the wavelet transform can be interpreted as a constant- Q Fourier transform [323, 383].

For a time series sampled at discrete times $n = 0, \dots, N - 1$, with a frequency f_s , equation (3.28) can be written as [291]:

$$W_n(s) = \sum_{n'=0}^{N-1} x_{n'} \psi^* \left(\frac{n' - n}{s f_s} \right). \quad (3.31)$$

By varying s and translating along time n , one can construct a form of time-scale (or frequency) representation, similar to the spectrogram, which is called the *scalogram* of the signal. In figure 3.20 the scalogram of a typical glitch is reported (bottom colour-map) together with its whitened time series (top plot). Details on the algorithm that has created the map, together with further information on the quantities shown therein, will be discussed in the next subsections.

⁴⁷Refer to the next section and to note 50 for the general relation between frequency and scale.

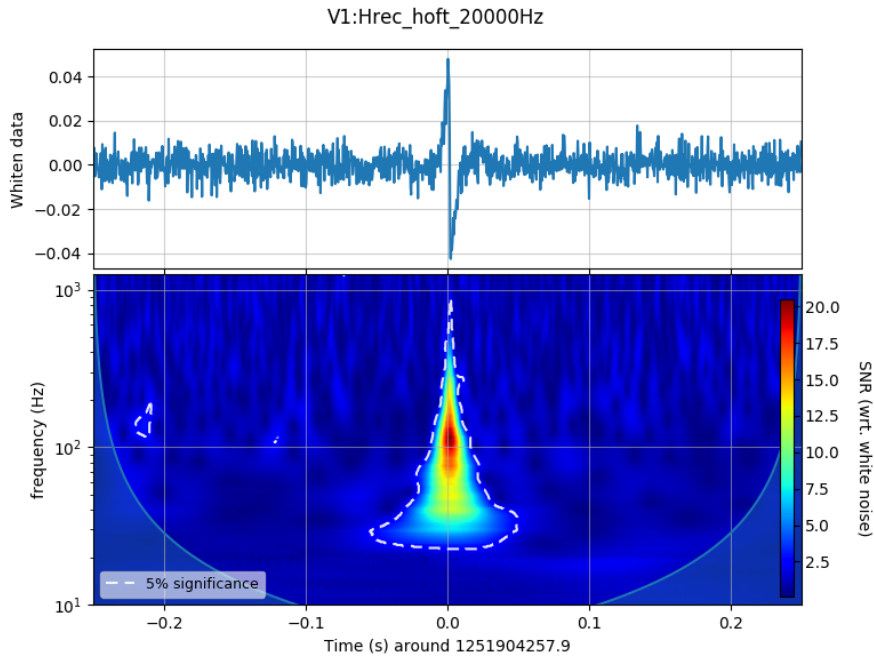


Figure 3.20: Example of “glitch” in the Virgo strain channel. In the top plot the *whitened* time series is reported, while the bottom map represents the time-frequency wavelet transform of the signal, reported in units of SNR with respect to the background (white) noise. The dashed contours represent the 5% significance with respect to white noise hypothesis.

Algorithm implementation

The map and the plot represented in figure 3.20 has been created with the Detector Characterisation tool WATERL00 (WAVELET coherence anaLysis t00l) [LVC45, LVC46].⁴⁸

For the study of glitches, with durations $\lesssim 1$ sec, chunks of $50 \div 200$ seconds of data around the time of the maximum excess of power were loaded, and used for whitening the corresponding time series. This process has already been described in equation (3.5); the ASD used was computed with the mean-median method described in section 2.1.6, with ffts of the size of the required window around the glitch, usually $0.5 \div 4$ sec, depending on its duration. This method helped to considerably speed up the computation of wavelet transforms, avoiding redundant ffts. Indeed, instead of performing the N convolutions of equation (3.31),⁴⁹ from the *convolution theorem* for the Fourier transform, one can compute them as multiplications from the frequency domain [384]. Indeed, for computing the ASD for the former whitening step, we already have computed the DFT of the data in the window

⁴⁸The author acknowledge the github project `aprovecharLab` by C. Brown for having constituted the starting point for the development of the WATERL00 tool, and L. Rei from INFN section of Genova for his help in adapting it to run on Virgo farm machines.

⁴⁹The choice of making one convolution per each time point of the series x_n is natural but arbitrary. One could have legitimately chosen to implement some sort of “downsampling”, skipping an arbitrary number points at each step [291].

of interest around the central time of the glitch:

$$\tilde{x}_k = \frac{1}{f_S} \sum_{n=0}^{N-1} x_n e^{-i\omega_k n/N}, \quad (3.32)$$

where $k = 0, 1, \dots, N-1$ is the frequency index. Then, from the convolution theorem, equation (3.31) can be rewritten as:

$$W_n(s) = \frac{f_S}{N} \sum_{k=0}^{N-1} \tilde{x}_k \tilde{\psi}^*(s\omega_k) e^{i\omega_k n} \quad (3.33)$$

where the *normalised angular frequency* (in units of radians per sample) is defined as:

$$\omega_k = \begin{cases} \frac{2\pi k}{N} & \text{for } k \leq \frac{N}{2} \\ -\frac{2\pi k}{N} & \text{for } k > \frac{N}{2} \end{cases} \quad (3.34)$$

This “trick”, allows to compute, for each s , the wavelet transform simultaneously in parallel for all the times n .

The scales s involved in the previous computation can be *arbitrarily* chosen. The WATERL00 algorithm allows to choose them linearly spaced, useful for studying non-stationarities spanning only one order of magnitude of frequencies, or as fractional power of two (*octaves*), as typical from Acoustics applications:

$$N_{\text{octaves}} = \lfloor \log_2 N \rfloor, \quad s_j := 2^{j/N_{\text{notes}}} \quad j = 1, \dots, N_{\text{octaves}}. \quad (3.35)$$

The parameter N_{notes} is the one giving the frequency granularity, in terms of “notes per octave”: $f_j \approx f_S/s_j$. For optimising the computation of (3.33), the lower frequency region, say below 10 Hz, is usually discarded since not of interest for glitch studies and not even reconstructed in the strain channel. This choice is represented in panel (d) of figure 3.18.

Notice that the previous relation holds for a choice $\omega_0 = 2\pi$ for the Morlet mother wavelet (3.27). With other choices of this parameter, or for other mother wavelets, the relation between scales and frequencies should have been investigated following, for example, the method of Meyers *et al.* [321]; this consists into computing the wavelet power spectrum for a cosine wave signal of known frequency, and attributing to this frequency the scale corresponding to the maximum value of $|W_n(s)|^2$.⁵⁰

Wavelet power spectrum and significance level

In the top plot of figure 3.20, we reported the whitened time series of a glitch in Virgo strain channel, clearly visible in correspondence of the central time. The units are arbitrary or, to be more precise, the typical normalisation of $|\tilde{x}_k|$ with unit variance, as a consequence of the whitening, is implemented. The *wavelet power spectrum* $|W_n(s)|^2$ is reported in the main colour-map. Notice that as the Morlet wavelet is complex, also $W_n(s)$ is so; we will make use of the phase component of this transform in section 3.6.3. Notice that with the energy

⁵⁰For the Morlet wavelet this relation is given by $f^{-1} = 4\pi s \cdot (\omega_0 + \sqrt{2 + \omega_0^2})^{-1}$ [321]. Therefore, with $\omega_0 = 2\pi$, $f^{-1} \simeq 0.99 \cdot s$, while with $\omega_0 = 6$, $f^{-1} \simeq 1.03 \cdot s$.

normalisation discussed in the previous section, if $|\tilde{x}_k|^2$ has expectation value σ^2 (equals to one with our whitening convention), then, for a white noise process:

$$E[|W_n(s)|^2] = N\sigma^2, \quad \text{for all } s, n.$$

For convenience, it was given this normalisation, “relative to white noise”, also to the wavelet power spectrum in figure 3.20.

To assess the *significance* of this spectrum, it is assumed that, in the absence of any non-stationary noise, namely glitches, $|W_n(s)|^2$ should be randomly distributed around its expected average for a white noise process (null hypothesis). If x_n is a white Gaussian process, we have already commented in section 2.1.6 that $|\tilde{x}_k|^2$ should be distributed as Γ distribution with shape parameter $k = 1$ and scale σ^2 , the latter also equal to its mean and standard deviation (refer to appendix B.2), which is equivalent to a χ^2 distribution with two degrees of freedom if $\sigma = 1$. Then, according to the null hypothesis:

$$\frac{|W_n(s)|^2}{N\sigma^2} \sim \frac{1}{2}\chi_2^2$$

where the $1/2$ in front of χ_2^2 (necessary only for complex wavelets) scales it to the “reduced” χ^2 distribution [291]. The previous considerations are analogous to those discussed for the non-stationarity test based on orthogonal wavelets in section 3.1.3. The validity of the previous assumptions has been checked by means of both “quiet data”, that is, in the absence of glitches, for testing the normalisation of the whitening, and with white Gaussian simulated data. In both cases, the distribution of the (normalised) wavelet power spectra was compatible with a χ_2^2 distribution, according to a Kolmogorov-Smirnov test with significance $\alpha = 1\%$; refer to section 3.1.2 or to [304, 305]. In particular, our results are compatible with those in [385], where the authors describe a similar analysis. Of particular relevance was also to repeat the computation in the case of a *time-scale* averaged spectrum, as we will describe in the next section. Moreover, the same sanity check was implemented for the wavelet cross-spectrum, described in the section 3.6.2.

The dashed white contour in figure 3.20 represents the 95% confidence level for the null hypothesis of white Gaussian noise, corresponding to a significance level $\alpha = 5\%$, representing the probability of wrongly rejecting the the null hypothesis when it is true. As expected, the entire excess power of this “loud” glitch fall into this region.

Notice also the blue-shaded, “U-shaped” region at the borders of the map. This corresponds to the *cone of influence* for the analysed data segment (the fft around the central time of the glitch) and the chosen kind of wavelet. In this region, edge effects can occur due to the fact that the wavelets are not completely localized in time, and this is where part of the them fall outside the borders of the data segment. A common practice, implemented in WATERL00, is to zero-pad the data up to the next power of two, as already discussed in section 2.1.6, in order to speed up the fft computation and, in this case, to extend the support for the wavelets. Other possible choices, like *cosine damping* are discussed in [321]. The shape of the cone is chosen as the time from the edge at which the power of the wavelets has decreased by a factor e^{-2} ; the previous quantity is called the wavelet *e-folding time* [291]. This gives also a measure of the *decorrelation time* for a single “spike” in a time series; by comparing the width

of a peak in the wavelet power spectrum with this decorrelation time, one can distinguish between a spike in the data (possibly due to random noise) and a harmonic component at the equivalent Fourier frequency. We will exploit this property in the next section.

Averaging in the time-scale plane

Notice that the previous definition for the wavelet power spectrum $|W_n(s)|^2$ is analogous, besides the normalisation f_s/N to a *density*, to that of the *periodogram* PSD estimate $\hat{S}(f) \propto |\tilde{X}(f)|^2$. As commented about the latter in section 2.1.6, and as expressed by the χ^2 distribution, its standard deviation is equal to its expected value. The “background texture” of figure 3.20 (with a different colour palette they would have resembled *flames*, as often they are colloquially referred to) represents these random variations. This has no effects when studying a single glitch, which is “by definition” more intense than the background. However, when in the next sections we will compare the wavelet transforms of two signals, these fluctuations may overlap, producing spurious coincidences.

Similarly to what was done with the Welch’s method in equation (2.55), the idea is to define some kind of averaging procedure, in order to suppress the effects of these random fluctuations. Unfortunately, there is no obvious way to do that since, in this case, we have only a single realisation of a non-stationary process (the glitch); the Ergodic Theorem is clearly not applicable in this context, and the consequent meaning to attribute to an averaging is not trivial.⁵¹

Again, if there is a clear distinction of *scales* (in the general sense of the term) between the features we want to represent and those of the background randomness, some authors [386] suggest the “compromise” of simply computing the local averages of consequent wavelet spectrum in both times and scales. This allows to suppress the random fluctuation of the background at the price of loosing part of the resolution about the glitch. In [291, 387] the authors make some quantitative statements about the natural shapes and sizes one should choose for these averages, in relation to the adopted wavelets. What they suggest is to make a time smoothing with a filter given by the absolute value of the wavelet function at each scale, normalized to have a total weight of unity, which for the Morlet wavelet (3.27) is just a Gaussian function: $\propto e^{-n^2/2s^2}$. The reason for that is related to the *e*-folding time and the cone of influence described before. Similarly, the scale smoothing can be done using a *boxcar filter* of width equals to the *scale-decorrelation length* of the particular wavelet (and parameter ω_0 , in the case of the Morlet) chosen. We will represent the corresponding time-scale averaged wavelet power spectrum by means of double angle brackets: $\langle\langle |W_n(s)|^2 \rangle\rangle$. These filters are the best “compromise” as they provide the minimal amount of smoothing necessary to include two independent points in both the time and scale dimensions. Larger windows, covering more decorrelation lengths, can be chosen if an even smaller estimation variance is required, in spite of the time-scale resolution. The characteristic width for these filters, in both time and scales, for various wavelet choices can be found in [291, Table 2].

⁵¹Notice that here and in the following discussion, when talking about averaging we only refer to the amplitude of the wavelet power spectrum. Indeed, for the same considerations about the inapplicability of the Ergodic Theorem, for no stationary signals it doesn’t have any physical meaning to compute the time-scale averages of the changing phases. These phases will be therefore excluded from all the following considerations about averages.

This strategy has been implemented in the WATERL00 algorithm in the computation of the wavelet cross-spectrum and, most importantly, the wavelet coherence, that will be described in the next section. Indeed, being the latter the ratio of stochastically varying quantities, analogous to (2.67), it is important for these quantities to have a moderate variance, otherwise the value of the resulting variable will be completely spoiled by their variations.

3.6.2 Wavelet cross-spectrum and wavelet coherence

Firstly, let's comment the logic that has led to the development of the water100 tool, and its usage for noise studies. The typical situation is that the cross-correlation analysis described in section 3.5 has highlighted two channels (the target and an auxiliary sensor) having an excess of power at the same time. We want then to gather more information on them studying the time-frequency morphology of these power excesses, inferring in particular if the two can be related. The available tools, like Omega [323] and Omicron [LVC30], and also the one described in the previous section, allow just a "visual" comparison of the two. UPV [326] partially tries to overcome this dividing this comparison in a small number of frequency bands, and computing in there the percentage of the time the two channels "glitch together".

The idea was then to improve upon this, extending to the wavelet representation the previously described concepts of correlation and coherence, with the purpose of providing new insights into the time-scale dependency between signals.⁵²

We have already noticed that the coherence analysis is inapplicable (or at least misleading) if non-stationarities occur during the time scales of estimations of the auto- and cross-spectral content of the signals involved. A similar, but quite more moderate, consideration has been made for the cross-correlation. If the signals are non-stationary, the Ergodic Theorem is not applicable; hence, strictly speaking, we are not really estimating the corresponding distribution moment of the underlying processes, but just a quantity representing the similarity between the records of the two. Wavelet coherence and cross-correlation represent a practical framework to overcome the previous issues, but not without some limitations, as we will discuss later.

Wavelet cross-spectrum and phase angle spectrum

Given two equally sampled time series, x_n and y_n , by means of equation (3.33) we can compute their *complex* (if using a complex wavelet, like the one of Morlet) wavelet transforms: $W_n^X(s)$, $W_n^Y(s)$. We can then define their *wavelet cross-spectrum* as [388]:

$$W_n^{XY}(s) := W_n^X(s) * W_n^Y(s). \quad (3.36)$$

⁵²In fact, this is more or less the same idea implemented by cWB for the search of coherent GW signals between two detectors [21]. The focus here is on the noise studies of the non-stationarities in the detectors strain channel, which motivated some choices different to those implemented in the cWB pipeline.




Figure 3.21: Wavelet cross-spectrum between Virgo Hrec_hoft_20000Hz strain channel and ENV_CEB_MAG_, that is, the environmental sensor corresponding to a magnetometer in the Virgo Central Building. The colour-map represent the amplitude of the spectrum, while the small black arrows its phase. The dashed white line is the 5% significance region for the hypothesis of two uncorrelated white Gaussian processes. The bottom plot shows the instantaneous time delay statistic defined in (3.47), together with its $\pm\sigma$ and $\pm 2\sigma$ regions.

This quantity is complex, analogously to the usual cross-spectrum $S_{xy}(f)$ defined in (2.64c), and can be rewritten as:

$$\begin{aligned} W_n^{XY}(s) &= |W_n^X(s)|e^{-i\theta_n^X(s)} |W_n^Y(s)|e^{+i\theta_n^Y(s)} \\ &= |W_n^{XY}(s)|e^{i(\theta_n^Y(s)-\theta_n^X(s))}. \end{aligned} \quad (3.37)$$

Similarly to what commented in section 3.5, the *wavelet phase angle spectrum* reflects the phase difference by which y_n “leads” x_n at the given scale and time. This aspect is potentially very important for understanding the relation between the two signals. It will be exploited in the next section, while for now we will mainly focus on its absolute value.

An example of wavelet cross-spectrum is reported in the colour-map in figure 3.21. This is referred to a glitch in the Virgo strain channel, already shown in figure 3.20, and a quasi coincident excess of power in an environmental sensor corresponding to a magnetometer in the Virgo Central Building (ENV_CEB_MAG_N). Apparently, a large amplitude of their wavelet

cross-spectra occurs when the two signals have large power at similar scales (frequencies) and around the same times.

The normalisation, as well as the 95% confidence level are obtained assuming the null hypothesis that both the wavelet spectra are realisations of white Gaussian noise processes, and therefore are distributed individually as χ_2^2 variables. The analytical form of the resulting distribution for the (squared) wavelet cross-spectrum, given by the product of two χ^2 distribution, has been obtained in [389, 390] and discussed, in the context of wavelet analysis, in [291, 391]. For convenience and for checking the various normalisations, in particular in the presence of time-scale averaging, as described in the previous section, we have proceeded to double check the consistency with the previous distribution and, most importantly, to derive the numerical values corresponding to its percentiles. The technique used is the same as that described in section 3.3.1 for the Rayleigh Gaussianity test [LVC24].

Notice that for the cross-spectrum in figure 3.21, a single correlation length averaging has been applied, to reduce the estimation variance, in both times and scales, as described in the previous section. Nonetheless, notice some spurious, although significant, fluctuation of its value, especially at high frequency.

The information about the wavelet phase spectrum,

$$\theta_n^{XY}(s) = (\theta_n^Y(s) - \theta_n^X(s)) := \tan^{-1} \left(\frac{\Im(W_n^{XY}(s))}{\Re(W_n^{XY}(s))} \right), \quad (3.38)$$

computed before any averaging procedure was applied to $W_n^{XY}(s)$, is represented in the same map by the small black arrows, for only those regions with significance smaller than 10%.

Wavelet coherence

From the previous wavelet cross-spectrum is not clear if its high values correspond to coincident high values in both the signals or just one very intense component in only one of them and some random background noise in the other. To distinguish these two cases, and quantify only the degree of time-scale similarity between the signals, we can export the concepts of coherence and correlation function already discussed. The first naive attempt to define a *wavelet coherence* was [388]:

$$\begin{aligned} C_n^{XY}(s) &= \frac{W_n^{XY}(s)}{\sqrt{|W_n^X(s)|^2 \cdot |W_n^Y(s)|^2}} = \frac{|W_n^{XY}(s)|}{\sqrt{|W_n^X(s)|^2 \cdot |W_n^Y(s)|^2}} e^{i\theta_n^{XY}(s)} \\ &= \frac{\sqrt{(W_n^{X*}(s) W_n^Y(s)) (W_n^X(s) W_n^{Y*}(s))}}{\sqrt{|W_n^X(s)|^2 \cdot |W_n^Y(s)|^2}} e^{i\theta_n^{XY}(s)} \\ &= \frac{\sqrt{|W_n^X(s)|^2 \cdot |W_n^Y(s)|^2}}{\sqrt{|W_n^X(s)|^2 \cdot |W_n^Y(s)|^2}} e^{i\theta_n^{XY}(s)} = 1 \cdot e^{i\theta_n^{XY}(s)}. \end{aligned} \quad (3.39)$$

This is a consequence of having made use of just one realisation of the process, and doesn't give any additional information with respect to the already introduced phase angle spectrum.⁵³ Many authors have tried to overcome this *impasse*, modifying the previous definition

⁵³The same issue would have occurred with a periodogram estimation of the usual coherence function.




Figure 3.22: Time-scale averaged wavelet coherence, as defined in (3.40), for the wavelet cross-spectrum represented in figure 3.21. Only the regions with 5% of significance of the latter are shown.

or recurring to related quantities [376, 392, 393], although some of them too distant from the original spirit of the coherence analysis and difficult to manipulate to make analytical predictions. The approach that we preferred was that of [381, 387], which redefined (3.39) in terms of time-scaled averaged quantities discussed in the previous section, analogously to how the standard coherence function is estimated in practice:

$$C_n^{XY}(s) = \frac{|\langle\langle s^{-1}W_n^{XY}(s) \rangle\rangle|}{\sqrt{\langle\langle s^{-1}|W_n^X(s)|^2 \rangle\rangle \cdot \langle\langle s^{-1}|W_n^Y(s)|^2 \rangle\rangle}} e^{i\theta_n^{XY}(s)}. \quad (3.40)$$

The factor s^{-1} (present in the definition in [387] but not in [381]) is used to convert nominator and denominator to units of energy density, in order to reinforce the similarity with (2.67). Of course this choice was irrelevant for (3.39). The authors also notice that, since the wavelet transform, with the normalisation described in section 3.6.1, conserves variance [291], the wavelet coherence is an accurate representation of the (normalized) covariance between the two time series [387]. Its normalisation is the one we expect: $0 \leq |C_n^{XY}(s)| \leq 1$, where the latter condition is true only for identical, and coincident, signal. The latter aspect represent a difference with respect to the usual coherence for stationary signal; this is a consequence of the time-scale localisation of implemented by the wavelet transform, as we will discuss in section 3.6.3.

In figure 3.22 we have reported the (magnitude squared) wavelet coherence, corresponding to the same signals of figure 3.21, in only those regions where the significance of $|W_n^{XY}(s)|$, with respect to the null hypothesis of two white Gaussian noise processes, was smaller than 5% (dashed regions of figure 3.21). This choice was made in order to focus only on those

regions that the previous wavelet cross-spectrum analysis found relevant, and not being distracted by other random fluctuations that, recalling the considerations at the beginning of this section, were likely to produce high spurious fluctuation of this quantity.

Two important things emerge from this figure. First of all, the central region, which appeared to possess large values of the wavelet cross-spectrum, has instead a not that high value of it if compared to the individual ones of the two wavelet power-spectra. This corresponds indeed only to a wavelet coherence $\lesssim 60\%$, which is only moderately relevant; that is, maybe the glitch in h_{rec} is not witnessed by that magnetometer. Moreover, the other small “drop-like” regions, mostly at high frequency, which exhibited only moderate values of the wavelet cross-spectrum in figure 3.21, are now indicating a very high value of the coherence. This is clearly a spurious effect, corresponding to coincident small variations of the background amplitude, where the fact that they were coincident implies a non-negligible value of the numerator of equation (3.40), while their small amplitudes boost it from the denominator.

A clearer picture of the wavelet coherence could have come from increasing the lengths of the averages used for computing (3.40), at the price of a worst time-frequency resolution. This doesn’t constitute a solution to this issue though, which is rooted in the fact of estimating statistical quantities, related to non-stationary processes, from just one realisation of them.

Wavelet cross-correlation

It is worth mentioning another related quantity to describe the time-frequency similarity between two signals. This is the *wavelet cross-correlation*, firstly introduced in [394]. Analogously to 3.9, or more specifically to the *biased* cross-correlation estimator (2.48), it can be defined as:⁵⁴

$$R_l^{XY}(s) := \frac{1}{N} \sum_{n=1}^{N-|l|} W_n^{X*}(s) W_{n+l}^Y(s). \quad (3.41)$$

This quantity admits an interesting interpretation. Rewriting both of the wavelet transforms in the form of equation (3.33), and applying the convolution theorem, we obtain that the previous quantity can be rewritten as:

$$R_l^{XY}(s) = \left(\frac{f_S}{N}\right)^2 \cdot \sum_{k=0}^{N-1} \tilde{x}_k^* \tilde{y}_k |\tilde{\psi}(s\omega_k)|^2 e^{i\omega_k l}$$

and, passing a factor $f_S/N = 1/T$ inside summation, we can interpret the term $\tilde{x}_k^* \tilde{y}_k / T$ as the periodogram estimate of the cross-power spectral density:

$$= \frac{f_S}{N} \sum_{k=0}^{N-1} \hat{S}_k^{XY} |\tilde{\psi}(s\omega_k)|^2 e^{i\omega_k l}. \quad (3.42)$$

Compare it with the periodogram definition in (2.49). The product of this cross-spectrum estimate and the kernel constituted by the squared modulus of the wavelet projects it to a

⁵⁴Other analogous definitions for a wavelet cross-correlation have been discussed in [375, 376, 392]. In [395] for example, the authors define *wavelet cross-correlation* a quantity identical to the coherence we defined in (3.40) except for the presence at the numerator of the real part of $W_n^{XY}(s)$ only. This clearly doesn’t give any additional information with respect to what already introduced here.

specific scale component, providing a term that can be interpreted as a time-scale “localised” cross-power spectral density. This is similar for many aspects to the windowing operation in the Welch’s method in (2.55) (without averaging though). Equation (3.41) is therefore the (inverse) DFT of the previous quantity (refer to the normalisation convention described in the Notation section).

Being quite redundant with respect to the notions already encountered, we have not directly exploited this concept in our analyses. More specifically, we are interested in an independent quantification of the relations between the two signals from their relative wavelet phase spectrum, as we will discuss in a moment in section 3.6.3. This information is summed over in the (3.42) and therefore no more available for this purpose.

3.6.3 Instantaneous time delay through wavelet phase angle spectrum

So far, we have never explicitly made use of the phase information provided by the cross-spectra. In section 2.3.3 we have briefly shown that the phase angle $\theta_{xy}(f)$, defined in (2.73c), can be used to measure the delay between two (stationary) signals, which are one the retarded copy of the other, plus some additive uncorrelated noise. However, it is unlikely to discover, in practical situations, two such signals and them to be of any help for Detector Characterisation purposes.

The phase angle of the wavelet cross-spectrum can provide a natural time-scale generalisation of the previous concept for non-stationary signals. We will proceed then to verify the previous condition for two such signals, in the framework provided by the wavelets, and we will introduce then a new measure of the *causal* relationship between them.

Phase angle spectrum between retarded signals

For simplicity, let’s consider impulsive signals, like the glitch represented in figure 3.20, and in particular let’s start for simplicity considering two *unit pulse* signals, $x_n = \delta_{n0}$ (equals 1 when $n = 0$) and $y_n = \delta_{nm}$ (equals 1 after m time units), with the former “leading” (for $m > 0$) the latter by m samples: $y_n = \delta_{n-m0} = x_{n-m}$. Making use of the definition (3.31), we can compute the two wavelet transforms:

$$W_n^X(s) = \sum_{n'=0}^{N-1} \delta_{n'0} \psi^* \left(\frac{n' - n}{s f_S} \right) = \psi^* \left(-\frac{n}{s f_S} \right), \quad (3.43a)$$

$$W_n^Y(s) = \sum_{n'=0}^{N-1} \delta_{n'm} \psi^* \left(\frac{n' - n}{s f_S} \right) = \psi^* \left(\frac{m - n}{s f_S} \right). \quad (3.43b)$$

The wavelet cross-spectrum of the two pulses consists in a comparison between the wavelet and a complex conjugate, time shifted version of it. Its form is therefore dependent on the particular wavelet choice. For the Morlet one (3.27), we have:

$$W_n^{XY}(s) := W_n^{X*}(s) W_n^Y(s) = \frac{1}{(\pi^{1/4} \sqrt{s f_S})^2} \left(e^{-\frac{n^2}{2s^2 f_S^2}} e^{-2\pi i \frac{-n}{s f_S}} \right)^* \left(e^{-\frac{(n-m)^2}{2s^2 f_S^2}} e^{-2\pi i \frac{m-n}{s f_S}} \right)$$

$$= \left(\frac{1}{\sqrt{\pi s} f_s} e^{-\frac{(n-m)^2 + n^2}{2s^2 f_s^2}} \right) e^{-2\pi i \frac{m}{s f_s}}. \quad (3.44)$$

The exponential within parentheses is the product of two Gaussians, centred in 0 and in m respectively, which represents the loss of coherence between the two pulses; this is due to the already described decorrelation length for the Morlet wavelet, and it constitutes a difference with respect to the previous coherence analysis of stationary signals. What interests us is the last phase term, which, *at all scales*, is proportional to the delay between the two signals:

$$\theta_n^{XY}(s) = -2\pi m/s f_s. \quad (3.45)$$

This means that the arrows in figures 3.21 and 3.22 should be “curled” in such a way to rotate their angle proportionally to the frequency if a similar delay relation holds for the two signals. The same result holds also for other choices of complex wavelet with a phase linear with time, although the Morlet one is arguably the most common of this kind.

