

GRAVITATIONAL WAVES

Editor

N. Tamanini

Authors

Q. Baghi, W. Chaibi, F. Gulminelli, H. Halloin,
N. Korsakova, F. Marion, S. Marsat, S. Mastrogiovanni,
G. Pierra, J. D. Romano, S. Speziale, D. Steer, L. Suleiman,
G. Theureau, M. Trad-Nery, M. Turconi, N. Webb

June 18, 2026

DRAFT

Contents

Chapter 1. The theoretical foundations of gravitational waves: overview and linearized theory	13
Simone SPEZIALE and Danièle STEER	
1.1. Overview and characteristic scales	14
1.1.1. Aims	14
1.1.2. On wave-like solutions and relativity	15
1.1.3. Detectors and GW frequencies	17
1.1.4. Compact binary systems: orders of magnitude and characteristic scales	18
1.1.4.1. The chirp signal	20
1.1.4.2. Merger frequency	21
1.1.4.3. Time to merger	22
1.1.4.4. Amplitude and distance	22
1.1.4.5. Distance between objects at merger	23
1.1.5. Roadmap	24
1.2. Einstein's equations: general covariance, Noether's theorem and gauge transformations	24
1.2.1. Einstein's equations and general covariance	24
1.2.2. Noether's theorem and diffeomorphisms as gauge symmetries	29
1.2.3. Gauge and observers	33
1.3. Perturbative treatment of Einstein's equations	35
1.3.1. The idea: general background spacetime	35
1.3.2. Weak-field approximation	39
1.3.3. Spin and helicity	40
1.3.4. De Donder and TT gauges	46
1.3.5. Vacuum solutions	48

1.3.6. Gauge-invariant description: independent and constrained degrees of freedom	53
1.3.7. Gauge-fixing with sources	57
1.4. Appendix: Second order action for perturbations around any background solution	61
1.5. Appendix: Green's functions	68
1.6. Bibliography	69
Chapter 2. The theoretical foundations of gravitational waves: detection and emission	73
Simone SPEZIALE and Danièle STEER	
2.1. Detection of GWs	74
2.1.1. Coordinate displacements versus physical displacements	74
2.1.2. Interferometers	79
2.2. Generation of GWs from sources	81
2.2.1. Introducing sources	81
2.2.2. Source multipoles	84
2.2.3. Solving the wave equation with sources	86
2.3. Dissipation by gravitational waves	91
2.3.1. Energy of gravitational waves	91
2.3.2. Dissipation equations	96
2.3.3. On the validity of the quadrupole formula	98
2.3.4. Back-reaction	99
2.4. GWs from binary systems: elliptical, circular and hyperbolic orbits 100	
2.4.1. Newtonian equations	100
2.4.2. Energy and angular momentum fluxes	103
2.4.3. Back-reaction and waveform: circular orbits	109
2.4.4. Back-reaction: Elliptical orbits	111
2.4.5. Frequency content of elliptical orbits	113
2.4.6. Hyperbolic orbits	116
2.4.6.1. Linear memory effect and low-frequency GWs	117
2.4.6.2. Capture due to GW emission	120
2.5. GWs in curved space-time, cosmology	122
2.5.1. General background metric	122
2.5.1.1. Linearised equations	122
2.5.1.2. WKB approximation	123
2.5.2. FLRW metric: background	125
2.5.3. FLRW metric: gravitational waveforms	126
2.5.4. Measuring cosmological parameters with GWs: outline of GW cosmology	130
2.6. Acknowledgements	131
2.7. Appendix: Landau-Lifshitz approach	131
2.8. Bibliography	133

Chapter 3. Panorama of gravitational-wave sources 137

Joseph D. ROMANO

3.1. Introduction	138
3.2. Signal properties	138
3.3. Deterministic signals	139
3.3.1. Continuous wave sources	140
3.3.2. Compact binary coalescence	141
3.4. Stochastic signals and sources	144
3.4.1. Signal properties	144
3.4.2. Characterization of stochastic signals in terms of power spectra	146
3.4.3. Astrophysical sources	147
3.4.4. Cosmological sources	148
3.5. Acknowledgments	151
3.6. Bibliography	151

Chapter 4. Gravitational wave detectors from an experimental perspective 155

Marina TRAD-NERY, Margherita TURCONI and Walid CHAIBI

4.1. The noise for an experimental physicist	156
4.1.1. Noise definition	157
4.1.2. Noise characterization	157
4.1.3. Auto-correlation	158
4.1.4. Stationary noise	158
4.1.5. Temporal average and ergodicity	159
4.1.6. Stationary and ergodic process	161
4.1.7. Harmonic analysis: The Fourier transform	162
4.1.7.1. Finite energy process	162
4.1.7.2. Continuous process	162
4.1.7.3. The Wiener-Khintchine Theorem	163
4.2. Detection principle of an interferometric ground based gravitational wave detector	164
4.2.1. Gravitational wave effect on a light field	164
4.2.2. The Michelson interferometer, its output and tuning	167
4.3. Sensitivity enhancement with Fabry-Perot cavities	171
4.3.1. Intracavity power	173
4.3.2. Gaussian beams	176
4.3.3. Cavity transmission and mode cleaner cavities	178
4.3.4. Cavity reflectivity	179
4.3.5. Cavity as a frequency reference: feedback loop	181
4.4. Power recycled interferometer	182
4.4.1. Frequency stabilization	184
4.4.2. Power stabilization	186

4.4.3. Shot noise sensitivity curve	187
4.5. Other noise contributions	187
4.5.1. Length noise contributions	187
4.5.1.1. Harmonic oscillator	188
4.5.1.2. Seismic noise	190
4.5.2. Thermal noise	193
4.5.2.1. Suspension thermal noise	195
4.5.2.2. Mirror thermal noise	196
4.5.3. Quantum noise	197
4.5.3.1. Quantification of the electromagnetic field	197
4.5.3.2. Semi-classical approach : The power quantum fluctu- ations	198
4.5.3.3. The quantum noise in the interferometer	200
4.6. Conclusion: full sensitivity of a GWD	205
4.7. Bibliography	207
Chapter 5. Pulsar timing arrays (PTAs)	209
Gilles THEUREAU and Joseph D. ROMANO	
5.1. Pulsars as gravitational wave detectors	210
5.1.1. Pulsar timing array principle	210
5.1.2. First evidence for a low-frequency gravitational-wave signal	210
5.1.3. Expected source populations	211
5.1.4. Measurement challenges and foreground noise	213
5.1.5. Status of PTA results in 2025	216
5.2. Searches for a stochastic gravitational-wave background	218
5.2.1. Recovery of the predicted Hellings and Downs correlation	220
5.2.2. Detection confidence	222
5.3. Bibliography	223
Chapter 6. Laser Interferometer Space Antenna	227
Quentin BAGHI and Hubert HALLOIN	
6.1. The LISA mission concept	228
6.1.1. A little history	228
6.1.2. Description of the instrument	233
6.1.2.1. Mission profile and measurement principle	233
6.1.2.2. Overview of the Laser Interferometer Space Antenna (LISA) payload	239
6.1.2.3. LISA Pathfinder heritage	240
6.2. The science enabled by LISA	241
6.2.1. Landscape of gravitational wave sources in the LISA band	242
6.2.2. Astrophysics of compact objects	243
6.2.2.1. Compact binary stars	244
6.2.2.2. Black holes history	247

6.2.2.3. Black hole properties and environments	251
6.2.3. Fundamental physics: The nature of gravity and black holes	252
6.2.4. Cosmology	254
6.2.4.1. Measuring the Universe's expansion rate	254
6.2.4.2. Peering into the early Universe with stochastic GW backgrounds	255
6.3. From data to science	257
6.3.1. Noise reduction	257
6.3.2. Source detection and inference	258
6.3.2.1. Continuous sources	259
6.3.2.2. Transient sources	261
6.3.2.3. Stochastic sources	262
6.3.2.4. The global fit orchestration	263
6.4. Bibliography	264
Chapter 7. Introduction to gravitational wave data analysis	281
Sylvain MARSAT and Frédérique MARION	
7.1. Introduction	282
7.2. Notations and preliminaries	282
7.2.1. Fourier transforms: continuous and discrete domain	282
7.2.2. Gaussian distributions	283
7.3. The noise as a stochastic process	283
7.3.1. Stochastic processes	283
7.3.2. The power spectral density	284
7.3.3. The noise autocorrelation function	285
7.3.4. Stationary and Gaussian noise	287
7.3.5. Noise in discrete data	287
7.3.6. Noise-weighted inner product	290
7.3.7. PSD estimation	291
7.4. Matched filtering	292
7.4.1. The matched filter: signal and noise	292
7.4.2. The Wiener filter	293
7.4.3. The matched filter statistics	294
7.4.4. Matched filter SNR optimization	296
7.5. Template bank searches	298
7.5.1. Template banks	298
7.5.2. Towards searches in realistic noise	299
7.6. Bayesian inference	302
7.6.1. Bayesian statistics	302
7.6.2. Sampling from a posterior	304
7.6.3. Markov chain Monte Carlo methods	305
7.6.4. Sampling techniques and parallel tempering	307
7.6.5. Nested sampling	309

7.7. Parameter estimation for gravitational wave signals	309
7.7.1. The Whittle likelihood: stationary and Gaussian noise . . .	309
7.7.2. The Fisher matrix approach	310
7.7.3. Noise, errors and biases	311
7.7.4. An example parameter estimation result: GW150914	312
7.8. Other topics	313
7.8.1. Bursts and unmodelled signals	314
7.8.2. Continuous waves and long-lived signals	315
7.8.3. Stochastic backgrounds	317
7.8.4. Hierarchical inference	320
7.9. Future challenges in data analysis	323
7.10. Acknowledgements	324
7.11. Bibliography	324
Chapter 8. Artificial Intelligence for data analysis	327
Natalia KORSAKOVA	
8.1. Foreword	328
8.2. Basic machine learning background	329
8.2.1. Simple network architectures	330
8.2.2. Optimization function	332
8.2.3. Gradient descent	333
8.2.4. Backpropagation	334
8.3. Point parameter estimation	334
8.4. Bayesian parameter estimation	339
8.4.1. Normalizing flows	341
8.4.1.1. Objective function	342
8.4.1.2. Conditioning of the flow on the simulated data	343
8.4.1.3. Example with Galactic Binaries for LISA	343
8.5. Bibliography	344
Chapter 9. The astrophysics of gravitational wave sources	347
Natalie WEBB	
9.1. Introduction	348
9.2. Compact objects and related phenomena as gravitational wave sources	348
9.3. White dwarfs	349
9.3.1. Formation of white dwarfs	350
9.3.1.1. Electron degeneracy	350
9.3.2. White dwarf masses and radii	354
9.3.2.1. The Chandrasekhar mass	354
9.3.3. White dwarf structure and matter	356
9.3.4. Observing and understanding white dwarfs	357
9.4. Neutron stars	358

9.4.1. Formation of neutron stars	359
9.4.1.1. The neutron drip and baryonic degeneracy	359
9.4.2. Neutron star structure and matter	360
9.4.3. The equation of state	362
9.4.4. Neutron star masses and radii	363
9.4.4.1. Mass	363
9.4.4.2. Radius	365
9.4.5. Observing and understanding neutron stars	366
9.5. Black holes	367
9.5.1. Formation of black holes	367
9.5.2. Black hole masses and radii	369
9.5.3. The mass	369
9.5.4. The radius	372
9.5.5. Black hole structure	374
9.5.5.1. Static black holes	374
9.5.5.2. Rotating black holes	374
9.5.6. Observing and understanding black holes	375
9.6. Future observations	376
9.7. Bibliography	377
Chapter 10. Gravitational Wave Cosmology: an introduction	385
Simone MASTROGIOVANNI and Gregoire PIERRA	
10.1. Introduction	386
10.2. Background for standard sirens cosmology	386
10.2.1. An introduction to the cosmic expansion	387
10.2.2. Compact binary coalescence at cosmological distance	394
10.3. Bayesian inference for cosmological and population properties of resolved sources	400
10.3.1. A brief recap of Bayesian theory	400
10.3.2. Hierarchical Bayesian Inference with gravitational waves	403
10.4. Spectral sirens: Cosmology with GW sources and their source mass spectrum	406
10.5. Adding galaxy surveys to dark sirens	411
10.6. Bright sirens: The “easiest” case	416
10.7. Conclusions	421
10.8. Bibliography	421
Chapter 11. Gravitational waves as a probe of dense matter	425
Lami SULEIMAN and Francesca GULMINELLI	
11.1. Introduction to dense matter physics in the Universe	426
11.1.1. Motivation	426
11.1.2. Sites of ultra-dense matter	428
11.1.3. Neutron star observations	430

11.2. Neutron star modelling and the equation of state	431
11.2.1. Equation of state	431
11.2.2. Modeling the core and the crust of neutron stars	433
11.2.3. Effective models of the strong interaction	437
11.2.4. EoS and astrophysical observables	439
11.2.5. Bayesian inference of the EoS	442
11.3. The gravitational wave probe for dense matter	444
11.3.1. Relativistic tidal deformation in binary merger	445
11.3.2. Impact of the tidal deformation on the gravitational wave signal	447
11.3.3. Observations of the tides	449
11.4. Prospects for next generation of GW detectors	452
11.4.1. Detecting the tidal deformability with precision	453
11.4.2. Beyond tidal deformation to constrain dense matter physics	455
11.4.3. Challenges in the analysis	456
11.5. Bibliography	457

Foreword

Nicola Tamanini

*Laboratoire des 2 infinis - Toulouse (L2IT), Université de Toulouse,
CNRS/IN2P3, Toulouse, France*

Here goes the forword...

DRAFT

1

The theoretical foundations of gravitational waves: overview and linearized theory

Simone Speziale¹ and Danièle Steer²

¹*Aix Marseille Univ., Univ. de Toulon, CNRS, CPT, UMR 7332, 13288
Marseille, France*

²*Laboratoire de Physique de l'ENS, Université Paris Cité, Ecole Normale
Supérieure, Université PSL, Sorbonne Université, CNRS, 75005 Paris*

1.1. Overview and characteristic scales

1.1.1. Aims

Indirect evidence of gravitational waves (GWs) has been available for many years (see e.g. (Will 2014) for a review). Direct detection on the other hand had to wait until 2015 (Abbott *et al.* 2016), that is nearly 100 years after GWs were predicted to exist in general relativity (GR). These detections by the LVK collaboration, consisting of the network of LIGO (Aasi *et al.* 2015), Virgo (Acernese *et al.* 2015) and more recently KAGRA (Akutsu *et al.* 2021) GW interferometers, are ongoing with new GWs signals being observed on a weekly basis (LIGO n.d.). In the future more sensitive detectors on earth, together with ones working in different frequency bands such as the Laser Interferometer Space Antenna (LISA) (Colpi *et al.* 2024) as well as Pulsar Timing Arrays (PTAs), will lead to new observations of the universe, potential new discoveries, and unprecedented tests of general relativity, cosmology and astrophysics.

The aim of these first two chapters is to present the basic introductory material required to understand GWs. We will address some of following questions:

- What are GWs? How do they emerge from GR? How does one deal with the symmetries (diffeomorphism invariance) of GR to fix gauges and coordinates, and what do they imply for the stress energy tensor of GWs?
- To what GW frequencies f_{GW} are current and future GW detectors sensitive? Why are those detectors designed to be sensitive to particular GW frequency ranges?
- For a source consisting of two *bound* compact binaries objects (such as black holes) of masses m_1 and m_2 at some distance R from an observer, what is the characteristic frequency, amplitude etc of the GWs emitted? Up to what distances R can such sources be detected?
- Using the quadrupole formula (which we derive) what is the waveform of the emitted GWs and how does it depend for example on the ellipticity of the bound orbit?
- What sources correspond to the GW events detected by LVK? Are there other possible GW sources? We give an example of compact binary sources on unbound orbits and discuss the GW memory effect.
- If we consider sources on cosmological distance scales, how are their amplitude, frequency e.t.c. affected by the cosmological expansion?

In particular, the first chapter will contain a general overview, and the theoretical aspects of gravitational waves. The second chapter will discuss detection and emission from binary systems, addressing most of the questions above.

1.1.2. *On wave-like solutions and relativity*

Gravitational waves are a natural expectation from GR, simply because it is a relativistic theory of gravity. To understand why, let us first take a step back to *non-relativistic* Newtonian gravity. When the famous apple drops on Newton's head, the mass distribution of the Earth changes, and so does the gravitational field created. In Newton's time, this variation, however negligible, was assumed to be the effect of some *instantaneous* "action at a distance". After the discovery that the *speed of light is finite*, and that all effects in our universe appear to follow this causal limitation, it seems natural to expect that also the variations of the gravitational field will *not* be felt instantaneously in the whole universe, but will rather be propagated at the speed of light — or less.¹ *The propagation of this perturbation of the gravitational field is intuitively what we call a gravitational wave.*

Conceptually, a gravitational wave is similar to a water wave or an electromagnetic (EM) wave. However, while those propagate a modification in the depth of water or the intensities of the electromagnetic field, a GW propagates a modification of the structure of spacetime itself. As for producing one, it is natural to expect that the Earth emits GWs when orbiting the sun, thus carrying away energy and making the orbit decay, just like a charged particle emits EM waves when accelerating.

In practise, however, understanding GWs is very subtle for a number of reasons. First of all, the concept itself of propagation makes reference to a background spacetime, and in GR there is no fixed background structure. Splitting the dynamical spacetime into a reference background and a perturbation on top of it is a delicate process in which potential ambiguities have to be dealt with. In fact many decades passed before a consensus was reached, to the point that, famously, Einstein himself initially doubted the physical existence of GW, for reasons that we will briefly review and clarify below. Secondly, if we think of GWs as waves propagating in a medium, this medium is extraordinarily rigid: the waves go as fast as possible and have very *tiny* amplitudes. To give an idea, the power emitted by the Earth-Sun system in the form of GW is around

1. We will see that Einstein's relativistic theory of gravity, GR, predicts that — whatever their wavelength — these variations propagate at exactly the speed of light, and if some future experiment shows that they propagate at a lesser speed, then this would be an explicit violation of GR.

200W! This rigidity has to do with the weakness of the gravitational coupling constant. One may think that gravity is strong when for e.g. trying to beat a high jump record, or when skydiving, but this strength is ridiculously small compared to the much much stronger electro-magnetic force that dominates our daily life. These two points — background independence and weakness of the signal — are typical issues that one has to face when studying GWs.

Indeed only very massive and energetic objects can produce GWs of amplitudes that are actually detectable. Amongst the most massive and compact astrophysical objects known are black holes (BH), neutron stars (NS) and white dwarfs (WD). The GW sources detected to date by the LVK collaboration are all ‘compact binary systems’ made of a bound pair of BH and/or NS on closed orbits. As a result of the energy lost through GW emission, the two bodies making up the bound system approach closer to each other, inspiralling inwards, and eventually merging into one final object. In fact the GW signals detected by the LVK collaboration correspond to the last moments in the life of these systems including their merger — they are known as ‘compact binary coalescences’ (CBC’s). For comparison with the earth-sun system mentioned above, the energy emitted in GWs by the very first detected GW event GW150914 (Abbott *et al.* 2016), which was due to the coalescence of two BHs of masses $m_1 \sim 36M_\odot$ and $m_2 \sim 29M_\odot$, was almost 10^{48} Joules in 0.2 seconds.

The direct detection of GWs can be used to test many aspects of gravity, for instance in the strong field regime, see e.g. (Yunes and Siemens 2013), as well as to probe cosmology, as will be discussed later. Indeed, the gravitational interaction is so weak that the universe is almost completely transparent to a gravitational wave. As a consequence, one can potentially collect pristine information about any cosmological era through GWs, and in particular through the detection and characterisation of a stochastic gravitational wave background. Sources relevant to cosmology include primordial GWs produced during inflation but there are also potential new sources to be discovered, such as primordial black holes, cosmic strings, and other exotic objects, see e.g. (Caprini and Figueroa 2018 ; Caprini *et al.* 2025) for reviews.

The aim of these first two chapters is not to provide an introduction to the broad set of fascinating GW sources, confirmed or hypothetical, nor to the many creative ideas to detect them that have been proposed, investigated and realised in practise; but only to provide an introduction to the field, and to that end, we decided to focus on the most common type of sources, and most common type of detectors: CBCs and laser interferometers. In the rest of this overview section we review the characteristic properties of GWs emitted by CBCs and the relevant frequency bands of laser interferometers, in particular explaining why LVK detectors are sensitive to the merger of stellar mass BH,

whilst LISA for example to that of supermassive BHs. The rest of the chapter will present the theoretical derivation of GWs from GR.

1.1.3. Detectors and GW frequencies

The LVK interferometers and future LISA detector are essentially *Michelson-Morley interferometers*, designed to be as sensitive as possible to time-varying changes in the separation between two freely falling test-masses — mirrors in the case of interferometers. The invariant distance between the test masses varies when a GW passes (see Section 2.1), leading to a change in the observed interference pattern in the detector.

- The LVK interferometers are on earth (in Livingston and Handford in the USA, in Pisa in Europe, and in Kamioka in Japan) and have a typical arm length $L \sim 3\text{km}$. They are sensitive to GWs with frequency of order

$$10\text{Hz} \lesssim f_{\text{GW}} \lesssim 5\text{kHz} \quad (\text{LVK}). \quad [1.1]$$

- The LISA interferometer (Colpi *et al.* 2024) was adopted by ESA on the 25th january 2024, and should be operational in 2037. The distance between the spacecraft which make up arms of LISA is $L \sim 2.5 \cdot 10^6\text{km}$. LISA will be sensitive to GWs with frequencies in the range

$$10^{-4}\text{Hz} \lesssim f_{\text{GW}} \lesssim 1\text{Hz} \quad (\text{LISA}). \quad [1.2]$$

- There are plans to build new interferometers on earth beyond LVK. These include the Einstein Telescope in Europe (Sathyaprakash *et al.* 2012) and Cosmic Explorer in the USA (Reitze *et al.* 2019), both of which should have $L \sim 10\text{km}$, and

$$\text{few Hz} \lesssim f_{\text{GW}} \lesssim 10^4\text{Hz} \quad (\text{ET, CE...}). \quad [1.3]$$

- An alternative to interferometers are PTAs which search for GWs by exploiting the variation in distance $L \sim 10^{17}\text{km}$ between the earth and a typical distant galactic pulsar due to GWs. Pulsars emit EM pulses with extreme regularity Δt , typically of the order of milliseconds. If GWs are present, then as the EM pulses propagate from the pulsar to the earth, the observed Δt_{obs} will be modulated. PTA experiments searching for these modulations are sensitive to GWs in the frequency band of the inverse year,

$$10^{-7}\text{Hz} \lesssim f_{\text{GW}} \lesssim 10^{-9}\text{Hz} \quad (\text{PTA}). \quad [1.4]$$

In 2023 different PTA experiments presented strong evidence for the existence of a stochastic GW background (Antoniadis *et al.* 2023 ; Agazie *et al.* 2023 ; Reardon *et al.* 2023 ; Xu *et al.* 2023).

- Prior to the success of interferometers, there was an effort pioneered by Weber in the 60's to use resonant bars to detect GWs, building material bars whose acoustic modes would resonate at a frequency as near as possible to that expected from the optimal sources (Weber 1960). These experiments would typically have a narrow-band sensitivity around 10^3Hz . In spite of constant experimental evolution throughout the 90's, no observation has occurred in this way.

Table 1.1 summarises the different characteristics of the existing experiments, and in particular the ratio of their characteristic size L to the GW wavelength $\lambda_{\text{GW}} = c/f_{\text{GW}}$.

	Characteristic detector size (km)	GW frequency detectability range (Hz)	$f_{\text{GW}}L$	L vs λ_{GW}
LVK	~ 1	$10^1 - 10^4$	$f_{\text{GW}}L \ll 1$	$L \ll \lambda_{\text{GW}}$
LISA	$\sim 10^6$	$10^{-4} - 10^{-1}$	$f_{\text{GW}}L \sim 1$	$L \sim \lambda_{\text{GW}}$
PTA	$\sim 10^{17}$	$10^{-9} - 10^{-7}$	$f_{\text{GW}}L \gg 1$	$L \gg \lambda_{\text{GW}}$

Table 1.1: Characteristics of different GW detectors and the corresponding GW wavelength.

1.1.4. Compact binary systems: orders of magnitude and characteristic scales

LVK and LISA were conceived in order to be sensitive to the particular range of frequencies that are not only within experimental reach, but also that are likely to constitute a rich source according to the known astrophysical data. Amongst those GW sources are compact binary systems. We now focus on orders of magnitude and characteristic scales for such compact binary system, consisting of two masses $m_{1,2}$ at a distance R from the detectors, see figure 1.1. The expressions given here will be derived later in section 2.4. Furthermore, the expansion of the universe, neglected here, is considered in Section 2.5.

As shown in figure 1.2, as a consequence of GW emission, the two masses $m_{1,2}$ approach each other — the *inspiral phase* — until they merge — the *merger phase* — and form a single object. This object will keep radiating GWs, in the so-called *ringdown phase*, until it settles down to an equilibrium state (which for BHs is expected to be represented by the Kerr or Schwarzschild




Figure 1.1: Sketch of a binary system of masses $m_{1,2}$ with conserved orbital angular momentum \vec{J} and inclination ν , at a distance R from different detectors (LVK, ET and LISA). The approximate frequency bands of each detector are indicated.

solutions, according to theoretical and numerical evidence) after which no further emission occurs. The typical corresponding waveform, related to the GW amplitude, is shown in figure 1.2 as a function of time.

- The *inspiral phase* can be understood with perturbation theory (the “post-Newtonian (PN) expansion” of the Einstein equations) presented below, more details in (Thorne 1980 ; Blanchet 2006 ; Poisson and Will 2014).
- The *merger phase* generally requires numerical relativity, or other techniques such as effective one-body techniques, see e.g. (Deruelle and Uzan 2018) for an introduction. These techniques can also be applied to the late inspiral phase, in particular to understand accurately the GW signal as the merger is reached.
- The *ringdown phase* can also be approached with perturbative methods, namely BH perturbation theory, see e.g. (Kokkotas and Schmidt 1999).




Figure 1.2: Sketch of the inspiral, merger and ringdown phases of a CBC with their corresponding gravitational waveform as a function of time. There is no analytic method that can reproduce this signal entirely, and different approximation schemes are used. Linearized GR and the PN expansion that will be explained here can be used for the initial inspiral phase, and extrapolated to get a first estimate of the merging amplitude.

1.1.4.1. The chirp signal

During the inspiral phase the GW frequency increases with time according to the well-known *chirp* signal. Using the dominant quadrupolar mode contribution for point masses m_1 and m_2 (with spins set to zero), and assuming circular orbits, the time dependence of the frequency is given by

$$f_{\text{GW}} = \frac{1}{\pi} \left(\frac{GM}{c^3} \right)^{-5/8} \left(\frac{5}{256\tau} \right)^{3/8} \quad [1.5]$$

see Eq. [2.101], section 2.4.3. Here the *chirp mass* is

$$\mathcal{M} \equiv \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}} \quad [1.6]$$

and

$$\tau = t - t_c \quad [1.7]$$

is the time to coalescence, with t_c the coalescence time. Clearly Eq. [1.5] will break down before $\tau = 0$ where formally f_{GW} diverges. We thus define a “merger time” $t_{\text{merger}} < t_c$ up to which Eq. [1.5] is assumed valid, and whose meaning we now discuss.

1.1.4.2. Merger frequency

We now assume that the two objects are Schwarzschild BHs, and that merger occurs at the innermost stable circular orbit (ISCO) namely a distance $a = 6GM/c^2$ with $m = m_1 + m_2$. It then follows from Keplers laws (see Sec. 2.4) together with Eq. [1.5] that

$$f_{\text{merger}} = \frac{1}{6^{3/2}\pi} \left(\frac{c^3}{Gm} \right). \quad [1.8]$$

(Note given a length scale a and a mass m , $\sqrt{Gm/a^3}$ has dimensions of frequency. Setting $a = 6Gm/c^2$ gives, modulo factors of 2π , Eq. [1.8].)

- For a binary neutron stars (**BNS**) system, with say $m_{1,2} \sim 1.4M_\odot$ then Eq. [1.8] gives

$$f_{\text{merger}} \simeq 1.5\text{kHz} \quad (\text{BNS}) \quad [1.9]$$

This is in the upper part of the LVK frequency band.

- For a stellar mass binary black hole (**BBH**) system with for instance $m_{1,2} \sim 35M_\odot$,

$$f_{\text{merger}} \simeq 60\text{Hz} \quad (\text{stellar mass BBH}). \quad [1.10]$$

This is right in the frequency band of LVK.

- For a massive black hole binary (**MBHB**) system with for instance $m_{1,2} \sim 10^6M_\odot$

$$f_{\text{merger}} \simeq 10^{-3}\text{Hz} \quad (\text{supermassive Binary BHs}) \quad [1.11]$$

which is in the frequency band of LISA.

- Notice that PTA frequencies do not correspond to the merger frequency of any know astrophysical system. Rather, they correspond to the inspiral phase of super MBHB at times much before merger, as can be seen from Eq. [1.8]. Hence these are on broad orbits, with periods of the order of years.

Detailed figures for the Ligo-Virgo, CE and ET sensitivities as a function of frequency can be found for instance in (Maggiore *et al.* 2024). LISA sensitivities can be found for instance in (Colpi *et al.* 2024).

1.1.4.3. Time to merger

If GWs emitted during the inspiral enter the frequency band of a given detector at frequency f_{low} , then it is straightforward to integrate Eq. [1.5] from f_{low} to f_{merger} to find the total duration of the GW signal as will be observed by the experiment. Assuming $f_{\text{merger}} \gg f_{\text{low}}$ for simplicity, one finds that the total duration of the signal is

$$T \sim 10^{-3} f_{\text{low}}^{-8/3} \left(\frac{c^3}{GM} \right)^{5/3} \quad [1.12]$$

- For BNS entering the LVK band with $f_{\text{low}} \sim 20\text{Hz}$, this gives $T \sim 4$ minutes.
- For BNS entering the ET band with $f_{\text{low}} \sim 1\text{Hz}$, then $T \sim 5$ days. (This implies for example that effects of the rotation of the earth cannot be neglected when calculating the GW properties in more detail, see e.g. (Iacovelli *et al.* 2022) and references within. Also one might expect other GW signals to be produced in such a long period, overlapping with the BNS one. This makes data analysis more complex (Samajdar *et al.* 2021).)
- For stellar mass BHs, with say $m_{1,2} \sim 35M_{\odot}$ entering the LVK band with $f_{\text{low}} \sim 20\text{Hz}$, then $T \sim 0.1$ seconds.
- For stellar mass BHs, with say $m_{1,2} \sim 35M_{\odot}$ entering the ET band with $f_{\text{low}} \sim 1\text{Hz}$, then $T \sim 300$ seconds.
- For MBHB with $m_{1,2} \sim 10^6 M_{\odot}$ entering the LISA band with $f_{\text{low}} \sim 10^{-4}\text{Hz}$, then $T \sim 1$ month. (The orbital motion of LISA will thus also be non-negligible and e.g. Doppler effects must be considered. Furthermore other LISA sources will overlap with the MBHB signal.)

1.1.4.4. Amplitude and distance

The dimensionless amplitude of the GW signal scales with distance R to the source and GW frequency f_{GW} as

$$h \sim \frac{4}{R} \left(\frac{GM}{c^2} \right)^{5/3} \left(\frac{\pi f_{\text{GW}}}{c} \right)^{2/3} \quad [1.13]$$

As an example, consider say stellar mass BBH with $m_{1,2} \sim 35M_{\odot}$ for which $f_{\text{merger}} \sim 60\text{Hz}$. In order to generate (at merger) a signal with amplitude $h \sim 10^{-21}$, which is accessible to LVK, requires

$$R \sim 400 \text{ Mpc} \quad [1.14]$$

which is of the order of cosmological scales (for comparison, the observable universe has a scale of $c/H_0 \sim \text{Gpc}$, where H_0 is the Hubble constant).

Clearly from Eq. [1.13], given \mathcal{M} and f_{GW} , the more sensitive a detector, namely the smaller h can be detected, the further one can detect a given GW source. The “detection volume” of LVK has been steadily increasing with the different observing runs of LVK obviously leading to increasing numbers of detected GW events.

Notice that if such a GW signal is detected, then from the time dependence of the GW frequency one can directly obtain chirp mass Eq. [1.5]. With that, from the amplitude one can obtain the distance through Eq. [1.13]. Distance measurements can thus directly be obtained with GW observations from binaries, hence their name *standard sirens* (Schutz 1986 ; Holz and Hughes 2005). This should be contrasted with the case of EM observations (*standard candles*) for which the determination of the distance is particularly difficult. See Sec. 2.5 for more information about GWs as distance indicators and their use in cosmology.

1.1.4.5. Distance between objects at merger

When GWs are emitted with frequency f_{GW} , the two bodies in the compact binary are separated by a characteristic scale

$$r \sim \left(\frac{Gm}{f_{\text{GW}}^2} \right)^{1/3} \quad [1.15]$$

(see also the discussion after Eq. [1.8]). Since, from Eq. [1.5], the GW frequency increases during inspiral, the distance r between the two bodies decreases. The minimum distance is at the merger frequency f_{merger} . For example, for stellar mass BBH with $m_{1,2} \sim 35M_{\odot}$ and $f_{\text{merger}} \sim 60\text{Hz}$ then from Eq. [1.15] $r \sim \mathcal{O}(100)\text{km}$.

A distance $r \sim \mathcal{O}(100)\text{km}$ is tiny compared to the characteristic size of a star. Some of the most dense stars in the universe — for instance WDs — have a size $\sim 10^3\text{km}$. Main sequence stars have a size which can go up to millions of km. Thus if GW signals are seen from objects which reach minimum approach

distances $\sim \mathcal{O}(100)\text{km}$, those objects cannot be stars as they would already have collided. We must be dealing with BH (or possibly NS) for which the minimum distance will be determined by the Schwarzschild radius.

1.1.5. Roadmap

Having gone through the overview and discussed these orders of magnitude, the remainder of this chapter aims to derive formal results on GWs starting from Einstein's equations. Many introductions and reviews on GWs already exist, see for instance (Poisson and Will 2014 ; Maggiore 2007, 2018 ; Blanchet 2006 ; Andersson 2019 ; Deruelle and Uzan 2018) to mention a few. It is a rich and intricate topic, and each of these reviews tends to have a different angle on it, whose mutual compatibility may not always be clear to somebody entering the field. These first two chapters strive at presenting the material in a way that allows one to understand how the different approaches relate to one another. They also strive to spend time on some of the subtleties and delicate conceptual aspects of GR and GWs which are often left to the side in gravitational wave introductions, such as gauge dependencies, asymptotic charges and memory effects, and which are becoming more and more relevant as theoretical research and experiments advance into more accurate comparisons. In these first two chapters, a reader will therefore find discussions of questions of such as coordinate invariance, diffeomorphisms and gauge invariance, spin and helicity, the controversies about the stress energy tensor of GWs, GW memory effects, and Noether charges. The hope is that although most of the more advanced material is not needed for a first introduction, its inclusion here will stimulate the reader, and provide a useful reference for delving further into the topic.

1.2. Einstein's equations: general covariance, Noether's theorem and gauge transformations

1.2.1. Einstein's equations and general covariance

Einstein's great discovery about gravitation was that it can be understood as the manifestation of the curvature of spacetime. In Wheeler's words, spacetime tells matter how to move, matter tells spacetime how to bend.² Understanding gravity as a dynamical spacetime has changed profoundly our understanding of inertia. If we go back to Galilean relativity, an inertial observer is defined as one moving on a straight line at constant velocity. Special relativity introduces a

2. While pictorially charming, this statement is not exactly true: spacetime can be extraordinarily bent even in the absence of matter, as black hole solutions show.

non-trivial mixing of space and time, but leaves this notion unaffected: Inertial observers are still moving on a straight line, even though they are now related by Poincaré transformations as opposed to Galilean transformations, so to account for the experimental invariance of the speed of light. But in a curved spacetime, straight lines may no longer exist. The notion that encompasses them is the one of geodesics, which describe free-falling observers. Constant motion on a straight line is simply the flat-spacetime version of free falling. The understanding offered by general relativity thus has the merit of not only explaining gravity, but also explaining the origin of inertia. On the other hand, it changes the perspective on it radically: You reading these notes at your desk are inertial in Newton's terms, but accelerated in Einstein's, since you are being held by the ground against Earth's gravitational attraction and not following a geodesic.

Another profound consequence of a dynamical spacetime metric is that the field equations of gravity and matter are *covariant under general coordinate transformation*, as we will review below, introducing a new paradigm that goes under the name of *principle of general covariance*. In a curved spacetime, there are no more preferred Cartesian coordinates, no more Poincaré transformations relating inertial observers, and familiar physical concepts such as time evolution and energy become surprisingly subtle. These aspects of GR are often glossed over in lectures aiming at introducing gravitational waves, where one can blissfully rely on the background spacetime introduced by the weak field approximation and ignore most of them. However we believe they are important in order to better appreciate some of the properties of gravitational waves, provide an understanding that is more conceptual and less application-driven. We will briefly review these aspects below, and use them as benchmark to discuss some conceptual aspects of gravitational waves. For instance, the lack of preferred clocks in a curved spacetime is relieved in the weak field approximation, where one can use the flat Minkowski background to introduce a class of Cartesian observers, and select their proper time as preferred time. But the lack of well-defined notion of energy density is a subtlety that persists also in the weak-field approximation, and has to be dealt with.

Let us start by recalling Einstein's equations

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}, \quad [1.16]$$

where $G_{\mu\nu} := R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu}$ is the Einstein tensor with $R_{\mu\nu}$ the Ricci tensor and R the Ricci scalar, and $T_{\mu\nu}$ is the (symmetric) stress-energy tensor of matter. We use the definitions and conventions of (Poisson and Will 2014), in particular mostly-plus convention for the spacetime metric $g_{\mu\nu}$. The constants

G/c^4 and Λ are respectively the relativistic gravitational coupling constant and the cosmological constant. The first can be determined from local gravitational experiments to be

$$\frac{8\pi G}{c^4} \simeq 10^{-43} \text{ kg}^{-1} \text{ m}^{-1} \text{ s}^2. \quad [1.17]$$

This value is stupendously small, and it is the origin of the ‘rigidity’ of space-time mentioned in the overview section. The smallness of this parameter has, on the other hand, a positive side: the gravitational force is so weak that many of the observed phenomena, and virtually all solar system experiments, can be studied using the *weak field approximation*, namely a perturbative expansion around the Minkowski metric. This is quite helpful because Einstein’s equations are non-linear and it is in general very difficult to find exact solutions. Strong gravity effects occur only near very compact objects, and to study them one has to resort to numerical techniques, or be able to push the perturbative treatment to high orders.

The cosmological constant Λ can be determined from the observed acceleration of the expansion of the universe assuming homogeneity and isotropy on large scales, and turns then out to be $\Lambda \simeq 10^{-52} \text{ m}^{-2}$. This coupling constant can also be interpreted as a sort of averaged ‘vacuum energy’ density, often referred to as *dark energy* since it is not associated to visible matter, and whose value is $\rho_{DE} = \Lambda c^2/G \sim 10^{-28} \text{ kg/m}^3$. The presence of Λ affects the propagation of gravitational waves on cosmological distances, but it can be ignored for a first understanding of the perturbative treatment. We will set $\Lambda = 0$ for now, and restore it below in Section 2.5 when discussing cosmological effects.

Analysis of the 10 field equations in Eq. [1.16] shows (see subsection 1.3.6) that: four are redundant, because of the Bianchi identities; four are elliptic, hence describe gravitational degrees of freedom constrained by the sources; two are hyperbolic, hence contain independent degrees of freedom. This three-sided structure is a common feature to Maxwell and Yang-Mills theories, with the role of the Gauss constraint generating gauge transformation replaced by the so-called Hamiltonian and vector constraints generating diffeomorphisms, and it is our first indication that coordinate transformations are a gauge symmetry. A second indication comes from Noether’s theorem, but before talking about it, let us review how coordinate transformations act.

Recall that a tensor is a quantity that transforms homogeneously under general coordinate transformations $x^\mu \rightarrow x'^\mu(x^\nu)$. For instance a scalar field, a vector field and the metric transform respectively as

$$\phi'(x') = \phi(x), \quad v'^\mu(x') = \frac{\partial x'^\mu}{\partial x^\nu} v^\nu(x), \quad g'_{\mu\nu}(x') = \frac{\partial x^\rho}{\partial x'^\mu} \frac{\partial x^\sigma}{\partial x'^\nu} g_{\rho\sigma}(x). \quad [1.18]$$

The transformation law of a scalar is such that its value at one point P is the same after the diffeomorphism, since both x and x' identify the same point, just in different coordinates.³ The vector and metric do the same, but furthermore their indices are mixed up using the Jacobian of the coordinate transformation, or its inverse. A vector field is said to transform as a contravariant tensor of order one, and the metric as a covariant tensor of order two. Since coordinate transformations are typically restricted to be differentiable, namely continuous and connected to the identity, they are also invertible, and correspond to mathematical transformations called diffeomorphisms. In this language, [1.18] is a diffeomorphism of the metric.

If the coordinate transformation is infinitesimal, we can write it as $x'^{\mu} = x^{\mu} + \xi^{\mu}(x)$, and approximate the transformation rules in [1.18] using the Taylor expansion (applied to both the field's argument and the Jacobian). This defines the infinitesimal transformations

$$\delta_{\xi}\phi := \phi(x) - \phi'(x) = \xi^{\mu}\partial_{\mu}\phi \equiv \mathcal{L}_{\xi}\phi, \quad [1.19]$$

$$\delta_{\xi}v^{\mu} := v^{\mu}(x) - v'^{\mu}(x) = \xi^{\nu}\partial_{\nu}v^{\mu} - v^{\nu}\partial_{\nu}\xi^{\mu} \equiv \mathcal{L}_{\xi}v^{\mu}, \quad [1.20]$$

$$\delta_{\xi}g_{\mu\nu} := g_{\mu\nu}(x) - g'_{\mu\nu}(x) = \xi^{\rho}\partial_{\rho}g_{\mu\nu} + 2g_{\rho(\mu}\partial_{\nu)}\xi^{\rho} = 2\nabla_{(\mu}\xi_{\nu)} \equiv \mathcal{L}_{\xi}g_{\mu\nu}, \quad [1.21]$$

where we introduced the use of round brackets for index symmetrization (and we will later on also use square brackets for index anti-symmetrization). Notice that in all cases we recover as infinitesimal transformation the Lie derivative. This is a general result valid for any tensor. The third equality in [1.21] is on the other hand special to the metric tensor, and follows from the expression of the connection in terms of the metric.

Being written in terms of tensors, Einstein's equations are automatically covariant under general coordinate transformations. This is the principle of general covariance, that played a key role in guiding Einstein to formulate his theory. One immediate implication is that locally we can always find coordinates such that the metric takes the Minkowski expression at a point, which is one version of the principle of equivalence. A more subtle implication is that coordinate transformations must be *symmetries* of the theory, in other words a

3. This is sometimes misstated by saying that scalars are invariant under coordinate transformations, which is not true. A quantity is invariant under coordinate transformations if it satisfies the stronger property that its value does not depend on the coordinates used, which for the scalar field would be the equation $\phi(x') = \phi(x)$. This is not true in general, but only for isometries — more on this below. A typical example of coordinate invariance is the integral over the whole manifold of a scalar times the volume form, as experience from solving integrals via change of coordinates should show.

solution can be equivalently written in any coordinate system. To understand this point, let us consider the Lagrangian description of the dynamics. The field equations [1.16] are Euler-Lagrange equations of $\mathcal{L} = \mathcal{L}_{\text{EH}} + \mathcal{L}_{\text{M}}$, where

$$\mathcal{L}_{\text{EH}} = \frac{c^3}{16\pi G}(R - 2\Lambda)\sqrt{-g} \quad [1.22]$$

is the Einstein-Hilbert Lagrangian density, and \mathcal{L}_{M} the matter contribution, left arbitrary for the moment. Here $g = \det(g_{\mu\nu})$, and the word density has a double meaning: in the physical sense, since $c\mathcal{L}$ has the dimensions J m^{-3} of an energy density, but also in the mathematical sense, since $\sqrt{-g}$ makes it transform not as a scalar but as a scalar density of weight 1. This has the following consequence. Recall that the general formula for the variation of a determinant is $\delta g = gg^{\mu\nu}\delta g_{\mu\nu}$. This implies that $\mathcal{L}_\xi\sqrt{-g} = \frac{1}{2}\sqrt{-g}g^{\mu\nu}\mathcal{L}_\xi g_{\mu\nu} = \sqrt{-g}\text{NewA}_\mu\xi^\mu$, whence $\mathcal{L}_\xi(\sqrt{-g}\phi) = \partial_\mu(\sqrt{-g}\xi^\mu\phi)$ for any scalar ϕ . Thus, a Lagrangian density transforms as a total derivative under diffeomorphisms:

$$\mathcal{L}_\xi\mathcal{L} = \partial_\mu(\xi^\mu\mathcal{L}). \quad [1.23]$$

Since total derivative do not affect the field equations, the transformed solutions are still solutions. We conclude that *any diffeomorphism is a symmetry of a general covariant Lagrangian*.

We stress that this result relies crucially on the fact that in a general covariant theory the metric is a dynamical field, and not a fixed background. To appreciate this point and the difference with non-general relativistic physics, let us consider the matter Lagrangian $\mathcal{L}_{\text{M}}(g, \psi)$, which depends on both the metric $g_{\mu\nu}$ and the matter fields, which we denote collectively as ψ . Applying the chain rule, we find

$$\begin{aligned} \delta_\xi\mathcal{L}_{\text{M}} &= \frac{\delta\mathcal{L}_{\text{M}}}{\delta\psi}\delta_\xi\psi + \frac{\delta\mathcal{L}_{\text{M}}}{\delta g_{\mu\nu}}\delta_\xi g_{\mu\nu} + \partial_\mu\tilde{\theta}^\mu = \frac{\delta\mathcal{L}_{\text{M}}}{\delta\psi}\mathcal{L}_\xi\psi + \frac{\delta\mathcal{L}_{\text{M}}}{\delta g_{\mu\nu}}\mathcal{L}_\xi g_{\mu\nu} + \partial_\mu\tilde{\theta}^\mu \\ &= \mathcal{L}_\xi\mathcal{L}_{\text{M}} + \partial_\mu\tilde{\theta}^\mu = \partial_\mu(\xi^\mu\mathcal{L}_{\text{M}} + \tilde{\theta}^\mu), \end{aligned} \quad [1.24]$$

where the second equality follows from [1.19] and [1.21], and $\tilde{\theta}^\mu$ is the boundary term that arises because the Lagrangian depends on derivatives of the fields as well. Since the result is a total derivative, the field equations are unchanged, and this means that diffeomorphisms are symmetries also of the matter sector. But this conclusion relies crucially on treating the metric as a dynamical variable! In non-general relativistic physics the metric is a non-dynamical, ‘background’ field. Accordingly, there is no variation with respect to the metric, and no

second term after the first equality of [1.24]. Lacking this term the second equality breaks down, and $\delta_\xi \mathcal{L}_M$ is no longer a boundary term. We find instead

$$\delta_\xi \mathcal{L}_M = \partial_\mu (\xi^\mu \mathcal{L}_M + \tilde{\theta}^\mu) - \frac{\delta \mathcal{L}_M}{\delta g_{\mu\nu}} \mathcal{L}_\xi g_{\mu\nu}. \quad [1.25]$$

This means that only those diffeomorphisms which are isometries are symmetries of the non-general relativistic physics, e.g. the familiar Poincaré invariance of special relativity. Whereas an arbitrary diffeomorphism does not map a solution into a new solution in non-general relativistic physics, and there is no invariance under general coordinate transformations. The discussion highlights why general covariance is often referred to as background independence, namely the absence of any fixed background metric in the theory, or as diffeomorphism invariance, since every physical observable should be independent of the coordinate used to describe it. General covariance, background independence, or diffeomorphism invariance, are thus different terms used to capture the same underlying property of general relativity.

The fact that every solution can be equivalently described in any coordinate system has a useful analogy with electromagnetism, where every solution can be described in any choice of gauge for the Maxwell potential. It is actually much more than an analogy, there is in fact a precise mathematical sense in which gauge transformations in Maxwell and Yang–Mills theories have the same property of coordinate transformations in general relativity, which we discuss next. Making this analogy precise is also useful in the context of gravitational waves as it will allow us to understand the origin of their gauge dependence and distinction between physical and unphysical modes.

1.2.2. Noether's theorem and diffeomorphisms as gauge symmetries

Noether's theorem proves that every differentiable⁴ symmetry defines a current j^μ which is conserved on solutions, namely $\text{New}A_\mu j^\mu \hat{=} 0$. Here the symbol $\hat{=}$ means an equality valid only for solutions, or 'on-shell', in theoretical physics jargon. Conserved currents are extremely useful to study the properties of the dynamics of the system, and to extract general physical predictions. We have seen above that diffeomorphisms are symmetries of a general covariant Lagrangian, continuous and connected to the identity. Therefore Noether's theorem guarantees that there will be a conserved current associated to any diffeomorphism. However, the application of Noether's theorem to general relativity is quite subtle. Let us first recall the difference between 'proper symmetries' and 'gauge symmetries'. Both map solutions of the field equations into

⁴. Namely, continuous and including the identity transformation.

new solutions. If the new solution is physically distinguishable, we say that it is a proper symmetry, or a physical symmetry. If the new solution is on the other hand physically indistinguishable, we say that it is a gauge symmetry: This typically occurs when there is a redundancy of the field equations, which leaves some quantities undetermined but irrelevant for the physics. Noether's theorem provides a simple test to distinguish the two cases: in the latter, the Noether current itself vanishes on shell, and not just its divergence. This is precisely the case with diffeomorphisms in general relativity. In fact, the conserved current associated with diffeomorphisms of the Einstein-Hilbert Lagrangian [1.22] is given by

$$j_{\xi}^{\mu} = \frac{c^3}{8\pi G} \left((G^{\mu}{}_{\nu} + \Lambda \delta^{\mu}_{\nu}) \xi^{\nu} - \text{NewA}_{\nu} \nabla^{[\mu} \xi^{\nu]} \right). \quad [1.26]$$

One can immediately verify using the Bianchi identities that $\nabla_{\mu} j_{\xi}^{\mu} \doteq 0$. On the other hand, the first term above vanishes on-shell, and the second term is a total derivative. Therefore, the Noether current itself vanishes on-shell, as anticipated, and there are no conserved quantities (in the absence of boundaries). Trivial conserved quantities is a hallmark of gauge symmetries as opposed to physical symmetries, hence the result provides a precise mathematical sense in which coordinate transformations in general relativity have the same status as gauge transformations in Maxwell and Yang-Mills theories.⁵ For this reason, diffeomorphisms are also referred to as the gauge symmetry of general relativity, and fixing a coordinate choice as *fixing the gauge* in general relativity.

Having said so, there is a special situation that stands out: when the diffeomorphism corresponds to an *isometry*, namely a transformation that does not change the metric. This occurs when [1.21] vanishes, and the corresponding equation $\text{NewA}_{(\mu} \xi_{\nu)} = 0$ is called Killing equation, and ξ a Killing vector. One should keep in mind that for a generic metric, this equation does not admit any solutions: Isometries occur only for very special metrics. These special metrics are, however, important for physical applications, hence Killing vectors play an important role. First of all, anyone who is familiar with the study of geodesics on spacetimes with isometries knows that there are conserved quantities associated with the Killing vectors, and which can be derived as Noether charges for the test particles' dynamics.

5. A more rigorous approach is to look at the symplectic 2-form, and show that it is degenerate along gauge transformations and diffeomorphisms.

More importantly for us, isometries play also an important role in the study of gravitational waves, because the dynamics of perturbations on a given background is such that the isometries of the background induce proper symmetries for the perturbations. We will see this in details in Section 1.3 below.

Before moving on, let us also mention another aspect in which isometries are important for the full theory. This is a more advanced topic, and will not be needed in the following, but it allows us to give a more complete picture, and also a first intuition of how boundaries introduce non-vanishing Noether charges for diffeomorphisms. If the spacetime has isometries, the Noether current [1.26] gives rise to a useful conservation law analogue to the Gauss law in electromagnetism, which we recall states that the total charge in a region is equal to the flux of the electric field. To see this, we first observe that

$$\text{NewA}_\nu \text{NewA}^{[\mu} \xi^{\nu]} = \frac{1}{2} (R^\mu{}_\nu \xi^\nu - \square \xi^\mu + \text{NewA}^\mu \text{NewA}_\nu \xi^\nu), \quad [1.27]$$

an identity which follows from the definition of the Riemann tensor as the commutator of two covariant derivatives. If ξ^ν is a Killing vector, the second term gives $-R^\mu{}_\nu \xi^\nu$ and the last term vanishes. Then integrating both sides of the equation over a 3d portion of space V delimited by two boundaries S_1 and S_2 , and using Stokes' theorem, we find

$$Q_\xi[S] = \oint_S \text{NewA}_\nu \text{NewA}^{[\mu} \xi^{\nu]} dS_\mu, \quad [1.28]$$

$$Q_\xi[S_2] - Q_\xi[S_1] = \int_V R^\mu{}_\nu \xi^\nu dV_\mu \doteq \frac{8\pi G}{c^4} \int_V \left(T^{\mu\nu} \xi_\nu - \left(\Lambda + \frac{T}{2} \right) \xi^\mu \right) dV_\mu. \quad [1.29]$$

The Noether charge [1.28] obtained in this way is known as Komar charge. If the right-hand side of [1.29] vanishes, the Komar charge is conserved in the sense that it has the same value no matter which surface S is used, and its value changes only when the deformations of S include some source terms. If the right-hand side does not vanish, the Noether charge varies by an amount determined by the total quantity of energy-momentum in the enclosed region, see Fig 1.3. As an example, one can consider the Kerr solution, which possesses two Killing vectors corresponding to stationarity and axial symmetry. Evaluating [1.28] on an arbitrary 2-sphere S encompassing the singularity gives respectively the mass and angular momentum (up to numerical coefficients to be fixed), independently of the coordinate used and independently of deformations of S .

While Komar charges are limited to isometries, it is possible to generalize the construction of Noether charges and canonical generators to arbitrary




Figure 1.3: Conservation laws on stationary spacetimes. Evaluating the surface integral [1.28] on the innermost surface S_1 gives a quantity proportional to the energy-momentum of the planet encompassed. Evaluating it on S_2 gives the total energy-momentum of all stars and planets. The difference between the two surface integrals is proportional to the energy-momentum of the region between them. Finally since there is no source outside S_2 , integrating on S_2 or S_3 gives the same result.

spacetimes, at least in so far as they admit boundaries with non-trivial residual diffeomorphisms.⁶ Let us mention three important examples. First, spacetimes that are asymptotically flat at spatial infinity. The residual diffeomorphisms compatible with the boundary conditions are the Poincaré transformations of the flat boundary metric. One can construct Noether charges and canonical generators for these boundary diffeomorphisms (see e.g. (Iyer and Wald 1994)), and the result coincides with the Arnowit-Deser-Misner (ADM) charges that were previously derived with canonical methods. Second, spacetimes that are asymptotically flat at null infinity. This case is particularly relevant to understand gravitational waves at the non-perturbative level. The residual transformations are a generalization of Poincaré transformations in which translations are angle-dependent, an infinite-dimensional extension known as Bondi-Van der Burg-Metzner-Sachs (BMS) transformations.⁷ Noether charges for the BMS symmetry were constructed in (Ashtekar and Streubel 1981 ; Dray and Streubel

6. There is also ongoing research on constructing charges on arbitrary regions in arbitrary spacetimes.

7. Intuitively, the extension comes about because the induced metric on a null hypersurface is degenerate, hence any deformation along that direction leaves the system invariant.

1984 ; Wald and Zoupas 2000), and there is a large body of recent literature on the subject motivated by ongoing applications and developments. Among these, the application of the Noether approach to horizons and more generally null boundaries, see e.g. (Chandrasekaran *et al.* 2018 ; Ashtekar *et al.* 2022).

1.2.3. Gauge and observers

Even though every physical phenomenon can be described in any coordinate system, some choices can stand out because they simplify the description, or because they are naturally associated with a class of observers of interest. For instance in flat spacetime, Cartesian coordinates make the Christoffel's symbols vanish, thus simplifying many calculations, and can be associated to inertial observers, in the sense that they label their rods and clocks. In curved spacetimes the situation is more complicated because there are no Cartesian coordinates. To take a simple example, consider the Schwarzschild black hole. The most common coordinates used to describe this solution are the so-called 'static ones', which make time-independence manifest since the metric is independent of the time coordinate t . Or in better terms, since the time-translation Killing vector is simply ∂_t . The vector field ∂_t describes a family of non-inertial observers static at a fixed distance outside the black hole, whose time delays are related by $\sqrt{-g_{tt}}$. The Schwarzschild static coordinates are naturally associated to these observers. An alternative choice are Gullstrand-Painlevé coordinates, whose t describes the proper time of observers radially free-falling into the black hole. Another choice is the 'temporal gauge' defined by

$$g_{0\mu} = (-1, 0, 0, 0), \quad [1.30]$$

in which the radially free-falling observers are all synchronized,⁸ whence the alternative name of 'synchronous gauge'. For Schwarzschild, this gauge can be achieved using Lemaitre coordinates. The synchronization may look like a nice feature, but in these coordinates the spatial part of the Schwarzschild metric is explicitly dependent on the coordinate t ! Therefore its staticity is hidden, and has to be verified by the existence of a time-translational Killing vector. In this example we can see the analogy between coordinate choices and gauge choices in electromagnetism very clearly. An electrostatic potential is more conveniently described in the Coulomb gauge because it makes the potential manifestly time-independent. But it can be described in any other gauge, and if we use the temporal gauge $A_0 = 0$, the potential acquires an inconvenient time dependence which is pure gauge and hides the staticity of the system.

⁸. This requires giving them non-zero energy, as opposed to the Gullstrand-Painlevé observers that have zero energy.

The temporal gauge [1.30] can be chosen for *any* spacetime in a given coordinate chart. If it is done, it fixes completely the 4-dimensional diffeomorphism freedom in the chart. It is thus an example of complete gauge fixing,⁹ and the resulting coordinates describe free-falling observers with synchronised clocks. It means that in the temporal gauge the coordinates are attached to free-falling test bodies, and this can be a very useful choice in many situations.

One shortcoming of the temporal gauge is that it may ‘squeeze’ physical information in field components one would not naturally expect, as the Schwarzschild example in Lemaitre coordinates shows. Another one is that it is non-symmetric among the components, in the sense that it relies on an initial choice of coordinate to be taken as the time. An example of symmetric gauge is the harmonic gauge, which is satisfied by coordinates such that

$$\square x^\mu = \Gamma_{\nu\rho}^\mu g^{\nu\rho} = \partial_\nu(\sqrt{-g}g^{\mu\nu}) = 0, \quad [1.31]$$

where $\square = \text{New}A_\nu \text{New}A^\nu$ is the curved spacetime d’Alembertian. The symmetry makes this a convenient choice in many dynamical situations. For instance, it is the gauge in which it is easiest to see that the initial value problem is well-posed (Choquet-Bruhat 1952), and the one in which it is easiest to study gravitational waves, where it gives rise to a gauge preserving the Lorentz covariance of the Minkowski background. Notice that [1.31] is not a complete gauge fixing of the 4-dimensional diffeomorphism freedom, as there are infinitely many solutions of the wave equation, hence infinitely many choices of harmonic coordinates for a given metric. To obtain a complete gauge fixing one has to specify a unique set of harmonic coordinates with additional conditions.

One version of the equivalence principle states that purely local experiments can not distinguish the presence of gravity. In the formalism of general relativity, this is embodied in the fact that at any given point in spacetime it is possible to find coordinates so that the metric is flat and its first derivative vanish. A coordinate system that achieves this is called a *local inertial frame*. Only an experiment that can probe second-order variations in the metric would be able to see the effect of gravity in a coordinate-independent way, and these variations are the tidal forces that show up in the geodesic deviation equation. An example of local inertial frame is provided by Riemann normal coordinates,

9. Complete here refers to the 4-dimensional picture. There remains the freedom of time-independent 3-dimensional diffeomorphisms, namely of choosing the coordinates on one – and one only – given hypersurface.

which are constructed around a given point so that the Taylor expansion of the metric components around that point taken as the origin gives

$$g_{\mu\nu} = \eta_{\mu\nu} + c_{\mu\nu} R_{\mu\rho\nu\sigma} x^\rho x^\sigma + O(x^3), \quad [1.32]$$

where $c_{\mu\nu} = -1/3 \forall \mu, \nu$.¹⁰ This gauge fixing specifies coordinates only in the neighbourhood of a point and not in a full coordinate chart, but it is still perfectly sufficient to describe the physics of a local inertial frame around that point. It is also possible to find coordinates such that the Christoffel's symbols vanish everywhere along a chosen time-like geodesic. Such coordinates are known as Fermi normal coordinates, and describe a free-falling local inertial frame. The metric takes the same form [1.32], but with $c_{00} = -1, c_{0a} = -2/3, c_{ab} = -1/3$. The difference between temporal gauge and Fermi normal coordinates is that in the first case every observer at constant spatial coordinates is free falling, whereas in the second case only the observer going through the origin along the ∂_t geodesic is free falling.

1.3. Perturbative treatment of Einstein's equations

1.3.1. The idea: general background spacetime

Perturbation theory can be set up choosing a background metric $\bar{g}_{\mu\nu}$ and writing

$$g_{\mu\nu} = \bar{g}_{\mu\nu} + h_{\mu\nu}, \quad [1.33]$$

with the assumption that $|h_{\mu\nu}| \ll 1$ in some chosen coordinate system, so that it can be treated as a perturbation. One can then systematically Taylor-expand all metric functionals around the background, starting with the inverse metric $g^{\mu\nu} = \bar{g}^{\mu\nu} - h^{\mu\nu} + O(h^2)$ and Levi-Civita connection

$$\Gamma_{\nu\rho}^\mu = \bar{\Gamma}_{\nu\rho}^\mu + \Gamma_{\nu\rho}^{(1)\mu} + O(h^2), \quad \Gamma_{\nu\rho}^{(1)\mu} = \frac{1}{2} \bar{g}^{\mu\sigma} (2\partial_{(\nu} g_{\rho)\sigma} - \partial_\sigma g_{\nu\rho}), \quad [1.34]$$

and attempt to solve the field equations order by order:

$$\bar{G}_{\mu\nu} + G_{\mu\nu}^{(1)} + G_{\mu\nu}^{(2)} + \dots = \frac{8\pi G}{c^4} (\bar{T}_{\mu\nu} + T_{\mu\nu}^{(1)} + T_{\mu\nu}^{(2)} + \dots) \quad [1.35]$$

¹⁰. This formula is only valid in certain coordinate choices, hence the non-covariant notation with the $\mu\nu$ indices repeated but not summed over.

Here $\bar{G}_{\mu\nu} = G_{\mu\nu}(\bar{g})$, $\bar{T} = T(\bar{g})$, $G_{\mu\nu}^{(1)} = G_{\mu\nu}(\bar{g}; h)$, and so on. Explicit expressions will not be needed here, but are given in Appendix 1.4 for completeness. The lowest order of the procedure is straightforward, one just has to be consistent with the treatment of the matter energy-momentum tensor: if we want to expand around a background which is a vacuum solution, we need $\bar{G}_{\mu\nu} = 0$ hence we should treat matter as a first order perturbation. Then the first order equation to solve is

$$G_{\mu\nu}^{(1)} = \frac{8\pi G}{c^4} \bar{T}_{\mu\nu}. \quad [1.36]$$

Solving this equation determines $h_{\mu\nu}$ in terms of its independent degrees of freedom, the background solution and the matter content. To go to second order, we add a second perturbation, writing

$$g_{\mu\nu} = \bar{g}_{\mu\nu} + h_{\mu\nu} + h_{\mu\nu}^{(2)}. \quad [1.37]$$

Then $G_{\mu\nu}^{(2)} = G_{\mu\nu}^{(2)}(h) + G_{\mu\nu}^{(1)}(h^{(2)})$ has two contributions, and we solve for $h^{(2)}$ using

$$G_{\mu\nu}^{(1)}(h^{(2)}) = \frac{8\pi G}{c^4} (T_{\mu\nu}^{(1)} + t_{\mu\nu}^G), \quad t_{\mu\nu}^G := -\frac{c^4}{8\pi G} G_{\mu\nu}^{(2)}(h). \quad [1.38]$$

Notice that the first order solution feeds back as a source for the second order solution, a standard procedure from perturbatively solving non-linear equations. The same procedure applies also to the case when matter fields contribute to the background as well. Although we wrote the perturbative expansion using dimensionless quantities only, it is always possible to rescale the metric perturbation by \sqrt{G} , so that the kinetic term of the free-field Lagrangian is canonically normalized. Then the expansion is a power series in G , and it is also referred to as *post-Minkowskian* (PM) expansion, when the background is flat. In many cases the equations are still too hard to solve at each order, and one needs to look for additional approximations. A very common one is the non-relativistic approximation, also known as *post-Newtonian* (PN) expansion, in which one starts from a source that moves at small velocity with respect to the background flat metric, namely with $v \ll c$, and expands in powers of v^2/c^2 . Alternative approximation schemes include the Bondi asymptotic expansion which is perturbative in the inverse distance from the source but valid at all orders in G and v^2/c^2 (Bondi *et al.* 1962 ; Sachs 1962); the extremal mass ratio inspiral (EMRI) and self-force expansion for a two-body system where the small parameter is the mass-ratio (Barack and Pound 2019); the effective one-body approach based on resummed PN results (Buonanno and Damour 1999), and

more recently the effective field theory approach based on tools from quantum field theory (Goldberger and Rothstein 2006).

On top of the technical difficulties, there is also a conceptual one: *truncating at any order beyond the lowest breaks covariance*. Expanding both sides of [1.21] with [1.33], we obtain

$$\delta_\xi \bar{g}_{\mu\nu} + \delta_\xi h_{\mu\nu} = \mathcal{L}_\xi \bar{g}_{\mu\nu} + \mathcal{L}_\xi h_{\mu\nu}. \quad [1.39]$$

We assume that the background metric is fixed once and for all and unaffected by diffeomorphisms, namely $\delta \bar{g}_{\mu\nu} = 0$ and $\delta_\xi \bar{g}_{\mu\nu} = 0$,¹¹ and define the effect of the linearized diffeomorphism on the perturbation to be

$$\delta_\xi h_{\mu\nu} = \mathcal{L}_\xi \bar{g}_{\mu\nu} + \mathcal{L}_\xi h_{\mu\nu}. \quad [1.40]$$

The key point is that the right-hand side contains terms of different order in $h_{\mu\nu}$. This is why different orders of the perturbative expansion must be included in order for diffeomorphisms to be a symmetry, and conversely truncating at a fixed order breaks covariance. It is instructive to prove this in detail, because it highlights the special features that occur for the quadratic Lagrangian, namely the free theory, for which covariance is on the other hand possible.

To study the symmetries of the perturbative expansion, we write the perturbed Lagrangian as follows

$$\mathcal{L}(\bar{g} + h) = \bar{\mathcal{L}} + \bar{\mathcal{L}}^{(1)\mu\nu} h_{\mu\nu} + \frac{1}{2} h_{\mu\nu} \bar{\mathcal{L}}^{(2)\mu\nu\rho\sigma} h_{\rho\sigma} + \dots, \quad [1.41]$$

where barred quantities only depend on \bar{g} and not on the perturbation. Using [1.40], the infinitesimal variation gives

$$\begin{aligned} \delta_\xi \mathcal{L} = & \bar{\mathcal{L}}^{(1)\mu\nu} \mathcal{L}_\xi \bar{g}_{\mu\nu} + \bar{\mathcal{L}}^{(1)\mu\nu} \mathcal{L}_\xi h_{\mu\nu} + h_{\mu\nu} \bar{\mathcal{L}}^{(2)\mu\nu\rho\sigma} \mathcal{L}_\xi \bar{g}_{\rho\sigma} \\ & + h_{\mu\nu} \bar{\mathcal{L}}^{(2)\mu\nu\rho\sigma} \mathcal{L}_\xi h_{\rho\sigma} + \frac{1}{2} h_{\mu\nu} h_{\rho\sigma} \bar{\mathcal{L}}^{(3)\mu\nu\rho\sigma\tau\lambda} \mathcal{L}_\xi \bar{g}_{\tau\lambda} + \dots \end{aligned} \quad [1.42]$$

11. It is also possible to interpret [1.39] as $\delta_\xi \bar{g}_{\mu\nu} = \mathcal{L}_\xi \bar{g}_{\mu\nu}$ and $\delta_\xi h_{\mu\nu} = \mathcal{L}_\xi h_{\mu\nu}$. The proof that this is a symmetry is identical. However this definition of perturbed transformations is less interesting physically, because it is more natural to compare perturbations when they are defined with respect to the same background.

All terms can be collected into Lie derivatives:

$$\bar{\mathcal{L}}^{(1)\mu\nu} \mathcal{L}_\xi \bar{g}_{\mu\nu} = \mathcal{L}_\xi \bar{\mathcal{L}}, \quad [1.43]$$

$$\bar{\mathcal{L}}^{(1)\mu\nu} \mathcal{L}_\xi h_{\mu\nu} + h_{\mu\nu} \bar{\mathcal{L}}^{(2)\mu\nu\rho\sigma} \mathcal{L}_\xi \bar{g}_{\rho\sigma} = \mathcal{L}_\xi (\bar{\mathcal{L}}^{(1)\mu\nu} h_{\mu\nu}), \quad [1.44]$$

$$h_{\mu\nu} \bar{\mathcal{L}}^{(2)\mu\nu\rho\sigma} \mathcal{L}_\xi h_{\rho\sigma} + \frac{1}{2} h_{\mu\nu} h_{\rho\sigma} \bar{\mathcal{L}}^{(3)\mu\nu\rho\sigma\tau\lambda} \mathcal{L}_\xi \bar{g}_{\tau\lambda} = \frac{1}{2} \mathcal{L}_\xi (\bar{\mathcal{L}}^{(2)\mu\nu\rho\sigma} h_{\mu\nu} h_{\rho\sigma}), \quad [1.45]$$

and so on. Each Lie derivative gives a boundary term through [1.23], hence [1.40] is indeed a symmetry of the full Lagrangian. However, in every case except the lowest one, getting a boundary term for $\bar{\mathcal{L}}^{(n)}$ requires both $\bar{\mathcal{L}}^{(n)}$ and $\bar{\mathcal{L}}^{(n+1)}$, hence if we truncate the series at a fixed order, we lose covariance.

There are two special features that occur at the quadratic order. First, if we take the background to be a solution, the first term in [1.44] vanishes, hence

$$\delta_\xi h_{\mu\nu} = \mathcal{L}_\xi \bar{g}_{\mu\nu} \quad [1.46]$$

is a symmetry of the quadratic Lagrangian. Second, if the background has isometries, then the term proportional to $\bar{\mathcal{L}}^{(3)}$ in [1.45] drops out, and then

$$\delta_\xi h_{\mu\nu} = \mathcal{L}_\xi h_{\mu\nu} \quad [1.47]$$

is a symmetry of the quadratic Lagrangian. This shows that isometries play a special role in perturbation theory. For a generic on-shell background the symmetry of the quadratic Lagrangian is [1.46]. But if the background has isometries, we have two different realization of the diffeomorphism symmetry in the quadratic Lagrangian: [1.46] for a generic diffeomorphism, and [1.47] for a Killing vector. To make this consistent with the perturbative expansion, we treat a Killing ξ as zero-th order, and a non-Killing ξ as first order.

We stress that [1.46] and [1.47] are symmetries only for the quadratic Lagrangian. From the cubic Lagrangian onwards, there is no symmetry at fixed order in $h_{\mu\nu}$. The only symmetry is the combined [1.40] and requires two different perturbative orders of the Lagrangian. This fact has immediate consequences for the expanded Einstein tensor [1.35]. Since the leading order $G^{(1)}$ comes from the quadratic Lagrangian, it is invariant under [1.46]. But $G_{\mu\nu}^{(2)}$ which is derived from the cubic Lagrangian is not invariant. *As a consequence, the quantity $t_{\mu\nu}^G$ appearing in [1.38] is not gauge invariant.* This means that it does not provide a meaningful notion of gravitational energy-momentum tensor. We will come back to this important point below.

Even though both symmetries [1.46] and [1.47] descend from the same diffeomorphism invariance of the theory, which is a gauge symmetry, they have a different status at the perturbative level. Applying Noether's theorem to the quadratic Lagrangian, we find that the generic diffeomorphisms still have vanishing conserved current, and therefore maintain their status of gauge symmetries. However the diffeomorphisms corresponding to isometries of the background have non-vanishing Noether charges, indeed just like a standard theory on flat spacetime. A related difference is that the field equations are *invariant* under [1.46], hence there are linear dependencies in the equations and some field components are left undetermined, and *covariant* under [1.47], and there are no constraints associated with them.

1.3.2. Weak-field approximation

The weak-field approximation is a special case of the perturbative expansion in which the background spacetime is Minkowski,

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}. \quad [1.48]$$

For instance in the solar system spacetime is approximately flat, hence it can be described using this approximation.¹² We further use Cartesian coordinates, so all covariant derivatives become partial derivatives. In this case the linearized Einstein equations [1.36] take the simple form

$$\square h_{\mu\nu} - 2\partial_{(\mu}\partial_{\rho}h^{\rho}_{\nu)} + \partial_{\mu}\partial_{\nu}h + \eta_{\mu\nu}(\partial_{\rho}\partial_{\sigma}h^{\rho\sigma} - \square h) = -\frac{16\pi G}{c^4}T_{\mu\nu}, \quad [1.49]$$

where \square is the d'Alembertian in flat spacetime.

Let us study the symmetries of [1.49]. The Minkowski background has ten isometries, the Poincaré transformations. We parametrize them with vector fields

$$\xi^{\mu} = a^{\mu}_{\nu}x^{\nu} + b^{\mu}, \quad [1.50]$$

where a^{μ}_{ν} and b^{μ} are constants, and $a_{(\mu\nu)} = 0$. According to the general discussion of the previous section, we expect two different types of symmetries

12. In first approximation, we can describe the solar system's metric using the Schwarzschild solution with M the mass of the sun. Then the curvature scale can be estimated writing the Kretschmann scalar in static coordinates, and this gives $R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma} = 48G^2M^2c^{-4}r^{-6} \sim 10^{-27}\text{m}^{-4}(R_{\odot}/r)^6$ which is extremely small.

in [1.49], both induced from diffeomorphism invariance of the full theory, and corresponding to those linearized diffeomorphisms that are isometries or not of the background. For diffeomorphisms corresponding to isometries, [1.47] gives

$$\delta_\xi h_{\mu\nu} = \mathcal{L}_\xi h_{\mu\nu} = (a^\rho{}_\sigma x^\sigma + b^\rho) \partial_\rho h_{\mu\nu} + 2a^\rho{}_{(\mu} h_{\nu)\rho}. \quad [1.51]$$

This is the transformation of a rank-2 tensor in Minkowski under Poincaré transformations. For generic diffeomorphisms that are not isometries, [1.46] gives

$$\delta_\xi h_{\mu\nu} = \mathcal{L}_\xi \eta_{\mu\nu} = 2\partial_{(\mu} \xi_{\nu)}. \quad [1.52]$$

Both transformations are symmetries of [1.49], the difference being that the first changes the equations but in a covariant way, whereas the second leaves them invariant. This can be checked easily on the left-hand side. For the right-hand side, one has to pay attention to the behaviour of the energy-momentum tensor. Assuming the background η to be a solution requires that we are treating the matter fields as first order in perturbation theory. Then it transforms as $\mathcal{L}_\xi T_{\mu\nu}$ for a Killing ξ , whereas its transformation under generic diffeomorphisms is second order, hence it does not affect the linearized equations.

The presence of the d’Alambertian in the linearized field equations suggests that wave solutions are indeed possible. However, there is an intricate tensorial structure that needs to be dealt with. Before doing so, let us discuss what it means for a wave to be ‘tensorial’. The waves that we are most familiar with, such as water waves or sound waves, are *scalar* waves, namely the quantity whose perturbation propagates following the wave equation is a scalar function, such as the height of the water or the pressure of the air. Electromagnetic waves on the other hand propagate changes in the electric and magnetic field which are described by vectors, and are thus ‘vectorial’ waves. The difference between scalar, vectorial and tensorial waves can be described in terms of *spin*. The reader already familiar with these concepts, or the reader interested in looking first at the explicit solutions without these details, can skip the next subsection.

1.3.3. Spin and helicity

Tensors in Minkowski spacetime belong to finite-dimensional representations of the Lorentz group. This property can be used to decompose each tensor into irreducible parts, namely parts that are not mixed with one another by a Lorentz transformation. The irreducible parts can be labelled by a pair of half-integers (j_1, j_2) , and contain $(2j_1 + 1)(2j_2 + 1)$ components. For instance,

a 4-vector v^μ transforms under the irreducible Lorentz representation $(\frac{1}{2}, \frac{1}{2})$ with 4 components, and a symmetric tensor $h_{\mu\nu}$ transforms under the reducible Lorentz representation $(\mathbf{1}, \mathbf{1}) \oplus (\mathbf{0}, \mathbf{0})$ with $9 + 1 = 10$ components.

The representations can be further subdivided if we pick a time direction τ^μ , and restrict attention to the rotation subgroup of the Lorentz group that preserves it. The subsets of the tensor which are irreducible with respect to the rotation subgroup are called *spin* representations, and their allowed values given by the Clebsch-Gordan addition rule $(j_1 + j_2, j_1 + j_2 - 1, \dots, |j_1 - j_2|)$. For instance a vector contains the two spin representations $\mathbf{1} \oplus \mathbf{0}$, and a symmetric tensor the four spin representations $\mathbf{2} \oplus \mathbf{1} \oplus \mathbf{0} \oplus \mathbf{0}$. To be more explicit let us take $\tau^\mu = (1, 0, 0, 0)$. The rotation subgroup that preserves it has the form

$$R^\mu{}_\nu = \begin{pmatrix} 1 & \vec{0} \\ \vec{0} & R^a{}_b \end{pmatrix}, \quad R^a{}_b \in \text{SO}(3). \quad [1.53]$$

It is then immediate to see that the spin-1 and spin-0 representation of the 4-vector are the spatial vector v^a and the spatial scalar v^0 respectively. Similarly for a 1-form v_μ , the spatial and time components v_a and v_0 . For a symmetric tensor, the four spin representations are

$$h_{\langle ab \rangle} = h_{ab} - \frac{1}{3} \delta_{ab} h^c{}_c, \quad h_{0a}, \quad h_{00}, \quad h_s := h^c{}_c. \quad [1.54]$$

The spin 0 representation has one component, the spin 1 has three, and the spin 2 has five. These different components can be classified choosing a reference spatial axis and looking at the eigenmodes of the rotation generator along that axis. To fix ideas let us choose the z axis. The rotation matrix that preserves it is

$$R^a{}_b = \begin{pmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad [1.55]$$

Inserting this in [1.53] and acting on a 4-vector v^μ we obtain $v'^\mu = R^\mu{}_\nu v^\nu$, where

$$\begin{aligned} v'^0 &= v^0, & v'^x &= \cos \theta v^x - \sin \theta v^y, \\ v'^y &= \sin \theta v^x + \cos \theta v^y, & v'^z &= v^z. \end{aligned} \quad [1.56]$$

The temporal and the z component, which is longitudinal along the rotation axis, are invariant. The components x and y , which are transverse to the rotation axis, transform among themselves, and we can diagonalize the action introducing

$$v^\pm := v^x \pm iv^y, \quad v'^\pm = e^{\pm i\theta} v^\pm. \quad [1.57]$$

In this way we have identified all 4 eigenmodes of the vector. We can equivalently describe the decomposition in terms of a basis of eigenvectors. If we start from the canonical basis $e_I^\mu := \delta_I^\mu$, where $I = 0, \dots, 3$, then the eigenvectors are

$$\epsilon_0^\mu = e_0^\mu, \quad \epsilon_L^\mu = e_3^\mu, \quad \epsilon_\pm^\mu = \frac{1}{\sqrt{2}}(e_1^\mu \mp ie_2^\mu), \quad [1.58]$$

and their corresponding eigenvalues

$$R^\mu{}_\nu \epsilon_0^\nu = \epsilon_0^\mu, \quad R^\mu{}_\nu \epsilon_L^\nu = \epsilon_L^\mu, \quad R^\mu{}_\nu \epsilon_\pm^\nu = e^{\pm i\theta} \epsilon_\pm^\mu. \quad [1.59]$$

The integers 0 and ± 1 that appear in front of the rotation angle θ are the projection of the spin along the z axis. Since our convention for vector rotations is anti-clockwise, the $+1$ (resp. -1) component can be also referred to as right-handed (resp. left-handed).

Acting on a symmetric tensor $h_{\mu\nu}$ we obtain $h'_{\mu\nu} = R^\rho{}_\mu R^\sigma{}_\nu h_{\rho\sigma}$, where

$$h'_{00} = h_{00}, \quad h'_{0z} = h_{0z}, \quad h'_s = h_s, \quad h'_{zz} = h_{zz} \quad [1.60a]$$

$$h'_{0x} = \cos\theta h_{0x} + \sin\theta h_{0y}, \quad h'_{0y} = -\sin\theta h_{0x} + \cos\theta h_{0y}, \quad [1.60b]$$

$$h'_{xz} = \cos\theta h_{xz} + \sin\theta h_{yz}, \quad h'_{yz} = -\sin\theta h_{xz} + \cos\theta h_{yz}, \quad [1.60c]$$

and

$$h'_{xx} = \cos\theta(h_{xx} \cos\theta + h_{xy} \sin\theta) + \sin\theta(h_{xy} \cos\theta + h_{yy} \sin\theta), \quad [1.60d]$$

$$h'_{xy} = -\sin\theta(h_{xx} \cos\theta + h_{xy} \sin\theta) + \cos\theta(h_{xy} \cos\theta + h_{yy} \sin\theta), \quad [1.60e]$$

$$h'_{yy} = -\sin\theta(h_{xy} \cos\theta - h_{xx} \sin\theta) + \cos\theta(h_{yy} \cos\theta - h_{xy} \sin\theta). \quad [1.60f]$$

The last three can be disentangled if we define

$$h_+ := \frac{1}{2}(h_{xx} - h_{yy}), \quad h_\times := h_{xy}, \quad h_s := h_{xx} + h_{yy} + h_{zz}. \quad [1.60g]$$

Then,

$$h'_+ = \cos 2\theta h_+ + \sin 2\theta h_\times, \quad h'_\times = -\sin 2\theta h_+ + \cos 2\theta h_\times, \quad [1.60h]$$

and

$$h^{\pm 2} := \frac{1}{\sqrt{2}}(h_+ \mp i h_\times), \quad h'^{\pm 2} = e^{\pm 2i\theta} h^{\pm 2}. \quad [1.60i]$$

We have thus identified all ten eigenmodes of a symmetric rank-2 tensor under the z rotations.

To write the eigenvectors, we introduce a canonical basis in the space of symmetric 4×4 matrices,

$$e_{\mu\nu}^{IJ} = \begin{cases} \delta_\mu^I \delta_\nu^J & I = J \\ \frac{1}{\sqrt{2}}(\delta_\mu^I \delta_\nu^J + \delta_\mu^J \delta_\nu^I) & I \neq J \end{cases} \quad [1.61]$$

Then

$$\begin{aligned} \epsilon_{\mu\nu}^0 &= e_{\mu\nu}^{00}, & \epsilon_{\mu\nu}^s &= \frac{1}{\sqrt{3}}(e_{\mu\nu}^{11} + e_{\mu\nu}^{22} + e_{\mu\nu}^{33}), \\ w_{\mu\nu}^\pm &= \frac{1}{\sqrt{2}}(e_{\mu\nu}^{01} \pm i e_{\mu\nu}^{02}), & w_{\mu\nu}^L &= e_{\mu\nu}^{03}, \end{aligned} \quad [1.62a]$$

$$\begin{aligned} \epsilon_{\mu\nu}^{\tau\pm} &= \frac{1}{2}(e_{\mu\nu}^{11} - e_{\mu\nu}^{22}) \pm \frac{i}{\sqrt{2}}e_{\mu\nu}^{12}, & \epsilon_{\mu\nu}^{L\pm} &= \frac{1}{\sqrt{2}}(e_{\mu\nu}^{13} \pm i e_{\mu\nu}^{23}), \\ \epsilon_{\mu\nu}^{LL} &= \frac{1}{\sqrt{6}}(e_{\mu\nu}^{11} + e_{\mu\nu}^{22} - 2e_{\mu\nu}^{33}) \end{aligned} \quad [1.62b]$$

with eigenvalues

$$R^\rho{}_\mu R^\rho{}_\nu \epsilon_{\rho\sigma}^i = \epsilon_{\mu\nu}^i, \quad \text{for } i = 0, s, LL \quad R^\rho{}_\mu R^\rho{}_\nu w_{\rho\sigma}^L = w_{\mu\nu}^L, \quad [1.63]$$

$$R^\rho{}_\mu R^\rho{}_\nu \epsilon_{\rho\sigma}^\pm = e^{\pm 2i\theta} \epsilon_{\mu\nu}^\pm, \quad R^\rho{}_\mu R^\rho{}_\nu w_{\rho\sigma}^\pm = e^{\pm i\theta} w_{\mu\nu}^\pm, \quad R^\rho{}_\mu R^\rho{}_\nu \epsilon_{\rho\sigma}^{L\pm} = \epsilon^{\pm i\theta} w_{\mu\nu}^{L\pm}. \quad [1.64]$$

The spin-2 components $e^{\tau\pm}$, $e^{L\pm}$ and e^{LL} carry respectively the $\pm 2, \pm 1$ and 0 modes of the z -projection. The spin-1 components w^\pm and w^L carry ± 1 and 0 modes as in [1.59], and $e^{0,s}$ are the remaining two spin-0 modes from [1.54].

The discussion so far concerned global vectors and tensors. In the case of electromagnetism and gravity linearized around Minkowski we work with vector and tensor *fields*, $A_\mu(x)$ and $h_{\mu\nu}(x)$. The mode decompositions described

above are then applied locally in Fourier space. Each Fourier mode $\tilde{A}_\mu(k)$ and $\tilde{h}_{\mu\nu}(k)$ is characterized by a wave vector k , and we can use its spatial direction to classify the spin components. With some abuse of language, we will often refer to the wave vector as the momentum of the wave. The eigenvalues of the projection of the spin along the momentum are called *helicities*, hence when this basis is chosen the mode decomposition is called helicity decomposition. The notion of helicity is closely related to the notion of polarization. More precisely, modes of helicity ± 1 describe waves of right-handed and left-handed circular polarization respectively, and different linear combinations can be taken to describe for instance linear or elliptic polarizations. The helicity-1 linear polarizations are simply the x and y components, and the helicity-2 linear polarizations are

$$\epsilon_{\mu\nu}^+ = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \epsilon_{\mu\nu}^\times = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad [1.65]$$

The labels stand for ‘plus’ (not to be confused with the plus used in the circular polarizations, and which will not be used in the following) and ‘cross’, and the reason will become clear below.

The helicity decomposition with an arbitrary momentum \vec{k} can be conveniently described introducing the transverse and longitudinal projectors

$$T_b^a := \delta_b^a - \frac{k^a k_b}{\vec{k}^2}, \quad L_b^a := \frac{k^a k_b}{\vec{k}^2}. \quad [1.66]$$

Using these, we can decompose a spin-1 vector as follows,

$$A^a = A_T^a + A_L^a, \quad A_T^a = P_T^{(1)a} A^b, \quad A_L^a = P_L^{(1)a} A^b, \quad [1.67]$$

where

$$P_T^{(1)a} = T_b^a, \quad P_L^{(1)a} = L_b^a, \quad P_T^{(1)} + P_L^{(1)} = P^{(1)}; \quad [1.68]$$

and a spin-2 tensor as follows,

$$\begin{aligned} h_{(ab)} &= h_{ab}^{\text{TT}} + h_{ab}^{\text{L}} + h_{ab}^{\text{LL}}, & h_{ab}^{\text{TT}} &= P_{\text{TT}}^{(2)cd} h_{cd} \\ h_{ab}^{\text{L}} &= P_{\text{L}}^{(2)cd} h_{cd}, & h_{ab}^{\text{LL}} &= P_{\text{LL}}^{(2)cd} h_{cd}, \end{aligned} \quad [1.69]$$

where

$$\begin{aligned}
P_{\text{TT}}^{(2)ab} &= T_{(c}^a T_{d)}^b - \frac{1}{2} T^{ab} T_{cd}, & P_{\text{L}}^{(2)ab} &= T_{(c}^a L_{d)}^b + T_{(c}^b L_{d)}^a, \\
P_{\text{LL}}^{(2)ab} &= \frac{1}{3} \left(\frac{1}{2} T^{ab} T_{cd} + 2L^{ab} L_{cd} - T^{ab} L_{cd} - L^{ab} T_{cd} \right), \\
P^{(2)} &= P_{\text{TT}}^{(2)} + P_{\text{L}}^{(2)} + P_{\text{LL}}^{(2)}.
\end{aligned} \tag{1.70}$$

We can then identify the helicities of the different projectors studying how they transform under a rotation with axis \vec{k} . For simplicity let us consider the case when \vec{k} is in the z direction, so that we can use the formulas already derived for the eigenvectors. In this case,

$$T^a_b = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad L^a_b = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{1.71}$$

Then [1.67] reduces to $A_a^{\text{T}} = (A_x, A_y, 0)$ and $A_a^{\text{L}} = (0, 0, A_z)$. Comparing these to the earlier decomposition [1.59], we conclude that the transverse projector contains the ± 1 helicity modes of a spin-1 field, and the longitudinal projector the 0 helicity mode. The spin-2 projectors based on [1.71] give

$$\begin{aligned}
P_{\text{TT}}^{(2)ab} h^{cd} &= (ThT)^{ab} - \frac{1}{2} \text{Tr}(Th) T^{ab} \\
&= \begin{pmatrix} \frac{1}{2}(h_{xx} - h_{yy}) & h_{xy} & 0 \\ h_{xy} & -\frac{1}{2}(h_{xx} - h_{yy}) & 0 \\ 0 & 0 & 0 \end{pmatrix} = \begin{pmatrix} h_+ & h_{\times} & 0 \\ h_{\times} & -h_+ & 0 \\ 0 & 0 & 0 \end{pmatrix}, \tag{1.72a}
\end{aligned}$$

$$P_{\text{L}}^{(2)ab} h^{cd} = \begin{pmatrix} 0 & 0 & h_{xz} \\ 0 & 0 & h_{yz} \\ 0 & 0 & 0 \end{pmatrix}, \quad P_{\text{LL}}^{(2)ab} h^{cd} = \frac{1}{6} (h_{xx} + h_{yy} - 2h_{zz}) \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -2 \end{pmatrix}. \tag{1.72b}$$

Comparing these to the earlier decomposition [1.62], we conclude that the spin-2 part of the gravitational perturbation can be decomposed into 5 helicity modes $\pm 2, \pm 1, 0$, which are carried respectively by the TT, L and LL components.

The spin-helicity interpretation of the components of $h_{\mu\nu}$ is a kinematical classification based on a choice of reference frame given by the time direction¹³

13. Which we have chosen to coincide with t of the background Cartesian coordinates. However the whole spin-helicity description is Lorentz covariant. The spin projectors for an arbitrary time-direction τ^μ , $\tau^2 = -1$, are

$$P_{\nu}^{(1)\mu} = q_{\nu}^{\mu} := \delta_{\nu}^{\mu} + \tau^{\mu} \tau_{\nu}, \quad P_{\nu}^{(0)\mu} = -\tau^{\mu} \tau_{\nu}, \quad 1^{(\frac{1}{2}, \frac{1}{2})} = P^{(1)} + P^{(0)}, \tag{1.73}$$

and the spatial direction \vec{k} . The next question is which of these components are dynamical. If the field satisfied the simple wave equation $\square h_{\mu\nu} = 0$, then all components would be dynamical, and the field would carry 10 independent degrees of freedom corresponding to all the helicity states described above. But this is not the case of [1.49], because of the additional derivative operators present, and the gauge redundancy. One way to identify the degrees of freedom is to fix the gauge and study the resulting solutions. This is what we do next.

1.3.4. De Donder and TT gauges

The gauge symmetry can be exploited to simplify the linearized field equations. This can be done using [1.52] to put the metric perturbation in a form where it satisfies

$$\partial_\mu h^\mu{}_\nu - \frac{1}{2}\partial_\nu h = 0. \quad [1.75]$$

Doing so eliminates all terms containing divergences. The linearized field equations [1.49] are thus equivalent to

$$\square \bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} T_{\mu\nu}, \quad \partial_\mu \bar{h}^{\mu\nu} = 0, \quad \bar{h}_{\mu\nu} := h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h. \quad [1.76]$$

The short-hand notation $\bar{h}_{\mu\nu}$ is introduced for convenience, and it is known as trace-reversed perturbation, since $\bar{h} = -h$. The condition [1.75] is called De Donder gauge in the literature, but also Lorenz gauge, in analogy with electromagnetism, and harmonic gauge, because it preserves harmonic coordinates, as can be seen linearizing [1.31]. The compatibility of the coupled system is guaranteed by the fact that $\partial_\mu T^{\mu\nu}$ vanishes on solutions of the matter field equations, more on this in Section 2.2.1 below.

for a vector, and

$$\begin{aligned} P^{(2)\mu\nu}{}_{\rho\sigma} &= \delta^\mu{}_\rho \delta^\nu{}_\sigma + \tau^\mu \tau_\rho \delta^\nu{}_\sigma - \frac{1}{2} q^{\mu\nu} q_{\rho\sigma}, & P^{(1)\mu\nu}{}_{\rho\sigma} &= -\tau^\mu \tau_\rho q^\nu{}_\sigma, \\ P^{(0)\mu\nu}{}_{\rho\sigma} &= \tau^\mu \tau_\rho \tau^\nu \tau_\sigma, & P_s^{(0)\mu\nu}{}_{\rho\sigma} &= \frac{1}{2} q^{\mu\nu} q_{\rho\sigma}, \\ \mathbb{1}^{(1,1)\oplus(0,0)} &= P^{(2)} + P^{(1)} + P^{(0)} + P_s^{(0)}, \end{aligned} \quad [1.74]$$

for a symmetric tensor. One can also write the helicity projectors as 4d covariant objects encoding τ in a choice of null vector transverse to k^μ , see e.g. (Weinberg 2005).

The condition [1.75] does *not* fix completely the 4-dimensional diffeomorphism symmetry. Intuitively, this occurs because we are not fixing metric coefficients, but only their derivatives. To make the residual freedom explicit, observe that the diffeomorphism required to put an arbitrary metric perturbation in De Donder form is given by a solution of the equation

$$\square\xi^\mu = -\partial_\nu h^{\mu\nu} + \frac{1}{2}\partial^\mu h. \quad [1.77]$$

But for a given initial metric, this equation admits infinitely many solutions, parametrized by the zero modes of the d'Alembertian operator. In other words, once the De Donder condition is satisfied, there remains a residual freedom of gauge transformations that satisfy $\square\xi^\mu = 0$. Thus the De Donder gauge contains in fact an infinite family of inequivalent gauge fixings. For this reason, it may be more appropriate to refer to [1.75] not as a gauge fixing condition, but rather a *family* of gauge fixings. The simpler albeit vaguer term De Donder gauge is, however, the one in use in most of the literature.

If need be, a unique representative of the De Donder gauge can be specified fixing the solution of $\square\xi^\mu = 0$ in terms of initial data at a reference $t = 0$ hypersurface. Such initial data can always be chosen such that any four components of $h_{\mu\nu}$ and their time derivatives vanish there. For vacuum solutions, this means that the chosen components vanish everywhere. It is convenient to apply this procedure and set to zero the components \bar{h} , so that $\bar{h}_{\mu\nu} = h_{\mu\nu}$, and h_{0a} . Then the De Donder condition implies that also $\partial_a h^{ab}$ and $\partial_0 h_{00}$ vanish everywhere. The vacuum equations then imply $\bar{\partial}^2 h_{00} = 0$, hence the only solution with vanishing boundary conditions is $h_{00} = 0$. We conclude that in this gauge, vacuum solutions satisfy

$$h = h_{00} = h_{0a} = \partial_a h^{ab} = 0, \quad h_{ab} = h_{ab}^{\text{TT}}. \quad [1.78]$$

This shows explicitly that there are only two independent degrees of freedom. We call this choice the *transverse-traceless gauge* (TT gauge in short), since in this gauge, vacuum solutions coincide with the transverse-traceless perturbations. We stress that this property and the equations [1.78] are only valid for vacuum solutions. In other words, [1.78] is not the definition of a gauge condition, but rather the specific value that solutions take in a certain gauge. The analogue for solutions with sources will be discussed below. Notice also that $h_{0\mu} = 0$ is the linearized approximation of [1.30]. Therefore the TT gauge of vacuum solutions implies the temporal gauge, hence the TT coordinates describe free-falling observers.

In the presence of sources on the other hand, even if we set to zero 4 components and their time derivatives on the chosen initial hypersurface, they will

no longer be zero in the future causal domain of the sources. In this case there are more convenient choices to specialize the De Donder gauge condition, as we will discuss below.

1.3.5. Vacuum solutions

Let us flesh out these considerations by looking at the explicit form of the vacuum solutions. This will make it clear that the identification of the degrees of freedom with the transverse-traceless modes is a gauge-invariant statement. The vacuum equations can be solved straightforwardly taking linear combination of plane waves via the Fourier transform

$$h_{\mu\nu}(x) = \text{Re} \int d^4k \tilde{h}_{\mu\nu}(k) e^{ik \cdot x}. \quad [1.79]$$

Imposing the vacuum equations and the De Donder condition [1.76] requires

$$k^2 = 0, \quad k^\mu \tilde{h}_{\mu\nu} = \frac{1}{2} k_\nu \tilde{h}. \quad [1.80]$$

The first equation is solved by $k^0 = \pm |\vec{k}|$, and we denote $\omega := k^0 c > 0$ for a future-pointing 4-momentum, thus

$$k \cdot x = -\omega(t - z/c). \quad [1.81]$$

The second equation gives a linear system of 4 conditions on the 10 components of the matrix. In the frame where \vec{k} is along the z axis, we take as independent components

$$\tilde{h}_{0\mu}, \quad h_+ = \tilde{h}_{xx}, \quad h_\times = \tilde{h}_{xy}, \quad [1.82]$$

and then the system is solved by

$$\tilde{h}_{xz} = -\tilde{h}_{0x}, \quad \tilde{h}_{yz} = -\tilde{h}_{0y}, \quad \tilde{h}_{zz} = -\tilde{h}_{00} - 2\tilde{h}_{0z}, \quad \tilde{h}_{yy} = -\tilde{h}_+. \quad [1.83]$$

Therefore the general solution for a real, monochromatic wave with frequency ω and propagating along the z axis is a 6-parameter family given by

$$h_{\mu\nu}(x) = \begin{pmatrix} \tilde{h}_{00} & \tilde{h}_{0x} & \tilde{h}_{0y} & \tilde{h}_{0z} \\ & h_+ & h_\times & -\tilde{h}_{0x} \\ & & -h_+ & -\tilde{h}_{0y} \\ & & & -\tilde{h}_{00} - 2\tilde{h}_{0z} \end{pmatrix} \cos(k \cdot x). \quad [1.84]$$

The solution, however, contains gauge redundancy. Recall in fact that there is residual gauge freedom in the form of diffeomorphisms ξ^μ that satisfy the vacuum wave equation. In Fourier space the gauge transformation [1.52] read $\delta_\xi \tilde{h}_{\mu\nu} = 2ik_{(\mu}\tilde{\xi}_{\nu)}$, where $\tilde{\xi}^\mu(k)$ is the Fourier transform of $\xi^\mu(x)$ with $k^2 = 0$ in order to satisfy the vacuum wave equation and be an admissible residual gauge. Under the residual gauge transformation,

$$\tilde{h}_{0\mu} \rightarrow \tilde{h}_{0\mu} + ik_0\tilde{\xi}_\mu + ik_\mu\tilde{\xi}_0, \quad h_{+,\times} \rightarrow h_{+,\times}. \quad [1.85]$$

The second property follows immediately from the fact that k_μ has no transverse components. We conclude that only the two components $h_{+,\times}$ are gauge-invariant and thus physically relevant. These can be recognized as the TT components. The four coefficients $\tilde{h}_{0\mu}$ can be changed arbitrarily, and have no physical meaning. We can in particular set them to zero simultaneously, choosing $i\tilde{\xi}_\mu = \frac{c}{\omega}(\frac{1}{2}\tilde{h}_{00}, \tilde{h}_{0x}, \tilde{h}_{0y}, \tilde{h}_{0z} - \frac{1}{2}\tilde{h}_{00})$. Then

$$h_{\mu\nu}(x) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \cos(\omega(t - z/c)). \quad [1.86]$$

We have thus achieved the TT form [1.78]. In this gauge, the only components of the vacuum solutions are the gauge-invariant transverse-traceless ones. Accordingly, the solution can also be written as a sum over the two TT polarization tensors [1.65], namely

$$\tilde{h}_{\mu\nu}(k) = \sum_p A_p(k) \epsilon_{\mu\nu}^p(k), \quad [1.87]$$

where $p = +, \times$ and $A_p = h_p$.¹⁴

The resulting perturbed metric is

$$ds^2 = -dt^2 + (1 + h_+ \cos k \cdot x)dx^2 + (1 - h_+ \cos k \cdot x)dy^2 + 2h_\times \cos k \cdot x dx dy + dz^2. \quad [1.88]$$

14. The polarization decomposition can also be used to solve the wave equation. In this approach one writes [1.87] with the sum including all ten polarization modes [1.62] and ten arbitrary coefficients A_i . Then the relations [1.83] between Cartesian components are replaced by relations between the polarization coefficients. The algebra is slightly more involved but the end result is the same. We will see the use of the the polarization mode procedure on cosmological backgrounds in Sec 2.5.

Let us pause for a short historical digression, about some of the confusion that hindered historically the understanding of gravitational waves. Let us consider any member of the general solution [1.84] with $h_+ = h_\times = 0$. This looks like a genuine wave, and it is a solution of the linearized Einstein's equations. However, it is a pure gauge solution, and can be set to vanish identically without loss of physical information. In other words, there are wave solutions which are in the end only coordinate artefacts. For the same reason, it is also possible to change coordinates so that the argument $ct - z$ in its cosine is replaced by $vt - z$ for an arbitrary constant v . Hence the pure gauge modes don't really propagate, and if a gauge is chosen so that they look like they are propagating, well one can do this with an arbitrary speed, the speed is not constrained in any way by the dynamics. To use Eddington's words, the non-physical gauge modes propagate at the "speed of thought". This initial confusion was clarified by the identification of gauge-invariant components, and the fact that for the physical modes, the propagation speed is fixed to be the speed of light by Einstein's equations. However, additional doubts persisted, because the metric [1.88] is *not* a solution of the exact Einstein's equations; only of the linearized theory. In other words, there are vacuum solutions that take approximately the form [1.88] in some regions of spacetime, but none that has that exact form everywhere in spacetime. This raised the issue of whether gravitational waves existed in the full theory, or were only an artefact of the linearized approximation. This more complicated issue was solved only much later, when a non-perturbative identification of the wave degrees of freedom and their energy was made possible by the work of Bondi, Sachs and many others (Bondi 1960 ; Bondi *et al.* 1962 ; Sachs 1962 ; Newman and Penrose 1968 ; Ashtekar and Streubel 1981 ; Dray and Streubel 1984 ; Wald and Zoupas 2000). For a historical overview, see (Kennefick 1997).

Coming back to the physical solution [1.86], we see that the resulting wave tensor is transverse to the direction of propagation. Comparing with [1.62], we conclude that the physical gauge-invariant modes have helicity ± 2 . All modes of helicity ± 1 and 0 have dropped out, either by gauge fixing, or by solving the vacuum field equations. Because the physical modes have helicity ± 2 , it takes a rotation of an angle $\pi/4$ to turn one mode into the other. This can be compared with the electromagnetic case, whose modes have helicity ± 1 , and it takes a rotation by $\pi/2$ to turn one mode into the other, see Fig. 2.1.¹⁵ For the electromagnetic case, this angle $\pi/2$ corresponds to the orthogonality of the oscillations of the electric and magnetic fields. For the gravitational case, it corresponds to the antipodal symmetry in tidal forces. We will see below in Section 2.1 how this intuition can be made precise by studying the effect of a gravitational wave on test particles, see in particular Fig. 2.1.

15. The general formula is that two modes of helicity $\pm j$ are related by $\pi/2j$.

The solution [1.86] describes a monochromatic wave. The most general solution has ∞^3 Fourier components (one per choice of \vec{k}), and two independent degrees of freedom per component (the values of $h_{+, \times}$), therefore it is described by $2 \times \infty^3$ arbitrary numbers. These are the independent degrees of freedom of gravitational waves. It is the same number of the full theory, so the linearized approximation simplifies the dynamics but preserves the number of independent variations of the gravitational field that can occur in the full theory. Notice that the number of independent degrees of freedom is the same of electromagnetism, or of two scalar fields. Apart from the dynamical behaviour, what changes is also the behaviour of these degrees of freedom under Lorentz transformations, because of their different spins.

The analysis has also shown that only the transverse-traceless modes are gauge invariant. We did this using the partial gauge fixing provided by the De Donder condition, but it can be proved in full generality starting from the projector $P^{(2)\text{TT}}$ on transverse-traceless modes defined in [1.70], and observing that it annihilates gauge transformations. Since we are going to use this projector often, we drop the label **(2)** from now on. We also introduce the notation $\hat{k} = \vec{k}/k^0$, using which we can write the explicit form of the projector as

$$P^{\text{TT}ab}(\hat{k}) = \delta_{(c}^a \delta_{d)}^b - \frac{1}{2} \delta^{ab} \delta_{cd} - \delta_{(c}^a \hat{k}^b \hat{k}_{d)} - \hat{k}^a \hat{k}_{(c} \delta_{d)}^b + \frac{1}{2} (\delta^{ab} \hat{k}_c \hat{k}_d + \hat{k}^a \hat{k}^b \delta_{cd} + \hat{k}^a \hat{k}^b \hat{k}_c \hat{k}_d). \quad [1.89]$$

Note that the symmetrization on the indices here and in [1.70] is omitted in some books (Poisson and Will 2014 ; Maggiore 2007), under the premises that one is applying it to symmetric tensors only anyway. In conclusion, we do not need to use the TT gauge, nor even the De Donder gauge, in order to identify the independent degrees of freedom. Whatever gauge we are using, we can always extract them via

$$h_{ab}^{\text{TT}} = P^{\text{TT}cd}_{ab} h_{cd}, \quad [1.90]$$

and the result of this projection is gauge invariant. If we align the frame so that \hat{k} coincides with the z axis, it is given by [1.72a]. To treat a general direction, we parametrize it using polar coordinates (θ, φ) on the sphere, and write

$$\hat{k}(\theta, \varphi) = (\sin \theta \cos \varphi, \sin \theta \sin \varphi, \cos \theta) = R_z(\varphi) R_y(\theta) \hat{e}_z =: R(\theta, \varphi) \hat{e}_z. \quad [1.91]$$

We then pick a basis in the plane orthogonal to \hat{k} , given by

$$\hat{m}_1 = (\cos \theta \cos \varphi, \cos \theta \sin \varphi, -\sin \theta), \quad \hat{m}_2 = (-\sin \varphi, \cos \varphi, 0). \quad [1.92]$$

We use this basis to distinguish the two polarizations, and define

$$h_{ab}^{\text{TT}} = (\hat{m}_{1a}\hat{m}_{1b} - \hat{m}_{2a}\hat{m}_{2b})h_+ + (\hat{m}_{1a}\hat{m}_{2b} + \hat{m}_{2a}\hat{m}_{1b})h_\times. \quad [1.93]$$

We obtain in this way

$$h_+ = \frac{1}{2} \left(h_{xx}(\cos^2 \theta \cos^2 \varphi - \sin^2 \varphi) + h_{yy}(\cos^2 \theta \sin^2 \varphi - \cos^2 \varphi) + h_{zz} \sin^2 \theta \right. \\ \left. + h_{xy}(1 + \cos^2 \theta) \sin 2\varphi - h_{xz} \cos \varphi \sin 2\theta - h_{yz} \sin \varphi \sin 2\theta \right), \quad [1.94a]$$

$$h_\times = h_{xy} \cos 2\varphi \cos \theta + h_{xz} \sin \varphi \sin \theta - h_{yz} \cos \varphi \sin \theta - \frac{h_{xx} - h_{yy}}{2} \sin 2\varphi \cos \theta, \quad [1.94b]$$

which coincides with (11.46) of (Poisson and Will 2014). If the wave is travelling along the x axis for example, then $(\theta, \phi) = (\frac{\pi}{2}, 0)$, and $h_+ = \frac{1}{2}(h_{zz} - h_{yy})$, $h_\times = -h_{yz}$.

We can also use this formula to relate the polarizations in an arbitrary direction \hat{k} to the polarizations in the \hat{z} direction. In this case $h_{xx} = -h_{yy} = h_+(\hat{z})$ and $h_{xy} = h_\times(\hat{z})$, while all other components vanish. The expressions then simplify to

$$h_+(\hat{k}) = \frac{1 + \cos^2 \theta}{2} (h_+(\hat{z}) \cos 2\varphi + h_\times(\hat{z}) \sin 2\varphi), \quad [1.95a]$$

$$h_\times(\hat{k}) = \cos \theta (h_\times(\hat{z}) \cos 2\varphi - h_+(\hat{z}) \sin 2\varphi). \quad [1.95b]$$

In physical applications, the coordinate system may be set from characteristics of the system, for instance for binaries, adapted so that the orbital plane coincides with the (x, y) plane, and x is in the direction of the pericenter. Then our rotation identifies θ with the inclination ι of the orbital plane, and the y axis is along the ‘line of nodes’, namely the intersection between the orbital plane and the plane of sight, pointing in the direction of the descendent node. On the other hand, astronomers often choose the convention to align the line of nodes along the x axis, pointing in the direction of the ascendent node. This means that in our reference system, the ‘longitude of ascending node’ Ω is $-\pi/2$, and the ‘longitude of pericenter’ ω taken from the ascending node is $\varphi - \pi/2$. That is,

$$\iota = \theta, \quad \omega = \frac{\pi}{2} - \varphi, \quad \Omega = -\frac{\pi}{2}. \quad [1.96]$$

This maps our basis (1.91-1.92) to (3.44-3.45) of (Poisson and Will 2014). See Fig. 1.4.




Figure 1.4: Rotation between the orbital plane of a binary system and the direction of the observer. The black reference frame is adapted to the orbit, with \hat{z} aligned with the angular momentum, and \hat{x} with the pericenter P . The first rotation along \hat{y} by an angle θ rotates the frame to the blue one, and the second rotation along \hat{z} by an angle φ rotates it to the final red frame. The red frame has \hat{z} pointing in the direction of the observer, and \hat{y} axis along the line of nodes, pointing towards the descending node. In the red frame, the longitude of the ascending node A is $\Omega = 3\pi/2$, and the longitude of pericenter is ω .

1.3.6. Gauge-invariant description: independent and constrained degrees of freedom

The analysis of vacuum solutions has allowed us to identify the independent degrees of freedom of the gravitational field. There are also dependent degrees of freedom, namely components of the field that are gauge invariant, but uniquely determined from the sources. These can be exposed looking at the constraints, namely the 00 and $0a$ components of [1.49] which give rise to elliptic equations, as opposed to hyperbolic ones. The intricacies of the tensorial structure make it however convenient to do first a kinematical analysis of gauge invariance using the transverse and longitudinal projectors. To that end, we first observe that the projectors can be described also in configuration space without doing the Fourier transform. In this case [1.66] is replaced by

$$T_b^a = \delta_b^a - \frac{\partial^a \partial_b}{\bar{\partial}^2}, \quad L_b^a := \frac{\partial^a \partial_b}{\bar{\partial}^2}. \quad [1.97]$$

These expressions are somewhat implicit because one needs to specify boundary conditions in order to have a well-defined inverse of the Laplace operator

$\vec{\partial}^2$. Requiring the fields to vanish at spatial infinity makes this definition equivalent to the one in momentum space. Notice that the transverse-longitudinal decomposition in configuration space can be recognized as the Helmholtz decomposition of a 3d vector field into solenoidal and irrotational parts. This representation also makes it clear that the projectors are local operators in Fourier space, but non-local in spacetime. It has the important consequence that it is not possible to identify exactly the gauge-invariant modes with local observations only.

It is also convenient to dispose of the projectors for the components with non-maximal helicities, by introducing auxiliary fields with smaller spin. For instance, we denote $W_a := h_{0a}$ the spin-1 part of the gravitational perturbation, and write its longitudinal part as the gradient of a scalar:

$$W_a^L := P_L^{(1)a} W^b = \partial_a W, \quad W = \vec{\partial}^{-2} \partial_a W^a. \quad [1.98]$$

Similarly for the spin-2 part, we can write the mixed transverse-longitudinal and fully longitudinal modes introducing a transverse vector B^a and a scalar B ,

$$\begin{aligned} P_L^{(2)ab} h^{cd} &= 2\partial^{(a} T_c^{b)} \frac{\partial_d h^{cd}}{\vec{\partial}^2} = 2\partial^{(a} B^{b)}, \quad B^b := \frac{2}{\vec{\partial}^2} T_c^b \partial_d h^{cd}, \\ P_{LL}^{(2)ab} h^{cd} &= \frac{3}{2} \left(\frac{\partial^a \partial^b}{\vec{\partial}^2} - \frac{1}{3} \delta^{ab} \right) \frac{\partial_{(c} \partial_{d)} h^{cd}}{\vec{\partial}^2} = (\partial_a \partial_b - \frac{1}{3} \delta_{ab} \vec{\partial}^2) B, \\ B &:= \frac{3}{2} \frac{\partial_{(c} \partial_{d)} h^{cd}}{\vec{\partial}^4}. \end{aligned}$$

Let us summarize. With this new notation at hand, the gravitational perturbation around Minkowski can be decomposed into

$$\text{spin 0} \quad h_s \quad [1.99a]$$

$$\text{spin 0} \quad h_{00} \quad [1.99b]$$

$$\text{spin 1} \quad h_{0a} = W_a = W_a^T + \partial_a W \quad [1.99c]$$

$$\text{spin 2} \quad h_{ab} = h_{ab}^{TT} + 2\partial_{(a} B_{b)} + (\partial_a \partial_b - \frac{1}{3} \delta_{ab} \vec{\partial}^2) B + \frac{1}{3} \delta_{ab} h_s, \quad [1.99d]$$

with

$$W_a^T = T_a^b W_b \quad \text{helicity } \pm 1 \quad [1.100a]$$

$$W = \vec{\partial}^{-2} \partial_a W^a \quad \text{helicity } 0, \quad [1.100b]$$

and

$$h_{ab}^{\text{TT}} := P_{\text{TT}}^{(2)ab} h^{cd} = \left(T_c^a T_d^b - \frac{1}{2} T^{ab} T_{cd} \right) h^{cd} \quad \text{helicity } \pm 2 \quad [1.101a]$$

$$B^b := \frac{2}{\bar{\partial}^2} T_c^b \partial_a h^{cd} \quad \text{helicity } \pm 1 \quad [1.101b]$$

$$B := \frac{3}{2} \frac{\partial_{(c} \partial_{d)}}{\bar{\partial}^4} h^{cd} \quad \text{helicity } 0 \quad [1.101c]$$

Next, we look at the behaviour of these different helicities under gauge transformations. We have used the Poincaré symmetry of the Minkowski background to organize the ten components of $h_{\mu\nu}$ in terms of spin and helicity. But Minkowski is not invariant under the general diffeomorphism symmetry [1.52], hence there is no reason to expect that this decomposition be gauge-invariant. Consider for instance the spin-0 part h_{00} : this is just a metric component, and manifestly not invariant under diffeomorphisms. Only the TT component is gauge-invariant, as can be seen explicitly replacing $h_{\mu\nu}$ in [1.101a] with its gauge transformation [1.52] and observing that it vanishes thanks to the transverse projection on both indices. This is the gauge-invariance of the two TT components already observed in (the second equation in) [1.83]. The remaining components transform as follows:

$$\begin{aligned} h'_{00} &= h_{00} + 2\partial_0 \xi_0, & h'_s &= h_s + 2\bar{\partial}^2 \xi^L, & \xi^L &:= \bar{\partial}^{-2} \partial_a \xi^a, \\ W'_a &= W_a + \partial_0 \xi_a^T, & W' &= W + \xi_0 + \partial_0 \xi^L, \\ B'_a &= B_a + \xi_a^T, & B' &= B + 2\xi^L. \end{aligned}$$

It is possible to combine them to find a maximal number of 6 gauge-invariant quantities, given by

$$\begin{aligned} h_{ab}^{\text{TT}}, & \quad c^{-2} \Phi := -\frac{1}{2} h_{00} + \partial_0 W - \frac{1}{2} \partial_0^2 B, \\ c^{-3} \Phi_a &:= \frac{1}{4} (W_a^T - \partial_0 B_a), & c^{-2} \Psi &:= \frac{1}{6} (\bar{\partial}^2 B - h_s). \end{aligned} \quad [1.102]$$

The numerical factors and powers of c have been chosen for later convenience. We will also use a dot for the derivative with respect to the time coordinate $t = x^0/c$, e.g. $\partial_0 W = c^{-1} \dot{W}$. Our conventions for the physical dimensions are summarized as follows:

$$\begin{aligned} [x^\mu] &= \mathbf{m} & [t = x^0/c] &= \mathbf{s} & [\partial_\mu] &= \mathbf{m}^{-1} & [\partial_t = c\partial_0] &= \mathbf{s}^{-1} \\ [g_{\mu\nu}] &= [h_{\mu\nu}] = [h] = [h_s] = [W^T] & & & & & & = 1 \\ [W] = [B_a] &= \mathbf{m} & [B] &= \mathbf{m}^2 & [\Phi] = [\Psi] &= \mathbf{m}^2 \mathbf{s}^{-2} & [\Phi_a] &= \mathbf{m}^3 \mathbf{s}^{-3} \end{aligned}$$

Inserting the parametrization [1.99] in the linearized Einstein's equation [1.49] one finds that gauge dependent quantities drop out and only the gauge-independent ones remain. This allows us to decouple the tensorial equations into two sets, an hyperbolic one featuring the d'Alambertian operator alone, and an elliptic one featuring the Laplace operator alone:

$$\square h_{ab}^{\text{TT}} = -\frac{16\pi G}{c^4} \sigma_{ab}, \quad [1.103a]$$

$$\bar{\partial}^2 \Psi = 4\pi G \rho, \quad \bar{\partial}^2 \Phi_a = 4\pi G s_a, \quad \bar{\partial}^2 (\Phi - \Psi) = \frac{12\pi G}{c^2} \left(\dot{s} + \frac{1}{3} \tau \right). \quad [1.103b]$$

The sources on the right-hand side of these equations are the components of $T_{\mu\nu}$ projected in the same way as [1.99], namely

$$T_{00} = c^2 \rho, \quad T_{0a} = c(s_a + \partial_a s), \\ T_{ab} = \sigma_{ab} + 2\partial_{(a} \sigma_{b)} + \left(\partial_a \partial_b - \frac{1}{3} \delta_{ab} \bar{\partial}^2 \right) \sigma + \frac{1}{3} \delta_{ab} \tau.$$

Rewriting the linearized Einstein equations [1.49] in the equivalent form [1.103] makes their three-sided structure, mentioned in subsection 1.2.1, manifest. Four equations were redundant and have dropped out, hence the field has four undetermined components, which can be assigned arbitrarily choosing a specific gauge. Four equations are elliptic, and describe constrained degrees of freedom, namely dynamical components of the gravitational field which are uniquely determined by the sources. Finally, two equations are hyperbolic, hence they contain free data, and describe how these independent degrees of freedom propagate and react to the sources. This analysis therefore establishes that the gravitational field has two independent degrees of freedom, which are carried by the two h_{ab}^{TT} components of the metric. These are gauge-invariant, and describe the ± 2 helicities of a spin-2 wave.