A similar condition on the wavelet cross-spectrum phase is expected to hold approximately also for other signal shapes, thanks indeed to the decorrelation term within parentheses in equation (3.44). We have tested this with simulated glitches, as it will be discussed in the next subsection.

Instantaneous time delay statistic

We want to exploit the previous idea that the phase angle spectrum is a measure of the time delay between the signals: $\tau_0 := m/f_s$. This represents the condition for one signal to lead the other, and therefore it can be interpreted as a *necessary* condition for one being “a possible cause” of the other. We will come back to the concept of causality in section 3.7. We want to define then an estimator for this delay but, again, we have to face the problem that our signals are single realisations of the corresponding non-stationary stochastic processes. To make our estimator able to distinguish this expected value from random noise fluctuations, it is important to define a sort of averaging procedure. Of course, we can’t rely on time averages and the standard recipe provided by the Ergodic Theorem exploited in section 2.1.6. However, for the particular form of (3.45), proportional to the frequency (inversely to the scale), we can think of defining this estimator averaging the difference

$$\theta_n^{XY}(1/f + \delta f) - \theta_n^{XY}(1/f) \approx -2\pi \tau_0 \delta f \quad (3.46)$$

over all the frequencies. Notice that δf is not constant in general, as for example with the choice in (3.35); to keep track of this, we restore the discrete label k to the frequencies.

Putting together the previous ideas, we can define the *instantaneous time delay* statistic:

$$\Theta_n^{XY} := \frac{1}{2\pi} \sum_k \frac{\theta_n^{XY}(1/f_k + \delta f_k) - \theta_n^{XY}(1/f_k)}{\delta f_k} \quad (3.47)$$

which, at each time n , is representative of the “instantaneous delay” between the signals. Its

sign can tell us which one is leading the other, and the amplitude by which extent. For example, in the previous case the first signal x_n is leading the retarded copy y_n by a time τ_0 ; this statistic is then identical (in expectation value) to $-\tau_0$ at all times, even if we will not be able to resolve them far from 0 and τ_0 for the Gaussian decorrelation factor in equation (3.44).⁵⁵

The value of this statistic, together with its $\pm\sigma$ and $\pm 2\sigma$ regions, is reported in the bottom plot of figure 3.21. To avoid spurious contributions from regions we know in advance from the modulus of $W_n^{XY}(1/f)$ to be compatible with uncorrelated white Gaussian noise processes, equation (3.47) is evaluated only at those frequencies where the significance of the null hypothesis is smaller than 5% for the wavelet cross-spectrum value. The contributions from the other regions should average to zero, with a certain variance, as expected for two Gaussian noise processes with random phases. Notice also that with this choice, small regions of averaged frequencies tend to have a higher variance than bigger regions, if the latter possess the coherent structure described in equation (3.45). This aspect is visible in the bottom plot of figure 3.21 furthest from the central time of the glitch.

Focussing to the centre of that figure, at about the central time of the glitch, we notice a small interval where the previous instantaneous time delay statistic is different from zero at least by two σ 's, and, in particular, it is positive. This suggests that the first signal, the strain `Hrec_hof t_20000Hz`, is led by the Central Building magnetometer `ENV_CEB_MAG_N`. We will come back to comment this specific application of the method described in this section momentarily.

It is important to notice that the previous conclusion has been obtained independently from that on the magnitude squared coherence of the two signals, exploiting all the information contained in its complex value. Making use of both these information at the same time gives a deeper insight into the relationships between two signals, and can be of great help for the search of the causes of the non-stationary noise. This constitute a second advantage with respect to previous methods based on wavelets or Q -transforms of one channel at a time.

Tests on simulated glitches

In order to test some of the techniques described in the previous sections, including normalisations, significance tests, and, most importantly, the new method provided by the instantaneous time delay statistic, we proceeded to verify them on simulated glitches. Several shapes have been tested, from sine Gaussians, which are represented by bivariate Gaussian envelopes in the scalogram with the Morlet wavelet, to deltas and broader structures.⁵⁶

As an example, a choice that resembles that in figure 3.20 is provided by the function:

$$g(t) = A \cdot \text{sign}(t_0 - t) e^{-|t-t_0|/d} \quad (3.48)$$

⁵⁵Notice here the reason behind defining the statistic with a negative sign if the first signal *leads* the other. This is motivated by the fact that in noise studies we are most likely concerned with the opposite situation, that is to test whether the first "target" channel, say the strain, is *led* by some other, auxiliary channel.

⁵⁶Notice that for sine Gaussian glitches, which basically contain only a single frequency contribution, besides those reconstructed by means of the wavelet transform, the previous technique of instantaneous time delay statistic based on frequency averages is not applicable.




Figure 3.23: Wavelet cross-spectrum for the glitch model (3.48), for an individual SNR of 10 and a time delay $\tau_0 = 1$ ms. The normalisation is given by the expected cross-spectrum for two uncorrelated white noise processes; the dashed white contour represent the 5% significance region for this (null) hypothesis. The arrows represent the phase angle of the complex cross-spectrum.

where A is the amplitude factor. The discontinuity provided by the “sign” function makes it broadband, while the characteristic decay time d gives its duration. This is bandpassed, say between 20 and 500 Hz, in order to avoid aliasing and low frequency effects, and superimposed to a white Gaussian noise with zero mean and variance σ^2 . In figure 3.23 we reported the wavelet cross-spectrum of such a glitch with a retarded copy of it, superimposed to a different noise realisations; in this particular example we choose a “moderate” SNR $A/\sigma^2 = 10$ and a very short time delay $\tau_0 = 1$ ms. Focussing on the the phase arrows about the centre of the figure, is evident their “clockwise curly behaviour” with the frequency, characteristic of a negative time delay, according to the interpretation given in the previous section.

In figure 3.24 we reported the value of the phase angle $\theta_n^{XY}(1/f)$, as a function of the frequency, for various n (gray lines) within 10 ms about the centre time of the target glitch. The blue line represents their median, with respect to n , while the red line is the theoretical value for $\tau_0 = 1$ ms (the intercept of this line, corresponding to the random phase at $f = 0$, has been fitted from the data). The correspondence between the median measured phase angle and its theoretical value is remarkable below 250 Hz; after that, the coherence between the two signals becomes negligible, as visible from figure 3.23. Moreover, from frequencies comparable to $1/\tau_0$, equals to 1 kHz in our case, depending also on the spacing of the sampled frequencies, the phases start to “spin” so fast it may become hard to accurately reconstruct their differences (3.46). This in a certain sense limits our method to small delays between signals, but this is not a real issue with the noise coupling we are likely to find in a GW interferometric detector.

The corresponding instantaneous time delay statistic Θ_n^{XY} has been computed from (3.47) and shown in figure 3.25. Besides the two “wings” on the sides, compatible with a null value




Figure 3.24: Wavelet cross-spectrum phase angle about the centre time of the glitch. The light grey lines represent the values of the angle as a function of the frequency in a 10 ms window around the time of the glitch. The blue line is their median average, while the dashed red line is the theoretical curve corresponding to a 1 ms delay. The constant phase (intercept at $f = 0$ Hz) of the previous curve has been fitted from the data.

of the statistic, in the centre region around the time of the glitch (zoomed inset plot), its value is similar to $\tau_0 = 1$ ms with which the two glitches have been generated.

3.6.4 Discussion and applications

The methods described in this section aim at extending the notions of correlation and coherence analysis, as previously described in sections 2.3.3 and 3.5, to a time-frequency representation of the signals, particularly suitable for the study of glitches. The idea to make use of the wavelet transform to this purpose is not new in GW community, as thoroughly exploited by the *Omicron* pipeline [LVC30], although this doesn't implement any measure of correlation between different signals. This is only partially overcome in UPV [326], which also includes a measure of the significance of this on the base of its occurrence for different glitch triggers, even if it can't be consider an actual time-frequency representation, as it comprises only a single time interval (that covering the duration of the glitch) and a very small number of frequency bands (typically $5 \div 9$). A proper representation of such a kind, on the base of which to compare the strain signals from different detectors, is that implemented in cWB [21] for the unmodelled search of astrophysical "burst-like" signals. The method described in section 3.6.2 is an application of similar ideas to the study of glitches in the signals within a single detector.

Differently from the previous methods, the technique developed by us exploits the phase of the complex wavelet cross-spectrum between two signals. This provides additional information, independent on that of the magnitude of this spectrum. Hence, it can give us further insight on the relations between two signals, and in particular on their relative delay. We will come back to discuss about this in a moment.




Figure 3.25: Instantaneous time delay statistic for the simulated glitches reported in figure 3.23. The blue line corresponds to the value of this statistic around the time of occurrence of the glitch in the main channel, while the blue region is its $\pm\sigma$ confidence band. At the centre of the figure, it is visible a small region, due to the correlation factor, where this statistic is compatible (within $\sim\sigma$) with the actual delay of 1 ms of the glitches in the two channels.

To achieve a proper reconstruction of the energy content of a signal, the *Omicron* and *Omega* [323, 396] pipelines make use of multiple wavelets with different Q -factors (equivalent to different choices of ω_0 for the Morlet wavelet (3.27)) to *tile* the time-frequency space.⁵⁷ We didn't implement a similar strategy in our algorithm, where the value of ω_0 should be chosen in advance on the base of the duration and the desired frequency resolution of the signals we want to investigate. Refer to the interpretation of this parameter as discussed at the beginning of section 3.6.1. Our decision is motivated by the application of the instantaneous time delay method of equation (3.47), which is feasible only for a fixed choice of ω_0 . The comparison of phase angles between time-frequency planes with different ω_0 is not of trivial interpretation, hence the extension to multiple Q -factors of the technique developed by us.

On the possibility of a novel wavelet based causality test

Coming back to the interpretation to give to this new statistic, this is meant to test, at each time, whether one signal is leading the other. The frequency averaging makes it robust with respect to the uncertainties we have with only one realisation of the processes, and allows to define "confidence regions" based on its standard deviation. Moreover, its evaluation only on those regions found significant for the wavelet cross-spectrum is an additional aid to get rid of spurious effects due to noise. Tests on simulated glitches, even with moderate SNR and small delays between the signals (usually hardly appreciable by eye), have shown its ability to correctly identify their relation. The Gaussian decorrelation factor in (3.44) is a

⁵⁷They do this in a slightly differently way, with results visually a little bit different. For different Q 's, *Omicron* produces multiple time-frequency histograms, corresponding to the regions where each wavelet is localised, and with values corresponding to the energy content in that wavelet component. Then, it stacks the various tiles of different histograms with on top those with the highest SNR. *Omega* produces similar time-frequency planes, tiled for different Q 's, and then it selects the most significant non-overlapping pixels, which can also be clustered, among all planes using a simple exclusion algorithm.

consequence of the locality of the wavelet representation, and makes the previous statistic to assume the value of that delay in correspondence of the occurrence of the glitches in the two channels. If these are close to each other, as in the example in figure 3.25, this becomes actually a band for the statistic to be correctly measured as the time delay between the two.

Being the temporal delay a necessary condition for one signal to constitute the “cause” of the other, it is tempting to interpret the previous instantaneous time delay, in fact, as a hint for a causality relation. To be precise, if a signal x_n is found to lead a second one y_n , it is natural to exclude the possibility that the latter has caused the former. Of course, there are exceptions against considering this true in general. For example, there could be an intrinsic delay between the actual cause and the time this shows up in a signal. This is the typical case of missing explanatory signals, which should have witnessed the cause for first. This case is not uncommon, and constitutes a limit for this kind of analysis. More details on the concept and interpretations of causality will be provided momentarily in the next section.

The idea to use wavelets to infer causality is not new, although of very recent ($\gtrsim 2013$) development, mostly in the Econometrics context, and in many cases not mature enough nor universally accepted. For example, in [395] the authors have proposed to infer the causality relation between two signals observing the sign of the *wavelet cross-correlation*, defined analogously to the complex wavelet coherence we introduced in (3.40), but with only the real part of $W_n^{XY}(s)$ at the numerator, and therefore assuming value between -1 and $+1$. The authors have then concluded that the variable were to be considered causal (anti-causal) if the sign of the previous function was positive (negative) and its absolute value close to 1. A similar idea was refined in [397, 398], where, very likely inspired by [213, Figure 5.4], the authors introduced a multiplicative *indicator function* that takes the value one if the variables were in phase ($\theta_n^{XY}(s) \in [0, \pi/2] \cup [-\pi/2, 0]$) and the value zero if not. Both this pictures are misleading (or wrong) if applied for example to the case in (3.46), where the phases change sign and quadrant with the frequency f . The method described in [399, 400] is probably one of the best founded and most widely used to-date, judging for example its significant follow up. It basically consists in a spectral Granger-Geweke causality test [401, 402] (to be introduced in the next section), where the local spectral estimations are made with the wavelet power spectrum and cross-spectrum. Its application to the study of non-stationary noise in GW detectors is probably not feasible, in the form it is presented. First of all, it is computationally very demanding, requiring at each time-scale point to computation the spectral factorisation of the spectral density matrix of the two signals [403] to compute their cross-correlation and transfer function matrices. Moreover, the estimations of these spectra are affected by the uncertainties inherent in having at our disposal only one realisation of the processes; time-scale averaging is only a partial solution. The authors overcome this difficulty averaging their statistic on 5000 simulations of the same physical system. In our case, it is unlikely to have this many realisations of the same identical glitch.

Analysis of “mystery glitches” and BNS range drops

The techniques described in this section and the WATERL00 analysis tool have been applied to some specific noise studies in Virgo. In particular, during O3a it was reported of some occurrences of an unknown family of very loud glitches that caused drops in the BNS range of

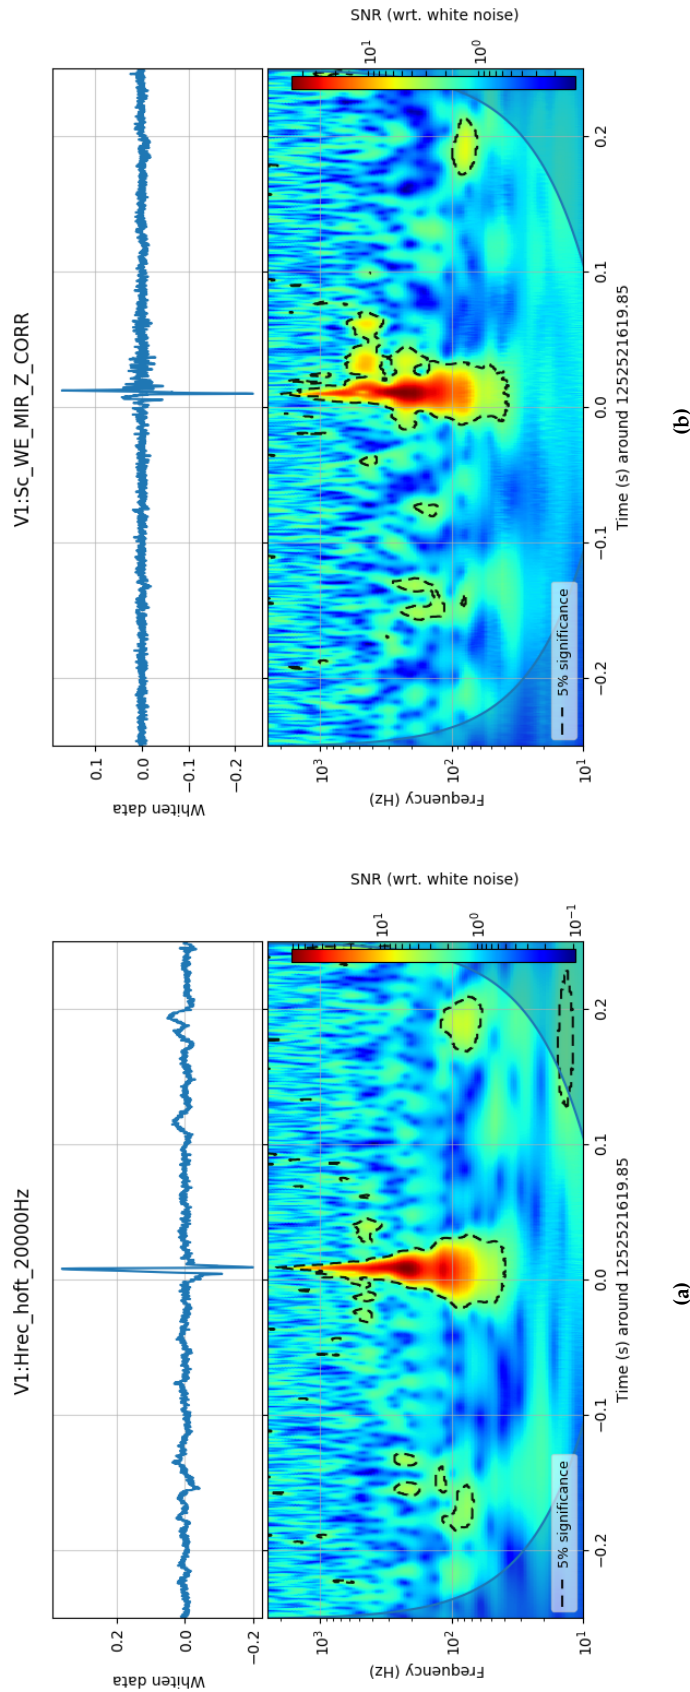


Figure 3.26: Scalograms obtained with WATERLOO of (a) `Hrec_hoft_20000Hz` and (b) `Sc_WE_MIR_Z_CORR` Virgo channels (bottom colour-maps), and corresponding whitened time series (top). In both of them is clearly visible a glitch at the centre of the figures, with approximately a similar shape. Notice that due to the high SNR with respect to white noise, the scalograms have logarithmic colour-map, while those in figures 3.20 and 2.9 were linear. In these figures the contours for the 5% significance level of the null hypothesis of white noise are also reported.

even 20 Mpc [LVC47]. The scalogram of one of them is reported in figure 3.26a where, analogously to 3.20, it is shown its wavelet power-spectrum and whitened time series. For convenience, the colour-map is logarithmic, to easily accommodate its high SNR. These glitches were identified by *Omicron* and they were present, besides in the strain channel and in all those related with it (like LSC_DARM and other photodiodes and actuators used for its reconstruction), in many other auxiliary channels, although with quite diverse morphologies in the time-frequency maps. A visual inspection of all of them was not very helpful, unless for the expert commissioner able to recognise shapes and hidden relations within channels, not evident from the scalograms alone.

UPV [326] helped a lot in discriminating between those channels that manifested glitches in coincidence with the strain channel with high significance. In particular, it pointed to some controls in the *long suspension tower* containing the beam splitter (Sc_BS) and other mirrors.

By means of WATERL00 we obtained the detailed correlation of their time-frequency energy content [LVC45]. In figure 3.26b we have reported the scalogram of a control channel of the *z* axis of the West End mirror (Sc_WE_MIR_Z_CORR). As visible from both the scalograms in 3.26, the glitches in these two channels present a very similar time-frequency morphology. To quantify this similarity, in figure 3.27 we reported their wavelet coherence (only for those regions where the null hypothesis of white noise was less significant than 5%) and instantaneous time delay statistic. From the former we had the confirmation that their coherence was indeed very close to 1 in a wide region of frequencies in correspondence of the glitch. Most noticeably, the time delay statistic revealed a certain region, with significance greater than 2σ , in correspondence of the two glitches where *Hrec was led* by Sc_WE_MIR_Z_CORR with a delay of about 0.2 ms [LVC46]. This information was clearly not available with the previous tools.

We must be careful about the interpretation to give to this result. Two aspects should be taken into account. First of all, this data records correspond to a single realisation of the corresponding stochastic processes. Hence, the value of the time delay statistic, as well as that of the magnitude of the coherence, could be so just by chance. The averaging procedures implemented in both of them help in reducing the effects of noise fluctuations, which are however unavoidable. Second aspect is whether to interpret the observed delay as a real causation relation. This is the most subtle part, as we will thoroughly elaborate in the next section. What seems reasonable to conclude is just that the glitch in the mirror control is unlikely *to be caused* by that in the strain channel, as it occurs before it. However, some caution is also needed in similar statements. In general, from a single realisation it is never safe enough to make conclusions about possible causality relations. This however provides a deeper insight, with respect to the scalograms alone, to people doing noise studies on multiple channel signals. For example, an expert commissioner can exploit this and judge with his/her experience if the possibility of a causation is actually plausible or not. Moreover, in some circumstances, acting on the controls highlighted by this method, one can also directly verify its claims. The latter approach is based on the concept of *physical influence*, where the *controlled* manipulation of the cause changes the effects, as a sufficient condition for claiming causation [32]. An example of this is provided by noise injections, as discussed in section 2.3.2.




Figure 3.27: Wavelet cross-spectrum between Virgo Hrec_hoft_20000Hz strain channel and ENV_CEB_MAG_, that is, the environmental sensor corresponding to a magnetometer in the Virgo Central Building. The colour-map represent the amplitude of the spectrum, while the small black arrows its phase. The dashed white line is the 5% significance region for the hypothesis of two uncorrelated white Gaussian processes. The bottom plot shows the instantaneous time delay statistic defined in (3.47), together with its $\pm\sigma$ and $\pm 2\sigma$ regions.

3.7 On the concept of causality

Except for the instantaneous time delay described in section 3.6.3, in our studies we have mostly focus on the concepts of *correlation*, or its frequency domain analogue, the *coherence*, and the previous wavelet generalisation, as ways to measure the *similarity* between two time series. Only seldom we have tried to extrapolate this, wandering about the possibility that a signal was indeed *responsible* of the observed variation of the other.⁵⁸ Indeed, a more profound insight could come if we are actually able to test whether one signal, representing a certain subsystem or device, can be actually considered the *cause* of the variations of another one, namely the strain. Unfortunately, the mathematical definition of the concept of *causation* is not a simple task in itself, and only in the last fifty years this issue has been formalised in a

⁵⁸Noise injections described in section 2.3.2 are, for many aspects, the only exception to this. In that case, the operator injecting the known noise waveform is, by the definition, the cause, and we are mostly concerned with the measurement of the effects.

consistent way [30–32]. The common habit to associate the latter to the concept of correlation leads to a *logical fallacy*, with notorious paradoxical consequences.

To the purpose of noise studies, we have faced this problem in two different ways, each of them with a specific “operative definition” of causality. To the best of our knowledge, as of October 2019, this is the first attempt in GW literature of a similar application of these techniques, mostly developed in the context of Climatology [33] and Neuroscience [34–36]. We shall present them, and the corresponding studies, anticipating that due to the complexity of the subject and the mechanisms governing the functioning of a GW detector, with many feedback controls and cross-talking parts, no significant new insight has been achieved so far. An extensive testing of the proposed methods is currently in process in order to convince us about the applicability and reliability of the predictions obtained by them. This is currently done by means of the analysis of the hardware noise injections performed for Detector Characterisation purposes in September 2019 [LVC48], and on specific noise issues under control, aided by the experience of the commissioners in charge of these subsystems.

The following discussion has greatly benefited from the collaboration with the Virgo group at the INFN section of Genova, and in particular of Dr. L. Rei.

3.7.1 Granger-Geweke causality

After the discussion in section 3.5.3, the most natural paradigm to attempt some inference about a casual relation between two (or more) stochastic processes $X_1(t)$ and $X_2(t)$ is that provided by *Granger causality* [401] and its frequency domain version due to Geweke [402]. This is probably also the most widely used approach in the literature, although it possesses some significant limitations to its applicability to the studies regarding non-stationary noise in GW detectors.

The measure of causality in the sense of Granger (therefore sometimes referred to as “Granger-” or “G-causality”) relies on the ability of the *driving*, or *causal*, variable to predict future values assumed by the *effect*, or *driven*, variable. This is verified in the context of *linear regression models* (refer to section 3.5.3), assuming the two variables are *model stationary*, in the sense described in section 2.1.3, and representable by Auto Regressive (AR) models.⁵⁹ The inclusion in their respective models of past values of the other variable yields:

$$X_1(t) = \sum_{t'=1}^p A_{11}(t') X_1(t-t') + \sum_{t'=1}^p A_{12}(t') X_2(t-t') + e_1(t) \quad (3.49a)$$

$$X_2(t) = \sum_{t'=1}^p A_{21}(t') X_1(t-t') + \sum_{t'=1}^p A_{22}(t') X_2(t-t') + e_2(t) \quad (3.49b)$$

which can be readily extended to the case of $n > 2$ variables, as shown in [34]. The residuals $e_{1,2}(t)$ are Gaussian distributed under the stationary model hypothesis (which is important for the correct statistical interpretation of the test results). The *model order* p , that is, the maximum number of lagged variables to included in the previous equations, is typically determined by some information based criterion, such as the AIC [364] or the BIC [365] already encountered in the context of linear regression. Then, the G-causality relation between

⁵⁹A sufficient condition for that is therefore that the variables are covariance-stationary, as explained in section 2.1.2.

the variables is inferred from the reduction in the variance of $e_1(t)$ (or $e_2(t)$) due to the inclusion in the corresponding model of the variable X_2 (X_1 , respectively). Similarly to what discussed in the context of the regression, the test is formally made on the cross-coefficients A_{12} (or A_{21}) performing an F -test of the null hypothesis that $A_{12}(t) = 0 \forall t$ (analogously for $A_{21}(t)$). The “strength” of the corresponding G -causality interaction is then typically measured by the logarithm of the corresponding F -test statistic [402]:

$$F_{2 \rightarrow 1} := \log \left(\frac{\text{Var}[e_{1|A_{12}=0}(t)]}{\text{Var}[e_{1|A_{12} \neq 0}(t)]} \right) \quad (3.50)$$

where $e_{1|A_{12}=0}$ are the residuals for the null hypothesis of autoregressive model for $X_1(t)$, while $e_{1|A_{12} \neq 0}$ accounts for the inclusion in the model (3.49a) of the contribution from X_2 . This statistic is larger than zero if the variable X_2 causes, in the Granger sense, X_1 . To assess its significance, we can use the critical values for an F -variable with p and $2p$ degrees of freedom, which should represent the distribution of the argument of the logarithm if the two variables are stationary ($e(t)$ Gaussian). Analogous definition holds for $F_{1 \rightarrow 2}$, and, in general, with the previous definitions we can construct a “net” of *conditional G-interactions* among a set of n signals [34].

Geweke formulation, limits and extensions

As noted by Geweke [402], and already discussed in section 2.1.2, two AR processes, and in particular two stationary ones, admit a spectral representation. Indeed, taking the Fourier(-Stieltjes) transform [210] of equation (3.49), and using the convenient matricial form [35]:

$$\begin{pmatrix} A_{11}(f) & A_{12}(f) \\ A_{21}(f) & A_{22}(f) \end{pmatrix} \begin{pmatrix} X_1(f) \\ X_2(f) \end{pmatrix} = \begin{pmatrix} e_1(f) \\ e_2(f) \end{pmatrix} \quad (3.51)$$

we can define the *transfer matrix* H_{ij} as the inverse of the matrix of the coefficients A_{ij} ,

$$\begin{pmatrix} X_1(f) \\ X_2(f) \end{pmatrix} := \begin{pmatrix} H_{11}(f) & H_{12}(f) \\ H_{21}(f) & H_{22}(f) \end{pmatrix} \begin{pmatrix} e_1(f) \\ e_2(f) \end{pmatrix}.$$

Left multiplying the previous vector by its complex transpose, and averaging on the space of its possible realisations (angle brackets “ $\langle \dots \rangle$ ”), we can define the PSD matrix:

$$S_{ij}(f) := \langle X_i^*(f) X_j(f) \rangle = \langle H_{ik}^*(f) \Sigma_{k\ell} H_{\ell j}(f) \rangle.$$

This definition is analogous to the one already introduced in equation (2.17) of the PSD of a stochastic process as the Fourier(-Stieltjes) transform of its (auto-)correlation function (2.11). Likewise, $\Sigma_{k\ell}$ is the transform of the cross-correlation matrix of the residuals: $\langle e_i^*(f) e_j(f) \rangle$. In the last term, the repeated indices summation convention, over the various processes, is understood. The spectral G -causality statistic is then defined, for a generic number of

signals, as:

$$I_{j \rightarrow i}(f) := -\log \left(1 - \frac{\left(\Sigma_{jj} - \frac{\Sigma_{ij}^2}{\Sigma_{ii}} \right) |H_{ij}(f)|^2}{S_{ii}(f)} \right). \quad (3.52)$$

The advantage of this formulation, with respect to (3.50), is that now the Granger causality test has become manifestly frequency dependent, as appropriate for the non-trivial spectral structure of GW strain data. One of the causality tests [399, 400] that we have mentioned at the end of section 3.6.3 relies indeed on the spectral formulation by Geweke, and it attempts to extend it to a time-scale representation.

Unfortunately, the simplicity of this technique, either in the Granger or Geweke formulation, comes with two major assumptions that limits its applicability to noise studies. First of all, it relies on the hypothesis that the signals are model stationary, that is, their AR coefficients shouldn't vary with time. Of course, this is not suitable for studying generic non-stationarities, such as glitches. Some authors have tried to overcome this issue, in the case of slow non-stationarities, considering the signals as locally or piecewise stationary [404, 405]. In our noise studies, we have followed this approach by choosing an appropriate model order p .