As for the remaining gauge-invariant components, we have seen that they satisfy Poisson equations, hence these are degrees of freedom that are entirely determined by the sources. For this reason they are sometimes called 'Coulombic' degrees of freedom. Their meaning can be elucidated looking at the post-Newtonian expansion, in which sources are moving slowly with respect to the speed of light. To begin with, let us first consider perfectly static sources.

For static sources in a given frame, $T_{00} = c^2 \rho$ is the only non-vanishing component on the right-hand side of Einstein's equations. Then the second

and third equations in [1.103b] imply¹⁶ $\Phi_a = 0$ and $\Psi = \Phi$, whilst the first gives Newton's equation, and we can identify Φ with Newton's potential. This is also the way in which one fixes the coupling constant of the full Einstein's equations in terms of G and c . How about the other potentials? They are sourced by moving bodies, and their existence is a consequence of relativistic invariance, akin to the electromagnetic occurrence of the vector potential next to the Coulomb potential. They describe effects that collectively go under the name of 'gravito-magnetism'. These include additional contributions to the precessions of equinoxes, light bending, and frame dragging or Lense-Thirring effect. Their effect can be studied looking at the Lagrangian for a massive test particle, which gives

$$\begin{aligned}\mathcal{L} &= -mc\sqrt{-g_{\mu\nu}u^\mu u^\nu} = -mc\sqrt{-(\eta_{\mu\nu} + h_{\mu\nu})u^\mu u^\nu} \\ &= -mc^2 + \frac{1}{2}mv^2 - m\Phi + \\ &\quad + \frac{m}{c^2} \left(\frac{1}{8}v^4 + \frac{1}{2}\Phi^2 - \frac{1}{2}v^2\Phi - \Psi v^2 + 4\Phi_a v^a + c^2 h_{ab}^{\text{TT}} v^a v^b \right) + O(c^{-4})\end{aligned}$$

where $v^a = dx^a/dt$. The term in round bracket is the *first post-Newtonian correction*. As we will see below, $h^{\text{TT}} \sim c^{-4}$, hence the last term there is higher order: dissipative effects for massive particles only appear at 2PN.

1.3.7. Gauge-fixing with sources

Solving the decoupled equations [1.103] determines the gauge-invariant quantities [1.102] in terms of the sources and the initial conditions. These solutions *do not determine a metric*. Doing so requires the additional step of specifying the coordinates to be used. Only after the coordinates are given, or in other words only after a gauge is chosen, the physical degrees of freedom can be described in terms of a metric tensor.

The gauge-invariant approach is conceptually satisfying because it identifies the physical degrees of freedom and decouples the equations, making them easier to solve in principle. However, it is very limited in applicability. Firstly, the decoupling and simple identification of gauge-invariant quantities occur only for very special backgrounds, such as flat spacetime or homogeneous and

¹⁶ Assuming trivial boundary conditions, see App 1.5.

isotropic.¹⁷ Secondly, even when the background is Minkowski, a general identification of gauge-invariant quantities at a fixed order in perturbation theory is only possible at the linear level, as explained earlier.

To go beyond these limitations, it is easier to put to the side the gauge-invariant description, and work instead in a fixed gauge. In the gauge-fixed approach, one chooses coordinates that impose restrictions on the metric, and then solves for individual metric components in that gauge, like we did in Section 1.3.4. Namely we do not solve [1.103], but the original system [1.49], coupled to additional equations fixing the gauge. The additional equations remove the problem that the field equations are redundant and do not determine all metric components.

A simple example of gauge fixing is the temporal gauge

$$h_{00} = h_{0a} = 0. \quad [1.104]$$

This is the linearized version of the non-perturbative temporal gauge described in Section 1.2.3, and provides a complete gauge fixing of the 4-dimensional diffeomorphism symmetry. In this gauge the only non-trivial components of the metric are the spatial ones, and [1.102] reduces to

$$c^{-2}\Phi = -\frac{1}{2}\partial_0^2 B, \quad c^{-3}\Phi_a = -\frac{1}{4}\partial_0 B_a, \quad c^{-2}\Psi = \frac{1}{6}(\partial^2 B - h_s). \quad [1.105]$$

All gauge-invariant potentials are encoded in the components of the spin-2 mode [1.99d]. A related choice would be to replace the condition on h_{00} with the trace condition $h = 0$. In this case

$$\begin{aligned} c^{-2}(\Phi + 3\Psi) &= -h_{00} + \frac{1}{2}\square B, & c^{-3}\Phi_a &= -\frac{1}{4}\partial_0 B_a, \\ c^{-2}(\Phi - 3\Psi) &= -\frac{1}{2}(\partial_0^2 + \partial^2)B. \end{aligned} \quad [1.106]$$

This time it is the four metric components B , B_a and h_{00} that are fixed in terms of the sources, while h_{0a} and h vanish.

17. On a general background, one can still build the analogue of the spin-helicity decomposition replacing the partial derivatives with covariant derivatives, although care is needed to invert the Laplacian and handle its non-commutativity with the covariant derivative. However, the metric now enters explicitly the decomposition of the energy-momentum tensor, hence the decoupling will be lost in general. The non-commutativity of covariant derivatives also hinders the identification of gauge-invariant quantities.

Another simple gauge fixing is the gravitational equivalent of the Coulomb gauge, defined so that both spin-2 and spin-1 parts are purely transverse, namely

$$\partial_a h^a{}_0 = 0, \quad \partial_a h^a{}_b = \frac{1}{3} \partial_b h_s \quad \Leftrightarrow \quad B_a = B = W = 0. \quad [1.107]$$

In this gauge the non-trivial components of the metric are

$$h_{00}, \quad h_{0a}^T = W_a^T, \quad h_{ab} = h_{ab}^{TT} + \frac{1}{3} \delta_{ab} h_s, \quad [1.108]$$

and their relation to the gauge invariant potential is given by

$$c^{-2} \Phi = -\frac{1}{2} h_{00}, \quad c^{-3} \Phi_a = \frac{1}{4} W_a^T, \quad c^{-2} \Psi = -\frac{1}{6} h_s. \quad [1.109]$$

These examples have the advantage that the gauge-invariant potentials are identified with individual components of the metric. On the other hand, the radiative gauge-invariant modes h_{ab}^{TT} are complicated functions of the metric components. When fixing the coordinate gauge, one should keep in mind that some choices may be better than others, as we discussed with the Schwarzschild example in Section 1.2.3.¹⁸ The examples above are non-covariant with respect to the Lorentz symmetry of the background, because they make reference to a given time foliation, and treat time and space components differently. In the presence of radiation, it is best to use a covariant gauge, because it simplifies the analysis of the solutions. We can, in fact, remark that while the above examples simplify the relation between the potentials and individual metric components, the remaining field equations for the propagating degrees of freedom are complicated. Whereas with the covariant De Donder gauge, all field equations took the simpler form [1.76]. However while the De Donder condition [1.75] can be also imposed in the presence of sources, we can no longer select a unique representative satisfying the TT condition [1.78]. If need be, a unique representative of the De Donder gauge in the presence of sources can be specified as follows. The general solution of [1.76] is

$$\bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} \int d^4 x' G(x, x') T_{\mu\nu}(x') + \bar{h}_{\mu\nu}^{\circ}, \quad [1.110]$$

18. If $GM/c^2 \ll r$ in the static spherical coordinates, we can treat it as a perturbative solution with $h_{\mu\nu} = 2GM/(c^2 r)(\delta_\mu^t \delta_\nu^t + \delta_\mu^r \delta_\nu^r)$. We then see that these coordinates correspond to the Coulomb gauge fixing. Changing coordinates so to have temporal gauge with $h_{00} = 0$ would push the Newton potential in the B component, and introduce a dependence on the t coordinate, as mentioned earlier.

where $G(x, x')$ is the d'Alembertian's Green function (see Appendix 1.5), $\bar{h}_{\mu\nu}^\circ$ any solution of the homogeneous equation, and the gauge condition is maintained via $\partial_\mu T^{\mu\nu} = \partial_\mu \bar{h}^{\circ\mu\nu} = 0$. Since both $\bar{h}_{\mu\nu}^\circ$ and the residual diffeomorphism parameters ξ^μ satisfy the vacuum wave equation, we can use the residual freedom to set to zero any four components of $\bar{h}_{\mu\nu}^\circ$. For instance, we can choose $\bar{h}^\circ = 0$, so that $\bar{h}_{\mu\nu}^\circ = \bar{h}_{\mu\nu}^\circ$, as well as $\bar{h}_{0a}^\circ = 0$. Then the De Donder condition implies that four more components of the homogeneous solution also vanish everywhere, specifically \bar{h}_{00}° and $\partial_b \bar{h}^{\circ ab}$. At this stage the gauge is completely fixed, and the only components left in the homogeneous solution are the gauge-invariant ones $\bar{h}_{\mu\nu}^{\circ\text{TT}}$, which carry the independent degrees of freedom of gravitational waves. The general solution is

$$\bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} \int d^4x' G(x, x') T_{\mu\nu}(x') + h_{\mu\nu}^{\circ\text{TT}}. \quad [1.111]$$

It is a complete gauge fixing that singles out a unique element of the De Donder family, and reduces to [1.78] for vacuum solutions. For generic sources, all components of the metric perturbation are non-zero in this gauge.¹⁹

Alternatively, it is also possible to select a unique representative of the De Donder gauge requiring that the metric components $h_{00} = h_{0a}$ and their derivatives vanish on a given initial value surface, as considered in (Wald 1984). However these metric components will remain zero only in the region outside the causal domain of the sources from the initial value surface. This prescription thus achieves a TT gauge valid also in the presence of sources, but only in a finite region of spacetime. A similar construction of TT gauge in a local region of finite time outside the sources is presented in (Flanagan and Hughes 2005). While interesting in principle, the 'local TT gauge' seems to be of minor practical use. The procedure described in the previous paragraph achieves instead a structure of the solutions valid globally and at all times, and it is the analogue of the complete gauge-fixing used for instance in electro-magnetism by the Lienard-Wiechert potentials of a moving charge.

A special situation occurs if the sources are static. In this case, it is possible to specialize the De Donder gauge to the Coulomb gauge [1.107]. To prove this, we first observe that when rewritten in terms of the gauge-invariant quantities

19. More precisely, in the region causally connected to the sources, since the retarded Green function vanishes outside the light cone. So if the sources are present at all times, the metric perturbation is non-zero everywhere, whereas if the sources are 'turned on' at some initial time, then the perturbation vanishes outside the causal domain of the sources from that initial time.

[1.102], the De Donder condition [1.75] implies

$$\square W = -\frac{4}{c^3}\dot{\Psi}, \quad \square B = \frac{2}{c^2}(\Phi - \Psi), \quad \square B_a = \frac{4}{c^4}\dot{\Phi}_a. \quad [1.112]$$

For static sources $\Phi_a = \dot{\Phi} = 0$ and $\Psi = \Phi$, hence all right-hand sides vanish. It is then possible to specialize the De Donder gauge choosing $W = B = B_a = 0$ everywhere. Having done so, $h_{00} = -2c^{-2}\Phi = h_s/3$. We have thus achieved the Coulomb gauge described earlier, as opposed to the TT gauge, and again this is a complete gauge fixing singling out a member of the De Donder family. In other words, the De Donder gauge is compatible with the Coulomb gauge for static sources. Notice that this is what happens also in electromagnetism, where the Lorenz gauge is compatible with the Coulomb gauge for static sources.

We conclude with a word on another important gauge fixing that can be used to study radiation including sources, the Bondi, or Bondi-Sachs gauge. In this case the background Minkowski metric is not in Cartesian coordinates, but rather spherical coordinates, and time is replaced with retarded time $u = t - r$. In these coordinates the Minkowskian Christoffel symbols do not vanish. However a crucial advantage is that it is a gauge that can be defined also in the full theory, and allows a non-perturbative description of gravitational waves, at the price of introducing an asymptotic expansion away from the sources.

1.4. Appendix: Second order action for perturbations around any background solution

This appendix gives the details of the steps required to calculate the Einstein-Hilbert action of GR to second order in perturbations about a general background metric $\bar{g}_{\mu\nu}$:

$$g_{\mu\nu} = \bar{g}_{\mu\nu} + h_{\mu\nu}. \quad [1.113]$$

The starting point is

$$S_{\text{EH}} \equiv \int d^4x \sqrt{-g} R = S^{(0)} + S^{(1)} + S^{(2)} + \dots \quad [1.114]$$

where $S^{(i)}$ is of $\mathcal{O}(h^i)$. The extremisation of the second order action $S^{(2)}$ with respect to $h_{\mu\nu}$ will give the linearized equations of motion for $h_{\mu\nu}$, namely those we want to calculate. The background metric $\bar{g}_{\mu\nu}$ satisfies the background Einstein equations which follow from $S^{(1)}$.

To find $S^{(2)}$ there are two main steps: calculating the Ricci scalar R to second order, and then the determinant of the metric and hence $\sqrt{-g}$ to second order.

Perturbed Riemann tensor, Ricci tensor and scalar, Einstein tensor

- The inverse metric, or *contravariant metric tensor* corresponding to Eq. [1.113] is given at second order by

$$g^{\mu\nu} = \bar{g}^{\mu\nu} - h^{\mu\nu} + h^{\mu\rho}h_{\rho}^{\nu} \quad [1.115]$$

where indices of $h_{\mu\nu}$ are raised and lowered with $\bar{g}_{\mu\nu}$.

- The perturbed *Christoffel symbols* are given by

$$\Gamma_{\mu\nu}^{\rho} \equiv \frac{1}{2}g^{\rho\lambda}(\partial_{\mu}g_{\nu\lambda} + \partial_{\nu}g_{\mu\lambda} - \partial_{\lambda}g_{\mu\nu}) = \bar{\Gamma}_{\mu\nu}^{\rho} + \Gamma_{\mu\nu}^{\rho(1)} + \Gamma_{\mu\nu}^{\rho(2)}. \quad [1.116]$$

In the following we denote the covariant derivative with respect to \bar{g} by $\bar{\nabla}$ (with of course $\bar{\nabla}_{\mu}\bar{g}_{\nu\alpha} = 0$). Substitution of Eqs. [1.113] and [1.115] gives

$$\begin{aligned} \bar{\Gamma}_{\mu\nu}^{\rho} &= \frac{1}{2}\bar{g}^{\rho\lambda}(\partial_{\mu}\bar{g}_{\nu\lambda} + \partial_{\nu}\bar{g}_{\mu\lambda} - \partial_{\lambda}\bar{g}_{\mu\nu}) \\ \Gamma_{\mu\nu}^{\rho(1)} &= \frac{1}{2}\bar{g}^{\rho\lambda}(\bar{\nabla}_{\mu}h_{\nu\lambda} + \bar{\nabla}_{\nu}h_{\mu\lambda} - \bar{\nabla}_{\lambda}h_{\mu\nu}) \end{aligned} \quad [1.117]$$

$$\Gamma_{\mu\nu}^{\rho(2)} = -h_{\beta}^{\rho}\Gamma_{\mu\nu}^{\beta(1)}. \quad [1.118]$$

- Next we calculate the *Riemann tensor*. Let

$$\Gamma_{\mu\nu}^{\rho} = \bar{\Gamma}_{\mu\nu}^{\rho} + \delta\Gamma_{\mu\nu}^{\rho} \quad [1.119]$$

where, from (1.116),

$$\delta\Gamma_{\mu\nu}^{\rho} = \Gamma_{\mu\nu}^{\rho(1)} + \Gamma_{\mu\nu}^{\rho(2)}. \quad [1.120]$$

The definition of the Riemann tensor together with (1.119) gives

$$R^{\mu}{}_{\nu\rho\sigma} \equiv \partial_{\rho}\Gamma_{\sigma\nu}^{\mu} + \Gamma_{\rho\lambda}^{\mu}\Gamma_{\sigma\nu}^{\lambda} - (\rho \leftrightarrow \sigma) \quad [1.121]$$

$$\begin{aligned} &= \bar{R}^{\mu}{}_{\nu\rho\sigma} + \bar{\nabla}_{\rho}(\delta\Gamma_{\sigma\nu}^{\mu}) - \bar{\nabla}_{\sigma}(\delta\Gamma_{\rho\nu}^{\mu}) \\ &\quad + (\delta\Gamma_{\rho\lambda}^{\mu})(\delta\Gamma_{\sigma\nu}^{\lambda}) - (\delta\Gamma_{\sigma\lambda}^{\mu})(\delta\Gamma_{\rho\nu}^{\lambda}) \end{aligned} \quad [1.122]$$

On writing

$$R^\mu{}_{\nu\rho\sigma} \equiv \bar{R}^\mu{}_{\nu\rho\sigma} + R^{(1)\mu}{}_{\nu\rho\sigma} + R^{(2)\mu}{}_{\nu\rho\sigma}$$

and using (1.120) one can read off the different orders of the Riemann tensor. To first order

$$R^{(1)\mu}{}_{\nu\rho\sigma} = \bar{\nabla}_\rho \Gamma_{\sigma\nu}^{\mu(1)} - \bar{\nabla}_\sigma \Gamma_{\rho\nu}^{\mu(1)} \quad [1.123]$$

$$\begin{aligned} &= \frac{1}{2} [(\bar{\nabla}_\rho \bar{\nabla}_\sigma - \bar{\nabla}_\sigma \bar{\nabla}_\rho) h_\nu^\mu + (\bar{\nabla}_\rho \bar{\nabla}_\nu h_\sigma^\mu - \bar{\nabla}_\sigma \bar{\nabla}_\nu h_\rho^\mu) \\ &\quad - (\bar{\nabla}_\rho \bar{\nabla}^\mu h_{\sigma\nu} - \bar{\nabla}_\sigma \bar{\nabla}^\mu h_{\rho\nu})] \end{aligned} \quad [1.124]$$

where we have used Eq. [1.117]. The second order term follows from (1.122) and (1.120) and reads

$$R^{(2)\mu}{}_{\nu\rho\sigma} = (\bar{\nabla}_\rho \Gamma_{\sigma\nu}^{\mu(2)} - \bar{\nabla}_\sigma \Gamma_{\rho\nu}^{\mu(2)}) + (\Gamma_{\rho\lambda}^{\mu(1)} \Gamma_{\sigma\nu}^{\lambda(1)} - \Gamma_{\sigma\lambda}^{\mu(1)} \Gamma_{\rho\nu}^{\lambda(1)}). \quad [1.125]$$

• The *Ricci tensor* is then obtained by contraction:

$$R_{\nu\sigma} \equiv R^\mu{}_{\nu\mu\sigma} = \bar{R}_{\nu\sigma} + R_{\nu\sigma}^{(1)} + R_{\nu\sigma}^{(2)}. \quad [1.126]$$

From (1.124) it follows that

$$R_{\nu\sigma}^{(1)} = \frac{1}{2} [\bar{\nabla}_\mu \bar{\nabla}_\sigma h_\nu^\mu + \bar{\nabla}_\mu \bar{\nabla}_\nu h_\sigma^\mu - \bar{\square} h_{\nu\sigma} - \bar{\nabla}_\sigma \bar{\nabla}_\nu h], \quad [1.127]$$

whereas from (1.125)

$$R_{\nu\sigma}^{(2)} = \bar{\nabla}_\rho \Gamma_{\sigma\nu}^{\rho(2)} - \bar{\nabla}_\sigma \Gamma_{\rho\nu}^{\rho(2)} + \Gamma_{\rho\lambda}^{\rho(1)} \Gamma_{\sigma\nu}^{\lambda(1)} - \Gamma_{\sigma\lambda}^{\rho(1)} \Gamma_{\rho\nu}^{\lambda(1)} \quad [1.128]$$

Its explicit form is not required below, but for completeness we give it here:

$$\begin{aligned} \bar{R}_{\nu\sigma}^{(2)} &= \frac{1}{4} \bar{\nabla}_\nu h_{\alpha\beta} \bar{\nabla}_\sigma h^{\alpha\beta} + \frac{1}{2} \bar{\nabla}^\beta h_\sigma^\beta (\bar{\nabla}_\beta h_{\alpha\nu} - \bar{\nabla}_\alpha h_{\beta\nu}) \\ &\quad + \frac{1}{2} h^{\alpha\beta} (\bar{\nabla}_\nu \bar{\nabla}_\sigma h_{\alpha\beta} + \bar{\nabla}_\beta \bar{\nabla}_\alpha h_{\nu\sigma} - \bar{\nabla}_\beta \bar{\nabla}_\sigma h_{\alpha\nu} - \bar{\nabla}_\beta \bar{\nabla}_\nu h_{\alpha\sigma}) \\ &\quad - \frac{1}{2} (\bar{\nabla}_\beta h^{\alpha\beta} - \frac{1}{2} \bar{\nabla}^\alpha h) (\bar{\nabla}_\sigma h_{\nu\alpha} + \bar{\nabla}_\nu h_{\alpha\sigma} - \bar{\nabla}_\alpha h_{\nu\sigma}). \end{aligned}$$

- The *Ricci scalar* is obtained from Eqs. [1.115] and [1.126] and is given by

$$\begin{aligned} R &\equiv g^{\nu\sigma} R_{\nu\sigma} = (\bar{g}^{\nu\sigma} - h^{\nu\sigma} + h^{\nu\rho} h_{\rho}^{\sigma})(\bar{R}_{\nu\sigma} + R_{\nu\sigma}^{(1)} + R_{\nu\sigma}^{(2)}) \\ &= \bar{R} + R^{(1)} + R^{(2)}, \end{aligned} \quad [1.129]$$

where

$$R^{(1)} = \bar{g}^{\nu\sigma} R_{\nu\sigma}^{(1)} - \bar{R}^{\nu\sigma} h_{\nu\sigma} = \bar{\nabla}_{\mu} \bar{\nabla}_{\nu} h^{\mu\nu} - \bar{\square} h - \bar{R}^{\nu\sigma} h_{\nu\sigma} \quad [1.130]$$

(in the last line we have used (1.127)), and

$$R^{(2)} = h^{\nu\alpha} h_{\alpha}^{\sigma} \bar{R}_{\nu\sigma} - h^{\nu\sigma} R_{\nu\sigma}^{(1)} + \bar{g}^{\nu\sigma} R_{\nu\sigma}^{(2)}. \quad [1.131]$$

- For completeness we also give the first order perturbed *Einstein tensor*

$$G_{\mu\nu}^{(1)} = R_{\mu\nu}^{(1)} - \frac{1}{2} \bar{R} h_{\mu\nu} - \frac{1}{2} \bar{g}_{\mu\nu} R^{(1)} \quad [1.132]$$

$$\begin{aligned} &= -\frac{1}{2} \bar{\square} h_{\mu\nu} + \bar{\nabla}_{(\mu} \bar{\nabla}_{\rho} h^{\rho}_{\nu)} - \frac{1}{2} \bar{\nabla}_{\mu} \bar{\nabla}_{\nu} h - \frac{1}{2} \bar{g}_{\mu\nu} (\bar{\nabla}_{\mu} \bar{\nabla}_{\nu} h^{\mu\nu} - \bar{\square} h) \\ &\quad + \bar{G}_{\rho(\mu} h_{\nu)\rho} - (\bar{R}_{\mu\rho\nu\sigma} - \frac{1}{2} \bar{g}_{\mu\nu} \bar{R}_{\rho\sigma}) h^{\rho\sigma}. \end{aligned} \quad [1.133]$$

Perturbed metric determinant

To expand $\sqrt{-g}$ to second order, we write Eq. [1.113] $g_{\alpha\beta} = \bar{g}_{\alpha\mu} (\delta_{\beta}^{\mu} + M_{\beta}^{\mu})$ where $M_{\beta}^{\mu} = \bar{g}^{\mu\lambda} h_{\lambda\beta}$. Thus

$$\det(g) = \det(\bar{g}) \det(\mathbf{1} + \mathbf{M}) \quad [1.134]$$

where the matrix \mathbf{M} has components M_{β}^{μ} . To quadratic order

$$\det(\mathbf{1} + \mathbf{M}) = 1 + \text{tr}\mathbf{M} + \frac{1}{2} (2(\text{tr}\mathbf{M})^2 - \text{tr}(\mathbf{M}^2)) + \dots \quad [1.135]$$

Replacing $\text{tr}\mathbf{M} = h$ and $\text{tr}(\mathbf{M}^2) = h_{\mu\nu} h^{\mu\nu}$ in Eq. [1.134], and then taking the square root, gives

$$\sqrt{-\det(g_{\mu\nu})} = \sqrt{-\det(\bar{g}_{\mu\nu})} \left[1 + \frac{1}{2} h + \frac{1}{8} (h^2 - 2h_{\mu\nu}^2) \right]. \quad [1.136]$$

EH action to second order

Substituting (1.129) and (1.136) into the perturbed Einstein Hilbert action Eq. [1.114] gives

$$\begin{aligned} S_{(0)} &= \int d^4x \sqrt{-\bar{g}} \bar{R} \\ S_{(1)} &= \int d^4x \sqrt{-\bar{g}} \left(R^{(1)} + \frac{h}{2} \bar{R} \right) \\ S_{(2)} &= \int d^4x \sqrt{-\bar{g}} \left(R^{(2)} + \frac{h}{2} R^{(1)} + \frac{\bar{R}}{8} (h^2 - 2h_{\mu\nu} h^{\mu\nu}) \right). \end{aligned}$$

- The *first order action* is

$$S_{(1)} = \int d^4x \sqrt{-\bar{g}} \left(\bar{\nabla}_\mu \bar{\nabla}_\nu h^{\mu\nu} - \square h - \bar{R}^{\nu\sigma} h_{\nu\sigma} + \frac{h}{2} \bar{R} \right)$$

The first two terms are total derivatives. After integration by parts and dropping the boundary terms

$$S_{(1)} = \int d^4x \sqrt{-\bar{g}} (-\bar{G}^{\nu\sigma} h_{\nu\sigma}). \quad [1.137]$$

On including matter through the stress tensor, the variation of this gives the background Einstein equation.

- The *second order* part becomes, on substituting [1.131],

$$\begin{aligned} S_{(2)} &= \int d^4x \sqrt{-\bar{g}} \left(\bar{g}^{\nu\sigma} R_{\nu\sigma}^{(2)} + \frac{h}{2} R^{(1)} - h^{\nu\sigma} R_{\nu\sigma}^{(1)} \right. \\ &\quad \left. + \left[h^{\nu\alpha} h_\alpha^\sigma \bar{R}_{\nu\sigma} + \frac{\bar{R}}{8} (h^2 - 2h_{\mu\nu} h^{\mu\nu}) \right] \right) \end{aligned} \quad [1.138]$$

where $R^{(1)}$ and $R_{\nu\sigma}^{(1)}$ are given in [1.130] and [1.127] respectively.

The first term $\int d^4x \sqrt{-\bar{g}} \bar{g}^{\nu\sigma} R_{\nu\sigma}^{(2)}$ splits into four parts on using (1.128). The first two parts are total derivatives and do not contribute. The last two parts

give

$$\begin{aligned} \int d^4x \sqrt{-\bar{g}} (\bar{g}^{\nu\sigma} R_{\nu\sigma}^{(2)}) &= + \int d^4x \sqrt{-\bar{g}} \bar{g}^{\nu\sigma} \left(\Gamma_{\rho\lambda}^{\mu(1)} \Gamma_{\sigma\nu}^{\lambda(1)} - \Gamma_{\sigma\lambda}^{\mu(1)} \Gamma_{\rho\nu}^{\lambda(1)} \right) \\ &= \int d^4x \sqrt{-\bar{g}} \frac{1}{2} \left[(\bar{\nabla}^\lambda h) \left(\bar{\nabla}^\nu h_{\lambda\nu} - \frac{1}{2} \bar{\nabla}_\lambda h \right) \right. \\ &\quad \left. - (\bar{\nabla}^\nu h^{\mu\lambda}) \left(\bar{\nabla}_\mu h_{\lambda\nu} - \frac{1}{2} \bar{\nabla}_\nu h_{\mu\lambda} \right) \right]. \end{aligned}$$

Collecting the expressions together into [1.138] gives

$$\begin{aligned} S_{(2)} &= \frac{1}{2} \int d^4x \sqrt{-\bar{g}} \left\{ \left[(\bar{\nabla}^\lambda h) \left(\bar{\nabla}^\nu h_{\lambda\nu} - \frac{1}{2} \bar{\nabla}_\lambda h \right) \right. \right. \\ &\quad \left. \left. - (\bar{\nabla}^\nu h^{\mu\lambda}) \left(\bar{\nabla}_\mu h_{\lambda\nu} - \frac{1}{2} \bar{\nabla}_\nu h_{\mu\lambda} \right) \right] \right. \\ &\quad + [h \bar{\nabla}_\mu \bar{\nabla}_\nu h^{\mu\nu} - h \square h] \\ &\quad + [-2h^{\nu\sigma} \bar{\nabla}^\mu \bar{\nabla}_\sigma h_{\mu\nu} + h^{\mu\nu} \square h_{\mu\nu} + (\bar{\nabla}_\sigma \bar{\nabla}_\nu h) h^{\sigma\nu}] \\ &\quad \left. + \left[-\bar{R}^{\mu\nu} h h_{\mu\nu} + 2h^{\nu\alpha} h_\alpha^\sigma \bar{R}_{\nu\sigma} + \frac{\bar{R}}{4} (h^2 - 2h_{\mu\nu} h^{\mu\nu}) \right] \right\}. \end{aligned}$$

Finally, after integration by parts,

$$\begin{aligned} S_{(2)} &= \frac{1}{2} \int d^4x \sqrt{-\bar{g}} \left\{ -(\bar{\nabla}^\lambda h) (\bar{\nabla}^\nu h_{\lambda\nu}) + \frac{1}{2} (\bar{\nabla}_\lambda h) (\bar{\nabla}^\lambda h) \right. \\ &\quad \left. - \frac{1}{2} (\bar{\nabla}^\nu h^{\lambda\mu}) (\bar{\nabla}_\nu h_{\lambda\mu}) + (\bar{\nabla}^\nu h^{\lambda\mu}) (\bar{\nabla}_\mu h_{\lambda\nu}) \right. \\ &\quad \left. + \left[-\bar{R}^{\mu\nu} h h_{\mu\nu} + 2h^{\nu\alpha} h_\alpha^\sigma \bar{R}_{\nu\sigma} + \frac{\bar{R}}{4} (h^2 - 2h_{\mu\nu} h^{\mu\nu}) \right] \right\}. \end{aligned}$$

In flat space $\bar{g}_{\mu\nu} = \eta_{\mu\nu}$, the terms in square brackets all vanish and this reduces to the usual perturbed equations around Minkowski space.

In terms of the *trace reversed perturbation*

$$h_{\alpha\beta} = \bar{h}_{\alpha\beta} - \frac{1}{2} \bar{g}_{\alpha\beta} \bar{h} \quad [1.139]$$

this becomes

$$S_{(2)} = \frac{1}{2} \int d^4x \sqrt{-\bar{g}} \left\{ -\frac{1}{2} (\bar{\nabla}^\nu \bar{h}^{\lambda\mu}) (\bar{\nabla}_\nu \bar{h}_{\lambda\mu}) + \frac{1}{4} (\bar{\nabla}^\lambda \bar{h}) (\bar{\nabla}^\lambda \bar{h}) \right. \\ \left. + (\bar{\nabla}^\nu \bar{h}^{\lambda\mu}) (\bar{\nabla}_\mu \bar{h}_{\lambda\nu}) + \left[-\bar{R}^{\mu\nu} \bar{h} \bar{h}_{\mu\nu} + 2\bar{h}^{\nu\alpha} \bar{h}_\alpha^\sigma \bar{R}_{\nu\sigma} - \frac{\bar{R}}{2} \bar{h}_{\mu\nu} \bar{h}^{\mu\nu} + \frac{1}{4} \bar{R} \bar{h}^2 \right] \right\}$$

Linearised equations of motion

- *Equations of motion.* Variation of action [1.138] with respect to $h^{\alpha\beta}$ gives

$$0 = \int d^4x \sqrt{-\bar{g}} \delta h^{\alpha\beta} \left(-G_{\alpha\beta}^{(1)} - \frac{1}{2} h \bar{G}_{\alpha\beta} + h^{\mu\nu} (\bar{g}_{\nu\beta} \bar{G}_{\alpha\nu} + \bar{g}_{\nu\alpha} \bar{G}_{\beta\nu}) \right)$$

leading to the equations of motion

$$G_{\alpha\beta}^{(1)} = -\frac{1}{2} h \bar{G}_{\alpha\beta} + h^{\mu\nu} (\bar{g}_{\nu\beta} \bar{G}_{\alpha\nu} + \bar{g}_{\nu\alpha} \bar{G}_{\beta\nu}). \quad [1.140]$$

In terms of the trace reversed perturbation, after commuting covariant derivatives, for example,

$$\bar{\nabla}_\mu \bar{\nabla}_\beta \bar{h}_\alpha^\mu = \bar{\nabla}_\beta (\bar{\nabla}_\mu \bar{h}_\alpha^\mu) + \bar{R}_{\lambda\beta} \bar{h}_\alpha^\lambda - \bar{R}^\lambda_{\alpha\mu\beta} \bar{h}_\lambda^\mu \quad [1.141]$$

and then imposing the Lorenz gauge $\bar{\nabla}_\mu \bar{h}^{\mu\nu} = 0$ these read

$$\bar{\square} \bar{h}_{\mu\nu} + 2\bar{R}_{\mu\rho\nu\sigma} \bar{h}^{\rho\sigma} - 2\bar{G}_{\rho(\mu} \bar{h}_{\nu)}^\rho - \bar{g}_{\mu\nu} (\bar{R}_{\rho\sigma} \bar{h}^{\rho\sigma}) = 0 \quad [1.142]$$

which can be rewritten as

$$\bar{\square} \bar{h}_{\alpha\beta} + 2\bar{R}^\mu_{\alpha\nu\beta} \bar{h}_\mu^\nu + \bar{S}_{\mu\alpha\nu\beta} \bar{h}^{\mu\nu} = 0, \quad [1.143]$$

where

$$\bar{S}_{\mu\alpha\nu\beta} = [\bar{G}_{\mu\alpha} \bar{g}_{\beta\nu} + \bar{G}_{\mu\beta} \bar{g}_{\alpha\nu}] - \bar{R}_{\mu\nu} \bar{g}_{\alpha\beta}. \quad [1.144]$$

Including matter, and splitting its energy momentum tensor $T_{\mu\nu} = \bar{T}_{\mu\nu} + \delta T_{\mu\nu}$ into a background (determining $\bar{g}_{\mu\nu}$) plus first-order perturbation, the first-order field equations are

$$\bar{\square} \bar{h}_{\alpha\beta} + 2\bar{R}^\mu_{\alpha\nu\beta} \bar{h}_\mu^\nu + \bar{S}_{\mu\alpha\nu\beta} \bar{h}^{\mu\nu} = -\frac{16\pi G}{c^4} \delta T_{\mu\nu}. \quad [1.145]$$

1.5. Appendix: Green's functions

Given a linear differential operator Δ acting on functions on \mathbb{R}^n , its corresponding Green's function is a function on $\mathbb{R}^n \times \mathbb{R}^n$ satisfying

$$\Delta G(x, x') = \delta^{(n)}(x, x'). \quad [1.146]$$

Green's functions are useful to solve differential equations in the presence of sources, since they allow one to write the solutions of $\Delta\Phi = J$ as

$$\Phi(x) = \int d^n x' G(x, x') J(x') + \Phi^\circ(x), \quad [1.147]$$

where Φ° is a solution of the homogeneous problem with $J = 0$. For a given Δ the Green function is typically not unique, but a unique one can be selected via boundary conditions or other physical requirements.

For the Laplace equation, $n = 3$ and $\Delta = \partial^2$, there is a unique solution of [1.146] with vanishing boundary conditions at infinity, given by

$$G(\vec{x}, \vec{x}') = -\frac{1}{4\pi|\vec{x} - \vec{x}'|}. \quad [1.148]$$

Therefore

$$\Phi(\vec{x}) = \frac{1}{4\pi} \int d^3 x' \frac{J(\vec{x}')}{|\vec{x} - \vec{x}'|} + \Phi^\circ(\vec{x}). \quad [1.149]$$

For the d'Alembert equation, $n = 4$ and $\Delta = \square$, fixing vanishing boundary conditions at spatial infinity is not enough to have a single solution: there are infinitely many solutions of the homogeneous equation that can be added to any given G and still satisfy the defining equation. Two notable examples are the retarded and advanced ones, which are uniquely characterised by vanishing for x respectively in the past or the future of x' , and given by

$$\begin{aligned} G_\pm(x, x') &= -\frac{\delta(t - t' \mp |\vec{x} - \vec{x}'|)}{4\pi|\vec{x} - \vec{x}'|} \\ &= -\frac{1}{2\pi} \Theta(\pm(t - t')) \delta((t - t')^2 - |\vec{x} - \vec{x}'|^2), \end{aligned} \quad [1.150]$$

where Θ is Heaviside's step function. The retarded solution imposes no-incoming radiation boundary conditions, and it is the one relevant to study the emission of waves from a source.

1.6. Bibliography

- Aasi, J. *et al.* (2015), Advanced LIGO, *Class. Quant. Grav.*, 32, 074001.
- Abbott, B. P. *et al.* (2016), Observation of Gravitational Waves from a Binary Black Hole Merger, *Phys. Rev. Lett.*, 116(6), 061102.
- Acernese, F. *et al.* (2015), Advanced Virgo: a second-generation interferometric gravitational wave detector, *Class. Quant. Grav.*, 32(2), 024001.
- Agazie, G. *et al.* (2023), The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background, *Astrophys. J. Lett.*, 951(1), L8.
- Akutsu, T. *et al.* (2021), Overview of KAGRA: Detector design and construction history, *PTEP*, 2021(5), 05A101.
- Andersson, N. (2019), *Gravitational-Wave Astronomy*, Oxford Graduate Texts, Oxford University Press.
- Antoniadis, J. *et al.* (2023), The second data release from the European Pulsar Timing Array - III. Search for gravitational wave signals, *Astron. Astrophys.*, 678, A50.
- Ashtekar, A., Khera, N., Kolanowski, M., Lewandowski, J. (2022), Charges and fluxes on (perturbed) non-expanding horizons, *JHEP*, 02, 066.
- Ashtekar, A., Streubel, M. (1981), Symplectic Geometry of Radiative Modes and Conserved Quantities at Null Infinity, *Proc. Roy. Soc. Lond. A*, 376, 585–607.
- Barack, L., Pound, A. (2019), Self-force and radiation reaction in general relativity, *Rept. Prog. Phys.*, 82(1), 016904.
- Blanchet, L. (2006), Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries, *Living Rev. Rel.*, 9(1), 1–114.
- Bondi, H. (1960), Gravitational Waves in General Relativity, *Nature*, 186(4724), 535–535.
- Bondi, H., van der Burg, M. G. J., Metzner, A. W. K. (1962), Gravitational waves in general relativity. 7. Waves from axisymmetric isolated systems, *Proc. Roy. Soc. Lond. A*, 269, 21–52.
- Buonanno, A., Damour, T. (1999), Effective one-body approach to general relativistic two-body dynamics, *Phys. Rev. D*, 59, 084006.
- Caprini, C., Figueroa, D. G. (2018), Cosmological Backgrounds of Gravitational Waves, *Class. Quant. Grav.*, 35(16), 163001.
- Caprini, C., Pujolàs, O., Quelquejay-Leclere, H., Rompineve, F., Steer, D. A. (2025), Primordial gravitational wave backgrounds from phase transitions with next generation ground based detectors, *Class. Quant. Grav.*, 42(4), 045015.

- Chandrasekaran, V., Flanagan, E. E., Prabhu, K. (2018), Symmetries and charges of general relativity at null boundaries, *JHEP*, 11, 125.
- Choquet-Bruhat, Y. (1952), Théorème d'existence pour certains systèmes d'équations aux dérivées partielles non linéaires, *Acta Mathematica*, 88, 141–225.
- Colpi, M. *et al.* (2024), LISA Definition Study Report, .
- Deruelle, N., Uzan, J.-P. (2018), *Relativity in Modern Physics*, Oxford University Press.
- Dray, T., Streubel, M. (1984), Angular momentum at null infinity, *Class. Quant. Grav.*, 1(1), 15–26.
- Flanagan, E. E., Hughes, S. A. (2005), The Basics of gravitational wave theory, *New J. Phys.*, 7, 204.
- Goldberger, W. D., Rothstein, I. Z. (2006), An Effective field theory of gravity for extended objects, *Phys. Rev. D*, 73, 104029.
- Holz, D. E., Hughes, S. A. (2005), Using gravitational-wave standard sirens, *Astrophys. J.*, 629, 15–22.
- Iacovelli, F., Mancarella, M., Foffa, S., Maggiore, M. (2022), Forecasting the Detection Capabilities of Third-generation Gravitational-wave Detectors Using GWFIRST, *Astrophys. J.*, 941(2), 208.
- Iyer, V., Wald, R. M. (1994), Some properties of Noether charge and a proposal for dynamical black hole entropy, *Phys. Rev. D*, 50, 846–864.
- Kennefick, D. (1997), Controversies in the history of the radiation reaction problem in general relativity, *CalTech Library*, .
- Kokkotas, K. D., Schmidt, B. G. (1999), Quasinormal modes of stars and black holes, *Living Rev. Rel.*, 2, 2.
- LIGO (n.d.), '<https://gracedb.ligo.org>'.
- Maggiore, M. (2007), *Gravitational Waves. Vol. 1: Theory and Experiments*, Oxford University Press.
- Maggiore, M. (2018), *Gravitational Waves. Vol. 2: Astrophysics and Cosmology*, Oxford University Press.
- Maggiore, M., Iacovelli, F., Belgacem, E., Mancarella, M., Muttoni, N. (2024), Comparison of global networks of third-generation gravitational-wave detectors, *to appear*, .
- Newman, E. T., Penrose, R. (1968), New conservation laws for zero rest-mass fields in asymptotically flat space-time, *Proc. Roy. Soc. Lond. A*, 305, 175–204.
- Poisson, E., Will, C. M. (2014), *Gravity: Newtonian, Post-Newtonian, Relativistic*, Cambridge University Press.

- Reardon, D. J. *et al.* (2023), Search for an Isotropic Gravitational-wave Background with the Parkes Pulsar Timing Array, *Astrophys. J. Lett.*, 951(1), L6.
- Reitze, D. *et al.* (2019), Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO, *Bull. Am. Astron. Soc.*, 51(7), 035.
- Sachs, R. K. (1962), Gravitational waves in general relativity. 8. Waves in asymptotically flat space-times, *Proc. Roy. Soc. Lond. A*, 270, 103–126.
- Samajdar, A., Janquart, J., Van Den Broeck, C., Dietrich, T. (2021), Biases in parameter estimation from overlapping gravitational-wave signals in the third-generation detector era, *Phys. Rev. D*, 104(4), 044003.
- Sathyaprakash, B. *et al.* (2012), Scientific Objectives of Einstein Telescope, *Class. Quant. Grav.*, 29, 124013. [Erratum: *Class. Quant. Grav.* 30, 079501 (2013)].
- Schutz, B. F. (1986), Determining the Hubble Constant from Gravitational Wave Observations, *Nature*, 323, 310–311.
- Thorne, K. S. (1980), Multipole Expansions of Gravitational Radiation, *Rev. Mod. Phys.*, 52, 299–339.
- Wald, R. M. (1984), *General relativity*, Chicago Univ. Press, Chicago, IL.
- Wald, R. M., Zoupas, A. (2000), A General definition of ‘conserved quantities’ in general relativity and other theories of gravity, *Phys. Rev. D*, 61, 084027.
- Weber, J. (1960), Detection and generation of gravitational waves, *Physical Review*, 117(1), 306–313.
- Weinberg, S. (2005), *The Quantum theory of fields. Vol. 1: Foundations*, Cambridge University Press.
- Will, C. M. (2014), The Confrontation between General Relativity and Experiment, *Living Rev. Rel.*, 17, 4.
- Xu, H. *et al.* (2023), Searching for the Nano-Hertz Stochastic Gravitational Wave Background with the Chinese Pulsar Timing Array Data Release I, *Res. Astron. Astrophys.*, 23(7), 075024.
- Yunes, N., Siemens, X. (2013), Gravitational-Wave Tests of General Relativity with Ground-Based Detectors and Pulsar Timing-Arrays, *Living Rev. Rel.*, 16, 9.

DRAFT

2

The theoretical foundations of gravitational waves: detection and emission

Simone Speziale¹ and Danièle Steer²

¹*Aix Marseille Univ., Univ. de Toulon, CNRS, CPT, UMR 7332, 13288
Marseille, France*

²*Laboratoire de Physique de l'ENS, Université Paris Cité, Ecole Normale
Supérieure, Université PSL, Sorbonne Université, CNRS, 75005 Paris*

The aim of this second chapter is to provide details about the detection and emission of gravitational waves from sources, and derive the highlight formulas that were used in the initial overview section of the first Chapter.

2.1. Detection of GWs

2.1.1. Coordinate displacements versus physical displacements

The simplest way to study the effect of a gravitational wave, is to compute how it changes the physical distance between free-falling test masses. These follow time-like geodesics, whose tangent vector field u^μ satisfies the geodesic equation

$$u^\nu \text{New} A_\nu u^\mu = \frac{du^\mu}{d\tau} + \Gamma_{\nu\rho}^\mu u^\nu u^\rho = 0, \quad u^\mu \partial_\mu = \frac{d}{d\tau}. \quad [2.1]$$

Here τ is the proper time τ , and $u^2 = -c^2$. If the masses are initially at rest, we have $u^a = 0$ and

$$\frac{du^\mu}{d\tau} = -\Gamma_{00}^\mu c^2 = \left(\frac{1}{2} \partial^\mu h_{00} - \partial_0 h_0^\mu \right) c^2 + O(h^2). \quad [2.2]$$

For vacuum solutions in the TT gauge $h_{0\mu} = 0$. Hence \dot{u}^μ vanishes at lowest order, and the coordinate distance (as well as the coordinate time delay) between two nearby time-like geodesics remains the same during the passage of the wave. This result provides us with an interpretation of the TT gauge: it is a choice of coordinates which are labelled by the position of free-falling test masses. It is thus the linearized version of the temporal gauge in the full theory discussed earlier. Now, even though the *coordinate distance* between two test masses remains the same in this gauge, their *physical distance* does not. It is given by

$$L = \int_0^{L_0} d\lambda \sqrt{g_{ab} \hat{e}^a \hat{e}^b} = \int_0^{L_0} d\lambda \left(1 + \frac{1}{2} h_{ab}^{\text{TT}} \hat{e}^a \hat{e}^b \right) + O(h^2). \quad [2.3]$$

Here \hat{e}^a is the tangent to the curve connecting the test masses, and λ an arbitrary parametrization thereof. Since we are working in perturbation theory, we can consider as curve the background geodesic connecting the test masses, and since we are working in Cartesian coordinates of the background, we can choose \hat{e}^a to be constant along a coordinate axis, and take λ as one of the coordinates. Then L_0 is the coordinate distance. It coincides with the physical

distance in the background flat metric, but differs from it when the spacetime is perturbed. If we further assume that the wavelengths of the wave are much bigger than L_0 , we can ignore the space dependence of h_{ab}^{TT} and write the result as

$$L \simeq \left(1 + \frac{1}{2} h_{ab}^{\text{TT}} \hat{e}^a \hat{e}^b \right) L_0. \quad [2.4]$$

The approximation becomes of course exact if the direction of propagation of the wave is orthogonal to the axis connecting the masses.

The discussion offers an example of one of the most important lessons of general relativity, namely the importance of distinguishing coordinate effects from physical results. General covariance guarantees that all calculations can be performed in any coordinate system. But one has to always make sure that the physical consequences derived from the calculations are coordinate independent. In this example, we found that the coordinate position of test particles are unaffected by the passage of the wave. This is true only in the TT gauge. The relevant coordinate-independent quantity is the physical distance, and we found that it changes. Furthermore, it would change in the same way in any coordinate system preserving the endpoints of the integral: the geodesic distance between two physical points is a gauge-invariant observable.

As we have already discussed, even though any coordinate system can be chosen, choosing a good coordinate system is important to simplify calculations. Even though TT coordinates hide the passage of the wave, they are convenient because we can write the physical distance using fixed extrema in the integral at all times.¹ Another convenient gauge is the one of Fermi normal coordinates, which works in the opposite way: the metric is trivial at lowest order near two geodesics, see [1.32], hence the physical distance is entirely captured by the coordinate distance. To see this, let us consider the geodesic deviation equation. It is given by

$$u^\rho \text{New} A_\rho (u^\nu \text{New} A_\nu \xi^\mu) = R^\mu{}_{\nu\rho\sigma} u^\nu u^\rho \xi^\sigma, \quad [2.5]$$

1. The same effect is used when choosing the synchronous gauge in cosmology, it is often convenient to choose a coordinate system such that the values of the coordinate grid represent galaxies, so that their coordinate distance does not change, while the physical one does.

where ξ^μ is a vector connecting neighbouring geodesics, chosen such that $\xi \cdot u = [\xi, u] = 0$. At first order in h , and with the assumption of vanishing initial velocities, it reduces to

$$\frac{d^2 \xi^a}{dt^2} = -2\Gamma_{0\nu}^a u^0 \dot{\xi}^\nu - c^2 \xi^\nu \partial_\nu \Gamma_{00}^a. \quad [2.6]$$

In the TT gauge, the last term vanishes and the first one too if the initial velocity was zero. The coordinate distance between the geodesics stays constant, in agreement with the result already derived using the geodesic equation above.

The description changes completely if we use a gauge corresponding to a local inertial frame, such as the Fermi normal coordinates, which can be used to set to zero the Christoffel symbols all along a chosen geodesic. Doing so, the LHS of [2.5] becomes simply a second partial derivative, and the equation reads

$$\frac{d^2 \xi^a}{dt^2} = c^2 R^a{}_{00b} \xi^b = \frac{1}{2} \ddot{h}_{\text{TT}}^{ab} \xi_b + O(h^2), \quad [2.7]$$

where in the second equality we have neglected any contribution from the potentials. This shows that in a local inertial frame, the deviation between nearby geodesics can be described directly in terms of their coordinate separation. This is consistent with what previously seen in [1.32], the metric only changes at quadratic order in the coordinate distance from the origin, hence coordinate distances coincide with physical distance at first order. The geodesic deviation equation [2.7] shows that an apparatus that can detect tidal effects will be sourced only by the physical components of the GW, and not by the gauge ones. The equation can be easily solved at first order in h , with

$$\xi^a(t) = \xi^a(0) + \frac{1}{2} h_{\text{TT}}^{ab} \xi_b(0). \quad [2.8]$$

Since in this gauge the coordinate distance coincides with the physical distance at first order, we recover the gauge-invariant result [2.4] but where this time the metric is unchanged, and it is the extremum of the integral that has moved.

The change of proper distance [2.4] also shows the meaning of the wave's polarizations. Using the example [1.86] of a monochromatic wave propagating along the z axis, and setting the test masses at $z = 0$, we can write the relative change in physical distance as

$$\frac{\delta L}{L_0} := \frac{L - L_0}{L_0} \simeq \left(\frac{1}{2} h_+ (\hat{e}^x \hat{e}^x - \hat{e}^y \hat{e}^y) + h_\times \hat{e}^x \hat{e}^y \right) \cos \omega t. \quad [2.9]$$

An h_+ polarization would cause pairs of masses along the x and y axis to periodically approach and recede, hence drawing a $+$ -like pulse in time, see Figure 2.1. An h_\times polarization would cause the same effect but along the axis $\hat{e}^a = (1, 1, 0)/\sqrt{2}$, namely rotated by 45 degrees. This type of deformation is also called ‘shear’ of the congruence of time-like geodesics followed by the test masses. It can be visualized even more clearly if we use [2.8]. We consider a circular distribution of test masses centered around the origin in the plane perpendicular to the direction of propagation of the wave, see Fig. 2.1. Then we can identify the displacement vector with the coordinate vector of each mass (labelled by i), and the effect of a monochromatic wave of frequency ω is

$$\begin{aligned}x_i(t) &= x_i(0) + h_+(t)x_i(0) + h_\times(t)y_i(0), \\y_i(t) &= y_i(0) - h_+(t)y_i(0) + h_\times(t)x_i(0).\end{aligned}$$

The effect is shown in Fig. 2.1, where the period $T = 2\pi/\omega$. The spin-2, quadrupolar nature of the gravitational force is evident from the shape of the shear deformation, and can be compared with the tidal deformation of earth’s oceans.




Figure 2.1: The effect of the two polarizations on a circular distribution of test masses. In the upper panel (red circles) $h_+ \neq 0$ and $h_\times = 0$. The lower panel has $h_\times \neq 0$ and $h_+ = 0$.

If external forces are present, on top of the gravitational one, then [2.1] and [2.5] acquire additional terms on the right hand side. So in particular [2.7] now reads

$$\frac{d^2 \xi^a}{dt^2} = c^2 R^a{}_{00b} \xi^b + \frac{F^a}{m}. \quad [2.10]$$

as an example, consider a material bar. While the effect of the gravitational wave is to stretch spacetime changing the physical distance between the molecules of the bar, there are also electromagnetic forces that hold the bar together, and which are intrinsically much stronger. For instance, the Coulomb interaction between two electrons one angstrom apart is $\sim 10^{42}$ times stronger than its Newtonian counterpart. For this reason, one could in principle use simply a rigid ruler measuring the distance between two freely falling masses to detect gravitational waves. The problem with this idea is the weakness of the waves, and more sophisticated experiments are required to actually observe the waves.

A special case of external forces appearing on the right-hand side of [2.10] are the inertial forces. These, by the equivalence principle, can be reabsorbed in a coordinate transformation of the metric. For instance if the frame has both an acceleration \vec{a} and an angular velocity $\vec{\Omega}$ with respect to a local inertial frame, then (Ni and Zimmermann 1978)

$$\begin{aligned} ds^2 = & -c^2 dt^2 \left(\left(1 + \frac{1}{c^2} \vec{a} \cdot \vec{x}\right)^2 - \frac{1}{c^2} (\vec{\Omega} \times \vec{x})^2 + R_{0c0d} x^c x^d \right) \\ & + 2cdtdx^a \left(\frac{1}{c} \epsilon_{abc} \Omega^b x^c - \frac{2}{3} R_{0cad} x^c x^d \right) + dx^a dx^b \left(\delta_{ab} - \frac{1}{3} R_{acbd} x^c x^d \right) \\ & + O(x^2), \end{aligned} \quad [2.11]$$

and [2.10] becomes

$$\frac{d^2 \vec{\xi}}{dt^2} = -\vec{a} - 2\vec{\Omega} \times \vec{v} + \frac{\vec{F}}{m} + O(x^2). \quad [2.12]$$

All the gravitational effects as well as further non-inertial effects such as centrifugal acceleration are $O(x^2)$. So in order to be capable of detecting gravitational waves, a detector must first of all be freed from all the external forces that would otherwise drown the signal in noise.

In realistic physical systems, the emission will not be a plane wave, but rather a wave packet with finite temporal extension. The effect on the circular

distribution will then be a superposition of different frequencies and different helicities, each with their own (time-dependent) amplitude. The temporal finiteness of the signal can also lead to a new type of effect: after the wave has passed, the distribution will stop oscillating, but its shape will in general not be the same as before the wave's arrival. This effect is called displacement memory, and we will see below in Sec.2.4.6 an explicit example. The effect carries the memory of the wave, since it permits in principle to detect the passage of a gravitational wave even after the event. In practise though the detection is very difficult, because the external forces that make up the matter distribution will act and bring it back to its rest configuration. It is nonetheless one of the targets of future detectors (Grant and Nichols 2023).

2.1.2. Interferometers

Let us briefly describe how the formulas above are used in the most common type of detectors, laser interferometers. Other chapters in this collection will cover more details as well as the types of detectors. The basic idea of a laser interferometer is to detect physical changes like [2.3] from the time-of-flight of monochromatic light signals. This can be done easily in the linear theory if we make the additional approximations that gravitational potentials can be neglected,² and that the wavelength of the signal is much longer than the arms of the interferometer. The first approximation guarantees that the only source of curvature comes from the wave, hence the Riemann tensor scales like λ^{-2} . We can then set up a free falling frame say in Fermi normal coordinates centered on the beam splitter's geodesic. Thanks to the second approximation, the spatial projection of the null geodesics follows straight lines, hence the time of flight is directly related to the physical distance along the interferometer's arms. The latter is given by [2.4] regardless of the direction of the wave, thanks again to the assumption that the wavelength is much larger than the arms' length. Denoting $\hat{e}_{1,2}^a$ the two axis, we have

$$L_2 - L_1 = \frac{L_0}{2} h_{ab}^{\text{TT}} (\hat{e}_1^a \hat{e}_1^b - \hat{e}_2^a \hat{e}_2^b). \quad [2.13]$$

For a typical signal $h \sim 10^{-21}$ (see overview Section 1.1.4), hence the difference in arrival time would be $\Delta T = (L_2 - L_1)/c \sim 10^{-26} s$ which is way too small to be measurable. Two ingenious ideas come to the rescue. The first is Michelson-Morley's idea to measure not time but phase interferences, and the second it

2. The potentials generated by the source can be naturally neglected because they fall off faster than the radiative modes, so this approximation concerns mostly the potentials of the local gravitational field of the observer.

to increase the effective path of light through Fabry-Perot cavities. If we set the lasers so that the phases at the beam splitter are identical, the phase shift after the travel to and back from the mirrors will be

$$\Delta\phi = \frac{2\pi\nu}{c} N_p (2L_1 - 2L_2), \quad [2.14]$$

where $N_p = 1$ for a Michelson device and up to 300 for the Fabry-Pérot one used by LIGO and Virgo, and the laser frequency ν can go as high as $10\text{kHz} = 10^4\text{s}^{-1}$, thus improving the strength of a typical signal to 10^{-20} , which is actually – and remarkably – an observable phase difference.




Figure 2.2: World lines of photon trajectories between the beam splitter and the end mirrors in an interferometer with arms of equal length. Blue dashed lines: with no gravitational wave. Red lines: perturbed trajectories. In the long wavelength approximation $\lambda \gg L_0$ the red lines are straight (but still have different angles than the unperturbed blue lines), and $\Delta T = 2\Delta L/c$.

Combining the last two formulas we have

$$\Delta\phi = \frac{2\pi\nu}{c} N_p L_0 h_{ab}^{\text{TT}} (\hat{e}_1^a \hat{e}_1^b - \hat{e}_2^a \hat{e}_2^b), \quad [2.15]$$

where now \hat{e}^a are the unit vectors giving the direction of each arm. It is possible to express the result in terms of the two wave polarizations by replacing h_{ab}^{TT} with [1.93]. To compute the scalar products between the vectors corresponding to the detector's arms and the vectors defining the polarizations in the plane orthogonal to the propagation, we introduce a rotation from the detector's frame to the frame of propagation, plus a reflection to take into account the fact that the axis of propagation is opposite to the direction of the source acting also on y to keep right-handed orientation of the frame. If we denote (θ, φ) the angles identifying the direction of propagation, the result is

$$\Delta\phi = \frac{4\pi\nu}{c} N_p L_0 (F_+ h_+ + F_\times h_\times), \quad [2.16]$$

where the coefficients

$$F_+ = \frac{1}{2}(1 + \cos^2 \theta) \cos 2\varphi, \quad F_\times = \cos \theta \sin 2\varphi \quad [2.17]$$

are called detector's pattern functions.³ This shows that while a single two-armed interferometer is sensitive to both polarizations, it cannot distinguish them. It also shows that the sensitivity depends on the relative orientation with respect to the sources. The dependence is very strong, to the point that there are directions in which the detector is completely blind, like $(\theta, \varphi) = (\frac{\pi}{2}, \frac{\pi}{4})$. Hence the importance of multiple detectors in order to increase sensitivity in every direction and the possibility of distinguishing the polarizations. Multiple detectors also allow studying the localization of the source via triangulation.

If the approximation $\lambda \gg L_0$ is no longer valid, then one has to take into account the redshift changes during the time of flight, see for instance discussion in (Cornish and Rubbo 2003 ; Andersson 2019).

2.2. Generation of GWs from sources

2.2.1. Introducing sources

Following the principle of general covariance, the matter Lagrangian should satisfy the property [1.24], namely be written solely in terms of the dynamical matter fields and spacetime metric, and no additional background fields. The

3. This formula assumes that the arms are perpendicular, and a fixed polarization basis. It can be generalized to include an additional rotation of the two polarization basis, as well as a non-perpendicular angle between the arms.

simplest way to obtain a viable matter Lagrangian is then to start from the one used in the absence of gravity, and ‘covariantize’ it by the replacements

$$\eta_{\mu\nu} \rightarrow g_{\mu\nu}, \quad \partial_\mu \rightarrow \text{NewA}_\mu, \quad d^4x \rightarrow \sqrt{-g}d^4x. \quad [2.18]$$

Doing so introduces a minimal coupling of matter to the gravitational field. Additional interactions can be included if phenomenologically or theoretically motivated, provided they respect the principle of general covariance, embodied for instance by condition [1.24]. Having done so, we define the matter energy-momentum tensor

$$T_{\mu\nu} = -\frac{2c}{\sqrt{-g}} \frac{\delta \mathcal{L}_M}{\delta g^{\mu\nu}}. \quad [2.19]$$

Inserting this definition in [1.24], and using [1.21], we obtain

$$\xi_\nu \text{NewA}_\mu T^{\mu\nu} = \frac{c}{\sqrt{-g}} \left[\frac{\delta \mathcal{L}_M}{\delta \psi} \mathcal{L}_\xi \psi + \partial_\mu (\tilde{\theta}_M^\mu - \xi^\mu \mathcal{L}_M + \frac{\sqrt{-g}}{c} T^{\mu\nu} \xi_\nu) \right]. \quad [2.20]$$

The first term on the right-hand side is proportional to the matter’s equations of motion, and vanishes on-shell. The rest is a total derivative and vanishes in the absence of boundaries or with appropriate boundary conditions. Since the identity holds for any ξ , we conclude that on-shell,

$$\text{NewA}_\mu T^{\mu\nu} \doteq 0. \quad [2.21]$$

This equation replaces the familiar conservation of the energy-momentum tensor guaranteed by Noether’s theorem in flat spacetime. More precisely, the Noether current of the total Lagrangian $\mathcal{L}_{\text{EH}} + \mathcal{L}_M$ is

$$j_\xi^\mu = \frac{c^3}{8\pi G} \left(E^\mu{}_\nu \xi^\nu - \text{NewA}_\nu \text{NewA}^{[\mu} \xi^{\nu]} \right), \quad [2.22]$$

where E are Einstein’s equations [1.16], namely [1.26] with the vacuum equations replaced by the equations in the presence of matter), and whose conservation requires to be on-shell of both the Einstein’s and matter’s field equations:

$$\text{NewA}_\mu j_\xi^\mu = \frac{c^3}{8\pi G} E^{\mu\nu} \text{NewA}_\mu \xi_\nu - \frac{1}{c} \text{NewA}_\mu T^{\mu\nu} \xi_\nu \doteq 0. \quad [2.23]$$

Even though [2.21] is often referred to as the general covariant version of energy-momentum conservation, it is important to remark that it is *not* a

conservation equation in the usual sense. To understand this point, let us follow the usual procedure to obtain Noether charges from the current, and apply Stokes's theorem to a finite region M with boundary ∂M . To do so we need a scalar, which we obtain by contracting the left-hand side of [2.21] with a vector ξ^μ . After integrating by parts, we find

$$\int_M \text{NewA}_\mu T^{\mu\nu} \xi_\nu \sqrt{-g} d^4x = \oint_{\partial M} T^{\mu\nu} \xi_\nu n_\mu \sqrt{q} d^3y + \int_M T^{\mu\nu} \text{NewA}_\mu \xi_\nu \sqrt{-g} d^4x, [2.24]$$

for a boundary ∂M with induced metric q , normal n_μ coordinates y . If $\text{NewA}_{(\mu} \xi_{\nu)}$ vanishes, namely if the Killing equation is satisfied, then [2.21] can be turned into a conservation law. To do so, we consider the case in which ∂M consists of two space-like hypersurfaces Σ_1 and Σ_2 connected by a time-like boundary \mathcal{T} , see Fig.2.3. If the fields satisfy conservative boundary conditions at \mathcal{T} (typically \mathcal{T} asymptotically far away and fall-off conditions on the fields), then

$$Q_\xi := \int_{\Sigma_1} T^{\mu\nu} \xi_\nu n_\mu \sqrt{q} d^3y = \int_{\Sigma_2} T^{\mu\nu} \xi_\nu n_\mu \sqrt{q} d^3y [2.25]$$

for each Killing vector ξ . Therefore [2.21] gives as many conserved quantities as there are isometries in spacetime. For flat spacetime, these are the ten Poincaré charges. For a generic dynamical spacetime, there are none.




Figure 2.3: A region of spacetime bounded by two space-like hypersurfaces $\Sigma_{1,2}$ and a time-like one \mathcal{T} . With conservative boundary conditions on \mathcal{T} , [2.25] establishes as many conservation laws as there are Killing vectors.

The validity of [2.21] implies the matter equations of motion, as we have seen from its derivation. In particular, if matter consists of test particles, namely free motion without self-interaction and ignoring the back-reaction on the metric, this equation implies the geodesics equation in curved spacetime. This is for

instance how one can derive the relativistic corrections to the Kepler problem, by evaluating [2.21] on the Schwarzschild background. At lowest order in the weak-field expansion [1.48], [2.21] reduces to the energy-momentum conservation law in flat spacetime,

$$\partial_\mu T^{\mu\nu} = 0. \quad [2.26]$$

This means that at lowest order the matter can interact with itself, but not with the gravitational field: the sources follow geodesics in flat spacetime (that is, straight lines). To include the effect of gravity on the sources we must go beyond the lowest order. In other words, the linearized theory still describes gravity in the Newtonian way, namely as a force acting in flat spacetime. Of course, it already contains departures from Newton's theory, since it includes the special relativistic effects such as the gravito-magnetic interaction and radiation.

2.2.2. Source multipoles

Let us study the conserved quantities that arise on the Minkowski background. We choose Σ to be a global hypersurface of constant time t , write its unit normal as $n_\mu = -\partial_\mu t$, and ξ is one of the ten Poincaré Killing vectors [1.50]. We can then use [2.25] to identify ten conserved quantities. Four are the energy and momentum

$$c^2 M := \int d^3x T^{00}, \quad cP^a := \int d^3x T^{0a}, \quad [2.27]$$

corresponding to $a^\mu{}_\nu = 0$ and unit values of b^μ . The remaining six are the relativistic angular momentum

$$\begin{aligned} cL^a &= \frac{c}{2} \epsilon^a{}_{bc} L^{bc} := \epsilon^a{}_{bc} \int d^3x x^b T^{0c}, \\ c^2 K^a &:= \int d^3x (T^{0a} ct - T^{00} x^a), \end{aligned} \quad [2.28]$$

corresponding to $b^\mu = 0$ and unit values of $a^a{}_\nu$ and $a^0{}_\nu$ respectively. Here ϵ_{abc} is the completely anti-symmetric Levi-Civita symbol in flat spacetime. Their conservation can be easily checked. We start by separating [2.26] in time and space components,

$$c^{-1} \dot{T}^{00} + \partial_a T^{a0} = 0, \quad c^{-1} \dot{T}^{0a} + \partial_b T^{ab} = 0. \quad [2.29]$$

Then using Stokes' theorem and vanishing boundary conditions we immediately see that

$$\dot{M} = \dot{P}^a = \dot{L}^a = \dot{K}^a = 0. \quad [2.30]$$

The first of [2.27] is the total energy, but we followed the custom in the literature to denote it M and refer to it as 'mass', using Newtonian language. The first of [2.28] is the angular momentum with respect to the frame defined by $n_\mu = -\partial_\mu t$. The second conserved quantity is the 'boost charge', and can be rewritten as $K^a = I^a - tP^a$, where we introduce the center-of-mass position

$$I^a = \frac{1}{c^2} \int d^3x T^{00} x^a. \quad [2.31]$$

Conservation of K^a is thus the statement that the center of mass moves following the total momentum.

The conserved quantities can be used to fix a reference frame as follows. First, we can choose the rest frame, where $P^a = 0$. This removes the freedom of Lorentz boosts. Then, we can fix the origin to be in the center-of-mass, where $I^a = 0$. This removes the freedom of spatial translations. The rotation freedom can be fixed choosing the axis so that L^a has only one component (say z), and the remaining $\text{SO}(2)$ freedom is fixed choosing an axis in the plane perpendicular to L^a . Finally the time translation symmetry is fixed setting the zero value of the clock.

The quantity I^a is also called mass-dipole moment. The terminology comes about if we see $\rho = c^{-2}T_{00}$ as a distribution, then I^a is the first moment of that distribution. Following this logic, we introduce a multi-index notation for the higher multipole moments:

$$\begin{aligned} I^{ab\dots} &= \frac{1}{c^2} \int d^3x T^{00} x^a x^b \dots, & P^{a,b\dots} &= \frac{1}{c} \int d^3x T^{0a} x^b \dots, \\ S^{ab,c\dots} &= \int d^3x T^{ab} x^c \dots \end{aligned} \quad [2.32]$$

The conservation laws [2.29] together with integration by parts in the absence of boundary terms provide relations between multipole moments and time variations of higher multipoles, such as

$$\begin{aligned} P^a &= -\dot{I}^a, & S^{ab} &= \frac{1}{2} \ddot{I}^{ab}, \\ \dot{S}^{ab,c} &= \frac{1}{6} \dddot{I}^{abc} + \frac{1}{3} (\ddot{P}^{a,bc} + P^{b,ac} - P^{c,ab}), & \dot{P}^{a,b} &= S^{ab} \end{aligned} \quad [2.33]$$

and so on. The first one above is the conservation of K^a already seen, and relates the momentum monopole to the mass dipole time variation. The second one allows one to determine the total effect of the stresses in the matter in terms of the second time derivative of the mass quadrupole. These relations are useful because it is typically easier to measure and interpret the multipole moments of the mass and momentum distributions, rather than the spatial stresses.

When working with multipoles, it is typically convenient to organize them into irreducible representations of the rotation group, which are label by a an integer number l and have $2l + 1$ components each, as recalled earlier. This can be achieved expanding the distribution in spherical harmonics, e.g. $\rho = \sum_{l,m} \rho_{l,m} Y_{l,m}$, then the integrals of the modes $\rho_{l,m}$ are the irreducible multipoles. It is possible although more cumbersome to do this composition directly in Cartesian coordinates without introducing spherical harmonics. One then gets

$$M = \frac{1}{c^2} \int d^3x \rho, \quad D^a = I^a = \frac{1}{c^2} \int d^3x \rho x^a, \quad [2.34]$$

$$Q^{ab} = \frac{1}{c^2} \int d^3x \rho (x^a x^b - \frac{r^2}{3} \delta^{ab}), \quad O^{abc} = \frac{1}{c^2} \int d^3x \rho (15x^a x^b x^c - 9x^{(a} \delta^{bc)} r^2), \quad [2.35]$$

and so on. Notice that while the multipole moments are useful at all orders in perturbation theory, the conservation laws [2.30] are only valid at lowest order.

2.2.3. Solving the wave equation with sources

We are interested in the emission of gravitational waves from matter sources, without incoming radiation. This can be imposed choosing the retarded Green function and setting to zero the independent degrees of freedom $h_{\mu\nu}^{\text{TT}}$. The general solution is then

$$\begin{aligned} \bar{h}_{\mu\nu} &= -\frac{16\pi G}{c^4} \int d^4x' G(x, x') T_{\mu\nu}(x') \\ &= \frac{4G}{c^4} \int d^3x' \frac{T_{\mu\nu}(t - \frac{1}{c}|\vec{x} - \vec{x}'|, \vec{x}')}{|\vec{x} - \vec{x}'|}. \end{aligned} \quad [2.36]$$

using [1.150] and the specialized De Donder gauge with no homogeneous solution. Even in the linearized approximation, the integral is in general very complicated and there is no analytic solution. So we resort to approximation schemes. In particular, we introduce two independent approximations:

- (i) Wave-zone approximation: we assume to be very far away from the sources, that is $R := |\vec{x}| \gg |\vec{x}'|$. This allows us to expand the integrand in powers of $1/R \ll 1$. For the numerator, we have

$$|\vec{x} - \vec{x}'| = R - \vec{N} \cdot \vec{x}' + \dots \quad [2.37]$$

where $\vec{N} := \vec{x}/R$, and

$$T_{\mu\nu}(t - \frac{1}{c}|\vec{x} - \vec{x}'|, \vec{x}') \simeq T_{\mu\nu}(t_{\text{R}}, \vec{x}') + \frac{\vec{N} \cdot \vec{x}'}{c} \dot{T}_{\mu\nu}(t_{\text{R}}, \vec{x}') + \dots, [2.38]$$

where we introduce the *retarded time*⁴

$$t_{\text{R}} := t - \frac{R}{c}. \quad [2.39]$$

For the denominator, we have

$$\frac{1}{|\vec{x} - \vec{x}'|} = \frac{1}{R} + \frac{\vec{N} \cdot \vec{x}'}{R^2} + \frac{3}{2}(x'_a x'_b - \frac{r'^2}{3} \delta_{ab}) \frac{N^a N^b}{R^3} + \dots \quad [2.40]$$

Furthermore, the direction of propagation of the wave coincides with the direction from the source, namely $-\vec{N}$ if we take the origin of the coordinates inside the source. Hence the TT projector can be written in terms of \vec{N} instead of the wave vector.

- (ii) Slow dynamics: We assume that the dynamics of the source is slow, so that time derivatives in [2.38] are small corrections. To understand why, consider that the integration coordinate \vec{x}' spans at most the size of the source, and if this has a typical frequency scale ω_s (for instance in a binary, the frequency of the orbit), then $v_s := |\vec{x}'|\omega_s$ is the velocity scale of the source. It follows that

$$\frac{|\vec{x}'|}{c} \dot{T}_{\mu\nu} \sim \frac{|\vec{x}'|\omega_s}{c} T_{\mu\nu} \sim \frac{v_s}{c} T_{\mu\nu} \quad [2.41]$$

is suppressed by v/c . The Taylor expansion [2.38] is therefore controlled by the parameter $v/c \ll 1$, and it is called post-Newtonian expansion.

⁴ Namely the time at which a signal travelling at the speed of light was sent in order to arrive at t .

The approximated solution can thus be written as

$$\bar{h}_{\mu\nu}(x) = \frac{4G}{c^4 R} \int d^3x' \left(T_{\mu\nu}(t_R, \vec{x}') + \frac{N_a}{c} \dot{T}_{\mu\nu}(t_R, \vec{x}') x'^a + \frac{N_a}{R} T_{\mu\nu}(t_R, \vec{x}') x'^a + \dots \right) \quad [2.42]$$

The first term is the leading order; the second term is the first of the PN corrections; the third term is the first of the $1/R$ corrections. Using the multipole definitions [2.32] and their conservation laws [2.33], we can rewrite the different components of the solution [2.42] as

$$\bar{h}_{00} = \frac{4G}{c^2 R} \left(M - \frac{N^a}{c} P_a + \frac{N^a N^b}{2c^2} \ddot{I}_{ab} + \frac{N^a}{R} I_a + \dots \right) \Big|_{t_R}, \quad [2.43a]$$

$$\bar{h}_{0a} = -\frac{4G}{c^3 R} \left(P_a + \frac{N^b}{2c} \dot{I}_{ab} + \frac{N^b}{2Rc} (L_{ab} + \dot{I}_{ab}) + \dots \right) \Big|_{t_R}, \quad [2.43b]$$

$$\bar{h}_{ab} = \frac{4G}{c^4 R} \left(\frac{1}{2} \ddot{I}_{ab} + \frac{N^c}{3c} \left(\frac{1}{2} \ddot{I}_{abc} + \ddot{P}_{a,bc} + \ddot{P}_{b,ac} - \ddot{P}_{c,ab} \right) + \dots \right) \Big|_{t_R}. \quad [2.43c]$$

These are the first few terms of the double expansion in velocities and distance from the sources. We are not giving all metric components to the same higher order, this partial result is sufficient for our purposes. In the PN expansion this is the metric parametrization in the wave zone, with the no-incoming radiation condition, and completely gauge fixed. The lowest order of the time-time component reproduces the Newtonian result.⁵ The first PN correction is the movement of the source, and can always be set to zero by going to the rest frame. Doing so eliminates the lowest order of the \bar{h}_{0a} component. The first corrections in that component contain the gravito-magnetic effects relevant to the Lense-Thirring effect, for instance. Notice also that the angular momentum is sub-leading in R , as one could have expected from a large distance expansion of Kerr's metric.

The radiative degrees of freedom are in the spatial components [2.43c] and can be extracted acting with the projector [1.90]. We have the mass quadrupole at leading order, and the first PN correction features the mass octupole and momentum quadrupole. We can immediately remark the absence of monopole and dipole contributions to the emission of waves. This is a direct consequence

⁵ In particular the lowest order of the Schwarzschild metric is obtained as the special case with constant M and all the rest vanishing. Notice in fact that $h_{00} = \bar{h}_{00} + \bar{h}/2 = 2GM/c^2 R$.

of the conservation laws, since they imply that the mass monopole and dipole have vanishing second time derivatives. As a consequence, an oscillating spherical distribution would not emit gravitational waves, in agreement with Birkhoff theorem in the full theory, nor would a distribution with axial symmetry rotating at constant velocity, in agreement with Kerr's solution.

Applying the TT projector removes any trace, hence $h^{\text{TT}} = \bar{h}^{\text{TT}}$ and one can replace I_{ab} with the irreducible quadrupole moment Q_{ab} , and obtain at lowest order

$$h_{ab}^{\text{TT}}(t, \vec{x}) = \frac{2G}{c^4 R} \ddot{Q}_{ab}^{\text{TT}}(t_{\text{R}}). \quad [2.44]$$

This is the celebrated *first quadrupole formula*, derived by Einstein in 1918: The dominant radiation in the slow-motion approximation arises from the acceleration of the quadrupole moment of the mass distribution. From this we can also obtain the expressions for the two independent polarizations. If $\vec{k} = \hat{z}$, we can use [1.72a] and

$$h_{+}(t, r) = \frac{G}{c^4 R} (\ddot{Q}_{11} - \ddot{Q}_{22})|_{t_{\text{R}}}, \quad h_{\times} = \frac{2G}{c^4 R} \ddot{Q}_{12}|_{t_{\text{R}}}. \quad [2.45]$$

For a general \vec{k} it is obtained replacing $h_{ab} \rightarrow (G/c^4 R) \ddot{Q}_{ab}$ in [1.94]. Notice also that $P^{\text{TT}}(Q) = P^{\text{TT}}(I)$ since the projector removes the trace, hence we can replace Q_{ab} with I_{ab} in these expressions.

Let us make some order-of-magnitude estimates. By dimensional analysis, the mass multipoles scale like Mr^l , where r is the typical size of the source. If the dynamics of the system has a typical velocity scale v , then $Q \sim Mr^2$ and $\ddot{Q} \sim Mv^2$. This gives

$$h \sim \frac{GMv^2}{c^4 R} = 5 \times 10^{-19} \left(\frac{M}{10M_{\odot}} \right) \left(\frac{1\text{Mpc}}{R} \right) \frac{v^2}{c^2}. \quad [2.46]$$

For example, if we extrapolate this formula to relativistic speeds $v \sim c$ for the merger of two 10-solar-masses black holes, we get a 10^{-18} amplitude at galactic distances, and 10^{-21} at 100 Mpc where the Virgo cluster is located.

This estimate is the lowest order of various approximations, which is useful to recap here: (1) weak-field, PM expansion; (2) long-distance, multipolar expansion; (3) small velocities, PN expansion. To obtain more accurate results, one has to include higher order corrections. Doing so is actually far from simple. Not only do we have three different expansion parameters with non-trivial

hierarchies among them, we also have to face both technical and conceptual challenges. Let us list a few, and tools used to deal with them. The PN expansion is not a convergent series, but rather what is known as an asymptotic series. Its accuracy degrades as we increase R . Dealing with this mathematical problem requires techniques such as the matched asymptotic expansion. Related to this is also the more conceptual issue that the causal propagation determined by the Green's function at lowest order follows the null cones of the background Minkowski metric. But null cones are bent by the gravitational interaction, hence higher order corrections have to also modify the retarded time to the correct one. For instance for the Schwarzschild metric the correct retarded time is

$$\begin{aligned} u &= t - R/c - 2GM/c^2 \ln(R - 2GM/c^2) \\ &= t_{\text{R}} + \frac{2GM}{c^2} \ln R - \left(\frac{2GM}{c^2}\right)^2 \frac{1}{R} + O(R^{-2}). \end{aligned}$$

Hence higher orders change the notion of retarded time.⁶ Another tricky effect comes in at higher orders: the waves backscatter and self-interact, causing a delay in part of the signal, which starts travelling inside the light-cone, similar to light slowing down in a medium due to interactions with the medium. Then the total signal includes a ‘tail’ that comes after the main part of the signal. To take this into account one has to include effects that arise from integration over time. Another problem is that divergences appear after the first iteration, because convolution of Poisson integrals diverge even if the initial source has compact support. To regularize this unphysical divergence one has to split the integrals into near-zone and far-zone integrations.

These and other types of difficulties plagued the theory throughout most of the seventies, and were addressed thanks to the work of many brilliant researchers, including pioneers like Thorne, Will, and Damour. On the phenomenological side, people thought for a while that the lowest quadrupole order would have been enough to match experiments, given the weakness of the waves. Later theoretical work, e.g. the seminal paper (Cutler *et al.* 1993), clarified the observational sensitivity to the PN corrections and justified the importance of the endeavour. The task is very challenging, and researchers have come up with different approaches. We refer to the specialized literature (Thorne 1980 ; Blanchet 2006 ; Poisson and Will 2014 ; Goldberger and Rothstein 2006) for reviews of this more advanced topic. In the following we will

6. This issue can be solved non-perturbatively using the Bondi coordinates mentioned earlier, but then one is switching from a PN expansion to an asymptotic expansion of the full theory, and the approach is both conceptually and technically different.

content ourselves to stay at lowest order, which is enough to understand the basics of the physics, if not for a detailed match to observations.

2.3. Dissipation by gravitational waves

The first historical evidence of gravitational waves was the orbital decay of the Hulse-Taylor pulsar. Accurate measurements showed that the orbital decay was consistent with the prediction of general relativity. Indeed, general relativity predicts that gravitational waves carry energy away from a system that produces them. In this Section we describe how this prediction arises at the lowest order in perturbation theory, and in the next Section we show how it can be applied to predict the orbital decay.

2.3.1. Energy of gravitational waves

Let us look at a gravitational wave as a spin-2 field propagating on the Minkowski background. Thanks to the Poincaré invariance of the background, we can apply Noether's theorem and derive a conserved energy-momentum tensor for $h_{\mu\nu}$. An explicit calculation starting from the linearized Lagrangian gives

$$t_{\mu\nu}^N = \frac{c^4}{32\pi G} \left(\partial_\mu h^{\alpha\beta} \partial_\nu h_{\alpha\beta} - \frac{1}{2} \eta_{\mu\nu} \partial_\lambda h_{\rho\sigma} \partial^\lambda h^{\rho\sigma} \right), \quad [2.47]$$

where the label N stands for Noether, and we assumed here the De Donder condition to simplify the expression. This tensor is conserved, namely $\partial_\mu t^{N\mu\nu} \doteq 0$, but has no clear physical meaning, because it is *not* gauge-invariant: It changes under a linearized diffeomorphism [1.52], including those compatible with the De Donder condition, and consequently assigns a non-zero value of energy-momentum to pure gauge modes. Furthermore, we can make it vanish entirely at any point using Riemann normal coordinates, since in these coordinates the first derivatives of the metric vanish at that point. Since it is zero at one point in one coordinate system but not in others, it is not a tensor. It is usually referred to as 'pseudo-tensor'.

The only gauge invariant quantities that can be extracted from [2.47], at least at lowest order, are global ones, such as the *total* energy and momentum on a space-like hypersurface Σ of constant time,

$$E_{\text{GW}} = \int_\Sigma t^{N00} d\Sigma = \frac{c^2}{64\pi G} \int_\Sigma \dot{h}^{\alpha\beta} \dot{h}_{\alpha\beta} d\Sigma, \quad [2.48a]$$

$$P_{\text{GW}}^a = \int_\Sigma t^{Na0} d\Sigma = -\frac{c^3}{32\pi G} \int_\Sigma \partial^a h^{\alpha\beta} \dot{h}_{\alpha\beta} d\Sigma. \quad [2.48b]$$

In fact, at lowest order in perturbation theory, a gauge transformation acts as [1.52] with $\xi = O(h)$, under which [2.47] transforms as a total derivative. Spatial total derivatives do not contribute to [2.48] under the standard condition that the perturbation goes to zero at infinity, and temporal ones can be converted to spatial derivatives on shell of the wave equation, hence the same argument applies. Similarly, one can define the total angular momentum as

$$L_{\text{GW}}^a := -\frac{c^4}{32\pi G} \epsilon^{abc} \int_{\Sigma} (\dot{h}_{\mu\nu} x_b \partial_c h^{\mu\nu} + 2\dot{h}_{b\mu} h_c{}^\mu) d\Sigma. \quad [2.49]$$

For our applications below, the total energy-momentum and angular momentum will be sufficient. They are not sufficient however for iterating the perturbation theory beyond lowest order, and we now briefly describe some of the conceptual aspects of what is done in the literature in this respect. First of all, the reader may recall that Noether currents are not unique, being defined only up to adding total derivatives whose conservation is trivial. In the case at hand, it means that any quantity of the type

$$t_{\mu\nu}^N + \partial^\rho \partial^\sigma U_{\mu\rho\nu\sigma}, \quad [2.50]$$

where U has the same index symmetries as the Riemann tensor, is an equally valid Noether current. One may hope that there exists a representative in the equivalence class [2.50] that would be gauge invariant, but this is not the case: the lack of gauge-invariance is simply the linearized version of the fact that there cannot be any local tensor representing the energy of the gravitational field. Such quantity will have to be zero in a local free-falling frame where the effects of gravity are absent, and if it were a tensor, it would then be zero in any frame.⁷ This is therefore a direct consequence of the equivalence principle, and its mathematical implementation via general covariance.⁸

7. A tensorial quantity capturing *some* aspects of gravitational energy can be constructed using the Bel-Robinson tensor, but it is fourth-order in derivatives, therefore does not have the right physical dimensions, and will capture only higher-order terms of the gravitational energy.

8. It is instructive to put this problem in perspective with what happens in the electromagnetic case. If one computes the canonical energy-momentum tensor of Maxwell's theory using the Noether formula, one also finds a meaningless gauge-dependent expression. However, the Noether construction only defines the tensor up to total divergences, and it is possible to find one that gives a gauge-invariant expression, and which is furthermore symmetric and coincides with the one derived from the variation with respect to the metric. In gravity there is an analogue problem, but even adding total divergences it is not possible to find a local gauge invariant quantity.

The obstruction to identifying a local energy density for the gravitational field shows up very prominently in the full theory. We have already discussed how Noether charges for generic diffeomorphisms are all trivial in the bulk of spacetime. Similarly, the bulk Hamiltonian one finds from the Legendre transform of the Lagrangian is a sum of constraints, and thus identically zero when evaluated on solutions. Any attempt to work around these facts and define quasi-local observables representing the gravitational energy unavoidably run into trouble with ambiguities and dependence on coordinates or other unphysical background structures (Szabados 2009). The clearest well-defined resolution to this problem is to consider not the energy density, but only the total energy. This is useful when describing isolated systems, namely spacetimes that are fully dynamical in a certain region, but become well approximated by flat spacetime at large distances from this region. In this case, one can introduce a physically meaningful notion of boundary to the spacetime, and exploit the fact that the Hamiltonian picks up a boundary contribution which is non-vanishing on solutions. The resulting *surface charges* can be used to characterise the total energy momentum and angular momentum of the system, and can be derived as Noether charges as well. Examples of this construction are the ADM charges at spatial infinity, and the BMS charges at future null infinity, as mentioned in Sec. 1.2.2.

While looking at global quantities such as the total energy of a free GW [2.48], or the surface charges in asymptotically flat spacetimes mentioned at the beginning of the section, is the safest way to define energy in the full theory, the perturbative treatment offers an alternative, ‘quasi-local’ possibility. Perturbatively in fact, it is possible to construct gauge-invariant quantities by introducing a spacetime averaging procedure based on the properties of the background. We consider a region L whose size is much larger than the typical wavelength λ of the perturbation, but much smaller than the typical wavelength λ_B of the background (which is infinite for a flat background), and we define the averaging of a functional F as $\langle F \rangle := \frac{1}{L} \int_L F$. If applied to an expression quadratic in the Fourier modes like [2.47], the procedure suppresses combinations with different frequencies or different phases, in a way completely similar to how the total energy in a standard background-dependent theory comes mainly from positive interference superposition of waves. The difference is that in background-dependent theories averaging the energy is a choice, since the local energy density is theoretically also well defined. In gravity it is not a choice but mandatory, since there is no meaningful local energy density, and furthermore care is needed to define correctly the procedure in a way to make it compatible with general covariance. Detailed analysis (Isaacson 1968 a,b ; Burnett 1989) shows that the result of the procedure is that expressions under the averaging sign can be freely integrated by parts in space and, upon going

on-shell, also in time derivatives since a wave propagates on the light-cone. For instance,

$$\langle \partial_\mu h_{\alpha\beta} \partial^\mu h^{\alpha\beta} \rangle = -\langle h_{\alpha\beta} \square h^{\alpha\beta} \rangle = 0 \quad [2.51]$$

outside the sources. Under this procedure, we find

$$\langle t_{\mu\nu}^N \rangle = \frac{c^4}{32\pi G} \langle \partial_\mu h^{\alpha\beta} \partial_\nu h_{\alpha\beta} \rangle \quad [2.52]$$

in a region outside the sources. One can show that the averaging procedure makes the right-hand side gauge-invariant (Isaacson 1968*b*). This means that it can be expressed in terms of the TT projection and the gauge invariant potentials. The latter can be neglected if the sources variation (induced by the partial derivatives in the expression above) occurs over much longer time scales than the h^{TT} wavelengths.⁹ This motivates the definition of

$$t_{\mu\nu} := \frac{c^4}{32\pi G} \partial_\mu h_{\text{TT}}^{ab} \partial_\nu h_{ab}^{\text{TT}}, \quad t_{\mu\nu} = \langle t_{\mu\nu}^N \rangle. \quad [2.53]$$

This quantity is actually gauge-invariant at lowest order, since the only non-invariant terms are the partial derivatives, and these transform linearly in $\xi \sim O(h)$. It follows that in so far as only lowest order results are needed, we can use the simpler expression [2.53] as a proxy for the averaging procedure. Notice also that it coincides with [2.47] in the TT gauge.

The expression [2.47] in the TT gauge is the one used by Einstein to determine the energy carried away by gravitational waves. It has, however, a limited applicability. We have already discussed its gauge dependence. Another issue is that it relies heavily on the specific background chosen, and had we worked with a non-isometric one, then there would be no Noether charge to begin with. Furthermore, it is not clear how to extend this construction to treat higher orders in perturbation theory. These shortcomings can be addressed if we look at a different definition for the gravitational energy-momentum pseudo-tensor, based on the actual back-reaction on the metric caused by the gravitational waves. In fact the actual “effective” source that determines the second-order metric perturbation is not [2.47], but rather the second order expansion of the Einstein tensor, that we denoted $t_{\mu\nu}^G$ in [1.38]. The candidate gravitational

9. In (Poisson and Will 2014) this step is called short-wave approximation, and it is performed without the averaging, in the context of the Landau-Lifshitz approach described in Appendix 2.7.

energy-momentum pseudo-tensor $t_{\mu\nu}^G$, obtained in this way is also conserved. In fact, an explicit calculation shows that it differs from [2.47] precisely by a term like [2.50], with U a certain quadratic expression in derivatives of $h_{\mu\nu}$. It has the improved property that it depends on second derivatives of the metric, so it cannot be made to vanish at any given point. However, it is still not gauge-invariant. Therefore one has to invoke again the averaging procedure in order to extract gauge-invariant information. Upon doing so, one finds that the two prescriptions give a consistent answer (Isaacson 1968*b*):

$$\langle t_{\mu\nu}^G \rangle = \langle t_{\mu\nu}^N \rangle = t_{\mu\nu}. \quad [2.54]$$

This matching supports averaging as a viable way to extract unambiguous and gauge-independent quantities. The prescription t^G overcomes some limitations of the Einstein-Noether construction. It can be used in perturbation theory around an arbitrary background, and can be systematically extended to any order in perturbation theory, by computing higher order corrections $G_{\mu\nu}^{(n)}$ and evaluating them on the perturbed solution. This procedure is however not very handy, and a better scheme is the one proposed by the Landau-Lifshitz reformulation of Einstein's equations, see e.g. (Blanchet 2006 ; Poisson and Will 2014). There one changes variables from the metric to a densitized inverse metric $g^{\mu\nu} := \sqrt{-g}g^{\mu\nu}$. One advantage of this reformulation is that it provides a full non-perturbative expression for a candidate energy-momentum, known as Landau-Lifshitz pseudo-tensor. This has the usual limitations (gauge-dependence and vanishing at any point in a local inertial frame) dictated by the equivalence principle, but has the merit of being set up in a way that makes it very natural to develop a systematic perturbative expansion, since the pseudo-tensor is defined already at non-perturbative level, and does not need to be determined order by order as in the previous approach. Furthermore, it provides a prescription for the energy, momentum and angular momentum as surface charges that, even though restricted in validity to Cartesian coordinates in the region far away from the sources, can be evaluated including higher orders, and bypasses the need for the spatial averaging of volume integrals. For these reasons, the Landau-Lifshitz formulation is widely used by the community working in the post-Newtonian expansion. We review it briefly in Appendix 2.7. The lowest order in the weak-field approximation of the Landau-Lifshitz pseudo-tensor differs from the previous two options by a term like [2.50], and also matches the gauge-invariant result after averaging:

$$\langle t_{\mu\nu}^{LL} \rangle = t_{\mu\nu}. \quad [2.55]$$

Similar considerations apply also to define the angular momentum of gravitational waves. In this case, the result of the averaging procedure starting from

any of the three prescriptions described above motivates the following definition (DeWitt 2011 ; Thorne 1980 ; Poisson and Will 2014)

$$j^a := \frac{1}{2} \epsilon^a{}_{bc} j^{bc} = -\frac{c^4}{32\pi G} \epsilon^{abc} (\dot{h}_{de}^{\text{TT}} x_b \partial_c h_{\text{TT}}^{de} + 2\dot{h}_{bd}^{\text{TT}} h_{cd}^{\text{TT}}). \quad [2.56]$$

As before, one can drop the TT projector and still get a gauge-invariant quantity after global integration or averaging.

We conclude that in so far as one is interested only in lowest order results, any of these three choices are equally good. For a systematic perturbative expansion, the latest is the better one. In fact even if the construction of gauge invariant quasi-local quantities via averaging is conceptually useful to clarify how gauge-invariant information could be extracted in principle, it is not very practical. To set up a systematic perturbative expansion, it is easier to work with gauge-fixed quantities at all intermediate steps, and then extract only at the end the physical predictions in terms of gauge-invariant observables. For instance, there is no problem in working with the non-averaged notions of energy-momentum and angular momentum pseudo-tensors, as long as one does not attempt to give them a direct physical interpretation. The idea is to use them to perform calculations, and at the end read off the physical dynamics not from their evolution but from that of gauge-invariant quantities such as the amplitude and frequency of TT modes, or the evolution of relative distances such as the periastron of an orbit. This is the logic used in the PN expansion, and based on the Landau-Lifshitz reformulation (Blanchet 2006 ; Poisson and Will 2014).

2.3.2. Dissipation equations

The fact that [2.53] is conserved means that we can derive identities between time and spatial derivatives like those that led to the conservation laws [2.30] for the matter sources. The key difference however is that the matter sources had compact support, hence we could neglect boundary contributions when integrating by parts. This is no longer true for the gravitational contributions, since the waves have non-compact support. The non-vanishing of the boundary terms has the effect that the ‘charges’ corresponding to energy, momentum and angular momentum are no longer conserved. This dissipation is precisely the statement that gravitational waves carry energy and have a physical impact on the system.

At lowest order, we do not even need to use the conservation equation in order to study the dissipation, because of a special property of the explicit solution [2.43]. Each metric component has functional dependence on coordinate of the form $f(t_R, N^a)$. For such functions, it is easy to check that

$$\partial_a f = -\frac{N_a}{c} \dot{f} + O(R^{-1}). \quad [2.57]$$

The leading order of this approximation plays an important role in simplifying many formulas in the wave zone, where $R \gg 1$.

Let us begin our analysis from the flux of gravitational energy, namely the emitted power. This is given by

$$\dot{E}_{\text{GW}} = \int_{\Sigma} \dot{t}^{00} d^3x = -c \oint_{\partial\Sigma} t^{0a} N_a dS = \frac{c^4}{32\pi G} \oint_{\partial\Sigma} \dot{h}_{cd}^{\text{TT}} N^a \partial_a h_{\text{TT}}^{cd} dS, \quad [2.58]$$

Stokes theorem choosing as boundary a 2-sphere of radius R in the asymptotic region (hence the outgoing unit normal is simply N_a , and $dS = R^2 d^2\Omega$ where $d^2\Omega = \sin\theta d\theta d\phi$) and [2.53] in the last equality. The spatial derivative can be replaced at lowest order with a time derivative using again [2.57], and we arrive at

$$\dot{E}_{\text{GW}} = -\frac{c^3}{32\pi G} \oint_{\partial\Sigma} \dot{h}_{ab}^{\text{TT}} \dot{h}_{\text{TT}}^{ab} dS = -\frac{G}{8\pi c^5 R^2} \oint_{\partial\Sigma} \ddot{Q}_{ab}^{\text{TT}} \ddot{Q}^{\text{TT}ab} dS|_{t_R}, \quad [2.59]$$

where in the last step we used the explicit form [2.43] of the solution, in particular [2.44]. To evaluate the integral, we observe that the only angular dependence occurs in the TT projector [1.90]. Using the following formula,

$$\oint_{S^2} P^{\text{TT}cd}_{ab} d^2\Omega = \frac{8\pi}{5} \left(\delta_{(a}^c \delta_{b)}^d - \frac{1}{3} \delta_{ab} \delta^{cd} \right), \quad [2.60]$$

we find

$$\dot{E}_{\text{GW}} = -\frac{G}{5c^5} \ddot{Q}_{ab} \ddot{Q}^{ab}|_{t_R}. \quad [2.61]$$

This is the second famous quadrupole formula of Einstein (Einstein 1918). It gives the instantaneous power radiated at a distance R from the source and a time t , as a function of the quadrupole time variation at the retarded time $t - R/c$. Notice that the index contraction occurs over *all* indices of the (traceless) quadrupole moment. The effect of the *TT* projection goes into a numerical factor, after the integration [2.60].

For the linear momentum loss,

$$\begin{aligned}\dot{P}_{\text{GW}}^a &= \frac{1}{c} \int_{\Sigma} \dot{t}^{0a} d^3x = - \int_{\Sigma} \partial_b t^{ab} d^3x \\ &= - \oint_{\partial\Sigma} t^{ab} N_b dS = - \frac{c^2}{32\pi G} \oint_{\partial\Sigma} N^a \dot{h}_{cd}^{\text{TT}} \dot{h}_{\text{TT}}^{cd} dS,\end{aligned}\quad [2.62]$$

where we used twice [2.57] in the last equality. Since N^a is an odd function on the sphere, the integral vanishes: *there is no loss of momentum at lowest order* (namely at order G/c^6 , once we use the first quadrupole formula). A change in the total momentum of the system caused by the emission of GWs ('kick') occurs only at the next order G/c^7 , when mixing of multipoles of different parity occurs.

For the angular momentum loss,

$$\dot{L}_{\text{GW}}^a = \int_{\Sigma} d^3x \partial_t j^a = - \frac{1}{c} \oint_{\partial\Sigma} j^a dS \frac{c^3}{16\pi G} \epsilon^{abc} \oint_{\partial\Sigma} (\dot{h}_{de}^{\text{TT}} x_b \partial_c h_{\text{TT}}^{de} + 2\dot{h}_{bd}^{\text{TT}} h_c^{\text{TT}d}) dS.\quad [2.63]$$

Using the quadrupole formula [2.44] and performing the integrals using identities similar to [2.60], one arrives at

$$\dot{L}_{\text{GW}}^a = - \frac{2G}{5c^5} \epsilon^{abc} \ddot{Q}_{bd} \ddot{Q}_c^d |_{t_{\text{R}}}.\quad [2.64]$$

Angular momentum loss occurs at the same order as energy loss, and involves one lesser time derivative.

2.3.3. On the validity of the quadrupole formula

The quadrupole formula [2.44], and its application leading to the second quadrupole formula [2.61], were derived in the linear approximation. In this approximation, the energy-momentum tensor satisfies the flat spacetime conservation law [2.26], and the geodesics of matter are straight lines. It is thus valid only for systems whose gravitational interaction is negligible. It is not valid, in particular, for a binary system held together by gravity, even at the non-relativistic, Newtonian level. To treat a gravitational binary, one has to go beyond the linear approximation, using the iterative scheme described around [1.37]. In the first step, one should include only the kinetic and non-gravitational pieces of the dynamics in the moment of inertia sourcing the quadrupole formula, consistently with the motion being along straight lines. The result is then

used to source the effective energy-momentum tensor $t_{\mu\nu}$ needed for the second iteration of the field equations, which determines $h_{\mu\nu}^{(2)}$ of [1.37]. When this approach is applied to a gravitational binary, one finds that the second-order correction restores precisely the Newtonian contribution to the quadrupole moment. This is quite remarkable, and means that *the same quadrupole formula is also valid if one includes the Newtonian interaction*. See for instance (Poisson and Will 2014 ; Deruelle and Uzan 2018) for details.

We will make use of this fact below, and deduce the right results for gravitational binary systems using the quadrupole formula with the Newtonian potential included in the moment of inertia, without going into the technical details required to solve the second iteration, for which we refer to the cited literature. However, we should keep in mind that this formula can only be trusted *because it has been derived including the second iteration*, and not from the linearized approximation alone. We should also keep in mind that restoring the validity of the quadrupole formula after the second iteration in gravitational binaries is a special and remarkable fact, and not a general result. For instance, the linearized quadrupole formula gives the wrong answer for gravitational binaries in modified theories of gravity, see e.g. (Taherasghari and Will 2025).

The fact that the linearized approximation provides potentially correct answers beyond its regime of validity contributed to the controversy that heated the debate around the quadrupole formula, and that were ultimately solved only with the systematic and rigorous developments of the early 80's, see (Damour 1982 ; Kennefick 1997 ; Will 2011) for discussions.

2.3.4. Back-reaction

Observable consequences of the emission of gravitational waves can be studied looking at how they impact the dynamics of the source. This can be done for instance evaluating the first PM correction to the source trajectories, by solving

$$0 = \text{NewA}_\mu T^{\mu\nu} = \partial_\mu \bar{T}^{\mu\nu} + \Gamma_{\mu\rho}^{(1)\mu} \bar{T}^{\rho\nu} + \Gamma_{\mu\rho}^{(1)\nu} \bar{T}^{\mu\rho} + \partial_\mu \bar{T}^{(1)\mu\nu} + O(h^2) [2.65]$$

at various orders in the PN expansion. The details of this calculation can be found in (Poisson and Will 2014). In some cases however, it is possible to consider the following shortcut. At zeroth order, we have the Newtonian dynamics, and this comes with a clear identification of conserved quantities such as energy E and angular momentum J . We then *assume* that the first order correction is obtained allowing these quantities to be not conserved, and equating their change to the dissipation caused by gravitational waves. That is, we posit

$$\dot{E} = -\dot{E}_{\text{GW}}, \quad \dot{L}^a = -\dot{L}_{\text{GW}}^a, \quad [2.66]$$

insert the expressions [2.48] and [2.49] on the right-hand side, and solve the resulting equations. From the solutions we deduce how the source should change in time in order for its dynamics to be consistent with the dissipation caused by gravitational waves. This is what we do in the next Section, and for which the total expressions [2.48] and [2.49] are enough. Detailed calculations using the proper method [2.65], see e.g. (Poisson and Will 2014), confirm the validity of this shortcut, at least in so far as the lowest order in the PN expansion is concerned.

2.4. GWs from binary systems: elliptical, circular and hyperbolic orbits

We now apply the results of the previous section to determine the GW signal from binary systems. We will first consider the case of a bound system, with circular or elliptical orbits. These provide a simple yet realistic model of astrophysical sources that corresponds to the signals observed by LVK. We will see how one can express the two quadrupole formulas (and more generally the dissipation equations) in terms of the dynamics of the sources, compute the backreaction leading to orbital decay and increased wave emission, and produce analytic waveforms. We will also see explicitly the importance of the averaging procedure, which in the case of bound binary systems neatly separates the effects related to the two time-scales involved: the period of each orbit, and the ‘secular’ effects that cumulate over many orbits. We will then consider the case of unbounded, hyperbolic orbits, produce their waveforms. These orbits are interesting because they provide the simplest examples of displacement memory and gravitational capture. Throughout this Section, we will approximate the gravitational bodies with non-spinning point particles. This provides a good approximation at lowest order: detailed PN analysis shows that spinning and finite-size effect only enter at higher orders.

2.4.1. Newtonian equations

We first recall the Newtonian equations of motion for two non-spinning point-particles of masses $m_{1,2}$, with relative position $\vec{r} = \vec{x}_1 - \vec{x}_2$ and relative velocity $\vec{v} = \vec{v}_1 - \vec{v}_2$. In the center-of-mass (CM) frame, the dynamics can be described by a single particle with position \vec{r} and reduced mass

$$\mu := \frac{m_1 m_2}{m}, \quad \nu := \frac{m_1 m_2}{m^2}, \quad m = m_1 + m_2. \quad [2.67]$$

The dimensionless quantity ν is introduced here for later convenience. The total energy $E = \frac{1}{2}\mu v^2 - G\mu m/r$ and angular momentum $\vec{L} = \mu \vec{r} \times \vec{v}$ are

conserved. The latter implies that the motion is confined to a plane, and we choose coordinates so that this is the (x, y) plane. We then parametrize

$$\vec{r} = r\vec{n}, \quad \vec{n} = (\cos \psi, \sin \psi, 0), \quad [2.68]$$

and introduce a second vector to form an orthogonal basis in the plane of the dynamics:

$$\vec{v} = \dot{r}\vec{n} + r\dot{\psi}\vec{\lambda}, \quad \vec{\lambda} = (-\sin \psi, \cos \psi, 0). \quad [2.69]$$

The observer's detector is at position $\vec{R} = R\vec{N}$ where in spherical coordinates the unit vector is $\vec{N} = (\sin \varphi \cos \theta, \sin \varphi \sin \theta, \cos \varphi)$, see figure 2.4. Newton's




Figure 2.4: Left panel: Bound binary system in the centre-of-mass frame: basic quantities and angles. Right panel: Geometric quantities of an ellipse.

equations in the CM frame can be conveniently rewritten as

$$r(\psi) = \frac{p}{1 + e \cos(\psi)}, \quad [2.70]$$

$$\dot{\psi} = \sqrt{\frac{Gm}{p^3}} (1 + e \cos(\psi))^2 \quad [2.71]$$

where $e = \sqrt{1 - (b/a)^2}$ is the *eccentricity*, and $p = a(1 - e^2)$ the *semi-latus rectum*, of an ellipse with semi-major axis a and semi-minor axis b , see right panel of Fig. 2.4. These geometric quantities are related to the physical conserved quantities by

$$E = \frac{G\mu m}{2a} = \frac{G\mu m}{2p} (e^2 - 1) = \nu \frac{Gm^2}{2p} (e^2 - 1). \quad [2.72]$$

and

$$\vec{L} = L\vec{e}_z, \quad L = \mu\sqrt{Gmp} = \nu\sqrt{Gm^3p}. \quad [2.73]$$

It follows from Eqs. [2.68]-[2.71] that

$$\vec{v} = \sqrt{\frac{Gm}{p}} (-\sin \psi, e + \cos \psi, 0), \quad [2.74]$$

where we have chosen the origin $\psi = 0$ at periastron, and

$$\frac{|\vec{v}|^2}{c^2} = \left(\frac{Gm}{c^2p}\right) (1 + e^2 + 2e \cos \psi). \quad [2.75]$$

The Newtonian approximation requires $\vec{v}^2 \ll c^2$, and thus the dimensionless ratio $Gm/c^2p \ll 1$.

Bound systems have eccentricity $e < 1$, while unbound ones have $e > 1$. The border case $e = 1$ corresponds to parabolic orbits.

- **Elliptical Orbits** have $0 < e < 1$ with $-\pi \leq \psi < \pi$, $r_{\min} = p/(1 + e)$ and $r_{\max} = p/(1 - e)$. The orbital angular frequency ω_0 and period T satisfy Kepler's laws

$$\omega_0 = \sqrt{\frac{Gm(1 - e^2)^3}{p^3}} \quad \text{and} \quad T = \frac{2\pi}{\omega_0}. \quad [2.76]$$

- **Circular orbits** have $e = 0$ and radius $r = p$, and orbital frequency

$$\omega_0 = \sqrt{\frac{Gm}{p^3}}. \quad [2.77]$$

- **Hyperbolic orbits** have $e > 1$. Now $\psi_-(e) \leq \psi < \psi_+(e)$ where

$$\psi_{\pm} = \pm \arccos(1/e), \quad [2.78]$$

and correspondingly $\sin \psi_{\pm} = \pm e^{-1} \sqrt{e^2 - 1}$. These orbits are not periodic, but have a characteristic time-scale (a burst time scale) related to the characteristic frequency scale

$$\omega_c = \sqrt{\frac{Gm(e^2 - 1)^3}{p^3}}. \quad [2.79]$$

The closest distance of approach $r_{\min} = p/(1 + e)$ at $\psi = 0$. As $\psi \rightarrow \psi_{\pm}$, $v \rightarrow v_{\infty}$ with

$$\frac{v_{\infty}^2}{c^2} = \frac{Gm}{c^2 p} (e^2 - 1). \quad [2.80]$$

The GR modifications of the Newtonian equations of motion can be separated in two classes. First, the contributions of the potentials, of which the most famous is of course Einstein's initial precession calculation: over an orbital period the perihelion of elliptical orbits advances by $\Delta_h = 2\pi(3Gm/c^2 p)$, while for hyperbolic orbits $\Delta_h = (\Delta_r/3) \{6 \arccos(-1/e) + e^{-2} \sqrt{e^2 - 1} [2(2 + e^2) + 5\nu(e^2 - 1)]\}$, see e.g. (Damour and Deruelle 1985). Second, non-conservative or dissipative effects caused by the emission of gravitational waves. In the following we will ignore the first, and focus on the second. We will also ignore the precession introduced by dissipative effects, since this is higher order in the PN expansion.

2.4.2. Energy and angular momentum fluxes

In the quadrupole approximation, the TT component of the waveform is given by Eq. [2.44], in terms of the traceless quadrupole tensor $Q_{ij} = I_{ij} - \frac{1}{3}I\delta_{ij}$. The mass quadrupole moment of the two point particles is

$$I_{ab} = \mu r_a r_b = \nu m r_a r_b. \quad [2.81]$$

This is where the discussion of Section 2.3.3 becomes crucial. If we are working in the linear approximation, we are only allowed to include in I_{ab} the kinematical contribution from the velocities, and not the Newtonian gravitational potential, which would curve the geodesics away from the straight lines. However going through the second iteration has ultimately the neat effect of restoring the quadrupole formula with the Newtonian potential included. We will make use of this fact, and assume in the following that the quadrupole formula was derived already going through the second iteration of the field equations, and thus include the Newtonian potential in I_{ab} . Doing so, and using the equations of motion $d\vec{v}/dt = -Gm\vec{n}/r^2$, it follows that

$$\ddot{I}_{ab} = 2\nu m \left(v_a v_b - \frac{Gm}{r} n_a n_b \right). \quad [2.82]$$

Thus from Eqs. [2.68] and [2.74] the non-zero components of \ddot{I}_{ab} are

$$\ddot{I}_{11} = -2\nu m c^2 \left(\frac{Gm}{c^2 p} \right) [\cos(2\psi) + e \cos^3 \psi], \quad [2.83a]$$

$$\ddot{I}_{12} = -2\nu m c^2 \left(\frac{Gm}{c^2 p} \right) [\sin(2\psi) + e \sin \psi (1 + \cos^2 \psi)], \quad [2.83b]$$

$$\ddot{I}_{22} = 2\nu m c^2 \left(\frac{Gm}{c^2 p} \right) [\cos(2\psi) + e \cos \psi (1 + \cos^2 \psi) + e^2], \quad [2.83c]$$

and $\ddot{I} = 2\nu m c^2 \left(\frac{Gm}{c^2 p} \right) e(e + \cos \psi)$. We have written the coupling constants as

$$\frac{Gm_1 m_2}{p} = m c^2 \left(\frac{Gm}{c^2 p} \right) \nu, \quad [2.84]$$

in order to highlight that \ddot{I}_{ab} has the dimensions of energy. The third derivatives of the quadrupole tensor are straightforwardly obtained from Eq. [2.83] and [2.71] and read

$$\dddot{I}_{11} = 2\nu(m c^2) \frac{c}{p} \left(\frac{Gm}{c^2 p} \right)^{3/2} (1 + e \cos \psi)^2 [2 \sin(2\psi) + 3e \cos^2 \psi \sin \psi] \quad [2.85a]$$

$$\dddot{I}_{12} = 2\nu(m c^2) \frac{c}{p} \left(\frac{Gm}{c^2 p} \right)^{3/2} (1 + e \cos \psi)^2 [-2 \cos(2\psi) + e \cos \psi (1 - 3 \cos^2 \psi)] \quad [2.85b]$$

$$\dddot{I}_{22} = -2\nu(m c^2) \frac{c}{p} \left(\frac{Gm}{c^2 p} \right)^{3/2} (1 + e \cos \psi)^2 [2 \sin(2\psi) + e \sin \psi (1 + 3 \cos^2 \psi)] \quad [2.85c]$$




Figure 2.5: Surfaces of constant h_+ and h_\times polarisations for circular orbits $e = 0$. Here have chosen $\psi = \pi$.

The GW perturbation is given by substituting these expressions into Eq. [2.44]. In the direction $\vec{N} = \hat{z}$ the plus and cross polarisations are given by (see Eq. [2.45]),

$$h_+(t) = \frac{G}{c^4 R} (\ddot{I}_{11} - \ddot{I}_{22}) = -h_0 [2 \cos(2\psi) + e \cos \psi + 2e \cos^3 \psi + e^2] \Big|_{t_R} \quad [2.86]$$

$$h_\times(t) = \frac{2G}{c^4 R} \ddot{I}_{12} = -2h_0 [\sin(2\psi) + e \sin \psi (1 + \cos^2 \psi)] \Big|_{t_R}. \quad [2.87]$$

where the dimensionless amplitude is

$$h_0 = \frac{2G^2 m_1 m_2}{c^4 R p} = 2\nu \left(\frac{Gm}{c^2 R} \right) \left(\frac{Gm}{c^2 p} \right). \quad [2.88]$$

The time-dependence is determined from $\psi(t)$ which is a solution of Eq. [2.71]. The polarizations $h_{+,\times}(t, \theta, \varphi)$ in an arbitrary direction $\vec{N} = (\sin \theta \cos \varphi, \sin \theta \sin \varphi, \cos \theta)$ can be obtained plugging [2.83] in [1.94], or written directly in terms of (2.86-2.87) using [1.95]. Figure 2.5 shows surfaces of constant $h_{+,\times}$ as a function of (θ, φ) , for a fixed value of $\psi = \pi$ and $e = 0$. The quadrupolar nature is clearly visible.




Figure 2.6: Circular orbits. The plot shows h_+ and h_x polarisations, and emitted GW power (solid line), as a function of retarded time in units of T for 2.5 orbital periods. The emitted power is constant and given by Eq. [2.94]. The GW wavelength is $cT/2$.

For circular orbits ($e = 0$), only the terms in $\cos(2\psi)$ and $\sin(2\psi)$ remain, and furthermore the angular velocity is constant, in particular $\psi = \omega_0 t$ from [2.71]. Thus for circular orbits, the GW angular frequency is twice the orbital frequency:

$$\omega = 2\omega_0. \quad [2.89]$$

Figure 2.6 shows the waveforms as a function of retarded time, in units of T , for 2.5 orbital periods. For elliptical orbits additional frequencies are present, both larger and smaller than ω_0 . In fact since the angular velocity is not constant, infinitely many harmonics are emitted. These are shown in Fig. 2.7.

In the quadrupole approximation, the energy and angular momentum fluxes [2.61] and [2.64] read

$$P_{\text{GW}}(\psi) = \frac{G}{5c^5} \ddot{Q}_{ab} \ddot{Q}^{ab} = \frac{2G}{15c^5} \left[\ddot{I}_{11}^2 + \ddot{I}_{22}^2 + 3\ddot{I}_{12}\ddot{I}_{12} - \ddot{I}_{11}\ddot{I}_{22} \right] [2.90]$$

$$\dot{L}_{\text{GW}}^z(\psi) = \frac{2G}{5c^5} \epsilon^{zab} \ddot{I}_{ac} \ddot{I}_{b^c} = \frac{2G}{5c^5} \left[(\ddot{I}_{11} - \ddot{I}_{22}) \ddot{I}_{12} + \ddot{I}_{12} (\ddot{I}_{22} - \ddot{I}_{11}) \right] [2.91]$$




Figure 2.7: Elliptical orbit with $e = 0.3$. The plot shows h_+ and h_\times polarisations, and emitted GW power, as a function of retarded time in units of T for 2.5 orbital periods. The emitted power is largest at periastron $\psi = 0 \pmod{2\pi}$ where the orbital velocity is the largest.

(only the z -component of angular momentum is relevant since the binary is in the xy -plane). Substituting Eq. [2.85] gives

$$P_{\text{GW}}(\psi(t)) = P_{\text{GW}}^{e=0}(1 + e \cos \psi)^4 \left[1 + 2e \cos \psi + \frac{e^2}{12}(1 + 11 \cos^2 \psi) \right] \Big|_{t_{\text{R}}} \quad [2.92]$$

$$\dot{L}_{\text{GW}}^z(\psi(t)) = \dot{L}_{\text{GW}}^{e=0}(1 + e \cos \psi)^3 \left[1 + \frac{3}{2}e \cos \psi - \frac{e^2}{4}(1 - 3 \cos^2 \psi) \right] \Big|_{t_{\text{R}}} \quad [2.93]$$

where for circular orbits the constant rates of emission are given by

$$P_{\text{GW}}^{e=0} = \frac{32}{5} \frac{G^4}{c^5} \frac{\mu^2 m^3}{p^5} = \frac{32}{5} \left(\frac{c^5}{G} \right) \left(\frac{Gm}{c^2 p} \right)^5 \quad [2.94]$$

$$\dot{L}_{\text{GW}}^{e=0} = \frac{32}{5} \frac{G^{7/2}}{c^5} \frac{\mu^2 m^{5/2}}{p^{7/2}} = \frac{32}{5} \nu^2 (mc^2) \left(\frac{Gm}{c^2 p} \right)^{7/2} \quad [2.95]$$




Figure 2.8: *Hyperbolic orbits with $e = 2$. The h_+ and h_x polarisations and emitted GW power as a function of retarded time, for an initial value $\psi = \psi_-$ see Eq. [2.78]. The motion is no-longer periodic and a burst of GW energy emitted when $\psi = 0$ at $t = 0$. Note that the emitted power scales as e^6 for large e , see Eq. [2.92], and for that reason in the plot the power is normalised by an extra factor of e^6 . See section 2.4.6*

in terms of the dimensionless coefficient Gm/c^2p . The above expressions are valid for all $e \geq 0$ provided $\bar{v}^2 \ll c^2$. The power emitted is constant for circular orbits, and maximal at periastron for elliptical orbits. Figure 2.6 shows the waveforms Eq. [2.86]-[2.87] and power emitted Eq. [2.92] over 2.5 periods of a circular orbit with $e = 0$, in unit of t_R/T where T is the orbital period. Figure 2.7 shows the same for an elliptical orbit with $e = 0.3$. In both cases the periodic motion is clear. Over longer time-scales $t \gg T$, however, the emission of energy and angular momentum backreact on the orbital trajectories and must be considered. For a hyperbolic orbit, the corresponding plots are given in figure 2.8. The motion is obviously no-longer periodic and simply amounts to a fly-by: thus back-reaction effects do not accumulate over time and will be less significant (see subsection 2.4.6).

We now evaluate the effect of energy and angular momentum dissipation on the system's dynamics in order to study the back-reaction.

2.4.3. Back-reaction and waveform: circular orbits

The effect of the gravitational wave emission produces an effect that modifies the dynamics of the system, known as *radiation-reaction force*. As explained at the end of the previous Section, this effect can be deduced equating the non-conservation of the energy and angular momentum with the corresponding gravitational wave fluxes, see [2.66], and solving to find how the orbital element change in time. In order to do so we have to replace the constant orbital elements with functions of time, a method already used to compute perturbations to Keplerian orbits and known as ‘osculating orbits’. For circular orbits, there is only one independent quantity, it is thus sufficient to look at the energy equation. On the source side, we have [2.72]

$$\dot{E} = -\frac{E}{p}\dot{p}. \quad [2.96]$$

On the gravitational wave side, we have [2.94]. Equating the two gives

$$\dot{p} = \frac{64}{5} \frac{G^3}{c^5} \frac{\mu m^2}{p^3}, \quad [2.97]$$

and this tells us how the orbit evolves. Using Kepler’s law [2.77], we can obtain the evolution of the angular frequency as

$$\dot{\omega}_0 = A\omega_0^{11/3}, \quad A = \frac{96}{5} \frac{G^{5/3} \mu m^{2/3}}{c^5} = \frac{96}{5} \left(\frac{G\mathcal{M}}{c^3} \right)^{5/3}, \quad [2.98]$$

where \mathcal{M} is the chirp mass [1.6]. This in turns tells us the change in the frequency of the emitted waves, which as seen in [2.89] is twice the orbital frequency:

$$\dot{\omega} = 2^{-8/3} A \omega^{11/3} = \frac{12}{5} 2^{1/3} \left(\frac{G\mathcal{M}}{c^3} \right)^{5/3} \omega^{11/3}, \quad [2.99]$$

or equivalently in terms of $f = \omega/2\pi$,

$$\dot{f} = \frac{96}{5} \pi^{8/3} \left(\frac{G\mathcal{M}}{c^3} \right)^{5/3} f^{11/3}. \quad [2.100]$$

This shows that as gravitational waves are emitted, the orbit decays, the angular frequency increases, and so does the GW’s frequencies as well, leading

to an even greater orbital decay. This run-away effect stops when the orbit decays completely and the two body coalesce. Clearly before that happens, higher order effects become important and one should improve the calculation. It is however instructive to get the lowest-order approximation and a qualitative overall picture to assume that the quadrupole approximation is valid throughout the evolution. We can then integrate between an initial time t and the coalescing time t_c , where the frequency formally diverges. This gives

$$f(t) = \frac{1}{\pi} \left(\frac{5}{256(t_c - t)} \right)^{3/8} \left(\frac{GM}{c^3} \right)^{-5/8}. \quad [2.101]$$

Thus for a binary inspiral on a circular orbit $f^{-8/3}$ is linear in time, with a slope which determines directly the chirp mass. From this solution we can also immediately deduce the evolution of the orbital frequency ω_0 , the orbital radius p , and the GW amplitude [2.88].

To get the complete waveform, we only need the time dependence of ψ . This can be computed observing from [2.71] that $\dot{\psi} = \omega_0$, therefore

$$\Phi(t) := 2\psi = 2\pi \int_t^{t_c} dt' f(t') + \Phi_c = -2 \left(\frac{5GM}{c^3} \right)^{-5/8} (t_c - t)^{5/8} \quad [2.102]$$

where Φ_c is the phase at coalescing time. Combining these results, we get for the polarizations (2.86-2.87) in the direction perpendicular to the orbital plane

$$h_+(t) = -\frac{4}{R} \left(\frac{GM}{c^2} \right)^{5/3} \left(\frac{\pi f(t_R)}{c} \right)^{2/3} \cos \Phi(t_R), \quad [2.103]$$

$$h_\times(t) = -\frac{4}{R} \left(\frac{GM}{c^2} \right)^{5/3} \left(\frac{\pi f(t_R)}{c} \right)^{2/3} \sin \Phi(t_R). \quad [2.104]$$

For circular orbits, it is also easy to give the polarizations in an arbitrary direction $\vec{N} = (\sin \theta \cos \varphi, \sin \theta \sin \varphi, \cos \theta)$, using Eq. [1.95]:

$$h_+(t) = \frac{4}{R} \left(\frac{GM}{c^2} \right)^{5/3} \left(\frac{\pi f(t_R)}{c} \right)^{2/3} \cos[\Phi(t_R) - 2\varphi + \pi] \left(\frac{1 + \cos^2 \theta}{2} \right) \quad [2.105]$$

$$h_\times(t) = \frac{4}{R} \left(\frac{GM}{c^2} \right)^{5/3} \left(\frac{\pi f(t_R)}{c} \right)^{2/3} \sin[\Phi(t_R) - 2\varphi + \pi] \cos \theta. \quad [2.106]$$

As described after [1.95], the angle θ can be identified with the inclination ι of the source relative to the detector, see Fig. 1.1, whereas $2\omega = -2\varphi + \pi$ is the

longitude of pericenter. Notice also that ω can be absorbed into a redefinition of the initial time. This is the analytic expression of the curve plotted in blue in figure 2.9.

2.4.4. Back-reaction: Elliptical orbits

The case of elliptic orbits is more intricate, and offers a few interesting insight into the dynamics: there is a full spectrum of emission, and not a monochromatic one, which shows up in a non-symmetric waveform; there is loss of both energy and angular momentum, and at different rates, which shows up in the orbital back-reaction losing eccentricity faster than it decays; GW emission is not constant during the orbit, but stronger at periastron.

To analyse the system, it is convenient to distinguish the effects on two different time scales: short-time effects, namely the variations within a single orbit; and long-time, or *secular*, effects, namely the cumulative changes over many orbits. For instance, the dependence of the power emitted on the position ψ illustrated in Fig.2.7 is a short-time effect. The orbital decay on the other hand is a secular effect. To consider secular effects, we introduce the average over one orbital period:

$$\langle X \rangle = \frac{1}{T} \int_0^T dt X(t) = \frac{1}{T} \int_{-\pi}^{\pi} d\psi \frac{1}{\dot{\psi}} X(\psi), \quad [2.107]$$

where $\dot{\psi}$ is given in Eq. [2.71]. We then replace [2.66] with their time-averages,

$$\dot{E} = -\langle P_{\text{GW}} \rangle \quad \dot{L} = -\langle \dot{L}_{\text{GW}} \rangle. \quad [2.108]$$

The aim of this subsection is to solve these equations to determine the secular evolution of $e(t)$, $p(t)$, and thus $h_{+,\times}(t)$ with backreaction included.

Substituting [2.92] and integrating gives the *Peter-Mathews* formula (Peters and Mathews 1963):

$$\langle P_{\text{GW}} \rangle = P_{\text{GW}}^{e=0} (1 - e^2)^{3/2} \left[1 + \frac{73}{24} e^2 + \frac{37}{96} e^4 \right]. \quad [2.109]$$

(This expression is only valid for $e < 1$ as we are dealing with elliptical orbits.) Keeping p constant, the radiation increases from $e = 0$, to a maximum at $e \sim$

0.5 before decreasing and vanishing at $e = 1$. The averaged angular momentum radiation is similarly determined using [2.93] and gives

$$\langle \dot{L}_{\text{GW}} \rangle = \dot{L}_{\text{GW}}^{e=0} (1 - e^2)^{3/2} \left[1 + \frac{7}{8} e^2 \right]. \quad [2.110]$$

We now return to Eqs. [2.108], where on the left hand side the time-dependence is in $e(t)$ and $p(t)$. By definition, see [2.73], $L = \nu \sqrt{Gmp}$ from which

$$\frac{dL}{dt} = \frac{\nu c}{2} \sqrt{\frac{Gm}{c^2 p}} \frac{dp}{dt}. \quad [2.111]$$

This combined with Eqs. [2.108] and Eq. [2.110] gives

$$\frac{dp}{dt} = -\frac{64}{5} \nu c \left(\frac{Gm}{c^2 p} \right)^3 (1 - e^2)^{3/2} \left[1 + \frac{7}{8} e^2 \right]. \quad [2.112]$$

The energy of the orbit is given in Eq. [2.72], from which

$$\dot{e} = \frac{1}{\nu G m^2 e} p \dot{E} - \frac{\dot{p}}{2pe} (1 - e^2). \quad [2.113]$$

Then plugging in [2.109] and [2.112] gives

$$\frac{de}{dt} = -\frac{304}{15} \nu c \left(\frac{e}{p} \right) \left(\frac{Gm}{c^2 p} \right)^3 (1 - e^2)^{3/2} \left[1 + \frac{121}{304} e^2 \right]. \quad [2.114]$$

These coupled equations Eq. [2.112] and [2.114] can be solved using hypergeometric function to get $e(t)$ and $p(t)$, or alternatively combined to determine $p(e)$.

Observe that both $p(t)$ and $e(t)$ decrease with time. An elliptical orbit with initial eccentricity $e \neq 0$ will thus become more circular due to GW radiation. Whereas an initially circular orbit with $e = 0$ remains circular for all times. This is the reason why often it is a good approximation to consider circular orbits, particularly when studying the last moments before the merger of the binary system. (This is the case of the events observed by LVK.) At first order in e , we can approximate [2.114] with

$$\frac{de}{dt} \sim -\frac{e}{\tau_R}, \quad \tau_R := \frac{1}{\nu} \left(\frac{Gm}{c^2 p} \right)^{-5/2} \frac{T}{2\pi}, \quad [2.115]$$

where T is the orbital period [2.76]. This approximation provides a time-scale, τ_R , of the radiative decay of e and p . From Eq. [2.75], $\left(\frac{Gm}{c^2 p}\right) \sim |\vec{v}|^2/c^2 \ll 1$, and thus $\tau_R \sim (c/v)^{-5} T \gg T$.

The decrease of p and e also implies that T decreases with time. Indeed from Eqs. [2.76], [2.114] and [2.112]

$$\frac{dT}{dt} = -\frac{192}{5} \pi \left(\frac{GM}{c^3} \frac{2\pi}{T} \right)^{5/3} \left[\frac{1 + \frac{73}{24}e^2 + \frac{37}{96}e^4}{(1-e^2)^{7/2}} \right] \quad [2.116]$$

where \mathcal{M} is the chirp mass [1.6]. Thus the orbital frequency ω_0 increases, and GWs are emitted with increasing frequencies. The amplitude also increases, since as we can see from [2.88] it is inversely proportional to p .

To have the quantitative behaviour of the waveforms $h_{+, \times}$, we also need the time evolution of ψ . This is obtained solving [2.71] with $e(t)$ and $p(t)$ the solutions of Eqs. [2.114] and [2.112], and can be done numerically. The result is plotted in Fig. 2.9, for 4 different initial values of the eccentricity $e = 0, 0.3, 0.5$ and 0.7 , for the plus polarization in the z direction [2.86]. The upper waveform in Fig. 2.9 is for circular orbits. The increasing amplitude and frequency of the GWs is clearly visible and will be quantified in the discussion below. The waveform diverges when p reaches zero, though clearly this is beyond the regime of applicability of the quadrupole approximation which assumes $|\vec{v}|/c \ll 1$. Since $|\vec{v}| \sim 1/\sqrt{p}$ this is clearly violated as $p \rightarrow 0$. That is the reason why, in Section 1.1.4.2, we invoked the ISCO as a possible minimum distance, which then defined a merger frequency through Eq. [1.8]. The elliptic orbits are not symmetric, and the amplitude is maximal at the periastron, where we also have the maximum emitted power. The different plot also show that the waveform diverges at earlier and earlier times as e increases, in agreement with the fact that the emitted powers increase with e .

To get also a geometric intuition about the dynamics of the system during evolution, we plot in figure 2.10 the orbits for $e = 0$ and $e = 0.3$, showing the decrease in orbital radius and eccentricity.

2.4.5. Frequency content of elliptical orbits

As a side result, let us come back to elliptic orbits, without dissipation, and discuss how one can study the frequency content of the waves emitted by computing the contributions at each frequency to the total power [2.90]. To that end, we expand the dynamics into Fourier modes multiples of the fundamental




Figure 2.9: Four waveforms, in the lowest order PN expansion, with initial values of eccentricity given by $e = 0, 0.3, 0.5$ and 0.7 . Most GW power is emitted near the pericenter where the orbital velocity is the largest. Also since more GW radiation is emitted as e increases, the merger occurs earlier.




Figure 2.10: Corresponding to Fig. 2.9, the orbits $r(t)$ given by Eq. [2.68] for $e = 0$ and $e = 0.3$

harmonics given by the orbital frequency $\omega_0 = \sqrt{Gm/p^3}$. So for instance we write

$$I_{ab}(t) = \nu m p^2 \left\{ \tilde{A}_{ab}^{(0)} + \sum_{n=1}^{\infty} \left[\tilde{A}_{ab}^{(n)} \cos(n\omega_0 t) + \tilde{B}_{ab}^{(n)} \sin(n\omega_0 t) \right] \right\}, \quad [2.117]$$

with the dimension-full factor in front so to have dimension-less Fourier components

$$\tilde{A}_{ab}^{(0)} = \frac{1}{\nu m p^2 T} \int_0^T dt I_{ab}(t), \quad [2.118]$$

$$\tilde{A}_{ab}^{(n)} = \frac{2}{\nu m p^2 T} \int_0^T dt I_{ab}(t) \cos(\omega_0 n t), \quad \tilde{B}_{ab}^{(n)} = \frac{2}{\nu m p^2 T} \int_0^T dt I_{ab}(t) \sin(\omega_0 n t)$$

Recall that I_{ab} is given by [2.81] with \vec{r} parametrized as in [2.68]. To solve this integral one then needs the explicit time dependence of r and ψ , namely the solution of the (unperturbed) Keplerian orbit of ellipticity e . Details can be found in e.g. (Maggiore 2007), and one finds

$$0 = \tilde{B}_{11}^{(n)} = \tilde{A}_{12}^{(0)} = \tilde{A}_{12}^{(n)} = \tilde{B}_{22}^{(n)}, \quad \tilde{A}_{11}^{(0)} = \frac{1+4e^2}{2}, \quad \tilde{A}_{22}^{(0)} = \frac{1}{2} [2.120]$$

and

$$\begin{aligned}\tilde{A}_{11}^{(n)} &= \frac{1}{(1-e^2)^2} \cdot \frac{1}{n} [J_{n-2}(ne) - J_{n+2}(ne) - 2e(J_{n-1}(ne) - J_{n+1}(ne))] \\ \tilde{B}_{12}^{(n)} &= \frac{1}{(1-e^2)^{3/2}} \cdot \frac{1}{n} [J_{n+2}(ne) + J_{n-2}(ne) - e(J_{n+1}(ne) + J_{n-1}(ne))] \\ \tilde{A}_{22}^{(n)} &= \frac{1}{(1-e^2)} \cdot \frac{1}{n} [J_{n+2}(ne) - J_{n-2}(ne)],\end{aligned}\quad [2.123]$$

in terms of the Bessel's functions J . Taking three time derivatives leads to factors of $n^3\omega_0^3$. Substituting into [2.90] and averaging over one period to calculate $\langle P_{\text{GW}} \rangle$ leads to terms such as $\langle \sin n\omega_0 t \sin m\omega_0 t \rangle \sim \delta_{mn}$ meaning that the different harmonics do not interfere. In conclusion one finds

$$\langle P_{\text{GW}} \rangle = \sum_{n=1}^{\infty} \langle P_n \rangle \quad [2.124]$$

where

$$\langle P_n \rangle = P_{\text{GW}}^{e=0} \cdot \frac{n^6}{96} (1-e^2)^4 \left[\left(\tilde{A}_{11}^{(n)} \right)^2 + \left(\tilde{B}_{12}^{(n)} \right)^2 + 3 \left(\tilde{A}_{22}^{(n)} \right)^2 - \tilde{A}_{11}^{(n)} \tilde{B}_{12}^{(n)} \right], \quad [2.125]$$

where $P_{\text{GW}}^{e=0}$ is given in Eq. [2.94]. The $\langle P_n \rangle$ are plotted in figure 2.11 for different values of e . For circular orbits all power is emitted via the $n = 2$ mode. The power is the more and more distributed among other harmonics as e increases, and the frequency at which the maximum power is radiated also increases with e .

2.4.6. Hyperbolic orbits

While all GW detections to date are from *bound* elliptical/circular CBCs with $e < 1$, many other potential GW sources exist, for instance non-spherical spinning NSs and asymmetric core collapse Supernovae. In this brief subsection we discuss another possible source, namely *unbound binary systems on hyperbolic orbits*. That is, we consider cases in which the eccentricity $e > 1$ see Eq. [2.72] and Fig. 2.12.

Hyperbolic orbits are interesting not only because unbound orbits are expected to exist in nature (and hence the waveform for such events is and will be searched for by GW detectors (Agazie *et al.* 2024 ; Gasparotto *et al.* 2023 ;




Figure 2.11: Plot of $\langle P_n \rangle / P_{\text{GW}}^{e=0}$ for different values of $e = 0$ (red), $e = 0.2$ (blue), $e = 0.4$ (cyan) $e = 0.6$ (green).

Goncharov *et al.* 2024 ; Inchauspé *et al.* 2024)), but also because this simple system provides a first example of a *gravitational wave memory effect*. There are many different kinds of memory effects (see e.g. (Favata 2010)), the simplest of which is the linear memory effect which occurs already at the lowest order in the PN expansion and which is illustrated by hyperbolic orbits. Memory effects occur when there is a permanent change $\Delta h_{ab}^{\text{TT}}$ in the gravitational waveform, and thus leads to permanent displacement ΔL of the arms of GW detector for example, see Eq. [2.4].

2.4.6.1. Linear memory effect and low-frequency GWs

The two GW polarisations for hyperbolic orbits are shown in Fig. 2.8 for $\vec{N} = (0, 0, 1)$. The waveform is not periodic but rather burst-like, and h_{\times} has a non-zero variation between $t = \pm\infty$: this is the linear memory effect. More generally, the variation of the metric perturbation between $t = \pm\infty$ is given by

$$\Delta h_{cd}^{\text{TT}} = \int_{-\infty}^{\infty} dt \dot{h}_{cd}^{\text{TT}}(t) = h_{cd}^{\text{TT},+} - h_{cd}^{\text{TT},-}, \quad [2.126]$$




Figure 2.12: Sketch of a hyperbolic orbit in the CM, showing the angles $\psi_{\pm} = \pm \arccos(-1/e)$. The scattering angle is $\psi_+ - \psi_- - \pi = 2 \arccos(-1/e) - \pi = 2 \arcsin(1/e)$.

where we introduced the shorthand notation $h_{cd}^{\text{TT},\pm}$ for the evaluation at $t = \pm\infty$. It follows from [2.44] that there will be a linear memory effect $\Delta h_{ij}^{\text{TT}} \neq 0$ when there is a *net change in the second time derivatives of quadrupole moments* of the system. This is precisely the case for hyperbolic orbits, since at initial and final times the accelerations are zero, but the velocity vectors are different. At lowest PN order,

$$\begin{aligned} \Delta h_{cd}^{\text{TT}} &= \frac{2G}{c^4 R} P^{\text{TT}ab}_{cd}(\vec{N})(\ddot{I}_{ab}^+ - \ddot{I}_{ab}^-) \\ &= \frac{4G}{c^4 R} P^{\text{TT}ab}_{cd}(\vec{N}) \sum_{i=1}^2 m_i v_a^i v_b^i \Big|_{-\infty}^{\infty} \end{aligned} \quad [2.127]$$

$$= \frac{4G\mu}{c^4 R} v_{\infty}^2 P^{\text{TT}ab}_{cd}(\vec{N}) n_a n_b \Big|_{-\infty}^{\infty}, \quad [2.128]$$

where v_∞ is given in [2.80] and the unit-norm n_a at initial and final times refers to incoming and outgoing directions. From this one can compute the polarizations, for instance in the \hat{z} direction,

$$\Delta h_+(\hat{z}) = \frac{16G\mu v_\infty^2}{c^4 R} \frac{1-e^2}{e^4}, \quad \Delta h_\times(\hat{z}) = \frac{8G\mu v_\infty^2}{c^4 R} \frac{(e^2-2)\sqrt{e^2-1}}{e^4}. \quad [2.129]$$

See (Braginsky and Thorne 1987) for more details.

The time-scale over which the GW signal varies in Fig 2.8, namely the burst time-scale, is determined by the inverse of the characteristic frequency ω_c^{-1} given in Eq. [2.79]. Up to factors of eccentricity, this also determines the characteristic frequency scale of the emitted GWs on hyperbolic orbits. Indeed, contrary to the case of periodic elliptical orbits discussed in section 2.4.5, GWs of all continuous frequencies are emitted, and one can determine the GW power as a function of frequency by now Fourier transforming the power emitted. Using the convention $\tilde{I}_{ab}(\omega) = \int dt I_{ab}(t) e^{-i\omega t}$, as well as the quadrupole approximation, the total energy emitted in GWs is

$$\begin{aligned} E_{\text{GW}} &= \frac{G}{5c^5} \int_{-\infty}^{\infty} dt (\ddot{I}_{ab})^2 = \frac{G}{5\pi c^5} \int_0^{\infty} d\omega \omega^6 |I_{ab}(\omega)|^2 \\ &\equiv \int_0^{\infty} d\omega \mathcal{P}_{\text{GW}}(\omega). \end{aligned}$$

The emitted power in GWs is thus

$$\mathcal{P}_{\text{GW}}(\omega) = \frac{G}{5\pi c^5} [\omega^3 \tilde{I}_{ab}(\omega)] [\omega^3 \tilde{I}_{ab}^*(\omega)]. \quad [2.130]$$

Direct calculation analogous to that of section 2.4.5 (see e.g. (Brax and Steer 2024)) shows that $\mathcal{P}_{\text{GW}}(\omega)$ is peaked at a value fixed by ω_c but which increases with e , see figure 2.13. Notice that $\mathcal{P}_{\text{GW}}(0) \neq 0$. This is due to the linear memory effect: indeed Eq. [2.126], written in Fourier space reads

$$\Delta h_{cd}^{\text{TT}} = -i \frac{2G}{c^4 R} P^{\text{TT}ab}(\vec{N}) [\omega^3 \tilde{I}_{ab}(\omega)]|_{\omega=0}, \quad [2.131]$$

thus a non-vanishing linear memory effect implies $\mathcal{P}_{\text{GW}}(0) \neq 0$.




Figure 2.13: The energy spectrum $\mathcal{P}_{\text{GW}}(\tilde{\omega})/\mathcal{P}_{\text{GW}}(0)$ where $\tilde{\omega} = \omega/\omega_c$ for different eccentricities. Figure from (Braz and Steer 2024).

2.4.6.2. Capture due to GW emission

One can estimate the energy emitted in GWs between $\psi_- < \psi < 0$ by calculating

$$\Delta E_{\text{GW}} = \int_{\psi_-}^0 d\psi \frac{1}{\dot{\psi}} P_{\text{GW}}(\psi) \quad [2.132]$$

where $P_{\text{GW}}(\psi)$ is given in Eq. [2.92]. Let us consider an orbit which is only slightly unbound, thus $e = 1 + \epsilon$ with $0 < \epsilon \ll 1$ so that the orbital energy is $E \simeq \nu Gm^2 \epsilon/p$, see Eq. [2.72]. Then it is straightforward to determine ΔE_{GW} to lowest order in ϵ , leading to

$$\Delta E_{\text{GW}} = \frac{85}{3} mc^2 \nu^2 \left(\frac{Gm}{c^2 p} \right)^{7/2} \pi + \mathcal{O}(1). \quad [2.133]$$

Note that here we assumed p constant, that is we have neglected the back-reaction of the emitted GR on p : this is a reasonable approximation for this

individual burst process since p changes on a time-scale $\sim (c/v)^{-5}\omega_c^{-1}$ as discussed previously. The orbital energy $E \sim \nu Gm^2\epsilon/p$ will thus be reduced by ΔE_{GW} , and if

$$E - \Delta E_{\text{GW}} < 0 \quad [2.134]$$

then the GW energy loss will convert the hyperbolic orbit into a bound orbit before the pericenter. This can be rewritten as the condition

$$\epsilon < \frac{85\nu}{3} \left(\frac{Gm}{c^2 p} \right)^{5/2} \pi. \quad [2.135]$$

Instead of parametrising the orbit in terms of (p, ϵ) , for such scattering trajectories it is more convenient to work with (b, v_∞) where b is the impact parameter and v_∞ is the orbital velocity at infinite separation. From Eq. [2.80], $v_\infty^2 \simeq 2(Gm/p)\epsilon$ and substituting ϵ from Eq. [2.135] gives

$$\left(\frac{Gm}{c^2 p} \right)^{7/2} > \frac{3}{170\pi\nu} \left(\frac{v_\infty}{c} \right)^2. \quad [2.136]$$

Furthermore, basic trigonometry gives

$$b = \frac{p}{e \sin \psi_-} = \frac{p}{\sqrt{e^2 - 1}} \simeq \frac{p}{\sqrt{2\epsilon}} = \frac{\sqrt{Gmp}}{v_\infty}. \quad [2.137]$$

Saturating the bound in [2.136] identifies the impact parameter b_{capture} for which capture occurs. The corresponding capture cross-section $\sigma_{\text{GW}} := \pi b_{\text{capture}}^2$ is thus given by

$$\sigma_{\text{GW}} = \pi \left(\frac{170\pi\nu}{3} \right)^{2/7} \left(\frac{Gm}{c^2} \right)^2 \left(\frac{c}{v_\infty} \right)^{18/7}. \quad [2.138]$$

Such a process could play an important role for instance in dense star clusters and galactic nuclei, see e.g. (Capozziello *et al.* 2008 ; O'Leary *et al.* 2009 ; Hoang *et al.* 2020).

2.5. GWs in curved space-time, cosmology

In the previous sections we discussed GWs in Minkowski space; our aim is now to generalise the results presented there to a cosmological space-time. There is now a further scale of interest other than the characteristic size d of the source and the GW wavelength λ_{GW} , namely the *cosmological horizon*. The results of the previous sections are valid in the so-called the *local wavezone* of the source, namely at distances scales R which are large compared to the GW wavelength but *small* compared to the cosmological horizon, $d \ll \lambda_{\text{GW}} \ll R \ll \text{horizon}$. We now aim to extend them to cosmological scales: we will see that the expansion of the universe dampen the GW amplitude, and redshift frequencies and masses.

2.5.1. General background metric

2.5.1.1. Linearised equations

The linearized equations around an arbitrary background [1.36] are given explicitly in Appendix 1.4. We can simplify them again using the de Donder gauge [1.75], generalised to curved space-time as¹⁰

$$\bar{\nabla}^\mu \bar{h}_{\mu\nu} = 0, \quad \bar{h}_{\mu\nu} := h_{\mu\nu} - \frac{1}{2} \bar{g}_{\mu\nu} h. \quad [2.139]$$

In this gauge, we have (see Eq. [??] for details)

$$\bar{\square} \bar{h}_{\alpha\beta} + 2\bar{R}^\mu{}_{\alpha\nu\beta} \bar{h}^\nu{}_\mu + 2\bar{G}_{\mu(\alpha} \bar{h}_{\beta)}{}^\nu - \bar{g}_{\alpha\beta} \bar{R}_{\mu\nu} \bar{h}^{\mu\nu} = -\frac{16\pi G}{c^4} T_{\alpha\beta}^{(1)}. \quad [2.140]$$

Here $\bar{R}_{\mu\alpha\nu\beta}$ is the Riemann curvature tensor of the background space-time, and $\bar{R}_{\mu\nu}$ the Ricci-tensor. The terms containing the background curvature depend on the ratio of λ_{GW} relative to the scale of variation of the background metric. In the following we will assume that they are negligible. The justification for this is that below we will restrict attention to the homogeneous and isotropic FLRW metric, where the scale of variation of the background is the cosmological horizon, which is much greater than λ_{GW} .¹¹ The right-hand side has generically

10. We are using a bar to indicate both the background metric, and the trace-reversed perturbation. We believe that no confusion should arise, because of the different places where these quantities enter.

11. In a perturbed FLRW metric, there could be local bumps in the curvature of scale similar to λ_{GW} : we do not consider this case here. See also (Isaacson 1968a) for further details.

two contributions, one from the metric perturbations, and one from the matter perturbation. In cosmology, it is often useful to approximate matter using a perfect fluid description, for which

$$T_{\mu\nu} = (\rho + p)u_\mu u_\nu + pg_{\mu\nu}, \quad u^2 = -c^2, \quad [2.141]$$

where ρ and p are respectively the energy and pressure densities of the fluid. In this case, $T^\mu{}_\nu$ is metric-independent, and this allows us to separate the metric and matter perturbations using the identity $T_{\mu\nu}^{(1)} = \bar{g}_{\mu\rho}T^{(1)\rho}{}_\nu + h_\mu{}^\rho\bar{T}_{\rho\nu}$. On-shell, the second term is proportional to the background curvature and therefore can be neglected again by the same argument as given above. We conclude that in the absence of matter perturbations, [2.140] reduces to

$$\bar{\square}\bar{h}_{\alpha\beta} \simeq 0. \quad [2.142]$$

2.5.1.2. WKB approximation

For a general space-time, the solution of Eq. [2.142] can be obtained in the WKB approximation. The underlying physical assumption is that the amplitude of the wave is slowly varying with respect to the frequency of the wave, and hence we write

$$\bar{h}_{\mu\nu}(x) = \sum_p \Re \left[A_p(x) \epsilon_{\mu\nu}^p(x) e^{iS(x)/\delta} \right] \quad [2.143]$$

where the sum is over polarisations p with polarisation tensor $\epsilon_{\mu\nu}$ satisfying $\epsilon_{\mu\nu}\epsilon^{\mu\nu} = 1$ (recall that to this leading order in h , indices are raised and lowered with the background metric) and A is the corresponding amplitude. The parameter $\delta \rightarrow 0$, and we define

$$k^\mu = \frac{\partial_\mu S}{\delta}. \quad [2.144]$$

Now substituting [2.143] the de Donder condition Eq. 2.139 becomes (dropping the p sum for simplicity)

$$\bar{\nabla}^\mu \bar{h}_{\mu\nu} = \left[\bar{\nabla}^\mu (A\epsilon_{\mu\nu}) + iA\epsilon_{\mu\nu} \frac{\partial_\mu S}{\delta} \right] e^{iS(x)/\delta} = 0, \quad [2.145]$$

which to leading order in δ implies

$$k^\mu \epsilon_{\mu\nu} = 0 \quad [2.146]$$

and hence that $\epsilon_{\mu\nu}$ is the transverse polarisation tensor. Substitution into the equation of motion Eq. [2.142] leads to terms in δ^{-2} , δ^{-1} which are, respectively

$$\bar{g}^{\mu\nu} k_\nu k_\mu = 0, \quad [2.147]$$

$$2\bar{\nabla}_\mu (A\epsilon_{\alpha\beta})k^\mu + (A\epsilon_{\alpha\beta})\bar{\nabla}_\mu k^\mu = 0. \quad [2.148]$$

The first equation, the Eikonal equation (geometric optics limit), implies that GWs are massless with dispersion relation $\omega^2 = \vec{k}^2$ and propagate on null geodesics. Contracting the second equation [2.148] with $\epsilon^{\alpha\beta}$, and using that $\epsilon_{\mu\nu}\epsilon^{\mu\nu} = 1$ so that $(\bar{\nabla}^\alpha \epsilon_{\mu\nu})\epsilon^{\mu\nu} = 0$, leads to

$$2(\bar{\nabla}_\mu A)k^\mu + A\bar{\nabla}_\mu k^\mu = 0 \quad \Rightarrow \quad \bar{\nabla}_\mu (A^2 k^\mu) = 0. \quad [2.149]$$

This gives the decay of the GW amplitude A along the null geodesics. Finally, substituted back Eq. [2.149] into Eq. [2.148] gives

$$k^\mu (\bar{\nabla}_\mu \epsilon_{\alpha\beta}) = 0 \quad [2.150]$$

which implies that the polarization tensor $\epsilon_{\alpha\beta}$ of the GW is parallel propagated along the null geodesics.

To summarise, the solution of

$$\bar{g}^{\mu\nu} \bar{\nabla}_\mu \bar{\nabla}_\nu \bar{h}_{\alpha\beta} \simeq 0 \quad [2.151]$$

in the WBK approximation, and in the Lorenz gauge, is $\bar{h}_{\mu\nu}(x) = \Re [A(x)\epsilon_{\mu\nu}(x)e^{ik_\mu x^\mu}]$ with

$$\bar{g}^{\mu\nu} k_\nu k_\mu = 0 \quad [2.152]$$

$$k^\mu \epsilon_{\mu\nu} = 0 \quad [2.153]$$

$$\bar{\nabla}_\mu (A^2 k^\mu) = 0 \quad [2.154]$$

$$k^\mu (\bar{\nabla}_\mu \epsilon_{\alpha\beta}) = 0 \quad [2.155]$$

We now consider these equations in a FLRW metric.

2.5.2. FLRW metric: background

The flat Friedmann-Lemaitre-Robertson-Walker (FLRW) metric is

$$d\bar{s}^2 = -dt^2 + a^2(t)d\vec{x}^2 = -dt^2 + a^2(t)(dr^2 + r^2d\Omega^2) \quad [2.156]$$

where \vec{x} are comoving coordinates and $a(t)$ is the scale factor (normalised such that today, at $t = t_0$, $a(t_0) \equiv a_0 = 1$). For a perfect fluid source, Einsteins equations Eq. [1.16] reduce to the Friedmann equations

$$H^2 = \frac{8\pi G}{3}\rho + \frac{\Lambda}{3} \quad [2.157]$$

$$\frac{\ddot{a}}{a} = -\frac{8\pi G}{6}(\rho + 3P) + \frac{\Lambda}{3} \quad [2.158]$$

where ρ and P are respectively the energy density and pressure of the perfect fluid, and $H = \frac{\dot{a}}{a}$ is the Hubble parameter, whose value today is the Hubble constant H_0 . These two equations imply the conservation equation $\dot{\rho} + 3H(\rho + P) = 0 = \nabla_\nu T^{\mu\nu}$. In terms of conformal time η defined by $d\eta = dt/a(t)$, the metric in Eq. [2.156] is conformally related to the Minkowski metric

$$d\bar{s}^2 \equiv \bar{g}_{\mu\nu}dx^\mu dx^\nu = a^2(\eta)[-d\eta^2 + dr^2 + r^2d\Omega^2]. \quad [2.159]$$

Consider now a source (of photons or GWs) at fixed radial position $r = 0$, and an observer at r_o . On a (η, r) space-time diagram null radial geodesics propagate at 45 degrees. If two null geodesics are emitted at a conformal time interval $\delta\eta_s$ by the source, then they arrive at the observer with the same conformal time interval $\delta\eta_o = \delta\eta_s$. This implies the standard time-dilation relation

$$dt_o = \frac{a(t_o)}{a(t_s)} dt_s \equiv (1+z)dt_s \quad [2.160]$$

where z is the redshift of the source (the observer is at $t_o = t_0$), and equivalently that the emitted (or ‘source’) frequency f_s is related to the observed frequency f_o by

$$f_o = \frac{f_s}{1+z}. \quad [2.161]$$

The radial comoving distance R to an event with redshift z is given by solving $ds^2 = 0$, thus $dr = dt/a(t)$, leading to

$$R = \int dr = \int \frac{dt}{a(t)} = \int \frac{1}{a} \frac{dt}{da} da dz = \int_0^z dz' \frac{1}{H(z')} \quad [2.162]$$

where $H(z) = H_0 E(z)$ is the Hubble parameter expressed in terms of redshift, and from the Friedmann equation

$$E(z) = \sqrt{\Omega_m(1+z)^3 + \Omega_r(1+z)^4 + \Omega_\Lambda} \quad [2.163]$$

where $\Omega_{r,m} = \frac{8\pi G\rho_{r,m}}{3H_0^2}$, $\Omega_\Lambda = \frac{\Lambda}{3H_0^2}$ and $\Omega_r + \Omega_m + \Omega_\Lambda = 1$.

A crucial quantity is the luminosity distance $d_L(z)$. This relates the EM luminosity of the source and the luminosity measured by the observer. In the flat FRWL metric [2.156] it is given by

$$d_L(z) = a(t_o)(1+z)R = (1+z) \int_0^z dz' \frac{1}{H(z')}. \quad [2.164]$$

As we will see, this same distance scale determines the GW amplitude in an expanding universe.

2.5.3. FLRW metric: gravitational waveforms

Consider a GW propagating radially outwards from the source at $r = 0$ and redshift z_s with $k_\mu = \omega(1, -1, 0, 0)$. From Eq. [2.149] it is possible to determine how the GW amplitude decreases along the null GW geodesic. In a FLRW metric (in conformal time) Eq. [2.149] becomes

$$\partial_\nu(\sqrt{-\bar{g}}A^2k^\nu) = 0 = \partial_\nu(a(\eta)^2A^2r^2k^\nu). \quad [2.165]$$

Thus $A(\eta, r)a(\eta)r$ remains conserved during the propagation, and

$$A(\eta, r) = \frac{\text{const}}{a(\eta)r} \Big|_{\eta-r=\text{const}} \quad [2.166]$$

The constant is fixed by the known amplitude of the wave in the wave-zone approximation, close to the source, where the Minkowski results are valid. Then the remainder of the solution $\bar{h}_{\mu\nu}$ is obtained by parallel transporting this solution from the source to the observer. We now carry out these steps.

Before doing so we note that in a flat FLRW universe and focusing on the *spatial* TT components only, then in fact [2.140] reduces to

$$\bar{g}_{\mu\nu}\bar{\nabla}^\mu\bar{\nabla}^\nu\bar{h}_{ij}^{\text{TT}} = 0 = h_{ij}^{\text{TT}''} + 2\mathcal{H}h_{ij}^{\text{TT}'} + \partial_k\partial^k h_{ij}^{\text{TT}} \quad [2.167]$$

since the spatial components of the Riemann and Ricci tensors vanish identically. From here the scaling of the amplitude of GWs as $1/a(\eta)$ is also immediate.

We now consider a compact binary system on circular orbits, as discussed in section 2.4.3, but this time in a FLRW background. In the wave-zone approximation and at a physical distance $R = a(t_s)r$ from the source as measured by time t_s of the source clock, the plus and cross polarisations of the GW are given in [2.105] and [2.106]. Focusing on the cross polarisation,

$$h_{\times}(t_s, \iota) = \frac{4}{R} \left(\frac{GM}{c^2} \right)^{5/3} \left(\frac{\pi f_s(t_s^{\text{ret}})}{c} \right)^{2/3} \cos \iota \sin(2\Phi_s(t_s^{\text{ret}})) \quad [2.168]$$

where $t_s^{\text{ret}} = t_s - t_c$ is the time to coalescence at t_c and \mathcal{M} is the chirp mass [1.6]. The time dependence of the frequency is given in Eq. [2.100] namely

$$\frac{df_s}{dt_s} = \frac{96}{5} \pi^{8/3} \left(\frac{GM}{c^3} \right)^{5/3} f_s^{11/3} \quad [2.169]$$

leading to

$$f_s(t_s^{\text{ret}}) = \frac{1}{\pi} \left(\frac{5}{256 t_s^{\text{ret}}} \right)^{3/8} \left(\frac{GM}{c^3} \right)^{-5/8}, \quad [2.170]$$

so that the phase dependence is

$$\Phi_s(t_s) = \Phi_c + 2\pi \int_{t_c}^{t_s} dt'_s f_s(t'_s) = -2 \left(\frac{t_s^{\text{ret}} c^3}{5GM} \right)^{5/8} + \Phi_c. \quad [2.171]$$

We now parallel transport this solution [2.168] along a null geodesic to the observer. Along the geodesic the GW phase remains constant because the time dilation effects cancel the redshifting of the frequency. Thus at the observer whose clock measures $dt = dt_s(1+z)$, the observed GW frequency $f = f_s/(1+z)$ leading to $\Phi(t) = \Phi_s(t_s)$. However, at the observer, the GW amplitude is changed. From Eq. [2.166], and using [2.168]

$$h_{\times}(t, \iota) = \frac{4}{a(t)R} \left(\frac{GM}{c^2} \right)^{5/3} \left[\frac{\pi}{c} f(t^{\text{ret}})(1+z) \right]^{2/3} \cos \iota \sin(2\Phi(t^{\text{ret}})) \quad [2.172]$$

where we have included the redshifting of frequency. Let us now define the *redshifted chirp mass*

$$\mathcal{M}_z = (1+z)\mathcal{M} \quad [2.173]$$

Then [2.172] becomes

$$\begin{aligned} h_{\times}(t, \iota) &= \frac{4}{a(t)R(1+z)} \left(\frac{G\mathcal{M}_z}{c^2} \right)^{5/3} \left(\frac{\pi f(t^{\text{ret}})}{c} \right)^{2/3} \cos \iota \sin(2\Phi(t^{\text{ret}})) \\ &= \frac{4}{d_L(z)} \left(\frac{G\mathcal{M}_z}{c^2} \right)^{5/3} \left(\frac{\pi f(t^{\text{ret}})}{c} \right)^{2/3} \cos \iota \sin(2\Phi(t_o^{\text{ret}})) \end{aligned} \quad [2.175]$$

where in the second line we have used Eq. [2.164] defining the luminosity distance to the source (today $a(t_o) = 1$). The dependence of the observed frequency on time t is obtained by $f_s = (1+z)f$ into Eq. [2.169]:

$$(1+z) \frac{d[f(1+z)]}{dt} = \frac{96}{5} \pi^{8/3} \left(\frac{G\mathcal{M}}{c^3} \right)^{5/3} f^{11/3} (1+z)^{11/3}. \quad [2.176]$$

Assuming that changes in z are negligible during the observation time, then z can be taken as constant¹² leading to

$$\frac{df}{dt} = \frac{96}{5} \pi^{8/3} \left(\frac{G\mathcal{M}_z}{c^3} \right)^{5/3} f^{11/3}, \quad [2.177]$$

namely the GW phase depends on the *redshifted chirp mass*,

$$\Phi(t^{\text{ret}}) = -2 \left(\frac{t^{\text{ret}} c^3}{5G\mathcal{M}_z} \right)^{5/8} + \Phi_c. \quad [2.178]$$

To summarize, the GW frequency depends on the redshifted chirp mass \mathcal{M}_z which is therefore determined by measurements of the phase of an inspiral signal. The GW amplitude depends on both \mathcal{M}_z and $d_L(z)$. Given that the former is determined from the phase, measurements of the amplitude of the signal determine $d_L(z)$. Generally speaking therefore, GW observations from

¹². See (Bonvin *et al.* 2017) for a discussion of where this assumption may lead to biases

individual CBC events determine the luminosity distance $d_L(z)$ and the so-called ‘redshifted’ masses,

$$m_{1,2}^{\text{detected}} = (1+z)m_{1,2} \quad [2.179]$$

which are related to the ‘source’ masses $m_{1,2}$ by the same factor of $1+z$ as in Eq. [2.173].¹³

It is important to note that while the redshifted masses and the luminosity distance can be deduced from the waveform, the redshift z of the individual CBC remains *undetermined*. To deduce this redshift, one possibility is to assume a cosmological model — such as Λ CDM, with given values of H_0 , Ω_m etc, say from the Planck observations — so that z can be read off the luminosity distance, see Eq. [2.164]. Another possibility is to find, and measure, effects in the waveform that also depend on the *source-frame* masses as well as the redshifted masses. For binary black holes there are no such effects (even to higher orders in the PN expansion). For binaries including neutron stars, tidal effects with these properties enter, however at higher PN order (see e.g. (Del Pozzo *et al.* 2017 ; Messenger and Read 2012) and references within). Furthermore, the features depend on the — uncertain — equation of state of the nuclear matter making up the neutron star. In conclusion, to lowest order in the PN expansion for CBCs, *it is not possible to determine the redshift z of the source from GW observations alone*: there is a perfect degeneracy between source masses, redshift, as well as spins.

Note that *if* z is determined, then from the detected masses $m_{1,2}^{\text{detected}}$ one can obtain the value of the source masses $m_{1,2}$ via Eq. [2.179]. This was done in (Abbott *et al.* 2016), for example, assuming Λ CDM cosmological model with Planck values of H_0 , Ω_m to find z : this is how the source-frame values of the two black hole masses was determined. However, the values of the cosmological parameters H_0 etc are in fact *not* known precisely, and a source of tension in cosmology today, see e.g. (Di Valentino *et al.* 2021 ; Verde *et al.* 2024). For these reasons, it can be interesting to use GW observations in a different way, namely as a new observable with which to measure cosmological parameters.

13. To determine each redshifted mass individually, rather than in the combination of the chirp mass, requires the waveform beyond the lowest order quadrupolar form discussed here

2.5.4. Measuring cosmological parameters with GWs: outline of GW cosmology

In this last brief subsection we outline how GW observations can be used to measure cosmological parameters.

The luminosity distance $d_L(z)$ in Λ CDM is given in Eq. [2.164], and is a function of cosmological parameters such as H_0 and Ω_m . At low redshifts $z \ll 1$, the domain of the O3 measurements of the LVK collaboration (Abbott *et al.* 2023), only H_0 enters since Eq. [2.164] reduces to

$$d_L \sim \frac{cz}{H_0} \quad \Rightarrow \quad H_0 \sim \frac{cz}{d_L}. \quad [2.180]$$

Clearly, in order to *measure* H_0 , not only is d_L required (and obtained from GW observations, as we have discussed), but the redshift z of the source is also needed. However, as mentioned above, this cannot be determined from GW observations: extra non-gravitational information is necessary to determine z . Such information could, for example, be electro-magnetic (EM).

Indeed, the most straightforward way to determine z is use EM observations to uniquely identify the “host galaxy” of the GW signal, namely the galaxy in which GW event occurred. This was possible for GW event GW170817 which occurred on August 17th 2017 and which corresponded to the the merger of two neutron stars (Abbott *et al.* 2017b). This GW signal was observed the two LIGO and Virgo GWdetectors, and 1.7s following the GW merger, EM observers around the globe observed a subsequent gamma-ray burst as well as multiple EM signals in different frequency bands. Using this EM data it was possible to determine the host galaxy, namely NGC 4993, a galaxy in the Hydra constellation. This constellation is receding from us with a velocity $cz = 3327 \pm 72$ km/s, due to the expansion of the universe. Combining this with the distance $d_L = 43.8_{-6.9}^{+2.9}$ Mpc inferred from the GW signal led, using Eq. [2.180], to an estimated value for the Hubble constant of $H_0 = 70_{-8}^{+12}$ km/s/Mpc (Abbott *et al.* 2017a). This result, using one GW event only, is consistent with other measurements but is of course less accurate because of its larger error bars. Its interest is that it shows that the idea works. The errors would be reduced (with a $\sim 1/\sqrt{N}$ scaling) if N other measurements of this kind existed, but unfortunately GW170817 was an extremely rare event as since then no further GW events with associated EM counterparts (known as *standard sirens*) have been detected.

However, LVK has detected GWs from hundreds of BBHs and a few NS-BH, for each of which there is a measured d_L and $m_{1,2}^{\text{detected}}$ — but no EM counterpart. Even for these *dark sirens*, it is possible to obtain redshift information,

and therefore measure H_0 . Today, two pieces of information are used together to get a statistical redshift for GW events: (i) galaxy catalogues and (ii) astrophysical modelling of the formation channels of BBHs. We refer the reader to (Mastrogiovanni *et al.* 2023) as well as chapter 10 of the present book for a review of these methods and results.

2.6. Acknowledgements

We are grateful to Abhay Ashtekar, Tom Bertheas, Cédric Deffayet, Nathalie Deruelle, Eanna Flanagan, Alejandro Perez, Pierre Piovesan, Eric Poisson, Syksy Räsänen, Carlo Rovelli, Robert Wald for discussions on the topics of this review, as well as multiple Master 2 and doctoral students whom we have had the pleasure to teach. We also thank ChatGPT for technical support (not for the calculations though). This work was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Research, Innovation and Science. This work was also supported in part by CERN.

2.7. Appendix: Landau-Lifshitz approach

The Landau-Lifshitz formulation of Einstein's equations is a convenient approach to perturbative theory around Minkowski, and it is widely used by the community working on the PN and PM expansions. We provide here a brief description of the approach from the perspective of the main text, and refer the reader to (Blanchet 2006 ; Poisson and Will 2014) for more details. In the Landau-Lifshitz approach one uses a density-weighted inverse metric as fundamental variable,

$$\mathfrak{g}^{\mu\nu} := \sqrt{-g}g^{\mu\nu} = \bar{g}^{\mu\nu} - \bar{h}^{\mu\nu} + O(h^2). \quad [2.181]$$

The interest in doing so is that the quantity

$$H^{\alpha\mu\beta\nu} := \mathfrak{g}^{\alpha\beta}\mathfrak{g}^{\mu\nu} - \mathfrak{g}^{\alpha\nu}\mathfrak{g}^{\beta\mu} \quad [2.182]$$

is related to the Einstein tensor via ordinary derivatives. More precisely,

$$\partial_\mu\partial_\nu H^{\alpha\mu\beta\nu} = 2(-g)(G^{\alpha\beta} + \frac{8\pi G}{c^4}t_{LL}^{\alpha\beta}) \doteq \frac{16\pi G}{c^4}(-g)(T^{\alpha\beta} + t_{LL}^{\alpha\beta}), \quad [2.183]$$

where $t_{LL}^{\alpha\beta}$ is a pseudo-tensor of density-weight two, given explicitly by some lengthy expression in terms of first derivatives of the metric. Crucially, it is conserved on-shell $\partial_\alpha(-gt_{LL}^{\alpha\beta}) \doteq 0$. The dimension-full numerical factor in front of it is included in its definition for convenience when going on-shell in the last equality above.

This arrangement of Einstein's equations manifestly breaks covariance.¹⁴ Not only we have partial derivatives as opposed to covariant derivatives, but also tensor densities, aka pseudo-tensors, appearing. Its usefulness is limited to situations in which there are regions of spacetimes that are approximately flat, and where one can choose a Cartesian coordinate system, so that partial derivatives can be interpreted. This is precisely the case when studying a perturbative approximation around Minkowski. Within that context, the reformulation has two useful advantages.

The first is that it provides a prescription for a gravitational energy-momentum pseudo-tensor valid to all orders, given by $t_{LL}^{\alpha\beta}$. By analogy with the matter counterpart, one has a prescription to further split this quantity into contributions to energy, momentum and angular momentum. All these expressions are gauge-dependent; in particular, we have the usual problem that the pseudo-tensor can be made to vanish at any given point, using a local inertial frame. But the logic in this approach is that we assume to have a preferred coordinate system, the Cartesian ones of the fiducial flat metric, and that is the gauge we stick to. Furthermore, since the left-hand side is a total derivative, the total energy, momentum and angular momentum can be expressed as surface integrals. This provides a prescription for these quantities that can be evaluated in a region far from the sources, where one can safely assume that spacetime is approximately flat and use Cartesian coordinates, and computed to all order in perturbation theory.

The second advantage of the formulation is that it allows to set up an iteration scheme for the perturbative resolution of the field equations in a very practical way. To see that, we change variables to $\hat{h}^{\mu\nu} := \eta^{\mu\nu} - \mathbf{g}^{\mu\nu}$, where $\eta^{\mu\nu}$ is a fiducial background metric for which the coordinates are Cartesian. Then [2.183] are equivalent to

$$\square \hat{h}^{\mu\nu} \doteq -\frac{16\pi G}{c^4} \tau^{\mu\nu}, \quad \tau^{\mu\nu} := (-g)(T^{\mu\nu} + t_{LL}^{\mu\nu} + t_H^{\mu\nu} + t_{NH}^{\mu\nu}), \quad [2.184]$$

where \square is the flat spacetime d'Alembertian, and $t_H^{\mu\nu}$ are the remaining terms in the LHS of [2.183]. Specifically, $t_H^{\mu\nu}$ is quadratic in $h^{\mu\nu}$ and satisfies $\partial_\mu(-gt_H^{\mu\nu}) \equiv$

14. It is like rearranging the covariant geodesic equation [2.1] as in [2.2], where neither side of the equation is covariant by itself.

0, whereas $t_{\text{NH}}^{\mu\nu}$ does not but contains only terms that vanish in harmonic gauge. It follows that in the harmonic gauge,

$$\partial_\mu \tau^{\mu\nu} \triangleq 0, \quad [2.185]$$

and this equation includes the conservation of the matter energy-momentum tensor, namely the matter dynamics. We stress that [2.184] are the exact Einstein's equations, no approximation has been done yet. It is the introduction of a fiducial flat background and a fixed choice of Cartesian coordinates on it, that allows us to rewrite the exact equations in the form of a flat space-time wave equation with a complicated source satisfying the matter dynamics through [2.185]. In particular, all non-linearities are recasted on the RHS of the equations. To solve the equations perturbatively, we assume that h is small, and proceed iteratively as explained in [1.37]. The idea that makes this approach particularly convenient is to first solve the 'relaxed field equations' [2.184] alone, in harmonic gauge, and afterwards impose the gauge consistency condition [2.185] on the matter dynamics. Notice that in this scheme, the n -th iteration of the gravitational potentials $h_{\mu\nu}$ are sourced by matter fields satisfying the equations of motion that use the $(n-1)$ -th iteration of the gravitational field. See (Blanchet 2006 ; Poisson and Will 2014) for further details.

2.8. Bibliography

- Abbott, B. P. *et al.* (2016), Observation of Gravitational Waves from a Binary Black Hole Merger, *Phys. Rev. Lett.*, 116(6), 061102.
- Abbott, B. P. *et al.* (2017a), A gravitational-wave standard siren measurement of the Hubble constant, *Nature*, 551(7678), 85–88.
- Abbott, B. P. *et al.* (2017b), GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, *Phys. Rev. Lett.*, 119(16), 161101.
- Abbott, R. *et al.* (2023), Constraints on the Cosmic Expansion History from GWTC-3, *Astrophys. J.*, 949(2), 76.
- Agazie, G. *et al.* (2024), The NANOGrav 12.5 yr Data Set: Search for Gravitational Wave Memory, *Astrophys. J.*, 963(1), 61.
- Andersson, N. (2019), *Gravitational-Wave Astronomy*, Oxford Graduate Texts, Oxford University Press.
- Blanchet, L. (2006), Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries, *Living Rev. Rel.*, 9(1), 1–114.
- Bonvin, C., Caprini, C., Sturani, R., Tamanini, N. (2017), Effect of matter structure on the gravitational waveform, *Phys. Rev. D*, 95(4), 044029.
- Braginsky, V. B., Thorne, K. S. (1987), Gravitational-wave bursts with memory and experimental prospects, *Nature*, 327, 123–125.

- Brax, P., Steer, D. A. (2024), Scalar kicks and memory, *to appear*, .
- Burnett, G. A. (1989), The high frequency limit in general relativity, *Journal of Mathematical Physics*, 30(1), 90–96.
- Capozziello, S., De Laurentis, M., De Paolis, F., Ingrosso, G., Nucita, A. (2008), Gravitational waves from hyperbolic encounters, *Mod. Phys. Lett. A*, 23, 99–107.
- Cornish, N. J., Rubbo, L. J. (2003), The LISA response function, *Phys. Rev. D*, 67, 022001. [Erratum: *Phys. Rev. D* 67, 029905 (2003)].
- Cutler, C. *et al.* (1993), The Last three minutes: issues in gravitational wave measurements of coalescing compact binaries, *Phys. Rev. Lett.*, 70, 2984–2987.
- Damour, T. (1982), Gravitational Radiation And The Motion Of Compact Bodies, *in Les Houches Summer School on Gravitational Radiation*.
- Damour, T., Deruelle, N. (1985), General relativistic celestial mechanics of binary systems. i. the post-newtonian motion, *Annales de l'I.H.P. Physique théorique*, 43(1), 107–132.
- Del Pozzo, W., Li, T. G. F., Messenger, C. (2017), Cosmological inference using only gravitational wave observations of binary neutron stars, *Phys. Rev. D*, 95(4), 043502.
- Deruelle, N., Uzan, J.-P. (2018), *Relativity in Modern Physics*, Oxford University Press.
- DeWitt, B. S. (2011), *1971's Bryce DeWitt's Lectures on Gravitation*, vol. 826 of *Lecture Notes in Physics*, Springer, Berlin, Heidelberg.
- Di Valentino, E., Mena, O., Pan, S., Visinelli, L., Yang, W., Melchiorri, A., Mota, D. F., Riess, A. G., Silk, J. (2021), In the realm of the Hubble tension—a review of solutions, *Class. Quant. Grav.*, 38(15), 153001.
- Einstein, A. (1918), On gravitational waves, *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften (Berlin)*, pp. 154–167.
- Favata, M. (2010), The gravitational-wave memory effect, *Class. Quant. Grav.*, 27, 084036.
- Gasparotto, S., Vicente, R., Blas, D., Jenkins, A. C., Barausse, E. (2023), Can gravitational-wave memory help constrain binary black-hole parameters? A LISA case study, *Phys. Rev. D*, 107(12), 124033.
- Goldberger, W. D., Rothstein, I. Z. (2006), An Effective field theory of gravity for extended objects, *Phys. Rev. D*, 73, 104029.
- Goncharov, B., Donnay, L., Harms, J. (2024), Inferring Fundamental Spacetime Symmetries with Gravitational-Wave Memory: From LISA to the Einstein Telescope, *Phys. Rev. Lett.*, 132(24), 241401.

- Grant, A. M., Nichols, D. A. (2023), Outlook for detecting the gravitational-wave displacement and spin memory effects with current and future gravitational-wave detectors, *Phys. Rev. D*, 107(6), 064056. [Erratum: *Phys.Rev.D* 108, 029901 (2023)].
- Hoang, B.-M., Naoz, S., Kremer, K. (2020), Neutron Star–Black Hole Mergers from Gravitational-wave Captures, *Astrophys. J.*, 903(1), 8.
- Inchauspé, H., Gasparotto, S., Blas, D., Heisenberg, L., Zosso, J., Tiwari, S. (2024), Measuring gravitational wave memory with LISA, , .
- Isaacson, R. A. (1968a), Gravitational Radiation in the Limit of High Frequency. I. The Linear Approximation and Geometrical Optics, *Phys. Rev.*, 166, 1263–1271.
- Isaacson, R. A. (1968b), Gravitational Radiation in the Limit of High Frequency. II. Nonlinear Terms and the Effective Stress Tensor, *Phys. Rev.*, 166, 1272–1279.
- Kennefick, D. (1997), Controversies in the history of the radiation reaction problem in general relativity, *CalTech Library*, .
- Maggiore, M. (2007), *Gravitational Waves. Vol. 1: Theory and Experiments*, Oxford University Press.
- Mastrogiovanni, S., Laghi, D., Gray, R., Santoro, G. C., Ghosh, A., Karathanasis, C., Leyde, K., Steer, D. A., Perries, S., Pierra, G. (2023), Joint population and cosmological properties inference with gravitational waves standard sirens and galaxy surveys, *Phys. Rev. D*, 108(4), 042002.
- Messenger, C., Read, J. (2012), Measuring a cosmological distance-redshift relationship using only gravitational wave observations of binary neutron star coalescences, *Phys. Rev. Lett.*, 108, 091101.
- Ni, W.-T., Zimmermann, M. (1978), Inertial and gravitational effects in the proper reference frame of an accelerated, rotating observer, *Phys. Rev. D*, 17, 1473–1476.
- O’Leary, R. M., Kocsis, B., Loeb, A. (2009), Gravitational waves from scattering of stellar-mass black holes in galactic nuclei, *Mon. Not. Roy. Astron. Soc.*, 395(4), 2127–2146.
- Peters, P. C., Mathews, J. (1963), Gravitational radiation from point masses in a keplerian orbit, *Physical Review*, 131(1), 435–439.
- Poisson, E., Will, C. M. (2014), *Gravity: Newtonian, Post-Newtonian, Relativistic*, Cambridge University Press.
- Szabados, L. B. (2009), Quasi-Local Energy-Momentum and Angular Momentum in General Relativity, *Living Rev. Rel.*, 12, 4.
- Taherasghari, F., Will, C. M. (2025), Compact binary systems in Einstein–Ether gravity. II. Radiation reaction to 2.5 post-Newtonian order, *Phys. Rev. D*, 112(2), 024013.

- Thorne, K. S. (1980), Multipole Expansions of Gravitational Radiation, *Rev. Mod. Phys.*, 52, 299–339.
- Verde, L., Schöneberg, N., Gil-Marín, H. (2024), A Tale of Many H_0 , *Ann. Rev. Astron. Astrophys.*, 62(1), 287–331.
- Will, C. M. (2011), On the unreasonable effectiveness of the post-Newtonian approximation in gravitational physics, *Proc. Nat. Acad. Sci.*, 108, 5938.

3

Panorama of gravitational-wave sources

Joseph D. Romano

Department of Physics and Astronomy, University of Texas Rio Grande Valley, One West University Boulevard, Brownsville, TX 78520, USA

In this chapter, I give a brief, non-mathematical, “bird’s eye view” of gravitational-wave sources across a broad range of source types, signal classes, and frequency bands. You will encounter these sources again, in much more detail, in several of the subsequent chapters.

3.1. Introduction

These notes provide a brief overview of various sources of gravitational waves (GWs). Since this is a broad introduction to the topic, I will not go into much mathematical detail. You will get that detail in subsequent chapters from other instructors. Nonetheless, I believe it is useful to learn some of the basic vocabulary of GW sources and their corresponding signals early on, before diving into the heavy math.

In general, one can classify GW sources in terms of either: (i) the *physical objects or processes* that produce the GWs, or (ii) the *properties of the GW signals* that the sources produce in the detectors that we use to observe them. Being a data analyst, I often focus on the latter classification, asking what properties of the GW signals allow one to identify (and separate when necessary) the different sources.

But as we shall see, there isn’t a simple one-to-one mapping between GW signal properties and physical sources. The same GW signal may be produced by different physical sources, and the same physical source may produce different GW signals.

This chapter is organized as follows: We start in Sec. 3.2 by discussing GW signal properties. We will restrict attention to *deterministic* signals and some of their sources in Sec. 3.3. After that, we will discuss *stochastic* GW signals and their sources, both astrophysical and cosmological in origin, in Sec. 3.4. Throughout this chapter, we will make reference to material from: (Allen 1997) (Christensen 2018) (Caprini and Figueroa 2018), and (Romano and Christensen 2025). Readers are recommended to visit these sources to obtain more details.

3.2. Signal properties

Let’s start in the time domain, classifying the GW signals in terms of their duration, either short or long, relative to the observation time (typically, a year or so), and whether they have deterministic or non-deterministic waveforms. Short-duration signals are often called *transients*, while long-duration signals are said to be *persistent*. Deterministic waveforms have a “well-defined, predictable phase evolution”, while non-deterministic (also known as stochastic or random) signals have a “non-predictable phase evolution”.

With this classification, GW signals fall into one of four categories as shown in Fig. 3.1. An example of a short duration, deterministic signal is shown in the top left-hand panel, while an example of a long duration, non-deterministic signal is shown in the bottom right-hand panel. Possible sources for each of the signal classes are (from upper left to bottom right): (i) A pair of black holes undergoing the final inspiral, merger, and ringdown to a single black hole; (ii) a rotating (non-axisymmetric) neutron star; (iii) a supernova explosion; (iv) a population of pairs of supermassive black holes orbiting one another in the centers of millions of merging galaxies.




Figure 3.1: Classification of GW signals in terms of their duration (short or long) and phase evolution (deterministic or non-deterministic). Representative time-series data are shown for each class. [Credit: Adapted from Figure 4 in (Jenkins 2021).]

Probably not surprisingly, the GW search groups in the LIGO-Virgo-KAGRA (LVK) Collaboration are divided precisely according to these signal or source classes. Following the same order as above, these are the Compact Binary Coalescence (CBC) Group, the Continuous Wave (CW) Group, the Unmodeled Burst Group, and the Stochastic Gravitational-Wave Background (SGWB) Group, shown in Fig. 3.2.

3.3. Deterministic signals

As mentioned above, deterministic signals have a well-defined, coherent phase evolution in the time domain. The simplest example of a deterministic signal is a sinusoid, like that shown in the top right-hand panel of Fig. 3.1.

	short	long
deterministic	(i) Compact Binary Coalescence (CBC)	(ii) Continuous Wave (CW)
non-deterministic	(iii) Unmodeled Burst	(iv) Stochastic (SGWB)

Figure 3.2: Search groups in the LIGO-Virgo-KAGRA collaboration. They are divided according the signal classes shown in Fig. 3.1.

Another example of a deterministic signal is a *GW chirp*, which consists of oscillations that increases in amplitude and frequency over time until it reaches its peak amplitude. It then decays like an exponentially damped sinusoid, like that shown in the top left-hand panel of Fig. 3.1. The following subsections go into more detail regarding possible sources for each of these types of deterministic signal.

3.3.1. Continuous wave sources

The simplest example of a CW source of GWs is a rotating non-axisymmetric neutron star (NS) with a “bump” on it, which lies somewhere off the rotational axis of the NS. Due to the “bump”, the NS has a non-zero, time-varying quadrupole moment, which is needed for the production of GWs. The frequency of the emitted GWs will equal twice the rotational frequency of the NS. The frequency of the signal will be approximately constant, changing only if the rotational frequency of the NS itself changes, or due to relative motion between the source and the detector (so-called Doppler modulation).

Another example of a CW source is the *early* inspiral of a binary system consisting of NSs or black holes (BHs). During the early inspiral phase (which can last millions of years), the GWs produced by the orbiting objects are approximately monochromatic (i.e., constant) over the timescale of an observation (typically years to decades). The frequency of the GWs does increase over time, since the orbital radius decreases due to energy radiated from the system in the form of GWs. But during the early inspiral phase, the change in frequency is

much smaller than the width of an observational frequency bin ($\Delta f = 1/T_{\text{obs}}$), so it is effectively a single-frequency signal.

A pair of supermassive BHs (having masses $\sim 10^9$ times the mass of the Sun) in the centers of two merging galaxies, emitting GWs in the frequency band 10^{-9} Hz – 10^{-7} Hz (relevant for pulsar timing arrays (PTAs)), is a concrete example of this last type of CW source. The supermassive BH binary (SMBHB) system slowly evolves for millions of years before it exits the PTA band and merges at a higher frequency.

3.3.2. Compact binary coalescence

The canonical source for a GW *chirp* is the late-stage inspiral, merger, and ringdown of a pair of NSs or BHs, or a NS-BH binary. This was the case for GW150914, which was the first direct observation of GWs by the LIGO GW detectors on 14 Sep 2015 (Abbott *et al.* 2016). The binary system that produced the GWs consisted of two BHs each having a mass roughly 30 times the mass of the Sun. Such an inspiral and merger event is called a *compact binary coalescence*.

The word “compact” here refers to the relatively small size of the coalescing objects. NSs, BHs, and white dwarfs (WDs) are all examples of compact objects. They are the three end states of stellar evolution. Low mass stars, having masses between ~ 0.1 and 8 solar masses (one solar mass being $M_{\odot} \approx 2 \times 10^{30}$ kg) will end up as WDs. A WD has a mass $\sim 1.4M_{\odot}$ and radius roughly equal to that of Earth (≈ 6400 km). More massive stars, having masses between 8 and $25 M_{\odot}$, eventually collapse to form a NS after a violent supernova explosion, which again has a mass roughly equal to $1 M_{\odot}$ but with a radius ~ 10 km (the size of a small city). Finally, stars with masses $\gtrsim 25M_{\odot}$ evolve to form a BH. The masses of the BHs formed from stellar collapse range from about $5M_{\odot}$ to roughly $50M_{\odot}$. Black holes with masses $\gtrsim 50M_{\odot}$ can be formed from the collapse of very massive stars in the early Universe (so-called Population-III stars), or from successive mergers of smaller-mass BHs. The characteristic size of a BH is given by its *Schwarzschild radius*, $R_s \equiv 2GM/c^2$, which is the radius of the event horizon of the black hole. (Here, G is Newton’s gravitational constant, and c is the speed of light.) For reference, $R_s \approx 3$ km for $M = M_{\odot}$.

Primordial BHs (PBHs), which are created in the early Universe via the collapse of density perturbations of the primordial plasma, can have much smaller masses than BHs formed from the collapse of ordinary stars discussed above. But PBHs with masses $M \lesssim 10^{12}$ kg (which is the mass of a small asteroid) will have already evaporated due to Hawking radiation. The time needed for a BH to evaporate is $t_{\text{evap}} \approx 10^{64}(M/M_{\odot})^3$ yr.

Just like the relationship between the size of a musical instrument and the frequency of the sound waves that it produces, so too do more massive gravitational systems produce lower-frequency GWs. For binary BHs (BBHs), there is a simple inverse relationship between the merger frequency of the binary system and its total mass M ,

$$f_{\text{merger}} = \frac{c^3}{6^{3/2}\pi GM}, \quad [3.1]$$

which can be derived using Kepler's 3rd law $\omega^2 a^3 = GM$, where $\omega = 2\pi f_{\text{orb}} = \pi f_{\text{merger}}$, where f_{merger} is the GW frequency at $a = 6GM/c^2$, which is the radius of the innermost-stable circular orbit for a pair of non-spinning (i.e., Schwarzschild) BHs. This is the highest GW frequency that the BBH system will produce. Figure 3.3 shows the frequency bands and different GW detectors needed to detect GWs from compact binary systems having different component masses.




Figure 3.3: The frequency bands and detectors needed to detect GWs from compact binary systems having different masses.

NS / stellar-mass BH binaries: GW150914 was the first direct detection of GWs from the final inspiral and merger of two stellar-mass black holes (Abbott *et al.* 2016). Since that first observation on 14 Sep 2015, more than 100 additional binary black hole (BBH) mergers have been detected by the Advanced LIGO and Virgo detectors. All of these mergers have involved pairs of BHs having masses between 5 and 100 M_{\odot} , which is relevant for the frequency range (~ 10 -1000 Hz) of the Advanced LIGO and Virgo detectors.

On 17 Aug 2017, the two Advanced LIGO detectors together with the European GW detector, Virgo, again detected GWs. But this time it was from the final inspiral and merger of two NSs (Abbott *et al.* 2017). This event, denoted GW170817, was the first *multi-messenger* observation of a GW event, being also observed in various forms of light across the electromagnetic spectrum. Since NSs are made of matter, when they smashed into one another after

the final orbits of GW170817, electromagnetic radiation was produced in the process. In fact, some of the constituent neutrons were rapidly converted into heavier elements, most notably gold and platinum. It had long been conjectured that BNS mergers were the source of these heavy metals, but it wasn't until GW170817 that this conjecture could be confirmed.

As of the time of writing, there has been only one more confirmed detection of GWs from a BNS inspiral and merger, GW190425 (Abbott *et al.* 2020). Nonetheless, such BNS mergers are expected to be common throughout the Universe. Rate estimates predict roughly 1 BNS merger in the visible Universe every ~ 15 sec, corresponding to between 10^6 and 10^7 BNS mergers per year. On the other hand, stellar-mass BBH mergers are expected to occur every 5–10 minutes somewhere in the visible Universe, corresponding to approximately 10^5 BBH mergers per year. The reason why we are currently detecting more BBH mergers than BNS mergers is because the BBH systems are heavier (having larger component masses), so we can see them out to larger distances with the same signal-to-noise ratio as the less-massive systems.

White dwarf binaries: Galactic WD binaries (often denoted DWD for “double white dwarfs”) are a prime source for the planned space-based detector LISA (Laser Interferometer Space Antenna). In fact, the combined “confusion noise” signal from millions of unresolved DWDs radiating in the LISA band (10^{-4} Hz to 10^{-1} Hz) will dominate the LISA instrumental noise at low frequencies, see Fig. 3.4. This stochastic GW *foreground* will act much like detector noise when trying to detect other GW signals with LISA.

Loud individually resolvable DWDs can be seen above the unresolved background. One expects to be able to detect tens of thousands of such binaries over the four-year duration of LISA. Of these, approximately 20 DWDs have already been observed electromagnetically, so observing these systems via the GWs that they also produce will be a good initial test of LISA. As such, these systems are often called LISA “verification binaries” in the literature.

Supermassive-black-hole binaries: SMBHBs are on the very-low-frequency end of the GW spectrum. Such systems are expected to form in the centers of merging galaxies, with the orbiting SMBHs producing GWs in the process. As mentioned earlier, SMBHBs with masses of order $10^9 M_{\odot}$ will merge outside the PTA sensitivity band (10^{-9} Hz– 10^{-7} Hz), at frequencies $\gtrsim 10^{-6}$ Hz. Thus, after entering the PTA band, the SMBHBs will orbit one another for several million years before they coalesce, producing a monochromatic (CW) signal over the the timeframe of a typical observation (years to decades).




Figure 3.4: The confusion noise from the population of unresolved galactic white dwarf binaries (labeled as DWD for “double white dwarfs” in the figure). At low frequencies, this confusion noise exceeds the LISA instrumental noise. It is an example of an astrophysical foreground that must be contended with when searching for signals from other, weaker GW sources.

3.4. Stochastic signals and sources

Many different GW sources can give rise to a stochastic GW background. The only condition is that the GW signals that the sources produce should be individually unresolvable. This will be the case if the signals are either too weak or too numerous (interfering with one another in the time or frequency domain) to be individually detected. Since this statement depends on the sensitivity of the detectors, GW sources that are currently unresolvable become resolvable when the detector sensitivity improves, eventually standing out above the lower levels of instrumental and environmental noise.

In the following subsections, we first discuss the signal properties of stochastic GWs. We then describe several different sources, of both astrophysical and cosmological origin.

3.4.1. Signal properties

Stochastic signals, being random, do not have deterministic (i.e., phase-coherent) waveforms. They look much like noise in a single detector. Examples of some very simple, simulated stochastic signals in the time domain are shown in Fig. 3.5 and as excess power coming from different directions on the sky in Fig. 3.6.




Figure 3.5: Simulated time-series data corresponding to different stochastic processes: (i) continuous-in-time, (ii) intermittent (or “popcorn-like”), and (iii) a modulated foreground. The red traces show the GW signal; the black is simulated detector noise (assumed here to be white); and the vertical scales have arbitrary units.




Figure 3.6: Example skymaps of GW power for (i) purely isotropic, (ii) statistically isotropic, and (iii) anisotropic source distributions.

Let’s start with Fig. 3.5. (i) The left-hand panel is the time-domain representation of a stochastic signal that is *continuous in time*. Since it is weaker than the noise, it is called a “background” signal. A possible source of this type of signal is the population of SMBHBs in the centers of merging galaxies throughout the visible Universe, each pair producing approximately monochromatic GWs which interfere with one another. (ii) The middle panel represents a stochastic background signal that is *intermittent*, or “popcorn-like”. A possible source for this type of signal is the population of stellar-mass BBHs, relevant for searches using the ground-based LVK detectors. The individual merger signals are too weak to be individually detected; their durations ($\lesssim 1$ sec) are short relative to the average time between successive signals (~ 5 – 10 minutes); and the arrival time of the individual signals are random (Poisson distributed). So the combined signal is stochastic, even though the individual BBH mergers are deterministic chirps. In contrast, the population of BNS mergers gives rise to a continuous-in-time stochastic background signal since the typical time

in band for a BNS inspiral and merger is ~ 200 sec, while the average time between successive signals is only ~ 15 sec. (iii) The right-hand panel is an example of a “modulated foreground” stochastic signal, having a predictable time variation and being stronger than the noise. A possible source for this type of signal is the population of millions of DWDs in the Milky Way galaxy producing a “confusion-limited” foreground signal. The modulation is due to the orbital motion of the constellation of the three LISA spacecraft as it performs a “cartwheel” during its yearly motion around the Sun. As mentioned previously, the combined signal from these DWDs is so loud that it dominates the LISA noise at low frequencies (see Fig. 3.4).

The different panels of Fig. 3.6 correspond to skymaps of GW power on the sky for different GW source distributions. (i) The left-hand panel is for a purely *isotropic* signal, so having only a monopole component. For this case, there is no preferred direction on the sky, nor preferred angular scales for the distribution of GW power. (ii) The middle panel correspond to a *statistically-isotropic* signal, like that for the cosmic microwave background (CMB), which is a stochastic background of *electromagnetic* waves. On average, there is no preferred direction of the GW power on the sky, but now there are preferred angular scales (in this case centered around 20 degrees). (iii) The right-hand panel corresponds to a statistically *anisotropic* distribution of GW power on the sky, which traces the angular distribution of the GW sources giving rise to the signal. An example of a source that would produce a skymap like this is the population of unresolved DWDs in the Milky Way, relevant for LISA, which we described in the context of Fig. 3.4. The excess GW power traces the shape of the Milky Way galaxy, represented by a vertical band in this skymap, plotted in galactic coordinates.

3.4.2. Characterization of stochastic signals in terms of power spectra

Since stochastic signals do not have deterministic waveforms, they must be characterized instead in terms of their power spectra, their distribution on the sky (e.g., potential anisotropy), and / or the correlations they induce across multiple detectors. The power spectrum of a time-domain signal $h(t)$ is defined as

$$P_h(f) \equiv \frac{2}{T} \langle |\tilde{h}(f)|^2 \rangle, \quad [3.2]$$

where T denotes the duration of the data, $\tilde{h}(f) \equiv \int_{-\infty}^{\infty} dt h(t) e^{-i2\pi ft}$ is the Fourier transform of $h(t)$, and angle brackets $\langle \rangle$ denote an average over elements of an ensemble of possible universes, which can be thought of as draws from a probability distribution for a given set of parameters. Without the angle

brackets, $P_h(f)$ would be an *estimator* of the power spectrum of the random process constructed from the available data.

The above definition of the power spectrum can be generalized to *cross spectra* $C_{12}(f) \equiv \frac{2}{T} \langle \tilde{h}_1(f) \tilde{h}_2^*(f) \rangle$, where $\tilde{h}_{1,2}(f)$ are the Fourier transforms of the response of two detectors to a GWB. The units of $P_h(f)$ and $C_{12}(f)$ would be the square of the units of whatever the detector measures (e.g., dimensionless strain for a LIGO-like measurement or seconds for a PTA timing-residual measurement) divided by Hz (so strain²/Hz or sec²/Hz, respectively).

Isotropic stochastic backgrounds are typically characterized by the (dimensionless) *energy density spectrum*

$$\Omega_{\text{gw}}(f) \equiv \frac{1}{\rho_{\text{crit}}} \frac{d\rho_{\text{gw}}}{d(\ln f)} = \frac{f}{\rho_{\text{crit}}} \frac{d\rho_{\text{gw}}}{df}, \quad [3.3]$$

where $d\rho_{\text{gw}}$ is the energy density in GWs contained in the frequency interval f to $f + df$, and $\rho_{\text{crit}} \equiv 3H_0^2 c^2 / 8\pi G$ is the critical energy density needed to close the Universe today. Here, $H_0 \equiv \dot{a}(t)/a(t)|_{t=t_0}$ is the *Hubble expansion rate* evaluated today, and $a(t)$ is the *scale factor* describing the expansion of the Universe as a function of time. $\Omega_{\text{gw}}(f)$ is related to the strain power spectrum $S_h(f)$ associated with the Fourier components of the metric perturbations $h_{ij}(t, \vec{x})$ via

$$S_h(f) = \frac{3H_0^2}{2\pi^2} \frac{\Omega_{\text{gw}}(f)}{f^3}. \quad [3.4]$$

So it's a simple matter to go back and forth between the power spectrum $S_h(f)$ and the energy density spectrum $\Omega_{\text{gw}}(f)$ for a stochastic GW background. Similarly, the relationship between $P_h(f)$ and $S_h(f)$ is given by $P_h(f) = \mathcal{R}(f)S_h(f)$, where $\mathcal{R}(f)$ is the sky- and polarization-averaged detector response function to a plane-polarized GW.

3.4.3. Astrophysical sources

Astrophysical sources of a stochastic GW background are associated with *populations* of compact stars or stellar remnants which formed $\gtrsim 1$ billion years after the Big Bang. These populations include NS / stellar-mass BH binaries (relevant for ground-based detectors like LIGO, Virgo, KAGRA), DWDs (relevant for LISA), and SMBH binaries (relevant for PTAs), which we discussed previously in Sec. 3.3.2.

Using the definition of $\Omega_{\text{gw}}(f)$ given in [3.3], one can show that for a population of inspiraling binaries, $\Omega_{\text{gw}}(f) \propto f^{2/3}$. This result is a consequence of Kepler's third law $\omega^2 a^3 = \text{const}$, the relationship between orbital energy and separation $E_{\text{orb}} \propto -1/a$, and $dE_{\text{gw}}/d\omega = -dE_{\text{orb}}/d\omega$, where $\omega = 2\pi f_{\text{orb}} = \pi f$ is the orbital angular frequency of the binary, and a here is the semi-major axis of the elliptical orbit. Thus, loss of orbital energy is converted to GW energy emitted by a binary system.

3.4.4. Cosmological sources

Cosmological sources are associated with processes in the very early Universe, which take place well before the formation of stars and galaxies, see Fig. 3.7. The detection of cosmologically-generated GWs is a means for astronomers to observe the Universe mere fractions of a second after the Big Bang, much earlier than what we can do with light. Currently, the earliest “picture” that we have of the Universe is $\sim 380,000$ years after the Big Bang, when the Universe had cooled enough for neutral hydrogen atoms to form. This was the first time that photons could propagate freely, and those photons detected today constitute the cosmic microwave background radiation. Fluctuations in the temperature of the CMB are associated with density perturbations in the Universe when it was roughly 1000 times hotter and 1000 times smaller than it is today.

The potential detection of GWs produced by *inflation*, *first-order phase transitions*, and *cosmic strings* (discussed below) are three possible ways of probing much earlier in the evolution of the Universe.

Inflation: Inflation is the theory that the Universe experienced a period of rapid (exponential) expansion starting around 10^{-36} s after the Big Bang. This rapid expansion is thought to have stretched the Universe to macroscopic scales, making it incredibly smooth and flat. But quantum fluctuations in the vacuum energy density driving inflation would be amplified to macroscopic scales after inflation, becoming the “seeds” for the large-scale structure that we see today. Similarly, quantum fluctuations of the spacetime metric would also be amplified by inflation, leading to the production of GWs. These GWs would form a stochastic GW background that is essentially flat in energy density across the different frequency bands, see Fig. 3.8. Standard inflationary models would produce a background at the level of $\Omega_{\text{gw}}(f) \sim 10^{-15}$ to 10^{-17} (constrained by the high level of isotropy of the CMB), which would be difficult to observe. However, there are modified theories of inflation which might push the energy density to higher values at higher frequencies.




Figure 3.7: A schematic diagram showing the evolution of the Universe from the Big Bang until today, roughly 14 billion years later. Highlighted are the times and energy scales for some important events during the history of the Universe.




Figure 3.8: Representative energy density spectra for several different GW sources (both astrophysical and cosmological) across a large range of frequencies. The relevant frequency bands for different GW detectors are also shown.

First-order phase transitions (FOPTs): A first-order phase transition is associated with a discontinuity in certain thermodynamic properties of a material. For example, water boiling to produce steam is an example of a first order phase transition, since the entropy ($dS = dQ/T$) changes discontinuously with temperature due to the latent heat of vaporization at fixed $T = 100$ Celsius. In the same way that boiling water creates bubbles of the new phase (steam), which expand in the old phase (water) and collide with one another, creating sound waves and turbulent flows, so too would a first-order phase transition in the early Universe create colliding bubbles and sound waves, which are the primary source of GWs.

Electroweak (EW) symmetry-breaking in the Standard Model is not thought to be a first-order phase transition. However, in some modifications to the Standard Model, a first-order phase transition might occur, producing GWs that could potentially be detected in the mHz frequency band (relevant for LISA). Similarly, a first-order QCD phase transition (when quarks first combine to form protons, neutrons, ...) would produce a stochastic GW background that is potentially observable by PTAs. Either of these observations would provide information about *new physics*, going beyond the predictions of the Standard Model (see Fig. 3.8).

Cosmic strings: Cosmic strings are one-dimensional topological defects, expected to form during phase transitions that have spontaneously broken symmetries. They are predicted in the context of grand unified theories. Cosmic strings would be relics of an earlier more-symmetric phase of the Universe, decaying only via the emission of GWs. GWs are produced when the strings oscillate, form loops, cusps, and kinks, and interact with one another in a network of strings. Cosmic strings are usually quantified by the dimensionless string tension or mass-per-unit-length $G\mu/c^2$, where μ has dimensions of mass/length. A stochastic background of GWs from cosmic strings having $G\mu/c^2 \gtrsim 10^{-16}$ could be detectable with LISA.

In order to better understand what early-Universe timescales our current detectors are potentially able to observe, we ask the following question (see (Allen 1997)): at what cosmic time t was the characteristic size of the Universe, ct , equal to the wavelength $\lambda(t)$ of a GW that, when redshifted to today, would fall into the sensitive frequency band of one of our GW detectors? To answer this question, we need to know how the scale factor of the Universe, $a(t)$, has evolved over time, since the wavelength at time t is directly proportional to the scale factor at time t :

$$\lambda_0 = (1 + z)\lambda(t) = \lambda(t)/a(t), \quad [3.5]$$

where λ_0 is the wavelength observed today. Approximating the Universe as being radiation-dominated for all times¹, we have

$$a(t) = (t/t_0)^{1/2}, \quad [3.6]$$

where $t_0 \approx 14 \times 10^9$ yrs is the current age of the Universe. Combining these last two equations and equating $\lambda(t) = ct$ yields a very simple relationship for the cosmic time t_* that solves this equation:

$$\lambda_0 = c(t_*t_0)^{1/2} \quad \Leftrightarrow \quad f_0 = (t_*t_0)^{-1/2}. \quad [3.7]$$

See Fig. 3.9 for a graphical visualization of this solution.

For example, using [3.7] one can show that $t_* = 10^{-22}$ sec corresponds to a wavelength today of $\lambda_0 = 3000$ km or, equivalently, a frequency of $f_0 = 100$ Hz. This frequency falls right in the center of the LVK band. Similarly, $t_* = 10^{-11}$ s and $t_* = 10^{-5}$ s (which are approximate times for the EW and QCD phase transitions) correspond to observable frequencies today of approximately $f_0 = 5 \times 10^{-4}$ Hz and $f_0 = 5 \times 10^{-7}$ Hz, which fall at the edges of the LISA and PTA sensitivity bands, as seen in Fig. 3.8.

3.5. Acknowledgments

JDR acknowledges financial support from National Science Foundation (NSF) Grant PHY-2207270, Physics Frontier Center Award PFC-2020265, as well as start-up funds from the University of Texas Rio Grande Valley. He also thanks colleagues Nelson Christensen and Marie-Anne Bizouard for their support and hospitality while visiting the Artemis Lab, Observatoire Cote D'Azur during the summer for the past several years.

3.6. Bibliography

Abbott, B. P., Abbott, R., Abbott, T. D., Abraham, S., Acernese, F., Ackley, K., Adams, C., Adhikari, R. X., Adya, V. B., Affeldt, C., et al. (2020), Gw190425: Observation of a compact binary coalescence with total mass $\sim 3.4 m_\odot$, *The Astrophysical Journal Letters*, 892(1), L3. LIGO Scientific Collaboration and Virgo Collaboration.

URL: <https://doi.org/10.3847/2041-8213/ab75f5>

1. This is not a bad approximation since the Universe has been matter and vacuum-energy dominated for only about one-tenth (on a log scale) of the history of the Universe between the end of inflation until today.




Figure 3.9: Determining the cosmic time where the characteristic size of the Universe agrees with the wavelength of a GW that, when redshifted to today, falls within one of the current detector sensitivity bands.

- Abbott, B. P., Abbott, R., Abbott, T. D., Acernese, F., Ackley, K., Adams, C., Adams, T., Addesso, P., Adhikari, R. X., Adya, V. B., et al. (2017), GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, *Physical Review Letters*, 119(16), 161101. LIGO Scientific Collaboration and Virgo Collaboration.
URL: <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.119.161101>
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. (2016), Observation of gravitational waves from a binary black hole merger, *Phys. Rev. Lett.*, 116, 061102.
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.061102>
- Allen, B. (1997), The stochastic gravity-wave background: Sources and detection, in J.-A. Marck, J.-P. Lasota, (eds), *Relativistic Gravitation and Gravitational Radiation*, vol. 26 of *Cambridge Contemporary Astrophysics*, Cambridge University Press, pp. 373–417.
URL: <https://arxiv.org/abs/gr-qc/9604033>
- Caprini, C., Figueroa, D. G. (2018), Cosmological backgrounds of gravitational waves, *Classical and Quantum Gravity*, 35(16), 163001.
URL: <https://dx.doi.org/10.1088/1361-6382/aac608>
- Christensen, N. (2018), Stochastic gravitational wave backgrounds, *Reports on Progress in Physics*, 82(1), 016903.

URL: <https://dx.doi.org/10.1088/1361-6633/aae6b5>

Jenkins, A. (2021), *Cosmology and Fundamental Physics in the Era of Gravitational-Wave Astronomy*, Phd thesis, King's College London. Supervised by Mairi Sakellariadou.

Romano, J. D., Christensen, N. (2025), 'Stochastic gravitational wave backgrounds', https://www.scholarpedia.org/article/Stochastic_gravitational_wave_backgrounds. Scholarpedia article.

Chap

DRAFT

4

Gravitational wave detectors from an experimental perspective

Marina Trad-Nery, Margherita Turconi and Walid Chaibi

*Université Côte d'Azur, OCA, CNRS, ARTEMIS, Boulevard de
l'Observatoire F-06304 Nice, France*

GRAVITATIONAL WAVES,
coordinated by Nicola Tamanini (L2IT/CNRS). © ISTE Editions 2025.

Foreword

This chapter introduces the fundamental principles of gravitational wave detectors (GWDs) in a simple and comprehensive manner. Because these instruments aim for extremely high sensitivity, it is essential to understand their various noise sources, how such noise couples to the detector output, and the strategies used to mitigate them. We begin with a description of noise from an experimental physicist's perspective, including an introduction to spectral densities, which are widely used to describe the detector sensitivity and noise. In the second section, the effect of a gravitational wave on a light field is calculated from the modification of the space-time metric. We then show how a simple Michelson interferometer can be used to detect gravitational waves and which parameters can be modified to enhance the gravitational wave signal. This section also introduces how different noise sources couple into the interferometer readout. From the third section onward, we move from a basic Michelson interferometer to a more realistic model of a gravitational-wave detector. For that, we introduce Fabry-Perot cavities and show how they can be used in the interferometer arms to increase the signal of a gravitational wave (GW), and how they are used as a filter for the spatial profile of the laser beam. We also discuss their use as sensors in frequency-stabilization feedback loops, and provide a general introduction to control systems that will be used later in the chapter. The advanced detector configuration is then described: a Michelson interferometer equipped with Fabry-Perot arm cavities, a power recycling mirror, laser power and frequency stabilization, and a mode-cleaner cavity. Finally, we focus on the main noise sources that limit detector performance: seismic noise, thermal noise, and quantum noise. This description is made with a simple concept of a harmonic oscillator, described both in the classical and quantum approaches. To conclude, these noises contributions are computed in the frame of the Virgo detector and a sensitivity curve is calculated. Although a simplified layout of a gravitational wave detector is considered, it takes into account the most dominant effects and yields in a sensitivity estimate close to the what is observed in real detectors.

4.1. The noise for an experimental physicist

As for any metrology experiment, setting up a gravitational wave detector requires managing a variety of noise sources of different physical origins. Measuring those noise sources and understanding their coupling paths to the detector is as important as identifying their nature. It requires a deep understanding of the corresponding mathematical basis and its physical interpretation. In the following sections, we present the definition of time dependent noise (relevant to gravitational wave detectors) with the concern of relating it to the measurement process in classical physics.

4.1.1. Noise definition

Noise can be defined as a random (or unpredictable) signal that adds to the physical quantity to be measured, potentially masking or distorting its true value. The total signal \mathcal{S} (which is real) that is measured can be written as:

$$\mathcal{S} = x + \epsilon, \quad [4.1]$$

where x is the quantity to be measured and ϵ is the noise due to a random process and is characterized by its probability density $p(\epsilon)$. In the specific case where the noise is time-dependent, one can write:

$$\mathcal{S}(t) = x(t) + \epsilon(t), \quad [4.2]$$

where for each time t , $\epsilon(t)$ is a time dependent random process characterized by the time-dependent probability density $p_t(\epsilon)$.

In order to extract $x(t)$ out of the measurement $\mathcal{S}(t)$, one has to address the noise, which will be the focus of the next sections.

4.1.2. Noise characterization

From now on, only time dependent random processes will be considered. They are partially characterized by their corresponding time varying probability density. Hereafter, we give an infinite set of parameters defined for a single time t as an alternative to characterize the process:

- The mean value:

$$\mu_\epsilon(t) = \langle \epsilon(t) \rangle = \int_{-\infty}^{+\infty} \epsilon(t) \times p_t(\epsilon) d\epsilon, \quad [4.3]$$

- The variance:

$$\mu_{\epsilon,2}(t) = \sigma_\epsilon^2(t) = \langle (\epsilon(t) - \mu_\epsilon(t))^2 \rangle = \int_{-\infty}^{+\infty} (\epsilon(t) - \mu_\epsilon(t))^2 \times p_t(\epsilon) d\epsilon, [4.4]$$

- The moment of order n :

$$\mu_{\epsilon,n}(t) = \langle (\epsilon(t) - \mu_\epsilon(t))^n \rangle = \int_{-\infty}^{+\infty} (\epsilon(t) - \mu_\epsilon(t))^n \times p_t(\epsilon) d\epsilon, [4.5]$$

- For any function f the mean value of $f(\epsilon)$:

$$\mu_{f(\epsilon)}(t) = \langle f(\epsilon(t)) \rangle = \int_{-\infty}^{+\infty} f(\epsilon) \times p_t(\epsilon) d\epsilon. \quad [4.6]$$

In the last formulae, the brackets correspond to an average procedure on the outcomes of an ensemble of identical experiments, all evaluated at time t . Knowing this infinite set of parameters (infinite moments n) fully substitutes the probability density $p_t(\epsilon)$ for a given time t , but it doesn't give any information on how it evolves with time. We will see next that the auto-correlation functions will give this information.

4.1.3. Auto-correlation

We are now interested in characterizing the relationship between $\epsilon(t_1)$ and $\epsilon(t_2)$ which is not deterministic. One can analyze the so-called auto-correlation function Γ_ϵ , defined as:

$$\begin{aligned} \Gamma_\epsilon(t_1, t_2) &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \epsilon_1 \times \epsilon_2 \times p_{t_1, t_2}(\epsilon_1, \epsilon_2) d\epsilon_1 d\epsilon_2 \\ &= \langle (\epsilon(t_1) \epsilon(t_2)) \rangle, \end{aligned} \quad [4.7]$$

where $p_{t_1, t_2}(\epsilon_1, \epsilon_2)$ is the joint probability density function and $\epsilon_j = \epsilon(t_j)$, ($j = 1, 2$). Additional information can be obtained by analyzing the auto-correlation functions of order n involving $n - 2$ intermediate times $t'_1, t'_2 \dots t'_{n-2}$ between t_1 and t_2 :

$$\begin{aligned} &\Gamma_\epsilon^{(n)}(t_1, t'_1, \dots, t'_{n-2}, t_2) \\ &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \dots \epsilon_1 \times \epsilon'_1 \times \dots \times \epsilon_2 \times p_{t_1, t'_1, \dots, t'_2}(\epsilon_1, \epsilon'_1, \dots, \epsilon_2) d\epsilon_1 d\epsilon'_1 \dots d\epsilon_2 \\ &= \langle (\epsilon(t_1) \epsilon(t'_1) \dots \epsilon(t'_{n-2}) \epsilon(t_2)) \rangle. \end{aligned} \quad [4.8]$$

with $\epsilon'_j = \epsilon(t'_j)$. In practice, however, the second order auto-correlation function already gives enough information about the considered process.

4.1.4. Stationary noise

In the specific case of stationary noise, the noise characteristics (set of parameters given by Equations 4.3 to 4.6) do not change with time. Hence, for all sets of times t_1, t_2, \dots, t_n and any delay T , one has:

$$p_{t_1, t_2, \dots, t_n}(\epsilon_1, \epsilon_2, \dots, \epsilon_n) = p_{t_1+T, t_2+T, \dots, t_n+T}(\epsilon_1, \epsilon_2, \dots, \epsilon_n). \quad [4.9]$$

An interesting case is the 2nd order stationary and centered noise where the condition 4.9 stands up to $n = 2$:

$$p_t(\epsilon) = p_{t+\tau}(\epsilon) = p(\epsilon) \quad \text{and} \quad p_{t_1, t_2}(\epsilon_1, \epsilon_2) = p_{t_1+\tau, t_2+\tau}(\epsilon_1, \epsilon_2), \quad [4.10]$$

and the noise has a zero mean value:

$$\mu_\epsilon(t) = \langle \epsilon(t) \rangle = 0. \quad [4.11]$$

The standard deviation is $\sigma_\epsilon^2(t) = \sigma_\epsilon^2$, and the second order auto-correlation only depends on the delay $\tau = t_2 - t_1$. We then have :

$$\begin{cases} \Gamma_\epsilon(t, t + \tau) = \Gamma_\epsilon(0, \tau), \\ \sigma_\epsilon^2 = \Gamma_\epsilon(0, 0) \geq \Gamma_\epsilon(0, \tau), \end{cases} \quad [4.12]$$

where the inequality can be demonstrated by calculating $\langle (\epsilon(t) \epsilon(t + \tau))^2 \rangle$. It should be noted that the last equation only stands for classical process and is violated in quantum mechanics (Bachor-Ralph 2019). On Figure 4.1, two types of noise characterized by different behaviors of the auto-correlation functions are shown: white noise for which the correlation between two different times is strictly equal to zero, i.e. $\Gamma_\epsilon(\tau) = \sigma_\epsilon^2 \times \delta(\tau)$, and colored noise for which the auto-correlation function varies smoothly from its maximum down to zero.

4.1.5. Temporal average and ergodicity

Let us examine the situation in which only a single experiment is available. For any function f applied on the noise process $\epsilon(t)$, one can define the temporal average m_f :

$$m_f = \overline{f(\epsilon)} = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T f(\epsilon(t)) dt, \quad [4.13]$$

with the assumption that this limit actually exists. In particular, one can define the mean value, the standard deviation, and the auto-correlation function of the noise ϵ based on the temporal average:

$$m_\epsilon = \bar{\epsilon} = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T \epsilon(t) dt, \quad [4.14]$$




Figure 4.1: Two examples of noises: white and colored. On the temporal profile (left), these two noises are hardly distinguishable, whereas their auto-correlation functions (right) have different behaviors.

$$s_\epsilon^2 = \overline{\epsilon^2} = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T \epsilon^2(t) dt, \quad [4.15]$$

$$C_\epsilon(\tau) = \overline{\epsilon(t)\epsilon(t+\tau)} = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T \epsilon(t)\epsilon(t+\tau) dt, \quad [4.16]$$

$$\begin{aligned} C_\epsilon^{(n)}(\tau_1, \dots, \tau_n) &= \overline{\epsilon(t)\epsilon(t+\tau_1)\dots\epsilon(t+\tau_n)} \\ &= \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T \epsilon(t)\epsilon(t+\tau)\dots\epsilon(t+\tau_n) dt. \end{aligned} \quad [4.17]$$

Again, this is considered for a single experiment. In the case of multiple similar experiments in which all temporal averages remain quantitatively the same, the process is called ergodic. The specific case of 2nd order ergodicity:

$$m_\epsilon^{(i)} = m_\epsilon^{(j)} = m_\epsilon = \bar{\epsilon}, \quad [4.18]$$

$$s_\epsilon^{(i)2} = s_\epsilon^{(j)2} = s_\epsilon^2 = \overline{\epsilon^2}, \quad [4.19]$$

$$C_\epsilon^{(i)}(\tau) = C_\epsilon^{(j)}(\tau) = C_\epsilon(\tau) = \overline{\epsilon(t)\epsilon(t+\tau)}, \quad [4.20]$$

is of particular interest for the rest of this section ((*i*) and (*j*) label two different experiments).

4.1.6. Stationary and ergodic process

The main class of process one can deal with in experimental physics is the so called (2nd order) stationary and (2nd order) ergodic, for which both conditions described in sections 4.1.4 and 4.1.5 are verified. Additionally, we'll consider this process to be centered (zero mean) as it is the case for most of noises $\epsilon(t)$. For a given function $f(\epsilon)$, one can use either averages over different experiments or over time:

$$\mu_{f(\epsilon)}(t) = \langle f(\epsilon(t)) \rangle = \int_{-\infty}^{+\infty} f(\epsilon) \times p_t(\epsilon) d\epsilon, \quad [4.21]$$

$$m_{f(\epsilon)} = \overline{f(\epsilon)} = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T f(\epsilon(t)) dt. \quad [4.22]$$

Since ϵ is stationary and ergodic these two averages will be equal:

$$\begin{aligned} \mu_{f(\epsilon)} &= \overline{\mu_{f(\epsilon)}(t)} && \text{stationarity} \\ &= \overline{\langle f(\epsilon^{(i)}(t)) \rangle} \\ &= \langle \overline{f(\epsilon^{(i)}(t))} \rangle && \text{linearity} \\ &= \langle m_{f(\epsilon^{(i)})} \rangle \\ &= \langle m_{f(\epsilon)} \rangle && \text{ergodicity} \\ &= m_{f(\epsilon)}. \end{aligned} \quad [4.23]$$

More specifically:

$$m_\epsilon = \mu_\epsilon = 0 \quad ; \quad s_\epsilon^2 = \sigma_\epsilon^2 \quad ; \quad C_\epsilon(\tau) = \Gamma_\epsilon(\tau). \quad [4.24]$$

From now on, all noises will be considered as (2nd order) stationary and ergodic, and averages will be, case by case, computed following either of the two procedures.

4.1.7. Harmonic analysis: The Fourier transform

The Fourier transform is the optimal tool to perform harmonic analysis. We'll see in the following that depending on the type of process, it is applied differently with different physical interpretations.

4.1.7.1. Finite energy process

From a mathematical point of view, a finite energy (\mathcal{E}) real process satisfies the condition:

$$\mathcal{E} = \int_{-\infty}^{+\infty} \epsilon(t)^2 dt, \quad [4.25]$$

which guarantees the existence of its Fourier transform $\tilde{\epsilon}$:

$$\tilde{\epsilon}(f) = \int_{-\infty}^{+\infty} e^{-2i\pi ft} \epsilon(t) dt. \quad [4.26]$$

The Parseval equality relates the energy \mathcal{E} to the Fourier transform $\tilde{\epsilon}(f)$ as:

$$\mathcal{E} = \int_{-\infty}^{+\infty} |\tilde{\epsilon}(f)|^2 df. \quad [4.27]$$

Hence, $|\tilde{\epsilon}(f)|^2$ represents the energy density per spectral interval df and is called the Energy Spectral Density (ESD).

4.1.7.2. Continuous process

By definition, a continuous process $\epsilon(t)$ has an infinite energy since it lasts indefinitely (example: the energy of a continuous wave (cw) laser beam):

$$\int_{-\infty}^{+\infty} \epsilon(t)^2 dt = \infty. \quad [4.28]$$

However, it is expected to have a finite mean power:

$$\overline{P}_\epsilon = \lim_{T \rightarrow +\infty} \frac{1}{2T} \int_{-T}^T \epsilon(t)^2 dt. \quad [4.29]$$

For the case of a stationary and ergodic noise and according to Equation 4.15 and Equation 4.29, the noise mean power is directly given by the variance: $\overline{P}_\epsilon = \sigma_\epsilon^2$.

If it exists, we are interested in the quantity:

$$S_\epsilon(f) = \lim_{T \rightarrow +\infty} \frac{1}{2T} |\tilde{\epsilon}_T(f)|^2 = \lim_{T \rightarrow +\infty} \frac{1}{2T} \left| \int_{-T}^T e^{-i2\pi ft} \epsilon(t) dt \right|^2, \quad [4.30]$$

where $\tilde{\epsilon}_T$ is the Fourier transform of the truncated function $\epsilon(t)$ over $[-T, T]$. Summed over all frequencies f one obtains:

$$\begin{aligned} \int_{-\infty}^{+\infty} S_\epsilon(f) df &= \lim_{T \rightarrow +\infty} \frac{1}{2T} \int_{-\infty}^{+\infty} |\tilde{\epsilon}_T(f)|^2 df \\ &= \lim_{T \rightarrow +\infty} \frac{1}{2T} \int_{-\infty}^{+\infty} \epsilon_T(t)^2 dt \\ &= \lim_{T \rightarrow +\infty} \frac{1}{2T} \int_{-T}^T \epsilon(t)^2 dt \\ &= \sigma_\epsilon^2 = \overline{P}_\epsilon. \end{aligned} \quad [4.31]$$

Therefore, if $S_\epsilon(f)$ exists, it represents the power density per spectral interval df and is then called Power Spectral Density (PSD).

4.1.7.3. The Wiener-Khintchine Theorem

This theorem states that for a stationary and ergodic process (like the type of noise we are dealing with) the limit defining the PSD $S_\epsilon(f)$ (Equation 4.30) does exist and it is given by the Fourier transform of the corresponding auto-correlation function:

$$\begin{aligned} S_\epsilon(f) &= \lim_{T \rightarrow \infty} \frac{1}{2T} \left| \int_{-T}^T e^{-i2\pi ft} \epsilon(t) dt \right|^2 \\ &= \int_{-\infty}^{+\infty} e^{-2i\pi f\tau} \Gamma_\epsilon(\tau) d\tau. \end{aligned} \quad [4.32]$$

The auto-correlation form of the PSD is very useful for analytic calculations, whereas for experimental measurements, we limit the integration in Equation 4.30 on a properly chosen acquisition time T_{acq} and we average it over a number N of successive measurements:

$$S_\epsilon(\nu) \simeq \frac{1}{T_{\text{acq}}} \left\langle \left| \int_0^{T_{\text{acq}}} e^{-i2\pi \nu t} \epsilon(t) dt \right|^2 \right\rangle_N. \quad [4.33]$$

In practice, $T_{\text{acq}} \gg 1/f_{\text{min}}$ where f_{min} is the minimum frequency we might be interested in. It is usual to characterize the noise by its Amplitude Spectral Density (ASD) which is the square root of its PSD. It results in an usual unit, which can be seen on the simple case of a voltage noise $\delta V(t)$:

$$\begin{aligned}\delta V(t) &\rightarrow V, \\ \Gamma_{\delta V}(\tau) &\rightarrow V^2, \\ S_{\delta V}(f) &\rightarrow V^2/\text{Hz}, \\ \sqrt{S_{\delta V}(f)} &\rightarrow V/\sqrt{\text{Hz}}.\end{aligned}$$

4.2. Detection principle of an interferometric ground based gravitational wave detector

Ground-based gravitational-wave detectors employ a Michelson interferometer configuration that transduces the gravitational-wave-induced phase shift in the interferometer arms into a measurable change in optical power at its output. In this section, we will compute the output power of the interferometer as a function of a gravitational wave (GW) signal and of usual noise sources. We begin by providing a simplified description of the gravitational wave arriving at the detector, which will then allow us to compute its effect on the light propagating in the interferometer.

4.2.1. Gravitational wave effect on a light field




Figure 4.2: Usual representation of the GW polarizations

Far from the source, the GW can be considered as a plane wave. For the sake of simplicity, we will consider its propagation along the z axis orthogonal to the plane (x, y) in which the light propagates (see Figure 4.2). Its polarization in the plane (x, y) can be projected on the set of orthogonal polarizations (+) and (\times). As represented in Figure 4.2, the (+) polarization corresponds to the

elongation of the space time along one axis (here x) and the (\times) polarization corresponds to a similar elongation at 45° . Mathematically speaking, the GW is described by its modification of the local space-time metric:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad [4.34]$$

where $\eta_{\mu\nu}$ is the Minkowski metric describing a flat space time (without perturbation):

$$\eta_{\mu\nu} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}. \quad [4.35]$$

In the TT (Transverse Traceless) gauge (Reitze-Saulson-Grote 2019), the GW is described by the time dependent perturbation $h_{\mu\nu}$ which we simplified to a pure (+) cross polarization for the sake of simplicity:

$$h_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & +h(t) & 0 & 0 \\ 0 & 0 & -h(t) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad [4.36]$$

By definition, light follows the geodesic equation, which in Einstein's notations, is written as:

$$g_{\mu\nu} dx^\mu dx^\nu = 0, \quad [4.37]$$

where $dx^{\mu(\nu)} = (cdt, dx, dy, dz)$. Expanding Equation 4.37 gives:

$$c^2 dt^2 + dx^2(-1 + h(t)) + dy^2(-1 - h(t)) - dz^2 = 0. \quad [4.38]$$

Now, we consider a "classical" photon (as it was defined by Einstein) propagating along the x axis, i.e. $dy = 0$ and $dz = 0$, one obtains:

$$dx = \pm c \sqrt{\frac{1}{1 - h(t)}} dt \simeq \pm c \left(1 + \frac{1}{2} h(t) \right) dt. \quad [4.39]$$

We now consider two coordinates points 0 and L . The coordinates of these two points remains unchanged in the TT gauge. However, if a photon is launched from $x = 0$ at t_0 , it will arrive at $x = L$ at a time t_L that satisfies:

$$L = \int_0^L dx = c \int_{t_0}^{t_L} \left(1 + \frac{1}{2} h(t) \right) dt = c \left(\int_{t_0}^{t_L} dt + \frac{1}{2} \int_{t_0}^{t_L} h(t) dt \right). \quad [4.40]$$

In case the characteristic time of $h(t)$ is much longer than the unperturbed propagation time L/c , $h(t)$ can be taken out of the integral, leading to:

$$t_L \simeq t_0 + \frac{L}{c} \left(1 - \frac{1}{2} h(t) \right) \quad \text{for the } x \text{ axis}, \quad [4.41]$$

showing that the GW produces a time delay on the propagation of the photon. The same effect can be computed for the y direction but with an opposite delay:

$$t_L \simeq t_0 + \frac{L}{c} \left(1 + \frac{1}{2} h(t) \right) \quad \text{for the } y \text{ axis}, \quad [4.42]$$

which is characteristic of (+) polarization.

Let us apply these results obtained for a single photon to the more general case of an electromagnetic field. For a light field propagating along x , $E(x, t)$, one can write:

$$\begin{aligned} \text{without GW : } E(L, t) &= E \left(0, t - \frac{L}{c} \right), \\ \text{with GW : } E(L, t) &= E \left(0, t - \frac{L}{c} \left(1 - \frac{h(t)}{2} \right) \right) = E \left(L \left(1 - \frac{h(t)}{2} \right), t \right). \end{aligned} \quad [4.43]$$

So, without GW, the electric field that arrives at a time t at $x = L$ is the electric field that was at $x = 0$ at time $t - L/c$. The GW causes a change in the propagation time according to Equation 4.41 that can also be expressed as a change in the propagation length.

The complex electric field that describes a single frequency laser beam at $x = L$ in presence of a GW can be written as:

$$E_L(t) = E_0 e^{-i2\pi\nu_0 t} \cdot e^{i\frac{2\pi\nu_0 L}{c} \left(1 - \frac{h(t)}{2} \right)}, \quad [4.44]$$

where ν_0 is the laser frequency and E_0 the field amplitude. This equation shows that the GW acts on the phase of a light beam, and since the phase cannot be directly measured with a photodetector, interferometry is an ideal technique to measure gravitational waves.

4.2.2. The Michelson interferometer, its output and tuning

The detection of GWs using optical interferometry was first proposed in 1963 by Gertsenshtein and Pustovoit (Gertsenshtein and Pustovoit 1963). They proposed using a Michelson interferometer as a detector, since it measures the phase difference between two orthogonal arms. This makes it well suited for detecting a GW when the wave's polarization axes are aligned with the interferometer arms. In the following sections we compute the output power of the interferometer as a function of the gravitational wave signal, and also as a function of the main noise sources in GWs.

We begin by calculating the complex electric fields in transmission (E_t) and reflection (E_r) of a simple Michelson interferometer (see Figure 4.3a) as a function of the complex input field $E_i(t) = E_0 \cdot e^{-i(2\pi\nu_0 t + \phi_L(t))}$ and of the round-trip phase $\phi_{N(W)}$ acquired by the fields in the north(west) arm (north and west are the directions of the Virgo interferometer's arms). Here $\phi_L(t)$ is the time dependent phase fluctuation of the input field, which represents the phase noise of the laser source. For this calculation the so called real beamsplitter convention, illustrated on Figure 4.3b, will be used. In this convention, the complex field acquires a minus sign in reflection when it impinges on one of the surfaces of the beamsplitter, whereas it remains unchanged when reflected on the other surface. It also remains unchanged for both transmissions. A phase convention is required in order to satisfy the energy conservation condition. For a 50:50 beamsplitter, like in the case of GWs, the reflection (\sqrt{R}) and transmission (\sqrt{T}) field amplitude coefficients are equal to $\sqrt{R} = \sqrt{T} = 1/\sqrt{2}$. The transmitted field is the sum of the field transmitted into the west arm that, after a round-trip, is reflected towards the photodetector, and of the field reflected to the north arm that is subsequently transmitted by the beamsplitter. A similar logic applies to the reflected interferometer field. This will lead to the following equations:

$$\begin{aligned} E_t(t) &= E_0 e^{-i2\pi\nu_0 t} \cdot \sqrt{T} \cdot \sqrt{R} \cdot e^{i\phi_W(t)} - E_0 e^{-i2\pi\nu_0 t} \cdot \sqrt{R} \cdot \sqrt{T} \cdot e^{i\phi_N(t)} \\ &= \frac{1}{2} E_0 e^{-i2\pi\nu_0 t} (e^{i\phi_W(t)} - e^{i\phi_N(t)}), \end{aligned} \quad [4.45]$$

$$\begin{aligned} E_r(t) &= E_0 e^{-i2\pi\nu_0 t} \cdot T \cdot e^{i\phi_W(t)} + E_0 e^{-i2\pi\nu_0 t} \cdot R \cdot e^{i\phi_N(t)} \\ &= \frac{1}{2} E_0 e^{-i2\pi\nu_0 t} (e^{i\phi_W(t)} + e^{i\phi_N(t)}), \end{aligned} \quad [4.46]$$

with the round-trip phase on the interferometer arms given by:

$$\phi_W = \frac{4\pi\nu_0 \cdot L_W}{c} \cdot \left(1 - \frac{h(t)}{2}\right) - \phi_L(t - 2L_W/c), \quad [4.47]$$

$$\phi_N = \frac{4\pi\nu_0 \cdot L_N}{c} \cdot \left(1 + \frac{h(t)}{2}\right) - \phi_L(t - 2L_N/c). \quad [4.48]$$




Figure 4.3: a) Schematics of a Michelson interferometer. Light coming from a laser source is split by a beamsplitter into two orthogonal arms and fully reflected back to the beamsplitter. Depending on the interference, part of the light is reflected back to the laser and part is transmitted to a photodetector. b) Fields in reflection (green arrows) and transmission (orange arrows) of a beamsplitter represented by the real convention, where the reflected light field acquires a minus sign in one (here the left) of the surfaces of the beamsplitter.

Here $L_{N(W)}$ are the geometrical lengths of the interferometer arms in the absence of a gravitational wave. Note that, to first order, the time argument for ϕ_L takes into account only the travel time of the light without the gravitational wave effect. It is useful to re-write these equations as a function of the common ($\bar{\phi}$, \bar{L}) and differential ($\Delta\phi$, ΔL) phase and arm lengths, given by:

$$\bar{\phi} = \frac{\phi_W + \phi_N}{2} \quad \text{and} \quad \Delta\phi = \frac{\phi_W - \phi_N}{2}, \quad [4.49]$$

$$\bar{L} = \frac{L_W + L_N}{2} \quad \text{and} \quad \Delta L = \frac{L_W - L_N}{2}. \quad [4.50]$$

This will lead to:

$$E_t(t) = iE_0 e^{-i(2\pi\nu_0 t - \bar{\phi})} \cdot \sin(\Delta\phi), \quad [4.51]$$

$$E_r(t) = E_0 e^{-i(2\pi\nu_0 t - \bar{\phi})} \cdot \cos(\Delta\phi). \quad [4.52]$$

Since the power is proportional to the modulus square of the electric field, it will only depend on the differential phase $\Delta\phi$. For frequencies $f \ll c/(2\bar{L})$, we obtain:

$$\Delta\phi \simeq \frac{4\pi\nu_0 \Delta L}{c} - \frac{4\pi\nu_0 \bar{L}}{c} \cdot \frac{h(t)}{2} - \frac{4\pi\delta\nu(t)\Delta L}{c}. \quad [4.53]$$

The term $\delta\nu(t)$ is the frequency noise of the laser which has the following relation with the phase noise:

$$\delta\nu(t) = \frac{1}{2\pi} \frac{d\phi_L(t)}{dt}. \quad [4.54]$$

The transmitted power by the interferometer as a function of the input power P_0 is:

$$\begin{aligned} P_t(t) &= K \times |E_t(t)|^2 \\ &= P_0 \cdot \sin^2 \left(\frac{4\pi\Delta L}{\lambda_0} - \frac{2\pi\bar{L}}{\lambda_0} \cdot h(t) - \frac{4\pi\delta\nu(t)}{\nu_0} \frac{\Delta L}{\lambda_0} \right). \end{aligned} \quad [4.55]$$

where K is a constant depending on the beam shape that will be defined later. One can see that, in the absence of a GW wave, the mean laser power transmitted by the interferometer will be zero when $\Delta L = 0$. This operational point is called dark fringe, since it corresponds to a specific destructive interference in transmission (the full mean power is reflected back to the laser source) which is insensitive to the frequency noise. The response for a GW, however, is not linear at this operational point since P_t is quadratic with $h(t)$. For this and other reasons, GW detectors operate close to the dark fringe, by introducing a small arm length offset ΔL_{DC} on a technique called DC readout. With this offset, a small amount of mean power is transmitted by the interferometer, even in the absence of a GW. For this operational point, and considering that $\bar{L} \cdot h(t)/\lambda_0 \ll \Delta L_{\text{DC}}/\lambda_0 \ll 1$, Equation 4.55 can be approximated to:

$$P_t(t) \simeq \frac{16\pi^2 P_0}{\lambda_0^2} \cdot \left(\Delta L_{\text{DC}}^2 - \bar{L} \cdot \Delta L_{\text{DC}} \cdot h(t) - \frac{2\Delta L_{\text{DC}}^2 \cdot \delta\nu(t)}{\nu_0} \right), \quad [4.56]$$

which is linear with $h(t)$. This equation shows that the power variation due to the gravitational wave signal is increased by increasing the mean arm length of the interferometer and the laser power. For this reason, ground based GWDs have armlengths of kilometers, and several Watts at the interferometer input. The first term on this equation is the mean transmitted power impinging on the photodetector $\bar{P}_t = 16\pi^2 \Delta L_{\text{DC}}^2 \bar{P}_0 / \lambda_0^2$.