Second (major) limitation comes from the form itself of the model, and the admissible relations between the variables. Equation (3.49) accounts indeed only for linear interactions between them. Non-linear contributions are usually included in this framework by means of approximations to "locally linear" neighbours [36], or by means of data transformations [406, 407]. Related to what discussed in section 3.5.2, a more consistent approach to this issue is provided by the concept of *transfer entropy* [408, 409], similar to that of mutual information introduced in (3.11). The idea is that, if one process passes information to the other one, there should exist some kind of interaction between the two, which can be indeed considered a sort of (information based) causality relation. As it is written, equation (3.11) doesn't take into account of the flow of information. However, this can be easily inserted, exploiting the same idea of Granger, into the Kullback-Leibler divergence between the conditional p.d.f.s including or not the past values of the *cause* process (X_2 , in this case):

$$T_{2 \rightarrow 1} := \iint p_{X_1}(x_1(t) | x_1(t-1), \dots, x_1(t-p); x_2(t-1), \dots, x_2(t-p)) \times \log \left(\frac{p_{X_1}(x_1(t) | x_1(t-1), \dots, x_1(t-p); x_2(t-1), \dots, x_2(t-p))}{p_{X_1}(x_1(t) | x_1(t-1), \dots, x_1(t-p))} \right) dx_1 dx_2 \quad (3.53)$$

Then, if the information transfer has a net value in one direction, say $T_{2 \rightarrow 1} > T_{1 \rightarrow 2}$, we can conclude that the process X_2 *concurrs* to the causes of X_1 . It has been shown that this method is equivalent to the previous one by Granger if the two process satisfy the correspondent assumptions, and in particular they can be described by (3.49), with Gaussian residuals [410]. While solving one limitation of the Granger method, this solution presents one even more severe, requiring a long sequence of stationary data to obtain a reliable estimation of the p.d.f.s involved; this is the same limitation already encountered in section 3.5.2 for the study of non-stationarities by means of the mutual information. However, this information theoretic approach will be at the base of the other important causality inference method that we

will describe in section 3.7.2.

Application to the study of interferometer fast unlocks

For the reasons explained before, the Granger causality test, and its extensions, has limited applicability to generic non-stationary noise. However, we attempted to use it for a specific noise issue, of high importance for Virgo during O3, that is represented by the so-called *fast unlocks* [LVC49]. With this term we typically refer to “sudden” losses of the interferometer locking, usually starting at the level of the Injection system, and in particular with the lock-loss of the *Inpute Mode Cleaner* (IMC) cavity. In the instants immediately before these, apparently no instability or drift is evident in the strain channel.

This issue has a very long story, and the very first record of it dates back to 2004, at the very first operational phases of the Virgo detector. During O3 and the commissioning phase that preceded it, this was particularly studied by the Injection and Detector Characterisation groups. The occurrence of this kind of events was typically repeated in groups of variable number every few weeks. Since the origin of them has never been understood, it is not clear whether they are consequence of the same cause or not. In several occasions hardware interventions, with the substitution of some components (slave laser pumping diode driver, master laser, and video driver), and various systems tuning have been attempted (master laser frequency and IMC locking) to fix this, but apparently these unlocks keep on returning after a variable time.

As explained in section 2.4.4, associated to every lock-loss there is a time, usually ranging between thirty minutes to one hour, before the interferometer can return back to its working condition. This clearly impacts negatively GW searches, reducing the *duty cycle* of the detector. It was therefore important to investigate them with the analysis tools at our disposal. An important resource to study them is provided by the *Lock-loss Monitor* developed by D. Cohen, which collects the record of every unlock, together with the indication of those compatible with “fast” ones [LVC50]. We made use of the corresponding gps times for our investigations of this issue. As anticipated, the strain channel (or similarly LSC_DARM) didn’t exhibit any unusual behaviour before the occurrences of these unlocks, and no correlations with other channels were found with NonNA. We decide then to apply the Granger test described in the previous section to the quasi stationary regime before the moment of the unlocks.⁶⁰ We firstly attempted to model the strain as an $AR(p)$ process, then we added the past values of some channels of the Injection and Pre-Stabilised Laser subsystems as of in the model (3.49).⁶¹ By means of the BIC test we determined a model order typically $p \sim 10^3$.

Unfortunately, the AR model alone was sufficient to correctly model the strain before to times of the unlocks, and the inclusion of none of the other channels contributed to the reduction of the prediction error, as expressed by the test statistic (3.50).

We tried then to investigate the relations among the other channels used in the model for the strain: INJ_EOM_CORR, INJ_IMC_REFL_I_PRE and _POST, PSL_ML_AC, PSTAB_HF_CORR,

⁶⁰For these tests we made use of the `statsmodels` Python library for the modelling of the signals as AR processes, and for the corresponding `grangercausalitytests` class [411].

⁶¹To this purpose we acknowledge the strict collaboration with the experts of the Injection group, and in particular of Dr. G. Pillant, who helped the author to sort out among the many sensors and controls of this subsystem, and in the interpretation of the results of our analyses.

BsX_ML_PZT_CORR, and INJ_ML_PZT_CORR. These, differently from LSC_DARM, exhibited certain unusual behaviours before the time of the fast unlocks, which was interesting to understand better. However, the results on them were mostly null or “ambiguous”, with hints for causality in both the directions, and therefore not interpretable in the *linear* framework provided by the Granger test.⁶²

The reason for that relies in the inapplicability of equation (3.49) to the highly dynamical, and *non-separable* processes characterising these subsystems (in particular for the relation between INJ_ML_PZT_CORR and PSL_ML_AC; refer to note 62). The method described in the next section aims at overcoming this limitation.

3.7.2 Convergent Cross Mapping

The framework provided by Granger to infer causality relies on *linear* regression and the *separability* of variables. This is evident from equation (3.49). From the information theoretic perspective of equation (3.53), this means that the information content of one time series must be separable from the universe information set in order to test if this can be considered a cause (or not) in the model for the other variable. However, in GW detectors, as well as in Nature, this assumptions are often too restrictive, as the example of the relation among Injection channels described in the previous section has proven.

A possible solution to this limitation, which allows to model causal relations in systems that are most likely governed by highly dynamical non-linear connections, has been proposed in 2012 by G. Sugihara with the introduction of *Convergent Cross-Mappings* (CCMs) [33, 412]. This approach fundamentally differs from what we have discussed so far by modelling time series by means of *dynamical systems*, instead of realisations of stochastic processes. Here, a family of smooth *evolution functions* (either deterministic or not) describes the time evolution of the system in a d -dimensional manifold \mathcal{M} , called the *phase* (or *state*) *space* [413]. In this framework, two time series are said to be causally linked if they are generated from (the projections of) the same evolution functions.

The Takens’ reconstruction theorem

A region of particular importance in the system phase space is that occupied by the so called *attractor*, and its *basin of attraction*, for many aspects resembling the general relativistic idea of an *event horizon*, inside which the evolution of the system is doomed to be confined. If this is the case, the fundamental result of the *Takens’ reconstruction theorem* [414] states that we can approximately reconstruct the attractor from the values assumed by the time series $x(t)$ and its past story. The E -dimensional points

$$x(t) = (x(t), x(t - \tau), x(t - 2\tau), \dots, x(t - (E - 1)\tau))$$

constitute the embedding space \mathcal{M}_X for this reconstruction, also named *shadow manifold* by Sugihara, which depends on two parameters: its dimension E and the time delay τ . The former should equal the minimum sufficient number of independent variables on which the

⁶²To be more specific, the Granger test pointed in favour of both a causality relation from the piezo correction on the frequency of the master laser (INJ_ML_PZT_CORR) and the master laser AC power (PSL_ML_AC), and vice versa.

attractor in \mathcal{M} can be projected (at least two) [415], and can be estimated from the data using the *false nearest neighbours method* [416–418].⁶³ The time delay can instead be determined by the *average mutual information criterion* [419, 420], and usually corresponds to the sampling time of the time series. From the definition in (3.11) (in the case of discrete partitions):

$$I(\tau) = \sum_{i,j} p_{ij}(\tau) \log \left(\frac{p_{ij}(\tau)}{p_i p_j} \right)$$

where p_i is the probability to find $x(t)$ in the i -th partition of the observation set, and $p_{ij}(\tau)$ is the joint probability that an observation falls into the i -th interval and the one after a time τ into the j -th. If there exists a minimum of this function for a certain value of τ then this is a good candidate for a reasonable time delay for the reconstruction of the shadow manifold \mathcal{M}_X . We should stress however that, beyond the values suggested by the previous criteria, the misspecification of E and/or τ affects the results of the CCM inference on causation, with consequent unreliability of its predictions.

If a second variable Y is part of the same dynamical system, the same considerations as of X apply, and the corresponding shadow manifold \mathcal{M}_Y , from the Takens' theorem, is therefore diffeomorphic to that of X : a small region around $\mathbf{y}(t)$ will map onto a corresponding one around $x(t)$. The idea of the cross-mapping is then represented by the fact that if X has a causal influence on the dynamics of Y , we can use \mathcal{M}_Y to estimate future values of X :⁶⁴

$$\hat{x}(t; \mathcal{M}_Y) = \sum_{i=1}^{E+1} w_i x(t_i) \quad (3.54)$$

where the weights w_i are estimated from \mathcal{M}_Y :

$$w_i = \frac{1}{W} \cdot \exp \left(- \frac{\|\mathbf{y}(t) - \mathbf{y}(t_i)\|}{\min_k \|\mathbf{y}(t) - \mathbf{y}(t_k)\|} \right), \quad W := \sum_{i=1}^{E+1} w_i \quad (3.55)$$

with $\|\cdot\|$ the Euclidean norm in \mathbb{R}^E [415].

To test the predictions of this reconstruction, a library of L points $\mathbf{y}(t)$ should be considered, and the corresponding predictions $\hat{x}(t; \mathcal{M}_Y)$ compared to the actual values $x(t)$. If the hypothesis of the Takens' theorem are satisfied a convergence theorem exists for the Pearson correlation coefficient (3.10) between the actual values and their predictions: $r_{x\hat{x}}$. This indeed should converge to 1 as the number of points L sampled from \mathcal{M}_Y becomes larger, since this library will become a more accurate representation of the attractor, and the nearest neighbour points will cluster closer to $\mathbf{y}(t)$.

The convergence of the correlation, hence of the representation, with L is clearly visible in some subplots in figure 3.28, which we are going to comment momentarily.

⁶³Loosely speaking, this criterion is what guarantees that the measure inside exponential in equation (3.55) remains bounded, meaning that the points $x(t_i)$ and $x(t_j)$ keep on remaining inside the same neighbour for a sufficient dimension E .

⁶⁴Notice in this definition of causation a sort of "twist" with respect to what described in the previous section for the Granger test and its extensions. Here, $X(t)$ is said to be the cause of a second process $Y(t)$, in the CCM sense, if *past values of the "effect" can help to predict future ones of the "cause"*.

Application to the study of the effects of squeezing

In order to verify the potentialities of the previous method, we have proceeded to study the interactions of some processes that we knew in advance to be highly non-linear, and for which we already had a good understanding about what relations to expect. This study involved the effects of the injections of *squeezed light*, already mentioned in section 1.3.2, to the Virgo DARM signal [163]. The list of the interesting channels, related to the squeezing subsystem, which were worth to investigate has been provided by D. Bersanetti.⁶⁵

The software infrastructure has been provided by the Python package `skccm` by N. Cortale [421], based on `scikit-learn` [422], where the author have provided the porting of the original pipeline proposed by Sugihara *et al.* [33], converted to `Matlab` by D. Mønster for validation [423]. This includes also the methods for computing the dimension of the embedding E and the delay τ , as described in the previous section.

In figure 3.28 we reported the results of the analysis in terms of the correlation $r_{x\hat{x}}$ between shadow manifold prediction $\hat{x}(t; \mathcal{M}_Y)$ and actual value $x(t)$, as a functions of the library size L . We averaged each point over 1000 realisations in order to get rid of spurious detector noise effects. After that, all the correlation curves seem converge to specific values, as expected from the convergence theorem described in the previous section. The triangular subplot structure of this figure should be read as follows. With the usual convention for the two-dimensional Cartesian plane (x, y) , the channels written on the top of each column represent the X 's, and those on the right of each row the Y 's. Then, the blue line in the various subplots is $r_{x\hat{x}}$ as a function of the library size L . If this converges to 1, this means that $\hat{x}(t; \mathcal{M}_Y)$ models correctly $x(t)$, that is Y can predict X , or that X is the cause of Y in the CCM sense. Refer to comments in note 64. Conversely, the orange lines represent $r_{y\hat{y}}$, where $\hat{y}(t; \mathcal{M}_X)$ is defined analogously to (3.54) with x 's substituted to y 's; when converging to 1, this means that X can predict (be caused by) Y . For example, in the first row the blue lines represent the predictions of `LSC_DARM` by means of the channels written above; in orange the predictions of the latter by means of the former.

For the interpretation of these results we referred to the discussion about simulated data reported in [33, 415]. When only one of the two lines possesses high values of the correlation, this indicates a clear direction of the dynamical effect, and hence which one of the two channels is *causing* the other in the CCM sense. A similar statement can be made when one of the two lines is markedly higher than the other, although not converging to 1. This can be an effect of the noise, as commented in the previous references. When both of them are high, this is a symptom of some “feedback” mechanism, or *bidirectional causality*; Sugihara compares it to the “predator and pray” relation.⁶⁶ In those cases where both the curves exhibit moderate to low values of the correlation Sugihara interpret this as a complex model

⁶⁵Private communication dated February 2019. For the purpose of this analysis the author also acknowledges M. Vardaro and E. Capocasa for the precious advices about the role of the various channels and the interpretation of the results.

⁶⁶More preys can provide nutrition to the predators, but the increase of the latter causes the decimation of the former. In turn, a reduction of the preys diminishes the number of predators. In those situations where the two species reach a sort of equilibrium in their respective numbers, this can be described as a common dynamical system characterised by an *attractor*, as defined at the beginning of this section [424]. This may also be thought as a sound description of certain feedback mechanisms inside the operational principle and control of an interferometric detector of GW.

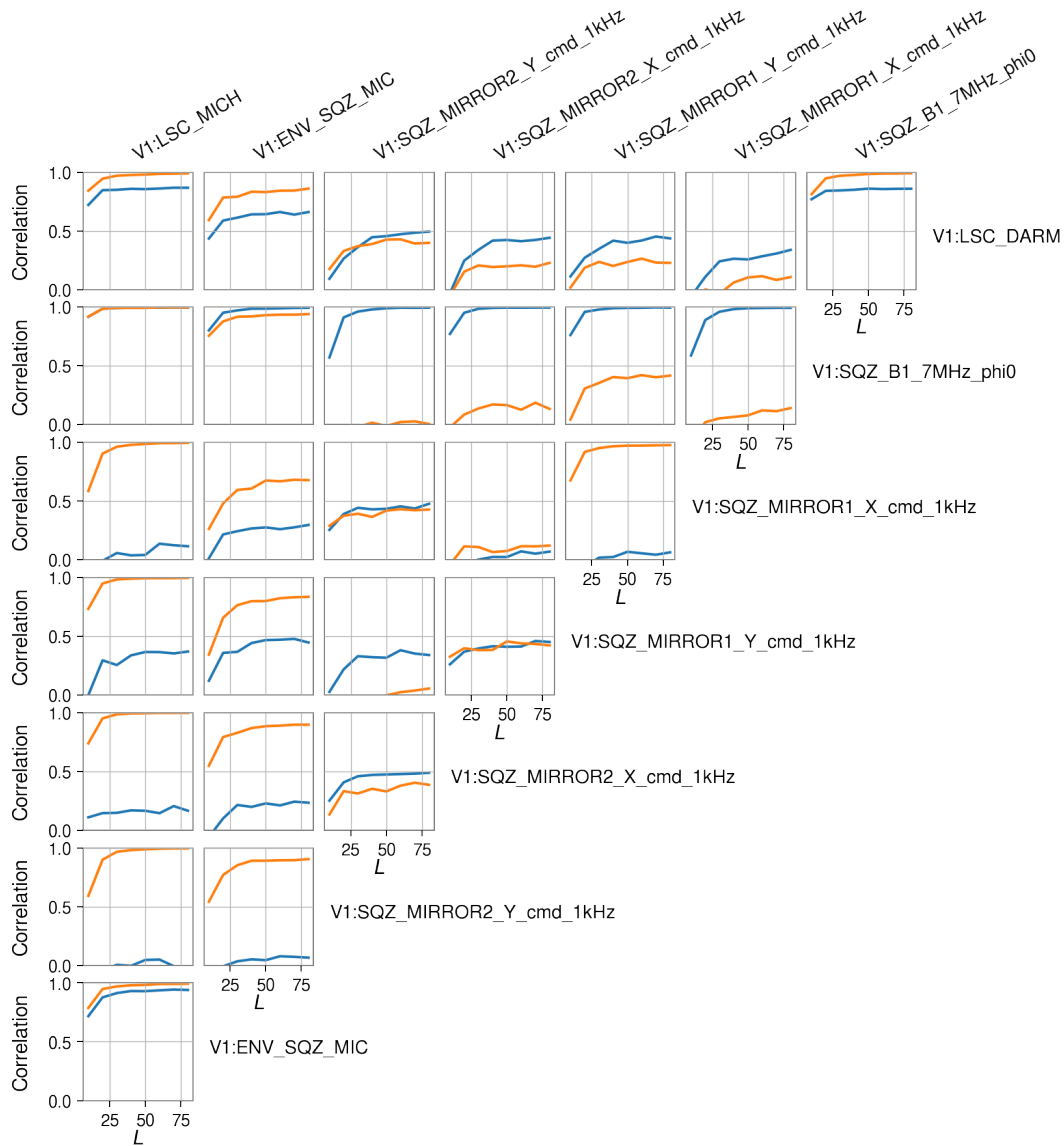


Figure 3.28: Example of Convergent Cross Maps with Advanced Virgo data related to the squeezing and to some important longitudinal degrees of freedom (DARM and MICH). The triangular structure should be read, in the usual directions of the Cartesian plane (x, y) , as the correlation of the values of the channel on the columns (X) with its prediction by means of that in the corresponding rows (Y): $r_{x\hat{x}}$, blue line. The opposite relation, $r_{y\hat{y}}$ is represented in orange. High values of the former can be interpreted in the CCM sense as a causation relation from X to Y (the opposite for the latter). Refer to the description in the main text for the various cases, or to the test on simulated data reported in [33, 415].

with an external forcing of non-coupled variables.⁶⁷

As regards the interpretation of the data at our disposal, probably the clearest thing to notice (and the easiest to interpret) is, in the first column, the longitudinal control of the short-Michelson degree of freedom, $V1:LSC_MICH$, that is predictable/ caused by almost all the other channels (high values of $r_{y\hat{y}}$, the orange curve), while it has almost no influence on them, in particular on those related to the squeezing. Also, in the second row the high values of $r_{x\hat{x}}$ (blue line) hint of a causation relation from the squeezing phase channel (SQZ_B1_7MHz_phi0) to those corresponding to the mirrors (SQZ_MIRROR*). Other relations are less evident and not obvious to interpret.

3.7.3 Some final thoughts about the application of causality studies to GW detector's data

The inference about causation is an advanced topic in time series analysis aimed at overcoming the limitations of those methods based on correlation, or its frequency domain formulation by means of coherence, or time-scale extension with wavelets. In this work we have described three different approaches to it for the analysis of GW detector's data, particularly focussed on the study of non-stationary noise. Although all of these presented some limitations, and none of the analysis described here can be trustworthy considered to provide original insight in the complex structure of the processes taking place inside a detector, this constitutes, to the best of our knowledge, the first documented attempt of such an inclusion, and the potentiality of it should be pointed out.

Large part of the difficulties related to similar studies are rooted in the definition itself of what we mean by causation. Indeed, all the methods that we described rely on a different interpretation of it. Our original instantaneous time delay statistic (3.47) aims at verifying a necessary condition for one process to cause the other, which is based on the common sense belief in *temporal precedence*, where the causes are assumed to precede their effects. We have studied it in the context of "impulsive" excesses of power (*glitches*), for which it represents a natural exploitation of the phase information contained in the complex wavelet transform (or of the cross-spectral density in the case of stationary signals, as described in section 2.3.3). Tests on simulated data has shown its ability to correctly reproduce the injected time delay. Further testing and possible improvements in the formulation of this statistic are currently under study. In particular, it is ongoing the analysis of some of the noise injections performed during O3 [LVC48] to test this technique in a controlled context. The optimisation of the pipeline is also envisioned, in order to make it suitable for fast analyses to be included in the DQRs.

The causality test by Granger, or its spectral version by Geweke and the related concept of transfer entropy, is the most widely adopted in the literature. This is due to the simplicity of its formulation in terms of linear prediction by means of a regressive model, or in terms of the mutual entropy in equation (3.53), which generalises to the non-linear case. It is also highly supported from the point of view of the software, with some standard Python packages reviewed and released [411]. However, its application "out of the box" to generic noise

⁶⁷This is the notable case studied by Sugihara *et al.* of the apparent correlation between sardine and anchovy in the California Current, which was proven to be due to shared climate forcing instead of a direct interaction [33].

studies, and in particular to non-stationary noise, has proven to be unsuccessful for the violation of the assumptions it is based on. As discussed, its applicability is limited to the case of stationary signals, or to those representable by AR processes, as described in 2.1.3.

The method provided by the CCMs is arguably the most complex and difficult to accurately handle. Also its results can be of not easy interpretation, consisting in multiple possible outcomes for the convergence of the correlations. Its formulation by G. Sugihara is quite recent (2012), and its follow up and applications are in constant development but not as mature as those of the Granger method. The most of its applications to-date have taken place in the biological and climatological context [33, 424], and the interpretation of the results took advantage of some common sense (as in the case of prays and predators) or of independent knowledge in the fields. A similar approach can be quite difficult to repeat for a GW detector, with many feedbacks and interacting parts in a non-linear way. Also, the most of its predictions and the related interpretations have been discussed on simulated data with very simplistic “toy models” only, as for example in [415]. These are presumably not suitable in the ambit of GW detectors. Still, even in this case some contradictory results has been observed, especially in the presence of noise and with strongly coupled processes, as discussed in [425]. From a practical point of view, the results of this method are sensitive to variation of some of its parameters, namely the dimension of the embedding E and the time delay τ , as discussed in the previous section. One should carefully verify the convergence of the results with the parameter choices recommended by the false nearest neighbours method and the average mutual information criterion. At a more profound level, the validity of the Takens’ Theorem is not guaranteed for generic shapes of the attractor. In particular, this aspect can’t be directly observed and the assumptions on it can’t be verified (both a priori or after the data has been recorded) with the consequence of the CCM method to return unreliable results. This aspect has been addressed in [425] with simulated data, exploring various configurations of the parameter space for certain coupled processes.

The previous considerations have led us to be careful about the inference to draw with the CCM method, and to present in this manuscript only an example of it in a context under control and with a certain amount prior knowledge on the relations between the various signals. Unfortunately, it was not possible yet to carry out the joint analysis of the “fast unlocks” [LVC49], described in section 3.7.1, in collaboration with the Injection group, as we did with the Granger method. Indeed, prior to this, a more profound knowledge on the method, its robustness, weakness and applicability to detector’s data should be taken into account. For the post O3 phase it is planned the analysis of the noise injections accomplished in September 2019 [LVC48], where we will verify the predictions of this already knowing the causes.

4 Noise artefacts mitigation

All the strategies presented so far are based on the identification of particular noise features, and then the investigation of what may have *caused* them. Consequent *mitigation* strategies have been implemented intervening directly on the latter and checking back if these changes have produced the sought after noise reduction. This is the typical approach adopted during the detectors commissioning phases, such as that from mid 2017 to April 2019, when most of this Thesis work has been developed. During an observation run, these mitigation operations should be postponed until the next weekly maintenance break or, if a specific noise issue is particularly detrimental, the interferometer status flag discussed in section 2.4.3 should be applied, intervening immediately on the noise issue. As a consequence, the corresponding data are not used for GW searches. This impacts the interferometer duty cycle and, in the long run, the number of detected events, as discussed in section 2.4.4.

Moreover, besides every effort, non-stationarities of various extents are naturally present in all data stream. In section 2.1.6 we have discussed how to make, in the presence of “fast non-stationarities” but a clear background noise, robust spectral estimations by means of the mean-median averaged ASD. This is useful for the correct application of the “standard” matched filter technique in GW searches. However, a more detrimental issue is given by the fact that with the improved sensitivities of the Advanced Detectors, and the consequently higher rate of detectable events, the possibility that some (fast) non-stationarity overlaps with one of them is now relevant. This is particularly true for “long” BNS signals, and indeed this was the case of LIGO Livingston detector data for GW170817 [14], and reported in figure 4.1. Gating is in general only a temporary solution since it produces a loss of potentially valuable information.

To cope with that, new post-offline strategies are under development to perform the mitigation of noise artefacts directly on the recorded data. An example of these was the application of the BayesWave algorithm to the removal of the glitch in LIGO Livingston data in correspondence of GW170817 [198]. We will elaborate on this in the last section, with also a comparison to the algorithm developed by us.

In this chapter, which constitutes a sort of “epilogue” about the noise investigation strategies described so far, and an outlook towards new advanced methods, we will briefly present a possible new “de-glitching” algorithm based on *deep neural network* (DNN), which we use for signal reconstruction. This project was started in 2018 by S. Kulkarni and M. Cavaglià from MS&T, which the author has started to collaborate with during his Thesis work. What follows is based on the original work and preliminary results obtained by them, to which goes the credit for the material hereafter proposed. New developments and tests are ongoing to improve this algorithm and check its reliability for the purposes of GW searches. These will constitute the subject of publication in preparation by the LIGO MS&T group and




Figure 4.1: GW170817 signal in LIGO Livingston detector, corresponding to a BNS coalescence overlapped with a DAC saturation glitch [14]. The top map reports its scalogram, obtained with the Omega pipeline [323], where the glitch and the chirp are clearly visible with high SNR (normalised energy). In the bottom plot the detail of the (whitened) strain time series (orange) is shown in correspondence of the glitch. It is also reported its *gated version* (blue) obtained by means of the inverse Tukey window shown in grey. Figure adapted from [14] under license CC BY 4.0.

the author.

4.1 NNetFix: a Neural Network to Fix glitches on signals

The aim of the `NNetFix` (Neural Network to Fix glitches on signals) algorithm is to cope with situations similar to that in figure 4.1, of a glitch overlapped with a GW signal. The idea is to perform the reconstruction of the parts of the signal corrupted by the glitch exploiting a Multi-layer Perceptron (MLP) artificial neural network structure [40]. The main targets have been BBH signals, but tests are ongoing to verify its performances on all kind of modelled searches. In the next sections, we will present the details of the method and some preliminary results on real and simulated data. These have already been presented at a conference by S. Kulkarni [LVC18], and are part of the contents of an article in preparation [285].

4.1.1 The Multi-layer Perceptron and method description

An Artificial NN constitutes a generalisation of the idea behind linear regression described in section 3.5.3, which allows for *non-linear* predictions. The most simple architecture of this kind is given by the MLP. This consists of a certain number (usually two, besides the input one) of *layers* of processing units (*neurons*), where the information is *feed-forward* propagated from the input nodes, through the hidden layer(s), to the output neurons. Several model order selection criteria exist for the optimal choice of the number of hidden layers [343].

Then, given d points of a certain time series x_n , represented as a vector in the corresponding d -dimensional *embedding*,

$$\mathbf{x}_n = (x_n, \dots, x_{n-d+1})^T, \quad (4.1)$$

the first (hidden) layer usually consists into a number of linear transforms of it, with weights w_{ij} ,

$$\mathbf{w}_j^T \mathbf{x}_n = \sum_{i=0}^{d-1} w_{ij} x_{n-i}, \quad j = 1, \dots, p \quad (4.2)$$

which are then operated on by a non-linear transformation:

$$h_j = f \left(\sum_{i=1}^d w_{ij} x_{n-i} \right), \quad j = 1, \dots, p. \quad (4.3)$$

The function f has the role of *activation function* for the next layer. Traditionally, it used to be a *sigmoid* function, similar to a “smooth step”, although it has now been shown that a more convenient choice is provided by a *rectifier* (or *ramp*) function [426]:

$$f(x) = x^+ = \max(0, x) \quad (4.4)$$

and the unit implementing this is called a *relu* (rectified linear unit). The output of the second layer (in the case of just two of them, otherwise the last) gives the non-linear prediction, which is just a linear combination of the hidden unit responses:

$$\hat{x}_{n+T} = \sum_{j=1}^p v_j h_j. \quad (4.5)$$

Then, given a data set of N embedded vectors \mathbf{x}_n and the corresponding future values x_{n+T} , the parameters of the model can be set as to minimise the prediction error (or loss function), provided for example by the mean-squared error:

$$MSE := \sum_{n=1}^N (x_{n+T} - \hat{x}_{n+T})^2. \quad (4.6)$$

This process is referred to as the *supervised training* of the neural net [40]. Its implementation is done by means of the method of *back-propagation* [427], which is based on the computation of the gradient of the previous quantity, and the corresponding *descent* towards the minimum.