Let us now analyze noise contributions to the transmitted power $P_t(t)$, which will compete with the GW signal that one wants to measure. Already expressed on the Equation 4.56 is the laser frequency noise, whose coupling to the output power is reduced by reducing the interferometer arm mismatch.

Another source of noise is laser power noise, which can be split into two type of sources: technical power fluctuations, which will be represented by $\delta P_0(t)$,

and fundamental power fluctuations (shot noise), represented by $\delta P_{\text{SN}}(t)$. Technical power fluctuations originates from different "classical" sources (coupled mechanical noise, electronic noise, and others) and its coupling can be obtained by substituting $P_0 = \bar{P}_0 + \delta P_0(t)$ in Equation 4.56. Shot noise is a fundamental noise source originated from the quantum nature of light, which will be tackled in Section 4.5.3. Here we will instead introduce its classical approach. An ideal laser with no technical noise will emit a number of photons $N(t)$ during a time Δt that can be detected randomly on a photodetector with a probability following a Poissonian law, i.e., with mean value \bar{N} and variance $\sigma_N^2 = \bar{N}$. Let us now determine the PSD of the corresponding power fluctuations due to shot noise. Over Δt , the laser power is $P_{\Delta t} = h_p \nu_0 N(t)/\Delta t$ and has a mean value $\bar{P}_{\Delta t} = h_p \nu_0 \bar{N}/\Delta t$, with h_p being the Planck's constant. Following the Poissonian law, the variance of $P_{\Delta t}$ is related to σ_N^2 by:

$$\sigma_{P_{\Delta t}}^2 = \left(\frac{h_p \nu_0}{\Delta t} \right)^2 \sigma_N^2 = \left(\frac{h_p \nu_0}{\Delta t} \right)^2 \bar{N} = \frac{h_p \nu_0}{\Delta t} \bar{P}_{\Delta t} = \int_{1/\Delta t} h_p \nu_0 \bar{P}_{\Delta t} df. \quad [4.57]$$

The last equality shows that the PSD of the shot noise is equal to:

$$S_{\text{SN}}(f) = h_p \nu_0 \bar{P}_{\Delta t}, \quad [4.58]$$

according to Equation 4.31, and that is proportional to the mean laser power. We will give in Section 4.5.3 a more realistic quantum model of light and infer the same expression.

As an example, for a typical detected mean power of 10 mW and for a wavelength of $\lambda_0 = 1064 \text{ nm}$, one finds the ASD for the shot noise equal to $4 \times 10^{-11} \text{ W} \cdot \text{Hz}^{-1/2}$ and a relative shot noise (shot noise divided by the mean laser power) of $4 \times 10^{-9} \text{ Hz}^{-1/2}$.

We now get back to the interferometer transmitted power. To calculate the PSD of the shot noise in transmission of the interferometer, one can simply substitute the transmitted mean power, which is given by the first term in Equation 4.56, into Equation 4.58.

Finally, any noise source that causes an unwanted differential motion $\delta L_-(t)$ of the interferometer mirrors will couple in transmission of the interferometer and cannot be distinguished from a GW signal. This differential motion could be induced by seismic noise, vibrational noise, or thermal noise, for example. This effect can be taken into account by making the substitution in Equation 4.56 $(\Delta L_{\text{DC}})^2 \rightarrow (\Delta L_{\text{DC}} + \delta L_-(t))^2 \simeq \Delta L_{\text{DC}}^2 + 2\Delta L_{\text{DC}} \cdot \delta L_-(t)$, for small

$\delta L_-(t) \ll \Delta L_{\text{DC}}$. Taking into account all discussed noise contributions, Equation 4.56 can be re-written as:

$$P_t(t) \simeq \underbrace{\frac{16\pi^2 \Delta L_{\text{DC}}^2}{\lambda_0^2} \bar{P}_0}_{\text{DC gain factor}} \cdot \left(1 + \underbrace{\frac{\delta P_0(t)}{\bar{P}_0}}_{\text{technical power noise}} + \underbrace{\frac{2\delta L_-(t)}{\Delta L_{\text{DC}}}}_{\text{displacement noise}} - \underbrace{\frac{\bar{L} \cdot h(t)}{\Delta L_{\text{DC}}}}_{\text{GW signal}} - \underbrace{\frac{2\delta\nu(t)}{\nu_0}}_{\text{frequency noise}} \right) + \underbrace{\delta P_{\text{SN}}(t)}_{\text{shot noise}}. \quad [4.59]$$

The signal to noise ratio can be obtained by dividing the PSD of the GW signal and the PSD of the considered noise sources. This leads to:

$$\rho^2(f) = \frac{\bar{L}^2 \cdot S_h(f)}{\Delta L_{\text{DC}}^2 S_{\text{rpn}}(f) + 4S_{\delta L_-}(f) + 4\Delta L_{\text{DC}}^2 \cdot S_\nu(f)/\nu_0^2 + \Delta L_{\text{DC}}^2 h_p \nu_0 / \bar{P}_t}, \quad [4.60]$$

where $S_{\text{rpn}}(f)$ is the PSD of $\delta P_0(t)/\bar{P}_0$ and represents the relative power noise. In practice one has to choose the adequate armlength offset ΔL_{DC} taking into account several factors. For our simple approach, we shall consider the sensitivity to be limited by the shot noise, which is actually true in practice for frequencies higher than few hundreds of Hz. The signal-to-shot-noise ratio is then given by:

$$\rho_{\text{SN}}^2(f) = \frac{16\pi^2 \bar{P}_0 \bar{L}^2}{h_p c \lambda_0} S_h(f). \quad [4.61]$$

One can notice that it does not depend on the arm length offset ΔL_{DC} provided the latter remains small enough to keep the used approximation valid. In a realistic situation, other criteria have to be considered. For example, the transmitted power should be high enough such that the detection is not limited by the photodetector dark (electronic) noise and, at the same time, the detected power should remain below the photodetector saturation power. Further constraints are to be considered related to the defects appearing in real detectors (asymmetries, high power effects, etc).

4.3. Sensitivity enhancement with Fabry-Perot cavities

We have shown that the phase shift induced by a GW in a Michelson interferometer increases with the interferometer's mean arm length \bar{L} (see Equation

4.53). In this section, we discuss how the GW signal can be further enhanced by introducing Fabry-Perot cavities in the interferometer arms. A Fabry-Perot cavity is composed of two mirrors placed face-to-face, like shown in Figure 4.4. An intuitive explanation why they can enhance the detector sensitivity is that the light is "trapped" in the cavity for a certain amount of time, making several round trips between the mirrors, and during which it interacts with the GW. This effectively increases the interferometer's arm length and thus amplifies the GW signal.

Let us now describe the Fabry-Perot behavior, first by considering a linear cavity injected by a laser beam such as the one depicted in Figure 4.4. The mirrors are considered infinitesimally thin and partially reflective with amplitude transmission and reflection coefficients $\sqrt{T_{1(2)}}$, $\sqrt{R_{1(2)}}$, and with the sign convention explained in Figure 4.3 b. We assume that the mirrors have reflection coefficients close to 1 and no absorption or scattering losses. Therefore the energy conservation laws at the mirrors' surfaces can be written as:

$$\begin{cases} T_1 + R_1 = 1 \\ T_2 + R_2 = 1 \end{cases} \quad [4.62]$$




Figure 4.4: Representation of a Fabry-Perot cavity. The mirrors are characterized by their transmission and reflection coefficients $\sqrt{T_1}$, $\sqrt{R_1}$ and $\sqrt{T_2}$, $\sqrt{R_2}$ and they are placed at a distance L . A polarized laser is injected in the cavity. E_i , E_r , E_c and E_t are respectively the input, reflected, circulating and transmitted complex electric fields.

By using the symbols and the sign convention represented in Figure 4.4, one can write the equations for the circulating and reflected time-dependent electric fields, $E_c(t)$ and $E_r(t)$, at the first mirror position and the equation of the transmitted field $E_t(t)$ at the second mirror position. For example, the circulating field results from the sum of the input field $E_i(t)$ transmitted by the first mirror and the circulating field itself after it traveled a round-trip length of $2L$.

Following this logic, one obtains:

$$E_c(t) = \sqrt{T_1}E_i(t) - \sqrt{R_2R_1}E_c\left(t - \frac{2L}{c}\right), \quad [4.63]$$

$$E_t(t) = \sqrt{T_2}E_c\left(t - \frac{L}{c}\right), \quad [4.64]$$

$$E_r(t) = -\sqrt{R_1}E_i(t) - \sqrt{R_2T_1}E_c\left(t - \frac{2L}{c}\right). \quad [4.65]$$

We still consider the electric fields in their complex form: $E_j(t) = E_{j0}(t)e^{-2i\pi\nu_0 t}$, with $E_{j0}(t)$ its complex amplitude. In the steady-state operation, the complex amplitudes of the electric fields don't depend on time, only the propagation time-delay matters which yields:

$$E_{c0} = \sqrt{T_1}E_{i0} - \sqrt{R_2R_1}E_{c0}e^{i\frac{4\pi\nu_0 L}{c}}, \quad [4.66]$$

$$E_{t0} = \sqrt{T_2}E_{c0}e^{i\frac{2\pi\nu_0 L}{c}}, \quad [4.67]$$

$$E_{r0} = -\sqrt{R_1}E_{i0} - \sqrt{R_2T_1}E_{c0}e^{i\frac{4\pi\nu_0 L}{c}}. \quad [4.68]$$

4.3.1. Intracavity power

Let us now focus on the circulating field and study the intracavity power. By solving Equation 4.66 for E_{c0} , the intracavity field writes as:

$$E_{c0} = \frac{\sqrt{T_1}}{1 + \sqrt{R_2R_1}e^{i\varphi}}E_{i0} = \Sigma(\varphi)E_{i0}, \quad [4.69]$$

where φ is the round-trip propagation phase:

$$\varphi = \frac{4\pi\nu_0 L}{c}, \quad [4.70]$$

and $\Sigma(\varphi)$ is called the enhancement factor. The intracavity power is then defined as:

$$P_c = K|E_{c0}|^2 = K|\Sigma(\varphi)|^2|E_{i0}|^2, \quad [4.71]$$

where $K = S/c\mu_0$, with μ_0 the vacuum permeability and S the beam cross section. Therefore the squared modulus of the enhancement factor determines

the intracavity power according to the round-trip phase. A conventional way to write it is:

$$|\Sigma(\varphi)|^2 = \frac{\Sigma_{\max}^2}{1 + \frac{4\mathcal{F}^2}{\pi^2} \cos^2\left(\frac{\varphi}{2}\right)}. \quad [4.72]$$

This function is periodic and it reaches a maximum value Σ_{\max}^2 when the cosine squared function in the denominator is equal to zero i.e. for $\varphi = \pi + 2q\pi$ with q an integer (see Figure 4.5). That is called a resonance condition. Depending on the application, it is also useful to write the resonant condition in terms of the laser frequency or the cavity length by replacing φ by its expression (Equation 4.70) and by isolating either ν_{res} or L_{res} :

$$\nu_{\text{res}} = \frac{c}{4L} + q\frac{c}{2L}, \quad [4.73]$$

$$L_{\text{res}} = \frac{\lambda_0}{4} + q\frac{\lambda_0}{2}. \quad [4.74]$$

We can then introduce the free spectral range (FSR) which represents the difference between two successive resonances. Depending on the chosen variable, it takes one of the following values:

$$\Delta\varphi_{\text{FSR}} = 2\pi; \quad \Delta\nu_{\text{FSR}} = \frac{c}{2L}; \quad \Delta L_{\text{FSR}} = \frac{\lambda_0}{2}.$$

The peak value, that is the cavity power gain, is given by:

$$\Sigma_{\max}^2 = \frac{T_1}{(1 - \sqrt{R_2 R_1})^2}. \quad [4.75]$$

The other key parameter of Fabry-Perot cavities is called *finesse* and it is noted as \mathcal{F} in the denominator of Equation 4.72. Its expression as function of the mirrors reflectivity is:

$$\mathcal{F} = \frac{\pi (R_1 R_2)^{1/4}}{1 - \sqrt{R_2 R_1}}. \quad [4.76]$$

The finesse is also defined as the ratio of the periodicity to the full width at half maximum FWHM of $|\Sigma(\varphi)|^2$ (Figure 4.5). For $\sqrt{R_1}$ and $\sqrt{R_2}$ close to 1, one can see in Equation 4.76 that $\mathcal{F} \gg 1$. Its typical range in real cavities is from 10 to a maximum of 10^6 and is, in practice, limited by intracavity losses which




Figure 4.5: Squared modulus of the enhancement function as function of the round-trip phase φ for $\sqrt{R_1} = 0.9$ and $\sqrt{R_2} = 0.914$. The spacing between the resonant peaks is noted as $\Delta\varphi_{\text{FSR}}$ and their full width at half maximum (FWHM) as $\delta\varphi_{\text{FWHM}}$.

are not considered here. The higher the finesse, the sharper are the resonance peaks whose FWHM are given by:

$$\delta\varphi_{\text{FWHM}} = \frac{2\pi}{\mathcal{F}}; \quad \delta\nu_{\text{FWHM}} = \frac{\Delta\nu_{\text{FSR}}}{\mathcal{F}}; \quad \delta L_{\text{FWHM}} = \frac{\lambda_0}{2\mathcal{F}}. \quad [4.77]$$

The peak height depends also on the finesse. We can re-write Equation 4.72 for Σ_{max} for the two Fabry-Perot cavity configurations that are used in GW detectors:

- a) $\sqrt{R_2} = 1$ and $\sqrt{R_1} \lesssim 1$ almost all light is reflected by the cavity,
- b) $\sqrt{R_1} = \sqrt{R_2} \lesssim 1$ almost all light is transmitted by the cavity.

In the first case the end mirror is considered perfectly reflective and the input mirror reflectivity close to 1. Then, almost all the light is reflected by the cavity, the finesse becomes $\mathcal{F} \simeq \frac{\pi}{1-\sqrt{R_1}}$ and $\Sigma_{\text{max}} \simeq \frac{2}{1-\sqrt{R_1}} \simeq \frac{2\mathcal{F}}{\pi}$. In the second case $\mathcal{F} \simeq \frac{\pi}{T_1}$ and $\Sigma_{\text{max}} \simeq \frac{1}{T_1} \simeq \frac{\mathcal{F}}{\pi}$. At resonance, the light is entirely transmitted by the cavity and the reflection is zero. The arm cavities of GW detectors are of type a), and their advantage will be discussed in Section 4.3.4. Cavities of type b) are used to clean the spatial profile of the laser, as it will be shown in Section 4.3.3.

4.3.2. Gaussian beams

So far, we have considered plane waves, which is a useful theoretical approximation. Now we will introduce a Gaussian beam description, which describes more accurately the spatial and phase properties of real laser beams. For that, consider an electric field propagating along the z axis (the optical axis) with a complex transverse amplitude that depends on the z coordinate and on the distance $r = x^2 + y^2$. For small distances and small angles with respect to the optical axis, such a wave obeys the paraxial Helmholtz Equation (Siegman 1986). A Gaussian beam is the fundamental solution of that equation. One can show that the electric field complex amplitude is given by:

$$E(x, y, z) = E_0 \frac{w_0}{w(z)} \cdot e^{i\left(\frac{2\pi}{\lambda_0}(z-z_0) + \psi_G(z)\right)} \cdot e^{-i\frac{2\pi}{\lambda_0} \frac{x^2+y^2}{2R(z)}} \cdot e^{-\frac{x^2+y^2}{w(z)^2}}. \quad [4.78]$$

As one can see from the last exponential term, the transverse profile is a gaussian function whose radius w depends on z . The gaussian peak amplitude is $E_0 w_0/w(z)$ where w_0 is the waist of the gaussian beam, i.e., its minimum radius obtained at the position $z = z_0$. The presence of the factor $w_0/w(z)$ in the amplitude can be understood by considering that the beam power must be the same for any z so the integral on the $x - y$ plane of $|E|^2$ shall not depend on z . The first exponential term in Equation 4.78 contains the propagation phase and an additional phase called Gouy phase $\psi_G(z)$ whereas the second exponential term describes the spherical phase front of radius $R(z)$. The quantities $w(z)$, $\psi_G(z)$, $R(z)$ are defined as:

$$w(z) = w_0 \sqrt{1 + \frac{(z - z_0)^2}{z_R^2}}, \quad [4.79]$$

$$\psi_G(z) = \tan^{-1} \left(\frac{z - z_0}{z_R} \right), \quad [4.80]$$

$$R(z) = z - z_0 + \frac{z_R^2}{z - z_0}, \quad [4.81]$$

where $z_R = \pi w_0^2/\lambda_0$ is a parameter called Rayleigh range characterizing the beam divergence. Hence a Gaussian beam is entirely defined by the three parameters z_0 , w_0 and λ_0 . Figure 4.6 a) shows a schematic of the gaussian beam evolution while propagating along the z axis. Figure 4.6 b) shows a plot of the Gouy phase as function of the distance from the waist normalized by the Rayleigh range.

A gaussian beam is resonant in a Fabry-Perot cavity when the radius of curvature of the beam matches the radius of curvature of the mirrors at their position and when the phase resonance condition is satisfied. By considering




Figure 4.6: a) Beam profile of a gaussian beam along the propagation axis z for $z_0 = 0$ (left) and at a given z along the x axis (right). The Rayleigh range z_R represents the distance between the two points where the beam radius is $\sqrt{2}w_0$. b) Plot of the Gouy phase as function of z/z_R .

the phase term in Equation 4.78, the previously defined resonance condition on the z -axis writes as:

$$\varphi + 2\Delta\psi_G = \pi + 2q\pi, \quad [4.82]$$

where $2\Delta\psi_G$ is the round-trip variation of the Gouy phase.

In reality the gaussian beam in Equation 4.78 is just one solution of the paraxial Helmholtz equation and it represents the fundamental mode of the complete orthogonal set of eigenfunctions. One useful set are the Hermite-Gauss (HG) modes:

$$E_{n,m}(x, y, z) = \frac{E_0 w_0}{w(z)} H_n \left(\frac{\sqrt{2}x}{w(z)} \right) H_m \left(\frac{\sqrt{2}y}{w(z)} \right) \times e^{i \left(\frac{2\pi}{\lambda_0} (z - z_0) + (n+m+1) \psi_G(z) \right)} \cdot e^{-i \frac{2\pi}{\lambda_0} \frac{x^2 + y^2}{2R(z)}} \cdot e^{-\frac{x^2 + y^2}{w^2(z)}}, \quad [4.83]$$

where H_j are Hermite polynomials of order j . The intensity profiles of some HG modes are shown in Figure 4.7, the fundamental gaussian mode corresponds to $m = n = 0$.




Figure 4.7: Intensity profiles of twelve consecutive Hermite-Gaussian modes. Each mode is defined by the integer numbers m and n that correspond to the number of zeros along the x and y axis. From (DrBob 2012).

HG modes can also resonate in a Farby-Perot cavity with the following resonance condition:

$$\varphi + 2(n + m + 1)\Delta\psi_G = \pi + 2q\pi. \quad [4.84]$$

The resonance peak of an HG mode of order $(m + n)$ is therefore shifted compared to the fundamental mode (Equation 4.82) and this shift depends on the mode order.

4.3.3. Cavity transmission and mode cleaner cavities

The transmitted field by a cavity is obtained by substituting Equation 4.69 and Equation 4.70 in Equation 4.67:

$$E_{t0} = \sqrt{T_2} e^{i\varphi/2} \Sigma(\varphi) E_{i0}. \quad [4.85]$$

The cavity transmission coefficient T_C is defined by the ratio between the transmitted power and the input power:

$$T_C(\varphi) = \left| \frac{E_{t0}}{E_{i0}} \right|^2 = \frac{T_{C,\max}}{1 + \frac{4\mathcal{F}^2}{\pi^2} \cos^2\left(\frac{\varphi}{2}\right)}. \quad [4.86]$$

Similarly to the intracavity power, the transmitted power as function of the round-trip phase φ is periodic, exhibiting a peak when the resonant condition is satisfied. The maximum transmitted power is given by $T_{C,\max} \times P_i$ with P_i the laser input power. If $\sqrt{R_2} = \sqrt{R_1}$ and $\sqrt{T_2} = \sqrt{T_1}$, i.e. configuration b) in Section 4.3.1, then $T_{C,\max} = T_2 \Sigma_{\max}^2 = 1$. That means that the cavity is transparent at resonance.

When one considers a realistic input beam with a given spatial profile, its electric field E_i can be written as a linear superposition of the cavity HG eigenmodes $E_{m,n}$:

$$E_i = \sum_{m,n} c_{m,n,i} \times E_{m,n}. \quad [4.87]$$

By scanning the laser frequency or the cavity length one observes a serie of resonances, each one corresponding to a specific spatial mode (see Equation 4.84). This allows a direct measurement of the modal content of the incident beam and if one could keep the cavity in a particular resonant condition with the laser, only one mode would be transmitted. For example, with a resonance on the fundamental mode, the transmitted power is $|E_t|_{0,0}^2 \propto |c_{0,0,i}|^2$ while the transmitted power of higher orders components is close to zero: $|E_t|_{(m,n) \neq (0,0)}^2 \propto \frac{|c_{m,n,i}|^2}{\mathcal{F}^2} \simeq 0$. In this case the cavity acts as a spatial mode-cleaner, which is used in GWs to ensure a close to pure, stable fundamental Gaussian mode in order to maximize contrast, and minimize noise coupling from higher order modes to the output power. A common technique to keep the resonant condition is discussed in Section 4.3.5.

4.3.4. Cavity reflectivity

From energy conservation, the laser power that is not transmitted by the Fabry-Perot cavity is reflected (in the absence of loss). The complex reflection coefficient of the cavity is given by $R(\varphi) = -E_{r0}/E_{i0}$ (the (-) sign is due to the convection of Figure 4.4) where E_{r0} can be retrieved from equations 4.68-4.69. For the specific configuration where $\sqrt{R_2} = 1$ and $\sqrt{T_2} = 0$, i.e. configuration a) mentioned in Section 4.3.1, the transmission is zero (according to Equations

4.85-4.86) and the reflectivity $|R(\varphi)|^2$ is equal to 1. Therefore the incident beam is fully reflected for any value of the round-trip phase.

The important modifications happen at the phase of the reflected field around resonance. For small and slow deviations $\delta\varphi$ (slow corresponds to Fourier frequencies $f \simeq 0$) of the round-trip phase around resonance, the phase of the reflected wave with respect to the input wave is given by:

$$\phi_r = \text{Arg}(R(\varphi)) \simeq \pi + 2 \tan^{-1} \left(\frac{\mathcal{F}}{\pi} \delta\varphi \right). \quad [4.88]$$

In Figure 4.8a) one can see the plot of the reflected phase for different values of finesse. The slope around $\delta\varphi = 0$ is equal to $2\mathcal{F}/\pi$. This is the main result that explains the enhancement of sensitivity in a Michelson interferometer with resonant Fabry-Perot cavities in its arms. Indeed, a small phase variation due to a change of the arm length is amplified by the factor $2\mathcal{F}/\pi$.




Figure 4.8: a) Phase of the reflected field ϕ_r as function of the phase deviation from resonance $\delta\varphi$ for three values of finesse and for $f \simeq 0$. b) Transfer function $H_r(f)$ for the same three finesse values. The pole frequencies correspond to an arbitrary cavity length of 7.5 km.

For $f \neq 0$, the Fabry-Perot behaves as a first-order low-pass filter with a pole frequency $f_p = \frac{\delta\nu_{\text{FWHM}}}{2}$ where $\delta\nu_{\text{FWHM}}$ is given in Equation 4.77. The corresponding transfer function is:

$$\mathcal{H}_{r,\varphi}(f) = \frac{\phi_r}{\varphi} \simeq \frac{2\mathcal{F}/\pi}{1 + i\frac{f}{f_p}}. \quad [4.89]$$

Its module, i.e. the gain, is shown in Figure 4.8b) for different values of the finesse. One can see that the higher the finesse is, the higher is the low frequency gain but lower is the amplification bandwidth. A trade-off between gain and

bandwidth has to be done in defining the cavity parameters according to the detection goals. On the other hand, a small phase variation ϕ_L introduced by the injected laser will be reflected following the transfer function $\mathcal{H}_{r,L}(f)$:

$$\mathcal{H}_{r,L}(f) = \frac{\phi_r}{\phi_L} = -\frac{1 - i\frac{f}{f_p}}{1 + i\frac{f}{f_p}}. \quad [4.90]$$

which describes a phase-shift across the cavity pole.

4.3.5. Cavity as a frequency reference: feedback loop

We will now describe how to lock the laser frequency, which is fluctuating, on a cavity resonance using a feedback control loop. For this we assume a cavity in which its length is extremely stable so that the cavity resonance can be used as a frequency reference, in order to determine the laser frequency noise.

Figure 4.9 shows a block diagram describing the feedback control loop. First, a laser emits on an optical frequency ν_0 to which the frequency noise $\delta\nu_n$ with respect to the cavity resonance is added. At the cavity input, the total laser frequency fluctuations $\delta\nu$ are compared with the cavity resonance, and their difference will result in what is called error signal V_e . An electronic servo amplifies and filters the error signal to produce the correction signal V_c which is sent to the laser frequency actuator. Let us call α the response of the laser




Figure 4.9: Block diagram describing a feedback control loop to stabilize the laser frequency to a cavity resonance.

frequency actuator, β the response of the error readout scheme including the

cavity, and let us assume that these responses are linear. The error signal V_e is inverted to create a negative back action. It is then sent to the electronic servo which has a transfer function H through which it generates the correction signal V_c . The parameters V_c , $\delta\nu$ and ϵ are linked by the following coupled equations:

$$V_c = -H \times V_e, \quad [4.91]$$

$$\delta\nu = \alpha \times V_c + \delta\nu_n, \quad [4.92]$$

$$V_e = \beta \times \delta\nu = \beta \times \delta\nu_n - \alpha\beta H \times V_e. \quad [4.93]$$

Solved, they give:

$$V_e = \frac{\beta \times \delta\nu_n}{1 + \alpha\beta H}. \quad [4.94]$$

H is designed such that, in the bandwidth of stabilization, $\alpha\beta H \gg 1$. Also, when $|\alpha\beta H| \simeq 1$, one needs $\arg(\alpha\beta H)$ far from 180° in order to avoid any loop oscillation, i.e. non zero denominator in Equation 4.94. For high loop gain, $V_e \simeq \frac{\delta\nu_n}{\alpha\beta H} \rightarrow 0$ which shows that the laser remains resonant with cavity. The resonance condition 4.73 being fulfilled, one can write:

$$\left| \frac{\delta\nu}{\nu_0} \right| = \left| \frac{\delta L}{L} \right|, \quad [4.95]$$

in the lock frequency bandwidth. One difficulty though is to generate a linear error signal. A well known adequate technique is the Pound Drever Hall (PDH)(Drever 1983). A typical setup of this technique is depicted in Figure 4.10. Before entering the cavity, the laser beam passes through a phase modulator driven by an RF (radio frequency) generator. For small modulation amplitudes only two sidebands are generated. If they are sufficiently apart from the carrier frequency ν_0 , the corresponding reflected field is not phase shifted when ν_0 is close to the cavity resonance whereas the carrier field undergoes an important phase change (see Figure 4.8a)). A photodetector in reflection of the cavity is then used to detect the beat note between the carrier leaving the cavity and the RF sidebands being directly reflected (not phase shifted). When the photodetector signal is demodulated at the RF modulation frequency, a linear error signal around the resonance is generated which is zero exactly at resonance, and has different signs depending if the laser frequency is above or below resonance. This error signal is plotted in Figure 4.11.

4.4. Power recycled interferometer

Now that a simple Michelson interferometer and optical cavities were introduced, we can consider an almost complete layout of a gravitational wave




Figure 4.10: Schematic of the experimental implementation of the PDH technique to lock the laser frequency to a cavity resonance. On the top of the figure one can see the sideband picture representation of the optical carrier frequency of the laser ν_0 and the RF sidebands $\nu_0 \pm \nu_{RF}$. More details in the text.




Figure 4.11: Error signal obtained by the PDH technique around the cavity resonance frequency. The PDH scheme generates an error signal (in Volts) which is zero at resonance, i.e., when $\delta\nu = 0$.

detector, with its most important characteristics. This layout is shown on Figure 4.12 in which we have included:

- a power recycling mirror which reflects the light from the bright port (interferometer reflection port) back to the interferometer, effectively increasing the power sent to the Michelson interferometer by a gain of $G \simeq 50$ (and consequently also increasing the transmitted power by a factor of G), and
- a Fabry-Perot cavity in each arm.

With those modifications, the PSD of the signal to noise ratio expression in Equation 4.60 will be:

$$\rho^2(f) = \frac{\bar{L}^2 \cdot |\mathcal{H}_{r,\varphi}|^2 \cdot S_h(f)}{\frac{4\mathcal{F}^2}{\pi^2} \Delta L_{DC}^2 \left(S_{\text{rpn}}(f) + 4|\mathcal{H}_{r,L}|^2 \frac{S_\nu(f)}{\nu_0^2} + \frac{h_p \nu_0}{G P_t} \right) + 4|\mathcal{H}_{r,\varphi}|^2 S_{\delta L_-}}. \quad [4.96]$$

With respect to Equation 4.60, the last equation sees the contribution of the cavity responses $\mathcal{H}_{r,\varphi}(f)$ and $\mathcal{H}_{r,L}(f)$, an enhancement of the arm length offset ΔL_{DC} by a factor $2\mathcal{F}/\pi$ and the enhancement of the incident power on the interferometer by the power recycling factor G .

Also are included in the figure triangular cavity (input mode cleaner) with the main role to increase the purity of the beam sent to the interferometer,¹ and frequency and power stabilization control loops to reduce their corresponding noise contribution. These will be detailed in the next two sections.

4.4.1. Frequency stabilization

As previously shown, frequency noise can limit the sensitivity of GWDs. It is therefore necessary to implement a feedback control loop to stabilize the laser frequency. In these detectors, one of the main control loops uses the interferometer's common mode as a frequency reference, as illustrated in Figure 4.12. To the first order of ΔL and according to Equation 4.52, only the phase of the reflected beam senses the interferometer:

$$\bar{\phi}(t) \simeq \frac{4\pi\nu_0\bar{L}}{c} - \frac{4\pi\nu_0\Delta L}{c} h(t) - \phi_L \left(t - \frac{2\bar{L}}{c} \right). \quad [4.97]$$

in which we can obviously neglect the GW effect to obtain:

$$\bar{\phi}(t) \simeq \frac{4\pi\nu_0\bar{L}}{c} - \phi_L \left(t - \frac{2\bar{L}}{c} \right). \quad [4.98]$$

1. The input mode cleaner also reduces the beam pointing noise which is responsible for coupling energy into high order modes.




Figure 4.12: Schematics of a simplified gravitational wave detector, composed of a Michelson interferometer with cavities in its arms and a power recycling mirror. The figure also shows the scheme of a power stabilization in which a photodetector senses the fluctuations of a pick-off beam, the scheme of frequency stabilization in which the frequency fluctuations are read in reflection of the interferometer, and finally a triangular mode cleaner cavity.

Therefore, the interferometer reflects the incident beam as a single arm of length \bar{L} . By placing identical arm cavities, and employing the PDH technique, one senses the frequency noise $\delta\nu$ with respect to a virtual cavity of length \bar{L} . After stabilization, within the control loop bandwidth, the residual frequency fluctuations follow Equation 4.95 and one gets:

$$\frac{S_\nu(f)}{\nu_0^2} \simeq \frac{S_{\delta L_+}(f)}{\bar{L}^2}, \quad [4.99]$$

where $S_{\delta L_+}(f)$ is the PSD of the sum of the length fluctuations from both interferometer's arms. Since (i) $S_{\delta L_+}(f) \simeq S_{\delta L_-}(f)$ (most of the length fluctuations are uncorrelated between the two arms), (ii) $\mathcal{F} \times \Delta L_{\text{DC}} \ll \bar{L}$ (iii) $|\mathcal{H}_{r,\varphi}|$ is of the same order of magnitude as $|\mathcal{H}_{r,L}|$ in the frequency band where $S_{\delta L_-}$ is dominant, we can neglect the frequency noise term in Equation 4.96 with respect to the length fluctuations. In a perfectly symmetric interferometer, frequency noise does not couple to the output port. The coupling factor due to unequal armlengths is of the order of $\mathcal{F} \times \Delta L_{\text{DC}} / \bar{L} \ll \mathcal{F} \times \lambda_0 / \bar{L} \simeq 3 \times 10^{-7}$. In practice, other asymmetries (mainly due to the finesse difference between the two arms) are to be taken into account and the coupling factor is of the order of 10^{-2} . This value is small enough for the stabilized frequency noise to not affect the detector sensitivity.

4.4.2. Power stabilization

Another necessary stabilization control loop in GWDs is the power stabilization one. The usual technique for power stabilization consists in sensing the power fluctuations of a pick-off beam directly with a photodetector, and use a feedback control loop (similar to the one described in Section 4.3.5) like shown on Figure 4.12. Let R be the power ratio of the pick off, and P the power of the main beam in transmission of the phase modulator. In the time domain, the power fluctuations δP_d sensed by the power stabilization photodetector are:

$$\delta P_d = R(\delta P_{\text{tech}} + \delta P_{\text{corr}}) + \delta P_{\text{SN}}(R\bar{P}), \quad [4.100]$$

where δP_{tech} represents the technical power fluctuations of the main beam, δP_{corr} are the power corrections injected by the stabilization loop and $\delta P_{\text{SN}}(R\bar{P})$ is the shot noise corresponding to the mean power $R\bar{P}$. In the bandwidth of stabilization, $\delta P_d \simeq 0$, the residual power fluctuations of the main beam are given by:

$$\delta P = \delta P_{\text{tech}} + \delta P_{\text{corr}} + \delta P_{\text{SN}}(\bar{P}) = -\frac{\delta P_{\text{SN}}(R\bar{P})}{R} + \delta P_{\text{SN}}(\bar{P}). \quad [4.101]$$

In the last equation, both shot noises are uncorrelated and since $R \ll 1$, we have in terms of PSD:

$$S_{\delta P}(f) = \frac{h_p \nu_0 R \bar{P}}{R^2} + h_p \nu_0 \bar{P} \simeq \frac{h_p \nu_0 \bar{P}}{R}. \quad [4.102]$$

Finally, the PSD of the relative power noise will be given by:

$$S_{\text{rpn}}(f) = \frac{h_p \nu_0}{R\bar{P}}. \quad [4.103]$$

This equation shows that, in order to reduce (and possibly neglect) S_{rpn} in Equation 4.96, the pick off power $R\bar{P}$ needs to be much larger than the power transmitted by the interferometer, i. e. $R\bar{P} \gg G\bar{P}_t$. In practice, there is one to two order of magnitudes between both.²

² In practice, this is true once the transmitted beam is filtered by an additional cavity called Output Mode Cleaner (OMC) which transmits the matched fundamental modes between the two arms.

4.4.3. Shot noise sensitivity curve

Thanks to the power and the frequency stabilization control loops, Equation 4.96 reduces to:

$$\rho^2(f) = \frac{\bar{L}^2 \cdot |\mathcal{H}_{r,\varphi}(f)|^2 \cdot S_h(f)}{|\mathcal{H}_{r,\varphi}(f)|^2 \times 4S_{\delta L_-}(f) + \frac{4\mathcal{F}^2}{\pi^2} \Delta L_{\text{DC}}^2 h_p \nu_0 / (G\bar{P}_t)}. \quad [4.104]$$

The signal-to-shot-noise ratio of the power recycled detector with arm cavities will be given by:

$$\rho_{\text{SN}}^2(f) = \frac{64\mathcal{F}^2 G\bar{P}_0 \bar{L}^2}{h_p c \lambda_0} \left| \frac{1}{1 + i \frac{f}{f_p}} \right|^2 S_h(f). \quad [4.105]$$

By setting the signal-to-shot-noise ratio equal to 1, i.e., $\rho_{\text{SN}}^2(f) = 1$, we infer the PSD of the shot noise limit in strain sensitivity:

$$S_{h,\text{SN}}(f) = \frac{h_p c \lambda_0}{64\mathcal{F}^2 G\bar{P}_0 \bar{L}^2} \left(1 + \frac{f^2}{f_p^2} \right). \quad [4.106]$$

Note that the sensitivity of the interferometer can be improved by increasing the interferometer power. In the following sections we will calculate the coupling of length noise and quantum noise to the strain sensitivity of a GWD, and finally calculate the full sensitivity of the detector described in this last section (Figure 4.12).

4.5. Other noise contributions

4.5.1. Length noise contributions

As already discussed, and also shown by Equation 4.104, differential length noise $S_{\delta L_-}$ can limit the interferometer sensitivity. In the following, we will estimate the contributions of different sources to differential length noise by adopting a simple approach based on the harmonic oscillator model.




Figure 4.13: Representation of an harmonic oscillator as a block of mass attached to a spring. x is the mass position, x_0 is the wall position and l_0 is the length of the unloaded spring at rest.

4.5.1.1. Harmonic oscillator

Figure 4.13 illustrates a harmonic oscillator which is represented by a point like mass m attached to a moving wall by a spring with spring constant k and with unloaded length at rest l_0 . This mass can be displaced along the spring axis by applying an external force F_0 or by changing the position of its attachment point x_0 which we will consider to be time dependent. We assume that this oscillator is damped, with a fluid damping coefficient proportional to the velocity. The dynamics of this damped harmonic oscillator subject to an external driving force F_0 can be described with Newton's second law in the time domain as:

$$m \frac{d^2 x(t)}{dt^2} = F_0 - k(x(t) - x_0(t) - l_0) - m\gamma \frac{dx(t)}{dt}. \quad [4.107]$$

The second term after the equality is the restoring spring force, which is proportional to the spring extension. The last term represents the damping, with γ being the viscous damping coefficient. Assuming that the system has a linear response to a sinusoidal external force, one can obtain the equation of motion in the frequency domain via a Fourier Transform. The system's frequency response is characterized by the mechanical susceptibility $\chi(f)$, in units of m/N , given by:

$$\chi(f) = \frac{x(f)}{F_0(f) + kx_0(f)} = \frac{1/m}{\omega_0^2 - \omega^2 + i\frac{\omega\omega_0}{Q}}, \quad [4.108]$$

where $\omega = 2\pi f$ is the angular frequency, $\omega_0 = 2\pi f_0 = \sqrt{k/m}$ is the fundamental angular resonance frequency, and $Q = \omega_0/\gamma$ is the mechanical quality factor. Equation 4.108 shows that the mass m can similarly be excited by the force F_0 or the motion of the fixation point x_0 . Figure 4.14 shows a plot of the magnitude of the mechanical susceptibility as a function of frequency. From this plot, one can see that the system response is frequency independent for low frequencies ($f \ll f_0$), it has a peak response at resonance, and then the

response decreases proportionally to $1/f^2$.

The momentum p and the total mechanical energy (Hamiltonian H) of the




Figure 4.14: Magnitude of the mechanical susceptibility of a harmonic oscillator with a fundamental resonance frequency f_0 equal to 1 kHz and a quality factor Q as a function of the Fourier frequency f .

harmonic oscillator are given by:

$$p = m \frac{dx}{dt}, \quad [4.109]$$

$$H = \frac{p^2}{2m} + m \frac{\omega_0^2 x^2}{2}, \quad [4.110]$$

where (x, p) represents the conjugate canonical variables (Landau 1976).

The quantization of the harmonic oscillator can be inferred by associating to (x, p) the conjugate observables (\hat{x}, \hat{p}) obeying the commutation relationship:

$$[\hat{x}, \hat{p}] = i\hbar \quad ; \quad \hbar = \frac{h_p}{2\pi}, \quad [4.111]$$

and the observable energy is given by the operator Hamiltonian:

$$\hat{H} = \frac{\hat{p}^2}{2m} + m \frac{\omega_0^2 \hat{x}^2}{2}. \quad [4.112]$$

With some algebra (Cohen-Tannoudji 1973), one can show that by defining the observable:

$$\hat{N} = \frac{1}{2} \left(\sqrt{\frac{m\omega_0}{\hbar}} \hat{x} + i \frac{1}{\hbar m \omega_0} \hat{p} \right) \left(\sqrt{\frac{m\omega_0}{\hbar}} \hat{x} - i \frac{1}{\hbar m \omega_0} \hat{p} \right), \quad [4.113]$$

the Hamiltonian can then be written as:

$$\hat{H} = \hbar\omega_0 \left(\hat{N} + \frac{1}{2} \right), \quad [4.114]$$

and that \hat{N} has integer eigenvalues $\{n \geq 0\}$. This shows that the harmonic oscillator has a set of discrete energy states given by (see Figure 4.15):

$$E_n = \hbar\omega_0 \left(n + \frac{1}{2} \right). \quad [4.115]$$

A fundamental comment is to be made here: the lowest energy state has a non zero energy $E_0 = \hbar\omega_0/2$ due to the Heisenberg uncertainty principle. We'll show later that this is responsible for quantum noise in the GWD.




Figure 4.15: Quantized energy states of a harmonic oscillator of resonance frequency ω_0 . The minimum energy state is non zero.

This summarizes all the properties of a classical and quantum harmonic oscillator that we need to set a simple description of different sources of differential length noise in GW detectors and how to mitigate them.

4.5.1.2. Seismic noise

In this section, we will analyse the seismic noise contribution to GWDs. We will refer to seismic noise as all types of vibrational noise that undergo the mirrors through their contact to the ground. The GWD mirrors are in free

fall at the detection bandwidth and are suspended as pendulum mirrors. In the frequency band of interest, a pendulum mirror behaves as an harmonic oscillator under small oscillations with a spring constant $k = mg/l_0$ where g stands for the local gravitational acceleration and l_0 its length. The pendulum acts like a passive seismic attenuator for frequencies larger than the fundamental resonance frequency (see Figure 4.14). Let us now calculate the horizontal displacement δx of the mirror substrate illustrated on Figure 4.16 a). For fre-




Figure 4.16: Representation of coupling from seismic noise displacement δx_{sis} at the top of a pendulum mirror to mirror displacement δx for: a) single pendulum driven by horizontal seismic noise, b) single pendulum driven by vertical seismic noise, and c) a chain of n pendulums.

quencies $f \geq 1$ Hz, the ASD of the horizontal acceleration noise $\delta a(f)$ and of the seismic horizontal displacement ground noise $\delta x_{\text{sis,h}}$ are typically equal to:

$$\delta a(f) \simeq 4 \times 10^{-6} \text{ ms}^{-2}/\sqrt{\text{Hz}}, \quad [4.116]$$

$$\delta x_{\text{sis}}(f) \simeq 10^{-7} \left(\frac{1 \text{ Hz}}{f} \right)^2 \text{ m}/\sqrt{\text{Hz}}. \quad [4.117]$$

The suspension for a Virgo mirror has a length of 70 cm, which will lead to a fundamental resonance frequency of 0.6 Hz. For an ideal suspension, the mirror's horizontal displacement $\delta x(f)$ can be calculated by using the mechanical susceptibility (Equation 4.108). This will lead to:

$$\delta x_{\text{h}}(f) = |\chi(f)| \times m\omega_0^2 \delta x_{\text{sis}}(f), \quad [4.118]$$

which, for $f \gg f_0$, can be approximated to:

$$\delta x_{1\text{pend,h}}(f) \simeq \frac{f_0^2}{f^2} \delta x_{\text{sis}}(f) \simeq 10^{-7} \left(\frac{f_0}{f} \right)^2 \times \left(\frac{1 \text{ Hz}}{f} \right)^2 \text{ m}/\sqrt{\text{Hz}}. \quad [4.119]$$

Cascading two pendulum suspensions enhances the attenuation even further. The mirror displacement $\delta x_{2\text{pend}}$ now can be calculated by inserting the attenuated displacement of the second mirror base (Equation 4.119) into Equation 4.118. This will lead to:

$$\delta x_{2\text{pend,h}}(f) \simeq \left(\frac{f_0}{f}\right)^4 \delta x_{\text{sis}}(f) \simeq 10^{-7} \left(\frac{f_0}{f}\right)^4 \times \left(\frac{1\text{ Hz}}{f}\right)^2 \text{ m}/\sqrt{\text{Hz}}. [4.120]$$

This equation can be generalized for a system with n pendulums (see Figure 4.16c)) as:

$$\delta x_{n\text{pend,h}}(f) \simeq \left(\frac{f_0}{f}\right)^{2n} \delta x_{\text{sis}}(f) \simeq 10^{-7} \left(\frac{f_0}{f}\right)^{2n} \times \left(\frac{1\text{ Hz}}{f}\right)^2 \text{ m}/\sqrt{\text{Hz}}. [4.121]$$

As shown in Figure 4.16b) seismic noise driving the top base of the mirror vertically (δz_{sis}) will also couple to horizontal mirror displacement by:

$$\delta x_v(f) = \alpha_c \times \delta z_{\text{sis}}(f), [4.122]$$

with a coupling coefficient of approximately $\alpha_c \simeq 1\%$ mainly related to mechanical asymmetries. To mitigate its effect, vertical isolation is required. As for the pendulums, vertical filters are based on passive vertical suspensions behaving like harmonic oscillators in the frequency of interest. In the case of the Virgo detector, the number of horizontal filters pendulums is $n = 7$ and the vertical filters are $n = 5$. The ASD of the total horizontal mirror displacement is an uncorrelated sum between the displacement due to horizontal and vertical seismic noise. In the case of the Virgo detector is given by:

$$\delta x(f) = \sqrt{\left[\delta x_{\text{sis}}(f) \times \left(\frac{f_0}{f}\right)^{14}\right]^2 + \left[\delta z_{\text{sis}}(f) \times \left(\frac{f_0}{f}\right)^{10} \times \alpha_c^5\right]^2}. [4.123]$$

The PSD of the equivalent strain due to seismic noise can again be computed by setting the signal to (seismic) noise ratio in Equation 4.96 equal to 1. Considering the 4 suspended mirrors of the arm cavities of the detector, the total seismic noise projected to strain sensitivity is given by Equation 4.104:

$$S_{h,\text{sis}}(f) \simeq \frac{4 \times 4}{\bar{L}^2} \left[\delta x_{\text{sis}}(f) \times \left(\frac{f_0}{f}\right)^{14}\right]^2 + \left[\delta z_{\text{sis}}(f) \times \left(\frac{f_0}{f}\right)^{10} \times \alpha_c^5\right]^2. [4.124]$$

In practice, $\delta z_{\text{sis}}(f)$ and $\delta x_{\text{sis}}(f)$ are of the same order of magnitude.

4.5.2. Thermal noise

Another important source of differential displacement noise is thermal noise, which sets a limitation in the degree to which the suspended mirror can stay at rest with the system in equilibrium at a temperature T . In order to introduce the physics underneath the thermal noise, we will follow the Langevin approach of the Brownian motion (Langevin 1908). Let us consider one free falling mirror




Figure 4.17: Representation of a free falling mirror surrounded by residual gas particles which imprints a random momentum in the mirror position.

surrounded by a gas of particles each transferring a random momentum to the mirror, as shown in Figure 4.17. This momentum transfer results into a driving stochastic force $F_L(t)$ called Langevin Force. The gas is also responsible for a viscous damping proportional to the mirror velocity as described in Equation 4.107. For the motion of the mirror along the x axis one can simply use the damped harmonic oscillator model driven by F_L :

$$m \frac{d^2 x(t)}{dt^2} = m \frac{dv(t)}{dt} = -m\gamma v(t) + F_L(t). \quad [4.125]$$

The same equation stands for the other directions. Thermal noise is a stochastic process in which the mean value is zero at all times, i.e., $\langle F_L(t) \rangle = 0$. This reflects the fact that the mirror is equally pushed in any direction by the kicks from the gas particles. It is also a Markovian process, i.e. $\Gamma_{F_L}(\tau) = \langle F_L(t) F_L(t + \tau) \rangle = \sigma_{F_L}^2 \times \Delta t \times \delta(\tau)$ where Δt is a very short amount of time above which the correlation vanishes³ (see Figure 4.1). This equation shows that the force values at different times are uncorrelated.

The damping coefficient γ and the driving force F_L have the same physical origin, but what is the relationship between them? To answer this question

³ the delta function can be approximated by a gate function of width Δt and height $1/\Delta t$

let us now solve Equation 4.125 for the velocity, that will lead to the general solution:

$$v(t) = v_0 e^{-\gamma t} + \frac{1}{m} \int_0^t e^{-\gamma(t-t')} F_L(t') dt', \quad [4.126]$$

where $t = 0$ was chosen such as $F_L(t = 0) = 0$ and $v(t = 0) = v_0$. From this equation, one can see that the velocity mean value will be $\langle v(t) \rangle = v_0 e^{-\gamma t}$ since $\langle F_L(t) \rangle = 0^4$. In addition to that, the variance of the velocity can be calculated:

$$\begin{aligned} \sigma_v^2(t) &= \left\langle \left(\frac{1}{m} \int_0^t e^{-\gamma(t-t')} F_L(t') dt' \right)^2 \right\rangle \\ &= \frac{1}{m^2} \int_0^t dt' \int_0^{t'} dt'' \langle F_L(t') F_L(t'') \rangle e^{-\gamma(t-t')} e^{-\gamma(t-t'')}, \\ &= \frac{\Delta t \sigma_{F_L}^2}{4\gamma m^2} (1 - e^{-2\gamma t}). \end{aligned} \quad [4.127]$$

The mean kinetic energy of the mirror in thermodynamic equilibrium can now be calculated, which, according to the equipartition theorem, will result in:

$$\langle U_k(t) \rangle = \frac{1}{2} m \langle v(t)^2 \rangle \xrightarrow{t \rightarrow \infty} \langle U_k(t) \rangle = \frac{1}{2} m \sigma_v^2(t \rightarrow \infty) = \frac{k_B T}{2}, \quad [4.128]$$

with k_B being the Boltzmann's constant. Now, substituting σ_v^2 of Equation 4.127 into Equation 4.128, will lead to:

$$\sigma_{F_L}^2 = \frac{1}{\Delta t} \times 4m\gamma k_B T = \int_{\frac{1}{\Delta t}} 4m\gamma k_B T df, \quad [4.129]$$

The Langevin force PSD is then:

$$S_{F_L}(f) = 4m\gamma k_B T. \quad [4.130]$$

Taking into account the mechanical susceptibility for a free falling mirror:

$$\chi(f) = \frac{1/m}{-\omega^2 + i\omega\gamma} = \frac{-\omega - i\gamma}{m\omega(\omega^2 + \gamma^2)}, \quad [4.131]$$

4. The time average is taken over times much shorter than $1/\gamma$

one can calculate the PSD of the Brownian motion of the mirror:

$$S_x(f) = |\chi(f)|^2 S_{F_L}(f) = \frac{4k_B T}{\omega} |\text{Im}(\chi(f))|. \quad [4.132]$$

This gives the relationship between the random fluctuations of the mirror position and the dissipation process. In 1952, this equation was generalized by Callen and Welton (Callen 1951) via the famous Fluctuation-Dissipation Theorem which states that if a system dissipates energy, it must also experience noise, and the larger the dissipation, the larger the noise. The PSD of the Langevin force was then generalized to any mechanical system and any dissipation mechanism as:

$$S_{F_L}(f) = \frac{4k_B T}{\omega} \left| \text{Im} \left(\frac{1}{\chi(f)} \right) \right|, \quad [4.133]$$

which is a macroscopic expression of a microscopic phenomena that is responsible for the noise. Now we apply this to compute two dominant thermal noises in the GWD: suspension thermal noise and mirror thermal noise.

4.5.2.1. Suspension thermal noise

The pendulum mirrors in GWDs are placed under high vacuum, which reduces considerably the contribution of the viscous damping. It is observed in this case that the mechanical losses of the mirrors suspensions are dominated by anelastic processes: in the frequencies of interest, each force applied to the suspension results in a motion of constant phase delay. This effect can be incorporated in the harmonic oscillator by adding an imaginary part to the resonance frequency $\omega_0^2 \rightarrow \omega_0^2 (1 + i\phi_p)$, with ϕ_p being the so called loss angle. In this case, the mechanical susceptibility of a single pendulum is:

$$\chi_p(f) = \frac{1/m}{\omega_0^2 - \omega^2 + i\phi_p \omega_0^2}. \quad [4.134]$$

To apply the Fluctuation-Dissipation Theorem, one has for $\omega \gg \omega_0$:

$$|\text{Im}(\chi_p(f))| = \frac{1}{m} \frac{\phi_p \omega_0^2}{(\omega_0^2 - \omega^2)^2 + \phi_p^2 \omega_0^4} \simeq \frac{1}{m} \frac{\omega_0^2}{\omega^4} \phi_p. \quad [4.135]$$

Now considering 4 suspensions for each of the 4 mirrors of the interferometer, the total suspension thermal noise projected to strain sensitivity is:

$$S_{h,\text{th},\text{sus}}(f) = \frac{4 \times 4}{L^2} \cdot 4 \cdot \frac{4k_B T f_0^2}{m (2\pi)^3 f^5} \phi_p. \quad [4.136]$$

The recoil force for a suspension is almost entirely gravitational, which is conservative. The only dissipation is related to the residual elastic recoil force through which a small amount of energy is stored. Hence, ϕ_p is expected to be very small.

4.5.2.2. Mirror thermal noise

The mirror (substrate plus coating) being a multiple resonances system, cannot be modelled as a single harmonic oscillator. Instead, each resonance can on its own be considered as a harmonic oscillator. For the sake of simplicity, we assume an anelastic dissipation for all the resonances with a constant loss angle ϕ_m . Considering only the coordinate along the optical axis x on the mirror surface, the total mechanical susceptibility is:

$$\chi(f) = \sum_j \frac{1/m_j}{\omega_j^2 - \omega^2 + i\phi_m\omega_j^2} = \frac{x_{\text{mir}}}{F}, \quad [4.137]$$

where m_j is the mode j mass involving its volume, i.e. the part of the mass actually moving along x when the mode is resonating, F is the amplitude of a uniform force applied to the mirror surface, and x_{mir} is the surface displacement of the mirror. The first mirror resonance frequency is much higher than the frequencies of interest, i.e. the detector bandwidth. Hence, for all j , $\omega \ll \omega_j$:

$$|\chi(f)| \simeq \sum_j \frac{1}{m_j\omega_j^2} \quad \text{and} \quad |\text{Im}(\chi(f))| \simeq \phi_m \sum_j \frac{1}{m_j\omega_j^2} = \phi_m |\chi(f)|, \quad [4.138]$$

which are frequency independent down to $f = 0$. We can then write the static stored elastic energy when a static force is applied to the mirror as:

$$W = \frac{1}{2} F_{\text{stat}} \times x_{\text{mir}}(f=0) = \frac{1}{2} F_{\text{stat}}^2 |\chi(f=0)|, \quad [4.139]$$

where F_{stat} is the force F at zero frequency. So we have:

$$|\text{Im}(\chi(f))| = \phi_m \frac{2W}{F_{\text{stat}}^2}. \quad [4.140]$$

For the sake of simplicity, we assume that the laser beam covers the whole surface area $S_m = \pi r_{\text{mir}}^2$ of a circular mirror with radius r_{mir} . When F_{stat} , which covers the mirror's surface, is applied, the mirror compresses by $\delta\xi \simeq$

$\xi F_{\text{stat}}/(S_m Y)$, where ξ is the mirror thickness and Y is the Young modulus. Having then $W = \frac{1}{2} F_{\text{stat}} \delta \xi$, we finally obtain for the mirror's surface displacement :

$$|\text{Im}(\chi(f))| \simeq \phi_m \times \frac{\xi}{Y S_m}, \quad S_{x,\text{th,mir}}(f) = \frac{2k_B T}{\pi f} |\text{Im}(\chi(f))| \quad [4.141]$$

Now considering the 4 mirrors, the mirror's thermal noise on the detector is:

$$S_{h,\text{th,mir}}(f) = \frac{4}{L^2} \cdot 4 \cdot \frac{2k_B T}{\pi f} \cdot \frac{\xi}{Y \pi r_{\text{mir}}^2} \phi_m. \quad [4.142]$$

A more general and accurate approach can be found in (Levin 1998).

4.5.3. Quantum noise

4.5.3.1. Quantification of the electromagnetic field

The classical treatment of quantum noise presented in the earlier sections was useful for an intuitive understanding of the coupling mechanism of shot noise in an interferometer. However, a rigorous and full description of the quantum noise coupling in the interferometer and how to reduce it can only be obtained within a quantum mechanics formalism, that takes into account vacuum fluctuations. For that we will treat each mode of the electromagnetic field as a quantum harmonic oscillator. Due to the scope of this book, our description will be short. A full mathematical description with most of the equations shown in this section can be found in many books of quantum mechanics, and we used as a reference the book (Bachor-Ralph 2019).

In classical physics, a single mode electric field in a plane wave approximation, propagating in vacuum, and with no technical noise, can be decomposed into a sum of two time independent quadratures named X_1 and X_2 that are oscillating 90° out of phase with each other:

$$\mathcal{E}(t) = \mathcal{E}_0 (X_1 \cos \omega_0 t + X_2 \sin \omega_0 t), \quad [4.143]$$

to which corresponds an energy:

$$E = \kappa \left(\frac{X_1^2}{4} + \frac{X_2^2}{4} \right), \quad [4.144]$$

with κ a constant that will be defined later. By analogy with the harmonic oscillator (see Equation 4.110), one can see that X_1 and X_2 are the conjugate

canonical variables of the position and momentum in a Hamiltonian representation, scaled to be dimensionless. Hence, a similar quantification process can be applied:

$$[\hat{X}_1, \hat{X}_2] = 2i, \quad [4.145]$$

$$\hat{H} = \kappa \left(\frac{\hat{X}_1^2}{4} + \frac{\hat{X}_2^2}{4} \right) = \kappa \left(\hat{N} + \frac{1}{2} \right). \quad [4.146]$$

Two main comments are to be made here:

- As for the harmonic oscillator, the single mode electric field is characterized by a set of discrete field eigenstates with energies $E_n = \kappa (n + 1/2)$, n being a positive integer. It can be demonstrated, but it seems obvious to set $\kappa = \hbar\omega_0$ so moving from the state E_n to the state E_{n+1} corresponds to the gain of one photon energy $\hbar\omega_0$.
- The vacuum state $|n = 0\rangle$ corresponds to a non zero energy $E_0 = \hbar\omega_0/2$. In the vacuum state (subscript v), the quadratures are time dependent and characterized by the quantum variances $\sigma_{\hat{X}_{1v(2v)}}^2$ that are linked with each other by the Heisenberg uncertainty principle:

$$\sigma_{\hat{X}_{1v}} \times \sigma_{\hat{X}_{2v}} = 1. \quad [4.147]$$

Since the definition of X_1 and X_2 is invariant with respect to an arbitrary phase, $\sigma_{\hat{X}_{1v}} = \sigma_{\hat{X}_{2v}} = 1$. In a phasor diagram, the vacuum state is represented by a circle centered at (0,0), like shown in Figure 4.18.a.

4.5.3.2. *Semi-classical approach : The power quantum fluctuations*

A full quantum treatment of the electric field requires the introduction of the Heisenberg formalism and goes beyond the scope of this book. We can however adopt a semi-classical approach in which the vacuum is represented by a classical fluctuating field of zero mean. A quantum noise limited laser field is called a coherent state (Bachor-Ralph 2019), and can be formally defined as the action of the displacement operator on the vacuum state in a full quantum mechanics formalism. Such a field is represented in Figure 4.18.b. The corresponding classical field can be written as a sum of the (already introduced) carrier field (with time independent quadratures) and the vacuum field:

$$\mathcal{E}(t) = \mathcal{E}_0 (X_1 \cos \omega_0 t + X_2 \sin \omega_0 t) + \epsilon_v (X_{1v}(t) \cos \omega_0 t + X_{2v}(t) \sin \omega_0 t), \quad [4.148]$$




Figure 4.18: Quadrature space representation for a) a vacuum state, b) coherent state, and c) for different squeezed vacuum states.

where ϵ_v is the amplitude of the vacuum electric field and $X_{1v(2v)}(t)$ the corresponding fluctuating quadratures. In order for this classical field to mimic the coherent state behavior, the following conditions must be satisfied:

- the average of the quadratures for the vacuum state is zero:

$$\langle 0 | \hat{X}_{1(2)} | 0 \rangle = \langle X_{1v(2v)}(t) \rangle = 0, \quad [4.149]$$

- the variance of the quadratures for the vacuum is equal to one:

$$\sigma_{\hat{X}_{1v(2v)}}^2 = \langle 0 | \hat{X}_{1(2)}^2 | 0 \rangle = \sigma_{X_{1v(2v)}}^2 = \langle X_{1v(2v)}^2(t) \rangle = 1, \quad [4.150]$$

- the vacuum electric field ϵ_v corresponds to the vacuum energy $\langle 0 | \hat{H} | 0 \rangle = \hbar\omega_0/2$.

To the first order of $X_{1v(2v)}$, the power averaged over the light period $2\pi/\omega_0$ is:

$$\begin{aligned} P(t) &= \frac{S}{2\mu_0 c} \mathcal{E}_0^2 (X_1^2 + X_2^2) + \frac{S}{\mu_0 c} \mathcal{E}_0 \epsilon_v (X_1 X_{1v}(t) + X_2 X_{2v}(t)) \\ &= \bar{P} + \delta P_{\text{QN}}(t), \end{aligned} \quad [4.151]$$

where S is the beam section and by definition we have $X_1^2 + X_2^2 = 1$. Now, we consider the arbitrary observation time Δt required to have for the vacuum energy given by the fluctuating term of Equation 4.148:

$$\frac{\hbar\omega_0}{2} \simeq \frac{S}{2\mu_0 c} \epsilon_v^2 \Delta t \langle X_{1v}^2 + X_{2v}^2 \rangle = \frac{S}{\mu_0 c} \epsilon_v^2 \Delta t, \quad [4.152]$$

to obtain:

$$\epsilon_v = \sqrt{\frac{\mu_0 c \hbar \omega_0}{2S \Delta t}}. \quad [4.153]$$

The power fluctuations due to quantum noise δP_{QN} are then given by Equation 4.151:

$$\delta P_{\text{QN}}(t) = \epsilon_0 \sqrt{\frac{S \hbar \omega_0}{2\mu_0 c \Delta t}} (X_1 X_{1v}(t) + X_2 X_{2v}(t)), \quad [4.154]$$

and since $S_{X_{1v}}(f) = S_{X_{2v}}(f) \equiv S_{X_v}(f)$, the corresponding PSD of power noise due to quantum noise is:

$$S_{\text{QN}}(f) = \hbar \omega_0 \cdot \epsilon_0^2 \frac{S}{2\mu_0 c} (X_1^2 + X_2^2) \cdot \frac{S_{X_v}(f)}{\Delta t} = \hbar \omega_0 \bar{P}, \quad [4.155]$$

which is frequency independent. The second equality is due to the fact that:

$$\int_{1/\Delta t} S_{X_v}(f) df = \sigma_{X_v}^2 = 1 \rightarrow \frac{S_{X_v}(f)}{\Delta t} = 1. \quad [4.156]$$

We then obtained the same PSD of the shot noise that was introduced based on a classical model in Section 4.2.2 for shot noise, i.e., $S_{\text{QN}}(f, \bar{P}) = S_{\text{SN}}(f, \bar{P})$ and which shows that the vacuum field is at its origin.

4.5.3.3. *The quantum noise in the interferometer*

The vacuum field couples into optical experiments through any channel that is lossy or open. The main coupling port of vacuum fluctuations in GWDs happens via the output port of the interferometer (Kimble 2001), which is dark (or close to dark), as shown in Figure 4.19a. As we will now derive, the vacuum fluctuations enters via the interferometer output and are completely (or almost completely) reflected back to the photodetector, resulting in the readout shot noise calculated in Equation 4.59. We will also show that the vacuum fluctuations coupling from the input port of the interferometer will be completely (or almost completely) reflected by the interferometer, as well as the laser technical power fluctuations.

Let us now derive the contribution of quantum noise in the Michelson interferometer shown in Figure 4.19a. For this calculation we will assume the interferometer is operated at the dark fringe and we will ignore technical noise




Figure 4.19: a) Vacuum fluctuations (dashed orange) coupling from the from the output of a gravitational wave detector, and a coherent field at its input. b) Squeezed vacuum field being injected via the interferometer output in order to reduce its quantum noise.

sources. We start by writing the interferometer input field $E_i(t)$, which is a coherent field, as the sum of the carrier and the vacuum fluctuations:

$$E_i(t) = \mathcal{E}_0 \cos(\omega_0 t) + \epsilon_v (X_{1v} \cos(\omega_0 t) + X_{2v} \sin(\omega_0 t)). \quad [4.157]$$

For simplicity, we will omit the time dependence on the vacuum quadratures. Similarly, we write the vacuum field coupling at the dark port of the interferometer as:

$$E_d(t) = \epsilon_v (Y_{1v} \cos(\omega_0 t) + Y_{2v} \sin(\omega_0 t)). \quad [4.158]$$

The field directly reflected and transmitted by the beamsplitter at the West and North arms will be given by:

$$\begin{aligned}
E_{iW}(t) &= \frac{1}{\sqrt{2}}\mathcal{E}_0 \cos(\omega_0 t) + \frac{1}{\sqrt{2}}\epsilon_v (X_{1v} \cos(\omega_0 t) + X_{2v} \sin(\omega_0 t)) \\
&\quad + \frac{1}{\sqrt{2}}\epsilon_v (Y_{1v} \cos(\omega_0 t) + Y_{2v} \sin(\omega_0 t)) \\
&= \frac{1}{\sqrt{2}}[\mathcal{E}_0 + \epsilon_v X_{1v} + \epsilon_v Y_{1v}] \cos(\omega_0 t) + \frac{1}{\sqrt{2}}[\epsilon_v X_{2v} + \epsilon_v Y_{2v}] \sin(\omega_0 t), \\
E_{iN}(t) &= \frac{1}{\sqrt{2}}[-\mathcal{E}_0 - \epsilon_v X_{1v} + \epsilon_v Y_{1v}] \cos(\omega_0 t) + \frac{1}{\sqrt{2}}[-\epsilon_v X_{2v} + \epsilon_v Y_{2v}] \sin(\omega_0 t).
\end{aligned} \tag{4.159}$$

For the sake of simplicity, we neglect the phase shifts in the small Michelson composed of the beamsplitter and the input cavity mirror with respect to the phase shifts $(\phi_{W,r}, \phi_{N,r})$ acquired by the beam in the arm cavities. The reflected fields are then:

$$\begin{aligned}
E_{rW}(t) &= \frac{1}{\sqrt{2}}[\mathcal{E}_0 + \epsilon_v X_{1v}^{(r)} + \epsilon_v Y_{1v}^{(r)}] \cos(\omega_0 t + \phi_{W,r}) \\
&\quad + \frac{1}{\sqrt{2}}[\epsilon_v X_{2v}^{(r)} + \epsilon_v Y_{2v}^{(r)}] \sin(\omega_0 t + \phi_{W,r}), \\
E_{rN}(t) &= \frac{1}{\sqrt{2}}[-\mathcal{E}_0 - \epsilon_v X_{1v}^{(r)} + \epsilon_v Y_{1v}^{(r)}] \cos(\omega_0 t + \phi_{N,r}) \\
&\quad + \frac{1}{\sqrt{2}}[-\epsilon_v X_{2v}^{(r)} + \epsilon_v Y_{2v}^{(r)}] \sin(\omega_0 t + \phi_{N,r}).
\end{aligned} \tag{4.160}$$

Where the superscript (r) stands for the fields in reflection of the optical cavities, in which we will consider that only the phase is changed, not the amplitude. Hence $(X_{1v}(t))^{(r)} \rightarrow X_{1v}(t)$. The intracavity fields can be written for $f \ll f_p$ as:

$$\begin{aligned}
E_{cav,W}(t) &= \sqrt{\frac{\mathcal{F}}{\pi}}[\mathcal{E}_0 + \epsilon_v X_{1v} + \epsilon_v Y_{1v}] \cos(\omega_0 t + \phi_{W,cav,i(e)}) \\
&\quad + \sqrt{\frac{\mathcal{F}}{\pi}}[\epsilon_v X_{2v} + \epsilon_v Y_{2v}] \sin(\omega_0 t + \phi_{W,cav,i(e)}), \\
E_{cav,N}(t) &= \sqrt{\frac{\mathcal{F}}{\pi}}[-\mathcal{E}_0 - \epsilon_v X_{1v} + \epsilon_v Y_{1v}] \cos(\omega_0 t + \phi_{N,cav,i(e)}) \\
&\quad + \sqrt{\frac{\mathcal{F}}{\pi}}[-\epsilon_v X_{2v} + \epsilon_v Y_{2v}] \sin(\omega_0 t + \phi_{N,cav,i(e)}).
\end{aligned} \tag{4.161}$$

where $\phi_{W,cav,i(e)}$ and $\phi_{N,cav,i(e)}$ are the intracavity beams' phases impinging on the input(end) mirrors. The detected field at the dark port is then:

$$\begin{aligned}
E_t &= \frac{1}{\sqrt{2}}E_{rW} + \frac{1}{\sqrt{2}}E_{rN} \\
&= \frac{1}{2}\mathcal{E}_0 [\cos(\omega_0 t + \phi_{W,r}) - \cos(\omega_0 t + \phi_{N,r})] \\
&\quad + \frac{1}{2}\epsilon_v X_{1v}^{(r)} [\cos(\omega_0 t + \phi_{W,r}) - \cos(\omega_0 t + \phi_{N,r})] \\
&\quad + \frac{1}{2}\epsilon_v X_{2v}^{(r)} [\sin(\omega_0 t + \phi_{W,r}) - \sin(\omega_0 t + \phi_{N,r})] \\
&\quad + \frac{1}{2}\epsilon_v Y_{1v}^{(r)} [\cos(\omega_0 t + \phi_{W,r}) + \cos(\omega_0 t + \phi_{N,r})] \\
&\quad + \frac{1}{2}\epsilon_v Y_{2v}^{(r)} [\sin(\omega_0 t + \phi_{W,r}) + \sin(\omega_0 t + \phi_{N,r})]. \quad [4.162]
\end{aligned}$$

As in Equation 4.49, we introduce the common and differential phases $\bar{\phi}_r = (\phi_{W,r} + \phi_{N,r})/2$ and $\Delta\phi_r = (\phi_{W,r} - \phi_{N,r})/2$. The field at the output can be written as:

$$\begin{aligned}
E_t &= -\mathcal{E}_0 \sin(\omega_0 t + \bar{\phi}_r) \sin \Delta\phi_r - \epsilon_v X_{1v}^{(r)} \sin(\omega_0 t + \bar{\phi}_r) \sin \Delta\phi_r \\
&\quad + \epsilon_v X_{2v}^{(r)} \cos(\omega_0 t + \bar{\phi}_r) \sin \Delta\phi_r + \epsilon_v Y_{1v}^{(r)} \cos(\omega_0 t + \bar{\phi}_r) \cos \Delta\phi_r \\
&\quad + \epsilon_v Y_{2v}^{(r)} \sin(\omega_0 t + \bar{\phi}_r) \cos \Delta\phi_r, \quad [4.163]
\end{aligned}$$

which, in the small differential phase approximation can be re-written as:

$$E_t \simeq -\mathcal{E}_0 \sin(\omega_0 t + \bar{\phi}_r) \sin \Delta\phi_r + \epsilon_v [Y_{1v}^{(r)} \cos(\omega_0 t + \bar{\phi}_r) + Y_{2v}^{(r)} \sin(\omega_0 t + \bar{\phi}_r)]. \quad [4.164]$$

Note that, as expected, the transmitted field depends only on the vacuum field coupling from the interferometer's dark port. The power at the interferometer output will be then given by:

$$\begin{aligned}
P_t(t) &= \frac{S}{\mu_0 c} \langle E_t^2(t) \rangle \\
&= \frac{S}{2\mu_0 c} \mathcal{E}_0^2 \sin^2 \Delta\phi_r - \frac{S}{\mu_0 c} \epsilon_v \mathcal{E}_0 Y_{2v}^{(r)}(t) \sin \Delta\phi_r, \quad [4.165]
\end{aligned}$$

Note that the power fluctuations are solely due to quantum noise, and depends only on the Y_2 (phase) quadrature, which is orthogonal to the input carrier at the interferometer. This coupling is also known as quantum readout noise,

since the noise couples at the same quadrature of the signal expected from a GW. Their power spectral density can be calculated as:

$$S_{P_i}(f) = \left[\frac{S^2}{\mu_0^2 c^2} \epsilon_v^2 \mathcal{E}_0^2 \sin^2 \Delta\phi_r \right] S_{Y_{2v}^{(r)}} = \hbar\omega_0 \bar{P}_t, \quad [4.166]$$

which gives the same result as in Equation 4.59. This quadrature Y_{2v} entirely couples to the output port of the detector and not to the input port. Hence the power recycling just enhances the input power by the factor G and one gets the same Equation 4.106. The strain sensitivity limited by the quantum readout noise is then:

$$S_{h,\text{QN}}(f) = \frac{h_p c \lambda_0}{64 \mathcal{F}^2 G \bar{P}_0 L^2} \left(1 + \frac{f^2}{f_p^2} \right). \quad [4.167]$$

The vacuum fluctuations will also result in the so-called quantum radiation pressure noise: the fluctuating power impinging on each suspended mirror applies a fluctuating force of quantum origin. The circulating power in the west and north arm cavities can be calculated from Equations 4.161 as:

$$P_{\text{cav}W}(t) \simeq \frac{S}{2\mu_0 c} \frac{\mathcal{F}}{\pi} (\mathcal{E}_0^2 + 2\mathcal{E}_0 \epsilon_v (X_{1v}(t) + Y_{1v}(t))), \quad [4.168]$$

$$P_{\text{cav}N}(t) \simeq \frac{S}{2\mu_0 c} \frac{\mathcal{F}}{\pi} (\mathcal{E}_0^2 + 2\mathcal{E}_0 \epsilon_v (X_{1v}(t) - Y_{1v}(t))), \quad [4.169]$$

which depends only on the quadrature of the vacuum fields that is aligned with the carrier (amplitude quadrature). The arm cavities length change due to quantum radiation pressure can be obtained from the equation of motion of a free falling mirror (and no damping) with a driving radiation pressure force $F_{\text{rp}}(t) = 2P(t)/c$:

$$\delta L_{W(N)}(t) = \frac{4}{mc} \int_{t'} \int_t P_{\text{cav}W(N)}(t) dt dt', \quad [4.170]$$

where an additional factor of 2 was inserted to account for the radiation pressure effect in both cavity mirrors. The differential displacement and its PSD can then be calculated:

$$\delta L_{\text{qrp}}(t) = \delta L_W(t) - \delta L_N(t) = -\frac{8S\mathcal{F}}{mc\mu_0\pi} \mathcal{E}_0 \epsilon_v \int_{t'} \int_t Y_{1v}(t) dt dt', \quad [4.171]$$

$$S_{\delta L_{\text{qrp}}}(f) = \frac{64S^2 \mathcal{F}^2}{(mc\mu_0\pi)^2} \frac{1}{(2\pi f)^4} \mathcal{E}_0^2 \epsilon_v^2 \cdot S_{Y_{1v}}(f). \quad [4.172]$$

From these equations, one can see that only the amplitude quadrature of the vacuum field coupling at the dark port Y_{1v} will contribute to the differential displacement since the contribution from X_{1v} results in common motion in the interferometer's arms. The strain sensitivity limited by quantum radiation pressure noise is then:

$$S_{h,\text{qrp}}(f) = \frac{4}{\bar{L}^2} \frac{4\mathcal{F}^2}{m^2\pi^6} \frac{h_P c G \bar{P}_0}{\lambda_0} \frac{1}{f^4} \quad [4.173]$$

Since quantum noise has its origin the vacuum, one can redistribute the noise uncertainty via a vacuum squeezing process (Bachor-Ralph 2019). This results in the two different ellipses illustrated on Figure 4.18.c. There one can see that the noise circle has been squeezed in one direction, resulting in a quantum state in which the uncertainty in one quadrature is reduced at the penalty of increasing the uncertainty in the orthogonal quadrature.

GWDs inject squeezed vacuum states at their output port (see Figure 4.19c) to improve their sensitivity (Ganapathy 2023). They inject the so called frequency dependent squeezing in which the quadrature that is squeezed depends on the frequency. At low frequencies, where the detector is limited by quantum radiation pressure noise, the squeezed quadrature is the one aligned with the input interferometer carrier (amplitude quadrature), i.e., Y_{1v} . At high frequencies, where the sensitivity is limited by the quantum readout noise, the squeezed quadrature is the one orthogonal to the input carrier, i.e., Y_{2v} .

4.6. Conclusion: full sensitivity of a GWD

Now that we calculated the contribution of the main noise sources to the interferometer output, we can finally compute the total interferometer strain sensitivity curve $S_{h,\text{tot}}(f)$. The total contribution from differential length noise can be calculated from seismic noise, suspension thermal noise, mirror thermal noise, and quantum radiation pressure noise (Equations 4.124, 4.136, 4.142 and 4.173):

$$S_{h,\delta L_-}(f) = S_{h,\text{sis}}(f) + S_{h,\text{th,sus}}(f) + S_{h,\text{th,mir}}(f) + S_{h,\text{qrp}}(f). \quad [4.174]$$

Hence, the total sensitivity will be given by:

$$S_{h,\text{tot}}(f) = S_{h,\text{QN}}(f) + S_{h,\delta L_-}(f). \quad [4.175]$$

As an example, we will plot the sensitivity considering current parameters of the Virgo GWD: input power of $\bar{P}_0 = 25$ W, an arm length of $\bar{L} = 3$ km,

a finesse for the arm cavities of $\mathcal{F} \simeq 400$, and the power recycling gain of $G \simeq 50$. We considered that the vertical seismic displacement noise is the same as the typical horizontal seismic displacement noise from Equation 4.117. For the thermal noise, we consider the loss angle for the pendulum and the mirror substrate to be $\phi_p \simeq 10^{-9}$ and $\phi_m \simeq 10^{-6}$. Finally, each mirror have a mass of $m = 40$ kg. The corresponding sensitivity curve is displayed in Figure 4.20 by the black curve, together with the individual contributions of the different noise sources. Even though based on simplified models, the result is quite similar to measured sensitivity curves that one can find for example in (Abbot 2020). From the figure, it is evident that the detector sensitivity is limited by different noise sources at different frequencies:

- seismic noise for $f \lesssim 5$ Hz,
- suspension thermal noise for $5 \text{ Hz} \lesssim f \lesssim 10 \text{ Hz}$,
- mirror thermal noise for $10 \text{ Hz} \lesssim f \lesssim 200 \text{ Hz}$ which, given the quality of the substrates, is mainly due to the coatings, and
- quantum readout noise for $f \gtrsim 200 \text{ Hz}$.

In current detectors, low-frequency sensitivity is in reality limited by noise from different control loops, which are not described here. In the following, we provide a non-exhaustive list of upgrades aimed at improving the sensitivity of future detectors:

- lowering the resonance frequency f_0 of the suspensions in order to shift both seismic noise and suspension thermal noise towards low frequencies. The options are, however, limited, since f_0 scales with the square root of the suspension length,
- increasing the input power to reduce the quantum readout noise contribution. This, however, enhances the contribution of quantum radiation pressure noise, whose ASD scales as $\sqrt{P_0}$ (see Equation 4.173). To mitigate this drawback, squeezing techniques have been proposed and are already implemented in current detectors (Ganapathy 2023), and
- reducing thermal noise by operating at cryogenic temperatures. While promising, this approach faces significant technical challenges (Ushiba 2021).




Figure 4.20: ASD of the total strain sensitivity of a gravitational wave detector, obtained from the simplified theoretical model developed in this book, which nevertheless is close to the sensitivity of real detectors. The noise contribution for different noise sources is shown separately. The minimum strain sensitivity is $10^{-23} \text{ Hz}^{-1/2}$ at $\simeq 200 \text{ Hz}$.

4.7. Bibliography

- Siegman, A. E. (1986). Lasers. University Science Books.
- Hans-A. Bachor, Timothy C. Ralph (2019) A Guide to Experiments in Quantum Optics. Wiley-VCH Verlag GmbH & Co. KGaA.
- 100 Years of General Relativity - Vol. 5: Advanced Interferometric Gravitational-Wave Detectors(2019). Editors : David Reitze, Peter Saulson & Hartmut Grote. World Scientific Publishing Co. Pte. Ltd.
- Drever, R.W.P., Hall, J.L., Kowalski, F.V. et al. Laser phase and frequency stabilization using an optical resonator. Appl. Phys. B 31, 97-105 (1983). <https://doi.org/10.1007/BF00702605>
- Landau, L. D., Lifshitz, E. M. (1976), Mechanics, Third Edition: Volume 1 (Course of Theoretical Physics), Butterworth-Heinemann.
- Don S. Lemons, Anthony Gythiel. Paul Langevin's 1908 paper "On the Theory of Brownian Motion" ["Sur la th orie du mouvement brownien" C. R. Acad. Sci. (Paris) 146, 530-533 (1908)]. Am. J. Phys. 1 November 1997; 65 (11): 1079-1081. <https://doi.org/10.1119/1.18725>
- Callen, H.B. and Welton, T.A.. Irreversibility and generalized noise. Physical Review, 83, 34-40 (1951). <http://dx.doi.org/10.1103/PhysRev.83.34>

- Mandel, L. and Wolf, E. (1995) *Optical Coherence and Quantum Optics*. Cambridge: Cambridge University Press.
- Born, M. and Wolf, E. (1999) *Principles of Optics: Electromagnetic Theory of Propagation, Interference, and Diffraction of Light*. Cambridge University Press, London. <http://dx.doi.org/10.1017/CBO9781139644181>
- Abbott, B.P., Abbott, R., Abbott, T.D. et al. Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA. *Living Rev Relativ* 23, 3 (2020). DOI: <https://doi.org/10.1007/s41114-020-00026-9>
- H. J. Kimble, Y. Levin, A. B. Matsko, K. S. Thorne, and S. P. Vyatchanin. Conversion of conventional gravitational-wave interferometers into quantum nondemolition interferometers by modifying their input and/or output optics. *Phys. Rev. D* 65, 022002 (2001) DOI: <https://doi.org/10.1103/PhysRevD.65.022002>
- Gertsenshtein, M.E. and Pustovoit, V.I. (1963) On the Detection of Low-Frequency Gravitational Waves. *Soviet Physics-JETP*, 16, 433-435
- D. Ganapathy and W. Jia and M. Nakano and V. Xu and N. Aritomi and T. Cullen and N. Kijbunchoo and S.E. Dwyer and A. Mullavey and et al. (LIGO O4 Detector Collaboration) (2023) Broadband Quantum Enhancement of the LIGO Detectors with Frequency-Dependent Squeezing. *Phys. Rev. X* 13, 041021
- DrBob at English Wikipedia, <https://commons.wikimedia.org/wiki/File:Hermite-gaussian.png>
- Levin, Y. (1998) Internal thermal noise in the LIGO test masses: A direct approach. *Phys. Rev. D*, 57, 659
- Cohen-Tannoudji, C. and Diu, B. and Laloë, F.(1973) *Quantum mechanics*. Collection Enseignement des Sciences.
- Ushiba, T et al. (2021) Cryogenic suspension design for a kilometer-scale gravitational-wave detector. *Class. Quantum Grav.* 38 085013. <https://doi.org/10.1088/1361-6382/abe9f3>

5

Pulsar timing arrays (PTAs)

Gilles Theureau¹ and Joseph D. Romano²

*¹LPC2E, OSUC, Univ. Orléans, CNRS, CNES, Observatoire de Paris,
Université PSL, F-45071 Orléans, France*

*²Department of Physics and Astronomy, University of Texas Rio Grande
Valley, One West University Boulevard, Brownsville, TX 78520, USA*

5.1. Pulsars as gravitational wave detectors

5.1.1. Pulsar timing array principle

A neutron star is a very compact and dense stellar remnant with a mass typically 1.5 to 2 times that of our Sun, but with a radius of just a dozen kilometers. Neutron stars are created by the gravitational collapse of a massive star (about ten times our Sun), whose outer layers are violently expelled in the form of a supernova, and whose core implodes and collapses in on itself before stabilising in the form of a degenerate gas of neutrons. Most neutron stars present rapid rotation and strong magnetic fields, producing beams of radiation that sweep through space.

The observational counterpart of a neutron star emitting a beam of radiation that sweeps across our line of sight is called a "pulsar" and it appears as a pulsating source of electromagnetic radiation. In 2025, there were more than 3,800 pulsars identified (see e.g., the ATNF catalogrotation periods range from a millisecond for the fastest pulsars to a few tens of seconds for the slowest. The former show exceptional properties as natural clocks, thanks to the stability of their rotation.

The study of pulsars has many applications, both in astrophysics and in fundamental physics. In particular they offer an alternative means of detecting and measuring very low-frequency gravitational waves. This idea was first proposed in 1977 by a Russian astronomer from the Moscow Astronomical Institute, Mikhail Sazhin (Sazhin 1978), and taken up again in 1979 by an American astronomer, Steven Detweiler (Detweiler 1979), from Yale University, who announced that measurements of the arrival times of pulsar pulses could be used to search for gravitational waves with periods of 1 to 10 years, corresponding to frequencies of 10^{-8} Hz. A few years later, in 1982, two other American astronomers, Ron Hellings and George Downs, from the Jet Propulsion Laboratory, formalised the method and gave an initial limit on the gravitational wave background using the four pulsars PSR B1133+16, PSR B1237+25, PSR B1604-00 and PSR B2045-16 (Hellings and Downs 1983). The discovery of millisecond pulsars that same year by Don Backer (Backer *et al.* 1982) (UC Berkeley), with much more stable rotation and potentially much more precise measurements, paved the way for what is now known as the Pulsar Timing Array (PTA) technique.

5.1.2. First evidence for a low-frequency gravitational-wave signal

The first detection of a gravitational signal by using an array of pulsars took place in 2023. It was carried out jointly by four independent consortia in

Europe+India (Antoniadis *et al.* 2023), the United States (Agazie *et al.* 2023), Australia (Reardon *et al.* 2023) and China (Xu *et al.* 2023), each using a different set of radio telescopes and a different pulsar array. In France, for example, the large decimetric radio telescope at Nançay (NRT) made a major contribution to this result, with almost 120,000 measurements between 2004 and 2024. At the time of writing, this announcement still needs to be confirmed by combining all the global data, to obtain greater sensitivity and better characterise the spectral and spatial distribution of the signal.

The principle of the measurement is to consider the pulsar and the Earth as test masses in free fall, sensitive to perturbations in space-time. The passage of a gravitational wave disturbs the local environment of each object and causes variations in the succession of arrival times of the observed pulses. This particular signature is analysed and separated from the other sources of noise in the data. This requires very precise modelling of the rotation and orbital motion of each pulsar, perfect control of the position of the Earth relative to the barycentre of the Solar System at the time of observation, very precise chronometric measurement instrumentation coupled to a reference atomic clock and, above all, real time correction of variations in signal propagation due to the passage through the interstellar medium.

Since the 2000s, thanks to the availability of both powerful real-time computers and fast electronics that make it possible to observe over wide frequency bands, most of the world's major radio telescopes have accumulated arrival time measurements for several tens of millisecond pulsars, with timing accuracies better than a microsecond, and even reaching tens of nanoseconds for some. Typically, PTAs are sensitive to frequencies of order $1/T$, where T is the total time span of the observations, and the gravitational wave amplitude they can detect goes as $\delta t/T$, where δt is the timing accuracy on pulse arrival times. A timing precision of the order of 100 ns, coupled with a few decades time span and a high observation cadence of a couple of weeks, make pulsar networks sensitive to frequencies between nanoHz and microHz, and can reach a gravitational signal amplitude of a few 10^{-16} .

5.1.3. *Expected source populations*

Such an amplitude of a few 10^{-16} is what we expect, for example, from the gravitational-wave signal produced by pairs of supermassive black holes (SMBHs) of a few billion solar masses orbiting at the core of large galaxies. We know that this population exists on the basis of models of the evolution of the large-scale structures of the universe and the formation of galaxies by successive mergers. Gravitational-wave measurements should provide a better understanding of the rates and characteristic times of mergers, as well as the

relative masses of black holes and their host galaxies throughout the history of the universe.

For a population of supermassive black-hole binaries (SMBHBs) in circular orbits, losing energy essentially by emitting gravitational waves, one expects a spectrum that follows a power law of the form (Phinney 2001):

$$h_c(f) = A \left(\frac{f}{\text{yr}^{-1}} \right)^{-2/3} \quad [5.1]$$

If one considers environmental effects such as dynamical friction with surrounding gas and stars, or stellar hardening (i.e. three body-interaction with individual stars), which accelerate the shrinking of the orbit, one expects a bending of the spectrum towards lower amplitude at the lowest frequencies. The same happens if one considers eccentricities instead of purely circular black holes' orbits: the power is then distributed on several harmonics at higher frequencies. Due to shorter lifetimes of pairs at frequencies above a μHz , one expects the low statistics to also produce an artificial bending at high frequencies. The phenomenological description of the spectrum detection is summarised in Fig. 5.1.

In this frequency range of gravitational waves, we are also expecting the signature of cosmological phenomena linked to the physics of the primordial universe. Cosmic strings, which are generic predictions of grand unification theories, are one-dimensional structures that form during phase transitions in the early universe. They form loops that oscillate and emit gravitational waves of greater or lesser amplitude depending on their tension or the number of discontinuities they carry. Similarly, during the period of inflation imagined to justify the current homogeneity and isotropy of the universe, quantum fluctuations in the space-time metric are amplified by the acceleration of the expansion of the universe and produce gravitational waves. The detection of this signal would provide unique proof of the existence of this epoch of inflation in the primordial universe. Finally, if turbulence existed at the time of the transition between quarks and hadrons, i.e. the formation of subatomic particles such as protons and neutrons, it would have produced bubbles whose collisions would in turn have emitted gravitational waves. The pulsar networks therefore make it possible to test the existence of this turbulence, whether it is due to a phase transition or to the existence of a pre-existing primordial magnetic field. This new measurement technique opens up a whole new astrophysics and cosmology of the origins of the universe and its major structures.




Figure 5.1: Expected spectral energy distribution from a population of SMBHBs and PTA sensitivity curve. The dashed line shows the sensitivity curve of a given PTA experiment. The blue straight line correspond to the expected spectrum for a population of super massive black hole binaries in circular orbits. The red line shows the expectation for a more realistic population, e.g. considering environmental effects and/or eccentric orbits. Arrows indicate an upper limit at the corresponding frequency, while vertical segments stand for error bars in the detection regime, where the stochastic background signal surpasses the PTA sensitivity. Credits: S.Chen 2018, PhD thesis

5.1.4. Measurement challenges and foreground noise

There are many difficulties involved in achieving a first detection. First and foremost, the entire processing chain must be mastered, from the sky to the gravitational-wave analysis of the time series collected, in order to identify all the systematics and biases present in the data. The first challenge consists in designing and implementing high-performance instrumentation at the output of the receiver of the radio telescopes, capable of processing in real time a wide frequency range (up to a couple of GHz of bandwidth). The waveform time-series are dedispersed in the Fourier domain to correct for the chromatic delays

imprinted by the passage through the charged interstellar and heliospheric media and further folded at the apparent rotation of the pulsar, given an a priori timing model. This is generally performed with a hybrid computer combining FPGA, GPU and CPU chips. Pulsar timing data is acquired at a rate of a few days to a few weeks and combined between the various instruments taking part in the international PTA programme, to form dense time-series spanning a couple or more decades.

An important step is the construction of a timing model to describe the apparent rotation of the pulsar as measured at the telescope.

We measure a series of pulsations with an observed period P due to the rotation of the neutron star, and we observe a variation δP of this period or of rotation frequency $\delta\nu$ with time. The integration of this $\delta\nu$ along the signal pathway from the pulse emission to its reception is called the timing residual and we write it as follows:

$$r(t) = \int_0^t \frac{\nu(t') - \nu_0}{\nu_0} dt' \quad [5.2]$$

In practice we assume that we know sufficiently well the pulsar, its environment, the material present along its line of sight and the Earth motion, that we can model any Doppler shift and dispersion of the signal, so that only the unknown, e.g. the gravitational perturbation we are looking for, remains. This model is deterministic and can be written as a sum of terms corresponding to each accounted effect. If $\tau_{\text{SSB}}^{\text{TM}}$ denotes the time of arrival of a given pulse at the Solar System Barycenter (SSB), one writes:

$$\begin{aligned} \tau_{\text{SSB}}^{\text{TM}} = & t_{\text{topo}} + t_{\text{clock}} - \delta D / f_{\text{obs}}^2 + \Delta R_{\odot} + \Delta\pi + \\ & \Delta S_{\odot} + \Delta E_{\odot} + \Delta R + \Delta S + \Delta E + \Delta A \end{aligned} \quad [5.3]$$

where t_{topo} is the time measured locally at the telescope's center of phase, t_{clock} is a local clock correction, $\delta D / f_{\text{obs}}^2$ is the chromatic dispersive delay due to the travel through the ISM, ΔR_{\odot} , $\Delta\pi$, ΔS_{\odot} , and ΔE_{\odot} are respectively the Römer, parallax, Shapiro, and Einstein delays inside the Solar system (the latter due essentially to the masses of the Sun and Jupiter), and ΔR , ΔS , and ΔE are the corresponding terms at the pulsar, if it is in a binary system, and the final term ΔA , called relativistic aberration, is an extra relativistic delay due to the rotation of the pulsar.

The final residual times series is obtained by taking the difference between the modeled times of arrival and the observed times of arrival. The remaining




Figure 5.2: Timing residuals for PSR J1909-3744. The top panel displays the time series after subtracting the timing model (Eq. 5.3). The bottom panel shows the "whitened" solution after modeling the different sources of stochastic noises (Eq. 5.4). In this example one obtains a 103 nanosecond rms over 15 years of data, leading to a gravitational-wave sensitivity of order $\delta t/T \sim 2.3 \times 10^{-16}$ (from (Liu *et al.* 2020))

structure in the noise (if any) contains the missing unmodeled information, e.g., a potential gravitational-wave signal (Fig. 5.2).

Then one performs a global statistical study of the combined time-series of each pulsar and decomposes them into a sum of different Gaussian processes modeled generally with power laws. The various sources of noise taken into account are the following: spin noise intrinsic to the neutron star (τ_{SN}), chromatic dispersion (τ_{DM}) and multi-propagation/scattering (τ_{SV}) noise variations related to the crossing of the interstellar medium, temporally uncorrelated white noise (τ_{WN}) due to the statistics of the pulsar emission or the measurement uncertainty, plus possible deterministic scintillation or interference events or sudden changes in the pulsar's magnetosphere, and finally reference clock noise (τ_{CN}) and Solar System Ephemeris systematics (τ_{SSE}). So, combining all these terms, the Time of Arrival (TOA) is expressed as follows:

$$\text{TOA} = \tau_{\text{SSB}}^{\text{TM}} + \tau_{\text{WN}} + \tau_{\text{SN}} + \tau_{\text{DM}} + \tau_{\text{SV}} + \tau_{\text{CN}} + \tau_{\text{SSE}} + \tau_{\text{GW}}, \quad [5.4]$$

where $\tau_{\text{SSB}}^{\text{TM}}$ is the deterministic part and the rest of the series is stochastic. The searched-for gravitational-wave signal τ_{GW} is thus one of the multiple components of the noise and we need to disentangle it from the others. Fortunately,

each has its own characteristics that helps to isolate it from the others. The spin noise is specific to each pulsar; so are the dispersion and scattering noises, which are also chromatic and depend on the observed radio frequency; reference clock noise is shared by all pulsars as a correlated monopolar signal; SSE noise is related to systematics in the position of the Solar Barycenter and is thus a correlated dipolar signal; and finally the stochastic gravitational-wave background (SGWB) noise is expected to be a correlated quasi-quadrupolar signal.

The proper gravitational-wave analysis proceeds with the analysis of the temporally correlated noise component common to all pulsars. This study takes the form of an analysis of a covariance matrix with up to $100,000 \times 100,000$ (or more) terms. The noises specific to each pulsar are grouped together on the diagonal (white noise) or in block-diagonal form (intrinsic pulsar red noise), and the correlated common signals that contain the possible quadrupolar signature are present in the cross terms. See Fig. 5.3 for a schematic representation of the covariance matrix.

In particular, the terms characterising the gravitational-wave signal depend both on the amplitude of the signal as a function of its sky position and on the antenna pattern through the so-called “overlap reduction function”. In the context of gravity described by General Relativity, and in the presence of an isotropic stochastic gravitational-wave background, the overlap reduction function has a unique and well-defined shape, known as the Hellings & Downs curve (Hellings and Downs 1983); see Sec. 5.2 for more details.

5.1.5. Status of PTA results in 2025

So far, a first evidence for the presence of a low-frequency gravitational-wave signal has been reported independently by the five active continental consortia: the North American group NANOGrav, the European and Indian groups EPTA+InPTA, the Australian group PPTA, the Chinese group CPTA and finally the MeerKAT international consortium MPTA. All of them detect common red noise, with a spatial correlation pattern compatible with a quadrupolar origin—the ultimate “smoking-gun” signature of a gravitational-wave signal (see e.g., Fig. 5.4).

The significance ranges from 2σ to 4σ , with none actually reaching the standard detection threshold of 5σ . The amplitude of the signal is about 3×10^{-15} at a frequency $f = 3 \times 10^{-8}$ Hz ($\approx 1/\text{yr}$), with a spectral index close to the canonical $-13/3$ expected for the emission gravitational waves from a population of supermassive black-hole binaries in circular orbits. If this signal is interpreted as originating from the SMBHB population, its relatively high amplitude would




Figure 5.3: Covariance Matrix. This matrix contains the correlation terms from all the TOA measurements and all the pulsars. One derives from it two main observables: the spectral energy distribution and the spatial or angular correlation of the signal. The significance of the results is estimated by calculating false alarm rates using various methods (see Sec. 5.2.2). Credits: (Taylor 2021)

favour a high galaxy merger rate, short time-scales for black-hole mergers, and a larger-than-expected mass of the central black hole relative to the bulge of its host galaxy. The effects of eccentricity and interaction with the environment are for the moment very poorly constrained, with a slight tendency for the spectrum to bend at low frequencies in this direction. As far as processes in the early Universe are concerned, the measurement imposes constraints on the cosmic string tension and on the level of turbulence developed by first-order phase transitions. Other processes would require non-standard scenarios, such as a blue-tilted inflationary spectrum or an excess in the primordial spectrum of long-wavelength scalar perturbations.

Finally, all collaborations have searched for the presence of a dominant individual source. The most significant candidate in this search was found to have a gravitational-wave frequency of 4-5 nHz. Such a signal could be generated by a supermassive binary black hole in the local Universe, compatible with a




Figure 5.4: Spectral energy distribution of the common correlated red noise. Combined results from EPTA+InPTA (blue), NANOGrav (orange) and PPTA (green) 2023 publications (Agazie *et al.* 2024).

luminosity distance of $d_L \sim 16.6(\mathcal{M}/10^9 M_\odot)^{5/3}$ Mpc, given a chirp mass \mathcal{M} in the range $[10^8 M_\odot, 10^{10} M_\odot]$ (EPTA Collaboration *et al.* 2024).

5.2. Searches for a stochastic gravitational-wave background

As mentioned in Sec. 5.1.3, a primary source of gravitational waves in the PTA band is the population consisting of millions of pairs of supermassive black holes (SMBHs) orbiting one another in the centers of merging galaxies. These SMBHs are expected to have masses of order $10^9 M_\odot$ and are separated by $\lesssim 0.01$ pc when they enter the nanohertz frequency band. Each pair emits an approximately monochromatic gravitational-wave signal, with the system slowly losing energy due to gravitational-wave emission. The binaries will proceed in this fashion for roughly a million years before they eventually merge at higher frequencies, outside of the PTA sensitivity band.

The signals from the different SMBH binaries overlap and hence interfere with one another, giving rise to an effectively random signal (a stochastic

gravitational-wave background) incident on the Earth-pulsar baselines. The stretching and squeezing of space induced by this gravitational-wave background slightly changes the arrival time of pulses from the various pulsars in the array. Although the perturbations to the pulse arrival times for each pulsar are random, they are *correlated* across pulsars in the array in a very specific fashion, allowing us to potentially extract the gravitational-wave signal from the other sources of noise (instrumental, intrinsic pulsar, disturbances from the interstellar medium, etc.) as described in Sec. 5.1.4.

For an unpolarized and isotropic distribution of gravitational waves on the sky, the correlation between the perturbations to the arrival times of pulses from two pulsars depends only on the angular separation between the two pulsars. The functional form of the expected correlation as predicted by general relativity is given by

$$\Gamma(\gamma_{ab}) = \frac{1}{2} - \frac{1}{4} \left(\frac{1 - \cos \gamma_{ab}}{2} \right) + \frac{3}{2} \left(\frac{1 - \cos \gamma_{ab}}{2} \right) \ln \left(\frac{1 - \cos \gamma_{ab}}{2} \right), [5.5]$$

where γ_{ab} is the angle between the lines of sight to two distinct pulsars labeled by a and b (see Fig. 5.5). This expression was first calculated by Hellings




Figure 5.5: The Hellings and Downs curve $\Gamma(\gamma_{ab})$ is the expected correlation in the timing residuals from a pair of pulsars with angular separation γ_{ab} induced by an unpolarized and isotropic gravitational-wave background.

and Downs in 1983 (Hellings and Downs 1983) and is appropriately called the “Hellings and Downs” curve.

An alternative derivation of the Hellings and Downs curve [5.5] was given by Cornish and Sesana in 2013 (Cornish and Sesana 2013). In their paper, Cornish and Sesana considered a single point-source of gravitational waves emitting a plane-polarized gravitational wave in the $-z$ direction. They then calculated a *pulsar-averaged correlation*, averaging over an assumed infinite collection of pairs of pulsars distributed uniformly over the sky, all having the same angular separation γ . The result of that calculation yields [5.5], with γ_{ab} replaced by γ .

This latter “pulsar average” approach is similar to what astronomers actually use when they bin together the measured correlation within a given angular separation bin. The main difference, of course, is that the number of pulsar pairs that PTAs have access to within a given angular separation bin is finite, not infinite. For example, for the NANOGrav 15-year analysis, they observed 67 pulsars (so $67 \times 66/2 = 2211$ distinct pairs), which they divided up into 15 angular separation bins. This yielded approximately 150 pairs per bin.

5.2.1. Recovery of the predicted Hellings and Downs correlation

Figure 5.6 shows the pulsar-averaged recovery of the Hellings and Downs curve for the NANOGrav 15-year analysis. The point estimates and error bars are calculated assuming maximum-a-posteriori pulsar noise parameters and a gravitational-wave background having a $-13/3$ spectral index (appropriate for binary inspiral) and an amplitude determined from an analysis assuming only a common uncorrelated red noise process (so no correlations assumed for that calculation). Specifically, the point estimates and error bars are given by

$$\hat{\Gamma}_j = \frac{\Gamma(\gamma_j) \sum_{ab \in j} \sum_{cd \in j} \Gamma_{ab} C_{ab,cd}^{-1} \rho_{cd}}{A_{\text{gw}}^2 \sum_{ef \in j} \sum_{gh \in j} \Gamma_{ef} C_{ef,gh}^{-1} \Gamma_{gh}} \quad [5.6]$$

with

$$\sigma_j^2 = \frac{\Gamma^2(\gamma_j)}{A_{\text{gw}}^4 \sum_{ef \in j} \sum_{gh \in j} \Gamma_{ef} C_{ef,gh}^{-1} \Gamma_{gh}}, \quad [5.7]$$

where j labels the angular separation bin with central angle γ_j and A_{gw}^2 is the measured amplitude of the gravitational-wave background. Here, $\rho_{ab} =$

$\delta t_a^T Q_{ab} \delta t_b$ denotes the cross-correlation of timing residual data from distinct pulsars a and b for some filter Q_{ab} (see (Hazboun *et al.* 2023)); it satisfies

$$\langle \rho_{ab} \rangle = A_{\text{gw}}^2 \Gamma_{ab}, \quad C_{ab,cd} \equiv \langle \rho_{ab} \rho_{cd} \rangle - \langle \rho_{ab} \rangle \langle \rho_{cd} \rangle. \quad [5.8]$$

Finally, $\Gamma_{ab} \equiv \Gamma(\gamma_{ab})(1 + \delta_{ab})$, with the Kronecker delta term allowing for possible autocorrelations, which are needed when evaluating $C_{ab,cd}$. The covariance matrix has contributions from the various noise sources (assumed to be uncorrelated across pulsars) and from the gravitational-wave background itself, with the latter being proportional to $\Gamma_{ac}\Gamma_{bd} + \Gamma_{ad}\Gamma_{bc}$, see (Allen and Romano 2023).




Figure 5.6: Recovery of the Hellings and Downs curve using NANOGrav’s 15-year data. See main text for details regarding the calculation of the point estimates and error bars. Figure taken from the NANOGrav 15-year analysis paper.

Since properties of the gravitational-wave signal are used in the calculation of this curve, this “recovery plot” should be thought of as demonstrating the *consistency* between the measured correlations and a model assuming the presence of a gravitational-wave background with the Hellings and Downs correlations predicted by General Relativity. As mentioned earlier, this is the “smoking gun” signature that pulsar timing astronomers were hoping to see for the measured cross-correlations. One can think of the Hellings and Downs curve as the stochastic equivalent of the deterministic “chirp signal” that LIGO data analysts observed for GW150914, which was the first direct observation of

gravitational waves from a pair of black holes. To assess detection confidence, one must appeal to a different analysis, as described in the following subsection.

5.2.2. Detection confidence

One way to assess the statistical significance of the presence of a gravitational-wave background in the observed data is to compare the value of a “detection statistic” (a particular combination of the data which is compared to a chosen threshold to determine whether or not one should claim detection of a signal) to its distribution in the presence of noise alone (the so-called “null hypothesis”). To be robust against possible mis-modeling of the pulsar noise, such a detection statistic is constructed using only cross correlations (i.e., distinct pairs of pulsars $a < b$):

$$S \equiv \sum_{a < b} \rho_{ab} w_{ab}, \quad [5.9]$$

where w_{ab} is a weight function which we will determine below so as to maximize the expected signal-to-noise $\langle S \rangle / N$ of S . Here, $\langle S \rangle$ denotes the expected value of S in the presence of a signal, and

$$N^2 \equiv \langle S^2 \rangle_0 - \langle S \rangle_0^2 \quad [5.10]$$

is the variance of S in the absence of a signal (indicated here with the subscript ‘0’ on the expectation values). If we assume that the measured correlations satisfy

$$\langle \rho_{ab} \rangle = A_{\text{gw}}^2 \Gamma_{ab}, \quad \langle \rho_{ab} \rangle_0 = 0, \quad \langle \rho_{ab}^2 \rangle_0 \equiv \sigma_{ab,0}^2, \quad [5.11]$$

then maximizing $\langle S \rangle / N$ yields the weight function up to an overall normalization factor, $w_{ab} \propto \Gamma_{ab} / \sigma_{ab,0}^2$. If we fix this overall normalization by setting $N^2 = 1$ yields, then the measured signal-to-noise ratio is given by

$$S/N = S = \frac{\sum_{ab} \rho_{ab} \Gamma_{ab} / \sigma_{ab,0}^2}{\sqrt{\sum_{ab} \Gamma_{ab}^2 / \sigma_{ab,0}^2}}. \quad [5.12]$$

Figure 5.7 shows 1 minus the cumulative distribution function (CDF) (the so-called “ p -value”) of the noise-marginalized S/N , for the null distribution of S/N detection statistic calculated three different ways: (i) by introducing artificial phase shifts to the timing residual data; (ii) via simulations, and (iii)

via an analytic expression for S/N (see (Hazboun *et al.* 2023)). The noise-marginalization is performed to take into account the uncertainty in the noise parameters which enter the calculation of S/N . From these plots, one can read off the probability that the measured noise-marginalized mean value of S/N of 4.9 was due to noise alone is approximately 1 in 10,000.




Figure 5.7: Determination of detection confidence for NANOGrav’s 15-year data. Plotted is 1 minus the cumulative distribution function (CDF) of the the null distribution of the noise-marginalized S/N calculated three different ways, as described in the main text. The probability that the measured noise-marginalized mean value of S/N of 4.9 was due to noise alone is approximately 1 in 10,000. Figure taken from the NANOGrav 15-year analysis paper.

This is a small probability, but it is not at the 1 in 3.5 million (i.e., 5σ) level needed in the physical sciences to claim detection. Hence, one says that several PTAs across the globe have observed “evidence for a gravitational-wave background”, not that they have “detected” the background.

5.3. Bibliography

- Agazie, G., Antoniadis, J., Anumarlapudi, A., Archibald, A. M. *et al.* (2024), Comparing Recent Pulsar Timing Array Results on the Nanohertz Stochastic Gravitational-wave Background, *Astrophysical Journal*, 966(1), 105.
- Agazie, G., Anumarlapudi, A., Archibald, A. M. *et al.* (2023), The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background, *Astrophysical Journal Letters*, 951(1), L8.

- Allen, B., Romano, J. D. (2023), Hellings and Downs correlation of an arbitrary set of pulsars, *Phys. Rev. D*, 108, 043026.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.108.043026>
- Antoniadis, J., Arumugam, P., Arumugam, S. *et al.* (2023), The second data release from the European Pulsar Timing Array: III. Search for gravitational wave signals, *Astronomy & Astrophysics*, 678, A50.
URL: <http://dx.doi.org/10.1051/0004-6361/202346844>
- Backer, D. C., Kulkarni, S. R., Heiles, C., Davis, M. M., Goss, W. M. (1982), A millisecond pulsar, *Nature*, 300(5893), 615–618.
- Cornish, N. J., Sesana, A. (2013), Pulsar timing array analysis for black hole backgrounds, *Classical and Quantum Gravity*, 30(22), 224005.
URL: <https://doi.org/10.1088/0264-9381/30/22/224005>
- Detweiler, S. (1979), Pulsar timing measurements and the search for gravitational waves, *ApJ*, 234, 1100–1104.
- EPTA Collaboration, InPTA Collaboration, Antoniadis, J., Arumugam, P., Arumugam, S. *et al.* (2024), The second data release from the European Pulsar Timing Array. V. Search for continuous gravitational wave signals, *A&A*, 690, A118.
- Hazboun, J. S., Meyers, P. M., Romano, J. D., Siemens, X., Archibald, A. M. (2023), [Analytic distribution of the optimal cross-correlation statistic for stochastic gravitational-wave-background searches using pulsar timing arrays](#), *Phys. Rev. D*, 108, 104050.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.108.104050>
- Hellings, R. W., Downs, G. S. (1983), Upper limits on the isotropic gravitational radiation background from pulsar timing analysis, *Astrophys. J.*, 265, L39–L42.
- Liu, K., Guillemot, L., Istrate, A. G., Shao, L., Tauris, T. M., Wex, N., Antoniadis, J., Chalumeau, A., Cognard, I., Desvignes, G., Freire, P. C. C., Kehl, M. S., Theureau, G. (2020), A revisit of PSR J1909-3744 with 15-yr high-precision timing, *MNRAS*, 499(2), 2276–2291.
- Phinney, E. S. (2001), A Practical Theorem on Gravitational Wave Backgrounds, *arXiv e-prints*, pp. astro-ph/0108028.
- Reardon, D. J., Zic, A., Shannon, R. M. *et al.* (2023), Search for an Isotropic Gravitational-wave Background with the Parkes Pulsar Timing Array, *ApJ*, 951(1), L6.
- Sazhin, M. V. (1978), Opportunities for detecting ultralong gravitational waves, *Soviet Ast.*, 22, 36–38.
- Taylor, S. R. (2021), The Nanohertz Gravitational Wave Astronomer, *arXiv e-prints*, p. arXiv:2105.13270.

Xu, H., Chen, S., Guo, Y. *et al.* (2023), Searching for the Nano-Hertz Stochastic Gravitational Wave Background with the Chinese Pulsar Timing Array Data Release I, *Research in Astronomy and Astrophysics*, 23(7), 075024.

DRAFT

DRAFT

6

Laser Interferometer Space Antenna

Quentin Baghi and Hubert Halloin

*Astroparticule et Cosmologie, Université Paris Cité, 75205 Paris Cedex 13,
France*

GRAVITATIONAL WAVES,
coordinated by Nicola Tamanini (L2IT/CNRS). © ISTE Editions 2025.

Foreword

On ground, the noise and constraints of the terrestrial environment limit the observable sources to frequencies above a few tens of Hz. The idea of building a space-based gravitational-wave observatory allows the measurement both to escape these terrestrial disturbances and, by significantly extending the length of the interferometer's arms, to observe at lower frequencies, corresponding in particular to the coalescence of much more massive objects, such as the super-massive black holes at the center of galaxies. The study of gravitational waves in this frequency range will provide many answers to some of the major questions still pending in the physics of the Universe: What is the origin of massive black holes? How do galaxies form? Is strong-field gravity correctly described by General Relativity? What is the nature of the primordial Universe's phase transition? Etc.

These scientific objectives (and many others) are behind the development of the LISA space project. This chapter describes the (long) genesis of this instrument, its key technical specifications, the science that the mission will enable, and how it will be extracted from its measurements.

6.1. The LISA mission concept

6.1.1. A little history

The LISA space mission is a European Space Agency (ESA) project to build a large-scale space interferometer to detect gravitational waves at around the millihertz range. The LISA mission was adopted as the third “Large” (L3) mission in the “Cosmic Vision” program by the ESA in January 2024, with a launch scheduled for 2035.

The history of LISA goes back much further, however. The first concepts for a space interferometer to detect gravitational waves were studied by NASA as early as 1974 (Bender *et al.* 1998, p. 43). The proposed instrument - Gravitational Wave Interferometer (GWI) - was based on the manufacture and deployment, from the Space Shuttle, of a rigid structure forming arms 1 km long, with a total mass of 16 t (see figure 6.1).

At the same time, these technical studies led to discussions on the feasibility of an interferometer with much larger arms, and therefore separate satellites. The idea, still valid today, was to exchange laser beams between test masses, protected from the solar wind and other external disturbances, to keep them inertial. Both approaches - monolithic structure and free masses with drag compensation - were presented at a workshop dedicated to gravitational waves by




Figure 6.1: Schematic diagram of the GWI space instrument as studied by National Aeronautics and Space Administration (NASA) in the 1970s. Left: structure after deployment in low-Earth orbit. Right: Principle of the interferometric measurement (Bender *et al.* 1998).

Weiss (1979). The concept of 3 distant satellites was studied and presented in greater detail shortly afterwards by Decher *et al.* (1980). The instrument thus consisted of 3 satellites, forming an isosceles right-angled triangle with sides of 10^6 km. These satellites are on heliocentric circular orbits; the one forming the right angle (i.e., at the emission and recombination of the laser beams) is regularly repositioned to maintain the shape of the triangle (see figure 6.2). To compensate for laser power attenuation (increased photon noise) and the Doppler effect, a transponder mode (copying the optical phase to distant satellites) and a heterodyne interferometric detection scheme (the two interfering lasers have slightly different frequencies) were proposed.

The idea of an interferometer with spacecraft orbits similar to those currently planned for LISA - and therefore not requiring frequent and extensive orbit corrections - was proposed as early as 1981. A fuller description of this concept, then called Gravitational-Wave Observatory in Space (LAGOS), is given by Faller *et al.* (1985). This mission concept proposed many of the ideas still valid for the LISA mission: Inertial satellites with drag compensation in heliocentric orbit located at 1 astronomical unit from the Sun, arms of 10^6 km (2.5×10^6 km for LISA), laser interferometry with heterodyne detection scheme and transponder satellites, and 50 cm diameter telescopes (30 cm for LISA).






Figure 6.2: Orbital concept of the instrument proposed by Decher *et al.* (1980). The two “sub-satellites” travel in the same circular heliocentric orbit. The S1 satellite also follows a generally circular, heliocentric orbit, but regular thrusts are required to maintain the constellation’s shape.



(a) Schematic diagram of the LAGOS main satellite.


(b) Schematic diagram of the LAGOS auxiliary satellites.




(c) Configuration of the LAGOS constellation.

Figure 6.3: First version of the LAGOS mission concept (from (Faller *et al.* 1985)).

The LISA mission was first proposed to ESA as part of the M3 missions of the Horizon 2000 program in the early 1990s, with a 4-satellite configuration (1 satellite per laser emitter), forming two arms of a single interferometer. However, it was considered to exceed the budget allocated to an “M” class mission. It was proposed in December 1993 and selected as the third “cornerstone” mission in the Horizon 2000 Plus program of the European Space Agency (ESA) (Danzmann 1996 ; Danzmann *et al.* 1997 ; Bender *et al.* 1998 ; Danzmann 2000, 2003), together with GAIA (astrometry), DARWIN (search for extra-solar planets), and a mission to Mercury. LISA then evolved into a configuration of 6, then 3 satellites - similar to the current configuration - with laser links on all three sides (extended to 5×10^6 km) of the constellation (cf. figure 6.4). The main elements of the satellites and payload are already present, and quite similar to the concept proposed by Faller *et al.* (1985).



(a) LISA configuration with 6 satellites (Danzmann 1996).



(b) LISA configuration with 3 satellites (Bender *et al.* 1998).

Figure 6.4: Evolution of the LISA orbital configuration in the late 1990s.

In early 1997, for reasons of scientific collaboration, as well as budgetary and scheduling constraints, LISA became a joint mission with NASA. A launch in the first decade of the 2000s was then considered, in particular to benefit from joint observations with the ground-based interferometers then in preparation : “A contribution by ESA in the range 50 - 200 MAU to a NASA/ESA collaborative LISA mission that could be launched considerably earlier than 2017 would fully satisfy the needs of the European scientific community. A launch within the time frame 2005 - 2010 would be ideal from the point of view of technological readiness of the payload and the availability of second-generation detectors in ground-based interferometers, making the detection of gravitational waves in the high-frequency band very likely.” (Bender *et al.* 1998, p. iv).

In the 2000s, LISA was finally identified as a candidate for the first Large mission (L1) in the new Cosmic Vision program (2015-2025) established by ESA.

At the same time, technical studies identified the need for a technological precursor to demonstrate inertial flight at the level required for LISA. The first concept for what would become LISA Pathfinder was proposed in 1998, under the name European Lisa Technology Experiment (ELITE). In 2000, in response to the SMART-2 mission call, a joint demonstrator for inertial flight for LISA and formation flying for the DARWIN mission was proposed (and selected), with the addition of a second satellite. Subsequent industrial studies led to the cancellation of the formation-flying demonstrator, and the remaining satellite, dedicated solely to inertial flight demonstration, was renamed LISA Pathfinder.

Technical studies continued in parallel on LISA, still considered a joint ESA-NASA mission, concluding in 2011 with a technical feasibility report. Unfortunately, due to budgetary difficulties at the time, NASA withdrew its support for joint projects with ESA, in particular the “L” mission and therefore LISA. Deprived of this important financial support, ESA asked the three missions previously candidates - LISA, International X-ray Observatory (IXO) and Europa Jupiter System Mission (EJSM) - to revise their plan to significantly reduce the cost of these projects (Samuel Reich 2011). For LISA, this mainly meant reducing the arms length (from 5 to 1 Mkm, which reduces the size and mass of the telescopes, and therefore the launch mass) and returning to a single interferometer (and thus four payloads instead of six) (Jennrich *et al.* 2011). The mission then took the name New Gravitational wave Observatory (NGO), then European Laser Interferometer Space Antenna (eLISA), and competed for the first L1 launch slot in 2022. Following evaluation by the Science Program Committee, the Jupiter Icy Moons Explorer (JUICE) mission (successor to EJSM) was selected.

Shortly afterwards, in March 2013, ESA launched a call for scientific themes for the L2 and L3 missions, with respective launch slots in 2028 and 2034. ESA received 32 proposals, including “The Gravitational Universe” (Danzmann 2013), associated with the mission concept eLISA. In October 2013, the Senior Survey Committee ultimately recommended “The Hot and Energetic Universe” (Athena mission) for L2 and “The Gravitational Universe” (LISA mission) for L3.

Following these recommendations, ESA published the call for L3 missions in October 2016, to which the LISA Consortium naturally responded (Danzmann 2017). With the American scientific community and NASA expressing their interest and willingness to support the mission, LISA returned to a configuration with 3 interferometers, 2.5×10^6 km between satellites, and also regained its original name. LISA was officially selected in the spring of 2017, and then adopted in January 2024. Since then, it has gone through the usual phases of a space project. The launch of the mission is currently expected in 2035. Following a cruise and commissioning phase of about 2 years, the scientific operations shall last for a nominal duration of 4 years, with a possible extension up to 6 additional years.

6.1.2. Description of the instrument

Here we summarize the main technical and operational features of LISA, in order to provide key insights into some of the main characteristics of the instrument. Thanks to the many years of studies that preceded its adoption and the Pathfinder experience, LISA entered its implementation phase with a mature payload architecture. The configuration presented here refers to that obtained at the end of Phase 0 (i.e., in spring 2018). At the level of detail required for this general description, this design is still valid today. However, this description only concerns the instrument (or payload), as the rest of the satellite (platform) is being developed by an industrial contractor.

6.1.2.1. Mission profile and measurement principle

As previously mentioned, the scientific measurement of LISA is based on determining the relative variation in the optical path length between the two arms of the interferometer. These arms are formed by laser links exchanged between the satellites. The three satellites are on heliocentric orbits, one astronomical unit from the Sun. The orbital parameters of each satellite are optimized so that the constellation maintains an almost constant shape throughout the mission, forming an equilateral triangle with sides of 2.5 Mkm. The constellation is located approximately 20 degrees behind (or, alternatively, ahead of) the Earth (in the same orbit) and drifts slowly away from it due to the pulling




Figure 6.5: LISA mission profile.

effect of the Earth's gravitational attraction. This mission profile is illustrated in Figure 6.5.

Each satellite hosts two free-falling masses, located at each end of the 3 arms of the constellation. The effect of a gravitational wave is to slightly modify the propagation time of the laser beam from one satellite to the next, and thus to alter the optical phase (or, equivalently, the frequency) of the received wave. These fluctuations are measured by heterodyne interferometry within the three satellites. The instrument's sensitivity is limited by the residual acceleration noise (i.e., the disturbing forces) exerted on the test masses, and by the stability of the interferometric measurement. The satellites surrounding the pairs of test masses act as shields. The position of the satellites is constantly corrected using micro-thrusters (with maximum thrusts in the order of μN) to maintain the test masses in free fall within them. The orientation of the satellites (and therefore of the test masses) is, however, constrained to ensure that laser beams can be both received and aimed at the distant satellites.

The changes in light distance between satellites caused by gravitational waves are very small, ranging from a few pm to a few nm at most, equivalent to residual accelerations of the order of 10^{-15} m/s^2 . Test mass trajectories, on the other hand, are bound by the laws of celestial mechanics and result in much larger variations in arm length, typically $\pm 35,000 \text{ km}$. Fortunately, the detection of gravitational waves remains possible thanks to the frequency separation of the phenomena: expected gravitational waves are around the millihertz, while orbital variations have periods of several months.

The expected amplitude of gravitational waves (see, for example, Danzmann (2017) for an overview of LISA's scientific objectives and the associated detection constraints) therefore imposes two specifications: one on the residual acceleration noise of the test masses, and the other on the precision of the

interferometric measurement. The requirements on the corresponding noise amplitude spectral densities are

$$S_{\text{link}}^{1/2} < 10 \text{ pm}/\sqrt{\text{Hz}} \times \left[1 + \left(\frac{2 \text{ mHz}}{f} \right)^4 \right]^{1/2}, \quad [6.1]$$

$$S_{\text{g}}^{1/2} < 3 \text{ fm}/\text{s}^2/\sqrt{\text{Hz}} \times \left[1 + \left(\frac{0.4 \text{ mHz}}{f} \right)^2 \right]^{1/2} \left[1 + \left(\frac{f}{8 \text{ mHz}} \right)^4 \right]^{1/2} \quad [6.2]$$

These noise limits, together with the instrumental response, enable us to calculate the interferometer's sensitivity and compare it with the scientific objectives (see Figure 6.6). A detailed description of the LISA scientific objectives is given in section 6.2.




Figure 6.6: Sensitivity of LISA compared with typical amplitudes for the main scientific targets. The characteristic amplitude is deduced from the amplitude spectral density by multiplying by the root of the frequency. This quantity can be used to calculate the signal-to-noise ratio of a quasi-monochromatic source, integrated over one period. (Danzmann 2017).

Optical path measurements between test masses are carried out continuously, using heterodyne interferometry, with laser beams propagating in both

directions of each constellation arm. The laser source used is a Nd:YAG laser emitting an optical beam at 1064 nm, of which approximately 1.5 W is emitted towards the distant satellites. After a propagation of 2.5×10^6 km, the power collected by the distant satellite is around 700 pW, which precludes any “passive” reflection using a mirror. Satellites are therefore used as active transponders, sending back an amplified signal which is phase-locked to the incident beam.

To compensate for the Doppler drift caused by the relative movements of the satellites ($\sim \pm 8$ MHz), and to maintain the optical beat frequencies within the bandwidth of the photodiodes (5-25 MHz), it is also necessary to phase-lock all the lasers together, and to introduce frequency deviations of a few MHz in these locking loops, pre-calculated according to orbital mechanics. In the end, only one of the 6 laser sources is free-running, with the others being directly or indirectly phase-locked to it.

Apart from this frequency control topology, the constellation is entirely symmetrical, with identical loads on each satellite and similar measurements. The measurement between test masses of a laser link is performed in 3 steps (see Figures 6.7 and 6.8) :

- Local interferometric measurement, on the first satellite, between the test mass and the optical bench attached to the satellite
- Long-distance measurement between optical benches
- Local interferometric measurement, on the second satellite, between the test mass and the optical bench attached to the satellite

The combination - a posteriori - of these 3 measurements, taking into account the appropriate propagation times, eliminates most of the positional noise between the satellites and their respective test mass. It is also necessary to measure the relative optical phase between two adjacent benches on the same satellite. This is achieved using a bi-directional fiber laser link between the two optical benches (and the associated interferometric measurements).

Another important and specific aspect of LISA is the need to synchronize the satellite clocks. Indeed, the stability of heterodyne measurement relies on the ability to measure, with negligible noise, the phase stability of a radio frequency signal at a few MHz. This measurement must be carried out relative to a “perfect” clock, in the sense of a clock with a constant oscillation rate. Thanks to the heterodyne interferometry, the optical phase is transferred to the RF signal phase, so that an optical path difference of 1 pm (i.e. $\sim 10^{-6}$ cycle at 1064




Figure 6.7: Principle of measurement between test masses on one arm of LISA




Figure 6.8: Schematic diagram of interferometric measurements in a LISA instrument

nm wavelength) results in an equivalent phase shift of the beat signal. Since the beat frequency is at most 25 MHz, a specification of $\sim 10^{-6}$ cycle/ $\sqrt{\text{Hz}}$ requires clock stability better than $\sim 4 \times 10^{-14}$ s/ $\sqrt{\text{Hz}}$ in the LISA band (0.1 mHz to 10 Hz). This is well beyond the stabilities of the clocks typically embedded in satellites (of the Oven Controlled Crystal Oscillator (OCXO) type). On the other hand, the desired stability is *relative* between the different arms of the interferometer. It is therefore sufficient to guarantee that the data from the 3 satellites are corrected for their *relative* drift.

The solution proposed in LISA is to use laser links to transfer clock frequencies from one satellite to another. To achieve this, the individual laser beams are phase-modulated at a frequency that is a multiple of the on-board clock frequency. This modulation creates sidebands at around ± 2.4 GHz from the optical carrier. When this transmitted beam interferes with the receiving satellite's local beam, 3 radio-frequency beats are measured :

- Between local and distant carriers: beat frequency equal to the Doppler shift plus the shift imposed during phase locking.
- Between lower sidebands: beat frequency equal to the heterodyne frequency between carriers added to the sideband frequency difference (local - distant).
- Between higher sidebands: beat frequency equal to the heterodyne frequency between carriers added to the sideband frequency difference (distant - local).

In this way, sideband beats encode the relative phase noise of the clocks, with an amplification factor of the order of 100 (ratio of modulation frequencies to main beat frequency). It is then possible (in post-processing) to correct the phase noise of the main beat for the noise induced by the instability of the clocks.

A posteriori signal reconstruction and noise reduction are largely based on the time-delay interferometry (TDI) algorithm (see Sec. 6.3.1), which combines the phases measured by the different beats (main and lateral) of all the interferometers present in the constellation. This algorithm, in fact, reconstructs a Michelson-like measurement, where the data collected is delayed by the propagation times required to synthesize an interferometer with equal arms. Knowledge of the absolute distances between satellites is therefore a prerequisite for this analysis method. These are obtained by slowly modulating the laser beams in phase, using a synchronization pattern that can be dated both in reception and in transmission, thus enabling propagation times to be determined at better than 1 meter over 2.5 million kilometers.

6.1.2.2. Overview of the LISA payload

The main equipment making up the LISA instrument and enabling interferometric measurements is as follows :

- A 30 cm diameter telescope to transmit and collect the laser beam to and from distant satellites.
- An inertial mass, with electrostatic readout and actuation devices to define the inertial measurement reference. This assembly is named Gravitational Reference Sensor (GRS).
- An optical bench combining local and distant beams to perform the various interferometric measurements.
- A phasemeter to reconstruct the relative phase of the beats described above, and to determine the time of flight (and hence the distance) between satellites.
- A Nd:YAG laser source, with a power of about 2 W, which can either be frequency-stabilized on a Fabry-Perot cavity (master laser), or phase-locked to another laser in the constellation (adjacent or distant).
- A set of thermal, magnetic, and radiation sensors for precise knowledge of the instrument's environment (and therefore to prevent any coupling effects with the measurement).

The most stringent stability constraints (from pm to nm) concern the 3 main elements involved in optical measurement: the telescope, the optical bench and the GRS. In addition, the optical axis of the beam (transmitted and/or received) must be adjustable to take account of the constellation's "breathing", i.e., the slight deformations of the triangle due to orbital mechanics. To meet these requirements, these three elements are mounted on a single structure, which itself can be rotated (on $\pm 1^\circ$) to follow the direction of the distant satellite. A LISA satellite contains two copies of this assembly, each pointing toward one of the two other satellites. The phasemeter, the stabilized laser source, the payload computer, the various environmental sensors and all the control electronics complete the instrument and are positioned on the satellite platform.

Among the various payload components, the optical bench and, above all, the GRS, benefit from the heritage of LISA Pathfinder (LPF). We briefly describe it in the next section.

6.1.2.3. *LISA Pathfinder heritage*

The LPF mission was launched on December 3, 2015 aboard a VEGA rocket from the Kourou space base in French Guiana. The satellite then reached the L1 Lagrange point of the Earth-Sun system, i.e. around 1.5 Mkm from Earth along the Sun-Earth axis. The main objective of the mission was to validate the possibility of maintaining an inertial mass in orbit, with a low residual acceleration - less, for this demonstrator, than 10 times the specification LISA. On the other hand, measuring this residual acceleration also requires the demonstration of optical benches stable enough to perform a measurement at $\text{pm}/\sqrt{\text{Hz}}$ between 10 mHz and 1 Hz.

We will not describe the instrument in detail here, but only its principle, as illustrated in Figure 6.9: two inertial masses are suspended without contact within two GRS, on either side of an ultra-stable Zerodur optical interferometry bench. This assembly forms the LISA Technology Package (LTP), located at the core of the satellite. The distance between the two inertial masses, each measuring 5 cm per side, is 38 cm. Fine fluctuations in this distance are measured by heterodyne interferometry. The distance between one of these test masses and the optical bench is also measured, and is used to control the satellite's position using cold gas or colloidal micro-thrusters placed on the satellite's external walls.

The first scientific measurements of the LPF instrument took place in March 2017, and the mission ended in July 2018. During this period, a large number of experiments were carried out to better characterize the various sources of acceleration noise, their coupling with the satellite environment and to set the instrument to the operating point that minimizes these couplings.

Data analysis has shown that the achieved performance far exceeds the specifications of LPF and even those required for LISA, cf. figure 6.10 (Armano *et al.* 2018). Compared with the noise spectra of terrestrial interferometers, no background noise lines are observed. The main contributors to the residual acceleration forces are all stochastic or broadband processes :

- Thermal fluctuations of the various elements, modifying optical paths and the response of the readout or electrostatic actuation systems.
- Residual motion of the rotating satellite, which creates acceleration coupling via inertial forces (Coriolis and Euler forces).
- Brownian motion of residual molecules on the surface of test masses, inducing momentum fluctuations




Figure 6.9: Transparent view of the instrument at the heart of LPF. Two inertial masses - each about 5 cm square - are held in vacuum chambers and between electrodes for actuation and electrostatic measurement (right and left of figure). The distance between these two test masses, and between a test mass and the satellite, is measured precisely by an interferometer located between the GRS. Credit: ESA/ATG medialab

- Photon noise related to quantum uncertainty about photon phase. This is white noise in phase, and therefore proportional to f^2 in acceleration. This effect limits performance at high frequencies. In LPF, the interferometric reading noise was $34.8 \pm 0.3 \text{ fm}\sqrt{\text{Hz}}$. It is higher in LISA in the long-arm interferometer ($5 \text{ pm}\sqrt{\text{Hz}}$ for photon noise only) due to the low power received from the distant satellite ($\sim 600 \text{ pW}$ instead of a few mW in LPF).

6.2. The science enabled by LISA

The mission profile and sensitivity requirements described in Sec. 6.5 are designed to meet several main scientific objectives that are detailed in the mission's Definition Study Report (Colpi 2024). They aim at answering pending questions in modern astrophysics that range from compact astrophysical source




Figure 6.10: Residual acceleration noise spectrum on a test mass of LISA Pathfinder. The blue curve corresponds to intermediate results, the red curve to final results, after various characterization campaigns and subtraction of environmental couplings (from (Armano *et al.* 2018)).

populations and evolution, cosmology, including the early Universe physics, to fundamental physics, including the nature of black holes and gravity. This section provides a summary of LISA’s expected scientific outcomes.

6.2.1. Landscape of gravitational wave sources in the LISA band

The LISA observatory is designed to observe gravitational waves in a frequency range that is not covered by ground-based detectors. This frequency window spans between 0.1 mHz and 1 Hz, which corresponds to frequencies ten thousand times lower than the lower bound of current terrestrial observatories. Similar to electromagnetic astronomy, studying different wavelengths enables us to detect various physical phenomena or sources of radiation. Notably, LISA provides access to either larger mass systems or systems in an early stage of their dynamic evolution, which are beyond the reach of ground-based detectors.

The gravitational-wave frequency evolution of a circular, non-spinning binary system during the inspiral phase can be approximated by Eq. [1.5] of Sec. 1.1.4.1. (using only the quadrupolar mode). This equation shows that the orbital frequency is proportional to $M_{\text{chirp}}^{-5/8}$. Hence the more massive the system,

the lower its frequency. Additionally, the frequency of the binary increases with time as it emits gravitational waves (GWs). If the system is very far from coalescence during observation, its frequency will be smaller than for a system that is about to merge.

As a result, LISA is designed to detect various types of sources (see Fig. 6.6). The closest ones are compact binary stars (with masses close to one solar mass) in their inspiral phase, mainly located in the Milky Way galaxy. Other systems that LISA can detect in their inspiral are stellar-mass black hole binaries, with masses between a few and 100 solar masses. These sources are located at redshifts $z < 1$ and become detectable by ground-based detectors during merger. A third type of potentially detectable inspiraling systems is extreme mass ratio inspirals (EMRIs), where a compact object like a compact star or a stellar-mass black hole is orbiting a much more massive black hole. The name of these sources stems from the large difference between the masses of the two bodies orbiting each other, with ratios ranging from 10^{-2} to 10^{-5} . LISA could detect them in their inspiral phase before the small object plunges into the bigger one, at redshifts up to $z \sim 4$.

LISA can also detect merging sources, the main ones being massive black hole binaries (MBHBs) with masses between 10^5 and 10^7 solar masses, which are mainly located at the centre of faraway galaxies. Due to their large chirp mass, they emit a relatively strong GW signal, allowing MBHBs to be detected up to redshifts $z \sim 12$ and possibly beyond.

Furthermore, some sources of gravitational waves are too far away or too faint to be individually resolvable. However, the cumulative contribution of many such sources can produce a detectable GW background. One usually distinguishes two types of backgrounds depending on their origin. The first type is astrophysical backgrounds, made by the addition of multiple faint binary sources (Galactic binaries, stellar-mass black hole binaries, MBHBs or EMRIs). The second type is cosmological backgrounds, originating from many GW sources in the early Universe, like local density fluctuations, forming a fossil radiation carrying information about the Universe's first instants.

In the following, we review the science that LISA will enable through the observation of these various GW sources.

6.2.2. Astrophysics of compact objects

Many compact objects in the Universe are faint or dark. Therefore, it is difficult and sometimes impossible to observe them through telescopes. However, a large fraction of stars (Whitworth and Lomax 2015) form or reside in binary

and triplet systems. This means that even though they cannot be detected by electromagnetic (EM) radiation, they can emit GWs. How compact binaries form, interact, and evolve is still poorly understood and an active area of research. Observing them through mHz gravitational waves and probing their populations will greatly enhance this understanding.

6.2.2.1. Compact binary stars

Main-sequence stars with initial masses up to $10M_{\odot}$ eventually shrink under their own gravity and reach a mass comparable to the Sun that fits into an Earth volume. These extremely dense dead stars are called white dwarfs. What prevents them from further collapsing is the degeneracy pressure of electrons in their interior. Current estimations (Napiwotzki 2009) suggest that the Milky Way contains about ten billions white dwarfs. A significant fraction (about a third according to observations) of them live in binary systems (Holberg 2009). Among these, about ten million binaries will emit GWs in LISA's sensitive band (Colpi 2024 ; Amaro-Seoane 2023), with typical orbital periods of tens of minutes. The detector will be able to individually detect 0.1% of them. The remainder will form a confusion foreground of unresolved sources, which nevertheless also carries information about the Galactic source population (Breivik *et al.* 2020).

More massive stars (with masses between $10M_{\odot}$ and $25M_{\odot}$) collapse into objects even denser than white dwarfs. They can evolve into a neutron star (Heger *et al.* 2003), whose mass of $1.4M_{\odot}$ is concentrated in a 10 km radius sphere, or into a stellar black hole (a few tens of solar masses). Binaries of neutron stars and black-hole-white dwarf binaries are also candidate sources for LISA, and could be detected in hundreds.

Fig. 6.11 represents LISA's detector sensitivity (red solid line) along with a simulation of detectable Galactic sources' amplitudes (blue dots) based on works from Korol *et al.* (2019) ; Wagg *et al.* (2022) ; Toonen *et al.* (2018). Although one could have represented sources below the cloud of points visible on the figure, we set the detectable threshold to be a signal-to-noise ratio (SNR) equal to 7.

High-cadence photometric surveys (Kupfer *et al.* 2021 ; Finch *et al.* 2023) and spectroscopic surveys (Brown *et al.* 2020) have discovered short-period compact binary systems in our Galaxy. Combined with astrometric measurements to derive distances, these observations have enabled more than fifty systems to be identified as potential GWs sources in the LISA band (Kupfer *et al.* 2024). The loudest of them are called verification binaries, as their parameters are already partially known and could help validate LISA observations. Although this multi-messenger measurement is promising, their utility as calibration sources remains uncertain (Littenberg and Lali 2024).

LISA measurements of compact Galactic binaries will bring unique information on binary interactions, whose many aspects are not well understood today. We give an overview of some promising scientific outcomes.

Compact stars evolution. First, LISA will improve our understanding of how compact stars evolve into close binaries, also leading to how many of them eventually merge. Indeed, the high number of Galactic sources LISA can detect will allow us to build a catalogue of binary systems, including their sky locations, frequencies, and chirp masses (when the frequency derivative is measurable). This catalogue will be compared to predictions of binary evolution models to determine the role of astrophysical processes such as mass transfer between the two bodies (Nelemans *et al.* 2001). Additionally, the detected population provides the number of systems as a function of frequency, which allows one to infer the binary merger rate based on the frequency evolution predicted by general relativity (GR) and given in Eq. [1.5]. In turn, one can constrain the number of explosive events such as Type-Ia supernovae and kilonovae (Webbink 1984).




Figure 6.11: Simulation of strain amplitudes of detectable Galactic binaries (blue dots) including verification Galactic binaries (red hexagons). The red solid line represents LISA’s sensitivity due to instrumental noise only. The green solid line includes instrumental noise and Galactic confusion noise for a 4.5 year mission lifetime. Plot produced based on the back-end software of the LISA Science Explorer (Le Jeune and Sartirana 2025).

Spatial distribution in the Galaxy. Second, LISA will probe the spatial distribution of white dwarf binaries in the Milky Way, thereby providing insight into the geometrical structure of the Galaxy. In addition to their sky location, that LISA can measure with precision of order 1 deg^2 depending on SNR, distances can be determined from the measurement of the GW frequency derivative when available. This leads to spatial distributions like the one depicted in Fig. 6.12. Such a mapping is very complementary to telescope observations, as it does not suffer from common obstacles encountered when using EM radiation, like dust extinction, stellar crowding in the Galactic centre, and strong distance selection effects.




Figure 6.12: Simulated spatial distribution of Galactic binary systems detected by LISA in Galactic Cartesian coordinates with origin in LISA location (0,0). The colour scale reflects precision forecasts of the sky location estimate. Plot produced based on the back-end software of the LISA Science Explorer (Le Jeune and Sartirana 2025).

Astrophysical interactions in binaries. Third, one can measure the impact of astrophysical effects such as tides and mass transfer. When a pair of white dwarfs gets close enough, tidal forces become significant, and one of them may

transfer mass to the other. The interplay of these phenomena with GW emission determines the system's final state, i.e., whether the binary eventually merges or evolves into a stable interacting binary. LISA measurements provide an estimate of the frequency derivative, which depends on both GW emission and astrophysical effects. Combined with EM observations of the system (which provides individual masses), this information can help disentangle the role played by GR and tidal effects (Shah and Nelemans 2014) in the binary's period evolution. Consequently, one can estimate the fate of several systems and assess the fraction of binaries bound to merge compared to the ones that stabilize into interacting binaries.

6.2.2.2. *Black holes history*

LISA will also provide unprecedented observations of black hole binaries. One distinguishes two families of black holes: stellar-mass black holes (sBHs) and massive black holes (MBHs). The first family comes from the collapse of massive stars at the end of their life, and have masses from a few to several tens of solar masses. When they come in binaries and are about to merge, they enter the detection band of terrestrial detectors like LIGO, Virgo or Kagra. LISA could nevertheless observe some of them earlier in their inspiral phase, as continuous sources of GWs. Massive black holes have masses between 10^4 and 10^7 solar masses, and live in the centre of most Galaxies. When two Galaxies merge, it is expected that the black holes in their centre eventually merge as well. However, how to form a massive black hole is a pending astrophysical question. The process of forming bonded binaries and their merger rate is also largely unknown. LISA is an ideal observatory to advance our understanding of these questions. Indeed, when MBHs become close enough, their GW frequencies enter the LISA sensitive band.

Note that the existence of a third family of black holes called intermediate-mass black holes (IBHs) is postulated (Greene *et al.* 2020 ; Straub 2023), with masses between $10^2 M_{\odot}$ and $10^5 M_{\odot}$, bridging the gap separating sBHs from MBHs. However, their existence as well as their possible formation process is under debate.

Today, we can detect black holes in several ways thanks to EM observations. One can model the kinematics of spatially resolved stars observed through astrometric measurements in the centre of our galaxy to reveal the black hole's presence (Ghez *et al.* 2008). Another method is to look for tidal disruption events (Gezari 2021) arising when a star ventures so close to a MBH that it is tidally ripped apart or captured, generating a burst of EM radiation in UV, optical, X-ray, or gamma-ray. Generally, one can detect black holes through the EM emission of the matter it accretes, producing what is called active galactic nuclei (AGNs) or its more luminous version, quasars. More recently,

direct imaging of the Milky Way and M87 galaxies' central black holes has even been made possible. This is a global array of radio telescopes, the Event Horizon Telescope (The Event Horizon Telescope Collaboration 2019, 2022), which produced this first image of a black hole. However, it is generally challenging to observe these dark objects with EM observations, especially the low-mass and high-redshift population. In this quest, mHz gravitational waves constitute an alternate messenger that has the potential of providing a very complementary insight.

As they merge, MBHs binaries produce chirping transient signals observable for several days to weeks. For massive enough systems ($M > 10^5 M_\odot$), LISA will observe the three main phases of the waveform: inspiral, merger, and ringdown. We show the strain amplitude spectral density of typical MBHs binary signals in Fig. 6.13. For each signal, the inspiral corresponds to the first relatively mild slope towards low frequencies, the merger approximately corresponds to the strain bump, and the ringdown to the steep final slope. As represented by the colour scale, MBHs binaries may have SNR larger than thousands. This makes them exceptional probes of gravity. For systems with a smaller total mass, the merger time happens at higher frequencies where LISA is less sensitive. However, it may still observe their inspiral phase.




Figure 6.13: Simulated amplitudes for possible MBH binary sources observed by LISA. Colours indicate the cumulated SNR as a function of frequency, from lowest (bluer) to highest (redder). Plot produced based on the back-end software of the LISA Science Explorer (Le Jeune and Sartirana 2025).

Searching the origin of black holes. We expect that MBHs are born with relatively small masses, and grow over time. Their origin, which we call black hole seeds, is unknown. LISA is an ideal tool to unveil this mystery, because it can peer into very large redshifts (beyond 10) when these seeds could have appeared.

We distinguish various types of black hole seeds that could have operated in parallel. One is referred to as primordial origin seeds, and would result from the collapse of non-linear density perturbations in the early Universe leading to black holes with masses up to $10^5 M_\odot$, but potentially much lower (Carr and Silk 2018). Another type is called light seeds, which designates black holes of masses between ten and hundreds of solar masses that would have formed from the collapse of metal-free stars inside the first dark matter halos, at redshifts between 15 and 30. Heavy seeds, on the contrary, refer to black holes with masses larger than $10^3 M_\odot$ forming from the direct collapse of massive stars or gas clouds at redshifts $z \in [10; 15]$ in later (and more massive) dark matter haloes (Regan and Volonteri 2024).

James Webb Space Telescope (JWST) observations of faraway quasars have already challenged our initial conceptions of black hole formation. They hint for the presence of very early supermassive black holes at $z > 6$ with millions solar masses and beyond, and also reveal high black hole-to-host galaxy mass ratios (Suh *et al.* 2025). These observations tend to support heavy seeds scenarios (Bogdan *et al.* 2024). However, more observations of lower mass black holes at high redshift would help to better understand the relationship between these observations and black hole seeds. This is where LISA comes into play.

By probing the binary black holes population at high redshift, LISA could help understand the pathways leading to the formation of the MBHs we observe today.

Growth mechanism and merger history. Black hole seeds are expected to grow to the present in order to explain the MBH we observe today. MBH growth can happen through two mechanisms: accretion of surrounding matter and coalescence with other black holes. It starts at cosmic reionisation ($z \sim 8$) when energetic photons emitted by the first stars reionize neutral hydrogen in the intergalactic medium.

As dark matter haloes assemble, galaxies merge across cosmological scales. However, the process of coalescing the MBH at the center of the two merging galaxies is largely unexplained today. It requires dissipating the orbital energy and transporting the angular momentum so that the two MBHs are brought from Galactic scales to kiloparsecs, and from kiloparsecs to microparsec separation. As a consequence, the delay between the merger of the galaxies and

their two central black holes is uncertain, ranging from $\sim 10^8$ years to the age of the Universe (Gualandris *et al.* 2022). A wealth of numerical simulations and semi-analytical models are proposed to describe the black hole population and merger rates (Amaro-Seoane 2023 ; Dayal *et al.* 2019 ; Klein *et al.* 2016). LISA will allow us to constrain them by measuring the masses, distances, and spins of MBHs, whose values are strongly related to the coalescence mechanisms.

Electromagnetic counterparts of massive black hole mergers. As MBHB coalescences should occur in a gaseous environment, dissipation of kinetic and magnetic field energy may lead to high energy EM emission during the pre-merger phase (d’Ascoli *et al.* 2018). This emission could even be periodically modulated by the orbital motion of the black holes, leaving a useful mark to relate the GW event to its counterpart (Tang *et al.* 2018). After the merger, the shock undergone by the circumbinary disc may lead to radio and X-ray radiation, revealing a newborn AGN (Rossi *et al.* 2010) or a collimated jet (Yuan *et al.* 2021).

By issuing alerts, LISA will enable multimessenger observations of MBHB merger events involving new-generation observatories like Vera Rubin (Ivezic *et al.* 2019), Roman space telescope (Akeson 2019), Athena (Cruise 2025), and SKAO (Carilli and Rawlings 2004). A few days before the merger, LISA can localize MBHBs at $z \sim 0.3$ with total masses in $[10^5 M_\odot; 10^6 M_\odot]$ within less




Figure 6.14: Average signal-to-noise ratio of MBH binary sources (indicated by the colour scale) as a function of redshift and source-frame total mass.




Figure 6.15: Three-dimensional representation of the orbit of a point mass in Kerr spacetime evolved during 5.56 hours, with a black hole mass of $10^6 M_\odot$ and a dimensionless spin of 0.9. One can view this illustration as an approximation of a short segment of a EMRI trajectory. Distances are expressed in units of the gravitational radius $r_g = GM/c^2$. Illustration extracted from Berry *et al.* (2019).

than 10 square degrees. A few hours before merging, the localisation uncertainty drops below 1 deg^2 , with a final localisation within 0.1 deg^2 (Lops *et al.* 2023 ; Lang and Hughes 2009 ; Colpi 2024). Such progressively updated sky locations will help plan the EM follow-up and find candidate host galaxies, prioritizing those featuring AGNs. Finding EM counterparts to LISA detections would bring unprecedented information about the accretion and jet mechanisms in a dynamical spacetime.

6.2.2.3. Black hole properties and environments

When a small compact object orbits around a MBH, it radiates GWs that encodes information about the local spacetime geometry and possibly about the immediate matter environment. If the smaller-to-larger mass ratio of the two bodies q is between 10^{-4} and 10^{-3} , the system is called an intermediate mass ratio inspiral (IMRI). When it is comprised between 10^{-6} and 10^{-4} , one call it an EMRI. For mass ratios below 10^{-6} , we resort to the designation of extremely mass ratio inspiral (XMRI). By detecting these source classes, LISA will hint at their environments, as EMRIs (or IMRIs when the secondary is an

intermediate-mass black hole) tend to form in galaxy centres while IMRIs tend to form in the centres of dense clusters.

Orbits of EMRIs and IMRIs are fully relativistic and may be highly precessing and eccentric, as depicted in Fig. 6.15. The resulting waveform exhibits many harmonics depending on orbital, azimuthal, polar, and radial motions. LISA can observe EMRIs for months to years before they plunge. During that time, the waveform may undergo hundreds of thousands of cycles. Extracting these waveforms from the data will allow accurate determination of the central black hole's properties at redshifts $z < 3$, like its mass and spin. The gaseous environment of the black hole and even the dark matter halo in which it is embedded can lead to measurable deviations of the GR vacuum solution (Speri *et al.* 2023 ; Gliorio *et al.* 2025).

6.2.3. Fundamental physics: The nature of gravity and black holes

LISA is expected to detect very loud sources of GWs like MBHB mergers with SNRs of thousands but also long-lived sources like EMRIs and IMRIs. They are formidable opportunities to test GR in the strong field regime. One can perform these tests in three ways. One way consists of studying the GW generation process, comparing measurements with waveforms derived from solutions of GR vacuum field equations. Another way is searching for beyond-GR emission channels through non-tensor GW polarizations. A third way is to study the propagation properties of GWs, measuring the wave velocity and dispersion, especially when lensing effects or EM counterparts are observed. We give some examples below (for a more detailed review, see Arun (2022)).

Probing the gravitational-wave ringdown. When two MBHs merge, they are expected to form a new massive black hole, called a remnant, that is perturbed and dissipates its energy by emitting GWs until it becomes completely gravitationally silent. Within the GR framework, the spacetime around the remnant is described as a perturbed Kerr solution, which depends only on the black hole mass and dimensionless spin¹. During the remnant de-excitation process, the waveform undergoes the ringdown phase, which features damped oscillations called quasi-normal modes. Each mode is identified by spherical angular indices l, m , and an overtone n . It is characterized by a damping time τ and a frequency ω which should only depend on the two black hole parameters. However, if the remnant is not a black hole (if it has no horizon) and the surrounding spacetime is not described by the Kerr metric, then the spectrum of

1. In principle, it also depends on charge, but it is neglected in the case of astrophysical black holes as it is expected to vanish quickly (Fabian and Lasenby 2019).




Figure 6.16: Relative uncertainty contours of frequency ω (left column) and damping time τ (right column) for ringdown quasi-normal modes measured by LISA as a function of source total mass M and redshift z . The upper row shows the mode $(l, m, n) = (3, 3, 0)$ and the lower row shows the mode $(l, m, n) = (4, 4, 0)$. The dotted black lines are ringdown SNR contours. Plot taken from Pitte *et al.* (2024).

quasi-normal modes will be modified. LISA could measure such modifications with high SNR sources, hereby testing the no-hair theorem (Toubiana *et al.* 2023 ; Pitte *et al.* 2024). Fig. 6.16 gives an assessment of the relative precision LISA could achieve on estimating modes' deviations from GR using ringdown signals, which ranges from sub-percent to ten percent depending on the source redshift and mass.

Mapping spacetime around massive black holes using EMRIs. The complex orbits a small compact object undergoes around a massive black hole provide a detailed mapping of the local spacetime. LISA may observe year-long EMRI signals with typical SNRs of about 50, which allows for investigations of spacetime geometry and tests for the presence of dark matter or new fundamental

fields. In the latter scenario, the primary (more massive) object can deviate from a GR black hole, featuring “hairs” (Doneva *et al.* 2024). Inspecting the multiple moments of the central MBH can enable testing the no-hair theorem. A new field (like ultralight scalars) can also affect the secondary object, which could carry a measurable charge (Barsanti *et al.* 2023).

Searching for emission channels beyond general relativity. Gravitational waves as predicted by GR can only be tensorial and transverse. Some alternate theories to GR like scalar-tensor theories (Will 2014 ; Callister *et al.* 2017) or massive graviton theories (Paula *et al.* 2004) predict the existence of non-tensor modes, to which LISA can be sensitive (Tinto and Alves 2010). Discovering such polarizations would strongly constrain extensions of GR. Another powerful testing tool would be multiband sources, i.e., sBH binary signals that would be detected first by LISA during inspiral and then by terrestrial observatories when they merge. The merging time compared to LISA first observation would be strongly affected by the presence of additional emission channels. Additionally, multiband sources would allow assessing GR consistency for the same system in very different dynamical regimes.

Testing gravitational waves propagation at mHz frequencies. Ground-based detectors’ measurements have provided strong constraints on the speed of GWs, thanks to the multi-messenger observation of the binary neutron star merger GW170817 (Feng *et al.* 2023). However, GW propagation properties could depend on frequency, and performing a similar measurement at mHz frequencies would test this hypothesis (Chamberlain and Yunes 2017). Like for testing alternate emission channels, multiband sources will be ideal probes of the speed of mHz GWs by comparing the coalescence time measured by terrestrial detectors relative to the GR prediction. Furthermore, inspecting the GW phase can help constrain the graviton mass in massive gravity theories (Will 1998) or reveal a modification of the wave dispersion relation (Mirshekari *et al.* n.d.).

6.2.4. Cosmology

6.2.4.1. Measuring the Universe’s expansion rate

Today’s cosmology faces a dilemma as two independent and different measurements of H_0 are in contradiction at the 5σ level, an issue known as Hubble tension. One relies on the cosmic microwave background (CMB) observation, assuming a standard Λ CDM cosmological model (Aghanim 2020). The other measurement depends on the distance ladder, using stars of known intrinsic brightness to measure distance (like Cepheids) in galaxies whose redshift can be determined with type-Ia supernovae (Riess 2022). Besides, recent observations made with JWST provide estimates in agreement with the CMB measurement (Freedman *et al.* 2025).

However, gravitational wave measurements can provide a direct estimation of a source's luminosity distance if it exhibits a detectable chirp. If the source's redshift can also be determined (for example, by identifying its host galaxy), it is possible to measure the Hubble constant, H_0 . We call such an object a standard siren, paralleling the denomination of standard candle in EM astronomy. This would provide a completely independent way of probing the Universe's expansion rate and help determine whether there is indeed a Hubble tension. At least three techniques can achieve this. The most accurate one is to use bright sirens (Mangiagli *et al.* 2022), if LISA detects MBHB mergers which can be associated with an EM counterpart. This way, it could provide an estimate of the expansion rate at $z \sim 2$ with 10% precision (Colpi 2024). If a counterpart is not observable, for example when detecting GWs from an EMRI (see Sec. 6.2.2.3), then the host galaxy can be inferred from the known distribution of galaxies in the sky region determined from the GW signal (Laghi *et al.* 2021). A third technique, known as spectral sirens, utilizes the mass spectrum of the assumed population of the observed object to infer its redshift (Ezquiaga and Holz 2022).

6.2.4.2. *Peering into the early Universe with stochastic GW backgrounds*

Today's EM observations of the CMB provide us with an instantaneous picture of the Universe as it was about 380,000 years after the Big Bang. Analysing the spectrum of its anisotropies was a major breakthrough that gave strong support to the Λ -CDM cosmological model. However, it does not grant us access to earlier times when the Universe was opaque to light. On the contrary, gravitational waves emitted before could have travelled to us almost unimpeded. They could be a new messenger carrying information about the youth of our Universe, in particular, the postulated cosmic inflation.

Multiple incoherent GW sources emitting simultaneously form what we call a stochastic background. Such sources in the early Universe could have formed a detectable fossil radiation, which may have various origins. First, the Universe may have undergone a phase transition at the end of the electroweak epoch, when the electromagnetic and weak interactions decoupled. The nature of this phase transition is currently unknown. Should it be first order, it could have given rise to a stochastic gravitational-wave background (SGWB) emanating from the expansion and collision of nucleating bubbles. The power spectrum of such a background is expected to peak in LISA's sensitivity sweet spot (Caprini *et al.* 2016 ; Caprini 2020). This makes LISA an ideal probe of first-order cosmological phase transitions and high-energy physics in general. Indeed, mHz frequencies correspond to beyond-TeV energy scales that we cannot probe with artificial particle colliders. They would therefore allow us to put the standard model of particle physics to the test at energies currently unexplored.

Second, very thin topological defects called cosmic strings could have formed as leftovers of a phase transition, similarly to the filamentary cracks that appear when water solidifies. As the Universe evolves, the network of cosmic strings would act as oscillating closed loops and source a SGWB that LISA could detect (Auclair *et al.* 2020).

Third, hypothetical primordial black holes could have formed in the early Universe, producing second-order density fluctuations at the end of inflation and thereby a relic background of GWs. Primordial black holes are candidates for dark matter. Although a significant part of possible masses has been experimentally ruled out, LISA could probe an uncharted mass range between $10^{-16} M_{\odot}$ and $10^{-10} M_{\odot}$ (Bartolo *et al.* 2019).




Figure 6.17: Power spectra of possible cosmological and astrophysical stochastic GW backgrounds (colored solid lines) compared to the noise level of one TDI channel (black dashed line). Green: example of a first-order phase transition signal; Red: primordial black hole signal; dark blue: stellar-mass black hole population signal; Orange: cosmic string signal; Light blue: Galactic white dwarf foreground; Dashed gray: detection threshold for a power-law background. Figure taken from Colpi (2024).

Fig 6.17 represents power spectra of possible sources of SGWBs that LISA could detect. This plot illustrates the variety of stochastic sources that could emit in the mHz band and the wide range of possible amplitudes. Based on the spectral features of these models, LISA should be able to determine the origin of the backgrounds and to set constraints on early-Universe physics. Note that although we do not represent it on this graph, a possible background from the population of extragalactic white dwarfs could also be present in LISA data, potentially larger than the background from stellar-mass black

hole binaries (Boileau *et al.* 2025). Likewise, unresolved EMRIs may lead to a detectable background, especially significant between 1 and 10 mHz (Pozzoli *et al.* 2023).

6.3. From data to science

LISA will be the first mission of its kind. To derive the expected scientific results described in Sec. 6.2 from the raw interferometric measurements that the satellites will send down to Earth as mentioned in Sec. 6.1, a very specific data processing will be developed. We outline its main aspects in this section.

6.3.1. Noise reduction

LISA will provide a set of interferometric measurements that are variations of the laser frequencies as a function of time. As described in Sec. 6.1.2.1, some of these measurements (called science interferometric data) contain the imprint of GWs passing through the constellation. However, they are completely dominated by laser frequency noise, i.e., the random fluctuations of the laser frequencies in time. This comes from the nature of the LISA instrument.

In a classic Michelson interferometer, a laser source sends light to a beam splitter, which separates the light into two beams that travel along two arms. These beams then reflect off mirrors located at the ends of the arms and are optically recombined at the beam splitter. There, a photodiode performs a differential measurement. The optimal path noise related to laser random fluctuations is directly proportional to the mismatch in the length of the arms. This is essentially what happens with ground-based GW detectors.

However, in the case of LISA, it is infeasible to use a single laser and have its light reflected by distant mirrors. The immense distances between the spacecraft (millions of kilometres) significantly reduce the light intensity and widen the beam to such an extent that it is impossible to construct a conventional Michelson apparatus. Instead, each one-way link formed by two spacecraft requires its own light source. A link measurement is made by comparing the outgoing local beam with the incoming distant beam from the other spacecraft. As the two beams being compared come from different sources, there is no cancellation of laser frequency noise. Additionally, the incoming beam travels a long distance, accumulating optical path differences and resulting in a laser frequency noise that dominates the GW signal by 8 to 9 orders of magnitude. Therefore, it would be impossible to detect any gravitational wave from a single spacecraft.

However, the constellation as a whole provides two measurements per pair of laser beams. Therefore, by adequately forming differences of such measurements delayed by the appropriate light travel times, one can cancel laser frequency noise. This processing is called TDI (Tinto and Dhurandhar 2020) and is performed on the ground, once the telemetry is downloaded. This is a crucial operation, without which GW detection is impossible. The output is a set of three laser noise-free variables that contain the GW signal information. Indeed, the TDI algorithm does not suppress the GW signal. Other sources of noise (called secondary noises) are not cancelled either, like the low-frequency acceleration noise and the high-frequency interferometric measurement noise mentioned in Sec. 6.1.2.1. Some TDI variables called “Michelson”, and labelled X , Y , Z , combine measurements in such a way as to synthetically reproduce the path of a photon in a Michelson interferometer formed by two LISA arms.

6.3.2. Source detection and inference

LISA scientific data can concentrate in 3 TDI observables, which are time series sampled every quarter of a second. As for ground-based detectors, there is no way of pointing the antenna towards particular regions in the sky. Rather, one can think of the data as continuous microphone recordings of several sounds and noises mixed. The constellation’s sensitivity to any sky location varies throughout the orbit.

Finding GW candidates relies on three ingredients: i) the input TDI data, ii) waveform templates (physical models of the signal produced by each source type) and iii) a likelihood function (i.e. a statistical model of the random fluctuations of data due to underlying stochastic processes). Even if certain machine learning algorithms may not need a likelihood, a physical and statistical description of the data is always required for training (Alvey *et al.* 2023).

The specificity of LISA data analysis is related to the abundance and the diversity of the sources. We expect future LISA data to contain millions of continuous Galactic binary signals, hundreds of transient MBH binary signals, possibly the same amount of EMRIs, a few sBH binaries, and the long-seek SGWBs. All these sources will emit simultaneously, adding up in the data, creating a GW cacophony. The objective of the data analysis is twofold: i) detect and identify all resolvable or stochastic sources of GWs with sufficient SNR and ii) characterize them through the estimation of their astrophysical parameters.

Contrary to LIGO-Virgo-KAGRA (LVK) data analysis, LISA will not measure isolated transient events that can be characterized independently. Instead, many sources overlap in time and frequency, and require a coordinated data

analysis approach called *global fit* where all sources are jointly analysed (Littenberg and Cornish 2023 ; Strub *et al.* 2024 ; Deng *et al.* 2025), see Sec. 6.3.2.4. To tackle this problem, Bayesian inference is preferred, as it allows for a better tracking of parameter correlations and source significance (although maximum-likelihood approaches have also been developed, see Strub *et al.* (2023)). In this framework, the goal of the search algorithms is to map the posterior probability distribution of astrophysical parameters, i.e., their probability of having a given value based on the observed data and the a priori information assumed before the analysis.

In general, waveforms generated by binary compact objects are determined by 17 parameters (Katz *et al.* 2021): the objects masses (2 parameters), each component's angular momentum (2×3), orbital parameters including eccentricity, separation and orbit inclination (3), initial phases (3), the source's sky location angles (2) and its distance (1). Many waveform models use simplifications or rearrangements of these parameters to achieve efficiency, but a full description of a relativistic system in vacuum would use the entire set. Obviously, including the effect of a third body or disturbances from the astrophysical environment requires additional parameters (Duque *et al.* 2024).

6.3.2.1. Continuous sources

Continuous sources usually designate the ones that produce long-lived and slowly chirping GW signals.




Figure 6.18: Example of waveform of a Galactic white dwarf binary with gravitational wave frequency around 10 mHz, observed over a few hours (left) and an entire year (right) in the TDI observable X (relative frequency deviations). We see the quasi-monochromatic nature on short time scales, with amplitude modulations due to the variation of the antenna pattern over one orbit.

Galactic binaries. The most common ones in LISA data are (comparable-mass) Galactic compact star binaries whose waveforms look monochromatic

on short time scales (see Fig. 6.18). Compared to the general full description of a relativistic binary system, Galactic binary waveforms can be simplified by ignoring the effect of component spins and describing their phase evolution with a Taylor expansion of their frequency about some initial time (Cornish and Porter 2005).

With frequencies ranging between a few 10^{-4} Hz and a few 10^{-3} Hz (see Fig. 6.11), their bandwidth is narrow and depends on the source frequency derivative, sky location, and observation duration (Cornish and Littenberg 2007). Searching for Galactic binary signals may require partitioning the frequency domain into short bands analysed in parallel (Littenberg 2011). The number of detectable sources within each band is unknown, so that any search algorithm must be able to jump between parameter dimensions, to allow for adjusting the number of Galactic binaries in the model (Littenberg and Cornish 2009). This can be performed by trans-dimensional samplers like reversible-jump Markov-chain Monte Carlo (MCMC) (Karnesis *et al.* 2023) or product-space approaches (Deng *et al.* 2025).

One challenge of detecting Galactic binaries is to deal with the millions of unresolvable signals, which form a confusion foreground that is non-Gaussian at high frequencies and non-stationary (Buscicchio *et al.* 2024). The detection threshold (to claim with high certainty that a source is present in the data) strongly depends on how this residual astrophysical noise is modelled.

Extreme-mass ratio inspirals. EMRIs are another type of continuous sources, which can be observed for months or years. These systems can be highly precessing and eccentric, leading to multiple harmonics (or voices) in the waveform (Hughes *et al.* 2021 ; Katz *et al.* 2021).

Due to their complexity, building waveform models that are both accurate and fast to compute is challenging. There has been continuous progress in this perspective since the early 2000s. The self-force formalism provides the most accurate models (Barack and Pound 2018), where the effect of the spacetime perturbation of the small compact object is treated using black hole perturbation theory, expanding equations for small mass ratios. Other models called “kludge” have focused on capturing the main physical features of the waveforms while achieving computational efficiency. Analytic kludge models extend the post-Newtonian formalism beyond its domain of validity (Barack and Cutler 2004). While still inaccurate, augmentations have been done to improve the model (Chua and Gair 2015). Another kludge, called numerical, treats the compact object orbit with Kerr geodesics, evolving the inspiral from geodesic to geodesic based on semi-analytic fits to strong-field radiation emission (Babak *et al.* 2007). Efforts have recently focused on speeding up fully relativistic EMRI




Figure 6.19: Simulation of an EMRI waveform in time domain (left) and time-frequency domain (right) observed through the Michelson TDI combination X. Simulated data from the LISA Data Challenge 1 (Laboratoire APC 2019).

waveforms based on self-force using order-reduction and deep-learning techniques, together with hardware acceleration (Chua *et al.* 2021 ; Katz *et al.* 2021). Further development is ongoing to take waveforms to the performance level needed to optimally analyse future LISA data.

Another active area of research concerns the detection and characterisation of EMRI signals, which faces two main challenges. First, the a priori intervals of parameter values are vast compared to the posterior distribution width, which requires an efficient and comprehensive exploration of a wide parameter space (Babak *et al.* 2009). Second, the likelihood function of physical parameters exhibits secondary non-local maxima, which means that a set of parameters potentially very different from the actual ones may describe the data with high probability (Chua and Cutler 2022). Efforts have started to tackle both problems, using phenomenological templates (Wang *et al.* 2012), time-frequency analysis (Gair *et al.* 2008*b,a*), and specific likelihood functions (Chua 2022). However, efficiently detecting one or several EMRI signals blended in all other LISA sources remains an open problem.

6.3.2.2. Transient sources

The fastest chirping sources LISA can observe are MBHBs, which can include inspiral, merger, and ringdown phases. GW templates for comparable-mass binary sources benefit from algorithm developments that have enabled LVK event detections. The main difference with the mergers that ground-based detectors observe is the event time scales. While LVK observations last a fraction of a second for black hole binaries and up to hundreds of seconds for neutron star binaries, LISA may observe MBH binaries for weeks to months. During that time, the constellation's position relative to the GW source changes and cannot be ignored when computing the antenna pattern. This leads to a




Figure 6.20: Simulation of an MBHB waveform in time domain (left) and time-frequency domain (right) observed through the Michelson TDI combination X. Simulated data from the LISA Data Challenge 1 (Laboratoire APC 2019).

time-varying instrumental response that must be included in the waveform. Wise approximations allow for the construction of fast frequency-domain templates for LISA (Marsat and Baker 2018) based on phenomenological waveforms featuring higher modes (London *et al.* 2018). Recently, the use of graphical processing unit acceleration allowed to build accurate and fast waveforms in the time domain (Katz *et al.* 2022 ; García-Quirós *et al.* 2025). Still, given the large SNR of some LISA sources, current model accuracy has to improve by two orders of magnitude (Colpi 2024) to enable the mission’s full scientific return, especially for what concerns GR tests.

The search phase of MBHB signals is usually based on waveform approximation and analytical maximization of the likelihood function with respect to amplitude parameters (Cornish and Shuman 2020). The full source characterization can then be done through Bayesian inference using MCMC sampling techniques (Marsat *et al.* 2021), nested sampling (Pitte *et al.* 2023), or simulation-based inference (Arredondo *et al.* 2024).

6.3.2.3. Stochastic sources

Stochastic GWs may arise from multiple sources incoherently emitting at the same time. The resulting signal is statistical and can be described by its power spectrum (second-order moment). Current terrestrial detectors already search for SGWBs, especially those originating from the population of stellar-mass black hole binaries in the Universe (Christensen 2019). The LISA mission will also carry out this search at mHz frequencies, especially with the hope of finding a SGWB of primordial origin (see Sec. 6.2.4). However, if LISA is the only flying detector in 2035, distinguishing a GW stochastic background from instrumental noise is particularly challenging. Indeed, while one can combine

ground-based observations coming from different detectors to reject instrumental noises (which are assumed uncorrelated across detectors, at least at current sensitivities), LISA must only rely on itself. One may still perform such a measurement by leveraging the spectral shapes and the transfer functions differences between the noise and the targetted SGWB. This requires a descriptive and flexible enough instrumental model.

Continuous progress has been made since the early 2000s to address the problem of SGWB detection with LISA data (Cornish 2001). Most of the works towards this goal rely on Bayesian parameter estimation of power spectrum templates. They are mainly based on parametrized instrumental noise models, sometimes including a Galactic foreground (Adams and Cornish 2014 ; Boileau *et al.* 2021). Given the variety of possible spectral shapes, template banks usable for data analysis have been developed (Blanco-Pillado *et al.* 2024 ; Braglia *et al.* 2024), while other studies have proposed template-free signal searches (Caprini *et al.* 2019 ; Poletti 2021 ; Pozzoli *et al.* 2024). Other works have focused on flexible noise modelling, highlighting the sensitivity of SGWB detection to noise characterization (Baghi *et al.* 2023 ; Hartwig *et al.* 2023). Recent analyses also started to incorporate the Galactic foreground non-stationarity (Hindmarsh *et al.* 2024 ; Criswell *et al.* 2024), and explore the detectability of anisotropic (Mentasti *et al.* 2024 ; Tian *et al.* 2024) or non-Gaussian (Buscicchio *et al.* 2023) backgrounds. Further research must be carried out to ensure the robustness of the detection of stochastic GWs with LISA. Model improvements are required to cope with realistic instrumental configurations and the presence of multiple other sources (Rosati and Littenberg 2024).

6.3.2.4. *The global fit orchestration*

LISA will observe all the sources reviewed in the previous sections simultaneously. Therefore, the data processing must disentangle these multiple signals but also determine how many of them are in the blend. Iterative procedures are usually adapted to this type of analysis, such as blocked-Gibbs sampling. Starting with the largest SNR sources, they can sequentially treat one block of parameter space at a time, such as a collection of GW sources belonging to the same type. This process can take advantage of the specific features of one block of signals, like their localisation in time or frequency. The process must be repeated until convergence. Ideally, each block of parameters is minimally coupled to the others, so that convergence is reached faster. The ultimate goal is to explore the posterior probability distribution of the full parameter space. The process that achieves this is designated as global fit.

The capability of data analysis pipelines to recover astrophysical GW sources has been tested thanks to mock data challenges that started in the early 2000s (Abbott 2017) and continued later on (Baghi 2022). The

principle of the LISA Data Challenges is to publish simulated LISA data to enable algorithm prototyping and testing. Some datasets contain GW sources of a single type while others mix several source types, gradually increasing complexity to reach realistic scenarios of future LISA data analysis. Fig. 6.21 provides an example of global fit results (Littenberg and Cornish 2023) obtained on simulated LISA data (Le Jeune and Babak 2022). The zoomed-in version of the right panel clearly shows the overlap between sources of the same type (Galactic binaries in purple) and of different types (Galactic binaries and merging MBHBs in magenta). The fit was performed by sequentially running three main inference blocks: Galactic binaries, MBHBs, and noise. In this case, “noise” refers to the fit residuals including instrumental noise and unresolved GW sources. More information can be found in Littenberg and Cornish (2023).




Figure 6.21: Reconstructed waveforms based on a global fit parameter posteriors based on the LISA Data Challenges dataset 2a. It shows Fourier transform amplitudes as a function of frequency for resolvable Galactic binaries (purple), verification Galactic binaries (orange), and massive black hole binaries (magenta). Fit residuals are represented in gray, and the estimated noise spectrum in light blue. The right panel is a zoomed version of the left panel between 6.210 and 6.245 mHz. Figures taken from Colpi (2024).

While many building blocks and prototypes have been developed, a complete global fit pipeline including all known GW sources has not yet been demonstrated. Developments are ongoing to achieve this, both within the scientific community and the teams in charge of building the mission data processing centres.

6.4. Bibliography

Abbott, B. P. a. o. (2017), Upper Limits on the Stochastic Gravitational-Wave

Background from Advanced LIGO's First Observing Run, *Physical Review Letters*, 118(12), 121101–121101. Publisher: American Physical Society.

URL: <https://link.aps.org/doi/10.1103/PhysRevLett.118.121101>

Adams, M. R., Cornish, N. J. (2014), Detecting a stochastic gravitational wave background in the presence of a galactic foreground and instrument noise, *Physical Review D*, 89(2), 022001. Publisher: American Physical Society.

URL: <https://link.aps.org/doi/10.1103/PhysRevD.89.022001>

Aghanim, N. a. o. (2020), Planck 2018 results - VI. Cosmological parameters, *Astronomy & Astrophysics*, 641, A6. Publisher: EDP Sciences.

URL: <https://www.aanda.org/articles/aa/abs/2020/09/aa33910-18/aa33910-18.html>

Akeson, R. a. o. (2019), 'The Wide Field Infrared Survey Telescope: 100 Hubbles for the 2020s'. ADS Bibcode: 2019arXiv190205569A.

URL: <https://ui.adsabs.harvard.edu/abs/2019arXiv190205569A>

Alvey, J., Bhardwaj, U., Domcke, V., Pieroni, M., Weniger, C. (2023), 'Simulation-based inference for stochastic gravitational wave background data analysis'. arXiv:2309.07954 [astro-ph, physics:gr-qc, physics:hep-ph].

URL: <http://arxiv.org/abs/2309.07954>

Amaro-Seoane, P. a. o. (2023), Astrophysics with the Laser Interferometer Space Antenna, *Living Reviews in Relativity*, 26(1), 2.

URL: <https://doi.org/10.1007/s41114-022-00041-y>

Armano, M., Audley, H., Baird, J., Binetruy, P., Born, M., Bortoluzzi, D., Castelli, E., Cavalleri, A., Cesarini, A., Cruise, A. M., Danzmann, K., de Deus Silva, M., Diepholz, I., Dixon, G., Dolesi, R., Ferraioli, L., Ferroni, V., Fitzsimons, E. D., Freschi, M., Gesa, L., Gibert, F., Giardini, D., Giusteri, R., Grimani, C., Grzymisch, J., Harrison, I., Heinzl, G., Hewitson, M., Hollington, D., Hoyland, D., Hueller, M., Inchauspé, H., Jennrich, O., Jetzer, P., Karnesis, N., Kaune, B., Korsakova, N., Killow, C. J., Lobo, J. A., Lloro, I., Liu, L., López-Zaragoza, J. P., Maarschalkerweerd, R., Mance, D., Meshksar, N., Martín, V., Martín-Polo, L., Martino, J., Martin-Porqueras, F., Mateos, I., McNamara, P. W., Mendes, J., Mendes, L., Nofrarias, M., Paczkowski, S., Perreux-Lloyd, M., Petiteau, A., Pivato, P., Plagnol, E., Ramos-Castro, J., Reiche, J., Robertson, D. I., Rivas, F., Russano, G., Slutsky, J., Sopuerta, C. F., Sumner, T., Texier, D., Thorpe, J. I., Vetrugno, D., Vitale, S., Wanner, G., Ward, H., Wass, P. J., Weber, W. J., Wissel, L., Wittchen, A., Zweifel, P. (2018), Beyond the Required LISA Free-Fall Performance: New LISA Pathfinder Results down to 20 μHz , *Physical review letters*, 120, 061101.

URL: <https://doi.org/10.1103/PhysRevLett.120.061101>

- Arredondo, J. N., Klein, A., Yunes, N. (2024), ‘Efficient Gravitational-Wave Model for Fully-Precessing and Moderately-Eccentric, Compact Binary Inspirals’. arXiv:2402.06804 [astro-ph, physics:gr-qc].
URL: <http://arxiv.org/abs/2402.06804>
- Arun, K. G. a. o. (2022), New horizons for fundamental physics with LISA, *Living Reviews in Relativity*, 25(1), 4.
URL: <https://doi.org/10.1007/s41114-022-00036-9>
- Auclair, P., Blanco-Pillado, J. J., Figueroa, D. G., Jenkins, A. C., Lewicki, M., Sakellariadou, M., Sanidas, S., Sousa, L., Steer, D. A., Wachter, J. M., Kuroyanagi, S. (2020), Probing the gravitational wave background from cosmic strings with LISA, *Journal of Cosmology and Astroparticle Physics*, 2020(04), 034–034. Publisher: IOP Publishing.
URL: <https://doi.org/10.1088/1475-7516/2020/04/034>
- Babak, S., Fang, H., Gair, J. R., Glampedakis, K., Hughes, S. A. (2007), “Kludge” gravitational waveforms for a test-body orbiting a Kerr black hole, *Physical Review D*, 75(2), 024005. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.75.024005>
- Babak, S., Gair, J. R., Porter, E. K. (2009), An algorithm for the detection of extreme mass ratio inspirals in LISA data, *Classical and Quantum Gravity*, 26(13), 135004–135004. Publisher: Institute of Physics Publishing.
URL: <https://iopscience.iop.org/article/10.1088/0264-9381/26/13/135004>
- Baghi, Q. (2022), ‘The LISA Data Challenges’. arXiv:2204.12142 [gr-qc].
URL: <http://arxiv.org/abs/2204.12142>
- Baghi, Q., Karnesis, N., Bayle, J.-B., Besançon, M., Inchauspé, H. (2023), Uncovering gravitational-wave backgrounds from noises of unknown shape with LISA, *Journal of Cosmology and Astroparticle Physics*, 2023(04), 066. Publisher: IOP Publishing.
URL: <https://dx.doi.org/10.1088/1475-7516/2023/04/066>
- Barack, L., Cutler, C. (2004), LISA capture sources: Approximate waveforms, signal-to-noise ratios, and parameter estimation accuracy, *Physical Review D*, 69(8), 082005. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.69.082005>
- Barack, L., Pound, A. (2018), Self-force and radiation reaction in general relativity, *Reports on Progress in Physics*, 82(1), 016904. Publisher: IOP Publishing.
URL: <https://dx.doi.org/10.1088/1361-6633/aae552>
- Barsanti, S., Maselli, A., Sotiriou, T. P., Gualtieri, L. (2023), Detecting Massive Scalar Fields with Extreme Mass-Ratio Inspirals, *Physical Review Letters*, 131(5), 051401. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.131.051401>

- Bartolo, N., De Luca, V., Franciolini, G., Peloso, M., Racco, D., Riotto, A. (2019), Testing primordial black holes as dark matter with LISA, *Physical Review D*, 99(10), 103521. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.99.103521>
- Bender, P., Brillet, A., Danzmann, K., Ciufolini, I., Cruise, M., Folkner, W., Fidecaro, F., Hough, J., McNamara, P. (1998), LISA - Pre-phase A report, Technical report, European Space Agency.
- Berry, C. P. L., Hughes, S. A., Sopuerta, C. F., Chua, A. J. K., Heffernan, A., Holley-Bockelmann, K., Mihaylov, D. P., Miller, M. C., Sesana, A. (2019), ‘The unique potential of extreme mass-ratio inspirals for gravitational-wave astronomy’. arXiv:1903.03686 [astro-ph].
URL: <http://arxiv.org/abs/1903.03686>
- Blanco-Pillado, J. J., Cui, Y., Kuroyanagi, S., Lewicki, M., Nardini, G., Pieroni, M., Rybak, I. Y., Sousa, L., Wachter, J. M. (2024), ‘Gravitational waves from cosmic strings in LISA: reconstruction pipeline and physics interpretation’. arXiv:2405.03740 [astro-ph, physics:gr-qc, physics:hep-th].
URL: <http://arxiv.org/abs/2405.03740>
- Bogdan, A., Goulding, A. D., Natarajan, P., Kovacs, O. E., Tremblay, G. R., Chadayammuri, U., Volonteri, M., Kraft, R. P., Forman, W. R., Jones, C., Churazov, E., Zhuravleva, I. (2024), Evidence for heavy-seed origin of early supermassive black holes from a $z=10$ X-ray quasar, *Nature Astronomy*, 8(1), 126–133. Publisher: Nature Publishing Group.
URL: <https://www.nature.com/articles/s41550-023-02111-9>
- Boileau, G., Bruel, T., Toubiana, A., Lamberts, A., Christensen, N. (2025), ‘Gravitational-Wave Background from Extragalactic Double White Dwarfs for LISA’.
- Boileau, G., Christensen, N., Meyer, R., Cornish, N. J. (2021), Spectral separation of the stochastic gravitational-wave background for LISA: Observing both cosmological and astrophysical backgrounds, *Physical Review D*, 103(10), 103529. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.103529>
- Braglia, M., Calcagni, G., Franciolini, G., Fumagalli, J., Nardini, G., Peloso, M., Pieroni, M., Renaux-Petel, S., Ricciardone, A., Tasinato, G., Vaskonen, V. (2024), ‘Gravitational waves from inflation in LISA: reconstruction pipeline and physics interpretation’. arXiv:2407.04356.
URL: <http://arxiv.org/abs/2407.04356>
- Breivik, K., Mingarelli, C. M. F., Larson, S. L. (2020), Constraining Galactic Structure with the LISA White Dwarf Foreground, *The Astrophysical Journal*, 901(1), 4. Publisher: The American Astronomical Society.
URL: <https://dx.doi.org/10.3847/1538-4357/abab99>

- Brown, W. R., Kilic, M., Kosakowski, A., Andrews, J. J., Heinke, C. O., Agüeros, M. A., Camilo, F., Gianninas, A., Hermes, J. J., Kenyon, S. J. (2020), The ELM Survey. VIII. Ninety-eight Double White Dwarf Binaries, *The Astrophysical Journal*, 889(1), 49.
- Buscicchio, R., Ain, A., Ballelli, M., Cella, G., Patricelli, B. (2023), Detecting non-Gaussian gravitational wave backgrounds: A unified framework, *Physical Review D*, 107(6), 063027. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.107.063027>
- Buscicchio, R., Klein, A., Korol, V., Renzo, F. D., Moore, C. J., Gerosa, D., Carzaniga, A. (2024), ‘A test for LISA foreground Gaussianity and stationarity. I. Galactic white-dwarf binaries’. arXiv:2410.08263.
URL: <http://arxiv.org/abs/2410.08263>
- Callister, T., Biscoveanu, A. S., Christensen, N., Isi, M., Matas, A., Minazzoli, O., Regimbau, T., Sakellariadou, M., Tasson, J., Thrane, E. (2017), Polarization-Based Tests of Gravity with the Stochastic Gravitational-Wave Background, *Physical Review X*, 7(4), 041058. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevX.7.041058>
- Caprini, C. a. o. (2020), Detecting gravitational waves from cosmological phase transitions with LISA: an update, *Journal of Cosmology and Astroparticle Physics*, 2020(03), 024.
URL: <https://dx.doi.org/10.1088/1475-7516/2020/03/024>
- Caprini, C., Figueroa, D. G., Flauger, R., Nardini, G., Peloso, M., Pieroni, M., Ricciardone, A., Tasinato, G. (2019), Reconstructing the spectral shape of a stochastic gravitational wave background with LISA, *Journal of Cosmology and Astroparticle Physics*, 2019(11), 017. Publisher: IOP Publishing.
URL: <https://iopscience.iop.org/article/10.1088/1475-7516/2019/11/017/meta>
- Caprini, C., Hindmarsh, M., Huber, S., Konstandin, T., Kozaczuk, J., Nardini, G., No, J. M., Petiteau, A., Schwaller, P., Servant, G., Weir, D. J. (2016), Science with the space-based interferometer eLISA. II: gravitational waves from cosmological phase transitions, *Journal of Cosmology and Astroparticle Physics*, 2016(04), 001.
URL: <https://dx.doi.org/10.1088/1475-7516/2016/04/001>
- Carilli, C., Rawlings, S. (2004), Science with the Square Kilometer Array: Motivation, Key Science Projects, Standards and Assumptions, *New Astronomy Reviews*, 48(11-12), 979–984. arXiv:astro-ph/0409274.
URL: <http://arxiv.org/abs/astro-ph/0409274>
- Carr, B., Silk, J. (2018), Primordial black holes as generators of cosmic structures, *Monthly Notices of the Royal Astronomical Society*, 478(3), 3756–3775.
URL: <https://doi.org/10.1093/mnras/sty1204>

- Chamberlain, K., Yunes, N. (2017), Theoretical physics implications of gravitational wave observation with future detectors, *Physical Review D*, 96(8), 084039. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.96.084039>
- Christensen, N. (2019), Stochastic gravitational wave backgrounds, *Reports on Progress in Physics*, 82(1), 016903–016903. Publisher: Institute of Physics Publishing.
URL: <https://iopscience.iop.org/article/10.1088/1361-6633/aae6b5>
- Chua, A. J., Cutler, C. J. (2022), Nonlocal parameter degeneracy in the intrinsic space of gravitational-wave signals from extreme-mass-ratio inspirals, *Physical Review D*, 106(12), 124046. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.106.124046>
- Chua, A. J. K. (2022), One-stop function for gravitational-wave detection, identification, and inference, *Phys. Rev. D*, 106, 104051.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.106.104051>
- Chua, A. J. K., Gair, J. R. (2015), Improved analytic extreme-mass-ratio inspiral model for scoping out eLISA data analysis, *Classical and Quantum Gravity*, 32(23), 232002. Publisher: IOP Publishing.
URL: <https://dx.doi.org/10.1088/0264-9381/32/23/232002>
- Chua, A. J., Katz, M. L., Warburton, N., Hughes, S. A. (2021), Rapid Generation of Fully Relativistic Extreme-Mass-Ratio-Inspiral Waveform Templates for LISA Data Analysis, *Physical Review Letters*, 126(5), 051102. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.126.051102>
- Colpi, M. a. o. (2024), ‘LISA Definition Study Report’. arXiv:2402.07571 [astro-ph].
URL: <http://arxiv.org/abs/2402.07571>
- Cornish, N. J. (2001), Detecting a stochastic gravitational wave background with the Laser Interferometer Space Antenna, , .
URL: <https://arxiv.org/pdf/gr-qc/0106058.pdf>
- Cornish, N. J., Littenberg, T. B. (2007), Tests of Bayesian model selection techniques for gravitational wave astronomy, *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 76(8).
URL: <https://journals.aps.org/prd/pdf/10.1103/PhysRevD.76.083006>
- Cornish, N. J., Porter, E. K. (2005), Detecting galactic binaries with LISA, vol. 22, pp. 927–933. Issue: 18.
URL: <http://iopscience.iop.org/article/10.1088/0264-9381/22/18/S06/pdf>
- Cornish, N. J., Shuman, K. (2020), Black hole hunting with LISA, *Physical Review D*, 101(12), 124008. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.101.124008>

- Criswell, A. W., Rieck, S., Mandic, V. (2024), ‘Templated Anisotropic Analyses of the LISA Galactic Foreground’. arXiv:2410.23260.
URL: <http://arxiv.org/abs/2410.23260>
- Cruise, M. a. o. (2025), The NewAthena mission concept in the context of the next decade of X-ray astronomy, *Nature Astronomy*, 9(1), 36–44. Publisher: Nature Publishing Group.
URL: <https://www.nature.com/articles/s41550-024-02416-3>
- Danzmann, K. (1996), LISA: laser interferometer space antenna for gravitational wave measurements, *Classical and Quantum Gravity*, 13, A247.
- Danzmann, K. (2000), LISA mission overview, *Advances in Space Research*, 25, 1129–1136.
URL: <https://linkinghub.elsevier.com/retrieve/pii/S0273117799009734>
- Danzmann, K. (2003), LISA - An ESA cornerstone mission for the detection and observation of gravitational waves, *Advances in Space Research*, 32, 1233–1242.
URL: <https://linkinghub.elsevier.com/retrieve/pii/S0273117703903231>
- Danzmann, K. a. o. (2013), The Gravitational Universe, *eprint arXiv*, 1305, 5720.
URL: <https://doi.org/10.48550/arXiv.1305.5720>
- Danzmann, K. a. o. (2017), Laser Interferometer Space Antenna, *eprint arXiv*, 1702, arXiv:1702.00786.
URL: <http://arxiv.org/abs/1702.00786>
- Danzmann, K. *et al.* (1997), LISA - an ESA cornerstone mission for a gravitational wave observatory, *Classical and Quantum Gravity*, 14, 1399–1404.
URL: <http://stacks.iop.org/0264-9381/14/i=6/a=002?key=crossref.1670e4f81d08cd491cfbd41058f0812f>
- d’Ascoli, S., Noble, S. C., Bowen, D. B., Campanelli, M., Krolik, J. H., Mewes, V. (2018), Electromagnetic emission from supermassive binary black holes approaching merger, *The Astrophysical Journal*, 865(2), 140. Publisher: The American Astronomical Society.
URL: <https://dx.doi.org/10.3847/1538-4357/aad8b4>
- Dayal, P., Rossi, E. M., Shiralilou, B., Piana, O., Choudhury, T. R., Volonteri, M. (2019), The hierarchical assembly of galaxies and black holes in the first billion years: predictions for the era of gravitational wave astronomy, *Monthly Notices of the Royal Astronomical Society*, 486(2), 2336–2350.
URL: <https://doi.org/10.1093/mnras/stz897>
- Decher, R., Randall, J. L., Bender, P. L., Faller, J. E. (1980), Design Aspects Of A Laser Gravitational Wave Detector In Space, *in* W. J. Cuneo, (ed.), Proc. SPIE 0228, Active Optical Devices and Applications, Washington, D.C., pp. 149–153.
URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=1229916>

- Deng, S., Babak, S., Jeune, M. L., Marsat, S., Plagnol, E., Sartirana, A. (2025), ‘Modular global-fit pipeline for LISA data analysis’. arXiv:2501.10277 [gr-qc].
URL: <http://arxiv.org/abs/2501.10277>
- Doneva, D. D., Ramazano??lu, F. M., Silva, H. O., Sotiriou, T. P., Yazadjiev, S. S. (2024), Spontaneous scalarization, *Reviews of Modern Physics*, 96(1), 015004. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/RevModPhys.96.015004>
- Duque, F., Kejriwal, S., Sberna, L., Speri, L., Gair, J. (2024), ‘Constraining accretion physics with gravitational waves from eccentric extreme-mass-ratio inspirals’. arXiv:2411.03436.
URL: <http://arxiv.org/abs/2411.03436>
- Ezquiaga, J. M., Holz, D. E. (2022), Spectral Sirens: Cosmology from the Full Mass Distribution of Compact Binaries, *Physical Review Letters*, 129(6), 061102. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.129.061102>
- Fabian, A. C., Lasenby, A. N. (2019), ‘Astrophysical Black Holes’. arXiv:1911.04305 [astro-ph].
URL: <http://arxiv.org/abs/1911.04305>
- Faller, J., Bender, P., Hall, J., Hills, D., Vincent, M. A. (1985), Space antenna for gravitational wave astronomy, in *ESA Colloquium in Kilometric Optical Arrays in Space*, vol. SP-226, ESA, Cargèse (Corsica), pp. 157–163.
URL: <http://adsabs.harvard.edu/abs/1985ESASP.226..157F>
- Feng, W.-F., Chen, J.-W., Wang, Y., Mohanty, S. D., Shao, Y. (2023), ‘Multi-messenger observations of double neutron stars in Galactic disk with gravitational and radio waves’. arXiv:2305.05202 [astro-ph, physics:gr-qc].
URL: <http://arxiv.org/abs/2305.05202>
- Finch, E., Bartolucci, G., Chucherko, D., Patterson, B. G., Korol, V., Klein, A., Bandopadhyay, D., Middleton, H., Moore, C. J., Vecchio, A. (2023), Identifying LISA verification binaries among the Galactic population of double white dwarfs, *Monthly Notices of the Royal Astronomical Society*, 522(4), 5358–5373.
URL: <https://doi.org/10.1093/mnras/stad1288>
- Freedman, W. L., Madore, B. F., Jang, I. S., Hoyt, T. J., Lee, A. J., Owens, K. A. (2025), ‘Status Report on the Chicago-Carnegie Hubble Program (CCHP): Measurement of the Hubble Constant Using the Hubble and James Webb Space Telescopes’. arXiv:2408.06153 [astro-ph].
URL: <http://arxiv.org/abs/2408.06153>
- Gair, J. R., Mandel, I., Wen, L. (2008a), Improved time???frequency analysis of extreme-mass-ratio inspiral signals in mock LISA data, *Classical and Quantum Gravity*, 25(18), 184031.

URL: <https://dx.doi.org/10.1088/0264-9381/25/18/184031>

Gair, J. R., Mandel, I., Wen, L. (2008*b*), Time-frequency analysis of extreme-mass-ratio inspiral signals in mock LISA data, *Journal of Physics: Conference Series*, 122(1), 012037.

URL: <https://dx.doi.org/10.1088/1742-6596/122/1/012037>

García-Quirós, C., Tiwari, S., Babak, S. (2025), ‘GPU-accelerated LISA parameter estimation with full time domain response’. arXiv:2501.08261 [gr-qc].

URL: <http://arxiv.org/abs/2501.08261>

Gezari, S. (2021), Tidal Disruption Events, *Annual Review of Astronomy and Astrophysics*, 59(Volume 59, 2021), 21–58. Publisher: Annual Reviews.

URL: <https://www.annualreviews.org/content/journals/10.1146/annurev-astro-111720-030029>

Ghez, A. M., Salim, S., Weinberg, N. N., Lu, J. R., Do, T., Dunn, J. K., Matthews, K., Morris, M. R., Yelda, S., Becklin, E. E., Kremenek, T., Milosavljevic, M., Naiman, J. (2008), Measuring Distance and Properties of the Milky Way’s Central Supermassive Black Hole with Stellar Orbits, *The Astrophysical Journal*, 689(2), 1044. Publisher: IOP Publishing.

URL: <https://iopscience.iop.org/article/10.1086/592738/meta>

Glorio, S., Berti, E., Maselli, A., Speeney, N. (2025), ‘Extreme mass ratio inspirals in dark matter halos: dynamics and distinguishability of halo models’. arXiv:2503.16649 [gr-qc].

URL: <http://arxiv.org/abs/2503.16649>

Greene, J. E., Strader, J., Ho, L. C. (2020), Intermediate-Mass Black Holes, *Annual Review of Astronomy and Astrophysics*, 58(Volume 58, 2020), 257–312. Publisher: Annual Reviews.

URL: <https://www.annualreviews.org/content/journals/10.1146/annurev-astro-032620-021835>

Gualandris, A., Khan, F. M., Bortolas, E., Bonetti, M., Sesana, A., Berczik, P., Holley-Bockelmann, K. (2022), Eccentricity evolution of massive black hole binaries from formation to coalescence, *Monthly Notices of the Royal Astronomical Society*, 511(4), 4753–4765.

URL: <https://doi.org/10.1093/mnras/stac241>

Hartwig, O., Lilley, M., Muratore, M., Pieroni, M. (2023), Stochastic gravitational wave background reconstruction for a non-equilateral and unequal-noise LISA constellation, *Physical Review D*, 107(12), 123531. arXiv:2303.15929 [astro-ph, physics:gr-qc].

URL: <http://arxiv.org/abs/2303.15929>

Heger, A., Fryer, C. L., Woosley, S. E., Langer, N., Hartmann, D. H. (2003), How Massive Single Stars End Their Life, *The Astrophysical Journal*, 591(1), 288. Publisher: IOP Publishing.

URL: <https://iopscience.iop.org/article/10.1086/375341/meta>

- Hindmarsh, M., Hooper, D. C., Minkinen, T., Weir, D. J. (2024), ‘Recovering a phase transition signal in simulated LISA data with a modulated galactic foreground’. arXiv:2406.04894 [astro-ph, physics:gr-qc].
URL: <http://arxiv.org/abs/2406.04894>
- Holberg, J. B. (2009), What fraction of white dwarfs are members of binary systems?, *Journal of Physics: Conference Series*, 172(1), 012022.
URL: <https://dx.doi.org/10.1088/1742-6596/172/1/012022>
- Hughes, S. A., Warburton, N., Khanna, G., Chua, A. J., Katz, M. L. (2021), Adiabatic waveforms for extreme mass-ratio inspirals via multivoice decomposition in time and frequency, *Physical Review D*, 103(10), 104014. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.104014>
- Ivezić, Z. *et al.* (2019), LSST: From Science Drivers to Reference Design and Anticipated Data Products, *The Astrophysical Journal*, 873(2), 111. Publisher: The American Astronomical Society.
URL: <https://dx.doi.org/10.3847/1538-4357/ab042c>
- Jennrich, O., Binetruy, P., Colpi, M., Danzmann, K., Jetzer, P., Lobo, A., Nelemans, G., Schutz, B., Stebbins, R., Sumner, T., Vitale, S., Ward, H. (2011), NGO Assessment Study Report (Yellow Book), Technical Report ESA/SRE(2011)19, ESA.
- Karnesis, N., Katz, M. L., Korsakova, N., Gair, J. R., Stergioulas, N. (2023), Eryn: a multipurpose sampler for Bayesian inference, *Monthly Notices of the Royal Astronomical Society*, 526(4), 4814–4830.
URL: <https://doi.org/10.1093/mnras/stad2939>
- Katz, M. L., Bayle, J.-B., Chua, A. J., Vallisneri, M. (2022), Assessing the data-analysis impact of LISA orbit approximations using a GPU-accelerated response model, *Physical Review D*, 106(10), 103001. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.106.103001>
- Katz, M. L., Chua, A. J., Speri, L., Warburton, N., Hughes, S. A. (2021), Fast extreme-mass-ratio-inspiral waveforms: New tools for millihertz gravitational-wave data analysis, *Physical Review D*, 104(6), 064047. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.064047>
- Klein, A., Barausse, E., Sesana, A., Petiteau, A., Berti, E., Babak, S., Gair, J., Aoudia, S., Hinder, I., Ohme, F., Wardell, B. (2016), Science with the space-based interferometer eLISA: Supermassive black hole binaries, *Physical Review D*, 93(2), 024003. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.93.024003>
- Korol, V., Rossi, E. M., Barausse, E. (2019), A multimessenger study of the Milky Way’s stellar disc and bulge with LISA, Gaia, and LSST, *Monthly*

Notices of the Royal Astronomical Society, 483(4), 5518–5533.

URL: <https://doi.org/10.1093/mnras/sty3440>

Kupfer, T., Korol, V., Littenberg, T. B., Shah, S., Savalle, E., Groot, P. J., Marsh, T. R., Le Jeune, M., Nelemans, G., Pala, A. F., Petiteau, A., Ramsay, G., Steeghs, D., Babak, S. (2024), LISA Galactic Binaries with Astrometry from Gaia DR3, *The Astrophysical Journal*, 963(2), 100.

Kupfer, T., Prince, T. A., van Roestel, J., Bellm, E. C., Bildsten, L., Coughlin, M. W., Drake, A. J., Graham, M. J., Klein, C., Kulkarni, S. R., Masci, F. J., Walters, R., Andreoni, I., Biswas, R., Bradshaw, C., Duev, D. A., Dekany, R., Guidry, J. A., Hermes, J. J., Laher, R. R., Riddle, R. (2021), Year 1 of the ZTF high-cadence Galactic plane survey: Strategy, goals, and early results on new single-mode hot subdwarf B-star pulsators, *Monthly Notices of the Royal Astronomical Society*, 505(1), 1254–1267.

Laboratoire APC (2019), LISA Data Challenge 1: Radler, , pp. <https://lisa-ldc.lal.in2p3.fr/ldc>.

URL: <https://lisa-ldc.lal.in2p3.fr/ldc>

Laghi, D., Tamanini, N., Del Pozzo, W., Sesana, A., Gair, J., Babak, S., Izquierdo-Villalba, D. (2021), Gravitational-wave cosmology with extreme mass-ratio inspirals, *Monthly Notices of the Royal Astronomical Society*, 508(3), 4512–4531.

URL: <https://doi.org/10.1093/mnras/stab2741>

Lang, R. N., Hughes, S. A. (2009), Advanced localization of massive black hole coalescences with LISA, *Classical and Quantum Gravity*, 26(9), 094035. arXiv:0810.5125 [astro-ph].

URL: <http://arxiv.org/abs/0810.5125>

Le Jeune, M., Babak, S. (2022), ‘LISA data challenge sangria (ldc2a)’.

URL: <https://doi.org/10.5281/zenodo.7132178>

Le Jeune, M., Sartirana, A. (2025), ‘fomweb’.

URL: <https://doi.org/10.5281/zenodo.11504201>

Littenberg, T. B. (2011), Detection pipeline for Galactic binaries in LISA data, *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 84(6), 063009–063009. Publisher: American Physical Society.

URL: <https://link.aps.org/doi/10.1103/PhysRevD.84.063009>

Littenberg, T. B., Cornish, N. J. (2009), A Bayesian Approach to the Detection Problem in Gravitational Wave Astronomy, *Physical Review D*, 80(6), 063007. arXiv:0902.0368 [gr-qc].

URL: <http://arxiv.org/abs/0902.0368>

Littenberg, T. B., Cornish, N. J. (2023), Prototype global analysis of LISA data with multiple source types, *Physical Review D*, 107(6), 063004. Publisher: American Physical Society.

URL: <https://link.aps.org/doi/10.1103/PhysRevD.107.063004>

- Littenberg, T. B., Lali, A. K. (2024), ‘Have any LISA verification binaries been found?’. arXiv:2404.03046 [astro-ph, physics:gr-qc].
URL: <http://arxiv.org/abs/2404.03046>
- London, L., Khan, S., Fauchon-Jones, E., García, C., Hannam, M., Husa, S., Jiménez-Forteza, X., Kalaghatgi, C., Ohme, F., Pannarale, F. (2018), First Higher-Multipole Model of Gravitational Waves from Spinning and Coalescing Black-Hole Binaries, *Physical Review Letters*, 120(16), 161102. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.161102>
- Lops, G., Izquierdo-Villalba, D., Colpi, M., Bonoli, S., Sesana, A., Mangiagli, A. (2023), Galaxy fields of LISA massive black hole mergers in a simulated universe, *Monthly Notices of the Royal Astronomical Society*, 519(4), 5962–5986.
URL: <https://doi.org/10.1093/mnras/stad058>
- Mangiagli, A., Caprini, C., Volonteri, M., Marsat, S., Vergani, S., Tamanini, N., Inchauspé, H. (2022), Massive black hole binaries in LISA: Multimessenger prospects and electromagnetic counterparts, *Physical Review D*, 106(10), 103017. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.106.103017>
- Marsat, S., Baker, J. G. (2018), Fourier-domain modulations and delays of gravitational-wave signals, . . .
URL: <https://arxiv.org/pdf/1806.10734.pdf>
- Marsat, S., Baker, J. G., Canton, T. D. (2021), Exploring the Bayesian parameter estimation of binary black holes with LISA, *Physical Review D*, 103(8), 083011. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.083011>
- Mentasti, G., Contaldi, C. R., Peloso, M. (2024), ‘Probing the galactic and extragalactic gravitational wave backgrounds with space-based interferometers’. arXiv:2312.10792 [astro-ph, physics:gr-qc].
URL: <http://arxiv.org/abs/2312.10792>
- Mirshekari, S., Yunes, N., Will, C. M. (n.d.), Constraining Lorentz-Violating, Modified Dispersion Relations with Gravitational Waves, Technical report.
- Napiwotzki, R. (2009), The galactic population of white dwarfs, *Journal of Physics: Conference Series*, 172(1), 012004.
URL: <https://dx.doi.org/10.1088/1742-6596/172/1/012004>
- Nelemans, G., Yungelson, L. R., Zwart, S. F. P. (2001), The gravitational wave signal from the Galactic disk population of binaries containing two compact objects, *Astronomy & Astrophysics*, 375(3), 890–898. Number: 3 Publisher: EDP Sciences.
URL: <https://www.aanda.org/articles/aa/abs/2001/33/aah2754/aah2754.html>

- Paula, W. L. S. d., Miranda, O. D., Marinho, R. M. (2004), Polarization states of gravitational waves with a massive graviton, *Classical and Quantum Gravity*, 21(19), 4595–4605. Publisher: IOP Publishing.
URL: <https://doi.org/10.1088/0264-9381/21/19/008>
- Pitte, C., Baghi, Q., Besançon, M., Petiteau, A. (2024), Exploring tests of the no-hair theorem with LISA, *Physical Review D*, 110(10), 104003. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.110.104003>
- Pitte, C., Baghi, Q., Marsat, S., Besançon, M., Petiteau, A. (2023), Detectability of higher harmonics with LISA, *Physical Review D*, 108(4), 044053. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.108.044053>
- Poletti, D. (2021), Measuring the primordial gravitational wave background in the presence of other stochastic signals, , .
URL: <http://arxiv.org/abs/2101.02713>
- Pozzoli, F., Babak, S., Sesana, A., Bonetti, M., Karnesis, N. (2023), Computation of stochastic background from extreme-mass-ratio inspiral populations for lisa, *Phys. Rev. D*, 108, 103039.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.108.103039>
- Pozzoli, F., Buscicchio, R., Moore, C. J., Haardt, F., Sesana, A. (2024), Weakly parametric approach to stochastic background inference in LISA, *Physical Review D*, 109(8), 083029. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.109.083029>
- Regan, J., Volonteri, M. (2024), Massive Black Hole Seeds, *The Open Journal of Astrophysics*, 7. Publisher: Maynooth Academic Publishing.
URL: <https://astro.theoj.org/article/123239-massive-black-hole-seeds>
- Riess, A. G. a. o. (2022), A Comprehensive Measurement of the Local Value of the Hubble Constant with 1 km s⁻¹ Mpc⁻¹ Uncertainty from the Hubble Space Telescope and the SH0ES Team, *The Astrophysical Journal Letters*, 934(1), L7. Publisher: The American Astronomical Society.
URL: <https://dx.doi.org/10.3847/2041-8213/ac5c5b>
- Rosati, R., Littenberg, T. B. (2024), ‘Prototype Stochastic Gravitational Wave Background Recovery in the LISA Global Fit Residual’. arXiv:2410.17180 [gr-qc].
URL: <http://arxiv.org/abs/2410.17180>
- Rossi, E. M., Lodato, G., Armitage, P. J., Pringle, J. E., King, A. R. (2010), Black hole mergers: the first light, *Monthly Notices of the Royal Astronomical Society*, 401(3), 2021–2035.
URL: <https://doi.org/10.1111/j.1365-2966.2009.15802.x>
- Samuel Reich, E. (2011), Europe makes do without NASA, *Nature*, .
URL: <http://www.nature.com/articles/471421a>

- Shah, S., Nelemans, G. (2014), CONSTRAINING PARAMETERS OF WHITE-DWARF BINARIES USING GRAVITATIONAL-WAVE AND ELECTROMAGNETIC OBSERVATIONS, *The Astrophysical Journal*, 790(2), 161. Publisher: The American Astronomical Society.
URL: <https://dx.doi.org/10.1088/0004-637X/790/2/161>
- Speri, L., Antonelli, A., Sberna, L., Babak, S., Barausse, E., Gair, J. R., Katz, M. L. (2023), Probing Accretion Physics with Gravitational Waves, *Physical Review X*, 13(2), 021035. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevX.13.021035>
- Straub, O. a. o. (2023), Where intermediate-mass black holes could hide in the Galactic Centre - A full parameter study with the S2 orbit, *Astronomy & Astrophysics*, 672, A63. Publisher: EDP Sciences.
URL: <https://www.aanda.org/articles/aa/abs/2023/04/aa45132-22/aa45132-22.html>
- Strub, S. H., Ferraioli, L., Schmelzbach, C., Stähler, S. C., Giardini, D. (2023), Accelerating global parameter estimation of gravitational waves from Galactic binaries using a genetic algorithm and GPUs, *Physical Review D*, 108(10), 103018. Publisher: American Physical Society.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.108.103018>
- Strub, S. H., Ferraioli, L., Schmelzbach, C., Stähler, S. C., Giardini, D. (2024), ‘Global Analysis of LISA Data with Galactic Binaries and Massive Black Hole Binaries’. arXiv:2403.15318 [gr-qc].
URL: <http://arxiv.org/abs/2403.15318>
- Suh, H., Scharwächter, J., Farina, E. P., Loiacono, F., Lanzuisi, G., Hasinger, G., Marchesi, S., Mezcua, M., Decarli, R., Lemaux, B. C., Volonteri, M., Civano, F., Yi, S. K., Han, S., Rawlings, M., Hung, D. (2025), A super-Eddington-accreting black hole ~ 1.5 Gyr after the Big Bang observed with JWST, *Nature Astronomy*, 9(2), 271–279. Publisher: Nature Publishing Group.
URL: <https://www.nature.com/articles/s41550-024-02402-9>
- Tang, Y., Haiman, Z., MacFadyen, A. (2018), The late inspiral of supermassive black hole binaries with circumbinary gas discs in the LISA band, *Monthly Notices of the Royal Astronomical Society*, 476(2), 2249–2257.
URL: <https://doi.org/10.1093/mnras/sty423>
- The Event Horizon Telescope Collaboration (2019), First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole, *The Astrophysical Journal Letters*, 875(1), L1. Publisher: The American Astronomical Society.
- The Event Horizon Telescope Collaboration (2022), First Sagittarius A* Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole in the Center of the Milky Way, *The Astrophysical Journal Letters*, 930(2), L12.

Publisher: The American Astronomical Society.

Tian, C., Ding, R., Kou, X.-X. (2024), ‘Estimating the gravitational wave background anisotropy: a Bayesian approach boosted by cross-correlation angular power spectrum’. arXiv:2412.01219.

URL: <http://arxiv.org/abs/2412.01219>

Tinto, M., Alves, M. E. d. S. (2010), LISA sensitivities to gravitational waves from relativistic metric theories of gravity, *Physical Review D*, 82(12), 122003.

Publisher: American Physical Society.

URL: <https://link.aps.org/doi/10.1103/PhysRevD.82.122003>

Tinto, M., Dhurandhar, S. V. (2020), Time-delay interferometry, *Living Reviews in Relativity*, 24(1), 1.

URL: <https://doi.org/10.1007/s41114-020-00029-6>

Toonen, S., Perets, H. B., Igoshev, A. P., Michaely, E., Zenati, Y. (2018), The demographics of neutron star white dwarf mergers - Rates, delay-time distributions, and progenitors, *Astronomy & Astrophysics*, 619, A53. Publisher: EDP Sciences.

URL: <https://www.aanda.org/articles/aa/abs/2018/11/aa33164-18/aa33164-18.html>

Toubiana, A., Pompili, L., Buonanno, A., Gair, J. R., Katz, M. L. (2023), ‘Measuring source properties and quasi-normal-mode frequencies of heavy massive black-hole binaries with LISA’. arXiv:2307.15086 [gr-qc].

URL: <http://arxiv.org/abs/2307.15086>

Wagg, T., Broekgaarden, F. S., de Mink, S. E., Frankel, N., van Son, L. A. C., Justham, S. (2022), Gravitational Wave Sources in Our Galactic Backyard: Predictions for BHBH, BHNS, and NSNS Binaries Detectable with LISA, *The Astrophysical Journal*, 937(2), 118. Publisher: The American Astronomical Society.

URL: <https://dx.doi.org/10.3847/1538-4357/ac8675>

Wang, Y., Shang, Y., Babak, S. (2012), EMRI data analysis with a phenomenological waveform, *Physical Review D*, 86(10), 104050. arXiv:1207.4956 [gr-qc].

URL: <http://arxiv.org/abs/1207.4956>

Webbink, R. F. (1984), Double white dwarfs as progenitors of R Coronae Borealis stars and type I supernovae., *The Astrophysical Journal*, 277, 355–360. Publisher: IOP ADS Bibcode: 1984ApJ...277..355W.

URL: <https://ui.adsabs.harvard.edu/abs/1984ApJ...277..355W>

Weiss, R. (1979), Gravitational Wave Detectors, in L. L. Smarr, (ed.), Sources of Gravitational Radiation, Cambridge University Press, Seattle, pp. 7–35.

Whitworth, A. P., Lomax, O. (2015), Are the majority of sun-like stars single?, *Monthly Notices of the Royal Astronomical Society*, 448(2), 1761–1766.

URL: <https://doi.org/10.1093/mnras/stv093>

Will, C. M. (1998), Bounding the mass of the graviton using gravitational-wave observations of inspiralling compact binaries, *Physical Review D*, 57(4), 2061–2068. Publisher: American Physical Society.

URL: <https://link.aps.org/doi/10.1103/PhysRevD.57.2061>

Will, C. M. (2014), The Confrontation between General Relativity and Experiment, *Living Reviews in Relativity*, 17(1), 4.

URL: <https://doi.org/10.12942/lrr-2014-4>

Yuan, C., Murase, K., Zhang, B. T., Kimura, S. S., Mészáros, P. (2021), Post-merger Jets from Supermassive Black Hole Coalescences as Electromagnetic Counterparts of Gravitational Wave Emission, *The Astrophysical Journal Letters*, 911(1), L15. Publisher: The American Astronomical Society.

URL: <https://dx.doi.org/10.3847/2041-8213/abee24>

DRAFT

7

Introduction to gravitational wave data analysis

Sylvain Marsat¹ and Frédérique Marion²

¹ *Laboratoire des 2 Infinis - Toulouse (L2IT-IN2P3), Université de
Toulouse, CNRS, UPS, F-31062 Toulouse Cedex 9, France*

² *Laboratoire d'Annecy de Physique des Particules, Univ. Savoie Mont
Blanc, CNRS, F-74940 Annecy, France*

7.1. Introduction

Data analysis bridges the gap between instrumental science and the scientific exploitation of gravitational wave (GW) observations. This chapter offers a somewhat reduced scope compared to what the title might suggest. Indeed, it focuses on describing the basics of analyzing the data of ground-based gravitational-wave detectors in order to detect resolved signals from compact binary coalescences (CBCs) and infer the properties of detected sources. Searches for other types of signals are only briefly introduced, as are the downstream analyses deriving astrophysical interpretations of observed sources, and challenges expected to arise when analyzing the data of future – space-borne or ground-based – detectors. We also refer the reader to chapter ?? for an introduction on analyzing the data of pulsar timing arrays.

The contents of this chapter is meant to complement the lectures and hands-on sessions given during the MaNiTou schools. It aims mainly at providing technical background for some aspects not explained in detail in the lectures, with no claim to exhaustivity. We therefore recommend that readers concurrently refer to lecture material, which in particular includes many more illustrative figures than this chapter.

Finding meaningful references for data analysis is sometimes difficult, as tools that are used by many articles are rarely explained from first principles. We will follow part of the presentation of the book (Maggiore 2007) (see Chapter 7), to which we refer for more details. The recent review (Chatziioannou *et al.* 2024) includes a good overview of data analysis for CBCs. Another useful reading is the guide to LVK data: (Abbott *et al.* 2020), and the Gravitational Waves Open Science Center (<https://gwosc.org/> n.d.) provides both access to public LVK data and useful tutorials.

7.2. Notations and preliminaries

7.2.1. Fourier transforms: continuous and discrete domain

The Fourier transform and inverse Fourier transform will be defined as

$$\tilde{F}(f) = \int_{-\infty}^{\infty} dt e^{-2\pi i f t} F(t), \quad F(t) = \int_{-\infty}^{\infty} df e^{2\pi i f t} \tilde{F}(f). \quad [7.1]$$

A useful identity will be the Plancherel theorem: for $F, G \in \mathbb{C}$

$$\int dt F^*(t)G(t) = \int df \tilde{F}^*(f)\tilde{G}(f). \quad [7.2]$$

In practice, we will always work with real signals. For $F(t) \in \mathbb{R}$, we have the symmetry $\tilde{F}(-f) = \tilde{F}^*(f)$.

The Discrete Fourier Transform (DFT) gives equivalent definitions in the discrete domain; it can be computed efficiently through the Fast Fourier Transform (FFT) algorithm. The time interval $\Delta t = 1/f_{\text{sample}}$ must be chosen so that the Fourier transform of the time-domain data does not have support beyond $f_{\text{Nyquist}} = f_{\text{sample}}/2$. Given a vector $\mathbf{F} = \{\mathbf{F}_j\}_{j=0,\dots,N-1}$ sampled at Δt , its DFT and inverse DFT are

$$\tilde{\mathbf{F}}_k = \Delta t \sum_{j=0}^{N-1} \omega^{-jk} \mathbf{F}_j, \quad \mathbf{F}_j = \Delta f \sum_{k=0}^{N-1} \omega^{jk} \tilde{\mathbf{F}}_k, \quad \omega = e^{2i\pi/N}, \quad [7.3]$$

where the frequency resolution $\Delta f = 1/(N\Delta t) = 1/T$ is the inverse of the total duration. Note that conventions differ in the literature, and might not include the same prefactors. Care must be taken with the output of the DFT: the first half of the vector $\tilde{\mathbf{F}}$ corresponds to the positive frequencies, and the second half to the negative frequencies.

7.2.2. Gaussian distributions

We recall here the expression for the distribution of a multidimensional Gaussian distribution for a random vector \mathbf{x} :

$$p(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^N \det \boldsymbol{\Sigma}}} \exp \left[-\frac{1}{2} (\mathbf{x} - \bar{\mathbf{x}})^T \cdot \boldsymbol{\Sigma}^{-1} \cdot (\mathbf{x} - \bar{\mathbf{x}}) \right],$$

$$\boldsymbol{\Sigma} = \langle (\mathbf{x} - \bar{\mathbf{x}})(\mathbf{x} - \bar{\mathbf{x}})^T \rangle. \quad [7.4]$$

Here, $\bar{\mathbf{x}}$ is the mean and $\boldsymbol{\Sigma}$ is the symmetric and positive-definite covariance matrix.

7.3. The noise as a stochastic process

7.3.1. Stochastic processes

At the heart of our statistical modelling of our experiments will be the concept of stochastic process. Our instruments are noisy, and their output will be affected by a myriad of microscopic (or quantum) processes composing together a random process, such that at each time t , the noise $n(t) \in \mathbb{R}$ is a random variable. Since our instruments are running for extended periods of time, for simplicity we will consider formal infinite-length processes. In reality,

the data will be sampled at a finite rate for a set duration, and will be digitized into a finite set of numbers. From this point of view, the noise can be represented as a random vector $\mathbf{n} \in \mathbb{R}^N$ for a certain data size N .

We will assume here that the process is *stationary*, meaning that its statistical properties do not change over time. If the process is further assumed to be *ergodic*, then we can interchange ensemble averages and time averages.

7.3.2. The power spectral density

The power spectral density (PSD) will be a central quantity in characterizing the noise process. Intuitively, it corresponds to a Fourier spectrum of the noise, telling us whether the process is dominated by low frequencies (red noise), high frequencies (blue noise), or has a flat spectral shape (white noise). Because the noise is a stochastic process, the PSD is defined in terms of power.

Let us first define the *mean power* of the stationary noise process:

$$P_n = \lim_{T \rightarrow +\infty} \frac{1}{T} \int_{-T}^T dt n^2(t), \quad [7.5]$$

where we introduced a limitation to a finite length segment $[-T, T]$ to make sure that e.g. the Fourier transform is well defined, assuming that the limit of infinite duration is well behaved. To ease notation, we can introduce $n_T(t) = \chi_{[-T, T]}(t)n(t)$ which is $n(t)$ inside the interval $[-T, T]$ and zero outside. We can then write in the Fourier domain

$$P_n = \lim_{T \rightarrow +\infty} \frac{1}{T} \int_{-\infty}^{+\infty} dt n_T^2(t) = \lim_{T \rightarrow +\infty} \frac{1}{T} \int_{-\infty}^{+\infty} df |\tilde{n}_T(f)|^2. \quad [7.6]$$

This leads us to introducing the one-sided PSD $S_n(f)$, defined by convention for positive f , as

$$S_n(f) = \lim_{T \rightarrow +\infty} \frac{2}{T} |\tilde{n}_T(f)|^2, \quad P_n = \int_0^{+\infty} df S_n(f), \quad [7.7]$$

where to obtain the second equality we interverted the limit and the integral. Note the factor 2, accounting for the contribution of negative frequencies. This gives us a first interpretation of the PSD: the mean power can be decomposed in frequency space, and the PSD represents the associated spectrum. The PSD can also be given a different interpretation in terms of the noise autocorrelation, as we will see now.

7.3.3. The noise autocorrelation function

At the heart of our statistical description of the noise will lie the correlations between different times, for which we will use the ensemble average $\langle \rangle$.

First, we note that our detectors always make relative measurement, so that a shift by a constant of the detector's output is meaningless. Without loss of generality, we can therefore assume that the noise is zero-mean: $\langle n(t) \rangle = 0$.

The noise autocorrelation function is defined in general as a function of any two distinct times as

$$C_n(t, t') = \langle n(t)n(t') \rangle. \quad [7.8]$$

If the process is Gaussian, as we will assume here, its statistic is entirely described by this variance; a more general process would be described by additional higher-order correlation functions.

If we now specialize to the case of stationary noise, we see that the autocorrelation $C_n(t, t')$ must depend on t, t' only through their difference, or lag $\tau = t' - t$:

$$C_n(\tau) = \langle n(t)n(t + \tau) \rangle. \quad [7.9]$$

For the stationary ergodic case that we are considering, we can exchange time and ensemble averages and write

$$\begin{aligned} C_n(\tau) &= \lim_{T \rightarrow +\infty} \frac{1}{T} \int_{-\infty}^{+\infty} dt n_T(t)n_T(t + \tau) \\ &= \lim_{T \rightarrow +\infty} \frac{1}{T} \int_{-\infty}^{+\infty} df' \tilde{n}_T^*(f') \tilde{n}_T(f') e^{2i\pi f' \tau} \end{aligned} \quad [7.10]$$

Computing the Fourier transform of $C_n(\tau)$ itself, we have

$$\int_{-\infty}^{+\infty} d\tau C_n(\tau) e^{-2i\pi f \tau} = \lim_{T \rightarrow +\infty} \frac{1}{T} \int_{-\infty}^{+\infty} df' \tilde{n}_T^*(f') \tilde{n}_T(f') \int_{-\infty}^{+\infty} d\tau e^{2i\pi(f' - f)\tau}. \quad [7.11]$$

The second integral collapses to a Dirac delta as $\delta(f' - f)$, and we can recognize the PSD $S_n(f)$ as introduced in [7.7]. We arrive at

$$\frac{1}{2} S_n(f) = \int d\tau e^{-2i\pi f \tau} C_n(\tau), \quad [7.12]$$

which gives us a useful and perhaps intuitive alternative definition of the noise PSD, as the Fourier transform of the autocorrelation function of a stationary process. Note the factor 2, which is a matter of convention and kept here so that S_n is the one-sided PSD defined on positive frequencies.

The PSD is real and positive, as shown by its definition in terms of power [7.7]: $S_n(f) \geq 0$. When seen as a Fourier transform as in [7.12], we have that as a FT of a real quantity $S_n(-f) = S_n(f)$; in practice one often considers the PSD as defined on the range $f > 0$.




Figure 7.1: Example PSDs for the instruments LIGO and Virgo during the runs O2 and O3. Figure taken from (Davis *et al.* 2021).

We show an example of PSDs from the LIGO and Virgo instruments during their observing runs O2 and O3 in Fig. 7.1. We see that they have a “bucket shape”, selecting a certain range of frequencies (roughly 20-2000Hz) as relevant. Far from being smooth functions, they feature numerous instrumental lines.

7.3.4. Stationary and Gaussian noise

The definition of the noise autocorrelation function [7.9] can be translated into Fourier-domain correlations. For stationary noise, assuming all Fourier transforms to be well defined (and removing bounds on the integrals), we can write:

$$\begin{aligned}
 \langle \tilde{n}(f) \tilde{n}^*(f') \rangle &= \int dt \int dt' e^{-2i\pi ft} e^{2i\pi f' t'} \langle n(t) n(t') \rangle \\
 &= \int dt \int d\tau e^{-2i\pi ft} e^{2i\pi f'(t+\tau)} \langle n(t) n(t+\tau) \rangle \\
 &= \int d\tau C_n(\tau) e^{2i\pi f' \tau} \int dt e^{2i\pi(f-f')t}. \tag{7.13}
 \end{aligned}$$

Recognizing the last integral as a Dirac delta, we obtain

$$\langle \tilde{n}(f) \tilde{n}^*(f') \rangle = \frac{1}{2} S_n(f) \delta(f - f'). \tag{7.14}$$

A similar computation would show that

$$\langle \tilde{n}(f) \tilde{n}(f') \rangle = \frac{1}{2} S_n(f) \delta(f + f'). \tag{7.15}$$

Equations [7.14] and [7.15] show a crucial point: *noise stationarity implies independence in the Fourier domain*. It is at the heart of a major simplification in our analysis, allowing us to treat Fourier bins as independent, meaning that $\tilde{n}(f)$ and $\tilde{n}(f')$ will be independent random variables if $f \neq f'$. although this statement needs to be made more precise: for finite-length data it relies on an additional assumption of *circularity* as we will discuss in the next section.

Gaussian noise is entirely described by its covariance, so [7.14] and [7.15] are enough to write down $p(n)$, the probability distribution of the noise. However, getting the normalization right will be easier by working with a discrete representation of the noise, as we will do in the next section. We will arrive at $p(n)$ in Eq. [7.19] below.

7.3.5. Noise in discrete data

Although writing the formalism with functions and distributions has the advantage of simplicity, in practice a real detector will always work with *sampled* processes, and we will always have discretized data at hand.




Figure 7.2: Example data from the LIGO observations around the first detection event GW150914 ((Abbott *et al.* 2016b), data from (<https://gwosc.org/n.d.>)). The first panel shows the raw strain for a segment of 32s around the event; the signal is completely invisible. The second panel shows whitened and band-passed data, as does the third panel, zooming on the event.

In the discrete point of view, the data is of finite length N , sampled with a fixed time interval δt , and represented by a vector $\mathbf{n} \in \mathbb{R}^N$. The random process is then replaced by a random vector, with its variance described by a covariance matrix $\Sigma \in \mathbb{R}^{N \times N}$ as

$$\langle \mathbf{n} \mathbf{n}^T \rangle = \Sigma. \quad [7.16]$$

The covariance Σ is, by construction, symmetric and positive-definite. In general, the time-domain covariance is dense, and is a very large matrix for large N .

The discrete representation equivalent of the stationarity property [7.9] is the requirement that Σ has a *Toeplitz* structure: there must exist a vector σ such that $\Sigma_{ij} = \langle n_j n_j \rangle = \sigma_{|i-j|}$. In practice, that means that the covariance matrix has a band structure, constant along each diagonal.

In order to get the equivalent of the Fourier-domain independence, eqs [7.14] and [7.15], the covariance needs to have the additional property of being *circulant*: $\sigma_j = \sigma_{N-j}$ for $j = 1, \dots, N-1$. This would mean that the process itself is periodic, which is never the case in practice. Without providing more detail here, we will accept that we can assume this circulant condition as an approximation, provided that we are working with data that is longer than the signal of interest, and that the data has been tapered at both ends.

In analogy with [7.14], the result for the Fourier-domain covariance matrix in the stationary (Toeplitz) circulant case is:

$$\tilde{\Sigma} = \langle \tilde{\mathbf{n}} \tilde{\mathbf{n}}^\dagger \rangle = \frac{1}{\Delta f} \text{Diag}(\tilde{\sigma}) = \text{Diag} \left(\frac{S_n}{2\Delta f} \right), \quad [7.17]$$

where $\tilde{\mathbf{n}}^\dagger$ is the conjugate transpose. Note that the covariance matrix $\tilde{\Sigma}$ has redundant information, its first quadrant corresponds to positive frequencies $f, f' > 0$. The diagonal is formed by values of the PSD $S_n(j\Delta f)$ for the first half, complemented by the same vector in reverse. Care must also be taken for the indices $j = 0$ and $j = N/2$.

Similarly to the result [7.15], one also finds that $\langle \tilde{\mathbf{n}} \tilde{\mathbf{n}}^T \rangle = 0$ ($\tilde{\mathbf{n}}^T$ is the normal transpose, not conjugate) on positive frequencies $f, f' > 0$. Separating the real and imaginary parts of each component $\tilde{\mathbf{n}}_j$, this allows to conclude that $\text{Re } \tilde{\mathbf{n}}_j, \text{Im } \tilde{\mathbf{n}}_j$ are all independent and distributed as

$$\text{Re } \tilde{\mathbf{n}}_j, \text{Im } \tilde{\mathbf{n}}_j \sim \mathcal{N} \left(0, \frac{S_n(j\Delta f)}{4\Delta f} \right) \quad \text{for } j = 1, \dots, \frac{N}{2} - 1. \quad [7.18]$$

In addition, for the special indices $j = 0$ and $j = N/2$, $\tilde{\mathbf{n}}_j$ is real by definition, and the real part is distributed as $\sim \mathcal{N}(0, S_n/(2\Delta f))$ with a factor 1/2 instead of 1/4. This specifies entirely the probability distribution of a Gaussian and stationary (and circulant) noise vector. From a given PSD, we could simulate Fourier-domain noise according to [7.18], and compute its inverse Fourier transform to obtain time-domain noise data \mathbf{n} . Comparing with [7.14], we see that

the Dirac delta $\delta(f - f')$ that we had when working with continuous functions has become a $1/\Delta f$ in the discrete representation.

The posterior distribution for the noise vector $\tilde{\mathbf{n}}$ can be obtained by multiplication of the probability distributions for the individual bins, and reads

$$\ln p(\tilde{\mathbf{n}}) = -2\Delta f \sum_{j=0}^{N/2} a_N^j \frac{|\tilde{\mathbf{n}}_j|^2}{S_n(j\Delta f)} - \sum_{j=0}^{N/2} a_N^j \frac{S_n(j\Delta f)}{\Delta t} - \frac{N}{2} \ln 2\pi. \quad [7.19]$$

Here, we introduced a new notation $a_N^j = 1$, except for $j = 0$ and $j = N/2$ where $a_N^j = 1/2$. The second term is a normalization that depends only on the PSD, and the third term is a pure constant.

We show in Fig. 7.3 an example of data from the LIGO observations around the first detection event GW150914 (Abbott *et al.* 2016b) (data from (<https://gwosc.org/> n.d.)). The signal is completely invisible in the times series of raw strain. We also show *band-passed* and *whitened* data, where the signal emerges above the noise. Band-passing is done by multiplying the Fourier-domain data $\tilde{d}(f) \rightarrow \tilde{d}(f) \times w(f)$ by a window function going smoothly to 0 outside of the frequency range of relevance for the signal (eliminating large noise power at low and high frequencies). Whitening amounts to dividing the data by the square root of the PSD, $\tilde{d}(f) \rightarrow \tilde{d}(f)/\sqrt{S_n(f)}$. Note that the fact that the signal can be detected by eye is atypical; GW150914 was especially loud, with an SNR of 24, and was relatively high-mass, making it a short transient. Longer-lived, quieter signals cannot be identified in the same way.

7.3.6. Noise-weighted inner product

A useful tool, that we will encounter multiple times, will be the following noise-weighted inner product: for real signals $a(t), b(t) \in \mathbb{R}$,

$$(a|b) = 4\text{Re} \int_0^{+\infty} df \frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n(f)}. \quad [7.20]$$

This inner product has all the properties of a scalar product over real functions a, b . It is conventionally written over positive frequencies, however we recall that $\tilde{a}(-f) = \tilde{a}^*(f)$ for $a \in \mathbb{R}$. A factor 2 is included as a matter of convention. The PSD of the underlying noise plays a crucial role in this *noise-weighted* inner product, penalizing frequency ranges where the instrumental noise is high and

highlighting the “bucket” of the instrumental sensitivity. In practice, one can often limit the integral to a finite range of relevant frequencies $[f_{\min}, f_{\max}]$.

When working with discrete data, the equivalent expression when restricting the range to $[f_{\min}, f_{\max}] = \Delta f \times [j_{\min}, j_{\max}]$ (thus ignoring the special cases $j = 0$ and $j = N/2$), is

$$\langle \mathbf{a} | \mathbf{b} \rangle = 4\Delta f \operatorname{Re} \sum_{j=j_{\min}}^{j_{\max}} \frac{\tilde{a}_j^* \tilde{b}_j}{S_n(j\Delta f)}. \quad [7.21]$$

Conversely, we can also translate the noise probability distribution [7.19] using this noise-weighted inner product. Removing constants that do not depend on n , we have the simple rewriting

$$\ln p(n) = -\frac{1}{2} \langle n | n \rangle + \text{const}. \quad [7.22]$$

This expression is also called the *Whittle likelihood*, and applies only to stationary Gaussian noise, provided that we also accept the circulant approximation for finite-length data.

An important property of this inner product that will simplify future derivations is the following:

$$\langle |n\rangle \langle n| \rangle = \mathbf{1}. \quad [7.23]$$

Namely, for any two signals $a, b \in \mathbb{R}$, writing the integrals over positive and negative frequencies we have

$$\begin{aligned} \langle (a|n)(n|b) \rangle &= 4 \int_{-\infty}^{+\infty} df \int_{-\infty}^{+\infty} df' \frac{1}{S_n(f)S_n(f')} \tilde{a}^*(f) \langle \tilde{n}(f) \tilde{n}^*(f') \rangle \tilde{b}(f') \\ &= 2 \int_{-\infty}^{+\infty} df \frac{\tilde{a}^*(f) \tilde{b}(f)}{S_n(f)} = \langle a | b \rangle, \end{aligned} \quad [7.24]$$

where we made use of [7.14].

7.3.7. PSD estimation

The noise PSD $S_n(f)$ will play a crucial role in all data analysis applications. Although we can build models for the noise based on the physical processes at

play, real instruments are of such complexity and undergo such variations over time that we will always need to estimate the PSD from the data. When working with the assumption of stationarity, we are in fact assuming that the process is stationary on a limited data stretch around the event of interest, hopefully longer than the signal itself.

A standard method to estimate the PSD is the *Welch method*. The idea is to divide the data of length N into K segments of length M each, with some overlap between segments. A smooth window tapering to zero at each end is applied to each segment, before taking an FFT. The resulting spectra of $|\tilde{n}(f)|^2$ for the K segments are then averaged to estimate $S_n(f)$ according to [7.17]. If segments are of lengths $M\Delta t$, the frequency resolution is $\Delta f = 1/(M\Delta t)$. This leads to a trade-off: the longer the individual segments, the better the frequency resolution Δf , but the smaller K becomes, leading to more statistical uncertainty when doing the averaging. Conversely, we can increase K for a better statistical averaging of the spectra, but this will degrade Δf since segments become shorter.

Note that this method assumes that the signal is subdominant, allowing us to ignore it in the PSD estimation. In future instruments, the data will not always be noise-dominated. In particular, in LISA, long-lived signals such as galactic binaries will always be present, meaning that we will not have noise-only data segments; in addition signals will dominate instrumental noise in some frequency bands and create confusion noise. In such a case, estimating the instrumental noise PSD becomes part of a global fit framework, see Chapter ??.

For a description of noise characterization in the recent O2 and O3 LVK runs, see (Davis *et al.* 2021). Real-life analysis with non-stationarity and glitches is significantly more involved than what we described here, and an important role is played by Bayesian methods using wavelets, such as BayesWave (Cornish and Littenberg 2015).