In the NNetFix algorithm, this idea has been used to reconstruct, after training, the evolution of the characteristic “chirp-like” sinusoidal pattern of a GW signal (1.47) in the presence of a corrupted portion of it, for example by a glitch (as in figure 4.1). Indeed, given a certain segment of strain data affected by this issue, we firstly proceeded to *gate* it, smoothly zeroing out its values by means of an *inverse Tukey window* (or *tapered cosine*) [250], as in the bottom plot in figure 4.1. Then, passed to the algorithm the chunk of data centred around the gated portion as input, we expect the trained MLP to provide an accurate reconstruction of the missing piece of data, as it would have been in the absence of the glitch. The number of nodes in the output layer is therefore set equal to the size of the zeroed interval of the glitch.

4.1.2 Training by means of simulated waveforms

We trained our NN with data generated superimposing to the Advanced LIGO reference noise curve the BBH coalescence waveforms obtained by means of the `LALSuite` data analysis routine [428], with `IMRPhenomD` waveform models [190]. This is the most delicate part of the whole process. Indeed, a general consideration about it is that predictions of any algorithm are only as good as the dataset we trained it on [429].¹ For this reason, in order to achieve the desirable performances, we needed to ensure that this training dataset was both dense, in terms of the granularity of the various parameters of the waveforms and the glitch (or gating), and in diversity of the possible scenarios presented, including mass ranges, sky position, SNR, etc. Of course, some sort of trade-off needed to be applied, in order to make the network training to converge within the given time and computing resources. On the other hand, including data waveforms that have highly degenerate features had the risk of *over-training* the network to produce a single kind of outcome all the time.

For these reasons, to optimise the training process, we made the following assumptions. Given the particular scenario of a GW signal with a glitch superimposed on it, we assume to know in advance the following properties: (i) an initial “rough” estimate of the binary component masses, provided for example by a preliminary analysis made with the glitch gated, (ii) the duration of the gating, and (iii) the time it occurs before merger. Then, we could generate a training set adequate for this scenario simulating component masses in the credible region around the previous preliminary estimate, and the other extrinsic parameters (sky position, angle of inclination, coalescence phase, and polarisation angle) uniformly distributed over their respective ranges. Moreover we uniformly simulated its SNR with respect to the background noise it is superimposed, randomly generating it from the Advanced LIGO reference PSD curve. A further precaution was to produce multiple copies of the previous waveforms, each of them superimposed to a different noise realisation, in order to be sure that the algorithm was trained to predict the signal without learning the noise properties at the same time. Moreover, a sanity check was made on real detector noise, by an amount of 10% of the whole training set.

This dataset, comprising $\mathcal{O}(10^4 \div 10^5)$ samples, was then used to train the algorithm, using the details about the glitch at points (ii-iii) to gate the corresponding data by means of an inverse Tukey window with roll-down parameter $\alpha = 0.2$.

The implementation of the code was done by means of the `scikit-learn` Python library [422]. The MLP architecture was constituted by one hidden layer of 200 neurons. The mean-squared error (4.6) was used as a loss function, together with the rectifier activation function (4.4) and the stochastic gradient descent optimisation method for the back-propagation. The dataset was split into a ratio of 60 : 30 : 10 to form the training, testing and cross-validation datasets respectively. By means of the latter, we have obtained a cross-validation accuracy of 95%.

¹Under many aspects, we already encountered this issue in section 3.5.3 when we tried to extrapolate the regression predictions to new segments of data, while trained on previous, non-stationary ones.




Figure 4.2: Performances of the NNetFix reconstruction as expressed by the percentage error in SNR recovery. In the left hand side plot, for a fixed glitch (or gate) duration of 50 ms, the comparison of the performance of the reconstruction versus the gate data is reported as a function of the time before merger of the glitch. On the right-hand side, both the duration and the time before merger are made varying; the different lines colours represent the reconstruction performances for various glitch durations, as a function of the time before merger. Images courtesy of S. Kulkarni.

4.1.3 Results on injected and real BBH signals

We tested the algorithm on both injected BBH signals, generated using PyCBC software [17], and those corresponding to actual events, available from the Gravitational Wave Open Science Center (GWOSC) [202]. The reconstruction performances have been evaluated using as a metric the (percentage) *error in SNR recovery*, defined as the difference between the SNR obtained without gating (that is, in the absence of a glitch) and that of the reconstruction with NNetFix, performing the gating, using the same “best matching” template of the original one.

In figure 4.3 we simulated the reconstruction of the signal corresponding to GW150914 [8], as measured by the LIGO Hanford detector, upon which we assumed to be superimposed a glitch of duration 50 ms, occurring 15 ms before the time of the merger. The training for this signal has been obtained simulating black holes with masses $m_1, m_2 \in [15, 45] M_\odot$. The scalogram in figure 4.3a represents the signal with the central part, corresponding to the hypothetical glitch, gated, as visible from the dark blue band. In figure 4.3b its reconstruction by means of the NNetFix algorithm is shown, which should be compared with the original GW150914 signal reported in figure 4.3c. Their similarity appears remarkable.

The accuracy of the reconstruction of GW150914 is reported in figure 4.2, where we made varying the two main parameters of the glitch (or gating): duration and time before merger. In figure 4.2a the percentage error in SNR for a glitch of duration 50 ms (reported for convenience as the horizontal black *whiskers*) at times before merger from 150 to 10 ms is shown for the gated data (blue line) and for the reconstructed ones by means of NNetFix (red line). Even in the very proximity of the merger, the error in the latter is just slightly above 20%, while with gating is almost 50%. In 4.2b both of the parameters are made to vary; the three


lines correspond to different values of the glitch duration, as functions of the time before merger at which they occur.

4.2 Discussion and future developments


Although what reported in the previous section are just preliminary studies about this possible new technique, the method provided by `NNetFix` has effectively proven to be able to perform the reconstruction of a GW signal in the presence of a glitch. One important point should be noted here; this reconstruction does not provide “new” information about the putative GW signal, besides that present in the non-gated part of it. For this reason, it can just be seen as a “sound way” to interpolate the data in the gap made with the gating. From the point of view of the SNR, this has proven to constitute an advantage with respect to relying on gated data. Another aspect we are investigating is the improvement it could provide to the fast sky localisation of the source. Indeed, being able to run it in quasi-real time on detectors data, the reconstruction it allows could provide a better sky localisation than that obtained by gating the corrupted portion of the data.

For the reasons discussed at the beginning of this chapter, the problem of “deglitching”, or in general “denoising” performed post-offline, is a very active field of research within the LIGO and Virgo collaborations. At the time of writing (December 2019) there are at least two other algorithms operating on this. One is the already mentioned `BayesWave` [195], successfully applied to the removal of the glitch in the LIGO Livingston data in correspondence of GW170817 [198], represented in figure 4.1. This was done exploiting a fit of the data, that is, glitch plus the GW signal, in terms of (real) Morlet wavelets (3.27). For the morphologies of the two were very different, the former impulsive and broadband and the latter slower and quasi-monochromatic for quite large portions of it, the wavelet reconstruction presented very disjoint components for them. Moreover, the authors exploited the coherence of these components between the two detectors, verifying that only the “signal part” of them was in fact coherent. This sufficed to resolve the signal, disregarding the presence of the glitch. In principle, in every other situation like that, with wavelet components clearly distinguishable between glitch and GW signal (therefore primarily for BNS signals), this method can be considered reliable for deglitching.


Another very recent approach, strictly related to the one presented for `NNetFix`, is that described in [284] by some members of the LIGO group of the University of Illinois at Urbana-Champaign. In this recent publication, the authors have adopted a more complex network architecture, particularly suited for the study of time series [40], constituted by a *Recurrent Neural Network* (RNN) [430]. They have presented the results about its application to the *denoising* problem, that is, the extraction of the GW signal from the noisy detector data. This, in a certain sense, is most closely related to the matched filtering technique than to the deglitching that we have presented. However, the authors are verifying this also for the case of general non-stationary noise.



(a) GW150914 gated.



(b) GW150914 reconstructed.



(c) GW150914 original.

Figure 4.3: Reconstruction of GW150914 signal in the LIGO Hanford detector by means of `NNetFix`. The top scalogram report the gated data, as for a glitch of duration 50 ms, occurring 15 ms before the time of the merger. The central scalogram is its reconstruction by means of `NNetFix`; this should be compared with the original version of it, reported in the bottom scalogram. Images obtained by means of `GWpy` [288], courtesy of S. Kulkarni.

5 Conclusions

In this dissertation, we have presented the implementation of several data analysis techniques for the study of noise in Advanced GW detectors, with particular attention to what are the needs of the search pipelines for different kind of GW sources. The arguments have been selected from various research fields, ranging from Digital Image Processing to Machine Learning, with the aim of providing state of the art methods for the study of specific noise issues, and in particular non-stationary noise. Two of them, the BLRMS based stationarity test and the instantaneous time delay statistic, constitute an original contribution by the author. The former has been conceived in order to overcome the limitations present in the most of the currently adopted tests devoted to the same goal. The latter wants to exploit better the information contained in the wavelet representation of the signals, adding the new concept of “causality” to their study.

This activity has been part of the work of the author within the Virgo collaboration for the period 2017-2019, in fulfilment of the commitments as a Ph.D. student. The relevance of this work has been motivated in particular in Chapter 2, and in section 1.4 in relation to the search of GWs. As explained, in the collaboration framework of experiments like Advanced LIGO and Advanced Virgo, the activity of Detector Characterisation has the role of bridging the gap between the actual functioning of the detectors, and the data they produce, and the consequent analysis of it, aimed at GW searches. In order to make reliable inference on the latter, we must possess profound knowledge of the properties of the detectors and their noises. In turn, improvements on the latter should be guided by the *quest for knowledge* about new, and possibly not yet observed, sources of GWs. For these reasons, we have investigated some new techniques able to tell the search pipelines when the data are valuable for being analysed, and when not, as in the case of the BLRMS based stationarity test (from O3a part of the Virgo DQR) and the Rayleigh test; sections 3.1.2 and B.1. At the same time, we have developed some tools aimed at the characterisation of the noise, with the goal of mitigating its effects, when possible, both intervening on the detectors and in post-offline phases.

We have chosen to organise the discussion about the author’s studies following, at least in part, the logic of the process of investigation typically adopted by a *noise hunter* or a *Detchar shifter*, especially in those circumstances when no prior hints or directions were available, or new paths for the understanding of known problems were sought for. In these situations, the starting point usually consists of the identification of the particular noise features that could affect GW searches. This aspect has been described in particular in sections 3.1 and 3.3, for non-stationary and non-Gaussian noise, respectively. Then, the strategies to relate them to the other information provided by the auxiliary channels have been discussed in sections 2.3.3 and 2.3.4, for the case of stationary noise, and from section 3.5 throughout the rest of this dissertation, for the non-stationary one. In the exposition of the latter, we

have presented the various strategies in order of complexity and connections to the previous ones. In particular, the last two methods, causality studies (sections 3.6.3 and 3.7) and deglitching (section 4.1), are arguably the most complex ones. Their implementations and predictions still require further checks and investigations. Nonetheless, we have decided to include them in this dissertation as an outlook of future activities for both the author and larger groups within the collaborations. In some circumstances, we have also presented methods that have proven not to be particularly useful, or even suitable, for the specific purposes they were originally thought for. We have chosen to include them in this dissertation for the relevance of the discussion about their limitations in relation to the particular noise features we were investigating. Also, their inapplicability for a particular analysis is a way of characterising the data involved, and, most importantly, the consequent discussion is sometimes a starting point for developing new methods. For example, this was the case of the PSR stationary test (section 3.1.1), which led the author to the development of the BLRMS based method described in 3.1.2. Moreover, the failure in many circumstances of the Granger-Geweke causality test (section 3.7.1) led us to investigate more advanced, non-linear methods, such as that provided by CCM in section 3.7.2. Again, the same thing happened for linear regression analysis, when adopted for predicting the values of a certain signal; we realised that the non-linear structure provided by an MLP was best suited for it (section 4.1).

With this said, the implementation of newer data analysis techniques, exploiting state of the art strategies in computing and *Machine Learning* will hopefully make available, in the not too distant future, several reliable algorithms and tools to get rid of the effects of non-stationary noise from GW data. This is the best wish the author can make about the future developments of the ideas presented in this Thesis, and the most desirable conclusion for it.

Lastly, but most importantly, what presented has been part of a collaborative work (as it should be within a Collaboration). All of the ideas and the concept presented have benefited, in some form, of the help, the suggestion and (often) the criticism of other members of it. Whenever possible, especially in those cases where the contribution did come from a specific person or conversation, we tried to acknowledge it in the text. Also, the authors of the tools we have made use of, and our collaborators in specific projects, have promptly been acknowledged in this dissertation.¹

¹With the sincere hope they will appreciate their mention.

A Examples of noise studies with Virgo data

In this appendix we present some further examples of noise studies where the tools described in this manuscript, and in particular the methods developed by us, have found application for the purposes of the characterisation of the noise of Advanced Virgo. We have decided to collect them here, after the presentation of the various arguments in the main text, in order not to “break” the theoretical links and the signal processing thread that have guided the development and presentation of the various methods, in response to the general needs of the various search pipelines and in relation to the characteristic detector noise features. Moreover, in the most of the following examples we will made use of more than one of the previous tools, and comment their connections and the complementary information we can gather from their joint application.

Several technical aspects of the functioning of the detector will be touched in the following examples. It is far from the purposes of this manuscript, and the author, to provide here the adequate contextualisation and description of them in order to fully understand the various noise issues. Hence, we have decided to heavily rely here on the corresponding references to various studies presented in the Virgo documenting system (TDS) and in the logbook. These are mostly meant for internal use, and some of them only accessible with LVC credentials. For this reason, these references has been included in the separate bibliography at page 233, and the corresponding entries labelled with the “LVC#” tag.

A.1 Beam splitter control noise at 150 Hz

From the 12th of May 2019, during the O3 science run of Advanced Virgo, it was noticed an increase of noise around the 150 Hz spectral line [LVC51]. It was Omicron to identify this issue for first as a *cluster* of high rate glitches, with central frequency compatible with this value and $\text{SNR} > 10$; some of them were enough intense to appear broadband in the spectrum and cause consequent range drops.¹ From the observation of the spectrogram it was evident that they corresponded indeed to an increased energy in the “bump” region around this line.

In figure A.1 we have reported the ASD of LSC_DARM in the low frequency region and in the band around 150 Hz for three different times. At 6:00 and 9:00 AM there were bad

¹This noise was also observed with our stationarity test BRiSTOL, as shown in the method presentation in [LVC22]. However, for the reasons explained in section 3.1.2, this tool was usually applied with glitch removed, hence the manifestations of this noise were clearly less evident than in Omicron, optimised for the opposite purpose.

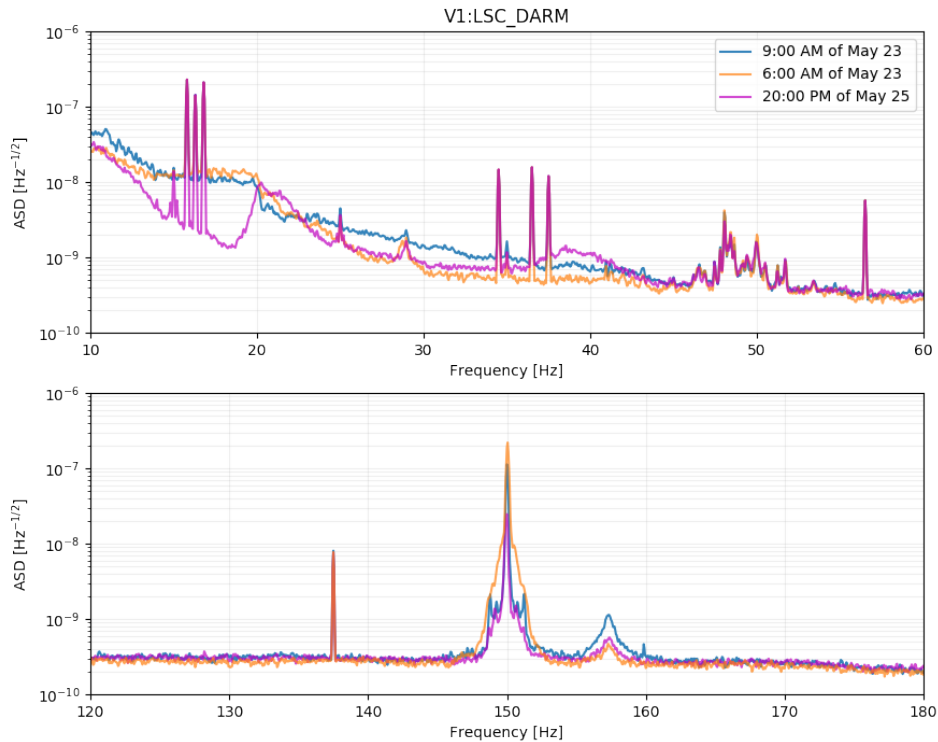


Figure A.1: ASD of LSC_DARM in the low frequency of the spectrum (top plot) and around the 150 Hz resonance for three different periods. Associated to the latter, it was observed a new family of glitches starting from the 12th of May 2019. The data at 6:00 and 9:00 AM were influenced by bad weather conditions, as visible in the higher values of ASD in the lower end of the spectrum. Notice also an increase of noise in the bump surrounding the 150 Hz line.

weather conditions, as manifest from the increased low frequency noise in the top plot. During the latter, there were ongoing adjustment operations to the detector, as we will describe momentarily. At 20:00 PM the weather was much quieter and also the noise at low frequency and around 150 Hz. Notice however that the ASD representation of the data is not meant to highlight noise non-stationarities, like the aforementioned glitches, but only the *average* noise power content.

As a consequence of this noise, a region of about 6 Hz needed to be excluded from the searches for gravitational waves. Moreover, in correspondence of the times of occurrence of the various glitches, it was necessary to apply vetos on the corresponding data, as discussed in section 2.4.4.

On the 23rd of May, during some adjustment operations for the tuning of the interferometer working point, it was noticed that these glitches seemed to disappear in correspondence of a specific control status of the beam splitter (BS) [LVC52]. In particular, they were absent when the adjustments were performed and the *beam splitter* (BS) set to “full bandwidth control” (with the corresponding status flag channel SAT_BS_AATX_FLAG equals 2). Instead, they were back when the BS returned to “drift control” (given by two independent controls for the low and high frequency movements of it, represented by SAT_BS_AAT*_FLAG equals 1).

A standard analysis of correlation, and in particular making use of the “full” version of




Figure A.2: Correlation between BS control mode and noise at 150 Hz. The blue line represents the value of the BLRMS of `Hrec_hoft_16384Hz` in the band [145,155] Hz, with line subtracted and glitch removed, as described in section 3.1.2. The black lines represent the control flags of the BS for the x (continuous) and y (dashed) directions.

the BLRMS written in equation (3.1), wasn't able to confirm this relation or to give more insight about it. This is due to the fact that in the region containing the 150 Hz line the energy (hence the rms) is largely dominated by the line itself, being its ASD two orders of magnitude greater than that of the surrounding bump where the glitches manifested [LVC53]. Moreover, in the presence of glitches the values of this BLRMS time series were rich of "spikes", corresponding to them, and hence not representative of the high/low level of noise in this region. We proceeded then to repeat the analysis with our modified version of the BLRMS in the band [145,155] Hz, with spectral lines (in this case, the one at 150 Hz) and glitches removed, as described in section 3.1.2. The corresponding plot is reported in the blue line in figure A.2. The analysis with the `NonNA` correlation tool (described in section 3.5) correctly identified what observed during the interferometer adjustment, with high values of the correlation of the previous BLRMS with the control flags of the BS, represented by the black lines in figure A.2: continuous for the x direction control and dashed for the y direction one respectively. In particular, it was clear that when the BS was in full bandwidth control (flags equal to 2) the 150 Hz region was less noisy, the opposite of when they are in drift control (flag 1).

This analysis was presented at the following Virgo commissioning meeting [LVC9] and convinced about the relation between the control mode of the BS and the increased noise around 150 Hz. Overall, the net effect of this issue, accounting for both the increased noise level and the larger amount of glitches causing range drops, has implied to the *average* BNS range (1.65) a reduction of 2 Mpc when the BS was in drift control with respect to when it was

in full bandwidth.² Moreover, the structures around this line had already been the subject of our analyses, as reported in section 2.3.3. In particular, they were found by MONET the result of the non-linear coupling of the angular control noise with this particular harmonic of the electric mains [LVC11]. Hence, the interpretation of this increased noise has the consequence of an higher level of the previous coupling when the BS was in drift control with respect to when it was in full bandwidth. This information helped the experts of the corresponding subsystems to intervene in the following months and limit this coupling still maintaining the drift control of the BS, more suitable for more stability and duty cycle of the interferometer with respect to the full bandwidth control (although less glitchy).

A.2 “Flat noise” investigations

The so-called “flat noise” has constituted the main limitation to the Advanced Virgo sensitivity during O3. Even after the end of it, this issue has not received a convincing explanation and, most importantly, been solved yet. This noise manifests itself as an unexplained, from the point of view of the noise budget (refer to section 2.2), additional contribution in the region of the “bucket”, between 80 and 300 Hz, with an ASD approximately flat, or proportional to $f^{1/4}$, as it seems compatible from recent noise studies and fit [LVC54, LVC55]. From a finer grain perspective, it also presents (presumably) a cut-off at 400 Hz, and a “bump” between 200 and 300 Hz [LVC56]. As this is the region of highest sensitivity of the detector for the search of GWs (refer to the discussion in section 1.3.3), the presence of this noise has limited the SNR of such signals and reduced the BNS range of the detector. The latter was estimated to be 5 Mpc ($\sim 10\%$ of the total) lower than the expected one from noise budgets [LVC57].

A large number of studies and direct tests have been carried out during O3 and in the previous commissioning phase in order to acquire better insight about it. In particular, the best of the current knowledge has been the result of the extensive investigation efforts of M. Was, A. Allocca, A. Chiummo, and P. Ruggi, with the collaboration of all the commissioning team. From the Detector Characterisation point of view, we have repeatedly investigated this issue, with almost all the instruments described in the main text of this document. We report here some of the results of our analysis, as presented in various logbook entries and at a dedicated commissioning meeting on the subject [LVC26].

In figure A.3 we have reported the ASD of the Virgo strain channel during December 2018 (dark blue line), for a particular value of the DARM offset. Further comments about the latter will follow momentarily. The various noise contributions are represented in different colours, and most noticeably the fit ($f^{1/4}$) of the residual unexplained noise, which we attribute to this flat noise, is shown in purple. Clearly visible the gap in the bucket region between the measured ASD (dark blue) and the expected noise budget expected for the thermal noise and the shot noise (light blue) and the suspension f^{-4} noise (yellow).

Although several mechanisms can in principle contribute as an additional frequency independent noise to the strain ASD, we are tempted to think that this, in the present sense

²This information however doesn’t take into account the better stability and higher duty cycle achievable when the BS is in the former control mode. This has imposed a trade-off between this and the presence of the glitches, thoroughly discussed by the commissioning team during the successive months.




Figure A.3: Effects of the *flat noise* on the ASD of the Advanced Virgo detector. The blue curve represents the measured ASD of the strain channel during December 2018, when some tests on the DARM offset were performed. The yellow, green and light blue line represent the noise budgets for the known sources of noise. In purple is shown the fit for the power law model $f^{1/4}$ of their residual with respect to the observed noise level. This is attributed to the flat noise. The orange line is the total fit of the observed noise, including the latter. This image has been taken from [LVC54].

of the term, is coming from a single though unknown source. The first occurrence of it was documented in October 2017, right after the end of the O2 science run, when the input power was increased to 26 W [LVC58]. In the spirit of what discussed in section 3.7, we believe that the increase of power *is not the direct cause* of this noise, since during O3 the intensity of the latter has varied while maintaining the same 26 W of power at the interferometer input. In particular, this noise was higher right after O2, when the input telescope was not tuned, then how it has been during O3 [LVC55]. In January 2019, it seemed to have reached its minimum as a consequence of the improved tuning of the alignment of the BS and of other angular controls [LVC59]. From this the idea that one of the possible causes of the flat noise was some scattered light; refer to its description in section 2.2 and 2.4.1. Conversely, it was observed that changing the DARM offset, hence the power exiting the antisymmetric port of the BS (refer to the description in section 1.3.1), this noise was increased approximately as the square root of the power arriving at the detection bench photodiode(s) SDB2_B1_DC [LVC54]. The data reported in figure A.3 corresponds indeed to a particularly high value of this offset, and a consequently high level of the flat noise. Another mechanism that we have recently found to increase its value is by changing the alignment of the *output mode cleaner* (OMC), either by means of tilting the detection bench SDB1 or by using pico-motors on the OMC steering [LVC60]. In particular, we have found that misaligning the OMC reduced the noise. This seems to be a consequence of the fact that a good OMC alignment facilitates the transmission of *higher order modes* of the incoming beam to the detection photodiodes, which resonate with the main mode carrier and cause an increased noise with a *flat* ASD shape.




Figure A.4: Application of the BRiSTOL stationarity test to the study of the flat noise. This time-frequency map represents the p -values for the Kolmogorov-Smirnov test used to assess stationarity, as described in section 3.1.2. With respect to figure 3.5, we have used different colours for the regions where the stationarity hypothesis should be rejected with significance $\alpha = 1\%$ (shades of red) and where not (green). The blue shaded regions correspond to the times where the operations on the OMC alignment have been executed.

After these tests, we had the clue that the DARM offset and the OMC alignment were related with the flat noise. The reason for this is still not clear, though. We have then proceeded to analyse the times the previous operations have been carried out, in search for correlations, increased number of glitches, or variations of their numbers, with the analysis tools described in the main text. Unfortunately, it was only possible to exclude some culprits (almost certainly the scattered light) without finding further hints about the causes of it [LVC26]. Nonetheless, this example is of particular relevance as a complement to the material described in this work since it reports the characteristic sequence of investigations typical of a noise study.

First of all, besides the previously described increase of the noise level in the bucket, the test of stationarity with BRiSTOL and the Rayleigh Gaussianity test described in sections 3.1.2 and 3.3 presented no variations in their results during these operations with respect to before and after them. For example, in figure A.4 we have reported the p -map produced by BRiSTOL for testing stationarity. Differently from figure 3.5, we have used two different colours to distinguish the time-frequency regions where the stationarity hypothesis should be rejected with significance $\alpha = 1\%$ (shades of red) and where not (green). Except for few insulated bins, and mostly clustered around the 50, 150 and 300 Hz spectral lines, during the operations of (mis)alignment of the OMC described before [LVC60], highlighted by the blue shades, it is not evident any difference in the stationarity of the data in the region of interest of the flat noise.³ This result supports the hypothesis that the flat noise is actually

³Notice that the marked non stationary regions at the beginning of the blue shades correspond to the times the OMC was misaligned, hence the corresponding data was clearly not stationary. Also, during the first few minutes this map appears of a uniform green colour. This is a consequence of the fact that the interferometer was not in




Figure A.5: Application of the Rayleigh test to the investigation of the flat noise in correspondence of the alignment operation on the OMC [LVC60]. The blue curve corresponds to the OMC misaligned while the orange one is for it aligned. The dashed black line represents the median of the Rayleigh test statistic (refer to the corresponding discussion in 3.3) under the null hypothesis of Gaussian data. Qualitatively, no difference is visible between the two curves in the frequency region of interest of the *flat noise*, where both of them are compatible with Gaussian noise (except for the usual spectral lines at multiples of the electric mains frequency).

stationary. Hence, a coherence investigation (section 2.3.3) is preferred with respect to one of its cross-correlations (section 3.5). We will come back to elaborate on this momentarily.

Moreover, the Rayleigh test described in equation (3.8), given by the ratio of the standard deviation of the signal DFT and its mean (Ry_3), is shown in figure A.5. The blue curve corresponds to a time interval where the OMC is misaligned while the orange one is for it aligned. The dashed black line corresponds to the median of the previous test statistic (for the corresponding number of averages M ; refer to the discussion in 3.3) under the null hypothesis of Gaussian data. What is evident from this figure is that, except in correspondence of known spectral features, like the lines at multiples of the mains frequency, the data are compatible with the hypothesis of Gaussianity.

Since these preliminary analyses support a stationary and Gaussian noise, some possible sources of it are automatically excluded. One example is the scattered light, or at least a dominant contribution from it. To further convince about it, one can observe the glitchgrams of Omicron during the times of the previous analyses, and notice that in none of them there was an increased amount of glitches. Also, from the detailed observation of the wavelet transforms computed in these periods of data one can notice a substantial absence of any hint for a dominant contribution due to the presence of glitches.

The investigation of a flat, stationary and Gaussian noise is quite difficult. However, as a last attempt, we applied the “brute force” correlation analysis described in 3.5, and in particular that of the coherence by means of Bruco [LVC8] (and our multiple channel

“low noise conditions” there, hence the calibrated strain not reconstructed, resulting in an “uniformly” absent data stream.