7.4. Matched filtering

7.4.1. The matched filter: signal and noise

The first task of data analysts is to detect the tiny gravitational wave signal in noise-dominated data streams. The matched filter approach to detection relies on leveraging our knowledge of what the signal might look like. Intuitively, if we cross-correlate such a signal template with data and if the template matches well the physical signal, it will accumulate positive correlation as time goes by, while the correlation with random instrumental noise will only fluctuate.

If our data stream is $s(t) = h(t) + n(t)$ with $h(t)$ the gravitational wave signal and $n(t)$ the noise which has a zero mean in the sense of the ensemble average $\langle n(t) \rangle = 0$, the idea is to introduce a filter $W(t)$ and to build its correlation with the data stream with and without signal as

$$\begin{aligned}\hat{s} &= \int dt W(t) s(t) \\ \hat{n} &= \int dt W(t) n(t),\end{aligned}\quad [7.25]$$

where we omit bounds on integrals, assuming we work with a finite data segment. The signal and noise estimators \hat{s} , \hat{n} are random variables, as they contain random noise, and by construction $\langle \hat{n} \rangle = 0$. The next step is to define the expectation values:

$$S = \langle \hat{s} \rangle, \quad N^2 = \langle \hat{n}^2 \rangle, \quad N = \sqrt{N^2}, \quad [7.26]$$

and the matched filtering approach consists in finding the optimal filter $W(t)$ to maximize the signal-to-noise ratio (SNR) S/N , assuming we know the physical signal $h(t)$.

7.4.2. The Wiener filter

Since $s(t) = h(t) + n(t)$ with $h(t)$ deterministic and $n(t)$ random and zero-mean, we can first write

$$\langle \hat{s} \rangle = \int dt W(t) h(t) = \int df \tilde{W}^*(f) \tilde{h}(f), \quad [7.27]$$

where we converted the integral from time-domain to Fourier-domain thanks to [7.2]; the Fourier-domain integrals are on the full frequency range, unless specified otherwise. For the noise, we have

$$\langle \hat{n}^2 \rangle = \int df \int df' \tilde{W}^*(f) \tilde{W}(f') \langle \tilde{n}(f) \tilde{n}^*(f') \rangle = \int df |\tilde{W}(f)|^2 \frac{S_n(f)}{2}, \quad [7.28]$$

where we used [7.14] for stationary noise.

A second helpful simplification is brought by introducing

$$\tilde{w}(f) = \frac{1}{2} S_n(f) \tilde{W}(f), \quad [7.29]$$

so that the previous results become

$$\langle \hat{s} \rangle = (w|h), \quad \langle \hat{n}^2 \rangle = (w|w), \quad \frac{S}{N} = \frac{(w|h)}{\sqrt{(w|w)}}. \quad [7.30]$$

Thanks to these redefinitions, the solution of the optimization problem is clear. The optimal value for S/N is obtained for w and h colinear, $w \propto h$. Our solution for the *matched filter* (up to a multiplicative constant that is a matter of convention) is therefore

$$\tilde{W}(f) = \frac{2\tilde{h}(f)}{S_n(f)}. \quad [7.31]$$

7.4.3. The matched filter statistics




Figure 7.3: Matched filter statistics for the event GW150914. The first panel shows the whitened band-passed data in blue and the best-matching template in black. The lower panel shows the time series of the matched filter SNR $|\hat{\rho}|$ as we slide the template in time.

With the result [7.31] at hand for the matched filter $\tilde{W}(f)$, we can come back to the expressions of \hat{s} and $\langle \hat{n}^2 \rangle$ in [7.25] and [7.28], and convert them to the Fourier domain to obtain

$$\hat{s} = (h|s), \quad N^2 = (h|h). \quad [7.32]$$

Applying our optimal matched filter on the observed data stream, which includes signal plus noise, we can build a statistic called the *matched filter SNR*:

$$\hat{\rho} = \frac{\hat{s}}{N} = \frac{(h|s)}{\sqrt{(h|h)}}. \quad [7.33]$$

Intuitively, this quantity is built by cross-correlating our template h with the data stream s , normalizing by the norm of the template itself, all in accordance with the noise-weighted inner product [7.20]. If the data has no noise at all, with a physical signal matching exactly the template we are using, then the value of the matched filter SNR would be the *optimal SNR*

$$\rho_{\text{opt}} = \sqrt{(h|h)}. \quad [7.34]$$

The matched filter SNR takes values that fluctuate due to the randomness of the noise. A crucial question is that of determining its probability distribution. Because we are assuming that the noise is a Gaussian process and because we can see $\hat{\rho}$ as a linear combination of Gaussian variables, it has to be Gaussian distributed itself, and we only have to determine its expectation value and variance. If we had a single data stream as in our previous notations, containing only noise and no signal, then $\langle \hat{\rho} \rangle = 0$ and we have

$$\text{In noise: } \hat{\rho} = \frac{\hat{n}}{N} \sim \mathcal{N}(0, 1). \quad [7.35]$$

In a data stream containing a signal matching exactly our template, $\langle \hat{\rho} \rangle = \rho_{\text{opt}}$ as defined in [7.34],

$$\text{In presence of signal: } \hat{\rho} = \rho_{\text{opt}} + \frac{\hat{n}}{N} \sim \mathcal{N}(\rho_{\text{opt}}, 1). \quad [7.36]$$

We could also work with the square of the matched filter SNR as a detection statistic; in that case, we would have that $\hat{\rho}^2$ is a χ^2 variable with one degree of freedom in the presence of noise only. In fact, when optimizing over the phase degree of freedom, two degrees of freedom arise, as we will see in the next section.

We show in Fig. 7.3 an example application of matched filtering, applied to data around the event GW150914 ((Abbott *et al.* 2016b)).

An important point is that of the expected *SNR distribution* of the sources. Eq. 7.33 shows that $\rho_{\text{opt}} \propto 1/d_L$, with d_L the luminosity distance (see Chapter ??). As long as we do observations in the local universe, we can identify the

luminosity distance with the comoving distance at low redshift, and a uniform comoving density of sources translates as $p(d_L) \propto d_L^2$. If the properties of the source population do not vary much within our reach in redshift, regardless of their distribution in mass, spin etc. we can change variable from d_L to the SNR variable ρ as follows:

$$p(\rho)|d\rho| = p(d_L)|dd_L|, \quad p(\rho) = p(d_L) \left| \frac{dd_L}{d\rho} \right| \propto \rho^{-4}, \quad [7.37]$$

which indicates a steep fall-off in SNR: if we detect sources above a certain threshold $\rho > \rho_{\text{thres}}$, most of the detections will be expected just above that threshold.

7.4.4. Matched filter SNR optimization

Up until this point, we always considered a single template $h(t)$, and assumed that it matches exactly the physical signal. However, even if we imagined knowing the intrinsic parameters of the gravitational wave source such as masses and spins, we would have to seep through the data by changing the merger time of the template, and possibly maximize our search statistic for phase degrees of freedom as well, which will change the effective number of degrees of freedom. In reality, none of the physical parameters of the source are known, and we have to rely on a bank of templates to cover the parameter space, and the matched filtering procedure is applied on each of these templates; this will be the subject of the next section.

First, a matched filter SNR can be maximized with respect to a shift in time via an inverse Fourier transform as follows. If we introduce a family of templates $h_{t_0}(t) = h(t - t_0)$ that are shifted in time by t_0 , their Fourier transform is

$$\tilde{h}_{t_0}(f) = e^{-2i\pi f t_0} \tilde{h}(f). \quad [7.38]$$

The normalization is unchanged by a time shift, $(h_{t_0}|h_{t_0}) = (h|h)$, and we have (coming back to a two-sided Fourier-domain integral)

$$(h_{t_0}|s) = 2 \int_{-\infty}^{+\infty} \frac{df}{S_n(f)} \tilde{s}^*(f) \tilde{h}(f) e^{-2i\pi f t_0}, \quad [7.39]$$

so that we can introduce an intermediate Fourier-domain integrand \tilde{g} and write

$$\tilde{g}(f) = 2\tilde{s}^*(f)\tilde{h}(f)/S_n(f), \quad (h_{t_0}|s) = g(t_0) \quad [7.40]$$

where g is the inverse Fourier transform of \tilde{g} . In practice, we are working with discrete data vectors, the inverse Fourier transform is discrete, and we obtain the value of $(h_{t_0}|s)$ on a grid of times $t_0 = k\Delta t$ with $k \in \mathbb{Z}$. We can then look for the optimum over k .

Another maximization can be done with respect to a phase degree of freedom, that is to say when there exist templates h_c , h_s and a phase φ so that

$$h(t) = \cos \varphi h_c(t) + \sin \varphi h_s(t). \quad [7.41]$$

The derivation is simpler if we can reparametrize the waveform so that the two template components are orthogonal, $(h_c|h_s) = 0$. In that case we can write

$$\hat{\rho}^2 = \frac{(h|s)^2}{(h|h)} = \frac{((h_c|s) + \tan \varphi (h_s|s))^2}{(h_c|h_c) + \tan^2 \varphi (h_s|h_s)}, \quad [7.42]$$

which can be maximized with respect to $\tan \varphi$ to obtain (with a star for quantities at the optimum)

$$\tan \varphi_* = \frac{(h_s|s)(h_c|h_c)}{(h_c|s)(h_s|h_s)}, \quad \hat{\rho}_*^2 = \frac{(h_c|s)^2}{(h_c|h_c)} + \frac{(h_s|s)^2}{(h_s|h_s)} = (\bar{h}_c|s)^2 + (\bar{h}_s|s)^2, [7.43]$$

where we introduced normalized templates $\bar{h}_c = h_c/\sqrt{(h_c|h_c)}$, $\bar{h}_s = h_s/\sqrt{(h_s|h_s)}$.

Comparing to our previous result [7.36], we see that the matched filter SNR statistic, when optimized over phase, is now the quadratic sum of two Gaussian variables,

$$(\bar{h}_c|s) \sim \mathcal{N}((\bar{h}_c|h), 1), \quad (\bar{h}_s|s) \sim \mathcal{N}((\bar{h}_s|h), 1), \quad [7.44]$$

with expectation values that become zero in the absence of signal.

We also have that those two Gaussian variables are independent, thanks to the orthogonality of the two templates and to the property [7.23]:

$$\langle (\bar{h}_c|n)(n|\bar{h}_s) \rangle = (\bar{h}_c|\bar{h}_s) = 0. \quad [7.45]$$

Thus, the phase-optimized matched filter statistic follows a Rayleigh distribution with two degrees of freedom, off-centered in presence of signal.

Armed with probability distributions for the matched filter statistic in presence of signal and of pure noise, one can adjust a detection threshold by balancing false alarms ($p(\hat{\rho}^2) > \rho_{\text{thres}}^2$ for noise) and false dismissals ($p(\hat{\rho}^2) < \rho_{\text{thres}}^2$ for the signal). In Gaussian and stationary noise and for an exactly known template, applying a threshold on the matched filtering statistic is optimal, i.e. the false alarm rate is guaranteed to be minimal for a given false dismissal rate. Such estimates are given in (Maggiore 2007), taking into account the fact that we need multiple templates to cover the parameter space. However, these estimates are too naive for realistic searches with non-Gaussian and non-stationary data, as we will see.

7.5. Template bank searches

7.5.1. Template banks

The matched filtering approach assumes that the intrinsic parameters of the source, which drive the dynamics of the system and determine the waveform, are known a priori, which is not the case. The parameter space of interest therefore needs to be scanned, with a finite set of representative waveforms called *templates*. The effectiveness of a discrete sampling of the parameters relies on the overlap between a given signal and 'nearby' templates leading to values of the matched filtering statistic 'close' to the optimal value (obtained when the parameters of the template match those of the signal). In practice, these notions of 'nearby' and 'close' need to be quantified in order to build appropriate template banks.

The match $M(a, b)$ between a signal a and a template b is defined as their normalized overlap maximized over all possible times and phases at coalescence and quantifies how well template b correlates with signal a . The fitting factor φ of a bank of templates $b(\theta_i)$, characterized by source parameters θ_i , for signal a is defined as the maximum match and quantifies how well the bank can recover this particular signal, i.e. the fraction of the optimal SNR captured by the matched filtering.

$$M(a, b) = \max_{t_c, \phi_c} \frac{(a|b)}{\sqrt{(a|a)}\sqrt{(b|b)}} \leq 1, \quad \varphi = \max_i M(a, b(\theta_i)) \quad [7.46]$$

Building a template bank over a target parameter space requires dedicated algorithms. The latter come in three broad classes: geometric, stochastic and hybrid (i.e., combining the geometric and stochastic approaches). The geometric approach relies on the possibility to compute a *metric* on the parameter space, i.e. to chart how the match between two neighboring points depends on

differences in parameters. Knowing the metric allows to pave the space in an optimal way (e.g., (Owen and Sathyaprakash 1999)). This approach is very efficient, but the metric may be difficult to compute – depending on the waveform model – and the complexity increases with dimensionality (from considering the two component masses to including spin parameters). The stochastic approach instead involves placing templates at random points in the parameter space in an iterative way, keeping new draws if their fitting factor with the previous points is smaller than the required minimum. This approach is simple and versatile, but tends to be slow and does not guarantee a complete coverage, which is why in practice hybrid algorithms are often used.

The density of a template bank is parameter-dependent and is related to the duration of the signal. Indeed, a signal with many cycles in the detector bandwidth requires a template with very similar parameters to maintain a good overlap and a good degree of correlation throughout the matched filtering integral, whereas a signal with fewer cycles will be more tolerant in parameter differences. This translates for instance into template banks that are denser at small masses and sparser at high masses. Similarly, positive spins (aligned with the orbital angular momentum) slow down the dynamics of the inspiral and make for longer signals, requiring a high template density, whereas negative (anti-aligned) spins shorten the signal and can accommodate a lower density. It is also worth noting that improving the detector sensitivity at low frequencies increases the number of signal cycles in band and requires an increase in the density of template banks.

7.5.2. Towards searches in realistic noise

In practice, the noise of actual detectors is neither strictly Gaussian nor stationary. The PSD shows variations on timescales of hours or days and needs to be estimated on shorter timescales in order to capture properly those variations and not bias the matched filter process. Moreover, the matched filter statistic is not perfectly Rayleigh distributed. Although the bulk of the distribution is usually consistent with a Rayleigh distribution, it typically includes a tail at high values of the statistic that does not match the expected decay of a Rayleigh distribution. This tail is due to non-Gaussian features in the detector noise, and varies according to the quality of the data, which is not constant over time. The non-Gaussian features are related to noise transients of instrumental or environmental origin, often referred to as *glitches*. The consequence is two-fold: On the one hand, methods need to be developed to suppress these features in the detection-statistic distribution and bring the search sensitivity as close to the ideal case as possible. On the other hand, the distribution of the search background needs to be estimated from the data, as it is not possible to




Figure 7.4: Illustration (from (Abbott *et al.* 2016a)) of how signal-consistency tests, such as the χ_r^2 (see equation 7.47) used by the PyCBC pipeline (top row) and the ξ^2 (based on the template autocorrelation function) used by the GstLAL pipeline (bottom row) help separate astrophysical signals from background. Points colored by their false alarm rate are simulated signals. Black dots are background events, that can have high SNR but also feature a high χ_r^2 or ξ^2 ; we see that one can draw contours in this 2D space to separate the two.

rely on the theoretical distribution derived under the assumptions of Gaussian and stationary noise.

There are two main strategies to suppress non-Gaussian features. The first is to assess the quality of the data a priori and flag times affected by elevated noise levels or by the presence of glitches witnessed by auxiliary channels indicating their instrumental or environmental origin. The flagged data are then either not analyzed (typically, when the quality is bad for an extended period of time) or analyzed after the problematic stretch of data has been excised (typically, when the data are affected by a strong but short glitch), or analyzed as is but vetoing triggers occurring at flagged times further down the line. The second strategy involves developing tests to assess the consistency of the data with the presence of a signal when the matched filter statistic takes a value high enough

to suggest a possible detection (such occurrences, when the statistic exceeds some threshold – typically around 4 – are usually called *triggers*).

The correlation involved in matched filtering provides an integrated quantity and the resulting statistic can take high values even in cases where the data match only part of the template. Signal-consistency tests typically seek to evaluate whether the data are consistent with the expected waveform in more details. This can be done in a variety of ways, including looking at the data in the frequency domain or in the time domain. In the frequency domain, a common test is to quantify how the SNR is distributed over frequency and compare to expectations for the template, computing a χ^2 statistic:

$$\chi_r^2 = \frac{p}{2p-2} \sum_{i=1}^p \left(\hat{\rho}_i - \frac{\hat{\rho}}{p} \right)^2 \quad [7.47]$$

where the frequency band has been divided into p intervals expected to contribute equal SNR, $\hat{\rho}_i$ is the matched filter statistic computed over frequency interval i (Allen 2005). In the time domain, a typical approach consists in comparing the time series of the matched filter statistic around the time of the trigger with the time series of the template autocorrelation function. Figure 7.4 provides an illustration of how consistency tests can help separating astrophysical signals from noise.

These quantities, or others resulting from different signal-consistency tests (e.g., evaluating if the data are consistent with a signal in multiple detectors), can be used in various ways: they can be compared to a threshold to discard triggers with elevated values indicating that they fail the test and are likely not of astrophysical origin; they can be used to weight the matched filter statistic and construct a modified, better-behaved statistic with a distribution less affected by a non-Gaussian tail; they can be used to construct a multi-parameter likelihood and a detection statistic based on a likelihood ratio comparing the odds of the measured quantities for signal and noise.

As already mentioned, non-ideal detector noise means that the ranking statistic distribution for background events must be estimated from the data, in order to reliably assess the significance of detection candidates. This is most easily done when the data of multiple detectors are available and detection candidates come in the form of *coincidences* (triggers in different detectors with matching time and template parameters). Evaluating the false-alarm rate associated with a trigger requires measuring the rate of fake coincidences expected to occur from noise, with a ranking statistic as high as that of the trigger. There are several methods for this measurement. A common approach involves

performing *time slides*, i.e. shifting the data of one detector with respect to another – by an amount of time suppressing coincidences for astrophysical signals –, which provides a handle on random noise coincidences. Statistical accuracy follows from repeating the process with many different values of the time shift. Estimating the background distribution for a single detector can also be done but is more difficult and involves being more aggressive at vetoing likely noise triggers, performing a degree of extrapolation and being more conservative.

An important aspect to keep in mind, when assessing the significance of a detection candidate, is to take into account the fact that the search is not done using the one template providing the candidate but is conducted on many other templates as well. This *look elsewhere* effect requires including a *trials factor* in the false alarm rate associated to the candidate. Care must be taken in the approach used to this end, since not all templates are equivalent, background-wise, as some (typically, short ones) are more susceptible to pick up glitches. Some strategies are needed to avoid that the elevated background in some parts of the parameter space limits the sensitivity of the search in parts of the parameter space with better-behaved background.

Many more details, and references, about realistic searches can be found in the review article (Chatziioannou *et al.* 2024). The LVK uses several pipelines to search for signals from compact binary coalescences (e.g., (Ewing *et al.* 2024), (Dal Canton *et al.* 2021), (All  n   *et al.* 2025)).

7.6. Bayesian inference

7.6.1. Bayesian statistics

Statistical inference in gravitational wave astronomy is frequently formulated from a Bayesian point of view. In fact, while searches often use a frequentist setting with quantities such as false alarm probabilities, leveraging techniques such as time slides to mimic repeated experiments, tasks such as inferring the physical parameters of the sources are set in the same context as most of astronomy and cosmology: we observe non-reproducible data in one-time detections, and we need to infer the physical parameters of the source; the prior information that we assume we have on these physical parameters will play an explicit role. This is the framework of *Bayesian inference*. Starting from the joint probability of any two events,

$$p(A \cap B) = p(B|A)p(A) = p(A|B)p(B), \quad [7.48]$$

we can write Bayes' theorem as follows, using notation that we will explicit below:

$$p(\boldsymbol{\theta}|\mathbf{d}, \mathcal{M}) = \frac{p(\mathbf{d}|\boldsymbol{\theta}, \mathcal{M})p(\boldsymbol{\theta}|\mathcal{M})}{p(\mathbf{d}|\mathcal{M})}. \quad [7.49]$$

Here, the vector \mathbf{d} represents the observed data; it is of dimension N , the data size. The vector $\boldsymbol{\theta}$ gathers the signal parameters that we want to infer. If those are the parameters of a compact binary coalescence, those would be masses, spins, sky position, \dots , and θ is of dimension d , the dimensionality of the parameter space. We have added in our notation a generic context \mathcal{M} , gathering all our models and assumptions about the signals and the instrument.

In the left-hand side, $p(\boldsymbol{\theta}|\mathbf{d}, \mathcal{M})$ is the *posterior*, that we want to access: the probability distribution of signal parameters $\boldsymbol{\theta}$ given the observed data \mathbf{d} (and the model \mathcal{M}). In the right-hand side, the term $p(\mathbf{d}|\boldsymbol{\theta}, \mathcal{M})$ is the *likelihood*: the probability distribution of data \mathbf{d} given the parameters $\boldsymbol{\theta}$. Evaluating the likelihood requires a statistical model of the detector noise, as we will see below. It is also where a model for the gravitational wave signals enters. The term $p(\boldsymbol{\theta}|\mathcal{M})$ is the *prior*, representing our pre-existing information on the signal parameters. Some priors are easy to justify physically, for instance assuming sources to be evenly distributed in volume for an homogeneous and isotropic universe, while others are more difficult to motivate, an example being the priors for masses and spins. As a rule, when quoting results for a posterior, one should always mention which priors were used in the analysis.

Finally, the denominator $p(\mathbf{d}|\mathcal{M})$ is the *evidence*, which can be seen as an overall normalization of the posterior, as

$$p(\mathbf{d}|\mathcal{M}) = \int d\boldsymbol{\theta} p(\mathbf{d}|\boldsymbol{\theta}, \mathcal{M})p(\boldsymbol{\theta}|\mathcal{M}). \quad [7.50]$$

When producing samples for the posterior, the evidence plays the role of a normalization constant, independent on $\boldsymbol{\theta}$, and can be ignored. Furthermore, as a high-dimensional integral, the evidence is difficult to compute in general. However, it plays an important role in model comparison (also named hypothesis testing).

Assume that we wish to compare two models \mathcal{M}_1 and \mathcal{M}_2 , given the observed data \mathbf{d} , to determine which model fits better the data. For instance, the two models might include different physics in their waveform templates,

or different assumptions about the instrumental noise. Using Bayes theorem again, we can write an odds ratio as

$$\frac{p(\mathcal{M}_1|\mathbf{d})}{p(\mathcal{M}_2|\mathbf{d})} = B_{12} \frac{p(\mathcal{M}_1)}{p(\mathcal{M}_2)}, \quad \text{with} \quad B_{12} = \frac{p(\mathbf{d}|\mathcal{M}_1)}{p(\mathbf{d}|\mathcal{M}_2)}, \quad [7.51]$$

where we introduced the notation B_{12} for the ratio of evidences, which is called the *Bayes factor*. The second term incorporates the priors that we have for each of the models.

The Bayes factor is often given in logarithmic scale. A Bayes factor $B_{12} \gg 1$ indicates that the data “prefers” model 1 over model 2. The question of how high a Bayes factor should be in order to be considered significant is somewhat open.

An important point is that the Bayes factor incorporates an “Occam’s razor” argument: if model 2 is more complex (has more parameters) than model 1 but does not achieve a significantly better fit to the observed data, the Bayes factor will favor model 1. This behaviour is to be relativized to some extent: as shown by [7.50], if extra parameters are introduced that do not affect the likelihood in any way, they will not affect the evidence; it is only when the extra parameters are constrained by fitting the data that the penalty occurs.

7.6.2. Sampling from a posterior

Although the posterior $p(\boldsymbol{\theta}|\mathbf{d}, \mathcal{M})$ can be seen as a function of the parameters $\boldsymbol{\theta}$, it is rarely advisable to directly evaluate it, except when dealing with one- or perhaps two-dimensional parameter spaces. The reason is the *curse of dimensionality*: the computational complexity of a brute-force approach such as a direct gridding of the parameter space is growing exponentially with the dimension. Namely, if we need n points in each dimension to cover the parameter space, a d -dimensional grid will require n^d points in total.

This is why Bayesian inference relies on *sampling* the posterior distribution. For any random variable \mathbf{X} with a probability density $p(\mathbf{x})$, a set of independent discrete samples is a set of values $\{\mathbf{x}_i\}_{i=1,\dots,N}$ representing independent draws of \mathbf{X} . The defining property for a set of samples to faithfully represent a random variable is to ensure that for any test function $F(\mathbf{x})$ on parameter space,

$$\frac{1}{N} \sum_{i=1}^N F(\mathbf{x}_i) \sim \int d\mathbf{x} p(\mathbf{x}) F(\mathbf{x}) \quad [N \rightarrow +\infty]. \quad [7.52]$$

Another important property for these samples is to be independent from each other, although that is often only achieved approximately, as we will see below.

An important distinction is to be made between marginal and conditional posteriors: for a joint distribution of two variables $p(x, y)$, the marginal distribution for x is $p(x) = \int dy p(x, y)$ (an integral over y of the joint distribution) while the conditional distribution for a given value of y is $p(x|y)$ (a slice of the joint distribution). In practice, we almost always report marginal distributions in scientific papers, simply because we cannot visualize a multidimensional posterior. Marginalizing is transparent for samples representing a distribution, as one can just ignore the extra parameters to obtain samples of the marginalized posterior.

In order to summarize the information contained in a posterior, marginalized posterior samples can be readily used to obtain confidence intervals in individual parameters. Summarizing further, the *max-posterior* point is the optimum of the posterior (one often uses the *max-likelihood*, which is justified when the prior does not make too much difference). Note that the max-posterior point is not necessarily the maximum of the marginal posterior, due to the integration over the marginalized axis. The statistical mean of the distribution can also be used, although this notion is less representative e.g. for multimodal posteriors.

Another important qualitative feature of multidimensional probability distributions is the relative weight of the central regions and of the tails of the distribution. This can be illustrated by considering a Gaussian distribution in d dimensions. If we were to draw samples from such a multivariate normal, the distribution of $\xi \equiv \ln(p/p_{\text{peak}})$ itself, which is related to the half square distance from the peak (normalized by the covariance) would be

$$p(\xi) = \frac{1}{\Gamma(d/2)} (-\xi)^{d/2-1} e^{\xi}, \quad [7.53]$$

with an average of $\bar{\xi} = d/2$. Thus, the higher the dimension, the further away from the peak are the typical samples: the relative probability weight of the tails is higher. This effect is already apparent when going from 2 to 1 dimensions. Consider the 1-sigma level for the 1D marginals: it contains 68% of the probability in each marginal. However, in 2D, the ellipse corresponding to this level curve in p , at $\xi = \ln 1/2$, contains only 39% of the total probability.

7.6.3. Markov chain Monte Carlo methods

We will give here an abridged presentation of the simplest and most widely used method for producing samples from a posterior distribution, provided

that we can evaluate it as a function (as a product of the likelihood and prior, unnormalized): the Metropolis-Hastings algorithm for a Markov Chain Monte-Carlo (MCMC) approach.

A Markov chain is a sequence with no memory; if $\{x_i\}$ is the sequence, the algorithm determining the next step x_{K+1} depends on x_K but is fully independent from the set $\{x_i\}_{i \leq K-1}$. This is particularly simple to implement in code, and such an algorithm can be restarted at any point. The aim here is to obtain samples $\{x_i\}$ representative of a target probability density $p(x)$ as in [7.52] (x can be multidimensional, but we do not use boldface here). We will denote by $T(x, y)$ the transition probability, that is to say the probability to jump to y if the chain is in x , at each step of the chain.

The condition that we will require for the chain is the *stationarity* of the distribution:

$$\int dx p(x) T(x, y) = p(y), \quad [7.54]$$

Intuitively, this condition states that if we draw a starting position x according to the target distribution $p(x)$ and if we perform a random jump according to $T(x, y)$, then the landing point y is still distributed according to the target $p(y)$. We also need to chain to be *ergodic*, but we will leave this question aside here.

We will in fact focus on another condition, the *detailed balance* condition:

$$p(x) T(x, y) = p(y) T(y, x). \quad [7.55]$$

This condition will be easier to prove, and is a sufficient condition to ensure the stationarity: if verified,

$$\int dx p(x) T(x, y) = \int dx p(y) T(y, x) = p(y), \quad [7.56]$$

since $\int dx T(y, x) = 1$ as T is a probability distribution in its second argument.

In the Metropolis-Hastings algorithm for MCMC, one first specifies a jump *proposal* $q(x, y)$ which is a probability density in y given x (q is not necessarily symmetric in its arguments), and then an *acceptance* probability $\alpha(x, y)$. Starting from x , we draw a proposed new position y according to $q(x, y)$; with the probability $\alpha(x, y)$, the chain moves to the new position, and with the

probability $1 - \alpha(x, y)$, we reject the jump and the chain stays in place. The Metropolis-Hastings rule for the acceptance probability is

$$\alpha(x, y) = \min \left(1, \frac{q(y, x) p(y)}{q(x, y) p(x)} \right). \quad [7.57]$$

If the proposal is chosen to be symmetric, $q(x, y) = q(y, x)$, the acceptance is $\alpha(x, y) = \min(1, p(y)/p(x))$. In other words, if the proposed point is higher in probability, $p(y) > p(x)$, we always accept the jump; if it is lower, $p(y) < p(x)$, we accept the jump with probability $p(y)/p(x)$. The chain will always accept to go up, and sometimes accept to go down.

We can verify the detailed balance relation for $x \neq y$, for example assuming that $\alpha(x, y) = p(y)q(y, x)/p(x)q(x, y) < 1$. Then we have

$$p(x)T(x, y) = p(x)q(x, y)\alpha(x, y) = p(y)q(y, x) = p(y)T(y, x), \quad [7.58]$$

since $\alpha(y, x) = 1$ in that case.

7.6.4. Sampling techniques and parallel tempering

Although the MCMC-MH algorithm is a remarkably simple formal solution, in practice efficient sampling entails a lot of cooking recipes to ensure convergence in a reasonable number of steps, with an adaptation required to each target probability distribution. For a pedagogical exposition on MCMC, see (Hogg and Foreman-Mackey 2018).

The design of the proposal in the MCMC Metropolis-Hastings approach is left to the user, and will in general be tailored to the specific problem addressed. The simplest idea for $q(x, y)$ could be to use a Gaussian proposal in y centered on x , with a given covariance. The choice of this covariance illustrates an important underlying tradeoff: we wish to achieve a good *acceptance rate* for the chain (in order not to waste computation time), while also achieving a good *mixing* of the chain (exploring well the parameter space). Indeed, if the proposal was too wide, many jumps would go to irrelevant regions of low probability and be rejected; if the proposal was too narrow, the jumps would be often accepted but the chain would need many iterations to move away from its current position and explore the full parameter space.

Several methods exist to build adaptive proposals, meant to avoid fine-tuning issues. These include *ensemble sampling*, where a group of walkers is

evolved in parallel, with proposals being built from the other walkers; and *differential evolution*, where the history of the chain is used. Another approach is called *Hamiltonian Monte-Carlo*, where one relies on gradients of the posterior surface to integrate trajectories.

Another important limitation of the method is that samples from an MCMC chain are not, in fact, independent from each other. One useful diagnostic here is to compute an autocorrelation length for the chain, and to run the chain until one obtains a predetermined number of effectively independent samples, measured as the length of the chain divided by this autocorrelation length.

Importantly, due to the acceptance/mixing tradeoff, it is difficult to ensure that the chain would not remain stuck in secondary maxima of the distribution, missing the max-probability region entirely. One generic and simple technique addressing the issue of secondary maxima is *parallel tempering* (PT), which is a particular case of replica exchange. The idea is to add high-temperature chains that will explore a tempered (flattened) probability distribution, moving more easily between secondary maxima, while exchanging information with our original chain.

One can generalize the chain to be K replica chains in parallel, i.e. $x = (x_1, \dots, x_K)$. Separating the target posterior p as a prior π and a likelihood \mathcal{L} , we write a joint probability that is independent across replicas,

$$p(x) = \prod_{i=1}^K p_k(x_k), \quad p_k(x_k) = \pi(x_k) \mathcal{L}(x_k)^{\beta_k}, \quad \beta_k = \frac{1}{T_k}. \quad [7.59]$$

Here the T_k are the chain temperatures. $T_1 = 1$ for the original, “cold” chain, while the others $T_2 < \dots < T_K$ must be chosen to ensure a good exploration of the parameter space. The MCMC chain for x advances either by steps that evolve each of the x_k independently, or by steps that propose an exchange between replicas. For a swap proposal $(\dots, x_i, \dots, x_j, \dots) \rightarrow (\dots, x_j, \dots, x_i, \dots)$ (symmetric by construction), the MH rule for the acceptance is

$$\alpha_{ij} = \min \left(1, \frac{p_i(x_j) p_j(x_i)}{p_i(x_i) p_j(x_j)} \right) = \min \left(1, \left(\frac{\mathcal{L}(x_j)}{\mathcal{L}(x_i)} \right)^{\beta_i - \beta_j} \right). \quad [7.60]$$

This swap acceptance respects the detailed balance, and the accepted swaps move information between high- and low-temperature chains. One typically implements swaps between adjacent temperatures.

In the context of gravitational wave astronomy, a sampler implementing ensemble sampling, parallel tempering and many other features is **eryn** (Karnesis *et al.* 2023).

7.6.5. Nested sampling

So far, we have not described how the proposal distribution is to be chosen. This choice is free in principle, but needs in practice to be tailored to the problem at hand. The simplest idea would be to use a Gaussian proposal (automatically symmetric). The covariance illustrates an *acceptance versus mixing tradeoff*. With a covariance that would be too narrow, we would have a good acceptance rate (most jumps would be accepted), but poor mixing, as the chain would not explore enough of the parameter space; a covariance too broad, and the acceptance rate would decrease, as we would often propose jumps to irrelevant parts of the parameter space.

A different approach to sampling is provided by *nested sampling*; we will not give any details here and refer to (Ashton *et al.* 2022) for a review. This method allows to compute the evidence, but can also produce posterior samples. The method evolves an ensemble of live points; at each iteration, a new sample is drawn from the prior volume constrained to be of higher likelihood than the lowest-likelihood point, which is then replaced. Applying a weight to the live point replacements turns them into samples. Nested sampling has become of paramount importance in LVK data analysis, and is used in the Bilby pipeline (Ashton *et al.* 2019) responsible for parameter estimation in the current observing runs.

7.7. Parameter estimation for gravitational wave signals

7.7.1. The Whittle likelihood: stationary and Gaussian noise

We had seen in [7.22] that the probability density for stationary Gaussian noise (using the circulant approximation for finite-length data) can be written as $p(n) = e^{-(n|n)/2}$. In this section we use notations for functions, but we would use discrete data in real applications. From this probability distribution for the noise, which encodes our statistical understanding of the instrument, we can build the likelihood function. Assume that the gravitational wave signal is $h(\boldsymbol{\theta})$, with $\boldsymbol{\theta}$ the parameters that we want to infer, and assume that we have perfect knowledge of $h(\boldsymbol{\theta})$ (this is not true in reality, we always have to resort to approximate waveform templates); we can then write

$$d = h(\boldsymbol{\theta}) + n, \quad p(d|\boldsymbol{\theta}) = p(n = d - h(\boldsymbol{\theta})), \quad [7.61]$$

which given the probability distribution of the noise gives us

$$\ln p(d|\boldsymbol{\theta}) = -\frac{1}{2}(d - h(\boldsymbol{\theta})|d - h(\boldsymbol{\theta})) + \text{const}. \quad [7.62]$$

In other words, the likelihood is the probability that the noise realization explains the residuals between the observed data and the template. This Whittle likelihood amounts to computing the square norm of these residuals according to the noise-weighted inner product [7.20], which is done in practice using the sum [7.21] for discrete Fourier-domain data. The likelihood can then be used as the basis for Bayesian inference as described in Sec. ??.

Note that there will be applications where one wants to infer parameters of the noise PSD itself, rather than those of a deterministic signal in the data stream. In that case, we can still use the Whittle likelihood for Bayesian inference, but we need to include all PSD-dependent terms in [7.19].

7.7.2. The Fisher matrix approach

Sampling a posterior distribution can be costly. Whenever we want to simulate the process of parameter inference, for instance to explore the capabilities of future instruments, simulating many signals and performing Bayesian inference for each individual source. An important tool for simulating a simplified version of parameter estimation is the *Fisher matrix* approach (sometimes called the Information matrix). As we will see, it is also useful in shedding light on the effect of the noise, or of waveform inaccuracies.

The Fisher matrix approach is formally valid in the limit of a large $\text{SNR}^2 = (h(\boldsymbol{\theta}_0)|h(\boldsymbol{\theta}_0))$. We will not give the general definition of the Fisher matrix here, and directly use the simple form for our Whittle likelihood [7.66]. Ignoring the presence of the noise, if the data is $d = h_{\boldsymbol{\theta}_0}$, the likelihood peaks at $\boldsymbol{\theta} = \boldsymbol{\theta}_0$. At high SNR, we can then consider only the vicinity of the peak and use a linearized signal approximation, at first order in deviations from the peak parameters $\Delta\boldsymbol{\theta} = \boldsymbol{\theta} - \boldsymbol{\theta}_0$. The Whittle likelihood then takes the following form (with indices j, k spanning parameter dimensions, with implicit summation):

$$\begin{aligned} h(\boldsymbol{\theta}) &= h(\boldsymbol{\theta}_0) + \Delta\theta^j \partial_j h(\boldsymbol{\theta}_0) + \mathcal{O}(\Delta\theta^2), \\ \ln p(d|\boldsymbol{\theta}) &= -\frac{1}{2} \Delta\theta^j \Delta\theta^k (\partial_j h|\partial_k h) + \mathcal{O}(\Delta\theta^3). \end{aligned} \quad [7.63]$$

We see that at leading order, the likelihood becomes a Gaussian in $\Delta\boldsymbol{\theta}$, with a covariance $\boldsymbol{\Gamma}^{-1}$ which is the inverse of the Fisher (or Information) matrix

$$\boldsymbol{\Gamma}_{jk} = (\partial_j h|\partial_k h). \quad [7.64]$$

The diagonal elements of the covariance $\sigma_j = \sqrt{\boldsymbol{\Gamma}_{jj}^{-1}}$ give the $1\text{-}\sigma$ Gaussian errors for the different signal parameters. If we were to rescale the signal $h(\boldsymbol{\theta}_0)$

by a factor λ , we would have $\mathbf{\Gamma} \propto \lambda^2$; as a result, the Fisher uncertainties are inversely proportional to the SNR, or proportional to the luminosity distance¹, $\sigma_j \propto 1/\text{SNR} \propto d_L$.

A word of warning is needed regarding the use of the Fisher matrix (Vallisneri 2008). It applies formally in the limit of large SNRs; it is conceptually a local approximation, and will miss by construction any multimodality in parameter space with distant secondary peaks. In addition, the numerical computation of the Fisher matrix is often a delicate task. The matrix inversion needed to compute the covariance is vulnerable to degenerate or quasi-degenerate parameters: if there are directions in parameter space where one could move without changing the signal much, the Fisher matrix will have near-zero eigenvalues and its inversion will be problematic (while such nuisance parameters are often uninteresting physically). The formalism requires computing parameter space derivatives of the waveforms, which is simple for some parameters (e.g. the luminosity distance) but often needs to be done numerically; in the latter case, one needs to carefully check that the results are stable with respect to changes in the steps chosen for the discrete derivatives. Finally, it is inherently a Gaussian approximation to the likelihood, and is limited as such. The performance of the Fisher matrix can sometimes be improved by a change of variables, choosing a set of parameters in which the likelihood looks as Gaussian as possible.

7.7.3. Noise, errors and biases

The linearized signal approximation that we used in the Fisher matrix approach can also help us understand errors when including the effect of the noise in the inference. With a non-zero noise realization,

$$\ln p(d|\boldsymbol{\theta}) = -\frac{1}{2}(h(\boldsymbol{\theta}_0) - h(\boldsymbol{\theta}) + n|h(\boldsymbol{\theta}_0) - h(\boldsymbol{\theta}) + n) + \text{const}. \quad [7.65]$$

Applying the same linear signal expansion as in [7.63]², we have

$$\ln p(d|\boldsymbol{\theta}) = -\frac{1}{2}\Delta\theta^j\Delta\theta^k(\partial_j h|\partial_k h) + \Delta\theta^j(\partial_j h|n) - \frac{1}{2}(n|n) + \text{const}. \quad [7.66]$$

1. Care must be taken for the luminosity distance; in that case it is the relative error σ_{d_L}/d_L that scales as $\propto 1/\text{SNR}$.

2. Note that our derivations here are a bit simplified and improper, we refer to (Cutler and Vallisneri 2007) for more rigorous treatment.

Here the term $\langle n|n \rangle$ is a constant and plays no role. Maximizing with respect to $\Delta\theta$, we obtain that the maximum likelihood is achieved for a non-zero parameter shift $\Delta\theta_n^j$, with

$$\Delta\theta_n^j = \Gamma_{jk}^{-1}(\partial_k h|n), \quad \ln p(d|\theta) = -\frac{1}{2}(\Delta\theta^j - \Delta\theta_n^j)(\Delta\theta^k - \Delta\theta_n^k)\Gamma_{jk} + \text{const.} \quad [7.67]$$

Thus, in this approximation the likelihood is shifted to $\theta_0 + \Delta\theta_n$, and retains the same covariance Γ^{-1} .

The noise-induced shift (or error) $\Delta\theta_n$ depends linearly on n and has therefore zero mean, $\langle \Delta\theta_n \rangle = 0$. We can compute its variance using the property [7.23] as

$$\langle \Delta\theta_n^j \Delta\theta_n^k \rangle = \Gamma_{jl}^{-1} \Gamma_{km}^{-1} \langle (\partial_l h|n)(n|\partial_m h) \rangle = \Gamma_{jl}^{-1} \Gamma_{km}^{-1} (\partial_l h|\partial_m h) = \Gamma_{jk}^{-1}, \quad [7.68]$$

which shows that the variance of this shift coincides with the Fisher covariance itself.

This leads to an important observation, telling us what would happen if we were to repeat simulations of parameter inference for different noise realizations: in presence of noise, the posterior is shifted with respect to the true parameters, and this noise-induced error is random, zero-mean, and scales as the width of the posterior itself. At leading order, the noise does not affect the posterior width.

The same formalism can also be applied to the case of *waveform systematics*, when the templates h_m that we use in the likelihood differ from the true signals h in the data as $h_m = h + \delta h$. In that case, we obtain similarly a *bias* $\Delta\theta_{\delta h}$ given by

$$\Delta\theta_{\delta h}^j = -\Gamma_{jk}^{-1}(\partial_k h|\delta h). \quad [7.69]$$

This bias is often called the Cutler-Vallisneri bias (Cutler and Vallisneri 2007); contrarily to the noise-induced errors, it is deterministic. The objective of waveform modellers is to ensure that such systematic biases remain small enough that they are subdominant with respect to statistical uncertainties, that is to say $|\Delta\theta_{\delta h}| \ll \sigma$ (schematically).

7.7.4. An example parameter estimation result: GW150914

We give in Fig. 7.5 an example corner plot representing the posterior on the physical parameters of GW150914, downloaded from GWOSC. In such a




Figure 7.5: Posterior samples for GW150914, with only a subset of the 15 parameters: masses (m_1 , m_2), spin magnitudes (a_1 , a_2), luminosity distance (d_L), inclination (θ_{JN}), and sky position (ra, dec).

corner plot, each panel represents a marginal 2D distribution for a pair of parameters, while the diagonal panels give the marginal 1D distributions. Due to lack of space, we only show a subset of parameters (6 out of 15): masses, spins, distance, inclination, and sky position.

7.8. Other topics

7.8.1. Bursts and unmodelled signals

In the previous sections, we discussed the search for gravitational wave signals and the recovery of the source parameters using templates representing (to our best knowledge) the expected signal. As such, we have left aside an important question: what can we do if we do not dispose of templates at all, if we are looking for *unmodelled signals* ?

There are multiple motivations for considering this case. First, as a sanity check, we want to verify that we can also do science without relying on our theoretical predictions, possibly flawed. Second, there are also astrophysical classes of sources with a complicated physics, for which building templates is challenging; primary candidates here are signals from core collapse supernovae, that can produce bursts of gravitational radiation among the targets of ground-based detectors. Third, gravitational wave science is a discovery science, and we do not want to exclude signals from exotic objects and/or unknown astrophysical processes.

Unmodelled searches are run in parallel of template-based searches for LVK observations. One such example is *CWB* (Coherent Wave Burst) (Drago *et al.* 2020). Due to the lack of templates, they are not as sensitive to extract long-lived low-amplitude signals. As a result, they achieve their best sensitivity for high-mass binary black holes. In fact, the first detection of GW150914 was made concurrently by template-based and unmodelled searches!

Even in the absence of detailed templates $h_+(t), h_\times(t)$, one can still exploit the minimal physical assumptions that the gravitational wave contains only the two polarizations plus and cross, and propagates at the speed of light. Armed with these hypotheses, we can hope to distinguish between noise transients and GW bursts by leveraging the fact that the former should be incoherent between detectors, while the latter signals should be coherent (with a small time delay related to sky position) and consistent in amplitude and phase with the known detector response to a h_+, h_\times signal.

Unmodelled burst analyses typically take place in the time-frequency domain, using a wavelet representation. In this representation, the signal will show as high-power regions, especially around merger for a CBC. Clustering algorithms are used to identify the time-frequency pixels corresponding to the signal. To distinguish between noise and signals, one builds a coherent energy E_c from cross-correlations between detectors, and a noise energy E_n from the auto-correlation in each detector after the correlated signal is subtracted. The detection statistic is then built as

$$\eta = \sqrt{\frac{2E_c^2}{E_c + E_n}}. \quad [7.70]$$

The information about the time shift and relative signal amplitude between the two detectors can also be used to reconstruct the position of the source, by scanning the sky.

7.8.2. Continuous waves and long-lived signals

In the standard picture that we took so far, the signal was assumed to be a CBC with a short duration. In the context of LVK, where even BNS signals last at most a couple of minutes, this allowed us to model the detector noise as stationary, and to ignore the motion of the detector itself during the observations. Crucially, such CBC signals have a limited number of orbital cycles in the detector's sensitive frequency band.

However, we know of astrophysical systems that could emit almost-monochromatic GW signals, that would last for the entirety of our observing runs: as shown by observations of pulsars, neutron stars can rotate to high velocities, and could emit gravitational waves if they presented a deviation from axisymmetry. We observe the spin-down of many pulsars, and this puts an upper limit on the possible angular momentum loss through GW emission. The expected signal is extremely low in amplitude, and we can only hope to detect it by leveraging the power of matched filtering, accumulating coherence between template and signal over a long time.

However, using matched filtering for these *continuous waves* sources comes with significant challenges. The first is simply the size of the data at hand: at a sampling frequency of 1kHz, a single day of data represents $\sim 10^8$ samples and a 3-months run, $\sim 10^{10}$. The second is the size of the parameter space to cover, in the meaning of the number of templates that would be needed to reliably catch the signal.

The signal from such a continuous source (see Chapter ??) will be characterized by a main frequency f_0 , possibly its derivative \dot{f}_0 (and maybe second derivative \ddot{f}_0). On long timescales, we also have to take into account the relative motion of the source and detector, which creates a varying Doppler effect in the observed frequency, and can also be represented as a varying delay (the Roemer delay):

$$f_{\text{obs}}(t) = f_{\text{source}}(t) \left(1 - \frac{\mathbf{k} \cdot \mathbf{v}(t)}{c} \right), \phi_{\text{obs}}(t) = \phi_{\text{source}} \left(t + \frac{\mathbf{k} \cdot \mathbf{x}(t)}{c} \right), [7.71]$$

where \mathbf{k} is the propagation vector, \mathbf{v} the relative velocity, and $\mathbf{x}(t) = \mathbf{x}_{\text{source}}(t) - \mathbf{x}_{\text{det}}(t)$ the separation between source and detector. The motion of the Earth around the Sun, as well as the motion of the detector due

to the Earth's rotation, give large modulations contributing a large number of cycles (see estimates in Chapter 7.6 of (Maggiore 2007)). If the pulsar is in a binary system, there is an additional delay due to the binary's orbit.³

This creates a strong contrast between *blind*, *directed* and *targeted* searches. In directed searches, we aim for a known candidate source (such as the Vela or Crab pulsars), which gives us the position of the source in the sky with a great accuracy. This eliminates two degrees of freedom from the problem, and we can de-modulate the signal, effectively removing the delays in [7.71]. If we have radio observations of a pulsar, we also know the frequency with great accuracy, considerably reducing the search space in a targeted search. In blind searches, all those parameters are left free. Because of the very large number of cycles observed for over months of signal, the frequency resolution is tiny. The orbital modulations in [7.71] also contribute a large number of cycles, which leads to a tiny resolution in the sky position. In other words: a very small change in f_0 or in the sky position will make our template go out of phase with the signal, and we will need a very large number of templates for such a blind search. Estimates on these resolutions can be found in Chapter 7.6 of (Maggiore 2007), and the result is that a blind search with a full data stream is *unfeasible*. Importantly, the duration of observation T is a strong driver for the cost of the search $\propto T^n$, with n at least 4 and possibly larger if e.g. the pulsar spin-down is included.

This motivates introducing a different approach: *semi-coherent searches*. The idea is to segment the data stream and use templates that are only coherent with the signal on each segment, instead of the full data. The price to pay in such methods is an increase in false alarms, as semi-coherent templates are more likely to pick up correlations with noise; as a result, the threshold of the minimal detectable signal amplitude is larger.

We can get qualitative scalings as follows. If the search statistic is not the matched-filter SNR [??], but instead a semi-coherent search statistic built as a sum over $k = 1, \dots, K$ segments as

$$\hat{\Lambda}_K = \sum_{k=1}^K \hat{\rho}_k^2, \quad \hat{\rho}_{*k}^2 = \max_{\phi_k} \hat{\rho}_k^2(\phi_k), \quad [7.72]$$

where we optimize the matched filter SNR in each segment over a free phase degree of freedom ϕ_k . This new search statistic $\hat{\Lambda}_K$ would then follow an off-centered χ^2 distribution with $2K$ degrees of freedom, which for large K behaves

3. Note that real pulsars can also present glitches, which will cause jumps in the frequency evolution.

as a normal distribution with

$$\text{In noise: } \hat{\Lambda}_K \sim \mathcal{N}(2K, 4K) \quad [K \rightarrow +\infty],$$

$$\text{In presence of signal: } \hat{\Lambda}_K \sim \mathcal{N}(\rho_{\text{opt}}^2 + 2K, 4\rho_{\text{opt}}^2 + 4K) \quad [K \rightarrow +\infty], \quad [7.73]$$

We write schematically here: in reality, one would need to go through the optimization procedure as in Sec. ??; and for complex signals, the above is only valid in the large- K limit. The important point is the following: on one hand, by allowing the use of templates that only need to remain coherent with the signal over shorter segments of length T/K , if the size of the template bank scales as $\sim T^n$, we can now use $\sim (T/K)^n$ templates on each segment. On the other hand, the variance is now linear in K , which means that the detection threshold allowing to reject the noise-only hypothesis at a given confidence level is $\rho_{\text{thres}}^2 \sim K$. In summary:

$$\text{Computational cost: } K \times (T/K)^n \sim T^n / K^{n-1},$$

$$\text{Detection threshold: } \rho_{\text{thres}} \sim K^{1/4}. \quad [7.74]$$

With $n \geq 4$, we can see that by increasing K , one can achieve indispensable computational gains, at the expense of a moderate loss of sensitivity.

Finally, an important aspect of such analyses is that the detector cannot be considered stationary on long timescales, and the data comes with interruptions due to the duty cycle of the detectors. This is tackled by segmenting the data stream in independent segments. Analyzing the whole data from a given run also requires to handle data quality (glitches, ...) everywhere and not just around CBC events.

For further reading on continuous waves, see the review (Prix 2009).

7.8.3. Stochastic backgrounds

So far, we have taken a deterministic view of the signal, representing it as $h(\boldsymbol{\theta})$, the emission from a single source with certain parameters $\boldsymbol{\theta}$. However, the gravitational wave signal could also be *stochastic* in nature. There are different scenarios: first, a stochastic signal can be formed by the superposition of individual signals generated by many inspiralling binaries. We already know of such a signal, indirectly: the population of binary black holes and neutron stars observed by LVK leaves such a signal after we identify only the loudest, nearby sources. In the case of LISA, the number of galactic binaires emitting in band is so large that we expect them to create a stochastic *foreground* (see Chapter ??), contributing a confusion noise to be considered alongside instrumental noise. A gravitational wave background could also be a relic of processes

in the early universe, such as phase transitions or inflation, that would produce anisotropic stress, source of tensor perturbations in the metric. In all cases, stochastic signals are of great physical importance and one of the targets of future GW astronomy.

A comprehensive review on data analysis for stochastic backgrounds of gravitational waves (SGWBs) is given in (Romano and Cornish 2017).

One can represent a SGWB as a superposition, for each polarization $A = +, \times$, of plane waves in a Fourier representation, propagating in each direction \mathbf{n} :

$$h_{ij}(t, \mathbf{x}) = \sum_{A=+, \times} \int df \int d^2\mathbf{n} \tilde{h}_A(f, \mathbf{n}) e^{-2i\pi f(t - \mathbf{n} \cdot \mathbf{x}/c)}. \quad [7.75]$$

Here, the signal \tilde{h}_A is considered itself as a stochastic process. Crucial simplifying hypotheses are to assume that the background is *stationary*, *isotropic* and *Gaussian*. Isotropy is justified by the cosmological principle, while stationarity and Gaussianity are justified if we think of the signal as being the superposition of a large number of weak individual signals, using the law of large numbers. Under those assumptions we have

$$\langle \tilde{h}_A(f, \mathbf{n}) \tilde{h}_B(f', \mathbf{n}') \rangle = \frac{1}{2} S_h(f) \delta(f - f') \frac{\delta^2(\mathbf{n}, \mathbf{n}')}{4\pi}. \quad [7.76]$$

The spectral density of the GW background $S_h(f)$ plays a role similar to the noise PSD $S_n(f)$ that we introduced in Sec. ??.

There are different approaches to the detection and characterization of such a signal. In a single detector, telling apart the stochastic signal S_h from instrumental noise S_n can be challenging, since the noise PSD must also be learned from the data. One avenue is to exploit a known (parametrized) spectral shape for S_h , as well as assuming a parametrized model for S_n motivated by the physics of the instrument. If the two components differ in spectral shape, one can use the likelihood formed by [??] with $S_n \rightarrow S_n + S_h$ (removing the deterministic $h(\boldsymbol{\theta})$) and infer the parameters of both models. For instance, this can be used for a loud background, or for a background with distinct spectral features (e.g. a bump in its spectrum).

A more informative approach is based on cross-correlation between detectors. Assuming that the noise in the detectors is independent (for instance, if they are far apart geographically), we can exploit the fact that the GW stochastic signal is coherent, in a sense to be clarified. By accumulating data over long

periods of time, we can tell apart this coherent signal from the incoherent instrumental noise.

Based on this idea, one can introduce a filter as in Sec. 7.4.2:

$$\hat{s} = \int dt \int dt' s_1(t) s_2(t') Q(t - t'), \quad [7.77]$$

where $s_1(t)$, $s_2(t)$ are the data streams from the two detectors, and Q is the filter we want to choose to optimize the signal-to-noise ratio, following a similar procedure to Sec. 7.4.3. When correlating the signal at the two detector locations \mathbf{x}_1 and \mathbf{x}_2 , we have to take into account the phase factors in [7.75], as well as the detector response $s_j = \sum_A F_A^j h_A^j(t, \mathbf{x}_j)$, with the detector label $j = 1, 2$ and F_+^j, F_\times^j the detector pattern functions (for LVK-like instruments). This gives rise to the following *overlap reduction function*:

$$\Gamma_{12}(f) = \int \frac{d^2\mathbf{n}}{4\pi} \left[\sum_A F_A^1(\mathbf{n}) F_A^2(\mathbf{n}) \right] e^{2i\pi f \mathbf{n} \cdot \Delta \mathbf{x} / c}, \quad \gamma_{12}(f) = \frac{\Gamma_{12}(f)}{\Gamma_{12}(0)}. \quad [7.78]$$

with $\Delta \mathbf{x} = \mathbf{x}_1 - \mathbf{x}_2$ the detector separation. The normalized overlap $\gamma_{12}(f)$ is 1 at low frequencies, but oscillates and drops to 0 at high frequencies, representing the fact that the gravitational waves become out-of-phase due to the propagation between the two detectors – in this regime, correlations are not informative anymore. How fast the fall-off occurs depends on the distance between detectors; for instance, for the pair of LIGO detectors in Livingston and Hanford, the fall-off is around ~ 50 Hz.

We will only quote here the result for the optimal filter and SNR, without demonstration:

$$\tilde{Q}(f) = \frac{\gamma_{12}(f) S_h(f)}{S_n^2(f)}, \quad \left(\frac{S}{N} \right)^2 = 2T \int_0^{+\infty} df \frac{\gamma_{12}^2(f) S_h^2(f)}{S_n^2(f)}. \quad [7.79]$$

This optimal filter, in analogy with the CBC case using templates, requires knowing the spectral density of the signal $S_h(f)$. The prefactor T is the total observation time; we can see that the SNR increases as \sqrt{T} as we integrate more data, exploiting the coherence of the SGWB versus the incoherence of the noise across detectors.

Note that the assumptions of stationarity, Gaussianity and isotropy have limitations. This is particularly true for the background formed by unresolved CBCs: in amplitude, the combined signal might have a strong component from

the loudest unresolved signals, just below the SNR detection threshold, which will appear as intermittent transients in the stochastic process – a phenomenon well described as *pop-corn noise*. New methods are currently being explored to go beyond these simplifying assumptions. Anisotropies in the background can also be targeted, although they could be challenging to detect as a sub-dominant component.

7.8.4. Hierarchical inference

As current detectors keep accumulating higher-quality data and detections, gravitational wave astronomy is evolving, from the discovery science that it was at the time of the first detections, to a statistical science now and in its future. With close to 300 CBC detections (as of March 2025, during the run O4), LVK observations now allow us to probe the populations of astrophysical sources, and not only the parameters of individual events, thus allowing us to infer the distribution of masses, spins etc. of the compact object binaries in our universe. Another such application is gravitational wave cosmology, where we infer cosmological parameters such as the Hubble constant H_0 from the collective observations of luminosity distance of CBCs combined with information about their redshift (see Chapter ??). Inferring hyperparameters, that is to say parameters that affect our set of observations as a whole, is a process called *hierarchical inference*: we assume that we can model the parameter distribution of the population $\pi(\boldsymbol{\theta}|\Lambda)$ and the rate of events N_s , and we want to sample from a posterior $p(\Lambda|\{d\})$ given a set of observations $\{d\}$.

The detection process of a GW detector can be characterized as an inhomogeneous Poisson process, incorporating selection effects, as described in details in (Vitale *et al.* 2020). The occurrence of CBCs in the universe can be treated as stochastic, and on any given time interval, we expect the number of events to follow a Poisson statistic: the probability of N events occurring when the expected value is \bar{N} is

$$p(N|\bar{N}) = \frac{\bar{N}^N}{N!} e^{-\bar{N}}. \quad [7.80]$$

One needs to distinguish the rate of astrophysical events from the rate of detection. Crucially, the detector's ability to detect signals depends on their parameters. For instance, high-mass systems are intrinsically louder and visible to larger distances than low-mass events; if we want the intrinsic distribution of masses in the population, this *selection effect* has to be taken into account. To be complete, one would also need to include false detections, when a random noise fluctuation causes a spurious trigger.




Figure 7.6: CBC masses reported by LVK in the GWTC-3 catalog. The top panel shows posteriors in m_1 and m_2 for individual events. The bottom panel shows the result of a hierarchical inference for the rate and distribution of m_1 in GWTC-2 and GWTC-3. Figures taken from (Abbott *et al.* 2023a) and (Abbott *et al.* 2023b).

Under the assumptions that we have at most one trigger per data segment (verified so far for LVK), and ignoring the false alarms in the formalism for simplicity, one can derive the following *hierarchical posterior*, giving the

posterior distribution for the hyperparameters Λ given the set of observations $\{d_j\}_{j=1,\dots,N_{\text{det}}}$ for N_{det} detections as

$$p(\Lambda|\{d\}) = \frac{p(\Lambda)}{p(\{d\})} e^{-N_{\text{exp}} N_{\text{exp}}^{N_{\text{det}}} \alpha(\Lambda)^{-N_{\text{det}}}} \prod_{j=1}^{N_{\text{det}}} \int d\boldsymbol{\theta} p(d_j|\boldsymbol{\theta}) \pi(\boldsymbol{\theta}|\Lambda). \quad [7.81]$$

Here, $p(\Lambda)$ is our prior on hyperparameters, $p(\{d\})$ is the evidence, and N_{exp} is the expected number of detections, taking selection into account: if N_s is the expected number of sources, $N_{\text{exp}} = \alpha(\Lambda) N_s$ with the selection fraction $\alpha(\Lambda)$ depending on the hyperparameters as

$$\alpha(\Lambda) = \int d\boldsymbol{\theta} p_{\text{det}}(\boldsymbol{\theta}) \pi(\boldsymbol{\theta}|\Lambda), \quad p_{\text{det}}(\boldsymbol{\theta}) = \int_{\rho(d) > \rho_{\text{thres}}} dd p(d|\boldsymbol{\theta}). \quad [7.82]$$

The selection fraction $\alpha(\Lambda)$ is integrated over the source parameters predicted by the population model for the given hyperparameters $\pi(\boldsymbol{\theta}|\Lambda)$. We have written the detection probability for given source parameters $p_{\text{det}}(\boldsymbol{\theta})$ schematically with a threshold in the SNR ρ . In reality, it is of crucial importance to compute selection effects accurately, in accordance with the detection metrics that were used in the actual instruments; this is done in practice through the generation of large *injection sets*, simulating many detections or non-detections across parameter space, using fully realistic search pipelines.

In practice, the integrals in [7.81] can be computed thanks to the samples $\boldsymbol{\theta}_k$ of $p(\boldsymbol{\theta}|d_j)$ produced by parameter estimation for the individual detections, discretizing the integrals as a sum over those samples. Since these samples were produced with a prior $p(\boldsymbol{\theta})$ that differs from $\pi(\boldsymbol{\theta}|\Lambda)$, one has to apply a reweighting factor $\pi(\boldsymbol{\theta}|\Lambda)/p(\boldsymbol{\theta})$ to each sample. Ignoring the evidence $p(d_j)$,

$$\int d\boldsymbol{\theta} p(d_j|\boldsymbol{\theta}) \pi(\boldsymbol{\theta}|\Lambda) = \int d\boldsymbol{\theta} \frac{p(\boldsymbol{\theta}|d_j) p(d_j)}{p(\boldsymbol{\theta})} \pi(\boldsymbol{\theta}|\Lambda) \propto \sum_{\text{samples } k} \frac{\pi(\boldsymbol{\theta}_k|\Lambda)}{p(\boldsymbol{\theta}_k)}. \quad [7.83]$$

The tools of hierarchical inference are very general, and they are also used in order to infer constraints on deviations from GR using a collection of observations. They are crucial to extract information about the sources as a whole, and will play an ever increasing role, with the sensitivity of future detectors allowing them to probe a much larger fraction of our universe.

We show in Fig. 7.6 examples of mass posteriors reported by LVK in GWTC-3 (Abbott *et al.* 2023a), as well as the result of a result of a hierarchical inference for the rate and distribution of the primary mass (Abbott *et al.* 2023b).

7.9. Future challenges in data analysis

As current gravitational wave detectors keep improving their sensitivity, data analysis will have to shift some of its paradigms. The increased number of detections will pose a first computational challenge, requiring a higher level of automatization of the pipelines and motivating improvements to the efficiency of our algorithms.

Future instruments such as the third generation ground based detectors Einstein Telescope and Cosmic Explorer and the LISA instrument will drastically change the landscape of gravitational wave science. The detection horizon will include most of binary black holes, and a large fraction of binary neutron stars in our universe, leading to high rates of CBC detections per day (see the data analysis section of (Abac *et al.* 2025)). Source superposition will start to become a challenge, as most signals will overlap in time with their neighbours; BNS signals will last for several hours in band. Elusive signals such as continuous waves and astrophysical stochastic backgrounds will be within reach.

LISA data (see (Colpi *et al.* 2024) and Chapter ??) will come with specific challenges: the data stream will be signal-dominated, with a large number of quasi-monochromatic signals creating a confusion foreground, in superposition with transient signals. Because of the timescales involved, the non-stationarity of the instrument will need to be taken into account. The data stream will also be interrupted by data gaps. Due to the absence of noise-only data and to the superposition of GW signals together, a *global fit* strategy will need to be applied, inferring simultaneously the properties of the GW sources and of the instrumental noise.

For both ET/CE and LISA, the much higher SNRs will turn gravitational wave science into a precision science, allowing to extract detailed information about the sources, and test deviations from GR with great precision; but this will come with a challenge for the accuracy requirements on our waveform models, to keep systematics in check – an issue that is in fact already present for our current loudest signals. A fine understanding of the instruments and of their calibration will also become more important.

Machine learning approaches (see Chapter ??) have been proposed to tackle some of these challenges and are under active development. These necessary developments make gravitational wave data analysis a very active field overall, of primary importance for the exciting scientific discoveries to come.

7.10. Acknowledgements

This research has made use of data or software obtained from the Gravitational Wave Open Science Center (gwosc.org), a service of the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA.

7.11. Bibliography

- Abac, A. *et al.* (2025), The Science of the Einstein Telescope, , .
- Abbott, B. P. *et al.* (2016a), GW150914: First results from the search for binary black hole coalescence with Advanced LIGO, *Phys. Rev. D*, 93(12), 122003.
- Abbott, B. P. *et al.* (2016b), Observation of Gravitational Waves from a Binary Black Hole Merger, *Phys. Rev. Lett.*, 116(6), 061102.
- Abbott, B. P. *et al.* (2020), A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals, *Class. Quant. Grav.*, 37(5), 055002.
- Abbott, R. *et al.* (2023a), GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run, *Phys. Rev. X*, 13(4), 041039.
- Abbott, R. *et al.* (2023b), Population of Merging Compact Binaries Inferred Using Gravitational Waves through GWTC-3, *Phys. Rev. X*, 13(1), 011048.
- Allen, B. (2005), χ^2 time-frequency discriminator for gravitational wave detection, *Phys. Rev. D*, 71, 062001.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.71.062001>
- Alléné, C. *et al.* (2025), The MBTA pipeline for detecting compact binary coalescences in the fourth LIGO-Virgo-KAGRA observing run, *Class. Quant. Grav.*, 42(10), 105009.
- Ashton, G. *et al.* (2019), BILBY: A user-friendly Bayesian inference library for gravitational-wave astronomy, *Astrophys. J. Suppl.*, 241(2), 27.
- Ashton, G. *et al.* (2022), Nested sampling for physical scientists, *Nature*, 2.
- Chatziioannou, K., Dent, T., Fishbach, M., Ohme, F., Pürrer, M., Raymond, V., Veitch, J. (2024), Compact binary coalescences: gravitational-wave astronomy with ground-based detectors, , .
- Colpi, M. *et al.* (2024), LISA Definition Study Report, , .
- Cornish, N. J., Littenberg, T. B. (2015), BayesWave: Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches, *Class. Quant. Grav.*, 32(13), 135012.
- Cutler, C., Vallisneri, M. (2007), LISA detections of massive black hole inspirals: Parameter extraction errors due to inaccurate template waveforms, *Phys. Rev. D*, 76, 104018.

- Dal Canton, T., Nitz, A. H., Gadre, B., Cabourn Davies, G. S., Villa-Ortega, V., Dent, T., Harry, I., Xiao, L. (2021), Real-time search for compact binary mergers in advanced ligo and virgo's third observing run using pycbc live, *The Astrophysical Journal*, 923(2), 254.
URL: <http://dx.doi.org/10.3847/1538-4357/ac2f9a>
- Davis, D. *et al.* (2021), LIGO detector characterization in the second and third observing runs, *Class. Quant. Grav.*, 38(13), 135014.
- Drago, M. *et al.* (2020), Coherent WaveBurst, a pipeline for unmodeled gravitational-wave data analysis, , .
- Ewing, B. *et al.* (2024), Performance of the low-latency GstLAL inspiral search towards LIGO, Virgo, and KAGRA's fourth observing run, *Phys. Rev. D*, 109(4), 042008.
- Hogg, D. W., Foreman-Mackey, D. (2018), Data analysis recipes: Using Markov Chain Monte Carlo, *Astrophys. J. Suppl.*, 236(1), 11.
<https://gwosc.org/> (n.d.).
- Karnesis, N., Katz, M. L., Korsakova, N., Gair, J. R., Stergioulas, N. (2023), Eryn: a multipurpose sampler for Bayesian inference, *Mon. Not. Roy. Astron. Soc.*, 526(4), 4814–4830.
- Maggiore, M. (2007), *Gravitational Waves. Vol. 1: Theory and Experiments*, Oxford University Press.
- Owen, B. J., Sathyaprakash, B. S. (1999), Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement, *Phys. Rev. D*, 60, 022002.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.60.022002>
- Prix, R. (2009), Gravitational Waves from Spinning Neutron Stars, *Astrophys. Space Sci. Lib.*, 357, 651–685.
- Romano, J. D., Cornish, N. J. (2017), Detection methods for stochastic gravitational-wave backgrounds: a unified treatment, *Living Rev. Rel.*, 20(1), 2.
- Vallisneri, M. (2008), Use and abuse of the Fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects, *Phys. Rev. D*, 77, 042001.
- Vitale, S., Gerosa, D., Farr, W. M., Taylor, S. R. (2020), Inferring the properties of a population of compact binaries in presence of selection effects, , .

DRAFT

8

Artificial Intelligence for data analysis

Natalia Korsakova¹

¹ *Université Côte d'Azur, ARTEMIS, Observatoire de la Côte d'Azur,
CNRS, Bd de l'Observatoire, F-06304 Nice, France*

8.1. Foreword

In this chapter, we are going to review the basics of artificial intelligence (AI) and study its applications to the data analysis of Gravitational Waves (GW). The topic is extremely broad, as AI can be useful in many places where there is a need for speeding up the computations or generalizing non-linear functions. However, here we will focus on one subject which is the cornerstone of the data analysis, namely, parameter estimation. Over the course of this chapter we will try to understand the variety of methods that can be used for this task, comprehend the difference between them, and apply them to the realistic simulated data.

Often, data analysis for GWs is divided into the search and parameter estimation parts. The basic difference is that the search provides an approximate solution which is located within the measurement error to the true value but does not give the error on the estimation of the parameter itself. Typically it is implemented with the methods that can return the maximum likelihood value. On the contrary the parameter estimation part has to provide the full Bayesian posterior. In the case when we can perform the observation only once we should rely on the Bayesian approach to correctly take the measurement error into account.

Bayesian inference is governed by the Bayes equation for the parameter estimation. This equation cannot be solved exactly, therefore we need to perform some sort of approximation or simplification of the model. The classical approach to GW data analysis searches for the solution using approximate sampling techniques, namely, the variety of Markov chain Monte Carlo (MCMC) or nested sampling methods. The advantages of these methods is that they have been studied extensively over the past years. However, they rely on simplifications to the noise model and can be slow due to the sequential evaluation of the likelihood, which can become a bottleneck in terms of performance.

Recent developments in the field of generative modeling in machine learning have brought a new perspective on the Bayesian parameter estimation and opened the opportunities for alternative methods to perform inference. There are different techniques in the field of machine learning which can propose a solution to the problem of approximating the intractable distribution which appears in the Bayes equation. We can approximate the intractable distribution with the simpler distribution and put the upper bound on the estimation with the Variational inference, that will be discussed in the sections later [NK link section](#). The other approach is to study the ration of the probabilities and to train the classifier that will help to accelerate the sampling by providing the fast likelihood estimation [NK cite](#). Another way will be to use implicit models that transforms the distribution of interest to a simple base distribution using parameterized transform and change of variable equation. An example, of such model is invertible differentiable transforms which are commonly referred to

as normalizing flows. We will study them in detail in the section [NK reference](#) and look at their applications to GW data analysis.

Advantages of such methods is that the inference is done separately from the training, which allows to do all time consuming processing in advance and do parameter estimation in seconds. At the same time these methods do not rely on the particular noise model for the likelihood, but are based on the access to the samples from the joint distribution of the data and parameters, we can produce simulated data with various noise assumptions not restricted by the Gaussian model. The disadvantage of these methods is that in the case the data has something we did not expect in advance we will need to retrain the models.

8.2. Basic machine learning background




Figure 8.1: Artificial intelligence, machine learning and deep learning relation.

Let us refresh the knowledge on the basic principles of machine learning, core architectures of the neural networks, optimization functions and training process. We will revise some concepts that the reader might be familiar with but to make this chapter accessible to anyone we want to introduce a minimum background knowledge that will be needed to understand more advanced ideas. We recommend following books as a broader introduction to machine learning (Bishop 2006) and deep learning (Goodfellow *et al* 2016).

First of all let us start with some definitions which are sometimes used interchangeably in the literature and can be confusing for a reader.

The roots of this discipline lie in scientists' attempts to understand how the human brain functions. Foundational work by Rosenblatt (Rosenblatt 1958) introduced the concept of the perceptron – an attempt to represent how a human discriminate between shapes and characters, i.e. how the human perception works. The overarching goal of the field is to mathematically model the functioning of the human brain, a pursuit known as Artificial Intelligence (AI).

Machine Learning (ML) can be viewed as a subfield of AI, focusing on the development of mathematical methods for tasks such as regression, classification, and pattern recognition. A major breakthrough came with the introduction of deep neural networks consisting of multiple layers, processed efficiently using Graphical Processing Units (GPUs). This advancement significantly improved performance and gave rise to what is now known as Deep Learning (DL) (see Figure 8.1). In this text, we will use the terms Artificial Intelligence, Machine Learning, and Deep Learning interchangeably. At the same time many concepts in ML come from the related fields of mathematics, such as probability theory, set theory, optimal transport theory, linear algebra, geometry etc and can have names that they inherit from different field signifying similar things. We will try to clarify this in the context of GW physics.

ML is typically subdivided into major subsections based on the way the training is done, – supervised, unsupervised and reinforcement learning. In reinforcement learning, we learn from the response of the external environment to the actions that the agent perform within it. This approach is outside of the scope of this chapter and we refer the interested reader to the following book (Sutton, Barto 2015). The boundary between supervised and unsupervised learning is not very rigid. Generally speaking in the case of supervised learning we can say that we are estimating the conditional probability $p(\mathbf{y}|\mathbf{x})$ and we have access to the training data both for \mathbf{x} and for \mathbf{y} . Whereas in unsupervised case we try to learn the distribution of the data $p(\mathbf{x})$ by the other means of optimization, for example, by finding the patterns in the data. There are some cases where the separation between supervised and unsupervised learning can be ambiguous, for example, in fo sequence modeling, where the distribution of the data can be presented as a conditional distribution of the samples on the previous observations $p(\mathbf{x}) = \prod_{i=1}^N p(x_i|x_1, \dots, x_{i-1})$.

8.2.1. Simple network architectures

Neural networks can be viewed as way for approximating complex non-linear functions. These functions are built from simple components known as artificial neurons (see Figure 8.2), and are defined by millions of parameters referred to as weights. In graphical representations, these weights are typically shown as edges, while the transformations performed by the neurons are represented as nodes.

We can summarize the Figure 8.2 with the equation

$$\phi \left(b + \sum_{i=1}^n w_i x_i \right) = \hat{y}, \quad [8.1]$$




Figure 8.2: Artificial neuron, building block of neural networks.

where $\mathbf{x} = \{x_i\}$ is an input vector which is multiplied by the weights w_i and is summed with a bias b . In addition to that a nonlinearity ϕ is applied to the sum. For the nonlinearities the choice of the function is usually dictated by the type of the required output and the performed task (for example, regression or classification). The typical choice of the functions are sigmoid function, hyperbolic tangent, rectified linear unit, etc. The output \hat{y} is propagated as the input to the next node.

We can combine artificial neurons in different architectures. The most basic original architecture is fully connected networks, which is called multi-layer perceptron (MLP) and is constructed of the layers of neurons when each neuron in a layer is connected to all neurons in the previous and the following layers (these layers in the common ML frameworks are called, for example, Linear or Dense). This is a very simple architecture, it is therefore quite robust for training and can serve as an approximation to a function to a good precision. However the big disadvantage is that these networks occupy a lot of memory due to connections present between each neuron and, for the same reason, these networks cannot be very deep.

In (Lecun Y. *et al* 1998) the authors first time apply backpropagation to the Convolution Neural Network (CNN) architecture. In this case we can make the network deeper keeping a reasonable number of the weight because we do not train the connection with each node of the network but just with the filters which are convolved with the data and the hidden layers.

To improve the stability of the training for the deep networks residual neural networks were proposed (He K. *et al* 2016). The adjustment to the architecture is simple, but has shown to be powerful and is still widely used. In this architecture the input is added to the hidden layers at a certain cadence, establishing a residual connection, which allows to propagate the initial information through the network.

If our observation is represented as a time series it is advantageous to take

a sequential nature of the data into account. The architectures that were developed to capture such structure of data was first introduced as a recurrent neural networks (RNNs) and further developed as more modern variants like Long-Short Term Memory (LSTM) (Hochreiter and Schmidhuber 1997) models, which deal with the problem of vanishing gradients.

However, modern state of the art architectures are based on the same principles both for the sequential data, images and other data representations. This type of architecture is called Transformers (Vaswani *et al* 2017) and is based on the principle of attention.

Last but not the least graph neural networks (Bruna *et al* 2014) has gained popularity in recent years that can be useful to represent the graphs and ordered structures in space.

8.2.2. Optimization function

One of the most important parts of designing machine learning algorithm is the choice of the optimization function. Optimization function is sometimes also called in the literature an error function, a cost function or an objective function.

Generally speaking we can say that we are trying to maximize probability distribution of the data or in other words find the maximum likelihood distribution. Note that in the ML literature the maximum likelihood of the data is formulated in terms of the parameters of the network (which we will define θ here).

$$\begin{aligned}
 \hat{\theta} &= \arg \max_{\theta} \mathcal{L}(\theta) \\
 &= \arg \max_{\theta} p(\mathbf{X}, \theta) \\
 &= \arg \max_{\theta} \prod_{i=1}^N p(\mathbf{x}^{(i)}, \theta) \\
 &= \arg \max_{\theta} \sum_{i=1}^N \log p(\mathbf{x}^{(i)}, \theta) \\
 &= \arg \max_{\theta} \mathbb{E}_{x \sim \hat{p}} \log p(\mathbf{x}, \theta),
 \end{aligned} \tag{8.2}$$

where \hat{p} is the empirical distribution (i.e. some training dataset) and $p(\theta)$ is the model parameterized by neural networks with parameters θ which we want to optimize.

In the case of unsupervised learning we will be maximizing $p(\mathbf{x}, \boldsymbol{\theta})$ and in the case of supervised learning $p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})$.

Depending on the task at question the error function will be different. In the case of the point parameter estimation we will use some sort of deterministic measure of how close the estimated value is to the true value for the training data. Whereas in the case of Bayesian parameter estimation we use a measure like Kullback-Leibler divergence to make sure that the two distributions are close to each other.

8.2.3. Gradient descent

The search for the parameters of the network is typically done with the stochastic gradient descent (SGD). At the moment the extension such as Adam are widely used instead of the classic gradient descent algorithm. In real dataset we have to deal with the millions of training samples which is not possible to process at once. To overcome this problem mini-batched gradient descent is used, that means that for each batch of data the weights are updates and overall loss is calculated as following

$$\mathcal{L}_m(\boldsymbol{\theta}) = \frac{1}{m} \sum_{i=1}^m \mathcal{L}(\mathbf{x}^{(i)}, \boldsymbol{\theta}). \quad [8.3]$$




Figure 8.3: Illustration of gradient descent algorithm with varying learning rate.

Thereafter the gradient of the loss is calculated and the parameters are updates using this gradient multiplied by the learning rate η

$$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \eta \nabla_{\boldsymbol{\theta}} \mathcal{L}_m(\boldsymbol{\theta}). \quad [8.4]$$

For the learning rate it is common to not fix a particular value but use the schedule to decrease the value in the process of training (see Figure 8.3).

8.2.4. Backpropagation

The gradient of the loss has to be propagated to each parameter of the network. Backpropagation is implemented by chain rule. Here is demonstration of the backpropagation for a very simple case (see Figure 8.4).




Figure 8.4: Chain rule.

$$\begin{aligned}\frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_2} &= \frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \hat{y}} \cdot \frac{\partial \hat{y}}{\partial \theta_2} \\ \frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_1} &= \frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \hat{y}} \cdot \frac{\partial \hat{y}}{\partial z} \cdot \frac{\partial z}{\partial \theta_1}\end{aligned}\tag{8.5}$$

Fortunately, with the advances of modern frameworks, one does not need to implement the backpropagation by hand. With the pioneering `theano` framework, and the modern packages such as, for example, `pytorch` and `jax`, the user can rely on the already implemented backpropagation. This gave a big push to the field of ML. In this tutorial we will be using `pytorch` as a reference framework, it is very close in syntactics to `python` and very intuitive to use for people who are accustomed to `python`.

8.3. Point parameter estimation

The classical machine learning techniques, which perform ML classification and regression can be useful as detection and point parameter estimation techniques respectively for GW data analysis. Classification can tell us if the signal is present in the data or not. That can be germane for the ground based detectors, whereas for signal dominated space-based detectors solving this task is not as relevant. Regression can provide a point estimation of the signal parameters which is useful as a search technique for all type of instruments.

In this section we will review the methodology for this approach and an example of the network architecture. We will also focus on the quantitative ways to assess the performance of the method, as it is very important to understand the error of the measurement in the absence of the natural way to incorporate the error which is present in the Bayesian techniques.

Let us consider that we have a waveform together with a response of the detector for the gravitational wave source $\mathbf{h}(\boldsymbol{\lambda})$ parameterized by a set of the physical parameters $\boldsymbol{\lambda}$. At the same time we can simulate the measurement noise \mathbf{n} , which can be linearly added to the signal to form a sample of the simulated data $\mathbf{x} = \mathbf{h}(\boldsymbol{\lambda}) + \mathbf{n}$. Both for detection and classification we will rely on supervised learning algorithms with the input data \mathbf{X} constructed from the simulated data samples $\mathbf{X} = \mathbf{x}^{(i)}$ and output data label \mathbf{y} being either a set $\{0, 1\}^{(i)}$ for the case of detection or a set of physical parameters $\Lambda = \boldsymbol{\lambda}^{(i)}$ for the case of point estimation.

Detection with ML

Detection of the GWs can be implemented as classification. Here by detection we mean identifying if the signal is present in the data or not and do not care about the particular values of parameters. As was already mention this a typical supervised learning task, where the label \mathbf{y} for the training data set is $\mathbf{0}$ if data does not contain signal $\mathbf{x} = \mathbf{n}$ and $\mathbf{1}$ if it does $\mathbf{x} = \mathbf{h}(\boldsymbol{\lambda}) + \mathbf{n}$.

In the case of ML, classification is usually implemented as logistic regression. Any type of architecture can be chosen for the network, with a typical choice being CNN (here are examples for LVK (Gabbard *et al* 2018) and LISA (Zhang *et al* 2022)).

The important part in the design of such types of an algorithm is a choice of the last non-linearity (activation function) and the loss function. This activation function have to map the output to the probability distribution therefore it is commonly chosen to be logistic function or it's special case, – sigmoid function (see Figure 8.5).

To derive optimization function we start as usual with the maximum likelihood. In the case of supervised learning we are interested in the conditional probability $p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})$. It follows the Bernoulli distribution, that takes the value 1 with the probability p and value 0 with the probability $q = 1 - p$. In our case the probabilities are the outputs of the last activation function which is a sigmoid function. Therefore we can write the distribution as follows

$$\begin{aligned} p(y | \mathbf{x}, \boldsymbol{\theta}) &= \text{Ber} \left(y \mid \sigma \left(\boldsymbol{\theta}^\top \mathbf{x} \right) \right) \\ &= \sigma(\boldsymbol{\theta}^\top \mathbf{x})^y \cdot (1 - \sigma(\boldsymbol{\theta}^\top \mathbf{x}))^{1-y}. \end{aligned} \quad [8.6]$$

Which makes log-likelihood to be

$$\log \mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^N \mathbf{y}^{(i)} \log \left(f \left(\mathbf{x}^{(i)}; \boldsymbol{\theta} \right) \right) + \left(1 - \mathbf{y}^{(i)} \right) \log \left(1 - f \left(\mathbf{x}^{(i)}; \boldsymbol{\theta} \right) \right), [8.7]$$




Figure 8.5: Last layer nonlinear activation function for logistic regression.

where function f represents the neural network with the last nonlinearity being sigmoid function.

In practice we minimize negative log-likelihood which makes a loss function $\mathcal{J}(\theta) = -\log \mathcal{L}(\theta)$. This loss function is usually called binary cross-entropy loss in literature.

To evaluate the performance of the method one has to evaluate how many cases were correctly classified. The frequentist approach in this case given by the receiver operator characteristics (ROC) curve is typically employed in GW data analysis (see Figure ??), where we plot true alarm probability against false alarm probability.

Point parameter estimation with ML

Simple classifier which decides if the signal is present in the data or not seems to be not very useful. Luckily it can be very easily extended to be a point parameter estimator. To make it work we need to change the last activation function in the network and the loss function. Interestingly we can reuse the other parts of the network between the detection and parameter estimation tasks.

In the case of point parameter estimation for the labels we are going to take values of the physical parameters λ which correspond to the waveforms in the simulated data. To term the problem from classification to regression




Figure 8.6: Example representation on the ROC curve of the different type of classifiers to provide the understanding where the good classifier should be located.




Figure 8.7: Activation function for the last layer in regression problems.

we will set the last activation function to be a linear positive function. In ML terminology it is usually called ReLU (Rectifier Linear Unit) (see Figure 8.7).

Let us derive the loss function from the maximum likelihood principle like in the previous task. Again we want to estimate conditional probability $p(y | \mathbf{x}, \boldsymbol{\theta})$, where in this case $y = \boldsymbol{\lambda}$ is a set of the physical parameters. Here the likelihood

derivation is based on the assumption that the noise has a Gaussian distribution. Therefore we can write for the log-likelihood the following expression

$$\begin{aligned}\log \mathcal{L}(\boldsymbol{\theta}) &= \sum_{i=1}^n \log p(\mathbf{y}^{(i)} | \mathbf{x}^{(i)}, \boldsymbol{\theta}) \\ &= \sum_{i=1}^n \left(-\frac{1}{2} \log(2\pi\sigma^2) - \frac{(\mathbf{y}^{(i)} - f(\mathbf{x}^{(i)}, \boldsymbol{\theta}))^2}{2\sigma^2} \right),\end{aligned}\tag{8.8}$$

where f represents the neural network, \mathbf{y} is a set of physical parameters and σ the variance of the Gaussian noise.

$$\mathcal{J}(\boldsymbol{\theta}) = -\log \mathcal{L}(\boldsymbol{\theta}) \propto \sum_{i=1}^n (\mathbf{y}^{(i)} - f(\mathbf{x}^{(i)}, \boldsymbol{\theta}))^2\tag{8.9}$$

We can have a point estimate and we can also get access of the error of this estimate by approximating the confidence interval in the frequentist way. This means that we can make many samples for the test dataset (see Figure 8.8) and access the error for the slice that correspond to the value that we inferred.




Figure 8.8: Example of the representation of the point parameter estimation with the error estimate at the slice corresponding to the inferred parameter value.

Example of the point parameter estimation for one Galactic Binary can be found in the following repository (Korsakova 2025).

This error, however, will be biased due to the presence of the error in the estimation of the weights of the network. One way to overcome this problem

is to initialize the weights not with the random number but with the distribution and marginalize over this distributions for the estimates of the physical parameters. This approach is called in literature *Bayesian networks*. However, it still will not provide the full Bayesian posterior. In the next section we are going to investigate different Bayesian DL techniques.

8.4. Bayesian parameter estimation

Point parameter estimation can provide in the best case the maximum likelihood estimation, however, because the measurement is contaminated by noise we need to correctly account for the noise which can be done by formulating the task in terms of probability distributions and using the Bayesian equation (Sivia 2006)

$$p(\lambda|x) = \frac{p(x|\lambda)p(\lambda)}{p(x)}, \quad [8.10]$$

where $p(\lambda|x)$ is the posterior distribution of parameters λ , $p(x|\lambda)$ is the likelihood, $p(\lambda)$ is the prior on the parameters and $p(x)$ is the marginal likelihood or evidence. The problem is that the marginal likelihood $p(x) = \int p(x|\lambda)p(\lambda)d\lambda$ does not have the exact solution.

Marginal likelihood does not have an exact solution that is why we need to find an approximate solution. There are two possible paths that we can take: (1) approximate inference or (2) simplified the model.

We have access to the waveforms that correspond to the parameters sampled from the prior space. At the same time we can simulate many realistic noise realizations, that means that we have access to the joint probability. We will use this principle in all generative methods instead on writing out the likelihood for the Gaussian noise.

NK Paragraph about naming: SBI, likelihood free, generative ,... NK Somewhere, here or in Bayesian PE write in more detail about all possible approaches. Also about the other generative models

Approximate inference

One of the way to solve the problem with the intractability of the marginal likelihood is to use sampling techniques such Markov Chain Monte Carlo (MCMC) or Nested Sampling. These methods are commonly used in GW data analysis. These methods can be also accelerated by the machine learning techniques but conceptually they all rely on the sequential

calculation of the likelihood values. The problem are twofold. First of all the likelihood evaluation can be time consuming and event it is speed up to be very fast still the sequential nature of this approach will result in the considerable times for the parameter estimation.

The other way to do simplified inference is a Variational inference. In particular used in the context of the deep latent variable models. The distribution of the data in this case in parametrized The method from which GW data analysis can benefit both in terms of the discriminative as well as the generative models is called Variational Autoencoders (VAE) (Kingma and Welling 2013 ; Rezende *et al* 2014). In variational inference variational lower bound is used to approximate the evidence.

Here z is the latent variable distribution.

$$p(x) = \int p(x|z)p(z)dz \quad [8.11]$$

First we apply the identity trick, i.e. we are going to approximate the integral by the expectation under another distribution. Usually this distribution is chosen to be well know and simple.

$$p(x) = \int p(x, z) \frac{q_\phi(z|x)}{q_\phi(z|x)} dz = \quad [8.12]$$

$$\int \frac{p(x, z)}{q_\phi(z|x)} q_\phi(z|x) dz = \quad [8.13]$$

$$\mathbb{E}_{q_\phi(z|x)} \left[\frac{p(x, z)}{q_\phi(z|x)} \right] \quad [8.14]$$

Here we can apply the property from the convex analysis that for a concave function the function of an expectation is bounded by the expectation of a function $f(\mathbb{E}[x]) \geq \mathbb{E}[f(x)]$.

$$\log p(x) \geq \mathbb{E}_{q_\phi(z|x)} \left[\log \frac{p(x|z)p(z)}{q_\phi(z|x)} \right] = \quad [8.15]$$

$$\mathbb{E}_{q_\phi(z|x)} [\log p(x|z)] - \int q_\phi(z|x) \log \frac{p(z)}{q_\phi(z|x)} = \quad [8.16]$$

$$\mathbb{E}_{q_\phi(z|x)} [\log p(x|z)] - D_{KL}(q_\phi(z|x)||p(z)) \quad [8.17]$$

The two part of the objective function can be considered as *reconstruction* and *regularisation* terms that keep the balance between reconstruction of the data and the continuity of the latent space.

Another important aspect is a *reparameterization trick*. We cannot back-propagate the the gradient of the loss function through the rando numbers. Instead we can fix a distribution, fot example, Gaussian and make mean and standard deviation of this distribution part of the network training process. **NK plot for reparameterisation trick**

This principle is used in Variational Autoencoders **NK cite** which can be considered from two different perspectives. It can be used for both discriminative (notably data representation) and generative tasks.

Simplified models

Simplified models is when the model for the posterior distribution is not generic but are approximated by the simplified distribution. The most simple case will be to assume that the posterior distribution will be Gaussian.

The easiest way will be to approximate the posterior distribution with the Gaussian distribution. In this case we can learn the mean and the variance of this distribution.

8.4.1. Normalizing flows

Normalizing from is the name for the invertible transform with trackable Jacobian. This is another approach when we simplify the model for the distribution. In this case we define the distribution as the a map from the latent Normal distribution, which provides the name for the method.

To understand how this principle works, let us look at the simple one dimensional distribution, its probability distribution function, cumulative distribution function and their connection with each other.

$$f_X(x) = \frac{d}{dx} F_X(x) = \frac{d}{dx} F_Z(f^{-1}(x)) \quad [8.18]$$

$$= f_Z(f^{-1}(x)) \left| \frac{d}{dx} f^{-1}(x) \right| \quad [8.19]$$




Figure 8.9: Probability distribution function can be constructed by sampling a random uniformly variable and mapping it with the cumulative distribution function. **NK Increase font size**

This was a very simple one dimensional case. If we write it for a multidimensional distribution, we can use the change of variable equation which can map any arbitrary distribution to another one.

$$p(x) = q(f^{-1}(x)) |\det(J_{f^{-1}}(x))|, \quad [8.20]$$

here J is the Jacobian of the transformation. It is required because we want to preserve the probability during the transformation. **NK Picture explaining this**

This transform just have to fulfill several conditions. The transformation have to be a diffeomorphism, which means that it is invertible and at the same the transform f and its inverse f^{-1} are continuously differentiable.

8.4.1.1. Objective function

The objective function is constructed in the similar way as for the other problems by maximizing the marginal likelihood $p(y | x)$.

We optimize the parameters of the network such as

$$\arg \min_{\theta} \mathbb{E}_{\mathbf{x}} [-\log p_{\theta}(\mathbf{x})] = \mathbb{E}_{\mathbf{x}} \left[-\log p(f_{\theta}(\mathbf{x})) - \log \det \left| \frac{\partial f_{\theta}(\mathbf{x})}{\partial \mathbf{x}} \right| \right] \quad [8.21]$$

Maximising likelihood estimation is similar to minimising KL divergence.

$$J_f(\mathbf{z}) = \begin{bmatrix} \frac{\partial f_1}{\partial z_1} & \dots & \frac{\partial f_1}{\partial z_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_n}{\partial z_1} & \dots & \frac{\partial f_n}{\partial z_n} \end{bmatrix} \quad [8.22]$$

We combine several transforms to get a good fit for a distribution

$$\mathbf{z} = f(\mathbf{x}) = f_K \circ \dots \circ f_1(\mathbf{x}). \quad [8.23]$$




Figure 8.10: Flow is a combination of transforms NK Check names of the variables

Let us look at several algorithms. The simple and easy to understand methods is Real Non Volume Preserving flow (Real NVP).

The other example that is widely use in the GW data analysis is Neural Spline Flows (NSF).

In addition to that let us look at continuous normalising flows.

There many other proposed architectures.

There are other generative models which can be also used for Bayesian inference, such as, score matching, flow matching, diffusion models, GANs and many others, which we are not going to touch in this review.

8.4.1.2. Conditioning of the flow on the simulated data

Methods that we distributed are usually formulated in perms of the distribution $p(x)$.

It can be generalized to the conditional distributions $p(y | x)$.

Embedding. The simplest way to do the embedding is to do SVD.

8.4.1.3. Example with Galactic Binaries for LISA

The tutorial can be found at the following address.

8.5. Bibliography

- Christopher M. Bishop (2006). *Pattern Recognition and Machine Learning*. MIT Press. [En ligne]. Disponible à l'adresse : <https://www.microsoft.com/en-us/research/wp-content/uploads/2006/01/Bishop-Pattern-Recognition-and-Machine-Learning-2006.pdf>
- Goodfellow I., Bengio Y., Courville A. (2016). *Deep Learning*. MIT Press. [En ligne]. Disponible à l'adresse : <http://www.deeplearningbook.org>
- Rosenblatt, F. (1958). *The perceptron: A probabilistic model for information storage and organization in the brain*. *Psychological Review*, 65(6), 386-408. [En ligne]. Disponible à l'adresse : <https://doi.org/10.1037/h0042519>
- Richard S. Sutton and Andrew G. Barto (2015). *Reinforcement Learning: An Introduction*. MIT Press. [En ligne]. Disponible à l'adresse : <https://web.stanford.edu/class/psych209/Readings/SuttonBartoIPRLBook2ndEd.pdf>
- Lecun Y., Bottou L., Bengio Y. and Haffner P. (1998). *Gradient-based learning applied to document recognition*. *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278-2324. [En ligne]. Disponible à l'adresse : <https://hal.science/hal-03926082v1/document>
- He K., Zhang X., Ren S. and Sun J. (2016). *Deep Residual Learning for Image Recognition*. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. [En ligne]. Disponible à l'adresse : <https://arxiv.org/abs/1512.03385>
- Sepp Hochreiter and Jürgen Schmidhuber (1997). *Long Short-Term Memory*. *Neural Comput.* 9, 8, 1735-1780. Disponible à l'adresse : <https://doi.org/10.1162/neco.1997.9.8.1735>
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin (2017). *Attention is all you need*. In *Proceedings of NIPS'17*. Disponible à l'adresse : <https://arxiv.org/abs/1706.03762>
- Bruna, J., Zaremba, W., Szlam, A., and Lecun, Y. (2014). *Spectral networks and locally connected networks on graphs*. *ICLR2014*. Disponible à l'adresse : <https://arxiv.org/abs/1312.6203>
- Gabbard, Hunter and Williams, Michael and Hayes, Fergus and Messenger, Chris (y2018). *Matching Matched Filtering with Deep Networks for Gravitational-Wave Astronomy*. *PhysRevLett*, 120, 141103. [En ligne]. Disponible à l'adresse : <http://dx.doi.org/10.1103/PhysRevLett.120.141103>
- Zhang, Xue-Ting and Messenger, Chris and Korsakova, Natalia and Chan, Man Leong and Hu, Yi-Ming and Zhang, Jian-dong (2022). *Detecting gravitational waves from extreme mass ratio inspirals using convolutional neural networks*. *PhysRevD*, 105, 123027. Disponible à l'adresse :

<http://dx.doi.org/10.1103/PhysRevD.105.123027>

- Korsakova (2025). *Point parameter estimation for galactic binary*. Disponible à l'adresse : https://github.com/NataliaKor/tutorial/blob/main/point_PE.ipynb
- Sivia D. S., Skilling J. (2006). *Data Analysis: A Bayesian Tutorial*. Oxford University Press. Disponible à l'adresse : <https://academic.oup.com/book/53014>
- Kingma D. P., Welling M. (2013). *Auto-Encoding Variational Bayes*. Disponible à l'adresse : <https://arxiv.org/abs/1312.6114>
- Rezende, D. J., Mohamed, S. and Wierstra, D. (2014). *Stochastic Backpropagation and Approximate Inference in Deep Generative Models*. Proceedings of the 31st International Conference on Machine Learning, PMLR 32(2):1278-1286. Disponible à l'adresse : <http://proceedings.mlr.press/v32/rezende14.pdf>
- Papamakarios G., Nalisnick E., Rezende D.J., Mohamed S. and Lakshminarayanan B. (2021). *Normalizing Flows for Probabilistic Modeling and Inference*. Journal of Machine Learning Research 22 1-64 . Disponible à l'adresse : <https://jmlr.org/papers/volume22/19-1028/19-1028.pdf>

DRAFT

9

The astrophysics of gravitational wave sources

Natalie Webb

*IRAP, Université de Toulouse, CNRS, CNES, 9 avenue du Colonel Roche,
31028 Toulouse, France*

GRAVITATIONAL WAVES,
coordinated by Nicola Tamanini (L2IT/CNRS). © ISTE Editions 2025.

9.1. Introduction

Observational constraints on our understanding of the Universe were largely limited to electromagnetic radiation until the first detection of a gravitational wave (GW) event on 14th September 2015 (GW150914), when *LIGO* observed the coalescence of two stellar mass black holes, $36_{-4}^{+5} M_{\odot}$ and $29_{-4}^{+4} M_{\odot}$, resulting in a black hole of $62_{-4}^{+4} M_{\odot}$ (Abbott et al 2016). Gravitational waves had been predicted almost one hundred years earlier, as a result of the theory of general relativity (Einstein 1916). Prior to this detection, radio observations of the binary pulsar system PSR B1913+16 (Hulse & Taylor 1975 ; Taylor & Weisberg 1982) had observationally confirmed the existence of gravitational waves by showing that the separation of the neutron stars in the binary was decreasing by exactly the amount predicted by the loss of energy from the system due to the emission of gravitational waves. This exceptional discovery was rewarded with the 1993 Nobel prize for Russell Hulse and Joseph Taylor. Nonetheless, other multi-messenger observations had provided observational constraints before 2015, including cosmic ray observations of astrophysical sources since 1952 (Galbraith & Jelley 1953) and the detection of neutrinos from astrophysical sources in 1965 (Reines et al. 1965 ; Achar et al. 1965). These messengers, joined the electromagnetic observations (radio, sub-millimetre, infra-red, optical, ultra-violet, X-ray and gamma-rays) to widen our understanding of the Universe.

9.2. Compact objects and related phenomena as gravitational wave sources

To radiate gravitational waves, sources must be asymmetric. They are often associated with compact objects (white dwarfs, neutron stars, black holes), which are formed at the end of the lives of stars. The nature of the compact object depends mostly on the initial mass of the star, but also on the mass loss over the lifetime of the star, its chemical composition, its rotation and whether it is part of a multiple system (binary, triple, etc, e.g. Meynet 1991 ; Belczynski et al. 2002). More massive compact objects also exist, namely massive black holes, though whether these are formed from stellar mass black holes or in another way, is still unclear.

The name *compact object* is derived from the star's compactness, which is proportional to the ratio of the compact object's mass (M) and radius (R). The compactness (\mathcal{C}) is defined as

$$\mathcal{C} = \frac{GM}{Rc^2}, \quad [9.1]$$

where G is the universal gravitational constant, and c is the speed of light in a vacuum. By definition, the maximum value for compactness is unity. A black hole is defined to have $\mathcal{C} = 1$, though formally, this is for a maximally rotating black hole, see Sec. 9.5.4. The compactness is not the same as density, where for a spherical compact object, the density is proportional to M/R^3 . Massive black holes can therefore have very low densities, even though their compactness is unity. Compact objects are generally symmetrical, although it has been proposed that neutron stars may have slight asymmetry in certain cases and could be a source of gravitational waves, for instance, during accretion (e.g. Andersson et al. 1999). As such, most compact objects do not radiate gravitational waves. However, two unequal mass compact objects in a tight binary will give rise to gravitational radiation.

Other phenomena are capable of producing gravitational waves. These include asymmetric supernova explosions or quantum fluctuations in the early Universe. Tidal disruption events (TDEs, when a star approaches sufficiently close to a massive black hole that the tidal forces will overcome the star's self-gravity, causing it to become disrupted and resulting in approximately half of the matter being accreted onto the black hole, Hills 1975) are not necessarily two compact objects, but the deformation of the star can cause it to become compact enough to emit gravitational waves (Kobayashi et al. 2004).

This chapter will focus largely on gravitational wave emission stemming from phenomena associated with compact objects, as the majority of the detections to date have involved these systems. The chapter concerning cosmology with gravitational waves discusses many of the other sources of gravitational waves.

9.3. White dwarfs

The first compact object to be found was the white dwarf. William Herschel detected the first white dwarf, 40 Eridani B in 1783 (Herschel 1785). Adams (1915) showed that the companion to Sirius was similar in mass to the Sun, but its luminosity was significantly lower, similar to that of 40 Eridani B (Adams 1914). These objects were clearly different from the majority of the observed stars. The fairly flat optical spectrum ($\sim 4000\text{-}7000\text{\AA}$) and low luminosity, indicating a small radius, led Luyten (1922) to refer to these stars as white dwarfs. White dwarfs are the most common endpoints of stellar evolution. Today, we know of almost a hundred thousand, many in the Solar neighbourhood (e.g. El-Badry et al. 2021).

9.3.1. Formation of white dwarfs

White dwarfs form during a planetary nebula at the end of the lives of stars with initial masses of up to $\sim 8 M_{\odot}$ (Maeder and Meynet 1989 ; Meynet 1991). Stars up to $10.6 M_{\odot}$ can also form white dwarfs, depending on the initial stellar metallicity (e.g. Woosley and Heger 2015). Stars undergo hydrogen fusion in their core over the majority of their life. This period of their lives is referred to as the main sequence, see Fig. 9.1. Fusion provides a force that equals the gravitational force due to the mass of the star, so it is in *hydrostatic equilibrium*. Hydrogen fusion produces helium. When insufficient hydrogen remains in the core, fusion ceases and the star leaves the main sequence with a helium-flash. Subsequently, for the lowest mass stars, there is insufficient pressure and temperature to allow the helium to undergo fusion and the core collapses under gravity until it is stabilised by another force resulting from the electron degeneracy pressure, see Section 9.3.1.1. The outer parts of the star are blown away in the planetary nebula. These resulting white dwarfs are either predominantly hydrogen or helium in nature. Stars with masses above $\sim 1.85 M_{\odot}$ (Meynet 1991) are able to sustain stable helium burning and enter the red giant phase, see Fig. 9.1. Again, once sufficient helium is exhausted, fusion stops, the core collapses and produces a carbon-oxygen white dwarf. Further cycles can also produce oxygen-neon white dwarfs.

9.3.1.1. Electron degeneracy

For white dwarfs, the force opposing further gravitational collapse is the electron degeneracy. The pressure (P_e) that the electrons (e^-) can exert is a function of their density (ρ_e), i.e. $P_e = f(\rho_e)$. This pressure is calculated by considering the number of free-electron states in a volume (V), with length (L), so that $V = L^3$. Replicating this cube to fill all the space, the wave vectors of the free-electron quantum states can only take discrete values. If the electron wave function is $\psi \propto e^{(ik \cdot x)}$, for the three dimensions, x, y, z , $k = (k_x, k_y, k_z)$. Replicating the cube gives:

$$k_x = n_x \frac{2\pi}{L} \quad \text{where} \quad n_x = 1, 2, \dots \quad [9.2]$$

Allowed states are separated by $2\pi/L$. The density of the states in k space (N) is

$$dN = g \frac{L^3}{(2\pi)^3} d^3k \quad (d^3k \equiv dk_x dk_y dk_z), \quad [9.3]$$

where g is the degeneracy factor for spin (spin is either up (\uparrow) or down (\downarrow)). Using the de Broglie relationship for the electron's momentum (p) and its wave




Figure 9.1: Hertzsprung Russell Diagram showing the stellar luminosity as a function of temperature. The main sequence (hydrogen burning stage), giant (and supergiant) phases are shown, as well as the white dwarf cooling sequence. Image credit : ESO.

vector: $p = \hbar k$, where \hbar is the Planck constant divided by 2π , the density of states in k space can be converted to momentum space:

$$dN = g \frac{L^3}{(2\pi\hbar)^3} d^3p. \quad [9.4]$$

The number density of particles per unit volume (n) with momentum states in d^3p is

$$dn = g \frac{1}{(2\pi\hbar)^3} f(p) d^3p, \quad [9.5]$$

where $f(p)$ is the mode occupation number, i.e. the number of particles in the cube with a particular wave function. For bosons (e.g. photons), $f(p)$ is unrestricted, but fermions (e.g. electrons with spin angular momentum $\hbar/2$) must obey the Pauli exclusion principle. This states

$$f(p) \leq 1. \quad [9.6]$$

The electron gas is therefore different from a classical one. The Pauli exclusion principle limits how dense the electron gas can be. Particle momenta are assumed to obey a Maxwellian distribution, where each velocity (v) component has a Gaussian distribution with standard deviation:

$$\Psi(v) = \frac{1}{(2\pi\sigma^2)^{3/2}} e^{(-v^2/2\sigma^2)} d^3v. \quad [9.7]$$

The velocity dispersion (σ) is related to the temperature (T) by energy equipartition: $\sigma^2 = k_B T/m$ (k_B is the Boltzmann constant and m the particle mass). Rewriting Equation 9.7 in terms of particle number density in a given range of momentum space, by multiplying by the total number density of particles and using $p = mv$, gives,

$$dn = \frac{n}{(2\pi mk_B T)^{3/2}} e^{(-p^2/2mk_B T)} d^3p. \quad [9.8]$$

Comparing Equations 9.8 and 9.5 implies a critical density (n_{crit}) where the classical law would violate the Pauli exclusion principle (Equation 9.6) for $p = 0$.

$$n_{crit} = \frac{g}{(2\pi)^{3/2}} \frac{(mk_B T)^{3/2}}{\hbar^3}. \quad [9.9]$$

The gas is in the classical regime at high temperatures and fixed density, but quantum effects become important as $T \rightarrow 0$. Integrating Equation 9.5 over all momentum states gives the total number density of particles:

$$n = \frac{g}{(2\pi\hbar)^3} \int f(p) d^3p. \quad [9.10]$$

As $T \rightarrow 0$, the states are occupied only up to the Fermi momentum (p_F). Hence:

$$n = \frac{g}{(2\pi\hbar)^3} \int_0^{p_F} d^3p = \frac{g}{(2\pi\hbar)^3} \frac{4\pi}{3} p_F^3. \quad [9.11]$$

The Fermi momentum is then related to the particle density as in Equation 9.12,

$$p_F = 2\pi\hbar \left(\frac{3}{4\pi g} \right)^{1/3} n^{1/3}. \quad [9.12]$$

As the gas density increases, the Fermi momentum also increases. Additional particles are forced to fill higher momentum states as the lower states are fully occupied. The pressure from the electrons can be determined using the fact that pressure is also a flux of momentum and therefore, the flux of electrons in the x -direction is just the number of electrons crossing a unit area per unit time, or $n_e v_x$. The pressure is then $\sim P \sim p_x n_e v_x$. Using Equation 9.5 the pressure in the x -direction (which by isotropy must be equal to the pressure in any direction) is

$$P = P_x = \frac{g}{(2\pi\hbar)^3} \int p_x v_x f(p) d^3p. \quad [9.13]$$

Using spherical polar coordinates in momentum space

$$\begin{aligned} \int p_x v_x dp_x dp_y dp_z &= \frac{1}{3} \int (p_x v_x + p_y v_y + p_z v_z) dp_x dp_y dp_z \\ &= \frac{1}{3} \int p \cdot v \cdot 4\pi p^2 dp, \end{aligned} \quad [9.14]$$

so

$$P = \frac{g}{3} \frac{1}{(2\pi\hbar)^3} \int_0^\infty p v f(p) 4\pi p^2 dp. \quad [9.15]$$

As $T \rightarrow 0$ (Eq. 9.11), electron speeds are non-relativistic so that $pv = p^2/m_e$, and

$$P_e = \frac{4\pi g}{3(2\pi\hbar)^3} \int_0^{p_F} \frac{p^2}{m_e} p^2 dp = \frac{g}{30\pi^2 \hbar^3 m_e} p_F^5. \quad [9.16]$$

Rewriting the pressure as a function of density $\rho_e = m_e n_e$ using Equation 9.12

$$p_F = \left(\frac{6\pi\hbar\rho_e}{gm_e} \right)^{1/3}, \quad [9.17]$$

and substituting into Equation 9.16 gives the pressure for non-relativistic electrons,

$$P_e = \frac{g}{30\pi^2 \hbar^3} \left(\frac{6\pi^2 \hbar^3}{g} \right)^{5/3} \rho^{5/3} m_e^{-8/3} = K_1 \rho_e^{5/3} \quad [9.18]$$

$$\text{where } K_1 = \frac{\pi^2 \hbar^2}{5m_e^{8/3}} \left(\frac{6}{g\pi} \right)^{2/3}.$$

However, as the densities increase (Eq. 9.12), the Fermi momentum can reach relativistic values. Some particles are forced into momentum states where the velocities approach the speed of light (c). Replacing v with c gives the relativistic expression

$$P_e = \frac{4\pi g}{3(2\pi\hbar)^3} \int_0^{p_F} pc p^2 dp = \frac{g c}{24\pi^2 \hbar^3} p_F^4, \quad [9.19]$$

or

$$P_e = \frac{gc}{24\pi^2 \hbar^3} \left(\frac{6\pi \hbar^3 \rho_e}{gm_e} \right)^{4/3} = K_2 \rho_e^{4/3} \quad [9.20]$$

$$\text{where } K_2 = \frac{\pi \hbar c}{4m_e^{4/3}} \left(\frac{6}{g\pi} \right)^{1/3}.$$

Equations 9.18 and 9.20 show that in a degenerate gas, the pressure only depends on the density, and is independent of temperature. In a partially degenerate gas, there is a residual temperature dependence. Only the electrons were considered in determining the pressure. In Equations 9.18 and 9.20, the mass of the particle is in the denominator. As the proton is 1836 times more massive than the electron, its contribution to the pressure is negligible.

9.3.2. White dwarf masses and radii

9.3.2.1. The Chandrasekhar mass

If the white dwarf is supported by electron degeneracy, there is a maximum possible pressure as there is a limit to the density of the electrons. Consequently, there must be a maximum mass for the white dwarf, which can be deduced. To estimate the energy density (U) of the degenerate gas, in a similar way as was done for the pressure, it is necessary to integrate over momentum space and include a term, $\epsilon(p)$, to denote the energy per mode:

$$U = \frac{g}{(2\pi\hbar)^3} \int_0^\infty \epsilon(p) f(p) 4\pi p^2 dp. \quad [9.21]$$

As $T \rightarrow 0$, $f(p) = 1$, but otherwise $f(p) = 0$. In the relativistic case, $\epsilon(p) = pc$. Using this in Equation 9.21 and integrating gives

$$U_e = \frac{g}{(2\pi\hbar)^3} \frac{1}{4} 4\pi c p_F^4, \quad [9.22]$$

and using Equation 9.12 to express the electron energy density (U_e) in terms of electron number density we find

$$U_e = \frac{3}{4} \left(\frac{6\pi^2}{g} \right)^{1/3} \hbar c n_e^{4/3}. \quad [9.23]$$

The non-relativistic case where $\epsilon(p) = p^2/2m_e$ gives

$$U_e = \frac{3\hbar^2}{10m_e} \left(\frac{6\pi^2}{g} \right)^{2/3} n_e^{5/3}. \quad [9.24]$$

The total kinetic energy (E_K) of the degenerate electrons is proportional to their energy density multiplied by the volume. In the relativistic case, $E_K \propto U_e V \propto n_e^{4/3} V \propto M^{4/3}/R$, where M is the mass and R is the radius of the white dwarf. The gravitational energy (E_p) is proportional to M^2/R . The total energy (E_{tot}) is the sum of the two terms: $E_{tot} = (AM^{4/3} - BM^2)/R$, where A and B are constants. For a set of masses, the total energy is positive. The total energy is reduced as the radius increases (the star expands) until the electrons reach the mildly relativistic regime, and the star becomes a stable white dwarf. But if the mass increases, the total energy is negative and the radius diminishes until gravitational collapse becomes inevitable. This is believed to be the mechanism responsible for producing Type Ia supernovae.

To find the critical mass at which the electron degeneracy pressure can no longer support the white dwarf, the constants A and B must be determined (Chandrasekhar 1931). These are related to the molecular weight per electron of the white dwarf matter and the density profile. The critical mass, first determined by Subrahmanyan Chandrasekhar (Chandrasekhar 1931), is named after him, the *Chandrasekhar mass*, and was determined to be $1.4 M_\odot$ (Ostriker 1966). By adding mass to the approximately spherical white dwarf which is supported by electron degeneracy pressure, the white dwarf radius decreases to sufficiently increase the density to support the additional mass. The radius of a solar mass white dwarf is approximately the radius of the Earth (~ 6000 km). Less massive white dwarfs have larger radii, for example, a $0.7 M_\odot$ white dwarf

has a radius of ~ 7500 km. Chandrasekhar mass white dwarfs have smaller radii of ~ 1200 km. The compact nature implies that the gravitational force holding the white dwarf together (GMm/R^2), where m is the mass of a particle on the surface of the white dwarf, is very high. They can then rotate extremely rapidly without flying apart due to the centripetal force (mv^2/R).

9.3.3. White dwarf structure and matter

At its formation, the white dwarf temperature is high. The initial temperature can be estimated using the contraction of the thermally unsupported stellar core down to the radius at which the contraction is halted by degeneracy pressure. According to the virial theorem, this is reached just before the final point of equilibrium, when the thermal energy ($E_{th} = \frac{3}{2}Nk_B T$) equals half of the potential energy (GM^2/R). For a pure helium, the number of nuclei in the core is $M/4m_p$, and the number of electrons is $M/2m_p$, so

$$E_{th} = \frac{3}{2} \frac{M}{m_p} \left(\frac{1}{2} + \frac{1}{4} \right) k_B T = \frac{9}{8} \frac{M}{m_p} k_B T. \quad [9.25]$$

With the virial theorem,

$$k_B T \sim \frac{4}{9} \frac{GMm_p}{R}. \quad [9.26]$$

For a solar mass white dwarf, the initial temperature is $\sim 1 \times 10^9$ K. This high temperature implies that the emission peaks in the X-ray domain. When the stellar core becomes an exposed white dwarf, its radiation ionises the gas blown off during the asymptotic giant branch phase, creating the planetary nebula. The white dwarf will then begin to cool. Neglecting the matter around the white dwarf and assuming a uniform temperature throughout, the maximal cooling rate can be estimated. The radiative energy loss is determined by equating the blackbody luminosity (L), given by the Stefan-Boltzmann law, to the rate of change of thermal energy (Equation 9.25)

$$L = 4\pi R^2 \sigma T^4 \sim \frac{dE_{th}}{dt} = \frac{3}{8} \frac{Mk_B}{m_p} \frac{dT}{dt} \quad [9.27]$$

where σ is the Stefan-Boltzmann constant. Separating the variables T and t and integrating, the cooling time to a temperature T is

$$\tau_{cool} \sim \frac{3}{8} \frac{Mk_B}{m_p} \frac{1}{4\pi R^2 \sigma} \frac{1}{3} \frac{1}{T^3} \sim 3 \times 10^9 \text{ yr} \left(\frac{T}{10^3 \text{ K}} \right)^{-3} \quad [9.28]$$

This estimate is too rapid as the outer, non-degenerate envelope acts as an insulator, and the inner regions crystallise, modifying the cooling with time. Understanding how white dwarfs cool can help estimate the age of the Universe, as the cooling times are comparable to that age.

The *Montreal White Dwarf Database* (Dufour et al. 2017) contains data for ~ 144800 spectroscopically confirmed white dwarfs published in the literature. The average temperature for white dwarfs in this database is ~ 7500 K, but with a large range, ~ 3000 - 180000 K, see Fig. 9.1. The sample mean of the mass is $0.57 M_{\odot}$, but spanning the range from very low $\sim 0.01 M_{\odot}$ to just above the Chandrasekhar mass $\sim 1.49 M_{\odot}$. The average surface gravity is $\log(g) \sim 8$, but spanning the range from 3.1-10. The magnetic fields span the range 1×10^5 - 9×10^8 G. These values are very similar to the almost complete sample of white dwarfs detected out to 100 pc with *Gaia* (Jiménez-Esteban et al. 2018). The average white dwarf radius is $\sim 0.012 R_{\odot}$ (8346 km) in this survey, with a range 0.002-0.02 R_{\odot} (1391-13910 km), but with a possible bimodal distribution peaking at 0.01 R_{\odot} (6955 km) and 0.0125 R_{\odot} (8694 km), which remains to be explained.