Figure A.6: ASD of Virgo LSC_DARM in the region corresponding to the 83 Hz wandering line during November 2018. The purple line refers to the 10th of November, when this line was centred at 83 Hz, while the blue one is of the 23rd of November, when this line appears to have moved to 85 Hz. Image obtained by means of Virgo `dataDisplay` [LVC13].

extension described in section 2.3.3). Unfortunately, none of these analysis has produced valuable hints about where this noise excess may come from, as reported at the dedicated commissioning meeting [LVC26].⁴ In fact, both of these analysis were likely to fail from the beginning in finding anything interesting related to this noise. The correlation analysis used to find time domain similarities is not very informative for a stationary signal, especially for one with a flat spectrum and a delta as autocorrelation. For the same reason, the frequency domain coherence is likewise not very effective when the spectrum of the signal under study is flat, with no characteristic spectral features.

The discussion included in this section represents a summary of the state of the art of the knowledge and (most of) the investigation performed until the end of O3 as regards this annoying noise source that has constituted arguably the most relevant limit for the Advanced Virgo sensitivity during O3, and which will certainly be one of the main targets for the next commissioning phase before O4.

A.3 The 83 Hz wandering line

Another thoroughly investigated, and still not fully understood, noise issue in the Advanced Virgo data during O3, and the commissioning phase before it, was the “infamous” 83 Hz wandering line. This has been firstly documented by the author in December 2019 with a two-month long survey of it [LVC61], plus some possible references to it during the 10th Virgo commissioning run (C10) in August 2018 and to its “stationary” behaviour back in 2017, before that Virgo joined O2 [LVC62].

⁴As regards Bruco this result was already evident from the daily reports automatically produced. Even finer grin investigations, with different resolutions, and in the time interval of the previous tests, have revealed no hint of coherence.

This line, in its “normal”, stationary position appears as a spectral line of moderate intensity, barely visible with respect to the surrounding noise floor, at 83 Hz, and with about $0.5 \div 1$ Hz of width. In figure A.6 we reported the ASD of Virgo LSC_DARM, where it is visible as a small “bump” centred at 83 Hz on the 10th of November 2018 (purple line). On the 14th of November it disappeared and after 8 days (22nd) reappeared at 85 Hz, as visible from the blue line in the previous figure, referred to the 23rd of November.⁵ Nothing very harmful for the Virgo sensitivity yet. The appearance of a spectral line with similar frequencies was already documented in 2017 to be associated with the activation of the stepping motors of the *suspended power recycling bench* (SPRB) [LVC62].

From the 2nd of December its movement became more significant, involving variations of about 1 Hz every ~ 1000 seconds. A similar change was not suitable to be studied, by means of NoEMi [273], optimised for the identification of stationary or very slowly moving lines, and the automatic rejection of anything faster.

As a first thing to observe, this line is not present in the spectrograms of other commonly studied auxiliary channels, which, witnessing the same noise issue, could have led our investigations in certain directions rather than others. Then we proceeded to its analysis with the instruments described in this manuscript. For example, in [LVC35] it was studied by the author its excursion from ~ 85 to 90 Hz during the 5th of December. We proceeded with the tracking of the line by means of the Line Tracker algorithm described in section 3.4, than to its cross-correlation analysis with NonNA (section 3.5). A Pearson’s correlation coefficient of 67%, and visually a remarkable similarity, was found with some accelerometers (ACC_) on some filter levels of the long suspensions (Sc_) of the BS and the SR, and in general of the environmental monitor (ENV_) in the central building.⁶

With a frequency evolution not easily recognisable, due to the long and frequent periods of unlock, this line “jumped” to 98 Hz in the first days of January, and then it stabilized back to 83 Hz on the 13th. Another jump on the 15th took this to 100 Hz in one day, and back to 83 in other three days [LVC61, LVC63]. the spectrogram of the first part of this evolution has already been presented in figure 3.10. It is remarkable how this time the frequency evolution appeared smoother than before, and over longer time scales, especially during the descent. As a consequence, the correlation analysis performed on its tracking didn’t identify the same channels as before, and in general it favoured slow trends, as it was approximately the behaviour of the line during the descent to 83 Hz. In particular, it was found correlated with the temperature of various environmental monitors all around the detector.

Similar evolutions of this line, where, from a stationary behaviour at 83 Hz, it suddenly starts to move, going up to 110 Hz and back, were repeatedly observed during 2019 [LVC16]. Two aspects are not understood about this. First of all, it is not clear what have triggered the departure of the line from its usual position at 83 Hz, say, on the 15th of January (nor what has made it to drift back there on the 20th). To investigate this we tried two approaches.

⁵It is not excluded the possibility that it didn’t actually disappeared, but it just get masked by the nearby 87.1 Hz calibration line during its frequency evolution between the 14th and the 22nd. Due to missing data in correspondence of periods of unlocks, and other noise features in the data, it was however impossible to track it to confirm this claim.

⁶Apparently, it was not found anything directly related to the PR, as previously reported in [LVC62], but only close to it.




Figure A.7: Spectrogram representing a wandering line in the Virgo strain channel, plus other noise features (glitches, lines at 87.1 and 100 Hz, and a bump at 106), together with the tracking of the former (red dashed line). Refer to section 3.4 for the description of the algorithm used to obtain it.

Firstly, we inspected both the Virgo Process Monitoring page and the Virgo logbook in order to check if some changes or interventions were done on the detector in that moment. The only guess that we obtained was related to the refill of some cryotrap for the vacuum system, both in the period of January described before and in others [LVC64]. By itself, this is not a very clear indication, since in other occasions when there were this kind of interventions the line was absent (or not drifting). Also, it is not evident the link between the cryotrap and the line. The second approach that we tried was to study the coherence budget when this line was stationary at 83 Hz and right after it, when it had drifted away, as in the example reported in figure 3.10 and in the study described in [LVC61]. We executed this analysis with Bruco for several similar cases during 2019 but we didn't find any differences in coherence (or coherence at all at 83 Hz) in any of them. As a consequence, the cause that make this line to drift away from 83 Hz is still unknown.

Second aspect that is not clear is weather, once the line has been “unlocked” from its rest position at 83 Hz, it is moved always by the same driving mechanism or by different causes. To make the investigation easier, one would prefer the former case, of course. However, this line has shown in different occasions visibly different evolution patterns; sometimes it was smoother and slower, as in figure 3.10, or in the past cases of thermally driven drifting lines reported in figures 2.12 and 3.9, and in other occasions it was faster and more jagged. If multiple driving causes are indeed the culprit of its evolution, its study by means of the cross-correlation analysis becomes quite difficult since they may overlap or change all the times.

In figure A.7 we reported a final example of this line and the corresponding tracking. We have used the same masking procedure by means of the Hough transform as described in section 3.4 but, differently from figure 3.10, we have preserved the corresponding regions of the spectrogram (glitches, lines at 87.1 and 100 Hz, and part of the bump at 106), not showing them masked. Notice that in the region of the bump at 106 Hz, which appears also

non-stationary from the point of view of its amplitude, the Hough transform wasn't able to entirely mask it, and the consequent line tracking was presumably affected by it. During its descent, the frequency evolution seems more jagged than in the case studied in [LVC61], and not very similar to a drift due to thermal causes. This reinforces the previous point of view that, although the "unlocking" of the line from 83 Hz may come from the same cause, the mechanisms that make it drift may be multiple. The correlation analysis of this more recent occurrence of the line, since September 2019, has been reported in [LVC16]. Different correlated channels have been found all the time, which could be just a coincidence or a manifestation of the previous hypothesis.

What described here represents the current understanding of this noise issue, obtained in about one year of investigations by means of the line tracking and the correlation analysis described in this manuscript. Moreover, almost all the tools and the strategies at disposal of the Detector Characterisation and Commissioning groups were exploited. Some possible causes of it are excluded. For example, there is no clear coupling between a thermometer sensor and it, as it was the case in figure 2.12. Also some hints about what may cause it to abandon its position at 83 Hz, namely the cryotrap refill or the activation of some step motors, have been hypothesised. However, at this point it is reasonable to suspect that the only way to gather further insight about it is through some direct searches, and the placement of some sensors where currently are absent. One hint about where to do that is at the PR-BS link, as many correlation analysis results have seemed to point to.

B Some notable distributions

In this appendix we report some useful notions related to two distribution function we made use of in the discussion of spectral estimations in section 2.1.6 and for the development of the various forms of the Rayleigh test in section 3.3.

B.1 Rayleigh

The Rayleigh distribution is a continuous probability distribution for positive-valued random variables. It is parametrised by a *scale* parameter σ , and it can be shown to describe the distribution of the modulus of two independent and identically distributed, zero-mean Gaussian variables: if $X \sim N(0, \sigma^2)$ and $Y \sim N(0, \sigma^2)$, $R \equiv \sqrt{X^2 + Y^2} \sim Ry(\sigma)$. It is essentially the square root of a chi-squared distribution with two degrees of freedom; if $\sigma = 1$, $R^2 \equiv Q \sim \chi_2^2$.

Its p.d.f., as derived in section 2.1.6, is:

$$p(r) = \frac{r}{\sigma^2} e^{-\frac{r^2}{2\sigma^2}}, \quad \text{for } r \geq 0.$$

Its mean is:

$$\int_0^{+\infty} r p(r) dr = \sqrt{\frac{\pi}{2}} \sigma \simeq 1.25\sigma$$

and rms:

$$\left(\int_0^{+\infty} r^2 p(r) dr \right)^{1/2} = \sqrt{2} \sigma \simeq 1.41\sigma.$$

The percentiles of this distribution can be computed to be:

$$p.p.f.(p; \sigma) = \sigma \sqrt{-2 \log \left(1 - \frac{p}{100} \right)}$$

and in particular the median:

$$p.p.f.(50\%; \sigma) = \sigma \sqrt{2 \log 2} \simeq 1.18\sigma.$$

These are the quantities used to define the Rayleigh test described in section 3.3.

Given a sample of M realisations of a Rayleigh distributed random variable, r_i , the estimator:

$$\hat{\sigma}^2 = \frac{1}{2M} \sum_{i=1}^M r_i^2$$

is the maximum likelihood estimate for the squared distribution parameter σ^2 , and is also unbiased (the estimator $\sqrt{\hat{\sigma}}$ is a biased estimator of σ , instead) [330].

B.2 Gamma

The Gamma distribution is a two-parameter family of continuous probability distributions, including the exponential distribution, and chi-squared distribution as special cases of it.

If R_i are a set of M independent and identically distributed Rayleigh random variables of parameter σ , $G = \sum_{i=1}^M R_i^2$ has a Gamma distribution with *shape parameter* M and *scale parameter* $2\sigma^2$: $G \sim \Gamma(M, 2\sigma^2)$ [331]. For large M the Gamma distribution converges to the normal distribution $N(2M\sigma^2, 4M\sigma^2)$.

Its p.d.f., according to the previous parametrisation on the base of shape k and scale θ , is:

$$p(g) = \frac{1}{\Gamma(k)\theta^k} g^{k-1} e^{-g/\theta}, \quad \text{for } g, \theta \text{ and } k > 0$$

where $\Gamma(k)$ is the *Gamma function*, equals to $(k-1)!$ for every positive integer k .

Its mean and standard deviation can be calculated as:

$$\int_0^{+\infty} g p(g) dg = k\theta$$

and:

$$\left(\int_0^{+\infty} r^2 p(r) dr - (k\theta)^2 \right)^{1/2} = \sqrt{k}\theta.$$

The median does not have an easy closed form equation (and indeed we didn't make use of it for our definitions of the Rayleigh test statistics).

Similarly to what described for the Rayleigh distribution in appendix B.1, given M realisations of the same Γ -distributed variable, g_i , the estimator for the parameter θ (assuming to know k) is:

$$\hat{\theta} = \frac{1}{kM} \sum_{i=1}^M g_i.$$

Bibliography

- [1] B.P. Abbott *et al.* “GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs”. In: *Phys. Rev. X* 9 (3 2019), p. 031040. DOI: 10.1103/PhysRevX.9.031040. URL: <https://link.aps.org/doi/10.1103/PhysRevX.9.031040>.
- [2] LIGO Scientific Collaboration, Virgo Collaboration, et al. “A guide to LIGO-Virgo detector noise and extraction of transient gravitational-wave signals”. In: *arXiv preprint arXiv:1908.11170* (2019).
- [3] J.D.E. Creighton and W.G. Anderson. *Gravitational-Wave Physics and Astronomy: An Introduction to Theory, Experiment and Data Analysis*. Wiley Series in Cosmology. Wiley, 2012. ISBN: 9783527636044. URL: https://books.google.it/books?id=W_TVS_6JYJcC.
- [4] C.W. Misner, K.S. Thorne, and J.A. Wheeler. *Gravitation*. Gravitation pt. 3. W. H. Freeman, 1973. ISBN: 9780716703440. URL: <http://books.google.it/books?id=w4Gigq3tY1kC>.
- [5] P. Jaranowski and A. Krolak. *Analysis of Gravitational-Wave Data*. Cambridge, UK: Cambridge University Press, Aug. 2009. URL: <https://ui.adsabs.harvard.edu/abs/2009agwd.book.....J>.
- [6] M.B. Priestley. *Spectral analysis and time series*. Vol. 1. Academic press London, 1981.
- [7] J. Aasi *et al.* “Searching for stochastic gravitational waves using data from the two colocated LIGO Hanford detectors”. In: *Phys. Rev. D* 91 (2 2015), p. 022003. DOI: 10.1103/PhysRevD.91.022003. URL: <http://link.aps.org/doi/10.1103/PhysRevD.91.022003>.
- [8] B.P. Abbott *et al.* “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Phys. Rev. Lett.* 116 (6 2016), p. 061102. DOI: 10.1103/PhysRevLett.116.061102. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.061102>.
- [9] A. Einstein. “Approximative Integration of the Field Equations of Gravitation”. In: *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* 1916 (1916), pp. 688–696. URL: <https://einsteinpapers.press.princeton.edu/vol6-doc/375>.
- [10] A. Einstein. “On Gravitational Waves”. In: *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* 1918 (1918), pp. 154–167. URL: <https://einsteinpapers.press.princeton.edu/vol7-trans/25>.
- [11] B.P. Abbott *et al.* “Binary black hole mergers in the first advanced LIGO observing run”. In: *Physical Review X* 6.4 (2016), p. 041015.
- [12] F. Acernese *et al.* “Advanced Virgo: a second-generation interferometric gravitational wave detector”. In: *Classical and Quantum Gravity* 32.2, 024001 (2015), p. 024001. DOI: 10.1088/0264-9381/32/2/024001. arXiv: 1408.3978 [gr-qc].
- [13] B.P. Abbott *et al.* “GW170814: a three-detector observation of gravitational waves from a binary black hole coalescence”. In: *Physical review letters* 119.14 (2017), p. 141101.
- [14] B. P. Abbott *et al.* “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. In: *Phys. Rev. Lett.* 119 (16 2017), p. 161101. DOI: 10.1103/PhysRevLett.119.161101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.119.161101>.
- [15] B.P. Abbott *et al.* “Multi-messenger observations of a binary neutron star merger”. In: *Astrophys. J. Lett* 848.2 (2017), p. L12.
- [16] A. Albert *et al.* “Search for High-energy Neutrinos from Binary Neutron Star Merger GW170817 with ANTARES, IceCube, and the Pierre Auger Observatory”. In: 850.2 (2017), p. L35. DOI: 10.3847/2041-8213/aa9aed.

- [17] S.A. Usman *et al.* "The PyCBC search for gravitational waves from compact binary coalescence". In: *Classical and Quantum Gravity* 33.21, 215004 (2016), p. 215004. DOI: 10.1088/0264-9381/33/21/215004. arXiv: 1508.02357 [gr-qc].
- [18] S. Sachdev *et al.* "The GstLAL Search Analysis Methods for Compact Binary Mergers in Advanced LIGO's Second and Advanced Virgo's First Observing Runs". In: *arXiv preprint arXiv:1901.08580* (2019).
- [19] C. Messick *et al.* "Analysis framework for the prompt discovery of compact binary mergers in gravitational-wave data". In: *Physical Review D* 95.4 (2017), p. 042001.
- [20] L.A. Weinstein and V.D. Zubakov. *Extraction of signals from noise*. Prentice-Hall, 1962.
- [21] S. Klimenko *et al.* "Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors". In: *Physical Review D* 93.4 (2016), p. 042004.
- [22] J. Aasi *et al.* "The characterization of Virgo data and its impact on gravitational-wave searches". In: *Classical and Quantum Gravity* 29.15, 155002 (2012), p. 155002. DOI: 10.1088/0264-9381/29/15/155002. arXiv: 1203.5613 [gr-qc].
- [23] B.P. Abbott *et al.* "Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914". In: *Classical and Quantum Gravity* 33.13, 134001 (2016), p. 134001. DOI: 10.1088/0264-9381/33/13/134001. arXiv: 1602.03844 [gr-qc].
- [24] B.P. Abbott *et al.* "Effects of data quality vetoes on a search for compact binary coalescences in Advanced LIGO's first observing run". In: *Classical and Quantum Gravity* 35.6 (2018), p. 065010. ISSN: 1361-6382. DOI: 10.1088/1361-6382/aaaafa. URL: <http://dx.doi.org/10.1088/1361-6382/aaaafa>.
- [25] M. Branchesi. "Multi-messenger astronomy: gravitational waves, neutrinos, photons, and cosmic rays". In: *Journal of Physics: Conference Series*. Vol. 718. 2. IOP Publishing, 2016, p. 022004.
- [26] B.P. Abbott *et al.* "Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA". In: *Living Reviews in Relativity* 21.1, 3 (2018), p. 3. DOI: 10.1007/s41114-018-0012-9. arXiv: 1304.0670 [gr-qc].
- [27] D. Bersanetti. "Status of the Virgo gravitational-wave detector and the O3 Observing Run". In: EPS-HEP2019. URL: <https://indi.co.cern.ch/event/577856/contributions/3422625/>.
- [28] G. Vajente. "Analysis of sensitivity and noise sources for the Virgo gravitational wave interferometer". PhD thesis. Scuola Normale Superiore di Pisa, 2008.
- [29] I. Fiori and Virgo Environmental group. "Environmental noise studies for the Virgo O3 science run". 2020. to be submitted to *Galaxies* (expected in June 2020).
- [30] J.L. Rodgers and W.A. Nicewander. "Thirteen Ways to Look at the Correlation Coefficient". In: *The American Statistician* 42.1 (1988), pp. 59–66. ISSN: 00031305. URL: <http://www.jstor.org/stable/2685263>.
- [31] N. Wiener. "The theory of prediction. Modern mathematics for engineers". In: *New York* (1956), pp. 165–190.
- [32] J. Pearl. *Causality*. Cambridge university press, 2009.
- [33] G. Sugihara *et al.* "Detecting causality in complex ecosystems". In: *science* 338.6106 (2012), pp. 496–500.
- [34] M. Ding, Y. Chen, and S.L. Bressler. "Granger causality: basic theory and application to neuroscience. 2006". In: *arXiv preprint q-bio/0608035* (2006).
- [35] M. Kamiński *et al.* "Evaluating causal relations in neural systems: Granger causality, directed transfer function and statistical assessment of significance". In: *Biological cybernetics* 85.2 (2001), pp. 145–157.
- [36] W. Freiwald *et al.* "Testing non-linearity and directedness of interactions between neural groups in the macaque inferotemporal cortex". In: *Journal of neuroscience methods* 94.1 (1999), pp. 105–119.
- [37] F. Di Renzo. "BRiSTOL: a Band-limited RMS stationarity test for gravitational wave detectors data". 2020. in preparation.
- [38] L. Nuttall *et al.* "Improving the data quality of Advanced LIGO based on early engineering run results". In: *Classical and Quantum Gravity* 32.24, 245005 (Dec. 2015), p. 245005. DOI: 10.1088/0264-9381/32/24/245005. arXiv: 1508.07316 [gr-qc].

- [39] P. Shawhan and E. Ochsner. "A new waveform consistency test for gravitational wave inspiral searches". In: *Classical and Quantum Gravity* 21.20 (2004), S1757. URL: <http://stacks.iop.org/0264-9381/21/i=20/a=018>.
- [40] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer Series in Statistics. New York, NY, USA: Springer New York Inc., 2001.
- [41] J.-M. Lévy-Leblond. "One more derivation of the Lorentz transformation". In: *American Journal of Physics* 44.3 (1976), pp. 271–277.
- [42] G. Cusin *et al.* "First Predictions of the Angular Power Spectrum of the Astrophysical Gravitational Wave Background". In: *Phys. Rev. Lett.* 120 (23 2018), p. 231101. DOI: 10.1103/PhysRevLett.120.231101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.231101>.
- [43] C.M. DeWitt and D. Rickles. *The role of gravitation in physics: Report from the 1957 Chapel Hill Conference*. Vol. 5. epubli, 2011.
- [44] F.A.E. Pirani. "On the Physical significance of the Riemann tensor". In: *Acta Phys. Polon.* 15 (1956). [Gen. Rel. Grav.41,1215(2009)], pp. 389–405. DOI: 10.1007/s10714-009-0787-9.
- [45] H. Bondi. "Plane Gravitational Waves in General Relativity". In: *Nature* 179 (1957), pp. 1072–1073.
- [46] J.L. Cervantes-Cota, S. Galindo-Uribarri, and G.F. Smoot. "A brief history of gravitational waves". In: *Universe* 2.3 (2016), p. 22.
- [47] C.-M. Chen, J.M. Nester, and W.-T. Ni. "A brief history of gravitational wave research". In: *Chinese Journal of Physics* 55.1 (2017), pp. 142–169. ISSN: 0577-9073. DOI: <https://doi.org/10.1016/j.cjph.2016.10.014>. URL: <http://www.sciencedirect.com/science/article/pii/S057790731630572X>.
- [48] D.J. Kennefick. *Traveling at the speed of thought: Einstein and the quest for gravitational waves*. Princeton university press, 2016.
- [49] M. Maggiore. *Gravitational Waves: Theory and Experiments*. Vol. 1. Gravitational Waves. Oxford: Oxford University Press, 2008. ISBN: 9780198570745. URL: <http://ukcatalogue.oup.com/product/9780198570745.do>.
- [50] M. Maggiore. *Gravitational Waves: Astrophysics and Cosmology*. Vol. 2. Oxford University Press, 2018.
- [51] H. Minkowski. "Raum und zeit". In: *Jahresbericht der Deutschen Mathematiker-Vereinigung, vol.18, p.75-88* 18 (1909), pp. 75–88.
- [52] S.M. Carroll. *Spacetime and Geometry: An Introduction to General Relativity*. Addison-Wesley Longman, Incorporated, 2004. ISBN: 9780805387322. URL: <http://books.google.it/books?id=1SKFQgAACAAJ>.
- [53] Y. Choquet-Bruhat. *General Relativity and the Einstein Equations*. Oxford Mathematical Monographs. United Kingdom: Oxford University Press, 2009. ISBN: 0199230722, 9780199230723. URL: <http://www.oup.com/us/catalog/general/subject/Mathematics/AppliedMathematics/?view=usa&ci=9780199230723>.
- [54] K. Schwarzschild. "On the Gravitational Field of a Mass Point According to Einstein's Theory". In: *Abh. Konigl. Preuss. Akad. Wissenschaften Jahre 1906,92, Berlin,1907 1916* (1916), pp. 189–196.
- [55] R.P. Kerr. "Gravitational Field of a Spinning Mass as an Example of Algebraically Special Metrics". In: *Phys. Rev. Lett.* 11 (5 1963), pp. 237–238. DOI: 10.1103/PhysRevLett.11.237. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.11.237>.
- [56] H.P. Robertson. "Kinematics and World-Structure". In: *Astrophysical Journal* 82 (1935), p. 284. DOI: 10.1086/143681.
- [57] H. Stephani *et al.* *Exact solutions of Einstein's field equations*. Cambridge university press, 2009.
- [58] A. Krasinski. *Inhomogeneous Cosmological Models*. 1997.
- [59] M. A.H. MacCallum. "Exact solutions of Einstein's equations". In: *Scholarpedia* 8.12 (2013). revision #138070, p. 8584. DOI: 10.4249/scholarpedia.8584.
- [60] V.A. Belinski and V.E. Zakharov. "Integration of the Einstein equations by the method of the inverse scattering problem and calculation of exact soliton solutions". In: *Zh. Eksp. Teor. Fiz* 75 (1953), p. 1978.
- [61] F.A.E. Pirani. "Invariant Formulation of Gravitational Radiation Theory". In: *Phys. Rev.* 105 (3 1957), pp. 1089–1099. DOI: 10.1103/PhysRev.105.1089. URL: <http://link.aps.org/doi/10.1103/PhysRev.105.1089>.

- [62] R.A. Isaacson. "Gravitational Radiation in the Limit of High Frequency. I. The Linear Approximation and Geometrical Optics". In: *Phys. Rev.* 166 (5 1968), pp. 1263–1271. DOI: 10.1103/PhysRev.166.1263. URL: <http://link.aps.org/doi/10.1103/PhysRev.166.1263>.
- [63] R.A. Isaacson. "Gravitational Radiation in the Limit of High Frequency. II. Nonlinear Terms and the Effective Stress Tensor". In: *Phys. Rev.* 166 (5 1968), pp. 1272–1280. DOI: 10.1103/PhysRev.166.1272. URL: <http://link.aps.org/doi/10.1103/PhysRev.166.1272>.
- [64] É.É. Flanagan and S.A. Hughes. "The basics of gravitational wave theory". In: *New Journal of Physics* 7.1 (2005), p. 204.
- [65] J. Stewart. *Advanced General Relativity*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 1993. ISBN: 9780521449465. URL: <http://books.google.it/books?id=LZrjJ2P1J1sC>.
- [66] J.D. Jackson. *Classical electrodynamics*. 3rd ed. Wiley, 1999. ISBN: 9780471309321.
- [67] A. Blaut. "Angular and frequency response of the gravitational wave interferometers in the metric theories of gravity". In: *Phys. Rev. D* 85 (4 2012), p. 043005. DOI: 10.1103/PhysRevD.85.043005. URL: <http://link.aps.org/doi/10.1103/PhysRevD.85.043005>.
- [68] A. Blaut. "Gauge independent response of a laser interferometer to gravitational waves". In: *Classical and Quantum Gravity* 36.5 (2019), p. 055004.
- [69] M.J. Koop and L.S. Finn. "Physical response of light-time gravitational wave detectors". In: *Phys. Rev. D* 90 (6 2014), p. 062002. DOI: 10.1103/PhysRevD.90.062002. URL: <http://link.aps.org/doi/10.1103/PhysRevD.90.062002>.
- [70] B.P. Abbott *et al.* "GW170817: Measurements of neutron star radii and equation of state". In: *Physical review letters* 121.16 (2018), p. 161101.
- [71] K. Danzmann for the LISA Study Team. "LISA - an ESA cornerstone mission for a gravitational wave observatory". In: *Classical and Quantum Gravity* 14.6 (1997), p. 1399. URL: <http://stacks.iop.org/0264-9381/14/i=6/a=002>.
- [72] G. Hobbs and S. Dai. "Gravitational wave research using pulsar timing arrays". In: *National Science Review* 4.5 (2017), pp. 707–717.
- [73] C.J. Moore, R.H. Cole, and C.P.L. Berry. "Gravitational-wave sensitivity curves". In: *Classical and Quantum Gravity* 32.1 (2014), p. 015014. ISSN: 1361-6382. DOI: 10.1088/0264-9381/32/1/015014. URL: <http://dx.doi.org/10.1088/0264-9381/32/1/015014>.
- [74] B. Allen. "The Stochastic gravity wave background: Sources and detection". In: *In Les Houches 1995, Relativistic gravitation and gravitational radiation 373-417* (1997). Ed. by J.-A. Marck and J.-P. Lasota, p. 373. eprint: [gr-qc/9604033](http://adsabs.harvard.edu/abs/1997rggr.conf..373A). URL: <http://adsabs.harvard.edu/abs/1997rggr.conf..373A>.
- [75] M. Maggiore. "Gravitational wave experiments and early universe cosmology". In: *Physics Reports* 331.6 (2000), pp. 283–367. ISSN: 0370-1573. DOI: 10.1016/S0370-1573(99)00102-7. URL: <http://www.sciencedirect.com/science/article/pii/S0370157399001027>.
- [76] C. Caprini and D.G. Figueroa. "Cosmological backgrounds of gravitational waves". In: *Classical and Quantum Gravity* 35.16 (2018), p. 163001.
- [77] K.A. Postnov. "Astrophysical sources of stochastic gravitational radiation in the universe". In: *arXiv preprint astro-ph/9706053* (1997).
- [78] T. Regimbau. "The astrophysical gravitational wave stochastic background". In: *Research in Astronomy and Astrophysics* 11 (2011), pp. 369–390. DOI: 10.1088/1674-4527/11/4/001. arXiv: 1101.2762 [astro-ph.CO].
- [79] S.L. Shapiro and S.A. Teukolsky. *Black holes, white dwarfs, and neutron stars: The physics of compact objects*. John Wiley & Sons, 2008.
- [80] B.S. Sathyaprakash and B.F. Schutz. "Physics, Astrophysics and Cosmology with Gravitational Waves". In: *Living Reviews in Relativity* 12.1 (2009), p. 2. ISSN: 1433-8351. DOI: 10.12942/lrr-2009-2. URL: <https://doi.org/10.12942/lrr-2009-2>.
- [81] L. Blanchet. "Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries". In: *Living Reviews in Relativity* 17.1 (2014), p. 2. ISSN: 1433-8351. DOI: 10.12942/lrr-2014-2. URL: <https://doi.org/10.12942/lrr-2014-2>.