9.3.4. Observing and understanding white dwarfs

As mentioned in Sec. 9.3.3, understanding the cooling of white dwarfs can help constrain the age of the Universe. They are also very useful for constraining the initial mass function of the stars, mostly in our galaxy, but also neighbouring galaxies, as they are tracers of their progenitors (e.g. Quirk and Tinsley 1973). This is also important for comprehending stellar evolution. Determining the structure and matter in white dwarfs is also necessary for understanding type Ia supernovae. These occur when a white dwarf accretes matter and reaches the Chandrasekhar mass, above which electron degeneracy can no longer support the mass of the star, causing it to collapse further, see Sec. 9.4.1. Before this can occur, carbon burning can take place. This process creates an instability that can become a runaway in which burning proceeds very rapidly, disrupting it completely within seconds. Depending on how the burning proceeds, a type Ia supernova can occur, (e.g. Goobar and Leibundgut 2011). A standard luminosity (L) can be determined for the supernova, depending on the supernova model, but which is highly energetic ($\sim 10^{46}$ J). They can therefore be detected out to cosmological distances. Measuring the flux (f) observed from a type Ia supernova, and given the standard luminosity, the distance can be determined through $f = L/(4\pi d^2)$. The distances can then be used to understand the local expansion rate of the Universe, as well as its structure. It is therefore important to understand the type Ia supernova process, whether it is due to accretion onto the white dwarf or through the coalescence of two white

dwarfs (e.g. Iben and Tutukov 1985), and how the burning proceeds (Goobar and Leibundgut 2011), among other factors.

Should they be due to merging white dwarfs, the European Space Agency's *Laser Interferometer Space Antenna* (LISA, Amaro-Seoane et al. 2017), to be launched in the 2030s, that will detect gravitational waves in the frequency range from tens of micro-hertz to 1 Hz, will constrain the potential rates (Korol et al. 2024). The most prevalent sources that are expected to be detected with LISA will be Galactic double white dwarfs, which are about a hundred thousand years before merger (e.g. Lamberts et al. 2019). Using cosmological simulations of our galaxy along with binary population synthesis models, Lamberts et al. (2019) suggest that as many as 12000 double white dwarfs could be detected with LISA, although this estimate is still uncertain as it strongly depends on the population model and the instrument sensitivity. These will be dominated by Helium-Helium systems, Helium-Carbon/Oxygen systems and Carbon/Oxygen-Carbon/Oxygen systems (Lamberts et al. 2019). It is important to locate as many of these double white dwarf systems as possible in advance of the launch of LISA as this will help include them in modelling the background and therefore, improve the signal to noise of other sources. 1400 double white dwarf binaries within ~ 1 kpc of the Sun have been found in the Gaia catalogue. However, these have separations of a few astronomical units to 1 pc (El-Badry et al. 2021), and will not be strong gravitational wave emitters at the time of LISA. Munday et al. (2024) recently discovered 34 double-lined double white dwarf systems by following up 117 Gaia white dwarfs with the 4.2 m William Herschel Telescope, which have periods between a few hours and a few days. Further spectroscopy of other Gaia candidates is planned and could reveal closer binaries. The catalogue of cataclysmic binaries (Ritter and Kolb 2003) contains 1429 white dwarf binaries. Of these, 346 binaries have periods below 1.5 hours, with the shortest, HW Cnc, having a period of just 5.4 minutes, and are likely to be strong LISA sources. It is clear that there is a lot more work to do to identify double white dwarf systems using electromagnetic observations before the launch of LISA, to find verification binaries with which to validate detector performance and better understand the interactions between the two compact objects, to help place constraints on tides. All of these systems will help provide further insight into the evolution of binary systems.

9.4. Neutron stars

Neutron stars were first proposed in 1934 (Baade and Zwicky 1934), just two years after the neutron was discovered (Chadwick 1932). However, it wasn't until the rocket flight of *American Science and Engineering* (AS&E), initially designed to observe the Moon, that the Galactic X-ray source, Scorpius X-1, was discovered (Giacconi et al. 1962), although, it was five years later that

its accreting neutron star nature was elucidated (Shklovskii 1967). During the same year, periodic radio signals (1.337 s) were detected by Jocelyn Bell during a search for scintillation from radio quasars with the *Mullard Radio Astronomy Observatory*. These pulsations were found to be coming from a neutron star (Hewish et al. 1968), called a *pulsar* (the contraction of 'pulsating star'). The detection of PSR B1919+21, resulted in a Nobel prize for Antony Hewish. It was shared with Martin Ryle for his role in developing the telescope and methods used. Today, almost 4000 neutron stars are known (Manchester et al. 2005).

9.4.1. Formation of neutron stars

Stars with initial masses of approximately $>9 M_{\odot}$ (Heger et al. 2003) may finish their lives in type II supernova explosions, leaving behind a neutron star. The nature of the compact object depends on the metallicity of the progenitor star, which dictates the mass loss and therefore influences the final masses of the helium core and hydrogen envelope (Heger et al. 2003). Nonetheless, some stars up to $\sim 10.6 M_{\odot}$ can form white dwarfs (Section 9.3.1) if they lose their envelopes, due to higher metallicities which imply higher mass loss and thus lower mass compact objects.

Given that there is a maximum mass that can be supported by electron degeneracy, the Chandrasekhar mass (Section 9.3.2.1), exceeding this mass causes the core to shrink further and the density to increase. At high densities, $\sim 10^7 \text{ g cm}^{-3}$, the highest energy electrons have energies greater than the difference between the masses of the neutrons (m_n) and the protons (m_p), $m_n c^2 - m_p c^2 = 1.294 \text{ MeV}$. This means that it becomes energetically favourable for the protons to capture electrons via the weak-interaction i.e. for inverse β -decay to occur (Equation 9.29)



where p is the proton, n is the neutron and ν_e is the electron-neutrino. The inverse β -decay produces more neutrons than electrons and protons, as the existing electrons in the neutron star are degenerate, filling all the available states and halting the creation of new electrons via the β -decay of the neutrons. The kinetic energy of the ν_e removes energy from the star allowing it to cool.

9.4.1.1. The neutron drip and baryonic degeneracy

At densities higher than $\sim 3 \times 10^{11} \text{ g cm}^{-3}$, the nuclear neutron to proton ratio becomes so high that the neutrons are forced out of the nucleus, which is called the *neutron drip*. The degenerate free neutrons can then oppose the

force due to gravity and stop the collapse of the star in the same way as the degenerate electrons halt the collapse of the white dwarf (see Section 9.3.1.1). If the chemical potential of the electrons exceeds the rest mass of muons (μ_e), $m_\mu \sim 205.7 \text{ MeV } c^{-2}$, the electrons can disintegrate into muons, producing further electron-neutrinos and anti-muon neutrinos ($\bar{\nu}_m$) in the general β -decay (Equation 9.30)

$$e^- \rightarrow \mu_e + \nu_e + \bar{\nu}_m \quad [9.30]$$

The escaping neutrinos contribute to the cooling and the collapse of the proto-neutron star. If the interior neutron star density reaches two or three times the nuclear density ($\rho_0 \sim 2.3 \times 10^{14} \text{ g cm}^{-3}$), the strong force between nucleons (N) can lead to the appearance of new particles, such as the hyperon (baryons with three quarks with at least one strange quark) lambda (L),

$$N + N \rightarrow N + L + K \quad [9.31]$$

If the particle does not undergo a phase change to a condensate (considering that they are bosons), kaons (K) can decay to form leptons and photons. The photons and neutrinos leave the star, further reducing the temperature, accelerating the collapse and increasing the baryon degeneracy. Once all the energy levels are filled (as a result of the Pauli exclusion principle, see Eq. 9.6), the hyperons are trapped in a *Fermi sea* (i.e. all the states are occupied) and become stable to decay (*Pauli blocking*), in the same way as the neutrons could no longer undergo β -decay, see Section 9.4.1. The conservation of strangeness, which is initially zero, is violated because the negative strangeness of kaons is lost through their decay (Prakash et al. 1995). This process can continue until the whole star becomes degenerate and no further strong interaction occurs. The phase transition can then produce a plasma of free quarks, where matter is in the fundamental state (i.e. has minimal energy).

9.4.2. Neutron star structure and matter

Neutron stars have a very thin atmosphere, with a depth of only $\sim 0.1\text{-}10$ cm, and densities between $10^{-2}\text{-}10^2 \text{ g cm}^{-3}$ (Zavlin and Pavlov 2002). These atmospheres are subjected to strong gravitational stratification, so the lightest element is always the uppermost layer. Neutron star stratification occurs in only a few seconds (Romani 1987). Even with as little as 1 g cm^{-2} of hydrogen, the radiation emitted is indistinguishable from that emitted by a purely hydrogen atmosphere (Zavlin and Pavlov 2002). The observed atmosphere is therefore often hydrogen, unless the neutron star is accreting from a star with a different




Figure 9.2: A summary of the neutron star structure.

composition atmosphere. However, the atmosphere can be modified by the strong magnetic field of the neutron star, changing the opacity.

Beneath the atmosphere is the solid outer crust, where the density is $\sim 10^6$ - 10^{11} g cm $^{-3}$. The crust is formed from a body-centered cubic lattice of iron (see e.g. Fig. 9.2) enveloped in a degenerate electron gas (Weber 2005). For densities $> 10^7$ g cm $^{-3}$, the nuclei become more neutron rich through electron capture, see Section 9.4.1. The thickness of the crust is ~ 1 km (Haensel 2001). Between the outer and the inner crust, the density reaches $3-4 \times 10^{11}$ g cm $^{-3}$. The neutrons have been forced out of the nucleus to form the degenerate matter opposing the force due to gravity. The number of free neutrons rises with density. The limit of the inner crust is where matter reaches nuclear density, $\sim 2.3 \times 10^{14}$ g cm $^{-3}$. At this point, some theories predict non-spherical nuclei, such as slabs or cylinders, which are often referred to as *pasta phases*, and the lengthening or flattening of the nuclei is known as the *spaghetti phase* and the *lasagne phase* (see e.g. Fig. 9.2) respectively (e.g. Chamel and Haensel 2008).

Below the crust is the neutron star core. This can be separated into the outer and the inner core, each one several kilometres deep. The core is responsible for the majority of the mass of the neutron star, as it fills most of the neutron star volume and the central densities can reach more than ten times the density of nuclear matter (Weber 2005). As the density increases, new exotic

particles (see Section 9.4.1.1) can appear, including strange quarks which may result in strange quark stars, a configuration of matter more stable than the most stable atomic nucleus, ^{62}Ni (Weber 2005). The neutrons in the core can become superfluid, and protons at the top of the *Fermi sea* are superconducting. Rotation causes the neutrons to form vortices threaded through the core into the crust, and proton flux tubes weave through the vortices, carrying the magnetic field. The number of vortices is inversely proportional to the rotation period, thus increasing in number at high rotation rates (Ruderman 1972).

9.4.3. The equation of state

A non-rotating (for simplicity) neutron star can be considered in an asymptotically flat, spherically symmetric, static space-time frame $(\mathcal{M}, g_{\alpha\beta})$. The coordinate system $x^\alpha = (t, r, \theta, \phi)$ can be used to write the metric coefficients,

$$g_{\alpha\beta} dx^\alpha dx^\beta = -N^2 c^2 dt^2 + A^2 dr^2 + r^2(d\theta^2 + \sin^2\theta d\phi^2) \quad [9.32]$$

where N and A are functions of r , determined using the Einstein equation, and $g_{\alpha\beta}$ is the tensor metric. The Einstein equation relates the space-time curvature to the mass-energy content represented by the stress-energy tensor, $T_{\alpha\beta}$. The neutron star interior can be considered to be an ideal liquid, where the effects of viscosity and thermal conductivity are negligible, so

$$T_{\alpha\beta} = (\epsilon + P)u^\alpha u^\beta + P g_{\alpha\beta} \quad [9.33]$$

where P is the pressure, ϵ the total energy per unit volume of a fluid measured in its reference frame and u^α is the quadri-velocity of the fluid. The different components of the Einstein equation can then be rewritten as a system of first order differential equations called the Tolman-Oppenheimer-Volkoff (TOV) equations, after the authors that first developed the system (Tolman 1939 ; Oppenheimer and Volkoff 1939). The TOV equations are given in Equations 11.11-9.36.

$$\frac{dm}{dr} = 4\pi r^2 \rho(r) \quad [9.34]$$

$$\frac{dP}{dr} = -\frac{d\Phi}{dr} \left(\rho(r) + \frac{P(r)}{c^2} \right) \quad [9.35]$$

$$\frac{d\Phi}{dr} = \frac{Gm(r)}{r^2} \left(1 - \frac{2Gm(r)}{rc^2} \right)^{-1} \left(1 + \frac{4\pi P(r)r^3}{m(r)c^2} \right) \quad [9.36]$$

where $\rho = \epsilon/c^2$, $m(r)$ and $\Phi(r)$ are general relativistic terms for the mass content in a sphere of radius r and related to the metric coefficients $N(r)$ and $A(r)$ by

$$N(r) = e^{\Phi(r)/c^2} \quad A(r) = \left(1 - \frac{2Gm(r)}{rc^2}\right)^{-0.5}. \quad [9.37]$$

This system is completed by the equation of state, linking the pressure (P) to the density (ρ) at zero temperature, $P = P(\rho)$, to be able to resolve the equations. The limits used to resolve them are

$$\rho(0) = \rho_c \quad m(0) = 0 \quad \Phi(r) = \frac{c^2}{2} \ln \left(1 - \frac{2GM}{Rc^2}\right) \quad [9.38]$$

where R is the radius of the star and corresponds to $P(r = R) = 0$, and the gravitational mass M , which is the mass of a body determined by its response to gravity, is given by $M = m(R)$. For any equation of state and central density, there is a corresponding unique mass and radius, so that varying the central density gives rise to a family of mass-radius points, see Figure 9.3. The TOV equations are non-linear making them difficult to resolve, so this is often done numerically. However, a polytropic equation of state, of the type $P = \kappa \rho^\Gamma$, where P is the pressure, κ a constant and Γ the polytropic index, is often used as the solutions are amongst the simplest to derive (e.g. Lattimer and Prakash 2016). Different types of matter and potentials can be considered to determine the different equations of state (e.g. Lattimer and Prakash 2001, and references therein). The equation of state of neutron stars is still unclear. To constrain it observationally, constraints on the mass and the radius of different neutron stars are required. Identifying the neutron star equation of state would complete our understanding of the extremes of the Quantum Chromodynamics phase diagram, which can not currently be explored using first principle calculations due to numerical challenges. Heavy ion collision experiments explore the high temperature and lower density parts of the phase diagram, and neutron stars are the unique laboratories for the study of strong and weak force physics in cold, ultra-dense matter.

9.4.4. Neutron star masses and radii

9.4.4.1. Mass

Resolving the TOV equations for a specific central density, ρ_{max} , gives a maximal mass

$$\left. \frac{d\rho_c}{dM} \right|_{\rho_c = \rho_{max}} < 0. \quad [9.39]$$




Figure 9.3: A mass radius diagram showing different families of neutron star equations of state (black, green, red, dark orange, purple and blue solid lines). SQM denotes Strange Quark Matter models, PAL describes equations of state including neutrons and protons in a schematic potential, GM indicates equations of state including neutrons, protons and hyperons using a field theoretical approach, etc. The regions in the top left are excluded due to causality, see Section 9.4.4.1, the condition that the central pressure remains finite, $P_c < 1$, or $R > (9/4)GM/c^2$ or general relativity, $R > 2GM/c^2$, i.e. that the star's surface does not lie within its own event horizon. Orange contours indicate fixed values of R_∞ . Adapted from Lattimer and Prakash (2016) (<https://doi.org/10.1016/j.physrep.2015.12.005>)

At higher masses, the baryon degeneracy is no longer able to oppose gravity, and any radial perturbation would cause the neutron star to collapse to a black hole, see Section 9.5. The maximal mass depends on the equation of state. Analytical solutions of the TOV equations, supposing an ideal neutron gas, lead to maximal masses of $\sim 0.7 M_\odot$ (Oppenheimer and Volkoff 1939). This is lower than the observed masses of neutron stars, e.g. Lattimer and Prakash (2016). However, the strong force between the nucleons also plays a role in supporting the neutron stars. Introducing this gives a much more realistic maximal masses (Cameron 1959). Most contemporary neutron star equations of state have maximal masses of 1.5–2.8 M_\odot (Lattimer and Prakash 2016, and see Figure 9.3). However, masses as high as 3.2 M_\odot can also be predicted using Equations 11.11–9.36, and assuming that the matter is less stiff than that

of non-interacting degenerate neutron matter up to a fiducial mass (energy) density $(\rho_f)\epsilon_f$, chosen to be $(\rho_f) = \epsilon_f/c^2 = 4.6 \times 10^{14} \text{ g cm}^{-3}$, and limited by causality¹ at higher densities. Neutron stars have been observed with masses as low as $\sim 1 M_\odot$, e.g. one of the neutron stars in PSR J0453+1559 has an observed mass of $1.174 \pm 0.004 M_\odot$ (1σ error, Martinez et al. 2015). Neutron stars with masses as high as $1.908 \pm 0.016 M_\odot$ (1σ error, PSR J1614-2230 Arzoumanian 2018) have been measured. The relativistic Shapiro delay measurements of PSR J0740+6620 have revealed an even higher mass of $2.08^{+0.07}_{-0.07} M_\odot$ (1σ error, Fonseca et al. 2021).

9.4.4.2. Radius

Masses of 1.0 - $2.8 M_\odot$ imply radii between ~ 8 - 15 km for common equations of state (e.g. Lattimer and Prakash 2016, and Figure 9.3). Observationally, radii are much harder to constrain than mass. Measuring the quiescent neutron star X-ray flux for sources where both the distance (d) and the temperature can be estimated reliably, enables the radiation radius (R_∞) to be determined, using the redshifted Kirchoff's law for a blackbody, $F_\infty = (R_\infty/d)^2 \sigma T_\infty^4$, where the quantities are redshifted to the Earth and $R_\infty = R_{NS}/\sqrt{1 - 2GM_{NS}/R_{NS}c^2}$. The energy reservoir giving rise to the emission is thought to be the heat deposited deep in the neutron star crust, which is re-radiated outwards during quiescence. The high temperature means that the emission is seen predominantly in the X-rays domain. The spectrum depends on the atmospheric composition, and the strength and the structure of the magnetic field. Recent constraints using several neutron stars and Bayesian methods constrain the radii to 11.5 - 13 km (90% confidence, Baillot d'Étivaux et al. 2020).

Radii (along with masses) can also be deduced through modelling the X-ray waveform of a neutron star with a fairly low magnetic field, where the emission from the polar caps is modulated as the star rotates and the polar caps rotate in and out of the line of sight. Millisecond pulsars have magnetic fields, $B \lesssim 10^9$ G, which are weak enough to limit a modification of the opacity of the neutron star atmosphere. Loose constraints have been obtained for long *NICER* exposures in conjunction with *XMM-Newton* data for the pulsar PSR J0030+0451, where the radii using different models span 11.5 and 14.5 km, although supposing a single model provides errors of only ~ 1 km (Vinciguerra et al. 2024).

1. Causality in this case is the fact that the speed of sound in dense matter must be less than the speed of light

9.4.5. Observing and understanding neutron stars

One of the current open questions concerning neutron stars is the nature of the neutron star equation of state, which is important not only for astrophysics, but also for nuclear physics. Whilst electromagnetic radiation has allowed some constraints to be made on the nature of the neutron star matter, see Section 9.4.4, gravitational waves will also play a crucial role in improving our knowledge of dense matter. Coalescing neutron stars have now been detected directly using gravitational wave detectors (Abbott et al. 2017). GW 170817 was detected by the LIGO and Virgo gravitational wave (GW) detectors on August 17th 2017. Modelling the waveform resulted in constraints on the tidal deformability of the neutron stars from the quadrupole moment in the space-time surrounding the neutron star merger. The constraints on the radii made from the gravitational wave measurements cluster around 12 km, similar to the electromagnetic constraints, see Sec. 9.4.4.2, where the radii (R) of both stars were constrained to be $R = 11.9 \pm 1.4$ km (error 90% confidence) using the requirement that the equation of state allows neutron stars with masses greater than $1.97 M_{\odot}$ (Abbott et al. 2018), as observed with electromagnetic observations, see Sec. 9.4.4.1. These values cluster around neutron stars with equations of state that are composed of normal nucleonic matter (neutrons and protons), but do not exclude softer equations of state that include strange quarks, nor do they exclude stiffer equations of state or those that include more exotic particles such as hadrons and kaons. Constraints on nuclear physics were also possible. Abbott et al. (2018) determined the pressure to be twice that of nuclear saturation density, $3.5^{+2.7}_{-1.7} \times 10^{34}$ dyn cm⁻² (error 90% confidence). Subsequent work has combined results from electromagnetic observations with the gravitational wave data, but so far the constraints have not been improved. Future observations of merging neutron stars, or maybe a neutron star and black hole, should improve constraints.

Continuous gravitational waves are expected to be produced by spinning neutron stars with an asymmetry around their rotation axis. So far, searches with the current ground-based detectors are yet to make a statistically significant detection of asymmetric neutrons (e.g. Abbott et al. 2022).

However, neutron stars that are observed as pulsars are also used as gravitational wave detectors to detect the lowest frequency gravitational waves coming from inspiralling supermassive black hole binaries, and signals from a cosmic string network or other physical processes occurring in the early Universe. This is done using ~ 100 pulsars with excellent timing precision, which are monitored for time delays in the arrival times of pulses on Earth. Correlated residuals point towards the passage of gravitational waves. The sensitivity of these *Pulsar Timing Arrays* (PTAs) is improving, and recent observations point towards the detection of the gravitational wave background or a supermassive black

hole binary system, or even something more exotic (e.g. EPTA Collaboration et al. 2023 ; Arzoumanian et al. 2023).

It is expected that double neutron stars will be observable with LISA. By following the gravitational radiation-driven evolution of double neutron stars generated from rapid population synthesis of massive binary stars, Lau et al. (2020) estimated that around 35 systems will be detectable with a signal-to-noise ratio above 8 over a four year LISA mission lifetime. These systems are mostly in our own galaxy ($\sim 94\%$), but others may be detected in nearby galaxies such as the Large Magellanic Cloud ($\sim 5\%$) or the Small Magellanic Cloud ($\sim 1\%$ Lau et al. 2020).

9.5. Black holes

The idea of black holes was first proposed around the same time the first white dwarfs were detected (Michell 1784 ; Laplace 1799). However, a proper description of these objects requires general relativity, which was not proposed until more than a century later (Einstein 1916). The term *black hole* itself was invented during comments to a talk given by John Wheeler at the *NASA Goddard Institute of Space Studies* in 1967, but the author of the term is unknown. Whilst general relativity predicts the existence of black holes, it doesn't prove that they exist, as the theory does not describe the astrophysical processes by which a black hole may form. It wasn't until the X-ray mission AS&E detected the source Cyg X-1 (Gursky et al. 1963), which was later suggested to be an accreting stellar mass black hole because of its high luminosity and rapid irregular variability (Pringle and Rees 1972), that there was any observational evidence for these objects. Today, more than 70 stellar mass black holes are known in our own galaxy (Corral-Santana et al. 2016), as well as many in other galaxies. At the same time as Galactic sources were being identified as stellar mass black holes, (Lynden-Bell and Rees 1971) suggested that the many of the known galaxies could have a supermassive black hole in their centres. Today, millions of galaxies are known in which supermassive black holes reside.

9.5.1. Formation of black holes

The formation mechanism of black holes depends on the mass (see Section 9.5.3). The formation of stellar mass black holes, $\sim 3-100 M_{\odot}$, is probably best understood. They can form at the end of the lives of massive stars or through the coalescence of two neutron stars. It is also possible that a black hole can form through accretion onto a neutron star. As we saw in Section 9.4.1, the nature of the compact object formed as a massive star comes to the end of its life is not only a function of the initial mass of the star. Stars with initial

masses $\gtrsim 25 M_{\odot}$ with fairly low metallicity can form a black hole if material from the supernova explosion falls back onto the central compact object. Stars with initial masses $\gtrsim 40 M_{\odot}$, again with fairly low metallicity will collapse directly into a black hole at the end of their lives. Stars of a similar mass but with higher metallicities, will nonetheless end their lives as a neutron star (Heger et al. 2003).

As we saw in Section 9.4.4.1, the neutron star has a maximum mass above which the baryon degeneracy is no longer able to oppose gravity. A compact object exceeding this maximum mass will collapse to form a black hole. Although the maximum neutron star mass remains unclear, as it depends on the still unconstrained neutron star equation of state, compact objects with masses above about $3 M_{\odot}$ are generally thought to be black holes. The gravitational collapse can be spherical or asymmetric, but in all cases, the equilibrium solutions depend on only three physical parameters, with the others being *washed out* in the collapse. These parameters are the mass (M), the electric charge (Q) and the angular momentum (J). This is known as the *no-hair theorem*, initially proposed by Israel (1967).

Given that an astrophysical black hole is fully characterised by just three parameters, there are only four exact solutions of Einstein's equations. The four solutions are (1) The Schwarzschild solution (Schwarzschild 1916) where the black hole has only mass i.e. it is static and spherically symmetric, (2) The Reissner-Nordström solution (Reissner 1916 ; Nordström 1918), where the black hole is static, spherically symmetric, and depends on mass and electric charge, (3) The Kerr solution (Kerr 1963), where the black hole is stationary, axisymmetric, and depends on mass and angular momentum and (4) The Kerr-Newman solution (Newman and Janis 1965), where the black hole is stationary and axisymmetric, and depends on all three parameters, mass, electric charge and angular momentum.

The black hole parameters are not arbitrary. Electric charge and angular momentum cannot exceed values that correspond to the disappearance of the event horizon, see Section 9.5.4. The relation

$$a^2 + Q^2 \leq M^2 \quad [9.40]$$

must be satisfied, where $a = Jc/GM$. In general, the dimensionless spin expression is used: $a^* = a/M$ so that $-1 \leq a^* \leq 1$. If $a^* > 0$, the black hole rotates in a prograde direction. If $a^* < 0$, the black hole rotation is retrograde. If $a^* = 0$, the black hole is static i.e. without rotation. If the condition in Equation 9.40 is violated, the event horizon disappears and the solution describes a naked singularity rather than a black hole. Naked singularities are not expected to exist as

astrophysical objects. In realistic astrophysical solutions, the black hole is not expected to have significant charge due to the extreme weakness of the gravitational interaction compared to electromagnetic interaction. If for example an astrophysical black hole formed with a positive charge, which is not isolated in empty space but surrounded by the interstellar medium full of charged particles, e.g. protons and electrons, it would attract electrons with a charge e^- and repel protons with charge e^+ . It would also attract protons of mass m_p thanks to its gravitational field. The repulsive electrostatic force on protons is larger than the gravitational attraction by a factor $eQ/m_p M \approx e/m_p \approx 10^{18}$. Therefore, the black hole would become neutral (without charge) almost immediately. Therefore, the Kerr solution can be used for any astrophysical black holes.

Black holes may also be formed through another mechanism proposed by Hawking (1971), but building on work by other authors at the time. In the very early Universe, there were large random fluctuations on all length scales. For some volumes, the gravitational energy would have exceeded the kinetic energy of expansion.

$$\frac{GM^2}{R} \simeq G\rho^2 R^5 \geq pR^3 \quad [9.41]$$

During the radiation era, $p \simeq \rho c^2$, so the Equation 9.41 would be the equivalent of $GM/c^2 \geq R$, where R is the size of the fluctuation. These regions would then not have continued to expand with the rest of the Universe, but would have collapsed. Black holes formed through this mechanism are called *primordial black holes*.

9.5.2. Black hole masses and radii

9.5.3. The mass

Black holes therefore have a wide range of masses. As in Section 9.5.1, stellar mass black holes have masses between ~ 3 and $100 M_\odot$. The lowest mass depends on the physical nature of the matter in the compact object, as described in Section 9.4.4.1. The maximal mass is uncertain and is an estimate of the most massive black hole that could form from the earliest stars. However, the pair-instability supernova process predicts maximum masses that could be significantly below these values, but the actual mass is unclear. Accurate mass constraints of stellar mass black holes are difficult to determine. One way is through optical observations of a companion star to a black hole in an X-ray binary system. Radial velocity measurements, along with constraints on

the inclination through modelling of the spectral lines or through ellipsoidal modulation measurements, in the case of the low mass X-ray binaries, can provide reasonable mass estimates. The masses of galactic black holes found in binaries are fairly low. GRO J0422+32 has a minimal mass estimate of $2.22 M_{\odot}$ (Webb et al. 2000), although the mass is likely to be somewhat higher. Cygnus X-1 has the highest accurately measured stellar mass black hole in an X-ray binary in our galaxy, with a mass of $21.2 \pm 2.2 M_{\odot}$ (Miller-Jones et al. 2021).

The first detections of gravitational waves from coalescing black holes have revealed more massive black holes than those detected with electromagnetic radiation. The masses of the black holes in the first three runs span $5.9 \pm_{1.3}^{4.4} M_{\odot}$ and $85 \pm_{14}^{21} M_{\odot}$ (Abbott et al. 2023). How these higher mass black holes came to exist is still unclear, but they may have been formed in the early Universe when the metallicity was very low, allowing more massive stars to form and then to create more massive black holes (e.g. Abbott et al 2016).

More massive than stellar mass black holes are intermediate mass black holes with masses in the range $\sim 10^2$ - $10^5 M_{\odot}$. Their origin is unclear, as very massive stars are needed to form such massive black holes. Such massive stars are not observed today, where the most massive observed have initial stellar masses of less than about $300 M_{\odot}$. The earliest stars (population III stars) with very low metallicities may have had much higher masses, allowing them to end their lives as intermediate mass black holes. Alternatively, intermediate mass black holes may originate from the earliest, low-metallicity dust clouds with low angular momentum. These could have collapsed directly to form massive black holes. Alternatively, accretion of matter onto stellar mass black holes, or mergers or even a combination of both mechanisms could lie at the origin of intermediate mass black holes.

Few good intermediate mass black hole candidates are known. 2XMM J011028.1-460421, more commonly known as Hyper Luminous X-ray source 1 (HLX-1, Farrell et al. 2009 ; Godet et al. 2009 ; Webb et al. 2010) has been widely studied, and the observational evidence points towards its intermediate mass nature. The mass has been determined to be $\sim 10^4 M_{\odot}$ (Godet et al. 2012 ; Webb et al. 2012). It is thought to be a partial tidal disruption event, accreting when it tidally strips a companion star at periastron in a highly elliptical orbit (Lasota et al. 2011 ; Godet et al. 2014 ; Webb et al. 2014). Intermediate mass black holes are expected to reside in the centres of lower mass galaxies. They are usually faint and hard to detect. However, if they tidally disrupt a passing star, the system can become brighter by several decades in X-ray luminosity, as well as at other wavelengths, making them easier to find. Examples of such systems include the black hole in the centre of the inactive galaxy IC 4765-f01-1504, which

showed a tidal disruption event in 2006. The black hole mass is estimated to be $6 \times 10^4 - 4 \times 10^6 M_{\odot}$ (Lin et al. 2011). Another similar event occurred around a black hole in a dwarf galaxy orbiting 6dFGS gJ215022.2-055059. The mass was estimated to be $5.3 \times 10^4 < M_{BH} < 1.2 \times 10^5 M_{\odot}$ (Lin et al. 2018). Another intermediate mass black hole of $\sim 5 \times 10^4 M_{\odot}$ has been proposed in the galaxy RGG 118 (Baldassare et al. 2015).

Supermassive black holes, with masses in the range $10^6 - 10^{10} M_{\odot}$ must have formed in a different way to the intermediate mass black holes, as even continuously accreting at the maximal rate (the Eddington limit), it is difficult to reach masses as high as $\sim 10^9 M_{\odot}$ by $z \sim 7.1$ (Mortlock et al. 2011) or of $8 \times 10^8 M_{\odot}$ for the black hole found at $z=7.54$ (0.69 Gyr, Bañados et al. 2018). These black holes may have been formed from intermediate mass black holes that either merged and/or accreted matter to create supermassive black holes. Supermassive black holes are found in the centres of all massive galaxies and have mass ranges from the maximal mass of intermediate mass black holes up to tens of billions of solar masses. Masses can be estimated using radial velocity measurements of orbiting stars for nearby black holes such as Sgr A* in the centre of the Milky Way. Similarly, the dispersion of spectral lines give insight into the stellar dynamics in the bulge of more distant galaxies and thus the black hole mass. Alternatively, reverberation mapping in fairly close galaxies using the broad line region (BLR) can be used to determine the mass by measuring the radius of this region. This is done by assuming the time lag between variations in the continuum and line emission are due to the light travel time. Assuming that the velocity width, W , of the permitted emission lines is due to Keplerian motion in the black hole potential, the mass can be estimated to be $\sim W^2 R_{BLR}/G$. Other methods for estimating the black hole mass include scaling relations between the mass of the black hole and for example the bulge mass, the spheroid luminosity or the pitch angle of the spiral arms.

Primordial black holes can be formed with different masses, without a specific upper mass limit due to their non-stellar origin. Using the relationship between density and time $G\rho \approx t^{-2}$ in an Einstein-DeSitter model of the early Universe, the maximum mass of a collapsing fluctuation is related to the cosmic time by $M(\text{grams}) \approx 10^{38} t(\text{seconds})$. Therefore, at the Planck time, $t \approx 10^{-43}$ s, very small black holes would have formed with the Planck mass $\approx 10^{-5}$ g. At $t \approx 10^{-4}$ s, black holes of $\sim 1 M_{\odot}$ could have formed. Later still at $t \approx 100$ s, supermassive black holes may have formed. However, no good observational evidence for primordial black holes has yet been discovered.

9.5.4. The radius

Considering a static and spherically symmetric black hole of mass M (a Schwarzschild black hole), the square of the invariant interval, a solution to Einstein's equation is given by

$$ds^2 = - \left(1 - \frac{2GM}{Rc^2} \right) c^2 dt^2 + \frac{dr^2}{1 - 2GM/Rc^2} + r^2 (d\theta^2 + \sin^2 \theta d\phi^2) [9.42]$$

In the case when $2GM/Rc^2 = 1$, the proper time, $d\tau$, term disappears, and the proper distance, dl , term diverges. The critical distance, R is then $\frac{2GM}{c^2}$ and is known as the *Schwarzschild radius*. At this radius, a clock would appear to stop for an observer, so nothing can be seen falling through the Schwarzschild radius (or *event horizon*) from anywhere outside.

Whilst it is not rigorously correct, it may be instructive to consider Newtonian approximation to understand the *Schwarzschild radius*. The gravitational potential energy required for a mass m to leave a gravitating body of mass M is $\frac{GMm}{R}$. To determine the velocity, v , required to leave the gravitational potential, the kinetic energy for a mass m , $\frac{1}{2}mv^2$ can be equated to the gravitational potential, so that $\frac{GM}{R} = \frac{1}{2}v^2$, which gives the velocity, $v = \sqrt{\frac{2GM}{R}}$. Calculating the radius R at which light can escape, as nothing can travel faster than light in a vacuum, gives the *Schwarzschild radius*. This radius is inside most astrophysical bodies, but for the gravitationally collapsed black hole, the Schwarzschild radius is outside the true radius. This implies that inside R , nothing has a velocity high enough to escape the black hole, including light, which explains why we can not see a black hole. For a black hole of $3 M_{\odot}$, the *Schwarzschild radius* is almost 9 km. As $R \propto M$, then as the mass increases, so does the *Schwarzschild radius*. Therefore, a $100 M_{\odot}$ black hole has an almost 300 km *Schwarzschild radius*. As no information can escape the *Schwarzschild radius*, it is impossible to determine the radius of the black hole and therefore we refer to the Schwarzschild radius, or the *event horizon*.

As indicated in Section 9.5.1, there are other astrophysical solutions to Einstein's equation, namely the *Kerr* solution for a spinning black hole. This solution includes the effect of the angular momentum, which is zero for the *Schwarzschild* solution. The radius is then defined as $r = (1 + \sqrt{1 - a^2})GM/c^2$, so for an extreme rotating black hole, $a = 1$ and $r = GM/c^2$, i.e. half the size of the *Schwarzschild* radius. For a static black hole, we again find the *Schwarzschild* radius. These radii describe the limit to the radius of the black hole for light. However, a second radius is defined for particles. This is known as the innermost stable circular orbit, or ISCO. In Newtonian gravity, all circular

orbits are stable such that if an orbit were perturbed slightly, it would become elliptical. In general relativity, there is an innermost stable circular orbit, where at and within, a particle on a circular orbit is no longer stable. If this particle becomes perturbed, it will spiral towards the black hole. For a *Schwarzschild* black hole, $R_{ISCO} = 6GM/c^2$, and for a *Kerr* black hole it depends non-trivially on the value of the black hole spin and the inclination and direction of the matter around it. For maximum spin and a prograde, equatorial orbit, the ISCO has the same radius as the event horizon, $R_{ISCO} = GM/c^2$.

Constraints have been made on black hole radii, allowing the spin to be estimated. Many methods used for determining the radius rely on measuring the inner accretion flow, usually in X-rays, as this hot part of the disc is closest to the black hole. The black body emission depends on the inner disc radius. Alternatively, the shape of the X-ray iron emission line (Fe K_α) is deformed by the effects of strong gravity, so it can be used to determine the inner radius of the accretion disc and constrain the radius of the black hole.

Observations of the stellar mass black hole in the X-ray binary system GRS 1915+105 indicate that it appears to be rotating at maximal spin, i.e. a Kerr black hole. M33 X-7 also has a well constrained high spin of $a_* = 0.84 \pm 0.05$ (McClintock et al. 2011), whereas the X-ray binaries A0620-00 and XTE J1550-564 show very low spins of ~ 0.1 and ~ 0.3 respectively, and are considered to be Schwarzschild like black holes (McClintock et al. 2011). Gravitational wave events have given some of the best spin constraints (typically $\sim 20\%$ of the spin value, Abbott et al. 2023), although these are only for the black hole produced through the merger. Measurements of the spin of the initial black holes are not yet possible. Spin values of the final black hole are almost all in the mid to high range, $0.4 < a_* < 0.89$ (Abbott et al. 2023). Spin constraints on supermassive black holes have all been made using the Fe K_α line as described above. Spins of the 19 supermassive black holes with such measurements lie in a very similar spin range as those measured for the stellar mass black holes formed through coalescence, with values ranging from 0.5-0.995 (no errors considered), but error values are much higher than those determined from gravitational waves, typically 10-100%.

Sub millimetre *Very Long Baseline Interferometry* (VLBI) observations taken with five telescopes across almost one hemisphere of the Earth (known as the *Event Horizon Telescope*), have measured the *shadow* of the supermassive black hole in the galaxy M 87 (Event Horizon Telescope Collaboration et al. 2019) and for the supermassive black hole in the Milky Way, Sgr A* (Event Horizon Telescope Collaboration et al. 2024). This shadow is cast when photons escape from unstable circular orbits near the event horizon of the black hole, creating the observational signature which is a bright ring of emission surrounding the shadow, called a *crescent*. In the

case of M 87, the diameter of the shadow as seen by a distant observer was predicted to be $\sim 9.6\text{-}10.4 GM/c^2$, where the range arises from the combination of possible spin values and the inclination angle to the observer. A crescent radius of $42 \pm 3 \mu\text{as}$ was measured, giving a mass of $M = 6.5 \pm 0.20|_{stat} \pm 0.7|_{sys} \times 10^9 M_\odot$ for the black hole and the angular size of one gravitational radius $\theta_g = \frac{GM}{c^2 D} = 3.8 \pm 0.4 \mu\text{as}$, where D was measured to be $16.8 \pm_{0.7}^{0.8}$ Mpc. The quasi-circular shape of the shadow is consistent with the presence of a central Kerr black hole (Event Horizon Telescope Collaboration et al. 2019).

9.5.5. Black hole structure

As detailed in Section 9.5.1, there are four solutions to Einstein's equations. These give rise to slightly different structures.

9.5.5.1. Static black holes

Einstein's equations describe two types of static black holes, i.e. with no angular momentum (or spin). Only one of these describes an astrophysical black hole, the *Schwarzschild* solution.




Figure 9.4: Schematic of a rotating (Kerr) black hole in two dimensions.

9.5.5.2. Rotating black holes

Of the two solutions to Einstein's equations describing rotating black holes, i.e. with angular momentum, only one describes an astrophysical black hole, the *Kerr* solution. The most important feature of this solution, compared to the *Schwarzschild* solution, are the $d\phi dt$ terms. These indicate a relation between time and azimuthal angle. The azimuthal angle of space around the black hole changes with time. This is called *frame dragging*, where spacetime is dragged in the direction of the black hole rotation. Frame dragging is also known as the Lense-Thirring effect or Lense-Thirring precession, named after the authors of

the paper describing the effect (Lense and Thirring 1918). This phenomenon is the relativistic correction to a gyroscope near an astrophysical body. Due to the Lense-Thirring effect, a particle with no angular velocity placed close to a rotating black hole would gain angular velocity as it nears the black hole. Matter close to the black hole (but outside the event horizon) will rotate in the same direction as the black hole. This implies a region called the *ergosphere*. This region touches the event horizon at the poles and has a radius equal to the Schwarzschild radius at the equator, see Figure 9.4. A consequence of frame dragging is that the ISCO radius increases for retrograde orbits of matter orbiting a rotating black holes, see Section 9.5.4. The accretion efficiency also decreases. This also implies that if gas spirals around a black hole on prograde orbits, the energy emitted, and hence the accretion efficiency increases with increasing spin.

9.5.6. Observing and understanding black holes

As indicated in Section 9.5.3 the black holes detected with gravitational wave detectors have revealed very massive black holes compared to those detected in our own galaxy. However, *Gaia* has been detecting a new population of black holes, thanks to astrometry, revealing binary populations due to the movement of one or both stellar components as they orbit their common centres of mass. Through follow-up spectroscopy to determine the radial velocity, the masses of the two components of the binary can be determined. Three black holes have been identified in this way, see e.g. Gaia Collaboration et al. (2024) and references therein. They are all in wide binaries, so that no accretion is occurring onto the black hole and therefore, can not be detected through X-ray observations. These black holes have masses of $9.62 \pm 0.18 M_{\odot}$ (El-Badry et al. 2023), $8.9 \pm 0.3 M_{\odot}$ (El-Badry et al. 2023) and $32.7 \pm 0.82 M_{\odot}$ (Gaia Collaboration et al. 2024) for black hole 1 (BH1), BH2 and BH3 respectively. This latter black hole falls firmly in the mass range of the black holes detected with gravitational waves. The companion star has low metallicity, indicating that the black hole progenitor may also have had low metallicity, as is expected for the formation of such a massive stellar mass black hole. Its existence indicates that other black holes of similar mass should exist in our galaxy. However, the mass of the companion is low $0.76 \pm 0.05 M_{\odot}$, so this system is dissimilar to the gravitational wave progenitors, but the system may have been formed through dynamical interactions, rather than from a primordial binary (Gaia Collaboration et al. 2024).

Whilst there has been weak evidence for intermediate mass black holes, gravitational wave detectors have identified the lowest mass intermediate mass black hole, GW 190521 (Abbott et al. 2020). This event may have originated from the coalescence of a $85 \pm_{14}^{21} M_{\odot}$ and a $66 \pm_{18}^{17} M_{\odot}$ black hole that merged to

create a $142^{+28}_{-16} M_{\odot}$ black hole, slightly above the lower limit for an intermediate mass black hole. Black holes with masses of 66 and 85 M_{\odot} are not expected, as stars with masses between ~ 65 and 135 M_{\odot} should undergo pair instability, when the production of free electrons and positrons in the collision between atomic nuclei and energetic gamma rays, temporarily reduces the internal radiation pressure. This leads to a partial collapse of the star, followed by an accelerated runaway thermonuclear explosion that leaves no compact remnant. More massive stars are thought to directly collapse to intermediate mass black holes. The formation of these massive black holes is therefore unclear and they may have originated through previous mergers. Since the launch of the *James Webb Space Telescope* in 2021, hundreds of *little red dots* have been discovered. They are a class of small, red, very compact galaxies (if they are dominated by star formation) at high redshift ($z > 5$) and their nature is unknown. They may represent a class of intermediate mass black hole, and it is possible that many of them are accreting at a high rate. Recent work has suggested that they may even be tidal disruption events of intermediate mass black holes (Bellovary 2025).

Recently, a feature has been associated with some tidal disruption events, quasi periodic eruptions (QPEs, Miniutti et al. 2019). They are seen in the X-ray lightcurve as the source progresses towards quiescence. The first was identified by Miniutti et al. (2019) who observed the X-ray count rate increase by up to two orders of magnitude over an hour, repeating every 9 h in the Seyfert 2 galaxy GSN 069. The eruptions show a transition over the burst, where the soft thermal emission heats up. A dozen other systems have also shown QPEs. Their nature is still unclear, but one interpretation is that QPEs are due to extreme mass-ratio inspirals (EMRIs), where a stellar mass object spirals towards the massive black hole as it loses energy through the emission of gravitational waves. More than one star could be orbiting the black hole, which may account for pairs of bursts. Alternatively, an inspiralling object could repeatedly impact an existing accretion disc, possibly created through a tidal disruption event, which would give rise to QPEs. If these sources are EMRIs, they may be detectable with *LISA* observations, depending on their orbital periods and eccentricities. Such detections would provide constraints on the massive black hole, as well as measure general-relativistic and Lense-Thirring precession. However, further sources need to be observed and further investigation carried out in order to validate (or refute) the EMRI theory.

9.6. Future observations

The coming decade or so will be rich with new observations helping to understand compact objects. Currently, the Space-based multi-band astronomical Variable Objects Monitor (SVOM) mission (Wei et al. 2016) is searching for

new gamma-ray bursts and other transients, and will help identify neutron star mergers (kilonovae). In conjunction with a ground-based gravitational wave detection, constraints on the neutron star equation of state will be possible, as for GW 170817. The Vera C. Rubin observatory (Ivezić et al. 2019) will start to observe the sky systematically and repeatedly in late 2025. Using the bands u, g, r, i, z and y , it will reach magnitudes over the 10 years as deep as 27.5 for the r band, and as deep as 25th magnitude in g for a single visit. The high cadence observations will allow thousands of new accreting white dwarfs to be discovered across the galaxy, and the frequent observations will allow orbital periods to be determined and outbursts to be studied. 2000 double degenerates are also expected to be discovered through the identification of periodicities. The study of their evolution will provide an insight into the type Ia supernovae and provide binaries that will be observable with *LISA*. *Rubin* will also detect merging massive black holes, identifying a population that will be detectable with *LISA* or even with the PTAs, as well as tidal disruption events. Some tidal disruptions should also be detectable with gravitational waves (e.g. Toscani et al. 2025), although they are less likely to be detected with *LISA*, and more probably with future deci-Hertz gravitational wave detectors. The majority of these sources can also be detected with the next generation X-ray mission from the *European Space Agency*, *NewAthena* (Cruise et al. 2025).

9.7. Bibliography

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. (2016). Observation of Gravitational Waves from a Binary Black Hole Merger. *Physics Review Letters*, 116, 061102.
- Abbott, B. P., and 199 colleagues (2017). GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Physical Review Letters* 119, 161101.
- Abbott, B. P., and 199 colleagues (2018). GW170817: Measurements of Neutron Star Radii and Equation of State. *Physical Review Letters* 121, 161101.
- Abbott, R., and 199 colleagues (2020). GW190521: A Binary Black Hole Merger with a Total Mass of $150\odot$. *Physical Review Letters* 125, 101102.
- Abbott, R., and 199 colleagues (2022). All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data. *Physical Review D* 106, 102008.
- Abbott, R., and 199 colleagues (2023). GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run. *Physical Review X* 13, 041039.

- Achar, C. V., Narasimham, V. S., Ramana Murthy, P. V., Creed, D. R., Pattison, J. B. M., Wolfendale, A. W. (1965). The intensity and angular distribution of cosmic rays far underground. *International Cosmic Ray Conference* 1, 989.
- Adams, W.S. (1914). An A-Type Star of Very Low Luminosity. *PASP*, 26, 198.
- Adams, W.S. (1915). The Spectrum of the Companion of Sirius. *PASP*, 27, 236.
- Amaro-Seoane, P., and 83 colleagues (2017). Laser Interferometer Space Antenna. *arXiv e-prints* arXiv:1702.00786.
- Andersson, N., Kokkotas, K., Schutz, B. F. (1999). Gravitational Radiation Limit on the Spin of Young Neutron Stars. *The Astrophysical Journal* 510, 846-853.
- Arzoumanian, Z., and 78 colleagues (2023). The NANOGrav 12.5 yr Data Set: Bayesian Limits on Gravitational Waves from Individual Supermassive Black Hole Binaries. *The Astrophysical Journal* 951, L28.
- Arzoumanian, Z. (2018). NASA's Neutron star Interior Composition Explorer (NICER): mission overview and initial results. *42nd COSPAR Scientific Assembly* 42, E1.10-4-18.
- Baade, W. and F. Zwicky (1934). Remarks on Super-Novae and Cosmic Rays. *Physical Review* 46, 76.
- Bailloy d'Étivaux, N., J. Margueron, S. Guillot, N. Webb, M. Catelan, and A. Reisenegger (2020). Constraining the equation of state of dense nuclear matter using thermal emission of neutron stars. *Journal of Physics Conference Series* 1667, 012001.
- Baldassare, V. F., A. E. Reines, E. Gallo, and J. E. Greene (2015). A $\sim 50,000 M_{\odot}$ Solar Mass Black Hole in the Nucleus of RGG 118. *The Astrophysical Journal* 809, L14.
- Bañados, E., and 17 colleagues (2018). An 800-million-solar-mass black hole in a significantly neutral Universe at a redshift of 7.5. *Nature* 553, 473.
- Belczynski, K., Bulik, T., Kluźniak, W. (2002). Population Synthesis of Neutron Stars, Strange (Quark) Stars, and Black Holes. *The Astrophysical Journal* 567, L63-L66.
- Bellovary, J. (2025). Little Red Dots are Tidal Disruption Events in Runaway-Collapsing Clusters. *arXiv e-prints* arXiv:2501.03309.
- Cameron, A. G. (1959). Neutron Star Models.. *The Astrophysical Journal* 130, 884.
- Chadwick, J. (1932). Possible Existence of a Neutron. *Nature*, 129, 312.
- Chamel, N. and P. Haensel (2008). Physics of Neutron Star Crusts. *Living Reviews in Relativity* 11, 10.

- Chandrasekhar, S. (1931). The Maximum Mass of Ideal White Dwarfs. *The Astrophysical Journal* 74, 81.
- Corral-Santana, J. M., J. Casares, T. Muñoz-Darias, F. E. Bauer, I. G. Martínez-Pais, and D. M. Russell (2016). BlackCAT: A catalogue of stellar-mass black holes in X-ray transients. *Astronomy and Astrophysics* 587, A61.
- Cruise, M., and 18 colleagues (2025). The NewAthena mission concept in the context of the next decade of X-ray astronomy. *Nature Astronomy* 9, 36.
- Dufour, P., S. Blouin, S. Coutu, M. Fortin-Archambault, C. Thibeault, P. Bergeron, and G. Fontaine (2017). The Montreal White Dwarf Database: A Tool for the Community. *20th European White Dwarf Workshop* 509, 3.
- Einstein, A. (1916). Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik* 354, 769-822.
- El-Badry, K., H.-W. Rix, and T. M. Heintz (2021). A million binaries from Gaia eDR3: sample selection and validation of Gaia parallax uncertainties. *Monthly Notices of the Royal Astronomical Society* 506, 2269.
- El-Badry, K., and 20 colleagues (2023). A Sun-like star orbiting a black hole. *Monthly Notices of the Royal Astronomical Society* 518, 1057.
- El-Badry, K., and 20 colleagues (2023). A red giant orbiting a black hole. *Monthly Notices of the Royal Astronomical Society* 521, 4323.
- EPTA Collaboration, and 99 colleagues (2023). The second data release from the European Pulsar Timing Array. III. Search for gravitational wave signals. *Astronomy and Astrophysics* 678, A50.
- Event Horizon Telescope Collaboration, and 199 colleagues (2019). First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole. *The Astrophysical Journal* 875, L1.
- Event Horizon Telescope Collaboration, and 199 colleagues (2024). The persistent shadow of the supermassive black hole of M 87. I. Observations, calibration, imaging, and analysis. *Astronomy and Astrophysics* 681, A79.
- Farrell, S. A., N. A. Webb, D. Barret, O. Godet, and J. M. Rodrigues (2009). An intermediate-mass black hole of over 500 solar masses in the galaxy ESO243-49. *Nature* 460, 73.
- Fonseca, E., and 44 colleagues (2021). Refined Mass and Geometric Measurements of the High-mass PSR J0740+6620. *The Astrophysical Journal* 915, L12.
- Gaia Collaboration, and 199 colleagues (2024). Discovery of a dormant 33 solar-mass black hole in pre-release Gaia astrometry. *Astronomy and Astrophysics* 686, L2.
- Galbraith, W., J. V. Jelley, J. V. (1953). Light Pulses from the Night Sky associated with Cosmic Rays. *Nature* 171, 349-350.

- Giacconi, R., H. Gursky, F. R. Paolini, and B. B. Rossi (1962). Evidence for x Rays From Sources Outside the Solar System. *Physical Review Letters* 9, 439.
- Godet, O., D. Barret, N. A. Webb, S. A. Farrell, and N. Gehrels (2009). First Evidence for Spectral State Transitions in the ESO 243-49 Hyperluminous X-Ray Source HLX-1. *The Astrophysical Journal* 705, L109.
- Godet, O., B. Plazolles, T. Kawaguchi, J.-P. Lasota, D. Barret, S. A. Farrell, V. Braito, M. Servillat, N. Webb, and N. Gehrels (2012). Investigating Slim Disk Solutions for HLX-1 in ESO 243-49. *The Astrophysical Journal* 752, 34.
- Godet, O., J. C. Lombardi, F. Antonini, D. Barret, N. A. Webb, J. Vingless, and M. Thomas (2014). Implications of the Delayed 2013 Outburst of ESO 243-49 HLX-1. *The Astrophysical Journal* 793, 105.
- Goobar, A. and B. Leibundgut (2011). Supernova Cosmology: Legacy and Future. *Annual Review of Nuclear and Particle Science* 61, 251.
- Gursky, H., R. Giacconi, F. R. Paolini, and B. B. Rossi (1963). Further Evidence for the Existence of Galactic x Rays. *Physical Review Letters* 11, 530.
- Hawking, S. (1971). Gravitationally collapsed objects of very low mass. *Monthly Notices of the Royal Astronomical Society* 152, 75.
- Heger, A., C. L. Fryer, S. E. Woosley, N. Langer, and D. H. Hartmann (2003). How Massive Single Stars End Their Life. *The Astrophysical Journal* 591, 288.
- Herschel, W. (1785). Catalogue of Double Stars. By William Herschel, Esq. F. R. S.. *Philosophical Transactions of the Royal Society of London Series I* 75, 40.
- Haensel, P. (2001). Neutron Star Crusts. *Physics of Neutron Star Interiors* 578, 127.
- Hewish, A., S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins (1968). Observation of a Rapidly Pulsating Radio Source. *Nature* 217, 709.
- Hills, J. G. (1975). Possible power source of Seyfert galaxies and QSOs. *Nature* 254, 295-298.
- Hulse, R. A., Taylor, J. H. (1975). Discovery of a pulsar in a binary system. *The Astrophysical Journal* 195, L51-L53.
- Iben, I. and A. V. Tutukov (1985). On the evolution of close binaries with components of initial mass between 3 M and 12 M.. *The Astrophysical Journal Supplement Series* 58, 661.
- Israel, W. (1967). Event Horizons in Static Vacuum Space-Times *Physical Review* 164, 1776.
- Ivezić, Ž., and 199 colleagues (2019). LSST: From Science Drivers to Reference Design and Anticipated Data Products. *The Astrophysical Journal* 873, 111.

- Jiménez-Esteban, F. M., S. Torres, A. Rebassa-Mansergas, G. Skorobogatov, E. Solano, C. Cantero, and C. Rodrigo (2018). A white dwarf catalogue from Gaia-DR2 and the Virtual Observatory. *Monthly Notices of the Royal Astronomical Society* 480, 4505.
- Kerr, R. P. (1963). Gravitational Field of a Spinning Mass as an Example of Algebraically Special Metrics. *Physical Review Letters* 11, 237.
- Kobayashi, S., Laguna, P., Phinney, E. S., Mészáros, P. (2004). Gravitational Waves and X-Ray Signals from Stellar Disruption by a Massive Black Hole. *The Astrophysical Journal* 615, 855-865.
- Korol, V., R. Buscicchio, R. Pakmor, J. Morán-Fraile, C. J. Moore, and S. E. de Mink (2024). Expected insights into Type Ia supernovae from LISA's gravitational wave observations. *Astronomy and Astrophysics* 691, A44.
- Lamberts, A., S. Blunt, T. B. Littenberg, S. Garrison-Kimmel, T. Kupfer, and R. E. Sanderson (2019). Predicting the LISA white dwarf binary population in the Milky Way with cosmological simulations. *Monthly Notices of the Royal Astronomical Society* 490, 5888.
- Laplace, P. S. (1799). Beweis des Satzes, dass die anziehende Kraft bey einem Weltkörper so groß seyn könne, dass das Licht davon nicht ausströmen kann. *Allgemeine Geographische Ephemeriden* 4, 1.
- Lasota, J.-P., T. Alexander, G. Dubus, D. Barret, S. A. Farrell, N. Gehrels, O. Godet, and N. A. Webb (2011). The Origin of Variability of the Intermediate-mass Black-hole ULX System HLX-1 in ESO 243-49. *The Astrophysical Journal* 735, 89.
- Lattimer, J. M. and M. Prakash (2001). Neutron Star Structure and the Equation of State. *The Astrophysical Journal* 550, 426.
- Lattimer, J. M. and M. Prakash (2016). The equation of state of hot, dense matter and neutron stars. *Physics Reports* 621, 127.
- Lau, M. Y. M., I. Mandel, A. Vigna-Gómez, C. J. Neijssel, S. Stevenson, and A. Sesana (2020). Detecting double neutron stars with LISA. *Monthly Notices of the Royal Astronomical Society* 492, 3061.
- Lense, J. and H. Thirring (1918). Über den Einfluß der Eigenrotation der Zentralkörper auf die Bewegung der Planeten und Monde nach der Einsteinschen Gravitationstheorie. *Physikalische Zeitschrift* 19, 156.
- Lin, D., E. R. Carrasco, D. Grupe, N. A. Webb, D. Barret, and S. A. Farrell (2011). Discovery of an Ultrasoft X-Ray Transient Source in the 2XMM Catalog: A Tidal Disruption Event Candidate. *The Astrophysical Journal* 738, 52.
- Lin, D., and 15 colleagues (2018). A luminous X-ray outburst from an intermediate-mass black hole in an off-centre star cluster. *Nature Astronomy* 2, 656.

- Luyten, W. J. (1922). Note on Some Faint Early Type Stars with Large Proper Motions. *Publications of the Astronomical Society of the Pacific* 34, 54.
- Lynden-Bell, D. and M. J. Rees (1971). On quasars, dust and the galactic centre. *Monthly Notices of the Royal Astronomical Society* 152, 461.
- Manchester, R. N., G. B. Hobbs, A. Teoh, and M. Hobbs (2005). VizieR Online Data Catalog: ATNF Pulsar Catalog (Manchester+, 2005). *VizieR Online Data Catalog* 7245, VII/245.
- Maeder, A. and G. Meynet (1989). Grids of evolutionary models from 0.85 to 120M : observational tests and the mass limits.. *Astronomy and Astrophysics* 210, 155.
- Martinez, J. G., K. Stovall, P. C. C. Freire, J. S. Deneva, F. A. Jenet, M. A. McLaughlin, M. Bagchi, S. D. Bates, and A. Ridolfi (2015). Pulsar J0453+1559: A Double Neutron Star System with a Large Mass Asymmetry. *The Astrophysical Journal* 812, 143.
- McClintock, J. E., R. Narayan, S. W. Davis, L. Gou, A. Kulkarni, J. A. Orosz, R. F. Penna, R. A. Remillard, and J. F. Steiner (2011). Measuring the spins of accreting black holes. *Classical and Quantum Gravity* 28, 114009.
- Meynet, G. (1991). The initial stellar masses for the formation of white dwarfs, neutron stars and black holes.. *Annales de Physique Colloque Supplement* 16, 459-479.
- Michell, J. (1784). On the Means of Discovering the Distance, Magnitude, &c. of the Fixed Stars, in Consequence of the Diminution of the Velocity of Their Light, in Case Such a Diminution Should be Found to Take Place in any of Them, and Such Other Data Should be Procured from Observations, as Would be Farther Necessary for That Purpose. By the Rev. John Michell, B. D. F. R. S. In a Letter to Henry Cavendish, Esq. F. R. S. and A. S.. *Philosophical Transactions of the Royal Society of London Series I* 74, 35.
- Miller-Jones, J. C. A., and 25 colleagues (2021). Cygnus X-1 contains a 21-solar mass black hole—Implications for massive star winds. *Science* 371, 1046.
- Miniutti, G., and 13 colleagues (2019). Nine-hour X-ray quasi-periodic eruptions from a low-mass black hole galactic nucleus. *Nature* 573, 381.
- Mortlock, D. J., and 16 colleagues (2011). A luminous quasar at a redshift of $z = 7.085$. *Nature* 474, 616.
- Munday, J., and 10 colleagues (2024). The DBL Survey I: discovery of 34 double-lined double white dwarf binaries. *Monthly Notices of the Royal Astronomical Society* 532, 2534.
- Newman, E. T. and A. I. Janis (1965). Note on the Kerr Spinning-Particle Metric. *Journal of Mathematical Physics* 6, 915.
- Nordström, G. (1918). On the Energy of the Gravitation field in Einstein's Theory. *Koninklijke Nederlandse Akademie van Wetenschappen Proceedings*

- Series B Physical Sciences* 20, 1238.
- Oppenheimer, J. R. and G. M. Volkoff (1939). On Massive Neutron Cores. *Physical Review* 55, 374.
- Ostriker, J. P. (1966). The Effect of Rotation on the Structure of White Dwarfs.. *The Astronomical Journal* 71, 394.
- Prakash, M., J. R. Cooke, and J. M. Lattimer (1995). Quark-hadron phase transition in proton-neutron stars. *Physical Review D* 52, 661.
- Pringle, J. E. and M. J. Rees (1972). Accretion Disc Models for Compact X-Ray Sources. *Astronomy and Astrophysics* 21, 1.
- Quirk, W. J., Tinsley, B. M. (1973). Star Formation and Evolution in Spiral Galaxies. *The Astrophysical Journal* 179, 69-84. doi:10.1086/151847
- Reines, R. et al. (1965). Evidence for high energy cosmic ray neutrino interactions. *International Cosmic Ray Conference* 1, 1051.
- Reissner, H. (1916). Über die Eigengravitation des elektrischen Feldes nach der Einsteinschen Theorie. *Annalen der Physik* 355, 106.
- Ritter, H. and U. Kolb (2003). Catalogue of cataclysmic binaries, low-mass X-ray binaries and related objects (Seventh edition). *Astronomy and Astrophysics* 404, 301.
- Romani, R. W. (1987). Model Atmospheres for Cooling Neutron Stars. *The Astrophysical Journal* 313, 718.
- Ruderman, M. (1972). Pulsars: Structure and Dynamics. *Annual Review of Astronomy and Astrophysics* 10, 427.
- Saffer, A., and 10 colleagues (2024). A Lower Mass Estimate for PSR J0348+0432 Based on CHIME/Pulsar Precision Timing. *arXiv e-prints* arXiv:2412.02850.
- Schwarzschild, K. (1916). Über das Gravitationsfeld eines Massenpunktes nach der Einsteinschen Theorie. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften* 189.
- Shklovskii, I. S. (1967). The Nature of the X-Ray Source Sco X-1.. *Astronomicheskii Zhurnal* 44, 930
- Taylor, J. H., Weisberg, J. M. (1982). A new test of general relativity - Gravitational radiation and the binary pulsar PSR 1913+16. *The Astrophysical Journal* 253, 908-920.
- Tolman, R. C. (1939). Static Solutions of Einstein's Field Equations for Spheres of Fluid. *Physical Review* 55, 364.
- Toscani, M., L. Broggi, A. Sesana, and E. M. Rossi (2025). Updated predictions for gravitational wave emission from TDEs for next generation observatories. *arXiv e-prints* arXiv:2505.22516.

- Vinciguerra, S., and 14 colleagues (2024). An Updated Mass-Radius Analysis of the 2017-2018 NICER Data Set of PSR J0030+0451. *The Astrophysical Journal* 961, 62.
- Weber, F. (2005). Strange quark matter and compact stars. *Progress in Particle and Nuclear Physics* 54, 193.
- Webb, N. A., T. Naylor, Z. Ioannou, P. A. Charles, and T. Shahbaz (2000). A TiO study of the black hole binary GRO J0422+32 in a very low state. *Monthly Notices of the Royal Astronomical Society* 317, 528.
- Webb, N. A., D. Barret, O. Godet, M. Servillat, S. A. Farrell, and S. R. Oates (2010). Chandra and Swift Follow-up Observations of the Intermediate-mass Black Hole in ESO 243-49. *The Astrophysical Journal* 712, L107.
- Webb, N., D. Cseh, E. Lenc, O. Godet, D. Barret, S. Corbel, S. Farrell, R. Fender, N. Gehrels, and I. Heywood (2012). Radio Detections During Two State Transitions of the Intermediate-Mass Black Hole HLX-1. *Science* 337, 554.
- Webb, N. A., O. Godet, K. Wiersema, J.-P. Lasota, D. Barret, S. A. Farrell, T. J. Maccarone, and M. Servillat (2014). Optical Variability of the Accretion Disk around the Intermediate-mass Black Hole ESO 243-49 HLX-1 during the 2012 Outburst. *The Astrophysical Journal* 780, L9.
- Wei, J., and 79 colleagues (2016). The Deep and Transient Universe in the SVOM Era: New Challenges and Opportunities - Scientific prospects of the SVOM mission. *arXiv e-prints* arXiv:1610.06892.
- Woosley, S. E. and A. Heger (2015). The Remarkable Deaths of 9-11 Solar Mass Stars. *The Astrophysical Journal* 810, 34.
- Zavlin, V. E. and G. G. Pavlov (2002). Modeling Neutron Star Atmospheres. *Neutron Stars, Pulsars, and Supernova Remnants* 263.

10

Gravitational Wave Cosmology: an introduction

Gregoire Pierra¹ and Simone Mastrogiovanni¹

¹*INFN, Sezione di Roma, I-00185 Roma, Italy*

10.1. Introduction

Gravitational-wave (GW) observations have opened a new era in cosmology by providing an independent way to study the cosmic expansion. Compact binary coalescences (CBCs) – the mergers of black holes (BHs) and neutron stars (NSs) – emit GWs that carry direct information about their luminosity distance. When combined with information about the redshift of the source, these events can serve as *standard sirens*, providing a new means to measure cosmological parameters such as the Hubble constant, the matter density, and the dark energy equation of state. This method is fundamentally different from traditional EM(EM) techniques, and therefore it is a valuable addition to the tools that we have to study the cosmos.

In this chapter, we develop the necessary tools to understand and implement GW-based cosmological measurements. We aim to provide a soft introduction to all the tools required for GW cosmology. We begin by summarizing some basic aspects of CBCs and GWs at cosmological distances in Sec. 10.2. In Sec. 10.3 we will be setting the core framework of hierarchical Bayesian inference (HBI), a statistical method used to infer population and cosmological properties of CBCs. This approach allows us to infer both the intrinsic properties of individual events and the underlying distributions governing the population, while carefully accounting for selection biases introduced by detector sensitivity. We will show how the merger rate of CBCs is formulated and how cosmological information naturally enters through the relationship between the detector-frame and source-frame parameters. Henceforth, this chapter proceeds with Secs. 10.6-10.5, in which we will delve into the details of GW cosmology for sources observed with an EM counterpart (Bright sirens) and without an EM counterpart (dark sirens). We will also discuss how it is possible to apply these methods to currently public gravitational-wave data products, while referencing external resources for a more involved reading on the current results.

With this chapter, our goal is to provide a comprehensive and accessible road map to GW cosmology, emphasizing the conceptual flow from the fundamental GW measurements to their applications in constraining the expansion history of the Universe.

10.2. Background for standard sirens cosmology

Direct measurements of the cosmic expansion exploit the observation of sources for which their distances and the recessional velocities can be measured. CBCs at cosmological scales are unique; they allow us to directly measure the luminosity distance d_L through their GW signal emitted. Unlike standard candles, which require a known intrinsic luminosity to infer distance, CBCs are

self-calibrating sources, earning them the name *standard sirens*. This property allows us to bypass the traditional cosmological distance ladder for cosmology purposes. However, to fully exploit CBCs for measuring the cosmic expansion, the redshift z of the source must also be known to establish the distance–redshift relation. For a more detailed discussion of the theoretical and experimental aspects of GW sources, we refer the interested reader to (Maggiore 2007 ; Sathyaprakash and Schutz 2009)

In this section, we will derive the expression for the luminosity distance $d_L(z)$ as a function of redshift, essential for GW cosmology. To do so, we start by revisiting the cosmological framework of an expanding Universe, which is governed by the Friedmann–Lemaître–Robertson–Walker (FLRW) metric and described by a scale factor $a(t)$. This geometric framework provides the foundation for understanding distances in a Universe that is not static but evolving over time. By relating the observed flux from a source to its intrinsic luminosity and incorporating the effects of cosmic expansion, we will ultimately arrive at the key expression for $d_L(z)$, which ties the observed distance to the redshift of the source. This equation is pivotal for utilizing CBCs as standard sirens and measuring the expansion of the Universe with GWs. We recommend (Dodelson 2003) for a more dedicated introduction for the field of cosmology.

10.2.1. An introduction to the cosmic expansion

Modern cosmology rests on the assumption of the cosmological principle, which states that the Universe is both *homogeneous* and *isotropic* when viewed on sufficiently large scales. Homogeneity means that the Universe has the same properties at every point, while isotropy implies that it looks the same in every direction. Although the Universe exhibits rich structures such as stars, galaxies, and clusters on small scales, observations show that on sufficiently large scales, the matter distribution becomes statistically uniform. A major observational confirmation of this principle is the expansion of the Universe. In the 1920s, Edwin Hubble and Georges Lemaître independently discovered that distant galaxies are receding from us, with a recession velocity that increases with distance. This relationship is known as the Hubble–Lemaître law:

$$v = H_0 d, \tag{10.1}$$

where v is the recession velocity, d is the distance to the galaxy, and H_0 is the Hubble constant, quantifying the present-day expansion rate of the Universe in $\text{km s}^{-1} \text{Mpc}^{-1}$. This quantity plays a central role in cosmology and will be discussed further in this section. One of the major open tensions in cosmology concerns the value of the Hubble constant. Observations of the Cosmic Microwave

Background yield a value of $H_0 = 67.49 \pm 0.53 \text{ km, s}^{-1}, \text{ Mpc}^{-1}$ —representing an early–Universe measurement (Planck Collaboration 2020)—whereas local measurements based on Cepheid–calibrated Type Ia supernovae indicate a higher value of $H_0 = 73.04 \pm 1.04 \text{ km, s}^{-1}, \text{ Mpc}^{-1}$ (Riess *et al.* 2022).

The Hubble–Lemaître law in Eq. 10.1 is in fact a product of Einstein’s General Relativity when supplied with a homogeneous and isotropic metric, the Friedmann–Lemaître–Robertson–Walker (FLRW) metric. This metric reads:

$$ds^2 = -c^2 dt^2 + a^2(t) \left[\frac{dr^2}{1 - kr^2} + r^2(d\theta^2 + \sin^2\theta d\phi^2) \right], \quad [10.2]$$

where ds^2 is the spacetime interval, c is the speed of light, $a(t)$ is the scale factor that evolves with cosmic time t , dr^2 the radial component of the metric, often called comoving coordinate, k the curvature of the Universe and θ, ϕ the usual angles for polar coordinates. As we will see later, the scale factor $a(t)$ is responsible for the expansion of the Universe, while the comoving coordinate dr^2 represents only the fixed distances defined on the spacetime metric. Let us now restrict ourselves to the case of a flat Universe ($k = 0$), as supported by most cosmological observations (see (Turner 2022) for a review). We can then center our reference frame on the observer and consider a galaxy emitting photons at some comoving distance r , with the condition $ds^2 = 0$. In other words, we are now working with a simplified version of the FLRW metric.

$$0 = -c^2 dt^2 + a^2(t) dr^2, \quad [10.3]$$

from which we can learn some interesting properties of cosmological sources.

- **Cosmological sources have a recessional velocity due to the Universe expansion:** By integrating Eq. 10.2, we can show that the physical distance traveled by the photon is

$$d_p(t) \equiv \int c dt = a(t) \int dr = a(t) d_c, \quad [10.4]$$

where $d_p(t)$ is the physical distance, $a(t)$ is the scale factor, and d_c is the comoving distance, which remains fixed for objects moving with the expansion of the Universe (i.e., with zero peculiar velocity). The scale factor $a(t)$ describes the relative expansion of the Universe over time, representing how distances between objects increase as the Universe expands. Fig. 10.1 provides a visual representation of the interplay between scale factor, comoving distance and physical distance. Since the Universe




Figure 10.1: Visual representation of the interplay between scale factor, comoving distance and physical distance. The comoving distance is the distance identified by the coordinates of the square vertexes, the physical distance is the length of the line connecting two vertexes, and the scale factor is how much the line stretches over time.

does not change its scales on the human lifespan, for simplicity, the scale factor is set at 1 today. In this way, comoving and physical distance coincide.

Differentiating Eq. [10.4] with respect to time gives back the recession velocity with respect to an observer, which can be expressed as:

$$\begin{aligned} \dot{d}_p(t) &= \dot{a}(t)d_c + a(t)\dot{d}_c, & [10.5] \\ &\approx \frac{\dot{a}(t)}{a(t)}a(t)d_c, \\ &= H(t)d_p(t), \end{aligned}$$

where we define the Hubble parameter $H(t)$ as the ratio between the time derivative of the scale factor and the scale factor itself:

$$H(t) \equiv \frac{\dot{a}(t)}{a(t)}, \quad [10.6]$$

and \dot{d}_c the proper motion of the galaxy. The proper motion of galaxies is usually much smaller $\dot{d}_c \ll H(t)d_p(t)$ for far away galaxies, but it might be important to take into account for close by galaxies. For the rest of the chapter, we will assume that this term is negligible. In Eq. 10.6, the Hubble parameter $H(t)$ describes the rate of expansion of the Universe at any cosmic time t . At the present time t_0 , or equivalently very close to us, we define $H(t_0) = H_0$ (the Hubble constant), and Eq. [10.5] becomes

$$\dot{d}_p(t) = H_0 d_p(t). \quad [10.7]$$

We emphasize that Eq. [10.7] is equivalent to Eq. [10.1]. This equation is telling us that we expect cosmological objects (such as galaxies) to present a recessional velocity from us that is due to the expansion. For close by objects, the Hubble constant can be simply measured by obtaining the physical distance of cosmological sources and measuring their recessional velocity. The recessional velocity is “trivial” to measure for photons as it can be measured from the redshift of light of known elements, on the other hand, the distance of sources a bit less, as we will see later in the chapter.

- **Cosmological sources are redshifted:** As we briefly anticipated in the previous point, since cosmological sources have a recessional velocity from us, they are redshifted. We will now see how the redshift is connected to the scale factor and the definition of the distance.

Imagine the same cosmological source at a comoving distance r emits two light pulses, one at time t_s and the other at time $t_s + dt_s$. These two signals arrive on Earth at times t_d and $t_d + dt_d$, after traveling through an expanding Universe. As the position of the galaxy (source) does not change during this small time interval, we can use Eq. 10.3 to equate

$$\int_{t_s+dt_s}^{t_d+dt_d} \frac{dt'}{a(t')} = \int_{t_s}^{t_d} \frac{dt'}{a(t')} \quad [10.8]$$

$$\int_{t_s}^{t_s+dt_s} \frac{dt'}{a(t')} = \int_{t_d}^{t_d+dt_d} \frac{dt'}{a(t')}, \quad [10.9]$$

and if we assume that the scale factor does not change between the two light emissions¹, we can write

$$\frac{dt_s}{a(t_s)} = \frac{dt_d}{a(t_d)} \quad [10.10]$$

$$dt_d = \frac{a(t_d)}{a(t_s)} dt_s. \quad [10.11]$$

The equation above means that two photons emitted with a time interval dt_s are detected with a delayed time interval dt_d that is proportional to the ratio of the scale factors at the epochs of detection and emission. If the detection happens today, with $a(t_d) = 1$ and $a(t_s) < 1$, then their ratio is greater than one, and the photons are redshifted. It follows that

1. This is a reasonable assumption, as light pulses are generated by physical processes occurring on timescales significantly smaller than cosmic times

we can define a relation between the scale factor and the observed redshift z as

$$a(t_s) = \frac{1}{1+z}, \quad [10.12]$$

and similarly, signals emitted at a given frequency f_s are observed on Earth at a redshifted frequency

$$f_d = \frac{f_s}{1+z}. \quad [10.13]$$

To conclude our discussion on the redshift, let us argue that in cosmology, the redshift is often used as an equivalent of time and distance measures. To understand why, we can take Eq. 10.3, invert it and calculate the comoving distance

$$d_c(t_s) = \int_0^{t_s} \frac{cdt}{a(t)}. \quad [10.14]$$

We can now recall the definition in Eq. 10.12, and differentiate to obtain

$$dz = -\frac{\dot{a}(t)}{a^2(t)} dt \quad [10.15]$$

$$dt = \frac{dz}{H(z)(1+z)}, \quad [10.16]$$

and rewrite Eq. 10.14 as

$$d_c(t_s) = \int_0^z \frac{cdz}{H(z)(1+z)} \approx \frac{cz}{H_0}, \quad [10.17]$$

where we have used the fact that very close to us (at redshifts $z \ll 1$), the Hubble parameter is just the Hubble constant H_0 . Eq. 10.17 is yet another expression for the Hubble–Lemaître law, where instead of the recessional velocity, we use the redshift of the source.

To summarize our discussion so far, in order to measure the cosmic expansion from individual sources, we need to measure the distance of the object and its recessional velocity (or redshift). For EM sources, the redshift can be obtained thanks to the measurement of the redshift of spectral emission lines of elements, while the distance is less trivial to obtain. To obtain distances for EM sources at cosmological distances, we need *standard candles*, namely a class of sources for which we know the intrinsic luminosity L . A standard candle allows

us to measure the flux F and then obtain the distance by inverting its flux equation given by

$$F = \frac{L}{4\pi d^2}. \quad [10.18]$$

In a static Euclidean Universe, this would define the physical distance d_p seen earlier. However, in an expanding Universe, this relation is modified by two key effects. First, the energy of each photon is redshifted by a factor $(1+z)$, and the rate at which photons are received is also reduced by $(1+z)$, leading to a total flux suppression of $(1+z)^2$. The proper observed flux is then corrected such that

$$F = \frac{L}{4\pi d_c^2 (1+z)^2}. \quad [10.19]$$

Comparing this with the definition in Eq. [10.18], we identify the *luminosity distance* as

$$d_L(z) = (1+z)d_c(z), \quad [10.20]$$

which effectively absorbs the redshift dependency, and it is the distance that we measure from standard candles. Here we do not aim for a dedicated review of standard candles (we refer the reader to references in (Moresco *et al.* 2022) for more details), let us just note that standard candles have some intrinsic limitations. Their intrinsic luminosity needs to be calibrated on closer standard candles, i.e. they need a *cosmological ladder*, and this can introduce systematics in their calibration. Moreover, only a very restricted ensemble of sources, such as Cepheid and Supernova Type IA, can be standardized.

The relationship between luminosity distance and redshift (comoving distance) in Eq. 10.20 is the main ingredient required for *standard candles* and *standard sirens* cosmology. Now, we have seen that at low redshifts, d_c can be expressed in terms of the Hubble constant; however, at higher redshifts, the scale factor and the universe expansion are dominated by other cosmological parameters. At this point, we need to introduce the first Friedmann's equation

$$\left(\frac{H(z)}{H_0}\right)^2 = \Omega_m(1+z)^3 + \Omega_\Lambda, \quad [10.21]$$

that connects the evolution of the Hubble parameter (and hence the scale factor) to the Hubble constant, the dark matter energy density Ω_m and the dark

energy density Ω_Λ^2 . The first Friedmann's equation is a direct product of Einstein's field equations, supplied with the FRLW metric and a stress-energy tensor of an ideal fluid, where some polytropic relations are assumed for all the energy density components of the stress-energy tensor. We refer the reader to (Dodelson 2003) for a complete derivation, here we just note that with Eq. 10.21, we describe a flat Universe dominated by dark matter and dark energy. When deriving the Friedmann equation under the assumption of an ideal fluid, one can define the critical density of the Universe as

$$\rho_c = \frac{3H_0^2}{8\pi G}, \quad [10.22]$$

which represents the energy density required for a spatially flat Universe. All density parameters are then expressed relative to this critical density. Remembering the expression of the comoving distance in Eq. 10.17 and Friedmann's Eq. 10.21, we can rewrite the full expression for the luminosity distance as

$$d_L(z) = (1+z)d_c(z) = (1+z) \int_0^z \frac{dz'}{H_0 \sqrt{\Omega_m(1+z')^3 + \Omega_\Lambda}}. \quad [10.23]$$

Another relevant quantity that we will need for GW cosmology is the notion of differential of comoving volume dV_c . The comoving volume for a flat universe, given by

$$V_c = \frac{4}{3}\pi d_c^3(z), \quad [10.24]$$

is often used to calculate the number density of compact objects, or even galaxies present in a certain volume. Its differential as a function of redshift is hence given by

$$\frac{dV_c}{dz} = 4\pi \frac{c}{H_0} \frac{d_c^2(z)}{\sqrt{\Omega_m(1+z')^3 + \Omega_\Lambda}}. \quad [10.25]$$

To conclude our introduction to standard cosmology, it is instructive to compare the typical distances and volumes relevant for GW observations. Figure 10.2 illustrates the various distance definitions introduced in this chapter for a flat Λ CDM Universe, along with the typical distances at which GW sources are observed. As we will discuss later, GWs allow for a direct measurement of the source's luminosity distance. For the most distant GW events, such as GW190521, it becomes necessary to use the full cosmological expression for the luminosity distance in terms of redshift and cosmological parameters.

2. Here, remember that we assume that the Universe is flat, i.e. $\Omega_k = 0$.

Flat Λ CDM with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_m = 0.3$




Figure 10.2: **Top panel:** Different definition of distances in cosmology compared to the median distance reported by three of the most discussed GW sources. **Bottom panel:** Comoving volume and its differential as a function of redshift. Both the panels have been generated using a flat Λ CDM model, with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $\Omega_m = 0.3$.

10.2.2. Compact binary coalescence at cosmological distance

In this section we introduce the relevant ingredients required for GW cosmology with CBCs. We refer the reader to the GW theory section of this book for a more in-depth introduction about GWs and their emission from CBCs.

A CBC is characterized by a set of parameters that can be classified as either *intrinsic* or *extrinsic*. The intrinsic parameters describe properties that are inherent to the binary system, including the masses (m_i), spins (χ_i and θ_i) of the individual objects, their tidal deformability (Λ_i) if we are talking about NSs, and the orbital eccentricity. Usually, the two spins (and their orientations)




Figure 10.3: Schematic view of a binary system of compact objects with their main intrinsic parameters such as the primary and secondary masses (m_1, m_2), their spin vectors (S_1, S_2) and their inclination angles (θ_1, θ_2).

are combined in two spin parameters; the effective spin parameter χ_{eff} and the precession spin parameter χ_p . The former is maximized when the two spins are aligned to the orbital angular momentum, the latter when one of the two components lies on the orbital plane. A schematic representation of the binary intrinsic parameters is provided in Fig. 10.3. The extrinsic parameters depend instead on the observer's perspective. These include the binary luminosity distance (or redshift), orbital inclination relative to the line-of-sight (ι), the time of merger as observed at the detector (t_m), the sky position (right ascension α and declination δ), and the coalescence phase ϕ_c and polarization angle ψ . Together, these parameters define the key features of a CBC and play an essential role in interpreting GW signals, especially in the context of cosmological studies.

The GW strain observed at the detector, which refers to the spatial deformation caused by the passage of GWs through the interferometer, can be expressed as a combination of both intrinsic and extrinsic parameters in the detector frame. The strain at a given time, t_m , is given by the equation:

$$h(t_m) = F_+(\alpha, \delta, \psi, t_m)h_+(t_m) + F_x(\alpha, \delta, \psi, t_m)h_x(t_m), \quad [10.26]$$

where the functions F_+ and F_x represent the detector's response to the two GW polarizations, h_+ and h_x . These response functions depend on several factors,

including the geometry of the detector, the sky position of the source, the time of arrival of the signal, and the polarisation angle. For ground-based GW detectors, CBC signals typically last from a few minutes to a few milliseconds. As a result, the response functions F_+ and F_x can be considered constant over the duration of the signal, as their variation occurs over time scales comparable to the sidereal day.

In Fourier space, under the quadrupole approximation and assuming circular orbits, the stationary phase approximation, negligible spin effects, and at the leading order of the Post-Newtonian (PN) expansion, the two GW polarizations can be expressed as:

$$\tilde{h}_+(f) = \mathcal{A}(f, \mathcal{M}) \frac{1 + \cos^2 \iota}{2} e^{i\Psi(f, \mathcal{M})}, \quad [10.27]$$

and

$$\tilde{h}_\times(f) = \mathcal{A}(f, \mathcal{M}) \cos \iota e^{i(\frac{\pi}{2} + \Psi(f, \mathcal{M}))}, \quad [10.28]$$

where the phase $\Psi(f, \mathcal{M})$ is expressed as

$$\Psi(f, \mathcal{M}) = 2\pi f t_m - \frac{\pi}{4} - \phi_c + \frac{3}{128} \left(\frac{\pi G \mathcal{M}}{c^3} \right)^{-5/3} \frac{1}{f^{5/3}}, \quad [10.29]$$

and the amplitude $\mathcal{A}(f, \mathcal{M})$ as

$$\mathcal{A}(f, \mathcal{M}) = \frac{1}{d_p} \frac{5}{24\pi^{4/3}} \frac{(G\mathcal{M})^{5/6}}{c^{3/2}} \frac{1}{f^{7/6}}. \quad [10.30]$$

In Eq. 10.29 and Eq. 10.30, f is the frequency of the GW and d_p is the physical distance of the source to the detector and the chirp mass \mathcal{M} is defined as

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}, \quad [10.31]$$

computed as a combination between the primary (m_1) and secondary (m_2) masses of the coalescing binary system. Usually, the convention between m_1 and m_2 is that the primary mass is larger or equal to the secondary so that $q = m_2/m_1 \leq 1$. The chirp mass of the two objects is typically well-measured from the GW phase, while the distance is measured from the amplitude term.

As discussed previously, when we consider GW signals originating from the coalescences of binary systems located at cosmological distances, we have to account for the effect of the cosmological redshift due to the expansion of the Universe. From now on, we will use the subscript “d” when we refer to detector quantities and “s” when we refer to source quantities. The first aspect that we need to consider is that now the Fourier space must be computed at the detector, meaning that we should consider the cosmic expansion. This can be done easily using the properties of the Fourier transforms. At the detector, the signal in time domain can be calculated from the signal in time domain at the source, considering

$$h_d(t_d) = h_s\left(\frac{t_d}{1+z}\right). \quad [10.32]$$

We can now apply the Fourier transform in detector time domain \mathcal{F}_d ,

$$\begin{aligned} \mathcal{F}_d[h_d(t_d)] &= \mathcal{F}_d\left[h_s\left(\frac{t_d}{1+z}\right)\right] \\ &= (1+z)\mathcal{F}_d[h_s(t_d)](f_d(1+z)) \\ &= (1+z)\tilde{h}_s(f_d(1+z)). \end{aligned} \quad [10.33]$$

In the second step above, we have used the properties of the Fourier transforms. In writing Eq. 10.33, we have taken into account the fact that the frequency is redshifted. We have also dropped the polarization index to indicate that this equation is valid for both polarizations. We can see that in the detector, the signal is the original signal, multiplied by a redshift factor and calculated at frequencies that are the redshifted ones. Eq. 10.33 is a standard relation in cosmology, in the limit that $z \rightarrow \infty$, the spectrum at the detector will be concentrated at low frequencies with a large amplitude.

Let us now explicitly expand Eq. 10.33. We will start with the phase term in Eq. 10.29. If no additional phase will arise from the amplitude of the GW in Eq. 10.30 (as we will show later), then the phase of the GW signal at the detector is

$$\Psi_d(f_d) = \Psi_s(f_d/(1+z)). \quad [10.34]$$

Given the above equation, to obtain the phase at the detector, we simply need to replace $f_s = f_d(1+z)$ and also remember that the inverse is true for times

(i.e. $t_s = t_d/(1+z)$). Eq. 10.29 can then be written as

$$\begin{aligned}\Psi_d(f_d, \mathcal{M}_d) &= 2\pi f_d(1+z) \frac{t_{m,d}}{1+z} - \frac{\pi}{4} - \phi_c + \frac{3}{128} \left(\frac{\pi G \mathcal{M}}{c^3} \right)^{-5/3} \frac{(1+z)^{-5/3}}{f_d^{5/3}} \\ &= 2\pi f_d t_{m,d} - \frac{\pi}{4} - \phi_c + \frac{3}{128} \left(\frac{\pi G \mathcal{M}_d}{c^3} \right)^{-5/3} \frac{1}{f_d^{5/3}}.\end{aligned}\quad [10.35]$$

The above equation expresses the GW source phase in terms of detector frequency and a newly defined quantity, the redshifted chirp mass $\mathcal{M}_d = (1+z)\mathcal{M}_s$. It is interesting to note that in the detector, the GW phase is mathematically equivalent to the phase of a GW signal with redshifted masses $\mathcal{M}_d = (1+z)\mathcal{M}_s$. Let us now see what happens to the GW amplitude in Eq. 10.30. Again, we replace $f_s = f_d(1+z)$, and again we redefine the redshifted chirp mass $\mathcal{M}_s = \mathcal{M}_d/(1+z)$. With these definitions, we have

$$\mathcal{A}(f_d, \mathcal{M}_d) = \frac{1}{d_p(1+z)^2} \frac{5}{24\pi^{4/3}} \frac{(G\mathcal{M}_d)^{5/6}}{c^{3/2}} \frac{1}{f_d^{7/6}}.\quad [10.36]$$

We can now write explicitly Eq. 10.33 using Eq. 10.36 and Eq. 10.35. In doing so, we will remember that the luminosity distance is defined as $d_L = d_p(1+z)$, so that the two polarizations now become

$$\tilde{h}_{d,+}(f_d) = \frac{1}{d_L} \frac{5}{24\pi^{4/3}} \frac{(G\mathcal{M}_d)^{5/6}}{c^{3/2}} \frac{1}{f_d^{7/6}} \frac{1 + \cos^2 \iota}{2} e^{i\Psi(f_d, \mathcal{M}_d)}\quad [10.37]$$

$$\tilde{h}_{d,\times}(f_d) = \frac{1}{d_L} \frac{5}{24\pi^{4/3}} \frac{(G\mathcal{M}_d)^{5/6}}{c^{3/2}} \frac{1}{f_d^{7/6}} \cos \iota e^{i(\frac{\pi}{2} + \Psi(f_d, \mathcal{M}_d))}\quad [10.38]$$

These equations reveal several crucial information.

- The observed GW signal retains the same functional form as in the source frame, but with two important substitutions: the physical distance d_p is replaced by the luminosity distance d_L , and the source-frame chirp mass \mathcal{M}_s is replaced by the redshifted chirp mass $\mathcal{M}_d = (1+z)\mathcal{M}_s$.
- As a result, the waveform detected by ground-based interferometers provides two key observables: the redshifted chirp mass (from the phase evolution) and the luminosity distance (from the signal amplitude).




Figure 10.4: From top to bottom: Plus and cross GW emission for an equal-masses binary with redshifted chirp mass of $32 M_{\odot}$ located at 500 Mpc. The figures on the right show the inclination of the binary with respect to the line of sight.

- However, these quantities are degenerate with redshift: one cannot disentangle the redshift from the mass based on the waveform alone. This fundamental limitation implies that the redshift of the source is not directly measurable from the GW signal, and additional information is needed to exploit standard sirens for cosmology.

Since for CBCs we can measure directly the distance via GW emission, we typically refer to this type of source as *standard sirens*. Standard sirens are self-calibrated sources, as in contrast to EM sources, they do not require an intrinsic luminosity (calibration) to provide a distance measure.

An important feature of GW signals is the degeneracy between the signal amplitude and the inclination angle of the binary's orbit relative to the observer. This degeneracy is the motivation for which standard sirens typically do not provide a precise measure of their distance. Specifically, the GW amplitude depends strongly on $\cos i$, see Eq. 10.37-10.38, making it difficult to distinguish between a nearby system viewed edge-on and a more distant system viewed face-on. A depiction of this effect from the point of view of the waveform at the detector is given in Fig. 10.4. As we can see from the plots, the emission of GWs is stronger perpendicularly to the orbital plane (as both polarizations are maximized), while it is weaker parallel to the orbital plane. This effect greatly impacts the overall amplitude of the resulting signal, as the

luminosity distance would. This degeneracy limits our ability to precisely infer the luminosity distance d_L from the observed signal. Fortunately, it can be mitigated by using waveform models that include additional physical effects, such as precession or higher-order modes. These effects imprint weak signatures in the signal that help to lift the degeneracy between inclination and distance, thereby improving our estimates of the luminosity distance. Accurate distance measurements are particularly important in GW cosmology, as they directly affect the precision with which we can constrain key cosmological parameters.

In summary, measuring the expansion of the universe requires both the luminosity distance d_L and the redshift z of each source. CBCs detected through GWs are the only known astrophysical sources that provide a direct measurement of the luminosity distance through the amplitude of their signal. From the waveform, we also extract detector-frame masses, but the redshift remains inaccessible due to its degeneracy with the source mass. Therefore, while GW sources have the potential to become powerful cosmological probes, additional information or assumptions are needed to determine z . The following sections review current approaches to assigning redshifts to GW sources.

10.3. Bayesian inference for cosmological and population properties of resolved sources

In this section, we introduce the hierarchical Bayesian inference (HBI) framework, the main statistical tool currently used for standard sirens cosmology. We will start by introducing the basics of Bayesian statistics and then derive the hierarchical likelihood in the case of GW standard sirens.

10.3.1. A brief recap of Bayesian theory

The main scope of Bayesian inference is to estimate a probability density function for a certain parameter θ given the some observations $\{x\}$. Given a certain model M , the Bayes' theorem allows us to write this probability as

$$P(\theta|\{x\}, M) = \frac{\mathcal{L}(\{x\}|\theta, M) P(\theta|M)}{P(\{x\}|M)}, \quad [10.39]$$

$\mathcal{L}(\{x\}|\theta, M)$ is the likelihood – the probability of observing $\{x\}$ given θ ; $P(\theta|M)$ is the prior – our knowledge about θ before seeing the data; $P(\{x\}|M)$ is the evidence – a normalization factor used for model comparison; and $P(\theta|\{x\}, M)$ is the posterior – the updated probability of θ given the data. Here, we do not wish to go into the details of Bayesian statistics; we refer the reader to

(MacKay 2003) for a more detailed introduction. We rather provide an operational example to understand its application.

The mean of a gaussian: Let's assume we are provided with a generator of random Gaussian numbers x . We are tasked to observe samples of x and estimate the mean of the gaussian μ . The first step is to write down the likelihood. The likelihood for obtaining the sample x given the mean μ and standard deviation σ of the normal distribution is

$$\mathcal{L}(x|\mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}}. \quad [10.40]$$

As different samples of x are independent from each other, the overall likelihood can be written as

$$\mathcal{L}(\{x\}|\mu, \sigma) = \prod_{i=1}^N \mathcal{L}(x_i|\mu, \sigma) = \prod_{i=1}^N \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x_i-\mu)^2}{2\sigma^2}}. \quad [10.41]$$

It can be shown that, if we define $\bar{\mu} = \sum_i x_i/N$ and $\bar{\sigma} = \sigma/\sqrt{N}$, then the above likelihood can be rewritten as

$$\mathcal{L}(\{x\}|\mu, \sigma) = \frac{1}{\sqrt{2\pi}\bar{\sigma}} e^{-\frac{(x_i-\bar{\mu})^2}{2\bar{\sigma}^2}}. \quad [10.42]$$

We now use this likelihood in the Bayes' theorem (Eq. 10.39) and write explicitly the posterior

$$p(\mu|\{x\}, \sigma) = \frac{\mathcal{L}(\{x\}|\mu, \sigma)p(\mu)}{\int \mathcal{L}(\{x\}|\mu, \sigma)p(\mu)d\mu} \propto \frac{1}{\sqrt{2\pi}\bar{\sigma}} e^{-\frac{(x_i-\bar{\mu})^2}{2\bar{\sigma}^2}}. \quad [10.43]$$

In the last step, we have assumed a uniform prior on μ (constant) and neglected the denominator as it is just a normalization constant. We have also assumed to have perfect knowledge of σ since we condition on it. We have discovered that the posterior on μ is itself a gaussian, centered around the mean of the dataset $\{x\}$ and with a standard deviation that becomes more and more precise as we have more observation (\sqrt{N}). This example helps us to understand the connection with frequentist statistics, where it is usually said that “errors improve as \sqrt{N} ”.

Our results can be easily validated numerically, we can, for instance, write a code that draws 50 samples from a normal distribution ($\mu = 0, \sigma = 1$) and then




Figure 10.5: Posterior distributions (different colors) for the values of μ given 50 observations x from a gaussian process. The vertical dashed line indicates the injected value of μ .

loop on the values of possible μ in the prior to calculate Eq. 10.43. Fig. 10.5 shows the result for this code, where multiple *realizations* of the 50 samples have been generated. One important aspect to notice is that the posterior does not have to present its mode at the true value of μ , since the measure of μ fluctuates due to the noise realization. Although, for a large number of observations, $N \rightarrow \infty$, we expect the posterior to converge to a Dirac-function centered on μ . These realizations, in virtue of their independence, can be later combined to obtain the posterior that would be obtained using all the 250 samples.

Introducing the selection bias: Let us make use of the previous example to also introduce the concept of selection bias. This time, our experiment is able to register only values of $x > 0$ (this is a threshold). Then the likelihood $\mathcal{L}(x|\mu, \sigma)$ of observing a value of x should be re-normalized by the number of x that you can effectively see. In other words, the likelihood of observing a data point x should effectively be normalized considering all the possible values of x that you can detect. This can be easily done as follows

$$\mathcal{L}(x|\mu, \sigma) = \frac{e^{[-(x-\mu)^2/(2\sigma^2)]}}{\int_{x_{thr}}^{\infty} e^{[-(x-\mu)^2/(2\sigma^2)]} dx} = \frac{e^{[-(x-\mu)^2/(2\sigma^2)]}}{I(\mu, x_{thr})} \quad [10.44]$$

Because the range of samples is restricted, the normalization at the denominator is equal to an unknown constant called $I(\mu, x_{thr})$. So the denominator has become a function of μ . The denominator approaches the value of the usual




Figure 10.6: **Left:** Posterior distributions of the estimated value of μ with 50 observations of x from a gaussian process without accounting for the selection bias. The vertical dashed line indicates the injected value of μ . **Right:** Same, but this time accounting for the selection bias.

normalization of the gaussian distribution when $\mu \gg x_{\text{thr}}$, meaning that we have a process for which we can see all the possible realizations, otherwise, it is smaller. Fig. 10.6 shows us the difference of a Bayesian inference with and without selection bias for this case. When the selection bias is not accounted for, we would find a posterior supporting higher values of μ . This is the consequence of the fact that we are not inferring the inference that observing values with $x < 0$ is forbidden.

10.3.2. Hierarchical Bayesian Inference with gravitational waves

We have seen that the first step to using Bayesian statistics is to write a likelihood. Different from the previous case, where we wanted to estimate a parameter μ common to all the observations, now we have an extra level of inference. First, we need to infer the binary parameters for each observation (which are different among sources), then we want to infer population-level parameters such as the Hubble constant. Because of that, this kind of inference is called hierarchical, meaning we have a hierarchy of parameters. We are now going to construct the hierarchical likelihood for that scenario, following (Mandel *et al.* 2019), and a more motivated reader can also follow a more accurate and detailed derivation in (Vitale *et al.* 2020).

We start our derivation by modeling the probability of generating a binary with parameters θ given a set of population parameters Λ .

$$p_{\text{pop}}(\theta|\Lambda) = \frac{T}{N} \frac{dN}{d\theta dt_d} \quad [10.45]$$

For this derivation, N is the total number of sources in the universe, T the observing time at the detector and $\frac{dN}{d\theta dt_d}$ the differential number of sources

(rate) per binary parameter (at the detector) and detector time. In writing the above identity, we have also assumed that the rate is constant in detector time (a reasonable assumption for cosmological sources as their rate do not change on human timescales), and the factor T is basically the result of the integration in dt_d . It follows that the probability of having a population of GW sources with parameters $\{\theta\}$ is

$$p_{\text{pop}}(\{\theta\}|\Lambda) = \prod_i^N \frac{p_{\text{pop}}(\theta_i|\Lambda)}{\int d\theta_i p_{\text{pop}}(\theta_i|\Lambda)}. \quad [10.46]$$

Now, let us imagine that we again have the presence of a selection bias (finite detector sensitivity), as for the case of the gaussian example, then we should also normalize the likelihood accordingly

$$p_d(\{\theta\}|\Lambda) = \prod_i^N \frac{p_{\text{pop}}(\theta_i|\Lambda)}{\int d\theta p_{\text{pop}}(\theta|\Lambda) p(\text{DET} = 1|\theta)}, \quad [10.47]$$

where the detection probability

$$p(\text{DET} = 1|\theta) = \int_{x \in \text{DET}} \mathcal{L}(x|\theta) dx, \quad [10.48]$$

is the integral of the GW likelihood over all the possible realizations of data that are detectable. If we would have been able to measure *perfectly* the binary parameters θ , we could have used directly 10.47 for our inference as likelihood. However, we do not have the luxury to perfectly measure θ_i , as we also have a GW likelihood $\mathcal{L}(x_i|\theta)$. As such, what we measure for the single event is actually

$$p(x_i|\Lambda) = \int d\theta \mathcal{L}(x_i|\theta) p_d(\theta|\Lambda) = \frac{\int d\theta \mathcal{L}(x_i|\theta) p_{\text{pop}}(\theta|\Lambda)}{\int d\theta p_{\text{pop}}(\theta|\Lambda) p(\text{DET} = 1|\theta)}. \quad [10.49]$$

Finally, as all the GW events are independent from each other, we can write the hierarchical likelihood

$$\mathcal{L}(\{x\}|\Lambda) = \prod_i^N p(x_i|\Lambda) = \prod_i^N \frac{\int d\theta \mathcal{L}(x_i|\theta) p_{\text{pop}}(\theta_i|\Lambda)}{\int d\theta p_{\text{pop}}(\theta|\Lambda) p(\text{DET} = 1|\theta)}. \quad [10.50]$$

The hierarchical likelihood can also be written in terms of physical rates. To do so, we need to introduce the expected number of detections

$$N_{\text{exp}}(\Lambda) = T \int d\theta \frac{dN}{d\theta dt_d} p(\text{DET} = 1|\theta), \quad [10.51]$$

and add a Poisson term to the likelihood

$$\mathcal{L}(\{x\}|\Lambda, N_{\text{exp}}) \propto e^{-N_{\text{exp}}(\Lambda)} [N_{\text{exp}}(\Lambda)]^N \prod_i^N \frac{\int d\theta \mathcal{L}(x_i|\theta) p_{\text{pop}}(\theta_i|\Lambda)}{\int d\theta p_{\text{pop}}(\theta|\Lambda) p(\text{DET} = 1|\theta)} \quad [10.52]$$

Using the rate parameterization, the hierarchical likelihood in Eq. 10.45 can be expressed as

$$\mathcal{L}(\{x\}|\Lambda, N_{\text{exp}}) \propto e^{-N_{\text{exp}}(\Lambda)} \prod_i^N T \int d\theta \mathcal{L}(x_i|\theta) \frac{dN}{d\theta dt_d}(\Lambda). \quad [10.53]$$

There also exist an equivalent version of that likelihood called “scale-free” likelihood. One can marginalize analytically the N_{exp} term using a scale-free prior $\pi(N_{\text{exp}}) \propto 1/N_{\text{exp}}$ to obtain a formally equivalent expression to Eq. 10.53.

$$\mathcal{L}(\{x\}|\Lambda) \propto \prod_i^N \frac{\int d\theta \mathcal{L}(x_i|\theta) \frac{dN}{d\theta dt_d}(\Lambda)}{\int d\theta p(\text{DET} = 1|\theta) \frac{dN}{d\theta dt_d}(\Lambda)}. \quad [10.54]$$

Eqs. 10.53-10.54 are the hierarchical likelihood used for GW cosmology.

In the context of GW cosmology, the binary parameters must be written in terms of *detector* binary parameters, namely luminosity distance, detector frame masses and spin parameters. The information on cosmological parameters enters either from an EM external data through the redshift (see later), or by modeling the CBC merger rate in terms of redshift and source masses. In the latter case, the cosmological information comes from the fact that the rate carries information about the source masses and redshift, while the GW likelihood measures the redshifted masses and luminosity distance. We can express the detector-frame merger rate as a function of the source-frame parameters through a change of variables:

$$\frac{dN}{dt_d d\theta} = \frac{dN}{dt_s d\theta_s} \frac{dt_s}{dt_d} \frac{1}{\det J_{d \rightarrow s}} \quad [10.55]$$

$$= \frac{dN}{dt_s d\theta_s} \frac{1}{1+z} \frac{1}{\det J_{d \rightarrow s}}, \quad [10.56]$$

where the time differential transforms as $1/(1+z)$, and the Jacobian accounts for the change of variables. It is important to notice that Eqs. 10.53-10.54 can either be written in detector variables using and dealing with the Jacobian to calculate the source rate, or directly in source variables without having to deal with the Jacobian. The two are mathematically equivalent, but in the second case, we will need to explicitly remember that the GW likelihood and detection probability are conditioned on the cosmological parameters (as they depend on detector quantities).

Since only masses and distances are affected by the expansion of the Universe, the Jacobian for this transformation is given by

$$\frac{1}{\det J_{d \rightarrow s}} = \frac{\partial d_L}{\partial z} (1+z)^2, \quad [10.57]$$

and using Eq. 10.23, it can be further written as

$$\frac{1}{\det J_{d \rightarrow s}} = \left(\frac{d_L(z)}{1+z} + c \frac{1+z}{H_0} \frac{1}{E(z)} \right) (1+z)^2. \quad [10.58]$$

Thus, the detector-frame CBC merger rate becomes

$$\frac{dN}{dt_d d\theta} = \frac{dN_{\text{CBC}}}{dt_s d\theta_s} \left(\frac{d_L(z)}{1+z} + c \frac{1+z}{H_0} \frac{1}{E(z)} \right) (1+z). \quad [10.59]$$

From this point, three different methodologies for population inference can be pursued, depending on the goals of the analysis and the available data. Each approach re-parameterizes the source-frame CBC rate accordingly.

10.4. Spectral sirens: Cosmology with GW sources and their source mass spectrum

One of the possibilities for obtaining redshift information from GWs alone is to make assumptions about the shape of the source-frame mass distribution of CBCs. The idea is the following: since the source and detector masses are related by a redshift factor of $(1+z)$, and since GW signals allow us to measure detector masses, we can statistically infer the redshift of each source by assuming a model for the source-frame mass distribution (with some free parameters). This approach typically involves modeling the population distribution of CBCs in the source frame. In the following, let us explicitly separate the redshift z from the rest of the source binary parameters θ_s , we can write:

$$\frac{dN_{\text{CBC}}}{dt_s dz d\theta_s} = \frac{dN_{\text{CBC}}}{dt_s dV_c d\theta_s} \frac{dV_c}{dz} = \mathcal{R}_0 \psi(z; \Lambda) p_{\text{pop}}(m_{1,s}, m_{2,s} | \Lambda) \frac{dV_c}{dz}, \quad [10.60]$$




Figure 10.7: Representation of some mass models currently in vogue for spectral sirens cosmology to describe the mass spectrum of stellar-mass BHs, namely the Broken – PowerLaw, the PowerLaw – Peak and the Multi – Peak.

where \mathcal{R}_0 is the local ($z = 0$) CBC merger rate density per time and per co-moving volume, typically expressed in $[\text{Gpc}^{-3}\text{yr}^{-1}]$. The function $\psi(z; \Lambda)$ is a phenomenological model for the redshift evolution of the merger rate, usually inspired by the star formation rate. The term $p_{\text{pop}}(m_{1,s}, m_{2,s} | \Lambda)$ represents the population model for the joint distribution of primary and secondary source masses. Note that Eq. 10.60 already involves a key assumption: the mass distribution is independent of redshift. While this may not strictly hold in reality—since the population could evolve over cosmic time—current observations have not revealed any significant redshift dependence. For mass distribution, one can assume flexible analytical models like the ones represented in Fig. 10.7, whose parameters are fit alongside the cosmological parameters. Here, we do not enter into the details of the modeling of the mass distribution and we refer the reader to (Palmese and Mastrogiovanni 2025) for a review with more in-depth discussion.

To understand why the source mass distribution can provide information about cosmology, it is helpful to consider the impact of redshift on the observed masses. GW detectors measure masses in the detector frame, which are redshifted relative to the intrinsic source masses by a factor of $(1 + z)$. If we had independent knowledge of the typical source masses of binary BHs (for instance, if we assumed they are all formed with primary masses around $35 M_{\odot}$), then any deviation from this scale in the observed detector masses could be attributed to redshift. In this way, a characteristic mass scale in the source population effectively acts as a standard ruler in mass space, allowing the redshift to be inferred statistically. Then, with this redshift measure and the GW luminosity distance measure from the waveform, we can infer the cosmic expansion. Of course, we do not expect the binary source masses to have all the same values. Instead, we expect a distribution of source masses like the ones in




Figure 10.8: Detector mass distribution as a function of luminosity distance and the Hubble constant assuming a source mass model constructed with a powerlaw and a Gaussian peak, *fixed* change in redshift. This figure is reproduced from (Chen *et al.* 2024).

Fig. 10.7. Through the rate Eq. 10.60, this distribution is converted to a d_L -dependent detector mass distribution as displayed in Fig. 10.8. For a different combination of cosmological parameters H_0, Ω_m , the resulting distribution of detector masses and function of d_L will change because the luminosity distance is a function of z and the cosmological parameters. HBI evaluates the fitting factor (the hierarchical) likelihood for this detector mass, luminosity distance distribution with real observed GW events.

An example with real GW data: It is interesting to provide here a simplified example of HBI in action for spectral sirens with current GW data. We will refer the reader to (Palmese and Mastrogiovanni 2025) for a more complete description of the results obtained from data in the current literature.

To evaluate the hierarchical likelihood in Eq. 10.52, we mainly need two ingredients. The integrals of the rate models over the GW likelihood and the evaluation of the expected number of detections. The former can be obtained via Monte Carlo integration by summing over posterior samples of the GW

binary parameters, namely:

$$\int \mathcal{L}(x_i|\theta) \frac{dN}{dt_d d\theta}(\Lambda) d\theta \approx \frac{1}{dN_{s,i}} \sum_{j=1}^{N_{s,i}} \frac{1}{\pi_{\text{PE}}(\theta_{i,j}|\Lambda)} \frac{dN}{dt_d d\theta}(\Lambda) \Big|_{i,j} \quad [10.61]$$

$$\equiv \frac{1}{N_{s,i}} \sum_{j=1}^{N_{s,i}} w_{i,j}, \quad [10.62]$$

where the index i refers to the event and the index j to the posterior samples of the events. We have also defined a weight $w_{i,j}$ of dimension equal to the number of events generated per unit of time. The latter can be evaluated as:

$$N_{\text{exp}}(\Lambda) = T \int p(\text{DET} = 1|\theta) \frac{dN}{dt_d d\theta} d\theta. \quad [10.63]$$

Typically, we do not have access to an analytical form of the detection probability (see Gair *et al.* (2023) for an introductory example in the context of GW cosmology with galaxy catalogs). The current approach to evaluate selection biases is to use Monte Carlo simulations of injected and detected events, often shortly referred to as *injections*. The injections are used to evaluate the volume that can be explored in the parameter space and correct for selection biases. Therefore, their occurrence is proportional to $p(\text{DET} = 1|\theta)$ and the population model used to generate them. Eq. [10.63] can also be approximated using Monte Carlo integration:

$$N_{\text{exp}} \approx \frac{T}{N_{\text{gen}}} \sum_{j=1}^{N_{\text{det}}} \frac{1}{\pi_{\text{inj}}(\theta_j)} \frac{dN}{dt_d d\theta} \Big|_j \equiv \frac{T}{N_{\text{gen}}} \sum_{j=1}^{N_{\text{det}}} s_j = R_0 \langle \text{VT} \rangle. \quad [10.64]$$

Here we have again defined a weight s_j with the dimension of a rate of events. Note that there is one fundamental difference with Eq. [10.62]. The injection prior $\pi_{\text{inj}}(\theta)$ must be properly normalized to obtain a reasonable value of N_{exp} , while a wrong normalization of $\pi_{\text{PE}}(\theta)$ (which is used in Eq. [10.62]) will only result in an overall normalization factor to the overall hierarchical likelihood. Note that Eq. 10.64 can also be used to define an interesting quantity, the explorable spacetime volume $\langle \text{VT} \rangle$ that quantifies how many Gpc^3 per year we are able to explore with our GW detectors. The explorable space-time volume




Figure 10.9: **Left:** Explorable spacetime volume for BBH signals with $\text{SNR} > 11$ during the third LIGO–Virgo–KAGRA observing run (O3), as function of H_0 . **Right:** Posterior distributions on H_0 calculated from 42 BBH signals from GWTC-3 with $\text{SNR} > 11$, and using three different population models.

can be obtained as

$$\begin{aligned} \langle VT \rangle &\equiv \frac{T}{R_0 N_{\text{gen}}} \sum_{j=1}^{N_{\text{det}}} s_j & [10.65] \\ &= \frac{T}{N_{\text{gen}}} \sum_{j=1}^{N_{\text{det}}} \psi(z^j; \Lambda) p_{\text{pop}}(m_{1,s}^j, m_{2,s}^j | \Lambda) \frac{1}{1+z_i} \frac{1}{\det J_{d \rightarrow s}|_i} \left. \frac{dV_c}{dz} \right|_i, & [10.66] \end{aligned}$$

where we have factorised the CBC merger rate at the detector using Eq. 10.60. Of course, coding up all the machinery required to perform HBI is not trivial, as it requires the implementation of cosmological and population models. For these lectures, however, we refer the reader to one of the python packages available for HBI with GWs, ICAROGW (Mastrogiovanni *et al.* 2024) and its dedicated tutorials (Mastrogiovanni 2023).

Fig. 10.9 shows two of the relevant quantities for HBI that can be computed with ICAROGW. On the left side of the figure, we see the average spacetime volume as a function of H_0 , while on the right, the posterior on H_0 for different choices of the mass distribution parameters. These plots have been generated by choosing a rate $\psi(z; \Lambda)$ that increases in redshift a power law plus peak model for the source masses³. The plots use a set of 42 BBHs with signal-to-noise ratio higher than 12 from the third GW transient catalog (GWTC-3) (Abbott *et al.* 2023). As shown in the plots, the average detectable space–time volume

³ the reader can check the tutorials for the actual values of the population parameters

increases with the Hubble constant. This scaling does not arise solely from the ability to detect events at higher redshifts—although it is true that the redshift detection horizon scales as $z_{\text{hor}} \propto H_0$ —but also because the differential comoving volume scales as $\propto H_0^{-3}$. The increase in detectable volume is further amplified by the fact that higher values of the Hubble constant correspond to probing regions of the Universe where the CBC merger rate is higher, as encoded in $\psi(z; \Lambda)$.

The right-hand plot illustrates another crucial point: population assumptions, particularly the mass model, significantly impact spectral siren cosmology. The displayed posteriors differ because they were generated using mass models with different Gaussian peak locations. This highlights the necessity of marginalizing over population parameters—such as those describing the mass spectrum—alongside cosmological parameters in current GW cosmology

10.5. Adding galaxy surveys to dark sirens

The first methodology that was actually proposed by (Schutz 1986) to use dark sirens for cosmology relates to the use of galaxy surveys. This idea involves using redshift information from galaxy surveys reported in the localization area of the GW event to identify potential hosts for each CBC event. In practice, what one can do is to build a redshift galaxy density profile for all the possible directions in which the source is localized, as shows in Fig. 10.10. The galaxy redshift density profile is not really uniform in comoving volume, especially at low redshifts, and can present over densities and under densities. The most probable redshift for the GW is on the over densities. Fig. 10.10 displays the logic behind this methodology. It is interesting to note that, from a historical point of view, this was the first method proposed for cosmology with GW dark sirens. Nowadays, we know that the spectral sirens method (cosmology with the source mass) and the galaxy catalog method are actually part of the same method, in fact, for both the methods it is necessary to model the CBC merger rate.

In this textbook, we will use the “rate” approach to see how to implement a galaxy survey for GW cosmology. We will refer the reader to (Gair *et al.* 2023) for a more complete step-by-step mathematical introduction on this topic from the point of view of probabilities. To parameterized the CBC merger rate as a function of the galaxy catalog, we can make two assumptions: first, that all CBC mergers occur within galaxies, and second, that the number of mergers per galaxy can be proportional to some galactic property, for instance the absolute magnitude M and redshift. The CBC merger rate, now expressed in terms of




Figure 10.10: Top plots: Galaxy number density profile as a function of redshift from the simulated galaxy catalog used in (Gair *et al.* 2023) for two random lines of sight. The distributions have been normalized in the corresponding redshift range. The black dashed line indicates a uniform in comoving volume redshift distribution. **Bottom plots:** Galaxy over/under density defined as the ratio between the galaxy number density profile and the uniform in comoving volume distribution.

the number of mergers per galaxy, is therefore rewritten as:

$$\frac{dN_{\text{CBC}}}{dt_s dz d\theta_s d\vec{\Omega}} = \int dM \frac{dN_{\text{CBC}}}{dt_s dz d\tilde{m}_s d\vec{\Omega} dM} \quad [10.67]$$

$$= \int dM \frac{dN_{\text{CBC}}}{dN_{\text{gal}} d\tilde{m}_s dt_s} \frac{dN_{\text{gal}}}{dz d\vec{\Omega} dM}, \quad [10.68]$$

where N_{gal} is the number of CBCs per galaxy, and $\vec{\Omega} = (\alpha, \delta)$ denotes the sky localization of the GW events. This new parameterization of the CBC merger rate can be understood as the product of two terms. The first one closely resembles the expression used in classic spectral sirens and can be written as:

$$\frac{dN_{\text{CBC}}}{dN_{\text{gal}} d\tilde{m}_s d\vec{\chi} dt_s} = \mathcal{R}_{\text{gal},0}^* \Psi(z, M; \Lambda) p_{\text{pop}}(\vec{m}_s | z, M; \Lambda). \quad [10.69]$$

Here, $\mathcal{R}_{\text{gal},0}^*$ replaces the usual local merger rate density and now represents the local CBC merger rate per galaxy per year. The function $\Psi(z, M; \Lambda)$ models how the rate of CBCs per galaxy evolves with redshift and magnitude, while the two population terms describe the source mass and spin distributions. It is a custom choice to model

$$\Psi(z, M; \Lambda) = \psi(z; \Lambda) 10^{0.4\epsilon(M_* - M)}, \quad [10.70]$$

where $\psi(z; \Lambda)$ is the usual merger rate as function of redshift used for spectral sirens and the second term simply introduces a GW hosting probability proportional to the galaxy intrinsic luminosity, namely $\propto L^\epsilon$. All of these term rates, besides the one dependent on M , can be effectively taken out the integral in Eq. 10.68.

The second term in Eq. 10.68, is the galaxy number density per redshift, steradian and absolute magnitude. If we were supplied with a galaxy catalog containing all the galaxies in the universe, this term would simply be a “histogram” in redshift, sky position and absolute magnitude of all the galaxies. Instead, realistic galaxy catalogs are flux-limited, meaning that they only report galaxies brighter than a certain apparent magnitude threshold m_{thr} . It follows that we need to apply a completeness correction by defining

$$\frac{dN_{\text{gal}}}{dzd\tilde{\Omega}dM} = \frac{dN_{\text{gal,CAT}}}{dzd\tilde{\Omega}dM} + \frac{dN_{\text{gal,OUT}}}{dzd\tilde{\Omega}dM}, \quad [10.71]$$

where $dN_{\text{gal,CAT}}$ and $dN_{\text{gal,OUT}}$ correspond to the number of galaxies reported by the galaxy catalog and brighter than m_{thr} , and to the completeness correction. The completeness correction is nonetheless given by the number density of missing galaxies in the redshift shell that can be computed as

$$\frac{dN_{\text{gal,OUT}}}{dzd\tilde{\Omega}dM} = \Theta(M - M_{\text{thr}}(m_{\text{thr}}, z, H_0)) \text{Sch}(M) \frac{1}{4\pi} \frac{dV_c}{dz}, \quad [10.72]$$

where Θ is a Heaviside step that is non-null when the absolute magnitude is fainter than the threshold absolute magnitude, and $\text{Sch}(M)$ is the Schechter function, namely a relation modelling the number density of galaxies in comoving volume per absolute magnitude.

Finally, the total CBC merger rate parameterization for the spectral siren analysis with a galaxy catalog becomes:

$$\begin{aligned} \frac{dN_{\text{CBC}}}{dt_s dz d\theta_s d\tilde{\Omega}} &= \mathcal{R}_{\text{gal},0}^* p_{\text{pop}}(m_{1,s}, m_{2,s} | \Lambda) \psi(z; \Lambda) \times \\ &\int dM 10^{0.4\epsilon(M_* - M)} \left(\frac{dN_{\text{gal,CAT}}}{dzd\tilde{\Omega}dM} + \frac{dN_{\text{gal,OUT}}}{dzd\tilde{\Omega}dM} \right). \quad [10.73] \end{aligned}$$

Before moving to a real example with some data, let us comment on two limiting cases of Eq. 10.73. In the limit that the galaxy catalog is 100% incomplete ($m_{\text{thr}} = -\infty$), then the integral in Eq. 10.73 is just the overall integral of the Schechter function (multiplied by the host probability). This integral can be computed analytically, and it results in Euler's gamma function, which acts as a normalization function. In this limiting case we see that (besides a normalizing factor), the rate is equivalent to the one of the spectral sirens. It follows that the mass distribution is fully dominating the implicit redshift information. In the other limiting case, the galaxy catalog is 100% complete and in Eq. 10.73 only the catalog term contributes. In this case, if the galaxies are provided with perfect redshift and sky position measures, Eq. 10.73 becomes

$$\frac{dN_{\text{CBC}}}{dt_s dz d\theta_s d\tilde{\Omega}} = \mathcal{R}_{\text{gal},0}^* p_{\text{pop}}(m_{1,s}, m_{2,s} | \Lambda) \psi(z; \Lambda) \times \sum_i^{\text{gal}} \delta(\tilde{\Omega} - \Omega_i) \delta(z - z_i) 10^{0.4\epsilon(M_* - M)}. \quad [10.74]$$

In the second scenario, the redshift and sky positions are entirely dictated by the galaxies, and the mass model acts as a weight factor between them. This shows us that the mass spectrum always enters the inference of galaxy catalogs.

Galaxy catalogs in action: As in the case of spectral sirens, the full inference of mass spectrum properties and the cosmological background parameters, with the addition of galaxy catalogs, is not an easy task. In this case, we have the additional difficulty of adding to the inference a galaxy survey that typically contains billions of galaxies that can make the computation of the hierarchical likelihood computationally inefficient. For this introductory text, we will follow (Gair *et al.* 2023) and work under some main assumptions

- We will neglect all the rate modeling associated to source masses and redshift of GW events (the function $\psi(z; \Lambda)$). We will further assume that GW events are at low redshifts and the $1/1+z$ factor coming from dt_s/dt_d is not important.
- We will assume that the galaxy catalog is complete and it measures perfectly the galaxies.
- We will assume that galaxies are all equally likely to host GW events, whatever their properties.

Under these three main assumptions, that could be valid for very well localized and close GW events, the CBC merger rate in detector frame can be written

as

$$\frac{dN_{\text{CBC}}}{dt_d dz d\tilde{\Omega}} = \mathcal{R}_{\text{gal},0}^* \sum_i^{N_{\text{gal}}} \delta(\tilde{\Omega} - \Omega_i) \delta(z - z_i). \quad [10.75]$$

If we note that the total number of CBC is $N_{\text{CBC}} = \mathcal{R}_{0,\text{gal}}^* T N_{\text{gal}}$ with N_{gal} the total number of galaxies in the universe, and recalling the definition for the population probability term in Eq. 10.45, we can also write that

$$p_{\text{pop}}(z, \tilde{\Omega} | \Lambda) = \frac{1}{N_{\text{gal}}} \sum_i^{N_{\text{gal}}} \delta(\tilde{\Omega} - \Omega_i) \delta(z - z_i). \quad [10.76]$$

Namely, we have obtained that the distribution of CBC sources can be described by a probability following the distribution of galaxies. We can now use the scale-free version of the hierarchical likelihood in Eq. 10.50 and analytically solve the integrals in $\tilde{\Omega}$ and z . In doing so, we will remember that the GW likelihood and detection probability actually depend on $d_L(z, H_0)$ and the sky direction, so that

$$\mathcal{L}(\{x\} | H_0) = \prod_i^{N_{\text{CBC}}} \frac{\sum_j^{N_{\text{gal}}} \mathcal{L}(x_i | d_L(z_j, H_0), \tilde{\Omega}_j)}{\sum_j^{N_{\text{gal}}} p(\text{DET} = 1 | d_L(z_j, H_0), \tilde{\Omega}_j)}. \quad [10.77]$$

This equation is telling us that for well localized GW events, the hierarchical likelihood can be simply computed by evaluating the GW likelihood (or posterior in luminosity distance and sky location) with the reported redshifts and sky positions of galaxies. It is also interesting to note that the selection bias

$$\sum_j^{N_{\text{gal}}} p(\text{DET} = 1 | d_L(z_j, H_0), \tilde{\Omega}_j) \propto H_0^3, \quad [10.78]$$

is proportional to the H_0^3 . This is a consequence of the fact that the detection probability for GWs is a function of d_L and as we change the Hubble constant, more and more galaxies enter inside the GW horizon. The number of galaxies that enter inside the luminosity distance horizon scales as a volume with redshift radius $z \propto H_0 d_{L,\text{horizon}}$, and hence the H_0^3 scaling.

In Fig. 10.11, we report a set of H_0 posteriors evaluated for Eq. 10.77 with the codes⁴ released in (Gair *et al.* 2023).

4. https://github.com/simone-mastrogiovanni/hitchhiker_guide_dark_sirens




Figure 10.11: From top to bottom Hubble constant posteriors generated with 200 simulated GW events with varying sky error budget on d_L and localized within 1 deg^2 and 5 deg^2 for two different lines of sight. The vertical dashed line marks the injected H_0 value. Posteriors are generated with codes released for (Gair *et al.* 2023).

10.6. Bright sirens: The “easiest” case

We conclude our overview of methods for GW cosmology with the case for which an EM counterpart is observed jointly with the GW emission. This scenario, as demonstrated by GW170817 (Abbott *et al.* 2017), are extremely unlikely, due to the rarity of GW detection from BNSs and the difficulty in identifying the transient EM counterpart. Naively, one can use these bright sirens to rapidly estimate cosmological parameters as follows: the GW provides an estimation of the luminosity distance, the EM counterpart a redshift and then the cosmological parameters are fit with a relation linking $d_L(z)$ and minimizing a χ^2 function. This is for instance what is reported in Fig. 10.12. However, while this treatment gives a rapid idea of the estimation of the cosmological parameters, it does not take into account the full complexity of selection biases at play with that type of detection. In this last section, we will see how to deal with EM counterparts within the hierarchical framework that we have discussed in the previously.




Figure 10.12: A χ^2 -like fit of the Hubble constant from the bright siren GW170817. The luminosity distance is inferred from the GW data and the redshift from the spectroscopic redshift of its host, NGC 4993. The orange contours is the $d_L(z)$ relation reconstructed assuming $d(z) = cz/H_0$.

For bright sirens, the parameterization of the CBC merger rate is equivalent to that one of the spectral sirens in Eq. 10.60. The difference now is that the likelihood in Eq. 10.52 is not anymore composed by solely the GW likelihood, but it is multiplied by the a new likelihood for the EM signal. From now on, let us work in source frame parameters, as it will be less tedious for calculations. The new likelihood can be written as follows

$$\mathcal{L}_{\text{GW+EM}}(x_i|z, \vec{\Omega}, \vec{m}_s, \cos \iota) = \mathcal{L}_{\text{GW}}(x_i|d_L(z, H_0), \vec{\Omega}, \vec{m}(\vec{m}_s, z), \cos \iota) \times \mathcal{L}_{\text{EM}}(x_i|z, \vec{\Omega}, \cos \iota), \quad [10.79]$$

where we explicitly have two terms, one likelihood for the GW part and one for the EM part. Eq. 10.79 tells us how well we measure the sky position, redshift, masses and inclination angle (and we will see why soon we included it back). At the same time, also the detection probability is now composed of a GW part

and EM part,

$$p_{\text{GW+EM}}(\text{DET} = 1|z, \vec{\Omega}, \vec{m}_s, \cos \iota) = p_{\text{GW}}(\text{DET} = 1|d_L(z, H_0), \vec{\Omega}, \vec{m}(\vec{m}_s, z), \cos \iota) \times p_{\text{EM}}(\text{DET} = 1|z, \vec{\Omega}, \cos \iota). \quad [10.80]$$

Before continuing, let us inspect Eqs.10.79-10.80. In both equations, we have added back the contribution of $\cos \iota$, since if the EM counterpart is collimated with respect to the orbital plane (e.g. a cone-shaped gamma ray burst), then the detection probability of the EM counterpart inherits a dependency on $\cos \iota$. In most of the current applications, it is assumed that the detection probability of an EM counterpart is more far reaching than the one of GW detections, meaning that $p_{\text{GW+EM}}(\text{DET} = 1|\dots) \approx p_{\text{GW}}(\text{DET} = 1|\dots)$. While this was a reasonable assumption for past GW observing runs, it is not anymore with the extended sensitivity range of GW detectors. Moreover, it is possible to try to model the EM counterpart, for instance the gamma ray burst afterglow, to try to measure $\cos \iota$ from EM data from the likelihood $\mathcal{L}_{\text{EM}}(x_i|z, \vec{\Omega}, \cos \iota)$. Indeed, this is of crucial importance given the strong degeneracy on the GW side between d_L and $\cos \iota$ (see Fig. 10.4). With this digression, we just wanted to argue that the inclusion of the EM counterpart is extremely rewarding in terms of (i) redshift localization and (ii) breaking the luminosity distance, $\cos \iota$ degeneracy. However, it carries the burden of having to model the geometry of the EM counterpart and its detection probability.

Let us now derive the hierarchical likelihood under some simplified assumptions. We will assume that

- There is no relevant information to obtain from the EM side on $\cos \iota$, hence that we can drop it from our equations.
- The GW detection probability dominates the EM detection probability, namely $p_{\text{GW+EM}}(\text{DET} = 1|\dots) \approx p_{\text{GW}}(\text{DET} = 1|\dots)$.
- The determination of masses from the GW likelihood, does not strongly correlate with the determination of the luminosity distance and sky position, namely

$$\mathcal{L}_{\text{GW}}(x_i|d_L(z, H_0), \vec{\Omega}, \vec{m}(\vec{m}_s, z)) = \mathcal{L}_{\text{GW}}(x_i|d_L(z, H_0), \vec{\Omega}) \times \mathcal{L}_{\text{GW}}(x_i|\vec{m}(\vec{m}_s, z)). \quad [10.81]$$

This is a reasonable assumption as masses are mostly dominated from the phase of the GW signal while the distance and sky position from

the amplitude and arrival time at the detectors. It is also reasonable since the redshift of the EM counterpart is measured with a precision scale where the GW likelihood is typically not correlated in terms of luminosity distances and masses (it is like exposing the GW likelihood around a small interval in redshift).

- The EM counterpart is perfectly localized in redshift

$$\mathcal{L}_{\text{EM}}(x_i|z, \vec{\Omega}) = \delta(z - z_i)\delta(\Omega - \Omega_i). \quad [10.82]$$

We now want to use Eq. 10.79-10.80, supplied with these assumptions to calculate Eq. 10.54 using detectors variables z, \vec{m}_s, Ω and the rate parameterization in Eq. 10.60.

We start with the denominator, this is formally equivalent to what we would have for the spectral siren case, namely the expected number of detections in Eq. 10.64. It can also be noted that, as in the galaxy catalog case, if we neglect mass information the selection bias term scales as H_0^3 as more and more possible host galaxies enter in the GW detection horizon. Considering a GW event i with host galaxy j and skipping some tedious steps for computation, we finally obtain that

$$I_{i,j} = \mathcal{L}_{\text{GW}}(x_i|d_L(z_j, H_0), \vec{\Omega}_j) R_0 \left. \frac{\psi(z_i|\Lambda)}{1+z_i} \frac{dV_c}{dz} \right|_{z=z_i} \times \int \mathcal{L}_{\text{GW}}(x_i|\vec{m}(\vec{m}_s, z_j)) p_{\text{pop}}(\vec{m}_s|\Lambda) d\vec{m}_s. \quad [10.83]$$

In the above equation, everything besides $\mathcal{L}_{\text{GW}}(x_i|d_L(z_j, H_0), \vec{\Omega}_j)$ acts as a normalization constant with respect to a varying H_0 . This means that, even if the mass model that we apply is wrong (but still supports the values of masses found for the events), the inference on H_0 will still be unbiased. Statistically speaking, this is a consequence of the fact that, the determination of the GW luminosity distance from the signal amplitude and the masses from the signal phase are conditionally independent given the redshift. This means that no information related to source masses can propagate to H_0 . Fig. 10.13 shows a depiction in terms of Bayesian direct acyclic graphs for this simplified bright siren scenario.

The hierarchical likelihood for the bright siren case can then be approximated as

$$\mathcal{L}(\{x\}|H_0, \Lambda) \propto \prod_i^N \frac{\mathcal{L}_{\text{GW}}(x_i|d_L(z_j, H_0))}{H_0^3}, \quad [10.84]$$




Figure 10.13: Schematic of a direct acyclic graph for the Hubble constant inference in a simplified bright siren case. The arrows represent conditional probabilities between nodes (random variables). The colored circles are observed quantities on which to condition. In this particular case, the redshift is perfectly measured and we estimate the luminosity distance and detector masses from amplitude and phase of the GW signal. According to the laws of probability, the posterior on H_0 given the redshift, detector, phase and amplitude is conditionally independent from the rest of the random variables to infer (including rate parameters)

assuming GW observatories sensitive only to low redshifts GW events. We warrant that for a real analysis, we would still need to consider possible uncertainties related to the population model. If we want to include uncertainties on the redshift of the EM counterpart, we can still do it by using

$$\mathcal{L}(\{x\}|H_0, \Lambda) \propto \prod_i^N \frac{\int dz \mathcal{L}_{\text{GW}}(x_i | d_L(z_j, H_0)) \mathcal{L}_{\text{EM}}(z_j | z)}{H_0^3}, \quad [10.85]$$

where $\mathcal{L}_{\text{EM}}(z_j|z)$ is a gaussian centered around the observed redshift. Eq. 10.85 is valid in the limiting case that the redshift uncertainties from the EM side are much smaller than the typical redshift scale on which the CBC merger rate evolves so that we can still consider the rate as a normalization constant.

10.7. Conclusions

We conclude this chapter with a bit of historical discussion about GW cosmology. In 1929, Henrietta Leavitt discovered the first standard candles (Leavitt and Pickering 1912), the Cepheid, that were later used by Edwin Hubble (Hubble 1929) to provide the first measure of the cosmic expansion (resulting in a biased $H_0 = 500$ km/s/Mpc). The error in the first Hubble's estimation of H_0 was later discovered to be a wrong calibration for galaxies distance. In the following 60 years, more standard candles were discovered, most notably Supernova type Ia, and new astrophysical calibrations for them were studied. Nowadays, the precision (and accuracy) on H_0 is of the order of a few per cent.

Almost 100 years later from the seminal discovery of Henrietta Leavitt, a new type of self-calibrating cosmological sources has been observed: GWs from compact binary coalescence. This exciting field started to be studied in 1986 by (Schutz 1986) and it has seen its dawn with the advent of the binary neutron star GW170817 in 2017, the first (and currently last) GW source with EM counterpart. Currently, GW cosmology is in the same situation as standard candles cosmology in the '50s. We have already explored the first GW sources to measure the cosmic expansion of the Universe and we are also beginning to learn how to properly calibrate the CBC merger rate to obtain an implicit redshift for dark sirens, in contrast to the luminosity distance of standard candles. As more and more GW detection are collected, the precision of the cosmological expansion parameters will improve, and in analogy to standard candles, we will also have to prove an excellent control of the systematics related to the calibration (statistical and astrophysical of our analyses).

In summary, these lectures only provides an introduction for GW cosmology, we expect more developments on the statistical and astrophysical aspects for GW cosmology will happen in the next years.

10.8. Bibliography

- Abbott, B. P. *et al.* (2017), A gravitational-wave standard siren measurement of the Hubble constant, *Nature*, 551(7678), 85–88.
- Abbott, R. *et al.* (2023), GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run, *Phys. Rev. X*, 13(4), 041039.

- Chen, H.-Y., Ezquiaga, J. M., Gupta, I. (2024), Cosmography with next-generation gravitational wave detectors, *Class. Quant. Grav.*, 41(12), 125004.
- Dodelson, S. (2003), *Modern Cosmology*, Academic Press, Amsterdam.
- Gair, J. R. *et al.* (2023), The Hitchhiker's Guide to the Galaxy Catalog Approach for Dark Siren Gravitational-wave Cosmology, *Astron. J.*, 166(1), 22.
- Hubble, E. (1929), A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae, *Proceedings of the National Academy of Science*, 15(3), 168–173.
- Leavitt, H. S., Pickering, E. C. (1912), Periods of 25 Variable Stars in the Small Magellanic Cloud., *Harvard College Observatory Circular*, 173, 1–3.
- MacKay, D. J. (2003), *Information Theory, Inference and Learning Algorithms*, Cambridge University Press, Cambridge, UK.
URL: <http://www.inference.org.uk/mackay/itila/>
- Maggiore, M. (2007), *Gravitational Waves. Vol. 1: Theory and Experiments*, Oxford University Press.
- Mandel, I., Farr, W. M., Gair, J. R. (2019), Extracting distribution parameters from multiple uncertain observations with selection biases, *Mon. Not. Roy. Astron. Soc.*, 486(1), 1086–1093.
- Mastrogiovanni, S. (2023), ‘Tutorials associated with "icarogw: A python package for inference of population properties of noisy, heterogeneous and incomplete observations"’.
URL: <https://doi.org/10.5281/zenodo.10135401>
- Mastrogiovanni, S., Pierra, G., Perriès, S., Laghi, D., Caneva Santoro, G., Ghosh, A., Gray, R., Karathanasis, C., Leyde, K. (2024), ICAROGW: A python package for inference of astrophysical population properties of noisy, heterogeneous, and incomplete observations, *Astron. Astrophys.*, 682, A167.
- Moresco, M. *et al.* (2022), Unveiling the Universe with emerging cosmological probes, *Living Rev. Rel.*, 25(1), 6.
- Palmese, A., Mastrogiovanni, S. (2025), Gravitational Wave Cosmology, , .
- Planck Collaboration (2020), Planck 2018 results. VI. Cosmological parameters, *A&A*, 641, A6.
- Riess, A. G., Yuan, W., Macri, L. M., Scolnic, D., Brout, D., Casertano, S., Jones, D. O., Murakami, Y., Anand, G. S., Breuval, L., Brink, T. G., Filippenko, A. V., Hoffmann, S., Jha, S. W., D’arcy Kenworthy, W., Mackenty, J., Stahl, B. E., Zheng, W. (2022), A Comprehensive Measurement of the Local Value of the Hubble Constant with $1 \text{ km s}^{-1} \text{ Mpc}^{-1}$ Uncertainty from the Hubble Space Telescope and the SH0ES Team, *ApJ*, 934(1), L7.
- Sathyaprakash, B. S., Schutz, B. F. (2009), Physics, Astrophysics and Cosmology with Gravitational Waves, *Living Rev. Rel.*, 12, 2.

Schutz, B. F. (1986), Determining the Hubble constant from gravitational wave observations, *Nature*, 323(6086), 310–311.

Turner, M. S. (2022), *The Road to Precision Cosmology*, , .

Vitale, S., Gerosa, D., Farr, W. M., Taylor, S. R. (2020), Inferring the properties of a population of compact binaries in presence of selection effects, , .

DRAFT

DRAFT

11

Gravitational waves as a probe of dense matter

Lami Suleiman¹ and Francesca Gulminelli²

¹*Nicholas and Lee Begovich Center for Gravitational Wave Physics and Astronomy, California State University Fullerton, Fullerton, California 92831, USA.*

²*Université de Caen Normandie, ENSICAEN, CNRS/IN2P3, LPC Caen UMR6534, 14000 Caen, France*




Figure 11.1: Schematic temperature-density plot showing the regions where the different states of matter and equations of state dominate, the boundaries representing isobars. The dashed line covers the structure of the sun from its surface (lower left) to its center (upper right) [Adapted from M. Schwarzschild, *Structure and Evolution of the Stars*, Princeton University Press 1958.]

11.1. Introduction to dense matter physics in the Universe

11.1.1. Motivation

The standard states of matter, as they can be accessed in terrestrial laboratories, only comprise four possible phases, that we all study in our undergraduate studies: solid, liquid, gas and plasma. Going towards astrophysical objects, the situation becomes more complex because both ultra-dilute and ultra-dense matter can be formed (Salaris and Cassisi 2005), as schematically represented in Fig. 11.1.

If gas and plasma rule the structure of classical stars like our sun, radiation dominates the pressure budget of the hottest stars (and of the early universe). Increasing compactness, new states of matter appear because of the increasing importance of quantum effects. The degeneracy pressure of electrons induces a new equilibrium state (solid plasma) that stabilizes the white dwarfs against gravitational collapse making their lifetime virtually infinite: the equilibrium is preserved even in the limit of zero temperature with zero radiative energy loss. This electron degeneracy pressure leads to a first limit of existence of dense matter, the density at which the Chandrasekhar mass is reached for a fully degenerate $e - H^+$ system (Chandrasekhar 1931). Beyond this limit, which is close to the transition between non-relativistic and relativistic kinematics for

the electrons in Fig.11.1, no equilibrium state exists if particles are interacting only gravitationally. However, the strong interaction comes into play when the average distance between the ions becomes comparable to the range of the nuclear force, corresponding to mass densities of the order of 10^{12} g/cm³ or more. The phase structure of such ultra-dense matter is still highly controversial, and the same is true concerning its physical properties and even its correct degrees of freedom. In the standard model of particle physics, the phase diagram of strongly interacting matter is governed by Quantum ChromoDynamics (QCD). However, the non-perturbative nature of QCD and the sign-problem of exact lattice calculations make direct analytical and numerical QCD calculations unfeasible (Nagata 2022). Because of these limitations, our theoretical predictions rely on effective models, that typically use hadronic degrees of freedom and are optimized via empirical observations.

The expected phase structure of ultra-dense matter is better represented using chemical potentials instead of density, as schematically shown in Fig.11.2. As a reference, the ground state of terrestrial atomic nuclei corresponding to densities of the order of $10^{13} - 10^{14}$ g/cm³, is depicted as a light blue region in the Figure. This representation underlines the fact that the study of ultra-dense matter cannot be disentangled from the notion of isospin asymmetry, because isospin is a good quantum number of the strong interaction. Because of that, ultra-dense matter cannot be modelled as a simple fluid like ordinary matter: two densities n_B and n_I (or equivalently: two chemical potentials μ_B and μ_I) must be specified to describe the state of matter ¹. In the case of a pure nucleonic composition (e.g. if the baryonic component is solely made of protons p and neutrons n), the baryonic and isospin chemical potentials are linked to the usual particle ones ², $\mu_q = \partial\mathcal{F}/\partial n_q$, with \mathcal{F} the free energy density, by the simple relations

$$\mu_B = \mu_n ; \mu_I = \mu_n - \mu_p. \quad [11.2]$$

In other words, the fundamental properties of QCD impose that the state and properties of matter are not only determined by the global baryon density

1. If strangeness is additionally considered, this introduces an extra chemical potential, and an extra dimensionality to the phase diagram.

2. If we suppose the presence of exotic particles (hyperons, kaons, deconfined matter..) the relations linking μ_B, μ_I to the particle densities are immediately worked out by imposing:

$$\sum_{C=B,I,L,S} \mu_C n_C = \sum_q \mu_q n_q, \quad [11.1]$$

where C are the conserved charges, the index q runs over the matter constituents, and $n_C = \sum_q C_q n_q$.




Figure 11.2: Schematic representation of the theoretically predicted phase diagram of ultra-dense matter in the temperature-chemical potentials space. Solid (dashed) lines represent first order (continuous) phase transitions, while circles are critical points. The approximate regions explored by atomic nuclei, neutron stars and neutron star mergers are represented as colored areas, while arrows give the thermodynamic trajectory followed in high energy heavy-ion collisions in different acceleration facilities [Taken from NuPECC Long Range Plan 2022 (https://www.esf.org/fileadmin/user_upload/esf/Nupecc-LRP2217.pdf), 'Perspectives for Nuclear Physics'.]

and temperature, but also on the matter composition. The evolution of the latter as a function of density and temperature depends on the astrophysical phenomenon under study as well as on the modelling itself, as we will discuss in the next section. As it can be seen in Fig.11.2, a complex phase diagram is expected when matter density overcomes the equilibrium density of atomic nuclei, the most relevant for actual astrophysical objects occurring in nature being the deconfinement transition towards elementary quarks and gluons ³.

11.1.2. Sites of ultra-dense matter

The best known astrophysical phenomenon that leads to the formation of ultra-dense matter is Core-Collapse SuperNovae (CCSN), occurring when the

³. Notice that the deconfinement transition might not coincide with chiral symmetry restoration, leading to different possible exotic phases of ultra-dense matter.

iron core of a massive star overcomes the Chandrasekhar limit $M_{\text{core}} \geq 1.4 M_{\odot}$, inducing the star gravitational collapse. If the initial core mass is of the order of $30\text{-}50 M_{\odot}$ or more, the collapse is so fast that it is thought to directly produce a stellar black hole. For lower masses, the free fall of matter onto itself is stopped when the core density overcomes the equilibrium density of nuclei, and the nuclear interaction becomes strongly repulsive. This leads to a core bounce and the formation of a shock wave which ejects the outer layers of the star in a CCSN event, leaving behind a hot and dense remnant. The latter further evolves into a neutron star or a black hole, depending on the initial progenitor mass and on the details of the dynamical evolution, that are not yet fully understood (Janka *et al.* 2007). Since the explosion dynamics breaks spherical symmetry, gravitational waves are expected to be emitted from CCSN, and such a possible detection by the Ligo-Virgo-Kagra (LVK) laser interferometer is eagerly awaited by the scientific community. However, the small signal-to-noise ratio (SNR) expected except for galactic events, and the complex waveform modelling that requires three-dimensional hydrodynamical simulations with neutrino transport and convection in full General Relativity (GR), limit the interest of this source as a direct probe of ultra-dense matter. The supernova core-collapse mechanism is believed to be the most important path for the formation of the other main site of ultra-dense matter, namely neutron stars (NS) (Haensel *et al.* 2007). In the initial collapse phase, the process of electron capture on nuclei and free protons leads to an important loss of electrons ("deleptonization") and a massive neutrino production. Neutrinos are initially trapped within the ultra-dense matter in this first post-supernova phase, corresponding to the formation of a proto-neutron star (PNS). This first stage, inaccessible to the observation since the PNS is hidden by the expanding supernova envelope for several years, is followed in a minute timescale by a rapid cooling of the newborn star via neutrino, and subsequently gamma emission.

At the stage where the star is observable, the evolution is sufficiently slow for all strong and weak interactions to be in equilibrium. In terms of chemical potentials, chemical equilibrium can be written as

$$\mu_n - \mu_p = \mu_e \leftrightarrow \mu_I = \mu_L. \quad [11.3]$$

This equation, together with the charge neutrality condition that links the leptonic and baryonic sector, imposes a relation between μ_B and μ_I , or equivalently between n_B and n_I , that is schematically represented in Fig.11.2.

It is important to remark that even if the surface temperature of observable neutron stars is never below $T_S \approx 2 \times 10^5$ K, the density is such that the star can always be considered as fully degenerate (see Fig.11.1). This means that the astrophysical observation of these objects provides a unique information on the ground state properties of the most extreme state of matter that can be reached in nature before the black hole singularity.

11.1.3. Neutron star observations

Neutron stars are observed in various wavelength of the electromagnetic spectrum. They were first identified by Jocelyn Bell Burnell and Antony Hewish in 1967 (Hewish *et al.* 1968) as radio pulsars, that is regularly pulsating emitters. The steady pulsating radio signal is well understood as due to the combined effects of a strong magnetic field and a fast rotation. In the simplest representation of a purely dipole field, the magnetic moment of the star being inclined to the spin axis, the electromagnetic radiation is beamed along the magnetic axis and thus rotates with the star, producing a lighthouse effect.

The most accurate mass measurements of NS's mass have been achieved with pulsar radio timing (Shapiro 1964 ; Bell 1996), and its extremely precise measurement of post-Keplerian parameters in binaries involving neutron stars, particularly double neutron star ones (Özel and Freire 2016 ; Alsing *et al.* 2018 ; Freire 2021). Pulsar radio timing has also permitted robust tests of gravity theories, particularly with the famous (and only known) double pulsar binary PSR J0737–3039 A/B for which ~ 16 years of radio data was accumulated and showed perfect agreement with general relativity (Kramer and *et al.* 2021). Neutron stars also emit in the X-ray wavelength throughout their life. Recently, The X-ray detector Neutron star Interior Composition ExploRer (NICER) (Gendreau *et al.* 2012) has exploited the relativistic effects of strong gravity on the X-ray emission from hot spots at the surface of millisecond pulsars to estimate the compactness ($C = M/R$) and simultaneously the mass (often with the help of radio information) and in turn deduce the radius of a few neutron star sources. This ensemble of observations confirms the expectation that the most extreme phases of matter should be present in the NS core: with masses ranging in the interval $1M_{\odot} \lesssim M \lesssim 2M_{\odot}$ and radii of the order of $9\text{km} \lesssim R \lesssim 15\text{km}$, the central density of the star is expected to be several times larger than the equilibrium density of atomic nuclei.

In the recent years, a new exciting probe of the inner structure of NSs has emerged as given by Gravitational Waves (GWs), which are the main object of this chapter. Binary Neutron Stars (BNS) are of particular importance in this context, as they have given the first historical indirect evidence of GWs (Hulse and Taylor 1975). This spectacular discovery of the binary pulsar PSR B1913+16 by Hulse and Taylor, that were awarded the Nobel price in physics in 1993, concerns the measurement of the rate of orbital decay of the binary, with an acceleration in perfect agreement with the hypothesis of GW emission in Einstein's theory of general relativity.

The first direct measurement of GWs from a BNS merger was performed by the LVK collaboration 25 years later with the famous GW170817 event (Abbott *et al.* 2017a), that stays still today the most precise GW measurement of a

BNS merger. This source is also the only multi-messenger event to date, where the GW signal was very precisely correlated to the electromagnetic emission in the full frequency band, from a gamma-ray burst that occurred in coincidence with the "chirp" signal associated to the merger, up to the optical emission followed for several days after the event. This optical counterpart, the so-called Kilonova, is interpreted as the first observational proof that BNS mergers are an important, if not the main site of the elusive r-process, the reaction process responsible of the synthesis of most elements heavier than iron. With the sensitivity of present interferometers (O4 run of the LVK collaboration), the GW signal of a BNS merger is detectable only in the last $\approx 10s$ of the inspiral phase, up to the merging. In this dynamical phase, the neutron stars are only perturbed via tidal effects as we will discuss at length in the next Section, and matter can still be considered as fully degenerate. Temperature is expected to rapidly build up at the merger though. If the mass of the merging stars is not as high as to lead to an immediate collapse into a black hole, a complex gravitational wave signal is expected from the oscillation modes of the hot post-merger metastable state (hypermassive neutron star).

The inspiral GW signal will be detailed in Section 11.3, while the prospects of detection of the post-merger will be touched upon in Section 11.4. To understand why and how these observations represent a unique observational window to access the properties of ultra-dense matter, we however need to first address the theoretical modelling of neutron star via its equation of state (EoS). Most interestingly, we will in particular explain why the inspiral gravitational waveform encodes the information on the EoS in a model-independent way. This is the subject of the next Section 11.2.

11.2. Neutron star modelling and the equation of state

11.2.1. Equation of state

An Equation of State (EoS) expresses the relation among thermodynamic variables such as temperature, free energy, entropy, enthalpy, chemical potential, density or pressure, thus describing the properties of matter in mechanical and thermal equilibrium. In a simple fluid at the thermodynamic limit, only two state variables are independent, and all the others can be expressed as a function of them via the EoS. The choice of the independent variables is arbitrary, but the common choice is to take temperature T and particle number density n (number of particles per unit volume) or mass density ρ as independent variables, and define the EoS as the relation $p(n, T)$ that expresses pressure as a function of density and temperature. This functional relation is characteristic of the material under study and of its thermodynamic phase. Once one EoS is known, *e.g.* $p(n, T)$ is given, any other state variable can be computed using universal thermodynamic relations.

As discussed in Section 11.1, in the case of charge neutral baryonic matter ruled by QCD, an additional degree of freedom has to be added, and a three-dimensional EoS must be considered in the form $f(P, n_B, n_I, T) = 0$ or $f(P, \mu_B, \mu_I, T) = 0$ as in the representation of Fig.11.2. This is notably the case to characterize matter in CCSN, PNS, or during the post-merger phase of binary coalescences involving neutron stars. For these applications, the most common choice for the independent variables is the set (n_B, y_L, T) corresponding to baryonic density n_B , electron fraction $y_e = n_L/n_B$ ⁴ and temperature T . Huge three-dimensional EoS tables must be provided, see Typel and et al. (2022), to describe the matter dynamical evolution with hydrodynamical codes, and are generally referred to as "general purpose equations of state". However, the EoS dimensionality reduces if we limit our attention to isolated NS or to the inspiral phase of BNS mergers. Indeed, even though neutron stars are born hot, with $T \simeq 10^{11}$ K, they cool and thermally equilibrate to $T \sim 10^8$ K in some ~ 100 years and to $T \simeq 10^5$ K in $\sim 10^6$ years (Potekhin and Chabrier 2017), which is still a short time compared to a NS lifetime. In a degenerate star such as a NS, thermal effects can be neglected when $k_B T \ll e_F$, e_F being the Fermi energy. A simple numerical application reveals that in all regions of the star both the electron and nucleon Fermi temperatures overcome $10^9 - 10^{10}$ K. This means that in isolated neutron stars (or in binaries before fusion) thermal effects can be safely neglected and one can assume complete thermodynamic equilibrium leading to the lowest energy levels possible in the matter (catalyzed matter hypothesis).⁵ Under the cold and catalyzed matter approximation the only remaining independent variable is the baryon number density n_B , see discussion in Section 11.1.2. The EoS is then 1-dimensional, and it can be entirely determined by the relation linking pressure P to n_B as

$$f(P, n_B) = 0. \quad [11.4]$$

In a completely equivalent way to Eq. [11.4], two other thermodynamic quantities are commonly used to describe a 1-dimensional EoS: (i) the rest-mass density $\rho_B = n_B \times m_B$, with m_B the baryon mass; (ii) the energy density ϵ that encompasses the contribution of the mass and momentum of particles as well as their interactions, and is related to P and n_B through the first law of

4. Here, n_L is the net lepton (electrons and muons) density, and $n_L = n_p$ if the baryonic component only contains protons and neutrons, such that y_L can also be identified with the proton fraction.

5. This hypothesis is not justified in the case of accretion onto the neutron star as it may significantly change the chemical equilibrium in the crust (Haensel and Zdunik 1990*b,a*). It might also be questioned in the NS crust, if the cooling is sufficiently fast for chemical equilibrium to be inhibited once the crust crystallizes (Dinh Thi, H. *et al.* 2023).




Figure 11.3: General picture of neutron star's structure in layers. [Taken from Kumar *et al.* (2024)]

thermodynamics. This latter quantity is particularly used by astrophysicists when constructing NS macroscopic parameters with relativistic hydrostatic solutions (Typel and *et al.* 2022), see next section.

11.2.2. Modeling the core and the crust of neutron stars

As we will explicitly see in Section 11.2.4, hydrostatic equilibrium implies an inhomogeneous density distribution inside any self-gravitation object, matter being denser in the core and more diluted close to the surface. Because of that, similar to the case for classical stars, the structure of a NS is supposed to be composed of different layers, see Fig. 11.3. To accurately study NS physics, it is then important to design the EoS for all the different parts of its internal structure.

While the physics of the atmosphere is important to study the thermal evolution of NSs, it contributes very little to the total mass and radius of the star (at most a few tens of centimeters), and as such it is usually neglected when modeling the EoS to solve hydrostatics equations and compute NS bulk properties. Underneath the atmosphere is the crust of the star, divided between the outer crust and the inner crust, for a full review see (Chamel and Haensel 2008). The crust is composed of inhomogeneous matter, with nuclei distributed on a lattice structure and immersed in a gas of degenerate fermions (electrons,

accompanied by unbound neutrons in the inner part). A common approach to describe the lattice consists in using the Wigner-Seitz cell approximation (Baym *et al.* 1971): each lattice point hosts one spherical nucleus surrounded by an electron gas which together compose a cell, that is periodically replicated in space. The crust composition (atomic Z and mass A number of the ion, and density of the neutron gas) is determined by the least Gibbs energy per cell at given density, a condition that is imposed by a ground state matter in the cold and catalyzed matter approximation. In all regions of the star, the electron energy density can be analytically computed as the one of a free degenerate relativistic Fermi gas:

$$\epsilon_e = \frac{m_e^4 c^5}{8\pi^2 \hbar^3} (x_r(2x_r^2 + 1)\gamma_r - \ln(x_r + \gamma_r)) , \quad [11.5]$$

with \hbar the reduced Planck constant, c the light velocity, $\gamma_r = \sqrt{x_r^2 + 1}$, $x_r = \hbar k_F^e / (m_e c)$, and the electron Fermi momentum is linked to the electron number density n_e by $k_F^e = (3\pi^2 n_e)^{1/3}$. The relation between n_e and baryon density n_B is determined by imposing charge neutrality and chemical equilibrium with respect to β -decay (β -equilibrium), see Eq.11.3. In the outer crust, the electrons can be considered ultra-relativistic to a very good approximation that fails (at $\sim 10\%$) only for the first shell of the outer crust, giving the simple expression

$$P_e(n_e) = \hbar c \frac{(3\pi^2)^{1/3}}{4} n_e^{4/3} . \quad [11.6]$$

Moreover, the electron pressure dominates the total pressure contribution while the lattice contribution is negligible such that $P(n_B) \simeq P_e(n_B)$. Also the composition in terms of ion species $Z(n_B)$, $N(n_B)$ is fairly under control, because the only unknown ingredient in the minimization of the Gibbs energy is the nuclear mass, and these latter are very precisely measured except in the case of very neutron rich nuclei (Wang *et al.* 2021). The electron chemical potential is a fast increasing function of density. As a consequence, the chemical equilibrium condition $\mu_n - \mu_p = \mu_e$ leads to increasingly neutron rich matter diving deeper into the star. The delimitation between the outer crust and the inner crust is defined by the neutron drip density point $\rho_{\text{nd}} \simeq 4 \times 10^{11} \text{ g/cm}^3$: at a sufficiently high density, matter is so neutron rich that it becomes energetically favorable for neutrons to spill outside of the nucleus. In the inner crust, the energy of the cell must include the contribution of the electrons, the nucleus and also the unbound neutrons, and the lattice contribution cannot be neglected. The inner crust pressure is heavily contributed to by the unbound neutrons and the dripline nuclei, that in terrestrial laboratories are unstable towards neutron emission and therefore cannot be experimentally produced. Because of that, the

inner crust EoS depends strongly on the model of the effective nucleon-nucleon interaction, as well as on the many-body treatment used to approximate the intractable quantum many-body problem: the most commonly used ones in the literature are, in order of increasing complexity and sophistication, the compressible liquid drop model (Mackie and Baym 1977 ; Douchin and Haensel 2001a), the extended Thomas-Fermi model (Centelles *et al.* 1993 ; Onsi *et al.* 2008) and the Hartree-Fock-Bogoliubov approach (Negele and Vautherin 1973 ; Pastore *et al.* 2011 ; Pearson *et al.* 2011). As we progress deeper in the inner crust, the nucleus structure is strained by the increasingly larger density, leading to nuclei cluster structures that are non spherical: we refer to it as pasta phases (Schneider *et al.* 2013 ; Caplan *et al.* 2018), and a large number of shapes (spaghetti, lasagna, gnocchi, bucatini, antispaghetti ...) might exist in a thin layer on the deeper end of the inner crust. Their possible presence might affect transport properties of stellar matter (see Schmitt and Shternin (2018) for a review), but it does not modify the EoS in a relevant way (Dinh Thi *et al.* 2022).

The core of NSs is composed of homogeneous matter: under sufficiently high densities, the nucleus structure is no longer energetically favorable and dissolves into an homogeneous fluid phase of baryons and leptons. Building an EoS model in the assumption of infinite homogeneous matter is analytically easier and not as numerically challenging as modeling the crust: a theory of strong interaction between particles is required, but not a model for complex cluster (nucleus) structures. In particular, a huge progress was recently achieved in Chiral Effective Field Theory (chEFT) (Machleidt and Sammarruca 2024), which now produces theoretical estimations of the energy density of homogeneous matter with controlled uncertainties, up to densities slightly above nuclear saturation⁶, $n_B \sim 0.2 \text{ fm}^{-3}$. However, the EoS of NS's core is still the least understood part of the star for two reasons. On the one hand, the density involved with NS cores (several times nuclear saturation density) cannot be achieved in nuclear laboratories and go beyond the theoretical limitations of chEFT, leading to a strong model dependence of the EoS. On the other hand, the composition of the core is unknown: in addition to neutrons, protons, electrons and muons (which appear in the early core), the deep parts of the core can host more "exotic" particles such as strange (hyperons) or non-strange (Δ s) heavy baryons (Sedrakian *et al.* 2023) or the density conditions can be such that a transition to a possibly super-conducting deconfined quark phase occurs (see Fig.11.2)

6. The saturation density is defined as the equilibrium density of homogeneous symmetric nuclear matter, and it is approximately equal to the central density of atomic nuclei, $n_{sat} \approx 0.16 \text{ fm}^{-3}$.




Figure 11.4: Adiabatic index Γ as a function of the mass density ρ for the nucleonic EoS DD2 (Typel *et al.* 2010), with inclusion of hyperonic degrees of freedom or of a first-order phase transition towards quark matter.

(Alford and Reddy 2003 ; Ferreira *et al.* 2020)⁷. For ideal degenerate fluids, a new particle q is formed in the matter as soon as the density is high enough such that its chemical potential $\mu_q = \sqrt{m_q^2 c^4 + k_F^2 c^2}$ overcomes the rest mass, $\mu_q \geq m_q c^2$, and its abundancy is governed by the relation between densities and chemical potentials (see Section 11.1). The emergence of a new (fermionic or bosonic) species leads to a reduction of the pressure, due to the availability of low-momentum single-particle states. This applies to standard model particles such as hyperons (Oertel *et al.* 2015), but also to hypothetical dark matter candidates (Das *et al.* 2020). However, the result can be strongly modified by the self-interactions as well as interactions with the nucleons, that are largely unconstrained. The situation is still more unclear in the case of a phase transition towards a deconfined phase, since no lattice QCD solution exists at finite chemical potential, and effective QCD models (bag model, PNJL, ...) are very much unconstrained (see Buballa (2005) for a review).

⁷. Because of the unknown composition, the deepest part of the core is often called "inner" core to be distinguished with respect to the "outer" core where the purely nucleonic composition is well established.

The qualitative effect of the possible emergence of a new degree of freedom or a phase transition in the NS core is depicted in Figure 11.4, that shows the behavior of the adiabatic index $\Gamma = d \ln P / d \ln \rho$ for a representative nucleonic model, as new degrees of freedom are introduced in the EoS model.

11.2.3. Effective models of the strong interaction

As discussed in Section 11.2.2, the EoS modelling of the NS core is particularly challenging because effective models of QCD in the non-perturbative regime are necessarily phenomenological in nature, but their parameters are largely unconstrained. As a consequence, a multitude of possible dense matter EoSs have been proposed in the literature, see the CompOSE data base (Typel and et al. 2022). In quantum many-body theory, the description of strongly interacting matter at the thermodynamic limit starts from the definition of an effective Hamiltonian or an effective Lagrangian for the strongly interacting particles. The basic quantity that is computed is then the energy density ϵ or the energy per particle e as a function of the relevant particle densities n_q , $q = n, p, \Lambda, \dots$. Once this energy functional is specified, the NS core EoS can be computed by adding the leptonic component that is model independent, and applying the conditions of charge neutrality and weak equilibrium.

There are two existing approaches to construct models of the energy functional: phenomenological approaches that rely on the density functional theory, and microscopic models that attempt a quasi-exact solution of the many-body problem starting from two-(or few-)body interactions in vacuum. Microscopic methods are conceptually the most physically founded approaches, as they rely on a limited set of free parameters that are fixed on precisely measured properties of light nuclear systems. The nuclear interaction is mostly taken from chEFT, an effective field theory approach where elementary degrees of freedom are nucleons and pions interacting via a force governed by spontaneously broken chiral symmetry. chEFT allows for a systematic low-momentum expansion known as chiral perturbation theory (ChPT). Contributions are analyzed in terms of powers of a typical momentum scale Q over a (large) breakdown scale $\Lambda \approx 1$ GeV at which the short distance structure becomes important and cannot be absorbed into simple contact terms. In this scheme, the various contributions to the potential are calculated systematically order by order, and theoretical uncertainties due to the truncation of the expansion can be computed. Once the nuclear interaction is defined, different techniques are employed to solve the many-body problem and finally compute the EoS. They include many-body perturbation theory, resummation schemes, or quantum Monte Carlo methods (see Gandolfi *et al.* (2015) ; Burgio *et al.* (2021) ; Machleidt and Sammarruca (2024) for recent reviews). Microscopic models, though being very appealing

conceptually, have also some drawbacks. First, they are limited to a pure nucleonic composition of the EoS, and to a restricted (low) density regime, because of the increase with density of the number of coupling constants linked to effective many-body interactions, and to the breakdown of the theory approaching the hard scale. Moreover, since chEFT is not a renormalizable theory, regulators must be introduced at each order to avoid divergencies, leading to a residual model dependence of the results that cannot be easily quantified. For these reasons, phenomenological approaches based on a parameterized form of the energy functional embedded in the mean field theory are also very popular. Non-relativistic formulations are based on fully phenomenological zero-range (Skyrme) or finite range (Gogny) effective interactions, augmented of density dependent terms to account for higher order correlations in an effective way (see Douchin and Haensel (2001a) ; Fantina *et al.* (2013) for two popular EoS models). With non relativistic approaches, the EoS is not causal by construction; however most models available in the literature respect causality up to a sufficiently high density to appropriately describe NS's interior. Relativistic Mean Field (RMF) approaches formulate the nuclear interaction with a relativistic Lagrangian density and assume that the nuclear force is mediated by effective mesons that are treated as classical fields (Glendenning 2012). To increase flexibility, non-linear or density-dependent couplings to the effective mesons are introduced. Phenomenological RMF approaches can also be used to model the introduction of new degrees of freedom such as hyperons (Oertel *et al.* 2015), as well as a possible deconfined phase with quark degrees of freedom (e.g. Nambu-Jona-Lasinio models) (Buballa 2005 ; Pereira *et al.* 2016). For both relativistic and non-relativistic formulations, the energy functionals depend on parameters that can be adjusted to a large pool of experimental nuclear data. They include a large number of nuclear masses and radii, as well as isoscalar and isovector giant resonances, electric dipole polarizability, parity violating asymmetry for some selected nuclei, see Klausner *et al.* (2025) for a recent comprehensive study with Skyrme forces. Because of their phenomenological nature, these effective EoS models do not suffer of the theoretical uncertainties associated to effective models of QCD; however, the empirical data they are optimized to only cover densities close to nuclear saturation, and compositions close to symmetric matter with almost equal proportion of neutrons and protons. Because of that, the different models produce widely different predictions for the NS EoS, particularly at high density.

While both microscopic and phenomenological approaches provide a way to build nuclear interaction and consequently compute the EoS, the astrophysical community in the last few years has favored so called agnostic approaches. The goal is to build a large ensemble of EoSs, not particularly based on a theory of strong interaction but rather on a simple formulation of the pressure-density relation such as spectral representations (Lindblom 2018), piecewise polytropes

(Read *et al.* 2009), sound speed models (Somasundaram *et al.* 2023) or Gaussian processes (Landry and Essick 2019). By randomly throwing the parameters of those formulation, one can build a large ensemble of EoSs that is then used in a Bayesian inference (see Section 11.2.4) to construct a posterior EoS ensemble constrained by observational data. The advantage of this approach is that it usually explores a large pressure-density plane (although some parametrizations have biases) but the disadvantage is that the posterior does not inform directly the strong interaction theory.

11.2.4. EoS and astrophysical observables

NS astrophysical parameters can be calculated by solving the hydrodynamics or hydrostatics equations. Given the compactness of NSs, the effects of strong gravitation cannot be neglected as they would be for the Sun, and the relativistic version of those equations must be considered. In the assumption that the space-time metric inside and outside of the star is isotropic, static and spherically symmetric, the corresponding line element can be written as

$$ds^2 = -e^{2\Phi(r)}c^2dt^2 + e^{2\nu(r)}dr^2 + r^2(d\theta^2 + \sin^2(\theta)d\phi^2), \quad [11.7]$$

with $\Phi(r)$ and $\nu(r)$ arbitrary functions of the radial coordinate r , and c the light velocity. In 1939, Tolmann (Tolman 1939) and Oppenheimer and Volkoff (Oppenheimer and Volkoff 1939) found an explicit solution to Einstein's equation (thus specifying $\Phi(r)$ and $\nu(r)$) by considering a star in equilibrium which matter forms an isotropic ideal fluid. The mass and radius of a NS are solutions of the resulting relativistic hydrostatic equations, generally referred to as the Tolman-Oppenheimer-Volkoff (TOV) equations.

For the exterior of the star, the metric is the simple vacuum Schwarzschild solution with

$$\Phi(r) = \frac{1}{2} \ln \left(1 - \frac{2Gm(r)}{rc^2} \right) = -\nu(r) \quad , \quad [11.8]$$

with $m(r)$ the gravitational mass contained in a radius r and G the gravitational constant. Outside the star's surface $r \geq R$ with R the total radius of the star, the gravitational mass is constant $m(r \geq R) = M$ with a value defined by the total gravitational mass of the star M . The functions $\Phi(r)$ and $\nu(r)$ can be understood respectively as the gravitational potential and redshift at radius r . The solution for $\Phi(r)$ and $\nu(r)$ inside the star is relatively simple because the stress-energy momentum tensor $T^{\mu\nu}$ of an ideal fluid contains only diagonal non-zero terms. Though NS matter is not an ideal fluid, at equilibrium only the diagonal part of $T^{\mu\nu}$ enters in the equilibrium equations, and viscosity does

not play any role. The time-time component of the tensor defines the energy density of the matter inside the star $T^{tt} = \epsilon$ and the space-space components define the pressure $T^{ii} = P$. On the one hand, the time-time component of Einstein equations specifies the form of the function $\nu(r)$ by introducing the gravitational mass of the star in terms of the energy density profile $\epsilon(r)$

$$m(r) = 4\pi \int_0^r (r')^2 \epsilon(r') dr' , \quad [11.9]$$

$$\nu(r) = -\frac{1}{2} \ln \left(1 - \frac{2Gm(r)}{rc^2} \right) , \quad [11.10]$$

with the integral constant in Eq. [11.9] defined by the boundary condition at the center of the star $m(r=0) = 0$. On the other hand, the space-space components relate $\Phi(r)$ and $\nu(r)$ (and consequently $m(r)$) in a ordinary differential equation

$$\frac{d\Phi}{dr} = \frac{2G}{rc^2} \frac{m(r)c^2 + 4\pi r^3 P(r)}{rc^2 - 2Gm(r)} . \quad [11.11]$$

Finally, the conservation of the stress-energy tensor leads to

$$\frac{dP}{dr} = -\frac{1}{2} \left(\epsilon(r)c^2 + P(r) \right) \frac{d\Phi(r)}{dr} . \quad [11.12]$$

The combination of Eqs. [11.9], [11.11] and [11.12] must be closed by the equation of state model $P(\epsilon)$. The total radius R and mass M of the star can be modeled by solving the TOV equations using the central pressure P_c as the boundary condition. By varying the central pressure, an $M - R$ sequence is built for a given EoS model, as presented in Fig. 11.5. This derivation underlines the important fact that the relativistic hydrostatics equations lead to a one-to-one correspondence between the EoS and the macroscopic observables of NSs.

As the EoS is at the center of solving the TOV equations and the deep layers of NSs remain a mystery, one can use observations to test EoS models for dense matter. The first and most constraining observation for the dense matter EoS in the deep layers of the NS core is the TOV mass limit. In general relativity, increasing the central pressure of the star does not increase indefinitely the NS mass. For each EoS model, there exists a limiting central pressure value beyond which the star collapses to a black hole (Oppenheimer and Volkoff 1939). Very accurate mass measurements of high mass NSs such as PSR J0740+6620 with a measured mass $M = 2.08 \pm 0.07 M_\odot$ (Fonseca and et al. 2021) therefore allow a first selection of EoS models, excluding those with a M_{TOV} lower than the observed maximum mass, e.g. EoS BSk19 (Goriely *et al.* 2010) as presented in




Figure 11.5: Neutron star mass-radius sequence for various EoS models (BSk19 (Goriely *et al.* 2010), SLy4 (Douchin and Haensel 2001*b*), BSk24 (Goriely *et al.* 2013), DD2 (Typel *et al.* 2010) with only nucleons, with hyperons and with quarks, and NL3 (Horowitz and Piekarewicz 2001)) compared to radio timing mass measurement of PSR J0740+6620 (Fonseca and *et al.* 2021) and PSR J1614-2230 (Arzoumanian and *et al.* 2018) and to NICER x-ray data for PSR J0740+6620 (Salmi *et al.* 2024), PSR J0437-4715 (Choudhury *et al.* 2024) and PSR J0614-3329 (Mauviard *et al.* 2025) (contours presented at $1\text{-}\sigma$ and $2\text{-}\sigma$). The finite pressure and causality condition limits are shown in yellow and green respectively.

Fig. 11.5. Concerning NS radii, as already mentioned the NICER X-ray detector (Gendreau *et al.* 2012) has recently achieved the simultaneous measurement of mass and radii of a few pulsars **three of which are reported as $1\text{-}\sigma$ and $2\text{-}\sigma$ contours in Fig.11.5.**

The TOV equations propose a simplified solution for the mass and radius of neutron stars with no or very little spin. However, when considering fastly rotating NSs in equilibrium, the metric is no longer static but stationary and the $M - R_{\text{eq}}$ sequence, with R_{eq} the equatorial radius, is drastically modified by the change in the neutron star shape. The star is no longer spherical due to the centrifugal force and for the same central density, the rotating solution will

present a higher mass than the TOV solution ; for details, see the introductory course (Gourgoulhon 2010) and Refs. (Haensel *et al.* 2008, 2009). The structure of the star is only affected for sub-millisecond rotational periods, corresponding to very high rotation frequency for pulsars, irrelevant for the inspiral of BNS mergers. However, these extreme rotations are expected in the dynamical evolution of the post-merger.

We note also that hydrostatics equations can be derived with a theory of gravitation that is not General Relativity, and that mass and radius of modified NSs can be strongly affected by the theory of gravitation (Olmo *et al.* 2020). However, general relativity has been thoroughly tested in the famous double pulsar binary PSR J0737–3039 with ~ 16 years of accumulated radio data (Kramer and et al. 2021), and nothing suggests a deviation from GR on the distance scales probed by NSs.

11.2.5. Bayesian inference of the EoS

To combine the information on dense matter that comes from astrophysical observations, with the one associated to nuclear theory and nuclear empirical data, a complete Bayesian analysis of the EoS parameter space has to be performed. To this aim, parametric energy functionals (also called meta-models) are employed, with a sufficiently flexible functional form to be able to reproduce existing nuclear models and interpolate between them (Margueron *et al.* 2018 ; Char, P. and Mondal, C. and Gulminelli, F. and Oertel, M. 2023). In the nucleonic sector, this can be achieved considering a Taylor expansion of the baryonic energy per particle e_B in the rescaled baryon density $x = (n_B - n_{sat})/3n_{sat}$ and the isospin asymmetry $\delta = (n_n - n_p)/n_B$ around the saturation density of symmetric nuclear matter n_{sat} :

$$e_B(n_B, \delta) = E_{sat} + \frac{1}{2!}K_{sat}x^2 + \frac{1}{3!}Q_{sat}x^3 + \frac{1}{4!}Z_{sat}x^4 \quad [11.13]$$

$$+ \delta^2 \left(E_{sym} + L_{sym}x + \frac{1}{2!}K_{sym}x^2 + \frac{1}{3!}Q_{sym}x^3 + \frac{1}{4!}Z_{sym}x^4 \right) + \dots$$

The coefficients of this expansion, also known as the nuclear empirical parameters (NEPs), entirely determine the nuclear density functional as long as it remains an analytic function, which is the case as long as no phase transition occurs and the EoS remains nucleonic in nature. If we choose for the meta-model an analytical function of x and δ , as for instance in Margueron *et al.* (2018) ; Huth *et al.* (2021) ; Lim and Schwenk (2024), using Eq.(11.13) the parameters of the meta-model can be mapped to the NEPs \mathbf{X} that are randomly generated, thus producing large non-informative priors of the possible energy functionals $\mathcal{M}(\mathbf{X})$, and consequently of the associated equations

of state. The approach can be extended to more general EoS models including phase transitions or hyperonic degrees of freedom, by coupling the nucleonic energy functional to a parametrized description of the deconfined phase, and correspondingly increase the size of the parameter set \mathbf{X} (Tews *et al.* 2018 ; Pfaff *et al.* 2022). Using Monte-Carlo techniques such as Markov Chain MC or nested sampling (Mackay 2003). Using the Bayes theorem (see Chapter 7, Eq. 7.49), the posterior probability for a given EoS is:

$$P(\mathbf{X}|\mathbf{c}) \propto \prod_k \mathcal{L}_k(\mathbf{X}) = \prod_k P(c_k|\mathbf{X}), \quad [11.14]$$

where $P(c_k|\mathbf{X})$ is the conditional probability of reproducing the data c_k assuming the meta-model instance $\mathcal{M}(\mathbf{X})$, and the index k runs over all the aforementioned constraints. Clearly, $P(\mathbf{X}|\mathbf{c})$ is automatically also the likelihood of all the stellar properties (e.g., mass-radius relation, mode frequencies) that can be derived by assuming the matter model $\mathcal{M}(\mathbf{X})$. The posterior probability distributions for observables $Y(\mathbf{X})$ are obtained through marginalization over the parameter space:

$$P(Y|\mathbf{c}) = \prod_{j=1}^N \int_{X_j^{min}}^{X_j^{max}} dX_j P(\mathbf{X}|\mathbf{c}) \delta(Y - Y(\mathbf{X})). \quad [11.15]$$

The advantage of modelling the (two-dimensional) energy functional instead of the simpler barotropic EoS through meta-modelling is that the Bayesian approach allows treating on an equal footing constraints coming from nuclear theory, nuclear experiments, as well as astrophysical observables, as far as we dispose of a theory or procedure $\mathcal{F} : \mathcal{M}(\mathbf{X}) \rightarrow \mathbf{c}$ that takes as input the values of the nuclear matter parameters \mathbf{X} and outputs the experimental, observational, or theoretical observable. Different likelihood models must be introduced for the different constraints, see Koehn *et al.* (2025) for a very recent and complete review. In particular, constraints coming both from nuclear structure (Klausner *et al.* 2025) and heavy-ion collisions (Sorensen *et al.* 2024) can be modelled by gaussian likelihoods. Even if the experimental measurements are relatively precise, the theoretical uncertainty for each $\mathcal{M}(\mathbf{X})$ realization coming from the many-body theory used to predict the observable dominates over the experimental uncertainty, which limits the influence of this type of data. Conversely, the theoretical calculations of the energy per particle $e_N(n)$ of infinite pure neutron matter with chEFT ab-initio methods (see Section 11.2.3) give the most constraining information to date, as far as the low density EoS is concerned. In these calculations, the uncertainty inherent to the truncation in the power counted can be directly estimated from the theory. The additional residual model dependence is accounted for by combining the uncertainty bands given

by the different approaches (Huth *et al.* 2021) into a single “conflated” band, that is often interpreted as a 90% confidence interval for $e(n)$. Finally, the relation between the EoS model and static astrophysical observables being direct through the solution of the TOV equation, the model likelihood of observational data can be directly taken from the observational posterior distribution that comes from the astrophysical parameter estimation. For example, in the case of mass-radius observations it reads :

$$\mathcal{L}_{NICER}(\mathbf{X}) = \prod_i \int_{0.7M_\odot}^{M_{TOV}(\mathbf{X})} P_i(m, R_{\mathbf{X}}(m)) dm, \quad [11.16]$$

where P_i is the measured joint probability distribution of mass and radius of the object i , and $M_{TOV}(\mathbf{X})$, $R_{\mathbf{X}}(m)$ are the TOV mass and the radius of the mass m star, respectively, calculated with the $\mathcal{M}(\mathbf{X})$ model (Montefusco *et al.* 2025). However, the most constraining observation to date comes from the GW probe, that we turn to examine in the next section.

11.3. The gravitational wave probe for dense matter

The gravitational force only depends on the space-time distribution of mass, and not on the nature, properties and phase of the matter under study. Because of this very basic statement, we might a-priori expect that GW cannot be a probe of the internal structure of neutron stars and therefore of ultra-dense matter. In spherical symmetry, Gauss theorem allows us replacing the mass distribution of a compact star by a point-like mass, and under this hypothesis the GWs signal from the inspiral phase of a compact binary coalescence, detailed in [Chapter XXX](#) of this book, can be analytically worked out in the post-newtonian (PN) expansion. Such a waveform only depends on the redshifted mass and spins of the merging stars, and the only indication that we are observing merging NSs instead of merging BHs comes from the mass estimation that can be extracted from the signal. Specifically, in the initial phase of the inspiral, for low orbital frequencies, the time evolution of the gravitational wave frequency is essentially determined by a specific combination of the two masses, the so-called chirp mass $\mathcal{M}_c = (m_1^3 m_2^3 / M)^{1/5}$, $M = m_1 + m_2$ being the total mass of the binary, [see Chapter XXX](#) for details. However, similarly to ocean tides on Earth brought forth by the presence of the moon’s gravitational field, a neutron star can be significantly tidally deformed by the presence of the other compact object in close proximity during a binary merger. In linear response theory, the importance of this deformation is solely determined by the matter EoS, as we work out in Section 11.3.1, and it is imprinted in the GW signal (Section 11.3.2), thus making the inspiral GW signal from BNS an essential probe, if not the most sensitive one, to the EoS of ultra-dense matter.

11.3.1. Relativistic tidal deformation in binary merger

In a BNS merger, compact object (1) feels the presence of compact object (2) as an external gravitational field in its vicinity that may perturb its space-time metric outside and inside the star. The external tidal field originating from (2) expresses the change in gravity acceleration per distance unit. It elicits an internal response from compact object (2) that necessarily depends on the neutron star interior, and therefore on dense matter equation of state. In this section, we present the fully relativistic derivation of gravitational Love numbers, that is to say the proportionality coefficients between the external tidal moment and the internal response, as discussed by (Damour and Nagar 2009 ; Binnington and Poisson 2009).

The simplest approach to evaluate the metric $g_{\alpha\beta}$ of a star deformed by an external tidal field is to consider a linear perturbation $h_{\alpha\beta}$ on a spherically symmetric and static background metric (Eq. [11.7]) and adiabatic tidal perturbation. The angular dependence of the perturbation treated on a spherical harmonics basis $Y_{\ell m}(\theta, \phi)$ as discussed in Thorne and Campolattaro (1967) leads to the perturbed metric

$$g_{\alpha\beta} = g_{\alpha\beta}^{(0)} + \sum_{\ell} \sum_m h_{\alpha\beta}^{\ell m} Y_{\ell m}(\theta, \phi), \quad [11.17]$$

and allows one to consider the internal response to this perturbation as a multipole moments expansion. The expression of the perturbed stress-energy momentum tensor $\delta T^{\mu\nu}$ of an ideal fluid (in the assumption that the companion does not contribute to it), together with the perturbed metric Eq. [11.17] is sufficient to design and solve the perturbed Einstein equations. The seminal works of Damour and Nagar (2009) ; Binnington and Poisson (2009) present in great detail the perturbed metric solutions for the interior and the exterior of the star and discuss the continuity at the surface to identify the tidally induced multipole moments in the solution. In practice, it is common to solve this system of coupled equations separately for the even and odd parity solutions, from which the so called gravito-electric and gravito-magnetic sectors (a parallel with electro-magnetism) arise respectively, and to distinguish between two types of multipole moments: the mass multipole moment \mathcal{M} (electric type) and the spin multipole moment \mathcal{S} (magnetic type)⁸ which are related to the gravito-electric \mathcal{E}_{el} and gravito-magnetic \mathcal{E}_{mag} relativistic tidal moments as

$$\mathcal{M}^{\ell m} = \lambda_{\ell} \mathcal{E}_{\text{el}}^{\ell m}, \quad \mathcal{S}^{\ell m} = \mu_{\ell} \mathcal{E}_{\text{mag}}^{\ell m}. \quad [11.18]$$

8. The lowest order in the multipole expansion for the gravito-electric and gravito-magnetic sectors lead to the gravitational mass and spin respectively, hence the commonly used "mass" and "spin" multipole qualifier.

The tidal polarizability coefficients λ_ℓ and μ_ℓ are directly related to the relativistic generalization of the tidal Love numbers developed in Newtonian gravity by (Love 1911 ; Murray and Dermott 1999) and they explicitly appear in the Einstein equations solution. The mass quadrupole moment and its associated second order gravito-electric tidal Love number k_2 are to date the most relevant quantity to explore dense matter physics in a neutron star merger. The set of equations solving the quantity k_2 has been discussed in a seminal work by Flanagan and Hinderer (2008) ; Hinderer (2008), where the authors propose to solve a singular second order differential field equations for the quantity y

$$r \frac{dy}{dr} + y(r)^2 + F(r)y(r) + Q(r) = 0, \quad [11.19]$$

where $F(r)$ and $Q(r)$ depend directly on the background metric solution (implying Eq. [11.19] must be solved simultaneously with the TOV equations) and the equation of state as

$$F(r) = e^{2\nu(r)} \left[1 - \alpha(r)(\epsilon(r) - P(r)) \right], \quad [11.20]$$

$$Q(r) = e^{2\nu(r)} \left[\alpha(r) \left(\epsilon(r) \left[5 + \frac{d\epsilon}{dP} \right] + P(r) \left[9 + \frac{d\epsilon}{dP} \right] - \frac{6}{\alpha(r)} \right) \right. \quad [11.21]$$

$$\left. - 4e^{2\nu(r)} (C(r) + \alpha(r)P(r))^2 \right], \quad [11.22]$$

$$\text{with } \alpha(r) = \frac{4\pi G r^2}{c^4} \quad \text{and} \quad C(r) = \frac{Gm(r)}{rc^2} \quad \text{the compactness.} \quad [11.23]$$

We note the presence of the inverse of the normalized sound speed $c^2/c_s^2(r) = d\epsilon/dP$: a correction is in order to avoid a divergence of this term in the case of an equation of state that includes a first order phase transition (zero pressure gradient for a jump in energy density), see (Postnikov *et al.* 2010 ; Pereira *et al.* 2020). The mass quadrupole tidal Love number is related to the solution y as

$$\begin{aligned} k_2(r) &= \frac{8C(r)^5}{5} e^{-4\nu(r)} [2 + 2C(r)(y(r) - 1) - y(r)] \\ &\quad \times \left\{ 2C(r)[6 - 3y(r) + 3C(r)(5y(r) - 8)] \right. \\ &\quad + 4C(r)^3 [13 - 11y(r) + C(r)(3y(r) - 2) + 2C(r)^2(1 + y(r))] \\ &\quad \left. - 3e^{-4\nu(r)} [2 - y(r) + 2C(r)(y(r) - 1)] 2\nu(r) \right\}^{-1}. \quad [11.24] \end{aligned}$$

Finally, the mass quadrupole dimensionless tidal deformability λ is the most commonly discussed quantity and is related to the $l = 2$ gravito-electric tidal Love number and the star compactness C as

$$\lambda(r) = \frac{2}{3}k_2(r)C^{-5}(r). \quad [11.25]$$

As it can be appreciated from Eq.[11.25], for a given stellar mass M the total mass quadrupole tidal deformability $\Lambda = \lambda(r = R)$, similar to the stellar radius R solution of the TOV equation and discussed in Section 11.2.4, solely depend on the EoS. Its measurement can therefore be considered as a direct probe of the properties of dense matter.

The mass quadrupole is not the only multipole moment to consider in the expansion: octupole, hexadecapole (etc.) mass multipole moments and spin multipole moments can also be calculated (Damour and Nagar 2009 ; Pappas *et al.* 2019). Particularly, the external tidal field can induce a spin multipole moment in a non-rotating neutron star (Cunningham *et al.* 1978), and similarly the spin itself induces a multipole moment related to rotational Love numbers. There are various types of multipolar responses from the star to a various number of external fields, but not all of them are easily detectable. For a publicly available solver of the TOV and k_2 Love number equations, we refer the reader to the LIGO Algorithm Library (LIGO Scientific Collaboration *et al.* 2018).

11.3.2. Impact of the tidal deformation on the gravitational wave signal

The deformation induced by an external field affects the binary merger and therefore impacts its gravitational wave emission. In simple terms, tidal interactions induce a loss of energy and momentum from the merging binary system, thus accelerating the inspiraling orbital motion of the two stars and affecting the amplitude and phase of the waveform, see Fig. 11.6. Different multipole moments induced in the star enter the gravitational waveform in different frequency bands. Let us focus on the effect of the mass quadrupole moment induced by an external tidal field on the gravitational waveform. To estimate the tidal contribution to the energy flux due to the gravitational radiation for a quasi-circular orbit (point particle stars), it is convenient to use the average weighted tidal deformation of the binary, called the effective tidal deformability

$$\tilde{\Lambda} = \frac{1}{16} \left(\frac{m_1 + 12m_2}{m_1} \Lambda(m_1) + \frac{m_2 + 12m_1}{m_2} \Lambda(m_2) \right), \quad [11.26]$$

with m_1 and m_2 the neutron star masses of the primary and secondary compact object in the binary, see (Flanagan and Hinderer 2008). The tidal deformation




Figure 11.6: Left part: schematic representation of the tidal effects. [Adapted from <https://www.nature.com/articles/s41550-024-02323-7>]. Right part: Plus polarisation of the gravitational wave strain as a function of time for a $1.4\text{-}1.4 M_{\odot}$ compact binary merger, calculated with the SEOB-NRV4T_SURROGATE waveform model. The strain is presented for no tidal deformation (gray), and for non vanishing $\Lambda(M)$ relations for various EoSs: SLy4, DD2 and NL3 leading to small, medium and strong mass quadrupole tidal deformation respectively.

terms in waveform models used to analyse BNS mergers are directly expressed in terms of $\tilde{\Lambda}$: it is a direct output of the intrinsic parameter estimation. The phase Φ of the gravitational waveform is the most affected by the tidal deformation in the binary (and the most easily detectable compared to the amplitude); the phase shift due to the deformation related to adiabatic tides is (Wade *et al.* 2014)

$$\delta\Phi = -\frac{117}{256} \frac{M^2 (\pi G M f)^{5/3}}{m_1 m_2 c^5} \tilde{\Lambda} \quad [11.27]$$

with $M = m_1 + m_2$ the binary total mass. The decrease of the phase shift at high frequency implied by Eq. [11.27] can be appreciated looking at the right part of Fig.11.6, for the model leading to the highest tidal effect. This phase correction occurs at fifth post Newtonian order, the next order would include a second parameter $\delta\tilde{\Lambda}$ corresponding to a different combination of Λ_1 and Λ_2 (see Eq.(6) of (Wade *et al.* 2014) for the formula), that would allow discriminating the individual tidal polarizabilities of the two stars ⁹. However, the sensitivity of existing interferometers is not sufficient to allow evaluating

⁹. In actual parameter estimation, for the determination of individual polarizabilities, it is generally assumed that the two neutron stars of the binary merger are governed by the same equation of state, such that the $\Lambda(m)$ relation is the same for the two.

this further correction, meaning that only the effective tidal deformability $\tilde{\Lambda}$ is presently measurable.

The mass quadrupole tidal response is the first and most important correction to the gravitational wave inspiral, and Eq. [11.27] is a pertinent estimation of the effects of tidal deformation on the gravitational wave signal given the current sensitivity of detectors. Higher order corrections on the gravitational waveform beyond the mass quadrupole (Henry *et al.* 2020) such as effects of spin-tidal couplings (particularly relevant for high spinning neutron stars) (Castro *et al.* 2022), induced spins related to the gravito-magnetic multipole moment response to external tides (Gupta *et al.* 2021), dynamical tides (in contrast to adiabatic ones) and oscillation modes (Schmidt and Hinderer 2019), or even viscosity (Ripley *et al.* 2023) can also be considered, especially in the context of future facility upgrades and the next generation gravitational wave detectors.

The LIGO/Virgo/KAGRA collaboration uses different gravitational waveform models to analyse the signal detected. In order to account for higher order tidal deformation or cross terms between tidal deformability and spin, several waveform models currently used rely on "quasi-universal relations", that is to say relations between chosen macroscopic parameters of neutron stars that were empirically shown to only marginally depend on the star interior description (and the EoS). To give an example, the so called I-Love-Q relations (Yagi and Yunes 2013) are able to break the degeneracy between the mass quadrupole moment and the spin in the waveform analysis: this universality is explicitly used by the IMR_NRTidal waveform model. Similarly, quasi-universal relations between the mass quadrupole moment and the frequency modes can be used to include oscillation modes in waveform models that include dynamical tides (Abac *et al.* 2024). For more details on waveforms and waveform modelling, see [Chapter XXX](#).

11.3.3. Observations of the tides

On the 17th of August 2017, during the second observing run of the LIGO/Virgo collaboration, the LIGO Hanford, LIGO Livingston and Virgo laser interferometers detected for the first time the gravitational wave emitted by the merger of two neutron stars (Abbott *et al.* 2017a), the famous GW170817 event. The signal, with a Signal to Noise Ratio (SNR) of 32.4, very loud for a BNS merger, lasted for ~ 100 s in the detectors sensitive band up to the coalescence time, followed 1.7 s later by a gamma-ray burst (GRB). The precise sky localization achieved by triangulation with the three detectors allowed an electromagnetic follow-up campaign, culminating in the observation by different terrestrial telescopes of an optical kilonova signal




Figure 11.7: Left figure: Effective tidal deformability $\tilde{\lambda}$ as a function of the chirp mass \mathcal{M}_c at a single mass-ratio value (peak of the posterior distribution for GW170817) for various EoS models, and compared to 1- σ and 2- σ contours for the sources GW170817 and GW190425 (data released by the LVK collaboration: <https://dcc.ligo.org/P2000223/public>). Figure adapted from J. Read [<https://github.com/jsread/APSP1ots2024>]. Right figure: distribution of the mass ratio and effective tidal deformability for GW170817 [Taken from (Abbott *et al.* 2019)].

around half a day after the merger. This remarkable multi-messenger detection, the unique of this kind up to now, has fostered a wealth of research and brought valuable information on a variety of interdisciplinary topics including cosmology, general relativity tests, and nucleosynthesis (Abbott *et al.* 2017b). For the purpose of the present chapter, the GW170817 event stands alone because the tidal information could be extracted with a Bayesian analysis by matched filtering the data with waveform models including the tidal information, and it was checked that the existing uncertainties in the waveform modelling only slightly modified the parameter estimation. The posterior distribution for the two relevant quantities which can be extracted from the Bayesian analysis, namely the effective tidal polarizability $\tilde{\lambda}$ and the chirp mass \mathcal{M}_c , are shown in Fig.11.7, in comparison with predictions obtained with different EoS models. The mass ratio $q = m_1/m_2$ (see Fig. 11.7), and consequently the two individual masses of the merging NSs, could also be estimated but with a larger uncertainty due to the degeneracy in the waveform between q and the aligned spin components of the stars.

Though a few other BNS and NS-BH mergers were detected in the successive years, the last one dating May 2023, the SNR of these events has not yet been high enough to provide constraining information besides the one of GW170817, see the comparison between GW170817 and GW190425 (Abbott *et al.* 2020) in Fig. 11.7. The EoS models reported in the figure are the same

already presented in Fig.11.5. The comparison between these figures shows that the constraints extracted from the tides are well compatible with the ones extracted from the NICER measurements presented in Fig.11.5, though none of these measurements are yet very constraining for the dense matter EoS. As it can be appreciated from this figure, the GW170817 observation is constraining enough to exclude some extreme EoS models, notably those assuming a purely nucleonic composition and very repulsive interactions at densities beyond nuclear saturation, such as the NL3 functional in the figure. However, to extract physically meaningful information on the actual structure of dense matter and on its EoS, Bayesian studies must be done. Using agnostic modelling, specifically a spectral parametrization of the $P(\rho)$ function (Lindblom 2010)¹⁰, the LVK collaboration has extracted a posterior distribution of the EoS (Abbott *et al.* 2018), that is reported in Fig. 11.8. To add the information of the complementary constraints from nuclear theory and experiment discussed in Section 11.2.3, and from other multi-messenger observations presented in Section 11.2.4, nuclear physics informed Bayesian modelling has been subsequently performed by different authors (see Mondal *et al.* (2023); Char, P. and Mondal, C. and Gulminelli, F. and Oertel, M. (2023); Davis *et al.* (2024); Scurto *et al.* (2024); Malik *et al.* (2024); Montefusco *et al.* (2025) for some recent works) in the line of Section 11.2.5. The use of a metamodel representation that starts from the energy functional as a function of the relevant degrees of freedom Eq.[11.13] also allows to progress in the interpretation of this EoS in terms of nuclear matter properties.

An example of this EoS inference is presented in Fig.11.8. This figure presents the EoS prior with a purely nucleonic metamodeling partially informed by nuclear experiments, and different posteriors including the theoretical constraint provided by ab-initio calculations of pure neutron matter (see Section 11.2.3) and the astrophysical information. The heavy mass pulsar constraint and the GW170817 observations are taken into account introducing directly the posterior parameter estimation of the observations in the likelihoods. In particular, for the tidal observation we have:

$$\mathcal{L}_{LVK} = \int P(\tilde{\Lambda}, q) dq, \quad [11.28]$$

where $P(\tilde{\Lambda}, q)$ is the observational joint posterior distribution reported in (Abbott *et al.* 2019). We can see that the nuclear physics and astrophysical constraints are very complementary in informing the EoS behavior. Interestingly,

10. The generality of the approach was validated by comparing to another agnostic representation based on piecewise polytropes (Read *et al.* 2009)




Figure 11.8: Bayesian inference of the EoS with the effect of the different constraints: ab-initio calculation of neutron matter EoS (upper right), TOV mass limit and GW170817 (lower left) and full posterior additionally including NICER mass-radius observations (lower right). The agnostic 90% probability contour from (Abbott *et al.* 2018) is also reported as dashed lines. [Taken from (Dinh Thi *et al.* 2021).]

the nucleonic modelling of Fig.11.8 agrees with the agnostic LVK estimation (Abbott *et al.* 2018), but the nucleonic hypothesis is more restrictive, and therefore leads to a more precisely defined EoS behavior. This implies that the GW170817 observation is consistent with a pure nucleonic composition of dense baryonic matter, but this hypothesis might be challenged if more precise tidal observations were to be available in the future.

11.4. Prospects for next generation of GW detectors

The increase in number of events and in the accuracy of the detection with future upgrades of the current interferometers and with the third generation of gravitational wave detectors will have the technological capability to constrain dense matter physics significantly. The next generation of detectors are discussed in details in [Chapter XXX](#), but in this last section, we specifically focus on the prospect for dense matter exploration.

The third generation of gravitational wave detectors (3G) is comprised of three different projects to date. Two of them are ground-based detectors and will be our primary focus in this section: Cosmic Explorer to be located in the




Figure 11.9: Average Bayes factor as a function of luminosity distance \mathcal{D}_L for different injection models, averaged over the respective class of hybrid meta-models. [Taken from ref.(Mondal *et al.* 2023)].

United States (Evans and *et. al.* 2021) and Einstein Telescope in Europe (Maggiore *et al.* 2020 ; Branchesi *et al.* 2023 ; Abac *et al.* 2025). The space-based Laser Interferometer Space Antenna (LISA) mission (Gong *et al.* 2011) led by the European Space Agency is scheduled to observe in the low frequency band $0.1 - 10^3$ ms: it will not be able to detect the late inspiral (and higher frequencies) part of the signal where the tidal deformation impacts the waveform. Dense matter exploration will be limited with LISA, although knowledge in NS physics will play an important role in cataloging the population of neutron star involved in binaries in our galaxy.

11.4.1. Detecting the tidal deformability with precision

Cosmic Explorer and Einstein Telescope will increase significantly the range of observations compared to current facilities, thus able to reach a maximum redshift $z \simeq 10$ for a canonical $1.4-1.4 M_\odot$ BNS merger. If both 3G detectors succeed in operating concurrently, we can expect $\sim 10^3 - 10^4$ detections of BNS mergers per year that are sufficiently precise to constrain the equation of state of dense matter. Recently, several studies have utilized the projected sensitivities of Cosmic Explorer and Einstein Telescope to estimate the equation of state recovery in the context of 3G, using different EoS injection models to simulate BNS mergers. A case study of GW170817 in the context of 3G has shown that a similar event could lead to a SNR of the order of 10^3 , and will be able to measure not only $\tilde{\Lambda}$ quite precisely, but also the next order correction

$\delta\tilde{\Lambda}$ and provide a measure of individual component tidal deformations (Smith *et al.* 2021). More complete estimations of the 3G capabilities were performed recently by simulating a population of BNS mergers to estimate to what extent an accumulation of events over long periods of observing time may constrain indirectly (through a full EoS inference) the radius of neutron stars, and showed that a few tens of meters to a few hundred will be achieved depending on the number of detectors concurrently observing (Huxford *et al.* 2024). In Iacovelli *et al.* (2023), the authors also consider a population study in the context of 3G, and discuss in detail the effects of the Einstein Telescope proposed configurations on the direct inference of dense matter nuclear empirical parameters defined in Eq.[11.13].

An example of a prospective study concerning the capabilities of third generation interferometers to detect possible phase transitions in ultra-dense matter, is reported in Fig.11.9. This figure displays the Bayes factor comparing two general modellings of the EoS, namely a purely nucleonic EoS posterior informed by the present nuclear and astrophysical constraints, and a more general parametrization (hybrid metamodels) where a first order phase transition occurs at a given baryonic density that is ad-hoc varied from $n_B = 0.3$ ("PT03") to $n_B = 0.5$ ("PT05") fm^{-3} . A merging event with characteristics similar to GW170817 is simulated at different luminosity distances using an hybrid injection EoS model, and the signal is filtered through the sensitivity curve of 3G interferometers. The Bayes factor measures the relative compatibility of the simulated signal with an EoS representation that does not contain the phase transition, with respect to the EoS family from which the injection model is taken. A value $\log(\langle B \rangle) > 1 - 2$ corresponds to a reliable discrimination between the two model classes. We can see that only if an event occurs analogous to GW170817 at $D = 40$ Mpc or closer, the sensitivity of 3G detectors might be enough to detect the early phase transition ("PT03") by only using the tidal information. Conversely, if the phase transition occurs at a very high baryonic density ("PT05"), the proportion of the deconfined phase in the star is negligible and the two models are equivalent; for more details, see Mondal *et al.* (2023). In all cases, Fig. 11.9 shows that a firm proof of a deconfined phase will be very hard to achieve with a single merging event within the present knowledge of the dense matter EoS. However, a notably increased discrimination power is expected in the case of multiple detections, as foreseen with 3G detectors, and with the addition of the analysis of oscillation modes of post-merger signals, as discussed above. For a more complete set of prospective studies, see the two recent bluebooks of the Einstein Telescope collaboration, (Branchesi *et al.* 2023 ; Abac *et al.* 2025).

11.4.2. *Beyond tidal deformation to constrain dense matter physics*

While the detection of the tidal deformation in the late inspiral of BNS mergers has provided the most compelling constraints on dense matter with gravitational waves so far, further observations and the accuracy increase of next generation of detectors offers a variety of prospects to constrain dense matter. For a review of the prospects in the context of Einstein Telescope, see e.g. (Abac *et al.* 2025).

The detection of stellar oscillations can help uncover the structure inside the neutron star and also give us a direct way of probing the star composition. Any sort of perturbations of the star can give rise to quasi-normal oscillation modes and radial (or non radial) pulsations, and in the case of non-axisymmetric oscillations is accompanied by gravitational wave emission. Oscillation modes can be calculated in a similar manner than Section 11.3, with or without the assumption of a weak gravitational field (Cowling 1941). A large plethora of oscillation modes can exist in neutron stars, for a review see (Kokkotas and Schmidt 1999): f (fundamental)-modes which present the strongest coupling with tides and are directly related to the fluid oscillation, g (gravity)-modes that are driven by the gravity response to inhomogeneities in the star (through buoyancy), p (pressure)-modes or acoustic modes that are driven by the pressure fluctuations and the local sound speed, r (rotation)-modes on rotating neutron stars risen by the effects of the Coriolis force, s (shear)-modes related to the shear viscosity in the crust of neutron stars, w (wave)-modes risen by the coupling of the fluid oscillation to the space-time metric, etc., see e.g. (Andersson 1998 ; Finn 1986 ; Kokkotas and Schutz 1992). For example, if the orbital frequency in the BNS merger matches the frequency of one of the modes of oscillation in the star, it may lead to a fracture of the crust and provide valuable insight into the nuclear empirical parameters (Neill *et al.* 2022). The excitation of the modes in core-collapse supernovae and proto-neutron stars, neutron stars or in BNS mergers may leave their imprint on the gravitational waveform.

Aside from detecting violent events of compact object coalescences, LVK interferometers and 3G may also detect continuous gravitational waves when non-axisymmetric deformations persist on rotating neutron stars and form mountains (due to accretion or starquakes). Current searches have only produced an upper limit on the mountain size (Abbott *et al.* 2022), but additional data and more precise detections may provide a measurement of the mountain size which directly relates to the neutron star equation of state and transport properties, such as the shear modulus. Post-merger physics will also be accessible to the next generation of gravitational wave detectors thanks to an increase in sensitivity in the high frequency band. The analysis of this part of the signal however is challenging as it involves complex modeling that includes finite temperature

and non-equilibrated matter effect, fast rotation and high magnetic fields (Shibata *et al.* 2021 ; Murguia-Berthier *et al.* 2021 ; Sarin and Lasky 2021). Finally, a multi interferometer detection of a BNS or BH-NS merger leading to a precise enough sky localization favors an analysis of the electromagnetic counterpart. By analysing the kilonova spectrum, one can determine the types of reactions occurring in the remnant matter and determine the role of compact object coalescences in r-process nucleosynthesis, and in turn provide insight in the origin of heavy elements in the Universe, see (Arcones and Thielemann 2023) for a review.

11.4.3. Challenges in the analysis

Given the significant increase in the precision with 3G, the community of gravitational wave analysts will have to carefully consider the battle between systematic and statistical uncertainties. For example, waveform models will need to include higher order correction than the mass quadrupole tidal response at 5PN order, see Section 11.3, if the detectors become sensitive enough to detect those effects. To analyse a precisely detected waveform, it will also be necessary to increase the sampling frequency for a signal that will be longer than detected in current facilities; these complexities paired to the addition of $\delta\tilde{\Lambda}$ in the Bayesian analysis, will lead to important numerical challenges in parameter estimation (PE), see [Chapter XXX](#). It will also be important to reevaluate the use of quasi-universal relations (see Section 11.3.2) in the context of decreased statistical errors: indeed, a small EoS dependence is observed in all those relations, which can be neglected only if the detection uncertainties dominate in the PE. For example, the estimation of the neutron star radius based on the quasi-universal relation between the compactness and the mass quadrupole tidal deformation performed for the GW170817 will be obsolete in the context of 3G (Suleiman and Read 2024). Finally, great precision on \mathcal{M}_c , q , $\tilde{\Lambda}$, $\delta\tilde{\Lambda}$ will constrain the equation of state and nuclear empirical parameters to the point that the Bayesian analysis will require sampling the EoS concurrently with the Bayesian inference: a pre-computed EoS prior containing $\sim 10^4$ EoS (as was used for the most recent analysis GW230529) will likely be decimated with an accumulation of sufficiently precise BNS mergers, with great loss of resolution to the posterior sample. In this context, several techniques to sample the EoS "on the fly" start to be available. They are notably based on nested sampling or other Monte-Carlo methods (Feroz and Hobson 2008 ; Wofford *et al.* 2023) , with great care of numerical efficiency particularly to ensure that the astrophysical quantities of neutron stars are calculated sufficiently fast (with optimized TOV and k_2 Love number solvers, employing ML techniques such as emulators or neural networks (Krastev *et al.* 2021 ; Reed *et al.* 2024)).

11.5. Bibliography

- Abac, A., Dietrich, T., Buonanno, A., Steinhoff, J., Ujevic, M. (2024), New and robust gravitational-waveform model for high-mass-ratio binary neutron star systems with dynamical tidal effects, *Physical Review D*, 109(2), 024062.
- Abac, A. *et al.* (2025), The Science of the Einstein Telescope, *arXiv e-prints*, p. arXiv:2503.12263.
- Abbott, B. P. *et al.* (2019), Properties of the Binary Neutron Star Merger GW170817, *Physical Review X*, 9(1), 011001.
- Abbott, B. P. *et al.* (2020), GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4 M_{\odot}$, *The Astrophysical Journal Letters*, 892(1), L3.
- Abbott, B. P., LIGO Scientific Collaboration, Virgo Collaboration *et al.* (2017a), GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, *Physical Review Letters*, 119(16), 161101.
- Abbott, B. P., LIGO Scientific Collaboration, Virgo Collaboration *et al.* (2017b), Multi-messenger Observations of a Binary Neutron Star Merger, *The Astrophysical Journal Letters*, 848(2), L12.
- Abbott, B. P., LIGO Scientific Collaboration, Virgo Collaboration, Kagra Collaboration *et al.* (2018), GW170817: Measurements of Neutron Star Radii and Equation of State, *Physical Review Letters*, 121(16), 161101.
- Abbott, R. *et al.* (2022), All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data, *Physical Review D*, 106(10), 102008.
- Alford, M., Reddy, S. (2003), Compact stars with color superconducting quark matter, *Phys. Rev. D*, 67, 074024.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.67.074024>
- Alsing, J., Silva, H. O., Berti, E. (2018), Evidence for a maximum mass cutoff in the neutron star mass distribution and constraints on the equation of state, *Monthly Notices of the Royal Astronomical Society*, 478(1), 1377–1391.
- Andersson, N. (1998), A New Class of Unstable Modes of Rotating Relativistic Stars, *The Astrophysical Journal*, 502(2), 708–713.
- Arcones, A., Thielemann, F.-K. (2023), Origin of the elements, *The Astronomy and Astrophysics Review*, 31(1), 1.
- Arzoumanian, Z., *et al.* (2018), The NANOGrav 11-year data set: High-precision timing of 45 millisecond pulsars, *The Astrophysical Journal Supplement Series*, 235(2), 37.
URL: <https://doi.org/10.3847/1538-4365/aab5b0>
- Baym, G., Pethick, C., Sutherland, P. (1971), The Ground State of Matter at High Densities: Equation of State and Stellar Models, *The Astrophysical*

- Journal*, 170, 299.
- Bell, J. F. (1996), Radio Pulsar Timing, *arXiv e-prints*, pp. astro-ph/9610145.
- Binnington, T., Poisson, E. (2009), Relativistic theory of tidal Love numbers, *Physical Review D*, 80(8), 084018.
- Branchesi, M. *et al.* (2023), Science with the Einstein Telescope: a comparison of different designs, *Journal of Cosmology and Astroparticle Physics*, 2023(7), 068.
- Buballa, M. (2005), Njl-model analysis of dense quark matter, *PHYSICS REPORTS-REVIEW SECTION OF PHYSICS LETTERS*, 407(4-6), 205–376.
- Burgio, G. F., Schulze, H. J., Vidana, I., Wei, J. B. (2021), Neutron stars and the nuclear equation of state, *PROGRESS IN PARTICLE AND NUCLEAR PHYSICS*, 120.
- Caplan, M. E., Schneider, A. S., Horowitz, C. J. (2018), Elasticity of Nuclear Pasta, *Physical Review Letters*, 121(13), 132701.
- Castro, G., Gualtieri, L., Maselli, A., Pani, P. (2022), Impact and detectability of spin-tidal couplings in neutron star inspirals, *Physical Review D*, 106(2), 024011.
- Centelles, M., Vinas, X., Barranco, M., Schuck, P. (1993), A Semiclassical Approach to Relativistic Nuclear Mean Field Theory, *Annals of Physics*, 221(1), 165–204.
- Chamel, N., Haensel, P. (2008), Physics of Neutron Star Crusts, *Living Reviews in Relativity*, 11(1), 10.
- Chandrasekhar, S. (1931), The maximum mass of ideal white dwarfs, *Astrophys. J.*, 74, 81–82.
- Char, P. and Mondal, C. and Gulminelli, F. and Oertel, M. (2023), Generalized description of neutron star matter with a nucleonic relativistic density functional, *Physical Review D*, 108(10), 103045.
- Choudhury, D., Salmi, T., Vinciguerra, S., Riley, T. E., Kini, Y., Watts, A. L., Dorsman, B., Bogdanov, S., Guillot, S., Ray, P. S., Reardon, D. J., Remillard, R. A., Bilous, A. V., Huppenkothen, D., Lattimer, J. M., Rutherford, N., Arzoumanian, Z., Gendreau, K. C., Morsink, S. M., Ho, W. C. G. (2024), A NICER View of the Nearest and Brightest Millisecond Pulsar: PSR J0437–4715, *The Astrophysical Journal Letters*, 971(1), L20.
- Cowling, T. G. (1941), The non-radial oscillations of polytropic stars, *Monthly Notices of the Royal Astronomical Society*, 101, 367.
- Cunningham, C. T., Price, R. H., Moncrief, V. (1978), Radiation from collapsing relativistic stars. I. Linearized odd-parity radiation., *The Astrophysical Journal*, 224, 643–667.

- Damour, T., Nagar, A. (2009), Relativistic tidal properties of neutron stars, *Physical Review D*, 80(8), 084035.
- Das, H. C., Kumar, A., Kumar, B., Biswal, S. K., Nakatsukasa, T., Li, A., Patra, S. K. (2020), Effects of dark matter on the nuclear and neutron star matter, *Monthly Notices of the Royal Astronomical Society*, 495(4), 4893–4903.
URL: <https://doi.org/10.1093/mnras/staa1435>
- Davis, P. J., Thi, H. D., Fantina, A. F., Gulminelli, F., Oertel, M., Suleiman, L. (2024), Inference of neutron-star properties with unified crust-core equations of state for parameter estimation, *Astronomy & Astrophysics*, 687, A44.
- Dinh Thi, H., Fantina, A. F., Gulminelli, F. (2022), Properties of pasta phases in catalyzed neutron stars, *arXiv e-prints*, p. arXiv:2206.07969.
- Dinh Thi, H., Fantina, A. F., Gulminelli, F. (2023), The proto-neutron star inner crust in a multi-component plasma approach, *Astronomy & Astrophysics*, 677, A174.
URL: <https://doi.org/10.1051/0004-6361/202346606>
- Dinh Thi, H., Mondal, C., Gulminelli, F. (2021), The Nuclear Matter Density Functional under the Nucleonic Hypothesis, *Universe*, 7(10), 373.
- Douchin, F., Haensel, P. (2001a), A unified equation of state of dense matter and neutron star structure, *Astronomy & Astrophysics*, 380, 151–167.
- Douchin, F., Haensel, P. (2001b), A unified equation of state of dense matter and neutron star structure, *Astronomy & Astrophysics*, 380, 151–167.
- Evans, M., et. al. (2021), A Horizon Study for Cosmic Explorer: Science, Observatories, and Community, *arXiv e-prints*, p. arXiv:2109.09882.
- Fantina, A. F., Chamel, N., Pearson, J. M., Goriely, S. (2013), Neutron star properties with unified equations of state of dense matter, *Astronomy & Astrophysics*, 559, A128.
- Feroz, F., Hobson, M. P. (2008), Multimodal nested sampling: an efficient and robust alternative to markov chain monte carlo methods for astronomical data analyses, *Monthly Notices of the Royal Astronomical Society*, 384(2), 449–463.
URL: <https://doi.org/10.1111/j.1365-2966.2007.12353.x>
- Ferreira, M., Pereira, R. C., Providência, C. (2020), Neutron stars with large quark cores, *Physical Review D*, 101(12), 123030.
- Finn, L. S. (1986), G-modes of non-radially pulsating relativistic stars - The slow-motion formalism, *Monthly Notices of the Royal Astronomical Society*, 222, 393–416.
- Flanagan, É. É., Hinderer, T. (2008), Constraining neutron-star tidal Love numbers with gravitational-wave detectors, *Physical Review D*, 77(2), 021502.

- Fonseca, E., et al. (2021), Refined Mass and Geometric Measurements of the High-mass PSR J0740+6620, *Astrophysical Journal, Letters*, 915(1), L12.
- Freire, P. (2021), ‘Pulsar mass measurements and tests of general relativity’.
URL: https://www3.mpi-fr-bonn.mpg.de/staff/pfreire/NS_masses.html
- Gandolfi, S., Gezerlis, A., Carlson, J. (2015), Neutron Matter from Low to High Density, *Annual Review of Nuclear and Particle Science*, 65(1), 303–328.
- Gendreau, K., Arzoumanian, Z., Okajima, T. (2012), The neutron star interior composition explorer (nicer): An explorer mission of opportunity for soft x-ray timing spectroscopy, *Proceedings of SPIE - The International Society for Optical Engineering*, 8443, 13.
- Glendenning, N. (2012), *Compact Stars: Nuclear Physics, Particle Physics and General Relativity*, Astronomy and Astrophysics Library, Springer New York.
URL: <https://books.google.fr/books?id=cCDlBwAAQBAJ>
- Gong, X. et al. (2011), A scientific case study of an advanced LISA mission, *Classical and Quantum Gravity*, 28(9), 094012.
- Goriely, S., Chamel, N., Pearson, J. M. (2010), Further explorations of Skyrme-Hartree-Fock-Bogoliubov mass formulas. XII. Stiffness and stability of neutron-star matter, *Physical Review C*, 82(3), 035804.
- Goriely, S., Chamel, N., Pearson, J. M. (2013), Further explorations of Skyrme-Hartree-Fock-Bogoliubov mass formulas. XIII. The 2012 atomic mass evaluation and the symmetry coefficient, *Physical Review C*, 88(2), 024308.
- Gourgoulhon, E. (2010), An introduction to the theory of rotating relativistic stars, *arXiv e-prints*, p. arXiv:1003.5015.
- Gupta, P. K., Steinhoff, J., Hinderer, T. (2021), Relativistic effective action of dynamical gravitomagnetic tides for slowly rotating neutron stars, *Physical Review Research*, 3(1), 013147.
- Haensel, P., Potekhin, A. Y., Yakovlev, D. G. (2007), *Neutron Stars 1 : Equation of State and Structure*, vol. 326, Springer Science & Business Media, 2007.
- Haensel, P., Zdunik, J. L. (1990a), Equation of state and structure of the crust of an accreting neutron star, *Astronomy and Astrophysics*, 229(1), 117–122.
- Haensel, P., Zdunik, J. L. (1990b), Non-equilibrium processes in the crust of an accreting neutron star, *Astronomy and Astrophysics*, 227(2), 431–436.
- Haensel, P., Zdunik, J. L., Bejger, M. (2008), Fast rotation of neutron stars and equation of state of dense matter, *New Astronomy Reviews*, 51(10-12), 785–790.
- Haensel, P., Zdunik, J. L., Bejger, M., Lattimer, J. M. (2009), Keplerian frequency of uniformly rotating neutron stars and strange stars, *Astronomy and Astrophysics*, 502(2), 605–610.

- Henry, Q., Faye, G., Blanchet, L. (2020), Tidal effects in the gravitational-wave phase evolution of compact binary systems to next-to-next-to-leading post-Newtonian order, *Physical Review D*, 102(4), 044033.
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., Collins, R. A. (1968), Observation of a Rapidly Pulsating Radio Source, *Nature*, 217(5130), 709–713.
- Hinderer, T. (2008), Tidal Love Numbers of Neutron Stars, *Astrophysical Journal*, 677(2), 1216–1220.
- Horowitz, C. J., Piekarewicz, J. (2001), Neutron Star Structure and the Neutron Radius of ^{208}Pb , *Physical Review Letters*, 86(25), 5647–5650.
- Hulse, R. A., Taylor, J. H. (1975), Discovery of a pulsar in a binary system., *The Astrophysical Journal*, 195, L51–L53.
- Huth, S., Wellenhofer, C., Schwenk, A. (2021), New equations of state constrained by nuclear physics, observations, and qcd calculations of high-density nuclear matter, *Physical Review C*, 103(2), 025803.
- Huxford, R., Kashyap, R., Borhanian, S., Dhani, A., Gupta, I., Sathyaprakash, B. S. (2024), Accuracy of neutron star radius measurement with the next generation of terrestrial gravitational-wave observatories, *Physical Review D*, 109(10), 103035.
- Iacovelli, F., Mancarella, M., Mondal, C., Puecher, A., Dietrich, T., Gulminelli, F., Maggiore, M., Oertel, M. (2023), Nuclear physics constraints from binary neutron star mergers in the Einstein Telescope era, *Physical Review D*, 108(12), 122006.
- Janka, H. T., Langanke, K., Marek, A., Martinez-Pinedo, G., Muller, B. (2007), Theory of core-collapse supernovae, *Physics Reports*, 442(1-6), 38–74.
- Klausner, P., Colò, G., Roca-Maza, X., Vigezzi, E. (2025), Impact of ground-state properties and collective excitations on the skyrme ansatz: A bayesian study, *Physical Review C*, 111, 014311.
URL: <https://link.aps.org/doi/10.1103/PhysRevC.111.014311>
- Koehn, H., Rose, H., Pang, P. T. H., Somasundaram, R., Reed, B. T., Tews, I., Abac, A., Komoltsev, O., Kunert, N., Kurkela, A., Coughlin, M. W., Healy, B. F., Dietrich, T. (2025), ‘From existing and new nuclear and astrophysical constraints to stringent limits on the equation of state of neutron-rich dense matter’.
URL: <https://arxiv.org/abs/2402.04172>
- Kokkotas, K. D., Schmidt, B. G. (1999), Quasi-Normal Modes of Stars and Black Holes, *Living Reviews in Relativity*, 2(1), 2.
- Kokkotas, K. D., Schutz, B. F. (1992), W-modes - A new family of normal modes of pulsating relativistic stars, *Monthly Notices of the Royal Astronomical Society*, 255, 119–128.

- Kramer, M., et al. (2021), Strong-Field Gravity Tests with the Double Pulsar, *Physical Review X*, 11(4), 041050.
- Krastev, P. G., Gill, K., Villar, V. A., Berger, E. (2021), Detection and parameter estimation of gravitational waves from binary neutron-star mergers in real ligo data using deep learning, *PHYSICS LETTERS B*, 815.
- Kumar, R., Dexheimer, V., Jahan, J., Noronha, J., Noronha-Hostler, J., Ratti, C., Yunes, N., Acuna, A., Alford, M., Anik, M., Chatterjee, D., Chatziioannou, K., Chen, H.-Y., Clevinger, A., Conde, C., Cruz Camacho, N., Dore, T., Drischler, C., Elfner, H., Zhang, Z. (2024), Theoretical and experimental constraints for the equation of state of dense and hot matter, *Living Reviews in Relativity*, 27.
- Landry, P., Essick, R. (2019), Nonparametric inference of the neutron star equation of state from gravitational wave observations, *Physical Review D*, 99(8), 084049.
- LIGO Scientific Collaboration, Virgo Collaboration, KAGRA Collaboration (2018), 'LVK Algorithm Library - LALSuite', Free software (GPL).
- Lim, Y., Schwenk, A. (2024), Symmetry energy and neutron star properties constrained by chiral effective field theory calculations, *Phys. Rev. C*, 109(3), 035801.
- Lindblom, L. (2010), Spectral representations of neutron-star equations of state, *Physical Review D*, 82, 103011.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.82.103011>
- Lindblom, L. (2018), Causal representations of neutron-star equations of state, *Physical Review D*, 97(12), 123019.
- Love, A. E. H. (1911), *Some Problems of Geodynamics*.
- Machleidt, R., Sammarruca, F. (2024), Recent advances in chiral eft based nuclear forces and their applications, *Progress in Particle and Nuclear Physics*, 137, 104117.
URL: <https://www.sciencedirect.com/science/article/pii/S0146641024000218>
- Mackay, D. J. C. (2003), *Information Theory, Inference and Learning Algorithms*.
- Mackie, F. D., Baym, G. (1977), Compressible liquid drop nuclear model and mass formula, *Nuclear Physics A*, 285(2), 332–348.
- Maggiore, M. et al. (2020), Science case for the Einstein telescope, *Journal of Cosmology and Astroparticle Physics*, 2020(3), 050.
- Malik, T., Pais, H., Providência, C. (2024), Unified neutron star equations of state calibrated to nuclear properties, *Astronomy & Astrophysics*, 689, A242.
- Margueron, J., Hoffmann Casali, R., Gulminelli, F. (2018), Equation of state for dense nucleonic matter from metamodeling. I. Foundational aspects, *Physical Review C*, 97(2), 025805.

- Mauviard, L., Guillot, S., Salmi, T., Choudhury, D., Dorsman, B., González-Jez-Caniulef, D., Hoogkamer, M., Huppenkothen, D., Kazantsev, C., Kini, Y., Olive, J.-F., Stammler, P., Watts, A. L., Mendes, M., Rutherford, N., Schwenk, A., Svensson, I., Bogdanov, S., Kerr, M., Ray, P. S., Guillemot, L., Cognard, I., Theureau, G. (2025), ‘A nicer view of the 1.4 solar-mass edge-on pulsar psr j0614–3329’.
URL: <https://arxiv.org/abs/2506.14883>
- Mondal, C., Antonelli, M., Gulminelli, F., Mancini, M., Novak, J., Oertel, M. (2023), Detectability of a phase transition in neutron star matter with third-generation gravitational wave interferometers, *Mon. Not. Roy. Astron. Soc.*, 524(3), 3464–3473.
- Montefusco, G., Antonelli, M., Gulminelli, F. (2025), Frozen and beta-equilibrated f and p modes of cold neutron stars: Nuclear metamodel predictions, *Astron. Astrophys.*, 694, A150.
URL: <https://doi.org/10.1051/0004-6361/202452727>
- Murguia-Berthier, A., Noble, S. C., Roberts, L. F., Ramirez-Ruiz, E., Wernicke, L. R., Kolacki, M., Etienne, Z. B., Avara, M., Campanelli, M., Ciolfi, R., Cipolletta, F., Drachler, B., Ennoggi, L., Faber, J., Fiacco, G., Giacomazzo, B., Gupte, T., Ha, T., Kelly, B. J., Krolik, J. H., Lopez Armengol, F. G., Margalit, B., Moon, T., O’Shaughnessy, R., Rueda-Becerril, J. M., Schnittman, J., Zenati, Y., Zlochower, Y. (2021), HARM3D+NUC: A New Method for Simulating the Post-merger Phase of Binary Neutron Star Mergers with GRMHD, Tabulated EOS, and Neutrino Leakage, *The Astrophysical Journal*, 919(2), 95.
- Murray, C. D., Dermott, S. F. (1999), *Solar System Dynamics*.
- Nagata, K. (2022), Finite-density lattice QCD and sign problem: Current status and open problems, *Prog. Part. Nucl. Phys.*, 127, 103991.
- Negele, J. W., Vautherin, D. (1973), Neutron star matter at sub-nuclear densities, *Nuclear Physics A*, 207(2), 298–320.
- Neill, D., Tsang, D., van Eerten, H., Ryan, G., Newton, W. G. (2022), Resonant shattering flares in black hole-neutron star and binary neutron star mergers, *Monthly Notices of the Royal Astronomical Society*, 514(4), 5385–5402.
- Oertel, M., Providencia, C., Gulminelli, F., Raduta, A. R. (2015), Hyperons in neutron star matter within relativistic mean-field models, *JOURNAL OF PHYSICS G-NUCLEAR AND PARTICLE PHYSICS*, 42(7).
- Olmo, G. J., Rubiera-Garcia, D., Wojnar, A. (2020), Stellar structure models in modified theories of gravity: Lessons and challenges, *Physics Reports*, 876, 1–75. Stellar structure models in modified theories of gravity: Lessons and challenges.
URL: <https://www.sciencedirect.com/science/article/pii/S0370157320302507>

- Onsi, M., Dutta, A. K., Chatri, H., Goriely, S., Chamel, N., Pearson, J. M. (2008), Semi-classical equation of state and specific-heat expressions with proton shell corrections for the inner crust of a neutron star, *Physical Review C*, 77, 065805.
URL: <https://link.aps.org/doi/10.1103/PhysRevC.77.065805>
- Oppenheimer, J. R., Volkoff, G. M. (1939), On Massive Neutron Cores, *Physical Review*, 55(4), 374–381.
- Özel, F., Freire, P. (2016), Masses, Radii, and the Equation of State of Neutron Stars, *Annual Review of Astronomy and Astrophysics*, 54, 401–440.
- Pappas, G., Doneva, D. D., Sotiriou, T. P., Yazadjiev, S. S., Kokkotas, K. D. (2019), Multipole moments and universal relations for scalarized neutron stars, *Physical Review D*, 99(10), 104014.
- Pastore, A., Baroni, S., Losa, C. (2011), Superfluid properties of the inner crust of neutron stars, *Physical Review C*, 84(6).
- Pearson, J. M., Goriely, S., Chamel, N. (2011), Properties of the outer crust of neutron stars from Hartree-Fock-Bogoliubov mass models, *Physical Review C*, 83(6), 065810.
- Pereira, J. P., Bejger, M., Andersson, N., Gittins, F. (2020), Tidal Deformations of Hybrid Stars with Sharp Phase Transitions and Elastic Crusts, *The Astrophysical Journal*, 895(1), 28.
- Pereira, R. C., Costa, P., Providência, C. (2016), Two-solar-mass hybrid stars: A two model description using the Nambu-Jona-Lasinio quark model, *Physical Review D*, 94(9), 094001.
- Pfaff, A., Hansen, H., Gulminelli, F. (2022), Bayesian analysis of the properties of hybrid stars with the nambu-jona-lasinio model, *Physical Review C*, 105, 035802.
URL: <https://link.aps.org/doi/10.1103/PhysRevC.105.035802>
- Postnikov, S., Prakash, M., Lattimer, J. M. (2010), Tidal Love numbers of neutron and self-bound quark stars, *Physical Review D*, 82(2), 024016.
- Potekhin, A., Chabrier, G. (2017), Magnetic neutron star cooling and microphysics, *Astronomy & Astrophysics*, 609.
- Read, J. S., Lackey, B. D., Owen, B. J., Friedman, J. L. (2009), Constraints on a phenomenologically parametrized neutron-star equation of state, *Physical Review D*, 79(12), 124032.
- Reed, B. T., Somasundaram, R., De, S., Armstrong, C. L., Giuliani, P., Capano, C., Brown, D. A., Tews, I. (2024), Toward accelerated nuclear-physics parameter estimation from binary neutron star mergers: Emulators for the tolmán-oppenheimer-volkoff equations, *The Astrophysical Journal*, 974(2), 285.
URL: <https://dx.doi.org/10.3847/1538-4357/ad737c>

- Ripley, J. L., Abhishek Hegade K., R., Yunes, N. (2023), Probing internal dissipative processes of neutron stars with gravitational waves during the inspiral of neutron star binaries, *Physical Review D*, 108(10), 103037.
- Salaris, M., Cassisi, S. (2005), *Evolution of Stars and Stellar Populations*, Wiley.
URL: <https://books.google.fr/books?id=r1dNzr8viRYC>
- Salmi, T., Choudhury, D., Kini, Y., Riley, T. E., Vinciguerra, S., Watts, A. L., Wolff, M. T., Arzoumanian, Z., Bogdanov, S., Chakrabarty, D., Gendreau, K., Guillot, S., Ho, W. C. G., Huppenkothen, D., Ludlam, R. M., Morsink, S. M., Ray, P. S. (2024), The Radius of the High-mass Pulsar PSR J0740+6620 with 3.6 yr of NICER Data, *The Astrophysical Journal*, 974(2), 294.
- Sarin, N., Lasky, P. D. (2021), The evolution of binary neutron star post-merger remnants: a review, *General Relativity and Gravitation*, 53(6), 59.
- Schmidt, P., Hinderer, T. (2019), Frequency domain model of f-mode dynamic tides in gravitational waveforms from compact binary inspirals, *Physical Review D*, 100(2), 021501.
- Schmitt, A., Shternin, P. (2018), Reaction Rates and Transport in Neutron Stars, in L. Rezzolla, P. Pizzochero, D. I. Jones, N. Rea, I. Vidaña, (eds), *Astrophysics and Space Science Library*, vol. 457 of *Astrophysics and Space Science Library*, p. 455.
- Schneider, A. S., Horowitz, C. J., Hughto, J., Berry, D. K. (2013), Nuclear “pasta” formation, *Physical Review C*, 88, 065807.
URL: <https://link.aps.org/doi/10.1103/PhysRevC.88.065807>
- Scurto, L., Pais, H., Gulminelli, F. (2024), General predictions of neutron star properties using unified relativistic mean-field equations of state, *Physical Review D*, 109(10), 103015.
- Sedrakian, A., Li, J. J., Weber, F. (2023), Heavy baryons in compact stars, *Progress in Particle and Nuclear Physics*, 131, 104041.
URL: <https://www.sciencedirect.com/science/article/pii/S0146641023000224>
- Shapiro, I. I. (1964), Fourth Test of General Relativity, *Physical Review Lett.*, 13, 789–791.
- Shibata, M., Fujibayashi, S., Sekiguchi, Y. (2021), Long-term evolution of a merger-remnant neutron star in general relativistic magnetohydrodynamics: Effect of magnetic winding, *Phys. Rev. D*, 103, 043022.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.043022>
- Smith, R., Borhanian, S., Sathyaprakash, B., Hernandez Vivanco, F., Field, S. E., Lasky, P., Mandel, I., Morisaki, S., Ottaway, D., Slagmolen, B. J. J., Thrane, E., Töyrä, D., Vitale, S. (2021), Bayesian Inference for Gravitational Waves from Binary Neutron Star Mergers in Third Generation Observatories,

- Physical Review Letters*, 127(8), 081102.
- Somasundaram, R., Tews, I., Margueron, J. (2023), Investigating signatures of phase transitions in neutron-star cores, *Physical Review C*, 107(2).
- Sorensen, A. *et al.* (2024), Dense nuclear matter equation of state from heavy-ion collisions, *Progress in Particle and Nuclear Physics*, 134, 104080.
URL: <https://www.sciencedirect.com/science/article/pii/S0146641023000613>
- Suleiman, L., Read, J. (2024), Quasiuniversal relations in the context of future neutron star detections, *Physical Review D*, 109(10), 103029.
- Tews, I., Margueron, J., Reddy, S. (2018), Critical examination of constraints on the equation of state of dense matter obtained from gw170817, *Physical Review C*, 98, 045804.
URL: <https://link.aps.org/doi/10.1103/PhysRevC.98.045804>
- Thorne, K. S., Campolattaro, A. (1967), ‘Non-Radial Pulsation of General-Relativistic Stellar Models. I. Analytic Analysis for $L \geq 2$ ’.
- Tolman, R. C. (1939), Static Solutions of Einstein’s Field Equations for Spheres of Fluid, *Physical Review*, 55(4), 364–373.
- Typel, S., *et al.* (2022), CompOSE Reference Manual, *arXiv e-prints*, p. arXiv:2203.03209.
- Typel, S., Röpke, G., Klähn, T., Blaschke, D., Wolter, H. H. (2010), Composition and thermodynamics of nuclear matter with light clusters, *Physical Review C*, 81(1), 015803.
- Wade, L., Creighton, J. D. E., Ochsner, E., Lackey, B. D., Farr, B. F., Littenberg, T. B., Raymond, V. (2014), Systematic and statistical errors in a bayesian approach to the estimation of the neutron-star equation of state using advanced gravitational wave detectors, *Physical Review D*, 89, 103012.
URL: <https://link.aps.org/doi/10.1103/PhysRevD.89.103012>
- Wang, M., Huang, W. J., Kondev, F. G., Audi, G., Naimi, S. (2021), The AME 2020 atomic mass evaluation (II). Tables, graphs and references, *Chinese Physics C*, 45(3), 030003.
- Wofford, J., Yelikar, A. B., Gallagher, H., Champion, E., Wysocki, D., Delfavero, V., Lange, J., Rose, C., Valsan, V., Morisaki, S., Read, J., Henshaw, C., O’Shaughnessy, R. (2023), Improving performance for gravitational-wave parameter inference with an efficient and highly-parallelized algorithm, *Physical Review D*, 107(2), 024040.
- Yagi, K., Yunes, N. (2013), I-Love-Q: Unexpected Universal Relations for Neutron Stars and Quark Stars, *Science*, 341(6144), 365–368.