- [82] E. Poisson and C.M. Will. *Gravity: Newtonian, post-newtonian, relativistic*. Cambridge University Press, 2014.
- [83] F. Pretorius. *Physics of Relativistic Objects in Compact Binaries: From Birth to Coalescence*. 2009.
- [84] K.D. Kokkotas and B.G. Schmidt. "Quasi-Normal Modes of Stars and Black Holes". In: *Living Reviews in Relativity* 2.1 (1999). ISSN: 1433-8351. DOI: 10.12942/lrr-1999-2. URL: <http://dx.doi.org/10.12942/lrr-1999-2>.
- [85] T. Damour and A. Nagar. "The Effective-One-Body Approach to the General Relativistic Two Body Problem". In: *Astrophysical Black Holes*. Ed. by V. Haardt and Gorini et al. Cham: Springer International Publishing, 2016, pp. 273–312. ISBN: 978-3-319-19416-5. DOI: 10.1007/978-3-319-19416-5_7. URL: https://doi.org/10.1007/978-3-319-19416-5_7.
- [86] J.S. Read et al. "Measuring the neutron star equation of state with gravitational wave observations". In: *Physical Review D* 79.12 (2009), p. 124033.
- [87] J. Clark et al. "Prospects for joint gravitational wave and short gamma-ray burst observations". In: *Astrophys. J.* 809.1 (2015), p. 53. DOI: 10.1088/0004-637X/809/1/53. arXiv: 1409.8149 [astro-ph.HE].
- [88] B.P. Abbott et al. "Properties of the binary neutron star merger GW170817". In: *Phys. Rev. X* 9.1 (2019), p. 011001. DOI: 10.1103/PhysRevX.9.011001. arXiv: 1805.11579 [gr-qc].
- [89] B.P. Abbott et al. "Tests of General Relativity with GW170817". In: *Phys. Rev. Lett.* 123 (1 2019), p. 011102. DOI: 10.1103/PhysRevLett.123.011102. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.011102>.
- [90] A. Loeb. "Electromagnetic counterparts to black hole mergers detected by LIGO". In: *The Astrophysical Journal Letters* 819.2 (2016), p. L21.
- [91] P.A. Mazzali et al. "A common explosion mechanism for type Ia supernovae". In: *Science* 315.5813 (2007), pp. 825–828.
- [92] D.E. Holz, S.A. Hughes, and B.F. Schutz. "Measuring cosmic distances with standard sirens". In: *Physics Today* 71.12 (2018), pp. 34–40. DOI: 10.1063/PT.3.4090.
- [93] S. Nissanke et al. "Determining the Hubble constant from gravitational wave observations of merging compact binaries". In: *arXiv preprint arXiv:1307.2638* (2013).
- [94] LIGO Scientific Collaboration et al. "A gravitational-wave standard siren measurement of the Hubble constant". In: *Nature* 551.7678 (2017), pp. 85–88.
- [95] M. Bassan, ed. *Advanced Interferometers and the Search for Gravitational Waves*. Vol. 404. Astrophysics and Space Science Library. 2014. DOI: 10.1007/978-3-319-03792-9.
- [96] K. Belczynski et al. "On the rarity of double black hole binaries: Consequences for gravitational wave detection". In: *The Astrophysical Journal* 662.1 (2007), p. 504.
- [97] B.P. Abbott et al. "Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo". In: *The Astrophysical Journal* 882.2 (2019), p. L24. DOI: 10.3847/2041-8213/ab3800. URL: <https://doi.org/10.3847/2041-8213/ab3800>.
- [98] J. Abadie et al. "TOPICAL REVIEW: Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors". In: *Classical and Quantum Gravity* 27.17, 173001 (2010), p. 173001. DOI: 10.1088/0264-9381/27/17/173001. arXiv: 1003.2480 [astro-ph.HE].
- [99] J. Abadie et al. "Search for Gravitational Waves from Low Mass Compact Binary Coalescence in LIGO's Sixth Science Run and Virgo's Science Runs 2 and 3". In: *Phys. Rev. D* 85 (2012), p. 082002. DOI: 10.1103/PhysRevD.85.082002. arXiv: 1111.7314 [gr-qc].
- [100] C.J. Horowitz and K. Kadau. "Breaking strain of neutron star crust and gravitational waves". In: *Physical Review Letters* 102.19 (2009), p. 191102.
- [101] N. Chamel and P. Haensel. "Physics of neutron star crusts". In: *Living Reviews in Relativity* 11.1 (2008), p. 10.
- [102] D.I. Jones. "Gravitational waves from rotating strained neutron stars". In: *Classical and Quantum Gravity* 19.7 (2002), p. 1255.

- [103] C. Cutler. "Gravitational waves from neutron stars with large toroidal B fields". In: *Physical Review D* 66.8 (2002), p. 084025.
- [104] N. Andersson *et al.* "Gravitational waves from neutron stars: promises and challenges". In: *General Relativity and Gravitation* 43.2 (2011), pp. 409–436.
- [105] B.P. Abbott *et al.* "Narrow-band search for gravitational waves from known pulsars using the second LIGO observing run". In: *Physical Review D* 99.12 (2019), p. 122002.
- [106] B.P. Abbott *et al.* "Searches for gravitational waves from known pulsars at two harmonics in 2015–2017 LIGO data". In: *The Astrophysical Journal* 879.1 (2019), p. 10.
- [107] B.P. Abbott *et al.* "All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO O2 data". In: *Physical Review D* 100.2 (2019), p. 024004.
- [108] P. Jaranowski, A. Królak, and B.F. Schutz. "Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection". In: *Phys. Rev. D* 58 (6 1998), p. 063001. DOI: 10 . 1103 / PhysRevD . 58 . 063001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.58.063001>.
- [109] M. Dine, W. Fischler, and M. Srednicki. "A simple solution to the strong CP problem with a harmless axion". In: *Physics letters B* 104.3 (1981), pp. 199–202.
- [110] A. Arvanitaki, M. Baryakhtar, and X. Huang. "Discovering the QCD axion with black holes and gravitational waves". In: *Phys. Rev. D* 91 (8 2015), p. 084011. DOI: 10 . 1103 / PhysRevD . 91 . 084011. URL: <https://link.aps.org/doi/10.1103/PhysRevD.91.084011>.
- [111] N. Siemonsen and W.E. East. "Gravitational wave signatures of ultralight vector bosons from black hole superradiance". In: *arXiv preprint arXiv:1910.09476* (2019).
- [112] R. Brito *et al.* "Gravitational wave searches for ultralight bosons with LIGO and LISA". In: *Physical Review D* 96.6 (2017). ISSN: 2470-0029. DOI: 10 . 1103 / physrevd . 96 . 064050. URL: <http://dx.doi.org/10.1103/PhysRevD.96.064050>.
- [113] H.-T. Janka *et al.* "Theory of core-collapse supernovae". In: *Physics Reports* 442.1-6 (2007), pp. 38–74.
- [114] C.L. Fryer and K.C.B. New. "Gravitational waves from gravitational collapse". In: *Living Reviews in Relativity* 14.1 (2011), p. 1.
- [115] H. Dimmelmeier, J.A. Font, and E. Müller. "Relativistic simulations of rotational core collapse II. Collapse dynamics and gravitational radiation". In: *Astronomy & Astrophysics* 393.2 (2002), pp. 523–542.
- [116] R. Diehl *et al.* "Radioactive ^{26}Al from massive stars in the Galaxy". In: *Nature* 439.7072 (2006), p. 45.
- [117] B.P. Abbott *et al.* *An Optically Targeted Search for Gravitational Waves emitted by Core-Collapse Supernovae during the First and Second Observing Runs of Advanced LIGO and Advanced Virgo*. 2019. arXiv: 1908 . 03584 [astro-ph.HE].
- [118] T. Zwerger and E. Mueller. "Dynamics and gravitational wave signature of axisymmetric rotational core collapse." In: *Astronomy and Astrophysics* 320 (1997), pp. 209–227.
- [119] H. Dimmelmeier *et al.* "Gravitational wave burst signal from core collapse of rotating stars". In: *Physical Review D* 78.6 (2008), p. 064056.
- [120] A. Vilenkin and E.P.S. Shellard. *Cosmic strings and other topological defects*. Cambridge University Press, 2000.
- [121] B.P. Abbott *et al.* "Constraints on cosmic strings using data from the first Advanced LIGO observing run". In: *Physical Review D* 97.10 (2018), p. 102002.
- [122] T. Damour and A. Vilenkin. "Gravitational radiation from cosmic (super)strings: Bursts, stochastic background, and observational windows". In: *Phys. Rev. D* 71 (6 2005), p. 063510. DOI: 10 . 1103 / PhysRevD . 71 . 063510. URL: <https://link.aps.org/doi/10.1103/PhysRevD.71.063510>.
- [123] N.J. Cornish and J.D. Romano. "When is a gravitational-wave signal stochastic?" In: *Phys. Rev. D* 92 (4 2015), p. 042001. DOI: 10 . 1103 / PhysRevD . 92 . 042001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.92.042001>.
- [124] B.P. Abbott *et al.* "GW150914: Implications for the Stochastic Gravitational-Wave Background from Binary Black Holes". In: *Phys. Rev. Lett.* 116 (13 2016), p. 131102. DOI: 10 . 1103 / PhysRevLett . 116 . 131102. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.131102>.

- [125] B.P. Abbott *et al.* "GW170817: Implications for the Stochastic Gravitational-Wave Background from Compact Binary Coalescences". In: *Phys. Rev. Lett.* 120 (9 2018), p. 091101. DOI: 10.1103/PhysRevLett.120.091101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.091101>.
- [126] T. Regimbau. "Cosmic background of gravitational waves from rotating neutron stars". In: *Astronomy & Astrophysics* 376.2 (2001), pp. 381–385.
- [127] T. Regimbau. "Gravitational wave background from magnetars". In: *Astronomy & Astrophysics* 447.1 (2006), pp. 1–7.
- [128] G. Nelemans, L.R. Yungelson, and S.F.P. Zwart. "The gravitational wave signal from the Galactic disk population of binaries containing two compact objects". In: *Astronomy & Astrophysics* 375.3 (2001), pp. 890–898.
- [129] A.J. Farmer and E.S. Phinney. "The gravitational wave background from cosmological compact binaries". In: *Monthly Notices of the Royal Astronomical Society* 346.4 (2003), pp. 1197–1214.
- [130] M.S. Turner. "Detectability of inflation-produced gravitational waves". In: *Phys. Rev. D* 55 (2 1997), R435–R439. DOI: 10.1103/PhysRevD.55.R435. URL: <http://link.aps.org/doi/10.1103/PhysRevD.55.R435>.
- [131] R. Easther, J.T. Giblin, and E.A. Lim. "Gravitational Waves From the End of Inflation: Computational Strategies". In: *Phys.Rev. D* 77 (2008), p. 103519. DOI: 10.1103/PhysRevD.77.103519. arXiv: 0712.2991 [astro-ph].
- [132] M. Hindmarsh *et al.* "Gravitational waves from the sound of a first order phase transition". In: *Physical review letters* 112.4 (2014), p. 041301.
- [133] A. Kosowsky, M.S Turner, and R. Watkins. "Gravitational waves from first-order cosmological phase transitions". In: *Physical review letters* 69.14 (1992), p. 2026.
- [134] X. Siemens, V. Mandic, and J. Creighton. "Gravitational-Wave Stochastic Background from Cosmic Strings". In: *Phys. Rev. Lett.* 98 (11 2007), p. 111101. DOI: 10.1103/PhysRevLett.98.111101. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.98.111101>.
- [135] S. Sarangi and S.-H.H. Tye. "Cosmic string production towards the end of brane inflation". In: *Physics Letters B* 536.3 (2002), pp. 185–192.
- [136] N. Christensen. "Stochastic gravitational wave backgrounds". In: *Reports on Progress in Physics* 82.1 (2018), p. 016903.
- [137] T.L. Smith, E. Pierpaoli, and M. Kamionkowski. "New Cosmic Microwave Background Constraint to Primordial Gravitational Waves". In: *Phys. Rev. Lett.* 97 (2 2006), p. 021301. DOI: 10.1103/PhysRevLett.97.021301. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.97.021301>.
- [138] B.P. Abbott *et al.* "Search for the isotropic stochastic background using data from Advanced LIGO's second observing run". In: *Phys. Rev. D* 100 (6 2019), p. 061101. DOI: 10.1103/PhysRevD.100.061101. URL: <https://link.aps.org/doi/10.1103/PhysRevD.100.061101>.
- [139] B.P. Abbott *et al.* "Directional limits on persistent gravitational waves using data from Advanced LIGO's first two observing runs". In: *Phys. Rev. D* 100 (6 2019), p. 062001. DOI: 10.1103/PhysRevD.100.062001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.100.062001>.
- [140] R. Weiss and D. Muehlner. "Electronically coupled broadband gravitational antenna". In: *Quarterly Progress Report, Research Laboratory of Electronics* 105 (1972).
- [141] M.E. Gertsenshtein and V.I. Pustovoit. "ZhETF, 43 (1962), 605". In: *Sov. Phys. JETP* 16 (1963), p. 433.
- [142] P.R. Saulson. *Fundamentals of interferometric gravitational wave detectors*. World Scientific, 1994.
- [143] M. Punturo *et al.* "The Einstein Telescope: a third-generation gravitational wave observatory". In: *Classical and Quantum Gravity* 27.19 (2010), p. 194002. URL: <http://stacks.iop.org/0264-9381/27/i=19/a=194002>.
- [144] C. Bond *et al.* "Interferometer techniques for gravitational-wave detection". In: *Living Reviews in Relativity* 19.1 (2017), p. 3. ISSN: 1433-8351. DOI: 10.1007/s41114-016-0002-8. URL: <https://doi.org/10.1007/s41114-016-0002-8>.

- [145] G. Hammond, S. Hild, and M. Pitkin. "Advanced technologies for future ground-based, laser-interferometric gravitational wave detectors". In: *Journal of Modern Optics* 61.sup1 (2014), S10–S45. ISSN: 1362-3044. DOI: 10.1080/09500340.2014.920934. URL: <http://dx.doi.org/10.1080/09500340.2014.920934>.
- [146] S. Detweiler. "Pulsar timing measurements and the search for gravitational waves". In: *Am. Phys. Journ.* 234 (Dec. 1979), pp. 1100–1104. DOI: 10.1086/157593.
- [147] G. Hobbs *et al.* "The international pulsar timing array project: using pulsars as a gravitational wave detector". In: *Classical and Quantum Gravity* 27.8 (2010), p. 084013.
- [148] E. Mauceli *et al.* "The Allegro gravitational wave detector: Data acquisition and analysis". In: *Physical Review D* 54.2 (1996), p. 1264.
- [149] M. Cerdonio *et al.* "The ultracryogenic gravitational-wave detector AURIGA". In: *Classical and Quantum Gravity* 14.6 (1997), p. 1491.
- [150] P. Astone *et al.* "The gravitational wave detector NAUTILUS operating at T= 0.1 K". In: *Astroparticle Physics* 7.3 (1997), pp. 231–243.
- [151] P. Astone *et al.* "The EXPLORER gravitational wave antenna: recent improvements and performances". In: *Classical and Quantum Gravity* 19.7 (2002), p. 1905.
- [152] I.S. Heng *et al.* "Long term operation of a niobium resonant bar gravitational wave antenna". In: *Physics Letters A* 218.3-6 (1996), pp. 190–196.
- [153] J. Weber. "Gravitational-Wave-Detector Events". In: *Phys. Rev. Lett.* 20 (23 1968), pp. 1307–1308. DOI: 10.1103/PhysRevLett.20.1307. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.20.1307>.
- [154] R.W.P. Drever *et al.* "Laser phase and frequency stabilization using an optical resonator". In: *Applied Physics B* 31.2 (1983), pp. 97–105.
- [155] E.D. Black and R.N. Gutenkunst. "An introduction to signal extraction in interferometric gravitational wave detectors". In: *American Journal of Physics* 71.4 (2003), pp. 365–378.
- [156] T.T. Fricke *et al.* "DC readout experiment in Enhanced LIGO". In: *Classical and Quantum Gravity* 29.6 (2012), p. 065005. URL: <http://stacks.iop.org/0264-9381/29/i=6/a=065005>.
- [157] P. Fritschel *et al.* "Arm cavity finesse for Advanced LIGO". In: *LIGO Document* 70303 (2007), p. 2007.
- [158] B.J. Meers. "Recycling in laser-interferometric gravitational-wave detectors". In: *Physical Review D* 38.8 (1988), p. 2317.
- [159] J. Abadie *et al.* "A gravitational wave observatory operating beyond the quantum shot-noise limit". In: *Nature Physics* 7.12 (2011), p. 962.
- [160] J. Aasi *et al.* "Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light". In: *Nature Photonics* 7.8 (2013), p. 613.
- [161] H. Grote *et al.* "First Long-Term Application of Squeezed States of Light in a Gravitational-Wave Observatory". In: *Phys. Rev. Lett.* 110 (18 2013), p. 181101. DOI: 10.1103/PhysRevLett.110.181101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.110.181101>.
- [162] M. Tse *et al.* "Quantum-Enhanced Advanced LIGO Detectors in the Era of Gravitational-Wave Astronomy". In: *Phys. Rev. Lett.* 123 (23 2019), p. 231107. DOI: 10.1103/PhysRevLett.123.231107. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.231107>.
- [163] F. Acernese *et al.* "Increasing the Astrophysical Reach of the Advanced Virgo Detector via the Application of Squeezed Vacuum States of Light". In: *Phys. Rev. Lett.* 123 (23 2019), p. 231108. DOI: 10.1103/PhysRevLett.123.231108. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.231108>.
- [164] H.J. Kimble *et al.* "Conversion of conventional gravitational-wave interferometers into quantum non-demolition interferometers by modifying their input and/or output optics". In: *Phys. Rev. D* 65 (2 2001), p. 022002. DOI: 10.1103/PhysRevD.65.022002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.65.022002>.
- [165] M. Evans *et al.* "Realistic filter cavities for advanced gravitational wave detectors". In: *Phys. Rev. D* 88 (2 2013), p. 022002. DOI: 10.1103/PhysRevD.88.022002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.88.022002>.

- [166] T. Akutsu *et al.* "Construction of KAGRA: an underground gravitational-wave observatory". In: *Progress of Theoretical and Experimental Physics* 2018.1 (Jan. 2018). 013F01. ISSN: 2050-3911. DOI: 10.1093/ptep/ptx180. eprint: <http://oup.prod.sis.lan/ptep/article-pdf/2018/1/013F01/23570266/ptx180.pdf>. URL: <https://doi.org/10.1093/ptep/ptx180>.
- [167] D. Reitze *et al.* "Cosmic explorer: the US Contribution to gravitational-wave astronomy beyond LIGO". In: *arXiv preprint arXiv:1907.04833* (2019).
- [168] T. Accadia *et al.* "Reconstruction of the gravitational wave signal $h(t)$ during the Virgo science runs and independent validation with a photon calibrator". In: *Classical and Quantum Gravity* 31.16, 165013 (2014), p. 165013. DOI: 10.1088/0264-9381/31/16/165013. arXiv: 1401.6066 [gr-qc].
- [169] F. Acernese *et al.* "Calibration of Advanced Virgo and Reconstruction of the Gravitational Wave Signal $h(t)$ during the Observing Run O2". In: *ArXiv e-prints* (2018). arXiv: 1807.03275 [gr-qc].
- [170] L.S. Finn and D.F. Chernoff. "Observing binary inspiral in gravitational radiation: One interferometer". In: *Physical Review D* 47.6 (1993), p. 2198.
- [171] C. Cutler and É.E. Flanagan. "Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral waveform?" In: *Phys. Rev. D* 49 (6 1994), pp. 2658–2697. DOI: 10.1103/PhysRevD.49.2658. URL: <http://link.aps.org/doi/10.1103/PhysRevD.49.2658>.
- [172] B. Iyer and IndIGO Consortium *et al.* "LIGO-India". In: *Report No. LIGO-M1100296 (Indian Initiative in Gravitational-wave Observations, 2011)* (2011).
- [173] P. Saulson. "If light waves are stretched by gravitational waves, how can we use light as a ruler to detect gravitational waves?" In: *American Journal of Physics* 65 (1997), pp. 501–505. DOI: 10.1119/1.18578.
- [174] B.F. Schutz. "Networks of gravitational wave detectors and three figures of merit". In: *Classical and Quantum Gravity* 28.12 (2011), p. 125023.
- [175] T.A. Callister *et al.* "Observing Gravitational Waves with a Single Detector". In: *Class. Quant. Grav.* 34.15 (2017), p. 155007. DOI: 10.1088/1361-6382/aa7a76. arXiv: 1704.00818 [astro-ph, IM].
- [176] E. Capocasa *et al.* "Estimation of losses in a 300 m filter cavity and quantum noise reduction in the KAGRA gravitational-wave detector". In: *Physical Review D* 93.8 (2016), p. 082004.
- [177] N. Aggarwal *et al.* "Exploring the sensitivity of next generation gravitational wave detectors". In: (2017).
- [178] LIGO Scientific Collaboration *et al.* "Gravitational wave astronomy with LIGO and similar detectors in the next decade". In: *arXiv preprint arXiv:1904.03187* (2019).
- [179] P. Amaro-Seoane *et al.* "eLISA/NGO: Astrophysics and cosmology in the gravitational-wave millihertz regime". In: *GW Notes* 6 (2013), pp. 4–110. arXiv: 1201.3621 [astro-ph, CO].
- [180] A. Sesana. "Prospects for multiband gravitational-wave astronomy after GW150914". In: *Physical Review Letters* 116.23 (2016), p. 231102.
- [181] S. Vitale. "Multiband gravitational-wave astronomy: parameter estimation and tests of general relativity with space-and ground-based detectors". In: *Physical review letters* 117.5 (2016), p. 051102.
- [182] E. Barausse, N. Yunes, and K. Chamberlain. "Theory-agnostic constraints on black-hole dipole radiation with multiband gravitational-wave astrophysics". In: *Physical review letters* 116.24 (2016), p. 241104.
- [183] The Royal Swedish Academy of Sciences. *Press release: The Nobel Prize in Physics 2017*. 2017. URL: <https://www.nobelprize.org/prizes/physics/2017/press-release/>.
- [184] B.P. Abbott *et al.* "GW150914: The Advanced LIGO detectors in the era of first discoveries". In: *Physical review letters* 116.13 (2016), p. 131103.
- [185] B.P. Abbott *et al.* "Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914". In: *Physical Review D* 95.6, 062003 (2017), p. 062003. DOI: 10.1103/PhysRevD.95.062003. arXiv: 1602.03845 [gr-qc].
- [186] B.P. Abbott *et al.* "GW150914: first results from the search for binary black hole coalescence with Advanced LIGO. *Physical Review D*, 93 (12). ISSN 1550-2368". In: *PHYSICAL REVIEW D Phys Rev D* 93 (2016), p. 122003.

- [187] B.P. Abbott *et al.* “ASTROPHYSICAL IMPLICATIONS OF THE BINARY BLACK HOLE MERGER GW150914”. In: *The Astrophysical Journal* 818.2 (2016), p. L22. DOI: 10.3847/2041-8205/818/2/L22. URL: <https://doi.org/10.3847/2041-8205/818/2/L22>.
- [188] B.P. Abbott *et al.* “Tests of general relativity with GW150914”. In: *arXiv preprint arXiv:1602.03841* (2016).
- [189] M. Pürrer. “Frequency domain reduced order model of aligned-spin effective-one-body waveforms with generic mass ratios and spins”. In: *Physical Review D* 93.6 (2016). ISSN: 2470-0029. DOI: 10.1103/PhysRevD.93.064041. URL: <http://dx.doi.org/10.1103/PhysRevD.93.064041>.
- [190] S. Khan *et al.* “Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era”. In: *Physical Review D* 93.4 (2016). ISSN: 2470-0029. DOI: 10.1103/PhysRevD.93.044007. URL: <http://dx.doi.org/10.1103/PhysRevD.93.044007>.
- [191] B.P. Abbott *et al.* “Properties of the Binary Black Hole Merger GW150914”. In: *Phys. Rev. Lett.* 116 (24 2016), p. 241102. DOI: 10.1103/PhysRevLett.116.241102. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.241102>.
- [192] B.P. Abbott *et al.* “GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence”. In: *Phys. Rev. Lett.* 116 (24 2016), p. 241103. DOI: 10.1103/PhysRevLett.116.241103. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.241103>.
- [193] A.H. Nitz *et al.* “1-OGC: The First Open Gravitational-wave Catalog of Binary Mergers from Analysis of Public Advanced LIGO Data”. In: *The Astrophysical Journal* 872.2 (2019), p. 195. ISSN: 1538-4357. DOI: 10.3847/1538-4357/ab0108. URL: <http://dx.doi.org/10.3847/1538-4357/ab0108>.
- [194] W.M. Farr *et al.* “Distinguishing spin-aligned and isotropic black hole populations with gravitational waves”. In: *Nature* 548.7667 (2017), pp. 426–429. DOI: 10.1038/nature23453. arXiv: 1706.01385 [astro-ph.HE].
- [195] N.J. Cornish and T.B. Littenberg. “Bayeswave: Bayesian inference for gravitational wave bursts and instrument glitches”. In: *Classical and Quantum Gravity* 32.13 (2015), p. 135012. DOI: 10.1088/0264-9381/32/13/135012. URL: <https://doi.org/10.1088/0264-9381/32/13/135012>.
- [196] B.P. Abbott *et al.* “Observing gravitational-wave transient GW150914 with minimal assumptions”. In: *Physical Review D* 93.12 (2016). ISSN: 2470-0029. DOI: 10.1103/PhysRevD.93.122004. URL: <http://dx.doi.org/10.1103/PhysRevD.93.122004>.
- [197] R. Lynch *et al.* “Information-theoretic approach to the gravitational-wave burst detection problem”. In: *Phys. Rev. D* 95 (10 2017), p. 104046. DOI: 10.1103/PhysRevD.95.104046. URL: <https://link.aps.org/doi/10.1103/PhysRevD.95.104046>.
- [198] C. Pankow *et al.* “Mitigation of the instrumental noise transient in gravitational-wave data surrounding GW170817”. In: *ArXiv e-prints* (Aug. 2018). arXiv: 1808.03619 [gr-qc].
- [199] A.H. Prestwich *et al.* “The Orbital Period of the Wolf-Rayet Binary IC 10 X-1: Dynamic Evidence that the Compact Object Is a Black Hole”. In: *The Astrophysical Journal* 669.1 (2007), L21–L24. ISSN: 1538-4357. DOI: 10.1086/523755. URL: <http://dx.doi.org/10.1086/523755>.
- [200] B.P. Abbott *et al.* “GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2”. In: *Phys. Rev. Lett.* 118.22, 221101 (June 2017), p. 221101. DOI: 10.1103/PhysRevLett.118.221101. arXiv: 1706.01812 [gr-qc].
- [201] B.P. Abbott *et al.* “Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A”. In: *The Astrophysical Journal* 848.2 (2017), p. L13. DOI: 10.3847/2041-8213/aa920c. URL: <https://doi.org/10.3847/2041-8213/aa920c>.
- [202] LIGO Laboratory Gravitational Wave Open Science Center (GWOSC) and LIGO Scientific Collaboration. *The LIGO Laboratory’s Data Management Plan*, <https://dcc.ligo.org/public/0009/M1000066/025/LIG0-M1000066-v25.pdf>.
- [203] B. Zackay *et al.* “A Highly Spinning and Aligned Binary Black Hole Merger in the Advanced LIGO First Observing Run”. In: *arXiv preprint arXiv:1902.10331* (2019).
- [204] T. Venumadhav *et al.* “New Binary Black Hole Mergers in the Second Observing Run of Advanced LIGO and Advanced Virgo”. In: *arXiv preprint arXiv:1904.07214* (2019).
- [205] M. Cabero *et al.* “Blip glitches in Advanced LIGO data”. In: *Classical and Quantum Gravity* 36.15, 155010 (2019), p. 155010. DOI: 10.1088/1361-6382/ab2e14. arXiv: 1901.05093 [physics.ins-det].

- [206] D.R. Cox and H.D. Miller. *The Theory of Stochastic Processes*. Methuen's monographs on applied probability and statistics. Taylor & Francis, 1977. ISBN: 9780412151705.
- [207] J. Lamperti. *Stochastic processes: a survey of the mathematical theory*. Applied mathematical sciences. Springer-Verlag, 1977. ISBN: 9780387902753.
- [208] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. Communications and signal processing. McGraw-Hill, 1991. ISBN: 9780070484771. URL: <https://books.google.it/books?id=4IwQAQAIAAJ>.
- [209] W. Rudin. *Real and complex analysis*. Tata McGraw-hill education, 2006.
- [210] T.-J. Stieltjes. "Recherches sur les fractions continues". In: *Annales de la Faculté des sciences de Toulouse: Mathématiques*. Vol. 8. 4. 1894, J1–J122.
- [211] D.E. Myers. "To be or not to be... stationary? That is the question". In: *Mathematical Geology* 21.3 (1989), pp. 347–362.
- [212] A. G. Miamiee and M. Pourahmadi. "Wold decomposition, prediction and parameterization of stationary processes with infinite variance". In: *Probability Theory and Related Fields* 79.1 (1988), pp. 145–164. ISSN: 1432-2064. DOI: 10.1007/BF00319110. URL: <https://doi.org/10.1007/BF00319110>.
- [213] J.S. Bendat and A.G. Piersol. *Random Data: Analysis and Measurement Procedures*. Wiley Series in Probability and Statistics. Wiley, 2011. ISBN: 9781118210826. URL: <https://books.google.it/books?id=qYSViFRNMLwC>.
- [214] P. Masani *et al.* "Wiener's contributions to generalized harmonic analysis, prediction theory and filter theory". In: *Bulletin of the American Mathematical Society* 72 (1966), pp. 73–125.
- [215] H. Wold. "A study in the analysis of stationary time series". PhD thesis. Almqvist & Wiksell, 1938.
- [216] W. Rudin. *Fourier analysis on groups*. Vol. 121967. Wiley Online Library, 1962.
- [217] A. Papoulis. *Signal analysis*. McGraw-Hill electrical and electronic engineering series. McGraw-Hill, 1977. ISBN: 9780070484603. URL: <http://books.google.it/books?id=Re5SAAAAAAAJ>.
- [218] C.B. Mehr and J.A. McFadden. "Certain Properties of Gaussian Processes and Their First-Passage Times". In: *Journal of the Royal Statistical Society: Series B (Methodological)* 27.3 (1965), pp. 505–522.
- [219] J.L. Doob. *Stochastic processes*. Vol. 101. New York Wiley, 1953.
- [220] E. Cuoco *et al.* "Noise parametric identification and whitening for LIGO 40-m interferometer data". In: *Physical Review D* 64.12 (2001), p. 122002.
- [221] E. Cuoco *et al.* "On-line power spectra identification and whitening for the noise in interferometric gravitational wave detectors". In: *Classical and Quantum Gravity* 18.9 (2001), p. 1727.
- [222] E. Cuoco, G. Cella, and G.M. Guidi. "Whitening of non-stationary noise from gravitational wave detectors". In: *Classical and quantum gravity* 21.5 (2004), S801.
- [223] M.B. Priestley. "Evolutionary spectra and non-stationary processes". In: *Journal of the Royal Statistical Society: Series B (Methodological)* 27.2 (1965), pp. 204–229.
- [224] R. Silverman. "Locally stationary random processes". In: *IRE Transactions on Information Theory* 3.3 (1957), pp. 182–187. DOI: 10.1109/TIT.1957.1057413.
- [225] R. Dahlhaus. "Asymptotic statistical inference for nonstationary processes with evolutionary spectra". In: *Athens conference on applied probability and time series analysis*. Springer, 1996, pp. 145–159.
- [226] R. Dahlhaus. "Locally Stationary Processes". In: *arXiv e-prints*, arXiv:1109.4174 (2011), arXiv:1109.4174. arXiv: 1109.4174 [math.ST].
- [227] G.E.P. Box and G. Jenkins. *Time Series Analysis, Forecasting and Control*. San Francisco, CA, USA: Holden-Day, Inc., 1990. ISBN: 0816211043.
- [228] P.V. Bliokh *et al.* *Schumann Resonances in the Earth-Ionosphere Cavity*. Ed. by NASA Astrophysic Data System. IEE electromagnetic waves series. Peter Peregrinus, 1980. ISBN: 9780906048337.
- [229] E. Thrane, N. Christensen, and R.M.S. Schofield. "Correlated magnetic noise in global networks of gravitational-wave detectors: Observations and implications". In: *Physical Review D* 87.12 (2013). ISSN: 1550-2368. DOI: 10.1103/physrevd.87.123009. URL: <http://dx.doi.org/10.1103/PhysRevD.87.123009>.

- [230] J.D. Romano and N.J. Cornish. "Detection methods for stochastic gravitational-wave backgrounds: a unified treatment". In: *Living Reviews in Relativity* 20.1 (2017), p. 2. ISSN: 1433-8351. DOI: 10.1007/s41114-017-0004-1. URL: <https://doi.org/10.1007/s41114-017-0004-1>.
- [231] E. Thrane *et al.* "Correlated noise in networks of gravitational-wave detectors: Subtraction and mitigation". In: *Physical Review D* 90.2 (2014). ISSN: 1550-2368. DOI: 10.1103/PhysRevD.90.023013. URL: <http://dx.doi.org/10.1103/PhysRevD.90.023013>.
- [232] T. Bayes. "Essay towards solving a problem in the doctrine of chances". In: *Biometrika* 45 (1958), pp. 293–315.
- [233] J. Von Neumann and O. Morgenstern. "Theory of games and economic behavior, 2nd rev". In: (1947).
- [234] J. Neyman and E.S. Pearson. "On the Problem of the Most Efficient Tests of Statistical Hypotheses". English. In: *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* 231 (1933), pp. 289–337. ISSN: 02643952. URL: <http://www.jstor.org/stable/91247>.
- [235] B. Allen and J.D. Romano. "Detecting a stochastic background of gravitational radiation: Signal processing strategies and sensitivities". In: *Phys. Rev. D* 59.10, 102001 (1999), p. 102001. DOI: 10.1103/PhysRevD.59.102001. eprint: [gr-qc/9710117](http://adsabs.harvard.edu/abs/1999PhRvD..59j2001A). URL: <http://adsabs.harvard.edu/abs/1999PhRvD..59j2001A>.
- [236] N. Christensen. "Measuring the stochastic gravitational-radiation background with laser-interferometric antennas". In: *Phys. Rev. D* 46 (12 1992), pp. 5250–5266. DOI: 10.1103/PhysRevD.46.5250. URL: <http://link.aps.org/doi/10.1103/PhysRevD.46.5250>.
- [237] E.E. Flanagan. "Sensitivity of the Laser Interferometer Gravitational Wave Observatory to a stochastic background, and its dependence on the detector orientations". In: *Phys. Rev. D* 48 (6 1993), pp. 2389–2407. DOI: 10.1103/PhysRevD.48.2389. URL: <http://link.aps.org/doi/10.1103/PhysRevD.48.2389>.
- [238] C. Röver. "Student-*t* based filter for robust signal detection". In: *Phys. Rev. D* 84 (12 2011), p. 122004. DOI: 10.1103/PhysRevD.84.122004. URL: <https://link.aps.org/doi/10.1103/PhysRevD.84.122004>.
- [239] T. Yamamoto *et al.* "Characterization of non-Gaussianity in gravitational wave detector noise". In: *Physical Review D* 93.8 (2016), p. 082005.
- [240] Student. "The probable error of a mean". In: *Biometrika* (1908), pp. 1–25.
- [241] B. Allen *et al.* "Robust statistics for deterministic and stochastic gravitational waves in non-Gaussian noise: Frequentist analyses". In: *Phys. Rev. D* 65.12 (2002), p. 122002.
- [242] B. Allen *et al.* "Robust statistics for deterministic and stochastic gravitational waves in nonGaussian noise. 2. Bayesian analyses". In: *Phys. Rev. D* 67 (2003), p. 122002. DOI: 10.1103/PhysRevD.67.122002. arXiv: [gr-qc/0205015](https://arxiv.org/abs/gr-qc/0205015).
- [243] B. Zackay *et al.* *Detecting Gravitational Waves in Data with Non-Gaussian Noise*. 2019. arXiv: 1908.05644 [astro-ph.IM].
- [244] P. Ajith *et al.* "Template bank for gravitational waveforms from coalescing binary black holes: Nonspinning binaries". In: *Physical Review D* 77.10 (2008), p. 104017.
- [245] C. Van Den Broeck *et al.* "Template banks to search for compact binaries with spinning components in gravitational wave data". In: *Phys. Rev. D* 80.2 (2009), p. 024009.
- [246] T. Dal Canton and I.W. Harry. "Designing a template bank to observe compact binary coalescences in Advanced LIGO's second observing run". In: *arXiv preprint arXiv:1705.01845* (2017).
- [247] E. Parzen. "Mathematical considerations in the estimation of spectra". In: *Technometrics* 3.2 (1961), pp. 167–190.
- [248] M.C. Schaerf. *Estimation of the covariance and autoregressive structure of a stationary time series*. Tech. rep. STANFORD UNIV CALIF, 1964.
- [249] A. Schuster. "On the investigation of hidden periodicities with application to a supposed 26 day period of meteorological phenomena". In: *Terrestrial Magnetism* 3.1 (1898), pp. 13–41. DOI: 10.1029/TM003i001p00013. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/TM003i001p00013>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1029/TM003i001p00013>.
- [250] J.G. Proakis. *Digital signal processing: principles algorithms and applications*. Pearson Education India, 2001.

- [251] C. Lanczos. *Applied analysis*. Courier Corporation, 1988.
- [252] J. López-Bonilla *et al.* "A note on Dirichlet and Fejér kernels". In: *Revista de Matemática: Teoría y Aplicaciones* 14 (Jan. 2007), pp. 101–104. DOI: 10.15517/rmta.v14i1.283.
- [253] M.S. Bartlett. "Periodogram analysis and continuous spectra". In: *Biometrika* 37.1/2 (1950), pp. 1–16.
- [254] J.W. Cooley and J.W. Tukey. "An algorithm for the machine calculation of complex Fourier series". In: *Mathematics of computation* 19.90 (1965), pp. 297–301.
- [255] B. Allen *et al.* "FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries". In: *Physical Review D* 85.12, 122006 (2012), p. 122006. DOI: 10.1103/PhysRevD.85.122006. arXiv: gr-qc/0509116 [gr-qc].
- [256] P. Welch. "The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms". In: *IEEE Transactions on audio and electroacoustics* 15.2 (1967), pp. 70–73.
- [257] G. Heinzel, A. Rüdiger, and R. Schilling. "Spectrum and spectral density estimation by the Discrete Fourier transform (DFT), including a comprehensive list of window functions and some new at-top windows". In: (2002). URL: <http://edoc.mpg.de/395068>.
- [258] L.K. Nuttall. "Characterizing transient noise in the LIGO detectors". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376.2120 (2018), p. 20170286.
- [259] B.P. Abbott *et al.* "First all-sky upper limits from LIGO on the strength of periodic gravitational waves using the Hough transform". In: *Physical Review D* 72.10 (2005), p. 102004.
- [260] M. Abramowitz and I.A. Stegun. "Handbook of mathematical functions". In: *Washington: National Bureau of Standards* (1999), p. 923.
- [261] D.G. Blair *et al.* *Advanced Gravitational Wave Detectors*. 2012.
- [262] P.B. Covas and LSC Instrument Authors. "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, 082002 (2018), p. 082002. DOI: 10.1103/PhysRevD.97.082002. arXiv: 1801.07204 [astro-ph.IM].
- [263] P. Billingsley. *Probability and measure*. Wiley series in probability and mathematical statistics. Probability and mathematical statistics. Wiley, 1979. ISBN: 9780471031734. URL: <http://books.google.it/books?id=6hRRAAAAAAAJ>.
- [264] F. Acernese *et al.* "Noise from scattered light in Virgo second science run data". In: *Classical and Quantum Gravity* 27.19 (2010), p. 194011. DOI: 10.1088/0264-9381/27/19/194011.
- [265] F. Ardhuin *et al.* "Ocean wave sources of seismic noise". In: *Journal of Geophysical Research: Oceans* 116.C9 (2011).
- [266] F. Acernese *et al.* "Properties of seismic noise at the Virgo site". In: *Classical and Quantum Gravity* 21.5 (2004), S433.
- [267] P.R. Saulson. "Terrestrial gravitational noise on a gravitational wave antenna". In: *Physical Review D* 30.4 (1984), p. 732.
- [268] G. Cella. "Off-line subtraction of seismic Newtonian noise". In: *Recent Developments in General Relativity*. Springer, 2000, pp. 495–503.
- [269] M.W. Coughlin *et al.* "Implications of Dedicated Seismometer Measurements on Newtonian-Noise Cancellation for Advanced LIGO". In: *Physical Review Letters* 121.22 (2018). ISSN: 1079-7114. DOI: 10.1103/physrevlett.121.221104. URL: <http://dx.doi.org/10.1103/PhysRevLett.121.221104>.
- [270] H.B. Callen and R.F. Greene. "On a theorem of irreversible thermodynamics". In: *Physical Review* 86.5 (1952), p. 702.
- [271] T. Akutsu *et al.* "First cryogenic test operation of underground km-scale gravitational-wave observatory KAGRA". In: *Classical and Quantum Gravity* 36.16 (2019), p. 165008. ISSN: 1361-6382. DOI: 10.1088/1361-6382/ab28a9. URL: <http://dx.doi.org/10.1088/1361-6382/ab28a9>.

- [272] J. Agresti *et al.* "Optimized multilayer dielectric mirror coatings for gravitational wave interferometers". In: *Advances in Thin-Film Coatings for Optical Applications III*. Vol. 6286. International Society for Optics and Photonics. 2006, p. 628608.
- [273] T. Accadia *et al.* "The NoEMi (Noise Frequency Event Miner) framework". In: *Journal of Physics: Conference Series* 363.1 (2012), p. 012037. URL: <http://stacks.iop.org/1742-6596/363/i=1/a=012037>.
- [274] V. Tiwari and M. Drago *et al.* "Regression of environmental noise in LIGO data". In: *Classical and Quantum Gravity* 32.16 (2015), p. 165014. DOI: 10.1088/0264-9381/32/16/165014. URL: <https://doi.org/10.1088/0264-9381/32/16/165014>.
- [275] D. Davis and T.J. Massinger *et al.* "Improving the sensitivity of Advanced LIGO using noise subtraction". In: *Classical and Quantum Gravity* 36.5 (2019), p. 055011. ISSN: 1361-6382. DOI: 10.1088/1361-6382/ab01c5. URL: <http://dx.doi.org/10.1088/1361-6382/ab01c5>.
- [276] J. Smith *et al.* "Linear projection of technical noise for interferometric gravitational-wave detectors". In: *Classical and Quantum Gravity* 23 (2006), pp. 527–537. DOI: 10.1088/0264-9381/23/2/016.
- [277] S.M. Kay. *Modern Spectral Estimation*. Englewood Cliffs, NJ: Prentice-Hall, 1988.
- [278] A.G. González *et al.* "Multiple coherence method in time domain for the analysis of the transmission paths of noise and vibrations with non stationary signals". In: ().
- [279] W.A. Gardner, A. Napolitano, and L. Paura. "Cyclostationarity: Half a century of research". In: *Signal Processing* 86.4 (2006), pp. 639–697. ISSN: 0165-1684. DOI: <https://doi.org/10.1016/j.sigpro.2005.06.016>. URL: <http://www.sciencedirect.com/science/article/pii/S0165168405002409>.
- [280] A. Grinsted, J. C. Moore, and S. Jevrejeva. "Application of the cross wavelet transform and wavelet coherence to geophysical time series". In: *Nonlinear Processes in Geophysics* 11.5/6 (2004), pp. 561–566. DOI: 10.5194/npg-11-561-2004. URL: <https://www.nonlin-processes-geophys.net/11/561/2004/>.
- [281] G. Matz and F. Hlawatsch. "Time-frequency coherence analysis of nonstationary random processes". In: *Proceedings of the Tenth IEEE Workshop on Statistical Signal and Array Processing (Cat. No.00TH8496)*. 2000, pp. 554–558. DOI: 10.1109/SSAP.2000.870186.
- [282] L. B. White and B. Boashash. "Cross spectral analysis of nonstationary processes". In: *IEEE Transactions on Information Theory* 36.4 (1990), pp. 830–835. ISSN: 0018-9448. DOI: 10.1109/18.53742.
- [283] M. Zevin *et al.* "Gravity Spy: integrating advanced LIGO detector characterization, machine learning, and citizen science". In: *Classical and Quantum Gravity* 34.6 (2017), p. 064003. URL: <http://stacks.iop.org/0264-9381/34/i=6/a=064003>.
- [284] W. Wei and E.A. Huerta. "Gravitational wave denoising of binary black hole mergers with deep learning". In: *Physics Letters B* 800 (2020), p. 135081.
- [285] S. Kulkarni *et al.* "NNETFIX: A Neural Network to "fix" Gravitational Wave data affected by short duration noise transients". 2020. in preparation.
- [286] J. Aasi *et al.* "Characterization of the LIGO detectors during their sixth science run". In: *Classical and Quantum Gravity* 32.11 (2015), p. 115012. ISSN: 1361-6382. DOI: 10.1088/0264-9381/32/11/115012. URL: <http://dx.doi.org/10.1088/0264-9381/32/11/115012>.
- [287] B.P. Abbott *et al.* "Upper Limits on the Stochastic Gravitational-Wave Background from Advanced LIGO's First Observing Run". In: *Phys. Rev. Lett.* 118 (12 2017), p. 121101. DOI: 10.1103/PhysRevLett.118.121101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.118.121101>.
- [288] D. Macleod *et al.* *gwpy/gwpy: 1.0.0*. Version v1.0.0. 2019. DOI: 10.5281/zenodo.3522047. URL: <https://doi.org/10.5281/zenodo.3522047>.
- [289] L.S. Finn, G. Gonzalez, and P.J. Sutton. *Rayleigh Monitor (Gaussianity test for LIGO data)*. 2001. URL: http://gallatin.physics.lsa.umich.edu/~keithr/lscdc/sutton_may01.html.
- [290] D.V. Martynov *et al.* "Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy". In: *Physical Review D* 93.11 (2016). ISSN: 2470-0029. DOI: 10.1103/physrevd.93.112004. URL: <http://dx.doi.org/10.1103/PhysRevD.93.112004>.

- [291] C. Torrence and G. P. Compo. "A Practical Guide to Wavelet Analysis." In: *Bulletin of the American Meteorological Society* 79 (1998), pp. 61–78. DOI: 10.1175/1520-0477(1998)079<061T:APGTWA>2.0.CO;2. URL: <https://ui.adsabs.harvard.edu/abs/1998BAMS...79...61T>.
- [292] W.A. Fuller. *Introduction to statistical time series*. 1976.
- [293] M.P. Priestley and T. Subba Rao. "A Test for Non-Stationarity of Time-Series". In: *Journal of the Royal Statistical Society. Series B (Methodological)* 31.1 (1969), pp. 140–149. ISSN: 00359246. URL: <http://www.jstor.org/stable/2984336>.
- [294] M.S. Bartlett. "The use of transformations". In: *Biometrics* 3.1 (1947), pp. 39–52.
- [295] G.M. Jenkins and M.B. Priestley. "The Spectral Analysis of Time-Series". In: *Journal of the Royal Statistical Society: Series B (Methodological)* 19.1 (1957), pp. 1–12. DOI: 10.1111/j.2517-6161.1957.tb00240.x. eprint: <https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.2517-6161.1957.tb00240.x>. URL: <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.2517-6161.1957.tb00240.x>.
- [296] J.L. Doob. "The limiting distributions of certain statistics". In: *The Annals of Mathematical Statistics* 6.3 (1935), pp. 160–169.
- [297] B. Mitchell. "A Comparison of Chi-Square and Kolmogorov-Smirnov Tests". In: *Area* 3.4 (1971), pp. 237–241. ISSN: 00040894, 14754762. URL: <http://www.jstor.org/stable/20000590>.
- [298] A. Kolmogoroff. "Confidence Limits for an Unknown Distribution Function". In: *Ann. Math. Statist.* 12.4 (1941), pp. 461–463. DOI: 10.1214/aoms/1177731684. URL: <https://doi.org/10.1214/aoms/1177731684>.
- [299] F. Acernese *et al.* "A simple line detection algorithm applied to Virgo data". In: *Classical and Quantum Gravity* 22.18 (2005), S1189.
- [300] T.B. Littenberg and N.J. Cornish. "Bayesian inference for spectral estimation of gravitational wave detector noise". In: *Physical Review D* 91.8 (2015), p. 084034.
- [301] D.C. Hoaglin, F. Mosteller, and J.W. Tukey. *Understanding robust and exploratory data analysis*. Vol. 3. Wiley New York, 1983.
- [302] The pandas development team. *pandas-dev/pandas: Pandas*. Version latest. Feb. 2020. DOI: 10.5281/zenodo.3509134. URL: <https://doi.org/10.5281/zenodo.3509134>.
- [303] M.H. van Emden. "Algorithms 402: Increasing the Efficiency of Quicksort". In: *Commun. ACM* 13.11 (1970), pp. 693–694. ISSN: 0001-0782. DOI: 10.1145/362790.362803. URL: <http://doi.acm.org/10.1145/362790.362803>.
- [304] A. Kolmogorov. "Sulla determinazione empirica di una legge di distribuzione". In: *Inst. Ital. Attuari, Giorn.* 4 (1933), pp. 83–91.
- [305] N. Smirnov. "Sur les écarts de la courbe de distribution empirique". In: *Matematicheskii Sbornik* 48.1 (1939), pp. 3–26.
- [306] H. Cramér. "On the composition of elementary errors: First paper: Mathematical deductions". In: *Scandinavian Actuarial Journal* 1928.1 (1928), pp. 13–74.
- [307] R. Von Mises. *Wahrscheinlichkeit, Statistik und Wahrheit: Einführung in d. neue Wahrscheinlichkeitslehre u. ihre Anwendung*. Vol. 3. Springer-Verlag, 1928.
- [308] T.W. Anderson. "On the distribution of the two-sample Cramer-von Mises criterion". In: *The Annals of Mathematical Statistics* (1962), pp. 1148–1159.
- [309] F. Schmid and M. Tiede. "An L^1 -variant of the Cramér-von Mises test". In: *Statistics & probability letters* 26.1 (1996), pp. 91–96.
- [310] Y. Xiao, A. Gordon, and A. Yakovlev. "The L^1 -Version of the Cramér-von Mises Test for Two-Sample Comparisons in Microarray Data Analysis". In: *EURASIP Journal on Bioinformatics and Systems Biology* 2006.1 (2006), p. 85769.
- [311] F.J. Massey Jr. "The Kolmogorov-Smirnov test for goodness of fit". In: *Journal of the American statistical Association* 46.253 (1951), pp. 68–78.
- [312] P. Virtanen *et al.* "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: *Nature Methods* 17 (2020), pp. 261–272. DOI: <https://doi.org/10.1038/s41592-019-0686-2>.

- [313] S. Engmann and D. Cousineau. "Comparing distributions: the two-sample Anderson-Darling test as an alternative to the Kolmogorov-Smirnov test". In: *Journal of Applied Quantitative Methods* 6.3 (2011), pp. 1–17.
- [314] H.W. Lilliefors. "On the Kolmogorov-Smirnov test for normality with mean and variance unknown". In: *Journal of the American statistical Association* 62.318 (1967), pp. 399–402.
- [315] H.W. Lilliefors. "On the Kolmogorov-Smirnov test for the exponential distribution with mean unknown". In: *Journal of the American Statistical Association* 64.325 (1969), pp. 387–389.
- [316] D.A. Darling. "The Kolmogorov-Smirnov, Cramer-von Mises Tests". In: *Ann. Math. Statist.* 28.4 (1957), pp. 823–838. DOI: 10.1214/aoms/1177706788. URL: <https://doi.org/10.1214/aoms/1177706788>.
- [317] A.N. Pettitt. "A two-sample Anderson-Darling rank statistic". In: *Biometrika* 63.1 (1976), pp. 161–168.
- [318] O. Cucconi. "Un nuovo test non parametrico per il confronto fra due gruppi di valori campionari". In: *Giornale degli Economisti e Annali di Economia* (1968), pp. 225–248.
- [319] M. Farge. "Wavelet transforms and their applications to turbulence". In: *Annual review of fluid mechanics* 24.1 (1992), pp. 395–458.
- [320] M.A.G. Izquierdo *et al.* "Signal-to-noise ratio enhancement based on the whitening transformation of colored structural noise". In: *Ultrasonics* 38.1 (2000), pp. 500–502. ISSN: 0041-624X. DOI: [https://doi.org/10.1016/S0041-624X\(99\)00073-6](https://doi.org/10.1016/S0041-624X(99)00073-6). URL: <http://www.sciencedirect.com/science/article/pii/S0041624X99000736>.
- [321] S.D. Meyers, B.G. Kelly, and J.J. O'Brien. "An introduction to wavelet analysis in oceanography and meteorology: With application to the dispersion of Yanai waves". In: *Monthly weather review* 121.10 (1993), pp. 2858–2866.
- [322] T.W. Anderson and D.A. Darling *et al.* "Asymptotic theory of certain 'goodness of fit' criteria based on stochastic processes". In: *The annals of mathematical statistics* 23.2 (1952), pp. 193–212.
- [323] S.K. Chatterji. "The Search for Gravitational Wave Bursts in Data from the Second LIGO Science Run". PhD thesis. PhD Thesis: Massachusetts Institute of Technology, 2005. URL: "https://gwic.ligo.org/thesisprize/2006/Chatterji_Thesis.pdf".
- [324] T. Adams *et al.* "Low-latency analysis pipeline for compact binary coalescences in the advanced gravitational wave detector era". In: *Classical and Quantum Gravity* 33.17 (2016), p. 175012.
- [325] J. Smith *et al.* "A hierarchical method for vetoing noise transients in gravitational-wave detectors". In: *Classical and Quantum Gravity* 28.23 (2011), p. 235005. ISSN: 1361-6382. DOI: 10.1088/0264-9381/28/23/235005. URL: <http://dx.doi.org/10.1088/0264-9381/28/23/235005>.
- [326] T. Isogai and the Ligo Scientific Collaboration a Collaboration. "Used percentage veto for LIGO and virgo binary inspiral searches". In: *Journal of Physics: Conference Series* 243 (2010), p. 012005. DOI: 10.1088/1742-6596/243/1/012005. URL: <https://doi.org/10.1088/1742-6596/243/1/012005>.
- [327] J.C. Driggers *et al.* "Active noise cancellation in a suspended interferometer". In: *Review of Scientific Instruments* 83.2 (2012), p. 024501. DOI: 10.1063/1.3675891. eprint: <https://doi.org/10.1063/1.3675891>. URL: <https://doi.org/10.1063/1.3675891>.
- [328] G. Vajente *et al.* *Machine-learning non-stationary noise out of gravitational wave detectors*. 2019. arXiv: 1911.09083 [gr-qc].
- [329] P. Kwee *et al.* "Stabilized high-power laser system for the gravitational wave detector advanced LIGO". In: *Opt. Express* 20.10 (2012), pp. 10617–10634. DOI: 10.1364/OE.20.010617. URL: <http://www.opticsexpress.org/abstract.cfm?URI=oe-20-10-10617>.
- [330] M.M. Siddiqui. "Statistical inference for Rayleigh distributions". In: *Sec. D* 68.9 (1964), p. 1007.
- [331] T.P. Minka. *Estimating a Gamma distribution*. 2002. URL: <https://tminka.github.io/papers/minka-gamma.pdf>.
- [332] J.J.A. Moors. "The meaning of kurtosis: Darlington reexamined". In: *The American Statistician* 40.4 (1986), pp. 283–284.
- [333] P.H. Westfall. "Kurtosis as peakedness, 1905–2014. RIP". In: *The American Statistician* 68.3 (2014), pp. 191–195.

- [334] R.C. Gonzalez and R.E. Woods *et al.* *Digital image processing*. 2002.
- [335] B. Jähne. *Digital Image Processing*. EngineeringPro collection. Springer Berlin Heidelberg, 2005. ISBN: 9783540275633. URL: <https://books.google.it/books?id=GvWYBrDXMnkC>.
- [336] Irwin Sobel. "An Isotropic 3x3 Image Gradient Operator". In: *Presentation at Stanford A.I. Project 1968* (Feb. 2014).
- [337] R.O. Duda and P.E. Hart. *Use of the Hough transformation to detect lines and curves in pictures*. Tech. rep. Sri International Menlo Park Ca Artificial Intelligence Center, 1971.
- [338] J.M.S. Prewitt. "Object enhancement and extraction". In: *Picture processing and Psychopictorics* 10.1 (1970), pp. 15–19.
- [339] T.M. Cover and J.A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [340] D.J.C. MacKay and D.J.C. Mac Kay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003.
- [341] S. Kullback and R.A. Leibler. "On information and sufficiency". In: *The annals of mathematical statistics* 22.1 (1951), pp. 79–86.
- [342] J.L.W.V. Jensen. "Sur les fonctions convexes et les inégalités entre les valeurs moyennes". In: *Acta Mathematica* 30.1 (1906), pp. 175–193. ISSN: 1871-2509. DOI: 10.1007/BF02418571. URL: <https://doi.org/10.1007/BF02418571>.
- [343] C.M. Bishop *et al.* *Neural networks for pattern recognition*. Oxford university press, 1995.
- [344] C. Spearman. "The proof and measurement of association between two things." In: (1961).
- [345] A. Kraskov, H. Stögbauer, and P. Grassberger. "Estimating mutual information". In: *Physical review E* 69.6 (2004), p. 066138.
- [346] A.S. Weigend. *Time Series Prediction: Forecasting The Future And Understanding The Past*. Taylor & Francis, 2018. ISBN: 9780429972270. URL: <https://books.google.it/books?id=uUdaDwAAQBAJ>.
- [347] J. Neter *et al.* *Applied Linear Statistical Models*. Chicago: Irwin, 1996.
- [348] A. Sen and M. Srivastava. *Regression Analysis: Theory, Methods, and Applications*. Springer Texts in Statistics. Springer New York, 1997. ISBN: 9780387972114. URL: <https://books.google.it/books?id=Jh86k9DAtr0C>.
- [349] S.M. Kay. *Fundamentals of Statistical Signal Processing: Estimation theory*. Prentice-Hall PTR, 1993.
- [350] T.K. Moon and W.C. Stirling. *Mathematical Methods and Algorithms for Signal Processing*. Prentice Hall, 2000. ISBN: 9780201361865. URL: <https://books.google.it/books?id=1xUfAQAIAAJ>.
- [351] I.T. Jolliffe. "A Note on the Use of Principal Components in Regression". In: *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 31.3 (1982), pp. 300–303. ISSN: 00359254, 14679876. URL: <http://www.jstor.org/stable/2348005>.
- [352] N. Halko, P.-G. Martinsson, and J.A. Tropp. "Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions". In: *ArXiv e-prints* (Sept. 2009). arXiv: 0909.4061 [math.NA].
- [353] D. Dereniowski and M. Kubale. "Cholesky Factorization of Matrices in Parallel and Ranking of Graphs". In: *Parallel Processing and Applied Mathematics*. Ed. by R. Wyrzykowski *et al.* Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 985–992. ISBN: 978-3-540-24669-5.
- [354] E. Bair *et al.* "Prediction by Supervised Principal Components". In: *Journal of the American Statistical Association* 101.473 (2006), pp. 119–137. DOI: 10.1198/016214505000000628. eprint: <https://doi.org/10.1198/016214505000000628>. URL: <https://doi.org/10.1198/016214505000000628>.
- [355] E. Barshan *et al.* "Supervised principal component analysis: Visualization, classification and regression on subspaces and submanifolds". In: *Pattern Recognition* 44.7 (2011), pp. 1357–1371. ISSN: 0031-3203. DOI: <https://doi.org/10.1016/j.patcog.2010.12.015>. URL: <http://www.sciencedirect.com/science/article/pii/S0031320310005819>.
- [356] M. Walker *et al.* "Identifying correlations between LIGO's astronomical range and auxiliary sensors using lasso regression". In: *ArXiv e-prints* (2018). arXiv: 1807.02592 [astro-ph.IM].

- [357] H. Zou and T. Hastie. "Regularization and variable selection via the elastic net". In: *Journal of the royal statistical society: series B (statistical methodology)* 67.2 (2005), pp. 301–320.
- [358] S.S. Shapiro and M.B. Wilk. "An analysis of variance test for normality (complete samples)". In: *Biometrika* 52.3-4 (1965), pp. 591–611. DOI: 10.1093/biomet/52.3-4.591. eprint: /oup/backfile/content_public/journal/biomet/52/3-4/10.1093/biomet/52.3-4.591/2/52-3-4-591.pdf. URL: <http://dx.doi.org/10.1093/biomet/52.3-4.591>.
- [359] N. Smirnov. "Table for Estimating the Goodness of Fit of Empirical Distributions". In: *Ann. Math. Statist.* 19.2 (1948), pp. 279–281. DOI: 10.1214/aoms/1177730256. URL: <https://doi.org/10.1214/aoms/1177730256>.
- [360] C.M. Jarque and A.K. Bera. "Efficient tests for normality, homoscedasticity and serial independence of regression residuals". In: *Economics Letters* 6.3 (1980), pp. 255–259. ISSN: 0165-1765. DOI: [https://doi.org/10.1016/0165-1765\(80\)90024-5](https://doi.org/10.1016/0165-1765(80)90024-5). URL: <http://www.sciencedirect.com/science/article/pii/0165176580900245>.
- [361] J. Durbin and G.S. Watson. "Testing for serial correlation in least squares regression." In: *Biometrika* 37.3-4 (1950), pp. 409–428. DOI: 10.1093/biomet/37.3-4.409. eprint: /oup/backfile/content_public/journal/biomet/37/3-4/10.1093/biomet/37.3-4.409/2/37-3-4-409.pdf. URL: <http://dx.doi.org/10.1093/biomet/37.3-4.409>.
- [362] R.F. Engle. "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation". In: *Econometrica* 50.4 (1982), pp. 987–1007. ISSN: 00129682, 14680262. URL: <http://www.jstor.org/stable/1912773>.
- [363] R. Bellman, Rand Corporation, and Karreman Mathematics Research Collection. *Dynamic Programming*. Rand Corporation research study. Princeton University Press, 1957. ISBN: 9780691079516. URL: <https://books.google.it/books?id=wtdoPwAACAAJ>.
- [364] H. Akaike. "A new look at the statistical model identification". In: *IEEE Transactions on Automatic Control* 19.6 (1974), pp. 716–723. ISSN: 0018-9286. DOI: 10.1109/TAC.1974.1100705.
- [365] G. Schwarz. "Estimating the Dimension of a Model". In: *Ann. Statist.* 6.2 (1978), pp. 461–464. DOI: 10.1214/aos/1176344136. URL: <https://doi.org/10.1214/aos/1176344136>.
- [366] T.M. Mitchell. "Does machine learning really work?" In: *AI magazine* 18.3 (1997), pp. 11–11.
- [367] V. Volterra. *Sopra le funzioni che dipendono da altre funzioni*. Tip. della R. Accademia dei Lincei, 1887.
- [368] V. Necula, S. Klimenko, and G. Mitselmakher. "Transient analysis with fast Wilson-Daubechies time-frequency transform". In: *Journal of Physics: Conference Series*. Vol. 363. 1. IOP Publishing, 2012, p. 012032.
- [369] J.S. Walker. "Fourier analysis and wavelet analysis". In: *Notices of the AMS* 44.6 (1997), pp. 658–670.
- [370] D. Gabor. "Theory of communication. Part 1: The analysis of information". In: *Journal of the Institution of Electrical Engineers-Part III: Radio and Communication Engineering* 93.26 (1946), pp. 429–441.
- [371] G. Kaiser. *A friendly guide to wavelets*. Springer Science & Business Media, 2010.
- [372] Maria Principe and Innocenzo M. Pinto. "Modeling the impulsive noise component and its effect on the operation of a simple coherent network algorithm for detecting unmodeled gravitational wave bursts". In: *Class. Quant. Grav.* 25 (2008), p. 075013. DOI: 10.1088/0264-9381/25/7/075013. arXiv: 0806.4574 [gr-qc].
- [373] P. Goupillaud, A. Grossmann, and J. Morlet. "Cycle-octave and related transforms in seismic signal analysis". In: *Geoexploration* 23.1 (1984), pp. 85–102.
- [374] W.A. Gardner. "A unifying view of coherence in signal processing". In: *Signal Processing* 29.2 (1992), pp. 113–140.
- [375] M. Onorato *et al.* "Application of a wavelet cross-correlation analysis to DNS velocity signals". In: *European journal of mechanics series b fluids* 16 (1997), pp. 575–597.
- [376] D. Labat, R. Ababou, and A. Mangin. "Analyse multirésolution croisée de pluies et débits de sources karstiques". In: *Comptes Rendus Géoscience* 334.8 (2002), pp. 551–556.

- [377] I. Daubechies. "The wavelet transform, time-frequency localization and signal analysis". In: *IEEE Transactions on Information Theory* 36 (Sept. 1990), pp. 961–1005. URL: <https://ui.adsabs.harvard.edu/abs/1990ITIT...36..961D>.
- [378] J.O. Smith. *Spectral Audio Signal Processing*. online book, 2011 edition. W3K Publishing, 2011. ISBN: 978-0-9745607-3-1. URL: <http://ccrma.stanford.edu/~jos/sasp/>.
- [379] D. Zhang. "Wavelet Transform". In: *Fundamentals of Image Data Mining: Analysis, Features, Classification and Retrieval*. Cham: Springer International Publishing, 2019, pp. 35–44. DOI: 10.1007/978-3-030-17989-2_3. URL: https://doi.org/10.1007/978-3-030-17989-2_3.
- [380] A. Grinsted, J.C. Moore, and S. Jevrejeva. "Application of the cross wavelet transform and wavelet coherence to geophysical time series". In: *Nonlinear processes in geophysics* 11.5/6 (2004), pp. 561–566.
- [381] B. Cazelles *et al.* "Wavelet analysis of ecological time series". In: *Oecologia* 156.2 (2008), pp. 287–304.
- [382] D.T.L. Lee and A. Yamamoto. "Wavelet analysis: theory and applications". In: *Hewlett Packard journal* 45 (1994), pp. 44–44.
- [383] J.C. Brown. "Calculation of a constant Q spectral transform". In: *The Journal of the Acoustical Society of America* 89.1 (1991), pp. 425–434.
- [384] J.C. Brown and M.S. Puckette. "An efficient algorithm for the calculation of a constant Q transform". In: *The Journal of the Acoustical Society of America* 92.5 (1992), pp. 2698–2701.
- [385] L. Qiu and M.H. Er. "Wavelet spectrogram of noisy signals". In: *International journal of electronics* 79.5 (1995), pp. 665–677.
- [386] D.A. Dickey. *The analysis of time series: an introduction*. 1991.
- [387] C. Torrence and P.J. Webster. "Interdecadal Changes in the ENSO–Monsoon System". In: *Journal of Climate* 12.8 (1999), pp. 2679–2690. DOI: 10.1175/1520-0442(1999)012<2679:ICITEM>2.0.CO;2. eprint: [https://doi.org/10.1175/1520-0442\(1999\)012<2679:ICITEM>2.0.CO;2](https://doi.org/10.1175/1520-0442(1999)012<2679:ICITEM>2.0.CO;2). URL: [https://doi.org/10.1175/1520-0442\(1999\)012<2679:ICITEM>2.0.CO;2](https://doi.org/10.1175/1520-0442(1999)012<2679:ICITEM>2.0.CO;2).
- [388] P.C. Liu. "Wavelet spectrum analysis and ocean wind waves". In: *Wavelet Analysis and Its Applications*. Vol. 4. Elsevier, 1994, pp. 151–166.
- [389] W.T. Wells, R.L. Anderson, and J.W. Cell *et al.* "The distribution of the product of two central or non-central chi-square variates". In: *The Annals of Mathematical Statistics* 33.3 (1962), pp. 1016–1020.
- [390] S. Kotz and R. Srinivasan. "Distribution of product and quotient of Bessel function variates". In: *Annals of the Institute of Statistical Mathematics* 21.1 (1969), pp. 201–210.
- [391] Z. Ge. "Significance tests for the wavelet cross spectrum and wavelet linear coherence". In: *Annales geophysicae: atmospheres, hydrospheres and space sciences*. Vol. 26. 12. 2008, p. 3819.
- [392] S. Sello and J. Bellazzini. "Wavelet cross-correlation analysis of turbulent mixing from large-eddy-simulations". In: *arXiv preprint physics/0003029* (2000).
- [393] D. Labat. "Recent advances in wavelet analyses: Part 1. A review of concepts". In: *Journal of Hydrology* 314.1-4 (2005), pp. 275–288.
- [394] H. Li and T. Nozaki. "Application of wavelet cross-correlation analysis to a plane turbulent jet". In: *JSME International Journal Series B Fluids and Thermal Engineering* 40.1 (1997), pp. 58–66.
- [395] A. Rua. "Worldwide synchronization since the nineteenth century: a wavelet-based view". In: *Applied Economics Letters* 20.8 (2013), pp. 773–776.
- [396] S. Chatterji *et al.* "Multiresolution techniques for the detection of gravitational-wave bursts". In: *Classical and Quantum Gravity* 21.20 (2004), S1809.
- [397] O.R. Olayeni. "Causality in continuous wavelet transform without spectral matrix factorization: theory and application". In: *Computational Economics* 47.3 (2016), pp. 321–340.
- [398] L. Aguiar-Conraria and M.J. Soares. "Oil and the macroeconomy: using wavelets to analyze old issues". In: *Empirical Economics* 40.3 (2011), pp. 645–655.
- [399] M. Dhamala, G. Rangarajan, and M. Ding. "Analyzing information flow in brain networks with non-parametric Granger causality". In: *Neuroimage* 41.2 (2008), pp. 354–362.

- [400] M. Dhamala, G. Rangarajan, and M. Ding. "Estimating Granger causality from Fourier and wavelet transforms of time series data". In: *Physical review letters* 100.1 (2008), p. 018701.
- [401] C.W.J. Granger. "Testing for causality and feedback". In: *Econometrica* 37.3 (1969), pp. 424–438.
- [402] J. Geweke. "Measurement of Linear Dependence and Feedback between Multiple Time Series". In: *Journal of the American Statistical Association* 77.378 (1982), pp. 304–313. DOI: 10.1080/01621459.1982.10477803. eprint: <https://www.tandfonline.com/doi/pdf/10.1080/01621459.1982.10477803>. URL: <https://www.tandfonline.com/doi/abs/10.1080/01621459.1982.10477803>.
- [403] G.T. Wilson. "A convergence theorem for spectral factorization". In: *Journal of Multivariate Analysis* 8.2 (1978), pp. 222–232.
- [404] W. Hesse *et al.* "The use of time-variant EEG Granger causality for inspecting directed interdependencies of neural assemblies". In: *Journal of neuroscience methods* 124.1 (2003), pp. 27–44.
- [405] M. Ding *et al.* "Short-window spectral analysis of cortical event-related potentials by adaptive multivariate autoregressive modeling: data preprocessing, model validation, and variability assessment". In: *Biological cybernetics* 83.1 (2000), pp. 35–45.
- [406] Y. Chen *et al.* "Analyzing multiple nonlinear time series with extended Granger causality". In: *Physics Letters A* 324.1 (2004), pp. 26–35.
- [407] N. Ancona, D. Marinazzo, and S. Stramaglia. "Radial basis function approach to nonlinear Granger causality of time series". In: *Physical Review E* 70.5 (2004), p. 056221.
- [408] T. Schreiber. "Measuring information transfer". In: *Physical review letters* 85.2 (2000), p. 461.
- [409] K. Hlaváčková-Schindler *et al.* "Causality detection based on information-theoretic approaches in time series analysis". In: *Physics Reports* 441.1 (2007), pp. 1–46.
- [410] L. Barnett, A.B. Barrett, and A.K. Seth. "Granger causality and transfer entropy are equivalent for Gaussian variables". In: *Physical review letters* 103.23 (2009), p. 238701.
- [411] S. Seabold and J. Perktold. "statsmodels: Econometric and statistical modeling with python". In: *9th Python in Science Conference*. 2010.
- [412] G. Sugihara and R.M. May. "Nonlinear forecasting as a way of distinguishing chaos from measurement error in time series". In: *Nature* 344.6268 (1990), p. 734.
- [413] A. Katok and B. Hasselblatt. *Introduction to the modern theory of dynamical systems*. Vol. 54. Cambridge university press, 1997.
- [414] F. Takens. "The reconstruction theorem for endomorphisms". In: *Bulletin of the Brazilian Mathematical Society* 33.2 (2002), pp. 231–262.
- [415] D. Mønster *et al.* "Causal inference from noisy time-series data—Testing the Convergent Cross-Mapping algorithm in the presence of noise and external influence". In: *Future Generation Computer Systems* 73 (2017), pp. 52–62.
- [416] M.B. Kennel, R. Brown, and H.D.I. Abarbanel. "Determining embedding dimension for phase-space reconstruction using a geometrical construction". In: *Physical review A* 45.6 (1992), p. 3403.
- [417] H. Abarbanel. *Analysis of observed chaotic data*. Springer Science & Business Media, 2012.
- [418] A. Krakovská, K. Mezeiová, and H. Budáčová. "Use of false nearest neighbours for selecting variables and embedding parameters for state space reconstruction". In: *Journal of Complex Systems* 2015 (2015).
- [419] H. Kantz and T. Schreiber. "Determinism and predictability". In: *Nonlinear time series analysis* (1997), pp. 42–57.
- [420] S. Wallot and D. Mønster. "Calculation of average mutual information (AMI) and false-nearest neighbors (FNN) for the estimation of embedding parameters of multidimensional time series in Matlab". In: *Frontiers in psychology* 9 (2018), p. 1679.
- [421] N. Crtale. *Convergent Cross Mapping in Scikit Learn's style*. 2018. URL: <https://github.com/nickc1/skccm>.
- [422] F. Pedregosa *et al.* "Scikit-learn: Machine Learning in Python". In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.

- [423] D. Mønster. *XMAP*. 2013. URL: <https://github.com/danm0nster/xmap>.
- [424] E. Sanchez-Palencia and P. Lherminier. "Paradoxes of vulnerability to predation in biological dynamics and mediate versus immediate causality". In: *Discrete & Continuous Dynamical Systems-S* (2018), p. 0.
- [425] J.M. McCracken and R.S. Weigel. "Convergent cross-mapping and pairwise asymmetric inference". In: *Physical Review E* 90.6 (2014), p. 062903.
- [426] X. Glorot, A. Bordes, and Y. Bengio. "Deep sparse rectifier neural networks". In: *Proceedings of the fourteenth international conference on artificial intelligence and statistics*. 2011, pp. 315–323.
- [427] D.E. Rumelhart, G.E. Hinton, and R.J. Williams. "Learning representations by back-propagating errors". In: *nature* 323.6088 (1986), pp. 533–536.
- [428] LIGO Scientific Collaboration. *LIGO Algorithm Library - LALSuite*. free software (GPL). 2018. DOI: 10.7935/GT1W-FZ16.
- [429] C.M. Bishop. *Pattern recognition and machine learning*. springer, 2006.
- [430] S. Hochreiter and J. Schmidhuber. "Long Short-Term Memory". In: *Neural Computation* 9.8 (1997). cited By 13600, pp. 1735–1780. DOI: 10.1162/neo.1997.9.8.1735.

LVC internal bibliography

This bibliography contains the references to the internal documenting systems of the LIGO and Virgo collaborations, that is the DCC and TDS. The corresponding entries are mostly accessible with LVC credentials at the provided links. Moreover, we collected here the links to various entries to the LIGO and Virgo logbooks. These resources are mainly meant for internal use, hence we have decided to separate them from the main bibliography.

- [LVC1] The Virgo detector characterization group. *Data quality model for Advanced Virgo. Tech. rep. 2015*. URL: <https://tds.virgo-gw.eu/ql/?c=10982>..
- [LVC2] G. Vajente. "At the border of commissioning and DetChar: brute force coherence and non stationary noise analysis". LSC-Virgo meeting - Pasadena, March 2015, LIGO-G1500230.
- [LVC3] L. Salconi EGO IT Department. *VIR-0037B-20: DAQ data storage and flows*. URL: <https://tds.virgo-gw.eu/?content=3&r=16658>.
- [LVC4] The Data Analysis Software and Computing Group of the Virgo experiments. *The Virgo computing farm*. Tech. rep. <https://scientists.virgo-gw.eu/computing/clusters/>.
- [LVC5] G. Losurdo for the Virgo Collaboration. *Advanced Virgo Technical Design Report, VIR-0128A-12, available at https://tds.ego-gw.it/ql/?c=8940*. 2012.
- [LVC6] G. Hemming and D. Verkindt. *VIR-0546A-16: Virgo Interferometer Monitor (VIM) Web User Interface (WUI) User Guide*. 2016. URL: <https://tds.virgo-gw.eu/ql/?c=11869>.
- [LVC7] G. Vajente. "Measurement of control noise budget". Virgo Internal note, VIR-003A-08.
- [LVC8] G. Vajente. *Brute Force Coherence (bruco) documentation*. Tech. rep. Virgo Internal note, <https://wiki.virgo-gw.eu/DetChar/BrucoDoc>, 2015.
- [LVC9] F. Di Renzo and A. Romero. *Detchar Shift Report: VIR-0553A-19*. 2019. URL: https://wiki.virgo-gw.eu/DetChar/DetcharShift2019_May21_May28.
- [LVC10] G. Cella and B. Patricelli. *VIR-0575B-19: Tools for modulated noise study*. 2019. URL: <https://tds.virgo-gw.eu/?content=3&r=15729>.
- [LVC11] F. Di Renzo and B. Patricelli. *150 Hz side-bands*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=45974>.
- [LVC12] G. Vajente. "NonStatMoni technical description". Virgo Internal note, VIR-0004A-08, <https://tds.virgo-gw.eu/ql/?c=1958>. 2008. URL: <https://tds.virgo-gw.eu/ql/?c=1958>.
- [LVC13] D. Verkindt. *VIR-0237A-15: Advanced Virgo DataDisplay tutorial*. 2015. URL: <https://tds.virgo-gw.eu/?content=3&r=11672>.
- [LVC14] S. Koley *et al*. *Correlation of noise lines drifting in frequency with Env channels*. 2017. URL: <https://logbook.virgo-gw.eu/virgo/?r=38249>.
- [LVC15] S. Koley. *VIR-0553A-17: Noise Hunting - Drifting Lines Analysis*. 2017. URL: <https://tds.virgo-gw.eu/ql/?c=12484>.
- [LVC16] F. Di Renzo. *Wandering line jumping from 83.5 Hz to 108 Hz*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=46952>.
- [LVC17] B. Swinkels. *VIR-0420A-18: Brute force correlation of drifting lines*. 2018. URL: <https://tds.virgo-gw.eu/?content=3&r=14433>.

- [LVC18] S. Kulkarni and M. Cavaglià. *NNETFIX: A neural network to “fix” Gravitational Wave signals overlapping with short-duration glitches in LIGO-Virgo data*. 2019. URL: <https://dcc.ligo.org/LIGO-G1802348>.
- [LVC19] A. Matas. *Non-stationarity cut for stochastic analyses in early O3*. 2019. URL: <https://dcc.ligo.org/LIGO-G1901254>.
- [LVC20] A. Matas. *Non-stationarity problems in O3 data*. 2019. URL: <https://stochastic-alog.ligo.org/aLOG/index.php?callRep=339844>.
- [LVC21] F. Di Renzo and G. Vajente. *VIR-0511A-19: BLRMS computation and stationarity tests update*. 2019. URL: <https://tds.virgo-gw.eu/?content=3&r=15638>.
- [LVC22] F. Di Renzo. *VIR-0676A-19: A new BLRMS based stationarity test, method description and results*. 2019. URL: <https://tds.virgo-gw.eu/?content=3&r=15869>.
- [LVC23] D. Verkindt. *Virgo Spectrogram tools*. 2019. URL: <https://tds.virgo-gw.eu/ql/?c=13835>.
- [LVC24] F. Di Renzo and N. Sorrentino. *VIR-0089A-20: BRISTOL and RAGOUT: two tools for investigating noise stationarity and Gaussianity, for Detchar studies and Data Analysis*. 2020. URL: <https://tds.virgo-gw.eu/?content=3&r=16715>.
- [LVC25] I. Fiori *et al.* *VIR-031B-07: Analysis of frequency lines in C7 run*. 2007. URL: <https://tds.virgo-gw.eu/?content=3&r=1888>.
- [LVC26] F. Di Renzo. *VIR-1171A-19: Flat noise studies - November 19, Commissioning meeting report*. 2019. URL: <https://tds.virgo-gw.eu/ql/?c=15005>.
- [LVC27] A. Magazzù, M. Vardaro, and F. Di Renzo. *Alignment of PR/BS/DIFFp*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=47549>.
- [LVC28] S. Mazzon, L.K. Nuttal, and A. Lundgren. *G1900224: PSD misestimation*. URL: <https://dcc.ligo.org/LIGO-G1900224>.
- [LVC29] N.J. Cornish. *Tests for Gaussianity and Stationarity*. 2018. URL: <https://dcc.ligo.org/LIGO-G1801886>.
- [LVC30] F. Robinet. *Omicron: a tool for detector characterization*. Tech. rep. Virgo note: VIR-0296A-12, <https://tds.virgo-gw.eu/ql/?c>, 2012.
- [LVC31] LIGO Detchar group. *LIGO detector status summary pages*. URL: <https://ldas-jobs.ligo.caltech.edu/~detchar/summary/>.
- [LVC32] D. Bersanetti. *VIR-015A-19: Feed-Forward Noise Cancellation in Virgo (talk Machine Learning Workshop - CA17137)*. URL: <https://tds.virgo-gw.eu/?content=3&r=15055>.
- [LVC33] T. Hardwick and Y. Michimura. *Comment to Further violin mode excitations: WI,NE,WE*. 2017. URL: <https://logbook.virgo-gw.eu/virgo/?r=41905>.
- [LVC34] F. Di Renzo. *NonNA: Non-stationary Noise Analysis tool update - LVC meeting, Joint LIGO-Virgo DetChar session*. 2019. URL: <https://dcc.ligo.org/LIGO-G1900516>.
- [LVC35] I. Fiori, B. Swinkels, and F. Di Renzo. *Line at 85 Hz now at 88 Hz and moving*. 2018. URL: <https://logbook.virgo-gw.eu/virgo/?r=43906>.
- [LVC36] I. Fiori, F. Paoletti, and F. Di Renzo *et al.* *Glitches around 50Hz in SSFS and DARM*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=44296>.
- [LVC37] F. Di Renzo. *NonNA web area*. 2018. URL: <https://scientists.virgo-gw.eu/DataAnalysis/NonNA/>.
- [LVC38] F. Di Renzo and M. Cieslar. *Glitches are back*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=47590>.
- [LVC39] A. Chiummo *et al.* *Glitches in the range and PMC alignment*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=47591>.
- [LVC40] F. Di Renzo and M. Cieslar. *Noise bumps between 155 and 170 Hz*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=47602>.
- [LVC41] A. Magazzù *et al.* *Operator Report - Afternoon Shift 17 February 2019*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=45116>.
- [LVC42] B. Swinkels *et al.* *Channel naming conventions for AdV signals*. Tech. rep. Virgo internal note. VIR-0233A-14, 2014.

- [LVC43] F. Di Renzo. *VIR-0573A-18: NonNA tools analysis of Virgo C10 data*. 2018. URL: <https://tds.virgo-gw.eu/ql/?c=13470>.
- [LVC44] F. Di Renzo. *VIR-0722A-18: NonNA tools analysis of C11 data*. 2018. URL: <https://tds.virgo-gw.eu/ql/?c=13619>.
- [LVC45] F. Di Renzo and L. Rei. *VIR-0918A-19: Application of Wavelets to Glitch investigations*. 2019. URL: <https://tds.virgo-gw.eu/ql/?c=14752>.
- [LVC46] F. Di Renzo and L. Rei. *VIR-0976A-19: Instantaneous causality via wavelet transform*. 2019. URL: <https://tds.virgo-gw.eu/ql/?c=14810>.
- [LVC47] F. Robinet and D. Verkindt. *Mystery glitches (BNS range drops)*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=46939>.
- [LVC48] D. Verkindt *et al.* *Detchar hardware injections*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=46900>.
- [LVC49] A. Chiummo and the ISYS team. *VIR-0194A-19: The Fast Unlock Saga: recent developments*. 2019. URL: <https://tds.virgo-gw.eu/ql/?c=14020>.
- [LVC50] D. Cohen and D. Arnaud. *LockLoss monitor*. 2019. URL: <https://wiki.virgo-gw.eu/DetChar/LockLossMonitor>.
- [LVC51] D. Verkindt and B. Swinkels. *Glitches around 150 Hz*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=45871>.
- [LVC52] L. Aiello, I. Nardecchia, and J. Casanueva. *ITF working point tuning*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=45945>.
- [LVC53] F. Di Renzo and D. Verkindt. *Band rms in interesting frequency bands*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=46411>.
- [LVC54] P. Ruggi. *Evaluation of noise level for different DARM offsets*. 2018. URL: <https://logbook.virgo-gw.eu/virgo/?r=44198>.
- [LVC55] M. Was. *VIR-1170A-19: Flat noise history and ideas*. 2019. URL: <https://tds.virgo-gw.eu/ql/?c=15004>.
- [LVC56] M. Was. *VIR-0803A-18: Flat noise investigations*. 2018. URL: <https://tds.virgo-gw.eu/ql/?c=13700>.
- [LVC57] A. Allocca *et al.* *VIR-0054A-19: "Flat" noise investigation*. 2019. URL: <https://tds.virgo-gw.eu/ql/?c=13877>.
- [LVC58] M. Was. *Flat noise in the bucket is increasing*. 2017. URL: <https://logbook.virgo-gw.eu/virgo/?r=40147>.
- [LVC59] J. Casanueva *et al.* *52 Mpc!!!!* 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=44544>.
- [LVC60] M. Was. *SDB1 alignment offset impact*. 2019. URL: [SDB1alignmentoffsetimpact](https://logbook.virgo-gw.eu/virgo/?r=44544).
- [LVC61] F. Di Renzo and L. Rei. *Detchar Shift: 08 Jan 2019 to 15 Jan 2019*. 2019. URL: https://wiki.virgo-gw.eu/DetChar/DetcharShift2019_Jan08_Jan15?validation_key=4ef2238d5db17ea063be4a2f22dc5cde.
- [LVC62] M. Mantovani *et al.* *84 Hz noise solved*. 2017. URL: <https://logbook.virgo-gw.eu/virgo/?r=38901>.
- [LVC63] M. Was and f. Di Renzo. *Wandering line moved to fixed 83 Hz?* 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=45037>.
- [LVC64] R. Bonnard *et al.* *Summary of SDB1 intervention*. 2019. URL: <https://logbook.virgo-gw.eu/virgo/?r=44446>.